
Imitation Learning in Discounted Linear MDPs without exploration assumptions

Luca Viano¹ Stratis Skoulakis¹ Volkan Cevher¹

Abstract

We present a new algorithm for imitation learning in infinite horizon linear MDPs dubbed ILARL which greatly improves the bound on the number of trajectories that the learner needs to sample from the environment. In particular, we remove exploration assumptions required in previous works and we improve the dependence on the desired accuracy ϵ from $\mathcal{O}(\epsilon^{-5})$ to $\mathcal{O}(\epsilon^{-4})$. Our result relies on a connection between imitation learning and online learning in MDPs with adversarial losses. For the latter setting, we present the first result for infinite horizon linear MDP which may be of independent interest. Moreover, we are able to provide a strengthened result for the finite horizon case where we achieve $\mathcal{O}(\epsilon^{-2})$. Numerical experiments with linear function approximation shows that ILARL outperforms other commonly used algorithms.

1. Introduction

Imitation Learning (IL) is of extreme importance for all applications where designing a reward function is cumbersome while collecting demonstrations from an expert policy π_E is easy. Examples are autonomous driving (Knox et al., 2021), robotics (Osa et al., 2018), and economics/finance (Charpentier et al., 2020). The goal is to learn a policy which competes with the expert policy under the true unknown cost function of the Markov Decision Process (MDP) (Puterman, 1994).

Imitation learning relies on two data resources: expert demonstrations collected acting with π_E and data that can be collected interacting in the MDP with policies chosen by the learning algorithm. The first approach known as behavioural cloning (BC) solves the problem applying supervised learning. That is, it requires no interaction in the MDP but it requires knowledge of a class Π such that $\pi_E \in \Pi$ and

*Equal contribution ¹LIONS, EPFL, Lausanne, Switzerland. Correspondence to: Luca Viano <luca.viano@epfl.ch>.

$\tilde{\mathcal{O}}\left(\frac{\log|\Pi|}{(1-\gamma)^4\epsilon_E^2}\right)$ expert demonstrations to ensure with high probability that the output policy is at most ϵ_E -suboptimal.

The quartic dependence on the effective horizon term $((1-\gamma)^{-1})$ is problematic for long horizon problems. Moreover, the dependence on Π requires to make prior assumption on the expert policy structure to provide bounds which do not scale with the number of states in the function approximation setting. Thankfully, the dependence on the effective horizon can be improved resorting to MDP interaction. There exists an interesting line of works achieving this goal considering an interacting setting where the learner has the possibility to query the expert policy at any state visited during the MDP interaction (Ross & Bagnell, 2010; Ross et al., 2011) or that require a generative model to implement efficiently the moment matching procedure (Swamy et al., 2022). Another recent work requires a generative model to sample the initial state of the trajectory from the expert occupancy measure (Swamy et al., 2023). In this work, we considered a different scenario which is adopted in most of applied imitation learning (Ho et al., 2016; Ho & Ermon, 2016; Fu et al., 2018; Reddy et al., 2019; Dadashi et al., 2021; Watson et al., 2023; Garg et al., 2021; Ni et al., 2021). In this case, the expert policy can not be queried but only a dataset of expert demonstrations collected beforehand is available.

The setting has received scarce theoretical attention so far. The only results we are aware of are: (Shani et al., 2021; Xu et al., 2023) that focus on the tabular, finite horizon case, (Liu et al., 2022) in the finite horizon linear mixture MDP setting and (Viano et al., 2022) in the infinite horizon Linear MDP setting. In all these works bound the number of required expert demonstrations scale as $(1-\gamma)^{-2}$ which improves considerably over the quartic dependence attained by BC. However, (Viano et al., 2022) made the following assumption on the features that greatly simplifies the exploration in the MDP.

Assumption. Persistent excitation *It holds that for any policy π^k in the sequence of policies generated by the algorithm adopted by the learner*
 $\lambda_{\min}\left(\mathbb{E}_{s,a\sim d^{\pi^k}}[\phi(s,a)\phi(s,a)^\top]\right) \geq \beta > 0.$

Despite similar assumptions being commonly used in infinite horizon function approximation setting in imitation (Wu

et al., 2023; Zeng et al., 2022b;a; 2023) and reinforcement learning (see for example (Abbasi-Yadkori et al., 2019a; Hao et al., 2021; Duan et al., 2020; Lazic et al., 2020; Abbasi-Yadkori et al., 2019b; Agarwal et al., 2020a)), the persistent excitation assumption is very restrictive as it can be easily violated by deterministic policies with tabular features.

Our contribution We propose a new algorithm that improves the results of (Viano et al., 2022) in two important aspects: it **bypasses the persistent excitation assumption** (i.e. $\beta = 0$ does not cause the bound to blow up) and **it improves the dependence on ϵ** . In particular, the new proposed algorithm Algorithm 3 only requires $\mathcal{O}\left(\frac{d^3}{(1-\gamma)^8 \epsilon^4}\right)$ MDP interactions which greatly improves upon the bound $\tilde{\mathcal{O}}\left(\frac{d^2}{\beta^6(1-\gamma)^9 \epsilon^5}\right)$ proven by (Viano et al., 2022). Moreover, it holds that $\beta \leq d^{-1}$. Therefore, the bound for PPIL scales at least as $\tilde{\mathcal{O}}\left(\frac{d^8}{(1-\gamma)^9 \epsilon^5}\right)$ and, consequently, the bound of ILARL is better in all the relevant parameters d , $(1-\gamma)$ and ϵ . The design is different from (Viano et al., 2022) and it builds on a connection between imitation learning and online learning in MDP with full information. Therefore, **we design** as a submodule of our algorithm **the first algorithm for adversarial infinite horizon linear MDPs** which achieves $\mathcal{O}(K^{3/4})$ pseudo-regret. We also consider the finite horizon version of this algorithm which obtains a regret bound $\tilde{\mathcal{O}}\left(d^{3/4}H^{3/2}K^{3/4}\right)$ which improves by a factor $H^{1/2}$ the first result in this setting proven in (Zhong & Zhang, 2023). Concurrently to our work (Sherman et al., 2023a) derived a further improvement with optimal dependence on K .

Finally, **we provide a stronger result for the finite horizon setting**. Key for this result is realizing that in the regret decomposition of (Shani et al., 2021) one of the two players can in fact play the best response rather than a more conservative no regret strategy. This observations leads to Algorithm 4 which only requires $\mathcal{O}(H^4 d^3 \epsilon^{-2})$ MDP interactions.

Related Works Early works in behavioural cloning (BC) (Pomerleau, 1991) popularized the framework showing its success in driving problem and (Ross & Bagnell, 2010; Ross et al., 2011) show that the problem can be analyzed via a reduction to supervised learning which provides an expert trajectories bound of order $\frac{H^4 \log|\Pi|}{\epsilon^2}$. In practice, it is difficult to choose a class Π such that simultaneously contains the expert policy and is small enough to make the bound meaningful. Other algorithms like Dagger (Ross et al., 2011) and Logger (Li & Zhang, 2022) need to query the expert interactively. In this case, the expert trajectories improve to $\frac{H^2 \max_{s,a}(A^*(s,a))^2 \log|\Pi|}{\epsilon^2}$ where A^* is the optimal advantage. Recent works (Rajaraman et al., 2020)

showed that in the worst case Dagger does not improve over BC but also that both can use only $\tilde{\mathcal{O}}\left(\frac{H^2|S|}{\epsilon}\right)$ in the tabular case. Moreover, when transitions and initial distribution are known and the expert is deterministic, the result can be improved to $\mathcal{O}\left(\frac{H^{3/2}|S|}{\epsilon}\right)$ using Mimic-MD (Rajaraman et al., 2020). Later, (Xu et al., 2023) introduced MB-TAIL that having trajectory access to the MDP attains the same bound. This shows that the traditional bound obtained matching occupancy measure (Syed & Schapire, 2007) adopted in (Shani et al., 2021) is suboptimal in the tabular setting. For the linear function approximation, the works in (Swamy et al., 2022; Rajaraman et al., 2021) introduced algorithms that uses $\mathcal{O}\left(\frac{H^{3/2}d}{\epsilon}\right)$ expert trajectories with knowledge of the transitions but those require strong assumptions such as linear expert (Rajaraman et al., 2021, Definition 4), particular choice of features, linear reward and uniform expert occupancy measure. (Rajaraman et al., 2021) also proves an improved result for BC but under the linear expert assumption which implies that the expert is deterministic. While one can notice that there exists an optimal policy in a Linear MDP which is a linear expert, in our work we do not impose assumption on the expert policy and we require $\mathcal{O}\left(\frac{H^2 d}{\epsilon^2}\right)$ demonstrations. Under the same setting, the best known bound for BC is $\frac{H^2 \log|\Pi|}{d}$ times larger which makes our algorithm preferable whenever $|\Pi| \geq \exp(dH^{-2})$. We report a comparison with existing IL theory work in Table 1.

All the works we mentioned so far focused on the finite horizon, however the infinite horizon setting is the most common in practice (Ho et al., 2016; Ho & Ermon, 2016; Fu et al., 2018; Reddy et al., 2019; Dadashi et al., 2021; Watson et al., 2023; Garg et al., 2021). The practical advantage is that in the infinite horizon setting the optimal policy can be sought in the class of stationary policies which are much easier to store in memory than the nonstationary ones. Despite this fact, there are only few previous result studying IL in infinite horizon MDP and all of them operate under limiting exploration assumptions that we describe next.

Exploration assumptions in the literature As mentioned, (Viano et al., 2022) requires the persistent excitation assumption, (Wu et al., 2023) requires a uniformly good evaluation policy evaluation error which is possible only if the policies generated by the algorithm visits every state with high probability. (Zeng et al., 2022b) provided the first guarantees for IL with non linear reward functions but it assumes ergodic dynamics and that the the soft action value function of every policy can be perfectly evaluated at every state action pair. The latter assumption has been relaxed later in (Zeng et al., 2022a) and in (Zeng et al., 2023) that allows for a uniformly bounded policy evaluation error. In the latter case, the policies are evaluated under the transitions learned from expert data. However, their bound on

Table 1: **Comparison with related algorithms** We report the number of expert trajectories and MDP interactions needed for the various algorithms to be ϵ -suboptimal according to Definition 1. Our algorithms provide guarantees for the number of expert trajectories independent on \mathcal{S} and Π without assumptions on the expert policy. By **Linear Expert**, we mean that the expert policy is $\pi(s) = \max_{a \in \mathcal{A}} \phi(s, a)^\top \theta$ for some unknown vector θ .

Algorithm	Setting	Expert Traj.	MDP Traj.
Behavioural Cloning	Function Approximation, Offline (Agarwal et al., 2019)	$\mathcal{O}\left(\frac{H^4 \log \Pi }{\epsilon^2}\right)$	-
	Tabular, Offline (Rajaraman et al., 2020)	$\tilde{\mathcal{O}}\left(\frac{H^2 \mathcal{S} }{\epsilon}\right)$	-
	Linear Expert, Offline (Rajaraman et al., 2021)	$\tilde{\mathcal{O}}\left(\frac{H^2 d}{\epsilon}\right)$	-
Mimic-MD (Rajaraman et al., 2020)	Tabular, Known Transitions, Deterministic Expert	$\mathcal{O}\left(\frac{H^{3/2} \mathcal{S} }{\epsilon}\right)$	-
OAL (Shani et al., 2021)	Tabular	$\mathcal{O}\left(\frac{H^2 \mathcal{S} }{\epsilon^2}\right)$	$\mathcal{O}\left(\frac{H^4 \mathcal{S} ^2 \mathcal{A} }{\epsilon^2}\right)$
MB-TAIL (Xu et al., 2023)	Tabular, Deterministic Expert	$\mathcal{O}\left(\frac{H^{3/2} \mathcal{S} }{\epsilon}\right)$	$\mathcal{O}\left(\frac{H^3 \mathcal{S} ^2 \mathcal{A} }{\epsilon^2}\right)$
OGAIL (Liu et al., 2022)	Linear Mixture MDP	$\mathcal{O}\left(\frac{H^3 d^2}{\epsilon^2}\right)$	$\mathcal{O}\left(\frac{H^4 d^3}{\epsilon^2}\right)$
PPIL (Viano et al., 2022)	Linear MDP, Persistent Excitation	$\mathcal{O}\left(\frac{d}{(1-\gamma)^2 \epsilon^2}\right)$	$\mathcal{O}\left(\frac{d^2}{\beta^6 (1-\gamma)^9 \epsilon^5}\right)$
ILARL (Algorithm 3)	Linear MDP	$\mathcal{O}\left(\frac{d}{(1-\gamma)^2 \epsilon^2}\right)$	$\mathcal{O}\left(\frac{d^3}{(1-\gamma)^8 \epsilon^4}\right)$
BRIG (Algorithm 4)	Episodic Linear MDP	$\mathcal{O}\left(\frac{dH^2}{\epsilon^2}\right)$	$\mathcal{O}\left(\frac{d^3 H^4}{\epsilon^2}\right)$

the number of expert trajectories scale as $(1 - \gamma)^{-4}$ and with the number of states visited by the expert while our bound leverages online access to the MDP to obtain a better horizon and to avoid the dependence on the number of states. Moreover, the bounds in (Zeng et al., 2023; 2022a; Wu et al., 2023) depend on the number of states which can be prohibitively large in the function approximation setting.

2. Background and Notation

In imitation learning (Osa et al., 2018), the environment is abstracted as Markov Decision Process (MDP) (Puterman, 1994) which consists of a tuple $(\mathcal{S}, \mathcal{A}, P, c, \nu_0)$ where \mathcal{S} is the state space, \mathcal{A} is the action space, $P : \mathcal{S} \times \mathcal{A} \rightarrow \Delta_{\mathcal{S}}$ is the transition kernel, that is, $P(s'|s, a)$ denotes the probability of landing in state s' after choosing action a in state s . Moreover, ν_0 is a distribution over states from which the initial state is sampled. Finally, $c : \mathcal{S} \times \mathcal{A} \rightarrow [-1, 1]$ is the cost function. In the infinite horizon setting, we endow the MDP tuple with an additional element called the discount factor $\gamma \in [0, 1)$. Alternatively, in the finite horizon setting we append to the MDP tuple the horizon $H \in \mathbb{N}$ and we consider possibly inhomogenous transitions or costs function. That is, they depend on the stage within the episode. The agent plays action in the environment sampled from a policy $\pi : \mathcal{S} \rightarrow \Delta_{\mathcal{A}}$. The learner is allowed to adopt an algorithm to update the policy across episodes given the previously observed history. We will see that imitation learning has a strong connection with MDPs with adversarial costs. The latter setting allows the cost function to change each time the learner samples a new episode in the MDP. For clarity, we include the pseudocode for the interaction in Protocol 1 in Appendix B.

Value functions and occupancy measures We define the state value function at state $s \in \mathcal{S}$ for the policy π under the cost function c as $V^\pi(s; c) \triangleq \mathbb{E}[\sum_{h=0}^{\infty} \gamma^h c(s_h, a_h) | s_1 = s]$. In the finite horizon case, the state value function also depends on the stage index h , that is $V_h^\pi(s; c) \triangleq \mathbb{E}[\sum_{\ell=h}^H c(s_\ell, a_\ell) | s_h = s]$ ¹. In both cases, the expectation over both the randomness of the transition dynamics and the one of the learner’s policy. Another convenient quantity is the occupancy measure of a policy π denoted as $d^\pi \in \Delta_{\mathcal{S} \times \mathcal{A}}$ and defined as follows $d^\pi(s, a) \triangleq (1 - \gamma) \sum_{h=0}^{\infty} \gamma^h \mathbb{P}[s, a \text{ is visited after } h \text{ steps acting with } \pi]$. We can also define the state occupancy measure as $d^\pi(s) \triangleq (1 - \gamma) \sum_{h=0}^{\infty} \gamma^h \mathbb{P}[s \text{ is visited after } h \text{ steps acting with } \pi]$. In the finite horizon setting, the occupancy measure depends on the stage h and its defined simply as $d_h^\pi(s, a) \triangleq \mathbb{P}[s, a \text{ is visited after } h \text{ steps acting with } \pi]$. The state occupancy measure is defined analogously.

Imitation Learning In imitation learning, the learner is given a dataset $\mathcal{D}_E \triangleq \{\tau^k\}_{k=1}^{\tau_E}$ containing τ_E trajectories collected in the MDP by an expert policy π_E according to Protocol 1. By trajectory τ^k , we mean the sequence of states and actions sampled at the k^{th} iteration of Protocol 1, that is $\tau^k = \{(s_h^k, a_h^k)\}_{h=1}^H$ for finite horizon case. For the infinite horizon case, the trajectories have random length sampled from the distribution $\text{Geometric}(1 - \gamma)$. Given, \mathcal{D}_E the learner adopts an algorithm \mathcal{A} to learn a policy π^{out} such that is ϵ -suboptimal according to the next definition.

Definition 1. An algorithm \mathcal{A} is said ϵ -suboptimal

¹In the finite horizon case we may use $V^\pi(s; c)$ as a shortcut for $V_1^\pi(s; c)$

if it outputs a policy π whose value function with respect to the unknown true cost \mathbf{c}_{true} satisfies $\mathbb{E}_{\mathcal{A}} \mathbb{E}_{s_1 \sim \nu_0} [V^\pi(s_1; \mathbf{c}_{\text{true}}) - V^{\pi_E}(s_1; \mathbf{c}_{\text{true}})] \leq \epsilon$ where the first expectation is on the randomness of the algorithm \mathcal{A} .

2.1. Setting

We study imitation learning in the linear MDP setting popularized by (Jin et al., 2019) and studied in imitation learning in (Viano et al., 2022). When studying finite horizon problems we consider possible inhomogeneous transition dynamics and cost function. That is, we work under the following assumptions.

Assumption 1. Episodic Linear MDP There exist a feature matrix $\Phi \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}| \times d}$ known to the learner, an unknown sequence of vectors $\mathbf{w}_h^k \in \mathbb{R}^d$ and an unknown matrix sequences $M_h \in \mathbb{R}^{d \times |\mathcal{S}|}$ such that the transition matrices P_h factorize as $P_h = \Phi M_h$ and the sequence of adversarial costs c_h^k can be written as $c_h^k = \Phi \mathbf{w}_h^k$. Moreover, it holds for all $k \in [K], h \in [H]$ and for all state action pairs $s, a \in \mathcal{S} \times \mathcal{A}$ that $\|\Phi\|_{1,\infty} \leq 1$, $\|M_h\|_{1,\infty} \leq 1$, $\|\mathbf{w}_h^k\|_2 \leq 1$.

Assumption 2. Linear MDP There exist a feature matrix $\Phi \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}| \times d}$ known to the learner, an unknown sequence of vectors $\mathbf{w}^k \in \mathbb{R}^d$ and an unknown matrix $M \in \mathbb{R}^{d \times |\mathcal{S}|}$ such that the transition matrices P factorize as $P = \Phi M$ and the sequence of adversarial costs c^k can be written as $c^k = \Phi \mathbf{w}^k$. Moreover, it holds for all $k \in [K]$ and for all state action pairs $s, a \in \mathcal{S} \times \mathcal{A}$ that $\|\Phi\|_{1,\infty} \leq 1$, $\|M\|_{1,\infty} \leq 1$, $\|\mathbf{w}^k\|_2 \leq 1$.

In the context of imitation learning, we also need to assume that the true unknown cost is realizable.

Assumption 3. Realizable cost The learner has access to a feature matrix $\Phi \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}| \times d}$ such that $\mathbf{c}_{\text{true}} = \Phi \mathbf{w}_{\text{true}}$.

3. Main Results and techniques

We provide our main results for the infinite horizon case in Theorem 1 and the stronger result for the finite horizon in Theorem 2.

Theorem 1. Under Assumptions 2,3, there exists an algorithm, i.e. Algorithm 3, such that after using $\tilde{\mathcal{O}}\left(\frac{\log|\mathcal{A}|d^3}{(1-\gamma)^8\epsilon^4}\right)$ state action pairs from the MDP and using $\tilde{\mathcal{O}}\left(\frac{2d \log(2d)}{(1-\gamma)^2\epsilon_E^2}\right)$ expert demonstrations is $\epsilon + \epsilon_E$ -suboptimal.

Theorem 2. Under Assumptions 1,3, there exists an algorithm, i.e. Algorithm 4, such that after sampling $\mathcal{O}(H^4 d^3 \log(dH)/(\epsilon)\epsilon^{-2})$ trajectories and having access to a dataset of $\tau_E = \tilde{\mathcal{O}}\left(\frac{2H^2 d \log(2d)}{\epsilon_E^2}\right)$ expert demonstrations is $\epsilon + \epsilon_E$ -suboptimal.

Remark 1. The results are proven via the high probability bounds in Theorems 5 and 6 respectively and apply the high probability to expectation conversion lemma in Lemma 7.

3.1. Technique overview

Online-to-batch conversion The core idea is to extract the policy achieving the sample complexity guarantees above via an online-to-batch conservation. That is the output policy is sampled uniformly from a collection of K policies $\{\pi^k\}_{k=1}^K$. The sample complexity result is proven, showing that the policies $\{\pi^k\}_{k=1}^K$ produced by the algorithms under study have sublinear pseudo regret in high probability. That is, in the infinite horizon discounted setting with Algorithm 3 guarantees

$$\text{Regret}(K) \triangleq \frac{1}{1-\gamma} \sum_{k=1}^K \langle \mathbf{c}_{\text{true}}, d^{\pi^k} - d^{\pi_E} \rangle \leq \mathcal{O}(K^{3/4})$$

with high probability and Algorithm 4 ensures $\text{Regret}(K)$

$$\sum_{h=1}^H \sum_{k=1}^K \langle \mathbf{c}_{\text{true},h}, d_h^{\pi^k} - d_h^{\pi_E} \rangle \leq \mathcal{O}(\sqrt{K}) \quad (1)$$

with high probability in the finite horizon setting. The next section presents the regret decomposition giving the crucial insights for the design of Algorithms 3 and 4.

Regret decomposition To obtain both regret bounds, we decompose the pseudo regret in 3 terms. We present it for the infinite horizon case, where $(1-\gamma)\text{Regret}(K)$ is equal to

$$\underbrace{\sum_{k=1}^K \langle \mathbf{c}^k, d^{\pi^k} - d^{\pi_E} \rangle}_{\text{Regret}_\pi(K; d^{\pi_E})} + \underbrace{\sum_{k=1}^K \langle \mathbf{w}_{\text{true}} - \mathbf{w}^k, \Phi^\top d^{\pi^k} - \Phi^\top d^{\pi_E} \rangle}_{\text{Regret}_w(K; \mathbf{w}_{\text{true}})} \quad (2)$$

This decomposition is inspired from (Shani et al., 2021) but it applies also to the infinite horizon setting and exploits the linear structure using Assumptions 2,3 to write $c^k = \Phi \mathbf{w}^k$ and $\mathbf{c}_{\text{true}} = \Phi \mathbf{w}_{\text{true}}$.

$\text{Regret}_w(K; \mathbf{w}_{\text{true}})$ is the pseudo regret of a player updating a sequence of cost functions and having \mathbf{c}_{true} as comparator while $\text{Regret}_\pi(K; d^{\pi_E})$ is the pseudo regret in a Linear MDP with adversarial costs $\{c^k\}_{k=1}^K$ and having the expert occupancy measure as a comparator.

Imitation Learning via no-regret algorithms The decomposition in Equation (2) suggests that imitation learning algorithm can be designed chaining one algorithm that updates the sequence \mathbf{w}^k to make sure that $\text{Regret}_w(K; \mathbf{w}_{\text{true}})$ grows sublinearly and a second one that updates the policy sequence to control $\text{Regret}_\pi(K; d^{\pi_E})$. Controlling

$\text{Regret}_w(K; w_{\text{true}})$ can be easily done via projected online gradient descent (Zinkevich, 2003).

Unfortunately, controlling $\text{Regret}_\pi(K; d^{\pi^\epsilon})$ is way more challenging because we have no knowledge of the transition dynamics. Therefore, we can not project on the feasible set of occupancy measures. To circumvent this issue we rely on the recent literature (Luo et al., 2021; Sherman et al., 2023b; Dai et al., 2023) that however focuses on bandit feedback. In our case, the π player has full information on the cost vector c^k . Thus, we design a simpler algorithm Algorithm 1 which achieves a better regret bound in the easier full information case. Algorithm 1 improves over the regret bound in (Zhong & Zhang, 2023) and easily extends to the infinite horizon setting (see Algorithm 2).

Improved algorithm for finite horizon The techniques explained so far do not allow to get the better bound of order $\mathcal{O}(\sqrt{K})$ in the finite horizon setting (see Equation (1)). The idea is to let the w player update first, then the π player can update their policy knowing in advance the loss that they will suffer. This allows to use LSVI-UCB (Jin et al., 2019) for the π -player which has been originally designed for a fixed cost but we show that it still guarantees $\mathcal{O}(\sqrt{T})$ regret against an arbitrary sequence of costs when the learner knows in advance the cost function at the next episode. On the other hand, LSVI-UCB suffers linear regret if the adversarial loss is not known in advance so letting the w player update first is crucial. This result is provided in Section 7.

4. Warm up: Online Learning in Adversarial Linear MDP

We start by presenting our result in full information episodic linear MDP with adversarial costs that improves over (Zhong & Zhang, 2023) by a factor $H^{1/2}$. The algorithm is quite simple. We apply a policy iteration like method with two important twist: (i) in the policy improvement step, we update the policy with a no regret algorithm rather than a greedy step. Moreover, the policy is updated only every τ episodes using as loss vector the average Q value over the last batch of collected episodes, (ii) in the policy evaluation step, we compute an optimistic estimate of the Q function for the current policy using only on-policy data.

The last part is crucial because the use of off-policy data makes the covering argument for Linear MDP problematic. Indeed, one would need to cover the space of stochastic policy when computing the covering number of the value function class but this leads to the undesirable dependence on the number of states and actions for the log covering number (see for example (Abbasi-Yadkori et al., 2013)). An alternative bound on the covering number shown in (Zhong & Zhang, 2023) would instead lead to linear regret.

Instead, using data collected on-policy allows to apply the covering argument to a much smaller function class avoiding the dependence on the number of states and actions. The first twist is at this point necessary to make the policy updates more rare giving the possibility to collect more on-policy episodes with a fixed policy. The algorithm pseudocode is in Algorithm 1.

4.1. Analysis

Theorem 3. *Under Assumption 1, run Algorithm 1 with exploration parameter $\beta = \tilde{\mathcal{O}}(dH)$, dataset size $\tau = \frac{5\beta}{2} \sqrt{\frac{Kd}{\log|\mathcal{A}|}}$ and step size $\eta = \sqrt{\frac{\tau \log|\mathcal{A}|}{KH^2}}$. Then, it holds with probability $1 - \delta$, that $\text{Regret}(K; \pi^*) = \sum_{k=1}^K V_1^{\pi^*,k}(s_1) - V_1^{\pi^*,k}(s_1)$ is upper bounded by*

$$\tilde{\mathcal{O}}\left(d^{3/4} H^{3/2} \log^{1/4} |\mathcal{A}| K^{3/4} \log \frac{K}{\delta}\right) \quad (3)$$

where we use the compact notation $V_h^{\pi^*,k}(\cdot) \triangleq V_h^\pi(\cdot; c^k)$.

Proof. Sketch Adding and subtracting the term $\sum_{k=1}^K V_1^k(s_1)$ in the definition of regret, we have that defining $\delta_h^k(s, a) \triangleq \mathbf{c}_h^k(s, a) + P_h V_{h+1}^k(s, a) - Q_h^k(s, a)$

$$\begin{aligned} \text{Regret}(K; \pi^*) &\leq \\ &\sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{s \sim d_h^{\pi^*}} [\langle Q_h^k(s, \cdot), \pi_h^k(s) - \pi_h^*(s) \rangle] \\ &\quad - \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{s, a \sim d_h^{\pi^*}} [\delta_h^k(s, a)] \\ &\quad + \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{s, a \sim d_h^{\pi^*}} [\delta_h^k(s, a)] \end{aligned}$$

where the last inequality holds by the extended performance difference lemma (Cai et al., 2020; Shani et al., 2020). At this point we can invoke Lemma 3 (see Appendix E) to obtain

$$-2b_h^k(s, a) \leq Q_h^k(s, a) - \mathbf{c}_h^k(s, a) - P_h V_{h+1}^k(s, a) \leq 0 \quad (4)$$

for all $(s, a) \in \mathcal{S} \times \mathcal{A}$, $h \in [H]$, $k \in [K]$ with probability $1 - \delta$. This implies that with probability $1 - \delta$

$$\begin{aligned} \text{Regret}(K; \pi^*) &\leq \\ &\sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{s \sim d_h^{\pi^*}} [\langle Q_h^k(s, \cdot), \pi_h^k(s) - \pi_h^*(s) \rangle] \\ &\quad + 2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{s, a \sim d_h^{\pi^*}} [b_h^k(s, a)] \end{aligned}$$

Algorithm 1 On-policy MDP-E with unknown transitions and adversarial costs.

```

1: Input: Dataset size  $\tau$ , Exploration parameter  $\beta$ , Step size  $\eta$ , initialize  $\pi_0$  as uniform distribution over  $\mathcal{A}$ 
2: for  $j = 1, \dots, \lfloor K/\tau \rfloor$  do
3:   Denote the indices interval  $T_j \triangleq [(j-1)\lfloor K/\tau \rfloor, j\lfloor K/\tau \rfloor]$ .
4:   // Collect on-policy data
5:   Collect  $\tau$  trajectories with policy  $\pi^{(j)}$  and store them in the dataset  $\mathcal{D}_h^{(j)} = \{(s_h^i, a_h^i, \mathbf{c}_h^i, s_{h+1}^i)\}_{i \in T_j}$ .
6:   Denote global dataset  $\mathcal{D}^{(j)} = \cup_{h=1}^H \mathcal{D}_h^{(j)}$ .
7:   for  $k \in T_j$  do
8:     // Optimistic policy evaluation
9:     Initialize  $V_{H+1}^k = 0$ 
10:    for  $h = H, \dots, 1$  do
11:       $\Lambda_h^k = \sum_{(s,a) \in \mathcal{D}_h^{(j)}} \phi(s,a)\phi(s,a)^\top + I$  //  $\phi(s,a)$  is the  $(s,a)$ th row of the matrix  $\Phi$ .
12:       $\mathbf{v}_h^k = (\Lambda_h^k)^{-1} \sum_{(s,a,s') \in \mathcal{D}_h^{(j)}} \phi(s,a)V_{h+1}^k(s')$ 
13:       $b_h^k(s,a) = \beta \|\phi(s,a)\|_{(\Lambda_h^k)^{-1}}$ 
14:       $Q_h^k = [\mathbf{c}_h^k + \Phi \mathbf{v}_h^k - b_h^k]_{-H+h-1}^{H-h+1}$ 
15:       $V_h^k(s) = \langle \pi_h^k(s), Q_h^k(s, \cdot) \rangle$  (with  $\pi^k = \pi^{(j)}$ ).
16:    end for
17:  end for
18:  // Policy Improvement Step
19:  Compute average  $Q$  value  $\bar{Q}_h^{(j)}(s,a) = \frac{1}{\tau} \sum_{k \in T_j} Q_h^k(s,a)$ .
20:  Update policy  $\pi_h^{(j+1)}(a|s) \propto \exp\left(-\eta \sum_{i=1}^j \bar{Q}_h^{(i)}(s,a)\right)$ 
21: end for
    
```

Then applying the mirror descent with blocking result given in Lemma 9 we obtain

$$\begin{aligned} \text{Regret}(K; \pi^*) &\leq \frac{\tau \log |\mathcal{A}|}{\eta} + \tau H + \eta K H^2 \\ &\quad + 2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{s,a \sim d^{\pi^k}} [b_h^k(s,a)] \end{aligned}$$

. Then, by Lemma 10, it holds that with probability $1 - 2\delta$ that

$$\begin{aligned} \text{Regret}(K; \pi^*) &\leq \frac{\tau \log |\mathcal{A}|}{\eta} + \tau H + \eta K H^2 \\ &\quad + \frac{14KH\sqrt{d}\beta \log(2\tau/\delta)}{\sqrt{\tau}}. \end{aligned}$$

The proof is concluded plugging in the values specified in the theorem statement. \square

4.2. Extension to the Infinite Horizon Setting.

We show our proposed extension to the infinite horizon in Algorithm 2. The main difference in the analysis is to handle the fact that in the infinite horizon setting we can not run a backward recursion to compute the optimistic value functions as done in Steps 10-16 of Algorithm 1. Instead, we use the optimistic estimate at the previous iterate Q^k to build an approximate optimistic estimate at the next iterate (see

Steps 10-12 in Algorithm 2). This error can be controlled thanks to the regularization in the policy improvement step as noticed in (Moulin & Neu, 2023).

Theorem 4. *Under Assumption 2, consider K iterations of Algorithm 2 run with $\tau \leq \frac{K}{\sqrt{\tau}}$ and $\beta = \tilde{\mathcal{O}}(dH)$, then it holds for any comparator policy π^* that with probability $1 - 2\delta$ by*

$$\begin{aligned} (1-\gamma)\text{Regret}(K; \pi^*) &\triangleq \sum_{k=1}^K \left\langle d^{\pi^k} - d^{\pi^*}, \mathbf{c}^k \right\rangle \\ &\leq \frac{\tau \log |\mathcal{A}|}{\eta} + \frac{\tau+1}{1-\gamma} + \frac{\eta K}{(1-\gamma)^2} \\ &\quad + 12\beta K \sqrt{\frac{d}{\tau}} \log\left(\frac{2Kd}{\tau\delta}\right) + \frac{\sqrt{2\eta}K}{(1-\gamma)^2\tau}. \end{aligned}$$

Proof. Sketch The proof is based on the following decomposition that holds in virtue of Lemma 2. Denoting $\delta^k(s,a) \triangleq \mathbf{c}^k(s,a) + \gamma P V^k(s,a) - Q^{k+1}(s,a)$ and $g^k(s,a) \triangleq Q^{k+1}(s,a) - Q^k(s,a)$, we have that $(1 -$

Algorithm 2 Infinite Horizon Linear MDP with adversarial losses.

- 1: **Input:** Dataset size τ , Exploration parameter β , Step size η , Initial policy π_0 (uniform over \mathcal{A}), initialize $V^1 = 0$.
- 2: **for** $j = 1, \dots, \lfloor K/\tau \rfloor$ **do**
- 3: **// Collect on-policy data**
- 4: Denote the indices interval $T_j \triangleq [(j-1)\lfloor K/\tau \rfloor, j\lfloor K/\tau \rfloor]$.
- 5: Sample $\mathcal{D}^{(j)} = \{(s^i, a^i, s'^i, \mathbf{c}^i)\}_{i \in T_j} \sim d^{\pi^{(j)}}$ using (Agarwal et al., 2020b, Algorithm 1).
- 6: Compute $\Lambda^{(j)} = \sum_{(s,a) \in \mathcal{D}^{(j)}} \phi(s,a)\phi(s,a)^\top + I$.
- 7: Compute $b^{(j)}(s,a) = \beta \|\phi(s,a)\|_{(\Lambda^{(j)})^{-1}}$.
- 8: **// Optimistic Policy Evaluation**
- 9: **for** $k \in T_j$ **do**
- 10: $\mathbf{v}^k = (\Lambda^{(j)})^{-1} \sum_{(s,a,s') \in \mathcal{D}^{(j)}} \phi(s,a)V^k(s')$
- 11: $Q^{k+1} = [\mathbf{c}^k + \gamma \Phi \mathbf{v}^k - b^{(j)}]_{-(1-\gamma)^{-1}}^{(1-\gamma)^{-1}}$
- 12: $V^{k+1}(s) = \langle \pi^{(j)}(a|s), Q^{k+1}(s,a) \rangle$
- 13: **end for**
- 14: **// Policy Improvement Step**
- 15: Compute average Q value $\bar{Q}^{(j)}(s,a) = \frac{1}{\tau} \sum_{k \in T_j} Q^k(s,a)$.
- 16: Update policy: $\pi^{(j+1)}(a|s) \propto \exp\left(-\eta \sum_{i=1}^j \bar{Q}^{(i)}(s,a)\right)$
- 17: **end for**

γ)Regret($K; \pi^*$) equals

$$\sum_{k=1}^K \mathbb{E}_{s \sim d^{\pi^*}} [\langle Q^k(s, \cdot), \pi^k(s) - \pi^*(s) \rangle] \quad (\text{OMD})$$

$$+ \sum_{k=1}^K \sum_{s,a} [d^{\pi^k}(s,a) - d^{\pi^*}(s,a)] \delta^k(s,a) \quad (\text{Opt})$$

$$+ \sum_{k=1}^K \sum_{s,a} [d^{\pi^k}(s,a) - d^{\pi^*}(s,a)] g^k(s,a) \quad (\text{Shift})$$

Then, we have that **Opt** can be bounded similarly to the finite horizon case using Lemma 4 while for **Shift** we rely on the regularization of the policy improvement step and on the fact that the policy is updated only every τ steps. All in all, we have that the first term in **(Shift)** can be bounded as $\frac{1}{(1-\gamma)^2} \sum_{j=2}^{\lfloor K/\tau \rfloor} \sqrt{2\eta} = \frac{\sqrt{2\eta}K}{(1-\gamma)^2\tau}$. The second term just telescopes therefore **(Shift)** $\leq \frac{\sqrt{2\eta}K}{(1-\gamma)^2\tau} + \frac{1}{1-\gamma}$. Finally, **(OMD)** can be bounded as in the finite horizon case. \square

5. Imitation Learning in Infinite Horizon MDPs

In this section, we apply Theorem 4 to imitation learning. Indeed, we design Algorithm 3 using the insights from the

Algorithm 3 Imitation Learning via Adversarial Reinforcement Learning (**ILARL**) for Infinite Horizon Linear MDPs.

- 1: **Input:** Access to Algorithm 2 (with inputs τ, β, η, π_0), Step size for Cost Update α , Expert dataset $\mathcal{D}_{\pi_E} = \{\tau_i\}_{i=1}^{\tau_E}$.
- 2: Estimate for expert feature visitation $\widehat{\Phi}^\top d^{\pi_E} \triangleq \frac{(1-\gamma)}{\tau_E} \sum_{\tau \in \mathcal{D}_{\pi_E}} \sum_{s_h, a_h \in \tau} \gamma^h \phi(s_h, a_h)$.
- 3: **for** $k = 1, \dots, K$ **do**
- 4: **// Cost Update (to control Regret $_w(K; w_{\text{true}})$)**
- 5: Estimate $\widehat{\Phi}^\top d^{\pi^k} \triangleq \tau^{-1} \sum_{s,a \in \mathcal{D}^{(j)}} \phi(s,a)$ where $\mathcal{D}^{(j)}$ is defined in Step 6 in Algorithm 2.
- 6: $\mathbf{w}^{k+1} = \Pi_{\mathcal{W}} \left[\mathbf{w}^k - \alpha (\widehat{\Phi}^\top d^{\pi_E} - \widehat{\Phi}^\top d^{\pi^k}) \right]$ with $\mathcal{W} = \{\mathbf{w} : \|\mathbf{w}\|_2 \leq 1\}$.
- 7: **// Policy Update (to control Regret $_\pi(K; \pi_E)$)**
- 8: The cost $\Phi \mathbf{w}^k$ is revealed to the learner.
- 9: The learner updates their policy π^k performing one iteration of Algorithm 2.
- 10: **end for**

decomposition in Equation (2): we use a no regret algorithm to update the cost at each round and we update the learner's policy using a no regret algorithm for infinite horizon full information adversarial Linear MDP, of which Algorithm 2 is the first example in the literature. The guarantees for Algorithm 3 are given in the following theorem.

Theorem 5. Under Assumptions 2,3, let us consider K iterations of Algorithm 3 with $K \geq \tilde{\mathcal{O}}\left(\frac{\log|\mathcal{A}|d\beta^2 \log^2(1/\delta)}{(1-\gamma)^6 \epsilon^4}\right)$ where β is chosen as in Lemma 6 (i.e. $\beta = \tilde{\mathcal{O}}(d(1-\gamma))^{-1}$). Moreover, let consider the following choices $\alpha = \frac{1}{\sqrt{2K}}$, $\tau = \mathcal{O}\left(\frac{\beta(1-\gamma)\sqrt{dK} \log(2dK/\delta)}{\sqrt{\log|\mathcal{A}|}}\right)$, expert trajectories $\tau_E = \frac{8d \log(d/\delta)}{(1-\gamma)^2 \epsilon^2}$ and $\eta = \sqrt{\frac{\tau \log|\mathcal{A}|(1-\gamma)^2}{K}}$. Then, the above conditions ensure $\frac{1}{1-\gamma} \left\langle \mathbf{c}_{\text{true}}, d^{\pi_E} - \frac{1}{K} \sum_{k=1}^K d^{\pi^k} \right\rangle \leq \epsilon + \epsilon_E$ with probability $1 - 4\delta$.

The proof included in Appendix E starts with the decomposition in Equation (2). Then, we control the term $\text{Regret}_w(K; w_{\text{true}})$ with the standard online gradient descent analyses and the term $\text{Regret}_\pi(K; \pi_E)$ with Theorem 4. Finally, we control the statistical estimation error for the losses seen by the w -player with an application of Lemma 8.

Remark 2. The resulting algorithm improves over (Viano et al., 2022) in two ways: (i) We bypass all kind of exploration assumptions, such as the persistent excitation assumption. We remark that this a qualitative improvement. Indeed, the persistent excitation assumption is easily violated by deterministic policies with tabular features. (ii) Moreover, the sample complexity improves from $\mathcal{O}(\epsilon^{-5})$ to $\mathcal{O}(\epsilon^{-4})$.

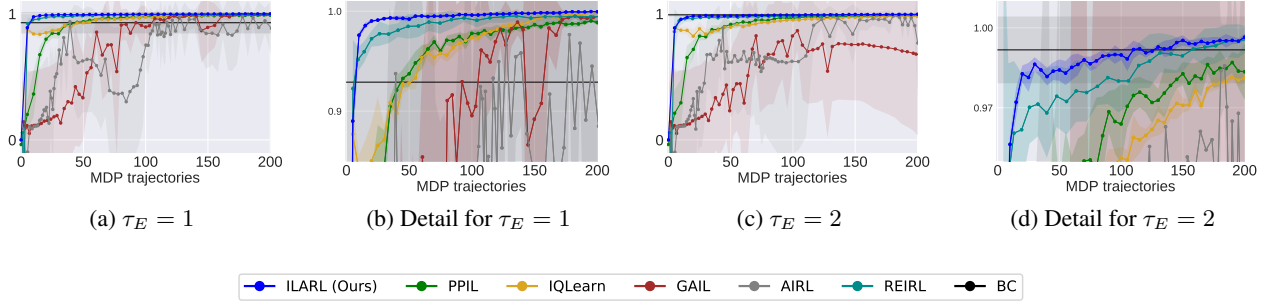


Figure 1: Experiments on a continuous gridworld with a stochastic expert. The y -axis reports the normalized return. 1 corresponds to the expert performance and 0 to the uniform policy one.

6. Empirical evaluation

We numerically verify the main theoretical insights derived in the previous sections (i) We aim to verify that for a general stochastic expert, the efficiency in terms of expert trajectories improves upon behavioural cloning. (ii) ILARL is more efficient in terms of MDP trajectories compared to PPIL (Viano et al., 2022) which has worst theoretical guarantees and with popular algorithms that are widely used in practice but do not enjoy theoretical guarantees: GAIL (Ho et al., 2016), AIRL (Fu et al., 2018), REIRL (Boularias et al., 2011) and IQLearn (Garg et al., 2021). The experiments are run in a continuous state MDP explained in Appendix G.

Expert trajectory efficiency with stochastic expert For the first claim, we use a stochastic expert obtained following with equal probability either the action taken by a deterministic experts previously trained with LSVI-UCB or an action sampled uniformly at random. We collect with such policy τ_E trajectories. From Figure 1, we observe that all imitation learning we tried have a final performance improving over behavioural cloning for the case $\tau_E = 1$ while only REIRL and ILARL do so for $\tau_E = 2$. In both cases, ILARL achieves the highest return that even matches the expert performance.

MDP trajectories efficiency For the second claim, we can see in Figure 1 that ILARL is the most efficient algorithm in terms of MDP trajectories for both values of τ_E .

7. An improvement for the finite horizon case

We notice that in Algorithm 3 we missed an opportunity. In fact, we could use prior knowledge of the cost function \mathbf{c}_h^k to update the policy π_h^k . In other words, in Step 8 of Algorithm 3 we could reveal \mathbf{w}^{k+1} to the algorithm that controls $\text{Regret}_\pi(K; d^{\pi_E})$. From an online learning perspective we can play best response to control more effectively the regret term $\text{Regret}_\pi(K; d^{\pi_E})$ using LSVI-UCB (Jin et al., 2019). This idea leads to the Algorithm 4. Theorem 6 proves that Algorithm 4 improves the required number of interaction to $K = \tilde{\mathcal{O}}(H^4 d^3 \epsilon^{-2} \log^2(1/\delta))$ which greatly improves over $K = \tilde{\mathcal{O}}(H^6 \log |\mathcal{A}| d^3 \epsilon^{-4} \log^2(1/\delta))$ achieved by Al-

Algorithm 4 Best Response Imitation learningG (BRIG).

- 1: **Input:** Exploration parameter β , Step size η , Reward step size α , expert dataset $\mathcal{D}_{\pi_E} = \cup_{h \in [H]} \mathcal{D}_{\pi_E, h}$.
- 2: Initialize π_0 as uniform distribution over \mathcal{A} .
- 3: Estimate expert features expectation vectors $\widehat{\Phi}^\top d_h^{\pi_E} = \frac{1}{|\mathcal{D}_{\pi_E, h}|} \sum_{s, a \in \mathcal{D}_{\pi_E, h}} \phi(s, a)$ for all $h \in [H]$.
- 4: **for** $k = 1, \dots, K$ **do**
- 5: Collect one episodes with policy π^k denoted as $\tau^k = \{(s_h^k, a_h^k, s_{h+1}^k)\}_{h=1}^{H-1}$ and for every $h \in [H]$ append data $(s_h^k, a_h^k, s_{h+1}^k)$ to \mathcal{D}_h .
- 6: // Cost Update (to control $\text{Regret}_w(K; w_{\text{true}})$)
- 7: $\mathbf{w}_h^{k+1} = \Pi_{\mathcal{W}} \left[\mathbf{w}_h^k - \alpha (\widehat{\Phi}^\top d_h^{\pi_E} - \phi(s_h^k, a_h^k)) \right]$ with $\mathcal{W} = \{\mathbf{w} : 0 \leq \mathbf{w} \leq 1\}$ for all $h \in [H]$.
- 8: // Full information LSVI-UCB (to control $\text{Regret}_\pi(K; \pi_E)$)
- 9: Initialize $V_{H+1}^k = 0$
- 10: **for** $h = H, \dots, 1$ **do**
- 11: $\Lambda_h^k = \sum_{(s, a) \in \mathcal{D}_h} \phi(s, a) \phi(s, a)^\top + I$
- 12: $\mathbf{v}_h^k = (\Lambda_h^k)^{-1} \sum_{(s, a, s') \in \mathcal{D}_h} \phi(s, a) V_{h+1}^k(s')$
- 13: $b_h^k(s, a) = \beta \|\phi(s, a)\|_{(\Lambda_h^k)^{-1}}$.
- 14: // Q-update using the future loss \mathbf{w}_h^{k+1}
- 15: $Q_h^k = [\widehat{\Phi} \mathbf{w}_h^{k+1} + \Phi \mathbf{v}_h^k - b_h^k]_{-H+h-1}^{H-h+1}$.
- 16: // Greedy policy update, Best Response
- 17: $\pi_h^{k+1}(s) = \arg \min (Q_h^k(s, \cdot))$.
- 18: $V_h^k(s) = \langle \pi_h^{k+1}(s), Q_h^k(s, \cdot) \rangle$.
- 19: **end for**
- 20: **end for**

gorithm 3 applied to finite horizon problems which does not use the best response observation. A core step in the proof is to show that the regret of LSVI-UCB is still $\mathcal{O}(\sqrt{K})$ if the cost function is not fixed but it is observed in advanced by the agent.

Theorem 6. Let us consider $K = \mathcal{O}\left(\frac{H^4 d^3 \log(dH/(\epsilon\delta))}{\epsilon^2}\right)$ iterations of Algorithm 4 run with $\alpha = \sqrt{\frac{1}{2K}}$ and ex-

pert demonstrations $\tau_E = \frac{2H^2 d \log(2d/\delta)}{\epsilon_E^2}$. Moreover, let \hat{k} be an iteration index sampled uniformly at random from $\{1, 2, \dots, K\}$, then it holds that with probability $1 - 3\delta$, it holds that $\mathbb{E}_{\hat{k}} \left[V_1^{\pi^{\hat{k}}}(s_1; \mathbf{c}_{\text{true}}) - V_1^{\pi^E}(s_1; \mathbf{c}_{\text{true}}) \right] \leq \mathcal{O}(\epsilon_E + \epsilon)$.

Remark 3. Unfortunately, the best response idea does not help improving the infinite horizon result because the use of greedy policies makes the term (Shift) impossible to control.

8. Future directions

Reducing the number of expert trajectories (Rajaraman et al., 2021) showed that under known transitions for a particular choice of features and when the linear expert occupancy measure is uniform over the state space the necessary expert trajectories are $\mathcal{O}\left(\frac{H^{3/2} d^2}{\epsilon}\right)$.

In the linear MDP case, under the same assumption on the features we can show that the same amount of expert trajectories is sufficient even for the unknown transition case. Moreover, the algorithm we propose is computationally efficient while it is unclear how the output policy in linear Mimic-MD (Rajaraman et al., 2021) can be computed with complexity independent on the number of states and actions. We detail this in Appendix H.

However, the assumptions under which this result is obtained are clearly too strong. Hence, an interesting open question on this front is to see if the same improvement in terms of expert trajectory can be achieved under weaker conditions.

A first step has been already made in (Swamy et al., 2022). They still requires a linear expert model but avoids the uniform expert occupancy measure assumption replacing it with the bounded density assumption (see (Swamy et al., 2022, Assumption 9)). A natural follow up would be to investigate if the same expert trajectory bound can be obtained bringing the Linear MDP assumption but dropping the linear expert and the other assumptions used in (Rajaraman et al., 2021; Swamy et al., 2022). An intermediate step could benefit from using the persistent excitation assumption which is the natural counterpart of the bounded density assumption (Swamy et al., 2022, Assumption 9) in Linear MDPs.

Improving dependence on ϵ , d and H For what concerns the finite horizon case we presented BRIG which has a dependence on ϵ which can not be improved further if not bypassing the reduction to online learning in MDPs. However the dependence in d and H can be improved. To circumvent this problem, we could think that one could apply LSVI-UCB++ (He et al., 2023). However, it turns out that LSVI-UCB++ fails if the cost changes adversarially even if the learner knows the cost at the next round. Therefore, a

new algorithm design is needed to improve by a factor dH upon the MDP trajectory complexity of Algorithm 4.

Results for Linear Mixture MDPs Analogous ideas can be used for the case of Linear Mixture MDPs. In particular, one can use the same structure as in Algorithm 3 but replacing an algorithm that deals with adversarial losses in Linear Mixture MDPs. For the infinite horizon case one can use (Moulin & Neu, 2023) while for the finite horizon case one can choose (He et al., 2022). In the latter case one can improve the Sample Complexity of OGAIL (Liu et al., 2022) to $\mathcal{O}(H^3 d^2 \epsilon^{-2})$. Moreover, we do not think that the Linear Expert assumption (Rajaraman et al., 2021) is meaningful in Linear Mixture MDPs because this would require the learner to know in advance the features $\int_{\mathcal{S}} \phi(s, a, s') V^*(s') ds'$ where V^* is the optimal state value function.

Extension to Bilinear Classes The current results can be extended to Bilinear Classes (Du et al., 2021) at least in the finite horizon case using the trick of updating first the cost, i.e. the same trick that allowed us to obtain $\mathcal{O}(\epsilon^{-2})$ MDP trajectories bound in the finite horizon case using LSVI-UCB.

9. Conclusions

In this paper, we proposes ILARL which greatly reduces the number of MDP trajectories in imitation learning in Linear MDP and BRIG that provides a further improvement for the finite horizon case. Both results build on the connection between imitation learning and MDPs with adversarial losses.

As mentioned, there is a number of exciting future directions. In particular, the estimation of $\widehat{\Phi^T d^{\pi_E}}$ could be carried out with **fewer expert trajectories** using trajectory access to the MDP. This observation has been proven successful having access to the exact transitions of the MDP in the tabular case (Rajaraman et al., 2020) or under linear function approximation with further assumption on the expert policy and the feature distribution (Swamy et al., 2022; Rajaraman et al., 2021). Whether the same is possible for general stochastic experts in Linear MDP is an interesting open question.

Finally, a **better sample complexity** can be achieved designing better no regret algorithm for infinite horizon adversarial discounted linear MDP with full-information feedback and apply them in Step 9 of Algorithm 3 building for example on the recent result for the finite horizon case in (Sherman et al., 2023a).

Acknowledgements

Luca Viano is thankful to Antoine Moulin, Gergely Neu and Adrian Müller for insightful and enjoyable discussions. This work is funded (in part) through a PhD fellowship of the Swiss Data Science Center, a joint venture between EPFL and ETH Zurich. Innovation project supported by Innosuisse (contract agreement 100.960 IP-ICT). This work was supported by Hasler Foundation Program: Hasler Responsible AI (project number 21043), Research was sponsored by the Army Research Office and was accomplished under Grant Number W911NF-24-1-0048 This work was supported by the Swiss National Science Foundation (SNSF) under grant number 200021_205011. Luca Viano acknowledges travel support from ELISE (GA no 951847).

Impact Statement

The work is mainly theoretical therefore we do not foresee any direct impact in the society.

References

- Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2011.
- Abbasi-Yadkori, Y., Bartlett, P. L., and Szepesvari, C. Online learning in markov decision processes with adversarially chosen transition probability distributions, 2013.
- Abbasi-Yadkori, Y., Bartlett, P., Bhatia, K., Lazic, N., Szepesvari, C., and Weisz, G. Politex: Regret bounds for policy iteration using expert prediction. In *International Conference on Machine Learning (ICML)*, 2019a.
- Abbasi-Yadkori, Y., Lazic, N., Szepesvari, C., and Weisz, G. Exploration-enhanced politex. *arXiv:1908.10479*, 2019b.
- Agarwal, A., Jiang, N., Kakade, S. M., and Sun, W. Reinforcement learning: Theory and algorithms. *CS Dept., UW Seattle, Seattle, WA, USA, Tech. Rep.*, 32, 2019.
- Agarwal, A., Kakade, S., Krishnamurthy, A., and Sun, W. Flambe: Structural complexity and representation learning of low rank MDPs. *Advances in neural information processing systems (NeurIPS)*, 2020a.
- Agarwal, A., Kakade, S. M., Lee, J. D., and Mahajan, G. On the theory of policy gradient methods: Optimality, approximation, and distribution shift, 2020b.
- Boularias, A., Kober, J., and Peters, J. Relative entropy inverse reinforcement learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 182–189. JMLR Workshop and Conference Proceedings, 2011.
- Cai, Q., Yang, Z., Jin, C., and Wang, Z. Provably efficient exploration in policy optimization. In *International Conference on Machine Learning (ICML)*, 2020.
- Cesa-Bianchi, N. and Lugosi, G. *Prediction, learning, and games*. Cambridge university press, 2006.
- Charpentier, A., Elie, R., and Remlinger, C. Reinforcement learning in economics and finance. *arXiv:20031004*, 2020.
- Cohen, A., Kaplan, H., Mansour, Y., and Rosenberg, A. Near-optimal regret bounds for stochastic shortest path, 2020.
- Dadashi, R., Hussenot, L., Geist, M., and Pietquin, O. Primal Wasserstein imitation learning. In *International Conference on Learning Representations (ICLR)*, 2021.
- Dai, Y., Luo, H., Wei, C.-Y., and Zimmert, J. Refined regret for adversarial mdps with linear function approximation. *arXiv preprint arXiv:2301.12942*, 2023.
- Du, S., Kakade, S., Lee, J., Lovett, S., Mahajan, G., Sun, W., and Wang, R. Bilinear classes: A structural framework for provable generalization in rl. In *International Conference on Machine Learning*, pp. 2826–2836. PMLR, 2021.
- Duan, Y., Jia, Z., and Wang, M. Minimax-optimal off-policy evaluation with linear function approximation. In *International Conference on Machine Learning (ICML)*, 2020.
- Fu, J., Luo, K., and Levine, S. Learning robust rewards with adversarial inverse reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2018.
- Garg, D., Chakraborty, S., Cundy, C., Song, J., and Ermon, S. IQ-learn: Inverse soft-Q learning for imitation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Hao, B., Lattimore, T., Szepesvári, C., and Wang, M. Online sparse reinforcement learning. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2021.
- He, J., Zhou, D., and Gu, Q. Near-optimal policy optimization algorithms for learning adversarial linear mixture mdps. In *International Conference on Artificial Intelligence and Statistics*, pp. 4259–4280. PMLR, 2022.
- He, J., Zhao, H., Zhou, D., and Gu, Q. Nearly minimax optimal reinforcement learning for linear markov decision processes. In *International Conference on Machine Learning*, pp. 12790–12822. PMLR, 2023.

- Ho, J. and Ermon, S. Generative adversarial imitation learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- Ho, J., Gupta, J. K., and Ermon, S. Model-free imitation learning with policy optimization. In *International Conference on Machine Learning (ICML)*, 2016.
- Jiang, N., Krishnamurthy, A., Agarwal, A., Langford, J., and Schapire, R. E. Contextual decision processes with low bellman rank are pac-learnable. In *International Conference on Machine Learning*, pp. 1704–1713. PMLR, 2017.
- Jin, C., Yang, Z., Wang, Z., and Jordan, M. I. Provably efficient reinforcement learning with linear function approximation, 2019.
- Knox, W. B., Allievi, A., Banzhaf, H., Schmitt, F., and Stone, P. Reward (mis)design for autonomous driving, 2021.
- Lazic, N., Yin, D., Farajtabar, M., Levine, N., Gorur, D., Harris, C., and Schuurmans, D. A maximum-entropy approach to off-policy evaluation in average-reward MDPs. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Li, Y. and Zhang, C. On efficient online imitation learning via classification. *Advances in Neural Information Processing Systems*, 35:32383–32397, 2022.
- Liu, Z., Zhang, Y., Fu, Z., Yang, Z., and Wang, Z. Learning from demonstration: Provably efficient adversarial policy imitation with linear function approximation. In *International Conference on Machine Learning (ICML)*, 2022.
- Luo, H., Wei, C.-Y., and Lee, C.-W. Policy optimization in adversarial mdps: Improved exploration via dilated bonuses. *Advances in Neural Information Processing Systems*, 34:22931–22942, 2021.
- Moulin, A. and Neu, G. Optimistic planning by regularized dynamic programming. *arXiv preprint arXiv:2302.14004*, 2023.
- Munos, R. Error bounds for approximate policy iteration. In *Proceedings of the Twentieth International Conference on International Conference on Machine Learning*, pp. 560–567, 2003.
- Ni, T., Sikchi, H., Wang, Y., Gupta, T., Lee, L., and Eysenbach, B. f-irl: Inverse reinforcement learning via state marginal matching. In *Conference on Robot Learning*, pp. 529–551. PMLR, 2021.
- Orabona, F. A modern introduction to online learning, 2023.
- Osa, T., Pajarinen, J., Neumann, G., Bagnell, J., Abbeel, P., and Peters, J. An algorithmic perspective on imitation learning. *Foundations and Trends in Robotics*, 2018.
- Pomerleau, D. A. Efficient training of artificial neural networks for autonomous navigation. *Neural Computation*, 3(1):88–97, 1991.
- Puterman, M. L. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., USA, 1st edition, 1994.
- Rajaraman, N., Yang, L., Jiao, J., and Ramchandran, K. Toward the fundamental limits of imitation learning. *Advances in Neural Information Processing Systems*, 33: 2914–2924, 2020.
- Rajaraman, N., Han, Y., Yang, L., Liu, J., Jiao, J., and Ramchandran, K. On the value of interaction and function approximation in imitation learning. *Advances in Neural Information Processing Systems*, 34:1325–1336, 2021.
- Reddy, S., Dragan, A. D., and Levine, S. SQIL: imitation learning via regularized behavioral cloning. *arXiv:1905.11108*, 2019.
- Ross, S. and Bagnell, D. Efficient reductions for imitation learning. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2010.
- Ross, S., Gordon, G., and Bagnell, D. A reduction of imitation learning and structured prediction to no-regret online learning. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011.
- Shani, L., Efroni, Y., Rosenberg, A., and Mannor, S. Optimistic policy optimization with bandit feedback. In *International Conference on Machine Learning*, pp. 8604–8613. PMLR, 2020.
- Shani, L., Zahavy, T., and Mannor, S. Online apprenticeship learning. *arXiv:2102.06924*, 2021.
- Sherman, U., Cohen, A., Koren, T., and Mansour, Y. Rate-optimal policy optimization for linear markov decision processes. *arXiv preprint arXiv:2308.14642*, 2023a.
- Sherman, U., Koren, T., and Mansour, Y. Improved regret for efficient online reinforcement learning with linear function approximation, 2023b.
- Swamy, G., Rajaraman, N., Peng, M., Choudhury, S., Bagnell, J., Wu, S. Z., Jiao, J., and Ramchandran, K. Minimax optimal online imitation learning via replay estimation. *Advances in Neural Information Processing Systems*, 35: 7077–7088, 2022.

- Swamy, G., Wu, D., Choudhury, S., Bagnell, D., and Wu, S. Inverse reinforcement learning without reinforcement learning. In *International Conference on Machine Learning*, pp. 33299–33318. PMLR, 2023.
- Syed, U. and Schapire, R. E. A game-theoretic approach to apprenticeship learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2007.
- Viano, L., Huang, Y.-T., Kamalaruban, P., Weller, A., and Cevher, V. Robust inverse reinforcement learning under transition dynamics mismatch. *Advances in Neural Information Processing Systems*, 34:25917–25931, 2021.
- Viano, L., Kamoutsis, A., Neu, G., Krawczuk, I., and Cevher, V. Proximal point imitation learning. *Advances in Neural Information Processing Systems*, 35:24309–24326, 2022.
- Watson, J., Huang, S. H., and Heess, N. Coherent soft imitation learning, 2023.
- Wu, F., Ke, J., and Wu, A. Inverse reinforcement learning with the average reward criterion. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=YFSrf8aciU>.
- Xu, T., Li, Z., Yu, Y., and Luo, Z.-Q. Provably efficient adversarial imitation learning with unknown transitions. *arXiv preprint arXiv:2306.06563*, 2023.
- Zeng, S., Hong, M., and Garcia, A. Structural estimation of markov decision processes in high-dimensional state space with finite-time guarantees. *arXiv preprint arXiv:2210.01282*, 2022a.
- Zeng, S., Li, C., Garcia, A., and Hong, M. Maximum-likelihood inverse reinforcement learning with finite-time guarantees. *Advances in Neural Information Processing Systems*, 35:10122–10135, 2022b.
- Zeng, S., Li, C., Garcia, A., and Hong, M. Understanding expertise through demonstrations: A maximum likelihood framework for offline inverse reinforcement learning. *arXiv preprint arXiv:2302.07457*, 2023.
- Zhong, H. and Zhang, T. A theoretical analysis of optimistic proximal policy optimization in linear markov decision processes. *arXiv preprint arXiv:2305.08841*, 2023.
- Zinkevich, M. Online convex programming and generalized infinitesimal gradient ascent. In Fawcett, T. and Mishra, N. (eds.), *Machine Learning, Proceedings of the Twentieth International Conference (ICML 2003), August 21-24, 2003, Washington, DC, USA*, pp. 928–936. AAAI Press, 2003.

A. On the β of the persistent excitation assumption

Lemma 1. *Let the persistent excitation assumption holds, then it holds that $\beta \leq d^{-1}$.*

Proof.

$$\begin{aligned}
 \beta &\leq \lambda_{\min}(\mathbb{E}_{s,a \sim d^{\pi^k}} \phi(s,a)\phi(s,a)^T) \leq \frac{1}{d} \text{Trace}(\mathbb{E}_{s,a \sim d^{\pi^k}} \phi(s,a)\phi(s,a)^T) \\
 &= \frac{1}{d} \mathbb{E}_{s,a \sim d^{\pi^k}} \text{Trace}(\phi(s,a)\phi(s,a)^T) \\
 &= \frac{1}{d} \mathbb{E}_{s,a \sim d^{\pi^k}} \text{Trace}(\phi(s,a)^T \phi(s,a)) \\
 &= \frac{1}{d} \mathbb{E}_{s,a \sim d^{\pi^k}} \|\phi(s,a)\|_2^2 \\
 &\leq \frac{1}{d} \mathbb{E}_{s,a \sim d^{\pi^k}} \|\phi(s,a)\|_1^2 \\
 &\leq \frac{1}{d} \max_{s,a} \|\phi(s,a)\|_1^2 \\
 &= \frac{1}{d} \|\Phi\|_{1,\infty}^2 \leq \frac{1}{d}.
 \end{aligned}$$

where the last step follows from $\|\Phi\|_{1,\infty} \leq 1$ as assumed in Assumptions 1 and 2. \square

Using this result, we obtain that the dimension dependence in the bound in Viano et al. 2022 is in the best case d^8 . Therefore, our new algorithm improves the dimension dependence as well.

B. Interaction Protocol

Protocol 1 Interaction in Adversarial MDPs

```

1: for Episode index  $k \in [1, K]$  do
2:   Sample initial state  $s_1^k \sim \nu_0$ 
3:   if Finite Horizon then
4:     for stage  $h \in [1, H]$  do
5:       The learner plays an action sampled from the policy  $a_h^k \sim \pi_h^k(\cdot | s_h^k)$ .
6:       The environment sample next state  $s_h^k \sim P_h(\cdot | s_h^k, a_h^k)$ .
7:       The agent observes the vector  $c_h^k$ .
8:     end for
9:   end if
10:  if Infinite Horizon then
11:    Initialize  $Z = 0, i = 1, s^1 \sim \nu_0$ .
12:    while  $Z == 0$  do
13:      The learner plays an action sampled from the policy  $a^i \sim \pi^k(\cdot | s^i)$ .
14:      The environment sample next state  $s'^i \sim P(\cdot | s^i, a^i), s^{i+1} = s'^i$ .
15:      // Restart with probability  $1 - \gamma$ .
16:      Sample  $Z \sim \text{Bernoulli}(1 - \gamma)$ .
17:    end while
18:    The agent observes the vector  $c^k$ .
19:  end if
20:  The learner chooses her next policy, i.e.  $\pi^{k+1}$ .
21: end for
    
```

C. Additional Future directions: Extension to Bilinear Classes

The current results can be extended to Bilinear Classes (Du et al., 2021) at least in the finite horizon case. To this end we would need the observation that we can update the cost first and then using an algorithm which is allowed to see the next cost vector one round in advance. This is the same fact that allowed us to obtain $\mathcal{O}(\epsilon^{-2})$ sample complexity bound in the finite horizon case using LSVI-UCB.

In the following we present an informal discussion of the proof technique that would prove polynomial sample complexity for Imitation Learning in bilinear classes.

If the cost at round k is known before the learner needs to take an action the agent can form the discrepancy function

$$\ell^k(s_h, a_h, s_{h+1}, g) = Q_{h,g}(s_h, a_h) + c_h^k - V_{h+1,g}(s_{h+1})$$

where the upper script k highlights the fact that the discrepancy function depends on the adversarial cost c^k . At each round, we can then compute

$$\operatorname{argmax}_{g \in \mathcal{H}} V_{0,g}(s_1) \quad \text{s.t.} \quad \frac{1}{m} \sum_{\tau=1}^k \sum_{i=1}^m \ell^k(s_h^i, a_h^i, s_{h+1}^i, g) \leq k \epsilon_{\text{gen},k}^2(m, \delta) \quad \forall h \in [H]$$

Where the generalization error at round k denoted $\epsilon_{\text{gen},k}$ satisfies that

$$\sup_{g \in \mathcal{H}} \left| \frac{1}{m} \sum_{i=1}^m \ell^k(s_h^i, a_h^i, s_{h+1}^i, g) - \mathbb{E}_{s,a \sim d^{\pi^k}, s' \sim P(\cdot|s,a)} [\ell^k(s_h, a_h, s_{h+1}, g)] \right| \leq \epsilon_{\text{gen},k}(m, \delta)$$

with probability at least $1 - \delta$.

At this point, we modify the Bilinear Classes assumption to keep into account the adversarial costs setting as follows. We assume that all the adversarial costs belongs to a convex set \mathcal{C} . Then we consider an MDP for which there exists a function $f^* \in \mathcal{H}$ such that for all $c \in \mathcal{C}$, it holds that

$$\left| \mathbb{E}_{s,a \sim d_h^{\pi_f}} [Q_{h,f}(s_h, a_h) + c_h(s_h, a_h) - V_{h+1,f}] \right| \leq |\langle W_h^c(f) - W_h^c(f^*), X_h^c(f) \rangle|$$

and

$$\left| \mathbb{E}_{s,a \sim d_h^{\pi_f}} [Q_{h,g}(s_h, a_h) + c_h(s_h, a_h) - V_{h+1,g}] \right| \leq |\langle W_h^c(g) - W_h^c(f^*), X_h^c(f) \rangle|$$

This modified assumption for Bilinear classes with time changing rewards implies that the comparator hypothesis f^* is realizable for all k , that is, for all k it holds that

$$\frac{1}{m} \sum_{\tau=1}^k \sum_{i=1}^m \ell^k(s_h^i, a_h^i, s_{h+1}^i, f^*) \leq k \epsilon_{\text{gen},k}^2(m, \delta)$$

so optimism holds and with the same steps in (Du et al., 2021) it can be proven that:

$$\begin{aligned} \text{Regret}(K, \pi^*) &\leq \sum_{k=1}^K \left| \mathbb{E}_{s,a \sim d_h^{\pi_{f^k}}} [Q_{h,f^k}(s_h, a_h) + c_h(s_h, a_h) - V_{h+1,f^k}] \right| \\ &\leq \sum_{k=1}^K \left| \langle W_h^{r^k}(f^k) - W_h^{r^k}(f^*), X_h^{r^k}(f^k) \rangle \right| \end{aligned}$$

Then, using (Du et al., 2021, Equation 8) and the elliptical potential lemma we obtain

$$\frac{1}{K} \text{Regret}(K, \pi^*) = \mathcal{O} \left(\sqrt{K \max_{k \in [K]} \epsilon_{\text{gen},k}^2(m, \delta) \left(\exp \frac{\log K}{K} - 1 \right)} \right)$$

Then, denoting $\epsilon_{\text{gen}}(m, \delta) = \max_{k \in [K]} \epsilon_{\text{gen},k}(m, \delta)$ and choosing $K = \mathcal{O}(\log \epsilon_{\text{gen}}^{-2}(m, \delta))$ gives $\frac{1}{K} \text{Regret}(K, \pi^*) \leq \mathcal{O}(\epsilon_{\text{gen}}(m, \delta))$.

This concludes the regret proof for the π -player. The player updating the sequences of cost can still use OGD projecting on the set \mathcal{C} .

Our algorithm for the finite horizon setting, BRIG, uses greedy policies. In this case, (Du et al., 2021) showed that the generalization error can be controlled effectively in many instances such as Linear Q^*/V^* , Low Occupancy measure models (Du et al., 2021), Bellman complete models (Munos, 2003) and finite Bellman Rank (Jiang et al., 2017).

However in the infinite horizon case we need to use regularization for which it is currently not known if $\text{gen}(m, \delta)$ can be controlled effectively. This is again related to the issue with the covering number of softmax policies in linear MDPs (see (Sherman et al., 2023a; Zhong & Zhang, 2023)). This is an interesting open question.

A final comment is that the algorithm proposed in (Du et al., 2021) for bilinear classes is not computationally efficient is general. Therefore also its adversarial extension presented above will have this drawback. In this paper we focused on the smaller class of Linear MDP for which we provide a computationally efficient algorithm.

D. Omitted Proofs

D.1. Proof of Lemma 2

Lemma 2. Consider the MDP $M = (\mathcal{S}, \mathcal{A}, \gamma, P, \mathbf{c})$ and two policies $\pi, \pi' : \mathcal{S} \rightarrow \Delta_{\mathcal{A}}$. Then consider for any $\widehat{Q} \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}|}$ and $\widehat{V}^\pi(s) = \langle \pi(\cdot|s), \widehat{Q}(s, \cdot) \rangle$ and $Q^{\pi'}, V^{\pi'}$ be respectively the state action and state value function of the policy π in MDP M and let $E \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{A}| \times |\mathcal{S}|}$ be the matrix such that for an arbitrary vector $f \in \mathbb{R}^{|\mathcal{S}|}$ it holds $(Ef)(s, a) = f(s)$. Then, it holds that

$$\langle \boldsymbol{\nu}_0, \widehat{V}^\pi - V^{\pi'} \rangle = \frac{1}{1-\gamma} \left(\langle d^{\pi'}, \widehat{Q} - \mathbf{c} - \gamma P \widehat{V}^\pi \rangle + \langle d^{\pi'}, E \widehat{V}^\pi - \widehat{Q} \rangle \right)$$

Proof. Consider the Bellman equation in vector form, i.e. $Q^{\pi'} = \mathbf{c} + \gamma P V^{\pi'}$. Then, let us add and subtract the term $\mathbf{c} + \gamma P \widehat{V}^\pi$ and let us consider on both sides the inner product with the occupancy measure $d^{\pi'}$.

$$\begin{aligned} \langle d^{\pi'}, \widehat{Q} \rangle &= \langle d^{\pi'}, \widehat{Q} - \mathbf{c} - \gamma P \widehat{V}^\pi \rangle + \langle d^{\pi'}, \mathbf{c} + \gamma P \widehat{V}^\pi \rangle \\ &= \langle d^{\pi'}, \widehat{Q} - \mathbf{c} - \gamma P \widehat{V}^\pi \rangle + (1-\gamma) \langle \boldsymbol{\nu}_0, V^{\pi'} \rangle + \langle \gamma P^\top d^{\pi'}, \widehat{V}^\pi \rangle \\ &= \langle d^{\pi'}, \widehat{Q} - \mathbf{c} - \gamma P \widehat{V}^\pi \rangle + (1-\gamma) \langle \boldsymbol{\nu}_0, V^{\pi'} - \widehat{V}^\pi \rangle + \langle E^\top d^{\pi'}, \widehat{V}^\pi \rangle \end{aligned}$$

Rearranging the terms leads to the conclusion. □

D.2. Proof of Theorem 4

Proof. Let $V^{\pi^k, k} \in [-(1-\gamma)^{-1}, (1-\gamma)^{-1}]^{|S|}$ be a compact notation for $V^{\pi^k}(\cdot; \mathbf{c}^k)$, then

$$\begin{aligned}
 \text{Regret}(K; \pi^*) &\triangleq \sum_{k=1}^K \left\langle d^{\pi^k} - d^{\pi^*}, \mathbf{c}^k \right\rangle \\
 &= (1-\gamma) \sum_{k=1}^K \left\langle \boldsymbol{\nu}_0, V^{\pi^k, k} - V^{\pi^*, k} \right\rangle \\
 &= \sum_{k=1}^K \mathbb{E}_{s \sim d^{\pi^*}} \left[\langle Q^k(s, \cdot), \pi^k(s) - \pi^*(s) \rangle \right] && \text{(OMD)} \\
 &\quad + \sum_{k=1}^K \mathbb{E}_{s, a \sim d^{\pi^*}} \left[Q^{k+1}(s, a) - \mathbf{c}^k(s, a) - \gamma P V^k(s, a) \right] && \text{(Optimism 1)} \\
 &\quad + \sum_{k=1}^K \mathbb{E}_{s, a \sim d^{\pi^k}} \left[\mathbf{c}^k(s, a) + \gamma P V^k(s, a) - Q^{k+1}(s, a) \right] && \text{(Optimism 2)} \\
 &\quad + \sum_{k=1}^K \mathbb{E}_{s, a \sim d^{\pi^k}} \left[Q^{k+1}(s, a) - Q^k(s, a) \right] && \text{(Shift 1)} \\
 &\quad + \sum_{k=1}^K \mathbb{E}_{s, a \sim d^{\pi^*}} \left[Q^k(s, a) - Q^{k+1}(s, a) \right] && \text{(Shift 2)}
 \end{aligned}$$

Then, we have that [Optimism 1](#), [Optimism 2](#) can be bounded using [Lemma 4](#) while for [Shift 1](#) we crucially rely on the regularization of the policy improvement step and on the fact that the policy is updated only every τ steps. Both these observations allow to derive

$$\begin{aligned}
 \sum_{k=1}^K \mathbb{E}_{s, a \sim d^{\pi^k}} \left[Q^k(s, a) - Q^{k+1}(s, a) \right] &\leq \sum_{k=2}^K \sum_{s, a} Q^k(s, a) \left(d^{\pi^k}(s, a) - d^{\pi^{k-1}}(s, a) \right) + \sum_{s, a} Q^1(s, a) d^{\pi^1}(s, a) \\
 &\leq \sum_{k=2}^K \|Q^k(s, a)\|_{\infty} \|d^{\pi^k} - d^{\pi^{k-1}}\|_1 \\
 &\leq \frac{1}{1-\gamma} \sum_{k=2}^K \|d^{\pi^k} - d^{\pi^{k-1}}\|_1 \\
 &= \frac{1}{1-\gamma} \sum_{j=2}^{\lfloor K/\tau \rfloor} \|d^{\pi^{(j)}} - d^{\pi^{(j-1)}}\|_1
 \end{aligned}$$

At this point applying Pinsker's inequality and ([Moulin & Neu, 2023, Lemma A.1](#)) we obtain

$$\begin{aligned}
 \|d^{\pi^{(j)}} - d^{\pi^{(j-1)}}\|_1 &\leq \sqrt{2D_{KL}(d^{\pi^{(j)}}, d^{\pi^{(j-1)}})} \leq \sqrt{\frac{2}{1-\gamma} \mathbb{E}_{s \sim d^{\pi^{(j)}}} D_{KL}(\pi^{(j)}(\cdot|s), \pi^{(j-1)}(\cdot|s))} \\
 &\leq \sqrt{\frac{2\eta}{1-\gamma} \mathbb{E}_{s, a \sim d^{\pi^{(j)}}} \bar{Q}^{(j-1)}} \\
 &\leq \frac{\sqrt{2\eta}}{1-\gamma}
 \end{aligned}$$

All in all, we have that

$$\text{(Shift 1)} \leq \frac{1}{(1-\gamma)^2} \sum_{j=2}^{\lfloor K/\tau \rfloor} \sqrt{2\eta} = \frac{\sqrt{2\eta}K}{(1-\gamma)^2\tau}$$

For the second shift term, we can use a trivial telescoping argument to obtain that [\(Shift 2\)](#) $\leq (1-\gamma)^{-1}$.

The terms (Optimism 1) and (Optimism 2) can be bounded exactly as in the finite horizon case thanks to Lemma 4 we have that with probability $1 - \delta$

$$(\text{Optimism 1}) + (\text{Optimism 2}) \leq 2 \sum_{k=1}^K \mathbb{E}_{s,a \sim d^{\pi^k}} [b^k(s, a)]. \quad (5)$$

Therefore, we just need to adapt the argument for bounding the exploration term $2 \sum_{k=1}^K \mathbb{E}_{s,a \sim d^{\pi^k}} [b^k(s, a)]$ in the infinite horizon case. We start by exploiting the fact that both b^k and π^k change in fact only every τ updates. So we have

$$\begin{aligned} \sum_{k=1}^K \mathbb{E}_{s,a \sim d^{\pi^k}} [b^k(s, a)] &= \tau \sum_{j=1}^{K/\tau} \mathbb{E}_{s,a \sim d^{\pi^{(j)}}} [b^{(j)}(s, a)] \\ &= \sum_{j=1}^{K/\tau} \mathbb{E}_{\mathcal{D}^{(j)} \sim d^{\pi^{(j)}}} \left[\sum_{s,a \in \mathcal{D}^{(j)}} b^{(j)}(s, a) \right] \\ &= \sum_{j=1}^{K/\tau} \mathbb{E}_{\mathcal{D}^{(j)} \sim d^{\pi^{(j)}}} \left[\sum_{s,a \in \mathcal{D}^{(j)}} \beta \|\phi(s, a)\|_{(\Lambda^{(j)})^{-1}} \right] \end{aligned}$$

At this point, fixing an arbitrary order for the state action pairs in $\mathcal{D}^{(j)}$ we can define for any $i \in [1, \tau]$ the matrix

$$\Lambda_i^{(j)} = \sum_{\ell=1}^i \phi(s^\ell, a^\ell) \phi(s^\ell, a^\ell)^T + \lambda I$$

Then, it holds that

$$\begin{aligned} \sum_{j=1}^{K/\tau} \mathbb{E}_{\mathcal{D}^{(j)} \sim d^{\pi^{(j)}}} \left[\sum_{s,a \in \mathcal{D}^{(j)}} \beta \|\phi(s, a)\|_{(\Lambda^{(j)})^{-1}} \right] &= \sum_{j=1}^{K/\tau} \mathbb{E}_{\mathcal{D}^{(j)} \sim d^{\pi^{(j)}}} \left[\sum_{\ell=1}^{\tau} \beta \|\phi(s^\ell, a^\ell)\|_{(\Lambda^{(j)})^{-1}} \right] \\ &\leq \sum_{j=1}^{K/\tau} \mathbb{E}_{\mathcal{D}^{(j)} \sim d^{\pi^{(j)}}} \left[\sum_{\ell=1}^{\tau} \beta \|\phi(s^\ell, a^\ell)\|_{(\Lambda_\ell^{(j)})^{-1}} \right] \end{aligned}$$

At this point, we can notice that for any index pair ℓ, j it holds that

$$\|\phi(s^\ell, a^\ell)\|_{(\Lambda_\ell^{(j)})^{-1}} \leq \sqrt{\lambda_{\max}((\Lambda_\ell^{(j)})^{-1})} \|\phi(s^\ell, a^\ell)\| \leq \sqrt{\frac{1}{\lambda_{\min}(\Lambda_\ell^{(j)})}} \leq 1$$

where the norm $\|\phi(s^\ell, a^\ell)\|$ is upper bounded by 1 thanks to the Linear MDP assumption. Then, via (Sherman et al., 2023b, Lemma F.1) (where the random variable X_i is in this context $\sum_{\ell=1}^{\tau} \beta \|\phi(s^\ell, a^\ell)\|_{(\Lambda_\ell^{(j)})^{-1}}$ which is supported in $[0, \beta\tau]$) we can continue upper bounding with probability $1 - \delta$ the bonus sum as

$$\begin{aligned} \sum_{k=1}^K \mathbb{E}_{s,a \sim d^{\pi^k}} [b^k(s, a)] &\leq 2 \sum_{j=1}^{K/\tau} \sum_{\ell=1}^{\tau} \beta \|\phi(s^\ell, a^\ell)\|_{(\Lambda_\ell^{(j)})^{-1}} + 4\beta\tau \log(2K/(\tau\delta)) \\ &= 2 \sum_{j=1}^{K/\tau} \sum_{\ell=1}^{\tau} \beta \sqrt{\phi(s^\ell, a^\ell)^T (\Lambda_\ell^{(j)})^{-1} \phi(s^\ell, a^\ell)} + 4\beta\tau \log(2K/(\tau\delta)) \\ &\leq 2 \sum_{j=1}^{K/\tau} \beta \sqrt{\tau \sum_{\ell=1}^{\tau} \phi(s^\ell, a^\ell)^T (\Lambda_\ell^{(j)})^{-1} \phi(s^\ell, a^\ell)} + 4\beta\tau \log(2K/(\tau\delta)) \\ &\leq 2 \sum_{j=1}^{K/\tau} \beta \sqrt{\tau \sum_{i=1}^d \log(1 + \lambda_i)} + 4\beta\tau \log(2K/(\tau\delta)) \end{aligned}$$

where the last inequality uses (Cesa-Bianchi & Lugosi, 2006, Lemma 11.11) and the notation $\{\lambda_i\}_{i=1}^d$ stands for the eigenvalues of the matrix $\Lambda^{(j)} - I$. At this point we can recognize the determinant inside the log and use the determinant trace inequality.

$$\begin{aligned}
 \sum_{j=1}^{K/\tau} \beta \sqrt{\tau \sum_{i=1}^d \log(1 + \lambda_i)} &= \sum_{j=1}^{K/\tau} \beta \sqrt{\tau \log \left(\prod_{i=1}^d 1 + \lambda_i \right)} \\
 &= \sum_{j=1}^{K/\tau} \beta \sqrt{\tau \log(\det(\Lambda^{(j)}))} \\
 &\leq \sum_{j=1}^{K/\tau} \beta \sqrt{\tau d \log \left(\frac{\text{Trace}(\Lambda^{(j)})}{d} \right)} \\
 &\leq \sum_{j=1}^{K/\tau} \beta \sqrt{\tau d \log \left(\frac{d + \text{Trace}(\sum_{\ell=1}^{\tau} \phi(s^\ell, a^\ell) \phi(s^\ell, a^\ell)^T)}{d} \right)} \\
 &\leq \sum_{j=1}^{K/\tau} \beta \sqrt{\tau d \log \left(\frac{d + \tau \max_{\ell} \text{Trace}(\phi(s^\ell, a^\ell) \phi(s^\ell, a^\ell)^T)}{d} \right)} \\
 &= \sum_{j=1}^{K/\tau} \beta \sqrt{\tau d \log \left(\frac{d + \tau \max_{\ell} \phi(s^\ell, a^\ell)^T \phi(s^\ell, a^\ell)}{d} \right)} \\
 &\leq \sum_{j=1}^{K/\tau} \beta \sqrt{\tau d \log \left(\frac{d + \tau}{d} \right)} \\
 &\leq \beta K \sqrt{\frac{d \log 2\tau d}{\tau}}.
 \end{aligned}$$

Hence, with probability $1 - 2\delta$ (union bound between the event under which Equation (5) holds and the the application of the concentration result (Sherman et al., 2023b, Lemma F.1) in bounding the exploration bonuses sum), it holds that

$$\begin{aligned}
 \text{Regret}(K; \pi^*) &\leq \sum_{k=1}^K \mathbb{E}_{s \sim d^{\pi^*}} [\langle Q^k(s, \cdot), \pi^k(s) - \pi^*(s) \rangle] + 4\beta K \sqrt{\frac{d \log(2\tau d)}{\tau}} + 8\beta\tau \log(2K/(\tau\delta)) \\
 &\quad + \frac{\sqrt{2\eta}K}{(1-\gamma)^2\tau} + \frac{1}{(1-\gamma)}
 \end{aligned}$$

The last step is to bound (OMD) invoking Lemma 9 and noticing that the gradient norm is upper bounded by $\frac{1}{1-\gamma}$. This gives

$$\sum_{k=1}^K \mathbb{E}_{s \sim d^{\pi^*}} [\langle Q^k(s, \cdot), \pi^k(s) - \pi^*(s) \rangle] \leq \frac{\tau \log |\mathcal{A}|}{\eta} + \frac{\tau}{1-\gamma} + \frac{\eta K}{(1-\gamma)^2}$$

Putting all together, we have that for $\tau \leq \frac{K}{\sqrt{\tau}}$

$$\text{Regret}(K; \pi^*) \leq \frac{\tau \log |\mathcal{A}|}{\eta} + \frac{\tau + 1}{1-\gamma} + \frac{\eta K}{(1-\gamma)^2} + 12\beta K \sqrt{\frac{d}{\tau}} \log \left(\frac{2Kd}{\tau\delta} \right) + \frac{\sqrt{2\eta}K}{(1-\gamma)^2\tau}$$

□

D.3. Proof of Theorem 5

To improve the readability of the proof we restate Algorithm 3 hereafter.

Algorithm 5 ILARL (detailed version).

- 1: **Input:** Dataset size τ , Exploration parameter β , Step size η , Expert dataset $\mathcal{D}_{\pi_E} = \{\tau_i\}_{i=1}^{\tau_E}$.
- 2: Estimate for expert feature visitation $\widehat{\Phi^\top d^{\pi_E}} \triangleq \frac{(1-\gamma)}{\tau_E} \sum_{\tau \in \mathcal{D}_{\pi_E}} \sum_{s_h, a_h \in \tau} \gamma^h \phi(s_h, a_h)$.
- 3: Initialize π_0 as uniform distribution over \mathcal{A}
- 4: Initialize $V^1 = 0$
- 5: **for** $j = 1, \dots, \lfloor K/\tau \rfloor$ **do**
- 6: **// Collect on-policy data**
- 7: Denote the indices interval $T_j \triangleq [(j-1)\lfloor K/\tau \rfloor, j\lfloor K/\tau \rfloor]$.
- 8: Sample $\mathcal{D}^{(j)} = \{s^i, a^i, s'^i\}_{i \in T_j} \sim d^{\pi^{(j)}}$
- 9: Compute $\Lambda^{(j)} = \sum_{(s,a) \in \mathcal{D}^{(j)}} \phi(s,a)\phi(s,a)^\top + I$.
- 10: Compute $b^{(j)}(s,a) = \beta \|\phi(s,a)\|_{(\Lambda^{(j)})^{-1}}$.
- 11: **for** $k \in T_j$ **do**
- 12: **// Cost update**
- 13: Estimate features expectation vector $\widehat{\Phi^\top d^{\pi^k}}$ as $\tau^{-1} \sum_{s,a \in \mathcal{D}^{(j)}} \phi(s,a)$.
- 14: $\mathbf{w}^{k+1} = \Pi_{\mathcal{W}} \left[\mathbf{w}^k - \alpha (\widehat{\Phi^\top d^{\pi_E}} - \widehat{\Phi^\top d^{\pi^k}}) \right]$ with $\mathcal{W} = \{\mathbf{w} : \|\mathbf{w}\|_2 \leq 1\}$.
- 15: **// Optimistic Policy Evaluation**
- 16: $\mathbf{v}^k = (\Lambda^{(j)})^{-1} \sum_{(s,a,s') \in \mathcal{D}^{(j)}} \phi(s,a) V^k(s')$
- 17: $Q^{k+1} = [\Phi \mathbf{w}^k + \gamma \Phi \mathbf{v}^k - b^{(j)}]_{-(1-\gamma)^{-1}}^{(1-\gamma)^{-1}}$
- 18: $V^{k+1}(s) = \langle \pi^{(j)}(a|s), Q^{k+1}(s,a) \rangle$ (notice that $\pi^{(j)} = \pi^{k+1}$).
- 19: **end for**
- 20: **// Policy Improvement Step**
- 21: Compute average Q value $\bar{Q}^{(j)}(s,a) = \frac{1}{\tau} \sum_{k \in T_j} Q^k(s,a)$.
- 22: Update policy

$$\pi^{(j+1)}(a|s) \propto \exp \left(-\eta \sum_{i=1}^j \bar{Q}^{(i)}(s,a) \right)$$

23: **end for**

Proof. Consider the following decomposition

$$\sum_{k=1}^K \langle \mathbf{c}_{\text{true}}, d^{\pi_E} - d^{\pi^k} \rangle = \sum_{k=1}^K \langle w_{\text{true}} - w^k, \widehat{\Phi^\top d^{\pi_E}} - \Phi^\top d^{\pi^k} \rangle + \sum_{k=1}^K \langle c^k, d^{\pi_E} - d^{\pi^k} \rangle \quad (6)$$

$$+ \sum_{k=1}^K \langle w_{\text{true}} - w^k, \widehat{\Phi^\top d^{\pi_E}} - \widehat{\Phi^\top d^{\pi^k}} \rangle \quad (7)$$

For the first term we can use the following steps.

$$\sum_{k=1}^K \langle w_{\text{true}} - w^k, \widehat{\Phi^\top d^{\pi_E}} - \Phi^\top d^{\pi^k} \rangle \leq \sum_{k=1}^K \langle w_{\text{true}} - w^k, \widehat{\Phi^\top d^{\pi_E}} - \widehat{\Phi^\top d^{\pi^k}} \rangle + \sum_{k=1}^K \langle w_{\text{true}} - w^k, \widehat{\Phi^\top d^{\pi^k}} - \Phi^\top d^{\pi^k} \rangle$$

Now, using the regret bound for OMD (Orabona, 2023, Theorem 6.10) we can bound the first term in the decomposition above as

$$\begin{aligned} \sum_{k=1}^K \langle w_{\text{true}} - w^k, \widehat{\Phi^\top d^{\pi_E}} - \widehat{\Phi^\top d^{\pi^k}} \rangle &\leq \frac{\max_{w \in \mathcal{W}} \|w - w^1\|_2^2}{2\alpha} + \frac{\alpha}{2} \sum_{k=1}^K \left\| \widehat{\Phi^\top d^{\pi_E}} - \widehat{\Phi^\top d^{\pi^k}} \right\|_2^2 \\ &\leq \frac{\max_{c \in \mathcal{C}} \|c - c^1\|_2^2}{2\alpha} + \frac{\alpha}{2} \sum_{k=1}^K \left\| \widehat{\Phi^\top d^{\pi_E}} - \widehat{\Phi^\top d^{\pi^k}} \right\|_1^2 \\ &\leq \frac{1}{2\alpha} + 2\alpha K \end{aligned}$$

Then, for $\alpha = \frac{1}{2\sqrt{K}}$, then $\sum_{k=1}^K \langle w_{\text{true}} - w^k, \widehat{\Phi^\top d^{\pi_E}} - \widehat{\Phi^\top d^{\pi^k}} \rangle \leq 2\sqrt{K}$. For the estimation term,

$$\begin{aligned} \sum_{k=1}^K \langle w_{\text{true}} - w^k, \Phi^\top d^{\pi^k} - \widehat{\Phi^\top d^{\pi^k}} \rangle &\leq \sum_{k=1}^K \|w - w^k\|_1 \left\| \Phi^\top d^{\pi^k} - \widehat{\Phi^\top d^{\pi^k}} \right\|_\infty \\ &\leq \sqrt{d} \sum_{k=1}^K \|w - w^k\|_2 \left\| \Phi^\top d^{\pi^k} - \widehat{\Phi^\top d^{\pi^k}} \right\|_\infty \\ &\leq 2\sqrt{d} \sum_{k=1}^K \left\| \Phi^\top d^{\pi^k} - \widehat{\Phi^\top d^{\pi^k}} \right\|_\infty \\ &\leq 2\sqrt{d} \sum_{k=1}^K \sqrt{\frac{2 \log(2dK/\delta)}{\tau}} = 2K \sqrt{\frac{2d \log(2dK/\delta)}{\tau}}, \end{aligned}$$

where the last inequality holds with probability $1 - \delta$ thanks to Azuma-Hoeffding inequality. Therefore, using Theorem 4 to control the second term in Equation (7) and a union bound we obtain that with probability $1 - 3\delta$.

$$\begin{aligned} \sum_{k=1}^K \langle \mathbf{c}_{\text{true}}, d^{\pi_E} - d^{\pi^k} \rangle &\leq 2\sqrt{K} + 2K \sqrt{\frac{2d \log(2dK/\delta)}{\tau}} + \frac{\tau \log |\mathcal{A}|}{\eta} + \frac{\tau + 1}{1 - \gamma} + \frac{\eta K}{(1 - \gamma)^2} \\ &\quad + 12\beta K \sqrt{\frac{d}{\tau}} \log\left(\frac{2Kd}{\tau\delta}\right) + \frac{\sqrt{2\eta}K}{(1 - \gamma)^2\tau} + \sum_{k=1}^K \langle w_{\text{true}} - w^k, \widehat{\Phi^\top d^{\pi_E}} - \Phi^\top d^{\pi^k} \rangle \end{aligned}$$

and using that for the empirical expert an application of Lemma 8 gives that with probability $1 - \delta$.

$$\sum_{k=1}^K \langle w_{\text{true}} - w^k, \widehat{\Phi^\top d^{\pi_E}} - \Phi^\top d^{\pi^k} \rangle \leq 2K\sqrt{d} \left\| \Phi^\top d^{\pi_E} - \widehat{\Phi^\top d^{\pi_E}} \right\|_\infty \leq 2K \sqrt{\frac{2d \log(d/\delta)}{\tau_E}}$$

Therefore, selecting $\tau_E \geq \frac{8d \log(d/\delta)}{\epsilon_E^2}$ and using a last union bound gives that with probability $1 - 4\delta$

$$\begin{aligned} \left\langle \mathbf{c}_{\text{true}}, d^{\pi_E} - \frac{1}{K} \sum_{k=1}^K d^{\pi^k} \right\rangle &\leq \epsilon_E + \frac{2}{\sqrt{K}} + \sqrt{\frac{8d \log(2dK/\delta)}{\tau}} + \frac{\tau \log |\mathcal{A}|}{\eta K} + \frac{\tau + 1}{(1 - \gamma)K} + \frac{\eta}{(1 - \gamma)^2} \\ &\quad + 12\beta \sqrt{\frac{d}{\tau}} \log\left(\frac{2K}{\tau\delta}\right) + \frac{\sqrt{2\eta}}{(1 - \gamma)^2\tau} \end{aligned}$$

Using $\eta = \sqrt{\frac{\tau \log |\mathcal{A}| (1 - \gamma)^2}{K}}$,

$$\begin{aligned} \left\langle \mathbf{c}_{\text{true}}, d^{\pi_E} - \frac{1}{K} \sum_{k=1}^K d^{\pi^k} \right\rangle &\leq \epsilon_E + \frac{2}{\sqrt{K}} + \sqrt{\frac{8d \log(2dK/\delta)}{\tau}} + \frac{2}{(1 - \gamma)} \sqrt{\frac{\tau \log |\mathcal{A}|}{K}} + \frac{\tau + 1}{(1 - \gamma)K} \\ &\quad + 12\beta \sqrt{\frac{d}{\tau}} \log\left(\frac{2K}{\tau\delta}\right) + \frac{\sqrt{2}}{(1 - \gamma)^2\tau} \sqrt{\frac{\tau \log |\mathcal{A}| (1 - \gamma)^2}{K}} \end{aligned}$$

Neglecting lower order terms we obtain

$$\begin{aligned} \left\langle \mathbf{c}_{\text{true}}, d^{\pi_E} - \frac{1}{K} \sum_{k=1}^K d^{\pi^k} \right\rangle &\leq \epsilon_E + \mathcal{O}\left(\sqrt{\frac{8d \log(2dK/\delta)}{\tau}} + \frac{2}{(1 - \gamma)} \sqrt{\frac{\tau \log |\mathcal{A}|}{K}} + 12\beta \sqrt{\frac{d}{\tau}} \log\left(\frac{2K}{\tau\delta}\right)\right) \\ &\leq \epsilon_E + \mathcal{O}\left(\frac{2}{(1 - \gamma)} \sqrt{\frac{\tau \log |\mathcal{A}|}{K}} + 15\beta \sqrt{\frac{d}{\tau}} \log\left(\frac{2dK}{\tau\delta}\right)\right) \end{aligned}$$

Therefore, choosing $\tau = \mathcal{O}\left(\frac{\beta(1-\gamma)\sqrt{dK}\log(2dK/\delta)}{\sqrt{\log|\mathcal{A}|}}\right)$, gives

$$\left\langle \mathbf{c}_{\text{true}}, d^{\pi_E} - \frac{1}{K} \sum_{k=1}^K d^{\pi^k} \right\rangle \leq \epsilon_E + \tilde{\mathcal{O}}\left(\frac{\log^{1/4}|\mathcal{A}| d^{1/4} \sqrt{\beta}}{\sqrt{1-\gamma} K^{1/4}}\right)$$

Therefore, choosing $K \geq \tilde{\mathcal{O}}\left(\frac{\log|\mathcal{A}|d\beta^2\log^2(2dK/\delta)}{(1-\gamma)^2\epsilon^4}\right)$ which is attained by $K = \mathcal{O}\left(\frac{\log|\mathcal{A}|d\beta^2}{(1-\gamma)^2\epsilon^4} \log^2\left(\frac{d\beta\log|\mathcal{A}|}{\delta(1-\gamma)\epsilon}\right)\right)$ ensures $\left\langle \mathbf{c}_{\text{true}}, d^{\pi_E} - \frac{1}{K} \sum_{k=1}^K d^{\pi^k} \right\rangle \leq \epsilon + \epsilon_E$ with probability $1 - 4\delta$. \square

E. Technical Lemmas

Lemma 3. Assume that in Algorithm 1, we set $\beta = \mathcal{O}(dH \log(RdH\delta^{-1}))$ with $R = \tau^2\sqrt{d}H$. Then, with probability $1 - \delta$ it holds that

$$-2b_h^k(s, a) \leq Q_h^k(s, a) - \mathbf{c}_h^k(s, a) - PV_{h+1}^k(s, a) \leq 0 \quad \forall s, a \in \mathcal{S} \times \mathcal{A}, k \in [K], h \in [H]$$

Proof. We first see that Lemma 5 holds for every j . Then, thanks to union bound we have that with probability $1 - \delta$ it holds that for any state action pairs we have that

$$|\phi(s, a)^T \mathbf{v}_h^k - PV_{h+1}^k(s, a)| \leq \beta \|\phi(s, a)\|_{(\Lambda_h^{(j)})^{-1}} = b_h^k(s, a)$$

From this fact the conclusion follows immediately if no truncation happens. That is, if $Q_h^k = \mathbf{c}_h^k + \Phi \mathbf{v}_h^k - b_h^k$. Now, we consider the case where a lower truncation takes place, in this case, we have

$$Q_h^k = -H + h - 1 \leq \mathbf{c}_h^k + PV_{h+1}^k$$

If a truncation from above takes place, it holds that

$$Q_h^k \leq \mathbf{c}_h^k + \Phi \mathbf{v}_h^k - b_h^k \leq \mathbf{c}_h^k + PV_{h+1}^k + b_h^k - b_h^k = \mathbf{c}_h^k + PV_{h+1}^k$$

To show the lower bound in the lemma in case of a lower truncation, we have that

$$Q_h^k \geq \mathbf{c}_h^k + \Phi \mathbf{v}_h^k - b_h^k \geq \mathbf{c}_h^k + PV_{h+1}^k - b_h^k - b_h^k = \mathbf{c}_h^k + PV_{h+1}^k - 2b_h^k$$

finally, if the truncation from above is triggered, we have that

$$Q_h^k = H - h + 1 \geq \mathbf{c}_h^k + PV_{h+1}^k \geq \mathbf{c}_h^k + PV_{h+1}^k - 2b_h^k$$

\square

Lemma 4. For any $k \in [K]$, let the bonus b^k be defined as in Algorithm 2 with $\beta = \mathcal{O}(d(1-\gamma)^{-1} \log(Rd(1-\gamma)^{-1}\delta^{-1}))$ with $R = \tau^2\sqrt{d}(1-\gamma)^{-1}$. Then, with probability $1 - \delta$ it holds that

$$-2b^k(s, a) \leq Q^{k+1}(s, a) - \mathbf{c}^k(s, a) - \gamma PV^k(s, a) \leq 0 \quad \forall s, a \in \mathcal{S} \times \mathcal{A}, k \in [K]$$

Proof. We first see that Lemma 6 holds for every j . Then, thanks to union bound we have that with probability $1 - \delta$ it holds that for any state action pairs we have that

$$|\phi(s, a)^T \mathbf{v}^k - PV^k(s, a)| \leq \beta \|\phi(s, a)\|_{(\Lambda^{(j)})^{-1}} = b^k(s, a)$$

From this fact the conclusion follows immediately if no truncation happens. That is, if $Q^{k+1} = \mathbf{c}^k + \Phi \mathbf{v}^k - b^k$. Now, we consider the case where a upper truncation takes place, in this case, we have

$$Q^{k+1} = \frac{1}{1-\gamma} = 1 + \frac{\gamma}{1-\gamma} \geq \mathbf{c}^k + \gamma PV^k - 2b^k$$

While for the upper bound, we have that

$$Q^{k+1} \leq \mathbf{c}^k + \gamma \Phi \mathbf{v}^k - b^k \leq \mathbf{c}^k + \gamma PV^k + b^k - b^k = \mathbf{c}^k + \gamma PV^k$$

Now, we handle the case of a lower truncation in this case

$$Q^{k+1} \geq \mathbf{c}^k + \gamma \Phi \mathbf{v}^k - b^k \geq \mathbf{c}^k + \gamma PV^k - b^k - b^k = \mathbf{c}^k + \gamma PV^k - 2b^k$$

for the upper bound we have that

$$Q^{k+1} = -\frac{1}{1-\gamma} = -1 - \frac{\gamma}{1-\gamma} \leq \mathbf{c}^k + \gamma PV^k$$

□

Lemma 5. Let V_h^k be the sequence of value functions generated by Algorithm 1, fix a batch index j and let T_j denote the set of indices in the j^{th} batch. Then, it holds that for $\beta = \tilde{O}(dH)$, the estimator

$$\mathbf{v}_h^k = (\Lambda^{(j)})^{-1} \sum_{(s,a,s') \in \mathcal{D}_h^{(j)}} \phi(s,a) V_{h+1}^k(s')$$

satisfies

$$|\phi(s,a)^T \mathbf{v}_h^k - PV_{h+1}^k(s,a)| \leq \beta \|\phi(s,a)\|_{(\Lambda_h^{(j)})^{-1}} = b^k(s,a) \quad \forall k \in T_j, h \in [H], \forall s,a \in \mathcal{S} \times \mathcal{A}$$

with probability $1 - \delta\tau/K$.

Proof. The proof is analogous to the proof of Lemma 6 but invoking Theorem 7 with $B = H$ and applying a further union bound over the set $[H]$. Thus, the proof is skipped for brevity. □

Lemma 6. Let V^k be the sequence of value functions generated by Algorithm 2, fix a batch index j and let T_j denote the set of indices in the j^{th} batch. Then, it holds that for $\beta = \tilde{O}\left(\frac{d}{1-\gamma}\right)$, the estimator

$$\mathbf{v}^k = (\Lambda^{(j)})^{-1} \sum_{(s,a,s') \in \mathcal{D}^{(j)}} \phi(s,a) V^k(s')$$

satisfies

$$|\phi(s,a)^T \mathbf{v}^k - PV^k(s,a)| \leq \beta \|\phi(s,a)\|_{(\Lambda^{(j)})^{-1}} = b^k(s,a) \quad \forall k \in T_j, \forall s,a \in \mathcal{S} \times \mathcal{A}$$

with probability $1 - \delta\tau/K$.

Proof. With standard manipulation one can prove that

$$PV^k = \Phi MV^k = \Phi (\Lambda^{(j)})^{-1} MV^k + \Phi (\Lambda^{(j)})^{-1} \sum_{s,a,s' \in \mathcal{D}^{(j)}} \phi(s,a) PV^k(s,a)$$

and by definition

$$\Phi \mathbf{v}^k = \Phi (\Lambda^{(j)})^{-1} \sum_{s,a,s' \in \mathcal{D}^{(j)}} \phi(s,a) V^k(s')$$

Therefore,

$$PV^k - \Phi \mathbf{v}^k = \Phi (\Lambda^{(j)})^{-1} MV^k + \Phi (\Lambda^{(j)})^{-1} \sum_{s,a,s' \in \mathcal{D}^{(j)}} \phi(s,a) (PV^k(s,a) - V^k(s'))$$

Then, for any state action pair (s,a) , we have

$$|PV^k(s,a) - \Phi \mathbf{v}^k(s,a)| = \phi(s,a)^T (\Lambda^{(j)})^{-1} MV^k + \phi(s,a)^T (\Lambda^{(j)})^{-1} \sum_{s,a,s' \in \mathcal{D}^{(j)}} \phi(s,a) (PV^k(s,a) - V^k(s'))$$

by applying Holder's inequality,

$$\begin{aligned}
 |PV^k(s, a) - \Phi \mathbf{v}^k(s, a)| &\leq \|\phi(s, a)\|_{(\Lambda^{(j)})^{-1}} \left\| (\Lambda^{(j)})^{-1} MV^k \right\|_{(\Lambda^{(j)})} \\
 &\quad + \|\phi(s, a)\|_{(\Lambda^{(j)})^{-1}} \left\| (\Lambda^{(j)})^{-1} \sum_{s, a, s' \in \mathcal{D}^{(j)}} \phi(s, a) (PV^k(s, a) - V^k(s')) \right\|_{(\Lambda^{(j)})} \\
 &= \|\phi(s, a)\|_{(\Lambda^{(j)})^{-1}} \|MV^k\|_{(\Lambda^{(j)})^{-1}} \\
 &\quad + \|\phi(s, a)\|_{(\Lambda^{(j)})^{-1}} \left\| \sum_{s, a, s' \in \mathcal{D}^{(j)}} \phi(s, a) (PV^k(s, a) - V^k(s')) \right\|_{(\Lambda^{(j)})^{-1}}
 \end{aligned}$$

where in the equality we used that for a symmetric matrix A we have that $\|Ax\|_{A^{-1}} = \|x\|_A$. Then, we can use that $\|MV^k\|_{(\Lambda^{(j)})^{-1}} \leq \frac{1}{1-\gamma}$ to obtain

$$|PV^k(s, a) - \Phi \mathbf{v}^k(s, a)| \leq \|\phi(s, a)\|_{(\Lambda^{(j)})^{-1}} \left(\frac{1}{1-\gamma} + \left\| \sum_{s, a, s' \in \mathcal{D}^{(j)}} \phi(s, a) (PV^k(s, a) - V^k(s')) \right\|_{(\Lambda^{(j)})^{-1}} \right)$$

To handle the second term in brackets we use that $V^k \in \mathcal{V}^{\pi^{(j)}}$ defined as

$$\mathcal{V}^{\pi^{(j)}} = \left\{ \left\langle \pi^{(j)}(\cdot|s), Q(s, \cdot) \right\rangle \mid Q(s, a) \in \mathcal{Q}(\beta, \Lambda, \mathbf{w}, \mathbf{v}) \right\}$$

where $\mathcal{Q}(\beta, \Lambda, \mathbf{w}, \mathbf{v})$ is defined as in Theorem 8. Denote as \mathcal{N}_ϵ^j the $\|\cdot\|_\infty$ -covering number of the class $\mathcal{V}^{\pi^{(j)}}$ and notice that $\mathcal{V}^{\pi^{(j)}}$ and $\mathcal{D}^{(j)}$ are conditionally independent given $\pi^{(j)}$.

Under this setting we can use Theorem 7 with $B = (1-\gamma)^{-1}$ to obtain that with probability $1 - \delta\tau/K$

$$\begin{aligned}
 |PV^k(s, a) - \Phi \mathbf{v}^k(s, a)| &\leq \|\phi(s, a)\|_{(\Lambda^{(j)})^{-1}} \left(\frac{1}{1-\gamma} + \sqrt{\frac{2d}{(1-\gamma)^2} \log \left(\frac{K(\tau+1)}{\delta\tau} \right)} + \frac{4}{(1-\gamma)^2} \log \mathcal{N}_\epsilon^j + 8\tau^2\epsilon^2 \right) \\
 &\leq \frac{\|\phi(s, a)\|_{(\Lambda^{(j)})^{-1}}}{1-\gamma} \left(1 + \sqrt{2d \log \left(\frac{K(\tau+1)}{\delta\tau} \right)} + 2\sqrt{\log \mathcal{N}_\epsilon^j} + 2\sqrt{2}\tau\epsilon \right).
 \end{aligned}$$

Then, we can conclude that the covering number is upper bounded by Theorem 8 with $L = \frac{\tau}{1-\gamma}$ since

$$\|\mathbf{v}^k\| \leq \frac{1}{(1-\gamma)} \left\| (\Lambda^{(j)})^{(-1)} \right\| \left\| \sum_{s, a, s' \in \mathcal{D}^{(j)}} \phi(s, a) \right\| \leq \frac{\tau}{(1-\gamma)}$$

we obtain that

$$\log \mathcal{N}_\epsilon^j \leq d \log \left(1 + \frac{4}{\epsilon} \sqrt{1 + \frac{\gamma^2 \tau^2}{(1-\gamma)^2}} \right) + d^2 \log(1 + 8\sqrt{d}\beta^2\epsilon^{-2})$$

Therefore,

$$\begin{aligned}
 |PV^k(s, a) - \Phi \mathbf{v}^k(s, a)| &\leq \frac{\|\phi(s, a)\|_{(\Lambda^{(j)})^{-1}}}{1-\gamma} \left[1 + \sqrt{2d \log \left(\frac{K(\tau+1)}{\delta\tau} \right)} + \sqrt{2d \log \left(1 + \frac{4}{\epsilon} \sqrt{1 + \frac{\gamma^2 \tau^2}{(1-\gamma)^2}} \right)} \right. \\
 &\quad \left. + \sqrt{2}d\sqrt{\log(1 + 8\sqrt{d}\beta^2\epsilon^{-2})} + 2\sqrt{2}\tau\epsilon \right]
 \end{aligned}$$

At this point, using $\epsilon = \tau^{-1}$, we obtain

$$|PV^k(s, a) - \Phi \mathbf{v}^k(s, a)| \leq \frac{\|\phi(s, a)\|_{(\Lambda^{(j)})^{-1}}}{1 - \gamma} \left[1 + \sqrt{2d \log \left(\frac{K(\tau + 1)}{\delta\tau} \right)} + \sqrt{2d \log \left(1 + 4\tau \sqrt{1 + \frac{\gamma^2 \tau^2}{(1 - \gamma)^2}} \right)} + \sqrt{2d} \sqrt{\log(1 + 8\sqrt{d}\beta^2\tau^2)} + 2\sqrt{2} \right]$$

To simplify the above expression, we notice that there exists a constant c such that

$$|PV^k(s, a) - \Phi \mathbf{v}^k(s, a)| \leq c \frac{\|\phi(s, a)\|_{(\Lambda^{(j)})^{-1}}}{1 - \gamma} d \sqrt{\log \left(\frac{\tau K \beta \sqrt{d}}{\delta(1 - \gamma)} \right)}$$

Now, using (Sherman et al., 2023b, Lemma D.2), we have that $\beta = \mathcal{O}(d(1 - \gamma)^{-1} \log [Rd(1 - \gamma)^{-1}])$ with $R = \tau K \sqrt{d} \delta^{-1} (1 - \gamma)^{-1}$ ensures

$$\beta \geq c \frac{d}{(1 - \gamma)} \log \left(\frac{\tau^2 \beta \sqrt{d}}{\delta(1 - \gamma)} \right),$$

and therefore

$$|PV^k(s, a) - \Phi \mathbf{v}^k(s, a)| \leq \beta \|\phi(s, a)\|_{(\Lambda^{(j)})^{-1}}.$$

□

Theorem 7. For a fixed policy π consider a function class \mathcal{V}^π and a state action pair dataset \mathcal{D} collected with a fixed policy π such that \mathcal{D} and \mathcal{V}^π are conditionally independent given π . Then, for any $f \in \mathcal{V}^\pi$ such that $\|f\|_\infty \leq B$, it holds with probability $1 - \delta\tau/K$

$$\left\| \sum_{s, a, s' \in \mathcal{D}} \phi(s, a) (Pf(s, a) - f(s')) \right\|_{(\Lambda^{(j)})^{-1}}^2 \leq 2dB^2 \log \left(\frac{K(\tau + 1)}{\delta\tau} \right) + 4B^2 \log \mathcal{N}_\epsilon + 8\tau^2 \epsilon^2$$

where N_ϵ is the $(\epsilon, \|\cdot\|_\infty)$ -covering number of the class \mathcal{V}^π .

Proof. Consider the decomposition in (Jin et al., 2019, Lemma D.4). In particular, let $\mathcal{C}_\epsilon(\mathcal{V}^\pi)$ denote the $(\epsilon, \|\cdot\|_\infty)$ -covering set of \mathcal{V}^π and pick $\tilde{f} \in \mathcal{C}_\epsilon(\mathcal{V}^\pi)$ such that $\|f - \tilde{f}\|_\infty \leq \epsilon$. The existence of \tilde{f} is guaranteed by the properties of covering sets. Then, we have

$$\begin{aligned} \left\| \sum_{s, a, s' \in \mathcal{D}} \phi(s, a) (Pf(s, a) - f(s')) \right\|_{(\Lambda^{(j)})^{-1}}^2 &\leq 2 \left\| \sum_{s, a, s' \in \mathcal{D}} \phi(s, a) (Pf(s, a) - \tilde{f}(s')) \right\|_{(\Lambda^{(j)})^{-1}}^2 \\ &\quad + 2 \left\| \sum_{s, a, s' \in \mathcal{D}} \phi(s, a) (P(f - \tilde{f})(s, a) - (f - \tilde{f})(s')) \right\|_{(\Lambda^{(j)})^{-1}}^2 \end{aligned}$$

The second term can be bounded by $8\tau^2 \epsilon^2$ as in (Jin et al., 2019) so now we focus on the first term via a uniform bound over the set $\mathcal{C}_\epsilon(\mathcal{V}^\pi)$. We need to index the dataset \mathcal{D} , i.e. $\mathcal{D} = \{(s^\ell, a^\ell)\}_{\ell=1}^{|\mathcal{D}|}$ and consider the filtration $\mathcal{F}_j = \{(s^\ell, a^\ell)\}_{\ell=1}^j$. Since the features mapping is deterministic, $\phi(s^\ell, a^\ell)$ is \mathcal{F}_ℓ -measurable. Then, notice that by assumption \mathcal{D} and \mathcal{V}^π are conditionally independent given π . Therefore, we also have that \mathcal{D} and $\mathcal{C}_\epsilon(\mathcal{V}^\pi)$ are conditionally independent given π . So for any function $\tilde{f} \in \mathcal{C}_\epsilon(\mathcal{V}^\pi)$ we have that $\mathbb{E}[\tilde{f}(s^{\ell+1}) | \mathcal{F}_\ell] = P\tilde{f}(s^\ell, a^\ell)$. Finally, from the assumption $\|\tilde{f}\|_\infty \leq B$ we have

that \bar{f} is B^2 -subgaussian. Therefore, all the conditions of (Jin et al., 2019, Theorem D.3) are met and via a union bound over the covering set allows to conclude that with probability $1 - \delta\tau/K$

$$2 \left\| \sum_{s,a,s' \in \mathcal{D}} \phi(s,a)(P\bar{f}(s,a) - \bar{f}(s')) \right\|_{(\Lambda^{(j)})^{-1}}^2 \leq 2dB^2 \log \left(\frac{K(\tau+1)}{\delta\tau} \right) + 4B^2 \log \mathcal{N}_\epsilon \quad \forall \bar{f} \in \mathcal{C}_\epsilon(\mathcal{V}),$$

and since $\tilde{f} \in \mathcal{C}_\epsilon(\mathcal{V}^\pi)$,

$$2 \left\| \sum_{s,a,s' \in \mathcal{D}} \phi(s,a)(P\tilde{f}(s,a) - \tilde{f}(s')) \right\|_{(\Lambda^{(j)})^{-1}}^2 \leq 2dB^2 \log \left(\frac{K(\tau+1)}{\delta\tau} \right) + 4B^2 \log \mathcal{N}_\epsilon.$$

□

Theorem 8. Let us consider the function class \mathcal{Q} defined as follows

$$\begin{aligned} \mathcal{Q}(\beta, \Lambda, \mathbf{w}, \mathbf{v}) &= \{Q(s, a; \beta, \Lambda, \mathbf{w}, \mathbf{v}) \mid \beta \in \mathbb{R}, \lambda_{\min}(\Lambda) \geq 1, \|\mathbf{w}\| \leq 1, \|\mathbf{v}\| \leq L\} \\ \text{where } Q(s, a; \beta, \Lambda, \mathbf{w}, \mathbf{v}) &= [\phi(s, a)^\top (\mathbf{w} + \gamma \mathbf{v}) + \beta \|\phi(s, a)\|_{\Lambda^{-1}}]_0^{(1-\gamma)^{-1}} \end{aligned}$$

and the classes

$$\mathcal{V}^\pi = \{\langle \pi(\cdot|s), Q(s, \cdot) \rangle \mid Q(s, a) \in \mathcal{Q}(\beta, \Lambda, \mathbf{w}, \mathbf{v})\}$$

for any $\pi : \mathcal{S} \rightarrow \Delta_{\mathcal{A}}$.

Then, it holds that for any $\pi : \mathcal{S} \rightarrow \Delta_{\mathcal{A}}$

$$\mathcal{N}_\epsilon(\mathcal{V}^\pi) \leq \mathcal{N}_\epsilon(\mathcal{Q}) = (1 + 4\sqrt{1 + \gamma^2 L^2}/\epsilon)^d (1 + 8\sqrt{d}\beta^2 \epsilon^{-2})^{d^2}$$

Proof. Let us remove clipping that can only decreasing the covering number of the function class and let us consider the matrix $A = \beta^2 \Lambda^{-1}$, then we can rewrite the function class of interest as parameterized only by A rather than β and Λ separately. In addition, let us consider a vector $\mathbf{z} = \mathbf{w} + \gamma \mathbf{v}$

$$\mathcal{Q}(A, \mathbf{z}) = \left\{ Q(s, a; A, \mathbf{w}, \mathbf{v}) \mid \lambda_{\min}(\Lambda) \geq \beta^2, \|\mathbf{z}\|^2 \leq 2 + 2\gamma^2 L^2 \right\}$$

with

$$Q(s, a; A, \mathbf{z}) = \phi(s, a)^\top \mathbf{z} + \|\phi(s, a)\|_A$$

Then, we have that

$$\begin{aligned} |Q(s, a; A_1, \mathbf{z}_1) - Q(s, a; A_2, \mathbf{z}_2)| &\leq \|\phi(s, a)\| \|\mathbf{z}_1 - \mathbf{z}_2\| + \left| \sqrt{\phi(s, a)^\top A_1 \phi(s, a)} - \sqrt{\phi(s, a)^\top A_2 \phi(s, a)} \right| \\ &\leq \|\phi(s, a)\| \|\mathbf{z}_1 - \mathbf{z}_2\| + \sqrt{|\phi(s, a)^\top (A_1 - A_2) \phi(s, a)|} \\ &\leq \|\mathbf{z}_1 - \mathbf{z}_2\| + \sqrt{\sup_{\phi: \|\phi\| \leq 1} |\phi(s, a)^\top (A_1 - A_2) \phi(s, a)|} \\ &\leq \|\mathbf{z}_1 - \mathbf{z}_2\| + \sqrt{\|A_1 - A_2\|} \\ &\leq \|\mathbf{z}_1 - \mathbf{z}_2\| + \sqrt{\|A_1 - A_2\|_F} \end{aligned}$$

where $\|A_1 - A_2\|$ is the spectral norm of the matrix $A_1 - A_2$ and $\|A_1 - A_2\|_F$ is the Frobenius norm. We also used the inequality $|\sqrt{x} - \sqrt{y}| \leq \sqrt{|x - y|}$ that holds for any $x, y \geq 0$. At this point we can constructing an ϵ -covering set for $\mathcal{Q}(A, \mathbf{z})$ as product of the $\epsilon^2/4$ covering set for the set $\mathcal{Y} = \left\{ A \in \mathbb{R}^{d \times d} \mid \|A\|_F \leq \sqrt{d}\beta^{-2} \right\}$ which has cardinality

$\mathcal{N}_\epsilon(\mathcal{Y}) = (1 + 8\sqrt{d}\beta^2\epsilon^{-2})^{d^2}$ while the covering for the set $\mathcal{Z} = \{\mathbf{z} \in \mathbb{R}^d : \|\mathbf{z}\|^2 \leq 1 + \gamma^2 L^2\}$ satisfies $\mathcal{N}_\epsilon(\mathcal{Z}) = (1 + 4\sqrt{1 + \gamma^2 L^2}/\epsilon)^d$. Hence, taking the product, we have that

$$\mathcal{N}_\epsilon(\mathcal{Q}) = (1 + 4\sqrt{1 + \gamma^2 L^2}/\epsilon)^d (1 + 8\sqrt{d}\beta^2\epsilon^{-2})^{d^2}$$

.

At this point, let us consider the set

$$\mathcal{V}^\pi = \{\langle \pi(\cdot|s), Q(s, \cdot) \rangle \mid Q(s, a) \in \mathcal{Q}(A, \mathbf{z})\}$$

Since the policy π is fixed and averaging is a non expansive operation, we have that $\mathcal{N}_\epsilon(\mathcal{V}^\pi) \leq \mathcal{N}_\epsilon(\mathcal{Q})$.

However, for the set

$$\mathcal{V} = \{\langle \pi(\cdot|s), Q(s, \cdot) \rangle \mid \pi \in \Pi, Q(s, a) \in \mathcal{Q}(\beta, \Lambda, \mathbf{w}, \mathbf{v})\}$$

the averaging is not wrt to a fixed distribution therefore we would need to proceed as follow

$$\begin{aligned} |\langle \pi_1(\cdot|s), Q_1(s, \cdot) \rangle - \langle \pi_2(\cdot|s), Q_2(s, \cdot) \rangle| &\leq \frac{1}{(1-\gamma)} \|\pi_1(\cdot|s) - \pi_2(\cdot|s)\|_1 + 2 \|Q_1(s, \cdot) - Q_2(s, \cdot)\|_\infty \\ &\leq \frac{1}{(1-\gamma)} \|\pi_1(\cdot|s) - \pi_2(\cdot|s)\|_1 + 2 \|Q_1(s, \cdot) - Q_2(s, \cdot)\|_\infty \\ &\leq \frac{1}{(1-\gamma)} \max_{s \in \mathcal{S}} \|\pi_1(\cdot|s) - \pi_2(\cdot|s)\|_1 + 2 \|\mathbf{z}_1 - \mathbf{z}_2\| + \sqrt{\|A_1 - A_2\|_F} \\ &\leq \frac{1}{(1-\gamma)} \|\pi_1(\cdot|s) - \pi_2(\cdot|s)\|_{\infty,1} + 2 \|\mathbf{z}_1 - \mathbf{z}_2\| + \sqrt{\|A_1 - A_2\|_F} \end{aligned}$$

Therefore, we can conclude that $\mathcal{N}_\epsilon(\mathcal{V}) = \mathcal{N}_\epsilon(\Pi, \|\cdot\|_{\infty,1}) \mathcal{N}_\epsilon(\mathcal{Q})$. \square

Next, we prove the Lemma that we use to state Theorems 1 and 2 using $\delta = \epsilon$.

Lemma 7. High probability to expectation conversion for a bounded random variable *Let us consider a random variable X such that $-X_{\max} \leq X \leq X_{\max}$ almost surely and that $\mathbb{P}[X \geq \mu] \leq \delta$, then it holds that*

$$\mathbb{E}[X] \leq \mu + \delta(X_{\max} - \mu)$$

Proof.

$$\mathbb{E}[X] = (1 - \delta)\mathbb{E}[X|X \leq \mu] + \delta\mathbb{E}[X|X \geq \mu] \leq (1 - \delta)\mu + \delta X_{\max}$$

\square

Lemma 8. Expert concentration (Syed & Schapire, 2007, Theorem 1) *Let $\mathcal{D}_{\pi_E} \triangleq \{(s_0^\ell, a_0^\ell, s_1^\ell, a_1^\ell, \dots, s_H^\ell, a_H^\ell)\}_{\ell=1}^{n_E}$ be a finite set of i.i.d. truncated sample trajectories collected with an expert policy π_E . We consider the empirical expert feature expectation vector $\widehat{\Phi}^\top d^{\pi_E}$ by taking sample averages, i.e.,*

$$\widehat{\Phi}^\top d^{\pi_E} \triangleq (1 - \gamma) \frac{1}{n_E} \sum_{t=0}^H \sum_{\ell=1}^{n_E} \gamma^t \phi_i(s_t^\ell, a_t^\ell), \quad \forall i \in [d].$$

Suppose the trajectory length is $H \geq \frac{1}{1-\gamma} \log(\frac{1}{\epsilon})$, and the number of of expert trajectories is $n_E \geq \frac{2 \log(\frac{2d}{\delta})}{\epsilon^2}$. Then, with probability at least $1 - \delta$, it holds that $\left\| \widehat{\Phi}^\top d^{\pi_E} - \widehat{\Phi}^\top d^{\pi_E} \right\|_\infty \leq \epsilon$.

Lemma 9. Bound on OMD *Let us consider a maximum number of allowed iterations of OMD K , and update frequency every τ round in an online learning problem with $|\mathcal{A}|$ actions. The total number of updates is hence equal to $J = \lfloor T/\tau \rfloor$ and*

for each $j \in 0, 1, \dots, J-1$ we denote the set of indices where the learner decision is kept fixed as $T_j = [\tau j, \dots, \tau(j+1))$. Then, consider a sequence of losses in $[-B, B]^{|A|}$ denoted $\{\ell^k\}_{k=1}^K$ and the updates

$$\begin{aligned}\bar{\ell}^{(j)} &= \tau^{-1} \sum_{k \in T_j} \ell^k \\ x^{(j)} &\propto \exp\left(\eta \sum_{i=0}^{j-1} \bar{\ell}^{(i)}\right)\end{aligned}$$

Let us consider $x^k = x^{(j)}$ for all $k \in T_j$, then it holds that the regret is bounded as

$$\sum_{k=1}^K \langle \ell^k, x^k - x^* \rangle \leq \frac{\tau \log |\mathcal{A}|}{\eta} + \tau B + \eta K B^2$$

for any $x^* \in \Delta_{\mathcal{A}}$.

Proof. The proof is basically equivalent to (Sherman et al., 2023b, Lemma F.5). Let us first consider the case $K = J\tau$. Then, we have that, by the standard bound on the regret of OMD (see for example (Orabona, 2023))

$$\sum_{j=1}^J \langle \bar{\ell}^j, x^{(j)} - x^* \rangle \leq \frac{\log |\mathcal{A}|}{\eta} + \eta \sum_{j=1}^J \sum_{a \in \mathcal{A}} x^{(j)}(a) (\bar{\ell}^{(j)}(a))^2 \quad (8)$$

Then, we can notice that by Jensen's inequality $(\bar{\ell}^{(j)}(a))^2 \leq \tau^{-1} \sum_{k \in T_j} (\ell^k(a))^2$. Moreover,

$$\sum_{j=1}^J \langle \bar{\ell}^j, x^{(j)} - x^* \rangle = \sum_{j=1}^J \left\langle \tau^{-1} \sum_{k \in T_j} \ell^k, x^{(j)} - x^* \right\rangle = \tau^{-1} \sum_{k=1}^K \langle \ell^k, x^k - x^* \rangle.$$

Therefore, plugging in Equation (8) we obtain

$$\sum_{k=1}^K \langle \ell^k, x^k - x^* \rangle \leq \frac{\tau \log |\mathcal{A}|}{\eta} + \eta \sum_{k=1}^K \sum_{a \in \mathcal{A}} x^k(a) (\ell^k(a))^2$$

and using that the losses are bounded we have $(\ell^k(a))^2 \leq B^2$. We get,

$$\sum_{k=1}^K \langle \ell^k, x^k - x^* \rangle \leq \frac{\tau \log |\mathcal{A}|}{\eta} + \eta K B^2.$$

Finally, if $K \geq \tau J$ we can upper bound the regret of the first $\lfloor K/\tau \rfloor$ rounds as before. Then, we have at most τ additional rounds between $K' = \lfloor K/\tau \rfloor \tau$ and K where the regret increases at most by B at each round. Hence, we have

$$\sum_{k=1}^K \langle \ell^k, x^k - x^* \rangle = \sum_{k=1}^{K'} \langle \ell^k, x^k - x^* \rangle + \tau B \leq \frac{\tau \log |\mathcal{A}|}{\eta} + \eta K B^2 + \tau B.$$

□

Lemma 10. For any $k \in [K]$, $h \in [H]$, let us denote as b_h^k the bonus function as constructed in Step 13 of Algorithm 1, i.e. $b_h^k(s, a) = \beta \|\phi(s, a)\|_{(\Lambda_h^k)^{-1}}$ where the matrix Λ_h^k is as in Step 11 of Algorithm 1. Then, it holds that with probability $1 - \delta$,

$$\mathbb{E}_{s, a \sim d_h^{\pi^k}} [b_h^k(s, a)] \leq \frac{7\sqrt{d}\beta \log(2\tau/\delta)}{\sqrt{\tau}}$$

Proof. We start by rewriting the expectation as

$$\mathbb{E}_{s,a \sim d_h^{\pi^k}} [b_h^k(s, a)] = \frac{1}{\tau} \mathbb{E}_{\mathcal{D}_h^k \sim d_h^{\pi^k}} \left[\sum_{s,a \in \mathcal{D}_h^k} b_h^k(s, a) \right]$$

Then, by (Cohen et al., 2020, Lemma D.4), we have that with probability $1 - \delta$, the following chain of inequality holds.

$$\begin{aligned} \frac{1}{\tau} \mathbb{E}_{\mathcal{D}_h^k \sim d_h^{\pi^k}} \left[\sum_{s,a \in \mathcal{D}_h^k} b_h^k(s, a) \right] &\leq \frac{2}{\tau} \sum_{s,a \in \mathcal{D}_h^k} b_h^k(s, a) + 4\beta \log(2\tau/\delta) \\ &= \frac{2\beta}{\tau} \sum_{s,a \in \mathcal{D}_h^k} \|\phi(s, a)\|_{(\Lambda_h^k)^{-1}} + 4\beta \log(2\tau/\delta) \\ &\leq \frac{2\beta}{\tau} \sqrt{\tau \sum_{s,a \in \mathcal{D}_h^k} \phi(s, a)^T (\Lambda_h^k)^{-1} \phi(s, a)} + 4\beta \log(2\tau/\delta) \\ &\leq \frac{2\beta}{\tau} \sqrt{2d\tau \log(1 + \tau/d)} + 4\beta \log(2\tau/\delta) \end{aligned}$$

where the last inequality follows from the Elliptical Potential Lemma (Abbasi-Yadkori et al., 2011). Overall, further upper bounding the last display, we get

$$\mathbb{E}_{s,a \sim d_h^{\pi^k}} [b_h^k(s, a)] \leq \frac{7\beta\sqrt{d} \log(2\tau/\delta)}{\sqrt{\tau}}$$

□

F. Omitted proofs for Best Response Imitation Learning

F.1. Proof of Theorem 6

Proof.

$$\begin{aligned} \text{Regret}(K, \pi^*) &= \sum_{k=1}^K V_1^{\pi^k, k}(s_1) - V_1^{\pi^*, k} \\ &= \sum_{k=1}^K V_1^{\pi^k, k}(s_1) - V_1^{k-1}(s_1) + V_1^{k-1}(s_1) - V_1^{\pi^*, k} \\ &= \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{s,a \sim d_h^{\pi^k}} [\mathbf{c}_h^k(s, a) + P_h V_{h+1}^{k-1}(s, a) - Q_h^{k-1}(s, a)] \\ &\quad + \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{s,a \sim d_h^{\pi^*}} [Q_h^{k-1}(s, a) - \mathbf{c}_h^k(s, a) - P_h V_{h+1}^{k-1}(s, a)] \\ &\quad + \sum_{h=1}^H \mathbb{E}_{s \sim d_h^{\pi^*}} \left[\sum_{k=1}^K \langle \pi_h^k(\cdot|s) - \pi_h^*(\cdot|s), Q_h^{k-1}(s, a) \rangle \right] \\ &\leq \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{s,a \sim d_h^{\pi^k}} [\mathbf{c}_h^k(s, a) + P_h V_{h+1}^{k-1}(s, a) - Q_h^{k-1}(s, a)] \\ &\quad + \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{s,a \sim d_h^{\pi^*}} [Q_h^{k-1}(s, a) - \mathbf{c}_h^k(s, a) - P_h V_{h+1}^{k-1}(s, a)] \end{aligned}$$

where the last inequality is due to the use of the best response (greedy policy) in Step 17 of Algorithm 4. At this point we can prove the optimistic properties of the estimator (that follows combining Lemmas 4 and 11), i.e. for any $h = H, \dots, 1$, it holds that

$$\mathbf{c}_h^k(s, a) + P_h V_{h+1}^{k-1}(s, a) - 2b_h^{k-1}(s, a) \leq Q_h^{k-1}(s, a) \leq \mathbf{c}_h^k(s, a) + P_h V_{h+1}^{k-1}(s, a) \quad \forall s, a \in \mathcal{S} \times \mathcal{A} \quad w.p. \quad 1 - \delta.$$

Thus, it holds with probability $1 - \delta$ that

$$\text{Regret}(K, \pi^*) \leq 2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{s, a \sim d_h^{\pi^k}} [b_h^{k-1}(s, a)]$$

and then, using Cauchy-Schwartz and the elliptical potential lemma (see Lemma 12), we obtain that with probability $1 - 2\delta$,

$$\text{Regret}(K, \pi^*) \leq \mathcal{O} \left(H^2 d^{3/2} \sqrt{K \log(K\delta^{-1})} \right)$$

Then, we apply this result in the imitation learning setting. We start with our usual decomposition

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \langle \mathbf{c}_{\text{true}, h}, d_h^{\pi^k} - d_h^{\widehat{\pi}^E} \rangle &= \sum_{k=1}^K \sum_{h=1}^H \langle \mathbf{c}_h^k, d_h^{\pi^k} - d_h^{\widehat{\pi}^E} \rangle + \sum_{k=1}^K \sum_{h=1}^H \langle \mathbf{c}_{\text{true}, h} - \mathbf{c}_h^k, d_h^{\pi^k} - d_h^{\widehat{\pi}^E} \rangle \\ &= \sum_{k=1}^K \sum_{h=1}^H \langle \mathbf{c}_h^k, d_h^{\pi^k} - d_h^{\widehat{\pi}^E} \rangle + \sum_{k=1}^K \sum_{h=1}^H \langle \mathbf{c}_{\text{true}, h} - \mathbf{c}_h^k, \mathbb{1} [s_k^h, a_k^h] - d_h^{\widehat{\pi}^E} \rangle \\ &\quad - \sum_{h=1}^H \sum_{k=1}^K \langle \mathbf{c}_{\text{true}, h} - \mathbf{c}_h^k, \mathbb{1} [s_k^h, a_k^h] - d_h^{\pi^k} \rangle \end{aligned}$$

Then, notice that $z_h^k = - \langle \mathbf{c}_{\text{true}, h} - \mathbf{c}_h^k, \mathbb{1} [s_k^h, a_k^h] - d_h^{\pi^k} \rangle$ is a martingale difference sequence adapted to the filtration $\mathcal{F}_h^k = \{(\tau^k, \mathbf{c}_h^k)\}$ almost surely bounded by 4. Therefore, by Azuma Hoeffding inequality, we obtain $\sum_{h=1}^H \sum_{k=1}^K z_h^k \leq H \sqrt{8K \log \delta^{-1}}$. Therefore, via a union bound, we have that with probability $1 - 3\delta$, it holds

$$\begin{aligned} \sum_{k=1}^K \sum_{h=1}^H \langle \mathbf{c}_{\text{true}, h}, d_h^{\pi^k} - d_h^{\widehat{\pi}^E} \rangle &\leq \tilde{\mathcal{O}} \left(H^2 d^{3/2} \sqrt{K \log(K\delta^{-1})} \right) + \frac{H}{\alpha} + 2\alpha KH + H \sqrt{8K \log \delta^{-1}} \\ &= \tilde{\mathcal{O}} \left(H^2 d^{3/2} \sqrt{K \log(K\delta^{-1})} \right) + 4H\sqrt{K} + H \sqrt{8K \log \delta^{-1}} \end{aligned}$$

where last step follows from choosing $\alpha = \frac{1}{\sqrt{2K}}$. At this point, the conclusion holds plugging in the value for K in the statement of the main theorem which is $K = \mathcal{O} \left(\frac{H^4 d^3 \log(dH/(\epsilon\delta))}{\epsilon^2} \right)$. Finally, we need to control the error in the estimation of the expert occupancy measure that can be done as in the proof for Algorithm 1.

$$\begin{aligned} V_1^{\pi^E}(s_1; \mathbf{c}_{\text{true}}) - V_1^{\widehat{\pi}^E}(s_1; \mathbf{c}_{\text{true}}) &= \sum_{h=1}^H \langle \Phi^\top d_h^{\pi^E} - \Phi^\top d_h^{\widehat{\pi}^E}, w_{\text{true}, h} \rangle \\ &\leq H \sqrt{d} \max_{h \in [H]} \left\| \Phi^\top d_h^{\pi^E} - \Phi^\top d_h^{\widehat{\pi}^E} \right\|_\infty \\ &\leq H \sqrt{\frac{2d \log(2d/\delta)}{\tau_E}} \end{aligned}$$

where the last inequality holds with probability $1 - \delta$. Therefore, the choice of τ_E in the theorem statement ensures that $V_1^{\pi^E}(s_1; \mathbf{c}_{\text{true}}) - V_1^{\widehat{\pi}^E}(s_1; \mathbf{c}_{\text{true}}) \leq \epsilon_E$. \square

Lemma 11. For $\beta = \mathcal{O} \left(dH \log \left(\frac{dT}{\delta} \right) \right)$, the estimator used in Algorithm 4

$$\mathbf{v}_h^k = (\Lambda_h^k)^{-1} \sum_{l=1}^k \phi(s_h^l, a_h^l) V_h^k(s_{h+1}^l)$$

satisfies for any $h, k \in [H] \times [K]$ and for any state action pair $(s, a) \in \mathcal{S} \times \mathcal{A}$.

$$\left| \phi(s, a)^\top \mathbf{v}_h^k - P_h V_{h+1}^k(s, a) \right| \leq \beta \|\phi(s, a)\|_{(\Lambda_h^k)^{-1}} \quad (9)$$

with probability $1 - \delta$.

Proof. With analogous steps to the proof of Lemma 6 that

$$\begin{aligned} |\phi(s, a)^\top \mathbf{v}_h^k - P_h V_{h+1}^k(s, a)| &\leq \|MV_{h+1}^k\|_{(\Lambda_h^k)^{-1}} \|\phi(s, a)\|_{(\Lambda_h^k)^{-1}} \\ &\quad + \left\| \sum_{l=1}^k \phi(s_h^l, a_h^l) (V_{h+1}^k(s_{h+1}^l) - P_h V_{h+1}^k(s_h^l, a_h^l)) \right\|_{(\Lambda_h^k)^{-1}} \|\phi(s, a)\|_{(\Lambda_h^k)^{-1}} \end{aligned}$$

Then, using the fact that by assumption on M and by the clipping of the value function, we have that $\|MV_{h+1}^k\|_{(\Lambda_h^k)^{-1}} \leq H$. Then, using (Jin et al., 2019, Lemma B.3) it holds that with probability $1 - \delta$

$$\left\| \sum_{l=1}^k \phi(s_h^l, a_h^l) (V_{h+1}^k(s_{h+1}^l) - P_h V_{h+1}^k(s_h^l, a_h^l)) \right\|_{(\Lambda_h^k)^{-1}} \leq \mathcal{O} \left(dH \log \left(\frac{dK}{\delta} \right) \right)$$

□

Then, noticing that this is the main term we conclude that

$$|\phi(s, a)^\top \mathbf{v}_h^k - P_h V_{h+1}^k(s, a)| \leq \mathcal{O} \left(dH \log \left(\frac{dK}{\delta} \right) \right) \|\phi(s, a)\|_{(\Lambda_h^k)^{-1}}.$$

Lemma 12. *It holds that with probability $1 - \delta$*

$$\sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{s, a \sim d_h^{\pi^k}} [b_h^{k-1}(s, a)] \leq \mathcal{O} \left(d^{3/2} H^2 \sqrt{K \log \left(\frac{2K}{\delta} \right)} \right)$$

Proof. We have that

$$\begin{aligned} 2 \sum_{k=1}^K \sum_{h=1}^H \mathbb{E}_{s, a \sim d_h^{\pi^k}} [b_h^{k-1}(s, a)] &\leq 2 \sum_{k=1}^K \sum_{h=1}^H b_h^{k-1}(s_h^k, a_h^k) + \beta H \sqrt{K \log \delta^{-1}} \quad (\text{Azuma-Hoeffding}) \\ &\leq 2 \sum_{h=1}^H \sqrt{K \sum_{k=1}^K (b_h^{k-1}(s_h^k, a_h^k))^2} + \beta H \sqrt{K \log \delta^{-1}} \quad (\text{Cauchy-Schwartz}) \\ &= 2\beta \sum_{h=1}^H \sqrt{K \sum_{k=1}^K \phi(s_h^k, a_h^k)^\top (\Lambda_h^{k-1}) \phi(s_h^k, a_h^k)} + \beta H \sqrt{K \log \delta^{-1}} \\ &\leq 2\beta \sum_{h=1}^H \sqrt{dK \log(2K)} + \beta H \sqrt{K \log \delta^{-1}} \\ &= \mathcal{O} \left(d^{3/2} H^2 \sqrt{K \log \left(\frac{2K}{\delta} \right)} \right) \end{aligned}$$

□

G. Experiments

G.1. Experiments with deterministic expert

We also run an experiment where the expert is deterministic and see if ILARL can compete with BC in this setting. The results are provided in Figure 2. The parameter σ is the probability at which the system does not evolve according to the agent's action but in an adversarial way. We experiment with $\sigma = \{0, 0.05, 0.1\}$. The details about the transition dynamics are given in Appendix G.2. From Figure 2, we can see that ILARL and REIRL are again the most efficient algorithms in terms of MDP trajectories and they are able to match the performance of behavioural cloning despite the fact it has better guarantees for the case of deterministic experts. For Figure 1, we used $\sigma = 0.1$.




Figure 2: Experiments on the continuous gridworld with one trajectory from a deterministic expert.


G.2. Environment description

We run the experiment in the following MDP with continuous states space. We consider a 2D environment, where we denote the horizontal coordinate as $x \in [-1, 1]$ and vertical one as $y \in [-1, 1]$. The agent starts in the upper left corner, i.e., the coordinate $[-1, 1]^\top$ and should learn to reach the opposite corner (i.e. $[1, -1]^\top$) while avoiding the central high cost area depicted in Figure 3. The reward function is given by: $c_{\text{true}}(s, a) = c_{\text{true}}([x, y]^\top, a) = (x - 1)^2 + (y + 1)^2 + 80 \cdot e^{-8(x^2 + y^2)} - 100 \cdot \mathbb{1}\{x \in [0.95, 1], y \in [-1, -0.95]\}$. The action space for the agent is given by

$$\mathcal{A} = \left\{ \underbrace{[0.01, 0]^\top}_{\triangleq A_1}, \underbrace{[0, 0.01]^\top}_{\triangleq A_2}, \underbrace{[-0.01, 0]^\top}_{\triangleq A_3}, \underbrace{[0, -0.01]^\top}_{\triangleq A_4} \right\}, \text{ and the transition dynamics are given by:}$$

$$s_{t+1} = \begin{cases} \Pi_{[-1,1]^2} \left[s_t + \frac{a_t}{10} \right] & \text{w.p. } 1 - \sigma \\ \Pi_{[-1,1]^2} \left[s_t - \frac{s_t}{10 \|s_t\|_2} \right] & \text{w.p. } \sigma \end{cases}$$

Thus, with probability σ , the environment does not respond to the action taken by the agent, but it takes a step towards the


 Figure 3: Graphical representation of $-c_{\text{true}}$ of the linear MDP used in Figures 1 and 2.

low reward area centered at the origin, i.e., $-\frac{s_t}{10\|s_t\|_2}$. The agent should therefore pass far enough from the origin. Consider

$$\phi(s, a) = \phi([x, y], a) = [x^2, y^2, x, y, e^{-8(x^2+y^2)}, \mathbf{1}\{x \in [0.95, 1], y \in [-1, -0.95]\}, \mathbf{e}_a^\top]$$

with

$$\mathbf{e}_a = [\mathbf{1}\{a = A_1\}, \mathbf{1}\{a = A_2\}, \mathbf{1}\{a = A_3\}, \mathbf{1}\{a = A_4\}]^\top.$$

Notice that Assumption 2 holds only for the cost $\mathbf{c}_{\text{true}} = \Phi[1, 1, -2, -2, 80, -100, 2, 2, 2, 2]^\top$ while for the dynamics the linearity assumption does not hold.

G.3. Numerical verification of the finite horizon improvement.

We test BRIG (Algorithm 4) in a toy finite horizon problem. In particular, we consider a linear bandits problem ($H = 1$) with true cost function $\mathbf{c}_{\text{true}} = \Phi w_{\text{true}}$ where Φ entries are sampled from a normal distribution. For w_{true} we choose $w_{\text{true}}(i) = 0$ if i is odd and $w_{\text{true}}(i) = 1$ otherwise. We generate the expert dataset sampling 10 actions from a softmax expert. The results are shown in Figure 4. They confirm the theoretical findings that BRIG outperforms ILARL for finite horizon problems in terms of MDP trajectories.




Figure 4: Experiment in finite horizon setting to assess the better efficiency of BRIG.

G.4. Hyperparameters

For the experiments in Figures 1 and 2 we used $\eta = 1$, $\tau = 5$ and $\beta = 8$. For IQlearn, we also collect 5 trajectory to perform each update on the Q -function, and we use again $\eta = 1$ and 0.005 as stepsize for the Q -function weights. For PPIL, we use batches of 5 trajectories, 20 gradient updates between each batch collection, $\eta = 1$ and 0.005 as stepsize for the Q -function weights. For GAIL and AIRL, we use the default hyperparameters in <https://github.com/Khrylx/PyTorch-RL> but we obtained a better performance with a larger batch size of 6144 states and we use linear models rather than neural networks. For REIRL, we used the implementation in (Viano et al., 2021) but again we increased the batch size equal to 6144 states for achieving a better performance.

H. Reducing the number of expert trajectories.

In this section, we show that the number of required expert trajectories can be further reduced at the price of additional assumption on the expert policy, features and expert occupancy measure. The estimator we use is build on the ideas underling Mimic-MD in the linear case (Rajaraman et al., 2021).

Remark 4. Using such an estimator in ILARL or BRIG allows to improve upon Mimic-MD in two ways. Indeed ILARL and BRIG are provably computationally efficient algorithms and do not require knowledge of the dynamics. On the other hand, Mimic-MD requires perfect knowledge of the transition dynamics and it is unclear if the output policy can be computed efficiently in Linear MDPs.

To this goal, we need to consider the following estimator for $\Phi^\top d^{\pi_E}$, where we denote via (s_h^τ, a_h^τ) the state action pair encountered at step h in the trajectory τ

$$\begin{aligned} \widetilde{\Phi^\top d^{\pi_E}} &= (1 - \gamma) \mathbb{E}_{\tau \sim \pi_E} \left[\sum_{h=1}^{\infty} \gamma^h \phi(s_h^\tau, a_h^\tau) \mathbb{1}[s_t \in \mathcal{K} \quad \forall s_t \in \tau] \right] \\ &\quad + (1 - \gamma) \mathbb{E}_{\tau \sim \text{Unif}(D_1)} \left[\sum_{h=1}^{\infty} \gamma^h \phi(s_h^\tau, a_h^\tau) \mathbb{1}[\exists s_h \in \tau \text{ s.t. } s_h \notin \mathcal{K}] \right] \end{aligned} \quad (10)$$

where we split the expert dataset \mathcal{D}_{π_E} in two disjoint halves D_0, D_1 . The first D_0 is used to compute the set \mathcal{K} which is according to (Rajaraman et al., 2021, Definition 7) the set where the policy π_{BC} learned via Behavioural Cloning on the input dataset D_0 equals the expert policy. That is, $\mathcal{K} = \{s \in \mathcal{S} \text{ s.t. } \pi_{BC}(s) = \pi_E(s)\}^2$. The other half denoted via D_1 is used for the second term in 10. In the analysis of (Rajaraman et al., 2021) the first term can be computed thanks to the perfect knowledge of the dynamics. In our case, we have only trajectory access so we use the estimator

$$\begin{aligned} \underline{\Phi^\top d^{\pi_E}} &= (1 - \gamma) \mathbb{E}_{\tau \sim \text{Unif}(\mathcal{D}_{\pi_{BC}})} \left[\sum_{h=1}^{\infty} \gamma^h \phi(s_h^\tau, a_h^\tau) \mathbb{1}[s_t \in \mathcal{K} \quad \forall s_t \in \tau] \right] \\ &\quad + (1 - \gamma) \mathbb{E}_{\tau \sim \text{Unif}(D_1)} \left[\sum_{h=1}^{\infty} \gamma^h \phi(s_h^\tau, a_h^\tau) \mathbb{1}[\exists s_h \in \tau \text{ s.t. } s_h \notin \mathcal{K}] \right] \end{aligned} \quad (11)$$

where the dataset $\mathcal{D}_{\pi_{BC}}$ contains trajectories sampled according to π_{BC} .

Lemma 13. *Let us consider the estimator $\underline{\Phi^\top d^{\pi_E}}$ with the set \mathcal{K} be the confidence set for a binary linear classifier as defined in (Rajaraman et al., 2021, Section 4.1), let the expert policy be deterministic and satisfy the Linear Expert Assumption (Rajaraman et al., 2021, Definition 4). Moreover consider features that satisfy $-\phi(s, 1) = \phi(s, 0) = s/2$ for all $s \in \mathcal{S}$ where the state space is chosen to be \mathbb{R}^d . Finally, let consider that $\sum_{a \in \mathcal{A}} d^{\pi_E}(\cdot, a)$ is the uniform distribution $\text{Unif}(\mathcal{S})$, then it holds that for any $\delta > 0$*

$$\mathbb{E} \|\underline{\Phi^\top d^{\pi_E}} - \Phi^\top d^{\pi_E}\|_\infty \leq \frac{1}{1 - \gamma} \sqrt{\frac{\log(d/\delta)}{2|\mathcal{D}_{\pi_{BC}}|}} + \frac{\delta}{1 - \gamma} + \mathcal{O}\left(\frac{d^{5/4} \log d}{(1 - \gamma)^{3/2} |\mathcal{D}_{\pi_E}|}\right) \quad (12)$$

Remark 5. *The Lemma above follows the construction in (Rajaraman et al., 2021) to show that there exists one example under which ILARL used with estimator $\underline{\Phi^\top d^{\pi_E}}$ requires only $\tilde{\mathcal{O}}(d^{5/4} \epsilon^{-1} (1 - \gamma)^{-3/2})$ expert trajectories. However, it remains open to prove that the same holds true for general expert in Linear MDPs without further assumptions on the features and expert occupancy measure.*

Proof. The error can be controlled as follow

$$E \triangleq \mathbb{E}_{\tau \sim \pi_E} \left[\sum_{h=1}^{\infty} \gamma^h \phi(s_h^\tau, a_h^\tau) \mathbb{1}[s_t \in \mathcal{K} \quad \forall s_t \in \tau] \right] - \mathbb{E}_{\tau \sim \text{Unif}(\mathcal{D}_{\pi_{BC}})} \left[\sum_{h=1}^{\infty} \gamma^h \phi(s_h^\tau, a_h^\tau) \mathbb{1}[s_t \in \mathcal{K} \quad \forall s_t \in \tau] \right]$$

so denoting $X(\tau) \triangleq \sum_{h=1}^{\infty} \gamma^h \phi(s_h^\tau, a_h^\tau) \mathbb{1}[s_t \in \mathcal{K} \quad \forall s_t \in \tau]$ and noticing that by definition of \mathcal{K} we have that

$$\mathbb{E}_{\tau \sim \text{Unif}(\mathcal{D}_{\pi_{BC}})} \left[\sum_{h=1}^{\infty} \gamma^h \phi(s_h^\tau, a_h^\tau) \mathbb{1}[s_t \in \mathcal{K} \quad \forall s_t \in \tau] \right] = \mathbb{E}_{\tau \sim \text{Unif}(\mathcal{D}_{\pi_E})} \left[\sum_{h=1}^{\infty} \gamma^h \phi(s_h^\tau, a_h^\tau) \mathbb{1}[s_t \in \mathcal{K} \quad \forall s_t \in \tau] \right],$$

²Notice that we consider a deterministic expert in this section as done in (Rajaraman et al., 2021). Therefore, we consider policies as mapping from states to actions, i.e. $\pi : \mathcal{S} \rightarrow \mathcal{A}$

we can rewrite E as a martingale difference sequence

$$\begin{aligned} \|E\|_\infty &= \left\| \mathbb{E}_{\tau \sim \pi_E} [X(\tau)] - \frac{1}{|\mathcal{D}_{\pi_{BC}}|} \sum_{\tau \in \mathcal{D}_{\pi_{BC}}} X(\tau) \right\|_\infty \\ &= \left\| \frac{1}{|\mathcal{D}_{\pi_{BC}}|} \sum_{\tau \in \mathcal{D}_{\pi_{BC}}} (X(\tau) - \mathbb{E}_{\tau \sim \pi_E} [X(\tau)]) \right\|_\infty \\ &\leq \frac{1}{1-\gamma} \sqrt{\frac{\log(d/\delta)}{2|\mathcal{D}_{\pi_{BC}}|}} \quad \text{w.p. } 1-\delta \end{aligned}$$

Therefore choosing $|\mathcal{D}_{\pi_{BC}}| = \frac{\log(d/\delta)}{2\epsilon^2(1-\gamma)^2}$ ensures $E \leq \epsilon$ with probability at least $1-\delta$. Therefore by Lemma 7,

$$\mathbb{E} \|E\|_\infty \leq \frac{1}{1-\gamma} \sqrt{\frac{\log(d/\delta)}{2|\mathcal{D}_{\pi_{BC}}|}} + \frac{\delta}{1-\gamma}$$

These trajectories can be simulated in the MDP therefore the latter it is not a requirement on the expert dataset size. The number of expert trajectories is crucial to control the error due to the trajectories containing trajectories not in \mathcal{K} , i.e. Equation (11). Denoting this error as E_2 we have

$$\begin{aligned} &\mathbb{E}_{D_0, D_1} \|E_2\|_\infty \\ &= \mathbb{E}_{D_0, D_1} \left[\left\| \mathbb{E}_{\tau \sim \pi_E} \left[\sum_{h=1}^{\infty} \gamma^h \phi(s_h^\tau, a_h^\tau) \mathbb{1}[\exists s_h \in \tau \text{ s.t. } s_h \notin \mathcal{K}] \right] \right. \right. \\ &\quad \left. \left. - \mathbb{E}_{\tau \sim \text{Unif}(D_1)} \left[\sum_{h=1}^{\infty} \gamma^h \phi(s_h^\tau, a_h^\tau) \mathbb{1}[\exists s_h \in \tau \text{ s.t. } s_h \notin \mathcal{K}] \right] \right\| \right] \\ &\leq \frac{1}{(1-\gamma)} \sqrt{\frac{d}{|D_1|} \mathbb{E}_{\tau \sim D_0} [\mathbb{1}[\exists s_h \in \tau \text{ s.t. } s_h \notin \mathcal{K}]]} \\ &= \frac{1}{(1-\gamma)} \sqrt{\frac{d}{|D_1|} \mathbb{E}_{\text{lenght}(\tau)} [\mathbb{E}_{\tau \sim D_0 | \text{lenght}(\tau)} [\mathbb{1}[\exists s_h \in \tau \text{ s.t. } s_h \notin \mathcal{K}]]]} \quad \text{Tower Property of Expectation} \\ &\leq \frac{1}{(1-\gamma)} \sqrt{\frac{d}{|D_1|} \mathbb{E}_{\text{lenght}(\tau)} \left[\sum_{h=1}^{\text{lenght}(\tau)} \mathbb{E}_{\tau \sim D_0 | \text{lenght}(\tau)} [\mathbb{1}[s_h \notin \mathcal{K}]] \right]} \quad \text{Union Bound} \\ &\leq \frac{1}{(1-\gamma)} \sqrt{\frac{d}{|D_1|} \mathbb{E}_{\text{lenght}(\tau)} \left[\sum_{h=1}^{\text{lenght}(\tau)} \mathcal{O}\left(\frac{d^{3/2} \log d}{|D_0|}\right) \right]} \quad \text{Thanks to (Rajaraman et al., 2021, Theorem 7)} \\ &\leq \mathcal{O}\left(\frac{1}{(1-\gamma)} \sqrt{\frac{d^{5/2} \log d}{|D_1|^2} \mathbb{E}_{\text{lenght}(\tau)} [\text{lenght}(\tau)]}\right) \quad \text{Using that } |D_0| = |D_1| \text{ by construction} \\ &\leq \mathcal{O}\left(\frac{1}{(1-\gamma)} \sqrt{\frac{d^{5/2} \log d}{|D_1|^2} \frac{1}{1-\gamma}}\right) \\ &= \mathcal{O}\left(\frac{d^{5/4} \log d}{(1-\gamma)^{3/2} |D_1|}\right) \end{aligned}$$

Where we used (Rajaraman et al., 2021, Theorem 7) to bound

$$\mathbb{E}_{\tau \sim D_0 | \text{lenght}(\tau)} [\mathbb{1}[s_h \notin \mathcal{K}]] \leq \mathcal{O}\left(\frac{d^{3/2} \log d}{|D_0|}\right)$$

so overall

$$\mathbb{E}_{D_0, D_1} \|E_2\|_\infty \leq \mathcal{O}\left(\frac{d^{5/4} \log d}{(1-\gamma)^{3/2} |D_1|}\right)$$

□