# Uniform-PAC Bounds for Reinforcement Learning with Linear Function Approximation 

Jiafan He<br>Department of Computer Science<br>University of California, Los Angeles<br>CA 90095, USA<br>jiafanhe19@ucla.edu

Dongruo Zhou<br>Department of Computer Science<br>University of California, Los Angeles<br>CA 90095, USA<br>drzhou@cs.ucla.edu

Quanquan Gu<br>Department of Computer Science<br>University of California, Los Angeles<br>CA 90095, USA<br>qgu@cs.ucla.edu


#### Abstract

We study reinforcement learning (RL) with linear function approximation. Existing algorithms for this problem only have high-probability regret and/or Probably Approximately Correct (PAC) sample complexity guarantees, which cannot guarantee the convergence to the optimal policy. In this paper, in order to overcome the limitation of existing algorithms, we propose a new algorithm called FLUTE, which enjoys uniform-PAC convergence to the optimal policy with high probability. The uniform-PAC guarantee is the strongest possible guarantee for reinforcement learning in the literature, which can directly imply both PAC and high probability regret bounds, making our algorithm superior to all existing algorithms with linear function approximation. At the core of our algorithm is a novel minimax value function estimator and a multi-level partition scheme to select the training samples from historical observations. Both of these techniques are new and of independent interest.


## 1 Introduction

Designing efficient reinforcement learning (RL) algorithms for environments with large state and action spaces is one of the main tasks in the RL community. To achieve this goal, function approximation, which uses a class of predefined functions to approximate either the value function or transition dynamic, has been widely studied in recent years. Specifically, a series of recent works [11, 13, 18, 24, 3, 27] have studied RL with linear function approximation with provable guarantees. They show that with linear function approximation, one can either obtain a sublinear regret bound against the optimal value function [13, 24, 3, 27] or a polynomial sample complexity bound [14] (Probably Approximately Correct (PAC) bound for short) in finding a near-optimal policy [11, 18].
However, neither the regret bound or PAC bound is a perfect performance measure. As discussed in detail by [7], these two measures fail to guarantee the convergence to the optimal policy. Therefore, an algorithm with high probability regret and/or PAC bound guarantees do not necessarily learn the optimal policy, and can perform badly in practice. In detail, one can face the following two situations:

- An algorithm with a sublinear regret suggests that the summation of the suboptimality gaps $\Delta_{t}$ (the gap between the values of the current adapted policy and optimal policy, see Definition 3.3
for details.) in the first $T$ rounds is bounded by $o(T)$. However, this algorithm may be arbitrarily suboptimal infinitely times ${ }^{1}$, thus it fails to converge to the optimal policy.
- An algorithm is $(\epsilon, \delta)$-PAC suggests that with probability at least $1-\delta$, the number of suboptimality gaps $\Delta_{t}$ that are greater than $\epsilon$ will be at most polynomial in $\epsilon$ and $\log (1 / \delta)$. The formal definition of $(\epsilon, \delta)$-PAC can be found in Definition 3.4. However, this algorithm may still have gaps satisfying $\epsilon / 2<\Delta_{t}<\epsilon$ infinitely often, thus fails to converge to the optimal policy.

To overcome the limitations of regret and PAC guarantees, Dann et al. [7] proposed a new performance measure called uniform-PAC, which is a strengthened notion of the PAC framework. Specifically, an algorithm is uniform-PAC if there exists a function of the target accuracy $\epsilon$ and the confidence parameter $\delta$ that upper bounds the number of suboptimality gaps satisfying $\Delta_{t}>\epsilon$ simultaneously for all $\epsilon>0$ with probability at least $1-\delta$. The formal definition of uniform-PAC can be found in Definition 3.6. Algorithms that are uniform-PAC converge to an optimal policy with high probability, and yield both PAC and high probability regret bounds. In addition, they proposed a UBEV algorithm for learning tabular MDPs, which is uniform-PAC. Nevertheless, UBEV is designed for tabular MDPs, and it is not clear how to incorporate function approximation into UBEV to scale it up for large (or even infinite) state and action space. Therefore, a natural question arises:

## Can we design a provable efficient uniform-PAC RL algorithm with linear function approximation?

In this work, we answer this question affirmatively. In detail, we propose new algorithms for both contextual linear bandits and linear Markov decision processes (MDPs) [22, 13]. Both of them are uniform-PAC, and their sample complexity is comparable to that of the state-of-the-art algorithms which are not uniform-PAC. Our key contributions are highlighted as follows.

- We begin with contextual linear bandits problem as a "warm-up" example of the RL with linear function approximation (with horizon length equals 1). We propose a new algorithm called uniformPAC OFUL (UPAC-OFUL), and show that our algorithm is uniform-PAC with $\widetilde{O}\left(d^{2} / \epsilon^{2}\right)$ sample complexity, where $d$ is the dimension of contexts and $\epsilon$ is the accuracy parameter. In addition, this result also implies an $\widetilde{O}(d \sqrt{T})$ regret in the first $T$ round and matches the result of OFUL algorithm [1] up to a logarithmic factor. The key idea of our algorithm is a novel minimax linear predictor and a multi-level partition scheme to select the training samples from past observations. To the best of our knowledge, this is the first algorithm with a uniform PAC-bound for contextual bandits problems.
- We also consider RL with linear function approximation in episodic linear MDPs, where the transition kernel admits a low-rank factorization. We propose an algorithm dubbed uniForm-PAC Least-sqUare value iTEration (FLUTE), which adapts the novel techniques we developed in the contextual linear bandits setting, and show that our algorithm is uniform-PAC with $\widetilde{O}\left(d^{3} H^{5} / \epsilon^{2}\right)$ sample complexity, where $d$ is the dimension of the feature mapping, $H$ is the length of episodes and $\epsilon$ is the accuracy parameter. This result further implies an $\widetilde{O}\left(\sqrt{d^{3} H^{4} T}\right)$ regret in the first $T$ steps and matches the result of LSVI-UCB algorithm [13] up to a $\sqrt{H}$-factor, while LSVI-UCB is not uniform-PAC. Again, FLUTE is the first uniform-PAC RL algorithm with linear function approximation.

Notation We use lower case letters to denote scalars, and use lower and upper case bold face letters to denote vectors and matrices respectively. For any positive integer $n$, we denote by $[n]$ the set $\{1, \ldots, n\}$. For a vector $\mathbf{x} \in \mathbb{R}^{d}$, we denote by $\|\mathbf{x}\|_{1}$ the Manhattan norm and denote by $\|\mathbf{x}\|_{2}$ the Euclidean norm. For a vector $\mathbf{x} \in \mathbb{R}^{d}$ and matrix $\boldsymbol{\Sigma} \in \mathbb{R}^{d \times d}$, we define $\|\mathbf{x}\|_{\boldsymbol{\Sigma}}=\sqrt{\mathbf{x}^{\top} \boldsymbol{\Sigma} \mathbf{x}}$. For two sequences $\left\{a_{n}\right\}$ and $\left\{b_{n}\right\}$, we write $a_{n}=O\left(b_{n}\right)$ if there exists an absolute constant $C$ such that $a_{n} \leq C b_{n}$. We use $\widetilde{O}(\cdot)$ to further hide the logarithmic factors. For logarithmic regret, we use $\widetilde{O}(\cdot)$ to hide all logarithmic terms except $\log T$.

[^0]
## 2 Related Work

### 2.1 Linear bandits

There is a series of works focusing on the stochastic linear bandits problem. These works can be categorized into two groups: the works aim at providing sublinear regret guarantee and the works providing PAC bound for linear best-arm identification problem. In detail, for finite action set with $K$ arms, Auer [2] proposed a SupLinRel algorithm which achieves an $O\left(\sqrt{d T \log ^{3}(T K)}\right)$ regret, where $K$ is the number of arms. Chu et al. [5] proposed a SupLinUCB algorithm which has the same regret bound as SupLinRel, but is easier to implement. Li et al. [17] proposed a Variable-Confidence-Level (VCL) SupLinUCB algorithm and improved the regret bound to $O(\sqrt{d T \log T \log K})$. For infinite action set, Dani et al. [6] proposed a Confidence Ball algorithm with an $O\left(d \sqrt{T \log ^{3} T}\right)$ regret and proved an $\Omega(d \sqrt{T})$ lower bound. Abbasi-Yadkori et al. [1] proposed OFUL algorithm and improved the regret bound to $O\left(d \sqrt{T \log ^{2} T}\right)$. When the reward has a bounded variance, Zhou et al. [27], Zhang et al. [25] proposed algorithms with variance-aware confident sets, and obtained tight variance-dependent regret bounds. For the best-arm identification problem, to find an arm which is $\epsilon$-suboptimal, Soare et al. [19] proposed a G-allocation strategy with an $\widetilde{O}\left(d / \epsilon^{2}\right)$ sample complexity. Karnin [15] proposed an Explore-Verify framework which improves the sample complexity by some logarithmic factors. Xu et al. [21] proposed a LinGapE algorithm whose sample complexity matches the lower bound up to some $K$ factors. Tao et al. [20] proposed an ALBA algorithm which improves the sample complexity to have a linear dimension dependence. Fiez et al. [8] studied transductive linear bandits and proposed an algorithm with sample complexity similar to linear bandits. Compared with best-arm identification, the contextual bandits setting we focus on is more challenging since the action set will change at each round.

### 2.2 RL with linear function approximation

Recently, a line of work focuses on analyzing RL with linear function approximation. To mention a few, Jiang et al. [11] studied MDPs with low Bellman rank and proposed an OLIVE algorithm, which has the PAC guarantee. Yang and Wang [22] studied the linear transition model and proposed a sample-optimal Q-learning method with a generative model. Jin et al. [13] studied the linear MDP model and proposed an LSVI-UCB algorithm under the online RL setting (without a generative model) with $\widetilde{O}\left(\sqrt{d^{3} H^{3} T}\right)$ regret. Later, Zanette and Brunskill [23] studied the low inherent Bellman error model and proposed an ELEANOR algorithm with a better regret $\widetilde{O}(d H \sqrt{T})$ using a global planning oracle. Modi et al. [18] studied the linearly combined model ensemble and proposed a provable sample-efficient algorithm. Jia et al. [10], Ayoub et al. [3] studied the linear mixture MDPs and proposed a UCRL-VTR algorithm with an $\widetilde{O}\left(d \sqrt{H^{3} T}\right)$ regret. Recently Zhou et al. [27] improved the regret bound to $\widetilde{O}(d H \sqrt{T})$ with a new algorithm design and a new Bernstein inequality. However, all of these works aim at deriving PAC sample complexity guarantee or regret bound, and none of them has the uniform PAC guarantee for learning MDPs with linear function approximation. Our work will fill this gap in the linear MDP setting [22, 13].

## 3 Preliminaries

We consider episodic Markov Decision Processes (MDPs) in this work. Each episodic MDP is denoted by a tuple $M\left(\mathcal{S}, \mathcal{A}, H,\left\{r_{h}\right\}_{h=1}^{H},\left\{\mathbb{P}_{h}\right\}_{h=1}^{H}\right)$. Here, $\mathcal{S}$ is the state space, $\mathcal{A}$ is the finite action space, $H$ is the length of each episode, $r_{h}: \mathcal{S} \times \mathcal{A} \rightarrow[0,1]$ is the reward function at stage $h$ and $\mathbb{P}_{h}\left(s^{\prime} \mid s, a\right)$ is the transition probability function at stage $h$ which denotes the probability for state $s$ to transfer to state $s^{\prime}$ with action $a$ at stage $h$. A policy $\pi: \mathcal{S} \times[H] \rightarrow \mathcal{A}$ is a function which maps a state $s$ and the stage number $h$ to an action $a$. For any policy $\pi$ and stage $h \in[H]$, we define the action-value function $Q_{h}^{\pi}(s, a)$ and value function $V_{h}^{\pi}(s)$ as follows

$$
Q_{h}^{\pi}(s, a)=r_{h}(s, a)+\mathbb{E}\left[\sum_{h^{\prime}=h+1}^{H} r_{h^{\prime}}\left(s_{h^{\prime}}, \pi\left(s_{h^{\prime}}, h^{\prime}\right)\right) \mid s_{h}=s, a_{h}=a\right], V_{h}^{\pi}(s)=Q_{h}^{\pi}(s, \pi(s, h))
$$

where $s_{h^{\prime}+1} \sim \mathbb{P}_{h}\left(\cdot \mid s_{h^{\prime}}, a_{h^{\prime}}\right)$. We define the optimal value function $V_{h}^{*}$ and the optimal actionvalue function $Q_{h}^{*}$ as $V_{h}^{*}(s)=\max _{\pi} V_{h}^{\pi}(s)$ and $Q_{h}^{*}(s, a)=\max _{\pi} Q_{h}^{\pi}(s, a)$. By definition, the
value function $V_{h}^{\pi}(s)$ and action-value function $Q_{h}^{\pi}(s, a)$ are bounded in $[0, H]$. For any function $V: \mathcal{S} \rightarrow \mathbb{R}$, we denote $\left[\mathbb{P}_{h} V\right](s, a)=\mathbb{E}_{s^{\prime} \sim \mathbb{P}_{h}(\cdot \mid s, a)} V\left(s^{\prime}\right)$. Therefore, for each stage $h \in[H]$ and policy $\pi$, we have the following Bellman equation, as well as the Bellman optimality equation:

$$
\begin{equation*}
Q_{h}^{\pi}(s, a)=r_{h}(s, a)+\left[\mathbb{P}_{h} V_{h+1}^{\pi}\right](s, a), Q_{h}^{*}(s, a)=r_{h}(s, a)+\left[\mathbb{P}_{h} V_{h+1}^{*}\right](s, a) \tag{3.1}
\end{equation*}
$$

where $V_{H+1}^{\pi}=V_{H+1}^{*}=0$. At the beginning of the episode $k$, the agent determines a policy $\pi_{k}$ to be followed in this episode. At each stage $h \in[H]$, the agent observes the state $s_{h}^{k}$, chooses an action following the policy $\pi_{k}$ and observes the next state with $s_{h+1}^{k} \sim \mathbb{P}_{h}\left(\cdot \mid s_{h}^{k}, a_{h}^{k}\right)$.
We consider linear function approximation in this work. Therefore, we make the following linear MDP assumption, which is firstly proposed in [22, 13].
Assumption 3.1. $\operatorname{MDP} \mathcal{M}\left(\mathcal{S}, \mathcal{A}, H,\left\{r_{h}\right\}_{h=1}^{H},\left\{\mathbb{P}_{h}\right\}_{h=1}^{H}\right)$ is a linear MDP such that for any stage $h \in[H]$, there exists an unknown vector $\boldsymbol{\mu}_{h}$, an unknown measure $\boldsymbol{\theta}_{h}(\cdot): \mathcal{S} \rightarrow \mathbb{R}^{d}$ and a known feature mapping $\phi: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^{d}$, such that for each $(s, a) \in \mathcal{S} \times \mathcal{A}$ and $s^{\prime} \in \mathcal{S}$,

$$
\mathbb{P}_{h}\left(s^{\prime} \mid s, a\right)=\left\langle\boldsymbol{\phi}(s, a), \boldsymbol{\theta}_{h}\left(s^{\prime}\right)\right\rangle, r_{h}(s, a)=\left\langle\boldsymbol{\phi}(s, a), \boldsymbol{\mu}_{h}\right\rangle .
$$

For simplicity, we assume that $\boldsymbol{\mu}_{h}, \boldsymbol{\theta}_{n}(\cdot)$ and $\boldsymbol{\phi}(\cdot, \cdot)$ satisfy $\|\boldsymbol{\phi}(s, a)\|_{2} \leq 1$ for all $s, a,\left\|\boldsymbol{\mu}_{h}\right\|_{2} \leq \sqrt{d}$ and $\left\|\boldsymbol{\theta}_{h}(\mathcal{S})\right\|_{2} \leq \sqrt{d}$. The linear MDP assumption automatically suggests that for any policy $\pi$, the action-value function $Q_{h}^{\pi}$ is always a linear function of the given feature mapping $\phi$, which is summarized in the following proposition.
Proposition 3.2 (Proposition 2.3, [13]). For any policy $\pi$, there exist weights $\left\{\mathbf{w}_{h}^{\pi}\right\}_{h=1}^{H}$ such that for any $s, a, h \in \mathcal{S} \times \mathcal{A} \times[H], Q_{h}^{\pi}(s, a)=\left\langle\phi(s, a), \mathbf{w}_{h}^{\pi}\right\rangle$.

Next we define the regret and $(\epsilon, \delta)$-PAC formally.
Definition 3.3. For an RL algorithm Alg , we define its regret on learning an $\operatorname{MDP} M(\mathcal{S}, \mathcal{A}, H, r, \mathbb{P})$ in the first $K$ episodes as the sum of the suboptimality for episode $k=1, \ldots, K$,

$$
\operatorname{Regret}(K)=\sum_{k=1}^{K} V_{1}^{*}\left(s_{1}^{k}\right)-V_{1}^{\pi_{k}}\left(s_{1}^{k}\right)
$$

where $\pi_{k}$ is the policy in the $k$-th episode.
Definition 3.4. For an RL algorithm Alg and a fixed $\epsilon$, let $\pi_{1}, \pi_{2}, \ldots$ be the policies generated by Alg. Let $N_{\epsilon}=\sum_{k=1}^{\infty} \mathbb{1}\left\{V_{1}^{*}\left(s_{1}^{k}\right)-V_{1}^{\pi_{k}}\left(s_{1}^{k}\right)>\epsilon\right\}$ be the number of episodes whose suboptimality gap is greater than $\epsilon$. Then we say Alg is $(\epsilon, \delta)$-PAC with sample complexity $f(\epsilon, \delta)$ if

$$
\mathbb{P}\left(N_{\epsilon}>f(\epsilon, \delta)\right) \leq \delta
$$

Remark 3.5. Dann et al. [7] suggested that an algorithm with a sublinear regret is not necessarily to be an $(\epsilon, \delta)$-PAC algorithm. However, with some modification, Jin et al. [12] showed that any algorithm with a sublinear regret can be converted to a new algorithm which is $(\epsilon, \delta)$-PAC, which does not contradict with the claim by [7]. For example, Ghavamzadeh et al. [9] and Zhang et al. [26] proposed algorithms with $\widetilde{O}(\sqrt{S A H T})$ regret, and both algorithms can be converted into new algorithms which are $(\epsilon, \delta)$-PAC with sample complexity $\widetilde{O}\left(S A H^{2} / \epsilon^{2}\right)$.

Both regret and PAC guarantees are not perfect. As Dann et al. [7] showed, an algorithm with sub-linear regret or $(\delta, \epsilon)$-PAC bound may fail to converge to the optimal policy. For an $(\delta, \epsilon)$-PAC algorithm with $\Delta_{t}=\epsilon / 2(t \in \mathbb{N})$, it still has linear regret $O(\epsilon T)$ and will never converge to the optimal policy. For an algorithm with a sub-linear regret bound, a constant sub-optimality gap may still occur infinite times. Therefore, Dann et al. [7] proposed uniform-PAC algorithms, which are defined formally as follows.
Definition 3.6. For an RL algorithm Alg, let $\pi_{1}, \pi_{2}, \ldots$ be the policies generated by Alg. Let $N_{\epsilon}=\sum_{k=1}^{\infty} \mathbb{1}\left\{V_{1}^{*}\left(s_{1}^{k}\right)-V_{1}^{\pi_{k}}\left(s_{1}^{k}\right)>\epsilon\right\}$ be the number of episodes whose suboptimality gap is greater than $\epsilon$. We say Alg is uniform-PAC for some $\delta \in(0,1)$ with sample complexity $f(\epsilon, \delta)$ if

$$
\mathbb{P}\left(\exists \epsilon>0, N_{\epsilon}>f(\epsilon, \delta)\right) \leq \delta
$$

The following theorem suggests that a uniform-PAC algorithm is automatically a PAC algorithm and an algorithm with sublinear regret.

```
Algorithm 1 Uniform-PAC OFUL (UPAC-OFUL)
Require: Regularization parameter \(\lambda\), confidence radius \(\beta_{l}(l \in \mathbb{N})\)
    Set \(\mathcal{C}^{l} \leftarrow \emptyset, l \in \mathbb{N}\) and the total level \(S_{1}=1\)
    for round \(k=1,2, .\). do
        for all level \(l \in\left[S_{k}\right]\) do
            Set \(\boldsymbol{\Sigma}_{k}^{l}=\lambda \mathbf{I}+\sum_{i \in \mathcal{C}^{l}} \mathbf{x}_{i} \mathbf{x}_{i}^{\top}\)
            Set \(\mathbf{b}_{k}^{l}=\sum_{i \in \mathcal{C}^{l}} \mathbf{x}_{i} r_{i}\) and \(\mathbf{w}_{k}^{l}=\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1} \mathbf{b}_{k}^{l}\)
        end for
        Receive the action set \(\mathcal{D}_{k}\)
        Choose action \(\mathbf{x}_{k} \leftarrow \operatorname{argmax}_{\mathbf{x} \in \mathcal{D}_{k}} \min _{1 \leq l \leq S_{k}}\left(\mathbf{w}_{k}^{l}\right)^{\top} \mathbf{x}+\beta_{l} \sqrt{\mathbf{x}^{\top}\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1} \mathbf{x}}\)
        Set level \(l_{k}=1\)
        while \(\sqrt{\mathbf{x}_{k}^{\top}\left(\boldsymbol{\Sigma}_{k}^{l_{k}}\right)^{-1} \mathbf{x}_{k}} \leq 2^{-l_{k}}\) and \(l_{k} \leq S_{k}\) do
            \(l_{k} \leftarrow l_{k}+1\)
        end while
        Add the new element \(k\) to the set \(\mathcal{C}^{l_{k}}\) and receive the reward \(r_{k}\)
        Set the total level \(S_{k+1}\) as \(S_{k+1}=\max _{l:\left|\mathcal{C}^{l}\right|>0} l\)
    end for
```

Theorem 3.7 (Theorem 3, [7]). If an algorithm Alg is uniform-PAC for some $\delta \geq 0$, with sample complexity $\widetilde{O}\left(C_{1} / \epsilon+C_{2} / \epsilon^{2}\right)$, where $C_{1}, C_{2}$ are constant and depend only on $S, A, H, \log (1 / \delta)$. Then, we have the following results:

- 1: Alg will converge to optimal policies with high probability at least $1-\delta: \mathbb{P}\left(\lim _{k \rightarrow+\infty} V_{1}^{*}\left(s_{1}^{k}\right)-\right.$ $\left.V_{1}^{\pi_{k}}\left(s_{1}^{k}\right)=0\right) \geq 1-\delta$
- 2: With probability at least $1-\delta$, for each $K \in \mathbb{N}$, the regret for Alg in the first $K$ episodes is upper bounded by $\widetilde{O}\left(\sqrt{C_{2} K}+C_{1}+C_{2}\right)$.
- 3: For each $\epsilon \geq 0$, Alg is also $(\epsilon, \delta)$-PAC with the same sample complexity $\widetilde{O}\left(C_{1} / \epsilon+C_{2} / \epsilon^{2}\right)$.

Theorem 3.7 suggests that uniform-PAC is stronger than both the PAC and regret guarantees. In the remainder of this paper, we aim at developing uniform-PAC RL algorithms with linear function approximation.

## 4 Warm up: Uniform-PAC Bounds for Linear Bandits

To better illustrate the idea of our algorithm, in this section, we consider a contextual linear bandits problem, which can be regarded as a special linear MDP with $H=1$. Let $\left\{\mathcal{D}_{k}\right\}_{k=1}^{\infty}$ be a fixed sequence of decision/action sets. At round $k$, the agent selects an action $\mathbf{x}_{k} \in \mathcal{D}_{k}$ by the algorithm $\mathcal{H}$ and then observes the reward $r_{k}=\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}_{k}\right\rangle+\epsilon_{k}$, where $\boldsymbol{\mu}^{*} \in \mathbb{R}^{d}$ is a vector unknown to the agent and $\epsilon_{k}$ is a sub-Gaussian random noise. $\mathbf{x}_{k}, \epsilon_{k}, \boldsymbol{\mu}^{*}$ satisfy the following properties:

$$
\begin{equation*}
\forall k \in \mathbb{N}, \lambda \in \mathbb{R}, \mathbb{E}\left[e^{\lambda \epsilon_{k}} \mid \mathbf{x}_{1: k}, \epsilon_{1: k-1}\right] \leq \exp \left(\lambda^{2} / 2\right),\left\|\mathbf{x}_{k}\right\|_{2} \leq 1,\left\|\boldsymbol{\mu}^{*}\right\|_{2} \leq 1 \tag{4.1}
\end{equation*}
$$

Our goal is to design an $(\epsilon, \delta)$-uniform-PAC algorithm with sample complexity $f(\epsilon, \delta)$ such that

$$
\mathbb{P}\left(\exists \epsilon>0, \sum_{k=1}^{\infty} \mathbb{1}\left\{\Delta_{k}:=\max _{\mathbf{x} \in \mathcal{D}_{k}}\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}\right\rangle-\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}_{k}\right\rangle>\epsilon\right\}>f(\epsilon, \delta)\right)<\delta,
$$

where $\Delta_{k}:=\max _{\mathbf{x} \in \mathcal{D}_{k}}\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}\right\rangle-\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}_{k}\right\rangle$ denotes the suboptimality at round $k$.
Here we assume the weight vector $\boldsymbol{\mu}^{*}$ satisfies $\left\|\boldsymbol{\mu}^{*}\right\|_{2} \leq 1$, to be consistent with the assumption made in Abbasi-Yadkori et al. [1]. Our assumption can be easily relaxed to the general $\left\|\boldsymbol{\mu}^{*}\right\|_{2} \leq B$ case with an additional $\log B$ factor in the sample complexity, as can be seen in the following analysis.
Why existing algorithms fail to be uniform-PAC? Before proposing our algorithm, it is natural to ask whether existing methods have already been uniform-PAC. We take OFUL [1] for example,
which is the state-of-the-art linear bandit algorithm in our setting. At round $k$, OFUL constructs an optimistic estimation of the true linear function $\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}\right\rangle$, by doing linear regression over all past $k$ selected actions $\mathbf{x}_{i}, 1 \leq i \leq k$ and their corresponding rewards. The optimistic estimation has a closed-form as the summation of the linear regression predictor and a quadratic confidence bound $\mathbf{w}_{k}^{\top} \mathbf{x}+\alpha \sqrt{\mathbf{x}^{\top} \boldsymbol{\Sigma}_{k}^{-1} \mathbf{x}}$, where $\boldsymbol{\Sigma}_{k}=\lambda \mathbf{I}+\sum_{i=1}^{k-1} \mathbf{x}_{i} \mathbf{x}_{i}^{\top}$ [16]. Following the standard analysis of OFUL in [1], we obtain the following upper confidence bound of the suboptimality gap $\Delta_{k}$ :

$$
\begin{equation*}
\text { With probability at least } 1-\delta, \forall k>0, \Delta_{k}=O\left(\sqrt{d \log (k / \delta)}\left\|\mathbf{x}_{k}\right\|_{\boldsymbol{\Sigma}_{k}^{-1}}\right) \text {, } \tag{4.2}
\end{equation*}
$$

where the $\log k$ is due to the fact that OFUL makes use of all past $k$ observed actions. Since the agent can only say whether an arm is good or not based on the confidence bound of $\Delta_{k}$, due to the existence of the $\log k$ term in (4.2), the bounds on the suboptimality gap for the "good" arms may be large (since $\log k$ grows as $k$ increases). That makes the agent fail to recognize those "good" arms and instead pull the "bad" arms infinite times, which suggests that OFUL is not a uniform-PAC algorithm. For other algorithms, they either need to know the total round $T$ before running the algorithm [5], or need to assume that the decision sets $\mathcal{D}_{k}$ are identical (e.g., algorithms for best-arm identification [19]), thus none of them fits into our setting.

Key techniques of our algorithm. In order to address the aforementioned issue, we proposed UPAC-OFUL in Algorithm 1. The key idea of Algorithm 1 is to divide all the historical observed data into non-overlapping sets $\mathcal{C}^{l}$, while each $\mathcal{C}^{l}$ only includes finite past historical observed actions. This helps successfully avoid the $\log k$ term appearing in the confidence bound in [1]. Then at round $k$, Algorithm 1 only constructs optimistic estimation of $\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}_{k}\right\rangle$ over the first $S_{k}$ sets $\mathcal{C}^{l}$ individually, where $S_{k}$ is the number of non-empty sets $\mathcal{C}^{l}$. In detail, the optimistic estimation over $\mathcal{C}^{l}$ has the form

$$
\begin{equation*}
\left(\mathbf{w}_{k}^{l}\right)^{\top} \mathbf{x}+\beta_{l} \sqrt{\mathbf{x}^{\top}\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1} \mathbf{x}} \tag{4.3}
\end{equation*}
$$

where $\Sigma_{k}^{l}$ is the covariance matrix for actions in set $\mathcal{C}^{l}$ (Line 4), and $\mathbf{w}_{k}^{l}$ is the estimation of $\boldsymbol{\mu}^{*}$ obtained by ridge regression defined in Line 5. Meanwhile, for a newly selected action $\mathbf{x}_{k}$, Algorithm 1 needs to decide which $\mathcal{C}^{l}$ it should be added to. Inspired by [5], Algorithm 1 tests the "uncertainty" of $\mathbf{x}_{k}$ against $\mathcal{C}^{l}$, by calculating its confidence bound $\left\|\mathbf{x}_{k}\right\|_{\left(\Sigma_{k}^{l}\right)^{-1}}$. Then Algorithm 1 adds $\mathbf{x}_{k}$ to the lowest possible level where the "uncertainty" is larger than a certain threshold (i.e., $\left.\left\|\mathbf{x}_{k}\right\|_{\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1}}>2^{-l}\right)$. Such a selection rule guarantees two things simultaneously. First, it ensures that the cardinality of each $\mathcal{C}^{l}$ is finite, due to the fact that the summation of $\left\|\mathbf{x}_{k}\right\|_{\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1}}$ can be properly bounded. Second, it also guarantees that the "uncertainty" of the reward corresponding to $\mathbf{x}_{k}$ is still small, since by the level selection rule we have $\left\|\mathbf{x}_{k}\right\|_{\left(\Sigma_{k}^{l-1}\right)^{-1}} \leq 2^{-(l-1)}$. Lastly, to make use of all $S_{k}$ optimistic estimations, Algorithm 1 constructs the final predictor as the minimal value of $S_{k}$ individual predictor (4.3) over $\mathcal{C}^{l}$ (Line 8). Since each individual predictor is a valid upper bound of the true function, the minimum of them is still valid and tighter than each of them (except the smallest one), which makes it possible to provide a stronger uniform-PAC guarantee.

The following theorem shows that Algorithm 1 is indeed uniform-PAC.
Theorem 4.1. For any $\delta \in(0,1)$, if we set $\lambda=1$ and $\beta_{l}=6 \sqrt{d l \log (d l / \delta)}$ for every level $l \in \mathbb{N}$, then there exists a constant $C$ such that with probability at least $1-\delta$, for all $\epsilon>0$, the number of rounds in Algorithm 1 which have sub-optimality no less than $\epsilon$ is bounded by

$$
\sum_{k=1}^{\infty} \mathbb{1}\left\{\max _{\mathbf{x} \in \mathcal{D}_{k}}\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}\right\rangle-\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}_{k}\right\rangle>\epsilon\right\} \leq \frac{C d^{2} \log ^{3}(d /(\delta \epsilon))}{\epsilon^{2}}
$$

Remark 4.2. Theorem 4.1 suggests that Algorithm 1 is uniform-PAC with sample complexity $O\left(d^{2} \log ^{3}(d /(\delta \epsilon)) / \epsilon^{2}\right)$. According to Theorem 3.7, this new algorithm will converge to the optimal policy. Theorem 4.1 also implies an $\widetilde{O}(d \sqrt{T})$ regret for infinite-arm linear bandit problem. This result matches the lower bound $\Omega(d \sqrt{T})$ [6] up to a logarithmic factor. Furthermore, Theorem 4.1 implies that Algorithm 1 is an $(\epsilon, \delta)$-PAC algorithm with sample complexity $\widetilde{O}\left(d^{2} / \epsilon^{2}\right)$. Specifically, if we set $\epsilon=\Delta_{\min }$, Theorem 4.1 implies an $\widetilde{O}\left(d^{2} / \Delta_{\text {min }}^{2}\right)$ sample complexity to identify the best $\operatorname{arm}^{2}$, which matches the sample complexity $\widetilde{O}\left(d \log K / \Delta_{\text {min }}^{2}\right)$ in [19] when $K=\Theta\left(2^{d}\right)$.

[^1]
## 5 Uniform-PAC Bounds for Linear MDPs

In this section, we propose our new FLUTE algorithm (Algorithm 2) for learning linear MDPs, and provide its theoretical guarantee.

Intuition behind FLUTE At a high level, FLUTE inherits the structure of Least-Square Value Iteration with UCB (LSVI-UCB) proposed in [13]. The Bellman optimality equation gives us the following equation:

$$
\begin{equation*}
r_{h}(s, a)+\left[\mathbb{P}_{h} V_{h+1}^{*}\right](s, a)=Q_{h}^{*}(s, a)=\left\langle\boldsymbol{\theta}_{h}^{*}, \boldsymbol{\phi}(s, a)\right\rangle \tag{5.1}
\end{equation*}
$$

where the second equality holds due to Proposition 3.2. (5.1) suggests that in order to learn $Q_{h}^{*}$, it suffices to learn $\boldsymbol{\theta}_{h}^{*}$, which can be roughly regarded as the unknown vector of a linear bandits problem with actions $\phi(s, a)$ and rewards $r_{h}(s, a)+V_{h+1}^{*}\left(s^{\prime}\right)$, where $\left(s, a, s^{\prime}\right)$ belongs to some set $\mathcal{C}$. Since $V_{h+1}^{*}$ is unknown, we use its estimation $V_{h+1}$ to replace it. Therefore, we can apply Algorithm 1 to this equivalent linear bandits problem to obtain our uniform-PAC RL algorithm FLUTE.

Details of FLUTE We now describe the details of FLUTE. For each stage $h$, FLUTE maintains non-overlapping index set $\left\{\mathcal{C}_{h}^{l}\right\}_{l}$, each $\mathcal{C}_{h}^{l}$ contains state-action-next-state triples $\left(s_{h}^{i}, a_{h}^{i}, s_{h+1}^{i}\right)$. Let $S_{1}=1$ and $S_{k}$ denote the number of non-empty sets $\left\{\mathcal{C}_{1}^{l}\right\}_{l}$ at episode $k$ for $k \geq 2$. Instead of maintaining only one estimated optimal value function $V_{k, h}$ and action-value function $Q_{k, h}$ [13], FLUTE maintains a group of estimated value functions $\left\{V_{k, h}^{l}\right\}_{l}$ and action-value functions $\left\{Q_{k, h}^{l}\right\}_{l}$. In detail, at stage $h$, given $\left\{V_{k, h+1}^{l}\right\}_{l}$, FLUTE calculates $\mathbf{w}_{k, h}^{l}$ as the minimizer of the ridge regression problem with training dataset $\left(s_{h}^{i}, a_{h}^{i}, s_{h+1}^{i}\right) \in \mathcal{C}_{h}^{l}$ and targets $V_{k, h+1}^{l}\left(s_{h+1}^{i}\right)$ (Line 8), and defines $Q_{k, h}^{l}$ as the summation of the linear predictor $\left(\mathbf{w}_{k, h}^{l}\right)^{\top} \phi(s, a)$ and a quadratic confidence bonus $\beta_{l} \sqrt{\phi(s, a)^{\top}\left(\boldsymbol{\Sigma}_{k, h}^{l}\right)^{-1} \phi(s, a)}$ (Line 9), where $\beta_{l}$ is the confidence radius for level $l$ and $\boldsymbol{\Sigma}_{k, h}^{l}$ is the covariance matrix for contexts in set $\mathcal{C}_{h}^{l}$. Then FLUTE defines value function $V_{k, h}^{l}$ as the maximum of the minimal value over the first $l$ action-value functions (Line 12). The max-min structure is similar to its counterpart for linear bandits in Algorithm 1, which provides a tighter estimation of the optimal value function and is pivotal to achieve uniform-PAC guarantee.
After constructing action-value functions $\left\{Q_{k, h}^{l}\right\}_{l}$, FLUTE executes the greedy policy induced by the minimal of action-value function $Q_{k, h}^{l}$ over $1 \leq l \leq l_{h-1}^{k}-1$, where $l_{0}^{k}=S_{k}+1$ and $l_{h-1}^{k}$ is the level of the set $\mathcal{A}_{h-1}^{l}$ that we add the triple $\left(s_{h-1}^{k}, a_{h-1}^{k}, s_{h}^{k}\right)$. After obtaining $\left(s_{h}^{k}, a_{h}^{k}, s_{h+1}^{k}\right)$, to decide which set $\mathcal{C}_{h}^{l}$ should this triple be added, FLUTE calculates the confidence bonus $\sqrt{\boldsymbol{\phi}\left(s_{h}^{k}, a_{h}^{k}\right)^{\top}\left(\boldsymbol{\Sigma}_{k, h}^{l}\right)^{-1} \boldsymbol{\phi}\left(s_{h}^{k}, a_{h}^{k}\right)}$ and puts it into the $l$-th set if the confidence bonus is large (Line 19), similar to that of Algorithm 1.
The following theorem shows that FLUTE is uniform-PAC for learning linear MDPs.
Theorem 5.1. Under Assumption 3.1, there exists a positive constant $C$ such that for any $\delta \in(0,1)$, if we set $\lambda=1$ and $\beta_{l}=C d H l \sqrt{\log (d l H / \delta)}$, then with probability at least $1-\delta$, for all $\epsilon>0$, we have

$$
\sum_{k=1}^{\infty} \mathbb{1}\left\{V_{1}^{*}\left(s_{1}^{k}\right)-V_{1}^{\pi_{k}} \geq \epsilon\right\}=O\left(d^{3} H^{5} \log ^{4}(d H /(\delta \epsilon)) / \epsilon^{2}\right) .
$$

Remark 5.2. Theorem 5.1 suggests that algorithm FLUTE is uniform-PAC with sample complexity $O\left(d^{3} H^{5} \log ^{4}(d H /(\delta \epsilon)) / \epsilon^{2}\right)$. According to Theorem 3.7, FLUTE will converge to the optimal policy with high probability. Theorem 5.1 also implies an $\widetilde{O}\left(\sqrt{d^{3} H^{4} T}\right)$ regret for linear MDPs. This result matches the regret bound $\widetilde{O}\left(\sqrt{d^{3} H^{3} T}\right)$ of LSVI-UCB [13] up to a $\sqrt{H}$-factor. Furthermore, Theorem 5.1 also implies FLUTE is an $(\epsilon, \delta)$-PAC with sample complexity $\widetilde{O}\left(d^{3} H^{5} / \epsilon^{2}\right)$, which matches the $O\left(d^{3} H^{3} / \epsilon^{2}\right)$ sample complexity of LSVI-UCB up to an $H$ factor.

Computational complexity As shown in Jin et al. [13], the time complexity of LSVI-UCB is $O\left(d^{2} A H K^{2}\right)$. Compared with the LSVI-UCB algorithm, Algorithm 2 maintains non-overlapping index sets $\left\{\mathcal{C}_{h}^{l}\right\}$ and computes the corresponding optimistic value function for each level $\ell$. Without further assumption on the norm of $\phi(s, a)$, the number of different levels in the first $K$ episodes is at

```
Algorithm 2 Uniform PAC Least-Square Value Iteration (FLUTE)
Require: Regualarization parameter \(\lambda\), confidence radius \(\beta_{l}\)
    Set \(\mathcal{C}_{h}^{l} \leftarrow \emptyset, l \in \mathbb{N}, h \in[H]\) and set the total level \(S_{1}=1\)
    for episode \(k=1,2\), .. do
        Set \(V_{k, H+1}^{l}(s, a)=0\) for all state-action pair \((s, a) \in \mathcal{S} \times \mathcal{A}\) and all level \(l \in\left[S_{k}\right]\)
        for stage \(h=H, H-1, . ., 1\) do
            for all level \(l \in\left[S_{k}\right]\) do
                Set \(\boldsymbol{\Sigma}_{k, h}^{l}=\lambda \mathbf{I}+\sum_{i \in \mathcal{C}_{h}^{l}} \boldsymbol{\phi}\left(s_{h}^{i}, a_{h}^{i}\right) \boldsymbol{\phi}\left(s_{h}^{i}, a_{h}^{i}\right)^{\top}\)
                Set \(\mathbf{b}_{k, h}^{l}=\sum_{i \in \mathcal{C}_{h}^{l}} \boldsymbol{\phi}\left(s_{h}^{i}, a_{h}^{i}\right)\left[r_{h}\left(s_{h}^{i}, a_{h}^{i}\right)+V_{k, h+1}^{l}\left(s_{h+1}^{i}\right)\right]\)
                \(\mathbf{w}_{k, h}^{l} \leftarrow\left(\boldsymbol{\Sigma}_{k, h}^{l}\right)^{-1} \mathbf{b}_{k, h}^{l}\)
                \(Q_{k, h}^{l}(s, a) \leftarrow \min \left\{H,\left(\mathbf{w}_{k, h}^{l}\right)^{\top} \boldsymbol{\phi}(s, a)+\beta_{l} \sqrt{\boldsymbol{\phi}(s, a)^{\top}\left(\boldsymbol{\Sigma}_{k, h}^{l}\right)^{-1} \boldsymbol{\phi}(s, a)}\right\}\)
            end for
            for all level \(l \in\left[S_{k}\right]\) do
                \(V_{k, h}^{l}(s) \leftarrow \max _{a} \min _{1 \leq i \leq l} Q_{k, h}^{i}(s, a)\)
            end for
        end for
        Receive the initial state \(s_{1}^{k}\) and set the current level \(l_{0}^{k}=S_{k}+1\)
        for stage \(h=1,2, . ., H\) do
            Take action \(a_{h}^{k} \leftarrow \operatorname{argmax}_{a} \min _{1 \leq i \leq l_{h-1}^{k}-1} Q_{k, h}^{i}\left(s_{h}^{k}, a\right)\)
            Set level \(l_{h}^{k}=1\)
            while \(\sqrt{\boldsymbol{\phi}\left(s_{h}^{k}, a_{h}^{k}\right)^{\top}\left(\boldsymbol{\Sigma}_{k, h}^{l k}\right)^{-1} \boldsymbol{\phi}\left(s_{h}^{k}, a_{h}^{k}\right)} \leq 2^{-l_{h}^{k}}\) and \(l_{h}^{k} \leq l_{h-1}^{k}-1\) do
                \(l_{h}^{k} \leftarrow l_{h}^{k}+1\)
            end while
            Add element \(k\) to the set \(\mathcal{C}_{h}^{l_{h}^{k}}\)
            Receive the reward \(r_{h}\left(s_{h}^{k}, a_{h}^{k}\right)\) and the next state \(s_{h+1}^{k}\)
        end for
        Set the total level \(S_{k+1}\) as \(S_{k+1}=\max _{l:\left|\mathcal{C}_{1}^{l}\right|>0} l\)
    end for
```

most $K$, which incurs an additional factor of $K$ in the computational complexity in the worst case. However, if we assume the norm of $\phi(s, a)$ equals 1 , then the number of levels in the first $K$ episodes is $O(\log K)$. Thus, since the computational complexity of our algorithm at each level can be bounded by that of LSVI-UCB, the computational complexity of Algorithm 2 is $O\left(d^{2} A H K^{2} \log K\right)$ under the above assumption.

## 6 Proof Outline

In this section, we show the proof roadmap for Theorem 5.1, which consists of three key steps.

## Step 1: Linear function approximates the optimal value function well

We first show that with high probability, for each level $l$, our constructed linear function $\left\langle\mathbf{w}_{k, h}^{l}, \phi(s, a)\right\rangle$ is indeed a "good" estimation of the optimal action-value function $Q_{h}^{*}(s, a)$. By the uniform self-normalized concentration inequality over a specific function class, for any policy $\pi$, any level $l \in \mathbb{N}$ and any state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have the following concentration property:

$$
\left(\mathbf{w}_{k, h}^{l}\right)^{\top} \boldsymbol{\phi}(s, a)-Q_{h}^{\pi}(s, a)=\left[\mathbb{P}_{h}\left(V_{k, h+1}^{l}-V_{h+1}^{\pi}\right)\right](s, a)+\Delta
$$

where $|\Delta| \leq \beta_{l} \sqrt{\phi(s, a)^{\top}\left(\Sigma_{k, h}^{l}\right)^{-1} \phi(s, a)}$. Then, taking a backward induction for each stage $h \in[H]$, let $\Omega=\left\{Q_{k, h}^{l}(s, a) \geq Q_{h}^{*}(s, a), V_{k, h}^{l}(s) \geq V_{h}^{*}(s), \forall(s, a) \in \mathcal{S} \times \mathcal{A}, k, l \in \mathbb{N}, h \in[H]\right\}$ denote the event that the estimated value function $Q_{k, h}^{l}$ and $V_{k, h}^{l}$ upper bounds the optimal value function $Q_{h}^{*}$ and $V_{h}^{*}$. We can show that event $\Omega$ holds with high probability. (More details can be found in Lemmas C. 4 and C.5)

## Step 2: Approximation error decomposition

On the event $\Omega$, the sub-optimality gap in round $k$ is upper bounded by the function value gap between our estimated function $Q_{k, 1}^{l}$ and the value function of our policy $\pi_{k}$.

$$
V_{1}^{*}\left(s_{1}^{k}\right)-V_{1}^{\pi_{k}}\left(s_{1}^{k}\right) \leq \max _{a} \min _{1 \leq l \leq l_{0}^{k}-1} Q_{k, 1}^{l}\left(s_{1}^{k}, a\right)-Q_{1}^{\pi_{k}}\left(s_{1}^{k}, a_{1}^{k}\right) \leq Q_{k, 1}^{l_{1}^{k}-1}\left(s_{1}^{k}, a_{1}^{k}\right)-Q_{1}^{\pi_{k}}\left(s_{1}^{k}, a_{1}^{k}\right)
$$

From now we only focus on the function value gap for level $l_{h}^{k}$. Some elementary calculation gives us

$$
\begin{align*}
& Q_{k, h}^{l_{h}^{k}-1}\left(s_{h}^{k}, a_{h}^{k}\right)-Q_{h}^{\pi_{k}}\left(s_{h}^{k}, a_{h}^{k}\right) \\
& \leq \underbrace{\left.2 \beta_{l_{h}^{k}-1} \sqrt{\boldsymbol{\phi}\left(s_{h}^{k}, a_{h}^{k}\right)^{\top}\left(\boldsymbol{\Sigma}_{k, h}^{l_{h}^{k}-1}\right)^{-1} \phi\left(s_{h}^{k}, a_{h}^{k}\right.}\right)}_{I_{h}^{k}}+Q_{k, h+1}^{l_{h+1}^{k}-1}\left(s_{h+1}^{k}, a_{h+1}^{k}\right)-Q_{h+1}^{\pi_{k}}\left(s_{h+1}^{k}, a_{h+1}^{k}\right) \\
& \quad+\underbrace{\left[\mathbb{P}_{h}\left(V_{k, h+1}^{l_{h}^{k}-1}-V_{h+1}^{\pi_{k}}\right)\right]\left(s_{h}^{k}, a_{h}^{k}\right)-\left(V_{k, h+1}^{l_{h}^{k}-1}\left(s_{h+1}^{k}\right)-V_{h+1}^{\pi_{k}}\left(s_{h+1}^{k}\right)\right)}_{\Delta_{k, h}} \tag{6.1}
\end{align*}
$$

Therefore, by telescoping (6.1) from stage $h=1$ to $H$, we conclude that the sub-optimality gap $V_{1}^{*}\left(s_{1}^{k}\right)-V_{1}^{\pi_{k}}\left(s_{1}^{k}\right)$ is upper bounded by the summation of the bonus $I_{h}^{k}$ and $\Delta_{k, h}$. The summation of the bonus $\sum I_{h}^{k}$ is the dominating error term. According to the rule of level $l$, if $k \in \mathcal{C}_{h}^{l}$ at stage $h \in[H]$, then $I_{h}^{k}$ satisfies $\sqrt{\boldsymbol{\phi}\left(s_{h}^{k}, a_{h}^{k}\right)^{\top}\left(\boldsymbol{\Sigma}_{k, h}^{l-1}\right)^{-1} \boldsymbol{\phi}\left(s_{h}^{k}, a_{h}^{k}\right)} \leq 2^{-(l-1)}$. Furthermore, the number of elements added into set $\mathcal{C}_{h}^{l}$ can be upper bounded by $\left|\mathcal{C}_{h}^{l}\right| \leq 17 d l H 4^{l}$ (See Lemma C.1). Thus we can bound the summation of $I_{h}^{k}$. For $\sum \Delta_{k, h}$, it is worth noting that $\Delta_{k, h}$ forms a martingale difference sequence, therefore by the standard Azuma-Hoeffding inequality, $\sum \Delta_{k, h}$ can be bounded by some non-dominating terms. Both of these two bounds will be used in the next step.

## Step 3: From upper confidence bonus to uniform-PAC sample complexity

In Step 2 we have already bounded the sub-optimality gap by the summation of bonus terms. In this step, we show how to transform the gap into the final uniform-PAC sample complexity. Instead of studying any accuracy $\epsilon$ directly, we focus on a special case where $\epsilon=H / 2^{i}(i \in \mathbb{N})$, which can be easily generalized to the general case. For each fixed $\epsilon=H / 2^{i}(i \in \mathbb{N})$, let $\mathcal{K}$ denote the set $\mathcal{K}=\left\{k \mid V_{1}^{*}\left(s_{1}^{k}\right)-V_{1}^{\pi_{k}} \geq \epsilon\right\}$ and $m=|\mathcal{K}|$. On the one hand, according to the definition of set $\mathcal{K}$, the summation of regret in episode $k(k \in \mathcal{K})$ is lower bounded by $m \epsilon$. On the other hand, according to Step 2, the summation of sub-optimality gaps of episode $k(k \in \mathcal{K})$, is upper bound by

$$
\begin{equation*}
\sum_{k \in \mathcal{K}}\left[V_{1}^{*}\left(s_{1}^{k}\right)-V_{1}^{\pi_{k}}\right] \leq \underbrace{\sum_{k \in \mathcal{K}} \sum_{h=1}^{H} 2 \beta_{l_{h}^{k}-1} 2^{-\left(l_{h}^{k}-1\right)}}_{J_{1}}+\underbrace{\sum_{k \in \mathcal{K}} \sum_{h=1}^{H} \Delta_{k, h}}_{J_{2}} \tag{6.2}
\end{equation*}
$$

To further bound $J_{1}$, we divide those episode-stage pairs $(k, h) \in \mathbb{N} \times[H]$ into two categories: $\mathcal{S}_{1}=\left\{(k, h) \mid 2 \beta_{l_{h}^{k}-1} 2^{-\left(l_{h}^{k}-1\right)} \leq \epsilon /(2 H)\right\}$ and $\mathcal{S}_{2}=\left\{(k, h) \mid 2 \beta_{l_{h}^{k}-1} 2^{-\left(l_{h}^{k}-1\right)}>\epsilon /(2 H)\right\}$. For the first category $\mathcal{S}_{1}$, the summation of terms $I_{h}^{k}$ in this category is upper bound by

$$
\begin{equation*}
\sum_{k \in \mathcal{K}} \sum_{h=1}^{H} \mathbb{1}\left\{(k, h) \in \mathcal{S}_{1}\right\} 2 \beta_{l_{h}^{k}-1} 2^{-\left(l_{h}^{k}-1\right)} \leq \sum_{k \in \mathcal{K}} \sum_{h=1}^{H} \frac{\epsilon}{2 H}=\frac{m \epsilon}{2} \tag{6.3}
\end{equation*}
$$

For any episode-stage pair $(k, h)$ in the second category $\mathcal{S}_{2}$, the level $l_{h}^{k}$ satisfies $2^{l^{k}} \leq \widetilde{O}\left(d H^{2} / \epsilon\right)$ due to the choice of $\beta_{l_{h}^{k}-1}$. Suppose $l^{\prime}$ is the maximum level that satisfies $2^{l} \leq \widetilde{O}\left(d H^{2} / \epsilon\right)$ and for each level $l \leq l^{\prime}$, the cardinality of set $\mathcal{C}_{h}^{l}$ can be upper bounded by $\left|\mathcal{C}_{h}^{l}\right| \leq 17 d l H 4^{l}$. Thus, the summation of terms $I_{h}^{k}$ in category $\mathcal{S}_{2}$ is upper bound by

$$
\sum_{k \in \mathcal{K}} \sum_{h=1}^{H} \mathbb{1}\left\{(k, h) \in \mathcal{S}_{2}\right\} 2 \beta_{l_{h}^{k}-1} 2^{-\left(l_{h}^{k}-1\right)} \leq \sum_{k \in \mathcal{K}} \sum_{h=1}^{H} \sum_{l=1}^{l^{\prime}} \mathbb{1}\left\{l_{h}^{k}=l\right\} 2 \beta_{l-1} 2^{-(l-1)}
$$

$$
\begin{align*}
& =\sum_{h=1}^{H} \sum_{l=1}^{l^{\prime}} 2 \beta_{l-1} 2^{-(l-1)} \sum_{k \in \mathcal{K}} \mathbb{1}\left\{l_{h}^{k}=l\right\} \\
& \leq \sum_{h=1}^{H} \sum_{l=1}^{l^{\prime}} 2 \beta_{l-1} 2^{-(l-1)} 17 d l H 4^{l} \\
& =\widetilde{O}\left(d^{3} H^{5} / \epsilon\right) \tag{6.4}
\end{align*}
$$

Back to (6.2), for the second term $J_{2}$, according to Azuma-Hoeffding inequality, it can be controlled by $\widetilde{O}(H \sqrt{H m})$. Therefore, combining (6.3), (6.4) with the bound of $J_{2}$, we have

$$
m \epsilon \leq \sum_{k \in \mathcal{K}} V_{1}^{*}\left(s_{1}^{k}\right)-V_{1}^{\pi_{k}} \leq m \epsilon / 2+\widetilde{O}\left(d^{3} H^{5} / \epsilon\right)+\widetilde{O}(H \sqrt{H m})
$$

and it implies that the number of episodes with a sub-optimality gap greater than $\epsilon$ is bounded by $\widetilde{O}\left(d^{3} H^{5} / \epsilon^{2}\right)$. This completes the proof.

## 7 Conclusion and Future Work

In this work, we proposed two novel uniform-PAC algorithms for linear bandits and RL with linear function approximation, with the nearly state-of-the-art sample complexity. To the best of our knowledge, these are the very first results to show that linear bandits and RL with linear function approximation can also achieve uniform-PAC guarantees, similar to the tabular RL setting. We leave proving their corresponding lower bounds and proposing algorithms with near-optimal uniform-PAC sample complexity as future work.

## Acknowledgments and Disclosure of Funding

We thank the anonymous reviewers for their helpful comments. Part of this work was done when JH, DZ and QG participated the Theory of Reinforcement Learning program at the Simons Institute for the Theory of Computing in Fall 2020. JH, DZ and QG are partially supported by the National Science Foundation CAREER Award 1906169, IIS-1904183, BIGDATA IIS-1855099 and AWS Machine Learning Research Award. The views and conclusions contained in this paper are those of the authors and should not be interpreted as representing any funding agencies.

## References

[1] Abbasi-Yadkori, Y., PÁl, D. and Szepesvári, C. (2011). Improved algorithms for linear stochastic bandits. In Advances in Neural Information Processing Systems.
[2] Auer, P. (2002). Using confidence bounds for exploitation-exploration trade-offs. Journal of Machine Learning Research 3 397-422.
[3] Ayoub, A., Jia, Z., Szepesvari, C., Wang, M. and Yang, L. F. (2020). Model-based reinforcement learning with value-targeted regression. arXiv preprint arXiv:2006.01107 .
[4] Cesa-Bianchi, N. and Lugosi, G. (2006). Prediction, learning, and games. Cambridge university press.
[5] Chu, W., Li, L., Reyzin, L. and Schapire, R. (2011). Contextual bandits with linear payoff functions. In Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics. JMLR Workshop and Conference Proceedings.
[6] Dani, V., Hayes, T. P. and Kakade, S. M. (2008). Stochastic linear optimization under bandit feedback.
[7] Dann, C., Lattimore, T. and Brunskill, E. (2017). Unifying pac and regret: Uniform pac bounds for episodic reinforcement learning. In Advances in Neural Information Processing Systems.
[8] Fiez, T., Jain, L., Jamieson, K. and Ratliff, L. (2019). Sequential experimental design for transductive linear bandits. arXiv preprint arXiv:1906.08399.
[9] Ghavamzadeh, M., Kappen, H., Azar, M. and Munos, R. (2011). Speedy q-learning. Advances in neural information processing systems 24 2411-2419.
[10] Jia, Z., Yang, L., Szepesvari, C. and Wang, M. (2020). Model-based reinforcement learning with value-targeted regression. In Learning for Dynamics and Control. PMLR.
[11] Jiang, N., Krishnamurthy, A., Agarwal, A., Langford, J. and Schapire, R. E. (2017). Contextual decision processes with low bellman rank are pac-learnable. In Proceedings of the 34th International Conference on Machine Learning-Volume 70. JMLR. org.
[12] Jin, C., Allen-Zhu, Z., Bubeck, S. and Jordan, M. I. (2018). Is q-learning provably efficient? In Advances in Neural Information Processing Systems.
[13] Jin, C., Yang, Z., Wang, Z. and Jordan, M. I. (2019). Provably efficient reinforcement learning with linear function approximation. arXiv preprint arXiv:1907.05388.
[14] Kakade, S. M. et al. (2003). On the sample complexity of reinforcement learning. Ph.D. thesis, University of London London, England.
[15] Karnin, Z. (2016). Verification based solution for structured mab problems. In Proceedings of the 30th International Conference on Neural Information Processing Systems.
[16] Li, L., Chu, W., LANGFORD, J. and SChapire, R. E. (2010). A contextual-bandit approach to personalized news article recommendation. In Proceedings of the 19th international conference on World wide web.
[17] Li, Y., WANG, Y. and Zhou, Y. (2019). Nearly minimax-optimal regret for linearly parameterized bandits. In Conference on Learning Theory. PMLR.
[18] Modi, A., Jiang, N., Tewari, A. and Singh, S. (2020). Sample complexity of reinforcement learning using linearly combined model ensembles. In International Conference on Artificial Intelligence and Statistics. PMLR.
[19] Soare, M., Lazaric, A. and Munos, R. (2014). Best-arm identification in linear bandits. arXiv preprint arXiv:1409.6110.
[20] Tao, C., Blanco, S. and Zhou, Y. (2018). Best arm identification in linear bandits with linear dimension dependency. In International Conference on Machine Learning. PMLR.
[21] Xu, L., Honda, J. and Sugiyama, M. (2018). A fully adaptive algorithm for pure exploration in linear bandits. In International Conference on Artificial Intelligence and Statistics. PMLR.
[22] Yang, L. and Wang, M. (2019). Sample-optimal parametric q-learning using linearly additive features. In International Conference on Machine Learning.
[23] Zanette, A. and Brunskill, E. (2019). Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. arXiv preprint arXiv:1901.00210 .
[24] Zanette, A., LaZaric, A., Kochenderfer, M. and Brunskill, E. (2020). Learning near optimal policies with low inherent bellman error. arXiv preprint arXiv:2003.00153 .
[25] Zhang, Z., Yang, J., Ji, X. and Du, S. S. (2021). Variance-aware confidence set: Variancedependent bound for linear bandits and horizon-free bound for linear mixture mdp. arXiv preprint arXiv:2101.12745.
[26] Zhang, Z., ZHOU, Y. and Ji, X. (2020). Almost optimal model-free reinforcement learning via reference-advantage decomposition. arXiv preprint arXiv:2004.10019.
[27] Zhou, D., Gu, Q. and Szepesvari, C. (2020). Nearly minimax optimal reinforcement learning for linear mixture markov decision processes. arXiv preprint arXiv:2012.08507.

## Checklist

1. For all authors...
(a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
(b) Did you describe the limitations of your work? [Yes]
(c) Did you discuss any potential negative societal impacts of your work? [N/A] Our work studies the uniform-PAC bounds for RL with function approximation, which is a pure theoretical problem and does not have any negative social impact.
(d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
(a) Did you state the full set of assumptions of all theoretical results? [Yes]
(b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
(a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
(a) If your work uses existing assets, did you cite the creators? [N/A]
(b) Did you mention the license of the assets? [N/A]
(c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

## A OFUL Algorithm is not Uniform-PAC

In this section, we consider a variant of the OFUL algorithm [1]. Then we will present a hard-to-learn linear bandit instance and show that the variant of OFUL algorithm cannot have the uniform-PAC guarantee for this instance.
In the original OFUL algorithm [1], following their notation, the agent selects the action by $\mathbf{x}_{k}=$ $\operatorname{argmax}_{(\mathbf{x}, \boldsymbol{\theta}) \in \mathcal{D}_{k} \times \Theta_{k-1}}\langle\mathbf{x}, \boldsymbol{\theta}\rangle$. Here we consider a variant of OFUL, where the agent selects the action by $\mathbf{x}_{k}=\operatorname{argmax}_{(\mathbf{x}, \boldsymbol{\theta}) \in \mathcal{D}_{k} \times \Theta_{k-1} \cap B(1)}\langle\mathbf{x}, \boldsymbol{\theta}\rangle$, where $B(1)$ is a unit ball centered at zero.
We consider a special contextual linear bandit instance with dimension $d=2, \boldsymbol{\theta}^{*}=(0,1)$, and zero noise. The action set in the first $K$ ( $K$ is an arbitrary parameter that can be chosen later) rounds is $\{(1,0),(-1,0)\}$ and the action set in the following $\log K$ rounds is $\{(0,1),(0,-1)\}$. So the reward in each step can only be 1 or -1 . The agent will randomly choose one action if both actions attain $\operatorname{argmax}_{(\mathbf{x}, \boldsymbol{\theta}) \in \mathcal{D}_{k} \times \Theta_{k-1} \cap B(1)}\langle\mathbf{x}, \boldsymbol{\theta}\rangle$. We can show that, in the first $K$ round, the confidence radius increases since the determinant of the covariance matrix increases, and it will not provide any information about the second dimension of the vector $\boldsymbol{\theta}^{*}$ since the two actions are orthogonal to $\boldsymbol{\theta}^{*}=(0,1)$. After the first $K$ rounds, the confidence radius will be in the order of $\log K$, and the covariance matrix $\boldsymbol{\Sigma}_{K}$ is a diagonal matrix and in the order of $\operatorname{diag}(K, \log K)$. We can show that both $\boldsymbol{\theta}=(0,1)$ and $\boldsymbol{\theta}=(0,-1)$ belong to $\Theta_{k-1} \cap B(1)$, and thus attain the maximum of $\operatorname{argmax}_{(\mathbf{x}, \boldsymbol{\theta}) \in \mathcal{D}_{k} \times \Theta_{k-1} \cap B(1)}\langle\mathbf{x}, \boldsymbol{\theta}\rangle$. Therefore, the agent will almost 'randomly' pick one of the two actions in the later $\log K$ rounds. The random selection leads to a 1-suboptimality gap for about half of the $\log K$ rounds, which indicates that OFUL cannot be uniform-PAC for any finite $f(\epsilon, \delta)$ on this bandit problem, by selecting $\log K>f(\epsilon, \delta)$.

The above reasoning can be extended to the original OFUL algorithm with a more involved argument.

## B Proof for Theorem 4.1

In this section, we provide the proof of Theorems 4.1 and for simplicity, let $\mathcal{C}_{k}^{l}$ denote the index set $\mathcal{C}^{l}$ at the beginning of round $k$. We first propose the following lemmas.
Lemma B.1. Suppose $\lambda \geq 1$, then for each level $l \in \mathbb{N}$ and round $k \in \mathbb{N}$, the number of elements in the index $\operatorname{set} \mathcal{C}_{k}^{l}$ is upper bounded by

$$
\left|\mathcal{C}_{k}^{l}\right| \leq 17 d l 4^{l}
$$

## Proof. See Appendix D.1.

Lemma B. 1 suggests that $\mathcal{C}^{l}$ is always a finite set.
Lemma B.2. If we set $\lambda=1$ and $\beta_{l}=6 \sqrt{d l \log (d l / \delta)}$ for every level $l \in \mathbb{N}$, then with probability at least $1-\delta$, for all level $l \in \mathbb{N}$ and all round $k \in \mathbb{N}$, we have

$$
\left\|\mathbf{w}_{k}^{l}-\boldsymbol{\mu}^{*}\right\|_{\boldsymbol{\Sigma}_{k}^{l}} \leq \beta_{l} .
$$

## Proof. See Appendix D.2.

For simplicity, let $\mathcal{E}$ denotes the event that the conclusion of Lemma B. 2 holds. Therefore, Lemma B. 2 suggests $\operatorname{Pr}(\mathcal{E}) \geq 1-\delta$.

Proof of Theorem 4.1. On the event $\mathcal{E}$, for all level $l \in \mathbb{N}$, round $k \in \mathbb{N}$ and action $\mathbf{x} \in \mathcal{D}_{k}$, we have

$$
\begin{align*}
\left(\mathbf{w}_{k}^{l}\right)^{\top} \mathbf{x}+\beta_{l} \sqrt{\mathbf{x}^{\top}\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1} \mathbf{x}}-\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}\right\rangle & =\left(\mathbf{w}_{k}^{l}-\boldsymbol{\mu}^{*}\right)^{\top} \mathbf{x}+\beta_{l} \sqrt{\mathbf{x}^{\top}\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1} \mathbf{x}} \\
& \geq \beta_{l} \sqrt{\mathbf{x}^{\top}\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1} \mathbf{x}}-\left\|\mathbf{w}_{k}^{l}-\boldsymbol{\mu}^{*}\right\|_{\boldsymbol{\Sigma}_{k}^{l}}\|\mathbf{x}\|_{\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1}} \\
& \geq \beta_{l} \sqrt{\mathbf{x}^{\top}\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1}}-\beta_{l} \sqrt{\mathbf{x}^{\top}\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1}} \\
& =0 \tag{B.1}
\end{align*}
$$

where the first inequality holds due to Cauchy-Schwarz inequality and the second inequality holds due to the definition of event $\mathcal{E}$. (B.1) implies that the estimated reward for each action $\mathbf{x} \in \mathcal{D}_{k}$ at level $l:\left(\mathbf{w}_{k}^{l}\right)^{\top} \mathbf{x}+\beta_{l} \sqrt{\mathbf{x}^{\top}\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1} \mathbf{x}}$ is an upper confidence bound of the expected reward $\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}\right\rangle$. Thus, for each action $\mathbf{x} \in \mathcal{D}_{k}$, we have

$$
\begin{equation*}
\min _{1 \leq l \leq S_{k}}\left(\mathbf{w}_{k}^{l}\right)^{\top} \mathbf{x}+\beta_{l} \sqrt{\mathbf{x}^{\top}\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1} \mathbf{x}} \geq \min _{1 \leq l \leq S_{k}}\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}\right\rangle=\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}\right\rangle \tag{B.2}
\end{equation*}
$$

Therefore, for the sub-optimality gap at round $k$, we have

$$
\begin{align*}
\max _{\mathbf{x} \in \mathcal{D}_{k}}\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}\right\rangle-\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}_{k}\right\rangle & \leq \max _{\mathbf{x} \in \mathcal{D}_{k}} \min _{1 \leq l \leq S_{k}}\left(\mathbf{w}_{k}^{l}\right)^{\top} \mathbf{x}+\beta_{l} \sqrt{\mathbf{x}^{\top}\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1} \mathbf{x}}-\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}_{k}\right\rangle \\
& =\min _{1 \leq l \leq S_{k}}\left(\mathbf{w}_{k}^{l}\right)^{\top} \mathbf{x}_{k}+\beta_{l} \sqrt{\mathbf{x}_{k}^{\top}\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1} \mathbf{x}_{k}}-\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}_{k}\right\rangle, \tag{B.3}
\end{align*}
$$

where the first inequality holds due to (B.2) and the second equality holds due to the policy in Algorithm 1 (line 8). Thus, for each round $k \in \mathbb{N}$, if the level $l_{k}>1$, we have

$$
\begin{aligned}
\max _{\mathbf{x} \in \mathcal{D}_{k}}\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}\right\rangle-\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}_{k}\right\rangle & \leq\left(\mathbf{w}_{k}^{l_{k}-1}\right)^{\top} \mathbf{x}_{k}+\beta_{l_{k}-1} \sqrt{\mathbf{x}_{k}^{\top}\left(\boldsymbol{\Sigma}_{k}^{l_{k}-1}\right)^{-1} \mathbf{x}_{k}}-\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}_{k}\right\rangle \\
& =\left(\mathbf{w}_{k}^{l_{k}-1}-\boldsymbol{\mu}^{*}\right)^{\top} \mathbf{x}_{k}+\beta_{l_{k}-1} \sqrt{\mathbf{x}_{k}^{\top}\left(\boldsymbol{\Sigma}_{k}^{l_{k}-1}\right)^{-1} \mathbf{x}_{k}} \\
& \leq\left\|\mathbf{w}_{k}^{l_{k}-1}-\boldsymbol{\mu}^{*}\right\|_{\boldsymbol{\Sigma}_{k}^{l_{k}-1}}\left\|\mathbf{x}_{k}\right\|_{\left(\boldsymbol{\Sigma}_{k}^{l_{k}-1}\right)^{-1}}+\beta_{l_{k}-1} \sqrt{\mathbf{x}_{k}^{\top}\left(\boldsymbol{\Sigma}_{k}^{l_{k}-1}\right)^{-1} \mathbf{x}_{k}} \\
& \leq 2 \beta_{l_{k}-1} \sqrt{\mathbf{x}_{k}^{\top}\left(\boldsymbol{\Sigma}_{k}^{l_{k}-1}\right)^{-1} \mathbf{x}_{k}} \\
& \leq 2 \beta_{l_{k}-1} \times 2^{-\left(l_{k}-1\right)}
\end{aligned}
$$

where the first inequality holds due to (B.3) with the fact that $l_{k}-1 \leq S_{k}$, the second inequality holds due to Cauchy-Schwarz inequality, the third inequality holds due to the definition of event $\mathcal{E}$ and the last inequality holds due to the definition of level $l_{k}$ in Algorithm 1 (line 10 to line 11). Since we set the parameter $\beta_{l}=6 \sqrt{d l \log (d l / \delta)}$, there exists a large constant $C$ such that for any level $l$ satisfied $2^{l} \geq C \sqrt{d \log ^{2}(d /(\delta \epsilon))} / \epsilon$, we have $2 \beta_{l-1} \times 2^{-(l-1)} \leq \epsilon$. For simplicity, we denote the minimum level $m=\left[\log \left(C \sqrt{d \log ^{2}(d /(\delta \epsilon))} / \epsilon\right)\right]$. Then for each round $k$, if level $l_{k}>m$, we have

$$
\begin{equation*}
\max _{\mathbf{x} \in \mathcal{D}_{k}}\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}\right\rangle-\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}_{k}\right\rangle \leq 2 \beta_{l_{k}-1} \times 2^{-\left(l_{k}-1\right)} \leq \epsilon \tag{B.4}
\end{equation*}
$$

Thus, for any $\epsilon>0$, we have

$$
\begin{aligned}
\sum_{k=1}^{\infty} \mathbb{1}\left\{\max _{\mathbf{x} \in \mathcal{D}_{k}}\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}\right\rangle-\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}_{k}\right\rangle>\epsilon\right\} & \leq \sum_{k=1}^{\infty} \mathbb{1}\left\{l_{k} \leq m\right\} \\
& =\sum_{k=1}^{\infty} \sum_{l=1}^{m} \mathbb{1}\left\{l_{k}=l\right\} \\
& =\sum_{l=1}^{m} \sum_{k=1}^{\infty} \mathbb{1}\left\{l_{k}=l\right\}
\end{aligned}
$$

where the inequality holds due to (B.4). According to Lemma B.1, the number of rounds with sub-optimality more than $\epsilon$ can be further bounded by

$$
\begin{aligned}
\sum_{k=1}^{\infty} \mathbb{1}\left\{\max _{\mathbf{x} \in \mathcal{D}_{k}}\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}\right\rangle-\left\langle\boldsymbol{\mu}^{*}, \mathbf{x}_{k}\right\rangle>\epsilon\right\} & \leq \sum_{l=1}^{m} \sum_{k=1}^{\infty} \mathbb{1}\left\{l_{k}=l\right\} \\
& \leq \sum_{l=1}^{m} 17 d l 4^{l} \\
& \leq C^{\prime} d^{2} \log ^{3}(d /(\delta \epsilon)) / \epsilon^{2}
\end{aligned}
$$

where the second inequality holds due to Lemma B. 1 and the last inequality holds due to the definition of $m$ with the fact that $\sum_{l=1}^{m} l 4^{l} \leq m 4^{m+1}$. Thus, we finish the proof of Theorem 4.1.

## C Proof of Theorem 5.1

In this section, we provide the proof of Theorems 5.1 and for simplicity, let $\mathcal{C}_{k, h}^{l}$ denote the index set $\mathcal{C}_{h}^{l}$ at the beginning of episode $k$. We first propose the following lemmas.
Lemma C.1. Suppose the parameter $\lambda$ satisfies $\lambda \geq 1$, then for each level $l \in \mathbb{N}$ each stage $h \in[H]$ and each episode $k \in \mathbb{N}$, the number of elements in the set $\mathcal{C}_{k, h}^{l}$ is upper bounded by

$$
\left|\mathcal{C}_{k, h}^{l}\right| \leq 17 d l h 4^{l}
$$

Proof. See Appendix E.1.
Lemma C.2. Under Assumption 3.1, for each stage $h \in[H]$, each level $l \in \mathbb{N}$ and each episode $k \in \mathbb{N}$, the norm of weight vector $\mathbf{w}_{k, h}^{l}$ can be upper bounded by

$$
\left\|\mathbf{w}_{k, h}^{l}\right\|_{2} \leq \frac{9 d 2^{l} \sqrt{H^{3} l}}{\sqrt{\lambda}}
$$

Proof. See Appendix E.2.
Lemma C.3. Suppose the parameter $\lambda=1$, then there exists a large constant $C$, such that with probability $1-\delta / 2$, for all stage $h \in[H]$, all episode $k \in \mathbb{N}$ and all level $l \in \mathbb{N}$, we have

$$
\left\|\sum_{i \in \mathcal{C}_{k, h}^{l}} \phi\left(s_{h}^{i}, a_{h}^{i}\right)\left[V_{k, h+1}^{l}\left(s_{h+1}^{i}\right)-\left[\mathbb{P}_{h} V_{k, h+1}^{l}\right]\left(s_{h}^{i}, a_{h}^{i}\right)\right]\right\|_{\left(\boldsymbol{\Sigma}_{k, h}^{l}\right)^{-1}} \leq C d H l \sqrt{\log \left(d l H \beta_{l}^{2} / \delta\right)}
$$

Proof. See Appendix E.3.
For simplicity, let $\mathcal{E}$ denote the event that the conclusion of Lemma C. 3 holds. Therefore, Lemma C. 3 shows that $\operatorname{Pr}(\mathcal{E}) \geq 1-\delta / 2$.

Lemma C.4. Suppose $\lambda=1$ and $\beta_{l}=C d H l \sqrt{\log (d l H / \delta)}$ with a large constant $C$, then on the event $\mathcal{E}$, for all state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, stage $h \in[H]$, episode $k \in \mathbb{N}$, level $l \in \mathbb{N}$ and any policy $\pi$, we have

$$
\left(\mathbf{w}_{k, h}^{l}\right)^{\top} \boldsymbol{\phi}(s, a)-Q_{h}^{\pi}(s, a)=\left[\mathbb{P}_{h}\left(V_{k, h+1}^{l}-V_{h+1}^{\pi}\right)\right](s, a)+\Delta
$$

where $|\Delta| \leq \beta_{l} \sqrt{\phi(s, a)^{\top}\left(\Sigma_{k, h}^{l}\right)^{-1} \phi(s, a)}$.
Proof. See Appendix E.4.
Lemma C.5. On the event $\mathcal{E}$, for all state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, stage $h \in[H]$, episode $k \in \mathbb{N}$ and level $l \in \mathbb{N}$, we have

$$
Q_{k, h}^{l}(s, a) \geq Q_{h}^{*}(s, a), V_{k, h}^{l}(s) \geq V_{h}^{*}(s)
$$

Proof. See Appendix E.5.
Now we begin to prove Theorem 5.1.
Proof of Theorem 5.1. Firstly, we focus on the special case that $\epsilon=H 2^{-i}(i \in \mathbb{N})$. Since $\epsilon=H 2^{-i}$ and we set the parameter $\beta_{l}=C d H l \sqrt{\log (d l H / \delta)}$, there exists a large constant $C^{\prime}$ such that for any level $l$ satisfied $2^{l} \geq C^{\prime} d H^{2} \log ^{1.5}(d H /(\delta \epsilon)) / \epsilon$, we have

$$
4 \beta_{l-1} 2^{-l}=C d H l \sqrt{\log \left(d l H \beta_{l-1}^{2} / \delta\right)} 2^{-l} \leq \epsilon /(2 H)
$$

For simplicity, we denote the maximum level $l^{\prime}$ as $l^{\prime}=\left[\log \left(C^{\prime} d H^{2} \log ^{1.5}(d H /(\delta \epsilon)) / \epsilon\right)\right]$.

Now, let $k_{0}=0$, and for each $i \in \mathbb{N}$, we denote $k_{i}$ as the minimum index of the episode where the sub-optimality gap is more than $\epsilon$, such that

$$
\begin{equation*}
k_{i}=\min \left\{k: k>k_{i-1}, V_{1}^{*}\left(s_{1}^{k}\right)-V_{1}^{\pi_{k}}\left(s_{1}^{k}\right) \geq \epsilon\right\} . \tag{C.1}
\end{equation*}
$$

Now, we denote the set $K=\left\{k_{i}: i \in \mathbb{N}, k_{i}<+\infty\right\}$ and we assume $K=\left\{k_{1}, . ., k_{m}\right\}$. According to the definition of $k_{i}$ in (C.1), we have

$$
\begin{equation*}
\sum_{i=1}^{m} V_{1}^{*}\left(s_{1}^{k_{i}}\right)-V_{1}^{\pi_{k_{i}}}\left(s_{1}^{k_{i}}\right) \geq m \epsilon \tag{C.2}
\end{equation*}
$$

On the other hand, for each episode $k \in \mathbb{N}$ with total level $S_{k}$ at the beginning of episode $k$, we have

$$
\begin{align*}
V_{1}^{*}\left(s_{1}^{k}\right)-V_{1}^{\pi_{k}}\left(s_{1}^{k}\right) & =\max _{a} Q_{1}^{*}\left(s_{1}^{k}, a\right)-Q_{1}^{\pi_{k}}\left(s_{1}^{k}, a_{1}^{k}\right) \\
& \leq \max _{a} \min _{1 \leq l \leq S_{k}} Q_{k, 1}^{l}\left(s_{1}^{k}, a\right)-Q_{1}^{\pi_{k}}\left(s_{1}^{k}, a_{1}^{k}\right) \\
& =\min _{1 \leq l \leq S_{k}} Q_{k, 1}^{l}\left(s_{1}^{k}, a_{1}^{k}\right)-Q_{1}^{\pi_{k}}\left(s_{1}^{k}, a_{1}^{k}\right) \\
& \leq Q_{k, 1}^{l_{1}^{k}-1}\left(s_{1}^{k}, a_{1}^{k}\right)-Q_{1}^{\pi_{k}}\left(s_{1}^{k}, a_{1}^{k}\right) \tag{C.3}
\end{align*}
$$

where the first inequality holds due to Lemma C.5, the third equation holds due to the policy in Algorithm 2 and the last inequality holds due to the fact that $l_{1}^{k}-1 \leq S_{k}$. Furthermore, for each stage $h \in[H]$ and each episode $k \in \mathbb{N}$, we have

$$
\begin{align*}
& Q_{k, h}^{l_{h}^{k}-1}\left(s_{h}^{k}, a_{h}^{k}\right)-Q_{h}^{\pi_{k}}\left(s_{h}^{k}, a_{h}^{k}\right) \\
& \leq\left(\mathbf{w}_{k, h}^{l_{h}^{k}-1}\right)^{\top} \boldsymbol{\phi}\left(s_{h}^{k}, a_{h}^{k}\right)+\beta_{l_{h}^{k}-1} \sqrt{\boldsymbol{\phi}\left(s_{h}^{k}, a_{h}^{k}\right)^{\top}\left(\boldsymbol{\Sigma}_{k, h}^{l_{h}^{k}-1}\right)^{-1} \boldsymbol{\phi}(s, a)}-Q_{h}^{\pi_{k}}\left(s_{h}^{k}, a_{h}^{k}\right) \\
& \leq 2 \beta_{l_{h}^{k}-1} \sqrt{\boldsymbol{\phi}\left(s_{h}^{k}, a_{h}^{k}\right)^{\top}\left(\boldsymbol{\Sigma}_{k, h}^{l_{h}^{k}-1}\right)^{-1} \boldsymbol{\phi}\left(s_{h}^{k}, a_{h}^{k}\right)}+\left[\mathbb{P}_{h}\left(V_{k, h+1}^{l}-V_{h+1}^{\pi_{k}}\right)\right]\left(s_{h}^{k}, a_{h}^{k}\right) \\
& \leq 2 \beta_{l_{h}^{k}-1} 2^{-\left(l_{h}^{k}-1\right)}+\left[\mathbb{P}_{h}\left(V_{k, h+1}^{l_{h}^{k}-1}-V_{h+1}^{\pi_{k}}\right)\right]\left(s_{h}^{k}, a_{h}^{k}\right) \\
& = \\
& 4 \beta_{l_{h}^{k}-1} 2^{-l_{h}^{k}}+\underbrace{\left[\mathbb{P}_{h}\left(V_{k, h+1}^{l_{h}^{k}-1}-V_{h+1}^{\pi_{k}}\right)\right]\left(s_{h}^{k}, a_{h}^{k}\right)-\left(V_{k, h+1}^{l_{h}^{k}-1}\left(s_{h+1}^{k}\right)-V_{h+1}^{\pi_{k}}\left(s_{h+1}^{k}\right)\right)}_{\Delta_{k, h}}  \tag{C.4}\\
& \quad+V_{k, h+1}^{l_{h}^{k}-1}\left(s_{h+1}^{k}\right)-V_{h+1}^{\pi_{k}}\left(s_{h+1}^{k}\right),
\end{align*}
$$

where the first inequality holds due to the definition of value function $Q_{k, h}^{l}$ in Algorithm 2, the second inequality holds due to Lemma C. 4 and the last inequality holds due to the definition of level $l_{h}^{k}$ in Algorithm 2. Furthermore, for the term $V_{k, h+1}^{l_{h}^{k}-1}\left(s_{h+1}^{k}\right)$, it can be upper bounded by

$$
\begin{align*}
V_{k, h+1}^{l_{h}^{k}-1}\left(s_{h+1}^{k}\right) & =\max _{a} \min _{1 \leq l \leq l_{h}^{k}-1} Q_{k, h+1}^{l}\left(s_{h+1}^{k}, a\right) \\
& =\min _{1 \leq l \leq l_{h}^{k}-1} Q_{k, h+1}^{l}\left(s_{h+1}^{k}, a_{h+1}^{k}\right) \\
& \leq Q_{k, h+1}^{l_{h+1}^{k}-1}\left(s_{h+1}^{k}, a_{h+1}^{k}\right) \tag{C.5}
\end{align*}
$$

where the inequality holds due to the fact that $l_{h+1}^{k}-1 \leq l_{h}^{k}-1$.Substituting (C.5) in to (C.4) and taking a summation of (C.4) with all stage $h \in[H]$, we have

$$
\begin{equation*}
V_{1}^{*}\left(s_{1}^{k}\right)-V_{1}^{\pi_{k}}\left(s_{1}^{k}\right) \leq Q_{k, 1}^{l_{1}^{k}-1}\left(s_{1}^{k}, a_{1}^{k}\right)-Q_{1}^{\pi_{k}}\left(s_{1}^{k}, a_{1}^{k}\right) \leq \sum_{h=1}^{H} 4 \beta_{l_{h}^{k}-1} 2^{-l_{h}^{k}}+\sum_{h=1}^{H} \Delta_{k, h} \tag{C.6}
\end{equation*}
$$

Taking a summation of (C.6) over all episode $k_{i} \in K$, we have

$$
\begin{equation*}
\sum_{i=1}^{m} V_{1}^{*}\left(s_{1}^{k_{i}}\right)-V_{1}^{\pi_{k_{i}}}\left(s_{1}^{k_{i}}\right) \leq \underbrace{\sum_{i=1}^{m} \sum_{h=1}^{H} 4 \beta_{l_{h}^{k_{i}}-1} 2^{-l_{h}^{k_{i}}}}_{I_{1}}+\underbrace{\sum_{i=1}^{m} \sum_{h=1}^{H} \Delta_{k_{i}, h}}_{I_{2}} \tag{C.7}
\end{equation*}
$$

Since $4 \beta_{l-1} 2^{-l} \leq \epsilon /(2 H)$ holds for all level $l>l^{\prime}$, the term $I_{1}$ can be upper bounded by

$$
\begin{aligned}
I_{1} & =\sum_{i=1}^{m} \sum_{h=1}^{H} 4 \beta_{l_{h}^{k_{i}}-1} 2^{-l_{h}^{k_{i}}} \\
& \leq \sum_{i=1}^{m} \sum_{h=1}^{H}\left(\mathbb{1}\left\{l_{h}^{k_{i}} \leq l^{\prime}\right\} 4 \beta_{l_{h}^{k_{i}}-1} 2^{-l_{h}^{k_{i}}}+\frac{\epsilon}{2 H}\right) \\
& =\sum_{i=1}^{m} \sum_{h=1}^{H}\left(\sum_{l=1}^{l^{\prime}} \mathbb{1}\left\{l_{h}^{k_{i}}=l\right\} 4 \beta_{l-1} 2^{-l}+\frac{\epsilon}{2 H}\right) \\
& =\frac{m \epsilon}{2}+\sum_{h=1}^{H} \sum_{l=1}^{l^{\prime}} 4 \beta_{l-1} 2^{-l} \sum_{i=1}^{m} \mathbb{1}\left\{l_{h}^{k_{i}}=l\right\}
\end{aligned}
$$

According to Lemma C.1, the number of elements added into the set $\mathcal{C}_{h}^{l}$ is upper bounded by $\left|\mathcal{C}_{h}^{l}\right| \leq 17 d l h 4^{l}$ and it implies that $\sum_{i=1}^{m} \mathbb{1}\left\{l_{h}^{k_{i}}=l\right\} \leq \sum_{k=1}^{+\infty} \mathbb{1}\left\{l_{h}^{k}=l\right\} \leq 17 d l h 4^{l}$. Thus, we have

$$
\begin{align*}
I_{1} & \leq \frac{m \epsilon}{2}+\sum_{h=1}^{H} \sum_{l=1}^{l^{\prime}} 4 \beta_{l-1} 2^{-l} \sum_{i=1}^{m} \mathbb{1}\left\{l_{h}^{k_{i}}=l\right\} \\
& \leq \frac{m \epsilon}{2}+\sum_{h=1}^{H} \sum_{l=1}^{l^{\prime}} 4 \beta_{l-1} 2^{-l} \times 17 d l h 4^{l} \\
& \leq \frac{m \epsilon}{2}+136 \beta_{l^{\prime}-1} 2^{l^{\prime}} d l^{\prime} H^{2} \tag{C.8}
\end{align*}
$$

According to the definition of level $l^{\prime}$ and parameter $\beta_{l^{\prime}}$, there exist a large constant $C^{\prime \prime}$ such that $I_{1} \leq m \epsilon / 2+C^{\prime \prime} d^{3} H^{5} \log ^{4}(d H /(\delta \epsilon)) / \epsilon$.
For the term $I_{2}$, according to Lemma F.1, for any fixed number $n \in \mathbb{N}$ and $\epsilon=H / 2^{i}$, with probability at least $1-\delta /(2 i(i+1) n(n+1))$, we have

$$
\sum_{i=1}^{n} \sum_{h=1}^{H} \Delta_{k_{i}, h} \leq H \sqrt{2 H n \log \frac{2 i(i+1) n(n+1)}{\delta}}
$$

Therefore, taking a union bound, with probability at least $1-\delta /(2 i(i+1))$, for all $n \in \mathbb{N}$, we have

$$
\sum_{i=1}^{n} \sum_{h=1}^{H} \Delta_{k_{i}, h} \leq H \sqrt{2 H n \log \frac{2 i(i+1) n(n+1)}{\delta}}
$$

Thus, for the term $I_{2}$ and $\epsilon=H / 2^{i}$, with probability at least $1-\delta /(2 i(i+1))$, we have

$$
\begin{equation*}
I_{2}=\sum_{i=1}^{m} \sum_{h=1}^{H} \Delta_{k_{i}, h} \leq H \sqrt{2 H m \log \frac{2 i(i+1) m(m+1)}{\delta}} . \tag{C.9}
\end{equation*}
$$

Substituting (C.8) and (C.9) into (C.7), for $\epsilon=H / 2^{i}$, we have

$$
\begin{aligned}
m \epsilon & \geq \sum_{i=1}^{m} V_{1}^{*}\left(s_{1}^{k_{i}}\right)-V_{1}^{\pi_{k_{i}}}\left(s_{1}^{k_{i}}\right) \\
\geq & \frac{m \epsilon}{2}+C^{\prime \prime} d^{3} H^{5} \log ^{4}(d H /(\delta \epsilon)) / \epsilon \\
& \quad+H \sqrt{2 m \log \frac{2 i(i+1) m(m+1)}{\delta}}
\end{aligned}
$$

which implies $m \leq O\left(d^{3} H^{5} \log ^{4}(d H /(\delta \epsilon)) / \epsilon^{2}\right)$. Finally, taking an union bound with the event $\mathcal{E}$ and (C.9), with probability at least $1-\delta / 2-\sum_{i=1}^{\infty} \delta /(2 i(i+1))=1-\delta$, for all $\epsilon=H / 2^{i}(i \in \mathbb{N})$, we have

$$
\sum_{k=1}^{\infty} \mathbb{1}\left\{V_{1}^{*}\left(s_{1}^{k}\right)-V_{1}^{\pi_{k}}\left(s_{1}^{k}\right) \geq \epsilon\right\} \leq O\left(d^{3} H^{5} \log ^{4}(d H /(\delta \epsilon)) / \epsilon^{2}\right)
$$

Finally, we extend the result to general $\epsilon>0$. For any $H / 2^{i} \leq \epsilon \leq H / 2^{i-1}$, we have

$$
\begin{aligned}
\sum_{k=1}^{\infty} \mathbb{1}\left\{V_{1}^{*}\left(s_{1}^{k}\right)-V_{1}^{\pi_{k}}\left(s_{1}^{k}\right) \geq \epsilon\right\} & \leq \sum_{k=1}^{\infty} \mathbb{1}\left\{V_{1}^{*}\left(s_{1}^{k}\right)-V_{1}^{\pi_{k}}\left(s_{1}^{k}\right) \geq H / 2^{i}\right\} \\
& \leq O\left(d^{3} H^{5} \log ^{4}\left(d H /\left(\delta \epsilon^{\prime}\right)\right) /\left(H / 2^{i}\right)^{2}\right) \\
& =O\left(d^{3} H^{5} \log ^{4}(d H /(\delta \epsilon)) / \epsilon^{2}\right)
\end{aligned}
$$

Thus, we finish the proof of Theorem 5.1.

## D Proof of Lemma in Section B

## D. 1 Proof of Lemma B. 1

Lemma $\mathbf{D} .1$ (Lemma 11, [1]). For any vector sequence $\left\{\mathbf{x}_{k}\right\}_{k=1}^{K}$ in $\mathbb{R}^{d}$, We denote $\boldsymbol{\Sigma}_{0}=\lambda \mathbf{I}$ and $\boldsymbol{\Sigma}_{k}=\boldsymbol{\Sigma}_{0}+\sum_{i=1}^{k} \mathbf{x}_{i} \mathbf{x}_{i}^{\top}$. If $\lambda \geq \max \left(1, L^{2}\right)$ and $\left\|\mathbf{x}_{k}\right\|_{2} \leq L$ holds for all $k \in[K]$, then we have

$$
\sum_{k=1}^{K}\left\|\mathbf{x}_{k}\right\|_{\boldsymbol{\Sigma}_{k-1}^{-1}}^{2} \leq 2 d \log \frac{d \lambda+K L^{2}}{d \lambda}
$$

Proof of Lemma B.1. We focus on round $k$ and we suppose set $\mathcal{C}_{k}^{l}=\left\{k_{1}, . ., k_{m}\right\}$ at that time, where $1 \leq k_{1}<k_{2}<. .<k_{m}<k$. According to the update rule of set $\mathcal{C}^{l}$ in Algorithm 1 (line 9 to line 13), for each $2 \leq i \leq m$, we have $S_{k_{i}} \geq l$ and it implies that

$$
\begin{equation*}
\mathbf{x}_{k_{i}}^{\top}\left(\boldsymbol{\Sigma}_{k_{i}}^{l}\right)^{-1} \mathbf{x}_{k_{i}} \geq 4^{-l} \tag{D.1}
\end{equation*}
$$

where $\boldsymbol{\Sigma}_{k_{i}}^{l}=\lambda \mathbf{I}+\sum_{j=1}^{i-1} \mathbf{x}_{k_{j}} \mathbf{x}_{k_{j}}^{\top}$. Therefore, taking a summation for (D.1) over all $2 \leq i \leq m$, we have

$$
\begin{equation*}
\sum_{i=1}^{m} \mathbf{x}_{k_{i}}^{\top}\left(\boldsymbol{\Sigma}_{k_{i}}^{l}\right)^{-1} \mathbf{x}_{k_{i}} \geq \sum_{i=2}^{m} \mathbf{x}_{k_{i}}^{\top}\left(\boldsymbol{\Sigma}_{k_{i}}^{l}\right)^{-1} \mathbf{x}_{k_{i}} \geq(m-1) 4^{-l} \tag{D.2}
\end{equation*}
$$

where the first inequality holds due to $\mathbf{x}_{k_{1}}^{\top}\left(\boldsymbol{\Sigma}_{k_{1}}^{l}\right)^{-1} \mathbf{x}_{k_{1}} \geq 0$ and the second inequality holds due to (D.1). On the other hand, according to Lemma D.1, this summation is upper bounded by

$$
\begin{equation*}
\sum_{i=1}^{m} \mathbf{x}_{k_{i}}^{\top}\left(\boldsymbol{\Sigma}_{k_{i}}^{l}\right)^{-1} \mathbf{x}_{k_{i}} \leq 2 d \log \frac{d \lambda+m}{d \lambda} \leq 2 d \log (1+m / d) \tag{D.3}
\end{equation*}
$$

where the first inequality holds due to Lemma D. 1 with $\left\|\mathrm{x}_{k_{i}}\right\|_{2} \leq 1$ and the second inequality holds due to $\lambda \geq 1$. Combining (D.2) and (D.3), we have

$$
(m-1) 4^{-l} \leq 2 d \log (1+m / d)
$$

which implies that the size of set $\left|\mathcal{C}_{k}^{l}\right|$ is upper bounded by $\left|\mathcal{C}_{k}^{l}\right|=m \leq 17 d l 4^{l}$ for each $k \in \mathbb{N}$. Therefore, we finish the proof of Lemma B.1.

## D. 2 Proof of Lemma B. 2

Lemma D. 2 (Theorem 2, [1]). Let $\left\{\epsilon_{t}\right\}_{t=1}^{\infty}$ be a real-valued stochastic process with corresponding filtration $\left\{\mathcal{F}_{t}\right\}_{t=0}^{\infty}$ such that $\epsilon_{t}$ is $\mathcal{F}_{t}$-measure and $\epsilon_{t}$ is conditionally $R$-sub-Gaussian, i.e.

$$
\forall \lambda \in \mathbb{R}, \mathbb{E}\left[e^{\lambda \epsilon_{t}} \mid \mathcal{F}_{t-1}\right] \leq \exp \left(\frac{\lambda^{2} R^{2}}{2}\right)
$$

Let $\left\{\mathbf{x}_{t}\right\}_{t=1}^{\infty}$ be an $\mathbb{R}^{d}$-valued stochastic process where $\mathbf{x}_{t}$ is $\mathcal{F}_{t-1}$-measurable and we define $y_{t}=$ $\left\langle\mathbf{x}_{t}, \boldsymbol{\theta}^{*}\right\rangle+\epsilon_{t}$. With this notation, for any $t \geq 0$, we further define

$$
\boldsymbol{\Sigma}_{t}=\lambda I+\sum_{i=1}^{t} \mathbf{x}_{t} \mathbf{x}_{t}^{\top}, \mathbf{b}_{t}=\sum_{i=1}^{t} \mathbf{x}_{t} y_{t}, \mathbf{w}_{t}=\left(\boldsymbol{\Sigma}_{t}\right)^{-1} \mathbf{b}_{t}
$$

If we assume $\left\|\boldsymbol{\theta}^{*}\right\| \leq S$ and $\left\|\mathrm{x}_{t}\right\| \leq L$ holds for all $t \in \mathbb{N}$, then with probability at least $1-\delta$, for all $t \geq 0$, we have

$$
\left\|\boldsymbol{\theta}^{*}-\mathbf{w}_{t}\right\|_{\boldsymbol{\Sigma}_{t}} \leq R \sqrt{d \log \left(\frac{1+t L^{2} / \lambda}{\delta}\right)}+\sqrt{\lambda} S
$$

Proof of Lemma B.2. In this proof, we first focus on a fixed level $l \in \mathbb{N}$ and then turn back to all level $l$. For a fixed level $l \in[N]$, we denote $k_{0}=0$, and for $i \in \mathbb{N}$, we denote $k_{i}$ as the minimum index of the round where the action is added to the set $\mathcal{C}^{l}$ :

$$
\begin{equation*}
k_{i}=\min \left\{k: k>k_{i-1}, l_{k}=l\right\} . \tag{D.4}
\end{equation*}
$$

Under this notation, for all round $k\left(k_{i}<k \leq k_{i+1}\right)$, we have

$$
\begin{equation*}
\boldsymbol{\Sigma}_{k}^{l}=\lambda \mathbf{I}+\sum_{j=1}^{i} \mathbf{x}_{k_{j}} \mathbf{x}_{k_{j}}^{\top}, \mathbf{b}_{k}^{l}=\sum_{j=1}^{i} \mathbf{x}_{k_{j}} r_{k_{j}}, \mathbf{w}_{k}^{l}=\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1} \mathbf{b}_{k}^{l} \tag{D.5}
\end{equation*}
$$

Now, we consider the $\sigma$-algebra filtration $\mathcal{F}_{i}=\sigma\left(\mathbf{x}_{1}, . ., \mathbf{x}_{k_{i+1}}, r_{1}, . ., r_{k_{i+1}-1}\right)$ that contains all randomness before receiving the reward $r_{k_{i+1}}$ at round $k_{i+1}$. By the definition of $\mathcal{F}_{i-1}$, vector $\mathbf{x}_{k_{i}}$ is $\mathcal{F}_{i-1}$-measurable and the noise $\epsilon_{k_{i}}=r_{k_{i}}-\left\langle\mathbf{x}_{k_{i}}, \boldsymbol{\mu}^{*}\right\rangle$ is $\mathcal{F}_{i}$-measurable. Since we choose the level $l_{k}$ and add element $k$ to the corresponding set $\mathcal{C}^{l_{k}}$ before receiving the reward $r_{k}$ at round $k$, the noise $\epsilon_{k_{i}}$ is conditionally 1-Sub-Gaussian. According to Lemma D.2, with probability at least $1-\delta(l(l+1))$, for all $i \geq 0$, we have

$$
\begin{equation*}
\left\|\boldsymbol{\mu}^{*}-\mathbf{w}_{k_{i+1}}^{l}\right\|_{\boldsymbol{\Sigma}_{k_{i+1}}^{l}} \leq \sqrt{d \log \left(\frac{i+1}{\delta /(l(l+1))}\right)}+1 \tag{D.6}
\end{equation*}
$$

Combining (D.5) and (D.6), for all round $k\left(k_{i}<k \leq k_{i+1}\right)$, we have

$$
\begin{equation*}
\left\|\boldsymbol{\mu}^{*}-\mathbf{w}_{k}^{l}\right\|_{\boldsymbol{\Sigma}_{k}^{l}} \leq \sqrt{d \log \left(\frac{i+1}{\delta /(l(l+1))}\right)}+1 \tag{D.7}
\end{equation*}
$$

Furthermore, Lemma B. 1 suggests that the size of set $\left|\mathcal{C}^{l}\right|$ is upper bounded by $\left|\mathcal{C}^{l}\right| \leq 17 d l 4^{l}$, which implies that $k_{17 d l 4^{l}+1}=+\infty$. Thus, (D.7) implies that with probability at least $1-\delta /(l(l+1))$, for all round $k \in \mathbb{N}$, we have

$$
\begin{equation*}
\left\|\boldsymbol{\mu}^{*}-\mathbf{w}_{k}^{l}\right\|_{\boldsymbol{\Sigma}_{k}^{l}} \leq \sqrt{d \log \left(\frac{17 d l 4^{l}+1}{\delta /(l(l+1))}\right)}+1 \leq \beta_{l} \tag{D.8}
\end{equation*}
$$

Finally, taking a union bound for (D.8) over all level $l \in \mathbb{N}$, with probability at least $1-$ $\sum_{l=1}^{\infty}(\delta /(l(l+1)))=1-\delta$, for all level $l \in[N]$ and all round $k \in \mathbb{N}$, we have

$$
\left\|\mathbf{w}_{k}^{l}-\boldsymbol{\mu}^{*}\right\|_{\boldsymbol{\Sigma}_{k}^{l}} \leq \beta_{l}
$$

Thus, we finish the proof of Lemma B. 2

## E Proof of Lemma in Section C

## E. 1 Proof of Lemma C. 1

Proof of Lemma C.1. Similar to the proof of Lemma B.1, we focus on episode $k$ and we suppose set $\mathcal{C}_{k, h}^{l}=\left\{k_{1}, . ., k_{m}\right\}$ at that time, where $1 \leq k_{1}<k_{2}<. .<k_{m}<k$. For simplicity, we further define the auxiliary sets $\mathcal{B}_{k, h}^{l}$ as

$$
\mathcal{B}_{k, h}^{l}=\left\{i \mid 1 \leq i<k ; l_{h}^{i}=l ;\left(h=1 \text { or } l_{h}^{i}<l_{h-1}^{i}\right)\right\} .
$$

Notice that for each stage $h \geq 2$ and episode $i \in[k]$, there are two stopping rules for the while loop in Algorithm 2 (line 19 to line 20) and $\mathcal{B}_{k, h}^{l}$ consists of all episode $i \in[k]$ that stop with the first rule.

Furthermore, for all element $k_{i} \in \mathcal{C}_{k, h}^{l}$ with the second rule stopping rule, we have $l_{h-1}^{k}=l_{h}^{k}=l$ and it implies that $k_{i} \in \mathcal{C}_{k, h-1}^{l}$. Combining these two cases, we have $\mathcal{C}_{h}^{l} \subseteq \mathcal{B}_{k, h}^{l} \cup \mathcal{C}_{k, h-1}^{l}$ and it implies that

$$
\begin{equation*}
m=\left|\mathcal{C}_{k, h}^{l}\right| \leq \sum_{j=1}^{h}\left|\mathcal{B}_{k, j}^{l}\right| \tag{E.1}
\end{equation*}
$$

where the inequality holds due to $|\mathcal{C} \cup \mathcal{B}| \leq|\mathcal{C}|+|\mathcal{B}|$ and the fact that $\mathcal{C}_{k, 1}^{l}=\mathcal{B}_{k, 1}^{l}$.
Now, we only need to control the size of $\mathcal{B}_{k, h}^{l}$ for each episode $k \in \mathbb{N}$. For simplicity, we suppose set $\mathcal{B}_{k, h}^{l}=\left\{k_{1}, . ., k_{n}\right\}$, where $1 \leq k_{1} \leq k_{2} \leq \ldots \leq k_{n}<k$. According to the definition of level $l_{h}^{k}$ in Algorithm 2 (line 19 to line 20), for $2 \leq i \leq n$, we have

$$
\boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)^{\top}\left(\boldsymbol{\Sigma}_{k_{i}, h}^{l}\right)^{-1} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right) \geq 4^{-l}
$$

Since $\mathcal{B}_{k, h}^{l} \subseteq \mathcal{C}_{k, h}^{l}$ holds for all stage $h \in[H]$, all level $l \in \mathbb{N}$ and all episode $k \in \mathbb{N}$, we have $\boldsymbol{\Sigma}_{k_{i}, h}^{l} \succeq \lambda \mathbf{I}+\sum_{j=1}^{i-1} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right) \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)^{\top}$ and it implies that

$$
\begin{equation*}
\boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)^{\top}\left(\boldsymbol{\Gamma}_{k_{i}, h}^{l}\right)^{-1} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right) \geq \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)^{\top}\left(\boldsymbol{\Sigma}_{k_{i}, h}^{l}\right)^{-1} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right) \geq 4^{-l} \tag{E.2}
\end{equation*}
$$

where $\boldsymbol{\Gamma}_{k_{i}, h}^{l}=\lambda \mathbf{I}+\sum_{j=1}^{i-1} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right) \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)^{\top}$. Thus, taking a summation for (E.2) over all $2 \leq i \leq n$, we have

$$
\begin{equation*}
\sum_{i=1}^{n} \phi\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)^{\top}\left(\boldsymbol{\Gamma}_{k_{i}, h}^{l}\right)^{-1} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right) \geq \sum_{i=2}^{n} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)^{\top}\left(\boldsymbol{\Gamma}_{k_{i}, h}^{l}\right)^{-1} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right) \geq(n-1) 4^{-l} \tag{E.3}
\end{equation*}
$$

where the first inequality holds due to $\boldsymbol{\phi}\left(s_{h}^{k_{1}}, a_{h}^{k_{1}}\right)^{\top}\left(\boldsymbol{\Gamma}_{k_{1}, h}^{l}\right)^{-1} \boldsymbol{\phi}\left(s_{h}^{k_{1}}, a_{h}^{k_{1}}\right) \geq 0$ and the second inequality holds due to (E.2). On the other hand, according to Lemma D.1, this summation is upper bounded by

$$
\begin{equation*}
\sum_{i=1}^{n} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)^{\top}\left(\boldsymbol{\Gamma}_{k_{i}, h}^{l}\right)^{-1} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right) \leq 2 d \log \frac{d \lambda+n}{d \lambda} \leq 2 d \log (1+n / d) \tag{E.4}
\end{equation*}
$$

where the first inequality holds due to Lemma D. 1 with the fact that $\|\phi(s, a)\|_{2} \leq 1$ and the second inequality holds due to $\lambda \geq 1$. Combining (E.3) and (E.4), we have

$$
\begin{equation*}
(n-1) 4^{-l} \leq 2 d \log (1+n / d) \tag{E.5}
\end{equation*}
$$

which implies that $\left|\mathcal{B}_{k, h}^{l}\right|=n \leq 17 d l 4^{l}$. Finally, substituting (E.5) into (E.1), we have.

$$
\begin{equation*}
m=\left|\mathcal{C}_{k, h}^{l}\right| \leq \sum_{j=1}^{h}\left|\mathcal{B}_{k, j}^{l}\right| \leq 17 d h l 4^{l} \tag{E.6}
\end{equation*}
$$

Therefore, we finish the proof of Lemma C.1.

## E. 2 Proof of Lemma C. 2

Proof of Lemma C.2. In this proof, we only need to show that the norm of vector $\mathbf{w}_{k, h}^{l}$ is bounded for each fixed episode $k \in \mathbb{N}$ and fixed level $l \in \mathbb{N}$. For simplicity, let $\mathcal{C}_{k, h}^{l}=\left\{k_{1}, . ., k_{m}\right\}$ denote the index set $\mathcal{C}_{h}^{l}$ at the beginning of episode $k$, where $1 \leq k_{1}<k_{2}<. .<k_{m}<k$. According to the definition of weight vector $\mathbf{w}_{k, h}^{l}$ in Algorithm 2 (line 6 to line 8), we have

$$
\begin{aligned}
\boldsymbol{\Sigma}_{k, h}^{l} & =\lambda \mathbf{I}+\sum_{i=1}^{m} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right) \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)^{\top} \\
\mathbf{b}_{k, h}^{l} & =\sum_{i=1}^{m} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\left[r_{h}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+V_{k, h+1}^{l}\left(s_{h+1}^{k_{i}}\right)\right],
\end{aligned}
$$

$$
\mathbf{w}_{k, h}^{l}=\left(\boldsymbol{\Sigma}_{k, h}^{l}\right)^{-1} \mathbf{b}_{k, h}^{l}
$$

For simplicity, we omit the subscript $h$ and denote $r_{k_{i}}=r_{h}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+V_{k, h+1}^{l}\left(s_{h+1}^{k_{i}}\right)$. Then for the norm $\left\|\mathbf{w}_{k}^{l}\right\|_{2}$, we have the following inequality

$$
\begin{align*}
\left\|\mathbf{w}_{k}^{l}\right\|_{2}^{2} & =\left\|\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1} \sum_{i=1}^{m} \boldsymbol{\phi}\left(s^{k_{i}}, a^{k_{i}}\right) r_{k_{i}}\right\|_{2}^{2} \\
& \leq m \sum_{i=1}^{m}\left\|\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1} \boldsymbol{\phi}\left(s^{k_{i}}, a^{k_{i}}\right) r_{k_{i}}\right\|_{2}^{2} \\
& \leq 4 m H^{2} \sum_{i=1}^{m}\left\|\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1} \boldsymbol{\phi}\left(s^{k_{i}}, a^{k_{i}}\right)\right\|_{2}^{2} \\
& \leq \frac{4 m H^{2}}{\lambda} \sum_{i=1}^{m} \boldsymbol{\phi}\left(s^{k_{i}}, a^{k_{i}}\right)^{\top}\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1} \boldsymbol{\phi}\left(s^{k_{i}}, a^{k_{i}}\right) \\
& =\frac{4 m H^{2}}{\lambda} \operatorname{tr}\left(\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1} \sum_{i=1}^{m} \boldsymbol{\phi}\left(s^{k_{i}}, a^{k_{i}}\right)^{\top} \boldsymbol{\phi}\left(s^{k_{i}}, a^{k_{i}}\right)\right) \tag{E.7}
\end{align*}
$$

where the first inequality holds due to Cauchy-Schwarz inequality, the second inequality holds due to $\left|r_{k_{i}}\right| \leq 2 H$ and the last inequality holds due to $\Sigma_{k}^{l} \succeq \lambda I$. Now, we assume the eigen-decomposition of matrix $\sum_{i=1}^{m} \phi\left(s^{k_{i}}, a^{k_{i}}\right)^{\top} \phi\left(s^{k_{i}}, a^{k_{i}}\right)$ is $Q^{\top} \Lambda Q$ and we have

$$
\begin{align*}
\operatorname{tr}\left(\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1} \sum_{i=1}^{m} \phi\left(s^{k_{i}}, a^{k_{i}}\right)^{\top} \boldsymbol{\phi}\left(s^{k_{i}}, a^{k_{i}}\right)\right) & =\operatorname{tr}\left(\left(Q^{\top} \Lambda Q+\lambda I\right)^{-1} Q^{\top} \Lambda Q\right) \\
& =\operatorname{tr}\left((\Lambda+\lambda I)^{-1} \Lambda\right) \\
& =\sum_{i=1}^{d} \frac{\Lambda_{i}}{\Lambda_{i}+\lambda} \\
& \leq d \tag{E.8}
\end{align*}
$$

Substituting (E.8) into (E.7), we have

$$
\begin{aligned}
\left\|\mathbf{w}_{k}^{l}\right\|_{2}^{2} & \leq \frac{4 m H^{2}}{\lambda} \operatorname{tr}\left(\left(\boldsymbol{\Sigma}_{k}^{l}\right)^{-1} \sum_{i=1}^{m} \boldsymbol{\phi}\left(s^{k_{i}}, a^{k_{i}}\right)^{\top} \boldsymbol{\phi}\left(s^{k_{i}}, a^{k_{i}}\right)\right) \\
& \leq \frac{4 m H^{2} d}{\lambda} \\
& \leq \frac{68 d^{2} H^{3} l 4^{l}}{\lambda}
\end{aligned}
$$

where the first inequality holds due to (E.7), the second inequality holds due to (E.8) and the last inequality holds due to Lemma C.1. Thus, we finish the proof of Lemma C. 2

## E. 3 Proof of Lemma C. 3

In this section, we provide the proof of Lemma C.3. For each level $l \in \mathbb{N}$, we first denote the function class $\mathcal{V}_{l}$ as

$$
\begin{gather*}
\mathcal{V}_{l}=\left\{V \mid V(\cdot)=\max _{a} \min _{1 \leq i \leq l} \min \left(H, \mathbf{w}_{i}^{\top} \boldsymbol{\phi}(\cdot, a)+\beta_{l} \sqrt{\boldsymbol{\phi}(\cdot, a)^{\top} \boldsymbol{\Sigma}_{i}^{-1} \boldsymbol{\phi}(\cdot, a)}\right)\right. \\
\left.\left\|\mathbf{w}_{i}\right\|_{2} \leq 9 d 2^{l} \sqrt{H^{3} l}, \boldsymbol{\Sigma}_{i} \succeq I\right\} \tag{E.9}
\end{gather*}
$$

Therefore, for all episode $k \in K$ and stage $h \in[H]$, according to Lemma C.2, we have $\left\|\mathbf{w}_{k, h}^{l}\right\| \leq$ $9 d 2^{l} \sqrt{H^{3} l}$ and it implies that the estimated value function $V_{k, h}^{l} \in \mathcal{V}_{l}$. For any function $V \in \mathcal{V}_{l}$, we have the following concentration property.

Lemma E.1. (Lemma D.4, [13]) Let $\left\{x_{k}\right\}_{k=1}^{\infty}$ be a real-valued stochastic process on state space $\mathcal{S}$ with corresponding filtration $\left\{\mathcal{F}_{k}\right\}_{k=1}^{\infty}$. Let $\left\{\boldsymbol{\phi}_{k}\right\}_{k=1}^{\infty}$ be an $\mathbb{R}^{d}$-valued stochastic process where $\phi_{k} \in \mathcal{F}_{k-1}$ and $\left\|\phi_{k}\right\|_{2} \leq 1$. For any $k \geq 0$, we define $\boldsymbol{\Sigma}_{k}=I+\sum_{i=1}^{k} \phi_{i} \boldsymbol{\phi}_{i}^{\top}$. Then with probability at least $1-\delta$, for all $k \in \mathbb{N}$ and all function $V \in \mathcal{V}$ with $\max _{s}|V(x)| \leq H$, we have

$$
\left\|\sum_{i=1}^{k} \phi_{i}\left\{V\left(x_{i}\right)-\mathbb{E}\left[V\left(x_{i}\right) \mid \mathcal{F}_{i-1}\right]\right\}\right\|_{\boldsymbol{\Sigma}_{k}^{-1}}^{2} \leq 4 H^{2}\left[\frac{d}{2} \log (k+1)+\log \frac{\mathcal{N}_{\epsilon}}{\delta}\right]+8 k^{2} \epsilon^{2}
$$

where $\mathcal{N}_{\epsilon}$ is the $\epsilon$-covering number of the function class $\mathcal{V}$ with respect to the distance function $\operatorname{dist}\left(V_{1}, V_{2}\right)=\max _{s}\left|V_{1}(s)-V_{2}(s)\right|$.

Furthermore, for each function class $\mathcal{V}_{l}$, the covering number $\mathcal{N}_{\epsilon}$ of $\mathcal{V}_{l}$ can be upper bounded by following Lemma.
Lemma E.2. For each function class $\mathcal{V}_{l}$, we define the distance between two function $V_{1}$ and $V_{2}$ as $V_{1}, V_{2} \in \mathcal{V}_{l}$ as $\operatorname{dist}\left(V_{1}, V_{2}\right)=\max _{s}\left|V_{1}(s)-V_{2}(s)\right|$. With respect to this distance function, the $\epsilon$-covering number $\mathcal{N}_{\epsilon}$ of the function class $\mathcal{V}_{l}$ can be upper bounded by

$$
\log \mathcal{N}_{\epsilon} \leq d l \log \left(1+36 d 2^{l} \sqrt{H^{3} l} / \epsilon\right)+d^{2} l \log \left(1+8 \sqrt{d} \beta_{l}^{2} / \epsilon^{2}\right)
$$

Proof. See Appendix F.1.

Proof of Lemma C.3. Similar to the proof of Lemma B.2, we first focus on a fixed level $l \in \mathbb{N}$ and a fixed stage $h \in[H]$. Now, we denote $k_{0}=0$, and for $i \in \mathbb{N}$, we denote $k_{i}$ as the minimum index of the episode where the action is added to the $\operatorname{set} \mathcal{C}_{h}^{l}$ :

$$
\begin{equation*}
k_{i}=\min \left\{k: k>k_{i-1}, l_{h}^{k}=l\right\} . \tag{E.10}
\end{equation*}
$$

Now, we consider the $\sigma$-algebra filtration $\mathcal{F}_{i}=\sigma\left(s_{1}^{1}, . ., s_{H}^{1}, s_{1}^{2}, . ., s_{H}^{2}, . ., s_{1}^{k_{i+1}}, . ., s_{h}^{k_{i+1}}\right)$ that contain all randomness before receive the reward $r_{h}\left(s_{h}^{k_{i+1}}, a_{h}^{k_{i+1}}\right)$ and next state $s_{h+1}^{k_{i+1}}$ at episode $k_{i+1}$. By this definition, $\boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)$ is $\mathcal{F}_{i-1}$-measurable and the next state $s_{h+1}^{k_{i}}$ is $\mathcal{F}_{i}$-measurable. Since the randomness in this filtration only comes from the stochastic state transition process $s_{h+1} \sim$ $\mathbb{P}_{h}\left(\cdot \mid s_{h}, a_{h}\right)$ and we determine the level $l_{h}^{k}$ before receive the reward $r_{h}\left(s_{h}^{k}, a_{h}^{k}\right)$ and next state $s_{h+1}^{k}$ at episode $k$, for any fixed value function $V \in \mathcal{V}_{l}$, we have

$$
\begin{equation*}
\mathbb{E}\left[V\left(s_{h+1}^{k_{i}}\right) \mid \mathcal{F}_{i-1}\right]=\left[\mathbb{P}_{h} V\right]\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right) \tag{E.11}
\end{equation*}
$$

According to Lemma E. 1 with probability at least $1-\delta /(H 2 l(l+1))$, for all number $i \in \mathbb{N}$ and all function $V \in \mathcal{V}_{l}$, we have

$$
\begin{align*}
& \left\|\sum_{j=1}^{i} \phi\left(s_{h}^{k_{j}}, a_{h}^{k_{j}}\right)\left[V\left(s_{h+1}^{k_{j}}\right)-\left[\mathbb{P}_{h} V\right]\left(s_{h}^{k_{j}}, a_{h}^{k_{j}}\right)\right]\right\|_{\left(\boldsymbol{\Sigma}_{k_{i+1}^{l}, h}^{l}\right)^{-1}} \\
& =\| \sum_{j=1}^{i} \phi\left(s_{h}^{k_{j}}, a_{h}^{k_{j}}\right)\left[V\left(s_{h+1}^{k_{j}}\right)-\mathbb{E}\left[V\left(s_{h+1}^{k_{i+1}}\right) \mid \mathcal{F}_{i-1}\right] \|_{\left(\boldsymbol{\Sigma}_{k_{i+1}, h}^{l}\right)^{-1}}\right. \\
& \leq 4 H^{2}\left[\frac{d}{2} \log (i+1)+\log \frac{\mathcal{N}_{\epsilon}}{\delta /(H 2 l(l+1))}\right]+8 i^{2} \epsilon^{2} \\
& \leq 4 H^{2}\left[\frac{d}{2} \log (i+1)+d l \log \left(1+36 d 2^{l} \sqrt{H^{3} l} / \epsilon\right)+d^{2} l \log \left(1+8 \sqrt{d} \beta_{l}^{2} / \epsilon^{2}\right)\right. \\
& \quad+\log (2 H l(l+1) / \delta)]+8 i^{2} \epsilon^{2} \tag{E.12}
\end{align*}
$$

where the first inequality holds due to Lemma E. 1 and the second inequality holds due to Lemma E.2. Furthermore, Lemma C. 1 show that the size of set $\left|\mathcal{C}_{h}^{l}\right|$ is upper bounded by $\left|\mathcal{C}_{h}^{l}\right| \leq 17 d l H 4^{l}$. This reuslt implies that $k_{17 d l H 4^{l}+1}=\infty$ and we only need to consider episode $k_{i}$ for $i \leq 17 d l H 4^{l}+1$.

Now, we choose $\epsilon=1 /\left(17 l 4^{l}\right)$, then for all episode $k_{i}<k \leq k_{i+1}$ and function $V=V_{k, h+1}^{l} \in \mathcal{V}_{l}$, we have

$$
\begin{align*}
& \left\|\sum_{j=1}^{i} \boldsymbol{\phi}\left(s_{h}^{k_{j}}, a_{h}^{k_{j}}\right)\left[V_{k, h+1}^{l}\left(s_{h+1}^{k_{j}}\right)-\left[\mathbb{P}_{h} V_{k, h+1}^{l}\right]\left(s_{h}^{k_{j}}, a_{h}^{k_{j}}\right)\right]\right\|_{\left(\Sigma_{k, h}^{l}\right)^{-1}}^{2} \\
& \leq 4 H^{2}\left[\frac{d l}{2} \log (69 d l H)+2 d l^{2} \log \left(1+2448 d \sqrt{H^{3} l^{3}}\right)\right. \\
& \left.\quad+d^{2} l^{2} \log \left(1+36992 \sqrt{d} l^{2} \beta_{l}^{2}\right)+2 \log (4 l H / \delta)\right]+8 d^{2} H^{2} \tag{E.13}
\end{align*}
$$

where the first inequality holds due to (E.12) with the fact that $\boldsymbol{\Sigma}_{k, h}^{l}$ does not change for $k_{i}<k \leq$ $k_{i+1}$ and $i \leq 17 d l H 4^{l}+1$. Finally, taking an union bound for all level $l \in \mathbb{N}$ and all stage $h \in[H]$, with probability at $1-\delta / 2$, for all level $l \in \mathbb{N}$, all stage $h \in[H]$ and all episode $k \in \mathbb{N}$, we have

$$
\left\|\sum_{i \in \mathcal{C}_{k, h}^{l}} \phi\left(s_{h}^{i}, a_{h}^{i}\right)\left[V_{k, h+1}^{l}\left(s_{h+1}^{i}\right)-\left[\mathbb{P}_{h} V_{k, h+1}^{l}\right]\left(s_{h}^{i}, a_{h}^{i}\right)\right]\right\|_{\left(\boldsymbol{\Sigma}_{k, h}^{l}\right)^{-1}}^{2} \leq C d^{2} H^{2} l^{2} \log \left(d l H \beta_{l}^{2} / \delta\right)
$$

where $C$ is a large absolute constant. Thus, we finish the proof of Lemma C.3.

## E. 4 Proof of Lemma C. 4

Lemma E.3. [Lemma B.1, [13]] Under Assumption 3.1, for any fixed policy $\pi$, there exists a series of vectors $\left\{\mathbf{w}_{h}^{\pi}\right\}_{h=1}^{H}$, such that for all state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$ and all stage $h \in[H]$, we have

$$
Q_{h}^{\pi}(s, a)=\left(\mathbf{w}_{h}^{\pi}\right)^{\top} \boldsymbol{\phi}(s, a),\left\|\mathbf{w}_{h}^{\pi}\right\| \leq 2 H \sqrt{d}
$$

Proof of Lemma C.4. For simplicity, let $\mathcal{C}_{k, h}^{l}=\left\{k_{1}, . ., k_{m}\right\}$ denote the index set $\mathcal{C}_{h}^{l}$ at the beginning of episode $k$, where $1 \leq k_{1}<k_{2}<. .<k_{m}<k$. According to Lemma E.3, for each fixed policy $\pi$, there exists a vector $\mathbf{w}_{h}^{\pi}$ such that

$$
\begin{equation*}
\left(\mathbf{w}_{h}^{\pi}\right)^{\top} \boldsymbol{\phi}(s, a)=Q_{h}^{\pi}(s, a)=r_{h}(s, a)+\left[\mathbb{P}_{h} V_{h+1}^{\pi}\right](s, a) \tag{E.14}
\end{equation*}
$$

According to the definition of vector $\mathbf{w}_{k, h}$ in Algorithm 2 (line 6 to line 8), we have

$$
\begin{align*}
\boldsymbol{\Sigma}_{k, h}^{l} & =\lambda \mathbf{I}+\sum_{i=1}^{m} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right) \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)^{\top} \\
\mathbf{b}_{k, h}^{l} & =\sum_{i=1}^{m} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\left[r_{h}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+V_{k, h+1}^{l}\left(s_{h+1}^{k_{i}}\right)\right] \\
\mathbf{w}_{k, h}^{l} & =\left(\boldsymbol{\Sigma}_{k, h}^{l}\right)^{-1} \mathbf{b}_{k, h}^{l} \tag{E.15}
\end{align*}
$$

For simplicity, we omit the subscript $l$ and combining (E.14) and (E.15), we have

$$
\begin{aligned}
\mathbf{w}_{k, h}-\mathbf{w}_{h}^{\pi}= & \boldsymbol{\Sigma}_{k, h}^{-1} \sum_{i=1}^{m} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\left[r_{h}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+V_{k, h+1}\left(s_{h+1}^{k_{i}}\right)\right]-\mathbf{w}_{h}^{\pi} \\
= & \boldsymbol{\Sigma}_{k, h}^{-1}\left[-\lambda \mathbf{w}_{h}^{\pi}-\sum_{i=1}^{m} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right) \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)^{\top} \mathbf{w}_{h}^{\pi}\right. \\
& \left.\quad+\sum_{i=1}^{m} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\left[r_{h}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)+V_{k, h+1}\left(s_{h+1}^{k_{i}}\right)\right]\right] \\
= & \boldsymbol{\Sigma}_{k, h}^{-1}\left[-\lambda \mathbf{w}_{h}^{\pi}+\sum_{i=1}^{m} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\left(V_{k, h+1}\left(s_{h+1}^{k_{i}}\right)-\left[\mathbb{P}_{h} V_{h+1}^{\pi}\right]\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\right)\right]
\end{aligned}
$$

$$
\begin{aligned}
= & \underbrace{-\lambda \boldsymbol{\Sigma}_{k, h}^{-1} \mathbf{w}_{h}^{\pi}}_{I_{1}}+\underbrace{\boldsymbol{\Sigma}_{k, h}^{-1} \sum_{i=1}^{m} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\left(V_{k, h+1}\left(s_{h+1}^{k_{i}}\right)-\left[\mathbb{P}_{h} V_{k, h+1}\right]\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\right)}_{I_{2}} \\
& +\underbrace{\boldsymbol{\Sigma}_{k, h}^{-1} \sum_{i=1}^{m} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\left[\mathbb{P}_{h}\left(V_{k, h+1}-V_{h+1}^{\pi}\right)\right]\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)}_{I_{3}},
\end{aligned}
$$

where the third equality holds due to (E.14). For the term $I_{1}$ and any state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have

$$
\begin{align*}
\left|\left\langle I_{1}, \boldsymbol{\phi}(s, a)\right\rangle\right| & =\left|\lambda \boldsymbol{\phi}(s, a)^{\top} \boldsymbol{\Sigma}_{k, h}^{-1} \mathbf{w}_{h}^{\pi}\right| \\
& \leq \lambda\left\|\boldsymbol{\phi}(s, a)^{\top} \boldsymbol{\Sigma}_{k, h}^{-1}\right\|_{2}\left\|\mathbf{w}_{h}^{\pi}\right\|_{2} \\
& \leq \sqrt{\lambda}\left\|\mathbf{w}_{h}^{\pi}\right\|_{2} \sqrt{\boldsymbol{\phi}(s, a)^{\top} \boldsymbol{\Sigma}_{k, h}^{-1} \boldsymbol{\phi}(s, a)} \\
& \leq 2 H \sqrt{d \lambda} \sqrt{\boldsymbol{\phi}(s, a)^{\top} \boldsymbol{\Sigma}_{k, h}^{-1} \boldsymbol{\phi}(s, a)}, \tag{E.16}
\end{align*}
$$

where the first inequality holds due to Cauchy-Schwarz inequality, the second inequality holds due to $\boldsymbol{\Sigma}_{k, h} \succeq \lambda I$ and the third inequality holds due to Lemma E.3. For the term $I_{2}$ and any state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, according to Lemma C.3, we have

$$
\begin{align*}
\left|\left\langle I_{2}, \boldsymbol{\phi}(s, a)\right\rangle\right| \leq & \sqrt{\boldsymbol{\phi}(s, a)^{\top} \boldsymbol{\Sigma}_{k, h}^{-1} \boldsymbol{\phi}(s, a)} \\
& \cdot\left\|\sum_{i=1}^{m} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\left[V_{k, h+1}\left(s_{h+1}^{k_{i}}\right)-\left[\mathbb{P}_{h} V_{k, h+1}\right]\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\right]\right\|_{\boldsymbol{\Sigma}_{k, h}^{-1}} \\
\leq & C d H l \sqrt{\log \left(d l H \beta_{l}^{2} / \delta\right)} \sqrt{\boldsymbol{\phi}(s, a)^{\top} \boldsymbol{\Sigma}_{k, h}^{-1} \boldsymbol{\phi}(s, a)} \tag{E.17}
\end{align*}
$$

where the first inequality holds due to Cauchy-Schwarz inequality and the second inequality holds due to Lemma C.3. For the term $I_{3}$ and any state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have

$$
\begin{aligned}
\left\langle\boldsymbol{\phi}(s, a), I_{3}\right\rangle= & \left\langle\boldsymbol{\phi}(s, a), \boldsymbol{\Sigma}_{k, h}^{-1} \sum_{i=1}^{m} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\left[\mathbb{P}_{h}\left(V_{k, h+1}-V_{h+1}^{\pi}\right)\right]\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)\right\rangle \\
= & \underbrace{\left\langle\boldsymbol{\phi}(s, a), \boldsymbol{\Sigma}_{k, h}^{-1} \sum_{i=1}^{m} \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right) \boldsymbol{\phi}\left(s_{h}^{k_{i}}, a_{h}^{k_{i}}\right)^{\top} \int\left(V_{k, h+1}-V_{h+1}^{\pi}\right)\left(s^{\prime}\right) d \boldsymbol{\mu}_{h}\left(s^{\prime}\right)\right\rangle}_{J_{1}} \begin{aligned}
= & \underbrace{\left\langle\boldsymbol{\phi}(s, a), \int\left(V_{k, h+1}-V_{h+1}^{\pi}\right)\left(s^{\prime}\right) d \boldsymbol{\mu}_{h}\left(s^{\prime}\right)\right\rangle}_{J_{2}} \\
& -\lambda\left\langle\boldsymbol{\phi}(s, a), \boldsymbol{\Sigma}_{k, h}^{-1} \int\left(V_{k, h+1}-V_{h+1}^{\pi}\right)\left(s^{\prime}\right) d \boldsymbol{\mu}_{h}\left(s^{\prime}\right)\right\rangle
\end{aligned}
\end{aligned}
$$

For term $J_{1}$, we have

$$
\begin{align*}
J_{1} & =\left\langle\boldsymbol{\phi}(s, a), \int\left(V_{k, h+1}-V_{h+1}^{\pi}\right)\left(s^{\prime}\right) d \boldsymbol{\mu}_{h}\left(s^{\prime}\right)\right\rangle \\
& =\int\left\langle\boldsymbol{\phi}(s, a),\left(V_{k, h+1}-V_{h+1}^{\pi}\right)\left(s^{\prime}\right)\right\rangle d \boldsymbol{\mu}_{h}\left(s^{\prime}\right) \\
& \left.=\int \mathbb{P}_{h}\left(s^{\prime} \mid s, a\right)\left(V_{k, h+1}-V_{h+1}^{\pi}\right)\left(s^{\prime}\right)\right\rangle d s^{\prime} \\
& =\left[\mathbb{P}_{h}\left(V_{k, h+1}-V_{h+1}^{\pi}\right)\right](s, a) \tag{E.18}
\end{align*}
$$

For term $J_{2}$, we have

$$
\left|J_{2}\right|=\lambda\left|\left\langle\phi(s, a), \boldsymbol{\Sigma}_{k, h}^{-1} \int\left(V_{k, h+1}-V_{h+1}^{\pi}\right)\left(s^{\prime}\right) d \boldsymbol{\mu}_{h}\left(s^{\prime}\right)\right\rangle\right|
$$

$$
\begin{align*}
& \leq \lambda\left\|\boldsymbol{\phi}(s, a)^{\top} \boldsymbol{\Sigma}_{k, h}^{-1}\right\|_{2}\left\|\int\left(V_{k, h+1}-V_{h+1}^{\pi}\right)\left(s^{\prime}\right) d \boldsymbol{\mu}_{h}\left(s^{\prime}\right)\right\|_{2} \\
& \leq \sqrt{d} \lambda\left\|\boldsymbol{\phi}(s, a)^{\top} \boldsymbol{\Sigma}_{k, h}^{-1}\right\|_{2} \max _{s^{\prime}}\left|\left(V_{k, h+1}-V_{h+1}^{\pi}\right)\left(s^{\prime}\right)\right| \\
& \leq 2 H \sqrt{d} \lambda\left\|\boldsymbol{\phi}(s, a)^{\top} \boldsymbol{\Sigma}_{k, h}^{-1}\right\|_{2} \\
& \leq 2 H \sqrt{d \lambda} \sqrt{\boldsymbol{\phi}(s, a)^{\top} \boldsymbol{\Sigma}_{k, h}^{-1} \boldsymbol{\phi}(s, a)} \tag{E.19}
\end{align*}
$$

where the first inequality holds due to Cauchy-Schwarz inequality, the second inequality holds due to Assumption 3.1, the third inequality holds because of $\left|\left(V_{k, h+1}-V_{h+1}^{\pi}\right)\left(s^{\prime}\right)\right| \leq 2 H$ and the last inequality holds due to $\boldsymbol{\Sigma}_{k, h} \succeq \lambda I$. Combining (E.16),(E.17),(E.18),(E.19) with the fact that $\lambda=1$, we have

$$
\begin{aligned}
& \left|\left\langle\boldsymbol{\phi}(s, a), \mathbf{w}_{k, h}\right\rangle-Q_{h}^{\pi}(s, a)-\left[\mathbb{P}_{h}\left(V_{k, h+1}-V_{h+1}^{\pi}\right)\right](s, a)\right| \\
& =\left|J_{2}+\left\langle I_{1}, \boldsymbol{\phi}(s, a)\right\rangle+\left\langle I_{2}, \boldsymbol{\phi}(s, a)\right\rangle\right| \\
& \leq\left(C d H l \sqrt{\log \left(d l H \beta_{l}^{2} / \delta\right)}+4 H \sqrt{d}\right) \sqrt{\boldsymbol{\phi}(s, a)^{\top}\left(\boldsymbol{\Sigma}_{k, h}\right)^{-1} \boldsymbol{\phi}(s, a)}
\end{aligned}
$$

Notice that there exists a large constant $C^{\prime}$ such that for all level $l \in \mathbb{N}$ with parameter $\beta_{l}=$ $C^{\prime} d H l \sqrt{\log (d l H / \delta)}$, the following inequality hods:

$$
\begin{equation*}
C d H l \sqrt{\log \left(d l H \beta_{l}^{2} / \delta\right)}+4 H \sqrt{d} \leq C^{\prime} d H l \sqrt{\log (d l H / \delta)} \tag{E.20}
\end{equation*}
$$

When (E.20) holds, we further have

$$
\begin{aligned}
& \left|\left\langle\boldsymbol{\phi}(s, a), \mathbf{w}_{k, t}\right\rangle-Q_{h}^{\pi}(s, a)-\left[\mathbb{P}_{h}\left(V_{k, h+1}-V_{h+1}^{\pi}\right)\right](s, a)\right| \\
& \leq\left(C d H l \sqrt{\log \left(d l H \beta_{l}^{2} / \delta\right)}+3 H \sqrt{d}\right) \sqrt{\boldsymbol{\phi}(s, a)^{\top}\left(\boldsymbol{\Sigma}_{k, h}\right)^{-1} \boldsymbol{\phi}(s, a)} \\
& \leq C^{\prime} d H l \sqrt{\log (d l H / \delta)} \sqrt{\boldsymbol{\phi}(s, a)^{\top}\left(\boldsymbol{\Sigma}_{k, h}\right)^{-1} \boldsymbol{\phi}(s, a)} \\
& =\beta_{l} \sqrt{\boldsymbol{\phi}(s, a)^{\top}\left(\boldsymbol{\Sigma}_{k, h}\right)^{-1} \boldsymbol{\phi}(s, a)} .
\end{aligned}
$$

Thus, we finish the proof of Lemma C.4.

## E. 5 Proof of Lemma C. 5

Proof of Lemma C.5. Now, we use induction to prove this lemma. First, we prove the base case. For all state $s \in \mathcal{S}$ and level $l \in \mathbb{N}$, we have $V_{k, H+1}^{l}(s)=0=V_{H+1}^{*}(s)$. Second, if $V_{k, h+1}^{l}(s) \geq V_{h+1}^{*}(s)$ holds for all state $s \in \mathcal{S}$ and level $l \in \mathbb{N}$ at stage $h+1$, then for any state $s \in \mathcal{S}$ and level $l \in \mathbb{N}$ at stage $h$, we have

$$
\left(\mathbf{w}_{k, h}^{l}\right)^{\top} \boldsymbol{\phi}(s, a)+\beta_{l} \sqrt{\boldsymbol{\phi}(s, a)^{\top}\left(\Sigma_{k, h}^{l}\right)^{-1} \phi(s, a)}-Q_{h}^{*}(s, a) \geq\left[\mathbb{P}_{h}\left(V_{k, h+1}^{l}-V_{h+1}^{*}\right)\right](s, a) \geq 0
$$

where the first inequality holds due to Lemma C. 4 and the second inequality holds due to the induction assumption. Furthermore, the optimal value function is upper bounded by $Q_{h}^{*}(s, a) \leq H$ and it implies that

$$
\begin{equation*}
Q_{h}^{*}(s, a) \leq \min \left(\left(\mathbf{w}_{k, h}^{l}\right)^{\top} \boldsymbol{\phi}(s, a)+\beta_{l} \sqrt{\boldsymbol{\phi}(s, a)^{\top}\left(\Sigma_{k, h}^{l}\right)^{-1} \boldsymbol{\phi}(s, a)}, H\right)=Q_{k, h}^{l}(s, a) \tag{E.21}
\end{equation*}
$$

Thus, for each level $l \in \mathbb{N}$ and state $s \in \mathcal{S}$, we have

$$
\begin{aligned}
V_{k, h}^{l}(s) & =\max _{a} \min _{1 \leq i \leq l} Q_{k, h}^{i}(s, a) \\
& \geq \max _{a} \min _{1 \leq i \leq l} Q_{h}^{*}(s, a) \\
& =\max _{a} Q_{h}^{*}(s, a) \\
& =V_{h}^{*}(s)
\end{aligned}
$$

where the inequality holds due to (E.21). Finally, by induction, we finish the proof of Lemma C.5.

## F Auxiliary Lemmas

Lemma F. 1 (Azuma-Hoeffding inequality, [4]). Let $\left\{x_{i}\right\}_{i=1}^{n}$ be a martingale difference sequence with respect to a filtration $\left\{\mathcal{G}_{i}\right\}$ satisfying $\left|x_{i}\right| \leq M$ for some constant $M, x_{i}$ is $\mathcal{G}_{i+1}$-measurable, $\mathbb{E}\left[x_{i} \mid \mathcal{G}_{i}\right]=0$. Then for any $0<\delta<1$, with probability at least $1-\delta$, we have

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

## F. 1 Proof of Lemma E. 2

We need the following Lemma:
Lemma $\mathbf{F} .2$ (Lemma D.5, [13]). For an Euclidean ball with radius $R$ in $\mathbb{R}^{d}$, the $\epsilon$-covering number of this ball is upper bounded by $(1+2 R / \epsilon)^{d}$.

Proof of Lemma E.2. For any two function $V_{1}, V_{2} \in \mathcal{V}_{l}$, according to the definition of function class $\mathcal{V}_{l}$, we have

$$
\begin{aligned}
& V_{1}(\cdot)=\max _{a} \min _{1 \leq i \leq l} \min \left(H, \mathbf{w}_{1, i}^{\top} \boldsymbol{\phi}(\cdot, a)+\beta_{l} \sqrt{\boldsymbol{\phi}(\cdot, a)^{\top} \boldsymbol{\Gamma}_{1, i} \boldsymbol{\phi}(\cdot, a)}\right), \\
& V_{2}(\cdot)=\max _{a} \min _{1 \leq i \leq l} \min \left(H, \mathbf{w}_{2, i}^{\top} \boldsymbol{\phi}(\cdot, a)+\beta_{l} \sqrt{\boldsymbol{\phi}(\cdot, a)^{\top} \boldsymbol{\Gamma}_{2, i} \boldsymbol{\phi}(\cdot, a)}\right),
\end{aligned}
$$

where $\left\|\mathbf{w}_{1, i}\right\|_{2},\left\|\mathbf{w}_{2, i}\right\|_{2} \leq 9 d 2^{l} \sqrt{H^{3} l}$ and $\boldsymbol{\Gamma}_{1, i}, \boldsymbol{\Gamma}_{2, i} \preceq \mathbf{I}$. Since all of the functions $\max _{a}, \min _{1 \leq i \leq l}$ and $\min (H, \cdot)$ are contraction functions, we have

$$
\begin{align*}
\operatorname{dist}\left(V_{1}, V_{2}\right)= & \max _{s \in \mathcal{S}}\left|V_{1}(s)-V_{2}(s)\right| \\
\leq & \max _{1 \leq i \leq l, s \in \mathcal{S}, a \in \mathcal{A}} \mid \mathbf{w}_{1, i}^{\top} \boldsymbol{\phi}(s, a)+\beta_{l} \sqrt{\boldsymbol{\phi}(s, a)^{\top} \boldsymbol{\Gamma}_{1, i} \boldsymbol{\phi}(s, a)} \\
& \quad-\mathbf{w}_{2, i}^{\top} \boldsymbol{\phi}(s, a)-\beta_{l} \sqrt{\boldsymbol{\phi}(s, a)^{\top} \boldsymbol{\Gamma}_{2, i} \boldsymbol{\phi}(s, a)} \mid \\
\leq & \beta_{l} \max _{1 \leq i \leq l, s \in \mathcal{S}, a \in \mathcal{A}}\left|\sqrt{\boldsymbol{\phi}(s, a)^{\top} \boldsymbol{\Gamma}_{1, i} \boldsymbol{\phi}(s, a)}-\sqrt{\boldsymbol{\phi}(s, a)^{\top} \boldsymbol{\Gamma}_{2, i} \boldsymbol{\phi}(s, a)}\right| \\
& \quad+\max _{1 \leq i \leq l, s \in \mathcal{S}, a \in \mathcal{A}}\left|\left(\mathbf{w}_{1, i}-\mathbf{w}_{2, i}\right)^{\top} \boldsymbol{\phi}(s, a)\right| \\
\leq & \beta_{l} \max _{1 \leq i \leq l, s \in \mathcal{S}, a \in \mathcal{A}}\left|\sqrt{\boldsymbol{\phi}(s, a)^{\top}\left(\boldsymbol{\Gamma}_{1, i}-\boldsymbol{\Gamma}_{2, i}\right) \boldsymbol{\phi}(s, a)}\right| \\
& +\max _{1 \leq i \leq l, s \in \mathcal{S}, a \in \mathcal{A}}\left|\left(\mathbf{w}_{1, i}-\mathbf{w}_{2, i}\right)^{\top} \boldsymbol{\phi}(s, a)\right| \\
\leq & \beta_{l} \max _{1 \leq i \leq l} \sqrt{\left\|\boldsymbol{\Gamma}_{1, i}-\boldsymbol{\Gamma}_{2, i}\right\|_{F}}+\max _{1 \leq i \leq l}\left\|\mathbf{w}_{1, i}-\mathbf{w}_{2, i}\right\|_{2}, \tag{F.1}
\end{align*}
$$

where the first inequality holds due to the contraction property, the second inequality holds due to the fact that $\max _{x}|f(x)+g(x)| \leq \max _{x}|f(x)|+\max _{x}|g(x)|$, the third inequality holds due to $|\sqrt{x}-\sqrt{y}| \geq|\sqrt{x}-\sqrt{y}|$ and the last inequality holds due to the fact that $\|\phi(s, a)\|_{2} \leq 1$. Now, we denote $\mathcal{C}_{\mathbf{w}}$ as a $\epsilon / 2$-cover of the set $\left\{\mathbf{w} \in \mathbb{R}^{d} \mid\|\mathbf{w}\|_{2} \leq 9 d 2^{l} \sqrt{H^{3} l}\right\}$ and $\mathcal{C}_{\boldsymbol{\Gamma}}$ as a $\epsilon^{2} /\left(4 \beta_{l}^{2}\right)$-cover of the set $\left\{\boldsymbol{\Gamma} \in \mathbb{R}^{d \times d} \mid\|\boldsymbol{\Gamma}\|_{F} \leq \sqrt{d}\right\}$ with respect to the Frobenius norms. Thus, according to Lemma F.2, we have following property:

$$
\begin{equation*}
\left|\mathcal{C}_{\mathbf{w}}\right| \leq\left(1+36 d 2^{l} \sqrt{H^{3} l} / \epsilon\right)^{d},\left|\mathcal{C}_{\boldsymbol{\Gamma}}\right| \leq\left(1+8 \sqrt{d} \beta_{l}^{2} / \epsilon^{2}\right)^{d^{2}} \tag{F.2}
\end{equation*}
$$

By the definition of covering number, for any function $V_{1} \in \mathcal{V}_{l}$ with parameters $\mathbf{w}_{1, i}, \boldsymbol{\Gamma}_{1, i}(1 \leq$ $i \leq l)$, there exists other parameters $\mathbf{w}_{2, i}, \boldsymbol{\Gamma}_{2, i}(1 \leq i \leq l)$ such that $\mathbf{w}_{2, i} \in \mathcal{C}_{\mathbf{w}}, \boldsymbol{\Gamma}_{2, i} \in \mathcal{C}_{\boldsymbol{\Gamma}}$ and $\left\|\mathbf{w}_{2, i}-\mathbf{w}_{1, i}\right\|_{2} \leq \epsilon / 2,\left\|\boldsymbol{\Gamma}_{2, i}-\boldsymbol{\Gamma}_{1, i}\right\|_{F} \leq \epsilon^{2} /\left(4 \beta_{l}^{2}\right)$. Thus, we have

$$
\operatorname{dist}\left(V_{1}, V_{2}\right) \leq \beta_{l} \max _{1 \leq i \leq l} \sqrt{\left\|\boldsymbol{\Gamma}_{1, i}-\boldsymbol{\Gamma}_{2, i}\right\|_{F}}+\max _{1 \leq i \leq l}\left\|\mathbf{w}_{1, i}-\mathbf{w}_{2, i}\right\|_{2} \leq \epsilon,
$$

where the inequality holds due to (F.1). Therefore, the $\epsilon$-covering number of function class $\mathcal{V}_{l}$ is bounded by $\mathcal{N}_{\epsilon} \leq\left|\mathcal{C}_{\mathbf{w}}\right|^{l} \cdot\left|\mathcal{C}_{\boldsymbol{\Gamma}}\right|^{l}$ and it implies

$$
\log \mathcal{N}_{\epsilon} \leq d l \log \left(1+36 d 2^{l} \sqrt{H^{3} l} / \epsilon\right)+d^{2} l \log \left(1+8 \sqrt{d} \beta_{l}^{2} / \epsilon^{2}\right)
$$

where the first inequality holds due to (F.2). Thus, we finish the proof of Lemma E.2.


[^0]:    ${ }^{1}$ Suppose the suboptimality gaps satisfy $\Delta_{t}=\mathbb{1}\left\{t=i^{2}, i=1, \ldots\right\}$, then the regret in the first $T$ rounds is upper bounded by $O(\sqrt{T})$, and the constant 1-gap will appear infinitely often

[^1]:    ${ }^{2}$ Soare et al. [19] denoted by $\Delta_{\min }$ the gap of the rewards between the best arm and the second-best arm. In this setting, the sample complexity to find the best arm is identical to the sample complexity to find an $\epsilon<\Delta_{\text {min }}$ sub-optimal arm.

