On the Effective Horizon of Inverse Reinforcement Learning

Yiqing Xu¹ Finale Doshi-velez² David Hsu¹

Abstract

Inverse reinforcement learning (IRL) algorithms often rely on (forward) reinforcement learning or planning over a given time horizon to compute an approximately optimal policy for a hypothesized reward function and then match this policy with expert demonstrations. The time horizon plays a critical role in determining both the accuracy of reward estimates and the computational efficiency of IRL algorithms. Interestingly, an effective time horizon shorter than the ground-truth value often produces better results faster. This work formally analyzes this phenomenon and provides an explanation: the time horizon controls the complexity of an induced policy class and mitigates overfitting with limited data. This analysis serves as a guide for the principled choice of the effective horizon for IRL. It also prompts us to re-examine the classic IRL formulation: it is more natural to learn jointly the reward and the effective horizon rather than the reward alone with a given horizon. To validate our findings, we implement a crossvalidation extension and the experimental results confirm the theoretical analysis.

1. Introduction

Inverse reinforcement learning (IRL) (Ng & Russell, 2000) aims to infer the underlying task objective from expert demonstrations. One common approach is to estimate a reward function that induces a policy matching closely the demonstrated expert trajectories. This model-based approach holds the promise of generalizing the learned reward function and the associated policy over states not seen in the demonstrations (Osa et al., 2018).

Many existing IRL algorithms follow the classic formulation and assume a known ground-truth discount factor (or equivalently, time horizon) for the expert demonstrations (Ng & Russell, 2000; Abbeel & Ng, 2004; Ramachandran & Amir, 2007; Ziebart et al., 2008; Boularias et al., 2011; Levine et al., 2011; Wulfmeier et al., 2016; Pirotta & Restelli, 2016; Finn et al., 2016b;a; Ho & Ermon, 2016; Fu et al., 2018; Ni et al., 2020; Ke et al., 2020; Ramponi et al., 2020; Metelli et al., 2021; Hoshino et al., 2022). They estimate the reward function based on this time horizon. However, we have the flexibility to choose a different time horizon when learning the reward function and optimizing the policy.

Surprisingly, we find that using a *horizon* shorter than the ground-truth value often produces better results faster, especially when expert data is scarce. Why? Intuitively, the time horizon controls the complexity of an induced policy class. With limited data, a shorter time horizon is preferred, as the induced policy class is simpler and mitigates overfitting. We refer to the horizon or discount factor used during learning as the *effective horizon* or *effective discount factor*, as it holds promise to enhance learning effectiveness.

We present a formal analysis showing that, with limited expert data, using a shorter discount factor or time horizon improves the generalization of the learned reward function to unseen states. In IRL settings where the effective horizon varies, the performance gap between the induced policy and the expert policy arises from two sources: (i) reward estimation error due to limited data during IRL, and (ii) policy optimization error from using an effective horizon shorter than the ground-truth. We prove that the effective horizon controls the complexity of the approximated policy class during IRL. As the horizon increases, reward estimation error grows-overfitting occurs because we estimate policies from a more complex class using limited expert data. On the other hand, as the horizon approaches the ground-truth, the policy optimization error decreases. These opposing errors suggest that an intermediate effective horizon balances this trade-off and produces the most expert-like policy.

Based on our theoretical findings, we propose a more natural and higher-performing formulation of the IRL problem: jointly learning the reward function and the effective horizon/discount factor. To validate our theoretical results, we extend the linear programming IRL algorithm (LP-IRL) (Ng & Russell, 2000) and the maximum entropy IRL algorithm (MaxEnt-IRL) (Ziebart et al., 2008) using cross-validation. Our experimental evaluation of both LP-IRL and MaxEnt-IRL across four different tasks corroborates our theoretical analysis.

¹National University of Singapore ²Harvard University. Correspondence to: Yiqing Xu <xuyiqing@comp.nus.edu.sg>.

Preliminary work. Under review.

This work presents the first formal analysis of the relationship between the time horizon and the performance of the learned reward function in IRL. Previous studies have examined the impact of the time horizon in Reinforcement Learning (RL) (Guo et al., 2022; Hu et al., 2022; Laidlaw et al., 2023; Amit et al., 2020), approximate dynamic programming (Petrik & Scherrer, 2008), and planning (Jiang et al., 2015; Mattingley et al., 2011), but the effects when the reward function is unknown and inferred from data remain under-explored. Our work addresses this gap by providing a theoretical foundation for how changes in the horizon affect policy performance in IRL. This analysis is timely, as the IRL community is increasingly interested in leveraging different horizons to enhance policy outcomes. Some IRL algorithms employ smaller effective time horizons for computational efficiency (MacGlashan & Littman, 2015; Lee et al., 2022; Xu et al., 2022), while others learn discount factors from data to better align with demonstrations, sometimes incorporating multiple feedback types (Giwa & Lee, 2021; Ghosal et al., 2023). By formally examining the role of varying time horizons, our work complements existing methods and offers valuable insights to guide future IRL research.

2. Related works

Effective Horizon of Imitation Learning Imitation learning learns desired behaviors by imitating expert demonstrations and comprises two classes of methods: model-free behavior cloning (BC) and model-based inverse reinforcement learning (IRL) (Osa et al., 2018). The primary distinction between BC and IRL lies in the horizons used to align the learned behaviors with expert data. BC matches step-wise expert actions, resulting in poor generalization to unseen states. In contrast, IRL addresses this issue by either matching multi-step trajectory distributions (Ziebart et al., 2008; Boularias et al., 2011; Levine et al., 2011; Wulfmeier et al., 2016; Finn et al., 2016b;a; Pirotta & Restelli, 2016; Ramponi et al., 2020) or their marginalized approximations (Ho & Ermon, 2016; Fu et al., 2018; Ni et al., 2020; Ke et al., 2020; Ghasemipour et al., 2019; Hoshino et al., 2022). The former employs a double-loop structure to interleave the policy optimization and reward function update, while the latter learns a discriminator to distinguish expert-like behaviors. Both approaches utilize the ground-truth horizon/discount factor for policy optimization, ensuring global temporal consistency between the learned policy and expert. Notably, few IRL methods adopt receding horizons to reduce computational cost(MacGlashan & Littman, 2015; Lee et al., 2022; Xu et al., 2022), claiming shorter optimization horizons yields sub-optimal policies. Moreover, several works focus on finding the optimal discount factors in an IRL context, sometimes incorporating different human feedback types (Ghosal et al., 2023; Giwa & Lee, 2021). However, a theoretical analysis on the impact of the horizon in IRL is lacking, making it an important yet overlooked consideration in the field.

Theoretical Analysis on Effective Horizon There has been significant research analyzing the impact of planning horizons on Reinforcement Learning (RL) (Guo et al., 2022; Hu et al., 2022; Laidlaw et al., 2023; Amit et al., 2020), approximate dynamic programming (Petrik & Scherrer, 2008), and planning (Mattingley et al., 2011), especially with transition models estimated from samples (Jiang et al., 2015). In these studies, where the reward function is known, policy performance is directly affected by differences in planning horizons. However, the impact of horizons when the reward function is learned from data remains under-explored. Our work examines how changes in the horizon affect policy performance when the reward function is derived from expert demonstrations. Under our setting, the policy optimization depends not only on the horizon but also on reward function estimation, which is in turn also horizon-dependent.

This dual dependency introduces unique analytical challenges not present in forward RL or planning. The key challenge lies in the second part: analyzing how varying discount factors impact reward function estimates when expert data is limited. While Metelli et al. (2021) defines a set of reward functions compatible with limited expert demonstrations under a fixed, known discount factor, we extend this formulation to settings with unknown and varying horizons. By considering this extended class of compatible reward functions, we analyze how horizon affects reward estimation and overall policy optimization performance, particularly when expert demonstrations are scarce.

3. Problem formulation

We consider an MDP (S, A, P, R_0, γ_0) , where S and A represent the state and action spaces, respectively. The transition function is denoted by $P: S \times A \times S \rightarrow [0, 1]$, and the ground-truth reward function is $R_0: S \times A \rightarrow [0, R_{max}]$. The discount factor, γ_0 , implicitly determines the value of future rewards at the current time step. The optimal policy, $\pi^*_{R_0,\gamma_0}$, maximizes the total discounted reward based on R_0 and γ_0 . In our setting, we are given the MDP without the reward function R_0 or the discount factor γ_0 . Instead, we have a set of N expert demonstrated state-action pairs $D = \{(s_0, a_0), (s_1, a_1), \dots (s_{N-1}, a_{N-1})\}$ sampled from $\pi^*_{R_0,\gamma_0}$. We assume the expert policy is deterministic. Moreover, the MDP is assumed to be ergodic, such that any state is reachable from any other state by following a suitable policy. This underlying interconnectivity within the state space allows for the deduction of unobserved states' value by propagating information obtained from observed states through the transition function. Our analysis employs the

discount factor for simplicity; however, the findings are equally applicable to the planning horizon since both determine how much future rewards are valued. A smaller discount factor effectively limits the agent's planning horizon by diminishing the importance of distant rewards.

We examine the scenario where both the reward function and discount factor $(R, \hat{\gamma})$ are jointly learned from the limited expert demonstrations. The scarcity of data suggests that $(\widehat{R}, \widehat{\gamma})$ is susceptible to approximation errors, which consequently affects the induced optimal policy $\pi_{\widehat{R}}^*$. We measure the quality of the $(\widehat{R},\widehat{\gamma})$ pair by comparing the performance of its induced policy $\pi^*_{\widehat{R},\widehat{\gamma}}$ with that of the ground-truth optimal policy $\pi^*_{R_0,\gamma_0}$, both evaluated under the ground-truth (R_0,γ_0) for fair comparison. Formally, we define the loss as the performance difference between the induced and ground-truth optimal policies: $\left\| V_{R_0,\gamma_0}^{\pi^*_{R_0,\gamma_0}} - V_{R_0,\gamma_0}^{\pi^*_{R,\hat{\gamma}}} \right\|_{\infty}$, where $V_{R,\gamma}^{\pi}$ represents the value function of policy π evaluated under (R, γ) . The "best" policy $\pi^*_{\widehat{B},\widehat{\gamma}}$ is the one that minimizes this loss. We use this loss to theoretically study the inter-dependency between the amount of expert data, the effective discount factor utilized for reward learning, and the performance of the policy induced by these factors. In practical terms, this loss can also guide our choice of the $(\hat{R}, \hat{\gamma})$ pair. The existing IRL works either use ground-truth γ_0 , or a smaller one to reduce the computation burden. In this work, we investigate how to choose the discount factor $\hat{\gamma} \leq \gamma_0$ that minimizes the performance loss defined above. We define the optimal effective discount factor as:

$$\widehat{\gamma}^* = \underset{0 \le \widehat{\gamma} \le \gamma_0}{\arg\min} \left\| V_{R_0, \gamma_0}^{\pi^*_{R_0, \gamma_0}} - V_{R_0, \gamma_0}^{\pi^*_{R, \widehat{\gamma}}} \right\|_{\infty}.$$
(1)

4. Analysis

4.1. Overview

We define the *effective horizon* as the horizon (or discount factor) used in learning the reward function and optimizing policies, which may differ from the ground-truth horizon. We formally analyze how this effective horizon influences the quality of the learned reward function under varying amounts of expert data. Our main result, Theorem 4.1, shows that when expert data is limited, using a discount factor smaller than the ground-truth value enables IRL methods to learn reward functions that induce policies more closely aligned with the expert.

Theorem 4.1. Let (S, A, P) be a controlled Markov process shared by two MDPs: the ground-truth MDP (S, A, P, R_0, γ_0) with reward function $R_0 : S \times A \rightarrow$ $[0, R_{\max}]$ and discount factor $\gamma_0 \in (0, 1)$; and the estimated MDP $(S, A, P, \hat{R}, \hat{\gamma})$ with reward function $\hat{R} : S \times A \rightarrow$ $[0, R_{\max}]$ and discount factor $\hat{\gamma} \in (0, 1)$, estimated from N expert state-action pairs. Let $|\Pi_{\widehat{\gamma}}|$ denote the complexity of the policy class induced by the estimated effective horizon $\widehat{\gamma}$, and suppose the optimal policy $\pi^*_{\widehat{R},\widehat{\gamma}}$ derived from $(\widehat{R},\widehat{\gamma})$ belongs to this class, i.e., $\pi^*_{\widehat{R},\widehat{\gamma}} \in \Pi_{\widehat{\gamma}}$. Then, for the optimal policies $\pi^*_{R_0,\gamma_0}$ and $\pi^*_{\widehat{R},\widehat{\gamma}}$ induced by the ground-truth and estimated parameters, respectively, the difference between their value functions evaluated under R_0 and γ_0 is bounded with probability at least $1 - \delta$ by

$$\left\| V_{R_{0},\gamma_{0}}^{\pi_{R_{0},\gamma_{0}}^{*}} - V_{R_{0},\gamma_{0}}^{\pi_{\widehat{R},\widehat{\gamma}}^{*}} \right\|_{\infty} \leq \frac{2R_{\max}}{(1-\widehat{\gamma})^{2}} \sqrt{\frac{1}{2N} \log \frac{|S||\Pi_{\widehat{\gamma}}|}{2\delta}} + \frac{\gamma_{0} - \widehat{\gamma}}{(1-\gamma_{0})(1-\widehat{\gamma})} R_{max}$$
(2)

Intuitively, Theorem 4.1 bounds the performance disparity between the policy induced by the learned $(\hat{R}, \hat{\gamma})$ and the expert policy as a sum of two terms: the first term bounds the *reward function estimation error* that arises from using a limited dataset and an effective $\hat{\gamma}$ during IRL, while the second term measures the *policy performance gap* when optimized using $\hat{\gamma} < \gamma_0$ using that estimated reward function. When $\hat{\gamma}$ increases, the first term, *i.e.*, the *reward function* estimation error, grows due to overfitting arising from estimating a policy from an increasingly complex class $\Pi_{\widehat{\gamma}}$ using only limited expert data, while the second term, *i.e.*, the policy optimization error, diminishes to encourage fidelity to the ground-truth γ_0 and approaches 0 when $\hat{\gamma} \to \gamma_0$. Consequently, these opposing error terms imply that an intermediate value of $\hat{\gamma}$ yields a better reward function that induces the most expert-like policy.

We build up intermediate theorems and lemmas to formally prove Theorem 4.1 in the remaining subsections. The overall strategy is the following: The overall bound in Theorem 4.1 measures the performance gap between the policy induced by the learned $(\hat{R}, \hat{\gamma})$ pair and the optimal policy. This gap arises from two sources: (i) differences in the reward functions, and (ii) differences in the horizons over which the policies are optimized. Therefore, we decompose the overall bound into two error terms: (i) the *reward function estimation error* resulting from limited expert data during IRL, and (ii) the *policy optimization error* due to using an effective horizon different from the ground truth.

After deriving bounds for both error terms, we combine them to prove Theorem 4.1 in Section 4.6. The second term is straightforward to bound (see Section 4.6), The main challenge lies in bounding the reward function estimation error in the first term (Sections 4.2 to 4.5). We summarize our strategy on bounding the first error term—reward function estimation error—in terms of the effective horizon $\hat{\gamma}$ and the number of expert state-action pairs N. IRL methods learn the reward function by minimizing the discrepancy between the induced policy and the expert data. Therefore, we first bound the reward function estimation error using the *expert policy estimation error*, which measures the gap between the policy induced by the learned reward function and the ground-truth expert policy (see Section 4.3 and 4.4). The expert policy estimation error, in turn, depends on the complexity of the policy class $\Pi_{\hat{\gamma}}$ and the number of expert pairs N used to fit it (see Section 4.5). Furthermore, we prove that the policy class complexity $|\Pi_{\hat{\gamma}}|$ is controlled by the effective horizon $\hat{\gamma}$: as $\hat{\gamma}$ increases, the complexity of the policy class $\Pi_{\hat{\gamma}}$ rises (see Section 4.2). By combining these results, we derive the bound for the reward function estimation error in terms of $\hat{\gamma}$ and N in Section 4.5. This completes the proof.

The remainder of this section is structured as follows. In Section 4.2 we derive how the *effective horizon* controls the complexity measure of the policy class and prove the monotonicity. In Section 4.3, we establish an correspondence between the feasible reward function set and an expert policy. This correspondence is crucial for bounding the error in reward function estimation by the expert policy estimation error. In Section 4.4, we bound the reward function estimation error by the expert policy estimation error. In Section 4.5, we derive a probabilistic bound on the expert policy estimation error, based on the number of samples, policy complexity, and effective horizon. In Section 4.6, we derive the final error bound by combining the results above.

4.2. Complexity of the Policy Class

We propose using the number of potentially optimal policies, given a fixed state space, action space, and transition function, but any unknown reward functions under mild conditions, as the complexity measure of the policy class for different γ s.

Definition 4.2 (Complexity Measure). The complexity measure of the policy class under a specific γ is defined as the number of optimal policies under the fixed state space S, action space A, and transition function P, but with arbitrary reward function $R \in F_R$ that satisfies our assumption described later. Formally, we define the class of optimal policy corresponding to the given γ as:

$$\Pi_{\gamma} = \{\pi : \exists R \in F_R \text{ s.t. } \pi \text{ is optimal in } (S, A, P, R, \gamma)\},\$$

where F_R denotes the set of reward functions satisfying the following condition: for each state $s \in S$, there exists a specific action $a^*(s)$ such that $R(s, a^*(s))$ is strictly greater than R(s, a) for all other actions $a \neq a^*(s)$. In other words, for every $R \in F_R$, the action $a^*(s)$ yields the highest reward at state s, although the exact reward values may vary among different functions in F_R . This form of reward function is not overly restrictive, as it merely excludes reward functions that assign equal rewards to multiple actions for any given state. In systems with uniquely higher step-wise rewards, we can always construct a reward function R from F_R . We assume this specific form of reward function to ensure that discussions about the policy class remain meaningful, as any policy could be considered optimal when arbitrary reward functions are allowed. The complexity of the policy class w.r.t. γ is the number of optimal policies in the corresponding policy class, which is $|\Pi_{\gamma}|$.

Next, we prove that the complexity of the policy class defined in Definition 4.2 increases monotonically as the discount factor $\hat{\gamma}$ increases. We refer the readers to Appendix A for the full proof.

Theorem 4.3. Under a specific MDP $M = (S, A, P, \cdot, \cdot)$ with fixed state space S, action space A, and transition function P, we define the optimal policy class according to Definition 4.2, then we have the following claims:

- 1. $\forall \gamma, \gamma' \in [0, 1)$, if $\gamma < \gamma'$, then $\Pi_{\gamma} \subseteq \Pi_{\gamma'}$.
- 2. When $\gamma = 0$, $|\Pi_0| = 1$.
- 3. If $\gamma \to 1$, $|\Pi_{\gamma}| \ge (|A|-1)^{|S|-1} |S|$ under mild conditions.

Intuitively, claim 1 asserts that as the discount factor γ grows, the number of potentially optimal policies increases monotonically. Together with claim 2 and 3, Theorem 4.3 demonstrates that there is a steep increase in policy complexity associated with the increment of γ . Specifically, when the discount factor is at its lowest ($\gamma = 0$), there is only one optimal policy, as the reward function has a unique maximum state-action pair for each state. However, as γ increases and approaches the largest value ($\gamma \rightarrow 1$), the optimal policy class can eventually encompass nearly all possible policies, with $|\Pi_{\gamma}| = (|A|-1)^{|S|-1}|S|$. In essence, γ effectively controls the complexity of the policy class.

4.3. Feasible Reward Function Set

In this section, we establish a correspondence between the given expert demonstrations and the *feasible reward function set*—the set of all reward functions consistent with the expert data. This correspondence is essential for later bounding the reward function estimation error by the expert policy estimation error.

To create an algorithm-agnostic mapping from the fixed set of expert data and the effective horizon to the learned reward functions, we extend Metelli et al.'s (Metelli et al., 2021) definition of feasible reward function sets—originally formulated for known discount factors—to variable and unknown effective discount factors. We begin by implicitly defining the feasible reward function set based on the foundational IRL formulation (Ng & Russell, 2000), which includes all reward functions whose induced policies match the expert data. From this implicit definition, we derive an explicit characterization of the feasible reward function set as a function of the expert policy.

We start by first implicitly defining the feasible reward set based on the IRL definition (Ng & Russell, 2000), adapting for the flexible discount factor.

Definition 4.4 (IRL Problem, adapted to the setting of varying discount factors). Let $\mathcal{M} = (S, A, P)$ be the MDP without the reward function or discount factor. An IRL problem, denoted as $\Re = (\mathcal{M}, \pi^E)$, consists of the MDP and an expert's policy π^E . A reward $\widehat{R} \in \mathbb{R}^{S \times A}$ is feasible for \Re if there exists a $\widehat{\gamma}$ such that π^E is optimal for the MDP $\mathcal{M} \cup (\widehat{R}, \widehat{\gamma}), i.e., \pi^E \in \Pi^*_{\widehat{R}, \widehat{\gamma}}$. We use \mathcal{R}_{\Re} to denote the set of feasible rewards for \Re .

This IRL formulation implies an implicit correspondence between the expert policy π^E and a *feasible* $(\hat{R}, \hat{\gamma})$ pair through their advantage function $A_{\hat{R},\hat{\gamma}}^{\pi^E}(s,a) = Q_{\hat{R},\hat{\gamma}}^{\pi^E}(s,a) - V_{\hat{R},\hat{\gamma}}^{\pi^E}(s)$. Specifically, the optimal policy induced by $(\hat{R},\hat{\gamma})$ matches π^E when the following two conditions on the advantage function are met:

1. if
$$\pi^{E}(a|s) > 0$$
, then $A_{\widehat{R},\widehat{\gamma}}^{\pi^{E}}(s,a) = 0$,
2. if $\pi^{E}(a|s) = 0$, then $A_{\widehat{R},\widehat{\gamma}}^{\pi^{E}}(s,a) \leq 0$.

The first condition ensures the expert's chosen actions have zero advantage, while the second guarantees unchosen actions have non-positive advantages.

Next, we establish an explicit correspondence between the estimated $(\hat{R}, \hat{\gamma})$ and their compatible expert policy π^E by enforcing these two conditions on the advantage function $A_{\hat{R},\hat{\gamma}}^{\pi^E}$. To do this, we introduce two operators for any given policy π :

- 1. the *expert-filter*: $(B^{\pi}A)(s,a) = A(s,a)\mathbb{1}\{\pi(a|s) > 0\}$, that retains the advantage A(s,a) values for actions taken by the expert policy $\pi^{E}(a|s)$,
- 2. the *expert-filter-complement*: $(\bar{B}^{\pi}A)(s,a) = A(s,a)\mathbb{1}\{\pi(a|s) = 0\}$, that preserves the advantage values for actions **not** taken by the expert.

Using the Bellman equation, we express the advantage functions in terms of the estimated $(\hat{R}, \hat{\gamma})$. We then apply the two filters to the advantage function $A_{\hat{R},\hat{\gamma}}^{\pi^E}$ to enforce the optimality conditions, simplifying the expression in the process. This allows us to derive the explicit expression for the feasible reward set presented in Lemma 4.5. The detailed derivation is provided in Appendix B.1. **Lemma 4.5** (Feasible Reward Set, extended from Metelli et al. (2021)). Let $\Re = (\mathcal{M}, \pi^E)$ be an IRL problem. Let $\widehat{R} \in \mathbb{R}^{S \times A}$ and $0 < \widehat{\gamma} < 1$, then \widehat{R} is a feasible reward, i.e., $\widehat{R} \in \mathcal{R}_{\Re}$ if and only if there exists $\zeta \in \mathbb{R}_{\geq 0}^{S \times A}$ and $V \in \mathbb{R}^S$ such that:

$$\widehat{R} = -\overline{B}^{\pi^{E}}\zeta + (E - \widehat{\gamma}P)V, \qquad (3)$$

whereas $E : \mathbb{R}^{|S|} \to \mathbb{R}^{|S| \times |A|}$ is an operator that marginalizes the action for each state on a function $f(\cdot)$ such that (Ef)(s, a) = f(s).

The feasible reward function in Lemma 4.5 comprises two terms: the first term depends on the expert policy π^E and the second term depends on the underlying MDP's transition function. The first term, $-\bar{B}^{\pi^E}\zeta$, is derived using the *expert-filter-complement* on a non-negative function ζ . This ensures that actions taken by the expert (i.e., $\pi^E(a|s) > 0$) are assigned a value of zero, while actions not taken (i.e., $\pi^E(a|s) = 0$) have non-positive values. The second term represents the policy's temporal effect that relies on the MDP's transition function. This can be viewed as reward shaping through the value function, which preserves the expert policy's optimality.

We have thus established an explicit correspondence between the feasible $(\hat{R}, \hat{\gamma})$ and the expert policy π^E in Lemma 4.5. This explicit expression will later be used to bound the difference in reward functions by the expert policy estimation error.

4.4. Reward Function Estimation Error from Expert Policy Estimation Error

In this section, we establish a bound on the reward function estimation error in terms of the *expert policy estimation error*—the discrepancy between the estimated expert policy (from limited data) and the true expert policy. Since IRL methods learn the reward function by matching the induced policy to expert data, establishing this bound is crucial, as it allows us to relate the reward estimation error to the amount of expert data and the induced policy complexity in later analysis.

Let's consider two IRL problems, $\Re = (\mathcal{M}, \pi^E)$ and $\hat{\Re} = (\mathcal{M}, \hat{\pi}^E)$, which differ only in expert policies: \Re utilizes the ground-truth expert policy, while $\hat{\Re}$ employs an estimated policy from samples. Since an IRL algorithm aligns its induced policy with the estimated expert policy, its feasible sets will be equivalent to that of the estimated expert policy. Intuitively, inaccuracies in estimating the expert policy π^E lead to errors in estimating the feasible sets \mathcal{R}_{\Re} . Our goal is to obtain a reward function \hat{R} with a feasible set "close" to the ground-truth R_0 's feasible set. Specifically, "closeness" is determined by the distance between the nearest reward functions in each set. The estimated \mathcal{R}_{\Re} is considered close

to the exact \mathcal{R}_{\Re} if, for every reward $R_0 \in \mathcal{R}_{\Re}$, there exists an estimated reward $\widehat{R} \in \mathcal{R}_{\hat{\Re}}$ with a small $|R_0 - \widehat{R}|$ value.

Given the form of the feasible reward functions corresponding to an expert policy as derived in Lemma 4.5, we express the estimation error $|R_0 - \hat{R}|$ as a function of the groundtruth expert policy π^E and the estimated expert policy $\hat{\pi}^E$ from limited data. The bound on the reward function estimation error is shown below.

Theorem 4.6 (Extension of Theorem 3.1 in Metelli et al. (2021)). Let $\Re = (\mathcal{M}, \pi^E)$ and $\hat{\Re} = (\mathcal{M}, \hat{\pi}^E)$ be two IRL problems. Then for any $R_0 \in \mathcal{R}_{\Re}$ such that $R_0 = -\bar{B}^{\pi^E}\zeta + (E - \gamma_0 P)V$ and $||R_0||_{\infty} \leq R_{\max}$ there exist $\hat{R} \in \mathcal{R}_{\hat{\Re}}$ and $0 < \hat{\gamma} < 1$, such that element-wise it holds that:

$$\left| R_0 - \widehat{R} \right| \le \bar{B}^{\pi^E} B^{\hat{\pi}^E} \zeta. \tag{4}$$

Furthermore, $\|\zeta\|_{\infty} \leq \frac{R_{\max}}{1-\gamma_0}$.

The theorem above bounds the reward function estimation error by the discrepancy between the true expert policy π^E and the estimated expert policy $\hat{\pi}^E$ derived from limited data. Intuitively, it states that there exists a reward function \hat{R} in the estimated feasible set $\mathcal{R}_{\hat{\mathfrak{R}}}$ whose estimation error is controlled by the expert policy estimation error, under the corresponding $\hat{\gamma}$. Specifically, the error term on the right-hand side is non-zero only for state-action pairs where $\pi^E(a \mid s) = 0$ but $\hat{\pi}^E(a \mid s) > 0$. This means the reward estimation error is zero for state-action pairs observed in the expert demonstrations, while errors arise where the estimated expert policy incorrectly assigns positive probability to actions the expert did not take. We refer readers to Appendix B.2 for the detailed proof.

4.5. Expert Policy Estimation Error from Limited Data

In this section, we derive a bound on the expert policy estimation error in terms of the amount of expert data N, the effective horizon $\hat{\gamma}$, and the induced policy complexity $|\Pi_{\hat{\gamma}}|$. We measure this error using the term $\bar{B}^{\pi^E}B^{\hat{\pi}^E}\zeta$, which quantifies how much the estimated policy $\hat{\pi}^E$ deviates from the true expert policy π^E .

Our derivation proceeds in three steps. First, we derive the expected value of the expert policy estimation error $\mathbb{E}[\bar{B}^{\pi^E}B^{\hat{\pi}^E}\zeta]$ under our estimation strategy. Then, we apply McDiarmid's inequality to obtain a probabilistic bound on how likely this error deviates from its expected value. Finally, we derive a uniform bound on the estimation error by substituting appropriate threshold values and simplifying the expression.

Expected Value of the Expert Policy Estimation Error To estimate the expert policy $\hat{\pi}^E \in \{0,1\}^{|S| \times |A|}$ from N independent state-action samples, we aggregate the demonstrated pairs by summing their counts into $\widehat{\Pi}_N^E$. For each state *s*, we estimate $\widehat{\pi}^E(s, a)$ by selecting the action *a* with the highest count:

$$\hat{a}_s = \arg\max_{a} \widehat{\Pi}_N^E(s, a).$$
(5)

If the count of \hat{a}_s is uniquely the highest, we set $\hat{\pi}^E(s, a)$ to 1 for \hat{a}_s and 0 for other actions. If no action has a unique maximum count, we set all entries to 0, indicating uncertainty. As the number of samples N increases, the estimated policy is more likely to reflect the expert's true decisions. For further details, see Appendix B.4.1.

Next, based on the estimation strategy described above, we first compute the expected value for the approximated expert policy $\mathbb{E}[\hat{\pi}^{E}]$ estimated from N expert samples, and in turn that of the estimation error $\mathbb{E}[\bar{B}^{\pi^{E}}B^{\hat{\pi}^{E}}\zeta]$.

We first compute the expected value of $\hat{\pi}^E(s, a)$ given N expert state-action pairs. The expected number state s is sampled across the N samples is $N \cdot \frac{1}{|S|}$. Moreover, since the expert policy is deterministic, the probability that a_s^* is the unique maximum at state s is the same as the probability that state s is sampled at least once. Therefore, the expectation of the estimated policy $\mathbb{E}[\hat{\pi}^E]$ for each state s is given by:

$$\mathbb{E}[\hat{\pi}^{E}(s,a)] = \begin{cases} 1 - \left(1 - \frac{1}{|S|}\right)^{N} & \text{for } a = a_{s}^{*}, \\ 0 & \text{otherwise.} \end{cases}$$

Next, we compute the expected value of the expert policy estimation error $\mathbb{E}[\bar{B}^{\pi^E}B^{\hat{\pi}^E}\zeta]$. Recall the operator \bar{B}^{π^E} retains the values of $\zeta(s,a)$ for actions a that are not part of the expert policy. The true expert policy π^E deterministically selects action a_s^* for each state s, so the expert-filter complement will only retain the values for actions $a \neq a_s^*$. Thus, after applying the expert-filter complement, the expected value becomes:

$$\mathbb{E}[\bar{B}^{\pi^E}B^{\hat{\pi}^E}\zeta(s,a)] = \mathbb{E}[B^{\hat{\pi}^E}\zeta(s,a)] \cdot \mathbb{E}[\mathbb{1}\{\pi^E(a|s)=0\}]$$

Moreover, the operator $B^{\hat{\pi}^E}$ preserves the function values for actions that are chosen by the estimated policy $\hat{\pi}^E(s, a)$. Thus, the expected value of $B^{\hat{\pi}^E}\zeta(s, a)$ is:

$$\mathbb{E}[B^{\hat{\pi}^E}\zeta(s,a)] = \zeta(s,a) \cdot \mathbb{E}[\mathbb{1}\{\hat{\pi}^E(a|s) > 0\}] \quad (6)$$

Simplifying the expressions above, we have $\mathbb{E}[\bar{B}^{\pi^E}B^{\hat{\pi}^E}\zeta(s,a)] = 0$, since $\pi^E(a|s) = 0$ for all $a \neq a_s^*$. Intuitively, an expected expert policy estimation error of 0 implies that, on average, the estimated policy $\hat{\pi}^E$ correctly matches the expert policy π^E . This means it does not assign non-zero probabilities to actions the expert would not take, indicating no systematic bias in the estimation process.

Applying McDiarmid's Inequality Next, we obtain a probabilistic bound on how likely the expert policy estimation error deviates from its expected value $\mathbb{E}[\bar{B}^{\pi^E}B^{\hat{\pi}^E}\zeta(s,a)]$. We do so by applying McDiarmid's Inequality to the term $\bar{B}^{\pi^E}B^{\hat{\pi}^E}\zeta$. Since ζ is bounded by $\frac{R_{\max}}{1-\hat{\gamma}}$, altering any single state-action sample affects at most one row of the estimated policy $\hat{\pi}^E(s,:)$, which leads to a bounded change in the value of $\bar{B}^{\pi^E}B^{\hat{\pi}^E}\zeta$ by at most $\frac{R_{\max}}{1-\hat{\gamma}}$. This satisfies the bounded difference condition necessary for McDiarmid's Inequality.

By applying McDiarmid's Inequality, we derive that the probability of the expert policy estimation error exceeding a threshold t is bounded by:

$$\Pr\left(\bar{B}^{\pi^{E}}B^{\hat{\pi}^{E}}\zeta(s,a) \ge t\right) \le \exp\left(\frac{-2Nt^{2}(1-\hat{\gamma})^{2}}{R_{\max}^{2}}\right).$$
(7)

This bound quantifies how likely this estimation error exceeds a threshold t. As the number of expert-demonstrated state-action pairs N increases, the probability of significant deviations from the true policy decreases exponentially.

Uniform Bound on the Expect Policy Estimation Error To obtain a uniform bound on the expert policy estimation error across all (s, π) pairs, we apply the union bound and set the right-hand side of the bound to $\frac{\delta}{|S||\Pi_{\hat{\gamma}}|}$. Solving for t, we derive the threshold that bounds the error with probability at least $1 - \delta$:

$$t = \frac{R_{\max}}{1 - \widehat{\gamma}} \sqrt{\frac{1}{2N} \ln\left(\frac{|S||\Pi_{\widehat{\gamma}}|}{\delta}\right)}.$$
 (8)

Since the reward function estimation error is bounded by the expert policy estimation error, $|R_0 - \hat{R}| \leq \bar{B}^{\pi^E} B^{\hat{\pi}^E} \zeta$, the error bound threshold t applies to the reward function estimation error with probability at least $1 - \delta$. As the amount of expert data N increases, the estimation error decreases, reflecting improved accuracy with more data. Moreover, the effective horizon $\hat{\gamma}$ plays an important role: smaller values of $\hat{\gamma}$ reduce both $\frac{1}{1-\hat{\gamma}}$ and $|\Pi_{\hat{\gamma}}|$, tightening the bound and better controlling the estimation error.

4.6. Error Decomposition and Deriving the Overall Bound

In this section, we use the reward function estimation error bound from earlier to establish the final bound on $\left\| V_{R_0,\gamma_0}^{\pi^*_{R_0,\gamma_0}} - V_{R_0,\gamma_0}^{\pi^*_{\widehat{R},\widehat{\gamma}}} \right\|_{\infty}$ presented in Theorem 4.1. Recall that $\pi^*_{R_0,\gamma_0} \in \Pi_{\gamma_0}$ is the optimal policy derived from the ground-truth parameters (R_0,γ_0) , while $\pi^*_{\widehat{R},\widehat{\gamma}} \in \Pi_{\widehat{\gamma}}$ is the optimal policy derived from the learned parameters $(\widehat{R},\widehat{\gamma})$.

We aim to bound the difference between their value functions evaluated under the ground-truth R_0 and γ_0 .

To simplify the analysis, we decompose the overall error into two manageable terms: the first accounts for the difference in value functions caused by the reward function estimation error, $||R_0 - \hat{R}||_{\infty}$; the second accounts for the value difference due to optimizing the policy under different horizons (γ_0 versus $\hat{\gamma}$).

Theorem 4.7 (Value Function Difference Bound). Let $\mathcal{M} = (S, A, P)$ be a partial Markov Decision Process shared by two MDPs. Let $R_0 : S \times A \rightarrow [0, R_{\max}]$ be the ground-truth reward function, with $\gamma_0 \in (0, 1)$ as the ground-truth discount factor. Let $\widehat{R} : S \times A \rightarrow [0, R_{\max}]$ and $\widehat{\gamma} \in (0, 1)$ be the estimated reward function and discount factor obtained from data. Consider the optimal policies π_{R_0,γ_0}^* and $\pi_{\widehat{R},\widehat{\gamma}}^*$ induced by (R_0,γ_0) and $(\widehat{R},\widehat{\gamma})$, respectively. Then, the difference between their value functions, evaluated using the ground-truth reward R_0 and discount factor γ_0 , is bounded above by

$$\left\| V_{R_0,\gamma_0}^{\pi_{R_0,\gamma_0}^*} - V_{R_0,\gamma_0}^{\pi_{\widehat{R},\widehat{\gamma}}^*} \right\|_{\infty} \leq \frac{2}{1-\widehat{\gamma}} \left\| R_0 - \widehat{R} \right\|_{\infty} + \frac{|\gamma_0 - \widehat{\gamma}|}{(1-\gamma_0)(1-\widehat{\gamma})} R_{\max}.$$
(9)

Proof of Theorem 4.7 is in Appendix B.3.

Recall that we have established a bound on the reward estimation error $\left\|R_0 - \hat{R}\right\|_{\infty}$ in terms of the number of expert data N, the effective horizon $\hat{\gamma}$, and the policy complexity $|\Pi_{\widehat{\gamma}}|$ in Section 4.5. Specifically, with probability at least $1 - \delta$, the reward function estimation error is bounded by $\frac{R_{\max}}{1-\widehat{\gamma}}\sqrt{\frac{1}{2N}\ln\left(\frac{|S||\Pi_{\widehat{\gamma}}|}{\delta}\right)}$. Substituting this bound into the first error term on the right-hand side of Theorem 4.7, we obtain:

$$\left\| V_{R_0,\gamma_0}^{\pi^*_{R_0,\gamma_0}} - V_{R_0,\gamma_0}^{\pi^*_{\widehat{R},\widehat{\gamma}}} \right\|_{\infty} \leq \frac{R_{\max}}{1 - \widehat{\gamma}} \sqrt{\frac{1}{2N} \ln\left(\frac{|S||\Pi_{\widehat{\gamma}}|}{\delta}\right)} + \frac{|\gamma_0 - \widehat{\gamma}|}{(1 - \gamma_0)(1 - \widehat{\gamma})} R_{\max}.$$

$$(10)$$

This completes the proof of Theorem 4.1. The overall error consists of two terms. The first term arises from the reward function estimation error, which increases with larger $\hat{\gamma}$. A larger $\hat{\gamma}$ leads to a more complex induced policy class, making overfitting more likely when expert data is limited. Consequently, the reward function estimation error increases as the expert policy estimation error grows significantly. The second term represents the performance loss from using

a smaller discount factor ($\hat{\gamma} < \gamma_0$), and decreases as $\hat{\gamma}$ approaches γ_0 . These opposing dependencies on $\hat{\gamma}$ suggest that there exists an intermediate value $0 < \hat{\gamma} < \gamma_0$ that minimizes the overall loss. We will empirically demonstrate this in the following section.

5. Experiments

In this section, we empirically examine Theorem 4.1 by exploring how discount factors influence IRL performance with varying amounts of expert data. We adapt Linear Programming IRL (LP-IRL) (Ng & Russell, 2000) and Maximum Entropy IRL (MaxEnt-IRL) (Ziebart et al., 2008) to accommodate different discount factor settings and expert data sizes (implementation details are provided in Appendices D and E).

To jointly learn $(\hat{R}, \hat{\gamma})$, we incorporate cross-validation into our modified IRL methods to optimize the discount factor (details in Section 5.2). By evaluating policies across a range of $\hat{\gamma}$ s, cross-validation allows us to directly observe how varying $\hat{\gamma}$ affects IRL performance with different amounts of expert data. While specialized IRL methods exist for optimizing $\hat{\gamma}$ (Giwa & Lee, 2021; Ghosal et al., 2023), we choose cross-validation because it not only identifies the optimal discount factor, but also reveals performance trends across various settings, which is crucial for validating the nuanced implications of our theoretical findings. Specifically, we answer the following questions through cross-validation:

- Q.1 Can a lower $\hat{\gamma} < \gamma_0$ improve IRL policy performance?
- Q.2 How does $\hat{\gamma}^*$ change with increasing expert data N?
- Q.3 Is the cross-validation extension effective in finding $\hat{\gamma}^*$?

We evaluate the performance of LP-IRL and MaxEnt-IRL on four Gridworld and Objectworld tasks with varying reward complexity. We use the number of incorrectly induced actions as a proxy for value estimation error in the overall error bound. This measure counts the number of states where the induced policy $\hat{\pi}^E$ differs from π^E in action selection. For Q.1, we measure the number of incorrectly induced actions under varying discounted factors and different amount of expert data. Our findings show that the optimal $\hat{\gamma}$ s across all expert amount are smaller than γ_0 for both algorithms. For Q.2, we plot how the optimal discount factors change as the number of expert data increases. The consistent U-shaped curves observed in all cases align with the anticipated overfitting effect implied by the second error term in Equation 2. For Q.3, we compare the performance of policies selected via cross-validation with the best policy learnable from the available expert data. Our results indicate that the discrepancy in performance is negligible for all

tasks, demonstrating the effectiveness of cross-validation in selecting $\hat{\gamma}^*$.

5.1. Task Setup

We design four tasks of varying complexity in reward functions: Gridworld-simple, Gridworld-hard, Objectworld-linear, and Object-world-nonlinear, adapted from (Levine et al., 2011) and (Ng & Russell, 2000). We illustrate each task instance in the first three columns of Tables 1 and more details on the task specification are in Appendix C. The ground-truth discount factor is $\gamma_0 = 0.99$.

The Gridworld tasks provide sparse rewards only at randomly sampled goals: Gridworld-simple has fewer goals (4) and a smaller state space (10×10 states), while Gridworldhard has more goals (6) and a larger state space (15×15 states). On the other hand, the Objectworld tasks have denser ground-truth rewards that are functions of nearby object features. The reward function for Objectworld-linear is linear with respect to the features of nearby objects, while that of Objectworld-nonlinear is non-linear. Intuitively, learning a complex reward function may be more susceptible to overfitting, especially when expert data is sparse compared to the state space.

We vary the amount of expert data by adjusting the number of state-action pairs included, measuring this as a percentage relative to the state space size. Since states may be sampled multiple times in the trajectories, the total number of state-action pairs N can exceed the number of states |S|. Given expert trajectories $D = \{\tau_0, \tau_1, \ldots\}$, we define the percentage of expert data as $K\% = \frac{N}{|S|} \times 100\%$. We evaluate the performance of the induced policy by counting the number of states where it selects a different action from the expert policy—referred to as the *state error count*.

5.2. Cross Validation Extension

We use cross-validation to determine the optimal $\widehat{\gamma}^*$ from the expert data D containing N state-action pairs. We split D into non-overlapping training (80%) and validation (20%) sets. We uniformly sample M = 20 discount factors from the interval $(0, \gamma_0)$. For each sampled γ , we learn R_{γ} using the training set and evaluate the induced policy on the validation set by counting the number of states where it selects a different action from the expert policy-the state error *count*. The γ that minimizes this error count is selected as the optimal $\hat{\gamma}$. For all tasks, we randomly sample 10 environments per task and report the mean and standard deviation of errors. To assess the effectiveness of cross-validation, we employ an oracle representing the best policy learnable from all available expert data. This oracle, considered "cheating," uses both the training and validation sets for training and validates on the entire state space (both observed and unob-



Figure 1: Summary of LP-IRL with varying discount factors across four tasks. The *error counts* measure the number of states for which a policy's action selection deviates from the expert's actions. Each task displays the ground-truth value function (column 1), reward function (column 2), expert policy (column 3), error count curves for different amount of expert data in a **single instance** (columns 4-8), and the error count curve summary for a **batch** of 10 MDPs across varying amount of expert data (column 9). In all four tasks, $\gamma_0 = 0.99$. The optimal discount factor $\hat{\gamma}^* < \gamma_0$ for varying amount of expert data. MaxEnt-IRL has similar curves in Figure 4.



Figure 2: Optimal $\hat{\gamma}^*$ for LP-IRL at varying amount of expert data. For each task, we select $\hat{\gamma}^*$ for all 10 sampled environments through cross-validation (see Section 5.2). The orange curves illustrate how the optimal discount factor $\hat{\gamma}^*$ changes with the amount of expert data, while the green curves show the corresponding error counts. The ground-truth $\gamma_0 = 0.99$ is depicted in grey, with its error counts displayed in blue. As the amount of expert data increases, $\hat{\gamma}^*$ initially decreases sharply and then gradually increases, indicating that overfitting is prominent when expert data is scarce.

served states). We use this oracle only to compare whether the $\hat{\gamma}^*$ chosen by cross-validation corresponds to that of the oracle policy, not for selecting the optimal $\hat{\gamma}$.

5.3. Results

We assess the impact of the *effective horizon* on IRL by evaluating LP-IRL and MaxEnt-IRL across four tasks. For

simplicity, we treat LP-IRL's discount factor γ and MaxEnt-IRL's horizon T interchangeably, with findings for γ also applying to T unless specified otherwise. Policy performance results are presented in Tables 1 (LP-IRL) and 4 (MaxEnt-IRL). We report *state error counts*, measuring discrepancies between induced and expert policies by counting the states where their action selections differ.

Q.1 Optimal $\hat{\gamma}^*$ is Lower than Ground-Truth.

As shown in Figures 1 and 4, the optimal discount factor $\hat{\gamma}^* < \gamma_0$ for all four tasks and across various amounts of expert data in both LP-IRL and MaxEnt-IRL. With limited expert data, the error count curves are generally U-shaped: discrepancies with the expert policy decrease as $\hat{\gamma}$ increases to a "sweet spot" and then rise sharply. This pattern confirms our error bounds in Theorem 4.1: for small $\hat{\gamma}$, the overfitting-related error (second term in Equation 2) is less significant, and increasing $\hat{\gamma}$ allows temporal extrapolation, reducing the overall error. However, beyond the optimal $\hat{\gamma}$, overfitting becomes more pronounced, and the overall error increases as the first error term outweighs the benefits.

With abundant expert data, error counts remain low (in LP-IRL) or initially decrease (in MaxEnt-IRL) for small $\hat{\gamma}$ and then increase as $\hat{\gamma}$ grows, indicating that $\hat{\gamma}^* < \gamma_0$ yields the most expert-like policy, confirming our theoretical results. Interestingly, in LP-IRL, the error counts do not initially



Figure 3: The cross-validation results for LP-IRL on four tasks are shown. The *x*-axis represents the amount of expert data; the *y*-axis shows policy error count differences. We compare discount factors $\hat{\gamma}_{cv}^*$ (learned from cross-validation) and $\hat{\gamma}_{oracle}^*$ (chosen by the oracle). Orange dots depict error differences between policies induced by $\hat{\gamma}_{cv}^*$ and $\hat{\gamma}_{oracle}^*$; blue dots show differences between policies induced by $\hat{\gamma}_{cv}^*$ and the ground-truth γ_0 . The orange curves near zero indicate that cross-validation effectively selects $\hat{\gamma}^*$, while the positive blue curves show that cross-validation consistently yields better policies than using γ_0 .

drop as $\hat{\gamma}$ increases. This is because, with dense expert data, LP-IRL accurately matches step-wise behaviors, making performance gains from temporal reasoning negligible. This observation supports Spencer et al. (2021)'s insight that naive behavioral cloning can outperform IRL algorithms when ample expert data is available. In contrast, MaxEnt-IRL does not exhibit low error counts for small $\hat{\gamma}$ because its reward function is parameterized linearly in state features, limiting its ability to precisely replicate actions even with abundant data.

Q.2 Optimal $\hat{\gamma}^*$ s Vary with the Amount of Expert Data.

Figures 2 and 5 show how the optimal discount factor $\hat{\gamma}^*$ varies with increasing amounts of expert data for LP-IRL and MaxEnt-IRL, respectively. When data is scarce, $\hat{\gamma}^*$ is high because the benefits of temporal reasoning outweigh overfitting concerns—the estimation error of the expert policy remains high regardless of overfitting. A large $\hat{\gamma}^*$ enables extrapolation of actions to nearby unobserved states, reducing the first error term in Theorem 4.1 and improving overall error reduction.

With a moderate amount of data, $\hat{\gamma}^*$ decreases: although expert policy estimation improves with more data, it is still limited. The overall error bound favors a smaller $\hat{\gamma}^*$ to mitigate overfitting, reducing the second error term even at the cost of some temporal reasoning benefits. As data increases further, $\hat{\gamma}^*$ rises again since overfitting becomes less significant, and larger values enhance temporal reasoning.

Overall, as expert data increases, error counts for $\hat{\gamma}^*$ strictly decrease and remain below those for γ_0 , indicating that $\hat{\gamma}^*$ enables IRL to learn more effectively from additional expert data.

Q.3 Cross-Validation Effectively Selects Optimal $\hat{\gamma}^*$ s.

Figures 3 and 6 summarize the cross-validation results for LP-IRL and MaxEnt-IRL, respectively. The performance discrepancy between policies induced by $\hat{\gamma}_{cv}^*$ and the oracle $\hat{\gamma}_{oracle}^*$ (orange curves) is consistently near zero across all four tasks, indicating that cross-validation effectively selects $\hat{\gamma}^*$ similar to the oracle. Moreover, the blue curves represent the error differences between policies induced by the ground-truth γ_0 and $\hat{\gamma}_{cv}^*$, which are significantly higher than zero. This suggests that $\hat{\gamma}_{cv}^*$ yields better-performing policies than using the ground-truth γ_0 , confirming our theoretical findings in Theorem 4.1.

6. Conclusion

In this paper, we present a theoretical analysis on IRL that unveils the potential of a reduced horizon in inducing a more expert-like policy, particularly in data-scarce situations. Our findings reveals an important insight on role of the horizon in IRL: it controls the complexity of the induced policy class, therefore reduces overfitting to the limited expert data. We, therefore, propose a more natural IRL function class that jointly learns reward-horizon pairs and empirically substantiate our analysis using a cross-validation extension for the existing IRL algorithms. As overfitting remains a challenge for IRL, especially with scarce expert data, we believe our findings offer valuable insights for the IRL community on better IRL formulations.

References

- Abbeel, P. and Ng, A. Apprenticeship Learning via Inverse Reinforcement Learning. In *Proc. Int. Conf. on Machine Learning*, 2004.
- Amit, R., Meir, R., and Ciosek, K. Discount factor as a regularizer in reinforcement learning. In Proc. Int. Conf. on Machine Learning, 2020.
- Boularias, A., Kober, J., and Peters, J. Relative Entropy Inverse Reinforcement Learning. In *Proc. Int. Conf. on Artificial Intelligence & Statistics*, 2011.
- Finn, C., Christiano, P., Abbeel, P., and Levine, S. A Connection Between Generative Adversarial Networks, Inverse Reinforcement Learning, and Energy-Based Models. In Advances in Neural Information Processing Systems, 2016a.
- Finn, C., Levine, S., and Abbeel, P. Guided Cost Learning: Deep Inverse Optimal Control via Policy Optimization. In Proc. Int. Conf. on Machine Learning, 2016b.
- Fu, J., Luo, K., and Levine, S. Learning Robust Rewards with Adversarial Inverse Reinforcement Learning. In *Proc. Int. Conf. on Learning Representations*, 2018.

- Ghasemipour, S. K. S., Zemel, R., and Gu, S. A Divergence Minimization Perspective on Imitation Learning Methods. In Proc. Conference on Robot Learning, 2019.
- Ghosal, G. R., Zurek, M., Brown, D. S., and Dragan, A. D. The effect of modeling human rationality level on learning rewards from multiple feedback types. In *Proceedings* of the AAAI Conference on Artificial Intelligence, 2023.
- Giwa, B. H. and Lee, C.-G. Estimation of discount factor in a model-based inverse reinforcement learning framework. In *Bridging the Gap Between AI Planning and Reinforcement Learning Workshop at ICAPS*, 2021.
- Guo, X., Hu, A., and Zhang, J. Theoretical guarantees of fictitious discount algorithms for episodic reinforcement learning and global convergence of policy gradient methods. In *Proc. AAAI Conf. on Artificial Intelligence*, 2022.
- Ho, J. and Ermon, S. Generative Adversarial Imitation Learning. In Advances in Neural Information Processing Systems, 2016.
- Hoshino, H., Ota, K., Kanezaki, A., and Yokota, R. OPIRL: Sample Efficient Off-Policy Inverse Reinforcement Learning via Distribution Matching. In *Proc. IEEE Int. Conf. on Robotics & Automation*, 2022.
- Hu, H., Yang, Y., Zhao, Q., and Zhang, C. On the role of discount factor in offline reinforcement learning. In *Proc. Int. Conf. on Machine Learning*, 2022.
- Jiang, N., Kulesza, A., Singh, S., and Lewis, R. The Dependence of Effective Planning Horizon on Model Accuracy. In Proc. Int. Conf. on Autonomous Agents & Multiagent Systems, 2015.
- Ke, L., Choudhury, S., Barnes, M., Sun, W., Lee, G., and Srinivasa, S. Imitation Learning as f-Divergence Minimization. In Algorithmic Foundations of Robotics XIV— Proc. Int. Workshop on the Algorithmic Foundations of Robotics (WAFR), 2020.
- Laidlaw, C., Russell, S. J., and Dragan, A. Bridging rl theory and practice with the effective horizon. *Advances* in Neural Information Processing Systems, 36:58953– 59007, 2023.
- Lee, K., Isele, D., Theodorou, E. A., and Bae, S. Spatiotemporal Costmap Inference for MPC via Deep Inverse Reinforcement Learning. In *IEEE Robotics & Automation Letters*, 2022.
- Levine, S., Popovic, Z., and Koltun, V. Nonlinear Inverse Reinforcement Learning with Gaussian Processes. In Advances in Neural Information Processing Systems, 2011.

- MacGlashan, J. and Littman, M. L. Between Imitation and Intention Learning. In Proc. Int. Jnt. Conf. on Artificial Intelligence, 2015.
- Mattingley, J., Wang, Y., and Boyd, S. Receding horizon control. *IEEE Control Systems Magazine*, 31(3):52–65, 2011.
- Metelli, A. M., Ramponi, G., Concetti, A., and Restelli, M. Provably Efficient Learning of Transferable Rewards. In *Proc. Int. Conf. on Machine Learning*, 2021.
- Ng, A. Y. and Russell, S. J. Algorithms for Inverse Reinforcement Learning. In *Proc. Int. Conf. on Machine Learning*, 2000.
- Ng, A. Y., Harada, D., and Russell, S. Policy Invariance Under Reward Transformations: Theory and Application to Reward Shaping. In *Proc. Int. Conf. on Machine Learning*, 1999.
- Ni, T., Sikchi, H., Wang, Y., Gupta, T., Lee, L., and Eysenbach, B. f-IRL: Inverse Reinforcement Learning via State Marginal Matching. In *Proc. Conference on Robot Learning*, 2020.
- Osa, T., Pajarinen, J., Neumann, G., Bagnell, J. A., Abbeel, P., Peters, J., et al. An Algorithmic Perspective on Imitation Learning. *Foundations and Trends*® *in Robotics*, 7 (1-2):1–179, 2018.
- Petrik, M. and Scherrer, B. Biasing approximate dynamic programming with a lower discount factor. In *Advances in neural information processing systems*, 2008.
- Pirotta, M. and Restelli, M. Inverse Reinforcement Learning through Policy Gradient Minimization. In Proc. AAAI Conf. on Artificial Intelligence, 2016.
- Ramachandran, D. and Amir, E. Bayesian Inverse Reinforcement Learning. In Proc. Int. Int. Conf. on Artificial Intelligence, 2007.
- Ramponi, G., Drappo, G., and Restelli, M. Inverse Reinforcement Learning from a Gradient-Based Learner. *Advances in Neural Information Processing Systems*, 33: 2458–2468, 2020.
- Ratliff, N. D., Bagnell, J. A., and Zinkevich, M. A. Maximum Margin Planning. In *Proc. Int. Conf. on Machine Learning*, 2006.
- Spencer, J., Choudhury, S., Venkatraman, A., Ziebart, B. D., and Bagnell, J. A. Feedback in Imitation Learning: The Three Regimes of Covariate Shift. *ArXiv*, abs/2102.02872, 2021.

- Wulfmeier, M., Ondruska, P., and Posner, I. Maximum Entropy Deep Inverse Reinforcement Learning. In Advances in Neural Information Processing Systems, 2016.
- Xu, Y., Gao, W., and Hsu, D. Receding Horizon Inverse Reinforcement Learning. In *Advances in Neural Information Processing Systems*, 2022.
- Ziebart, B. D., Maas, A., Bagnell, J. A., and Dey, A. K. Maximum Entropy Inverse Reinforcement Learning. In *Proc. AAAI Conf. on Artificial Intelligence*, 2008.

A. Complexity Measure of Policy

Definition A.1 (Reward and Policy Equivalence). For an MDP $M = (S, A, P, ., \gamma)$, we define two bounded reward functions R and R' to be equivalent, i.e. $R \equiv R'$, if and only if they induce the same set of optimal policies.

Lemma A.2 (Potential-based Reward Shaping (Ng et al., 1999)). For an MDP $M = (S, A, P, .., \gamma)$, two bounded reward functions R and R' are equivalent, i.e. $R \equiv R'$, if and only if there exists a bounded potential function $\phi : S \to \mathbb{R}$ such that for all $s, s' \in S$ and $a \in A$:

$$R'(s, a, s') = R(s, a, s') + \gamma \phi(s') - \phi(s)$$
(11)

Remark A.3. We extend the potential-based reward shaping presented in lemma A.2, which utilizes a reward function parameterized by R(s, a, s'), to one that does not rely on the subsequent state, namely, R(s, a). For an MDP $M = (S, A, P, ., \gamma)$, two bounded reward functions R(s, a) and R'(s, a) are equivalent, if and only if there exists a bounded potential function $\phi : S \to \mathbb{R}$ such that for all $s, s' \in S$ and $a \in A$:

$$R'(s,a) = R(s,a) + \gamma \sum_{s'} P(s'|a,s)\phi(s') - \phi(s)$$
(12)

Proof. Consider the reward function R(s, a, s'), its R(s, a) counterpart is defined as follows:

$$R(s,a) = \sum_{s'} P(s'|a,s) R(s,a,s')$$
(13)

Now, we consider $R'(s, a, s') \equiv R(s, a, s')$ by potential-based shaping with potential function ϕ , we have the following reward-equivalent shaping for R'(s, a):

$$R'(s,a) = \sum_{s'} P(s'|a,s)R'(s,a,s')$$

= $\sum_{s'} P(s'|a,s)(R(s,a,s') + \gamma\phi(s') - \phi(s))$
= $\sum_{s'} P(s'|a,s')R(s,a,s') + \gamma \sum_{s'} P(s'|a,s)\phi(s') - \phi(s) \sum_{s'} P(s'|a,s)$
= $R(s,a) + \gamma \sum_{s'} P(s'|a,s)\phi(s') - \phi(s)$ (14)

Proof. Proof of theorem 1, claim 1. Given $\gamma, \gamma' \in [0, 1), \gamma < \gamma'$, we prove if $\pi \in \Pi_{\gamma}$, then $\pi \in \Pi_{\gamma'}$ as well. Formally, let π be the optimal policy in MDP $M = (S, A, P, R, \gamma)$, we construct a reward function $R' \in F_R$ such that π is still optimal in MDP (S, A, P, R', γ') .

Given policy π , let P^{π} be the transition function matrix of size $|S| \times |S|$ such that $[P^{\pi}](s, s') = P(s'|\pi(s), s)$, and R^{π} be the reward vector of size |S| such that $[R^{\pi}](s) = R(s, \pi(s))$. We can write the value function $V_{R,\gamma}^{\pi}$ of policy π evaluated under the reward function R and discount factor γ as follows:

$$V_{R,\gamma}^{\pi} = R^{\pi} + \gamma P^{\pi} V_{R,\gamma}^{\pi} \tag{15}$$

Next, we apply a potential-based shaping to the original reward function R^{π} . Let $R^{'\pi}$ be the shaped reward vector of size |S| such that $[R^{'\pi}](s) = R'(s, \pi(s))$, we have:

$$R^{'\pi} = R^{\pi} + \gamma P^{\pi} \phi - \phi \tag{16}$$

where ϕ is the potential vector defined as follows:

$$\phi = \frac{\gamma' - \gamma}{\gamma} V_{R',\gamma}^{\pi} \tag{17}$$

With this shaped reward function R' and γ , we have the new value function:

$$\begin{aligned} V_{R',\gamma}^{\pi} &= R'^{\pi} + \gamma P^{\pi} V_{R',\gamma}^{\pi} \\ &= R^{\pi} + \gamma P^{\pi} \phi - \phi + \gamma P^{\pi} V_{R',\gamma}^{\pi} \\ &= R^{\pi} + \gamma \frac{\gamma' - \gamma}{\gamma} P^{\pi} V_{R',\gamma}^{\pi} - \frac{\gamma' - \gamma}{\gamma} V_{R',\gamma}^{\pi} + \gamma P^{\pi} V_{R',\gamma}^{\pi} \\ &= R^{\pi} + \gamma' P^{\pi} V_{R',\gamma}^{\pi} - \frac{\gamma' - \gamma}{\gamma} V_{R',\gamma}^{\pi} \end{aligned}$$

Rearranging, we have:

$$V_{R',\gamma}^{\pi} = (I - \gamma' P^{\pi} + \frac{\gamma' - \gamma}{\gamma})^{-1} R^{\pi}$$
(18)

Since $R \equiv R'$, their respective value function $V_{R,\gamma}^{\pi}$ and $V_{R',\gamma}^{\pi}$ induce the same set of optimal policies. We emphasize that R' is not necessary from F_R . To prove that any optimal policy $\pi \in \Pi_{\gamma}$ is still optimal in $\Pi_{\gamma'}$, we construct a $\hat{R} \in F_R$ which can make its value function $V_{\hat{R}\gamma'}^{\pi}$ evaluated with larger γ' the same as $V_{R',\gamma}^{\pi}$.

We construct \hat{R} as follows:

$$\hat{R}^{\pi} = R^{\pi} - \frac{\gamma' - \gamma}{\gamma} V^{\pi}_{\hat{R},\gamma'} \tag{19}$$

The value function $V^{\pi}_{\hat{R},\gamma'}$ for \hat{R} and γ' is:

$$V^{\pi}_{\hat{R},\gamma'} = \hat{R}^{\pi} + \gamma' P^{\pi} V^{\pi}_{\hat{R},\gamma'} \tag{20}$$

$$= R^{\pi} - \frac{\gamma' - \gamma}{\gamma} V^{\pi}_{\hat{R},\gamma'} + \gamma' P^{\pi} V^{\pi}_{\hat{R},\gamma'}$$
(21)

Rearrange, we have:

$$V_{\hat{R},\gamma'}^{\pi} = (I - \gamma' P^{\pi} + \frac{\gamma' - \gamma}{\gamma})^{-1} R^{\pi}$$
(22)

$$=V_{R',\gamma}^{\pi} \tag{23}$$

It suffices to show that the construction for \hat{R} in equation (19) satisfies $\hat{R} \in F_R$. We now write the construction for $\hat{R}(s, a)$ for every (s, a) pair:

$$\hat{R}(s,a) = R(s,a) - \frac{\gamma' - \gamma}{\gamma} V^{\pi}_{\hat{R},\gamma'}(s)$$
(24)

We notice that the second term is a factor of the value function, which only depends on the current state s. Therefore, $R(s, a^*) > R(s, a)$ iff $\hat{R}(s, a^*) > \hat{R}(s, a)$ for all $s \in S$. That is, $\hat{R} \in F_R$.

Proof. Proof of theorem 1, claim 2. $\gamma = 0$ is the special case where the planning only performs one-step look ahead and optimize the immediate reward greedily. When $\gamma = 0$, the planning objective reduces to

$$\pi_{R,\gamma=0}^* = \underset{\pi \in \Pi}{\arg\max} \underset{a_t \sim \pi(s_t)}{\mathbb{E}} [R(s_t, a_t)]$$
(25)

Given the assumption that $\forall s \in S$, $\underset{a \in A}{\operatorname{arg\,max}} R(s, a)$ is unique, then π^* is also unique and $|\Pi_0| = 1$. \Box

Proof. Proof of theorem 1, claim 3. This proof is by construction. We consider a fully connected state space with the transition function P(s'|s, a) defined below. Recall that for each $s \in S$, there exists a unique a_s^* that maximizes R(s, a), we first define the transition for each state for taking this a_s^* :

$$\forall s \in S, \ P(s'|s, a_s^*) = \begin{cases} 1 & \text{if } s' = s \\ 0 & \text{otherwise} \end{cases}$$
(26)

For all other actions $a \in A/\{a_s^*\}$ in state $s \in S$, the transition function is defined as follows:

$$P(s'|s,a) = \frac{1}{|S|-1}, \forall s' \in S/\{s\}$$
(27)

The given transition model corresponds to a fully connected state space where (s, a_s^*) creates a self-loop, and any other action has an equal probability of transitioning to different states. We construct $R_{s^*} \in F_R$ in this manner: for a state $s^* \in S$, let $R_{s^*}(s^*, a_{s^*}^*) > 2|S|R_{s^*}(s, a_s^*)$ for any other state $s \in S/\{s^*\}$. Consider an arbitrary policy π with the constraints $\pi(s^*) = a_{s^*}^*$ and $\pi(s) \neq a_s^*$ for other states. We demonstrate that this policy π is optimal in R_{s^*} and P. The optimality of π at s^* is apparent since this state is absorbing and π selects the action maximizing immediate reward. For any other state s, we show that π is optimal by calculating the optimal Q-value of $(s, \pi(s))$ compared to any other action a. Remember that for $s \neq s^*$, we constrain π not to choose a_s^* , so the alternative choice of a is to precisely select a_s^* . Therefore, we have:

$$Q^*(s,\pi(s)) = R(s,\pi(s)) + \gamma(\frac{1}{|S|-1}\frac{1}{1-\gamma}R(s^*,a^*_{s^*})$$
(28)

$$+\frac{1}{|S|-1}\sum_{s'\in S/\{s^*\}}Q(s',\pi(s'))$$
(29)

$$Q^*(s, a_s^*) = R(s, a_s^*) + \frac{\gamma}{1 - \gamma} R(s, a_s^*)$$
(30)

We have:

$$Q^{*}(s,\pi(s)) > R(s,\pi(s)) + \frac{\gamma}{1-\gamma} \frac{1}{|S|-1} R(s^{*},a_{s^{*}}^{*})$$

> $R(s,\pi(s)) + \frac{\gamma}{1-\gamma} \frac{2|S|}{|S|-1} R(s,a_{s}^{*})$
> $R(s,\pi(s)) + \frac{\gamma}{1-\gamma} 2R(s,a_{s}^{*})$ (31)

Since $2R(s, a_s^*) - R(s, a_s^*) = R(s, a_s^*) > 0$, and as γ approaches one, $\frac{\gamma}{1-\gamma}$ tends to infinity, so for sufficiently large γ we can guarantee that $Q^*(s, \pi(s)) \ge Q^*(s, a)$. Given each s^* and its corresponding R_{s^*} , under our constraints for π , there are $(|A|-1)^{|S|-1}$ such policies. In addition, since the choice of s^* is arbitrary, we can form |S| of such R_{s^*} , therefore, the total number of such policy is $(|A|-1)^{|S|-1}|S|$.

B. Proofs for Section 4

B.1. Proof for Lemma 4.5

Proof. This proof is adapted from the proof of Lemma B.1 and Lemma 3.2 in Metelli et al. (2021). Recall that we use the advantage function to derive two conditions such that the expert policy π^E is optimal under the reward function \hat{R} and $\hat{\gamma}$. Specifically,

$$Q_{\hat{R},\hat{\gamma}}^{\pi^{E}}(s,a) - V_{\hat{R},\hat{\gamma}}^{\pi^{E}}(s) = 0 \qquad \text{if } \pi^{E}(a|s) > 0,$$
(32)

$$Q_{\hat{R},\hat{\gamma}}^{\pi^{E}}(s,a) - V_{\hat{R},\hat{\gamma}}^{\pi^{E}}(s) \le 0 \qquad \text{if } \pi^{E}(a|s) = 0.$$
(33)

Consider an IRL problem $\Re = (\mathcal{M}, \pi^E)$. A Q-function satisfies the specified conditions if and only if there exist $\zeta \in \mathbb{R}^{S \times A}_{\geq 0}$ and $V \in \mathbb{R}^{|S|}$ such that:

$$Q_{\widehat{R},\widehat{\gamma}} = -\bar{B}^{\pi^E} \zeta + EV. \tag{34}$$

Given that $\pi^E \bar{B}^{\pi^E} = \mathbf{0}_S$ and $\pi^E E = I_S$, the corresponding value function is $V_{\hat{R},\hat{\gamma}} = \pi^E Q_{\hat{R},\hat{\gamma}} = V$. For any $s \in S$ and $a \in A$ with $\pi^E(a|s) > 0$, we obtain $Q_{\hat{R},\hat{\gamma}}(s,a) = V(s) = V_{\hat{R},\hat{\gamma}}(s)$. This establishes the first condition in equation 32. If

 $a \in A$ has $\pi^E(a|s) = 0$, then $Q_{\widehat{R},\widehat{\gamma}}(s,a) = -\zeta(s,a) + V(s) = -\zeta(s,a) + V_{\widehat{R},\widehat{\gamma}}(s) \leq V_{\widehat{R},\widehat{\gamma}}(s)$. This verifies the second condition in equation 33. Conversely, if $Q_{\widehat{R},\widehat{\gamma}}$ fulfills the two conditions, we set $V = V_{\widehat{R},\widehat{\gamma}}$ and $\zeta = EV_{\widehat{R},\widehat{\gamma}} - Q_{\widehat{R},\widehat{\gamma}} \leq 0$.

Next, recall that $Q_{\widehat{R},\widehat{\gamma}} = \widehat{R} + \widehat{\gamma}P\pi^E Q_{\widehat{R},\widehat{\gamma}}$. The Q-function can be written as the fixed point of the above Bellman equation: $Q_{\widehat{R},\widehat{\gamma}} = (I_{S \times A} - \widehat{\gamma}P\pi^E)^{-1}\widehat{R}$, and for $\widehat{\gamma} < 1$, the matrix is invertible. In other words, with fixed π^E , P, and $\widehat{\gamma} < 1$, there is a one-to-one correspondence between Q-functions and rewards. From equation 34, we obtain:

$$\widehat{R} = (I_{S \times A} - \widehat{\gamma} P \pi^E) (-\bar{B}^{\pi^E} \zeta + EV)$$

$$= -\bar{B}^{\pi^E} \zeta + \widehat{\gamma} P \pi^E \bar{B}^{\pi^E} \zeta + (E - \widehat{\gamma} P \pi^E E) V$$

$$= -\bar{B}^{\pi^E} \zeta + (E - \widehat{\gamma} P) V, \qquad (35)$$

since $\pi^E \bar{B}^{\pi^E} = \mathbf{0}_S$ and $\pi^E E = I_S$.

B.2. Proof for Theorem 4.6

Proof. This proof is adapted from Theorem 3.1 in Metelli et al. (2021). Note that R_0 in Theorem 4.6 is the ground-truth reward function and has the corresponding ground-truth discount factor γ_0 . Using Lemma 4.5, we express reward functions $R_0 \in \mathcal{R}_{\Re}$ and $\hat{R} \in \mathcal{R}_{\hat{\Re}}$ as:

$$R_0 = -\bar{B}^{\pi^E} \zeta + (E - \gamma_0 P) V, \tag{36}$$

$$\widehat{R} = -\overline{B}^{\widehat{\pi}^{E}}\widehat{\zeta} + (E - \widehat{\gamma}P)\widehat{V},\tag{37}$$

where $V, \hat{V} \in \mathbb{R}^S$ and $\zeta, \hat{\zeta} \in \mathbb{R}_{\geq 0}^{S \times A}$. To find the existence of $\hat{R} \in \mathcal{R}_{\hat{\Re}}$, we choose $\hat{V} = (E - \hat{\gamma}P)^{-1}(E - \gamma_0 P)V$ and $\hat{\zeta} = \bar{B}^{\pi^E} \zeta$. Then:

$$R_{0} - \hat{R} = -(\bar{B}^{\pi^{E}}\zeta - \bar{B}^{\hat{\pi}^{E}}\bar{B}^{\pi^{E}}\zeta) + (E - \gamma_{0}P)V - (E - \hat{\gamma}P)(E - \hat{\gamma}P)^{-1}(E - \gamma_{0}P)V$$

$$= -(I_{S \times A} - \bar{B}^{\hat{\pi}^{E}})\bar{B}^{\pi^{E}}\zeta$$

$$= -B^{\hat{\pi}^{E}}\bar{B}^{\pi^{E}}\zeta$$
(38)

As the expert-filter-complement $\bar{B}^{\hat{\pi}^E}$ is linear and sums with the expert-filter to unity, i.e., $B^{\hat{\pi}^E} + \bar{B}^{\hat{\pi}^E} = I_{S \times A}$, we get:

$$|R_0 - \widehat{R}| \le B^{\hat{\pi}^E} \bar{B}^{\pi^E} \zeta.$$
(39)

Finally, we obtain $\|\zeta\|_{\infty} \leq \frac{R_{\max}}{1-\gamma_0}$ by using the condition $\|R_0\| \leq R_{\max}$.

B.3. Proof of Theorem 4.7

To prove Theorem 4.7, we need three lemmas.

Lemma B.1 (lemma 1 from (Jiang et al., 2015)). For any MDP M with rewards in $[0, R_{max}]$, $\forall \pi : S \to A$ and $\widehat{\gamma} \leq \gamma_0$,

$$V_{R_0,\widehat{\gamma}}^{\pi} \le V_{R_0,\gamma_0}^{\pi} \le V_{R_0,\widehat{\gamma}}^{\pi} + \frac{\gamma_0 - \widehat{\gamma}}{(1 - \gamma_0)(1 - \widehat{\gamma})} R_{max}$$

$$\tag{40}$$

Hence, we have the following upper bound:

$$\left\| V_{R_{0},\gamma_{0}}^{\pi_{R_{0},\gamma_{0}}^{*}} - V_{R_{0},\widehat{\gamma}}^{\pi_{R_{0},\gamma_{0}}^{*}} \right\|_{\infty} \leq \frac{\gamma_{0} - \widehat{\gamma}}{(1 - \gamma_{0})(1 - \widehat{\gamma})} R_{max}$$
(41)

Intuitively, Lemma B.1 measures the performance discrepancy of a policy π when evaluated under two different discount factors: $\hat{\gamma}$ and γ_0 .

Lemma B.2 (Adapted from lemma 3 in (Jiang et al., 2015)). For any $M = (S, A, P, \hat{R}, \hat{\gamma})$ with \hat{R} bounded by $[0, R_{max}]$,

$$\left\| V_{R_{0},\widehat{\gamma}}^{\pi_{R_{0},\widehat{\gamma}}^{*}} - V_{R_{0},\widehat{\gamma}}^{\pi_{\widehat{R},\widehat{\gamma}}^{*}} \right\|_{\infty} \leq 2 \max_{\pi \in \Pi_{\widehat{\gamma}}} \left\| V_{R_{0},\widehat{\gamma}}^{\pi} - V_{\widehat{R},\widehat{\gamma}}^{\pi} \right\|_{\infty}$$
(42)

Proof. $\forall s \in S$,

$$V_{R_{0},\widehat{\gamma}}^{\pi_{R_{0},\widehat{\gamma}}}(s) - V_{R_{0},\widehat{\gamma}}^{\pi_{\widehat{R},\widehat{\gamma}}}(s)$$

$$= (V_{R_{0},\widehat{\gamma}}^{\pi_{R_{0},\widehat{\gamma}}}(s) - V_{\widehat{R},\widehat{\gamma}}^{\pi_{R_{0},\widehat{\gamma}}}(s)) - (V_{R_{0},\widehat{\gamma}}^{\pi_{\widehat{R},\widehat{\gamma}}}(s) - V_{\widehat{R},\widehat{\gamma}}^{\pi_{\widehat{R},\widehat{\gamma}}}(s)) + (V_{\widehat{R},\widehat{\gamma}}^{\pi_{R_{0},\widehat{\gamma}}}(s) - V_{\widehat{R},\widehat{\gamma}}^{\pi_{\widehat{R},\widehat{\gamma}}}(s))$$

$$\leq (V_{R_{0},\widehat{\gamma}}^{\pi_{R_{0},\widehat{\gamma}}}(s) - V_{\widehat{R},\widehat{\gamma}}^{\pi_{R_{0},\widehat{\gamma}}}(s)) - (V_{R_{0},\widehat{\gamma}}^{\pi_{\widehat{R},\widehat{\gamma}}}(s) - V_{\widehat{R},\widehat{\gamma}}^{\pi_{\widehat{R},\widehat{\gamma}}}(s))$$

$$\leq 2\max_{\pi \in \Pi_{\widehat{\gamma}}} |V_{R_{0},\widehat{\gamma}}^{\pi}(s) - V_{\widehat{R},\widehat{\gamma}}^{\pi}(s)|.$$
(43)

Intuitively, Lemma B.2 measures the difference between V-functions of two policies: the optimal policy under the groundtruth reward function R_0 , and the one under the estimated reward function \hat{R} . Both are evaluated using R_0 and the effective discount factor $\hat{\gamma}$. Lemma B.2 shows that this difference is at most double the highest V-function difference among policies in $\Pi_{\hat{\gamma}}$, when evaluated with R_0 and \hat{R} . This connects the value difference between two optimal policies evaluated under the ground truth reward to the difference in the same policy evaluated using the ground truth and estimated rewards respectively.

Lemma B.3. Let $\mathcal{M} = (S, A, P)$ be a partial MDP, and let $R_0 \in \mathbb{R}_{\geq 0}^{S \times A}$ be the ground-truth reward function and $\widehat{R} \in \mathbb{R}_{\geq 0}^{S \times A}$ be the estimated reward function with the discount factor $\widehat{\gamma}$, and let π be a policy. Then the following inequality holds:

$$\left\| V_{R_0,\widehat{\gamma}}^{\pi} - V_{\widehat{R},\widehat{\gamma}}^{\pi} \right\|_{\infty} \le \frac{1}{1 - \widehat{\gamma}} \left\| R_0 - \widehat{R} \right\|_{\infty}$$
(44)

Proof. From the bellman equation, we have:

$$V_{R_0,\widehat{\gamma}}^{\pi} = \pi R_0 + \widehat{\gamma} \pi P V_{R_0,\widehat{\gamma}}^{\pi}.$$
(45)

Rearrange this, we will have $V_{R_0,\widehat{\gamma}}^{\pi} = (I_S - \widehat{\gamma}\pi P)^{-1}\pi R_0$. Therefore,

$$V_{R_0,\widehat{\gamma}}^{\pi} - V_{\widehat{R},\widehat{\gamma}}^{\pi} = (I_S - \widehat{\gamma}\pi P)^{-1}\pi R_0 - (I_S - \widehat{\gamma}\pi P)^{-1}\pi \widehat{R}$$
$$= (I_S - \widehat{\gamma}\pi P)^{-1}\pi (R_0 - \widehat{R})$$
(46)

For the L_{∞} inequality, we simply observe:

$$\begin{aligned} \left\| V_{R_{0},\widehat{\gamma}}^{\pi} - V_{\widehat{R},\widehat{\gamma}}^{\pi} \right\|_{\infty} &= \left\| (I_{S} - \widehat{\gamma}\pi P)^{-1}\pi (R_{0} - \widehat{R}) \right\|_{\infty} \\ &\leq \left\| (I_{S} - \widehat{\gamma}\pi P)^{-1} \right\|_{\infty} \left\| \pi \right\|_{\infty} \left\| R_{0} - \widehat{R} \right\|_{\infty} \\ &\leq \frac{1}{1 - \widehat{\gamma}} \left\| R_{0} - \widehat{R} \right\|_{\infty} \end{aligned}$$

$$(47)$$

where we exploited the fact that $\|(I_S - \hat{\gamma}\pi P)^{-1}\|_{\infty} = \frac{1}{1-\hat{\gamma}}$ and that $\|\pi\|_{\infty} \leq 1$.

Using these three lemmas, we can now proceed to prove Theorem 4.7.

Proof. **Proof of Theorem 4.7** $\forall s \in S$

$$V_{R_{0},\gamma_{0}}^{\pi_{R_{0},\gamma_{0}}^{*}}(s) - V_{R_{0},\gamma_{0}}^{\pi_{\hat{R},\hat{\gamma}}^{*}}(s) = (V_{R_{0},\gamma_{0}}^{\pi_{R_{0},\gamma_{0}}^{*}}(s) - V_{R_{0},\hat{\gamma}}^{\pi_{R_{0},\gamma_{0}}^{*}}(s)) + (V_{R_{0},\hat{\gamma}}^{\pi_{R_{0},\gamma_{0}}^{*}}(s) - V_{R_{0},\gamma_{0}}^{\pi_{\hat{R},\hat{\gamma}}^{*}}(s)).$$

$$(48)$$

Using Lemma B.1, we obtain:

$$\left\| V_{R_{0},\gamma_{0}}^{\pi_{R_{0},\gamma_{0}}^{*}} - V_{R_{0},\widehat{\gamma}}^{\pi_{R_{0},\gamma_{0}}^{*}} \right\|_{\infty} \leq \frac{\gamma_{0} - \widehat{\gamma}}{(1 - \gamma_{0})(1 - \widehat{\gamma})} R_{max}.$$
(49)

For the second term, we derive the following:

$$V_{R_{0},\widehat{\gamma}}^{\pi_{R_{0},\gamma_{0}}^{*}}(s) - V_{R_{0},\gamma_{0}}^{\pi_{\widehat{R},\widehat{\gamma}}^{*}}(s) \leq V_{R_{0},\widehat{\gamma}}^{\pi_{R_{0},\gamma_{0}}^{*}}(s) - V_{R_{0},\widehat{\gamma}}^{\pi_{\widehat{R},\widehat{\gamma}}^{*}}(s) \\ \leq V_{R_{0},\widehat{\gamma}}^{\pi_{R_{0},\widehat{\gamma}}^{*}}(s) - V_{R_{0},\widehat{\gamma}}^{\pi_{\widehat{R},\widehat{\gamma}}^{*}}(s) \\ \leq 2\max_{\pi \in \Pi_{\widehat{\gamma}}} |V_{R_{0},\widehat{\gamma}}^{\pi}(s) - V_{\widehat{R},\widehat{\gamma}}^{\pi}(s)|.$$
(50)

The first line results from $\hat{\gamma} \leq \gamma_0$, the second line from $V_{R_0,\hat{\gamma}}^{\pi^*_{R_0,\hat{\gamma}'}}(s) \geq V_{R_0,\hat{\gamma}'}^{\pi^*_{R_0,\hat{\gamma}'}}(s)$ for any $\gamma' \in [0, 1]$, and the last line utilizes Lemma B.2.

Using Lemma B.3, for any policy π , we get:

$$\left\| V_{R_0,\widehat{\gamma}}^{\pi} - V_{\widehat{R},\widehat{\gamma}}^{\pi} \right\|_{\infty} \le \frac{1}{1 - \widehat{\gamma}} \left\| R_0 - \widehat{R} \right\|_{\infty}$$
(51)

By combining Equations 50 and 51, we arrive at the inequality:

$$\left\| V_{R_0,\widehat{\gamma}}^{\pi^*_{R_0,\gamma_0}} - V_{R_0,\gamma_0}^{\pi^*_{\widehat{R},\widehat{\gamma}}} \right\|_{\infty} \le \frac{2}{1-\widehat{\gamma}} \left\| R_0 - \widehat{R} \right\|_{\infty}$$
(52)

B.4. Expert Policy Estimation Error Bound

In this section, we establish a bound on the expert policy estimation error based on the coverage of expert demonstrations. Specifically, we use $\bar{B}^{\pi^E} B^{\hat{\pi}^E} \zeta$ to measure the expert policy estimation error and constrain the reward function estimation error:

$$\left|R_0 - \hat{R}\right| \le \bar{B}^{\pi^E} B^{\hat{\pi}^E} \zeta.$$
(53)

We bound this error term $\bar{B}^{\pi^E} B^{\hat{\pi}^E} \zeta$ by the number of expert-demonstrated state-action pairs. The derivation involves three key steps:

- 1. State the strategy for estimating the expert policy $\hat{\pi}^E$ from N samples of expert-demonstrated state-action pairs,
- 2. Compute the expectation of the estimated expert policy $\hat{\pi}^E$ based on N samples and subsequently the expectation of the expert policy estimation error $\bar{B}^{\pi^E} B^{\hat{\pi}^E} \zeta$.
- 3. Apply McDiarmid's Inequality to $\bar{B}^{\pi^E} B^{\hat{\pi}^E} \zeta$ as a function of N expert-demonstrated samples to bound the expert policy estimation error.

After simplification, the bound on the expert policy estimation error, in terms of the number of expert-demonstrated samples N and the effective $\hat{\gamma}$, is given by:

$$\Pr\left(\bar{B}^{\pi^{E}}B^{\hat{\pi}^{E}}\zeta(s,a) \ge t\right) \le \exp\left(\frac{-2t^{2}(1-\gamma)^{2}}{NR_{\max}^{2}}\right)$$
(54)

B.4.1. EXPERT POLICY ESTIMATION STRATEGY

Setup We estimate the expert policy $\hat{\pi}^E$ from N independently sampled state-action pairs from the expert policy π^E . For simplicity, we assume a discrete state and action space. Thus, the expert policies can be represented as a matrix of size $|S| \times |A|$, where rows correspond to states and columns to actions. Given that π^E is deterministic, each row of $\hat{\pi}^E$ contains at most one entry with a value of 1, indicating that action a is chosen at state s; all other entries are 0. Formally, the estimated π^E can be represented as:

$$\pi^{E} = \begin{pmatrix} x_{1,1} & x_{1,2} & \cdots & x_{1,A} \\ x_{2,1} & x_{2,2} & \cdots & x_{2,A} \\ \vdots & \vdots & \ddots & \vdots \\ x_{S,1} & x_{S,2} & \cdots & x_{S,A} \end{pmatrix}$$

where $x_{s,a} \in \{0,1\}$, and for each $s \in \{1, ..., |S|\}$, exactly one entry in the row is 1, while the rest are 0. For each demonstrated state-action pair, we represent them as a matrix of the same size $\widehat{\pi}_i^E \in \{0,1\}^{|S| \times |A|}$, and there is one row s_i such that:

$$\widehat{\pi}_i^E(s_i,:) = \pi^E(s_i,:)$$

That is, for each matrix $\hat{\pi}_i^E$, there exists a row s_i where $\hat{\pi}_i^E(s_i, :) = \pi^E(s_i, :)$. The remaining rows of $\hat{\pi}_i^E$ are not necessarily useful and contain value 0. We want to devise an estimation strategy that maps the N observed matrices $\hat{\pi}_1^E, \hat{\pi}_2^E, \ldots, \hat{\pi}_N^E$) to the estimated matrix $\hat{\pi}^E$.

Estimation We describe the estimation strategy below. We first take the element-wise sum for all observed partial expert policy matrices:

$$\widehat{\Pi}_{N}^{E}(s,a) = \sum_{i=1}^{N} \widehat{\pi}_{i}^{E}(s,a)$$
(55)

Specifically, each of the entry of $\widehat{\Pi}_N^E(s, a)$ counts how many times action a was chosen for state s across the N observed matrices. We estimate $\widehat{\pi}^E$ from the aggregated $\widehat{\Pi}_N^E$.

For each row s, estimate the action a that corresponds to the column where the value is 1 by finding the column with the uniquely maximum count:

$$\hat{a}_s = \arg \max_a \widehat{\Pi}_N^E(s, a) \text{ such that } \widehat{\Pi}_N^E(s, \hat{a}_s) > \widehat{\Pi}_N^E(s, a), \forall a \neq \hat{a}_s$$

Thus, the estimated row $\hat{\pi}^{E}(s, :)$ is given by placing a 1 in column \hat{a}_{s} and 0 in all other columns:

$$\hat{\pi}^{E}(s,a) = \begin{cases} 1, & \text{if } a = \hat{a}_{s} \\ 0, & \text{otherwise} \end{cases}$$

For rows that do not have a single action with the highest count, set all entries in the row to 0. Thus the final $\hat{\pi}^E$ is estimated by choosing the column with the highest vote from the observed matrices.

B.4.2. CALCULATING THE EXPECTED VALUE

Given that there are |S| states and each of the N demonstrated states is uniformly and independently sampled, we first compute of the expectation of the estimated policy $\mathbb{E}[\hat{\pi}^E]$, then the expectation of the expert policy estimation error $\bar{B}^{\pi^E}B^{\hat{\pi}^E}\zeta$.

Expected Value of Estimated Policy The count for each action in the aggregated matrix $\widehat{\Pi}_{N}^{E}(s, a)$ depends on how often state s is sampled in N independent demonstrations. Since the states are sampled uniformly, the probability of observing state s in any single demonstration is $\frac{1}{|S|}$. Let X_s be the random variable representing the number of times state s is sampled in N independent demonstrations. X_s follows a binomial distribution: $X_s \sim \text{Binomial}(N, \frac{1}{|S|})$. Thus, the expected number of times state s is sampled across the N demonstrations is:

$$\mathbb{E}[X_s] = N \cdot \frac{1}{|S|}.$$

For each time state s is sampled, the expert chooses the correct action a_s^* . Therefore, the expected count for action a_s^* at state s is:

$$\mathbb{E}[\widehat{\Pi}_{N}^{E}(s, a_{s}^{*})] = \mathbb{E}[X_{s}] = N \cdot \frac{1}{|S|}$$

For any action $a \neq a_s^*$, the expert never chooses it, so the expected count for those actions is 0:

$$\mathbb{E}[\widehat{\Pi}_N^E(s,a)] = 0 \quad \text{for all } a \neq a_s^*$$

The probability that a_s^* is the unique maximum can be approximated by the probability that state s is sampled at least once. This probability is given by $1 - \left(1 - \frac{1}{|S|}\right)^N$. Therefore, the probability that the correct action a_s^* is the unique maximum for state s is approximately:

$$P_s^{\rm unique} \approx P_s^{\rm sampled} = 1 - \left(1 - \frac{1}{|S|}\right)^N$$

The expectation of the estimated policy $\mathbb{E}[\hat{\pi}^E]$ for each state s is given by:

$$\mathbb{E}[\hat{\pi}^E(s,:)] = P_s^{\text{unique}} \cdot \pi^E(s,:)$$

Substituting P_s^{unique} with the expression derived above:

$$\mathbb{E}[\hat{\pi}^E(s,a)] = \begin{cases} 1 - \left(1 - \frac{1}{|S|}\right)^N & \text{for } a = a_s^*, \\ 0 & \text{otherwise.} \end{cases}$$

The expectation of the estimated expert policy $\mathbb{E}[\hat{\pi}^E]$ is scaled by the probability that each state s is sampled at least once across the N demonstrations. This probability increases as N grows, and in the limit as $N \to \infty$, $\mathbb{E}[\hat{\pi}^E]$ converges to the true expert policy π^E .

Expected Value of Expert Policy Estimation Error We compute the expected value of the operation $\bar{B}^{\pi^E} B^{\hat{\pi}^E} \zeta$ given N demonstrated expert state-action pairs.

We recap the key definitions for the ease of reading:

- $B^{\hat{\pi}^{E}}(.)$ represents the *expert-filter* that preserves function values only for actions taken by the approximate expert policy $\hat{\pi}^{E}(a|s)$,
- $\bar{B}^{\pi^{E}}(.)$ is the *expert-filter-complement* that preserves values for actions not taken by the ground-truth expert policy $\pi^{E}(a|s)$,
- ζ is a non-negative function bounded by $\frac{R_{\text{max}}}{1-\gamma}$.

We are interested in computing the expectation of the operation $\mathbb{E}[\bar{B}^{\pi^E}B^{\hat{\pi}^E}\zeta]$. This involves two steps:

In step 1, the expert filter $B^{\hat{\pi}^E} \zeta$ keeps the values of $\zeta(s, a)$ for actions a that are taken by the estimated expert policy $\hat{\pi}^E(a|s)$. From the expectation of $\hat{\pi}^E$, we know that:

$$\mathbb{E}[\hat{\pi}^E(s,a)] = P_s^{\text{unique}} \cdot \mathbb{1}\{a = a_s^*\}$$

Thus, the expected value of $B^{\hat{\pi}^E} \zeta(s, a)$ is:

$$\mathbb{E}[B^{\hat{\pi}^E}\zeta(s,a)] = \zeta(s,a) \cdot \mathbb{E}[\mathbb{1}\{\hat{\pi}^E(a|s) > 0\}]$$
(56)

$$= P_s^{\text{unique}} \cdot \zeta(s, a_s^*) \cdot \mathbb{1}\{a = a_s^*\}$$
(57)

For any other action $a \neq a_s^*$, $\mathbb{E}[B^{\hat{\pi}^E}\zeta(s,a)] = 0$, since $\hat{\pi}^E(s,a)$ will not take that action.

In step 2, we apply the expert-filter complement \bar{B}^{π^E} , which preserves the function values for actions not taken by the true expert policy π^E . The true expert policy π^E deterministically selects action a_s^* for each state *s*, so the expert-filter complement will only retain the values for actions $a \neq a_s^*$.

Thus, after applying the expert-filter complement, the expected value becomes:

$$\mathbb{E}[\bar{B}^{\pi^E}B^{\hat{\pi}^E}\zeta(s,a)] = \mathbb{E}[B^{\hat{\pi}^E}\zeta(s,a)] \cdot \mathbb{1}\{\pi^E(a|s) = 0\}$$

Since $\pi^E(a|s) = 0$ for all $a \neq a_s^*$, we only preserve values for actions other than a_s^* . Therefore, we have:

$$\mathbb{E}[\bar{B}^{\pi^E}B^{\hat{\pi}^E}\zeta(s,a)] = \begin{cases} P_s^{\text{unique}} \cdot \zeta(s,a_s^*) \cdot \mathbb{1}\{a = a_s^*\} & \text{if } a \neq a_s^* \\ 0 & \text{otherwise} \end{cases}$$

We now combine the two steps to compute the full expectation $\mathbb{E}[\bar{B}^{\pi^E}B^{\hat{\pi}^E}\zeta]$. Since the expert-filter complement \bar{B}^{π^E} retains only the values for actions $a \neq a_s^*$, and $B^{\hat{\pi}^E}$ retains only the values for action a_s^* , the result of applying both filters will be 0 for all s and a, because no action can satisfy both conditions simultaneously.

Thus, the expected value of the operation is:

$$\mathbb{E}[\bar{B}^{\pi^E}B^{\hat{\pi}^E}\zeta(s,a)] = 0$$

B.4.3. APPLYING MCDIARMID'S INEQUALITY ON THE EXPERT POLICY ESTIMATION ERROR

We apply McDiarmid's Inequality to bound the error in the expert policy estimation, specifically $\mathbb{E}[\bar{B}^{\pi^E}B^{\hat{\pi}^E}\zeta]$, where ζ is bounded by $\frac{R_{\max}}{1-\gamma}$.

McDiarmid's Inequality Let X_1, X_2, \ldots, X_N be independent random variables taking values in a space \mathcal{X} , and let $f : \mathcal{X}^N \to \mathbb{R}$ be a function that satisfies a bounded difference property. Specifically, if there exist constants c_i such that:

$$\sup_{1,\ldots,x_N,x'_i} |f(x_1,\ldots,x_i,\ldots,x_N) - f(x_1,\ldots,x'_i,\ldots,x_N)| \le c_i$$

then McDiarmid's Inequality states that:

x

$$\Pr\left(f(X_1,\ldots,X_N) - \mathbb{E}[f(X_1,\ldots,X_N)] \ge t\right) \le \exp\left(\frac{-2t^2}{\sum_{i=1}^N c_i^2}\right)$$

The Function $f(\cdot)$ **in Our Case** Let the function $f(\hat{\pi}_1^E, \hat{\pi}_2^E, \dots, \hat{\pi}_N^E)$ represent $\bar{B}^{\pi^E} B^{\hat{\pi}^E} \zeta$, which is determined by N independent state-action pairs sampled from the expert policy. Based on the boundedness of ζ by $\frac{R_{\text{max}}}{1-\gamma}$, we want to bound the deviation of this function from its expectation.

Sensitivity (Bounded Difference Property) For each sample $\hat{\pi}_i^E$, changing it affects at most one row of the estimated policy $\hat{\pi}^E(s,:)$, and consequently, the value of the function $\bar{B}^{\pi^E}B^{\hat{\pi}^E}\zeta$ is only affected in that row. The magnitude of this change is limited by the bound on ζ , which is $\frac{R_{\text{max}}}{1-\gamma}$.

Thus, for any change in a single sample $\hat{\pi}_i^E$, the change in the value of the function is bounded by:

$$\left|f(\widehat{\pi}_{1}^{E},\ldots,\widehat{\pi}_{i}^{E},\ldots,\widehat{\pi}_{N}^{E})-f(\widehat{\pi}_{1}^{E},\ldots,(\widehat{\pi}_{i}^{E})',\ldots,\widehat{\pi}_{N}^{E})\right| \leq \frac{R_{\max}}{1-\gamma}$$

This gives us the sensitivity constant $c_i = \frac{R_{\text{max}}}{1-\gamma}$ for each sample $\hat{\pi}_i^E$.

Applying McDiarmid's Inequality Now, we apply McDiarmid's Inequality with $c_i = \frac{R_{\text{max}}}{1-\gamma}$ for all i = 1, 2, ..., N. The inequality becomes:

$$\Pr\left(\bar{B}^{\pi^{E}}B^{\hat{\pi}^{E}}\zeta - \mathbb{E}[\bar{B}^{\pi^{E}}B^{\hat{\pi}^{E}}\zeta] \ge t\right) \le \exp\left(-\frac{2t^{2}}{N\left(\frac{R_{\max}}{1-\gamma}\right)^{2}}\right)$$

Simplifying the bound, the probability that the expert policy estimation error, measured by $\bar{B}^{\pi^E} B^{\hat{\pi}^E} \zeta$, deviates from its expected value by more than ϵ is bounded by:

$$\Pr\left(\bar{B}^{\pi^{E}}B^{\hat{\pi}^{E}}\zeta(s,a) \ge t\right) \le \exp\left(\frac{-2t^{2}(1-\gamma)^{2}}{NR_{\max}^{2}}\right)$$
(58)

As N increases, the probability of significant deviations from the expert policy estimation error diminishes exponentially, providing a high-confidence guarantee that the estimated policy will converge to the true expert policy.

Given the state space S and policy space $\Pi_{\hat{\gamma}}$ for the effective horizon $\hat{\gamma}$, we want the bound to hold uniformly over all (s, π) pairs. Using the union bound, we ensure the total failure probability remains below $\frac{\delta}{|S||\Pi_{\hat{\gamma}}|}$, hence each (s, π) pair has a small enough failure probability so that, collectively, the probability of any pair violating the bound is at most δ . By setting the RHS of the bound to $\frac{\delta}{|S||\Pi_{\hat{\gamma}}|}$, we solve for t:

$$\exp\left(\frac{-2t^{2}(1-\gamma)^{2}}{NR_{\max}^{2}}\right) = \frac{\delta}{|S||\Pi_{\widehat{\gamma}}|}$$

$$\frac{2t^{2}(1-\gamma)^{2}}{NR_{\max}^{2}} = -\ln\left(\frac{\delta}{|S||\Pi_{\widehat{\gamma}}|}\right)$$

$$t^{2} = \frac{NR_{\max}^{2}}{2(1-\gamma)^{2}}\ln\left(\frac{|S||\Pi_{\widehat{\gamma}}|}{\delta}\right)$$

$$t = \frac{R_{\max}}{1-\widehat{\gamma}}\sqrt{\frac{N}{2}\ln\left(\frac{|S||\Pi_{\widehat{\gamma}}|}{\delta}\right)}$$
(59)

We obtain the threshold that bounds the expert policy estimation error with probability at least $1 - \delta$, uniformly across all state-policy pairs.

At first glance, it may seem counterintuitive that the error bound threshold t increases with the number of samples N, as indicated in the expression of t. However, this bound should be interpreted in terms of the probability of the estimation error exceeding t, i.e., Pr(estimation error > t). As N increases, the threshold t grows, but the probability of the error exceeding this larger bound decreases exponentially.

This means that, though larger deviations are allowed with more samples, these deviations become increasingly unlikely as N increases. Therefore, for a fixed level of confidence (i.e., fixed δ), the overall estimation error *actually decreases* with more data, as the model becomes more accurate, and the likelihood of significant errors diminishes. The increasing threshold merely reflects the fact that we are accounting for possible rare deviations, but these deviations are becoming less probable as N grows.

Interpreting the Error Bound We rearrange the original bound to focus on how the estimation error behaves as N grows, for a fixed probability δ . We can rewrite the bound to emphasize the dependence on N and δ more explicitly:

$$t = \frac{R_{\max}}{1 - \gamma} \sqrt{\frac{N}{2} \ln\left(\frac{|S||\Pi_{\widehat{\gamma}}|}{\delta}\right)}$$

We have the following key observations:

Dependence on N: The term \sqrt{N} shows that the error bound increases as N increases, but the rate of increase is proportional to the square root of N. This means that to significantly reduce the error, you need a large increase in N, as the error decreases inversely with \sqrt{N} .

Dependence on δ (confidence level): The logarithmic term $\ln\left(\frac{|S||\Pi_{\hat{\gamma}}|}{\delta}\right)$ captures the effect of the confidence level. As δ (the failure probability) decreases, the logarithmic term grows, making *t* larger. This means that requiring higher confidence (lower δ) results in a higher error bound, reflecting the trade-off between confidence and error tolerance.

Simplified Bound to Show Error Decrease:

For a fixed confidence level (i.e., for fixed δ and $|S||\Pi_{\hat{\gamma}}|$), you can express the bound more clearly as:

$$t = O\left(\frac{R_{\max}}{1 - \gamma}\sqrt{N}\right)$$

This shows that, for fixed confidence, the error bound increases with the square root of N, and as a result, the probability that the error exceeds this bound decreases. In practice, this means that while larger sample sizes allow for larger deviations, these deviations become increasingly unlikely as N increases. By writing the bound in this direct form, we make it clear that as N increases, the error bound becomes more controlled (grows slowly compared to N) while the probability of exceeding this bound shrinks exponentially. This formulation provides an intuitive and direct interpretation of the trade-off between error, sample size, and confidence level.

C. Details of the Tasks

All tasks share a common action space, transition function with 9 actions per state, and a 10% chance of moving randomly.

Gridworld-simple. In each Gridworld-simple environment, there are 10×10 states with 4 randomly selected goal states. The ground-truth reward function assigns +1 for goal states and 0 for all others. The expert policy moves greedily towards the nearest goal.

Gridworld-hard. The Gridworld-hard environment consists of 15×15 states and 6 random goal states. It has a larger state space and more goal states than Gridworld-simple, but shares the same reward function.

Objectworld-linear. Each Objectworld-linear environment has 10×10 states and 10 randomly placed objects, each with a randomly assigned inner and outer color from the set $C = \{c_0, c_1\}$. Object states receive a reward of +3 for outer color c_0 and +1 for outer color c_1 . All other states have a 0 reward.

Objectworld-nonlinear. Objectworld-nonlinear differs from its linear counterpart only in the reward function. The reward depends on the outer colors of surrounding objects: +1 for states within 3 grids of c_0 and 2 grids of c_1 , -1 for states within 3 grids of c_1 , and 0 otherwise.

D. Modifications to the Linear Programming IRL method (LP-IRL)

The original objective function of LP-IRL (Ng & Russell, 2000) is formulated on the full state coverage of the expert policy and the ground-truth γ_0 . We modified the objective function to extend it to our setting with varying partial state coverage and $\hat{\gamma}$. W.L.O.G, let a_1 be the uniquely optimal action for all states. The transition probability vector at state s for the expert action is thus $P_{a_1}(s) = P(.|s, a_1)$, while that for the rest of the actions $a \in \{a_2, ..., a_k\}$ be $P_a(s) = P(.|s, a)$. Our modified objective function for LP-IRL is as follows: Given expert demonstrations covering N states,

$$\max_{\widehat{R}} \sum_{i=1}^{N} \min_{a \in \{a_2, \dots, a_k\}} (P_{a_1}(i) - P_a(i)) (I - \widehat{\gamma} \widehat{P}_{a_1})^{-1} \widehat{R} \\
\text{subject to} \quad (P_{a_1}(i) - P_a(i)) (I - \widehat{\gamma} \widehat{P}_{a_1})^{-1} \widehat{R} \ge b \\
\quad b > 0, \forall a \in \{a_2, \dots, a_k\} \\
\quad |\widehat{R}| \le R_{max}, \text{ for } i = 1, \dots, N,$$
(60)

whereas P_{a_1} is full expert policy transition matrix of size $|S| \times |S|$ such that $[P_{a_1}](s, s') = P(s'|a_1, s)$, while \hat{P}_{a_1} is an estimate of P_{a_1} from partial expert demonstrations covering only $N \leq |S|$ states. Essentially, the transition matrix of the expert policy P_{a_1} is the combination of the transition function P (known) and the expert policy π^E (unknown). Therefore, estimating transition matrix of the expert policy \hat{P}_{a_1} is equivalent to estimating the expert policy from expert demonstrations. Our modifications are highlighted below.

Estimate Expert Transition Matrix. We estimate \hat{P}_{a_1} by setting the transitions of the undemonstrated states to 0. Using this estimation, the action choices at the undemonstrated states do not affect the constraints on the other states. In addition, the value of the undemonstrated state myopically reduced to the state reward function.

Use Different Discount Factors. Our objective function also allows $\hat{\gamma}$ to differ from the ground-truth γ_0 . That is, with an estimated \hat{P}_{a_1} , we allow the mapping $\hat{F}_{i,\hat{\gamma}} = (P_{a_1}(i) - P_a(i))(I - \hat{\gamma}\hat{P}_{a_1})^{-1}$ to take different $\hat{\gamma}$ s to account the temporal effect of the policy at different horizons.

Remove L1 Regularization. The original LP-IRL uses the L_1 -regularization on the reward function, i.e., |R|, to resolve the ambiguity among the feasible reward functions. This assumption on the preferred reward function form can be a confounding factor to the performance of the induced policy, therefore we remove it to focus on the effect of horizons.

Enforce Uniquely Optimal Expert Policy. To enforce the induced policy to be uniquely optimal, we modified the constraints to be strictly positive, i.e., b > 0.

E. Modifications to the Maximum Entropy IRL method (MaxEnt-IRL)

Our theoretical analysis of the effective horizon stems from the error bound between two feasible reward function sets: the ground truth expert policy set and the approximated expert policy set derived from limited expert demonstrations. However, most IRL algorithms, such as Maximum Entropy IRL (Ziebart et al., 2008) and Maximum Margin IRL (Ratliff et al., 2006), obtain a single reward function from the feasible set using specific criteria. The error bound between the learned reward function and the ground-truth R_0 may deviate from the bound between their corresponding feasible sets. In this section, we extend our theoretical insights to Maximum Entropy Inverse Reinforcement Learning (Ziebart et al., 2008) and demonstrate that our conclusions on the effective horizon are still applicable when choosing the reward function based on specific criteria.

E.1. Maximum Entropy IRL

MaxEnt-IRL is formulated for finite horizon MDPs with horizon T_0 and no discounting. The reward function is linear in the state feature, i.e., $r = \theta^{T_0} \phi(s)$. The reward of a trajectory $\tau = (s_0, ..., s_{T_0})$ is the undiscounted sum of all state rewards, where f_{T_0} is the total feature count for the trajectory τ . MaxEnt IRL aims to match the feature counts between the policy and the expert demonstrations while maximizing the entropy of the induced policy. This leads to a distribution over behaviors constrained to match feature expectations, as shown in Equation (61).

$$p(\tau|\theta, P, T_0) \approx \frac{1}{Z(\theta, P, T_0)} \exp^{\theta^T f_{T_0}} \prod_{s_t, a_t, s_{t+1} \in \tau} P(s_{t+1}|s_t, a_t)$$
(61)

The MaxEnt IRL objective is to choose the reward parameter θ that maximizes the probability of expert demonstrations under the distribution in Equation (61):

$$\theta^* := \arg\max_{\theta} L(\theta) = \arg\max_{\theta} \sum_{\tau_i \in demo} \log p(\tau_i | \theta, P, T_0)$$
(62)

On the Effective Horizon of Inverse Reinforcement Learning



Figure 4: Summary of MaxEnt-IRL with different horizons for the four tasks. For each task, we present the ground-truth value function (column 1), ground-truth reward function (column 2), expert policy (column 3), error count curves in all states for different amount of expert data for a **single instance** of the task (columns 4-8), and finally, a summary of the error curves for a **batch** of 10 MDPs (column 9). In all four tasks, the ground-truth horizon $T_0 = 20$. The optimal horizon $\hat{T}^* < T_0$ for varying amount of expert data.

The objective function in equation (62) can be optimized using gradient-based methods, where the gradient is the difference between the empirical and learner's expected feature counts, expressed in terms of expected state visitation frequencies, D_{s_i,T_0} .

$$\nabla L(\theta) = \tilde{f} - \sum_{\tau \in demo} p(\tau|\theta, P) f_{T_0} = \tilde{f} - \sum_{s_i} D_{s_i, T_0} f_{s_i}$$
(63)

E.2. Maximum Entropy IRL with varying Effective Horizon

To extend MaxEnt IRL to varying effective horizons, we optimize trajectories with horizon $\hat{T} \leq T_0$. The corresponding optimal trajectory distribution is given in Equation (64), where $f_{\hat{T}}$ is the undiscounted sum of \hat{T} consecutive states in trajectory $\tau = (s_0, ..., s_{\hat{T}})$.

$$p(\tau|\theta, P, \widehat{T}) \approx \frac{1}{Z(\theta, P, \widehat{T})} \exp^{\theta^T f_{\widehat{T}}} \prod_{t=0}^{\widehat{T}-1} P(s_{t+1}|s_t, a_t)$$
(64)

whereas the feature factor $f_{\widehat{T}} = \sum_{t=0}^{\widehat{T}} \phi(s_t)$ is the undiscounted sum of \widehat{T} consecutive states in trajectory $\tau = (s_0, ..., s_{\widehat{T}})$.

To investigate the performance of varying effective horizons under the same amount of expert data, we adjust the expert trajectories' lengths to match the planning horizons \hat{T} , while maintaining the total number of demonstrated states constant. Given that the expert covers N states, for each effective horizon \hat{T} , we collect $N//\hat{T}$ trajectories with random initial states. Consequently, the gradient of the objective function is modified as shown in Equation (65):

$$\nabla L(\theta, \widehat{T}) = \widetilde{f}_{\widehat{T}} - \sum_{\tau \in demo} p(\tau|\theta, P, \widehat{T}) f_{\widehat{T}} = \widetilde{f}_{\widehat{T}} - \sum_{s_i} D_{s_i, \widehat{T}} f_{s_i}$$
(65)

where the expert expected feature vector $\tilde{f}_{\hat{T}}$ and the induced state visitation frequency Ds_i , \hat{T} consider only trajectories with length \hat{T} .

Our method for collecting expert demonstrations differs from the standard IRL approaches, where expert demonstrations are

fixed and given upfront. Instead, we generate expert trajectories independently for each horizon \hat{T} by rolling out trajectories with random initial states using the expert policy.

An alternative approach is to augment the fixed set of expert trajectories with length T_0 by breaking them into $T_0 - \hat{T} + 1$ shorter segments each with length \hat{T} . However, this method can slightly underestimate the visitation frequency of the final states of the original trajectories due to uniform sampling over the first $T_0 - \hat{T} + 1$ states. Our data collection method eliminates the potential bias that may arise from the quality of the expert demonstrations. In addition, our initial state distribution ρ_0 is a uniform distribution over the entire state space, which ensures a good approximation of the true state visitation frequency of the expert policy.

E.3. Results

The performance trends of MaxEnt-IRL align with those of LP-IRL, further supporting our theoretical results. Figure 4 provides a summary of the performance. Figure 5 demonstrates how the optimal horizon \hat{T}^* changes with increasing amount of expert data for MaxEnt-IRL. Figure 6 summarizes cross-validation results for MaxEnt-IRL.



Figure 5: Optimal horizons (\hat{T}^*) for MaxEnt-IRL at varying amount of expert data. We select the optimal \hat{T}^* for each of 10 sampled task environments using the algorithm in Section 5.2, based on the amount of expert data. Orange curves show how \hat{T}^* changes with the amount of expert data, while green curves display corresponding error counts. The ground-truth $T_0 = 20$ is depicted by a grey line, with corresponding error lines in blue. The trends are consistent with LP-IRL.



Figure 6: Cross-validation results for four tasks on MaxEnt-IRL. The x-axis represents the amount of expert data, while the y-axis shows error count differences in all states for various \hat{T}^* s: \hat{T}^*_{cv} learned via cross-validation, and \hat{T}^*_{oracle} chosen optimally using the oracle. Orange dots depict error count differences between the induced policies of \hat{T}^*_{cv} and \hat{T}^*_{oracle} , while blue dots represent differences between the induced polices of \hat{T}^*_{cv} and \hat{T}^*_{oracle} , while