
Minimax Optimal Online Imitation Learning via Replay Estimation

Gokul Swamy*
Carnegie Mellon University
gswamy@cmu.edu

Nived Rajaraman*
UC Berkeley
nived.rajaraman@berkeley.edu

Matthew Peng
UC Berkeley

Sanjiban Choudhury
Cornell University

J. Andrew Bagnell
Aurora Innovation and Carnegie Mellon University

Zhiwei Steven Wu
Carnegie Mellon University

Jiantao Jiao
UC Berkeley

Kannan Ramchandran
UC Berkeley

Abstract

Online imitation learning is the problem of how best to mimic expert demonstrations, given access to the environment or an accurate simulator. Prior work has shown that in the *infinite* sample regime, exact moment matching achieves value equivalence to the expert policy. However, in the *finite* sample regime, even if one has no optimization error, empirical variance can lead to a performance gap that scales with H^2/N_{exp} for behavioral cloning and $H/\sqrt{N_{\text{exp}}}$ for online moment matching, where H is the horizon and N_{exp} is the size of the expert dataset. We introduce the technique of *replay estimation* to reduce this empirical variance: by repeatedly executing cached expert actions in a stochastic simulator, we compute a *smoother* expert visitation distribution estimate to match. In the presence of parametric function approximation, we prove a meta theorem reducing the performance gap of our approach to the *parameter estimation error* for offline classification (i.e. learning the expert policy). In the tabular setting or with linear function approximation, our meta theorem shows that the performance gap incurred by our approach achieves the optimal $\tilde{O}(\min(H^{3/2}/N_{\text{exp}}, H/\sqrt{N_{\text{exp}}})$ dependency, under significantly weaker assumptions compared to prior work Rajaraman et al. [2021]. We implement multiple instantiations of our approach on several continuous control tasks and find that we are able to significantly improve policy performance across a variety of dataset sizes.

1 Introduction

In *online* imitation learning (IL), one is given access to (a) a fixed set of expert demonstrations and (b) an environment or simulator to perform rollouts in. Many online IL approaches fall under the umbrella of solving a *moment matching* problem between learner and expert trajectory distributions Ziebart et al. [2008], Ho and Ermon [2016],

$$\min_{\pi \in \Pi} \sup_{f \in \mathcal{F}} \mathbb{E}_{\pi} [f(s, a)] - \mathbb{E}_{\pi^E} [f(s, a)], \quad (1)$$

where $\mathbb{E}_{\pi}[\cdot]$ denotes the expectation over a random trajectory $\{(s_1, a_1), \dots, (s_H, a_H)\}$ generated by rolling out π . Swamy et al. [2021] show that for an appropriate choice of \mathcal{F} , approximate solutions

*Equal contribution. Correspondence to gswamy@cmu.edu and nived.rajaraman@berkeley.edu.

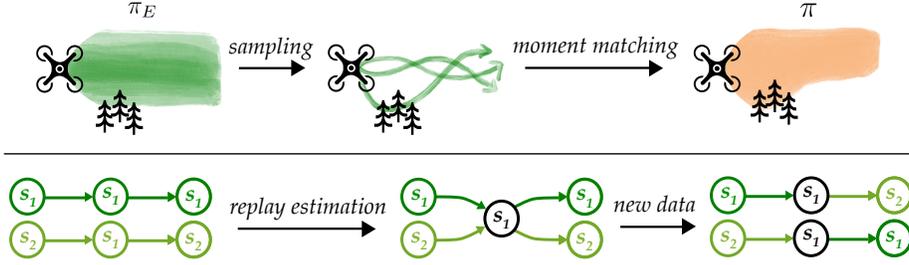


Figure 1: *Top*: Attempting to exactly match a finite-sample approximation of expert moments can cause a learner to reproduce chance occurrences (e.g. the relatively unlikely flight through the trees). This can lead to policies that perform poorly at test time (e.g. because the learner flies through the trees relatively often). *Bottom*: Replay estimation reduces the empirical variance in expert demonstrations by repeatedly executing observed expert actions in a stochastic simulator. By generating new trajectories (e.g. $s_1 \rightarrow s_1 \rightarrow s_2$ on the right) that are consistent with expert actions, one can augment the original demonstration set and compute expert moments more accurately.

of (1) have a performance gap linear in the horizon, the gold standard for sequential problems. However, a key assumption in their work is that expert moments ($\mathbb{E}_{\pi_E}[f(s, a)]$) can be estimated arbitrarily well from the available demonstrations. The resulting bounds are therefore purely a function of *optimization error*. Moving to the finite-sample regime introduces an additional concern: the *statistical error* that stems from the randomness in the data-generating process. When a small set of demonstrations are used in an adversarial optimization procedure like (1), the learner may choose to take incorrect actions in order to match noisy moments estimated from the dataset, leading to policies that perform poorly at test-time. Ideally, one would solve this problem by querying the expert for more demonstrations, as in the work of Ross et al. [2011]. However, when we are unable to do so, we still have to grapple with the question of “*how can we smooth out a noisy empirical estimate of expert moments?*”

Our answer to this question is the technique of *replay estimation*. In its most basic form, replay estimation consists of repeatedly executing observed expert actions within a stochastic simulator, terminating rollouts whenever one ventures out of the support of the expert demonstrations. Effectively, this approach stitches together parts of different trajectories to generate a smoothed estimate of expert moments. By using the simulator where we know the expert’s actions, we can generate more diverse training data that is nevertheless consistent with the expert demonstrations. We argue that this technique is at once conceptually simple, practically feasible, and minimax optimal in several settings. Formally, we prove that in the worst case, behavioral cloning has a performance gap $\propto H^2/N_{\text{exp}}$, online empirical moment matching, $\propto H/\sqrt{N_{\text{exp}}}$, and our approach of replay estimation $\propto \min\{H^{3/2}/N_{\text{exp}}, H/\sqrt{N_{\text{exp}}}\}$ in the tabular setting as well as with linear function approximation. Our key insight is that ***we can use a combination of simulated and empirical rollouts to optimally estimate expert moments***. We can then plug this improved estimate into a variety of moment-matching algorithms, for strong test-time policy performance. More explicitly, our work makes the following three contributions:

1. We extend replay estimation (RE) Rajaraman et al. [2020] beyond the tabular and deterministic setting by introducing the notion of a *soft membership oracle* and *prefix weights*.
2. We show how to instantiate the membership oracles for IL with parametric function approximation and prove a meta-theorem relating the imitation gap of RE to the parameter estimation error for offline classification on the dataset (Theorem 3). Instantiating our main result in the case of linear function approximation, we show how to achieve the best known imitation gap of $\tilde{O}(H^{3/2}d^{5/4}/N_{\text{exp}})$ under significantly weaker assumptions compared to prior work Rajaraman et al. [2021].
3. We give multiple practical options for constructing performant membership oracles. We then use these approximate oracles to significantly improve the performance of online IL on several continuous control tasks across a variety of dataset sizes. We also investigate the differences between our proposed oracles.

We sketch the benefits of replay estimation before providing theoretical and empirical evidence to support our claims.

2 The Replay Estimator

We begin with a tabular vignette to illustrate our key insight in greater detail. We compare two algorithms: offline behavioral cloning (BC) Pomerleau [1989] and online moment matching (MM) Swamy et al. [2021]. Throughout, we focus on learning policies from finite samples.

Suboptimality of Empirical Moment Matching. Consider the MDP in Fig. 2, where the expert always takes the green action. Doing so puts them in s_1 or s_2 with equal probability. Given that the expert is deterministic and there are few states, BC could easily recover the expert’s policy by learning to simply output the observed green action on both states, even when there are very few demonstrations.

Now, what would happen if we tried to match moments of the expert’s state-action visitation distribution for this problem? It is rather unlikely that we see *exactly* equal probabilities for both states in the observed data. If by chance we see s_2 more than we see s_1 , the learner might realize that the only way to match the observed state distribution (a prerequisite for matching the observed state-action distribution) is to occasionally take the red action at s_2 . In general, this could cause the learner to spend an unnecessary amount of time in s_2 which may be undesirable (e.g. if s_2 corresponds to the tree-filled area in Fig. 1 (top)). The core issue we hope to illustrate in this example is that by treating the empirical estimate of the expert’s behavior as perfectly accurate, distribution matching can force the learner to take incorrect actions to minimize training error, leading to test-time performance degradation. As we will discuss in Sec. 3, this can lead to slow statistical rates $\propto H/\sqrt{N_{\text{exp}}}$.

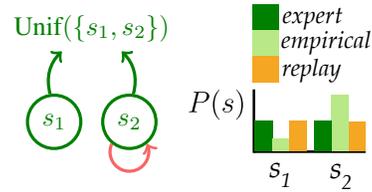


Figure 2: An MDP where the expert always takes the green action that puts them in the uniform distribution over s_1 and s_2 . Because of full expert support, BC will learn to always take this action at both states. However, if the empirical state distribution is more tilted towards s_2 , MM will take the incorrect red action.

Suboptimality of Behavioral Cloning. Because it does not account for the covariate shift that results from policy action choices, behavioral cloning can lead to a quadratic compounding of errors and poor test time performance Ross et al. [2011]. Consider, for example, the MDP in Fig. 3.

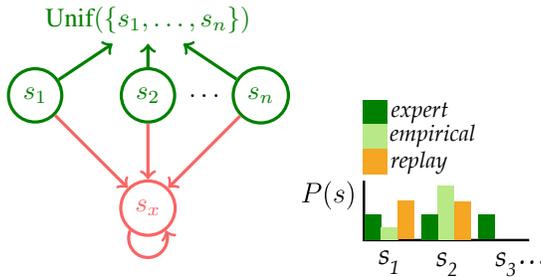


Figure 3: The expert always takes the green action, which places it in a uniform distribution over s_1, \dots, s_n . At states where we have demonstrations (e.g. s_1, s_2), both BC and MM will take the same, correct action. However, at states where we have no demonstrations (e.g. s_3), MM will correctly take the green action to get back to states with demonstration support, while BC might not.

Let us assume that the expert always takes the green action, dropping them in a state in the top row with uniform probability. In a small demonstration set, we might not see expert actions at some states in the top row. At all such states, BC will have no idea of what to do. In contrast, MM will take the green action as doing so might send the learner back to a state with positive demonstration support. Thus for this problem, MM will recover the optimal policy while BC will not. As we will discuss in Sec. 3, this leads to errors $\propto H^2/N_{\text{exp}}$ in the worst case.

Replay Estimation. The previous two examples show us that there exist simple MDPs for which BC or MM will not recover the expert’s policy. This begs the question: is it possible to do

better than both worlds and recover the optimal policy on both problems with a single algorithm? It turns out it is indeed possible to do so, via the technique of *replay estimation*. In its simplest form,

replay estimation involves playing a cached expert action whenever possible and re-starting the rollout if one ventures out of the support of the demonstrations. Then, one appends these rollouts to the demonstration set, treating them as additional training data – while biased, they are consistent with observed expert behavior. Intuitively, repeated simulation has a *smoothing* effect on the training data as doing so marginalizes out the statistical error that comes from the stochasticity of the dynamics. We can see this point more explicitly by considering the above two MDP examples: in Fig. 2, repeatedly playing the green action and appending these rollouts to the expert dataset would bring us much closer to a uniform distribution over s_1 and s_2 . Similarly, in Fig. 3, replay estimation would bring us toward a uniform distribution over the states $\{s_1, \dots, s_n\}$ in the expert demonstrations.

We could then plug in this improved distribution estimate into the MM procedure (1). Notice how doing so would cause MM to be highly likely to recover the optimal policy on both MDPs. For example, in Fig. 2, replay estimation would make the learner much less likely to play the red action in s_2 . It turns out this fact is sufficient to establish *statistical optimality* in the tabular setting and with linear function approximation, with an error rate $\propto \min(H^{3/2}/N_{\text{exp}}, H/\sqrt{N_{\text{exp}}})$. In short, replay estimation is a practical technique for reducing some of the finite-sample variance in expert demonstrations that enables MM to perform optimally in the finite sample regime. We now provide some intuition on how to generalize this approach to beyond the tabular setting.

Leaving the Tabular Setting. The prior work of Rajaraman et al. [2020] considers the tabular setting; this characteristic makes it easy to answer the question of “*on what states do we know the expert’s action?*” To enable us to answer this question more generally, we introduce the notion of a *membership oracle* $\mathcal{M} : \mathcal{S} \rightarrow \{0, 1\}$. Explicitly, $\mathcal{M}(s) = 1$ for states where we know the expert’s action well (e.g. states where we have lots of similar demonstrations) and $\mathcal{M}(s) = 0$ otherwise. Alternatively, $\mathcal{M}(s) = 1$ on states where BC, which attempts to directly output expert actions, is accurate and 0 otherwise.

We can then compute expert moments by splitting on the output of the membership oracle:

$$\mathbb{E}_{\pi^E}[f(s, a)] = \underbrace{\mathbb{E}_{\pi^E}[f(s, a)\mathbb{1}(\mathcal{M}(s) = 1)]}_{(i)} + \underbrace{\mathbb{E}_{\pi^E}[f(s, a)\mathbb{1}(\mathcal{M}(s) = 0)]}_{(ii)} \quad (2)$$

Note that the indicators in (i) and (ii) are complements of each other, rendering the above sum a valid estimate of the expert moment. As we know the expert action well wherever $\mathcal{M}(s) = 1$, simulated rollouts of the BC policy approximates (i) well; on the other hand we resort to a naive empirical estimate to approximate (ii), as we do not know enough about the expert’s action at these states to accurately generate additional demonstrations via BC rollouts. In general, we relax \mathcal{M} to a *soft membership oracle* Zadeh [1965], in order to handle uncertainty in how well we know the expert’s action at a given state. We proceed by first analyzing the statistical properties of applying MM to this bipartite estimator before discussing practical constructions of performant membership oracles.

3 Theoretical Analysis

The proofs of all results in this section are deferred to appendix A. We begin by introducing some notation.

Notation. Let $\Delta(X)$ denote the probability simplex over set X and let $\gtrsim, \lesssim, \asymp$ respectively denote greater than, lesser than and equality up to constants. We study the IL problem in the episodic MDP setting with state space \mathcal{S} , action space \mathcal{A} and horizon H . We assume that the transition, reward function and policies can be non-stationary. The MDP transition is denoted $P = \{\rho, P_1, \dots, P_{H-1}\}$, where ρ is the initial state distribution and $P_t : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$, while the reward function is denoted $r = \{r_1, \dots, r_H\}$ where $r_t : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$. In the online imitation learning setting, the learner has access to a finite dataset D of N_{exp} trajectories (i.e. the sequences of states visited and actions played) generated by rolling out expert policy π^E . Importantly, the learner *does not observe rewards during rollouts*. The fundamental goal of the learner is to learn a policy $\hat{\pi}$ such that the

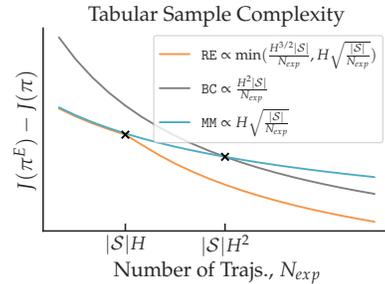


Figure 4: In the tabular setting, RE inherits the superior low-data performance of moment-matching approaches and is able to perform better than both MM and BC with enough data.

imitation gap, $J(\pi^E) - J(\hat{\pi})$ is small. Here, $J(\pi)$ denotes the expected value of the policy π , $\mathbb{E}_\pi \left[\sum_{t=1}^H r_t(s_t, a_t) \right]$. $\mathbb{E}_D[\cdot]$ denotes the empirical expectation computed using trajectories from the dataset D .

Behavior cloning. A standard approach for imitation learning is behavioral cloning (BC) Pomerleau [1989], which trains a classifier from expert states to expert actions. More formally, for the 0-1 loss ℓ (or a continuous proxy) BC minimizes the empirical classification error,

$$\pi^{\text{BC}} \leftarrow \arg \min_{\pi \in \Pi} \mathbb{E}_D \left[\frac{1}{H} \sum_{t=1}^H \ell(\pi_t(s_t), a_t) \right]. \quad (3)$$

It is known from Ross et al. [2011] and Rajaraman et al. [2020] that in the tabular setting, the expected imitation gap for BC is always $\lesssim |\mathcal{S}|H^2/N_{\text{exp}}$. The fast $1/N_{\text{exp}}$ rate comes from the fact that on states where the learner observes the deterministic expert’s actions, BC simply replays them. Thus, the performance gap is proportional to the mass of unseen states, which decays as $1/N_{\text{exp}}$. The H^2 dependence comes from the fact that a single mistake can take the learner out of the distribution of the expert, causing it to make mistakes for the rest of the horizon. Importantly, this bound is tight – i.e. there exists an MDP instance on which BC incurs this imitation gap (See Appendix A.3).

Theorem 1 (Theorem 6.1 of Rajaraman et al. [2020]). *Then, there exists a tabular MDP instance such that the BC incurs, $\mathbb{E} [J(\pi^E) - J(\pi^{\text{BC}})] \gtrsim \min \{H, |\mathcal{S}|H^2/N_{\text{exp}}\}$.*

Moment matching. Many standard algorithms in the IL literature fall under the empirical moment matching framework (e.g. GAIL Ho and Ermon [2016], MaxEnt IRL Ziebart et al. [2008]); see Table 3 of Swamy et al. [2021] for more examples. Reward moment matching corresponds to finding a policy which best matches the state-action visitation measure of π^E , in the sense of minimizing an Integral Probability Metric (IPM) Müller [1997]. In the finite sample setting, the *empirical moment matching* learner π^{MM} attempts to best match the empirical state-visitation measure. Namely,

$$\pi^{\text{MM}} \in \arg \min_{\pi \in \Pi} \sup_{f \in \mathcal{F}} \mathbb{E}_\pi \left[\frac{\sum_{t=1}^H f_t(s_t, a_t)}{H} \right] - \mathbb{E}_D \left[\frac{\sum_{t=1}^H f_t(s_t, a_t)}{H} \right]. \quad (4)$$

First we show an upper bound on the imitation gap incurred by empirical moment matching.

Theorem 3.1. *Consider the empirical moment matching learner π^{MM} (eq. (4)), instantiated with an appropriate discriminator class \mathcal{F} . The imitation gap satisfies $\mathbb{E} [J(\pi^E) - J(\pi^{\text{MM}})] \lesssim H\sqrt{|\mathcal{S}|/N_{\text{exp}}}$.*

The proof of this result can be found in Appendix A.2. It turns out that this guarantee is essentially tight for empirical moment matching, answering an open question from Rajaraman et al. [2020].

Theorem 3.2. *If $H \geq 4$, there is a tabular IL instance with 2 states and actions on which with constant probability, the empirical moment matching learner (eq. (4)) incurs, $J(\pi^E) - J(\pi^{\text{MM}}) \gtrsim H/\sqrt{N_{\text{exp}}}$.*

The proof of this result is deferred to appendix A.4. The proof of this lower bound exploits the fact that the data generation process in the dataset is inherently random. Consider a slight modification of the MDP instance shown in fig. 2, where the reward function is 0 for $t = 1$. For $t \geq 2$, the transition function is absorbing at both states; the reward function equals 1 at the state s_1 for any action and is 0 everywhere else. Then, the expert state distribution at time 2 and every time thereon is in uniform across the two states, $\{1/2, 1/2\}$. However, in the dataset D , the learner sees a noisy realization of this distribution in the dataset of the form $\{1/2 - \delta, 1/2 + \delta\}$ for $|\delta| \approx \pm 1/\sqrt{N_{\text{exp}}}$. Because of this noise, the empirical moment matching learner may be encouraged to deviate from the expert’s observed behavior and pick the red action at s_2 as this results in a better match to the empirical state visitation measures at every point in the rest of the episode - a prerequisite to matching the empirical state-action visitation measure. The learner is willing to pick an action different from what the expert played in order to better match the *inherently noisy* empirical state-action visitation distribution.

Remark 1. *Theorems 1 and 3.2 are separate lower bound IL instances against the performance of BC and empirical moment matching. On the uniform mixture of the two MDPs (i.e. deciding the underlying MDP based on the outcome of a fair coin), with constant probability, both $J(\pi^E) - J(\pi^{\text{BC}}) \gtrsim |\mathcal{S}|H^2/N_{\text{exp}}$ and $J(\pi^E) - J(\pi^{\text{MM}}) \gtrsim H/\sqrt{N_{\text{exp}}}$. On this mixture instance, training both BC and empirical moment matching and choosing the better of the two is also statistically suboptimal.*

3.1 Replay Estimation

A natural question at this point is whether there is an algorithm that is *better than both worlds*, i.e. algorithm which can outperform the worst case imitation gap of empirical moment matching and BC . The answer to this question in the tabular setting was recently provided by Rajaraman et al. [2020] who propose Mimic-MD , achieving better performance than both BC and MM . This improvement is possible because BC does not use any dynamics information and MM does not leverage the knowledge of where expert actions are known. However, it is unclear how to extend this approach beyond the tabular setting as the algorithm relies on a large measure of states being visited in the demonstrations.

Algorithm 1: Replay Estimation (RE)

Input: Expert demonstrations D , policy class Π , moment class $\mathcal{F} = \bigoplus_{t=1}^H \mathcal{F}_t$, simulator SIM, ALG which returns a membership oracle given a dataset;

- 1 Partition the dataset D into D_1 and D_2 ;
- 2 Using ALG, train the membership oracle \mathcal{M} on D_1 ;
- 3 Train π^{BC} using behavior cloning on D_1 ;
- 4 Roll out π^{BC} in SIM N_{replay} times to construct a new dataset, D_{replay} ;
- 5 Define prefix weights $\mathcal{P}(s_{1\dots t}) = \prod_{t'=1}^t \mathcal{M}(s_{t'}, t')$;
- 6 Define,

$$7 \quad \widehat{E}(f) = \mathbb{E}_{D_{\text{replay}}} \left[\frac{1}{H} \sum_{t=1}^H f_t(s_t, a_t) (\mathcal{P}(s_{1\dots t})) \right] + \mathbb{E}_{D_2} \left[\frac{1}{H} \sum_{t=1}^H f_t(s_t, a_t) (1 - \mathcal{P}(s_{1\dots t})) \right].$$

Output: π^{RE} , a solution to the moment-matching problem:

$$\arg \min_{\pi \in \Pi} \sup_{f \in \mathcal{F}} \mathbb{E}_{\pi} \left[\frac{1}{H} \sum_{t=1}^H f_t(s_t, a_t) \right] - \widehat{E}(f) \quad (5)$$

To handle this challenge, we introduce the notion of a *soft membership oracle* Zadeh [1965], $\mathcal{M} : \mathcal{S} \times [H] \rightarrow [0, 1]$ which captures the learner’s inherent uncertainty in the expert’s actions at a state at each point in an episode. The soft membership oracle assigns high weight to a state if BC is likely to closely agree with the expert policy and gives a lower weight to states where BC is likely to be inaccurate. By this definition, if the membership oracle is consistently large at all the states visited in a trajectory, we can be confident that *a trajectory generated by BC is as though it was a rollout from the expert policy*. Formally, for any function g and time $t = 1, \dots, H$, we have the decomposition,

$$\mathbb{E}_{\pi^E} [g(s_t, a_t)] = \underbrace{\mathbb{E}_{\pi^E} [g(s_t, a_t) \mathcal{P}(s_{1\dots t})]}_{(i)} + \underbrace{\mathbb{E}_{\pi^E} [g(s_t, a_t) (1 - \mathcal{P}(s_{1\dots t}))]}_{(ii)} \quad (6)$$

where $\mathcal{P}(s_{1\dots t})$ is defined as the *prefix weight* $\prod_{t'=1}^t \mathcal{M}(s_{t'}, t')$. We need to use prefix weights instead of the single-sample weights sketched in the previous section to account for the probability of BC getting to the current state in the same manner the expert would have. Because of the high accuracy of BC on segments with high prefix weights, in eq. (6), (i) can be approximated by replacing the expectation over π^E by that over π^{BC} , i.e. replay estimation. On the other hand, since the prefix weight is low on the remaining trajectories in (ii), we know that BC is inaccurate, so we resort to using a simple empirical estimate to estimate this term.

While we leave the particular choice of the soft membership oracle flexible, intuitively, states at which BC closely agrees with the expert policy should be given high weight while where those where BC is inaccurate should be weighted lower. In Section 4, we discuss several practical approaches to designing such a soft membership oracle. We first prove a generic policy performance guarantee for the outputs of our algorithm as a function of the choice of \mathcal{M} .

Theorem 2. Consider the policy π^{RE} returned by Algorithm 1. Assume that $\pi^E \in \Pi$ and the ground truth reward function $r_t \in \mathcal{F}_t$, which is assumed to be symmetric ($f_t \in \mathcal{F}_t \iff -f_t \in \mathcal{F}_t$) and bounded (For all $f_t \in \mathcal{F}_t$, $\|f_t\|_{\infty} \leq 1$). Choose $|D_1|, |D_2| = \Theta(N_{\text{exp}})$ and suppose $N_{\text{replay}} \rightarrow \infty$.

With probability $\geq 1 - 3\delta$,

$$J(\pi^E) - J(\pi^{RE}) \lesssim \mathcal{L}_1 + \mathcal{L}_2 + \frac{\log(F_{\max}H/\delta)}{N_{\text{exp}}} \quad (7)$$

where $F_{\max} \triangleq \max_{t \in [H]} |\mathcal{F}_t|$, and,

$$\mathcal{L}_1 \triangleq H^2 \mathbb{E}_{\pi^E} \left[\frac{\sum_{t=1}^H \mathcal{M}(s_t, t) \text{TV}(\pi_t^E(\cdot|s_t), \pi_t^{BC}(\cdot|s_t))}{H} \right], \quad (8)$$

$$\mathcal{L}_2 \triangleq H^{3/2} \sqrt{\frac{\log(F_{\max}H/\delta)}{N_{\text{exp}}} \frac{\sum_{t=1}^H \mathbb{E}_{\pi^E} [1 - \mathcal{M}(s_t, t)]}{H}}.$$

We discuss a proof of this result in Appendix B and include bounds when N_{replay} is finite.

Remark 2. Note that Theorem 2 can be extended to infinite function families using the standard technique of replacing $|\mathcal{F}_t|$ by the ϵ log-covering number of \mathcal{F}_t in say, the L_2 norm, for $\epsilon = \frac{1}{N_{\text{exp}}}$. For ease of exposition here, we stick to the case where \mathcal{F}_t is finite.

The term \mathcal{L}_1 measures how accurate BC is on states from expert trajectories where $\mathcal{M}(s_t, t)$ is large. Intuitively, if we set $\mathcal{M}(s_t, t) = 1$ on states where BC is accurate and $\mathcal{M}(s_t, t) = 0$ elsewhere, we would expect this term to be small. \mathcal{L}_2 can be thought of a measure of BC's coverage: it tells us how much of the expert's visitation distribution we believe BC to be inaccurate on. If BC has good coverage (i.e. $1 - \mathcal{M}(s_t, t)$ is small on expert trajectories), we expect this term to be small.

Prima facie, one might think that because \mathcal{L}_1 resembles the imitation gap of BC and \mathcal{L}_2 resembles that of MM, RE can only perform as well as the best of BC ($\propto H^2/N_{\text{exp}}$) and MM ($\propto H/\sqrt{N_{\text{exp}}}$) on a given instance. However, with a careful choice of \mathcal{M} , one can achieve "better than both worlds" statistical rates. In particular, since RE is a generalization of Mimic-MD of Rajaraman et al. [2020], in the tabular setting, an appropriately initialized version of RE achieves the optimal imitation gap of $\min \left\{ \frac{|\mathcal{S}|H^{3/2}}{N_{\text{exp}}}, H\sqrt{\frac{|\mathcal{S}|}{N_{\text{exp}}}} \right\} \log \left(\frac{|\mathcal{S}|H}{\delta} \right)$ and strictly improves over both BC and MM.

We now show how to extend this result and instantiate the membership oracle for parametric function approximation and provide a statistical guarantee under a particular margin assumption.

3.2 Parametric Function Approximation: Reduction to Offline Classification

While BC can be thought of as a reduction of IL to offline classification, the algorithm does not take into account the knowledge of the transition of the MDP. This is reflected in the quadratic dependency in the horizon, error compounding Ross et al. [2011]. In this section, we study a novel reduction of IL with parametric function approximation to *parameter estimation in offline classification*, which we define formally. We provide a provable guarantee, assuming the learner has access to a classification oracle and the underlying function class admits a Lipschitz parameterization.

Definition 1 (IL with function-approximation). *In this setting, for each $t \in [H]$, there is a parameter class $\Theta_t \subseteq \mathbb{B}_2^d$, the unit L_2 ball in d dimensions, and an associated function class $\{f_{\theta_t} : \theta_t \in \Theta_t\}$. For each $t \in [H]$ there exists an unknown $\theta_t^E \in \Theta_t$ such that $\forall s \in \mathcal{S}$,*

$$\pi_t^E(s) = \arg \max_{a \in \mathcal{A}} f_{\theta_t^E}(s, a). \quad (9)$$

Definition 2 (Lipschitz parameterization). *A function class $\mathcal{G} = \{g_{\theta} : \theta \in \Theta\}$ where $g_{\theta}(\cdot) : \mathcal{X} \rightarrow \mathbb{R}$ is said to satisfy L -Lipschitz parameterization if, $\|g_{\theta}(\cdot) - g_{\theta'}(\cdot)\|_{\infty} \leq L\|\theta - \theta'\|_2$. In other words, for each $x \in \mathcal{X}$, $g_{\theta}(x)$ is an L -Lipschitz function in θ , in the L_2 norm.*

Assumption 1. *For each t , the class $\{f_{\theta_t} : \theta_t \in \Theta_t\}$ is L -Lipschitz in its parameterization, $\theta_t \in \Theta_t$.*

To deal with parametric function approximation, we assume that the learner has access to an *offline classification oracle* which, given a dataset of classification examples, approximately returns the underlying ground truth parameter. More formally,

Assumption 2 (Offline classification oracle). *We assume that the learner has access to a multi-class classification oracle, which given n examples of the form, (s^i, a^i) where $s^i \stackrel{i.i.d.}{\sim} \mathcal{D}$ and $a^i = \arg \max_{a \in \mathcal{A}} f_{\theta^*}(s^i, a)$, returns a $\hat{\theta} \in \Theta$ such that, with probability $\geq 1 - \delta$, $\|\hat{\theta} - \theta^*\|_2 \leq \mathcal{E}_{\Theta, n, \delta}$.*

We assume that this classification oracle is used by RE to train the BC policy in Line 3 of Algorithm 1.

A careful reader might note that Assumption 2 asks for a slightly stronger requirement than just finding a classifier with small generalization error (which need not be close to the ground truth θ^*). The latter problem was studied in Daniely et al. [2013] who show that the *Natarajan dimension*, up to log-factors in the number of classes (i.e. number of actions) captures the generalization error of the best learner, which scales as $\Theta(1/n)$ given n classification examples. Under some mild regularity (coverage) assumptions on the input distribution \mathcal{D} , we show that the generalization guarantee carries over to learning the parameter θ , for example in linear classification (Appendix B.4). In particular, we show that $\mathcal{E}_{\Theta, n, \delta} \lesssim \frac{d + \log(1/\delta)}{n}$ when $f_{\theta} = \langle \theta, \cdot \rangle$. The membership oracle we study is defined below,

$$\mathcal{M}(s, t) = \begin{cases} +1 & \text{if } \exists a \in \mathcal{A} \text{ such that, } \forall a' \in \mathcal{A}, f_{\hat{\theta}_t^{\text{BC}}}(s, a) - f_{\hat{\theta}_t^{\text{BC}}}(s, a') \geq 2L\mathcal{E}_{\Theta_t, N_{\text{exp}}, \delta/H} \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

Intuitively, \mathcal{M} assigns a state as +1 if BC classifies it with a significant margin as some action.

Finally, we introduce the main assumption on the IL instances we study. We assume that the classification problems solved by BC at each $t \in [H]$ satisfy a margin condition.

Assumption 3 (Weak margin condition). *The weak margin condition assumes that for each t , there is no classifier $\theta \in \Theta_t$ such that for a large mass of states, $f_{\theta}(s_t, \pi_t^E(s_t)) - \max_{a \neq \pi_t^E(s_t)} f_{\theta}(s_t, a)$, i.e. the “margin” from the nearest classification boundary, is small. Formally, the weak-margin condition with parameter μ states that, for any $\theta \in \Theta_t$ and $\eta \leq 1/\mu$,*

$$\Pr_{\pi^E} \left(f_{\theta}(s_t, \pi_t^E(s_t)) - \max_{a \neq \pi_t^E(s_t)} f_{\theta}(s_t, a) \geq \eta \right) \geq e^{-\mu\eta}. \quad (11)$$

The weak margin condition only assumes that there is at least an exponentially small (in η) mass of states with margin at least η . Smaller μ indicates a larger mass away from any decision boundary. It suffices to assume that eq. (11) is only true for θ as the classifier in Assumption 2 for our guarantees (Theorem 3) to hold.

Under these three assumptions and using the membership oracle defined above, we can provide a strong guarantee for RE :

Theorem 3. *For IL with parametric function approximation, under Assumptions 1 to 3, appropriately instantiating RE ensures that with probability $\geq 1 - 4\delta$,*

$$J(\pi^E) - J(\pi^{\text{RE}}) \lesssim H^{3/2} \sqrt{\frac{\mu L \log(F_{\max} H / \delta)}{N_{\text{exp}}} \frac{\sum_{t=1}^H \mathcal{E}_{\Theta_t, N_{\text{exp}}, \delta/H}}{H}} + \frac{\log(F_{\max} H / \delta)}{N_{\text{exp}}}. \quad (12)$$

In Appendix B.4, we explicitly instantiate these guarantees for the special case of where the expert follows a linear classifier acting on a known set of feature representations of the state-action pairs. The prior work of Rajaraman et al. [2021] shows an imitation gap of $\tilde{O}(H^{3/2} d^{5/4} / N_{\text{exp}})$ for RE in the linear expert setting, under the restrictive assumptions of a binary action space, and a certain uniform distribution assumption on the feature distribution. In contrast, the guarantee of Theorem 3 can be used to achieve the same guarantee, for general action spaces and under a significant weakening of the uniform feature distribution constraint to a multi-class analog of the *strong distribution assumption* Audibert and Tsybakov [2007]. Under this assumption, we show that the best achievable parameter recovery error, $\mathcal{E}_{\Theta_t, N_{\text{exp}}, \delta/H}$ and the best achievable 0-1 generalization error ($\mathbb{E}_{s \sim \mathcal{D}}[\max_{a \in \mathcal{A}} f_{\hat{\theta}}(s, a) \neq \max_{a \in \mathcal{A}} f_{\theta^*}(s, a)]$, in the notation of Assumption 2) for classification, $\mathcal{E}_{\Theta_t, N_{\text{exp}}, \delta/H}^{\text{class}}$, match up to problem dependent constants.

The best known statistical guarantees on imitation gap for BC (Ross et al. [2011]) and MM are,

$$J(\pi^E) - J(\pi^{\text{BC}}) \leq \text{Gap}(\pi^{\text{BC}}) \triangleq H^2 \frac{\sum_{t=1}^H \mathcal{E}_{\Theta_t, N_{\text{exp}}, \delta/H}^{\text{class}}}{H}, \text{ and} \quad (13)$$

$$J(\pi^E) - J(\pi^{\text{MM}}) \leq \text{Gap}(\pi^{\text{MM}}) \triangleq H \sqrt{\frac{\log(F_{\max} H / \delta)}{N_{\text{exp}}}} \quad (14)$$

The guarantee in Theorem 3 for RE therefore grows as $\text{Gap}(\pi^{\text{MM}}) \sqrt{\text{Gap}(\pi^{\text{BC}})/H}$, whenever we are able to establish that $\mathcal{E}_{\Theta_t, N_{\text{exp}}, \delta/H} \asymp \mathcal{E}_{\Theta_t, N_{\text{exp}}, \delta/H}^{\text{class}}$ up to problem dependent constants. Our guarantees for RE in the presence of parametric function approximation thus give us reasons to expect RE to improve the performance of MM. We now turn our attention to validating this in practice.

4 Practical Algorithm

When considering RE (Alg. 1), two main questions arise: (i) How does one construct a membership oracle in practice, especially when action spaces may be continuous?, and (ii) How does one solve the moment matching problem (eq. (5))? We now provide potential answers to both of these questions.

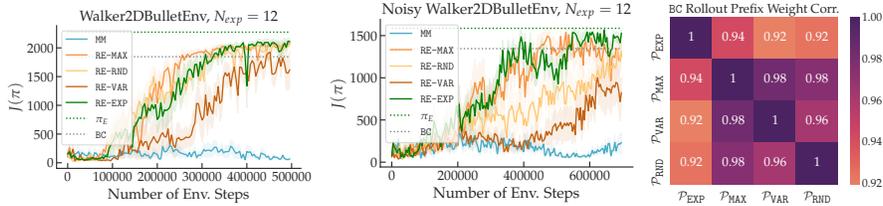


Figure 5: **Left:** All variants of RE are able to nearly match expert performance while MM struggles to make any progress. **Center:** We add i.i.d. noise to the environment to make the control problem more challenging. RE is still able to match expert performance, unlike MM. **Right:** We compute correlations between the idealized prefix weights of \mathcal{M}_{EXP} and the other oracles and see \mathcal{M}_{MAX} correlate most.

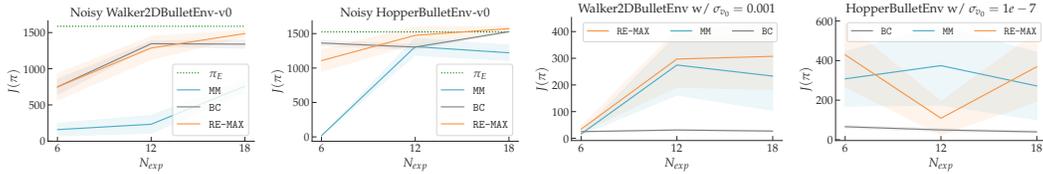


Figure 6: We see RE with \mathcal{M}_{MAX} improve the performance of MM on the Noisy Walker2DBulletEnv and HopperBulletEnv tasks. We see RE (and MM) out-perform BC on the initial-state-perturbed Walker2DBulletEnv and HopperBulletEnv tasks.

Membership Oracle. \mathcal{M} represents how uncertain BC is about the expert’s action. Ideally, we would be able to query the expert policy for an action and weight by the distance between the BC and expert actions, e.g. for sigmoid function σ and constants μ, β ,

$$\mathcal{M}_{\text{EXP}}(s, t) = \sigma \left(\frac{\mu - \|\pi^{\text{BC}}(s) - \pi^E(s)\|_2}{\beta} \right), \quad (15)$$

However, in the non-interactive setting, we can only approximate this quantity. The first approximation we propose using is inspired by Random Network Distillation (RND) [Burda et al., 2018], used by Wang et al. [2019] to estimate the support of the expert policy. We instead propose using RND as a measure of how uncertain BC is about expert actions. That is,

$$\mathcal{M}_{\text{RND}}(s, t) = \sigma \left(\frac{\mu - \|\pi^{\text{BC}}(s) - \widehat{\pi}^{\text{BC}}(s)\|_2}{\beta} \right), \quad (16)$$

where $\widehat{\pi}^{\text{BC}}$ is a network trained to imitate π^{BC} on expert states. To train $\widehat{\pi}^{\text{BC}}$, we evaluate π^{BC} on expert states and plug this new dataset into their standard IL pipeline.

We can also utilize other uncertainty measures like the disagreement of an ensemble, which has shown success on various simulated sequential decision making tasks [Pathak et al., 2019]. Past work by Brantley et al. [2019] utilizes the variance of a set of BC learners as a regularizer on top of standard

BC error. We instead propose using it to *weight* the two halves of our estimator. Instead, we define,

$$\mathcal{M}_{\text{VAR}}(s, t) = \sigma \left(\frac{\mu - \text{Var}(\pi^{\text{BC}(1)}(s), \dots, \pi^{\text{BC}(k)}(s))}{\beta} \right), \quad (17)$$

where $\{\pi^{\text{BC}(1)}, \dots, \pi^{\text{BC}(k)}\}$ are BC policies trained with different initializations, which produces sufficient diversity when using deep networks as function approximators. Lastly, we can also use the maximum difference in the ensemble Kidambi et al. [2021] as our disagreement measure:

$$\mathcal{M}_{\text{MAX}}(s, t) = \sigma \left(\frac{\mu - \max_{i, j \in [k]} \|\pi_i^{\text{BC}}(s) - \pi_j^{\text{BC}}(s)\|_2}{\beta} \right), \quad (18)$$

For computing prefix weights, we use the average of distances up till the current timestep instead of the sum of distances that would follow directly from the idealized algorithm presented in Sec. 3. This modification serves to improve the numerical stability of our method.

Moment Matching. We implement approximate Nash equilibrium computation of (5) by running a no-regret learner against a best-response counterpart [Swamy et al., 2021]. Our approach is related to the GAIL algorithm of Ho and Ermon [2016] which we improve in 4 ways: (i) we use a general Integral Probability Metric Müller [1997] instead of the Jensen-Shannon Divergence used in the original paper which improves the representation power of the the discriminator, (ii) we add in gradient penalties to the discriminator, which improves convergence rates Gulrajani et al. [2017], (iii) we solve the entropy-regularized forward problem via Soft-Actor Critic Haarnoja et al. [2018] as the policy optimizer, for improved sample efficiency, and (iv) we use optimistic mirror descent instead of gradient descent as our optimization algorithm for both players, giving us faster convergence to Nash equilibria, both in theory Syrgkanis et al. [2015] and in practice Daskalakis et al. [2017]. Together, these changes lead to an implementation which *significantly out-performs the original*, giving us a strong baseline to compare against. We include an ablation to confirm this fact in Appendix C. We emphasize that the RE technique can be used to improve *any* online moment matching algorithm and that the above description is merely the approach we chose for this paper.

5 Experimental Results

We now quantify the empirical benefits of RE on several continuous control tasks from the the PyBullet suite Coumans and Bai [2016–2019]. All the task we consider have long horizons ($H \approx 1000$) and we use relatively few demonstrations. ($N_{\text{exp}} \leq 20$). We set N_{replay} as 100 BC rollouts (Line 4 of Algorithm 1). We test all four membership oracles from the previous section (\mathcal{M}_{EXP} as an idealized target, \mathcal{M}_{RND} , \mathcal{M}_{VAR} , and \mathcal{M}_{MAX} as practical solutions). In Fig. 5 (left), we see that with only twelve trajectories, RE is able to reliably match expert performance for all oracles considered, while MM is not. The environment considered in this experiment is nearly deterministic, indicating that RE can help even when the environment is not stochastic. We hypothesize that the randomness in the initial state is sufficient for replay estimation to generate a significant improved estimate of the state-action visitation measure. This improvement is especially interesting considering both of the examples we study in Sec. 3 were heavily stochastic. We see a similar result in Fig. 5 (center), where we add in i.i.d. noise to the environment dynamics at each timestep. This makes the problem significantly more challenging than the standard version of the Walker task. RE is still able to match expert performance, with \mathcal{M}_{MAX} working notably well. The correlation plot in Fig. 5 (right) shows us \mathcal{M}_{MAX} appears to be best correlated with the idealized prefix weights, \mathcal{M}_{EXP} under the state distribution induced by BC. Because of its superior performance, we use \mathcal{M}_{MAX} for the rest of our experiments. In the left half of Fig. 6, we see RE improve the performance of MM. In the right half, we see RE out-perform BC in responding to an *extremely tiny* amount of noise added to the initial velocity of the agent (similar to the experiments of Reddy et al. [2019]) – we defer more details to Appendix C). These results indicate that RE can out-perform MM and BC, agreeing with our theory. We release our code at https://github.com/gkswamy98/replay_est.

6 Acknowledgements

ZSW is supported in part by the NSF FAI Award #1939606, a Google Faculty Research Award, a J.P. Morgan Faculty Award, a Facebook Research Award, an Okawa Foundation Research Grant, and a Mozilla Research Grant. GS is supported by a GPU award from NVIDIA. NR and JJ were partially supported by NSF Grants IIS-1901252, and CCF-1909499. KR is supported by ARO fund 051242-002.

References

- Jean-Yves Audibert and Alexandre B. Tsybakov. Fast learning rates for plug-in classifiers. *The Annals of Statistics*, 35(2):608 – 633, 2007. doi: 10.1214/009053606000001217. URL <https://doi.org/10.1214/009053606000001217>.
- Kianté Brantley, Wen Sun, and Mikael Henaff. Disagreement-regularized imitation learning. In *International Conference on Learning Representations*, 2019.
- Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov. Exploration by random network distillation. *arXiv preprint arXiv:1810.12894*, 2018.
- Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>, 2016–2019.
- Amit Daniely and Shai Shalev-Shwartz. Optimal learners for multiclass problems. *CoRR*, abs/1405.2420, 2014. URL <http://arxiv.org/abs/1405.2420>.
- Amit Daniely, Sivan Sabato, Shai Ben-David, and Shai Shalev-Shwartz. Multiclass learnability and the ERM principle. *CoRR*, abs/1308.2893, 2013. URL <http://arxiv.org/abs/1308.2893>.
- Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training gans with optimism. *arXiv preprint arXiv:1711.00141*, 2017.
- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of wasserstein gans, 2017.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor, 2018.
- Yanjun Han, Jiantao Jiao, and Tsachy Weissman. Minimax estimation of discrete distributions under ℓ_1 loss, 2015.
- Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning, 2016.
- Rahul Kidambi, Aravind Rajeswaran, Praneeth Netrapalli, and Thorsten Joachims. Morel : Model-based offline reinforcement learning, 2021.
- Enno Mammen and Alexandre B. Tsybakov. Smooth discrimination analysis. *The Annals of Statistics*, 27(6):1808–1829, 1999. ISSN 00905364. URL <http://www.jstor.org/stable/2673938>.
- David A. McAllester and Robert E. Schapire. On the convergence rate of good-turing estimators. In *Proceedings of the Thirteenth Annual Conference on Computational Learning Theory, COLT '00*, page 1–6, San Francisco, CA, USA, 2000. Morgan Kaufmann Publishers Inc. ISBN 155860703X.
- Alfred Müller. Integral probability metrics and their generating classes of functions. *Advances in Applied Probability*, 29(2):429–443, 1997.
- Deepak Pathak, Dhiraj Gandhi, and Abhinav Gupta. Self-supervised exploration via disagreement, 2019.
- Dean A Pomerleau. Alvin: An autonomous land vehicle in a neural network. Technical report, CARNEGIE-MELLON UNIV PITTSBURGH PA ARTIFICIAL INTELLIGENCE AND PSYCHOLOGY ..., 1989.
- Antonin Raffin. Rl baselines3 zoo. <https://github.com/DLR-RM/rl-baselines3-zoo>, 2020.
- Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann. Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 22(268):1–8, 2021. URL <http://jmlr.org/papers/v22/20-1364.html>.
- Nived Rajaraman, Lin F Yang, Jiantao Jiao, and Kannan Ramachandran. Toward the fundamental limits of imitation learning. *arXiv preprint arXiv:2009.05990*, 2020.

- Nived Rajaraman, Yanjun Han, Lin Yang, Jingbo Liu, Jiantao Jiao, and Kannan Ramchandran. On the value of interaction and function approximation in imitation learning. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=Xa9Ba6NsJ6>.
- Siddharth Reddy, Anca D. Dragan, and Sergey Levine. Sqil: Imitation learning via reinforcement learning with sparse rewards, 2019.
- Stephane Ross and Drew Bagnell. Efficient reductions for imitation learning. In Yee Whye Teh and Mike Titterton, editors, *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pages 661–668, Chia Laguna Resort, Sardinia, Italy, 13–15 May 2010. PMLR. URL <https://proceedings.mlr.press/v9/ross10a.html>.
- Stephane Ross, Geoffrey J. Gordon, and J. Andrew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning, 2011.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *CoRR*, abs/1707.06347, 2017. URL <http://arxiv.org/abs/1707.06347>.
- Gokul Swamy, Sanjiban Choudhury, J. Andrew Bagnell, and Zhiwei Steven Wu. Of moments and matching: A game-theoretic framework for closing the imitation gap. 2021. URL <https://arxiv.org/abs/2103.03236>.
- Vasilis Syrgkanis, Alekh Agarwal, Haipeng Luo, and Robert E Schapire. Fast convergence of regularized learning in games. *arXiv preprint arXiv:1507.00407*, 2015.
- Ruohan Wang, Carlo Ciliberto, Pierluigi Vito Amadori, and Yiannis Demiris. Random expert distillation: Imitation learning via expert policy support estimation. In *International Conference on Machine Learning*, pages 6536–6544. PMLR, 2019.
- L.A. Zadeh. Fuzzy sets. *Information and Control*, 8(3):338–353, 1965. ISSN 0019-9958. doi: [https://doi.org/10.1016/S0019-9958\(65\)90241-X](https://doi.org/10.1016/S0019-9958(65)90241-X). URL <https://www.sciencedirect.com/science/article/pii/S001999586590241X>.
- Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, Anind K Dey, et al. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, pages 1433–1438. Chicago, IL, USA, 2008.

A Proofs

A.1 Notation

In this appendix, we use the notation $d_t^\pi(\cdot, \cdot)$ to indicate the state-action visitation measure induced by the policy π at time t . We overload the notation $d_t^\pi(\cdot)$ to denote the state-visitation measure induced by the policy π at time t . Likewise, the notations $d_t^D(\cdot, \cdot)$ and $d_t^D(\cdot)$ indicate the empirical visitation measures in the dataset D . For a function $g : \mathcal{X} \rightarrow \mathbb{R}$, the norm $\|g\|_\infty \triangleq \sup_{x \in \mathcal{X}} |g(x)|$.

Before discussing the proofs of the results, we also explain the instantiation of the function class in the tabular setting below.

Remark 3. *In the tabular setting, we instantiate the discriminator class as $\mathcal{F}_t = \{f_t : \|f_t\|_\infty \leq 1\}$ for each t , as the set of all 1-bounded functions, and the policy class Π as the set of all policies. eq. (4) corresponds to finding a policy which best matches the empirical state-action visitation measure observed in the dataset D in total variation (TV) distance (see Appendix A.2 for a proof).*

A.2 Imitation gap upper bound on empirical moment matching (Theorem 3.1)

Below we restate Theorem 3.1 and provide a proof of this result. The key observation is that since the learner π^{MM} best matches the empirical distribution in the dataset, which is in turn close to the population visitation measure induced by π^E , we can expect the visitation measure induced by π^E and π^{MM} to be close. This in turns implies that both policies will collect a similar value under any reward function. Precisely characterizing the rates at which these distributions converge to one another results in the final bound.

Theorem 3.1. *Consider the empirical moment matching learner π^{MM} (eq. (4)), instantiated with an appropriate discriminator class \mathcal{F} . The imitation gap satisfies $\mathbb{E} [J(\pi^E) - J(\pi^{\text{MM}})] \lesssim H \sqrt{|\mathcal{S}|/N_{\text{exp}}}$.*

Proof. Recall that the learner π^{MM} is the solution to the following optimization problem,

$$\arg \min_{\pi} \sup_{f \in \mathcal{F}} \left\{ \mathbb{E}_{\pi} \left[\frac{\sum_{t=1}^H f_t(s_t, a_t)}{H} \right] - \mathbb{E}_D \left[\frac{\sum_{t=1}^H f_t(s_t, a_t)}{H} \right] \right\} \quad (19)$$

Exchanging the summation and maximization operators and recalling from Remark 3 that in the tabular setting, the discriminator class \mathcal{F} is instantiated as the set of all 1-bounded functions $\bigoplus_{t=1}^H \{f_t : \|f_t\|_\infty \leq 1\}$, π^{MM} is a solution to

$$\arg \min_{\pi} \frac{1}{H} \sum_{t=1}^H \left(\sup_{f: \|f\|_\infty \leq 1} \mathbb{E}_{\pi} [f_t(s_t, a_t)] - \mathbb{E}_D [f_t(s_t, a_t)] \right) = \arg \min_{\pi} \frac{1}{H} \sum_{t=1}^H \text{TV} (d_t^\pi, d_t^D) \quad (20)$$

where the equation follows by the variational definition of the total variation distance, and where d_t^π is the state-action visitation measure induced by π^E and d_t^D is the empirical state-action visitation measure in the dataset D . The imitation gap of this policy can be upper bounded by,

$$J(\pi^E) - J(\pi^{\text{MM}}) = \mathbb{E}_{\pi^E} \left[\sum_{t=1}^H r_t(s_t, a_t) \right] - \mathbb{E}_{\pi^{\text{MM}}} \left[\sum_{t=1}^H r_t(s_t, a_t) \right] \quad (21)$$

$$\stackrel{(i)}{\leq} \sum_{t=1}^H \sup_{r_t: \|r_t\|_\infty \leq 1} \left(\mathbb{E}_{\pi^E} [r_t(s_t, a_t)] - \mathbb{E}_{\pi^{\text{MM}}} [r_t(s_t, a_t)] \right) \quad (22)$$

$$\stackrel{(ii)}{=} \sum_{t=1}^H \text{TV} \left(d_t^{\pi^E}(\cdot, \cdot), d_t^{\pi^{\text{MM}}}(\cdot, \cdot) \right) \quad (23)$$

where (i) maximizes over the reward function which is assumed to lie in the interval $[0, 1]$ pointwise. (ii) again follows from the variational definition of total variation distance. This goes to show that in the tabular setting, MM is equivalent to finding the policy which best matches (in TV-distance) the empirical state-action distribution observed in the dataset.

By an application of triangle inequality,

$$J(\pi^E) - J(\pi^{\text{MM}}) \leq \sum_{t=1}^H \text{TV} \left(d_t^{\pi^E}(\cdot, \cdot), d_t^D(\cdot, \cdot) \right) + \text{TV} \left(d_t^D(\cdot, \cdot), d_t^{\pi^{\text{MM}}}(\cdot, \cdot) \right) \quad (24)$$

$$\leq 2 \sum_{t=1}^H \text{TV} \left(d_t^{\pi^E}(\cdot, \cdot), d_t^D(\cdot, \cdot) \right) \quad (25)$$

where (i) follows from eq. (20) which shows that π^{MM} is the policy which best approximates the empirical state-action visitation measure in total variation distance, and therefore $\text{TV} \left(d_t^{\pi^{\text{MM}}}(\cdot, \cdot), d_t^D(\cdot, \cdot) \right) \leq \text{TV} \left(d_t^{\pi^E}(\cdot, \cdot), d_t^D(\cdot, \cdot) \right)$. The final element is to identify the rate of convergence of the empirical visitation measure d_t^D , to the population $d_t^{\pi^E}$ in total variation distance. This result is known from Theorem 1 of Han et al. [2015], which shows that $\mathbb{E} \left[\text{TV} \left(d_t^{\pi^E}(\cdot, \cdot), d_t^D(\cdot, \cdot) \right) \right] \lesssim \sqrt{\frac{|\mathcal{S}|}{N_{\text{exp}}}}$ noting that $d_t^{\pi^E}$ is a distribution with support size $|\mathcal{S}|$ since π^E is deterministic. Putting it together with eq. (25) after taking expectations on both sides gives,

$$J(\pi^E) - J(\pi^{\text{MM}}) \lesssim \sum_{t=1}^H \sqrt{\frac{|\mathcal{S}|}{N_{\text{exp}}}} = H \sqrt{\frac{|\mathcal{S}|}{N_{\text{exp}}}}. \quad (26)$$

This completes the proof of the result. \square

A.3 Lower bounding the Imitation gap of Behavior Cloning

Since BC is an offline algorithm - namely, the learner does not interact with the MDP, any lower bound against offline algorithms applies for behavior cloning as well. The lower bound instance in Rajaraman et al. [2020] is one such example. One state in the MDP is labelled as the “bad” state, b , which is absorbing and offers no reward. The remaining states each have a single “good” action which re-initializes the policy in a particular distribution ρ on the set of good states, and offering a reward of 1. The other actions at these states are “bad” and transition the learner to the bad state b with probability 1.

The key idea in the lower bound is that any offline algorithm does not know (i) which action the expert would have chosen at states unvisited in the dataset, and (ii) which action does what at these states. At best, the learner can correctly guess the good actions at a state with probability $1/|\mathcal{A}| \leq 1/2$. So, at each state unvisited in the dataset, the learner has a constant probability of getting stuck at the bad state in the MDP and collecting no reward then on. On the other hand, the expert would never choose bad actions at states and collects the maximum possible reward. By carefully counting the probability mass on the unvisited states, the expected imitation gap of any offline IL algorithm (such as BC) can be lower bounded by $\Omega(|\mathcal{S}|H^2/N_{\text{exp}})$ on these instances. Thus we have the following theorem,

Theorem 4 (Theorem 6.1 of Rajaraman et al. [2020]). *Consider any learner $\hat{\pi}$ which carries out an offline imitation learning algorithm (such as behavior cloning). Then, there exists an MDP instance such that the expected imitation gap is lower bounded by,*

$$\mathbb{E} [J(\pi^E) - J(\hat{\pi})] \gtrsim \min \left\{ H, \frac{|\mathcal{S}|H^2}{N_{\text{exp}}} \right\}. \quad (27)$$

A.4 Lower bounding the imitation gap of Empirical Moment Matching

In this section, we show that in the tabular setting, empirical moment matching is suboptimal compared for imitation learning in the worst-case. The main result we prove in this section is,

Theorem 3.2. *If $H \geq 4$, there is a tabular IL instance with 2 states and actions on which with constant probability, the empirical moment matching learner (eq. (4)) incurs, $J(\pi^E) - J(\pi^{\text{MM}}) \gtrsim H/\sqrt{N_{\text{exp}}}$.*

Remark 4. *It is known from Rajaraman et al. [2020] that the Mimic-MD algorithm achieves an imitation gap of $\min \left\{ H, \frac{S|H^{3/2}}{N_{\text{exp}}}, H\sqrt{\frac{|\mathcal{S}|}{N_{\text{exp}}}} \right\}$. This is always better than the worst case error bound incurred by TV distribution matching from Theorem 3.2. In fact when $N_{\text{exp}} \gtrsim \sqrt{H}$ the bound $\frac{H^{3/2}}{N_{\text{exp}}}$ is significantly better than $\frac{H}{\sqrt{N_{\text{exp}}}}$ which decays as $1/\sqrt{N_{\text{exp}}}$. This is illustrated in Figure 4.*

First note that the learner π^{MM} carries out empirical moment matching (eq. (4)), with the discriminator class \mathcal{F} as initialized in Remark 3. As shown in eq. (20), the empirical moment matching learner can be redefined as the solution to a distribution matching problem,

$$\arg \min_{\pi} \frac{1}{H} \sum_{t=1}^H \text{TV} (d_t^{\pi}(\cdot, \cdot), d_t^D(\cdot, \cdot)) \quad (28)$$

Consider an MDP instance with 2 states and 2 actions with a non-stationary transition and reward structure as described in Figure fig. 7. State 1 effectively has a single action (i.e. two actions, a_1 and a_2 with both inducing the same next-state distribution and reward). One of the actions at state 2 induces the uniform distribution over next states. The other action deterministically keeps the learner at state 2. The reward function is 0 at $t = 1$, and the action a_2 at state 2 is the only one which offers a reward of 1. The initial state distribution is highly skewed toward the state $s = 1$ and places approximately $1/\sqrt{N_{\text{exp}}}$ mass on $s = 2$ and the remaining on $s = 1$.

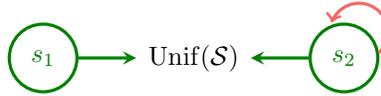


Figure 7: MDP instance which shows that L_1 distribution matching is suboptimal. Here the transition structure is illustrated for $t = 1$. Both states have one action which reinitializes in the uniform distribution. State 2 has an additional action which keeps the state the same. The reward function is 0 for $t = 1$. For $t \geq 2$ the transition function is absorbing at both states; the reward function equals 1 at the state $s = 1$ for any action and is 0 everywhere else.

MDP transition: The state 2 is the only one with two actions. Action a_1 induces the uniform distribution over states, while action a_2 transitions the learner to state 2 with probability 1. Namely,

$$P_1(\cdot | s = 1, a) = \text{Unif}(\mathcal{S}) \text{ for all } a \in \mathcal{A} \quad (29)$$

$$P_1(\cdot | s = 2, a_1) = \text{Unif}(\mathcal{S}) \quad (30)$$

$$P_1(\cdot | s = 2, a_2) = \delta_2 \quad (31)$$

From time $t = 2$ onward, the actions are all absorbing. Namely, for all $t \geq 2$, $s \in \mathcal{S}$ and $a \in \mathcal{A}$,

$$P_t(\cdot | s, a) = \delta_s. \quad (32)$$

Initial state distribution: The initial state distribution $\rho = \left(1 - \frac{1}{\sqrt{N_{\text{exp}}}}, \frac{1}{\sqrt{N_{\text{exp}}}}\right)$.

MDP reward function: The reward function of the MDP encourages the learner to stay at the state $s = 1$ from time $t = 2$ onward. Namely,

$$r_t(s, a) = \begin{cases} 1, & \text{if } t \geq 2 \text{ and } s = 1 \\ 0, & \text{otherwise.} \end{cases} \quad (33)$$

Expert policy: At both states in the MDP, the expert picks the action a_1 to play, which induces the uniform distribution over actions at the next state. Namely, for each $t \in [H]$ and $s \in \mathcal{S}$,

$$\pi_t^E(\cdot | s) = \delta_{a_1} \quad (34)$$

The intuition behind the lower bound is as follows. The only action which affects the value of a policy is the choice made at $s = 2$ at time $t = 1$. At all other states, we may assume that there is effectively only a single action.

By the absorbing nature of states for $t \geq 2$, it turns out that if the observed empirical distribution in the dataset at time 2 is skewed toward state 2 (which is possible because of the inherent randomness in the data generation process), the learner's behavior at time 1 may be to *ignore* the expert's action observed at state $s = 2$, and instead pick the action a_2 which moves the learner to the state $s = 2$ deterministically. The learner is willing to choose a different action because the loss function eq. (28) encourages the state-action distribution at time $t = 2$ also to be well matched with what is observed

in the dataset. Even if it comes at the cost of picking an action different from what the expert plays. By exploiting this fact, we are able to show that the error incurred by a learner which solves eq. (28) in this simple 2 state example must be $\Omega(H/\sqrt{N_{\text{exp}}})$.

Formally, we define 3 events,

- (i) \mathcal{E}_1 : All states in the MDP are visited in the dataset at each time $t = 1, 2, \dots, H$.
- (ii) \mathcal{E}_2 : State 2 is visited at most $\sqrt{N_{\text{exp}}}$ times at time 1 in the dataset D . In other words, $d_1^D(s=2) = \delta'$ where $\delta' \leq \frac{1}{\sqrt{N_{\text{exp}}}}$.
- (iii) \mathcal{E}_3 : At time 2 in the dataset D , the empirical distribution over states is of the form $(\frac{1}{2} - \delta, \frac{1}{2} + \delta)$ for some $\delta \geq \frac{2}{\sqrt{N_{\text{exp}}}}$.

Lemma 1. *Jointly, the events $\mathcal{E}_1, \mathcal{E}_2$ and \mathcal{E}_3 occur with at least constant probability.*

$$\Pr(\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3) \geq C, \quad (35)$$

for some constant $C > 0$.

Proof. By the absorbing nature of states for $t \geq 2$, it suffices for both states of the MDP to be visited in the dataset at time $t = 1, 2$. At time $t = 2$, the marginal state distribution under π^E is the uniform distribution. By binomial concentration, both states are observed in the dataset at time $t = 2$ with probability $\geq 1 - e^{-C_1 N_{\text{exp}}}$ for some constant $C_1 > 0$. On the other hand, at time $t = 1$, the marginal state distribution is $\rho = \left(1 - \frac{1}{\sqrt{N_{\text{exp}}}}, \frac{1}{\sqrt{N_{\text{exp}}}}\right)$. Yet again, by binomial concentration, both states are observed with probability $\geq 1 - e^{-C_2 \sqrt{N_{\text{exp}}}}$ for some constant $C_2 > 0$. By union bounding,

$$\Pr(\mathcal{E}_1) \geq 1 - e^{-C_1 N_{\text{exp}}} - e^{-C_2 \sqrt{N_{\text{exp}}}}. \quad (36)$$

Next we study \mathcal{E}_2 and \mathcal{E}_3 together. First of all, note that the state observed at $t = 1$ and $t = 2$ in a rollout of the expert policy are independent. This is because at both states at $t = 1$, the next state distribution under π^E is uniform. Because of this fact, \mathcal{E}_2 and \mathcal{E}_3 are independent. Next we individually bound the probability of the two events.

\mathcal{E}_2 : The number of times $s = 2$ is the initial state in trajectories the dataset D is distributed as a binomial random variable with distribution $\text{Bin}(N_{\text{exp}}, q)$ with $q = \rho(s=2) = \frac{1}{\sqrt{N_{\text{exp}}}}$. A median of a binomial random variable is $N_{\text{exp}}q = \sqrt{N_{\text{exp}}}$ (in fact any number in the interval $[\lfloor N_{\text{exp}}q \rfloor, \lceil N_{\text{exp}}q \rceil]$ is a median). Therefore, the probability that $s = 2$ is visited $\leq \sqrt{N_{\text{exp}}}$ times in the dataset at time 1 is at least 1/2. In summary,

$$\Pr(\mathcal{E}_2) \geq \frac{1}{2} \quad (37)$$

\mathcal{E}_3 : The marginal distribution over states at time 2 in the dataset is uniform. Therefore, we expect the states 1 and 2 to be visited roughly $N_{\text{exp}}/2$ times each in the dataset, but with a random variation of $\approx \sqrt{N_{\text{exp}}}$ around this average. In other words, the empirical distribution fluctuates as $(\frac{1}{2} - \delta, \frac{1}{2} + \delta)$ with $\delta \geq \frac{2}{\sqrt{N_{\text{exp}}}}$ with constant probability.

By the independence of \mathcal{E}_2 and \mathcal{E}_3 and union bounding to account for \mathcal{E}_1 , the statement of the lemma follows. \square

Lemma 2. *For each $t \geq 2$,*

$$\text{TV}(d_t^\pi(\cdot, \cdot), d_t^D(\cdot, \cdot)) \geq \text{TV}(d_2^\pi(\cdot), d_2^D(\cdot)). \quad (38)$$

The RHS is the TV distance between the state-visitation measure at time $t = 2$ under π and that empirically observed in the dataset D . Conditioned on the events $\mathcal{E}_1, \mathcal{E}_2$ and \mathcal{E}_3 occurring, equality is met in eq. (38) if and only if $\pi_t(\cdot|s) = \pi_t^E(\cdot|s)$ for all states $s \in \mathcal{S}$.

Proof. For any state $s \in \mathcal{S}$ and $t \geq 2$, observe that,

$$\sum_{a \in \mathcal{A}} |d_t^\pi(s, a) - d_t^D(s, a)| \quad (39)$$

$$= d_t^\pi(s)(1 - \pi_t(a^*|s)) + |d_t^\pi(s)\pi_t(a^*|s) - d_t^D(s, a^*)|, \text{ where } a^* = \pi_t^E(s), \quad (40)$$

$$\stackrel{(i)}{=} d_2^\pi(s)(1 - \pi_t(a^*|s)) + |d_2^\pi(s)\pi_t(a^*|s) - d_2^D(s, a^*)| \quad (41)$$

$$\stackrel{(ii)}{\geq} |d_2^\pi(s) - d_2^D(s)|, \quad (42)$$

where (i) follows by the fact that the states of the MDP are absorbing under π for $t \geq 2$. (ii) follows by triangle inequality and using the fact that π^E is deterministic, so $d_t^D(s, a^*) = d_t^D(s)$. Equality is met only if $\pi_t(a^*|s) = 1$ (since $d_t^D(s, a^*) > 0$ conditioned on \mathcal{E}_1). \square

The above lemma asserts the behavior of π^{MM} in eq. (28) for $t \geq 2$. Namely, conditioned on the event \mathcal{E}_1 which happens with very high probability, all states are visited in the MDP and therefore, $\pi_t^{\text{MM}}(\cdot|s) = \pi_t^E(\cdot|s)$ for each state $s \in \mathcal{S}$ and time $t \geq 2$.

The only thing left to study is the MM learner's behavior at $t = 1$. We wish to show that with constant probability, the learner may choose to deviate from the expert policy in order to better match empirical state-action visitation measures. Conditioned on \mathcal{E}_1 , the learner's policy at time $t = 1$ can be computed by solving the following optimization problem,

$$\text{TV}(d_1^\pi(\cdot, \cdot), d_1^D(\cdot, \cdot)) + (H - 1)\text{TV}(d_2^\pi(\cdot), d_2^D(\cdot)). \quad (43)$$

This follows directly by simplifying the learner's objective using Lemma 2.

Now, conditioned on the event \mathcal{E}_1 , at time $t = 1$, the learner policy only needs to be optimized at the state $s = 2$. At the state $s = 1$, we may assume that the learner picks the expert's action $\pi_1^E(s = 1)$. To this end, suppose the learner picks the action a_1 with probability p and the action a_2 with probability $1 - p$.

$$\text{TV}(d_1^\pi(\cdot, \cdot), d_1^D(\cdot, \cdot)) = \sum_{a \in \mathcal{A}} |d_1^{\pi^E}(s = 2, a) - d_1^D(s = 2, a)| \quad (44)$$

$$= |\rho(2)p - \delta'| + |\rho(2)(1 - p) - 0| \quad (45)$$

$$= \left| \frac{p}{\sqrt{N_{\text{exp}}}} - \delta' \right| + \frac{1 - p}{\sqrt{N_{\text{exp}}}}. \quad (46)$$

which follows by plugging in $\rho(2) = \frac{1}{\sqrt{N_{\text{exp}}}}$. And,

$$\text{TV}(d_2^\pi(\cdot), d_2^D(\cdot)) = \left| \left(\frac{1}{2} - \delta \right) - \frac{\rho(1)}{2} - \rho(2) \frac{p}{2} \right| + \left| \left(\frac{1}{2} + \delta \right) - \frac{\rho(1)}{2} - \rho(2) \left((1 - p) + \frac{p}{2} \right) \right|. \quad (47)$$

Plugging in $\rho(2) = \frac{1}{\sqrt{N_{\text{exp}}}}$ and $\rho(1) = 1 - \frac{1}{\sqrt{N_{\text{exp}}}}$, we get,

$$\text{TV}(d_2^\pi(\cdot), d_2^D(\cdot)) = \left| \frac{1}{2\sqrt{N_{\text{exp}}}} - \delta - \frac{p}{2\sqrt{N_{\text{exp}}}} \right| + \left| \frac{p}{2\sqrt{N_{\text{exp}}}} - \frac{1}{2\sqrt{N_{\text{exp}}}} + \delta \right|. \quad (48)$$

Summing up eqs. (46) and (48), p minimizes,

$$\underbrace{\left| \frac{p}{\sqrt{N_{\text{exp}}}} - \delta' \right| + \frac{1 - p}{\sqrt{N_{\text{exp}}}}}_{(i)} + \underbrace{(H - 1) \left(\left| \frac{p}{2\sqrt{N_{\text{exp}}}} + \delta - \frac{1}{2\sqrt{N_{\text{exp}}}} \right| + \left| \frac{1}{2\sqrt{N_{\text{exp}}}} - \delta - \frac{p}{2\sqrt{N_{\text{exp}}}} \right| \right)}_{(ii)}. \quad (49)$$

Intuitively, term (i) captures the error incurred by the learner in the loss eq. (28) by deviating from π^E at the first time step. Term (ii) captures the decrease in the error at every subsequent time step because of the same deviation, since the learner is able to better match the state distribution at future time steps. In the next lemma we show that under events that hold with at least constant probability, the empirical moment matching learner chooses to play the wrong action at time $t = 1$ at the state $s = 2$.

Lemma 3. *Conditioned on the events \mathcal{E}_2 and \mathcal{E}_3 , for $H \geq 4$, the unique minimizer of eq. (49) for $p \in [0, 1]$ is $p = 0$.*

Proof. The first term of eq. (49) is $\left| \frac{p}{\sqrt{N_{\text{exp}}}} - \delta' \right| + \frac{1-p}{\sqrt{N_{\text{exp}}}}$, the error from not picking the expert's action at state 1 at time 1 decreases at most linearly with a slope of $\frac{2}{\sqrt{N_{\text{exp}}}}$.

Conditioned on the event \mathcal{E}_3 , $\delta \geq \frac{2}{\sqrt{N_{\text{exp}}}}$. Therefore, $\left| \frac{p}{2\sqrt{N_{\text{exp}}}} + \delta - \frac{1}{2\sqrt{N_{\text{exp}}}} \right| = \frac{p}{2\sqrt{N_{\text{exp}}}} + \delta - \frac{1}{2\sqrt{N_{\text{exp}}}}$. Therefore, the decrease in error at future steps by deviating from π^E at the time $t = 1$, term (ii) in eq. (49) is,

$$2(H-1) \left(\frac{p}{2\sqrt{N_{\text{exp}}}} + \delta - \frac{1}{2\sqrt{N_{\text{exp}}}} \right) \quad (50)$$

which is an increasing function of p with slope $\frac{H-1}{\sqrt{N_{\text{exp}}}}$. For $H \geq 4$ and the argument from the previous paragraph, this implies that term (ii) increases more rapidly in p than the rate at which term (i) decreases. Therefore, the minimizer must be $p = 0$. \square

Thus, we conclude from Lemmas 2 and 3 that conditioned on the events $\mathcal{E}_1, \mathcal{E}_2$ and \mathcal{E}_3 , the learner π^{MM} perfectly mimics π^E at each time $t \geq 2$, but deviates from the action played by π^E at the state $s = 1$ at time $t = 1$.

Finally, we bound the difference in value between π^E and π^{MM} induced because of this deviation under the reward eq. (33).

Lemma 4. *Under the events $\mathcal{E}_1, \mathcal{E}_2$ and \mathcal{E}_3 , under the reward eq. (33), the empirical moment matching learner π^{MM} incurs imitation gap,*

$$J(\pi^E) - J(\pi^{\text{MM}}) = \frac{H}{2\sqrt{N_{\text{exp}}}}. \quad (51)$$

Proof. Recall that under the events $\mathcal{E}_1, \mathcal{E}_2$ and \mathcal{E}_3 , the learner π^{MM} is identical to π^E except at the state $s = 2$ where they perfectly deviate from each other. The state distribution induced by π^E at each time $t \geq 2$ is the uniform distribution over states $(\frac{1}{2}, \frac{1}{2})$. On the other hand, for $t \geq 2$, the state distribution induced by π^{MM} at each time $t \geq 2$ is $(\rho(1)\frac{1}{2}, \rho(1)\frac{1}{2} + \rho(2)) = \left(\frac{1-1/\sqrt{N_{\text{exp}}}}{2}, \frac{1+1/\sqrt{N_{\text{exp}}}}{2} \right)$. Since the reward function is 1 on state 1, the difference in value between the expert and learner policy is,

$$J(\pi^E) - J(\pi^{\text{MM}}) = \frac{H}{2} - H \left(\frac{1-1/\sqrt{N_{\text{exp}}}}{2} \right) = \frac{H}{2\sqrt{N_{\text{exp}}}}. \quad (52)$$

This completes the proof. \square

Since $\mathcal{E}_1, \mathcal{E}_2$ and \mathcal{E}_3 jointly occur with constant probability by Lemma 1, this completes the proof of Theorem 3.2.

B Imitation gap of RE : Proof of Theorem 2

In this section, we discuss a proof of a more general version of Theorem 2, where N_{replay} can be finite. We prove the following result,

Theorem 5. *Consider the policy π^{RE} returned by Algorithm 1. Assume that $\pi^E \in \Pi$ and the ground truth reward function $r_t \in \mathcal{F}_t$, which is assumed to be symmetric ($f_t \in \mathcal{F}_t \iff -f_t \in \mathcal{F}_t$) and bounded (For all $f_t \in \mathcal{F}_t$, $\|f_t\|_\infty \leq 1$). Choose $|D_1|, |D_2| = \Theta(N_{\text{exp}})$ and suppose $N_{\text{replay}} \rightarrow \infty$. With probability $\geq 1 - 3\delta$,*

$$J(\pi^E) - J(\pi^{\text{RE}}) \lesssim \mathcal{L}_1 + \mathcal{L}_2 + \mathcal{L}_3 + \frac{\log(F_{\max}H/\delta)}{N_{\text{exp}}} \quad (53)$$

where $F_{\max} \triangleq \max_{t \in [H]} |\mathcal{F}_t|$, and,

$$\mathcal{L}_1 \triangleq H^2 \mathbb{E}_{\pi^E} \left[\frac{\sum_{t=1}^H \mathcal{M}(s_t, t) \text{TV}(\pi_t^E(\cdot | s_t), \pi_t^{BC}(\cdot | s_t))}{H} \right], \quad (54)$$

$$\mathcal{L}_2 \triangleq H^{3/2} \sqrt{\frac{\log(F_{\max} H / \delta)}{N_{\text{exp}}} \frac{\sum_{t=1}^H \mathbb{E}_{\pi^E} [1 - \mathcal{M}(s_t, t)]}{H}}, \quad (55)$$

And,

$$\mathcal{L}_3 \triangleq H \sqrt{\frac{\log(F_{\max} H / \delta)}{N_{\text{replay}}}} + \frac{H \log(F_{\max} H / \delta)}{N_{\text{replay}}}. \quad (56)$$

Recall that the learner carrying out replay estimation returns the policy which minimizes the loss $\sup_{f \in \mathcal{F}} J_f(\pi) - \widehat{E}(f)$ over policies π , where $J_f(\pi) \triangleq \mathbb{E}_{\pi} \left[\frac{1}{H} \sum_{t=1}^H f_t(s_t, a_t) \right]$ and abbreviating the notation $f = (f_1, \dots, f_H)$. Note that,

$$J(\pi^E) - J(\pi^{\text{RE}}) \stackrel{(i)}{\leq} \sup_{f \in \mathcal{F}} J_f(\pi^E) - J_f(\pi^{\text{RE}}) \quad (57)$$

$$\leq \sup_{f \in \mathcal{F}} \left| J_f(\pi^E) - \widehat{E}(f) \right| + \sup_{f \in \mathcal{F}} \left| \widehat{E}(f) - J_f(\pi^{\text{RE}}) \right| \quad (58)$$

$$\stackrel{(ii)}{\leq} 2 \sup_{f \in \mathcal{F}} \left| J_f(\pi^E) - \widehat{E}(f) \right|. \quad (59)$$

where (i) uses the realizability assumption that the ground truth reward lies in \mathcal{F} , and (ii) uses the fact that π^{RE} is a minimizer of eq. (5) and the fact that \mathcal{F} is symmetric (this implies that $\sup_{f \in \mathcal{F}} J_f(\pi^E) - \widehat{E}(f) = \sup_{f \in \mathcal{F}} -J_f(\pi^E) + \widehat{E}(f) = \sup_{f \in \mathcal{F}} \left| J_f(\pi^E) - \widehat{E}(f) \right|$).

Note that $\widehat{E}(f)$ can be decomposed into a sum of two parts,

$$\widehat{E}^{(1)}(f) = \mathbb{E}_{D_{\text{replay}}} \left[\frac{1}{H} \sum_{t=1}^H f_t(s_t, a_t) (1 - \mathcal{P}(s_{1..t-1})) \right], \text{ and,} \quad (60)$$

$$\widehat{E}^{(2)}(f) = \mathbb{E}_{D_2} \left[\frac{1}{H} \sum_{t=1}^H f_t(s_t, a_t) (1 - \mathcal{P}(s_{1..t-1})) \right] \quad (61)$$

Likewise, we can decompose $J_f(\pi^E)$ into two terms,

$$J_f^{(1)}(\pi^E) \triangleq \mathbb{E}_{\pi^E} \left[\sum_{t=1}^H f_t(s_t, a_t) \mathcal{P}(s_{1..t-1}) \right], \text{ and} \quad (62)$$

$$J_f^{(2)}(\pi^E) \triangleq \mathbb{E}_{\pi^E} \left[\sum_{t=1}^H f_t(s_t, a_t) (1 - \mathcal{P}(s_{1..t-1})) \right] \quad (63)$$

Then, from eq. (59),

$$J(\pi^E) - J(\pi^{\text{RE}}) \leq 2 \sup_{f \in \mathcal{F}} \left| J_f(\pi^E) - \widehat{E}(f) \right| \quad (64)$$

$$\begin{aligned} &\leq 2 \underbrace{\sup_{f \in \mathcal{F}} \left| J_f^{(1)}(\pi^E) - \mathbb{E} \left[\widehat{E}^{(1)}(f) \mid D_1 \right] \right|}_{\text{(I)}} + 2 \underbrace{\sup_{f \in \mathcal{F}} \left| \mathbb{E} \left[\widehat{E}^{(1)}(f) \mid D_1 \right] - \widehat{E}^{(1)}(f) \right|}_{\text{(II)}} \\ &\quad + 2 \underbrace{\sup_{f \in \mathcal{F}} \left| J_f^{(2)}(\pi^E) - \widehat{E}^{(2)}(f) \right|}_{\text{(III)}}. \end{aligned} \quad (65)$$

where the last line follows by triangle inequality. We bound each of these terms in the next 3 lemmas, starting with (I).

Lemma 5.

$$\sup_{f \in \mathcal{F}} \left| J_f^{(1)}(\pi^E) - \mathbb{E} \left[\widehat{E}^{(1)}(f) \middle| D_1 \right] \right| \leq H \sum_{h=1}^H \mathbb{E}_{\pi^E} \left[\mathcal{M}(s_h, h) \text{TV}(\pi_h^E(\cdot|s_h), \pi_h^{\text{BC}}(\cdot|s_h)) \right] \quad (66)$$

Proof. The proof of this result closely follows the supervised learning reduction of BC (cf. Ross and Bagnell [2010]). Note that,

$$\mathbb{E} \left[\widehat{E}^{(1)}(f) \middle| D_1 \right] - V_f^{(1)}(\pi^E) = \sum_{t=1}^H \mathbb{E}_{\pi^{\text{BC}}} [f_t(s_t, a_t) \mathcal{P}(s_{1 \dots t-1})] - \mathbb{E}_{\pi^E} [f_t(s_t, a_t) \mathcal{P}(s_{1 \dots t-1})]. \quad (67)$$

Define $\pi^{(h)}$ as the policy which plays π^E until (and including) time h and π^{BC} after time h . Then, by cascading,

$$\begin{aligned} & \mathbb{E}_{\pi^{\text{BC}}} [f_t(s_t, a_t) \mathcal{P}(s_{1 \dots t-1})] - \mathbb{E}_{\pi^E} [f_t(s_t, a_t) \mathcal{P}(s_{1 \dots t-1})] \\ &= \sum_{h=0}^{t-1} \mathbb{E}_{\pi^{(h)}} [f_t(s_t, a_t) \mathcal{P}(s_{1 \dots t-1})] - \mathbb{E}_{\pi^{(h+1)}} [f_t(s_t, a_t) \mathcal{P}(s_{1 \dots t-1})] \end{aligned} \quad (68)$$

Define, the uncertainty weighted state visitation measure $d^{\mathcal{M}}$ and the uncertainty weighted look-forward reward $\rho^{\mathcal{M}}$ as follows,

$$d_{h+1}^{\mathcal{M}}(s') \triangleq \mathbb{E}_{\pi^E} \left[\mathbb{1}(s_{h+1} = s') \prod_{t'=1}^h \mathcal{M}(s_{t'}, t') \right] \quad (69)$$

$$\rho_{h+1}^{\mathcal{M}}(s', a') \triangleq \mathbb{E}_{\pi^{\text{BC}}} \left[f_t(s_t, a_t) \prod_{t'=h+2}^t \mathcal{M}(s_{t'}, t') \middle| s_{h+1} = s', a_{h+1} = a' \right] \quad (70)$$

By decomposing expectations along trajectories, using the fact that $\mathcal{P}(s_{1 \dots t-1}) = \prod_{t'=1}^H \mathcal{M}(s_{t'}, t')$ some simplification results in the following equation,

$$|\mathbb{E}_{\pi^{(h)}} [f_t(s_t, a_t) \mathcal{P}(s_{1 \dots t-1})] - \mathbb{E}_{\pi^{(h+1)}} [f_t(s_t, a_t) \mathcal{P}(s_{1 \dots t-1})]| \quad (71)$$

$$= \left| \sum_{s' \in \mathcal{S}} \sum_{a' \in \mathcal{A}} d_{h+1}^{\mathcal{M}}(s') \mathcal{M}(s', h+1) (\pi_{h+1}^E(a'|s') - \pi_{h+1}^{\text{BC}}(a'|s')) \rho_{h+1}^{\mathcal{M}}(s', a') \right| \quad (72)$$

$$\stackrel{(i)}{\leq} \sum_{s' \in \mathcal{S}} d_{h+1}^{\mathcal{M}}(s') \mathcal{M}(s', h+1) \text{TV}(\pi_{h+1}^E(\cdot|s'), \pi_{h+1}^{\text{BC}}(\cdot|s')) \quad (73)$$

$$= \mathbb{E}_{\pi^E} [\mathcal{M}(s_{h+1}, h+1) \text{TV}(\pi_{h+1}^E(\cdot|s_{h+1}), \pi_{h+1}^{\text{BC}}(\cdot|s_{h+1}))]. \quad (74)$$

where (i) uses the fact that the membership oracle is a function $\in [0, 1]$ and f is bounded and lies in the interval $[0, 1]$ (which implies that $\rho^{\mathcal{M}}$ also lies in $[0, 1]$ pointwise). Plugging into eq. (68) and subsequently into eq. (67) completes the proof. \square

Next we bound the 3rd term, (III). This follows by an application of Bernstein's inequality.

Lemma 6. *With probability $\geq 1 - \delta$,*

$$\sup_{f \in \mathcal{F}} \left| J_f^{(2)}(\pi^E) - \widehat{E}^{(2)}(f) \right| \leq H \sqrt{\frac{\log(F_{\max} H / \delta) \sum_{t=1}^{H-1} \mathbb{E}_{\pi^E} [1 - \mathcal{M}(s_t, t)]}{N_{\text{exp}}}} + \frac{H \log(F_{\max} H / \delta)}{N_{\text{exp}}} \quad (75)$$

Proof. First observe that,

$$J_f^{(2)}(\pi^E) - \widehat{V}_f^{(2)} = \sum_{t=1}^H \mathbb{E}_{\text{tr} \sim \text{Unif}(D_2)} [f_t(s_t, a_t) (1 - \mathcal{P}(s_{1 \dots t-1}))] - \mathbb{E}_{\pi^E} [f_t(s_t, a_t) (1 - \mathcal{P}(s_{1 \dots t-1}))] \quad (76)$$

For each t , note that $f_t(s_t, a_t) (1 - \mathcal{P}(s_{1\dots t-1}))$ is bounded in the range $[0, 1]$. Therefore, invoking Bernstein's inequality, with probability $\geq 1 - \delta$,

$$\left| \mathbb{E}_{\text{Unif}(D_2)} [f_t(s_t, a_t) (1 - \mathcal{P}(s_{1\dots t-1}))] - \mathbb{E}_{\pi^E} [f_t(s_t, a_t) (1 - \mathcal{P}(s_{1\dots t-1}))] \right| \quad (77)$$

$$\lesssim \sqrt{\frac{\text{Var}_{\pi^E} (f_t(s_t, a_t) (1 - \mathcal{P}(s_{1\dots t-1}))) \log(1/\delta)}{N_{\text{exp}}}} + \frac{\log(1/\delta)}{N_{\text{exp}}} \quad (78)$$

$$\leq \sqrt{\frac{\mathbb{E}_{\pi^E} [(f_t(s_t, a_t) (1 - \mathcal{P}(s_{1\dots t-1})))^2] \log(1/\delta)}{N_{\text{exp}}}} + \frac{\log(1/\delta)}{N_{\text{exp}}} \quad (79)$$

$$\stackrel{(i)}{\leq} \sqrt{\frac{\mathbb{E}_{\pi^E} [f_t(s_t, a_t) (1 - \mathcal{P}(s_{1\dots t-1}))] \log(1/\delta)}{N_{\text{exp}}}} + \frac{\log(1/\delta)}{N_{\text{exp}}} \quad (80)$$

$$\stackrel{(ii)}{\leq} \sqrt{\frac{\mathbb{E}_{\pi^E} [1 - \mathcal{P}(s_{1\dots t-1})] \log(1/\delta)}{N_{\text{exp}}}} + \frac{\log(1/\delta)}{N_{\text{exp}}} \quad (81)$$

where (i) uses the fact that $f_t(s_t, a_t) (1 - \mathcal{P}(s_{1\dots t-1}))$ is bounded in the range $[0, 1]$, and (ii) uses the fact that $0 \leq f_t(s_t, a_t) \leq 1$. Assuming $0 \leq x_i \leq 1$ for all $i \in [n]$, we have the inequality,

$$1 - \prod_{i=1}^n x_i \leq \sum_{i=1}^n 1 - x_i \quad (82)$$

Applying this to eq. (81) for $1 - \mathcal{P}(s_{1\dots t-1}) = 1 - \prod_{t'=1}^{t-1} \mathcal{M}(s_{t'}, t')$, we have,

$$\left| \mathbb{E}_{\text{Unif}(D_2)} [f_t(s_t, a_t) (1 - \mathcal{P}(s_{1\dots t-1}))] - \mathbb{E}_{\pi^E} [f_t(s_t, a_t) (1 - \mathcal{P}(s_{1\dots t-1}))] \right| \quad (83)$$

$$\leq \sqrt{\frac{\sum_{t'=1}^{t-1} \mathbb{E}_{\pi^E} [1 - \mathcal{M}(s_{t'}, t')] \log(1/\delta)}{N_{\text{exp}}}} + \frac{\log(1/\delta)}{N_{\text{exp}}} \quad (84)$$

Therefore, by union bounding, with probability $\geq 1 - \delta/H$, simultaneously for every $f_t \in \mathcal{F}_t$,

$$\left| \mathbb{E}_{\text{Unif}(D_2)} [f_t(s_t, a_t) (1 - \mathcal{P}(s_{1\dots t-1}))] - \mathbb{E}_{\pi^E} [f_t(s_t, a_t) (1 - \mathcal{P}(s_{1\dots t-1}))] \right| \quad (85)$$

$$\lesssim \sqrt{\frac{\log(|\mathcal{F}_t|H/\delta) \sum_{t'=1}^{t-1} \mathbb{E}_{\pi^E} [1 - \mathcal{M}(s_{t'}, t')]}{N_{\text{exp}}}} + \frac{\log(|\mathcal{F}_t|H/\delta)}{N_{\text{exp}}}. \quad (86)$$

This implies that the maximum over f_t of the LHS is upper bounded by the RHS. Union bounding over $t = 1, \dots, H$ and plugging into eq. (76), we have that with probability $\geq 1 - \delta$,

$$\sup_{f \in \mathcal{F}} \left| J_f^{(2)}(\pi^E) - \widehat{V}_f^{(2)} \right| \quad (87)$$

$$\leq \sum_{t=1}^H \left| \mathbb{E}_{\text{Unif}(D_2)} [f_t(s_t, a_t) (1 - \mathcal{P}(s_{1\dots t-1}))] - \mathbb{E}_{\pi^E} [f_t(s_t, a_t) (1 - \mathcal{P}(s_{1\dots t-1}))] \right| \quad (88)$$

$$\lesssim H \sqrt{\frac{\log(F_{\max}H/\delta) \sum_{t=1}^{H-1} \mathbb{E}_{\pi^E} [1 - \mathcal{M}(s_t, t)]}{N_{\text{exp}}}} + \frac{H \log(F_{\max}H/\delta)}{N_{\text{exp}}}. \quad (89)$$

□

Lemma 7. With probability $\geq 1 - \delta$,

$$\sup_{f \in \mathcal{F}} \left| \mathbb{E} \left[\widehat{V}_f^{(1)} \middle| D_1 \right] - \widehat{V}_f^{(1)} \right| \lesssim H \sqrt{\frac{\log(F_{\max}H/\delta)}{N_{\text{replay}}}} + \frac{H \log(F_{\max}H/\delta)}{N_{\text{replay}}}. \quad (90)$$

Proof. The proof follows essentially the same structure as Lemma 6 by decomposing $\widehat{V}_f^{(1)}$ into a sum of H terms of the form $f_t(s_t, a_t) \mathcal{P}(s_{1\dots t-1})$, applying Bernstein's inequality to bound the deviation of each term from its mean and finally union bounding over the rewards $f_t \in \mathcal{F}_t$ to get the uniform bound over all discriminators $f \in \mathcal{F}$. □

Putting together Lemmas 5 to 7 with eq. (65) completes the proof of Theorem 2.

B.1 Bound in the tabular setting (Theorem 6)

In this section, we provide an upper bound on the imitation gap of RE in the tabular setting when the expert is a deterministic policy. This recovers the bound on the imitation gap for RE proved in Rajaraman et al. [2020].

Theorem 6. *Consider an appropriately initialized version of RE, and let the size of the replay dataset $N_{\text{replay}} \rightarrow \infty$. For any tabular IL instance with $H \geq 10$, with probability $\geq 1 - 3\delta$,*

$$J(\pi^E) - J(\pi^{\text{RE}}) \lesssim \min \left\{ \frac{|\mathcal{S}|H^{3/2}}{N_{\text{exp}}}, H \sqrt{\frac{|\mathcal{S}|}{N_{\text{exp}}}} \right\} \log \left(\frac{|\mathcal{S}|H}{\delta} \right). \quad (91)$$

Below we describe the implementation of RE corresponding to Theorem 6 in more detail.

The membership oracle we use in this setting for RE is defined below,

$$\mathcal{M}(s, t) = \begin{cases} 1 & \text{if } s \text{ is visited in } D_1 \text{ at time } t \\ 0 & \text{otherwise.} \end{cases} \quad (92)$$

The function class \mathcal{F} which we use is identical to that for empirical moment matching, which is described in Remark 3.

Note that in the tabular setting, BC simply mimics the deterministic expert's actions at states visited in the dataset D_1 and plays an arbitrary deterministic action on the remaining states. As a consequence of this definition, if $\mathcal{M}(s, t) = 1 \iff \pi_t^{\text{BC}}(\cdot|s) = \pi_t^E(\cdot|s)$ and $\mathcal{M}(s, t) = 0$ otherwise. We instantiate the family of discriminators as in Remark 3, as $\mathcal{F} = \bigoplus_{t=1}^H \{f_t : \|f_t\|_\infty \leq 1\}$ and the set of policies Π optimized over is chosen as the set of all deterministic policies. While the guarantee of Theorem 2 depends on $F_{\text{max}} = \max_{t \in [H]} |\mathcal{F}_t|$ which is unbounded (or $\exp(|\mathcal{S}||\mathcal{A}|)$) by using a discretization of the reward space), note that we can improve the guarantee to effectively have $F_{\text{max}} \approx \exp(|\mathcal{S}|)$ noting the structure of the set of discriminators. Looking into the proof of Theorem 2 we bring out this dependence below. We note that there are many ways of bringing out this dependence, including a careful net argument directly on top of the guarantee of Theorem 2. We simply present one such argument below.

The critical step where the finiteness of the set of discriminators \mathcal{F} is used, is in union bounding the gap between the population and the empirical estimate of $f_t(s_t, a_t) (1 - \mathcal{P}(s_{1..t-1}))$ in eq. (84).

$$\left| \mathbb{E}_{\text{Unif}(D_2)} [f_t(s_t, a_t) (1 - \mathcal{P}(s_{1..t-1}))] - \mathbb{E}_{\pi^E} [f_t(s_t, a_t) (1 - \mathcal{P}(s_{1..t-1}))] \right| \quad (93)$$

In the next step of the proof of Theorem 2, we union bound over all $f_t \in \mathcal{F}_t$. However, note that for $\mathcal{F}_t = \{f_t : \|f_t\|_\infty \leq 1\}$, we have that,

$$\sup_{f_t : \|f_t\|_\infty \leq 1} \left| \mathbb{E}_{\text{Unif}(D_2)} [f_t(s_t, a_t) (1 - \mathcal{P}(s_{1..t-1}))] - \mathbb{E}_{\pi^E} [f_t(s_t, a_t) (1 - \mathcal{P}(s_{1..t-1}))] \right| \quad (94)$$

$$\stackrel{(i)}{=} \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \left| \mathbb{E}_{\text{Unif}(D_2)} [\mathbb{1}(s_t = s, a_t = a) (1 - \mathcal{P}(s_{1..t-1}))] - \mathbb{E}_{\pi^E} [\mathbb{1}(s_t = s, a_t = a) (1 - \mathcal{P}(s_{1..t-1}))] \right| \quad (95)$$

$$\stackrel{(ii)}{=} \sum_{s \in \mathcal{S}} \left| \mathbb{E}_{\text{Unif}(D_2)} [\mathbb{1}(s_t = s) (1 - \mathcal{P}(s_{1..t-1}))] - \mathbb{E}_{\pi^E} [\mathbb{1}(s_t = s) (1 - \mathcal{P}(s_{1..t-1}))] \right| \quad (96)$$

$$\leq \sum_{s \in \mathcal{S}} \left| \mathbb{E}_{\text{Unif}(D_2)} [\mathbb{1}(s_t = s) (1 - \mathcal{P}(s_{1..t-1}))] - \mathbb{E}_{\pi^E} [\mathbb{1}(s_t = s) (1 - \mathcal{P}(s_{1..t-1}))] \right| \quad (97)$$

where (i) follows similar to the equivalence between the variational representation of TV distance ($\text{TV}(P, Q) = \frac{1}{2} \sup_{f: \|f\|_\infty \leq 1} \mathbb{E}_P[f] - \mathbb{E}_Q[f]$) and the relationship to the L_1 distance, $\text{TV}(P, Q) = \frac{1}{2} L_1(P, Q)$. On the other hand, (ii) follows by noting that the expert is a deterministic policy (and D_2 is generated by rolling out π^E). (iii) follows by an application of Holder's inequality. By subgaussian

concentration, for each $s \in \mathcal{S}$, with probability $\geq 1 - \frac{\delta}{|\mathcal{S}|H}$,

$$\left| \mathbb{E}_{\text{Unif}(D_2)} [\mathbb{1}(s_t = s) (1 - \mathcal{P}(s_{1..t-1}))] - \mathbb{E}_{\pi^E} [\mathbb{1}(s_t = s) (1 - \mathcal{P}(s_{1..t-1}))] \right| \quad (98)$$

$$\lesssim \sqrt{\frac{\text{Var}_{\pi^E} (\mathbb{1}(s_t = s) (1 - \mathcal{P}(s_{1..t-1}))) \log\left(\frac{|\mathcal{S}|H}{\delta}\right) + \frac{\log\left(\frac{|\mathcal{S}|H}{\delta}\right)}{|D_2|}}{|D_2|}} \quad (99)$$

$$\stackrel{(i)}{\leq} \sqrt{\frac{\mathbb{E}_{\pi^E} [\mathbb{1}(s_t = s) (1 - \mathcal{P}(s_{1..t-1}))] \log\left(\frac{|\mathcal{S}|H}{\delta}\right) + \frac{\log\left(\frac{|\mathcal{S}|H}{\delta}\right)}{|D_2|}}{|D_2|}} \quad (100)$$

where (i) uses the fact that $0 \leq \mathbb{1}(s_t = s) (1 - \mathcal{P}(s_{1..t-1})) \leq 1$. Combining with eq. (97), union bounding and applying Cauchy Schwarz inequality, with probability $\geq 1 - \frac{\delta}{H}$,

$$\sup_{f_t: \|f_t\|_\infty \leq 1} \left| \mathbb{E}_{\text{Unif}(D_2)} [f_t(s_t, a_t) (1 - \mathcal{P}(s_{1..t-1}))] - \mathbb{E}_{\pi^E} [f_t(s_t, a_t) (1 - \mathcal{P}(s_{1..t-1}))] \right| \quad (101)$$

$$\lesssim \sqrt{|\mathcal{S}|} \sqrt{\sum_{s \in \mathcal{S}} \frac{\mathbb{E}_{\pi^E} [\mathbb{1}(s_t = s) (1 - \mathcal{P}(s_{1..t-1}))] \log\left(\frac{|\mathcal{S}|H}{\delta}\right) + \frac{|\mathcal{S}| \log\left(\frac{|\mathcal{S}|H}{\delta}\right)}{|D_2|}}{|D_2|}} \quad (102)$$

$$= \sqrt{|\mathcal{S}|} \sqrt{\frac{\mathbb{E}_{\pi^E} [1 - \mathcal{P}(s_{1..t-1})] \log\left(\frac{|\mathcal{S}|H}{\delta}\right) + \frac{|\mathcal{S}| \log\left(\frac{|\mathcal{S}|H}{\delta}\right)}{|D_2|}}{|D_2|}} \quad (103)$$

$$\stackrel{(i)}{\leq} \min \left\{ \sqrt{\frac{|\mathcal{S}| \log\left(\frac{|\mathcal{S}|H}{\delta}\right)}{|D_2|}}, \sqrt{\frac{|\mathcal{S}| \sum_{t=1}^{H-1} \mathbb{E}_{\pi^E} [1 - \mathcal{M}(s_t, t)] \log\left(\frac{|\mathcal{S}|H}{\delta}\right)}{|D_2|}} \right\} + \frac{|\mathcal{S}| \log\left(\frac{|\mathcal{S}|H}{\delta}\right)}{|D_2|} \quad (104)$$

where (i) follows by the same simplification as in eq. (82). Comparing with eq. (86), this roughly corresponds to setting $F_{\max} \approx \exp(|\mathcal{S}|)$. All in all, summing eq. (104) over $t \in [H]$ and plugging into eq. (88), with probability $\geq 1 - \delta$,

$$\begin{aligned} & \sup_{f \in \mathcal{F}} \left| J_f^{(2)}(\pi^E) - \widehat{V}_f^{(2)} \right| \quad (105) \\ & \lesssim H \left\{ \sqrt{\frac{|\mathcal{S}| \log\left(\frac{|\mathcal{S}|H}{\delta}\right)}{|D_2|}}, \sqrt{\frac{|\mathcal{S}| \sum_{t=1}^{H-1} \mathbb{E}_{\pi^E} [1 - \mathcal{M}(s_t, t)] \log\left(\frac{|\mathcal{S}|H}{\delta}\right)}{|D_2|}} \right\} + \frac{|\mathcal{S}| \log\left(\frac{|\mathcal{S}|H}{\delta}\right)}{|D_2|} \quad (106) \end{aligned}$$

Finally, we plug this into eq. (65), which is restated below,

$$J(\pi^E) - J(\pi^{\text{RE}}) \leq 2 \sup_{f \in \mathcal{F}} \left| J_f(\pi^E) - \widehat{E}(f) \right| \quad (107)$$

$$\leq 2 \underbrace{\sup_{f \in \mathcal{F}} \left| J_f^{(1)}(\pi^E) - \mathbb{E} \left[\widehat{E}^{(1)}(f) \mid D_1 \right] \right|}_{\text{(I)}} + 2 \underbrace{\sup_{f \in \mathcal{F}} \left| \mathbb{E} \left[\widehat{E}^{(1)}(f) \mid D_1 \right] - \widehat{E}^{(1)}(f) \right|}_{\text{(II)}} \quad (108)$$

$$+ 2 \underbrace{\sup_{f \in \mathcal{F}} \left| J_f^{(2)}(\pi^E) - \widehat{E}^{(2)}(f) \right|}_{\text{(III)}}. \quad (109)$$

For the chosen membership oracle in eq. (92), the term (I) is 0, since by Lemma 5 it is upper bounded by $H \sum_{h=1}^H \mathbb{E}_{\pi^E} [\mathcal{M}(s_h, h) \text{TV}(\pi_h^E(\cdot | s_h), \pi_h^{\text{BC}}(\cdot | s_h))]$. This is equal to 0 since $\mathcal{M}(s, t) = 0$ wherever $\pi_t^E(\cdot | s) \neq \pi_t^{\text{BC}}(\cdot | s)$. On the other hand, $N_{\text{replay}} \rightarrow \infty$ ensures that the term (III) goes to

0 by the strong law of large numbers. Therefore, with probability $\geq 1 - 2\delta$,

$$J(\pi^E) - J(\pi^{\text{RE}}) \tag{110}$$

$$\leq 2 \sup_{f \in \mathcal{F}} \left| \mathbb{E} \left[\widehat{E}^{(1)}(f) \middle| D_1 \right] - \widehat{E}^{(1)}(f) \right| \tag{111}$$

$$\lesssim H \left\{ \sqrt{\frac{|\mathcal{S}| \log \left(\frac{|\mathcal{S}|H}{\delta} \right)}{|D_2|}}, \sqrt{|\mathcal{S}| \frac{\sum_{t=1}^{H-1} \mathbb{E}_{\pi^E} [1 - \mathcal{M}(s_t, t)] \log \left(\frac{|\mathcal{S}|H}{\delta} \right)}{|D_2|}} \right\} + \frac{|\mathcal{S}|H \log \left(\frac{|\mathcal{S}|H}{\delta} \right)}{|D_2|} \tag{112}$$

Finally, we bound $\mathbb{E}_{\pi^E} [1 - \mathcal{M}(s_t, t)]$ for the membership oracle defined in eq. (92). By definition, this quantity is the same as $\Pr_{\pi^E} (s_t \text{ not visited in } D_1 \text{ at time } t)$. This is the probability that given N_{exp} samples from a distribution (the state visited at time t in an expert rollout), the probability that a new sample from the same distribution is not in the support of the observed samples. This is known as the missing mass McAllester and Schapire [2000]. In Lemma A.3 Rajaraman et al. [2020] it is shown that with probability $\geq 1 - \delta$,

$$\sum_{t=1}^{H-1} \Pr_{\pi^E} (s_t \text{ not visited in } D_1 \text{ at time } t) \lesssim \frac{|\mathcal{S}|H}{|D_1|} + \frac{\sqrt{|\mathcal{S}|H \log \left(\frac{|\mathcal{S}|H}{\delta} \right)}}{|D_1|} \tag{113}$$

Finally, combining with eq. (112) and using the fact that that $|D_1|, |D_2| = \Theta(N_{\text{exp}})$, with probability $\geq 1 - 3\delta$,

$$J(\pi^E) - J(\pi^{\text{RE}}) \lesssim \min \left\{ H \sqrt{\frac{|\mathcal{S}| \log \left(\frac{|\mathcal{S}|H}{\delta} \right)}{N_{\text{exp}}}}, \frac{|\mathcal{S}|H^{3/2} \log \left(\frac{|\mathcal{S}|H}{\delta} \right)}{N_{\text{exp}}} \right\}. \tag{114}$$

This completes the proof of Theorem 6.

B.2 Bound with parametric function approximation under Lipschitzness

In this section, we provide an upper bound on the imitation gap of RE in the presence of parametric function approximation under a Lipschitzness assumption on the function classes, and assuming access to a parameter estimation oracle for offline classification.

Definition (IL with function-approximation). *In this setting, for each $t \in [H]$, there is a parameter class $\Theta_t \subseteq \mathbb{B}_2^d$, the unit L_2 ball in d dimensions, and an associated function class $\{f_{\theta_t} : \theta_t \in \Theta_t\}$. For each $t \in [H]$ there exists an unknown $\theta_t^E \in \Theta_t$ such that $\forall s \in \mathcal{S}$,*

$$\pi_t^E(s) = \arg \max_{a \in \mathcal{A}} f_{\theta_t^E}(s, a). \tag{115}$$

Definition 3 (Policy induced by a classifier). *Consider a set of parameters $\theta = \{\theta_1, \dots, \theta_H\}$ where $\theta_t \in \Theta_t$ for each t . A policy π^θ is said to be induced by the set of classifiers defined by θ if for all $s \in \mathcal{S}$ and $t \in [H]$,*

$$\pi_t^\theta(s) = \arg \max_{a \in \mathcal{A}} f_{\theta_t}(s, a). \tag{116}$$

By this definition, $\pi^E = \pi^{\theta^E}$ where $\theta^E = \{\theta_1^E, \dots, \theta_H^E\}$.

Definition 4 (Lipschitz parameterization). *A function class $\mathcal{G} = \{g_\theta : \theta \in \Theta\}$ where $g_\theta(\cdot) : \mathcal{X} \rightarrow \mathbb{R}$ is said to satisfy L -Lipschitz parameterization if, $\|g_\theta(\cdot) - g_{\theta'}(\cdot)\|_\infty \leq L \|\theta - \theta'\|_2$ for all $\theta, \theta' \in \Theta$. In other words, for each $x \in \mathcal{X}$, $g_\theta(x)$ is an L -Lipschitz function in θ , in the L_2 norm.*

Assumption 4 (Assumption 1 restated). *For each t , the class $\{f_{\theta_t} : \theta_t \in \Theta_t\}$ is L -Lipschitz in its parameterization, $\theta_t \in \Theta_t$.*

To deal with parametric function approximation, we assume that the learner has access to a proper offline classification oracle, which given a dataset of classification examples, guarantees to approximately return the underlying ground truth parameter. Namely,

Assumption 5 (Assumption 2 restated). We assume that the learner has access to a multi-class classification oracle, which given n examples of the form, (s^i, a^i) where $s^i \stackrel{i.i.d.}{\sim} \mathcal{D}$ and $a^i = \arg \max_{a \in \mathcal{A}} f_{\theta^*}(s^i, a)$, returns a $\hat{\theta} \in \Theta$ such that, with probability $\geq 1 - \delta$, $\|\hat{\theta} - \theta^*\|_2 \leq \mathcal{E}_{\Theta, n, \delta}$.

This assumption implies that the parameter class Θ_t (and the associated function class $\{f_{\theta_t} : \theta_t \in \Theta_t\}$) admits finite sample complexity guarantees for learning the parameter θ_t^* given classification examples from the underlying ground truth function $f_{\theta_t^*}$. As we discuss in more detail later, we will assume that this classification oracle is used by RE to train the BC policy in Line 3 of Algorithm 1.

Finally, we introduce the main assumption on the IL instances we study. We assume that the classification problems solved by BC at each $t \in [H]$ satisfy a margin condition.

Assumption 6 (Assumption 3 restated). For $\theta \in \Theta_t$, define $a_s^\theta = \arg \max_{a \in \mathcal{A}} f_\theta(s, a)$ as the classifier output on the state s . The weak margin condition assumes that for each t , there is no classifier $\theta \in \Theta_t$ such that for a large mass of states, $f_\theta(s_t, a_{s_t}^\theta) - \max_{a \neq a_{s_t}^\theta} f_\theta(s_t, a)$, i.e. the “margin” from the nearest classification boundary, is small. Formally, the weak-margin condition with parameter μ states that, for any $\theta \in \Theta_t$ and $\eta \leq 1/\mu$,

$$\Pr_{\pi^E} \left(f_\theta(s_t, a_{s_t}^\theta) - \max_{a \neq a_{s_t}^\theta} f_\theta(s_t, a) \geq \eta \right) \geq e^{-\mu\eta}. \quad (117)$$

The weak margin condition only assumes that there is at least an exponentially small (in η) mass of states with margin at least η . Smaller μ indicates a larger mass away from any decision boundary. It suffices to assume that eq. (117) is only true for θ as the classifier in Assumption 5 for our guarantees (Theorem 7) to hold.

Remark 5. Note that the weak margin condition is the multi-class extension of the Tsybakov margin condition of Mammen and Tsybakov [1999], Audibert and Tsybakov [2007] defined for the binary case. In particular, in eq. (117), we may replace the RHS by $1 - \mu\eta$, or $1 - (\mu\eta)^\alpha$ for $\alpha > 0$ to get different analogs of the margin condition and the main guarantee, Theorem 3, as we discuss in Appendix B.2.

Theorem 7. For IL with parametric function approximation, under Assumptions 4 to 6, appropriately instantiating RE ensures that with probability $\geq 1 - 4\delta$,

$$J(\pi^E) - J(\pi^{RE}) \lesssim H^{3/2} \sqrt{\frac{\mu L \log(F_{\max} H / \delta)}{N_{\exp}} \frac{\sum_{t=1}^H \mathcal{E}_{\Theta_t, N_{\exp}, \delta / H}}{H}} + \frac{\log(F_{\max} H / \delta)}{N_{\exp}}. \quad (118)$$

Note that we assume the same conditions on \mathcal{F} as required in Theorem 2.

Remark 6. The classification oracle in Assumption 5 asks for a stronger condition than just finding a classifier with small generalization error, which need not be close to the ground truth θ^* in the parameter space. Learning classifiers with small generalization error is studied in Daniely et al. [2013] who show that the Natarajan dimension, up to log-factors in the number of classes (i.e. number of actions) is the right statistical complexity measure which characterizes the generalization error of the best learner. In the realizable case, the optimal generalization error guarantee scales as $\tilde{O}(1/n)$ where n is the number of classification examples. Under certain assumptions on the input distribution \mathcal{D} and the function family (e.g. for linear families), we later show that the generalization error guarantee can be extended to approximately learning the parameter as well (up to problem dependent constants). Generally, under two conditions,

1. Generalization error guarantees which scale as $\tilde{O}(1/N_{\exp})$ can be extended to parameter learning,
2. The dependence of the generalization error on the failure probability δ scales as $\text{polylog}(1/\delta)$,

the optimality gap for RE which we prove in Theorem 7 scales as $\tilde{O}(H^{3/2}/N_{\exp})$. The constants in $\tilde{O}(\cdot)$ here depend on the Natarajan dimension and the covering number of the function classes \mathcal{F}_t among problem dependent constants.

In proving this result, we first discuss the implementation of RE .

Implementation of RE (Algorithm 1) We discuss the instantiation of RE in the Lipschitz setting below. The underlying function class \mathcal{F} is chosen arbitrarily (note that the guarantee we prove depends on this function class, and the only constraints on \mathcal{F} are those in Theorem 2 - the ground truth reward must belong in $\mathcal{F} = \otimes_{t=1}^H \mathcal{F}_t$, the function class is symmetric, i.e., $f_t \in \mathcal{F}_t \iff -f_t \in \mathcal{F}_t$ for each t and for all $f_t \in \mathcal{F}_t, \|f_t\|_\infty \leq 1$) This requires specifying the choice of the membership oracle \mathcal{M} and describing the instantiation of BC .

Implementation of BC : Recall that in Algorithm 1, the learner trains BC on the dataset D_1 . In particular, under the offline classification oracle condition, Assumption 5, the learner trains H classifiers, one for each t , trained on the state-action pairs (i.e. state is the input, and the action at this state is the corresponding class) observed in the expert dataset at time t using the offline classifier in Assumption 5. We assume that each classifier is trained with the failure probability chosen as δ/H . Denoting this set of H classifiers as $\widehat{\theta}^{\text{BC}} = \{\widehat{\theta}_1^{\text{BC}}, \dots, \widehat{\theta}_H^{\text{BC}}\}$, this corresponds to the policy $\pi^{\text{BC}} = \pi^{\widehat{\theta}^{\text{BC}}}$ induced by the classifier $\widehat{\theta}^{\text{BC}}$ (Definition 3).

In particular, by union bounding, the classifiers $\widehat{\theta}^{\text{BC}}$ satisfy with probability $\geq 1 - \delta$ simultaneously for each time $t \in [H]$,

$$\|\theta_t^E - \widehat{\theta}_t\|_2 \leq \mathcal{E}_{\Theta_t, N_{\text{exp}}, \delta/H}. \quad (119)$$

Membership oracle: Fix a time-step $t \in [H]$. The membership oracle \mathcal{M} is defined in eq. (10) as,

$$\mathcal{M}(s, t) = \begin{cases} +1 & \text{if } \exists a \in \mathcal{A} \text{ such that, } \forall a' \in \mathcal{A}, f_{\widehat{\theta}_t^{\text{BC}}}(s, a) - f_{\widehat{\theta}_t^{\text{BC}}}(s, a') \geq 2L\mathcal{E}_{\Theta_t, N_{\text{exp}}, \delta/H} \\ 0 & \text{otherwise.} \end{cases} \quad (120)$$

We first show that on the states such that the membership oracle is 1, the expert policy perfectly matches the learner's policy.

Lemma 8. *At every state s such that $\mathcal{M}(s, t) = +1$, $\pi_t^E(s) = \pi_t^{\text{BC}}(s)$.*

Proof. Note that θ_t^E satisfies $\|\theta_t^E - \widehat{\theta}_t^{\text{BC}}\|_2 \leq \mathcal{E}_{\Theta_t, N_{\text{exp}}, \delta/H}$ with probability $1 - \delta$. Consider the action a played by the learner, for any $a' \in \mathcal{A}$,

$$f_{\theta_t^E}(s, a) - f_{\theta_t^E}(s, a') \geq f_{\widehat{\theta}_t^{\text{BC}}}(s, a) - \mathcal{E}_{\Theta_t, N_{\text{exp}}, \delta/H}L - f_{\widehat{\theta}_t^{\text{BC}}}(s, a') - \mathcal{E}_{\Theta_t, N_{\text{exp}}, \delta/H}L \quad (121)$$

$$\geq 0 \quad (122)$$

where the first inequality follows by Lipschitzness of $f(s, a)$ and the last inequality follows by definition of the set of states where $\mathcal{M}(s, t) = +1$: $\forall a' \in \mathcal{A}, f_{\widehat{\theta}_t^{\text{BC}}}(s, a) - f_{\widehat{\theta}_t^{\text{BC}}}(s, a') \geq 2\mathcal{E}_{\Theta_t, N_{\text{exp}}, \delta/H}L$.

Since for this action a , $f_{\theta_t^E}(s, a) - f_{\theta_t^E}(s, a') \geq 0$ for all other actions $a' \in \mathcal{A}$, a must be the action played by the expert policy. This completes the proof. \square

Note that π^{BC} always matches π^E wherever the membership oracle \mathcal{M} is non-zero. We run Algorithm 1. Therefore, from Theorem 2, with probability $\geq 1 - 4\delta$, the imitation gap of the learner is bounded by,

$$J_r(\pi^E) - J_r(\pi^{\text{RE}}) \lesssim H^{3/2} \sqrt{\frac{\log(F_{\text{max}}H/\delta) \sum_{t=1}^H \mathbb{E}_{\pi^E} [1 - \mathcal{M}(s_t, t)]}{N_{\text{exp}} H}} + \frac{\log(F_{\text{max}}H/\delta)}{N_{\text{exp}}}. \quad (123)$$

To complete the proof, we must bound $\mathbb{E}_{\pi^E} [1 - \mathcal{M}(s_t, t)]$, which is the measure of states s such that $\forall a \in \mathcal{A}, \exists a' \in \mathcal{A} : f_{\widehat{\theta}_t}(s, a) - f_{\widehat{\theta}_t}(s, a') \leq 2L\mathcal{E}_{\Theta_t, N_{\text{exp}}, \delta/H}$, i.e. the mass of states which are very close to a decision boundary. The probability of this set of states is upper bounded by the weak margin condition. Indeed, for each $t \in [H]$, defining $a_s^* = \arg \max_{a \in \mathcal{A}} f_{\widehat{\theta}_t^{\text{BC}}}(s, a)$,

$$\Pr_{\pi^E} \left(f_{\widehat{\theta}_t^{\text{BC}}}(s_t, a_{s_t}^*) - \max_{a \neq a_{s_t}^*} f_{\widehat{\theta}_t^{\text{BC}}}(s_t, a) \geq 2L\mathcal{E}_{\Theta_t, N_{\text{exp}}, \delta/H} \right) \geq e^{-\mu L \mathcal{E}_{\Theta_t, N_{\text{exp}}, \delta/H}} \quad (124)$$

$$\geq 1 - \mu L \mathcal{E}_{\Theta_t, N_{\text{exp}}, \delta/H}. \quad (125)$$

Therefore,

$$\mathbb{E}_{\pi^E} [1 - \mathcal{M}(s_t, t)] \lesssim \mu L \mathcal{E}_{\Theta_t, N_{\text{exp}}, \delta/H}. \quad (126)$$

Putting it together with eq. (123), and simplifying, with probability $\geq 1 - 4\delta$,

$$J(\pi^E) - J(\pi^{\text{RE}}) \lesssim H^{3/2} \sqrt{\frac{\mu L \log(F_{\max} H/\delta)}{N_{\text{exp}}} \frac{\sum_{t=1}^H \mathcal{E}_{\Theta_t, N_{\text{exp}}, \delta/H}}{H}} + \frac{\log(F_{\max} H/\delta)}{N_{\text{exp}}}. \quad (127)$$

Note that in Equation (125), we only use the fact that the probability mass of states which are η -close to any decision boundary is not too high. Similar to Audibert and Tsybakov [2007], we may consider relaxations of the weak margin condition, as below.

Assumption 7 (α -weak margin condition). *Consider any $t \in [H]$ and $\theta \in \Theta_t$. For each $s \in \mathcal{S}$, define $a_s^* = \arg \max_{a \in \mathcal{A}} f_\theta(s, a)$ as the classifier output under f_θ . The α weak margin condition with parameter μ assumes that, for any $\eta \leq 1/\mu$,*

$$\forall \theta \in \Theta_t, \Pr_{\pi^E} \left(f_\theta(s_t, a_{s_t}^*) - \max_{a \neq a_{s_t}^*} f_\theta(s_t, a) \geq \eta \right) \geq 1 - (\mu\eta)^\alpha. \quad (128)$$

When $\alpha = 1$, this condition is effectively equivalent to the weak margin condition in Assumption 6.

Following the proof of Theorem 7, we may obtain the following result under the α weak margin condition for $\alpha \neq 1$.

Theorem 8. *For IL with parametric function approximation, under Assumptions 4, 5 and 7, appropriately instantiating RE ensures that with probability $\geq 1 - 4\delta$,*

$$J(\pi^E) - J(\pi^{\text{RE}}) \lesssim H^{3/2} \sqrt{\frac{(\mu L)^\alpha \log(F_{\max} H/\delta)}{N_{\text{exp}}} \frac{\sum_{t=1}^H (\mathcal{E}_{\Theta_t, N_{\text{exp}}, \delta/H})^\alpha}{H}} + \frac{\log(F_{\max} H/\delta)}{N_{\text{exp}}}. \quad (129)$$

Once again, we assume the same conditions on \mathcal{F} as required in Theorem 2.

B.3 Extension to unbounded discriminator families

Note that when the family of discriminators \mathcal{F} does not have finite cardinality, it in fact suffices to just bound the imitation gap against a finite covering of \mathcal{F} . We spell out the details explicitly below.

In particular, we can replace F_{\max} by $\max_{t \in [H]} \mathcal{N}(\mathcal{F}_t, 1/N_{\text{exp}}, \|\cdot\|_\infty)$, where $\mathcal{N}(\mathcal{G}, \xi, \|\cdot\|)$ denotes the covering number of \mathcal{G} in the norm $\|\cdot\|$ as defined below.

Definition 5 (Covering number). *For a function class \mathcal{G} , tolerance ξ and norm $\|\cdot\|$, the covering number $\mathcal{N}(\mathcal{G}, \xi, \|\cdot\|)$ is defined as the cardinality of the smallest set of functions \mathcal{G}^ξ such that for each $g \in \mathcal{G}$, there exists a $g' \in \mathcal{G}^\xi$,*

$$\|g - g'\| \leq \xi. \quad (130)$$

Corollary 1. *When \mathcal{G} is chosen as the set of 1-bounded linear functions, $\mathcal{G} = \{\langle x, \theta \rangle : x \in \mathbb{B}_2^d\} : \theta \in \mathbb{B}_2^d\}$, where \mathbb{B}_2^d denotes the L_2 unit ball in \mathbb{R}^d , $\mathcal{N}(\mathcal{G}, \xi, \|\cdot\|_\infty) \leq \left(\frac{2\sqrt{d}}{\xi} + 1\right)^d$.*

Proof. For any $g, g' \in \mathcal{G}$, where g and g' correspond to parameters $\theta, \theta' \in \mathbb{B}_2^d$,

$$\|g - g'\|_\infty \leq \max_{x \in \mathcal{X}} \langle x, \theta - \theta' \rangle \quad (131)$$

$$\leq \|x\|_2 \|\theta - \theta'\|_2 \quad (132)$$

$$\leq \|\theta - \theta'\|_2. \quad (133)$$

Since the L_2 covering number of \mathbb{B}_2^d is bounded by $\left(\frac{2\sqrt{d}}{\xi} + 1\right)^d$, the result immediately follows by defining the covering of \mathcal{G} as $\{\langle \theta, \cdot \rangle : \theta \in \mathcal{K}\}$ where \mathcal{K} is the optimal covering of \mathbb{B}_2^d in L_2 norm. \square

Definition 6 (Discretization of discriminator space). Define \mathcal{F}_t^ξ as the optimal covering of \mathcal{F}_t under the L_∞ norm in the sense of Definition 5. The discretized family of discriminators we consider is, $\mathcal{F}^\xi = \otimes_{t=1}^H \mathcal{F}_t^\xi$.

Lemma 9. Suppose for all functions $f' \in \mathcal{F}_t^{\xi_1/H}$, simultaneously $J_{f'}(\pi^E) - J_{f'}(\pi^{RE}) \leq \xi_2$. Then, for all discriminators $f \in \mathcal{F}$, $J_f(\pi^E) - J_f(\pi^{RE}) \leq 2\xi_1 + \xi_2$.

Proof. Consider any discriminator $f \in \mathcal{F}$. By construction, there exists an $f' \in \mathcal{F}^{\xi_1/H}$ such that,

$$\|f - f'\|_\infty \leq \xi_1/H. \quad (134)$$

Since for any policy π , the value $J_f(\pi)$ under a discriminator $f \in \mathcal{F}$ is an H -Lipschitz function of f , we can make a statement about how well $J_{f'}(\pi)$ approximates $J_f(\pi)$ for an appropriately chosen $f' \in \mathcal{F}^{\xi_1/H}$. In particular, the nearest (in L_∞ norm) $f' \in \mathcal{F}^{\xi_1/H}$ to $f \in \mathcal{F}$ satisfies that for any policy π ,

$$|J_f(\pi) - J_{f'}(\pi)| \leq H \times \frac{\xi_1}{H}. \quad (135)$$

As a consequence, for any discriminator $f \in \mathcal{F}$,

$$J_f(\pi^E) - J_f(\pi^{RE}) \leq |J_f(\pi^E) - J_{f'}(\pi^E)| + J_{f'}(\pi^E) - J_{f'}(\pi^{RE}) + |J_{f'}(\pi^{RE}) - J_f(\pi^{RE})| \quad (136)$$

$$\leq \xi_1 + \xi_2 + \xi_1 = 2\xi_1 + \xi_2. \quad (137)$$

□

In particular, this means that if we minimize $J_{f'}(\pi^E) - J_{f'}(\pi^{RE}) \leq \xi_2$ for all $f' \in \mathcal{F}^{1/N_{\text{exp}}H}$, then we can ensure that for all $f \in \mathcal{F}$,

$$J_f(\pi^E) - J_f(\pi^{RE}) \leq \frac{2}{N_{\text{exp}}} + \xi_2. \quad (138)$$

This implies the following theorem,

Theorem 9. Consider the policy π^{RE} returned by Algorithm 1 where \mathcal{F} is instead chosen as $\mathcal{F}^{\frac{1}{HN_{\text{exp}}}}$ (as defined in Definition 6). Assume that $\pi^E \in \Pi$, the ground truth reward function $r_t \in \mathcal{F}_t$ which is assumed to be bounded (For all $f_t \in \mathcal{F}_t$, $\|f_t\|_\infty \leq 1$). Choose $|D_1|, |D_2| = \Theta(N_{\text{exp}})$ and suppose $N_{\text{replay}} \rightarrow \infty$. With probability $\geq 1 - 3\delta$,

$$J(\pi^E) - J(\pi^{RE}) \lesssim \mathcal{L}_1 + \mathcal{L}_2 + \frac{\log(\mathcal{N}_{\text{max}}H/\delta) + 1}{N_{\text{exp}}} \quad (139)$$

where $\mathcal{N}_{\text{max}} \triangleq \max_{t \in [H]} \mathcal{N}(\mathcal{F}_t, 1/HN_{\text{exp}}, \|\cdot\|_\infty)$ corresponds to the maximal covering number of the function classes \mathcal{F}_t , and,

$$\mathcal{L}_1 \triangleq H^2 \mathbb{E}_{\pi^E} \left[\frac{\sum_{t=1}^H \mathcal{M}(s_t, t) \text{TV}(\pi_t^E(\cdot|s_t), \pi_t^{BC}(\cdot|s_t))}{H} \right], \quad (140)$$

$$\mathcal{L}_2 \triangleq H^{3/2} \sqrt{\frac{\log(\mathcal{N}_{\text{max}}H/\delta)}{N_{\text{exp}}} \frac{\sum_{t=1}^H \mathbb{E}_{\pi^E} [1 - \mathcal{M}(s_t, t)]}{H}}.$$

Remark 7. Note that this line of reasoning can be extended to Theorem 3 and Theorem 8 to show that the same guarantees as eq. (12) and eq. (129) respectively hold, but with F_{max} replaced by \mathcal{N}_{max} .

B.4 Bounds on RE in the linear expert setting

In this section, we provide an upper bound on the imitation gap of RE in the presence of linear function approximation, which is studied in Rajaraman et al. [2021]. This is a special case of the case of IL under parametric function approximation with Lipschitzness. To avoid any ambiguity, we formally define IL with linear function approximation, which is the case when (i) the expert follows

an unknown linear classifier in a known set of features, and (ii) the reward function admits a linear parameterization.

The goal is to show that there exists a simple choice of the membership oracle such that the imitation gap of the resulting algorithm grows as $H^{3/2}$ and decay in the size of the expert dataset as $1/N_{\text{exp}}$ up to logarithmic factors, breaking the error compounding barrier and achieving the optimal dependency on these parameters. We first introduce the linear setting below.

Assumption 8 (Linear-expert setting). *For each (s, a, t) tuple, the learner is assumed to have a feature representation $\phi_t(s, a) \in \mathbb{R}^d$. For each time t , there exists an unknown vector $\theta_t^E \in \mathbb{R}^d$ such that $\forall s \in \mathcal{S}$,*

$$\pi_t^E(s) = \arg \max_{a \in \mathcal{A}} \langle \theta_t^E, \phi_t(s, a) \rangle. \quad (141)$$

i.e., the expert policy is deterministic and realized by a linear classifier. We assume that $\theta_t^E \in \mathbb{S}^{d-1}$ without loss of generality.

Definition 7 (Policy induced by a linear classifier). *Consider a set of vectors $\theta = \{\theta_1, \dots, \theta_H\}$ where each $\theta_t \in \mathbb{R}^d$. A policy π^θ is said to be induced by the set of linear classifiers defined by θ if for all $s \in \mathcal{S}$ and $t \in [H]$,*

$$\pi_t^\theta(s) = \arg \max_{a \in \mathcal{A}} \langle \theta_t, \phi_t(s, a) \rangle. \quad (142)$$

By this definition, $\pi^E = \pi^{\theta^E}$.

Definition 8 (Linear reward setting). *Define $\mathcal{R}_{\text{lin},t}$ as the family of linear reward functions (defined at the single time-step t) which takes the form of an unknown linear function of a set of the features,*

$$\mathcal{R}_{\text{lin},t} = \{ \{r_t(s, a) = \langle \omega, \phi_t(s, a) \rangle : s \in \mathcal{S}, a \in \mathcal{A}\} : \omega \in \mathbb{R}^d, \|\omega\|_2 \leq 1 \}. \quad (143)$$

For the rewards to be 1-bounded, we assume the features satisfy $\|\phi_t(s, a)\|_2 \leq 1$. Define $\mathcal{R}_{\text{lin}} = \otimes_{t=1}^H \mathcal{R}_{\text{lin},t}$. The linear reward setting assumes the true reward function of the MDP, $r \in \mathcal{R}_{\text{lin}}$.

Remark 8. *Note that our guarantees in Theorem 10 hold even if the set of features in the definition of $\mathcal{R}_{\text{lin},t}$ in Definition 8 differ from those used to define the expert classifier Assumption 8. Regardless, we assume that both sets of features are known to the learner.*

In the case of parametric function approximation with Lipschitzness, note that we assume both the weak margin condition (Assumption 6), as well as the existence of a linear classification oracle (Assumption 5). Below, in the linear expert case, we show a sufficient condition which implies both of these conditions. In particular, define the positive hemisphere with pole at θ , i.e. $\{x : \mathbb{B}_2^d : \langle \theta, x \rangle \geq 0\}$ as \mathbb{H}_θ^d . We abbreviate $\mathbb{H}_{\theta^E}^d$ as \mathbb{H}_t^d .

Assumption 9 (Bounded density assumption). *For each time $t \in [H]$, state $s \in \mathcal{S}$, action $a \in \mathcal{A}$ and $\theta \in \Theta_t$, define $\bar{\phi}_t(s, a) = \phi_t(s, a_s^\theta) - \phi_t(s, a)$ where $a_s^\theta = \arg \max_{a' \in \mathcal{A}} \langle \theta, \phi_t(s, a') \rangle$. Consider the measure $\Pr_{\pi^E}(\exists a \neq a_{s_t}^\theta : \bar{\phi}_t(s_t, a) \in \cdot)$. Let \bar{d}_t^E represent the Radon-Nikodym derivative of this measure against the uniform measure on \mathbb{H}_t^{d-1} . The bounded density assumption states that for each $t \in [H]$ there are constants $c_{\min} > 0$ and $c_{\max} < \infty$ such that for all $x \in \mathbb{H}_t^d$,*

$$c_{\min} \leq \bar{d}_t^E(x) \leq c_{\max}. \quad (144)$$

We now state the main result we prove for IL in the linear setting.

Theorem 10. *Under Assumptions 8 and 9, appropriately instantiating RE ensures that with probability $\geq 1 - \delta$,*

$$J(\pi^E) - J(\pi^{\text{RE}}) \lesssim \sqrt{\frac{c_{\max}}{c_{\min}}} \frac{H^{3/2} d^{5/4} \log^{\frac{3}{2}}(N_{\text{exp}} d H / \delta)}{N_{\text{exp}}}.$$

The proof of this result follows by showing that under Assumption 9, both the weak margin condition (Assumption 6) is satisfied, and the classification oracle (Assumption 5) can be constructed. We begin by showing the former.

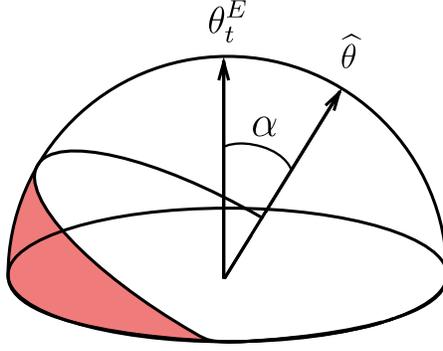


Figure 8: If at any time $t \in [H]$ and state s , $\phi_t(s, \pi_t^E(s)) - \phi_t(s, a)$ lies in the red shaded region for some action a , then, the action played by π^{BC} and π^E at this state are different.

Lemma 10. *Under Assumption 9, the α -weak margin condition is satisfied with $\alpha = 1$ and $\mu = 2c_{\max}\sqrt{d}$. In particular, for all $\theta \in \mathbb{S}^{d-1}$,*

$$\Pr_{\pi^E} \left(\langle \theta, \phi_t(s, a_{s_t}^\theta) \rangle - \max_{a \neq a_{s_t}^\theta} \langle \theta, \phi_t(s, a) \rangle \geq \eta \right) \geq 1 - \left(2c_{\max}\sqrt{d} \right) \eta. \quad (145)$$

where $a_{s_t}^\theta \triangleq \arg \max_{a \in \mathcal{A}} \langle \theta, \phi_t(s, a) \rangle$.

Proof. Observe that,

$$\Pr_{\pi^E} (\exists a \neq a_{s_t}^\theta : \langle \theta, \phi_t(s_t, a_{s_t}^\theta) \rangle - \langle \theta, \phi_t(s_t, a) \rangle \leq \eta) \quad (146)$$

$$= \Pr_{\pi^E} (\exists a \neq a_{s_t}^\theta : \phi_t(s_t, a_{s_t}^\theta) - \phi_t(s_t, a) \in \{x \in \mathbb{H}_\theta^d : \langle x, \theta \rangle \leq \eta\}) \quad (147)$$

$$\stackrel{(i)}{=} \Pr_{\pi^E} (\exists a \neq a_{s_t}^\theta : \bar{\phi}_t(s_t, a) \in \{x \in \mathbb{H}_\theta^d : \langle x, \theta \rangle \leq \eta\}) \quad (148)$$

$$\stackrel{(ii)}{\leq} c_{\max} \Pr(\langle U, \theta \rangle \leq \eta) \quad (149)$$

where in (i), $\bar{\phi}_t$ is as defined in Assumption 9 and in (ii), U is uniformly distributed on the unit hemisphere, \mathbb{H}_θ^d . Note that (ii) follows from the bounded density condition, Assumption 9. Note that the RHS essentially corresponds to the volume (probability measure) of a disc of height η cut out of a sphere from the center. Up to normalization factors, this can be upper bounded by the surface area of the base of the disc, multiplied by the height of the disc. Namely,

$$\frac{\eta \times \frac{\pi^{\frac{d-1}{2}}}{\Gamma(\frac{d-1}{2}+1)}}{\frac{1}{2} \frac{\pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2}+1)}} \quad (150)$$

Using Gautschi's inequality, for any $x \geq 0$ and $\ell \in (0, 1)$ $x^{1-\ell} \leq \frac{\Gamma(x+1)}{\Gamma(x+\ell)} \leq (1+x)^{1-\ell}$. With $\ell = \frac{1}{2}$, $\frac{\Gamma(\frac{d}{2}+1)}{\Gamma(\frac{d+1}{2})} \leq \sqrt{d}$. Combining with eq. (149) results in,

$$\Pr_{\pi^E} (\exists a \neq a_{s_t}^\theta : \langle \theta, \phi_t(s_t, a_{s_t}^\theta) \rangle - \langle \theta, \phi_t(s_t, a) \rangle \leq \eta) \leq 2c_{\max}\sqrt{d}\eta \quad (151)$$

Therefore the probability of the complement event is lower bounded by $1 - 2c_{\max}\sqrt{d}\eta$, completing the proof. \square

The final thing to show is that the bounded density assumption can also be used to construct a classification oracle in the sense of Assumption 5.

In particular, as discussed in the main paper, we show that algorithms for minimizing the generalization error, can be used to construct a classification oracle. The compression based algorithm of

Daniely and Shalev-Shwartz [2014] provides a guarantee on the generalization error. From Theorem 5 of Daniely and Shalev-Shwartz [2014], in the realizable setting, for linear classification, the resulting classifier $\hat{\theta}$ has expected 0-1 loss upper bounded by $(d + \log(1/\delta)) \log(n)/n$, given n classification examples. Namely, in the notation of Assumption 5, the resulting classifier $\hat{\theta}$ satisfies with probability $\geq 1 - \delta$,

$$\Pr_{s \sim \mathcal{D}} \left(\arg \max_{a \in \mathcal{A}} f_{\theta^*}(s^i, a) \neq \arg \max_{a \in \mathcal{A}} f_{\hat{\theta}}(s^i, a) \right) \leq \frac{(d + \log(1/\delta)) \log(n)}{n}. \quad (152)$$

Next we show that under Assumption 9, this equation can be used to bound the error in the parameter space, $\|\theta^* - \hat{\theta}\|_2$. Namely, in Assumption 5, we may choose $\mathcal{E}_{\mathbb{B}_2^d, n, \delta}$ as $\asymp \frac{(d + \log(1/\delta)) \log(n)}{n}$, up to constants depending on c_{\min} .

Lemma 11. *Consider the compression based learner $\hat{\theta}_t^{\text{BC}} = \hat{\theta}_t$ of Daniely and Shalev-Shwartz [2014] for multi-class linear classification. Then, under Assumption 9, with probability $\geq 1 - \delta$,*

$$\|\hat{\theta}_t^{\text{BC}} - \theta_t^*\|_2 \leq \frac{2\pi}{c_{\min}} \frac{(d + \log(1/\delta)) \log(N_{\text{exp}})}{N_{\text{exp}}} \quad (153)$$

Proof. Fix $t \in [H]$. The generalization error of $\hat{\theta}_t^{\text{BC}} = \hat{\theta}_t$ can be written as,

$$\begin{aligned} & \Pr_{\pi^E} \left(\arg \max_{a \in \mathcal{A}} \langle \theta_t^*, \phi_t(s_t, a) \rangle \neq \arg \max_{a \in \mathcal{A}} \langle \hat{\theta}_t^{\text{BC}}, \phi_t(s_t, a) \rangle \right) \\ &= \Pr_{\pi^E} \left(\exists a \neq \pi_t^E(s_t) : \phi_t(s_t, \pi_t^E(s_t)) - \phi_t(s_t, a) \in \mathcal{C} \right), \end{aligned} \quad (154)$$

where \mathcal{C} is illustrated in fig. 8 and is formally defined as,

$$\mathcal{C} \triangleq \{x \in \mathbb{H}_t^d : \langle x, \hat{\theta}_t^{\text{BC}} \rangle \leq 0\}. \quad (155)$$

On the states which “belong” to \mathcal{C} (i.e. at those states s where $\exists a \neq \pi_t^E(s_t) : \phi_t(s, \pi_t^E(s_t)) - \phi_t(s, a) \in \mathcal{C}$), there exists an action a such that $\hat{\theta}_t^{\text{BC}}$ is better correlated with this action than a_s^* . In other words, $\hat{\theta}_t^{\text{BC}}$ and θ_t^* play different actions at this state. Note that \mathcal{C} is essentially the set difference of two hemispheres with different poles. By the bounded density condition, Assumption 9, and eq. (154),

$$\Pr_{\pi^E} \left(\pi_t^E(s_t) \neq \arg \max_{a \in \mathcal{A}} \langle \hat{\theta}_t^{\text{BC}}, \phi(s, a) \rangle \right) \geq c_{\min} \Pr(U \in \mathcal{C}), \quad (156)$$

where U is uniformly distributed over \mathbb{H}_t^d . Referring to fig. 8, we have that,

$$\Pr(U \in \mathcal{C}) = \frac{\alpha}{\pi} \quad (157)$$

where α is the angle between $\hat{\theta}_t^{\text{BC}}$ and θ_t^E . In particular, from eq. (156),

$$\Pr_{\pi^E} \left(a_s^* \neq \arg \max_{a \in \mathcal{A}} \langle \hat{\theta}_t^{\text{BC}}, \phi(s, a) \rangle \right) \geq c_{\min} \frac{\alpha}{\pi} \geq c_{\min} \frac{\|\theta_t^* - \hat{\theta}_t^{\text{BC}}\|_2}{\pi}, \quad (158)$$

where in the last inequality, we use the fact that $\|\theta_t^*\|_2 = \|\hat{\theta}_t^{\text{BC}}\|_2 = 1$ without loss of generality. By the generalization error bound on $\hat{\theta}_t^{\text{BC}} = \hat{\theta}_t$ in eq. (152), with probability $\geq 1 - \delta$,

$$\|\theta_t^* - \hat{\theta}_t^{\text{BC}}\|_2 \leq \frac{\pi}{c_{\min}} \frac{(d + \log(1/\delta)) \log(N_{\text{exp}})}{N_{\text{exp}}} \quad (159)$$

□

Lemma 11 shows that under the bounded density condition Assumption 9, the compression based learner $\hat{\theta}$ of Daniely and Shalev-Shwartz [2014] essentially induces a classification oracle for linear classification with $\mathcal{E}_{\mathbb{B}_2^d, n, \delta} = \frac{\pi}{c_{\min}} \frac{(d + \log(1/\delta)) \log(n)}{n}$. Finally, from Corollary 1, we have a bound on the covering number of linear families. Putting together all of these results with Theorem 3 (noting that we can replace F_{\max} by \mathcal{N}_{\max} from Remark 7) results in Theorem 10.

C Additional experimental results

We include some additional experimental results in this section of the paper.

In fig. 9, we plot the distributions of the prefix weights generated by each membership oracle on simulated BC rollouts on WalkerBulletEnv. Note that \mathcal{M}_{VAR} is significantly overconfident in prefix weights compared to \mathcal{M}_{EXP} , as indicated by the heavier right-tail. On the other hand, \mathcal{M}_{RND} and \mathcal{M}_{MAX} are less overconfident and better overlap with the idealized prefix weights induced by \mathcal{M}_{EXP} . This aligns with the correlation plot between the various membership oracles in Fig. 5. Moreover, in terms of policy performance, this further justifies the superior behavior of \mathcal{M}_{MAX} compared to \mathcal{M}_{VAR} .

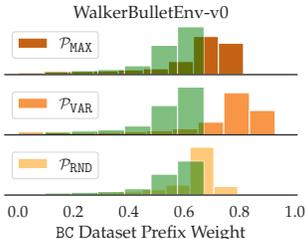


Figure 9: Histogram of prefix weights generated by rolling out trajectories from BC . The green superimposed histogram represents prefix weights generated by \mathcal{M}_{EXP}

In Fig. 10, we consider how each of the changes we described in the main section of the paper lead to improved performance of our RE baseline. The first, using a Wasserstein distance, leads to lower expected return but is required for solving the full moment-matching problem – see Swamy et al. [2021] for more details. Switching from PPO to the more sample-efficient SAC [Haarnoja et al., 2018] leads to fast learning. Adding in gradient penalties for discriminator stability [Swamy et al., 2021, Gulrajani et al., 2017] also improves final performance and learning speed. The last change we employ, using Optimistic Mirror Descent [Daskalakis et al., 2017] in both the discriminator and RL algorithm also improves performance. To our knowledge, we are the first to utilize this technique in the imitation learning literature and recommend it as best practice for future moment-matching algorithms. We refer interested readers to the work of Syrgkanis et al. [2015] for theoretical details of why OMD enables superior performance.

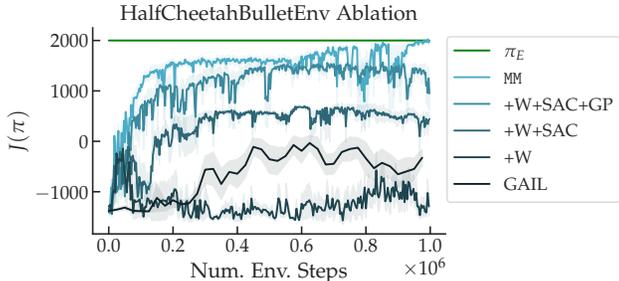


Figure 10: We ablate the four key changes we made to off-the-shelf GAIL to improve performance / theoretical guarantees. We see that each improved performance, with MM significantly out-performing options with fewer changes. Our improvements upon MM with the Replay Estimation technique are therefore improving upon an already strong baseline.

D Experimental Setup

We begin with the hyperparameters for our Standard Bullet and Noisy Bullet experiments.

D.1 Expert

We use the Stable Baselines 3 [Raffin et al., 2021] implementation of PPO [Schulman et al., 2017] or SAC [Haarnoja et al., 2018] to train experts for each environment. For the most part, we use already tuned hyperparameters from [Raffin, 2020] in the implementation. The modifications we used are shown in table 1.

PARAMETER	VALUE
BUFFER SIZE	300000
BATCH SIZE	256
γ	0.98
τ	0.02
TRAINING FREQ.	64
GRADIENT STEPS	64
LEARNING RATE	LIN. SCHED. 7.3E-4
POLICY ARCHITECTURE	256 X 2
STATE-DEPENDENT EXPLORATION	TRUE
TRAINING TIMESTEPS	1E6

Table 1: Expert hyperparameters for Walker Bullet Task and Hopper Bullet Task

D.1.1 Noisy Experts

In addition to the default Bullet Tasks, we test performance of algorithms on noisy environments. Namely, we generate noisy expert data by re-training expert policies with Gaussian noise added to the actions of the expert during the exploration phase while training. We vary the amount of noise of the average action values across each dimension. We then re-generate expert data by sampling from the expert policies trained on noisy data to analyze the performance of our method under stochasticity. Table 2 showcases the effect and new performance of the expert policies across different Bullet Tasks we used to evaluate our method as well as the standard deviation we applied to the actions of the different environments.

ENV.	NOISY EXPERT	STANDARD DEV.
HOPPER	1524	$\mathcal{N}(0, 0.5)$
WALKER	1586	$\mathcal{N}(0, 0.5)$

Table 2: Evaluated performance of the noise trained expert policies.

D.2 Baseline

We average over 5 runs and use a common architecture of 256 x 2 with ReLU activations for both our method and the improved version of GAIL we compare against. For each datapoint, the cumulative reward is averaged over 10 trajectories. For the noisy tasks (environments where the expert was trained with added Gaussian noise), we train on {6, 12, 18} expert trajectories with a maximum of 500k iterations of the optimization procedure. For the non-noisy tasks, we trained the different algorithms on a fixed number of trajectories – 12 for Walker, and 6 for Hopper. Table 3 shows the hyperparameters we used for our improved version of GAIL. Empirically, smaller learning rates, large batch sizes, and gradient penalties were critical for the stable convergence of our method.

We note that our version of GAIL requires careful tuning of f UPDATE FREQ. for strong performance. We searched over batch sizes of {1250, 2500, 5000} and selected the one which achieved the most stable updates. In practice, we recommend evaluating a trained policy on a validation set to set this parameter.

PARAMETER	VALUE
BATCH SIZE	*
LEARNING RATE	*
f UPDATE FREQ.	2500
f GRADIENT TARGET	0.4
f GRADIENT PENALTY WEIGHT	10

Table 3: Learner hyperparameters for MM . * indicates the parameter was dependent on the environment.

We also used similar parameters for training SAC, also from the Stable Baselines 3 [Raffin et al., 2021] implementation, as we did for training the expert policy. Table 4 shows the choice of hyperparameters we used for training SAC. We directly added in the optimistic mirror descent optimizers Daskalakis et al. [2017] for both the critic and actor objectives of SAC.

PARAMETER	VALUE
γ	0.98
τ	0.02
TRAINING FREQ.	64
GRADIENT STEPS	64
LEARNING RATE	LINEAR SCHEDULE OF 7.3E-4
POLICY ARCHITECTURE	256 X 2

Table 4: Learning hyperparameters for the SAC component of MM

Table 5 shows the learning hyperparameters for any BC policies used for generating simulated data for the membership oracles. Table 6 shows the number of training steps per task we used for both the baseline and our method.

PARAMETER	VALUE
ENTROPY WEIGHT	0
L2 WEIGHT	0
TRAINING TIMESTEPS	5E5

Table 5: Learner hyperparameters for Behavioral Cloning

D.3 Our Algorithm

In this section, we use **bold text** to highlight sensitive hyperparameters. We use the same network architecture choices as the improved GAIL baseline. For all environments, we generated 100 trajectories of simulated behavior cloning data to use with our method. We found that we had to tune **batch size** and **learning rate** for our improved version of GAIL for each environment.

For all tasks, we rolled out **100 trajectories** from a BC trained network to use with our membership oracle. Table 8 shows how we partitioned our dataset between the BC training set and the expert membership oracle dataset.

D.4 Membership Oracle Parameters

For both \mathcal{M}_{VAR} and \mathcal{M}_{MAX} , we use 5 BC networks in the ensemble. In training the BC imitator network for \mathcal{M}_{RND} , we generate a new set of expert data from the allowed number of expert trajectories for the oracle and included the original expert states combined with the predicted actions on those states from a BC trained policy. We followed the exact same parameters in Table 5 to train the BC imitator.

The choice of μ and β (as defined in the formula for \mathcal{M}_{EXP} , \mathcal{M}_{RND} , \mathcal{M}_{VAR} and \mathcal{M}_{MAX}) are pivotal in achieving successful results for our method. We tuned these parameters by visualizing the distribution

ENV.	TRAINING STEPS
WALKER (NO NOISE)	500000
WALKER (WITH NOISE)	400000
HOPPER (NO NOISE)	400000
HOPPER (WITH NOISE)	400000

Table 6: Number of training steps for the different tasks

ENV.	LEARNING RATE	BATCH SIZE
HOPPER	8E-5	512
WALKER	8E-4	2500

Table 7: Sensitive hyperparameters for our improved version of GAIL

EXPERT SIZE	D_1	D_2
6 TRAJS	5	1
12 TRAJS	10	2
18 TRAJS	14	4

Table 8: Partition of trajectories into D_1 and D_2 based on the number of expert trajectories provided

ENV	PARAMETER	\mathcal{M}_{EXP}	\mathcal{M}_{RND}	\mathcal{M}_{VAR}	\mathcal{M}_{MAX}
WALKER	β	0.1	0.1	0.01	0.1
WALKER	μ	0.33	0.22	0.015	0.35
HOPPER	β	0.8	0.25	0.08	0.5
HOPPER	μ	0.68	0.4	0.05	0.8

Table 9: Membership oracle hyperparameters across different environments

of prefix weights on BC rollouts and expert data and tuning based on the tails of the histograms. Table 9 shows the choice of μ and β values we used for each membership oracle.

For both \mathcal{M}_{VAR} and \mathcal{M}_{MAX} , we use **5 BC networks** in the ensemble. In training the BC imitator network $\widehat{\pi}^{BC}(s)$ for \mathcal{M}_{RND} , we generate a new set of expert data (within the allowed budget of expert trajectories) and include the original expert states combined with the predicted actions on those states from a BC trained policy. We follow the exact same choice of parameters in Table 5 to train the BC imitator network.

D.5 Initial State Shift Experiments

Unless stated below, the parameters for these experiments are the same as those listed above. For these experiments, we used demonstrations generated by an expert trained on the standard Bullet tasks but subject the learner (both at train and test time) to a initial velocity perturbation of a zero-mean Gaussian with variance ($\sigma = 1e - 3$ for Walker and $\sigma = 1e - 7$ for Hopper). We refer interested readers to our code for our precise method of injecting noise as we believe it might be of interest for future experiments. We note that in all the demonstrations, we see the expert start from rest. Despite this relatively small shift, we see BC performance drop significantly, as is characteristic of real-world problems where it significantly under-performs on-policy IL methods. All results are averaged over five seeds.

For all environments, we train BC for $1e5$ steps over the training set (as well as for the query policies for RE). We use the parameters in Table 10 for all environments for MM .

For RE , we train 3 policies and use the max-distance approximate membership oracle. We use the sample splits in Table 11 and the μ and β values in 12 for all experiments. We use the above parameters for MM for our base moment-matcher. We do not use prefix weights for these experiments

as we were able to achieve stable performance without them. We also use the full dataset for moment matching, not just D_2 , as we found this lead to slightly better performance on these experiments.

PARAMETER	VALUE
BATCH SIZE	2048
LEARNING RATE	8E-5
f UPDATE FREQ.	2500
f GRADIENT TARGET	0.4
f GRADIENT PENALTY WEIGHT	10

Table 10: Learner hyperparameters for MM for initial state shift experiments.

EXPERT SIZE	D_1	D_2
6 TRAJS	2	4
12 TRAJS	2	10
18 TRAJS	2	14

Table 11: Partition of trajectories into D_1 and D_2 based on the number of expert trajectories provided for initial state shift experiments.

ENV	PARAMETER	\mathcal{M}_{\max}
WALKER	β	0.01
WALKER	μ	0.01
HOPPER	β	0.01
HOPPER	μ	0.01

Table 12: Membership oracle hyperparameters across different initial state shift environments.