
A General Framework for Sample-Efficient Function Approximation in Reinforcement Learning

Zixiang Chen*

University of California, Los Angeles
chenzx19@cs.ucla.edu

Chris Junchi Li*

University of California, Berkeley
junchili@berkeley.edu

Angela Yuan*

University of California, Los Angeles
hzyuan@cs.ucla.edu

Quanquan Gu

University of California, Los Angeles
qgu@cs.ucla.edu

Michael I. Jordan

University of California, Berkeley
jordan@cs.berkeley.edu

Abstract

With the increasing need for handling large state and action spaces, general function approximation has become a key technique in reinforcement learning (RL). In this paper, we propose a general framework that unifies model-based and model-free RL, and an Admissible Bellman Characterization (ABC) class that subsumes nearly all Markov Decision Process (MDP) models in the literature for tractable RL. We propose a novel estimation function with decomposable structural properties for optimization-based exploration and the functional eluder dimension as a complexity measure of the ABC class. Under our framework, a new sample-efficient algorithm namely OPTimization-based ExpLOration with Approximation (OPERA) is proposed, achieving regret bounds that match or improve over the best-known results for a variety of MDP models. In particular, for MDPs with low Witness rank, under a slightly stronger assumption, OPERA improves the state-of-the-art sample complexity results by a factor of dH . Our framework provides a generic interface to design and analyze new RL models and algorithms.

1 Introduction

Reinforcement learning (RL) is a decision-making process that seeks to maximize the expected reward when an agent interacts with the environment [Sutton and Barto, 2018]. Over the past decade, RL has gained increasing attention due to its successes in a wide range of domains, including Atari games [Mnih et al., 2013], Go game [Silver et al., 2016], autonomous driving [Yurtsever et al., 2020], Robotics [Kober et al., 2013], etc. Existing RL algorithms can be categorized into value-based algorithms such as Q-learning [Watkins, 1989] and policy-based algorithms such as policy gradient [Sutton et al., 1999]. They can also be categorized as a model-free approach where one directly models the value function classes, or alternatively, a model-based approach where one needs to estimate the transition probability.

Due to the intractably large state and action spaces that are used to model the real-world complex environment, function approximation in RL has become prominent in both algorithm design and theoretical analysis. It is a pressing challenge to design sample-efficient RL algorithms with general function approximations. In the special case where the underlying Markov Decision Processes (MDPs) enjoy certain linear structures, several lines of works have achieved polynomial sample

*Equal contribution.

complexity and/or \bar{T} regret guarantees under either model-free or model-based RL settings. For linear MDPs where the transition probability and the reward function admit linear structure, Yang and Wang [2019] developed a variant of learning when granted access to a generative model, Jin et al. [2020] proposed an LSVI-UCB algorithm with $\mathcal{O}(d^3 H^3 \bar{T})$ regret bound and Zanette et al. [2020a] further extended the MDP model and improved the regret to $\mathcal{O}(dH \bar{T})$. Another line of work considers linear mixture MDPs Yang and Wang [2020], Modi et al. [2020], Jia et al. [2020], Zhou et al. [2021a], where the transition probability can be represented by a mixture of base models. In Zhou et al. [2021a], a $\mathcal{O}(dH \bar{T})$ minimax optimal regret was achieved with weighted linear regression and a Bernstein-type bonus. Other structural MDP models include the block MDPs [Du et al., 2019] and FLAMBE [Agarwal et al., 2020b] to mention a few.

In a more general setting, however, there is still a gap between the plethora of MDP models and sample-efficient RL algorithms that can learn the MDP model with function approximation. The question remains open as to what constitutes minimal structural assumptions that admit sample-efficient reinforcement learning. To answer this question, there are several lines of work along this direction. Russo and Van Roy [2013], Osband and Van Roy [2014] proposed a structural condition named eluder dimension, and Wang et al. [2020] extended the LSVI-UCB for general linear function classes with small eluder dimension. Another line of works proposed low-rank structural conditions, including Bellman rank [Jiang et al., 2017, Dong et al., 2020] and Witness rank [Sun et al., 2019]. Recently, Jin et al. [2021] proposed a complexity called Bellman eluder (BE) dimension, which unifies low Bellman rank and low eluder dimension. Concurrently, Du et al. [2021] proposed Bilinear Classes, which can be applied to a variety of loss estimators beyond vanilla Bellman error. Very recently, Foster et al. [2021] proposed Decision-Estimation Coefficient (DEC), which is a necessary and sufficient condition for sample-efficient interactive learning. To apply DEC to RL, they proposed a RL class named Bellman Representability, which can be viewed as a generalization of the Bilinear Class. Nevertheless, Sun et al. [2019] is limited to model-based RL, and Jin et al. [2021] is restricted to model-free RL. The only frameworks that can unify both model-based and model-free RL are Du et al. [2021] and Foster et al. [2021], but their sample complexity results when restricted to special MDP instances do not always match the best-known results. Viewing the above gap, we aim to answer the following question:

Is there a unified framework that includes all model-free and model-based RL classes while maintaining sharp sample efficiency?

In this paper, we tackle this challenging question and give a nearly affirmative answer to it. We summarize our contributions as follows:

- We propose a general framework called Admissible Bellman Characterization (ABC) that covers a wide set of structural assumptions in both model-free and model-based RL, such as linear MDPs, FLAMBE, linear mixture MDPs, kernelized nonlinear regulator [Kakade et al., 2020], etc. Furthermore, our framework encompasses comparative structural frameworks such as the low Bellman eluder dimension and low Witness rank.
- Under our ABC framework, we design a novel algorithm, OPTimization-based ExploRation with Approximation (OPERA), based on maximizing the value function while constrained in a small confidence region around the model minimizing the estimation function.
- We apply our framework to several specific examples that are known to be not sample-efficient with value-based algorithms. For the kernelized nonlinear regulator (KNR), our framework is the first general framework to derive a \bar{T} regret-bound result. For the witness rank, our framework yields a sharper sample complexity with a mild additional assumption compared to prior works.

We visualize and compare prevailing sample-efficient RL frameworks and ours in Figure 1. We can see that both the general Bilinear Class and our ABC frameworks capture most existing MDP classes, including the low Witness rank and the KNR models.

Notation. For a state-action sequence s_1, a_1, \dots, s_H in our given context, we use $\mathcal{H} := \{s_1, a_1, \dots, s_H\}$ to denote the σ -algebra generated by trajectories up to step H . Let π denote the policy of following the max-action strategy induced by hypothesis $W = f^i$ we write π^i as π^i for notational simplicity. We write $s_{1:H}$ to indicate the state-action sequence are generated by step H by following policy (π^j) and transition probabilities $P(s; a)$ of the

²In this paper, we use FLAMBE to refer to both the algorithm and the low-rank MDP with unknown feature mappings.

Figure 1: Venn-Diagram Visualization of Prevailing Sample-Efficient RL Classes. As by far the richest concept, the DEC framework is both a necessary and sufficient condition for sample-efficient interactive learning. BE dimension is a rich class that subsumes both low Bellman rank and low eluder dimension and addresses almost all model-free RL classes. The generalized Bilinear Class captures model-based RL settings including KNRs, linear mixture MDPs and low Witness rank MDPs, yet precludes some eluder-dimension based models. Bellman Representability is another unified framework that subsumes the vanilla bilinear classes but fails to capture KNRs and low Witness rank MDPs. Our ABC class encloses both generalized Bilinear Class and Bellman Representability and subsumes almost all known solvable MDP cases, with the exception of state-action aggregation and deterministic linear MDP models, which neither Bilinear Class nor our ABC class captures.

underlying MDP model M . We also write a_h to mean $a_h(j, s_h)$ for the h -th step. Let $\|\cdot\|_2$ denote the ℓ_2 -norm and $\|\cdot\|_1$ the ℓ_1 -norm of a given vector. Other notations will be explained at their first appearances.

2 Preliminaries

We consider a finite-horizon, episodic Markov Decision Process (MDP) defined by the tuple $(S; A; P; r; H)$, where S is the space of feasible states, A is the action space, H is the horizon in each episode defined by the number of action steps in one episode, and $P_{h,2[H]}$ is defined for every $h \in [H]$ as the transition probability from the current state-action pair $(s_h, a_h) \in S \times A$ to the next state $s_{h+1} \in S$. We use $r_h(s; a) \in [0, 1]$ to denote the reward received at step $h \in [H]$ when taking action a at states s and assume throughout this paper that for any possible trajectories, $\sum_{h=1}^H r_h(s_h; a_h) \in [0, 1]$.

A deterministic policy π is a sequence of functions $\pi_h : S \rightarrow A_{h,2[H]}$, where each π_h specifies a strategy at step h . Given a policy π , the action-value function is defined to be the expected cumulative rewards where the expectation is taken over the trajectory distribution generated by $(P_h(j, s_h; a_h); \pi_h(j, s_h))_{h \in [H]}$ as

$$Q_h(s; a) := \mathbb{E} \sum_{h^0=h}^H r_{h^0}(s_{h^0}; a_{h^0}) \mid s_h = s; a_h = a :$$

Similarly, we define the state-value function for policy π as the expected cumulative rewards as

$$V_h(s) := \mathbb{E} \sum_{h^0=h}^H r_{h^0}(s_{h^0}; a_{h^0}) \mid s_h = s :$$

We use π^* to denote the optimal policy that satisfies $Q_{\pi^*}(s) = \max_{\pi} V_{\pi}(s)$ for all $s \in S$ [Puterman, 2014]. For simplicity, we abbreviate Q_{π^*} as V_h and Q_h as Q_h . Moreover, for a sequence of value functions $\{Q_h\}_{h \in [H]}$, the Bellman operator at step h is defined as:

$$(T_h Q_{h+1})(s; a) = r_h(s; a) + \mathbb{E}_{s^0 \sim P_h(j; s; a)} \max_{a^0 \in A} Q_{h+1}(s^0; a^0)$$

We also call $Q_h - (T_h Q_{h+1})$ the Bellman error (or Bellman residual). The goal of an RL algorithm is to find an ϵ -optimal policy such that $V_1(s_1) - V_1^*(s_1) \leq \epsilon$. For an RL algorithm that updates the policy π^t for T iterations, the cumulative regret is defined as

$$\text{Regret}(T) := \sum_{t=1}^T \sum_h \sum_i V_1(s_1) - V_1^{\pi^t}(s_1);$$

Hypothesis Classes Following Du et al. [2021], we define the hypothesis class for both model-free and model-based RL. Generally speaking, a hypothesis class is a set of functions that are used to estimate the value functions (for model-free RL) or the transitional probability and reward (for model-based RL). Specifically, a hypothesis class on a finite-horizon MDP is the Cartesian product of H hypothesis classes $\mathcal{F} := \mathcal{F}_1 \times \dots \times \mathcal{F}_H$ in which each hypothesis $f = (f_h)_{h \in [H]} \in \mathcal{F}$ can be identified by a pair of value functions $(Q_{h,f}; V_{h,f})_{h \in [H]}$. Based on the value function pair, it is natural to introduce the greedy policy $\pi_f^h(s) = \arg \max_{a \in \mathcal{A}} Q_{h,f}(s; a)$ at each step $h \in [H]$, and the corresponding (s) as the sequence of time-dependent policies $\pi_f^{H:1}$. An example of a model-free hypothesis class is defined by a sequence of action-value functions $(Q_{h,f})_{h \in [H]}$. The corresponding state-value function is given by:

$$V_{h,f}(s) = E_{a \sim \pi_f^h} [Q_{h,f}(s; a)];$$

In another example that falls under the model-based RL setting, where for each hypothesis f we have the knowledge of the transition matrix and the reward function $(P; r)$. We define the value function $Q_{h,f}$ corresponding to hypothesis f as the optimal value function following $\pi_f := (\pi_f^h; r)$:

$$Q_{h,f}(s; a) = Q_{h;M_f}(s; a) \quad \text{and} \quad V_{h,f}(s) = V_{h;M_f}(s);$$

We also need the following realizability assumption that requires the true model (model-based RL) or the optimal value function (model-free RL) to belong to the hypothesis class

Assumption 1 (Realizability) For an MDP model M and a hypothesis class \mathcal{F} , we say that the hypothesis class \mathcal{F} is realizable with respect to M if there exists $f^* \in \mathcal{F}$ such that for any $h \in [H]$, $Q_h(s; a) = Q_{h,f^*}(s; a)$. We call such f^* an optimal hypothesis

This assumption has also been made in the Bilinear Classes [Du et al., 2021] and low Bellman eluder dimension frameworks [Jin et al., 2021]. We also define the covering number of \mathcal{F} under a well-defined metric ρ of a hypothesis class \mathcal{F} :

Definition 2 (ϵ -covering Number of Hypothesis Class) For any $\epsilon > 0$ and a hypothesis class \mathcal{F} , we use $N_{\mathcal{F}}(\epsilon)$ to denote the ϵ -covering number, which is the smallest possible cardinality of $\mathcal{F}' \subseteq \mathcal{F}$ such that for any $f \in \mathcal{F}$ there exists $f' \in \mathcal{F}'$ such that $\rho(f, f') \leq \epsilon$.

Functional Eluder Dimension. We proceed to introduce our new complexity measure, functional eluder dimension, which generalizes the concept of eluder dimension first proposed in bandit literature [Russo and Van Roy, 2013, 2014]. It has since become a widely used complexity measure for function approximations in RL [Wang et al., 2020, Ayoub et al., 2020, Jin et al., 2021, Foster et al., 2021]. Here we revisit its definition:

Definition 3 (Eluder Dimension) For a given space X and a class \mathcal{F} of functions defined on X , the eluder dimension $\text{dim}_E(\mathcal{F}; \epsilon)$ is the length of the existing longest sequence $x_1, \dots, x_n \in X$ satisfying for some $\epsilon_0 > 0$ and any $2 \leq t \leq n$, there exist $f_1, f_2 \in \mathcal{F}$ such that $\sum_{i=1}^{t-1} (f_1(x_i) - f_2(x_i))^2 \leq \epsilon_0$ while $|f_1(x_t) - f_2(x_t)| > \epsilon$.

The eluder dimension is usually applied to the state-action space $S \times A$ and the corresponding value function class $\mathcal{F} : S \times A \rightarrow \mathbb{R}$ [Jin et al., 2021, Wang et al., 2020]. We extend the concept of eluder dimension as a complexity measure of the hypothesis class, namely functional eluder dimension, which is formally defined as follows.

³For example for model-free cases where $(f; g) = \max_{h \in [H]} \sum_{p \in \mathcal{P}_h} k_f^h g_p^h k_1$. For model-based RL where $(P; Q) = \max_{h \in [H]} \sum_{p \in \mathcal{P}_h} (dP_h - dQ_h)^2$ which is the maximal (squared) Hellinger distance between two probability distribution sequences.

Definition 4 (Functional Eluder Dimension) For a given hypothesis class \mathcal{F} and a function G defined on $\mathcal{F} \times \mathcal{F}$, the functional eluder dimension (FE dimension) $\dim_{FE}(\mathcal{F}; G; \epsilon)$ is the length of the existing longest sequence $g_1, \dots, g_n \in \mathcal{F}$ satisfying for some θ_0 and any $2 \leq t \leq n$, there exists $g \in \mathcal{F}$ such that $\sum_{i=1}^{t-1} (G(g; f_i))^2 \leq \epsilon$ while $|G(g; f_t)| > \epsilon$. Function G is dubbed as the coupling function.

The notion of functional eluder dimension introduced in Definition 4 is generalizable in a straightforward fashion to a sequence $\mathcal{G} := \{G_h\}_{h \in [H]}$ of coupling functions: we simply set $\dim_{FE}(\mathcal{F}; \mathcal{G}; \epsilon) = \max_{h \in [H]} \dim_{FE}(\mathcal{F}; G_h; \epsilon)$ to denote the FE dimension of \mathcal{G} . The Bellman eluder (BE) dimension recently proposed by [Jin et al., 2021] is in fact a special case of FE dimension with a specific choice of coupling function sequence. As will be shown later, our framework based on FE dimension with respect to the corresponding coupling function captures many specific MDP instances such as the kernelized nonlinear regulator (KNR) [Kakade et al., 2020] and the generalized linear Bellman complete model [Wang et al., 2019], which are not captured by the framework of low BE dimension. As we will see in later sections, introducing the concept of FE dimension allows the coverage of a strictly wider range of MDP models and hypothesis classes.

3 Admissible Bellman Characterization Framework

In this section, we introduce the framework of admissible Bellman characterization.

3.1 Admissible Bellman Characterization

Given an MDP \mathcal{M} , a sequence of states and actions s_1, a_1, \dots, s_H , two hypothesis classes \mathcal{F} and \mathcal{G} satisfying the realizability assumption (Assumption 1) and a discriminator function class $\mathcal{V} = \{v(s; a; s') : S \times A \times S \rightarrow \mathbb{R}^d\}$, the estimation function $\hat{v} = \hat{v}_{h,f} : \mathcal{G}_{h2[H],f} \times \mathcal{F} \rightarrow \mathbb{R}^d$ is an \mathbb{R}^d -valued function defined on the set consisting of $\mathcal{G} := \{(s_h, a_h; s_{h+1}) \in S \times A \times S, f \in \mathcal{F}, g \in \mathcal{G}\}$ and serves as a surrogate loss function of the Bellman error. Note that our estimation function is a vector-valued function, and is more general than the scalar-valued estimation function (or discrepancy function) used in Foster et al. [2021], Du et al. [2021]. The discriminator v originates from the function class the Integral Probability Metrics (IPM) [Müller, 1997] is taken with respect to (as a metric between two distributions), and is also used in the definition of Witness rank [Sun et al., 2019]. We use a coupling function $G_{h,f}(f; g)$ defined on $\mathcal{F} \times \mathcal{F}$ to characterize the interaction between two hypotheses $f, g \in \mathcal{F}$. The subscript h is an indicator of the true model and is by default unchanged throughout the context. When the two hypotheses coincide, our characterization of the coupling function reduces to the Bellman error.

Definition 5 (Admissible Bellman Characterization) Given an MDP \mathcal{M} , two hypothesis classes $\mathcal{F}; \mathcal{G}$ satisfying the realizability assumption (Assumption 1) and \mathcal{G} , an estimation function $\hat{v}_{h,f} : (S \times A \times S) \times \mathcal{F} \times \mathcal{G} \rightarrow \mathbb{R}^d$, an operation policy π_{op} and a constant $\epsilon \in (0, 1]$, we say that \mathcal{G} is an admissible Bellman characterization of $(\mathcal{M}; \mathcal{F}; \mathcal{G}; \hat{v})$ if the following conditions hold:

(i) (Dominating Average Estimation Function) For any $f, g \in \mathcal{F}$

$$\max_{v \in \mathcal{V}} E_{s_h, a_h \sim \pi_{op}} E_{s_{h+1}} [\hat{v}_{h,f}(s_h, a_h; s_{h+1}; f, v) - \hat{v}_{h,f}(s_h, a_h; s_{h+1}; g, v)]^2 \leq \epsilon (G_{h,f}(f; g))^2;$$

(ii) (Bellman Dominance) For any $(h, f) \in [H] \times \mathcal{F}$,

$$E_{s_h, a_h \sim \pi_{op}} [Q_{h,f}(s_h, a_h) - r(s_h, a_h) - V_{h+1,f}(s_{h+1})] \leq \epsilon |G_{h,f}(f; f)|;$$

We further say $(\mathcal{M}; \mathcal{F}; \mathcal{G}; \hat{v}; \epsilon)$ is an ABC class if \mathcal{G} is an admissible Bellman characterization of $(\mathcal{M}; \mathcal{F}; \mathcal{G}; \hat{v})$.

In Definition 5, one can choose either $\pi_{op} = \pi_g$ or $\pi_{op} = \pi_f$. We refer readers to §D for further explanations on π_{op} . The ABC class is quite general and de facto covers many existing MDP models; see §3.2 for more details.

⁴Indeed, when the coupling function is chosen as the expected Bellman error $G_{h,f}(f; g) := E_{s_h, a_h} [Q_{h,g}(s_h, a_h) - r(s_h, a_h) - V_{h+1,g}(s_{h+1})]$ where T_h denotes the Bellman operator, we recover the definition of BE dimension [Jin et al., 2021], i.e. $\dim_{FE}(\mathcal{F}; \mathcal{G}; \epsilon) = \dim_{BE}(\mathcal{F}; \mathcal{G}; \epsilon)$.

⁵We assume $\mathcal{F} \subseteq \mathcal{G}$ throughout this paper and in the general case where $\mathcal{F} \not\subseteq \mathcal{G}$, we overload $\mathcal{G} := \mathcal{F} \cup \mathcal{G}$.

Comparison with Existing MDP Classes. Here we compare our ABC class with three recently proposed MDP structural classes: Bilinear Classes [Du et al., 2021], low Bellman eluder dimension [Jin et al., 2021], and Bellman Representability [Foster et al., 2021].

- **Bilinear Classes.** Compared to the structural framework of Bilinear Class in Du et al. [2021, Definition 4.3], Definition 5 of Admissible Bellman Characterization does not require a bilinear structure and recovers the Bilinear Class when we set $\psi(f; g) = \mathbb{E}_{h \sim W_h(g)} [W_h(f); X_h(f)]$. Our ABC class is strictly broader than the Bilinear Class since the latter does not capture low eluder dimension models, and our ABC class does. In addition, the ABC class admits an estimation function that is vector-valued and the corresponding algorithm achieves $\bar{\epsilon}$ -regret for KNR case while the BiLin-UCB algorithm for Bilinear Classes [Du et al., 2021] does not.
- **Low Bellman Eluder Dimension.** Definition 5 subsumes the MDP class of low BE dimension when $\psi_{h,f}(o_h; f_{h+1}; g_h; v) := Q_{h,g}(s_h; a_h) - r_h - V_{h+1,f}(s_{h+1})$. Moreover, our definition unifies the V-type and Q-type problems under the same framework by the notion of ψ . We will provide a more detailed discussion on this in §3.2. Our extension from the concept of the Bellman error to estimation function (i.e. the surrogate of the Bellman error) enables us to accommodate model-based RL for linear mixture MDPs, KNR model, and low Witness rank.
- **Bellman Representability.** Foster et al. [2021] proposed DEC framework which is another MDP class that unifies both the Bilinear Class and the low BE dimension. Indeed, our ABC framework introduced in Definition 5 shares similar spirits with the Bellman Representability Definition F.1 in Foster et al. [2021]. Nevertheless, our framework and theirs bifurcate from the base point: our work studies an optimization-based exploration instead of the posterior sampling-based exploration in Foster et al. [2021]. Structurally different from their DEC framework, our ABC requires estimation functions to be vector-valued, introduces the discriminator function, and imposes the weaker Bellman dominance property (i) in Definition 5 than the corresponding one as in Foster et al. [2021, Eq. (166)]. In total, this allows broader choices of coupling functions as well as our ABC class (with low FE dimension) to include as special instances both low Witness rank and KNR models, which are not captured in Foster et al. [2021].

Decomposable Estimation Function Now we introduce the concept of decomposable estimation function which generalizes the Bellman error in earlier literature and plays a pivotal role in our algorithm design and analysis.

Definition 6 (Decomposable Estimation Function) A decomposable estimation function $(S, A, S') \in \mathcal{F} \times \mathcal{G} \times \mathcal{V}$ is a function with bounded ℓ_2 -norm such that the following two conditions hold:

- (i) (Decomposability) There exists an operator T that maps between two hypothesis classes $\mathcal{F} \times \mathcal{G} \times \mathcal{V}$ such that for any $(f, g, v) \in \mathcal{F} \times \mathcal{G} \times \mathcal{V}$ and all possible o_h

$$\psi_{h,f}(o_h; f_{h+1}; g_h; v) - \mathbb{E}_{s_{h+1}} [\psi_{h,f}(o_h; f_{h+1}; g_h; v) | s_h; a_h] = \psi_{h,f}(o_h; f_{h+1}; T(f)_h; v):$$

Moreover, iff $\psi = f$, then $T(f) = f$ holds.

- (ii) (Global Discriminator Optimality) For any $(f, g, v) \in \mathcal{F} \times \mathcal{G} \times \mathcal{V}$ there exists a global maximum $\eta(f) \in \mathcal{V}$ such that for any $(h; f^0; g; v) \in \mathcal{H} \times \mathcal{F} \times \mathcal{G} \times \mathcal{V}$ and all possible o_h

$$\mathbb{E}_{s_{h+1}} [\psi_{h,f}(o_h; f_{h+1}; f_h; v_h(f)) | s_h; a_h] \leq \mathbb{E}_{s_{h+1}} [\psi_{h,f}(o_h; f_{h+1}; f_h; v) | s_h; a_h]:$$

Compared with the discrepancy function or estimation function used in prior work [Du et al., 2021, Foster et al., 2021], our estimation function (EF) admits the unique properties listed as follows:

- Our EF enjoys a decomposable property inherited from the Bellman error — intuitively speaking, the decomposability can be seen as a property shared by all functions in the form of the difference of a J_h -measurable function and a J_{h+1} -measurable function;
- Our EF involves a discriminator class and assumes the global optimality of the discriminator on all $(s_h; a_h)$ pairs;
- Our EF is a vector-valued function which is more general than a scalar-valued estimation function (or the discrepancy function).

⁶The decomposability item (i) in Definition 6 directly implies that a Generalized Completeness condition similar to Assumption 14 of Jin et al. [2021] holds.

We remark that when $\rho = g$, $E_{S_{h+1}} [\rho_{h,f}(O_h; f_{h+1}; g_h; v) | S_h; a_h]$ measures the discrepancy in optimality between ρ and f . In particular, when $\rho = f$, $E_{S_{h+1}} [\rho_{h,f}(O_h; f_{h+1}; f_h; v) | S_h; a_h] = 0$. Consider a special case when $\rho_{h,f}(O_h; f_{h+1}; g_h; v) := Q_{h,g}(S_h; a_h) - r(S_h; a_h) - V_{h+1,f}(S_{h+1})$. Then the decomposability (i) in Definition 6 reduces to

$$[Q_{h,g}(S_h; a_h) - r(S_h; a_h) - V_{h+1,f}(S_{h+1})] - [Q_{h,g}(S_h; a_h) - (T_h V_{h+1})(S_h; a_h)] = (T_h V_{h+1})(S_h; a_h) - r(S_h; a_h) - V_{h+1,f}(S_{h+1}).$$

In addition, we make the following Lipschitz continuity assumption on the estimation function.

Assumption 7 (Lipschitz Estimation Function) There exists $\alpha_L > 0$ such that for any $(h; f^0, f; g; v) \in [H] \times F \times G \times V$, $(\rho, g; \theta; \rho^0) \in 2F \times G \times V \times F$ and all possible θ_h ,

$$\begin{aligned} & \rho_{h,f}(; f; g; v) - \rho_{h,f}(; \rho; g; v) \leq L(f, \rho); \quad k_{h,f}(\rho; f; g; v) - \rho_{h,f}(; f; g; v) \leq L(g; g); \\ & k_{h,f}(\rho; f; g; v) - \rho_{h,f}(; f; g; \theta) \leq Lkv - \theta k_1; \quad \rho_{h,f}(; f; g; v) - \rho_{h,\rho^0}(; f; g; v) \leq L(f^0, \rho^0); \end{aligned}$$

Note that we have omitted the subscript hypotheses in Assumption 7 for notational simplicity. We further define the induced estimation function class as $\rho_{h,f}(; f; g; v) : (h; f^0, f; g; v) \in [H] \times F \times G \times V \times G$. We can show that under Assumption 7, the covering number of the induced estimation function class can be upper bounded as $\mathcal{N}_{[]}(\epsilon) \leq N_F^2(\frac{\epsilon}{4L}) N_G(\frac{\epsilon}{4L}) N_V(\frac{\epsilon}{4L})$, where $N_F(\cdot)$; $N_G(\cdot)$; $N_V(\cdot)$ are the ϵ -covering number of F , G and V , respectively. Later in our theoretical analysis in §4, our regret upper bound will depend on the growth rate of the covering number or the metric entropy $\log N_F(\cdot)$.

3.2 MDP Instances in the ABC Class

In this subsection, we present a number of MDP instances that belong to ABC class with low FE dimension. As we have mentioned before, for all special cases with $\rho_{h,f}(O_h; f_{h+1}; g_h; v) := Q_{h,g}(S_h; a_h) - r_h - V_{h+1,f}(S_{h+1})$, both conditions in Definition 5 are satisfied automatically with $G_{h,f}(f; g) = E_{S_h, g; a_h} [Q_{h,f}(S_h; a_h) - r_h - V_{h+1,f}(S_{h+1})]$. The FE dimension under this setting recovers the the BE dimension. Thus, all model-free RL models with low BE dimension [Jin et al., 2021] belong to our ABC class with low FE dimension. In the rest of this subsection, our focus shifts to the model-based RLs that belong to the ABC class: linear mixture MDPs, low Witness rank, and kernelized nonlinear regulator.

Linear Mixture MDPs. We start with a model-based RL with a linear structure called linear mixture MDP [Modi et al., 2020, Ayoub et al., 2020, Zhou et al., 2021b]. We defer the formal definition of linear mixture MDPs to §E.1 and only provide the following proposition, which shows that linear mixture MDPs belong to the ABC class with low FE dimension.

Proposition 8 (Linear Mixture MDP ABC with Low FE Dimension) The linear mixture MDP model belongs to the ABC class with estimation function

$$\rho_{h,f}(O_h; f_{h+1}; g_h; v) = \sum_{g_h} \rho_{h,g}(S_h; a_h) + \sum_{s^0} (S_h; a_h; s^0) V_{h+1,f}(s^0) - r_h - V_{h+1,f}(S_{h+1}); \quad (3.1)$$

and coupling function $G_{h,f}(f; g) = \sum_{g_h} \rho_{h,g}(S_h; a_h) + \sum_{s^0} (S_h; a_h; s^0) V_{h+1,f}(s^0)$. Moreover, it has a low FE dimension.

Low Witness Rank. The following definition is a generalized version of the witness rank in Sun et al. [2019], where we require the discriminator class to be complete meaning that the assemblage of functions by taking the value $\rho(S; a)$ from different functions also belongs to \mathcal{D} .

Definition 9 (Witness Rank) For an MDPM, a given symmetric and complete discriminator class $\mathcal{D} = \{V_h, g_h\}_{h \in [H]}$, $V_h \in \mathcal{S}^A \times \mathcal{S}^B$ and a hypothesis class \mathcal{F} , we define the Witness rank W_h as the smallest d such that for any two hypothesis sets $\mathcal{F}_1, \mathcal{F}_2 \subseteq \mathcal{F}$, there exist two mappings $\phi_h : \mathcal{F}_1 \rightarrow \mathcal{R}^d$ and $\psi_h : \mathcal{F}_2 \rightarrow \mathcal{R}^d$ and a constant $\beta \in (0, 1]$, the following inequalities hold for all $h \in [H]$:

$$\max_{V_h} E_{S_h, f; a_h} [E_{\theta \sim g_h} V_h(S_h; a_h; \theta) - E_{\theta \sim P_h} V_h(S_h; a_h; \theta)] \leq W_h(g); X_h(f); \quad (3.2)$$

$$E_{S_h, f; a_h} [E_{\theta \sim g_h} V_{h+1,g}(\theta) - E_{\theta \sim P_h} V_{h+1,g}(\theta)] \leq W_h(g); X_h(f); \quad (3.3)$$

The following proposition shows that low Witness rank belongs to our ABC class with low FE dimension.

Proposition 10 (Low Witness Rank ABC with Low FE Dimension) The low Witness rank model belongs to the ABC class with estimation function

$$\rho_{h,f}(a_h; f_{h+1}; g_h; v) = E_{\theta \sim g_h} v(s_h; a_h; \theta) - v(s_h; a_h; s_{h+1}); \quad (3.4)$$

and coupling function $G_{h,f}(f; g) = \mathbb{E} [v_h(g); X_h(f)]$. Moreover, it has a low FE dimension.

Kernelized Nonlinear Regulator. The kernelized nonlinear regulator (KNR) proposed recently by Mania et al. [2020], Kakade et al. [2020] models a nonlinear control dynamics on an RKHS finite or countably infinite dimensions. Under the KNR setting, given current state s_h at step $h \in [H]$ and a known feature mapping: $S \rightarrow A \times H$, the subsequent state obeys a Gaussian distribution with mean vector $U_h(s_h; a_h)$ and homoskedastic covariance Σ_h , where $U_h \in \mathbb{R}^{d_s \times H}$ and $\Sigma_h \in \mathbb{R}^{d_s \times d_s \times H}$ are true model parameters and d_s is the dimension of the state space. Mathematically, we have for each $h = 1; \dots; H$,

$$s_{h+1} = U_h(s_h; a_h) + \epsilon_{h+1}; \quad \text{where } \epsilon_{h+1} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0; \Sigma_h); \quad (3.5)$$

Furthermore, we assume bounded reward $r(s; a) \in [0, 1]$ and uniformly bounded feature map $\| \phi(s; a) \|_2 \leq B$. The following proposition shows that KNR belongs to the ABC class with low FE dimension.

Proposition 11 (KNR ABC with Low FE Dimension) KNR belongs to the ABC class with estimation function

$$\rho_{h,f}(a_h; f_{h+1}; g_h; v) = U_{h,g}(s_h; a_h) - v(s_{h+1}); \quad (3.6)$$

and coupling function $G_{h,f}(f; g) := \frac{1}{q} \frac{1}{E_{s_h; a_h} \sim g} \mathbb{E} [k(U_{h,f}(s_h; a_h) - U_h(s_h; a_h)) k_2^2]$. Moreover, it has a low FE dimension.

Although the dimension of the RKHS can be infinite, our complexity analysis depends solely on its effective dimension.

4 Algorithm and Main Results

In this section, we present an RL algorithm for the ABC class. Then we present the regret bound of this algorithm, along with its implications to several MDP instances in the ABC class.

4.1 Opera Algorithm

We first present the **OPTimization-based ExpLOration with Approximation (OPERA)** algorithm in Algorithm 1, which finds an ϵ -optimal policy in polynomial time. Following earlier algorithmic art in the same vein e.g., GOLF [Jin et al., 2021], the core optimization step of OPERA is optimization-based exploration under the constraint of an identified confidence region; we additionally introduce an estimation policy π_{est} sharing the similar spirit as in Du et al. [2021]. Due to space limit, we focus on the Q-type analysis here and defer the V-type results to §D in the appendix.

Pertinent to the constrained optimization subproblem in (Eq.) of our Algorithm 1, we adopt the confidence region based on a general DEF, extending the Bellman-error-based confidence region used in Jin et al. [2021]. As a result of such an extension, our algorithm can deal with more complex models such as low Witness rank and KNR. Similar to existing literature on RL theory with general function approximation, our algorithm is in general computationally inefficient. Yet OPERA is oracle efficient given the oracle for solving the optimization problem in Line 3 of Algorithm 1. We will discuss its computational issues in detail in §E.1, §E.2 and §E.3.

4.2 Regret Bounds

We are ready to present the main theoretical results of our ABC class with low FE dimension:

⁷Here and throughout our paper we consider $\pi_{\text{est}} = \pi^*$ for Q-type models. For V-type models, we instead consider $\pi_{\text{est}} = U(A)$ to be the uniform distribution over the action space. Such a representation of estimation policy allows us to unify the Q-type and V-type models in a single analysis.

Algorithm 1 OPTimization-based ExploRation with Approximation (OPERA)

- 1: Initialize: $D_h = ?$ for $h = 1; \dots; H$
 - 2: for iteration = 1; 2; ...; T do
 - 3: Set $t := t$ where t is taken as $\arg\max_{f \in \mathcal{F}} Q_{1,f}(s_1; f(s_1))$ subject to

$$\max_{v \in \mathcal{V}} \sum_{i=1}^X \sum_{h,f} \left(\mathcal{O}_h^i; f_{h+1}; f_h; v \right)^2 \inf_{g \in \mathcal{G}_h} \sum_{i=1}^X \left(\mathcal{O}_h^i; f_{h+1}; g_h; v \right)^2 \quad \text{for all } h \in [H] \quad (4.1)$$
 - 4: For any $h \in [H]$, collect tuple $(r_h; s_h; a_h; s_{h+1})$ by rolling in s_h t and executing g_h est
 - 5: Augment $D_h = D_h \cup \{(r_h; s_h; a_h; s_{h+1})\}$
 - 6: end for
 - 7: Output: out uniformly sampled from $t \in \mathcal{G}_1^T$
-

Theorem 12 (Regret Bound of OPERA) For an MDP \mathcal{M} , hypothesis class $\mathcal{F}; \mathcal{G}$, a Decomposable Estimation Function \mathcal{O} satisfying Assumption 7, an admissible Bellman characterization \mathcal{G} . Suppose $(\mathcal{M}; \mathcal{F}; \mathcal{G}; \mathcal{O}; \mathcal{G})$ is an ABC class with low functional eluder dimension. For any $\epsilon \in (0; 1)$, we choose $\beta = O(\log(T) \sqrt{N_L(1-\epsilon)})$ in Algorithm 1. Then for the on-policy case when $op = est = t$, with probability at least $1 - \epsilon$, the regret is upper bounded by

$$\text{Regret}(T) = O \left(\frac{H}{\epsilon} \sqrt{T \dim_{FE}(\mathcal{F}; \mathcal{G}; \frac{1}{\beta})} \right) :$$

We defer the proof of Theorem 12, together with a corollary for sample complexity analysis, to §C in the appendix. We observe that the regret bound of the OPERA algorithm is dependent on both the functional eluder dimension \dim_{FE} and the covering number of the induced DEF class $N_L(\frac{1}{\beta})$. In the special case when DEF is chosen as the Bellman error, the relation $\dim_{FE}(\mathcal{F}; \mathcal{G}; \frac{1}{\beta}) = \dim_{BE}(\mathcal{F}; \mathcal{G}; \frac{1}{\beta})$ holds with \mathcal{F} being the function class induced by $f; f \in \mathcal{F}$, and our Theorem 12 reduces to the regret bound in Jin et al. [2021] (Theorem 15). We will provide a detailed comparison between our framework and other related frameworks in §A when applied to different MDP models in the appendix.

4.3 Implication for Specific MDP Instances

Here we focus on comparing our results applied to model-based RLs that are hardly analyzable in the model-free framework in §3.2. We demonstrate how OPERA can find near-optimal policies and achieve a state-of-the-art sample complexity under our new framework. Regret-bound analyses of linear mixture MDPs and several other MDP models can be found in §B in the appendix.

We highlight that Algorithm 1 not only provides a simple optimization-based scheme, recovers previous near-optimal algorithms in literature (Algorithms 2 and 4 in §E) when applied to specific MDP instances, but also reduces to a novel Algorithm 3 for low witness rank MDPs with improved sample complexity.

Low Witness Rank. We first provide a sample complexity result for the low Witness rank model structure. Let M and V be the cardinality of the model class \mathcal{M} and discriminator class \mathcal{V} , respectively, and W be the witness rank (Definition 9) of the model. We have the following sample complexity result for low Witness rank models.

Corollary 13 (Finite Witness Rank) For an MDP model \mathcal{M} with finite witness rank structure and any $\epsilon \in (0; 1)$, we choose $\beta = O(\log(T) \sqrt{MjVj})$ in Algorithm 1. With probability at least $1 - \epsilon$, Algorithm 1 outputs an ϵ -optimal policy out within $T = O(H^2 jAj W^2 \epsilon^{-2})$ trajectories.

Proof of Corollary 13 is delayed to §E. Compared with previous best-known sample complexity result of $O(H^3 W^2 jAj \log(T) \sqrt{MjVj}) = O(H^2 \epsilon^{-2})$ due to Sun et al. [2019], our sample complexity is superior by a factor of H up to a polylogarithmic prefactor in model parameters.

Kernel Nonlinear Regulator. Now we turn to the implication of Theorem 12 for learning KNR models. We have the following regret bound result for KNR.

⁸Hypothesis class reduces to model class [Sun et al., 2019] when restricted to model-based setting.

⁹The definition of witness rank adopts V -type representation and hence we can only derive the sample complexity of our algorithm. For detailed discussion on W -type cases, we refer readers to §D in the appendix.

Corollary 14 (KNR). For the KNR model in Eq(3.5) and any fixed $\beta \in (0; 1)$, we choose $\epsilon = O\left(\frac{2d \log^2(TH)}{\beta}\right)$ in Algorithm 1. With probability at least $1 - \beta$, the regret is upper bounded by $O\left(\frac{H^2}{\beta} \sqrt{dT}\right)$.

We remark that neither the low BE dimension nor the Bellman Representability classes admit the KNR model with a sharp regret bound. Among earlier attempts, Du et al. [2021, §6] proposed to use a generalized version of Bilinear Classes to capture models including KNR, Generalized Linear Bellman Complete, and finite Witness rank. Nevertheless, their characterization requires imposing monotone transformations on the statistic and yields a suboptimal (\sqrt{T}) regret bound. Our ABC class with low FE dimension is free of monotone operators, albeit that the coupling function for the KNR model is not of a bilinear form.

5 Conclusion and Future Work

In this paper, we proposed a unified framework that subsumes all Markov Decision Process (MDP) models in existing literature from model-based and model-free RLs. For the complexity analysis, we propose a new type of estimation function with the decomposable property for optimization-based exploration and use the functional eluder dimension with respect to an admissible Bellman characterization function as the complexity measure of our model class. In addition, we proposed a new sample-efficient algorithm, OPERA, which matches or improves the state-of-the-art sample complexity (or regret) results.

Nevertheless, we notice that some MDP instances are not covered by our framework such as the state-action aggregation, and the deterministic linear models where only Q has a linear structure. We leave it as a future work to include these MDP models.

References

- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 2011.
- Alekh Agarwal, Daniel Hsu, Satyen Kale, John Langford, Lihong Li, and Robert Schapire. Taming the monster: A fast and simple algorithm for contextual bandits. *International Conference on Machine Learning*, pages 1638–1646. PMLR, 2014.
- Alekh Agarwal, Mikael Henaff, Sham Kakade, and Wen Sun. Pc-pg: Policy cover directed exploration for provable policy gradient learning. *Advances in neural information processing systems*, 13399–13412, 2020a.
- Alekh Agarwal, Sham Kakade, Akshay Krishnamurthy, and Wen Sun. Flambe: Structural complexity and representation learning of low rank mdp. *Advances in neural information processing systems* 33:20095–20107, 2020b.
- Akshay Agrawal, Robin Verschueren, Steven Diamond, and Stephen Boyd. A rewriting system for convex optimization problems. *Journal of Control and Decision*, 5(1):42–60, 2018.
- Shipra Agrawal and Randy Jia. Optimistic posterior sampling for reinforcement learning: worst-case regret bounds. *Advances in Neural Information Processing Systems*, 30:605–615, 2017.
- Brian DO Anderson and John B Moore. *Optimal control: linear quadratic methods*. Courier Corporation, 2007.
- Peter Auer, Thomas Jaksch, and Ronald Ortner. Near-optimal regret bounds for reinforcement learning. *Advances in neural information processing systems*, 21:2005–2013, 2008.
- Alex Ayoub, Zeyu Jia, Csaba Szepesvari, Mengdi Wang, and Lin Yang. Model-based reinforcement learning with value-targeted regression. *International Conference on Machine Learning*, pages 463–474. PMLR, 2020.
- Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 263–272. JMLR.org, 2017.

- Peter L Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research* 3(Nov):463–482, 2002.
- Steven Bradtko. Reinforcement learning applied to linear quadratic regulation. *Advances in neural information processing systems* 5, 1992.
- Ronen I Brafman and Moshe Tennenholtz. R-max: a general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research* 3(Oct):213–231, 2002.
- Qi Cai, Zhuoran Yang, Chi Jin, and Zhaoran Wang. Provably efficient exploration in policy optimization. In *International Conference on Machine Learning*, pages 1283–1294. PMLR, 2020.
- Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic linear optimization under bandit feedback. 2008.
- Christoph Dann and Emma Brunskill. Sample complexity of episodic fixed-horizon reinforcement learning. *Advances in Neural Information Processing Systems* 28, 2015.
- Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics* 20(4):633–679, 2020.
- Steven Diamond and Stephen Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research* 17(83):1–5, 2016.
- Omar Darwiche Domingues, Pierre Ménard, Emilie Kaufmann, and Michal Valko. Episodic reinforcement learning in finite mdps: Minimax lower bounds revisited. *Algorithmic Learning Theory*, pages 578–598. PMLR, 2021.
- Kefan Dong, Jian Peng, Yining Wang, and Yuan Zhou. Root-n-regret for learning in markov decision processes with function approximation and low bellman rank. *Conference on Learning Theory*, pages 1554–1557. PMLR, 2020.
- Shi Dong, Benjamin Van Roy, and Zhengyuan Zhou. Provably efficient reinforcement learning with aggregated states. *arXiv preprint arXiv:1912.06366*, 2019.
- Simon Du, Akshay Krishnamurthy, Nan Jiang, Alekh Agarwal, Miroslav Dudik, and John Langford. Provably efficient rl with rich observations via latent state decoding. *International Conference on Machine Learning*, pages 1665–1674. PMLR, 2019.
- Simon Du, Sham Kakade, Jason Lee, Shachar Lovett, Gaurav Mahajan, Wen Sun, and Ruosong Wang. Bilinear classes: A structural framework for provable generalization in linear. *International Conference on Machine Learning*, pages 2826–2836. PMLR, 2021.
- Dylan J Foster, Sham M Kakade, Jian Qian, and Alexander Rakhlin. The statistical complexity of interactive decision making. *arXiv preprint arXiv:2112.13487*, 2021.
- Zeyu Jia, Lin Yang, Csaba Szepesvari, and Mengdi Wang. Model-based reinforcement learning with value-targeted regression. *Learning for Dynamics and Control*, pages 666–686. PMLR, 2020.
- Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, and Robert E Schapire. Contextual decision processes with low bellman rank are pac-learnable. *International Conference on Machine Learning*, pages 1704–1713. PMLR, 2017.
- Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is Q-learning provably efficient? In *Advances in Neural Information Processing Systems*, pages 4868–4878, 2018.
- Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. *Conference on Learning Theory*, pages 2137–2143. PMLR, 2020.
- Chi Jin, Qinghua Liu, and Sobhan Miryoosefi. Bellman eluder dimension: New rich classes of RL problems, and sample-efficient algorithms. *Advances in Neural Information Processing Systems* 34, 2021.

- Sham Kakade, Akshay Krishnamurthy, Kendall Lowrey, Motoya Ohnishi, and Wen Sun. Information theoretic regret bounds for online nonlinear control. *Advances in Neural Information Processing Systems* 33:15312–15325, 2020.
- Michael Kearns. Efficient noise-tolerant learning from statistical queries. *Journal of the ACM (JACM)* 45(6):983–1006, 1998.
- Jens Kober, J Andrew Bagnell, and Jan Peters. Reinforcement learning in robotics: A survey. *International Journal of Robotics Research* 32(11):1238–1274, 2013.
- Akshay Krishnamurthy, Alekh Agarwal, and John Langford. Pac reinforcement learning with rich observations. *Advances in Neural Information Processing Systems* 29:2016.
- Nick Littlestone. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine learning* 2(4):285–318, 1988.
- Horia Mania, Michael I Jordan, and Benjamin Recht. Active learning for nonlinear system identification with guarantees. *arXiv preprint arXiv:2006.10277*, 2020.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Aditya Modi, Nan Jiang, Ambuj Tewari, and Satinder Singh. Sample complexity of reinforcement learning using linearly combined model ensembles. *International Conference on Artificial Intelligence and Statistics* pages 2010–2020. PMLR, 2020.
- Alfred Müller. Integral probability metrics and their generating classes of functions. *Advances in Applied Probability* 29(2):429–443, 1997.
- Gergely Neu and Ciara Pike-Burke. A unifying view of optimism in episodic reinforcement learning. *Advances in Neural Information Processing Systems* 33:1392–1403, 2020.
- Ian Osband and Benjamin Van Roy. Model-based reinforcement learning and the eluder dimension. *Advances in Neural Information Processing Systems* 27:2014.
- David Pollard. *Convergence of stochastic processes*. Springer Science & Business Media, 2012.
- Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. Wiley & Sons, 2014.
- Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Online learning: Random averages, combinatorial parameters, and learnability. *Advances in Neural Information Processing Systems* 23, 2010.
- Daniel Russo and Benjamin Van Roy. Eluder dimension and the sample complexity of optimistic exploration. *Advances in Neural Information Processing Systems* 26:2013.
- Daniel Russo and Benjamin Van Roy. Learning to optimize via posterior sampling. *Mathematics of Operations Research* 39(4):1221–1243, 2014.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature* 529(7587):484–489, 2016.
- Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*, 2009.
- Wen Sun, Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, and John Langford. Model-based rl in contextual decision processes: Pac bounds and exponential improvements over model-free approaches. *Conference on learning theory* pages 2898–2933. PMLR, 2019.
- Richard S Sutton and Andrew G Barto. *Reinforcement Learning: An Introduction*. 2018.

- Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. *Advances in neural information processing systems* 12, 1999.
- Vladimir Vapnik. *The nature of statistical learning theory*. Springer science & business media, 1999.
- Nolan Wagener, Ching-An Cheng, Jacob Sacks, and Byron Boots. An online learning approach to model predictive control. *arXiv preprint arXiv:1902.08967*, 2019.
- Ruosong Wang, Russ R Salakhutdinov, and Lin Yang. Reinforcement learning with general value function approximation: Provably efficient approach via bounded eluder dimension. *Advances in Neural Information Processing Systems* 33:6123–6135, 2020.
- Yining Wang, Ruosong Wang, Simon S Du, and Akshay Krishnamurthy. Optimism in reinforcement learning with generalized linear function approximation. *arXiv preprint arXiv:1912.04136*, 2019.
- Christopher John Cornish Hellaby Watkins. *Learning from delayed rewards*. Thesis, Cambridge University, 1989.
- Grady Williams, Andrew Aldrich, and Evangelos Theodorou. Model predictive path integral control using covariance variable importance sampling. *arXiv preprint arXiv:1509.01149*, 2015.
- Lin Yang and Mengdi Wang. Sample-optimal parametric q-learning using linearly additive features. In *International Conference on Machine Learning*, pages 6995–7004. PMLR, 2019.
- Lin Yang and Mengdi Wang. Reinforcement learning in feature space: Matrix bandit, kernels, and regret bound. In *International Conference on Machine Learning*, pages 10746–10756. PMLR, 2020.
- Zhuoran Yang, Chi Jin, Zhaoran Wang, Mengdi Wang, and Michael I Jordan. On function approximation in reinforcement learning: Optimism in the face of large state space. *arXiv preprint arXiv:2011.04622*, 2020.
- Ekim Yurtsever, Jacob Lambert, Alexander Carballo, and Kazuya Takeda. A survey of autonomous driving: Common practices and emerging technologies. *IEEE Access* 8:58443–58469, 2020.
- Andrea Zanette and Emma Brunskill. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. *arXiv preprint arXiv:1901.00210*, 2019.
- Andrea Zanette, Alessandro Lazaric, Mykel Kochenderfer, and Emma Brunskill. Learning near optimal policies with low inherent bellman error. In *International Conference on Machine Learning*, pages 10978–10989. PMLR, 2020a.
- Andrea Zanette, Alessandro Lazaric, Mykel J Kochenderfer, and Emma Brunskill. Provably efficient reward-agnostic navigation with linear value iteration. *Advances in Neural Information Processing Systems* 33:11756–11766, 2020b.
- Zihan Zhang, Yuan Zhou, and Xiangyang Ji. Almost optimal model-free reinforcement learning via reference-advantage decomposition. *Advances in Neural Information Processing Systems* 33:15198–15207, 2020.
- Dongruo Zhou, Quanquan Gu, and Csaba Szepesvari. Nearly minimax optimal reinforcement learning for linear mixture markov decision processes. *Conference on Learning Theory*, pages 4532–4576. PMLR, 2021a.
- Dongruo Zhou, Jiafan He, and Quanquan Gu. Provably efficient reinforcement learning for discounted mdps with feature mapping. In *International Conference on Machine Learning*, pages 12793–12802. PMLR, 2021b.

Appendix

The appendix is organized as follows. §A discusses the related work, providing comparisons with previous frameworks based on both coverage and sharpness of sample complexity. §B compares our regret bound and sample complexity on specific examples and discusses several additional examples including reactive POMDPs, FLAMBE, LQR, and the generalized linear Bellman complete model. §C proves the main results (Theorem 12 and Corollary 27 on sample complexity of OPERA). §D explains the ϵ -type setting and the corresponding results. §E discusses the OPERA algorithm when being applied to special examples (linear mixture MDPs, low Witness rank MDPs, KNRs). §F details the delayed proofs of technical lemmas. §G details the proofs relevant to FE dimension. Finally in §H we provide an experimental result.

A Related Work

Tabular RL. Tabular RL considers MDPs with finite state space and action space. This setting has been extensively studied [Auer et al., 2008, Dann and Brunskill, 2015, Brafman and Tennenholtz, 2002, Agrawal and Jia, 2017, Azar et al., 2017, Zanette and Brunskill, 2019, Zhang et al., 2020] and the minimax-optimal regret bound is proved to be $\tilde{O}(\sqrt{H^2 S_j A_j T})$ [Jin et al., 2018, Domingues et al., 2021]. The minimax optimal bounds suggests that the tabular RL is information-theoretically hard for large S_j and A_j . Therefore, in order to deal with high-dimensional state-action space arose in many real-world applications, more advanced structural assumptions that enable function approximation are in demand.

Complexity Measures for Statistical Learning. In classic statistical learning, a variety of complexity measures have been proposed to upper bound the sample complexity required for achieving a certain accuracy, including VC Dimension [Vapnik, 1999], covering number [Pollard, 2012], Rademacher Complexity [Bartlett and Mendelson, 2002], sequential Rademacher complexity [Rakhlin et al., 2010] and Littlestone dimension [Littlestone, 1988]. However, for reinforcement learning, it is a major challenge to find such general complexity measures that can be used to analyze the sample complexity under a general framework.

RL with Linear Function Approximation. A line of work studied the MDPs that can be represented as a linear function of some given feature mapping. Under certain completeness conditions, the proposed algorithms can enjoy sample complexity/regret scaling with the dimension of the feature mapping rather than S_j and A_j . One such class of MDPs is linear MDPs [Jin et al., 2020, Wang et al., 2019, Neu and Pike-Burke, 2020], where the transition probability function and reward function are linear in some feature mapping over state-action pairs. Zanette et al. [2020a,b] studied MDPs under a weaker assumption called low inherent Bellman error, where the value functions are nearly linear w.r.t. the feature mapping. Another class of MDPs is linear mixture MDPs [Modi et al., 2020, Jia et al., 2020, Ayoub et al., 2020, Zhou et al., 2021b, Cai et al., 2020], where the transition probability kernel is a linear mixture of a number of basis kernels. The above paper assumed that feature vectors are known in the MDPs with linear approximation while Agarwal et al. [2020b] studied a harder setting where both the feature and parameters are unknown in the linear model.

RL with General Function Approximation. Beyond the linear setting, a recent line of research attempted to unify existing sample-efficient approaches with general function approximation. Osband and Van Roy [2014] proposed an structural condition named eluder dimension. Wang et al. [2020] further proposed an efficient algorithm LSVI-UCB for general linear function classes with small eluder dimension. Another line of works proposed low-rank structural conditions, including Bellman rank [Jiang et al., 2017, Dong et al., 2020] and Witness rank [Sun et al., 2019]. Yang et al. [2020] studied the MDPs with a structure where the action-value function can be represented by a kernel function or an over-parameterized neural network. Recently, Jin et al. [2021] proposed a complexity called Bellman eluder (BE) dimension. The RL problems with low BE dimension subsume the problems with low Bellman rank and low eluder dimension. Simultaneously Du et al. [2021] proposed Bilinear Classes, which can be applied to a variety of loss estimators beyond vanilla Bellman error, but with possibly worse sample complexity. Very recently, Foster et al. [2021] proposed Decision-Estimation Coefficient (DEC), which is a necessary and sufficient condition for sample-efficient interactive learning. To apply DEC to reinforcement learning, Foster et al. [2021] further proposed a RL class named Bellman Representability, which can be viewed as a generalization of the Bilinear Class.

	Bilinear Class	Low BE Dimension	DEC and Bellman Representability	ABC Class (with Low FE Dimension)
Linear MDPs [Yang and Wang, 2019, Jin et al., 2020]	$d^3 H^4 = 2$	$d^2 H^2 = 2$	$d^3 H^3 = 2$	$d^2 H^2 = 2$
Linear Mixture MDPs [Modi et al., 2020]	$d^3 H^4 = 2$	8	$d^3 H^3 = 2$	$d^2 H^2 = 2$
Bellman Rank [Jiang et al., 2017]	$d^2 H^5 j A_j = 2$	$d H^2 j A_j = 2$	$d^2 H^3 j A_j = 2$	$d H^2 j A_j = 2$
Eluder Dimension [Wang et al., 2020]	8	$\dim_E H^2 = 2$	$\dim_E^2 H^3 = 2$	$\dim_E H^2 = 2$
Witness Rank [Sun et al., 2019]	—	8	—	$W H^2 j A_j = 2$
Low Occupancy Complexity [Du et al., 2021]	$d^3 H^4 = 2$	$d^2 H^2 = 2$	$d^3 H^3 = 2$	$d^2 H^2 = 2$
Kernelized Nonlinear Regulator [Kakade et al., 2020]	—	8	—	$d^2 d_s H^4 = 2$
LinearQ =V [Du et al., 2021]	$d^3 H^4 = 2$	$d^2 H^2 = 2$	$d^3 H^3 = 2$	$d^2 H^2 = 2$

Table 1: Comparison of sample complexity for different MDP models under different RL frameworks. “—” indicates that the original work of framework does not provide an explicit sample complexity result for that model (although can be computed in principle). “0” indicates the model is not included in the framework for complexity analysis. For models with the linear structure in a d -dimensional space, we present the sample complexity in terms of d . For models with their own complexity measures, we use W to denote the witness rank, \dim_E the eluder dimension, d the dimension of H in KNR and d_s the dimension number of the state space of KNR. The dependency covering number is deliberately ignored for Bellman rank, eluder dimension, and the witness rank.

In Table 1, we compare our ABC framework with other structural RL frameworks in terms of the model coverage and sample complexity.

B Additional Examples

In this section, we compare our work with other results in the literature in terms of regret bounds/sample complexity. First of all, as we mentioned earlier in §3 when taking DEF as $\mathbb{E}_{h,f} \langle \mathbf{a}_h; \mathbf{f}_{h+1}; \mathbf{g}_h; \mathbf{v} \rangle = \mathbb{Q}_{h,g}(\mathbf{s}_h; \mathbf{a}_h) - \mathbf{r}_h - V_{h+1,f}(\mathbf{s}_{h+1})$ the ABC function reduces to the average Bellman error, and our ABC framework recovers the low Bellman eluder dimension framework for all cases compatible with such an estimation function. On several model-free structures, our regret bound is equivalent to that of the GOLF algorithm [Jin et al., 2021]. For example for linear MDPs, OPERA exhibits a $\mathcal{O}(dH^2 \bar{T})$ regret bound that matches the state-of-the-art result on linear function approximation provided in Zanette et al. [2020a]. For low eluder dimension models, the dependency on the eluder dimension in our regret analysis is $\mathcal{O}(\bar{d})$ while the dependency in Wang et al. [2020] is $\mathcal{O}(d)$. Also, for models with low Bellman rank, our sample complexity scales linearly with d as in Jin et al. [2021] while complexity in Jiang et al. [2017] scales quadratically.

For model-based RL settings with linear structure that are not within the low BE dimension framework such as the linear mixture MDPs, our OPERA algorithm obtains a $\mathcal{O}(dH^2 \bar{T})$ regret bound and $d_{FE}^2 H^2 = 2$ sample complexity result. In comparison, Jia et al. [2020], Modi et al. [2020] proposed an UCRL-VTR algorithm on linear mixture MDPs with a $\mathcal{O}(dH^2 \bar{T})$ regret bound, and Zhou et al. [2021a] improves this result by \bar{H} via a Bernstein-type bonus for exploration. The Bilinear Classes [Du et al., 2021] is a general framework that covers linear mixture MDPs as a special case. The sample complexity of the BiLin-UCB algorithm when constrained to linear mixture models is $\mathcal{O}(d^2 H^4 = 2)$, which is dH^2 worse than that of OPERA in this work. In terms of lower bound, when specialized to linear mixture MDPs, our result $\mathcal{O}(dH^2 \bar{T})$ matches the lower bound provided in Zhou et al. [2021a] up to a factor of $dH^{1=2}$.

In the rest of this section, we compare on six additional examples: the linear model [Du et al., 2021], the low occupancy complexity model [Du et al., 2021], kernel reactive POMDPs, FLAMBE/Feature Selection, Linear Quadratic Regulator, and partially Generalized Linear Bellman Complete. Moreover, we added a discussion on the state-action aggregation model.

B.1 Linear Q / V

The linear Q / V model was proposed in Du et al. [2021]. In addition to the linear structure of the optimal action-value function Q , we further assume linear structure of the optimal state-value function V . We formally define the linear Q / V model as follows:

Definition 15 (Linear Q / V, Definition 4.5 in Du et al. 2021.) A linear Q / V model satisfies for two Hilbert spaces $H_1; H_2$ and two given feature mappings $\phi; \psi : S \rightarrow H_1, (\phi^0; \psi^0) : S \rightarrow H_2$, there exist $w_h \in H_1; v_h \in H_2$ such that

$$Q_h(s; a) = \langle w_h; \phi(s; a) \rangle \quad \text{and} \quad V_h(s^0) = \langle v_h; \psi^0(s^0) \rangle$$

for any $h \in [H]$ and $(s; a; s^0) \in S \times A \times S$.

Suppose that H_1 and H_2 has dimension number d_1 and d_2 , separately, Du et al. [2021] shows that linear Q / V model belongs to the Bilinear Class with dimension $d_1 + d_2$ and BiLin-UCB algorithm achieves $\mathcal{O}(\frac{d^3 H^4}{2})$ sample complexity. On the other hand the sample complexity of OPERA is $\mathcal{O}(\frac{d^2 H^2}{2})$.

B.2 Low Occupancy Complexity

The low occupancy complexity model assumes linearity on the state-action distribution and has been proposed in Du et al. [2021]. We recap its definition formally as follows:

Definition 16 (Low Occupancy Complexity, Definition 4.7 in Du et al. 2021) A low occupancy complexity model is an MDP satisfying for some hypothesis class \mathcal{F} a Hilbert space H and feature mappings $\phi; \psi : S \times A \rightarrow H; \phi^0; \psi^0 : S \rightarrow H$ that there exists a function on hypothesis classes $f : \mathcal{F} \rightarrow H$ such that

$$Q_h^f(s_h; a_h) = \langle f; \phi(s_h; a_h) \rangle; \quad \phi^0; \psi^0 : S \rightarrow H$$

Du et al. [2021] proved that the low occupancy complexity model belongs to the Bilinear Classes and has a sample complexity $\mathcal{O}(H^4)$ under the BiLin-UCB algorithm. In the meantime, the low occupancy complexity model admits an improved sample complexity $\mathcal{O}(H^2)$ under the OPERA algorithm.

B.3 Kernel Reactive POMDPs

The Reactive POMDP [Krishnamurthy et al., 2016] is a partially observable MDP (POMDP) model that can be described by the tuple $(S; A; O; T; O; r; H)$, where S and A are the state and action spaces respectively, O is the observation space, T is the transition matrix that maps each $(s; a) \in S \times A$ to a probability measure \mathcal{P} and determines the dynamics of the next state as $T(\cdot | s; a)$, O is the emission measure that determines the observation $O(\cdot | s)$ given current state s . The reactivity of a POMDP refers to the property that the optimal value function depends only on the current observation and action. In other words, for each $h \in [H]$ there exists $f_h : O \times A \rightarrow [0; 1]$ such that for any given trajectory $\tau_h = [o_1; a_1; \dots; o_n]$ and a_h , we have

$$Q_h(\tau_h; a_h) = f_h(o_n; a_h)$$

Given the definition of a reactive POMDP, we define the kernel reactive POMDP [Jin et al., 2021] as follows:

Definition 17 (Kernel Reactive POMDP) A kernel reactive POMDP is a reactive POMDP that satisfies for each $h \in [H]$ and a given separable Hilbert space H there exist feature mappings $\phi; \psi : S \times A \rightarrow H$ and $\phi^0; \psi^0 : S \rightarrow H$ such that the transition matrix $T_h(s^0; j; s; a) = \langle \phi_h(s; a); \phi_h(s^0) \rangle_{H_1}$ and ψ is bounded in the sense that for any $(s; a) : S \times A \rightarrow [0; 1]$, $\sum_{s^0 \in S} V(s^0) \langle \psi^0(s^0); \psi^0(s^0) \rangle_{H_2} < 1$.

In Jin et al. [2021], the authors showed that the kernel reactive POMDP with vanilla estimation function $\hat{V}_h(o_h; f_{h+1}; g_h; v) = Q_{h;g}(s_h; a_h) + r_h + V_{h+1;f}(s_{h+1})$ has V -type BE dimension bounded by the effective dimension. According to Proposition 34, the kernel reactive POMDP model also has low FE dimension bounded by the effective dimension.

B.4 FLAMBE/Feature Selection

For FLAMBE/feature selection model firstly introduced in Agarwal et al. [2020b], similarity is shared with the linear MDP setting but the main difference lies in that the feature mappings are unknown. We formally define the feature selection model as follows:

Definition 18 (Feature Selection) A low rank feature selection model is an MDP that satisfies for any $h \in [H]$ and a given Hilbert space \mathcal{H} , there exist unknown feature mappings $\phi_h: \mathcal{S} \rightarrow \mathcal{H}$ and $\psi_h: \mathcal{S} \rightarrow \mathcal{A}$ such that the transition probability satisfies:

$$P_h(s^0 | s; a) = \langle \phi_h(s^0), \psi_h(s; a) \rangle; \quad \forall (s; a; s^0) \in \mathcal{S} \times \mathcal{A} \times \mathcal{S}$$

We consider the feature selection model with $\text{DEF}(\phi_h; \psi_h; \mathbf{g}_h; \mathbf{v}) := Q_{h;g}(s_h; a_h) + r_h + V_{h+1;f}(s_{h+1})$. In Du et al. [2021] they have proved in Lemma A.1 that

$$E_{s_h; g; a_h} [Q_{h;f}(s_h; a_h) + r_h + V_{h+1;f}(s_{h+1})] = \langle W_h(f); X_h(g) \rangle; \quad (\text{B.1})$$

where

$$\begin{aligned} W_h(f) &:= \int_{s \in \mathcal{S}} \phi_h(s) V_{h;f}(s) + r(s; f(s)) E_{s^0 \sim P_h(\cdot | s; f(s))} [V_{h+1;f}(s^0)] ds; \\ X_h(g) &:= E_{s_h \sim \pi_{h-1}; a_{h-1} \sim g} [\psi_h(s_h; a_{h-1})]; \end{aligned}$$

We note that Eq(B.1) ensures condition (i) and (ii) in Definition 5 at the same time and the ABC of the feature selection setting has a bilinear structure that enables us to apply Proposition 35 to conclude low FE dimension.

B.5 Linear Quadratic Regulator

In a linear quadratic regulator (LQR) model [Bradtke, 1992, Anderson and Moore, 2007, Dean et al., 2020], we consider the d -dimensional state space $\mathcal{S} \subseteq \mathbb{R}^d$ and K -dimensional action space $\mathcal{A} \subseteq \mathbb{R}^K$. The transition dynamics of an LQR model can be written in matrix form so that the induced value function is quadratic [Jiang et al., 2017]. We formally define the LQR model as follows:

Definition 19 (Linear Quadratic Regulator) A linear quadratic regulator model is an MDP such that there exist unknown matrices $A \in \mathbb{R}^{d \times d}$, $B \in \mathbb{R}^{d \times K}$ and $Q \in \mathbb{R}^{d \times d}$ satisfying for $h \in [H]$ and zero-centered random variables s_h with $E[s_h] = 0$ and $E[s_h^2] = \sigma^2$ that

$$\begin{aligned} s_{h+1} &= A s_h + B a_h + \eta_h; \\ r_h &= s_h^T Q s_h + a_h^T a_h + \nu_h. \end{aligned}$$

The LQR model has been analyzed in Du et al. [2021] and proved to belong to the Bilinear Classes. Du et al. [2021] used the hypothesis class defined as

$$F_h = \{ (C_h; \eta_h; O_h) : C_h \in \mathbb{R}^{K \times d}; \eta_h \in \mathbb{R}^{d \times d}; O_h \in \mathbb{R}^{d \times d} \}_{h \in [H]}$$

For each hypothesis in the class F , the corresponding policy and value function are

$$f(s_h) = C_{h;f} s_h; \quad V_{h;f}(s_h) = s_h^T \eta_{h;f} s_h + O_{h;f};$$

Under the above setting, we use the DEF for LQR $\text{DEF}(\phi_h; \psi_h; \mathbf{g}_h; \mathbf{v}) := Q_{h;g}(s_h; a_h) + r_h + V_{h+1;f}(s_{h+1})$ and Lemma A.4 in Du et al. [2021] showed that

$$E_{s_h; a_h} [Q_{h;f}(s_h; a_h) + r_h + V_{h+1;f}(s_{h+1})] = \langle W_h(f); X_h(g) \rangle; \quad (\text{B.2})$$

where

$$\begin{aligned} W_h(f) &= \text{vec}(\eta_{h;f} + Q + C_{h;f}^T C_{h;f} + (A + B C_{h;f})^T \eta_{h+1;f} (A + B C_{h;f})); \\ &\quad O_{h;f} + O_{h+1;f} + \text{trace}(\eta_{h+1;f}); \\ X_h(f) &= \text{vec}(E_{s_h} [s_h s_h^T]); \end{aligned}$$

We note that Eq(B.2) ensures condition (i) and (ii) in Definition 5 simultaneously and the ABC of the LQR model setting admits a bilinear structure that enables us to apply Proposition 35 and conclude low FE dimension.

B.6 Generalized Linear Bellman Complete

Next we introduce the generalized linear Bellman complete model, showing that our ABC class with low FE dimension captures this model even without the monotone operator used in Du et al. [2021].

Definition 20 (Generalized Linear Bellman Complete) A generalized linear Bellman complete model consists of an inverse link function $\tau : \mathbb{R} \rightarrow \mathbb{R}^+$ and a hypothesis class $\mathcal{F}_h := \{f_h(\cdot; \mathbf{a}) : h \in \mathcal{H}; k_h \in \mathbb{R}^{h \times [H]}\}$ such that for any $f \in \mathcal{F}$ and $h \in \mathcal{H}$ the Bellman completeness condition holds:

$$r(\mathbf{s}; \mathbf{a}) + E_{\mathbf{s}^0 \in \mathcal{P}_h} \max_{\mathbf{a}^0 \in \mathcal{A}} (\tau_{h+1, f}(\mathbf{s}^0, \mathbf{a}^0)) \in \mathcal{H}_h$$

By the choice of the hypothesis class, we know that there exists a mapping $\tau_h : \mathcal{H} \rightarrow \mathcal{H}$ such that

$$\tau_h(\tau_{h+1, f})^\triangleright(\mathbf{s}; \mathbf{a}) = r(\mathbf{s}; \mathbf{a}) + E_{\mathbf{s}^0 \in \mathcal{P}_h} \max_{\mathbf{a}^0 \in \mathcal{A}} (\tau_{h+1, f}(\mathbf{s}^0, \mathbf{a}^0)); \quad (\text{B.3})$$

We note that in Du et al. [2021] they choose a discrepancy function dependent on a discriminator function v . In this work, we choose a different estimation function that allows much simpler calculation and sharper sample complexity result. We let

$$\tau_h(\mathbf{o}_h; \mathbf{f}_{h+1}; \mathbf{g}_h; v) := (\tau_{h, g}(\mathbf{s}_h; \mathbf{a}_h)) \quad r_h \max_{\mathbf{a}^0} \tau_{h+1, f}(\mathbf{s}_{h+1}; \mathbf{a}^0)$$

By Eq. (B.3), it is easy to check that the above DEF satisfies the decomposable condition. Assuming a $\tau_h(\mathbf{x}) = b$, Lemma 6.2 in Du et al. [2021] has already shown the Bellman dominance property that

$$\begin{aligned} & E_{\mathbf{s}_h; \mathbf{a}_h} \frac{1}{q} [Q_{h, f}(\mathbf{s}_h; \mathbf{a}_h) - r_h - V_{h+1, f}(\mathbf{s}_{h+1})] \\ & \leq \frac{1}{b} \text{vec}((\tau_{h, f} - \tau_h)(\tau_{h+1, f})^\triangleright); \text{vec} E_{\mathbf{s}_h; \mathbf{a}_h} (\mathbf{s}_h; \mathbf{a}_h)^\triangleright \\ & = \frac{1}{b} \overline{\text{HW}_h(f); X_h(f)} \end{aligned}$$

Next, we illustrate that the Dominating Average EF condition holds in our framework. We have

$$\begin{aligned} & E_{\mathbf{s}_h; \mathbf{g}; \mathbf{a}_h} \text{op} E_{\mathbf{s}_{h+1}} [(\tau_{h, g}(\mathbf{o}_h; \mathbf{f}_{h+1}; \mathbf{f}_h; v))^\triangleright | \mathbf{s}_h; \mathbf{a}_h] \frac{1}{2} \\ & = E_{\mathbf{s}_h; \mathbf{a}_h} \text{op} (\tau_{h, f}(\mathbf{s}_h; \mathbf{a}_h))^\triangleright (\tau_h(\tau_{h+1, f})^\triangleright(\mathbf{s}; \mathbf{a})) \frac{1}{2} \\ & \leq E_{\mathbf{s}_h; \mathbf{a}_h} \text{op} (\tau_{h, f}(\mathbf{s}_h; \mathbf{a}_h))^\triangleright (\mathbf{s}_h; \mathbf{a}_h) \frac{1}{2} \leq \overline{\text{HW}_h(f); X_h(g)} \end{aligned}$$

where

$$\begin{aligned} W_h(f) & := \text{vec}((\tau_{h, f} - \tau_h)(\tau_{h+1, f})^\triangleright); \\ X_h(f) & := \text{vec} E_{\mathbf{s}_h; \mathbf{a}_h} (\mathbf{s}_h; \mathbf{a}_h)^\triangleright; \end{aligned}$$

Analogous to the KNR case and the proof of Lemma 30, the aforementioned model with ABC function $\overline{\text{HW}_h(f); X_h(f)}$ has low FE dimension.

B.7 Q state-action aggregation

Finally, we consider the Q state-action aggregation model [Dong et al., 2019], which cannot be covered by the bilinear classes [Du et al., 2021]. We illustrate that our ABC framework covers this model with a nonlinear coupling function.

Definition 21 (Q state-action aggregation) We call an MDP a Q state-action model if there exists a $(\mathbf{s}; \mathbf{a}) : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{B}$ such that for any state-action pair $(\mathbf{s}; \mathbf{a}), (\mathbf{s}^0; \mathbf{a}^0) \in \mathcal{S} \times \mathcal{A}$, if $(\mathbf{s}; \mathbf{a}) = (\mathbf{s}^0; \mathbf{a}^0)$, then $Q(\mathbf{s}; \mathbf{a}) = Q(\mathbf{s}^0; \mathbf{a}^0)$. The dimension of the Q state-action aggregation model is defined as the cardinality of \mathcal{B} , i.e., $d = |\mathcal{B}|$.

For each element $b \in \mathcal{B}$, we note that for all $(\mathbf{s}; \mathbf{a}) \in \mathcal{S} \times \mathcal{A}$ satisfying $(\mathbf{s}; \mathbf{a}) = b$, they share the same value $Q(\mathbf{s}; \mathbf{a})$. For notational simplicity, we use $\mathcal{Q}(b)$ to denote this common value for all $(\mathbf{s}; \mathbf{a}) = b$. Moreover, we let \mathbf{w}_h be a d -dimensional vector of all aggregated values with $(\mathbf{w}_h)_b = Q(b)$ and $Q(\mathbf{s}; \mathbf{a})$ can be expressed as a linear function on the aggregated values:

$$Q_h(\mathbf{s}; \mathbf{a}) = \mathbf{w}_h; (\mathbf{s}; \mathbf{a})$$

where $(s; a) : S \times A \rightarrow \{0, 1\}^d$ is a one-hot vector satisfying

$$(s; a)_b = 1 \text{ when } (s; a) = b \text{ and } (s; a)_b = 0 \text{ otherwise}$$

For this case, we use the hypothesis class defined as

$$F_h = \{w_h \in \mathbb{R}^d\}_{h \in [H]}$$

where we take

$$Q_{h;g}(s; a) = \sum_{i \in A} w_{h;g}(s; a)_i v_i; \quad V_{h;g}(s^0) = \max_{a \in A} \sum_{i \in A} w_{h;g}(s^0; a)_i v_i$$

We define the DEF as the Bellman residual and conclude that

$$\begin{aligned} \mathcal{R}_h(\theta_h; g_{h+1}; g_h; v) &= Q_{h;g}(s_h; a_h) - r_h - V_{h;g}(s_{h+1}) \\ &= Q_{h;g}(s_h; a_h) - V_{h;g}(s_{h+1}) - [Q_h(s_h; a_h) - V_h(s_{h+1})] \\ &= \sum_{i \in A} w_{h;g}(s_h; a_h)_i v_i - \max_{a \in A} \sum_{i \in A} w_{h;g}(s_{h+1}; a^0)_i v_i - \max_{a \in A} \sum_{i \in A} w_h(s_{h+1}; a^0)_i v_i \end{aligned} \quad (\text{B.4})$$

Taking expectation over the distribution θ_h given by $s_h; a_h \sim f$ and $s_{h+1} \sim P_h$, we have

$$\begin{aligned} &E_{s_h; a_h \sim f} [\mathcal{R}_h(\theta_h; g_{h+1}; g_h; v)] \\ &= \sum_{i \in A} w_{h;g} v_i - E_{s_h; a_h \sim f} \left(\sum_{i \in A} w_h(s_h; a_h)_i v_i \right) \\ &= E_{s_h; a_h \sim f; s_{h+1} \sim P_h} \left(\max_{a \in A} \sum_{i \in A} w_{h;g}(s_{h+1}; a^0)_i v_i - \max_{a \in A} \sum_{i \in A} w_h(s_{h+1}; a^0)_i v_i \right) \end{aligned} \quad (\text{B.5})$$

If we define $G(g; f)$ by the right hand side of (B.5), it is obvious that $G_{h;f}(\cdot; \cdot)$ with \mathcal{R}_h defined in (B.4) serves as the coupling function and the estimation function of an ABC class, respectively. In the mean time, Equation (B.5) cannot be expressed as a inner product of $\sum_{i \in A} w_{h;g} v_i$; $X_h(g)$ and thus cannot be covered by Du et al. [2021]. Nevertheless, in our ABC framework, it is unclear if the FE dimension of F with respect to $G_{h;f}(\cdot; \cdot)$ can be bounded in a nontrivial way (i.e. $|S| \leq |A|$).

C Proof of Main Results

In this section, we provide proofs of our main result Theorem 12 and a sample complexity corollary of the OPERA algorithm. Originated from proof techniques widely used in confidence bound based RL algorithms Russo and Van Roy [2013] our proof steps generalizes that of the GOLF algorithm Jin et al. [2021] but admits general DEF and ABCs. We prove our main result as follows:

C.1 Proof of Theorem 12

Proof of Theorem 12 We recall that the objective of an RL problem is to find an optimal policy satisfying $V_1(s_1) = V_1^{\pi^*}(s_1)$. Moreover, the regret of an RL problem is defined as $\sum_{t=1}^T V_1(s_1) - V_1^{\pi^t}(s_1)$, where π^t is the output policy of an algorithm at time t .

Step 1: Feasibility of f . First of all, we show that the optimal hypothesis lies within the confidence region defined by Eq. (4.1) with high probability:

Lemma 22 (Feasibility of f). In Algorithm 1, given $\epsilon > 0$ and $\delta > 0$ we choose $\eta = c(\log(T) + \log(1/\delta)) + T$ for some large enough constant c . Then with probability at least $1 - \delta$, f satisfies for any $t \in [T]$:

$$\max_{v \in \mathcal{V}} \left(\sum_{i=1}^t \mathcal{R}_{h;f}(s_i^0; f_{h+1}; f_h; v) \right)^2 \leq \inf_{g_h \in \mathcal{G}_h} \left(\sum_{i=1}^t \mathcal{R}_{h;f}(s_i^0; f_{h+1}; g_h; v) \right)^2 + O(\epsilon)$$

Lemma 22 shows that at each round of updates the optimal hypothesis lies in the confidence region depicted by Eq. (4.1) with radius $O(\epsilon)$. We delay the proof of Lemma 22 to §F.2. Lemma 22 together with the optimization procedure Line 3 of Algorithm 1 implies an upper bound on $V_1^{\pi^t}(s_1)$ with probability at least $1 - \delta$ as follows:

$$V_1(s_1) - V_1^{\pi^t}(s_1) \leq V_{1;f^t}(s_1) - V_1^{\pi^t}(s_1) \quad (\text{C.1})$$

Step 2: Policy Loss Decomposition. The second step is to upper bound the regret by the summation of Bellman errors. We apply the policy loss decomposition lemma in Jiang et al. [2017].

Lemma 23 (Lemma 1 in Jiang et al. 2017). For $H \geq 2$,

$$V_{1,f^t}(s_1) - V_1^t(s_1) = \sum_{h=1}^{H-1} E_{S_h; a_h^t} [Q_{h,f^t}(S_h; a_h) - r_h - V_{h+1,f^t}(S_{h+1})]:$$

Combining Lemma 23 with Eq. (C.1) we have the following:

$$V_1(s_1) - V_1^t(s_1) - V_{1,f^t}(s_1) - V_1^t(s_1) = \sum_{h=1}^{H-1} E_{S_h; a_h^t} [Q_{h,f^t}(S_h; a_h) - r_h - V_{h+1,f^t}(S_{h+1})]: \quad (C.2)$$

Step 3: Small ABC Value in the Confidence Region. The third step is devoted to controlling the cumulative square of Admissible Bellman Characterization function. Recalling that the ABC function is upper bounded by the average DEF, where each feasible DEF stays in the confidence region that satisfies Eq. (4.1), we arrive at the following Lemma 24:

Lemma 24. In Algorithm 1, given $\epsilon > 0$ and $\delta > 0$ we choose $\eta = c(\log(\frac{1}{\delta}) + T)$ for some large enough constant c . Then with probability at least $1 - \delta$, for all $(t; h) \in [T] \times [H]$, we have

$$\sum_{i=1}^{H-1} G_{h,f^t}(f^i; f^i)^2 \leq O(\epsilon): \quad (C.3)$$

The proof of Lemma 24 makes use of Freedman's inequality (the precise version as in Agarwal et al. [2014]) and we delay the proof to §F.1.

Step 4: Bounding the Cumulative Bellman Error by Functional Eluder Dimension. In the fourth step, we aim to translate the upper bound of the cumulative squared ABC function in Eq. (C.3) to an upper bound of the cumulative ABC function. The following Lemma 25 is adapted from Lemma 41 in Jin et al. [2021] and Lemma 2 in Russo and Van Roy [2013]. Lemma 25 controls the sum of ABC functions by properties of the functional eluder dimension.

Lemma 25. For a hypothesis class \mathcal{F} and a given coupling function $G(\cdot; \cdot): \mathcal{F} \times \mathcal{F} \rightarrow \mathbb{R}$ with bounded image space $\mathcal{G}(\cdot; \cdot) \subset \mathbb{C}$. For any pair of sequences $f_t, g_t \in \mathcal{F}$ satisfying for all $t \in [T]$, $\sum_{i=1}^{t-1} (G(f_i; g_i))^2 \leq \epsilon$, the following inequality holds for all $t \in [T]$ and $\epsilon > 0$:

$$\sum_{i=1}^t |G(f_i; g_i)| \leq O\left(\sqrt{\frac{\epsilon}{\dim_{FE}(\mathcal{F}; \mathcal{G}; \epsilon)} t} + C \min\{t, \dim_{FE}(\mathcal{F}; \mathcal{G}; \epsilon) g + t\}\right):$$

The proof of Lemma 25 is in §F.3.

Step 5: Combining Everything. In the final step, we combine the regret bound decomposition argument, the cumulative ABC bound, and the Bellman dominance property together to derive our final regret guarantee.

For any $h \in [H]$, we take $G(\cdot; \cdot) = G_{h,f^t}(\cdot; \cdot)$, $g_i = f^i; f_t = f^t$ and $\epsilon = \frac{\epsilon}{T}$ in Lemma 25. By Eq. (C.3) in Lemma 24, we have for any $h \in [H]$ and $t \in [T]$,

$$\begin{aligned} \sum_{i=1}^t |G_{h,f^t}(f^i; f^i)| &\leq O\left(\sqrt{\frac{\epsilon}{\dim_{FE}(\mathcal{F}; G_{h,f^t}; \frac{\epsilon}{T})} t} + C \min\{t, \dim_{FE}(\mathcal{F}; G_{h,f^t}; \frac{\epsilon}{T}) g + t\}\right) \\ &\leq O\left(\sqrt{\frac{\epsilon}{\dim_{FE}(\mathcal{F}; G_{h,f^t}; \frac{\epsilon}{T})} t}\right): \end{aligned}$$

We recall our choice of $\epsilon = c(\log(\frac{1}{\delta}) + T)$. Taking $\epsilon = \frac{\epsilon}{T}$, we have

$$\begin{aligned} \sum_{i=1}^t |G_{h,f^t}(f^i; f^i)| &\leq O\left(\sqrt{\frac{r}{\dim_{FE}(\mathcal{F}; G_{h,f^t}; \frac{\epsilon}{T})} \log(\frac{1}{\delta})} t\right) \\ &\leq O\left(\sqrt{\frac{r}{\dim_{FE}(\mathcal{F}; G; \frac{\epsilon}{T})} \log(\frac{1}{\delta})} t\right): \end{aligned}$$

Combining this with property (ii) in Definition 5 and decomposition (C.2), we conclude our main result that with probability at least $1 - \frac{1}{T}$,

$$\frac{1}{T} \sum_{t=1}^T V_1(s_1) - V_1^*(s_1) \leq \frac{1}{T} \sum_{t=1}^T \sum_{h=1}^H \mathbb{E}_{\mathcal{S}} [G_{h,f}(f^t; f^t)] + \frac{H}{T} \frac{\log(\text{THN}_L(\frac{1}{T}))}{\dim_{\text{FE}}(F; G; \frac{1}{T})} + \frac{H!}{T}.$$

This completes the whole proof of Theorem 12. \square

C.2 Sample Complexity of OPERA

Corollary 26 (Sample Complexity of OPERA) For an MDP with hypothesis classes \mathcal{F}, \mathcal{G} that satisfies Assumption 1 and a Decomposable Estimation Function satisfying Assumption 7. If there exists an Admissible Bellman Characterization with low functional eluder dimension. For any $\epsilon \in (0, 1]$, we choose $\frac{1}{T} = c \log(\text{THN}_L(\frac{1}{T})) + T \frac{1}{\dim_{\text{FE}}(F; G; \frac{1}{T})}$ for some large enough constant c . For the on-policy case when $\mu_{\text{op}} = \mu_{\text{est}} = \mu^*$, with probability at least $1 - \frac{1}{T}$, Algorithm 1 outputs a ϵ -optimal policy μ_{out} within T trajectories where

$$T = \frac{\dim_{\text{FE}}(F; G; \frac{1}{T}) \log(\text{THN}_L(\frac{1}{T}))}{\frac{1}{2}} = H^2.$$

Proof of Corollary 26. By the policy loss decomposition (C.2), (C.3) in Lemma 24 and Lemma 25, we have that

$$\frac{1}{T} \sum_{t=1}^T V_1(s_1) - V_1^*(s_1) \leq \frac{1}{T} \sum_{t=1}^T \sum_{h=1}^H \mathbb{E}_{\mathcal{S}} [G_{h,f}(f^t; f^t)] + \frac{H}{T} \frac{\log(\text{THN}_L(\frac{1}{T}))}{\dim_{\text{FE}}(F; G; \frac{1}{T})} + \frac{H!}{T}. \quad (\text{C.4})$$

Taking $\frac{1}{T} = \frac{1}{H}$ and $\frac{1}{T} = \frac{1}{\dim_{\text{FE}}(F; G; \frac{1}{H})}$, the above Eq. (C.4) becomes

$$\frac{1}{T} \sum_{t=1}^T V_1(s_1) - V_1^*(s_1) \leq \frac{H}{T} \frac{\log(\text{THN}_L(\frac{1}{T}))}{\dim_{\text{FE}}(F; G; \frac{1}{H})} + \frac{H!}{T}.$$

Taking

$$T = \frac{\dim_{\text{FE}}(F; G; \frac{1}{H}) \log(\text{THN}_L(\frac{1}{T}))}{\frac{1}{2}} H^2$$

yields the desired result. \square

D Q-type and V-type Sample Complexity Analysis

In Definition 5, we note that there are two ways to calculate the ABC of an MDP model depending on the different choices of the operating policy μ_{op} . Specifically, if $\mu_{\text{op}} = \mu_g$, we call it the Q-type ABC. Otherwise, if $\mu_{\text{op}} = \mu_f$, we call it the V-type ABC. For example, when taking

$$G_{h,f}(f; g) = \mathbb{E}_{\mathcal{S}_h} [Q_{h,f}(s_h; a_h) - r(s_h; a_h) - V_{h+1;f}(s_{h+1})]$$

the FE dimension of $G_{h,f}(f; g)$ recovers the Q-type BE dimension (Definition 8 in Jin et al. [2021]). When taking

$$G_{h,f}(f; g) = \mathbb{E}_{\mathcal{S}_h} [Q_{h,f}(s_h; a_h) - r(s_h; a_h) - V_{h+1;f}(s_{h+1})]$$

the FE dimension of $G_{h,f}(f; g)$ recovers the V-type BE dimension (Definition 20 in Jin et al. [2021]). The algorithm for solving Q-type or V-type models slightly differs in the executing policy μ_{est} . We use $\mu_{\text{est}} = \mu^*$ for Q-type models in Algorithm 1, while $\mu_{\text{est}} = U(A)$ is the uniform distribution on action set for V-type models.

The Q-type characterization and the V-type characterization have respective applicable zones. For example, the reactive POMDP model belongs to ABC with low FE dimension with respect to type ABC while inducing large FE dimension with respect to type ABC. On the contrary, the low inherent bellman error problem in Zanette et al. [2020a] is more suitable for V-type characterization rather than Q-type characterization. For general RL models, we often prefer type ABC because the sample complexity of V-type algorithms scales with the dimension of the action space $|A_j|$. Due to the uniform executing policy, we will only be able to derive regret bound for Q-type characterizations, as is explained in Jin et al. [2021].

In §4 and §C, we have illustrated regret bound and sample complexity results for V-type cases where we let $\epsilon_{op} = \epsilon_{est} = \epsilon$ through Algorithm 1. In the following Corollary 27, we prove sample complexity result for V-type ABC models.

Corollary 27. For an MDPM with hypothesis class \mathcal{F}, \mathcal{G} that satisfies Assumption 1 and a Decomposable Estimation Function satisfying Assumption 7. If there exists an Admissible Bellman Characterization \mathcal{G} with low functional eluder dimension. For any $\epsilon \in (0, 1]$, if we choose $\epsilon = O(\log(\frac{1}{\epsilon}) \sqrt{H} + T^{-1})$. For V-type models when $\epsilon_{op} = \epsilon_{est} = \epsilon$, with probability at least $1 - \epsilon$ Algorithm 1 outputs a-optimal policy π_{out} within $T = \frac{|A_j| \dim_{FE}(\mathcal{F}; \mathcal{G}; \frac{1}{H}) \log(\frac{1}{\epsilon}) \sqrt{H}}{\epsilon^2}$ trajectories where $\epsilon = \frac{\epsilon^2}{\dim_{FE}(\mathcal{F}; \mathcal{G}; \frac{1}{H}) H^2}$.

Proof of Corollary 27. The proof of Corollary 27 basically follows the proof of Theorem 12 and Corollary 26. We again have feasibility of π and policy loss decomposition. However, due to different sampling policy, the proof of Lemma 24 differs at Eq. (F.5). Instead, we have

$$\begin{aligned} & \sum_{i=1}^X \max_{V \in \mathcal{V}} E_{S_h} \sum_{i; a_h} \epsilon_{S_{h+1}} X_i(h; f^t; v) j_{S_h; a_h} \\ &= \sum_{i=1}^X \max_{V \in \mathcal{V}} E_{S_h} \sum_{i; a_h} U(A) \frac{1(a_h^i = f^i(s_h^i))}{|A_j|} E_{S_{h+1}} X_i(h; f^t; v) j_{S_h; a_h} \\ &= \sum_{i=1}^X \max_{V \in \mathcal{V}} E_{S_h} \sum_{i; a_h} U(A) \frac{1(a_h^i = f^i(s_h^i))}{|A_j|} E_{S_{h+1}} \sum_{h; f^i} (Q_h; f_{h+1}^t; f_h^t; v) j_{S_h; a_h} \\ & O(|A_j| \epsilon + R\epsilon + R^2 \epsilon): \end{aligned} \tag{D.1}$$

Thus, Eq. (C.3) in Lemma 24 becomes

$$\sum_{i=1}^X G_{h;f} (f^t; f^i)^2 = O(|A_j| \epsilon):$$

The rest of the proof follow the proof of Corollary 26 with an additional $|A_j|$ factor. By the policy loss decomposition (C.2) and Lemma 25, we have that

$$\begin{aligned} & \frac{1}{T} \sum_{t=1}^T V_1(s_1) - V_1^t(s_1) = \frac{1}{T} \sum_{t=1}^T \sum_{s_{h=1}}^X G_{h;f} (f^t; f^t) \\ & O\left(\frac{H}{|A_j| \dim_{FE}(\mathcal{F}; \mathcal{G}; \frac{1}{H})} \frac{\log(\frac{1}{\epsilon}) \sqrt{H}}{\epsilon} + \frac{H}{\epsilon}\right); \end{aligned} \tag{D.2}$$

Taking $\epsilon = \frac{\epsilon^2}{H}$ and $\epsilon = \frac{\epsilon^2}{\dim_{FE}(\mathcal{F}; \mathcal{G}; \frac{1}{H}) H^2}$, the above Eq. (D.2) becomes

$$\frac{1}{T} \sum_{t=1}^T V_1(s_1) - V_1^t(s_1) = O\left(\frac{H}{|A_j| \dim_{FE}(\mathcal{F}; \mathcal{G}; \frac{1}{H}) \log(\frac{1}{\epsilon}) \sqrt{H}} + \frac{H}{\epsilon}\right):$$

Taking

$$T = \frac{|A_j| \dim_{FE}(\mathcal{F}; \mathcal{G}; \frac{1}{H}) \log(\frac{1}{\epsilon}) \sqrt{H}}{\epsilon^2}$$

yields the desired result. \square

E Proof for Specific Examples

In this section, we consider three specific examples: linear mixture MDPs, low Witness rank MDPs, and KNRs. We explain how our framework exhibits superior properties than other general frameworks on these three instances of MDPs. For reader's convenience, we summarize the conditions introduced in Items (i), (ii) in Definition 6 and also Items (i), (ii) in Definition 5, that are essential for any RL models to fit in our framework:

- Decomposability:

$$\mathbb{E}_{s_{h+1}} [\mathbb{E}_{h,f} (Q_{h,f}(O_h; f_{h+1}; g_h; v) | s_h; a_h)] = \mathbb{E}_{h,f} (Q_{h,f}(O_h; f_{h+1}; T(f)_h; v))$$

- Global Discriminator Optimality:

$$\mathbb{E}_{s_{h+1}} [\mathbb{E}_{h,f} (Q_{h,f}(O_h; f_{h+1}; f_h; v_h(f)) | s_h; a_h)] \geq \mathbb{E}_{s_{h+1}} [\mathbb{E}_{h,f} (Q_{h,f}(O_h; f_{h+1}; f_h; v) | s_h; a_h)]$$

- Dominating Average EF:

$$\max_{V \geq 0} \mathbb{E}_{s_h} \mathbb{E}_{g; a_h} \mathbb{E}_{op} [\mathbb{E}_{h,g} (Q_{h,g}(O_h; f_{h+1}; f_h; v) | s_h; a_h)] \geq (G_{h,f}(f; g))^2$$

- Bellman Dominance:

$$\mathbb{E}_{s_h; a_h} [Q_{h,f}(s_h; a_h) - r(s_h; a_h) - V_{h+1;f}(s_{h+1})] \leq G_{h,f}(f; f)$$

E.1 Linear Mixture MDPs

For known transition and reward feature mappings $(s; a; s^0) : S \times A \times S \rightarrow \mathbb{H}$, $(s; a) : S \times A \rightarrow \mathbb{H}$ taking values in a Hilbert space \mathbb{H} and an unknown $\mathbb{H} \rightarrow \mathbb{H}$, a linear mixture MDP assumes that for any $(s; a; s^0) \in S \times A \times S$ and $h \in [H]$, the transition probability $P_h(s^0 | s; a)$ and the reward function $r(s; a)$ are linearly parameterized as

$$P_h(s^0 | s; a) = \langle s; a; s^0 \rangle_{\mathbb{H}}; \quad r(s; a) = \langle s; a \rangle_{\mathbb{H}}$$

In this case, we choose $\mathbb{F}_h = G_h = f_h \in \mathbb{H}_g$. Thus, the hypothesis classes \mathbb{F}_h and \mathbb{G}_h consist of the set of parameters $\{ \langle \cdot; \cdot \rangle_{\mathbb{H}} \}_{\mathbb{H} \times \mathbb{H}}$. Moreover, for each hypothesis class $\mathbb{F}_h = \{ \langle \cdot; \cdot \rangle_{\mathbb{H}} \}_{\mathbb{H} \times \mathbb{H}}$, the value function with respect to s_{h+1} for any $h \in [H]$ that

$$Q_{h,f}(s; a) = \langle s; a \rangle_{\mathbb{H}} + V_{h+1;f}(s; a)$$

where $V_{h+1;f}(s; a) := \mathbb{E}_{s^0 \in S} \langle s; a; s^0 \rangle_{\mathbb{H}} V_{h+1;f}(s^0)$. It is natural to define the DEF by

$$\mathbb{E}_{h,f} (Q_{h,f}(O_h; f_{h+1}; g_h; v)) = \langle s_h; a_h \rangle_{\mathbb{H}} + V_{h+1;f}(s_h; a_h) - r_h - V_{h+1;f}(s_{h+1})$$

If we use \mathbf{t}_h^{t-1} to denote the matrix $(\langle s_h^1; a_h^1 \rangle_{\mathbb{H}}; \dots; \langle s_h^{t-1}; a_h^{t-1} \rangle_{\mathbb{H}})$ and \mathbf{y}_h^{t-1} to denote the vector $(V_{h+1;f}(s_{h+1}^1); \dots; V_{h+1;f}(s_{h+1}^{t-1}))$, Eq. (4.1) in Algorithm 1 under linear mixture setting can be written in a matrix form as:

$$\mathbb{E}_{h,f} \mathbf{t}_h^{t-1} \mathbf{y}_h^{t-1} \geq \inf_{\mathbf{t}_h^{t-1}} \mathbf{t}_h^{t-1} \mathbf{y}_h^{t-1} \quad (E.1)$$

Taking $\mathbf{b}_{h,t} = \arg \min_{\mathbf{t}_h^{t-1}} \mathbf{t}_h^{t-1} \mathbf{y}_h^{t-1} = \mathbf{t}_h^{t-1} \mathbf{t}_h^{t-1} \mathbf{y}_h^{t-1} \mathbf{y}_h^{t-1}$ and $\mathbf{t}_h^{t-1} := \mathbf{t}_h^{t-1} \mathbf{t}_h^{t-1}$, simple algebra yields

$$\mathbb{E}_{h,f} \mathbf{t}_h^{t-1} \mathbf{y}_h^{t-1} \geq \inf_{\mathbf{t}_h^{t-1}} \mathbf{t}_h^{t-1} \mathbf{y}_h^{t-1} = \mathbb{E}_{h,f} \mathbf{b}_{h,t} \mathbf{t}_h^{t-1} = \mathbb{E}_{h,f} \mathbf{b}_{h,t} \mathbf{t}_h^{t-1} \quad (E.2)$$

So Algorithm 1 reduces to Algorithm 2.

In particular, the confidence region defined by (E.2) in Algorithm 2 is the same as the confidence region used in the upper confidence RL with value-targeted regression (UCRL-VTR) algorithm [Jia et al., 2020, Ayoub et al., 2020]. While in UCRL-VTR, they perform step-by-step optimization within the confidence region, resulting in a confidence bonus added upon the value function, our Algorithm 2 follows a global optimization scheme, where the objective is the total expected return

Algorithm 2 OPERA (linear mixture MDPs)

- 1: Initialize: $D_h = ?$ for $h = 1; \dots; H$
- 2: for iteration = 1; 2; ...; T do
- 3: Set $t := t$ where t is taken as $\arg\max_{f \in \mathcal{F}} Q_{1,f}(s_1; f(s_1))$ subject to

$$b_{h,t} = \begin{matrix} t-1 \\ h \end{matrix} > \begin{matrix} t-1 \\ h \end{matrix} > \begin{matrix} t-1 \\ h \end{matrix} y_{h,t-1} > ; \quad b_{h,t} \leq \begin{matrix} 2 \\ h \end{matrix} \quad \text{for all } h \in [H] \quad (\text{E.3})$$

- 4: For any $h \in [H]$, collect tuple $(r_h; s_h; a_h; s_{h+1})$ by executing $s_h; a_h$
 - 5: Augment $D_h = D_h \cup \{f(r_h; s_h; a_h; s_{h+1})g\}$
 - 6: end for
 - 7: Output: π_{out} uniformly sampled from $\{g_{t=1}^T\}$
-

by following the optimal policy under the current hypothesis. The design principle of the algorithm is the same as the ELEANOR algorithm [Zanette et al., 2020a]. In fact, the difference between UCRL-VTR with Algorithm 2 is analogous to the difference between LSVI-UCB [Jin et al., 2020] with ELEANOR [Zanette et al., 2020a].

Algorithm 2 exhibits a \sqrt{H} regret bound and $d^2 H^2 = 2$ sample complexity result, as will be shown later in this subsection. Compared with the $d^3 H^4 = 2$ sample complexity in Du et al. [2021], our algorithm improves over the best-known results on general frameworks that subsumes linear mixture MDPs. We provide more comparisons on the linear mixture model in §B.

In terms of computation, assume that there exists a planning oracle for the optimization problem in Line 3 of Algorithm 2 that requires \mathcal{B} time complexity to solve. Then for each $t \in [T]; h \in [H]$, the computational complexity of the rest of the algorithm is dominated by the computation of $\begin{matrix} t-1 \\ h \end{matrix}$, and the total computational complexity would be $\mathcal{O}(\mathcal{B}T + d^2 HT)$.

Next, we proceed to prove that a linear mixture MDP belongs to ABC class with low FE dimension.

Proof of Proposition 8. In the linear mixture MDP model, we choose hypothesis class $\mathcal{G}_h = \{f_h\}$, and DEF function

$$\tilde{v}_{h,f}(o_h; f_{h+1}; g_h; v) = \begin{matrix} h \\ h;g \end{matrix} (s_h; a_h) + \sum_{s^0} (s_h; a_h; s^0) V_{h+1;f}(s^0) \quad r_h \quad V_{h+1;f}(s_{h+1}):$$

(a) **Decomposability.** Taking expectation over s_{h+1} and we obtain that

$$\mathbb{E}_{s_{h+1}} [\tilde{v}_{h,f}(o_h; f_{h+1}; g_h; v) \mid s_h; a_h] = \begin{matrix} h \\ h;g \end{matrix} (s_h; a_h) + \sum_{s^0} (s_h; a_h; s^0) V_{h+1;f}(s^0) :$$

Thus, we have

$$\begin{aligned} & \tilde{v}_{h,f}(o_h; f_{h+1}; g_h; v) - \mathbb{E}_{s_{h+1}} [\tilde{v}_{h,f}(o_h; f_{h+1}; g_h; v) \mid s_h; a_h] \\ &= \begin{matrix} h \\ h;g \end{matrix} (s_h; a_h) + \sum_{s^0} (s_h; a_h; s^0) V_{h+1;f}(s^0) - \begin{matrix} h \\ h;g \end{matrix} (s_h; a_h) \\ &= \tilde{v}_{h,f}(o_h; f_{h+1}; f_h; v): \end{aligned}$$

(b) **Global Discriminator Optimality** holds automatically since is independent of f .

(c) **Dominating Average EF.** We have the following inequality for linear mixture models:

$$\begin{aligned} & \mathbb{E}_{s_h; a_h} g \text{KE} [\begin{matrix} h \\ h;g \end{matrix} (o_h; f_{h+1}; f_h; v) \mid s_h; a_h] \leq \sum_{s^0} \begin{matrix} h \\ h;g \end{matrix} (s_h; a_h; s^0) V_{h+1;g}(s^0) \\ &= \mathbb{E}_{s_h; a_h} g \left(\begin{matrix} h \\ h;g \end{matrix} (s_h; a_h) + \sum_{s^0} (s_h; a_h; s^0) V_{h+1;g}(s^0) - \begin{matrix} h \\ h;g \end{matrix} (s_h; a_h) \right) \\ &= \begin{matrix} h \\ h;g \end{matrix} \mathbb{E}_{s_h; a_h} g \left(\sum_{s^0} (s_h; a_h; s^0) V_{h+1;g}(s^0) \right) : \quad (\text{E.4}) \end{aligned}$$

(d) Bellman Dominance. On the other hand, we know that

$$\begin{aligned}
 & E_{s_h; a_h} [Q_{h;f}(s_h; a_h) - r_h - V_{h+1;f}(s_{h+1})] \\
 &= E_{s_h; a_h} (Q_{h;f}(s_h; a_h) - r_h) + \sum_{s^0} (s_h; a_h; s^0) V_{h+1;f}(s^0) \\
 &= (Q_{h;f}(s_h; a_h) - r_h) + \sum_{s^0} (s_h; a_h; s^0) V_{h+1;f}(s^0) : \quad (E.5)
 \end{aligned}$$

(e) Low FE Dimension. Observe from Eqs.(E.4) and (E.5) that we can choose ABC function of a linear mixture MDP as

$$G_{h;f}(f;g) := (Q_{h;f}(s_h; a_h) - r_h) + \sum_{s^0} (s_h; a_h; s^0) V_{h+1;g}(s^0) : \quad (E.6)$$

The next Lemma 28 proves that the FE dimension with respect to the coupling function $G_{h;f}(f;g)$ is less than the effective dimension of the parameter space.

Lemma 28. The linear mixture MDP model has FE dimension $\Theta(d)$ with respect to the ABC defined in (E.6).

We prove Lemma 28 in §G.

Thus, we conclude our proof of Proposition 8. □

From the above Proof of Proposition 8, we see that linear mixture MDPs perfectly fit our framework. We apply Theorem 12 and Corollary 26 to linear mixture MDPs and conclude directly that Algorithm 2 has a regret upper bound $\tilde{O}(\sqrt{T})$ together with a sample complexity upper bound $\tilde{O}(T^2)$, matching the best-known results that uses a Hoeffding-type bonus for exploration.

E.2 Low Witness Rank MDPs

In this subsection, we provide a novel method for solving low Witness rank MDPs as a direct application of the OPERA algorithm. The witness rank is an important model-based assumption that covers several structural models including the factored MDPs [Kearns, 1998]. Also, all models with low Bellman rank structure belong to the class of low Witness rank models while the opposite does not hold [Sun et al., 2019]. Although the witness rank models can be solved in a model-free manner, model-free algorithms cannot find near-optimal solutions of general witness rank models in polynomial time. Meanwhile, existing frameworks [Sun et al., 2019, Du et al., 2021] with an efficient algorithm does not exhibit sharp sample complexity results. We recall that in low Witness rank settings, hypotheses on model-based parameters (transition kernel and reward function) are made. Based on this, there are two recent lines of related approaches. Sun et al. [2019] first proposed an algorithm that eliminates candidate models with high estimated witness model misfit. On the other hand, Du et al. [2021] proposed a general algorithmic framework that would imply an optimization-based algorithm on low Witness rank models.

We prove an improved sample complexity result over existing literature and illustrate the differences in design scheme of our algorithm. We present the pseudocode in Algorithm 3. Note that in Eq. (E.7), we replace the DEF in Eq(4.1) by (3.4). Next, we elaborate the design scheme of our algorithm in comparison with Sun et al. [2019] and Du et al. [2021]. Note that the DEF $E_{\theta} v(s_h; a_h; \theta) - v(s_h; a_h; s_{h+1})$ is similar with the discrepancy function used in Du et al. [2021] except for an importance sampling factor. Moreover, after taking over discriminator functions, the expected DEF equals the witnessed model misfit in Sun et al. [2019]. Although Du et al. [2021] did not explicitly give an algorithm for witness rank, we observe some general differences between OPERA and BiLin-UCB [Du et al., 2021]. The confidence region used in Algorithm 3 (simplified version for comparison) is $\mathbb{P}_i[(\hat{f}_i)^2 - \inf_g (\hat{g}_i)^2] \leq \frac{1}{m}$ centered at the optimal hypothesis, while the confidence region used in BiLin-UCB is $\mathbb{P}_i[\frac{1}{m} \sum_{j=1}^m \hat{f}_i(j)^2 - \inf_g \sum_{j=1}^m \hat{g}_i(j)^2] \leq \frac{1}{m}$ that bound an estimate of f centered at a_0 . Similarly as in BiLin-UCB, Sun et al. [2019] also attempts to bound a batched estimate of f . Their algorithm constantly eliminates out of range models, enforcing small witness model misfit on prior distributions. The analysis in Sun et al. [2019] and Du et al. [2021], however,

Algorithm 3 OPERA (Low Witness Rank MDPs)

- 1: Initialize: $D_h = ?$ for $h = 1; \dots; H$
 - 2: for iteration = 1; 2; \dots; T do
 - 3: Set $t := f_t$ where t is taken as $\arg\max_{f \in \mathcal{F}} Q_{1,f}(s_1; f(s_1))$ subject to

$$\left(\max_{v \in \mathcal{V}} \prod_{i=1}^{X-1} E_{\mathbf{e}_{f_h}} v(s_h^i; \mathbf{a}_h^i; \mathbf{e}) v(s_h^i; \mathbf{a}_h^i; s_{h+1}^i)^2 \right. \\ \left. \inf_{g_h \in \mathcal{G}_h} \prod_{i=1}^{X-1} E_{\mathbf{e}_{g_h}} v(s_h^i; \mathbf{a}_h^i; \mathbf{e}) v(s_h^i; \mathbf{a}_h^i; s_{h+1}^i)^2 \right) \quad \text{for all } h \in [H] \quad (\text{E.7})$$
 - 4: For any $h \in [H]$, collect tuple $(r_h; s_h; \mathbf{a}_h; s_{h+1})$ by rolling in $s_h \sim U(A)$ and executing g_h^t
 - 5: Augment $D_h = D_h \cup \{(r_h; s_h; \mathbf{a}_h; s_{h+1})\}$
 - 6: end for
 - 7: Output: s_{out} uniformly sampled from $\prod_{g=1}^T g_{g=1}^T$
-

does not enforce the additional assumption on the discriminator class; we obtain a sharper sample complexity as in Corollary 13.

If we assume that there exists a planning oracle for solving the optimization problem in Line 3 of Algorithm 3 with B time complexity. The computation in the rest of the algorithm is dependent on the structure of discriminator class and the hypothesis class. We omit the discussion here as the planning oracle with a total computational complexity $O(BT)$ is usually the dominating term. In the forthcoming, we prove that low Witness rank MDPs belongs to ABC class with low FE dimension.

Proof of Proposition 10. In the low Witness rank model, we choose hypothesis class $\mathcal{G}_h = \mathcal{M}$, and DEF function

$$\tilde{v}_h(\mathbf{O}_h; f_{h+1}; g_h; v) = E_{\mathbf{e}_{g_h}} v(s_h; \mathbf{a}_h; \mathbf{e}) v(s_h; \mathbf{a}_h; s_{h+1}): \quad (\text{E.8})$$

Without loss of generality, we assume that the discriminator class is rich enough in the sense that if $\exists s; a \in \mathcal{S} \times \mathcal{A}$, $v_{s;a}(s; a; s^0) \in \mathcal{V}$, then $v(s; a; s^0) := v_{s;a}(s; a; s^0) \in \mathcal{V}$ (if not, we can use a rich enough \mathcal{V}^0 induced by \mathcal{V}), an assumption generally satisfied by common discriminator classes. For example, Total variation, Exponential family, MMD, Factored MDP in Sun et al. [2019] all use a rich enough discriminator class. Also, $\forall f = f_v : kvk_1 \leq c$ for some absolute constant c , the function class is rich enough.

(a) **Decomposability.** Taking expectation over s_{h+1} of Eq. (E.8) and we obtain that

$$E_{s_{h+1}} [\tilde{v}_h(\mathbf{O}_h; f_{h+1}; g_h; v) | s_h; \mathbf{a}_h] = E_{\mathbf{e}_{g_h}} v(s_h; \mathbf{a}_h; \mathbf{e}) E_{\mathbf{e}_{P_h}} v(s_h; \mathbf{a}_h; \mathbf{e}): \quad (\text{E.9})$$

Thus, we have

$$\tilde{v}_h(\mathbf{O}_h; f_{h+1}; g_h; v) E_{s_{h+1}} [\tilde{v}_h(\mathbf{O}_h; f_{h+1}; g_h; v) | s_h; \mathbf{a}_h] = E_{\mathbf{e}_{P_h}} v(s_h; \mathbf{a}_h; \mathbf{e}) v(s_h; \mathbf{a}_h; s_{h+1}) \\ = \tilde{v}_h(\mathbf{O}_h; f_{h+1}; f_h; v):$$

(b) **Global Discriminator Optimality.** Eq. (E.9) implies that

$$E_{s_{h+1}} [\tilde{v}_h(\mathbf{O}_h; f_{h+1}; f_h) | s_h; \mathbf{a}_h] = \int v(s_h; \mathbf{a}_h; s) (f_h(s | s_h; \mathbf{a}_h) P_h(s | s_h; \mathbf{a}_h)) ds:$$

We define $v_h(f)(s; a; s^0) = v_{s;a}(s; a; s^0)$ where

$$v_{s;a} := \arg \max_{v \in \mathcal{V}} v(s; a; \mathbf{e}) (f_h(\mathbf{e} | s; a) P_h(\mathbf{e} | s; a)) ds:$$

It is easy to verify that $v_h(f)$ satisfies for all $h \in [H]$ and $(s_h; \mathbf{a}_h) \in \mathcal{S} \times \mathcal{A}$,

$$E_{s_{h+1}} [\tilde{v}_h(\mathbf{O}_h; f_{h+1}; f_h; v_h(f)) | s_h; \mathbf{a}_h] = E_{s_{h+1}} [\tilde{v}_h(\mathbf{O}_h; f_{h+1}; f_h; v) | s_h; \mathbf{a}_h]:$$

Finally, the symmetry of \tilde{v}_h concludes the global discriminator optimality.

(c) Dominating Average EF. We have the following inequality for low Witness rank model:

$$\begin{aligned} & \max_{v_2V} E_{S_h, g; a_h, f} k E [\text{h}(\theta_h; f_{h+1}; f_h; v) \text{j} S_h; a_h] k_2^2 \\ & = \max_{v_2V} E_{S_h, g; a_h, f} (E_{\theta} f_h v(S_h; a_h; \theta) - E_{\theta} P_h v(S_h; a_h; \theta))^2 \\ & \max_{v_2V} E_{S_h, g; a_h, f} [E_{\theta} f_h v(S_h; a_h; \theta) - E_{\theta} P_h v(S_h; a_h; \theta)]^2 \stackrel{(i)}{=} h W_h(f); X_h(g) i^2: \end{aligned} \tag{E.10}$$

where the last inequality (i) follows Definition 9 of witness rank.

(d) Bellman Dominance. On the other hand, by Definition 9 we know that

$$E_{S_h; a_h, f} [Q_{h,f}(S_h; a_h) - r_h - V_{h+1; f}(S_{h+1})] = h W_h(f); X_h(f) i: \tag{E.11}$$

(e) Low FE Dimension. We see from Eq.(E.10) and (E.11) that we can choose ABC function with low Witness rank RL model as

$$G_{h,f}(f; g) := h W_h(f); X_h(g) i: \tag{E.12}$$

The next Lemma 29 proves that the FE dimension of with respect to the coupling function $G_{h,f}(f; g)$ is less than the dimension of the witness model.

Lemma 29. The low Witness rank MDP model has FE dimension $\Theta(W)$ with respect to the ABC defined in (E.12).

We prove Lemma 29 in §G.

Thus, we conclude our proof of Proposition 10. □

By Proposition 10 we can straightforwardly derive the sample complexity by applying Corollary 27. For better understanding of the context, we present a complete proof of the sample complexity result of witness rank model in §E.4.

E.3 Kernelized Nonlinear Regulator

In the KNR setting introduced in §3.2, the norms $\| \cdot \|_1$ might be arbitrarily large if the random vector s_{h+1} is large in magnitude. On the contrary, our framework requires the boundedness of the DEF. To resolve this issue, we note the tail bound of one-dimensional Gaussian distribution indicates that for any given positive ϵ :

$$e^{-x^2/2} \int_x^{\infty} e^{-t^2/2} dt \leq e^{-x^2/2} \int_x^{\infty} \frac{t}{x} e^{-t^2/2} dt = \frac{1}{x}:$$

Thus, for T i.i.d. \mathbb{R}^{d_s} -valued random vectors $\{s_i\}_{i=1}^T \sim N(0, I)$ and a fixed $\epsilon \in (0, 1)$, there exists an even B with $P(B) \geq 1 - \epsilon$ such that $\|s_i\|_1 \leq \frac{1}{\epsilon} \sqrt{\log(T/d_s)}$ holds on even B .

We first provide the application of OPERA on the KNR model, the algorithm is written in Algorithm 4. Note that by similar algebra as in Eq. (E.2), the confidence set (E.13) is equivalent to

$$(U_{h,f} - \hat{\theta}_{h,f}) (\frac{1}{h})^{1=2} \leq \frac{2}{h};$$

where $\frac{1}{h} := \frac{1}{h} (\frac{1}{h})^>$ and $\hat{\theta}_{h,f}$ is the optimal solution to the least square problem $\arg \min_U \sum_{i=1}^T U(s_i; a_h) - s_{h+1}^2$. The OPERA algorithm reduces to the \mathcal{L} algorithm in Kakade et al. [2020] except that \mathcal{L} is under a homogeneous setting. The only difference between Algorithm 4 and \mathcal{L} is that in Eq.(E.13), \mathcal{L} sums over H and we can only sum over T because of the inhomogeneous setting.

Bringing in the choice of ϵ in Corollary 14 yields a regret bound $\mathcal{O} \left(\frac{1}{\epsilon^2} d_s H^4 T \right)$. In comparison, \mathcal{L} in Kakade et al. [2020] has a regret bound $\mathcal{O} \left(\frac{1}{\epsilon^2} (d_s + d) H^3 T \right)$. The improved factor of $\frac{1}{\epsilon^2} H$ is due to the reduction from the time-inhomogeneous setting to the time-homogeneous setting. Thus, our regret bound matches the state-of-the-art result on KNR instances [Kakade et al.,

Algorithm 4 OPERA (kernelized nonlinear regulator)

- 1: Initialize: $D_h = ?$ for $h = 1; \dots; H$
- 2: for iteration = 1; 2; \dots; T do
- 3: Set $\mathbf{t} := \mathbf{f}_t$ where \mathbf{t} is taken as $\arg\max_{\mathbf{f}} Q_{1,\mathbf{f}}(s_1; \mathbf{f}(s_1))$ subject to

$$\prod_{i=1}^X U_{h,\mathbf{f}}(s_h^i; \mathbf{a}_h^i) \prod_{i=1}^2 s_{h+1}^i \inf_{g_h \in \mathcal{G}_h} \prod_{i=1}^X U_{h,g}(s_h^i; \mathbf{a}_h^i) \prod_{i=1}^2 s_{h+1}^i \quad \text{for all } h \in [H] \quad (\text{E.13})$$

- 4: For any $h \in [H]$, collect tuple $(r_h; s_h; \mathbf{a}_h; s_{h+1})$ by executing $s_h; \mathbf{a}_h \xrightarrow{\mathbf{t}}$
 - 5: Augment $D_h = D_h \cup \{(r_h; s_h; \mathbf{a}_h; s_{h+1})g\}$
 - 6: end for
 - 7: Output: s_{out} uniformly sampled from $\prod_{t=1}^T g_t^T$
-

[2020] regarding the dependencies on $d_s; H$. However, $d^2 d_s$ in our result is slightly looser than $d(d_s + d)$ in Kakade et al. [2020] and can be possibly improved by instance-specific analysis of KNR. Similar to the linear mixture MDP case, if we assume that there exists a planning oracle for the optimization problem in Line 3 of Algorithm 4 that requires time complexity to solve, the rest of the algorithm can be solved efficiently $\mathcal{O}(d(d + d_s)HT)$ time complexity. So the total computational complexity of Algorithm 4 is $\mathcal{O}(BT + d(d + d_s)HT)$.

In addition, we would like to remark that we can adapt algorithms from the optimal control literature such as MPPI [Williams et al., 2015] and DMDMPC [Wagener et al., 2019] to solve the optimization problem in Line 3 of Algorithm 4. This approach has been used in Kakade et al. [2020], where they designed the L \mathcal{C} algorithm for solving KNRs. In particular, they leveraged the MPPI algorithm for the planning oracle and provided rich empirical results.

Proof of Proposition 11. In the KNR model, we choose hypothesis class $\mathcal{F}_s = \mathcal{G}_h = \mathbf{f} \cup \mathbf{2}^H$! $R^{d_s} : kUk_2$ Rg, and DEF function

$$\tilde{V}_h(\mathbf{o}_h; \mathbf{f}_{h+1}; \mathbf{g}_h; \mathbf{v}) = U_{h,g}(s_h; \mathbf{a}_h) \prod_{i=1}^2 s_{h+1}^i$$

- (a) **Decomposability.** Taking expectation over s_{h+1} and we obtain that

$$E_{s_{h+1}}[\tilde{V}_h(\mathbf{o}_h; \mathbf{f}_{h+1}; \mathbf{g}_h; \mathbf{v}) | s_h; \mathbf{a}_h] = (U_{h,g} \prod_{i=1}^2 U_h)(s_h; \mathbf{a}_h)$$

Thus, we have

$$\tilde{V}_h(\mathbf{o}_h; \mathbf{f}_{h+1}; \mathbf{g}_h; \mathbf{v}) = E_{s_{h+1}}[\tilde{V}_h(\mathbf{o}_h; \mathbf{f}_{h+1}; \mathbf{g}_h; \mathbf{v}) | s_h; \mathbf{a}_h] = U_h(s_h; \mathbf{a}_h) \prod_{i=1}^2 s_{h+1}^i = \tilde{V}_h(\mathbf{o}_h; \mathbf{f}_{h+1}; \mathbf{f}_h; \mathbf{v})$$

- (b) **Global Discriminator Optimality** holds automatically since \tilde{V}_h is independent of \mathbf{f}_h .

- (c) **Dominating Average EF.** We have the following inequality for the KNR model:

$$E_{s_h; \mathbf{a}_h} \int g \cdot kE[\tilde{V}_h(\mathbf{o}_h; \mathbf{f}_{h+1}; \mathbf{f}_h; \mathbf{v}) | s_h; \mathbf{a}_h]k_2^2 = E_{s_h; \mathbf{a}_h} \int g \cdot k(U_{h,\mathbf{f}} \prod_{i=1}^2 U_h)(s_h; \mathbf{a}_h)k_2^2: \quad (\text{E.14})$$

- (d) **Bellman Dominance.** On the other hand, we know that

$$E_{s_h; \mathbf{a}_h} \int [Q_{h,\mathbf{f}}(s_h; \mathbf{a}_h) - r_h - V_{h+1,\mathbf{f}}(s_{h+1})] \frac{2H}{\epsilon} E_{s_h; \mathbf{a}_h} \int g \cdot k(U_{h,\mathbf{f}} \prod_{i=1}^2 U_h)(s_h; \mathbf{a}_h)k_2: \quad (\text{E.15})$$

- (e) **Low FE Dimension.** We see from Eqs.(E.14) and (E.15) that we can choose ABC function of a linear mixture MDP as

$$G_{h,\mathbf{f}}(\mathbf{f}; \mathbf{g}) := \frac{q}{E_{s_h; \mathbf{a}_h} \int g \cdot k(U_{h,\mathbf{f}} \prod_{i=1}^2 U_h)(s_h; \mathbf{a}_h)k_2^2}; \quad (\text{E.16})$$

and KNR has an ABC with $\epsilon = \frac{2H}{q}$. The next Lemma 30 proves that the FE dimension of \mathcal{F} with respect to the coupling function $G_{h,\mathbf{f}}(\mathbf{f}; \mathbf{g})$ can be controlled by d :

Lemma 30. The KNR model has FE dimension $\mathcal{O}(d)$ with respect to the ABC defined in (E.16).

We prove Lemma 30 in §G.

Thus, we conclude our proof of Proposition 11. □

E.4 Proof of Corollary 13

In this subsection, we provide sample complexity guarantee for models with low Witness rank. In the main text in §4.3 we presented our Corollary 13 with finite cardinality for convenience of comparison with previous works. Here, we prove general result for model class and discriminator class with finite ϵ -covering.

Proof of Corollary 13. We start the proof by showing that $\|V_1(s_1) - V_1^t(s_1)\|$ can be upper bounded by a sum of Bellman errors, which is a simple deduction from the policy loss decomposition lemma in Jiang et al. [2017] and is the same as the equality in (E.4) in the proof of Theorem 12 in §C. Next, we verify that π^t satisfies constraint (E.7) so that taking $\pi^t = \arg \max_{\pi} V_{1,f}(s_1)$ in the confluence region yields $\|V_1(s_1) - V_1^t(s_1)\|$.

Lemma 31 (Feasibility of π^t). In Algorithm 3, given $\epsilon > 0$ and $\delta > 0$, we choose $\epsilon = c(\log(THjM) + T)$ for some large enough constant, then with probability at least $1 - \delta$, π^t satisfies for any $t \geq 2$ [T]:

$$\max_{\pi} \sum_{i=1}^{H-1} \mathbb{E}_{\mathbf{s}_h \sim f_h} v(\mathbf{s}_h^i; \mathbf{a}_h^i; \boldsymbol{\theta}) - v(\mathbf{s}_h^i; \mathbf{a}_h^i; \mathbf{s}_{h+1}^i)^2 \leq \epsilon$$

$$\inf_{g_h \in \mathcal{G}_h} \sum_{i=1}^{H-1} \mathbb{E}_{\mathbf{s}_h \sim g_h} v(\mathbf{s}_h^i; \mathbf{a}_h^i; \boldsymbol{\theta}) - v(\mathbf{s}_h^i; \mathbf{a}_h^i; \mathbf{s}_{h+1}^i)^2 \leq \delta$$

We prove Lemma 31 in §F.5. The next Lemma 32 is devoted to controlling the average squared DEF.

Lemma 32. In Algorithm 3, given $\epsilon > 0$ and $\delta > 0$, we choose $\epsilon = c(\log(THjM) + T)$ for some large enough constant, then with probability at least $1 - \delta$, for all $(t; h) \geq 2$ [T] [H], we have

$$\sum_{i=1}^{H-1} \max_{\pi} \mathbb{E}_{\mathbf{s}_h \sim \pi} \mathbb{E}_{\mathbf{a}_h \sim f_h} v(\mathbf{s}_h^i; \mathbf{a}_h^i; \boldsymbol{\theta}) - v(\mathbf{s}_h^i; \mathbf{a}_h^i; \boldsymbol{\theta})^2 = O(jAj):$$

Proof is delayed to §F.4. By Lemma 32 and properties of the witness rank in Definition 9, we have

$$\sum_{i=1}^{H-1} \mathbb{H}W_h(f); X_h(f_i) \leq \sum_{i=1}^{H-1} \max_{\pi} \mathbb{E}_{\mathbf{s}_h \sim \pi} \mathbb{E}_{\mathbf{a}_h \sim f_h} v(\mathbf{s}_h^i; \mathbf{a}_h^i; \boldsymbol{\theta}) - v(\mathbf{s}_h^i; \mathbf{a}_h^i; \boldsymbol{\theta})^2$$

$$\sum_{i=1}^{H-1} \max_{\pi} \mathbb{E}_{\mathbf{s}_h \sim \pi} \mathbb{E}_{\mathbf{a}_h \sim f_h} v(\mathbf{s}_h^i; \mathbf{a}_h^i; \boldsymbol{\theta}) - v(\mathbf{s}_h^i; \mathbf{a}_h^i; \boldsymbol{\theta})^2 = O(jAj):$$

Applying Lemma 25 with $G_{h,f}(\mathbf{f}; \mathbf{g}) := \mathbb{H}W_h(\mathbf{f}); X_h(\mathbf{g})$ and $\mathbf{g}_t = \mathbf{f}^t; \mathbf{f}_t = \mathbf{f}^t$, we have

$$\sum_{i=1}^t \mathbb{H}W_h(\mathbf{f}); X_h(\mathbf{f}_i) \leq O\left(\frac{q}{jAj \dim_{FE}(F; G_{h,f}; !)} t + t!\right):$$

Policy loss decomposition (C.2) yields

$$\frac{1}{T} \sum_{t=1}^T V_1(s_1) - V_1^t(s_1) \leq O\left(\frac{H}{jAj \dim_{FE}(F; G_{h,f}; !)} \frac{\log(THjM) + T}{T} + \frac{H!}{T}\right):$$

Taking $\frac{1}{T} = \frac{1}{H}$ and $\frac{1}{T} = \frac{2}{\dim_{FE}(F; G; \frac{1}{H})H^2}$, the above Eq. (C.4) becomes

$$\frac{1}{T} \sum_{t=1}^T V_1(s_1) - V_1^t(s_1) \leq O\left(\frac{H}{jAj} \frac{\dim_{FE}(F; G; \frac{1}{H}) \log(THjM) + T}{T}\right):$$

Taking

$$T = \frac{jAj \dim_{FE}(F; G; \frac{1}{H}) \log(THjM) + T}{2 \frac{2}{H+1}}$$

yields the desired result. \square

E.5 Proof of Corollary 14

We can directly apply Theorem 12 to the KNR model based on Proposition 11 to obtain the regret bound result. For better understanding of our framework, we illustrate the main features in the proof of Corollary 14 that are different from the proof of Theorem 12.

Proof of Corollary 14. To resolve the unboundedness issue, we unfold the analysis of KNR case and conclude a high-probability event analogous to the argument in §E.3. However, doing so would impose an additional \bar{d}_s factor induced by estimating the ℓ_2 -norm of multivariate Gaussians. In lieu to this, we present a sharper convergence analysis that incorporates KNR instance-specific structures. We recall the DEF of the KNR model:

$$\ell_h(\alpha_h; f_{h+1}; g_h; v) = U_{h,g}(s_h; a_h) - s_{h+1} :$$

We first define an auxiliary random variable

$$\begin{aligned} X_t(h; f; v) &:= \ell_h(\alpha_h^t; f_{h+1}; g_h; v)^2 - \ell_h(\alpha_h^t; f_{h+1}; T(f)_h; v)^2 \\ &= U_{h,f}(s_h^t; a_h^t) - s_{h+1}^t - \frac{1}{2} (U_h(s_h^t; a_h^t) - s_{h+1}^t)^2 \\ &= (U_{h,f} - U_h)(s_h^t; a_h^t) + \frac{1}{2} (U_{h,f} - U_h)(s_h^t; a_h^t) - \frac{1}{2} s_{h+1}^t \\ &= (U_{h,f} - U_h)(s_h^t; a_h^t) - \frac{1}{2} s_{h+1}^t : \end{aligned}$$

By the boundedness of operator $U_{h,f}$, U_h and uniform boundedness of $(s; a)$, we obtain that $k(U_{h,f} - U_h)(s_h^t; a_h^t)k_2^2 \leq 4B_U^2 B^2$. The conditional distribution of $(U_{h,f} - U_h)(s_h^t; a_h^t); s_{h+1}^t$ is a zero-mean Gaussian with variance $2k(U_{h,f} - U_h)(s_h^t; a_h^t)k_2^2 + 4B_U^2 B^2$. By the tail bound of Gaussian distributions along with standard union bound, we know that with probability at least $1 - \frac{\delta}{2}$,

$$(U_{h,f} - U_h)(s_h^t; a_h^t); s_{h+1}^t \leq O\left(\sqrt{\frac{\log(TH)}{\delta}}\right)$$

holds uniformly for all $t \in [T]$ and $h \in [H]$. Thus, we bound the absolute value of the auxiliary variable X_t by $|X_t| \leq R$ where R is positive and of order $O\left(\sqrt{\frac{\log(TH)}{\delta}}\right)$. Taking expectation with respect to s_{h+1} , we have

$$\mathbb{E}_{s_{h+1}} [X_t(h; f; v) | s_h; a_h] = (U_{h,f} - U_h)(s_h^t; a_h^t)^2 :$$

On the other hand,

$$\begin{aligned} &\mathbb{E}_{s_{h+1}} (X_t(h; f; v))^2 | s_h; a_h \\ &= \mathbb{E}_{s_{h+1}} (U_{h,f} - U_h)(s_h^t; a_h^t) - \frac{1}{2} s_{h+1}^t - \frac{1}{2} (U_{h,f} - U_h)(s_h^t; a_h^t) - \frac{1}{2} s_{h+1}^t)^2 | s_h; a_h \\ &= \mathbb{E}_{s_{h+1}} (U_{h,f} - U_h)(s_h^t; a_h^t)^2 + 4 (U_{h,f} - U_h)(s_h^t; a_h^t) - \frac{1}{2} s_{h+1}^t)^2 | s_h; a_h \\ &= \mathbb{E}_{s_{h+1}} (U_{h,f} - U_h)(s_h^t; a_h^t)^2 + 4 (U_{h,f} - U_h)(s_h^t; a_h^t) - \frac{1}{2} s_{h+1}^t)^2 | s_h; a_h \\ &\leq 2R^2 \mathbb{E} [X_t(h; f; v) | s_h; a_h] : \end{aligned}$$

By taking $Z_t = X_t(h; f; v) - \mathbb{E}_{s_{h+1}} [X_t(h; f; v) | s_h; a_h]$ with $|Z_t| \leq 2R$ in Freedman's inequality (F.1) in Lemma 33, we have for any satisfying $0 < \delta < \frac{1}{2R^2 - 2}$ almost surely, with probability at least $1 - \frac{\delta}{2}$:

$$\sum_{i=1}^t Z_i \leq O\left(R^2 \sqrt{t}\right) + \sum_{i=1}^t \mathbb{E}_{s_{h+1}} [X_i(h; f; v) | s_h; a_h] + \frac{\log(1/\delta)}{\delta} :$$

Optimizing over t , we have

$$\sum_{i=1}^t Z_i \leq O\left(R^2 \sqrt{t}\right) + \sum_{i=1}^t \mathbb{E}_{s_{h+1}} [X_i(h; f; v) | s_h; a_h] \log(1/\delta) + R^2 \sqrt{t} \log(1/\delta) : \quad (E.17)$$

Following the same Freedman's inequality (Lemma 33) and covering argument as in the proof of Theorem 12 with derivations detailed in §F.1, we have with probability $1 - \delta$ and $\epsilon = O\left(\frac{1}{T}\right)$:

$$\sum_{i=1}^T \mathbb{E}_{s_h; a_h} \left[k(U_{h,f}^t - U_h) (s_h; a_h) k_2 \right]^2 \leq \sum_{i=1}^T \mathbb{E}_{s_h; a_h} \left[k(U_{h,f}^t - U_h) (s_h; a_h) k_2 \right]^2 = O\left(\frac{1}{T}\right);$$

Feasibility of \hat{g} can be derived by taking the same auxiliary random variable and analyze on $\sum_{i=1}^t X_i(h; f; v)$ as in the proof of Lemma 31.

As explained in §E.3, we can apply Lemma 25 with $\frac{q}{T}, \frac{1}{T} = \frac{1}{T}$,

$$G_{h,f}(f; g) = \frac{q}{\mathbb{E}_{s_h; a_h} \left[k(U_{h,f} - U_h) (s_h; a_h) k_2 \right]^2};$$

and have

$$\sum_{i=1}^T \mathbb{E}_{s_h; a_h} \left[k(U_{h,f}^t - U_h) (s_h; a_h) k_2 \right]^2 \leq \frac{r}{\dim_{FE}(F; G; \frac{p}{1-T})} \frac{p}{1-T} \log(\text{THN}_L(1-T)) \leq t;$$

The rest of the proof follows by applying Bellman dominance, policy loss decomposition and calculating the FE dimension based $G_{h,f}(f; g)$, which is shown in Lemma 30. We therefore obtain that

$$\begin{aligned} \sum_{t=1}^T V_1(s_t) - V_1^*(s_t) &\leq \sum_{t=1}^T \sum_{h=1}^H j_{G_{h,f}}(f^t; f^t) \\ &\leq O\left(\frac{H}{T} \frac{q}{\dim_{FE}(F; G; \frac{p}{1-T})} \frac{p}{1-T} \log(\text{THN}_L(1-T))\right) \\ &= O\left(H^2 \frac{q}{d^2 d_s T}\right); \end{aligned}$$

□

F Proof of Technical Lemmas

We start with introducing the Freedman's inequality that are crucial in proving concentration properties in our main results.

Lemma 33 (Freedman-Style Inequality, Agarwal et al. 2014) Consider an adapted sequence $\{Z_t; J_{t-1}\}_{t=1;2;\dots;T}$ that satisfies $\mathbb{E}[Z_t | J_{t-1}] = 0$ and $Z_t \in \mathbb{R}$ for any $t = 1; 2; \dots; T$. Then for any $\epsilon > 0$ and $\delta \in [0; \frac{1}{2}]$, it holds with probability at least $1 - \delta$ that

$$\sum_{t=1}^T Z_t \leq (e - 2) \sum_{t=1}^T \mathbb{E}[Z_t^2 | J_{t-1}] + \frac{\log(\frac{1}{\delta})}{\epsilon}; \tag{F.1}$$

Before proving our technical lemmas, we note that for notational simplicity we use the expectation $\mathbb{E}_{s_h; a_h}[\cdot | s_h; a_h]$ to denote the conditional expectation with respect to the transition probability of the true model \mathbb{P}_h . The value of $s_h; a_h$ is data dependent (might be $s_h^t; a_h^t$ or $s_h^{t+1}; a_h^{t+1}$ depending on the function inside the expectation).

F.1 Proof of Lemma 24

Proof of Lemma 24 We recall that $\hat{\cdot}$ has a bounded ℓ_2 -norm in Definition 6 and assume that $\hat{k}_{h,f}(\cdot; f_{h+1}; g_h; v) k_2 \leq R$ for $h \in [H]; f \in F; g \in G; v \in V$ throughout the paper. For a sequence of data $\mathcal{D}_h = \{r_h^t; s_h^t; a_h^t; s_{h+1}^t\}_{t=1;2;\dots;T}$, we first build an auxiliary random variable defined for every $(t; h; f; v) \in [T] \times [H] \times F \times V$ and consider

$$X_i(h; f; v) := \hat{k}_{h,f}(\alpha_h^i; f_{h+1}; f_h; v) \frac{1}{2} - \hat{k}_{h,f}(\alpha_h^i; f_{h+1}; T(f)_h; v) \frac{1}{2};$$

where the randomness is due to uniformly sampling the data sequence. We know that $X_t(h; f) \leq R^2$. Take conditional expectation \mathbb{E}_i with respect to $s_h; a_h$, we have by definition that

$$\begin{aligned} \mathbb{E}_{S_{h+1}} [X_i(h; f; v) | s_h; a_h] \\ = \mathbb{E}_{S_{h+1}} \left[\mathbb{E}_{h,f} (\alpha_h^i; f_{h+1}; f_h; v) \right] - \mathbb{E}_{h,f} (\alpha_h^i; f_{h+1}; T(f)_h; v) \end{aligned}$$

Using the fact that $\|a - b\|^2 = \|a\|^2 + \|b\|^2 - 2\langle a, b \rangle$ for arbitrary vectors $a; b$ and property (i) in Definition 6 we have

$$\begin{aligned} \mathbb{E}_{S_{h+1}} [X_i(h; f; v) | s_h; a_h] &= \mathbb{E}_{h,f} (\alpha_h^i; f_{h+1}; f_h; v) - \mathbb{E}_{h,f} (\alpha_h^i; f_{h+1}; T(f)_h; v); \\ &= \mathbb{E}_{S_{h+1}} \left[\mathbb{E}_{h,f} (\alpha_h^i; f_{h+1}; f_h; v) + \mathbb{E}_{h,f} (\alpha_h^i; f_{h+1}; T(f)_h; v) | s_h; a_h \right] \\ &= \mathbb{E}_{S_{h+1}} \left[\mathbb{E}_{h,f} (\alpha_h^i; f_{h+1}; f_h; v) | s_h; a_h \right] \end{aligned}$$

On the other hand,

$$\begin{aligned} \mathbb{E}_{S_{h+1}} \left[(X_i(h; f; v))^2 | s_h; a_h \right] &= \mathbb{E}_{S_{h+1}} \left[\mathbb{E}_{h,f} (\alpha_h^i; f_{h+1}; f_h; v) - \mathbb{E}_{h,f} (\alpha_h^i; f_{h+1}; T(f)_h; v) \right]^2 \\ &= \mathbb{E}_{S_{h+1}} \left[\mathbb{E}_{h,f} (\alpha_h^i; f_{h+1}; f_h; v) + \mathbb{E}_{h,f} (\alpha_h^i; f_{h+1}; T(f)_h; v) \right]^2 \\ &\leq 4 \mathbb{E}_{S_{h+1}} \left[\mathbb{E}_{h,f} (\alpha_h^i; f_{h+1}; f_h; v) | s_h; a_h \right]^2 + 4 \mathbb{E}_{S_{h+1}} \left[\mathbb{E}_{h,f} (\alpha_h^i; f_{h+1}; T(f)_h; v) | s_h; a_h \right]^2 \\ &= 4R^2 \mathbb{E}_{S_{h+1}} [X_i(h; f; v) | s_h; a_h] \end{aligned}$$

By taking $Z_t = X_t(h; f; v) - \mathbb{E}_{S_{h+1}} [X_t(h; f; v) | s_h; a_h]$ with $|Z_t| \leq 2R^2$ in Freedman's inequality (F.1) in Lemma 33, we have for any satisfying $0 < \epsilon < \frac{1}{2R^2}$, with probability at least $1 - \epsilon$:

$$\begin{aligned} \sum_{i=1}^t Z_i &\leq \sum_{i=1}^t \mathbb{E}_{S_{h+1}} [X_i(h; f; v) | s_h; a_h] + \frac{\log(1/\epsilon)}{2R^2} \\ &\leq \sum_{i=1}^t \mathbb{E}_{S_{h+1}} [X_i^2(h; f; v) | s_h; a_h] + \frac{\log(1/\epsilon)}{2R^2} \\ &\leq 4R^2 \sum_{i=1}^t \mathbb{E}_{S_{h+1}} [X_i(h; f; v) | s_h; a_h] + \frac{\log(1/\epsilon)}{2R^2} \end{aligned}$$

Taking $\epsilon = \frac{1}{2R^2} \frac{\log(1/\epsilon)}{\sum_{i=1}^t \mathbb{E}_{S_{h+1}} [X_i(h; f; v) | s_h; a_h]}$, we have

$$\sum_{i=1}^t Z_i \leq \frac{\log(1/\epsilon)}{2R^2} \sum_{i=1}^t \mathbb{E}_{S_{h+1}} [X_i(h; f; v) | s_h; a_h] + \frac{\log(1/\epsilon)}{2R^2} \quad (F.2)$$

Similarly by applying Freedman's inequality to $\sum_{i=1}^t Z_t$ and combining with Eq(F.2), we have that for any three-tuple $(\epsilon; h; f)$, the following holds with probability at least $1 - \epsilon$:

$$\sum_{i=1}^t Z_i \leq \frac{\log(1/\epsilon)}{2R^2} \sum_{i=1}^t \mathbb{E}_{S_{h+1}} [X_i(h; f; v) | s_h; a_h] + \frac{\log(1/\epsilon)}{2R^2} \quad (F.3)$$

We note that in §3 we have that admits a ϵ -covering of $F; G; V$, meaning that for any $(f; g; v) \in F \times G \times V$ and a $\epsilon > 0$ there exists $(f^\epsilon; g^\epsilon; v^\epsilon) \in F_\epsilon \times G_\epsilon \times V_\epsilon$ such that $\|f - f^\epsilon\| \leq \epsilon; \|g - g^\epsilon\| \leq \epsilon; \|v - v^\epsilon\| \leq \epsilon$; where $F_\epsilon; G_\epsilon; V_\epsilon$ are ϵ -covers of $F; G; V$ respectively.

This is denoted by $(F_\epsilon; G_\epsilon; V_\epsilon) \in \mathcal{L}_\epsilon$. In definition of X_t , g is always taken as f^ϵ or a function of $T(f^\epsilon)$. Then if T is Lipschitz, as it is mostly the expectation operator, we omit the tuple and use $(f^\epsilon; g^\epsilon; v^\epsilon) \in \mathcal{L}_\epsilon$ to denote an element in the covering. By taking a union bound over \mathcal{L}_ϵ , we

have with probability at least $1 - 2^{-L}$ that the following holds for any $(\mathbf{f}^a; \mathbf{f}^t; \mathbf{v}) \in \mathcal{L}$,

$$\sum_{i=1}^h \mathbb{E}_{S_{h+1}} \left[\sum_{j \in S_h; a_h} \mathcal{X}_i(h; \mathbf{f}^a; \mathbf{v}) \right] \leq \sum_{i=1}^h \mathbb{E}_{S_{h+1}} \left[\sum_{j \in S_h; a_h} \mathcal{X}_i(h; \mathbf{f}^a; \mathbf{v}) \right] + 2R^2 A; \quad (\text{F.3})$$

where $\mathcal{X}_i(h; \mathbf{f}^a; \mathbf{v}) := \frac{1}{2} \sum_{h, \mathbf{f}^a} (\mathcal{O}_h^i; \mathbf{f}_{h+1}^a; \mathbf{f}_h^a; \mathbf{v})^2 - \frac{1}{2} \sum_{h, \mathbf{f}^a} (\mathcal{O}_h^i; \mathbf{f}_{h+1}^a; T(\mathbf{f}_h^a; \mathbf{v}))^2$ and $A = \log \frac{HTNL(\cdot)}{2L}$. Further for any $X_i(h; \mathbf{f}^t; \mathbf{v})$, we choose the three-tuple $(\mathbf{f}^a; \mathbf{f}^t; \mathbf{v}) := \arg \min_{(\mathbf{f}^a; \mathbf{f}^t; \mathbf{v}) \in \mathcal{L}} X_i(h; \mathbf{f}^t; \mathbf{v}) - \mathcal{X}_i(h; \mathbf{f}^a; \mathbf{v})$ and by the ϵ -covering argument, we arrive at

$$\sum_{i=1}^h \mathcal{X}_i(h; \mathbf{f}^a; \mathbf{v}) = \sum_{i=1}^h \left[\frac{1}{2} \sum_{h, \mathbf{f}^a} (\mathcal{O}_h^i; \mathbf{f}_{h+1}^a; \mathbf{f}_h^a; \mathbf{v})^2 - \frac{1}{2} \sum_{h, \mathbf{f}^a} (\mathcal{O}_h^i; \mathbf{f}_{h+1}^a; T(\mathbf{f}_h^a; \mathbf{v}))^2 \right] \\ + \sum_{i=1}^h \left[\frac{1}{2} \sum_{h, \mathbf{f}^t} (\mathcal{O}_h^i; \mathbf{f}_{h+1}^t; \mathbf{f}_h^t; \mathbf{v})^2 - \frac{1}{2} \sum_{h, \mathbf{f}^t} (\mathcal{O}_h^i; \mathbf{f}_{h+1}^t; T(\mathbf{f}_h^t; \mathbf{v}))^2 \right] + O(Rt) \stackrel{(i)}{=} O(\epsilon + Rt); \quad (\text{F.4})$$

where (i) comes from the constraint (4.1) of Algorithm 1.

Combining (F.3) with (F.4), we derive the following

$$\sum_{i=1}^h \mathbb{E}_{S_{h+1}} \left[\sum_{j \in S_h; a_h} \mathcal{X}_i(h; \mathbf{f}^a; \mathbf{v}) \right] \leq O(\epsilon + Rt + R^2);$$

Applying the ϵ -covering argument as in before, we conclude

$$\max_{\mathbf{v} \in \mathcal{V}} \sum_{i=1}^h \mathbb{E}_{S_{h+1}} \left[\sum_{j \in S_h; a_h} X_i(h; \mathbf{f}^t; \mathbf{v}) \right] \leq O(\epsilon + Rt + R^2);$$

Global optimality of the discriminator in (ii) of Definition 6 implies that \mathbf{a}_h is the optimal discriminator under any distribution or summation $\mathbb{E}_{\mathbf{a}_h}$ (and thus \max is interchangeable with summation):

$$\sum_{i=1}^h \mathbb{E}_{S_h; \mathbf{a}_h} \left[\sum_{j \in S_h; \mathbf{a}_h} \mathbb{E}_{S_{h+1}} \left[\sum_{h, \mathbf{f}^t} (\mathcal{O}_h^i; \mathbf{f}_{h+1}^t; \mathbf{f}_h^t; \mathbf{v}_h(\mathbf{f}^t)) \right] \right] \leq \frac{1}{2} \\ \sum_{i=1}^h \max_{\mathbf{v} \in \mathcal{V}} \mathbb{E}_{S_h; \mathbf{a}_h} \left[\sum_{j \in S_h; \mathbf{a}_h} \mathbb{E}_{S_{h+1}} \left[\sum_{h, \mathbf{f}^t} (\mathcal{O}_h^i; \mathbf{f}_{h+1}^t; \mathbf{f}_h^t; \mathbf{v}) \right] \right] \leq \frac{1}{2} \cdot 8v \cdot 2V;$$

Thus, we have

$$\sum_{i=1}^h \max_{\mathbf{v} \in \mathcal{V}} \mathbb{E}_{S_h; \mathbf{a}_h} \left[\sum_{j \in S_h; \mathbf{a}_h} \mathbb{E}_{S_{h+1}} \left[\sum_{h, \mathbf{f}^t} (\mathcal{O}_h^i; \mathbf{f}_{h+1}^t; \mathbf{f}_h^t; \mathbf{v}) \right] \right] \leq \frac{1}{2} \\ = \sum_{i=1}^h \mathbb{E}_{S_h; \mathbf{a}_h} \left[\sum_{j \in S_h; \mathbf{a}_h} \mathbb{E}_{S_{h+1}} \left[\sum_{h, \mathbf{f}^t} (\mathcal{O}_h^i; \mathbf{f}_{h+1}^t; \mathbf{f}_h^t; \mathbf{v}_h(\mathbf{f}^t)) \right] \right] \leq \frac{1}{2};$$

and also

$$\sum_{i=1}^h \mathbb{E}_{S_h; \mathbf{a}_h} \left[\sum_{j \in S_h; \mathbf{a}_h} \mathbb{E}_{S_{h+1}} \left[\sum_{h, \mathbf{f}^t} (\mathcal{O}_h^i; \mathbf{f}_{h+1}^t; \mathbf{f}_h^t; \mathbf{v}_h(\mathbf{f}^t)) \right] \right] \leq \frac{1}{2} \\ = \max_{\mathbf{v} \in \mathcal{V}} \sum_{i=1}^h \mathbb{E}_{S_h; \mathbf{a}_h} \left[\sum_{j \in S_h; \mathbf{a}_h} \mathbb{E}_{S_{h+1}} \left[\sum_{h, \mathbf{f}^t} (\mathcal{O}_h^i; \mathbf{f}_{h+1}^t; \mathbf{f}_h^t; \mathbf{v}) \right] \right] \leq \frac{1}{2} \\ = \max_{\mathbf{v} \in \mathcal{V}} \sum_{i=1}^h \mathbb{E}_{S_h; \mathbf{a}_h} \left[\sum_{j \in S_h; \mathbf{a}_h} X_i(h; \mathbf{f}^t; \mathbf{v}) \right] \leq O(\epsilon + Rt + R^2); \quad (\text{F.5})$$

We apply property (i) in Definition 5 and conclude that

$$\sum_{i=1}^X \mathbb{G}_{h,f}(f^t; f^i)^2 \leq O(\cdot);$$

which finishes the proof of Lemma 24. \square

F.2 Proof of Lemma 22

Proof of Lemma 22 For a data set $\mathcal{D}_h = \{f_h^t; s_h^t; a_h^t; s_{h+1}^t; g_{t=1;2;\dots;T}\}$, we first build an auxiliary random variable defined for every $(t; h; f; v) \in [T] \times [H] \times F \times V$

$$X_i(h; f; v) := \mathbb{E}_{h,f} \left[\mathbb{G}_{h,f}(o_h^i; f_h; f_h; v)^2 \right] - \mathbb{E}_{h,f} \left[\mathbb{G}_{h,f}(o_h^i; f_h; f_h; v) \right]^2;$$

By similar derivations as in the proof of Lemma 24, we have

$$\begin{aligned} \mathbb{E}_{s_{h+1}} [X_i(h; f; v) | s_h; a_h] &= \mathbb{E}_{s_{h+1}} \left[\mathbb{E}_{h,f} \left[\mathbb{G}_{h,f}(o_h^i; f_h; f_h; v) | s_h; a_h \right]^2 \right]; \\ \mathbb{E}_{s_{h+1}} (X_i(h; f; v))^2 | s_h; a_h &\leq 4R^2 \mathbb{E}_{s_{h+1}} [X_i(h; f; v) | s_h; a_h]; \end{aligned}$$

Take $Z_t = X_t(h; f; v) - \mathbb{E}_{s_{h+1}} [X_t(h; f; v) | s_h; a_h]$ with $|Z_t| \leq 2R^2$ in Freedman's inequality (F.1) in Lemma 33. Then via the same procedure as in the proof of Lemma 24 we have that for any four-tuple $(t; h; f; v)$, the following holds with probability at least $1 - \delta$:

$$\begin{aligned} \sum_{i=1}^X X_i(h; f; v) &\leq \sum_{i=1}^X \mathbb{E}_{s_{h+1}} [X_i(h; f; v) | s_h; a_h] \\ &\leq O\left(\frac{\sum_{i=1}^X \mathbb{E}_{s_{h+1}} [X_i(h; f; v) | s_h; a_h] \log(1/\delta) + 2R^2 \log(1/\delta) A}{\sum_{i=1}^X \mathbb{E}_{s_{h+1}} [X_i(h; f; v) | s_h; a_h]}\right) \end{aligned}$$

Thus, we have

$$\sum_{i=1}^X X_i(h; f; v) \leq O(R^2 \log(1/\delta));$$

By the same ϵ -covering argument as in the proof of Lemma 24, there exists a covering of \mathcal{L} such that we can take a union bound over \mathcal{L} and have $\sum_{i=1}^t X_i(h; f; v) \leq O(R^2 + Rt)$ where $\epsilon = \log \frac{HTN_L(\cdot)}{\delta}$. Then for $f \in F$ and any $v \in V$, we can use the nearest three-tuple $(f^i; s^i; a^i)$ in the ϵ -covering and conclude that

$$\max_{v \in V} \sum_{i=1}^X \mathbb{E}_{h,f} \left[\mathbb{G}_{h,f}(o_h^i; f_h; f_h; v)^2 \right] - \mathbb{E}_{h,f} \left[\mathbb{G}_{h,f}(o_h^i; f_h; f_h; v) \right]^2 = \max_{v \in V} \sum_{i=1}^X X_i(h; f; v) \leq O(\cdot);$$

This in sum finishes our proof of Lemma 22 with $\epsilon = O(R^2 + Rt)$. \square

F.3 Proof of Lemma 25

Proof of Lemma 25. The proof basically follows Appendix §C of Russo and Van Roy [2013] and Appendix §D of Jin et al. [2021]. We first prove that for all $[T]$,

$$\sum_{k=1}^X \mathbb{1}(jG(f_k; g_k)j > \epsilon) \leq (\epsilon^{-2} + 1) \dim_{FE}(F; G; \cdot); \tag{F.6}$$

Let $m := \sum_{k=1}^X \mathbb{1}(jG(f_k; g_k)j > \epsilon)$, then there exists $s_1; \dots; s_m$ which is a subsequence of $[T]$ such that $G(f_{s_1}; g_{s_1}); \dots; G(f_{s_m}; g_{s_m}) > \epsilon$.

We first show that for the sequence $f_{s_1}; \dots; f_{s_m} \in F$, there exists $j \in [m]$ such that f_{s_j} is ϵ -independent on at least $d(m-1) = \dim_{FE}(F; G; \cdot)$ disjoint sequences $f_{s_1}; \dots; f_{s_{j-1}}$ and $f_{s_{j+1}}; \dots; f_{s_m}$ [Russo and Van Roy, 2013]. We will prove this by following procedure. Starting with singleton sequence $\mathcal{B}_1 = \{f_{s_1}; g_{s_1}\}; \dots; \mathcal{B}_L = \{f_{s_L}; g_{s_L}\}$ and $j = L + 1$. For each j , if f_{s_j} is ϵ -independent on

B_1, \dots, B_L we already achieved our goal and the process stops. Otherwise, there exists i such that f_{s_j} is ϵ -dependent on B_i and update $B_i = B_i \cup \{f_{s_j}, g\}$. Then we add increment j by 1 and continue the process. By the definition of FE dimension, the cardinality of each set B_i cannot be larger than $\dim_{FE}(F; G; \epsilon)$ at any point in this process. Therefore, by pigeonhole principle the process stops by step $\leq L \dim_{FE}(F; G; \epsilon) + 1$.

Therefore, we have proved that there exists such that $|G(f_{s_j}; g_{s_j})| > \epsilon$ and f_{s_j} is ϵ -independent with at least $L = d(m-1) = \dim_{FE}(F; G; \epsilon)$ disjoint sequences $f_{s_1}, \dots, f_{s_j}, \dots, f_{s_{j-1}}, g$. For each of the sequences $f_{s_1}, \dots, f_{s_j}, \dots, f_{s_{j-1}}, g$, by definition of the FE dimension in Definition 3 we have that

$$\sum_{k=1}^L |G(f_k; g_{s_j})|^2 \geq \epsilon^2. \quad (F.7)$$

Summing all of bounds (F.7) for disjoint sequences together we have that

$$\sum_{k=1}^L |G(f_k; g_{s_j})|^2 \geq L \epsilon^2 = d(m-1) \epsilon^2. \quad (F.8)$$

The left hand side of (F.8) can be upper bounded by $\sum_{k=1}^L |G(f_k; g_{s_j})|^2 \leq \sum_{k=1}^L |G(f_k; g_{s_j})|^2$ due to the condition of lemma. Therefore, we have proved that $\sum_{k=1}^L |G(f_k; g_{s_j})|^2 \geq L \epsilon^2$ which completes the proof of (F.6).

Now let $d = \dim_{FE}(F; G; \epsilon)$ and sort $|G(f_1; g_1)|, \dots, |G(f_t; g_t)|$ in a nonincreasing order, denoted by e_1, \dots, e_t . Then we have that

$$\sum_{k=1}^t |G(f_k; g_k)|^2 = \sum_{k=1}^t e_k^2 = \sum_{k=1}^t e_k^2 \mathbb{1}(e_k > \epsilon) + \sum_{i=1}^t e_i^2 \mathbb{1}(e_i > \epsilon) \leq t \epsilon^2 + \sum_{i=1}^t e_i^2 \mathbb{1}(e_i > \epsilon). \quad (F.9)$$

For $k \geq t$, we want to give an upper bound for $\sum_{i=1}^k e_i^2 \mathbb{1}(e_i > \epsilon)$. Assume $e_k > \epsilon$, then for any such that $e_k > \epsilon$, by (F.6), we have that

$$\sum_{i=1}^k |G(f_i; g_i)|^2 \leq (\epsilon^2 + 1) \dim_{FE}(F; G; \epsilon) \leq (\epsilon^2 + 1) d;$$

which implies that $\sum_{i=1}^k |G(f_i; g_i)|^2 \leq (\epsilon^2 + 1) d$. Taking the limit $\epsilon \rightarrow e_k$, we have that $\min_{\epsilon} \sum_{i=1}^k |G(f_i; g_i)|^2 \leq (\epsilon^2 + 1) d$. Finally, we have that

$$\sum_{k=1}^t e_k^2 \mathbb{1}(e_k > \epsilon) \leq \min_{\epsilon} \sum_{i=1}^t |G(f_i; g_i)|^2 \leq C + \sum_{i=d+1}^t \frac{e_i^2}{k-d} \leq \min_{\epsilon} \sum_{i=1}^t |G(f_i; g_i)|^2 \leq C + 2 \int_0^{\epsilon} z^{p-1} dz \leq \min_{\epsilon} \sum_{i=1}^t |G(f_i; g_i)|^2 \leq C + 2 \frac{\epsilon^p}{p}. \quad (F.10)$$

Plugging (F.10) into (F.9) completes the proof. □

F.4 Proof of Lemma 32

Proof of Lemma 32. We assume that $\kappa_1 \leq B$ and treat B as an absolute constant ($\kappa_1 \leq 2$ in Sun et al. [2019]) in the following derivations. For a data set $\mathcal{D}_t = \{r_h^t, s_h^t, a_h^t, s_{h+1}^t, g_{t=1,2,\dots,T}\}$, we first build an auxiliary random variable defined for every $(h; f; v) \in [T] \times [H] \times \mathcal{V}$

$$X_t(h; f; v) := \sum_{\theta} E_{\theta} \left[v(s_h^t; a_h^t; \theta) - v(s_h^t; a_h^t; s_{h+1}^t) \right]^2 = E_{\theta} \left[v(s_h^t; a_h^t; \theta) - v(s_h^t; a_h^t; s_{h+1}^t) \right]^2;$$

where the randomness lies in the sampling of the data set \mathcal{D}_t . We know that $|X_t(h; f; v)| \leq 4B^2$ almost surely. Take conditional expectation E_{θ} with respect to s_h, a_h , we have by definition that

$$E_{s_{h+1}} [X_t(h; f; v) | s_h, a_h] = E_{s_{h+1}} \left[\sum_{\theta} E_{\theta} \left[v(s_h^i; a_h^i; \theta) - v(s_h^i; a_h^i; s_{h+1}^i) \right]^2 \right] = \sum_{\theta} E_{\theta} \left[v(s_h^i; a_h^i; \theta) - v(s_h^i; a_h^i; s_{h+1}^i) \right]^2 | s_h, a_h;$$

Using the fact that $a^2 - b^2 = (a - b)(a + b)$ and $E_{\mathbf{s}_h} f v(s_h^i; a_h^i; \boldsymbol{\theta}) - E_{\mathbf{s}_h} f v(s_h^i; a_h^i; \boldsymbol{\theta})$ is nonrandom given $\mathbf{s}_h; a_h$, we have

$$\begin{aligned} & E_{\mathbf{s}_{h+1}} [X_i(h; f; v) | \mathbf{s}_h; a_h] \\ &= E_{\mathbf{s}_h} f v(s_h^i; a_h^i; \boldsymbol{\theta}) - E_{\mathbf{s}_h} f v(s_h^i; a_h^i; \boldsymbol{\theta}) \\ &\quad E_{\mathbf{s}_{h+1}} E_{\mathbf{s}_h} f v(s_h^i; a_h^i; \boldsymbol{\theta}) + E_{\mathbf{s}_h} f v(s_h^i; a_h^i; \boldsymbol{\theta}) - 2v(s_h^i; a_h^i; s_{h+1}^i) | \mathbf{s}_h; a_h \\ &= E_{\mathbf{s}_h} f v(s_h^i; a_h^i; \boldsymbol{\theta}) - E_{\mathbf{s}_h} f v(s_h^i; a_h^i; \boldsymbol{\theta})^2 : \end{aligned}$$

On the other hand,

$$\begin{aligned} E_{\mathbf{s}_{h+1}} X_i(h; f; v)^2 | \mathbf{s}_h; a_h &= E_{\mathbf{s}_{h+1}} E_{\mathbf{s}_h} f v(s_h^i; a_h^i; \boldsymbol{\theta}) - E_{\mathbf{s}_h} f v(s_h^i; a_h^i; \boldsymbol{\theta}) - 4B^2 | \mathbf{s}_h; a_h \\ &= 16B^2 E_{\mathbf{s}_h} f v(s_h^i; a_h^i; \boldsymbol{\theta}) - E_{\mathbf{s}_h} f v(s_h^i; a_h^i; \boldsymbol{\theta})^2 \\ &\quad 16B^2 E_{\mathbf{s}_{h+1}} [X_i(h; f; v) | \mathbf{s}_h; a_h] : \end{aligned}$$

By taking $Z_t = X_t(h; f; v) - E_{\mathbf{s}_{h+1}} [X_t(h; f; v) | \mathbf{s}_h; a_h]$ with $|Z_t| \leq 8B^2$ a.s. in Freedman's inequality (F.1) in Lemma 33, by the same procedure as in the proof of Lemma 24, we have that for any four-tuple $(t; h; f; v)$, the following holds with probability at least $1 - \delta$:

$$\begin{aligned} & \sum_{i=1}^t X_i(h; f; v) - \sum_{i=1}^t E_{\mathbf{s}_{h+1}} [X_i(h; f; v) | \mathbf{s}_h; a_h] \\ & \leq O\left(\frac{1}{\delta} \sum_{i=1}^t E_{\mathbf{s}_{h+1}} [X_i(h; f; v) | \mathbf{s}_h; a_h] \log\left(\frac{1}{\delta}\right) + 8B^2 \log\left(\frac{1}{\delta}\right) A\right) : \quad (\text{F.11}) \end{aligned}$$

Let \mathcal{M} be a δ -cover of \mathcal{M} and \mathcal{V} a δ -cover of \mathcal{V} . By taking a union bound over $(t; h; f^0; v^0) \in [T] \times [H] \times \mathcal{M} \times \mathcal{V}$, we have with probability at least $1 - \delta$ that the following holds for any $f^0 \in \mathcal{M}, v^0 \in \mathcal{V}$,

$$\begin{aligned} & \sum_{i=1}^t X_i(h; f^0; v^0) - \sum_{i=1}^t E_{\mathbf{s}_{h+1}} [X_i(h; f^0; v^0) | \mathbf{s}_h; a_h] \\ & \leq O\left(\frac{1}{\delta} \sum_{i=1}^t E_{\mathbf{s}_{h+1}} [X_i(h; f^0; v^0) | \mathbf{s}_h; a_h] + 8B^2 A\right) ; \quad (\text{F.12}) \end{aligned}$$

where $A = \log\left(\frac{HTjM}{jVj}\right)$. Further for any t calculated at $2/T$ and any $v^0 \in \mathcal{V}$, we choose $f^0 = \arg \min_{f \in \mathcal{M}} \text{dist}(f, f^0)$ where dist is the distance measure $\text{dist}(f, f^0) = \min_{v \in \mathcal{V}} (v^0, v)$ and conclude

$$\begin{aligned} & \sum_{i=1}^t X_i(h; f^0; v^0) \\ &= \sum_{i=1}^t E_{\mathbf{s}_h} f^0 v^0(s_h^i; a_h^i; \boldsymbol{\theta}) - v^0(s_h^i; a_h^i; s_{h+1}^i)^2 - E_{\mathbf{s}_h} f^0 v^0(s_h^i; a_h^i; \boldsymbol{\theta}) - v^0(s_h^i; a_h^i; s_{h+1}^i)^2 \\ &\quad \sum_{i=1}^t E_{\mathbf{s}_h} f^0 v^0(s_h^i; a_h^i; \boldsymbol{\theta}) - v^0(s_h^i; a_h^i; s_{h+1}^i)^2 - E_{\mathbf{s}_h} f^0 v^0(s_h^i; a_h^i; \boldsymbol{\theta}) - v^0(s_h^i; a_h^i; s_{h+1}^i)^2 + O(Bt) \\ &\stackrel{(i)}{=} O\left(\frac{1}{\delta} + Bt\right) ; \quad (\text{F.13}) \end{aligned}$$

where (i) is due to the constraint of Algorithm 3. Combining (F.12) with (F.13), we derive the following

$$\sum_{i=1}^t E_{\mathbf{s}_{h+1}} [X_i(h; f^0; v^0) | \mathbf{s}_h; a_h] \leq O\left(\frac{1}{\delta} + Bt + B^2\right) :$$

Note that f^0 is chosen as the nearest model in the ϵ -covering of \mathcal{M} and for any f there exists a nearest f^0 in the ϵ -covering of \mathcal{V} , we conclude

$$\max_{v \in \mathcal{V}} \sum_{i=1}^T E_{S_{h+1}} [X_i(h; f^t; v) | S_h; a_h] \leq O(\epsilon + Bt + B^2);$$

Note we also have proved property (ii) in Definition 6 in §E.2, and we apply the global optimality of the discriminator as in the proof of Lemma 24 and obtains

$$\sum_{i=1}^T \max_{v \in \mathcal{V}} E_{S_{h+1}} [X_i(h; f^t; v) | S_h; a_h] \leq O(\epsilon + Bt + B^2);$$

Multiplying $E_{\mathbf{s}} [v(s_h^i; a_h^i; \boldsymbol{\theta}) - E_{\mathbf{s}} [v(s_h^i; a_h^i; \boldsymbol{\theta})]^2]$ by $\frac{1(a_h^i = f(s_h^i))}{1-jA_j}$, taking expectation over $\mathbf{s}_h^i; a_h^i$ and again using the global discriminator optimality, we arrive at

$$\begin{aligned} & \sum_{i=1}^T \max_{v \in \mathcal{V}} E_{S_h^i} [v(s_h^i; a_h^i; \boldsymbol{\theta}) - E_{\mathbf{s}} [v(s_h^i; a_h^i; \boldsymbol{\theta})]^2] \\ &= \sum_{i=1}^T \max_{v \in \mathcal{V}} E_{S_h^i} [v(s_h^i; a_h^i; \boldsymbol{\theta}) - E_{\mathbf{s}} [v(s_h^i; a_h^i; \boldsymbol{\theta})]^2] \\ & \leq O(jA_j + Bt + B^2); \end{aligned}$$

which concludes the proof. \square

F.5 Proof of Lemma 31

Proof of Lemma 31 For a dataset $\mathcal{D}_h = \{r_h^t, s_h^t, a_h^t, s_{h+1}^t, g_{t=1;2;\dots;T}\}$, we first build an auxiliary random variable defined for every $(t; h; f; v) \in [T] \times [H] \times \mathcal{F} \times \mathcal{V}$

$$X_t(h; f; v) := E_{\mathbf{s}} [v(s_h^t; a_h^t; \boldsymbol{\theta}) - v(s_h^t; a_h^t; s_{h+1}^t)]^2 - E_{\mathbf{s}} [v(s_h^t; a_h^t; \boldsymbol{\theta}) - v(s_h^t; a_h^t; s_{h+1}^t)]^2;$$

By Eq. (F.11), with probability at least $1 - \delta$,

$$\begin{aligned} & \sum_{i=1}^T X_i(h; f; v) \leq \sum_{i=1}^T E_{S_{h+1}} [X_i(h; f; v) | S_h; a_h] \\ & \leq O(4B^2 \sum_{i=1}^T E_{S_{h+1}} [X_i(h; f; v) | S_h; a_h] \log(1/\delta) + 8B^2 \log(1/\delta) A); \end{aligned}$$

Let \mathcal{M} be a ϵ -cover of \mathcal{M} and \mathcal{V} a ϵ -cover of \mathcal{V} . By taking a union bound over $(t; h; f^0; v^0) \in [T] \times [H] \times \mathcal{M} \times \mathcal{V}$, we have with probability at least $1 - \delta$ that the following holds for any $f^0 \in \mathcal{M}, v^0 \in \mathcal{V}$,

$$\begin{aligned} & \sum_{i=1}^T X_i(h; f^0; v^0) \leq \sum_{i=1}^T E_{S_{h+1}} [X_i(h; f^0; v^0) | S_h; a_h] \\ & \leq O(4B^2 \sum_{i=1}^T E_{S_{h+1}} [X_i(h; f^0; v^0) | S_h; a_h] + 8B^2 A); \end{aligned}$$

where $\epsilon = \log \frac{HTjM}{jVj}$. Thus, we have

$$\sum_{i=1}^T X_i(h; f^0; v^0) \leq O(B^2);$$

Further for any $f \in F$ and any $v \in V$, we choose $v^0 = \arg \min_{v \in V} \text{dist}(f, v)$ where dist is the distance measure $\text{dist}(f, v) = \min_{v^0 \in V} (v^0, v)$ and have

$$\sum_{i=1}^n X_i(h; f; v) = \sum_{i=1}^n E_{\theta} \left[v(s_h^t; a_h^t; \theta) - v(s_h^t; a_h^t; s_{h+1}^t) \right]^2 \leq \sum_{i=1}^n E_{\theta} \left[v(s_h^t; a_h^t; \theta) - v(s_h^t; a_h^t; s_{h+1}^t) \right]^2$$

Thus,

$$\max_{v \in V} \sum_{i=1}^n E_{\theta} \left[v(s_h^i; a_h^i; \theta) - v(s_h^i; a_h^i; s_{h+1}^i) \right]^2 \leq \inf_{g \in \mathcal{G}} \sum_{i=1}^n E_{\theta} \left[v(s_h^i; a_h^i; \theta) - v(s_h^i; a_h^i; s_{h+1}^i) \right]^2$$

which concludes the proof. \square

G Proof for Functional Eluder Dimension

In the following proposition, we prove that the Bellman eluder (BE) dimension [Jin et al., 2021] is a special case of the FE dimension where $G_h(g; f) := E_{\theta} \left[(g_h - T_h g_{h+1}) \right]$.

Proposition 34. For any hypothesis class \mathcal{F} , taking coupling function G to be the union of G_h : $\mathcal{F}_h \times \mathcal{F}_{h+1} \rightarrow \mathbb{R}$ with each $G_h(g; f) := E_{\theta} \left[(g_h - T_h g_{h+1}) \right]$.

$$\dim_{FE}(\mathcal{F}; G) = \dim_{BE}(\mathcal{F}; \{G_h\})$$

Proof of Proposition 34. By definition of the functional eluder dimension,

$$\dim_{FE}(\mathcal{F}; G) = \max_{h \in [H]} \dim_{FE}(\mathcal{F}; G_h)$$

where $\dim_{FE}(\mathcal{F}; G_h)$ is the length of the longest sequence satisfying for every $t \in [n]$, $\sum_{i=1}^t (G_h(g_i; f_i))^2 \leq \epsilon$ and $|G_h(g_t; f_t)| > \epsilon$. Bringing in $G_h(g; f) := E_{\theta} \left[(g_h - T_h g_{h+1}) \right]$, we have f_1, \dots, f_n is also the longest sequence that satisfies for some θ that

$$\sum_{i=1}^t E_{\theta} \left[(g_{t,h} - T_h g_{t,h+1}) \right]^2 \leq \epsilon, \quad \text{and} \quad E_{\theta} \left[(g_{t,h} - T_h g_{t,h+1}) \right] > \epsilon.$$

Thus, $\dim_{FE}(\mathcal{F}; G_h) \leq n$. Taking maximum over $h \in [H]$, we have

$$\dim_{FE}(\mathcal{F}; G) = \max_{h \in [H]} \dim_{FE}(\mathcal{F}; G_h) \leq \max_{h \in [H]} \dim_{DE}(\mathcal{F}; G_h) = \dim_{BE}(\mathcal{F}; \{G_h\})$$

which concludes our proof. \square

Combining Proposition 34 with Proposition 29 in Jin et al. [2021], it is straightforward to conclude that FE dimension is smaller than the effective dimension. In particular, Proposition 30 says \dim_{FE} is controlled by \dim_{BE} , Proposition 29 in Jin et al. [2021] says \dim_{BE} is controlled by the effective dimension \dim_{eff} , therefore low effective dimension would imply ABC with low FE dimension.

In the following paragraphs and Proposition 35 we prove this conclusion from sketch to grant a better understanding of the FE dimension.

The effective dimension [Jin et al., 2021] (or equivalently, critical information gain [Du et al., 2021]) $d_{\text{eff}}(X; \mathcal{G})$ of a set X is defined as the smallest integer $n > 0$ such that

$$n > e \sup_{x_1, \dots, x_n \in X} \log \det \left(I + \frac{1}{2} \sum_{i=1}^n x_i x_i^\top \right)$$

Remark 5.2 in Du et al. [2021] showed that for finite dimensional setting X with \mathbb{R}^d and $\|x\|_2 \leq B$, $d_{\text{eff}}(X; \mathcal{G}) = \Theta(d)$. Moreover, the effective dimension can be small even for infinite dimensional RKHS case.

In the next proposition, we prove that when the coupling function exhibits a bilinear structure $G(f; g) = \langle W(f); X(g) \rangle_H$ with feature space $\mathcal{X} := \{X(g) \in H : g \in \mathcal{F}\}$ and $\|X(g)\|_H \leq \bar{B}$, the functional eluder dimension in Definition 4 is always less than the effective dimension of

Proposition 35. For any hypothesis class F and coupling function $G(\cdot, \cdot) : F \times F \rightarrow \mathbb{R}$ that can be expressed in bilinear form $\langle hW(f), X(g) \rangle_{\mathcal{H}}$, we have

$$\dim_{\text{FE}}(F, G, \epsilon) \leq d_{\text{eff}}(X, \epsilon/\sqrt{B}).$$

Proof of Proposition 35. The proof basically follows the proof of Proposition 29 in Jin et al. [2021] with modifications specified for the functional eluder dimension. Given a hypothesis class F and a coupling function $G(\cdot, \cdot) : F \times F \rightarrow \mathbb{R}$. Suppose there exists an ϵ' -independent sequence $f_1, \dots, f_n \in F$ such that there exist $g_1, \dots, g_n \in F$,

$$\begin{aligned} \sum_{t=1}^n \frac{1}{\epsilon'} \sum_{i=1}^t (G(g_t, f_i))^2 &\leq \epsilon', \quad t \geq [n], \\ \sum_{i=1}^t |G(g_t, f_i)| &> \epsilon', \quad t \geq [n]. \end{aligned} \quad (\text{G.1})$$

When $G(f, g) := \langle hW(f), X(g) \rangle_{\mathcal{H}}$, the above becomes

$$\begin{aligned} \sum_{t=1}^n \frac{1}{\epsilon'} \sum_{i=1}^t \langle hW(g_t), X(f_i) \rangle_{\mathcal{H}}^2 &\leq \epsilon', \quad t \geq [n], \\ \sum_{i=1}^t \langle hW(g_t), X(f_i) \rangle_{\mathcal{H}} &> \epsilon', \quad t \geq [n]. \end{aligned} \quad (\text{G.2})$$

Defining $K_t = \sum_{i=1}^{t-1} X(f_i)X(f_i)^\top + \frac{\alpha}{B} I$, we have by Eq. (G.2) that $\|kW(g_t)\|_{K_t} \leq \sqrt{2\epsilon'}$. Furthermore,

$$\epsilon' \leq \sum_{i=1}^t \langle hW(g_t), X(f_i) \rangle_{\mathcal{H}} \leq \|kW(g_t)\|_{K_t} \|kX(f_t)\|_{K_t} \leq \sqrt{2\epsilon'} \|kX(f_t)\|_{K_t}.$$

Thus, we have $\|kX(f_t)\|_{K_t} \geq \frac{\epsilon'}{\sqrt{2\epsilon'}} = \frac{\sqrt{\epsilon'}}{2}$ for any $t \geq [n]$. By applying the log-determinant argument, we have

$$\sum_{t=1}^n \log(1 + \|kx_t\|_{K_t}^2) = \log \frac{\det(K_{n+1})}{\det(K_1)} = \log \det \left(I + \frac{B}{\epsilon'^2} \sum_{i=1}^n x_i x_i^\top \right).$$

The above equality implies

$$\frac{1}{2} \leq \min_{t \in [n]} \|kx_t\|_{K_t} \leq \exp \left(\frac{1}{n} \log \det \left(I + \frac{B}{\epsilon'^2} \sum_{i=1}^n x_i x_i^\top \right) \right) \leq 1. \quad (\text{G.3})$$

Taking $n = d_{\text{eff}}(X, \epsilon/\sqrt{B})$ yields

$$\exp \left(\frac{1}{n} \log \det \left(I + \frac{B}{\epsilon'^2} \sum_{i=1}^n x_i x_i^\top \right) \right) \leq \frac{1}{n} \sup_{x_1, \dots, x_n \in \mathcal{X}} \log \det \left(I + \frac{B}{\epsilon'^2} \sum_{i=1}^n x_i x_i^\top \right) \leq e^{-1},$$

which contradicts with the inequality (G.3) and concludes our proof. \square

We now provide the detailed proofs of Lemmas 28, 29 and 30.

G.1 Proof of Lemma 28

Proof of Lemma 28. Taking

$$G_{h;f}(f, g) := (\theta_{h;f} - \theta_h^*)^\top \mathbb{E}_{S_h; a_h \sim g} [\psi(s_h, a_h) + \sum_{s'} \phi(s_h, a_h, s') V_{h+1;g}(s')] = \langle hW_h(f), X_h(g) \rangle_{\mathcal{H}},$$

where $W_h(f) := \theta_{h;f} - \theta_h^*$, $X_h(g) := \mathbb{E}_{S_h; a_h \sim g} [\psi(s_h, a_h) + \sum_{s'} \phi(s_h, a_h, s') V_{h+1;g}(s')]$ in Proposition 35. Properties of the effective dimension yield that the FE dimension of the linear mixture MDP model is $\Theta(d)$. \square

G.2 Proof of Lemma 29

Proof of Lemma 29. Taking $G_{h;f}(f, g) := \langle hW_h(f), X_h(g) \rangle_{\mathcal{H}}$ in Proposition 35, and properties of the effective dimension yields the conclusion that the FE dimension of low Witness rank MDP model is $\Theta(W)$. \square

G.3 Proof of Lemma 30

We first introduce two auxillary lemmas:

Lemma 36. Let random variable $x_i \in \mathbb{R}^d$ and $\mathbb{E}[kx_ik^2] \leq B^2$. Then we have that

$$\frac{1}{n} \log \det \left(I + \frac{1}{\lambda} \sum_{t=0}^{T-1} \mathbb{E}[x_t x_t^\top] \right) \geq \frac{d \log \left(1 + \frac{nB^2}{d} \right)}{n}.$$

Proof. We first have

$$\text{trace} \left(I + \frac{1}{\lambda} \sum_{t=0}^{T-1} \mathbb{E}[x_t x_t^\top] \right) = d + \frac{1}{\lambda} \sum_{t=0}^{T-1} \mathbb{E}[kx_t k^2] \leq d + \frac{nB^2}{\lambda}.$$

Therefore, using the Determinant-Trace inequality, we get the first result,

$$\log \det \left(I + \frac{1}{\lambda} \sum_{t=0}^{T-1} \mathbb{E}[x_t x_t^\top] \right) \geq d \log \frac{\text{trace} \left(I + \frac{1}{\lambda} \sum_{t=0}^{T-1} \mathbb{E}[x_t x_t^\top] \right)}{d} \geq d \log \left(1 + \frac{nB^2}{d\lambda} \right).$$

Dividing n from the both side of the inequality completes the proof. \square

The following lemma is a variant of the well-known Elliptical Potential Lemma [Dani et al., 2008, Srinivas et al., 2009, Abbasi-Yadkori et al., 2011, Agarwal et al., 2020a].

Lemma 37 (Randomized elliptical potential). Consider a sequence of random vectors $\{x_0, \dots, x_{T-1}\}$. Let $\lambda > 0$ and $\Sigma_t = \lambda I + \sum_{i=0}^{t-1} \mathbb{E}[x_i x_i^\top]$, we have that

$$\min_{t \in [T]} \log \left(1 + \mathbb{E}[kx_t k^2] \right) \geq \frac{1}{T} \log \frac{\det(\Sigma_T)}{\det(\lambda I)}.$$

Proof. By definition of Σ_t we have that

$$\log \det(\Sigma_{t+1}) = \log \det(\Sigma_t) + \log \det \left(I + \frac{1}{\Sigma_t} \mathbb{E}[x_t x_t^\top] \Sigma_t \right). \quad (\text{G.4})$$

Denote $\Sigma_t^{-1} \mathbb{E}[x_t x_t^\top] \Sigma_t$ with eigenvalue $\lambda_1, \dots, \lambda_d \geq 0$, we have that

$$\det \left(I + \Sigma_t^{-1} \mathbb{E}[x_t x_t^\top] \Sigma_t \right) = \prod_{i=1}^d (\lambda_i + 1) \geq 1 + \frac{\text{trace}(\Sigma_t^{-1} \mathbb{E}[x_t x_t^\top] \Sigma_t)}{d} = 1 + \mathbb{E}[kx_t k^2]. \quad (\text{G.5})$$

Plugging (G.5) into (G.4) gives that

$$\log \det(\Sigma_{t+1}) \geq \log \det(\Sigma_t) + \log \left(1 + \mathbb{E}[kx_t k^2] \right) \geq \log \det(\Sigma_t) + \min_{t \in [T]} \log \left(1 + \mathbb{E}[kx_t k^2] \right).$$

Taking telescope sum from $t = 0$ to $t = T - 1$ completes the proof. \square

Proof of Lemma 30. Given a hypothesis class F and a coupling function $G(\cdot, \cdot) : F \times F \rightarrow \mathbb{R}$. Let n to be defined as follows,

$$n := \min \{ n \in \mathbb{N} : n \geq \frac{d \log(1 + 4nd_s R^4 / (d \epsilon'^2))}{\epsilon'} \}.$$

Then we have that $n = \Theta(d \log(1 + 4nd_s R^4 / (d \epsilon'^2)))$. We will prove $\dim_{FE}(F, G, \epsilon) \leq n$ by contradiction. Suppose that $\dim_{FE}(F, G, \epsilon) > n$, there exists an ϵ' -independent (where $\epsilon' = \epsilon$) sequence $f_1, \dots, f_n \in F$ such that there exist $g_1, \dots, g_n \in F$,

$$\begin{aligned} \sum_{t=1}^n \mathbb{E} \left[\sum_{i=1}^n (G(g_t, f_i))^2 \right] &\leq \epsilon', \quad t \in [n], \\ \mathbb{E} \left[\sum_{i=1}^n |G(g_t, f_i)| \right] &> \epsilon', \quad t \in [n]. \end{aligned} \quad (\text{G.6})$$

Recall that the ABC function of KNR model is defined as,

$$\begin{aligned} G_{h;f}(f, g) &= \frac{1}{\mathbb{E}_{S_h; a_h \sim g} \text{vec}((U_{h;f} \ U_h^*)^\top (U_{h;f} \ U_h^*)), \text{vec} \ E_{S_h; a_h \sim g} \phi(s_h, a_h) \phi(s_h, a_h)^\top} \\ &= \frac{1}{\mathbb{E}_{S_h; a_h \sim g} k(U_{h;f} \ U_h^*) \phi(s_h, a_h) k_2^2}. \end{aligned}$$

Therefore, condition (G.6) can be reduced to

$$\begin{aligned} \frac{1}{\mathbb{E}_{S_h; a_h \sim f_t} k(U_{h;g_t} \ U_h^*) \phi(s_h, a_h) k_2^2} &\leq \epsilon', \quad t \geq [n], \\ \frac{1}{\mathbb{E}_{S_h; a_h \sim f_t} k(U_{h;g_t} \ U_h^*) \phi(s_h, a_h) k_2^2} &> \epsilon', \quad t \geq [n]. \end{aligned} \quad (\text{G.7})$$

Denote $U_{h;g_t;j}, j \geq [d_s]$ and $U_{h;j}^*, j \geq [d_s]$ to be the rows of $U_{h;g_t}$ and U_h^* . Taking square over both side of the inequalities in (G.7) gives that

$$\begin{aligned} \frac{1}{\mathbb{E}_{S_h; a_h \sim f_t} [(U_{h;g_t;j} \ U_{h;j}^*) \phi(s_h, a_h)]^2} &\leq \epsilon'^2, \quad t \geq [n], \\ \frac{1}{\mathbb{E}_{S_h; a_h \sim f_t} [(U_{h;g_t;j} \ U_{h;j}^*) \phi(s_h, a_h)]^2} &> \epsilon'^2, \quad t \geq [n]. \end{aligned} \quad (\text{G.8})$$

Define $t = \sum_{i=1}^{t-1} \mathbb{E}_{S_h; a_h \sim f_i} [\phi(s_h, a_h) \phi(s_h, a_h)^\top] + (\epsilon'^2/4d_s R^2) \ I$. Then by (G.8), we have that

$$\begin{aligned} &\sum_{j=1}^{d_s} k(U_{h;g_t;j} \ U_{h;j}^*) k_2^2 \\ &= \sum_{i=1}^{t-1} \sum_{j=1}^{d_s} \mathbb{E}_{S_h; a_h \sim f_i} [(U_{h;g_t;j} \ U_{h;j}^*) \phi(s_h, a_h)]^2 + (\epsilon'^2/4d_s R^2) \sum_{j=1}^{d_s} \mathbb{E}_{S_h; a_h \sim f_t} k(U_{h;g_t;j} \ U_{h;j}^*) k_2^2 \\ &\leq \sum_{i=1}^{t-1} \sum_{j=1}^{d_s} \mathbb{E}_{S_h; a_h \sim f_i} [(U_{h;g_t;j} \ U_{h;j}^*) \phi(s_h, a_h)]^2 + (\epsilon'^2/4d_s R^2) \ 2d_s \max_j k(U_{h;g_t;j}) k_2^2 + 2d_s \max_j k(U_{h;j}^*) k_2^2 \\ &\leq \epsilon'^2; \end{aligned}$$

where the first equality is by the Cauchy-Schwartz inequality and the last inequality is by $k(U_{h;g_t;j}) k_2 \leq R, k(U_{h;j}^*) k_2 \leq R$. Furthermore we have that

$$\begin{aligned} &\epsilon'^2 \sum_{j=1}^{d_s} \mathbb{E}_{S_h; a_h \sim f_t} [(U_{h;g_t;j} \ U_{h;j}^*) \phi(s_h, a_h)]^2 \\ &= \sum_{j=1}^{d_s} \mathbb{E}_{S_h; a_h \sim f_t} [(U_{h;g_t;j} \ U_{h;j}^*) \ \frac{1}{t} \ \frac{1}{t} \phi(s_h, a_h)]^2 \\ &\leq \sum_{j=1}^{d_s} \mathbb{E}_{S_h; a_h \sim f_t} k(U_{h;g_t;j} \ U_{h;j}^*) \ \frac{1}{t} k_2^2 \ \mathbb{E}_{S_h; a_h \sim f_t} k \ \frac{1}{t} \phi(s_h, a_h) k_2^2 \\ &= \mathbb{E}_{S_h; a_h \sim f_t} k \ \frac{1}{t} \phi(s_h, a_h) k_2^2 \ \sum_{j=1}^{d_s} k(U_{h;g_t;j} \ U_{h;j}^*) k_2^2, \end{aligned}$$

where the last inequality is by the Cauchy-Schwartz inequality for random variables. Thus, we have that $\mathbb{E}_{S_h; a_h \sim f_t} k \ \frac{1}{t} \phi(s_h, a_h) k_2^2 \leq 1/2$ for all $t \geq [n]$. By applying Lemma 37, we have that

$$\begin{aligned} &\min_{t \in [n]} \log(1 + \mathbb{E}_{S_h; a_h \sim f_t} k \ \frac{1}{t} \phi(s_h, a_h) k_2^2) \\ &\frac{1}{n} \log \frac{\det(n+1)}{\det(1)} = \frac{1}{n} \log \det \left(1 + \frac{4d_s R^2}{\epsilon'^2} \sum_{i=1}^{n-1} \mathbb{E}_{S_h; a_h \sim f_i} [\phi(s_h, a_h) \phi(s_h, a_h)^\top] \right). \end{aligned}$$

The above equation further implies that

$$\frac{1}{n} \log \det \left(\mathbf{I} + \frac{4d_s R^2}{\epsilon'^2} \sum_{i=1}^n \mathbb{E}_{s_h; a_h \sim \pi_i} [\phi(s_h, a_h) \phi(s_h, a_h)^\top] \right) \leq \min_{t \in [n]} \log \left(\mathbf{I} + \mathbb{E}_{s_h; a_h \sim \pi_t} k_t^{-1=2} \phi(s_h, a_h) k_t^2 \right) \leq \log(3/2).$$

On the other hand, Lemma 36 implies that

$$\frac{1}{n} \log \det \left(\mathbf{I} + \frac{4d_s R^2}{\epsilon'^2} \sum_{i=1}^n \mathbb{E}_{s_h; a_h \sim \pi_i} [\phi(s_h, a_h) \phi(s_h, a_h)^\top] \right) \geq \frac{d \log \left(\mathbf{I} + \frac{4nd_s R^4}{d \epsilon'^2} \right)}{n} \geq e^{-1}.$$

This leads to a contradiction because $\epsilon' \leq \epsilon$ and $\log(3/2) > e^{-1}$. We complete the proof of $\dim_{FE}(F, G, \epsilon) = \Theta(d)$. \square

H Experiment

In this section, we carry out experiments to evaluate the empirical performance of our algorithm OPERA for linear mixture MDPs (Algorithm 2). In this experiment, we construct an MDP M with dimension $d = 3$ and episode length $H = 5$. The state space S consists of $H + 2$ different states x_1, \dots, x_{H+2} and the action space $A = \{1, \dots, 2^{d-1}\}$ consists of 2^{d-1} different actions. For each step $h \geq [H]$ and episode $k \geq [K]$, we assume that the reward function $r(s, a)$ is known (so no need to introduce ψ as in Section E.1). In particular, for all $1 \leq h \leq H$, the reward function $r_h(s, a) = 1$ if and only if $s = x_{H+2}$ and $r_h(s, a) = 0$ otherwise. For each step $h \geq [H]$ and corresponding transition probability function \mathbb{P}_h , x_{H+1} and x_{H+2} are absorbing states. For other states $x_h (1 \leq h \leq H)$, the transition probability satisfies that

$$\begin{aligned} \mathbb{P}_h(x_{h+1} | x_h, \mathbf{a}) &= 0.95 + \theta_{0.01} \mathbf{1}_{d-1, \mathbf{a}}, \\ \mathbb{P}_h(x_{H+2} | x_h, \mathbf{a}) &= 0.05 + \theta_{0.01} \mathbf{1}_{d-1, \mathbf{a}}, \end{aligned}$$

where each $\mathbf{1}_{d-1}$ is a $(d - 1)$ -dimensional vector of all ones, $\mathbf{a} \in A$ is also a $(d - 1)$ -dimensional vector. Then we have that $\mathbb{P}_h(x_{h'} | x_h, \mathbf{a}) = \theta_{h'}^* \phi(x_{h'}, \mathbf{a}, x_h)$ where $\theta_{h'}^* = [0.01 \mathbf{1}_{d-1}, 1]$, and $\phi(x_{h'}, \mathbf{a}, x_h)$ are as follows,

- $\phi(x_{h'}, \mathbf{a}, x_h) = [\mathbf{a}, 0.95]$ if $1 \leq h \leq H$ and $h' = h + 1$,
- $\phi(x_{h'}, \mathbf{a}, x_h) = [\mathbf{a}, 0.05]$ if $1 \leq h \leq H$ and $h' = H + 2$,
- $\phi(x_{h'}, \mathbf{a}, x_h) = [\mathbf{0}_{d-1}, 1]$ if $h' = h = H + 1$,
- $\phi(x_{h'}, \mathbf{a}, x_h) = [\mathbf{0}_{d-1}, 1]$ if $h' = h = H + 2$,
- $\phi(x_{h'}, \mathbf{a}, x_h) = \mathbf{0}_d$, otherwise.

Here $\mathbf{0}_{d-1}$ is a $(d - 1)$ -dimensional vector of all zeros. We compare our algorithm OPERA with the following two baselines: *Optimal* (optimal policy) and *Random* (uniformly random policy which chooses actions uniformly from A). For numerical stability, we add λI to $\frac{\partial}{\partial \theta}$ with $\lambda = 1$ and use CVX [Diamond and Boyd, 2016, Agrawal et al., 2018] to approximately solve (E.3) when we implement Algorithm 2. The cumulative rewards of different algorithms averaged over 10 runs for the first 10000 episodes are plotted in Figure 2. We can see that OPERA performs much better than the random policy and can converge to the optimal policy after 10000 episodes.

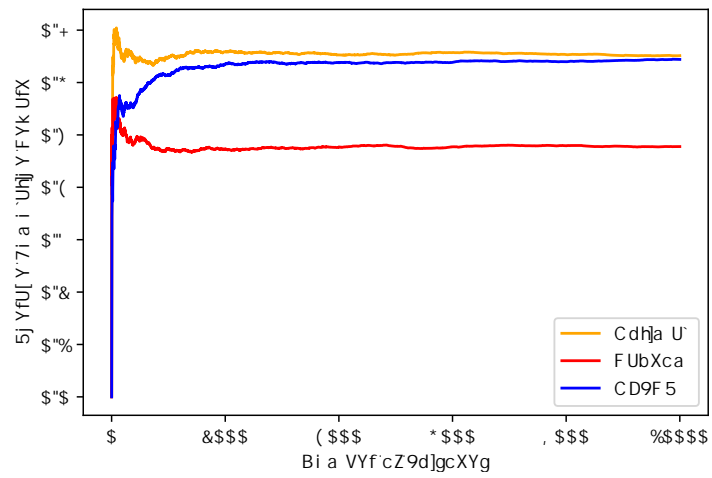


Figure 2: Cumulative regret comparison in the first 10000 episodes of different RL algorithms (i.e., OPERA, Optimal Policy, Random Policy). Results are averaged over 10 runs.