
Low-rank MDPs with Continuous Action Spaces

Andrew Bennett[†]
Morgan Stanley

Nathan Kallus[†]
Cornell Tech

Miruna Oprescu[†]
Cornell Tech

Abstract

Low-Rank Markov Decision Processes (MDPs) have recently emerged as a promising framework within the domain of reinforcement learning (RL), as they allow for provably approximately correct (PAC) learning guarantees while also incorporating ML algorithms for representation learning. However, current methods for low-rank MDPs are limited in that they only consider finite action spaces, and give vacuous bounds as $|\mathcal{A}| \rightarrow \infty$, which greatly limits their applicability. In this work, we study the problem of extending such methods to settings with continuous actions, and explore multiple concrete approaches for performing this extension. As a case study, we consider the seminal FLAMBE algorithm (Agarwal et al., 2020), which is a reward-agnostic method for PAC RL with low-rank MDPs. We show that, without any modifications to the algorithm, we obtain a similar PAC bound when actions are allowed to be continuous. Specifically, when the model for transition functions satisfies a Hölder smoothness condition w.r.t. actions, and either the policy class has a uniformly bounded minimum density or the reward function is also Hölder smooth, we obtain a polynomial PAC bound that depends on the order of smoothness.

1 INTRODUCTION

Reinforcement learning (RL) is an important framework for the study of automated decision making, where decisions are sequentially made over time with

context that evolves based on previous decisions. While RL gained popularity through large scale game-playing, this problem has many diverse practical applications, such as memory management for computing systems (Chinchali et al., 2018), online resource allocation in computer clusters (Lykouris and Vassilvtiskii, 2018; Tessler et al., 2021), personalized recommendations (Zheng et al., 2018), bidding/advertising (Cai et al., 2017), autonomous driving (Sallab et al., 2017), robotics (Kober et al., 2013), and automated trading (Li et al., 2019).

It is widely recognized that achieving provably correct RL for general Markov Decision Processes (MDPs) with arbitrary state spaces incurs a sample complexity of either $\Omega(|\mathcal{S}|)$ or $\Omega(|\mathcal{A}|^H)$, where $|\mathcal{S}|$, $|\mathcal{A}|$, and H are the sizes of the state space, action space, and horizon respectively (Krishnamurthy et al., 2016). This is problematic, because many of the most promising RL applications involve both long horizons and complex observations such as images or text, so both $|\mathcal{S}|$ and $|\mathcal{A}|^H$ may be intractably large. Fortunately, there has been extensive recent work on discovering more restricted MDP classes that are useful for modeling realistic problems with rich state spaces while still allowing for PAC learning results with polynomial sample complexity (Jiang et al., 2017; Du et al., 2019; Sun et al., 2019; Jin et al., 2020; Agarwal et al., 2020; Misra et al., 2020; Du et al., 2021; Chen et al., 2022).

A framework of particular interest is low-rank MDPs (Jiang et al., 2017), since they are extremely general and subsume many other such frameworks (e.g., linear MDPs, Jin et al., 2020, or block MDPs, Du et al., 2019). Unlike some of the more general frameworks, such as bounded Bellman-rank MDPs (Jiang et al., 2017), low-rank MDPs allow for the development of efficient and practical algorithms. This framework utilizes a low-rank structure for the MDP transitions based on state and action features, which can be estimated using conventional machine learning techniques, rendering them particularly suited for addressing challenges in environments with rich and complex action spaces. Despite the rich body of recent research on PAC algorithms for low-rank MDPs with polynomial

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s). [†]Alphabetical order.

sample complexity, the complexity invariably scales polynomially with $|\mathcal{A}|$. This is problematic, because many of the real world applications that motivate this line of work feature large or infinite action spaces; for example, actions in autonomous driving have continuous parts (*e.g.*, steering wheel angle), and actions in algorithmic trading have (roughly) continuous parts (*e.g.*, quantity of security to buy/sell). Furthermore, the low-rank RL framework allows for flexible modeling of generic action spaces via the state/action embedding function, so it seems natural that it should be able to tractably handle infinitely large action spaces.

In this paper, we study the problem of extending low-rank MDP PAC results to settings with continuous action spaces, where $|\mathcal{A}| = \infty$. We first discuss why the dependence on $|\mathcal{A}|$ occurs with these algorithms, and how this dependence may be alleviated. In particular, we provide a novel result motivated by Sobolev interpolation theory that can be used to improve existing analyses when certain problem parameters are sufficiently smooth. Then, as a specific case study, we consider the FLAMBE algorithm of Agarwal et al. (2020), which was the first tractable (computationally oracle efficient) PAC RL algorithm for low-rank MDPs, and allows for reward-agnostic learning. We discuss how our methods can be applied to the analysis of FLAMBE, in order to provide PAC bounds for continuous action spaces. These bounds apply when the transition operators allowed by the model are smooth w.r.t. actions, and either the policy class has a uniform bounded maximum density or the allowed reward functions are Hölder-continuous in actions. In addition, we provide a similar analysis for the more recent RAFFLE algorithm of Cheng et al. (2023), which was shown to have near-optimal sample complexity, in the appendix. Aside from the specific continuous action PAC bounds established for FLAMBE and RAFFLE, we hope this work will serve as a valuable template for generalizing low-rank MDP PAC algorithms to continuous action spaces more broadly.

1.1 Related Work

Low-rank MDPs Low-rank MDPs are a class of models at the interplay of representation learning and exploration in reinforcement learning. In recent years, these models have garnered a lot of interest due to efforts to provide algorithms that can leverage machine learning while also having provable guarantees on *e.g.* correctness and sample complexity (Jin et al., 2020; Yang and Wang, 2019, 2020; Agarwal et al., 2020; Modi et al., 2021; Uehara et al., 2021; Chen et al., 2022; Cheng et al., 2023). The seminal oracle-efficient algorithm in this space is FLAMBE (Agarwal et al., 2020), which performs reward-free exploration to learn a low-

rank model of the MDP. Subsequent works by have developed other algorithms improving on this in various ways, such as establishing tighter bounds for exploring with a particular reward function (Uehara et al., 2021), performing both model-free and reward-free exploration (Modi et al., 2021), extending to MDPs with low Bellman Eluder dimension (Chen et al., 2022), or using exploration-driven proxy rewards (Cheng et al., 2023). In addition to oracle-efficient algorithms for low-rank MDPs, there is past work on (computationally inefficient) algorithms for more general bounded Bellman/Witness Rank classes (Jiang et al., 2017; Sun et al., 2019; Du et al., 2021), or for the more restricted block MDP class (Du et al., 2019; Misra et al., 2020). However, these works generally do not allow for continuous action spaces.

One exception to the above is Agarwal and Zhang (2022), which considers an even more general setting (linear embeddability), and allows for arbitrary action spaces. However, this approach has two important drawbacks. First, as is common in the literature that considers similarly general settings such as bounded Bellman rank, the algorithms presented are at face value computationally intractable; in the case of Agarwal and Zhang (2022), their proposed methodology requires posterior sampling over all possible Q functions. In contrast, our paper follows the line of work on the more restricted low-rank MDP setting, for which tractable algorithms (given an oracle for performing MLE) can be designed. Second, given the lower bound in Cheng et al. (2023) for low-rank MDPs that depends on $|\mathcal{A}|$, it is clear that the dependence on action space size *cannot* be avoided in general without additional assumptions, such as smoothness. Thus, the assumptions in Agarwal and Zhang (2022) implicitly constrain the low-rank MDP setting in some other way.

Continuous action spaces in RL Some efforts have been made to adapt PAC RL algorithms designed for discrete action spaces to continuous ones via the discretization of the action space, as seen in the work of Song and Sun (2019); Domingues et al. (2020); Sinclair et al. (2021). However, discretization has several drawbacks, such as computational inefficiency and exponential performance degradation in high dimensions, due to its limited capacity to exploit function smoothness beyond Lipschitz. In contrast, our approach can leverage any degree of smoothness to counteract the curse of dimension and, as we will see, our bounds approach those for finite action spaces when the smoothness of the transition function grows. This cannot be achieved by discretization. To the best of our knowledge, there has been no other effort in directly adapting such RL algorithms to continuous action spaces without discretizing the action space.

Leveraging Smoothness in Machine Learning
 Smooth functions (continuous functions that admit higher order derivatives) have become popular in machine learning, since improved learning rates can often be achieved for sufficiently smooth classes. For example, it is well known that learning d -dimensional, α -times differentiable functions have a minimax-optimal convergence rate of $O(n^{-2\alpha/(2\alpha+d)})$ (Stone, 1982). Many algorithms can leverage such smoothness, including kernel methods, Gaussian processes, and deep neural networks (Giné and Nickl, 2021). While smooth functions have also found some practical application in reinforcement learning, particularly when considering smooth rewards in multi-armed bandit ($H = 1$) problems (Liu et al., 2021; Jia et al., 2023; Hu et al., 2022), its application to reinforcement learning more broadly remains a promising area for further exploration.

1.2 Overview of Paper

First, in Section 2 we formally introduce the low-rank MDP setting as in Jiang et al. (2017), and provide a brief overview of existing PAC results for this setting. Then, in Section 3 we discuss why the $|\mathcal{A}|$ dependence tends to occur for these algorithms and corresponding PAC bounds, and in Section 4 we provide some theoretical results and discussion that are helpful for extending these results to continuous \mathcal{A} . Expanding upon this, in Section 5 we consider as a case study the FLAMBE algorithm of Agarwal et al. (2020), and we apply our results from Section 4 to obtain PAC bounds for FLAMBE that hold when \mathcal{A} is continuous, *without any modifications to the algorithm*. We also provide a similar analysis for the more recent RAFFLE algorithm of Cheng et al. (2023) in Appendix B. Finally, we draw some conclusions in Section 6. We relegate all proofs to Appendix A.

1.3 Notation

We let \mathbb{N} denote the space of non-negative integers $\{0, 1, 2, \dots\}$. We let \mathbb{N}^+ and \mathbb{R}^+ denote the space of positive integers and real numbers, respectively (i.e., $\mathbb{N} \setminus \mathbb{N}^+ = \{0\}$). For any $k \in \mathbb{N}^+$, we let $[k] = \{0, 1, \dots, k-1\}$. For any $s \in \mathbb{R}^+$, we let $\lfloor s \rfloor = \{i \in \mathbb{N} : i < s\}$ be the largest integer strictly smaller than s (note $\lfloor 5 \rfloor = 4$ unlike the usual floor function) and $\{s\} = s - \lfloor s \rfloor$ be the remainder; that is, we have the unique decomposition $s = \lfloor s \rfloor + \{s\}$, where $\lfloor s \rfloor \in \mathbb{N}$ and $\{s\} \in (0, 1]$. For any $f : \mathbb{R}^d \rightarrow \mathbb{R}$ and $\beta \in \mathbb{N}^d$, we define $D^\beta f = \frac{\partial^{|\beta|}}{\partial x_1^{\beta_1} \partial x_2^{\beta_2} \dots \partial x_d^{\beta_d}} f(x)$

whenever this derivative exists, where $|\beta| := \sum_{i=1}^d \beta_i$. Whenever all derivatives of order j exist, we let $\nabla^j f$ denote the vector-valued function of all $D^\beta f$ such that $|\beta| = j$, and we let $\nabla^0 f = f$. Finally, for any

$p \in [1, \infty]$, positive integer d , and Borel measurable $\mathcal{D} \subseteq \mathbb{R}^d$, we let $\|\cdot\|_{L_p(\mathcal{D})}$ denote the standard L_p norm w.r.t. the Borel measure on \mathcal{D} . We abbreviate $\|\cdot\|_{L_\infty(\mathcal{D})}$ (i.e., the essential supremum of the absolute value) by $\|\cdot\|_\infty$.

2 LOW-RANK MDP SETTING

We consider a finite-horizon episodic MDP defined by a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, T, H, \rho)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, $T = (T_0, \dots, T_{H-1})$, where $T_h : \mathcal{S} \times \mathcal{A} \mapsto \Delta(\mathcal{S})$ is the transition function at time h , H is the finite time horizon, and $\rho \in \Delta(\mathcal{S})$ is the initial state distribution. We assume that \mathcal{S} may be some generic measurable set (with measure μ), such that $T(\cdot | s, a)$ is continuous w.r.t. μ for every (s, a) .¹ We note that this MDP definition does *not* include reward functions $R_h : \mathcal{S} \times \mathcal{A} \mapsto [0, 1]$ for each $h \in [H]$, since we will consider reward-agnostic algorithms. In addition, past definitions of this setting such as in Jiang et al. (2017) often dictate that $\mathcal{A} = [K]$ for some finite $K \in \mathbb{N}^+$, but we do not impose this restriction.

The *low-rank* assumption is that the transition operator T has a low-rank representation, and can be summarized by the following definition leveraged from Agarwal et al. (2020):

Definition 1 (Low-rank MDP). *An operator $T : \mathcal{S} \times \mathcal{A} \mapsto \Delta(\mathcal{S})$ admits a low-rank decomposition with dimension $d \in \mathbb{N}$ if there exists two embedding functions $\phi^* : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}^d$ and $\psi^* : \mathcal{S} \mapsto \mathbb{R}^d$ such that*

$$T(s' | s, a) = \phi^*(s, a)^\top \psi^*(s') \quad \forall s, s' \in \mathcal{S},$$

where the embeddings ϕ^* and ψ^* satisfy the (without loss of generality) normalization conditions $\|\phi^*(s, a)\|_2 \leq 1 \forall s, a$, and $\|\int \psi^*(s)g(s)d\mu(s)\|_2 \leq \sqrt{d} \forall g$ such that $g : \mathcal{S} \mapsto [0, 1]$. An MDP \mathcal{M} is a low-rank MDP with rank d if T_h admits a low-rank decomposition with dimension d for every $h \in [H]$. We let ϕ_h^* and ψ_h^* denote the true embeddings for the decomposition of T_h .

For learning within a low-rank MDP structure, we assume that we are given a set of hypothesis embeddings $\Phi \subseteq \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}^d$ and $\Psi \subseteq \mathcal{S} \mapsto \mathbb{R}^d$ such that we have the following *realizability* assumption:

Assumption 1 (Realizability). *For every $h \in [H]$, we have $\phi_h^* \in \Phi$ and $\psi_h^* \in \Psi$.*

¹That is, we can interpret the function $T(s' | s, a) \in \mathbb{R}^+$ as the conditional density of s' given (s, a) , with respect to base measure μ , which may be, e.g., a Borel measure (if \mathcal{S} is continuous) or a counting measure (if \mathcal{S} is discrete).

3 LIMITATIONS OF EXISTING LOW-RANK MDP RESULTS

As highlighted in Section 1.1, a range of recent algorithms offer PAC guarantees for low-rank MDPs. In general, all of these algorithms have sample complexity that is $\tilde{\mathcal{O}}(\text{Poly}(d, H, |\mathcal{A}|, \epsilon, \log(1/\delta)))$ for learning ϵ -optimal policies with probability at least $1 - \delta$. Additionally, in every instance, these bounds exhibit a super-linear dependence on $|\mathcal{A}|$, leading to significantly poor scalability when $|\mathcal{A}|$ becomes combinatorially large or infinite.

Fundamentally, the dependence on $|\mathcal{A}|$ arises because these algorithms tend to employ some form of uniform exploration, with actions being sampled uniformly at random at least sometimes. In the linear MDP setting, where the ϕ_h^* embeddings are known, it is possible to bypass uniform exploration by exploring strategically within the embedding space structure (Jin et al., 2020), thus eliminating the sample complexity’s dependence on $|\mathcal{A}|$. However, low-rank MDPs present the added complexity of needing to learn these embeddings. For example, low-rank MDP papers leverage MLE bounds (e.g., see Agarwal et al., 2020, Theorem 21) to learn embeddings that are accurate across the distribution of state-action pairs (s, a) that the training data is sampled from: if we ensure that such error is small when a is sampled uniformly at random for a given state s , then we can ensure that it is small for *all* policies applied at the same s , based on the following importance sampling lemma.

Lemma 1. *For any given $f : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^+$, any distribution ρ over states, and any policy π , we have*

$$\mathbb{E}_{s \sim \rho, a \sim \pi(\cdot | s)}[f(s, a)] \leq |\mathcal{A}| \mathbb{E}_{s \sim \rho, a \sim \text{unif}_{\mathcal{A}}}[f(s, a)].$$

This lemma is trivial to prove by noting that $\pi(a | s) / \text{unif}_{\mathcal{A}}(a) \leq 1 / \text{unif}_{\mathcal{A}}(a) = |\mathcal{A}|$ for all discrete sets \mathcal{A} . That is, the importance sampling lemma, which is implicitly used in the analysis of most of these papers, allows us to bound the errors of learned embeddings (or the probabilities of other “bad” events) under any action distribution, given that we perform some exploration uniformly at random.

Unfortunately, the implicit use of Lemma 1 in the theoretical analysis of these works introduces $|\mathcal{A}|$ factors within error bounds, which translates into polynomial dependence on $|\mathcal{A}|$ in the sample complexity.

4 ERROR BOUNDS WITH CONTINUOUS ACTIONS

Given the limitations discussed in Section 3, a natural question to ask is whether we can provide bounds like

Lemma 1 when \mathcal{A} is continuous. In order to make this discussion concrete, we will consider the continuous action space $\mathcal{A} = [0, 1]^m \subset \mathbb{R}^m$ for some $m \in \mathbb{N}^+$, which is equipped with the standard Borel measure for the Euclidean space \mathbb{R}^m . Although this choice of continuous action space may at a first glance seem restrictive, it is a natural way to parameterize many kinds of continuous actions that have m independent degrees of freedom, and the action space can typically be written this way, *e.g.*, by a simple change of variables. Note as well that that this assumed normalization of action components to be on the scale $[0, 1]$ is important for interpreting smoothness of functions w.r.t. a ; for example, if the “natural” scale of actions for a given problem were $[0, A_{\max}]$, for some $A_{\max} \gg 1$, then we would need to scale up all order- α derivatives w.r.t. actions by a factor of A_{\max}^α in order to normalize and be scale invariant. Additionally, our results easily extend to more general \mathcal{A} with Lipschitz boundaries, as the main property of $[0, 1]^m$ used in the proofs of the lemmas and theorem below are that it is Lipschitz. See our proofs in the appendix for more detail.

In the next sections, we explore two strategies to address the challenges outlined. In our case study examining the FLAMBE algorithm of Agarwal et al. (2020) in Section 5, we will use both approaches in different parts of the analysis.

4.1 Utilizing Smoothness of Error Functions

In many cases where we need to apply Lemma 1, the function $f(s, a)$ corresponds to an “error” term associated with state-action pairs. For many problems, we may be able to justify that such errors are smooth with respect to a given, *e.g.*, assumptions on the smoothness of transitions in terms of a . Specifically, we will provide a generalization of Lemma 1 that can be applied to a fairly general class of smooth functions.

Definition 2 (α -smooth functions). *For any $\alpha \in (0, \infty)$ and $\mathcal{D} \subseteq \mathbb{R}^m$, $f : \mathcal{D} \rightarrow \mathbb{R}$ is an α -smooth function if the $\lfloor \alpha \rfloor$ -order mixed derivatives of f exist, and there exists a constant L such that: (1) $\|\nabla^j f\|_\infty \leq L$ for all $j \leq \lfloor \alpha \rfloor$; and (2) all $\lfloor \alpha \rfloor$ -order derivatives are $\{\alpha\}$ -Hölder continuous with norm at most L , i.e.,*

$$\sup_{a, a' \in \mathcal{D}, |\beta| = \lfloor \alpha \rfloor} \frac{|D^\beta f(a) - D^\beta f(a')|}{\|a - a'\|_2^{\{\alpha\}}} \leq L.$$

That is, α -smooth functions have mixed derivatives up to the order $\lfloor \alpha \rfloor$ that are bounded. Moreover, these mixed derivatives have bounded $\{\alpha\}$ fractional derivatives. For example, when $\alpha \in (0, 1]$, this definition corresponds to α -Hölder continuity (Leoni, 2017); and, in particular, when $\alpha = 1$, this definition corresponds to Lipschitz continuity. Furthermore, for any

$\alpha \in (0, \infty)$ the set of α -smooth functions corresponds to the α -Hölder space (Tsybakov, 2008). However, implicit in our definition is a particular norm, which we explicitly define below.

Definition 3 (α -smoothness norm). *For any α -smooth function f , we define $\|f\|_{C^\alpha}$ as the minimum $L \geq 0$ such that the conditions of Definition 2 hold.*

For example, $\|f\|_{C^\alpha}$ corresponds to the Hölder-continuity norm for all $\alpha \leq 1$, and $\|f\|_{C^1}$ corresponds to Lipschitz norm. More generally, for any $\alpha \in (0, \infty)$ we have that $\|\cdot\|_{C^\alpha}$ is a norm on the α -Hölder space. We also note that these definitions are clearly nested, with $\|f\|_{C^\alpha} \leq \|f\|_{C^\beta}$ for any $\alpha < \beta$.

Given this smoothness definition, we can provide a weaker version of Lemma 1 under the condition that f is α -smooth in a for each s . This bound is based on the following theorem, which bounds the maximum of α -smooth functions by some power of their mean.

Theorem 2 (Uniform bound on α -smooth functions). *Let $f : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^+$ be α -smooth in the action a such that $\|f\|_{C^\alpha} \leq L$ for any $s \in \mathcal{S}$ and some $\alpha \in (0, \infty)$. Then, for any distribution ρ over states, and any policy π , we have*

$$\begin{aligned} & \mathbb{E}_{s \sim \rho, a \sim \pi(\cdot | s)}[f(s, a)] \\ & \leq cL^{\frac{m}{m+\alpha}} \left(\mathbb{E}_{s \sim \rho, a \sim \text{unif}_{\mathcal{A}}}[f(s, a)] \right)^{\frac{\alpha}{m+\alpha}}, \end{aligned}$$

for some constant c that depends only on m and α .

The proof of Theorem 2 is based on a result that allows us to bound the maximum value of any non-negative, α -smooth function on a Lipschitz subset of \mathbb{R}^d by its L_1 norm on that domain to the power of $\alpha/(m+\alpha)$; this follows from Sobolev interpolation theory (see e.g., Fiorenza et al., 2021) and may be of independent interest. Full proof details are provided in the appendix.

In terms of exponents on the mean, Theorem 2 gives a weaker version of the inequality in Lemma 1, and the two exponents coincide only as $\alpha \rightarrow \infty$. Often, this more relaxed bound can substitute Lemma 1, trading $|\mathcal{A}|$ dependency for less favorable exponents in downstream sample complexity calculations.

4.2 Using Smoothed Policies

An alternative approach to extending Lemma 1 to continuous \mathcal{A} is to restrict attention to policies with uniformly bounded density. Specifically, if we restrict to policies π such that $\pi(a | s)/\text{unif}_{\mathcal{A}}(a)$ is bounded by some constant K for all $s \in \mathcal{S}$, then it is trivial to verify that Lemma 1 holds with $|\mathcal{A}|$ replaced by K . Therefore, in cases where this restriction in policy class is feasible, we may be able to produce identical PAC results under this restriction with $|\mathcal{A}|$ replaced by K .

Nevertheless, this restriction may be unsuitable in practice, since we might want to be able to learn approximately optimal deterministic policies, which have unbounded density. Fortunately, in many cases it may be possible to extend PAC results based on this restriction to unrestricted PAC results when reward and transition functions are Hölder continuous, based on the following lemma.

Lemma 3. *Let $V(\pi; R) = \mathbb{E}_\pi[\sum_{h \in H} R(s_h, a_h)]$ denote the value of policy π for reward function R , where \mathbb{E}_π denotes expectation over data sampled following policy π for all H time steps. In addition, assume that there exists some $\alpha \in (0, 1]$ and $L < \infty$ such that: (1) $R(s, a)$ is uniformly α -Hölder continuous in a with norm at most L all $s \in \mathcal{S}$; and (2) $\phi_h^*(s, a)^\top \psi_h^*(\cdot)$ is uniformly α -Hölder continuous in a under total variation distance with norm at most L for all $s \in \mathcal{S}$ and $h \in [H]$. Then, for any policy π , there exists a corresponding policy π_K such that $\sup_{s \in \mathcal{S}, a \in \mathcal{A}} \pi_K(a | s)/\text{unif}_{\mathcal{A}}(a) \leq K$, and*

$$|V(\pi) - V(\pi_K)| \leq 2\sqrt{m}LHK^{-\alpha/m}.$$

The proof of this lemma follows by comparing the value of π with the value of a “smoothed” version of π that has density ratio at most K , using the simulation lemma (Agarwal et al., 2019). We provide details in the appendix.

Lemma 3 implies that learning an approximately optimal policy under this density restriction gives an approximately optimal unrestricted policy when K is sufficiently large, and when rewards and transitions are sufficiently smooth. Similarly, for model-based algorithms under the same smoothness assumptions, if a learned model $\widehat{\mathcal{M}}$ of the MDP gives ϵ -accurate value estimates for all policies with density ratio at most K , then $\widehat{\mathcal{M}}$ also gives $\epsilon + 2LHK^{-\alpha/m}$ -accurate estimates for all unrestricted policies. In either case, this can be used to extend PAC bounds under the policy density restriction to unrestricted PAC bounds, by setting K appropriately. We will use the latter model-based approach for the reward-agnostic setting of FLAMBE in the following section.

5 CASE STUDY: FLAMBE

We consider a case study on implementing our proposed strategies to extend PAC results for low-rank MDPs to continuous action spaces, by applying these to the FLAMBE algorithm of Agarwal et al. (2020). We first present the FLAMBE algorithm along with the existing PAC result that holds for discrete actions, and then we extend this result to continuous actions under appropriate smoothness conditions. Then, at

the end of the section we provide a brief discussion of implementation.

5.1 FLAMBE Algorithm

We summarize the seminal FLAMBE algorithm of Agarwal et al. (2020) in Algorithm 1. This algorithm is computationally efficient assuming that we have access to two oracles: (1) a maximum likelihood estimation (MLE) oracle, which, given a dataset of n tuples $\{(s_i, a_i, s'_i) : i \in [n]\}$ of state, action, and successor state, returns the embeddings that maximize the empirical log likelihood $\sum_{i=1}^n \log(\phi(s_i, a_i)^\top \psi(s'_i))$; and (2) a sampling oracle, which, given a pair of embeddings $\phi \in \Phi$ and $\psi \in \Psi$, and a state-action pair (s, a) , can sample successor states $s' \sim \phi(s, a)^\top \psi(\cdot)$. The MLE oracle is used to solve the maximum likelihood step of the algorithm, and the sampling oracle is used to solve the approximate policy optimization problem in the elliptical planner.² We note that for practical applications (Φ, Ψ) could be parameterized directly as a generative model, in which case the MLE oracle could be implemented using standard methods for fitting generative models, and the sampling oracle would be trivial. In contrast to Agarwal et al. (2020), our presentation of FLAMBE only performs approximate policy optimization in the planner. This approach aligns more closely with their analysis using the sampling oracle and does not impact the theoretical analysis.³

Ultimately, FLAMBE returns a learnt model $\widehat{\mathcal{M}}$ for \mathcal{M} and estimates $\hat{\phi}_h, \hat{\psi}_h$ for ϕ_h^* and ψ_h^* , respectively, for each $h \in [H]$. These learnt embeddings can then be used to perform planning for arbitrary reward functions, with the PAC guarantee for *sparse* reward functions defined as follows:

Definition 4. Let $V(\pi; R, \mathcal{M})$ denote the value (expected sum of rewards) of policy π under MDP model \mathcal{M} , with reward function $R = (R_0, \dots, R_{H-1})$ such that $R_h : \mathcal{S} \times \mathcal{A} \mapsto [0, 1]$ for each $h \in [H]$. We say that R is sparse if $V(\pi; R, \mathcal{M}) \leq 1$ for every policy π and model \mathcal{M} .

As a simple example, reward functions that can only take non-zero values at a single time step are sparse. The authors gave the following guarantee on FLAMBE’s performance, which is adapted from Agarwal et al. (2020, Theorem 2 and Appendix B).

Theorem 4. For appropriate values of β, J_{\max}, n ,

²By Agarwal et al. (2020, Lemma 15) this step can be solved in $\tilde{\mathcal{O}}(d^3 H^6 (1/\beta)^4 \log(1/\delta))$ calls to the sampling oracle, with probability at least $1 - \delta$.

³This is because at the termination of the elliptical planner we still have $\mathbb{E}[\phi_{\tilde{h}}(s_{\tilde{h}}, a_{\tilde{h}})^\top \Sigma_T^{-1} \phi_{\tilde{h}}(s_{\tilde{h}}, a_{\tilde{h}}) \mid \pi, \widetilde{\mathcal{M}}] \leq \beta$ for every policy π , which is all that is required.

FLAMBE returns a model $\widehat{\mathcal{M}}$ that, with probability at least $1 - \delta$, satisfies that for every policy π and sparse reward function R ,

$$|V(\pi; R, \widehat{\mathcal{M}}) - V(\pi; R, \mathcal{M})| \leq \epsilon,$$

and the number of trajectories collected by the algorithm is

$$\tilde{\mathcal{O}}\left(H^{22} |\mathcal{A}|^9 d^7 \left(\frac{1}{\epsilon}\right)^{10} \ln\left(\frac{|\Phi||\Psi|}{\delta}\right)\right).$$

That is, with a polynomial sample complexity, the algorithm learns a model $\widehat{\mathcal{M}}$ that can be used for policy optimization under *any* sparse reward function. Although the assumption that R is sparse may seem restrictive, as argued in Agarwal et al. (2020), this is sufficient to ensure that we can fit Q functions using the estimated embeddings $\hat{\phi}_h$ with error at most ϵ . We also note that this result can be easily extended to a similar PAC guarantee for *all* (non-sparse) reward functions taking values in $[0, 1]$, by replacing ϵ with ϵ/H in the sample complexity bound, since any reward function can be written as a sum of H sparse reward functions.

We also note that the exact same FLAMBE algorithm can be used with continuous action spaces as long as we can sample actions uniformly from \mathcal{A} : for example, when $\mathcal{A} = [0, 1]^m$ we can sample uniformly random actions by sampling m i.i.d. uniformly random variables from $[0, 1]$. All of the results below apply to the same FLAMBE algorithm without any special modifications for continuous actions.

5.2 Smoothness Assumption for Extension

We now provide a concrete smoothness assumption that can be used to extend Theorem 4 to continuous action spaces. As in Section 4, we consider continuous action spaces of the form $\mathcal{A} = [0, 1]^m \subset \mathbb{R}^m$ for some $m \in \mathbb{N}^+$. These assumptions will allow us to use the methods detailed in Section 4 in order to obtain generalized PAC bounds, which we present in Section 5.3.

For this extension, we require that the embedding classes Φ and Ψ satisfy a smoothness condition for transition function errors. Specifically, we require that the errors of transition functions allowed by Φ and Ψ versus the true transition functions are smooth w.r.t. actions, as detailed below:

Assumption 2 (α -smooth Transition Errors). *There exists some error function $\mathcal{E} : \Delta(\mathcal{S}) \times \Delta(\mathcal{S}) \rightarrow [0, 1]$ on distributions of states, $\alpha_E \in (0, \infty)$, and $L_E < \infty$ such that the following hold: (1) $\text{TV}(\rho, \rho') \leq \mathcal{E}(\rho, \rho') \leq \mathcal{H}(\rho, \rho')$ for all distributions $\rho, \rho' \in \Delta(\mathcal{S})$,*

Algorithm 1 FLAMBE: Feature Learning and Model-Based Exploration

Input: Environment \mathcal{M} , function classes Φ, Ψ , subroutines MLE and SAMP, parameters β, n .

Set ρ_0 to be the random policy, which takes all actions uniformly at random.

Set $D_h = \emptyset$ for each $h \in \{0, \dots, H-1\}$.

for $j = 1, \dots, J_{\max}$ **do**

for $h = 0, \dots, H-1$ **do**

 Collect n samples (s_h, a_h, s_{h+1}) by rolling into s_h with ρ_{j-1} and taking $a_h \sim \text{unif}(\mathcal{A})$.

 Add these samples to D_h .

 Solve maximum likelihood problem: $(\hat{\phi}_h, \hat{\psi}_h) \leftarrow \text{MLE}(D_h)$.

 Set $\hat{T}_h(s_{h+1} | s_h, a_h) = \langle \hat{\phi}_h(s_h, a_h), \hat{\psi}_h(s_{h+1}) \rangle$

end for

For each h , call the planner (Algorithm 2) with h step model $\hat{T}_{0:h-1}$ and β to obtain ρ_h^{pre} .

Set $\rho_j = \text{unif}(\{\rho_h^{\text{pre}} \circ \text{unif}(\mathcal{A})\}_{h=0}^{H-1})$ to be uniform over the discovered h -step policies, each of which we augment with random actions from time step h onwards.

end for

Algorithm 2 Elliptical Planner

Input: MDP $\widetilde{\mathcal{M}} = (\phi_{0:\tilde{h}}, \psi_{0:\tilde{h}})$, subroutine SAMP, parameter $\beta > 0$. Initialize $\Sigma_0 = I_{d \times d}$.

for $t = 1, 2, \dots$ **do**

 Compute π_t such that

$$\mathbb{E} \left[\phi_{\tilde{h}}(s_{\tilde{h}}, a_{\tilde{h}})^T \Sigma_{t-1}^{-1} \phi_{\tilde{h}}(s_{\tilde{h}}, a_{\tilde{h}}) \mid \pi_t, \widetilde{\mathcal{M}} \right] \geq \sup_{\pi} \mathbb{E} \left[\phi_{\tilde{h}}(s_{\tilde{h}}, a_{\tilde{h}})^T \Sigma_{t-1}^{-1} \phi_{\tilde{h}}(s_{\tilde{h}}, a_{\tilde{h}}) \mid \pi, \widetilde{\mathcal{M}} \right] - \frac{\beta}{2}$$

If the objective is at most $\beta/2$, halt and output $\rho = \text{unif}(\{\pi_{\tau}\}_{\tau < t})$.

Compute $\Sigma_{\pi_t} = \mathbb{E} \left[\phi_{\tilde{h}}(s_{\tilde{h}}, a_{\tilde{h}}) \phi_{\tilde{h}}(s_{\tilde{h}}, a_{\tilde{h}})^T \mid \pi_t, \widetilde{\mathcal{M}} \right]$. Update $\Sigma_t \leftarrow \Sigma_{t-1} + \Sigma_{\pi_t}$.

end for

where $\text{TV}(\cdot, \cdot)$ and $\mathcal{H}(\cdot, \cdot)$ denote total variation distance and Hellinger distance, respectively; and (2) for any $s \in \mathcal{S}$, $\phi \in \Phi$, $\psi \in \Psi$, and $h \in [H]$, the function $a \mapsto \mathcal{E}(\phi(s, a)^T \psi(\cdot), \phi_h^*(s, a)^T \psi_h^*(\cdot))$ is α_E -smooth with C^{α_E} -norm at most L_E .

In other words, this assumption requires that the errors (in terms of some metric that lies between the total variation and Hellinger distances) of transition operators allowed by Φ, Ψ are smooth in the action a .

We note that this assumption is used to apply Theorem 2 in place of Lemma 1 in the FLAMBE analysis. While Agarwal et al. (2020) only use the total variation distance error function in their analysis, we decided to weaken this assumption and allow for various error functions. This relaxation is due to the fact that the smoothness assumption might be more feasible for certain error functions than others, depending on the context. We elaborate on two particular cases below.

Using Total Variation Distance

In cases where the embeddings $\phi \in \Phi$ themselves are α_E -Hölder continuous for some $\alpha_E \in (0, 1]$, we can easily justify Assumption 2 using total variation distance, as shown in the following lemma. Note that this

case includes Lipschitz-continuous embeddings with $\alpha_E = 1$.

Lemma 5. Suppose that $\phi(s, a)$ is α -Hölder continuous in a with norm L_Φ for every $\phi \in \Phi$ and $s \in \mathcal{S}$ (under L_2 norm) for some $\alpha \in (0, 1]$, and that $\int \|\psi(s)\|_2 d\mu(s) \leq U$ for all $\psi \in \Psi$, for some fixed $U < \infty$. Then, the conditions of Assumption 2 hold with $\mathcal{E}(\cdot, \cdot) = \text{TV}(\cdot, \cdot)$, $\alpha_E = \alpha$, and $L_E = 2UL_\Phi$.

We defer the proof of Lemma 5 to the appendix. Note that the assumption on Ψ in Lemma 5 is slightly stronger than the assumed normalization condition in Definition 1, but it can also be easily justified in many settings; for example, if all embedding values are non-negative then it is easy to argue that $U \leq d$.⁴

Using Hellinger Distance

Unfortunately, Lemma 5 only allows us to leverage smoothness up to Lipschitz continuity ($\alpha_E = 1$). In order to exploit higher order smoothness, we can instead consider the Hellinger distance. Recall that the

⁴This holds since $\|\psi(s')\|_1 = \sum_{i=1}^d \psi(s')_i$ for non-negative embeddings, therefore we have $\int \|\psi(s)\|_2 d\mu(s) \leq \int \|\psi(s)\|_1 d\mu(s) = \|\int \psi(s) d\mu(s)\|_1 \leq d$, given the normalization in Definition 1 and $\|\cdot\|_2 \leq \|\cdot\|_1 \leq \sqrt{d} \|\cdot\|_2$.

Hellinger distance is defined according to $\mathcal{H}(\rho, \rho') = \frac{1}{\sqrt{2}} \int (\sqrt{\rho(s)} - \sqrt{\rho'(s)})^2 d\mu(s)^{1/2}$, for distributions ρ, ρ' that are continuous w.r.t. μ . Since this corresponds to L_2 distance in the Hilbert space of square root densities, it is evident that for any integer $\alpha_E \in \mathbb{N}^+$, Assumption 2 holds with Hellinger distance as long as all order α_E mixed derivatives of $\sqrt{\phi(s, a)^\top \psi(\cdot)}$ w.r.t. a exist, and that these derivatives have $L_2(\mathcal{S})$ norm at most $L_E/2$.

Noting that $\frac{d}{dx} \sqrt{f(x)} = f(x)^{-1/2} \frac{d}{dx} f(x)$, one could easily justify this as long as $\phi(s, a)$ is uniformly α_E -smooth in a for all $\phi \in \Phi$ and $s \in \mathcal{S}$, and $\sup_{\phi, \psi, s, a} \int (\phi(s, a)^\top \psi(s'))^{1-2\alpha_E} d\mu(s') < \infty$. The latter of these conditions holds for any α_E if all transition operators allowed by Φ, Ψ have a uniformly bounded minimum density. Similarly, Assumption 2 may be justified for non-integer $\alpha_E > 1$ using analogous reasoning in terms of fractional derivatives.

Compared with using total variation distance, this approach allows us to leverage higher order smoothness of the embeddings $\phi(s, a)$, at the cost of requiring that the allowed transition densities $\phi(s, a)^\top \psi(\cdot)$ cannot be too concentrated around zero.

5.3 Bound for Restricted Policies

We now present the first of our PAC results for FLAMBE in the continuous action setting under Assumption 2. For this result, we use both the methods described in Section 4.1 and Section 4.2. In some parts of the FLAMBE analysis, Lemma 1 is applied to policies returned by the elliptical planner (which may have arbitrarily large density) and with the transition function error (which is smooth under Assumption 2) so we can apply the approach described in Section 4.1. In other places, Lemma 1 is applied to the policies that we want to bound estimation error for with discontinuous error functions, in which case we can apply the simple bound described in Section 4.2. With this, the sample complexity of our result depends on the density ratio K of the policies being evaluated.

Theorem 6. *Suppose \mathcal{M} has rank d , and let Assumption 1 and Assumption 2 be given. Then, if we set the hyperparameters β , n , and J_{\max} appropriately, with probability at least $1 - \delta$, FLAMBE returns a model $\widehat{\mathcal{M}}$ that satisfies*

$$|V(\pi; R, \widehat{\mathcal{M}}) - V(\pi; R, \mathcal{M})| \leq \epsilon$$

for every sparse reward function R , and every policy π that is continuous with respect to the measure on \mathcal{A} and satisfies $\sup_{s, a} \pi(a | s) \leq K$. The number of

trajectories collected by the algorithm is

$$\tilde{O} \left(H^{22+16\tau} K^{5+4\tau} d^{7+4\tau} \left(\frac{1}{\epsilon} \right)^{10+8\tau} \cdot L_E^{(9+8\tau)\kappa} \ln \left(\frac{|\Phi||\Psi|}{\delta} \right) \right),$$

where $\tau = m/\alpha_E$, and $\kappa = m/(m + \alpha_E)$.

The exact hyperparameter values that yield this result are given in the proof of this theorem (see Appendix A). In this bound, K corresponds to the maximum probability density (or sharpness) of policies we are guaranteed to accurately evaluate, and τ corresponds (inversely) to the assumed level of smoothness relative to the dimensionality of the action space. We note that as $\tau \rightarrow 0$ (*i.e.*, we assume a very high order of smoothness relative to m) we recover the same dependence on H , d , and $1/\epsilon$ as in Theorem 4 (L_E and K together play a similar role as $|\mathcal{A}|$, so comparing exponents of these parameters is less straightforward). However, when we have a lower level of smoothness, the polynomial factors in this bound can be significantly worse than in Theorem 4. A simple intuition for this is that the smoother the transition operators allowed by the embedding classes Φ and Ψ , and the lower the number of degrees of freedom in the action space, the easier it is to explore \mathcal{A} .

5.4 Bound for Unrestricted Policies

Unfortunately, in many applications we would like to consider deterministic policies, for which $K = \infty$ and Theorem 6 becomes vacuous. In order to address this, we take the approach described in Section 4.2. Here we require an additional assumption on the smoothness of the actual transition operator, as follows:

Assumption 3. *(α -smooth Transitions) There exists some fixed $\alpha_T \in (0, 1]$ and $L_T < \infty$ such that $\text{TV}(\phi_h^*(s, a)^\top \psi_h^*(\cdot), \phi_h^*(s, a')^\top \psi_h^*(\cdot)) \leq L_T \|a' - a\|^{\alpha_T}$, for every $h \in [H]$, $s \in \mathcal{S}$, and $a, a' \in \mathcal{A}$.*

This assumption is easy to justify as long as the true embeddings $\phi_h^*(s, a)$ are uniformly α_T -Hölder continuous for all s, h by applying Lemma 5. Given this, we are ready to provide our second PAC result.

Theorem 7. *Let the conditions and definitions of Theorem 6 be given, as well as Assumption 3. Then, if we set the hyperparameters β , n , and J_{\max} appropriately, with probability at least $1 - \delta$, FLAMBE returns a model $\widehat{\mathcal{M}}$ that satisfies*

$$|V(\pi; R, \widehat{\mathcal{M}}) - V(\pi; R, \mathcal{M})| \leq \epsilon$$

for every policy π , and every sparse reward function R such that $R_h(s, a)$ is α_R -Hölder continuous with norm

at most L_R for all $h \in [H]$ and $s \in \mathcal{S}$. The number of trajectories collected by the algorithm is

$$\tilde{O} \left(H^{22+16\tau+(4\tau+5)\sigma} d^{7+4\tau} \left(\frac{1}{\epsilon} \right)^{10+8\tau+(4\tau+5)\sigma} \cdot L_E^{(9+8\tau)\kappa} L^{(4\tau+5)\sigma} \ln \left(\frac{|\Phi||\Psi|}{\delta} \right) \right),$$

where $\sigma = m / \min(\alpha_T, \alpha_R)$, and $L = \max(L_T, L_R)$.

We note that this theorem similarly gives us a sample complexity that is polynomial in all problem variables for fixed action space dimension and order of smoothness. We can see that the less smooth the reward functions under consideration (*i.e.*, the smaller α_R is and the greater L_R is), the greater the number of trajectories needed to learn a model that can accurately estimate policy values. Similarly, the less smooth the transition function is, the more trajectories are needed. Note as well that, since $\alpha_T \leq 1$, we cannot have $\sigma \rightarrow 0$ for very smooth problem instances, so this bound gives worse exponents than Theorem 6 for smooth problem instances, at the benefit of allowing for arbitrarily sharp policies. The proof of Theorem 7 is a fairly straightforward application of Lemma 3 to Theorem 6, along the lines discussed in Section 4.2, and we defer details to the appendix.

5.5 Discussion of Implementation

FLAMBE with continuous actions can mostly be implemented identically as with discrete actions. The only slight complication comes with implementing the elliptical planner. As in the discrete case, we can implement this planner using backward dynamic programming, since under MDP linearity Q functions are linear in $\phi(s, a)$ (Jin et al., 2020). However, implementing this requires computing $\max_a w^\top \hat{\phi}(s, a)$ for various vectors w , states s , and estimated embeddings $\hat{\phi}$. With discrete actions this optimization is trivial, whereas when \mathcal{A} is continuous we require some kind of non-trivial optimization procedure. Below we discuss some possible approaches.

Using Concave Embeddings If we ensure that by design $\phi(s, a)$ is concave in a for all $\phi \in \Phi$, then the optimization problem $\max_a w^\top \hat{\phi}(s, a)$ is tractable (equivalent to minimizing a convex function) as long as w is non-negative. This latter condition may be additional assumptions on a problem by problem basis depending on the embedding structure.

Using Grid Search We could approximately solve the problem $\max_a w^\top \hat{\phi}(s, a)$ by performing grid search over $a \in [0, 1]^m$. When the embeddings $\hat{\phi}(s, a)$ are

relatively smooth in a , and m is not too large, then this approach should be reasonably tractable.

Using Blackbox Non-convex optimization A third approach, which might be the most useful in many problems in practice, is to simply treat $\max_a w^\top \hat{\phi}(s, a)$ as a blackbox optimization problem. For example, for some fixed iteration of dynamic programming, where w and $\hat{\phi}$ are fixed, we could learn a mapping from s to $\operatorname{argmax}_a w^\top \hat{\phi}(s, a)$ using machine learning procedures such as deep learning, treating this as a general nonconvex policy optimization problem. While this kind of approach may be difficult to theoretically analyze, it may be effective in ensuring accurate elliptical planning performance in practice.

6 Conclusion

We studied the problem of generalizing low-rank MDP methods to settings with continuous action spaces. We first provided a detailed discussion of exactly how/why low-rank MDP PAC bounds depend on the size of the action space, and we discussed multiple techniques that could be used to extend existing PAC results to continuous action spaces by utilizing smoothness of problems parameters w.r.t. actions. Then, as a case study, we considered the seminal reward-agnostic FLAMBE algorithm of Agarwal et al. (2020), and investigated how to extend their theory to continuous action spaces using these methods. Based on these insights, we provided novel PAC bounds for low-rank MDPs with continuous action spaces, and we derived sample complexities that are polynomial in all problem parameters except for the assumed degree of smoothness relative to the dimension of the action space. We hope that this work is useful and impactful in motivating further research on low-rank MDPs, as well as other practical PAC RL settings, with continuous action spaces. In particular, an interesting and relevant avenue of future research would be to construct lower bounds on sample complexity under appropriate smoothness assumptions, further advancing the theoretical foundations and algorithmic approaches in these areas.

Acknowledgements

We thank the anonymous reviewers for their helpful comments. This material is based upon work supported by the National Science Foundation under Grant No. 1846210 and the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, under Award Number DE-SC0023112.

References

- A. Agarwal and T. Zhang. Non-linear reinforcement learning in large action spaces: Structural conditions and sample-efficiency of posterior sampling. In *Conference on Learning Theory*, pages 2776–2814. PMLR, 2022.
- A. Agarwal, N. Jiang, and S. M. Kakade. Reinforcement learning: Theory and algorithms. 2019.
- A. Agarwal, S. Kakade, A. Krishnamurthy, and W. Sun. Flambe: Structural complexity and representation learning of low rank mdps. In *Advances in Neural Information Processing Systems*, volume 33, pages 20095–20107, 2020.
- A. Behzadan and M. Holst. Multiplication in sobolev spaces, revisited. *Arkiv för Matematik*, 59(2):275–306, 2021.
- H. Brezis and P. Mironescu. Gagliardo–nirenberg inequalities and non-inequalities: the full story. In *Annales de l’Institut Henri Poincaré C, Analyse non linéaire*, volume 35, pages 1355–1376. Elsevier, 2018.
- H. Brezis and P. Mironescu. Where sobolev interacts with gagliardo–nirenberg. *Journal of functional analysis*, 277(8):2839–2864, 2019.
- H. Cai, K. Ren, W. Zhang, K. Malialis, J. Wang, Y. Yu, and D. Guo. Real-time bidding by reinforcement learning in display advertising. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pages 661–670, 2017.
- J. Chen, A. Modi, A. Krishnamurthy, N. Jiang, and A. Agarwal. On the statistical efficiency of reward-free exploration in non-linear rl. *Advances in Neural Information Processing Systems*, 35:20960–20973, 2022.
- Y. Cheng, R. Huang, J. Yang, and Y. Liang. Improved sample complexity for reward-free reinforcement learning under low-rank mdps. *arXiv preprint arXiv:2303.10859*, 2023.
- S. Chinchali, P. Hu, T. Chu, M. Sharma, M. Bansal, R. Misra, M. Pavone, and S. Katti. Cellular network traffic scheduling with deep reinforcement learning. In *Thirty-second AAAI conference on artificial intelligence*, 2018.
- O. D. Domingues, P. Ménard, M. Pirotta, E. Kauffmann, and M. Valko. Regret bounds for kernel-based reinforcement learning. *arXiv preprint arXiv:2004.05599*, 2020.
- S. Du, A. Krishnamurthy, N. Jiang, A. Agarwal, M. Dudik, and J. Langford. Provably efficient rl with rich observations via latent state decoding. In *International Conference on Machine Learning*, pages 1665–1674. PMLR, 2019.
- S. S. Du, S. M. Kakade, J. D. Lee, S. Lovett, G. Mahajan, W. Sun, and R. Wang. Bilinear classes: A structural framework for provable generalization in rl. *arXiv preprint arXiv:2103.10897*, 2021.
- A. Fiorenza, M. R. Formica, T. G. Roskovec, and F. Soudský. Detailed proof of classical gagliardo–nirenberg interpolation inequality with historical remarks. *Zeitschrift für Analysis und ihre Anwendungen*, 40(2):217–236, 2021.
- E. Giné and R. Nickl. *Mathematical foundations of infinite-dimensional statistical models*. Cambridge university press, 2021.
- Y. Hu, N. Kallus, and X. Mao. Smooth contextual bandits: Bridging the parametric and nondifferentiable regret regimes. *Operations Research*, 70(6):3261–3281, 2022.
- S. Jia, Q. Xie, N. Kallus, and P. I. Frazier. Smooth non-stationary bandits. *arXiv preprint arXiv:2301.12366*, 2023.
- N. Jiang, A. Krishnamurthy, A. Agarwal, J. Langford, and R. E. Schapire. Contextual decision processes with low bellman rank are pac-learnable. In *International Conference on Machine Learning*, pages 1704–1713. PMLR, 2017.
- C. Jin, Z. Yang, Z. Wang, and M. I. Jordan. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pages 2137–2143. PMLR, 2020.
- J. Kober, J. A. Bagnell, and J. Peters. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274, 2013.
- A. Krishnamurthy, A. Agarwal, and J. Langford. Pac reinforcement learning with rich observations. *arXiv preprint arXiv:1602.02722*, 2016.
- G. Leoni. *A first course in Sobolev spaces*. American Mathematical Soc., 2017.
- Y. Li, W. Zheng, and Z. Zheng. Deep robust reinforcement learning for practical algorithmic trading. *IEEE Access*, 7:108014–108022, 2019.
- J.-G. Liu and J. Wang. On the best constant for gagliardo–nirenberg interpolation inequalities. *arXiv preprint arXiv:1712.10208*, 2017.
- Y. Liu, Y. Wang, and A. Singh. Smooth bandit optimization: generalization to holder space. In *International Conference on Artificial Intelligence and Statistics*, pages 2206–2214. PMLR, 2021.
- T. Lykouris and S. Vassilvtiskii. Competitive caching with machine learned advice. In *International Conference on Machine Learning*, pages 3296–3305. PMLR, 2018.

- D. Misra, M. Henaff, A. Krishnamurthy, and J. Langford. Kinematic state abstraction and provably efficient rich-observation reinforcement learning. In *International conference on machine learning*, pages 6961–6971. PMLR, 2020.
- A. Modi, J. Chen, A. Krishnamurthy, N. Jiang, and A. Agarwal. Model-free representation learning and exploration in low-rank mdps. *arXiv preprint arXiv:2102.07035*, 2021.
- A. E. Sallab, M. Abdou, E. Perot, and S. Yogamani. Deep reinforcement learning framework for autonomous driving. *Electronic Imaging*, 2017(19):70–76, 2017.
- S. R. Sinclair, S. Banerjee, and C. L. Yu. Adaptive discretization in online reinforcement learning. *arXiv preprint arXiv:2110.15843*, 2021.
- Z. Song and W. Sun. Efficient model-free reinforcement learning in metric spaces. *arXiv preprint arXiv:1905.00475*, 2019.
- E. M. Stein. *Singular integrals and differentiability properties of functions*. Princeton university press, 1970.
- C. J. Stone. Optimal global rates of convergence for nonparametric regression. *The annals of statistics*, pages 1040–1053, 1982.
- W. Sun, N. Jiang, A. Krishnamurthy, A. Agarwal, and J. Langford. Model-based rl in contextual decision processes: Pac bounds and exponential improvements over model-free approaches. In *Conference on learning theory*, pages 2898–2933. PMLR, 2019.
- C. Tessler, Y. Shpigelman, G. Dalal, A. Mandelbaum, D. H. Kazakov, B. Fuhrer, G. Chechik, and S. Mannor. Reinforcement learning for datacenter congestion control. *arXiv preprint arXiv:2102.09337*, 2021.
- A. B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Publishing Company, Incorporated, 1st edition, 2008. ISBN 0387790519.
- M. Uehara, X. Zhang, and W. Sun. Representation learning for online and offline rl in low-rank mdps. *arXiv preprint arXiv:2110.04652*, 2021.
- L. Yang and M. Wang. Sample-optimal parametric q-learning using linearly additive features. In *International Conference on Machine Learning*, pages 6995–7004. PMLR, 2019.
- L. Yang and M. Wang. Reinforcement learning in feature space: Matrix bandit, kernels, and regret bound. In *International Conference on Machine Learning*, pages 10746–10756. PMLR, 2020.
- G. Zheng, F. Zhang, Z. Zheng, Y. Xiang, N. J. Yuan, X. Xie, and Z. Li. Drn: A deep reinforcement learning framework for news recommendation. In *Proceedings of the 2018 World Wide Web Conference*, pages 167–176, 2018.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. **Yes**
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. **Yes**
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. **Not Applicable**
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. **Yes**
 - (b) Complete proofs of all theoretical results. **Yes**
 - (c) Clear explanations of any assumptions. **Yes**
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). **Not Applicable**
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). **Not Applicable**
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). **Not Applicable**
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). **Not Applicable**
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. **Not Applicable**
 - (b) The license information of the assets, if applicable. **Not Applicable**
 - (c) New assets either in the supplemental material or as a URL, if applicable. **Not Applicable**
 - (d) Information about consent from data providers/curators. **Not Applicable**

- (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. **Not Applicable**
- 5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. **Not Applicable**
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. **Not Applicable**
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. **Not Applicable**

A PROOFS

A.1 Proof of Theorem 2

The proof is based on a recent contribution (Lemma 9) from the field of functional analysis. Before delving into it, we introduce a prior, more established lemma to develop a better intuition about the final result.

Lemma 8 (Gagliardo–Nirenberg interpolation inequality). *Let $1 \leq p, q \leq \infty$, $s \in \mathbb{N}$, $k \in \mathbb{N}^+$ such that $0 \leq k < s$ and let θ be such that*

$$0 \leq \theta \leq 1 - k/s$$

and

$$(1 - \theta) \left(\frac{1}{p} - \frac{s - k}{m} \right) + \theta \left(\frac{1}{q} + \frac{k}{m} \right) = \frac{1}{r} \in (-\infty, 1]$$

Then there exists a constant $C_{s,m,k,p,q,r} > 0$ that depends only on s, m, k, p, q, r such that:

$$\|\nabla^k f\|_{L^r(\mathbb{R}^m)} \leq C_{s,m,k,p,q,r} \|\nabla^s f\|_{L^p(\mathbb{R}^m)}^{1-\theta} \|f\|_{L^q(\mathbb{R}^m)}^\theta$$

for every $f : \mathbb{R}^m \rightarrow \mathbb{R}$ with the following properties: (1) f is s times differentiable; (2) $\|\nabla^s f\|_{L^p(\mathbb{R}^m)}$, $\|f\|_{L^q(\mathbb{R}^m)} < \infty$; and (3) f vanishes at infinity.

The Gagliardo–Nirenberg (GN) interpolation inequality was originally proposed independently by Emilio Gagliardo and Louis Nirenberg in 1958, but the generality of the statement, as well as the details of the proof have been subject to refinement until recently. The version presented in Lemma 8 is based on Theorem 12.87 in Leoni (2017). The GN interpolation inequality has found numerous applications in various branches of mathematics, including the theory of elliptic and parabolic partial differential equations, harmonic analysis, and functional analysis.

The GN inequality will be central to our proof of Theorem 2. To see this, let $g : \mathbb{R}^m \rightarrow \mathbb{R}$ be a function defined on all of \mathbb{R}^m that is α -smooth with $\alpha \geq 1$ and vanishes at ∞ . Setting $p, r = \infty$, $s = \lfloor \alpha \rfloor$, $k = 0$, $q = 1$ in Lemma 8, we immediately obtain $\theta = \frac{\lfloor \alpha \rfloor}{m + \lfloor \alpha \rfloor}$ and a bound of the form:

$$\|g\|_\infty \leq C_{\lfloor \alpha \rfloor, m, 0, \infty, 1, \infty} L^{\frac{m}{m + \lfloor \alpha \rfloor}} \|g\|_{L^1(\mathbb{R}^m)}^{\frac{\lfloor \alpha \rfloor}{m + \lfloor \alpha \rfloor}} \quad (1)$$

Remark 1 (GN Inequality Constant). *In (1), the GN constant $C_{\lfloor \alpha \rfloor, m, 0, \infty, 1, \infty}$ depends on $p, r = \infty$ which can potentially lead to a large constant value. Finding the best values for the GN constant is an area of active research with several recent results finding tight values for these constants in specific scenarios. For example, Liu and Wang (2017) find that for $s = 1$ (i.e. $\lfloor \alpha \rfloor = 1$) the constant is $C_{1, m, 0, \infty, 1, \infty} = 1$ (see Eq. 1.20 from their paper). We assume that the GN constant for $s > 1$ will similarly have values that are not too large.*

Two challenges arise when applying Lemma 8 directly to our scenario: (1) the GN inequality applies to functions over the entire \mathbb{R}^m , while our functions are defined on a bounded domain, and (2) the bound lacks precision for non-integer α values, requiring an adjustment similar to Theorem 2. Traditionally, the first issue has been mitigated using linear extension operators for domains with Lipschitz boundaries (see Stein (1970)). The second challenge has been fully resolved only recently, through the work of Ha"im Brezis and Petru Mironescu (Brezis and Mironescu, 2018, 2019). This result relies on the Sobolev norm, which we define here for α -smooth functions:

Definition 5 (Sobolev norm of α -smooth functions). *For two integers $s \in \mathbb{N}$ and $p \in \mathbb{N}^+$, the Sobolev norm $\|\cdot\|_{W^{s,p}}$ of an s -smooth function defined on \mathcal{D} is defined as:*

$$\|f\|_{W^{s,p}(\mathcal{D})} = \sum_{j=0}^s \|\nabla^j f\|_{L^p(\mathcal{D})}$$

For α -smooth functions with a potentially non-integer $\alpha \in (0, \infty)$, the fractional Sobolev norm for the (α, ∞) indices is given by:

$$\|f\|_{W^{\alpha, \infty}(\mathcal{D})} = \|f\|_{W^{\lfloor \alpha \rfloor, \infty}(\mathcal{D})} + \max_{\beta: |\beta| = \lfloor \alpha \rfloor} |D^\beta f|_{C^0, \{\alpha\}}$$

We now present the GN inequality for fractional smoothness from Brezis and Mironescu (2019).

Lemma 9. [Fractional GN inequality, Theorem 1 in Brezis and Mironescu (2019)] Let \mathcal{D} be \mathbb{R}^m , a half space or a Lipschitz bounded domain in \mathbb{R}^m , and let $s_1, s_2, r, p_1, p_2, q, \theta, m$ satisfy:

$$\begin{aligned} s_1, s_2, r &\in \mathbb{R}^+, p_1, p_2, p \in \mathbb{N}, 1 \leq p_1, p_2, p \leq \infty, \\ r &< s := \theta s_1 + (1 - \theta) s_2, \text{ and} \\ \frac{1}{q} &= \left(\frac{\theta}{p_1} + \frac{1 - \theta}{p_2} \right) - \frac{s - r}{m}. \end{aligned}$$

Then, the following inequality holds

$$\|f\|_{W^{r,q}(\mathcal{D})} \leq C_{m,r,q,s_1,s_2,p_1,p_2} \|f\|_{W^{s_1,p_1}(\mathcal{D})}^\theta \|f\|_{W^{s_2,p_2}(\mathcal{D})}^{1-\theta}$$

for every $f: \mathcal{D} \rightarrow \mathbb{R}$ with $\|f\|_{W^{s_1,p_1}(\mathcal{D})}, \|f\|_{W^{s_2,p_2}(\mathcal{D})} < \infty$ if and only if at least one of the following statements

$$\text{is false: } \begin{cases} s_2 \in \mathbb{N} \text{ and } s_2 \geq 1, \\ p_2 = 1, \\ 0 < s_2 - s_1 \leq 1 - \frac{1}{p_1}. \end{cases}.$$

Here $C_{m,r,q,s_1,s_2,p_1,p_2}$ is a constant that depends only on $m, r, q, s_1, s_2, p_1, p_2$ and $\|f\|_{W^{s,p}(\mathcal{D})}$ is the fractional Sobolev norm.

We are now ready to provide a proof for Theorem 2.

Proof of Theorem 2. Let $g(a) := \mathbb{E}_{s \sim \rho}[f(s, a)]$. Since $f(s, a)$ is α -smooth in the action a with $\alpha \in (0, \infty)$ and $\|f\|_{C^\alpha} \leq L$, it immediately follows that $g(a)$ is also α -smooth with $\|g\|_{C^\alpha} \leq L$. Furthermore, since \mathcal{A} is a Lipschitz domain, we can apply Lemma 9 directly by setting the following parameters: $q = \infty, r = 0, p_1 = 1, s_1 = 0, p_2 = \infty, s_2 = \alpha$. This yields a coefficient $\theta = \frac{\alpha}{m+\alpha}$ and

$$\|g\|_{W^{0,\infty}(\mathcal{A})} \leq C_{m,0,\infty,0,\alpha,1,\infty} \|g\|_{W^{0,1}(\mathcal{A})}^{\frac{\alpha}{m+\alpha}} \|g\|_{W^{\alpha,\infty}(\mathcal{A})}^{\frac{m}{m+\alpha}} \quad (2)$$

We note that the conditions $\|g\|_{W^{s_1,p_1}(\mathcal{A})} = \|g\|_{W^{0,1}(\mathcal{A})} < \infty, \|g\|_{W^{s_2,p_2}(\mathcal{A})} = \|g\|_{W^{\alpha,\infty}(\mathcal{A})} < \infty$ hold from the properties of α -smooth functions (Definition 2). In particular, $\|g\|_{W^{\alpha,\infty}(\mathcal{A})} = \sum_{j=0}^{\lfloor \alpha \rfloor} \|\nabla^j g\|_{L^p(\mathcal{D})} + \max_{\beta: |\beta|=\lfloor \alpha \rfloor} |D^\beta g|_{C^{0,\{\alpha\}}} \leq (\lfloor \alpha \rfloor + 2)L$. Absorbing all constants into one constant C and rearranging the terms, we can rewrite (2) as:

$$\|g\|_\infty \leq CL^{\frac{\alpha}{m+\alpha}} \|g\|_{L^1(\mathcal{A})}^{\frac{\alpha}{m+\alpha}}$$

We further write $\|g\|_{L^1(\mathcal{A})} = \int_{\mathcal{A}} g(a) d\mu(a) = \mu(\mathcal{A}) \int_{\mathcal{A}} g(a) \frac{1}{\mu(\mathcal{A})} d\mu(a) = \mathbb{E}_{a \sim \text{unif}_{\mathcal{A}}} g(a)$ since $\mu(\mathcal{A}) = 1$. Putting everything together, we obtain following inequality:

$$\sup_{a \in \mathcal{A}} g(a) \leq CL^{\frac{\alpha}{m+\alpha}} [\mathbb{E}_{a \sim \text{unif}_{\mathcal{A}}} g(a)]^{\frac{\alpha}{m+\alpha}}$$

The desired conclusion immediately follows by noting that $\mathbb{E}_{s \in \rho, a \sim \pi(\cdot|s)} [f(s, a)] \leq \mathbb{E}_{s \sim \rho} [\sup_{a \in \mathcal{A}} f(s, a)] = \sup_{a \in \mathcal{A}} g(a)$ and $\mathbb{E}_{s \sim \rho, a \sim \text{unif}_{\mathcal{A}}} [f(s, a)] = \mathbb{E}_{a \sim \text{unif}_{\mathcal{A}}} \mathbb{E}_{s \sim \rho} [f(s, a)] = \mathbb{E}_{a \sim \text{unif}_{\mathcal{A}}} g(a)$. This concludes our proof. \square

Remark 2 (Generality of $\mathcal{A} = [0, 1]^n$). We note that the result in Theorem 2 is scale-invariant. That is, if we consider $\mathcal{A} = [0, 1]^n$ to be the scaling of a larger space $\mathcal{A}' = [0, A_{\max}]^n$ with $f': \mathcal{A}' \rightarrow \mathbb{R}$ such that $f'(a') = f(a'/A_{\max})$, the bound in Theorem 2 is the same across both \mathcal{A} and \mathcal{A}' . Intuitively, this is because the bound on the derivatives increases as the space shrinks, but the measure of the space decreases proportionally. Let L' be bound in Definition 2 of f' . Then $L = \frac{L'}{A_{\max}^\alpha}$. Noting that $\sup_{a \in \mathcal{A}} f(a) = \sup_{a' \in \mathcal{A}'} f'(a')$, we have:

$$\begin{aligned} \sup_{a \in \mathcal{A}} f(a) &\leq CL^{\frac{\alpha}{m+\alpha}} [\mathbb{E}_{a \sim \text{unif}_{\mathcal{A}}} f(a)]^{\frac{\alpha}{m+\alpha}} \\ &\Rightarrow \sup_{a' \in \mathcal{A}'} f'(a') \leq C \left(\frac{L'}{A_{\max}^\alpha} \right)^{\frac{\alpha}{m+\alpha}} \left(\mu(\mathcal{A}') \int_{\mathcal{A}'} f'(a') \frac{1}{\mu(\mathcal{A}')} d\mu(a') \right)^{\frac{\alpha}{m+\alpha}} \\ &= C(L')^{\frac{\alpha}{m+\alpha}} (A_{\max})^{-\frac{\alpha m}{m+\alpha}} (A_{\max})^{\frac{m\alpha}{m+\alpha}} [\mathbb{E}_{a' \sim \text{unif}_{\mathcal{A}'}} f'(a')]^{\frac{\alpha}{m+\alpha}} \\ &= C(L')^{\frac{\alpha}{m+\alpha}} [\mathbb{E}_{a' \sim \text{unif}_{\mathcal{A}'}} f'(a')]^{\frac{\alpha}{m+\alpha}} \end{aligned}$$

Thus, as long as \mathcal{A} is bounded, the result in Theorem 2 holds in all generality.

A.2 Proof of Lemma 3

Consider any policy $\pi(a \mid s)$. The core of this proof relies on the construction of a “smoothed”, intermediary policy $\pi_K(a \mid s)$. We define π_K as follows:

Definition 6. [Intermediary policy] Let $\pi(\cdot \mid s) \in \Delta(\mathcal{A})$ be a given policy that we wish to evaluate. Then, for all $s \in \mathcal{S}$ and $K \in \mathbb{R}^+$, define the intermediary policy $\pi_K(\cdot \mid s) \in \Delta(\mathcal{A})$ as the distribution that samples $a \in \mathcal{A}$ by first sampling $a' \sim \pi(\cdot \mid s)$, and then $a \sim \text{unif}_{B_\infty(a', K^{-1/m}/2) \cap \mathcal{A}}$, where $B_\infty(a', K^{-1/m}/2) = \{a \in \mathbb{R}^m : \|a - a'\|_\infty \leq K^{-1/m}/2\}$ is the ℓ_∞ ball of radius $K^{-1/m}/2$ in \mathbb{R}^m .

In other words, the intermediary policy π_K is given by sampling in an m -dimensional hypercube of side length $K^{-1/m}$ centered at action a' sampled from the original policy π . π_K is trivially continuous in the action a . Furthermore, since the volume of the hypercube is $1/K$, π_K has a density bounded by K . We now introduce a result that describes the behavior of $\mathbb{E}_{a \sim \pi_K(\cdot \mid s)}[f(a)]$, where $f(a)$ is an α -Hölder continuous function.

Lemma 10. Let $C^{\alpha, L}$ be the class of α -Hölder continuous functions $f : \mathcal{A} \rightarrow [0, 1]$ with $\alpha \in (0, 1]$ and norm at most L as in Definition 2. Furthermore, let $\pi(\cdot \mid s) \in \Delta(\mathcal{A})$ be a given policy and let $\pi_K(\cdot \mid s)$ be its intermediary policy as in Definition 6. Then, the following inequality holds:

$$\sup_{f \in C^{\alpha, L}} (\mathbb{E}_{a \sim \pi_K(\cdot \mid s)}[f(a)] - \mathbb{E}_{a \sim \pi(\cdot \mid s)}[f(a)]) \leq \sqrt{m} L K^{-\frac{\alpha}{m}}.$$

Proof of Lemma 10. For α -Hölder continuous functions with Hölder norm bounded by L we have that $|f(a) - f(a')| \leq L \|a - a'\|_2^\alpha$ for any $a, a' \in \mathcal{A}$. This allows us to bound the difference in expectations under a given policy $\pi(\cdot \mid a)$ and the corresponding intermediary policy $\pi_K(\cdot \mid a)$:

$$\begin{aligned} & \sup_{f \in C^{\alpha, L}} (\mathbb{E}_{a \sim \pi_K(\cdot \mid s)}[f(a)] - \mathbb{E}_{a \sim \pi(\cdot \mid s)}[f(a)]) \\ &= \sup_{f \in C^{\alpha, L}} \left(\mathbb{E}_{a' \sim \pi(\cdot \mid s)} \left[\mathbb{E}_{a \sim \text{unif}_{B_{\mathcal{A}}(a', K^{-1/m}) \cap \mathcal{A}}} f(a) \right] - \mathbb{E}_{a \sim \pi(\cdot \mid s)}[f(a)] \right) \\ &\leq \sup_{f \in C^{\alpha, L}} \left(\mathbb{E}_{a' \sim \pi(\cdot \mid s)} \left[\mathbb{E}_{a \sim \text{unif}_{B_{\mathcal{A}}(a', K^{-1/m}) \cap \mathcal{A}}} [f(a') + \sqrt{m} L K^{-\alpha/m}] \right] - \mathbb{E}_{a \sim \pi(\cdot \mid s)}[f(a)] \right) = \sqrt{m} L K^{-\alpha/m} \end{aligned}$$

where the inequality comes from $|f(a) - f(a')| \leq L \|a - a'\|_2^\alpha \leq L (\sqrt{m} \|a - a'\|_\infty)^\alpha \leq \sqrt{m} L K^{-\alpha/m}$ for $a, a' \in B_\infty(a', K^{-1/m}/2) \cap \mathcal{A}$. \square

Remark 3. For $\alpha > 1$, we cannot easily leverage the additional smoothness properties of the class $C^{\alpha, L}$. To understand this, let $m = 1$, $r \ll 1$, and consider the Taylor expansion of $f(a')$ around $f(a)$ such that $\|a - a'\|_2 = r$. The Taylor expansion contains a term that is bounded by Lr which is of leading order given $r \ll 1$. In other words, for small r , functions of smoothness $\alpha \geq 1$ are locally linear with derivatives bounded by L . Thus, for $\alpha > 1$, we cannot obtain a bound better than Lr using similar arguments as above. Note that Lr is still a valid bound in this regime since the functions are still Lipschitz with constant L .

Using the result of Lemma 10, we can obtain the conclusion of Lemma 3 by applying the performance difference lemma, and bounding the difference between $V(\pi; R, \tilde{\mathcal{M}})$ and $V(\pi_K; R, \tilde{\mathcal{M}})$ for any model $\tilde{\mathcal{M}}$ allowed by (Φ, Ψ) , where π_K is the intermediary policy of π (Definition 6). We give details below:

$$\begin{aligned} |V(\pi; R, \tilde{\mathcal{M}}) - V(\pi_K; R, \tilde{\mathcal{M}})| &\leq \sum_{h=0}^{H-1} \mathbb{E}_{s_h \sim \mathbb{P}_h^{\pi_K}} \left| \mathbb{E}_{a_h \sim \pi(s_h)}[Q^{\pi, \tilde{\mathcal{M}}}(s_h, a_h)] - \mathbb{E}_{a_h \sim \pi_K(s_h)}[Q^{\pi, \tilde{\mathcal{M}}}(s_h, a_h)] \right| \\ &\leq \sum_{h=0}^{H-1} \mathbb{E}_{s_h \sim \mathbb{P}_h^{\pi_K}} \left| \mathbb{E}_{a_h \sim \pi(s_h)}[R_h(s_h, a_h)] - \mathbb{E}_{a_h \sim \pi_K(s_h)}[R_h(s_h, a_h)] \right. \\ &\quad \left. + \mathbb{E}_{a_h \sim \pi(s_h)} \mathbb{E} \left[V_{h+1}^{\pi, \tilde{\mathcal{M}}}(s_{h+1}) \mid s_h, a_h, \tilde{\mathcal{M}} \right] \right. \\ &\quad \left. - \mathbb{E}_{a_h \sim \pi_K(s_h)} \mathbb{E} \left[V_{h+1}^{\pi, \tilde{\mathcal{M}}}(s_{h+1}) \mid s_h, a_h, \tilde{\mathcal{M}} \right] \right|. \end{aligned}$$

Now, by assumption R_h is α -Hölder continuous in a with Hölder norm at most L . In addition, R is a sparse reward so $V_{h+1}^{\pi, \tilde{\mathcal{M}}} : \mathcal{S} \mapsto [0, 1]$, which along with the assumption that $\phi_h^*(s, a)\psi_h^*(\cdot)$ is also α -Hölder continuous

under the total variation distance, implies that $\mathbb{E} \left[V_{h+1}^{\pi, \tilde{\mathcal{M}}}(s_{h+1}) \mid s_h, a_h, \tilde{\mathcal{M}} \right]$ is also α -Hölder continuous in a with constant L . Therefore, we can apply the result of Lemma 10 to control the terms above and obtain the desired result:

$$\begin{aligned} |V(\pi; R, \tilde{\mathcal{M}}) - V(\pi_K; R, \tilde{\mathcal{M}})| &\leq \sum_{h=0}^{H-1} \mathbb{E}_{s_h \sim \mathbb{P}_h^{\pi_K}} \left(\sqrt{m} L K^{-\alpha/m} + \sqrt{m} L K^{-\alpha/m} \right) \\ &\leq 2\sqrt{m} H L K^{-\alpha/m}. \end{aligned}$$

Remark 4. While the assumption of the theorem is that both R and $\phi_h^*(s, a)\psi_h^*(\cdot)$ are α -Hölder continuous in a with Hölder norm at most L , this is without loss of generality. To see this, assume instead that R is α_R -smooth with $\alpha_R \in (0, \infty)$ and Hölder norm L_R , and similarly $\phi_h^*(s, a)\psi_h^*(\cdot)$ is α_T -smooth with $\alpha_T \in (0, \infty)$ and Hölder norm L_T . Then, we could simply set $\alpha := \min(\alpha_R, \alpha_T, 1)$ and $L := \max(L_R, L_T)$ and the result of Lemma 3 applies due to the nesting of α -smooth classes.

A.3 Proof of Lemma 5

Proof. For any given h, ϕ, ψ, s, a, a' we have

$$\begin{aligned} &|\text{err}_h^{\phi, \psi}(s, a') - \text{err}_h^{\phi, \psi}(s, a)| \\ &\leq \left\| \left(\phi(s, a')^T \psi(\cdot) - \phi_h^*(s, a')^T \psi_h^*(\cdot) \right) \right. \\ &\quad \left. - \left(\phi(s, a)^T \psi(\cdot) - \phi_h^*(s, a)^T \psi_h^*(\cdot) \right) \right\|_{\text{TV}} \\ &\leq 2 \sup_{\phi \in \Phi, \psi \in \Psi} \left\| \phi(s, a')^T \psi(\cdot) - \phi_h(s, a)^T \psi_h(\cdot) \right\|_{\text{TV}} \\ &\leq 2U \sup_{\phi \in \Phi} \|\phi(s, a') - \phi(s, a)\|_2 \\ &\leq 2UL_\Phi \|a' - a\|^\alpha, \end{aligned}$$

where the first inequality follows from the triangle inequality, the second follows from Assumption 1, the third follows from Hölder's inequality and the assumed bound on Ψ , and the final inequality follows from Hölder continuity of all $\phi \in \Phi$. The conclusion immediately follows. \square

A.4 Proof of Theorem 6

Proof of Theorem 6. Since $m \geq 1$, we cannot simply use Theorem 2 to replace the simple importance sampling bound applied by Agarwal et al. (2020), and conclude by simply replacing $|\mathcal{A}|$ in the result of Agarwal et al. (2020) with some complex term involving the constants c, L_E, α_E as in Assumption 2. Instead, we must carefully percolate the weaker bound given by Theorem 2 through the arguments of Agarwal et al. (2020, Appendix B). In this section, we briefly summarize this process.

Let us refer to Agarwal et al. (2020, Lemma x) by AKKS-Lx, for each x , and similarly AKKS-Cx for corollaries. We will briefly describe how updated versions of the lemmas and corollaries in Agarwal et al. (2020, Appendix B) can be obtained under our assumptions, by using Theorem 2 instead of the simple importance sampling bound. We sketch these arguments and results, with formal details omitted when the implications are straightforward. In these arguments we will assume without further comment that $\lambda \leq 1/d$.

AKKS-L17 First, we note that at termination we have

$$\begin{aligned} &\sup_{\pi} \mathbb{E} \left[\phi_{H-1}(s_{H-1}, a_{H-1})^T \Sigma_T^{-1} \phi_{H-1}(s_{H-1}, a_{H-1}) \mid \tilde{\mathcal{M}}, \pi \right] \\ &\leq \sup_{\pi} \mathbb{E} \left[\phi_{H-1}(s_{H-1}, a_{H-1})^T \Sigma_{T-1}^{-1} \phi_{H-1}(s_{H-1}, a_{H-1}) \mid \tilde{\mathcal{M}}, \pi \right] \\ &\leq \mathbb{E} \left[\phi_{H-1}(s_{H-1}, a_{H-1})^T \Sigma_{T-1}^{-1} \phi_{H-1}(s_{H-1}, a_{H-1}) \mid \tilde{\mathcal{M}}, \pi_T \right] + \frac{\beta}{2} \\ &\leq \beta, \end{aligned}$$

where the first inequality follows because $\Sigma_T - \Sigma_{T-1}$ is positive semi-definite (PSD), the second equality follows because inequality follows due to the approximate optimization condition, and the third follows from the termination condition. Therefore, since $\Sigma_\rho = \frac{1}{T} \sum_{t=1}^T \Sigma_{\pi_t}$ and $\Sigma_T = \sum_{t=1}^T \Sigma_{\pi_t} + I$, we have

$$\forall \pi : \mathbb{E} \left[\phi_{H-1}(s_{H-1}, a_{H-1})^\top (\Sigma_\rho + I/T)^{-1} \phi_{H-1}(s_{H-1}, a_{H-1}) \mid \widetilde{\mathcal{M}}, \pi \right] \leq T\beta.$$

Finally, we note that since at every non-terminating iteration the objective is at least $\beta/2$, we can apply an identical potential-based argument as in the original proof of AKKS-L17, except with β replaced with $\beta/2$, which gives us the bound

$$T \leq 8d \log(1 + 8/\beta)/\beta.$$

That is, by using our alternative elliptical planner that only approximately optimizes the objective each iteration, we obtain the same guarantee as in AKKS-L17, except the bound on T is a constant-factor weaker.

AKKS-C5 The statement and proof of this corollary is identical since no aspect of it relies on \mathcal{A} being discrete.

AKKS-L6 For this lemma, we can obtain a similar result, by relying on both Assumption 2 and Theorem 2 instead of the simple importance sampling bound. We simplify notation by letting $U(\alpha_E, L_E) := cL_E^{\frac{m}{m+\alpha_E}}$ where c is the constant in Theorem 2. Thus, we can obtain the bound

$$\begin{aligned} & \mathbb{E}_{s_h \sim P_h^\pi, a_h \sim \pi(s_h)} [\mathbf{1}\{s_h \neq s_{\text{absorb}}\} \text{err}(s_h, a_h)] \\ & \leq \mathbb{E}_{s_h \sim P_h^\pi} \left[\mathbf{1}\{s_h \neq s_{\text{absorb}}\} U(\alpha_E, L_E) \mathbb{E}_{a_h \sim \text{unif}_{\mathcal{A}}} [\text{err}(s_h, a_h)]^{\alpha_E/(m+\alpha)} \right] \\ & \leq U(\alpha_E, L_E) \mathbb{E}_{s_h \sim P_h^\pi, a_h \sim \text{unif}_{\mathcal{A}}} [\mathbf{1}\{s_h \neq s_{\text{absorb}}\} \text{err}(s_h, a_h)]^{\alpha_E/(m+\alpha_E)} \\ & \leq U(\alpha_E, L_E) \epsilon_{\text{TV}}^{\alpha_E/(2(m+\alpha_E))}, \end{aligned}$$

where the first inequality follows from Assumption 2, the second follows from Jensen's inequality, and the last following the same logic as in AKKS-L6. Therefore, following the remaining steps as in AKKS-L6, we obtain the same result except with $U(\alpha_E, L_E) \epsilon_{\text{TV}}^{\alpha_E/(2(m+\alpha_E))}$ in place of $K\sqrt{\epsilon_{\text{TV}}}$.

AKKS-L7 and AKKS-C8 Plugging in the previous result from AKKS-L6, we obtain an almost-identical results, except with $K\sqrt{\epsilon_{\text{TV}}}$ replaced everywhere by $U(\alpha_E, L_E) \epsilon_{\text{TV}}^{\alpha_E/(2(m+\alpha_E))}$. In particular, this gives us a bound on the maximum number of rounds

$$J_{\max} = \frac{4Hd}{\lambda U(\alpha_E, L_E) \epsilon_{\text{TV}}^{\alpha_E/(2(m+\alpha_E))}} \log \left(1 + \frac{4H}{\lambda U(\alpha_E, L_E) \epsilon_{\text{TV}}^{\alpha_E/(2(m+\alpha_E))}} \right),$$

in order to obtain a round where

$$\mathbb{P}[(s_h, a_h) \notin \mathcal{K}(\Sigma_{j,j}) \mid \pi, \widetilde{\mathcal{M}}^{(j)}] \leq 2HU(\alpha_E, L_E) \epsilon_{\text{TV}}^{\alpha_E/(2(m+\alpha_E))}$$

AKKS-L9 We can obtain a very similar lemma here, except that we need to interpret K differently. Here, we will interpret K as some constant such that the density ratio $\pi(a \mid s)/\text{unif}_{\mathcal{A}}(a)$ is bounded by K for every s, a . Then, the same result follows using an identical argument involving importance sampling, except that the value K in the bound is π -dependent. We note that this implicitly requires π to be non-deterministic.

AKKS-L10 We will obtain a similar lemma here, except that the lemma will only apply to policies π that have a bounded density ratio $\sup_{s,a} \pi(a \mid s)/\text{unif}_{\mathcal{A}}(a) \leq K$ for some K , and the bounds will be K -dependent.

First, applying successively our updated versions of AKKS-L6, AKKS-L9, and AKKS-C8, we have

$$\begin{aligned} & \mathbb{P}[(s_h, a_h) \notin \mathcal{K}(\Sigma_{j,j}) \mid \pi, \widetilde{\mathcal{M}}^{(j)}] \\ & \leq \alpha H^2 K U(\alpha_E, L_E) \epsilon_{\text{TV}}^{\alpha_E/(2(m+\alpha_E))} + \frac{T\beta}{2\alpha} + \frac{\alpha d}{2T} + H U(\alpha_E, L_E) \epsilon_{\text{TV}}^{\alpha_E/(2(m+\alpha_E))} \triangleq \epsilon_{\text{escape}}. \end{aligned}$$

Then, tracing the rest of the proof with the updated version of AKKS-L6, we get a final result of

$$\max_{\pi, R} \left| V(\pi; R, \widetilde{\mathcal{M}}) - V(\pi; R, \mathcal{M}) \right| \leq H U(\alpha_E, L_E) \epsilon_{\text{TV}}^{\alpha_E/(2(m+\alpha_E))} + H \epsilon_{\text{escape}}.$$

Final steps. We put together the updated lemmas and corollaries to obtain the final K -dependent bound on the sample complexity in order to accurately evaluate all policies π that satisfy $\sup_{s,a} \pi(a \mid s) / \text{unif}_{\mathcal{A}}(a) \leq K$. To simplify our notation, we introduce the following abbreviations: $\tau = m/\alpha_E$ and $\kappa = m/(m + \alpha_E)$.

First, we note that from AKKS-10, the policy evaluation error is given by

$$\begin{aligned} \text{error} &= \alpha H^3 K U(\alpha_E, L_E) \epsilon_{\text{TV}}^{\alpha_E/(2(m+\alpha_E))} + \frac{HT\beta}{2\alpha} + \frac{H\alpha d}{2T} + (H^2 + H) U(\alpha_E, L_E) \epsilon_{\text{TV}}^{\alpha_E/(2(m+\alpha_E))} \\ &= \alpha H^3 K U(\alpha_E, L_E) \epsilon_{\text{TV}}^{1/(2(1+\tau))} + \frac{HT\beta}{2\alpha} + \frac{H\alpha d}{2T} + (H^2 + H) U(\alpha_E, L_E) \epsilon_{\text{TV}}^{1/(2(1+\tau))} \\ &\leq \alpha H^3 K U(\alpha_E, L_E) \epsilon_{\text{TV}}^{1/(2(1+\tau))} + \frac{HT\beta}{2\alpha} + \frac{H\alpha d}{2T} + 2H^2 U(\alpha_E, L_E) \epsilon_{\text{TV}}^{1/(2(1+\tau))}, \end{aligned}$$

Next, optimizing over α in this bound and setting

$$\alpha = \sqrt{\frac{HT\beta}{2} \left(H^3 K U(\alpha_E, L_E) \epsilon_{\text{TV}}^{1/(2(1+\tau))} + \frac{Hd}{2T} \right)^{-1}}$$

gives us

$$\text{error} \leq \sqrt{\left(2H^4 K T \beta U(\alpha_E, L_E) \epsilon_{\text{TV}}^{1/(2(1+\tau))} + H^2 d \beta \right) + 2H^2 U(\alpha_E, L_E) \epsilon_{\text{TV}}^{1/(2(1+\tau))}}.$$

Next, we note that β is a free parameter of the algorithm, which we may set arbitrarily low at the cost of greater sample complexity. We balance the first term by setting β according to

$$\beta = \frac{2H^2 K U(\alpha_E, L_E) \epsilon_{\text{TV}}^{1/(2(1+\tau))}}{d},$$

which gives us

$$\text{error} \leq \sqrt{2} H^2 \sqrt{T\beta + 1} \sqrt{K U(\alpha_E, L_E) \epsilon_{\text{TV}}^{1/(4(1+\tau))}} + 2H^2 U(\alpha_E, L_E) \epsilon_{\text{TV}}^{1/(2(1+\tau))}.$$

We note that $U(\alpha_E, L_E) + \sqrt{U(\alpha_E, L_E)} \leq 2 \max(U(\alpha_E, L_E), 1)$. For now, assume that $U(\alpha_E, L_E) \geq 1$. We further apply the fact that $K \geq 1$, $\epsilon_{\text{TV}} \leq 1$, and $T\beta \leq 8d \log(1 + 8/\beta)$ in order to get

$$\text{error} \leq 8\sqrt{d} \sqrt{\log(1 + 8/\beta)} H^2 \sqrt{K U(\alpha_E, L_E) \epsilon_{\text{TV}}^{1/(4(1+\tau))}}.$$

Now, we would like to set ϵ_{TV} sufficiently small such that the error is less than ϵ . We can achieve this if we set

$$\epsilon_{\text{TV}} = 8^{-4-4\tau} \frac{H^{-8-8\tau} U(\alpha_E, L_E)^{-4-4\tau} K^{-2-2\tau} d^{-2-2\tau} \epsilon^{4+4\tau}}{\log(1 + 8/\beta')^{2+2\tau}},$$

as long as β' satisfies

$$\begin{aligned} \beta' &\leq \beta = \frac{2H^2 K U(\alpha_E, L_E) \epsilon_{\text{TV}}^{1/(2(1+\tau))}}{d} = \frac{\epsilon^2}{2^5 d^2 H^2 U(\lambda_T, L_E) \log(1 + 8/\beta')} \\ &\iff 8/\beta' \geq 2^8 d^2 H^2 U(\alpha_E, L_E) \epsilon^{-2} \log(1 + 8/\beta'). \end{aligned}$$

We note that $\log(1 + x) \leq x/\sqrt{x+1}$ for $x \geq 0$, so this condition on β' is clearly satisfied if we set

$$\beta' = \frac{8}{2^{16} d^4 H^4 U^2(\alpha_E, L_E) \epsilon^{-4} - 1}.$$

Next, we need to reason about how small λ and n must be set to achieve this. Recall that

$$\epsilon_{\text{TV}} \leq \lambda d + \frac{2 \log(J_{\max} H |\Phi| |\Psi| / \delta)}{n}.$$

This implies that we can achieve the required bound as long as we set

$$\lambda = \frac{8^{-4-4\tau}}{2} \frac{H^{-8-8\tau} U(\alpha_E, L_E)^{-4-4\tau} K^{-2-2\tau} d^{-3-2\tau} \epsilon^{4+4\tau}}{\log(1 + 8/\beta')^{2+2\tau}}$$

$$n = 8^{4+4\tau} \frac{H^{8+8\tau} U(\alpha_E, L_E)^{4+4\tau} K^{2+2\tau} d^{2+2\tau} \log(1+8/\beta')^{2+2\tau}}{\epsilon^{4+4\tau}} \log(J_{\max} H |\Phi||\Psi|/\delta)$$

Furthermore, recalling the definition of J_{\max} and disregarding the log terms, we have

$$\begin{aligned} J_{\max} &= \frac{4Hd}{\lambda U(\alpha_E, L_E) \epsilon_{\text{TV}}^{1/(2(1+\tau))}} \log \left(1 + \frac{4H}{\lambda U(\alpha_E, L_E) \epsilon_{\text{TV}}^{1/(2(1+\tau))}} \right) \\ &= \tilde{\mathcal{O}} \left(\frac{Hd}{\lambda U(\alpha_E, L_E) \epsilon_{\text{TV}}^{1/(2(1+\tau))}} \right) \\ &= \tilde{\mathcal{O}} \left(\frac{H^{13+8\tau} U(\alpha_E, L_E)^{5+4\tau} K^{3+2\tau} d^{5+2\tau}}{\epsilon^{6+4\tau}} \right). \end{aligned}$$

Finally, the total sample complexity is given by $nJ_{\max}H$. Plugging in the above values for n , λ , and ϵ_{TV} , and disregarding log terms, we have

$$nJ_{\max}H = \tilde{\mathcal{O}} \left(H^{22+16\tau} K^{5+4\tau} d^{7+4\tau} \left(\frac{1}{\epsilon} \right)^{10+8\tau} U(\alpha_E, L_E)^{9+8\tau} \ln \left(\frac{|\Phi||\Psi|}{\delta} \right) \right). \quad (3)$$

Furthermore, recalling that $U(\alpha_E, L_E) = cL_E^{\frac{m}{m+\alpha_E}} = cL_E^\kappa$, we obtain the desired sample complexity:

$$nJ_{\max}H = \tilde{\mathcal{O}} \left(H^{22+16\tau} K^{5+4\tau} d^{7+4\tau} \left(\frac{1}{\epsilon} \right)^{10+8\tau} L_E^{(9+8\tau)\kappa} \ln \left(\frac{|\Phi||\Psi|}{\delta} \right) \right). \quad (4)$$

□

A.5 Proof of Theorem 7

Proof of Theorem 7. The proof of this theorem is an immediate application of Theorem 6 for the smoothed, intermediary policy π_K as in Definition 6. First, note that we can decompose the error of the policy π in terms of errors of the policy π_K :

$$\begin{aligned} |V(\pi; R, \mathcal{M}) - V(\pi; R, \widehat{\mathcal{M}})| &\leq |V(\pi; R, \mathcal{M}) - V(\pi_K; R, \mathcal{M})| \\ &\quad + |V(\pi_K; R, \mathcal{M}) - V(\pi_K; R, \widehat{\mathcal{M}})| \\ &\quad + |V(\pi_K; R, \widehat{\mathcal{M}}) - V(\pi; R, \widehat{\mathcal{M}})| \end{aligned}$$

We can bound the first and third term on the RHS of the above inequality using the result of Lemma 3. To see this, let $\sigma = m/\min(\alpha_T, \alpha_R)$ and $L = \max(L_T, L_R)$. Then, by setting $K = \left(\frac{8\sqrt{m}HL}{\epsilon} \right)^\sigma$, Lemma 3 yields:

$$|V(\pi; R, \mathcal{M}) - V(\pi; R, \widehat{\mathcal{M}})| \leq |V(\pi_K; R, \mathcal{M}) - V(\pi_K; R, \widehat{\mathcal{M}})| + \frac{\epsilon}{2}$$

Since π_K has bounded density, we can control the remaining term on the RHS by leveraging the result of Theorem 6 with $K = \left(\frac{8\sqrt{m}HL}{\epsilon} \right)^\sigma$ and ϵ replaced with $\epsilon/2$. This would ensure that $|V(\pi_K; R, \mathcal{M}) - V(\pi_K; R, \widehat{\mathcal{M}})| \leq \epsilon/2$ with probability at least $1 - \delta$. Therefore, putting this together and plugging in this value of K into our result from Theorem 6, we get our desired performance result of

$$|V(\pi; R, \mathcal{M}) - V(\pi; R, \widehat{\mathcal{M}})| \leq \epsilon,$$

with a total number of trajectories collected given by

$$\tilde{\mathcal{O}} \left(H^{22+16\tau+(4\tau+5)\sigma} d^{7+4\tau} \left(\frac{1}{\epsilon} \right)^{10+8\tau+(4\tau+5)\sigma} L_E^{(9+8\tau)\kappa} L^{(4\tau+5)\sigma} \ln \left(\frac{|\Phi||\Psi|}{\delta} \right) \right),$$

where $\sigma = m/\min(\alpha_T, \alpha_R)$, and $L = \max(L_T, L_R)$. Furthermore, the corresponding values of the hyperparameters β , n , and J_{\max} that achieve these results are the same as those from the proof of Theorem 6, plugging in the above value of K . □

B CASE STUDY: RAFFLE

In this section, we apply our framework to RAFFLE (Cheng et al., 2023), a recently proposed reward-free low-rank MDP algorithm which was shown to achieve near optimal sample complexity $\tilde{O}\left(\frac{H^3 d^2 K(d^2+K)}{\epsilon^2}\right)$ (Theorem 2 in Cheng et al. (2023)). The significant improvement in sample complexity over FLAMBE is attributed to a novel exploration-driven pseudo-reward mechanism which targets sample collection towards high-error regions.

We now seek to derive a similar sample complexity for the RAFFLE algorithm when the action space is continuous, specifically $\mathcal{A} = [0, 1]^m$. We first note that the size K of the discrete action space in RAFFLE enters the bound in their Theorem 2 via the application of importance sampling (IS) bounds, as in our Lemma 1. The IS bound is applied with two different choices of $f(s, a)$: (1) the (squared) TV error of transition distributions (as in our Assumption 2), denoted by $f_h^{(n)}(s, a)$; and (2) their (squared) bonus reward function terms, given by $\hat{b}_h^{(n)}(s, a)$. Thus, in order to provide a similar bound, it suffices to make $f_h^{(n)}(s, a)$ and $\hat{b}_h^{(n)}(s, a)$ α -smooth in a and replace the IS bounds with the bound in our Theorem 2.

The TV error smoothness assumption is already outlined in Assumption 2 which we take to hold for RAFFLE as well. Next, we see that the bonus reward function is the matrix norm of $\phi_h^{(n)}(s, a)$ induced by a symmetric matrix with eigenvalues smaller than λ_n^{-1} . Thus, we introduce the following assumption:

Assumption 4 (α -smooth feature vectors). *There exists $\alpha_\Phi \in (0, \infty)$ such that the scalar components $\phi_i(s, a)$, $i \in [d]$ are α_Φ -smooth in $a \in \mathcal{A}$ with norm L_Φ for every $\phi \in \Phi$ and $s \in \mathcal{S}$.*

Assumption 4 is similar to the assumption that enables smooth TV errors in Lemma 5. However, Assumption 4 is more general in that it allows for higher degrees of smoothness when $\alpha_\Phi \geq 1$.

Remark 5. *RAFFLE applied IS bounds to $f(s, a)^2 = \|\phi(s, a)^\top \psi(\cdot) - \phi_h^*(s, a)^\top \psi_h^*(\cdot)\|_{\text{TV}}^2$ and $g(s, a)^2 = \|\hat{\phi}_h^{(n)}(s, a)\|_U^2$, involving sums over $U_{ji} \hat{\phi}_{h,i}^{(n)}(s, a) \hat{\phi}_{h,j}^{(n)}(s, a)$. While Assumption 2 and Assumption 4 only impose smoothness on $f(s, a)$ and each of the $\hat{\phi}_{h,i}^{(n)}(s, a)$'s, respectively, the space of α -smooth functions as defined in Definition 2 is closed under multiplication (Behzadan and Holst, 2021, Theorem 4.2 and Proposition 4.3), making $f(s, a)^2$ α_E -smooth and $\hat{\phi}_{h,i}^{(n)}(s, a) \hat{\phi}_{h,j}^{(n)}(s, a)$ α_Φ -smooth. In the remainder of this section, we extend our notation, denoting the α -smooth norms of $f^2(s, a)$ and the $\hat{\phi}_{h,i}^{(n)}(s, a) \hat{\phi}_{h,j}^{(n)}(s, a)$ products by L_E and L_Φ , respectively.*

With Assumption 2 and Assumption 4, we can replace the IS bounds with the bound in our Theorem 2 and carefully propagate these arguments through sample complexity proof in Cheng et al. (2023, Appendix A). In the remainder of this section, we sketch this process and results, with formal details omitted for brevity. We provide the sample complexity of the RAFFLE algorithm with continuous action spaces in the following theorem:

Theorem 11 (RAFFLE guarantees for continuous \mathcal{A}). *Assume \mathcal{M} has rank d and Assumption 1, Assumption 2, and Assumption 4 hold. Let $\alpha = \min(\alpha_E, \alpha_\Phi)$, $L = \max(L_E, L_\Phi)$, and $\tau = m/(m + \alpha)$. Then, if $\tau < 1/2$ and we set $\hat{\alpha}_n = \tilde{O}\left(\sqrt{n^\tau (L^\tau + (dL)^\tau + d^2)}\right)$, $\lambda_n = \tilde{O}(d)$, with probability at least $1 - \delta$, RAFFLE returns a model $\widehat{\mathcal{M}}$ that satisfies*

$$|V(\pi; R, \widehat{\mathcal{M}}) - V(\pi; R, \mathcal{M})| \leq \epsilon$$

for every sparse reward function R , and every policy π . The total number of trajectories collected by RAFFLE is upper bounded by

$$\tilde{O}\left(H^{1+\frac{2}{1-2\tau}} d^{\frac{2}{1-2\tau}} \left(\frac{1}{\epsilon}\right)^{\frac{2}{1-2\tau}} L^{\frac{\tau}{1-2\tau}} (L^\tau + (dL)^\tau + d^2)^{\frac{1}{1-2\tau}}\right).$$

We note that as $\tau \rightarrow 0$ (indicating that ϵ and $\phi(s, a)$ have a very high order of smoothness relative to m) we achieve identical dependence on H , d , and $1/\epsilon$ as in Cheng et al. (2023).

Remark 6. *Although this theorem is not applicable to ϵ and $\phi(s, a)$ with a lower degree of smoothness ($\tau \geq 1/2$), it is worth noting that a more refined bound could potentially be achieved through an alternative proof strategy, instead of naively replacing the IS arguments in Cheng et al. (2023)'s proof with our Theorem 2. We leave this for future work.*

Proof of Theorem 11. Similarly to the FLAMBE analysis, we refer to Cheng et al. (2023, Lemma x , Appendix A) by CHYL-L x , for each x , and similarly CHYL-C x and CHYL-P x for corollaries and propositions, respectively. Whenever possible, we keep notation consistent with that in Cheng et al. (2023).

CHYL-L1 The statement and proof of this lemma remains unchanged, as its validity does not depend on \mathcal{A} being discrete.

CHYL-C1 Let $\alpha = \min(\alpha_E, \alpha_\Phi)$ and $L = \max(L_E, L_\Phi)$. Setting

$$\hat{\alpha}_n = 5\sqrt{\beta_3 n \zeta_n^{\frac{\alpha}{m+\alpha}} (L^{\frac{m}{m+\alpha}} + (dL)^{\frac{m}{m+\alpha}} + d^2)}$$

in the definition of the bonus reward term $\hat{b}_h(s, a)$, the statement and result of this lemma stay the same.

CHYL-L2 We add an additional condition to the statement of the lemma that g^2 is α -smooth with norm L . Then, for all $h \geq 2$ and for all policies π_h , there exists some $c = O(1)$ such that the following inequality holds:

$$\begin{aligned} & \mathbb{E}_{\substack{s_h \sim P_{h-1} \\ a_h \sim \pi_h}} [g(s_h, a_h) | s_{h-1}, a_{h-1}] \\ & \leq \|\phi_{h-1}(s_{h-1}, a_{h-1})\|_{(M_{h-1, \phi})^{-1}} \times \\ & \quad \sqrt{cnL^{\frac{m}{m+\alpha}} \mathbb{E}_{\substack{s_h \sim (P^*, \Pi) \\ a_h \sim \mathcal{U}}} [g^2(s_h, a_h)]^{\frac{\alpha}{m+\alpha}} + \lambda_n dB^2 + nB^2 \mathbb{E}_{\substack{s_{h-1} \sim (P^*, \Pi) \\ a_{h-1} \sim \Pi}} [f_{h-1}(s_{h-1}, a_{h-1})^2]}. \end{aligned}$$

Here, $\mathcal{U} = \text{unif}_{\mathcal{A}}$. To obtain this result, we replace the importance sampling argument in step (iv) of the proof in (Cheng et al., 2023) with the result in our Theorem 2.

CHYL-L3 We now propagate the result of CHYL-L2 to obtain new bounds for Eq. (11-14) in Cheng et al. (2023, Appendix A). We first note that for any h , we have

$$\begin{aligned} & \mathbb{E}_{\substack{s_{h-1} \sim P_{h-1}^* (\cdot | s_{h-2}, a_{h-2}) \\ a_{h-1} \sim \Pi_n(s_{h-1})}} \left[f_{h-1}^{(n)}(s_{h-1}, a_{h-1})^2 \mid s_{h-2}, a_{h-2} \right] \\ & = \int_{s_{h-1}} \int_{a_{h-1}} f^{(n)}(s_{h-1}, a_{h-1})^2 \phi_{h-2}^*(s_{h-2}, a_{h-2})^\top \mu_{h-2}^*(s_{h-1}) \Pi_n(a_{h-1} \mid s_{h-1}) da_{h-1} ds_{h-1} \\ & = \sum_{i=1}^d \Omega_i \phi_{h-2, i}^*(s_{h-2}, a_{h-2}), \end{aligned} \tag{5}$$

where

$$\Omega_i = \int_{s_{h-1}} \int_{a_{h-1}} f^{(n)}(s_{h-1}, a_{h-1})^2 \mu_{h-2, i}^*(s_{h-1}) \Pi_n(a_{h-1} \mid s_{h-1}) da_{h-1} ds_{h-1}.$$

Now, by the normalization assumption on μ_{h-2}^* , and the fact that $f^{(n)}$ is uniformly bounded by 1, it easily follows that $\|\Omega\|_2 \leq \sqrt{d}$, and therefore $\|\Omega\|_1 \leq d$. Therefore, given Assumption 4, it easily follows that (5) is uniformly α_Φ -smooth in a_{h-2} with norm at most dL_Φ . Then, applying Theorem 2 it follows that

$$\begin{aligned} & \mathbb{E}_{\substack{s_{h-1} \sim (P^*, \Pi_n) \\ a_{h-1} \sim \mathcal{U}}} \left[f_{h-1}^{(n)}(s_{h-1}, a_{h-1})^2 \right] \\ & = \mathbb{E}_{\substack{s_{h-2} \sim (P^*, \Pi_n) \\ a_{h-2} \sim \Pi_n(s_{h-1})}} \left[\mathbb{E}_{\substack{s_{h-1} \sim P_{h-1}^* (\cdot | s_{h-2}, a_{h-2}) \\ a_{h-1} \sim \Pi_n(s_{h-1})}} \left[f_{h-1}^{(n)}(s_{h-1}, a_{h-1})^2 \mid s_{h-2}, a_{h-2} \right] \right] \\ & \leq d^{\frac{m}{m+\alpha_\Phi}} L_\Phi^{\frac{m}{m+\alpha_\Phi}} \mathbb{E}_{\substack{s_{h-2} \sim (P^*, \Pi_n) \\ a_{h-2} \sim \mathcal{U}}} \left[\mathbb{E}_{\substack{s_{h-1} \sim P_{h-1}^* (\cdot | s_{h-2}, a_{h-2}) \\ a_{h-1} \sim \Pi_n(s_{h-1})}} \left[f_{h-1}^{(n)}(s_{h-1}, a_{h-1})^2 \mid s_{h-2}, a_{h-2} \right] \right]^{\frac{\alpha_\Phi}{m+\alpha_\Phi}} \end{aligned}$$

$$= d^{\frac{m}{m+\alpha_\Phi}} L_\Phi^{\frac{m}{m+\alpha_\Phi}} \mathbb{E}_{\substack{s_{h-2} \sim (P^*, \Pi_n) \\ a_{h-2}, a_{h-1} \sim \mathcal{U} \\ s_{h-1} \sim P_{h-1}^*(\cdot | s_{h-2}, a_{h-2})}} \left[f_{h-1}^{(n)}(s_{h-1}, a_{h-1})^2 \right]^{\frac{\alpha_\Phi}{m+\alpha_\Phi}}$$

Given this observation, we can follow the steps as in the original proof of CHYL-L3, using the the modified CHYL-L2 above, as well replacing the importance sampling argument at time step $h-2$ with Theorem 2 using the above observation. Specifically, we have

$$\begin{aligned} & \mathbb{E}_{\substack{s_h \sim P_{h-1}^{(n)} \\ a_h \sim \pi}} \left[f_h^{(n)}(s_h, a_h) \middle| s_{h-1}, a_{h-1} \right] \\ & \leq \left\| \hat{\phi}_{h-1}^{(n)}(s_{h-1}, a_{h-1}) \right\|_{(U_{h-1, \hat{\phi}}^{(n)})^{-1}} \times \\ & \quad \sqrt{c n L_E^{\frac{m}{m+\alpha_E}} \mathbb{E}_{\substack{s_{h-1} \sim (P^*, \Pi_n) \\ a_{h-1}, a_h \sim \mathcal{U} \\ s_h \sim P_h^*}} \left[f_h^{(n)}(s_h, a_h)^2 \right]^{\frac{\alpha_E}{m+\alpha_E}} + \lambda_n d + n(dL_\Phi)^{\frac{m}{m+\alpha_\Phi}} \mathbb{E}_{\substack{s_{h-2} \sim (P^*, \Pi_n) \\ a_{h-2}, a_{h-1} \sim \mathcal{U} \\ s_{h-1} \sim P_{h-1}^*}} \left[f_{h-1}^{(n)}(s_{h-1}, a_{h-1})^2 \right]^{\frac{\alpha_\Phi}{m+\alpha_\Phi}}} } \quad (\text{above bound and CHYL-L2}) \\ & \leq \left\| \hat{\phi}_{h-1}^{(n)}(s_{h-1}, a_{h-1}) \right\|_{(U_{h-1, \hat{\phi}}^{(n)})^{-1}} \sqrt{c n L_E^{\frac{m}{m+\alpha_E}} \zeta_n^{\frac{\alpha_E}{m+\alpha_E}} + c n (dL_\Phi)^{\frac{m}{m+\alpha_\Phi}} \zeta_n^{\frac{\alpha_\Phi}{m+\alpha_\Phi}} + \beta_3 n \zeta_n d^2} \quad (\text{CHYL-L1 and } \lambda_n = \beta_3 n d \zeta_n) \\ & \leq \left\| \hat{\phi}_{h-1}^{(n)}(s_{h-1}, a_{h-1}) \right\|_{(U_{h-1, \hat{\phi}}^{(n)})^{-1}} \sqrt{\beta_3 n \left(L_E^{\frac{m}{m+\alpha_E}} \zeta_n^{\frac{\alpha_E}{m+\alpha_E}} + (dL_\Phi)^{\frac{m}{m+\alpha_\Phi}} \zeta_n^{\frac{\alpha_\Phi}{m+\alpha_\Phi}} + \zeta_n d^2 \right)} \\ & \leq \alpha_n \left\| \hat{\phi}_{h-1}^{(n)}(s_{h-1}, a_{h-1}) \right\|_{(U_{h-1, \hat{\phi}}^{(n)})^{-1}} \end{aligned}$$

where we set $\alpha_n = \hat{\alpha}_n/5 = \sqrt{\beta_3 n \zeta_n^{\frac{\alpha}{m+\alpha}} (L_E^{\frac{m}{m+\alpha}} + (dL)^{\frac{m}{m+\alpha}} + d^2)}$. Note that without loss of generality we can let the constant c from Theorem 2 be absorbed into β_3 since they are both $O(1)$.

Eq. (12) follows from similar arguments such that

$$\mathbb{E}_{\substack{s_h \sim P_{h-1}^* \\ a_h \sim \pi}} \left[f_h^{(n)}(s_h, a_h) \middle| s_{h-1}, a_{h-1} \right] \leq \alpha_n \left\| \phi_{h-1}^*(s_{h-1}, a_{h-1}) \right\|_{(U_{h-1, \phi^*})^{-1}}.$$

For Eq. (13), we have $\hat{b}_h^{(n)}(s, a) = \min \left\{ \hat{\alpha}_n \left\| \hat{\phi}_h^{(n)}(s, a) \right\|_{(\hat{U}_h^{(n)})^{-1}}, 1 \right\} \leq 3\hat{\alpha}_n \left\| \hat{\phi}_h^{(n)}(s, a) \right\|_{(U_{h, \hat{\phi}}^{(n)})^{-1}}$ from CHYL-L1.

We can then apply CHYL-L2 (before the IS argument):

$$\mathbb{E}_{\substack{s_h \sim P_{h-1}^* \\ a_h \sim \pi}} \left[\hat{b}_h^{(n)}(s_h, a_h) \middle| s_{h-1}, a_{h-1} \right] \leq \left\| \phi_{h-1}^*(s_{h-1}, a_{h-1}) \right\|_{(W_{h-1, \phi^*})^{-1}} \sqrt{n \mathbb{E}_{\substack{s_h \sim (P^*, \Pi_n) \\ a_h \sim \pi_n}} [\{\hat{b}_h^{(n)}(s_h, a_h)\}^2] + \lambda_n d}.$$

We note that $\hat{b}_h^{(n)}(s, a)$ is not α -smooth and thus we cannot directly apply our Theorem 2 results here. Instead, we will have to bound this term by $\left\| \hat{\phi}_h^{(n)}(s, a) \right\|_{(U_{h, \hat{\phi}}^{(n)})^{-1}}$ which we now show is α_Φ -smooth with norm at most L_Φ :

$$\begin{aligned} \left\| \hat{\phi}_h^{(n)}(s, a) \right\|_{(U_{h, \hat{\phi}}^{(n)})^{-1}}^2 &= \text{tr} \left\{ \hat{\phi}_h^{(n)}(s_h, a_h) \hat{\phi}_h^{(n)}(s_h, a_h)^T (U_{h, \hat{\phi}}^{(n)})^{-1} \right\} \\ &= \sum_{i=1}^d \sum_{j=1}^d \{ (U_{h, \hat{\phi}}^{(n)})^{-1} \}_{ji} \hat{\phi}_{h,i}^{(n)}(s_h, a_h) \hat{\phi}_{h,j}^{(n)}(s_h, a_h) \\ &\leq \frac{1}{\lambda_n} \sum_{i=1}^d \sum_{j=1}^d \hat{\phi}_{h,i}^{(n)}(s_h, a_h) \hat{\phi}_{h,j}^{(n)}(s_h, a_h) \quad (\text{from the definition of } (U_{h, \hat{\phi}}^{(n)})^{-1}) \end{aligned}$$

By Assumption 4, we have that $\|\hat{\phi}_h^{(n)}(s, a)\|_{(U_{h, \hat{\phi}}^{(n)})^{-1}}^2$ is the sum of d^2 α_Φ -smooth functions with norm at most L_Φ/λ_n . Since the set of α -smooth function is closed under addition, it immediately follows that $\|\hat{\phi}_h^{(n)}(s, a)\|_{(U_{h, \hat{\phi}}^{(n)})^{-1}}^2$ is α_Φ -smooth with norm at most $d^2 L_\Phi/\lambda_n = dL_\Phi/(\beta_3 \log(2|\Phi||\Psi|nH/\delta)) < dL_\Phi$. Then, we can control the $\hat{b}_h^{(n)}(s, a)$ expression as follows:

$$\begin{aligned}
 n \mathbb{E}_{\substack{s_h \sim (P^\star, \Pi_n) \\ a_h \sim \pi_n}} [\{\hat{b}_h^{(n)}(s_h, a_h)\}^2] &\leq 9n\hat{\alpha}_n^2 \mathbb{E}_{\substack{s_h \sim (P^\star, \Pi_n) \\ a_h \sim \pi_n}} \left[\left\| \hat{\phi}_h^{(n)}(s_h, a_h) \right\|_{(U_{h, \hat{\phi}}^{(n)})^{-1}}^2 \right] \\
 &\leq 9n\hat{\alpha}_n^2 c(dL_\Phi)^{\frac{m}{m+\alpha_\Phi}} \mathbb{E}_{\substack{s_h \sim (P^\star, \Pi_n) \\ a_h \sim \mathcal{U}}} \left[\left\| \hat{\phi}_h^{(n)}(s_h, a_h) \right\|_{(U_{h, \hat{\phi}}^{(n)})^{-1}}^2 \right]^{\frac{\alpha_\Phi}{m+\alpha_\Phi}} \\
 &\quad \text{(from Theorem 2 on } \left\| \hat{\phi}_h^{(n)}(s_h, a_h) \right\|_{(U_{h, \hat{\phi}}^{(n)})^{-1}}^2 \text{)} \\
 &\leq 9c\hat{\alpha}_n^2 (dL_\Phi)^{\frac{m}{m+\alpha_\Phi}} n^{\frac{m}{m+\alpha_\Phi}} \mathbb{E}_{\substack{s_h \sim (P^\star, \Pi_n) \\ a_h \sim \mathcal{U}}} \left[n \left\| \hat{\phi}_h^{(n)}(s_h, a_h) \right\|_{(U_{h, \hat{\phi}}^{(n)})^{-1}}^2 \right]^{\frac{\alpha_\Phi}{m+\alpha_\Phi}} \\
 &\leq 9c\hat{\alpha}_n^2 (dL_\Phi)^{\frac{m}{m+\alpha_\Phi}} n^{\frac{m}{m+\alpha_\Phi}} d^{\frac{\alpha_\Phi}{m+\alpha_\Phi}} \\
 &= 9c\hat{\alpha}_n^2 dL_\Phi^{\frac{m}{m+\alpha_\Phi}} n^{\frac{m}{m+\alpha_\Phi}} \quad \text{(same argument as the original CHYL-L3)}
 \end{aligned}$$

Putting everything together, we obtain

$$\begin{aligned}
 &\mathbb{E}_{\substack{s_h \sim P_{h-1}^\star \\ a_h \sim \pi_n}} \left[\hat{b}_h^{(n)}(s_h, a_h) \middle| s_{h-1}, a_{h-1} \right] \\
 &\leq \|\phi_{h-1}^\star(s_{h-1}, a_{h-1})\|_{(W_{h-1, \phi^\star}^{(n)})^{-1}} \sqrt{9c\hat{\alpha}_n^2 dL_\Phi^{\frac{m}{m+\alpha_\Phi}} n^{\frac{m}{m+\alpha_\Phi}} + \lambda_n d} \\
 &\leq \|\phi_{h-1}^\star(s_{h-1}, a_{h-1})\|_{(W_{h-1, \phi^\star}^{(n)})^{-1}} \sqrt{225\beta_3 n^{1+\frac{m}{m+\alpha_\Phi}} \zeta_n^{\frac{\alpha}{m+\alpha}} L_\Phi^{\frac{m}{m+\alpha_\Phi}} d (L^{\frac{m}{m+\alpha}} + (dL)^{\frac{m}{m+\alpha}} + d^2) + \beta_3 n \zeta_n d^2} \\
 &\leq \|\phi_{h-1}^\star(s_{h-1}, a_{h-1})\|_{(W_{h-1, \phi^\star}^{(n)})^{-1}} \sqrt{450\beta_3 n^{1+\frac{m}{m+\alpha}} \zeta_n^{\frac{\alpha}{m+\alpha}} L^{\frac{m}{m+\alpha}} d (L^{\frac{m}{m+\alpha}} + (dL)^{\frac{m}{m+\alpha}} + d^2)} \\
 &\leq \gamma_n \|\phi_{h-1}^\star(s_{h-1}, a_{h-1})\|_{(W_{h-1, \phi^\star}^{(n)})^{-1}},
 \end{aligned}$$

where we set $\gamma_n = \sqrt{450\beta_3 n^{1+\frac{m}{m+\alpha}} \zeta_n^{\frac{\alpha}{m+\alpha}} L^{\frac{m}{m+\alpha}} d (L^{\frac{m}{m+\alpha}} + (dL)^{\frac{m}{m+\alpha}} + d^2)}$. We note that this coefficient has a similar to the original γ_n , but with the first K replaced by $L^{\frac{m}{m+\alpha}}$, the second K replaced by $L^{\frac{m}{m+\alpha}}$ and different exponents for ζ_n and n .

For $h = 1$, we similarly have:

$$\begin{aligned}
 \mathbb{E}_{a_1 \sim \pi_n} [f_1^{(n)}(s_1, a_1)] &\leq \sqrt{cL_E^{\frac{m}{m+\alpha_E}} \mathbb{E}_{a_1 \sim \mathcal{U}} [f_1^{(n)}(s_1, a_1)^2]^{\frac{\alpha_E}{m+\alpha_E}}} \leq \sqrt{cL_E^{\frac{m}{m+\alpha_E}} \zeta_n^{\frac{\alpha_E}{m+\alpha_E}}} \\
 &\leq \sqrt{cL^{\frac{m}{m+\alpha}} \zeta_n^{\frac{\alpha}{m+\alpha}}}, \\
 \mathbb{E}_{a_1 \sim \pi_n} [\hat{b}(s_1, a_1)] &\leq 3\hat{\alpha}_n \sqrt{\mathbb{E}_{a_1 \sim \pi_n} \left[\left\| \hat{\phi}_1^{(n)}(s_1, a_1) \right\|_{(U_{1, \hat{\phi}}^{(n)})^{-1}}^2 \right]} \\
 &\leq 15\alpha_n \sqrt{cL_\Phi^{\frac{m}{m+\alpha_\Phi}} \mathbb{E}_{a_1 \sim \mathcal{U}} \left[\left\| \hat{\phi}_1^{(n)}(s_1, a_1) \right\|_{(U_{1, \hat{\phi}}^{(n)})^{-1}}^2 \right]^{\frac{\alpha_\Phi}{m+\alpha_\Phi}}} \\
 &\leq 15\alpha_n \sqrt{c \frac{L_\Phi^{\frac{m}{m+\alpha_\Phi}}}{n^{\frac{\alpha_\Phi}{m+\alpha_\Phi}}} d^{\frac{\alpha_\Phi}{m+\alpha_\Phi}}} \quad \text{(Theorem 2)}
 \end{aligned}$$

$$\leq 15\alpha_n \sqrt{cd \frac{L_{\Phi}^{\frac{m}{m+\alpha_\Phi}}}{n^{\frac{\alpha_\Phi}{m+\alpha_\Phi}}}} \leq 15\alpha_n \sqrt{cd \frac{L^{\frac{m}{m+\alpha}}}{n^{\frac{\alpha}{m+\alpha}}}}$$

This concludes the updated CHYL-L3. We propagate these results through Propositions 4-6 of Cheng et al. (2023).

CHYL-P4 For all $n \in [N]$, policy π and reward r , given that the event \mathcal{E} occurs, we obtain the inequality:

$$\left| V_{P^*,r}^\pi - V_{\hat{P}^{(n)},r}^\pi \right| \leq \hat{V}_{\hat{P}^{(n)},\hat{b}^{(n)}}^\pi + \sqrt{cL^{\frac{m}{m+\alpha}} \zeta_n^{\frac{\alpha}{m+\alpha}}}.$$

The proof of this proposition only changes when bounding $\mathbb{E}_{a_1 \sim \pi_n} [f_1^{(n)}(s_1, a_1)]$ which we instead bound using the updated CHYL-L3 for $h = 1$. This concludes the proof.

CHYL-P5 Replacing the IS arguments and the statements of CHYL-L3 and CHYL-P4 with the updated ones, we obtain:

$$\begin{aligned} V_{P^*,\hat{b}^{(n)}}^{\pi_n} &\leq \sum_{h=1}^H \mathbb{E}_{\substack{s_h \sim (P^*, \pi_n) \\ a_h \sim \pi_n}} \left[\alpha_n \|\phi_h^*(s_h, a_h)\|_{(U_{h,\phi^*}^{(n)})^{-1}} \right] + \sqrt{cL^{\frac{m}{m+\alpha}} \zeta_n^{\frac{\alpha}{m+\alpha}}}, \\ V_{P^*,f^{(n)}}^{\pi_n} &\leq \sum_{h=1}^H \mathbb{E}_{\substack{s_h \sim (P^*, \pi_n) \\ a_h \sim \pi_n}} \left[\alpha_n \|\phi_h^*(s_h, a_h)\|_{(U_{h,\phi^*}^{(n)})^{-1}} \right] + 15\alpha_n \sqrt{cd \frac{L^{\frac{m}{m+\alpha}}}{n^{\frac{\alpha}{m+\alpha}}}}. \end{aligned}$$

Taking the sum of these terms and noting that $\zeta = N\zeta_N$, we have:

$$\begin{aligned} &\sum_{n \in [N]} V_{P^*,f^{(n)}+\hat{b}^{(n)}}^{\pi_n} + \sqrt{cL^{\frac{m}{m+\alpha}} \zeta_n^{\frac{\alpha}{m+\alpha}}} \\ &\leq 15\alpha_N N \sqrt{cd \frac{L^{\frac{m}{m+\alpha}}}{N^{\frac{\alpha}{m+\alpha}}}} + 2N \sqrt{cL^{\frac{m}{m+\alpha}} \zeta_N^{\frac{\alpha}{m+\alpha}}} + H\gamma_N \sqrt{Nd\zeta} + H\alpha_N \sqrt{cNL_{\Phi}^{\frac{m}{m+\alpha_\Phi}} (d\zeta)^{\frac{\alpha_\Phi}{m+\alpha_\Phi}}} \\ &\leq 15\alpha_N \sqrt{cdL^{\frac{m}{m+\alpha}} N^{1+\frac{m}{m+\alpha}}} + \frac{2\alpha_N \sqrt{N}}{\sqrt{\beta_3(L^{\frac{m}{m+\alpha}} + (dL)^{\frac{m}{m+\alpha}} + d^2)}} \\ &\quad + H\sqrt{Nd\zeta} \sqrt{450\beta_3 N^{1+\frac{m}{m+\alpha}} \zeta_N^{\frac{\alpha}{m+\alpha}} L^{\frac{m}{m+\alpha}} d (L^{\frac{m}{m+\alpha}} + (dL)^{\frac{m}{m+\alpha}} + d^2)} \\ &\quad + H\sqrt{\beta_3 N \zeta_N^{\frac{\alpha}{m+\alpha}} (L^{\frac{m}{m+\alpha}} + (dL)^{\frac{m}{m+\alpha}} + d^2)} \sqrt{cNL^{\frac{m}{m+\alpha}} (d\zeta)^{\frac{\alpha}{m+\alpha}}} \\ &\leq 17\zeta H d \sqrt{\beta_3 N^{1+\frac{2m}{m+\alpha}} L^{\frac{m}{m+\alpha}} (L^{\frac{m}{m+\alpha}} + (dL)^{\frac{m}{m+\alpha}} + d^2)} \\ &\quad + \zeta H d \sqrt{450\beta_3 N^{1+\frac{2m}{m+\alpha}} L^{\frac{m}{m+\alpha}} (L^{\frac{m}{m+\alpha}} + (dL)^{\frac{m}{m+\alpha}} + d^2)} \\ &\quad + \zeta H d \sqrt{\beta_3 N^{1+\frac{m}{m+\alpha}} L^{\frac{m}{m+\alpha}} (L^{\frac{m}{m+\alpha}} + (dL)^{\frac{m}{m+\alpha}} + d^2)} \\ &\leq (17 + \sqrt{450} + 1)\zeta H d \sqrt{\beta_3 N^{1+\frac{2m}{m+\alpha}} L^{\frac{m}{m+\alpha}} (L^{\frac{m}{m+\alpha}} + (dL)^{\frac{m}{m+\alpha}} + d^2)} \\ &\leq 64\zeta H d \sqrt{\beta_3 N^{1+\frac{2m}{m+\alpha}} L^{\frac{m}{m+\alpha}} (L^{\frac{m}{m+\alpha}} + (dL)^{\frac{m}{m+\alpha}} + d^2)} \end{aligned}$$

Note that in first inequality above we apply Theorem 2 first, then we follow the steps as in the original proof of CHYL-L3, and after apply Cauchy Schwartz we use Jensen's inequality to move the $\alpha_E/(m + \alpha_E)$ power out the front, before applying their Lemma 10.

CHYL-P6 We can now leverage the result in CHYL-P6 to obtain the following expression for N_ϵ :

$$\frac{\epsilon N_\epsilon}{2} = 64\zeta H d \sqrt{\beta_3 N_\epsilon^{1+\frac{2m}{m+\alpha}} L^{\frac{m}{m+\alpha}} (L^{\frac{m}{m+\alpha}} + (dL)^{\frac{m}{m+\alpha}} + d^2)}$$

$$\Rightarrow N_\epsilon = \left(\frac{2^{14} \beta_3 H^2 d^2 L^{\frac{m}{m+\alpha}} (L^{\frac{m}{m+\alpha}} + (dL)^{\frac{m}{m+\alpha}} + d^2) \zeta^2}{\epsilon^2} \right)^{\frac{1}{1-\frac{2m}{m+\alpha}}}$$

Thus, letting $\tau = m/(m + \alpha) < 1/2$, the sample complexity of the algorithm is given by HN_ϵ which can be written as:

$$\tilde{O} \left(H^{1+\frac{2}{1-2\tau}} d^{\frac{2}{1-2\tau}} \left(\frac{1}{\epsilon} \right)^{\frac{2}{1-2\tau}} L^{\frac{\tau}{1-2\tau}} (L^\tau + (Ld)^\tau + d^2)^{\frac{1}{1-2\tau}} \right).$$

We note that as in the case of FLAMBE, when $\alpha_E, \alpha_\Phi \rightarrow \infty$, *i.e.* $\tau \rightarrow 0$, we recover the sample complexity of RAFFLE with a discrete action space and $K = 1$. \square