

---

# Improved Algorithms for Misspecified Linear Markov Decision Processes

---

Daniel Vial  
UIUC

Advait Parulekar  
UT Austin

Sanjay Shakkottai  
UT Austin

R. Srikant  
UIUC

## Abstract

For the misspecified linear Markov decision process (MLMDP) model of [Jin et al. \(2020\)](#), we propose an algorithm with three desirable properties. **(P1)** Its regret after  $K$  episodes scales as  $K \max\{\varepsilon_{\text{mis}}, \varepsilon_{\text{tol}}\}$ , where  $\varepsilon_{\text{mis}}$  is the degree of misspecification and  $\varepsilon_{\text{tol}}$  is a user-specified error tolerance. **(P2)** Its space and per-episode time complexities are bounded as  $K \rightarrow \infty$ . **(P3)** It does not require  $\varepsilon_{\text{mis}}$  as input. To our knowledge, this is the first algorithm satisfying all three properties. For concrete choices of  $\varepsilon_{\text{tol}}$ , we also improve existing regret bounds (up to log factors) while achieving either **(P2)** or **(P3)** (existing algorithms satisfy neither). At a high level, our algorithm generalizes (to MLMDPs) and refines the `Sup-Lin-UCB` algorithm, which [Takemura et al. \(2021\)](#) recently showed satisfies **(P3)** for contextual bandits. We also provide an intuitive interpretation of their result, which informs the design of our algorithm.

## 1 INTRODUCTION

Due to the large state spaces of modern reinforcement learning applications, practical algorithms must generalize across states. To understand generalization on a theoretical level, recent work has studied linear Markov decision processes (LMDPs), among other models (see [Section 1.1](#)). The LMDP model assumes the next-state distribution and reward are linear in known  $d$ -dimensional features, which enables tractable generalization when  $d$  is small. Of course, this assumption most likely fails in practice, which motivates the misspecified LMDP (MLMDP) model. Here linearity

holds up to some misspecification error  $\varepsilon_{\text{mis}}$  in total variation and absolute value for the next-state distribution and reward, respectively (see [Assumption 1](#)).

In this work, we consider episodic finite-horizon MLMDPs, i.e., for each of  $K$  episodes, the algorithm interacts with the MLMDP for  $H$  steps. We assume the action space  $\mathcal{A}$  is finite, though the state space  $\mathcal{S}$  may be infinite. We measure performance in terms of regret  $R(K)$ , i.e., the additive loss in expected cumulative reward compared to the optimal policy (see [\(1\)](#)). We seek an algorithm with three basic properties:

**(P1) Asymptotically non-trivial regret:** Given a user-specified error tolerance  $\varepsilon_{\text{tol}} > 0$  and a failure probability  $\delta > 0$ , the algorithm should ensure that with probability at least  $1 - \delta$ ,

$$\limsup_{K \rightarrow \infty} \frac{R(K)}{K} \leq \text{poly} \left( d, H, \log \frac{1}{\delta \varepsilon_{\text{tol}}} \right) \max\{\varepsilon_{\text{mis}}, \varepsilon_{\text{tol}}\}.$$

Hence, in terms of  $K$ ,  $\varepsilon_{\text{mis}}$ , and  $\varepsilon_{\text{tol}}$ , we desire regret that scales as  $K \max\{\varepsilon_{\text{mis}}, \varepsilon_{\text{tol}}\}$ . The  $K\varepsilon_{\text{mis}}$  term is unavoidable due to misspecification, while the  $K\varepsilon_{\text{tol}}$  term can be controlled by the user.

**(P2) Bounded complexity:** The space and per-episode time complexities should both be independent of  $K$ , so that the algorithm can be implemented for arbitrarily large  $K$ .

**(P3) Parameter free:** The algorithm should not require knowledge of the degree of misspecification  $\varepsilon_{\text{mis}}$ , which is unavailable in practice.

These properties seem benign, but to the best of our knowledge, *no existing algorithm satisfies all three*. We note in particular that **(P1)** often fails because regret guarantees include  $K\varepsilon_{\text{mis}}\text{polylog}(K)$  terms. We would argue such super-linear (in  $K$ ) bounds are asymptotically trivial, since the regret of any policy is linear in  $K$  (for bounded rewards). However, even if  $K$  is small and one can tolerate failure of the asymptotically-motivated **(P1)** and **(P2)**, there are (essentially) no algorithms with  $K\varepsilon_{\text{mis}}\text{polylog}(K)$ -type regret that satisfy **(P3)**. (An exception is model selection, which violates a stronger version of **(P3)**; see [Section 1.1](#).)

Table 1: Summary of MLMDP algorithms, where  $\varepsilon_{\text{mis}}$  is the degree of misspecification and  $\varepsilon_{\text{tol}}$  is a user-chosen input. Regret bounds hide constants and log terms, except those that yield super-linear in  $K$  terms. See Section 1.1 for other existing algorithms that require more restrictive assumptions.

Algorithm	(P1)	(P2)	(P3)	Regret bound
Ours, $\varepsilon_{\text{tol}} \in (0, 1)^*$	Yes	Yes	Yes	$\sqrt{d^3 H^4 \min\{(\frac{d}{\varepsilon_{\text{tol}}})^2, K\}} + \sqrt{H^3 K} + \sqrt{d} H^2 K \max\{\varepsilon_{\text{mis}}, \varepsilon_{\text{tol}}\}$
Ours, $\varepsilon_{\text{tol}} = d/\sqrt{K}$	No	No	Yes	$\sqrt{d^3 H^4 K} + \sqrt{d} H^2 K \varepsilon_{\text{mis}} \sqrt{\log K}$
Ours, $\varepsilon_{\text{tol}} = \varepsilon_{\text{mis}}$	Yes	Yes	No	$\sqrt{d^3 H^4 \min\{(\frac{d}{\varepsilon_{\text{tol}}})^2, K\}} + \sqrt{H^3 K} + \sqrt{d} H^2 K \varepsilon_{\text{mis}}$
Jin et al. (2020)	No	No	No	$\sqrt{d^3 H^4 K} + d H^2 K \varepsilon_{\text{mis}} \sqrt{\log K}$
Zanette et al. (2020a) <sup>†</sup>	Yes	No	No	$\sqrt{d^4 H^5 K} + d H^2 K \varepsilon_{\text{mis}}$
Zanette et al. (2020b) <sup>† ‡</sup>	Yes	No	No	$\sqrt{d^2 H^4 K} + \sqrt{d} H^2 K \varepsilon_{\text{mis}}$

\*To satisfy (P1) and avoid a super-linear regret bound, this row assumes we choose  $\varepsilon_{\text{tol}}$  independent of  $K$ .

<sup>†</sup>These papers report regret with  $\tilde{O}$  notation but (as far as we can tell) do not hide super-linear terms.

<sup>‡</sup>The algorithm from this paper is defined in terms of an optimization problem but no solution to that problem is provided, so it is unclear how to implement it efficiently.

The situation is better if we restrict to misspecified linear contextual bandits (MLCBs), which are the special case  $H = 1$  (there, states are called contexts; we use the terms interchangeably). In particular, Takemura et al. (2021) showed that a Sup-Lin-UCB (Auer, 2002; Chu et al., 2011) variant that satisfies (P3) has  $K\varepsilon_{\text{mis}}\text{polylog}(K)$  regret. This result is important because the simpler Lin-UCB (Abbasi-Yadkori et al., 2011) can suffer  $\Omega(K)$  regret when  $\varepsilon_{\text{mis}}$  is unknown (Lattimore et al., 2020). In light of this, and because Sup-Lin-UCB was originally motivated by technical issues *seemingly unrelated to misspecification*, Takemura et al. (2021)’s result is also rather surprising. However, Sup-Lin-UCB is a complicated algorithm, so it was not intuitively clear (at least to us) why it should adapt to the misspecified setting better than Lin-UCB. As will be seen, one of our contributions is to provide a new interpretation of Sup-Lin-UCB that intuitively explains this. Furthermore, our interpretation is a key building block that leads to improved results for MLMDPs.

**Contributions:** Our contributions are as follows.

- **An intuitive Sup-Lin-UCB variant:** In Section 3, we show that Sup-Lin-UCB’s success for MLCBs is not an accident; rather, it can be *derived* from the perspective of misspecification. More precisely, we first propose an MLCB algorithm called EXPL3, which explicitly decides to explore or exploit. EXPL3 is simple and intuitive but requires  $\varepsilon_{\text{mis}}$  as input to perform well. We overcome this requirement by constructing an intuitive ensemble of EXPL3 algorithms, one for each possible  $\varepsilon_{\text{mis}}$  value (see Proposition 2). The ensemble closely resembles Sup-Lin-UCB and sheds light onto Takemura et al. (2021)’s result.
- **The Sup-LSVI-UCB algorithm:** In Section 4, we leverage the insights developed for MLCBs to de-

sign an MLMDP algorithm called Sup-LSVI-UCB.<sup>1</sup> At a high level, Sup-LSVI-UCB combines our Sup-Lin-UCB variant with a backward induction procedure, analogous to how LSVI-UCB (Jin et al., 2020) was derived from Lin-UCB. However, because Sup-Lin-UCB is more complicated than Lin-UCB, we encounter new technical issues when generalizing from MLCBs to MLMDPs, which requires some new algorithmic ideas; see Remarks 3 and 4.

- **Improved guarantees:** In Section 5, we show that when the input  $\varepsilon_{\text{tol}}$  is chosen independent of  $K$ , Sup-LSVI-UCB is the first algorithm to satisfy (P1), (P2), and (P3) (see Theorem 1). If instead  $\varepsilon_{\text{tol}} = d/\sqrt{K}$ , Sup-LSVI-UCB improves existing regret bounds (up to log factors), while simultaneously removing the requirement that  $\varepsilon_{\text{mis}}$  is known (Corollary 1). Finally, if  $\varepsilon_{\text{mis}}$  is known, we can set  $\varepsilon_{\text{tol}} = \varepsilon_{\text{mis}}$  to improve existing bounds while simultaneously avoiding unbounded complexity (Corollary 2). See Table 1.

Finally, though somewhat orthogonal to our main results, we also revisit Lin-UCB for MLCBs. While it is known that this algorithm can be modified to obtain  $K\varepsilon_{\text{mis}}\text{polylog}(K)$  regret when  $\varepsilon_{\text{mis}}$  is known (Lattimore et al., 2020; Jin et al., 2020), we are not aware of any bounds that satisfy (P1). Hence, we propose a new modification that satisfies (P1) and sharpens existing bounds when  $K \geq \varepsilon_{\text{mis}}^{-2}$ . See Section 6.

## 1.1 Related work

**MLMDP:** Jin et al. (2020) proposed the aforementioned LSVI-UCB, which fails to satisfy all of (P1), (P2), and (P3). Zanette et al. (2020a) analyzed the Thompson sampling-based randomized LSVI (Osband

<sup>1</sup>“LSVI” stands for “least-squares value iteration” and “UCB” stands for “upper confidence bound.”

et al., 2019). Their regret bound uses  $\tilde{O}(\cdot)$  notation, so (P1) is a bit ambiguous, but we believe it holds. However, (P2) and (P3) again fail.

**Regarding sample complexity:** Several papers (including Jin et al. (2020)) provide sample complexity bounds, i.e., number of samples to learn an approximately optimal policy, though to our knowledge, all violate (P3). On the other hand, such bounds yield algorithms that satisfy (P2) (and possibly (P1), though not for Jin et al. (2020)): one can just fix the approximately optimal policy after finding it. However, this requires well-behaved initial states, e.g., fixed across episodes as in Jin et al. (2020). In contrast, we allow for arbitrary initial states, as in the regret analyses from Jin et al. (2020); Zanette et al. (2020a).

**Low inherent Bellman error:** Zanette et al. (2020b) proposed the low inherent Bellman error (LIBE) model, which generalizes MLMDPs while retaining a linear flavor. For this model, Zanette et al. (2020b) improved the regret of Jin et al. (2020), while Hu et al. (2021) studied multi-task learning. However, these algorithms are defined in terms of optimization problems but no solutions are provided, so the algorithms lack explicit implementation. Zanette et al. (2020c) proved a sample complexity bound, which yields a regret minimization algorithm that satisfies (P2) but requires i.i.d. initial states. All violate (P3).

**Linear mixture:** The linear mixture model (LMM) assumes the transition kernel is a linear combination of  $d$  known measures (see, e.g., Jia et al. (2020); Modi et al. (2020); Zhang et al. (2021); Zhou et al. (2021a,b)), which is distinct from our Assumption 1. To our knowledge, the only regret bound for misspecified LMMs is from Ayoub et al. (2020); the algorithm satisfies (P1) but violates (P2) and (P3).

**Nonlinear generalizations:** Some nonlinear generalizations of LMDPs have been proposed, such as the case where the state-action value function belongs to a class of bounded eluder dimension (Russo and Van Roy, 2013) or can be represented by a kernel function or neural network. While such generalization is important, these works (see, e.g., Chowdhury and Oliveira (2020); Ishfaq et al. (2021); Kong et al. (2021); Wang et al. (2020a,b); Yang et al. (2020a,b)) fail to improve over Jin et al. (2020); Zanette et al. (2020a) in terms of (P1), (P2), or (P3) (or regret).

**Model selection:** To overcome the fact that existing MLMDP algorithms require  $\varepsilon_{\text{mis}}$  as input, one could alternatively use a model selection algorithm (see, e.g., Cutkosky et al. (2021); Pacchiano et al. (2020a,b)). In our context, these initialize  $M$  base algorithms (e.g., LSVI-UCB) with respective inputs  $\varepsilon_{\text{mis}}(1), \dots, \varepsilon_{\text{mis}}(M)$ . Then at each episode, the bases

compute policies and a master algorithm uses past data to choose one of the policies. To our knowledge, the only explicit results use Zanette et al. (2020b) as the base, which lacks an efficient implementation, and while the resulting master achieves (P3), it violates (P1) and (P2) (see Appendix D.4 in Cutkosky et al. (2021) and Section 6.4 in Pacchiano et al. (2020a)). Another downside is that the master algorithms in Cutkosky et al. (2021); Pacchiano et al. (2020a) require a regret bound for each base, so while (P3) holds, the stronger “parameter free” property that neither  $\varepsilon_{\text{mis}}$  nor regret bounds are known (which we satisfy) is violated. On the other hand, Pacchiano et al. (2020b) do not require such regret bounds, but their paper only considers the special case of MLCBs ( $H = 1$ ).

**MLCBs:** For MLCBs, Gopalan et al. (2016) showed (unmodified) Lin-UCB can achieve sublinear regret when  $\varepsilon_{\text{mis}}$  is very small. Foster and Rakhlin (2020) proved regret bounds more generally but (P3) fails. Foster et al. (2021) provided expected regret bounds for an algorithm that satisfies (P3). As mentioned above, Takemura et al. (2021)’s algorithm satisfies (P3), and their bounds hold with high probability. In the noncontextual case, Lattimore et al. (2020) proposed an algorithm that achieves (P1) or (P3), but not both. Ghosh et al. (2017) proved regret bounds that may be polynomial in  $|\mathcal{A}|$ .

**Other related work:** Dong et al. (2019) considered a misspecified state aggregation model; their algorithm satisfies (P1) and (P2) but not (P3). Lattimore et al. (2020) proved sample complexity bounds for discounted MDPs where the  $Q$ -function is approximately linear (more general than us), but they require a simulator/generative model. Yin et al. (2021) considered a similar setting, though only requires “local” simulator access. Wang et al. (2021) only assumed the transition kernel is linear but requires a simulator.

## 2 PRELIMINARIES

**Finite-horizon MDP:** We use the standard notation.  $\mathcal{S}$  is the state space,  $\mathcal{A}$  is the finite action space,  $H \in \mathbb{N}$  is the horizon,  $\{r_h\}_{h=1}^H$  are the mean rewards, and  $\{P_h\}_{h=1}^{H-1}$  are the transition kernels. We assume  $r_h(s, a) \in [0, 1]$  for each  $h \in [H] = \{1, \dots, H\}$ ,  $s \in \mathcal{S}$ , and  $a \in \mathcal{A}$ . We let  $\Pi$  denote the set of policies, i.e., the set of sequences  $\pi = (\pi_h)_{h=1}^H$  with  $\pi_h : \mathcal{S} \rightarrow \mathcal{A}$  for each  $h$ . For any  $\pi \in \Pi$  and  $h \in [H]$ ,  $V_h^\pi : \mathcal{S} \rightarrow [H - h + 1]$  denotes the value function

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

where  $s_{h'+1} \sim P_{h'}(\cdot | s_{h'}, \pi_{h'}(s_{h'}))$  for each  $h'$ . We let  $Q_h^\pi : \mathcal{S} \times \mathcal{A} \rightarrow [H - h + 1]$  denote the state-action value

function (or  $Q$ -function) given by

$$Q_h^\pi(s, a) = r_h(s, a) + \mathbb{E}[V_{h+1}^\pi(s_{h+1}) | s_h = s, a_h = a],$$

where  $s_{h+1} \sim P_h(\cdot | s, a)$  and  $V_{H+1}^\pi(s_{H+1}) = 0$  by convention. It is well known that there exists an optimal policy  $\pi^* = (\pi_h^*)_{h=1}^H$ , i.e.,  $V_h^*(s) \triangleq V_h^{\pi^*}(s) = \max_{\pi \in \Pi} V_h^\pi(s)$  for all  $h \in [H]$  and  $s \in \mathcal{S}$ . Also, for any  $h \in [H]$ ,  $\pi_h^*$  is greedy with respect to  $Q_h^* \triangleq Q_h^{\pi^*}$ , i.e.,  $V_h^*(s) = \max_{a \in \mathcal{A}} Q_h^*(s, a)$  for each  $s \in \mathcal{S}$ .

**MLMDP:** As discussed in Section 1, we make the following linearity assumption.

**Assumption 1** (MLMDP). *For some known  $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$  and all  $h \in [H]$ , there exists unknown  $\theta_h \in \mathbb{R}^d$  and  $d$  unknown measures  $\mu_h = (\mu_{h,1}, \dots, \mu_{h,d})$  over  $\mathcal{S}$  such that, for all  $s \in \mathcal{S}$  and  $a \in \mathcal{A}$ ,*

$$\begin{aligned} |r_h(s, a) - \phi(s, a)^\top \theta_h| &\leq \varepsilon_{\text{mis}}, \\ \|P_h(\cdot | s, a) - \phi(s, a)^\top \mu_h\|_1 &\leq \varepsilon_{\text{mis}}. \end{aligned}$$

We also have  $\max_{(s,a) \in \mathcal{S} \times \mathcal{A}} \|\phi(s, a)\|_2 \leq 1$  and  $\max_{h \in [H]} \max\{\|\theta_h\|_2, \int_{s' \in \mathcal{S}} \|\mu_h(s')\|_2\} \leq \sqrt{d}$ .

**Remark 1** (Comparison to prior work). *Assumption 1 matches Assumption B of Jin et al. (2020) and Assumption 1 of Zanette et al. (2020a), except the latter has general  $\ell_2$  norm bounds (e.g.,  $\|\phi(s, a)\|_2 \leq L_\phi$  for some  $L_\phi$ ). Our analysis can be similarly generalized.*

A key consequence is that the  $Q$ -function is approximately linear. (See Proposition 2.3 of Jin et al. (2020) or Corollary B.3 of Zanette et al. (2020a) for a proof.)

**Proposition 1** (MLMDP  $Q$ -function). *For any  $\pi \in \Pi$  and  $h \in [H]$ , there exists  $w_h^\pi \in \mathbb{R}^d$  such that  $|Q_h^\pi(s, a) - \phi(s, a)^\top w_h^\pi| \leq (H - h + 1)\varepsilon_{\text{mis}}$  for any  $(s, a) \in \mathcal{S} \times \mathcal{A}$ .*

**Regret:** We follow the standard episodic framework. At episode  $k$ , we choose policy  $\pi^k = (\pi_h^k)_{h=1}^H$  and begin at an arbitrary initial state  $s_1^k \in \mathcal{S}$ . For each  $h \in [H]$ , we take action  $a_h^k = \pi_h^k(s_h^k)$ , observe noisy reward  $r_h^k(s_h^k, a_h^k) = r_h(s_h^k, a_h^k) + \eta_h^k$  (where  $\eta_h^k$  is conditionally zero-mean noise, i.e.,  $\mathbb{E}[\eta_h^k | s_h^k, a_h^k] = 0$ ), and (when  $h < H$ ) transition to  $s_{h+1}^k \sim P_h(\cdot | s_h^k, a_h^k)$ . We assume  $r_h(s_h^k, a_h^k)$  and  $r_h^k(s_h^k, a_h^k)$  lie in  $[0, 1]$ , so  $\eta_h^k \in [-1, 1]$ .<sup>2</sup> We measure performance in terms of the regret

$$R(K) = \sum_{k=1}^K (V_1^*(s_1^k) - V_1^{\pi^k}(s_1^k)). \quad (1)$$

**MLCB:** For MLCBs, we use the notation above with  $H = 1$  and discard subscripts  $h$ . So, for each  $k \in [K]$ , we observe context  $s^k \in \mathcal{S}$ , take action  $a^k \in \mathcal{A}$ , and receive reward  $r^k(s^k, a^k) = r(s^k, a^k) + \eta^k$ . As above,  $\eta^k$  is conditionally zero-mean,  $r^k(s^k, a^k)$  and  $r(s^k, a^k)$  are  $[0, 1]$ -valued, and  $|r(s, a) - \phi(s, a)^\top \theta| \leq \varepsilon_{\text{mis}}$ .

<sup>2</sup>The results extend with minor modification to bounded mean rewards and subgaussian noise.

---

**Algorithm 1: EXPL3(thres)**


---

```

1  $\Psi^0 = \emptyset$ 
2 for episode  $k = 1, \dots, K$  do
3   Observe  $s^k \in \mathcal{S}$ 
4    $\Lambda^k = I + \sum_{\tau \in \Psi^{k-1}} \phi(s^\tau, a^\tau) \phi(s^\tau, a^\tau)^\top$ 
5    $w^k = (\Lambda^k)^{-1} \sum_{\tau \in \Psi^{k-1}} \phi(s^\tau, a^\tau) r^\tau(s^\tau, a^\tau)$ 
6   if  $\max_{a \in \mathcal{A}} \|\phi(s^k, a)\|_{(\Lambda^k)^{-1}} > \text{thres}$  then
7      $a^k = \arg \max_{a \in \mathcal{A}} \|\phi(s^k, a)\|_{(\Lambda^k)^{-1}}$ 
8      $\Psi^k = \Psi^{k-1} \cup \{k\}$ 
9   else
10     $a^k = \arg \max_{a \in \mathcal{A}} \phi(s^k, a)^\top w^k$ ,  $\Psi^k = \Psi^{k-1}$ 
11   Play  $a^k$ , observe  $r^k(s^k, a^k)$ 

```

---

### 3 MLCB ALGORITHMS

In this section, we restrict to MLCBs and discuss EXPL3 and our Sup-Lin-UCB variant. We will later leverage the insights developed in this section to design our MLMDP algorithm (see, e.g., Remark 4).

**Warm-up: noncontextual, known  $\varepsilon_{\text{mis}}$ :** Assume momentarily that  $s^1 = \dots = s^K$  and  $\varepsilon_{\text{mis}}$  is known. In this case, we propose a simple two stage algorithm. For the first stage, we play the action with the highest noise/uncertainty in order to reduce it. Mathematically, we quantify the uncertainty in the usual way for LCBs: let  $\Lambda^k = I + \sum_{\tau \leq k} \phi(s^\tau, a^\tau) \phi(s^\tau, a^\tau)^\top$ ,

$$\|\phi(s^k, a)\|_{(\Lambda^k)^{-1}} = \sqrt{\phi(s^k, a)^\top (\Lambda^k)^{-1} \phi(s^k, a)}, \quad (2)$$

and play  $a^k = \arg \max_a \|\phi(s^k, a)\|_{(\Lambda^k)^{-1}}$ . After  $\approx \varepsilon_{\text{mis}}^{-2}$  such episodes, the uncertainty falls below  $\varepsilon_{\text{mis}}$ , which means  $|\phi(s^k, a)^\top (w^k - \theta)| = O(\varepsilon_{\text{mis}})$  for the least-squares estimate  $w^k$  of  $\theta$ . Hence,  $\arg \max_a \phi(s^k, a)^\top w^k$  is an  $\varepsilon_{\text{mis}}$ -suboptimal policy, which is the best possible under Assumption 1. Accordingly, the second stage exploits by playing  $a^k = \arg \max_a \phi(s^k, a)^\top w^k$ .

**General case:** For the contextual setting and unknown  $\varepsilon_{\text{mis}}$ , EXPL3 (Algorithm 1) generalizes this approach. Note uncertainty now depends on the context  $s^k$ , and we can no longer define “high” as “ $\geq \varepsilon_{\text{mis}}$ ”. Thus, at episode  $k$ , EXPL3 checks if the uncertainty at the current context  $s^k$  exceeds the input  $\text{thres}$  (Line 6). If so, it explores a high uncertainty action (Line 7); otherwise, it exploits the estimated best action (Line 10). Here  $\Lambda^k$  and  $w^k$  are only computed from episodes  $\Psi^{k-1}$  where EXPL3 explored. When  $\text{thres} \approx \varepsilon_{\text{mis}}$ , EXPL3(thres) is essentially a contextual version of the previous approach and should perform well. However, if  $\text{thres} \not\approx \varepsilon_{\text{mis}}$ , it will fail for one of two reasons:

- (A) If  $\text{thres} \gg \varepsilon_{\text{mis}}$ , EXPL3(thres) stops exploring as soon as the uncertainty falls below  $\text{thres}$ , so it only learns a  $\text{thres}$ -suboptimal policy.

**Algorithm 2:** Sup-Lin-UCB-Var

---

```

1  $\Psi_l^0 = \emptyset \forall l \in [L]$ 
2 for episode  $k = 1, \dots, K$  do
3   Observe  $s^k \in \mathcal{S}$ , set  $\mathcal{A}_1^k = \mathcal{A}$ 
4   for phase  $l = 1, \dots, L$  do
5      $\Lambda_l^k = I + \sum_{\tau \in \Psi_l^{k-1}} \phi(s^\tau, a^\tau) \phi(s^\tau, a^\tau)^\top$ 
6      $w_l^k = (\Lambda_l^k)^{-1} \sum_{\tau \in \Psi_l^{k-1}} \phi(s^\tau, a^\tau) r^\tau(s^\tau, a^\tau)$ 
7     if  $\max_{a \in \mathcal{A}_l^k} \|\phi(s^k, a)\|_{(\Lambda_l^k)^{-1}} > 2^{-l}$  then
8        $a^k = \arg \max_{a \in \mathcal{A}_l^k} \|\phi(s^k, a)\|_{(\Lambda_l^k)^{-1}}$ 
9        $\Psi_l^k = \Psi_l^{k-1} \cup \{k\}$ ,  $\Psi_{l'}^k = \Psi_{l'}^{k-1} \forall l' \neq l$ 
10      break
11     else if  $l < L$  then
12        $\mathcal{A}_{l+1}^k = \{a \in \mathcal{A}_l^k : \phi(s^k, a)^\top w_l^k \geq \max_{a' \in \mathcal{A}_l^k} \phi(s^k, a')^\top w_l^k - \Theta(2^{-l})\}$ 
13     else
14        $a^k = \arg \max_{a \in \mathcal{A}_L^k} \phi(s^k, a)^\top w_L^k$ 
15        $\Psi_{l'}^k = \Psi_{l'}^{k-1} \forall l'$ 
16   Play  $a^k$ , observe  $r^k(s^k, a^k)$ 

```

---

(B) If  $thres \ll \varepsilon_{\text{mis}}$ , **EXPL3**( $thres$ ) explores too much – for roughly  $thres^{-2}$  episodes when  $\varepsilon_{\text{mis}}^{-2}$  suffice. This may cause  $\text{sp}(r)thres^{-2}$  additional regret, where  $\text{sp}(r) = \max_a r(s^k, a) - \min_a r(s^k, a)$ .

**Ensemble approach:** We next show these failures can be overcome by carefully interconnecting the ensemble  $\{\text{EXPL3}(2^{-l})\}_{l=1}^L$ . Roughly, the  $l$ -th will explore until its uncertainty is  $2^{-l}$ , as above. Thereafter, we know its greedy policy is only  $2^{-l}$ -suboptimal, and since  $\varepsilon_{\text{mis}}$  is unknown, we may have  $2^{-l} \gg \varepsilon_{\text{mis}}$ . Hence, instead of choosing the best action – which causes failure (A) – *we only ask it to eliminate  $2^{-l}$ -suboptimal actions*. By the same token, the  $(l-1)$ -th algorithm has already eliminated  $2^{1-l}$ -suboptimal actions, so *the  $l$ -th sees a more favorable problem instance*, with  $\text{sp}(r) = O(2^{-l})$ . Thus, if instead  $2^{-l} \ll \varepsilon_{\text{mis}}$ , the extra regret in (B) scales as  $2^{-l}(2^{-l})^{-2} = 2^l$ , which is tolerable if we choose  $L = \log_2(O(\sqrt{K}))$ .

More precisely, our **Sup-Lin-UCB** variant is given in Algorithm 2. At episode  $k$ , it chooses an action via a phased elimination procedure that lasts at most  $L$  phases. Generalizing **EXPL3**, it computes  $\Lambda_l^k$  and  $w_l^k$  using  $\Psi_l^k$ , which are the exploratory episodes at phase  $l$ . As discussed above, the  $l$ -th phase chooses a high uncertainty action if one exists (Line 8) and otherwise eliminates actions with estimated reward  $\Theta(2^{-l})$  less than the maximal (Line 12). Finally, if phase  $L$  is reached and an action was never chosen in Line 8, it exploits the estimated best action (Line 14).

**Formal interpretation:** To complement this intuition, we provide a formal result. Roughly, it shows

that if  $\mathcal{S}$  is rich enough, then for any phase  $l$  and any contexts in **Sup-Lin-UCB-Var**, there are contexts in **EXPL3**( $2^{-l}$ ) such that the latter learns the same policy as the  $l$ -th phase of **Sup-Lin-UCB-Var**. In other words, **Sup-Lin-UCB-Var** runs  $\{\text{EXPL3}(2^{-l})\}_{l=1}^L$ . In this way, **Sup-Lin-UCB** is akin to model selection, but unlike those approaches, does not attempt to learn the best **EXPL3**( $2^{-l}$ ). See Appendix D.1 for a proof.

**Proposition 2** (**Sup-Lin-UCB** = **EXPL3** ensemble). *Assume that for any  $\{\phi_a\}_{a \in \mathcal{A}} \subset \mathbb{R}^d$  and  $\{r_a\}_{a \in \mathcal{A}} \subset \mathbb{R}$ , there exists  $s \in \mathcal{S}$  such that  $\phi(s, a) = \phi_a$  and  $r(s, a) = r_a \forall a \in \mathcal{A}$ . Then for any  $l \in [L]$  and  $\{s^k\}_{k=1}^K \subset \mathcal{S}$ , there exists  $\{\tilde{s}^k\}_{k=1}^K \subset \mathcal{S}$  such that, if Algorithms 1 and 2 are run with contexts  $\{\tilde{s}^k\}_{k=1}^K$  and  $\{s^k\}_{k=1}^K$ , respectively, and if both face the same noise sequence  $\{\eta^k\}_{k=1}^K$ , then  $(w^k, \Lambda^k) = (w_l^k, \Lambda_l^k) \forall k \in [K]$ .*

**Remark 2** (Historical note). *Sup-Lin-UCB-Var* simplifies Takemura et al. (2021)'s algorithm (see Appendix A), which modifies *Sup-Lin-UCB* (Chu et al., 2011), which builds upon *Sup-Lin-Rel* (Auer, 2002). The latter three set  $L = \log_2(O(\sqrt{K}))$  as above, but we keep it general, which is crucial for (P1) and (P2).

## 4 MLMDP ALGORITHM

We can now leverage the intuition of the previous section to discuss **Sup-LSVI-UCB** (Algorithm 3). To begin, it initializes the aforementioned parameter  $L$ , an exploration parameter  $\alpha$ , a rounding parameter  $\varepsilon_{\text{rnd}}$  (to be discussed shortly), and the sets  $\Psi_{h,l} = \emptyset$  (now indexed by step  $h \in [H]$  but similar to  $\Psi_l$  in Algorithm 2). The  $k$ -th episodes then contains two parts:

- **Policy update (Alg. 3, Lines 4-10):** Starting at  $h = H$  and inducting backward, for each  $l \in [L]$ , Lines 6 and 7 use the episodes  $\Psi_{h,l}^k$  to compute a least-squares estimate  $w_{h,l}^k$  of the vector  $w_h^\pi$  from Proposition 1 (and the matrix  $\Lambda_{h,l}^k$ ). This is the same approach used by **LSVI-UCB**; see Section 4 of Jin et al. (2020) for intuition. The difference is that the next-state value estimate  $V_{h+1}^k(s_{h+1}^\tau)$  in Line 7 is computed via Algorithm 4, which is essentially **Sup-Lin-UCB-Var** and will be discussed soon. In contrast, **LSVI-UCB** uses **Lin-UCB**-style estimates

$$\max_{a \in \mathcal{A}} \left( \phi(s, a)^\top w_{h+1}^k + \alpha \|\phi(s, a)\|_{(\Lambda_{h+1}^k)^{-1}} \right), \quad (3)$$

where  $w_{h+1}^k$  and  $\Lambda_{h+1}^k$  are computed using all data (not just  $\Psi_{h,l}^k$ ). Additionally, Lines 8 and 9 elementwise round  $w_{h,l}^k$  and  $(\Lambda_{h,l}^k)^{-1}$  to  $\tilde{w}_{h,l}^k$  and  $(\tilde{\Lambda}_{h,l}^k)^{-1}$  for reasons discussed in Remark 3.<sup>3</sup>

<sup>3</sup>We emphasize  $(\tilde{\Lambda}_{h,l}^k)^{-1}$  elementwise rounds  $(\Lambda_{h,l}^k)^{-1}$ ; we have *not* defined  $\tilde{\Lambda}_{h,l}^k$  and inverted it.

**Algorithm 3:** Sup-LSVI-UCB( $\varepsilon_{\text{tol}}, \delta$ )

---

```

1  $L = \lceil \log_2(\sqrt{d}/\varepsilon_{\text{tol}}) \rceil$ ,  $\alpha = 42dHL\sqrt{\log(\frac{3dHL}{\delta})}$ 
2  $\varepsilon_{\text{rnd}} = 2^{-4L}/d$ ,  $\Psi_{h,l}^0 = \emptyset \forall l \in [L], h \in [H]$ 
3 for episode  $k = 1, \dots, K$  do
4   for step  $h = H, \dots, 1$  do
5     for phase  $l = 1, \dots, L$  do
6        $\Lambda_{h,l}^k = 16I + \sum_{\tau \in \Psi_{h,l}^{k-1}} \phi(s_h^\tau, a_h^\tau) \phi(s_h^\tau, a_h^\tau)^\top$ 
7        $w_{h,l}^k = (\Lambda_{h,l}^k)^{-1} \sum_{\tau \in \Psi_{h,l}^{k-1}} \phi(s_h^\tau, a_h^\tau) \times$ 
8          $(r_h^\tau(s_h^\tau, a_h^\tau) + V_{h+1}^k(s_h^\tau, a_h^\tau))$ 
9          $(\tilde{\Lambda}_{h,l}^k)^{-1} = \varepsilon_{\text{rnd}} \lceil (\Lambda_{h,l}^k)^{-1} / \varepsilon_{\text{rnd}} \rceil$ 
10         $\tilde{w}_{h,l}^k = \varepsilon_{\text{rnd}} \lceil w_{h,l}^k / \varepsilon_{\text{rnd}} \rceil$ 
11         $(\pi_h^k(\cdot), V_h^k(\cdot), l_h^k(\cdot)) = \text{Sup-Lin-UCB-Var}(\cdot)$ 
12        (see Algorithm 4)
13      Observe  $s_1^k \in \mathcal{S}$ 
14      for step  $h = 1, \dots, H$  do
15        Play  $a_h^k = \pi_h^k(s_h^k)$ , observe  $r_h^k(s_h^k, a_h^k)$ 
16        Transition to  $s_{h+1}^k \sim P_h(\cdot | s_h^k, a_h^k)$ 
17        if  $l_h^k(s_h^k) \leq L$  then
18           $\Psi_{h,l_h^k(s_h^k)}^k = \Psi_{h,l_h^k(s_h^k)}^{k-1} \cup \{k\}$ 
19           $\Psi_{h,l}^k = \Psi_{h,l}^{k-1} \forall l \in [L] \setminus \{l_h^k(s_h^k)\}$ 
20        else  $\Psi_{h,l}^k = \Psi_{h,l}^{k-1} \forall l \in [L]$ 

```

---

- **Policy execution (Alg. 3, Lines 11-18):** After updating the policy, we execute it by computing its relevant entries  $\{\pi_h^k(s_h^k)\}_{h=1}^H$  via Algorithm 4. (We discuss the  $\Psi_{h,l}$  update in Lines 16, 17, and 18 soon.)

**Subroutine (Alg. 4):** As mentioned above, Algorithm 4 implements Sup-Lin-UCB-Var logic to choose  $\pi_h^k(s)$  for a given  $s \in \mathcal{S}$ , though using the rounded  $\tilde{w}_{h,l}^k$  and  $(\tilde{\Lambda}_{h,l}^k)^{-1}$ . A small technical issue is that the analogue of (2) may be ill-defined, so instead we let<sup>4</sup>

$$\|\phi(s, a)\|_{(\tilde{\Lambda}_{h,l}^k)^{-1}} = \sqrt{|\phi(s, a)^\top (\tilde{\Lambda}_{h,l}^k)^{-1} \phi(s, a)|}.$$

Algorithm 4 also returns  $V_h^k(s)$ , which is the value estimate used in the backward induction, and  $l_h^k(s)$ , where  $l_h^k(s) - 1$  is the number of eliminations conducted (see Lines 5 and 13). For the value estimate, we use  $\mathcal{P}_{[0,H]}$  to project onto  $[0, H]$ , i.e.,  $\mathcal{P}_{[0,H]}(x) = 0$ ,  $x$ , and  $H$  when  $x < 0$ ,  $\in [0, H]$ , and  $> H$ , respectively. This is typical for MLMDPs and ensures boundedness of the random variables. See Remark 4 for further discussion of  $V_h^k$ . Also note  $l_h^k(s) \leq L$  only when an exploratory action is chosen in Line 4, in which case Sup-LSVI-UCB adds the episode to  $\Psi_{h,l}$  (Line 16 of Algorithm 3).

**Remark 3 (Rounding).** *The rounding in Lines 8 and 9 of Algorithm 3 ensures that  $V_{h+1}^k : \mathcal{S} \rightarrow [0, H]$  defined*

<sup>4</sup>While this need not be a norm,  $\varepsilon_{\text{rnd}}$  is small enough that it behaves like one (at least enough for our purposes).

**Algorithm 4:** Sup-Lin-UCB-Var( $s$ )

---

```

1  $\mathcal{A}_{h,1}^k(s) = \mathcal{A}$ ,  $V_{h,0}^k(s) = H$ 
2 for phase  $l = 1, \dots, L$  do
3   if  $\max_{a \in \mathcal{A}_{h,l}^k(s)} \|\phi(s, a)\|_{(\tilde{\Lambda}_{h,l}^k)^{-1}} > 2^{-l}$  then
4      $\pi_h^k(s) = \arg \max_{a \in \mathcal{A}_{h,l}^k(s)} \|\phi(s, a)\|_{(\tilde{\Lambda}_{h,l}^k)^{-1}}$ 
5      $V_h^k(s) = \mathcal{P}_{[0,H]}(V_{h,l-1}^k(s) + 2^{1-l}\alpha)$ ,  $l_h^k(s) = l$ 
6     return  $(\pi_h^k(s), V_h^k(s), l_h^k(s))$ 
7   else if  $l < L$  then
8      $\pi_{h,l}^k(s) = \arg \max_{a \in \mathcal{A}_{h,l}^k(s)} \phi(s, a)^\top \tilde{w}_{h,l}^k$ 
9      $V_{h,l}^k(s) = \phi(s, \pi_{h,l}^k(s))^\top \tilde{w}_{h,l}^k$ 
10     $\mathcal{A}_{h,l+1}^k(s) = \{a \in \mathcal{A}_{h,l}^k(s) : \phi(s, a)^\top \tilde{w}_{h,l}^k \geq$ 
11       $V_{h,l}^k(s) - 2^{1-l}\alpha\}$ 
12    else
13       $\pi_h^k(s) = \pi_{h,L}^k(s)$ 
14       $V_h^k(s) = \mathcal{P}_{[0,H]}(V_{h,L}^k(s))$ ,  $l_h^k(s) = L + 1$ 
15    return  $(\pi_h^k(s), V_h^k(s), l_h^k(s))$ 

```

---

by Algorithm 4 belongs to a finite function class. This enables a union bound over the function class in our concentration lemma, which is needed because  $V_{h+1}^k$  is a random function that depends on past data. In contrast, Jin et al. (2020) shows (3) is close to a function class with a small covering number, then takes a union bound over the cover. This relies on the fact that their  $\mathcal{S}$ -dimensional value function estimate is itself a continuous function of the poly( $d$ )-dimensional  $w_h^k$  and  $\Lambda_h^k$ . In our case, Lines 3 and 10 of Algorithm 4 introduce discontinuities that cause this to fail.

**Remark 4 (Off-policy estimates).** *Algorithm 4's (unprojected) value estimate takes one of two forms. In Line 13, it is  $\phi(s, \pi_h^k(s))^\top \tilde{w}_{h,L}^k$ , which is the Q-function estimate at the chosen action  $\pi_h^k(s)$ . In this sense, it is the usual "on-policy" estimate used in LSVI-UCB and most other algorithms. In contrast, Line 5 uses*

$$\phi(s, \pi_{h,l-1}^k(s))^\top \tilde{w}_{h,l-1}^k + \Theta(2^{-l}). \quad (4)$$

Since  $\pi_{h,l-1}^k(s) \neq \pi_h^k(s)$  in general, this estimate is off-policy. We use such estimates in Line 5 because Lines 4-6 correspond to the explicit exploration discussed in Section 3. When such exploration is needed, the on-policy estimate may be quite low, so we use an off-policy estimate to incentivize the algorithm to visit  $s$  (after which we can conduct the exploration). The  $2^{-l}$  term in (4) is also motivated by Section 3. In essence, since  $\text{EXPL}(2^{-l})$  learns up to  $2^{-l}$  noise, it is an uncertainty bonus that makes (4) the highest statistically-plausible value from the perspective of  $\text{EXPL}(2^{-l})$ .

**Remark 5 (Adversarial corruptions).** *Lykouris et al. (2021) assume the MDP is linear ( $\varepsilon_{\text{mis}} = 0$ ) except for a small number of episodes, where it changes arbitrarily. While quite different from MLMDP, they also*

combine backward induction and phased elimination, with each phase using a subset of episodes. The commonality is high level, though; in that work, each phase runs its own backward induction subroutine, which episodes each phase uses differ, and the algorithm is model-based (among other differences).

## 5 MAIN RESULTS

We can now present Theorem 1, which provides regret and complexity guarantees for Sup-LSVI-UCB. This is our most general result, and we will soon examine some special cases of  $\varepsilon_{\text{tol}}$  to build further intuition. For now, we mention that the theorem (and the fact that  $\varepsilon_{\text{mis}}$  does not appear in our algorithms) guarantees (P1), (P2), and (P3) hold when  $\varepsilon_{\text{tol}}$  is independent of  $K$ , which (to our knowledge) is a first for MLMDPs.

**Theorem 1** (General result). *If Assumption 1 holds and we run Algorithm 3 with inputs  $\varepsilon_{\text{tol}} \in (0, 1)$  and  $\delta \in (0, 1)$ , then with probability at least  $1 - \delta$ ,*

$$R(K) = O\left(\sqrt{d^3 H^4 \min\{(d/\varepsilon_{\text{tol}})^2, K\} \log^5\left(\frac{d}{\varepsilon_{\text{tol}}}\right)\iota} + \sqrt{H^3 K \iota} + \sqrt{d} H^2 K \max\{\varepsilon_{\text{mis}}, \varepsilon_{\text{tol}}\} \sqrt{\log^3\left(\frac{d}{\varepsilon_{\text{tol}}}\right)\iota}\right),$$

where  $\iota = \log(3dH/\delta)$ . Furthermore, Algorithm 3's space complexity is

$$O\left(d^2 H \log\left(\frac{d}{\varepsilon_{\text{tol}}}\right) + dH|\mathcal{A}| \min\left\{d^2 \log\left(\frac{d}{\varepsilon_{\text{tol}}}\right)/\varepsilon_{\text{tol}}^2, K\right\}\right), \quad (5)$$

and its per-episode time complexity is

$$O\left(d^2 H |\mathcal{A}| \min\left\{d^2 \log\left(\frac{d}{\varepsilon_{\text{tol}}}\right)/\varepsilon_{\text{tol}}^2, K\right\} \log\left(\frac{d}{\varepsilon_{\text{tol}}}\right)\right). \quad (6)$$

Alternatively, if  $K$  is small (so (P1) and (P2) are less relevant), we can choose  $\varepsilon_{\text{tol}}$  in terms of  $K$  to obtain the following corollary. Here (P3) still holds, though (like many existing algorithms) (P1) and (P2) fail.

**Corollary 1** (Unknown  $\varepsilon_{\text{mis}}$ , small  $K$ ). *If Assumption 1 holds and we run Algorithm 3 with inputs  $\varepsilon_{\text{tol}} = \frac{d}{\sqrt{K}}$  and  $\delta \in (0, 1)$ , then with probability at least  $1 - \delta$ ,*

$$R(K) = O\left(\sqrt{d^3 H^4 K \log^5(K)\iota} + \sqrt{d} H^2 K \varepsilon_{\text{mis}} \sqrt{\log^3(K)\iota}\right),$$

where  $\iota = \log(3dH/\delta)$ . Furthermore, Algorithm 3's space complexity is  $O(d^2 H \log(K) + dHK|\mathcal{A}|)$  and its per-episode time complexity is  $O(d^2 HK|\mathcal{A}| \log(K))$ .

Finally, given knowledge of  $\varepsilon_{\text{mis}}$  as in Jin et al. (2020); Zanette et al. (2020a), we can set  $\varepsilon_{\text{mis}} = \varepsilon_{\text{tol}}$  to ensure (P1) and (P2) hold with the following regret bound.

**Corollary 2** (Known  $\varepsilon_{\text{mis}}$ , large  $K$ ). *If Assumption 1 holds and we run Algorithm 3 with inputs  $\varepsilon_{\text{tol}} = \varepsilon_{\text{mis}}$  and  $\delta \in (0, 1)$ , then with probability at least  $1 - \delta$ ,*

$$R(K) = O\left(\sqrt{d^3 H^4 \min\{(d/\varepsilon_{\text{mis}})^2, K\} \log^5\left(\frac{d}{\varepsilon_{\text{mis}}}\right)\iota} + \sqrt{H^3 K \iota} + \sqrt{d} H^2 K \varepsilon_{\text{mis}} \sqrt{\log^3\left(\frac{d}{\varepsilon_{\text{mis}}}\right)\iota}\right),$$

where  $\iota = \log(3dH/\delta)$ . Furthermore, Algorithm 3's space and per-episode time complexities are bounded by (5) and (6), respectively, with  $\varepsilon_{\text{tol}}$  replaced by  $\varepsilon_{\text{mis}}$ .

**Remark 6** (Comparison to prior work). *Up to log factors, the “linear” terms in the corollaries improve the results from Jin et al. (2020); Zanette et al. (2020a) by a  $\sqrt{d}$  factor; the sublinear term in Corollary 1 matches those results, while Corollary 2 improves them for  $K \geq (d/\varepsilon_{\text{tol}})^2$  (see Table 1). The complexity bounds in Corollary 1 match Jin et al. (2020); Zanette et al. (2020a), while Corollary 2 improves these bounds (also for  $K \geq (d/\varepsilon_{\text{tol}})^2$ ). The regret bound in Zanette et al. (2020b) beats ours by a  $\sqrt{d}$  factor in the sublinear term, but it is unclear how to implement their algorithm efficiently (again, see Table 1).*

**Remark 7** (Linear term). *Again neglecting logs, the linear terms in the corollaries are  $\sqrt{d} H^2 K \varepsilon_{\text{mis}}$ . Lattimore et al. (2020) shows the  $\sqrt{d}$  “blow-up” is unavoidable and conjectures that for  $\gamma$ -discounted MDPs,  $(1 - \gamma)^{-2}$  dependence (the analogue of  $H^2$  in the discounted setting) is optimal.*

**Remark 8** (Choice of  $L$ ). *When  $\varepsilon_{\text{tol}} = \frac{d}{\sqrt{K}}$  in Corollary 1,  $L = \log_2\left(\frac{K}{\sqrt{d}}\right)$  in Algorithm 4, which is the choice used in Takemura et al. (2021). When  $\varepsilon_{\text{tol}} = \varepsilon_{\text{mis}}$  in Corollary 2,  $L = \log_2\left(\frac{\sqrt{d}}{\varepsilon_{\text{mis}}}\right)$ , which roughly means the ensemble  $\{\text{EXPL3}(2^{-l})\}_{l=1}^L$  explores until the lowest noise level matches the misspecification bias.*

**Remark 9** (MLCB). *When  $H = 1$  and  $L = \log_2\left(\frac{K}{\sqrt{d}}\right)$ , Corollary 1 can be sharpened to  $d\sqrt{K} + \sqrt{d} K \varepsilon_{\text{mis}}$  (see Remark 11 in Appendix B), which matches Takemura et al. (2021)'s result when  $|\mathcal{A}|$  is large.*

We prove Theorem 1 in Appendix B. At a high level, the proof generalizes that of Takemura et al. (2021)'s Theorem 1 from MLCBs to MLMDPs so is structurally similar. However, there are some key differences and challenges that are unique to the MLMDP setting:

- As discussed in Remark 3, we use rounding to cope with dependent noise that arises when  $H > 1$ ; see Lemma 1 for details. The downside is that rounding introduces additional errors. However, we generalize the proof in such a way that we can define an “effective misspecification” that accounts for both rounding error and misspecification, and that plays the same role the latter alone plays in Takemura et al. (2021) (see discussion preceding Lemma 2).

- As discussed in Remark 4, we need to ensure the backward induction uses optimistic estimates  $V_h^k(s)$  despite the fact that the algorithm occasionally takes exploratory (non-optimistic) actions. We show in Lemmas 3 and 4 that our definition of  $V_h^k(s)$  in Algorithm 4 judiciously balances two competing forces: estimating the value of the algorithm’s policy (i.e., ensuring  $V_h^k \approx V_h^{\pi^k}$ ) and remaining optimistic with respect to the optimal policy (i.e.,  $V_h^* \approx V_h^k$ ), which together imply low regret (i.e.,  $V_h^* \approx V_h^{\pi^k}$ ).
- In general, we are more careful with log terms that Takemura et al. (2021) simply bound by  $\log K$  (see, e.g., discussion before Claim 3), as this leads to super-linear bounds that our analysis avoids.
- Along these lines, the proof shows that at each phase of Algorithm 4, the misspecification may cause all  $\varepsilon_{\text{mis}}$ -suboptimal actions to be eliminated (see Claim 6, which generalizes Takemura et al. (2021)’s Lemma 4). Hence, after  $\Omega(L)$  phases, Algorithm 4 may recommend  $\Omega(L\varepsilon_{\text{mis}})$ -suboptimal actions, which leads to super-linear regret bounds when  $L$  grows with  $K$ . This is why we need to choose  $\varepsilon_{\text{tol}}$  (and subsequently  $L$ ) independent of  $K$  in order to achieve (P1) in Theorem 1 and Corollary 2.
- Finally, when  $H = 1$ , Takemura et al. (2021) separately bounds regret when (1)  $l_1^k(s_1^k) = 1$ , (2)  $l_1^k(s_1^k) \in \{2, \dots, L\}$ , and (3)  $l_1^k(s_1^k) = L + 1$  (though they do not use this notation). For general  $H \in \mathbb{N}$ , we have an entire sequence  $\{l_h^k(s_h^k)\}_{h=1}^H$ , which renders this case-based analysis intractable. Instead, we streamline their approach by showing (1) never occurs for our parameter choices (see Corollary 6) and by introducing the  $l_h^k$  notation to treat (2) and (3) in a more unified manner (see, e.g., Claim 5).

## 6 OTHER RESULTS

Finally, we return to discuss Lin-UCB. Recall we assume  $\eta^k$  is zero-mean,  $r(s^k, a^k)$  and  $r^k(s^k, a^k) = r(s^k, a^k) + \eta^k$  lie in  $[0, 1]$  (so  $\eta^k \in [-1, 1]$ ),  $|r(s, a) - \phi(s, a)^\top \theta| \leq \varepsilon_{\text{mis}}$ , and  $\|\phi(s, a)\|_2 \leq 1$ ; we strengthen the assumption  $\|\theta\|_2 \leq \sqrt{d}$  to  $\|\theta\|_2 \leq 1$  in this section.<sup>5</sup> For this setting, the regret definition (1) simplifies to

$$R(K) = \sum_{k=1}^K (r(s^k, a_\star^k) - r(s^k, a^k)),$$

where  $a_\star^k = \arg \max_{a \in \mathcal{A}} r(s^k, a)$  and  $a^k$  is the chosen action. In our notation, Lin-UCB chooses

$$a^k = \arg \max_{a \in \mathcal{A}} (\phi(s^k, a)^\top w^k + \alpha \|\phi(s^k, a)\|_{(\Lambda^k)^{-1}}), \quad (7)$$

<sup>5</sup>To prove (8), Lattimore et al. (2020) assumes 1-subgaussian noise (see their Section 5),  $\|\phi(s, a)\|_2 \leq 1$ , and  $|\phi(s, a)^\top \theta| \leq 1$  (see their Appendix E), which is similar.

where  $\Lambda^k = \lambda I + \sum_{\tau=1}^{k-1} \phi(s^\tau, a^\tau) \phi(s^\tau, a^\tau)^\top$  and  $w^k = (\Lambda^k)^{-1} \sum_{\tau=1}^{k-1} \phi(s^\tau, a^\tau) r^\tau(s^\tau, a^\tau)$ . Lattimore et al. (2020) (building upon Jin et al. (2020)) show that choosing  $\delta \in (0, 1)$ ,  $\lambda = 1$ , and  $\alpha = O(\sqrt{d \log(K/\delta)} + \sqrt{K} \varepsilon_{\text{mis}})$  ensures that with probability at least  $1 - \delta$ ,<sup>6</sup>

$$R(K) = O\left(d\sqrt{K} \log(K/\delta) + \sqrt{d} K \varepsilon_{\text{mis}} \sqrt{\log K}\right) \quad (8)$$

(see their Lemma E.1). While  $\sqrt{d} K \varepsilon_{\text{mis}} \sqrt{\log K}$  is optimal up to the log term, it violates (P1). We show this can be remedied (and, when  $K \gg \varepsilon_{\text{mis}}^{-2}$ , (8) improved) by choosing a different regularizer  $\lambda$ .

**Proposition 3.** *Let  $\delta \in (0, 1)$ ,  $\lambda = 1 + K\varepsilon_{\text{mis}}^2$ , and  $\alpha = 1 + \sqrt{2d \log((\lambda + K)/(\lambda\delta))} + 2\sqrt{K} \varepsilon_{\text{mis}}$ . Under the assumptions of Section 6, with probability at least  $1 - \delta$ , Lin-UCB (7) satisfies*

$$R(K) = O\left(d\sqrt{K} \log(\min\{K, \varepsilon_{\text{mis}}^{-2}\}/\delta) + \sqrt{d} K \varepsilon_{\text{mis}} \sqrt{\log(\min\{K, \varepsilon_{\text{mis}}^{-2}\})}\right). \quad (9)$$

*Proof idea.* The proof is standard; see Appendix D.2 for a sketch. The key step is the sum-of-bonuses calculation, which (ignoring lower order terms) shows

$$\alpha \sum_{k=1}^K \|\phi(s^k, a^k)\|_{(\Lambda^k)^{-1}} = O\left(\varepsilon_{\text{mis}} K \sqrt{d \log(K/\lambda)}\right).$$

Hence, when  $\lambda = O(1)$  and  $\lambda = O(K\varepsilon_{\text{mis}}^2)$ , respectively, we obtain a super-linear term like (8) and a linear term like (9), respectively.  $\square$

**Remark 10 (Intuition).** *Consider the case  $\mathcal{S} = \{1\}$ ,  $\mathcal{A} = \{1, \dots, d\}$ , and  $\phi(1, i) = e_i$  (the  $i$ -th standard basis vector). For  $\lambda = O(1)$  and  $\alpha = \Omega(\sqrt{K} \varepsilon_{\text{mis}})$  as in Lattimore et al. (2020); Jin et al. (2020), we have  $(\alpha \|\phi(1, i)\|_{(\Lambda^k)^{-1}})^2 = O(K\varepsilon_{\text{mis}}^2/N_k(i))$ , where  $N_k(i)$  is the number of times action  $i$  was played in the first  $k$  episodes. By (7), this means the algorithm needs to explore uniformly for the first  $\Omega(K\varepsilon_{\text{mis}}^2)$  steps to drive the exploration bonuses down to  $O(1)$ . In contrast,  $(\alpha \|\phi(1, i)\|_{(\Lambda^k)^{-1}})^2 = O(K\varepsilon_{\text{mis}}^2/(K\varepsilon_{\text{mis}}^2 + N_k(i))) = O(1)$  holds right away (i.e., for  $k = 1$ ) with our  $\lambda$ .*

## 7 DISCUSSION

Before closing, we discuss some additional important aspects of our algorithms.

**Practicality of Sup-Lin-UCB and Sup-LSVI-UCB:** Recall from the introduction that, in the setting of MLCBs with unknown  $\varepsilon_{\text{mis}}$ , Lattimore et al. (2020)

<sup>6</sup>The lemma actually bounds  $\mathbb{E}[R(K)]$  for a refined algorithm, but (8) can be similarly proven for (7).



devised a problem instance for which Lin-UCB incurs  $\Omega(K)$  regret and Takemura et al. (2021) showed Sup-Lin-UCB can achieve  $K\varepsilon_{\text{mis}}\text{polylog}(K)$  regret. For MLMDPs, the instance of Lattimore et al. (2020) similarly shows that LSVI-UCB can incur  $\Omega(K)$  regret and we devised Sup-LSVI-UCB and proved it has  $K\max\{\varepsilon_{\text{mis}}, \varepsilon_{\text{tol}}\}$  regret. While this suggests the superiority of Sup-Lin-UCB and Sup-LSVI-UCB over Lin-UCB and LSVI-UCB, respectively, the latter two algorithms are much simpler and easier to deploy in practice. Thus, an important open problem is to devise algorithms with the strong theoretical guarantees of the former but the practical advantages of the latter. Alternatively, it would be interesting to understand the minimal assumptions under which the latter obtain the guarantees of the former (namely, assumptions which rule out the  $\Omega(K)$  instance from Lattimore et al. (2020), which is somewhat pathological).

**Empirical evaluation:** As shown in Table 1, our algorithm incurs lower regret than those from Jin et al. (2020); Zanette et al. (2020a), and it would be useful to see if this improvement also occurs empirically. Unfortunately, both of these existing algorithms have parameters that the papers only specify in a big- $O$  sense ( $\beta$  in Jin et al. (2020) and  $\sigma$  in Zanette et al. (2020a)), so it is unclear how to choose these parameters to meaningfully compare the algorithms to ours. Of course, one would typically optimize the parameters for all algorithms, then use the optimized parameters to compare. But in our case,  $\beta$  and  $\sigma$  are defined in terms of  $\varepsilon_{\text{mis}}$ , so this hyperparameter optimization implicitly exploits knowledge of  $\varepsilon_{\text{mis}}$ . In short, since the existing work does not provide guidance on how to choose their parameters, it is unclear how to choose them without exploiting knowledge of  $\varepsilon_{\text{mis}}$ . This in turn renders a comparison to our algorithm (which does not require knowledge of  $\varepsilon_{\text{mis}}$  for the theoretical results) less meaningful. Nevertheless, we believe numerical comparison is important and leave it as an important point for future work.

## 8 CONCLUSION

In this work, we proposed the Sup-LSVI-UCB algorithm and showed it is the first to achieve (P1), (P2), and (P3). Our algorithm is motivated by a new interpretation of Sup-Lin-UCB, which also helps explain the results of Takemura et al. (2021) intuitively. Additionally, we improved existing regret bounds for MLMDPs when only (P3), or (P1) and (P2), are required. We also showed Lin-UCB can be improved in terms of (P1) when  $\varepsilon_{\text{mis}}$  is known, which should extend to LSVI-UCB.

*Broader societal impact:* While this work is theoretical, it is motivated by the very practical issue of model

misspecification. In practice, automated decision making algorithms like ours should be actively monitored to mitigate the risk of biased decisions (which may arise from biased training data, for example).

## Acknowledgements

This work was partially supported by ONR Grant N00014-19-1-2566, ARO Grant ARO W911NF-19-1-0379, NSF/USDA Grant AG 2018-67007-28379, and NSF Grants 1910112, 2019844, 2112471, 1704970, and 1934986.

## References

- Yasin Abbasi-Yadkori, David Pal, and Csaba Szepesvari. Improved algorithms for linear stochastic bandits. In *NIPS*, volume 11, pages 2312–2320, 2011.
- 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 Yang. Model-based reinforcement learning with value-targeted regression. In *International Conference on Machine Learning*, pages 463–474. PMLR, 2020.
- Sayak Ray Chowdhury and Rafael Oliveira. No-regret reinforcement learning with value function approximation: A kernel embedding approach. *arXiv preprint arXiv:2011.07881*, 2020.
- 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*, pages 208–214. JMLR Workshop and Conference Proceedings, 2011.
- Ashok Cutkosky, Christoph Dann, Abhimanyu Das, Claudio Gentile, Aldo Pacchiano, and Manish Purohit. Dynamic balancing for model selection in bandits and rl. In *International Conference on Machine Learning*, pages 2276–2285. PMLR, 2021.
- Shi Dong, Benjamin Van Roy, and Zhengyuan Zhou. Provably efficient reinforcement learning with aggregated states. *arXiv preprint arXiv:1912.06366*, 2019.
- Dylan Foster and Alexander Rakhlin. Beyond ucb: Optimal and efficient contextual bandits with regression oracles. In *International Conference on Machine Learning*, pages 3199–3210. PMLR, 2020.
- Dylan J Foster, Claudio Gentile, Mehryar Mohri, and Julian Zimmert. Adapting to misspecification in contextual bandits. *arXiv preprint arXiv:2107.05745*, 2021.

- Avishek Ghosh, Sayak Ray Chowdhury, and Aditya Gopalan. Misspecified linear bandits. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.
- Aditya Gopalan, Odalric-Ambrym Maillard, and Mohammadi Zaki. Low-rank bandits with latent mixtures. *arXiv preprint arXiv:1609.01508*, 2016.
- Jiachen Hu, Xiaoyu Chen, Chi Jin, Lihong Li, and Liwei Wang. Near-optimal representation learning for linear bandits and linear rl. In *International Conference on Machine Learning*, pages 4349–4358. PMLR, 2021.
- Haqeeb Ishaq, Qiwen Cui, Viet Nguyen, Alex Ayoub, Zhuoran Yang, Zhaoran Wang, Doina Precup, and Lin F Yang. Randomized exploration for reinforcement learning with general value function approximation. *arXiv preprint arXiv:2106.07841*, 2021.
- Zeyu Jia, Lin Yang, Csaba Szepesvari, and Mengdi Wang. Model-based reinforcement learning with value-targeted regression. In *Learning for Dynamics and Control*, pages 666–686. PMLR, 2020.
- Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pages 2137–2143. PMLR, 2020.
- Dingwen Kong, Ruslan Salakhutdinov, Ruosong Wang, and Lin F Yang. Online sub-sampling for reinforcement learning with general function approximation. *arXiv preprint arXiv:2106.07203*, 2021.
- Tor Lattimore, Csaba Szepesvari, and Gellert Weisz. Learning with good feature representations in bandits and in rl with a generative model. In *International Conference on Machine Learning*, pages 5662–5670. PMLR, 2020.
- Thodoris Lykouris, Max Simchowitz, Alex Slivkins, and Wen Sun. Corruption-robust exploration in episodic reinforcement learning. In *Conference on Learning Theory*, pages 3242–3245. PMLR, 2021.
- 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*, pages 2010–2020. PMLR, 2020.
- Ian Osband, Benjamin Van Roy, Daniel J Russo, and Zheng Wen. Deep exploration via randomized value functions. *Journal of Machine Learning Research*, 20(124):1–62, 2019.
- Aldo Pacchiano, Christoph Dann, Claudio Gentile, and Peter Bartlett. Regret bound balancing and elimination for model selection in bandits and rl. *arXiv preprint arXiv:2012.13045*, 2020a.
- Aldo Pacchiano, My Phan, Yasin Abbasi-Yadkori, Anup Rao, Julian Zimmert, Tor Lattimore, and Csaba Szepesvari. Model selection in contextual stochastic bandit problems. *Advances in Neural Information Processing Systems*, 33, 2020b.
- Daniel Russo and Benjamin Van Roy. Eluder dimension and the sample complexity of optimistic exploration. In *NIPS*, pages 2256–2264. Citeseer, 2013.
- Kei Takemura, Shinji Ito, Daisuke Hatano, Hanna Sumita, Takuro Fukunaga, Naonori Kakimura, and Ken-ichi Kawarabayashi. A parameter-free algorithm for misspecified linear contextual bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 3367–3375. PMLR, 2021.
- Bingyan Wang, Yuling Yan, and Jianqing Fan. Sample-efficient reinforcement learning for linearly-parameterized mdps with a generative model. *arXiv preprint arXiv:2105.14016*, 2021.
- Ruosong Wang, Russ R Salakhutdinov, and Lin Yang. Reinforcement learning with general value function approximation: Provably efficient approach via bounded eluder dimension. *Advances in Neural Information Processing Systems*, 33, 2020a.
- Yining Wang, Ruosong Wang, Simon Shaolei Du, and Akshay Krishnamurthy. Optimism in reinforcement learning with generalized linear function approximation. In *International Conference on Learning Representations*, 2020b.
- Zhuoran Yang, Chi Jin, Zhaoran Wang, Mengdi Wang, and Michael Jordan. Provably efficient reinforcement learning with kernel and neural function approximations. *Advances in Neural Information Processing Systems*, 33, 2020a.
- Zhuoran Yang, Chi Jin, Zhaoran Wang, Mengdi Wang, and Michael I Jordan. On function approximation in reinforcement learning: Optimism in the face of large state spaces. *arXiv preprint arXiv:2011.04622*, 2020b.
- Dong Yin, Botao Hao, Yasin Abbasi-Yadkori, Nevena Lazić, and Csaba Szepesvári. Efficient local planning with linear function approximation. *arXiv preprint arXiv:2108.05533*, 2021.
- Andrea Zanette, David Brandfonbrener, Emma Brunskill, Matteo Pirota, and Alessandro Lazaric. Frequentist regret bounds for randomized least-squares value iteration. In *International Conference on Artificial Intelligence and Statistics*, pages 1954–1964. PMLR, 2020a.
- Andrea Zanette, Alessandro Lazaric, Mykel Kochenderfer, and Emma Brunskill. Learning near optimal policies with low inherent bellman error. In *International Conference on Machine Learning*, pages 10978–10989. PMLR, 2020b.

Andrea Zanette, Alessandro Lazaric, Mykel J Kochenderfer, and Emma Brunskill. Provably efficient reward-agnostic navigation with linear value iteration. *arXiv preprint arXiv:2008.07737*, 2020c.

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. *arXiv preprint arXiv:2101.12745*, 2021.

Dongruo Zhou, Quanquan Gu, and Csaba Szepesvari. Nearly minimax optimal reinforcement learning for linear mixture markov decision processes. In *Conference on Learning Theory*, pages 4532–4576. PMLR, 2021a.

Dongruo Zhou, Jiafan He, and Quanquan Gu. Provably efficient reinforcement learning for discounted mdps with feature mapping. In *International Conference on Machine Learning*, pages 12793–12802. PMLR, 2021b.

---

## Supplementary Material: Improved Algorithms for Misspecified Linear Markov Decision Processes

---

### A Takemura et al. (2021)'s ALGORITHM

Algorithm 5 is the Sup-Lin-UCB variant from Takemura et al. (2021) (in our notation). The key differences from Algorithm 2 are twofold. First, at each phase  $l$ , Algorithm 5 either chooses an optimistic action (Line 7), eliminates suboptimal actions (Line 9), or chooses an exploratory action (Line 11). In contrast, Algorithm 2 either explores (Line 8) or eliminates (Line 12) for phases  $l < L$  and either explores (Line 8) or exploits (Line 14) in phase  $l = L$ . Second, Algorithm 5 uses Lin-UCB-style exploration bonuses in Lines 7 and 9, which the corresponding lines of Algorithm 2 do not. In both cases, we made these changes to simplify the algorithm and unify the presentation with EXPL3, and we found this does not worsen regret in an order sense.

---

#### Algorithm 5: Sup-Lin-UCB-Var

---

```

1  $\Psi_l^0 = \emptyset \forall l \in [L]$ 
2 for episode  $k = 1, \dots, K$  do
3   Observe  $s^k \in \mathcal{S}$ , set  $l = 1$  and  $\mathcal{A}_l^k = \mathcal{A}$ 
4   repeat
5      $\Lambda_l^k = I + \sum_{\tau \in \Psi_l^{k-1}} \phi(s^\tau, a^\tau) \phi(s^\tau, a^\tau)^\top$ ,  $w_l^k = (\Lambda_l^k)^{-1} \sum_{\tau \in \Psi_l^{k-1}} \phi(s^\tau, a^\tau) r^\tau(s^\tau, a^\tau)$ 
6     if  $\max_{a \in \mathcal{A}_l^k} \|\phi(s^k, a)\|_{(\Lambda_l^k)^{-1}} \leq \sqrt{d/K}$  then
7        $a^k = \arg \max_{a \in \mathcal{A}_l^k} (\phi(s^k, a)^\top w_l^k + \alpha \|\phi(s^k, a)\|_{(\Lambda_l^k)^{-1}})$ ,  $\Psi_{l'}^k = \Psi_{l'}^{k-1} \forall l'$ 
8     else if  $\max_{a \in \mathcal{A}_l^k} \|\phi(s^k, a)\|_{(\Lambda_l^k)^{-1}} \leq 2^{-l}$  then
9        $\mathcal{A}_{l+1}^k = \{a \in \mathcal{A}_l^k : \phi(s^k, a)^\top w_l^k + \alpha \|\phi(s^k, a)\|_{(\Lambda_l^k)^{-1}} \geq$   

         $\max_{a' \in \mathcal{A}_l^k} (\phi(s^k, a')^\top w_l^k + \alpha \|\phi(s^k, a')\|_{(\Lambda_l^k)^{-1}}) - 2^{1-l} \alpha\}$ 
10    else
11     $a^k \in \{a \in \mathcal{A}_l^k : \|\phi(s^k, a)\|_{(\Lambda_l^k)^{-1}} > 2^{-l}\}$ ,  $\Psi_l^k = \Psi_l^{k-1} \cup \{k\}$ ,  $\Psi_{l'}^k = \Psi_{l'}^{k-1} \forall l' \neq l$ 
12  until  $a^k$  is chosen
13  Play  $a^k$ , observe  $r^k(s^k, a^k)$ 

```

---

### B THEOREM 1 PROOF

In this appendix, we prove Theorem 1. We begin with some basic inequalities in Appendix B.1. We then prove our main concentration result in Appendix B.2. Next, Appendix B.3 provides a general result for the  $Q$ -function estimates in Algorithm 4. Using this result, Appendices B.4 and B.5 bound the differences  $V_h^k(s) - V_h^{\pi^k}(s)$  and  $V_h^*(s) - V_h^k(s)$ , respectively. This yields a bound on the episode  $k$  regret  $V_1^*(s_1^k) - V_1^{\pi^k}(s_1^k)$ , which we use in Appendix B.6 to prove the regret guarantee. Along the way, we defer some proof details to Appendix C, which also contains the complexity analysis.

#### B.1 Simple results

We first bound the error that arises from the rounding performed in Algorithm 3.

**Claim 1** (Rounding error). *For any  $k \in [K]$ ,  $h \in [H]$ ,  $l \in [L]$ ,  $s \in \mathcal{S}$ , and  $a \in \mathcal{A}$ , we have*

$$|\phi(s, a)^\top (w_{h,l}^k - \tilde{w}_{h,l}^k)| \leq \sqrt{d\varepsilon_{rd}}, \quad \left| \|\phi(s, a)\|_{(\Lambda_{h,l}^k)^{-1}} - \|\phi(s, a)\|_{(\tilde{\Lambda}_{h,l}^k)^{-1}} \right| \leq \sqrt{d\varepsilon_{rd}}.$$

Next, we have the following bounds for the bonus terms.

**Claim 2** (Bonuses). *For any  $k \in [K]$ ,  $h \in [H]$ ,  $s \in \mathcal{S}$ ,  $l \in [k_h(s) - 1]$ , and  $a \in \mathcal{A}_{h,l}^k(s)$ , we have  $\|\phi(s, a)\|_{(\tilde{\Lambda}_{h,l}^k)^{-1}} \leq 2^{-l}$  and  $\|\phi(s, a)\|_{(\Lambda_{h,l}^k)^{-1}} \leq 2^{1-l}$ .*

*Proof.* The first bound holds by definition in Algorithm 4. For the second bound, we use the first, Claim 1, and  $\varepsilon_{\text{rnd}} \leq 2^{-2l}/d$  in Algorithm 3 to obtain  $\|\phi(s, a)\|_{(\Lambda_{h,l}^k)^{-1}} \leq \|\phi(s, a)\|_{(\tilde{\Lambda}_{h,l}^k)^{-1}} + 2^{-l} \leq 2^{-l} + 2^{-l} = 2^{1-l}$ .  $\square$

Finally, we bound the cardinality of  $\Psi_{h,l}^k$ . This is an analogue of Takemura et al. (2021)'s Lemma 1, which shows  $|\Psi_{h,l}^k| = O(4^l d \log(K/d))$ . With a more careful argument, we obtain a bound that is independent of  $K$  (for any fixed  $l$ ), which will be crucial in achieving (P1) and (P2).

**Claim 3** (Dataset bound). *For any  $k \in [K]$ ,  $h \in [H]$ , and  $l \in [L]$ , we have  $|\Psi_{h,l}^k| \leq 40 \cdot 4^l d l \leq 2^{3l+5} d$ .*

*Proof sketch.* By Lines 4-5 of Algorithm 4 and Line 16 of Algorithm 3, for each  $\tau \in \Psi_{h,l}^k$ , we know that  $\|\phi(s_h^\tau, a_h^\tau)\|_{(\tilde{\Lambda}_{h,l}^\tau)^{-1}} > 2^{-l}$ . This implies  $|\Psi_{h,l}^k| \leq 4^l \sum_{\tau \in \Psi_{h,l}^k} \|\phi(s_h^\tau, a_h^\tau)\|_{(\tilde{\Lambda}_{h,l}^\tau)^{-1}}^2 = \tilde{O}(4^l d)$ , where the equality follows from Claim 1 and Abbasi-Yadkori et al. (2011). See Appendix B for details.  $\square$

## B.2 Concentration

For any  $k \in [K]$ ,  $h \in [H]$ ,  $l \in [L]$ , and  $V : \mathcal{S} \rightarrow \mathbb{R}$ , define the bad event

$$\mathcal{B}(k, h, l, V) = \left\{ \left\| \sum_{\tau \in \Psi_{h,l}^k} \phi(s_h^\tau, a_h^\tau) \left( \eta_h^\tau + V(s_{h+1}^\tau) - \mathbb{E}_{s_{h+1}^\tau} V(s_{h+1}^\tau) \right) \right\|_{(\Lambda_{h,l}^k)^{-1}} > \beta \right\},$$

where  $\mathbb{E}_{s_{h+1}^\tau} V(s_{h+1}^\tau) = \int_{s' \in \mathcal{S}} V(s') P_h(s' | s_h^\tau, a_h^\tau)$  only averages over  $s_{h+1}^\tau$  (even if  $V$  is random, in particular, if  $V = V_{h+1}^k$ ) and  $\beta = 13dHL\sqrt{\log(3dHL/\delta)}$ . Also define the good event

$$\mathcal{G} = \cap_{k=1}^K \cap_{h=1}^H \cap_{l=1}^L \mathcal{B}(k, h, l, V_{h+1}^k)^C.$$

As discussed in Remark 3, a similar event is analyzed in Jin et al. (2020) using covering arguments. In contrast, here  $V_{h+1}^k$  belongs to a finite function class, which allows us to show that  $\mathcal{G}$  occurs with high probability via a direct union bound over the function class.

**Lemma 1** (Concentration). *The good event  $\mathcal{G}$  occurs with probability at least  $1 - \delta/2$ .*

*Proof.* We fix  $h \in [H]$  and  $l \in [L]$  and show  $\mathbb{P}(\cup_{k=1}^K \mathcal{B}(k, h, l, V_{h+1}^k)) \leq \frac{\delta}{2HL}$ , which (by the union bound) completes the proof. Toward this end, we introduce some notation. For any (ordered) sets  $\mathcal{X} = \{x_{l'}\}_{l'=1}^L \subset \mathbb{R}^d$  and  $\mathcal{Y} = \{Y_{l'}\}_{l'=1}^L \subset \mathbb{R}^{d \times d}$ , let  $V_{\mathcal{X}, \mathcal{Y}} : \mathcal{S} \rightarrow [0, H]$  be the function that results from running Algorithm 4 with  $\tilde{w}_{h+1, l'}^k$  and  $(\tilde{\Lambda}_{h+1, l'}^k)^{-1}$  replaced by  $x_{l'}$  and  $Y_{l'}$ , respectively. Hence, if  $\mathcal{X} = \{\tilde{w}_{h+1, l'}^k\}_{l'=1}^L$  and  $\mathcal{Y} = \{(\tilde{\Lambda}_{h+1, l'}^k)^{-1}\}_{l'=1}^L$ , then  $V_{\mathcal{X}, \mathcal{Y}} = V_{h+1}^k$ . Next, define

$$\begin{aligned} \mathcal{X} &= \{[\varepsilon_{\text{rnd}} i_j]_{j=1}^d : i_j \in \{-\lceil (2^L dH)^4 / \varepsilon_{\text{rnd}} \rceil, \dots, \lceil (2^L dH)^4 / \varepsilon_{\text{rnd}} \rceil\} \forall j \in [d]\} \subset \mathbb{R}^d, \\ \mathcal{Y} &= \{[\varepsilon_{\text{rnd}} i_{j_1, j_2}]_{j_1, j_2=1}^d : i_{j_1, j_2} \in \{-\lceil 1/(16\varepsilon_{\text{rnd}}) \rceil, \dots, \lceil 1/(16\varepsilon_{\text{rnd}}) \rceil\} \forall (j_1, j_2) \in [d]^2\} \subset \mathbb{R}^{d \times d}. \end{aligned}$$

By Claim 10 in Appendix B (which shows  $\|w_{h,l}^k\|_\infty \leq (2^L dH)^4$ ), we have  $\{\tilde{w}_{h+1, l'}^k\}_{l'=1}^L \in \mathcal{X}^L$ . Further, by a standard matrix norm inequality and the fact that the eigenvalues of  $\Lambda_{h+1, l'}^k$  are at least 16, we have  $\max_{j_1, j_2} |(\Lambda_{h+1, l'}^k)^{-1}|_{j_1, j_2} \leq \|(\Lambda_{h+1, l'}^k)^{-1}\|_2 \leq 1/16$ , so  $\{(\tilde{\Lambda}_{h+1, l'}^k)^{-1}\}_{l'=1}^L \in \mathcal{Y}^L$ . Finally, we know  $V_{h+1}^k : \mathcal{S} \rightarrow [0, H]$  by Lines 5 and 13 of Algorithm 4. Thus, if we define  $\mathcal{V} = \{V_{\mathcal{X}, \mathcal{Y}} : \mathcal{S} \rightarrow [0, H] | \mathcal{X} \in \mathcal{X}^L, \mathcal{Y} \in \mathcal{Y}^L\}$ , then  $V_{h+1}^k \in \mathcal{V} \forall k \in [K]$ , which implies  $\cup_{k=1}^K \mathcal{B}(k, h, l, V_{h+1}^k) \subset \cup_{V \in \mathcal{V}} \cup_{k=1}^K \mathcal{B}(k, h, l, V)$ . Hence, taking another union

bound, it suffices to show that for any  $V \in \mathcal{V}$ ,  $\mathbb{P}(\cup_{k=1}^K \mathcal{B}(k, h, l, V)) \leq \frac{\delta}{2HL|\mathcal{V}|}$ . Let  $\mathbf{1}(\cdot)$  denote the indicator function, and for each  $\tau \in [K]$ , define the following:

$$\underline{s}_{\tau+1} = s_{h+1}^\tau, \quad \underline{\phi}_\tau = \phi(s_h^\tau, a_h^\tau) \mathbf{1}(\tau \in \Psi_{h,l}^\tau), \quad \underline{\Lambda}_k = 16I + \sum_{\tau=1}^{k-1} \underline{\phi}_\tau \underline{\phi}_\tau^\top, \quad v_\tau = \eta_h^\tau + V(\underline{s}_{\tau+1}) - \mathbb{E}_{s_{h+1}^\tau} V(s_{h+1}^\tau).$$

Then by definition, for any  $k \in [K]$ , we have

$$\Lambda_{h,l}^k = \underline{\Lambda}_k, \quad \left\| \sum_{\tau \in \Psi_{h,l}^{k-1}} \phi(s_h^\tau, a_h^\tau) \left( \eta_h^\tau + V(s_{h+1}^\tau) - \mathbb{E}_{s_{h+1}^\tau} V(s_{h+1}^\tau) \right) \right\|_{(\Lambda_{h,l}^k)^{-1}} = \left\| \sum_{\tau=1}^{k-1} \underline{\phi}_\tau v_\tau \right\|_{\underline{\Lambda}_k^{-1}}. \quad (10)$$

Also let  $\mathcal{F}_0 = \emptyset$  and  $\mathcal{F}_\tau = \sigma(\mathcal{F}_{\tau-1} \cup \sigma(s_1^\tau, a_1^\tau, \eta_1^\tau, \dots, s_{h-1}^\tau, a_{h-1}^\tau, \eta_{h-1}^\tau, s_h^\tau, a_h^\tau))$  for each  $\tau \in \mathbb{N}$ , where  $\sigma(\cdot)$  is the generated  $\sigma$ -algebra. Hence, in words,  $\mathcal{F}_\tau$  contains all randomness until the random reward and next state are realized at step  $h$  of episode  $\tau$ . Note  $\underline{\phi}_\tau$  is  $\mathcal{F}_\tau$ -measurable and  $v_\tau$  is  $\mathcal{F}_{\tau+1}$ -measurable with  $\mathbb{E}[v_\tau | \mathcal{F}_\tau] = 0$ . Furthermore, since  $\eta_h^\tau \in [-1, 1]$  by assumption (see Section 2) and  $V : \mathcal{S} \rightarrow [0, H]$  by definition, we have  $v_\tau \in [-2H, 2H]$ , so  $v_\tau$  is  $(2H)$ -subgaussian. Therefore,

$$\begin{aligned} \mathbb{P}(\cup_{k=1}^K \mathcal{B}(k, h, l, V)) &= \mathbb{P} \left( \cup_{k=1}^K \left\{ \left\| \sum_{\tau=1}^{k-1} \underline{\phi}_\tau v_\tau \right\|_{\underline{\Lambda}_k^{-1}} > \beta \right\} \right) \\ &\leq \mathbb{P} \left( \cup_{k=1}^K \left\{ \left\| \sum_{\tau=1}^{k-1} \underline{\phi}_\tau v_\tau \right\|_{\underline{\Lambda}_k^{-1}} > \sqrt{8H^2 \log \left( \frac{\det(\Lambda_{h,l}^k) 2HL|\mathcal{V}|}{\det(16I) \delta} \right)} \right\} \right) \leq \frac{\delta}{2HL|\mathcal{V}|}, \end{aligned}$$

where the equality uses (10), the first inequality is a simple calculation (see Claim 11 in Appendix B for details), and the second inequality is Theorem 1 from Abbasi-Yadkori et al. (2011).  $\square$

**Remark 11** (MLCB). When  $H = 1$ , we simply have  $v_\tau = \eta_h^\tau$  in the proof of Lemma 1, so we do not require a union bound over  $\mathcal{V}$ . This union bound makes  $\beta$ , and subsequently  $\alpha$ , have linear (instead of square root) dependence on  $d$ , which in turn gives the  $\sqrt{K}$  term in our regret bound  $d^{3/2}$  (instead of  $d$ ) dependence.

### B.3 Estimation error

For the remainder of the proof, we bound regret on the good event  $\mathcal{G}$ . We first show that on  $\mathcal{G}$ , the least-squares estimate  $w_{h,l}^k$  is close to  $\bar{w}_h^k \triangleq \theta_h + \int_{s' \in \mathcal{S}} V_{h+1}^k(s') \mu_h(s')$  in a certain sense. For this, it will be convenient to introduce the following notation:

$$\Delta_h^r(s, a) = r_h(s, a) - \phi(s, a)^\top \theta_h, \quad \Delta_h^P(s' | s, a) = P_h(s' | s, a) - \phi(s, a)^\top \mu_h.$$

(Note  $|\Delta_h^r(s, a)|, \|\Delta_h^P(\cdot | s, a)\|_1 \leq \varepsilon_{\text{mis}}$  by Assumption 1.) We can now prove a generalization of Takemura et al. (2021)'s Lemma 2 using an approach somewhat similar to Jin et al. (2020)'s Lemma B.4.

**Claim 4** (Least-squares error). *On the event  $\mathcal{G}$ , for any  $k \in [K]$ ,  $h \in [H]$ ,  $s \in \mathcal{S}$ ,  $l \in [l_h^k(s) - 1]$ , and  $a \in \mathcal{A}_{h,l}^k(s)$ , we have  $|\phi(s, a)^\top (w_{h,l}^k - \bar{w}_h^k)| \leq 2^{-l} \alpha + 26H \sqrt{dl} \varepsilon_{\text{mis}}$ .*

*Proof.* By definition of  $\Lambda_{h,l}^k$ , we have

$$\bar{w}_h^k = (\Lambda_{h,l}^k)^{-1} \Lambda_{h,l}^k \bar{w}_h^k = 16(\Lambda_{h,l}^k)^{-1} \bar{w}_h^k + (\Lambda_{h,l}^k)^{-1} \sum_{\tau \in \Psi_{h,l}^{k-1}} \phi(s_h^\tau, a_h^\tau) \phi(s_h^\tau, a_h^\tau)^\top \bar{w}_h^k.$$

By definition of  $\bar{w}_h^k$  and Assumption 1, for any  $\tau \in \Psi_{h,l}^{k-1}$ , we know

$$\phi(s_h^\tau, a_h^\tau)^\top \bar{w}_h^k = r_h(s_h^\tau, a_h^\tau) + \mathbb{E}_{s_{h+1}^\tau} V_{h+1}^k(s_{h+1}^\tau) - \Delta_h^r(s_h^\tau, a_h^\tau) - \int_{s' \in \mathcal{S}} V_{h+1}^k(s') \Delta_h^P(s' | s_h^\tau, a_h^\tau).$$

Additionally, recall that in Algorithm 3, we have

$$w_{h,l}^k = (\Lambda_{h,l}^k)^{-1} \sum_{\tau \in \Psi_{h,l}^{k-1}} \phi(s_h^\tau, a_h^\tau) (r_h(s_h^\tau, a_h^\tau) + \eta_h^\tau + V_{h+1}^k(s_{h+1}^\tau)).$$

It follows that  $\phi(s, a)^\top (w_{h,l}^k - \bar{w}_h^k) = \sum_{i=1}^3 \phi(s, a)^\top (\Lambda_{h,l}^k)^{-1} z_i$ , where we define

$$\begin{aligned} z_1 &= \sum_{\tau \in \Psi_{h,l}^{k-1}} \phi(s_h^\tau, a_h^\tau) \left( \eta_h^\tau + V_{h+1}^k(s_{h+1}^\tau) - \mathbb{E}_{s_{h+1}^\tau} V_{h+1}^k(s_{h+1}^\tau) \right), \quad z_2 = -16\bar{w}_h^k, \\ z_3 &= \sum_{\tau \in \Psi_{h,l}^{k-1}} \phi(s_h^\tau, a_h^\tau) \left( \Delta_h^r(s_h^\tau, a_h^\tau) + \int_{s' \in \mathcal{S}} V_{h+1}^k(s') \Delta_{h,s'}^P(s_h^\tau, a_h^\tau) \right). \end{aligned}$$

Hence, we aim to bound  $|\phi(s, a)^\top (\Lambda_{h,l}^k)^{-1} z_i|$  for each  $i \in [3]$ . By Cauchy-Schwarz, on the event  $\mathcal{G}$ ,

$$|\phi(s, a)^\top (\Lambda_{h,l}^k)^{-1} z_1| \leq \|z_1\|_{(\Lambda_{h,l}^k)^{-1}} \|\phi(s, a)\|_{(\Lambda_{h,l}^k)^{-1}} \leq \beta \|\phi(s, a)\|_{(\Lambda_{h,l}^k)^{-1}}. \quad (11)$$

Again using Cauchy-Schwarz, we have

$$|\phi(s, a)^\top (\Lambda_{h,l}^k)^{-1} z_2| \leq 16 \|\bar{w}_h^k\|_{(\Lambda_{h,l}^k)^{-1}} \|\phi(s, a)\|_{(\Lambda_{h,l}^k)^{-1}} \leq 8\sqrt{d}H \|\phi(s, a)\|_{(\Lambda_{h,l}^k)^{-1}}, \quad (12)$$

where the second inequality holds because, by Claim 9 in Appendix B (a simple norm equivalence), Assumption 1, and the fact that  $V_{h+1}^k : \mathcal{S} \rightarrow [0, H]$  in Algorithm 4,

$$\|\bar{w}_h^k\|_{(\Lambda_{h,l}^k)^{-1}} \leq \|\bar{w}_h^k\|_2 / 4 \leq \sqrt{d}(1+H)/4 \leq \sqrt{d}H/2.$$

For  $z_3$ , first note that by Assumption 1 and Algorithm 4,

$$\left| \Delta_h^r(s_h^\tau, a_h^\tau) + \int_{s' \in \mathcal{S}} V_{h+1}^k(s') \Delta_{h,s'}^P(s_h^\tau, a_h^\tau) \right| \leq \left( 1 + \max_{s' \in \mathcal{S}} V_{h+1}^k(s') \right) \varepsilon_{\text{mis}} \leq (1+H)\varepsilon_{\text{mis}} \leq 2H\varepsilon_{\text{mis}}.$$

Furthermore, since  $l \leq l_h^k(s) - 1$  and  $a \in \mathcal{A}_{h,l}^k(s)$ , we can use Claims 2 and 3 to obtain

$$|\Psi_{h,l}^{k-1}| \times \phi(s, a)^\top (\Lambda_{h,l}^k)^{-1} \phi(s, a) \leq 40 \cdot 4^l dl \times 4^{1-l} = 160dl < 13^2 dl.$$

By the previous two bounds, Cauchy-Schwarz, and positive-semidefiniteness, we obtain

$$\begin{aligned} |\phi(s, a)^\top (\Lambda_{h,l}^k)^{-1} z_3| &\leq 2H\varepsilon_{\text{mis}} \sum_{\tau \in \Psi_{h,l}^{k-1}} |\phi(s, a)^\top (\Lambda_{h,l}^k)^{-1} \phi(s_h^\tau, a_h^\tau)| \\ &\leq 2H\varepsilon_{\text{mis}} \sqrt{|\Psi_{h,l}^{k-1}| \phi(s, a)^\top (\Lambda_{h,l}^k)^{-1} \sum_{\tau \in \Psi_{h,l}^{k-1}} \phi(s_h^\tau, a_h^\tau) \phi(s_h^\tau, a_h^\tau) (\Lambda_{h,l}^k)^{-1} \phi(s, a)} \\ &= 2H\varepsilon_{\text{mis}} \sqrt{|\Psi_{h,l}^{k-1}| \phi(s, a)^\top (\Lambda_{h,l}^k)^{-1} (\Lambda_{h,l}^k - 16I) (\Lambda_{h,l}^k)^{-1} \phi(s, a)} \\ &\leq 2H\varepsilon_{\text{mis}} \sqrt{|\Psi_{h,l}^{k-1}| \phi(s, a)^\top (\Lambda_{h,l}^k)^{-1} \phi(s, a)} \leq 26H\varepsilon_{\text{mis}} \sqrt{dl}. \end{aligned} \quad (13)$$

Hence, combining (11), (12), and (13), we obtain

$$|\phi(s, a)^\top (w_{h,l}^k - \bar{w}_h^k)| \leq \sum_{i=1}^3 |\phi(s, a)^\top (\Lambda_{h,l}^k)^{-1} z_i| \leq (\beta + 8\sqrt{d}H) \|\phi(s, a)\|_{(\Lambda_{h,l}^k)^{-1}} + 26H\varepsilon_{\text{mis}} \sqrt{dl}.$$

This completes the proof, because  $\|\phi(s, a)\|_{(\Lambda_{h,l}^k)^{-1}} \leq 2^{1-l}$  by Claim 2, and by definition,

$$\beta + 8\sqrt{d}H = 13dHL\sqrt{\log(3dHL/\delta)} + 8\sqrt{d}H \leq 21dHL\sqrt{\log(3dHL/\delta)} = \alpha/2. \quad \square$$

We conclude this subsection by using Claim 4 to show the  $Q$ -function estimates in Algorithm 4 are close to the function  $\bar{Q}_h^k : \mathcal{S} \times \mathcal{A} \rightarrow [0, 2H]$  defined by

$$\bar{Q}_h^k(s, a) = r_h(s, a) + \int_{s' \in \mathcal{S}} V_{h+1}^k(s') P_h(s'|s, a) \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}.$$

It will also be convenient to define  $\varepsilon_{\text{eff}} = 2\alpha\sqrt{d\varepsilon_{\text{rnd}}} + 28H\sqrt{dL}\varepsilon_{\text{mis}}$ , which is the effective misspecification (true misspecification and rounding error) that we carry through the next portion of the proof.

**Lemma 2** ( $Q$ -function error). *On the event  $\mathcal{G}$ , for any  $k \in [K]$ ,  $h \in [H]$ ,  $s \in \mathcal{S}$ ,  $l \in [l_h^k(s) - 1]$ , and  $a \in \mathcal{A}_{h,l}^k(s)$ , we have  $|\phi(s, a)^\top \tilde{w}_{h,l}^k - \bar{Q}_h^k(s, a)| \leq 2^{-l}\alpha + \varepsilon_{\text{eff}}$ .*

*Proof.* By the triangle inequality, we have

$$|\phi(s, a)^\top \tilde{w}_{h,l}^k - \bar{Q}_h^k(s, a)| \leq |\phi(s, a)^\top (\tilde{w}_{h,l}^k - w_{h,l}^k)| + |\phi(s, a)^\top (w_{h,l}^k - \bar{w}_h^k)| + |\phi(s, a)^\top \bar{w}_h^k - \bar{Q}_h^k(s, a)|.$$

For the first term, by Claim 1, and since  $\alpha \geq 1$  and  $\varepsilon_{\text{rnd}} \in (0, 1)$ ,

$$|\phi(s, a)^\top (\tilde{w}_{h,l}^k - w_{h,l}^k)| \leq \sqrt{d}\varepsilon_{\text{rnd}} \leq \alpha\sqrt{d\varepsilon_{\text{rnd}}}.$$

For the second term, by Claim 4 and since  $l \leq l_h^k(s) - 1 \leq L$ ,

$$|\phi(s, a)^\top (w_{h,l}^k - \bar{w}_h^k)| \leq 2^{-l}\alpha + 26H\sqrt{dL}\varepsilon_{\text{mis}} \leq 2^{-l}\alpha + 26H\sqrt{dL}\varepsilon_{\text{mis}}$$

For the third term, by Assumption 1 and Algorithm 4,

$$|\bar{Q}_h^k(s, a) - \phi(s, a)^\top \bar{w}_h^k| = \left| \Delta_h^r(s, a) + \int_{s' \in \mathcal{S}} \Delta_h^P(s'|s, a) V_{h+1}^k(s') \right| \leq 2H\varepsilon_{\text{mis}} \leq 2H\sqrt{dL}\varepsilon_{\text{mis}}.$$

Hence, combining all of the above, we obtain

$$|\phi(s, a)^\top \tilde{w}_{h,l}^k - \bar{Q}_h^k(s, a)| \leq 2^{-l}\alpha + 2\alpha\sqrt{d\varepsilon_{\text{rnd}}} + 28H\sqrt{dL}\varepsilon_{\text{mis}} = 2^{-l}\alpha + \varepsilon_{\text{eff}}. \quad \square$$

#### B.4 Algorithm policy error

Our next goal is to bound the difference between the value function estimate  $V_h^k$  and the true value function  $V_h^{\pi^k}$  of the algorithm's policy. We begin with an intermediate result. This is roughly an analogue of Takemura et al. (2021)'s Lemma 5 and 8, though our streamlined approach yields a single result. Additionally, we have to deal with the projection in Algorithm 4, which complicates the proof.

**Claim 5** (Algorithm error, one-step). *On the event  $\mathcal{G}$ , for any  $k \in [K]$ ,  $h \in [H]$ , and  $s \in \mathcal{S}$ , we have  $V_h^k(s) - \bar{Q}_h^k(s, \pi_h^k(s)) \leq 8\alpha \cdot 2^{-l_h^k(s)} + \varepsilon_{\text{eff}}$ .*

*Proof.* Let  $l = l_h^k(s) - 1$ . By Corollary 6 from Appendix B, we know that  $l \in [L]$ . We first assume  $l \in [L - 1]$ , which implies  $V_h^k(s) = \mathcal{P}_{[0,H]}(V_{h,l}^k(s) + 2^{-l}\alpha)$  in Algorithm 4. Hence, if  $V_{h,l}^k(s) < -2^{-l}\alpha$ , then  $V_h^k(s) = 0$ , which immediately yields the desired bound (since  $\bar{Q}_h^k(s, \pi_h^k(s)) \geq 0$ ). If instead  $V_{h,l}^k(s) \geq -2^{-l}\alpha$ , then  $V_h^k(s) \leq V_{h,l}^k(s) + 2^{-l}\alpha$ , so it suffices to prove the bound with  $V_h^k(s)$  replaced by  $V_{h,l}^k(s) + 2^{-l}\alpha$ . Toward this end, first observe that by Lemma 2 and since  $\pi_h^k(s) \in \mathcal{A}_{h,l+1}^k(s) \subset \mathcal{A}_{h,l}^k(s)$ ,

$$\phi(s, \pi_h^k(s))^\top \tilde{w}_{h,l}^k \leq \bar{Q}_h^k(s, \pi_h^k(s)) + 2^{-l}\alpha + \varepsilon_{\text{eff}}.$$

On the other hand, again using  $\pi_h^k(s) \in \mathcal{A}_{h,l+1}^k(s)$ , we know

$$V_{h,l}^k(s) + 2^{-l}\alpha \leq \phi(s, \pi_h^k(s))^\top \tilde{w}_{h,l}^k + 2^{1-l}\alpha + 2^{-l}\alpha.$$

Hence, combining the inequalities, we obtain

$$V_{h,l}^k(s) + 2^{-l}\alpha - \bar{Q}_h^k(s, \pi_h^k(s)) \leq 2^{1-l}\alpha + 2^{-l}\alpha + 2^{-l}\alpha + \varepsilon_{\text{eff}} = 8\alpha \cdot 2^{-l_h^k(s)} + \varepsilon_{\text{eff}}.$$



For  $l = L$ , we have  $V_h^k(s) = \mathcal{P}_{[0,H]}(V_{h,L}^k(s))$  in Algorithm 4. If  $V_{h,L}^k(s) < 0$ , the bound is again immediate. If instead  $V_{h,L}^k(s) \geq 0$ , we know  $V_h^k(s) \leq V_{h,L}^k(s)$ , so we can prove the bound with  $V_h^k(s)$  replaced by  $V_{h,L}^k(s)$ . By Algorithm 4 and Lemma 2, we have

$$V_{h,L}^k(s) - \bar{Q}_h^k(s, \pi_h^k(s)) = \phi(s, \pi_h^k(s))^\top \tilde{w}_{h,L}^k - \bar{Q}_h^k(s, \pi_h^k(s)) \leq 2^{-L}\alpha + \varepsilon_{\text{eff}} < 8\alpha \cdot 2^{-l_h^k(s)} + \varepsilon_{\text{eff}},$$

where the final inequality is  $2^{-L} = 2^{1-l_h^k(s)} < 8 \cdot 2^{-l_h^k(s)}$ .  $\square$

Next, for any  $k \in [K]$  and  $h \in [H]$ , define the martingale noise term

$$\gamma_h^k = \mathbb{E}_{s_{h+1}^k} (V_{h+1}^k(s_{h+1}^k) - V_{h+1}^{\pi_k}(s_{h+1}^k)) - (V_{h+1}^k(s_{h+1}^k) - V_{h+1}^{\pi_k}(s_{h+1}^k)).$$

Using the previous claim and a simple inductive argument, we can prove the following lemma (see Appendix B for details). In essence, similar to Jin et al. (2020)'s Lemma B.6, this lemma shows that the noise  $\gamma_h^k$  in the backward induction yields a martingale difference sequence.

**Lemma 3** (Algorithm error, multi-step). *On the event  $\mathcal{G}$ , for any  $k \in [K]$  and any  $h \in [H]$ , we have*

$$V_h^k(s_h^k) - V_h^{\pi_k}(s_h^k) \leq 8\alpha \sum_{h'=h}^H 2^{-l_{h'}^k(s_{h'}^k)} + \sum_{h'=h}^H \gamma_{h'}^k + (H-h+1)\varepsilon_{\text{eff}}. \quad (14)$$

## B.5 Optimal policy error

Next, we bound the difference between the optimal value function  $V_h^*$  and the value function estimate  $V_h^k$ . We start with two intermediate results. First, for each  $k \in [K]$ ,  $h \in [H]$ ,  $s \in \mathcal{S}$ , and  $l \in [l_h^k(s) \wedge L]$ , let  $\bar{\pi}_{h,l}^k(s) = \arg \max_{a \in \mathcal{A}_{h,l}^k(s)} \bar{Q}_h^k(s, a)$ . We can then generalize Takemura et al. (2021)'s Lemma 4.

**Claim 6** (Error across phases). *On the event  $\mathcal{G}$ , for any  $k \in [K]$ ,  $h \in [H]$ ,  $s \in \mathcal{S}$ , and  $l \in [l_h^k(s) \wedge L]$ , we have  $\bar{Q}_h^k(s, \bar{\pi}_{h,l}^k(s)) - \bar{Q}_h^k(s, \bar{\pi}_{h,l}^k(s)) \leq 2(l-1)\varepsilon_{\text{eff}}$ .*

*Proof.* We use induction on  $l$ . For  $l = 1$ , the bound holds with equality. Assuming it holds for  $l \in [(l_h^k(s) \wedge L) - 1]$ , we prove it for  $l + 1$ . By the inductive hypothesis, it suffices to show

$$\bar{Q}_h^k(s, \bar{\pi}_{h,l}^k(s)) - \bar{Q}_h^k(s, \bar{\pi}_{h,l+1}^k(s)) \leq 2\varepsilon_{\text{eff}}. \quad (15)$$

If  $\bar{\pi}_{h,l}^k(s) \in \mathcal{A}_{h,l+1}^k(s)$ , then since  $\mathcal{A}_{h,l+1}^k(s) \subset \mathcal{A}_{h,l}^k(s)$ , we have  $\bar{\pi}_{h,l+1}^k(s) = \bar{\pi}_{h,l}^k(s)$  by definition, so (15) is immediate. Hence, it only remains to prove (15) when  $\bar{\pi}_{h,l}^k(s) \notin \mathcal{A}_{h,l+1}^k(s)$ . Since  $\bar{\pi}_{h,l}^k(s) \in \mathcal{A}_{h,l}^k(s)$  in Algorithm 4, the definition of  $\bar{\pi}_{h,l+1}^k(s)$  implies

$$\bar{Q}_h^k(s, \bar{\pi}_{h,l}^k(s)) - \bar{Q}_h^k(s, \bar{\pi}_{h,l+1}^k(s)) \leq \bar{Q}_h^k(s, \bar{\pi}_{h,l}^k(s)) - \bar{Q}_h^k(s, \pi_{h,l}^k(s)).$$

By  $l \leq l_h^k(s) - 1$ , Lemma 2, the assumption that  $\bar{\pi}_{h,l}^k(s) \notin \mathcal{A}_{h,l+1}^k(s)$ , and Algorithm 4, we have

$$\bar{Q}_h^k(s, \bar{\pi}_{h,l}^k(s)) - \bar{Q}_h^k(s, \pi_{h,l}^k(s)) \leq \phi(s, \bar{\pi}_{h,l}^k(s))^\top \tilde{w}_{h,l}^k - \phi(s, \pi_{h,l}^k(s))^\top \tilde{w}_{h,l}^k + 2^{1-l}\alpha + 2\varepsilon_{\text{eff}} < 2\varepsilon_{\text{eff}}.$$

Combining the previous two inequalities, we obtain the desired bound (15).  $\square$

As an immediate corollary, we have the following.

**Corollary 3** (Error across phases). *On the event  $\mathcal{G}$ , for any  $k \in [K]$ ,  $h \in [H]$ ,  $s \in \mathcal{S}$ , and  $l \in [l_h^k(s) \wedge L]$ , we have  $\bar{Q}_h^k(s, \bar{\pi}_{h,l}^k(s)) - \bar{Q}_h^k(s, \bar{\pi}_{h,l}^k(s)) \leq 2(L-l)\varepsilon_{\text{eff}}$ .*

We can now bound the difference between  $V_h^*$  and  $V_h^k$  in terms of the difference at the next step, i.e., between  $V_{h+1}^*$  and  $V_{h+1}^k$ . This is similar in spirit to Jin et al. (2020)'s Lemma B.5.

**Claim 7** (Optimal error, one-step). *On the event  $\mathcal{G}$ , for any  $k \in [K]$ ,  $h \in [H]$ , and  $s \in \mathcal{S}$ ,*

$$V_h^*(s) - V_h^k(s) \leq \max \left\{ \max_{a \in \mathcal{A}} \int_{s' \in \mathcal{S}} (V_{h+1}^*(s') - V_{h+1}^k(s')) P_h(s'|s, a) + 2L\varepsilon_{\text{eff}} + 2^{-L}\alpha, 0 \right\}.$$

*Proof.* Let  $l = l_h^k(s) - 1 \in [L]$ . Note that if (1)  $l \in [L-1]$  and  $V_{h,l}^k(s) + 2^{-l}\alpha > H$  or (2)  $l = L$  and  $V_{h,l}^k(s) > H$ , then  $V_h^k(s) = H$ , so the bound  $V_h^*(s) - V_h^k(s) \leq 0$  holds. Hence, we assume for the remainder of the proof that either (3)  $l \in [L-1]$  and  $V_{h,l}^k(s) + 2^{-l}\alpha \leq H$  or (4)  $l = L$  and  $V_{h,l}^k(s) \leq H$ . By definition of  $V_h^*$  and  $\bar{Q}_h^k$ , we have

$$\begin{aligned} V_h^*(s) &= \bar{Q}_h^k(s, \pi_h^*(s)) + \int_{s' \in \mathcal{S}} (V_{h+1}^*(s') - V_{h+1}^k(s')) P_h(s'|s, \pi_h^*(s)) \\ &\leq \bar{Q}_h^k(s, \pi_h^*(s)) + \max_{a \in \mathcal{A}} \int_{s' \in \mathcal{S}} (V_{h+1}^*(s') - V_{h+1}^k(s')) P_h(s'|s, a) \end{aligned}$$

By definition of  $\bar{\pi}_{h,1}^k(s)$ , since  $\mathcal{A}_{h,1}^k(s) = \mathcal{A}$ , and by Corollary 3,

$$\bar{Q}_h^k(s, \pi_h^*(s)) \leq \bar{Q}_h^k(s, \bar{\pi}_{h,1}^k(s)) \leq \bar{Q}_h^k(s, \bar{\pi}_{h,l}^k(s)) + 2(L-1)\varepsilon_{\text{eff}}.$$

Again using the definition of  $\bar{\pi}_{h,l}^k(s)$ , along with Lemma 2, we know that

$$\bar{Q}_h^k(s, \bar{\pi}_{h,l}^k(s)) = \max_{a \in \mathcal{A}_{h,l}^k(s)} \bar{Q}_h^k(s, a) \leq \max_{a \in \mathcal{A}_{h,l}^k(s)} \phi(s, a)^\top \tilde{w}_{h,l}^k + 2^{-l}\alpha + \varepsilon_{\text{eff}} = V_{h,l}^k(s) + 2^{-l}\alpha + \varepsilon_{\text{eff}}.$$

Hence, stringing together the inequalities, we obtain

$$V_h^*(s) \leq \max_{a \in \mathcal{A}} \int_{s' \in \mathcal{S}} (V_{h+1}^*(s') - V_{h+1}^k(s')) P_h(s'|s, a) + V_{h,l}^k(s) + 2^{-l}\alpha + 2L\varepsilon_{\text{eff}}. \quad (16)$$

Now in case (3), we have  $V_h^k(s) = \mathcal{P}_{[0,H]}(V_{h,l}^k(s) + 2^{-l}\alpha)$  and  $V_{h,l}^k(s) + 2^{-l}\alpha \leq H$ , which together imply

$$V_{h,l}^k(s) + 2^{-l}\alpha \leq V_h^k(s) < V_h^k(s) + 2^{-L}\alpha. \quad (17)$$

In case (4), we have  $V_h^k(s) = \mathcal{P}_{[0,H]}(V_{h,l}^k(s))$  and  $V_{h,l}^k(s) \leq H$ , which implies  $V_{h,l}^k(s) \leq V_h^k(s)$ . Hence, because  $L = l$  in case (4), we again have (17). Therefore, combining (16) and (17), we obtain

$$V_h^*(s) - V_h^k(s) \leq \max_{a \in \mathcal{A}} \int_{s' \in \mathcal{S}} (V_{h+1}^*(s') - V_{h+1}^k(s')) P_h(s'|s, a) + 2L\varepsilon_{\text{eff}} + 2^{-L}\alpha. \quad \square$$

Finally, a simple inductive argument yields the following. See Appendix B for details.

**Lemma 4** (Optimal error, multi-step). *On the event  $\mathcal{G}$ , for any  $k \in [K]$ ,  $h \in [H]$ , and  $s \in \mathcal{S}$ , we have  $V_h^*(s) - V_h^k(s) \leq (2L\varepsilon_{\text{eff}} + 2^{-L}\alpha)(H - h + 1)$ .*

## B.6 Regret bound

First observe that by Algorithms 3-4 and Corollary 6 from Appendix C, for any  $h \in [H]$ , we have

$$\begin{aligned} \sum_{k=1}^K 2^{-l_h^k(s_h^k)} &= \sum_{k=1}^K \sum_{l=2}^{L+1} 2^{-l} \mathbf{1}(l_h^k(s_h^k) = l) = \sum_{l=2}^L 2^{-l} \sum_{k=1}^K \mathbf{1}(l_h^k(s_h^k) = l) + 2^{-(L+1)} \sum_{k=1}^K \mathbf{1}(l_h^k(s_h^k) = L+1) \\ &= \sum_{l=2}^L 2^{-l} |\Psi_{h,l}^K| + 2^{-(L+1)} \sum_{k=1}^K \mathbf{1}(l_h^k(s_h^k) = L+1) \leq \sum_{l=2}^L 2^{-l} |\Psi_{h,l}^K| + 2^{-(L+1)} K. \end{aligned}$$

Combined with Lemmas 3 and 4, on the event  $\mathcal{G}$ , we obtain

$$\begin{aligned} R(K) &\leq 8\alpha \sum_{h=1}^H \sum_{k=1}^K 2^{-l_h^k(s_h^k)} + \sum_{k=1}^K \sum_{h=1}^H \gamma_h^k + (2^{-L}\alpha + 3L\varepsilon_{\text{eff}})HK \\ &\leq 8\alpha \sum_{h=1}^H \sum_{l=2}^L 2^{-l} |\Psi_{h,l}^K| + \sum_{k=1}^K \sum_{h=1}^H \gamma_h^k + (2^{2-L}\alpha + 2^{-L}\alpha + 3L\varepsilon_{\text{eff}})HK. \end{aligned} \quad (18)$$

For the first summation, by Claim 3 and a simple geometric series computation, for any  $h \in [H]$ ,

$$\sum_{l=2}^L 2^{-l} |\Psi_{h,l}^K| \leq 40dL \sum_{l=2}^L 2^{-l} \cdot 4^l = 40dL \sum_{l=2}^L 2^l < 80 \cdot 2^L dL.$$

Alternatively, we can use Cauchy-Schwarz and Claim 3 to obtain

$$\sum_{l=2}^L 2^{-l} |\Psi_{h,l}^K| = \sum_{l=2}^L 2^{-l} \sqrt{|\Psi_{h,l}^K|} \sqrt{|\Psi_{h,l}^K|} \leq \sqrt{40dL} \sum_{l=1}^L \sqrt{|\Psi_{h,l}^K|} \leq \sqrt{40dL^2 \sum_{l=1}^L |\Psi_{h,l}^K|} \leq \sqrt{40dL^2 K}.$$

Hence, combining the previous two inequalities, we have shown

$$\sum_{l=2}^L 2^{-l} |\Psi_{h,l}^K| \leq \min\{80 \cdot 2^L dL, \sqrt{40dL^2 K}\} = \sqrt{40dL^2 \min\{160 \cdot 2^{2L} d, K\}}. \quad (19)$$

Returning to (18), since  $\{\gamma_h^k\}_{k \in [K] h \in [H]}$  is a martingale difference sequence with  $|\gamma_h^k| \leq 2H$ , the Azuma-Hoeffding inequality implies that with probability at least  $1 - \delta/2$ ,

$$\sum_{k=1}^K \sum_{h=1}^H \gamma_h^k \leq \sqrt{8H^3 K \log(2/\delta)}. \quad (20)$$

For the last term in (18), by definition  $\varepsilon_{\text{eff}} = 2\alpha\sqrt{d\varepsilon_{\text{rnd}}} + 28H\sqrt{dL}\varepsilon_{\text{mis}}$  and  $\varepsilon_{\text{rnd}} = 2^{-4L}/d$ , and since  $L \leq 2^{L-1} \forall L \in \mathbb{N}$ , we have

$$L\varepsilon_{\text{eff}} = 2^{1-2L}L\alpha + 28\sqrt{dL^3}H\varepsilon_{\text{mis}} \leq 2^{-L}\alpha + 28\sqrt{dL^3}H\varepsilon_{\text{mis}}. \quad (21)$$

Hence, when  $\mathcal{G}$  and the event (20) both occur (which happens with probability at least  $1 - \delta$  by Lemma 1 and Azuma-Hoeffding), we can combine (18), (19), (20), and (21) to obtain

$$R(K) = O\left(\alpha\sqrt{dH^2L^2 \min\{2^{2L}d, K\}} + \sqrt{H^3K \log(\delta^{-1})} + 2^{-L}\alpha HK + \sqrt{dL^3}H^2K\varepsilon_{\text{mis}}\right).$$

Recall  $\alpha = 42dHL\sqrt{\log(3dHL/\delta)}$  and  $\iota = \log(3dH/\delta)$ , so  $\alpha = O(\sqrt{d^2H^2L^3\iota})$ . Substituting above,

$$R(K) = O\left(\sqrt{d^3H^4L^5 \min\{2^{2L}d, K\}}\iota + \sqrt{H^3K\iota} + 2^{-L}\sqrt{d^2L^3\iota}H^2K + \sqrt{dL^3}H^2K\varepsilon_{\text{mis}}\right).$$

The regret bound in Theorem 1 follows by definition  $L = \lceil \log_2(\sqrt{d}/\varepsilon_{\text{tol}}) \rceil$ .

## C THEOREM 1 PROOF DETAILS

### C.1 Regret bound details

*Proof of Claim 1.* The first bound follows from Holder's inequality, Algorithm 3, a standard norm equivalence, and Assumption 1:

$$\|\phi(s, a)^\top (w_{h,l}^k - \tilde{w}_{h,l}^k)\| \leq \|\phi(s, a)\|_1 \|w_{h,l}^k - \tilde{w}_{h,l}^k\|_\infty \leq \sqrt{d} \|\phi(s, a)\|_2 \varepsilon_{\text{rnd}} \leq \sqrt{d} \varepsilon_{\text{rnd}}.$$

For the second bound, by similar logic, we have

$$\begin{aligned} |\phi(s, a)^\top ((\Lambda_{h,l}^k)^{-1} - (\tilde{\Lambda}_{h,l}^k)^{-1}) \phi(s, a)| &\leq \sum_{i,j=1}^d |\phi_i(s, a)| |\phi_j(s, a)| |((\Lambda_{h,l}^k)^{-1} - (\tilde{\Lambda}_{h,l}^k)^{-1})_{ij}| \\ &\leq \|\phi(s, a)\|_1^2 \varepsilon_{\text{rnd}} \leq d \|\phi(s, a)\|_2^2 \varepsilon_{\text{rnd}} \leq d \varepsilon_{\text{rnd}}, \end{aligned}$$

which implies that

$$\begin{aligned} \|\phi(s, a)\|_{(\Lambda_{h,l}^k)^{-1}} &\leq \sqrt{|\phi(s, a)^\top (\tilde{\Lambda}_{h,l}^k)^{-1} \phi(s, a)| + |\phi(s, a)^\top ((\Lambda_{h,l}^k)^{-1} - (\tilde{\Lambda}_{h,l}^k)^{-1}) \phi(s, a)|} \\ &\leq \|\phi(s, a)\|_{(\tilde{\Lambda}_{h,l}^k)^{-1}} + \sqrt{d \varepsilon_{\text{rnd}}}. \end{aligned}$$

By symmetry,  $\|\phi(s, a)\|_{(\tilde{\Lambda}_{h,l}^k)^{-1}} \leq \|\phi(s, a)\|_{(\Lambda_{h,l}^k)^{-1}} + \sqrt{d \varepsilon_{\text{rnd}}}$  as well, which completes the proof.  $\square$

**Claim 8.** For any  $k \in [K]$ ,  $h \in [H]$ , and  $l \in [L]$ , we have

$$\sum_{\tau \in \Psi_{h,l}^k} \|\phi(s_h^\tau, a_h^\tau)\|_{(\Lambda_{h,l}^\tau)^{-1}}^2 \leq 2 \log(\det(\Lambda_{h,l}^k) / \det(16I)) \leq 2d \log(1 + |\Psi_{h,l}^k| / (16d)).$$

*Proof.* The first bound is a restatement of Lemma D.2 from Jin et al. (2020). The second follows from Lemma 10 of Abbasi-Yadkori et al. (2011) and Assumption 1.  $\square$

*Proof of Claim 3.* Since  $\Psi_{h,l}^0 \subset \dots \subset \Psi_{h,l}^K$ , it suffices to prove the bound for  $k = K$ . By Algorithm 3,  $\Psi_{h,l}^K$  is the set of episodes  $k \in [K]$  for which  $l_h^k(s_h^k) = l$ . By Algorithm 4,  $l_h^k(s_h^k) = l$  implies that  $\|\phi(s_h^k, a_h^k)\|_{(\bar{\Lambda}_{h,l}^k)^{-1}} > 2^{-l}$ . Combined with Claim 1, and since  $\varepsilon_{\text{rnd}} \leq 2^{-4l}/d \leq 2^{-2(l+1)}/d$  in Algorithm 3, we obtain

$$1 < 2^{2l} \|\phi(s_h^k, a_h^k)\|_{(\bar{\Lambda}_{h,l}^k)^{-1}}^2 \leq 2^{2l+1} \|\phi(s_h^k, a_h^k)\|_{(\Lambda_{h,l}^k)^{-1}}^2 + 2^{2l+1} d \varepsilon_{\text{rnd}} \leq 2^{2l+1} \|\phi(s_h^k, a_h^k)\|_{(\Lambda_{h,l}^k)^{-1}}^2 + 1/2,$$

or, after rearranging,  $1 < 2^{2(l+1)} \|\phi(s_h^k, a_h^k)\|_{(\Lambda_{h,l}^k)^{-1}}^2$ . Combined with Claim 8, we obtain

$$|\Psi_{h,l}^K| = \sum_{k \in \Psi_{h,l}^K} 1 < 2^{2(l+1)} \sum_{k \in \Psi_{h,l}^K} \|\phi(s_h^k, a_h^k)\|_{(\Lambda_{h,l}^k)^{-1}}^2 \leq 2^{2l+3} d \log(1 + |\Psi_{h,l}^K| / (16d)). \quad (22)$$

Multiplying and dividing the right side of (22) by 2, we get

$$|\Psi_{h,l}^K| \leq 2^{2(l+2)} d \log\left(\sqrt{1 + |\Psi_{h,l}^K| / (16d)}\right) \leq 2^{2(l+2)} d \log\left(1 + \sqrt{|\Psi_{h,l}^K| / (16d)}\right) \leq 2^{2(l+1)} \sqrt{d |\Psi_{h,l}^K|},$$

or, after rearranging,

$$|\Psi_{h,l}^K| \leq (2^{2(l+1)} \sqrt{d})^2 = 2^{4(l+1)} d. \quad (23)$$

Plugging (23) into the right side of (22), we obtain

$$|\Psi_{h,l}^K| \leq 2^{2l+3} d \log(1 + 2^{4l}) < 2^{2l+3} d \log(2 \cdot 2^{4l}) \leq 2^{2l+3} d \log 2^{5l} < 2^{2l+3} d \cdot 5l = 40 \cdot 4^l dl.$$

Finally, since  $40 < 64 = 2^6$  and  $l \leq 2^{l-1}$  for any  $l \in \mathbb{N}$ , we have  $40 \cdot 4^l \cdot l < 2^6 \cdot 4^l \cdot 2^{l-1} = 2^{3l+5}$ .  $\square$

**Corollary 4.** For any  $k \in [K]$ ,  $h \in [H]$ , and  $l \in [L]$ , we have  $\det(\Lambda_{h,l}^k) / \det(16I) \leq 2^{5dl}$ .

*Proof.* Combining Claims 8 and 3, we obtain

$$\det(\Lambda_{h,l}^k) / \det(16I) \leq (1 + |\Psi_{h,l}^k| / (16d))^d \leq (1 + 2^{3l+1})^d < (2 \cdot 2^{3l+1})^d = 2^{(3l+2)d} \leq 2^{5dl}. \quad \square$$

**Corollary 5.** For any  $k \in [K]$ ,  $h \in [H]$ , and  $l \in [L]$ , we have  $\sum_{\tau \in \Psi_{h,l}^{k-1}} \|\phi(s_h^\tau, a_h^\tau)\|_{(\Lambda_{h,l}^k)^{-1}} \leq (2^l d)^4$ .

*Proof.* By Cauchy-Schwarz, Claim 3, and Lemma D.1 of Jin et al. (2020),

$$\sum_{\tau \in \Psi_{h,l}^{k-1}} \|\phi(s_h^\tau, a_h^\tau)\|_{(\Lambda_{h,l}^k)^{-1}} \leq \sqrt{|\Psi_{h,l}^{k-1}| \times \sum_{\tau \in \Psi_{h,l}^{k-1}} \|\phi(s_h^\tau, a_h^\tau)\|_{(\Lambda_{h,l}^k)^{-1}}^2} \leq \sqrt{2^{3l+5} d \times d} \leq 2^{4l} d \leq (2^l d)^4. \quad \square$$

**Claim 9.** For any  $k \in [K]$ ,  $h \in [H]$ ,  $l \in [L]$ , and  $x \in \mathbb{R}^d$ , we have  $\|x\|_{(\Lambda_{h,l}^k)^{-1}} \leq \|x\|_2 / 4$ .

*Proof.* Let  $\lambda_i \in [16, \infty)$  and  $q_i \in \mathbb{R}^d$  denote the eigenvalues and eigenvectors of  $\Lambda_{h,l}^k$ . Then

$$\|x\|_{(\Lambda_{h,l}^k)^{-1}}^2 = \sum_{i=1}^d (q_i^\top x)^2 / \lambda_i \leq \sum_{i=1}^d (q_i^\top x)^2 / 16 = \|[q_1 \cdots q_d]^\top x\|_2^2 / 16 = \|x\|_2^2 / 16. \quad \square$$

**Corollary 6.** For any  $k \in [K]$ ,  $h \in [H]$ , and  $s \in \mathcal{S}$ , we have  $l_h^k(s) \in \{2, \dots, L+1\}$ .

*Proof.* Since  $l_h^k(s) \in [L+1]$  by definition in Algorithm 4, it suffices to show  $l_h^k(s) = 1$  cannot occur. If it does, then by Algorithm 4, Claim 1, Claim 9, Assumption 1, and definition  $\varepsilon_{\text{rnd}} = 2^{-4L}/d$ , we can find  $a \in \mathcal{A}$  to obtain the following contradiction:

$$1/2 < \|\phi(s, a)\|_{(\tilde{\Lambda}_{h,l}^k)^{-1}} \leq \|\phi(s, a)\|_{(\Lambda_{h,l}^k)^{-1}} + \sqrt{d\varepsilon_{\text{rnd}}} \leq (\|\phi(s, a)\|_2/4) + 2^{-2L} \leq 1/2. \quad \square$$

**Claim 10.** For any  $k \in [K]$ ,  $h \in [H]$ , and  $l \in [L]$ , we have  $\|w_{h,l}^k\|_2 \leq (2^l dH)^4$ .

*Proof.* For any  $x \in \mathbb{R}^d$ , Claim 9 implies

$$\|(\Lambda_{h,l}^k)^{-1}x\|_2 = \|(\Lambda_{h,l}^k)^{-1/2}x\|_{(\Lambda_{h,l}^k)^{-1}} \leq \|(\Lambda_{h,l}^k)^{-1/2}x\|_2/4 = \|x\|_{(\Lambda_{h,l}^k)^{-1}}/4.$$

Combined with the triangle inequality and Corollary 5, we obtain

$$\begin{aligned} \|w_{h,l}^k\|_2 &\leq \frac{1}{4} \sum_{\tau \in \Psi_{h,l}^{k-1}} \|\phi(s_h^\tau, a_h^\tau)\|_{(\Lambda_{h,l}^k)^{-1}} |r_h^\tau(s_h^\tau, a_h^\tau) + V_{h+1}^k(s_{h+1}^\tau)| \\ &< H \sum_{\tau \in \Psi_{h,l}^{k-1}} \|\phi(s_h^\tau, a_h^\tau)\|_{(\Lambda_{h,l}^k)^{-1}} \leq H(2^l d)^4 \leq (2^l dH)^4, \end{aligned}$$

where the second inequality holds because  $r_h^\tau(s_h^\tau, a_h^\tau) \in [0, 1]$  by assumption and  $V_{h+1}^k(s_{h+1}^\tau) \in [0, H]$  by definition in Algorithm 4 (so  $|r_h^\tau(s_h^\tau, a_h^\tau) + V_{h+1}^k(s_{h+1}^\tau)| \leq 1 + H \leq 2H < 4H$ ).  $\square$

**Claim 11.** Define  $\mathcal{V}$  as in the proof of Lemma 1. Then for any  $k \in [K]$ ,  $h \in [H]$ , and  $l \in [L]$ , we have

$$\sqrt{8H^2 \log \left( \frac{\det(\Lambda_{h,l}^k) 2HL|\mathcal{V}|}{\det(16I) \delta} \right)} \leq \beta.$$

*Proof.* We first bound  $|\mathcal{V}|$ . Clearly,  $|\mathcal{V}| \leq |\mathcal{X}|^L |\mathcal{Y}|^L$ . Next, observe

$$|\mathcal{Y}| = (1 + 2\lceil 1/(16\varepsilon_{\text{rnd}}) \rceil)^{d^2} \leq (3 + 1/(8\varepsilon_{\text{rnd}}))^{d^2} < 1/\varepsilon_{\text{rnd}}^{d^2} = 2^{4Ld^2} d^{d^2},$$

where the second inequality holds since  $\varepsilon_{\text{rnd}} < 7/24$  in Algorithm 3. For  $\mathcal{X}$ , we have

$$|\mathcal{X}| \leq (3 + 2(2^L dH)^4/\varepsilon_{\text{rnd}})^d \leq ((3 \cdot 2^{-4} + 2)(2^L dH)^4/\varepsilon_{\text{rnd}})^d < (4(2^L dH)^4/\varepsilon_{\text{rnd}})^d = 2^{(8L+2)d} d^{5d} H^{4d},$$

where the second inequality uses  $\varepsilon_{\text{rnd}} \leq 1$  in Algorithm 3. Hence, we have shown

$$|\mathcal{V}| \leq (2^{4Ld^2} d^{d^2} \cdot 2^{(8L+2)d} d^{5d} H^{4d})^L \leq (2^{14} d^6 H^4)^{L^2 d^2}.$$

Furthermore, by Corollary 4, we have  $\det(\Lambda_{h,l}^k)/\det(16I) \leq 2^{5dl} \leq 2^{5L^2 d^2}$ . Combining,

$$\log \left( \frac{2HL}{\delta} \cdot \frac{\det(\Lambda_{h,l}^k)|\mathcal{V}|}{\det(16I)} \right) \leq \log \left( \frac{2HL}{\delta} \cdot (2^{19} d^6 H^4)^{L^2 d^2} \right) \leq 20L^2 d^2 \log(2dHL/\delta).$$

Together with the fact that  $8 \cdot 20 = 160 < 169 = 13^2$ , we obtain

$$\sqrt{8H^2 \log \left( \frac{\det(\Lambda_{h,l}^k) 2HL|\mathcal{V}|}{\det(16I) \delta} \right)} < 13dHL \sqrt{\log(3dHL/\delta)} = \beta. \quad \square$$

*Proof of Lemma 3.* We fix  $k$  and use induction on  $h$ . When  $h = H$ , since  $V_{H+1}^k(\cdot) = 0$ , we have  $\gamma_H^k = 0$  and  $V_H^{\pi_k}(s) = r_H(s, \pi_H^k(s)) = \bar{Q}_H^k(s, \pi_H^k(s))$  (for any  $s$ ). Therefore, by Claim 5, for any  $s$ ,

$$V_H^k(s) - V_H^{\pi_k}(s) = V_H^k(s) - \bar{Q}_H^k(s, \pi_H^k(s)) \leq 8\alpha \cdot 2^{-l_H^k(s)} + \varepsilon_{\text{eff}}.$$

Hence, choosing  $s = s_H^k$  yields the bound. Now assume (14) holds for  $h + 1 \in \{2, \dots, H\}$ . Then for any  $s \in \mathcal{S}$ ,

$$\begin{aligned} V_h^k(s) - V_h^{\pi^k}(s) &= V_h^k(s) - \bar{Q}_h^k(s, \pi_h^k(s)) + \int_{s' \in \mathcal{S}} (V_{h+1}^k(s') - V_{h+1}^{\pi^k}(s')) P_h(s'|s, \pi_h^k(s)) \\ &\leq 8\alpha \cdot 2^{-l_h^k(s)} + \varepsilon_{\text{eff}} + \int_{s' \in \mathcal{S}} (V_{h+1}^k(s') - V_{h+1}^{\pi^k}(s')) P_h(s'|s, \pi_h^k(s)), \end{aligned}$$

where the inequality again uses Claim 5. Choosing  $s = s_h^k$  and using the fact that  $a_h^k = \pi_h^k(s_h^k)$  and the inductive hypothesis, we thus obtain

$$\begin{aligned} V_h^k(s_h^k) - V_h^{\pi^k}(s_h^k) &\leq 8\alpha \cdot 2^{-l_h^k(s_h^k)} + \varepsilon_{\text{eff}} + \gamma_h^k + V_{h+1}^k(s_{h+1}^k) - V_{h+1}^{\pi^k}(s_{h+1}^k) \\ &\leq 8\alpha \sum_{h'=h}^H 2^{-l_{h'}^k(s_{h'}^k)} + \sum_{h'=h}^H \gamma_{h'}^k + (H - h + 1)\varepsilon_{\text{eff}}. \quad \square \end{aligned}$$

*Proof of Lemma 4.* We again use induction on  $h$ . For  $h = H$ , the bound follows immediately from Claim 7 and  $V_{H+1}^*(\cdot) = V_{H+1}^k(\cdot) = 0$ . Now assume that for some  $h + 1 \in \{2, \dots, H\}$  and all  $s' \in \mathcal{S}$ ,

$$V_{h+1}^*(s') - V_{h+1}^k(s') \leq (2L\varepsilon_{\text{eff}} + 2^{-L}\alpha)(H - h).$$

Combining this bound with Claim 7, for any  $s \in \mathcal{S}$ , we obtain

$$\begin{aligned} V_h^*(s) - V_h^k(s) &\leq \max \left\{ \max_{a \in \mathcal{A}} \int_{s' \in \mathcal{S}} (V_{h+1}^*(s') - V_{h+1}^k(s')) P_h(s'|s, a) + 2L\varepsilon_{\text{eff}} + 2^{-L}\alpha, 0 \right\} \\ &\leq \max \left\{ (2L\varepsilon_{\text{eff}} + 2^{-L}\alpha)(H - h + 1), 0 \right\} = (2L\varepsilon_{\text{eff}} + 2^{-L}\alpha)(H - h + 1). \quad \square \end{aligned}$$

## C.2 Complexity analysis

The time complexity of Algorithm 4 is dominated by the computation of the set of induced norms  $\{\|\phi(s, a)\|_{(\bar{\Lambda}_{h,l}^k)^{-1}} : a \in \mathcal{A}_{h,l}^k(s), l \in [l_h^k(s) \wedge L]\}$ , which requires at most  $O(d^2 L |\mathcal{A}|)$  time. At episode  $k$  of Algorithm 3, Algorithm 4 is called to compute  $\{V_{h+1}^k(s_{h+1}^\tau) : \tau \in \Psi_{h,l}^{k-1}, l \in [L], h \in \{2, \dots, H\}\}$  in Line 7 and  $\{\pi_h^k(s_h^k) : h \in [H]\}$  in Line 13, for a total number of calls

$$\sum_{h=2}^H \sum_{l=1}^L |\Psi_{h,l}^{k-1}| + H \leq 40dHL \sum_{l=1}^L 4^l + H = O(4^L dHL),$$

where the inequality uses Claim 3. Alternatively, since  $\cup_{l=1}^L \Psi_{h,l}^{k-1} \subset [K]$  for each  $h \in [H]$ , we can simply bound the number of calls by  $O(HK)$ . Hence, during each episode, the time complexity of all Algorithm 4 calls is bounded above by

$$O(d^2 HL |\mathcal{A}| \min\{4^L dL, K\}). \quad (24)$$

Besides these calls, by a similar argument, computing the summations in Line 7 has runtime at most  $d \min\{\sum_{h=1}^H \sum_{l=1}^L |\Psi_{h,l}^{k-1}|, HK\} = O(dH \min\{4^L dL, K\})$ , which is dominated by (24). Iterative updates of  $\Lambda_{h,l}^k$  and  $(\Lambda_{h,l}^k)^{-1}$  and the rounding in Lines 8-9 both have complexity  $O(d^2 HL)$  per episode, and maintaining  $\Psi_{h,l}^k$  has complexity  $O(H)$ ; both quantities are dominated by (24). Hence, Algorithm 3's per-episode runtime is

$$O(d^2 HL |\mathcal{A}| \min\{4^L dL, K\}) = O\left(d^2 H |\mathcal{A}| \min\left\{d^2 \log\left(\frac{d}{\varepsilon_{\text{tol}}}\right) / \varepsilon_{\text{tol}}^2, K\right\} \log\left(\frac{d}{\varepsilon_{\text{tol}}}\right)\right).$$

At episode  $k$ , Algorithm 3 uses  $\{\Lambda_{h,l}^k\}_{h \in [H], l \in [L]}$  and  $\{w_{h,l}^k\}_{h \in [H], l \in [L]}$ , and their rounded versions, which requires  $O(d^2 HL)$  storage. Note these can be overwritten across episodes, so the total storage is  $O(d^2 HL)$  as well. Additionally, at episode  $k$ , it needs to store  $r_h^\tau(s_h^\tau, a_h^\tau)$  and  $\{\phi(s_h^\tau, a)\}_{a \in \mathcal{A}}$ , for each  $\tau \in \Psi_{h,l}^{k-1}$ ,  $h \in [H]$ , and  $l \in [L]$ . Similar to above, this storage can be bounded by either  $d |\mathcal{A}| HK$  or  $d |\mathcal{A}| \sum_{h=1}^H \sum_{l=1}^L |\Psi_{h,l}^{k-1}| = O(4^L d^2 HL |\mathcal{A}|)$ . Hence, the total space complexity is

$$O(d^2 HL + dH |\mathcal{A}| \min\{4^L dL, K\}) = O\left(d^2 H \log\left(\frac{d}{\varepsilon_{\text{tol}}}\right) + dH |\mathcal{A}| \min\left\{d^2 \log\left(\frac{d}{\varepsilon_{\text{tol}}}\right) / \varepsilon_{\text{tol}}^2, K\right\}\right).$$

## D PROOFS OF PROPOSITIONS

### D.1 Proposition 2 proof

We use induction on  $k$ . For  $k = 1$ , since  $\Psi^0 = \Psi_1^0 = \emptyset$ , we have  $(w^k, \Lambda^k) = (w_1^k, \Lambda_1^k) = (0, 0)$ . For the inductive hypothesis, suppose  $(w^k, \Lambda^k) = (w_l^k, \Lambda_l^k)$  for some  $k \in [K]$ . By assumption, we can find  $\tilde{s}^k \in \mathcal{S}$  such that

$$\phi(\tilde{s}^k, a) = \begin{cases} \phi(s^k, a), & k \in \Psi_l^K, a \in \mathcal{A}_l^k \\ 0, & \text{otherwise} \end{cases}, \quad r(\tilde{s}^k, a) = r(s^k, a) \quad \forall a \in \mathcal{A}. \quad (25)$$

Now consider two cases. First, if  $k \notin \Psi_l^K$ , then  $(w_l^{k+1}, \Lambda_l^{k+1}) = (w_l^k, \Lambda_l^k)$ . On the other hand, (25) implies  $\max_{a \in \mathcal{A}} \|\phi(\tilde{s}^k, a)\|_{(\Lambda^k)^{-1}} = 0$ , so  $\Psi^k = \Psi^{k-1}$  and  $(w^{k+1}, \Lambda^{k+1}) = (w^k, \Lambda^k)$ . Hence,  $(w^{k+1}, \Lambda^{k+1}) = (w_l^{k+1}, \Lambda_l^{k+1})$  follows from the inductive hypothesis. Next, assume  $k \in \Psi_l^K$ . Then  $\max_{a \in \mathcal{A}_l^k} \|\phi(s^k, a)\|_{(\Lambda_l^k)^{-1}} > 2^{-l}$  in **Sup-Lin-UCB-Var**, and it plays  $a^k = \arg \max_{a \in \mathcal{A}_l^k} \|\phi(s^k, a)\|_{(\Lambda_l^k)^{-1}}$ , observes  $r^k(s^k, a^k) = r(s^k, a^k) + \eta^k$ , and updates

$$\Lambda_l^{k+1} = I + \sum_{\tau \in \Psi_l^k} \phi(s^\tau, a^\tau) \phi(s^\tau, a^\tau)^\top = \Lambda_l^k + \phi(s^k, a^k) \phi(s^k, a^k)^\top, \quad (26)$$

$$w_l^{k+1} = (\Lambda_l^{k+1})^{-1} \sum_{\tau \in \Psi_l^k} \phi(s^\tau, a^\tau) r^\tau(s^\tau, a^\tau) = (\Lambda_l^{k+1})^{-1} (\Lambda_l^k w_l^k + \phi(s^k, a^k) r^k(s^k, a^k)). \quad (27)$$

On the other hand,  $\max_{a \in \mathcal{A}} \|\phi(\tilde{s}^k, a)\|_{(\Lambda^k)^{-1}} = \max_{a \in \mathcal{A}_l^k} \|\phi(s^k, a)\|_{(\Lambda_l^k)^{-1}} > 2^{-l}$  by (25) and the inductive hypothesis. Hence, if we write  $\tilde{a}^k$  (instead of  $a^k$ ) for the action chosen by **EXPL**( $2^{-l}$ ), we have

$$\tilde{a}^k = \arg \max_{a \in \mathcal{A}} \|\phi(\tilde{s}^k, a)\|_{(\Lambda^k)^{-1}} = \arg \max_{a \in \mathcal{A}_l^k} \|\phi(s^k, a)\|_{(\Lambda_l^k)^{-1}} = a^k.$$

By (25), this implies  $\phi(\tilde{s}^k, \tilde{a}^k) = \phi(s^k, a^k)$ , and by (25) and the noise assumption, we also have  $r^k(\tilde{s}^k, \tilde{a}^k) = r^k(s^k, a^k)$ . Hence, similar to (26) and (27), we see that **EXPL**( $2^{-l}$ ) updates

$$\Lambda^{k+1} = \Lambda^k + \phi(s^k, a^k) \phi(s^k, a^k)^\top, \quad w^{k+1} = (\Lambda^{k+1})^{-1} (\Lambda^k w^k + \phi(s^k, a^k) r^k(s^k, a^k)). \quad (28)$$

Combining (26), (27), (28), and the inductive hypothesis completes the proof.

### D.2 Proposition 3 proof

We essentially follow the existing proof. First, define the good event

$$\mathcal{G}' = \bigcap_{k=1}^K \left\{ \left\| \sum_{\tau=1}^{k-1} \phi(s^\tau, a^\tau) \eta^\tau \right\|_{(\Lambda^k)^{-1}} \leq \sqrt{2d \log((\lambda + K)/(\lambda \delta))} \right\}.$$

By Theorem 1 of [Abbasi-Yadkori et al. \(2011\)](#), the assumption on  $\eta^\tau$ , Lemma 10 of [Abbasi-Yadkori et al. \(2011\)](#), and the  $\ell_2$  norm assumption on  $\phi(s^\tau, a^\tau)$ ,  $\mathbb{P}(\mathcal{G}') \geq 1 - \delta$ . We bound regret on  $\mathcal{G}'$ . Let  $\Delta^\tau = r(s^\tau, a^\tau) - \phi(s^\tau, a^\tau)^\top \theta$ . Then if we define

$$w_{k,1} = (\Lambda^k)^{-1} \sum_{\tau=1}^{k-1} \phi(s^\tau, a^\tau) \phi(s^\tau, a^\tau)^\top \theta, \quad w_{k,2} = (\Lambda^k)^{-1} \sum_{\tau=1}^{k-1} \phi(s^\tau, a^\tau) \eta^\tau, \quad w_{k,3} = (\Lambda^k)^{-1} \sum_{\tau=1}^{k-1} \phi(s^\tau, a^\tau) \Delta^\tau,$$

we have  $w_k = \sum_{i=1}^3 w_{k,i}$ , which implies

$$|\phi(s, a)^\top (w_k - \theta)| \leq |\phi(s, a)^\top (w_{k,1} - \theta)| + |\phi(s, a)^\top w_{k,2}| + |\phi(s, a)^\top w_{k,3}|.$$

Similar to the analysis of  $z_2$  in the proof of Claim 4, we have

$$|\phi(s, a)^\top (w_{k,1} - \theta)| = \lambda |\phi(s, a)^\top (\Lambda^k)^{-1} \theta| \leq \sqrt{\lambda} \|\phi(s, a)\|_{(\Lambda^k)^{-1}}.$$

On the event  $\mathcal{G}'$ , we obtain

$$|\phi(s, a)^\top w_{k,2}| \leq \sqrt{2d \log((\lambda + K)/(\lambda\delta))} \|\phi(s, a)\|_{(\Lambda^k)^{-1}}.$$

Similar to analysis of  $z_3$  in the proof of Claim 4, we know

$$|\phi(s, a)^\top w_{k,3}| \leq \varepsilon_{\text{mis}} \sqrt{K} \|\phi(s, a)\|_{(\Lambda^k)^{-1}}.$$

Hence, combining the previous four inequalities, we have shown that on  $\mathcal{G}'$ ,

$$|\phi(s, a)^\top (w_k - \theta)| \leq \left( \sqrt{\lambda} + \sqrt{2d \log((\lambda + K)/(\lambda\delta))} + \varepsilon_{\text{mis}} \sqrt{K} \right) \|\phi(s, a)\|_{(\Lambda^k)^{-1}} \leq \alpha \|\phi(s, a)\|_{(\Lambda^k)^{-1}},$$

where the second inequality is by choice of  $\lambda$  and  $\alpha$ . Thus, by the misspecification assumption,

$$|r(s, a) - \phi(s, a)^\top w_k| \leq |r(s, a) - \phi(s, a)^\top \theta| + |\phi(s, a)^\top (\theta - w_k)| \leq \varepsilon_{\text{mis}} + \alpha \|\phi(s, a)\|_{(\Lambda^k)^{-1}}.$$

Hence, by definition of the optimal policy and Lin-UCB,

$$r(s^k, a_\star^k) \leq \max_{a \in \mathcal{A}} (\phi(s^k, a)^\top w_k + \alpha \|\phi(s^k, a)\|_{(\Lambda^k)^{-1}} + \varepsilon_{\text{mis}}) \leq r(s^k, a^k) + 2\alpha \|\phi(s^k, a^k)\|_{(\Lambda^k)^{-1}} + 2\varepsilon_{\text{mis}}.$$

Substituting into the regret definition, then using Cauchy-Schwarz, Lemma 11 from [Abbasi-Yadkori et al. \(2011\)](#), and some simple eigenvalue bounds, we obtain

$$R(K) \leq 2\alpha \sum_{k=1}^K \|\phi(s^k, a^k)\|_{(\Lambda^k)^{-1}} + 2\varepsilon_{\text{mis}} K \leq 2\alpha \sqrt{Kd \log((\lambda + K)/K)} + 2\varepsilon_{\text{mis}} K.$$

Since  $(\lambda + K)/\lambda = 1 + K/(1 + \varepsilon_{\text{mis}}^2 K) \leq 1 + \min\{K, \varepsilon_{\text{mis}}^{-2}\}$ , by definition of  $\alpha$ , we have

$$\alpha = O\left(\sqrt{2d \log(\min\{K, \varepsilon_{\text{mis}}^{-2}\}/\delta)} + \varepsilon_{\text{mis}} \sqrt{K}\right).$$

Combining the previous two bounds and again using  $(\lambda + K)/\lambda \leq 1 + \min\{K, \varepsilon_{\text{mis}}^{-2}\}$  yields the desired bound.