

Leveraging Unlabeled Data Sharing through Kernel Function Approximation in Offline Reinforcement Learning

Yen-Ru Lai

Graduate Institute of Communication Engineering, National Taiwan University

r09942079@ntu.edu.tw

Fu-Chieh Chang

*Graduate Institute of Communication Engineering, National Taiwan University
Mediatek Research*

d09942015@ntu.edu.tw

Pei-Yuan Wu

Graduate Institute of Communication Engineering, National Taiwan University

peiyuanwu@ntu.edu.tw

Reviewed on OpenReview: <https://openreview.net/forum?id=78N9tCL6Ly>

Abstract

Offline reinforcement learning (RL) learns policies from a fixed dataset, but often requires large amounts of data. The challenge arises when labeled datasets are expensive, especially when rewards have to be provided by human labelers for large datasets. In contrast, unlabelled data tends to be less expensive. This situation highlights the importance of finding effective ways to use unlabelled data in offline RL, especially when labelled data is limited or expensive to obtain. In this paper, we present the algorithm to utilize the unlabeled data in the offline RL method with kernel function approximation and give the theoretical guarantee. We present various eigenvalue decay conditions of the RKHS \mathcal{H}_k induced by kernel k which determine the complexity of the algorithm. In summary, our work provides a promising approach for exploiting the advantages offered by unlabeled data in offline RL, whilst maintaining theoretical assurances.

1 Introduction

Reinforcement learning (RL) algorithms have demonstrated empirical success in a variety of domains, including the defeat of Go champions (Silver et al., 2016), robot control (Kalashnikov et al., 2018), and the development of large language models such as ChatGPT (Stiennon et al., 2020). In particular, these achievements are largely associated with online reinforcement learning, characterized by dynamic data collection. However, the widespread adoption of online RL faces significant challenges. In many scenarios, active exploration is impractical due to factors such as the high cost of data collection (Levine et al., 2020). In many scenarios, active exploration is impractical due to factors such as the high cost of data collection (Levine et al., 2020). To this end, in this paper we explore offline reinforcement learning - a fully data-driven framework similar to supervised learning. Unfortunately, fully data-driven offline RL demands large datasets. In more realistic scenarios, offline reinforcement learning (RL) could allow us to use a smaller amount of task-specific data along with a significant amount of task-agnostic data. This data is not labeled with task rewards, and some of it may not be directly relevant to the task at hand.

Prior works use learned classifiers that discriminate between successes and failures for reward labeling (Fu et al., 2018; Singh et al., 2019) in the online RL setting. However, these approaches are unsuitable for the offline RL setting since they require real-time interaction. Alternatively, some research focuses on learning from data without explicit reward labels by directly imitating expert trajectories (Ho & Ermon,

2016; Kostrikov et al., 2019) or deriving the reward function through inverse reinforcement learning using an expert dataset (Fu et al., 2017; Finn et al., 2016). However, in real-world scenarios, these approaches may face challenges due to the resource-intensive and costly nature of the expert trajectory acquisition and reward labelling process.

Yu et al. (2022) has revealed the challenges associated with learning to predict rewards, highlighting the surprising efficacy of setting the reward to zero. Despite these findings, the impact of reward prediction methods on performance and the potential demonstrable benefits of reward-free data in offline reinforcement learning (RL) remain unclear. In response to this, Hu et al. (2023) have introduced a novel model-free approach named Provable Data Sharing (PDS). PDS incorporates uncertainty penalties into the learned reward functions, maintaining a conservative algorithm. This method allows PDS to take advantage of unlabeled data for offline RL, especially in linear MDPs. However, the linear MDP assumption is inflexible and rarely is fulfilled in practice. This question naturally arises.

How can we enhance the performance of offline RL algorithms that use kernel function approximation by effectively using reward-free data?

This work focuses on the episodic Markov decision process (MDP). The reward function and value function are both represented by kernel functions. Inspired by the Provable Data Sharing (PDS) (Hu et al., 2023) framework, we propose a new algorithm. The PDS algorithm has two main components. First, it pessimistically estimates rewards by applying additional penalties to the reward function learned from labeled data. This augmentation is designed to prevent overestimation, thus ensuring a conservative algorithm. The second part of the PDS algorithm uses the Pessimistic Value Iteration (PEVI) algorithm introduced by Jin et al. (2021) to derive the policy. Our main contribution is that

- **Extension of PDS framework:** We expand the applicability of the Provable Data Sharing (PDS) framework, initially introduced by Hu et al. (2023). This extension goes beyond the original linear Markov Decision Process (MDP) setting, incorporating kernel function approximation. This expansion enhances the versatility of the PDS framework, making it applicable to a broader range of scenarios. Our derivation is influenced by methodologies proposed for kernelized contextual bandits (Chowdhury & Gopalan, 2017; Valko et al., 2013; Srinivas et al., 2009), as well as techniques such as pessimistic value iteration (PEVI) (Jin et al., 2021) and the kernel optimum least squares value iteration algorithm (KOVI) (Yang et al., 2020).
- **Focus on finite-horizon MDPs:** While Hu et al. (2023) focuses on discounted infinite-horizon MDPs, our work shifts attention to finite-horizon MDPs, which introduce unique challenges due to their horizon-dependent reward and transition functions. Unlike the discounted setting, where strong coverage assumptions such as finite concentrability coefficients are often required, our framework does not rely on such uniform data coverage.
- **Feature coverage assessment via concentrability coefficient:** In contrast to Hu et al. (2023), which relies on a bounded concentrability coefficient to assess coverage over the state-action space, we introduce a global coverage assumption (Assumption 4.7) based on the spectrum of feature covariance matrices. This approach ((Wang et al., 2020a)), widely used in supervised learning, is particularly well-suited for settings with linear function approximation. Unlike single-policy coverage, our assumption requires sufficient data coverage across all policies in the considered class, ensuring broader applicability.
- **Enhance the suboptimality:** By employing the data-splitting technique discussed by Xie et al. (2021), the suboptimality can be enhanced by a factor of \sqrt{d} , which depends on the choice of kernel. This enhancement comes at the cost of \sqrt{H} , a constant inherent in the MDP. Nevertheless, with an appropriate selection of kernel, the overall algorithmic performance can be significantly improved.

Our research provides a theoretical guarantee for effectively utilizing the benefits of reward-free data in offline RL. We aim to enhance the robustness of offline RL methods by maintaining theoretical guarantees, which offers a valuable contribution to the ongoing development of more resilient and efficient RL frameworks.

2 Related Works

The issue of suboptimality in discounted and episodic MDP with a model has been considered in linear and kernel settings. The results are presented in Table 1. In the episodic MDP setting, we have the dataset with N trajectories of horizon H , and the suboptimality dependent on N and H . On the other hand, in a discounted MDP setting, we have the dataset with length N , and suboptimality dependent on N . The PEVI algorithm (Jin et al., 2021) serves as the foundational algorithm within Hu et al. (2023) and our work. If we assume that the infinite horizon MDP should conclude within H steps (referred to as the effective horizon) (Yan et al., 2022), we can set the discount factor γ such that $H = 1/(1 - \gamma)$. Consequently, the suboptimality for the PDS algorithm is expressed as $\tilde{O}(dH^2N_2^{-\frac{1}{2}})$ where N_2 is the number of trajectories for the unlabeled dataset. Similar to Hu et al. (2023), we incorporate unsupervised data sharing to enhance the offline RL algorithm. The linear setting is a special case of the kernel setting with a linear kernel. In this case, we can recover the suboptimality as $\tilde{O}(Hd^{\frac{1}{2}}N_1^{-\frac{1}{2}})$, where N_1 is the number of trajectories for the labeled dataset, as provided in Hu et al. (2023). A notable difference between PEVI and PDS lies in PDS’s utilization of data sharing to improve the suboptimality through an unlabeled dataset. It’s important to note that $N_2 > N_1$ in general. When comparing PDS with our approach in a linear setting, the H -folds data splitting in our algorithm enhances the suboptimality by a factor of \sqrt{d} . However, this improvement comes with a tradeoff, as our algorithm introduces a suboptimality increment by a factor of \sqrt{H} because we need to partition the data set into H folds. As a result, each estimated value function is derived from only N_2/H episodes of data.

Algorithm	MDP	Setting	SubOpt
PEVI (Jin et al., 2021)	Episodic	Linear	$\tilde{O}(dH^2N_1^{-\frac{1}{2}})$
PDS (Hu et al., 2023)	Discounted	Linear	$\tilde{O}(d^{\frac{1}{2}}(1 - \gamma)^{-1}N_1^{-\frac{1}{2}}) + \tilde{O}(d(1 - \gamma)^{-2}N_2^{-\frac{1}{2}})$
Our work	Episodic	kernel-based, d -finite spectrum	$\tilde{O}(Hd^{\frac{1}{2}}N_1^{-\frac{1}{2}}) + \tilde{O}(H^{\frac{5}{2}}d^{\frac{1}{2}}N_2^{-\frac{1}{2}})$
Our work	Episodic	kernel-based, general setting	$\tilde{O}(H\sqrt{G(N_1, \nu)\zeta_{\mathcal{D}_1}}) + \tilde{O}(H^2\sqrt{G(\frac{N}{H}, \lambda)\zeta_{\tilde{\mathcal{D}}}})$

Table 1: The existing suboptimality under weak convergence (see Assumption 4.7)(except for the last row), discussed in Section 2. Here, the labeled dataset represented as $\{(s'_h{}^\tau, a'_h{}^\tau, r_h{}^\tau)\}_{\tau, h=1}^{N_1, H}$, unlabeled dataset represented as $\{(s'_h{}^{\tau+N_1}, a'_h{}^{\tau+N_1})\}_{\tau, h=1}^{N_2, H}$, and \mathcal{D}^θ , which is a combination of labeled dataset and unlabeled dataset with $N = N_1 + N_2$ trajectories. We partition the dataset \mathcal{D}^θ into H disjoint and equally sized sub dataset $\{\tilde{\mathcal{D}}_h^\theta\}_{h=1}^H$. Denote γ as the discount factor for discounted MDP, $G(N, \lambda)$ is the maximum information gain, $\zeta_{\mathcal{D}} = \max_{h \in [H]} \zeta_h(\mathcal{D}', \mathcal{D})$ represents a maximum amount of information from the dataset \mathcal{D} and \mathcal{D}' , where \mathcal{D}' is the combination of \mathcal{D} and observed data z , and $\zeta_{\tilde{\mathcal{D}}} = \max_{h \in [H]} \zeta_h((\tilde{\mathcal{D}}_h^\theta)', \tilde{\mathcal{D}}_h^\theta)$. Note that $\nu = 1 + \frac{1}{N_1}$ and $\lambda = 1 + \frac{1}{N}$. In a linear MDP setting, it is stated that the transition probability can be represented linearly in a feature map of state-action with d dimensions.

2.1 Offline Reinforcement Learning

In offline reinforcement learning (RL), the goal is to learn a policy from a static data set collected previously without interacting with the environment. Current approaches in offline RL (Levine et al., 2020) can be broadly classified into dynamic programming methods and model-based methods. Dynamic programming methods aim to learn a state action value function, known as the Q function. Subsequently, this value function is used either to directly find the optimal policy or, in the case of actor-critic methods, to estimate a gradient for the expected returns of a policy. The offline dynamic programming algorithm operates in a tabular setting (Jin et al., 2018). However, algorithms designed for tabular settings have limitations when applied to function approximation settings with a large number of effective states. Recent work has centered around the functional approximation setting, especially in the linear setting, where the value function (or transition model) can be represented using a linear function of a known feature mapping (Jin et al., 2021; Cai et al., 2020; Zanette et al., 2021). As the linear Markov decision process (MDP) assumption is rigid and

rather restrictive in practice, Wang et al. (2020b) explores the kernel optimal least squares value iteration (KOVI) algorithm (Yang et al., 2020) for general function approximation. In contrast, model-based methods rely on their ability to estimate the transition function using a parameterized model, such as a neural network. Instead of employing dynamic programming methods to fit the model, model-based approaches leverage their ability to effectively utilize large and diverse datasets to estimate the transition function (Yu et al., 2021b; Janner et al., 2019; Uehara & Sun, 2021). Both of the methods presented above require a large amount of data to learn a state-action or transition function. In our work, we use reward-free data (i.e., unlabeled data) to improve the performance of learning a state-action function. On the theoretical front, Yin et al. (2022b) explore offline reinforcement learning with differentiable function class approximations, extending to non-linear function approximation. Blanchet et al. (2024) investigate distributionally robust offline reinforcement learning (robust offline RL), which aims to identify an optimal policy from offline datasets that remains effective in perturbed environments. Meanwhile, Hu et al. (2024) addresses the fundamental challenge of transitioning from offline learning to online fine-tuning.

2.2 Offline Data Sharing

Data sharing strategies in multi-task reinforcement learning (RL) have shown effectiveness, as observed in works such as Yu et al. (2021a); Eysenbach et al. (2020); Chen et al. (2021). This involves reusing data across different tasks by relabeling rewards, thereby enhancing performance in multi-task offline RL scenarios. Prior work has employed various relabeling strategies. These include uniform labeling (Kalashnikov et al., 2021), labeling based on metrics such as estimated Q -values (Yu et al., 2021a), and labeling based on distances to states in goal-conditioned settings (Chen et al., 2021). However, these approaches either necessitate access to the functional form of the reward for relabeling or are confined to goal-conditioned settings. On the other hand, Yu et al. (2022) proposes a straightforward strategy by assigning zero rewards to unlabeled data. On the other hand, Hu et al. (2023) employs linear regression to label rewards for unlabeled data. These approaches present alternative and potentially simpler methods for relabeling, especially in scenarios where direct access to the reward function is challenging or unavailable. In our work, we propose kernel ridge regression to exploit unlabeled data which under certain conditions can be reduced to linear regression.

3 Background

3.1 Episodic Markov Decision Process

Consider an episodic MDP (Yang et al., 2020; Sutton & Barto, 2018), denoted as $\mathcal{M} = (\mathcal{S}, \mathcal{A}, H, \mathcal{P}, r)$ with state space \mathcal{S} , action space \mathcal{A} , horizon H , transition function $\mathcal{P} = \{\mathcal{P}_h\}_{h \in [H]}$, and reward function $r = \{r_h\}_{h \in [H]}$. We assume that the reward function is bounded, that is, $r_h \in [0, 1]$. For any policy $\pi = \{\pi_h\}_{h \in [H]}$ and $h \in [H]$, we define the state-value function $V_h^\pi : \mathcal{S} \rightarrow \mathbb{R}$ and the action-valued function (Q-function) $Q_h^\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ as $V_h^\pi(s) = \mathbb{E}_\pi \left[\sum_{t=h}^H r_t(s_t, a_t) | s_h = s \right]$ and $Q_h^\pi(s, a) = \mathbb{E}_\pi \left[\sum_{t=h}^H r_t(s_t, a_t) | s_h = s, a_h = a \right]$. These two functions satisfy the well-known Bellman equation: $V_h^\pi(s) = \langle Q_h^\pi(s, \cdot), \pi_h(\cdot | s) \rangle_{\mathcal{A}}$ and $Q_h^\pi(s, a) = \mathbb{E} [r_h(s_h, a_h) + V_{h+1}^\pi(s_{h+1}) | s_h = s, a_h = a]$. For any function $f : \mathcal{S} \rightarrow \mathbb{R}$, we define the transition operator at each step $h \in [H]$ as $(\mathbb{P}_h f)(s, a) = \mathbb{E} [f(s_{h+1}) | s_h = s, a_h = a]$, and define the Bellman operator as $(\mathbb{B}_h f)(s, a) = \mathbb{E} [r_h(s_h, a_h) | s_h = s, a_h = a] + (\mathbb{P}_h f)(s, a)$. Similarly, for all $h \in [H]$, the Bellman optimality equations defined as $V_h^*(s) = \sup_{a \in \mathcal{A}} Q_h^*(s, a)$ and $Q_h^*(s, a) = (\mathbb{B}_h V_{h+1}^*)(s, a)$. Meanwhile, the optimal policy π^* satisfies $\pi_h^*(\cdot | s) = \operatorname{argmax}_{\pi_h} \langle Q_h^*(s, \cdot), \pi_h(\cdot | s) \rangle_{\mathcal{A}}$ and $V_h^*(s) = \langle Q_h^*(s, \cdot), \pi_h^*(\cdot | s) \rangle_{\mathcal{A}}$. Reinforcement learning aims to learn a policy maximizing expected cumulative reward. Accordingly, we define the performance metric (i.e., suboptimality) as

$$\text{SubOpt}(\pi; s) = V_1^{\pi^*}(s) - V_1^\pi(s). \quad (1)$$

3.2 Assumption of Offline Data

In offline RL setting, a learner uses pre-collected dataset \mathcal{D} , which consists of N trajectories $\{(s_h^\tau, a_h^\tau, r_h^\tau)\}_{\tau, h=1}^{N, H}$, generated by some fixed but unknown MDP \mathcal{M} under the behavior policy π^b in the following manner: $s_1^\tau \sim \rho^b$, $a_h^\tau \sim \pi_h^b(\cdot | s_h^\tau)$ and $s_{h+1}^\tau \sim \mathcal{P}_h(\cdot | s_h^\tau, a_h^\tau)$, $1 \leq h \leq H$. Here ρ^b represents a predetermined initial

state distribution associated with the static dataset. The learner may also have partial observations of the reward in addition to the above state-action observations. More elaborately, we assume access to both a labeled dataset $\mathcal{D}_1 = \{(s_h^\tau, a_h^\tau, r_h^\tau)\}_{\tau, h=1}^{N_1, H}$, and an unlabeled dataset $\mathcal{D}_2 = \{(s_h^{\tau+N_1}, a_h^{\tau+N_1})\}_{\tau, h=1}^{N_2, H}$. We utilize the estimated reward function with parameter θ , as determined in section 4, to relabel dataset \mathcal{D}_2 . The relabeled dataset, denoted as $\mathcal{D}_2^\theta = \{(s_h^{\tau+N_1}, a_h^{\tau+N_1}, \widehat{r}_h^\theta(s_h^{\tau+N_1}, a_h^{\tau+N_1}))\}_{\tau, h=1}^{N_2, H}$.

3.3 Reproducing Kernel Hilbert Space

Consider a reproducing kernel Hilbert space (RKHS) as a function space. For simplicity, let $z = (s, a)$ denote a state-action pair and denote $\mathcal{Z} = \mathcal{S} \times \mathcal{A}$. Without loss of generality, we regard \mathcal{Z} as a compact subset of \mathbb{R}^m , where the dimension m is fixed. Let $k : \mathcal{Z} \times \mathcal{Z} \rightarrow \mathbb{R}$ be a positive definite continuous kernel and its corresponding kernel matrix $[\mathbf{K}]_{i,j} = k(z_i, z_j), \forall i, j \in [m]$. Note that \mathbf{K} is positive semi-definite. Define \mathcal{H}_k as the RKHS induced by k , containing a family of functions defined in \mathcal{Z} . Let $\langle \cdot, \cdot \rangle_{\mathcal{H}_k} : \mathcal{H}_k \times \mathcal{H}_k \rightarrow \mathbb{R}$ and $\|\cdot\|_{\mathcal{H}_k} : \mathcal{H}_k \rightarrow \mathbb{R}$ denote the inner product and the norm on \mathcal{H}_k , respectively. According to the reproducing property, for all $f \in \mathcal{H}_k$, and $z \in \mathcal{Z}$, holds $f(z) = \langle f, k(\cdot, z) \rangle_{\mathcal{H}_k}$. For more details and different characterizations of RKHS, see Aronszajn (1950); Berlinet & Thomas-Agnan (2011). Without loss of generality, we assume that $\sup_{z \in \mathcal{Z}} k(z, z) \leq 1$.

Let $\mathcal{L}^2(\mathcal{Z})$ be the set of square-integrable functions on \mathcal{Z} with respect to the Lebesgue measure and let $\langle \cdot, \cdot \rangle_{\mathcal{L}^2}$ be the inner product on $\mathcal{L}^2(\mathcal{Z})$. The kernel function k induces an integral operator $T_k : \mathcal{L}^2(\mathcal{Z}) \rightarrow \mathcal{L}^2(\mathcal{Z})$ defined as $T_k f(z) = \int_{\mathcal{Z}} k(z, z') f(z') dz'$ for all $f \in \mathcal{L}^2(\mathcal{Z})$. By Mercer's theorem (Steinwart & Christmann, 2008), the integral operator T_k has countable and positive eigenvalues $\{\sigma_i\}_{i \geq 1}$ and the corresponding eigenfunctions $\{\psi_i\}_{i \geq 1}$. Then, the kernel function admits a spectral expansion $k(z, z') = \sum_{i=1}^{\infty} \sigma_i \psi_i(z) \psi_i(z')$. Moreover, the RKHS \mathcal{H}_k can be written as a subset of $\mathcal{L}^2(\mathcal{Z})$ such that $\mathcal{H}_k = \{f \in \mathcal{L}^2(\mathcal{Z}) : \sum_{i=1}^{\infty} \frac{\langle f, \psi_i \rangle_{\mathcal{L}^2}^2}{\sigma_i} < \infty\}$, and the inner product of \mathcal{H}_k also can be written as $\langle f, g \rangle_{\mathcal{H}_k} = \sum_{i=1}^{\infty} (1/\sigma_i) \langle f, \psi_i \rangle_{\mathcal{L}^2} \langle g, \psi_i \rangle_{\mathcal{L}^2}$ for all $f, g \in \mathcal{H}_k$. With the above construction, the scaled eigenfunctions $\{\sqrt{\sigma_i} \psi_i\}_{i \geq 1}$ form an orthonormal basis for \mathcal{H}_k . We define the mapping $\phi : z \mapsto k(z, \cdot)$ to transform data from $\mathcal{Z} = \mathcal{S} \times \mathcal{A}$ to the (possibly infinite-dimensional) RKHS \mathcal{H}_k , which satisfies $k(z, z') = \langle \phi(z), \phi(z') \rangle_{\mathcal{H}_k}$ for all $z, z' \in \mathcal{Z}$ (Steinwart & Christmann, 2008, Lemma 4.19). We define the maximum information gain (Srinivas et al., 2009) to describe the complexity of \mathcal{H}_k :

$$G(n, \lambda) = \sup \left\{ \frac{1}{2} \log \det (I + K_{\mathcal{D}}/\lambda) : \mathcal{D} \subset \mathcal{Z}, |\mathcal{D}| \leq n \right\}, \quad (2)$$

where $K_{\mathcal{D}}$ is the kernel matrix for the set \mathcal{D} . Furthermore, the magnitude of maximal information gain $G(n, \lambda)$ depends on how rapidly the eigenvalues decay to zero, serving as a proxy dimension of \mathcal{H} in the case of an infinite-dimensional space. If \mathcal{H}_k is of finite rank, we have that $G(n, \lambda) = \mathcal{O}(d \log n)$ (Yang et al., 2020), where d is the rank of \mathcal{H}_k – referred as the d -finite spectrum. In the following, we present several conditions that are often used in the analysis of the RKHS property of \mathcal{H}_k (Yang et al., 2020; Vakili et al., 2021; Yeh et al., 2023) characterizing the eigenvalue decay of \mathcal{H}_k .

Assumption 3.1. *The integral operator T_K has eigenvalues $\{\sigma_j\}_{j \geq 1}$ and the associated eigenfunctions $\{\psi_j\}_{j \geq 1}$. We assume that $\{\sigma_j\}_{j \geq 1}$ satisfies one of the following conditions for some constant $d > 0$.*

- *d -finite spectrum: $\sigma_j = 0, \forall j > d$, where d is a positive integer.*
- *d -exponential decay: there exists some constants $C_1, C_2 > 0$ such that $\sigma_j \leq C_1 \cdot \exp(-C_2 \cdot j^d)$, $\forall j \geq 1$, where $d > 0$.*
- *d -polynomial decay: there exists some constants $C_1 > 0$ such that $\sigma_j \leq C_1 \cdot j^{-d} \forall j \geq 1$, where $d > 1$.*

For both d -exponential decay and d -polynomial decay, we assume that there exists $C_\psi > 0$ such that $\sup_{z \in \mathcal{Z}} \sigma_j^\tau \cdot |\psi_j(z)| \leq C_\psi$ holds for all $j \geq 1$ and $\tau \in [0, 1/2)$.

For instance, let $\mathcal{D} \subset \mathbb{R}^d$ and assume the kernel function $k(z, z') \leq 1$, the Squared Exponential kernel, defined as

$$k(z, z') = \exp\left(-\frac{|z - z'|^2}{2l^2}\right) \quad (3)$$

where l is a lengthscale parameter, exhibits d -exponential decay. Similarly, the Matérn kernel, given by

$$k(z, z') = \frac{2^{1-\nu}}{\Gamma(\nu)} r^\nu B_\nu(r), \quad r = \frac{\sqrt{2\nu}}{l} |z - z'|$$

is characterized by d -polynomial decay. Here, ν determines the smoothness of sample paths (with smaller ν producing rougher paths), and B_ν is the modified Bessel function. Theorem 5 in Srinivas et al. (2009) provides detailed proof for these kernel properties.

We assume that the Bellman operator maps any bounded function onto a bounded RKHS norm ball, which is the common assumption used in the function approximation (Yang et al., 2020; Jin et al., 2018).

Assumption 3.2. Define the function class $\mathcal{Q}^* = \{f \in \mathcal{H}_k : \|f\|_{\mathcal{H}_k} \leq R_Q H\}$ for some fixed constant $R_Q > 0$. Then, for any $h \in [H]$ and any $Q : \mathcal{S} \times \mathcal{A} \rightarrow [0, H]$, it holds that $\mathbb{B}_h V \in \mathcal{Q}^*$ for $V(s) = \max_{a \in \mathcal{A}} Q(s, a)$.

A sufficient condition for Assumption 3.2 to hold is when $\mathcal{S} = [0, 1]^m$ and that $r_h(\cdot, \cdot), \mathcal{P}_h(s' | \cdot, \cdot) \in \{f \in \mathcal{H}_k : \|f\|_{\mathcal{H}_k} \leq 1\}$ for all $h \in [H], \forall s' \in \mathcal{S}$. To see this, suppose this condition holds, then for any integrable $V : \mathcal{S} \rightarrow [0, H]$ holds,

$$\begin{aligned} \|r_h + \mathbb{P}_h V\|_{\mathcal{H}_k} &\leq \|r_h\|_{\mathcal{H}_k} + \|\mathbb{P}_h V\|_{\mathcal{H}_k} \leq 1 + \left\| \int_{s' \in \mathcal{S}} \mathcal{P}_h(s' | \cdot, \cdot) V(s') ds' \right\|_{\mathcal{H}_k} \\ &\leq 1 + \int_{s' \in \mathcal{S}} \|\mathcal{P}_h(s' | \cdot, \cdot) V(s')\|_{\mathcal{H}_k} ds' = 1 + \int_{s' \in \mathcal{S}} \|\mathcal{P}_h(s' | \cdot, \cdot)\|_{\mathcal{H}_k} \|V(s')\|_{\mathcal{H}_k} ds' \\ &\leq 1 + H \int_{s' \in \mathcal{S}} ds' = H + 1. \end{aligned}$$

Note that under the assumptions of measurability and boundedness on the kernel k , $\|\mathbb{P}_h V\|_{\mathcal{H}_k} \in \mathcal{H}_k$, which is given in Muandet et al. (2017, section 3.1). Thus, Assumption 3.2 holds with $R_Q = 2$. This assumption is mild and is also used in Yang et al. (2020). Similar assumptions are used in linear MDP's, which are much stricter (Jin et al., 2021; Zanette et al., 2020).

3.4 Pessimistic Value Iteration and Kernel Setting

We consider the pessimistic value iteration, i.e., PEVI (Jin et al., 2021) algorithm, described in Algorithm 2, as the backbone algorithm. This is a model-free, theoretically guaranteed offline algorithm. The fundamental insight of PEVI lies in the incorporation of a penalty function, which essentially introduces a sense of pessimism, into the value iteration algorithm. The key challenge to extend PEVI to kernel setting is that the dimension (even effective dimension) of the kernel based model (when interpreted as linear model) is divergent. In addition, we apply the data splitting method (Rashidinejad et al., 2021; Xie et al., 2021). As introduced in Rashidinejad et al. (2021), data splitting makes sure that the estimated value \hat{V}_{h+1} and estimated Bellman operator $\hat{\mathbb{B}}_h$ are estimated using different subsets of \mathcal{D} , this yields conditional independence that is required in bounding concentration terms of the form $(\hat{\mathbb{B}}_h - \mathbb{B}_h) \hat{V}_{h+1}$, and hence the suboptimality can be reduced by a factor of \sqrt{d} . However, applied naively, this data splitting induces one undesired \sqrt{H} factor in the optimality as we need to split \mathcal{D} into H folds and thus each \mathbb{B}_h is estimated using only N/H episodes of data. Further details of the PEVI algorithm can be found in Appendix B.

4 Unsupervised Data Sharing

Our algorithm comprises two main components. The first part involves employing kernel ridge regression to learn the reward function using the labeled dataset and constructing the confidence set. Next, to mitigate

overestimation in reward prediction, we construct the pessimism reward parameter $\tilde{\theta}$ within the confidence set. Section 4.1 discusses this in more detail. The second part involves using the pessimistic reward estimator $\tilde{\theta}$ to relabel the entire dataset, which is a combination of the labeled dataset and the relabeled dataset. Following this, we employ the PEVI algorithm with kernel approximation and data splitting (refer to Algorithm 3) to determine the optimal policy. The detailed steps of the algorithm are outlined in Algorithm 1.

Algorithm 1 Data Sharing, Kernel Approximation

- 1: **Data:** Labeled dataset \mathcal{D}_1 , and unlabeled dataset \mathcal{D}_2 .
- 2: **Input:** Parameter $\beta_h(\delta), \delta, B, \nu, \lambda$.
- 3: Define \mathcal{D}^θ , which is a combination of the labeled dataset \mathcal{D}_1 and the unlabeled dataset \mathcal{D}_2 , and partition the dataset \mathcal{D}^θ into H disjoint and equally sized sub datasets $\{\tilde{\mathcal{D}}_h^\theta\}_{h=1}^H$.
- 4: Learn the reward function $\hat{\theta}_1, \dots, \hat{\theta}_H$ from \mathcal{D}_1 with

$$\hat{\theta}_h = \operatorname{argmin}_{\theta_h \in \mathcal{H}_k} \sum_{\tau=1}^{N_1} \left[r_h^\tau - \hat{r}_h^{\theta_h}(s_h^\tau, a_h^\tau) \right]^2 + \nu \|\theta_h\|_{\mathcal{H}_k}^2.$$

- 5: Construct the pessimistic reward function with parameter $\tilde{\theta} := \{\tilde{\theta}_h\}_{h=1}^H$ satisfy

$$\tilde{r}_h^{\tilde{\theta}}(s, a) = \max \left\{ \left\langle \hat{\theta}_h, \phi(s, a) \right\rangle_{\mathcal{H}_k} - \beta_h(\delta) \left\| (\Lambda_h^{\mathcal{D}_1})^{-\frac{1}{2}} \phi(s, a) \right\|_{\mathcal{H}_k}, 0 \right\}. \quad (4)$$

- 6: Annotate the reward in \mathcal{D}^θ with parameter $\theta = \tilde{\theta}$.
- 7: Learn the policy from the relabeled dataset $\mathcal{D}^{\tilde{\theta}}$ using Algorithm 3 in Appendix.

$$\{\hat{\pi}_h\}_{h=1}^H \leftarrow \text{PEVI}(\mathcal{D}^{\tilde{\theta}}, B, \lambda).$$

- 8: **Result:** $\hat{\pi} = \{\hat{\pi}_h\}_{h=1}^H$.
-

4.1 Pessimistic Reward Estimation

We utilize labeled dataset \mathcal{D}_1 to train a reward function $\hat{r}_h^{\theta_h}$, using it to label the unlabeled data. Assume that the observed reward is generated as $r_h^\tau = r_h(s_h^\tau, a_h^\tau) + \epsilon_h^\tau$ where $r_h : (s, a) \mapsto \langle \theta_h^*, \phi(s, a) \rangle_{\mathcal{H}_k}$ satisfies $r_h(s, a) \in [0, 1]$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$, and ϵ_h^τ are i.i.d. centered 1-SubGaussian noise. Here $\theta_h^* \in \mathcal{H}_k$ is an unknown parameter, and $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{H}_k$ is a known feature map defined in Section 3.3. Furthermore, we assume that $\|\theta_h^*\|_{\mathcal{H}_k} \leq \mathcal{S}$. We learn the reward function from labeled data through a kernel ridge regression problem. Using the feature representation, we write

$$\hat{\theta}_h = \operatorname{argmin}_{\theta_h \in \mathcal{H}_k} \sum_{\tau=1}^N \left[r_h^\tau - \hat{r}_h^{\theta_h}(s_h^\tau, a_h^\tau) \right]^2 + \nu \|\theta_h\|_{\mathcal{H}_k}^2, \quad (5)$$

where $\hat{r}_h^{\theta_h}(s, a) = \langle \phi(s, a), \theta_h \rangle_{\mathcal{H}_k}$ with parameter θ_h . However, this method leads to an overestimation of predicted reward values, as highlighted in Yu et al. (2022). A novel algorithm called Provable Data Sharing (PDS) is introduced in Hu et al. (2023) to mitigate this problem. PDS incorporates uncertainty penalties into the learned reward functions and integrates seamlessly with existing offline RL algorithms in a linear MDP setting. We extend the application of this algorithm to the kernel setting.

To address the problem of overestimating predicted rewards, we analyze the uncertainty in the learned reward function. The previous solution defines the center of the ellipsoidal confidence set:

$$\mathcal{C}_h(\delta) = \left\{ \theta \in \mathcal{H}_k : \left\| \theta - \hat{\theta}_h \right\|_{\Lambda_h^{\mathcal{D}_1}} \leq \beta_h(\delta) \right\}, \quad (6)$$

where $\|\theta\|_{\Lambda_h^{\mathcal{D}_1}}^2 = \langle \theta, \Lambda_h^{\mathcal{D}_1} \theta \rangle_{\mathcal{H}_k}$ and $\Lambda_h^{\mathcal{D}_1} = \sum_{\tau=1}^{N_1} \phi(s_h^\tau, a_h^\tau) \phi(s_h^\tau, a_h^\tau)^\top + \nu I_{\mathcal{H}_k}$ is a positive definite operator, and $\beta_h(\delta)$ is its radius which follows Proposition 4.1.

Proposition 4.1. *We define $\beta_h(\delta)$ with the labeled data set \mathcal{D}_1 by $\beta_h(\delta) = \sqrt{\nu} \mathcal{S} + \sqrt{\log \frac{\det[\nu I + K_h^{\mathcal{D}_1}]}{\delta^2}}$, where $K_h^{\mathcal{D}_1}$ is the Gram matrix constructed from the dataset \mathcal{D}_1 as $[K_h^{\mathcal{D}_1}]_{\tau, \tau'} = k(z_h^\tau, z_h^{\tau'})$, where $z_h^\tau = (s_h^\tau, a_h^\tau)$ for $\tau, \tau' \in [N_1]$ and for each $h \in [H]$ and $\delta \in (0, 1)$. Then, with probability at least $1 - \delta$ we have $\|\hat{\theta}_h - \theta_h^*\|_{\Lambda_h^{\mathcal{D}_1}} \leq \beta_h(\delta)$, where $\hat{\theta}_h$ is the solution of equation (5). Furthermore, consider the information gain $G(N, \nu)$, defined in equation (2) of the matrix $K_h^{\mathcal{D}_1}$ and set $\nu = 1 + 1/N_1$, $\beta_h(\delta)$ is rewritten as*

$$\sqrt{\nu} \mathcal{S} + \sqrt{2G(N_1, \nu) + 1 + \log \frac{1}{\delta^2}}. \quad (7)$$

Moreover, define $\mathcal{C}_h(\delta) = \left\{ \theta \in \mathcal{H}_k : \|\theta - \hat{\theta}_h\|_{\Lambda_h^{\mathcal{D}_1}} \leq \beta_h(\delta) \right\}$, we have $\mathbb{P}(\theta_h^* \in \mathcal{C}_h(\delta)) \geq 1 - \delta$.

Proof. Please refer to Appendix B.1 for detailed proof. \square

In Proposition 4.1, the uncertainty of the learned reward function depends on the maximum information gain of the kernel matrix $K_h^{\mathcal{D}_1}$. However, finding the optimal parameter within the confidence set is computationally inefficient. To address this challenge, Hu et al. (2023) proposes an approach that preserves the pessimistic property of the offline algorithm. This method uses pessimistic estimation, allowing the algorithm to remain pessimistic while mitigating computational challenges. Formally, we construct the pessimistic reward function $\tilde{r}_h^{\hat{\theta}_h}(s, a)$ for the parameter $\tilde{\theta}_h$ as

$$\tilde{r}_h^{\hat{\theta}_h}(s, a) = \max \left\{ \langle \hat{\theta}_h, \phi(s, a) \rangle_{\mathcal{H}_k} - \beta_h(\delta) \|(\Lambda_h^{\mathcal{D}_1})^{-\frac{1}{2}} \phi(s, a)\|_{\mathcal{H}_k}, 0 \right\}. \quad (8)$$

The equation (8) is guaranteed by the following lemma derived from Cauchy-Schwarz inequalities.

Lemma 4.2. $\left| \langle \theta_h - \hat{\theta}_h, \phi(s, a) \rangle_{\mathcal{H}_k} \right| \leq \beta_h(\delta) \|(\Lambda_h^{\mathcal{D}_1})^{-\frac{1}{2}} \phi(s, a)\|_{\mathcal{H}_k}$ for any $\theta_h \in \mathcal{C}_h(\delta)$, $h \in [H]$.

The equation (8) provides a lower bound for the reward function within the confidence set $\mathcal{C}(\delta)$. When the labeled data is scarce, or when there is a significant shift in the distribution between the labeled and unlabeled data, the confidence interval becomes wider and then the equation (8) degenerates to 0, which is reduced to the UDS algorithm (Yu et al., 2022; Hu et al., 2023).

4.2 Theoretical Analysis

The suboptimality of the Algorithm 1 is characterized by the following theorem.

Theorem 4.3. *Consider the MDP described in Section 3.1. Under Assumption 3.1 and Assumption 3.2, and suppose the labeled dataset \mathcal{D}_1 and unlabeled dataset \mathcal{D}_2^θ are defined in Section 3.2. Define $\mathcal{D}^\theta = \{(s_h^\tau, a_h^\tau, \hat{r}_h^{\theta_h}(s_h^\tau, a_h^\tau))\}_{\tau, h=1}^{N, H}$, which is a combination of labeled dataset \mathcal{D}_1 and unlabeled dataset \mathcal{D}_2^θ with $N = N_1 + N_2$. We partition dataset \mathcal{D}^θ into H disjoint and equally sized sub dataset $\{\tilde{\mathcal{D}}_h^\theta\}_{h=1}^H$, where $|\tilde{\mathcal{D}}_h^\theta| = N_h = N/H$. Let $\mathcal{I}_h = \{N_h \cdot (h-1) + 1, \dots, N_h \cdot h\} = \{\tau_{h,1}, \dots, \tau_{h,N_h}\}$ satisfy $\tilde{\mathcal{D}}_h^\theta = \{(s_h^\tau, a_h^\tau, r_h^\tau)\}_{\tau \in \mathcal{I}_h}$. We set $\lambda = 1 + \frac{1}{N}$, $\nu = 1 + \frac{1}{N_1}$ in Algorithm 1, where*

$$\beta_h(\delta) = \begin{cases} \sqrt{1 + \frac{1}{N_1}} \mathcal{S} + \sqrt{C_1 \cdot d \cdot \log N_1 + \log(\frac{1}{\delta^2})} & d\text{-finite spectrum,} \\ \sqrt{1 + \frac{1}{N_1}} \mathcal{S} + \sqrt{C_1 \cdot (\log N_1)^{1+\frac{1}{d}} + \log(\frac{1}{\delta^2})} & d\text{-exponential decay,} \\ \sqrt{1 + \frac{1}{N_1}} \mathcal{S} + \sqrt{C_1 \cdot (N_1)^{\frac{m+1}{d+m}} \cdot \log(N_1) + \log(\frac{1}{\delta^2})} & d\text{-polynomial decay.} \end{cases}$$

$$B = \begin{cases} C_2 \cdot H \cdot \sqrt{d \log(N/\delta)} & d\text{-finite spectrum,} \\ C_2 \cdot H \cdot \sqrt{(\log N/\delta)^{1+1/d}} & d\text{-exponential decay,} \\ C_2 \cdot N^{\frac{m+1}{2(d+m)}} H^{1-\frac{m+1}{2(d+m)}} \cdot \sqrt{\log(N/\delta)} & d\text{-polynomial decay.} \end{cases}$$

Here, $C_1, C_2 > 0$ are absolute constants that does not depend on N_1, N , nor H . Then, for fixed initial state $s_0 \in \mathcal{S}$, with probability $1 - 2\delta$, the policy $\hat{\pi}$ generated by Algorithm 1 satisfies

$$\begin{aligned} \text{SubOpt}(\hat{\pi}; s_0) &\leq 2 \sum_{h=1}^H \beta_h(\delta) \mathbb{E}_{\pi^*} \left[\|\phi(s_h, a_h)\|_{(\Lambda_h^{\mathcal{D}_1})^{-1}} \mid s_1 = s_0 \right] \\ &\quad + 2B \sum_{h=1}^H \mathbb{E}_{\pi^*} \left[\|\phi(s_h, a_h)\|_{(\tilde{\Lambda}_h^{\tilde{\mathcal{D}}_h})^{-1}} \mid s_1 = s_0 \right], \end{aligned}$$

where $\tilde{\theta}$ is defined in equation 4.

Proof. For a detailed proof, see Appendix B.2. \square

Two key terms express the suboptimality bound. The first term is the reward bias introduced by uncertainties in estimating rewards. This term reflects the challenges and inaccuracies associated with predicting or estimating rewards in a given environment. The second term represents the offline algorithm and optimal policy π^* error.

Remark 4.4. We use the Lemma C.2 to rewrite the term of $\beta_h(\delta)$ and B in the Theorem 4.3 as $\beta_h(\delta) = \tilde{\mathcal{O}}(\sqrt{G(N_1, 1 + \frac{1}{N_1})})$ and $B = \tilde{\mathcal{O}}(H\sqrt{G(N, 1 + \frac{1}{N})})$.

By Remark 4.4, both terms $\beta_h(\delta)$ and B depend on the kernel function class. It is worth noting that the term $\|\phi(s_h, a_h)\|_{(\Lambda_h^{\mathcal{D}})^{-1}}$ can be expressed as an information quantity for the dataset \mathcal{D} , as outlined in Lemma 4.5.

Proposition 4.5. For all $h \in [H]$, we partition dataset \mathcal{D} into H disjoint and equally sized sub datasets $\{\tilde{\mathcal{D}}_h\}_{h=1}^H$, where $\tilde{\mathcal{D}}_h = \{(s_h^\tau, a_h^\tau, r_h^\tau)\}_{\tau \in \mathcal{I}_h}$ with $\mathcal{I}_h = \{N_h \cdot (h-1) + 1, \dots, N_h \cdot h\} = \{\tau_{h,1}, \dots, \tau_{h,N_h}\}$ and $N_h = N/H$. Denote the operator $\Phi_h^{\tilde{\mathcal{D}}_h} : \mathcal{H}_k \rightarrow \mathbb{R}^{N_h}$, and $\Lambda_h^{\tilde{\mathcal{D}}_h} : \mathcal{H}_k \rightarrow \mathcal{H}_k$ as

$$\Phi_h^{\tilde{\mathcal{D}}_h} = \begin{pmatrix} \phi(z_h^{\tau_{h,1}})^\top \\ \vdots \\ \phi(z_h^{\tau_{h,N_h}})^\top \end{pmatrix} = \begin{pmatrix} k(\cdot, z_h^{\tau_{h,1}})^\top \\ \vdots \\ k(\cdot, z_h^{\tau_{h,N_h}})^\top \end{pmatrix}, \quad \Lambda_h^{\tilde{\mathcal{D}}_h} = \lambda \cdot I_{\mathcal{H}} + (\Phi_h^{\tilde{\mathcal{D}}_h})^\top \Phi_h^{\tilde{\mathcal{D}}_h}.$$

Define gram matrix $K_h^{\tilde{\mathcal{D}}_h} = \Phi_h^{\tilde{\mathcal{D}}_h} (\Phi_h^{\tilde{\mathcal{D}}_h})^\top$. Then, for any $z \in \mathcal{Z}$, we have

$$\phi(z)^\top (\Lambda_h^{\tilde{\mathcal{D}}_h})^{-1} \phi(z) \leq 2 \cdot \left[\log \det \left(I + K_h^{\tilde{\mathcal{D}}'_h} / \lambda \right) - \log \det \left(I + K_h^{\tilde{\mathcal{D}}_h} / \lambda \right) \right], \quad (9)$$

where $\tilde{\mathcal{D}}'_h$ is the combination of dataset $\tilde{\mathcal{D}}_h$ and z which satisfies $\Lambda_h^{\tilde{\mathcal{D}}'_h} = \Lambda_h^{\tilde{\mathcal{D}}_h} + \phi(z)\phi(z)^\top$.

Proof. For a detailed proof, see Appendix B.3. \square

Remark 4.6. In Proposition 4.5, \mathcal{H}_k can be infinite dimensional. However, for the sake of clarity, we represent $\Phi_h^{\tilde{\mathcal{D}}_h}$ as a matrix and $\phi(z_h^\tau)$ as a column vector for all $\tau \in \mathcal{I}_h$.

Here, we define

$$\zeta_h(\mathcal{D}', \mathcal{D}) = 2 \left[\log \det \left(I + K_h^{\mathcal{D}'} / \lambda \right) - \log \det \left(I + K_h^{\mathcal{D}} / \lambda \right) \right], \quad (10)$$

as the maximal information amount between the dataset \mathcal{D}' and \mathcal{D} . Proposition 4.5 states that if the training data set is well known about z , then equation (10) will be close to zero. On the other hand, if the training data set is not well known about z , then equation (10) will be large.

We specialize the d -finite spectrum case of Theorem 4.3 under a weak data coverage assumption to better understand the convergence of Algorithm 1.

Assumption 4.7 (Weak Convergence). *Suppose the dataset $\mathcal{D} = \{(s_h^\tau, a_h^\tau, r_h^\tau)\}_{\tau,h=1}^{N,H}$ consists of N trajectories, for all $h \in [H]$, the trajectories are drawn independently and identically from distributions induced by some fixed behavior policy $\bar{\pi}$ such that there exists a constant $c_{\min} > 0$ satisfying $\inf_{\|f\|_{\mathcal{H}_k}=1} \langle f, \mathbb{E}_{\bar{\pi}} [\phi(z_h)\phi(z_h)^\top] f \rangle \geq c_{\min}$ for any $h \in [H]$.*

Intuitively, Assumption 4.7 posits that the collected data should be relatively well distributed throughout the state action space. Notably, assumption 4.7 shares similarities with other explorability assumptions common in reinforcement learning literature, such as those in Yin et al. (2022a); Wagenmaker & Pacchiano (2023).

Corollary 4.8 (Well-Explored Dataset). *In the d -finite spectrum case, assume that the Assumption 4.7 holds under the same conditions as Theorem 4.3. Then for $N_1 \geq \Omega(\log(dH/\delta))$ and $N \geq H \cdot \Omega(\log(dH/\delta))$, with probability at least $1 - \delta$, we have*

$$\begin{aligned} \text{SubOpt}(\hat{\pi}; s) &\leq 2\beta_h(\delta) \cdot H \cdot c' / \sqrt{N_1} + 2B \cdot H \cdot c' / \sqrt{N_h} \\ &\leq \tilde{\mathcal{O}}(H \sqrt{\frac{d}{N_1}}) + \tilde{\mathcal{O}}(H^{\frac{5}{2}} \sqrt{\frac{d}{N_2}}). \end{aligned} \quad (11)$$

In the d -finite spectrum case, a significant difference between our present study and previous work (Hu et al., 2023) lies in the incorporation of factors \sqrt{d} and \sqrt{H} , introduced by the implementation of the data splitting technique (Xie et al., 2021). This technique plays a crucial role in the linear case, influencing the overall convergence behavior of the learned policy. If we aim to transform the feature mapping from a dimensionality of d to d' , where $d' > d$. In this context, the data partitioning method can help mitigate the convergence of the error bound. Finally, we combine the result in Theorem 4.3, Remark 4.4, and Corollary 4.8 to get the Table 1.

Remark 4.9. *For Assumption 4.7, the scenarios involving d -exponential and d -polynomial decay are not generally valid. If Assumption 4.7 holds true, by integrating Lemma B.2, Lemma B.3, and equation (89), we can deduce the explicit form of $\text{SubOpt}(\hat{\pi}; s)$ as*

$$\text{SubOpt}(\hat{\pi}; s) = \begin{cases} \tilde{\mathcal{O}}\left(HN_1^{-\frac{1}{2}}\right) + \tilde{\mathcal{O}}\left(H^{\frac{5}{2}}N_2^{-\frac{1}{2}}\right) & d\text{-exponential decay,} \\ \tilde{\mathcal{O}}\left(HN_1^{-\frac{1}{2} + \frac{m+1}{d+m}}\right) + \tilde{\mathcal{O}}\left(H^{\frac{5}{2} - \frac{m+1}{2(d+m)}}N_2^{-\frac{1}{2} + \frac{m+1}{2(d+m)}}\right) & d\text{-polynomial decay.} \end{cases} \quad (12)$$

Nonetheless, Assumption 4.7 does not hold under scrutiny. To demonstrate this, let's assume that Assumption 4.7 is true. It means that for every f within the set $\{\|f\|_{\mathcal{H}_k} = 1\}$, it satisfies $\mathbb{E}_{\bar{\pi}}[\langle f, \phi(z_h)\phi(z_h)^\top f \rangle] \geq c_{\min}$. Then, we express ϕ and f as $\phi = \sum_{i=1}^{\infty} a_i \psi_i$ and $f = \sum_{i=1}^{\infty} b_i \psi_i$ respectively, where $\{\psi_i\}_{i=1}^{\infty}$ is orthonormal basis of \mathcal{H}_k . Given that f can represent any function satisfying $\|f\|_{\mathcal{H}_k} = 1$, let f be any vector such that $f = b_j \psi_j$ for an arbitrary j . Consequently, for all j , the expectation $\mathbb{E}_{\bar{\pi}}[\langle f, \phi(z_h)\phi(z_h)^\top f \rangle] = a_j^2 \geq c_{\min}$ is satisfied, which results in a paradox because the norm should be finite; however, $\|\phi\|_{\mathcal{H}_k} = \sum_{j=1}^{\infty} a_j^2 \geq \sum_{j=1}^{\infty} c_{\min} = \infty$.

5 Experiments

In the following experiments, we execute Algorithm 1 to obtain a policy π and the corresponding value function $V_1^\pi(s)$, where s is sampled from the initial distribution $\rho(s)$. For our experiment, we select the Squared Exponential kernel from equation (3) as the kernel function K . With this choice of kernel, the pessimistic reward function $\tilde{r}_h^{\theta_h}(s, a)$, as defined in equation (8), can be computed using the following expression:

$$\begin{aligned} \tilde{r}_h^{\theta_h}(s, a) &= k_h^{\mathcal{D}_1}(s, a)^\top \left(K_h^{\mathcal{D}_1} + \nu I \right)^{-1} y_h \\ &\quad - \beta_h(\delta) \cdot \nu^{-1/2} \cdot \left(K\left((s, a), (s, a)\right) - k_h^{\mathcal{D}_1}(s, a)^\top \left(K_h^{\mathcal{D}_1} + \nu I \right)^{-1} k_h^{\mathcal{D}_1}(s, a) \right)^{1/2} \end{aligned}$$

We implemented our algorithm using Python's NumPy library. To ensure reproducibility, our experimental code is publicly accessible¹.

¹https://github.com/d09942015ntu/leveraging_unlabeled_offline_rl

5.1 Asymptotic Behavior of $V_1^\pi(s)$

In this experiment, we investigate the asymptotic behavior of $V_1^\pi(s)$. Although we are unable to theoretically validate the correctness of equation (12) due to the violation of Assumption 4.7, we provide empirical evidence that, when the kernel function is the Squared Exponential kernel satisfying d -exponential decay, its asymptotic behavior closely aligns with equation (12). We create a toy example of RL environment which meets Assumption 3.2, ensuring bounded RKHS norms for $r_h(\cdot, \cdot)$ and $\mathcal{P}(s'|\cdot, \cdot)$. The MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, H, \mathcal{P}, r)$ for this toy example is as follows.

$$\begin{aligned}\mathcal{S} &= [0, C], \quad \mathcal{A} = \{0, 1, \dots, C\}, \quad H = C, \\ \mathcal{P}_h(s' | s, a) &= \exp[-\alpha (s' - ((s + a) \bmod C))^2] / \sqrt{\pi/\alpha}, \\ r(s, a) &= \exp[-\alpha (s - C/2)^2] / \sqrt{\pi/\alpha},\end{aligned}$$

where α and C are constants, set to $\alpha = 3$ and $C = 8$ in our experiment. To ensure that s' remains within the state space \mathcal{S} , we replace s' by " $s' \bmod C$ " after sampling $s' \sim \mathcal{P}_h(\cdot | s, a)$. To examine the asymptotic behavior of $V_1^\pi(s)$, we varied the parameters N_1 and N_2 and plotted the resulting values of $V_1^\pi(s)$ in the left column of Fig. 1. For asymptotic approximation, using equation (12), $V_1^\pi(s)$ can be expressed in the following form:

$$V_1^\pi(s) = c - aN_1^{-\frac{1}{2}} - bN_2^{-\frac{1}{2}},$$

where a, b , and c are constants. By applying linear regression to the experimental data (left column of Fig. 1), we estimated these constants and plotted the corresponding curve in the right column of Fig. 1. The strong agreement between the experimental results and the asymptotic approximation validates our hypothesis that the asymptotic behavior of $V_1^\pi(s)$ conforms to our theoretical predictions.

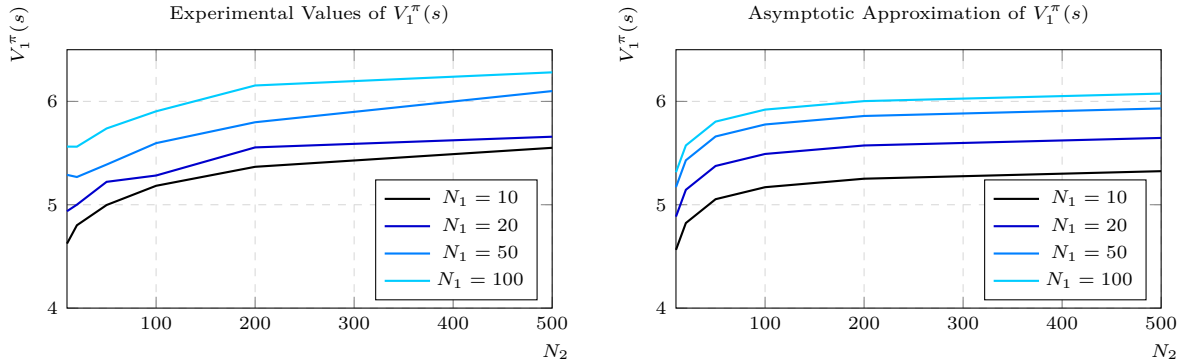


Figure 1: Comparison of experimental values and asymptotic approximation of $V_1^\pi(s)$.

5.2 Comparison between Finite Dimensional and Kernel Features

In this experiment, we present a real-world example demonstrating the superior performance of our kernel-based offline RL method compared to finite-dimensional feature representations ϕ , which follows the setting of Hu et al. (2023) but in a finite-horizon scenario. We select the CartPole environment from OpenAI Gym² as our test environment. To implement the ϕ function used in their work, we consider three different realizations of $\phi(s, a)$ —linear, quadratic, and cubic—denoted as ϕ_{lin} , ϕ_{quad} , and ϕ_{cubic} , respectively, which are defined as follows:

$$\begin{aligned}\phi_{\text{lin}} &= (1, x_1, x_2, \dots, x_n) \quad \text{where } x_1, \dots, x_n \text{ is the elements in the concatenation of } (s, a) \\ \phi_{\text{quad}} &= (1, x_1, x_2, \dots, x_n, x_1^2, x_1x_2, \dots, x_{n-1}x_n, x_n^2) \\ \phi_{\text{cubic}} &= (1, x_1, x_2, \dots, x_n, x_1^2, x_1x_2, \dots, x_{n-1}x_n, x_n^2, x_1^3, x_1^2x_2, x_1x_2x_3, \dots, x_{n-2}x_{n-1}x_n^2, x_{n-1}x_n^2, x_n^3)\end{aligned}$$

²https://www.gymnasium.dev/environments/classic_control/cart_pole/

We compare these three finite-dimensional feature representations with our kernel-based RL method using the Squared Exponential kernel feature, denoted as ϕ_K , under various configurations of N_1 and N_2 . The results are presented in Fig. 2. Our findings suggest that, the kernel-based ϕ_K generally outperforms alternative finite-dimensional ϕ approaches when $N_2 \leq 500$. The performance gap between ϕ_K and other ϕ becomes more pronounced when $N_2 \leq 100$. This indicates that kernel methods, with their more flexible function approximation, better capture the reward and transition dynamics—particularly when unlabeled data is scarce—resulting in higher value estimates compared to finite-dimensional ϕ representations.

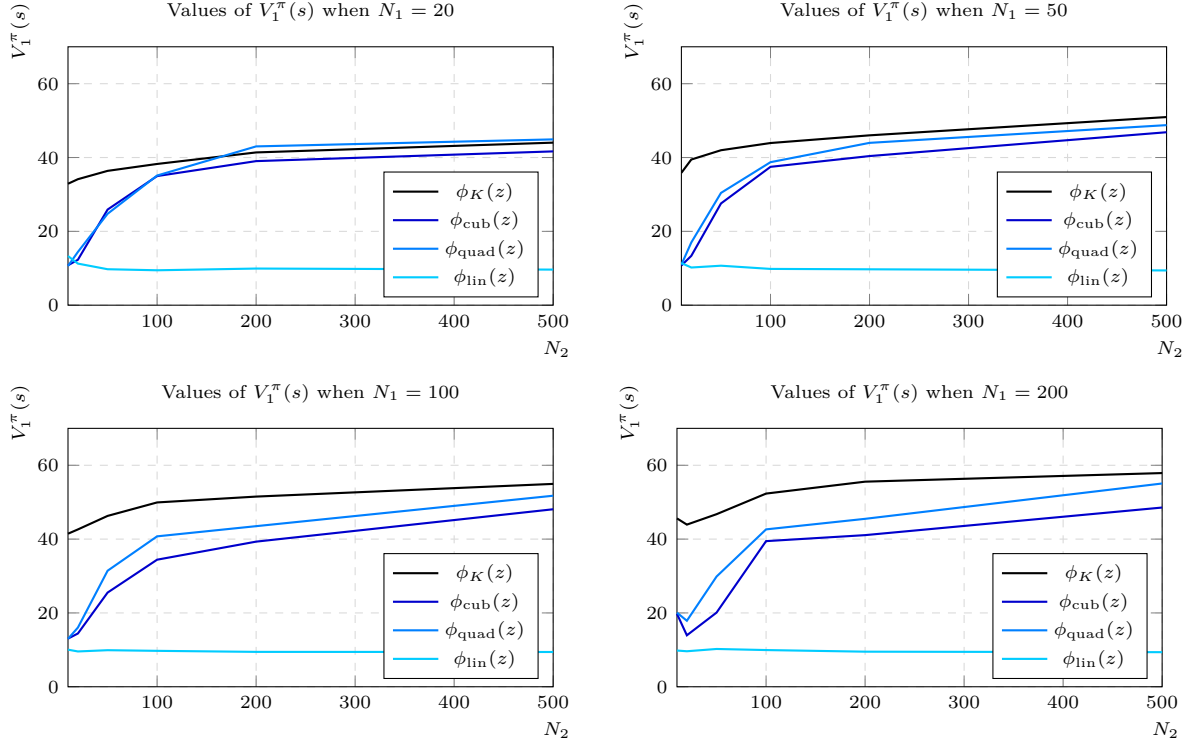


Figure 2: Comparison the values of $V_1^\pi(s)$ between finite dimensional features and kernel features.

6 Conclusion

In this paper, we demonstrate that incorporating unlabeled data into offline RL can greatly improve offline RL performance. Our theoretical analysis shows how unlabeled data can improve the performance of offline RL, especially in a more general function approximation setting, in contrast to the results in Hu et al. (2023). Our analysis is based on the common offline RL assumption about the dataset, providing a comprehensive examination of the algorithm’s performance under these conditions. In future work, it may be interesting to extend to the discounted MDP setting to deal with more category problems and the low-rank MDP (Uehara et al., 2021).

References

- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24, 2011.
- Nachman Aronszajn. Theory of reproducing kernels. *Transactions of the American mathematical society*, 68 (3):337–404, 1950.

- Alain Berlinet and Christine Thomas-Agnan. *Reproducing kernel Hilbert spaces in probability and statistics*. Springer Science & Business Media, 2011.
- Jose Blanchet, Miao Lu, Tong Zhang, and Han Zhong. Double pessimism is provably efficient for distributionally robust offline reinforcement learning: Generic algorithm and robust partial coverage. *Advances in Neural Information Processing Systems*, 36, 2024.
- Qi Cai, Zhuoran Yang, Chi Jin, and Zhaoran Wang. Provably efficient exploration in policy optimization. In *International Conference on Machine Learning*, pp. 1283–1294. PMLR, 2020.
- Annie S Chen, Suraj Nair, and Chelsea Finn. Learning generalizable robotic reward functions from "in-the-wild" human videos. *Robotics: Science and Systems (RSS)*, 2021.
- Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *International Conference on Machine Learning*, pp. 844–853. PMLR, 2017.
- Yaqi Duan, Zeyu Jia, and Mengdi Wang. Minimax-optimal off-policy evaluation with linear function approximation. In *International Conference on Machine Learning*, pp. 2701–2709. PMLR, 2020.
- Ben Eysenbach, Xinyang Geng, Sergey Levine, and Russ R Salakhutdinov. Rewriting history with inverse rl: Hindsight inference for policy improvement. *Advances in neural information processing systems*, 33: 14783–14795, 2020.
- Chelsea Finn, Sergey Levine, and Pieter Abbeel. Guided cost learning: Deep inverse optimal control via policy optimization. In *International conference on machine learning*, pp. 49–58. PMLR, 2016.
- Justin Fu, Katie Luo, and Sergey Levine. Learning robust rewards with adversarial inverse reinforcement learning. In *International Conference on Learning Representations*, 2017.
- Justin Fu, Avi Singh, Dibya Ghosh, Larry Yang, and Sergey Levine. Variational inverse control with events: A general framework for data-driven reward definition. *Advances in neural information processing systems*, 31, 2018.
- Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. *Advances in neural information processing systems*, 29, 2016.
- Hao Hu, Yiqin Yang, Qianchuan Zhao, and Chongjie Zhang. The provable benefits of unsupervised data sharing for offline reinforcement learning. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=MTTPLcwwqTt>.
- Hao Hu, Yiqin Yang, Jianing Ye, Chengjie Wu, Ziqing Mai, Yujing Hu, Tangjie Lv, Changjie Fan, Qianchuan Zhao, and Chongjie Zhang. Bayesian design principles for offline-to-online reinforcement learning. *arXiv preprint arXiv:2405.20984*, 2024.
- Michael Janner, Justin Fu, Marvin Zhang, and Sergey Levine. When to trust your model: Model-based policy optimization. *Advances in neural information processing systems*, 32, 2019.
- Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is q-learning provably efficient? *Advances in neural information processing systems*, 31, 2018.
- Ying Jin, Zhuoran Yang, and Zhaoran Wang. Is pessimism provably efficient for offline rl? In *International Conference on Machine Learning*, pp. 5084–5096. PMLR, 2021.
- Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, Ethan Holly, Mrinal Kalakrishnan, Vincent Vanhoucke, et al. Scalable deep reinforcement learning for vision-based robotic manipulation. In *Conference on Robot Learning*, pp. 651–673. PMLR, 2018.
- Dmitry Kalashnikov, Jacob Varley, Yevgen Chebotar, Benjamin Swanson, Rico Jonschkowski, Chelsea Finn, Sergey Levine, and Karol Hausman. Mt-opt: Continuous multi-task robotic reinforcement learning at scale. *arXiv preprint arXiv:2104.08212*, 2021.

- Ilya Kostrikov, Ofir Nachum, and Jonathan Tompson. Imitation learning via off-policy distribution matching. In *International Conference on Learning Representations*, 2019.
- Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.
- Krikamol Muandet, Kenji Fukumizu, Bharath Sriperumbudur, Bernhard Schölkopf, et al. Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends® in Machine Learning*, 10(1-2): 1–141, 2017.
- Paria Rashidinejad, Banghua Zhu, Cong Ma, Jiantao Jiao, and Stuart Russell. Bridging offline reinforcement learning and imitation learning: A tale of pessimism. *Advances in Neural Information Processing Systems*, 34:11702–11716, 2021.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- Avi Singh, Larry Yang, Kristian Hartikainen, Chelsea Finn, and Sergey Levine. End-to-end robotic reinforcement learning without reward engineering. *environment (eg, by placing additional sensors)*, 2019.
- Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the 27th International Conference on Machine Learning*, 2009.
- Ingo Steinwart and Andreas Christmann. *Support vector machines*. Springer Science & Business Media, 2008.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33:3008–3021, 2020.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- Joel A Tropp et al. An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 8(1-2):1–230, 2015.
- Masatoshi Uehara and Wen Sun. Pessimistic model-based offline reinforcement learning under partial coverage. *arXiv preprint arXiv:2107.06226*, 2021.
- Masatoshi Uehara, Xuezhou Zhang, and Wen Sun. Representation learning for online and offline rl in low-rank mdps. In *International Conference on Learning Representations*, 2021.
- Sattar Vakili, Kia Khezeli, and Victor Picheny. On information gain and regret bounds in gaussian process bandits. In *International Conference on Artificial Intelligence and Statistics*, pp. 82–90. PMLR, 2021.
- Michal Valko, Nathaniel Korda, Rémi Munos, Ilias Flaounas, and Nelo Cristianini. Finite-time analysis of kernelised contextual bandits. In *Uncertainty in Artificial Intelligence*, 2013.
- Andrew Wagenmaker and Aldo Pacchiano. Leveraging offline data in online reinforcement learning. In *International Conference on Machine Learning*, pp. 35300–35338. PMLR, 2023.
- Ruosong Wang, Dean P Foster, and Sham M Kakade. What are the statistical limits of offline rl with linear function approximation? In *International Conference on Learning Representations*, 2020a.
- Ruosong Wang, Ruslan Salakhutdinov, and Lin F Yang. Provably efficient reinforcement learning with general value function approximation. *arXiv preprint arXiv:2005.10804*, 2020b.
- Tengyang Xie, Nan Jiang, Huan Wang, Caiming Xiong, and Yu Bai. Policy finetuning: Bridging sample-efficient offline and online reinforcement learning. *Advances in neural information processing systems*, 34: 27395–27407, 2021.

- Yuling Yan, Gen Li, Yuxin Chen, and Jianqing Fan. Model-based reinforcement learning is minimax-optimal for offline zero-sum markov games. *arXiv preprint arXiv:2206.04044*, 2022.
- Zhuoran Yang, Chi Jin, Zhaoran Wang, Mengdi Wang, and Michael I Jordan. On function approximation in reinforcement learning: Optimism in the face of large state spaces. *arXiv preprint arXiv:2011.04622*, 2020.
- Sing-Yuan Yeh, Fu-Chieh Chang, Chang-Wei Yueh, Pei-Yuan Wu, Alberto Bernacchia, and Sattar Vakili. Sample complexity of kernel-based q-learning. In *International Conference on Artificial Intelligence and Statistics*, pp. 453–469. PMLR, 2023.
- Ming Yin, Yaqi Duan, Mengdi Wang, and Yu-Xiang Wang. Near-optimal offline reinforcement learning with linear representation: Leveraging variance information with pessimism. In *International Conference on Learning Representations*, 2022a.
- Ming Yin, Mengdi Wang, and Yu-Xiang Wang. Offline reinforcement learning with differentiable function approximation is provably efficient. *arXiv preprint arXiv:2210.00750*, 2022b.
- Tianhe Yu, Aviral Kumar, Yevgen Chebotar, Karol Hausman, Sergey Levine, and Chelsea Finn. Conservative data sharing for multi-task offline reinforcement learning. *Advances in Neural Information Processing Systems*, 34:11501–11516, 2021a.
- Tianhe Yu, Aviral Kumar, Rafael Rafailov, Aravind Rajeswaran, Sergey Levine, and Chelsea Finn. Combo: Conservative offline model-based policy optimization. *Advances in neural information processing systems*, 34:28954–28967, 2021b.
- Tianhe Yu, Aviral Kumar, Yevgen Chebotar, Karol Hausman, Chelsea Finn, and Sergey Levine. How to leverage unlabeled data in offline reinforcement learning. In *International Conference on Machine Learning*, pp. 25611–25635. PMLR, 2022.
- Andrea Zanette, David Brandfonbrener, Emma Brunskill, Matteo Pirotta, and Alessandro Lazaric. Frequentist regret bounds for randomized least-squares value iteration. In *International Conference on Artificial Intelligence and Statistics*, pp. 1954–1964. PMLR, 2020.
- Andrea Zanette, Ching-An Cheng, and Alekh Agarwal. Cautiously optimistic policy optimization and exploration with linear function approximation. In *Conference on Learning Theory*, pp. 4473–4525. PMLR, 2021.

A Pessimistic Value Iteration

The Pessimistic Value Iteration (Jin et al., 2021) (PEVI) algorithm constructs an estimated Bellman operator $\widehat{\mathbb{B}}_h$ based on the dataset \mathcal{D} so that $\widehat{\mathbb{B}}_h \widehat{V}_{h+1}^{\mathcal{D}} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ approximates $\mathbb{B}_h \widehat{V}_{h+1}^{\mathcal{D}} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. Here $\widehat{V}_{h+1}^{\mathcal{D}} : \mathcal{S} \rightarrow \mathbb{R}$ is an estimated value function based on \mathcal{D} . Define an uncertainty quantifier with the confidence parameter $\xi \in (0, 1)$ as follows.

Definition A.1 (ξ -Uncertainty Quantifier). *We say $\{\Gamma_h\}_{h=1}^H$ ($\Gamma_h : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$) is a ξ -uncertainty quantifier if the event*

$$\mathcal{E} = \left\{ \left| \left(\widehat{\mathbb{B}}_h \widehat{V}_{h+1}^{\mathcal{D}} \right) (s, a) - \left(\mathbb{B}_h \widehat{V}_{h+1}^{\mathcal{D}} \right) (s, a) \right| \leq \Gamma_h(s, a), \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}, h \in [H] \right\} \quad (13)$$

satisfies $\mathbb{P}(\mathcal{E}) \geq 1 - \xi$.

Algorithm 2 Pessimistic Value Iteration (PEVI): General MDP

- 1: **Input:** Dataset $\mathcal{D} = \{(s_h^\tau, a_h^\tau, r_h^\tau)\}_{\tau, h=1}^{K, H}$.
 - 2: Initialization: Set $\widehat{V}_{H+1}^{\mathcal{D}}(\cdot) \leftarrow 0$.
 - 3: **for** step $h = H, H-1, \dots, 1$ **do**
 - 4: Construct $\left(\widehat{\mathbb{B}}_h \widehat{V}_{h+1}^{\mathcal{D}} \right) (\cdot, \cdot)$ and $\Gamma_h(\cdot, \cdot)$ based on \mathcal{D} .
 - 5: Set $\bar{Q}_h^{\mathcal{D}}(\cdot, \cdot) \leftarrow \left(\widehat{\mathbb{B}}_h \widehat{V}_{h+1}^{\mathcal{D}} \right) (\cdot, \cdot) - \Gamma_h(\cdot, \cdot)$.
 - 6: Set $\widehat{Q}_h^{\mathcal{D}}(\cdot, \cdot) \leftarrow \min \{ \bar{Q}_h^{\mathcal{D}}(\cdot, \cdot), H - h + 1 \}^+$.
 - 7: Set $\widehat{\pi}_h(\cdot | s) \leftarrow \arg \max_{\pi_h} \left\langle \widehat{Q}_h^{\mathcal{D}}(s, \cdot), \pi_h(\cdot | s) \right\rangle_{\mathcal{A}}$.
 - 8: Set $\widehat{V}_h^{\mathcal{D}}(\cdot) \leftarrow \left\langle \widehat{Q}_h^{\mathcal{D}}(\cdot, \cdot), \widehat{\pi}_h(\cdot | \cdot) \right\rangle_{\mathcal{A}}$.
 - 9: **end for**
 - 10: **Output:** $\text{Pess}(\mathcal{D}) = \{\widehat{\pi}_h\}_{h=1}^H, \{\widehat{V}_h^{\mathcal{D}}\}_{h=1}^H$.
-

By equation (13), Γ_h quantifies the uncertainty, which allows us to develop the meta-algorithm in Algorithm 2. Now we introduce the PEVI with the data splitting and the kernel setting. Suppose we have the dataset

Algorithm 3 PEVI: Kernel Approximation with Data Split

- 1: **Input:** Dataset $\mathcal{D} = \{(s_h^\tau, a_h^\tau, r_h^\tau)\}_{\tau, h=1}^{K, H}$ and parameter B and λ .
 - 2: Data split: Randomly split dataset \mathcal{D} into H disjoint and equally sub-datasets $\{\widetilde{\mathcal{D}}_h\}_{h=1}^H$.
 - 3: Initialization: Set $\widehat{V}_{H+1}^{\widetilde{\mathcal{D}}_{H+1:H}}(\cdot) \leftarrow 0$.
 - 4: **for** step $h = H, \dots, 1$ **do**
 - 5: Compute the Gram matrix $K_h^{\widetilde{\mathcal{D}}_h}$, function $k_h^{\widetilde{\mathcal{D}}_h}$ and response vector y_h defined in equation (17) and (18), respectively.
 - 6: Set $\Gamma_h(\cdot, \cdot) \leftarrow B \cdot \lambda^{-1/2} \cdot \left(k(\cdot, \cdot; \cdot, \cdot) - k_h^{\widetilde{\mathcal{D}}_h}(\cdot, \cdot)^\top \left(K_h^{\widetilde{\mathcal{D}}_h} + \lambda I \right)^{-1} k_h^{\widetilde{\mathcal{D}}_h}(\cdot, \cdot) \right)^{1/2}$.
 - 7: Set $\bar{Q}_h^{\widetilde{\mathcal{D}}_{h:H}}(\cdot, \cdot) \leftarrow k_h^{\widetilde{\mathcal{D}}_h}(\cdot, \cdot)^\top \left(K_h^{\widetilde{\mathcal{D}}_h} + \lambda I \right)^{-1} y_h - \Gamma_h(\cdot, \cdot)$.
 - 8: Set $\widehat{Q}_h^{\widetilde{\mathcal{D}}_{h:H}}(\cdot, \cdot) \leftarrow \min \{ \bar{Q}_h^{\widetilde{\mathcal{D}}_{h:H}}(\cdot, \cdot), H - h + 1 \}^+$.
 - 9: Set $\widehat{\pi}_h(\cdot | s) \leftarrow \arg \max_{\pi_h} \left\langle \widehat{Q}_h^{\widetilde{\mathcal{D}}_{h:H}}(s, \cdot), \pi_h(\cdot | s) \right\rangle_{\mathcal{A}}$.
 - 10: Set $\widehat{V}_h^{\widetilde{\mathcal{D}}_{h:H}}(\cdot) \leftarrow \left\langle \widehat{Q}_h^{\widetilde{\mathcal{D}}_{h:H}}(\cdot, \cdot), \widehat{\pi}_h(\cdot | \cdot) \right\rangle_{\mathcal{A}}$.
 - 11: **end for**
 - 12: **Output:** $\text{Pess}(\mathcal{D}) = \{\widehat{\pi}_h\}_{h=1}^H, \{\widehat{V}_h^{\widetilde{\mathcal{D}}_{h:H}}\}_{h=1}^H$.
-

$\mathcal{D} = \{(s_h^\tau, a_h^\tau, r_h^\tau)\}_{h, \tau=1}^{H, N}$, we partition the dataset \mathcal{D} into H disjoint and equally sub-datasets $\{\widetilde{\mathcal{D}}_h\}_{h=1}^H$ and

$|\tilde{\mathcal{D}}_h| := N_h = N/H$ for all $h \in [H]$. Consider the index set $\mathcal{I}_h = \{N_h \cdot (h-1) + 1, \dots, N_h \cdot h\} = \{\tau_{h,1}, \dots, \tau_{h,N_h}\}$ such that $\tilde{\mathcal{D}}_h = \{(s_h^\tau, a_h^\tau, r_h^\tau)\}_{\tau \in \mathcal{I}_h}$. To simplify notation, let $\tilde{\mathcal{D}}_{h:H} = \bigcup_{t=h}^H \tilde{\mathcal{D}}_t$ for $h \in [H]$. Note that $\mathcal{D}_{h:H} = \emptyset$ if $h > H$. Then, we construct the pessimistic value iterations (Jin et al., 2021) with $\mathbb{B}_h \hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}}, \Gamma_h$, and $\hat{V}_h^{\tilde{\mathcal{D}}_{h:H}}$. Note that $\hat{V}_{H+1}^{\tilde{\mathcal{D}}_{H+1:H}} = 0$. Define the empirical mean squared Bellman error (MSBE) as

$$M_h(f) = \sum_{\tau \in \mathcal{I}_h} \left(r_h^\tau + \hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}}(s_{h+1}^\tau) - f(s_h^\tau, a_h^\tau) \right)^2 \quad (14)$$

at each step $h \in [H]$ and for all $f \in \mathcal{H}_k$. Corresponding, we set

$$\left(\mathbb{B}_h \hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}} \right)(z) = \hat{f}_h(z), \quad \text{where } \hat{f}_h = \arg \min_{f \in \mathcal{H}_k} M_h(f) + \lambda \cdot \|f\|_{\mathcal{H}_k}^2 \quad (15)$$

for $\lambda > 0$. Moreover, we construct Γ_h via

$$\Gamma_h(z) = B \cdot \lambda^{-1/2} \cdot \left(K(z, z) - k_h^{\tilde{\mathcal{D}}_h}(z)^\top \left(K_h^{\tilde{\mathcal{D}}_h} + \lambda I \right)^{-1} k_h^{\tilde{\mathcal{D}}_h}(z) \right)^{1/2} \quad (16)$$

where $B > 0$ is a scaling parameter. Note that it is a bonus function defined in (Yang et al., 2020) and that it is clearly a ξ quantifier. Here, the kernel matrix $K_h^{\tilde{\mathcal{D}}_h} \in \mathbb{R}^{N_h \times N_h}$, and the function $k_h^{\tilde{\mathcal{D}}_h} : \mathcal{Z} \rightarrow \mathbb{R}^{N_h}$ as

$$\left[K_h^{\tilde{\mathcal{D}}_h} \right]_{\tau, \tau'} = k \left(z_h^\tau, z_h^{\tau'} \right), \quad k_h^{\tilde{\mathcal{D}}_h}(z) = \begin{pmatrix} k(z_h^{\tau_{h,1}}, z) \\ \vdots \\ k(z_h^{\tau_{h,N_h}}, z) \end{pmatrix} \in \mathbb{R}^{N_h}, \quad (17)$$

for $\tau, \tau' \in \mathcal{I}_h$. The entry of $y_h \in \mathbb{R}^{N_h}$ corresponding to $\tau \in \mathcal{I}_h$ is

$$[y_h]_\tau = r_h^\tau + \hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}}(s_{h+1}^\tau) \quad (18)$$

We construct the pessimistic value iteration with kernel approximation with kernel k by

$$\begin{aligned} \bar{Q}_h^{\tilde{\mathcal{D}}_{h+1:H}}(z) &= \mathbb{B}_h \hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}}(z) - \Gamma_h(z), \\ \hat{Q}_h^{\tilde{\mathcal{D}}_{h+1:H}}(z) &= \min \left\{ \bar{Q}_h^{\tilde{\mathcal{D}}_{h+1:H}}(z), H - h + 1 \right\}^+, \\ \hat{\pi}_h(\cdot | s) &= \arg \max_{\pi_h} \left\langle \hat{Q}_h^{\tilde{\mathcal{D}}_{h+1:H}}(s, \cdot), \pi_h(\cdot | s) \right\rangle_{\mathcal{A}}, \\ \hat{V}_h^{\tilde{\mathcal{D}}_{h+1:H}}(s) &= \left\langle \hat{Q}_h^{\tilde{\mathcal{D}}_{h+1:H}}(s, \cdot), \hat{\pi}_h(\cdot | s) \right\rangle_{\mathcal{A}}, \end{aligned}$$

The algorithm 3 summarizes the entire PEVI algorithm with data splitting.

B Proof of Main Result

B.1 Proof of Proposition 4.1

We present a generalization of (Abbasi-Yadkori et al., 2011) (Theorem 1). Its proof closely mirrors that of the special case where \mathcal{H}_k has a linear kernel. To simplify notation, denote the labeled dataset as $\mathcal{D}_1 = \{s_h^\tau, a_h^\tau, r_h^\tau\}_{\tau, h=1}^{N_1, H}$. Subsequently, we address the following problem:

$$\hat{\theta}_h^t \in \arg \min_{\theta \in \mathcal{H}_k} \left\{ \sum_{\tau=1}^t \left(\langle \theta, \phi(s_h^\tau, a_h^\tau) \rangle_{\mathcal{H}_k} - r_h^\tau \right)^2 + \nu \|\theta\|_{\mathcal{H}}^2 \right\}. \quad (19)$$

Here, $\phi(s_h^t, a_h^t)$ is a column vector for all $h \in [H]$ and $t > 0$. The solution for the above equation is that $\hat{\theta}_h^t = (\Lambda_h^t)^{-1}(\Phi_h^t)^\top Y_h^t$, where $\Lambda_h^t = (\Phi_h^t)^\top \Phi_h^t + \nu I_{\mathcal{H}}$, $\Phi_h^t = [\phi(s_h^1, a_h^1)^\top, \dots, \phi(s_h^t, a_h^t)^\top]^\top$, and $Y_h^t = [r_h^1, \dots, r_h^t]$. Denote $H_h^t = [\epsilon_h^1, \dots, \epsilon_h^t]$ and $K_h^t = \Phi_h^t (\Phi_h^t)^\top$. We then determine the upper bound with $\|\hat{\theta}_h^t - \theta_h^*\|_{\Lambda_h^t}$. We write

$$\begin{aligned}\hat{\theta}_h^t &= ((\Phi_h^t)^\top \Phi_h^t + \nu I_{\mathcal{H}})^{-1} (\Phi_h^t)^\top (\Phi_h^t \theta_h^* + H_h^t) \\ &= ((\Phi_h^t)^\top \Phi_h^t + \nu I_{\mathcal{H}})^{-1} (((\Phi_h^t)^\top \Phi_h^t + \nu I_{\mathcal{H}}) \theta_h^* + (\Phi_h^t)^\top H_h^t) \\ &= ((\Phi_h^t)^\top \Phi_h^t + \nu I_{\mathcal{H}})^{-1} (\Phi_h^t)^\top H_h^t + \theta_h^* - \nu ((\Phi_h^t)^\top \Phi_h^t + \nu I_{\mathcal{H}})^{-1} \theta_h^* \\ &= (\Lambda_h^t)^{-1} (\Phi_h^t)^\top H_h^t + \theta_h^* - \nu (\Lambda_h^t)^{-1} \theta_h^*\end{aligned}\tag{20}$$

That implies

$$\begin{aligned}x^\top \hat{\theta}_h^t - x^\top \theta_h^* &= x^\top (\Lambda_h^t)^{-1} (\Phi_h^t)^\top H_h^t - \nu x^\top (\Lambda_h^t)^{-1} \theta_h^* \\ &= \langle x, (\Phi_h^t)^\top H_h^t \rangle_{(\Lambda_h^t)^{-1}} - \nu \langle x, \theta_h^* \rangle_{(\Lambda_h^t)^{-1}}\end{aligned}\tag{21}$$

where $\langle x, y \rangle_{(\Lambda_h^t)^{-1}} = x^\top (\Lambda_h^t)^{-1} y$ and Λ_h^t is positive definite, then $(\Lambda_h^t)^{-1/2} \leq \nu^{-1/2}$. Using Cauchy-Schwarz inequality, we get

$$\begin{aligned}|x^\top \hat{\theta}_h^t - x^\top \theta_h^*| &\leq \|x\|_{(\Lambda_h^t)^{-1}} (\|(\Phi_h^t)^\top H_h^t\|_{(\Lambda_h^t)^{-1}} + \nu \|\theta_h^*\|_{(\Lambda_h^t)^{-1}}) \\ &\leq \|x\|_{(\Lambda_h^t)^{-1}} (\|(\Phi_h^t)^\top H_h^t\|_{(\Lambda_h^t)^{-1}} + \sqrt{\nu} \|\theta_h^*\|_{\mathcal{H}_k})\end{aligned}\tag{22}$$

where the second inequality uses the fact that

$$\|\theta_h^*\|_{(\Lambda_h^t)^{-1}} = \|(\Lambda_h^t)^{-1/2} \theta_h^*\|_{\mathcal{H}_k} \leq \nu \frac{1}{\sqrt{\nu}} \|\theta_h^*\|_{\mathcal{H}_k} \leq \sqrt{\nu} \|\theta_h^*\|_{\mathcal{H}_k}$$

Next, we show that with probability at least $1 - \delta$

$$\|(\Phi_h^t)^\top H_h^t\|_{(\Lambda_h^t)^{-1}}^2 \leq H^2 \cdot \log \det [\nu I + K_h^t] + 2H^2 \cdot \log(1/\delta)\tag{23}$$

Following (Valko et al., 2013), we will use the following identities:

$$\begin{aligned}((\Phi_h^t)^\top \Phi_h^t + \nu I_{\mathcal{H}}) (\Phi_h^t)^\top &= (\Phi_h^t)^\top (\Phi_h^t (\Phi_h^t)^\top + \nu I) \\ \Rightarrow \Lambda_h^t (\Phi_h^t)^\top &= (\Phi_h^t)^\top (K_h^t + \nu I) \\ \Rightarrow (\Phi_h^t)^\top (K_h^t + \nu I)^{-1} &= (\Lambda_h^t)^{-1} (\Phi_h^t)^\top.\end{aligned}$$

With the basic operation, we get

$$\begin{aligned}\|(\Phi_h^t)^\top H_h^t\|_{(\Lambda_h^t)^{-1}}^2 &= (H_h^t)^\top \Phi_h^t (\Lambda_h^t)^{-1} (\Phi_h^t)^\top H_h^t \\ &= (H_h^t)^\top \Phi_h^t (\Phi_h^t)^\top (K_h^t + \nu I)^{-1} H_h^t \\ &= (H_h^t)^\top K_h^t (K_h^t + \nu I)^{-1} H_h^t\end{aligned}\tag{24}$$

Setting $\nu = 1 + \eta$, for some $\eta > 0$, we have

$$\begin{aligned}(K_h^t + \eta \cdot I) [K_h^t + (1 + \eta) \cdot I]^{-1} &= (K_h^t + \eta \cdot I) [I + (K_h^t + \eta \cdot I)]^{-1} \\ &= [(K_h^t + \eta \cdot I)^{-1} + I]^{-1}\end{aligned}\tag{25}$$

, which implies

$$\begin{aligned}(H_h^t)^\top K_h^t (K_h^t + \nu I)^{-1} H_h^t &\leq (H_h^t)^\top (K_h^t + \eta \cdot I) (K_h^t + (1 + \eta) I)^{-1} H_h^t \\ &= (H_h^t)^\top [(K_h^t + \eta \cdot I)^{-1} + I]^{-1} H_h^t\end{aligned}\tag{26}$$

Applying Lemma C.1 with $E_t = H_h^t$ and $\sigma^2 = 1$, we obtain the following result: with a probability of at least $1 - \delta$, the given condition holds simultaneously for all $t \geq 1$

$$(H_h^t)^\top \left[(K_h^t + \eta \cdot I)^{-1} + I \right]^{-1} H_h^t \leq \log \det [(1 + \eta) \cdot I + K_h^t] + 2 \log(1/\delta)$$

for any $\eta > 0$ and $\delta \in (0, 1)$. Combine equation (24) and (26), for any $t > 0$, we get

$$\|(\Phi_h^t)^\top H_h^t\|_{(\Lambda_h^t)^{-1}} \leq \sqrt{\log \frac{\det [(1 + \eta) \cdot I + K_h^t]}{\delta^2}} \quad (27)$$

with probability at least $1 - \delta$. Therefore, combine the equation (22) and (27), one also has

$$|x^\top \hat{\theta}_h^t - x^\top \theta_h^*| \leq \|x\|_{(\Lambda_h^t)^{-1}} \left(\sqrt{\log \frac{\det [(1 + \eta) \cdot I + K_h^t]}{\delta^2}} + \sqrt{\nu} \|\theta_h^*\|_{\mathcal{H}_k} \right) \quad (28)$$

for all $t \geq 0$. Plugging in $x = \Lambda_h^t (\hat{\theta}_h^t - \theta_h^*)$, we get

$$\|\hat{\theta}_h^t - \theta_h^*\|_{\Lambda_h^t}^2 \leq \|\Lambda_h^t (\hat{\theta}_h^t - \theta_h^*)\|_{(\Lambda_h^t)^{-1}} \left(\sqrt{\log \frac{\det [(1 + \eta) \cdot I + K_h^t]}{\delta^2}} + \sqrt{\nu} \|\theta_h^*\|_{\mathcal{H}_k} \right) \quad (29)$$

Now, $\|\hat{\theta}_h^t - \theta_h^*\|_{\Lambda_h^t} = \|\Lambda_h^t (\hat{\theta}_h^t - \theta_h^*)\|_{(\Lambda_h^t)^{-1}}$ dividing both sides by $\|\hat{\theta}_h^t - \theta_h^*\|_{\Lambda_h^t}$, we get

$$\|\hat{\theta}_h^t - \theta_h^*\|_{\Lambda_h^t} \leq \left(\sqrt{\log \frac{\det [(1 + \eta) \cdot I + K_h^t]}{\delta^2}} + \sqrt{\nu} \|\theta_h^*\|_{\mathcal{H}_k} \right) \quad (30)$$

Finally, we fix $t = N_1$ and let $\hat{\theta}_h = \hat{\theta}_h^{N_1}$, $\Lambda_h = \Lambda_h^{N_1}$ and $K_h^{\mathcal{D}_1} = K_h^{N_1}$, where \mathcal{D}_1 is offline dataset, we get

$$\|\hat{\theta}_h - \theta_h^*\|_{\Lambda_h} \leq \sqrt{\nu} \|\theta_h^*\|_{\mathcal{H}_k} + \sqrt{\log \frac{\det [\nu I + K_h^{\mathcal{D}_1}]}{\delta^2}} \quad (31)$$

Furthermore, observed that $\det (\nu I + K_h^{\mathcal{D}_1}) = \det (I + \nu^{-1} K_h^{\mathcal{D}_1}) \det (\nu I)$. Thus, we have

$$\log \left(\det (\nu I + K_h^{\mathcal{D}_1}) \right) = \log \left(\det (I + \nu^{-1} K_h^{\mathcal{D}_1}) \right) + N_1 \log \nu \leq 2G(N_1, \nu) + N_1(\nu - 1) \quad (32)$$

where $\nu > 1$. Thus, set $\nu = 1 + 1/N_1$, we have

$$\|\hat{\theta}_h - \theta_h^*\|_{\Lambda_h} \leq \sqrt{\nu} \|\theta_h^*\|_{\mathcal{H}_k} + \sqrt{2G(N_1, 1 + 1/N_1) + 1 + \log \frac{1}{\delta^2}} \quad (33)$$

B.2 Proof of Theorem 4.3

The proof of Theorem 4.3 and the related supporting lemmas are given in this part. Recall that we denote the labeled dataset \mathcal{D}_1 , unlabeled dataset \mathcal{D}_2^θ , and $\mathcal{D}^\theta = \{(s_h^\tau, a_h^\tau, \hat{r}_h^{\theta_h}(s_h^\tau, a_h^\tau))\}_{\tau, h=1}^{N, H}$, which is a combination of labeled dataset \mathcal{D}_1 and unlabeled dataset \mathcal{D}_2^θ and $N = N_1 + N_2$. We partition dataset \mathcal{D}^θ into H disjoint and equally sized sub dataset $\{\tilde{\mathcal{D}}_h^\theta\}_{h=1}^H$, where $|\tilde{\mathcal{D}}_h^\theta| = N_h = N/H$. Let $\mathcal{I}_h = \{N_h \cdot (h - 1) + 1, \dots, N_h \cdot h\} = \{\tau_{h,1}, \dots, \tau_{h,N_h}\}$ satisfy $\tilde{\mathcal{D}}_h^\theta = \{(s_h^\tau, a_h^\tau, \hat{r}_h^{\theta_h}(s_h^\tau, a_h^\tau))\}_{\tau \in \mathcal{I}_h}$. Denote $\{\hat{V}_h^{\tilde{\mathcal{D}}_h^\theta}\}_{h=1}^H$ as the estimated value function constructed by PEVI. These estimations are based on datasets $\tilde{\mathcal{D}}_{h:H}$, where rewards have been relabeled using the parameter $\theta := \{\theta_h\}_{h=1}^H$. Furthermore, let $V_1^{\pi, \theta}$ be the value function associated with policy π and the estimated reward function with parameter θ . From equation (4) in Algorithm 1, for all $h \in [H]$, we have

$$\hat{V}_h^{\tilde{\mathcal{D}}_h^\theta}(s_0) \leq \hat{V}_h^{\tilde{\mathcal{D}}_{h:H}^\theta}(s_0), \quad \forall \theta_h \in \mathcal{C}_h(\delta) \quad (34)$$

where s_0 is a unique initial state and $\tilde{\theta} := \{\tilde{\theta}_h\}_{h=1}^H$ is the pessimistic estimation of θ .

Let \mathcal{E}_1 be the event $\theta_h^* \in \mathcal{C}_h(\delta)$, then we have $\mathbb{P}(\mathcal{E}_1) \geq 1 - \delta$ from Proposition 4.1

Denote $\{\widehat{V}_h^{\widetilde{\mathcal{D}}_{h:H}^\theta}\} = \text{Pess}(\mathcal{D})$ from algorithm 3. Let \mathcal{E}_2 be the event where the following inequality holds for each dataset $\widetilde{\mathcal{D}}_h^\theta$ and

$$\left| \left(\widehat{\mathbb{B}}_h \widehat{V}_{h+1}^{\widetilde{\mathcal{D}}_{h+1:H}^\theta} \right) (s, a) - \left(\mathbb{B}_h \widehat{V}_{h+1}^{\widetilde{\mathcal{D}}_{h+1:H}^\theta} \right) (s, a) \right| \leq \Gamma_h(s, a), \forall (s, a) \in \mathcal{S} \times \mathcal{A}, h \in [H], \quad (35)$$

where $\Gamma_h(s, a) = B \|\phi(s, a)\|_{(\Lambda_h^{\widetilde{\mathcal{D}}_h^\theta})^{-1}}$, and

$$B = \begin{cases} C_2 \cdot H \cdot \sqrt{d \log(N/\delta)} & d\text{-finite spectrum,} \\ C_2 \cdot H \cdot \sqrt{(\log N/\delta)^{1+1/d}} & d\text{-exponential decay} \\ C_2 \cdot N^{\frac{d+1}{2(d+m)}} H^{1-\frac{d+1}{2(d+m)}} \cdot \sqrt{\log(N/\delta)} & d\text{-polynomial decay} \end{cases} \quad (36)$$

where $C_1 > 0$ does not depend on N_1 nor H and $C_2 > 0$ does not depend on N nor H . Then, apply Lemma B.3 such that $\mathbb{P}(\mathcal{E}_2) \geq 1 - \delta$.

Condition on $\mathcal{E}_1 \cap \mathcal{E}_2$, we have

$$\begin{aligned} V_1^{\pi^*, \theta^*}(s_0) - \widehat{V}_1^{\pi, \theta^*}(s_0) &= V_1^{\pi^*, \theta^*}(s_0) - \widehat{V}_1^{\widetilde{\mathcal{D}}_{1:H}^{\theta^*}}(s_0) + \widehat{V}_1^{\widetilde{\mathcal{D}}_{1:H}^{\theta^*}}(s_0) - \widehat{V}_1^{\pi, \theta^*}(s_0) \\ &\leq V_1^{\pi^*, \theta^*}(s_0) - \widehat{V}_1^{\widetilde{\mathcal{D}}_{1:H}^{\theta^*}}(s_0) \\ &= V_1^{\pi^*, \theta^*}(s_0) - V_1^{\pi^*, \tilde{\theta}}(s_0) + V_1^{\pi^*, \tilde{\theta}}(s_0) - \widehat{V}_1^{\widetilde{\mathcal{D}}_{1:H}^{\theta}}(s_0) + \widehat{V}_1^{\widetilde{\mathcal{D}}_{1:H}^{\theta}}(s_0) - \widehat{V}_1^{\widetilde{\mathcal{D}}_{1:H}^{\theta^*}}(s_0) \\ &\leq V_1^{\pi^*, \theta^*}(s_0) - V_1^{\pi^*, \tilde{\theta}}(s_0) + V_1^{\pi^*, \tilde{\theta}}(s_0) - \widehat{V}_1^{\widetilde{\mathcal{D}}_{1:H}^{\theta}}(s_0) \\ &= V_1^{\pi^*, \theta^*}(s_0) - V_1^{\pi^*, \widehat{\theta}}(s_0) + V_1^{\pi^*, \widehat{\theta}}(s_0) - V_1^{\pi^*, \tilde{\theta}}(s_0) + V_1^{\pi^*, \tilde{\theta}}(s_0) - \widehat{V}_1^{\widetilde{\mathcal{D}}_{1:H}^{\theta}}(s_0) \\ &\leq \left| V_1^{\pi^*, \theta^*}(s_0) - V_1^{\pi^*, \widehat{\theta}}(s_0) \right| + \left| V_1^{\pi^*, \widehat{\theta}}(s_0) - V_1^{\pi^*, \tilde{\theta}}(s_0) \right| + V_1^{\pi^*, \tilde{\theta}}(s_0) - \widehat{V}_1^{\widetilde{\mathcal{D}}_{1:H}^{\theta}}(s_0) \\ &\leq 2 \sum_{h=1}^H \beta_h(\delta) \mathbb{E}_{\pi^*} \left[\|\phi(s_h, a_h)\|_{(\Lambda_h^{\mathcal{D}_1})^{-1}} \mid s_1 = s_0 \right] \\ &\quad + 2B \sum_{h=1}^H \mathbb{E}_{\pi^*} \left[\|\phi(s_h, a_h)\|_{(\Lambda_h^{\widetilde{\mathcal{D}}_h^\theta})^{-1}} \mid s_1 = s_0 \right] \end{aligned} \quad (37)$$

for s_0 is the unique initial state and $\beta_h(\delta)$ is

$$\beta_h(\delta) = \begin{cases} \sqrt{1 + \frac{1}{N_1}} \mathcal{S} + \sqrt{C_1 \cdot d \cdot \log N_1 + \log(\frac{1}{\delta^2})} & d\text{-finite spectrum,} \\ \sqrt{1 + \frac{1}{N_1}} \mathcal{S} + \sqrt{C_1 \cdot (\log N_1)^{1+\frac{1}{d}} + \log(\frac{1}{\delta^2})} & d\text{-exponential decay} \\ \sqrt{1 + \frac{1}{N_1}} \mathcal{S} + \sqrt{C_1 \cdot (N_1)^{\frac{d+1}{d+m}} \cdot \log(N_1) + \log(\frac{1}{\delta^2})} & d\text{-polynomial decay} \end{cases}$$

Note that the first inequality follows Lemma B.1, while the second inequality follows directly from equation (34), and the last inequality follows Lemma B.2 and Theorem B.4.

Lemma B.1. *Under the event \mathcal{E}_2 , we have*

$$\widehat{V}_1^{\widetilde{\mathcal{D}}_{1:H}^{\theta^*}}(s_0) - \widehat{V}_1^{\pi, \theta^*}(s_0) \leq 0$$

where s_0 is the unique initial state.

Proof. For simplicity, we denote $\widehat{V}_1^{\mathcal{D}}(s_0) = \widehat{V}_1^{\widetilde{\mathcal{D}}^{1:H}}(s_0)$ and $V_1^{\widehat{\pi}}(s_0) = V_1^{\widehat{\pi}, \theta^*}(s_0)$. By Lemma C.3 with $\pi = \pi' = \widehat{\pi}$ and $\{\widehat{Q}_h^{\mathcal{D}}\}_{h=1}^H$ is constructed by PEVI (Algorithm 3), we have

$$\widehat{V}_1^{\mathcal{D}}(s_0) - V_1^{\widehat{\pi}}(s_0) = \sum_{h=1}^H \mathbb{E}_{\widehat{\pi}} \left[\widehat{Q}_h^{\mathcal{D}}(s_h, a_h) - \left(\mathbb{B}_h \widehat{V}_{h+1}^{\mathcal{D}} \right)(s_h, a_h) \mid s_1 = s_0 \right].$$

Recall that $\bar{Q}_h^{\mathcal{D}}(s, a)$ is defined in line 6 in Algorithm 3. For all $h \in [H]$ and all $(s, a) \in \mathcal{S} \times \mathcal{A}$, if $\bar{Q}_h^{\mathcal{D}}(s, a) < 0$, implies $\widehat{Q}_h^{\mathcal{D}}(s, a) = 0$. Then

$$\widehat{Q}_h^{\mathcal{D}}(s, a) - \left(\mathbb{B}_h \widehat{V}_{h+1}^{\mathcal{D}} \right)(s, a) = - \left(\mathbb{B}_h \widehat{V}_{h+1}^{\mathcal{D}} \right)(s, a) < 0$$

as $r_h \in [0, 1]$. Otherwise, $\bar{Q}_h^{\mathcal{D}}(s, a) \geq 0$, we have

$$\begin{aligned} \widehat{Q}_h^{\mathcal{D}}(s, a) - \left(\mathbb{B}_h \widehat{V}_{h+1}^{\mathcal{D}} \right)(s, a) &\leq \bar{Q}_h^{\mathcal{D}}(s, a) - \left(\mathbb{B}_h \widehat{V}_{h+1}^{\mathcal{D}} \right)(s, a) \\ &= \left(\mathbb{B}_h \widehat{V}_{h+1}^{\mathcal{D}} \right)(s, a) - \left(\mathbb{B}_h \widehat{V}_{h+1}^{\mathcal{D}} \right)(s, a) - \Gamma_h(s, a) \leq 0. \end{aligned}$$

Finally, we have

$$\widehat{V}_1^{\mathcal{D}}(s_0) - V_1^{\widehat{\pi}}(s_0) \leq 0 \quad (38)$$

□

Lemma B.2. For policy π^* , and offline dataset $\mathcal{D} = \{(s_h^\tau, a_h^\tau, r_h^\tau)\}_{\tau, h=1}^{N, H}$, and any reward function parameter $\theta_h \in C_h(\delta)$, we have

$$|V_1^{\pi^*, \theta}(s_0) - V_1^{\pi^*, \widehat{\theta}}(s_0)| \leq \sum_{h=1}^H \beta_h(\delta) \mathbb{E}_{\pi^*} \left[\|\phi(s_h, a_h)\|_{(\Lambda_h^{\mathcal{D}})^{-1}} \mid s_1 = s_0 \right] \quad (39)$$

where s_0 is the unique initial state and $\beta_h(\delta)$ is defined in equation (7). Moreover, suppose Assumption 3.1 holds, then $\beta_h(\delta)$ can be written as

$$\beta_h(\delta) = \begin{cases} \sqrt{1 + \frac{1}{N} \mathcal{S}} + \sqrt{C_1 \cdot d \cdot \log N + 1 + \log(\frac{1}{\delta^2})} & d\text{-finite spectrum,} \\ \sqrt{1 + \frac{1}{N} \mathcal{S}} + \sqrt{C_1 \cdot (\log N)^{1+\frac{1}{d}} + 1 + \log(\frac{1}{\delta^2})} & d\text{-exponential decay} \\ \sqrt{1 + \frac{1}{N} \mathcal{S}} + \sqrt{C_1 \cdot (N)^{\frac{m+1}{d+m}} \cdot \log(N) + 1 + \log(\frac{1}{\delta^2})} & d\text{-polynomial decay} \end{cases} \quad (40)$$

Proof. From the definition, we have

$$\begin{aligned} |\widehat{r}_h^{\theta_h}(s, a) - \widehat{r}_h^{\widehat{\theta}_h}(s, a)| &= \left| \phi(s, a)^\top \theta_h - \phi(s, a)^\top \widehat{\theta}_h \right| \\ &\leq \|\theta - \widehat{\theta}_h\|_{\Lambda_h^{\mathcal{D}}} \cdot \|\phi(s, a)\|_{(\Lambda_h^{\mathcal{D}})^{-1}} \\ &\leq \beta_h(\delta) \sqrt{\phi(s, a)^\top (\Lambda_h^{\mathcal{D}})^{-1} \phi(s, a)} \end{aligned} \quad (41)$$

Then, we have

$$\begin{aligned} V_1^{\pi^*, \theta}(s_0) - V_1^{\pi^*, \widehat{\theta}}(s_0) &= \sum_{h=1}^H \mathbb{E}_{\pi^*} \left[\widehat{r}_h^{\theta_h}(s, a) - \widehat{r}_h^{\widehat{\theta}_h}(s, a) \mid s_1 = s_0 \right] \\ &\leq \sum_{h=1}^H \mathbb{E}_{\pi^*} \left[\left| \widehat{r}_h^{\theta_h}(s, a) - \widehat{r}_h^{\widehat{\theta}_h}(s, a) \right| \mid s_1 = s_0 \right] \\ &\leq \sum_{h=1}^H \beta_h(\delta) \mathbb{E}_{\pi^*} \left[\sqrt{\phi(s_h, a_h)^\top (\Lambda_h^{\mathcal{D}})^{-1} \phi(s_h, a_h)} \mid s_1 = s_0 \right] \\ &= \sum_{h=1}^H \beta_h(\delta) \mathbb{E}_{\pi^*} \left[\|\phi(s_h, a_h)\|_{(\Lambda_h^{\mathcal{D}})^{-1}} \mid s_1 = s_0 \right]. \end{aligned} \quad (42)$$

Recall that $\beta_h(\delta) = \sqrt{\nu}\mathcal{S} + \sqrt{2G(N, \nu) + 1 + \log \frac{1}{\delta^2}}$ defined in equation (7). Setting $\nu = 1 + 1/N$ and applying Lemma C.2, we obtain

$$\beta_h(\delta) = \begin{cases} \sqrt{1 + \frac{1}{N}}\mathcal{S} + \sqrt{C_1 \cdot d \cdot \log N + \log(\frac{1}{\delta^2})} & d\text{-finite spectrum,} \\ \sqrt{1 + \frac{1}{N}}\mathcal{S} + \sqrt{C_1 \cdot (\log N)^{1+\frac{1}{d}} + \log(\frac{1}{\delta^2})} & d\text{-exponential decay} \\ \sqrt{1 + \frac{1}{N}}\mathcal{S} + \sqrt{C_1 \cdot (N)^{\frac{m+1}{d+m}} \cdot \log(N) + \log(\frac{1}{\delta^2})} & d\text{-polynomial decay} \end{cases} \quad (43)$$

for some sufficient large C_1 and C_1 is an absolute constant that does not depend on N_1 nor H . \square

Lemma B.3. Suppose Assumption 3.2 and 3.1 hold, with dataset $\mathcal{D} = \{(s_h^\tau, a_h^\tau, r_h^\tau)\}_{h,\tau=1}^{H,N}$, we set $\lambda = 1 + 1/N$ and $B > 0$ satisfies

$$2(1 + \frac{1}{N})R_Q^2 + 8G(N/H, 1 + 1/N) + 2/H + 8\log(H/\xi) \leq (B/H)^2 \quad (44)$$

in Algorithm 3. Then $\Gamma_h(s, a) = B \cdot \|\phi(s, a)\|_{(\Lambda_h^{\mathcal{D}_h})^{-1}}$ is a ξ -quantifier where $\tilde{\mathcal{D}}_h$ is defined in Theorem 4.3. That is, for dataset \mathcal{D} , the following inequality holds,

$$\left| \left(\widehat{\mathbb{B}}_h \widehat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}} \right) (s, a) - \left(\mathbb{B}_h \widehat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}} \right) (s, a) \right| \leq \Gamma_h(s, a), \forall (s, a) \in \mathcal{S} \times \mathcal{A}, h \in [H] \quad (45)$$

with $\mathbb{P}(\mathcal{E}_2) \geq 1 - \xi$, where \mathcal{E}_2 is defined in equation(35). In particular, B is given by

$$B = \begin{cases} C \cdot H \cdot \sqrt{d \log(N/\xi)} & d\text{-finite spectrum,} \\ C \cdot H \cdot \sqrt{(\log N/\xi)^{1+1/d}} & d\text{-exponential decay} \\ C \cdot N^{\frac{m+1}{2(d+m)}} H^{1-\frac{m+1}{2(d+m)}} \cdot \sqrt{\log(N/\xi)} & d\text{-polynomial decay} \end{cases} \quad (46)$$

for some absolute constant C that does not depend on N nor H .

Proof. We present the offline reinforcement setting of (Yang et al., 2020) and combine the data split skill from (Xie et al., 2021) (Rashidinejad et al., 2021). Recall that we partition dataset \mathcal{D} into H disjoint and equally sized sub dataset $\{\tilde{\mathcal{D}}_h\}_{h=1}^H$, where $|\tilde{\mathcal{D}}_h| = N_h = N/H$. Let $\mathcal{I}_h = \{N_h \cdot (h-1) + 1, \dots, N_h \cdot h\} = \{\tau_{h,1}, \dots, \tau_{h,N_h}\}$ satisfy $\tilde{\mathcal{D}}_h = \{(s_h^\tau, a_h^\tau, r_h^\tau)\}_{\tau \in \mathcal{I}_h}$. Denote the operator $\Phi_h^{\tilde{\mathcal{D}}_h} : \mathcal{H}_k \rightarrow \mathbb{R}^{N_h}$ and $\Lambda_h^{\tilde{\mathcal{D}}_h} : \mathcal{H}_k \rightarrow \mathcal{H}_k$ as

$$\Phi_h^{\tilde{\mathcal{D}}_h} = \begin{pmatrix} \phi(z_h^{\tau_{h,1}})^\top \\ \vdots \\ \phi(z_h^{\tau_{h,N_h}})^\top \end{pmatrix} = \begin{pmatrix} k(\cdot, z_h^{\tau_{h,1}})^\top \\ \vdots \\ k(\cdot, z_h^{\tau_{h,N_h}})^\top \end{pmatrix}, \Lambda_h^{\tilde{\mathcal{D}}_h} = \lambda \cdot I_{\mathcal{H}} + (\Phi_h^{\tilde{\mathcal{D}}_h})^\top \Phi_h^{\tilde{\mathcal{D}}_h} \quad (47)$$

For notation simplicity, let $K_h = K_h^{\tilde{\mathcal{D}}_h}$, $k_h(z) = k_h^{\tilde{\mathcal{D}}_h}(z)$, and $\Lambda_h = \Lambda_h^{\tilde{\mathcal{D}}_h}$. For $h \in [H]$, the solution \hat{f}_h of the kernel ridge regression in equation (15) is that

$$\hat{f}_h(\cdot) = \sum_{\tau \in \mathcal{I}_h} [\hat{\alpha}_h]_\tau k(z_h^\tau, \cdot) = \Phi_h^\top \hat{\alpha}_h \quad (48)$$

where $\hat{\alpha}_h = (K_h + \lambda \cdot I)^{-1} y_h$. Note that both matrix $\Phi_h \Phi_h^\top + \lambda \cdot I$ and $\Phi_h^\top \Phi_h + \lambda \cdot I_{\mathcal{H}}$ are positive definite, following (Valko et al., 2013), we have

$$\Phi_h^\top (\Phi_h \Phi_h^\top + \lambda \cdot I)^{-1} = (\Phi_h^\top \Phi_h + \lambda \cdot I_{\mathcal{H}})^{-1} \Phi_h^\top \quad (49)$$

Then, the fitted value function $\widehat{\mathbb{B}}_h \widehat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}}$ is that

$$\hat{f}_h(z) = \langle \hat{f}_h, k(\cdot, z) \rangle_{\mathcal{H}_k} = k_h(z)^\top \hat{\alpha}_h \quad (50)$$

Furthermore, in combination with the equation (49), $k_h(z)$ can be written as $k_h(z) = \Phi_h \phi(z)$, we have

$$\begin{aligned}
\phi(z) &= \Lambda_h^{-1} \Lambda_h \phi(z) \\
&= \Lambda_h^{-1} [\Phi_h^\top \Phi_h + \lambda \cdot I_{\mathcal{H}_k}] \phi(z) \\
&= \Lambda_h^{-1} \Phi_h^\top \Phi_h \phi(z) + \lambda \Lambda_h^{-1} \phi(z) \\
&= \Phi_h^\top (K_h + \lambda \cdot I)^{-1} \Phi_h \phi(z) + \lambda \Lambda_h^{-1} \phi(z) \\
&= \Phi_h^\top (K_h + \lambda \cdot I)^{-1} k_h(z) + \lambda \Lambda_h^{-1} \phi(z)
\end{aligned} \tag{51}$$

where the forth equality follows equation (49). Recall that $\tilde{Q}_h^{\mathcal{D}_{h+1:H}}(z) = \min\{\tilde{Q}_h^{\mathcal{D}_{h+1:H}}(z), H - h + 1\}^+ = \min\{k_h(z)^\top \hat{\alpha}_h - \Gamma_h(z), H - h + 1\}^+$ in Algorithm 3. Since $\tilde{Q}_{h+1}^{\mathcal{D}_{h+1:H}} \in [0, H]$ for all $h \in [H]$, by Assumption 3.2, we have $\mathbb{B}_h \hat{V}_{h+1}^{\mathcal{D}_{h+1:H}} \in \mathcal{Q}^*$, i.e., $\left| \mathbb{B}_h \hat{V}_{h+1}^{\mathcal{D}_{h+1:H}} \right|_{\mathcal{H}} \leq R_Q H$. There exists $f_h \in \mathcal{Q}^*$ such that $f_h = \mathbb{B}_h \hat{V}_{h+1}^{\mathcal{D}_{h+1:H}}$ and $\mathbb{B}_h \hat{V}_{h+1}^{\mathcal{D}_{h+1:H}}(z) = \langle f_h, k(\cdot, z) \rangle_{\mathcal{H}_k} = \phi(z)^\top f_h$ by the feature representation of RKHS. For any $h \in [H]$,

$$\begin{aligned}
&\left| \left(\hat{\mathbb{B}}_h \hat{V}_{h+1}^{\mathcal{D}_{h+1:H}} \right) (s, a) - \left(\mathbb{B}_h \hat{V}_{h+1}^{\mathcal{D}_{h+1:H}} \right) (s, a) \right| \\
&= \left| \hat{f}_h(s, a) - \phi(s, a)^\top f_h \right| \\
&= \left| k_h(z)^\top (K_h + \lambda \cdot I)^{-1} y_h - k_h^\top(z) (K_h + \lambda \cdot I)^{-1} \Phi_h f_h - \lambda \phi(z)^\top \Lambda_h^{-1} f_h \right| \\
&\leq |k_h(z)^\top (K_h + \lambda \cdot I)^{-1} (y_h - \Phi_h f_h)| + |\lambda \phi(z)^\top \Lambda_h^{-1} f_h| \\
&\leq (A) + (B)
\end{aligned} \tag{52}$$

where the second equality follows equation (51). Next, we bound (A) and (B) separately. By Cauchy-Schwarz inequality,

$$\begin{aligned}
(B) &= |\lambda \phi(z)^\top \Lambda_h^{-1} f_h| = \lambda \langle \Lambda_h^{-1} \phi(z), f_h \rangle_{\mathcal{H}_k} \\
&\leq \lambda \cdot \left\| \Lambda_h^{-1} \phi(z) \right\|_{\mathcal{H}_k} \cdot \|f_h\|_{\mathcal{H}_k} \\
&= \lambda \cdot \left\| \Lambda_h^{-1/2} \Lambda_h^{-1/2} \phi(z) \right\|_{\mathcal{H}_k} \cdot \|f_h\|_{\mathcal{H}_k} \\
&\leq R_Q H \cdot \lambda^{1/2} \cdot \left\| \Lambda_h^{-1/2} \phi(z) \right\|_{\mathcal{H}_k} \\
&= R_Q H \cdot \lambda^{1/2} \cdot \left\| \phi(z) \right\|_{\Lambda_h^{-1}}
\end{aligned} \tag{53}$$

Furthermore, y_h is defined in equation (18) and Section 3.1, the τ -th entry of $(y_h - \Phi_h f_h)$ can be written as

$$\begin{aligned}
[y_h]_\tau - [\Phi_h f_h]_\tau &= r_h^\tau + \hat{V}_{h+1}^{\mathcal{D}_{h+1:H}}(s_{h+1}^\tau) - \phi(s_h^\tau, a_h^\tau) f_h \\
&= r_h^\tau + \hat{V}_{h+1}^{\mathcal{D}_{h+1:H}}(s_{h+1}^\tau) - \mathbb{B}_h \hat{V}_{h+1}^{\mathcal{D}_{h+1:H}}(s_h^\tau, a_h^\tau) \\
&= \hat{V}_{h+1}^{\mathcal{D}_{h+1:H}}(s_{h+1}^\tau) - \mathbb{P}_h \hat{V}_{h+1}^{\mathcal{D}_{h+1:H}}(s_h^\tau, a_h^\tau) + \epsilon_h^\tau
\end{aligned} \tag{54}$$

With equation (49) and $k_h(z) = \Phi_h \phi(z)$, (A) can be written as

$$\begin{aligned}
(A) &= |k_h(z)^\top (K_h + \lambda \cdot I)^{-1} (y_h - \Phi_h f_h)| \\
&= |\phi(s, a)^\top \Phi_h^\top (K_h + \lambda \cdot I)^{-1} (y_h - \Phi_h f_h)| \\
&= |\phi(s, a)^\top \Lambda_h^{-1} \Phi_h^\top (y_h - \Phi_h f_h)| \\
&= \left| \phi(s, a)^\top \Lambda_h^{-1} \sum_{\tau \in \mathcal{I}_h} \phi(s_h^\tau, a_h^\tau) \left[\hat{V}_{h+1}^{\mathcal{D}_{h+1:H}}(s_{h+1}^\tau) - \mathbb{P}_h \hat{V}_{h+1}^{\mathcal{D}_{h+1:H}}(s_h^\tau, a_h^\tau) + \epsilon_h^\tau \right] \right| \\
&\leq \left\| \phi(s, a) \right\|_{\Lambda_h^{-1}} \cdot \left\| \sum_{\tau \in \mathcal{I}_h} \phi(s_h^\tau, a_h^\tau) \left[\hat{V}_{h+1}^{\mathcal{D}_{h+1:H}}(s_{h+1}^\tau) - \mathbb{P}_h \hat{V}_{h+1}^{\mathcal{D}_{h+1:H}}(s_h^\tau, a_h^\tau) + \epsilon_h^\tau \right] \right\|_{\Lambda_h^{-1}},
\end{aligned} \tag{55}$$

where the last inequality uses Cauchy-Schwarz inequality.
For $h \in [H - 1]$, we define the filtration

$$\mathcal{F}_h = \sigma\left(\tilde{\mathcal{D}}_1 \cup \dots \cup \tilde{\mathcal{D}}_h\right),$$

where $\sigma(\cdot)$ is the σ -algebra generated by the set of random variables. Let

$$\varepsilon_h^\tau = (\hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}}(s_{h+1}^\tau) + \epsilon_h^\tau) - \mathbb{P}_h \hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}}(s_h^\tau, a_h^\tau)$$

is adapted to the filtration $\{\mathcal{F}_{h+1}\}_{h=1}^{H-1}$. Then

$$\mathbb{E}\left[\hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}}(s_{h+1}^\tau) + \epsilon_h^\tau \mid \mathcal{F}_h\right] = \mathbb{E}\left[\hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}}(s_{h+1}) + \epsilon_h^\tau \mid s_h = s_h^\tau, a_h = a_h^\tau\right] = \mathbb{P}_h \hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}}(s_h^\tau, a_h^\tau)$$

Thus, we have $\mathbb{E}[\varepsilon_h^\tau \mid \mathcal{F}_h] = 0$. Applying Lemma C.1 to $\epsilon_\tau = \varepsilon_h^\tau$ and $\sigma^2 = 2H^2$ as

$$\hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}}(s_{h+1}^\tau) - \mathbb{P}_h \hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}}(s_h^\tau, a_h^\tau) \in [-H, H] \quad (56)$$

and ϵ_h^τ is 1-sub Gaussian noise, for any $\eta > 0$ and $\xi > 0$, it holds probability at least $1 - \xi/H$ that

$$\begin{aligned} E_h^\top \left[(K_h + \eta \cdot I)^{-1} + I \right]^{-1} E_h \\ \leq 2H^2 \cdot \log \det[(1 + \eta) \cdot I + K_h] + 4H^2 \cdot \log(H/\xi) \end{aligned} \quad (57)$$

where $E_h = \begin{pmatrix} \varepsilon_h^{\tau_{h,1}} \\ \vdots \\ \varepsilon_h^{\tau_{h,N_h}} \end{pmatrix}$. Using the equation (57), we get

$$\begin{aligned} \left\| \sum_{\tau \in \mathcal{I}_h} \phi(s_h^\tau, a_h^\tau) \varepsilon_h^\tau \right\|_{\Lambda_h^{-1}}^2 &= E_h^\top \Phi_h (\Phi_h^\top \Phi_h + \lambda \cdot I_{\mathcal{H}})^{-1} \Phi_h^\top E_h \\ &= E_h^\top \Phi_h \Phi_h^\top (\Phi_h \Phi_h^\top + \lambda \cdot I)^{-1} E_h \\ &= E_h^\top K_h (K_h + \lambda \cdot I)^{-1} E_h \end{aligned} \quad (58)$$

For $\lambda = \eta + 1 > 1$ and $\eta > 0$, we have

$$\begin{aligned} E_h^\top K_h (K_h + \lambda \cdot I)^{-1} E_h &= E_h^\top K_h (K_h + (\eta + 1) \cdot I)^{-1} E_h \\ &\leq E_h^\top (K_h + \eta \cdot I) (K_h + (\eta + 1) \cdot I)^{-1} E_h \\ &= E_h^\top \left[(K_h + \eta \cdot I)^{-1} + I \right]^{-1} E_h \end{aligned} \quad (59)$$

where the first equality follows the fact $\left((K_h + \eta \cdot I)^{-1} + I\right)^{-1} = (K_h + \eta \cdot I) (K_h + (1 + \eta) \cdot I)^{-1}$. For any fixed $\xi > 0$, combining equation (57), (58), and (59), we get

$$\left\| \sum_{\tau \in \mathcal{I}_h} \phi(s_h^\tau, a_h^\tau) \varepsilon_h^\tau \right\|_{\Lambda_h^{-1}}^2 \leq 2H^2 \cdot \log \det[\lambda \cdot I + K_h] + 4H^2 \cdot \log(H/\xi) \quad (60)$$

holds simultaneously for all $h \in [H]$ with probability at least $1 - \xi$. Clearly, $\lambda \cdot I + K_h = (\lambda \cdot I) (I + K_h/\lambda)$, then

$$\begin{aligned} \log \det(\lambda \cdot I + K_h) &= N_h \log \lambda + \log \det(I + K_h/\lambda) \\ &\leq N_h(\lambda - 1) + \log \det(I + K_h/\lambda), \end{aligned} \quad (61)$$

where $N_h = |\mathcal{I}_h| = N/H$. Hence, for any $\varepsilon > 0$ and $\lambda > 1$,

$$\left\| \sum_{\tau \in \mathcal{I}_h} \phi(s_h^\tau, a_h^\tau) \varepsilon_h^\tau \right\|_{\Lambda_h^{-1}}^2 \leq 2H^2 \cdot \log \det [I + K_h/\lambda] + H^2 N_h (\lambda - 1) + 4H^2 \cdot \log(H/\xi) \quad (62)$$

holds simultaneously for all $h \in [H]$ with probability at least $1 - \xi$. Finally, combine equation (52), (53), (55), and (60) and take $\lambda = 1 + \frac{1}{N}$, we get

$$\begin{aligned} & \left| \left(\widehat{\mathbb{B}}_h \widehat{V}_{h+1}^{\widetilde{\mathcal{D}}_{h+1:H}} \right)(z) - \left(\mathbb{B}_h \widehat{V}_{h+1}^{\widetilde{\mathcal{D}}_{h+1:H}} \right)(z) \right| \\ & \leq \|\phi(z)\|_{\Lambda_h^{-1}} \left[R_Q H \cdot \sqrt{\lambda} + \left\| \sum_{\tau \in \mathcal{I}_h} \phi(s_h^\tau, a_h^\tau) \varepsilon_h^\tau \right\|_{\Lambda_h^{-1}} \right] \\ & \leq \|\phi(z)\|_{\Lambda_h^{-1}} \left[2\lambda R_Q^2 H^2 + 2 \left\| \sum_{\tau \in \mathcal{I}_h} \phi(s_h^\tau, a_h^\tau) \varepsilon_h^\tau \right\|_{\Lambda_h^{-1}}^2 \right]^{1/2} \\ & \leq \|\phi(z)\|_{\Lambda_h^{-1}} \left[2\lambda R_Q^2 H^2 + 8H^2 \cdot G(N/H, \lambda) + 2H^2 N_h (\lambda - 1) + 8H^2 \cdot \log(H/\xi) \right]^{1/2} \\ & \leq \|\phi(z)\|_{\Lambda_h^{-1}} \left[2(1 + \frac{1}{N}) R_Q^2 H^2 + 8H^2 \cdot G(N/H, 1 + 1/N) + 2H + 8H^2 \cdot \log(H/\xi) \right]^{1/2} \\ & \leq B \cdot \|\phi(z)\|_{\Lambda_h^{-1}} = \Gamma_h(z), \quad \forall z \in \mathcal{S} \times \mathcal{A}, \end{aligned} \quad (63)$$

where the second inequality follows from $\sqrt{x} + \sqrt{y} \leq \sqrt{2(x^2 + y^2)}$. Thus, $\{\Gamma_h\}_{h \in [H]}$ is a ξ -uncertainty quantifier. To give explicit expressions for B , we distinguish three cases according to the spectrum of k .

- Case I (d -finite spectrum): By Lemma C.2, we have $G(N/H, 1 + 1/N) \leq C_k \cdot d \log(N/H)$, where C_k is absolute constant that depends on m , and d . Then, B^2 equals to

$$\begin{aligned} & 2(1 + \frac{1}{N}) R_Q^2 H^2 + 8H^2 \cdot G(N/H, 1 + 1/N) + 2H + 8H^2 \cdot \log(H/\xi) \\ & \leq 4R_Q^2 H^2 + 8H^2 \cdot C_k \cdot d \log(N/H) + 2H + 8H^2 \cdot \log(H/\xi) \\ & \leq C^2 \cdot H^2 \cdot d \log(N/\xi) \end{aligned} \quad (64)$$

for sufficient large $C > 0$. Hence, we take $B = C \cdot H \cdot \sqrt{d \log(N/\xi)}$.

- Case II (d -exponential decay): By Lemma C.2, we get

$$G(N/H, 1 + 1/N) \leq C_k \cdot (\log(N/H))^{1+1/d} \quad (65)$$

where C_k is absolute constant that only depends on m and d . Then, B^2 equals to

$$\begin{aligned} & 2(1 + \frac{1}{N}) R_Q^2 H^2 + 8H^2 \cdot G(N/H, 1 + 1/N) + 2H + 8H^2 \cdot \log(H/\xi) \\ & \leq 4R_Q^2 H^2 + 8H^2 \cdot C_k \cdot (\log(N/H))^{1+1/d} + 2H + 8H^2 \cdot \log(H/\xi) \\ & \leq C^2 H^2 \cdot \left[(\log(N/H))^{1+1/d} + \log(H/\xi) \right] \\ & \leq C^2 H^2 \cdot [\log(N/H) + \log(H/\xi)]^{1+1/d} \\ & \leq C^2 \cdot H^2 \cdot [\log(N/\xi)]^{1+1/d} \end{aligned} \quad (66)$$

for sufficient large $C > 0$. Hence, we take $B = C \cdot H \cdot \sqrt{(\log N/\xi)^{1+1/d}}$.

- Case III: (d -polynomial decay) By Lemma C.2, we get

$$G(N/H, 1 + 1/N) \leq C_k \cdot (N/H)^{\frac{m+1}{d+m}} \cdot \log(N/H) \quad (67)$$

where C_k is absolute constant that only depends on m and d . Then, B^2 equals to

$$\begin{aligned} & 2(1 + \frac{1}{N})R_Q^2 H^2 + 8H^2 \cdot G(N/H, 1 + 1/N) + 2H + 8H^2 \cdot \log(H/\xi) \\ & \leq 4R_Q^2 H^2 + 8H^2 \cdot C_k \cdot (N/H)^{\frac{d+1}{d+m}} \cdot \log(N/H) + 2H + 8H^2 \cdot \log(H/\xi) \\ & \leq C^2 H^{2-\frac{m+1}{d+m}} N^{\frac{m+1}{d+m}} \cdot [\log(N/H) + \log(H/\xi)] \\ & \leq C^2 H^{2-\frac{m+1}{d+m}} N^{\frac{m+1}{d+m}} \cdot [\log(N/\xi)] \end{aligned} \quad (68)$$

for sufficient large $C > 0$. Thus, it suffices to choose $B = C \cdot N^{\frac{m+1}{2(d+m)}} H^{1-\frac{m+1}{2(d+m)}} \cdot \sqrt{\log(N/\xi)}$.

□

Theorem B.4. Under Assumption 3.2 and 3.1, we set $\lambda = 1 + 1/N$ and B is defined as equation (46). Then with probability at least $1 - \xi$, it holds that

$$V_1^{\pi^*, \tilde{\theta}}(s_0) - \hat{V}_1^{\tilde{\mathcal{D}}_{1:H}^{\tilde{\theta}}}(s_0) \leq 2B \sum_{h=1}^H \mathbb{E}_{\pi^*} \left[\left\| \phi(s_h, a_h) \right\|_{(\Lambda_h^{\tilde{\mathcal{D}}_{h:H}^{\tilde{\theta}}})^{-1}} \mid s_1 = s_0 \right] \quad (69)$$

Proof. Recall that the ξ -quantifier satisfies the following inequality with probability at least $1 - \xi$:

$$\left| \left(\mathbb{B}_h \hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}^{\tilde{\theta}}} \right) (s, a) - \left(\mathbb{B}_h \hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}^{\tilde{\theta}}} \right) (s, a) \right| \leq \Gamma_h(s, a), \quad (70)$$

for all $(s, a) \in \mathcal{S} \times \mathcal{A}, h \in [H]$. Define $\hat{\pi} = \{\hat{\pi}_h\}_{h=1}^H$ as the policy such that $\hat{V}_h^{\tilde{\mathcal{D}}_{h:H}^{\tilde{\theta}}}(s) = \left\langle \hat{Q}_h^{\tilde{\mathcal{D}}_{h:H}^{\tilde{\theta}}}(s, \cdot), \hat{\pi}_h(\cdot \mid s) \right\rangle_{\mathcal{A}}$. For simplicity, we denote $\delta_h(s, a) = \left(\mathbb{B}_h \hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}^{\tilde{\theta}}} \right) (s, a) - \hat{Q}_h^{\tilde{\mathcal{D}}_{h:H}^{\tilde{\theta}}}(s, a)$. Applying Lemma C.3 with $\pi = \hat{\pi}$, and $\pi' = \pi^*$, we have

$$\begin{aligned} \hat{V}_1^{\tilde{\mathcal{D}}_{1:H}^{\tilde{\theta}}}(s_0) - V_1^{\pi^*, \tilde{\theta}}(s_0) &= \sum_{h=1}^H \mathbb{E}_{\pi^*} \left[\left\langle \hat{Q}_h^{\tilde{\mathcal{D}}_{h:H}^{\tilde{\theta}}}(s_h, \cdot), \hat{\pi}_h(\cdot \mid s_h) - \pi_h^*(\cdot \mid s_h) \right\rangle_{\mathcal{A}} \mid s_1 = s_0 \right] \\ &+ \sum_{h=1}^H \mathbb{E}_{\pi^*} \left[\hat{Q}_h^{\tilde{\mathcal{D}}_{h:H}^{\tilde{\theta}}}(s_h, a_h) - \left(\mathbb{B}_h \hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}^{\tilde{\theta}}} \right) (s_h, a_h) \mid s_1 = s_0 \right] \\ &= \sum_{h=1}^H \mathbb{E}_{\pi^*} \left[\left\langle \hat{Q}_h^{\tilde{\mathcal{D}}_{h:H}^{\tilde{\theta}}}(s_h, \cdot), \hat{\pi}_h(\cdot \mid s_h) - \pi_h^*(\cdot \mid s_h) \right\rangle_{\mathcal{A}} \mid s_1 = s_0 \right] \\ &- \sum_{h=1}^H \mathbb{E}_{\pi^*} [\delta_h(s_h, a_h) \mid s_1 = s_0] \end{aligned} \quad (71)$$

where \mathbb{E}_{π^*} is taken with respect to the trajectory generated by π^* . Since $\hat{\pi}$ is greedy with respect to $\hat{Q}_h^{\tilde{\mathcal{D}}_{h:H}^{\tilde{\theta}}}$, then

$$\begin{aligned} V_1^{\pi^*, \tilde{\theta}}(s_0) - \hat{V}_1^{\tilde{\mathcal{D}}_{1:H}^{\tilde{\theta}}}(s_0) &= \sum_{h=1}^H \mathbb{E}_{\pi^*} [\delta_h(s_h, a_h) \mid s_1 = s_0] \\ &+ \sum_{h=1}^H \mathbb{E}_{\pi^*} \left[\left\langle \hat{Q}_h^{\tilde{\mathcal{D}}_{h:H}^{\tilde{\theta}}}(s_h, \cdot), \pi_h^*(\cdot \mid s_h) - \hat{\pi}_h(\cdot \mid s_h) \right\rangle_{\mathcal{A}} \mid s_1 = s_0 \right] \\ &= \sum_{h=1}^H \mathbb{E}_{\pi^*} [\delta_h(s_h, a_h) \mid s_1 = s_0] \end{aligned} \quad (72)$$

Recall that the construction of $\tilde{Q}_h^{\tilde{\mathcal{D}}_{h:H}^\theta}$ in Line 5 in Algorithm 2. For all $h \in [H]$ and all $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have

$$\begin{aligned}\tilde{Q}_h^{\tilde{\mathcal{D}}_{h:H}^\theta}(s, a) &= \mathbb{B}_h \hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}^\theta}(s, a) - \Gamma_h(s, a) \\ &\leq \mathbb{B}_h \hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}^\theta}(s, a) \leq H - h + 1\end{aligned}\quad (73)$$

where the first inequality follows the definition of $\Gamma_h(s, a)$ and the second inequality follows that $r_h \in [0, 1]$ and $\hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1:H}^\theta}(s, a) \in [H - h]$. Then, we have

$$\tilde{Q}_h^{\tilde{\mathcal{D}}_{h:H}^\theta}(s, a) = \left\{ \tilde{Q}_h^{\tilde{\mathcal{D}}_{h:H}^\theta}(s, a), 0 \right\}^+ = \max \left\{ \tilde{Q}_h^{\tilde{\mathcal{D}}_{h:H}^\theta}(s, a), 0 \right\} \geq \tilde{Q}_h^{\tilde{\mathcal{D}}_{h:H}^\theta}(s, a) \quad (74)$$

Then, $\delta_h(s, a)$ can be written as

$$\begin{aligned}\delta_h(s, a) &= \left(\mathbb{B}_h \hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1}^\theta} \right)(s, a) - \tilde{Q}_h^{\tilde{\mathcal{D}}_{h:H}^\theta}(s, a) \\ &\leq \left(\mathbb{B}_h \hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1}^\theta} \right)(s, a) - \tilde{Q}_h^{\tilde{\mathcal{D}}_{h:H}^\theta}(s, a) \\ &\leq \left(\mathbb{B}_h \hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1}^\theta} \right)(s, a) - \left(\mathbb{B}_h \hat{V}_{h+1}^{\tilde{\mathcal{D}}_{h+1}^\theta} \right)(s, a) + \Gamma_h(s, a) \leq 2\Gamma_h(s, a)\end{aligned}\quad (75)$$

where the inequality follows line 4 in Algorithm 2. Hence we have $\delta_h(s, a) \leq 2\Gamma_h(s, a)$. Combine Lemma B.3, equation (72), and equation (75), we get

$$\begin{aligned}V_1^{\pi^*, \tilde{\theta}}(s_0) - \hat{V}_1^{\tilde{\mathcal{D}}_{1:H}^\theta}(s_0) &\leq 2 \sum_{h=1}^H \mathbb{E}_{\pi^*} [\Gamma_h(s_h, a_h) \mid s_1 = s_0] \\ &\leq 2B \sum_{h=1}^H \mathbb{E}_{\pi^*} \left[\|\phi(s_h, a_h)\|_{(\Lambda_h^{\tilde{\mathcal{D}}_{h:H}^\theta})^{-1}} \mid s_1 = s_0 \right]\end{aligned}\quad (76)$$

where B satisfies equation (44). \square

B.3 Proof of Proposition 4.5

Denote $\Lambda_h = \Lambda_h^{\tilde{\mathcal{D}}_h}$, $\Lambda'_h = \Lambda_h^{\tilde{\mathcal{D}}'_h}$, and $K'_h = K_h^{\tilde{\mathcal{D}}'_h}$. Note that Λ_h is a self-adjoint operator, then

$$\Lambda'_h = \Lambda_h^{1/2} \left(I_{\mathcal{H}_k} + \Lambda_h^{-1/2} \phi(z) \phi(z)^\top \Lambda_h^{-1/2} \right) \Lambda_h^{1/2} \quad (77)$$

We take log det on both sides with the equation (77), then

$$\begin{aligned}\log \det (\Lambda'_h) &= \log \det (\Lambda_h) + \log \det \left(I_{\mathcal{H}_k} + \Lambda_h^{-1/2} \phi(z) \phi(z)^\top \Lambda_h^{-1/2} \right) \\ &= \log \det (\Lambda_h) + \log (1 + \phi(z)^\top \Lambda_h^{-1} \phi(z))\end{aligned}\quad (78)$$

Note that $\det (\Lambda_h) = \det (\lambda I + K_h)$, and $\det (\Lambda'_h) = \det (\lambda I + K'_h)$ for $\lambda \geq 1$ because $\phi(z)^\top \Lambda_h^{-1} \phi(z) \leq 1$, we have

$$\begin{aligned}\phi(z)^\top \Lambda_h^{-1} \phi(z) &\leq 2 \log (1 + \phi(z)^\top \Lambda_h^{-1} \phi(z)) \\ &= 2 \cdot [\log \det (\Lambda'_h) - \log \det (\Lambda_h)] \\ &= 2 \cdot [\log \det (I + K'_h/\lambda) - \log \det (I + K_h/\lambda)]\end{aligned}$$

Moreover, by equation (2), we have

$$\phi(z)^\top \Lambda_h^{-1} \phi(z) \leq 2 \cdot [\log \det (I + K'_h/\lambda) - \log \det (I + K_h/\lambda)] \quad (79)$$

B.4 Proof of Corollary 4.8

The proof is inspired from (Jin et al., 2021; Duan et al., 2020). Recall that we denote $\mathcal{D}^\theta = \{(s_h^\tau, a_h^\tau, \hat{r}_h^{\theta_h}(s_h^\tau, a_h^\tau))\}_{\tau, h=1}^{N, H}$, which is a combination of labeled dataset \mathcal{D}_1 and unlabeled dataset \mathcal{D}_2^θ . We partition dataset \mathcal{D}^θ into H disjoint and equally sized sub dataset $\{\tilde{\mathcal{D}}_h^\theta\}_{h=1}^H$, where $|\tilde{\mathcal{D}}_h^\theta| = N_h = N/H$. Let $\mathcal{I}_h = \{N_h \cdot (h-1) + 1, \dots, N_h \cdot h\} = \{\tau_{h,1}, \dots, \tau_{h,N_h}\}$ satisfy $\tilde{\mathcal{D}}_h^\theta = \{(s_h^\tau, a_h^\tau, \hat{r}_h^{\theta_h}(s_h^\tau, a_h^\tau))\}_{\tau \in \mathcal{I}_h}$. Define

$$Z_h = \sum_{\tau \in \mathcal{I}_h} A_h^\tau, \quad A_h^\tau = \phi(s_h^\tau, a_h^\tau) \phi(s_h^\tau, a_h^\tau)^\top - \Sigma_h, \quad (80)$$

where $\Sigma_h = \mathbb{E}_{\bar{\pi}} [\phi(s_h, a_h) \phi(s_h, a_h)^\top]$ for all $h \in [H]$. Clearly, $\mathbb{E}_{\bar{\pi}}[A_h^\tau] = 0$ from equation (80). Note that $\mathbb{E}_{\bar{\pi}}$ is taken with respect to the trajectory induced by the fixed behavior policy $\bar{\pi}$ in the underlying MDP, and the set $\{A_h^\tau\}_{\tau \in \mathcal{I}_h}$ is i.i.d. and centered for all $h \in [H]$.

As shown in Section 3.3, The feature mapping $\phi : \mathcal{Z} \rightarrow \mathcal{H}$ satisfies

$$\phi(z) = \sum_{j=1}^{\infty} \sigma_j \cdot \psi_j(z) \cdot \psi_j = \sum_{j=1}^{\infty} \sqrt{\sigma_j} \cdot \psi_j(z) \cdot (\sqrt{\sigma_j} \cdot \psi_j) \quad (81)$$

Let t be any positive integer and let $\Pi_t : \mathcal{H} \rightarrow \mathcal{H}$ denote the projection onto the subspace spanned by $\{\psi_j\}_{j \in [t]}$, i.e., $\Pi_t[\phi(z)] = \sum_{j=1}^t \sigma_j \cdot \psi_j(z) \cdot \psi_j$.

For d -finite Spectrum case, consider the case where $\sigma_j = 0$ for all $j > d$. Then, $\phi(z) = \Pi_d[\phi(z)]$ for any $z \in \mathcal{Z}$. That is, A_h^τ can be written as

$$A_h^\tau := W_h^\tau - W_h, \quad Z_h = \sum_{\tau \in \mathcal{I}_h} A_h^\tau = \sum_{\tau \in \mathcal{I}_h} W_h^\tau - W_h, \quad (82)$$

where $W_h^\tau = \phi(z_h^\tau) \phi(z_h^\tau)^\top$ as a $d \times d$ matrix, and $W_h = \mathbb{E}_{\bar{\pi}} [\phi(z_h) \phi(z_h)^\top]$. By the boundness of kernel (i.e. $\sup_{z \in \mathcal{Z}} k(z, z) \leq 1$), we have $\|\phi(z)\|_{\mathcal{H}_k} \leq 1, \forall z \in \mathcal{Z}$. By Jensen's inequality, we have

$$\|\Sigma_h\|_{\text{op}} \leq \mathbb{E}_{\bar{\pi}} \left[\left\| \phi(s_h, a_h) \phi(s_h, a_h)^\top \right\|_{\text{op}} \right] \leq 1$$

Similarly, for all $h \in [H]$ and all $\tau \in \mathcal{I}_h$, as it holds that

$$\|A_h^\tau\|_{\text{op}} \leq \|W_h^\tau\|_{\text{op}} + \|W_h\|_{\text{op}} \leq 2$$

we have

$$\begin{aligned} \|\mathbb{E}_{\bar{\pi}} [Z_h^\top Z_h]\|_{\text{op}} &= N_h \left\| \mathbb{E}_{\bar{\pi}} [(A_h^\tau)^\top A_h^\tau] \right\|_{\text{op}} \\ &\leq N_h \left\| \mathbb{E}_{\bar{\pi}} [A_h^\tau (A_h^\tau)^\top] \right\|_{\text{op}} \\ &\leq 4N_h \end{aligned}$$

Similarly, we have

$$\left\| (A_h^\tau)^\top A_h^\tau \right\|_{\text{op}} \leq \left\| (A_h^\tau)^\top \right\|_{\text{op}} \cdot \|A_h^\tau\|_{\text{op}} \leq 4 \text{ and } \|\mathbb{E}_{\bar{\pi}} [Z_h^\top Z_h]\|_{\text{op}} \leq 4N_h$$

Applying Lemma C.4 for Z_h defined in equation (80), for any fixed $h \in [H]$ and any $l \geq 0$, we have

$$\mathbb{P}(\|Z_h\|_{\text{op}} > l) = \mathbb{P}\left(\left\| \sum_{\tau \in \mathcal{I}_h} A_h^\tau \right\|_{\text{op}} > l\right) \leq 2d \cdot \exp\left(-\frac{l^2/2}{4N_h + 2l/3}\right)$$

For all $\delta \in (0, 1)$, we set $l = \sqrt{10N_h \log(4dH/\delta)}$, for sufficiently large $N_h \geq 5 \log(4dH/\delta)$, we obtain $\|Z_h\|_{\text{op}} \leq \sqrt{10N_h \log(4dH/\delta)}$ holds with probability at least $1 - \delta/2H$.

Moreover, Z_h defined in equation (80) can be written as

$$Z_h = \sum_{\tau \in \mathcal{I}_h} \phi(x_h^\tau, a_h^\tau) \phi(x_h^\tau, a_h^\tau)^\top - N_h \cdot \Sigma_h = (\Lambda_h - \lambda \cdot I) - N_h \cdot \Sigma_h \quad (83)$$

Recall that there exist positive constant c_{\min} such that $\inf_{\|f\|_{\mathcal{H}_k}=1} \langle f, \Sigma_h f \rangle_{\mathcal{H}_k} \geq c_{\min}$. For sufficiently large $N_h \geq \frac{4C^2}{c_{\min}^2} \log(4dH/\delta)$, we have

$$\begin{aligned} \inf_{\|f\|_{\mathcal{H}_k}=1} \langle f, \frac{\Lambda_h}{N_h} f \rangle_{\mathcal{H}_k} &\geq \inf_{\|f\|_{\mathcal{H}_k}=1} \langle f, \left(\frac{Z_h}{N_h} + \Sigma_h + \frac{\lambda}{N_h} I \right) f \rangle_{\mathcal{H}_k} \\ &\geq \frac{1}{N_h} \inf_{\|f\|_{\mathcal{H}_k}=1} \langle f, Z_h f \rangle_{\mathcal{H}_k} + \inf_{\|f\|_{\mathcal{H}_k}=1} \langle f, \Sigma_h f \rangle_{\mathcal{H}_k} \\ &\geq c_{\min} - \frac{1}{N_h} \|Z_h\|_{\text{op}} \geq c_{\min} - C \sqrt{\frac{\log(4HG(N_h, \lambda)/\delta)}{N_h}} \geq \frac{c_{\min}}{2} \end{aligned} \quad (84)$$

Hence, it holds that

$$\|\Lambda_h^{-1}\|_{\text{op}} \leq \frac{2}{N_h \cdot c_{\min}} \quad (85)$$

for all $h \in [H]$. This implies that

$$\|\Lambda_h^{-\frac{1}{2}} \phi(s, a)\|_{\mathcal{H}_k} \leq \|\phi(s, a)\|_{\mathcal{H}_k} \|\Lambda_h^{-1}\|_{\text{op}}^{1/2} \leq c' / \sqrt{N_h} \quad (86)$$

where $c' = \sqrt{2/c_{\min}}$ and using the fact that $\|\phi(s, a)\|_{\mathcal{H}_k} \leq 1$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$.

We define the event

$$\mathcal{E}_1^* = \left\{ \|\Lambda_h^{-\frac{1}{2}} \phi(s, a)\|_{\mathcal{H}_k} \leq c' / \sqrt{N_h} \text{ for all } (s, a) \in \mathcal{S} \times \mathcal{A} \text{ and all } h \in [H] \right\}$$

By equation (86), we have $\mathbb{P}(\mathcal{E}_1^*) \geq 1 - \delta/2$ for $N_h \geq \frac{4C^2}{c_{\min}^2} \log(4dH/\delta)$. Recall that for d -finite spectrum case, we have

$$\begin{aligned} \beta_h(\delta) &= \sqrt{1 + \frac{1}{N_1} \mathcal{S}} + \sqrt{C_1 \cdot d \cdot \log N_1 + \log(\frac{1}{\delta^2})} \\ B &= C_2 \cdot H \cdot \sqrt{d \log(N/\delta)} \end{aligned} \quad (87)$$

Use big tilde O notation, they can be written as

$$\begin{aligned} \beta_h(\delta) &= \tilde{\mathcal{O}}(\sqrt{d}) \\ B &= \tilde{\mathcal{O}}(H\sqrt{d}) \end{aligned} \quad (88)$$

Combining the result in Theorem 4.3 and equation (88) with $\delta = \delta/4$, we have

$$\begin{aligned} \text{SubOpt}(\hat{\pi}; s) &\leq 2 \sum_{h=1}^H \beta_h(\delta) \mathbb{E}_{\pi^*} \left[\|\phi(s_h, a_h)\|_{(\Lambda_h^{\mathcal{P}_1})^{-1}} \mid s_1 = s \right] \\ &\quad + 2B \sum_{h=1}^H \mathbb{E}_{\pi^*} \left[\|\phi(s_h, a_h)\|_{(\tilde{\Lambda}_h^{\mathcal{P}_1})^{-1}} \mid s_1 = s \right] \\ &\leq 2\beta_h(\delta) \cdot H \cdot c' / \sqrt{N_1} + 2B \cdot H \cdot c' / \sqrt{N_h} \\ &= \tilde{\mathcal{O}}(H\sqrt{\frac{d}{N_1}}) + \tilde{\mathcal{O}}(H^2\sqrt{\frac{d}{N_h}}) \end{aligned} \quad (89)$$

where the last equality follows from the fact that $N_h = N/H$ and $N = N_1 + N_2$ for all $h \in [H]$.

C Sufficient Lemma

Lemma C.1 (Concentration of Self-Normalized Processes in RKHS (Chowdhury & Gopalan, 2017)). *Let \mathcal{H} be an RKHS defined over $\mathcal{X} \subseteq \mathbb{R}^d$ with kernel function $K(\cdot, \cdot) : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$. Let $\{x_\tau\}_{\tau=1}^\infty \subset \mathcal{X}$ be a discrete time stochastic process that is adapted to the filtration $\{\mathcal{F}_t\}_{t=0}^\infty$. Let $\{\epsilon_\tau\}_{\tau=1}^\infty$ be a real-valued stochastic process such that (i) ϵ_τ is \mathcal{F}_τ -measurable and (ii) ϵ_τ is zero-mean and σ -sub-Gaussian conditioning on $\mathcal{F}_{\tau-1}$, i.e.,*

$$\mathbb{E}[\epsilon_\tau \mid \mathcal{F}_{\tau-1}] = 0, \quad \mathbb{E}[e^{\lambda \epsilon_\tau} \mid \mathcal{F}_{\tau-1}] \leq e^{\lambda^2 \sigma^2 / 2}, \quad \forall \lambda \in \mathbb{R}$$

Moreover, for any $t \geq 2$, let $E_t = (\epsilon_1, \dots, \epsilon_{t-1})^\top \in \mathbb{R}^{t-1}$ and $K_t \in \mathbb{R}^{(t-1) \times (t-1)}$ be the Gram matrix of $\{x_\tau\}_{\tau \in [t-1]}$. Then for any $\eta > 0$ and any $\delta \in (0, 1)$, with probability at least $1 - \delta$, it holds simultaneously for all $t \geq 1$ that

$$E_t^\top \left[(K_t + \eta \cdot I)^{-1} + I \right]^{-1} E_t \leq \sigma^2 \cdot \log \det [(1 + \eta) \cdot I + K_t] + 2\sigma^2 \cdot \log(1/\delta)$$

Proof. Please refer to Theorem 1 in (Chowdhury & Gopalan, 2017) □

Lemma C.2 (Lemma D.5 in (Yang et al., 2020)). *Let \mathcal{Z} be a compact subset of \mathbb{R}^d and $K : \mathcal{Z} \times \mathcal{Z} \rightarrow \mathbb{R}$ be the RKHS kernel of \mathcal{H} . We assume K is a bounded kernel so that $\sup_{z \in \mathcal{Z}} K(z, z) \leq 1$, and K is continuously differentiable on $\mathcal{Z} \times \mathcal{Z}$. Moreover, let T_K be the integral operator induced by K and the Lebesgue measure on \mathcal{Z} , such that $T_K f(z) = \int_{\mathcal{Z}} k(z, z') \cdot f(z') dz'$, $\forall f \in \mathcal{L}^2(\mathcal{Z})$. Let $\{\sigma_j\}_{j \geq 1}$ be the non-increasing sequence of eigenvalues of T_K . Recall the definition of maximal information gain in equation (2). We assume $\{\sigma_j\}_{j \geq 1}$ satisfies one of the following eigenvalue decay conditions:*

- γ -finite spectrum: $\sigma_j = 0$ for all $j > \gamma$, where γ is a positive integer.
- γ -exponential decay: there exists some constants $C_1, C_2 > 0$ such that $\sigma_j \leq C_1 \cdot \exp(-C_2 \cdot j^\gamma)$ for all $j \geq 1$, where $\gamma > 0$ is a positive constant.
- γ -polynomial decay: there exists some constants $C_1 > 0, \tau \in [0, 1/2)$ and $C_\psi > 0$ such that $\sigma_j \leq C_1 \cdot j^{-\gamma}$ and $\sup_{z \in \mathcal{Z}} \sigma_j^\tau \cdot |\psi_j(z)| \leq C_\psi$ for all $j \geq 1$, where $\gamma > 1$.

Suppose $\lambda \in [c_1, c_2]$ for absolute constants c_1, c_2 . Then we have

$$G(K, \lambda) \leq \begin{cases} C \cdot \gamma \cdot \log K & \gamma\text{-finite spectrum} \\ C \cdot (\log K)^{1+1/\gamma} & \gamma\text{-exponential decay} \\ C \cdot K^{(d+1)/(\gamma+d)} \cdot \log K & \gamma\text{-polynomial decay} \end{cases}$$

where C is an absolute constant that only depends on $d, \gamma, C_1, C_2, C, c_1$ and c_2 .

Proof. Please refer to Lemma D.5 in (Yang et al., 2020) for a detailed proof. □

Lemma C.3 (Extended Value Difference (Cai et al., 2020)). *Let $\pi = \{\pi_h\}_{h=1}^H$ and $\pi' = \{\pi'_h\}_{h=1}^H$ be any two policies and let $\{\widehat{Q}_h\}_{h=1}^H$ be any estimated Q -functions. For all $h \in [H]$, we define the estimated value function $\widehat{V}_h : \mathcal{S} \rightarrow \mathbb{R}$ by setting $\widehat{V}_h(x) = \left\langle \widehat{Q}_h(x, \cdot), \pi_h(\cdot \mid x) \right\rangle_{\mathcal{A}}$ for all $x \in \mathcal{S}$. For all $x \in \mathcal{S}$, we have*

$$\begin{aligned} \widehat{V}_1(x) - V_1^{\pi'}(x) &= \sum_{h=1}^H \mathbb{E}_{\pi'} \left[\left\langle \widehat{Q}_h(s_h, \cdot), \pi_h(\cdot \mid s_h) - \pi'_h(\cdot \mid s_h) \right\rangle_{\mathcal{A}} \mid s_1 = x \right] \\ &\quad + \sum_{h=1}^H \mathbb{E}_{\pi'} \left[\widehat{Q}_h(s_h, a_h) - (\mathbb{B}_h \widehat{V}_{h+1})(s_h, a_h) \mid s_1 = x \right] \end{aligned}$$

where $\mathbb{E}_{\pi'}$ is taken with respect to the trajectory generated by π' , while \mathbb{B}_h is the Bellman operator defined in Section 3.1.

Proof. Fix $h \in [H]$. Denote that $\iota_i = \widehat{Q}_i - \mathbb{B}_i \widehat{V}_{i+1}$. For all $i \in [h, H]$ and $s \in \mathcal{S}$, we have

$$\begin{aligned}
& \mathbb{E}_{\pi'} \left[\widehat{V}_i(s_i) - V_i^{\pi'}(s_i) \mid s_h = s \right] \\
&= \mathbb{E}_{\pi'} \left[\left\langle \widehat{Q}_i(s_i, \cdot), \pi_i(\cdot \mid s_i) \right\rangle - \left\langle \widehat{Q}_i^{\pi'}(s_i, \cdot), \pi_i'(\cdot \mid s_i) \right\rangle \mid s_h = s \right] \\
&= \mathbb{E}_{\pi'} \left[\left\langle \widehat{Q}_i(s_i, \cdot), \pi_i(\cdot \mid s_i) - \pi_i'(\cdot \mid s_i) \right\rangle + \left\langle \widehat{Q}_i(s_i, \cdot) - Q_i^{\pi'}(s_i, \cdot), \pi_i'(\cdot \mid s_i) \right\rangle \mid s_h = s \right] \\
&= \mathbb{E}_{\pi'} \left[\left\langle \widehat{Q}_i(s_i, \cdot), \pi_i(\cdot \mid s_i) - \pi_i'(\cdot \mid s_i) \right\rangle \mid s_h = s \right] \\
&+ \mathbb{E}_{\pi'} \left[\left\langle \iota_i(s_i, \cdot) + \mathbb{B}_i \widehat{V}_{i+1}(s_i, \cdot) - \left(r_i(s_i, \cdot) + \mathbb{P}_i V_{i+1}^{\pi'}(s_i, \cdot) \right), \pi_i'(\cdot \mid s_i) \right\rangle \mid s_h = s \right] \\
&= \mathbb{E}_{\pi'} \left[\left\langle \widehat{Q}_i(s_i, \cdot), \pi_i(\cdot \mid s_i) - \pi_i'(\cdot \mid s_i) \right\rangle \mid s_h = s \right] + \mathbb{E}_{\pi'} [\iota_i(s_i, a_i) \mid s_h = s] \\
&+ \mathbb{E}_{\pi'} \left[\mathbb{P}_i \left(\widehat{V}_{i+1} - V_{i+1}^{\pi'} \right)(s_i, a_i) \mid s_h = s \right] \\
&= \mathbb{E}_{\pi'} \left[\left\langle \widehat{Q}_i(s_i, \cdot), \pi_i(\cdot \mid s_i) - \pi_i'(\cdot \mid s_i) \right\rangle \mid s_h = s \right] + \mathbb{E}_{\pi'} [\iota_i(s_i, a_i) \mid s_h = s] \\
&+ \mathbb{E}_{\pi'} \left[\widehat{V}_{i+1}(s_{i+1}) - V_{i+1}^{\pi'}(s_{i+1}) \mid s_h = s \right]
\end{aligned} \tag{90}$$

where \mathbb{P}_i is the transition operator defined in Section 3.1. Rewrite equation (90), we have

$$\begin{aligned}
& \mathbb{E}_{\pi'} \left[\widehat{V}_i(s_i) - V_i^{\pi'}(s_i) \mid s_h = s \right] - \mathbb{E}_{\pi'} \left[\widehat{V}_{i+1}(s_{i+1}) - V_{i+1}^{\pi'}(s_{i+1}) \mid s_h = s \right] \\
&= \mathbb{E}_{\pi'} \left[\left\langle \widehat{Q}_i(s_i, \cdot), \pi_i(\cdot \mid s_i) - \pi_i'(\cdot \mid s_i) \right\rangle \mid s_h = s \right] + \mathbb{E}_{\pi'} [\iota_i(s_i, a_i) \mid s_h = s]
\end{aligned} \tag{91}$$

Taking $\sum_{i=h}^H$ on equation (91), then

$$\begin{aligned}
\widehat{V}_h(s) - V_h^{\pi'}(s) &= \sum_{i=h}^H \mathbb{E}_{\pi'} \left[\left\langle \widehat{Q}_i(s_i, \cdot), \pi_i(\cdot \mid s_i) - \pi_i'(\cdot \mid s_i) \right\rangle \mid s_h = s \right] \\
&+ \sum_{i=h}^H \mathbb{E}_{\pi'} \left[\widehat{Q}_i(s_i, a_i) - \mathbb{B}_i \widehat{V}_{i+1}(s_i, a_i) \mid s_h = s \right]
\end{aligned} \tag{92}$$

letting $h = 1$ completes the proof. \square

Lemma C.4 (Matrix Bernstein Inequality (Tropp et al., 2015)). *Suppose that $\{A_k\}_{k=1}^n$ are independent and centered random matrices in $\mathbb{R}^{d_1 \times d_2}$, that is, $\mathbb{E}[A_k] = 0$ for all $k \in [n]$. Also, suppose that such random matrices are uniformly upper bounded in the matrix operator norm, that is, $\|A_k\|_{op} \leq L$ for all $k \in [n]$. Let $Z = \sum_{k=1}^n A_k$ and*

$$v(Z) = \max \left\{ \left\| \mathbb{E}[ZZ^\top] \right\|_{op}, \left\| \mathbb{E}[Z^\top Z] \right\|_{op} \right\} = \max \left\{ \left\| \sum_{k=1}^n \mathbb{E}[A_k A_k^\top] \right\|_{op}, \left\| \sum_{k=1}^n \mathbb{E}[A_k^\top A_k] \right\|_{op} \right\}.$$

For all $t \geq 0$, we have

$$\mathbb{P}(\|Z\|_{op} \geq t) \leq (d_1 + d_2) \cdot \exp \left(-\frac{t^2/2}{v(Z) + L/3 \cdot t} \right)$$

Proof. See Tropp et al. (2015, Theorem 1.6.2) for a detailed proof. \square