IMPROVING GENERATIVE ADVERSARIAL IMITATION LEARNING WITH NON-EXPERT DEMONSTRATIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

Imitation learning aims to learn an optimal policy from expert demonstrations and its recent combination with deep learning has shown impressive performance. However, collecting a large number of expert demonstrations for deep learning is time-consuming and requires much expert effort. In this paper, we propose a method to improve generative adversarial imitation learning by using additional information from non-expert demonstrations which are easier to obtain. The key idea of our method is to perform multiclass classification to learn discriminator functions where non-expert demonstrations are regarded as being drawn from an extra class. Experiments in continuous control tasks demonstrate that our method learns optimal policies faster and has more stable performance than the generative adversarial imitation learning baseline.

1 INTRODUCTION

The goal of sequential decision making problems is to learn an optimal policy that exhibits task-solving behavior. Reinforcement learning (RL) is a powerful approach to find such a policy by maximizing rewards computed by a reward function (Puterman, 1994; Sutton & Barto, 1998). While RL has achieved great success in solving challenging tasks (Mnih et al., 2015; Silver et al., 2017), its performance depends heavily on a *good* reward function which well captures the concept of task-solving behavior. Unfortunately, designing such a good reward function is a difficult trial-and-error process and often time-consuming. This difficulty is one of the major limitations of RL for many real-world applications.

Imitation learning (IL) (Schaal, 1999) is an alternative approach to learn an optimal policy. In contrast to RL, IL has access to expert demonstrations, which are task-solving trajectories collected from experts who have mastered the task, and IL finds a policy that generates trajectories similar to these expert demonstrations. IL has been a long-studied problem and is attracting more attention recently, particularly in robotics (Duan et al., 2017; Stadie et al., 2017) and games (Ross et al., 2011). However, traditional IL methods rely on extensive feature engineering which makes their applicability quite limited.

Many approaches were proposed to overcome the difficulty of feature engineering in IL. Among them, the most successful approach is to use deep neural networks to learn representative features in an end-to-end manner (Wulfmeier et al., 2015; Finn et al., 2016a; Ho & Ermon, 2016; Fu et al., 2018). In particular, *generative adversarial imitation learning* (GAIL) (Ho & Ermon, 2016) is a state-of-the-art method that uses generative adversarial training to perform IL with deep neural networks. While the combination of IL and deep learning led to impressive performance improvement, it has introduced a new limitation regarding sample efficiency as training a deep neural network generally requires a large amount of data. This is a severe limitation in IL since collecting a large number of expert demonstrations can be expensive and time-consuming, and requires much expert effort. Moreover, the generative adversarial training procedure is known to be highly unstable (Mescheder et al., 2018), and this issue becomes more severe when only a small amount of data is available.

While expert demonstrations may be expensive, *non-expert* demonstrations collected from non-experts who have not mastered the task are often much cheaper to obtain. For instance, demonstrations from amateur-level players in the game of Go are much cheaper to obtain than those from master-level players. Leveraging additional information from a large number of non-expert demonstrations to improve IL is the key idea of semi-supervised inverse RL (SSIRL) (Valko et al., 2012)

and IRL from failure (IRLF) (Shiarlis et al., 2016). Both SSIRL and IRLF have shown to perform well for low-dimensional problems with proper feature engineering. However, they are extensions of traditional IL methods and are not capable of efficiently training deep neural networks which are needed for handling high-dimensional problems. Moreover, both of them rely on rather restrictive assumptions about data generating processes of non-expert demonstrations which we cannot control in practice.

In this paper, we propose a novel method to leverage non-expert demonstrations without the aforementioned weaknesses of SSIRL and IRLF. Our method is built upon the generative adversarial training procedure where we perform multiclass classification to learn discriminator functions with non-expert demonstrations regarded as being drawn from an extra class. Our method uses both expert and non-expert demonstrations in the discriminator learning objective, and this leads to a better feature representation of discriminator functions. We show that the minimax formulation commonly used in generative adversarial training does not guarantee the optimality of policies for our method, and we alternatively propose a modified optimization procedure which provides such a guarantee. We further show that our method learns an approximated solution to maximum entropy IRL with two reward functions. Experiments on benchmark continuous control tasks show that our method learns the expert policies faster and achieves more stable performances overall than GAIL when only a small number of expert demonstrations is available.

2 Related Work

IL has been a long-studied problem and there are many approaches to solve this problem, including behavior cloning (Pomerleau, 1988), occupancy measure matching (Syed et al., 2008) and IRL (Russell, 1998; Ng & Russell, 2000). Recently, IL methods that use a generative adversarial training procedure (Goodfellow et al., 2016) have gained a great deal of interest thanks to its effectiveness at training deep neural networks (Ho & Ermon, 2016; Finn et al., 2016a; Fu et al., 2018). Despite such success, deep neural networks are well-known to have poor data efficiency and require a large amount of data to train. For this reason, these methods may not perform well when only a small number of expert demonstrations are available.

Semi-supervised learning (SSL) (Chapelle et al., 2010) improves sample efficiency by utilizing a large amount of unlabeled data and it has shown promising results in deep learning (Ranzato & Szummer, 2008; Weston et al., 2012). While mainly developed for supervised learning, SSL can be applied to improve some IL methods as well. In particular, *semi-supervised IRL* (SSIRL) (Valko et al., 2012) improves the IRL method of Abbeel & Ng (2004) by using semi-supervised support vector machines to classify between expert demonstrations and trajectories generated by the agent. However, this approach is not suitable due to the difference in data generating processes between SSL and IRL. More specifically, SSL methods generally assume that an unlabeled dataset is a mixture of positive and negative samples. For SSIRL, this implies that an unlabeled demonstration dataset is a mixture of expert demonstrations and the agent's trajectories. However, collecting such an unlabeled dataset is quite difficult in practice since we do not know the agent's policies beforehand. This issue has been remedied to some extent by Audiffren et al. (2015) where the authors proposed using a manifold regularization technique which relies on a milder assumption on the unlabeled dataset. However, manifold regularization requires an appropriate similarity function to perform well. Moreover, both methods are unsuitable in high-dimensional problems due to its dependent on the linearity of reward functions and good feature engineering.

IRL from failure (IRLF) (Shiarlis et al., 2016) also utilizes an additional demonstration dataset which is assumed to consist of demonstrations collected by non-expert who *failed* to solve the task. Using this dataset, the authors proposed an IRL method that encourages the agent to be dissimilar to non-expert while learning an expert policy. While IRLF is technically sounded, collecting strictly failure demonstrations can be expensive on tasks such as autonomous driving where failures are catastrophic. Moreover, this method is still restricted by the linear reward assumption and is not applicable to train deep neural networks.

Using additional datasets to improve generative adversarial training was explored in the context of semi-supervised generative modeling (Salimans et al., 2016; Li et al., 2017a) and multi-modal generative modeling (Liu & Tuzel, 2016). However, if we want to apply these methods to our IL setting, then we are required to generate trajectories to imitate non-experts. This is inefficient since

generating trajectories requires interactions with environment and we would like to keep the number of such interactions as small as possible. In contrast, our method only learns the expert policy and only generate trajectories to imitate the expert.

Our proposal of using an additional dataset as an extra class in multiclass classification resembles the idea of *universum learning* (Vapnik, 1998; 2006; Zhang & LeCun, 2017). So far, universum learning has been applied only to supervised learning problems, especially for discriminative learning with support vector machines. Thus, our contribution may be regarded as the first attempt to apply the idea of universum learning to IL and also to generative adversarial learning.

3 BACKGROUND

In this section, we provide backgrounds of RL, IL, and GAIL.

3.1 REINFORCEMENT LEARNING (RL)

An RL problem is formulated as a discrete-time Markov decision process (MDP) which is defined by a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, p(\mathbf{s}'|\mathbf{s},\mathbf{a}), p_0(\mathbf{s}), r(\mathbf{s},\mathbf{a}), \gamma)$, where $\mathcal{S} \subseteq \mathbb{R}^{d_\mathbf{s}}$ is the (continuous) state space, $\mathbf{s} \in \mathcal{S}$ is a state, $\mathcal{A} \subseteq \mathbb{R}^{d_\mathbf{a}}$ is the (continuous) action space, $\mathbf{a} \in \mathcal{A}$ is an action, $p(\mathbf{s}'|\mathbf{s},\mathbf{a})$ is the transition probability density from $\mathbf{s}$ to $\mathbf{s}'$ when $\mathbf{a}$ is taken, $p_0(\mathbf{s})$ is the initial state probability density, $r(\mathbf{s},\mathbf{a})$ is the reward function, and $0 < \gamma \leqslant 1$ is the discount factor[1]. In each discrete time-step $t \geqslant 0$, an agent in a state $\mathbf{s}_t$ chooses an action $\mathbf{a}_t$ according to a policy $\pi(\mathbf{a}|\mathbf{s}_t)$, which is a conditional probability density. Then, the agent transits to a next state $\mathbf{s}_{t+1} \sim p(\mathbf{s}'|\mathbf{s}_t,\mathbf{a}_t)$ and receives a reward $r(\mathbf{s}_t,\mathbf{a}_t)$. We call a sequence of states and actions a trajectory $\boldsymbol{\tau}$. The goal of RL is to find an optimal policy that maximizes the expected discounted cumulative rewards (also called return) defined as

$$\mathbb{E}_{p_0(\mathbf{s}_0),\pi(\mathbf{a}_t|\mathbf{s}_t)_{t\geqslant 0},p(\mathbf{s}_{t+1}|\mathbf{s}_t,\mathbf{a}_t)_{t\geqslant 0}}\left[\sum_{t=0}^{\infty}\gamma^t r(\mathbf{s}_t,\mathbf{a}_t)\right] = \mathbb{E}_\pi\left[r(\mathbf{s},\mathbf{a})\right], \tag{1}$$

where the expectation is taken over the probability densities for all time steps and we use the notation $\mathbb{E}_\pi$ for brevity. When the policy is a parameterized function with parameter $\boldsymbol{\theta}$, a locally optimal policy can be found by optimization methods such as policy gradients (Williams, 1992).

We also use an equivalent formulation in terms of *occupancy measures* (Puterman, 1994; Altman, 1999; Syed et al., 2008). A state-action occupancy measure defines the expected discounted (unnormalized) visitation density of each state-action pair and is denoted by $\rho_\pi(\mathbf{s},\mathbf{a}) = \mathbb{E}_{p_0(\mathbf{s}_0),\pi(\mathbf{a}_t|\mathbf{s}_t)_{t\geqslant 0},p(\mathbf{s}_{t+1}|\mathbf{s}_t,\mathbf{a}_t)_{t\geqslant 0}}[\sum_{t=0}^{T}\gamma^t\delta(\mathbf{s}_t-\mathbf{s},\mathbf{a}_t-\mathbf{a})]$, where δ is the Dirac delta function[2]. An important property of the occupancy measure is that if it satisfies the Bellman flow constraints, $\int\rho_\pi(\mathbf{s}',\mathbf{a}')\mathrm{d}\mathbf{a}' = p_0(\mathbf{s}') + \gamma\iint p(\mathbf{s}'|\mathbf{s},\mathbf{a})\rho_\pi(\mathbf{s},\mathbf{a})\mathrm{d}\mathbf{s}\mathrm{d}\mathbf{a}$, then there is one-to-one correspondence between the occupancy measure and a policy given as $\pi(\mathbf{a}|\mathbf{s}) = \rho_\pi(\mathbf{s},\mathbf{a})/\rho_\pi(\mathbf{s})$, where $\rho_\pi(\mathbf{s}) = \int\rho_\pi(\mathbf{s},\mathbf{a})\mathrm{d}\mathbf{a}$ is a state occupancy measure. This property allows us to rewrite the RL objective to be maximized as $\mathbb{E}_\pi\left[r(\mathbf{s},\mathbf{a})\right] = \iint\rho_\pi(\mathbf{s},\mathbf{a})r(\mathbf{s},\mathbf{a})\mathrm{d}\mathbf{s}\mathrm{d}\mathbf{a}$.

The optimal policy of the above MDP is deterministic (Puterman, 1994) and a stochastic policy should be reduced to a deterministic policy at an optimum. However, a deterministic policy suffers from an exploration issue. In many tasks, it is beneficial to consider maximum entropy RL (Ziebart et al., 2008; 2010) whose optimal policy is a stochastic policy maximizing

$$\mathbb{E}_\pi\left[r(\mathbf{s},\mathbf{a})\right] + \beta\mathcal{H}(\pi), \tag{2}$$

where $\beta \geqslant 0$ and $\mathcal{H}(\pi) = -\mathbb{E}_\pi\left[\log\pi(\mathbf{a}|\mathbf{s})\right]$ is a discounted causal entropy (Ziebart et al., 2010). The advantage of maximum entropy RL is that it encourages exploration and allows the agent to find a better policy when compared to the standard RL formulation (Haarnoja et al., 2017).

3.2 IMITATION LEARNING (IL)

The goal of IL, or apprenticeship learning, is to learn a parameterized policy $\pi_{\boldsymbol{\theta}}$, with policy parameter $\boldsymbol{\theta}$, such that $\pi_{\boldsymbol{\theta}}$ exhibits the same behavior as an expert policy π_E. We assume that π_E is unknown.

[1] $\gamma = 1$ is only allowed for finite horizon setting.

[2] For discrete $\mathcal{S}$ and $\mathcal{A}$, δ is replaced by the indicator function.

We instead have access to expert demonstrations $\mathcal{D}_{\mathrm{E}} = \{(\mathbf{s}_i, \mathbf{a}_i)\}_{i=1}^N$ which are trajectories generated by executing the expert policy under an MDP. IL methods can be categorized into interactive methods and non-interactive methods. Interactive methods such as structured prediction (Ross et al., 2011) allow the agent to query for expert demonstrations during learning. Despite their strong theoretical guarantees and great empirical performances, these methods require the expert to be available during learning which is not always possible in reality. On the other hand, non-interactive methods, such as occupancy measure matching (Syed et al., 2008) and IRL (Ng & Russell, 2000; Abbeel & Ng, 2004), only require a pre-collected demonstration dataset for learning. In this paper, we focus on the non-interactive IL setting due to its high practicality.

3.3 Generative Adversarial Imitation Learning (GAIL)

GAIL (Ho & Ermon, 2016) is a state-of-the-art non-interactive IL method that performs occupancy measure matching to learn the parameterized policy. In occupancy measure matching (Syed et al., 2008), the policy parameter is learned to minimize a distance measure ℓ between occupancy measures of π_{E} and $\pi_{\boldsymbol{\theta}}$. More formally, occupancy measure matching methods solve an optimization problem $\min_{\boldsymbol{\theta}} \ell(\rho_{\pi_{\mathrm{E}}}, \rho_{\pi_{\boldsymbol{\theta}}}) - \beta\mathcal{H}(\pi_{\boldsymbol{\theta}})$, where $\mathcal{H}$ is the causal entropy regularizer with $\beta \geqslant 0$.

The key idea of GAIL is to use generative adversarial training to estimate the distance and to minimize the estimated distance. Briefly speaking, the distance measure is the Jensen-Shannon (JS) divergence defined as $\mathrm{JS}(\rho_{\pi_{\mathrm{E}}}, \rho_{\pi_{\boldsymbol{\theta}}}) = \frac{1}{2}\left(\mathrm{gKL}(\rho_{\pi_{\mathrm{E}}}||(\rho_{\pi_{\mathrm{E}}} + \rho_{\pi_{\boldsymbol{\theta}}})/2) + \mathrm{gKL}(\rho_{\pi_{\boldsymbol{\theta}}}||(\rho_{\pi_{\mathrm{E}}} + \rho_{\pi_{\boldsymbol{\theta}}})/2)\right)$, where $\mathrm{gKL}(\rho||q) = \iint \rho(\mathbf{s},\mathbf{a}) \log \frac{\rho(\mathbf{s},\mathbf{a})}{q(\mathbf{s},\mathbf{a})} \mathrm{d}\mathbf{s}\mathrm{d}\mathbf{a} - \iint \rho(\mathbf{s},\mathbf{a})\mathrm{d}\mathbf{s}\mathrm{d}\mathbf{a} + \iint q(\mathbf{s},\mathbf{a})\mathrm{d}\mathbf{s}\mathrm{d}\mathbf{a}$ is the generalized Kullback-Leibler (gKL) divergence defined for unnormalized densities[3] (Dikmen et al., 2015). Since both occupancy measures are unknown, the divergence is approximated via a binary classification problem: $\max_{\boldsymbol{\phi}} \mathbb{E}_{\pi_{\boldsymbol{\theta}}}[\log D_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})] + \mathbb{E}_{\pi_{\mathrm{E}}}[\log(1 - D_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}))]$. The function $D_{\boldsymbol{\phi}} : \mathcal{S} \times \mathcal{A} \to (0,1)$ is called a discriminator and is often parameterized by a deep neural network as $D_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}) = \exp(d_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}))/(\exp(d_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})) + 1)$. It can be shown that if the discriminator has infinite capacity, the global maximum of this binary classification problem corresponds to the JS divergence up to a constant. Based on this fact, GAIL minimizes an approximated JS divergence by solving the following minimax optimization problem:

$$\min_{\boldsymbol{\theta}} \max_{\boldsymbol{\phi}} \mathbb{E}_{\pi_{\boldsymbol{\theta}}}\left[\log D_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\right] + \mathbb{E}_{\pi_{\mathrm{E}}}\left[\log(1 - D_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}))\right] - \beta\mathcal{H}(\pi_{\boldsymbol{\theta}}). \tag{3}$$

In practice, this optimization problem is solved by alternately performing gradient ascent for $\boldsymbol{\phi}$ and gradient descent for $\boldsymbol{\theta}$. Notice that the gradient descent step for $\boldsymbol{\theta}$ is equivalent to performing policy gradient ascent with reward function $r(\mathbf{s},\mathbf{a}) = -\log(D_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}))$ and a causal entropy.

4 Imitation Learning with Non-expert Demonstrations

GAIL is an effective method that achieves state-of-the-art performance for high-dimensional IL problems. However, the generative adversarial training procedure described above is known to be highly unstable (Mescheder et al., 2018). This instability issue becomes more apparent when only a small amount of data is available. One of the reasons is due to inaccuracy of estimating the divergence via a discriminator learned by using a small number of demonstrations. A classical approach to improve binary classification is to use an unlabeled dataset in the context of SSL. However, as discussed previously, SSL is not suitable for IL due to its assumption of unlabeled data.

Note that the discriminator in GAIL is learned using two sets of data samples; expert demonstrations and trajectories collected by agent. For this reason, even when the number of expert demonstrations is small, GAIL may still learn well after observing a sufficiently large number of agent's trajectories. However, this is not desirable since in practice we often prefer data efficient methods that can learn well even with a small number of agent's trajectories.

To improve discriminator learning with a milder assumption on an additional demonstration dataset, we propose to perform multiclass classification using the additional dataset as a new class. In the following, we first formally describe our problem setting, and then present our IL method that learns a multiclass classifier for generative adversarial training. Lastly, we show an interesting relation

[3] An occupancy measure is not a density since it is not summed to one. Therefore, gKL is used instead.

between our multiclass classification objective and maximum entropy IRL with two reward functions (Ziebart et al., 2010).

4.1 Problem Setting

We consider an IL problem with an additional non-expert demonstration dataset as follows. We assume that expert demonstrations $\mathcal{D}_{\mathrm{E}} = \{(\mathbf{s}_i, \mathbf{a}_i)\}_{i=1}^{N}$ are trajectories collected by executing expert policy π_{E} in an MDP with reward function $r_{\mathrm{E}}(\mathbf{s}, \mathbf{a})$. We assume that π_{E} is an optimal policy, i.e., $\pi_{\mathrm{E}} = \mathrm{argmax}_{\pi} \mathbb{E}_{\pi}\left[r_{\mathrm{E}}(\mathbf{s}, \mathbf{a})\right]$. In addition, we also have non-expert demonstrations $\mathcal{D}_{\mathrm{N}} = \{(\mathbf{s}_j, \mathbf{a}_j)\}_{j=1}^{M}$ collected by executing non-expert policy π_{N} in the same MDP. The non-expert is assumed to be sub-optimal and gives worse expected rewards than the expert policy, i.e., $\mathbb{E}_{\pi_{\mathrm{N}}}\left[r_{\mathrm{E}}(\mathbf{s}, \mathbf{a})\right] < \mathbb{E}_{\pi_{\mathrm{E}}}\left[r_{\mathrm{E}}(\mathbf{s}, \mathbf{a})\right]$. We assume that $\mathcal{D}_{\mathrm{N}}$ is easier to obtain than $\mathcal{D}_{\mathrm{E}}$ and thus $M \gg N$. Our goal is to learn parameterized policy $\pi_{\boldsymbol{\theta}}$ that generates trajectories similar to those from π_{E} using both $\mathcal{D}_{\mathrm{E}}$ and $\mathcal{D}_{\mathrm{N}}$.

The main challenge for solving this problem is to efficiently leverage information brought by non-expert demonstrations to learn the expert policy. Under a very weak assumption that $\mathbb{E}_{\pi_{\mathrm{N}}}\left[r_{\mathrm{E}}(\mathbf{s}, \mathbf{a})\right] < \mathbb{E}_{\pi_{\mathrm{E}}}\left[r_{\mathrm{E}}(\mathbf{s}, \mathbf{a})\right]$, the non-expert policy can be totally unrelated to the expert policy and contains little to none of useful information that can improve learning. In the worst case, the non-expert policy could just be a random policy that randomly generates trajectories.

Therefore, to make efficient learning possible, we need a stronger but not too restrictive assumption on the non-expert demonstrations. Recall that both expert and non-expert demonstrations are collected under the same MDP, and this makes it intuitive to assume that state-action pairs of both the expert and non-expert lie on similar low-dimensional manifolds. To make this statement more explicit, we assume that there exists a feature map $\psi : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}^b$ such that state-action pairs from the expert and non-expert are linearly separable in this feature space. While this assumption itself still does not allow us to always escape from the random policy case, it allows us to utilize reasonably generated non-expert demonstrations to improve representation learning when training a discriminator. Furthermore, this assumption also naturally implies that the quality of learned features is improved as the non-expert policy becomes more related to the expert policy since their low-dimensional manifolds becomes more similar to each other. This is the opposite to IRLF where the failure demonstrations should make performance as worse as possible.

4.2 Multiclass Classification for Discriminator Learning

Our generative adversarial method alternates between the discriminator learning step and policy learning step. In the discriminator learning step, our goal is to learn a multiclass probabilistic classifier that classifies state-action pairs into three classes; the expert class with label $y = \mathrm{E}$, the non-expert class with label $y = \mathrm{N}$, and the agent class with label $y = \mathrm{A}$. Let the following softmax models be estimates of the class posterior of the three classes:

$$F_{\boldsymbol{\phi}}(\mathbf{s}, \mathbf{a}) = \frac{\exp(f_{\boldsymbol{\phi}}(\mathbf{s}, \mathbf{a}))}{Z_{\boldsymbol{\phi}}(\mathbf{s}, \mathbf{a}))}, \quad G_{\boldsymbol{\phi}}(\mathbf{s}, \mathbf{a}) = \frac{\lambda \exp(g_{\boldsymbol{\phi}}(\mathbf{s}, \mathbf{a}))}{Z_{\boldsymbol{\phi}}(\mathbf{s}, \mathbf{a}))}, \quad H_{\boldsymbol{\phi}}(\mathbf{s}, \mathbf{a}) = \frac{\exp(h_{\boldsymbol{\phi}}(\mathbf{s}, \mathbf{a}))}{Z_{\boldsymbol{\phi}}(\mathbf{s}, \mathbf{a}))}, \quad (4)$$

where $F_{\boldsymbol{\phi}}(\mathbf{s}, \mathbf{a})$ estimates the expert class posterior $p(y{=}\mathrm{E}|\mathbf{s}, \mathbf{a})$, $G_{\boldsymbol{\phi}}(\mathbf{s}, \mathbf{a})$ estimates the non-expert class posterior $p(y{=}\mathrm{N}|\mathbf{s}, \mathbf{a})$, $H_{\boldsymbol{\phi}}(\mathbf{s}, \mathbf{a})$ estimates the agent class posterior $p(y{=}\mathrm{A}|\mathbf{s}, \mathbf{a})$, and $Z_{\boldsymbol{\phi}}(\mathbf{s}, \mathbf{a}) = \exp(f_{\boldsymbol{\phi}}(\mathbf{s}, \mathbf{a})) + \lambda \exp(g_{\boldsymbol{\phi}}(\mathbf{s}, \mathbf{a})) + \exp(h_{\boldsymbol{\phi}}(\mathbf{s}, \mathbf{a}))$ is the normalization factor. $\boldsymbol{\phi}$ is the parameter to be learned and $\lambda > 0$ is a weight parameter specified by the user to control the influence of non-expert demonstrations. To learn parameter $\boldsymbol{\phi}$ for given policy parameter $\boldsymbol{\theta}$, we maximize the weighted log-likelihood objective defined as

$$\mathcal{L}_{\lambda}(\boldsymbol{\phi}, \boldsymbol{\theta}) = \mathbb{E}_{\pi_{\mathrm{E}}}\left[\log F_{\boldsymbol{\phi}}(\mathbf{s}, \mathbf{a})\right] + \lambda \mathbb{E}_{\pi_{\mathrm{N}}}\left[\log G_{\boldsymbol{\phi}}(\mathbf{s}, \mathbf{a})\right] + \mathbb{E}_{\pi_{\boldsymbol{\theta}}}\left[\log H_{\boldsymbol{\phi}}(\mathbf{s}, \mathbf{a})\right]. \quad (5)$$

Notice that the same λ is used for both $G_{\boldsymbol{\phi}}(\mathbf{s}, \mathbf{a})$ and $\mathcal{L}_{\lambda}(\boldsymbol{\phi}, \boldsymbol{\theta})$, and this choice is particularly important for our analyses in the following sections. We call the functions $f_{\boldsymbol{\phi}}$, $g_{\boldsymbol{\phi}}$, and $h_{\boldsymbol{\phi}}$ discriminators. For smooth and differentiable discriminators, this weighted-log-likelihood objective can be maximized by stochastic gradient ascent or its adaptive step-size variants. The expectations over π_{E} and π_{N} are approximated using mini-batch samples from $\mathcal{D}_{\mathrm{E}}$ and $\mathcal{D}_{\mathrm{N}}$, respectively, while the expectation over $\pi_{\boldsymbol{\theta}}$ is approximated using trajectories $\{(\mathbf{s}_k, \mathbf{a}_k)\}_{k=1}^{K}$ collected by executing policy $\pi_{\boldsymbol{\theta}}$.

We explicitly assume that the discriminators are represented by deep neural networks with shared hidden layers and three outputs where each output represents $f_{\boldsymbol{\phi}}$, $g_{\boldsymbol{\phi}}$, and $h_{\boldsymbol{\phi}}$. This parameterization

allows us to explicitly make use of our assumption of using common feature map ψ. That is, the shared hidden layers learns a feature map which allows the three classes to be linearly separable, or closed to be linearly separable. A large number of non-expert demonstrations allow us to accurately learn such a feature map which leads to more reliable discriminators and a better classification accuracy, as experimentally demonstrated in Section 5.

4.3 Occupancy Measure Matching with Multiclass Classifier

Our remaining task is to learn the parametrized policy in the policy learning step. Here, we present our occupancy measure matching method where a divergence is approximated from discriminators. First, we show that the discriminator learning objective approximates a JS divergence among three occupancy measures. This actually also implies that a commonly used minimax formulation is unsuitable since we do not recover the expert policy even in an infinite capacity setting. To cope with this problem, we propose a modified objective which allows us to recover the expert policy.

Let $\rho_{\pi_{\mathrm{E}}}(\mathbf{s},\mathbf{a})$, $\rho_{\pi_{\mathrm{N}}}(\mathbf{s},\mathbf{a})$, and $\rho_{\pi_{\boldsymbol{\theta}}}(\mathbf{s},\mathbf{a})$ be occupancy measures of the expert, non-expert, and agent, respectively. A JS divergence among the three occupancy measures with corresponding weights $\left(\frac{1}{2+\lambda}, \frac{\lambda}{2+\lambda}, \frac{1}{2+\lambda}\right)$ is defined as

$$\mathrm{JS}_{\lambda}(\rho_{\pi_{\mathrm{E}}}, \rho_{\pi_{\mathrm{N}}}, \rho_{\pi_{\boldsymbol{\theta}}}) = \frac{1}{2+\lambda}\left(\mathrm{gKL}(\rho_{\pi_{\mathrm{E}}}||\bar{q}) + \lambda \mathrm{gKL}(\rho_{\pi_{\mathrm{N}}}||\bar{q}) + \mathrm{gKL}(\rho_{\pi_{\boldsymbol{\theta}}}||\bar{q})\right), \tag{6}$$

where $\bar{q}(\mathbf{s},\mathbf{a}) = (\rho_{\pi_{\mathrm{E}}}(\mathbf{s},\mathbf{a}) + \lambda\rho_{\pi_{\mathrm{N}}}(\mathbf{s},\mathbf{a}) + \rho_{\pi_{\boldsymbol{\theta}}}(\mathbf{s},\mathbf{a}))/(2+\lambda)$ and $\mathrm{gKL}(\rho||\bar{q})$ is the generalized KL divergence. Note that Eq.(6) may be defined by either KL or gKL since the extra terms in gKL cancel out. Then, under the assumption that the discriminators have infinite capacity, the optimal parameter $\boldsymbol{\phi}^{\star}$ which maximizes $\mathcal{L}_{\lambda}(\boldsymbol{\phi},\boldsymbol{\theta})$ satisfies

$$\mathcal{L}_{\lambda}(\boldsymbol{\phi}^{\star},\boldsymbol{\theta}) = (2+\lambda)\mathrm{JS}_{\lambda}(\rho_{\pi_{\mathrm{E}}}, \rho_{\pi_{\mathrm{N}}}, \rho_{\pi_{\boldsymbol{\theta}}}) + \log\left(\lambda^{\lambda}(2+\lambda)^{-(2+\lambda)}\right). \tag{7}$$

The proof is given in Appendix A.1. This implies that minimizing $\mathcal{L}_{\lambda}(\boldsymbol{\phi}^{\star},\boldsymbol{\theta})$ is equivalent to minimizing $\mathrm{JS}_{\lambda}(\rho_{\pi_{\mathrm{E}}}, \rho_{\pi_{\mathrm{N}}}, \rho_{\pi_{\boldsymbol{\theta}}})$ up to a constant. This relation is similar to the one between binary classification and the JS divergence in GAIL, where a minimax formulation in Eq.(3) is used for occupancy measure matching.

However, solving the minimax problem, $\min_{\boldsymbol{\theta}}\max_{\boldsymbol{\phi}}\mathcal{L}_{\lambda}(\boldsymbol{\phi},\boldsymbol{\theta})$, is inappropriate in our method since the JS divergence in Eq.(6) is minimized by $\rho_{\pi_{\boldsymbol{\theta}}} = (\rho_{\pi_{\mathrm{E}}} + \lambda\rho_{\pi_{\mathrm{N}}})/(1+\lambda)$ instead by $\rho_{\pi_{\boldsymbol{\theta}}} = \rho_{\pi_{\mathrm{E}}}$. This minimizer corresponds to a policy mixture, $\pi_{\boldsymbol{\theta}}(\mathbf{a}|\mathbf{s}) = (\pi_{\mathrm{E}}(\mathbf{a}|\mathbf{s})\rho_{\pi_{\mathrm{E}}}(\mathbf{s}) + \lambda\pi_{\mathrm{N}}(\mathbf{a}|\mathbf{s})\rho_{\pi_{\mathrm{N}}}(\mathbf{s}))\,/\,(\rho_{\pi_{\mathrm{E}}}(\mathbf{s}) + \lambda\rho_{\pi_{\mathrm{N}}}(\mathbf{s}))$, which has a positive probability depending on λ to select non-expert actions instead of expert actions. Therefore, the above minimax formulation does not lead to occupancy measure matching in our method.

Here, we propose to perform occupancy measure matching by directly minimizing an approximation of $\mathrm{JS}(\rho_{\pi_{\mathrm{E}}}, \rho_{\pi_{\boldsymbol{\theta}}})$. We can approximate this divergence based on the fact that with $\lambda = 0$, we have $\mathrm{JS}_0(\rho_{\pi_{\mathrm{E}}}, \rho_{\pi_{\mathrm{N}}}, \rho_{\pi_{\boldsymbol{\theta}}}) = \mathrm{JS}(\rho_{\pi_{\mathrm{E}}}, \rho_{\pi_{\boldsymbol{\theta}}})$. Since minimizing $\mathcal{L}_{\lambda}(\boldsymbol{\phi}^{\star},\boldsymbol{\theta})$ is equivalent to minimizing $\mathrm{JS}_{\lambda}(\rho_{\pi_{\mathrm{E}}}, \rho_{\pi_{\mathrm{N}}}, \rho_{\pi_{\boldsymbol{\theta}}})$ for any λ, we can perform occupancy measure matching by solving $\min_{\boldsymbol{\theta}}\mathcal{L}_0(\boldsymbol{\phi}^{\star},\boldsymbol{\theta})$ with $\lambda = 0$. In practice, we instead minimize an approximation of the divergence based on an intermediate solution $\boldsymbol{\phi}$ since the discriminators do not have infinite capacity and finding a global optimum $\boldsymbol{\phi}^{\star}$ is not possible in practice. By ignoring constant terms in $\min_{\boldsymbol{\theta}}\mathcal{L}_0(\boldsymbol{\phi},\boldsymbol{\theta})$ and including a causal entropy regularization with $\beta \geqslant 0$, we obtain the following minimization problem for learning $\boldsymbol{\theta}$:

$$\min_{\boldsymbol{\theta}} \mathbb{E}_{\pi_{\boldsymbol{\theta}}}\left[\log H_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\right] - \beta\mathcal{H}(\pi_{\boldsymbol{\theta}}). \tag{8}$$

This minimization problem is equivalent to solving a maximum entropy RL problem with reward function $r(\mathbf{s},\mathbf{a}) = -\log H_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})$. Eq.(8) also resembles the minimization problem of GAIL in Eq.(3). The main difference between them is that we learn discriminator $H_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})$ by solving a multiclass classification problem, while GAIL learns discriminator $D_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})$ by solving a binary classification problem. Due to this similarity, we name our method *multiclass GAIL* (M-GAIL).

The algorithmic summary of our method is given in Appendix B. Our method alternates between updating $\boldsymbol{\phi}$ for maximization using an estimated gradient of $\mathcal{L}_{\lambda}(\boldsymbol{\phi},\boldsymbol{\theta})$ with $\lambda > 0$, and updating $\boldsymbol{\theta}$ for minimization using an estimated gradient of $\mathbb{E}_{\pi_{\boldsymbol{\theta}}}\left[\log H_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\right] - \beta\mathcal{H}(\pi_{\boldsymbol{\theta}})$. The computation complexity of our method is similar to GAIL, with a small additional cost of computing gradients of $G_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})$ in the discriminator learning step.

4.4 RELATION TO MAXIMUM ENTROPY INVERSE REINFORCEMENT LEARNING

Recently, Finn et al. (2016a) and Fu et al. (2018) showed an interesting result that a log-likelihood objective for solving a binary classification problem in generative adversarial training is related to a maximum entropy IRL objective (Ziebart et al., 2008; 2010). Here, we show a similar result that the weighted log-likelihood objective for solving a multiclass classification problem is related to a maximum entropy IRL objective with two reward functions. However, this relation is not equivalence and there is small discrepancy between the two objectives.

We consider the goal of learning the expert reward function $r_{\mathrm{E}}(\mathbf{s},\mathbf{a})$ and the non-expert reward function $r_{\mathrm{N}}(\mathbf{s},\mathbf{a})$ by parameterized functions $f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})$, and $g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})$, respectively. $r_{\mathrm{N}}(\mathbf{s},\mathbf{a})$ is a reward function whose π_{N} is optimal[4]. The parameter $\boldsymbol{\phi}$ of the two reward functions may be learned independently using maximum entropy IRL. However, we are interested in jointly learn both reward functions by maximizing the following objective:

$$\mathcal{M}_{\lambda}(\boldsymbol{\phi}) = \mathbb{E}_{p_{\mathrm{E}}(\boldsymbol{\tau})}\left[\log p_f(\boldsymbol{\tau})\right] + \lambda \mathbb{E}_{p_{\mathrm{N}}(\boldsymbol{\tau})}\left[\log p_g(\boldsymbol{\tau})\right]. \tag{9}$$

The first and second terms are the maximum entropy IRL objectives for learning the expert reward and the non-expert reward, respectively. The weight parameter $\lambda > 0$ controls the influence of the second objective. The trajectory density $p_f(\boldsymbol{\tau})$ is defined as $p_f(\boldsymbol{\tau}) \propto p_0(\mathbf{s}_0)\prod_{t=0}^{T} p(\mathbf{s}_{t+1}|\mathbf{s}_t,\mathbf{a}_t)\exp(f_{\boldsymbol{\phi}}(\mathbf{s}_t,\mathbf{a}_t))$ and $p_g(\boldsymbol{\tau})$ is also defined similarly. The following proposition shows that $\mathcal{L}_{\lambda}(\boldsymbol{\phi},\boldsymbol{\theta})$ is a surrogate of $\mathcal{M}_{\lambda}(\boldsymbol{\phi})$.

Proposition 1. *Under the assumption that $h_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}) = \pi_{\boldsymbol{\theta}}(\mathbf{a}|\mathbf{s})$ for fixed $\boldsymbol{\theta}$, the gradient of $\mathcal{L}_{\lambda}(\boldsymbol{\phi},\boldsymbol{\theta})$ is a biased gradient of $\mathcal{M}_{\lambda}(\boldsymbol{\phi})$. The bias vanishes when $\lambda = 0$ and $\pi_{\boldsymbol{\theta}}$ is an optimal maximum entropy policy for reward function $f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})$.*

The proof is based on the proof by Fu et al. (2018) for their adversarial IRL (AIRL) method, but with additional terms from non-expert demonstrations in our case. The detailed proof is given in Appendix A.2 and a proof sketch is as follow. We show that the gradient of $\mathcal{M}_{\lambda}(\boldsymbol{\phi})$ with importance weights which depend on $\pi_{\boldsymbol{\theta}}$ has the same form as the gradient of $\mathcal{L}_{\lambda}(\boldsymbol{\phi},\boldsymbol{\theta})$ with incorrect importance weights. Then, we show that the importance weights of the latter are corrected only when $\lambda = 0$ and the policy is optimal for reward $f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})$ under the maximum entropy RL framework.

This proposition indicates that local maxima of $\mathcal{M}_{\lambda}(\boldsymbol{\phi})$ are approximated by local maxima of $\mathcal{L}_{\lambda}(\boldsymbol{\phi},\boldsymbol{\theta})$, and that approximation errors decrease as λ decreases to zero. Due to approximation errors, this result is less attractive than the occupancy measure matching method which we proposed. However, it still suggests that we may use $f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})$ and $g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})$ as approximations of the reward functions which could be useful in the IRLF framework of Shiarlis et al. (2016) when the non-expert demonstrations are clearly failure demonstrations.

5 EXPERIMENTS

We compare M-GAIL against GAIL on four continuous control tasks from OpenAI gym (Brockman et al., 2016) with the PyBullet physics simulator (Coumans & Bai, 2018). These tasks are accompanied with true reward functions which are used for evaluation. We consider an episodic setting where a trajectory consists of 1000 transition samples. The discriminator and policy are neural networks with the same architecture used by GAIL (Ho & Ermon, 2016). We use Adam (Kingma & Ba, 2014) and TRPO (Schulman et al., 2016) to optimize the discriminator and policy, respectively. More details such as hyper-parameter settings are provided in Appendix C.

Setup: To collect expert demonstrations, we train a TRPO agent with the true reward function until convergence. Then, we treat a policy at the final iteration as expert and use it to collect two expert datasets with $N = 1000$ (i.e., 1 trajectory) and $N = 10000$. To collect non-expert demonstrations, we choose two intermediate policies which achieve approximately 50% and 70% performance of expert as non-expert. The 0% performance policy corresponds to initial random policy. Then, we use these non-expert policies to collect two non-expert datasets denoted by $\mathcal{D}_{\mathrm{N}}^{50\%}$ and $\mathcal{D}_{\mathrm{N}}^{70\%}$ with $M = 100000$. We train IL agents for 5 trials with different random seeds using fixed datasets. In each update iteration, the agents collect transition samples of size $K = 2000$ for learning.

[4]The optimal reward function always exists since for any trajectory there is at least one reward function that makes the trajectory optimal (Ng & Russell, 2000).

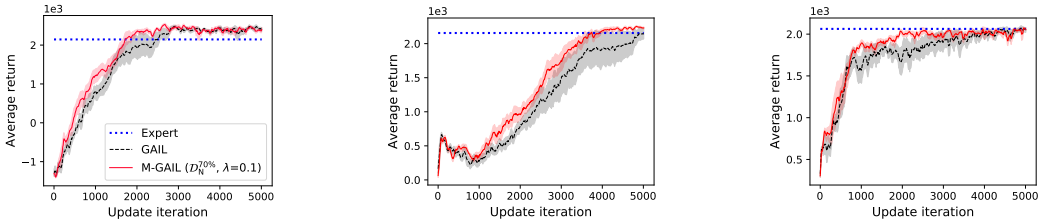

(a) HalfCheetah (b) Ant (c) Hopper

Figure 1: Returns averaged over 5 trials for $N = 1000$. The shaded area indicates standard error.

Table 1: The mean and standard error of total returns over 5 trials for $N = 1000$ (higher is better). The best and comparable methods according to paired t-test with p-value > 0.01 are in boldface.

Tasks	**GAIL**	**M-GAIL** with $\mathcal{D}_{\mathrm{N}}^{50\%}$		**M-GAIL** with $\mathcal{D}_{\mathrm{N}}^{70\%}$	
		$\lambda = 0.1$	$\lambda = 0.5$	$\lambda = 0.1$	$\lambda = 0.5$
Cheetah	1621.6(22.6)	1644.1(21.3)	1725.0(21.2)	**1788.2(20.7)**	**1787.2(19.7)**
Ant	1158.0(15.2)	1180.3(14.6)	**1342.8(15.1)**	**1351.5(14.9)**	1281.1(15.5)
Hopper	1744.1(8.8)	1702.3(9.9)	1821.4(8.3)	**1834.8(7.9)**	**1844.0(7.4)**
Walker	1036.9(5.6)	1056.4(5.3)	1046.9(5.3)	**1067.7(5.3)**	**1069.6(5.2)**

Results: We consider M-GAIL with $\lambda \in \{0.1, 0.5\}$. Three representative learning curves of GAIL and M-GAIL with $N = 1000$ are shown in Figure 1. The learning curves of four tasks with $N = 1000$ and $N = 10000$ are presented in Appendix D. This result shows that the proposed training procedure of M-GAIL indeed learns and converges to the expert policy without diverging. It also shows that M-GAIL learns faster than GAIL and requires collecting less agent's trajectories. While GAIL eventually learns the expert policy after collecting enough trajectories, it has quite unstable learning performance across trials as seen by large standard errors. We conjecture that a large number of non-expert demonstrations makes discriminator learning more stable and less prone to overfitting. On the other hand, discriminator learning in GAIL overfits to poor local optima in each trial, and it requires a sufficient amount of agent's trajectories to escape from these optima.

We also consider the total returns, computed as an area under learning curves divided by 5000, as an evaluation metric for comparing the overall performance given the same number of update iteration. Table 1 presents the total returns for $N = 1000$. We can see that M-GAIL with $\mathcal{D}_{\mathrm{N}}^{70\%}$ achieves higher total returns than GAIL on all tasks, with $\lambda = 0.1$ performs the best overall. M-GAIL with $\mathcal{D}_{\mathrm{N}}^{50\%}$ also tends to outperform GAIL, and it performs slightly worse than M-GAIL with $\mathcal{D}_{\mathrm{N}}^{70\%}$. This is consistent with our assumption that as π_{N} becomes more related to π_{E}, we can learn a better shared representation using non-expert demonstrations which leads to more accurate discriminator and faster learning. For M-GAIL with $\mathcal{D}_{\mathrm{N}}^{70\%}$ and $N = 1000$, the two values of λ achieve similar performance, and this this shows that M-GAIL is quite robust against different values λ for this setting. However, the choice of λ becomes more crucial for $N = 10000$ as shown in Appendix D. Nonetheless, M-GAIL still achieves higher total returns than GAIL overall.

6 Conclusion

We presented M-GAIL, a method that improves GAIL by using non-expert demonstration as an extra class in discriminator learning. Compared to related methods that use an additional dataset for IL, M-GAIL relies on a less restrictive assumption on the dataset and can efficiently train deep neural networks. Our experiments showed that M-GAIL effectively uses non-expert demonstrations to achieve faster and more stable learning than GAIL.

We developed M-GAIL based on JS-divergence. However, recent research suggests that other distances such as Wasserstein distance are more suitable when learning from image data (Arjovsky et al., 2017; Li et al., 2017b). Extending our method to use such distances is an important future work for applying our idea to solve image-based tasks such as Atari games (Mnih et al., 2015).

REFERENCES

Pieter Abbeel and Andrew Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the 21st International Conference on Machine Learning*, 2004. doi: 10.1145/1015330.1015430.

Eitan Altman. *Constrained Markov Decision Processes*. Chapman and Hall, 1999.

Martín Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein Generative Adversarial Networks. In *Proceedings of the 34th International Conference on Machine Learning*, pp. 214–223, 2017.

Julien Audiffren, Michal Valko, Alessandro Lazaric, and Mohammad Ghavamzadeh. Maximum entropy semi-supervised inverse reinforcement learning. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI*, pp. 3315–3321, 2015.

Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI Gym. *CoRR*, abs/1606.01540, 2016.

Olivier Chapelle, Bernhard Schlkopf, and Alexander Zien. *Semi-Supervised Learning*. The MIT Press, 1st edition, 2010. ISBN 0262514125, 9780262514125.

Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. `http://pybullet.org`, 2018.

Onur Dikmen, Zhirong Yang, and Erkki Oja. Learning the information divergence. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(7):1442–1454, 2015.

Yan Duan, Marcin Andrychowicz, Bradly C. Stadie, Jonathan Ho, Jonas Schneider, Ilya Sutskever, Pieter Abbeel, and Wojciech Zaremba. One-shot imitation learning. In *Advances in Neural Information Processing Systems 30*, pp. 1087–1098, 2017.

Chelsea Finn, Paul F. Christiano, Pieter Abbeel, and Sergey Levine. A connection between generative adversarial networks, inverse reinforcement learning, and energy-based models. *CoRR*, 2016a.

Chelsea Finn, Sergey Levine, and Pieter Abbeel. Guided cost learning: Deep inverse optimal control via policy optimization. In *Proceedings of the 33nd International Conference on Machine Learning*, pp. 49–58, 2016b.

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

Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.

Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. Reinforcement Learning with Deep Energy-Based Policies. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, 2017.

Jonathan Ho and Stefano Ermon. Generative Adversarial Imitation Learning. In *Advances in Neural Information Processing Systems 29*, pp. 4565–4573, 2016.

Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *CoRR*, abs/1412.6980, 2014.

Chongxuan Li, Taufik Xu, Jun Zhu, and Bo Zhang. Triple generative adversarial nets. In *Advances in Neural Information Processing Systems 30*, pp. 4091–4101, 2017a.

Yunzhu Li, Jiaming Song, and Stefano Ermon. Infogail: Interpretable imitation learning from visual demonstrations. In *Advances in Neural Information Processing Systems 30*, pp. 3815–3825, 2017b.

Ming-Yu Liu and Oncel Tuzel. Coupled generative adversarial networks. In *Advances in Neural Information Processing Systems 29*, pp. 469–477, 2016.

Lars M. Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *Proceedings of the 35th International Conference on Machine Learning*, pp. 3478–3487, 2018.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-Level Control Through Deep Reinforcement Learning. *Nature*, 518(7540):529–533, February 2015. ISSN 00280836.

Andrew Y. Ng and Stuart J. Russell. Algorithms for Inverse Reinforcement Learning. In *Proceedings of the 17th International Conference on Machine Learning*, pp. 663–670, 2000.

Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017.

Dean Pomerleau. ALVINN: an autonomous land vehicle in a neural network. In *Advances in Neural Information Processing Systems 1*, pp. 305–313, 1988.

Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York, NY, USA, 1st edition, 1994. ISBN 0-471-61977-9.

Marc'Aurelio Ranzato and Martin Szummer. Semi-supervised learning of compact document representations with deep networks. In *Proceedings of the Twenty-Fifth International Conference in Machine Learning*, pp. 792–799, 2008.

Stéphane Ross, Geoffrey J. 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, 2011.

Stuart J. Russell. Learning agents for uncertain environments (extended abstract). In *Proceedings of the Eleventh Annual Conference on Computational Learning Theory*, pp. 101–103, 1998.

Tim Salimans, Ian J. Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in Neural Information Processing Systems 29*, pp. 2226–2234, 2016.

Stefan Schaal. Is imitation learning the route to humanoid robots? *Trends in Cognitive Sciences*, 3 (6):233–242, 1999.

John Schulman, Sergey Levine, Philipp Moritz, Michael Jordan, and Pieter Abbeel. Trust Region Policy Optimization. In *Proceedings of the 32nd International Conference on Machine Learning*, 2015.

John Schulman, Philipp Moritz, Sergey Levine, Michael I. Jordan, and Pieter Abbeel. High-Dimensional Continuous Control Using Generalized Advantage Estimation. In *International Conference on Learning Representation*, 2016.

Kyriacos Shiarlis, João V. Messias, and Shimon Whiteson. Inverse reinforcement learning from failure. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pp. 1060–1068, 2016.

David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. Mastering the Game of Go Without Human Knowledge. *Nature*, 550(7676):354–359, October 2017. ISSN 00280836.

Bradly C. Stadie, Pieter Abbeel, and Ilya Sutskever. Third person imitation learning. In *International Conference on Learning Representation*, 2017.

Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning - an Introduction*. Adaptive computation and machine learning. MIT Press, 1998.

Umar Syed, Michael H. Bowling, and Robert E. Schapire. Apprenticeship learning using linear programming. In *Proceedings of the 25th International Conference on Machine Learning*, pp. 1032–1039, 2008. doi: 10.1145/1390156.1390286.

Michal Valko, Mohammad Ghavamzadeh, and Alessandro Lazaric. Semi-Supervised Apprenticeship Learning. In *Proceedings of the 10th European Workshop on Reinforcement Learning*, pp. 131–142, 2012.

Vladimir N. Vapnik. *Statistical Learning Theory*. Wiley-Interscience, 1998.

Vladimir N. Vapnik. *Estimation of Dependences Based on Empirical Data*. Springer-Verlag, 2006.

Jason Weston, Frédéric Ratle, Hossein Mobahi, and Ronan Collobert. Deep learning via semi-supervised embedding. In *Neural Networks: Tricks of the Trade - Second Edition*, pp. 639–655. 2012.

Ronald J. Williams. Simple Statistical Gradient-Following Algorithms for Connectionist Reinforcement Learning. *Machine Learning*, 8(3):229–256, 1992. doi: 10.1007/BF00992696.

Markus Wulfmeier, Peter Ondruska, and Ingmar Posner. Maximum Entropy Deep Inverse Reinforcement Learning. *CoRR*, 2015.

Xiang Zhang and Yann LeCun. Universum prescription: Regularization using unlabeled data. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pp. 2907–2913, 2017.

Brian D. Ziebart, Andrew L. Maas, J. Andrew Bagnell, and Anind K. Dey. Maximum entropy inverse reinforcement learning. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence*, pp. 1433–1438, 2008.

Brian D. Ziebart, J. Andrew Bagnell, and Anind K. Dey. Modeling Interaction via the Principle of Maximum Causal Entropy. In *Proceedings of the 27th International Conference on Machine Learning*, 2010.

A PROOFS

A.1 MAXIMUM OF WEIGHT LOG-LIKELIHOOD OBJECTIVE APPROXIMATES JENSEN-SHANNON DIVERGENCE

If we assume that the functions $F_{\boldsymbol{\phi}}$, $G_{\boldsymbol{\phi}}$, and $H_{\boldsymbol{\phi}}$ are non-parametric and have infinite capacity, we can rewrite the multiclass classification objective as

$$\mathcal{L}_\lambda(F, G, H) = \mathbb{E}_{\pi_{\mathrm{E}}}\left[\log F(\mathbf{s}, \mathbf{a})\right] + \lambda \mathbb{E}_{\pi_{\mathrm{N}}}\left[\log G(\mathbf{s}, \mathbf{a})\right] + \mathbb{E}_{\pi_{\boldsymbol{\theta}}}\left[\log H(\mathbf{s}, \mathbf{a})\right]. \tag{10}$$

Recall that $\mathbb{E}_\pi\left[r(\mathbf{s}, \mathbf{a})\right] = \iint \rho_\pi(\mathbf{s}, \mathbf{a}) r(\mathbf{s}, \mathbf{a}) \mathrm{d}\mathbf{s}\mathrm{d}\mathbf{a}$ where ρ_π is the occupancy measure of π. This objective is a concave function and we can find its global maximizer under constraints $F(\mathbf{s}, \mathbf{a}) \geqslant 0$, $G(\mathbf{s}, \mathbf{a}) \geqslant 0$, $H(\mathbf{s}, \mathbf{a}) \geqslant 0$ and $F(\mathbf{s}, \mathbf{a}) + G(\mathbf{s}, \mathbf{a}) + H(\mathbf{s}, \mathbf{a}) = 1, \forall \mathbf{s} \forall \mathbf{a}$, by using the method of Lagrange multipliers. The global maximizer corresponds to

$$F^\star(\mathbf{s}, \mathbf{a}) = \frac{\rho_{\pi_{\mathrm{E}}}(\mathbf{s}, \mathbf{a})}{q(\mathbf{s}, \mathbf{a})}, \quad G^\star(\mathbf{s}, \mathbf{a}) = \frac{\lambda \rho_{\pi_{\mathrm{N}}}(\mathbf{s}, \mathbf{a})}{q(\mathbf{s}, \mathbf{a})}, \quad H^\star(\mathbf{s}, \mathbf{a}) = \frac{\rho_{\pi_{\boldsymbol{\theta}}}(\mathbf{s}, \mathbf{a})}{q(\mathbf{s}, \mathbf{a})}, \tag{11}$$

where $q(\mathbf{s}, \mathbf{a}) = \rho_{\pi_{\mathrm{E}}}(\mathbf{s}, \mathbf{a}) + \lambda \rho_{\pi_{\mathrm{N}}}(\mathbf{s}, \mathbf{a}) + \rho_{\pi_{\boldsymbol{\theta}}}(\mathbf{s}, \mathbf{a})$. Let $\bar{q}(\mathbf{s}, \mathbf{a}) = \frac{q(\mathbf{s}, \mathbf{a})}{2+\lambda}$, by substituting these quantities back to $\mathcal{L}_\lambda$, we obtain the maximum objective value as

$$\begin{aligned}
\mathcal{L}_\lambda(F^\star, G^\star, H^\star) &= \mathbb{E}_{\pi_{\mathrm{E}}}\left[\log \frac{\rho_{\pi_{\mathrm{E}}}(\mathbf{s}, \mathbf{a})}{q(\mathbf{s}, \mathbf{a})}\right] + \lambda \mathbb{E}_{\pi_{\mathrm{N}}}\left[\log \frac{\lambda \rho_{\pi_{\mathrm{N}}}(\mathbf{s}, \mathbf{a})}{q(\mathbf{s}, \mathbf{a})}\right] + \mathbb{E}_{\pi_{\boldsymbol{\theta}}}\left[\log \frac{\rho_{\pi_{\boldsymbol{\theta}}}(\mathbf{s}, \mathbf{a})}{q(\mathbf{s}, \mathbf{a})}\right] \\
&= \mathbb{E}_{\pi_{\mathrm{E}}}\left[\log \frac{\rho_{\pi_{\mathrm{E}}}(\mathbf{s}, \mathbf{a})}{\bar{q}(\mathbf{s}, \mathbf{a})}\right] + \lambda \mathbb{E}_{\pi_{\mathrm{N}}}\left[\log \frac{\rho_{\pi_{\mathrm{N}}}(\mathbf{s}, \mathbf{a})}{\bar{q}(\mathbf{s}, \mathbf{a})}\right] + \mathbb{E}_{\pi_{\boldsymbol{\theta}}}\left[\log \frac{\rho_{\pi_{\boldsymbol{\theta}}}(\mathbf{s}, \mathbf{a})}{\bar{q}(\mathbf{s}, \mathbf{a})}\right] \\
&\quad + \lambda \log \lambda - (2 + \lambda) \log(2 + \lambda).
\end{aligned} \tag{12}$$

Notice that the first, second, and third terms look similar to Kullback-Leibler (KL) divergences which are defined as $\mathrm{KL}(p_1 || p_2) = \iint p_1(\mathbf{s}, \mathbf{a}) \log \frac{p_1(\mathbf{s}, \mathbf{a})}{p_1(\mathbf{s}, \mathbf{a})} \mathrm{d}\mathbf{s}\mathrm{d}\mathbf{a}$, for probability densities p_1 and p_2. However, the occupancy measures are unnormalized densities and theoretically we cannot use KL divergences to rewrite Eq.(12). For unnormalized densities $\rho(\mathbf{s}, \mathbf{a})$ and $\bar{q}(\mathbf{s}, \mathbf{a})$, we consider a generalized KL divergence defined as

$$\begin{aligned}
\mathrm{gKL}(\rho || \bar{q}) &= \iint \rho(\mathbf{s}, \mathbf{a}) \log \frac{\rho(\mathbf{s}, \mathbf{a})}{\bar{q}(\mathbf{s}, \mathbf{a})} \mathrm{d}\mathbf{s}\mathrm{d}\mathbf{a} - \iint \rho(\mathbf{s}, \mathbf{a}) \mathrm{d}\mathbf{s}\mathrm{d}\mathbf{a} + \iint \bar{q}(\mathbf{s}, \mathbf{a}) \mathrm{d}\mathbf{s}\mathrm{d}\mathbf{a} \\
&= \mathrm{KL}(\rho || \bar{q}) - \iint \rho(\mathbf{s}, \mathbf{a}) \mathrm{d}\mathbf{s}\mathrm{d}\mathbf{a} + \iint \bar{q}(\mathbf{s}, \mathbf{a}) \mathrm{d}\mathbf{s}\mathrm{d}\mathbf{a}.
\end{aligned} \tag{13}$$

We can show that for Eq.(12), gKL and KL are equivalent since the extra terms cancel out:

$$\begin{aligned}
&\mathrm{gKL}(\rho_{\pi_{\mathrm{E}}} || \bar{q}) + \lambda \mathrm{gKL}(\rho_{\pi_{\mathrm{N}}} || \bar{q}) + \mathrm{gKL}(\rho_{\pi_{\boldsymbol{\theta}}} || \bar{q}) \\
&\quad = \mathrm{KL}(\rho_{\pi_{\mathrm{E}}} || \bar{q}) + \lambda \mathrm{KL}(\rho_{\pi_{\mathrm{N}}} || \bar{q}) + \mathrm{KL}(\rho_{\pi_{\boldsymbol{\theta}}} || \bar{q}) \\
&\qquad - \iint \left(\rho_{\pi_{\mathrm{E}}}(\mathbf{s}, \mathbf{a}) + \lambda \rho_{\pi_{\mathrm{N}}}(\mathbf{s}, \mathbf{a}) + \rho_{\pi_{\boldsymbol{\theta}}}(\mathbf{s}, \mathbf{a})\right) \mathrm{d}\mathbf{s}\mathrm{d}\mathbf{a} + (2 + \lambda) \iint \bar{q}(\mathbf{s}, \mathbf{a}) \mathrm{d}\mathbf{s}\mathrm{d}\mathbf{a} \\
&\quad = \mathrm{KL}(\rho_{\pi_{\mathrm{E}}} || \bar{q}) + \lambda \mathrm{KL}(\rho_{\pi_{\mathrm{N}}} || \bar{q}) + \mathrm{KL}(\rho_{\pi_{\boldsymbol{\theta}}} || \bar{q}) \\
&\qquad - \iint \left(\rho_{\pi_{\mathrm{E}}}(\mathbf{s}, \mathbf{a}) + \lambda \rho_{\pi_{\mathrm{N}}}(\mathbf{s}, \mathbf{a}) + \rho_{\pi_{\boldsymbol{\theta}}}(\mathbf{s}, \mathbf{a})\right) \mathrm{d}\mathbf{s}\mathrm{d}\mathbf{a} + \iint q(\mathbf{s}, \mathbf{a}) \mathrm{d}\mathbf{s}\mathrm{d}\mathbf{a} \\
&\quad = \mathrm{KL}(\rho_{\pi_{\mathrm{E}}} || \bar{q}) + \lambda \mathrm{KL}(\rho_{\pi_{\mathrm{N}}} || \bar{q}) + \mathrm{KL}(\rho_{\pi_{\boldsymbol{\theta}}} || \bar{q}).
\end{aligned} \tag{14}$$

Then, the optimal value in Eq.(12) can be rewritten in terms of gKL divergences as

$$\begin{aligned}
\mathcal{L}_\lambda(F^\star, G^\star, H^\star) = &\ \mathrm{gKL}(\rho_{\pi_{\mathrm{E}}} || \bar{q}) + \lambda \mathrm{gKL}(\rho_{\pi_{\mathrm{N}}} || \bar{q}) + \mathrm{gKL}(\rho_{\pi_{\boldsymbol{\theta}}} || \bar{q}) \\
&+ \lambda \log \lambda - (2 + \lambda) \log(2 + \lambda)
\end{aligned} \tag{15}$$

Next, we consider a JS divergence among $\rho_{\pi_{\mathrm{E}}}$, $\rho_{\pi_{\mathrm{N}}}$, and $\rho_{\pi_{\boldsymbol{\theta}}}$ with corresponding weights $1/(2 + \lambda)$, $\lambda/(2 + \lambda)$, and $1/(2 + \lambda)$, which can be defined based on (generalized) KL divergence as

$$\mathrm{JS}_\lambda(\rho_{\pi_{\mathrm{E}}}, \rho_{\pi_{\mathrm{N}}}, \rho_{\pi_{\boldsymbol{\theta}}}) = \frac{1}{2 + \lambda}\left(\mathrm{gKL}(\rho_{\pi_{\mathrm{E}}} || \bar{q}) + \lambda \mathrm{gKL}(\rho_{\pi_{\mathrm{N}}} || \bar{q}) + \mathrm{gKL}(\rho_{\pi_{\boldsymbol{\theta}}} || \bar{q})\right). \tag{16}$$

Note that we use the fact that the sum of gKL and the sum of KL are equivalent to define the JS divergence. Alternatively, we may use a definition of JS in terms of a Shannon entropy, $\mathrm{JS}_\lambda(\rho_{\pi_\mathrm{E}}, \rho_{\pi_\mathrm{N}}, \rho_{\pi_{\boldsymbol{\theta}}}) = -\frac{1}{2+\lambda}\left(H(\rho_{\pi_\mathrm{E}}) + \lambda H(\rho_{\pi_\mathrm{N}}) + H(\rho_{\pi_{\boldsymbol{\theta}}})\right) + H\left(\frac{1}{2+\lambda}\left(\rho_{\pi_\mathrm{E}} + \lambda\rho_{\pi_\mathrm{N}} + \rho_{\pi_{\boldsymbol{\theta}}}\right)\right)$, where $H(\rho) = -\iint \rho_\pi(\mathbf{s},\mathbf{a}) \log \rho_\pi(\mathbf{s},\mathbf{a}) \mathrm{d}\mathbf{s}\mathrm{d}\mathbf{a}$, to obtain the same definition of JS divergence.

By comparing Eq.(15) and Eq.(16), we have that the maximum value of $\mathcal{L}_\lambda$ corresponds to JS_λ with a positive scaling and a constant shift:

$$\mathcal{L}_\lambda(F^\star, G^\star, H^\star) = (2+\lambda)\mathrm{JS}_\lambda(\rho_{\pi_\mathrm{E}}, \rho_{\pi_\mathrm{N}}, \rho_{\pi_{\boldsymbol{\theta}}}) + \log\left(\lambda^\lambda (2+\lambda)^{-(2+\lambda)}\right). \tag{17}$$

Notice that $\mathrm{JS}_\lambda(\rho_{\pi_\mathrm{E}}, \rho_{\pi_\mathrm{N}}, \rho_{\pi_{\boldsymbol{\theta}}})$ is a convex function of $\rho_{\pi_{\boldsymbol{\theta}}}$. By setting its derivative to zero, we have that the minimum of $\mathrm{JS}_\lambda(\rho_{\pi_\mathrm{E}}, \rho_{\pi_\mathrm{N}}, \rho_{\pi_{\boldsymbol{\theta}}})$ is obtained when $\rho_{\pi_{\boldsymbol{\theta}}} = \frac{\rho_{\pi_\mathrm{E}} + \lambda\rho_{\pi_\mathrm{N}}}{1+\lambda}$. Then, the corresponding mixture policy can be obtained by using the one-to-one correspondence between occupancy measure and policy given as $\pi(\mathbf{a}|\mathbf{s}) = \frac{\rho_\pi(\mathbf{s},\mathbf{a})}{\rho_\pi(\mathbf{s})}$.

A.2 Proposition 1

Here, we give the proof of Proposition 1. We follows the proof for AIRL by Fu et al. (2018), but with additional terms from non-expert demonstrations in our case. That is, we show that the gradient of a maximum entropy IRL objective proposed by Finn et al. (2016b) with specific importance weights only differs from the gradient of the discriminator learning objective in the importance weights. Similar to the case of AIRL, we only consider the case of non-discounted, finite horizon MDP.

Let $r_\mathrm{E}(\mathbf{s},\mathbf{a})$ be the expert reward function whose π_E is the optimal policy. Similarly, let $r_\mathrm{N}(\mathbf{s},\mathbf{a})$ be the non-expert reward function whose π_N is the optimal policy. The maximum entropy IRL objective for estimating the two rewards by $f_{\boldsymbol{\phi}}$ and $g_{\boldsymbol{\phi}}$ is

$$\mathcal{M}_\lambda(\boldsymbol{\phi}) = \mathbb{E}_{p_\mathrm{E}(\boldsymbol{\tau})}\left[\log p_f(\boldsymbol{\tau})\right] + \lambda \mathbb{E}_{p_\mathrm{N}(\boldsymbol{\tau})}\left[\log p_g(\boldsymbol{\tau})\right], \tag{18}$$

where $p_f(\boldsymbol{\tau}) \propto p_0(\mathbf{s}_0)\prod_{t=0}^T p(\mathbf{s}_{t+1}|\mathbf{s}_t,\mathbf{a}_t)\exp(f_{\boldsymbol{\phi}}(\mathbf{s}_t,\mathbf{a}_t))$ and $p_g(\boldsymbol{\tau})$ is defined similarly by $g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})$. $p_\mathrm{E}(\boldsymbol{\tau})$ and $p_\mathrm{N}(\boldsymbol{\tau})$ denotes (unknown) trajectory densities determined by expert and non-expert policies, respectively. The gradient of this objective w.r.t. $\boldsymbol{\phi}$ is

$$\begin{aligned}
\nabla_{\boldsymbol{\phi}}\mathcal{M}_\lambda(\boldsymbol{\phi}) &= \mathbb{E}_{p_\mathrm{E}(\boldsymbol{\tau})}\left[\sum_{t=0}^T \nabla_{\boldsymbol{\phi}} f_{\boldsymbol{\phi}}(\mathbf{s}_t,\mathbf{a}_t)\right] - \nabla_{\boldsymbol{\phi}} \log \int p_0(\mathbf{s}_0)\prod_{t=0}^T p(\mathbf{s}_{t+1}|\mathbf{s}_t,\mathbf{a}_t)\exp(f_{\boldsymbol{\phi}}(\mathbf{s}_t,\mathbf{a}_t))\mathrm{d}\boldsymbol{\tau} \\
&\quad + \lambda\mathbb{E}_{p_\mathrm{N}(\boldsymbol{\tau})}\left[\sum_{t=0}^T \nabla_{\boldsymbol{\phi}} g_{\boldsymbol{\phi}}(\mathbf{s}_t,\mathbf{a}_t)\right] - \lambda\nabla_{\boldsymbol{\phi}} \log \int p_0(\mathbf{s}_0)\prod_{t=0}^T p(\mathbf{s}_{t+1}|\mathbf{s}_t,\mathbf{a}_t)\exp(g_{\boldsymbol{\phi}}(\mathbf{s}_t,\mathbf{a}_t))\mathrm{d}\boldsymbol{\tau} \\
&= \mathbb{E}_{p_\mathrm{E}(\boldsymbol{\tau})}\left[\sum_{t=0}^T \nabla_{\boldsymbol{\phi}} f_{\boldsymbol{\phi}}(\mathbf{s}_t,\mathbf{a}_t)\right] - \mathbb{E}_{p_f(\boldsymbol{\tau})}\left[\sum_{t=0}^T \nabla_{\boldsymbol{\phi}} f_{\boldsymbol{\phi}}(\mathbf{s}_t,\mathbf{a}_t)\right] \\
&\quad + \lambda\mathbb{E}_{p_\mathrm{N}(\boldsymbol{\tau})}\left[\sum_{t=0}^T \nabla_{\boldsymbol{\phi}} g_{\boldsymbol{\phi}}(\mathbf{s}_t,\mathbf{a}_t)\right] - \lambda\mathbb{E}_{p_g(\boldsymbol{\tau})}\left[\sum_{t=0}^T \nabla_{\boldsymbol{\phi}} g_{\boldsymbol{\phi}}(\mathbf{s}_t,\mathbf{a}_t)\right].
\end{aligned} \tag{19}$$

Let $p_\mathrm{E}^t(\mathbf{s},\mathbf{a}) = \int p_\mathrm{E}(\boldsymbol{\tau})\mathrm{d}\boldsymbol{\tau}_{t'\neq t}$ be a state-action marginal density at time t obtained by marginalizing $p_\mathrm{E}(\boldsymbol{\tau})$ over states and actions at all time-step except at t. The state-action marginal densities $p_\mathrm{N}^t(\mathbf{s},\mathbf{a})$, $p_f^t(\mathbf{s},\mathbf{a})$, and $p_g^t(\mathbf{s},\mathbf{a})$ are defined similarly. We can rewrite the expectations over trajectories as a sum of expectations over state-action marginals as follows:

$$\begin{aligned}
\nabla_{\boldsymbol{\phi}}\mathcal{M}_\lambda(\boldsymbol{\phi}) = \sum_{t=0}^T \Big\{ &\mathbb{E}_{p_\mathrm{E}^t(\mathbf{s},\mathbf{a})}\Big[\nabla_{\boldsymbol{\phi}} f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\Big] - \mathbb{E}_{p_f^t(\mathbf{s},\mathbf{a})}\Big[\nabla_{\boldsymbol{\phi}} f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\Big] \\
&+ \lambda\mathbb{E}_{p_\mathrm{N}^t(\mathbf{s},\mathbf{a})}\Big[\nabla_{\boldsymbol{\phi}} g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\Big] - \lambda\mathbb{E}_{p_g^t(\mathbf{s},\mathbf{a})}\Big[\nabla_{\boldsymbol{\phi}} g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\Big]\Big\}.
\end{aligned} \tag{20}$$

Let $\mu_{\boldsymbol{\theta}}^t(\mathbf{s},\mathbf{a}) = p_\mathrm{E}^t(\mathbf{s},\mathbf{a}) + \lambda p_\mathrm{N}^t(\mathbf{s},\mathbf{a}) + p_{\boldsymbol{\theta}}^t(\mathbf{s},\mathbf{a})$ be a (unnormalized) sampling distribution where $p_{\boldsymbol{\theta}}^t(\mathbf{s},\mathbf{a})$ is a state-action marginal of a trajectory density $p_{\boldsymbol{\theta}}(\boldsymbol{\tau}) =$

$p_0(\mathbf{s}_0)\prod_{t=0}^{T} p(\mathbf{s}_{t+1}|\mathbf{s}_t,\mathbf{a}_t)\pi_{\boldsymbol{\theta}}(\mathbf{a}_t|\mathbf{s}_t)$. By rearranging terms and using importance weights for the expectations over $p_f^t(\mathbf{s},\mathbf{a})$ and $p_g^t(\mathbf{s},\mathbf{a})$, we can rewrite the gradient as

$$\nabla_{\boldsymbol{\phi}}\mathcal{M}_\lambda(\boldsymbol{\phi}) = \sum_{t=0}^{T}\Big\{\mathbb{E}_{p_{\mathrm{E}}^t(\mathbf{s},\mathbf{a})}\Big[\nabla_{\boldsymbol{\phi}} f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\Big] + \lambda\mathbb{E}_{p_{\mathrm{N}}^t(\mathbf{s},\mathbf{a})}\Big[\nabla_{\boldsymbol{\phi}} g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\Big] - \mathbb{E}_{\mu_{\boldsymbol{\theta}}^t(\mathbf{s},\mathbf{a})}\Big[\frac{p_f^t(\mathbf{s},\mathbf{a})}{\mu_{\boldsymbol{\theta}}^t(\mathbf{s},\mathbf{a})}\nabla_{\boldsymbol{\phi}} f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}) + \lambda\frac{p_g^t(\mathbf{s},\mathbf{a})}{\mu_{\boldsymbol{\theta}}^t(\mathbf{s},\mathbf{a})}\nabla_{\boldsymbol{\phi}} g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\Big]\Big\}. \tag{21}$$

Next, we show that $\nabla_{\boldsymbol{\phi}}\mathcal{L}_\lambda(\boldsymbol{\phi},\boldsymbol{\theta})$ has the same form as $\nabla_{\boldsymbol{\phi}}\mathcal{M}_\lambda(\boldsymbol{\phi})$ but with incorrect importance weights. First, we rewrite $\mathcal{L}_\lambda(\boldsymbol{\phi},\boldsymbol{\theta})$ using the state-action marginal densities as

$$\begin{aligned}\mathcal{L}_\lambda(\boldsymbol{\phi},\boldsymbol{\theta}) &= \mathbb{E}_{\pi_{\mathrm{E}}}\left[\log F_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\right] + \lambda\mathbb{E}_{\pi_{\mathrm{N}}}\left[\log G_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\right] + \mathbb{E}_{\pi_{\boldsymbol{\theta}}}\left[\log H_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\right]\\ &= \mathbb{E}_{p_{\mathrm{E}}(\boldsymbol{\tau})}\left[\sum_{t=0}^{T}\log F_{\boldsymbol{\phi}}(\mathbf{s}_t,\mathbf{a}_t)\right] + \lambda\mathbb{E}_{p_{\mathrm{N}}(\boldsymbol{\tau})}\left[\sum_{t=0}^{T}\log G_{\boldsymbol{\phi}}(\mathbf{s}_t,\mathbf{a}_t)\right] + \mathbb{E}_{p_{\boldsymbol{\theta}}(\boldsymbol{\tau})}\left[\sum_{t=0}^{T}\log H_{\boldsymbol{\phi}}(\mathbf{s}_t,\mathbf{a}_t)\right]\\ &= \sum_{t=0}^{T}\Big\{\mathbb{E}_{p_{\mathrm{E}}^t(\mathbf{s},\mathbf{a})}\left[\log F_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\right] + \lambda\mathbb{E}_{p_{\mathrm{N}}^t(\mathbf{s},\mathbf{a})}\left[\log G_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\right] + \mathbb{E}_{p_{\boldsymbol{\theta}}^t(\mathbf{s},\mathbf{a})}\left[\log H_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\right]\Big\},\end{aligned} \tag{22}$$

where in the second line we use the definition of the expected return for non-discounted finite-horizon case. Replacing $F_{\boldsymbol{\phi}}$, $G_{\boldsymbol{\phi}}$, and $H_{\boldsymbol{\phi}}$ by their parameterization in Eq.(4) with $h(\mathbf{s},\mathbf{a}) = \log\pi_{\boldsymbol{\theta}}(\mathbf{a}|\mathbf{s})$ gives

$$\begin{aligned}\mathcal{L}_\lambda(\boldsymbol{\phi},\boldsymbol{\theta}) &= \sum_{t=0}^{T}\Big\{\mathbb{E}_{p_{\mathrm{E}}^t(\mathbf{s},\mathbf{a})}\left[f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\right] + \lambda\mathbb{E}_{p_{\mathrm{N}}^t(\mathbf{s},\mathbf{a})}\left[g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}) + \log\lambda\right] + \mathbb{E}_{p_{\boldsymbol{\theta}}^t(\mathbf{s},\mathbf{a})}\left[\log\pi_{\boldsymbol{\theta}}(\mathbf{a}|\mathbf{s})\right]\\ &\quad - \mathbb{E}_{p_{\mathrm{E}}^t(\mathbf{s},\mathbf{a})+\lambda p_{\mathrm{N}}^t(\mathbf{s},\mathbf{a})+p_{\boldsymbol{\theta}}^t(\mathbf{s},\mathbf{a})}\left[\log\left(\exp(f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})) + \lambda\exp(g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})) + \pi_{\boldsymbol{\theta}}(\mathbf{a}|\mathbf{s})\right)\right]\Big\}\\ &= \sum_{t=0}^{T}\Big\{\mathbb{E}_{p_{\mathrm{E}}^t(\mathbf{s},\mathbf{a})}\left[f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\right] + \lambda\mathbb{E}_{p_{\mathrm{N}}^t(\mathbf{s},\mathbf{a})}\left[g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}) + \log\lambda\right] + \mathbb{E}_{p_{\boldsymbol{\theta}}^t(\mathbf{s},\mathbf{a})}\left[\log\pi_{\boldsymbol{\theta}}(\mathbf{a}|\mathbf{s})\right]\\ &\quad - \mathbb{E}_{\mu_{\boldsymbol{\theta}}^t(\mathbf{s},\mathbf{a})}\left[\log\left(\exp(f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})) + \lambda\exp(g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})) + \pi_{\boldsymbol{\theta}}(\mathbf{a}|\mathbf{s})\right)\right]\Big\},\end{aligned} \tag{23}$$

where we use $\log(\lambda\exp(g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}))) = g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}) + \log\lambda$ in the first equality and use $\mu_{\boldsymbol{\theta}}^t(\mathbf{s},\mathbf{a})$ defined above in the second equality. Then, the gradient of $\mathcal{L}_\lambda(\boldsymbol{\phi},\boldsymbol{\theta})$ w.r.t. $\boldsymbol{\phi}$ is given by

$$\begin{aligned}\nabla_{\boldsymbol{\phi}}\mathcal{L}_\lambda(\boldsymbol{\phi},\boldsymbol{\theta}) &= \sum_{t=0}^{T}\Big\{\mathbb{E}_{p_{\mathrm{E}}^t(\mathbf{s},\mathbf{a})}\left[\nabla_{\boldsymbol{\phi}} f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\right] + \lambda\mathbb{E}_{p_{\mathrm{N}}^t(\mathbf{s},\mathbf{a})}\left[\nabla_{\boldsymbol{\phi}} g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\right]\\ &\quad - \mathbb{E}_{\mu_{\boldsymbol{\theta}}^t(\mathbf{s},\mathbf{a})}\left[\nabla_{\boldsymbol{\phi}}\log\left(\exp(f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})) + \lambda\exp(g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})) + \pi_{\boldsymbol{\theta}}(\mathbf{a}|\mathbf{s})\right)\right]\Big\}\\ &= \sum_{t=0}^{T}\Big\{\mathbb{E}_{p_{\mathrm{E}}^t(\mathbf{s},\mathbf{a})}\left[\nabla_{\boldsymbol{\phi}} f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\right] + \lambda\mathbb{E}_{p_{\mathrm{N}}^t(\mathbf{s},\mathbf{a})}\left[\nabla_{\boldsymbol{\phi}} g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\right]\\ &\quad - \mathbb{E}_{\mu_{\boldsymbol{\theta}}^t(\mathbf{s},\mathbf{a})}\left[\frac{\exp(f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}))}{Z_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})}\nabla_{\boldsymbol{\phi}} f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}) + \lambda\frac{\exp(g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}))}{Z_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})}\nabla_{\boldsymbol{\phi}} g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\right]\Big\},\end{aligned} \tag{24}$$

where $Z_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}) = \exp(f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})) + \lambda\exp(g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})) + \pi_{\boldsymbol{\theta}}(\mathbf{a}|\mathbf{s})$. Let $p_{\boldsymbol{\theta}}^t(\mathbf{s}) = \int p_{\boldsymbol{\theta}}^t(\mathbf{s},\mathbf{a})\mathrm{d}\mathbf{a}$ be a state marginal of the policy trajectory density. By multiplying both the nominators and denominators of the third and fourth terms by $p_{\boldsymbol{\theta}}^t(\mathbf{s})$, we obtain

$$\begin{aligned}\nabla_{\boldsymbol{\phi}}\mathcal{L}_\lambda(\boldsymbol{\phi},\boldsymbol{\theta}) &= \sum_{t=0}^{T}\Big\{\mathbb{E}_{p_{\mathrm{E}}^t(\mathbf{s},\mathbf{a})}\left[\nabla_{\boldsymbol{\phi}} f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\right] + \lambda\mathbb{E}_{p_{\mathrm{N}}^t(\mathbf{s},\mathbf{a})}\left[\nabla_{\boldsymbol{\phi}} g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\right]\\ &\quad - \mathbb{E}_{\mu_{\boldsymbol{\theta}}^t(\mathbf{s},\mathbf{a})}\left[\frac{p_{\boldsymbol{\theta}}^t(\mathbf{s})\exp(f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}))}{p_{\boldsymbol{\theta}}^t(\mathbf{s})Z_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})}\nabla_{\boldsymbol{\phi}} f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}) + \lambda\frac{p_{\boldsymbol{\theta}}^t(\mathbf{s})\exp(g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}))}{p_{\boldsymbol{\theta}}^t(\mathbf{s})Z_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})}\nabla_{\boldsymbol{\phi}} g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})\right]\Big\},\end{aligned} \tag{25}$$

By comparing Eq.(21) and Eq.(25), we can see that $\nabla_{\boldsymbol{\phi}}\mathcal{L}_\lambda(\boldsymbol{\phi},\boldsymbol{\theta})$ and $\nabla_{\boldsymbol{\phi}}\mathcal{M}_\lambda(\boldsymbol{\phi})$ have the same form and they only differs in the importance weights in the third and fourth terms. The importance weights used in $\nabla_{\boldsymbol{\phi}}\mathcal{M}_\lambda(\boldsymbol{\phi})$ are correct and performing gradient ascent with $\nabla_{\boldsymbol{\phi}}\mathcal{M}_\lambda(\boldsymbol{\phi})$ leads to a local maxima of $\mathcal{M}_\lambda(\boldsymbol{\phi})$ (assuming exact computation of the expectations). On the other hand, the importance weights used in $\nabla_{\boldsymbol{\phi}}\mathcal{L}_\lambda(\boldsymbol{\phi},\boldsymbol{\theta})$ lead to bias in terms of the objective $\mathcal{M}_\lambda(\boldsymbol{\phi})$, and performing gradient ascent with $\nabla_{\boldsymbol{\phi}}\mathcal{L}_\lambda(\boldsymbol{\phi},\boldsymbol{\theta})$ does not lead to a local maxima of $\mathcal{M}_\lambda(\boldsymbol{\phi})$. Thus, the global maximizer of $\mathcal{L}_\lambda(\boldsymbol{\phi},\boldsymbol{\theta})$ does not recover the true reward functions regardless of $\boldsymbol{\theta}$.

The bias can be corrected only when $\lambda = 0$ and $\pi_{\boldsymbol{\theta}}$ is the optimal maximum entropy policy of a reward $f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})$. This is because when $\lambda = 0$, $\mathcal{M}_0(\boldsymbol{\phi})$ is the maximum entropy IRL objective for learning the expert reward function and $\mathcal{L}_0(\boldsymbol{\phi},\boldsymbol{\theta})$ us the binary classification objective of GAIL and AIRL. As shown by Fu et al. (2018), the gradients of these two objectives are equivalent under the assumption that $\pi_{\boldsymbol{\theta}}$ is the optimal maximum entropy policy for reward $r(\mathbf{s},\mathbf{a}) = f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})$.

We emphasize that for our case it is insufficient to only assume that $\pi_{\boldsymbol{\theta}}$ is the optimal maximum entropy policy without assuming that $\lambda = 0$. This is because to make the nominators equal, we need that the two equalities, $p^t_{\boldsymbol{\theta}}(\mathbf{s})\exp(f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})) = p^t_f(\mathbf{s},\mathbf{a})$ and $p^t_{\boldsymbol{\theta}}(\mathbf{s})\exp(g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})) = p^t_g(\mathbf{s},\mathbf{a})$, hold. The former holds when $\pi_{\boldsymbol{\theta}}$ is (maximum entropy) optimal for a reward $f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})$, while the latter holds when $\pi_{\boldsymbol{\theta}}$ is (maximum entropy) optimal for a reward $g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})$. Both cannot hold at the same time unless $f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}) = \alpha g_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})$ with a positive scaling α, which implies non-expert is expert himself. Moreover, to have that the partition functions are equivalent, $p^t_{\boldsymbol{\theta}}(\mathbf{s})Z_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}) = \mu^t_{\boldsymbol{\theta}}(\mathbf{s},\mathbf{a})$, we also need the policy $\pi_{\boldsymbol{\theta}}$ to be optimal for a reward mixture. These contradictions imply that the two gradients can be made identical only when $\lambda = 0$.

B PSEUDOCODE

Algorithm 1 Multiclass Generative Adversarial Imitation Learning (M-GAIL)

1: **Input:** Demonstration datasets $\mathcal{D}_{\mathrm{E}}$ and $\mathcal{D}_{\mathrm{N}}$, initial parameters $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$, weight parameter λ.
2: **while** not converge **do**
3: Collect trajectory samples $\{(\mathbf{s}_k,\mathbf{a}_k)\}_{k=1}^K$ using $\pi_{\boldsymbol{\theta}}(\mathbf{a}|\mathbf{s})$.
4: Sample mini-batches data from $\mathcal{D}_{\mathrm{E}}$ and $\mathcal{D}_{\mathrm{N}}$.
5: Update $\boldsymbol{\phi}$ to maximize $\mathcal{L}_\lambda(\boldsymbol{\phi},\boldsymbol{\theta})$ in Eq.(5), by e.g., Adam (Kingma & Ba, 2014).
6: Update $\boldsymbol{\theta}$ to solve the RL problem in Eq.(8), by e.g., TRPO (Schulman et al., 2015).
7: **end while**
8: **Output:** Parameterized policy $\pi_{\boldsymbol{\theta}}$.

C MORE EXPERIMENT DETAILS

We implement GAIL and M-GAIL with PyTorch deep learning framework (Paszke et al., 2017) (Our code will be publicly available.). Both GAIL and M-GAIL use a common setting as follows. The discriminator and policy are represented by neural networks with two hidden layers and 100 hyperbolic tangent units in each layer, as proposed by the original GAIL paper (Ho & Ermon, 2016). We use a Gaussian policy with state-independent and diagonal covariance for all tasks. The network parameters are initialized randomly. The discriminator is optimized by Adam (Kingma & Ba, 2014) with step-size 3×10^{-4}, $\beta_1 = 0.9$, and $\beta_2 = 0.999$. In each update iteration, we sample mini-batch of size 64 from $\mathcal{D}_{\mathrm{E}}$ and $\mathcal{D}_{\mathrm{N}}$ to update the discriminator.

We tried to make the implementation of GAIL and M-GAIL as close as possible. In particular, we let $f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a}) = 0$ be a constant function in M-GAIL (This makes $\exp(f_{\boldsymbol{\phi}}(\mathbf{s},\mathbf{a})) = 1$). That is, the discriminator networks give two outputs; $g_{\boldsymbol{\phi}}$ and $h_{\boldsymbol{\phi}}$. This choice is valid since using three estimators to estimate three class posteriors is over-parameterize and we only need two estimators to compute an estimate of the third class posterior. With this implementation choice, M-GAIL with $\lambda = 0$ is exactly the same as GAIL where $h_{\boldsymbol{\phi}}$ in M-GAIL is the same as $d_{\boldsymbol{\phi}}$ in GAIL.

The policy is optimized by trust-region policy optimization (TRPO) with generalized advantage estimation (GAE) (Schulman et al., 2016) with KL bound $\epsilon = 0.01$ and damping coefficient for the Fisher information matrix 0.1. We do not add causal entropy and set $\beta = 0$ for GAIL and M-GAIL. For GAE, we use the above network architecture to learn the value function with $\gamma = 0.995$

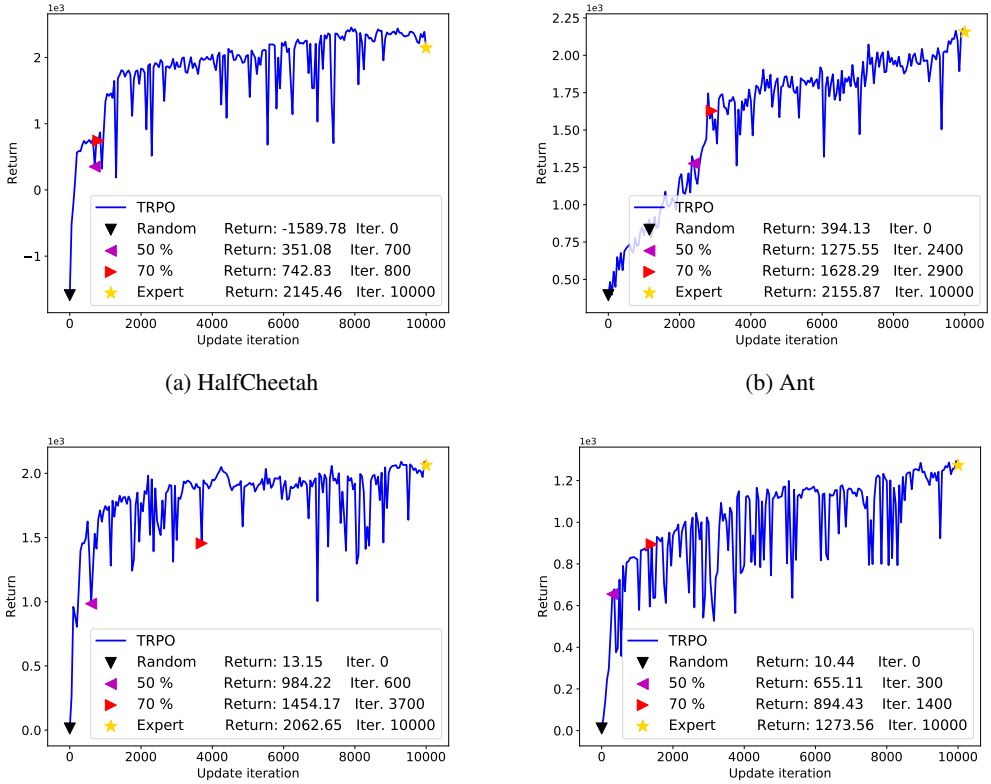

(a) HalfCheetah

(b) Ant

(c) Hopper

(d) Walker2D

Figure 2: The return in each update iteration of TRPO expert agent. The return is computed using Gaussian policy with exploration noise. "Random" is the initial random policy. "Iter." denotes the update iteration to achieve the corresponding return values.

and $\lambda_{\mathrm{GAE}} = 0.97$. In each update iteration, we use Adam, with step-size 3×10^{-4}, $\beta_1 = 0.9$, $\beta_2 = 0.999$, and ℓ_2 regularizer of 10^{-3}, to update the value function network for 3 epochs with minibatch size 128. The GAE value is standardized to have zero mean and unit variance for TRPO update. In each update iteration, we collect transition samples of size $K = 2000$ using the current policy. This hyper-parameter setting is chosen for TRPO since it works well for learning the expert policy, as described below.

We consider four tasks; Half-Cheetah, Hopper, Walker2D, and Ant, simulated by PyBullet physics simulator (Coumans & Bai, 2018). (At the time of submission, OpenAI gym with Pybullet physics does not have the Swimmer task, and we could not successfully train a Humanoid expert agent. The two inverted-pendulum tasks are too simple with all methods perform equally well.) To train the expert policy, we use TRPO with GAE with the above setting, except for Walker2D where we require an entropy regularizer of $\beta = 0.0001$ to learn a walking policy. Without this regularizer, the agent simply stands still to receive a reward of 1 in each time-step. The same scenario does not happen when we train the IL agents and we do not include the entropy regularizer for GAIL and M-GAIL. For each task, we train a TRPO agent for 5 trials with different random seeds and choose the best policy among 5 trials that gives the highest return at the 10000-th iteration as expert. The return of the chosen TRPO agent, as well as the returns at approximately 50% and 70% of the expert's return are shown in Figure 2. We can clearly see that non-expert policies can be learned using much less samples than the expert policy.

After obtaining the expert and non-expert policies, we use the learned Gaussian policies to collect demonstration trajectories with exploration noise. The same datasets are used in all 5 trials of GAIL and M-GAIL training.

Table 2: The mean and standard error of total returns over 5 trials for $N = 1000$ and $N = 10000$. The best and comparable methods according to paired t-test with p-value > 0.01 are in boldface.

Tasks (N)	**GAIL**	**M-GAIL** with $\mathcal{D}_{\mathrm{N}}^{50\%}$		**M-GAIL** with $\mathcal{D}_{\mathrm{N}}^{70\%}$	
		$\lambda = 0.1$	$\lambda = 0.5$	$\lambda = 0.1$	$\lambda = 0.5$
Cheetah (1000)	1621.6(22.6)	1644.1(21.3)	1725.0(21.2)	**1788.2(20.7)**	**1787.2(19.7)**
Ant (1000)	1158.0(15.2)	1180.3(14.6)	**1342.8(15.1)**	**1351.5(14.9)**	1281.1(15.5)
Hopper (1000)	1744.1(8.8)	1702.3(9.9)	1821.4(8.3)	**1834.8(7.9)**	**1844.0(7.4)**
Walker (1000)	1036.9(5.6)	1056.4(5.3)	1046.9(5.3)	**1067.7(5.3)**	**1069.6(5.2)**
Cheetah (10000)	**1502.0(21.3)**	1444.0(21.0)	**1477.4(20.5)**	**1503.7(20.6)**	1391.4(21.5)
Ant (10000)	1240.6(15.3)	1295.2(15.7)	1308.0(15.2)	1194.6(14.4)	**1412.0(15.4)**
Hopper (10000)	1679.3(8.9)	1786.2(8.5)	1775.8(8.1)	1824.3(8.0)	**1849.8(7.9)**
Walker (10000)	1031.7(5.1)	992.6(5.5)	1041.6(5.2)	1038.0(5.3)	**1060.8(5.0)**

D MORE EXPERIMENTAL RESULTS

We train GAIL and M-GAIL agents for 5 trials with different random seeds. All trials use the same datasets collected by the TRPO agent as explained previously. We evaluate the learned policies by generating 10 test trajectories using the learned policies to select actions without exploration noise. Then, we compute the averaged return over 10 trajectories using the true reward function without discount factor, and this gives us a learning curve for each trial. The mean and standard error of these learning curves over 5 trials for $N = 1000$ and $N = 10000$ are shown in Figure 3. Table 2 shows the total return (area under learning curves divided by 5000) of GAIL and M-GAIL with $N = 1000$ and $N = 10000$. The comparison between GAIL and M-GAIL for $N = 10000$ is similar to that for $N = 1000$ in the main text, except that GAIL performs quite well on HalfCheetah and achieves comparable performance.

We can also see that M-GAIL with $\lambda = 0.5$ performs better overall than M-GAIL with $\lambda = 0.1$ for $N = 10000$. A possible explanation of this phenomena is that λ can be regarded as a class prior of the non-expert class which determines the ratio of non-expert samples against samples from the other two classes. This ratio becomes close to one when the number of expert demonstrations increases while the number of non-expert demonstrations remains constant. This suggests that an appropriate value of the weight parameter should depend on the number of demonstrations. Developing a systematic procedure to select this tuning parameter is one of our future work.

Figure 4 shows the learning curves after we let the agents train for more update iteration. This result confirms that all methods converge to the expert policy without diverging. We only report the results with $N = 1000$ until the 5000-th update iteration of GAIL and the best performing M-GAIL in the main text for better clarity.

We can also see that for HalfCheetah, the total returns of GAIL and M-GAIL with $N = 10000$ is significantly worse than those with $N = 1000$. This is quite unintuitive since using more samples should yield faster learning. A possible explanation is that for HalfCheetah the expert policy may generate relatively poor trajectories which have poor performance, and the expert dataset with $N = 10000$ includes these trajectories which makes learning more difficult. Nonetheless, both GAIL and M-GAIL converge to the expert policy as seen in Figure 4b.

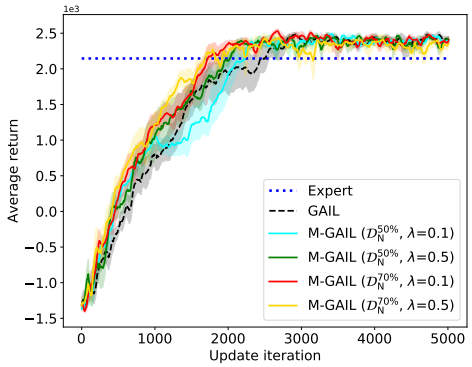

(a) HalfCheetah ($N = 1000$)

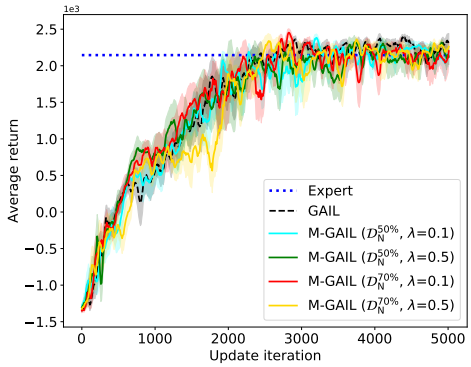

(b) HalfCheetah ($N = 10000$)

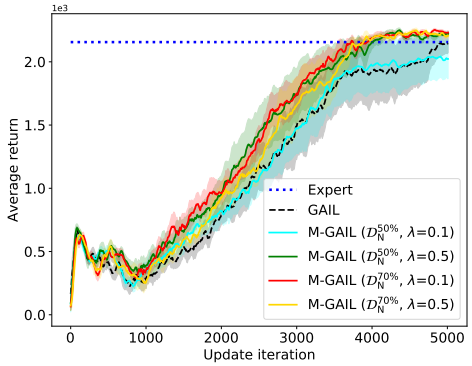

(c) Ant ($N = 1000$)

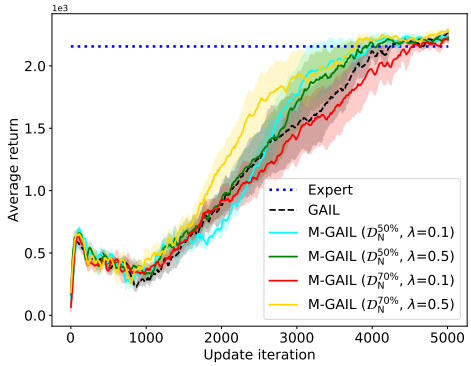

(d) Ant ($N = 10000$)

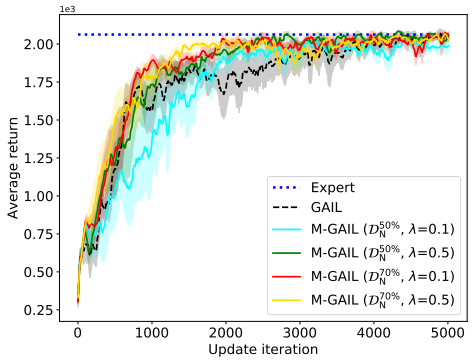

(e) Hopper ($N = 1000$)

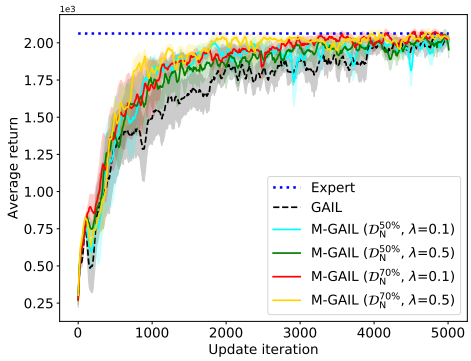

(f) Hopper ($N = 1000$)

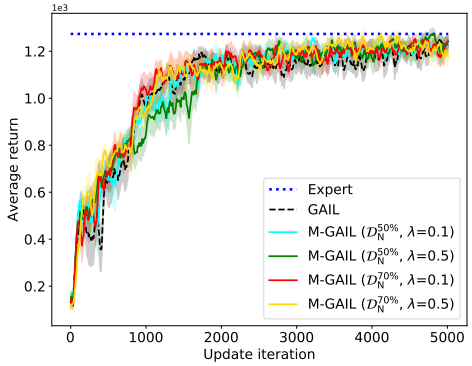

(g) Walker2D ($N = 1000$)

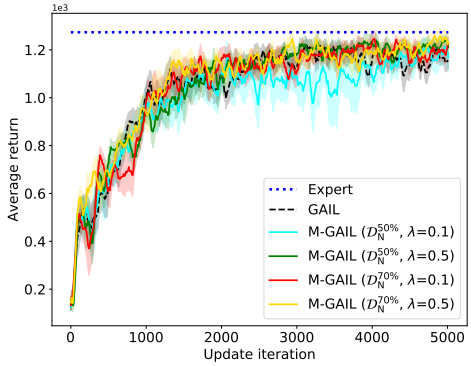

(h) Walker2D ($N = 10000$)

Figure 3: Returns averaged over 5 trials where the shaded area indicates standard error. The results for $N = 1000$ are in the left column and for $N = 10000$ in the right column. Returns are computed without exploration noise and discount factor.

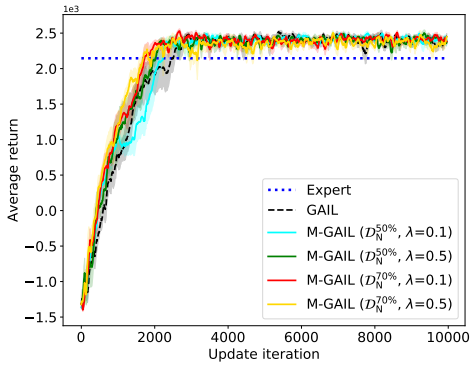

(a) HalfCheetah ($N = 1000$)

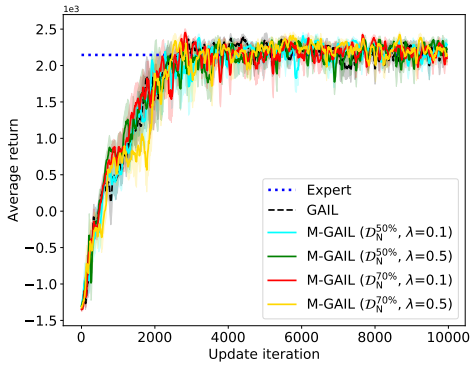

(b) HalfCheetah ($N = 10000$)

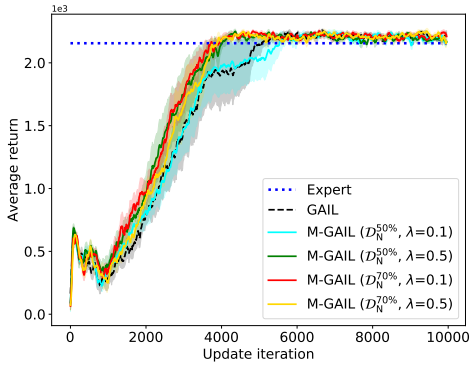

(c) Ant ($N = 1000$)

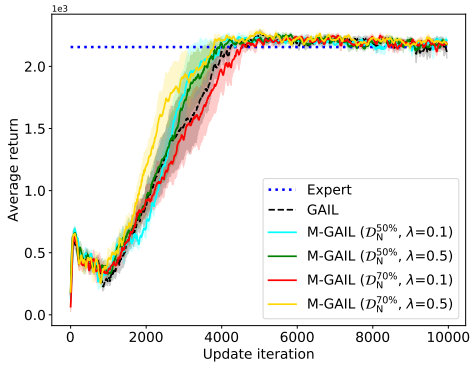

(d) Ant ($N = 10000$)

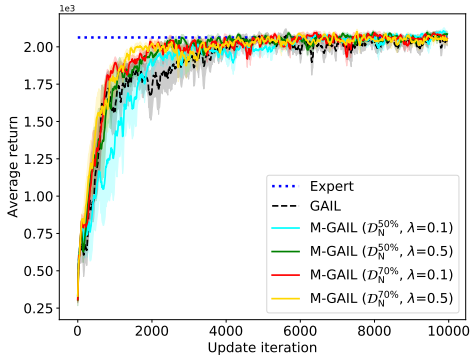

(e) Hopper ($N = 1000$)

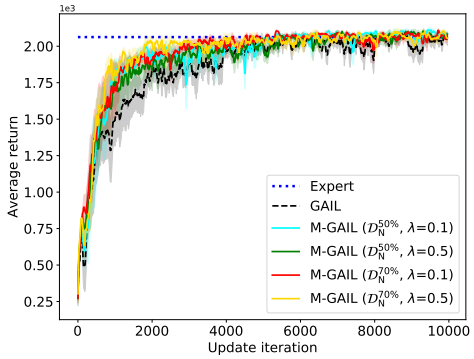

(f) Hopper ($N = 10000$)

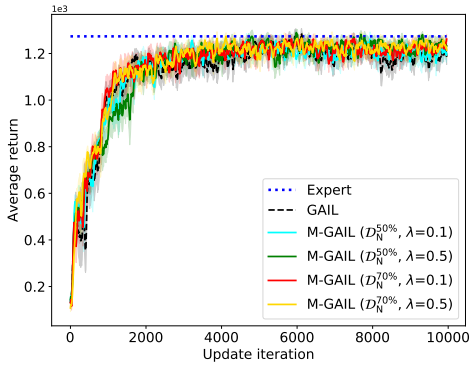

(g) Walker2D ($N = 1000$)

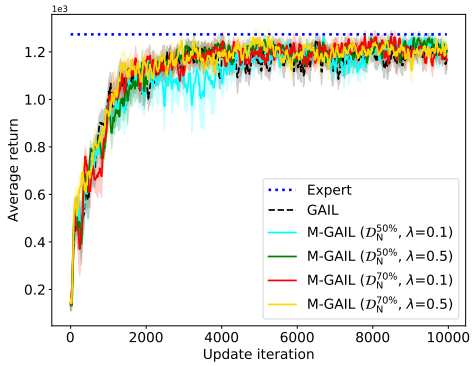

(h) Walker2D ($N = 10000$)

Figure 4: The results in Figure 3 with a longer update iteration.