# RILE: REINFORCED IMITATION LEARNING

**Mert Albaba**[*1,2] **Sammy Christen**[1] **Thomas Langarek**[1] **Christoph Gebhardt**[1]
**Otmar Hilliges**[1] **Michael J. Black**[2]
[1] ETH Zürich [2] Max Planck Institute for Intelligent Systems
{balbaba, sammyc, thomalan, cgebhard, otmarh}@ethz.ch, black@tue.mpg.de

## ABSTRACT

Acquiring complex behaviors is essential for artificially intelligent agents, yet learning these behaviors in high-dimensional settings poses a significant challenge due to the vast search space. Traditional reinforcement learning (RL) requires extensive manual effort for reward function engineering. Inverse reinforcement learning (IRL) uncovers reward functions from expert demonstrations but relies on an iterative process that is often computationally expensive. Imitation learning (IL) provides a more efficient alternative by directly comparing an agent's actions to expert demonstrations; however, in high-dimensional environments, such direct comparisons offer insufficient feedback for effective learning. We introduce RILe (Reinforced Imitation Learning), a framework that combines the strengths of imitation learning and inverse reinforcement learning to learn a dense reward function efficiently and achieve strong performance in high-dimensional tasks. RILe employs a novel trainer–student framework: the trainer learns an adaptive reward function, and the student uses this reward signal to imitate expert behaviors. By dynamically adjusting its guidance as the student evolves, the trainer provides nuanced feedback across different phases of learning. Our framework produces high-performing policies in high-dimensional tasks where direct imitation fails to replicate complex behaviors. We validate RILe in challenging robotic locomotion tasks, demonstrating that it significantly outperforms existing methods and achieves near-expert performance across multiple settings.

## 1 INTRODUCTION

Over the years, reinforcement learning (RL) has emerged as a powerful framework for teaching agents to perform sophisticated tasks, yet it often requires extensive manual reward function design, which is both time-consuming and error-prone.

There are two ways to address the reward engineering problem. First, Inverse Reinforcement Learning (IRL) (Ng & Russell, 2000; Ziebart et al., 2008) offers a remedy by inferring the reward function from expert demonstrations, thus reducing the burden of manual reward engineering. IRL proceeds iteratively: it first trains a policy (the learning agent's decision-making mechanism) using the current reward function, observes how well the agent's behavior aligns with the expert's, and then refines the reward function to better guide the policy toward expert-like behaviors. Repeating this process eventually yields a reward function capable of providing nuanced feedback at different stages of learning. However, this iterative procedure is computationally expensive (Zheng et al., 2022), especially in high-dimensional environments where both the reward and the policy must explore a large state-action space.

Second, Imitation learning (IL) bypasses explicit reward design by directly comparing learned behaviors to expert demonstrations via a comparison mechanism. Traditional IL approaches such as Behavioral Cloning (BC) (Bain & Sammut, 1995) match the learned actions to expert demonstrations directly, requiring a substantial amount of expert data in high-dimensional tasks. To improve data efficiency, Adversarial Imitation Learning (AIL) methods, such as GAIL (Ho & Ermon, 2016), introduce a discriminator that operates as a comparison mechanism and judges how expert-like the learned behaviors are. However, both traditional IL and AIL lack a reward function that emphasizes specific subgoals or partial improvements. Instead, they rely on a comparison metric, often a distance measure or binary classification, that merely checks whether the agent's behavior is (or is not) similar

to the expert. Such comparison-based signals offer no fine-grained guidance on which specific actions or sub-strategies to prioritize. Consequently, both traditional IL and AIL struggle in high-dimensional environments (Peng et al., 2018; Garg et al., 2021), where the agent needs more granular and adaptive feedback than these mechanisms provide.

Adversarial Inverse Reinforcement Learning (AIRL) (Fu et al., 2018) attempts to remedy IRL's inefficiency by integrating a learned reward function within a discriminator. However, AIRL tightly couples the reward function to the discriminator's output, causing it to inherit AIL's limitations in high-dimensional settings where more fine-grained guidance is needed.

Real-life learning scenarios suggest a different approach: think of parents and children, or a pet owner and their dog. The *teacher* also refines how they teach as the student progresses. Each success or failure in the student's understanding informs the teacher's approach, creating a *positive-sum* relationship: lessons learned from suboptimal behaviors ultimately yield better trainers, which, in turn, guide the student's learning progress more effectively. By contrast, existing approaches lack this continuous cooperative synergy. Adversarial Imitation Learning (AIL) does update a discriminator alongside the policy, but the discriminator's sole role is to distinguish expert-like behavior from non-expert behavior. Consequently, the student attempts to fool this judge into classifying its behavior as expert-like, resulting in a *competitive* process rather than a *cooperative trainer* that dynamically shapes rewards based on suboptimal behaviors. Meanwhile, IRL methods only refine the reward *after* the policy converges, missing the opportunity for continuous co-evolution throughout training.

To address these issues, inspiring by these insights, we propose Reinforced Imitation Learning (RILe). RIle combines the adaptive reward benefits of IRL with the computational efficiency of AIL (Fig. 1-(d)). RILe is a novel *trainer-student* system that establishes a positive-sum relationship between the trainer and the student. Specifically, RILe is composed of:

- **Student Agent:** Learns a policy to imitate expert demonstrations using reinforcement learning.
- **Trainer Agent:** Simultaneously learns a reward function using reinforcement learning, leveraging an adversarial discriminator for continuous feedback on student performance.

RILe's trainer continuously updates the reward function in tandem with the student's policy updates, where IRL refines its reward function only after training a policy to convergence. Specifically, the trainer queries a discriminator to measure how expert-like the student's behavior is, then optimizes the reward function based on that feedback, without waiting for the policy to converge. RILe preserves nuanced reward shaping, while avoiding IRL's heavy computational loop. As a result, RILe is particularly effective in high-dimensional settings, where agents need fine-grained guidance at every stage of learning. Our contributions are two-fold:

1. **Efficient Reward-Function Learning via RL**: We introduce a reinforcement-learning-based approach for training a reward function simultaneously with the policy. This avoids IRL's repeated policy re-training and the purely discriminator-based rewards of AIL/AIRL. Reinforcement learning enables the trainer agent to explore multiple reward regimes, yielding a reward function that considers long-horizon effects.
2. **Dynamic Reward Customization:** RILe offers context-sensitive guidance at every stage of training, because the trainer agent updates the reward function as the student evolves. This dynamic shaping is especially valuable in high-dimensional tasks, where the learning agent requires different forms of encouragement during intermediate-stages than later-stages of the training. Consequently, RILe enables accurate imitation of the expert performance in high-dimensional tasks.

We evaluate RILe in comparison to state-of-the-art methods in AIL, IRL, and AIRL, specifically GAIL (Ho & Ermon, 2016) AIRL (Fu et al., 2018), GAIfO (Torabi et al., 2018b), BCO (Torabi et al., 2018a), IQ-Learn (Garg et al., 2021) and DRAIL (Lai et al., 2024). Our experiments spans four scenarios: (1) Empirically investigating how RILe's reward-learning strategy differs from AIL and AIRL in a maze task, (2) Quantifying the adaptability of the learned reward function in a humanoid locomotion task, (3) Assessing RILe's performance in both low- and high-dimensional continuous-control problems, and (4) Analyzing the impact of explicit expert-data usage within RILe's training process. Our results show RILe's superior performance, particularly in high-dimensional environments, and highlight RILe's ability to learn a dynamic reward function that effectively guides the student through multiple stages of training.

## 2 RELATED WORK

We review research on learning from expert demonstrations, focusing on Imitation Learning (IL) and Inverse Reinforcement Learning (IRL), the conceptual foundations of RILe.

**Imitation Learning** Early work in IL introduced Behavioral Cloning (BC) (Bain & Sammut, 1995), which frames policy learning as a supervised problem where the agent's actions are directly matched to expert demonstrations. DAgger (Ross et al., 2011) refines BC by aggregating data over multiple iterations to mitigate compounding errors. GAIL (Ho & Ermon, 2016) employs adversarial training: a discriminator learns to distinguish expert trajectories from the agent's, while the generator (agent) adapts to mimic expert-like behavior. BCO (Torabi et al., 2018a) extends BC, and GAIfO (Torabi et al., 2018b) extends GAIL, both to handle state-only observation scenarios. DQfD (Hester et al., 2018) introduces a two-stage approach with pre-training, while ValueDice (Kostrikov et al., 2020) aligns policy and expert distributions via a distribution-matching objective. More recently, DRAIL (Lai et al., 2024) leverages a diffusion-based discriminator to enhance learning efficiency in adversarial imitation. Despite these advances, IL methods face challenges in high-dimensional environments (Peng et al., 2018; Garg et al., 2021), where naive action matching or purely adversarial comparisons fail to provide sufficiently granular guidance. RILe addresses these limitations through an adaptive trainer–student framework, where a learned reward function provides more nuanced guidance than standard IL comparison mechanisms.

**Inverse Reinforcement Learning** Inverse Reinforcement Learning (IRL), introduced by Ng & Russell (2000), aims to uncover the expert's intrinsic reward function from demonstrations. Major developments include Apprenticeship Learning (Abbeel & Ng, 2004), Maximum Entropy IRL (Ziebart et al., 2008), and adversarial variants like AIRL (Fu et al., 2018). IQ-Learn (Garg et al., 2021) reformulates IRL by integrating the inverse reward learning process into Q-learning for better scalability. More recent work focuses on unstructured data (Chen et al., 2021) and cross-embodiment transfer (Zakka et al., 2022).

Nonetheless, IRL methods struggle with computational inefficiency and limited scalability (Arora & Doshi, 2021), particularly in high-dimensional tasks where repeated iterations of policy learning and reward refinement become costly. RILe mitigates these challenges by jointly learning the policy and reward function in a single process, avoiding IRL's iterative retraining loop and facilitating more efficient reward shaping for complex environments.

## 3 BACKGROUND

### 3.1 MARKOV DECISION PROCESS

A standard Markov Decision Process (MDP) is defined by $(S, A, R, T, K, \gamma)$. $S$ is the state space consisting of all possible environment states $s$, and $A$ is action space containing all possible environment actions $a$. $R = R(s, a) : S \times A \rightarrow \mathbb{R}$ is the reward function. $T = \{P(\cdot|s, a)\}$ is the transition dynamics where $P(\cdot|s, a)$ is an unknown state state transition probability function upon taking action $a \in A$ in state $s \in S$. $K(s)$ is the initial state distribution, i.e., $s_0 \sim K(s)$ and $\gamma$ is the discount factor. The policy $\pi = \pi(a|s) : S \rightarrow A$ is a mapping from states to actions. In this work, we consider $\gamma$-discounted infinite horizon settings. Following Ho & Ermon (2016), expectation with respect to the policy $\pi \in \Pi$ refers to the expectation when actions are sampled from $\pi(s)$: $\mathbb{E}_\pi[R(s, a)] \triangleq \mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t)]$, where $s_0$ is sampled from an initial state distribution K$(s)$, $a_t$ is given by $\pi(\cdot|s_t)$ and $s_{t+1}$ is determined by the unknown transition model as $P(\cdot|s_t, a_t)$. The unknown reward function $R(s, a)$ generates a reward given a state-action pair $(s, a)$. We consider a setting where $R = R(s, a)$ is parameterized by $\theta$ as $R_\theta(s, a) \in \mathbb{R}$ (Finn et al., 2016).

Our work considers an imitation learning problem from expert trajectories, consisting of states $s$ and actions $a$. The set of expert trajectories $\tau_E$ are sampled from an expert policy $\pi_E \in \Pi$, where $\Pi$ is the set of all possible policies. We assume that we have access to $m$ expert trajectories, all of which have $n$ time-steps, $\tau_E = \{(s_0^i, a_0^i), (s_1^i, a_1^i), \ldots, (s_n^i, a_n^i)\}_{i=1}^m$.

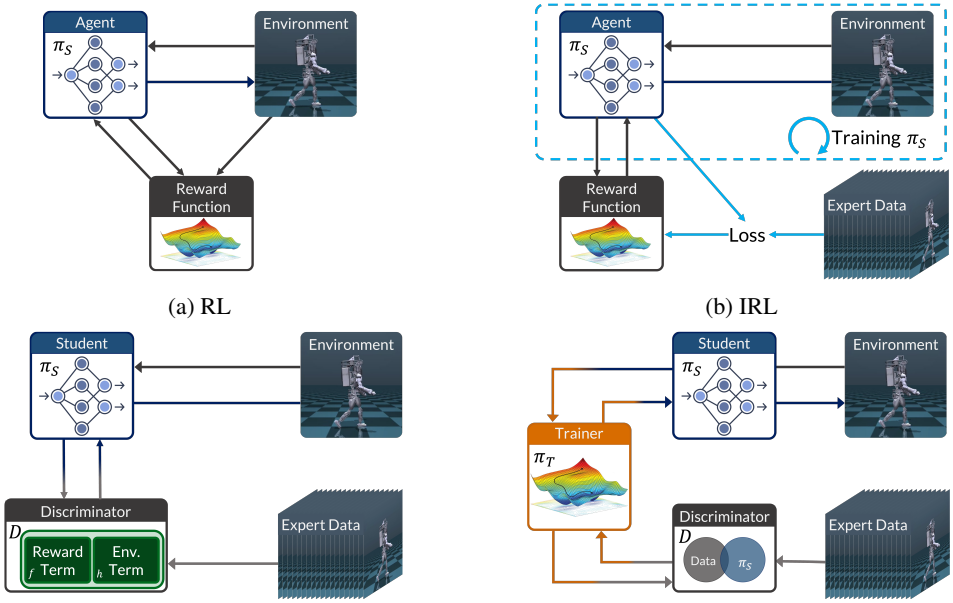


(a) RL

(b) IRL

(c) GAIL + AIRL (terms in green)

(d) RILe

Figure 1: **Overview of the related works. (a) Reinforcement Learning (RL):** learning a policy that maximizes hand-defined reward function; **(b) Inverse RL (IRL):** learning a reward function from data. IRL has two stages: 1. training a policy with frozen reward function, and 2. updating the reward function by comparing the converged policy with data. These stages repeated several times; **(C) Generative Adversarial Imitation Learning (GAIL) + Adversarial IRL (AIRL):** using discriminator as a reward function. GAIL trains both policy and the discriminator at the same time. AIRL implements a new structure on the discriminator, seperating reward from environment dynamics by using two networks under the discriminator (see additional terms in green). **(D) RILe:** similar to IRL, learning a reward function from data. RILe learns the reward function at the same time with the policy, using a discriminator as a guide for learning the reward.

### 3.2 Reinforcement Learning (RL)

Reinforcement learning seeks to find an optimal policy, $\pi^*$. that maximizes the discounted cumulative reward given from the reward function $R = R(s, a)$ (Fig. 1-(a)). In this work, we incorporate entropy regularization using the $\gamma$-discounted casual entropy function $H(\pi) = \mathbb{E}_\pi[-\log \pi(a|s)]$ (Ho & Ermon, 2016; Bloem & Bambos, 2014). The RL problem with a parameterized reward function and entropy regularization is defined as

$$\text{RL}(R_\theta(s, a)) = \pi^* = \arg\max_\pi \mathbb{E}_\pi[R_\theta(s, a)] + H(\pi). \tag{1}$$

### 3.3 Inverse Reinforcement Learning (IRL)

Given sample trajectories $\tau_E$ from an optimal expert policy $\pi_E$, inverse reinforcement learning aims to recover a reward function $R^*_\theta(s, a)$ that maximally rewards the expert's behavior (Fig. 1-(b)). Formally, IRL seeks a reward function, $R^*_\theta(s, a)$, satisfying: $\mathbb{E}_{\pi_E}[\sum_{t=0}^{\infty} \gamma^t R^*_\theta(s_t, a_t)] \geq \mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t R^*_\theta(s_t, a_t) + H(\pi)] \quad \forall \pi$. Optimizing this reward function with reinforcement learning yields a policy that replicates expert behavior: $\text{RL}(R^*_\theta(s, a)) = \pi^*$. Since only the expert's trajectories are observed, expectations over $\pi_E$ are estimated from samples in $\tau_E$. Incorporating entropy regularization $H(\pi)$, maximum causal entropy inverse reinforcement learning (Ziebart et al., 2008) is defined as

$$\text{IRL}(\tau_E) = \arg\max_{R_\theta(s,a) \in \mathbb{R}} \left( \mathbb{E}_{s,a \in \tau_E}[R_\theta(s, a)] - \max_\pi \left( \mathbb{E}_\pi[R_\theta(s, a)] + H(\pi) \right) \right). \tag{2}$$

### 3.4 Adversarial Imitation Learning (AIL) and Adversarial Inverse Reinforcement Learning (AIRL)

Imitation Learning (IL) aims to directly approximate the expert policy from given expert trajectory samples $\tau_E$. It can be formulated as $\text{IL}(\tau_E) = \arg\min_\pi \mathbb{E}_{(s,a)\sim\tau_E}[L(\pi(\cdot|s), a)]$, where $L$ is a loss function, that captures the difference between policy and expert data.

GAIL (Ho & Ermon, 2016) introduces an adversarial imitation learning setting by quantifying the difference between the agent and the expert with a discriminator $D_\phi(s, a)$, parameterized by $\phi$ (Fig. 1-(c)). The discriminator distinguishes between between expert-generated state-action pairs $(s, a) \sim \tau_E$ and non-expert ones $(s, a) \notin \tau_E$. The goal of GAIL is to find the optimal policy that fools the discriminator while maximizing an entropy constraint. The optimization is formulated as a zero-sum game between the discriminator $D_\phi(s, a)$ and the policy $\pi$:

$$\min_\pi \max_\phi \mathbb{E}_\pi[\log D_\phi(s, a)] + \mathbb{E}_{\tau_E}[\log (1 - D_\phi(s, a))] - \lambda H(\pi). \tag{3}$$

In other words, the reward function that is maximized by the policy is defined as a similarity function, expressed as $R(s, a) = -\log (D_\phi(s, a))$.

AIRL (Fu et al., 2018) extends AIL to inverse reinforcement learning, aiming to recover a reward function decoupled from environment dynamics (Fig. 1-(c)). AIRL structures the discriminator as:

$$D_{\phi,\psi}(s, a, s') = \frac{\exp(f_\phi(s, a, s'))}{\exp(f_\phi(s, a, s')) + \pi(a|s)}, \tag{4}$$

where $f_\phi(s, a, s') = r_\psi(s, a) + \gamma V_\phi(s') - V_\phi(s)$. Here, $r_\psi(s, a)$ represents the learned reward function that is decoupled from the environment dynamics, $\gamma V_\phi(s') - V_\phi(s)$. The AIRL optimization problem is formulated equivalently to GAIL (see Eqn. 3). The reward function $r_\psi(s, a)$ is learned through minimizing the cross-entropy loss inherent in this adversarial setup. Therefore, the reward function remains tightly coupled with the discriminator's learning process.

## 4 RILe: Reinforced Imitation Learning

We propose Reinforced Imitation Learning (RILe) to jointly learn a reward function and a policy that emulates expert-like behavior within a single learning process. RILe introduces a novel trainer–student dynamic, as illustrated in Figure 2.

In RILe, the student agent learns an action policy by interacting with the environment, while the trainer agent learns a reward function that effectively guides the student toward expert-like behavior. Both agents are trained simultaneously via reinforcement learning, with an assistance from an adversarial discriminator. Specifically, the trainer queries the discriminator, which judges how expert-like the student's behavior is, and then optimizes the reward function based on that feedback on-the-fly. Unlike traditional AIL, where the discriminator effectively is employed as the reward function for the student, RILe introduces a trainer agent to provide fine-grained feedback to the student, while avoiding IRL's iterative computational expense.

The trainer agent plays the key role in RILe. Trained via RL, the trainer explores different reward designs and learns to provide gradually tailored feedback to the student by maximizing the cumulative rewards it receives from the discriminator. This approach equips RILe with two key advantages over existing IRL/AIRL/AIL frameworks: (1) On-the-fly reward function learning via RL: The reward function is learned continuously with RL, enabling the trainer to explore different reward options and account for long-horizon effects of its signals, (2) Context-sensitive guidance: The trainer adjusts its reward outputs in response to the student's current policy, thereby encouraging the student to explore suboptimal actions that ultimately guide it closer to expert behavior. By providing tailored feedback at different stages of training, RILe addresses the limitations of prior methods, particularly in high-dimensional tasks.

In the remainder of this section, we define the components of RILe and explain how they jointly learn from expert demonstrations.

**Student Agent** The student agent learns a policy $\pi_S$ by interacting with an environment in a standard RL setting within an MDP. For each of its actions $a^S \in A$, the environment returns a new

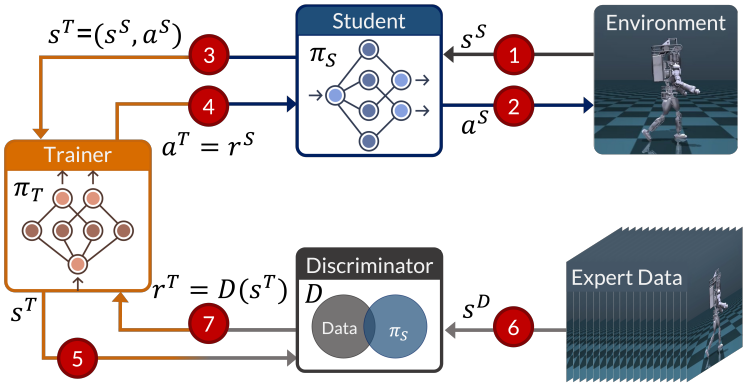


Figure 2: **Reinforced Imitation Learning (RILe)**. The framework consists of three key components: a student agent, a trainer agent, and a discriminator. The student agent learns a policy $\pi_S$ by interacting with an environment, and the trainer agent learns a reward function as a policy $\pi_T$. (1) The student receives the environment state $s^S$. (2) The student takes an action $a^S$, forwards it to the environment which is updated based on $a^S$. (3) The student forwards its state and action to the trainer, whose state is $s^T = (s^S, a^S)$. (4) Trainer, $\pi_T$, evaluates the state action pair of the student agent $s^T = (s^S, a^S)$ and chooses an action $a^T$ that then becomes the reward of the student agent $a^T = r^S$. (5) The trainer agent forwards the $s^T = (s^S, a^S)$ to the discriminator. (6) Discriminator compares student state-action pair with expert demonstrations ($s^D$). (7) Discriminator gives reward to the trainer, based on the similarity between student- and expert-behavior.

state $s^S \in S$. However, instead of using a handcrafted reward function, the student's reward comes from the trainer agent's policy, $\pi_T$. Therefore, the reward function is represented by the trainer policy. Thus, the student agent is guided by the actions of the trainer agent, i.e., the action of the trainer is the reward of the student: $r^S = \pi_T((s^S, a^S))$. The optimization problem of the student agent is then defined as

$$\min_{\pi_S} -\mathbb{E}_{(s^S,a^S)\sim\pi_S}[\pi_T\left((s^S, a^S)\right)]. \tag{5}$$

**Discriminator** The discriminator differentiates between expert-generated state-action pairs, $(s, a) \sim \tau_E$, and pairs from the student, $(s, a) \sim \pi_S$. In RILe, the discriminator is defined as a feed-forward deep neural network, parameterized by $\phi$. Its objective is:

$$\max_{\phi} \mathbb{E}_{(s,a)\sim\tau_E}[\log(D_\phi(s, a))] + \mathbb{E}_{(s,a)\sim\pi_S}[\log(1 - D_\phi(s, a))]. \tag{6}$$

To provide effective guidance, the discriminator must accurately identify whether a given state–action pair originates from the expert distribution $(s, a) \sim \tau_E$ or not $(s, a) \notin \tau_E$. GAIL (Ho & Ermon, 2016) established the feasibility of such a discriminator (see Appendix B for details).

**Trainer Agent** The trainer agent guides the student toward expert behavior by serving as its reward mechanism. Since the trainer does not directly observe the student's policy $\pi_S$, we model the trainer's environment as a Partially Observable MDP (POMDP): $\text{POMDP}_T = (S_T, A_T, \Omega_T, T_T, O_T, R_T, \gamma)$. The state space $S_T = S \times A \times \pi_S$ includes all possible state-action pairs from the standard MDP and the student's policy $\pi_S$, which is hidden from the trainer, introducing partial observability. The trainer's action space, $A_T$, consists of scalar values. Formally, $A_T$ is defined as a mapping from $S_T \rightarrow \mathbb{R}$. The observation space $\Omega_T = S \times A$ consists of the observable state-action pairs of the student. The transition dynamics $T_T$ and the observation function $O_T$ are defined formally in Appendix A. The reward function $R_T(s^T, a^T)$ evaluates the effectiveness of the trainer's action in guiding the student, where $s^T = (s^S, a^S)$ is the observation of the trainer. $\gamma$ is the discount factor.

Within this POMDP, the trainer learns a policy $\pi_T$ that produces helpful reward signals for $\pi_S$. The trainer observes only the student's state–action pair, $s^T = (s^S, a^S) \in S \times A$, not $\pi_S$ itself. It then outputs a scalar action $a^T \in [-1, 1]$, which is provided to the student as the reward, $r^S$.

If the trainer's reward depends only on the discriminator's output, the trainer receives the same reward regardless of whether it rewards or penalizes the student, yielding no immediate feedback on its choices. For instance, if the student behaves like the expert and the discriminator outputs $\approx 1$, the

trainer should ideally reward the student (action, $a^T$, $\approx 1$). But if the trainer's action is not factored into its own reward, it gains no immediate signal whether rewarding or punishing the student was effective, since it receives the same reward in either case. This ambiguity forces extensive trial and error. To address this, we define the trainer reward as:

$$R^T = \upsilon(D_\phi(s^T))(a^T) \tag{7}$$

where $\upsilon(x) = 2x - 1$ scales the discriminator's output, making it symmetric around zero. Including $a^T$ in the trainer's reward ensures the trainer effectively learns from its own actions. Formally, we define the trainer's objective as:

$$\max_{\pi_T} \mathbb{E}_{\substack{(s,a)\sim\pi_S \\ a^T\sim\pi_T}}[\upsilon(D_\phi(s,a))a^T]. \tag{8}$$

**RILe** RILe brings together these three components, student, trainer, and discriminator, to discover a student policy that imitates expert behaviors in $\tau_E$. Both $\pi_S$ and $\pi_T$ can be trained via any single-agent RL method. The overall training algorithm is detailed in Appendix J.

The student agent aims to recover the optimal policy $\pi_S^*$:

$$\pi_S^* = \arg\max_{\pi_S} \mathbb{E}_{(s^S,a^S)\sim\pi_S}\left[\sum_{t=0}^{\infty}\gamma^t[\pi_T\left((s_t^S, a_t^S)\right)]\right]. \tag{9}$$

Simultaneously, the trainer aims to recover $\pi_T^*$:

$$\pi_T^* = \arg\max_{\pi_T} \mathbb{E}_{\substack{s^T\sim\pi_S \\ a^T\sim\pi_T}}\left[\sum_{t=0}^{\infty}\gamma^t[\upsilon(D_\phi(s_t^T))a_t^T]\right]. \tag{10}$$

By optimizing these objectives together, RILe efficiently learns both a reward function and a policy in high-dimensional settings where traditional AIL or IRL methods often struggle. Details on specific training strategies are provided in Appendix C.

## 5 EXPERIMENTS

We evaluate the performance of RILe by addressing four key questions:

1. How does RILe's reward-learning strategy differs from AIL and AIRL?
2. How adaptive is RILe's learned reward function?
3. How does RILe perform in high-dimensional robotic imitation tasks compared to AIL/AIRL/IRL?
4. How explicitly using expert-data within RILe's training affects the context-sensitive reward learning?

**Baselines** We compare RILe with seven baseline methods: Behavioral cloning (BC (Bain & Sammut, 1995; Ross & Bagnell, 2010), BCO (Torabi et al., 2018a)), adversarial imitation learning (GAIL (Ho & Ermon, 2016), GAIfO (Torabi et al., 2018b) and DRAIL (Lai et al., 2024)), adversarial inverse reinforcement learning (AIRL (Fu et al., 2018)), and inverse reinforcement learning (IQ-Learn (Garg et al., 2021)). DRAIL (Lai et al., 2024) introduces a diffusion-based discriminator implementation, which is applied to both GAIL and RILe, and referred as DRAIL-GAIL and DRAIL-RILe. Additional experimental details are provided in the Appendix D.

### 5.1 EVOLVING REWARD FUNCTION

To answer the first question about RILe's reward-learning strategy, we compare RILe's performance with AI(R)L baselines in a maze setting. In this environment, the agent must navigate from a fixed start to a goal while avoiding static obstacles; we use a single expert demonstration.

Figure 3 shows how each method's learned reward function evolves during training. For RILe, we plot the trainer's learned reward function. For GAIL and AIRL, we visualize the discriminator outputs. The columns represent reward landscapes at 25%, 50%, 75%, and 100% of the total training process, and each subplot overlays the student's trajectory from the previous epoch.

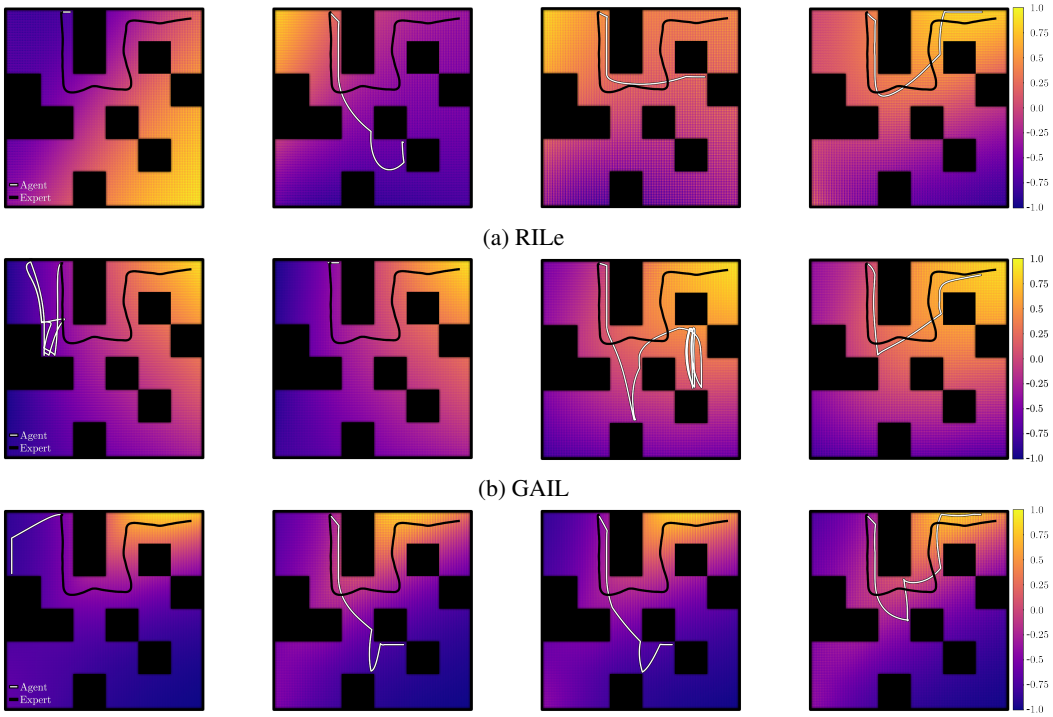


(a) RILe

(b) GAIL

(c) AIRL

Figure 3: **Reward Function Comparison**. Evolution of reward functions during training for (a) RILe, (b) GAIL, and (c) AIRL in a continuous maze environment. Columns show reward landscapes at 25%, 50%, 75%, and 100% of training completion (left to right). The expert's trajectory is shown in black, while the student agent's trajectory from the previous training epoch is in white. Color gradients represent reward values, with darker colors indicating lower rewards and brighter colors indicating higher rewards. Black squares represent obstacles. RILe demonstrates a dynamic reward function that adapts with the student's progress, while GAIL and AIRL maintain relatively static reward landscapes throughout training and struggle to adapt.

RILe's reward function dynamically adapts to the student's current policy, providing guidance that encourage suboptimal actions which eventually lead the student closer to the expert trajectory. By contrast, GAIL and AIRL's reward functions remain relatively static. Specifically, the first column in Figure 3 shows RILe's trainer encouraging exploration toward the bottom-right of the maze, which is initially suboptimal but helpful in the long run. As the student learns to reach the lower part of the maze, RILe shifts high-reward regions toward the top-left (second column), again encouraging incremental progress. The third column illustrates how RILe boosts rewards near the goal while still maintaining some incentive around top left areas to keep the agent from getting stuck. Overall, RILe's evolving reward function serves as a curriculum that promotes gradual improvement toward expert-like performance. This dynamic reward adaptation gets important in higher-dimensional tasks as we show in Section 5.3.

### 5.2 Reward Function Dynamics

To address the second question on adaptability, we quantify how the reward function evolves and its relationship to the student's performance. We compare RILe with GAIL, DRAIL-GAIL, and DRAIL-RILe in a high-dimensional robotic control scenario (learning to walk with UnitreeH1 robot).

We introduce three metrics (see Appendix D.2 for details): (1) RFDC (Reward Function Distribution Change): Wasserstein distance between reward distributions over consecutive training intervals, capturing overall shifts in reward space, (2) FS-RFDC (Fixed-State Reward Function Distribution Change): Mean absolute deviation of reward values at a fixed set of expert states over time, (3) CPR (Correlation between Performance and Reward): Pearson correlation between changes in the reward function and changes in the student's performance.

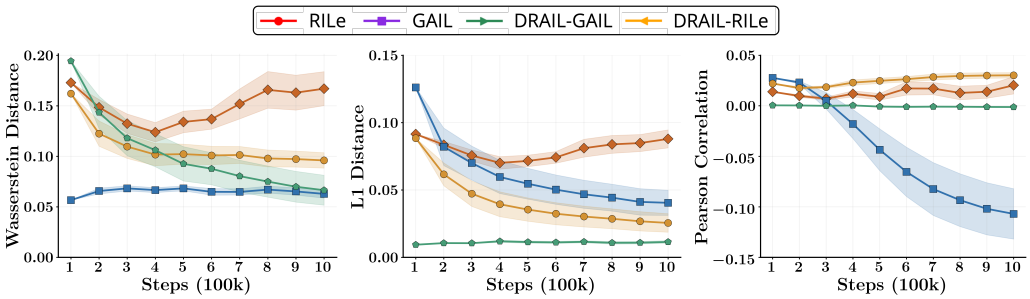


(a) RFDC (b) FS-RFDC (c) CPR

Figure 4: **Dynamics of Reward Functions**. **(a) Reward Function Distribution Change (RFDC):** Wasserstein distance between reward function distributions. **(b) Fixed-State Reward Function Distribution Change (FS-RFDC):** Mean absolute deviation of reward values for a fixed set of expert states. **(c) Correlation between Performance and Reward (CPR):** Pearson correlation between changes in the reward function and changes in the student's performance.

#### 5.2.1 Adaptability of the Learned Reward Function

We assess how adaptive the reward function learned by RILe is compared to that of AIL. Fig. 4a presents changes in reward distributions over 10,000 consecutive steps. RILe exhibits the highest adaptability in its reward function, aligning with our goal of having the reward function adapt based on the student's learning stage. The advanced discriminator in DRAIL reduces the need for drastic reward function changes, yet RILe remains more adaptive than GAIL. Since the changing student policy indirectly affects RFDC, we also show changes in reward values for the fixed set of states in Fig. 4b. Again, RILe's reward function is the most adaptive among all methods.

Higher adaptability of RILe ensures that the reward signal remains aligned with incremental performance improvements, enabling the student to receive more timely and effective guidance throughout training.

#### 5.2.2 Correlation between the Learned Reward and the Student Performance

We evaluate how changes in the reward function correlate with improvements in student performance. To this end, Fig. 4c presents the Pearson correlation between student's performance and reward updates. DRAIL-RILe achieves the highest positive correlation, indicating that it learns the most effective rewards for improving student performance. RILe ranks second, demonstrating that the trainer agent effectively helps the student achieve better scores even with the help of a naive-discriminator. In contrast, GAIL's correlation starts positive but soon turns negative and remains so throughout training. We hypothesize that this occurs because the discriminator in GAIL tends to saturate as training progresses. While the discriminator's reward signal effectively guides learning early on, its increasingly static nature at later stages fails to capture subtle performance improvements, leading to a negative correlation.

### 5.3 Motion-Capture Data Imitation for Robotic Continuous Control

To answer the third question, we test RILe in high-dimensional robotic locomotion tasks (Al-Hafez et al., 2023), where the agent must imitate motion-capture data for various robotic bodies. This benchmark is especially demanding due to its complexity and dimensionality.

Table 1 presents results for seven LocoMujoco tasks across different test seeds (see Appendix D.3 for details). Overall, these results underscore that RILe outperforms the AIL/IL/IRL baselines, and benefits even more from the enhancements provided by the DRAIL variants, achieving performance levels close to the expert. The performance variations across tasks indicate that although RILe's adaptive reward function is highly effective, task-specific factors also play a role. This overall performance aligns with our claim that an adaptive reward function is crucial for mastering complex, high-dimensional behaviors

Table 1: Test results on seven LocoMujoco tasks.

| | | RILe | GAIL | AIRL | IQ | BCO | GAIfO | DRAIL GAIL | DRAIL RILe | Expert |
|---|---|---|---|---|---|---|---|---|---|---|
| Walk | Atlas | 870.6 | 792.7 | 300.5 | 30.9 | 21.0 | 834.2 | 834.4 | **899.1** | 1000 |
| | Talos | 842.5 | 442.3 | 102.1 | 4.5 | 11.9 | 710.0 | 787.7 | **896.6** | 1000 |
| | UnitreeH1 | 966.2 | 950.2 | 568.1 | 8.8 | 34.8 | 526.8 | 940.8 | **995.8** | 1000 |
| | Humanoid | **831.3** | 181.4 | 80.1 | 4.5 | 3.5 | 706.5 | 814.6 | 527.6 | 1000 |
| Carry | Atlas | **850.8** | 669.3 | 256.4 | 36.8 | 20.3 | 810.1 | 516.6 | 317.1 | 1000 |
| | Talos | 220.1 | 186.3 | 134.2 | 10.5 | 10.3 | 212.5 | 836.7 | **840.5** | 1000 |
| | UnitreeH1 | 788.3 | 634.6 | 130.5 | 14.4 | 21.1 | 604.5 | 796.7 | **909.5** | 1000 |

### 5.4 Impact of Expert Data on Trainer-Student Dynamics

To answer the fourth question, we use MuJoCo's Humanoid environment (Todorov et al., 2012; Brockman et al., 2016) with a single expert trajectory from (Garg et al., 2021), varying the proportion of expert data in the replay buffers from 0% to 100% (e.g., 25% means a quarter expert data and 75% agent data; see Appendix D.4 for details).

Figure 5 shows that introducing expert data in both the trainer's and the student's replay buffers accelerates RILe's convergence (i.e., fewer training steps), but reduces the final performance. In the extreme case of 100% expert data, the student's performance drops substantially. This indicates that too much expert data hamper the trainer's ability to adapt to the student's real-time needs, disrupting RILe's context-sensitive reward customization. We also include IQ-Learn and BC results, both rely heavily on expert data, and find that neither matches RILe's performance, even when RILe uses a large portion of expert data.

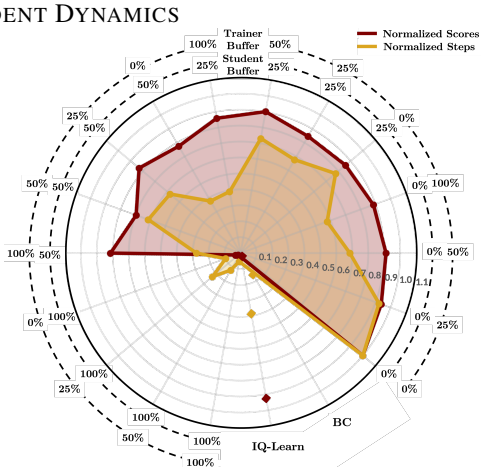


Figure 5: **Explicit Usage of Expert Data**. Red and yellow markers show normalized scores and steps, respectively. Expert data usage speeds the training of RILe but reduce final performance.

## 6 Discussion

As our experiments demonstrate, RILe consistently outperforms baseline models across various tasks, thanks to its adaptive learning approach, where the trainer agent continuously adjusts the reward based on the student's current learning stage.

Our maze experiments provide an interpretable example of how the trainer agent tailors its rewards. By encouraging actions that might seem suboptimal for immediate imitation but advantageous for long-term learning, RILe establishes a curriculum that ultimately boosts performance. This adaptive strategy helps RILe to achieve superior results in our continuous control experiments, where reward shaping becomes especially critical in high-dimensional settings.

Nonetheless, policy stability remains challenging when the reward function is constantly evolving. Freezing the trainer (see Appendix C) stabilizes learning but halts further adaptation, and the discriminator itself tends to overfit quickly. Future work could explore fully cooperative multi-agent RL techniques to allow ongoing adaptation, investigate ways to bound or regulate trainer updates, and consider discriminator-less formulations for reward learning.

Despite these challenges, RILe shows that *cooperatively learning* the policy and the reward function can offer significant advantages over static or iteratively updated methods. By providing dynamic and tailored rewards, RILe effectively guides the student through complex tasks. We believe this opens up new possibilities for responsive and adaptive learning frameworks in imitation learning and beyond.

## REFERENCES


Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, pp. 1, 2004.

Firas Al-Hafez, Guoping Zhao, Jan Peters, and Davide Tateo. Locomujoco: A comprehensive imitation learning benchmark for locomotion. In *6th Robot Learning Workshop, NeurIPS*, 2023.

Saurabh Arora and Prashant Doshi. A survey of inverse reinforcement learning: Challenges, methods and progress. *Artificial Intelligence*, 297:103500, 2021.

Michael Bain and Claude Sammut. A framework for behavioural cloning. In *Machine Intelligence 15*, pp. 103–129, 1995.

Michael Bloem and Nicholas Bambos. Infinite time horizon maximum causal entropy inverse reinforcement learning. *53rd IEEE Conference on Decision and Control*, pp. 4911–4916, 2014. URL `https://api.semanticscholar.org/CorpusID:14981371`.

Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym, 2016.

Annie S Chen, Suraj Nair, and Chelsea Finn. Learning generalizable robotic reward functions from" in-the-wild" human videos. In *Robotics: Science and Systems*, 2021.

Chelsea Finn, Sergey Levine, and Pieter Abbeel. Guided cost learning: Deep inverse optimal control via policy optimization. In *International conference on machine learning*, pp. 49–58. PMLR, 2016.

Justin Fu, Katie Luo, and Sergey Levine. Learning robust rewards with adverserial inverse reinforcement learning. In *International Conference on Learning Representations*, 2018.

Divyansh Garg, Shuvam Chakraborty, Chris Cundy, Jiaming Song, and Stefano Ermon. Iq-learn: Inverse soft-q learning for imitation. *Advances in Neural Information Processing Systems*, 34: 4028–4039, 2021.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

Todd Hester, Matej Vecerik, Olivier Pietquin, Marc Lanctot, Tom Schaul, Bilal Piot, Dan Horgan, John Quan, Andrew Sendonaris, Ian Osband, et al. Deep q-learning from demonstrations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. *Advances in neural information processing systems*, 29, 2016.

Ilya Kostrikov, Ofir Nachum, and Jonathan Tompson. Imitation learning via off-policy distribution matching. In *International Conference on Learning Representations*, 2020.

Chun-Mao Lai, Hsiang-Chun Wang, Ping-Chun Hsieh, Yu-Chiang Frank Wang, Min-Hung Chen, and Shao-Hua Sun. Diffusion-reward adversarial imitation learning. *arXiv preprint arXiv:2405.16194*, 2024.

Andrew Y Ng and Stuart J Russell. Algorithms for inverse reinforcement learning. In *Proceedings of the Seventeenth International Conference on Machine Learning*, pp. 663–670, 2000.

Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel Van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions On Graphics (TOG)*, 37(4):1–14, 2018.

Stéphane Ross and Drew Bagnell. Efficient reductions for imitation learning. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 661–668. JMLR Workshop and Conference Proceedings, 2010.

Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 627–635. JMLR Workshop and Conference Proceedings, 2011.

Voot Tangkaratt, Nontawat Charoenphakdee, and Masashi Sugiyama. Robust imitation learning from noisy demonstrations. In *International Conference on Artificial Intelligence and Statistics*, pp. 298–306. PMLR, 2021.

Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, pp. 5026–5033. IEEE, 2012.

Faraz Torabi, Garrett Warnell, and Peter Stone. Behavioral cloning from observation. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pp. 4950–4957, 2018a.

Faraz Torabi, Garrett Warnell, and Peter Stone. Generative adversarial imitation from observation. *arXiv preprint arXiv:1807.06158*, 2018b.

Yueh-Hua Wu, Nontawat Charoenphakdee, Han Bao, Voot Tangkaratt, and Masashi Sugiyama. Imitation learning from imperfect demonstration. In *International Conference on Machine Learning*, pp. 6818–6827. PMLR, 2019.

Yiqing Xu, Wei Gao, and David Hsu. Receding horizon inverse reinforcement learning. *Advances in Neural Information Processing Systems*, 35:27880–27892, 2022.

Kevin Zakka, Andy Zeng, Pete Florence, Jonathan Tompson, Jeannette Bohg, and Debidatta Dwibedi. Xirl: Cross-embodiment inverse reinforcement learning. In *Conference on Robot Learning*, pp. 537–546. PMLR, 2022.

Boyuan Zheng, Sunny Verma, Jianlong Zhou, Ivor W Tsang, and Fang Chen. Imitation learning: Progress, taxonomies and challenges. *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–16, 2022.

Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, Anind K Dey, et al. Maximum entropy inverse reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 8, pp. 1433–1438. Chicago, IL, USA, 2008.

# A POMDP OF THE TRAINER

Partially Observable Markov Decision Process (POMDP) of the trainer is defined as $\text{POMDP}_T = (S_T, A_T, \Omega_T, T_T, O_T, R_T, \gamma)$. Here, $T_T = \{P(. \mid f^T, a^T)\}$ is the transition dynamics where $P(. \mid f^T, a^T)$ is the state distribution upon taking action $a \in A_T$ in state $f \in S_T$. The transition function incorporates the student's policy $\pi_S$, which evolves in response to the rewards provided, reflecting the hidden dynamics due to the unobserved $\pi_S$. The observation function $O_T = \{P(s^T \mid f^T, a^T)\}$ defines the probability of observing $s^T \in \Omega_T$ given the state $(f^T, a^T)$. The trainer deterministically observes the student's state-action pair, so $P(s^T = (s^S, a^S) \mid f^T, a^T) = 1$, where $f^T = (s^S, a^S, \pi_S)$.

# B JUSTIFICATION OF RILE

**Assumptions:**

- The discriminator loss curve is complex and the discriminator function, $D_\phi(s, a)$, is sufficiently expressive since it is parameterized by a neural network with adequate capacity.
- For the trainer's and student's policy functions $(\pi^{\theta_T})$ and $(\pi^{\theta_S})$, and the Q-functions $(Q^{\theta_S})$, each is Lipschitz continuous with respect to its parameters with constants $(L_{\theta_T}), (L_{\theta_S}), and (L_Q)$, respectively. This means for all $(s, a)$ and for any pair of parameter settings $(\theta, \theta') : [|\pi^\theta(s, a) - \pi^{\theta'}(s, a)| \leq L_\theta |\theta - \theta'|, ][|Q^\theta(s, a) - Q^{\theta'}(s, a)| \leq L_Q |\theta - \theta'|.]$

To prove that the student agent can learn expert-like behavior, we need to show that the trainer agent learns to give higher rewards to student experiences that match with the expert state-action pair distribution, as this would enable a student policy to eventually mimic expert behavior.

## B.1 LEMMA 1:

Given the discriminator $D_\phi$, the trainer agent optimizes its policy $\pi^{\theta_T}$ via policy gradients to provide rewards that guide the student agent to match expert's state-action distributions.

**Proof for Lemma 1** The student agent, $\pi_S(a_t^S | s_t^S)$, interacts with the environment and generates state-action pairs as $(s_t^S, a_t^S)$. The trainer agent observes these pairs and provides a reward $r_t^S = a_t^T = \pi_T(a_t^T | (s_t^S, a_t^S))$ to the student, where $a_t^T \in [-1, 1]$ is the trainer's action. We have $D_\phi : \mathcal{S} \times \mathcal{A} \to [0, 1]$ as the discriminator, parameterized by $\phi$, which outputs the likelihood that a given state-action pair $(s, a)$ originates from the expert, as opposed to the student.

The trainer's reward at timestep $t$ is:

$$r_t^T = \upsilon(D_\phi(s_t^T))a_t^T \tag{11}$$

where $s_t^T = (s_t^S, a_t^S)$ is the trainer's observation, $D_\phi(s_t^T)$ is the discrimantor output that estimates the likelihood that $s_t^T$ comes from the expert data, and $\upsilon(D) = 2D - 1$ is a scaling function that maps discriminator's output to the range $[-1, 1]$.

The trainer maximizes the expected cumulative reward:

$$J_T(\pi_T) = \mathbb{E}_{\pi_T, \pi_S}\left[\sum_{t=0}^{\infty} \gamma^t r_t^T\right] \tag{12}$$

where $\gamma \in [0, 1)$ is the discount factor. In other words, trainer aims to find the policy that maximizes $J_T(\pi_T)$: $\pi^{*T} = \arg\max_{\pi^T} J_T(\pi_T)$.

From the policy gradient theorem, the gradient of the trainer's objective with respect to the policy parameters, $\theta_T$, is:

$$\nabla_{\theta_T} J_T(\pi_T) = \mathbb{E}_{\pi_T, \pi_S}\left[\nabla_{\theta_T} \log \pi_T(a_t^T | s_t^T) Q_T(s_t^T, a_t^T)\right] \tag{13}$$

where $Q_T(s_t^T, a_t^T)$ is the action-value function of the trainer. The action-value function, $Q_T(s_t^T, a_t^T)$, and the value function, $V_T(s_t^T)$ is defined by Bellman equation as:

$$Q_T(s_t^T, a_t^T) = r_t^T + \gamma \mathbb{E}_{s_{t+1}^T}\left[V_T(s_{t+1}^T)\right] \tag{14}$$

$$V_T(s_{t+1}^T) = \mathbb{E}_{a_t^T \sim \pi_T}\left[Q_T(s_t^T, a_t^T))\right] \tag{15}$$

The trainer aims to maximize $Q_T(s_t^T, a_t^T)$ to satisfy Equation 13. Since $r_t^T$ depends directly on $D_\phi(s_t^T)$ and $a_t^T$, the trainer learns to select $a_t^T$ that maximizes $Q_T(s_t^T, a_t^T)$ . Considering that $a_t^T \in [-1, 1]$, the immediate reward $r_t^T$ is maximized when $a_t^T$ has the same sign as $\upsilon(D_\phi(s_t^T))$ . Therefore, the optimal action $a_t^{*T}$ is:

$$\alpha_t^{*T} = \begin{cases} 1, & \text{if } D\phi(s_t^T) > 0.5, \\ -1, & \text{if } D\phi(s_t^T) < 0.5, \\ \text{any value in } [-1, 1], & \text{if } D\phi(s_t^T) = 0.5. \end{cases} \tag{16}$$

Equation 16 implies the trainer assigns positive rewards to student state-action pairs that the discriminator assesses as more likely to be from the expert ($D_\phi(s_t^T) > 0.5$) and negative rewards to those unlikely to be from the expert ($D_\phi(s_t^T) < 0.5$). By this mechanism, the trainer's policy optimization relies on the discriminator's assessment to assign rewards that encourage expert-like behavior.

All in all, the derivative of the trainer's expected reward, Equation 13, with respect to its policy parameters is rewritten as:

$$\nabla_{\theta_T} J_T(\pi_T) = \mathbb{E}_{\pi_T, \pi_S} \left[ \nabla_{\theta_T} \log \pi_T(a_t^T | s_t^T) \left( (2D_\phi(s_t^T) - 1) a_t^T + \gamma Q_T(s_{t+1}^T, a_{t+1}^T) \right) \right] \tag{17}$$

The trainer adjusts $\pi_T$ to output high rewards when $D_\phi(s_t^T)$ is high. Therefore the trainer learns to assign higher rewards to student behaviors that are more similar to expert behaviors, according to the discriminator.

### B.2 Lemma 2:

The discriminator $D_\phi$, parameterized by $\phi$ will converge to a function that estimates the probability of a state-action pair being generated by the expert policy, when trained on samples generated by both a student policy $\pi^{\theta_S}$ and an expert policy $\pi_E$.

**Proof for Lemma 2**: The discriminator's objective is to distinguish between state-action pairs generated by the expert and those generated by the student. The training objective for the discriminator is framed as a binary classification problem over expert demonstrations and student-generated trajectories. The discriminator's loss function $\mathcal{L}_D(\phi)$ is the binary cross-entropy loss, which is defined as:

$$L_D(\phi) = -\mathbb{E}_{(s,a)\sim p_E}[\log(D_\phi(s,a))] - \mathbb{E}_{(s,a)\sim p_{\pi_S}}[\log(1 - D_\phi(s,a))]. \tag{18}$$

where $p_E(s,a)$ is the state-action distribution of the expert policy, and $p_{\pi_S}(s,a)$ is the state-action distribution of the student agent. Considering that $x = (s,a)$, this loss can be rewritten as:

$$L_D(\phi) = -\int [p_E(s,a) \log D_\phi(s,a) + p_{\pi_S}(s,a) \log(1 - D_\phi(s,a))] \, ds \, da \tag{19}$$

$$L_D(\phi) = -\int [p_E(x) \log D_\phi(x) + p_{\pi_S}(x) \log(1 - D_\phi(x))] \, dx \, . \tag{20}$$

As presented in Goodfellow et al. (2014), the optimal discriminator that minimizes this loss, $D_\phi^*$, is:

$$D_\phi^*(x) = \frac{p_E(x)}{p_E(x) + p_{\pi_S}(x)}, \tag{21}$$

$$D_\phi^*(s,a) = \frac{p_E(s,a)}{p_E(s,a) + p_{\pi_S}(s,a)}. \tag{22}$$

This shows that the optimal discriminator estimates the probability that a state-action pair comes from the expert policy, normalized by the total probability from both expert and student policies.

## C Training Strategies

The introduction of the trainer agent into the AIL framework introduces instabilities that can hinder the learning process. To address these challenges, we employ three strategies.

**Freezing the Trainer Agent Midway:** Continuing to train the trainer agent throughout the entire process can lead to overfitting on minor fluctuations in the student's behavior. This overfitting causes the trainer to assign inappropriate negative rewards, which diverts the student away from expert behavior—especially since the student agent may fail to interpret these subtle nuances correctly in the later stages of training. To prevent this, we freeze the trainer agent once its critic network within the actor-critic framework converges during the training process.

We consider the trainer's critic network to have converged when the change in the exponential moving average (with a smoothing factor of 0.99) of the critic output and its variance over a window of 50000 training iterations fall below a certain threshold. In all our experiments, this threshold is set to 0.1, which we found empirically after our hyperparameter search (see Appendix H). This threshold works for all settings where the reward is bounded between $-1$ and $1$, which is the case for all our experiments.

**Reducing the Trainer Update Frequency:** We decrease the update frequency of the trainer agent to one forth of the student agent. We empirically found that updating at one fourth of the student agent's frequency works best. This adjustment aims to prevent overestimation bias in the trainer's value function and to slow down its learning pace. By updating less frequently, the trainer provides more consistent and reliable reward signals. This steadier guidance helps the student agent better understand and adapt to the trainer's rewards, facilitating more stable learning.

**Increasing the Student Agent's Exploration:** We increase the exploration rate of the student agent compared to standard AIL methods. We implement an epsilon-greedy strategy within the actor-critic framework, allowing the student to occasionally take random actions. This increased exploration enables the student to visit a wider range of state-action pairs. Consequently, the trainer agent receives diverse input, helping it learn a more effective reward function. This diversity is crucial for the trainer to observe the outcomes of various actions and to guide the student more effectively toward expert behavior.

# D Experimental Settings

## D.1 Evolving Reward Function

We use single expert demonstration in this experiment. For RILe, we plot the reward function learned by the trainer. For GAIL, we visualize the discriminator output, and for AIRL, the reward term under the discriminator.

## D.2 Reward Function Dynamics

In this experiment, we select the student agent's hyperparameters to be identical to those used in GAIL, ensuring that the only difference between the agents is the reward function. Therefore, we use the best hyperparameters identified for GAIL, applied to both GAIL and RILe, from our hyperparameter sweeps presented in Appendix H.

**RFDC:** We calculate the Wasserstein distance between reward distributions over consecutive 10,000-step training intervals, denoted as times $t$ and $t + 10,000$. This metric quantifies how much the overall reward distribution shifts over time. Changes in reward distributions depend both on the reward function and the student policy updates. Since we use the same student agent with the same hyperparameters, higher RFDC values still indicate that the reward function is adapting more dynamically in response to the student's learning progress.

**FS-RFDC:** We compute the mean absolute deviation of rewards between consecutive 10,000-step training intervals for a fixed set of states derived from expert data. As the fixed set, we use all the states in the expert data. Since the states used for calculating rewards are fixed, changes in this value purely depend on the reward function updates. This metric assesses how the reward values for specific states change over time.

**CPR:** We evaluate how changes in the reward function correlate with improvements in student performance. We store rewards from both the learned reward function and the environment-defined rewards in separate buffers. In other words, we collect samples from two reward functions: the learned reward function and the environment-defined reward function. The environment rewards

consider the agent's velocity and stability. Every 10,000 steps, we calculate the Pearson correlation between these rewards and empty the buffers. This metric evaluates whether increases in the learned rewards relate to performance enhancements.

### D.3 Motion-Capture Data Imitation for Robotic Continuous Control

During training, we use 8 different random seeds and 8 distinct initial positions for the robot. The validation setting mirrors the training conditions: we sample initial positions from the same set of 8 possibilities and use the same random seeds. In this setting, the student agent selects actions deterministically, allowing us to assess its performance under familiar conditions.

For the test setting, we evaluate the policy's ability to generalize to new, unseen scenarios. We modify the initial positions of the robot by randomly initializing it in stable configurations not included in the fixed set used during training. Additionally, we use different random seeds from those in training, introducing new random variations that affect the environment's dynamics during state transitions. This setup enables us to assess how well the learned policy performs when faced with novel initial conditions.

### D.4 Impact of Expert Data on Trainer-Student Dynamics

In this experiment, both seeds and initial positions in the test setting are different from the training one, and we report values from the test setting.

For every percentage of the expert-data in buffers, we continue trainings of both the trainer agent and the student agent of RILe. For instance, in 100% expert data in the trainer's buffer case, both the student and the discriminator are trained normally using samples from the student agent. However, we didn't include student's state-action pairs to the trainer's buffer, instead, we filled that buffer with a batch of expert data, and updated the trainer regularly using this modified buffer. Similarly, in 100% expert data in the student's buffer case , we trained the trainer agent and the discriminator normally, using samples from the student. However, student's state-actions pairs are not included in the student's buffer, and student agent is updated just by using expert state-action pairs, using rewards coming from the trainer agent for these expert pairs.

Regarding the normalizations, we trained Behavioral Cloning (BC) and RILe across various data leakage levels, selecting the highest-scoring run (0% leakage RILe) as the baseline. Other scores and convergence steps are normalized by dividing by the score and convergence steps of the baseline (0% leakage RILe). For IQLearn, we used their reported numbers in their paper, as we couldn't replicate their results with their code and hyperparameters.

## E Additional Experiments

### E.1 Robustness to Noise in the Expert Data

To evaluate the robustness of RILe and baseline methods to noise in the expert data, we conducted experiments in the MuJoCo Humanoid-v2 environment. Artificial noise sampled from a zero-mean Gaussian distribution with varying standard deviations ($\Sigma$) was added to a single expert trajectory, affecting either the actions or the states. The baselines used for comparison were GAIL (Ho & Ermon, 2016), AIRL (Fu et al., 2018), RIL-Co (Tangkaratt et al., 2021), IC-GAIL (Wu et al., 2019), and IQ-Learn (Garg et al., 2021).

As shown in Table 2, RILe consistently outperforms the baselines across different noise levels, demonstrating superior robustness even when a high amount of noise is present in the expert data ($\Sigma = 0.5$). These results indicate that RILe is less sensitive to imperfections in the expert demonstrations compared to existing methods.

### E.2 Robustness of the Learned Reward Function

We evaluated the robustness of the reward functions learned by RILe and AIRL (Fu et al., 2018) through an experiment similar to that conducted by Xu et al. (2022). Initially, both methods were trained to learn reward functions in a noise-free MuJoCo Humanoid-v2 environment. After training,

Table 2: Test results in MuJoCo Humanoid-v2 environment, where artificial noise sampled from a zero-mean Gaussian distribution is added to a single expert trajectory. Results are aggregated over 20 different-seed environments. IQ-Learn* is trained using the official code and hyperparameters of the IQ-Learn algorithm.

| | Noise-Free | Action Noise | | State Noise | |
|---|---|---|---|---|---|
| | $\Sigma = 0$ | $\Sigma = 0.2$ | $\Sigma = 0.5$ | $\Sigma = 0.2$ | $\Sigma = 0.5$ |
| RILe | **5681** | **5280** | **5154** | **5350** | **5205** |
| GAIL | 5430 | 5275 | 902 | 5147 | 917 |
| AIRL | 5276 | 4869 | 4589 | 4898 | 4780 |
| RIL-Co | 576 | 491 | 493 | 505 | 501 |
| IC-GAIL | 610 | 601 | 568 | 590 | 591 |
| IQ-Learn* | 312 | 192 | 153 | 243 | 277 |

these reward functions were frozen. Subsequently, new student agents were trained using these fixed reward functions in environments where Gaussian noise was added to the agents' actions, with varying noise levels.

Table 3 presents the results of this evaluation. The reward function learned by RILe demonstrates superior robustness to noise, maintaining high performance even under increased noise levels. In contrast, the performance of agents using the reward function learned by AIRL decreases more significantly as noise increases. These findings indicate that the reward function learned by RILe is more resilient to environmental noise, contributing to better agent performance in noisy conditions.

Table 3: We test the robustness of learned reward functions. After training reward functions in a noise-free setting, reward functions are frozen, and used to train a new agent in a noisy environment, where Gaussian noise is added to agent's actions in every step.

| | No Noise | Mild Noise | High Noise |
|---|---|---|---|
| | $\Sigma = 0$ | $\Sigma = 0.2$ | $\Sigma = 0.5$ |
| RILe | **5748** | **5201** | **5196** |
| AIRL | 5334 | 5005 | 4967 |

## F EXTENDED MUJOCO RESULTS

We present MuJoCo results for the test setting, with standard errors, in Table 4.

Table 4: Test results on four MuJoCo tasks with standard errors.

| | RILe | GAIL | AIRL | IQLearn | DRAIL |
|---|---|---|---|---|---|
| Humanoid-v2 | **5928** $\pm$ **188** | 5709 $\pm$ 63 | 5623 $\pm$ 252 | 327 $\pm$ 105 | 5755 $\pm$ 34 |
| Walker2d-v2 | 4435 $\pm$ 206 | **4906** $\pm$ **159** | 4823 $\pm$ 221 | 270 $\pm$ 43 | 4016 $\pm$ 127 |
| Hopper-v2 | **3417** $\pm$ **155** | 3361 $\pm$ 51 | 3014 $\pm$ 190 | 310 $\pm$ 47 | 1230 $\pm$ 73 |
| HalfCheetah-v2 | **5205** $\pm$ **31** | 4173 $\pm$ 94 | 3991 $\pm$ 126 | 755 $\pm$ 211 | 4133 $\pm$ 41 |

## G EXTENDED LOCOMUJOCO RESULTS

We present LocoMujoco results for the validation setting and test setting, with standard errors, in Table 5 and 6, respectively.

Table 5: Validation results on seven LocoMujoco tasks.

| | | RILe | GAIL | AIRL | IQ | BCO | GAIfO | DRAIL GAIL | DRAIL RILe | Expert |
|---|---|---|---|---|---|---|---|---|---|---|
| Walk | Atlas | 895.4 ±25 | 918.6 ±133 | 356.0 ±68 | 32.1 ±4 | 28.7 ±4 | 831.6 ±41 | 741.3 ±46 | 773.9 ±13 | 1000 |
| | Talos | 884.7 ±8 | 675.5 ±105 | 103.4 ±22 | 7.2 ±2 | 19.9 ±4 | 718.8 ±16 | 963.7 ±48 | 949.4 ±54 | 1000 |
| | UnitreeH1 | 980.7 ±15 | 965.1 ±20 | 716.2 ±124 | 12.5 ±6 | 43.7 ±8.4 | 586.6 ±102 | 954.7 ±20 | 973.5 ±8 | 1000 |
| | Humanoid | 970.3 ±101 | 216.2 ±18 | 78.2 ±6 | 6.8 ±1 | 8.3 ±1 | 345.7 ±34 | 550.8 ±148 | 595.3 ±73 | 1000 |
| Carry | Atlas | 889.7 ±44 | 974.2 ±80 | 271.9 ±30 | 39.5 ±8 | 42.7 ±9 | 306.2 ±9 | 654.1 ±109 | 344.1 ±28 | 1000 |
| | Talos | 503.3 ±72 | 338.5 ±48 | 74.1 ±8 | 11.7 ±3 | 8.1 ±1 | 444.5 ±96 | 889.8 ±163 | 874.3 ±174 | 1000 |
| | UnitreeH1 | 850.6 ±80 | 637.4 ±90 | 140.9 ±21 | 12.3 ±2 | 30.2 ±5 | 503.6 ±55 | 620.8 ±60 | 878.1 ±46 | 1000 |

Table 6: Test results on seven LocoMujoco tasks.

| | | RILe | GAIL | AIRL | IQ | BCO | GAIfO | DRAIL GAIL | DRAIL RILe | Expert |
|---|---|---|---|---|---|---|---|---|---|---|
| Walk | Atlas | 870.6 ±13 | 792.7 ±105 | 300.5 ±74 | 30.9 ±10 | 21.0 ±3 | 803.1 ±68 | 834.4 ±23 | **899.1** ±17 | 1000 |
| | Talos | 842.5 ±24 | 442.3 ±76 | 102.1 ±17 | 4.5 ±3 | 11.9 ±1 | 687.2 ±44 | 787.7 ±11 | **896.6** ±12 | 1000 |
| | UnitreeH1 | 966.2 ±14 | 950.2 ±13 | 568.1 ±156 | 8.8 ±3 | 34.8 ±10 | 526.8 ±72 | 940.8 ±20 | **995.8** ±6 | 1000 |
| | Humanoid | **831.3** ±98 | 181.4 ±24 | 80.1 ±9 | 4.5 ±2 | 3.5 ±2 | 292.1 ±25 | 814.6 ±80 | 527.6 ±39 | 1000 |
| Carry | Atlas | **850.8** ±62 | 669.3 ±55 | 256.4 ±47 | 36.8 ±14 | 20.3 ±1 | 402.9 ±39 | 516.6 ±60 | 317.1 ±19 | 1000 |
| | Talos | 220.1 ±88 | 186.3 ±28 | 134.2 ±18 | 10.5 ±3 | 10.3 ±2 | 212.5 ±32 | 836.7 ±160 | **840.5** ±133 | 1000 |
| | UnitreeH1 | 788.3 ±71 | 634.6 ±45 | 130.5 ±22 | 14.4 ±2 | 21.1 ±6 | 504.5 ±30 | 796.7 ±131 | **909.5** ±9 | 1000 |

## H HYPERPARAMETERS

We present hyperparameters in Table 7. For DRAIL, we replaced the discriminators with the implementation provided by DRAIL and adopted their hyperparameters for the HandRotate task.

Our experiments revealed that RILe's performance is particularly sensitive to certain hyperparameters. We highlight three key observations:

- RILe is more sensitive to the hyperparameters of the discriminator compared to other methods. Specifically, increasing the discriminator's capacity or training speed, by using a larger network architecture or increasing the number of updates per iteration, adversely affects RILe's performance. A powerful discriminator tends to overfit quickly to the expert data, resulting in high confidence when distinguishing between expert and student behaviors. This poses challenges for the trainer agent, as the discriminator's feedback becomes less informative.
- The update frequency of the trainer agent's target network influences the stability of the RILe framework. Lower update frequencies lead to improved stability. A slower-updating trainer provides more consistent reward signals, allowing the student agent to better adapt to the rewards. However, a lower update frequency slows down the learning process, as the trainer adapts more slowly to changes in the student's behavior. Therefore, there is a trade-off between stability and learning speed that needs to be balanced.
- Enhancing the exploration rate of the student agent benefits RILe more than it does baseline methods. By encouraging the student to explore more, through strategies like higher entropy regularization or implementing an epsilon-greedy policy, the student visits a broader range of state-action pairs. This increased diversity provides the trainer agent with more varied data, enabling it to learn a more effective and robust reward function. The additional exploration helps the trainer to better capture the effects of different actions.

## I COMPUTE RESOURCES

For the training of RILe and baselines, following computational sources are employed:

- AMD EPYC 7742 64-Core Processor
- 1 x Nvidia A100 GPU
- 32GB Memory

Table 7: Hyperparameter Sweeps and Best Hyperparameters for LocoMujoco and Humanoid Experiments

| | **Hyperparameters** | **RILe** | **GAIL** | **AIRL** | **IQ-Learn** |
|---|---|---|---|---|---|
| **Discriminator** | Updates per Round | **1**, 2, 8 | **1**, 2, 8 | **1**, 2, 8 | - |
| | Batch Size | **32**, 64, 128 | **32**, 64, 128 | **32**, 64, 128 | - |
| | Buffer Size | 8192, **16384**, 1e5 | 8192, **16384**, 1e5 | 8192, **16384**, 1e5 | - |
| | Network | [512FC, 512FC]<br>[256FC, 256FC]<br>**[64FC, 64FC]** | [512FC, 512FC]<br>[256FC, 256FC]<br>**[64FC, 64FC]** | [512FC, 512FC]<br>[256FC, 256FC]<br>**[64FC, 64FC]** | - |
| | Gradient Penalty | 0.5, **1** | 0.5, **1** | 0.5, **1** | - |
| | Learning Rate | 3e-4, 1e-4, **3e-5**, 1e-5 | 3e-4, 1e-4, **3e-5**, 1e-5 | 3e-4, 1e-4, **3e-5**, 1e-5 | - |
| **Student** | Buffer Size | 1e5, **1e6** | 1e5, **1e6** | 1e5, **1e6** | 1e5, **1e6** |
| | Batch Size | 32, **256** | 32, **256** | 32, **256** | 32, **256** |
| | Network | **[256FC, 256FC]** | **[256FC, 256FC]** | **[256FC, 256FC]** | **[256FC, 256FC]** |
| | Activation Function | **ReLU**, Tanh | **ReLU**, Tanh | **ReLU**, Tanh | **ReLU**, Tanh |
| | Discount Factor ($\gamma$) | **0.99**, 0.97, 0.95 | **0.99**, 0.97, 0.95 | **0.99**, 0.97, 0.95 | **0.99**, 0.97, 0.95 |
| | Learning Rate | **3e-4**, 1e-4, 3e-5, 1e-5 | **3e-4**, 1e-4, 3e-5, 1e-5 | **3e-4**, 1e-4, 3e-5, 1e-5 | **3e-4**, 1e-4, 3e-5, 1e-5 |
| | Tau ($\tau$) | 0.05, **0.01**, 0.005 | 0.05, **0.01**, 0.005 | 0.05, **0.01**, 0.005 | 0.05, **0.01**, 0.005 |
| | Epsilon-greedy | 0, 0.1, **0.2** | **0**, 0.1, 0.2 | **0**, 0.1, 0.2 | **0**, 0.1, 0.2 |
| | Entropy | **0.2**, 0.5, 1 | **0.2**, 0.5, 1 | **0.2**, 0.5, 1 | 0.05, 0.1, **0.2**, 0.5, 1 |
| **Trainer** | Buffer Size | 8192, **16384**, 1e5, 1e6 | - | - | - |
| | Batch Size | 32, **256** | - | - | - |
| | Network | **[256FC, 256FC]**<br>[64FC, 64FC] | - | - | - |
| | Activation Function | **ReLU**, Tanh | - | - | - |
| | Discount Factor ($\gamma$) | **0.99**, 0.97, 0.95 | - | - | - |
| | Learning Rate | **3e-4**, 1e-4, 3e-5, 1e-5 | - | - | - |
| | Tau ($\tau$) | 0.05, 0.01, **0.005** | - | - | - |
| | Entropy | **0.2**, 0.5, 1 | - | - | - |
| | Freeze Threshold | 1, 0.5, **0.1**, 0.01, 0.001 | - | - | - |

## J ALGORITHM

**Algorithm 1** RILe Training Process

Initialize student policy $\pi_S$ and trainer policy $\pi_T$ with random weights, and the discriminator $D$ with random weights.
Initialize an empty replay buffer $B$
**for** each iteration **do**
    Sample trajectory $\tau_S$ using current student policy $\pi_S$
    Store $\tau_S$ in replay buffer $B$
    **for** each transition $(s, a)$ in $\tau_S$ **do**
        Calculate student reward $R^S$ using trainer policy:

$$R^S \leftarrow \pi_T \tag{23}$$

        Update $\pi_S$ using policy gradient with reward $R^S$
    **end for**
    Sample a batch of transitions from $B$
    Train discriminator $D$ to classify student and expert transitions

$$\max_D E_{\pi_S}[\log(D(s,a))] + E_{\pi_E}[\log(1 - D(s,a))] \tag{24}$$

    **for** each transition $(s, a)$ in $\tau_S$ **do**
        Calculate trainer reward $R^T$ using discriminator:

$$R^T \leftarrow \upsilon(D(s,a))a^T \tag{25}$$

        Update $\pi_T$ using policy gradient with reward $R^T$
    **end for**
**end for**

**Algorithm 2** RILe Training Process with Off-policy RL

Initialize student policy $\pi_S$, trainer policy $\pi_T$, and the discriminator $D$ with random weights.
Initialize an empty replay buffers $B_D$, $B_S$, $B_T$ with different sizes
**for** each iteration **do**
    Sample trajectory $\tau_S$ using current student policy $\pi_S$
    Store $\tau_S$ in replay buffers $B_D$, $B_S$, $B_T$
    Sample a batch of transitions, $b_S$ from $B_S$
    **for** each transition $(s, a)$ in $b_S$ **do**
        Calculate student reward $R^S$ using trainer policy:

$$R^S \leftarrow \pi_T \tag{26}$$

        Update $\pi_S$ using calculated rewards
    **end for**
    Sample a batch of transitions $b_D$ from $B_D$
    Train discriminator $D$ to classify student and expert transitions

$$\max_D E_{\pi_S}[\log(D(s,a))] + E_{\pi_E}[\log(1 - D(s,a))] \tag{27}$$

    Sample a batch of transitions, $b_T$ from $B_T$
    **for** each transition $(s, a)$ in $b_T$ **do**
        Calculate trainer reward $R^T$ using discriminator:

$$R^T \leftarrow \upsilon(D(s,a))a^T \tag{28}$$

        Update $\pi_T$ using calculated rewards
    **end for**
**end for**