# Sample Complexity of Kernel-Based Q-Learning 

Sing-Yuan Yeh ${ }^{1,2}$ Fu-Chieh Chang ${ }^{1,3}$ Chang-Wei Yueh ${ }^{4}$ Pei-Yuan Wu ${ }^{3,4}$ Alberto Bernacchia ${ }^{1}$ Sattar Vakili ${ }^{1}$<br>${ }^{1}$ MediaTek Research<br>${ }^{2}$ Graduate Program of Data Science, National Taiwan University and Academia Sinica<br>${ }^{3}$ Graduate Institute of Communication Engineering, National Taiwan University<br>${ }^{4}$ Department of Electrical Engineering, National Taiwan University


#### Abstract

Modern reinforcement learning (RL) often faces an enormous state-action space. Existing analytical results are typically for settings with a small number of state-actions, or simple models such as linearly modeled Q-functions. To derive statistically efficient RL policies handling large state-action spaces, with more general Qfunctions, some recent works have considered nonlinear function approximation using kernel ridge regression. In this work, we derive sample complexities for kernel based Q-learning when a generative model exists. We propose a nonparametric Q-learning algorithm which finds an $\epsilon$-optimal policy in an arbitrarily large scale discounted MDP. The sample complexity of the proposed algorithm is order optimal with respect to $\epsilon$ and the complexity of the kernel (in terms of its information gain). To the best of our knowledge, this is the first result showing a finite sample complexity under such a general model.


## 1 INTRODUCTION

In recent years, Reinforcement Learning (RL) has been successfully applied to several fields, including gaming (Silver et al., 2016; Lee et al., 2018; Vinyals et al., 2019), autonomous driving (Kahn et al., 2017), microchip design (Mirhoseini et al., 2021), robot control (Kalashnikov et al., 2018), and algorithm search (Fawzi et al. 2022). Realworld problems usually contain an enormous state-action space, possibly infinite. For example, the game of Go has $10^{170}$ states (Silver et al. 2016), and the number of actions

Proceedings of the $26^{\text {th }}$ International Conference on Artificial Intelligence and Statistics (AISTATS) 2023, Valencia, Spain. PMLR: Volume 206. Copyright 2023 by the author(s).
in the space of algorithms for matrix multiplication is $10^{30}$ (Fawzi et al. 2022). It is currently not fully understood how RL algorithms are able to learn successful policies to solve these problems. Modern function approximators, such as kernel-based learning and deep neural networks, seem to be required for this success.

An important theoretical question is as follows. Consider a Markov decision process (MDP) with an unknown transition probability distribution. Suppose that a generative model (Kakade 2003) is available, which provides sample transitions from any state-action pair. How many samples are required to learn a sufficiently good policy? That is referred to as the sample complexity.
Previous works have derived theoretical bounds on the sample complexity, under certain simple settings such as tabular and linear MDPs. In the tabular setting, it was shown that the sample complexity of learning an $\epsilon$-optimal policy (that is the value function is at most $\epsilon$ away from the optimal value function) is in $\mathcal{O}\left(\frac{|\mathcal{S}||\mathcal{A}|}{\epsilon^{2}}\right)$, where $|\mathcal{S}|$ and $|\mathcal{A}|$ are the cardinality of the state and action sets, respectively (Kearns and Singh, 1998; Azar et al., 2013; Sidford et al. 2018ab b; Agarwal et al., 2020), implying that for a very large state-action space, a virtually infinite number of samples is required to obtain a good performance. Another line of work considers a linear MDP model, where the transition probability admits a linear representation in a $d$-dimensional state-action feature map. It was shown that the sample complexity is in $\mathcal{O}\left(\frac{d}{\epsilon^{2}}\right)$ in this case, that is independent of the size of state and action spaces Yang and Wang (2019). Unfortunately, the linear assumption is rather inflexible and not often the case in practice.

In order to address the limitations of small state-actions or simple models, arising from tabular and linear MDP assumptions, a few recent studies considered nonlinear function approximation over possibly infinite state-action domains using kernel ridge regression. In these works, the transition probability distribution (and sometimes the reward function) are flexibly represented using a kernel-
based model (Yang and Wang, 2020, Yang et al. 2020abb). The kernel-based models provide powerful regressor and uncertainty estimates, which can be leveraged to guide the RL algorithm. Furthermore, kernel-based models have a great representation capacity and can model a wide range of problems, considering that all continuous functions on compact subsets of $\mathbb{R}^{d}$ can be approximated using common kernels (Srinivas et al., 2010). The existing works, however, do not address the specific question of sample complexity considered in this work, and instead derive regret bounds in the setting of an episodic MDP. A more detailed comparison is provided in Section 1.2 .

The kernel-based approaches may be understood as a linear model with an infinite-dimensional state-action feature map, that corresponds to, e.g., the Mercer eigenfeatures of the kernel (see Section 2.2. In this sense, the linear model is a special case of the kernel-based model with a linear kernel. Nonetheless, the results on the sample complexity of linear MDPs do not extend to the kernel-based models, as those sample complexities scale with the dimension of the feature map (that is possibly infinite in the kernel setting). In contrast, in the kernel setting, the sample complexity depends on certain kernel-specific properties determined by the complexity of the kernel, which will be discussed in more detail.

### 1.1 Contributions

Considering a discounted MDP and the question of sample complexity (similar to Azar et al., 2013; Sidford et al., 2018ab; Yang and Wang, 2019), we extend and generalize the existing work as follows.

- We introduce Kernel-based Q-Learning, referred to as KQLearn, a sample collection algorithm, which returns an $\epsilon$-optimal policy with a finite sample complexity over a very general class of models. In comparison to tabular and linear MDP settings, KQLearn makes at least two innovative contributions. In the tabular setting, the samples are collected from all stateaction pairs that leads to an $|\mathcal{S}||\mathcal{A}|$ scaling of the sample complexity. In the linear setting, the samples are collected from a set of state-actions spanning the entire state-action space (leading to the scaling of the sample complexity with dimension $d$ ). Then, an estimation of the parameters of the linear model are updated through value iteration. Neither of these approaches are feasible in our case with an infinite state-action space and a non-parametric kernel-based model. KQLearn instead takes advantage of uncertainties provided by the kernel model to create a finite state-action set which is used for collecting the samples. These samples are then passed through an approximate Bellman operator using kernel ridge regression to update the value function (that is a continuous
function over the entire state-action space).
- We derive a finite sample complexity for KQLearn under a wide range of kernel models. In particular, we consider two classes of kernels with exponentially ( $\sigma_{m} \sim \exp \left(-m^{\beta_{e}}\right), \beta_{e}>0$ ) and polynomially ( $\sigma_{m} \sim m^{-\beta_{p}}, \beta_{p}>1$ ) decaying Mercer eigenvalues $\sigma_{m}$ (see Definition 2). We prove a sample complexity of $\tilde{\mathcal{O}}\left(\frac{1}{\epsilon^{2}}\right)$ and $\tilde{\mathcal{O}}\left(\left(\frac{1}{\epsilon}\right)^{\frac{2 \beta_{p}}{\beta_{p}-1}}\right)[1$ under these two settings, respectively. To the best of our knowledge, this is the first finite sample complexity, for all $\beta_{p}>1$, and the first order optimal sample complexity in $\epsilon$, under the setting of polynomially decaying eigenvalues. Comparison with the related work is discussed in more detail in Section 1.2 As a special case, we recover the $\tilde{\mathcal{O}}\left(\frac{d}{\epsilon^{2}}\right)$ sample complexity of the linear setting reported in Yang and Wang (2019).

We acknowledge that our bounds on the sample complexity of KQLearn may not be order optimal in the dependence on the discount factor $\gamma$. In particular, our bounds grow with $\frac{1}{(1-\gamma)^{7}}$ in the case of smooth kernels, similar to the PPQLearning algorithm proposed in Yang and Wang (2019) for the linear setting. Under the tabular and linear settings, however, this dependency was improved to $\frac{1}{(1-\gamma)^{3}}$, in Sidford et al. (2018a) and Yang and Wang (2019), respectively. It appears a challenging problem whether the same improvement is feasible here. As mentioned above, even establishing a finite sample complexity is a challenging problem, and the sample complexities in the existing work may diverge with difficult kernels (some polynomial kernels as discussed in Section 1.2.

### 1.2 Related Work

The specific problem of sample complexity in a discounted MDP using a generative model has been considered in tabular and linear settings. The results are summarized in Table 1 Other variants of the problem, consider MDPs in the absence of a generative model (e.g., see, Azar et al., 2017, Jin et al., 2018, 2020, Russo, 2019, Yang et al. 2020a b; Kakade et al., 2020; Zhou et al. 2021; Domingues et al. 2021, as representative works, as well as references therein), often episodic, with $T$ episodes of length $H$, and regret bounds depending on $T$ and $H$. The regret bounds can then be translated into sample complexities (e.g., see, Jin et al., 2018, Yang et al. 2020b). These results are also reported in Table 1. Other approaches to nonlinear function approximation in RL include models with bounded eluder dimension (Wang et al., 2020; Ayoub et al., 2020) and smoothing kernels (Domingues et al., 2021). Among these works the two most relevant ones to ours are Yang and Wang (2019) and Yang et al. (2020ab).

[^0]Table 1: The existing sample complexities under various settings, discussed in Section 1.2

| ALGORITHM | MDP | SETTING | SAMPLE COMPLEXITY |
| :---: | :---: | :---: | :---: |
| (Q-learning with UCB, Jin et al., 2018) | Episodic | Tabular | $\tilde{\mathcal{O}}\left(\frac{\|S\|\|\mathcal{A}\| H^{4}}{\epsilon^{2}}\right)$ |
| (LSVI-UCB, Jin et al. 2020) | Episodic | Linear | $\tilde{\mathcal{O}}\left(\frac{d^{3} H^{4}}{\epsilon^{2}}\right)$ |
| (KOVI, Yang et al. 2020b) | Episodic | Kernel-based, polynomial eigendecay | $\tilde{\mathcal{O}}\left(H^{4}\left(\frac{1}{\epsilon}\right)^{\frac{2 \beta_{p}}{\beta_{p}-2}}\right)$ |
|  |  | Kernel-based, exponential eigendecay | $\tilde{\mathcal{O}}\left(\frac{H^{4}}{\epsilon^{2}}\right)$ |
| Sidford et al. QVI) 2018a, Variance-Reduced | Discounted | Tabular | $\tilde{\mathcal{O}}\left(\frac{\|\mathcal{S}\|\|\mathcal{A}\|}{(1-\gamma){ }^{\text {c }}{ }^{2}}\right)$ |
| (PPQ-Learning, Yang and Wang, 2019) | Discounted | Linear | $\tilde{\mathcal{O}}\left(\frac{d}{(1-\gamma)^{7} \epsilon^{2}}\right)$ |
| (OPPQ-Learning, Yang and Wang, 2019) | Discounted | Linear | $\tilde{\mathcal{O}}\left(\frac{d}{(1-\gamma)^{3} \epsilon^{2}}\right)$ |
| KQLearn | Discounted | Kernel-based, polynomial eigendecay | $\tilde{\mathcal{O}}\left(\frac{1}{\epsilon^{\frac{2 \beta_{p}}{\beta_{p-1}}}(1-\gamma)^{\frac{7 \beta_{p}-1}{\beta_{p}-1}}}\right)$ |
|  |  | Kernel-based, exponential eigendecay | $\tilde{\mathcal{O}}\left(\frac{1}{\epsilon^{2}(1-\gamma)^{7}}\right)$ |

Similar to Yang and Wang (2019), we also consider sample complexity in a discounted MDP using a generative model. We consider a non-parametric kernel-based model, while they considered a parametric linear model. Thus, neither their algorithm nor their results extend to our setting. The linear setting is a special case of the kernel setting with a linear kernel, in which, we recover the $\tilde{\mathcal{O}}\left(\frac{d}{\epsilon^{2}}\right)$ sample complexity, given in Yang and Wang (2019), for two algorithms: PPQ-Learning and OPPQ-Learning. The latter improved the sample complexity with respect to the discount factor.
Similar to Yang et al. (2020b), we also consider a kernelbased model. We consider sample complexity in a discounted MDP with a generative model, while they primarily considered regret bounds in an episodic MDP. They also reported sample complexities as a direct consequence of their regret bounds. Specifically, under the two settings of exponentially and polynomially decaying eigenvalues, their sample complexities translate to $\tilde{\mathcal{O}}\left(\frac{1}{\epsilon^{2}}\right)$ and $\tilde{\mathcal{O}}\left(\left(\frac{1}{\epsilon}\right)^{\frac{2 \beta_{p}}{\beta_{p}-2}}\right)$, respectively. Under the polynomial setting, their sample complexity bound is larger than ours. In addition, their sample complexity is not always finite and may diverge when $1<\beta_{p} \leq 2$, that includes many cases of interest. This suboptimality and possibly trivial result is a consequence of the superlinear (thus, trivial) regret bounds when $1<\beta_{p} \leq 2$. See Vakili et al. (2021d), for a detailed discussion on the theoretical challenges related to this result.

For example, consider the Matérn family of kernels as one of the most commonly used (Snoek et al., 2012; Shahriari
et al. (2015) and theoretically interesting (Srinivas et al. 2010) family of kernels. For a Matérn kernel with smoothness parameter $\nu$ on a $d$ dimensional input domain, $\beta_{p}=$ $1+\frac{2 \nu}{d}$ (Yang et al. 2020b). That implies the sample complexity in Yang et al. (2020b) diverges when $d>2 \nu$ (that is often the case when using the Matérn kernel). We, however, emphasize that the discounted MDP with a generative model and the episodic MDPs are different settings, and cannot be compared directly. Nonetheless, we present the first always finite sample complexity under a very general setting covering all kernels with polynomially decaying eigenvalues.
Another related problem is the kernel-based bandit problem (Srinivas et al. 2010), which corresponds to a degenerate MDP with $|\mathcal{S}|=1$. The kernel-based bandit problem is a well studied problem with order optimal regret bounds (Salgia et al., 2021, Li and Scarlett. 2022) and sample complexities (Vakili et al. 2021a). The lower bounds on sample complexities for the squared exponential (SE) and Matérn kernels are reported in (Scarlett et al. 2017), which have the same scaling with $\epsilon$ (up to logarithmic factors) as in our results, showing the order optimality of our sample complexities with $\epsilon$ (see Section (4). To obtain the optimal order of sample complexities, Vakili et al. (2021a) proposed the Maximum Variance Reduction (MVR) algorithm, which selects observation points with the highest uncertainty in a greedy manner. We build upon their approach by extending it to the sample complexities of Q-learning. This extension is significantly more involved due to the presence of Markovian dynamics and the Bellman operator.

There exist other studies in the field of kernel-based RL that are related to our work, including Ormoneit and Sen (2002); Domingues et al. (2021). However, there are notable differences between their work and ours in terms of both the settings and the algorithms. Specifically, these works assumed a Lipschitz reward function and a Lipschitz transition probability distribution, which allowed them to use kernel smoothing methods. In contrast, our work uses kernel ridge regression with the assumption of a bounded RKHS norm of the transition probability distribution.

Paper structure: In Section 2, the problem is formalized, after an overview of the background on MDPs and kernel ridge regression. In Section 3, KQLearn is presented. The results are discussed in Section4. A high level analysis is provided in Section 5, while the details are deferred to the appendix.

## 2 PRELIMINARIES

In this section, we overview the background on MDPs and kernel ridge regression. We then formally state the problem of sample complexity for Q-learning under this setting.

### 2.1 Discounted Markov Decision Process

A discounted Markov Decision Process (MDP) can be described by the tuple $M=(\mathcal{S}, \mathcal{A}, P, r, \gamma)$, where $\mathcal{S}$ is the state space, $\mathcal{A}$ is the action space, $\gamma \in(0,1)$ is the discount factor, $r: \mathcal{S} \times \mathcal{A} \rightarrow[0,1]$ is the reward function and $P(\cdot \mid s, a)$ is the transition probability distribution ${ }^{2}$ on $\mathcal{S}$ for the next state from state-action pair $(s, a)$. We use the notation $\mathcal{Z}=\mathcal{S} \times \mathcal{A}$ to denote the state-action space. Our results generally hold true for (possibly very large and) finite $\mathcal{Z}$ or certain infinite $\mathcal{Z}$. For correctness, we assume that $\mathcal{Z}$ is a compact subset of $\mathbb{R}^{d}$.

The goal is to find a (possibly random) policy $\pi: \mathcal{S} \rightarrow$ $\mathcal{A}$, that maximizes the long-term expected reward, i.e., the value function,

$$
V^{\pi}(s):=\mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^{t} r\left(s_{t}, \pi\left(s_{t}\right)\right) \mid s_{0}=s\right]
$$

where $s_{t} \sim P\left(\cdot \mid s_{t-1}, \pi\left(s_{t-1}\right)\right)$ forms the trajectory of the states. It can be shown that (e.g., see Puterman, 2014), under mild assumptions (e.g., continuity of $P$, compactness of $\mathcal{Z}$, and boundedness of $r$ ) there exists an optimal policy $\pi^{*}$ which attains the maximal possible value $V^{*}$ at every state,

$$
\forall s \in \mathcal{S}: \quad V^{*}(s):=V^{\pi^{*}}(s)=\max _{\pi} V^{\pi}(s)
$$

[^1]To simplify the notation, for a value function $V: \mathcal{S} \rightarrow \mathbb{R}$, let

$$
[P V](s, a):=\mathbb{E}_{s^{\prime} \sim P(\cdot \mid s, a)}\left[V\left(s^{\prime}\right)\right]
$$

The Q-function, also sometimes referred to as the stateaction value function, of a policy $\pi$, and the optimal Q function are defined as

$$
\begin{aligned}
& Q^{\pi}(s, a)=r(s, a)+\gamma\left[P V^{\pi}\right](s, a), \text { and } \\
& Q^{*}(s, a)=Q^{\pi^{*}}(s, a)
\end{aligned}
$$

respectively. The Bellman operator $\mathcal{T}: \mathbb{R}^{\mathcal{S}} \rightarrow \mathbb{R}^{\mathcal{S}}$ is defined as

$$
\forall s \in \mathcal{S}: \quad[\mathcal{T} V](s)=\max _{a \in \mathcal{A}}\{r(s, a)+\gamma[P V](s, a)\}
$$

Sample complexity of $\epsilon$-optimal policies: An $\epsilon$-optimal policy is defined as follows.

Definition 1. ( $\epsilon$-optimal policy) A policy $\pi$ is called $\epsilon$ optimal if it achieves near optimal values from any initial state as follows:

$$
V^{\pi}(s) \geq V^{*}(s)-\epsilon, \quad \forall s \in \mathcal{S}
$$

or equivalently $\left\|V^{\pi}-V^{*}\right\|_{\infty} \leq \epsilon$.
We aim to learn $\epsilon$-optimal policies using a small number of samples. In this work, following Kearns and Singh (1998); Azar et al. (2013); Sidford et al. (2018a b); Yang and Wang (2019), we suppose that a generative model (Kakade, 2003) is given where the RL algorithm is able to query transition samples $s^{\prime} \sim P(\cdot \mid s, a)$ for any state-action pair $(s, a) \in \mathcal{Z}$. The sample complexity of an RL algorithm is defined as the number of such samples used by the algorithm to obtain an $\epsilon$-optimal policy.

### 2.2 RKHS and Kernel Ridge Regression

The existing work achieving finite sample complexity in the RL setting typically assumes a small state-action space or linearly modeled MDPs. These results can be generalized and extended using kernel-based learning. In particular, a natural approach is to use elements of a known reproducing kernel Hilbert space (RKHS) to model the transitions. In this section, we overview RKHSs and kernel ridge regression.
Let $K: \mathcal{Z} \times \mathcal{Z} \rightarrow \mathbb{R}$ be a known positive definite kernel with respect to a finite Borel measure. Let $\mathcal{H}_{K}$ be the RKHS induced by $K$, where $\mathcal{H}_{K}$ contains a family of functions defined on $\mathcal{Z}$. 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 of $\mathcal{H}_{K}$, respectively. The reproducing property implies that for all $f \in \mathcal{H}_{K}$, and $z \in \mathcal{Z},\langle f, K(\cdot, z)\rangle_{\mathcal{H}_{K}}=f(z)$. Without loss of generality, we assume $K(z, z) \leq 1$ for
all $z$. Mercer theorem implies, under certain mild conditions, $K$ can be represented using an infinite dimensional feature map:

$$
\begin{equation*}
K\left(z, z^{\prime}\right)=\sum_{m=1}^{\infty} \sigma_{m} \psi_{m}(z) \psi_{m}\left(z^{\prime}\right) \tag{1}
\end{equation*}
$$

A formal statement and the details are provided in Appendix A

Kernel ridge regression: Kernel-based models provide powerful regressor and uncertainty estimators (roughly speaking, surrogate posterior variances) which can be leveraged to guide the RL algorithm. In particular, consider an unknown function $f \in \mathcal{H}_{K}$. Consider a set $\mathcal{U}_{J}=\left\{z_{j}\right\}_{j=1}^{J} \subset \mathcal{Z}$ of $J$ inputs. Assume $J$ noisy observations $\left\{Y\left(z_{j}\right)=f\left(z_{j}\right)+\epsilon_{j}\right\}_{j=1}^{J}$ are provided, where $\epsilon_{j}$ are i.i.d. zero mean sub-Gaussian noise terms. Kernel ridge regression provides the following regressor and uncertainty estimate, respectively (see, e.g., Schölkopf et al. 2002),

$$
\begin{align*}
\hat{f}_{\mathcal{U}_{J}}(z) & =k_{\mathcal{U}_{J}}^{\top}(z)\left(K_{\mathcal{U}_{J}}+\lambda^{2} I_{J}\right)^{-1} Y_{\mathcal{U}_{J}} \\
\Sigma_{\mathcal{U}_{J}}^{2}(z) & =K(z, z)-k_{\mathcal{U}_{J}}^{\top}(z)\left(K_{\mathcal{U}_{J}}+\lambda^{2} I_{J}\right)^{-1} k_{\mathcal{U}_{J}}(z), \tag{2}
\end{align*}
$$

where $k_{\mathcal{U}_{J}}(z)=\left[K\left(z, z_{1}\right), \ldots, K\left(z, z_{J}\right)\right]^{\top}$ is a $J \times 1$ vector of the kernel values between $z$ and observations, $K_{\mathcal{U}_{J}}=\left[K\left(z_{i}, z_{j}\right)\right]_{i, j=1}^{J}$ is the $J \times J$ kernel matrix, $Y_{\mathcal{U}_{J}}=$ $\left[Y\left(z_{1}\right), \ldots, Y\left(z_{J}\right)\right]^{\top}$ is the $J \times 1$ observation vector, $I_{J}$ is the identity matrix of dimensions $J$, and $\lambda>0$ is a free regularization parameter.

Confidence intervals: The prediction and uncertainties provided by kernel ridge regression allow us to use standard confidence intervals in the algorithm and analysis. In particular, various results exist stating that with probability at least $1-\delta$, the prediction function satisfies $\left|f(z)-\hat{f}_{\mathcal{U}_{J}}(z)\right| \leq \beta(\delta) \Sigma_{\mathcal{U}_{J}}(z)$ (either for fixed $z$, or simultaneously for all $z$ ) where the confidence interval width multiplier $\beta(\delta)$ depends on the properties of the observation noise and the complexity of $f$ in terms of its RKHS norm (Srinivas et al., 2010, Abbasi-Yadkori, 2013, Vakili et al. 2021a, 2022). If the domain $\mathcal{Z}$ is finite, the uniform confidence bounds readily follow from a union bound over the confidence intervals for a fixed $z$. For continuous domains, a discretization argument is typically used considering the following continuity assumption.
Assumption 1. For each $n \in \mathbb{N}$, there exists a discretization $\mathbb{Z}$ of $\mathcal{Z}$ such that, for any $f \in \mathcal{H}_{K}$ with $\|f\|_{\mathcal{H}_{K}} \leq C_{K}$, we have $f(z)-f([z]) \leq \frac{1}{n}$, where $[z]=\arg \min _{z^{\prime} \in \mathbb{Z}} \| z^{\prime}-$ $z \|_{l^{2}}$ is the closest point in $\mathbb{Z}$ to $z$, and $|\mathbb{Z}| \leq c C_{K}^{d} n^{d}$, where $c$ is a constant independent of $n$ and $C_{K}$.

Assumption 1 is a technical and mild assumption that holds for typical kernels such as SE and Matérn with $\nu>$

1 (Srinivas et al. 2010; Chowdhury and Gopalan, 2017, Vakili et al., 2021a).

In our analysis, we use the following confidence interval for the RKHS elements.
Lemma 1 (Vakili et al. (2021a, 2022)). Consider a fixed design of observation points where $\mathcal{U}_{J}$ is independent of the observation noise. When the noise terms are sub-Gaussian with parameter $R{ }^{3}$ and $\|f\|_{\mathcal{H}_{K}} \leq C_{K}$, the following each hold uniformly in $z \in \mathcal{Z}$, with probability $1-\delta$,

$$
\begin{align*}
f(z) & \leq \hat{f}_{\mathcal{U}_{J}}(z)+\beta(\delta) \Sigma_{\mathcal{U}_{J}}(z) \\
f(z) & \geq \hat{f}_{\mathcal{U}_{J}}(z)-\beta(\delta) \Sigma_{\mathcal{U}_{J}}(z) \tag{3}
\end{align*}
$$

where $\beta(\delta)=\mathcal{O}\left(C_{K}+\frac{R}{\lambda} \sqrt{d \log \left(\frac{J C_{K}}{\delta}\right)}\right)$.
Maximal information gain: It is useful for our analysis to define maximal information gain $\Gamma_{K, \lambda}$, that is a kernel specific complexity term. It allows us to bound the total uncertainty in the kernel model using results similar to elliptical potential lemma (Carpentier et al. 2020). In particular, let us define

$$
\begin{equation*}
\Gamma_{K, \lambda}(J)=\sup _{\mathcal{U} \subset \mathcal{Z},|\mathcal{U}| \leq J} \frac{1}{2} \log \operatorname{det}\left(I_{J}+\frac{1}{\lambda^{2}} K_{\mathcal{U}}\right) \tag{4}
\end{equation*}
$$

Then, we have the following.
Lemma 2 (Srinivas et al. (2010)). For any set $\mathcal{U}_{J} \subset \mathcal{Z}$, we have

$$
\begin{equation*}
\sum_{j=1}^{J} \Sigma_{\mathcal{U}_{j-1}}^{2}\left(z_{j}\right) \leq \frac{2}{\log \left(1+1 / \lambda^{2}\right)} \Gamma_{K, \lambda}(J) \tag{5}
\end{equation*}
$$

### 2.3 Problem Formulation

Consider the discounted MDP described in Section 2.1. We are interested in designing an algorithm with a small sample complexity which obtains an $\epsilon$-optimal RL policy, under the assumption that the transition probability distribution lives in the RKHS of a known kernel. Without loss of generality, we assume its RKHS norm is bounded by 1.
Assumption 2. Assume that the transition probability distribution satisfies,

$$
\forall s^{\prime} \in \mathcal{S}: \quad\left\|P\left(s^{\prime} \mid \cdot, \cdot\right)\right\|_{\mathcal{H}_{K}} \leq 1
$$

This assumption is very flexible given the generality of the RKHSs. This is a standard assumption which is also used in Yang et al. (2020b). We do not make any explicit assumptions on the Q-function related to the policy. Recall the definition of $P V$ from Section 2.1. In Lemma 3, we prove that for any $V: \mathcal{S} \rightarrow\left[0, \frac{1}{1-\gamma}\right],\|P V\|_{\mathcal{H}_{K}} \leq \frac{c}{1-\gamma}$, as a consequence of Assumption 2, that is essential for our analysis.

[^2]Some generic notation: For any real number $x$, and real numbers $a, b$, the notation $\Pi_{[a, b]}[x]$ is used to denote the projection of $x$ onto $[a, b]$. For any integer $J, I_{J}$ denotes the $J \times J$ identity matrix, and $\mathbf{0}_{J}$ denotes the $J \times 1$ zero vector.

## 3 KERNEL BASED Q-LEARNING

In this section, we present a novel kernel based Q-learning algorithm, referred to as KQLearn. Recall $Q(s, a)=$ $r(s, a)+\gamma[P V](s, a)$. The transition probability distribution $P$ and the value function $V$ are both unknown to the algorithm. The algorithm, thus, recursively approximates $P V$, in rounds, using kernel ridge regression of $P V$ from the observations in the previous round. I.e., the algorithm performs updates based on an approximate Bellman operator using predictions for $P V$ provided by the kernel model. The samples are collected based on uncertainties for $P V$ in the kernel model. For this purpose, the algorithm first creates a maximum uncertainty set which is used to collect the samples.

Maximum Uncertainty Set $\left(\mathcal{U}_{J}\right)$ : The algorithm starts with creating a maximum uncertainty set with size $J$ referred to as $\mathcal{U}_{J} \subset \mathcal{Z}$. This set is created based on the uncertainties provided by the kernel model. In particular, each state-action is added to this set based on the following rule: choose the state-action with the highest uncertainty in the kernel model.

$$
\begin{equation*}
\left(s_{j}, a_{j}\right)=\underset{(s, a) \in \mathcal{Z}}{\arg \max } \Sigma_{\mathcal{U}_{j-1}}^{2}(s, a) \tag{6}
\end{equation*}
$$

Then, recursively, $\mathcal{U}_{j}=\mathcal{U}_{j-1} \cup\left\{\left(s_{j}, a_{j}\right)\right\}$, starting from $\mathcal{U}_{0}=\varnothing$. The set $\mathcal{U}_{J}$ is then used to collect samples from the generative model.

The algorithm proceeds in rounds indexed by $\ell=1, \ldots, L$. Each round $\ell$ receives noisy observations $Y_{\mathcal{U}_{J}}^{(\ell-1)}$ of $P V$ from the previous round, $\ell-1$. These observations are then used within kernel ridge regression to form a regressor of $P V$ over entire $\mathcal{Z}$, and obtain new observation $Y^{(\ell)}$. The observation vector is initialized to a zero vector $Y^{(0)}=\mathbf{0}_{J}$.

During each round $\ell$, for each state-action pair $\left(s_{j}, a_{j}\right) \in$ $\mathcal{U}_{J}$, a transition state $s_{j}^{\prime} \sim P\left(\cdot \mid s_{j}, a_{j}\right)$ is acquired from the generative model. The observation $Y^{(\ell)}\left(s_{j}, a_{j}\right)$ is then given as follows:

$$
\begin{align*}
& Y^{(\ell)}\left(s_{j}, a_{j}\right)=\Pi_{\left[0, \frac{1}{1-\gamma}\right]} \max _{a \in \mathcal{A}}\left\{r\left(s_{j}^{\prime}, a\right)\right. \\
& \left.\quad+\gamma k_{\mathcal{U}_{J}}\left(s_{j}^{\prime}, a\right)^{\top}\left(K_{\mathcal{U}_{J}}+\lambda^{2} I_{J}\right)^{-1} Y_{\mathcal{U}_{J}}^{(\ell-1)}\right\} . \tag{7}
\end{align*}
$$

The second term on the right hand side is the regressor in kernel ridge regression on $P V$, using $Y_{\mathcal{U}_{J}}^{(\ell-1)}$ as a vector
of observations. In the analysis, we show a high probability bound on the error of this regression. The vector $Y_{\mathcal{U}_{J}}^{(\ell)}=\left[Y^{(\ell)}\left(s_{1}, a_{1}\right), \ldots, Y^{(\ell)}\left(s_{J}, a_{J}\right)\right]^{\top}$ can be understood as updated noisy observations of $P V$ which is passed to the next round, $\ell+1$. By definition of the value function and the assumption of bounded rewards, it can be easily checked that $0 \leq V^{*}(s) \leq \frac{1}{1-\gamma}$, for all $s \in \mathcal{S}$. We thus project the value of $P V$ on $\left[0, \frac{1}{1-\gamma}\right]$ interval.

KQLearn collects $N=J L$ samples in total $\left[^{4}\right.$ A pseudocode is provided in Algorithm 1 .

After collecting all samples, the KQLearn algorithm returns an RL policy $\pi$ which selects the actions based on the following proxy $Q$-function:

$$
\begin{equation*}
\widehat{Q}^{(L)}(s, a)=r(s, a)+\gamma k_{\mathcal{U}_{J}}^{\top}(s, a)\left(K_{\mathcal{U}_{J}}+\lambda^{2} I_{J}\right)^{-1} Y_{\mathcal{U}_{J}}^{(L)} \tag{8}
\end{equation*}
$$

Specifically, when state $s$ is observed, the policy $\pi$ selects the action $\pi(s)=\arg \max _{a \in \mathcal{A}} \widehat{Q}^{(L)}(s, a)$. The second term on the right hand side is the kernel ridge regression of $P V$ using the observation in round $L$ of the KQlearn algorithm.

```
Algorithm 1 Kernel-based Q-learning (KQLearn)
    Input Discounted MDP with a generative model,
    kernel \(K\), regularization parameter \(\lambda>0\), and \(N>0\)
    Output \(\widehat{Q}^{(L)}: \mathcal{Z} \rightarrow \mathbb{R}\)
    Initialize \(L, J \in \mathbb{N}, N=L J\).
    Initialize \(Y^{(0)}=\mathbf{0}_{J}\) and the \(\operatorname{set} \mathcal{U}_{0}=\varnothing\).
    for all \(j=1, \ldots, J\) do
        Update the function \(\Sigma_{\mathcal{U}_{j-1}}(\cdot)\) using Equation 2
        Pick \(\left(s_{j}, a_{j}\right) \leftarrow \arg \max _{(s, a) \in \mathcal{Z}} \Sigma_{\mathcal{U}_{j-1}}^{2}(s, a)\).
        \(\mathcal{U}_{j} \leftarrow \mathcal{U}_{j-1} \cup\left\{\left(s_{j}, a_{j}\right)\right\}\).
    end for
    for all \(\ell=1, \ldots, L\) do \(\quad \triangleright\) round
        for all \(j=1, \ldots, J\) do
            Obtain a sample transition state
                \(s^{\prime} \sim P\left(\cdot \mid s_{j}, a_{j}\right)\)
            Update the \(Y^{(\ell)}\) as follows.
                \(Y^{(\ell)}\left(s_{j}, a_{j}\right) \leftarrow \Pi_{\left[0, \frac{1}{1-\gamma}\right]} \max _{a \in \mathcal{A}}\left\{r\left(s^{\prime}, a\right)\right.\)
                \(\left.+\gamma k_{\mathcal{U}_{J}}^{\top}\left(s^{\prime}, a\right)\left(K_{\mathcal{U}_{J}}+\lambda I_{J}\right)^{-1} Y_{\mathcal{U}_{J}}^{(\ell-1)}\right\}\)
        end for
    end for
    \(\widehat{Q}^{(L)}(\cdot)=r(\cdot)+\gamma k_{\mathcal{U}_{J}}^{\top}(\cdot)\left(K_{\mathcal{U}_{J}}+\lambda^{2} I_{J}\right)^{-1} Y_{\mathcal{U}_{J}}^{(L)}\).
```

[^3]
## 4 SAMPLE COMPLEXITY OF KQLEARN

In this section, we present our theoretical results. The following theorem establishes a bound on the error in the value function for the policy obtained in the KQLearn algorithm.
Theorem 1. Consider the discounted MDP described in Section 2.1. Consider the KQLearn algorithm described in Section 3. Under Assumptions 1 and 2 with probability at least $1-\delta$,

$$
\begin{aligned}
\left\|V^{\pi}-V^{*}\right\|_{\infty} \leq & 2 \beta(\delta)\left(\frac{\gamma}{1-\gamma}\right)^{2} \sqrt{\frac{2 \Gamma_{K, \lambda}(J)}{J}} \\
& +2 \gamma^{L-1}\left(\frac{1}{1-\gamma}\right)^{2}
\end{aligned}
$$

where $\beta(\delta)=\mathcal{O}\left(\frac{c}{1-\gamma}+\frac{1}{2 \lambda(1-\gamma)} \sqrt{d \log \left(\frac{J}{(1-\gamma) \delta}\right)}\right)$ and $c$ is a constant given in Lemma 3

When $J$ and $L$ are sufficiently large, both terms in the upper bound become arbitrarily small.

### 4.1 Sample Complexity

We can obtain explicit sample complexities for the KQLearn algorithm, using kernel specific bounds on $\Gamma_{K, \lambda}$, which depend on the decay rate of the Mercer eigenvalues of $K$. In particular, we define the following characteristic eigendecay profiles (which are similar to those outlined in Chatterji et al. (2019); Vakili et al. (2021c); Yang et al. (2020b)).
Definition 2. [Polynomial and Exponential Eigendecay] Consider the Mercer eigenvalues $\left\{\sigma_{m}\right\}_{m=1}^{\infty}$ of $K$ as given in Equation 1 in a decreasing order.
(i) For some $C_{p}>0, \beta_{p}>1, K$ is said to have a $\left(C_{p}, \beta_{p}\right)$ polynomial eigendecay, if for all $m \in \mathbb{N}$, we have $\sigma_{m} \leq C_{p} m^{-\beta_{p}}$.
(ii) For some $C_{e, 1}, C_{e, 2}, \beta_{e}>0, K$ is said to have a $\left(C_{e, 1}, C_{e, 2}, \beta_{e}\right)$ exponential eigendecay, if for all $m \in \mathbb{N}$, we have $\sigma_{m} \leq C_{e, 1} \exp \left(-C_{e, 2} m^{\beta_{e}}\right)$.

We are now ready to present explicit bounds on the sample complexity for the very general classes of kernels with polynomial and exponential decay of Mercer eigenvalues.

Theorem 2. Consider the discounted MDP described in Section 2.1. Consider the KQLearn algorithm described in Section 3. with $L=\Theta\left(\frac{\log \left(\epsilon(1-\gamma)^{2}\right)}{1-\gamma}\right)$ and $J=\frac{N}{L}$. Under Assumptions 1 and 2 KQLearn obtains an $\epsilon$-optimal policy with probability at least $1-\delta$, with a sample complexity at most

- In the case of a kernel with $\left(C_{p}, \beta_{p}\right)$ polynomial eigendecay,

$$
\begin{equation*}
N=\tilde{\mathcal{O}}\left(\frac{\left(\log \left(\frac{1}{\delta}\right)\right)^{\frac{\beta_{p}}{\beta_{p}-1}}}{\epsilon^{\frac{2 \beta_{p}}{\beta_{p}-1}}(1-\gamma)^{\frac{7 \beta_{p}-1}{\beta_{p}-1}}}\right) . \tag{9}
\end{equation*}
$$

- In the case of a kernel with $\left(C_{e, 1}, C_{e, 2}, \beta_{e}\right)$ exponential eigendecay,

$$
\begin{equation*}
N=\tilde{\mathcal{O}}\left(\frac{\log \left(\frac{1}{\delta}\right)}{\epsilon^{2}(1-\gamma)^{7}}\right) \tag{10}
\end{equation*}
$$

The detailed expressions for $J$ and $L$ including the implied constants and logarithmic factors in the $\tilde{\mathcal{O}}$ notation are provided in Appendix B. Equations 29 and 32

Specific Kernels: Our bounds on the sample complexity can be specialized for various kernels where the eigendecay or bounds on $\Gamma_{K, \lambda}$ is known (such as the ones in Srinivas et al. 2010, Vakili et al., 2021c a). Specifically, for the Matérn and SE kernels, we have, respectively,
$N=\tilde{\mathcal{O}}\left(\frac{\left(\log \left(\frac{1}{\delta}\right)\right)^{1+\frac{d}{2 \nu}}}{\epsilon^{2+\frac{d}{\nu}}(1-\gamma)^{7+\frac{3 d}{\nu}}}\right)$, and $N=\tilde{\mathcal{O}}\left(\frac{\log \left(\frac{1}{\delta}\right)}{\epsilon^{2}(1-\gamma)^{7}}\right)$.

### 4.2 Optimality of the Sample Complexities

The sample complexities given above are order optimal with respect to $\epsilon$. We compare them to the lower bounds on the sample complexity for kernel bandits (that is a special case of our setting when $|\mathcal{S}|=1$ ). In particular, Scarlett et al. (2017) proved $\Omega\left(\left(\frac{1}{\epsilon}\right)^{2+\frac{d}{\nu}}\right)$ and $\Omega\left(\frac{1}{\epsilon^{2}}\right)$ sample complexities for the Matérn and SE kernels, respectively. Our results are the first finite sample complexities for the RL problem under a very general case which includes all kernels with polynomially decaying eigenvalues.

In terms of the discount factor, our sample complexities scale with $\frac{1}{(1-\gamma)^{7}}$ in the case of smooth kernels (similar to the PPQ-Learning algorithm, in the linear setting, Yang and Wang, 2019). In the tabular and linear settings, however, this has been improved to $\mathcal{O}\left(\frac{1}{(1-\gamma)^{3}}\right)$. It remains an interesting problem for future investigation that whether the dependency of the sample complexity on the discount factor can be improved to $\mathcal{O}\left(\frac{1}{(1-\gamma)^{3}}\right)$, also in the kernel setting. As discussed in the introduction, in the kernel setting, neither observing all state-actions nor a parametric update of the model through value iteration is feasible. Thus, a different approach to algorithm design and analysis is required that is increasingly more challenging among these settings: tabular $\rightarrow$ linear $\rightarrow$ kernel-based.

## 5 ANALYSIS

Theorem 2 is a consequence of Theorem 1, and using the kernel specific bounds on $\Gamma_{K, \lambda}$. The proof of Theorem 1 builds on several components including tracking the approximation error in kernel ridge regression and convergence error of an approximate Bellman operator. In this section, we overview the main steps in the proof of Theorem 1 , while deferring the details to the appendix.

Approximate Bellman operator: Recall the Bellman operator defined in Section 2.1. The transition probability distribution and the value function are complex nonlinear functions on continuous domains, unknown to the algorithm. We thus define an approximate Bellman operator $\widehat{\mathcal{T}}$, which uses noisy observations of $P V$ on a fixed set $\mathcal{U}_{J}$, and takes advantage of kernel ridge regression, to perform an approximate Bellman operator update. In particular, for all $V: \mathcal{S} \rightarrow\left[0, \frac{1}{1-\gamma}\right]$, and a fixed $\operatorname{set} \mathcal{U}_{J} \subset \mathcal{Z}$, let us define

$$
\begin{align*}
{[\widehat{\mathcal{T}} V](s) } & =\max _{a \in \mathcal{A}}\{r(s, a) \\
& \left.+\gamma k_{\mathcal{U}_{J}}^{\top}(s, a)\left(K_{\mathcal{U}_{J}}+\lambda^{2} I_{J}\right)^{-1}[\widehat{P} V]_{\mathcal{U}_{J}}\right\} \tag{11}
\end{align*}
$$

where $[\widehat{P} V](s, a)=V\left(s^{\prime}\right)$ is a random variable, $s^{\prime} \sim$ $P(\cdot \mid s, a)$ is a random transition state, and

$$
[\widehat{P} V]_{\mathcal{U}_{J}}=\left[[\widehat{P} V]\left(s_{1}, a_{1}\right), \ldots,[\widehat{P} V]\left(s_{J}, a_{J}\right)\right]^{\top}
$$

In the KQLearn algorithm, define

$$
\begin{align*}
\widehat{V}^{(\ell)}(s)= & \max _{a \in \mathcal{A}}\{r(s, a) \\
& \left.+\gamma k_{\mathcal{U}_{J}}^{\top}(s, a)\left(K_{\mathcal{U}_{J}}+\lambda^{2} I_{J}\right)^{-1} Y_{\mathcal{U}_{J}}^{(\ell)}\right\} \tag{12}
\end{align*}
$$

which we refer to as proxy value function (similar to the proxy Q-function given in (8). We then have the following recursive relation over $\ell$.

$$
\begin{equation*}
\widehat{V}^{(\ell)}=\widehat{\mathcal{T}} \Pi_{\left[0, \frac{1}{1-\gamma}\right]}\left[\widehat{V}^{(\ell-1)}\right] . \tag{13}
\end{equation*}
$$

Error in proxy value function: In order to bound the error in the value function of the policy $\pi$ obtained by KQLearn, $\left\|V^{\pi}-V^{*}\right\|$, we need to bound the error in the proxy Q-function given in (8), which is used to obtain $\pi$. The error in proxy Q-function can be bounded based on the error in the proxy value function at round $L-1$. In particular, we have

$$
\begin{equation*}
\left\|\widehat{Q}^{(L)}-Q^{*}\right\|_{\infty} \leq\left\|\widehat{V}^{(L-1)}-V^{*}\right\|_{\infty} \tag{14}
\end{equation*}
$$

Therefore, we next bound the error in the proxy value function. We can write the error in proxy value function as the
sum of two terms: the error in approximate Bellman operator and the error in the value function using true Bellman operator. Specifically,

$$
\begin{align*}
& \left\|\widehat{V}^{(L-1)}-V^{*}\right\|_{\infty}=\left\|\widehat{\mathcal{T}} \Pi_{\left[0, \frac{1}{1-\gamma}\right.} \widehat{V}^{(L-2)}-V^{*}\right\|_{\infty} \\
& \quad \leq\left\|\widehat{\mathcal{T}} \Pi_{\left[0, \frac{1}{1-\gamma}\right]} \widehat{V}^{(L-2)}-\mathcal{T} \Pi_{\left[0, \frac{1}{1-\gamma}\right]} \widehat{V}^{(L-2)}\right\|_{\infty} \\
& \quad+\left\|\mathcal{T} \Pi_{\left[0, \frac{1}{1-\gamma}\right]} \widehat{V}^{(L-2)}-V^{*}\right\|_{\infty} \tag{15}
\end{align*}
$$

The second term can be recursively bounded which leads to the second term $\frac{2 \gamma^{L-1}}{(1-\gamma)^{2}}$ in the error bound in Theorem 1. The first term leads to an important step in the analysis which is based on the error in kernel ridge regression. Specifically, let $V: \mathcal{S} \rightarrow\left[0, \frac{1}{1-\gamma}\right]$ be any value function. We have, for all $s \in \mathcal{S}$,

$$
\begin{align*}
& {[\widehat{\mathcal{T}} V](s)-[\mathcal{T} V](s)} \\
& =\max _{a \in \mathcal{A}}\left\{r(s, a)+\gamma k_{\mathcal{U}_{J}}^{\top}(s, a)\left(K_{\mathcal{U}_{J}}+\lambda^{2} I_{J}\right)^{-1}[\widehat{P} V]_{\mathcal{U}_{J}}\right\} \\
& \quad-\max _{a \in \mathcal{A}}\{r(s, a)+\gamma[P V](s, a)\} \\
& \leq \gamma \max _{a \in \mathcal{A}}\left\{k_{\mathcal{U}_{J}}^{\top}(s, a)\left(K_{\mathcal{U}_{J}}+\lambda^{2} I_{J}\right)^{-1}[\widehat{P} V]_{\mathcal{U}_{J}}\right. \\
& \quad-[P V](s, a)\} \tag{16}
\end{align*}
$$

Error in kernel ridge regression: The term inside max in Equation 16 is the error in kernel ridge regression, where $P V$ is the target function, $[\widehat{P} V]_{\mathcal{U}_{J}}$ is a set of $J$ noisy observations, and $k_{\mathcal{U}_{J}}^{\top}(s, a)\left(K_{\mathcal{U}_{J}}+\lambda^{2} I_{J}\right)^{-1}[\widehat{P} V]_{\mathcal{U}_{J}}$ is the regressor. In order to apply Lemma 1, we need an upper bound on the RKHS norm of $P V$, as well as an upper bound on the sub-Gaussianity parameter of the observation noise in $\widehat{P} V$. These are established in the following lemmas.
Lemma 3. Consider an integrable value function $V: \mathcal{S} \rightarrow$ [ $\left.0, \frac{1}{1-\gamma}\right]$. Under Assumption 2, we have

$$
\begin{equation*}
\|P V\|_{\mathcal{H}_{K}} \leq \frac{c}{1-\gamma} \tag{17}
\end{equation*}
$$

where $c=\int_{\mathcal{S}} d s$ is the volume of $\mathcal{S}$.
Lemma 4. Consider a transition probability distribution $P$, and an integrable value function $V: \mathcal{S} \rightarrow\left[0, \frac{1}{1-\gamma}\right]$. We have, for all $(s, a), \mathbb{E}[[\widehat{P} V](s, a)]=P V(s, a)$. In addition, $[\widehat{P} V](s, a)$ is a sub-Gaussian random variable with parameter $\frac{1}{2(1-\gamma)}$.

Lemma 4 follows from the definition of $\widehat{P} V$, as well as Hoeffding lemma for bounded random variables. A proof of Lemma 3 is provided in Appendix C.

Applying Lemma 1, we obtain, with probability $1-\delta$, for all $(s, a) \in \mathcal{Z}$,

$$
\begin{align*}
& \left|k_{\mathcal{U}_{J}}^{\top}(s, a)\left(K_{\mathcal{U}_{J}}+\lambda^{2} I_{J}\right)^{-1}[\widehat{P} V]_{\mathcal{U}_{J}}-[P V](s, a)\right|  \tag{18}\\
& \quad \leq \beta(\delta) \Sigma_{\mathcal{U}_{J}}(s, a)
\end{align*}
$$

where $\beta(\delta)=\mathcal{O}\left(\frac{c}{1-\gamma}+\frac{1}{2(1-\gamma) \lambda} \sqrt{d \log \left(\frac{J c}{(1-\gamma) \delta}\right)}\right)$.
Eventually, using Lemma 2 on the total uncertainty, and by the design of $\mathcal{U}_{J}$, we bound $\Sigma_{\mathcal{U}_{J}}(s, a)$ on the right hand side.

We thus bounded the two terms in 21, that bounds the error in proxy value function. More details on the proof of theorems and the proof of lemmas are provided in Appendix $B$ and Appendix C, respectively.

## 6 CONCLUSION

Modern RL often faces an enormous state-action space and complex models. We considered the question of sample complexity in a discounted MDP with a generative model under the kernel setting, furthering a line of research in the literature (e.g., see Kearns and Singh, 1998, Azar et al., 2017; Sidford et al. 2018a b; Yang and Wang, 2019). We introduced a novel kernel-based Q learning algorithm referred to as KQLearn and proved a finite bound on its sample complexity for very general classes of kernels. That is to the best of our knowledge the first finite sample complexity result under the general kernel setting (including all kernels with polynomially decaying eigenvalues). In addition, compared to the lower bounds on the special case of the kernel bandit problem, our sample complexities are tight with respect to $\epsilon$ in finding an $\epsilon$-optimal policy. Our sample complexities, however, scale possibly suboptimally with respect to the discount factor, which remains an interesting open problem for future investigation.

## References

Abbasi-Yadkori, Y. (2013). Online learning for linearly parametrized control problems.

Agarwal, A., Jiang, N., Kakade, S. M., and Sun, W. (2019). Reinforcement learning: Theory and algorithms. CS Dept., UW Seattle, Seattle, WA, USA, Tech. Rep, pages 10-4.

Agarwal, A., Kakade, S., and Yang, L. F. (2020). Modelbased reinforcement learning with a generative model is minimax optimal. In Conference on Learning Theory, pages 67-83. PMLR.

Ayoub, A., Jia, Z., Szepesvari, C., Wang, M., and Yang, L. (2020). Model-based reinforcement learning with valuetargeted regression. In International Conference on Machine Learning, pages 463-474. PMLR.

Azar, M. G., Munos, R., and Kappen, H. J. (2013). Minimax pac bounds on the sample complexity of reinforcement learning with a generative model. Machine learning, 91(3):325-349.
Azar, M. G., Osband, I., and Munos, R. (2017). Minimax regret bounds for reinforcement learning. In International Conference on Machine Learning, pages 263-272. PMLR.
Carpentier, A., Vernade, C., and Abbasi-Yadkori, Y. (2020). The elliptical potential lemma revisited. arXiv preprint arXiv:2010.10182.

Chatterji, N., Pacchiano, A., and Bartlett, P. (2019). Online learning with kernel losses. In Proceedings of Machine Learning Research, volume 97, pages 971-980, Long Beach, California, USA. PMLR.
Chowdhury, S. R. and Gopalan, A. (2017). On kernelized multi-armed bandits. In International Conference on Machine Learning, pages 844-853. PMLR.

Christmann, A. and Steinwart, I. (2008). Support Vector Machines. Springer New York, NY.
Domingues, O. D., Ménard, P., Pirotta, M., Kaufmann, E., and Valko, M. (2021). Kernel-based reinforcement learning: A finite-time analysis. In International Conference on Machine Learning, pages 2783-2792. PMLR.

Fawzi, A., Balog, M., Huang, A., Hubert, T., RomeraParedes, B., Barekatain, M., Novikov, A., R Ruiz, F. J., Schrittwieser, J., Swirszcz, G., et al. (2022). Discovering faster matrix multiplication algorithms with reinforcement learning. Nature, 610(7930):47-53.
Jin, C., Allen-Zhu, Z., Bubeck, S., and Jordan, M. I. (2018). Is q-learning provably efficient? Advances in Neural Information Processing Systems, 31.
Jin, C., Yang, Z., Wang, Z., and Jordan, M. I. (2020). Provably efficient reinforcement learning with linear function approximation. In Conference on Learning Theory, pages 2137-2143. PMLR.

Kahn, G., Villaflor, A., Pong, V., Abbeel, P., and Levine, S. (2017). Uncertainty-aware reinforcement learning for collision avoidance. arXiv preprint arXiv:1702.01182.
Kakade, S., Krishnamurthy, A., Lowrey, K., Ohnishi, M., and Sun, W. (2020). Information theoretic regret bounds for online nonlinear control. Advances in Neural Information Processing Systems, 33:15312-15325.

Kakade, S. M. (2003). On the sample complexity of reinforcement learning. University of London, University College London (United Kingdom).
Kalashnikov, D., Irpan, A., Pastor, P., Ibarz, J., Herzog, A., Jang, E., Quillen, D., Holly, E., Kalakrishnan, M., Vanhoucke, V., et al. (2018). Scalable deep reinforcement learning for vision-based robotic manipulation. In Conference on Robot Learning, pages 651-673. PMLR.

Kearns, M. and Singh, S. (1998). Finite-sample convergence rates for q-learning and indirect algorithms. In Advances in Neural Information Processing Systems, volume 11. MIT Press.

Lee, K., Kim, S.-A., Choi, J., and Lee, S.-W. (2018). Deep reinforcement learning in continuous action spaces: a case study in the game of simulated curling. In International Conference on Machine Learning,, pages 29372946. PMLR.

Li, Z. and Scarlett, J. (2022). Gaussian process bandit optimization with few batches. In International Conference on Artificial Intelligence and Statistics, pages 92-107. PMLR.
Mercer, J. (1909). Functions of positive and negative type, and their connection with the theory of integral equations. Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character, 209:415-446.
Mirhoseini, A., Goldie, A., Yazgan, M., Jiang, J. W., Songhori, E., Wang, S., Lee, Y.-J., Johnson, E., Pathak, O., Nazi, A., et al. (2021). A graph placement methodology for fast chip design. Nature, 594(7862):207-212.

Ormoneit, D. and Sen, A. (2002). Kernel-based reinforcement learning. Machine learning, 49(2-3):161.
Puterman, M. L. (2014). Markov decision processes: discrete stochastic dynamic programming. John Wiley \& Sons.
Russo, D. (2019). Worst-case regret bounds for exploration via randomized value functions. Advances in Neural Information Processing Systems, 32.
Salgia, S., Vakili, S., and Zhao, Q. (2021). A domainshrinking based bayesian optimization algorithm with order-optimal regret performance. Advances in Neural Information Processing Systems, 34:28836-28847.
Scarlett, J., Bogunovic, I., and Cevher, V. (2017). Lower bounds on regret for noisy gaussian process bandit optimization. In Conference on Learning Theory, pages 1723-1742. PMLR.
Schölkopf, B., Smola, A. J., Bach, F., et al. (2002). Learning with kernels: support vector machines, regularization, optimization, and beyond. MIT press.
Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and De Freitas, N. (2015). Taking the human out of the loop: A review of bayesian optimization. Proceedings of the IEEE, 104(1):148-175.

Sidford, A., Wang, M., Wu, X., Yang, L., and Ye, Y. (2018a). Near-optimal time and sample complexities for solving markov decision processes with a generative model. Advances in Neural Information Processing Systems, 31 .

Sidford, A., Wang, M., Wu, X., and Ye, Y. (2018b). Variance reduced value iteration and faster algorithms for
solving markov decision processes. In Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, pages 770-787. SIAM.
Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. (2016). Mastering the game of go with deep neural networks and tree search. nature, 529(7587):484-489.
Snoek, J., Larochelle, H., and Adams, R. P. (2012). Practical bayesian optimization of machine learning algorithms. Advances in neural information processing systems, 25.

Srinivas, N., Krause, A., Kakade, S., and Seeger, M. (2010). Gaussian process optimization in the bandit setting: No regret and experimental design. In ICML 2010 Proceedings, 27th International Conference on Machine Learning, pages 1015-1022.

Vakili, S., Bouziani, N., Jalali, S., Bernacchia, A., and Shiu, D.-S. (2021a). Optimal order simple regret for gaussian process bandits. Advances in Neural Information Processing Systems, 34:21202-21215.
Vakili, S., Bromberg, M., Garcia, J., Shiu, D.-s., and Bernacchia, A. (2021b). Uniform generalization bounds for overparameterized neural networks. arXiv preprint arXiv:2109.06099.
Vakili, S., Khezeli, K., and Picheny, V. (2021c). On information gain and regret bounds in gaussian process bandits. In International Conference on Artificial Intelligence and Statistics, pages 82-90. PMLR.

Vakili, S., Scarlett, J., and Javidi, T. (2021d). Open problem: Tight online confidence intervals for rkhs elements. In Conference on Learning Theory, pages 4647-4652. PMLR.

Vakili, S., Scarlett, J., Shiu, D.-S., and Bernacchia, A. (2022). Improved convergence rates for sparse approximation methods in kernel-based learning. arXiv preprint arXiv:2202.04005.
Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., Choi, D. H., Powell, R., Ewalds, T., Georgiev, P., et al. (2019). Grandmaster level in starcraft ii using multi-agent reinforcement learning. Nature, 575(7782):350-354.
Wang, R., Salakhutdinov, R., and Yang, L. F. (2020). Provably efficient reinforcement learning with general value function approximation. arXiv preprint arXiv:2005.10804.

Yang, L. and Wang, M. (2019). Sample-optimal parametric q-learning using linearly additive features. In International Conference on Machine Learning, pages 69957004. PMLR.

Yang, L. and Wang, M. (2020). Reinforcement learning in feature space: Matrix bandit, kernels, and regret
bound. In International Conference on Machine Learning, pages 10746-10756. PMLR.
Yang, Z., Jin, C., Wang, Z., Wang, M., and Jordan, M. (2020a). Provably efficient reinforcement learning with kernel and neural function approximations. Advances in Neural Information Processing Systems, 33:1390313916.

Yang, Z., Jin, C., Wang, Z., Wang, M., and Jordan, M. I. (2020b). On function approximation in reinforcement learning: Optimism in the face of large state spaces. arXiv preprint arXiv:2011.04622.
Zhou, D., He, J., and Gu, Q. (2021). Provably efficient reinforcement learning for discounted mdps with feature mapping. In International Conference on Machine Learning, pages 12793-12802. PMLR.

In the appendix, we provide some details and proofs omitted from the main paper due to space limit. In Appendix A we provide a formal statement of Mercer theorem and a constructive definition of the RKHS. The proof of Theorems 1 and 2 , and auxiliary lemmas are provided in Appendix B and Appendix C respectively.

## A MERCER THEOREM

Mercer theorem (Mercer, 1909) provides a representation of the kernel in terms of an infinite dimensional feature map (see, e.g. Christmann and Steinwart (2008), Theorem 4.49). Let $\mathcal{Z}$ be a compact metric space and $\mu$ be a finite Borel measure on $\mathcal{Z}$ (we consider Lebesgue measure in a Euclidean space). Let $L_{\mu}^{2}(\mathcal{Z})$ be the set of square-integrable functions on $\mathcal{Z}$ with respect to $\mu$. We further say a kernel is square-integrable if

$$
\int_{\mathcal{Z}} \int_{\mathcal{Z}} K\left(z, z^{\prime}\right)^{2} d \mu(z) d \mu\left(z^{\prime}\right)<\infty
$$

Theorem 3. (Mercer Theorem) Let $\mathcal{Z}$ be a compact metric space and $\mu$ be a finite Borel measure on $\mathcal{Z}$. Let $K$ be a continuous and square-integrable kernel, inducing an integral operator $T_{K}: L_{\mu}^{2}(\mathcal{Z}) \rightarrow L_{\mu}^{2}(\mathcal{Z})$ defined by

$$
\left(T_{K} f\right)(\cdot)=\int_{\mathcal{Z}} K\left(\cdot, z^{\prime}\right) f\left(z^{\prime}\right) d \mu\left(z^{\prime}\right)
$$

where $f \in L_{\mu}^{2}(\mathcal{Z})$. Then, there exists a sequence of eigenvalue-eigenfunction pairs $\left\{\left(\sigma_{m}, \psi_{m}\right)\right\}_{m=1}^{\infty}$ such that $\sigma_{m}>0$, and $T_{K} \psi_{m}=\sigma_{m} \psi_{m}$, for $m \geq 1$. Moreover, the kernel function can be represented as

$$
K\left(z, z^{\prime}\right)=\sum_{m=1}^{\infty} \sigma_{m} \psi_{m}(z) \psi_{m}\left(z^{\prime}\right)
$$

where the convergence of the series holds uniformly on $\mathcal{Z} \times \mathcal{Z}$.
According to Mercer representation theorem (see, e.g., Christmann and Steinwart (2008), Theorem 4.51), the RKHS induced by $K$ can consequently be represented in terms of $\left\{\left(\sigma_{m}, \psi_{m}\right)\right\}_{m=1}^{\infty}$.
Theorem 4. (Mercer Representation Theorem) Let $\left\{\left(\sigma_{m}, \psi_{m}\right)\right\}_{i=1}^{\infty}$ be the Mercer eigenvalue eigenfunction pairs. Then, the RKHS of $K$ is given by

$$
\mathcal{H}_{K}=\left\{f(\cdot)=\sum_{i=1}^{\infty} w_{i} \sigma_{i}^{\frac{1}{2}} \psi_{i}(\cdot): w_{i} \in \mathbb{R},\|f\|_{\mathcal{H}_{K}}^{2}:=\sum_{i=1}^{\infty} w_{i}^{2}<\infty\right\}
$$

Mercer representation theorem indicates that the scaled eigenfunctions $\left\{\sqrt{\sigma_{i}} \psi_{i}\right\}_{i=1}^{\infty}$ form an orthonormal basis for $\mathcal{H}_{K}$.

## B PROOF OF THEOREMS

In this section, we provide the proof of main theorems.

## B. 1 Proof of Theorem 1 .

The proof of Theorem 1 builds on an approximate Bellman operator, that uses noisy observations of $P V$ within the rounds of KQLearn, and kernel ridge regression. We prove bounds on the error of this approximate Bellman operator. That is then used to bound the error in the value function of the policy obtained by KQLearn.

Approximate Bellman operator: Recall the Bellman operator defined in Section 2.1 The transition probability distribution and the value function are complex non-linear functions on continuous domains, unknown to the algorithm. We thus define an approximate Bellman operator $\widehat{\mathcal{T}}$, which uses noisy observations of $P V$ on a fixed set $\mathcal{U}_{J}$, and takes advantage of kernel ridge regression, to perform an approximate Bellman operator update. In particular, for all $V: \mathcal{S} \rightarrow\left[0, \frac{1}{1-\gamma}\right]$, and a fixed $\operatorname{set} \mathcal{U}_{J} \subset \mathcal{Z}$, let us define

$$
[\widehat{\mathcal{T}} V](s)=\max _{a \in \mathcal{A}}\left\{r(s, a)+\gamma k_{\mathcal{U}_{J}}^{\top}(s, a)\left(K_{\mathcal{U}_{J}}+\lambda^{2} I_{J}\right)^{-1}[\widehat{P} V]_{\mathcal{U}_{J}}\right\}
$$

where $[\widehat{P} V](s, a)=V\left(s^{\prime}\right)$ is a random variable, $s^{\prime} \sim P(\cdot \mid s, a)$ is a random transition state, and

$$
\left.[\widehat{P} V]_{\mathcal{U}_{J}}=\left[[\widehat{P} V]\left(s_{1}, a_{1}\right), \ldots,[\widehat{P} V]\left(s_{J}, a_{J}\right)\right]\right]^{\top}
$$

In the KQLearn algorithm, define

$$
\left.\widehat{V}^{(\ell)}(s)=\max _{a \in \mathcal{A}}\left\{r(s, a)+\gamma k_{\mathcal{U}_{J}}^{\top}(s, a)\left(K_{\mathcal{U}_{J}}+\lambda^{2} I_{J}\right)^{-1} Y_{\mathcal{U}_{J}}^{(\ell)}\right]\right\}
$$

which we refer to as proxy value function (similar to the proxy Q-function given in (8)). We then have the following recursive relation over $\ell$.

$$
\begin{equation*}
\widehat{V}^{(\ell)}=\widehat{\mathcal{T}} \Pi_{\left[0, \frac{1}{1-\gamma}\right]}\left[\widehat{V}^{(\ell-1)}\right] . \tag{19}
\end{equation*}
$$

Error in proxy value function: We next bound the error in the proxy value function. We can write the error in proxy value function as the sum of two terms: the error in approximate Bellman operator and the error in the value function using true Bellman operator. Specifically, for $l>1$,

$$
\begin{align*}
& \left\|\widehat{V}^{(l-1)}-V^{*}\right\|_{\infty}=\left\|\widehat{\mathcal{T}} \Pi_{\left[0, \frac{1}{1-\gamma}\right]} \widehat{V}^{(l-2)}-V^{*}\right\|_{\infty} \\
& \leq \underbrace{\left\|\mathcal{T} \Pi_{\left[0, \frac{1}{1-\gamma}\right]} \widehat{V}^{(l-2)}-\mathcal{T} \Pi_{\left[0, \frac{1}{1-\gamma}\right.} \widehat{V}^{(l-2)}\right\|_{\infty}}_{\text {Term I }} \\
& \quad+\underbrace{\left\|\mathcal{T} \Pi_{\left[0, \frac{1}{1-\gamma}\right]} \widehat{V}^{(l-2)}-V^{*}\right\|_{\infty}}_{\text {Term II }} . \tag{20}
\end{align*}
$$

We now bound the two terms on the right hand side of 20 .

Term I: The first term leads us to an important step in the analysis which is based on the error in kernel ridge regression. Specifically, let $V: \mathcal{S} \rightarrow\left[0, \frac{1}{1-\gamma}\right]$ be any value function. We have, for all $s \in \mathcal{S}$,

$$
\begin{align*}
{[\widehat{\mathcal{T}} V](s)-[\mathcal{T} V](s) } & =\max _{a \in \mathcal{A}}\left\{r(s, a)+\gamma k_{\mathcal{U}_{J}}^{\top}(s, a)\left(K_{\mathcal{U}_{J}}+\lambda^{2} I_{J}\right)^{-1}[\widehat{P} V]_{\mathcal{U}_{J}}\right\}-\max _{a \in \mathcal{A}}\{r(s, a)+\gamma[P V](s, a)\} \\
& \leq \gamma \max _{a \in \mathcal{A}}\left\{k_{\mathcal{U}_{J}}^{\top}(s, a)\left(K_{\mathcal{U}_{J}}+\lambda^{2} I_{J}\right)^{-1}[\widehat{P} V]_{\mathcal{U}_{J}}-[P V](s, a)\right\} \tag{21}
\end{align*}
$$

Error in kernel ridge regression: The term inside max in Equation 21 is the error in kernel ridge regression, where $P V$ is the target function, $[\widehat{P} V]_{\mathcal{U}_{J}}$ is a set of $J$ noisy observations, and $k_{\mathcal{U}_{J}}^{\top}(s, a)\left(K_{\mathcal{U}_{J}}+\lambda^{2} I_{J}\right)^{-1}[\widehat{P} V]_{\mathcal{U}_{J}}$ is the regressor. In order to apply Lemma 1, we need an upper bound on the RKHS norm of $P V$, as well as an upper bound on the subGaussianity parameter of the observation noise in $\widehat{P} V$. These are established in Lemmas 3 and 4 , respectively. Specifically, we have

$$
\begin{equation*}
\|P V\|_{\mathcal{H}_{K}} \leq \frac{c}{1-\gamma} \tag{22}
\end{equation*}
$$

And, $[\widehat{P} V](s, a)$ is a sub-Gaussian random variable with parameter $\frac{1}{2(1-\gamma)}$.
Applying Lemma 1 we obtain, with probability $1-\delta$, for all $(s, a) \in \mathcal{Z}$,

$$
\begin{equation*}
\left|k_{\mathcal{U}_{J}}^{\top}(s, a)\left(K_{\mathcal{U}_{J}}+\lambda^{2} I_{J}\right)^{-1}[\widehat{P} V]_{\mathcal{U}_{J}}-[P V](s, a)\right| \leq \beta(\delta) \Sigma_{\mathcal{U}_{J}}(s, a) \tag{23}
\end{equation*}
$$

where $\beta(\delta)=\mathcal{O}\left(\frac{c}{1-\gamma}+\frac{1}{2(1-\gamma) \lambda} \sqrt{d \log \left(\frac{J c}{(1-\gamma) \delta}\right)}\right)$.

Bounding $\Sigma_{\mathcal{U}_{J}}(s, a)$ : Conditioning on a smaller subset of observation reduces the variance $\Sigma_{\mathcal{U}_{j}}(s, a) \geq \Sigma_{\mathcal{U}_{J}}(s, a)$, for all $j \leq J$ (due to positive definiteness of the kernel matrix). By the selection rule of the observation points:

$$
\begin{equation*}
\left(s_{j}, a_{j}\right)=\underset{(s, a) \in \mathcal{Z}}{\arg \max } \Sigma_{\mathcal{U}_{j-1}}^{2}(s, a), \tag{24}
\end{equation*}
$$

we have $\Sigma_{\mathcal{U}_{j-1}}(s, a) \leq \Sigma_{\mathcal{U}_{j-1}}\left(s_{j}, a_{j}\right)$. Thus, for all $(s, a) \in \mathcal{Z}$,

$$
\begin{aligned}
\Sigma_{\mathcal{U}_{J}}^{2}(s, a) & \leq \frac{1}{J} \sum_{j=1}^{J} \Sigma_{\mathcal{U}_{j-1}}^{2}(s, a) \\
& \leq \frac{1}{J} \sum_{j=1}^{J} \Sigma_{\mathcal{U}_{j-1}}^{2}\left(s_{j}, a_{j}\right) \\
& \leq \frac{2 \Gamma_{K, \lambda}(J)}{\log \left(1+1 / \lambda^{2}\right) J}
\end{aligned}
$$

where the last line follows from Lemma 2 ,
Replacing the bound on $\Sigma_{\mathcal{U}_{J}}(s, a)$, we obtain, for all $V: \mathcal{S} \rightarrow\left[0, \frac{1}{1-\gamma}\right], s \in \mathcal{S}$,

$$
\begin{equation*}
[\widehat{\mathcal{T}} V](s)-[\mathcal{T} V](s) \leq \gamma \beta(\delta) \sqrt{\frac{2 \Gamma_{K, \lambda}(J)}{\log \left(1+1 / \lambda^{2}\right) J}} \tag{25}
\end{equation*}
$$

Thus,

$$
\begin{equation*}
\left\|\widehat{\mathcal{T}} \Pi_{\left[0, \frac{1}{1-\gamma}\right]} \widehat{V}^{(l-2)}-\mathcal{T} \Pi_{\left[0, \frac{1}{1-\gamma}\right]} \widehat{V}^{(l-2)}\right\|_{\infty} \leq \gamma \beta(\delta) \sqrt{\frac{2 \Gamma_{K, \lambda}(J)}{\log \left(1+1 / \lambda^{2}\right) J}} \tag{26}
\end{equation*}
$$

Term II: We now bound the second term on the right hand side of 20, by the contraction of Bellman operator.

$$
\begin{align*}
\left\|\mathcal{T} \Pi_{\left[0, \frac{1}{1-\gamma}\right]} \widehat{V}^{(l-2)}-V^{*}\right\|_{\infty} & \leq \gamma\left\|\Pi_{\left[0, \frac{1}{1-\gamma}\right]} \widehat{V}^{(l-2)}-V^{*}\right\|_{\infty} \\
& \leq \gamma\left\|\widehat{V}^{(l-2)}-V^{*}\right\|_{\infty} \tag{27}
\end{align*}
$$

The second inequality follows from the observation that $V^{*}(s) \in\left[0, \frac{1}{1-\gamma}\right]$ for all $s \in \mathcal{S}$.
Combing the bounds on Term I and Term II, we obtain

$$
\left\|\widehat{V}^{(l-1)}-V^{*}\right\|_{\infty} \leq \gamma \beta(\delta) \sqrt{\frac{2 \Gamma_{K, \lambda}(J)}{\log \left(1+1 / \lambda^{2}\right) J}}+\gamma\left\|\widehat{V}^{(l-2)}-V^{*}\right\|_{\infty}
$$

Recursively bounding the error in proxy value function at round $l$ using the error at round $l-1$ for $l=2, \ldots, L-1$, we have,

$$
\begin{aligned}
\left\|\widehat{V}^{(L-1)}-V^{*}\right\|_{\infty} & \leq \beta(\delta) \sqrt{\frac{2 \Gamma_{K, \lambda}(J)}{\log \left(1+1 / \lambda^{2}\right) J}}\left(\sum_{i=1}^{L-1} \gamma^{i}\right)+\gamma^{L-1}\left\|\widehat{V}^{(0)}-V^{*}\right\|_{\infty} \\
& \leq \beta(\delta) \sqrt{\frac{2 \Gamma_{K, \lambda}(J)}{\log \left(1+1 / \lambda^{2}\right) J}}\left(\sum_{i=1}^{L-1} \gamma^{i}\right)+\frac{\gamma^{L-1}}{1-\gamma}
\end{aligned}
$$

where the second inequality comes from $\left\|\widehat{V}^{(0)}-V^{*}\right\|_{\infty} \leq \frac{1}{1-\gamma}$.

Recall the definition of the proxy Q-function $\widehat{Q}^{(L)}$ given in (8). We bound the error in $\widehat{Q}^{(L)}$ as follows. For all $(s, a) \in \mathcal{Z}$,

$$
\begin{aligned}
\left\|\widehat{Q}^{(L)}(s, a)-Q^{*}(s, a)\right\|_{\infty} & =\left\|\left(r(s, a)+\gamma k_{\mathcal{U}_{J}}^{\top}(s, a)\left(K_{\mathcal{U}_{J}}+\lambda^{2} I_{J}\right)^{-1} Y_{\mathcal{U}_{J}}^{(L)}\right)-\left(r(s, a)+\gamma\left[P V^{*}\right](s, a)\right)\right\|_{\infty} \\
\leq & \| \gamma\left(k _ { \mathcal { U } _ { J } } ^ { \top } ( s , a ) ( K _ { \mathcal { U } _ { J } } + \lambda ^ { 2 } I _ { J } ) ^ { - 1 } \left[{\left.\left.\widehat{P} \Pi_{\left[0, \frac{1}{1-\gamma}\right]} \widehat{V}^{(L-1)}\right] \mathcal{U}_{J}-\left[P \Pi_{\left[0, \frac{1}{1-\gamma}\right.} \widehat{V}^{(L-1)}\right](s, a)\right) \|_{\infty}} \quad+\left\|\gamma\left(\left[P \Pi_{\left[0, \frac{1}{1-\gamma}\right]} \widehat{V}^{(L-1)}\right](s, a)-\left[P V^{*}\right](s, a)\right)\right\|_{\infty}\right.\right. \\
\leq & \gamma \beta(\delta) \sqrt{\frac{2 \Gamma_{K, \lambda}(J)}{\log \left(1+1 / \lambda^{2}\right) J}}+\gamma\left\|\mathbb{E}_{s^{\prime} \sim P(\cdot \mid s, a)}\left[\Pi_{\left[0, \frac{1}{1-\gamma}\right]} \widehat{V}^{(L-1)}\right]-\mathbb{E}_{s^{\prime} \sim P(\cdot \mid s, a)}\left[V^{*}\right]\right\|_{\infty} \\
\leq & \gamma \beta(\delta) \sqrt{\frac{2 \Gamma_{K, \lambda}(J)}{\log \left(1+1 / \lambda^{2}\right) J}}+\gamma\left\|\Pi_{\left[0, \frac{1}{1-\gamma}\right]} \widehat{V}^{(L-1)}-V^{*}\right\|_{\infty} \\
\leq & \gamma \beta(\delta) \sqrt{\frac{2 \Gamma_{K, \lambda}(J)}{\log \left(1+1 / \lambda^{2}\right) J}}+\gamma\left\|\widehat{V}^{(L-1)}-V^{*}\right\|_{\infty} \\
\leq & \gamma \beta(\delta) \sqrt{\frac{2 \Gamma_{K, \lambda}(J)}{\log \left(1+1 / \lambda^{2}\right) J}}+\gamma\left(\beta(\delta) \sqrt{\frac{2 \Gamma_{K, \lambda}(J)}{\log \left(1+1 / \lambda^{2}\right) J}}\left(\sum_{i=1}^{L-1} \gamma^{i}\right)+\frac{\gamma^{L-1}}{1-\gamma}\right) \\
& =\beta(\delta) \sqrt{\frac{2 \Gamma_{K, \lambda}(J)}{\log \left(1+1 / \lambda^{2}\right) J}}\left(\sum_{i=1}^{L} \gamma^{i}\right)+\frac{\gamma^{L}}{1-\gamma} \\
& \leq \frac{\gamma \beta(\delta)}{1-\gamma \sqrt{\frac{2 \Gamma_{K, \lambda}(J)}{\log \left(1+1 / \lambda^{2}\right) J}}+\frac{\gamma^{L}}{1-\gamma}} .
\end{aligned}
$$

The first inequality is a result of triangle inequality after adding and subtracting the term $\gamma\left[P \Pi_{\left[0, \frac{1}{1-\gamma}\right]} \widehat{V}^{(L-1)}\right](s, a)$. The second inequality is the error in kernel ridge regression bounded above in Term I. The third inequality bounds the difference in expectations with maximum difference. The fourth inequality is a consequence of the observation that $V^{*} \in\left[0, \frac{1}{1-\gamma}\right]$ interval. The fifth inequality is obtained using the bound on the error in the proxy value function given above.

The value function of $\pi$ obtained from the proxy $\mathbf{Q}$-function: The following lemma establishes that the error in the value function of $\pi$ can be bounded using the error in the proxy Q -function.
Lemma 5. Consider any $Q$-function $\tilde{Q}: \mathcal{S} \times \mathcal{A} \rightarrow\left[0, \frac{1}{1-\gamma}+1\right]$ satisfying $\left\|\tilde{Q}-Q^{*}\right\|_{\infty} \leq \epsilon$. Define policy $\pi_{\tilde{Q}}$ such that $\pi_{\tilde{Q}}(s)=\underset{a \in \mathcal{A}}{\arg \max } \tilde{Q}(s, a)$. Then, for all $s \in \mathcal{S}$,

$$
V^{*}(s)-V^{\pi_{\tilde{Q}}}(s) \leq \frac{2 \epsilon}{1-\gamma}
$$

Applying Lemma 5 to $\widehat{Q}$ returned by KQLearn, it follows that with probability $1-\delta$,

$$
\left\|V^{\pi}-V^{*}\right\|_{\infty} \leq \frac{2 \gamma \beta(\delta)}{(1-\gamma)^{2}} \sqrt{\frac{2 \Gamma_{K, \lambda}(J)}{\log \left(1+1 / \lambda^{2}\right) J}}+\frac{2 \gamma^{L}}{(1-\gamma)^{2}}
$$

That completes the proof.

## B. 2 Proof of Theorem 2

Theorem 2 is a consequence of Theorem 1 and the kernel specific bounds on $\Gamma_{K, \lambda}$. Several works have established bounds on $\Gamma_{K, \lambda}$ for various kernels (Srinivas et al., 2010, Vakili et al., 2021c|b). We use the reuslt in Vakili et al. (2021c) for kernels with polynomial and exponential eigendecay.

Polynomial Eigendecay: Consider a kernel with $\left(C_{p}, \beta_{p}\right)$ polynomial eigendecay. We have the following bound on $\Gamma_{K, \lambda}$ (Vakili et al. 2021c). For all $J \in \mathbb{N}$,

$$
\begin{equation*}
\Gamma_{K, \lambda}(J)=\mathcal{O}\left(J^{\frac{1}{\beta_{p}}} \log ^{1-\frac{1}{\beta_{p}}}(J)\right) \tag{28}
\end{equation*}
$$

We replace this bound in the error in the value function of $\pi$ obtained by KQLearn, in Theorem 1

$$
\left\|V^{\pi}-V^{*}\right\|_{\infty} \leq \frac{2 \gamma \beta(\delta)}{(1-\gamma)^{2}} \sqrt{\frac{2 \Gamma_{K, \lambda}(J)}{\log \left(1+\frac{1}{\lambda^{2}}\right) J}}+\frac{2 \gamma^{L}}{(1-\gamma)^{2}}
$$

We then obtain

$$
\left\|V^{\pi}-V^{*}\right\|_{\infty}=\mathcal{O}\left(\frac{\gamma}{(1-\gamma)^{3}}\left(c+\frac{1}{\lambda} \sqrt{d \log \left(\frac{J c}{(1-\gamma) \delta}\right)}\right) \sqrt{J^{\frac{1}{\beta_{p}}-1} \log ^{1-\frac{1}{\beta_{p}}}(J)}\right)+\mathcal{O}\left(\frac{\gamma^{L}}{(1-\gamma)^{2}}\right)
$$

We choose $J$ and $L$ large enough so that each term on the right hand side is bounded by $\epsilon / 2$. The choices of

$$
\begin{align*}
J & =\mathcal{O}\left(\left(\frac{1}{\epsilon}\right)^{\frac{2 \beta_{p}}{\beta_{p}-1}} \frac{\gamma^{\frac{2 \beta_{p}}{\beta_{p}-1}}}{(1-\gamma)^{\frac{6 \beta_{p}}{\beta_{p}-1}}}\left(c+\frac{1}{\lambda} \sqrt{d \log \left(\frac{c}{(1-\gamma) \delta}\right)}\right)^{\frac{2 \beta_{p}}{\beta_{p}-1}} \log ^{\frac{\beta_{p}}{\beta_{p}-1}}\left(\frac{1}{\epsilon(1-\gamma)}\right)\right) \\
L & =\mathcal{O}\left(\frac{\log \left(\frac{1}{\epsilon}\right)+\log \left(\frac{1}{1-\gamma}\right)}{(1-\gamma)}\right) \tag{29}
\end{align*}
$$

with proper constants ensures $\left\|V^{\pi}-V^{*}\right\|_{\infty} \leq \epsilon$. The expression can be simplified as $J=\tilde{\mathcal{O}}\left(\frac{\left(\log \left(\frac{1}{\delta}\right)\right)^{\frac{\beta_{p}}{\beta_{p}-1}}}{\epsilon^{\frac{2 \beta_{p}-1}{\beta_{p}-1}}(1-\gamma)^{\frac{6 \beta_{p}}{\beta_{p}-1}}}\right)$ and $L=\tilde{\mathcal{O}}\left(\frac{1}{1-\gamma}\right)$, omitting the logarithmic and constant terms. That leads to

$$
\begin{equation*}
N=\tilde{\mathcal{O}}\left(\frac{\left(\log \left(\frac{1}{\delta}\right)\right)^{\frac{\beta_{p}}{\beta_{p}-1}}}{\epsilon^{\frac{2 \beta_{p}}{\beta_{p}-1}}(1-\gamma)^{\frac{7 \beta_{p}-1}{\beta_{p}-1}}}\right) \tag{30}
\end{equation*}
$$

Exponential Eigendecay: Consider a kernel with $\left(C_{e_{1}}, C_{e_{2}}, \beta_{e}\right)$ polynomial eigendecay. We have the following bound on $\Gamma_{K, \lambda}$. For all $J \in \mathbb{N}$,

$$
\begin{equation*}
\Gamma_{K, \lambda}(J)=\mathcal{O}\left(\log ^{1+\frac{1}{\beta_{e}}}(J)\right) \tag{31}
\end{equation*}
$$

We replace this bound in the error in the value function of $\pi$ obtained by KQLearn, in Theorem 1 , and obtain

$$
\left\|V^{\pi}-V^{*}\right\|_{\infty}=\mathcal{O}\left(\frac{\gamma}{(1-\gamma)^{3}}\left(c+\frac{1}{\lambda} \sqrt{d \log \left(\frac{J c}{(1-\gamma) \delta}\right)}\right) \sqrt{\frac{\log ^{1+\frac{1}{\beta_{e}}}(J)}{J}}\right)+\mathcal{O}\left(\frac{\gamma^{L-1}}{(1-\gamma)^{2}}\right)
$$

We choose $J$ and $L$ large enough so that each term on the right hand side is bounded by $\epsilon / 2$. The choices of

$$
\begin{align*}
J & =\mathcal{O}\left(\left(\frac{1}{\epsilon}\right)^{2} \frac{\gamma^{2}}{(1-\gamma)^{6}}\left(c+\frac{1}{\lambda} \sqrt{d \log \left(\frac{c}{(1-\gamma) \delta}\right)}\right)^{2} \log ^{2+\frac{1}{\beta_{e}}}\left(\frac{1}{\epsilon(1-\gamma)}\right)\right) \\
L & =\mathcal{O}\left(\frac{\log \left(\frac{1}{\epsilon}\right)+\log \left(\frac{1}{1-\gamma}\right)}{(1-\gamma)}\right) \tag{32}
\end{align*}
$$

with proper constants ensures $\left\|V^{\pi}-V^{*}\right\|_{\infty} \leq \epsilon$. The expression can be simplified as $J=\tilde{\mathcal{O}}\left(\frac{\log \left(\frac{1}{\delta}\right)}{\epsilon^{2}(1-\gamma)^{6}}\right)$ and $L=$ $\tilde{\mathcal{O}}\left(\frac{1}{1-\gamma}\right)$, omitting the logarithmic and constant terms. That leads to

$$
\begin{equation*}
N=\tilde{\mathcal{O}}\left(\frac{\log \left(\frac{1}{\delta}\right)}{\epsilon^{2}(1-\gamma)^{7}}\right) \tag{33}
\end{equation*}
$$

## C PROOF OF LEMMAS

In this section, we provide the proof of auxiliary lemmas.

## C. 1 Proof of Lemma 3 [RKHS norm of $P V]$

We have

$$
\begin{align*}
\|P V\|_{\mathcal{H}_{K}} & =\left\|\int_{s^{\prime} \in \mathcal{S}} P\left(s^{\prime} \mid \cdot, \cdot\right) V\left(s^{\prime}\right) d s^{\prime}\right\|_{\mathcal{H}_{K}} \\
& \leq \int_{s^{\prime} \in \mathcal{S}}\left\|P\left(s^{\prime} \mid \cdot, \cdot\right) V\left(s^{\prime}\right)\right\|_{\mathcal{H}_{K}} d s^{\prime} \\
& =\int_{s^{\prime} \in \mathcal{S}}\left\|P\left(s^{\prime} \mid \cdot, \cdot\right)\right\|_{\mathcal{H}_{K}} V\left(s^{\prime}\right) d s^{\prime} \\
& \leq \int_{s^{\prime} \in \mathcal{S}} V\left(s^{\prime}\right) d s^{\prime} . \tag{34}
\end{align*}
$$

where the last inequality holds by Assumption 2 . We note that $\int_{s^{\prime} \in \mathcal{S}} V\left(s^{\prime}\right) d s^{\prime} \leq \frac{1}{1-\gamma} \int_{s^{\prime}} d s^{\prime} \leq \frac{c}{1-\gamma}$ where $c$ is the volume of $\mathcal{S}$.

## C. 2 Proof of Lemma4[Sub-Gaussianity of PV]

The first part, $\mathbb{E}[[\widehat{P} V](s, a)]=P V(s, a)$, follows from the definition of $\widehat{P} V$. For the second part note that $\widehat{P} V$ is a random variable with a bounded support in $\left[0, \frac{1}{1-\gamma}\right]$ by definition. Hoeffding lemma states that: let $X$ be any random variable such that $0 \leq X \leq a$, then for any $\zeta \in \mathbb{R}, \mathbb{E}[\exp (\zeta(X-\mathbb{E}[X]))] \leq \exp \left(\frac{\zeta^{2} a^{2}}{8}\right)$. Applying Hoeffding lemma, we can see that $\widehat{P} V$ is sub-Gaussian with parameter $\frac{1}{2(1-\gamma)}$.

## C. 3 Proof of Lemma 5 [Error in value function based on the error in proxy Q-function]

The proof follows similar steps as the proof of Lemma 1.11 in Agarwal et al. (2019) which considered finite state-actions. First, recall the following definitions from Section 2 ,

$$
Q^{\pi_{\tilde{Q}}}(s, a)=r(s, a)+\gamma\left[P V^{\pi_{\tilde{Q}}}\right](s, a) \quad \text { and } \quad Q^{*}(s, a)=\max _{\pi} Q(s, \pi(s))
$$

Fix state $s \in \mathcal{S}$ and let $\tilde{a}_{s}=\pi_{\tilde{Q}}(s)$. We have

$$
\begin{aligned}
V^{*}(s)-V^{\pi_{\tilde{Q}}}(s)= & Q^{*}\left(s, \pi^{*}(s)\right)-Q^{\pi_{\tilde{Q}}}\left(s, \tilde{a}_{s}\right) \\
= & Q^{*}\left(s, \pi^{*}(s)\right)-Q^{*}\left(s, \tilde{a}_{s}\right)+Q^{*}\left(s, \tilde{a}_{s}\right)-Q^{\pi_{\tilde{Q}}}\left(s, \tilde{a}_{s}\right) \\
= & Q^{*}\left(s, \pi^{*}(s)\right)-Q^{*}\left(s, \tilde{a}_{s}\right)+\gamma\left[P V^{*}\right]\left(s, \tilde{a}_{s}\right)-\gamma\left[P V^{\pi_{\tilde{Q}}}\right]\left(s, \tilde{a}_{s}\right) \\
= & Q^{*}\left(s, \pi^{*}(s)\right)-Q^{*}\left(s, \tilde{a}_{s}\right)+\gamma \mathbb{E}_{s^{\prime} \sim P\left(\cdot \mid s, \tilde{a}_{s}\right)}\left[V^{*}\left(s^{\prime}\right)-V^{\pi_{\tilde{Q}}}\left(s^{\prime}\right)\right] \\
\leq & Q^{*}\left(s, \pi^{*}(s)\right)-\tilde{Q}\left(s, \pi^{*}(s)\right)+\tilde{Q}\left(s, \tilde{a}_{s}\right)-Q^{*}\left(s, \tilde{a}_{s}\right) \\
& +\gamma \mathbb{E}_{s^{\prime} \sim P\left(s, \tilde{a}_{s}\right)}\left[V^{*}\left(s^{\prime}\right)-V^{\pi_{\tilde{Q}}}\left(s^{\prime}\right)\right] \\
\leq & 2\left\|\tilde{Q}-Q^{*}\right\|_{\infty}+\gamma\left\|V^{*}-V^{\pi_{\tilde{Q}}}\right\|_{\infty},
\end{aligned}
$$

where the first inequality uses $\tilde{Q}\left(s, \pi^{*}(s)\right) \leq \tilde{Q}\left(s, \tilde{a}_{s}\right)$ by definition of $\tilde{a}_{s}$. We thus have

$$
\begin{equation*}
\left\|V^{*}-V^{\pi_{\tilde{Q}}}\right\|_{\infty} \leq 2\left\|\tilde{Q}-Q^{*}\right\|_{\infty}+\gamma\left\|V^{*}-V^{\pi_{\tilde{Q}}}\right\|_{\infty} \tag{35}
\end{equation*}
$$

Rearranging this inequality, we arrive at the lemma

$$
\begin{equation*}
\left\|V^{*}-V^{\pi_{\tilde{Q}}}\right\|_{\infty} \leq \frac{2}{1-\gamma}\left\|\tilde{Q}-Q^{*}\right\|_{\infty} \tag{36}
\end{equation*}
$$


[^0]:    ${ }^{1}$ The notations $\mathcal{O}$ and $\tilde{\mathcal{O}}$ are used to denote the mathematical order, and that up to hiding logarithmic factors, respectively.

[^1]:    ${ }^{2}$ We intentionally do note use the standard term transition kernel for $P$, to avoid confusion with the term kernel in kernel-based learning.

[^2]:    ${ }^{3}$ A random variable $X$ is said to be sub-Gaussian with parameter $R$ if its moment generating function is bounded by that of a zero mean Gaussian with variance $R^{2}$.

[^3]:    ${ }^{4}$ For the simplicity of presentation, we assume $N=J L$. When $J$ does not divide $N$, we can ignore the samples in the last round.

