

---

# Robust Imitation Learning from Noisy Demonstrations

---

Voot Tangkaratt  
RIKEN  
voot.tangkaratt@riken.jp

Nontawat Charoenphakdee  
University of Tokyo & RIKEN  
nontawat@ms.k.u-tokyo.ac.jp

Masashi Sugiyama  
RIKEN & University of Tokyo  
sugi@k.u-tokyo.ac.jp

## Abstract

Robust learning from noisy demonstrations is a practical but highly challenging problem in imitation learning. In this paper, we first theoretically show that robust imitation learning can be achieved by optimizing a classification risk with a symmetric loss. Based on this theoretical finding, we then propose a new imitation learning method that optimizes the classification risk by effectively combining pseudo-labeling with co-training. Unlike existing methods, our method does not require additional labels or strict assumptions about noise distributions. Experimental results on continuous-control benchmarks show that our method is more robust compared to state-of-the-art methods.

## 1 INTRODUCTION

The goal of sequential decision making is to learn a good policy that makes good decisions (Puterman, 1994). Imitation learning (IL) is an approach that learns a policy from demonstrations (i.e., sequences of demonstrators’ decisions) (Schaal, 1999). Researchers have shown that a good policy can be learned efficiently from high-quality demonstrations collected from experts (Ng and Russell, 2000; Ho and Ermon, 2016). However, demonstrations in the real-world often have lower quality due to noise or insufficient expertise of demonstrators, especially when humans are involved in the data collection process (Mandlekar et al., 2018). This is problematic because low-quality demonstrations can reduce the efficiency of IL both in theory and practice (Tangkaratt et al., 2020). In this paper, we theoretically and experimentally show that IL can perform well even in the presence of noises.

In the literature, methods for IL from noisy demonstrations have been proposed, but they still have limitations as they require additional labels from experts or a strict assumption about noise (Brown et al., 2019, 2020; Wu et al., 2019; Tangkaratt et al., 2020). Specifically, methods of Brown et al. (2019, 2020) require noisy demonstrations to be ranked according to their relative performance. Meanwhile, methods of Wu et al. (2019) require some of noisy demonstrations to be labeled with a score determining the probability that demonstrations are collected from experts. On the other hand, the method of Tangkaratt et al. (2020) does not require these labels, but instead it assumes that noisy demonstrations are generated by Gaussian noise distributions. To sum up, these methods require either additional labels from experts or a strict assumption about noise distributions. Due to this, the practicality of these methods is still limited, and IL from noisy demonstrations is still highly challenging.

To overcome the above limitation, we propose a new method for IL from noisy demonstrations called *Robust IL with Co-pseudo-labeling* (RIL-Co). Briefly, we built upon the recent theoretical results of robust classification (Charoenphakdee et al., 2019), and prove that robust IL can be achieved by optimizing a classification risk with a *symmetric loss*. However, optimizing the proposed risk is not trivial because it contains a data density whose data samples are not observed. We show that pseudo-labeling (Chapelle et al., 2010) can be utilized to estimate the data density. However, naive pseudo-labeling may suffer from over-fitting and is not suitable in practice (Kingma et al., 2014). To remedy this issue, we propose *co-pseudo-labeling*, which effectively combines pseudo-labeling with co-training (Blum and Mitchell, 1998). Compare to prior work, RIL-Co does not require additional labels or assumptions about noise distributions. In addition, RIL-Co does not require an additional hyper-parameter tuning because an appropriate hyper-parameter value can be derived from the theory. Experiments on continuous-control benchmarks show that RIL-Co is more robust against noisy demonstrations when compared to state-of-the-art methods.

## 2 IMITATION LEARNING AND ROBUSTNESS

In this section, we firstly give backgrounds about reinforcement learning and imitation learning. Then, we describe the setting of imitation learning from noisy demonstrations. Lastly, we discuss the robustness of existing imitation learning methods.

### 2.1 Reinforcement Learning

Reinforcement learning (RL) aims to learn an optimal policy of a Markov decision process (MDP) (Puterman, 1994). We consider a discrete-time MDP denoted by  $\mathcal{M} = (\mathcal{S}, \mathcal{A}, p_{\mathbf{T}}(\mathbf{s}'|\mathbf{s}, \mathbf{a}), p_1(\mathbf{s}_1), r(\mathbf{s}, \mathbf{a}), \gamma)$  with state  $\mathbf{s} \in \mathcal{S}$ , action  $\mathbf{a} \in \mathcal{A}$ , transition probability density  $p_{\mathbf{T}}$ , initial state probability density  $p_1$ , reward function  $r$ , and discount factor  $0 < \gamma \leq 1$ . A policy function  $\pi(\mathbf{a}|\mathbf{s})$  determines the conditional probability density of an action in a state. An agent acts in an MDP by observing a state, choosing an action according to a policy, transiting to a next state according to the transition probability density, and possibly receiving an immediate reward according to the reward function. An optimal policy of an MDP is a policy maximizing the expected cumulative discounted rewards.

Formally, RL seeks for an optimal policy by solving the optimization problem  $\max_{\pi} \mathcal{J}(\pi)$ , where  $\mathcal{J}(\pi)$  is the expected cumulative discounted rewards defined as

$$\begin{aligned} \mathcal{J}(\pi) &= \mathbb{E}_{p_{\pi}} \left[ \sum_{t=1}^T \gamma^{t-1} r(\mathbf{s}_t, \mathbf{a}_t) \right] \\ &= \mathbb{E}_{\rho_{\pi}} [r(\mathbf{s}_t, \mathbf{a}_t)] / (1 - \gamma). \end{aligned} \quad (1)$$

Here,  $p_{\pi}(\tau) = p_1(\mathbf{s}_1) \prod_{t=1}^T p_{\mathbf{T}}(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t) \pi(\mathbf{a}_t|\mathbf{s}_t)$  is the probability density of trajectory  $\tau = (\mathbf{s}_1, \mathbf{a}_1, \dots, \mathbf{s}_{T+1})$  with length  $T$ ,  $\rho_{\pi}(\mathbf{s}, \mathbf{a}) = (1 - \gamma) \mathbb{E}_{p_{\pi}(\tau)} [\sum_{t=1}^T \gamma^{t-1} \delta(\mathbf{s}_t - \mathbf{s}, \mathbf{a}_t - \mathbf{a})]$  is the state-action density determining the probability density of the agent with  $\pi$  observing  $\mathbf{s}$  and executing  $\mathbf{a}$ , and  $\delta(\mathbf{s}_t - \mathbf{s}, \mathbf{a}_t - \mathbf{a})$  is the Dirac delta function. Note that the state-action density is a normalized occupancy measure and uniquely corresponds to a policy by a relation  $\pi(\mathbf{a}|\mathbf{s}) = \rho_{\pi}(\mathbf{s}, \mathbf{a}) / \rho_{\pi}(\mathbf{s})$ , where  $\rho_{\pi}(\mathbf{s}) = \int_{\mathcal{A}} \rho_{\pi}(\mathbf{s}, \mathbf{a}) d\mathbf{a}$  (Syed et al., 2008). Also note that the optimal policy is not necessarily unique since there can be different policies that achieve the same expected cumulative discounted rewards.

While RL has achieved impressive performance in recent years (Silver et al., 2017), its major limitation is that it requires a suitable reward function which may be unavailable in practice (Schaal, 1999).

### 2.2 Imitation Learning

Imitation learning (IL) is a well-known approach to learn an optimal policy when the reward function is

unavailable (Ng and Russell, 2000). Instead of learning from the reward function or reward values, IL methods learn an optimal policy from demonstrations that contain information about an optimal policy. IL methods typically assume that demonstrations (i.e., a dataset of state-action samples) are collected by using an expert policy that is similar to an optimal (or near optimal) policy, and they aim to learn the expert policy (Ng and Russell, 2000; Ziebart et al., 2010; Sun et al., 2019). More formally, the typical goal of IL is to learn an expert policy  $\pi_{\mathbf{E}}$  by using a dataset of state-action samples drawn from an expert state-action density:

$$\{(\mathbf{s}_n, \mathbf{a}_n)\}_{n=1}^N \stackrel{\text{i.i.d.}}{\sim} \rho_{\mathbf{E}}(\mathbf{s}, \mathbf{a}), \quad (2)$$

where  $\rho_{\mathbf{E}}$  is a state-action density of expert policy  $\pi_{\mathbf{E}}$ .

The density matching approach was shown to be effective in learning the expert policy from expert demonstrations (Syed et al., 2008; Ho and Ermon, 2016; Ghasemipour et al., 2020). Briefly, this approach seeks for a policy  $\pi$  that minimizes a divergence between the state-action densities of the expert and learning policies. Formally, this approach aims to solve the following optimization problem:

$$\min_{\pi} D(\rho_{\mathbf{E}}(\mathbf{s}, \mathbf{a}) || \rho_{\pi}(\mathbf{s}, \mathbf{a})), \quad (3)$$

where  $D$  is a divergence such as the Jensen-Shannon divergence<sup>1</sup>. In practice, the divergence, which contains unknown state-action densities, is estimated by using demonstrations and trajectories drawn from  $\rho_{\mathbf{E}}$  and  $\rho_{\pi}$ , respectively. A well-known density matching method is generative adversarial IL (GAIL) (Ho and Ermon, 2016), which minimizes an estimate of the Jensen-Shannon divergence by solving

$$\begin{aligned} \min_{\pi} \max_g \mathbb{E}_{\rho_{\mathbf{E}}} \left[ \log \left( \frac{1}{1 + \exp(-g(\mathbf{s}, \mathbf{a}))} \right) \right] \\ + \mathbb{E}_{\rho_{\pi}} \left[ \log \left( \frac{1}{1 + \exp(g(\mathbf{s}, \mathbf{a}))} \right) \right], \end{aligned} \quad (4)$$

where  $g : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$  is a discriminator function.

Density matching methods were shown to scale well to high-dimensional problems when combined with deep neural networks (Ho and Ermon, 2016; Ghasemipour et al., 2020). However, an issue of this approach is that it is not robust against noisy demonstrations in general, as will be described in Section 2.4.

### 2.3 Learning from Noisy Demonstrations

In this paper, we consider a scenario of *IL from noisy demonstrations*, where given demonstrations are

<sup>1</sup>For non-symmetric divergence, an optimization problem  $\min_{\pi} D(\rho_{\pi}(\mathbf{s}, \mathbf{a}) || \rho_{\mathbf{E}}(\mathbf{s}, \mathbf{a}))$  can be considered as well (Ghasemipour et al., 2020).

a mixture of expert and non-expert demonstrations. We assume that we are given a dataset of state-action samples drawn from a noisy state-action density:

$$\mathcal{D} = \{(\mathbf{s}_n, \mathbf{a}_n)\}_{n=1}^N \stackrel{\text{i.i.d.}}{\sim} \rho'(\mathbf{s}, \mathbf{a}), \quad (5)$$

where the noisy state-action density  $\rho'$  is a mixture of the expert and non-expert state-action densities:

$$\rho'(\mathbf{s}, \mathbf{a}) = \alpha \rho_E(\mathbf{s}, \mathbf{a}) + (1 - \alpha) \rho_N(\mathbf{s}, \mathbf{a}). \quad (6)$$

Here,  $0.5 < \alpha < 1$  is an unknown mixing coefficient and  $\rho_N$  is the state-action density of a non-expert policy  $\pi_N$ . The policy  $\pi_N$  is non-expert in the sense that  $\mathbb{E}_{\rho_N}[r(\mathbf{s}, \mathbf{a})] < \mathbb{E}_{\rho_E}[r(\mathbf{s}, \mathbf{a})]$ , where  $r$  is an unknown reward function of the MDP. Our goal is to learn the expert policy using the dataset in Eq. (5).

We emphasize that  $0.5 < \alpha < 1$  corresponds to an assumption that the majority of demonstrations are obtained by the expert policy. This is a typical assumption when learning from noisy data, i.e., the number of good quality samples should be more than that of low quality samples (Angluin and Laird, 1988; Natarajan et al., 2013). For notational brevity, we denote a state-action pair by  $\mathbf{x} = (\mathbf{s}, \mathbf{a})$ , where  $\mathbf{x} \in \mathcal{X}$  and  $\mathcal{X} = \mathcal{S} \times \mathcal{A}$ .

## 2.4 Robustness of Imitation Learning

It can be verified that the density matching approach is not robust against noisy demonstrations according to the data generation assumption in Eq. (6). Specifically, given demonstrations drawn from  $\rho'$ , the density matching approach would solve  $\min_{\pi} D(\rho'(\mathbf{x}) || \rho_{\pi}(\mathbf{x}))$ . By assuming that the space of  $\pi$  is sufficiently large, the solution of this optimization problem is

$$\begin{aligned} \pi^*(\mathbf{a}|\mathbf{s}) &= \pi_E(\mathbf{a}|\mathbf{s}) \left( \frac{\alpha \rho_E(\mathbf{x})}{\alpha \rho_E(\mathbf{x}) + (1 - \alpha) \rho_N(\mathbf{x})} \right) \\ &+ \pi_N(\mathbf{a}|\mathbf{s}) \left( \frac{(1 - \alpha) \rho_N(\mathbf{x})}{\alpha \rho_E(\mathbf{x}) + (1 - \alpha) \rho_N(\mathbf{x})} \right), \end{aligned} \quad (7)$$

which yields  $D(\rho'(\mathbf{x}) || \rho_{\pi^*}(\mathbf{x})) = 0$ . However, this policy is not equivalent to the expert policy unless  $\alpha = 1$ . Therefore, density matching is not robust against noisy demonstrations generated according to Eq. (6).

We note that the data generation assumption in Eq. (6) has been considered previously by Wu et al. (2019). In this prior work, the authors proposed a robust method that learns the policy by solving

$$\min_{\pi} D(\rho'(\mathbf{x}) || \alpha \rho_{\pi}(\mathbf{x}) + (1 - \alpha) \rho_N(\mathbf{x})). \quad (8)$$

The authors showed that this optimization problem yields the expert policy under the data generation assumption in Eq. (6). However, solving this optimization problem requires  $\alpha$  and  $\rho_N$  which are typically unknown. To overcome this issue, Wu et al. (2019) proposed to estimate  $\alpha$  and  $\rho_N$  by using additional noisy

demonstrations that are labeled with a score determining the probability that demonstrations are drawn from  $\rho_E$ . However, this method is not applicable in our setting since labeled demonstrations are not available.

## 3 ROBUST IMITATION LEARNING

In this section, we propose our method for robust IL. Briefly, in Section 3.1, we propose an IL objective which optimizes a classification risk with a symmetric loss and prove its robustness. Then, in Section 3.2, we propose a new IL method that utilizes co-pseudo-labeling to optimize the classification risk. Lastly, we discuss the choice of a hyper-parameter in Section 3.3 and the choice of symmetric losses in Section 3.4.

### 3.1 Imitation Learning via Risk Optimization

Classification risks are fundamental quantities in classification (Hastie et al., 2001). We are interested in a *balanced risk* for binary classification where the class prior is balanced (Brodersen et al., 2010). Specifically, we propose to perform IL by solving the following risk optimization problem:

$$\max_{\pi} \min_g \mathcal{R}(g; \rho', \rho_{\pi}^{\lambda}, \ell_{\text{sym}}), \quad (9)$$

where  $\mathcal{R}$  is the balanced risk defined as

$$\mathcal{R}(g; \rho', \rho_{\pi}^{\lambda}, \ell) = \frac{1}{2} \mathbb{E}_{\rho'} [\ell(g(\mathbf{x}))] + \frac{1}{2} \mathbb{E}_{\rho_{\pi}^{\lambda}} [\ell(-g(\mathbf{x}))],$$

and  $\rho_{\pi}^{\lambda}$  is a mixture density defined as

$$\rho_{\pi}^{\lambda}(\mathbf{x}) = \lambda \rho_N(\mathbf{x}) + (1 - \lambda) \rho_{\pi}(\mathbf{x}). \quad (10)$$

Here,  $0 < \lambda < 1$  is a hyper-parameter,  $\pi$  is a policy to be learned by maximizing the risk,  $g : \mathcal{X} \mapsto \mathbb{R}$  is a classifier to be learned by minimizing the risk, and  $\ell_{\text{sym}} : \mathbb{R} \mapsto \mathbb{R}$  is a *symmetric loss* satisfying

$$\ell_{\text{sym}}(g(\mathbf{x})) + \ell_{\text{sym}}(-g(\mathbf{x})) = c, \quad (11)$$

for all  $\mathbf{x} \in \mathcal{X}$ , where  $c \in \mathbb{R}$  is a constant. Appropriate choices of the hyper-parameter and loss will be discussed in Sections 3.3 and 3.4, respectively.

We note that the balanced risk assumes that the positive and negative class priors are equal to  $\frac{1}{2}$ . This assumption typically makes the balanced risk more restrictive than other risks, because a classifier is learned to maximize the balanced accuracy instead of the accuracy (Menon et al., 2013; Lu et al., 2019). However, the balanced risk is not too restrictive for IL, because the goal is to learn the expert policy and the classifier is discarded after learning. Moreover, existing methods such as GAIL can be viewed as methods that optimize the balanced risk, as will be discussed in Section 3.3.

Next, we prove that the optimization in Eq. (9) yields the expert policy under the following assumption.

**Assumption 1** (Mixture state-action density). *The state-action density of the learning policy  $\pi$  is a mixture of the state-action densities of the expert and non-expert policies with a mixing coefficient  $0 \leq \kappa(\pi) \leq 1$ :*

$$\rho_\pi(\mathbf{x}) = \kappa(\pi)\rho_E(\mathbf{x}) + (1 - \kappa(\pi))\rho_N(\mathbf{x}), \quad (12)$$

where  $\rho_\pi(\mathbf{x})$ ,  $\rho_E(\mathbf{x})$ , and  $\rho_N(\mathbf{x})$  are the state-action densities of the learning policy, the expert policy, and the non-expert policy, respectively.

This assumption is based on the following observation on a typical optimization procedure of  $\pi$ : At the start of learning,  $\pi$  is randomly initialized and generates data that are similar to those of the non-expert policy. This scenario corresponds to Eq. (12) with  $\kappa(\pi) \approx 0$ . As training progresses, the policy improves and generates data that are a mixture of those from the expert and non-expert policies. This scenario corresponds to Eq. (12) with  $0 < \kappa(\pi) < 1$ . Indeed, the scenario where the agent successfully learns the expert policy corresponds to Eq. (12) with  $\kappa(\pi) = 1$ .

We note that a policy uniquely corresponding to  $\rho_\pi$  in Eq. (12) is a mixture between  $\pi_E$  and  $\pi_N$  with a mixture coefficient depending on  $\kappa(\pi)$ . However, we cannot directly evaluate the value of  $\kappa(\pi)$ . This is because we do not directly optimize the state-action density  $\rho_\pi$ . Instead, we optimize the policy  $\pi$  by using an RL method, as will be discussed in Section 3.2.

Under Assumption 1, we obtain Lemma 1.

**Lemma 1.** *Letting  $\ell_{\text{sym}}(\cdot)$  be a symmetric loss that satisfies  $\ell_{\text{sym}}(g(\mathbf{x})) + \ell_{\text{sym}}(-g(\mathbf{x})) = c$ ,  $\forall \mathbf{x} \in \mathcal{X}$  and a constant  $c \in \mathbb{R}$ , the following equality holds.*

$$\begin{aligned} \mathcal{R}(g; \rho', \rho_\pi^\lambda, \ell_{\text{sym}}) &= (\alpha - \kappa(\pi)(1 - \lambda))\mathcal{R}(g; \rho_E, \rho_N, \ell_{\text{sym}}) \\ &\quad + \frac{1 - \alpha + \kappa(\pi)(1 - \lambda)}{2}c. \end{aligned} \quad (13)$$

The proof is given in Appendix A, which follows Charoenphakdee et al. (2019). This lemma indicates that, a minimizer  $g^*$  of  $\mathcal{R}(g; \rho', \rho_\pi^\lambda, \ell_{\text{sym}})$  is identical to that of  $\mathcal{R}(g; \rho_E, \rho_N, \ell_{\text{sym}})$ :

$$\begin{aligned} g^* &= \underset{g}{\operatorname{argmin}} \mathcal{R}(g; \rho', \rho_\pi^\lambda, \ell_{\text{sym}}) \\ &= \underset{g}{\operatorname{argmin}} \mathcal{R}(g; \rho_E, \rho_N, \ell_{\text{sym}}), \end{aligned} \quad (14)$$

when  $\alpha - \kappa(\pi)(1 - \lambda) > 0$ . This result enables us to prove that the maximizer of the risk optimization in Eq. (9) is the expert policy.

**Theorem 1.** *Given the optimal classifier  $g^*$  in Eq. (14), the solution of  $\max_\pi \mathcal{R}(g^*; \rho', \rho_\pi^\lambda, \ell_{\text{sym}})$  is equivalent to the expert policy.*

*Proof sketch.* It can be shown that the solution of  $\max_\pi \mathcal{R}(g^*; \rho', \rho_\pi^\lambda, \ell_{\text{sym}})$  is equivalent to the solution of  $\max_\pi \kappa(\pi) \left( \mathbb{E}_{\rho_E} [\ell_{\text{sym}}(-g^*(\mathbf{x}))] - \mathbb{E}_{\rho_N} [\ell_{\text{sym}}(-g^*(\mathbf{x}))] \right)$ . Since  $g^* = \operatorname{argmin}_g \mathcal{R}(g; \rho_E, \rho_N, \ell_{\text{sym}})$ , it follows that  $\mathbb{E}_{\rho_E} [\ell_{\text{sym}}(-g^*(\mathbf{x}))] - \mathbb{E}_{\rho_N} [\ell_{\text{sym}}(-g^*(\mathbf{x}))] > 0$ . Thus,  $\max_\pi \kappa(\pi) \left( \mathbb{E}_{\rho_E} [\ell_{\text{sym}}(-g^*(\mathbf{x}))] - \mathbb{E}_{\rho_N} [\ell_{\text{sym}}(-g^*(\mathbf{x}))] \right)$  is solved by increasing  $\kappa(\pi)$  to 1. Because  $\kappa(\pi) = 1$  if and only if  $\rho_\pi(\mathbf{x}) = \rho_E(\mathbf{x})$  under Assumption 1, we conclude that the solution of  $\max_\pi \mathcal{R}(g^*; \rho', \rho_\pi^\lambda, \ell_{\text{sym}})$  is equivalent to the expert policy.  $\square$

A detailed proof is given in Appendix A. This result indicates that robust IL can be achieved by optimizing the risk in Eq. (9). In a practice aspect, this is a significant advance compared to the prior work (Wu et al., 2019), because Theorem 1 shows that robust IL can be achieved *without* the knowledge of the mixing coefficient  $\alpha$  or additional labels to estimate  $\alpha$ . Next, we present a new IL method that empirically solves Eq. (9) by using co-pseudo-labeling.

### 3.2 Co-pseudo-labeling for Risk Optimization

While the risk in Eq. (9) leads to robust IL, we cannot directly optimize this risk in our setting. This is because the risk contains an expectation over  $\rho_N(\mathbf{x})$ , but we are given only demonstration samples drawn from  $\rho'(\mathbf{x})^2$ . We address this issue by using co-pseudo-labeling to approximately draw samples from  $\rho_N(\mathbf{x})$ , as described below.

Recall that the optimal classifier  $g^*(\mathbf{x})$  in Eq. (14) also minimizes the risk  $\mathcal{R}(g; \rho_E, \rho_N, \ell_{\text{sym}})$ . Therefore, given a state-action sample  $\tilde{\mathbf{x}} \in \mathcal{X}$ , we can use  $g^*(\tilde{\mathbf{x}})$  to predict whether  $\tilde{\mathbf{x}}$  is drawn from  $\rho_E$  or  $\rho_N$ . Specifically,  $\tilde{\mathbf{x}}$  is predicted to be drawn from  $\rho_E$  when  $g^*(\tilde{\mathbf{x}}) \geq 0$ , and it is predicted to be drawn from  $\rho_N$  when  $g^*(\tilde{\mathbf{x}}) < 0$ . Based on this observation, our key idea is to approximate the expectation over  $\rho_N$  in Eq. (9) by using samples that are predicted by  $g$  to be drawn from  $\rho_N$ .

To realize this idea, we firstly consider the following empirical risk with a semi-supervised learning technique called *pseudo-labeling* (Chapelle et al., 2010):

$$\begin{aligned} \hat{\mathcal{R}}(g) &= \frac{1}{2} \hat{\mathbb{E}}_{\mathcal{D}} [\ell_{\text{sym}}(g(\mathbf{x}))] + \frac{\lambda}{2} \hat{\mathbb{E}}_{\mathcal{P}} [\ell_{\text{sym}}(-g(\mathbf{x}))] \\ &\quad + \frac{1 - \lambda}{2} \hat{\mathbb{E}}_{\mathcal{B}} [\ell_{\text{sym}}(-g(\mathbf{x}))]. \end{aligned} \quad (15)$$

Here,  $\hat{\mathbb{E}}[\cdot]$  denotes an empirical expectation (i.e., the sample average),  $\mathcal{D}$  is the demonstration dataset in Eq. (5),  $\mathcal{B}$  is a dataset of trajectories collected by using  $\pi$ , and  $\mathcal{P}$  is a dataset of pseudo-labeled demonstra-

<sup>2</sup>The expectation over  $\rho_\pi(\mathbf{x})$  can be approximated using trajectories independently collected by the policy  $\pi$ .

tions obtained by choosing demonstrations in  $\mathcal{D}$  with  $g(\mathbf{x}) < 0$ , i.e., samples that are predicted to be drawn from  $\rho_N$ . This risk with pseudo-labeling enables us to empirically solve Eq. (9) in our setting. However, the trained classifier may perform poorly, mainly because samples in  $\mathcal{P}$  are labeled by the classifier itself. Specifically, the classifier during training may predict the labels of demonstrations incorrectly, i.e., demonstrations drawn from  $\rho_E(\mathbf{x})$  are incorrectly predicted to be drawn from  $\rho_N(\mathbf{x})$ . This may degrade the performance of the classifier, because using incorrectly-labeled data can reinforce the classifier to be over-confident in its incorrect prediction (Kingma et al., 2014).

To remedy the over-confidence of the classifier, we propose *co-pseudo-labeling*, which combines the ideas of pseudo-labeling and co-training (Blum and Mitchell, 1998). Specifically, we train two classifiers denoted by  $g_1$  and  $g_2$  by minimizing the following empirical risks:

$$\begin{aligned} \widehat{\mathcal{R}}_1(g_1) &= \frac{1}{2} \widehat{\mathbb{E}}_{\mathcal{D}_1} [\ell_{\text{sym}}(g_1(\mathbf{x}))] + \frac{\lambda}{2} \widehat{\mathbb{E}}_{\mathcal{P}_1} [\ell_{\text{sym}}(-g_1(\mathbf{x}))] \\ &\quad + \frac{1-\lambda}{2} \widehat{\mathbb{E}}_{\mathcal{B}} [\ell_{\text{sym}}(-g_1(\mathbf{x}))], \end{aligned} \quad (16)$$

$$\begin{aligned} \widehat{\mathcal{R}}_2(g_2) &= \frac{1}{2} \widehat{\mathbb{E}}_{\mathcal{D}_2} [\ell_{\text{sym}}(g_2(\mathbf{x}))] + \frac{\lambda}{2} \widehat{\mathbb{E}}_{\mathcal{P}_2} [\ell_{\text{sym}}(-g_2(\mathbf{x}))] \\ &\quad + \frac{1-\lambda}{2} \widehat{\mathbb{E}}_{\mathcal{B}} [\ell_{\text{sym}}(-g_2(\mathbf{x}))], \end{aligned} \quad (17)$$

where  $\mathcal{D}_1$  and  $\mathcal{D}_2$  are disjoint subsets of  $\mathcal{D}$ . Pseudo-labeled dataset  $\mathcal{P}_1$  is obtained by choosing demonstrations from  $\mathcal{D}_2$  with  $g_2(\mathbf{x}) < 0$ , and pseudo-labeled dataset  $\mathcal{P}_2$  is obtained by choosing demonstrations from  $\mathcal{D}_1$  with  $g_1(\mathbf{x}) < 0$ . With these risks, we reduce the influence of over-confident classifiers because  $g_1$  is trained using samples pseudo-labeled by  $g_2$  and vice-versa (Han et al., 2018). We call our proposed method *Robust IL with Co-pseudo-labeling* (RIL-Co).

We implement RIL-Co by using a stochastic gradient method to optimize the empirical risk where we alternately optimize the classifiers and policy. Recall from Eq. (9) that we aim to maximize  $\mathcal{R}(g; \rho', \rho_\pi^\lambda, \ell_{\text{sym}})$  w.r.t.  $\pi$ . After ignoring terms that are constant w.r.t.  $\pi$ , solving this maximization is equivalent to maximizing  $\mathbb{E}_{\rho_\pi} [\ell_{\text{sym}}(-g(\mathbf{x}))]$  w.r.t.  $\pi$ . This objective is identical to the RL objective in Eq. (1) with a reward function  $r(\mathbf{x}) = \ell_{\text{sym}}(-g(\mathbf{x}))$ . Therefore, we can train the policy in RIL-Co by simply using an existing RL method, e.g., the trust-region policy gradient (Wu et al., 2017). We summarize the procedure of RIL-Co in Algorithm 1. Sourcecode of our implementation is available at [https://github.com/voot-t/ril\\_co](https://github.com/voot-t/ril_co).

### 3.3 Choice of Hyper-parameter

We propose to use  $\lambda = 0.5$  for RIL-Co, because it makes the equality in Eq. (14) holds which is essential

for our theoretical result in Section 3.1. Specifically, recall that Theorem 1 relies on the equality in Eq. (14). At the same time, the equality in Eq. (14) holds if the the following inequality holds:

$$\alpha - \kappa(\pi)(1 - \lambda) > 0. \quad (18)$$

This inequality depends on  $\alpha$ ,  $\kappa(\pi)$ , and  $\lambda$ , where  $\alpha$  is unknown,  $\kappa(\pi)$  depends on the policy, and  $\lambda$  is the hyper-parameter. However, we cannot choose nor evaluate  $\kappa(\pi)$  since we do not directly optimize  $\kappa(\pi)$  during policy training. Thus, we need to choose an appropriate value of  $\lambda$  so that the inequality in Eq. (18) always holds. Recall that we assumed  $0.5 < \alpha < 1$  in Section 2.3. Under this assumption, the inequality in Eq. (18) holds regardless of the true value of  $\alpha$  when  $\kappa(\pi)(1 - \lambda) \leq 0.5$ . Since the value of  $\kappa(\pi)$  increases to 1 as the policy improves by training (see Assumption 1), the appropriate value of  $\lambda$  is  $0.5 \leq \lambda < 1$ .

However, a large value of  $\lambda$  may not be preferable in practice since it increases the influence of pseudo-labels on the risks (i.e., the second term in Eqs. (16) and (17)). These pseudo-labels should not have larger influence than real labels (i.e., the third term in Eqs. (16) and (17)). For this reason, we decided to use  $\lambda = 0.5$ , which is the smallest value of  $\lambda$  that ensures the inequality in Eq. (18) to be always held during training. Nonetheless, we note that while RIL-Co with  $\lambda = 0.5$  already yields good performance in the following experiments, the performance may still be improved by fine-tuning  $\lambda$  using e.g., grid-search.

**Remarks.** Choosing  $\lambda = 0$  corresponds to omitting co-pseudo-labeling, and doing so reduces RIL-Co to variants of GAIL which are not robust. Concretely, the risk optimization problem of RIL-Co with  $\lambda = 0$  is  $\max_{\pi} \min_g \mathcal{R}(g; \rho', \rho_\pi, \ell_{\text{sym}})$ . By using the logistic loss:  $\ell(z) = \log(1 + \exp(-z))$ , instead of a symmetric loss, we obtain the following risk:

$$\begin{aligned} 2\mathcal{R}(g; \rho', \rho_\pi, \ell) &= \mathbb{E}_{\rho'} [\log(1 + \exp(-g(\mathbf{x})))] \\ &\quad + \mathbb{E}_{\rho_\pi} [\log(1 + \exp(g(\mathbf{x})))], \end{aligned} \quad (19)$$

which is the negative of GAIL’s objective in Eq. (4)<sup>3</sup>. Meanwhile, we may obtain other variants of GAIL by using symmetric losses such as the sigmoid loss and the unhinged loss (van Rooyen et al., 2015). In particular, with the unhinged loss:  $\ell(z) = 1 - z$ , the risk becomes the negative of Wasserstein GAIL’s objective with an additive constant (Li et al., 2017; Xiao et al., 2019):

$$2\mathcal{R}(g; \rho', \rho_\pi, \ell) = \mathbb{E}_{\rho'} [-g(\mathbf{x})] + \mathbb{E}_{\rho_\pi} [g(\mathbf{x})] + \frac{1}{2}. \quad (20)$$

However, even when  $\ell(z)$  is symmetric, we conjecture that such variants of GAIL are not robust, because

<sup>3</sup>Here,  $\rho'$  replaces  $\rho_E$ . The sign flips since RIL-Co and GAIL solve max-min and min-max problems, respectively.

**Algorithm 1** RIL-Co: Robust Imitation Learning with Co-pseudo-labeling

- 1: **Input:** Demonstration dataset  $\mathcal{D}$ , initial policy  $\pi$ , and initial classifiers  $g_1$  and  $g_2$ .
- 2: Set hyper-parameter  $\lambda = 0.5$  (see Section 3.3) and batch-sizes ( $B = U = V = 640$  and  $K = 128$ ).
- 3: Split  $\mathcal{D}$  into two disjoint datasets  $\mathcal{D}_1$  and  $\mathcal{D}_2$ .
- 4: **while** Not converge **do**
- 5:   **while**  $|\mathcal{B}| < B$  with batch size  $B$  **do**
- 6:     Use  $\pi$  to collect and include transition samples into  $\mathcal{B}$
- 7:   **Co-pseudo-labeling:**
- 8:     Sample  $\{\mathbf{x}_u\}_{u=1}^U$  from  $\mathcal{D}_2$ , and choose  $K$  samples with  $g_2(\mathbf{x}_u) < 0$  in an ascending order as  $\mathcal{P}_1$ .
- 9:     Sample  $\{\mathbf{x}_v\}_{v=1}^V$  from  $\mathcal{D}_1$ , and choose  $K$  samples with  $g_1(\mathbf{x}_v) < 0$  in an ascending order as  $\mathcal{P}_2$ .
- 10:   **Train classifiers:**
- 11:     Train  $g_1$  by performing gradient descent to minimize the empirical risk  $\hat{R}_1(g_1)$  using  $\mathcal{D}_1$ ,  $\mathcal{P}_1$  and  $\mathcal{B}$ .
- 12:     Train  $g_2$  by performing gradient descent to minimize the empirical risk  $\hat{R}_2(g_2)$  using  $\mathcal{D}_2$ ,  $\mathcal{P}_2$  and  $\mathcal{B}$ .
- 13:   **Train policy:**
- 14:     Train the policy by an RL method with transition samples in  $\mathcal{B}$  and rewards  $r(\mathbf{x}) = \ell_{\text{sym}}(-g_1(\mathbf{x}))$ .

Table 1: Examples of losses and their symmetric property, i.e., whether  $\ell(z) + \ell(-z) = c$ . We denote normalized counterparts of non-symmetric losses by (N). The AP loss in Eq. (21) is a linear combination of the normalized logistic and sigmoid losses.

Loss name	$\ell(z)$	Symmetric
Logistic	$\log(1 + \exp(-z))$	✗
Hinge	$\max(1 - z, 0)$	✗
Sigmoid	$1/(1 + \exp(z))$	✓
Unhinged	$1 - z$	✓
Logistic (N)	$\frac{\log(1 + \exp(-z))}{\sum_{k \in \{-1, 1\}} \log(1 + \exp(-zk))}$	✓
Hinge (N)	$\frac{\max(1 - z, 0)}{\sum_{k \in \{-1, 1\}} \max(1 - zk, 0)}$	✓

$\lambda = 0$  does not make the inequality in Eq. (18) holds when  $\kappa(\pi) > 0.5$ .

### 3.4 Choice of Symmetric Loss

In our implementation of RIL-Co, we use the active-passive loss (AP loss) (Ma et al., 2020) defined as

$$\ell_{\text{AP}}(z) = \frac{0.5 \times \log(1 + \exp(-z))}{\log(1 + \exp(-z)) + \log(1 + \exp(z))} + \frac{0.5}{1 + \exp(z)}, \quad (21)$$

which satisfies  $\ell_{\text{AP}}(z) + \ell_{\text{AP}}(-z) = 1$ . This loss is a linear combination of two symmetric losses: the normalized logistic loss (the first term) and the sigmoid loss (the second term). It was shown that this loss suffers less from the issue of under-fitting when compared to each of the normalized logistic loss or the sigmoid loss (Ma et al., 2020). However, we emphasize that any symmetric loss can be used to learn the expert policy with RIL-Co, as indicated by our theo-

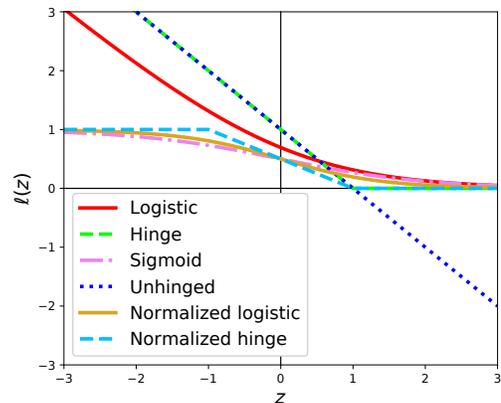


Figure 1: The value of losses in Table 1. Non-symmetric losses, i.e., the logistic and hinge losses, become symmetric after normalization.

retical result in Section 3.1. In addition, any loss can be made symmetric by using normalization (Ma et al., 2020). Therefore, the requirement of symmetric losses is not a severe limitation. Table 1 and Figure 1 show examples of non-symmetric and symmetric losses.

## 4 EXPERIMENTS

We evaluate the robustness of RIL-Co on continuous-control benchmarks simulated by PyBullet simulator (HalfCheetah, Hopper, Walker2d, and Ant) (Coumans and Bai, 2019). These tasks are equipped with the true reward functions that we use for the evaluation purpose. The learning is conducted using the true states (e.g., joint positions) and not the visual observations. We report the mean and standard error of the performance (cumulative true rewards) over 5 trials.

We compare RIL-Co with the AP loss against the following baselines: BC (Pomerleau, 1988), FAIRL (Ghasemipour et al., 2020), VILD (Tangkaratt et al., 2020), and three variants of GAIL where each

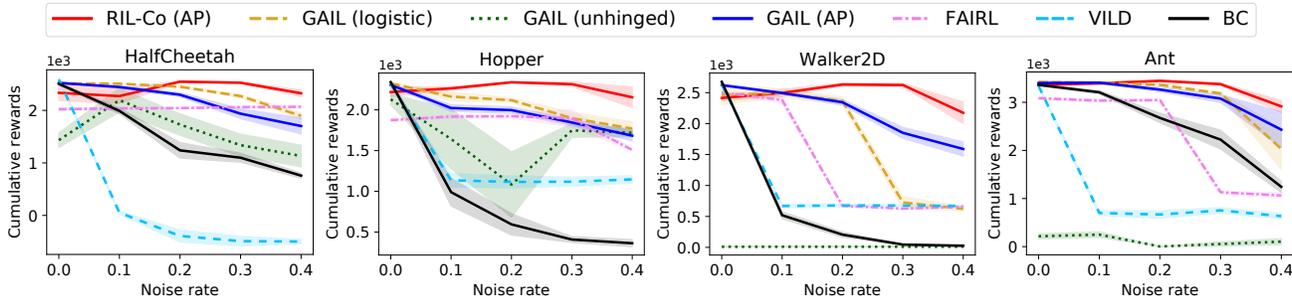


Figure 2: Final performance in continuous-control benchmarks with different noise rates. Vertical axes denote cumulative rewards obtained during the last 1000 training iterations. Shaded regions denote standard errors computed over 5 runs. RIL-Co performs well even when the noise rate increases. Meanwhile, the performance of other methods significantly degrades as the noise rate increases.

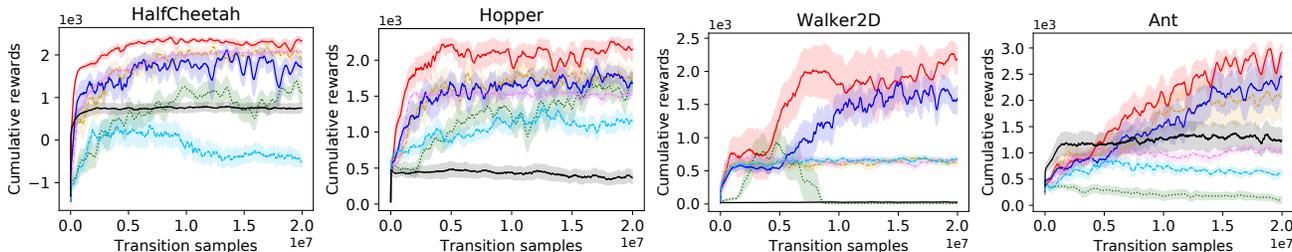


Figure 3: Performance against the number of transition samples in continuous-control benchmarks with noise rate  $\delta = 0.4$ . RIL-Co achieves better performances and uses less transition samples compared to other methods.

variant uses different losses: logistic, unhinged, and AP. As discussed in the remarks of Section 3.3, GAIL with the logistic loss denotes the original GAIL that performs density matching with the Jensen-Shannon divergence, GAIL with the unhinged loss denotes a variant of GAIL that performs density matching with the Wasserstein distance, and GAIL with the AP loss corresponds to RIL-Co without co-pseudo-labeling.

All methods use policy networks with 2 hidden-layers of 64 hyperbolic tangent units. We use similar networks with 100 hyperbolic tangent units for classifiers in RIL-Co and discriminators in other methods. The policy networks are trained by the trust region policy gradient (Wu et al., 2017) from a public implementation (Kostrikov, 2018). The classifiers and discriminators are trained by Adam (Kingma and Ba, 2015) with the gradient penalty regularizer with the regularization parameter of 10 (Gulrajani et al., 2017). The total number of transition samples collected by the learning policy is 20 million. More details of experimental setting can be found in Appendix B.

#### 4.1 Evaluation on Noisy Datasets with Different Noise Rates

In this experiment, we evaluate RIL-Co on noisy datasets generated with different noise rates. To obtain datasets, we firstly train policies by RL with the true reward functions. Next, we choose 6 policy snapshots where each snapshot is trained using different

numbers of transition samples. Then, we use the best performing policy snapshot (in terms of cumulative rewards) to collect 10000 expert state-action samples, and use the other 5 policy snapshots to collect a total of 10000 non-expert state-action samples. Lastly, we generate datasets with different noise rates by mixing expert and non-expert state-action samples, where noise rate  $\delta \in \{0, 0.1, 0.2, 0.3, 0.4\}$  approximately determines the number of randomly chosen non-expert state-action samples. Specifically, a dataset consisting of 10000 expert samples corresponds to a dataset with  $\delta = 0$  (i.e., no noise), whereas a dataset consisting of 10000 expert samples and 7500 randomly chosen non-expert samples corresponds to a dataset with  $\delta = 0.4$  approximately<sup>4</sup>. We note that the value of  $\delta$  approximately equals to the value of  $1 - \alpha$  in Eq. (6).

Figure 2 shows the final performance achieved by each method. We can see that RIL-Co outperforms comparison methods and achieves the best performance in high noise scenarios where  $\delta \in \{0.2, 0.3, 0.4\}$ . Meanwhile, in low noise scenarios where  $\delta \in \{0.0, 0.1\}$ , RIL-Co performs comparable to the best performing methods such as GAIL with the logistic and AP losses. Overall, the results show that RIL-Co achieves good performance in the presence of noises, while the other methods fail to learn and their performance degrades as the noise rate increases.

<sup>4</sup>The true noise rates of these datasets are as follows:  $\tilde{\delta} \in \{0, 1000/11000, 2500/12500, 5000/15000, 7500/17500\}$ .

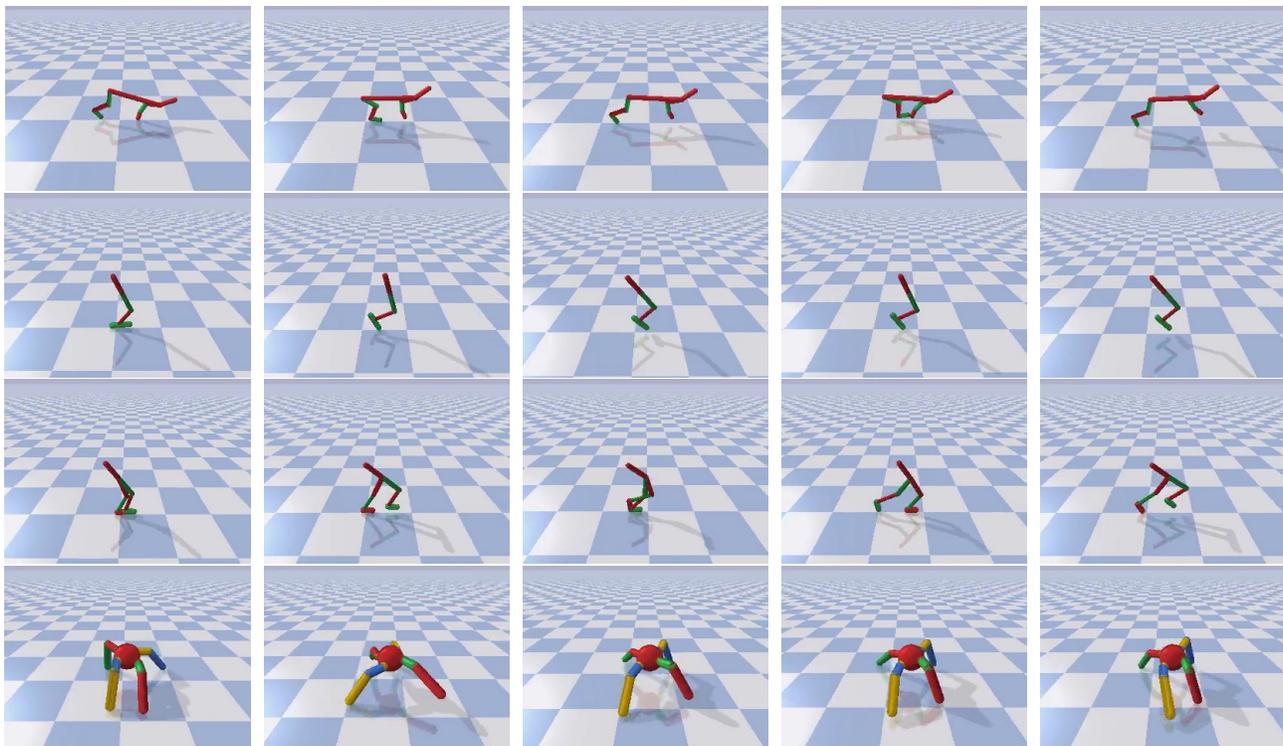


Figure 4: Visualization of the first 100 time steps of trajectories obtained by RIL-Co from datasets with noise rate  $\delta = 0.4$ . Time step increases from the leftmost figure ( $t = 0$ ) to the rightmost figure ( $t = 100$ ). RIL-Co agents successfully solve these tasks. The obtained trajectories also closely resemble expert demonstrations.

In contrast, density matching methods, namely FAIRL and GAIL with the logistic and unhinged losses, do not perform well. This is as expected, because these methods would learn a policy that is a mixture of the expert and non-expert policies and would not perform well. Also notice that GAIL with the unhinged loss performs very poorly on the Walker2D and Ant tasks even with noiseless demonstrations where  $\delta = 0$ . This is an intriguing result given that the unhinged loss is also symmetric similarly to the AP loss. We conjecture that the poor performance is due to the unboundedness from below of the unhinged loss (see Figure 1). This unboundedness may lead to a poorly behaved classifier that outputs values with very large magnitudes, as suggested by Charoenphakdee et al. (2019). With such a classifier, we expect that GAIL with the unhinged loss would require a strong regularization to perform well, especially for complex control tasks.

In addition, we can see that RIL-Co is more robust when compared to GAIL with the AP loss. This result supports our theorem which indicates that a symmetric loss alone is insufficient for robustness. Interestingly, with expert demonstrations (i.e.,  $\delta = 0$ ), GAIL tends to outperform RIL-Co. This is perhaps because co-pseudo-labeling introduces additional biases. This could be avoided by initially using  $\lambda = 0$  (i.e., per-

forming GAIL) and gradually increasing the value to  $\lambda = 0.5$  as learning progresses.

On the other hand, VILD performs poorly with noisy datasets even with a small noise rate of  $\delta = 0.1$ . We conjecture that this is because VILD could not accurately estimate the noise distributions due to the violation of its Gaussian assumption. Specifically, VILD assumes that noisy demonstrations are generated by adding Gaussian noise to actions drawn from the expert policy, and that expert demonstrations consist of low-variance actions. However, noisy demonstrations in this experiment are generated by using policy snapshots without adding any noise. In this case, non-expert demonstrations may consist of low-variance actions (e.g., non-expert policy may yield a constant action). Due to this, VILD cannot accurately estimate the noise distributions and performs poorly. Meanwhile, behavior cloning (BC) does not perform well. This is because BC assumes that demonstrations are generated by experts. It also suffers from the issue of compounding error which worsens the performance.

Next, Figure 3 shows the performance against the number of transition samples collected by the learning policy for  $\delta = 0.4$ . RIL-Co achieves better performances and uses less transition samples when compared to other methods. This result indicates that

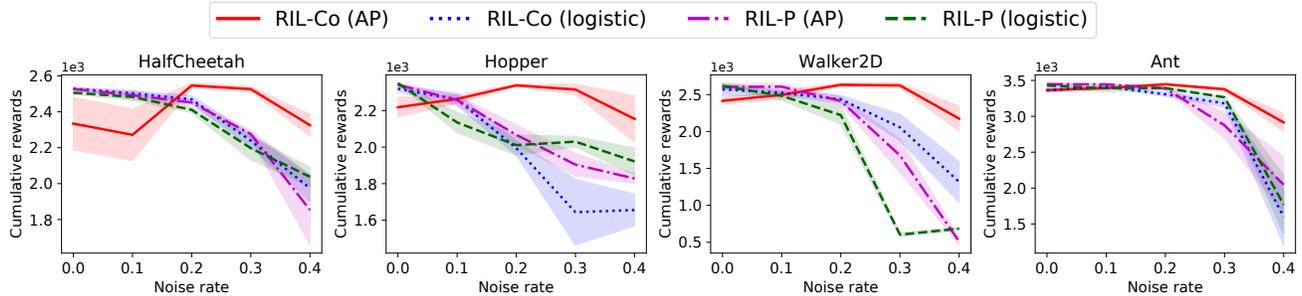


Figure 5: Final performance of variants of RIL-Co in the ablation study. RIL-Co with the AP loss performs the best. This shows that using a symmetric loss and co-pseudo-labeling together is important for robustness.

RIL-CO is data efficient. Results for different noise rates can be found in Appendix C, and they show a similar tendency.

Lastly, Figure 4 depicts visualization of trajectories obtained by the policy of RIL-Co. Indeed, RIL-Co agents successfully solve the tasks. The trajectories also closely resemble expert demonstrations shown in Appendix B. This qualitative result further verifies that RIL-Co successfully learns the expert policy.

## 4.2 Ablation Study

In this section, we conduct ablation study by evaluating different variants of RIL-Co. Specifically, we evaluate RIL-Co with the logistic loss to investigate the importance of symmetric loss. In addition, to investigate the importance of co-pseudo-labeling, we also evaluate Robust IL with Pseudo-labeling (RIL-P) which uses naive pseudo-labeling in Eq. (15) instead of co-pseudo-labeling. Experiments are conducted using the same datasets in the previous section.

The result in Figure 5 shows that RIL-Co with the AP loss outperforms the variants. This result further indicates that using a symmetric loss and co-pseudo-labeling together is important for robustness.

## 4.3 Evaluation on Gaussian Noise Dataset

Next, we evaluate RIL-Co with the AP loss in the Ant task with a noisy dataset generated by Gaussian noise. Specifically, we use a dataset with 10000 expert and 7500 non-expert state-action samples, where non-expert samples are obtained by adding Gaussian noise to action samples drawn from the expert policy. This dataset is generated according to the main assumption of VILD (Tangkaratt et al., 2020), and we expect VILD to perform well in this experiment.

Figure 6 depicts the performance against the number of transition samples. It can be seen that VILD performs much better in this setting compared to the previous setting in Figure 2. This is as expected, since the

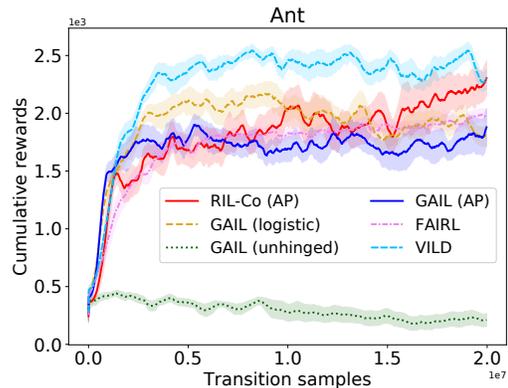


Figure 6: Performance with a Gaussian noise dataset. RIL-Co performs better than others except VILD.

Gaussian assumption of VILD is correct in this setting but incorrect in the previous setting. Still, RIL-Co achieves a performance comparable to that of VILD with 20 million samples, even though RIL-Co relies on a milder data generation assumption (see Section 2.3). Meanwhile, the other methods do not perform as well as RIL-Co and VILD.

Overall, the empirical results in our experiments indicate that RIL-Co is more robust against noisy demonstrations when compared to existing methods.

## 5 CONCLUSIONS

We presented a new method for IL from noisy demonstrations. We proved that robust IL can be achieved by optimizing a classification risk with a symmetric loss, and we proposed RIL-Co which optimizes the risk by using co-pseudo-labeling. We showed through experiments that RIL-Co is more robust against noisy demonstrations when compared to existing methods.

In this paper, we utilized co-pseudo-labeling to approximate data from non-expert densities. However, data from non-expert densities may be readily available (Grollman and Billard, 2011). Utilizing such data for robust IL is an interesting future direction.

## Acknowledgement

We thank the reviewers for their useful comments. NC was supported by MEXT scholarship and Google PhD Fellowship program. MS was supported by KAKENHI 17H00757.

## References

- Angluin, D. and Laird, P. (1988). Learning from noisy examples. *Machine Learning*.
- Blum, A. and Mitchell, T. (1998). Combining labeled and unlabeled data with co-training. In *International Conference on Computational Learning Theory*.
- Brodersen, K. H., Ong, C. S., Stephan, K. E., and Buhmann, J. M. (2010). The balanced accuracy and its posterior distribution. In *International Conference on Pattern Recognition*.
- Brown, D., Coleman, R., Srinivasan, R., and Niekum, S. (2020). Safe imitation learning via fast bayesian reward inference from preferences. In *International Conference on Machine Learning*.
- Brown, D. S., Goo, W., Nagarajan, P., and Niekum, S. (2019). Extrapolating beyond suboptimal demonstrations via inverse reinforcement learning from observations. In *International Conference on Machine Learning*.
- Chapelle, O., Schlkopf, B., and Zien, A. (2010). *Semi-Supervised Learning*. The MIT Press, 1st edition.
- Charoenphakdee, N., Lee, J., and Sugiyama, M. (2019). On symmetric losses for learning from corrupted labels. In *International Conference on Machine Learning*.
- Coumans, E. and Bai, Y. (2016–2019). Pybullet, a python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>.
- Ghasemipour, S. K. S., Zemel, R., and Gu, S. (2020). A divergence minimization perspective on imitation learning methods. In *Proceedings of the Conference on Robot Learning*.
- Grollman, D. H. and Billard, A. (2011). Donut as I do: Learning from failed demonstrations. In *IEEE International Conference on Robotics and Automation, ICRA*, pages 3804–3809.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. C. (2017). Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems*.
- Han, B., Yao, Q., Yu, X., Niu, G., Xu, M., Hu, W., Tsang, I. W., and Sugiyama, M. (2018). Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *Advances in Neural Information Processing Systems*.
- Hastie, T., Tibshirani, R., and Friedman, J. (2001). *The Elements of Statistical Learning*. Springer New York Inc.
- Ho, J. and Ermon, S. (2016). Generative adversarial imitation learning. In *Advances in Neural Information Processing Systems*.
- Kingma, D. P. and Ba, J. (2015). Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- Kingma, D. P., Mohamed, S., Jimenez Rezende, D., and Welling, M. (2014). Semi-supervised learning with deep generative models. In *Advances in Neural Information Processing Systems*.
- Kostrikov, I. (2018). Pytorch implementations of reinforcement learning algorithms. <https://github.com/ikostrikov/pytorch-a2c-ppo-acktr-gail>.
- Li, Y., Song, J., and Ermon, S. (2017). Infogail: Interpretable imitation learning from visual demonstrations. In *Advances in Neural Information Processing Systems*.
- Lu, N., Niu, G., Menon, A. K., and Sugiyama, M. (2019). On the minimal supervision for training any binary classifier from only unlabeled data. In *International Conference on Learning Representations*.
- Ma, X., Huang, H., Wang, Y., Romano, S., Erfani, S., and Bailey, J. (2020). Normalized loss functions for deep learning with noisy labels. In *International Conference on Machine Learning*.
- Mandlekar, A., Zhu, Y., Garg, A., Booher, J., Spero, M., Tung, A., Gao, J., Emmons, J., Gupta, A., Orbay, E., Savarese, S., and Fei-Fei, L. (2018). ROBO-TURK: A crowdsourcing platform for robotic skill learning through imitation. In *Conference on Robot Learning*.
- Menon, A., Narasimhan, H., Agarwal, S., and Chawla, S. (2013). On the statistical consistency of algorithms for binary classification under class imbalance. In *International Conference on Machine Learning*.
- Natarajan, N., Dhillon, I. S., Ravikumar, P. K., and Tewari, A. (2013). Learning with noisy labels. In *Advances in Neural Information Processing Systems*.
- Ng, A. Y. and Russell, S. J. (2000). Algorithms for inverse reinforcement learning. In *International Conference on Machine Learning*.
- Pomerleau, D. (1988). ALVINN: an autonomous land vehicle in a neural network. In *Advances in Neural Information Processing Systems*.
- Puterman, M. L. (1994). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*.

- Schaal, S. (1999). Is imitation learning the route to humanoid robots? *Trends in Cognitive Sciences*.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., Chen, Y., Lillicrap, T., Hui, F., Sifre, L., van den Driessche, G., Graepel, T., and Hassabis, D. (2017). Mastering the game of Go without human knowledge. *Nature*.
- Sun, W., Vemula, A., Boots, B., and Bagnell, D. (2019). Provably efficient imitation learning from observation alone. In *International Conference on Machine Learning*.
- Syed, U., Bowling, M. H., and Schapire, R. E. (2008). Apprenticeship learning using linear programming. In *International Conference on Machine Learning*.
- Tangkaratt, V., Han, B., Khan, M. E., and Sugiyama, M. (2020). Variational imitation learning with diverse-quality demonstrations. In *International Conference on Machine Learning*.
- van Rooyen, B., Menon, A., and Williamson, R. C. (2015). Learning with symmetric label noise: The importance of being unhinged. In *Advances in Neural Information Processing Systems 28*.
- Wu, Y., Charoenphakdee, N., Bao, H., Tangkaratt, V., and Sugiyama, M. (2019). Imitation learning from imperfect demonstration. In *International Conference on Machine Learning*.
- Wu, Y., Mansimov, E., Grosse, R. B., Liao, S., and Ba, J. (2017). Scalable trust-region method for deep reinforcement learning using kronecker-factored approximation. In *Advances in Neural Information Processing Systems*.
- Xiao, H., Herman, M., Wagner, J., Ziesche, S., Etesami, J., and Linh, T. H. (2019). Wasserstein adversarial imitation learning. *CoRR*, abs/1906.08113.
- Ziebart, B. D., Bagnell, J. A., and Dey, A. K. (2010). Modeling interaction via the principle of maximum causal entropy. In *International Conference on Machine Learning*.