

---

# On the Value of Interaction and Function Approximation in Imitation Learning

---

**Nived Rajaraman**

University of California, Berkeley  
nived.raajaraman@berkeley.edu

**Yanjun Han**

University of California, Berkeley  
yjhan@berkeley.edu

**Lin F. Yang**

University of California, Los Angeles  
linyang@ee.ucla.edu

**Jingbo Liu**

University of Illinois, Urbana-Champaign  
jingbol@illinois.edu

**Jiantao Jiao**

University of California, Berkeley  
jiantao@eecs.berkeley.edu

**Kannan Ramchandran**

University of California, Berkeley  
kannanr@eecs.berkeley.edu

## Abstract

We study the statistical guarantees for the Imitation Learning (IL) problem in episodic MDPs. [22] show an information theoretic lower bound that in the worst case, a learner which can even actively query the expert policy suffers from a suboptimality growing quadratically in the length of the horizon,  $H$ . We study imitation learning under the  $\mu$ -recoverability assumption of [27] which assumes that the difference in the  $Q$ -value under the expert policy across different actions in a state do not deviate beyond  $\mu$  from the maximum. We show that the reduction proposed by [25] is statistically optimal: the resulting algorithm upon interacting with the MDP for  $N$  episodes results in a suboptimality bound of  $\tilde{O}(\mu|\mathcal{S}|H/N)$  which we show is optimal up to log-factors. In contrast, we show that any algorithm which does not interact with the MDP and uses an offline dataset of  $N$  expert trajectories must incur suboptimality growing as  $\gtrsim |\mathcal{S}|H^2/N$  even under the  $\mu$ -recoverability assumption. This establishes a clear and provable separation of the minimax rates between the active setting and the no-interaction setting. We also study IL with *linear function approximation*. When the expert plays actions according to a linear classifier of known state-action features, we use the reduction to multi-class classification to show that with high probability, the suboptimality of behavior cloning is  $\tilde{O}(dH^2/N)$  given  $N$  rollouts from the optimal policy. This is optimal up to log-factors but can be improved to  $\tilde{O}(dH/N)$  if we have a linear expert with *parameter-sharing* across time steps. In contrast, when the MDP transition structure is known to the learner such as in the case of simulators, we demonstrate fundamental differences compared to the tabular setting in terms of the performance of an optimal algorithm, MIMIC-MD (Rajaraman et al. [22]) when extended to the function approximation setting. Here, we introduce a new problem called confidence set linear classification, that can be used to construct sample-efficient IL algorithms.

## 1 Introduction

In many practical sequential decision making problems it is difficult to manually design reward functions that capture the essence of carrying out the task “nicely”. Furthermore, many modern-day

reinforcement learning tasks operate in very large state and action spaces - with sparse reward feedback it is difficult to train good agents without additional feedback or supervision. This motivates the setting of Imitation Learning (IL) where the learner operates in a setting of unknown or unreliable rewards, but with an expert that provides demonstrations as to how to carry out the task in the desirable way. The work of [21] first showed that using expert demonstrations can significantly improve performance in autonomous driving applications. Imitation Learning approaches have found remarkable success in practice over the last decade since expert demonstrations are often available abundantly such as in game AI [13, 1], as well as more recently in autonomous-driving applications such as [5, 20]. In this paper, we study IL in the episodic Markov Decision Process (MDP) formalism.

**Notation:** An MDP  $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, \rho, \mathbf{r})$ , where  $\mathcal{S}$  is the state space,  $\mathcal{A}$  is the action space  $P$  is the MDP transition,  $\rho$  is initial state distribution and  $\mathbf{r}$  is the reward function. The *value*  $J_{\mathbf{r}}(\pi)$  of a policy  $\pi$  is defined as the expected cumulative reward accumulated over an episode of length  $H$ ,  $J_{\mathbf{r}}(\pi) = \mathbb{E}_{\pi}[\sum_{t=1}^H \mathbf{r}_t(s_t, a_t)]$ , where the notation  $\mathbb{E}_{\pi}[\cdot]$  denotes expectation with respect to a random trajectory  $\{(s_1, a_1), \dots, (s_H, a_H)\}$  obtained by rolling out the policy  $\pi = (\pi_1, \dots, \pi_H)$ , where the initial state  $s_1$  is sampled independently from an initial state distribution  $\rho(\cdot)$ . Here  $\pi_t$  denotes the policy at time step  $t$ . In the IL setting, we assume that the underlying reward function is *unknown and unobserved*. The reward function  $\mathbf{r} = \{\mathbf{r}_1, \dots, \mathbf{r}_H\}$  is assumed to be time-variant and pointwise bounded in  $[0, 1]$ , and the transition function  $P = (P_1, \dots, P_{H-1})$  of the MDP is also assumed to be time-variant. The simplest IL setting is the no-interaction setting [22].

**Definition 1** (IL in the no-interaction setting). *The learner is provided an offline dataset  $D$  of  $N$  trajectories (without rewards) drawn by independently rolling out an (unknown) expert policy  $\pi^*$  through the MDP. The learner is not allowed to interact with the MDP.*

The learner’s objective in IL is to construct a policy  $\hat{\pi}$  with small *suboptimality*, defined as the difference in the expert’s and learner’s values:  $J_{\mathbf{r}}(\pi^*) - J_{\mathbf{r}}(\hat{\pi})$ . In this paper we restrict the expert policy to be deterministic and define  $\pi_t^*(s)$  as the action played by the expert at time  $t$  at state  $s$ . An IL instance refers to the tuple  $(\mathcal{M}, \pi^*)$ . The  $Q$ -function of a policy  $\pi$  is defined as the expected reward-to-go,  $Q_t^{\pi}(s, a) = \mathbb{E}_{\pi}[\sum_{t'=t}^H \mathbf{r}_{t'}(s_{t'}, a_{t'}) | s_t = s, a_t = a]$ , and  $f_t^{\pi}(s)$  is defined as the distribution over states induced at time  $t$ , by rolling out the policy  $\pi$ .

Since the expert policy is a collection of actions played at different states visited by the expert, a natural approach to IL is to use any classification algorithm to learn a mapping from states to actions as the learner’s policy. This supervised learning approach has proved to be quite popular in practice and is known as behavior cloning (BC). [25] study BC from a theoretical point of view, and bound the suboptimality of a policy in terms of the 0-1 loss of the resulting classifier. More recently, in the tabular setting, [22] show that BC is statistically optimal in the no-interaction setting, achieving expected suboptimality  $\lesssim \frac{|\mathcal{S}|H^2}{N}$ . This  $H^2$  dependence is known as *error-compounding* and is shown to be necessary even if expert is optimal or the learner can *actively query the expert*.

**Definition 2** (IL in the active setting). *In this setting, the learner is not provided a dataset of expert demonstrations up front. The learner can instead interact with the MDP for  $N$  episodes. While interacting, the learner can query an oracle to return the expert’s action  $\pi_t^*(s)$  at the current state  $s$ .*

It begs the question as to why approaches such as DAGGER ([25]) and AGGREGATE ([26]) which actively query the expert often perform better than BC in practice and to explain this gap, additional assumptions must be imposed. To this end, we look at the minimax lower bound of [22] in the no-interaction setting. The key idea of the lower bound is to include an absorbing “bad” state in the MDP which is never visited in the expert dataset and offers no reward. Any policy which visits this state is doomed to incur a large suboptimality - in the absence of full information, the learner is forced to visit this often. The lower bound instance is pathological in the sense that even if the expert itself visits the bad state, it is never able to “recover” and return to the remaining states. Indeed in practical situations such as driving a car, experts often can recover and collect a high reward even if a mistake is made locally. [25] introduce an assumption to this effect, which we refer to as  $\mu$ -recoverability.

**Definition 3** ( $\mu$ -recoverability). *An IL instance is said to satisfy  $\mu$ -recoverability if for each  $t \in [H]$  and  $s \in \mathcal{S}$ ,  $\mathbb{E}_{a \sim \pi_t^*(\cdot|s)} [Q_t^{\pi^*}(s, a)] - Q_t^{\pi^*}(s, a) \leq \mu$  for all actions  $a \in \mathcal{A}$ . Informally, if the expert plays an “incorrect” action at any state  $s$  at a single time  $t$  and goes back to choosing the correct actions afterwards, the expected reward collected is less by at most  $\mu$ .*

Under the  $\mu$ -recoverability assumption, [25] show that a learner which minimizes the 0-1 loss under the learner’s own state distribution to  $\epsilon$  admits a suboptimality upper bound of  $\mu H \epsilon$ . However, it is

a-priori unclear how small  $\epsilon$  can be made as a function of the number of the size of expert dataset / number of MDP interactions,  $N$  in the no-interaction / active settings. This is a drawback of the reduction approach followed by [25, 27] since it cannot distinguish the power of learners in different interaction models. In this paper, we propose an policy in the active setting with expected 0-1 loss under the learner’s own state distribution bounded by  $|\mathcal{S}|/N$ . This “completes” the reduction in a sense, and establishes suboptimality bounds for the active setting as an explicit function of the number of states  $|\mathcal{S}|$ , interactions  $N$  and horizon  $H$ .

**Informal Theorem 1** (Formal version: Theorem 1 and 2). *In the active setting, under the  $\mu$ -recoverability assumption, there exists a learner  $\hat{\pi}$  which incurs expected suboptimality  $\mathbb{E}[J(\pi^*) - J(\hat{\pi})] \lesssim \mu H |\mathcal{S}|/N$ . Furthermore, if  $N \geq |\mathcal{S}|H$ , for any learner  $\hat{\pi}$  in the active setting, there exists an MDP such that the expected suboptimality  $\mathbb{E}[J(\pi^*) - J(\hat{\pi})] \gtrsim \mu H |\mathcal{S}|/N$ .*

The key challenge for a learner to minimize the 0-1 under its own state distribution is that the learner’s policy changes over the course of optimization. Note however that it is possible to compute an unbiased estimate of the 0-1 loss under the learner’s own state distribution by rolling out *just a single trajectory*. Thus the active sampling model plays a crucial role in this regard. idea is crucial towards constructing the learner policy discussed in Informal Theorem 1. Under the same  $\mu$ -recoverability assumption, we next consider learners in the no-interaction setting. In contrast to the active setting, we show that error compounding is unavoidable for no-interaction learners.

**Informal Theorem 2** (Formal version: Theorem 3). *In the no-interaction setting, for any learner  $\hat{\pi}$  there exists an IL instance which satisfies  $\mu$ -recoverability for  $\mu \geq 1$  such that the expected suboptimality is  $\mathbb{E}[J(\pi^*) - J(\hat{\pi})] \gtrsim H^2 |\mathcal{S}|/N$ .*

This is the first result to establish a clear *separation* in the statistical minimax rate of the suboptimality incurred by learners in the no-interaction setting such as BC, and learners which can interact with the MDP, such as DAGGER [27] and AGGRAVATE [26].

A common theme of the previous bounds in the tabular setting is that the suboptimality necessarily scales linearly in the number of states. In practical RL settings, state and actions spaces are often continuous, and thus additional assumptions are required to carry out efficient learning. In this paper, we study IL with function approximation, in particular in the linear-expert setting.

**Definition 4** (Linear-expert setting). *In this setting, for each  $(s, a, t)$  tuple, the learner is provided a feature representation  $\phi_t(s, a) \in \mathbb{R}^d$ . For each  $t \in [H]$  there exists an unknown vector  $\theta_t^* \in \mathbb{R}^d$  such that  $\forall s \in \mathcal{S}, \pi_t^*(s) = \arg \max_{a \in \mathcal{A}} \langle \theta_t^*, \phi_t(s, a) \rangle$ .*

As we discuss in Remark 1, the linear-expert setting generalizes several known settings such as when the expert is an optimal policy under the linear- $Q^*$  assumption as well as the tabular setting with an optimal expert. We first establish a bound on the expected suboptimality incurred by BC.

**Informal Theorem 3** (Formal version: Theorem 4). *Under the linear-expert setting, the policy  $\hat{\pi}$  returned by BC incurs suboptimality  $J(\pi^*) - J(\hat{\pi}) \lesssim \frac{(d + \log(1/\delta))H^2 \log(N)}{N}$  with probability  $\geq 1 - \delta$ .*

The presence of this error-compounding is not so surprising because the tabular setting with an optimal expert is a special case of the linear-expert setting where [22] show that error compounding is unavoidable for no-interaction learners. In order to break this  $H^2$ -dependence, we introduce a natural variant of the linear-expert setting known as linear-expert setting with parameter sharing.

**Definition 5** (Linear-expert with parameter sharing). *This setting is the same as the linear-expert setting (Definition 4), with the added constraint that for all  $t, \theta_t^* = \theta^*$  is shared across time.*

Our main contribution is to show that in the linear-expert setting with parameter sharing, IL can be reduced to *sequence multi-class linear classification* where we learn linear classifiers from  $\mathcal{S}^H \rightarrow \mathcal{A}^H$ . The supervised learning reduction of [25] posits to learn separate classifiers from  $\mathcal{S} \rightarrow \mathcal{A}$  or  $\mathcal{S} \times [H] \rightarrow \mathcal{A}$ : this fails to account for the shared parameter  $\theta^*$  across time. While in both cases the resulting policy is an ERM classifier, the suboptimality grows quadratically in  $H$  using the supervised learning reduction. In contrast, using the multi-class classification algorithm of [8], we also provide an algorithm  $\hat{\pi}$  with suboptimality *growing linearly in  $H$* .

**Informal Theorem 4** (Formal version: Theorem 5). *Under the linear-expert setting with parameter sharing, there exists a learner  $\hat{\pi}$  with suboptimality  $J(\pi^*) - J(\hat{\pi}) \lesssim \frac{(d + \log(1/\delta))H \log(N)}{N}$  with probability  $\geq 1 - \delta$ .*

With the additional linearity assumption on the expert, the learner can potentially infer the expert’s action on states that are *not observed* in the dataset. However in the absence of transition information or the parameter sharing assumption, a learner cannot even distinguish between different actions at the remaining states, which is what leads to catastrophic error compounding. To remedy this issue, we borrow from the work of [22, 23] who study IL in the *known-transition setting* in tabular MDPs where the learner exactly knows the Markov transition kernel and the initial state distribution of the MDP. The motivation for this setting stems from autonomous driving applications where policies are often learned in a simulated environment prior to fine-tuning in the real world [9, 35] and in the simulator the rewards functions are still difficult to specify. In such settings, the state and action spaces are indeed unbounded, which makes it ideal to study through the frame of function approximation.

**Definition 6** (IL in the known-transition setting). *The learner is provided an dataset  $D$  of  $N$  trajectories (without rewards) drawn by independently rolling out the expert policy  $\pi^*$  through the MDP. The learner also knows the MDP transition  $P$  and initial state distribution  $\rho$  exactly.*

In the tabular setting, the known transition setting has an interesting landscape: it is known from [22] that the quadratic- $H$  barrier can be broken - the authors propose the MIMIC-MD algorithm which achieves an expected suboptimality upper bound of  $|\mathcal{S}|H^{3/2}/N$  and this dependence on the horizon is optimal [23]. The key idea is that with access to the MDP transition structure, as long as the visited states are conditioned to be observed in the dataset (so the expert’s action is known), the learner can simulate artificial trajectories *according to the expert’s policy* to generate more training data.

While the approach of simulating artificial trajectories in MIMIC-MD achieves the minimax optimal bounds here, a natural question to ask is whether the approach is tailored to work only in the tabular setting. Indeed in the presence of continuous state spaces, the learner may observe but a measure-0 subset of the state-space in the expert dataset. In spite of this, to apply the approach of simulating artificial trajectories, the learner must be able to *infer the expert’s actions on a large fraction of the state-space*. To this end, in the known-transition setting, we propose a problem known as *confidence set linear classification* which extends multi-class linear classification and we prove that algorithms with small expected loss for confidence set linear classification can be used to construct policies with small suboptimality, using the approach of simulating artificial trajectories. At a high level, the objective of the learner is to not only output a classifier, but also a set of inputs (confidence set) where the classifier *certifiably* outputs the correct label.

**Definition 7** (Confidence set linear classification). *Consider a classification problem on  $X$  with input distribution  $\rho_X$ , output space  $Y$ , with features  $\phi : X \times Y \rightarrow \mathbb{R}^d$  and a dataset  $\mathcal{D}$  of  $N$  examples drawn i.i.d. as  $x_i \sim \rho_X$  and  $y_i \sim h^*(x_i)$ , where  $h^*$  is an unknown linear multi-class classifier mapping  $x \mapsto \arg \max_{y \in Y} (\theta^*, \phi(x, y))$ . Given the dataset  $\mathcal{D}$ , a confidence set linear classifier returns a tuple  $(\hat{h}, \mathcal{X})$  where  $\hat{h}$  is any classifier from  $X \rightarrow Y$  and  $\mathcal{X} \subseteq X$  is a measurable set of inputs (known as the confidence set) such that  $\forall x \in \mathcal{X}, \hat{h}(x) = h^*(x)$ . The learner’s objective is to minimize the expected loss  $\mathbb{E}[1 - \rho_X(\mathcal{X})]$ .*

Sample-efficient confidence set linear classification algorithms can be used to construct learners with small suboptimality. We prove such a reduction in the linear-expert setting with linear rewards: here, in addition to the linear-expert setting, for each  $t \in [H]$ , the reward function  $\mathbf{r}_t(\cdot, \cdot)$  is also constrained to be a linear function of the feature representations  $\phi_t(\cdot, \cdot)$ . The linear reward setting was considered previously in the known transition setting in [3]. We formally define it in Definition 8.

**Informal Theorem 5** (Formal version: Theorem 6). *Consider any algorithm  $\text{Alg}$  for confidence set linear classification and define  $\mathfrak{R}_{N,d}(\rho_X, Y, \psi)$  as the expected loss (Definition 7) incurred by  $\text{Alg}$  when (i) the input distribution is  $\rho_X$ , (ii) features are  $\psi : X \times Y \rightarrow \mathbb{R}^d$  and (iii) the learner is provided a dataset of  $N$  samples (with labels from an unknown multi-class linear classifier). In the linear-expert setting with linear rewards, there exists a learner policy  $\hat{\pi}$  with expected suboptimality:*

$$\mathbb{E}[J(\pi^*) - J(\hat{\pi})] \lesssim H^{3/2} \sqrt{\frac{d \sum_{t=1}^H \mathfrak{R}_{N,d}(f_t^{\pi^*}, \mathcal{A}, \phi_t)}{N H}} \quad (1)$$

Informal Theorem 5 shows that it suffices to find good algorithms for confidence set linear classification and bound  $\mathfrak{R}_{d,N}(\rho_X, Y, \phi)$  to carry out sample efficient IL. However, even in the case of binary output space  $Y = \{0, 1\}$  and uniformly distributed features, the answer to this question is quite challenging and admits a non-standard rate.

**Informal Theorem 6** (Formal version: Theorem 7). *Consider an instance of confidence set linear classification where  $Y = \{0, 1\}$ ,  $\rho_X = \text{Unif}(\mathbb{S}^{d-1})$  and  $\phi(x, 0) = -\phi(x, 1) = x/2 \in \mathbb{R}^d$ . Then, for sufficiently large  $N$ ,*

- (i) *For any algorithm  $\frac{d^{3/2}}{N\sqrt{\log(d)}} \lesssim \mathfrak{R}_{d,N,\mathcal{A}}(\rho_X, Y, \phi)$ .*
- (ii) *There exists an algorithm such that  $\mathfrak{R}_{d,N,\mathcal{A}}(\rho_X, Y, \phi) \lesssim \frac{d^{3/2} \log(d)}{N}$ .*

This result shows that the minimax risk for confidence set linear classification necessarily grows as  $\gtrsim d^{3/2}/N$ . This rate establishes a fundamental difference between function approximation and tabular settings. In the tabular setting, where the features for each state-action pair are orthogonal, the learner cannot conclude the labels at unobserved states. Thus, the minimax risk of confidence set linear classification corresponds to the expected probability mass on unobserved states, which is also known as the *missing mass* [18]. It is known from [22, Lemma A.20] that the expected missing mass is  $\lesssim |\mathcal{S}|/N \equiv d/N$  under any distribution over states and binary action space. Informal Theorem 6 establishes a fundamental difference between the tabular setting and the linear-expert setting with linear rewards, in terms of the suboptimality guarantees achieved by the approach of simulating artificial trajectories.

While the upper bound in Informal Theorem 6 (ii) only applies in the special case of binary classification with uniformly distributed features, we conjecture that the minimax expected loss for confidence set linear classification is  $\frac{d^{3/2}}{N}$ . If this conjecture is true, then Informal Theorem 6 shows that there exists a learner  $\hat{\pi}$  in the known-transition setting with linear-expert and linear rewards such that the expected suboptimality is  $\lesssim H^{3/2}d^{5/4}/N$ . For sufficiently large  $H$ , this improves the  $\Theta(dH^2/N)$  of BC and achieves the optimal dependence on the horizon [23].

## 1.1 Related Work

There is a long line of history studying the IL problem, [2, 34, 24, 37, 25, 10, 11, 19, 14]. A number of algorithms target the  $H^2$  error compounding issue, [12, 15, 17, 4, 36]. [25, 27, 4] study IL in the reduction framework, where they reduce the IL problem to a supervised learning problem and study the how the supervised learning error translates to the IL error. DAGGER [27], AGGREGATE [26], and AGGREGATED [30] learn policies by actively interacting with the environment and the expert during training. Going beyond the tabular setting, [3, 33] study IL in the presence of linear function approximation. [31] study IL in a setting where expert actions are not observed in the dataset. [6, 16] analyze DAGGER and dynamic regret under some regularity conditions, in comparison with the static regret reductions of [27], [7] propose a policy learning method called LOKI based on bootstrapping policy gradient methods using IL. [32] establish that many proposed algorithms for imitation learning indeed carry out some form of matching of distributions or features (namely, “moment matching” as studied in the paper). Our proposed algorithm for the linear-expert setting also lies in this framework.

## 2 IL with $\mu$ -recoverability

As defined in Definition 3, the  $\mu$ -recoverability assumption captures the ability of an expert to recover and collect a high reward at a state even upon locally deviating from its action distribution at states. The reduction in [27, Theorem 2] shows that under  $\mu$ -recoverability, a learner policy  $\hat{\pi}$  which minimizes the 0-1 loss with respect to the expert’s policy under the learner’s own state distribution,

$$\mathcal{L}(f^{\hat{\pi}}, \hat{\pi}, \pi^*) \triangleq \frac{1}{H} \sum_{t=1}^H \mathbb{E}_{s \sim f_t^{\hat{\pi}}(\cdot)} \left[ \mathbb{E}_{a \sim \hat{\pi}_t(\cdot|s)} [\mathbb{1}(a \neq \pi_t^*(s))] \right]. \quad (2)$$

to be less than  $\epsilon$ , incurs suboptimality upper bounded by  $\mu H \epsilon$ . However, in the active setting, it is a-priori unclear how small  $\epsilon$  can be made as a function of the number of times the learner interacts with the MDP,  $N$ . We address this question in the following theorem.

**Theorem 1.** *In the active setting it is possible to construct a learner policy  $\hat{\pi}$  such that  $\mathbb{E}[\mathcal{L}(f^{\hat{\pi}}, \hat{\pi}, \pi^*)] \lesssim |\mathcal{S}|/N$ . Furthermore, under  $\mu$ -recoverability,  $J(\pi^*) - \mathbb{E}[J(\hat{\pi})] \lesssim \mu |\mathcal{S}| H/N$ .*

Following the no-regret reduction of [27], it suffices for the learner to find a sequence of policies  $\hat{\pi}^1, \dots, \hat{\pi}^T$  such that the *online-learning regret*,

$$\frac{1}{N} \sum_{i=1}^N \mathcal{L}(f^{\hat{\pi}^i}, \hat{\pi}^i, \pi^*) - \min_{\pi} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(f^{\hat{\pi}^i}, \pi, \pi^*) \lesssim \frac{|\mathcal{S}|}{N}. \quad (3)$$

Then, the mixture policy  $\frac{1}{N} \sum_{i=1}^N \hat{\pi}^i$  satisfies  $\mathcal{L}(f^{\hat{\pi}}, \hat{\pi}, \pi^*) \lesssim \frac{|\mathcal{S}|}{N}$ . Note that in eq. (3), the oracle loss  $\min_{\pi} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(f^{\hat{\pi}^i}, \pi, \pi^*)$  is in fact 0, achieved by  $\pi = \pi^*$ . Suppose for each  $i$ , the learner rolls out a single trajectory according to  $\hat{\pi}^i$ . Denoting the empirical state-visitation distribution  $\hat{f}^{\pi^i} = (\hat{f}_1^{\pi^i}, \dots, \hat{f}_H^{\pi^i})$ , observe that,  $\frac{1}{N} \sum_{i=1}^N \mathcal{L}(\hat{f}^{\pi^i}, \hat{\pi}^i, \pi^*)$  is an unbiased estimate of  $\frac{1}{N} \sum_{i=1}^N \mathcal{L}(f^{\hat{\pi}^i}, \hat{\pi}^i, \pi^*)$  if  $\hat{\pi}^i$  is a measurable function the first  $i-1$  rolled out trajectories (according to  $\hat{\pi}^1, \dots, \hat{\pi}^{i-1}$ ). Thus, it suffices for the learner to find a sequence of policies  $\hat{\pi}^1, \dots, \hat{\pi}^T$  which minimize the *empirical online-learning regret*:  $\frac{1}{N} \sum_{i=1}^N \mathcal{L}(\hat{f}^{\hat{\pi}^i}, \hat{\pi}^i, \pi^*) - \min_{\pi} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(\hat{f}^{\hat{\pi}^i}, \pi, \pi^*)$  to be  $\lesssim \frac{|\mathcal{S}|}{N}$ . As we discuss in more detail in the Appendix, it is possible to construct a sequence of policies  $\hat{\pi}^1, \dots, \hat{\pi}^N$  using entropy-regularized mirror descent [29] which minimizes the empirical online-learning regret to be  $\lesssim |\mathcal{S}|/N$ . The resulting policy  $\hat{\pi} = \frac{1}{N} \sum_{i=1}^N \hat{\pi}^i$  minimizes the expected 0-1 loss under its own state distribution to be  $\lesssim |\mathcal{S}|/N$  in expectation. The guarantee on the expected suboptimality of this policy directly follows from [27, Theorem 2] under  $\mu$ -recoverability.

This suboptimality guarantee is optimal for any learner in the active setting. The lower bound instance essentially follows from that of [22] for the active tabular setting where if  $N \geq |\mathcal{S}|H$ , the expected suboptimality incurred is  $\gtrsim \frac{|\mathcal{S}|H^2}{N}$ . By scaling each reward by a factor of  $\mu/H$ , the same family of IL instances now satisfies  $\mu$ -recoverability and results in the lower bound for active learners.

**Theorem 2.** *In the active setting, if  $N \geq |\mathcal{S}|H$ , every learner  $\hat{\pi}$  incurs expected suboptimality  $\mathbb{E}[J(\pi^*) - J(\hat{\pi})] \gtrsim \min\{\mu, \mu|\mathcal{S}|H/N\}$ .*

Now, under the same  $\mu$ -recoverability assumption, we study learners in the no-interaction setting. We prove a lower bound that in the worst case, error compounding is unavoidable for such learners.

**Theorem 3.** *For  $|\mathcal{S}| \geq 3$  and  $|\mathcal{A}| \geq H$ , for any learner  $\hat{\pi}$ , in the no-interaction setting, there exists an IL instance which satisfies  $\mu$ -recoverability for  $\mu \geq 1$  such that the expected suboptimality incurred by the learner is lower bounded,  $\mathbb{E}[J(\pi^*) - J(\hat{\pi}(D))] \gtrsim \min\{H, |\mathcal{S}|H^2/N\}$ .*

The lower bound we consider is a modification of the lower bound of [22] where the MDP is constructed to have a “bad” state in the MDP never visited by the expert. We modify the instance to add a single “recovery” action at the bad state; the instance now satisfies  $\mu$ -recoverability for any  $\mu \geq 1$ . If the number of actions are large  $|\mathcal{A}| \geq H$ , any no-interaction learner still fails to identify the recovery action with constant probability. In essence this reduces the instance to the lower bound of [22] and any no-interaction learner incurs an expected suboptimality  $\gtrsim \min\{H, |\mathcal{S}|H^2/N\}$ .

The classical reduction formulations of [25, 27] prove upper bounds for IL based on minimizing a certain surrogate objective. However the statistical rate of minimizing different surrogate objectives as a function of the number of interactions (active setting) / size of the expert dataset (no-interaction setting) is unclear. As we show here, with  $\mu$ -recoverability, the surrogate objective of 0-1 loss under the learner’s policy can be minimized to  $|\mathcal{S}|/N$  in the active setting, but this is impossible in the no-interaction setting. Going beyond the reduction formulation, we thus distinguish between the statistical power of learners under different interaction models.

### 3 Linear function approximation in the no-interaction setting

In this section, we go beyond the tabular setting and study IL in the presence of function approximation. In practical settings, state and action spaces are often continuous or unbounded and carrying out efficient IL requires imposing additional assumptions. In this section we study the linear-expert setting (Definition 4) where  $\mathcal{S}$  and  $\mathcal{A}$  may be unbounded, but the learner is provided a set of feature representations of state-actions, and the expert policy is constrained to be realizable by a unknown linear (in the feature representations) classifier. The linear-expert setting generalizes several known settings as we discuss in the following remark.

**Remark 1.** *The linear-expert setting (Definition 4) generalizes the linear- $Q^*$  setting with an optimal expert. Under this assumption, the optimal expert policy plays actions according to  $\pi_t^*(s) = \arg \max_{a \in \mathcal{A}} Q_t^*(s, a) = \arg \max_{a \in \mathcal{A}} \langle \theta_t^*, \phi_t(s, a) \rangle$  for an unknown  $\theta_t^* \in \mathbb{R}^d$ . Thus the expert policy is realizable by a linear multi-class classifier. Since the tabular setting is a special case of the linear- $Q^*$  setting with  $d = |\mathcal{S}||\mathcal{A}|$ , with features for each  $t$  chosen as the standard basis vectors in  $\mathbb{R}^d$ , the linear-expert setting with  $d = |\mathcal{S}||\mathcal{A}|$  generalizes the tabular setting with an optimal expert.*

In the tabular setting, it is known that the expected suboptimality of behavior cloning is  $O(|\mathcal{S}|H^2/N)$  in the worst case which is minimax optimal [22]. We first establish an upper bounds on the suboptimality incurred by BC in the linear-expert setting.

**Theorem 4.** For  $t = 1, \dots, H$ , denote  $(D)_t$  as a collection of  $N$  state-action pairs visited at time  $t$  across trajectories in  $D$ . Consider a learner policy which trains a policy  $\hat{\pi}$  using BC as follows: for each  $t = 1, \dots, H$ , the learner trains a linear multi-class classifier  $\hat{h}_t : \mathcal{S} \rightarrow \mathcal{A}$  on the dataset  $(D)_t$  using the algorithm of [8] and plays the policy  $\hat{\pi}_t(s) = \hat{h}_t(s)$ . Then, in the linear-expert setting, with probability  $1 - \delta$ , the suboptimality of  $\hat{\pi}$  is upper bounded by  $J(\pi^*) - J(\hat{\pi}) \lesssim \frac{H^2(d+\log(1/\delta))\log(N)}{N}$ .

This result is in fact a special case of Theorem 5 where we prove guarantees in the *linear-expert setting with parameter sharing*, where the expert plays according to the same linear classifier shared across the episode. In Remark 2, we show that the linear-expert setting with parameter sharing with dimension  $dH$  generalizes the linear-expert setting with dimension  $d$ .

**Remark 2.** *Linear-expert setting with dimension  $d$  is a special case of the linear-expert setting with parameter sharing, with dimension  $dH$ . Define  $\theta^* = (\theta_1^*, \dots, \theta_H^*) \in \mathbb{R}^{dH}$  and  $\phi'_t(s, a) = (0^d, \dots, \phi_t(s, a), \dots, 0^d) \in \mathbb{R}^{dH}$  where  $\phi_t(s, a)$  is embedded in the coordinates  $td + 1$  to  $(t + 1)d$ . Then,  $\pi_t^*(s) = \arg \max_{a \in \mathcal{A}} \langle \theta_t^*, \phi_t(s, a) \rangle = \arg \max_{a \in \mathcal{A}} \langle \theta^*, \phi'_t(s, a) \rangle$ , satisfying Definition 5.*

### 3.1 Linear-expert with parameter sharing: Reducing IL to sequence classification

In this section, we demonstrate a reduction of IL to sequence multi-class linear classification from  $\mathcal{S}^H \rightarrow \mathcal{A}^H$ , in contrast to BC which learns a classifier from  $\mathcal{S} \rightarrow \mathcal{A}$ . First note that the expert’s policy can be thought of as a classifier from  $\mathcal{S}^H \rightarrow \mathcal{A}^H$ : for each input sequence of states  $(s_1, s_2, \dots, s_H)$  the expert “classifier” outputs the sequence of actions  $(\pi_1^*(s_1), \pi_2^*(s_2), \dots, \pi_H^*(s_H))$ . The learner obtains  $N$  i.i.d. trajectories from the expert which are examples in the classification training dataset, and the objective is to predict actions for each new trajectory from the expert. Define  $\Theta$  as the set of linear multi-class classifiers for sequences, of the form

$$\mathcal{S}^H \ni (s_1, \dots, s_H) \mapsto \arg \max_{a_1, \dots, a_H \in \mathcal{A}} \left\langle \theta, \sum_{t=1}^H \phi_t(s_t, a_t) \right\rangle \in \mathcal{A}^H. \quad (4)$$

for  $\theta \in \mathbb{R}^d$ . Note that under the linear-expert assumption with parameter sharing, the expert’s policy can be identified as a classifier in the family described above. At each state  $s$ , the expert plays the action according to  $\arg \max_{a \in \mathcal{A}} \langle \theta^*, \phi_t(s, a) \rangle$  at time  $t$ . Summing over any sequence of states  $s_1, \dots, s_H$ , the expert’s policy therefore satisfies  $(\pi_1^*(s_1), \dots, \pi_H^*(s_H)) = \arg \max_{a_1, \dots, a_H} \langle \theta^*, \sum_{t=1}^H \phi_t(s_t, a_t) \rangle$ .

Note that classifiers of the form eq. (4) indeed correspond to meaningful (Markovian) policies which drawn actions  $a_t$  from a policy which is a function of only the current state  $s_t$ . Indeed the map in eq. (4) is separable as  $\sum_{t=1}^H \arg \max_{a_t \in \mathcal{A}} \langle \theta, \phi_t(s_t, a_t) \rangle$  where we carry out the optimization for each variable  $a_1, \dots, a_H$  separately. By contradiction, the action played by the classifier at any state  $s_t$  at time  $t$  must be  $\arg \max_{a_t \in \mathcal{A}} \langle \theta, \phi_t(s_t, a_t) \rangle$  which is Markovian. Finally, we prove a bound on the suboptimality of the policy induced by  $\hat{\theta}$  by the expected 0-1 loss of  $\hat{\theta}$ . The intuition is that in any trajectory where the learner’s actions exactly match the expert’s actions, no suboptimality is incurred. In contrast, in any trajectory where the learner plays an action different from the expert at some time, the suboptimality incurred is  $\leq H$ .

**Lemma 1.** Consider any linear multi-class classifier  $\hat{\theta} : \mathcal{S}^H \rightarrow \mathcal{A}^H$  (in eq. (4)) with expected 0-1 loss,  $\mathbb{E}_{\pi^*} [\mathbb{1}(\hat{\theta}(s_1, \dots, s_H) \neq (a_1, \dots, a_H))] \leq \gamma$ . Then, the policy  $\hat{\pi}$  corresponding to the linear classifier  $\hat{\theta}$ , satisfies  $J(\pi^*) - J(\hat{\pi}) \leq H\gamma$ .

**Remark 3.** The linear dependence on the horizon can be interpreted in a different way: with parameter sharing, the learner can aggregate information across time steps in an episode to learn a single linear classifier with improved guarantees. The amount of training data the learner has access to is effectively larger by a factor of  $H$  since each trajectory provides  $H$  samples of data for learning a single classifier common across time. The reduction analysis of [25] shows a suboptimality gap of  $H^2\epsilon$  for learners with expected 0-1 loss under the expert state distribution upper bounded by  $\epsilon$ . If data can be aggregated across time to learn a single parameter, the expected 0-1 loss can be brought down by a factor of  $H$ , showing that it is possible to achieve a suboptimality bound with linear dependence on the length of the horizon.

[8] provide a compression based algorithm for linear multi-class classification in the realizable setting. Indeed, invoking [8, Theorem 5], it is possible to learn a linear classifier  $\hat{\theta} \in \Theta$  such that the expected 0-1 loss of the classifier is upper bounded by  $\frac{(d+\log(1/\delta)) \log(N)}{N}$  given  $N$  expert trajectories. In conjunction with Lemma 1 this results in an upper bound on the suboptimality of the resulting policy.

**Theorem 5.** *Consider a learner  $\hat{\pi}$  which trains a classifier  $\hat{\theta}$  from the family in eq. (4) on the expert dataset using the compression based learner of [8], and at each time  $t$  and state  $s$ ,  $\hat{\pi}_t(s) = \arg \max_{a \in \mathcal{A}} (\hat{\theta}, \phi_t(s, a))$ . Under the linear-expert assumption with parameter sharing, with probability  $\geq 1 - \delta$ , the suboptimality of the learner’s policy satisfies  $J(\pi^*) - J(\hat{\pi}) \lesssim \frac{H(d+\log(1/\delta)) \log(N)}{N}$ .*

## 4 Linear function approximation under known-transition assumption

The framework of [23] shows that IL under the known-transition setting can be reduced to the problem of *uniform expert value estimation*: the problem of estimating the value of the expert policy under all reward functions. The authors show that given a uniform expert value estimator  $\tilde{J}_{\mathbf{r}}(\pi^*)$ , which with probability  $1 - \delta$  (over the expert dataset and external randomness) for all reward functions  $\mathbf{r}$ , satisfies  $|J_{\mathbf{r}}(\pi^*) - \tilde{J}_{\mathbf{r}}(\pi^*)| \leq \epsilon$ , then the policy  $\hat{\pi}$  output by the following optimization problem,

$$\hat{\pi} \leftarrow \arg \min_{\pi} \max_{\mathbf{r}} \tilde{J}_{\mathbf{r}}(\pi) - J_{\mathbf{r}}(\pi) \quad (\text{OPT})$$

incurs suboptimality  $J(\pi^*) - J(\hat{\pi}) \leq 2\epsilon$  with the same probability  $1 - \delta$ . [23] also show that this objective is a convex program and can be efficiently solved approximately in the tabular setting.

In this context, to execute the approach of simulating artificial trajectories, observe that a learner can construct a good estimate of the expert’s value under some reward function  $\mathbf{r}$  by decomposing it as the sum of two parts:

$$J_{\mathbf{r}}^1(\pi^*) = \mathbb{E}_{\pi^*} \left[ \sum_{t=1}^H \mathbf{r}_t(s_t, a_t) \mathbb{1}(\mathcal{E}) \right], \text{ and } J_{\mathbf{r}}^2(\pi^*) = \mathbb{E}_{\pi^*} \left[ \sum_{t=1}^H \mathbf{r}_t(s_t, a_t) \mathbb{1}(\mathcal{E}^c) \right]. \quad (5)$$

where  $\mathcal{E}$  is the event that the all the states  $(s_1, \dots, s_H)$  visited in the trajectory are observed in the expert dataset. The first term,  $J_{\mathbf{r}}^1$ , can be estimated to an arbitrary level of accuracy for any reward function  $\mathbf{r}$  by rolling out many artificial trajectories using  $\pi^*$ , known at all states observed in the dataset. The remaining term,  $J_{\mathbf{r}}^2$  can be tackled using a simple empirical estimate, as explained below. The event  $\mathcal{E}^c$  guarantees that states in a trajectory are visited are observed in the dataset  $D$  (where the expert’s policy is known). Therefore, by holding out some trajectories in the dataset, the learner may carry out an empirical estimate of  $J_{\mathbf{r}}^2(\pi^*)$  using these trajectories. The error in uniform value estimation precisely stems from the error incurred by the empirical estimate, which is shown to be  $O(|\mathcal{S}|H^{3/2}/N)$  in [22], translating to the suboptimality of the policy  $\hat{\pi}$  in (OPT).

It is a natural question to ask whether this approach of simulating artificial trajectories can be applied when state and action spaces may be unbounded. To effectively use such an approach, the learner should be able to infer the expert’s action at a large fraction of states in spite of *observing the expert’s actions only on a measure-0 subset of states*. We show that if the learner is able to identify the expert’s policy on a known large measure of states (under the expert’s state distribution), then the approach of simulating artificial trajectories can be employed to give a policy with small suboptimality. We establish such a reduction under the linear-expert setting with an additional assumption on the linearity of rewards which we introduce below. The linear reward setting was first introduced in [3] in the known-transition and discounted setting. Here, imposing the linear reward assumption enables the learner to construct linear estimates of the expert value function, which is otherwise not possible.

**Definition 8** (Linear reward assumption). *Define  $\mathcal{R}_{\text{lin}}$  as the family of reward functions which take the form of an unknown linear function of the known feature representation of states. Namely  $\mathcal{R}_{\text{lin}} = \{ \{ \mathbf{r}_t(s, a) = \langle \omega_t, \phi_t(s, a) \rangle : t \in [H], s \in \mathcal{S}, a \in \mathcal{A} \} : \forall t \in [H], \omega_t \in \mathbb{R}^d, \|\omega_t\|_{\infty} \leq 1 \}$ . The features are assumed to satisfy  $\|\phi_t(s, a)\|_1 \leq 1$ . The linear reward assumption assumes the true reward function of the MDP belongs to  $\mathcal{R}_{\text{lin}}$ .*

We propose an extension of MIMIC-MD to the linear-expert setting with linear rewards. The algorithm is based on identifying a set of states  $\mathcal{X}_1, \dots, \mathcal{X}_H$  on which the expert’s policy is exactly known. The learner then constructs a uniform expert value estimator by simulating artificial trajectories using the expert policy conditioned on visiting only these states and uses an empirical estimate of the reward on the remaining states. The final policy is output using the minmax optimization problem in (OPT). A formal description is provided in Algorithm 1.



---

**Algorithm 1** MIMIC-MD under linear-expert and linear rewards assumption

- 1: **Input:** A dataset  $D$  of  $N$  expert policy rollouts; MDP transition  $P$ ; feature representations  $\{\phi_t(s, a) : t \in [H], s \in \mathcal{S}, a \in \mathcal{A}\}$ ; confidence set linear classification algorithm **Alg**.
- 2: Pick a uniformly random permutation of the trajectories of  $D$  and assign the first  $N/2$  as  $D_0$  and the remaining trajectories as  $D_1$ .
- 3: **For**  $t = 1, \dots, H$ : define  $(\hat{h}_t, \mathcal{X}_t(D_0))$  as the output of **Alg** $((D_0)_t)$ 
  - ▷  $(D_0)_t$  are state-action pairs at time  $t$  across trajectories in  $D_0$ .
  - ▷  $\hat{h}_t$  is a classifier from  $\mathcal{S} \rightarrow \mathcal{A}$  and can be identified as a policy.
  - ▷  $\mathcal{X}_t(D_0)$  captures a set of states on which expert's action is certifiably known.
- 4: Define event  $\mathcal{E}_{D_0} = \{\forall t \in [H], s_t \in \mathcal{X}_t(D_0)\}$ : all states in a trajectory belong to  $\{\mathcal{X}_t(D_0)\}_{t=1}^H$ .
- 5: Define the expert value estimator,

$$\tilde{J}_{\mathbf{r}}(\pi^*) = \mathbb{E}_{\pi^*} \left[ \sum_{t=1}^H \mathbf{r}_t(s_t, a_t) \mathbb{1}(\mathcal{E}_{D_0}) \right] + \mathbb{E}_{\text{tr} \sim \text{Unif}(D_1)} \left[ \sum_{t=1}^H \mathbf{r}_t(s_t, a_t) \mathbb{1}(\mathcal{E}_{D_0}^c) \right] \quad (6)$$

▷ The estimator is measurable: the first term can be estimated by rolling out many trajectories using the policy  $(\hat{h}_1, \dots, \hat{h}_t)$ , equal to  $(\pi_1^*, \dots, \pi_H^*)$  under the measurable event  $\mathcal{E}_{D_0}$

- 6: **Output:** Return  $\hat{\pi} \leftarrow \arg \min_{\pi} \max_{\mathbf{r} \in \mathcal{R}_{\text{lin}}} \tilde{J}_{\mathbf{r}}(\pi^*) - J_{\mathbf{r}}(\pi)$ . ▷  $\mathcal{R}_{\text{lin}}$  is defined in Definition 8
- 

**Theorem 6.** *The expected suboptimality of the policy  $\hat{\pi}$  returned by Algorithm 1 under the linear-expert setting with linear rewards can be upper bounded by,*

$$\mathbb{E}[J(\pi^*) - J(\hat{\pi})] \lesssim H^{3/2} \sqrt{\frac{d \sum_{t=1}^H \mathbb{E}[\text{Pr}_{\pi^*}(s_t \notin \mathcal{X}_t(D_0))]}{N H}} \quad (7)$$

Note that for each  $t = 1, \dots, H$ , the probability  $\text{Pr}_{\pi^*}(s_t \notin \mathcal{X}_t(D_0))$  is the loss (as defined in Definition 7) of the confidence set linear classifier **Alg** $((D_0)_t)$  in Algorithm 1.

As a consequence of Theorem 6, it suffices to upper bound the loss of the confidence set linear classification algorithm **Alg**. However, it is quite a challenging problem to compute the minimax risk for confidence set linear classification. Below, we discuss the case of binary classification.

#### 4.1 Confidence set linear classification with binary outputs

In this section, we study confidence set linear classification when the output space  $Y = \{0, 1\}$  is binary. Denote the dataset  $D$  provided to the learner as  $\{(x_1, y_1), \dots, (x_n, y_n)\}$  where  $y_i = 0$  if  $\langle \theta^*, \phi(x, 0) \rangle \geq \langle \theta^*, \phi(x, 1) \rangle$  and 1 otherwise for some unknown  $\theta^* \in \mathbb{R}^d$ . For each sample  $x_i$  observed in the dataset, the learner can conclude that  $\langle \theta^*, \phi(x_i, 0) \rangle - \langle \theta^*, \phi(x_i, 1) \rangle$  is non-negative if  $y_i = 0$  and is non-positive if  $y_i = 1$ . In other words, for each  $x_i$ , the learner can conclude that  $\langle \theta^*, \phi(x, y) - \phi(x, 1 - y) \rangle \geq 0$ . Incorporating the information from all samples in the dataset, the learner can localize  $\theta^*$  to a cone  $\Theta = \{\theta \in \mathbb{R}^d : \forall (x, y) \in D, \langle \theta, \phi(x, y) - \phi(x, 1 - y) \rangle \geq 0\}$ . This cone captures the maximum amount of information the learner can discern about  $\theta^*$ . Indeed, every linear classifier  $\theta \in \Theta$  correctly classifies every sample  $x$  observed in the dataset as the correct label 0 or 1 observed in the dataset. Given the cone  $\Theta$  which captures the uncertainty in  $\theta^*$ , one can construct a set of inputs  $\mathcal{C}$  such that each classifier  $\theta \in \Theta$  is consistent with their labelling of inputs in  $\mathcal{C}$ . We prove that largest such  $\mathcal{C}$  can be directly constructed from the dataset  $D$  as  $\mathcal{C} = \mathcal{K} \cup -\mathcal{K}$ , where  $\mathcal{K}$  is the conical hull of the set of points  $(\phi(x_i, 0) - \phi(x_i, 1))(-1)^{y_i}$  for  $i = 1, \dots, n$  (Figure 1).

**Lemma 2.** *The set  $\mathcal{C} \subseteq X$  as defined above satisfies the following two properties:*

- (i) For each  $x \in \mathcal{C}$ ,  $\text{sign}(\langle \theta, \phi(x, 0) - \phi(x, 1) \rangle) = \text{sign}(\langle \theta^*, \phi(x, 0) - \phi(x, 1) \rangle)$ .
- (ii) For any  $x \notin \mathcal{C}$ , any classifier  $\hat{h}$ , there exists  $\theta \in \Theta$  s.t.  $\hat{h}(x) \neq \mathbb{1}(\langle \theta, \phi(x, 0) - \phi(x, 1) \rangle < 0)$ .

Therefore, the learner can guarantee that the label was correctly predicted for any  $x \in \mathcal{C}$ . More importantly, Lemma 2 (ii) shows that for any classifier  $\hat{h}$ ,  $\mathcal{C}$  is indeed the largest set of inputs for which the learner can guarantee to correctly predict the same output as the true classifier  $\theta^*$ . Thus, in the case of  $\mathcal{A} = \{0, 1\}$ ,  $\mathbb{E}[\rho_X(\mathcal{C}^c)]$  is the minimum expected loss of confidence set linear classification.

Next, we study the special case where the input space  $X$  is the unit sphere  $\mathbb{S}^{d-1}$ , the distribution over inputs  $\rho_X$  is uniform over  $X$ , and the feature  $\phi(x, 0) = -\phi(x, 1) = x/2 \in \mathbb{R}^d$ . Then




Figure 1:  $\{x_1, \dots, x_N\}$  is the set of points observed in the dataset. For each  $x \in \mathbb{S}^{d-1}$ ,  $\Psi_x \triangleq (-1)^{y_i}(\phi(x, 0) - \phi(x, 1))$ . The gray region,  $\mathcal{K}$  is the conical hull of  $\{\Psi_{x_1}, \dots, \Psi_{x_N}\}$ .

$\rho_X(\mathcal{C}) = 2\rho_X(\mathcal{K})$  is the same as probability that a randomly sampled point on the surface of a hemisphere lies in the conical hull of  $n$  points sampled uniformly on the surface of the hemisphere.

**Theorem 7.** Recall that  $\rho_X$  is the uniform distribution over  $\mathbb{S}^{d-1}$ . Then, for sufficiently large  $N$ ,

$$\frac{d^{3/2}}{N\sqrt{\log(d)}} \lesssim \mathbb{E}[\rho_X(\mathcal{C}^c)] \lesssim \frac{d^{3/2} \log(d)}{N}. \quad (8)$$

The proof of this result is fairly involved and we defer it to the Appendix. The key approach is to represent  $\mathcal{K}$  in its dual representation and computing the probability in the dual space. The proof uses the Poissonization trick and a delicate covering argument to argue concentration in the absence of Lipschitzness. Tabular IL corresponds to the confidence set linear classification with orthogonal features. In particular,  $\langle \theta, \phi(x, 0) - \phi(x, 1) \rangle$  is known exactly for each  $x$ , and for each unobserved  $x$  nothing can be said about  $\text{sign}(\langle \theta, \phi(x, 0) - \phi(x, 1) \rangle)$ . There, the minimax risk translates to the expected probability mass on unobserved inputs - the *missing mass* [18] which is in expectation  $\lesssim d/N$  [22], which is tight over the worst-case choice of distribution on inputs,  $\rho_X$ . The  $\tilde{\Omega}(d^{3/2}/N)$  rate in thus Theorem 7 establishes differences between the tabular setting and the linear-expert setting with linear rewards in the context of the approach of simulating artificial trajectories.

## 5 Conclusion

We study IL in the presence of  $\mu$ -recoverability and under linear function approximation. In the former case, we establish a separation in the minimax expected suboptimality of learners in the no-interaction and active settings. We show upper bounds for BC under the linear expert setting and show that this quadratic dependence on  $H$  can be broken in the presence of parameter sharing. Finally, we study the known transition setting, and introduce a problem known as confidence set linear classification which extends the approach of simulating artificial trajectories to the function approximation setting.

## 6 Acknowledgments and Disclosure of Funding

Nived Rajaraman and Jiantao Jiao were partially supported by NSF Grants IIS-1901252, and CCF-1909499. Part of the work was done while Lin F. Yang was visiting the Simons Institute for the Theory of Computing at UC Berkeley (the Theory of Reinforcement Learning Program). Jingbo Liu was partially supported by the starting grant from the Department of Statistics, University of Illinois. Kannan Ramchandran was supported in part by the Army Research Office (ARO) under fund 62448.

## References

- [1] Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019. ISSN 1476-4687. doi: 10.1038/s41586-019-1724-z. URL <https://doi.org/10.1038/s41586-019-1724-z>.
- [2] Pieter Abbeel and Andrew Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the Twenty-First International Conference on Machine Learning*, ICML '04, page 1, New York, NY, USA, 2004. Association for Computing Machinery. ISBN 1581138385. doi: 10.1145/1015330.1015430. URL <https://doi.org/10.1145/1015330.1015430>.
- [3] Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, page 1, 2004.
- [4] Kianté Brantley, Wen Sun, and Mikael Henaff. Disagreement-regularized imitation learning. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rkgbYyHtwB>.
- [5] Jianyu Chen, Bodi Yuan, and Masayoshi Tomizuka. Deep imitation learning for autonomous driving in generic urban scenarios with enhanced safety. *CoRR*, abs/1903.00640, 2019. URL <http://arxiv.org/abs/1903.00640>.
- [6] Ching-An Cheng and Byron Boots. Convergence of value aggregation for imitation learning, 2018.
- [7] Ching-An Cheng, Xinyan Yan, Nolan Wagener, and Byron Boots. Fast policy learning through imitation and reinforcement, 2018.
- [8] Amit Daniely and S. Shalev-Shwartz. Optimal learners for multiclass problems. In *COLT*, 2014.
- [9] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. CARLA: An open urban driving simulator. In Sergey Levine, Vincent Vanhoucke, and Ken Goldberg, editors, *Proceedings of the 1st Annual Conference on Robot Learning*, volume 78 of *Proceedings of Machine Learning Research*, pages 1–16. PMLR, 13–15 Nov 2017. URL <http://proceedings.mlr.press/v78/dosovitskiy17a.html>.
- [10] Chelsea Finn, Sergey Levine, and Pieter Abbeel. Guided cost learning: Deep inverse optimal control via policy optimization. In *International conference on machine learning*, pages 49–58, 2016.
- [11] Justin Fu, Katie Luo, and Sergey Levine. Learning robust rewards with adversarial inverse reinforcement learning. *arXiv preprint arXiv:1710.11248*, 2017.
- [12] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *Advances in neural information processing systems*, pages 4565–4573, 2016.
- [13] Borja Ibarz, Jan Leike, Tobias Pohlen, Geoffrey Irving, Shane Legg, and Dario Amodei. Reward learning from human preferences and demonstrations in atari. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 8011–8023. Curran Associates, Inc., 2018. URL <http://papers.nips.cc/paper/8025-reward-learning-from-human-preferences-and-demonstrations-in-atari.pdf>.
- [14] Liyiming Ke, Matt Barnes, Wen Sun, Gilwoo Lee, Sanjiban Choudhury, and Siddhartha Srinivasa. Imitation learning as  $f$ -divergence minimization. *arXiv preprint arXiv:1905.12888*, 2019.
- [15] Michael Laskey, Jonathan Lee, Roy Fox, Anca Dragan, and Ken Goldberg. Dart: Noise injection for robust imitation learning. In *Conference on robot learning*, pages 143–156. PMLR, 2017.
- [16] Jonathan N Lee, Michael Laskey, Ajay Kumar Tanwani, Anil Aswani, and Ken Goldberg. A dynamic regret analysis and adaptive regularization algorithm for on-policy robot imitation learning. In *International Workshop on the Algorithmic Foundations of Robotics*, pages 212–227. Springer, 2018.
- [17] Yuping Luo, Huazhe Xu, and Tengyu Ma. Learning self-correctable policies and value functions from demonstrations with negative sampling. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rke-f6NKvS>.
- [18] David Mcallester, Luis Ortiz, Ralf Herbrich, and Thore Graepel. Concentration inequalities for the missing mass and for histogram rule error. In *Journal of Machine Learning Research*, pages 895–911, 2003.
- [19] Yunpeng Pan, Ching-An Cheng, Kamil Saigol, Keuntaek Lee, Xinyan Yan, Evangelos Theodorou, and Byron Boots. Agile autonomous driving using end-to-end deep imitation learning. *arXiv preprint arXiv:1709.07174*, 2017.
- [20] Yunpeng Pan, Ching-An Cheng, Kamil Saigol, Keuntaek Lee, Xinyan Yan, Evangelos A Theodorou, and Byron Boots. Imitation learning for agile autonomous driving. *The International Journal of Robotics Research*, 39(2-3):286–302, 2020. doi: 10.1177/0278364919880273. URL <https://doi.org/10.1177/0278364919880273>.

- [21] Dean A. Pomerleau. Alvin: An autonomous land vehicle in a neural network. In D. S. Touretzky, editor, *Advances in Neural Information Processing Systems 1*, pages 305–313. Morgan-Kaufmann, 1989. URL <http://papers.nips.cc/paper/95-alvin-an-autonomous-land-vehicle-in-a-neural-network.pdf>.
- [22] Nived Rajaraman, Lin F. Yang, Jiantao Jiao, and Kannan Ramchandran. Toward the fundamental limits of imitation learning. In *Advances in Neural Information Processing Systems*, 2020.
- [23] Nived Rajaraman, Yanjun Han, Lin F. Yang, Kannan Ramchandran, and Jiantao Jiao. Provably breaking the quadratic error compounding barrier in imitation learning, optimally, 2021.
- [24] Nathan D Ratliff, J Andrew Bagnell, and Martin A Zinkevich. Maximum margin planning. In *Proceedings of the 23rd international conference on Machine learning*, pages 729–736, 2006.
- [25] Stéphane 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 <http://proceedings.mlr.press/v9/ross10a.html>.
- [26] Stéphane Ross and J. Andrew Bagnell. Reinforcement and imitation learning via interactive no-regret learning. *ArXiv*, abs/1406.5979, 2014.
- [27] Stéphane Ross, Geoffrey J. Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In Geoffrey J. Gordon, David B. Dunson, and Miroslav Dudík, editors, *AISTATS*, volume 15 of *JMLR Proceedings*, pages 627–635. JMLR.org, 2011. URL <http://dblp.uni-trier.de/db/journals/jmlr/jmlr15.html#RossGB11>.
- [28] Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194, 2012. ISSN 1935-8237. doi: 10.1561/22000000018. URL <http://dx.doi.org/10.1561/22000000018>.
- [29] Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, USA, 2014. ISBN 1107057132.
- [30] Wen Sun, Arun Venkatraman, Geoffrey J Gordon, Byron Boots, and J Andrew Bagnell. Deeply aggravated: Differentiable imitation learning for sequential prediction. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3309–3318. JMLR. org, 2017.
- [31] Wen Sun, Anirudh Vemula, Byron Boots, and Drew Bagnell. Provably efficient imitation learning from observation alone. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 6036–6045, Long Beach, California, USA, 09–15 Jun 2019. PMLR. URL <http://proceedings.mlr.press/v97/sun19b.html>.
- [32] Gokul Swamy, Sanjiban Choudhury, J. Andrew Bagnell, and Steven Wu. Of moments and matching: A game-theoretic framework for closing the imitation gap. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 10022–10032. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/swamy21a.html>.
- [33] Umar Syed and Robert E Schapire. A game-theoretic approach to apprenticeship learning. In J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis, editors, *Advances in Neural Information Processing Systems 20*, pages 1449–1456. Curran Associates, Inc., 2008. URL <http://papers.nips.cc/paper/3293-a-game-theoretic-approach-to-apprenticeship-learning.pdf>.
- [34] Umar Syed, Michael Bowling, and Robert E Schapire. Apprenticeship learning using linear programming. In *Proceedings of the 25th international conference on Machine learning*, pages 1032–1039, 2008.
- [35] Yurong You, Xinlei Pan, Ziyang Wang, and Cewu Lu. Virtual to real reinforcement learning for autonomous driving. *CoRR*, abs/1704.03952, 2017. URL <http://arxiv.org/abs/1704.03952>.
- [36] Yufeng Zhang, Qi Cai, Zhuoran Yang, and Zhaoran Wang. Generative adversarial imitation learning with neural networks: Global optimality and convergence rate. *arXiv preprint arXiv:2003.03709*, 2020.
- [37] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, pages 1433–1438. Chicago, IL, USA, 2008.

## Checklist

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
  - (b) Did you describe the limitations of your work? [Yes] See section 5. We do not address general minimax rates of confidence set linear classification, which is an important and challenging open problem
  - (c) Did you discuss any potential negative societal impacts of your work? [Yes] See section 5
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
  - (b) Did you include complete proofs of all theoretical results? [Yes] All proofs are included in the supplementary material
3. If you ran experiments...
  - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
  - (a) If your work uses existing assets, did you cite the creators? [N/A]
  - (b) Did you mention the license of the assets? [N/A]
  - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
  - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
  - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
  - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
  - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

## 7 Appendix A

### 7.1 Proofs for Theorem 1

**Notation:** Throughout this section, we use  $\preceq$  and  $\succeq$  to indicate the partial ordering on vectors where  $a \preceq b$  if  $a$  is not larger than  $b$  co-ordinate wise ( $\succeq$  is defined analogously). We also use  $\mathbf{0}$  to denote the all 0's vector and  $\mathbf{1}$  denote the all 1's vector (where the dimension is inferred from context). For a vector  $\mathbf{0} \preceq w \in \mathbb{R}^d$ , the norm  $\|x\|_w^2 \triangleq \sum_{i=1}^d w_i x_i^2$  is the weighted-L2 norm.

Recall from eq. (2), the population expected 0-1 loss of a policy  $\pi$  is defined as

$$\mathcal{L}(f^\pi, \pi, \pi^*) = \frac{1}{H} \sum_{t=1}^H \mathbb{E}_{S_t \sim f_t^\pi(\cdot)} \left[ \mathbb{E}_{a \sim \hat{\pi}_t(\cdot|s)} [\mathbb{1}(a \neq \pi_t^*(s))] \right]. \quad (9)$$

**Lemma 3.** *Suppose there is an online learning algorithm which outputs policies  $\{\hat{\pi}^1, \dots, \hat{\pi}^N\}$  sequentially according to any procedure where the learner samples the policy  $\hat{\pi}_i$  from some distribution conditioned on  $\text{tr}_1, \dots, \text{tr}_{i-1}$ , subsequently samples a trajectory  $\text{tr}_i$  by rolling out  $\hat{\pi}_i$ , repeating this process for  $N$  iterations. Denote  $\hat{f}^{\hat{\pi}^i} = \{\hat{f}_1^{\hat{\pi}^i}, \dots, \hat{f}_H^{\hat{\pi}^i}\}$  where  $\hat{f}_t^{\hat{\pi}^i}$  denotes the empirical distribution over states induced by the single trajectory  $\text{tr}_i$  at time  $t$ . Denote  $\hat{\pi} = \frac{1}{N} \sum_{i=1}^N \hat{\pi}^i$  as the mixture policy. Then,*

$$\mathbb{E} \left[ \mathcal{L}(f^{\hat{\pi}}, \hat{\pi}, \pi^*) \right] = \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[ \mathcal{L}(\hat{f}^{\hat{\pi}^i}, \hat{\pi}^i, \pi^*) \right]. \quad (10)$$

*Proof.* Since the trajectory  $\text{tr}_i$  is rolled out using  $\hat{\pi}^i$ , conditioned on  $\hat{\pi}^i$ ,  $\hat{f}^{\hat{\pi}^i}$  is conditionally unbiased and in expectation equal to  $f^{\hat{\pi}^i}$  (conditioned on  $\hat{\pi}^i$ ). Therefore, for each  $i$ , since  $\mathcal{L}(f, \hat{\pi}, \pi^*)$  is a linear function in  $f$ ,

$$\mathbb{E} \left[ \mathcal{L}(\hat{f}^{\hat{\pi}^i}, \hat{\pi}^i, \pi^*) \middle| \hat{\pi}^i \right] = \mathcal{L}(f^{\hat{\pi}^i}, \hat{\pi}^i, \pi^*). \quad (11)$$

Summing across  $i = 1, \dots, N$  and using the fact that  $\mathcal{L}(f^{\hat{\pi}}, \hat{\pi}, \pi^*) = \frac{1}{N} \sum_{i=1}^N \mathcal{L}(f^{\hat{\pi}^i}, \hat{\pi}^i, \pi^*)$  and taking expectation completes the proof.  $\square$

The conclusion of this lemma is that it suffices to minimize the empirical 0-1 loss under the learner's own *one-trajectory* empirical state distribution  $\frac{1}{N} \sum_{i=1}^N \mathcal{L}(\hat{f}^{\hat{\pi}^i}, \hat{\pi}^i, \pi^*)$ . Note that for any policy  $\pi$ , the loss

$$\mathcal{L}(\hat{f}^{\hat{\pi}^i}, \pi, \pi^*) = \frac{1}{H} \sum_{t=1}^H \sum_{s \in \mathcal{S}} \left\langle \pi_t(\cdot|s), z_t^i(s) \right\rangle \quad (12)$$

where  $z_t^i(s) = \left\{ \hat{f}_t^{\hat{\pi}^i}(s)(1 - \pi^*(a|s)) \right\}_{a \in \mathcal{A}} \in \Delta_{\mathcal{A}}^1$  is a linear function in the policy  $\pi$ . The constraint on the policy variable  $\pi$  is that for each  $t \in [H]$  and  $s \in \mathcal{S}$ ,  $\pi_t(\cdot|s) \in \Delta_{\mathcal{A}}^1$ .

Define the loss  $\ell_{i,s,t}(\pi) = \sum_{s \in \mathcal{S}} \left\langle \pi_t(\cdot|s), z_t^i(s) \right\rangle$ . Then the variable  $\pi_t(\cdot|s)$  lies in the simplex  $\Delta_{\mathcal{A}}^1$  and the vector  $z_t^i(s)$  is co-ordinate wise  $\geq 0$  and  $\leq 1$ .

To learn the sequence of policies returned by the learner, we use the normalized-EG algorithm of [28] which is also known as Follow-the-regularized-leader / Online Mirror Descent with entropy regularization for online learning. Formally, the online learning problem and the algorithm are as defined in Section 2 of [28].

**Theorem 8** (Adapted from Theorem 2.22 in [28]). *Assume that the normalized EG algorithm is run on a sequence of linear loss functions  $\{\langle z_i, \cdot \rangle : i = 1, \dots, T\}$ , with  $\eta = 1/2$  to return a sequence of distributions  $w_1, \dots, w_T \in \Delta_{\mathcal{A}}^1$ . Assume that for all  $t \in [H]$ ,  $\mathbf{0} \preceq z_t \preceq \mathbf{1}$ . For any  $u$  such that  $\sum_{i=1}^T \langle z_i, u \rangle = 0$ ,*

$$\sum_{t=1}^T \langle w_t - u, z_t \rangle \leq 4 \log(|\mathcal{A}|). \quad (13)$$

This result is adapted from Theorem 2.22 in [28] by invoking the condition that  $\mathbf{0} \preceq z_t \preceq \mathbf{1}$ , so the local norm  $\|z_t\|_{w_t}^2$  can be upper bounded by  $\langle z_t, w_t \rangle$ . Choosing  $\eta = \frac{1}{2}$ , using the assumption that  $\sum_{i=1}^T \langle z_i, u \rangle = 0$  and simplifying results in the statement of Theorem 8.

Suppose the learner returns the sequence of policies  $\hat{\pi}^1, \dots, \hat{\pi}^N$  by running the normalized EG algorithm on the sequence of losses  $\ell_{1,s,t}, \dots, \ell_{N,s,t}$  for each  $s \in \mathcal{S}$  and  $t \in [H]$  to return a sequence of distributions  $\hat{\pi}_t^1(\cdot|s), \dots, \hat{\pi}_t^N(\cdot|s) \in \Delta_{\mathcal{A}}^1$ . Finally, for  $i = 1, \dots, N$ , the learner returns the policy  $\hat{\pi}^i$  as  $\{\{\hat{\pi}_t^i(\cdot|s) : s \in \mathcal{S}\} : t \in [H]\}$ .

Invoking the guarantee in Theorem 8 for the sequence of policies  $\hat{\pi}_t^1(\cdot|s), \dots, \hat{\pi}_t^N(\cdot|s)$  returned by a single instance of the normalized-EG algorithm,

$$\sum_{i=1}^T \langle z_t^i(s), \hat{\pi}_t^i(\cdot|s) \rangle \leq 4 \log(|\mathcal{A}|) \quad (14)$$

Averaging across  $t \in [H]$ , summing across  $s \in \mathcal{S}$  and recalling the definition of  $\mathcal{L}$  in eq. (12) results in the bound,

$$\frac{1}{N} \sum_{i=1}^N \mathcal{L}(\hat{f}^{\hat{\pi}^i}, \hat{\pi}^i, \pi^*) \leq \frac{4|\mathcal{S}| \log(|\mathcal{A}|)}{N}. \quad (15)$$

Finally invoking Lemma 3 shows that the resulting sequence of policies  $\hat{\pi}^1, \dots, \hat{\pi}^N$  and their mixtures  $\frac{1}{N} \sum_{i=1}^N \hat{\pi}^i$  satisfies,

$$\mathbb{E} \left[ \mathcal{L}(f^{\hat{\pi}}, \hat{\pi}, \pi^*) \right] \leq \frac{4|\mathcal{S}| \log(|\mathcal{A}|)}{N}. \quad (16)$$

Invoking [27, Theorem 2], under  $\mu$ -recoverability shows that the resulting policy  $\hat{\pi}$  satisfies,

$$\mathbb{E}[J(\pi^*) - J(\hat{\pi})] \leq \frac{4|\mathcal{S}| \log(|\mathcal{A}|)}{N}. \quad (17)$$

This completes the proof of Theorem 1.

## 7.2 Proof of Theorem 2

**Theorem 9** (Theorem 6.1 in [22]). *For any learner  $\hat{\pi}$ , there exists an MDP  $\mathcal{M}$  and a deterministic expert policy  $\pi^*$  such that the expected suboptimality of the learner is lower bounded in the active setting by,  $J_{\mathcal{M}}(\pi^*) - \mathbb{E}[J_{\mathcal{M}}(\hat{\pi})] \gtrsim \min \left\{ H, \frac{|\mathcal{S}|H^2}{N} \right\}$ .*

For each active learner  $\hat{\pi}$  and the worst-case IL instance  $(\pi^*, \mathcal{M})$  from Theorem 9, consider the IL instance  $(\pi^*, \mathcal{M}_\mu)$  where the only difference between  $\mathcal{M}$  and  $\mathcal{M}_\mu$  is that each reward is scaled by a factor of  $\mu/H \leq 1$ . Note that  $\mathcal{M}_\mu$  satisfies  $\mu$ -recoverability. Indeed, consider any state  $s$ . Since the rewards in  $\mathcal{M}_\mu$  are in the interval  $[0, \mu/H]$ , the total reward of any trajectory in  $\mathcal{M}_\mu$  lies in the interval  $[0, \mu]$ . Therefore, trivially, for each  $(s, a, t) \in \mathcal{S} \times \mathcal{A} \times [H]$  tuple,  $Q_t^{\pi^*}(s, \pi_t^*(s)) - Q_t^{\pi^*}(s, a) \leq \mu - 0 = \mu$  and the IL instance satisfies  $\mu$ -recoverability. More importantly the suboptimality of  $\hat{\pi}$  on the IL instance  $(\pi^*, \mathcal{M}_\mu)$  is  $\frac{\mu}{H}$  times the suboptimality under  $(\pi^*, \mathcal{M})$ . In other words,

$$\mathbb{E} [J_{\mathcal{M}_\mu}(\pi^*) - J_{\mathcal{M}_\mu}(\hat{\pi})] = \frac{\mu}{H} \mathbb{E} [J_{\mathcal{M}}(\pi^*) - J_{\mathcal{M}}(\hat{\pi})] \quad (18)$$

$$\gtrsim \frac{\mu}{H} \min \left\{ H, \frac{|\mathcal{S}|H^2}{N} \right\} \quad (19)$$

$$= \min \left\{ \mu, \frac{\mu|\mathcal{S}|H}{N} \right\}, \quad (20)$$

where the last inequality uses [22, Theorem 6.1]. This concludes the proof of Theorem 2.

## 7.3 Proof of Theorem 3

**Theorem 10.** *In the no-interaction setting, for any learner  $\hat{\pi}$  and  $|\mathcal{S}| \geq 3$ , there exists an MDP  $\mathcal{M}$  with state space  $\mathcal{S}$  and a deterministic expert policy  $\pi^*$  which (i) satisfies the  $\mu$ -recovery assumption for any  $\mu > 1$ , and (ii) such that the expected suboptimality of the learner is lower bounded by,*

$$J_{\mathcal{M}}(\pi^*) - \mathbb{E}[J_{\mathcal{M}}(\hat{\pi})] \gtrsim \min \left\{ H, |\mathcal{S}|H^2/N \right\}. \quad (21)$$

In this section we discuss the proof of the lower bound in Theorem 10.

Define  $\mathcal{S}_t(D)$  as the set of states observed in at least one trajectory at time  $t$  in the expert dataset  $D$ . In particular, the learner exactly knows the expert's policy  $\pi_t^*(\cdot|s)$  at all states  $s \in \mathcal{S}_t(D)$  for each  $t = 1, \dots, H$ .

The expert policy is deterministic in the lower bound instances we construct. Define  $\Pi_{\text{mimic}}(D)$  as the family of policies which mimics the expert on the states visited in  $D$ . Namely,

$$\Pi_{\text{mimic}}(D) \triangleq \left\{ \pi : \forall t \in [H], s \in \mathcal{S}_t(D), \pi_t(\cdot|s) = \pi_t^*(\cdot|s) \right\}, \quad (22)$$

Informally,  $\Pi_{\text{mimic}}(D)$  is the family of policies which are ‘‘compatible’’ with the expert dataset  $D$ .

In order to prove the lower bound on the worst-case expected suboptimality of any learner  $\hat{\pi}(D)$ , it suffices to lower bound the Bayes expected suboptimality and find a joint distribution  $\mathcal{P}$  over MDPs and expert policies satisfying  $\mu$ -recoverability, such that,  $\mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}} \left[ J_{\mathcal{M}}(\pi^*) - \mathbb{E} [J_{\mathcal{M}}(\hat{\pi}(D))] \right] \gtrsim \min \left\{ H, \frac{|\mathcal{S}|H^2}{N} \right\}$ .

**Construction of  $\mathcal{P}$ :** First the expert's policy is sampled uniformly from  $\Pi_{\text{det}}$ : for each  $t \in [H]$  and  $s \in \mathcal{S}$ ,  $\pi_t^*(s) \sim \text{Unif}(\mathcal{A})$ . Conditioned on  $\pi^*$ , the distribution over MDPs induced by  $\mathcal{P}$  is deterministic and given by the MDP  $\mathcal{M}[\pi^*]$  in fig. 2.  $\mathcal{M}[\pi^*]$  has a fixed initial distribution over states  $\rho = \{\zeta, \dots, \zeta, 1 - (|\mathcal{S}| - 2)\zeta, 0\}$  where  $\zeta = \frac{1}{N+1}$ . There is a special state  $b \in \mathcal{S}$  in the MDP which has behavior different from the remaining states. At each state  $s \in \mathcal{S}$ , choosing the expert's action renews the state in the initial distribution  $\rho$  providing a reward of 1 (except at state  $b$  it provides a reward of 0), while every other action deterministically transitions the learner to the bad state and provides no reward. That is,

$$P_t(\cdot|s, a) = \begin{cases} \rho, & s \in \mathcal{S}, a = \pi_t^*(s) \\ \delta_b, & \text{otherwise,} \end{cases} \quad (23)$$

and the reward function of the MDP is given by,

$$\mathbf{r}_t(s, a) = \begin{cases} 1, & s \in \mathcal{S} \setminus \{b\}, a = \pi_t^*(s) \\ 0, & \text{otherwise.} \end{cases} \quad (24)$$

We first state a simple consequence of the construction of the MDP instances and  $\mathcal{P}$ . Note that the expert never visits the bad state  $b$  by virtue of the distribution  $\rho$  placing no mass on  $b$ . Therefore, the value of  $\pi^*$  on the MDP  $\mathcal{M}[\pi^*]$  is  $H$ .

**Lemma 4.** *The value of  $\pi^*$  on the MDP  $\mathcal{M}[\pi^*]$  is  $H$ . Namely  $J_{\mathcal{M}[\pi^*]}(\pi^*) = H$ .*

*Proof.* Playing the expert's action at any state in  $\mathcal{S}$  is the only way to accrue non-zero reward, and in fact accrues a reward of 1. Thus the expert collects a reward of 1 at each time in any trajectory.  $\square$

At the states unvisited in the dataset  $D$ , the learner cannot infer the expert's policy or even the transitions induced under different actions. Intuitively, the learner cannot guess the expert's action with probability  $\geq 1/|\mathcal{A}|$  at such states, a statement which we prove by leveraging the Bayesian construction. In turn, the learner is forced to visit the bad state  $b$  at the next point in the episode. Since the bad state is never observed in the dataset, the learner is forced to guess the expert's action to be able to recover in the distribution  $\rho$  over the remaining states (lest it collects a reward of 0 for the rest of the episode). However by making  $|\mathcal{A}|$  large ( $\gtrsim H$ ), any learner, with constant probability fails to guess the expert's action at  $b$  at at least a constant fraction of the episode.

Using [22, Lemma A.14], the conditional distribution of the expert's policy given the expert dataset  $D$  can be characterized.

**Lemma 5.** *([22, Lemma A.14]) Conditioned on the dataset  $D$  collected by the learner, the expert's deterministic policy  $\pi^*$  is distributed  $\sim \text{Unif}(\Pi_{\text{mimic}}(D))$ . In other words, at each state visited in the expert dataset, the expert's choice of action is fixed as the one returned when the expert was queried at this state. At the remaining states, the expert's choice of action is sampled uniformly from  $\mathcal{A}$ .*






Figure 2: Upon playing the expert's (green) action at any state, the learner is renewed in the initial distribution  $\rho = \{\zeta, \dots, \zeta, 1 - (|\mathcal{S}|-2)\zeta, 0\}$  where  $\zeta = \frac{1}{N+1}$ . Any other choice of action (red) deterministically transitions the learner to  $b$ .

**Definition 9.** Define  $\mathcal{P}(D)$  as the joint distribution of  $(\pi^*, \mathcal{M})$  conditioned on the expert dataset  $D$ . Conditionally,  $\pi^* \sim \text{Unif}(\Pi_{\text{mimic}}(D))$  and  $\mathcal{M} = \mathcal{M}[\pi^*]$ .

From Lemma 5 and the definition of  $\mathcal{P}(D)$  in Definition 9, applying Fubini's theorem gives,

$$\mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}} [J_{\mathcal{M}}(\pi^*) - \mathbb{E}[J_{\mathcal{M}}(\hat{\pi})]] = \mathbb{E} [\mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} [H - J_{\mathcal{M}}(\hat{\pi}(D))]]. \quad (25)$$

Next we relate this to the first time the learner visits a state unobserved in  $D$ .

**Lemma 6.** Define the stopping time  $\tau$  as the first time  $t$  that the learner encounters a state  $s_t \neq b$  that has not been visited in  $D$  at time  $t$ . That is,

$$\tau = \begin{cases} \inf\{t : s_t \notin \mathcal{S}_t(D) \cup \{b\}\} & \exists t : s_t \notin \mathcal{S}_t(D) \cup \{b\} \\ H + 1 & \text{otherwise.} \end{cases} \quad (26)$$

Then, conditioned on the expert dataset  $D$ ,

$$\mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} [J(\pi^*) - \mathbb{E}[J(\hat{\pi})]] \geq \left(1 - \frac{1}{|\mathcal{A}|}\right)^{H+1} \mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} [\mathbb{E}_{\hat{\pi}(D)} [H - \tau]]. \quad (27)$$

We defer the proof of this result to the end of this section.

Plugging the result of Lemma 6 into eq. (25), and recalling the assumption that  $|\mathcal{A}| \geq H + 1$ ,

$$\mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}} [J(\pi^*) - \mathbb{E}[J(\hat{\pi})]] \geq \frac{1}{4} \mathbb{E} [\mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} [\mathbb{E}_{\hat{\pi}} [H - \tau]]], \quad (28)$$

$$\stackrel{(i)}{\geq} \frac{H}{8} \mathbb{E} [\mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} [\Pr_{\hat{\pi}} [\tau \leq \lfloor H/2 \rfloor]]], \quad (29)$$

$$= \frac{H}{8} \mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}} [\mathbb{E} [\Pr_{\hat{\pi}} [\tau \leq \lfloor H/2 \rfloor]]], \quad (30)$$

where (i) uses Markov's inequality, and the last equation uses Fubini's theorem.

The last remaining element of the proof is to indeed bound the probability that the learner visits a state unobserved in the dataset before time  $\lfloor H/2 \rfloor$  which immediately follows from [22, Lemma A.16] shows that for any learner  $\hat{\pi}$ ,  $\mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}} [\mathbb{E} [\Pr_{\hat{\pi}} [\tau \leq \lfloor H/2 \rfloor]]]$  is lower bounded by  $\gtrsim \min\{1, |\mathcal{S}|H/N\}$ . Therefore,

$$\mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}} [J(\pi^*) - \mathbb{E}[J(\hat{\pi})]] \gtrsim H \min \left\{ 1, \frac{|\mathcal{S}|H}{N} \right\}. \quad (31)$$

as long as  $|\mathcal{A}| \geq H + 1$ .

**Lemma 7.** ([22, Lemma A.16]) For any learner policy  $\hat{\pi}$ ,

$$\mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}} [\mathbb{E} [\Pr_{\hat{\pi}} [\tau \leq \lfloor H/2 \rfloor]]] \geq 1 - \left(1 - \frac{|\mathcal{S}| - 2}{e(N+1)}\right)^{\lfloor H/2 \rfloor} \gtrsim \min \left\{ 1, \frac{|\mathcal{S}|H}{N} \right\}. \quad (32)$$

Finally, we prove Lemma 6.

*Proof of Lemma 6.* Define the random time  $\tau_b$  to be the first time the learner encounters the state  $b$  while rolling out a trajectory. Formally,

$$\tau_b = \begin{cases} \inf\{t : s_t = b\} & \exists t : s_t = b \\ H + 1 & \text{otherwise.} \end{cases}$$

Furthermore, define  $\Gamma_b$  as the random variable which counts the number of time steps the trajectory stays in the state  $b$  after visiting it for the first time. Namely,

$$\Gamma_b = \begin{cases} \inf\{\Delta \geq 0 : s_{\tau_b + \Delta + 1} \neq b\} & \tau_b \leq H \\ 0 & \text{otherwise.} \end{cases} \quad (33)$$

Since the state  $b$  always dispenses 0 reward and since  $\mathbf{r}$  is bounded in  $[0, 1]$ , conditioned on the expert dataset  $D$ ,

$$H - \mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} [J(\hat{\pi})] = H - \mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} \left[ \mathbb{E}_{\hat{\pi}} \left[ \sum_{t=1}^H \mathbf{r}_t(s_t, a_t) \right] \right] \quad (34)$$

$$\geq \mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} [\mathbb{E}_{\hat{\pi}} [\Gamma_b]] \quad (35)$$

Fixing the expert dataset  $D$  and the expert's policy  $\pi^*$  (which determines the MDP  $\mathcal{M}[\pi^*]$ ), we under the distribution of  $\Gamma_b$ .

To this end, first observe that for any  $t \leq H - 1$  and state  $s \in \mathcal{S}$ ,

$$\Pr_{\hat{\pi}} [\Gamma_b \geq \Delta + 1, \Gamma_b \geq \Delta, \tau_b = t] \quad (36)$$

$$= \Pr_{\hat{\pi}} [\Gamma_b \geq \Delta + 1 | \Gamma_b \geq \Delta, \tau_b = t] \Pr_{\hat{\pi}} [\Gamma_b \geq \Delta, \tau_b = t] \quad (37)$$

$$= \left( 1 - \hat{\pi}_{t+\Delta}(\pi_{t+\Delta}^*(b)|b) \right) \Pr_{\hat{\pi}} [\Gamma_b \geq \Delta, \tau_b = t]. \quad (38)$$

where in the last equation, we use the fact that the learner must play an action other than  $\pi_{t+\Delta}^*(b)$  to stay in state  $b$  at time  $t + \Delta$ . Next, we take expectation with respect to the randomness of  $\pi^*$ . Conditioned on  $D$ ,  $\pi^*$  is sampled uniformly from the set of policies  $\Pi_{\text{mimic}}(D)$  (Lemma 5). In particular, conditioned on  $D$ , the expert policy is sampled independently at states. Conditioned on  $\pi^*$ , the underlying MDP is  $\mathcal{M}[\pi^*]$ . Observe that the dependence of the second term  $\Pr_{\hat{\pi}} [\tau = t, s_t = s]$  on  $\pi^*$  comes from the probability computed with the underlying MDP chosen as  $\mathcal{M}[\pi^*]$ . Observe that it only depends on the characteristics of  $\mathcal{M}[\pi^*]$  till time  $t - 1$  which are determined by  $\pi_1^*, \dots, \pi_{t+\Delta-1}^*$ . On the other hand, the first term  $(1 - \hat{\pi}_t(\pi_{t+\Delta}^*(b)|b))$  depends only on the random variable  $\pi_{t+\Delta}^*$ . As a consequence, the two terms depend on a disjoint set of random variables which are independent.

Taking expectation with respect to the randomness of  $\pi^* \sim \text{Unif}(\Pi_{\text{mimic}}(D))$  and  $\mathcal{M} = \mathcal{M}[\pi^*]$  (which defines the joint distribution  $\mathcal{P}(D)$  in eq. (22)), for  $0 \leq \Delta \leq H - t$  and  $t \in [H]$ ,

$$\begin{aligned} & \mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} \left[ \Pr_{\hat{\pi}} [\Gamma_b \geq \Delta + 1, \Gamma_b \geq \Delta, \tau_b = t] \right] \\ &= \mathbb{1}(t + \Delta \leq H) \cdot \mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} \left[ 1 - \hat{\pi}_{t+\Delta}(\pi_{t+\Delta}^*(b)|b) \right] \mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} \left[ \Pr_{\hat{\pi}} [\Gamma_b \geq \Delta, \tau_b = t] \right] \end{aligned} \quad (39)$$

$$= \mathbb{1}(t + \Delta \leq H) \cdot \left( 1 - \frac{1}{|\mathcal{A}|} \right) \mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} \left[ \Pr_{\hat{\pi}} [\Gamma_b \geq \Delta, \tau_b = t] \right], \quad (40)$$

where in the last equation we use the fact that the state  $b$  is never observed in the expert dataset. So conditioned on  $D$ ,  $\pi_{t+\Delta}^*(b)$  is sampled uniformly from  $\mathcal{A}$ . By upper bounding  $\Pr_{\hat{\pi}} [\Gamma_b \geq \Delta + 1, \Gamma_b \geq \Delta, \tau_b = t] \leq \Pr_{\hat{\pi}} [\Gamma_b \geq \Delta + 1, \tau_b = t]$  results in the inequality,

$$\begin{aligned} & \mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} \left[ \Pr_{\hat{\pi}} [\Gamma_b \geq \Delta + 1, \tau_b = t] \right] \\ & \geq \mathbb{1}(t + \Delta \leq H) \cdot \left( 1 - \frac{1}{|\mathcal{A}|} \right) \mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} \left[ \Pr_{\hat{\pi}} [\Gamma_b \geq \Delta, \tau_b = t] \right] \end{aligned} \quad (41)$$

Unrolling the equation, for each  $\Delta = 0, 1, \dots, H - t + 1$  we have,

$$\begin{aligned} & \mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} \left[ \Pr_{\hat{\pi}} [\Gamma_b \geq \Delta, \tau_b = t] \right] \\ & \geq \left( 1 - \frac{1}{|\mathcal{A}|} \right)^\Delta \mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} \left[ \Pr_{\hat{\pi}} [\tau_b = t] \right] \end{aligned} \quad (42)$$

Summing up over  $\Delta = 0, 1, \dots, H - t + 1$ ,

$$\begin{aligned} & \mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} \left[ \mathbb{E}_{\hat{\pi}} [\Gamma_b \mathbb{1}(\tau_b = t)] \right] \\ & \geq \sum_{\Delta=0}^{H-t+1} \left( 1 - \frac{1}{|\mathcal{A}|} \right)^\Delta \mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} \left[ \Pr_{\hat{\pi}} [\tau_b = t] \right] \end{aligned} \quad (43)$$

$$\geq (H - t + 1) \left( 1 - \frac{1}{|\mathcal{A}|} \right)^H \mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} \left[ \Pr_{\hat{\pi}} [\tau_b = t] \right]. \quad (44)$$

Finally summing across  $t = 1, \dots, H + 1$ ,

$$\mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} \left[ \mathbb{E}_{\hat{\pi}} [\Gamma_b] \right] \geq \left( 1 - \frac{1}{|\mathcal{A}|} \right)^H \mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} \left[ \mathbb{E}_{\hat{\pi}} [H - \tau_b + 1] \right]. \quad (45)$$

Finally, we invoke [22, Lemma A.15] and in particular, eq. (134) to arrive at the desired bound.

$$\mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} \left[ \mathbb{E}_{\hat{\pi}} [\Gamma_b] \right] \geq \left( 1 - \frac{1}{|\mathcal{A}|} \right)^{H+1} \mathbb{E}_{(\pi^*, \mathcal{M}) \sim \mathcal{P}(D)} \left[ \mathbb{E}_{\hat{\pi}} [H - \tau] \right] \quad (46)$$

Note that although the MDP family considered in [22, Lemma A.15] is different, until the state  $b$  is visited the two MDPs are identical and therefore  $\tau$  and  $\tau_b$  are distributed identically under either MDP family for the same policy.  $\square$

#### 7.4 Proof of Theorems 4 and 5

*Proof of Lemma 1.* Conditioned on the expert and the learner playing the same actions in the state, the error of the learner is exactly 0 since in such trajectories both policies collect the same reward. On the other hand, when the learner plays an action different from the expert at a visited state (and thus the 0-1 loss for this trajectory is 1), the maximum error the learner can incur is  $H$ .  $\square$

*Proof of Theorem 5.* Consider the compression based multi-class linear classification algorithm of [8]. This algorithm admits the following guarantee for multi-class sequence classification.

**Theorem 11** (Theorem 5 in [8]). *Consider any linear multi-class classification problem with features  $\phi : X \times Y \rightarrow \mathbb{R}$ . The learner is provided samples  $D = \{(x_1, y_1), \dots, (x_n, y_n)\}$ , where each  $x_i$  is sampled from an unknown distribution  $\rho$  and with label  $y_i = \arg \max_{y \in Y} \langle \phi(x, y), \theta^* \rangle$  for an unknown  $\theta^* \in \mathbb{R}^d$ . Then, if  $n \geq \frac{d \log(1/\epsilon) + \log(1/\delta)}{\epsilon}$ , with probability  $\geq 1 - \delta$  the compression algorithm of [8] returns a linear function  $\hat{\theta} \in \mathbb{R}^d$  such that, the expected 0-1 loss is bounded by  $\epsilon$ . Namely,*

$$\mathbb{E}_{x \sim \rho} \left[ \mathbb{1} \left( \arg \max_{y \in Y} \langle \phi(x, y), \hat{\theta} \rangle \neq \arg \max_{y \in Y} \langle \phi(x, y), \theta^* \rangle \right) \right] \leq \epsilon \quad (47)$$

Consider the dataset as a mapping from sequences of states to sequences of actions  $\mathcal{S}^H \rightarrow \mathcal{A}^H$ . In addition, the expert policy can be thought of as a classifier from  $\mathcal{S}^H \rightarrow \mathcal{A}^H$ . In the sense: for  $\mathcal{S}^H \ni (s_1, \dots, s_H) \mapsto (\pi_1^*(s_1), \dots, \pi_H^*(s_H)) \in \mathcal{A}^H$ . A sequence classifier is a mapping from  $\mathcal{S}^H \rightarrow \mathcal{A}^H$ .

A sequence linear classifier is defined as: For  $\theta \in \mathbb{R}^d$ , the corresponding linear sequence classifier is  $(s_1, \dots, s_H) \mapsto \arg \max_{a_1, \dots, a_H \in \mathcal{A}} \theta \mapsto \left\langle \theta, \sum_{t=1}^H \phi_t(s_t, a_t) \right\rangle$ . Define the set of linear sequence classifiers corresponding to  $\theta \in \mathbb{R}^d$ . Then, the following two propositions are true:

- i. The expert policy  $\pi^*$  is a linear sequence classifier under the linear-expert assumption. At any time  $t$  and state  $s_t$ , the expert chooses the action  $a_t = \arg \max_{a \in \mathcal{A}} \langle \theta_t^*, \phi_t(s, a) \rangle$ . Summing across any sequence of states  $(s_1, \dots, s_H)$  shows that the sequence of actions played by the expert satisfies:  $(a_1, \dots, a_H) = \arg \max_{a'_1, \dots, a'_H \in \mathcal{A}} \langle \theta, \phi_t(s_t, a_t) \rangle$  which proves the claim.

- ii. Every sequence linear classifier corresponds to a meaningful Markovian policy. Indeed, for some  $\theta \in \mathbb{R}^d$ , consider the sequence linear classifier corresponding to  $\theta$ . If at a state  $s_t$  at time  $t$ , the classifier does not choose the action  $a_t = \arg \max_{a \in \mathcal{A}} \langle \theta, \phi_t(s_t, a) \rangle$ , then on any sequence that visits the state  $s_t$  at time  $t$ ,  $(a_1, \dots, a_H) \neq \arg \max_{a'_1, \dots, a'_H \in \mathcal{A}} \langle \theta, \phi_t(s_t, a'_t) \rangle$  which leads to a contradiction. Therefore, the sequence linear classifier plays the action  $a_t = \arg \max_{a \in \mathcal{A}} \langle \theta, \phi_t(s_t, a) \rangle$  at each state  $s_t$  at each time  $t$ . It is therefore a Markovian policy.

The implication of these two points is that it suffices to find a sequence linear classification algorithm from  $\mathcal{S}^H \rightarrow \mathcal{A}^H$  with small expected 0-1 error, given a dataset of trajectories from the expert policy. Invoking the algorithm of [8] for linear multi-class classification and Theorem 11 completes the proof shows that indeed there is a linear sequence classifier with expected 0-1 loss upper bounded by  $\frac{(d+\log(1/\delta))\log(N)}{N}$  which completes the proof, invoking Lemma 1.

The proof of Theorem 4 follows immediately as a corollary of Theorem 5, by invoking Remark 2.  $\square$

## 7.5 Reduction of IL to Confidence set linear classification

*Proof of Theorem 6.* By decomposing as the sum of two parts,  $J(\pi^*) = \sum_{t=1}^H \mathbb{E}_{\pi^*} [\mathbf{r}_t(s_t, a_t) \mathbb{1}(\mathcal{E}_{D_0})] + \sum_{t=1}^H \mathbb{E}_{\pi^*} [\mathbf{r}_t(s_t, a_t) \mathbb{1}(\mathcal{E}_{D_0}^c)]$  and recalling from definition that  $\tilde{J}_r(\pi^*, D) = \mathbb{E}_{\pi^*} [\sum_{t=1}^H \mathbf{r}_t(s_t, a_t) \mathbb{1}(\mathcal{E}_{D_0})] + \mathbb{E}_{\text{tr} \sim \text{Unif}(D_1)} [\sum_{t=1}^H \mathbf{r}_t(s_t, a_t) \mathbb{1}(\mathcal{E}_{D_0}^c)]$ , we have that,

$$J(\pi^*) - \tilde{J}(\pi^*, D) = \sum_{t=1}^H \mathbb{E}_{\pi^*} [\mathbf{r}_t(s_t, a_t) \mathbb{1}(\mathcal{E}_{D_0}^c)] - \mathbb{E}_{\text{tr} \sim \text{Unif}(D_1)} [\mathbf{r}_t(s_t, a_t) \mathbb{1}(\mathcal{E}_{D_0}^c)]. \quad (48)$$

By the linear reward assumption,

$$J(\pi^*) - \tilde{J}(\pi^*, D) = \sum_{t=1}^H \langle \mathbb{E}_{\pi^*} [\phi_t(s_t, a_t) \mathbb{1}(\mathcal{E}_{D_0}^c)] - \mathbb{E}_{\text{tr} \sim \text{Unif}(D_1)} [\phi_t(s_t, a_t) \mathbb{1}(\mathcal{E}_{D_0}^c)], \theta_t \rangle \quad (49)$$

Therefore, by Holder's inequality,

$$\sup_{(\theta_1, \dots, \theta_t): \forall t, \|\theta_t\|_\infty \leq 1} \left| J(\pi^*) - \tilde{J}(\pi^*, D) \right| \quad (50)$$

$$\leq \sum_{t=1}^H \left\| \mathbb{E}_{\pi^*} [\phi_t(s_t, a_t) \mathbb{1}(\mathcal{E}_{D_0}^c)] - \mathbb{E}_{\text{tr} \sim \text{Unif}(D_1)} [\phi_t(s_t, a_t) \mathbb{1}(\mathcal{E}_{D_0}^c)] \right\|_1. \quad (51)$$

Taking expectation over  $D_0$  and  $D_1$ ,

$$\mathbb{E}_D \left[ \sup_{(\theta_1, \dots, \theta_t): \forall t, \|\theta_t\|_\infty \leq 1} \left| J(\pi^*) - \tilde{J}(\pi^*, D) \right| \right] \quad (52)$$

$$\leq \sum_{t=1}^H \mathbb{E}_{D_0} \left[ \mathbb{E}_{D_1} \left[ \left\| \mathbb{E}_{\pi^*} [\phi_t(s_t, a_t) \mathbb{1}(\mathcal{E}_{D_0}^c)] - \mathbb{E}_{\text{tr} \sim \text{Unif}(D_1)} [\phi_t(s_t, a_t) \mathbb{1}(\mathcal{E}_{D_0}^c)] \right\|_1 \middle| D_0 \right] \right] \quad (53)$$

$$\stackrel{(i)}{\leq} \sum_{t=1}^H \sum_{i=1}^d \mathbb{E}_{D_0} \left[ \mathbb{E}_{D_1} \left[ \left| \mathbb{E}_{\pi^*} [\langle \phi_t(s_t, a_t), e_i \rangle \mathbb{1}(\mathcal{E}_{D_0}^c)] - \mathbb{E}_{\text{tr} \sim \text{Unif}(D_1)} [\langle \phi_t(s_t, a_t), e_i \rangle \mathbb{1}(\mathcal{E}_{D_0}^c)] \right| \middle| D_0 \right] \right] \quad (54)$$

$$\stackrel{(ii)}{\leq} \sum_{t=1}^H \sum_{i=1}^d \mathbb{E}_{D_0} \left[ \left( \mathbb{E}_{D_1} \left[ \left( \mathbb{E}_{\pi^*} [\langle \phi_t(s_t, a_t), e_i \rangle \mathbb{1}(\mathcal{E}_{D_0}^c)] - \mathbb{E}_{\text{tr} \sim \text{Unif}(D_1)} [\langle \phi_t(s_t, a_t), e_i \rangle \mathbb{1}(\mathcal{E}_{D_0}^c)] \right)^2 \middle| D_0 \right] \right)^{1/2} \right] \quad (55)$$

where in (i)  $\{e_i : i = 1, \dots, d\}$  denotes the set of standard basis vectors in  $\mathbb{R}^d$  and (ii) follows by Jensen's inequality. Finally, we bound the variance of the expected feature under  $\pi^*$  and under the

empirical distribution in  $D_1$  (which is independent of  $D_0$ ). For each  $i = 1, \dots, d$  the variance of the co-ordinate  $i$  is upper bounded by  $\frac{\mathbb{E}_{\pi^*} [\langle \phi_t(s_t, a_t), e_i \rangle^2 \mathbb{1}(\mathcal{E}_{D_0}^c)]}{N}$ . Therefore,

$$\mathbb{E}_D \left[ \sup_{(\theta_1, \dots, \theta_t): \forall t, \|\theta_t\|_\infty \leq 1} \left| J(\pi^*) - \tilde{J}(\pi^*, D) \right| \right] \quad (56)$$

$$\leq \sqrt{\frac{1}{N}} \sum_{t=1}^H \sum_{i=1}^d \mathbb{E}_{D_0} \left[ \left( \mathbb{E}_{\pi^*} [\langle \phi_t(s_t, a_t), e_i \rangle^2 \mathbb{1}(\mathcal{E}_{D_0}^c)] \right)^{1/2} \right] \quad (57)$$

$$\stackrel{(i)}{\leq} \sum_{t=1}^H \sqrt{\frac{d}{N}} \left( \sum_{i=1}^d \mathbb{E}_{D_0} \left[ \mathbb{E}_{\pi^*} [|\langle \phi_t(s_t, a_t), e_i \rangle| \mathbb{1}(\mathcal{E}_{D_0}^c)] \right] \right)^{1/2} \quad (58)$$

$$= \sum_{t=1}^H \sqrt{\frac{d}{N}} \left( \mathbb{E}_{D_0} \left[ \mathbb{E}_{\pi^*} [\|\phi_t(s_t, a_t)\|_1 \mathbb{1}(\mathcal{E}_{D_0}^c)] \right] \right)^{1/2} \quad (59)$$

$$\leq H \sqrt{\frac{d}{N} \mathbb{E}_{D_0} [\Pr(\mathcal{E}_{D_0}^c)]} \quad (60)$$

where (i) follows by an application of Jensen's inequality and Cauchy Schwarz inequality and using the fact that  $\langle \phi_t(s_t, a_t), e_i \rangle^2 \leq |\langle \phi_t(s_t, a_t), e_i \rangle|$  since  $\|\phi_t(s_t, a_t)\|_1 \leq 1$ . By union bounding,  $\Pr(\mathcal{E}_{D_0}^c) \leq \sum_{t=1}^H \Pr_{\pi^*}(s_t \notin \mathcal{X}_t(D_0))$  which completes the proof of Theorem 6.  $\square$

## 7.6 Proof of Lemma 2

*Proof of Lemma 2.* (i) Observe that for any input  $x$ , if  $x \in \mathcal{K}$ , then  $x = \phi(x, 0) - \phi(x, 1)$  can almost surely be expressed as a positive linear combination  $\sum_{i=1}^N \omega_i (-1)^{y_i} x_i$  where  $\omega_i \geq 0$  for all  $i = 1, \dots, N$ . Since  $\langle \theta^*, (-1)^{y_i} x_i \rangle \geq 0$ , this implies that  $\langle \theta^*, x \rangle \geq 0$  and consequently the expert must have classified  $x$  as 0. Likewise, for inputs  $x \in -\mathcal{K}$ ,  $x$  can be expressed as a linear combination of  $(-1)^{y_i} x_i$ 's with non-positive coefficients. Consequently, the expert must have played the action 1 at these states.

(ii) We first prove that for any  $x \notin \mathcal{C}$ , one can find a  $\theta' \in \Theta$  such that  $\text{sgn}(\langle \theta', x \rangle) \neq \text{sgn}(\langle \theta^*, x \rangle)$ . We subsequently prove that this result implies the required statement.

This result can be proved in two cases: (a) assuming that  $\langle \theta^*, x \rangle \geq 0$  and (b)  $\langle \theta^*, x \rangle < 0$ . We prove the case (a) and argue that (b) follows similarly. For any input  $x$  such that  $\langle \theta^*, x \rangle \geq 0$ , we show that there exists a  $\theta \in \mathcal{K}$  such that  $\langle \theta, x \rangle < 0$ .

Indeed, in the first case, for any point  $v \in \mathcal{K}$ , the inner product  $\langle v, \theta \rangle \geq 0$  for all  $\theta \in \mathcal{K}$ . This is because for every  $\theta \in \mathcal{K}$ ,  $\langle \theta, (-1)^{y_i} x_i \rangle \geq 0$  and every point  $v$  in  $\mathcal{K}$  is a positive linear combination of the  $(-1)^{y_i} x_i$  vectors. Therefore if  $v \notin \mathcal{C}$ , there exists a  $\theta \in \mathcal{K}$  such that  $\langle v, \theta \rangle < 0$  and choosing  $v = x$  completes the proof.

Finally, observe that from (i), and since  $\theta^* \in \mathcal{K}$  (we can repeat the same argument for  $-\mathcal{K}$ ), the cone  $\mathcal{K}$  corresponds to  $\bigcap_{\theta \in \Theta} \mathcal{H}_\theta$ , where  $\mathcal{H}_\theta$  is the halfspace  $\{x \in \mathbb{R}^d : \langle x, \theta \rangle \geq 0\}$ . For any  $\hat{\theta}$ , by using Demorgan's laws it is easy to verify that this is a superset of  $(\bigcup_{\theta \in \Theta} (\mathcal{V} \cap \mathcal{H}_\theta)^c)^c$  for any  $\mathcal{V} \subseteq X$ . In particular, choosing  $\mathcal{V}$  as the set of inputs classified as 0 by  $\hat{h}$  (output classifier of the confidence set linear classification algorithm), we see that for every input  $x \notin \mathcal{K}$ , there exists a  $\theta \in \Theta$  such that  $x \in \mathcal{V}$  but  $x \notin \mathcal{H}_\theta$  or vice versa. Namely the output predicted by  $\hat{h}$  on the point  $x$  differs from the output predicted by some linear classifier  $\theta \in \Theta$ .  $\square$

## 7.7 Proof of Theorem 7

**Theorem 12.** Fix  $d > 1$ . Draw points  $S_0, S_1, S_2, \dots, S_{n+1}$  uniform random points on a half sphere in  $\mathbb{R}^d$ . Let  $\mathcal{K}$  be the convex cone spanned by  $S_1, \dots, S_n$ . Then there exists universal constants  $c, C > 0$  such that

$$\frac{cd^{3/2}}{n\sqrt{\log d}} \leq \mathbb{P}[S_0 \notin \mathcal{K}] \leq \frac{Cd^{3/2} \log d}{n} \quad (61)$$

holds for sufficiently large  $n$ .




Figure 3: The red points represent  $S_1, \dots, S_n$ . The green point  $w \in \mathcal{K}$  satisfies  $\langle z \cos \alpha - y \sin \alpha, w \rangle = 0$

*Proof.* First, observe that it is sufficient to prove the bound (61) (with possibly different constants  $c, C$ ) when  $S_1, \dots, S_n$  are replaced by points in a Poisson process the half sphere with intensity  $\frac{2n}{\text{vol}(\mathbb{S}^{d-1})}$ . Indeed, this can be show using the facts that  $\mathbb{P}[S_0 \notin \mathcal{K}]$  is monotonically decreasing in  $n$  and that there are more than  $2n$  or fewer than  $n/2$  Poisson points with exponentially small probability.

Henceforth, we shall consider the Poisson version of the problem. Note that the boundary of the half sphere is  $\mathbb{S}^{d-2}$ . Denote by  $z$  the north pole of the hemisphere. For each  $y \in \mathbb{S}^{d-2}$ , define

$$\alpha_y := \sup\{\alpha \in [0, \pi/2) : y \cos \alpha + z \sin \alpha \notin \mathcal{K}\}. \quad (62)$$

Moreover, define

$$\tilde{\alpha}_y := \sup\{\alpha \in [0, \pi/2) : \langle z \cos \alpha - y \sin \alpha, w \rangle \geq 0, \forall w \in \mathcal{K}\}. \quad (63)$$

Then the Hahn-Banach theorem, we see that  $\alpha_y$  is the smallest value such that

$$\langle y \cos \alpha_y + z \sin \alpha_y, -\tilde{y} \sin \tilde{\alpha}_{\tilde{y}} + z \cos \tilde{\alpha}_{\tilde{y}} \rangle \geq 0, \quad \forall \tilde{y}. \quad (64)$$

Equivalently,

$$\tan \alpha_y = \sup_{\tilde{y}} \langle y, \tilde{y} \rangle \tan \tilde{\alpha}_{\tilde{y}}. \quad (65)$$

The significance of  $\alpha_y$  is the following: suppose  $\kappa := \sup_y \tilde{\alpha}_y$ . Then  $\kappa = \sup_y \alpha_y$  in view of (65), and we have

$$\frac{2(1 - \cos^{d-1} \kappa)}{\text{vol}(\mathbb{S}^{d-1})} \int_{y \in \mathbb{S}^{d-2}} \alpha_y dy \leq \mathbb{P}[S_0 \notin \mathcal{K} | \mathcal{K}] \leq \frac{2}{\text{vol}(\mathbb{S}^{d-1})} \int_{y \in \mathbb{S}^{d-2}} \alpha_y dy. \quad (66)$$

Thus, as long as  $\kappa$  is small, the problem is reduced to estimating  $\int_{y \in \mathbb{S}^{d-2}} \alpha_y dy$ . Note that there exists a constant  $c_1 > 0$  (which is allowed to depend on  $d$ ) and a vanishing sequence  $\kappa_1, \kappa_2, \dots \rightarrow 0$  such that with probability at least  $1 - e^{-c_1 n}$ , we have  $\kappa < \kappa_n$ . Indeed, this can be shown using by choosing an  $\epsilon$ -net on  $\mathbb{S}^{d-2}$  with  $\epsilon = 0.1$  (say), and arguing that with high probability each of  $\epsilon$ -neighborhood in the hemisphere of a point in the net includes a Poisson point; the details are omitted.

Conditioned on  $\kappa < \kappa_n$ , (65) implies

$$\alpha_y = (1 + o_{\kappa_n}(1)) \sup_{\tilde{y}} \langle y, \tilde{y} \rangle \tilde{\alpha}_{\tilde{y}} \quad (67)$$

where explicitly,  $1 + o_{\kappa_n}(1)$  here is a number between  $\sup_{0 < \alpha < \kappa_n} \frac{\tan \alpha}{\alpha}$  and its inverse. Then

$$\mathbb{P}[S_0 \notin \mathcal{K}] = \mathbb{P}[S_0 \notin \mathcal{K} | \kappa < \kappa_n] + O(e^{-c_1 n}) \quad (68)$$

$$= \frac{2(1 + o_{\kappa_n}(1))}{\text{vol}(\mathbb{S}^{d-1})} \int_{y \in \mathbb{S}^{d-2}} \mathbb{E}[\alpha_y | \kappa < \kappa_n] dy + O(e^{-c_1 n}) \quad (69)$$

$$= \frac{2(1 + o_{\kappa_n}(1))}{\text{vol}(\mathbb{S}^{d-1})} \int_{y \in \mathbb{S}^{d-2}} \mathbb{E}[\sup_{\tilde{y}} \langle y, \tilde{y} \rangle \tilde{\alpha}_{\tilde{y}} | \kappa < \kappa_n] dy + O(e^{-c_1 n}). \quad (70)$$

Therefore, the theorem is proved if we show that

$$\frac{cd^{1.5}}{n\sqrt{\log d}} \leq \frac{1}{\text{vol}(\mathbb{S}^{d-1})} \int_{y \in \mathbb{S}^{d-2}} \mathbb{E}[\sup_{\tilde{y}} \langle y, \tilde{y} \rangle \tilde{\alpha}_{\tilde{y}} | \kappa < \kappa_n] dy \leq \frac{Cd^{1.5} \log d}{n} \quad (71)$$

for possibly different constants  $c, C$ .

Now let  $\phi$  be the map from the hemisphere to  $\mathbb{S}^{d-2} \times [0, \frac{\text{vol}(\mathbb{S}^{d-1})}{2\text{vol}(\mathbb{S}^{d-2})})$  which is measure preserving and satisfies  $\phi_1(y \cos \alpha + z \sin \alpha) = y$  for any  $y \in \mathbb{S}^{d-2}$  and  $\alpha \in [0, \pi/2)$ , where  $\phi_1, \phi_2$  denote the first and the second coordinates of the value of  $\phi$ . We have

$$(\cos^{d-1} \alpha) \cdot \alpha \leq \phi_2(y \cos \alpha + z \sin \alpha) \leq \alpha. \quad (72)$$

Under  $\phi$ , the images of the Poisson points on the hemisphere still behave as a Poisson point process with intensity  $\frac{2n}{\text{vol}(\mathbb{S}^{d-1})}$ .

Now define for each  $y \in \mathbb{S}^{d-2}$ ,

$$\tilde{\theta}_y := \sup\{t \geq 0 : \text{no } \phi\text{-image of Poisson point in } V(y, t)\} \quad (73)$$

where  $V(y, t)$  is defined as the set of  $(a, b) \in \mathbb{S}^{d-2} \times [0, \infty)$  such that  $b < t \cos \angle(y, a)$ . Define for each  $x \in \mathbb{S}^{d-2}$ ,

$$\theta_x := \sup_{y \in \mathbb{S}^{d-2}} \tilde{\theta}_y \cos \angle(x, y). \quad (74)$$

Using linear algebra, we can show that given  $\kappa < \kappa_n$ , we have

$$\tilde{\theta}_y = (1 + o_{\kappa_n}(1)) \tilde{\alpha}_n \quad (75)$$

Therefore, by the observation in (71), we prove the theorem if we show the following

$$\frac{cd^{1.5}}{n\sqrt{\log d}} \leq \frac{1}{\text{vol}(\mathbb{S}^{d-1})} \int_{y \in \mathbb{S}^{d-2}} \mathbb{E}[\sup_{\tilde{y}} \langle y, \tilde{y} \rangle \tilde{\theta}_{\tilde{y}} | \kappa < \kappa_n] dy \leq \frac{Cd^{1.5} \log d}{n}. \quad (76)$$

By symmetry and  $\text{vol}(\mathbb{S}^{d-2})/\text{vol}(\mathbb{S}^{d-1}) = \Theta(d^{1/2})$ , the above is equivalent to

$$\frac{cd}{n\sqrt{\log d}} \leq \mathbb{E}[\sup_{\tilde{y}} \langle y, \tilde{y} \rangle \tilde{\theta}_{\tilde{y}} | \kappa < \kappa_n] \leq \frac{Cd \log d}{n}. \quad (77)$$

Since  $\mathbb{P}[\kappa < \kappa_n] \geq 1 - e^{-c_1 n}$ , it suffices to show the above bound with two modifications of the problem: 1) in the definition of  $\tilde{\theta}_y$ , the Poisson process is extended to one with the same density  $\frac{2n}{\text{vol}(\mathbb{S}^{d-1})}$  but on  $\mathbb{S}^{d-2} \times [0, \infty)$ ; 2) uncondition on  $\kappa < \kappa_n$ , namely,

$$\frac{cd}{n\sqrt{\log d}} \leq \mathbb{E}[\sup_{\tilde{y}} \langle y, \tilde{y} \rangle \tilde{\theta}_{\tilde{y}}] \leq \frac{Cd \log d}{n}. \quad (78)$$

for possibly different universal constants  $c, C > 0$ . Now (78) is the content of Lemma 8; the difference in the factor  $n$  is because we consider a Poisson process with intensity  $\frac{2}{\text{vol}(\mathbb{S}^{d-1})}$  instead in the lemma. The Theorem is now proved.  $\square$

We shall establish the following lemma; the result for the original problem about random points on the half sphere will then easily follow. Note that a nice feature about the lemma is that there is no  $N$  in the problem; the problem is purely depending on the dimension  $d$ .

**Lemma 8.** Consider a Poisson point process on  $\mathbb{S}^{d-2} \times [0, \infty)$  with intensity  $\frac{2}{\text{vol}(\mathbb{S}^{d-1})}$ . For each  $y \in \mathbb{S}^{d-2}$ , define

$$\tilde{\theta}_y := \sup\{t \geq 0 : \text{no Poisson point in } V(y, t)\} \quad (79)$$

where  $V(y, t)$  is defined as the set of  $(a, b) \in \mathbb{S}^{d-2} \times [0, \infty)$  such that  $b < t \cos \angle(y, a)$ . Define for each  $x \in \mathbb{S}^{d-2}$ ,

$$\theta_x := \sup_{y \in \mathbb{S}^{d-2}} \tilde{\theta}_y \cos \angle(x, y). \quad (80)$$

Then

$$\frac{d}{\sqrt{\log d}} \lesssim \mathbb{E}[\theta_x] \lesssim d \log d. \quad (81)$$

*Proof.* We begin by first proving the upper bound.

$$\hat{\theta}_y := \sup\{t \geq 0 : \text{no Poisson point in } \hat{V}(y, t)\} \quad (82)$$

where  $\hat{V}(y, t)$  is defined as the set of  $(a, b) \in \mathbb{S}^{d-2} \times [0, \infty)$  such that  $b < t(\cos \angle(y, a))1\{\angle(y, a) < \frac{\pi}{2} - d^{-1/2}\}$ .

Next, let  $\mathcal{N}$  be an optimal  $\epsilon$ -covering of  $\mathbb{S}^{d-2}$ , where  $\epsilon = \frac{1}{4}d^{-1/2}$ . For each  $y$  and  $t$ ,

$$\mathbb{P}[\hat{\theta}_y > t] = \exp\left(-\text{vol}(\hat{V}(y, t)) \times \frac{2}{\text{vol}(\mathbb{S}^{d-1})}\right). \quad (83)$$

However,

$$\text{vol}(\hat{V}(y, t)) = \text{vol}((\cos d^{-1/2})B^{d-2}) = \cos^{d-2} d^{-1/2} \text{vol}(B^{d-2}). \quad (84)$$

Define  $c_d := \cos^{d-2} d^{-1/2}$ . Then

$$\mathbb{P}[\hat{\theta}_y > t] = \exp\left(-2tc_d \frac{\text{vol}(B^{d-2})}{\text{vol}(\mathbb{S}^{d-1})}\right) \quad (85)$$

$$= \exp\left(-\frac{2tc_d \text{vol}(\mathbb{S}^{d-3})}{d-2 \text{vol}(\mathbb{S}^{d-1})}\right) \quad (86)$$

$$= \exp\left(-\frac{tc_d}{\pi}\right). \quad (87)$$

Therefore by the union bound,

$$\mathbb{P}[\sup_{y \in \mathcal{N}} \hat{\theta}_y > t] \leq |\mathcal{N}| \exp\left(-\frac{tc_d}{\pi}\right) \quad (88)$$

$$\leq (3/\epsilon)^{d-2} \exp\left(-\frac{tc_d}{\pi}\right) \quad (89)$$

Let  $T: \mathbb{S}^{d-2} \rightarrow \mathcal{N}$  be such that  $\|T(y) - y\| \leq \epsilon$ . Define  $\hat{t} := \sup_{y \in \mathcal{N}} \hat{\theta}_y$ . For any  $y \in \mathbb{S}^{d-2}$ , by the definition of  $\hat{t}$ , we can find a Poisson point  $(a, b)$  such that

$$b \leq \hat{t} \cos \angle(T(y), a)1\{\angle(T(y), a) < \frac{\pi}{2} - d^{-1/2}\}. \quad (90)$$

Then, since

$$\frac{\cos \angle(T(y), a)1\{\cos \angle(T(y), a) < \frac{\pi}{2} - d^{-1/2}\}}{\cos \angle(y, a)} \leq \frac{\cos(\frac{\pi}{2} - d^{-1/2})}{\cos(\frac{\pi}{2} - d^{-1/2} + \epsilon')} \quad (91)$$

$$\leq \frac{\cos(\frac{\pi}{2} - d^{-1/2})}{\cos(\frac{\pi}{2} - \frac{1}{2}d^{-1/2})} \quad (92)$$

$$\leq 3 \quad (93)$$



for large  $d$ , where  $\epsilon'$  denotes the angle between two points with distance  $\epsilon$  on the sphere, we obtain

$$b \leq 3\hat{t} \cos \angle(y, a) \quad (94)$$

which implies that  $\tilde{\theta}_y \leq 3\hat{t}$ . Since this bound does not depend on  $y$ , we also obtain  $\theta_x \leq 3\hat{t}$ . This shows that

$$\mathbb{E}[\theta_x] \leq 3\mathbb{E}[\hat{t}] \leq 3 \int_0^\infty \min \left\{ (3/\epsilon)^{d-2} \exp(-\frac{t}{6\pi}), 1 \right\} dt \quad (95)$$

$$\lesssim \ln((3/\epsilon)^{d-2}) \quad (96)$$

$$\lesssim d \ln d \quad (97)$$

where we used  $c_d \geq e^{-1/2} \geq 1/3$  for large  $d$ .

Now we move on to proving the lower bound. First, note that for any  $t > 0$ , the number of Poisson points in  $\mathbb{S}^{d-2} \times [0, t]$  follows the Poisson distribution with strength bounded by:

$$\lambda = 2t \frac{\text{vol}(\mathbb{S}^{d-2})}{\text{vol}(\mathbb{S}^{d-1})} \quad (98)$$

$$\leq c\sqrt{dt} \quad (99)$$

where  $c > 0$  is a universal constant. In particular, if

$$t = 0.1d(c\sqrt{d})^{-1}, \quad (100)$$

then the expected number of Poisson points is  $\lambda = 0.1d$ . By the Markov inequality, the number is smaller than  $0.5d$  with probability at least 0.8.

Let  $\mathcal{E}_1$  be the event that there are at most  $0.5d$  Poisson points in  $\mathbb{S}^{d-2} \times [0, t]$ , and denote by  $\mathcal{P}_t$  the set of these Poisson points. Define

$$\mathcal{H} := \{h \in \mathbb{S}^{d-2} : \langle h, p \rangle \leq 0, \forall p \in \mathcal{P}_t\}. \quad (101)$$

We claim that there exists a universal constant  $c_1 \in (0, \pi/2)$  such that  $\mathcal{H}^{c_1}$  has normalized conic area at least 0.5, with probability at least 0.5 conditioned on  $\mathcal{E}_1$ ; call this conditional event  $\mathcal{E}_2$ . Here,  $\mathcal{H}^{c_1}$  is defined as the set of all points on  $\mathbb{S}^{d-2}$  whose angle with some point in  $\mathcal{H}$  is at most  $c_1$ .

In fact, the above claim can be proved if we show that the same claims holds when  $\mathcal{H}$  is replaced by

$$\mathcal{H}' := \{h \in \mathbb{S}^{d-2} : \langle h, p \rangle \leq 0, \forall p \in \mathcal{B}\}. \quad (102)$$

where  $\mathcal{B}$  is the set of  $0.5d$  uniformly randomly drawn points  $p_1, \dots, p_{0.5d}$  on  $\mathbb{S}^{d-2}$ . Let  $p_{0.5d+1}, \dots, p_{d-1}$  be an arbitrary basis of the orthogonal complement of the linear space spanned by  $p_1, \dots, p_{0.5d}$ . With probability at least 0.5, the nonzero singular values of the matrix  $[p_1, \dots, p_{0.5d}]$  are bounded in the interval  $(c_2, c_3)$  where  $c_2, c_3 > 0$  are universal constants; see e.g. <https://djalil.chafai.net/docs/sing.pdf>. It follows that the singular values of  $P := [p_1, p_2, \dots, p_{d-1}]$  is bounded in the same interval. Note that if  $r_1$  and  $r_2$  are two rays with a small angle, then their linear transforms  $Pr_1$  and  $Pr_2$  satisfy  $\frac{\angle(Pr_1, Pr_2)}{\angle(r_1, r_2)} \in (c_4, c_5)$ , where  $c_4, c_5 > 0$  are universal constants explicitly defined by  $c_2, c_3$ . It follows that under the linear transform, conic areas are changed by a factor within  $(c_5^{-(d-2)}, c_4^{-(d-2)})$ . Then we note that the cone of  $\mathcal{H}'$  equals  $(P^\top)^{-1}$  applied to  $\{s \in \mathbb{R}^{d-1} : \langle s, e_i \rangle \leq 0, i = 1, \dots, 0.5d\}$ ; indeed, for each  $i$ ,

$$\{s : \langle s, p_i \rangle = 0\} = \{s : \langle s, Pe_i \rangle = 0\} \quad (103)$$

$$= \{s : \langle P^\top s, e_i \rangle = 0\} \quad (104)$$

$$= \{(P^\top)^{-1}s : \langle s, e_i \rangle = 0\}. \quad (105)$$

Therefore, the normalized conic area of  $\mathcal{H}'$  is at least  $c_5^{-(d-2)} 2^{-0.5d}$ . It then follows by the isoperimetry on the sphere that there exists  $c_1 \in (0, \pi/2)$  such that  $\mathcal{H}'^{c_1}$  has normalized conic area at least 0.5. This proves the claim following (101).

Next, recall that  $x$  is a fixed point on  $\mathbb{S}^{d-2}$ . Conditioned on  $\mathcal{E}_1, \mathcal{E}_2$ , by the rotation invariance,  $x \in \mathcal{H}^{c_1}$  with probability at least 0.5; call this conditional event  $\mathcal{E}_3$ . Under  $\mathcal{E}_3$  Let  $y$  be a point in  $\mathcal{H}$  closest

to  $x$ . Then  $\angle(x, y) \leq c_1$ . Note that  $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3$  are determined by the Poisson point configurations in  $\mathbb{S}^{d-2} \times [0, t]$ , which is independent of the configurations in  $\mathbb{S}^{d-2} \times (t, \infty)$ . Now set

$$t' := \frac{\sqrt{d}}{10\sqrt{\log d}}t. \quad (106)$$

We claim that conditioned on  $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3$ ,

$$\theta_y \geq t' \quad (107)$$

with probability at least 0.9; call this conditional event  $\mathcal{E}_4$ . Indeed, the above bound on  $\theta_y$  is equivalent to the event that there is no Poisson point  $(a, b)$  satisfying

$$b < t' \cos \angle(y, a). \quad (108)$$

Note that by the definition of  $y$ , there is no Poisson point for which  $b \in [0, t]$ . Therefore (conditioned on  $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3$ ), (107) is equivalent to the event that there is no Poisson point in

$$\{(a, b) : b \in (t, t' \cos \angle(y, a))\}. \quad (109)$$

The Lebesgue measure of the above region equals  $t' \text{vol}(\sin \arccos \frac{t}{t'} B^{d-2})$ , where  $B^{d-2}$  is the unit ball. Therefore, no Poisson point in (109) with probability

$$\begin{aligned} & \exp\left(-t' \text{vol}(\sin \arccos \frac{t}{t'} B^{d-2}) \cdot \frac{2}{\text{vol}(\mathbb{S}^{d-1})}\right) \\ &= \exp\left(-\frac{t'}{\pi} (\sin \arccos \frac{t}{t'})^{d-2}\right) \end{aligned} \quad (110)$$

$$\geq \exp\left(-\frac{t'}{\pi} \left(1 - \frac{100 \log d}{d}\right)^{\frac{d-2}{2}}\right) \quad (111)$$

$$\geq 0.9 \quad (112)$$

for large enough  $d$ . This proves the claim around (107).

Finally, under all the events  $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3, \mathcal{E}_4$ , which happens with probability at least 0.1, we have

$$\theta_x \geq \cos c_1 \theta_y \geq t' \cos c_1 \quad (113)$$

which is  $\Theta(\frac{d}{\sqrt{\log d}})$ , thus the lower bound is established.  $\square$