
Mitigating Covariate Shift in Imitation Learning via Offline Data Without Great Coverage

Jonathan D. Chang*

Department of Computer Science
Cornell University
jdc396@cornell.edu

Masatoshi Uehara*

Department of Computer Science
Cornell University
mu223@cornell.edu

Dhruv Sreenivas

Department of Computer Science
Cornell University
ds844@cornell.edu

Rahul Kidambi†

Amazon Search & AI
rk773@cornell.edu

Wen Sun

Department of Computer Science
Cornell University
ws455@cornell.edu

Abstract

This paper studies offline Imitation Learning (IL) where an agent learns to imitate an expert demonstrator without additional online environment interactions. Instead, the learner is presented with a static offline dataset of state-action-next state transition triples from a potentially less proficient behavior policy. We introduce Model-based IL from Offline data (MILO): an algorithmic framework that utilizes the static dataset to solve the offline IL problem efficiently both in theory and in practice. In theory, even if the behavior policy is highly sub-optimal compared to the expert, we show that as long as the data from the behavior policy provides sufficient coverage on the expert state-action traces (and with no necessity for a global coverage over the entire state-action space), MILO can provably combat the covariate shift issue in IL. Complementing our theory results, we also demonstrate that a practical implementation of our approach mitigates covariate shift on benchmark MuJoCo continuous control tasks. We demonstrate that with behavior policies whose performances are less than half of that of the expert, MILO still successfully imitates with an extremely low number of expert state-action pairs while traditional offline IL methods such as behavior cloning (BC) fail completely. Source code is provided at <https://github.com/jdchang1/milo>.

1 Introduction

Covariate shift is a core issue in Imitation Learning (IL). Traditional IL methods like behavior cloning (BC) [49], while simple, suffer from covariate shift, learning a policy that can make arbitrary mistakes in parts of the state space not covered by the expert dataset. This leads to compounding errors in the agent’s performance [57], hurting the generalization capabilities in practice.

Prior works have presented several means to combat this phenomenon in IL. One line of thought utilizes an *interactive expert*, i.e. an expert that can be queried at an arbitrary state encountered during the training procedure. Interactive IL algorithms such as DAgger [59], LOLS [15], DART [40], and AggreVaTe(D) [58; 66] utilize a reduction to no-regret online learning and demonstrate that under certain conditions, they can successfully learn a policy that imitates the expert. These interactive IL algorithms, however, cannot provably avoid covariate shift if the expert is not *recoverable*. That is,

*Equal contribution

†Work done outside Amazon




Figure 1: (Left) Frames at timesteps 200, 400, 600, 800, and 1000 for Humanoid-v2 from policies trained with BC on 100 state-action pairs from the expert (blue), BC on 1M offline samples plus 100 expert samples (yellow), and our algorithm MILO (red). The expert has a performance of 3248 and the behavior policy used to collect the offline dataset has performance of 1505 ± 473 ($\approx 46\%$ of the expert’s). (Right) Expert performance normalized scores averaged across 5 seeds.

$A^{\pi_e}(s, a) = \Theta(H)$ where π_e is the expert, A^π is the usual (dis)advantage function,³ and H is the planning horizon [52; 4, Chapter 15]. A second line of work that avoids covariate shift utilizes either a known transition dynamics model [2; 86] or uses real world interactions [27; 10; 65; 38; 56; 33]. Prior works have shown that with known transition dynamics or real world interactions, agents can provably avoid covariate shift in both tabular and general MDPs [4; 52] even without a recoverable expert. While these results offer strong theoretical guarantees and empirical performance, online interactions are often costly and prohibitive for real world applications where active trial-and-error exploration in the environment could be unsafe or impossible. A third perspective towards addressing this issue is to assume that the expert visits *the entire state space* [62], where the expert effectively informs the learner what actions to take in every state. Unfortunately, such a full coverage expert distribution might be rare and holds only for special MDPs and expert policies (for e.g. an expert that induces ergodicity in the MDP).

In this work, we consider a new perspective towards handling the covariate shift issue in IL. In particular, we investigate a pure offline learning setting where the learner has access to neither the expert nor the environment for additional interactions. The learner, instead, has access to a small pre-collected dataset of state-action pairs sampled from the expert and a large batch offline dataset of state-action-next state transition triples sampled from a behavior policy that could be highly sub-optimal (see Figure 1 where BC on the offline data results in a low-quality policy). Unlike prior works that require online interactions, our proposed method, MILO performs high fidelity imitation in an offline, data-driven manner. Moreover, different from interactive IL, we do not require the expert to be present during learning, significantly relieving the expert’s burden. Finally, in contrast to the prior work [62] that assumes the expert distribution covers the entire state-action space (i.e., $\max_{\pi} \max_{s,a} d^{\pi}(s, a) / d^{\pi_e}(s, a) < \infty$ where d^{π} denotes the state-action distribution of policy π), we require the offline dataset to provide *partial coverage*, i.e., it only needs to cover the expert’s state-actions (i.e., $\max_{s,a} d^{\pi_e}(s, a) / \rho(s, a) < \infty$ where ρ is the offline distribution of some behavior policy).⁴

In summary, we list our main contributions below:

1. We propose Model based Imitation Learning from Offline data, MILO: a model-based framework that leverages offline batch data with only partial coverage (see Section 4.1 for definition) to overcome covariate shift in IL.

³in this work, we use cost instead of reward, thus we call A^π the disadvantage function.

⁴In our analysis, we refine the density ratio $d^{\pi_e}(s, a) / \rho(s, a)$ via the concept of relative conditional number which allows us to extend it to large MDPs where the ratio is infinite but the relative condition number is finite.

2. Our analysis is modular and covers common models such as discrete MDPs and linear models. Notably, our new result on non-parametric models (e.g. Gaussian Processes) with relative condition number is new even considering all existing results in offline RL (see Remark 4.7,11).
3. The practical instantiation of our general framework leverages neural network model ensembles, and demonstrates its efficacy on benchmark MuJoCo continuous control problems. Specifically, even under low-quality behavior policies, our approach can successfully imitate using an extremely small number of expert samples while algorithms like BC completely fail (Figure 1).

1.1 Related work

Imitation Learning As summarized above, avoiding covariate shift in IL is an important topic. Another relevant line of research is IL algorithms that use *offline* or *off-policy* learning. ValueDICE [38] presents a principled way to leverage off-policy data for IL. In theory, the techniques from ValueDICE (and more broadly, DICE [46; 84]) require the data provided to the agent to have global coverage. Moreover in practice, ValueDICE uses online interaction and maintains an increasing replay buffer which may eventually provide global coverage. Instead, we aim to study offline IL without any online interactions and are interested in the setting where offline data does not have global coverage. Another line of work [29; 14] studies IL in an offline setting by only using the expert dataset. In contrast to these works, our goal is to study the use of an additional offline dataset collected from a behavior policy to mitigate covariate shift, as information theoretically any algorithm that relies solely on expert data will still suffer from covariate shift in the worst case [52].

Similar in setting, Cascaded Supervised IRL (CSI) [35] performs imitation offline with both an expert dataset and a static offline dataset by first fitting a reward function that is then used in Least Squares Policy Iteration. CSI, however, requires the expert data to have global coverage and does not mitigate covariate shift with partial coverage like MLO does. Finally, Variational Model-Based Adversarial IL (V-MAIL) [50] learns a dynamics model from a static offline dataset and performs offline imitation within the model. V-MAIL, however, studies zero-shot IL where the static offline data is samples from a variety of source tasks, and the expert dataset is samples from the transfer target task. In contrast, MLO investigates avoiding covariate shift in IL with an offline dataset collected by a potentially suboptimal policy from the same task.

Offline RL In offline RL, algorithms such as FQI [22] have finite-sample error guarantees under the global coverage [45; 5]. Recently, many algorithms to tackle this problem have been proposed from both model-free [76; 68; 43; 26; 23; 39] and model-based perspectives [82; 34; 44] with some pessimism ideas. The idea of pessimism features in offline RL with an eye to penalize the learner from visiting unknown regions of the state-action space [54; 31; 79; 12]. We utilize pessimism within the IL context where, unlike RL, the learner does not have access to an underlying reward signal. Our work expands prior theoretical results by (a) formalizing the partial coverage condition using a notion of a relative condition number, and (b) offering distribution-dependent results when working with non-parametric models including gaussian processes. See Appendix C for a detailed literature review.

2 Setting

We consider an episodic finite-horizon Markov Decision Process (MDP), $\mathcal{M} = \{\mathcal{S}, \mathcal{A}, P, H, c, d_0\}$, where \mathcal{S} is the state space, \mathcal{A} is the action space, $P : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ is the MDP’s transition, H is the horizon, d_0 is an initial distribution, and $c : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is the cost function. A policy $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ maps from state to distribution over actions. We denote $d_P^\pi \in \Delta(\mathcal{S} \times \mathcal{A})$ as the average state-action distribution of π under transition kernel P , that is, $d_P^\pi = 1/H \sum_{t=1}^H d_{P,t}^\pi$, where $d_{P,t}^\pi \in \Delta(\mathcal{S} \times \mathcal{A})$ is the distribution of $(s^{(t)}, a^{(t)})$ under π at t . Given a cost function $f : \mathcal{S} \times \mathcal{A} \mapsto [0, 1]$, $V_{P,f}^\pi$ denotes the expected cumulative cost of π under the transition kernel P and cost function f . Following a standard IL setting, the ground truth cost function c is unknown. Instead, we have the demonstrations by the expert specified by $\pi_e : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ (potentially stochastic and not necessarily optimal). Concretely, we have an expert dataset in the form of i.i.d tuples $\mathcal{D}_e = \{s_i, a_i\}_{i=1}^{n_e}$ sampled from distribution $d_{P_e}^{\pi_e}$.

In our setting, we also have an offline static dataset consisting of i.i.d tuples $\mathcal{D}_o = \{s_i, a_i, s'_i\}_{i=1}^{n_o}$ s.t. $(s, a) \sim \rho(s, a), s' \sim P(s, a)$, where $\rho \in \Delta(\mathcal{S} \times \mathcal{A})$ is an *offline distribution* resulting from some behavior policies. Note behavior policy could be a much worse policy than the expert π_e . Our goal is

Algorithm 1 Framework for model-based Imitation Learning with offline data (MILO)

- 1: **Require:** IPM class \mathcal{F} , model class \mathcal{P} , policy class Π , datasets $\mathcal{D}_e = \{s, a\}$, $\mathcal{D}_o := \{s, a, s'\}$
- 2: **Train Dynamics Model and Bonus:** $\hat{P} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ and $b : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^+$ on offline data \mathcal{D}_o
- 3: **Pessimistic model-based min-max IL:** with \hat{P} , b , \mathcal{D}_e , obtain $\hat{\pi}_{\text{IL}}$ by solving the following:

$$\hat{\pi}_{\text{IL}} = \operatorname{argmin}_{\pi \in \Pi} \max_{f \in \mathcal{F}} \left[\mathbb{E}_{(s,a) \sim d_{\hat{P}}^{\pi}} [f(s, a) + b(s, a)] - \mathbb{E}_{(s,a) \sim \mathcal{D}_e} [f(s, a)] \right] \quad (1)$$

to only leverage $(\mathcal{D}_e + \mathcal{D}_o)$ to learn a policy π that performs as well as π_e with regard to optimizing the ground truth cost c . More specifically, our goal is to utilize the offline static data \mathcal{D}_o to combat covariate shift and learn a policy that can significantly outperform traditional offline IL methods such as Behavior cloning (BC), *without any interaction with the real world or expert*.

Function classes We introduce function approximation. Since we do not know the true cost function c and transition kernel P , we introduce a cost function class $\mathcal{F} \subset \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ and a transition model class $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$. We also need a policy class Π . For the analysis, we assume realizability:

Assumption 1. $c \in \mathcal{F}, P \in \mathcal{P}, \pi_e \in \Pi$.

We use Integral Probability Metric (IPM) as a distribution distance measure, i.e., given two distributions ρ_1 and ρ_2 , IPM with \mathcal{F} is defined as $\max_{f \in \mathcal{F}} [\mathbb{E}_{(s,a) \sim \rho_1} [f(s, a)] - \mathbb{E}_{(s,a) \sim \rho_2} [f(s, a)]]$.

3 Algorithm

The core idea of MILO is to imitate the expert by optimizing an IPM distance between the agent and the expert with a penalty term for pessimism over the policy class. MILO consists of three steps:

1. **Model learning:** fit a model \hat{P} from the offline data \mathcal{D}_o to learn P ,
2. **Pessimistic penalty design:** construct penalty function $b(s, a)$ such that there is a high penalty on state-action pairs that are not covered by the offline data distribution ρ .
3. **Offline min-max model-based policy optimization:** optimize Eq. (1)

Algorithm 1 provides the details of MILO. We explain each component in detail as follows.

Model learning and Penalty: Our framework assumes we can learn a calibrated model (\hat{P}, σ) from the dataset \mathcal{D}_o , in the sense that for any s, a , we have: $\left\| \hat{P}(\cdot|s, a) - P(\cdot|s, a) \right\|_1 \leq \min\{2, \sigma(s, a)\}$. Such model training is possible in many settings including classic discrete MDPs, linear models (KNR [32]), and non-parametric models such as GP. In practice, it is also common to train a model ensemble based on the idea of bootstrapping and then use the model-disagreement to approximate σ . With such a calibrated model, the penalty will simply be $b(s, a) = O(H\sigma(s, a))$. We will formalize this model learning assumption in Section 4. We give several examples below.

For any discrete MDP, we use the empirical distribution, i.e., $\hat{P}(s'|s, a) = N(s', s, a)/(N(s, a) + \lambda)$, where $N(s, a)$ is the number of (s, a) in \mathcal{D}_o , and $N(s', s, a)$ is the number of (s, a, s') in \mathcal{D}_o , and $\lambda \in \mathbb{R}^+$. In this case, we can set $\sigma(s, a) = \tilde{O}\left(\sqrt{|\mathcal{S}|/N(s, a)}\right)$. See example 1 for more details.

For continuous Kernelized Nonlinear Regulator (KNR [32]) model where the ground truth transition $P(s'|s, a)$ is defined as $s' = W^* \phi(s, a) + \epsilon$, $\epsilon \sim \mathcal{N}(0, \Sigma)$, with ϕ being a (nonlinear) feature mapping, we can learn \hat{P} by classic Ridge regression on offline dataset \mathcal{D}_o . Here we can set $\sigma(s, a) = \tilde{O}\left(\beta \sqrt{\phi(s, a)^\top \Sigma_{n_o}^{-1} \phi(s, a)}\right)$ for some $\beta \in \mathbb{R}^+$, where Σ_o is the data covariance matrix $\Sigma_{n_o} := \sum_{i=1}^{n_o} \phi(s_i, a_i) \phi(s_i, a_i)^\top + \lambda \mathbf{I}$. See example 2 for more details.

For non-parametric nonlinear model such as Gaussian Process (GP), under the assumption that P is in the form of $s' = g^*(s, a) + \epsilon$, $\epsilon \sim \mathcal{N}(0, \Sigma)$ (here $\mathcal{S} \subset \mathbb{R}^{d_s}$), we can simply represent \hat{P} using GP posteriors induced by \mathcal{D}_o , i.e., letting GP posterior be $GP(\hat{g}, k_{n_o})$, we have $\hat{P}(s'|s, a)$ being

represented as $s' = \hat{g}(s, a) + \epsilon$. Then, we can set $\sigma(s, a) = \tilde{O}(\beta k_{n_o}((s, a), (s, a)))$ with some parameter $\beta \in \mathbb{R}^+$ (see example 3 for more details). GP is a powerful model and has been being widely used in robotics problems, see [36; 18; 8; 71; 25] for examples.

In practice, we can also use a model ensemble of neural networks with the maximum disagreement between models as σ . This has been widely used in practice (e.g., [47; 6; 48]). We leave the details to Section 5 where we instantiate a practical version of MLO, and the experiment section. As we can see from the examples mentioned above, in general, the penalty $b(s, a) = O(H\sigma(s, a))$ is designed such that it has a high value in state-action space that is not covered well by the offline data \mathcal{D}_o , and has a low value in space that is covered by \mathcal{D}_o . Adding such a penalty automatically forces our policy to stay away from these regions where \hat{P} is not accurate. On the other hand, for regions where ρ has good coverage (thus \hat{P} is accurate), we force π to stay close to π^e .

Pessimistic model-based min-max IL: Note Eq. 1 is purely computational, i.e., we do not need any real world samples. To solve such min-max objective, we can iteratively (1) perform the best response on the max player, i.e., compute the argmax discriminator f given the current π , and (2) perform incremental update on the min player, e.g., use policy gradient (PG) methods (e.g. TRPO) inside the learned model \hat{P} with cost function $f(s, a) + b(s, a)$. We again leave the details to Section 5.

3.1 Specialization to offline RL

In RL, the cost function c is given. The goal is to obtain $\pi^* = \operatorname{argmax}_{\pi \in \Pi} V_{P, c}^\pi$. The pessimistic policy optimization procedure [81; 31] is $\hat{\pi}_{\text{RL}} = \operatorname{argmin}_{\pi \in \Pi} \mathbb{E}_{(s, a) \sim d_\pi} [c(s, a) + b(s, a)]$. While this is not our main contribution, we will show a byproduct of our result is a novel non-parametric analysis for offline RL which does not assume ρ has global coverage (see Remarks 4, 7, 11).

4 Analysis

Our algorithm depends on the model \hat{P} estimated from the offline data. We provide a unified analysis assuming that \hat{P} is calibrated in that its confidence interval is provided. Specifically, we assume:

Assumption 2. *With probability $1 - \delta$, the estimate model \hat{P} satisfies the following: $\|\hat{P}(\cdot|s, a) - P(\cdot|s, a)\|_1 \leq \min(\sigma(s, a), 2) \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}$. We set the penalty as $b(s, a) = H \min(\sigma(s, a), 2)$.*

We give the following three examples. For details, refer to appendix A.

Example 1 (Discrete MDPs). *Set uncertainty measure $\sigma(s, a) = \sqrt{\frac{|S| \log 2 + \log(2|S||\mathcal{A}|/\delta)}{2\{N(s, a) + \lambda\}}} + \frac{\lambda}{N(s, a) + \lambda}$.*

Example 2 (KNRs). *In KNRs, the ground truth model is $s' = W^* \phi(s, a) + \epsilon$, $\epsilon \sim \mathcal{N}(0, \zeta^2 \mathbf{I})$, where $s \in \mathbb{R}^{d_S}$, $a \in \mathbb{R}^{d_A}$, $\phi: \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}^d$ is some known state-action feature mapping. The estimator is*

$$\hat{g}(\cdot) = \hat{W} \phi(\cdot), \quad \hat{W} = \operatorname{argmin}_{W \in \mathbb{R}^{d_S \times d_A}} 1/n_o \sum_{(s, a) \in \mathcal{D}_o} [\|W \phi(s, a) - s'\|_2^2] + \lambda \|W\|_F^2,$$

where $\|\cdot\|_F$ is a frobenius norm. We set the uncertainty measure $\sigma(s, a)$:

$$\sigma(s, a) = (1/\zeta) \beta_{n_o} \sqrt{\phi^\top(s, a) \Sigma_{n_o}^{-1} \phi(s, a)}, \quad \Sigma_{n_o} = \sum_{i=1}^{n_o} \phi(s_i, a_i) \phi^\top(s_i, a_i) + \lambda \mathbf{I}$$

with $\beta_{n_o} = \{2\lambda \|W^*\|_2^2 + 8\zeta^2 [d_S \log(5) + \log(1/\delta) + \bar{\mathcal{I}}_{n_o}]\}^{1/2}$, where $\bar{\mathcal{I}}_{n_o} = \log(\det(\Sigma_{n_o}/\lambda \mathbf{I}))$.

Example 3 (GPs). *In GPs, the ground truth model is defined as $s' = g^*(s, a) + \epsilon$, $\epsilon \sim \mathcal{N}(0, \zeta^2 \mathbf{I})$ where g^* belongs to an RKHS \mathcal{H}_k with a kernel $k(\cdot, \cdot)$. Denote $x := (s, a)$, we have GP posterior as*

$$\hat{g}(\cdot) = S(\mathbf{K}_{n_o} + \zeta^2 \mathbf{I})^{-1} \bar{k}_{n_o}(\cdot), \quad S = [s'_1, \dots, s'_{n_o}] \in \mathbb{R}^{d_S \times n_o}, \quad \bar{k}_{n_o}(x) = [k(x_1, x), \dots, k(x_{n_o}, x)]^\top, \\ \{\mathbf{K}_{n_o}\}_{i,j} = k(x_i, x_j) \quad (1 \leq i \leq n_o, 1 \leq j \leq n_o), \quad k_{n_o}(x, x') = k(x, x') - \bar{k}_{n_o}(x)^\top (\mathbf{K}_{n_o} + \zeta^2 \mathbf{I})^{-1} \bar{k}_{n_o}(x'),$$

with $\sigma(\cdot) = \beta_{n_o} k_{n_o}(\cdot, \cdot)/\zeta$, $\beta_{n_o} = O((d_S \log^3(d_S n_o/\delta) \mathcal{I}_{n_o})^{1/2})$, $\mathcal{I}_{n_o} = \log(\det(\mathbf{I} + \zeta^{-2} \mathbf{K}_{n_o}))$.

General results We show our general error bound results. For the proof, refer to appendix B. For analytical simplicity, we assume $|\mathcal{F}|$ is finite (but the bound only depends on $\ln(|\mathcal{F}|)$)⁵.

⁵When $|\mathcal{F}|$ is infinite, we can show that the resulting error bound scales w.r.t its metric entropy.

Theorem 3 (Bound of MILD). *Suppose assumptions 1,2. Then, with probability $1 - 2\delta$,*

$$V_{P,c}^{\hat{\pi}_{\text{IL}}} - V_{P,c}^{\pi_e} \leq \text{Err}_o + \text{Err}_e, \text{Err}_o = 8H^2 \mathbb{E}_{(s,a) \sim d_P^{\pi_e}} [\min(\sigma(s,a), 1)], \text{Err}_e = 2H \sqrt{\frac{\log(2|\mathcal{F}|/\delta)}{2n_e}}.$$

We will show through a set of examples where $\mathbb{E}_{(s,a) \sim d_P^{\pi_e}} [\min(\sigma(s,a), 1)]$ shrinks to zero as $n_o \rightarrow \infty$ under the partial coverage, i.e., when ρ covers $d_P^{\pi_e}$. Asymptotically, Err_e will dominate the bound. Note that Err_e has two components, a linear H and a term that corresponds to the statistical error related to expert samples and function class complexity. Comparing to BC, which has a rate $O(H^2 \sqrt{\log(|\Pi|)/n_e})$ [4, Chapter 14], we see that the horizon dependence is improved.

Before going to each analysis of Err_o , we highlight two important points in our analysis. First, our bound requires only the partial coverage, i.e., it depends on π_e -concentrability coefficient which measures the discrepancy between the offline data and expert data. This is the first work deriving the bound with π_e -concentrability coefficient in IL with offline data. Second, our analysis covers non-parametric models. This is a significant contribution as previous pessimistic offline RL finite-sample error results have been limited to the finite-dimensional linear models or discrete MDPs [31; 54].

Remark 4 (Implications on offline RL). *As in theorem 3, we have $V_{P,c}^{\hat{\pi}_{\text{RL}}} - V_{P,c}^{\pi^*} = O(H^2 \mathbb{E}_{(s,a) \sim d_P^{\pi^*}} [\sigma(s,a)])$ (appendix B). Note similar results have been obtained in [82; 34]. Since this term is Err_o by just replacing π_e with π^* , this offline RL result is a by-product of our analysis.*

4.1 Analysis: Discrete MDPs

We start from discrete MDP as a warm up. Denote $C^{\pi_e} = \max_{(s,a)} d_P^{\pi_e}(s,a)/\rho(s,a)$.

Theorem 5. *Suppose $\lambda = \Omega(1)$ and the partial coverage $C^{\pi_e} < \infty$. With probability $1 - \delta$,*

$$\text{Err}_o \leq c_1 H^2 \left(\sqrt{\frac{C^{\pi_e} |\mathcal{S}|^2 |\mathcal{A}|}{n_o}} + \frac{C^{\pi_e} |\mathcal{S}| |\mathcal{A}|}{n_o} \right) \cdot \log(|\mathcal{S}| |\mathcal{A}| c_2 / \delta),$$

where c_1, c_2 are universal constants.

The error does not depend on $\sup_{\pi \in \Pi} C^\pi$ or $\bar{C} = \sup_{\pi \in \Pi} \max_{(s,a)} d_P^\pi(s,a)/d_P^{\pi_e}(s,a)$. We only require the *partial coverage* $C^{\pi_e} < \infty$, which is much weaker than $\sup_{\pi \in \Pi} C^\pi < \infty$ (ρ has global coverage) and $\bar{C} < \infty$ ($d_P^{\pi_e}$ has global coverage [62]). When C^{π_e} is small and n_o is large enough, $\text{Err}_e = O\left(H \sqrt{|\mathcal{S}| |\mathcal{A}| / n_e}\right)$ dominates Err_o in theorem 3. Then, the error is linear in horizon H .

4.2 Analysis: KNRs and GPs for Continuous MDPs

Now we move to continuous state-action MDPs. In continuous MDPs, assuming the boundedness of density ratio C^{π_e} is still a strong assumption. As we dive into the KNR and the nonparametric GP model, we will replace the density ratio with a more refined concept *relative condition number*.

KNRs Let $\Sigma_\rho = \mathbb{E}_{(s,a) \sim \rho} [\phi(s,a)\phi(s,a)^\top]$ and $\Sigma_{\pi_e} = \mathbb{E}_{(s,a) \sim d_P^{\pi_e}} [\phi(s,a)\phi(s,a)^\top]$. We define the *relative condition number* as $C^{\pi_e} = \sup_{x \in \mathbb{R}^d} \left(\frac{x^\top \Sigma_{\pi_e} x}{x^\top \Sigma_\rho x} \right)$. Even when density ratio is infinite, this number could still be finite as it concerns subspaces on $\phi(s,a)$ rather than the whole $\mathcal{S} \times \mathcal{A}$.

To further gain its intuition, we can consider discrete MDPs and the feature mapping $\phi(s,a) \in \mathbb{R}^{|\mathcal{S}| |\mathcal{A}|}$ which is a one-hot encoding vector that has zero everywhere except one at the entry corresponding to the pair (s,a) . In this case, the relative condition number is reduced to $\max_{(s,a)} d_P^{\pi_e}(s,a)/\rho(s,a)$.

Theorem 6 (Error for KNRs). *Suppose $\sup_{s,a} \|\phi(s,a)\| \leq 1$, $\lambda = \Omega(1)$, $\zeta^2 = \Omega(1)$, $\|W^*\|_2 = \Omega(1)$ and the partial coverage $C^{\pi_e} < \infty$. With probability $1 - \delta$,*

$$\text{Err}_o \leq c_1 H^2 \left(\text{rank}^2(\Sigma_\rho) + \text{rank}(\Sigma_\rho) \log\left(\frac{c_2}{\delta}\right) \right) \sqrt{\frac{d_{\mathcal{S}} C^{\pi_e}}{n_o}} \cdot \log^{1/2}(1 + n_o), \quad (2)$$

where c_1 and c_2 are some universal constants.

Theorem 6 suggests Err_o is $\tilde{O}(H^2 \text{rank}[\Sigma_\rho]^2 \sqrt{d_{|\mathcal{S}|} C^{\pi_e} / n_o})$. In other words, when C^{π_e} , $\text{rank}[\Sigma_\rho]$ are small and the offline sample size n_o is large enough, Err_e dominates Err_o in theorem 3. Again, in this case, $\text{Err}_e = O\left(H \sqrt{\ln(\mathcal{F}) / n_e}\right)$, and we see that it grows linearly w.r.t horizon H .

Our result is distribution dependent and captures the possible low-rankness of the offline data, i.e., $\text{rank}[\Sigma_\rho]$ depends on ρ and could be much smaller than the ambient dimension of feature $\phi(s, a)$. The quantity C^{π_e} corresponds to the discrepancy measured between the batch data and expert data. This is much smaller than the worst-case concentrability coefficient: $\tilde{C} = \sup_{\pi \in \Pi} C^\pi$.

Remark 7. In RL, a similar quantity has been analyzed in [31], which studies the error bound of linear FQI with pessimism. Comparing to our result only requiring partial coverage, [31, Corollary 4.5] assumes the global coverage, i.e., Σ_ρ is full-rank, which is stronger than $\tilde{C} < \infty$.

GPs Now we specialize our main theorem to non-parametric GP models. For simplicity, following [63], we assume $\mathcal{S} \times \mathcal{A}$ is a compact space. We also suppose the following. Recall $x := (s, a)$.

Assumption 8. $k(x, x) \leq 1, \forall x \in \mathcal{S} \times \mathcal{A}$. $k(\cdot, \cdot)$ is a continuous and positive semidefinite kernel.

Under the assumption 8, we can use Mercer’s theorem [73], which shows that there exists a set of pairs of eigenvalues and eigenfunctions $\{\mu_i, \psi_i\}_{i=1}^\infty$, where $\int \rho(x) \psi_i(x) \psi_i(x) dx = 1$ for all i and $\int \rho(x) \psi_i(x) \psi_j(x) dx = 0$ for $i \neq j$. Eigenfunctions and eigenvalues essentially defines an infinite-dimensional feature mapping $\phi(x) := [\sqrt{\mu_1} \psi_1(x), \dots, \sqrt{\mu_\infty} \psi_\infty(x)]^\top$. Here, $k(x, x) = \phi(x)^\top \phi(x)$, and any function $f \in \mathcal{H}_k$ can be represented as $f(\cdot) = \alpha^\top \phi(\cdot)$. Note that the eigenvalues and eigenfunctions are defined w.r.t the offline data ρ , thus our result here is still distribution dependent rather than a worst case analysis which often appears in online RL/IL settings [63; 32; 78; 17].

Assume eigenvalues $\{\mu_1, \dots, \mu_\infty\}$ are in non-increasing order, we define the effective dimension,

Definition 9 (Effective dimension). $d^* = \min\{j \in \mathbb{N} : j \geq B(j+1)n_o/\zeta^2\}$, $B(j) = \sum_{k=j}^\infty \mu_k$.

The effective dimensions d^* is widely used and calculated for many kernels [85; 7; 72; 28]. In finite-dimensional linear kernels $\{x \mapsto a^\top \phi(x); a \in \mathbb{R}^d\}$ ($k(x, x) = \phi^\top(x) \phi(x)$), we have $d^* \leq \text{rank}[\Sigma_\rho]$. Thus, d^* is considered to be a natural extension of $\text{rank}[\Sigma_\rho]$ to infinite-dimensional models.

Theorem 10 (Error for GPs). Let $\Sigma_{\pi_e} = \mathbb{E}_{x \sim d_P^{\pi_e}}[\phi(x) \phi(x)^\top]$, $\Sigma_\rho = \mathbb{E}_{x \sim \rho}[\phi(x) \phi(x)^\top]$. Suppose assumption 8, $\zeta^2 = \Omega(1)$ and the partial coverage $C^{\pi_e} = \sup_{\|x\|_2 \leq 1} (x^\top \Sigma_{\pi_e} x / x^\top \Sigma_\rho x) < \infty$. With probability $1 - \delta$,

$$\text{Err}_o \leq c_1 H^2 \left((d^*)^2 + d^* \log(c_2/\delta) \right) \sqrt{\frac{d_S C^{\pi_e}}{n_o}} \cdot \sqrt{\log^3(c_2 d_S n_o / \delta) \log(1 + n_o)}, \quad (3)$$

where c_1, c_2 are universal constants.

The theorem suggests that Err_o is $\tilde{O}(H^2 d^{*2} \sqrt{d_S C^{\pi_e} / n_o})$. Thus, when C^{π_e}, d^* are not so large and n_o is large enough, Err_e asymptotically dominates Err_o in theorem 3 (again Err_e is linear in H).

While we defer the detailed proof of the above theorem to Appendix C.3, we highlight some techniques we used here. The analysis is reduced to how to bound the information gain \mathcal{I}_{n_o} and $\mathbb{E}_{x \sim d_P^{\pi_e}}[k_{n_o}(x, x)]$. In both cases, we analyze them into two steps: transforming them into the variational representation and then bounding them via the uniform law with localization (Lemma 39).

Remark 11 (Implication to Offline RL). As related literature, in model-free offline RL, [70; 21] obtain the finite-sample error bounds using nonparametric models. Though their bounds can be characterized by the effective dimension, their bounds assume full coverage, i.e., $\max_{(s,a)} 1/\rho(s, a) < \infty$.

5 Practical Algorithm

In this section we instantiate a practical version of MILO using neural networks for the model class \mathcal{P} and policy class Π . We use the Maximum Mean Discrepancy (MMD) with a Radial Basis Function kernel as our discriminator class \mathcal{F} . Note using MMD as our discrepancy measure allows us to compute the exact maximum discriminator $\arg\max_{f \in \mathcal{F}}$ in closed form (and is detailed in appendix). We use a KL-based trust-region formulation for incremental policy update inside the learned model \hat{P} . Based on Eq. (1), we first formalize the following constrained optimization framework:

$$\min_{\pi \in \Pi} \max_{f \in \mathcal{F}} \left(\mathbb{E}_{(s,a) \sim d_P^\pi} [f(s, a) + b(s, a)] - \mathbb{E}_{(s,a) \sim \mathcal{D}_e} [f(s, a)] \right) \text{ s.t. } \mathbb{E}_{(s,a) \sim \mathcal{D}_e} [\ell(a, s, \pi)] \leq \delta,$$

where $\ell : \mathcal{A} \times \mathcal{S} \times \Pi \mapsto \mathbb{R}$ is a loss function (e.g., negative log-likelihood or any supervised learning loss one would use in BC). Essentially, since we have \mathcal{D}_e available, we use it together with any

Algorithm 2 A practical instantiation of MILO

- 1: **Require:** expert dataset $\mathcal{D}_e = \{s, a\}$, offline dataset $\mathcal{D}_o := \{s, a, s'\}$
 - 2: Train an ensemble of neural network models $\{\hat{g}_1, \dots, \hat{g}_n\}$ where each P_i starts with different random initialization;
 - 3: Set bonus $b(s, a) = \max_{i,j} \|g_i(s, a) - g_j(s, a)\|_2$ and initialize π_{θ_0} .
 - 4: **for** $t = 0 \rightarrow T - 1$ **do**
 - 5: $w_t = \arg \max_{\|w\|_2 \leq 1} w^\top \left(\mathbb{E}_{(s,a) \sim d_{\hat{P}}} [\phi(s, a)] - \mathbb{E}_{(s,a) \sim \mathcal{D}_e} [\phi(s, a)] \right)$, $f_t(s, a) := w_t^\top \phi(s, a)$
 - 6: $\theta_{t+1} = \theta_t - \eta F_{\theta_t}^{-1} \left(\mathbb{E}_{(s,a) \sim d_{\hat{P}}^{\pi_{\theta_t}}} \left[\nabla \ln \pi_{\theta_t}(a|s) A_{\hat{P}, f_t + b}^{\pi_{\theta_t}}(s, a) \right] + \lambda \mathbb{E}_{(s,a) \sim \mathcal{D}_e} [\nabla \ell(a, s, \pi_{\theta_t})] \right)$
 - 7: **end for**
-

supervised learning loss to constrain the policy hypothesis space. Note for a deterministic expert π_e , the expert policy is always a feasible solution. Thus adding this constraint reduces the complexity of the policy class but does not eliminate the expert policy, and our analysis in Section 4 still applies.

In our practical instantiation, we replace the hard constraint instead by a Lagrange multiplier, i.e. we use the behavior cloning objective as a regularization term when solving the min-max problem:

$$\min_{\pi \in \Pi} \max_{f \in \mathcal{F}} \left[\mathbb{E}_{(s,a) \sim d_{\hat{P}}} (f(s, a) + b(s, a)) - \mathbb{E}_{(s,a) \sim \mathcal{D}_e} [f(s, a)] \right] + \lambda \cdot \mathbb{E}_{(s,a) \sim \mathcal{D}_e} [\ell(a, s, \pi)].$$

Given policy π_{θ_t} (θ_t denotes the parameters), we update the discriminator f_t . Then, with a fixed f_t , in order to update policy π we use NPG as in line 6 in Algorithm 2, where $A_{\hat{P}, f+b}^{\pi_{\theta}}$ is the disadvantage function of π_{θ} and $F_{\theta_t} := \mathbb{E}_{(s,a) \sim d_{\hat{P}}^{\pi_{\theta_t}}} [\nabla \ln \pi_{\theta_t}(a|s) \nabla \ln \pi_{\theta_t}(a|s)^\top]$ is the fisher information matrix.

6 Experiments

We aim to answer the following questions with our experiments: (1) How does MILO perform relative to other offline IL methods, (2) What is the impact of pessimism on MILO’s performance? (3) How does the behavior policy’s coverage impact MILO’s performance? (4) How does MILO’s result vary when we increase the number of samples drawn from the expert policy?

We evaluate MILO on five environments from OpenAI Gym [11] simulated with MuJoCo [67]: Hopper-v2, Walker2d-v2, HalfCheetah-v2, Ant-v2, and Humanoid-v2. We compare MILO against the following baselines: (1) ValueDICE [38], a state-of-the-art off-policy IL method modified for the offline IL setting; (2) BC on the expert dataset; and (3) BC on both the offline and expert dataset.

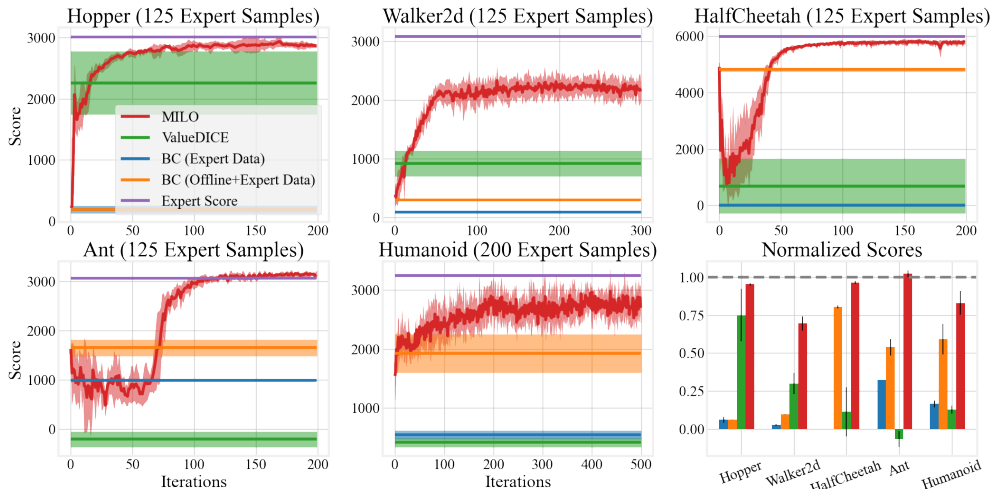


Figure 2: Learning curves across five seeds for MILO plotted against the best performance of BC after 1000 epochs of training on the expert/offline+expert data and the best performance of ValueDICE after 10 thousand iterations. The bottom right bar graph shows the expert performance normalized scores where we plot the performance at the last iteration for MILO.

For the expert dataset, we first train expert policies and then randomly sample (s, a) -pairs from a pool of 100 expert trajectories collected from these expert policies. We randomly sample to create very small expert (s, a) -pair datasets where BC struggles to learn. Note that BC is effective at imitating the expert for MuJoCo tasks even with a single trajectory; prior works [27; 37; 38] have used similar sub-sampling strategies to create expert datasets to make it harder for BC to learn. The offline datasets are collected building on prior Offline RL works [76; 34]; each dataset contains 1 million samples from the environment. We first train behavior policies with mean performances often less than half of the expert performance (Table 1, column 2). All results are averaged over five random seeds. See appendix for details on hyperparameters, environments, and dataset composition.

Table 1: Performance for expert and behavior policy used to collect expert and offline datasets respectively.

Environment	Expert Performance	Behavior Performance
Hopper-v2	3012	752 (25%)
Walker2d-v2	3082	1383 (45%)
HalfCheetah-v2	5986	3972 (66%)
Ant-v2	3072	1208 (40%)
Humanoid-v2	3248	1505 (46%)

6.1 Evaluation on MuJoCo Continuous Control Tasks

Figure 2 presents results comparing MILO against benchmarks. MILO is able to achieve close to expert level performance on three out of the five environments and outperforms both BC and ValueDICE on all five environments. Both MILO and ValueDICE were warmstarted with one epoch of BC on the offline dataset. We significantly outperform BC’s performance when trained on the expert dataset, suggesting MILO indeed mitigates covariate shift through the use of a static offline dataset of (s, a) -pairs. BC on both the offline and expert dataset does improve the performance, but this still cannot successfully imitate the expert since BC has no way of differentiating random/sub-optimal trajectories from the expert samples. ValueDICE, on the other hand, does explicitly aim to imitate the expert samples; however, in theory, it would require either the offline data (i.e. the replay buffer) or the expert samples to have full coverage over the state-action space. Since our offline dataset is mainly collected from a sub-optimal behavior policy and our expert samples are from a high quality expert, neither our offline nor our expert dataset is likely to have full coverage globally; thus potentially hurting the performance of algorithms like ValueDICE. Note that MILO is still able to perform reasonably well across environments even with these offline and expert datasets.

6.2 Ablation

Impact of Pessimism Figure 3 (Left 2) presents MILO’s performance on two representative environments with and without pessimism (i.e., setting penalty to be zero) added to the imitation objective. Pessimism stabilizes and improves the final performance for MILO. In general, MILO consistently outperforms benchmarks and/or achieves expert level performance for a given set of hyperparameters. See the Appendix for evaluation on other environments.

Behavior with more expert samples We investigate whether MILO is able to achieve expert performance with more expert samples in the two environments (walker and humanoid) that it did not




Figure 3: (Left 2) Learning curves for Hopper and Walker2d with (red) and without (blue) pessimism. MILO generally performs worse without pessimism. (Right 2) Learning curve for Walker2d and Humanoid with more expert samples.

solve with very small expert datasets in Figure 2. Figure 3 (Right 2) shows that with one trajectory worth of expert samples, MILO is able to achieve expert performance on walker and humanoid.

Impact of Coverage As our analysis suggests, MILO’s performance degrades as the offline data’s coverage over the expert’s state-action space decreases. We use the behavior policy’s value as a surrogate for lower coverage, i.e. a lower value suggests lower coverage. We generate two additional offline datasets for each environment by lowering the performance of the behavior policy. The three datasets are: (1) the original offline datasets used in Table 1 ($\approx 25\%$ for Hopper-v2 and $\approx 50\%$ for others); (2) ones that have roughly half the performance of (1) (12% for Hopper-v2 and $\approx 25\%$ for others); and (3) ones collected from a random behavior policy (Random). Table 2 shows that MILO performs reasonably on three environments even with a lower coverage dataset (second column) and achieves more than 20% of the expert performance on three environment even with the Random dataset.

Table 2: Expert performance normalized scores on three different offline datasets collected from behavior policies with approximately 50%, 25%, and random performance relative to the expert.

Environment	$\approx 50\%$	$\approx 25\%$	Random
Hopper-v2	0.95 ± 0.01	0.66 ± 0.33	0.42 ± 0.36
Walker2d-v2	0.72 ± 0.02	0.27 ± 0.06	0.23 ± 0.12
HalfCheetah-v2	0.96 ± 0.01	0.01 ± 0.02	0.01 ± 0.02
Ant-v2	1.02 ± 0.02	0.99 ± 0.01	0.21 ± 0.52
Humanoid-v2	0.88 ± 0.10	0.72 ± 0.03	0.08 ± 0.01

7 Conclusion

MILO investigates how to mitigate covariate shift in IL using an offline dataset of environment interactions that has partial coverage of the expert’s state-action space. We show the effectiveness of MILO both in theory and in practice. In future works, we hope to scale to image-based control.

We want to highlight the potential negative societal/ethical impacts our work. An IL algorithm is only as good as the expert that it is imitating, not only in terms of performance but also with regards to the negative biases and intentions that the demonstrator has. When designing real-world experiments/applications for MILO we believe the users should do their due diligence on removing any negative bias or malicious intent in the demonstrations that they provide.

References

- [1] Y. Abbasi-yadkori, D. Pál, and C. Szepesvári. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, volume 24. Curran Associates, Inc., 2011.
- [2] P. Abbeel and A. Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *ICML*, page 1. ACM, 2004.
- [3] A. Agarwal, M. Henaff, S. Kakade, and W. Sun. Pc-pg: Policy cover directed exploration for provable policy gradient learning. *NeurIPS*, 2020.
- [4] A. Agarwal, N. Jiang, S. M. Kakade, and W. Sun. Reinforcement learning: Theory and algorithms. *CS Dept., UW Seattle, Seattle, WA, USA, Tech. Rep*, 2019.
- [5] A. Antos, C. Szepesvári, and R. Munos. Learning near-optimal policies with bellman-residual minimization based fitted policy iteration and a single sample path. *Machine Learning*, 71:89–129, 2008.
- [6] K. Azzizadenesheli, E. Brunskill, and A. Anandkumar. Efficient exploration through bayesian deep q-networks. In *2018 Information Theory and Applications Workshop (ITA)*, pages 1–9. IEEE, 2018.
- [7] F. Bach. On the equivalence between kernel quadrature rules and random feature expansions. *Journal of machine learning research*, 18(21):1–38, 2017.
- [8] S. Bansal, R. Calandra, T. Xiao, S. Levine, and C. J. Tomiini. Goal-driven dynamics learning via bayesian optimization. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 5168–5173. IEEE, 2017.
- [9] P. L. Bartlett, O. Bousquet, and S. Mendelson. Local rademacher complexities. *Annals of Statistics*, 33(4):1497–1537, 08 2005.
- [10] K. Brantley, W. Sun, and M. Henaff. Disagreement-regularized imitation learning. In *International Conference on Learning Representations*, 2019.
- [11] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. Openai gym, 2016.
- [12] J. Buckman, C. Gelada, and M. G. Bellemare. The importance of pessimism in fixed-dataset policy optimization. *arXiv preprint arXiv:2009.06799*, 2020.
- [13] D. Calandriello, L. Carratino, A. Lazaric, M. Valko, and L. Rosasco. Gaussian process optimization with adaptive sketching: Scalable and no regret. In *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 533–557, 2019.
- [14] A. J. Chan and M. van der Schaar. Scalable bayesian inverse reinforcement learning. *arXiv preprint arXiv:2102.06483*, 2021.
- [15] K.-W. Chang, H. He, H. Daumé III, and J. Langford. Learning to search for dependencies. *arXiv preprint arXiv:1503.05615*, 2015.
- [16] J. Chen and N. Jiang. Information-theoretic considerations in batch reinforcement learning. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pages 1042–1051, 2019.
- [17] S. R. Chowdhury and A. Gopalan. Online learning in kernelized markov decision processes. In K. Chaudhuri and M. Sugiyama, editors, *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pages 3197–3205. PMLR, 16–18 Apr 2019.
- [18] M. Deisenroth and C. E. Rasmussen. PILCO: A model-based and data-efficient approach to policy search. In *International Conference on Machine Learning*, pages 465–472, 2011.

- [19] S. S. Du, S. M. Kakade, J. D. Lee, S. Lovett, G. Mahajan, W. Sun, and R. Wang. Bilinear classes: A structural framework for provable generalization in rl. *ICML*, 2021.
- [20] Y. Duan, Z. Jia, and M. Wang. Minimax-optimal off-policy evaluation with linear function approximation. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 2701–2709, 2020.
- [21] Y. Duan, C. Jin, and Z. Li. Risk bounds and rademacher complexity in batch reinforcement learning. *arXiv preprint arXiv:2103.13883*, 2021.
- [22] D. Ernst, P. Geurts, and L. Wehenkel. Tree-based batch mode reinforcement learning. *Journal of Machine Learning Research*, 6:503–556, 2005.
- [23] R. Fakoor, J. Mueller, P. Chaudhari, and A. J. Smola. Continuous doubly constrained batch reinforcement learning. *arXiv preprint arXiv:2102.09225*, 2021.
- [24] J. Fan, Z. Wang, Y. Xie, and Z. Yang. A theoretical analysis of deep q-learning. In *Proceedings of the 2nd Conference on Learning for Dynamics and Control*, volume 120 of *Proceedings of Machine Learning Research*, pages 486–489, 2020.
- [25] J. F. Fisac, A. K. Akametalu, M. N. Zeilinger, S. Kaynama, J. Gillula, and C. J. Tomlin. A general safety framework for learning-based control in uncertain robotic systems. *IEEE Transactions on Automatic Control*, 64(7):2737–2752, 2018.
- [26] S. Fujimoto, D. Meger, and D. Precup. Off-policy deep reinforcement learning without exploration. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2052–2062. PMLR, 09–15 Jun 2019.
- [27] J. Ho and S. Ermon. Generative adversarial imitation learning. In *NIPS*, 2016.
- [28] D. Janz, D. Burt, and J. Gonzalez. Bandit optimisation of functions in the matern kernel rkhs. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 2486–2495, 2020.
- [29] D. Jarrett, I. Bica, and M. van der Schaar. Strictly batch imitation learning by energy-based distribution matching. In *Advances in Neural Information Processing Systems*, volume 33, pages 7354–7365, 2020.
- [30] N. Jiang. Notes on tabular methods, 2020.
- [31] Y. Jin, Z. Yang, and Z. Wang. Is pessimism provably efficient for offline rl? *arXiv preprint arXiv:2012.15085*, 2020.
- [32] S. Kakade, A. Krishnamurthy, K. Lowrey, M. Ohnishi, and W. Sun. Information theoretic regret bounds for online nonlinear control. In *Advances in Neural Information Processing Systems*, volume 33, pages 15312–15325, 2020.
- [33] R. Kidambi, J. Chang, and W. Sun. Optimism is all you need: Model-based imitation learning from observation alone. *arXiv preprint arXiv:2102.10769*, 2021.
- [34] R. Kidambi, A. Rajeswaran, P. Netrapalli, and T. Joachims. Morel: Model-based offline reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 21810–21823. Curran Associates, Inc., 2020.
- [35] E. Klein, B. Piot, M. Geist, and O. Pietquin. A cascaded supervised learning approach to inverse reinforcement learning. In *ECML/PKDD*, 2013.
- [36] J. Ko, D. J. Klein, D. Fox, and D. Haehnel. Gaussian processes and reinforcement learning for identification and control of an autonomous blimp. In *Proceedings 2007 IEEE International Conference on Robotics and Automation*, pages 742–747. IEEE, 2007.
- [37] I. Kostrikov, K. K. Agrawal, D. Dwibedi, S. Levine, and J. Tompson. Discriminator-actor-critic: Addressing sample inefficiency and reward bias in adversarial imitation learning. In *ICLR*. OpenReview.net, 2019.

- [38] I. Kostrikov, O. Nachum, and J. Tompson. Imitation learning via off-policy distribution matching. *ICLR*, 2019.
- [39] A. Kumar, A. Zhou, G. Tucker, and S. Levine. Conservative q-learning for offline reinforcement learning. *arXiv preprint arXiv:2006.04779*, 2020.
- [40] M. Laskey, J. Lee, W. Y. Hsieh, R. Liaw, J. Mahler, R. Fox, and K. Goldberg. Iterative noise injection for scalable imitation learning. *CoRR*, abs/1703.09327, 2017.
- [41] L. Le Gratiet, L. Le Gratiet, J. Garnier, and J. Garnier. Asymptotic analysis of the learning curve for gaussian process regression. *Machine learning*, 98(3):407–433, 2015.
- [42] P. Liao, Z. Qi, and S. Murphy. Batch policy learning in average reward markov decision processes. *arXiv preprint arXiv:2007.11771*, 2020.
- [43] Y. Liu, A. Swaminathan, A. Agarwal, and E. Brunskill. Provably good batch off-policy reinforcement learning without great exploration. In *Advances in Neural Information Processing Systems*, volume 33, pages 1264–1274, 2020.
- [44] T. Matsushima, H. Furuta, Y. Matsuo, O. Nachum, and S. Gu. Deployment-efficient reinforcement learning via model-based offline optimization. *ICLR*, 2020.
- [45] R. Munos and C. Szepesvári. Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research*, 9(May):815–857, 2008.
- [46] O. Nachum, Y. Chow, B. Dai, and L. Li. Dualdice: Behavior-agnostic estimation of discounted stationary distribution corrections. *Advances in Neural Information Processing Systems 2019*, 2019.
- [47] I. Osband, J. Aslanides, and A. Cassirer. Randomized prior functions for deep reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- [48] D. Pathak, D. Gandhi, and A. Gupta. Self-supervised exploration via disagreement. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5062–5071. PMLR, 09–15 Jun 2019.
- [49] D. A. Pomerlau. ALVINN: An autonomous land vehicle in a neural network. In D. S. Touretzky, editor, *Advances in Neural Information Processing Systems*, pages 323–331, San Mateo, CA, 1989. Morgan Kaufmann Publishers inc.
- [50] R. Rafailov, T. Yu, A. Rajeswaran, and C. Finn. Visual adversarial imitation learning using variational models. *CoRR*, abs/2107.08829, 2021.
- [51] A. Rahimi and B. Recht. Random features for large-scale kernel machines. In J. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2008.
- [52] N. Rajaraman, L. F. Yang, J. Jiao, and K. Ramachandran. Toward the fundamental limits of imitation learning. *arXiv preprint arXiv:2009.05990*, 2020.
- [53] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine. Learning Complex Dexterous Manipulation with Deep Reinforcement Learning and Demonstrations. In *Proceedings of Robotics: Science and Systems (RSS)*, 2018.
- [54] P. Rashidinejad, B. Zhu, C. Ma, J. Jiao, and S. Russell. Bridging offline reinforcement learning and imitation learning: A tale of pessimism. *arXiv preprint arXiv:2103.12021*, 2021.
- [55] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.
- [56] S. Reddy, A. D. Dragan, and S. Levine. Sqil: Imitation learning via reinforcement learning with sparse rewards. In *ICLR*, 2020.

- [57] S. Ross and J. A. Bagnell. Efficient reductions for imitation learning. In *AISTATS*, pages 661–668, 2010.
- [58] S. Ross and J. A. Bagnell. Reinforcement and imitation learning via interactive no-regret learning. *arXiv preprint arXiv:1406.5979*, 2014.
- [59] S. Ross, G. J. Gordon, and J. Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *AISTATS*, 2011.
- [60] M. W. Seeger, S. M. Kakade, and D. P. Foster. Information consistency of nonparametric gaussian process methods. *IEEE Transactions on Information Theory*, 54(5):2376–2382, 2008.
- [61] P. Sollich and A. Halees. Learning curves for gaussian process regression: Approximations and bounds. *Neural computation*, 14(6):1393–1428, 2002.
- [62] J. Spencer, S. Choudhury, A. Venkatraman, B. Ziebart, and J. A. Bagnell. Feedback in imitation learning: The three regimes of covariate shift. *arXiv preprint arXiv:2102.02872*, 2021.
- [63] N. Srinivas, A. Krause, S. Kakade, and M. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the 27th International Conference on International Conference on Machine Learning, ICML’10*, page 1015–1022, 2010.
- [64] W. Sun, N. Jiang, A. Krishnamurthy, A. Agarwal, and J. Langford. Model-based rl in contextual decision processes: Pac bounds and exponential improvements over model-free approaches. In *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 2898–2933, 2019.
- [65] W. Sun, A. Vemula, B. Boots, and D. Bagnell. Provably efficient imitation learning from observation alone. In *International Conference on Machine Learning*, pages 6036–6045. PMLR, 2019.
- [66] W. Sun, A. Venkatraman, G. J. Gordon, B. Boots, and J. A. Bagnell. Deeply aggregated: Differentiable imitation learning for sequential prediction. *arXiv preprint arXiv:1703.01030*, 2017.
- [67] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *IROS*, pages 5026–5033. IEEE, 2012.
- [68] A. Touati, A. Zhang, J. Pineau, and P. Vincent. Stable policy optimization via off-policy divergence regularization. *arXiv preprint arXiv:2003.04108*, 2020.
- [69] M. Uehara, J. Huang, and N. Jiang. Minimax weight and q-function learning for off-policy evaluation. In *Proceedings of the 37th International Conference on Machine Learning*, pages 9659–9668, 2020.
- [70] M. Uehara, M. Imaizumi, N. Jiang, N. Kallus, W. Sun, and T. Xie. Finite sample analysis of minimax offline reinforcement learning: Completeness, fast rates and first-order efficiency. *arXiv preprint arXiv:2102.02981*, 2021.
- [71] J. Umlauft, L. Pöhler, and S. Hirche. An uncertainty-based control lyapunov approach for control-affine systems modeled by gaussian process. *IEEE Control Systems Letters*, 2(3):483–488, 2018.
- [72] M. Valko, N. Korda, R. Munos, I. Flaounas, and N. Cristianini. Finite-time analysis of kernelised contextual bandits. In *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence, UAI’13*, page 654–663, Arlington, Virginia, USA, 2013. AUAI Press.
- [73] M. J. Wainwright. *High-Dimensional Statistics : A Non-Asymptotic Viewpoint*. Cambridge University Press, New York, 2019.
- [74] R. Wang, D. P. Foster, and S. M. Kakade. What are the statistical limits of offline rl with linear function approximation?. *arXiv preprint arXiv:2010.11895*, 2020.
- [75] C. K. Williams and F. Vivarelli. Upper and lower bounds on the learning curve for gaussian processes. *Machine learning*, 40(1):77–102, 2000.

- [76] Y. Wu, G. Tucker, and O. Nachum. Behavior regularized offline reinforcement learning. *arXiv preprint arXiv:1911.11361*, 2019.
- [77] T. Xie and N. Jiang. Q* approximation schemes for batch reinforcement learning: A theoretical comparison. *UAI2020*, 2020.
- [78] Z. Yang, C. Jin, Z. Wang, M. Wang, and M. Jordan. Provably efficient reinforcement learning with kernel and neural function approximations. In *Advances in Neural Information Processing Systems*, volume 33, pages 13903–13916, 2020.
- [79] M. Yin, Y. Bai, and Y.-X. Wang. Near-optimal offline reinforcement learning via double variance reduction. *arXiv preprint arXiv:2102.01748*, 2021.
- [80] M. Yin and Y.-X. Wang. Asymptotically efficient off-policy evaluation for tabular reinforcement learning. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, pages 3948–3958, 2020.
- [81] T. Yu, A. Kumar, R. Rafailov, A. Rajeswaran, S. Levine, and C. Finn. Combo: Conservative offline model-based policy optimization. *arXiv preprint arXiv:2102.08363*, 2021.
- [82] T. Yu, G. Thomas, L. Yu, S. Ermon, J. Y. Zou, S. Levine, C. Finn, and T. Ma. Mopo: Model-based offline policy optimization. In *Advances in Neural Information Processing Systems*, volume 33, pages 14129–14142, 2020.
- [83] A. Zanette. Exponential lower bounds for batch reinforcement learning: Batch rl can be exponentially harder than online rl. *arXiv preprint arXiv:2012.08005*, 2020.
- [84] R. Zhang, B. Dai, L. Li, and D. Schuurmans. Gendice: Generalized offline estimation of stationary values. In *International Conference on Learning Representations*, 2020.
- [85] T. Zhang. Learning bounds for kernel regression using effective data dimensionality. *Neural computation*, 17(9):2077–2098, 2005.
- [86] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey. Maximum entropy inverse reinforcement learning. In *AAAI*, 2008.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? **[Yes]** See [Section 1](#)
 - (b) Did you describe the limitations of your work? **[Yes]** See [Section 4](#)
 - (c) Did you discuss any potential negative societal impacts of your work? **[Yes]** See [Section 7](#) for a discussion
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]** We have read through the ethical guidelines.
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? **[Yes]** See [Section 4](#)
 - (b) Did you include complete proofs of all theoretical results? **[Yes]** See [Section 4](#)
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[Yes]** Experimental details and exact hyperparameters are reported in the supplemental material. Code and data will be released along with the camera ready version with a URL to the github repository.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[Yes]** All details are enumerated and documented in the supplemental material.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[Yes]** All plots and tables report error bars of the average across five random seeds.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[No]** We use small neural networks and do not work with image inputs thus did not require GPU acceleration or specialized hardware. Each of our experiments were run on 2 cpu cores and 8GB of ram.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? **[Yes]** See [Section 6](#) where we cite the creators of the assets.
 - (b) Did you mention the license of the assets? **[No]** For OpenAI Gym, we use their opensource software and we payed for the proprietary MuJoCo license.
 - (c) Did you include any new assets either in the supplemental material or as a URL? **[N/A]**
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? **[N/A]**
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[N/A]**
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? **[N/A]**
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? **[N/A]**
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **[N/A]**

Contents

1	Introduction	1
1.1	Related work	3
2	Setting	3
3	Algorithm	4
3.1	Specialization to offline RL	5
4	Analysis	5
4.1	Analysis: Discrete MDPs	6
4.2	Analysis: KNRs and GPs for Continuous MDPs	6
5	Practical Algorithm	7
6	Experiments	8
6.1	Evaluation on MuJoCo Continuous Control Tasks	9
6.2	Ablation	9
7	Conclusion	10
A	Penalty Designs	19
A.1	Tabular models	19
A.2	KNRs	19
A.3	Gaussian processes	20
B	Proof of theorem 3	20
C	Finite sample error bound for each model	22
C.1	Discrete MDPs	22
C.2	KNRs	23
C.3	Gaussian processes	25
C.4	Missing Proofs	28
C.4.1	Missing proofs for tabular result	28
C.4.2	Missing proofs for KNR results	29
C.4.3	Missing proofs of non-parametric model	32
D	Auxiliary Lemmas	36
E	Implementation Details	41
E.1	Environment Details	41
E.2	Dynamics Ensemble Architecture and Model Learning	41
E.3	Policy Architecture and TRPO Details	42

E.4	Hyperparameter Selection	42
E.5	Discriminator Update and Cost Function Details	42
F	Additional Experiments	43

A Penalty Designs

We show that the penalty design in section 4 is valid, i.e, the model is well-calibrated for tabular MDPs, KNRs, and GPs.

A.1 Tabular models

Lemma 12. *With probability $1 - \delta$,*

$$\|\hat{P}(\cdot|s, a) - P(\cdot|s, a)\|_1 \leq \sqrt{\frac{|\mathcal{S}| \log 2 + \log(2|\mathcal{S}||\mathcal{A}|/\delta)}{2\{N(s, a) + \lambda\}}} + \frac{\lambda}{N(s, a) + \lambda} \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}.$$

Proof. When $N(s, a) > 0$, we use the concentration inequality of discrete distributions [30]. Then, with probability $1 - \delta$,

$$\left\| \frac{N(\cdot|s, a)}{N(s, a)} - P(\cdot|s, a) \right\|_1 \leq \sqrt{\frac{|\mathcal{S}| \log 2 + \log(2|\mathcal{S}||\mathcal{A}|/\delta)}{2N(s, a)}} \quad \forall (s, a) \in \{(s, a) : N(s, a) > 0\}.$$

Thus, noting $0 < N(s, a)/(N(s, a) + \lambda) < 1$, with probability $1 - \delta$, we have $\forall (s, a) \in \{(s, a) : N(s, a) > 0\}$,

$$\left\| \frac{N(\cdot|s, a)}{(N(s, a) + \lambda)} - P(\cdot|s, a) \times \frac{N(s, a)}{N(s, a) + \lambda} \right\|_1 \leq \sqrt{\frac{|\mathcal{S}| \log 2 + \log(2|\mathcal{S}||\mathcal{A}|/\delta)}{2\{N(s, a) + \lambda\}}}. \quad (4)$$

Besides, the above inequality is still well-defined and holds including the case $N(s, a) = 0$. Thus, with probability $1 - \delta$, we have $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$, we have eq. (4).

Recall the estimator \hat{P} is $N(s'|s, a)/(N(s, a) + \lambda)$. Therefore,

$$\begin{aligned} \|\hat{P}(\cdot|s, a) - P(\cdot|s, a)\|_1 &\leq \left\| \hat{P}(\cdot|s, a) - P(\cdot|s, a) \times \frac{N(s, a)}{N(s, a) + \lambda} \right\|_1 + \left\| P(\cdot|s, a) - P(\cdot|s, a) \times \frac{N(s, a)}{N(s, a) + \lambda} \right\|_1 \\ &\leq \sqrt{\frac{|\mathcal{S}| \log 2 + \log(2|\mathcal{S}||\mathcal{A}|/\delta)}{2\{N(s, a) + \lambda\}}} + \frac{\lambda}{N(s, a) + \lambda}. \end{aligned}$$

This concludes the proof. \square

A.2 KNRs

In KNRs, the ground truth model is $s' = W^* \phi(s, a) + \epsilon$, $\epsilon \sim \mathcal{N}(0, \zeta^2 \mathbf{I})$, where $s \in \mathbb{R}^{d_S}$, $a \in \mathbb{R}^{d_A}$, $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$. We define

$$\|\phi(s, a)\|_{\Sigma_{n_o}^{-1}} := \phi^\top(s, a) \Sigma_{n_o}^{-1} \phi(s, a).$$

Lemma 13. *With probability at least $1 - \delta$, we have:*

$$\left\| \hat{P}(\cdot|s, a) - P(\cdot|s, a) \right\|_1 \leq \min \left\{ \frac{\beta_{n_o}}{\zeta} \|\phi(s, a)\|_{\Sigma_{n_o}^{-1}}, 2 \right\} \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A},$$

where

$$\beta_{n_o} = \sqrt{2\lambda \|W^*\|_2^2 + 8\zeta^2 (d_S \ln(5) + \ln(1/\delta) + \bar{\mathcal{I}}_{n_o})}, \quad \bar{\mathcal{I}}_{n_o} = \ln(\det(\Sigma_{n_o})/\det(\lambda \mathbf{I})).$$

Proof. The proof directly follows the confidence ball construction and proof from [32]. Specifically, from Lemma B.5 in [32], we have that with probability at least $1 - \delta$,

$$\left\| (\widehat{W} - W^*) (\Sigma_{n_o})^{1/2} \right\|_2^2 \leq \beta_{n_o}^2.$$

Thus, with Lemma 34, we have:

$$\begin{aligned} \left\| \hat{P}(\cdot|s, a) - P(\cdot|s, a) \right\|_1 &\leq \frac{1}{\zeta} \left\| (\widehat{W} - W^*) \phi(s, a) \right\|_2 \\ &\leq \left\| (\widehat{W} - W^*) (\Sigma_{n_o})^{1/2} \right\|_2 \|\phi(s, a)\|_{\Sigma_{n_o}^{-1}} / \zeta \leq \frac{\beta_{n_o}}{\zeta} \|\phi(s, a)\|_{\Sigma_{n_o}^{-1}}. \end{aligned}$$

This concludes the proof. \square

A.3 Gaussian processes

Let \mathcal{H}_k be the RKHS with the kernel $k(\cdot, \cdot)$. We denote the associated norm and inner product by $\|\cdot\|_k$ and $\langle \cdot, \cdot \rangle_k$. In GPs, the ground truth model is defined as $s' = g^*(s, a) + \epsilon$, $\epsilon \sim \mathcal{N}(0, \zeta^2 \mathbf{I})$, where g^* belongs to an RKHS \mathcal{H}_k .

Lemma 14. *With probability $1 - \delta$,*

$$\|\hat{P}(\cdot|s, a) - P(\cdot|s, a)\|_1 \leq \min\left(\frac{\beta_{n_o}}{\zeta} \sqrt{k_{n_o}((s, a), (s, a))}, 2\right) \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A},$$

and

$$\beta_{n_o} = \sqrt{d_{\mathcal{S}}\{2 + 150 \log^3(d_{\mathcal{S}} n_o / \delta) \mathcal{I}_{n_o}\}}, \quad \mathcal{I}_{n_o} = \log(\det(\mathbf{I} + \zeta^{-2} \mathbf{K}_{n_o})).$$

Proof. Let \hat{g}_i and g^* be i -th component of \hat{g} and g^* . We have

$$\begin{aligned} \|\hat{P}(\cdot|s, a) - P(\cdot|s, a)\|_1 &\leq \frac{1}{\zeta} \|\hat{g}(s, a) - g^*(s, a)\|_2 && \text{(Lemma 34)} \\ &= \frac{1}{\zeta} \left(\sum_{i=1}^{d_{\mathcal{S}}} \{\hat{g}_i(s, a) - g_i^*(s, a)\}^2 \right)^{1/2} \\ &\leq \frac{1}{\zeta} \left(\sum_{i=1}^{d_{\mathcal{S}}} k_{n_o}((s, a), (s, a)) \|\hat{g}_i - g_i^*\|_{k_{n_o}}^2 \right)^{1/2}. \end{aligned}$$

(CS inequality and $g = \langle g(\cdot), k((s, a), \cdot) \rangle_{k_{n_o}}$)

By [63, Theorem 6], with probability $1 - \delta$, we have

$$\|\hat{g}_i(s, a) - g_i^*\|_{k_{n_o}} \leq \beta_{n_o} \quad \forall i \in [1, \dots, d_{\mathcal{S}}].$$

This concludes the statement. □

B Proof of theorem 3

In this section, we prove [Theorem 3](#). We also prove the RL version of [Theorem 3](#) when the cost c is given and the goal is policy optimization. Before that, we prepare several lemmas.

Lemma 15. *With probability $1 - \delta$, we have $\forall f \in \mathcal{F}$,*

$$|\mathbb{E}_{(s,a) \sim d^{\pi_e}} [f(s, a)] - \mathbb{E}_{D_e} [f(s, a)]| \leq \epsilon_{\text{stat}}, \quad \epsilon_{\text{stat}} = \sqrt{\log(2|\mathcal{F}|/\delta)/2n_e}.$$

Proof. From Hoeffding's inequality and a union bound over \mathcal{F} . □

Lemma 16 (Pessimistic Policy Evaluation 1). *Suppose [Assumption 2](#) holds and $\max_{f \in \mathcal{F}} \|f\|_{\infty} \leq 1$. With probability at least $1 - \delta$, $\forall \pi \in \Pi, \forall f \in \mathcal{F}$,*

$$0 \leq V_{\hat{P}, f+b}^{\pi} - V_{P, f}^{\pi}.$$

Proof of Lemma 16. We denote the expected total cost of π under \hat{P} and cost function f by $V_{\hat{P}, f; h}^{\pi}(s, a)$. In this proof, we condition on the event

$$\|\hat{P}(\cdot|s, a) - P(\cdot|s, a)\|_1 \leq \min(\sigma(s, a), 2) \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}.$$

We use the inductive hypothesis argument. We start from $h = H + 1$, where $V_{\hat{P}, f+b; H+1}^{\pi} = V_{P, f; H+1}^{\pi} = 0$. Assume the inductive hypothesis holds at $h + 1$, i.e.,

$$0 \leq V_{\hat{P}, f+b; h+1}^{\pi}(s) - V_{P, f; h+1}^{\pi}(s), \quad \forall s \in \mathcal{S}, \forall \pi \in \Pi, \forall f \in \mathcal{F}.$$

Then, $\forall \pi \in \Pi, \forall f \in \mathcal{F}$,

$$\begin{aligned}
& Q_{P,f,h}^\pi(s, a) - Q_{\hat{P},f+b,h}^\pi(s, a) \\
&= -b(s, a) + \mathbb{E}_{s' \sim \hat{P}(\cdot|s,a)}[V_{P,f,h+1}^\pi(s')] - \mathbb{E}_{s' \sim P(\cdot|s,a)}[V_{\hat{P},f+b,h+1}^\pi(s')] \\
&\leq -b(s, a) + \mathbb{E}_{s' \sim \hat{P}(\cdot|s,a)}[V_{P,f,h+1}^\pi(s')] - \mathbb{E}_{s' \sim P(\cdot|s,a)}[V_{P,f,h+1}^\pi(s')] \\
&\hspace{15em} \text{(Inductive hypothesis assumption)} \\
&\leq -b(s, a) + H \|\hat{P}(\cdot|s, a) - P(\cdot|s, a)\|_1 \hspace{10em} (\|\mathcal{F}\|_\infty \leq 1) \\
&\leq -H \min(\sigma(s, a), 2) + H \min(\sigma(s, a), 2) = 0. \hspace{5em} \text{(Bonus construction)}
\end{aligned}$$

Then, noting $Q_{P,f,h}^\pi(s, \pi(s)) - Q_{\hat{P},f+b,h}^\pi(s, \pi(s)) = V_{P,f,h}^\pi(s) - V_{\hat{P},f+b,h}^\pi(s)$, we have

$$V_{P,f,h}^\pi(s) - V_{\hat{P},f+b,h}^\pi(s) \leq 0 \quad \forall \pi \in \Pi, \forall f \in \mathcal{F}.$$

This concludes the induction step.

Then, we have

$$V_{P,f}^\pi - V_{\hat{P},f+b}^\pi = V_{P,f:1}^\pi - V_{\hat{P},f+b:1}^\pi \leq 0 \quad \forall \pi \in \Pi, \forall f \in \mathcal{F}.$$

□

Lemma 17 (Pessimistic Policy Evaluation 2). *Suppose [Assumption 2](#) holds and $\|\mathcal{F}\|_\infty \leq 1$. With probability at least $1 - \delta$, $\forall \pi \in \Pi, \forall f \in \mathcal{F}$,*

$$V_{\hat{P},f+b}^\pi - V_{P,f}^\pi \leq \text{Error}, \quad \text{Error} := (3H^2 + H)\mathbb{E}_{(s,a) \sim d_P^\pi}[\min(\sigma(s, a), 2)].$$

Proof of [Lemma 17](#). In this proof, we condition on the event

$$\|\hat{P}(\cdot|s, a) - P(\cdot|s, a)\|_1 \leq \min(\sigma(s, a), 2) \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}.$$

We invoke simulation [Lemma 33](#). Then, we have $\forall \pi \in \Pi, \forall f \in \mathcal{F}$

$$\begin{aligned}
V_{\hat{P},f+b}^\pi - V_{P,f}^\pi &= \sum_{h=1}^H \mathbb{E}_{(s,a) \sim d_P^\pi} [b(s, a) + \mathbb{E}_{s' \sim \hat{P}(\cdot|s,a)}[V_{\hat{P},f+b,h}^\pi(s')] - \mathbb{E}_{s' \sim P(\cdot|s,a)}[V_{\hat{P},f+b,h}^\pi(s')]] \\
&\leq \sum_{h=1}^H \mathbb{E}_{(s,a) \sim d_P^\pi} [b(s, a) + \|V_{\hat{P},f+b,h}^\pi\|_\infty \|\hat{P}(\cdot|s, a) - P(\cdot|s, a)\|_1] \\
&\leq H \mathbb{E}_{(s,a) \sim d_P^\pi} [H \min(\sigma(s, a), 2) + H(2H + 1) \min(\sigma(s, a), 2)] \\
&\hspace{15em} (\|V_{\hat{P},f+b,h}^\pi\|_\infty \leq H(2H + 1)) \\
&= (3H^2 + H)\mathbb{E}_{(s,a) \sim d_P^\pi} [\min(\sigma(s, a), 2)].
\end{aligned}$$

Here, we use $\|V_{\hat{P},f+b,h}^\pi\|_\infty \leq H(2H + 1)$ which is derived by $0 \leq f + b \leq 2H + 1$. □

By using the above lemmas, we prove our main result.

Proof of [Theorem 3](#). In this proof, we condition on the event

$$\|\hat{P}(\cdot|s, a) - P(\cdot|s, a)\|_1 \leq \min(\sigma(s, a), 2),$$

which holds with probability $1 - \delta$, and the event in [Lemma 15](#), which holds with probability $1 - \delta$.

Then, with probability $1 - 2\delta$, we have

$$\begin{aligned}
V_{P,c}^{\hat{\pi}_{\text{IL}}} - V_{P,c}^{\pi_e} &\leq V_{\hat{P},c+b}^{\hat{\pi}_{\text{IL}}} - V_{P,c}^{\pi_e} && \text{(Lemma 16)} \\
&\leq H \max_{f \in \mathcal{F}} \{ \mathbb{E}_{(s,a) \sim d_{\hat{P}}^{\hat{\pi}_{\text{IL}}}} [f(s,a) + b(s,a)] - \mathbb{E}_{(s,a) \sim d_P^{\pi_e}} [f(s,a)] \} && (c \in \mathcal{F}) \\
&\leq H \max_{f \in \mathcal{F}} \{ \mathbb{E}_{(s,a) \sim d_{\hat{P}}^{\hat{\pi}_{\text{IL}}}} [f(s,a) + b(s,a)] - \mathbb{E}_{\mathcal{D}_e} [f(s,a)] \} + H\epsilon_{\text{stats}} && \text{(Lemma 15)} \\
&\leq H \max_{f \in \mathcal{F}} \{ \mathbb{E}_{(s,a) \sim d_{\hat{P}}^{\pi_e}} [f(s,a) + b(s,a)] - \mathbb{E}_{\mathcal{D}_e} [f(s,a)] \} + H\epsilon_{\text{stats}} \\
&\hspace{15em} (\pi_e \in \Pi \text{ and the definition of } \hat{\pi}_{\text{IL}}) \\
&\leq H \max_{f \in \mathcal{F}} \{ \mathbb{E}_{(s,a) \sim d_{\hat{P}}^{\pi_e}} [f(s,a) + b(s,a)] - \mathbb{E}_{(s,a) \sim d_P^{\pi_e}} [f(s,a)] \} + 2H\epsilon_{\text{stats}} \\
&\hspace{15em} \text{(Lemma 15)} \\
&\leq \max_{f \in \mathcal{F}} \{ V_{\hat{P},f+b}^{\pi_e} - V_{P,f}^{\pi_e} \} + 2H\epsilon_{\text{stats}} \\
&\leq (3H^2 + H) \mathbb{E}_{(s,a) \sim d_P^{\pi_e}} [\min(\sigma(s,a), 2)] + 2H\epsilon_{\text{stats}} && \text{(Lemma 17)} \\
&\leq (6H^2 + 2H) \mathbb{E}_{(s,a) \sim d_P^{\pi_e}} [\min(\sigma(s,a), 1)] + 2H\epsilon_{\text{stats}}.
\end{aligned}$$

This concludes the proof. \square

Finally, we prove the finite-sample error bounds for the RL case. Similar results are obtained in [34; 82]. We use this theorem in the next section.

Theorem 18 (Bounds for RL). *Suppose $\pi^* \in \Pi$, $P \in \mathcal{P}$ and Assumption 2. With probability $1 - 2\delta$, we have*

$$V_{P,c}^{\hat{\pi}_{\text{RL}}} - V_{P,c}^{\pi^*} \leq (6H^2 + 2H) \mathbb{E}_{(s,a) \sim d_P^{\pi^*}} [\min(\sigma(s,a), 1)]. \quad (5)$$

Proof of Theorem 18.

$$\begin{aligned}
V_{P,c}^{\hat{\pi}_{\text{RL}}} - V_{P,c}^{\pi^*} &\leq V_{\hat{P},c+b}^{\hat{\pi}_{\text{RL}}} - V_{P,c}^{\pi^*} && \text{(Lemma 16)} \\
&= V_{\hat{P},c+b}^{\pi^*} - V_{P,c}^{\pi^*} && (\pi^* \in \Pi \text{ and the definition of } \hat{\pi}_{\text{RL}}) \\
&= (3H^2 + H) \mathbb{E}_{(s,a) \sim d_{\hat{P}}^{\pi^*}} [\min(\sigma(s,a), 2)] && \text{(Lemma 17)} \\
&\leq (6H^2 + 2H) \mathbb{E}_{(s,a) \sim d_{\hat{P}}^{\pi^*}} [\min(\sigma(s,a), 1)].
\end{aligned}$$

This concludes the proof. \square

C Finite sample error bound for each model

In this section, we analyze the bound for the following models: (1) discrete MDPs, (2) KNRs, (3) GPs. All of the proofs are deferred to Section C.4. We will also discuss the implication to the RL case using theorem 18.

C.1 Discrete MDPs

Recall π_e -concentratability coefficient is defined by

$$C^{\pi_e} = \max_{(s,a)} \frac{d_P^{\pi_e}(s,a)}{\rho(s,a)}.$$

Then, the error is calculated as follows.

Theorem 19 (Error of MLO for discrete MDPs).

- With probability $1 - \delta$, when $\lambda = \Omega(1)$,

$$\begin{aligned}
V_{P,c}^{\hat{\pi}_{\text{IL}}} - V_{P,c}^{\pi_e} &\leq \text{Err}_o + \text{Err}_e, \\
\text{Err}_o &= c_1 H^2 \log(|\mathcal{S}||\mathcal{A}|c_2/\delta) \left(\sqrt{\frac{C^{\pi_e} |\mathcal{S}|^2 |\mathcal{A}|}{n_o}} + \frac{C^{\pi_e} |\mathcal{S}| |\mathcal{A}|}{n_o} \right), \text{Err}_e = 2H \sqrt{\frac{\log(2|\mathcal{F}|/\delta)}{2n_e}},
\end{aligned}$$

where c_1 and c_2 are some universal constants.

- With probability $1 - \delta$, when $\lambda = \Omega(1)$,

$$V_{P,c}^{\hat{\pi}^{\text{RL}}} - V_{P,c}^{\pi_e} \leq c_1 H^2 \log(|\mathcal{S}||\mathcal{A}|c_2/\delta) \left(\sqrt{\frac{C^{\pi^*} |\mathcal{S}|^2 |\mathcal{A}|}{n_o}} + \frac{C^{\pi^*} |\mathcal{S}| |\mathcal{A}|}{n_o} \right), \quad C^{\pi^*} = \max_{(s,a)} \frac{d_P^{\pi^*}(s,a)}{\rho(s,a)}. \quad (6)$$

where c_1 and c_2 are some universal constants.

The quantity C^{π_e} measures the difference of distributions between the expert and the batch data. This is much smaller than the common concentratability coefficients in offline RL:

$$\max_{\pi \in \Pi} \max_{(s,a) \in \mathcal{S} \times \mathcal{A}} \frac{d_P^\pi(s,a)}{\rho(s,a)}, \quad \frac{1}{\min_{(s,a)} \rho(s,a)},$$

which measure the worst discrepancy between all policies in Π and the batch data [80]. These assumptions imply ρ has global coverage. We achieve this better bound via pessimism. In the RL case, the similar bound as (6) has been obtained in offline policy optimization based on FQI [54]. However, their work is limited to a tabular case. Hereafter, we will show our result is extended to more general continuous MDPs.

C.2 KNRs

As in Proposition 13, $\sigma(s,a)$ is given by $\beta_{n_o}/\zeta \|\phi(s,a)\|_{\Sigma_{n_o}^{-1}}$. Thus, from Theorem 3, the final error bound of $\hat{V}_{P,c}^{\hat{\pi}^{\text{IL}}} - V_{P,c}^{\pi_e}$ is

$$(6H^2 + 2H) \min(\mathbb{E}_{(s,a) \sim d_P^{\pi_e}} [\beta_{n_o}/\zeta \|\phi(s,a)\|_{\Sigma_{n_o}^{-1}}], 1) + 2H \sqrt{\log(2|\mathcal{F}|/\delta)/(2n_e)}.$$

Hereafter, we analyze β_{n_o} and $\mathbb{E}_{(s,a) \sim d_P^{\pi_e}} [\|\phi(s,a)\|_{\Sigma_{n_o}^{-1}}]$.

Analysis of information gain First, we analyze β_{n_o} . We need to upper-bound the information gain $\bar{\mathcal{I}}_{n_o}$ in β_{n_o} . Recall $\Sigma_\rho = \mathbb{E}_{(s,a) \sim \rho} [\phi(s,a)\phi^\top(s,a)]$ and $\phi(s,a) \in \mathbb{R}^d$.

Theorem 20 (Finite sample analysis of information gain in finite-dimensional linear models). *Assume $\|\phi(s,a)\|_2 \leq 1 \forall (s,a) \in \mathcal{S} \times \mathcal{A}$. Let c_1, c_2 be universal constants.*

1. When $\lambda = \Omega(1)$, with probability $1 - \delta$, we have

$$\bar{\mathcal{I}}_{n_o} = \log(\det(\Sigma_{n_o}/\lambda \mathbf{I})) \leq c_1 \text{rank}(\Sigma_\rho) \{ \text{rank}(\Sigma_\rho) + \log(c_2/\delta) \} \log(1 + n_o).$$

2. When $\lambda = \Omega(1)$ and $\zeta^2 = \Omega(1)$, With probability $1 - \delta$, we have

$$\beta_{n_o} \leq c_1 \sqrt{\|W^*\|_2 + d_S \text{rank}(\Sigma_\rho) \{ \text{rank}(\Sigma_\rho) + \log(c_2/\delta) \} \log(1 + n_o)}.$$

Theorem 20 states $\bar{\mathcal{I}}_{n_o} = O(\text{rank}[\Sigma_\rho]^2 \log(n_o))$. We highlight the novelty of our analysis comparing to the other literature. [60] analyze the expectation of the information gain in a fixed or random design setting. Following their discussion, we can prove

$$\mathbb{E}[\bar{\mathcal{I}}_{n_o}] \leq \text{rank}(\Sigma_\rho) \log(1 + n_o)$$

as Theorem 42 by Jensen's inequality. Going beyond the expectation, we derive the finite-sample result by leveraging the variational representation and the uniform law with localization in Lemma 39. The finite-sample analysis is much harder than calculating the bound of the expectation.

The worse case of $\bar{\mathcal{I}}_{n_o}$ referred to as the maximum information gain has been often used in online learning [63; 1; 32]. From their discussion, we always have $\bar{\mathcal{I}}_{n_o} = O(d \log(n))$. Here, we show that the information gain can be upper-bounded more tightly when $\text{rank}[\Sigma_\rho]^2 \leq d$ in offline RL (a random design setting). Comparing to the analysis of maximum information gain, our analysis takes the low-rankness of the design matrix Σ_ρ into consideration by fully utilizing the random design setting assumption.

Analysis of $\mathbb{E}_{(s,a) \sim d_P^{\pi_e}} [\|\phi(s, a)\|_{\Sigma_{n_o}^{-1}}]$ and the final bound

Next, we analyze $\mathbb{E}_{(s,a) \sim d_P^{\pi_e}} [\|\phi(s, a)\|_{\Sigma_{n_o}^{-1}}]$.

Theorem 21. *Suppose $\lambda = \Omega(1), \zeta^2 = \Omega(1), \|W^*\|_2 = \Omega(1)$. Let c_1, c_2 be some universal constants.*

1. With probability $1 - \delta$,

$$\mathbb{E}_{(s,a) \sim d_P^{\pi_e}} [\|\phi(s, a)\|_{\Sigma_{n_o}^{-1}}] \leq c_1 \sqrt{\frac{C^{\pi_e} \text{rank}[\Sigma_\rho] \{\text{rank}[\Sigma_\rho] + \log(c_2/\delta)\}}{n_o}}, \quad C^{\pi_e} = \sup_{x \in \mathbb{R}^d} \left(\frac{x^\top \Sigma_{\pi_e} x}{x^\top \Sigma_\rho x} \right),$$

where $\Sigma_{\pi_e} = \mathbb{E}_{(s,a) \sim d_P^{\pi_e}} [\phi(s, a)\phi(s, a)^\top]$.

2. With probability $1 - \delta$,

$$V_{P,c}^{\hat{\pi}_{\text{RL}}} - V_{P,c}^{\pi_e} \leq \text{Err}_o + \text{Err}_e, \quad \bar{R} = \text{rank}[\Sigma_\rho] \{\text{rank}[\Sigma_\rho] + \log(c_2/\delta)\}, \quad (7)$$

$$\text{Err}_o = c_1 H^2 \min(d^{1/2}, \bar{R}) \sqrt{\bar{R}} \sqrt{\frac{d_S C^{\pi_e} \log(1 + n_o)}{n_o}},$$

$$\text{Err}_e = 2H \sqrt{\log(2|\mathcal{F}|/\delta)/(2n_e)}.$$

3. With probability $1 - \delta$, let $C^{\pi^*} = \sup_{x \in \mathbb{R}^d} \left(\frac{x^\top \Sigma_{\pi^*} x}{x^\top \Sigma_\rho x} \right)$, $\Sigma_{\pi^*} = \mathbb{E}_{(s,a) \sim d_P^{\pi^*}} [\phi(s, a)\phi^\top(s, a)]$.

Then, we have

$$V_{P,c}^{\hat{\pi}_{\text{RL}}} - V_{P,c}^{\pi^*} \leq c_1 H^2 \{\text{rank}(\Sigma_\rho) + \log(c_2/\delta)\} \text{rank}(\Sigma_\rho) \sqrt{\frac{d_S C^{\pi^*} \log(1 + n_o)}{n_o}}.$$

The final bound (7) suggests Err_o is $\tilde{O}(H^2 \text{rank}[\Sigma_\rho]^2 \sqrt{d_S C^{\pi_e}/n_o})$. We can also get $\tilde{O}(H^2 \text{rank}[\Sigma_\rho] d^{1/2} \sqrt{d_S C^{\pi_e}/n_o})$, which implies $\tilde{O}(H^2 d^{3/2} \sqrt{d_S C^{\pi_e}/n_o})$. In other words, when $C^{\pi_e}, \text{rank}[\Sigma_\rho]$ are not so large and the offline sample size n_o is large enough, $O(H \sqrt{\log(|\mathcal{F}|)/n_e})$ is a dominating term and the covariate shift problem in BC can be avoided since the horizon dependence is just H . Recall the known BC error bound is $O(H^2 \sqrt{\log|\Pi|/n_e})$ [4, Chapter 14].

We see the implication of Err_o in more details, which also corresponds to the error of RL case. The rate regarding n_o is $n_o^{-1/2}$, which is the standard rate in parametric regression. Besides, we can see the bound depends on $\text{rank}[\Sigma_\rho], C^{\pi_e}$. Importantly, since we always have $\text{rank}[\Sigma_\rho] \leq d$, our final bound captures the possible low-rankness of the batch data. The quantity C^{π_e} corresponds to π_e -concentrability coefficient (*-concentrability in the RL case). This is much smaller than the worst case concentrability coefficients:

$$\sup_{\pi \in \Pi} C^\pi, \quad \tilde{C} = \sup_{(s,a)} \|\phi(s, a)\|_2^2 \|\Sigma_\rho^{-1}\|_2.$$

Finally, we note the technical novelty by comparing it to the techniques developed in the offline RL literature. A quantity that is similar to $\mathbb{E}_{(s,a) \sim d_P^{\pi_e}} [\|\phi(s, a)\|_{\Sigma_{n_o}^{-1}}]$ has been analyzed in [31]⁶, which studies the error bound of FQI with pessimism in linear MDPs. [31, Corollary 4.5] assumes that full coverage, i.e., Σ_ρ is full-rank and has lower bounded eigenvalues. Also the number of offline samples n_o depends on the smallest eigenvalue. Our analysis just uses partial coverage with the refined concept of relative condition number and thus does not require the full rank assumption on Σ_ρ . Moreover, our bound is distribution dependent, i.e., it depends on $\text{rank}[\Sigma_\rho]$ rather than the ambient dimension of the feature vector ϕ . Thus the bound is much tighter for benign cases where the offline data from ρ happens to concentrate on a low-dimensional subspace. Beyond model-based offline RL literature, one can potentially adapt the model-free offline policy evaluation results (e.g., [20; 74]) with linear function approximation to offline policy optimization (without pessimism). Such model-free results will also incur $\sup_{\pi \in \Pi} C^\pi, \tilde{C}$ and the ambient dimension d , instead of much more refined quantities C^{π_e} and $\text{rank}[\Sigma_\rho]$.

⁶They analyze $\mathbb{E}_{(s,a) \sim d_P^{\pi^*}} [\|\phi(s, a)\|_{\Sigma_{n_o}^{-1}}]$, which also appears in our RL result theorem 21.

C.3 Gaussian processes

In this section, we give details on GPs. Note that prior works on model-free and model-based offline IL do not have results for infinite-dimensional non-parametric models. Thus our techniques developed in this section are new and relevant even to the offline RL literature—a point that we will return to at the end of this section.

From [Theorem 3](#), the final error is

$$(6H^2 + 2H) \min(\mathbb{E}_{x \sim d_P^{\pi_e}} [\beta_{n_o} / \zeta \sqrt{k_{n_o}(x, x)}], 1) + 2H \sqrt{\log(2|\mathcal{F}|/\delta)/(2n_e)}.$$

where $x = (s, a)$. Hereafter, we analyze β_{n_o} and $\mathbb{E}_{x \sim d_P^{\pi_e}} [\sqrt{k_{n_o}(x, x)}]$. Before going into the details, we repeat several important notations below.

In this section, following [\[63\]](#), for simplicity, we suppose the following:

Assumption 22. $k(x, x) \leq 1, \forall x \in \mathcal{S} \times \mathcal{A}$. $k(\cdot, \cdot)$ is a continuous and positive semidefinite kernel. $\mathcal{S} \times \mathcal{A}$ is a compact space.

Recall we denote $x := (s, a)$ and we have orthonormal eigenfunctions and eigenvalues $\{\psi_i, \mu_i\}_{i=1}^{\infty}$ by Mercer's theorem. We denote the feature mapping $\phi(x) := [\sqrt{\mu_1} \psi_1(x), \dots, \sqrt{\mu_{\infty}} \psi_{\infty}(x)]^{\top}$.

Assume eigenvalues $\{\mu_1, \dots, \mu_{\infty}\}$ is in non-increasing order, we recall the effective dimension:

$$d^* = \min\{j \in \mathbb{N} : j \geq B(j+1)n_o/\zeta^2\}, \quad B(j) = \sum_{k=j}^{\infty} \mu_k.$$

We also introduce the empirical version of d^* , where $\hat{\mu}_i$ are eigenvalues of the gram matrix \mathbf{K}_{n_o} .

Definition 23 (Empirical effective dimension). $\hat{d} = \min\{j \in \mathbb{N} : j \geq B(j+1)/\zeta^2, \hat{B}(j) = \sum_{k=j}^{n_o} \hat{\mu}_k\}$.

Hereafter, for simplicity, we treat $\zeta^2 = 1$, that is, $\zeta^2 = \Omega(1)$. Then, since $n_o \leq B(n_o+1)n_o/\zeta^2$, we have $d^* \leq n_o$.

The effective dimensions \hat{d} and d^* are widely used in machine learning literature. The first quantity d^* is often referred to as the degree of freedom [\[85; 7\]](#). In finite-dimensional linear kernels $\{x \mapsto a^{\top} \phi(x), a \in \mathbb{R}^d\}$ ($k(x, x) = \phi^{\top}(x) \phi(x)$), d^* is $\text{rank}[\mathbb{E}_{x \sim \rho}[\phi(x) \phi^{\top}(x)]]$. Thus, d^* is considered to be a natural extension of $\text{rank}[\mathbb{E}_{x \sim \rho}[\phi(x) \phi^{\top}(x)]]$ to infinite-dimensional models. The worst case of the second quantity:

$$\max_{\{x_1 \in \mathcal{S} \times \mathcal{A}, \dots, x_{n_o} \in \mathcal{S} \times \mathcal{A}\}} \hat{d}$$

is often used in online learning literature [\[72; 28\]](#). Up to logarithmic factors, it is equal to the maximum information gain [\[63\]](#):

$$\max_{\{x_1 \in \mathcal{S} \times \mathcal{A}, \dots, x_{n_o} \in \mathcal{S} \times \mathcal{A}\}} \log \det(\mathbf{I} + \mathbf{K}_{n_o}).$$

as shown in [\[13; 72\]](#). Importantly, as we will see soon since our setting is offline (a random design setting), \hat{d} can be upper-bounded much tightly than their analysis.

Analysis of information gain With the above in mind, we first analyze β_{n_o} . To do that, we need to bound the information gain \mathcal{I}_{n_o} . From [\[60, Lemma 1\]](#), we can easily prove

$$\mathbb{E}[\mathcal{I}_{n_o}] \leq \log(1 + n_o)d^*.$$

as in [Theorem 43](#). Going beyond the expectation, we derive the finite-sample error bound.

Theorem 24 (Finite sample analysis of information gain in infinite-dimensional models). *Suppose [Assumption 22](#). Let c_1 and c_2 be universal constants.*

1. We have

$$\mathcal{I}_{n_o} = \log(\det(\mathbf{I} + \zeta^{-2} \mathbf{K}_{n_o})) \leq 2\hat{d}\{\log(1 + n_o/\zeta^2) + 1\}. \quad (8)$$

2. When $\zeta^2 = \Omega(1)$, with probability $1 - \delta$,

$$\mathcal{I}_{n_o} = \log(\det(\mathbf{I} + \zeta^{-2}\mathbf{K}_{n_o})) \leq c_1\{d^* + \log(c_2/\delta)\}d^* \log(1 + n_o).$$

3. When $\zeta^2 = \Omega(1)$, with probability $1 - \delta$,

$$\beta_{n_o} \leq c_1 \sqrt{d_S \log^3(c_2 d_S n_o / \delta) \{d^* + \log(c_2/\delta)\} d^* \log(1 + n_o)}.$$

Theorem 24 states $\mathcal{I}_{n_o} = O((d^*)^2 \log(n_o))$. Our bound in the offline (a random design) setting can be much tighter compared to the online setting, that is, the known upper bound of maximum information gain in [63] though we can always use this as the bound of \mathcal{I}_{n_o} with probability 1. We can see this situation in linear kernels as we see in the previous section. In d -linear dimensional linear kernels, the maximum information gain is d . On the other hand, $\{d^*\}^2 = \text{rank}[\Sigma_\rho]^2$ can be much smaller than d .

Analysis of learning curves and the final bound We bound $\mathbb{E}_{x \sim d_P^{\pi_e}}[\sqrt{k_{n_o}(x, x)}]$, where

$$k_{n_o}(x, x') = k(x, x') - \bar{k}_{n_o}(x)^\top (\mathbf{K}_{n_o} + \zeta^2 \mathbf{I})^{-1} \bar{k}_{n_o}(x'), \{x_i\}_{i=1}^{n_o} \sim \rho(x).$$

where $x = (s, a)$.

Recall the definition of eigenvalues $\{\mu_i\}$ and eigenfunctions $\{\psi_i\}$ (which are orthonormal), we define the feature mapping $\phi(x) = [\sqrt{\mu_1}\psi_1(x), \dots, \sqrt{\mu_\infty}\psi_\infty(x)]^\top$. Denote $\Phi \in \mathbb{R}^{n_o \times \infty}$ as a matrix where each row of Φ corresponds to $\phi(x_i)$. Since $k(x, x') = \phi(x)^\top \phi(x')$, we can rewrite the kernel $k_{n_o}(x, x)$ as follows:

$$\begin{aligned} k_{n_o}(x, x) &= \phi(x)^\top \phi(x) - \phi(x)^\top \Phi^\top (\Phi \Phi^\top + \zeta^2 \mathbf{I})^{-1} \Phi \phi(x) \\ &= \phi(x)^\top \left[\mathbf{I} - \Phi^\top (\Phi \Phi^\top + \zeta^2 \mathbf{I})^{-1} \Phi \right] \phi(x) \\ &= \phi(x)^\top (\mathbf{I} + \zeta^{-2} \Phi^\top \Phi)^{-1} \phi(x) \\ &= \phi(x)^\top \Sigma_{n_o}^{-1} \phi(x), \end{aligned}$$

where $\Sigma_{n_o} := \mathbf{I} + \zeta^{-2} \sum_{i=1}^{n_o} \phi(x_i) \phi(x_i)^\top$, and we use matrix inverse lemma in the third equality. Note the infinite-dimensional inverse lemma is formalized in the proof.

Now we can use the relative condition number definition and Lemma 30 for a distribution change, i.e.,

$$\begin{aligned} \mathbb{E}_{x \sim d_P^{\pi_e}}[\sqrt{k_{n_o}(x, x)}] &\leq \sqrt{\mathbb{E}_{x \sim d_P^{\pi_e}}[k_{n_o}(x, x)]} \\ &= \sqrt{\text{tr} \left(\mathbb{E}_{x \sim d_P^{\pi_e}} \phi(x) \phi(x)^\top \Sigma_{n_o} \right)} \leq \sqrt{C^{\pi_e} \text{tr} \left(\mathbb{E}_{x \sim \rho} \phi(x) \phi(x)^\top \Sigma_{n_o} \right)} = \sqrt{C^{\pi_e} \mathbb{E}_{x \sim \rho} k_{n_o}(x, x)}, \end{aligned}$$

where

$$C^{\pi_e} = \sup_{\|x\|_2 \leq 1} \frac{x \Sigma_{\pi_e} x}{x \Sigma_\rho x}, \quad \Sigma_{\pi_e} = \mathbb{E}_{x \sim d_P^{\pi_e}}[\phi(x) \phi(x)^\top], \quad \Sigma_\rho = \mathbb{E}_{x \sim \rho}[\phi(x) \phi(x)^\top].$$

Now we only need to focus on analyzing $\mathbb{E}_{x \sim \rho}[k_{n_o}(x, x)]$.

Before proceeding to the analysis, we introduce the critical radius [9]. Given some function class \mathcal{F} , consider the localized population Rademacher complexity:

$$\mathcal{R}_n(\delta; \mathcal{F}) = \mathbb{E} \left[\sup_{f \in \mathcal{F}, \mathbb{E}_{x \sim \rho}[f^2(x)] \leq \delta} \left| \frac{1}{n_o} \sum_{i=1}^{n_o} \epsilon_i f(x_i) \right| \right]$$

where $\{x_i\}$ are i.i.d samples following $\rho(x)$ and $\{\epsilon_i\}$ are i.i.d Rademacher variables taking values in $\{-1, +1\}$ equiprobably, independent of the sequence $\{x_i\}$. The critical radius is defined as the minimum solution to

$$\mathcal{R}_n(\xi; \mathcal{F}) \leq \xi^2 / b$$

w.r.t ξ where b is a value s.t. $\|\mathcal{F}\|_\infty \leq b$.

Theorem 25. Suppose Assumption 22. Let c_1 and c_2 be universal constants.

1. Let δ_{n_o} be the critical radius of the function class $\{f : f \in \mathcal{H}_k, \|f\|_k \leq 1\}$. With probability $1 - \delta$,

$$\mathbb{E}_{x \sim d_P^{\pi_e}} [\sqrt{k_{n_o}(x, x)}] \leq c_1 \zeta \delta'_{n_o} \sqrt{C^{\pi_e} d^*},$$

$$\text{where } \delta'_{n_o} = \delta_{n_o} + \sqrt{\log(c_2/\delta)/n_o}.$$

2. Assume $\zeta^2 = \Omega(1)$. With probability $1 - \delta$,

$$\delta_{n_o} \leq c_1 \sqrt{d^*/n_o}, \quad \mathbb{E}_{x \sim d_P^{\pi_e}} [\sqrt{k_{n_o}(x, x)}] \leq c_1 \sqrt{\frac{C^{\pi_e} d^* \{d^* + \log(c_2/\delta)\}}{n_o}}.$$

3. Assume $\zeta^2 = \Omega(1)$. With probability $1 - \delta$,

$$V_{P,c}^{\hat{\pi}_{\text{RL}}} - V_{P,c}^{\pi_e} \leq \text{Err}_o + \text{Err}_e \tag{9}$$

$$\text{Err}_o = c_1 H^2 \{d^* + \log(c_2/\delta)\} d^* \sqrt{\frac{d_{\mathcal{S}} C^{\pi_e} \log^3(c_2 d_{\mathcal{S}} n_o / \delta) \log(1 + n_o)}{n_o}}$$

$$\text{Err}_e = 2H \sqrt{\log(2|\mathcal{F}|/\delta)/(2n_e)}.$$

4. Assume $\zeta^2 = \Omega(1)$. For offline RL, with probability $1 - \delta$,

$$V_{P,c}^{\hat{\pi}_{\text{RL}}} - V_{P,c}^{\pi^*} \leq c_1 H^2 \{d^* + \log(c_2/\delta)\} d^* \sqrt{\frac{d_{\mathcal{S}} C^{\pi^*} \log^3(c_2 d_{\mathcal{S}} n_o / \delta) \log(1 + n_o)}{n_o}},$$

$$\text{where } C^{\pi^*} = \sup_{\|x\|_2 \leq 1} \frac{x \Sigma_{\pi^*} x}{x \Sigma_{\rho} x}.$$

The final bound in (9) suggests that Err_o is $\tilde{O}(H^2 \{d^*\}^2 \sqrt{d_{\mathcal{S}} C^{\pi_e} / n_o})$. In other words, when C^{π_e} , d^* are not so large and the offline sample size is large enough, Err_e dominates Err_o and the covariate shift problem in BC can be avoided since the horizon dependence is just H . Our bound is the natural extension of Theorem 21 to possibly infinite dimensional models.

The first and second statements in Theorem 25 are mainly proved in two steps: formulating $k_{n_o}(x, x)$ into the variational representation and utilizing the uniform law with localization. Note the critical radius can be upper-bounded more tightly than $O(\sqrt{d^*/n_o})$ depending on the kernels. Besides, C^{π_e} can be replaced with a tighter quantity:

$$\max_{i \in \mathbb{N}} \mathbb{E}_{(s,a) \sim d_P^{\pi_e}} [\psi_i^2(s, a)].$$

Since $\mathbb{E}_{(s,a) \sim \rho} [\psi_i^2(s, a)] = 1$, this quantity also measure the difference of batch data and expert. This is less than C^{π_e} noting that $\frac{x \Sigma_{\pi^*} x}{x \Sigma_{\rho} x} = \mathbb{E}_{(s,a) \sim d_P^{\pi_e}} [\psi_i^2(s, a)]$ when x is a vector s.t. only i -th element is 1 and the other elements are 0. The third statement in Theorem 25 is directly proved by combining the second statement in Theorem 25 and Theorem 24.

Implication to offline RL The final statement in Theorem 25 is the bound for the RL case. This is the first result showing the error bound for pessimistic offline RL with nonparametric models. As related literature, in model-free offline RL, [70; 21] obtained the finite-sample error bounds characterized by the critical radius for some minimax-type estimators called Modified RBM [5]. As we show in Theorem 25, since the critical radius of an RKHS ball is upper-bounded by the effective dimension d^* , their bounds are also characterized by the effective dimension. On top of that, several papers derived the bounds under the general function approximation setting: FQI [24; 21; 45; 16], marginal weighting based estimators [69], DICE methods [84; 46], policy based methods [42; 43] and MABO [77]. Comparing to our result, all of their bounds depend on

$$\sup_{\pi \in \Pi} \sup_{(s,a)} \frac{d_P^{\pi}(s, a)}{\rho(s, a)} \quad \text{or} \quad \sup_{(s,a)} \frac{1}{\rho(s, a)}.$$

The pessimistic bonus allows us to obtain the bound only depending on C^{π^*} but not the above constants. Besides, our C^{π^*} in [Theorem 25](#) is more refined quantity than the density ratios in the sense that it is defined in terms of the relative condition number. Note we can easily obtain the statements which replace C^{π^*} in [Theorem 25](#) with $\frac{d_P^{\pi^*}(s,a)}{\rho(s,a)}$.

Remark 26 (Relation with more general offline RL literature). *Due to the lack of exploration, it is known how to deal with the lack of the coverage of the offline data is a challenging problem [83; 74]. We use the penalty terms based on model-based RL. In the above, we explain how the penalty term in MIL0 (and its RL counterpart) is transferred to the final sample-error bounds. The idea of penalization has been utilized in a variety of other ways in offline RL. The first other way is imposing constraints on the policy class or Q-function class so that estimated policies are not too much far away from behavior policies. For example, we can use KL divergences, MMD distance, Wasserstein distance to measure the distance from behavior policies [76; 23; 44; 68; 26] and add $D(\pi, \pi_b)$ as penalty terms, where π_b is a behavior policy. Another way is explicitly estimating the lower bound of q-functions [39; 81; 82]. By doing so, we can avoid the overestimation of the q-functions in unknown (non-covered) regions.*

Remark 27 (Relation with GP literature). *The quantity $\mathbb{E}_{x \sim \rho(x)}[k_{n_o}(x, x)]$ is often referred to as the learning curve in GP literature [75; 61; 55]. Their analysis mainly focuses on the numerical viewpoints, that is, how to approximately calculate $\mathbb{E}_{x \sim \rho(x)}[k_{n_o}(x, x)]$. Though [41] analyzes the convergence property, their analysis is limited to the expectation and the result is asymptotic. As far as we know, our result is the first result showing the finite-sample error rate.*

Remark 28 (Duality between KNRs and GPs). *KNRs and GPs have a primal and dual relationship via Mercer's theorem. In fact, as we see, $k(\cdot, \cdot) = \langle \phi(\cdot), \phi(\cdot) \rangle$, we have $k_{n_o}(x, x) = \phi(x)^\top \Sigma_{n_o}^{-1} \phi(x)$. Thus, our result in GPs can be applied to the result for infinite-dimensional KNRs with $\phi : \mathcal{S} \times \mathcal{A} \mapsto \mathcal{H}$ where \mathcal{H} is some RKHS.*

Remark 29 (Online RL using RKHS). *There are several online RL literature using RKHS such as the model-based way [13] like our work and the model-free way [3; 78; 19]. In both cases, the final-sample error bounds incur the maximum information gain, i.e., a worse case quantity which is distribution independent. Comparing to that, our final bounds use distribution-dependent quantities d^* .*

C.4 Missing Proofs

Below, we provide missing proofs for tabular MDPs, KNRs, and non-parametric GP models.

C.4.1 Missing proofs for tabular result

We start by providing proof of the tabular MDP result.

Proof of [Theorem 19](#). We use [Theorem 3](#). Then, we have

$$V_{P,c}^{\hat{\pi}^e} - V_{P,c}^{\pi^e} \leq (6H^2 + 2H) \min(1, \mathbb{E}_{(s,a) \sim d_P^{\pi^e}}[\sigma(s, a)]) + H\epsilon_{\text{stat}}.$$

Hereafter, we show how to upper-bound $\mathbb{E}_{(s,a) \sim d_P^{\pi^e}}[\sigma(s, a)]$. We use [Lemma 35](#). Then, by letting $\xi = c_1 \log(|\mathcal{S}||\mathcal{A}|c_2/\delta)$, with probability $1 - \delta$, we have

$$\frac{1}{N(s, a) + \lambda} \leq \frac{\xi}{n_o \rho(s, a) + \lambda} \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}.$$

We condition on the above event. Then,

$$\begin{aligned} \mathbb{E}_{(s,a) \sim d_P^{\pi^e}}[\sigma(s, a)] &\leq \mathbb{E}_{(s,a) \sim d_P^{\pi^e}} \left[\sqrt{\frac{|\mathcal{S}| \log 2 + \log(2|\mathcal{S}||\mathcal{A}|/\delta)}{2\{N(s, a) + \lambda\}}} + \frac{\lambda}{N(s, a) + \lambda} \right] \\ &\leq \sqrt{\mathbb{E}_{(s,a) \sim d_P^{\pi^e}} \left[\frac{|\mathcal{S}| \log 2 + \log(2|\mathcal{S}||\mathcal{A}|/\delta)}{2\{N(s, a) + \lambda\}} \right]} + \mathbb{E}_{(s,a) \sim d_P^{\pi^e}} \left[\frac{\lambda}{N(s, a) + \lambda} \right]. \end{aligned}$$

From [Lemma 35](#), we have

$$\begin{aligned}
\mathbb{E}_{(s,a) \sim d_P^{\pi_e}}[\sigma(s,a)] &\leq \sqrt{\xi \mathbb{E}_{(s,a) \sim d_P^{\pi_e}} \left[\frac{|\mathcal{S}| \log 2 + \log(2|\mathcal{S}||\mathcal{A}|/\delta)}{\{n_o \rho(s,a) + \lambda\}} \right]} + \mathbb{E}_{(s,a) \sim d_P^{\pi_e}} \left[\frac{\lambda \xi}{n_o \rho(s,a) + \lambda} \right] \\
&\leq \sqrt{\xi C^{\pi_e} \mathbb{E}_{(s,a) \sim \rho} \left[\frac{|\mathcal{S}| \log 2 + \log(2|\mathcal{S}||\mathcal{A}|/\delta)}{\{n_o \rho(s,a) + \lambda\}} \right]} + C^{\pi_e} \mathbb{E}_{(s,a) \sim \rho} \left[\frac{\lambda \xi}{n_o \rho(s,a) + \lambda} \right] \\
&\leq \sqrt{\xi C^{\pi_e} \sum_{s,a} \left[\frac{\{|\mathcal{S}| \log 2 + \log(2|\mathcal{S}||\mathcal{A}|/\delta)\} \rho(s,a)}{\{n_o \rho(s,a) + \lambda\}} \right]} + C^{\pi_e} \sum_{s,a} \left[\frac{\rho(s,a) \lambda \xi}{n_o \rho(s,a) + \lambda} \right] \\
&\leq \sqrt{\xi C^{\pi_e} \{|\mathcal{S}| \log 2 + \log(2|\mathcal{S}||\mathcal{A}|/\delta)\} |\mathcal{S}||\mathcal{A}|/n_o} + \lambda C^{\pi_e} \xi |\mathcal{S}||\mathcal{A}|/n_o.
\end{aligned}$$

where again

$$C^{\pi_e} = \max_{(s,a)} \frac{d_P^{\pi_e}(s,a)}{\rho(s,a)}.$$

This concludes the proof. \square

C.4.2 Missing proofs for KNR results

Next we move to provide proofs for the KNR results.

Proof of Theorem 20. In the proof, we use two statements, eq. (11) and eq. (12), in the proof of theorem 21. We recommend readers to read the proof of theorem 21 first.

We denote the eigenvalues of $\sum_{i=1}^{n_o} \phi(s_i, a_i) \phi^\top(s_i, a_i)$ by $\{\hat{\mu}_i\}_{i=1}^d$ s.t. $\hat{\mu}_1 \geq \hat{\mu}_2 \geq \dots$. Since we assume $\|\phi(s, a)\|_2 \leq 1$, we have $\hat{\mu}_1 \leq n_o$.

First step We first show

$$\log(\det(\Sigma_{n_o}) / \det(\lambda \mathbf{I})) \leq \text{tr} \left[\Sigma_{n_o}^{-1} \sum_{i=1}^{n_o} \phi(s_i, a_i) \phi^\top(s_i, a_i) \right] \{\log(1 + n_o/\lambda) + 1\}.$$

Note this directly shows $\log(\det(\Sigma_{n_o}) / \det(\lambda \mathbf{I})) \leq d \log(1 + n_o/\lambda)$, $\phi(s, a) \in \mathbb{R}^d$. The above is proved as follows:

$$\begin{aligned}
\log(\det(\Sigma_{n_o}) / \det(\lambda \mathbf{I})) &= \sum_{i=1}^d \log \left(1 + \frac{\hat{\mu}_i}{\lambda} \right) = \sum_{i=1}^d \log \left(1 + \frac{\hat{\mu}_i}{\lambda} \right) \frac{\hat{\mu}_i/\lambda + 1}{\hat{\mu}_i/\lambda + 1} \\
&= \sum_{i=1}^d \log \left(1 + \frac{\hat{\mu}_i}{\lambda} \right) \frac{\hat{\mu}_i/\lambda}{\hat{\mu}_i/\lambda + 1} + \log \left(1 + \frac{\hat{\mu}_i}{\lambda} \right) \frac{1}{\hat{\mu}_i/\lambda + 1} \\
&\leq \log \left(1 + \frac{\hat{\mu}_1}{\lambda} \right) \sum_{i=1}^d \frac{\hat{\mu}_i/\lambda}{\hat{\mu}_i/\lambda + 1} + \sum_{i=1}^d \frac{\hat{\mu}_i/\lambda}{\hat{\mu}_i/\lambda + 1} \quad (\log(1+x) < x) \\
&\leq \{\log(1 + n_o/\lambda) + 1\} \sum_{i=1}^d \frac{\hat{\mu}_i/\lambda}{\hat{\mu}_i/\lambda + 1} \quad (\hat{\mu}_1 \leq n_o) \\
&= \{\log(1 + n_o/\lambda) + 1\} \text{tr} \left[\Sigma_{n_o}^{-1} \sum_{i=1}^{n_o} \phi(s_i, a_i) \phi^\top(s_i, a_i) \right].
\end{aligned}$$

In the last line, letting UVU^\top be the eigendecomposition of $\sum_{i=1}^{n_o} \phi(s_i, a_i) \phi^\top(s_i, a_i)$, we use

$$\text{tr} \left[\Sigma_{n_o}^{-1} \sum_{i=1}^{n_o} \phi(s_i, a_i) \phi^\top(s_i, a_i) \right] = \text{tr} [\{V + \lambda \mathbf{I}\}^{-1} V] = \sum_{i=1}^d \frac{\hat{\mu}_i/\lambda}{\hat{\mu}_i/\lambda + 1}.$$

Then, the first statement is proved.

Second step Next, we prove the second statement. We have

$$\text{tr} \left[\Sigma_{n_o}^{-1} \sum_{i=1}^{n_o} \phi(s_i, a_i) \phi^\top(s_i, a_i) \right] = \sum_{i=1}^{n_o} \phi^\top(s_i, a_i) \Sigma_{n_o}^{-1} \phi(s_i, a_i).$$

Then, from (11), with probability $1 - \delta$,

$$\sum_{i=1}^{n_o} \phi^\top(s_i, a_i) \Sigma_{n_o}^{-1} \phi(s_i, a_i) \lesssim c_1 \{ \text{rank}[\Sigma_\rho] + \log(c_2/\delta) \} \sum_{i=1}^{n_o} \phi^\top(s_i, a_i) \{ n_o \Sigma_\rho + \lambda \mathbf{I} \}^{-1} \phi(s_i, a_i). \quad (10)$$

Hereafter, we condition on the above event. To upper-bound $\sum_{i=1}^{n_o} \|\phi(s_i, a_i)\|_{\{n_o \Sigma_\rho + \lambda \mathbf{I}\}^{-1}}^2$, we use Bernstein's inequality:

$$\begin{aligned} & \left| \sum_{i=1}^{n_o} \phi^\top(s_i, a_i) \{ n_o \Sigma_\rho + \lambda \mathbf{I} \}^{-1} \phi(s_i, a_i) - n_o \mathbb{E}_{(s,a) \sim \rho} [\phi^\top(s, a) \{ n_o \Sigma_\rho + \lambda \mathbf{I} \}^{-1} \phi(s, a)] \right| \\ & \lesssim \sqrt{n_o \text{Var}_{(s,a) \sim \rho} [\phi^\top(s, a) \{ n_o \Sigma_\rho + \lambda \mathbf{I} \}^{-1} \phi(s, a)]} + 1/\lambda. \end{aligned}$$

since $\|\phi(s, a)\|_{\{n_o \Sigma_\rho + \lambda \mathbf{I}\}^{-1}}^2 \leq 1/\lambda \forall (s, a) \in \mathcal{S} \times \mathcal{A}$. Here, from (12),

$$n_o \mathbb{E}_{(s,a) \sim \rho} [\phi^\top(s, a) \{ n_o \Sigma_\rho + \lambda \mathbf{I} \}^{-1} \phi(s, a)] \leq \text{rank}[\Sigma_\rho].$$

Besides,

$$\begin{aligned} \text{Var}_{(s,a) \sim \rho} [\phi^\top(s, a) \{ n_o \Sigma_\rho + \lambda \mathbf{I} \}^{-1} \phi(s, a)] & \leq \mathbb{E}_{(s,a) \sim \rho} [\{\phi^\top(s, a) \{ n_o \Sigma_\rho + \lambda \mathbf{I} \}^{-1} \phi(s, a)\}^2] \\ & \leq 1/\lambda \mathbb{E}_{(s,a) \sim \rho} [\phi^\top(s, a) \{ n_o \Sigma_\rho + \lambda \mathbf{I} \}^{-1} \phi(s, a)] \\ & \leq \text{rank}[\Sigma_\rho] / (n_o \lambda). \end{aligned} \quad (\text{from (12)})$$

Thus,

$$\sum_{i=1}^{n_o} \phi^\top(s_i, a_i) \{ n_o \Sigma_\rho + \lambda \mathbf{I} \}^{-1} \phi(s_i, a_i) \lesssim \text{rank}[\Sigma_\rho] + \sqrt{\text{rank}[\Sigma_\rho] / \lambda} + 1/\lambda.$$

By combining (10) with the above, we have

$$\log(\det(\Sigma_{n_o}) / \det(\lambda \mathbf{I})) \leq c_1 \text{rank}(\Sigma_\rho) \{ \text{rank}(\Sigma_\rho) + \log(c_2/\delta) \} \log(1 + n_o c_3).$$

from $\lambda = \Omega(1)$. □

Before proving [Theorem 21](#), we first present some lemmas.

Lemma 30 (Distribution change). *Consider two distributions $\rho_1 \in \Delta(\mathcal{S} \times \mathcal{A})$ and $\rho_2 \in \Delta(\mathcal{S} \times \mathcal{A})$, and a feature mapping $\phi : \mathcal{S} \times \mathcal{A} \mapsto \mathcal{H}$ where \mathcal{H} is some Hilbert space (e.g., finite dimensional Euclidean space). Denote $C := \sup_{x \in \mathcal{H}} \frac{x^\top \mathbb{E}_{s,a \sim \rho_1} \phi(s,a) \phi(s,a)^\top x}{x^\top \mathbb{E}_{s,a \sim \rho_2} \phi(s,a) \phi(s,a)^\top x}$. Then for any positive definite linear matrix (operator) Λ , we have:*

$$\mathbb{E}_{s,a \sim \rho_1} \phi(s, a)^\top \Lambda \phi(s, a) \leq C \mathbb{E}_{s,a \sim \rho_2} \phi(s, a)^\top \Lambda \phi(s, a).$$

Proof. Denote the eigendecomposition of $\Lambda = U \Sigma U^\top$ where $\{\sigma_i, u_i\}$ as the eigenvalue-eigenvector pairs. We have:

$$\begin{aligned} \mathbb{E}_{s,a \sim \rho_1} \phi(s, a)^\top \Lambda \phi(s, a) & = \sum_{i=0}^{\infty} \sigma_i u_i^\top \mathbb{E}_{s,a \sim \rho_1} \phi(s, a) \phi(s, a)^\top u_i \\ & \leq \sum_{i=0}^{\infty} \sigma_i C u_i^\top \mathbb{E}_{s,a \sim \rho_2} \phi(s, a) \phi(s, a)^\top u_i \\ & = C \mathbb{E}_{s,a \sim \rho_2} \phi(s, a)^\top \Lambda \phi(s, a), \end{aligned}$$

which concludes the proof. □

Proof of Theorem 21. Here, we prove the first statement. We need to upper-bound

$$\mathbb{E}_{(s,a) \sim d_P^{\pi_e}} \left[\sqrt{\phi^\top(s, a) \Sigma_{n_o}^{-1} \phi(s, a)} \right].$$

As the first step, we use Jensen's inequality:

$$\mathbb{E}_{(s,a) \sim d_P^{\pi_e}} \left[\sqrt{\phi^\top(s, a) \Sigma_{n_o}^{-1} \phi(s, a)} \right] \leq \sqrt{\mathbb{E}_{(s,a) \sim d_P^{\pi_e}} [\phi^\top(s, a) \Sigma_{n_o}^{-1} \phi(s, a)]}.$$

Hereafter, we analyze $\mathbb{E}_{(s,a) \sim d_P^{\pi_e}} [\phi^\top(s, a) \Sigma_{n_o}^{-1} \phi(s, a)]$.

We first use the definition of the relative condition number C^{π_e} and Lemma 30 to change distribution from $d_P^{\pi_e}$ to ρ , i.e., via Lemma 30, we have:

$$\mathbb{E}_{s,a \sim d_P^{\pi_e}} \phi(s, a)^\top \Sigma_{n_o}^{-1} \phi(s, a) \leq C^{\pi_e} \mathbb{E}_{s,a \sim \rho} \phi(s, a)^\top \Sigma_{n_o}^{-1} \phi(s, a).$$

Thus, below we just need to bound $\mathbb{E}_{s,a \sim \rho} \phi(s, a)^\top \Sigma_{n_o}^{-1} \phi(s, a)$.

Concentration argument In this step, we consider how to bound $\mathbb{E}_{(s,a) \sim \rho} [\phi^\top(s, a) \Sigma_{n_o}^{-1} \phi(s, a)]$. To do that, we show with probability $1 - \delta$,

$$\phi^\top(s, a) \Sigma_{n_o}^{-1} \phi(s, a) \leq c_1 \{\text{rank}(\Sigma_\rho) + \log(c_2/\delta)\} \phi^\top(s, a) \{n_o \Sigma_\rho + \lambda \mathbf{I}\}^{-1} \phi(s, a) \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}. \quad (11)$$

We use the variational representation:

$$\begin{aligned} \phi^\top(s, a) \Sigma_{n_o}^{-1} \phi(s, a) &= \sup_{\{a \in \mathbb{R}^d : a^\top \Sigma_{n_o} a \leq 1\}} \{a^\top \phi(s, a)\}^2 \\ &= \sup_{\{a \in \mathbb{R}^d : a^\top \Sigma_{n_o} a \leq 1, \|a\|_2^2 \leq (1+\lambda)/\lambda, \|a^\top \phi\|_\infty \leq 1/\lambda\}} \{a^\top \phi(s, a)\}^2. \end{aligned}$$

Note that in the first line, we use

$$\sup_{\{a \in \mathbb{R}^d : a^\top \Sigma_{n_o} a \leq 1\}} a^\top \phi(s, a) = \sup_{\{b \in \mathbb{R}^d : b^\top b \leq 1\}} b^\top \Sigma_{n_o}^{-1/2} \phi(s, a) = \|\phi(s, a)\|_{\Sigma_{n_o}^{-1}}.$$

From the first line to the second line, we use the fact that the maximization regarding a is taken when $\tilde{a} = \Sigma_{n_o}^{-1} \phi(s, a) / \|\phi(s, a)\|_{\Sigma_{n_o}^{-1}}$ and

$$\begin{aligned} \|\tilde{a}\|_2^2 &= \phi^\top(s, a) \Sigma_{n_o}^{-2} \phi(s, a) / \{\phi^\top(s, a) \Sigma_{n_o}^{-1} \phi(s, a)\} = (n_o + \lambda) / \lambda^2, \\ \|\tilde{a}^\top \phi\| &\leq \|\phi(s, a)\|_{\Sigma_{n_o}^{-1}} \leq 1/\lambda \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}, \end{aligned}$$

noting $\|\phi(s, a)\|_2 \leq 1$. By defining $\bar{c} = (n_o + \lambda) / \lambda^2$, we have $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$,

$$\begin{aligned} \phi^\top(s, a) \Sigma_{n_o}^{-1} \phi(s, a) &= \sup_{\{a \in \mathbb{R}^d : a^\top \Sigma_{n_o} a \leq 1, \|a\|_2^2 \leq \bar{c}, \|a^\top \phi\|_\infty \leq 1/\lambda\}} \{a^\top \phi(s, a)\}^2 \\ &= \sup_{\{a \in \mathbb{R}^d : a^\top \lambda \mathbf{I} a + \sum_{i=1}^{n_o} \{a^\top \phi_i\}^2 \leq 1, \|a\|_2^2 \leq \bar{c}, \|a^\top \phi\|_\infty \leq 1/\lambda\}} \{a^\top \phi(s, a)\}^2. \end{aligned}$$

Next, we use Lemma 36, that is, with probability $1 - \delta$,

$$\frac{1}{n_o} \sum_{i=1}^{n_o} f^2(s_i, a_i) \geq 0.5 \mathbb{E}_{(s,a) \sim \rho} [f^2(s, a)] - 0.5 \{\delta'_{n_o}\}^2 \forall f \in \mathcal{F},$$

where

$$\mathcal{F} = \{(s, a) \mapsto a^\top \phi(s, a) : a^\top \Sigma_{n_o} a \leq 1, \|a\|_2^2 \leq \bar{c}, \|a^\top \phi\|_\infty \leq 1/\lambda, a \in \mathbb{R}^d\}.$$

Here, $\delta'_{n_o} = \delta_{n_o} + \sqrt{\log(c_2/\delta)/n_o}$, where δ_{n_o} is the critical radius of the function class \mathcal{F} . Noting $\lambda = \Omega(1)$, from Lemma 37, $\delta'_{n_o} = c_1 \sqrt{\text{rank}[\Sigma_\rho]/n_o} + \sqrt{\log(c_2/\delta)/n_o}$. By conditioning on the

above event, $\forall (s, a) \in \mathcal{S} \times \mathcal{A}$, we have

$$\begin{aligned}
\|\phi(s, a)\|_{\Sigma_{n_o}^{-1}}^2 &\leq \sup_{\{a \in \mathbb{R}^d: a^\top \lambda \mathbf{I} a + 0.5 n_o \mathbb{E}_{(s,a) \sim \rho} [\{a^\top \phi\}^2] \leq 1 + 0.5 n_o \delta_{n_o}'^2, \|a\|_2^2 \leq \bar{c}, \|a^\top \phi\|_\infty \leq 1/\lambda\}} \{a^\top \phi(s, a)\}^2 \\
&\leq \sup_{\{a \in \mathbb{R}^d: a^\top \{n_o \Sigma_\rho + \lambda \mathbf{I}\} a \leq 2 + n_o \delta_{n_o}'^2, \|a\|_2^2 \leq \bar{c}, \|a^\top \phi\|_\infty < 1/\lambda\}} \{a^\top \phi(s, a)\}^2 \\
&\leq \sup_{\{a \in \mathbb{R}^d: a^\top \{n_o \Sigma_\rho + \lambda \mathbf{I}\} a \leq 2 + n_o \delta_{n_o}'^2\}} \{a^\top \phi(s, a)\}^2 \\
&= (2 + n_o \delta_{n_o}'^2) \phi^\top(s, a) \{n_o \Sigma_\rho + \lambda \mathbf{I}\}^{-1} \phi(s, a) \\
&\leq c_1 \{\text{rank}[\Sigma_\rho] + \log(c_2/\delta)\} \phi^\top(s, a) \{n_o \Sigma_\rho + \lambda \mathbf{I}\}^{-1} \phi(s, a).
\end{aligned}$$

Last step Then, the final bound is

$$\begin{aligned}
\mathbb{E}_{(s,a) \sim d_P^{\pi_e}} [\|\phi(s, a)\|_{\Sigma_{n_o}^{-1}}] &= \sqrt{C^{\pi_e} \mathbb{E}_{(s,a) \sim \rho} [\phi^\top(s, a) \Sigma_{n_o}^{-1} \phi(s, a)]} \\
&\leq c_1 \sqrt{C^{\pi_e} n_o \{\text{rank}[\Sigma_\rho] + \log(c_2/\delta)\} \mathbb{E}_{(s,a) \sim \rho} [\phi^\top(s, a) \{\Sigma_\rho + \lambda \mathbf{I}\}^{-1} \phi(s, a)]}.
\end{aligned}$$

Let UVU^\top be the eigenvalue decomposition of Σ_ρ s.t. $V_{i,i} = \mu_i$. We have

$$\begin{aligned}
\mathbb{E}_{(s,a) \sim \rho} [\phi^\top(s, a) \{n_o \Sigma_\rho + \lambda \mathbf{I}\}^{-1} \phi(s, a)] &= \text{Tr}[\{n_o \Sigma_\rho + \lambda \mathbf{I}\}^{-1} \{\Sigma_\rho\}] = \text{Tr}[\{n_o V + \lambda \mathbf{I}\}^{-1} V] \\
&= \frac{1}{n_o} \sum_{i=1}^{n_o} \frac{\mu_i}{\mu_i + \lambda/n_o} \leq \frac{\text{rank}[\Sigma_\rho]}{n_o}. \tag{12}
\end{aligned}$$

By combining all things together, with probability $1 - \delta$,

$$\mathbb{E}_{(s,a) \sim d_P^{\pi_e}} [\|\phi(s, a)\|_{\Sigma_{n_o}^{-1}}] \leq c_1 \sqrt{\frac{C^{\pi_e} \text{rank}[\Sigma_\rho] \{\text{rank}[\Sigma_\rho] + \log(c_2/\delta)\}}{n_o}}.$$

□

C.4.3 Missing proofs of non-parametric model

Finally, we provide missing proofs for the non-parametric GP model.

Proof of Theorem 24. In the proof, we use two statements, (13) and (14), in the proof of Theorem 25. We recommend readers to read the proof of Theorem 25 first.

We denote the eigenvalues of \mathbf{K}_{n_o} by $\{\hat{\mu}_i\}_{i=1}^{n_o}$ s.t. $\hat{\mu}_1 \geq \hat{\mu}_2 \geq \dots$. From Assumption 22, we have

$$n_o = \text{tr}(\mathbf{K}_{n_o}) = \sum_{i=1}^{n_o} \hat{\mu}_i.$$

Thus implies $\hat{\mu}_1 \leq n_o$. Then,

$$\begin{aligned}
\log(\det(\mathbf{I} + \zeta^{-2} \mathbf{K}_{n_o})) &= \sum_{i=1}^{n_o} \log\left(1 + \frac{\hat{\mu}_i}{\zeta^2}\right) = \sum_{i=1}^{n_o} \log\left(1 + \frac{\hat{\mu}_i}{\zeta^2}\right) \frac{\hat{\mu}_i/\zeta^2 + 1}{\hat{\mu}_i/\zeta^2 + 1} \\
&= \sum_{i=1}^{n_o} \log\left(1 + \frac{\hat{\mu}_i}{\zeta^2}\right) \frac{\hat{\mu}_i/\zeta^2}{\hat{\mu}_i/\zeta^2 + 1} + \log\left(1 + \frac{\hat{\mu}_i}{\zeta^2}\right) \frac{1}{\hat{\mu}_i/\zeta^2 + 1} \\
&= \sum_{i=1}^{n_o} \log\left(1 + \frac{\hat{\mu}_i}{\zeta^2}\right) \frac{\hat{\mu}_i/\zeta^2}{\hat{\mu}_i/\zeta^2 + 1} + \log\left(1 + \frac{\hat{\mu}_i}{\zeta^2}\right) \frac{1}{\hat{\mu}_i/\zeta^2 + 1} \\
&\leq \log\left(1 + \frac{\hat{\mu}_1}{\zeta^2}\right) \sum_{i=1}^{n_o} \frac{\hat{\mu}_i/\zeta^2}{\hat{\mu}_i/\zeta^2 + 1} + \sum_{i=1}^{n_o} \frac{\hat{\mu}_i/\zeta^2}{\hat{\mu}_i/\zeta^2 + 1} \quad (\log(1+x) \leq x) \\
&\leq \{\log(1 + n_o/\zeta^2) + 1\} \sum_{i=1}^{n_o} \frac{\hat{\mu}_i/\zeta^2}{\hat{\mu}_i/\zeta^2 + 1} \quad (\hat{\mu}_1 \leq n_o) \\
&\leq \{\log(1 + n_o/\zeta^2) + 1\} \min_j \{j + \hat{B}(j+1)/\zeta^2\} \leq 2\{\log(1 + n_o/\zeta^2) + 1\} \hat{d},
\end{aligned}$$

where the last second inequality uses the fact that $\sum_{i=1}^{n_o} \frac{\hat{\mu}_i/\xi^2}{\hat{\mu}_i/\xi^2+1} \leq j + \sum_{i=j+1}^{n_o} \hat{\mu}_i/\xi^2$. Then, the first statement is proved.

Next, we prove the second statement. We use

$$\sum_{i=1}^{n_o} \frac{\hat{\mu}_i/\zeta^2}{\hat{\mu}_i/\zeta^2+1} = \frac{1}{\zeta^2} \sum_{i=1}^{n_o} k_{n_o}(x_i, x_i).$$

proved in [Lemma 40](#). Then, from [\(13\)](#), with probability $1 - \delta$,

$$\frac{1}{\zeta^2} \sum_{i=1}^{n_o} k_{n_o}(x_i, x_i) \lesssim \delta'_{n_o} \sum_{i=1}^{n_o} \sup_{\{f:\zeta^2/n_o\|f\|_k^2+\mathbb{E}_{x\sim\rho}[f^2(x)]\leq 1, f\in\mathcal{H}_k\}} f^2(x_i),$$

where $\delta'_n = \delta_n + \sqrt{\log(c_2/\delta)/n_o}$ and δ_n is the critical radius of $\{f \in \mathcal{H}_k : \|f\|_k \leq 1\}$. Hereafter, we condition on the above event.

Then, from Bernstein's inequality,

$$\begin{aligned} & \left| \left\{ \sum_{i=1}^{n_o} \sup_{\{f:\zeta^2/n_o\|f\|_k^2+\mathbb{E}_{x\sim\rho}[f^2(x)]\leq 1, f\in\mathcal{H}_k\}} f^2(x_i) \right\} - n_o \mathbb{E}_{x\sim\rho} \left[\sup_{\{f:\zeta^2/n_o\|f\|_k^2+\mathbb{E}_{x\sim\rho}[f^2(x)]\leq 1, f\in\mathcal{H}_k\}} f^2(x) \right] \right| \\ & \lesssim \sqrt{n_o \text{Var}_{x\sim\rho} \left[\sup_{\{f:\zeta^2/n_o\|f\|_k^2+\mathbb{E}_{x\sim\rho}[f^2(x)]\leq 1, f\in\mathcal{H}_k\}} f^2(x) \right]} + n_o. \end{aligned}$$

We use for f in \mathcal{H}_k s.t. $\|f\|_k \leq 1$

$$|f(x)| = |\langle f(\cdot), k(x, \cdot) \rangle_k| \leq \|f\|_k \|k(x, \cdot)\|_k \leq 1.$$

from assumption [22](#). Here, from [\(14\)](#), the expectation is upper-bounded by

$$\mathbb{E}_{x\sim\rho} \left[\sup_{\{f:\zeta^2/n_o\|f\|_k^2+\mathbb{E}_{x\sim\rho}[f^2(x)]\leq 1, f\in\mathcal{H}_k\}} f^2(x) \right] \leq d^*.$$

Besides, the variance is also upper-bounded by

$$\begin{aligned} & \text{Var}_{x\sim\rho} \left[\sup_{\{f:\zeta^2/n_o\|f\|_k^2+\mathbb{E}_{x\sim\rho}[f^2(x)]\leq 1, f\in\mathcal{H}_k\}} f^2(x) \right] \\ & \leq \mathbb{E}_{x\sim\rho} \left[\sup_{\{f:\zeta^2/n_o\|f\|_k^2+\mathbb{E}_{x\sim\rho}[f^2(x)]\leq 1, f\in\mathcal{H}_k\}} f^4(x) \right] \\ & \leq \mathbb{E}_{x\sim\rho} \left[\sup_{\{f:\zeta^2/n_o\|f\|_k^2+\mathbb{E}_{x\sim\rho}[f^2(x)]\leq 1, f\in\mathcal{H}_k\}} f^2(x) \right] \\ & \quad (f^2(x) \leq 1 \forall x \in \mathcal{S} \times \mathcal{A} \text{ from Assumption 22}) \\ & = d^*. \end{aligned} \tag{From (14)}$$

Thus, with probability $1 - \delta$,

$$\begin{aligned} \sum_{i=1}^{n_o} k_{n_o}(x_i, x_i) & \lesssim \{\delta'_{n_o}\}^2 n_o (d^* + \sqrt{d^*} + 1) \\ & \lesssim c_1 \{d^* + \log(c_2/\delta)\} d^*. \end{aligned}$$

noting $\delta'_{n_o} = \sqrt{d^*/n_o} + \sqrt{\log(c_2/\delta)/n_o}$ from [Theorem 25](#).

By combining all things together, with probability $1 - \delta$,

$$\begin{aligned} \log(\det(\mathbf{I} + \zeta^{-2} \mathbf{K}_{n_o})) & \leq \{\log(1 + n_o/\zeta^2) + 1\} \sum_{i=1}^{n_o} \frac{\hat{\mu}_i/\zeta^2}{\hat{\mu}_i/\zeta^2+1} \\ & = \{\log(1 + n_o/\zeta^2) + 1\} \frac{1}{\zeta^2} \sum_{i=1}^{n_o} k_{n_o}(x_i, x_i) \\ & \lesssim \{\log(1 + c_3 n_o)\} \{d^* + \log(c_2/\delta)\} d^*. \end{aligned}$$

This concludes the proof. □

Proof of [Theorem 25](#).

First Statement From Jensen's inequality, we have

$$\mathbb{E}_{x \sim d_P^{\pi_e}} [\sqrt{k_{n_o}(x, x)}] \leq \sqrt{\mathbb{E}_{x \sim d_P^{\pi_e}} [k_{n_o}(x, x)]}.$$

Thus, we focus how to bound $\mathbb{E}_{x \sim d_P^{\pi_e}} [k_{n_o}(x, x)]$. Before that, we show the following statement. With probability $1 - \delta$, we have for $\forall x \in \mathcal{S} \times \mathcal{A}$:

$$k_{n_o}(x, x) \leq c_1 \zeta^2 \delta'_{n_o} \times \sup_{\{f \in \mathcal{H}_k: \zeta^2/n_o \|f\|_k^2 + \mathbb{E}_{x \sim \rho}[f^2(x)] \leq 1\}} f^2(x), \quad (13)$$

where $\delta'_{n_o} = \delta_{n_o} + \sqrt{\log(c_2/\delta)/n_o}$ and δ_{n_o} is the critical radius of $\{f \in \mathcal{H}_k : \|f\|_k \leq 1\}$.

As the first step, we use [Lemma 38](#) and [Lemma 39](#).

$$k_{n_o}(x, x) = \sup_{\{f \in \mathcal{H}_{k_{n_o}} \|f\|_{k_{n_o}}^2 \leq 1\}} f^2(x) \quad (\text{From Lemma 38})$$

$$= \sup_{\{f \in \mathcal{H}_k \|f\|_k^2 + \zeta^{-2} \sum_{i=1}^{n_o} f(x_i)^2 \leq 1\}} f^2(x). \quad (\text{From Lemma 39})$$

Next invoke [Lemma 36](#), that is, with probability $1 - \delta$,

$$\frac{1}{n_o} \sum_{i=1}^{n_o} f^2(x_i) \geq 0.5 \mathbb{E}_{(s,a) \sim \rho} [f^2(x)] - 0.5 \{\delta'_{n_o}\}^2 \forall f \in \mathcal{F}$$

where

$$\mathcal{F} = \{f : f \in \mathcal{H}_k, \|f\|_k^2 = 1\}.$$

Here, $\delta'_{n_o} = \delta_{n_o} + \sqrt{\log(c_2/\delta)/n_o}$, where δ_{n_o} is the critical radius of the function class \mathcal{F} . Hereafter, we condition on the above event. Note the uniform boundedness assumption of \mathcal{F} for [Lemma 36](#) is satisfied noting

$$|f(x)| = |\langle f(\cdot), k(\cdot, x) \rangle_k| \leq \|f\|_k \|k(\cdot, x)\|_k \leq 1.$$

noting assumption [22](#). Then, we have

$$k_{n_o}(x, x) \leq \sup_{\{f \in \mathcal{H}_k \|f\|_k^2 + \zeta^{-2} n_o / 2 \mathbb{E}_{x \sim \rho} [f^2(x)] \leq 1 + n_o \delta'_{n_o} / 2\}} f^2(x).$$

$k_{n_o}(x, x)$ is further upper-bounded by

$$\begin{aligned} k_{n_o}(x, x) &\leq \sup_{\{f \in \mathcal{H}_k: 2\zeta^2/n_o \|f\|_k^2 + \mathbb{E}_{x \sim \rho} [f^2(x)] \leq 2\zeta^2/n_o + \zeta^2 \delta'_{n_o}\}} f^2(x) && (\text{Multiply } 2\zeta^2/n_o) \\ &\leq (2\zeta^2/n_o + \zeta^2 \delta'_{n_o}) \times \sup_{\{f \in \mathcal{H}_k: \zeta^2/n_o \|f\|_k^2 + \mathbb{E}_{x \sim \rho} [f^2(x)] \leq 1\}} f^2(x) \\ &\leq c_1 \zeta^2 \delta'_{n_o} \times \sup_{\{f \in \mathcal{H}_k: \zeta^2/n_o \|f\|_k^2 + \mathbb{E}_{x \sim \rho} [f^2(x)] \leq 1\}} f^2(x). \end{aligned}$$

This concludes [\(13\)](#).

Next, we show

$$\mathbb{E}_{x \sim d_P^{\pi_e}} \left[\sup_{\{f \in \mathcal{H}_k: \zeta^2/n_o \|f\|_k^2 + \mathbb{E}_{x \sim \rho} [f^2(x)] \leq 1\}} f^2(x) \right] \leq 2d^* \times \sup_{\|x\|_2 \leq 1} \frac{x^\top \Sigma_{\pi_e} x}{x^\top \Sigma_{\rho} x}.$$

For $f(\cdot) = a^\top \phi(\cdot)$ (recall $\phi(\cdot)$ is the feature mapping defined by the eigenvalues μ_i and eigenfunctions ϕ , s.t. $\phi = (\phi_1, \dots, \phi_\infty)$), we have

$$\|f\|_k^2 = a^\top a, \quad \mathbb{E}_{x \sim \rho} [f^2(x)] = a^\top M a.$$

where M is a diagonal matrix in $\mathbb{R}^{\infty \times \infty}$ s.t. $M_{i,i} = \mu_i$. Thus,

$$\mathbb{E}_{x \sim d_P^{\pi_e}} \left[\sup_{\{f: \zeta^2/n_o \|f\|_k^2 + \mathbb{E}_{x \sim \rho} [f^2(x)] \leq 1, f \in \mathcal{H}_k\}} f^2(x) \right] = \mathbb{E}_{x \sim d_P^{\pi_e}} \left[\sup_{\{a \in \mathbb{R}^\infty: a^\top (\zeta^2/n_o \mathbf{I} + M) a \leq 1\}} \{a^\top \phi(x)\}^2 \right].$$

Then, by letting Σ_ρ and Σ_{π_e} be $\mathbb{E}_{(s,a)\sim\rho}[\phi(s,a)\phi^\top(s,a)]$ and $\mathbb{E}_{(s,a)\sim d_P^{\pi_e}}[\phi(s,a)\phi^\top(s,a)]$,

$$\begin{aligned}\mathbb{E}_{x\sim d_P^{\pi_e}}\left[\sup_{\{a\in\mathbb{R}^d:a^\top(\zeta^2/n_o\mathbf{I}+M)\leq 1\}}\{a^\top\phi(x)\}^2\right] &= \mathbb{E}_{x\sim d_P^{\pi_e}}[\phi(x)\{\zeta^2/n_o\mathbf{I}+M\}^{-1}\phi(x)] \\ &= \text{tr}[\mathbb{E}_{x\sim d_P^{\pi_e}}[\phi(x)\phi(x)^\top]\{\zeta^2/n_o\mathbf{I}+M\}^{-1}] \\ &= \text{tr}[\mathbb{E}_{x\sim\rho}[\phi(x)\phi(x)^\top]\{\zeta^2/n_o\mathbf{I}+M\}^{-1}]\times\sup_{\|x\|_2\leq 1}\frac{x^\top\Sigma_{\pi_e}x}{x^\top\Sigma_\rho x} \\ &= \sum_{i=1}^{\infty}\frac{\mu_i}{\zeta^2/n_o+\mu_i}\times\sup_{\|x\|_2\leq 1}\frac{x^\top\Sigma_{\pi_e}x}{x^\top\Sigma_\rho x}.\end{aligned}$$

Then, by defining $C^{\pi_e} = \sup_{\|x\|_2\leq 1}\frac{x^\top\Sigma_{\pi_e}x}{x^\top\Sigma_\rho x}$, we have

$$\begin{aligned}\mathbb{E}_{x\sim d_P^{\pi_e}}\left[\sup_{\{a\in\mathbb{R}^d:a^\top(\zeta^2/n_o+M)\leq 1\}}\{a^\top\phi(x)\}^2\right] &\leq \min_j\{j+n_o/\zeta^2\sum_{i=j+1}^{\infty}\mu_i\}\times C^{\pi_e} \\ &\leq \min_j\{j+n_o/\zeta^2\sum_{i=j+1}^{\infty}\mu_i\}\times C^{\pi_e} \\ &\leq 2d^*\times C^{\pi_e}.\end{aligned}$$

By combining all things together ((13) and (14)), the statement is concluded, that is, with probability $1-\delta$:

$$\begin{aligned}\mathbb{E}_{d_P^{\pi_e}}[\sqrt{k_{n_o}(x,x)}] &\leq \sqrt{\zeta^2\delta_{n_o}'^2\times\mathbb{E}_{x\sim d_P^{\pi_e}}\left[\sup_{\{f\in\mathcal{H}_k:\zeta^2/n_o\|f\|_k^2+\mathbb{E}_{x\sim\rho}[f^2(x)]\leq 1\}}f^2(x)\right]} \\ &\leq \zeta\delta_{n_o}'\sqrt{C^{\pi_e}d^*}.\end{aligned}$$

where $C^{\pi_e} = \sup_{\|x\|_2\leq 1}\frac{x^\top\Sigma_{\pi_e}x}{x^\top\Sigma_\rho x}$.

Remark 31. Like the above, We can also prove

$$\mathbb{E}_{x\sim\rho}\left[\sup_{\{f\in\mathcal{H}_k:\zeta^2/n_o\|f\|_k^2+\mathbb{E}_{x\sim\rho}[f^2(x)]\leq 1\}}f^2(x)\right]\leq\sum_{i=1}^{\infty}\frac{\mu_i}{\zeta^2/n_o+\mu_i}\leq 2d^*. \quad (14)$$

This is used in the proof of theorem 24.

Remark 32. We can also use

$$\begin{aligned}\mathbb{E}_{x\sim d_P^{\pi_e}}\left[\sup_{\{a\in\mathbb{R}^d:a^\top(\zeta^2/n_o\mathbf{I}+M)\leq 1\}}\{a^\top\phi(x)\}^2\right] &= \mathbb{E}_{x\sim d_P^{\pi_e}}[\phi(x)\{\zeta^2/n_o\mathbf{I}+M\}^{-1}\phi(x)] \\ &= \text{tr}[\mathbb{E}_{x\sim d_P^{\pi_e}}[\phi(x)\phi(x)^\top]\{\zeta^2/n_o\mathbf{I}+M\}^{-1}] \\ &= \sum_{i=1}^{\infty}\frac{\mathbb{E}_{x\sim d_P^{\pi_e}}[\phi_i(x)\phi_i(x)^\top]}{\zeta^2/n_o+\mu_i} \\ &= \sum_{i=1}^{\infty}\frac{\mu_i}{\zeta^2/n_o+\mu_i}\times\left\{\frac{\mathbb{E}_{x\sim d_P^{\pi_e}}[\phi_i(x)\phi_i(x)^\top]}{\mu_i}\right\} \\ &= \sum_{j=1}^{\infty}\frac{\mu_j}{\zeta^2/n_o+\mu_j}\times\max_i(\mathbb{E}_{x\sim d_P^{\pi_e}}[\psi_i(x)\psi_i(x)^\top]).\end{aligned}$$

Then, C^{π_e} is replaced with $\max_i(\mathbb{E}_{x\sim d_P^{\pi_e}}[\psi_i(x)\psi_i(x)^\top])$.

Second statement We use [Lemma 41](#) to calculate the critical radius of the RKHS ball. The critical inequality is

$$\sqrt{1/n_o} \sqrt{\sum_{i=1}^{n_o} \min(y^2, \mu_j)} \leq y^2.$$

We show $y = \sqrt{d^*/n_o}$ satisfies the above. This is proved by

$$\begin{aligned} \sqrt{1/n_o} \sqrt{\sum_{i=1}^{n_o} \min(y^2, \mu_j)} &\leq \min_{1 \leq k \leq n_o} \{\sqrt{1/n} \sqrt{ky^2 + B(k+1)}\} \\ &\leq \sqrt{1/n_o} \sqrt{d^*y^2 + B(d^*+1)} && (d^* \leq n_o) \\ &\leq \sqrt{1/n_o} \sqrt{d^*y^2 + d^*/n_o} && (B(d^*+1) \leq d^*/n_o) \\ &\leq \sqrt{d^*y^2/n_o} \leq y^2. \end{aligned}$$

□

D Auxiliary Lemmas

Lemma 33 (Simulation Lemma). *Consider any two functions $f : \mathcal{S} \times \mathcal{A} \mapsto [0, 1]$ and $\hat{f} : \mathcal{S} \times \mathcal{A} \mapsto [0, 1]$, any two transitions P and \hat{P} , and any policy $\pi : \mathcal{S} \mapsto \Delta(\mathcal{A})$. We have:*

$$\begin{aligned} V_{P;f}^\pi - V_{\hat{P};\hat{f}}^\pi &= \sum_{h=0}^H \mathbb{E}_{s,a \sim d_P^\pi} \left[f(s,a) - \hat{f}(s,a) + \mathbb{E}_{s' \sim P(\cdot|s,a)} [V_{\hat{P};\hat{f};h}^\pi(s')] - \mathbb{E}_{s' \sim \hat{P}(\cdot|s,a)} [V_{\hat{P};\hat{f};h}^\pi(s')] \right] \\ &\leq \sum_{h=0}^H \mathbb{E}_{s,a \sim d_P^\pi} \left[f(s,a) - \hat{f}(s,a) + \|V_{\hat{P};\hat{f};h}^\pi\|_\infty \|P(\cdot|s,a) - \hat{P}(\cdot|s,a)\|_1 \right]. \end{aligned}$$

where $V_{P;f;h}^\pi$ denotes the value function at time step h , under π, P, f .

Such simulation lemma is standard in model-based RL literature and the derivation can be found, for instance, in the proof of [Lemma 10](#) from [\[64\]](#).

Lemma 34 (ℓ_1 Distance between two Gaussians). *Consider two Gaussian distributions $P_1 := \mathcal{N}(\mu_1, \zeta^2 I)$ and $P_2 := \mathcal{N}(\mu_2, \zeta^2 I)$. We have:*

$$\|P_1 - P_2\|_1 \leq \frac{1}{\zeta} \|\mu_1 - \mu_2\|_2.$$

This lemma is proved by Pinsker's inequality and the closed-form of the KL divergence between P_1 and P_2 .

Lemma 35 (Concentration on the inverse of state-action visitation). *We set $\lambda = \Omega(1)$. Then, with probability $1 - \delta$,*

$$\frac{1}{N(s,a) + \lambda} \leq \frac{c_1 \log(|\mathcal{S}||\mathcal{A}|c_2/\delta)}{n_o \rho(s,a) + \lambda} \quad \forall (s,a) \in \mathcal{S} \times \mathcal{A}.$$

The extension of this lemma to the linear models is stated in [eq. \(11\)](#).

Proof. We set $\xi = c_1 \log(|\mathcal{S}||\mathcal{A}|/\delta) + 1$ ($c_1 > 4/3 + 3$). First, we have

$$\frac{1}{N(s,a) + \lambda} \leq \frac{\xi}{N(s,a) + \xi\lambda}.$$

from $\xi \geq 1$. Here, by Bernsteins's inequality, with probability $1 - \delta$,

$$N(s,a) \geq n_o \rho(s,a) - 2\sqrt{2n_o \rho(s,a)(1 - \rho(s,a)) \log(|\mathcal{S}||\mathcal{A}|/\delta)} - 4 \log(|\mathcal{S}||\mathcal{A}|/\delta)/3, \quad \forall (s,a).$$

Thus, $\forall (s, a) \in \bar{V}$, we have

$$\begin{aligned}
N(s, a) + \xi\lambda &\geq n_o\rho(s, a) - 2\sqrt{2n_o\rho(s, a)(1 - \rho(s, a)) \log(|\mathcal{S}||\mathcal{A}|/\delta)} - 4\log(|\mathcal{S}||\mathcal{A}|/\delta)/3 + \xi\lambda \\
&\geq n_o\rho(s, a) - 2\sqrt{2n_o\rho(s, a)(1 - \rho(s, a)) \log(|\mathcal{S}||\mathcal{A}|/\delta)} + (c_1 - 4/3)\log(|\mathcal{S}||\mathcal{A}|/\delta) + \lambda \\
&\geq n_o\rho(s, a) - 2\sqrt{2n_o\rho(s, a) \log(|\mathcal{S}||\mathcal{A}|/\delta)} + (c_1 - 4/3)\log(|\mathcal{S}||\mathcal{A}|/\delta) + \lambda \\
&\geq 0.5n_o\rho(s, a) + (\sqrt{0.5n_o\rho(s, a)} - \sqrt{4\log(|\mathcal{S}||\mathcal{A}|/\delta)})^2 + (c_1 - 4/3 - 4)\log(|\mathcal{S}||\mathcal{A}|/\delta) + \lambda \\
&\geq 0.5n_o\rho(s, a) + 0.5\lambda.
\end{aligned}$$

This implies with $1 - \delta$,

$$\frac{1}{N(s, a) + \lambda} \leq \frac{2\xi}{n_o\rho(s, a) + \lambda} \forall (s, a).$$

Then, noting $c_1 \log(|\mathcal{S}||\mathcal{A}|/\delta) + 1 \leq c_1 \log(|\mathcal{S}||\mathcal{A}|c_2/\delta)$ for some c_2 , the proof is concluded. \square

Lemma 36 (A uniform law with localization: Theorem 14.1 in [73]). *Assume $\|\mathcal{F}\|_\infty \leq b$. Denote the critical radius of a function class \mathcal{F} by δ_n . The critical radius δ_n is defined as a solution to*

$$\mathcal{R}_{n_o}(y; \mathcal{F}) \leq y^2/b.$$

w.r.t y . Then, with probability $1 - \delta$

$$\frac{1}{n_o} \sum_{i=1}^{n_o} f(x_i)^2 \geq 1/2 \mathbb{E}_{x \sim \rho}[f^2(x)] - (\delta'_n)^2/2 \quad \forall f \in \mathcal{F},$$

where $\delta'_n = \delta_n + c_1 \sqrt{\log(c_2/\delta)/n_o}$.

Lemma 37 (Critical radius of linear models). *Assume $\|\phi(s, a)\|_2 \leq 1$ for any $(s, a) \in \mathcal{S} \times \mathcal{A}$. Then, the critical radius of function class $\mathcal{F} = \{(s, a) \mapsto a^\top \phi(s, a) : \|a\|_2^2 \leq \alpha, a^\top \phi \leq \beta, a \in \mathbb{R}^d\}$ is upper-bounded by*

$$c\sqrt{\beta \text{rank}(\Sigma_\rho)/n_o}.$$

where c is a universal constant.

We follow the proof of [73, Chapter 14]. Their argument depends on the assumption Σ_ρ is full rank. We need to change the proof so that the full-rank assumption is removed and the rank $\text{rank}[\Sigma_\rho]$ would appear in the final bound instead of d . Note that the final bound does not include α .

Proof. Unless otherwise noted, in this proof, $\mathbb{E}[\cdot]$ is taken w.r.t.

$$x_i = (s_i, a_i) \sim \rho(s, a), \quad \epsilon_i \sim 2\{\text{Ber}(0.5) - 1\}.$$

Note that x_i and ϵ_i are independent.

Noting $\mathbb{E}_{\rho \sim (s, a)}[(a^\top \phi(s, a))^2] = a^\top \Sigma_\rho a$, the localized Rademacher complexity of \mathcal{F} , $\mathcal{R}_{n_o}(\xi; \mathcal{F})$, is

$$\mathbb{E} \left[\sup_{\{b \in \mathbb{R}^d: \|b\|_2^2 \leq \alpha, \|b\|_{\Sigma_\rho} \leq \xi, b^\top \phi \leq \beta\}} \left| \frac{1}{n_o} \sum_{i=1}^{n_o} \epsilon_i \{b^\top \phi(s_i, a_i)\} \right| \right]$$

where $\{\epsilon_i\}_{i=1}^{n_o}$ is a set of independent Rademacher variables. This is upper-bounded by

$$\mathbb{E} \left[\sup_{\{b \in \mathbb{R}^d: \|b\|_2^2 \leq \alpha, \|b\|_{\Sigma_\rho} \leq \xi\}} \left| \frac{1}{n_o} \epsilon^\top \Phi b \right| \right]$$

where Φ is a $n_o \times d$ design matrix s.t. the i -th row is $\phi^\top(s_i, a_i)$ and $\epsilon = (\epsilon_1, \dots, \epsilon_{n_o})^\top$.

Here, we have $\mathbb{E}[\Phi^\top \Phi] = n_o \Sigma_\rho$. Let UVU^\top be the SVD of Σ_ρ , where U is a $n \times \text{rank}[\Sigma_\rho]$ matrix and V is a $\text{rank}[\Sigma_\rho] \times \text{rank}[\Sigma_\rho]$ diagonal matrix. Noting $b = UU^\top b + (\mathbf{I} - UU^\top)b$, we have

$$\begin{aligned}
& \mathbb{E} \left[\sup_{\{b \in \mathbb{R}^d: \|b\|_2^2 \leq \alpha, \|b\|_{\Sigma_\rho} \leq \xi\}} \left| \frac{1}{n_o} \epsilon^\top \Phi \{UU^\top b + (\mathbf{I} - UU^\top)b\} \right| \right] \\
& \leq \mathbb{E} \left[\sup_{\{b \in \mathbb{R}^d: \|b\|_2^2 \leq \alpha\}} \left| \frac{1}{n_o} \epsilon^\top \Phi (\mathbf{I} - UU^\top)b \right| \right] + \mathbb{E} \left[\sup_{\|b\|_{\Sigma_\rho} \leq \xi} \left| \frac{1}{n_o} \epsilon^\top \Phi UU^\top b \right| \right] \\
& \leq \mathbb{E} \left[\sup_{\{b \in \mathbb{R}^d: \|b\|_2^2 \leq \alpha\}} \left| \frac{1}{n_o} \epsilon^\top \Phi (\mathbf{I} - UU^\top)b \right| \right] + \mathbb{E} \left[\sup_{\|c\|_V \leq \xi} \left| \frac{1}{n_o} \epsilon^\top \Phi U c \right| \right] \quad (U^\top c = b) \\
& \leq \mathbb{E} \left[\frac{\alpha}{n_o} \|\epsilon^\top \Phi (\mathbf{I} - UU^\top)\|_2 \right] + \frac{\zeta}{n_o} \mathbb{E} [\|\epsilon^\top \Phi U\|_{V^{-1}}] \quad (\text{CS inequality}) \\
& \leq \frac{\alpha}{n_o} \sqrt{\mathbb{E} [\|\epsilon^\top \Phi (\mathbf{I} - UU^\top)\|_2^2]} + \frac{\zeta}{n_o} \sqrt{\mathbb{E} [\|\epsilon^\top \Phi U\|_{V^{-1}}^2]}. \quad (\text{Jensen's inequality})
\end{aligned}$$

We analyze the second term and first term respectively.

Regarding the second term, we have

$$\mathbb{E}_\epsilon [\|\epsilon^\top \Phi U\|_{V^{-1}}^2] = \mathbb{E}_\epsilon [\epsilon^\top \Phi UV^{-1}U^\top \Phi^\top \epsilon] = \text{tr}(\Phi UV^{-1}U^\top \Phi^\top),$$

where $\mathbb{E}_\epsilon[\cdot]$ is an expectation only regarding ϵ . Then, by the law of total expectation,

$$\begin{aligned}
\mathbb{E} [\|\epsilon^\top \Phi U\|_{V^{-1}}^2] &= \mathbb{E} [\text{tr}(\Phi UV^{-1}U^\top \Phi^\top)] \\
&= \mathbb{E} [\text{tr}(\Phi^\top \Phi UV^{-1}U^\top)] = \text{tr}(n_o \Sigma_\rho UV^{-1}U^\top) \\
&= n_o \text{tr}(UVU^\top UV^{-1}U^\top) \\
&= n_o \text{tr}(UU^\top) = n_o \text{tr}(U^\top U) = n_o \text{rank}(\Sigma_\rho).
\end{aligned}$$

Similarly,

$$\mathbb{E}_\epsilon [\|\epsilon^\top \Phi (\mathbf{I} - UU^\top)\|_2^2] = \text{tr}(\Phi (\mathbf{I} - UU^\top)(\mathbf{I} - UU^\top) \Phi^\top) = \text{tr}(\Phi^\top \Phi (\mathbf{I} - UU^\top)).$$

Then, by the law of total expectation,

$$\begin{aligned}
\mathbb{E} [\|\epsilon^\top \Phi (\mathbf{I} - UU^\top)\|_2^2] &= \mathbb{E} [\text{tr}(\Phi^\top \Phi (\mathbf{I} - UU^\top))] \\
&= n_o \text{tr}(\Sigma_\rho (\mathbf{I} - UU^\top)) = n_o \text{tr}(UVU^\top (\mathbf{I} - UU^\top)) = 0.
\end{aligned}$$

Combining all things together,

$$\mathcal{R}_n(\xi; \mathcal{F}) \leq \xi \sqrt{\text{rank}[\Sigma_\rho]/n_o}.$$

Then, the critical inequality becomes

$$y \sqrt{\text{rank}(\Sigma_\rho)/n_o} \leq y^2/\beta.$$

Thus, the critical radius of \mathcal{F} is

$$\sqrt{\beta \text{rank}(\Sigma_\rho)/n_o}.$$

□

Lemma 38 (Variatioanal representation of kernels). *We denote the RKHS associated with a kernel $k(\cdot, \cdot)$ by \mathcal{H}_k . Then,*

$$k(x, x) = \sup_{\{f: \|f\|_k \leq 1, f \in \mathcal{H}_k\}} f^2(x).$$

Proof. We have

$$\begin{aligned} \sup_{\{f: \|f\|_k \leq 1, f \in \mathcal{H}_k\}} f^2(x) &= \sup_{\{f: \|f\|_k \leq 1, f \in \mathcal{H}_k\}} \langle f, k(x, \cdot) \rangle_k^2 \\ &\leq \sup_{\{f: \|f\|_k \leq 1, f \in \mathcal{H}_k\}} \|f\|_k^2 k(x, x) \\ &= k(x, x). \end{aligned} \quad (\text{CS inequality})$$

Besides, the equality is satisfied when $f(\cdot) = k(x, \cdot) / \sqrt{k(x, x)}$ noting

$$f^2(x) = k^2(x, x) / k(x, x) = k(x, x), \quad \|f(\cdot)\|_k = \|k(x, \cdot)\|_k / k(x, x) = 1.$$

Thus,

$$k(x, x) = \sup_{\{f: \|f\|_k \leq 1, f \in \mathcal{H}_k\}} f^2(x).$$

□

Lemma 39 (Relation between $\mathcal{H}_{k_{n_o}}$ and \mathcal{H}_k). *We denote the RKHS associated with a kernel $k(\cdot, \cdot)$ by \mathcal{H}_k and the RKHS with a kernel $k_{n_o}(\cdot, \cdot)$ by $\mathcal{H}_{k_{n_o}}$. Then, we have $\mathcal{H}_k = \mathcal{H}_{k_{n_o}}$. Besides, for $f \in \mathcal{H}_k$, we have*

$$\|f\|_{k_{n_o}}^2 = \|f\|_k^2 + \zeta^{-2} \sum_{i=1}^{n_o} f(x_i)^2.$$

This is stated in [63, Appendix B] without the proof. For completeness, we provide the proof.

Proof. We use Mercer's theorem [73, Theorem 12.20]. Then, any element in the RKHS associated with the kernel $k(x, x)$ is represented by

$$f(x) = \sum_{i=1}^{\infty} f_i \psi_i(x).$$

where $\{\psi_i\}_{i=1}^{\infty}$ is an orthonormal basis for $L^2(\rho)$: $\mathbb{E}_{x \sim \rho}[\psi_i(x)\psi_j(x)] = I(i = j)$. Here, we have

$$k(x, x) = \psi^\top(x) \Lambda \psi(x) = \phi^\top(x) \phi(x), \quad \|f\|_k = \tilde{f}^\top \Lambda^{-1} \tilde{f},$$

where $\phi_i(x) = \sqrt{\mu_i} \psi_i(x)$ and $\tilde{f} = \{f_i\}_{i=1}^{\infty} \in \mathbb{R}^\infty$.

Then, by letting Φ be a $n \times d$ matrix s.t. the i -th row is $\phi^\top(s_i, a_i)$,

$$\begin{aligned} k_{n_o}(x, x) &= \phi^\top(x) \phi(x) - \phi^\top(x) \Phi^\top (\Phi \Phi^\top + \zeta^2 \mathbf{I})^{-1} \Phi \phi(x) \\ &= \phi^\top(x) \{ \mathbf{I} - \Phi^\top (\Phi \Phi^\top + \zeta^2 \mathbf{I})^{-1} \Phi \} \phi(x) \\ &= \phi^\top(x) \{ \mathbf{I} + \Phi^\top \Phi / \zeta^2 \}^{-1} \phi(x) \quad (\text{Woodbury matrix identity}) \\ &= \phi^\top(x) \left(\mathbf{I} + \sum_{i=1}^{n_o} \phi(x_i) \phi(x_i)^\top / \zeta^2 \right)^{-1} \phi(x). \end{aligned}$$

Here, let UVU^\top be the eigenvalue decomposition of $\{ \Lambda^{-1} + \sum_{i=1}^{n_o} \psi(x_i) \psi(x_i)^\top / \zeta^2 \}^{-1} = UVU^\top$. Then,

$$\begin{aligned} k_{n_o}(x, x) &= \psi^\top(x) \left(\Lambda^{-1} + \sum_{i=1}^{n_o} \psi(x_i) \psi(x_i)^\top / \zeta^2 \right)^{-1} \psi(x) \\ &= \psi^\top(x) UVU^\top \psi(x) \\ &= \psi'^\top(x) V \psi'(x). \quad (U^\top \psi = \psi') \end{aligned}$$

Then, any element $f(\cdot)$ in the RKHS associated with the kernel $k_{n_o}(x, x)$ is represented as

$$f(\cdot) = \tilde{g}^\top \psi'(\cdot), \quad \tilde{g} \in \mathbb{R}^\infty,$$

and the associated norm is $\|f\|_{k_{n_o}} = \tilde{g}^\top V^{-1} \tilde{g}$ since $\psi'(\cdot)$ is still an orthonormal basis for $L^2(\rho)$, i.e., $\mathbb{E}_{x \sim \rho}[\phi'_i(x) \phi'_j(x)] = I(i = j)$. This immediately implies $\mathcal{H}_k = \mathcal{H}_{k_{n_o}}$.

Finally, we check the relation of the norm:

$$\begin{aligned}
\|f\|_{k_{n_o}}^2 &= \left\| \sum_{i=1}^{n_o} f_i \psi_i \right\|_{k_{n_o}} = \|\tilde{f}^\top \psi\|_{k_{n_o}} & (\tilde{f} = \{f_1, f_2 \cdots\}^\top) \\
&= \|\{U^\top \tilde{f}\}^\top U^\top \psi\|_{k_{n_o}} \\
&= \|\{U^\top \tilde{f}\}^\top \psi'\|_{k_{n_o}} \\
&= \{U^\top \tilde{f}\}^\top V^{-1} U \tilde{f} \\
&= \tilde{f}^\top (\Lambda^{-1} + \sum_{i=1}^{n_o} \phi(x_i) \phi(x_i)^\top / \zeta^2) \tilde{f} \\
&= \|f\|_k + 1/\zeta^2 \sum_{i=1}^{n_o} \{f^\top \phi(x_i)\}^2 = \|f\|_k + \zeta^{-2} \sum_{i=1}^{n_o} f^2(x_i).
\end{aligned}$$

□

Lemma 40. Let $\{\hat{\mu}_i\}_{i=1}^{n_o}$ be the eigenvalues of \mathbf{K}_{n_o} . Then,

$$\sum_{i=1}^{n_o} \frac{\hat{\mu}_i / \zeta^2}{\hat{\mu}_i / \zeta^2 + 1} = \frac{1}{\zeta^2} \sum_{i=1}^{n_o} k_{n_o}(x_i, x_i).$$

Proof.

$$\begin{aligned}
\sum_{i=1}^{n_o} k_{n_o}(x_i, x_i) &= \sum_{i=1}^{n_o} k(x_i, x_i) - \bar{k}_{n_o}^\top(x_i) \{\mathbf{K}_{n_o} + \zeta^2 \mathbf{I}\}^{-1} \bar{k}_{n_o}(x_i) \\
&= \text{tr} \left(\sum_{i=1}^{n_o} k(x_i, x_i) - \bar{k}_{n_o}^\top(x_i) \{\mathbf{K}_{n_o} + \zeta^2 \mathbf{I}\}^{-1} \bar{k}_{n_o}(x_i) \right) \\
&= \text{tr}(\mathbf{K}_{n_o}) - \text{tr} \left(\sum_{i=1}^{n_o} \bar{k}_{n_o}(x_i) \bar{k}_{n_o}^\top(x_i) \{\mathbf{K}_{n_o} + \zeta^2 \mathbf{I}\}^{-1} \right) \\
&= \text{tr}(\mathbf{K}_{n_o} - \mathbf{K}_{n_o}^2 \{\mathbf{K}_{n_o} + \zeta^2 \mathbf{I}\}^{-1}) \\
&= \text{tr}(\{\mathbf{K}_{n_o}^2 + \zeta^2 \mathbf{K}_{n_o} - \mathbf{K}_{n_o}^2\} \{\mathbf{K}_{n_o} + \zeta^2 \mathbf{I}\}^{-1}) \\
&= \text{tr}(\zeta^2 \mathbf{K}_{n_o} \{\mathbf{K}_{n_o} + \zeta^2 \mathbf{I}\}^{-1}) = \sum_{i=1}^{n_o} \frac{\hat{\mu}_i}{\hat{\mu}_i / \zeta^2 + 1}.
\end{aligned}$$

□

Lemma 41 (Calculation of localized Rademacher complexity of RKHS balls: Corollary 14.5 in [73]). Let $\mathcal{F} = \{f \in \mathcal{H}_k : \|f\|_k \leq 1\}$ be the unit ball of an RKHS with eigenvalues $\{\mu_j\}_{j=1}^\infty$. Then, the localized population Rademacher complexity is upper-bounded by

$$\mathcal{R}_n(\delta; \mathcal{F}) \leq \sqrt{\frac{2}{n}} \sqrt{\sum_{j=1}^\infty \min(\mu_j, \delta^2)}.$$

Lemma 42 (Upper-bound of expectation of information gains: finite-dimensional models).

$$\mathbb{E}[\bar{\mathcal{I}}_{n_o}] \leq \text{rank}(\Sigma_\rho) \{\log(1 + n_o/\lambda) + 1\}.$$

Proof.

$$\begin{aligned}
\mathbb{E}[\bar{\mathcal{I}}_{n_o}] &= \mathbb{E}[\log(\det(\Sigma_{n_o}/\lambda))] \\
&\leq \log \det(\mathbb{E}[\Sigma_{n_o}/\lambda]) = \log \det(\mathbf{I} + n_o/\lambda \Sigma_\rho) && \text{(Jensen's inequality)} \\
&\leq \text{rank}(\Sigma_\rho) \{\log(1 + n_o/\lambda) + 1\}.
\end{aligned}$$

The final line is proved as in the proof of theorem 20. □

Lemma 43 (Upper-bound of expectation of information gains: RKHS).

$$\mathbb{E}[\mathcal{I}_{n_o}] \leq 2d^* \{\log(1 + n_o/\zeta^2) + 1\}.$$

Proof.

$$\begin{aligned} \mathbb{E}[\mathcal{I}_{n_o}] &= \mathbb{E}[\log(\det(I + \zeta^{-2}\mathbf{K}_{n_o}))] \\ &\leq \sum_{s=1}^{\infty} \log(1 + \zeta^{-2}\mu_s n_o) && \text{(Refer to [60, Lemma 1])} \\ &\leq \{\log(1 + n_o/\zeta^2) + 1\}2d^*. \end{aligned}$$

From the second line to the third line, we follow in the proof of theorem 24. \square

E Implementation Details

Here we detail all environment details and hyperparameters used for the experiments in the main text.

E.1 Environment Details

All environments have a maximum horizon length of 500 timesteps. We achieve this by reducing the data collection frequency of the base 1000 horizon environments. We also remove all contact information from the observation and the reward. Finally, to be able to compute the ground truth reward from the state, we add the velocity of the center of mass into the state.

Table 3: Observation and action space dimensions for each of the environments

Environment	Observation Space Dimension	Action Space Dimension
Hopper	12	3
Walker2d	18	6
HalfCheetah	18	6
Ant	29	8
Humanoid	47	17

Table 4: Ground truth environment reward function used to train the expert and behavior policies as well as evaluate the performance in the learning curves. At time t , \dot{x}_t is the velocity of the center of mass in the x -axis, \mathbf{a}_t is the action vector, and z_t is the position of the center of mass in the z -axis.

Environment	Ground Truth Reward Function
Hopper	$\dot{x}_t - 0.1\ \mathbf{a}_t\ _2^2 - 3.0 \times (z_t - 1.3)^2$
Walker2d	$\dot{x}_t - 0.1\ \mathbf{a}_t\ _2^2 - 3.0 \times (z_t - 0.57)^2$
HalfCheetah	$\dot{x}_t - 0.1\ \mathbf{a}_t\ _2^2$
Ant	$\dot{x}_t - 0.1\ \mathbf{a}_t\ _2^2 - 3.0 \times (z_t - 1.3)^2$
Humanoid	$1.25 \times \dot{x}_t - 0.1\ \mathbf{a}_t\ _2^2 + 5 \times \text{bool}(1.0 \leq z_t \leq 2.0)$

E.2 Dynamics Ensemble Architecture and Model Learning

For all of our experiments we use an ensemble of four dynamics models with each model parameterized by a feed-forward neural network with two hidden layers containing 1024 units. The learned model does not predict next state, but instead predicts the normalized difference between the next state and the current state, $s_{t+1} - s_t$. The activation function used at each layer is ReLU. We train all of our ensembles using Adam with learning rate 5×10^{-5} and otherwise default hyperparameters. We train each dynamics model for 300 epochs on just the offline dataset for all of our experiments. Please see Table 5 for all values.

Table 5: All hyperparameters used for dynamics model learning

Hyperparameter	Value
Hidden Layers	(1024, 1024)
Activation	ReLU
Optimizer	Adam
Learning Rate	5×10^{-5}
Batch Size	256
Epochs	300

E.3 Policy Architecture and TRPO Details

We use the open source NPG/TRPO implementation, MJRL [53]. The policy network and the value network are feedforward neural networks with two hidden layers containing 32 and 128 hidden units respectively. Both networks use a tanh activation function with the policy network outputting a Gaussian distribution $\mathcal{N}(\mu(s), \sigma^2)$ where σ is a trainable parameter. We use Generalized Advantage Estimator (GAE) to estimate the advantages. Please see Table 6 for all values.

Table 6: TRPO/NPG hyperparameter values used in experiments.

Hyperparameter	Value
Policy Hidden Layers	(32, 32)
Critic Hidden Layers	(128, 128)
Batch Size	40000
Max KL Divergence	0.01
Discount γ	0.995
CG Iterations	25
CG Damping	1×10^{-5}
GAE λ	0.97
Critic Update Epochs	2
Critic Optimizer	Adam
Critic Learning Rate	1×10^{-4}
Critic L2 Regularization	1×10^{-4}
Policy Init Log Std.	-0.25
Policy Min Log Std.	-2.0
BC Regularization λ_{BC}	0.1

E.4 Hyperparameter Selection

For our core results, we tuned our hyperparameters on a randomly selected seed for Hopper-v2 and then applied it for all environments. For TRPO, we tuned the conjugate gradient iterations from values 10, 25, and 50; and the conjugate gradient damping coefficients from values 1e-2, 1e-3, 1e-4, and 1e-5. All other hyperparameters were default ones in the MJRL repository [53]. For the BC regularization coefficient we tested values of 0.1, 1e-2, 1e-3, 1e-4, and 1e-5. For the dynamics model architecture we tested 3 different hidden layer sizes: 256, 512, and 1024. Beyond this we used the exact same Adam optimizer and training procedure as [34].

E.5 Discriminator Update and Cost Function Details

We parameterize our discriminator as a linear function $f(s, a) = w^\top \phi(s, a)$, where $\phi(s, a)$ are Random Fourier Features [51] and w is the vector of parameters for the discriminator. Recall our objective,

$$\min_{\pi \in \Pi} \max_{f \in \mathcal{F}} \left[\mathbb{E}_{(s,a) \sim d_{\pi}^P} (f(s, a) + b(s, a)) - \mathbb{E}_{(s,a) \sim \mathcal{D}_e} [f(s, a)] \right] + \lambda_{BC} \cdot \mathbb{E}_{(s,a) \sim \mathcal{D}_e} [\ell(a, s, \pi)].$$

Now given a policy π , we can compute a closed form update for the discriminator parameters w like so

$$\begin{aligned} \max_{w: \|w\|_2^2 \leq \eta} L(w; \pi, \hat{P}, b, \mathcal{D}_e) &:= \mathbb{E}_{(s,a) \sim d_{\hat{P}}^\pi} (f(s, a) + b(s, a)) - \mathbb{E}_{(s,a) \sim \mathcal{D}_e} [f(s, a)] \\ \equiv \max_w L_\eta(w; \pi, \hat{P}, b, \mathcal{D}_e) &= \mathbb{E}_{(s,a) \sim d_{\hat{P}}^\pi} (f(s, a) + b(s, a)) - \mathbb{E}_{(s,a) \sim \mathcal{D}_e} [f(s, a)] - \frac{1}{2} \cdot (\|w\|_2^2 - \eta) \\ \Rightarrow \partial_w L_\eta(w; \pi, \hat{P}, b, \mathcal{D}_e) &= \mathbb{E}_{(s,a) \sim d_{\hat{P}}^\pi} [\phi(s, a)] - \mathbb{E}_{(s,a) \sim \mathcal{D}_e} [\phi(s, a)] - w \end{aligned}$$

where $\partial_w L_\eta(w; \pi, \hat{P}, b, \mathcal{D}_e)$ denotes the partial derivative of $L_\eta(\cdot)$ wrt to w . Setting the above expression to 0 and solving for w gives us the closed form solution. Note that even with the BC regularization constraint added into the objective, the solution will still hold.

Now for a given updated w_t , we have our cost function $c(s, a) = w_t^\top \phi(s, a) + b(s, a)$ where our penalty, $b(s, a)$, is the maximum discrepancy of our model ensemble predictions. To balance our penalty term with our cost term, we introduce a parameter λ_{penalty} to get the cost

$$c(s, a) = (1 - \lambda_{\text{penalty}}) \cdot w_t^\top \phi(s, a) + \lambda_{\text{penalty}} \cdot b(s, a).$$

In our experiments, λ_{penalty} was the only parameter we varied across environments.

Table 7: λ_{penalty} values used for each environment.

Environment	λ_{penalty}
Hopper	2.5×10^{-4}
Walker2d	1.0×10^{-7}
HalfCheetah	1.0×10^{-4}
Ant	1.0×10^{-4}
Humanoid	5.0×10^{-4}

F Additional Experiments

Recall that in our main experiments, we create an extremely small expert dataset containing expert (s, a) pairs by randomly sampling state-action pairs from an expert dataset consisting of state-action pairs from many expert trajectories, and we did that for the purpose of creating an expert dataset where BC almost fails completely. One may wonder what MIMO would do if we feed MIMO a complete single expert trajectory. We conduct such experiments in this section. Figure 4 shows the performance of MIMO with one expert trajectory using the *same hyperparameters* as before. All plots are shown averaged across five seeds. Note that MIMO is still performs well with one expert trajectory—matching or nearly matching the expert performance across all 5 continuous control tasks.




Figure 4: Performance of MIMO with one expert trajectory. Note MIMO performance just as well with trajectory inputs as with state-action pair sample inputs.