
A Nearly Optimal and Low-Switching Algorithm for Reinforcement Learning with General Function Approximation

Heyang Zhao

Department of Computer Science
University of California, Los Angeles
Los Angeles, CA 90095
hyzhao@cs.ucla.edu

Jiafan He

Department of Computer Science
University of California, Los Angeles
Los Angeles, CA 90095
jiafanhe19@ucla.edu

Quanquan Gu

Department of Computer Science
University of California, Los Angeles
Los Angeles, CA 90095
qgu@cs.ucla.edu

Abstract

The exploration-exploitation dilemma has been a central challenge in reinforcement learning (RL) with complex model classes. In this paper, we propose a new algorithm, Monotonic Q-Learning with Upper Confidence Bound (MQL-UCB) for RL with general function approximation. Our key algorithmic design includes (1) a general deterministic policy-switching strategy that achieves low switching cost, (2) a monotonic value function structure with carefully controlled function class complexity, and (3) a variance-weighted regression scheme that exploits historical trajectories with high data efficiency. MQL-UCB achieves minimax optimal regret of $\tilde{O}(d\sqrt{HK})$ when K is sufficiently large and near-optimal policy switching cost of $\tilde{O}(dH)$, with d being the eluder dimension of the function class, H being the planning horizon, and K being the number of episodes. Our work sheds light on designing provably sample-efficient and deployment-efficient Q-learning with nonlinear function approximation.

1 Introduction

In reinforcement learning (RL), a learner interacts with an unknown environment and aims to maximize the cumulative reward. As one of the most mainstream paradigms for sequential decision-making, RL has extensive applications in many real-world problems (Kober et al., 2013; Mnih et al., 2015; Lillicrap et al., 2015; Zoph and Le, 2016; Zheng et al., 2018). Theoretically, the RL problem is often formulated as a Markov Decision Process (MDP) (Puterman, 2014). Achieving the optimal regret bound for various MDP settings has been a long-standing fundamental problem in RL research. In tabular MDPs where the state space \mathcal{S} and the action space \mathcal{A} are finite, the optimal regret bound has been well-established, ranging from the episodic setting (Azar et al., 2017; Zanette and Brunskill, 2019; Simchowitz and Jamieson, 2019; Zhang et al., 2020), average-reward setting (Zhang and Ji, 2019) to the discounted setting (He et al., 2021b). Nevertheless, these regret guarantees are intolerably large in many real-world applications, where the state space \mathcal{S} and the action space \mathcal{A} are often large or even infinite.

As is commonly applied in applications, *function approximation* has been widely studied by theorists to demonstrate the generalization across large state-action spaces, proving the performance guarantees of various RL algorithms for specific function classes. There are recent works on RL with linear function approximation under different assumptions such as linear MDPs (Yang and Wang, 2019; Jin et al., 2020), linear mixture MDPs (Modi et al., 2020; Ayoub et al., 2020; Zhou et al., 2021a). Among them, Zhou et al. (2021a) achieved nearly optimal regret bounds for linear mixture MDPs through a model-based approach adopting variance-weighted linear regression. Later, Hu et al. (2022) proposed LSVI-UCB+ algorithm, making an attempt to improve the regret for linear MDP through an over-optimistic value function approach. However, their analysis was later discovered to suffer from a technical issue (Agarwal et al., 2022; He et al., 2022). To fix this issue, Agarwal et al. (2022) introduced similar over-optimistic value functions to construct a monotonic variance estimator and a non-Markov policy, achieving the first statistically optimal regret for linear MDPs. They also proposed a novel algorithm dubbed VOQL for RL with general function approximation. Concurrently, He et al. (2022) proposed LSVI-UCB++, which takes a different approach and employs a rare-switching technique to obtain the optimal regret. In a parallel line of research, there has been a growing body of literature proposing more general frameworks to unify sample efficient RL algorithms, e.g., Bellman rank (Jiang et al., 2017), Witness rank (Sun et al., 2019), eluder dimension (Russo and Van Roy, 2013), Bellman eluder dimension (Jin et al., 2021), Bilinear Classes (Du et al., 2021), Decision-Estimation Coefficient (Foster et al., 2021), Admissible Bellman Characterization (Chen et al., 2022) and Decoupling Coefficient (Agarwal and Zhang, 2022a,b). However, when applying these frameworks to linear MDPs, none of them can achieve the minimax optimal regret. To our knowledge, Agarwal et al. (2022) is the only algorithmic framework achieving optimal regret for RL beyond linear function approximation. Since VOQL requires a non-Markov planning procedure, where the resulting policy does not act greedily with respect to a single optimistic value function, it is natural to ask

Can we develop a RL algorithm with Markov policy¹ to solve MDPs with general function approximation and achieve the optimal regret?

While the sample efficiency of RL algorithms for MDPs with nonlinear function classes has been comprehensively researched, *deployment efficiency* (Matsushima et al., 2020) is also a major concern in many real-world application scenarios. For example, in recommendation systems (Afsar et al., 2022), it may take several weeks to deploy a new recommendation policy. On the other hand, the system is capable of collecting an enormous amount of data every minute implementing a fixed policy. As a result, it is computationally inefficient to change the executed policy after each data point is collected as is demanded by most of the online RL algorithms in theoretical studies. To resolve this issue, Bai et al. (2019) first introduced the concept of switching cost, defined as the number of policy updates. Following this concept, a series of RL algorithms have been proposed on the theoretical side with low switching cost guarantees (e.g., Zhang et al., 2020; Wang et al., 2021; Qiao et al., 2022; Kong et al., 2021; Velegkas et al., 2022; Li et al., 2023). Xiong et al. (2023) considered low-switching RL with general function approximation, which achieves $\tilde{O}(dH)$ switching cost. However, their algorithm has an intractable planning phase and is not statistically optimal. Kong et al. (2021), Velegkas et al. (2022), and Li et al. (2023) considered RL with general function approximation, all of which achieved the switching cost of $O(d^2 H \text{polylog}(K))$, where d is the eluder dimension of the underlying function class, H is the planning horizon, and K is the number of episodes. In contrast, Gao et al. (2021) proved an $\Omega(dH/\log d)$ lower bound of switching cost for any deterministic algorithms in learning linear MDPs. Therefore, the following question remains open:

Can we design an efficient algorithm with $\tilde{O}(dH)$ switching cost for MDPs with bounded eluder dimension?

In this paper, we answer the above two questions simultaneously by proposing a novel algorithm *Monotonic Q-Learning with UCB (MQL-UCB)* with all the aforementioned appealing properties. At the core of our algorithmic design are the following innovative techniques:

¹A Markov policy means that the action chosen by the policy only depends on the current state instead of the prefix trajectory. It is more aligned with the empirical RL approaches, since the estimated value function is not well-defined under non-Markov policy.

Table 1: A comparison of existing algorithms in terms of regret and switching cost for linear MDP and general function class with bounded eluder dimension and Bellman completeness. The results hold for in-homogeneous episodic RL with horizon length H , number of episodes K where the total reward obtained in an episode is not larger than 1. For regret, we only present the leading term when K is large enough compared to other variables and hide poly-logarithmic factors in K , d or \dim , H and the constant. For linear MDPs, d is the dimension of the feature vectors. For general function class, \dim is a shorthand of the eluder dimension of the underlying function class, N is the covering number of the value function class, and $N_{S,A}$ is the covering number of the state-action space.

| Algorithm | Regret | # of Switches | Model Class |
|---|---|------------------------------------|--|
| LSVI-UCB (Jin et al., 2020) | $d^{3/2}H\sqrt{K}$ | K | Linear MDPs |
| LSVI-UCB-RareSwitch (Wang et al., 2021) | $d^{3/2}H\sqrt{K}$ | $\tilde{O}(dH)$ | |
| LSVI-UCB++ (He et al., 2022) | $d\sqrt{HK}$ | $\tilde{O}(dH)$ | |
| \mathcal{F} -LSVI (Wang et al., 2020a) | $\dim(\mathcal{F})\sqrt{\log N \log N_{S,A}} \cdot H\sqrt{K}$ | K | Bounded eluder dimension + Completeness |
| GOLF (Jin et al., 2021) | $\sqrt{\dim(\mathcal{F}) \log N} \cdot H\sqrt{K}$ | K | |
| VOQL (Agarwal et al., 2022) | $\sqrt{\dim(\mathcal{F}) \log N} \cdot HK$ | $\tilde{O}(\dim(\mathcal{F})^2 H)$ | |
| MQL-UCB (Theorem 4.1) | $\sqrt{\dim(\mathcal{F}) \log N} \cdot HK$ | $\tilde{O}(\dim(\mathcal{F})H)$ | |

- We propose a novel policy-switching strategy based on the cumulative sensitivity of historical data. To the best of our knowledge, this is the first computationally-tractable deterministic rare-switching algorithm for RL with general function approximation which achieves $\tilde{O}(dH)$ switching cost. We also prove a nearly matching lower bound for any algorithm with arbitrary policies including both deterministic and stochastic policies (See Theorem B.1). Previous approaches for low switching cost in RL with general function approximation are sampling-based (Kong et al., 2021; Velegkas et al., 2022; Li et al., 2023) and highly coupled with a sub-sampling technique used for regression, making it less efficient. When restricted to the linear case, sampling-based rare-switching has a worse switching cost of $\tilde{O}(d^2 H)$ (Section 3.3 in Kong et al. 2021; Theorem 3.3 in Velegkas et al. 2022; Section C in Li et al. 2023) than that in Wang et al. (2021).
- With the novel policy-switching scheme, we illustrate how to reduce the complexity of value function classes while maintaining a series of monotonic value functions, strictly generalizing the LSVI-UCB++ algorithm (He et al., 2022) to general function class with bounded eluder dimension. Based on the structure of the value functions, we further demonstrate how MQL-UCB achieves a variance-dependent regret bound that can gracefully reduce to a nearly constant regret for deterministic MDPs. While in the worst case, our algorithm still attains a nearly minimax optimal regret guarantee. Our work is the first work for RL with general function approximation that achieves the nearly minimax optimal regret when specialized to linear MDPs, while still enjoys simple Markov planning phase.

It is worth noting that recently, Qiao et al. (2023) also considered RL with low-switching cost beyond linear function approximation, i.e., MDPs with low inherent Bellman error (Zanette et al., 2020b) and generalized linear MDPs (Wang et al., 2020b). Their approach can be seen as a slight extension of RL with low-switching cost for linear MDPs, since in both settings, the covariance matrix still exists and they can still use the determinant of the covariance as a criterion for policy switching.

Notation. We use lower case letters to denote scalars and use lower and upper case bold face letters to denote vectors and matrices respectively. We denote by $[n]$ the set $\{1, \dots, n\}$. For two positive sequences $\{a_n\}$ and $\{b_n\}$ with $n = 1, 2, \dots$, we write $a_n = O(b_n)$ if there exists an absolute constant $C > 0$ such that $a_n \leq Cb_n$ holds for all $n \geq 1$ and write $a_n = \Omega(b_n)$ if there exists an absolute constant $C > 0$ such that $a_n \geq Cb_n$ holds for all $n \geq 1$. We use $\tilde{O}(\cdot)$ to further hide the polylogarithmic factors except log-covering numbers. We use $\mathbb{1}\{\cdot\}$ to denote the indicator function.

2 Preliminaries

2.1 Time-Inhomogeneous Episodic MDP

We consider a time-inhomogeneous episodic Markov Decision Process (MDP), denoted by a tuple $\mathcal{M} = M(\mathcal{S}, \mathcal{A}, H, \{\mathbb{P}_h\}_{h=1}^H, \{r_h\}_{h=1}^H)$. Here, \mathcal{S} and \mathcal{A} are the spaces of state and action, respectively, H is the length of each episode, $\mathbb{P}_h : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is the transition probability function at stage h which denotes the probability for state s to transfer to the next state s' with current action a , and $r_h : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is the deterministic reward function at stage h . A policy $\pi := \{\pi_h\}_{h=1}^H$ is a collection of mappings π_h from an observed state $s \in \mathcal{S}$ to the simplex of action space \mathcal{A} . For any policy $\pi = \{\pi_h\}_{h=1}^H$ and stage $h \in [H]$, we define the value function $V_h^\pi(s)$ and the action-value function $Q_h^\pi(s, a)$ as follows:

$$Q_h^\pi(s, a) = r_h(s, a) + \mathbb{E} \left[\sum_{h'=h+1}^H r_{h'}(s_{h'}, \pi_{h'}(s_{h'})) \middle| s_h = s, a_h = a \right], V_h^\pi(s) = Q_h^\pi(s, \pi_h(s)),$$

where $s_{h'+1} \sim \mathbb{P}_{h'}(\cdot | s_{h'}, a_{h'})$. Then, we further define the optimal value function V_h^* and the optimal action-value function Q_h^* as $V_h^*(s) = \max_{\pi} V_h^\pi(s)$ and $Q_h^*(s, a) = \max_{\pi} Q_h^\pi(s, a)$. For simplicity, we assume the total reward for each possible trajectory $(s_1, a_1, \dots, s_H, a_H)$ satisfies $\sum_{h=1}^H r_h(s_h, a_h) \leq 1$. Under this assumption, the value function $V_h^\pi(\cdot)$ and $Q_h^\pi(\cdot, \cdot)$ are bounded in $[0, 1]$. For any function $V : \mathcal{S} \rightarrow \mathbb{R}$ and stage $h \in [H]$, we define the following first-order Bellman operator \mathcal{T}_h and second-order Bellman operator \mathcal{T}_h^2 on function V :

$$\mathcal{T}_h V(s_h, a_h) = \mathbb{E}_{s_{h+1}} [r_h + V(s_{h+1}) | s_h, a_h], \mathcal{T}_h^2 V(s_h, a_h) = \mathbb{E}_{s_{h+1}} \left[(r_h + V(s_{h+1}))^2 | s_h, a_h \right],$$

where $s_{h+1} \sim \mathbb{P}_h(\cdot | s_h, a_h)$ and $r_h = r_h(s_h, a_h)$. For simplicity, we further define $[\mathbb{P}_h V](s, a) = \mathbb{E}_{s' \sim \mathbb{P}_h(\cdot | s, a)} V(s')$ and $[\mathbb{V}_h V](s, a) = \mathcal{T}_h^2 V(s_h, a_h) - (\mathcal{T}_h V(s_h, a_h))^2$. Using this notation, for each stage $h \in [H]$, the Bellman equation and Bellman optimality equation take the following forms:

$$Q_h^\pi(s, a) = \mathcal{T}_h V_{h+1}^\pi(s, a), \quad Q^*(s, a) = \mathcal{T}_h V_{h+1}^*(s, a),$$

where $V_{H+1}^\pi(\cdot) = V_{H+1}^*(\cdot) = 0$. At the beginning of each episode $k \in [K]$, the agent selects a policy π^k to be executed throughout the episode, and an initial state s_1^k is arbitrarily selected by the environment. For each stage $h \in [H]$, the agent first observes the current state s_h^k , chooses an action following the policy π_h^k , then transits to the next state with $s_{h+1}^k \sim \mathbb{P}_h(\cdot | s_h^k, a_h^k)$ and reward $r_h(s_h, a_h)$. Based on the protocol, we defined the suboptimality gap in episode k as the difference between the value function for selected policy π^k and the optimal value function $V_1^*(s_1^k) - V_1^{\pi^k}(s_1^k)$. Based on these definitions, we can define the regret in the first K episodes as follows:

Definition 2.1. For any RL algorithm Alg, the regret in the first K episodes is denoted by the sum of the suboptimality for episode $k = 1, \dots, K$,

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

where π^k is the agent's policy at episode k .

2.2 Function Classes and Covering Numbers

Assumption 2.2 (Completeness). Given $\mathcal{F} := \{\mathcal{F}_h\}_{h=1}^H$ which is composed of bounded functions $f_h : \mathcal{S} \times \mathcal{A} \rightarrow [0, L]$. We assume that for any function $V : \mathcal{S} \rightarrow [0, 1]$ there exists $f_1, f_2 \in \mathcal{F}_h$ such that for any $(s, a) \in \mathcal{S} \times \mathcal{A}$,

$$\mathbb{E}_{s' \sim \mathbb{P}_h(\cdot | s, a)} [r_h(s, a) + V(s')] = f_1(s, a), \text{ and } \mathbb{E}_{s' \sim \mathbb{P}_h(\cdot | s, a)} \left[(r_h(s, a) + V(s'))^2 \right] = f_2(s, a).$$

We assume that $L = O(1)$ throughout the paper.

Remark 2.3. Completeness is a fundamental assumption in RL with general function approximation, as recognized in Wang et al. (2020b); Jin et al. (2021); Agarwal et al. (2022). Our assumption is the same as that in Agarwal et al. (2022) and is slightly stronger than that in Wang et al. (2020a) and

Jin et al. (2021). More specifically, in Wang et al. (2020a), completeness is only required for the first-order Bellman operator. In contrast, we necessitate completeness with respect to the second-order Bellman operator, which becomes imperative during the construction of variance-based weights. Jin et al. (2021) only requires the completeness for the function class \mathcal{F}_{h+1} ($\mathcal{T}_h \mathcal{F}_{h+1} \subseteq \mathcal{F}_h$). However, the GOLF algorithm (Jin et al., 2021) requires solving an intricate optimization problem across the entire episode. In contrast, we employ pointwise exploration bonuses as an alternative strategy, which requires the completeness for function class $\mathcal{V} = \{V : \mathcal{S} \rightarrow [0, 1]\}$, i.e., $\mathcal{T}_h \mathcal{V} \subseteq \mathcal{F}_h$. The completeness assumption on the second moment is first introduced by Agarwal et al. (2022), and is crucial for obtaining a tighter regret bound. More specifically, making use of the variance of the value function at the next state is known to be crucial to achieve minimax-optimal regret bound in RL ranging from tabular MDPs (Azar et al., 2017) to linear mixture MDPs (Zhou et al., 2021a) and linear MDPs (He et al., 2022). In RL with general function approximation, the second-moment completeness assumption makes the variance of value functions computationally tractable.

Definition 2.4 (Generalized Eluder dimension, Agarwal et al. 2022). Let $\lambda \geq 0$, a sequence of state-action pairs $\mathbf{Z} = \{z_i\}_{i \in [T]}$ and a sequence of positive numbers $\boldsymbol{\sigma} = \{\sigma_i\}_{i \in [T]}$. The generalized Eluder dimension of a function class $\mathcal{F} : \mathcal{S} \times \mathcal{A} \rightarrow [0, L]$ with respect to λ is defined by $\dim_{\alpha, T}(\mathcal{F}) := \sup_{\mathbf{Z}, \boldsymbol{\sigma} : |\mathbf{Z}|=T, \sigma_i \geq \alpha} \dim(\mathcal{F}, \mathbf{Z}, \boldsymbol{\sigma})$,

$$\dim(\mathcal{F}, \mathbf{Z}, \boldsymbol{\sigma}) := \sum_{i=1}^T \min \left(1, \frac{1}{\sigma_i^2} D_{\mathcal{F}}^2(z_i; z_{[i-1]}, \sigma_{[i-1]}) \right),$$

$$D_{\mathcal{F}}^2(z; z_{[t-1]}, \sigma_{[t-1]}) := \sup_{f_1, f_2 \in \mathcal{F}} \frac{(f_1(z) - f_2(z))^2}{\sum_{s \in [t-1]} \frac{1}{\sigma_s^2} (f_1(z_s) - f_2(z_s))^2 + \lambda}.$$

We write $\dim_{\alpha, T}(\mathcal{F}) := H^{-1} \cdot \sum_{h \in [H]} \dim_{\alpha, T}(\mathcal{F}_h)$ for short when \mathcal{F} is a collection of function classes $\mathcal{F} = \{\mathcal{F}_h\}_{h=1}^H$ in the context.

Remark 2.5. The $D_{\mathcal{F}}^2$ quantity has been introduced in Agarwal et al. (2022) and Ye et al. (2023) to quantify the uncertainty of a state-action pair given a collected dataset with corresponding weights. It was inspired by Gentile et al. (2022) where an unweighted version of uncertainty has been defined for active learning. Prior to that, Wang et al. (2020a) introduced a similar ‘sensitivity’ measure to determine the sampling probability in their sub-sampling framework. As discussed in Agarwal et al. (2022), when specialized to linear function classes, $D_{\mathcal{F}}^2(z_t; z_{[t-1]}, \sigma_{[t-1]})$ can be written as the elliptical norm $\|z_t\|_{\Sigma_{t-1}}^2$, where Σ_{t-1} is the weighted covariance matrix of the feature vectors $z_{[t-1]}$.

Definition 2.6 (Bonus oracle $\bar{D}_{\mathcal{F}}^2$). In this paper, the bonus oracle is denoted by $\bar{D}_{\mathcal{F}}^2$, which computes the estimated uncertainty of a state-action pair $z = (s, a) \in \mathcal{S} \times \mathcal{A}$ with respect to historical data $z_{[t-1]}$ and corresponding weights $\sigma_{[t-1]}$. In detail, we assume that a computable function $\bar{D}_{\mathcal{F}}^2(z; z_{[t-1]}, \sigma_{[t-1]})$ satisfies $\frac{\bar{D}_{\mathcal{F}}^2(z; z_{[t-1]}, \sigma_{[t-1]})}{D_{\mathcal{F}}^2(z; z_{[t-1]}, \sigma_{[t-1]})} \in [1, C]$, where C is a fixed constant.

Remark 2.7. Agarwal et al. (2022) also assumed access to such a bonus oracle defined in Definition 2.6, where they assume that the bonus oracle finds a proper bonus from a finite bonus class (Definition 3, Agarwal et al. 2022). Our definition is slightly different in the sense that the bonus class is not assumed to be finite but with a finite covering number. Previous works by Kong et al. (2021) and Wang et al. (2020a) proposed a sub-sampling idea to compute such a bonus function efficiently in general cases, which is also applicable in our framework. In a similar nonlinear RL setting, Ye et al. (2023) assumed that the uncertainly $D_{\mathcal{F}}^2$ can be directly computed, which is a slightly stronger assumption. But essentially, these differences in bonus assumption only lightly affect the algorithm structure.

Definition 2.8 (Covering numbers of function classes). For any $\epsilon > 0$, we define the following covering numbers of the involved function classes:

1. For each $h \in [H]$, there exists an ϵ -cover $\mathcal{C}(\mathcal{F}_h, \epsilon) \subseteq \mathcal{F}_h$ with size $|\mathcal{C}(\mathcal{F}_h, \epsilon)| \leq N(\mathcal{F}_h, \epsilon)$, such that for any $f \in \mathcal{F}$, there exists $f' \in \mathcal{C}(\mathcal{F}_h, \epsilon)$, such that $\|f - f'\|_{\infty} \leq \epsilon$. For any $\epsilon > 0$, we define the uniform covering number of \mathcal{F} with respect to ϵ as $N_{\mathcal{F}}(\epsilon) := \max_{h \in [H]} N(\mathcal{F}_h, \epsilon)$.
2. There exists a bonus class $\mathcal{B} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ such that for any $t \geq 0$, $z_{[t]} \in (\mathcal{S} \times \mathcal{A})^t$, $\sigma_{[t]} \in \mathbb{R}^t$, the oracle defined in Definition 2.6 $\bar{D}_{\mathcal{F}}^2(\cdot; z_{[t]}, \sigma_{[t]})$ is in \mathcal{B} .
3. For bonus class \mathcal{B} , there exists an ϵ -cover $\mathcal{C}(\mathcal{B}, \epsilon) \subseteq \mathcal{B}$ with size $|\mathcal{C}(\mathcal{B}, \epsilon)| \leq N(\mathcal{B}, \epsilon)$, such that for any $b \in \mathcal{B}$, there exists $b' \in \mathcal{C}(\mathcal{B}, \epsilon)$, such that $\|b - b'\|_{\infty} \leq \epsilon$.

Remark 2.9. In general function approximation, it is common to introduce the additional assumption on the covering number of bonus function classes. For example, in [Ye et al. \(2023\)](#), [Agarwal and Zhang \(2022a\)](#), and [Di et al. \(2023\)](#), the covering number of the bonus function class is bounded.

3 Algorithm and Key Techniques

In this section, we will introduce our new algorithm, MQL-UCB. The detailed algorithm is provided in [Algorithm 1](#). Our algorithm’s foundational framework follows the Upper Confidence Bound (UCB) approach. In detail, for each episode $k \in [K]$, we construct an optimistic value function $Q_{k,h}(s, a)$ during the planning phase. Subsequently, in the exploration phase, the agent interacts with the environment, employing a greedy policy based on the current optimistic value function $Q_{k,h}(s, a)$. Once the agent obtains the reward r_h^k and transitions to the next state s_{h+1}^k , these outcomes are incorporated into the dataset, contributing to the subsequent planning phase. We will proceed to elucidate the essential components of our method.

3.1 Rare Policy Switching

For MQL-UCB algorithm, the value functions $Q_{k,h}, \check{Q}_{k,h}$, along with their corresponding policy π_k , undergo updates when the agent collects a sufficient number of trajectories within the dataset that could significantly diminish the uncertainty associated with the Bellman operator $\mathcal{T}_h V(\cdot, \cdot)$ through the weighted regression. In the context of linear bandits ([Abbasi-Yadkori et al., 2011](#)) or linear MDPs ([He et al., 2022](#)), the uncertainty pertaining to the least-square regression is quantified by the covariance matrix Σ_k . In this scenario, the agent adjusts its policy once the determinant of the covariance matrix doubles, employing a determinant-based criterion. Nevertheless, in the general function approximation setting, such a method is not applicable in the absence of the covariance matrix which serves as a feature extractor in the linear setting. Circumventing this issue, [Kong et al. \(2021\)](#) proposed a sub-sampling-based method to achieve low-switching properties in nonlinear RL. Their subsampling technique is inspired by [Wang et al. \(2021\)](#), which showed that one only needs to maintain a small subset of historical data to obtain a sufficiently accurate least-square estimator. Such a subset can be generated sequentially according to the *sensitivity* of a new coming data point. However, their approach leads to a switching cost of $\tilde{O}(d^2 H)$, which does not match the lower bound in linear MDPs proposed by [Gao et al. \(2021\)](#).

To resolve this issue, we proposed a more general deterministic policy-updating framework for nonlinear RL. In detail, we use $\bar{D}_{\mathcal{F}_h}^2(z_{i,h}; z_{[k_{last}-1],h}, \bar{\sigma}_{[k_{last}-1],h})$ to evaluate the information collected at the episode i , given the last updating k_{last} . Once the collected information goes beyond a threshold χ from last updating, i.e.,

$$\sum_{i=k_{last}}^{k-1} \frac{1}{\bar{\sigma}_{i,h}^2} \bar{D}_{\mathcal{F}_h}^2(z_{i,h}; z_{[k_{last}-1],h}, \bar{\sigma}_{[k_{last}-1],h}) \geq \chi. \quad (3.1)$$

the agent will perform updates on both the optimistic estimated value function and the pessimistic value function. Utilizing the $\bar{D}_{\mathcal{F}_h}^2$ -based criterion, we will show that the number of policy updates can be bounded by $O(H \cdot \dim_{\alpha,K}(\mathcal{F}))$. This further reduces the complexity of the optimistic value function class and removes additional factors from a uniform convergence argument over the function class. Specifically, we showcase under our rare-switching framework, the ϵ -covering number of the optimistic and the pessimistic value function class at episode k is bounded by

$$N_\epsilon(k) := [N_{\mathcal{F}}(\epsilon/2) \cdot N(\mathcal{B}, \epsilon/2\hat{\beta}_K)]^{l_k+1}, \quad (3.2)$$

where $\hat{\beta}_K$ is the maximum confidence radius as shown in [Algorithm 1](#), which will be specified in [Lemmas G.3](#) and [G.4](#), l_k is the number of policy switches before the end of the k -th episode according to [Algorithm 1](#).

3.2 Weighted Regression

The estimation of Q function in MQL-UCB extends LSVI-UCB proposed by [Jin et al. \(2020\)](#) to general function classes. While the estimators in LSVI-UCB are computed from the classic least

squares regression, we construct the estimated value function $\hat{f}_{k,h}$ for general function classes by solving the following weighted regression:

$$\hat{f}_{k,h} = \operatorname{argmin}_{f_h \in \mathcal{F}_h} \sum_{i \in [k-1]} \frac{1}{\bar{\sigma}_{i,h}^2} (f_h(s_h^i, a_h^i) - r_h^i - V_{k,h+1}(s_{h+1}^i))^2.$$

In the weighted regression, we set the weight $\bar{\sigma}_{k,h}$ as

$$\bar{\sigma}_{k,h} = \max \{ \sigma_{k,h}, \alpha, \gamma \cdot \bar{D}_{\mathcal{F}_h}^{1/2}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) \},$$

where $\sigma_{k,h}$ is the estimated variance for the stochastic transition process, $\bar{D}_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h})$ denotes the uncertainty of the estimated function $\hat{f}_{k,h}$ conditioned on the historical observations and

$$\gamma^2 := \log \left((2HK^2(2 \log(L^2k/\alpha^4) + 2) \cdot (\log(4L/\alpha^2) + 2) \cdot N_{\mathcal{F}}^4(\epsilon) \cdot N_\epsilon^2(K)) / \delta \right) \quad (3.3)$$

is used to properly balance the uncertainty across various state-action pairs. It is worth noting that [Ye et al. \(2023\)](#) also introduced the uncertainty-aware variance in the general function approximation with a distinct intention to deal with the adversarial corruption from the attacker. In addition, according to the weighted regression, with high probability, the Bellman operator $\mathcal{T}_h V_{k,h}$ satisfies:

$$\lambda + \sum_{i \in [k-1]} \bar{\sigma}_{i,h}^{-2} \left(\hat{f}_{k,h}(s_h^i, a_h^i) - \mathcal{T}_h V_{k,h+1}(s_h^i, a_h^i) \right)^2 \leq \hat{\beta}_k^2,$$

where $\hat{\beta}_k$ is the exploration radius for the Q functions, which will be specified in [Theorem 4.1](#). According to the definition of Generalized Eluder dimension, the estimation error between the estimated function $\hat{f}_{k,h}$ and the Bellman operator is upper bounded by:

$$|\hat{f}_{k,h}(s, a) - \mathcal{T}_h V_{k,h+1}(s, a)| \leq \hat{\beta}_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}).$$

Therefore, we introduce the exploration bonus $b_{k,h}$ and construct the optimistic value function $Q_{k,h}(s, a)$, i.e., $Q_{k,h}(s, a) \approx \hat{f}_{k,h}(s, a) + b_{k,h}(s, a)$, where $b_{k,h}(s, a) = \hat{\beta}_k \cdot \bar{D}_{\mathcal{F}}((s, a); z_{[k-1],h}, \sigma_{[k-1],h})$. Inspired by [Hu et al. \(2022\)](#); [He et al. \(2022\)](#); [Agarwal et al. \(2022\)](#), in order to estimate the gap between the optimistic value function $V_{k,h}(s)$ and the optimal value function $V_h^*(s)$, we further construct the pessimistic estimator $\check{f}_{k,h}$ by the following weighted regression

$$\check{f}_{k,h} = \operatorname{argmin}_{f_h \in \mathcal{F}_h} \sum_{i \in [k-1]} \frac{1}{\bar{\sigma}_{i,h}^2} (f_h(s_h^i, a_h^i) - r_h^i - \check{V}_{k,h+1}(s_{h+1}^i))^2,$$

and introduce a negative exploration bonus when generating the pessimistic estimator. $\check{Q}_{k,h}(s, a) \approx \check{f}_{k,h}(s, a) - \check{b}_{k,h}(s, a)$, where $\check{b}_{k,h}(s, a) = \check{\beta}_k \cdot \bar{D}_{\mathcal{F}}((s, a); z_{[k-1],h}, \sigma_{[k-1],h})$. Different from [Agarwal et al. \(2022\)](#), the pessimistic value function $\check{f}_{k,h}$ is computed from a similar weighted-regression scheme as in the case of the optimistic estimator, leading to a tighter confidence set.

3.3 Variance Estimator

In this subsection, we provide more details about the variance estimator $\sigma_{k,h}$, which measures the variance of the value function $V_{k,h+1}(s_{h+1}^k)$ caused by the stochastic transition from state-action pair (s_h^k, a_h^k) . According to the definition of $\hat{f}_{k,h}$, the difference between the estimator $\hat{f}_{k,h}$ and $\mathcal{T}_h V_{k,h+1}$ satisfies

$$\lambda + \sum_{i \in [k-1]} \frac{1}{\bar{\sigma}_{i,h}^2} (\hat{f}_{k,h}(s_h^i, a_h^i) - \mathcal{T}_h V_{k,h+1}(s_h^i, a_h^i))^2 \leq 2 \sum_{i \in [k-1]} \frac{1}{\bar{\sigma}_{i,h}^2} (f(s_h^i, a_h^i) - \hat{f}_k^*(s_h^i, a_h^i)) \cdot \hat{\eta}_h^*(V_{k,h+1}),$$

where the noise $\hat{\eta}_h^k(V) = r_h^k + V(s_{h+1}^k) - \mathbb{E}_{s' \sim \mathbb{P}_h(s_h^k, a_h^k)} [r_h(s_h^k, a_h^k, s') + V(s')]$ denotes the stochastic transition noise for the value function V . However, the generation of the target function $V_{k,h+1}$ relies on previously collected data $z_{[k_{\text{max}}]}$, thus violating the conditional independence property. Consequently, the noise term $\hat{\eta}_h^k(V_{k,h+1})$ may not be unbiased. To address this challenge, it becomes imperative to establish a uniform convergence property over the potential function class, which is first introduced in linear MDPs by [Jin et al. \(2020\)](#).

Following the previous approach introduced by [Azar et al. \(2017\)](#); [Hu et al. \(2022\)](#); [Agarwal et al. \(2022\)](#); [He et al. \(2022\)](#), we decompose the noise of optimistic value function $\hat{\eta}_h^k(V_{k,h+1})$ into the

Algorithm 1 Monotonic Q-Learning with UCB (MQL-UCB)

Require: Regularization parameter λ , confidence radius $\{\tilde{\beta}_k\}_{k \in [K]}$, $\{\hat{\beta}_k\}_{k \in [K]}$ and $\{\check{\beta}_k\}_{k \in [K]}$.

- 1: Initialize $k_{\text{last}} = 0$. For each stage $h \in [H]$ and state-action $(s, a) \in \mathcal{S} \times \mathcal{A}$, set $Q_{0,h}(s, a) \leftarrow H, \check{Q}_{0,h}(s, a) \leftarrow 0$.
 - 2: **for** episodes $k = 1, \dots, K$ **do**
 - 3: Received the initial state s_1^k .
 - 4: **for** stage $h = H, \dots, 1$ **do**
 - 5: **if** there exists a stage $h' \in [H]$ such that (3.1) holds **then**
 - 6: $\hat{f}_{k,h} \leftarrow \operatorname{argmin}_{f_h \in \mathcal{F}_h} \sum_{i \in [k-1]} \frac{1}{\bar{\sigma}_{i,h}^2} (f_h(s_h^i, a_h^i) - r_h^i - V_{k,h+1}(s_{h+1}^i))^2$.
 - 7: $\check{f}_{k,h} \leftarrow \operatorname{argmin}_{f_h \in \mathcal{F}_h} \sum_{i \in [k-1]} \frac{1}{\bar{\sigma}_{i,h}^2} (f_h(s_h^i, a_h^i) - r_h^i - \check{V}_{k,h+1}(s_{h+1}^i))^2$.
 - 8: $\tilde{f}_{k,h} \leftarrow \operatorname{argmin}_{f_h \in \mathcal{F}_h} \sum_{i \in [k-1]} (f_h(s_h^i, a_h^i) - (r_h^i + V_{k,h+1}(s_{h+1}^i))^2)^2$.
 - 9: $Q_{k,h}(s, a) \leftarrow \min \left\{ \hat{f}_{k,h}(s, a) + b_{k,h}(s, a), Q_{k-1,h}(s, a), 1 \right\}$.
 - 10: $\check{Q}_{k,h}(s, a) \leftarrow \max \left\{ \check{f}_{k,h}(s, a) - \check{b}_{k,h}(s, a), \check{Q}_{k-1,h}(s, a), 0 \right\}$.
 - 11: Set the last updating episode $k_{\text{last}} \leftarrow k$ and number of policies as $l_k \leftarrow l_{k-1} + 1$.
 - 12: **else**
 - 13: $Q_{k,h}(s, a) \leftarrow Q_{k-1,h}(s, a), \check{Q}_{k,h}(s, a) \leftarrow \check{Q}_{k-1,h}(s, a)$ and $l_k \leftarrow l_{k-1}$.
 - 14: **end if**
 - 15: Set the policy π^k as $\pi_h^k(\cdot) \leftarrow \operatorname{argmax}_{a \in \mathcal{A}} Q_{k,h}(\cdot, a)$. $V_{k,h}(s) \leftarrow \max_a Q_{k,h}(s, a)$, $\check{V}_{k,h}(s) \leftarrow \max_a \check{Q}_{k,h}(s, a)$.
 - 16: **end for**
 - 17: **for** stage $h = 1, \dots, H$ **do**
 - 18: Take action $a_h^k \leftarrow \pi_h^k(s_h^k)$ and receive next state s_{h+1}^k .
 - 19: Set the estimated variance $\sigma_{k,h}$ as in (3.4) and set $\bar{\sigma}_{k,h} \leftarrow \max \left\{ \sigma_{k,h}, \alpha, \gamma \cdot D_{\mathcal{F}_h}^{1/2}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) \right\}$.
 - 20: **end for**
 - 21: **end for**
-

noise of *optimal* value function $\hat{\eta}_h^k(V_{h+1}^*)$ and the noise $\hat{\eta}_h^k(V_{k,h+1} - V_{h+1}^*)$ to reduce the extra $\log(N_\epsilon(K))$ dependency in the confidence radius. With the noise decomposition, we evaluate the variances $[\mathbb{V}_h V_{h+1}^*](s, a)$ and $[\mathbb{V}_h(V_{k,h+1} - V_{h+1}^*)](s, a)$ separately.

For the variance of the *optimal* value function $[\mathbb{V}_h V_{h+1}^*](s, a)$, since the optimal value function V_{h+1}^* is independent with the collected data $z_{[k_{\text{last}}]}$, it prevents a uniform convergence-based argument over the function class. However, the optimal value function V_{h+1}^* is unobservable, and it requires several steps to estimate the variance. In summary, we utilize the optimistic function $V_{k,h+1}$ to approximate the optimal value function V_{h+1}^* and calculate the estimated variance $[\bar{\mathbb{V}}_h V_{k,h+1}]$ as the difference between the second-order moment and the square of the first-order moment of $V_{k,h}$

$$[\bar{\mathbb{V}}_{k,h} V_{k,h+1}] = \tilde{f}_{k,h} - \hat{f}_{k,h}^2.$$

Here, the approximate second-order moment $\tilde{f}_{k,h}$ and the approximate first-order moment $\hat{f}_{k,h}$ is generated by the least-square regression (Lines 6 and 8). In addition, we introduce the exploration bonus $E_{k,h}$ to control the estimation error between the estimated variance and the true variance of $V_{k,h+1}$ and $F_{k,h}$ to control the sub-optimality gap between $V_{k,h+1}$ and V_{h+1}^* :

$$E_{k,h} = (2L\beta_k + \tilde{\beta}_k) \min(1, \bar{D}_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h})),$$

$$F_{k,h} = (\log(N(\mathcal{F}, \epsilon) \cdot N_\epsilon(K))) \cdot \min(1, 2\hat{f}_{k,h}(s_h^k, a_h^k) - 2\check{f}_{k,h}(s_h^k, a_h^k) + 4\beta_k \bar{D}_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h})),$$

where

$$\tilde{\beta}_k = \sqrt{128 \log \frac{N_\epsilon(k) \cdot N(\mathcal{F}, \epsilon) \cdot H}{\delta}} + 64L\epsilon \cdot k, \beta_k = \sqrt{128 \cdot \log \frac{N_\epsilon(k) \cdot N(\mathcal{F}, \epsilon) H}{\delta}} + 64L\epsilon \cdot k / \alpha^2.$$

For the variance of the sub-optimality gap, $[\mathbb{V}_h(V_{k,h+1} - V_{h+1}^*)](s, a)$, based on the structure of optimistic and pessimistic value function, it can be approximate and upper bounded by

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

$$\leq 2[\mathbb{P}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \approx 2\hat{f}_{k,h}(s_h^k, a_h^k) - 2\check{f}_{k,h}(s_h^k, a_h^k),$$

where the approximate first-order moments $\hat{f}_{k,h}, \check{f}_{k,h}$ are generated by the least-square regression (Lines 6 and 7) and can be dominated by the exploration bonus $F_{k,h}$.

In summary, we construct the estimated variance $\sigma_{k,h}$ as:

$$\sigma_{k,h} = \sqrt{[\check{\mathbb{V}}_{k,h} V_{k,h+1}](s_h^k, a_h^k) + E_{k,h} + F_{k,h}}. \quad (3.4)$$

3.4 Monotonic Value Function

As we discussed in the previous subsection, we decompose the value function $V_{k,h}$ and evaluate the variance $[\mathbb{V}_h V_{h+1}^*](s, a)$, $[\mathbb{V}_h(V_{k,h+1} - V_{h+1}^*)](s, a)$ separately. However, for each state-action pair (s_h^k, a_h^k) and any subsequent episode $i > k$, the value function $V_{i,h}$ and corresponding variance $[\mathbb{V}_h(V_{i,h+1} - V_{h+1}^*)](s_h^k, a_h^k)$ may differ from the previous value function $V_{k,h}$ and variance $[\mathbb{V}_h(V_{k,h+1} - V_{h+1}^*)](s_h^k, a_h^k)$. Extending the idea proposed by He et al. (2022) for linear MDPs, we ensure that the pessimistic value function $\check{Q}_{k,h}$ maintains a monotonically increasing property during updates, while the optimistic value function $\check{Q}_{k,h}$ maintains a monotonically decreasing property. Leveraging these monotonic properties, we can establish an upper bound on the variance as follows:

$$[\mathbb{V}_h(V_{i,h+1} - V_{h+1}^*)](s_h^k, a_h^k) \leq 2[\mathbb{P}_h(V_{i,h+1} - \check{V}_{i,h+1})](s_h^k, a_h^k) \leq 2[\mathbb{P}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \leq F_{k,h}.$$

In this scenario, the previously employed variance estimator $\sigma_{k,h}$ offers a consistent and uniform upper bound for the variance across all subsequent episodes.

4 Main Results

In this section, we present the main theoretical results. In detail, we provide the regret upper bound of Algorithm MQL-UCB in Theorem 4.1. As a complement, in Section B.1, we provide a lower bound on the communication complexity for cooperative linear MDPs. Finally, in Section B.2, we discussed the connection between the generalized eluder dimension and the standard eluder dimension.

The following theorem provides the regret upper bound of Algorithm MQL-UCB.

Theorem 4.1. Suppose Assumption 2.2 holds for function classes $\mathcal{F} := \{\mathcal{F}_h\}_{h=1}^H$ and Definition 2.4 holds with $\lambda = 1$. If we set $\alpha = 1/\sqrt{KH}$, $\epsilon = (KLH)^{-1}$, and set $\hat{\beta}_k^2 = \check{\beta}_k^2 := O\left(\log \frac{2k^2(2\log(L^2k/\alpha^4)+2) \cdot (\log(4L/\alpha^2)+2)}{\delta/H}\right) \cdot [\log(N_{\mathcal{F}}(\epsilon)) + 1] + O(\lambda) + O(\epsilon kL/\alpha^2)$, then with probability $1 - O(\delta)$, the regret of MQL-UCB is upper bounded as follows:

$$\begin{aligned} \text{Regret}(K) &= \tilde{O}\left(\sqrt{\dim(\mathcal{F}) \log N \cdot H \text{Var}_K}\right) \\ &+ \tilde{O}\left(H^{2.5} \dim^2(\mathcal{F}) \sqrt{\log N \log(N \cdot N_b)}\right) \cdot \sqrt{H \log N + \dim(\mathcal{F}) \log(N \cdot N_b)} \end{aligned}$$

where $\text{Var}_K := \sum_{k=1}^K \sum_{h=1}^H [\mathbb{V}_h V_{h+1}^{\pi^k}](s_h^k, a_h^k) = \tilde{O}(K)$, we denote the covering number of bonus function class by N_b , the covering number of function class \mathcal{F} by N , and the dimension $\dim_{\alpha,K}(\mathcal{F})$ by $\dim(\mathcal{F})$. Meanwhile, the switching cost of Algorithm 1 is $O(\dim_{\alpha,K}(\mathcal{F}) \cdot H)$.

In the worst case, when the number of episodes K is sufficiently large, the leading term in our regret bound is $\tilde{O}\left(\sqrt{\dim(\mathcal{F}) \log N \cdot HK}\right)$. Our result matches the optimal regret achieved by Agarwal et al. (2022). While their proposed algorithm involves the execution of a complicated and non-Markovian policy with an action selection phase based on two series of optimistic value functions and the prefix trajectory, MQL-UCB only requires the knowledge of the current state and a single optimistic state-action value function Q to make a decision over the action space. In addition, our theorem also provides a variance-dependent regret bound, which is adaptive to the randomness of the underlying MDP encountered by the agent. Our definition of Var_K is inspired by Zhao et al. (2023) and Zhou et al. (2023), which studied variance-adaptive RL under tabular MDP setting and linear mixture MDP setting, respectively.

As a direct application, we also present the regret guarantee of MQL-UCB for linear MDPs.

Corollary 4.2. Under the same conditions of Theorem 4.1, assume that the underlying MDP is a linear MDP such that $\mathcal{F} := \{\mathcal{F}_h\}_{h \in [H]}$ is composed of linear function classes with a known feature mapping over the state-action space $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$. If we set $\lambda = 1$, $\alpha = 1/\sqrt{K}$, then with probability $1 - O(\delta)$, the following cumulative regret guarantee holds for MQL-UCB:

$$\text{Regret}(K) = \tilde{O}(d\sqrt{HK} + H^{2.5}d^5\sqrt{H + d^2}).$$

Remark 4.3. The leading term in our regret bound, as demonstrated in Corollary 4.2, matches the lower bound proved in Zhou et al. (2021a) for linear MDPs. Similar optimal regrets have also been accomplished by He et al. (2022) and Agarwal et al. (2022) for linear MDPs. Since we also apply weighted regression to enhance the precision of our pessimistic value functions, the lower order term (i.e., $H^{2.5}d^5\sqrt{H + d^2}$) in our regret has a better dependency on H than VOQL (Agarwal et al., 2022) and LSVI-UCB++ (He et al., 2022), which may be of independent interest when considering long-horizon MDPs. In addition, the switching cost of Algorithm 1 is bounded by $\tilde{O}(dH)$, which matches the lower bound in Gao et al. (2021) for deterministic algorithms and our new lower bound in Theorem B.1 for arbitrary algorithms up to logarithmic factors. For more details about our lower bound, please refer to Appendix E.

5 Conclusion and Future Work

In this paper, we delve into the realm of RL with general function approximation. We proposed the MQL-UCB algorithm with an innovative uncertainty-based rare-switching strategy in general function approximation. Notably, our algorithm only requires $\tilde{O}(dH)$ updating times, which matches with the lower bound established by Gao et al. (2021) up to logarithmic factors, and obtains a $\tilde{O}(d\sqrt{HK})$ regret guarantee, which is near-optimal when restricted to the linear cases.

Acknowledgments and Disclosure of Funding

We thank the anonymous reviewers for their helpful comments. HZ is partially supported by the research award from Cisco. JH and QG are partially supported by the research fund from UCLA-Amazon Science Hub. The views and conclusions contained in this paper are those of the authors and should not be interpreted as representing any funding agencies.

References

- ABBASI-YADKORI, Y., PÁL, D. and SZEPESVÁRI, C. (2011). Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems* **24** 2312–2320.
- AFSAR, M. M., CRUMP, T. and FAR, B. (2022). Reinforcement learning based recommender systems: A survey. *ACM Computing Surveys* **55** 1–38.
- AGARWAL, A., JIN, Y. and ZHANG, T. (2022). Voql: Towards optimal regret in model-free rl with nonlinear function approximation. *arXiv preprint arXiv:2212.06069*.
- AGARWAL, A. and ZHANG, T. (2022a). Model-based rl with optimistic posterior sampling: Structural conditions and sample complexity. *Advances in Neural Information Processing Systems* **35** 35284–35297.
- AGARWAL, A. and ZHANG, T. (2022b). Non-linear reinforcement learning in large action spaces: Structural conditions and sample-efficiency of posterior sampling. In *Conference on Learning Theory*. PMLR.
- AYOUB, A., JIA, Z., SZEPESVARI, C., WANG, M. and YANG, L. (2020). Model-based reinforcement learning with value-targeted regression. In *International Conference on Machine Learning*. PMLR.
- AZAR, M. G., OSBAND, I. and MUNOS, R. (2017). Minimax regret bounds for reinforcement learning. In *International Conference on Machine Learning*. PMLR.
- BAI, Y., XIE, T., JIANG, N. and WANG, Y.-X. (2019). Provably efficient q-learning with low switching cost. *Advances in Neural Information Processing Systems* **32**.

- CHEN, Z., LI, C. J., YUAN, H., GU, Q. and JORDAN, M. (2022). A general framework for sample-efficient function approximation in reinforcement learning. In *The Eleventh International Conference on Learning Representations*.
- DANN, C., JIANG, N., KRISHNAMURTHY, A., AGARWAL, A., LANGFORD, J. and SCHAPIRE, R. E. (2018). On oracle-efficient pac rl with rich observations. *Advances in neural information processing systems* **31**.
- DI, Q., ZHAO, H., HE, J. and GU, Q. (2023). Pessimistic nonlinear least-squares value iteration for offline reinforcement learning. *arXiv preprint arXiv:2310.01380*.
- DU, S., KAKADE, S., LEE, J., LOVETT, S., MAHAJAN, G., SUN, W. and WANG, R. (2021). Bilinear classes: A structural framework for provable generalization in rl. In *International Conference on Machine Learning*. PMLR.
- DU, S. S., KAKADE, S. M., WANG, R. and YANG, L. F. (2019). Is a good representation sufficient for sample efficient reinforcement learning? In *International Conference on Learning Representations*.
- FOSTER, D. J., KAKADE, S. M., QIAN, J. and RAKHLIN, A. (2021). The statistical complexity of interactive decision making. *arXiv preprint arXiv:2112.13487*.
- GAO, M., XIE, T., DU, S. S. and YANG, L. F. (2021). A provably efficient algorithm for linear markov decision process with low switching cost. *arXiv preprint arXiv:2101.00494*.
- GENTILE, C., WANG, Z. and ZHANG, T. (2022). Achieving minimax rates in pool-based batch active learning. In *International Conference on Machine Learning*. PMLR.
- HAN, Y., ZHOU, Z., ZHOU, Z., BLANCHET, J., GLYNN, P. W. and YE, Y. (2020). Sequential batch learning in finite-action linear contextual bandits. *arXiv preprint arXiv:2004.06321*.
- HE, J., ZHAO, H., ZHOU, D. and GU, Q. (2022). Nearly minimax optimal reinforcement learning for linear markov decision processes. *arXiv preprint arXiv:2212.06132*.
- HE, J., ZHOU, D. and GU, Q. (2021a). Logarithmic regret for reinforcement learning with linear function approximation. In *International Conference on Machine Learning*. PMLR.
- HE, J., ZHOU, D. and GU, Q. (2021b). Nearly minimax optimal reinforcement learning for discounted mdps. *Advances in Neural Information Processing Systems* **34** 22288–22300.
- HU, P., CHEN, Y. and HUANG, L. (2022). Nearly minimax optimal reinforcement learning with linear function approximation. In *International Conference on Machine Learning*. PMLR.
- HUANG, J., CHEN, J., ZHAO, L., QIN, T., JIANG, N. and LIU, T.-Y. (2022). Towards deployment-efficient reinforcement learning: Lower bound and optimality. *arXiv preprint arXiv:2202.06450*.
- JIA, Z., YANG, L., SZEPESVARI, C. and WANG, M. (2020). Model-based reinforcement learning with value-targeted regression. In *Learning for Dynamics and Control*. PMLR.
- JIANG, N., KRISHNAMURTHY, A., AGARWAL, A., LANGFORD, J. and SCHAPIRE, R. E. (2017). Contextual decision processes with low bellman rank are pac-learnable. In *International Conference on Machine Learning*. PMLR.
- JIN, C., ALLEN-ZHU, Z., BUBECK, S. and JORDAN, M. I. (2018). Is q-learning provably efficient? *Advances in neural information processing systems* **31**.
- JIN, C., LIU, Q. and MIRYOUSEFI, S. (2021). Bellman eluder dimension: New rich classes of rl problems, and sample-efficient algorithms. *Advances in neural information processing systems* **34** 13406–13418.
- JIN, C., YANG, Z., WANG, Z. and JORDAN, M. I. (2020). Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*. PMLR.

- KOBER, J., BAGNELL, J. A. and PETERS, J. (2013). Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research* **32** 1238–1274.
- KONG, D., SALAKHUTDINOV, R., WANG, R. and YANG, L. F. (2021). Online sub-sampling for reinforcement learning with general function approximation. *arXiv preprint arXiv:2106.07203* .
- LI, Y., WANG, Y., CHENG, Y. and YANG, L. (2023). Low-switching policy gradient with exploration via online sensitivity sampling. In *International Conference on Machine Learning*. PMLR.
- LILICRAP, T. P., HUNT, J. J., PRITZEL, A., HEES, N., EREZ, T., TASSA, Y., SILVER, D. and WIERSTRA, D. (2015). Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971* .
- MATSUSHIMA, T., FURUTA, H., MATSUO, Y., NACHUM, O. and GU, S. (2020). Deployment-efficient reinforcement learning via model-based offline optimization. In *International Conference on Learning Representations*.
- MNIH, V., KAVUKCUOGLU, K., SILVER, D., RUSU, A. A., VENESS, J., BELLEMARE, M. G., GRAVES, A., RIEDMILLER, M., FIDJELAND, A. K., OSTROVSKI, G. ET AL. (2015). Human-level control through deep reinforcement learning. *nature* **518** 529–533.
- 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*.
- PUTERMAN, M. L. (2014). *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.
- QIAO, D., YIN, M., MIN, M. and WANG, Y.-X. (2022). Sample-efficient reinforcement learning with loglog (t) switching cost. In *International Conference on Machine Learning*. PMLR.
- QIAO, D., YIN, M. and WANG, Y.-X. (2023). Logarithmic switching cost in reinforcement learning beyond linear mdps. *arXiv preprint arXiv:2302.12456* .
- RUAN, Y., YANG, J. and ZHOU, Y. (2021). Linear bandits with limited adaptivity and learning distributional optimal design. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*.
- RUSSO, D. and VAN ROY, B. (2013). Eluder dimension and the sample complexity of optimistic exploration. *Advances in Neural Information Processing Systems* **26**.
- RUSSO, D. and VAN ROY, B. (2014). Learning to optimize via posterior sampling. *Mathematics of Operations Research* **39** 1221–1243.
- SIMCHOWITZ, M. and JAMIESON, K. G. (2019). Non-asymptotic gap-dependent regret bounds for tabular mdps. *Advances in Neural Information Processing Systems* **32**.
- SUN, W., JIANG, N., KRISHNAMURTHY, A., AGARWAL, A. and LANGFORD, J. (2019). Model-based rl in contextual decision processes: Pac bounds and exponential improvements over model-free approaches. In *Conference on learning theory*. PMLR.
- VELEGKAS, G., YANG, Z. and KARBASI, A. (2022). Reinforcement learning with logarithmic regret and policy switches. *Advances in Neural Information Processing Systems* **35** 36040–36053.
- WANG, R., SALAKHUTDINOV, R. R. and YANG, L. (2020a). Reinforcement learning with general value function approximation: Provably efficient approach via bounded eluder dimension. *Advances in Neural Information Processing Systems* **33** 6123–6135.
- WANG, T., ZHOU, D. and GU, Q. (2021). Provably efficient reinforcement learning with linear function approximation under adaptivity constraints. *Advances in Neural Information Processing Systems* **34** 13524–13536.
- WANG, Y., WANG, R., DU, S. S. and KRISHNAMURTHY, A. (2020b). Optimism in reinforcement learning with generalized linear function approximation. In *International Conference on Learning Representations*.

- XIONG, N., YANG, Z. and WANG, Z. (2023). A general framework for sequential decision-making under adaptivity constraints. *arXiv preprint arXiv:2306.14468* .
- YANG, L. and WANG, M. (2019). Sample-optimal parametric q-learning using linearly additive features. In *International Conference on Machine Learning*. PMLR.
- YANG, L. and WANG, M. (2020). Reinforcement learning in feature space: Matrix bandit, kernels, and regret bound. In *International Conference on Machine Learning*. PMLR.
- YE, C., XIONG, W., GU, Q. and ZHANG, T. (2023). Corruption-robust algorithms with uncertainty weighting for nonlinear contextual bandits and markov decision processes. In *International Conference on Machine Learning*. PMLR.
- ZANETTE, A., BRANDFONBRENER, D., BRUNSKILL, E., PIROTTA, M. and LAZARIC, A. (2020a). Frequentist regret bounds for randomized least-squares value iteration. In *International Conference on Artificial Intelligence and Statistics*.
- ZANETTE, A. and BRUNSKILL, E. (2019). Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. In *International Conference on Machine Learning*. PMLR.
- ZANETTE, A., LAZARIC, A., KOCHENDERFER, M. and BRUNSKILL, E. (2020b). Learning near optimal policies with low inherent bellman error. In *International Conference on Machine Learning*. PMLR.
- ZHANG, Z. and JI, X. (2019). Regret minimization for reinforcement learning by evaluating the optimal bias function. *Advances in Neural Information Processing Systems* **32**.
- ZHANG, Z., JI, X. and DU, S. (2021). Is reinforcement learning more difficult than bandits? a near-optimal algorithm escaping the curse of horizon. In *Conference on Learning Theory*. PMLR.
- ZHANG, Z., JIANG, Y., ZHOU, Y. and JI, X. (2022). Near-optimal regret bounds for multi-batch reinforcement learning. *Advances in Neural Information Processing Systems* **35** 24586–24596.
- ZHANG, Z., ZHOU, Y. and JI, X. (2020). Almost optimal model-free reinforcement learning via reference-advantage decomposition. *Advances in Neural Information Processing Systems* **33** 15198–15207.
- ZHAO, H., HE, J., ZHOU, D., ZHANG, T. and GU, Q. (2023). Variance-dependent regret bounds for linear bandits and reinforcement learning: Adaptivity and computational efficiency. *arXiv preprint arXiv:2302.10371* .
- ZHENG, G., ZHANG, F., ZHENG, Z., XIANG, Y., YUAN, N. J., XIE, X. and LI, Z. (2018). Dnn: A deep reinforcement learning framework for news recommendation. In *Proceedings of the 2018 world wide web conference*.
- ZHOU, D. and GU, Q. (2022). Computationally efficient horizon-free reinforcement learning for linear mixture mdps. *arXiv preprint arXiv:2205.11507* .
- ZHOU, D., GU, Q. and SZEPESVARI, C. (2021a). Nearly minimax optimal reinforcement learning for linear mixture markov decision processes. In *Conference on Learning Theory*. PMLR.
- ZHOU, D., HE, J. and GU, Q. (2021b). Provably efficient reinforcement learning for discounted mdps with feature mapping. In *International Conference on Machine Learning*. PMLR.
- ZHOU, R., ZIHAN, Z. and DU, S. S. (2023). Sharp variance-dependent bounds in reinforcement learning: Best of both worlds in stochastic and deterministic environments. In *International Conference on Machine Learning*. PMLR.
- ZOPH, B. and LE, Q. V. (2016). Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578* .

A Additional Related Work

A.1 RL with Linear Function Approximation

In recent years, a substantial body of research has emerged to address the challenges of solving Markov Decision Processes (MDP) with linear function approximation, particularly to handle the vast state and action spaces (Jiang et al., 2017; Dann et al., 2018; Yang and Wang, 2019; Du et al., 2019; Sun et al., 2019; Jin et al., 2020; Wang et al., 2020b; Zanette et al., 2020a; Yang and Wang, 2020; Modi et al., 2020; Ayoub et al., 2020; Zhou et al., 2021a; He et al., 2021a; Zhou and Gu, 2022; He et al., 2022; Zhao et al., 2023). These works can be broadly categorized into two groups based on the linear structures applied to the underlying MDP. One commonly employed linear structure is known as the linear MDP (Jin et al., 2020), where the transition probability function \mathbb{P}_h and reward function r_h are represented as linear functions with respect to a given feature mapping $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$. Under this assumption, the LSVI-UCB algorithm (Jin et al., 2020) has been shown to achieve a regret guarantee of $\tilde{O}(\sqrt{d^3 H^4 K})$. Subsequently, Zanette et al. (2020a) introduced the RLSVI algorithm, utilizing the Thompson sampling method, to attain a regret bound of $\tilde{O}(\sqrt{d^4 H^5 K})$. More recently, He et al. (2022) improved the regret guarantee to $\tilde{O}(\sqrt{d^2 H^3 K})$ with the LSVI-UCB++ algorithm, aligning with the theoretical lower bound in Zhou et al. (2021a) up to logarithmic factors. Another line of research has focused on linear mixture MDPs (Modi et al., 2020; Yang and Wang, 2020; Jia et al., 2020; Ayoub et al., 2020; Zhou et al., 2021a), where the transition probability is expressed as a linear combination of basic models $\mathbb{P}_1, \mathbb{P}_2, \dots, \mathbb{P}_d$. For linear mixture MDPs, Jia et al. (2020) introduced the UCRL-VTR algorithm, achieving a regret guarantee of $\tilde{O}(\sqrt{d^2 H^4 K})$. Subsequently, Zhou et al. (2021a) enhanced this result to $\tilde{O}(\sqrt{d^2 H^3 K})$, reaching a nearly minimax optimal regret bound. Recently, several works focused on time-homogeneous linear mixture MDPs, and removed the dependency on the episode length (horizon-free) (Zhang and Ji, 2019; Zhou et al., 2021b; Zhao et al., 2023).

A.2 RL with General Function Approximation

Recent years have witnessed a flurry of progress on RL with nonlinear function classes. To explore the theoretical limits of RL algorithms, various complexity measures have been developed to characterize the hardness of RL instances such as Bellman rank (Jiang et al., 2017), Witness rank (Sun et al., 2019), eluder dimension (Russo and Van Roy, 2013), Bellman eluder dimension (Jin et al., 2021), Bilinear Classes (Du et al., 2021), Decision-Estimation Coefficient (Foster et al., 2021), Admissible Bellman Characterization (Chen et al., 2022), generalized eluder dimension (Agarwal et al., 2022). Among them, only Agarwal et al. (2022) yields a near-optimal regret guarantee when specialized to linear MDPs. In their paper, they proposed a new framework named generalized eluder dimension to handle the weighted objects in weighted regression, which can be seen as a variant of eluder dimension. In their proposed algorithm VOQL, they adopted over-optimistic and over-pessimistic value functions in order to bound the variance of regression targets, making it possible to apply a weighted regression scheme in the model-free framework.

A.3 RL with Low Switching Cost

Most of the aforementioned approaches necessitate updating both the value function and the corresponding policy in each episode, a practice that proves to be inefficient when dealing with substantial datasets. To overcome this limitation, a widely adopted technique involves dividing the time sequence into several epochs and updating the policy only between different epochs. In the context of the linear bandit problem, Abbasi-Yadkori et al. (2011) introduced the rarely-switching OFUL algorithm, where the agent updates the policy only when the determinant of the covariance matrix doubles. This method enables the algorithm to achieve near-optimal regret of $\tilde{O}(\sqrt{K})$ while maintaining policy updates to $\tilde{O}(d \log K)$ times. When the number of arms, denoted as $|\mathcal{D}|$, is finite, Ruan et al. (2021) proposed an algorithm with regret bounded by $\tilde{O}(\sqrt{dK})$ and a mere $\tilde{O}(d \log d \log K)$ policy-updating times. In the realm of episodic reinforcement learning, Bai et al. (2019) and Zhang et al. (2021) delved into tabular Markov Decision Processes, introducing algorithms that achieve a regret of $\tilde{O}(\sqrt{T})$ and $\tilde{O}(SA \log K)$ updating times. Wang et al. (2021) later extended these results to linear MDPs, unveiling the LSVI-UCB-RareSwitch algorithm. LSVI-UCB-RareSwitch delivers a regret bound

of $\tilde{O}(d^3 H^4 K)$ with a policy switching frequency of $\tilde{O}(d \log T)$, which matches the lower bound of the switching cost (Gao et al., 2021) up to logarithmic terms. Furthermore, if the policy is updated within fixed-length epochs (Han et al., 2020), the method is termed batch learning model, but this falls beyond the scope of our current work. In addition, with the help of stochastic policies, Zhang et al. (2022) proposed an algorithm with $\tilde{O}(\sqrt{K})$ regret guarantee and only $\tilde{O}(H)$ switching cost for learning tabular MDPs. Later, Huang et al. (2022) employed the stochastic policy in learning linear MDPs, which is able to find an ϵ -optimal policy with only $\tilde{O}(H)$ switching cost.

B Additional Results

B.1 Lower Bound for the Switching Cost

As a complement, we prove a new lower bound on the switching cost for RL with linear MDPs. Note that linear MDPs lie in the class of MDPs studied in our paper with bounded generalized eluder dimension. In particular, the generalized eluder dimension of linear MDPs is $\tilde{O}(d)$.

Theorem B.1. For any algorithm **Alg** with expected switching cost less than $dH/(16 \log K)$, there exists a hard-to-learn linear MDP, such that the expected regret of **Alg** is at least $\Omega(K)$.

Remark B.2. Theorem B.1 suggests that, to achieve a sublinear regret guarantee, an $\tilde{\Omega}(dH)$ switching cost is inevitable. This lower bound does not violate the minimax upper bound of $\tilde{O}(H)$ proved in Zhang et al. (2022); Huang et al. (2022), which additionally assume that the initial state s_1^k is either fixed or sampled from a fixed distribution. In contrast, our work and Gao et al. (2021) allow the initial state to be **adaptively** chosen by an adversarial environment, bringing more challenges to the learning of linear MDPs. When comparing our lower bound with the result in Gao et al. (2021), it is worth noting that their focus is solely on deterministic algorithms, and they suggest that an $\tilde{\Omega}(dH)$ switching cost is necessary. As a comparison, our result holds for any algorithm with arbitrary policies including both deterministic and stochastic policies.

B.2 Connection Between $D_{\mathcal{F}}^2$ -Uncertainty and Eluder Dimension

Previous work by Agarwal et al. (2022) achieved the optimal regret bound $\tilde{O}(\sqrt{\dim(\mathcal{F}) \log N \cdot HK})$, where $\dim(\mathcal{F})$ is defined as the generalized eluder dimension as stated in Definition 2.4. However, the connection between generalized eluder dimension and eluder dimension proposed by Russo and Van Roy (2013) is still under-discussed². Consequently, their results could not be directly compared with the results based on the classic eluder dimension measure (Wang et al., 2020a) or the more general Bellman eluder dimension (Jin et al., 2021).

In this section, we make a first attempt to establish a connection between generalized eluder dimension and eluder dimension in Theorem B.3.

Theorem B.3. For a function space \mathcal{G} , $\alpha > 0$, let \dim be defined in Definition 2.4. WLOG, we assume that for all $g \in \mathcal{G}$ and $z \in \mathcal{Z}$, $|g(z)| \leq 1$. Then the following inequality between $\dim(\mathcal{F}, \mathbf{Z}, \boldsymbol{\sigma})$ and $\dim_E(\mathcal{G}, 1/\sqrt{T})$ holds for all $\mathbf{Z} := \{z_i\}_{i \in [T]}$ with $z_i \in \mathcal{Z}$ and $\boldsymbol{\sigma} := \{\sigma_i\}_{i \in [T]}$ s.t. $\alpha \leq \sigma_i \leq M \ \forall i \in [T]$:

$$\sum_{i \in [T]} \min \left(1, \frac{1}{\sigma_i} D_{\mathcal{G}}^2(z_i; z_{[i-1]}, \sigma_{[i-1]}) \right) \leq O(\dim_E(\mathcal{F}, 1/\sqrt{T}) \log T \log \lambda T \log(M/\alpha) + \lambda^{-1}).$$

According to Theorem B.3, the generalized eluder dimension is upper bounded by eluder dimension up to logarithmic terms. When the number of episodes K is sufficiently large, the leading term in our regret bound in Theorem 4.1 is $\tilde{O}(\sqrt{\dim_E(\mathcal{F}) \log \mathcal{N} \cdot HK})$, where $\dim_E(\mathcal{F})$ is the eluder dimension of the function class \mathcal{F} .

C Proof of Theorem B.3

To start with, we first recall the concept of eluder dimension as follows.

²Agarwal et al. (2022) (Remark 4) also discussed the relationship between the generalized eluder dimension and eluder dimension. However, there exists a technique flaw in the proof and we will discuss it in Appendix D.

Definition C.1 (Definition of eluder dimension, [Russo and Van Roy 2013](#)). The eluder dimension of a function class \mathcal{G} with domain \mathcal{Z} is defined as follows:

- A point $z \in \mathcal{Z}$ is ϵ -dependent on $z_1, z_2, \dots, z_k \in \mathcal{Z}$ with respect to \mathcal{G} , if for all $g_1, g_2 \in \mathcal{G}$ such that $\sum_{i=1}^k [g_1(z_i) - g_2(z_i)]^2 \leq \epsilon$, it holds that $|g_1(z) - g_2(z)| \leq \epsilon$.
- Further z is said to be ϵ -independent of z_1, z_2, \dots, z_k with respect to \mathcal{G} , if z is not dependent on z_1, z_2, \dots, z_k .
- The eluder dimension of \mathcal{G} , denoted by $\dim_E(\mathcal{G}, \epsilon)$, is the length of the longest sequence of elements in \mathcal{Z} such that every element is ϵ' -independent of its predecessors for some $\epsilon' \geq \epsilon$.

With this definition, we can prove [Theorem B.3](#).

Proof of Theorem B.3. Let \mathcal{I}_j ($1 \leq j \leq \lceil \log_2 M/\alpha \rceil$) be the index set such that

$$\mathcal{I}_j = \{t \in [T] \mid \sigma_t \in [2^{j-1} \cdot \alpha, 2^j \alpha]\}.$$

Then we focus on the summation over \mathcal{I}_j for each j . For simplicity, we denote the subsequence $\{z_i\}_{i \in \mathcal{I}_j}$ by $\{x_i\}_{i \in \llbracket \mathcal{I}_j \rrbracket}$. Then we have

$$\sum_{i \in \mathcal{I}_j} \min \left(1, \frac{1}{\sigma_i^2} D_{\mathcal{G}}^2(z_i; z_{[i-1]}, \sigma_{[i-1]}) \right) \leq \sum_{i \in \llbracket \mathcal{I}_j \rrbracket} \min(1, 4D_{\mathcal{G}}^2(x_i; x_{[i-1]}, 1_{[i-1]})).$$

To bound $\sum_{i \in \llbracket \mathcal{I}_j \rrbracket} \min(1, 4D_{\mathcal{G}}^2(x_i; x_{[i-1]}, 1_{[i-1]}))$,

$$\begin{aligned} & \sum_{i \in \llbracket \mathcal{I}_j \rrbracket} \min(1, 4D_{\mathcal{G}}^2(x_i; x_{[i-1]}, 1_{[i-1]})) \\ & \leq 4/\lambda + \sum_{i \in \llbracket \mathcal{I}_j \rrbracket} 4 \int_{1/\lambda T}^1 \mathbb{1}\{D_{\mathcal{G}}^2(x_i; x_{[i-1]}, 1_{[i-1]}) \geq \rho\} d\rho \\ & \leq 4/\lambda + 4 \int_{1/\lambda T}^1 \sum_{i \in \llbracket \mathcal{I}_j \rrbracket} \mathbb{1}\{D_{\mathcal{G}}^2(x_i; x_{[i-1]}, 1_{[i-1]}) \geq \rho\} d\rho. \end{aligned} \quad (\text{C.1})$$

Then we proceed by bounding the summation $\sum_{i \in \llbracket \mathcal{I}_j \rrbracket} \mathbb{1}\{D_{\mathcal{G}}^2(x_i; x_{[i-1]}, 1_{[i-1]}) \geq \rho\}$ for each $\rho \geq 1/(\lambda T)$. For short, let $d := \dim_E(\mathcal{G}, 1/\sqrt{T})$. Essentially, it suffices to provide an upper bound of the cardinality of the subset $\mathcal{J}_j := \{i \in \mathcal{I}_j \mid D_{\mathcal{G}}^2(x_i; x_{[i-1]}, 1_{[i-1]}) \geq \rho\}$.

For each $i \in \mathcal{J}_j$, since $D_{\mathcal{G}}^2(x_i; x_{[i-1]}, 1_{[i-1]}) \geq \rho$, there exists g_1, g_2 in \mathcal{G} such that

$$\frac{(g_1(x_i) - g_2(x_i))^2}{\sum_{t \in [i-1]} (g_1(x_t) - g_2(x_t))^2 + \lambda} \geq \rho/2.$$

Here $(g_1(x_i) - g_2(x_i))^2 \geq \lambda \rho \geq 1/T$. Denote such value of $(g_1(x_i) - g_2(x_i))^2$ by $L(x_i)$. Then we consider split $c\mathcal{J}_j$ into $\lceil \log_2 T \rceil$ layers such that in each layer \mathcal{J}_j^k ($k \in \llbracket \log_2 T \rrbracket$), we have $1/T \leq \xi \leq L(x_i) \leq 2\xi$ for some ξ .

Denote the cardinality of \mathcal{J}_j^k by A and the subsequence $\{x_i\}_{i \in \mathcal{J}_j^k}$ by $\{y_i\}_{i \in [A]}$. For the elements in $\{y\}$, we dynamically maintain $\lfloor A/d \rfloor$ queues of elements. We enumerate through $\{y\}$ in its original order and put the current element y_i into the queue Q such that y_i is ξ -independent of the current elements in Q . By the Pigeonhole principle and the definition of eluder dimension d , we can find an element y_i in $\{y\}$ such that y_i is ξ -dependent of all the $\lfloor A/d \rfloor$ queues before i .

Then by the definition of $L(y_i)$ and ξ , we can choose g_1, g_2 such that

$$\frac{(g_1(y_i) - g_2(y_i))^2}{\sum_{t \in [i-1]} (g_1(y_t) - g_2(y_t))^2 + \lambda} \geq \rho/2, \quad 2\xi \geq (g_1(y_t) - g_2(y_t))^2 \geq \xi. \quad (\text{C.2})$$

By the ξ -dependencies, we have $\sum_{t \in [i-1]} (g_1(y_t) - g_2(y_t))^2 \geq \lfloor A/d \rfloor \cdot \xi$. At the same time, $\sum_{t \in [i-1]} (g_1(y_t) - g_2(y_t))^2 \leq 4\xi/\rho$ due to (C.2). Thus, we deduce that $A = |\mathcal{J}_j^k| \leq O(d/\rho)$ for all $k \in [\lceil \log_2 T \rceil]$. Substituting it into (C.1), we have

$$\begin{aligned} \sum_{i \in [\lceil \mathcal{I}_j \rceil]} \min(1, 4D_G^2(x_i; x_{[i-1]}, \mathbf{1}_{[i-1]})) &\leq 4/\lambda + 4 \int_{1/\lambda T}^1 O(d/\rho \cdot \log_2 T) d\rho \\ &= O(d \log_2 T \cdot \log \lambda T + \lambda^{-1}). \end{aligned}$$

Here j is chosen arbitrarily. Hence, the generalized eluder dimension can be further bounded as follows:

$$\sum_{i \in [T]} \min\left(1, \frac{1}{\sigma_i^2} D_G^2(z_i; z_{[i-1]}, \sigma_{[i-1]})\right) \leq O(\dim_E(\mathcal{F}, 1/\sqrt{T}) \log T \log \lambda T \log(M/\alpha) + \lambda^{-1}).$$

□

In the following part, we will also discuss the issue in Agarwal et al. (2022) about the relationship between the standard eluder dimension \dim_E and the generalized eluder dimension \dim . In detail, Agarwal et al. (2022) proposed the concept of Generalized Eluder dimension and made the following claim:

Lemma C.2 (Remark 4, Agarwal et al. 2022). If we set the weight $\sigma = 1$, then the Generalized Eluder dimension $\dim = \sup_{\mathbf{Z}, \sigma: |\mathbf{Z}|=T} \dim(\mathcal{F}, \mathbf{Z}, \mathbf{1})$ satisfied

$$\dim \leq \dim_E(\mathcal{F}, \sqrt{\lambda/T}) + 1,$$

where \dim_E denotes the standard Eluder dimension proposed in Russo and Van Roy (2013). In the proof of Remark 4, Agarwal et al. (2022) claims that given the standard Eluder dimension $\dim_E(\mathcal{F}, \sqrt{\lambda}) = n$, there are at most n different (sorted) indices $\{t_1, t_2, \dots\}$ such that $D_{\mathcal{F}}^2(z_{t_i}; z_{t_{i-1}}, \mathbf{1}) \geq \epsilon^2/\lambda$. However, according to the definition of $D_{\mathcal{F}}^2$ -uncertainty, we only have

$$D_{\mathcal{F}}^2(z; z_{[t-1]}, \mathbf{1}) := \sup_{f_1, f_2 \in \mathcal{F}} \frac{(f_1(z) - f_2(z))^2}{\sum_{s \in [t-1]} (f_1(z_s) - f_2(z_s))^2 + \lambda}.$$

However, z_{t_i} is ϵ -dependence with $z_1, \dots, z_{t_{i-1}}$ is only a sufficient condition for the uncertainty $D_{\mathcal{F}}^2(z; z_{[t-1]}, \mathbf{1})$ rather than necessary condition. In cases where both $(f_1(z) - f_2(z))^2 \geq \epsilon^2$ and $\sum_{s \in [t-1]} (f_1(z_s) - f_2(z_s))^2 \geq \epsilon^2$ hold, the uncertainty $D_{\mathcal{F}}^2(z; z_{[t-1]}, \mathbf{1})$ may exceed ϵ^2/λ , yet it will not be counted within the longest ϵ -independent sequence for the standard Eluder dimension.

D Proof of Theorem 4.1

In this section, we provide the proof of Theorem 4.1.

D.1 High Probability Events

In this subsection, we define the following high-probability events:

$$\underline{\mathcal{E}}_{k,h}^{\tilde{f}} = \left\{ \lambda + \sum_{i \in [k-1]} \left(\tilde{f}_{k,h}(s_h^i, a_h^i) - \mathcal{T}_h^2 V_{k,h+1}(s_h^i, a_h^i) \right)^2 \leq \tilde{\beta}_k^2 \right\}, \quad (\text{D.1})$$

$$\underline{\mathcal{E}}_{k,h}^{\hat{f}} = \left\{ \lambda + \sum_{i \in [k-1]} \frac{1}{\bar{\sigma}_{i,h}^2} \left(\hat{f}_{k,h}(s_h^i, a_h^i) - \mathcal{T}_h V_{k,h+1}(s_h^i, a_h^i) \right)^2 \leq \beta_k^2 \right\}, \quad (\text{D.2})$$

$$\underline{\mathcal{E}}_{k,h}^{\check{f}} = \left\{ \lambda + \sum_{i \in [k-1]} \frac{1}{\bar{\sigma}_{i,h}^2} \left(\check{f}_{k,h}(s_h^i, a_h^i) - \mathcal{T}_h \check{V}_{k,h+1}(s_h^i, a_h^i) \right)^2 \leq \beta_k^2 \right\}, \quad (\text{D.3})$$

where

$$\tilde{\beta}_k := \sqrt{128 \log \frac{N_\epsilon(k) \cdot N_{\mathcal{F}}(\epsilon) \cdot H}{\delta} + 64L\epsilon \cdot k}, \quad \beta_k := \sqrt{128 \cdot \log \frac{N_\epsilon(k) \cdot N_{\mathcal{F}}(\epsilon)H}{\delta} + 64L\epsilon \cdot k/\alpha^2}.$$

$\underline{\mathcal{E}}_{k,h}^{\tilde{f}}$, $\underline{\mathcal{E}}_{k,h}^{\hat{f}}$ and $\underline{\mathcal{E}}_{k,h}^{\check{f}}$ are the hoeffding-type concentration results for $\tilde{f}_{k,h}$, $\hat{f}_{k,h}$ and $\check{f}_{k,h}$ respectively.

Then we define the following bellman-type concentration events for $\hat{f}_{k,h}$, $\check{f}_{k,h}$, which implies a tighter confidence set due to carefully designed variance estimators and an inverse-variance weighted regression scheme for the general function classes.

$$\begin{aligned} \mathcal{E}_h^{\hat{f}} &= \left\{ \lambda + \sum_{i=1}^{k-1} \frac{1}{(\bar{\sigma}_{i,h'})^2} \left(\hat{f}_{k,h'}(s_{h'}^i, a_{h'}^i) - \mathcal{T}_{h'} V_{k,h'+1}(s_{h'}^i, a_{h'}^i) \right)^2 \leq \hat{\beta}_k^2, \forall h \leq h' \leq H, k \in [K] \right\}, \\ \mathcal{E}_h^{\check{f}} &= \left\{ \lambda + \sum_{i=1}^{k-1} \frac{1}{(\bar{\sigma}_{i,h'})^2} \left(\check{f}_{k,h'}(s_{h'}^i, a_{h'}^i) - \mathcal{T}_{h'} V_{k,h'+1}(s_{h'}^i, a_{h'}^i) \right)^2 \leq \check{\beta}_k^2, \forall h \leq h' \leq H, k \in [K] \right\}, \end{aligned} \quad (\text{D.4})$$

where

$$\begin{aligned} \hat{\beta}_k^2 = \check{\beta}_k^2 &:= O \left(\log \frac{2k^2 (2 \log(L^2 k / \alpha^4) + 2) \cdot (\log(4L/\alpha^2) + 2)}{\delta/H} \right) \cdot [\log(N_{\mathcal{F}}(\epsilon)) + 1] \\ &+ O(\lambda) + O(\epsilon k L / \alpha^2). \end{aligned}$$

We also define the following events which are later applied to prove the concentration of $\hat{f}_{k,h}$ and $\check{f}_{k,h}$ by induction.

$$\bar{\mathcal{E}}_{k,h}^{\hat{f}} = \left\{ \lambda + \sum_{i \in [k-1]} \frac{\hat{\mathbb{1}}_{i,h}}{(\bar{\sigma}_{i,h})^2} \left(\hat{f}_{k,h}(s_h^i, a_h^i) - \mathcal{T}_h V_{k,h+1}(s_h^i, a_h^i) \right)^2 \leq \hat{\beta}_k^2 \right\}, \quad (\text{D.5})$$

where $\hat{\mathbb{1}}_{i,h}$ is the shorthand for the following product of indicator functions

$$\begin{aligned} \hat{\mathbb{1}}_{i,h} &:= \mathbf{1} \left([\mathbb{V}_h V_{h+1}^*](s_h^i, a_h^i) \leq \bar{\sigma}_{i,h}^2 \right) \cdot \mathbf{1} \left([\mathcal{T}_h V_{i,h+1} - \mathcal{T} V_{h+1}^*] \leq (\log N_{\mathcal{F}}(\epsilon) + \log N_\epsilon(K))^{-1} \bar{\sigma}_{i,h}^2 \right) \\ &\cdot \mathbf{1} \left(V_{i,h+1}(s) \geq V_{h+1}^*(s) \quad \forall s \in \mathcal{S} \right). \end{aligned} \quad (\text{D.6})$$

$$\bar{\mathcal{E}}_{k,h}^{\check{f}} = \left\{ \lambda + \sum_{i \in [k-1]} \frac{\check{\mathbb{1}}_{i,h}}{(\bar{\sigma}_{i,h})^2} \left(\check{f}_{k,h}(s_h^i, a_h^i) - \mathcal{T}_h \check{V}_{k,h+1}(s_h^i, a_h^i) \right)^2 \leq \check{\beta}_k^2 \right\}, \quad (\text{D.7})$$

where $\check{\mathbb{1}}_{i,h}$ is the shorthand for the following product of indicator functions

$$\begin{aligned} \check{\mathbb{1}}_{i,h} &:= \mathbf{1} \left([\mathbb{V}_h V_{h+1}^*](s_h^i, a_h^i) \leq \bar{\sigma}_{i,h}^2 \right) \cdot \mathbf{1} \left([\mathcal{T}_h V_{h+1}^* - \mathcal{T}_h \check{V}_{i,h+1}] \leq (\log N_{\mathcal{F}}(\epsilon) + \log N_\epsilon(K))^{-1} \bar{\sigma}_{i,h}^2 \right) \\ &\cdot \mathbf{1} \left(\check{V}_{i,h+1}(s) \leq V_{h+1}^*(s) \quad \forall s \in \mathcal{S} \right). \end{aligned} \quad (\text{D.8})$$

The following Lemmas suggest that previous events hold with high probability.

Lemma D.1. Let $\tilde{f}_{k,h}$ be defined as in line 8 of Algorithm 1, we have $\underline{\mathcal{E}}^{\tilde{f}} := \bigcap_{k \geq 1, h \in [H]} \underline{\mathcal{E}}_{k,h}^{\tilde{f}}$ holds with probability at least $1 - \delta$, where $\underline{\mathcal{E}}_{k,h}^{\tilde{f}}$ is defined in (D.1).

Lemma D.2. Let $\hat{f}_{k,h}$ be defined as in line 6 of Algorithm 1, we have $\underline{\mathcal{E}}^{\hat{f}} := \bigcap_{k \geq 1, h \in [H]} \underline{\mathcal{E}}_{k,h}^{\hat{f}}$ holds with probability at least $1 - \delta$, where $\underline{\mathcal{E}}_{k,h}^{\hat{f}}$ is defined in (D.2).

Lemma D.3. Let $\check{f}_{k,h}$ be defined as in line 7 of Algorithm 1, we have $\underline{\mathcal{E}}^{\check{f}} := \bigcap_{k \geq 1, h \in [H]} \underline{\mathcal{E}}_{k,h}^{\check{f}}$ holds with probability at least $1 - \delta$, where $\underline{\mathcal{E}}_{k,h}^{\check{f}}$ is defined in (D.3).

Lemma D.4. Let $\hat{f}_{k,h}$ be defined as in line 6 of Algorithm 1, we have $\bar{\mathcal{E}}^{\hat{f}} := \bigcap_{k \geq 1, h \in [H]} \bar{\mathcal{E}}_{k,h}^{\hat{f}}$ holds with probability at least $1 - 2\delta$, where $\bar{\mathcal{E}}_{k,h}^{\hat{f}}$ is defined in (D.5).

Lemma D.5. Let $\check{f}_{k,h}$ be defined as in line 7 of Algorithm 1, we have $\bar{\mathcal{E}}^{\check{f}}$ holds with probability at least $1 - 2\delta$.

D.2 Optimism and Pessimism

Based on the high probability events, we have the following lemmas for the function $\widehat{f}_{k,h}(s, a)$ and $\check{f}_{k,h}(s, a)$

Lemma D.6. On the events $\mathcal{E}_h^{\widehat{f}}$ and $\mathcal{E}_h^{\check{f}}$, for each episode $k \in [K]$, we have

$$\begin{aligned} |\widehat{f}_{k,h}(s, a) - \mathcal{T}_h V_{k,h+1}(s, a)| &\leq \widehat{\beta}_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), \\ |\check{f}_{k,h}(s, a) - \mathcal{T}_h \check{V}_{k,h+1}(s, a)| &\leq \check{\beta}_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), \end{aligned}$$

where $z = (s, a)$ and $z_{[k-1],h} = \{z_{1,h}, z_{2,h}, \dots, z_{k-1,h}\}$.

Lemma D.7. On the events $\mathcal{E}_{h+1}^{\widehat{f}}$ and $\mathcal{E}_{h+1}^{\check{f}}$, for each stage $h \leq h' \leq H$ and episode $k \in [K]$, we have $Q_{k,h}(s, a) \geq Q_h^*(s, a) \geq \check{Q}_{k,h}(s, a)$. Furthermore, for the value functions $V_{k,h}(s)$ and $\check{V}_{k,h}(s)$, we have $V_{k,h}(s) \geq V_h^*(s) \geq \check{V}_{k,h}(s)$.

D.3 Monotonic Variance Estimator

In this subsection, we introduce the following lemmas to construct the monotonic variance estimator.

Lemma D.8. On the events $\underline{\mathcal{E}}^{\check{f}}$, $\underline{\mathcal{E}}^{\widehat{f}}$ and $\underline{\mathcal{E}}^{\check{f}}$, for each episode $k \in [K]$ and stage $h \in [H]$, we have

$$\begin{aligned} |\widehat{f}_{k,h}(s, a) - \mathcal{T}_h V_{k,h+1}(s, a)| &\leq \beta_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), \\ |\check{f}_{k,h}(s, a) - \mathcal{T}_h \check{V}_{k,h+1}(s, a)| &\leq \beta_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), \\ |\widetilde{f}_{k,h}(s, a) - \mathcal{T}_h^2 V_{k,h+1}(s, a)| &\leq \widetilde{\beta}_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), \end{aligned}$$

where $z = (s, a)$ and $z_{[k-1],h} = \{z_{1,h}, z_{2,h}, \dots, z_{k-1,h}\}$.

Lemma D.9. On the events $\underline{\mathcal{E}}^{\check{f}}$, $\underline{\mathcal{E}}^{\widehat{f}}$, $\underline{\mathcal{E}}^{\check{f}}$, $\mathcal{E}_{h+1}^{\widehat{f}}$, $\mathcal{E}_{h+1}^{\check{f}}$, for each episode $k \in [K]$, the empirical variance $[\widetilde{\mathbb{V}}_h V_{k,h+1}](s_h^k, a_h^k)$ satisfies the following inequalities:

$$\begin{aligned} |[\widetilde{\mathbb{V}}_h V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{V}_h V_{k,h+1}](s_h^k, a_h^k)| &\leq E_{k,h}, \\ |[\widetilde{\mathbb{V}}_h V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{V}_h V_{h+1}^*](s_h^k, a_h^k)| &\leq E_{k,h} + F_{k,h}. \end{aligned}$$

Lemma D.10. On the events $\underline{\mathcal{E}}^{\check{f}}$, $\underline{\mathcal{E}}^{\widehat{f}}$, $\mathcal{E}_{h+1}^{\widehat{f}}$, $\mathcal{E}_{h+1}^{\check{f}}$, for each episode $k \in [K]$ and $i > k$, we have

$$\begin{aligned} (\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)) \cdot [\mathbb{V}_h(V_{i,h+1} - V_{h+1}^*)](s_h^k, a_h^k) &\leq F_{i,h}, \\ (\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)) \cdot [\mathbb{V}_h(V_{h+1}^* - \check{V}_{i,h+1})](s_h^k, a_h^k) &\leq F_{i,h}, \end{aligned}$$

Based on previous lemmas, we can construct an optimistic estimator $\sigma_{k,h}$ for the transition variance. Under this situation, the weighted regression have the following guarantee.

Lemma D.11. If the events $\underline{\mathcal{E}}^{\check{f}}$, $\underline{\mathcal{E}}^{\widehat{f}}$, $\underline{\mathcal{E}}^{\check{f}}$, $\bar{\mathcal{E}}^{\widehat{f}}$ and $\bar{\mathcal{E}}^{\check{f}}$ hold, then events $\mathcal{E}^{\widehat{f}} = \mathcal{E}_1^{\widehat{f}}$ and $\mathcal{E}^{\check{f}} = \mathcal{E}_1^{\check{f}}$ hold.

D.4 Proof of Regret Bound

In the subsection, we first define the following high probability events to control the stochastic noise from the transition process:

$$\begin{aligned} \mathcal{E}_1 &= \left\{ \forall h' \in [H], \sum_{k=1}^K \sum_{h=h'}^H [\mathbb{P}_h(V_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) - \sum_{k=1}^K \sum_{h=h'}^H (V_{k,h+1}(s_{h+1}^k) - V_{h+1}^{\pi^k}(s_{h+1}^k)) \right. \\ &\leq 2 \sqrt{\sum_{k=1}^K \sum_{h=h'}^H [\mathbb{V}_h(V_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) \log(2K^2 H/\delta)} + 2\sqrt{\log(2K^2 H/\delta)} + 2\log(2K^2 H/\delta) \left. \right\}, \\ \mathcal{E}_2 &= \left\{ \forall h' \in [H], \sum_{k=1}^K \sum_{h=h'}^H [\mathbb{P}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) - \sum_{k=1}^K \sum_{h=h'}^H (V_{k,h+1}(s_{h+1}^k) - \check{V}_{k,h+1}(s_{h+1}^k)) \right. \end{aligned}$$

$$\leq 2 \sqrt{\sum_{k=1}^K \sum_{h'=h}^H [\mathbb{V}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \log(2K^2 H/\delta) + 2\sqrt{\log(2K^2 H/\delta)} + 2\log(2K^2 H/\delta)}.$$

According to Freedman inequality (Lemma H.5), we directly have the following results.

Lemma D.12. Events \mathcal{E}_1 and \mathcal{E}_2 hold with probability at least $1 - 2\delta$.

Proof. Fix an arbitrary $h \in [H]$. Applying Lemma H.5, we have the following inequality:

$$\begin{aligned} & \sum_{k=1}^K \sum_{h'=h}^H \mathbb{P}_h(V_{k,h+1} - V_{h+1}^{\pi^k})(s_h^k, a_h^k) - \sum_{k=1}^K \sum_{h'=h}^H (V_{k,h+1}(s_{h+1}^k) - V_{h+1}^{\pi^k}(s_{h+1}^k)) \\ & \leq 2 \sqrt{\sum_{k=1}^K \sum_{h'=h}^H [\mathbb{V}_h(V_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) \log(2K^2 H/\delta) + 2\sqrt{\log(2K^2 H/\delta)} + 2\log(2K^2 H/\delta)} \end{aligned} \quad (\text{D.9})$$

Applying a union bound for (D.9) across all $h \in [H], k \geq 0$, we have Event \mathcal{E}_1 holds with probability at least $1 - 2\delta$.

Similarly, we also have the corresponding high-probability bound for \mathcal{E}_2 . \square

Next, we need the following lemma to control the summation of confidence radius.

Lemma D.13. For any parameters $\beta \geq 1$ and stage $h \in [H]$, the summation of confidence radius over episode $k \in [K]$ is upper bounded by

$$\begin{aligned} & \sum_{k=1}^K \min\left(\beta D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), 1\right) \\ & \leq (1 + \beta\gamma^2) \dim_{\alpha,K}(\mathcal{F}_h) + 2\beta \sqrt{\dim_{\alpha,K}(\mathcal{F}_h)} \sqrt{\sum_{k=1}^K (\sigma_{k,h}^2 + \alpha^2)}, \end{aligned}$$

where $z = (s, a)$ and $z_{[k-1],h} = \{z_{1,h}, z_{2,h}, \dots, z_{k-1,h}\}$.

Then, we can decompose the total regret in the first K episode to the summation of variance $\sum_{k=1}^K \sum_{h=1}^H \sigma_{k,h}^2$ as the following lemma.

Lemma D.14. On the events $\mathcal{E}^{\hat{f}} = \mathcal{E}_1^{\hat{f}}$, $\mathcal{E}^{\check{f}} = \mathcal{E}_1^{\check{f}}$ and \mathcal{E}_1 , for each stage $h \in [H]$, the regret in the first K episodes can be decomposed and controlled as:

$$\begin{aligned} & \sum_{k=1}^K (V_{k,h}(s_h^k) - V_h^{\pi^k}(s_h^k)) \leq 2CH(1 + \chi)(1 + \hat{\beta}_k \gamma^2) \dim_{\alpha,K}(\mathcal{F}) \\ & + 4C(1 + \chi) \hat{\beta}_k \sqrt{\dim_{\alpha,K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)} \\ & + 2 \sqrt{\sum_{k=1}^K \sum_{h'=h}^H [\mathbb{V}_h(V_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) \log(2K^2 H/\delta) + 2\sqrt{\log(2K^2 H/\delta)} + 2\log(2K^2 H/\delta)} \end{aligned}$$

and for all stage $h \in [H]$, we further have

$$\begin{aligned} & \sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h(V_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) \\ & \leq 2CH^2(1 + \chi)(1 + \hat{\beta}_k \gamma^2) \dim_{\alpha,K}(\mathcal{F}) + 4CH(1 + \chi) \hat{\beta}_k \sqrt{\dim_{\alpha,K}(\mathcal{F}_h)} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)} \end{aligned}$$

$$+ H \left(2 \sqrt{\sum_{k=1}^K \sum_{h=1}^H [\mathbb{V}_h(V_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) \log(2K^2H/\delta) + 2\sqrt{\log(2K^2H/\delta)} + 2\log(2K^2H/\delta)} \right).$$

In addition, the gap between the optimistic value function $V_{k,h}(s)$ and pessimistic value function $\check{V}_{k,h}(s)$ can be upper bounded by the following lemma.

Lemma D.15. On the events $\mathcal{E}^{\hat{f}} = \mathcal{E}_1^{\hat{f}}$, $\mathcal{E}^{\check{f}} = \mathcal{E}_1^{\check{f}}$ and \mathcal{E}_2 , for each stage $h \in [H]$, the regret in the first K episodes can be decomposed and controlled as:

$$\begin{aligned} \sum_{k=1}^K (V_{k,h}(s_h^k) - \check{V}_{k,h}(s_h^k)) &\leq 4CH(1+\chi)(1 + \hat{\beta}_k \gamma^2) \dim_{\alpha,K}(\mathcal{F}_h) & \text{(D.10)} \\ &+ 8C(1+\chi)\hat{\beta}_k \sqrt{\dim_{\alpha,K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)} \\ &+ 2 \sqrt{\sum_{k=1}^K \sum_{h'=h}^H [\mathbb{V}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \log(2K^2H/\delta) + 2\sqrt{\log(2K^2H/\delta)} + 2\log(2K^2H/\delta)} \end{aligned}$$

and for all stage $h \in [H]$, we further have

$$\begin{aligned} &\sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \\ &\leq 4CH^2(1+\chi)(1 + \hat{\beta}_k \gamma^2) \dim_{\alpha,K}(\mathcal{F}_h) + 8CH(1+\chi)\hat{\beta}_k \sqrt{\dim_{\alpha,K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)} \\ &+ 2H \left(2 \sqrt{\sum_{k=1}^K \sum_{h'=h}^H [\mathbb{V}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \log(2K^2H/\delta) + 2\sqrt{\log(2K^2H/\delta)} + 2\log(2K^2H/\delta)} \right). \end{aligned}$$

We define the following high probability event \mathcal{E}_3 to control the summation of variance.

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

According to [Jin et al. \(2018\)](#) (Lemma C.5)³, with probability at least $1 - \delta$. Condition on this event, the summation of variance $\sum_{k=1}^K \sum_{h=1}^H \sigma_{k,h}^2$ can be upper bounded by the following lemma.

Lemma D.16. On the events $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3, \mathcal{E}^{\hat{f}} = \mathcal{E}_1^{\hat{f}}$ and $\mathcal{E}^{\check{f}} = \mathcal{E}_1^{\check{f}}$, the total estimated variance is upper bounded by:

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \sigma_{k,h}^2 &\leq (\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)) \times \tilde{O}((1 + \gamma^2)(\beta_k + H\hat{\beta}_k + \tilde{\beta}_k)H \dim_{\alpha,K}(\mathcal{F})) \\ &+ (\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K))^2 \times \tilde{O}((\beta_k + H\hat{\beta}_k + \tilde{\beta}_k)^2 H \dim_{\alpha,K}(\mathcal{F})) \\ &+ \tilde{O}(\text{Var}_K + KH\alpha^2). \end{aligned}$$

With all previous lemmas, we can prove Theorem 4.1

Proof of Theorem 4.1. The low switching cost result is given by Lemma G.1.

³Jin et al. (2018) showed that $\sum_{k=1}^K \sum_{h=1}^H [\mathbb{V}_h V_{h+1}^{\pi^k}](s_h^k, a_h^k) = \tilde{O}(KH^2 + H^3)$ when $\sum h = 1^H r_h(s_h, a_h) \leq H$. In this work, we assume the total reward satisfied $\sum_{h=1}^H r_h(s_h, a_h) \leq 1$, and the summation of variance is upper bounded by $\tilde{O}(K + H)$

After taking a union bound, the high probability events $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3, \underline{\mathcal{E}}^{\hat{f}}, \underline{\mathcal{E}}^{\tilde{f}}, \underline{\mathcal{E}}^{\tilde{f}}, \bar{\mathcal{E}}^{\hat{f}}$ and $\bar{\mathcal{E}}^{\tilde{f}}$ hold with probability at least $1 - 10\delta$. Conditioned on these events, the regret is upper bounded by

$$\begin{aligned}
& \text{Regret}(K) \\
&= \sum_{k=1}^K (V_1^*(s_1^k) - V_{k,1}^{\pi^k}(s_1^k)) \\
&\leq \sum_{k=1}^K (V_{k,1}(s_1^k) - V_{k,1}^{\pi^k}(s_1^k)) \\
&\leq 2CH(1 + \chi)(1 + \hat{\beta}_K \gamma^2) \dim_{\alpha,K}(\mathcal{F}) + \tilde{O} \left(\sqrt{\sum_{k=1}^K \sum_{h=1}^H [\mathbb{V}_h(V_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k)} \right) \\
&\quad + 4C(1 + \chi) \hat{\beta}_K \sqrt{\dim_{\alpha,K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)} \\
&\leq 2CH(1 + \chi)(1 + \hat{\beta}_K \gamma^2) \dim_{\alpha,K}(\mathcal{F}) + \tilde{O} \left(\sqrt{\sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h(V_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k)} \right) \\
&\quad + 4C(1 + \chi) \hat{\beta}_K \sqrt{\dim_{\alpha,K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)} \\
&\leq 2CH(1 + \chi)(1 + \hat{\beta}_K \gamma^2) \dim_{\alpha,K}(\mathcal{F}) + \tilde{O} \left(\hat{\beta}_K \sqrt{\dim_{\alpha,K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)} \right) \\
&\leq \tilde{O} \left(H \dim_{\alpha,K}(\mathcal{F}) \cdot (1 + \hat{\beta}_K \gamma^2) \right) + \tilde{O} \left(\hat{\beta}_K \sqrt{\dim_{\alpha,K}(\mathcal{F})} \cdot \sqrt{H \text{Var}_K} \right) \\
&\quad + \tilde{O} \left(H \dim_{\alpha,K}(\mathcal{F}) \hat{\beta}_K (\beta_K + H \hat{\beta}_K + \tilde{\beta}_K) \cdot \log[N_{\mathcal{F}}(\epsilon) \cdot N_{\epsilon}(K)] \right) + \tilde{O} \left(H \dim_{\alpha,K}(\mathcal{F}) \hat{\beta}_K (1 + \gamma^2) \right) \\
&= \tilde{O} \left(H^3 \dim_{\alpha,K}^2(\mathcal{F}) \cdot \log N_{\mathcal{F}} \left(\frac{1}{2KHL} \right) \cdot \log \left[N_{\mathcal{F}} \left(\frac{1}{2KHL} \right) \cdot N \left(\mathcal{B}, \frac{1}{2KHL \hat{\beta}_K} \right) \right] \right) \\
&\quad + \tilde{O} \left(H^{2.5} \dim_{\alpha,K}^{2.5}(\mathcal{F}) \cdot \sqrt{\log N_{\mathcal{F}} \left(\frac{1}{2KHL} \right)} \cdot \log^{1.5} \left[N_{\mathcal{F}} \left(\frac{1}{2KHL} \right) \cdot N \left(\mathcal{B}, \frac{1}{2KHL \hat{\beta}_K} \right) \right] \right) \\
&\quad + \tilde{O} \left(\sqrt{\dim_{\alpha,K}(\mathcal{F}) \log N_{\mathcal{F}} \left(\frac{1}{2KHL} \right)} \cdot H \text{Var}_K \right), \tag{D.11}
\end{aligned}$$

where the first inequality holds due to Lemma D.7, the second inequality holds due to Lemma D.14, the third inequality follows from Lemma D.7 and $\text{Var}[X] \leq \mathbb{E}[X] \cdot M$ if random variable X is in $[0, M]$, the fourth inequality holds due to (F.41) and $2ab \leq a^2 + b^2$, the last inequality holds due to Lemma D.16. Thus, we complete the proof of Theorem 4.1.

We can reorganize (D.11) into the following upper bound for regret,

$$\begin{aligned}
\text{Regret}(K) &= \tilde{O}(\sqrt{\dim(\mathcal{F}) \log N} \cdot H \text{Var}_K) \\
&\quad + \tilde{O} \left(H^{2.5} \dim^2(\mathcal{F}) \sqrt{\log N} \log(N \cdot N_b) \cdot \sqrt{H \log N + \dim(\mathcal{F}) \log(N \cdot N_b)} \right),
\end{aligned}$$

where we denote the covering number of bonus function class $N \left(\mathcal{B}, \frac{1}{2KHL \hat{\beta}_K} \right)$ by N_b , the covering number of function class \mathcal{F} by N and the dimension $\dim_{\alpha,K}(\mathcal{F})$ by $\dim(\mathcal{F})$. Since \mathcal{E}_3 occurs with high probability according to Jin et al. (2018) (Lemma C.5), we have $\text{Var}_K = \tilde{O}(K)$, which matches the worst-case minimax optimal regret bound of MDPs with general function approximation. \square

D.5 Proof Sketch of Regret Bound

In this subsection, we provide a proof sketch on the regret bound and show how our policy-switching strategy works.

First, we introduce the following lemma which illustrates the stability of the bonus function after using the uncertainty-based policy-switching condition.

Lemma D.17 (Restatement of Lemma G.2). If the policy is not updated at episode k , the uncertainty of all state-action pair $z = (s, a) \in \mathcal{S} \times \mathcal{A}$ and stage $h \in [H]$ satisfies the following stability property:

$$D_{\mathcal{F}_h}^2(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) \geq \frac{1}{1+\chi} D_{\mathcal{F}_h}^2(z; z_{[k_{last}-1],h}, \bar{\sigma}_{[k_{last}-1],h}).$$

With Lemma G.2, we can then convert $D_{\mathcal{F}_h}^2(z; z_{[k_{last}-1],h}, \bar{\sigma}_{[k_{last}-1],h})$ in the bonus term into $D_{\mathcal{F}_h}^2(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h})$ with a cost of an additional $1 + \chi$ constant.

As a direct consequence of Lemma G.2, We have the following lemma, which controls the gap between optimistic value functions and real value functions resulting from the exploration bonuses.

Lemma D.18 (Restatement of Lemma D.14). On the events $\mathcal{E}^{\hat{f}} = \mathcal{E}_1^{\hat{f}}$, $\mathcal{E}^{\check{f}} = \mathcal{E}_1^{\check{f}}$ and \mathcal{E}_1 , for each stage $h \in [H]$, the regret in the first K episodes can be decomposed and controlled as:

$$\begin{aligned} \sum_{k=1}^K (V_{k,h}(s_h^k) - V_h^{\pi^k}(s_h^k)) &\leq 2CH(1+\chi)(1+\hat{\beta}_k\gamma^2) \dim_{\alpha,K}(\mathcal{F}) \\ &+ 4C(1+\chi)\hat{\beta}_k \sqrt{\dim_{\alpha,K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)} \\ &+ 2\sqrt{\sum_{k=1}^K \sum_{h'=h}^H [\mathbb{V}_h(V_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) \log(2K^2H/\delta) + 2\sqrt{\log(2K^2H/\delta)} + 2\log(2K^2H/\delta)} \end{aligned}$$

and for all stage $h \in [H]$, we further have

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h(V_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) &\leq 2CH^2(1+\chi)(1+\hat{\beta}_k\gamma^2) \dim_{\alpha,K}(\mathcal{F}) + 4CH(1+\chi)\hat{\beta}_k \sqrt{\dim_{\alpha,K}(\mathcal{F}_h)} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)} \\ &+ H \left(2\sqrt{\sum_{k=1}^K \sum_{h=1}^H [\mathbb{V}_h(V_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) \log(2K^2H/\delta) + 2\sqrt{\log(2K^2H/\delta)} + 2\log(2K^2H/\delta)} \right). \end{aligned}$$

As a need to bound the sum of bonuses corresponding to a weighted regression target, we derive the following results on the sum of inverse weights.

Lemma D.19 (Informal version of Lemma D.16). With high probability, the total estimated variance is upper bounded by:

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \sigma_{k,h}^2 &\leq (\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)) \times \tilde{O}((1+\gamma^2)(\beta_k + H\hat{\beta}_k + \tilde{\beta}_k)H \dim_{\alpha,K}(\mathcal{F})) \\ &+ (\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K))^2 \times \tilde{O}((\beta_k + H\hat{\beta}_k + \tilde{\beta}_k)^2 H \dim_{\alpha,K}(\mathcal{F})) \\ &+ \tilde{O}(\text{Var}_K + KH\alpha^2). \end{aligned}$$

Proof sketch of Theorem 4.1. As the result of optimism, we have shown that with high probability, for all $k, h \in [K] \times [H]$, $V_{k,h}$ is an upper bound for V_h^* . Hence, we further have,

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

To further proceed, we apply Lemma D.14 and obtain

$$\begin{aligned} \text{Regret}(K) &\leq 2CH(1 + \chi)(1 + \widehat{\beta}_K \gamma^2) \dim_{\alpha, K}(\mathcal{F}) + \widetilde{O} \left(\sqrt{\sum_{k=1}^K \sum_{h=1}^H [\mathbb{V}_h(V_{k, h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k)} \right) \\ &\quad + 4C(1 + \chi) \widehat{\beta}_K \sqrt{\dim_{\alpha, K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k, h}^2 + \alpha^2)}, \end{aligned}$$

where the second term can be bounded by applying Lemma D.14 repeatedly and the third term can be controlled by Lemma D.16.

Then after substituting the value of $\widehat{\beta}_K, \gamma$ and α , we conclude that with high probability,

$$\begin{aligned} \text{Regret}(K) &= \widetilde{O}(\sqrt{\dim(\mathcal{F}) \log N \cdot H \text{Var}_K}) \\ &\quad + \widetilde{O} \left(H^{2.5} \dim^2(\mathcal{F}) \sqrt{\log N \log(N \cdot N_b)} \cdot \sqrt{H \log N + \dim(\mathcal{F}) \log(N \cdot N_b)} \right). \end{aligned}$$

Please refer to Subsection D.4 for the detailed calculation of this part. \square

E Proof of Theorem B.1

In this section, we provide the proof of Theorem B.1. To prove the lower bound, we create a series of hard-to-learn MDPs as follows. Each hard-to-learn MDP comprises $d/4$ distinct sub-MDPs denoted as $\mathcal{M}_1, \dots, \mathcal{M}_{d/4}$. Each sub-MDP \mathcal{M}_i is characterized by two distinct states, initial state $s_{i,0}$ and absorbing state $s_{i,1}$, and shares the same action set $\mathcal{A} = \{a_0, a_1\}$. Since the state and action spaces are finite, these tabular MDPs can always be represented as linear MDPs with dimension $|\mathcal{S}| \times |\mathcal{A}| = d$.

To generate each sub-MDP \mathcal{M}_i , for all stage $h \in [H]$, a special action $a_{i,h}$ is uniformly randomly selected from the action set $\{a_0, a_1\}$. Given the current state $s_{i,0}$, the agent transitions to the state $s_{i,0}$ if it takes the special action $a_{i,h}$. Otherwise, the agent transitions to the absorbing state $s_{i,1}$ and remains in that state in subsequent stages. The agent will receive the reward 1 if it takes the special action $a_{i,H}$ at the state $s_{i,0}$ during the last stage H . Otherwise, the agent always receives reward 0. In this scenario, for sub-MDP \mathcal{M}_i , the optimal policy entails following the special action sequence $(a_{i,1}, a_{i,2}, \dots, a_{i,H})$ to achieve a total reward of 1. In contrast, any other action sequence fails to yield any reward.

Now, we partition the K episodes to $d/4$ different distinct epochs. For each epoch (ranging from episodes $4(i-1)K/d + 1$ to episode $4iK/d$), we initialize the state as $s_{i,0}$ and exclusively focus on the sub-MDP \mathcal{M}_i . The regret in each epoch can be lower bounded separately as follows:

Lemma E.1. For each epoch $i \in [d/4]$ and any algorithm **Alg** capable of deploying arbitrary policies, if the expected switching cost in epoch i is less than $H/(2 \log K)$, the expected regret of **Alg** in the i -th epoch is at least $\Omega(K/d)$.

Proof of Lemma E.1. Given that each sub-MDP is independently generated, policy updates before epoch i only offer information for the sub-MDPs $\mathcal{M}_1, \dots, \mathcal{M}_{i-1}$ and do not provide any information for the current epoch i . In this scenario, there is no distinction between epochs and for simplicity, we only focus on the first epoch, encompassing episodes 1 to $4K/d$.

Now, let $k_0 = 0$ and we denote $\mathcal{K} = \{k_1, k_2, \dots\}$ as the set of episodes where the algorithm **Alg** updates the policy. If **Alg** does not update the policy i times, we set $k_i = 4K/d + 1$. For simplicity, we set $C = 2 \log K$ and for each $i \leq H/C$, we define the events \mathcal{E} as the algorithm **Alg** has not reached the state $s_{1,0}$ at the stage iC before the episode k_i .

Conditioned on the events $\mathcal{E}_1, \dots, \mathcal{E}_{i-1}$, the algorithm **Alg** does not gather any information about the special action $a_{1,h}$ for stage $h \geq (i-1)C + 1$. In this scenario, the special actions can still be considered as uniformly randomly selected from the action set a_0, a_1 . For each episode between k_{i-1} and k_i , the probability that a policy π arrives at state $s_{1,0}$ at stage iC is upper-bounded by:

$$\mathbb{E}_{a_{1,(i-1)C+1}, \dots, a_{1,iC}} \left[\prod_{h=1}^{iC} \Pr(\pi_h(s_{1,0}) = a_{1,h}) \right] \leq \mathbb{E}_{a_{1,(i-1)C+1}, \dots, a_{1,iC}} \left[\prod_{h=(i-1)C+1}^{iC} \Pr(\pi_h(s_{1,0}) = a_{1,h}) \right] = \frac{1}{2^C},$$

where the first inequality holds due to $\pi(s_{1,0}) = a_{1,h} \leq 1$ and the second equation holds due to the random generation process of the special actions. Notice that there are at most K episodes between the k_{i-1} and k_i and after applying an union bound for all episodes, we have

$$\Pr(\mathcal{E}_i | \mathcal{E}_1, \dots, \mathcal{E}_{i-1}) \leq 1 - \frac{K}{2C} = 1 - \frac{1}{K}.$$

Furthermore, we have

$$\Pr(\mathcal{E}_{H/C}) \leq \Pr(\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3 \cap \dots \cap \mathcal{E}_{H/C}) = \prod_{i=1}^{H/C} \Pr(\mathcal{E}_i | \mathcal{E}_1, \dots, \mathcal{E}_{i-1}) \leq (1 - \frac{1}{K})^{H/C} \leq 1 - \frac{H}{CK}. \quad (\text{E.1})$$

Notice that the agent cannot receive any reward unless it has reached the state $s_{i,0}$ at the last stage H . Therefore, the expected regret for the algorithm **Alg** for the sub-MDP \mathcal{M} can be bounded by the switching cost δ :

$$\begin{aligned} \mathbb{E}_{\mathcal{M}}[\text{Regret}(\mathbf{Alg})] &\geq \mathbb{E}_{\mathcal{M}}[\mathbb{1}(\mathcal{E}_{H/C}) \times (k_{H/C} - 1)] \\ &\geq \mathbb{E}_{\mathcal{M}}[k_{H/C} - 1] - \frac{H}{C} \\ &\geq \mathbb{E}_{\mathcal{M}}[\mathbb{1}(\delta < H/C) \cdot K/d] - \frac{H}{C} \\ &\geq 4K/d - \mathbb{E}_{\mathcal{M}}[\delta] \cdot \frac{4KC}{dH} - \frac{H}{C}, \end{aligned} \quad (\text{E.2})$$

where the first inequality holds due to the fact that the agent receives no reward before $k_{H/C}$ conditioned on the event $\mathcal{E}_{H/C}$, the second inequality holds due to (E.1) with $k_{H/C} - 1 \leq K$, the third inequality holds due to the definition of $k_{H/C}$ and the last inequality holds due to $\mathbb{E}[\mathbb{1}(x \geq a)] \leq \mathbb{E}[x]/a$ for any non-negative random variable x . According to the result in (E.1), if the expected switching $\mathbb{E}[\delta] \leq H/(2C)$, the expected regret is at least $\Omega(K)$, when K is large enough compared to H . \square

With Lemma B.1, we can prove Theorem B.1.

Proof of Theorem B.1. For these constructed hard-to-learn MDPs and any given algorithm **Alg**, we denote the expected switching cost for sub-MDP \mathcal{M}_i as δ_i .

According to Lemma E.1, we have

$$\begin{aligned} \mathbb{E}_{\mathcal{M}_i}[\text{Regret}(\mathbf{Alg})] &\geq \Pr(\delta_i < H/(2 \log K)) \cdot K/d \\ &\geq (1 - \mathbb{E}_{\mathcal{M}_i}[\delta_i] \cdot 2 \log K/H) \cdot K/d, \end{aligned} \quad (\text{E.3})$$

where the first inequality holds due to lemma B.1 and last inequality holds due to $\mathbb{E}[\mathbb{1}(x \geq a)] \leq \mathbb{E}[x]/a$ for any non-negative random variable x . Taking a summation of (E.3) for all sub-Mdps, we have

$$\begin{aligned} \mathbb{E}_{\mathcal{M}}[\text{Regret}(\mathbf{Alg})] &= \sum_{i=1}^{d/4} \mathbb{E}_{\mathcal{M}_i}[\text{Regret}(\mathbf{Alg})] \\ &\geq \sum_{i=1}^{d/4} (1 - \mathbb{E}_{\mathcal{M}_i}[\delta_i] \cdot 2 \log K/H) \cdot K/d \\ &= (d/4 - \mathbb{E}_{\mathcal{M}}[\delta] \cdot 2 \log K/H) \cdot K/d, \end{aligned}$$

where δ is the total expected switching cost. Therefore, for any algorithm with total expected switching cost less than $dH/(16 \log K)$, the expected is lower bounded by

$$\mathbb{E}_{\mathcal{M}}[\text{Regret}(\mathbf{Alg})] \geq (d/4 - \mathbb{E}_{\mathcal{M}}[\delta] \cdot 2 \log K/H) \cdot K/d \geq K/8.$$

Thus, we finish the proof of Theorem B.1. \square

F Proof of Lemmas in Appendix D

In this section, we provide the detailed proof of lemmas in Appendix D.

F.1 Proof of High Probability Events

F.1.1 Proof of Lemma D.1

Proof of Lemma D.1. We first prove that for an arbitrarily chosen $h \in [H]$, $\underline{\mathcal{E}}_{k,h}^{\tilde{f}}$ holds with probability at least $1 - \delta/H$ for all k .

For simplicity, in this proof, we denote $\mathcal{T}_h^2 V_{k,h+1}(s_h^i, a_h^i)$ as $\tilde{f}_k^*(s_h^i, a_h^i)$ where $\tilde{f}_k^* \in \mathcal{F}_h$ exists due to our assumption. For any function $V : \mathcal{S} \rightarrow [0, 1]$, let $\tilde{\eta}_h^k(V) = (r_h^k + V(s_{h+1}^k))^2 - \mathbb{E}_{s' \sim s_h^k, a_h^k} [(r(s_h^k, a_h^k, s') + V(s'))^2]$.

By simple calculation, for all $f \in \mathcal{F}_h$, we have

$$\begin{aligned} & \sum_{i \in [k-1]} \left(f(s_h^i, a_h^i) - \tilde{f}_k^*(s_h^i, a_h^i) \right)^2 + 2 \underbrace{\sum_{i \in [k-1]} \left(f(s_h^i, a_h^i) - \tilde{f}_k^*(s_h^i, a_h^i) \right) \cdot \tilde{\eta}_h^i(V_{k,h+1})}_{I(f, \tilde{f}_k^*, V_{k,h+1})} \\ &= \sum_{i \in [k-1]} \left[(r_h^i + V_{k,h+1}(s_h^i))^2 - f(s_h^i, a_h^i) \right]^2 - \sum_{i \in [k-1]} \left[(r_h^i + V_{k,h+1}(s_h^i))^2 - \tilde{f}_k^*(s_h^i, a_h^i) \right]^2. \end{aligned}$$

Due to the definition of $\tilde{f}_{k,h}$, we have

$$\sum_{(i,j) \in [k-1] \times [H]} \left(\tilde{f}_{k,h}(s_j^i, a_j^i) - \tilde{f}_k^*(s_j^i, a_j^i) \right)^2 + 2I(\tilde{f}_{k,h}, \tilde{f}_k^*, V_{k,h+1}) \leq 0. \quad (\text{F.1})$$

Then we give a high probability bound for $-I(f, \tilde{f}_k^*, V_{k,h+1})$ through the following calculation.

Applying Lemma H.4, for fixed f, \bar{f} and V , with probability at least $1 - \delta$,

$$\begin{aligned} -I(f, \bar{f}, V) &:= - \sum_{i \in [k-1]} \left(f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i) \right) \cdot \tilde{\eta}_h^i(V) \\ &\leq 8\lambda \sum_{i \in [k-1]} \left(f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i) \right)^2 + \frac{1}{\lambda} \cdot \log \frac{1}{\delta}. \end{aligned}$$

By the definition of V in line \cdot of Algorithm 1, $V_{k,h+1}$ lies in the optimistic value function class \mathcal{V}_k defined in (G.4). Applying a union bound for all the value functions V^c in the corresponding ϵ -net \mathcal{V}^c , we have

$$-I(f, \bar{f}, V^c) \leq \frac{1}{4} \sum_{i \in [k-1]} \left(f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i) \right)^2 + 32 \cdot \log \frac{N_\epsilon(k)}{\delta}$$

holds for all k with probability at least $1 - \delta$.

For all V such that $\|V - V^c\|_\infty \leq \epsilon$, we have $|\eta_h^i(V) - \eta_h^i(V^c)| \leq 4\epsilon$. Therefore, with probability $1 - \delta$, the following bound holds for $I(f, \bar{f}, V_{k,h+1})$:

$$-I(f, \bar{f}, V_{k,h+1}) \leq \frac{1}{4} \sum_{i \in [k-1]} \left(f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i) \right)^2 + 32 \cdot \log \frac{N_\epsilon(k)}{\delta} + 4\epsilon \cdot k.$$

To further bound $I(\tilde{f}_{k,h}, \tilde{f}_k^*, V_{k,h+1})$ in (F.1), we apply an ϵ -covering argument on \mathcal{F}_h and show that with probability at least $1 - \delta$,

$$-I(\tilde{f}_{k,h}, \tilde{f}_k^*, V_{k,h+1}) \leq \frac{1}{4} \sum_{i \in [k-1]} \left(\tilde{f}_{k,h}(s_h^i, a_h^i) - \tilde{f}_k^*(s_h^i, a_h^i) \right)^2 + 32 \cdot \log \frac{N_\epsilon(k) \cdot N_{\mathcal{F}}(\epsilon)}{\delta}$$

$$+ 16L\epsilon \cdot k \quad (\text{F.2})$$

for probability at least $1 - \delta$.

Substituting (F.2) into (F.1) and rearranging the terms,

$$\sum_{i \in [k-1]} \left(\tilde{f}_{k,h}(s_h^i, a_h^i) - \tilde{f}_k^*(s_h^i, a_h^i) \right)^2 \leq 128 \log \frac{N_\epsilon(k) \cdot N_{\mathcal{F}}(\epsilon) \cdot H}{\delta} + 64L\epsilon \cdot k$$

for all k with probability at least $1 - \delta/H$.

Finally, we apply a union bound over all $h \in [H]$ to conclude that $\mathcal{E}^{\tilde{f}}$ holds with probability at least $1 - \delta$. \square

F.1.2 Proof of Lemmas D.2 and D.3

Proof of Lemma D.2. Similar to the proof of Lemma D.1, we first prove that for an arbitrarily chosen $h \in [H]$, $\mathcal{E}_{k,h}^{\tilde{f}}$ holds with probability at least $1 - \delta/H$ for all k .

In this proof, we denote $\mathcal{T}_h V_{k,h+1}(s_h^i, a_h^i)$ as $\hat{f}_k^*(s_h^i, a_h^i)$ where $\hat{f}_k^* \in \mathcal{F}_h$ exists due to our assumption. For any function $V : \mathcal{S} \rightarrow [0, 1]$, let $\hat{\eta}_h^k(V) = (r_h^k + V(s_{h+1}^k)) - \mathbb{E}_{s' \sim s_h^k, a_h^k} [r_h(s_h^k, a_h^k, s') + V(s')]$.

For all $f \in \mathcal{F}_h$, we have

$$\begin{aligned} & \sum_{i \in [k-1]} \frac{1}{(\bar{\sigma}_{i,h})^2} \left(f(s_h^i, a_h^i) - \hat{f}_k^*(s_h^i, a_h^i) \right)^2 + 2 \underbrace{\sum_{i \in [k-1]} \frac{1}{(\bar{\sigma}_{i,h})^2} \left(f(s_h^i, a_h^i) - \hat{f}_k^*(s_h^i, a_h^i) \right) \cdot \hat{\eta}_h^i(V_{k,h+1})}_{I(f, \hat{f}_k^*, V_{k,h+1})} \\ &= \sum_{i \in [k-1]} \frac{1}{(\bar{\sigma}_{i,h})^2} [r_h^i + V_{k,h+1}(s_h^i) - f(s_h^i, a_h^i)]^2 - \sum_{i \in [k-1]} \frac{1}{(\bar{\sigma}_{i,h})^2} [r_h^i + V_{k,h+1}(s_h^i) - \hat{f}_k^*(s_h^i, a_h^i)]^2. \end{aligned}$$

By the definition of $\hat{f}_{k,h}$, we have

$$\sum_{i \in [k-1]} \frac{1}{(\bar{\sigma}_{i,h})^2} \left(\hat{f}_{k,h}(s_h^i, a_h^i) - \hat{f}_k^*(s_h^i, a_h^i) \right)^2 + 2I(\hat{f}_{k,h}, V_{k,h+1}) \leq 0. \quad (\text{F.3})$$

Then we give a high probability bound for $-I(\hat{f}_{k,h}, \hat{f}_k^*, V_{k,h+1})$ through the following calculation.

Applying Lemma H.4, for fixed f, \bar{f} and V , with probability at least $1 - \delta$,

$$\begin{aligned} -I(f, \bar{f}, V) &:= - \sum_{i \in [k-1]} \frac{1}{(\bar{\sigma}_{i,h})^2} \left(f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i) \right) \cdot \hat{\eta}_h^i(V) \\ &\leq 8\lambda \frac{1}{\alpha^2} \sum_{i \in [k-1]} \frac{1}{(\bar{\sigma}_{i,h})^2} \left(f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i) \right)^2 + \frac{1}{\lambda} \cdot \log \frac{1}{\delta}. \end{aligned}$$

By the definition of V in line \cdot of Algorithm 1, $V_{k,h+1}$ lies in the optimistic value function class \mathcal{V}_k defined in (G.4). Applying a union bound for all the value functions V^c in the corresponding ϵ -net \mathcal{V}^c , we have

$$-I(f, \bar{f}, V^c) \leq \frac{1}{4} \sum_{i \in [k-1]} \frac{1}{(\bar{\sigma}_{i,h})^2} \left(f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i) \right)^2 + \frac{32}{\alpha^2} \cdot \log \frac{N_\epsilon(k)}{\delta}$$

holds for all k with probability at least $1 - \delta$.

For all V such that $\|V - V^c\|_\infty \leq \epsilon$, we have $|\eta_h^i(V) - \eta_h^i(V^c)| \leq 4\epsilon$. Therefore, with probability $1 - \delta$, the following bound holds for $I(f, \bar{f}, V_{k,h+1})$:

$$-I(f, \bar{f}, V_{k,h+1}) \leq \frac{1}{4} \sum_{i \in [k-1]} \frac{1}{(\bar{\sigma}_{i,h})^2} \left(f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i) \right)^2 + \frac{32}{\alpha^2} \cdot \log \frac{N_\epsilon(k)}{\delta} + 4\epsilon \cdot k/\alpha^2.$$

To further bound $I(\widehat{f}_{k,h}, \widehat{f}_k^*, V_{k,h+1})$ in (F.3), we apply an ϵ -covering argument on \mathcal{F}_h and show that with probability at least $1 - \delta$,

$$-I(\widehat{f}_{k,h}, \widehat{f}_k^*, V_{k,h+1}) \leq \frac{1}{4} \sum_{i \in [k-1]} \frac{1}{(\bar{\sigma}_{i,h})^2} \left(\widehat{f}_{k,h}(s_h^i, a_h^i) - \widehat{f}_k^*(s_h^i, a_h^i) \right)^2 + 32 \cdot \log \frac{N_\epsilon(k) \cdot N_{\mathcal{F}}(\epsilon)}{\delta} + 16L\epsilon \cdot k/\alpha^2. \quad (\text{F.4})$$

Substituting (F.4) into (F.3) and rearranging the terms, we have

$$\sum_{i \in [k-1]} \frac{1}{(\bar{\sigma}_{i,h})^2} \left(\widehat{f}_{k,h}(s_h^i, a_h^i) - \widehat{f}_k^*(s_h^i, a_h^i) \right)^2 \leq 128 \cdot \log \frac{N_\epsilon(k) \cdot N_{\mathcal{F}}(\epsilon)H}{\delta} + 64L\epsilon \cdot k/\alpha^2$$

for all k with probability at least $1 - \delta/H$. Then we can complete the proof by using a union bound over all $h \in [H]$. \square

Proof of Lemma D.3. The proof is almost identical to the proof of Lemma D.2. \square

F.1.3 Proof of Lemmas D.4 and D.5

Proof of Lemma D.4. Similar to the proof of Lemma D.1, we first prove that for an arbitrary $h \in [H]$, $\bar{\mathcal{E}}_{k,h}^{\widehat{f}}$ holds with probability at least $1 - \delta/H$ for all k .

In this proof, we denote $\mathcal{T}_h V_{k,h+1}(s_h^i, a_h^i)$ as $\widehat{f}_k^*(s_h^i, a_h^i)$ where $\widehat{f}_k^* \in \mathcal{F}$ exists due to our assumption. For any function $V : \mathcal{S} \rightarrow [0, 1]$, let $\widehat{\eta}_h^k(V) = (r_h^k + V(s_{h+1}^k)) - \mathbb{E}_{s' \sim s_h^k, a_h^k} [r_h(s_h^k, a_h^k, s') + V(s')]$.

For all $f \in \mathcal{F}_h$, we have

$$\begin{aligned} & \sum_{i \in [k-1]} \frac{\widehat{\mathbf{1}}_{i,h}}{(\bar{\sigma}_{i,h})^2} \left(f(s_h^i, a_h^i) - \widehat{f}_k^*(s_h^i, a_h^i) \right)^2 + 2 \underbrace{\sum_{i \in [k-1]} \frac{\widehat{\mathbf{1}}_{i,h}}{(\bar{\sigma}_{i,h})^2} \left(f(s_h^i, a_h^i) - \widehat{f}_k^*(s_h^i, a_h^i) \right) \cdot \widehat{\eta}_h^i(V_{k,h+1})}_{I(f, \widehat{f}_k^*, V_{k,h+1})} \\ &= \sum_{i \in [k-1]} \frac{\widehat{\mathbf{1}}_{i,h}}{(\bar{\sigma}_{i,h})^2} [r_h^i + V_{k,h+1}(s_h^i) - f(s_h^i, a_h^i)]^2 - \sum_{i \in [k-1]} \frac{\widehat{\mathbf{1}}_{i,h}}{(\bar{\sigma}_{i,h})^2} [r_h^i + V_{k,h+1}(s_h^i) - \widehat{f}_k^*(s_h^i, a_h^i)]^2. \end{aligned}$$

Due to the definition of $\widehat{f}_{k,h}$, we have

$$\sum_{i \in [k-1]} \frac{\widehat{\mathbf{1}}_{i,h}}{(\bar{\sigma}_{i,h})^2} \left(\widehat{f}_{k,h}(s_h^i, a_h^i) - \widehat{f}_k^*(s_h^i, a_h^i) \right)^2 + 2I(\widehat{f}_{k,h}, V_{k,h+1}) \leq 0. \quad (\text{F.5})$$

Then it suffices to bound the value of $I(f, \bar{f}, V_{k,h+1})$ for all $f, \bar{f} \in \mathcal{F}$.

Unlike the proof for Lemma D.1, we decompose $I(f, \bar{f}, V_{k,h+1})$ into two parts:

$$\begin{aligned} I(f, \bar{f}, V_{k,h+1}) &= \sum_{i \in [k-1]} \frac{\widehat{\mathbf{1}}_{i,h}}{(\bar{\sigma}_{i,h})^2} (f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i)) \cdot \widehat{\eta}_h^i(V_{h+1}^*) \\ &\quad + \sum_{i \in [k-1]} \frac{\widehat{\mathbf{1}}_{i,h}}{(\bar{\sigma}_{i,h})^2} (f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i)) \cdot \widehat{\eta}_h^i(V_{k,h+1} - V_{h+1}^*). \quad (\text{F.6}) \end{aligned}$$

Then we bound the two terms separately.

For the first term, we first check the following conditions before applying Lemma H.2, which is a variant of Freedman inequality. For fixed f and \bar{f} , we have

$$\mathbb{E} \left[\frac{\widehat{\mathbf{1}}_{i,h}}{(\bar{\sigma}_{i,h})^2} (f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i)) \cdot \widehat{\eta}_h^i(V_{h+1}^*) \right] = 0,$$

since s_{h+1}^i is sampled from $\mathbb{P}_h(\cdot | s_h^i, a_h^i)$.

Next, we need to derive a bound for the maximum absolute value of each ‘weighted’ transition noise:

$$\begin{aligned}
& \max_{i \in [k-1]} \left| \frac{\hat{\mathbb{1}}_{i,h}}{(\bar{\sigma}_{i,h})^2} (f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i)) \cdot \hat{\eta}_h^i(V_{h+1}^*) \right| \\
& \leq 4 \max_{i \in [k-1]} \left| \frac{\hat{\mathbb{1}}_{i,h}}{(\bar{\sigma}_{i,h})^2} (f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i)) \right| \\
& \leq 4 \max_{i \in [k-1]} \frac{1}{(\bar{\sigma}_{i,h})^2} \sqrt{D_{\mathcal{F}_h}^2(z_{i,h}; z_{[i-1],h}, \bar{\sigma}_{[i-1],h}) \left(\sum_{\tau \in [k-1]} \frac{\hat{\mathbb{1}}_{\tau,h}}{(\bar{\sigma}_{\tau,h})^2} (f(s_h^\tau, a_h^\tau) - \bar{f}(s_h^\tau, a_h^\tau))^2 + \lambda \right)} \\
& \leq 4 \cdot \gamma^{-2} \sqrt{\sum_{\tau \in [k-1]} \frac{\hat{\mathbb{1}}_{\tau,h}}{(\bar{\sigma}_{\tau,h})^2} (f(s_h^\tau, a_h^\tau) - \bar{f}(s_h^\tau, a_h^\tau))^2 + \lambda}, \tag{F.7}
\end{aligned}$$

where the second inequality follows from the definition of $\mathcal{D}_{\mathcal{F}_h}$ in Definition 2.4, the last inequality holds since $\bar{\sigma}_{i,h} \geq \gamma \cdot D_{\mathcal{F}_h}^{1/2}(z_{i,h}; z_{[i-1],h}, \bar{\sigma}_{[i-1],h})$ according to line 19 in Algorithm 1. From the definition of $\hat{\mathbb{1}}_{i,h}$ in (D.6) we directly obtain the following upper bound of the variance:

$$\begin{aligned}
& \sum_{i \in [k-1]} \mathbb{E} \left[\frac{\hat{\mathbb{1}}_{i,h}}{(\bar{\sigma}_{i,h})^4} (f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i))^2 \cdot \hat{\eta}_h^i(V_{h+1}^*)^2 \right] \\
& \leq 4 \sum_{i \in [k-1]} \frac{\hat{\mathbb{1}}_{i,h}}{(\bar{\sigma}_{i,h})^2} (f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i))^2 \leq 4L^2k/\alpha^2 := V.
\end{aligned}$$

Applying Lemma H.2 with $V = 4L^2k/\alpha^2$, $M = 2L/\alpha^2$ and $v = m = 1$, for fixed f, \bar{f}, k , with probability at least $1 - \delta/(2k^2 \cdot N_{\mathcal{F}}(\epsilon) \cdot H)$,

$$- \sum_{i \in [k-1]} \frac{\hat{\mathbb{1}}_{i,h}}{(\bar{\sigma}_{i,h})^2} (f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i)) \cdot \hat{\eta}_h^i(V_{h+1}^*) \tag{F.8}$$

$$\begin{aligned}
& \leq \iota \sqrt{2 \left(8 \sum_{i \in [k-1]} \frac{\hat{\mathbb{1}}_{i,h}}{(\bar{\sigma}_{i,h})^2} (f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i))^2 + 1 \right)} \\
& \quad + \frac{2}{3} \iota^2 \left(8\gamma^{-2} \sqrt{\sum_{\tau \in [k-1]} \frac{\hat{\mathbb{1}}_{\tau,h}}{(\bar{\sigma}_{\tau,h})^2} (f(s_h^\tau, a_h^\tau) - \bar{f}(s_h^\tau, a_h^\tau))^2 + \lambda + 1} \right) \tag{F.9}
\end{aligned}$$

$$\leq \left(4\iota + \frac{16}{3} \iota^2 \gamma^{-2} \right) \sqrt{\sum_{\tau \in [k-1]} \frac{\hat{\mathbb{1}}_{\tau,h}}{(\bar{\sigma}_{\tau,h})^2} (f(s_h^\tau, a_h^\tau) - \bar{f}(s_h^\tau, a_h^\tau))^2 + \lambda + \sqrt{2}\iota + \frac{2}{3} \iota^2}, \tag{F.10}$$

where $\iota := \iota_1(k, h, \delta) = \sqrt{\log \frac{2k^2(2 \log(\frac{4L^2k}{\alpha^2}) + 2) \cdot (\log(\frac{4L}{\alpha^2}) + 2) \cdot N_{\mathcal{F}}(\epsilon)}{\delta/H}}$.

Using a union bound across all $f, \bar{f} \in \mathcal{C}(\mathcal{F}_h, \epsilon)$ and $k \geq 1$, with probability at least $1 - \delta/H$,

$$\begin{aligned}
& - \sum_{i \in [k-1]} \frac{\hat{\mathbb{1}}_{i,h}}{(\bar{\sigma}_{i,h})^2} (f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i)) \cdot \hat{\eta}_h^i(V_{h+1}^*) \\
& \leq \left(4\iota + \frac{16}{3} \iota^2 \gamma^{-2} \right) \sqrt{\sum_{\tau \in [k-1]} \frac{\hat{\mathbb{1}}_{\tau,h}}{(\bar{\sigma}_{\tau,h})^2} (f(s_h^\tau, a_h^\tau) - \bar{f}(s_h^\tau, a_h^\tau))^2 + \lambda + \sqrt{2}\iota + \frac{2}{3} \iota^2} \tag{F.11}
\end{aligned}$$

holds for all $f, \bar{f} \in \mathcal{C}(\mathcal{F}_h, \epsilon)$ and k . Due to the definition of ϵ -net, we deduce that for $\hat{f}_{k,h}$ and \hat{f}_k^* ,

$$- \sum_{i \in [k-1]} \frac{\hat{\mathbb{1}}_{i,h}}{(\bar{\sigma}_{i,h})^2} \left(\hat{f}_{k,h}(s_h^i, a_h^i) - \hat{f}_k^*(s_h^i, a_h^i) \right) \cdot \hat{\eta}_h^i(V_{h+1}^*)$$

$$\begin{aligned}
&\leq \left(4\iota + \frac{16}{3}\iota^2\gamma^{-2}\right) \sqrt{\sum_{\tau \in [k-1]} \frac{\hat{\mathbf{1}}_{\tau,h}}{(\bar{\sigma}_{\tau,h})^2} \left(\hat{f}_{k,h}(s_h^\tau, a_h^\tau) - \hat{f}_k^*(s_h^\tau, a_h^\tau)\right)^2} + \lambda + \sqrt{2}\iota + \frac{2}{3}\iota^2 \\
&\quad + \left(4\iota + \frac{16}{3}\iota^2\gamma^{-2}\right) \sqrt{16k\epsilon L/\alpha^2} + 8\epsilon k/\alpha^2. \tag{F.12}
\end{aligned}$$

For the second term in (F.6), the following inequality holds for all $V \in \mathcal{V}_{k,h+1}$,

$$\begin{aligned}
&\mathbb{E} \left[\frac{1}{(\bar{\sigma}_{i,h})^2} (f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i)) \cdot \hat{\eta}_h^i (V - V_{h+1}^*) \right] = 0, \\
&\max_{i \in [k-1]} \left| \frac{1}{(\bar{\sigma}_{i,h})^2} (f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i)) \cdot \hat{\eta}_h^i (V - V_{h+1}^*) \right| \\
&\leq 4 \max_{i \in [k-1]} \left| \frac{1}{(\bar{\sigma}_{i,h})^2} (f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i)) \right| \\
&\leq 4 \max_{i \in [k-1]} \frac{1}{(\bar{\sigma}_{i,h})^2} \sqrt{D_{\mathcal{F}_h}^2(z_{i,h}; z_{[i-1],h}, \bar{\sigma}_{[i-1],h}) \sum_{\tau \in [k-1]} \frac{1}{(\bar{\sigma}_{\tau,h})^2} (f(s_h^\tau, a_h^\tau) - \bar{f}(s_h^\tau, a_h^\tau))^2} + \lambda} \\
&= 4 \cdot \gamma^{-2} \sqrt{\sum_{\tau \in [k-1]} \frac{1}{(\bar{\sigma}_{\tau,h})^2} (f(s_h^\tau, a_h^\tau) - \bar{f}(s_h^\tau, a_h^\tau))^2} + \lambda,
\end{aligned}$$

where the calculation is similar to that in (F.7).

We denote the sum of variance by $\text{Var}(V - V_{h+1}^*)$ as follows for simplicity:

$$\text{Var}(V - V_{h+1}^*) := \sum_{i \in [k-1]} \mathbb{E} \left[\frac{\hat{\mathbf{1}}_{i,h}}{(\bar{\sigma}_{i,h})^4} (f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i))^2 \cdot \hat{\eta}_h^i (V - V_{h+1}^*)^2 \right] \tag{F.13}$$

$$\leq k \cdot L^2/\alpha^4. \tag{F.14}$$

For $V_{k,h+1}$, we have

$$\begin{aligned}
\text{Var}(V_{k,h+1} - V_{h+1}^*) &:= \sum_{i \in [k-1]} \mathbb{E} \left[\frac{\hat{\mathbf{1}}_{i,h}}{(\bar{\sigma}_{i,h})^4} (f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i))^2 \cdot \hat{\eta}_h^i (V_{k,h+1} - V_{h+1}^*)^2 \right] \\
&\leq \frac{4}{\log N_{\mathcal{F}}(\epsilon) + \log N_\epsilon(K)} \sum_{i \in [k-1]} \frac{\hat{\mathbf{1}}_{i,h}}{(\bar{\sigma}_{i,h})^2} (f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i))^2,
\end{aligned}$$

where the second inequality holds due to the definition of $\hat{\mathbf{1}}_{i,h}$ in (D.6).

With a similar argument as shown in (F.8)~(F.12), we have with probability at least $1 - \delta/(2k^2 \cdot N_{\mathcal{F}}(\epsilon) \cdot N_\epsilon(k-1) \cdot H)$, for a fixed f, \bar{f}, k and V , (applying Lemma H.2, with $V = k \cdot L^2/\alpha^4$ and $M = 2L/\alpha^2, v = (\log N_{\mathcal{F}}(\epsilon) + \log N_\epsilon(K))^{-1/2}, m = v^2$.)

$$\begin{aligned}
& - \sum_{i \in [k-1]} \frac{\hat{\mathbf{1}}_{i,h}}{(\bar{\sigma}_{i,h})^2} (f(s_h^i, a_h^i) - \bar{f}(s_h^i, a_h^i)) \cdot \hat{\eta}_h^i (V_{k,h+1} - V_{h+1}^*) \\
&\leq \iota \sqrt{2 \left(2\text{Var}(V - V_{h+1}^*) + (\log N_{\mathcal{F}}(\epsilon) + \log N_\epsilon(K))^{-1} \right)} \\
&\quad + \frac{2}{3}\iota^2 \left(8\gamma^{-2} \sqrt{\sum_{\tau \in [k-1]} \frac{\hat{\mathbf{1}}_{\tau,h}}{(\bar{\sigma}_{\tau,h})^2} (f(s_h^\tau, a_h^\tau) - \bar{f}(s_h^\tau, a_h^\tau))^2} + \lambda + (\log N_{\mathcal{F}}(\epsilon) + \log N_\epsilon(K))^{-1} \right).
\end{aligned}$$

where

$$\log \frac{2k^2 \left(2 \log \frac{L^2 k (\log N_{\mathcal{F}}(\epsilon) \cdot N_\epsilon(K))^{1/2}}{\alpha^4} + 2 \right) \cdot \left(\log \left(\frac{4L (\log N_{\mathcal{F}}(\epsilon) \cdot N_\epsilon(K))}{\alpha^2} \right) + 2 \right) \cdot N_{\mathcal{F}}(\epsilon) \cdot N_\epsilon(k)}{\delta/H}$$

$$\leq \log \frac{2k^2 \left(2 \log \frac{L^2 k}{\alpha^4} + 2\right) \cdot \left(\log \left(\frac{4L}{\alpha^2}\right) + 2\right) \cdot N_{\mathcal{F}}^4(\epsilon) \cdot N_{\epsilon}^2(K)}{\delta/H} := \iota_2^2(k, h, \delta).$$

Using a union bound over all $(f, \bar{f}, V) \in \mathcal{C}(\mathcal{F}_h, \epsilon) \times \mathcal{C}(\mathcal{F}_h, \epsilon) \times \mathcal{V}_{k, h+1}^c$ and all $k \geq 1$, we have the inequality above holds for all such f, \bar{f}, V, k with probability at least $1 - \delta/H$.

There exists a $V_{k, h+1}^c$ in the ϵ -net such that $\|V_{k, h+1} - V_{k, h+1}^c\|_{\infty} \leq \epsilon$. Then we have

$$\begin{aligned} & - \sum_{i \in [k-1]} \frac{\hat{\mathbb{1}}_{i, h}}{(\bar{\sigma}_{i, h})^2} \left(\hat{f}_{k, h}(s_h^i, a_h^i) - \hat{f}_k^*(s_h^i, a_h^i) \right) \cdot \hat{\eta}_h^i(V_{k, h+1} - V_{h+1}^*) \\ & \leq O \left(\frac{\iota_2(k, h, \delta)}{\sqrt{\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)}} + \frac{\iota_2(k, h, \delta)^2}{\gamma^{-2}} \right) \\ & \quad \cdot \sqrt{\sum_{\tau \in [k-1]} \frac{\hat{\mathbb{1}}_{\tau, h}}{(\bar{\sigma}_{\tau, h})^2} \left(\hat{f}_{k, h}(s_h^{\tau}, a_h^{\tau}) - \hat{f}_k^*(s_h^{\tau}, a_h^{\tau}) \right)^2} + \lambda} \\ & \quad + O(\epsilon k L / \alpha^2) + O \left(\frac{\iota_2^2(k, h, \delta)}{\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)} \right), \end{aligned} \tag{F.15}$$

for all k with at least probability $1 - \delta/H$.

Substituting (F.15) and (F.12) into (F.5), we can conclude that

$$\begin{aligned} & \sum_{i \in [k-1]} \frac{\hat{\mathbb{1}}_{i, h}}{(\bar{\sigma}_{i, h})^2} \left(\hat{f}_{k, h}(s_h^i, a_h^i) - \hat{f}_k^*(s_h^i, a_h^i) \right)^2 \\ & \leq O \left(\left(\frac{\iota_2(k, h, \delta)}{\sqrt{\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)}} + \iota_2(k, h, \delta)^2 \cdot \gamma^{-2} \right)^2 \right) + O \left((\iota_1(k, h, \delta) + \iota_1(k, h, \delta)^2 / \gamma^2)^2 \right). \end{aligned}$$

From the definition of γ in (3.3), we can rewrite the upper bound of the squared loss

$\sum_{i \in [k-1]} \frac{\hat{\mathbb{1}}_{i, h}}{(\bar{\sigma}_{i, h})^2} \left(\hat{f}_{k, h}(s_h^i, a_h^i) - \hat{f}_k^*(s_h^i, a_h^i) \right)^2$ as follows:

$$\begin{aligned} & \lambda + \sum_{i \in [k-1]} \frac{\hat{\mathbb{1}}_{i, h}}{(\bar{\sigma}_{i, h})^2} \left(\hat{f}_{k, h}(s_h^i, a_h^i) - \hat{f}_k^*(s_h^i, a_h^i) \right)^2 \\ & \leq O \left(\log \frac{2k^2 \left(2 \log \frac{L^2 k}{\alpha^4} + 2\right) \cdot \left(\log \left(\frac{4L}{\alpha^2}\right) + 2\right)}{\delta/H} \right) \cdot [\log(N_{\mathcal{F}}(\epsilon)) + 1] + O(\lambda) + O(\epsilon k L / \alpha^2). \end{aligned}$$

□

Proof of Lemma D.5. The proof is almost identical to the proof of Lemma D.4. □

F.2 Proof of Optimism and Pessimism

F.2.1 Proof of Lemma D.6

Proof of Lemma D.6. According to the definition of $D_{\mathcal{F}}^2$ function, we have

$$\begin{aligned} & \left(\hat{f}_{k, h}(s, a) - \mathcal{T}_h V_{k, h+1}(s, a) \right)^2 \\ & \leq D_{\mathcal{F}_h}^2(z; z_{[k-1], h}, \bar{\sigma}_{[k-1], h}) \times \left(\lambda + \sum_{i=1}^{k-1} \frac{1}{(\bar{\sigma}_{i, h})^2} \left(\hat{f}_{k, h}(s_h^i, a_h^i) - \mathcal{T}_h V_{k, h+1}(s_h^i, a_h^i) \right)^2 \right) \\ & \leq \hat{\beta}_k^2 \times D_{\mathcal{F}_h}^2(z; z_{[k-1], h}, \bar{\sigma}_{[k-1], h}), \end{aligned}$$

where the first inequality holds due the definition of $D_{\mathcal{F}}^2$ function with the Assumption 2.2 and the second inequality holds due to the events $\mathcal{E}_h^{\hat{f}}$. Thus, we have

$$|\hat{f}_{k,h}(s, a) - \mathcal{T}_h V_{k,h+1}(s, a)| \leq \hat{\beta}_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}).$$

With a similar argument, for the pessimistic value function $\check{f}_{k,h}$, we have

$$\begin{aligned} & (\check{f}_{k,h}(s, a) - \mathcal{T}_h \check{V}_{k,h}(s, a))^2 \\ & \leq D_{\mathcal{F}_h}^2(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) \times \left(\lambda + \sum_{i=1}^{k-1} \frac{1}{(\bar{\sigma}_{i,h})^2} \left(\check{f}_{k,h}(s_h^i, a_h^i) - \mathcal{T}_h \check{V}_{k,h+1}(s_h^i, a_h^i) \right)^2 \right) \\ & \leq \check{\beta}_k^2 \times D_{\mathcal{F}_h}^2(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), \end{aligned}$$

where the first inequality holds due to the definition of $D_{\mathcal{F}}^2$ function with the Assumption 2.2 and the second inequality holds due to the events $\mathcal{E}_h^{\check{f}}$. In addition, we have,

$$|\check{f}_{k,h}(s, a) - \mathcal{T}_h \check{V}_{k,h+1}(s, a)| \leq \hat{\beta}_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}).$$

Thus, we complete the proof of Lemma D.6. \square

F.2.2 Proof of Lemma D.7

Proof of Lemma D.7. We use induction to prove the optimistic and pessimistic property. First, we study the basic case with the last stage $H + 1$. In this situation, $Q_{k,H+1}(s, a) = Q_h^*(s, a) = \check{Q}_{k,h}(s, a) = 0$ and $V_{k,h}(s) = V_h^*(s) = \check{V}_{k,h}(s) = 0$ hold for all state $s \in \mathcal{S}$ and action $a \in \mathcal{A}$. Therefore, Lemma D.7 holds for the basic case (stage $H + 1$).

Second, if Lemma D.7 holds for stage $h + 1$, then we focus on the stage h . Notice that the event $\tilde{\mathcal{E}}_h$ directly implies the event $\tilde{\mathcal{E}}_{h+1}$. Therefore, according to the induction assumption, the following inequality holds for all state $s \in \mathcal{S}$ and episode $k \in [K]$.

$$V_{k,h+1}(s) \geq V_{h+1}^*(s) \geq \check{V}_{k,h}(s). \quad (\text{F.16})$$

Thus, for all episode $k \in [K]$ and state-action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have

$$\begin{aligned} & \hat{f}_{k,h}(s, a) + b_{k,h}(s, a) - Q_h^*(s, a) \\ & \geq \mathcal{T}_h V_{k,h+1}(s, a) - \hat{\beta}_k \cdot D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) + b_{k,h}(s, a) - Q_h^*(s, a) \\ & \geq \mathcal{T}_h V_{k,h+1}(s, a) - Q_h^*(s, a) \\ & = \mathbb{P}_h V_{k,h+1}(s, a) - \mathbb{P}_h V_h^*(s, a) \\ & \geq 0, \end{aligned} \quad (\text{F.17})$$

where the first inequality holds due to Lemma D.6, the second inequality holds due to the definition of the exploration bonus $b_{k,h}$ and the last inequality holds due to the (F.16). Therefore, the optimal value function $Q_h^*(s, a)$ is upper bounded by

$$Q_h^*(s, a) \leq \min \left\{ \min_{1 \leq i \leq k} \hat{f}_{i,h}(s, a) + b_{i,h}(s, a), 1 \right\} \leq Q_{k,h}(s, a), \quad (\text{F.18})$$

where the first inequality holds due to (F.17) with the fact that $Q_h^*(s, a) \leq 1$ and the second inequality holds due to the update rule of value function $Q_{k,h}$.

With a similar argument, for the pessimistic estimator $\check{f}_{k,h}$, we have

$$\begin{aligned} & \check{f}_{k,h}(s, a) - b_{k,h}(s, a) - Q_h^*(s, a) \\ & \leq \mathcal{T}_h \check{V}_{k,h+1}(s, a) + \check{\beta}_k \cdot D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) - b_{k,h}(s, a) - Q_h^*(s, a) \\ & \leq \mathcal{T}_h \check{V}_{k,h+1}(s, a) - Q_h^*(s, a) \\ & = \mathbb{P}_h \check{V}_{k,h+1}(s, a) - \mathbb{P}_h V_h^*(s, a) \\ & \leq 0, \end{aligned} \quad (\text{F.19})$$

where the first inequality holds due to Lemma D.6, the second inequality holds due to the definition of the exploration bonus $b_{k,h}$ and the last inequality holds due to the (F.19). Therefore, the optimal value function $Q_h^*(s, a)$ is lower bounded by

$$Q_h^*(s, a) \geq \max \left\{ \max_{1 \leq i \leq k} \check{f}_{i,h}(s, a) - b_{i,h}(s, a), 0 \right\} \geq \check{Q}_{k,h}(s, a), \quad (\text{F.20})$$

where the first inequality holds due to (F.19) with the fact that $Q_h^*(s, a) \geq 0$ and the second inequality holds due to the update rule of value function $\check{Q}_{k,h}$.

Furthermore, for the value functions $V_{k,h}$ and $\check{V}_{k,h}$, we have

$$\begin{aligned} V_{k,h}(s) &= \max_a Q_{k,h}(s, a) \geq \max_a Q_h^*(s, a) = V_h^*(s), \\ \check{V}_{k,h}(s) &= \max_a \check{Q}_{k,h}(s, a) \leq \max_a Q_h^*(s, a) = V_h^*(s), \end{aligned}$$

where the first inequality holds due to (F.18) and the second inequality holds due to (F.20). Thus, by induction, we complete the proof of Lemma D.7. \square

E.3 Proof of Monotonic Variance Estimator

E.3.1 Proof of Lemma D.8

Proof of Lemma D.8. According to the definition of $D_{\mathcal{F}}^2$ function, we have

$$\begin{aligned} & (\hat{f}_{k,h}(s, a) - \mathcal{T}_h V_{k,h+1}(s, a))^2 \\ & \leq D_{\mathcal{F}_h}^2(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) \times \left(\lambda + \sum_{i=1}^{k-1} \frac{1}{(\bar{\sigma}_{i,h})^2} \left(\hat{f}_{k,h}(s_h^i, a_h^i) - \mathcal{T}_h V_{k,h+1}(s_h^i, a_h^i) \right)^2 \right) \\ & \leq \beta_k^2 \times D_{\mathcal{F}_h}^2(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), \end{aligned}$$

where the first inequality holds due the definition of $D_{\mathcal{F}}^2$ function with the Assumption 2.2 and the second inequality holds due to the events $\underline{\mathcal{E}}^{\hat{f}}$. Thus, we have

$$|\hat{f}_{k,h}(s, a) - \mathcal{T}_h V_{k,h+1}(s, a)| \leq \beta_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}).$$

For the pessimistic value function $\check{f}_{k,h}$, we have

$$\begin{aligned} & (\check{f}_{k,h}(s, a) - \mathcal{T}_h \check{V}_{k,h}(s, a))^2 \\ & \leq D_{\mathcal{F}_h}^2(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) \times \left(\lambda + \sum_{i=1}^{k-1} \frac{1}{(\bar{\sigma}_{i,h})^2} \left(\check{f}_{k,h}(s_h^i, a_h^i) - \mathcal{T}_h \check{V}_{k,h+1}(s_h^i, a_h^i) \right)^2 \right) \\ & \leq \beta_k^2 \times D_{\mathcal{F}_h}^2(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), \end{aligned}$$

where the first inequality holds due the definition of $D_{\mathcal{F}}^2$ function with the Assumption 2.2 and the second inequality holds due to the events $\underline{\mathcal{E}}^{\check{f}}$. In addition, we have

$$|\check{f}_{k,h}(s, a) - \mathcal{T}_h \check{V}_{k,h+1}(s, a)| \leq \beta_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}).$$

With a similar argument, for the second-order estimator $\tilde{f}_{k,h}$, we have

$$\begin{aligned} & (\tilde{f}_{k,h}(s, a) - \mathcal{T}_h^2 V_{k,h}(s, a))^2 \\ & \leq D_{\mathcal{F}_h}^2(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) \times \left(\lambda + \sum_{i=1}^{k-1} \frac{1}{(\bar{\sigma}_{i,h})^2} \left(\tilde{f}_{k,h}(s_h^i, a_h^i) - \mathcal{T}_h^2 V_{k,h+1}(s_h^i, a_h^i) \right)^2 \right) \\ & \leq \tilde{\beta}_k^2 \times D_{\mathcal{F}_h}^2(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), \end{aligned}$$

where the first inequality holds due the definition of $D_{\mathcal{F}}^2$ function with the Assumption 2.2 and the second inequality holds due to the events $\underline{\mathcal{E}}^{\tilde{f}}$. Therefore, we have

$$|\tilde{f}_{k,h}(s, a) - \mathcal{T}_h^2 V_{k,h}(s, a)| \leq \tilde{\beta}_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}).$$

Now, we complete the proof of Lemma D.8. \square

F.3.2 Proof of Lemma D.9

Proof of Lemma D.9. First, according to Lemma D.8, we have

$$\begin{aligned}
& \left| [\bar{\mathbb{V}}_h V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{V}_h V_{k,h+1}](s_h^k, a_h^k) \right| \\
&= \left| \tilde{f}_{k,h} - \hat{f}_{k,h}^2 - [\mathbb{P}_h V_{k,h+1}^2](s_h^k, a_h^k) + ([\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k))^2 \right| \\
&\leq \left| \hat{f}_{k,h} - ([\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k))^2 \right| + \left| \tilde{f}_{k,h} - [\mathbb{P}_h V_{k,h+1}^2](s_h^k, a_h^k) \right| \\
&\leq 2L \left| \hat{f}_{k,h} - ([\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k))^2 \right| + \left| \tilde{f}_{k,h} - [\mathbb{P}_h V_{k,h+1}^2](s_h^k, a_h^k) \right| \\
&\leq (2L\beta_k + \tilde{\beta}_k) D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) \\
&= E_{k,h},
\end{aligned} \tag{F.21}$$

where the first inequality holds due to $|a + b| \leq |a| + |b|$, the second inequality holds due to $|a^2 - b^2| = |a - b| \cdot |a + b| \leq |a - b| \cdot 2 \max(|a|, |b|)$ and the last inequality holds due to Lemma D.8.

For the difference between variances $[\mathbb{V}_h V_{k,h+1}](s_h^k, a_h^k)$ and $[\mathbb{V}_h V_{h+1}^*](s_h^k, a_h^k)$, it can be upper bounded by

$$\begin{aligned}
& \left| [\mathbb{V}_h V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{V}_h V_{h+1}^*](s_h^k, a_h^k) \right| \\
&= \left| [\mathbb{P}_h V_{k,h+1}^2](s_h^k, a_h^k) - ([\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k))^2 - [\mathbb{P}_h (V_{h+1}^*)^2](s_h^k, a_h^k) + ([\mathbb{P}_h V_{h+1}^*](s_h^k, a_h^k))^2 \right| \\
&\leq \left| [\mathbb{P}_h V_{k,h+1}^2](s_h^k, a_h^k) - [\mathbb{P}_h (V_{h+1}^*)^2](s_h^k, a_h^k) \right| + \left| ([\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k))^2 - ([\mathbb{P}_h V_{h+1}^*](s_h^k, a_h^k))^2 \right| \\
&\leq 4 \left| [\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{P}_h V_{h+1}^*](s_h^k, a_h^k) \right| \\
&\leq ([\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{P}_h \check{V}_{k,h+1}](s_h^k, a_h^k)) \\
&\leq \hat{f}_{k,h}(s_h^k, a_h^k) - \check{f}_{k,h}(s_h^k, a_h^k) + 2\beta_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}),
\end{aligned} \tag{F.22}$$

where the first inequality holds due to $|a + b| \leq |a| + |b|$, the second inequality holds due to $1 \geq V_{k,h+1}(\cdot) \geq V_{h+1}^*(\cdot) \geq 0$ (Lemma D.7), the third inequality holds due to $V_{h+1}^*(\cdot) \geq \check{V}_{k,h+1}(\cdot)$ (Lemma D.7) and the last inequality holds due to Lemma D.8. Combining the results in (F.21) and (F.22) with the fact that $0 \leq V_{k,h+1}(\cdot), V_{h+1}^*(\cdot) \leq 1$, we have

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

Thus, we complete the proof of Lemma D.9. \square

F.3.3 Proof of Lemma D.10

Proof of Lemma D.10. On the events $\underline{\mathcal{E}}^{\hat{f}}$ and $\mathcal{E}_{h+1}^{\hat{f}}$, we have

$$\begin{aligned}
& [\mathbb{V}_h (V_{i,h+1} - V_{h+1}^*)](s_h^k, a_h^k) \\
&= [\mathbb{P}_h (V_{i,h+1} - V_{h+1}^*)^2](s_h^k, a_h^k) - ([\mathbb{P}_h (V_{i,h+1} - V_{h+1}^*)](s_h^k, a_h^k))^2 \\
&\leq [\mathbb{P}_h (V_{i,h+1} - V_{h+1}^*)^2](s_h^k, a_h^k) \\
&\leq 2 [\mathbb{P}_h (V_{i,h+1} - V_{h+1}^*)](s_h^k, a_h^k) \\
&\leq 2 ([\mathbb{P}_h V_{i,h+1}](s_h^k, a_h^k) - [\mathbb{P}_h \check{V}_{k,h+1}](s_h^k, a_h^k)) \\
&\leq 2 ([\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{P}_h \check{V}_{k,h+1}](s_h^k, a_h^k)),
\end{aligned} \tag{F.23}$$

where the first equation holds since the reward function is deterministic, the second inequality holds due to Lemma D.7 with the fact that $0 \leq V_{i,h+1}(s), V_{h+1}^*(s) \leq 1$ and the third inequality holds due to Lemma D.7 and the last inequality holds due to the fact that $V_{k,h+1} \geq V_{i,h+1}$.

With a similar argument, on the events $\underline{\mathcal{E}}^{\check{f}}$ and $\mathcal{E}_{h+1}^{\check{f}}$, we have

$$\begin{aligned}
& [\mathbb{V}_h(V_{h+1}^* - \check{V}_{i,h+1})](s_h^k, a_h^k) \\
&= [\mathbb{P}_h(V_{h+1}^* - \check{V}_{i,h+1})^2](s_h^k, a_h^k) - ([\mathbb{P}_h(V_{h+1}^* - \check{V}_{i,h+1})](s_h^k, a_h^k))^2 \\
&\leq [\mathbb{P}_h(V_{h+1}^* - \check{V}_{i,h+1})^2](s_h^k, a_h^k) \\
&\leq 2[\mathbb{P}_h(V_{h+1}^* - \check{V}_{i,h+1})](s_h^k, a_h^k) \\
&\leq 2([\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{P}_h \check{V}_{i,h+1}](s_h^k, a_h^k)) \\
&\leq 2([\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{P}_h \check{V}_{k,h+1}](s_h^k, a_h^k)), \tag{F.24}
\end{aligned}$$

where the first inequality holds since the reward function is deterministic, the second and third inequality holds due to Lemma D.7 with the fact that $0 \leq \check{V}_{i,h+1}(s), V_{h+1}^*(s) \leq H$, the last inequality the fact $\check{V}_{k,h+1}(s) \leq \check{V}_{i,h+1}(s)$. For both variances $[\mathbb{V}_h(V_{i,h+1} - V_{h+1}^*)](s_h^k, a_h^k)$ and $[\mathbb{V}_h(V_{h+1}^* - \check{V}_{i,h+1})](s_h^k, a_h^k)$, they are upper bounded by

$$\begin{aligned}
& 2([\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{P}_h \check{V}_{k,h+1}](s_h^k, a_h^k)) \\
&= 2\mathcal{T}_h V_{k,h+1}(s_h^k, a_h^k) - 2\mathcal{T}_h \check{V}_{k,h+1}(s_h^k, a_h^k) \\
&\leq 2\hat{f}_{k,h}(s_h^k, a_h^k) - 2\check{f}_{k,h}(s_h^k, a_h^k) + 4\beta_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), \tag{F.25}
\end{aligned}$$

where the first inequality holds due to Lemma D.8 and the second inequality holds due to the definition of $F_{k,h}$. Substituting the result in (F.25) into (F.23), (F.24) and combining the fact that $[\mathbb{V}_h(V_{i,h+1} - V_{h+1}^*)](s_h^k, a_h^k), [\mathbb{V}_h(V_{h+1}^* - \check{V}_{i,h+1})](s_h^k, a_h^k) \leq 1$, we finish the proof of Lemma D.10. \square

F.4 Proof of Lemmas in Section D.4

F.4.1 Proof of Lemma D.11

Proof of Lemma D.11. We use induction to shows that the conclusion in Lemma D.7 and events $\mathcal{E}_h^{\hat{f}}, \mathcal{E}_H^{\check{f}}$ hold for each stage $h \in [H]$. First, for the basic situation (stage $H + 1$), $Q_{k,H+1}(s, a) = Q_h^*(s, a) = \check{Q}_{k,h}(s, a) = 0$ and $V_{k,h}(s) = V_h^*(s) = \check{V}_{k,h}(s) = 0$ hold for all state $s \in \mathcal{S}$ and action $a \in \mathcal{A}$. Therefore, Lemma D.7 holds for the basic case (stage $H + 1$)

Second, if Lemma D.7 holds for stage $h + 1$, then we focus on the stage h . According to Lemmas D.10 and Lemma D.9, we have the following inequalities:

$$\begin{aligned}
\sigma_{i,h}^2 &= [\bar{\mathbb{V}}_h V_{i,h+1}](s_h^i, a_h^i) + E_{i,h} + D_{i,h} \geq [\bar{\mathbb{V}}_h V_{h+1}^*](s_h^i, a_h^i), \\
\sigma_{i,h}^2 &\geq F_{i,h} \geq (\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)) \cdot [\mathbb{V}_h(V_{k,h+1} - V_{h+1}^*)](s_h^i, a_h^i), \\
\sigma_{i,h}^2 &\geq F_{i,h} \geq (\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)) \cdot [\mathbb{V}_h(V_{h+1}^* - \check{V}_{k,h+1})](s_h^i, a_h^i),
\end{aligned}$$

where the first inequality holds due to Lemma D.9, the second and third inequality holds due to Lemma D.10. Thus, the indicator function in events $\bar{\mathcal{E}}_h^{\hat{f}}$ and $\bar{\mathcal{E}}_h^{\check{f}}$ hold, which implies events $\mathcal{E}_h^{\hat{f}}, \mathcal{E}_H^{\check{f}}$ hold. Furthermore, when events $\mathcal{E}_h^{\hat{f}}, \mathcal{E}_H^{\check{f}}$ hold, then Lemma D.7 holds for stage h . Thus, we complete the proof of Lemma D.11 by induction. \square

F.4.2 Proof of Lemma D.13

Proof of Lemma D.13. For each stage h , we divide the episodes $\{1, 2, \dots, K\}$ to the following sets:

$$\begin{aligned}
\mathcal{I}_1 &= \{k \in [K] : D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) / \bar{\sigma}_{k,h} \geq 1\}, \\
\mathcal{I}_2 &= \{k \in [K] : D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) / \bar{\sigma}_{k,h} < 1, \bar{\sigma}_{k,h} = \sigma_{k,h}\}, \\
\mathcal{I}_3 &= \{k \in [K] : D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) / \bar{\sigma}_{k,h} < 1, \bar{\sigma}_{k,h} = \alpha\}, \\
\mathcal{I}_4 &= \{k \in [K] : D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) / \bar{\sigma}_{k,h} < 1, \bar{\sigma}_{k,h} = \gamma \times D_{\mathcal{F}_h}^{1/2}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h})\}.
\end{aligned}$$

The number of episode in set \mathcal{I}_1 is upper bounded by

$$|\mathcal{I}_1| = \sum_{k \in \mathcal{I}_1} \min \left(D_{\mathcal{F}_h}^2(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) / \bar{\sigma}_{k,h}^2, 1 \right) \leq \dim_{\alpha,K}(\mathcal{F}_h),$$

where the equation holds due to $D_{\mathcal{F}_h}^2(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) / \bar{\sigma}_{k,h}^2 \geq 1$ and the inequality holds due to the definition of Generalized Eluder dimension. Thus, for set \mathcal{I}_1 , the summation of confidence radius is upper bounded by

$$\sum_{k \in \mathcal{I}_1} \min \left(\beta D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), 1 \right) \leq |\mathcal{I}_1| \leq \dim_{\alpha,K}(\mathcal{F}_h). \quad (\text{F.26})$$

For set \mathcal{I}_2 , the summation of confidence radius is upper bounded by

$$\begin{aligned} & \sum_{k \in \mathcal{I}_2} \min \left(\beta D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), 1 \right) \\ & \leq \sum_{k \in \mathcal{I}_2} \beta D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) \\ & \leq \beta \sqrt{\sum_{k \in \mathcal{I}_2} \sigma_{k,h}^2} \cdot \sqrt{\sum_{k \in \mathcal{I}_2} D_{\mathcal{F}_h}^2(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) / \bar{\sigma}_{k,h}^2} \\ & \leq \beta \sqrt{\dim_{\alpha,K}(\mathcal{F}_h)} \sqrt{\sum_{k \in \mathcal{I}_2} \sigma_{k,h}^2}, \end{aligned} \quad (\text{F.27})$$

where the second inequality holds due to Cauchy-Schwartz inequality with $\sigma_{k,h} = \bar{\sigma}_{k,h}$ and the last inequality holds due to the definition of Generalized Eluder dimension with the fact that $D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) / \bar{\sigma}_{k,h} < 1$.

With a similar argument, the summation of confidence radius over set \mathcal{I}_3 is upper bounded by

$$\begin{aligned} & \sum_{k \in \mathcal{I}_3} \min \left(\beta D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), 1 \right) \\ & \leq \sum_{k \in \mathcal{I}_3} \beta D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) \\ & \leq \beta \sqrt{\sum_{k \in \mathcal{I}_3} \alpha^2} \cdot \sqrt{\sum_{k \in \mathcal{I}_3} D_{\mathcal{F}_h}^2(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) / \bar{\sigma}_{k,h}^2} \\ & \leq \beta \sqrt{\dim_{\alpha,K}(\mathcal{F}_h)} \sqrt{\sum_{k \in \mathcal{I}_3} \alpha^2}, \end{aligned} \quad (\text{F.28})$$

where the second inequality holds due to Cauchy-Schwartz inequality with $\bar{\sigma}_{k,h} = \alpha$ and the last inequality holds due to the definition of Generalized Eluder dimension with the fact that $D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) / \bar{\sigma}_{k,h} < 1$.

Finally, the summation of confidence radius over set \mathcal{I}_4 is upper bounded by With a similar argument, the summation of confidence radius over set \mathcal{I}_3 is upper bounded by

$$\begin{aligned} & \sum_{k \in \mathcal{I}_4} \min \left(\beta D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), 1 \right) \\ & \leq \sum_{k \in \mathcal{I}_4} \beta D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) \\ & = \sum_{k \in \mathcal{I}_4} \beta \gamma^2 D_{\mathcal{F}_h}^2(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) / \bar{\sigma}_{k,h}^2 \\ & \leq \beta \gamma^2 \dim_{\alpha,K}(\mathcal{F}_h), \end{aligned} \quad (\text{F.29})$$

where the first equation holds due to $\bar{\sigma}_{k,h} = \gamma \times D_{\mathcal{F}_h}^{1/2}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h})$ and the last inequality holds due to the definition of Generalized Eluder dimension with $D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) / \bar{\sigma}_{k,h} < 1$.

Combining the results in (F.26), (F.27), (F.28) and (F.29), we have

$$\begin{aligned} & \sum_{k=1}^K \min \left(\beta D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), 1 \right) \\ & \leq (1 + \beta\gamma^2) \dim_{\alpha,K}(\mathcal{F}_h) + 2\beta \sqrt{\dim_{\alpha,K}(\mathcal{F}_h)} \sqrt{\sum_{k=1}^K (\sigma_{k,h}^2 + \alpha^2)}. \end{aligned}$$

Thus, we complete the proof of Lemma D.13. \square

F.4.3 Proof of Lemma D.14

Proof of Lemma D.14. First, for each stage $h \in [H]$ and episode $k \in [K]$, the gap between $V_{k,h}(s_h^k)$ and $V_h^{\pi^k}(s_h^k)$ can be decomposed as:

$$\begin{aligned} & V_{k,h}(s_h^k) - V_h^{\pi^k}(s_h^k) \\ & = Q_{k,h}(s_h^k, a_h^k) - Q_h^{\pi^k}(s_h^k, a_h^k) \\ & \leq \min \left(\hat{f}_{k_{\text{last}},h}(s, a) + b_{k_{\text{last}},h}(s, a), 1 \right) - \mathcal{T}_h V_{k,h+1}(s_h^k, a_h^k) \\ & \quad + \mathcal{T}_h V_{k,h+1}(s_h^k, a_h^k) - \mathcal{T}_h V_{h+1}^{\pi^k}(s_h^k, a_h^k) \\ & \leq [\mathbb{P}_h(V_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) + \min \left(\hat{\beta}_{k_{\text{last}}} D_{\mathcal{F}_h}(z; z_{[k_{\text{last}}-1],h}, \bar{\sigma}_{[k_{\text{last}}-1],h}), 1 \right) + \min(b_{k_{\text{last}},h}(s_h^k, a_h^k), 1) \\ & \leq [\mathbb{P}_h(V_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) + 2C \cdot \min \left(\hat{\beta}_{k_{\text{last}}} D_{\mathcal{F}_h}(z; z_{[k_{\text{last}}-1],h}, \bar{\sigma}_{[k_{\text{last}}-1],h}), 1 \right) \\ & \leq [\mathbb{P}_h(V_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) + 2C(1 + \chi) \cdot \min \left(\hat{\beta}_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), 1 \right) \\ & = V_{k,h+1}(s_{h+1}^k) - V_{h+1}^{\pi^k}(s_{h+1}^k) + [\mathbb{P}_h(V_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) - (V_{k,h+1}(s_{h+1}^k) - V_{h+1}^{\pi^k}(s_{h+1}^k)) \\ & \quad + 2C(1 + \chi) \cdot \min \left(\hat{\beta}_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), 1 \right), \tag{F.30} \end{aligned}$$

where the first inequality holds due to the definition of value function $Q_{k,h}(s_h^k, a_h^k)$, the second inequality holds due to Lemma D.6, the third inequality holds due to $b_{k_{\text{last}},h}(s_h^k, a_h^k) \leq C \cdot D_{\mathcal{F}_h}(z; z_{[k_{\text{last}}-1],h}, \bar{\sigma}_{[k_{\text{last}}-1],h})$ and the last inequality holds due to Lemma G.2. Taking a summation of (F.30) over all episode $k \in [K]$ and stage $h' \geq h$, we have

$$\begin{aligned} & \sum_{k=1}^K (V_{k,h}(s_h^k) - V_h^{\pi^k}(s_h^k)) \\ & \leq \sum_{k=1}^K \sum_{h'=h}^H \left([\mathbb{P}_h(V_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) - (V_{k,h+1}(s_{h+1}^k) - V_{h+1}^{\pi^k}(s_{h+1}^k)) \right) \\ & \quad + \sum_{k=1}^K \sum_{h'=h}^H 2C(1 + \chi) \cdot \min \left(\hat{\beta}_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), 1 \right) \\ & \leq \sum_{k=1}^K \sum_{h'=h}^H 2C(1 + \chi) \cdot \min \left(\hat{\beta}_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), 1 \right) \\ & \quad + 2 \sqrt{\sum_{k=1}^K \sum_{h'=h}^H [\mathbb{V}_h(V_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) \log(2K^2 H/\delta) + 2\sqrt{\log(2K^2 H/\delta)} + 2\log(2K^2 H/\delta)} \\ & \leq \sum_{h'=h}^H 2C(1 + \chi)(1 + \hat{\beta}_k \gamma^2) \dim_{\alpha,K}(\mathcal{F}_{h'}) + \sum_{h'=h}^H 4C(1 + \chi) \hat{\beta}_k \sqrt{\dim_{\alpha,K}(\mathcal{F}_{h'})} \sqrt{\sum_{k=1}^K (\sigma_{k,h'}^2 + \alpha^2)} \\ & \quad + 2 \sqrt{\sum_{k=1}^K \sum_{h'=h}^H [\mathbb{V}_h(V_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) \log(2K^2 H/\delta) + 2\sqrt{\log(2K^2 H/\delta)} + 2\log(2K^2 H/\delta)} \end{aligned}$$

$$\begin{aligned}
&\leq 2CH(1+\chi)(1+\widehat{\beta}_k\gamma^2)\dim_{\alpha,K}(\mathcal{F}_h) + 4C(1+\chi)\widehat{\beta}_k\sqrt{\sum_{h'=h}^H\dim_{\alpha,K}(\mathcal{F}_{h'})}\sqrt{\sum_{k=1}^K\sum_{h'=h}^H(\sigma_{k,h'}^2+\alpha^2)} \\
&\quad + 2\sqrt{\sum_{k=1}^K\sum_{h'=h}^H[\mathbb{V}_h(V_{k,h+1}-V_{h+1}^{\pi^k})](s_h^k,a_h^k)\log(2K^2H/\delta)+2\sqrt{\log(2K^2H/\delta)}+2\log(2K^2H/\delta)} \\
&\leq 2CH(1+\chi)(1+\widehat{\beta}_k\gamma^2)\dim_{\alpha,K}(\mathcal{F}_h) + 4C(1+\chi)\widehat{\beta}_k\sqrt{\dim_{\alpha,K}(\mathcal{F})}\sqrt{H\sum_{k=1}^K\sum_{h=1}^H(\sigma_{k,h}^2+\alpha^2)} \\
&\quad + 2\sqrt{\sum_{k=1}^K\sum_{h'=h}^H[\mathbb{V}_h(V_{k,h+1}-V_{h+1}^{\pi^k})](s_h^k,a_h^k)\log(2K^2H/\delta)+2\sqrt{\log(2K^2H/\delta)}+2\log(2K^2H/\delta)},
\end{aligned} \tag{F.31}$$

where the first inequality holds due to (F.30), the second inequality holds due to event \mathcal{E}_1 , the third inequality holds due to Lemma D.13, the fourth inequality holds due to Cauchy-Schwartz inequality and the last inequality holds due to $\sum_{h'=h}^H\dim_{\alpha,K}(\mathcal{F}_{h'})\leq\sum_{h'=1}^H\dim_{\alpha,K}(\mathcal{F}_{h'})=H\dim_{\alpha,K}(\mathcal{F})$. Furthermore, taking a summation of (F.31), we have

$$\begin{aligned}
&\sum_{k=1}^K\sum_{h=1}^H[\mathbb{P}_h(V_{k,h+1}-V_{h+1}^{\pi^k})](s_h^k,a_h^k) \\
&= \sum_{k=1}^K\sum_{h=1}^H(V_{k,h+1}(s_{h+1}^k)-V_{h+1}^{\pi^k}(s_{h+1}^k)) \\
&\quad + \sum_{k=1}^K\sum_{h=1}^H\left([\mathbb{P}_h(V_{k,h+1}-V_{h+1}^{\pi^k})](s_h^k,a_h^k)-(V_{k,h+1}(s_{h+1}^k)-V_{h+1}^{\pi^k}(s_{h+1}^k))\right) \\
&\leq \sum_{k=1}^K\sum_{h=1}^H(V_{k,h+1}(s_{h+1}^k)-V_{h+1}^{\pi^k}(s_{h+1}^k)) \\
&\quad + 2\sqrt{\sum_{k=1}^K\sum_{h=1}^H[\mathbb{V}_h(V_{k,h+1}-V_{h+1}^{\pi^k})](s_h^k,a_h^k)\log(2K^2H/\delta)+2\sqrt{\log(2K^2H/\delta)}+2\log(2K^2H/\delta)} \\
&\leq 2CH^2(1+\chi)(1+\widehat{\beta}_k\gamma^2)\dim_{\alpha,K}(\mathcal{F}) + 4CH(1+\chi)\widehat{\beta}_k\sqrt{\dim_{\alpha,K}(\mathcal{F}_h)}\sqrt{H\sum_{k=1}^K\sum_{h=1}^H(\sigma_{k,h}^2+\alpha^2)} \\
&\quad + H\left(2\sqrt{\sum_{k=1}^K\sum_{h=1}^H[\mathbb{V}_h(V_{k,h+1}-V_{h+1}^{\pi^k})](s_h^k,a_h^k)\log(2K^2H/\delta)+2\sqrt{\log(2K^2H/\delta)}+2\log(2K^2H/\delta)}\right),
\end{aligned}$$

where the first inequality holds due to event \mathcal{E}_1 and the second inequality holds due to (F.31). Thus, we complete the proof of Lemma D.14. \square

F.4.4 Proof of Lemma D.15

Proof of Lemma D.15. Similar to the proof of Lemma D.14, for each stage $h\in[H]$ and episode $k\in[K]$, the gap between $V_{k,h}(s_h^k)$ and $\check{V}_{k,h}(s_h^k)$ can be decomposed as:

$$\begin{aligned}
&V_{k,h}(s_h^k)-\check{V}_{k,h}(s_h^k) \\
&\leq Q_{k,h}(s_h^k,a_h^k)-\check{Q}_{k,h}(s_h^k,a_h^k) \\
&\leq \min\left(\widehat{f}_{k_{\text{last}},h}(s,a)+b_{k_{\text{last}},h}(s,a),1\right)-\mathcal{T}_hV_{k,h+1}(s_h^k,a_h^k) \\
&\quad -\max\left(\check{f}_{k_{\text{last}},h}(s,a)-b_{k_{\text{last}},h}(s,a),0\right)+\mathcal{T}_h\check{V}_{k,h+1}(s_h^k,a_h^k)
\end{aligned}$$

$$\begin{aligned}
& + \mathcal{T}_h V_{k,h+1}(s_h^k, a_h^k) - \mathcal{T}_h \check{V}_{k,h+1}(s_h^k, a_h^k) \\
\leq & \left[\mathbb{P}_h(V_{k,h+1} - \check{V}_{k,h+1}) \right](s_h^k, a_h^k) + 2 \cdot \min(\widehat{\beta}_{k_{\text{last}}} D_{\mathcal{F}_h}(z; z_{[k_{\text{last}}-1],h}, \bar{\sigma}_{[k_{\text{last}}-1],h}), 1) \\
& + 2 \cdot \min(b_{k_{\text{last}},h}(s_h^k, a_h^k), 1) \\
\leq & \left[\mathbb{P}_h(V_{k,h+1} - \check{V}_{k,h+1}) \right](s_h^k, a_h^k) + 4C \cdot \min(\widehat{\beta}_{k_{\text{last}}} D_{\mathcal{F}_h}(z; z_{[k_{\text{last}}-1],h}, \bar{\sigma}_{[k_{\text{last}}-1],h}), 1) \\
\leq & \left[\mathbb{P}_h(V_{k,h+1} - \check{V}_{k,h+1}) \right](s_h^k, a_h^k) + 4C(1 + \chi) \cdot \min(\widehat{\beta}_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), 1) \\
= & V_{k,h+1}(s_{h+1}^k) - \check{V}_{k,h+1}(s_{h+1}^k) + \left[\mathbb{P}_h(V_{k,h+1} - \check{V}_{k,h+1}) \right](s_h^k, a_h^k) - (V_{k,h+1}(s_{h+1}^k) - \check{V}_{k,h+1}(s_{h+1}^k)) \\
& + 4C(1 + \chi) \cdot \min(\widehat{\beta}_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), 1), \tag{F.32}
\end{aligned}$$

where the first and second inequalities hold due to the definition of $\check{V}_{k,h}(s_h^k)$ and $V_{k,h}(s_h^k)$, the third inequality holds due to Lemma D.6 with $\widehat{\beta}_{k_{\text{last}}} = \check{\beta}_{k_{\text{last}}}$, the fourth inequality holds due to $b_{k_{\text{last}},h}(s_h^k, a_h^k) \leq C \cdot D_{\mathcal{F}_h}(z; z_{[k_{\text{last}}-1],h}, \bar{\sigma}_{[k_{\text{last}}-1],h})$ and the last inequality holds due to Lemma G.2. Taking a summation of (F.32) over all episode $k \in [K]$ and stage $h' \geq h$, we have

$$\begin{aligned}
& \sum_{k=1}^K (V_{k,h}(s_h^k) - \check{V}_{k,h}(s_h^k)) \\
\leq & \sum_{k=1}^K \sum_{h'=h}^H \left(\left[\mathbb{P}_h(V_{k,h+1} - \check{V}_{k,h+1}) \right](s_h^k, a_h^k) - (V_{k,h+1}(s_{h+1}^k) - \check{V}_{k,h+1}(s_{h+1}^k)) \right) \\
& + \sum_{k=1}^K \sum_{h'=h}^H 4C(1 + \chi) \cdot \min(\widehat{\beta}_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), 1) \\
\leq & \sum_{k=1}^K \sum_{h'=h}^H 4C(1 + \chi) \cdot \min(\widehat{\beta}_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}), 1) \\
& + 2 \sqrt{\sum_{k=1}^K \sum_{h'=h}^H [\mathbb{V}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \log(2K^2 H/\delta) + 2\sqrt{\log(2K^2 H/\delta)} + 2\log(2K^2 H/\delta)} \\
\leq & \sum_{h'=h}^H 4C(1 + \chi)(1 + \widehat{\beta}_k \gamma^2) \dim_{\alpha,K}(\mathcal{F}_{h'}) + \sum_{h'=h}^H 8C(1 + \chi) \widehat{\beta}_k \sqrt{\dim_{\alpha,K}(\mathcal{F}_{h'})} \sqrt{\sum_{k=1}^K (\sigma_{k,h'}^2 + \alpha^2)} \\
& + 2 \sqrt{\sum_{k=1}^K \sum_{h'=h}^H [\mathbb{V}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \log(2K^2 H/\delta) + 2\sqrt{\log(2K^2 H/\delta)} + 2\log(2K^2 H/\delta)} \\
\leq & 4CH(1 + \chi)(1 + \widehat{\beta}_k \gamma^2) \dim_{\alpha,K}(\mathcal{F}_h) + 8C(1 + \chi) \widehat{\beta}_k \sqrt{\sum_{h'=h}^H \dim_{\alpha,K}(\mathcal{F}_{h'})} \sqrt{\sum_{k=1}^K \sum_{h'=h}^H (\sigma_{k,h'}^2 + \alpha^2)} \\
& + 2 \sqrt{\sum_{k=1}^K \sum_{h'=h}^H [\mathbb{V}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \log(2K^2 H/\delta) + 2\sqrt{\log(2K^2 H/\delta)} + 2\log(2K^2 H/\delta)} \\
\leq & 4CH(1 + \chi)(1 + \widehat{\beta}_k \gamma^2) \dim_{\alpha,K}(\mathcal{F}_h) + 8C(1 + \chi) \widehat{\beta}_k \sqrt{\dim_{\alpha,K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)} \\
& + 2 \sqrt{\sum_{k=1}^K \sum_{h'=h}^H [\mathbb{V}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \log(2K^2 H/\delta) + 2\sqrt{\log(2K^2 H/\delta)} + 2\log(2K^2 H/\delta)}, \tag{F.33}
\end{aligned}$$

where the first inequality holds due to (F.32), the second inequality holds due to event \mathcal{E}_2 , the third inequality holds due to Lemma D.13, the fourth inequality holds due to Cauchy-Schwartz inequality and the last inequality holds due to $\sum_{h'=h}^H \dim_{\alpha,K}(\mathcal{F}_{h'}) \leq \sum_{h'=1}^H \dim_{\alpha,K}(\mathcal{F}_{h'}) = H \dim_{\alpha,K}(\mathcal{F})$.

Furthermore, taking a summation of (F.33), we have

$$\begin{aligned}
& \sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \\
&= \sum_{k=1}^K \sum_{h=1}^H (V_{k,h+1}(s_{h+1}^k) - \check{V}_{k,h+1}(s_{h+1}^k)) \\
&\quad + \sum_{k=1}^K \sum_{h=1}^H \left([\mathbb{P}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) - (V_{k,h+1}(s_{h+1}^k) - \check{V}_{k,h+1}(s_{h+1}^k)) \right) \\
&\leq \sum_{k=1}^K \sum_{h=1}^H (V_{k,h+1}(s_{h+1}^k) - \check{V}_{k,h+1}(s_{h+1}^k)) \\
&\quad + H \left(2 \sqrt{\sum_{k=1}^K \sum_{h=1}^H [\mathbb{V}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \log(2k^2 H/\delta) + 2\sqrt{\log(2K^2 H/\delta)} + 2\log(2K^2 H/\delta)} \right) \\
&\leq 2CH^2(1 + \chi)(1 + \hat{\beta}_k \gamma^2) \dim_{\alpha, K}(\mathcal{F}_h) + 4CH(1 + \chi)\hat{\beta}_k \sqrt{\dim_{\alpha, K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)} \\
&\quad + 2H \left(2 \sqrt{\sum_{k=1}^K \sum_{h=1}^H [\mathbb{V}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \log(2K^2 H/\delta) + 2\sqrt{\log(2K^2 H/\delta)} + 2\log(2K^2 H/\delta)} \right),
\end{aligned}$$

where the first inequality holds due to event \mathcal{E}_2 and the last inequality holds due to (F.33). Thus, we complete the proof of Lemma D.15. \square

F.4.5 Proof of Lemma D.16

Proof of Lemma D.16. According to the definition of estimated variance $\sigma_{k,h}$, the summation of variance can be decomposed as following:

$$\begin{aligned}
\sum_{k=1}^K \sum_{h=1}^H \sigma_{k,h}^2 &= \sum_{k=1}^K \sum_{h=1}^H [\bar{\mathbb{V}}_{k,h} V_{k,h+1}](s_h^k, a_h^k) + E_{k,h} + F_{k,h} \\
&= \underbrace{\sum_{k=1}^K \sum_{h=1}^H ([\bar{\mathbb{V}}_{k,h} V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{V}_h V_{k,h+1}](s_h^k, a_h^k))}_{I_1} + \underbrace{\sum_{k=1}^K \sum_{h=1}^H E_{k,h}}_{I_2} + \underbrace{\sum_{k=1}^K \sum_{h=1}^H F_{k,h}}_{I_3} \\
&\quad + \underbrace{\sum_{k=1}^K \sum_{h=1}^H ([\mathbb{V}_h V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{V}_h V_{h+1}^{\pi^k}](s_h^k, a_h^k))}_{I_4} + \underbrace{\sum_{k=1}^K \sum_{h=1}^H [\mathbb{V}_h V_{h+1}^{\pi^k}](s_h^k, a_h^k)}_{I_5}.
\end{aligned} \tag{F.34}$$

For the term I_1 , it can be upper bounded by

$$I_1 = \sum_{k=1}^K \sum_{h=1}^H ([\bar{\mathbb{V}}_{k,h} V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{V}_h V_{k,h+1}](s_h^k, a_h^k)) \leq \sum_{k=1}^K \sum_{h=1}^H E_{k,h}, \tag{F.35}$$

where the inequality holds due to Lemma D.9.

For the second term $I_2 = \sum_{k=1}^K \sum_{h=1}^H E_{k,h}$, we have

$$\sum_{k=1}^K \sum_{h=1}^H E_{k,h} = \sum_{k=1}^K \sum_{h=1}^H (2L\beta_k + \tilde{\beta}_k) \min\left(1, D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h})\right)$$

$$\begin{aligned}
&\leq \sum_{h=1}^H (2L\beta_k + \tilde{\beta}_k) \times (1 + \gamma^2) \dim_{\alpha, K}(\mathcal{F}_h) \\
&\quad + \sum_{h=1}^H (2L\beta_k + \tilde{\beta}_k) \times 2\sqrt{\dim_{\alpha, K}(\mathcal{F}_h)} \sqrt{\sum_{k=1}^K (\sigma_{k,h}^2 + \alpha^2)} \\
&\leq (2L\beta_k + \tilde{\beta}_k)H \times (1 + \gamma^2) \dim_{\alpha, K}(\mathcal{F}) \\
&\quad + (2L\beta_k + \tilde{\beta}_k) \times 2\sqrt{\sum_{h=1}^H \dim_{\alpha, K}(\mathcal{F}_h)} \sqrt{\sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)} \\
&= (2L\beta_k + \tilde{\beta}_k)H \times (1 + \gamma^2) \dim_{\alpha, K}(\mathcal{F}) \\
&\quad + (2L\beta_k + \tilde{\beta}_k) \times 2\sqrt{\dim_{\alpha, K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)}, \tag{F.36}
\end{aligned}$$

where the inequality holds due to Lemma D.13 and the second inequality holds due to Cauchy-Schwartz inequality.

For the term I_3 , we have

$$\begin{aligned}
I_3 &= \sum_{k=1}^K \sum_{h=1}^H F_{k,h} \\
&= (\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)) \\
&\quad \times \sum_{k=1}^K \sum_{h=1}^H \min \left(1, 2\hat{f}_{k,h}(s_h^k, a_h^k) - 2\check{f}_{k,h}(s_h^k, a_h^k) + 4\beta_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) \right) \\
&\leq (\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)) \\
&\quad \times \sum_{k=1}^K \sum_{h=1}^H \min \left(1, 2\mathcal{T}_h V_{k,h+1}(s_h^k, a_h^k) - 2\mathcal{T}_h \check{V}_{k,h+1}(s_h^k, a_h^k)(s_h^k, a_h^k) + 8\beta_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) \right) \\
&\leq (\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)) \times \sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \\
&\quad + (\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)) \times \sum_{k=1}^K \sum_{h=1}^H \min \left(1, 8\beta_k D_{\mathcal{F}_h}(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) \right) \\
&\leq (\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)) \times \sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \\
&\quad + (\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)) \times (1 + 8\beta_k \gamma^2) H \dim_{\alpha, K}(\mathcal{F}) \\
&\quad + ((\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)) \times 16\beta_k \sqrt{\dim_{\alpha, K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)}) \tag{F.37}
\end{aligned}$$

where the first inequality holds due to Lemma D.8, the second inequality holds due to $V_{k,h+1}(\cdot) \geq V_{h+1}^*(\cdot) \geq \check{V}_{k,h+1}(\cdot)$, the third inequality holds due to Lemma D.13 with Cauchy-Schwartz inequality.

By Lemma D.15,

$$\begin{aligned}
&\sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \leq 4CH^2(1 + \chi)(1 + \hat{\beta}_k \gamma^2) \dim_{\alpha, K}(\mathcal{F}_h) \\
&\quad + 8CH(1 + \chi)\hat{\beta}_k \sqrt{\dim_{\alpha, K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)}
\end{aligned}$$

$$\begin{aligned}
& + 2H \left(2 \sqrt{\sum_{k=1}^K \sum_{h=1}^H [\mathbb{V}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \log(2K^2H/\delta) + 2\sqrt{\log(2K^2H/\delta)} + 2\log(2K^2H/\delta)} \right) \\
& \leq 4CH^2(1+\chi)(1+\hat{\beta}_k\gamma^2) \dim_{\alpha,K}(\mathcal{F}_h) + 8CH(1+\chi)\hat{\beta}_k \sqrt{\dim_{\alpha,K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)} \\
& \quad + \tilde{O} \left(H \sqrt{\sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k)} \right), \tag{F.38}
\end{aligned}$$

where the last inequality follows from Lemma D.7.

Notice that for each variable x , $x \leq a\sqrt{x} + b$ implies $x \leq a^2 + 2b$, from (F.38), we further have

$$\begin{aligned}
& \sum_{k=1}^K \sum_{h=1}^H [\mathbb{P}_h(V_{k,h+1} - \check{V}_{k,h+1})](s_h^k, a_h^k) \\
& \leq \tilde{O} \left(CH^2(1+\chi)(1+\hat{\beta}_k\gamma^2) \dim_{\alpha,K}(\mathcal{F}_h) + CH(1+\chi)\hat{\beta}_k \sqrt{\dim_{\alpha,K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)} \right). \tag{F.39}
\end{aligned}$$

Substituting (F.39) into (F.37), we obtain the following upper bound for I_3 ,

$$\begin{aligned}
I_3 & \leq \tilde{O} \left((\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)) H \hat{\beta}_k \sqrt{\dim_{\alpha,K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)} \right) \\
& \quad + \tilde{O} \left((\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)) \cdot H^2(1 + \hat{\beta}_k\gamma^2) \dim_{\alpha,K}(\mathcal{F}_h) \right). \tag{F.40}
\end{aligned}$$

For the term I_4 , we have

$$\begin{aligned}
I_4 & = \sum_{k=1}^K \sum_{h=1}^H ([\mathbb{V}_h V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{V}_h V_{h+1}^{\pi^k}](s_h^k, a_h^k)) \\
& = \sum_{k=1}^K \sum_{h=1}^H \left([\mathbb{P}_h V_{k,h+1}^2](s_h^k, a_h^k) - ([\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k))^2 - [\mathbb{P}_h (V_{h+1}^{\pi^k})^2](s_h^k, a_h^k) + ([\mathbb{P}_h V_{h+1}^{\pi^k}](s_h^k, a_h^k))^2 \right) \\
& \leq \sum_{k=1}^K \sum_{h=1}^H \left([\mathbb{P}_h V_{k,h+1}^2](s_h^k, a_h^k) - [\mathbb{P}_h (V_{h+1}^{\pi^k})^2](s_h^k, a_h^k) \right) \\
& \leq 2 \sum_{k=1}^K \sum_{h=1}^H \left([\mathbb{P}_h V_{k,h+1}](s_h^k, a_h^k) - [\mathbb{P}_h V_{h+1}^{\pi^k}](s_h^k, a_h^k) \right) \\
& \leq 8CH^2(1+\chi)(1+\hat{\beta}_k\gamma^2) \dim_{\alpha,K}(\mathcal{F}) + 16CH(1+\chi)\hat{\beta}_k \sqrt{\dim_{\alpha,K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)} \\
& \quad + 2H \left(2 \sqrt{\sum_{k=1}^K \sum_{h=1}^H [\mathbb{V}_h(V_{k,h+1} - V_{h+1}^{\pi^k})](s_h^k, a_h^k) \log(2K^2H/\delta) + 2\sqrt{\log(2K^2H/\delta)} + 2\log(2K^2H/\delta)} \right) \\
& \leq \tilde{O} \left(H^2(1 + \hat{\beta}_k\gamma^2) \dim_{\alpha,K}(\mathcal{F}) + H \hat{\beta}_k \sqrt{\dim_{\alpha,K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)} \right) \tag{F.41}
\end{aligned}$$

where the first inequality holds due to $V_{k,h+1}(\cdot) \geq V_{h+1}^*(\cdot) \geq V_{h+1}^{\pi^k}(\cdot)$, the second inequality holds due to $0 \leq V_{h+1}^*(\cdot), V_{h+1}^{\pi^k}(\cdot) \leq 1$, the third inequality holds due to Lemma D.14, and the last inequality follows from Lemma D.7 and the fact that $x \leq a\sqrt{x} + b$ implies $x \leq a^2 + 2b$.

For the term I_5 , according to the definition of Var_K , we have

$$I_5 = \sum_{k=1}^K \sum_{h=1}^H [\mathbb{V}_h V_{h+1}^{\pi^k}](s_h^k, a_h^k) = \text{Var}_K. \quad (\text{F.42})$$

Substituting the results in (F.35), (F.36), (F.40), (F.41) and (F.42) into (F.34), we have

$$\begin{aligned} & \sum_{k=1}^K \sum_{h=1}^H \sigma_{k,h}^2 \\ &= I_1 + I_2 + I_3 + I_4 + I_5 \\ &\leq \tilde{O} \left((4L\beta_k + 2\tilde{\beta}_k)H \times (1 + \gamma^2) \dim_{\alpha,K}(\mathcal{F}) \right. \\ &\quad + (4L\beta_k + 2\tilde{\beta}_k) \times 2\sqrt{\dim_{\alpha,K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2) + \text{Var}_K} \\ &\quad + 8CH^2(1 + \chi)(1 + \hat{\beta}_k\gamma^2) \dim_{\alpha,K}(\mathcal{F}) + 16CH(1 + \chi)\hat{\beta}_k \sqrt{\dim_{\alpha,K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)} \\ &\quad + (\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)) H \hat{\beta}_k \sqrt{\dim_{\alpha,K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)} \\ &\quad \left. + (\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)) \cdot H^2(1 + \hat{\beta}_k\gamma^2) \dim_{\alpha,K}(\mathcal{F}_h) \right) \\ &\leq \text{Var}_K + (\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)) \times \tilde{O}((1 + \gamma^2)(\beta_k + H\hat{\beta}_k + \tilde{\beta}_k)H \dim_{\alpha,K}(\mathcal{F})) \\ &\quad + (\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)) \times \tilde{O} \left((\beta_k + H\hat{\beta}_k + \tilde{\beta}_k) \sqrt{\dim_{\alpha,K}(\mathcal{F})} \sqrt{H \sum_{k=1}^K \sum_{h=1}^H (\sigma_{k,h}^2 + \alpha^2)} \right) \end{aligned}$$

Notice that for each variable x , $x \leq a\sqrt{x} + b$ implies $x \leq a^2 + 2b$. With this fact, we have

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \sigma_{k,h}^2 &\leq (\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K)) \times \tilde{O}((1 + \gamma^2)(\beta_k + H\hat{\beta}_k + \tilde{\beta}_k)H \dim_{\alpha,K}(\mathcal{F})) \\ &\quad + (\log N_{\mathcal{F}}(\epsilon) + \log N_{\epsilon}(K))^2 \times \tilde{O}((\beta_k + H\hat{\beta}_k + \tilde{\beta}_k)^2 H \dim_{\alpha,K}(\mathcal{F})) \\ &\quad + \tilde{O}(\text{Var}_K + KH\alpha^2). \end{aligned}$$

Thus, we complete the proof of Lemma D.16. \square

G Covering Number Argument

G.1 Rare Switching

Based on the policy-updating criterion, the following lemma provides an upper bound of the switching cost.

Lemma G.1. The number of episodes when the algorithm updates the value function is at most $O(\dim_{\alpha,K}(\mathcal{F}) \cdot H)$.

Proof. According to line 9, the policy is updated at episode k only when there exists a stage $h \in [H]$ such that

$$\sum_{i \in [k_{last}, k-1]} \frac{1}{\bar{\sigma}_{i,h}^2} D_{\mathcal{F}_h}^2(z_{i,h}; z_{[k_{last}-1],h}, \bar{\sigma}_{[k_{last}-1],h}) \geq \chi/C.$$

and

$$\sum_{i \in [k_{last}, k-2]} \frac{1}{\bar{\sigma}_{i,h}^2} D_{\mathcal{F}}^2(z_{i,h}; z_{[k_{last}-1],h}, \bar{\sigma}_{[k_{last}-1],h}) < \chi. \quad (\text{G.1})$$

Then the following inequality holds,

$$\sup_{f_1, f_2 \in \mathcal{F}_h} \frac{\sum_{i \in [1, k-2]} \frac{1}{\bar{\sigma}_{i,h}^2} (f_1(z_{i,h}) - f_2(z_{i,h}))^2 + \lambda}{\sum_{i \in [1, k_{last}-1]} \frac{1}{\bar{\sigma}_{i,h}^2} (f_1(z_{i,h}) - f_2(z_{i,h}))^2 + \lambda} \quad (\text{G.2})$$

$$\begin{aligned} &= 1 + \sup_{f_1, f_2 \in \mathcal{F}_h} \frac{\sum_{i \in [k_{last}, k-2]} \frac{1}{\bar{\sigma}_{i,h}^2} (f_1(z_{i,h}) - f_2(z_{i,h}))^2}{\sum_{i \in [1, k_{last}-1]} \frac{1}{\bar{\sigma}_{i,h}^2} (f_1(z_{i,h}) - f_2(z_{i,h}))^2 + \lambda} \\ &\leq 1 + \sum_{i \in [k_{last}, k-2]} \frac{1}{\bar{\sigma}_{i,h}^2} D_{\mathcal{F}_h}^2(z_{i,h}; z_{[k_{last}-1],h}, \bar{\sigma}_{[k_{last}-1],h}) \\ &\leq 1 + \chi, \end{aligned} \quad (\text{G.3})$$

where the first inequality holds due to the definition of $D_{\mathcal{F}_h}$ (Definition 2.4), the second inequality follows from (G.1).

(G.3) further gives a lower bound for the summation

$$\begin{aligned} &\sum_{i \in [k_{last}, k-1]} \frac{1}{\bar{\sigma}_{i,h}^2} D_{\mathcal{F}}^2(z_{i,h}; z_{[i-1],h}, \bar{\sigma}_{[i-1],h}) \\ &\geq \frac{1}{1 + \chi} \sum_{i \in [k_{last}, k-1]} \frac{1}{\bar{\sigma}_{i,h}^2} D_{\mathcal{F}}^2(z_{i,h}; z_{[k_{last}-1],h}, \bar{\sigma}_{[k_{last}-1],h}) \\ &\geq \frac{\chi/C}{1 + \chi}. \end{aligned}$$

Note that $\frac{\chi/C}{1 + \chi} \leq 1$, we also have

$$\sum_{i \in [k_{last}, k-1]} \min \left\{ 1, \frac{1}{\bar{\sigma}_{i,h}^2} D_{\mathcal{F}}^2(z_{i,h}; z_{[i-1],h}, \bar{\sigma}_{[i-1],h}) \right\} \geq \frac{\chi/C}{1 + \chi}.$$

Then we have an upper bound and lower bound for the following summation:

$$l_K \cdot \frac{\chi/C}{1 + \chi} \leq \sum_{k=1}^K \sum_{h=1}^H \min \left\{ 1, \frac{1}{\bar{\sigma}_{k,h}^2} D_{\mathcal{F}}^2(z_{k,h}; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) \right\} \leq \dim_{\alpha, K}(\mathcal{F}) \cdot H.$$

Therefore, the number of policy switching l_K is of order $O(\dim_{\alpha, K}(\mathcal{F}) \cdot H)$. \square

Lemma G.2 (Stability of uncertainty under rare switching strategy). If the policy is not updated at episode k , the uncertainty of all state-action pair $z = (s, a) \in \mathcal{S} \times \mathcal{A}$ and stage $h \in [H]$ satisfies the following stability property:

$$D_{\mathcal{F}_h}^2(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) \geq \frac{1}{1 + \chi} D_{\mathcal{F}_h}^2(z; z_{[k_{last}-1],h}, \bar{\sigma}_{[k_{last}-1],h}).$$

Proof. Due to the definition of k_{last} in Algorithm 1, we have

$$\sum_{i \in [k_{last}, k-1]} \frac{1}{\bar{\sigma}_{i,h}^2} D_{\mathcal{F}}^2(z_{i,h}; z_{[k_{last}-1],h}, \bar{\sigma}_{[k_{last}-1],h}) < \chi.$$

As is shown in (G.3), here we also have

$$\sup_{f_1, f_2 \in \mathcal{F}_h} \frac{\sum_{i \in [1, k-1]} \frac{1}{\bar{\sigma}_{i,h}^2} (f_1(z_{i,h}) - f_2(z_{i,h}))^2 + \lambda}{\sum_{i \in [1, k_{last}-1]} \frac{1}{\bar{\sigma}_{i,h}^2} (f_1(z_{i,h}) - f_2(z_{i,h}))^2 + \lambda} \leq 1 + \chi.$$

From the definition of $D_{\mathcal{F}_h}$,

$$D_{\mathcal{F}_h}^2(z; z_{[k-1],h}, \bar{\sigma}_{[k-1],h}) \geq \frac{1}{1+\chi} D_{\mathcal{F}_h}^2(z; z_{[k_{last}-1],h}, \bar{\sigma}_{[k_{last}-1],h}).$$

The proof is then completed due to the arbitrariness of h . \square

G.2 Value Function Class and Its Covering Number

The optimistic value functions at episode k and stage $h \in [H]$ in our construction belong to the following function class:

$$\mathcal{V}_{k,h} = \left\{ V \left| \max_{a \in \mathcal{A}} \min_{1 \leq i \leq l_k+1} \min(1, f_i(\cdot, a) + \beta \cdot b(\cdot, a)) \right. \right\}, \quad (\text{G.4})$$

where l_k is the number of updated policies as defined in Algorithm 1, $f_i \in \mathcal{F}_h$ and $b \in \mathcal{B}$.

Similarly, we also define the following pessimistic value function classes for all $k \geq 1$:

$$\check{\mathcal{V}}_{k,h} = \left\{ V \left| \max_{a \in \mathcal{A}} \max_{1 \leq i \leq l_k+1} \max(0, f_i(\cdot, a) - \beta \cdot b(\cdot, a)) \right. \right\}, \quad (\text{G.5})$$

Lemma G.3 (ϵ -covering number of optimistic value function classes). For optimistic value function class $\mathcal{V}_{k,h}$ defined in (G.4), we define the distance between two value functions V_1 and V_2 as $\|V_1 - V_2\|_\infty := \max_{s \in \mathcal{S}} |V_1(s) - V_2(s)|$. Then the ϵ -covering number with respect to the distance function can be upper bounded by

$$N_\epsilon(k) := [N_{\mathcal{F}}(\epsilon/2) \cdot N(\mathcal{B}, \epsilon/2\beta)]^{l_k+1}. \quad (\text{G.6})$$

Proof. By the definition of $N(\mathcal{F}, \epsilon)$, there exists an $\epsilon/2$ -net of \mathcal{F} , denoted by $\mathcal{C}(\mathcal{F}, \epsilon/2)$, such that for any $f \in \mathcal{F}$, we can find $f' \in \mathcal{C}(\mathcal{F}, \epsilon/2)$ such that $\|f - f'\|_\infty \leq \epsilon/2$. Also, there exists an $\epsilon/2\beta$ -net of \mathcal{B} , $\mathcal{C}(\mathcal{B}, \epsilon/2\beta)$.

Then we consider the following subset of \mathcal{V}_k ,

$$\mathcal{V}^c = \left\{ V \left| \max_{a \in \mathcal{A}} \min_{1 \leq i \leq l_k+1} \min(1, f_i(\cdot, a) + \beta \cdot b_i(\cdot, a)), f_i \in \mathcal{C}(\mathcal{F}_h, \epsilon/2), b_i \in \mathcal{C}(\mathcal{B}, \epsilon/2\beta) \right. \right\}.$$

Consider an arbitrary $V \in \mathcal{V}$ where $V = \max_{a \in \mathcal{A}} \min_{1 \leq i \leq l_k+1} \min(1, f_i(\cdot, a) + \beta \cdot b_i(\cdot, a))$. For each f_i , there exists $f_i^c \in \mathcal{C}(\mathcal{F}, \epsilon/2)$ such that $\|f_i - f_i^c\|_\infty \leq \epsilon/2$. There also exists $b^c \in \mathcal{C}(\mathcal{B}, \epsilon/2\beta)$ such that $\|b - b^c\|_\infty \leq \epsilon/2\beta$. Let $V^c = \max_{a \in \mathcal{A}} \min_{1 \leq i \leq l_k+1} \min(1, f_i^c(\cdot, a) + \beta \cdot b^c(\cdot, a)) \in \mathcal{V}^c$. It is then straightforward to check that $\|V - V^c\|_\infty \leq \epsilon/2 + \beta \cdot \epsilon/2\beta = \epsilon$.

By direct calculation, we have $|\mathcal{V}^c| = [N(\mathcal{F}_h, \epsilon/2) \cdot N(\mathcal{B}, \epsilon/2\beta)]^{l_k+1}$. \square

Lemma G.4 (ϵ -covering number of pessimistic value function classes). For pessimistic value function class $\check{\mathcal{V}}_{k,h}$ defined in (G.5), we define the distance between two value functions V_1 and V_2 as $\|V_1 - V_2\|_\infty := \max_{s \in \mathcal{S}} |V_1(s) - V_2(s)|$. Then the ϵ -covering number of $\check{\mathcal{V}}_k$ with respect to the distance function can be upper bounded by $N_\epsilon(k)$ defined in (G.6).

Proof. The proof is nearly the same as that of Lemma G.3. \square

H Auxiliary Lemmas

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

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

Lemma H.2 (Corollary 2, Agarwal et al. 2022). Let $M > 0, V > v > 0$ be constants, and $\{x_i\}_{i \in [t]}$ be stochastic process adapted to a filtration $\{\mathcal{H}_i\}_{i \in [t]}$. Suppose $\mathbb{E}[x_i | \mathcal{H}_{i-1}] = 0$, $|x_i| \leq M$ and $\sum_{i \in [t]} \mathbb{E}[x_i^2 | \mathcal{H}_{i-1}] \leq V^2$ almost surely. Then for any $\delta, \epsilon > 0$, let $\iota = \sqrt{\log \frac{(2 \log(V/v)+2) \cdot (\log(M/m)+2)}{\delta}}$ we have

$$\mathbb{P} \left(\sum_{i \in [t]} x_i > \iota \sqrt{2 \left(2 \sum_{i \in [t]} \mathbb{E}[x_i^2 | \mathcal{H}_{i-1}] + v^2 \right)} + \frac{2}{3} \iota^2 \left(2 \max_{i \in [t]} |x_i| + m \right) \right) \leq \delta.$$

Lemma H.3 (Lemma 7, Russo and Van Roy 2014). Consider random variables $(Z_n | n \in \mathbb{N})$ adapted to the filtration $(\mathcal{H}_n : n = 0, 1, \dots)$. Assume $\mathbb{E}[\exp\{\lambda Z_i\}]$ is finite for all λ . Define the conditional mean $\mu_i = \mathbb{E}[Z_i | \mathcal{H}_{i-1}]$. We define the conditional cumulant generating function of the centered random variable $[Z_i - \mu_i]$ by $\psi_i(\lambda) = \log \mathbb{E}[\exp(\lambda [Z_i - \mu_i]) | \mathcal{H}_{i-1}]$. For all $x \geq 0$ and $\lambda \geq 0$,

$$\mathbb{P} \left(\sum_1^n \lambda Z_i \leq x + \sum_1^n [\lambda \mu_i + \psi_i(\lambda)] \quad \forall n \in \mathbb{N} \right) \geq 1 - e^{-x}.$$

Lemma H.4 (Self-normalized bound for scalar-valued martingales). Consider random variables $(v_n | n \in \mathbb{N})$ adapted to the filtration $(\mathcal{H}_n : n = 0, 1, \dots)$. Let $\{\eta_i\}_{i=1}^\infty$ be a sequence of real-valued random variables which is \mathcal{H}_{i+1} -measurable and is conditionally σ -sub-Gaussian. Then for an arbitrarily chosen $\lambda > 0$, for any $\delta > 0$, with probability at least $1 - \delta$, it holds that

$$\sum_{i=1}^n \epsilon_i v_i \leq \frac{\lambda \sigma^2}{2} \cdot \sum_{i=1}^n v_i^2 + \log(1/\delta)/\lambda \quad \forall n \in \mathbb{N}.$$

Lemma H.5 (Lemma 10, Zhang et al. 2020). Let $(M_n)_{n \geq 0}$ be a martingale such that $M_0 = 0$ and $|M_n - M_{n-1}| \leq c$ for some $c > 0$ and any $n \geq 1$. Let $\text{Var}_n = \sum_{k=1}^n E[(M_k - M_{k-1})^2 | \mathcal{F}_{k-1}]$ for $n \geq 0$, where $\mathcal{F}_k = \sigma(M_1, M_2, \dots, M_k)$. Then for any positive integer n and any $\epsilon, p > 0$, we have that

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

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: In both abstract and introduction, we highlight the contribution in our paper. The proposed algorithm and the corresponding theoretical results are discussed in the followed sections.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We explicitly list all the necessary assumptions for our theoretical analysis.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We state all the assumptions in Section 2 and give the proof of theorems in the appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [NA]

Justification: The paper does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification: Paper does not include experiments requiring code.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.

- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [NA]

Justification: The paper does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: The paper does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [NA]

Justification: The paper does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.

- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: The research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: The paper is a theoretical work with no societal impact.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We have described the related works, especially those work which our work is based on with proper citations in corresponding sections.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [No]

Justification: This is a theoretical paper without experiments.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This paper does not include crowdsourcing or human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This paper does not include crowdsourcing or human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.