

VILD: VARIATIONAL IMITATION LEARNING WITH DIVERSE-QUALITY DEMONSTRATIONS

Anonymous authors

Paper under double-blind review

ABSTRACT

The goal of imitation learning (IL) is to learn a good policy from high-quality demonstrations. However, the quality of demonstrations in reality can be diverse, since it is easier and cheaper to collect demonstrations from a mix of experts and amateurs. IL in such situations can be challenging, especially when the level of demonstrators' expertise is unknown. We propose a new IL paradigm called Variational Imitation Learning with Diverse-quality demonstrations (VILD), where we explicitly model the level of demonstrators' expertise with a probabilistic graphical model and estimate it along with a reward function. We show that a naive estimation approach is not suitable to large state and action spaces, and fix this issue by using a variational approach that can be easily implemented using existing reinforcement-learning methods. Experiments on continuous-control benchmarks demonstrate that VILD outperforms state-of-the-art methods. Our work enables scalable and data-efficient IL under more realistic settings than before.

1 INTRODUCTION

The goal of sequential decision making is to learn a policy that makes good decisions (Puterman, 1994). As an important branch of sequential decision making, imitation learning (IL) (Schaal, 1999) aims to learn such a policy from demonstrations (i.e., sequences of decisions) collected from experts. However, high-quality demonstrations can be difficult to obtain in reality, since such experts may not always be available and sometimes are too costly (Osa et al., 2018). This is especially true when the quality of decisions depends on specific domain-knowledge not typically available to amateurs; e.g., in applications such as robot control (Osa et al., 2018) and autonomous driving (Silver et al., 2012).

In practice, demonstrations are often diverse in quality, since it is cheaper to collect demonstrations from mixed demonstrators, containing both experts and amateurs (Audiffren et al., 2015). Unfortunately, IL in such settings tends to perform poorly, since low-quality demonstrations often negatively affect the performance of IL methods (Shiarlis et al., 2016). For example, amateurs' demonstrations for robotics can be cheaply collected via a robot simulation (Mandlekar et al., 2018), but such demonstrations may cause damages to the robot which is catastrophic in the real-world (Shiarlis et al., 2016). Similarly, demonstrations for autonomous driving can be collected from drivers in public roads (Fridman et al., 2017), which may contain traffic-accident demonstrations. Learning a self-driving car from these low-quality demonstrations may cause traffic accidents.

When the level of demonstrators' expertise is known, multi-modal IL (MM-IL) can be used to learn a good policy with diverse-quality demonstrations (Li et al., 2017; Hausman et al., 2017; Wang et al., 2017). Specifically, MM-IL aims to learn a multi-modal policy, where each mode of the policy represents the decision making of each demonstrator. When knowing the level of demonstrators' expertise, good policies can be obtained by selecting modes that correspond to the decision making of high-expertise demonstrators. However, in practice, it is difficult to truly determine the level of demonstrators' expertise beforehand. Without knowing the level of expertise, it is difficult to distinguish the decision making of experts and amateurs, and learning a good policy is challenging.

To overcome the issue of MM-IL, pioneer works have proposed to estimate the quality of each demonstration using auxiliary information from experts (Audiffren et al., 2015; Wu et al., 2019; Brown et al., 2019). Specifically, Audiffren et al. (2015) inferred the demonstration quality using similarities between diverse-quality demonstrations and high-quality demonstrations, where the latter

are collected in a small number from experts. In contrast, Wu et al. (2019) proposed to estimate the demonstration quality using a small number of demonstrations with confidence scores. Namely, the score value given by an expert is proportion to the demonstration quality. Similarly, the demonstration quality can be estimated by ranked demonstrations, where ranking from an expert is evaluated due to the relative quality (Brown et al., 2019). To sum up, these methods rely on auxiliary information from experts, namely high-quality demonstrations, confidence scores, and ranking. In practice, these pieces of information can be scarce or noisy, which leads to a poor performance of these methods.

In this paper, we consider a novel but realistic setting of IL where only diverse-quality demonstrations are available. Meanwhile, the level of demonstrators’ expertise and auxiliary information from experts are fully absent. To tackle this challenging setting, we propose a new learning paradigm called variational imitation learning with diverse-quality demonstrations (VILD). The central idea of VILD is to model the level of demonstrators’ expertise via a probabilistic graphical model, and learn it along with a reward function that represents an intention of expert’s decision making. To scale up our model for large state and action spaces, we leverage the variational approach (Jordan et al., 1999), which can be implemented using reinforcement learning (RL) for flexibility (Sutton & Barto, 1998). To further improve data-efficiency of VILD when learning the reward function, we utilize importance sampling (IS) to re-weight a sampling distribution according to the estimated level of demonstrators’ expertise. Experiments on continuous-control benchmarks demonstrate that: 1) VILD is robust against diverse-quality demonstrations and outperforms existing methods significantly. 2) VILD with IS is data-efficient, since it learns the policy using a less number of transition samples.

2 IL FROM DIVERSE-QUALITY DEMONSTRATIONS AND ITS CHALLENGE

Before delving into our main contribution, we first give the minimum background about RL and IL. Then, we formulate a new setting in IL called *diverse-quality demonstrations*, discuss its challenge, and reveal the deficiency of existing methods.

Reinforcement learning. Reinforcement learning (RL) (Sutton & Barto, 1998) aims to learn an optimal policy of a sequential decision making problem, which is often mathematically formulated as a Markov decision process (MDP) (Puterman, 1994). We consider a finite-horizon MDP with continuous state and action spaces defined by a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, p(s'|s, \mathbf{a}), p_1(\mathbf{s}_1), r(\mathbf{s}, \mathbf{a}))$ with a state $\mathbf{s}_t \in \mathcal{S} \subseteq \mathbb{R}^{d_s}$, an action $\mathbf{a}_t \in \mathcal{A} \subseteq \mathbb{R}^{d_a}$, an initial state density $p_1(\mathbf{s}_1)$, a transition probability density $p(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t)$, and a reward function $r: \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$, where the subscript $t \in \{1, \dots, T\}$ denotes the time step. A sequence of states and actions, $(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})$, is called a trajectory. A decision making of an agent is determined by a policy $\pi(\mathbf{a}_t|\mathbf{s}_t)$, which is a conditional probability density of action given state. RL seeks for an optimal policy $\pi^*(\mathbf{a}_t|\mathbf{s}_t)$ which maximizes the expected cumulative reward: $\mathbb{E}_{p_\pi(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})}[\sum_{t=1}^T r(\mathbf{s}_t, \mathbf{a}_t)]$, where $p_\pi(\mathbf{s}_{1:T}, \mathbf{a}_{1:T}) = p_1(\mathbf{s}_1)\prod_{t=1}^T p(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t)\pi(\mathbf{a}_t|\mathbf{s}_t)$ is a trajectory probability density induced by π . RL has shown great successes recently, especially when combined with deep neural networks (Silver et al., 2017). However, a major limitation of RL is that it relies on the reward function which may be unavailable in practice (Schaal, 1999).

Imitation learning. To address the above limitation of RL, imitation learning (IL) was proposed (Schaal, 1999). Without using the reward function, IL aims to learn the optimal policy from demonstrations that encode information about the optimal policy. A common assumption in most IL methods is that, demonstrations are collected by $K \geq 1$ demonstrators who execute actions \mathbf{a}_t drawn from $\pi^*(\mathbf{a}_t|\mathbf{s}_t)$ for every states \mathbf{s}_t . A graphical model describing this data collection process is depicted in Figure 1(a), where a random variable $k \in \{1, \dots, K\}$ denotes each demonstrator’s identification number and $p(k)$ denotes the probability of collecting a demonstration from the k -th demonstrator. Under this assumption, demonstrations $\{(\mathbf{s}_{1:T}, \mathbf{a}_{1:T}, k)_n\}_{n=1}^N$ (i.e., observed random variables in Figure 1(a)) are called *expert demonstrations* and are regarded to be drawn independently from a probability density $p^*(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})p(k) = p(k)p_1(\mathbf{s}_1)\prod_{t=1}^T p(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t)\pi^*(\mathbf{a}_t|\mathbf{s}_t)$. We note that k does not affect the trajectory density $p^*(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})$ and can be omitted. We assume a common assumption that $p_1(\mathbf{s}_1)$ and $p(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t)$ are unknown but we can sample states from them.

IL has shown great successes in benchmark settings (Ho & Ermon, 2016; Fu et al., 2018; Peng et al., 2019). However, practical applications of IL in the real-world is relatively few (Schroecker et al., 2019). One of the main reasons is that most IL methods aim to learn with expert demonstrations. In practice, such demonstrations are often too costly to obtain due to a limited number of experts, and even when we obtain them, the number of demonstrations is often too few to accurately learn the optimal policy (Audiffren et al., 2015; Wu et al., 2019; Brown et al., 2019).

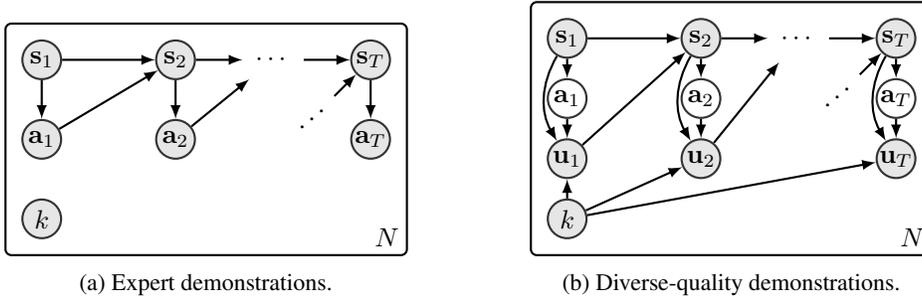


Figure 1: Graphical models describe expert demonstrations and diverse-quality demonstrations. Shaded and unshaded nodes indicate observed and unobserved random variables, respectively. Plate notations indicate that the sampling process is repeated for N times. $\mathbf{s}_t \in \mathcal{S}$ is a state with transition densities $p(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t)$, $\mathbf{a}_t \in \mathcal{A}$ is an action with density $\pi^*(\mathbf{a}_t|\mathbf{s}_t)$, $\mathbf{u}_t \in \mathcal{A}$ is a noisy action with density $p(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k)$, and $k \in \{1, \dots, K\}$ is an identification number with distribution $p(k)$.

New setting in IL: Diverse-quality demonstrations. To improve practicality of IL, we consider a new learning paradigm called *IL with diverse-quality demonstrations*, where demonstrations are collected from demonstrators with different level of expertise. Compared to expert demonstrations, diverse-quality demonstrations can be collected more cheaply, e.g., via crowdsourcing (Mandlekar et al., 2018). The graphical model in Figure 1(b) depicts the process of collecting such demonstrations from $K > 1$ demonstrators. Formally, we select the k -th demonstrator according to a distribution $p(k)$. After selecting k , for each time step t , the k -th demonstrator observes state \mathbf{s}_t and samples action \mathbf{a}_t using $\pi^*(\mathbf{a}_t|\mathbf{s}_t)$. However, the demonstrator may not execute \mathbf{a}_t in the MDP if this demonstrator is not expertised. Instead, he/she may sample an action $\mathbf{u}_t \in \mathcal{A}$ with another probability density $p(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k)$ and execute it. Then, the next state \mathbf{s}_{t+1} is observed with a probability density $p(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{u}_t)$, and the demonstrator continues making decision until time step T . We repeat this process for N times to collect *diverse-quality demonstrations* $\mathcal{D}_d = \{(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, k)_n\}_{n=1}^N$. These demonstrations are regarded to be drawn independently from a probability density

$$p_d(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}|k)p(k) = p(k)p(\mathbf{s}_1) \prod_{t=1}^T p_1(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{u}_t) \int_{\mathcal{A}} \pi^*(\mathbf{a}_t|\mathbf{s}_t)p(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k)d\mathbf{a}_t. \quad (1)$$

We refer to $p(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k)$ as a noisy policy of the k -th demonstrator, since it is used to execute a noisy action \mathbf{u}_t . Our goal is to learn the optimal policy π^* using diverse-quality demonstrations \mathcal{D}_d .

The deficiency of existing methods. We conjecture that existing IL methods are not suitable to learn with diverse-quality demonstrations according to p_d . Specifically, these methods always treat observed demonstrations as if they were drawn from p^* . By comparing p^* and p_d , we can see that existing methods would learn $\pi(\mathbf{u}_t|\mathbf{s}_t)$ such that $\pi(\mathbf{u}_t|\mathbf{s}_t) \approx \sum_{k=1}^K p(k) \int_{\mathcal{A}} \pi^*(\mathbf{a}_t|\mathbf{s}_t)p(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k)d\mathbf{a}_t$. In other words, they learn a policy that averages over decisions of all demonstrators. This would be problematic when amateurs are present, as averaged decisions of all demonstrators would be highly different from those of all experts. Worse yet, state distributions of amateurs and experts tend to be highly different, which often leads to the unstable learning: The learned policy oscillated between well-performed policy and poorly-performed policy. For these reasons, we believe that existing methods tend to learn a policy that achieves average performances, and are not suitable for handling the setting of diverse-quality demonstrations.

3 VILD: A ROBUST METHOD FOR DIVERSE-QUALITY DEMONSTRATIONS

This section presents VILD, namely a robust method for tackling the challenge from diverse-quality demonstrations. Specifically, we build a probabilistic model that explicitly describes the level of demonstrators' expertise and a reward function (Section 3.1), and estimate its parameters by a variational approach (Section 3.2), which can be implemented easily by RL (Section 3.3). We also improve data-efficiency by using importance sampling (Section 3.4). Mathematical derivations are provided in Appendix A.

3.1 MODELING DIVERSE-QUALITY DEMONSTRATIONS

This section describes a model which enables estimating the level of demonstrators’ expertise. We first describe a naive model, whose parameters can be estimated trivially via supervised learning, but suffers from the issue of compounding error. Then, we describe our proposed model, which avoids the issue of the naive model by learning a reward function.

Naive model. Based on p_d , one of the simplest models to handle diverse-quality demonstrations is $p_{\theta, \omega}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, k) = p(k)p(\mathbf{s}_1)\prod_{t=1}^T p(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{u}_t) \int_{\mathcal{A}} \pi_{\theta}(\mathbf{a}_t|\mathbf{s}_t)p_{\omega}(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k)d\mathbf{a}_t$, where π_{θ} and p_{ω} are learned to estimate the optimal policy and the noisy policy, respectively. The parameters θ and ω can be learned by minimizing the Kullback-Leibler (KL) divergence from the data distribution to the model. This naive model can be regarded as an extension of a model proposed by Raykar et al. (2010) for handling diverse-quality data in supervised learning.

The main advantage of this naive model is that its parameters can be estimated trivially via supervised learning. However, this naive model suffers from the issue of compounding error (Ross & Bagnell, 2010) and tends to perform poorly. Specifically, supervised-learning methods assume that data distributions during training and testing are identical. However, data distributions during training and testing are different in IL, since data distributions depend on policies (Puterman, 1994). A discrepancy of data distributions causes compounding errors during testing, where prediction errors increase further in future predictions. Due to this issue, supervised-learning methods often perform poorly in IL (Ross & Bagnell, 2010). The issue becomes even worse with diverse-quality demonstrations, since data distributions of different demonstrators tend to be highly different. For these reasons, this naive model is not suitable for our setting.

Proposed model. To avoid the issue of compounding error, our method utilizes the inverse RL (IRL) approach (Ng & Russell, 2000), where we aim to learn a reward function from diverse-quality demonstrations¹. IL problems can be solved by a combination of IRL and RL, where we learn a reward function by IRL and then learn a policy from the reward function by RL. This combination avoids the issue of compounding error, since the policy is learned by RL which generalizes to states not presented in demonstrations (Ho & Ermon, 2016).

Specifically, our proposed model is based on a model of maximum entropy IRL (MaxEnt-IRL) (Ziebart et al., 2010). Briefly speaking, MaxEnt-IRL learns a reward function from expert demonstrations by using a model $p_{\phi}(\mathbf{s}_{1:T}, \mathbf{a}_{1:T}) \propto p(\mathbf{s}_1)\prod_{t=1}^T p_1(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t) \exp(r_{\phi}(\mathbf{s}_t, \mathbf{a}_t))$. Based on this model, we propose to learn the reward function *and* the level of expertise by a model

$$p_{\phi, \omega}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, k) \propto p(k)p_1(\mathbf{s}_1) \prod_{t=1}^T p(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{u}_t) \int_{\mathcal{A}} \exp(r_{\phi}(\mathbf{s}_t, \mathbf{a}_t)) p_{\omega}(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k)d\mathbf{a}_t, \quad (2)$$

where ϕ and ω are parameters. We denote a normalization term of this model by $Z_{\phi, \omega}$. By comparing the proposed model $p_{\phi, \omega}$ to the data distribution p_d , the reward parameter ϕ should be learned so that the cumulative reward is proportion to a joint probability density of actions given by the optimal policy, i.e., $\exp(\sum_{t=1}^T r_{\phi}(\mathbf{s}_t, \mathbf{a}_t)) \propto \prod_{t=1}^T \pi^*(\mathbf{a}_t|\mathbf{s}_t)$. In other words, the cumulative reward is large for trajectories induced by the optimal policy. Therefore, the optimal policy can be learned by maximizing the cumulative reward. Meanwhile, the density $p_{\omega}(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k)$ is learned to estimate the noisy policy $p(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k)$. In the remainder, we refer to ω as an expertise parameter.

To learn parameters of this model, we propose to minimize the KL divergence from the data distribution to the model: $\min_{\phi, \omega} \text{KL}(p_d(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}|k)p(k)||p_{\phi, \omega}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, k))$. By rearranging terms and ignoring constant terms, minimizing this KL divergence is equivalent to solving an optimization problem $\max_{\phi, \omega} f(\phi, \omega) - g(\phi, \omega)$, where $f(\phi, \omega) = \mathbb{E}_{p_d(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}|k)p(k)}[\sum_{t=1}^T \log(\int_{\mathcal{A}} \exp(r_{\phi}(\mathbf{s}_t, \mathbf{a}_t))p_{\omega}(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k)d\mathbf{a}_t)]$ and $g(\phi, \omega) = \log Z_{\phi, \omega}$. To solve this optimization, we need to compute the integrals over both state space \mathcal{S} and action space \mathcal{A} . Computing these integrals is feasible for small state and action spaces, but is infeasible for large state and action spaces. To scale up our model to MDPs with large state and action spaces, we leverage a variational approach in the followings.

¹We emphasize that IRL is different from RL; IRL learns a reward function from demonstrations, whereas RL learns an optimal policy from a known reward function.

3.2 VARIATIONAL APPROACH FOR PARAMETER ESTIMATION

The central idea of the variational approach is to lower-bound an integral by the Jensen inequality and a variational distribution (Jordan et al., 1999). The main benefit of the variational approach is that the integral can be indirectly computed via the lower-bound, given an optimal variational distribution. However, finding the optimal distribution often requires solving a sub-optimization problem.

Before we proceed, notice that $f(\phi, \omega) - g(\phi, \omega)$ is not a joint concave function of the integrals, and this prohibits using the Jensen inequality. However, we can separately lower-bound f and g by the Jensen inequality, since they are concave functions of their corresponding integrals. Specifically, let $l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k) = r_{\phi}(\mathbf{s}_t, \mathbf{a}_t) + \log p_{\omega}(\mathbf{u}_t | \mathbf{s}_t, \mathbf{a}_t, k)$. By using a variational distribution $q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k)$ with parameter ψ , we obtain an inequality $f(\phi, \omega) \geq \mathcal{F}(\phi, \omega, \psi)$, where

$$\mathcal{F}(\phi, \omega, \psi) = \mathbb{E}_{p_d(\mathbf{s}_{1:T}, \mathbf{u}_{1:T} | k) p(k)} \left[\sum_{t=1}^T \mathbb{E}_{q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k)} [l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k)] + H_t(q_{\psi}) \right], \quad (3)$$

and $H_t(q_{\psi}) = -\mathbb{E}_{q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k)} [\log q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k)]$. It is trivial to verify that the equality $f(\phi, \omega) = \max_{\psi} \mathcal{F}(\phi, \omega, \psi)$ holds (Jordan et al., 1999), where the maximizer ψ^* of the lower-bound yields $q_{\psi^*}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k) \propto \exp(l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k))$. Therefore, the function $f(\phi, \omega)$ can be substituted by $\max_{\psi} \mathcal{F}(\phi, \omega, \psi)$. Meanwhile, by using a variational distribution $q_{\theta}(\mathbf{a}_t, \mathbf{u}_t | \mathbf{s}_t, k)$ with parameter θ , we obtain an inequality $g(\phi, \omega) \geq \mathcal{G}(\phi, \omega, \theta)$, where

$$\mathcal{G}(\phi, \omega, \theta) = \mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)} \left[\sum_{t=1}^T l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k) - \log q_{\theta}(\mathbf{a}_t, \mathbf{u}_t | \mathbf{s}_t, k) \right], \quad (4)$$

and $\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k) = p(k) p_1(\mathbf{s}_1) \prod_{t=1}^T p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{u}_t) q_{\theta}(\mathbf{a}_t, \mathbf{u}_t | \mathbf{s}_t, k)$. The lower-bound \mathcal{G} resembles an objective function of maximum entropy RL (MaxEnt-RL) (Ziebart et al., 2010). By using the optimality results of MaxEnt-RL (Haarnoja et al., 2018), we have an equality $g(\phi, \omega) = \max_{\theta} \mathcal{G}(\phi, \omega, \theta)$. Therefore, the function $g(\phi, \omega)$ can be substituted by $\max_{\theta} \mathcal{G}(\phi, \omega, \theta)$.

By using these lower-bounds, we have that $\max_{\phi, \omega} f(\phi, \omega) - g(\phi, \omega) = \max_{\phi, \omega, \psi} \mathcal{F}(\phi, \omega, \psi) - \max_{\theta} \mathcal{G}(\phi, \omega, \theta) = \max_{\phi, \omega, \psi} \min_{\theta} \mathcal{F}(\phi, \omega, \psi) - \mathcal{G}(\phi, \omega, \theta)$. Solving the max-min problem is often feasible even for large state and action spaces, since $\mathcal{F}(\phi, \omega, \psi)$ and $\mathcal{G}(\phi, \omega, \theta)$ are defined as an expectation and can be optimized straightforwardly. Nevertheless, in practice, we represent the variational distributions by parameterized functions, and iteratively solve the sub-optimization (w.r.t. ψ and θ) by stochastic optimization methods. However, in this scenario, the equalities $f(\phi, \omega) = \max_{\psi} \mathcal{F}(\phi, \omega, \psi)$ and $g(\phi, \omega) = \max_{\theta} \mathcal{G}(\phi, \omega, \theta)$ may not hold for two reasons. First, the optimal variational distributions may not be in the space of our parameterized functions. Second, stochastic optimization methods may yield local solutions. Nonetheless, when the variational distributions are represented by deep neural networks, the obtained variational distributions are often reasonably accurate and the equalities approximately hold (Ranganath et al., 2014).

3.3 MODEL SPECIFICATION

In practice, we are required to specify models for q_{θ} and p_{ω} . We propose to use $q_{\theta}(\mathbf{a}_t, \mathbf{u}_t | \mathbf{s}_t, k) = q_{\theta}(\mathbf{a}_t | \mathbf{s}_t) \mathcal{N}(\mathbf{u}_t | \mathbf{a}_t, \Sigma)$ and $p_{\omega}(\mathbf{u}_t | \mathbf{s}_t, \mathbf{a}_t, k) = \mathcal{N}(\mathbf{u}_t | \mathbf{a}_t, \mathbf{C}_{\omega}(k))$. As shown below, the choice for $q_{\theta}(\mathbf{a}_t, \mathbf{u}_t | \mathbf{s}_t, k)$ enables us to solve the sub-optimization w.r.t. θ by using RL with reward function r_{ϕ} . Meanwhile, the choice for $p_{\omega}(\mathbf{u}_t | \mathbf{s}_t, \mathbf{a}_t, k)$ incorporates our prior assumption that the noisy policy tends to Gaussian, which is a reasonable assumption for actual human motor behavior (van Beers et al., 2004). Under these model specifications, solving $\max_{\phi, \omega, \psi} \min_{\theta} \mathcal{F}(\phi, \omega, \psi) - \mathcal{G}(\phi, \omega, \theta)$ is equivalent to solving $\max_{\phi, \omega, \psi} \min_{\theta} \mathcal{H}(\phi, \omega, \psi, \theta)$, where

$$\begin{aligned} \mathcal{H}(\phi, \omega, \psi, \theta) = & \mathbb{E}_{p_d(\mathbf{s}_{1:T}, \mathbf{u}_{1:T} | k) p(k)} \left[\sum_{t=1}^T \mathbb{E}_{q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k)} \left[r_{\phi}(\mathbf{s}_t, \mathbf{a}_t) - \frac{1}{2} \|\mathbf{u}_t - \mathbf{a}_t\|_{\mathbf{C}_{\omega}^{-1}(k)}^2 \right] + H_t(q_{\psi}) \right] \\ & - \mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})} \left[\sum_{t=1}^T r_{\phi}(\mathbf{s}_t, \mathbf{a}_t) - \log q_{\theta}(\mathbf{a}_t | \mathbf{s}_t) \right] + \frac{T}{2} \mathbb{E}_{p(k)} [\text{Tr}(\mathbf{C}_{\omega}^{-1}(k) \Sigma)]. \end{aligned} \quad (5)$$

Here, $\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{a}_{1:T}) = p_1(\mathbf{s}_1) \prod_{t=1}^T \int_{\mathcal{A}} p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{u}_t) \mathcal{N}(\mathbf{u}_t | \mathbf{a}_t, \Sigma) d\mathbf{u}_t q_{\theta}(\mathbf{a}_t | \mathbf{s}_t)$ is a noisy trajectory density induced by a policy $q_{\theta}(\mathbf{a}_t | \mathbf{s}_t)$, where $\mathcal{N}(\mathbf{u}_t | \mathbf{a}_t, \Sigma)$ can be regarded as an approximation of the noisy policy in Figure 1(b). Minimizing \mathcal{H} w.r.t. θ resembles solving a MaxEnt-RL problem with a reward function $r_{\phi}(\mathbf{s}_t, \mathbf{a}_t)$, except that trajectories are collected according to the noisy trajectory density. In other words, this minimization problem can be solved using RL, and $q_{\theta}(\mathbf{a}_t | \mathbf{s}_t)$ can be regarded as an approximation of the optimal policy. The hyper-parameter Σ determines the quality of this approximation: smaller value of Σ gives a better approximation. Therefore, by choosing

a reasonably small value of Σ , solving the max-min problem in Eq. (5) yields a reward function $r_\phi(\mathbf{s}_t, \mathbf{a}_t)$ and a policy $q_\theta(\mathbf{a}_t|\mathbf{s}_t)$. This policy imitates the optimal policy, which is the goal of IL.

The model specification for p_ω incorporates our prior assumption about the noisy policy. Namely, $p_\omega(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k) = \mathcal{N}(\mathbf{u}_t|\mathbf{a}_t, \mathbf{C}_\omega(k))$ assumes that the noisy policy tends to Gaussian, where $\mathbf{C}_\omega(k)$ gives an estimated expertise of the k -th demonstrator: High-expertise demonstrators have small $\mathbf{C}_\omega(k)$ and vice-versa for low-expertise demonstrators. Note that VILD is not restricted to this choice. Different choices of p_ω incorporate different prior assumptions. For example, a Laplace distribution incorporates a prior assumption about demonstrators who tend to execute outlier actions (Murphy, 2013). In such a case, the squared error in \mathcal{H} is replaced by the absolute error (see Appendix A.3).

It should be mentioned that $q_\psi(\mathbf{a}_t|\mathbf{s}_t, \mathbf{u}_t, k)$ maximizes the immediate reward and minimizes the weighted squared error between \mathbf{u}_t and \mathbf{a}_t . The trade-off between the reward and squared-error is determined by $\mathbf{C}_\omega(k)$. Specifically, for demonstrators with a small $\mathbf{C}_\omega(k)$ (i.e., high-expertise demonstrators), the squared error has a large magnitude and q_ψ tends to minimize the squared error. Meanwhile, for demonstrators with a large value of $\mathbf{C}_\omega(k)$ (i.e., low-expertise demonstrators), the squared error has a small magnitude and q_ψ tends to maximize the immediate reward.

We implement VILD with deep neural networks where we iteratively update ϕ , ω , and ψ by stochastic gradient methods, and update θ by policy gradient methods. A pseudo-code of VILD and implementation details are given in Appendix B. In our implementation, we include a regularization term $L(\omega) = T\mathbb{E}_{p(k)}[\log|\mathbf{C}_\omega^{-1}(k)|]/2$, to penalize large value of $\mathbf{C}_\omega(k)$. Without this regularization, $\mathbf{C}_\omega(k)$ can be overly large which makes learning degenerate. We note that \mathcal{H} already includes such a penalty via the trace term: $\mathbb{E}_{p(k)}[\text{Tr}(\mathbf{C}_\omega^{-1}(k)\Sigma)]$. However, the strength of this penalty tends to be too small, since we choose Σ to be small.

3.4 IMPORTANCE SAMPLING FOR REWARD LEARNING

To improve the convergence rate of VILD when updating ϕ , we use importance sampling (IS). Specifically, by analyzing the gradient $\nabla_\phi \mathcal{H} = \nabla_\phi \{ \mathbb{E}_{p_d(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}|k)p(k)} [\sum_{t=1}^T \mathbb{E}_{q_\psi(\mathbf{a}_t|\mathbf{s}_t, \mathbf{u}_t, k)} [r_\phi(\mathbf{s}_t, \mathbf{a}_t)]] - \mathbb{E}_{\tilde{q}_\theta(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})} [\sum_{t=1}^T r_\phi(\mathbf{s}_t, \mathbf{a}_t)] \}$, we can see that the reward function is updated to maximize the expected cumulative reward obtained by demonstrators and q_ψ , while minimizing the expected cumulative reward obtained by q_θ . However, low-quality demonstrations often yield low reward values. For this reason, stochastic gradients estimated by these demonstrations tend to be uninformative, which leads to slow convergence and poor data-efficiency.

To avoid estimating such uninformative gradients, we use IS to estimate gradients using high-quality demonstrations which are sampled with high probability. Briefly, IS is a technique for estimating an expectation over a distribution by using samples from a different distribution (Robert & Casella, 2005). For VILD, we propose to sample k from a distribution $\tilde{p}(k) \propto \|\text{vec}(\mathbf{C}_\omega^{-1}(k))\|_1$. This distribution assigns high probabilities to demonstrators with high estimated level of expertise (i.e., demonstrators with a small $\mathbf{C}_\omega(k)$). With this distribution, the estimated gradients tend to be more informative which leads to a faster convergence. To reduce a sampling bias, we use a truncated importance weight: $w(k) = \min(p(k)/\tilde{p}(k), 1)$ (Ionides, 2008), which leads to an IS gradient: $\nabla_\phi \mathcal{H}_{\text{IS}} = \nabla_\phi \{ \mathbb{E}_{p_d(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}|k)\tilde{p}(k)} [w(k)\sum_{t=1}^T \mathbb{E}_{q_\psi(\mathbf{a}_t|\mathbf{s}_t, \mathbf{u}_t, k)} [r_\phi(\mathbf{s}_t, \mathbf{a}_t)]] - \mathbb{E}_{\tilde{q}_\theta(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})} [\sum_{t=1}^T r_\phi(\mathbf{s}_t, \mathbf{a}_t)] \}$. Computing $w(k)$ requires $p(k)$, which can be estimated accurately since k is a discrete random variable. For simplicity, we assume that $p(k)$ is a uniform distribution.

4 RELATED WORK

In this section, we will discuss a related area of supervised learning with diverse-quality data. Besides, we will discuss existing IL methods that use the variational approach.

Supervised learning with diverse-quality data. In supervised learning, diverse-quality data has been extensively studied, e.g. learning with noisy labels (Angluin & Laird, 1988). This task assumes that human labelers may assign incorrect labels to training inputs. With such labelers, the obtained dataset consists of high-quality data with correct labels and low-quality data with incorrect labels. To handle this setting, many methods were proposed (Natarajan et al., 2013; Han et al., 2018). The most related methods are probabilistic models, which aim to infer correct labels and the level of labelers' expertise (Raykar et al., 2010; Khetan et al., 2018). Specifically, Raykar et al. (2010) proposed a method based on a two-coin model which enables estimating the correct labels and level of expertise. Recently, Khetan et al. (2018) proposed a method based on weighted loss functions, where the weight

is determined by the estimated labels and level of expertise. Methods for supervised learning with diverse-quality data can be leveraged to learn a policy in our setting. However, they tend to perform poorly due to the issue of compounding error, as discussed previously in Section 3.1.

Variational approach in IL. The variational approach has been previously utilized in IL to perform MM-IL and reduce over-fitting. Specifically, MM-IL aims to learn a multi-modal policy from diverse demonstrations collected by many experts (Li et al., 2017), where each mode of the policy represents decision making of each expert². A multi-modal policy is commonly represented by a context-dependent policy, where each context represents each mode of the policy. The variational approach has been used to learn such contexts, i.e., by learning a variational auto-encoder (Wang et al., 2017) and maximizing a variational lower-bound of mutual information (Li et al., 2017; Hausman et al., 2017). Meanwhile, variational information bottleneck (VIB) (Alemi et al., 2017) has been used to reduce over-fitting in IL (Peng et al., 2019). Specifically, VIB aims to compress information flow by minimizing a variational bound of mutual information. This compression filters irrelevant signals, which leads to less over-fitting. Unlike these existing works, we utilize the variational approach to aid computing integrals in large state-action spaces, but not for learning a variational auto-encoder or optimizing a variational bound of mutual information.

5 EXPERIMENTS

In this section, we experimentally evaluate the performance of VILD (with and without IS) in continuous-control benchmark tasks (Brockman et al., 2016). Performance is evaluated using a cumulative ground-truth reward along trajectories (i.e., higher is better) (Ho & Ermon, 2016), which is computed using 10 test trajectories generated by learned policies (i.e., $q_{\theta}(\mathbf{a}_t | \mathbf{s}_t)$). We repeat experiments for 5 trials with different random seeds and report the mean and standard error.

Baselines & data generation. We compare VILD against GAIL (Ho & Ermon, 2016), AIRL (Fu et al., 2018), VAIL (Peng et al., 2019), MaxEnt-IRL (Ziebart et al., 2010), and InfoGAIL (Li et al., 2017). These are online IL methods which collect transition samples to learn policies. We use trust region policy optimization (TRPO) (Schulman et al., 2015) to update policies, except for the Humanoid task where we use soft actor-critic (SAC) (Haarnoja et al., 2018). To generate demonstrations from π^* (pre-trained by TRPO) according to Figure 1(b), we use two types of noisy policy $p(\mathbf{u}_t | \mathbf{a}_t, \mathbf{s}_t, k)$: Gaussian noisy policy: $\mathcal{N}(\mathbf{u}_t | \mathbf{a}_t, \sigma_k^2 \mathbf{I})$ and time-signal-dependent (TSD) noisy policy: $\mathcal{N}(\mathbf{u}_t | \mathbf{a}_t, \text{diag}(\mathbf{b}_k(t) \times \|\mathbf{a}_t\|_1 / d_{\mathbf{a}}))$, where $\mathbf{b}_k(t)$ is sampled from a noise process. We use 10 demonstrators with different σ_k and noise processes for $\mathbf{b}_k(t)$. Notice that for TSD, the noise variance depends on time and magnitude of actions. This characteristic of TSD has been observed in human motor control (van Beers et al., 2004). More details of data generation are given in Appendix C.

Results against online IL methods. Figure 2 shows learning curves of VILD and existing methods against the number of transition samples in HalfCheetah and Ant³, whereas Table 1 reports the performance achieved in the last 100 iterations. Clearly, VILD with IS overall outperforms existing methods in terms of both data-efficiency and final performance, i.e., VILD with IS learns better policies using less numbers of transition samples. VILD without IS tends to outperform existing methods in terms of the final performance. However, it is less data-efficient when compared to VILD with IS, except on Humanoid with the Gaussian noisy policy, where VILD without IS tends to perform better than VILD with IS. We conjecture that this is because IS slightly biases gradient estimation, which may have a negative effect on the performance. Nonetheless, the overall good performance of VILD with IS suggests that it is an effective method to handle diverse-quality demonstrations.

On the contrary, existing methods perform poorly as expected, except on the Humanoid task. For the Humanoid task, VILD tends to perform the best in terms of the mean performance. Nonetheless, all methods except GAIL achieve statistically comparable performance according to t-test. This is perhaps because amateurs in this task perform relatively well compared to amateurs in other tasks, as seen from demonstrators’ performance given in Table 2 and 3 (Appendix C). Since amateurs perform relatively well, demonstrations from these amateurs do not severely affect the performance of IL methods in this task when compared to the other tasks.

²We emphasize that diverse demonstrations are different from diverse-quality demonstrations. Diverse demonstrations are collected by experts who execute equally good policies, while diverse-quality demonstrations are collected by mixed demonstrators; The former consists of demonstrations that are equally high-quality but diverse in behavior, while the latter consists of demonstrations that are diverse in both quality and behavior.

³Learning curves of other tasks are given in Figure 4 in Appendix D.

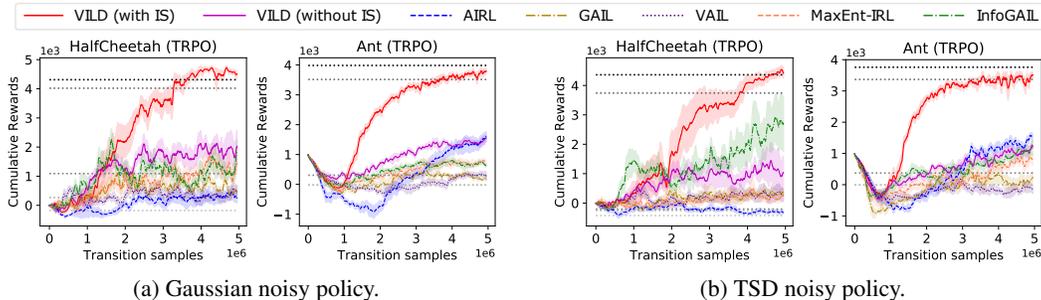


Figure 2: Performance averaged over 5 trials in terms of the mean and standard error. Demonstrations are generated by 10 demonstrators using (a) Gaussian and (b) TSD noisy policies. Horizontal dotted lines indicate performance of $k = 1, 3, 5, 7, 10$ demonstrators. IS denotes importance sampling.

Table 1: Performance in the last 100 iterations in terms of the mean and standard error of cumulative rewards over 5 trials (higher is better). Boldfaces indicate best and comparable methods according to t-test with significance level 0.01. (G) denotes Gaussian noisy policy and (TSD) denotes time-signal-dependent noisy policy. The performance of VAIL is similar to that of GAIL and is omitted.

Task	VILD (IS)	VILD (w/o IS)	AIRL	GAIL	MaxEnt-IRL	InfoGAIL
HalfCheetah (G)	4559 (43)	1848 (429)	341 (177)	551 (23)	1192 (245)	1244 (210)
HalfCheetah (TSD)	4394 (136)	1159 (594)	-304 (51)	318 (134)	177 (132)	2664 (779)
Ant (G)	3719 (65)	1426 (81)	1417 (184)	209 (30)	731 (93)	675 (36)
Ant (TSD)	3396 (64)	1072 (134)	1357 (59)	97 (161)	775 (135)	1076 (140)
Walker2d (G)	3470 (300)	2132 (64)	1534 (99)	1410 (115)	1795 (172)	1668 (82)
Walker2d (TSD)	3115 (130)	1244 (132)	578 (47)	834 (84)	752 (112)	1041 (36)
Humanoid (G)	3781 (557)	4840 (56)	4274 (93)	284 (24)	3038 (731)	4047 (653)
Humanoid (TSD)	4600 (97)	3610 (448)	4212 (121)	203 (31)	4132 (651)	3962 (635)

We found that InfoGAIL, which learns a context-dependent policy, can achieve good performance when the policy is conditioned on specific contexts. However, its performance is quite poor on average when using contexts from a uniform prior. These results support our conjecture that MM-IL methods are not suitable for our setting where the level of demonstrators’ expertise is absent.

It can be seen that VILD without IS performs better for the Gaussian noisy policy when compared to the TSD noisy policy. This is because the model of VILD is correctly specified for the Gaussian noisy policy, but the model is incorrectly specified for the TSD noisy policy; misspecified model indeed leads to the reduction in performance. Nonetheless, VILD with IS still performs well for both types of noisy policy. This is perhaps because negative effects of a misspecified model are not too severe for learning expertise parameters, which are required to compute $\tilde{p}(k)$.

We also conduct the following evaluations. Due to space limitation, figures are given in Appendix D.

Results against offline IL methods. We compare VILD against offline IL methods based on supervised learning, namely behavior cloning (BC) (Pomerleau, 1988), Co-Teaching which is based on a method for learning with noisy labels (Han et al., 2018), and BC from diverse-quality demonstrations (BC-D) which optimizes the naive model described in Section 3.1. Results in Figure 5 show that these methods perform worse than VILD overall; BC performs the worst since it severely suffers from both the compounding error and low-quality demonstrations. Compared to BC, BC-D and Co-teaching are quite robust against low-quality demonstrations, but they still perform worse than VILD with IS.

Accuracy of estimated expertise parameter. To evaluate accuracy of estimated expertise parameter, we compare the ground-truth value of σ_k under the Gaussian noisy policy against the learned covariance $C_\omega(k)$. Figure 6 shows that VILD learns an accurate ranking of demonstrators’ expertise. The values of these parameters are also quite accurate compared to the ground-truth, except for demonstrators with low-level of expertise. A reason for this phenomena is that low-quality demonstrations are highly dissimilar, which makes learning the expertise more challenging.

6 CONCLUSION AND FUTURE WORK

In this paper, we explored a practical setting in IL where demonstrations have diverse-quality. We showed the deficiency of existing methods, and proposed a robust method called VILD, which learns both the reward function and the level of demonstrators’ expertise by using the variational approach. Empirical results demonstrated that our work enables scalable and data-efficient IL under this practical setting. In future, we will explore other approaches to efficiently estimate parameters of the proposed model except the variational approach.

REFERENCES

- Alexander A. Alemi, Ian Fischer, Joshua V. Dillon, and Kevin Murphy. Deep variational information bottleneck. In *International Conference on Learning Representations*, 2017.
- Dana Angluin and Philip Laird. Learning from noisy examples. *Machine Learning*, 1988.
- Devansh Arpit, Stanislaw K. Jastrzebski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S. Kanwal, Tegan Maharaj, Asja Fischer, Aaron C. Courville, Yoshua Bengio, and Simon Lacoste-Julien. A closer look at memorization in deep networks. In *International Conference on Machine Learning*, 2017.
- Julien Audiffren, Michal Valko, Alessandro Lazaric, and Mohammad Ghavamzadeh. Maximum entropy semi-supervised inverse reinforcement learning. In *International Joint Conferences on Artificial Intelligence*, 2015.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI Gym. *CoRR*, abs/1606.01540, 2016.
- Daniel S. Brown, Wonjoon Goo, Prabhat Nagarajan, and Scott Niekum. Extrapolating beyond suboptimal demonstrations via inverse reinforcement learning from observations. In *International Conference on Machine Learning*, 2019.
- Chelsea Finn, Sergey Levine, and Pieter Abbeel. Guided cost learning: Deep inverse optimal control via policy optimization. In *International Conference on Machine Learning*, 2016.
- Lex Fridman, Daniel E. Brown, Michael Glazer, William Angell, Spencer Dodd, Benedikt Jenik, Jack Terwilliger, Julia Kindelsberger, Li Ding, Sean Seaman, Hillary Abraham, Alea Mehler, Andrew Sipperley, Anthony Pettinato, Bobbie Seppelt, Linda Angell, Bruce Mehler, and Bryan Reimer. MIT autonomous vehicle technology study: Large-scale deep learning based analysis of driver behavior and interaction with automation. *CoRR*, abs/1711.06976, 2017.
- Justin Fu, Katie Luo, and Sergey Levine. Learning robust rewards with adversarial inverse reinforcement learning. In *International Conference on Learning Representation*, 2018.
- Shixiang Gu, Timothy Lillicrap, Richard E Turner, Zoubin Ghahramani, Bernhard Schölkopf, and Sergey Levine. Interpolated policy gradient: Merging on-policy and off-policy gradient estimation for deep reinforcement learning. In *Advances in Neural Information Processing Systems*, 2017.
- Ishaan Gulrajani, Faruk Ahmed, Martín Arjovsky, Vincent Dumoulin, and Aaron C. Courville. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems*, 2017.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning*, 2018.
- Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor W. Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *Advances in Neural Information Processing Systems*, 2018.
- Karol Hausman, Yevgen Chebotar, Stefan Schaal, Gaurav S. Sukhatme, and Joseph J. Lim. Multi-modal imitation learning from unstructured demonstrations using generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2017.
- Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *Advances in Neural Information Processing Systems*, 2016.
- Matthew D. Hoffman and David M. Blei. Stochastic structured variational inference. In *International Conference on Artificial Intelligence and Statistics*, 2015.

- Edward L. Ionides. Truncated importance sampling. *Journal of Computational and Graphical Statistics*, 2008.
- E. T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 1957.
- Michael I. Jordan, Zoubin Ghahramani, Tommi S. Jaakkola, and Lawrence K. Saul. An introduction to variational methods for graphical models. *Machine Learning*, 1999.
- Ashish Kheta, Zachary C. Lipton, and Animashree Anandkumar. Learning from noisy singly-labeled data. In *International Conference on Learning Representations*, 2018.
- Yunzhu Li, Jiaming Song, and Stefano Ermon. Infogail: Interpretable imitation learning from visual demonstrations. In *Advances in Neural Information Processing Systems*, 2017.
- Ajay Mandelkar, Yuke Zhu, Animesh Garg, Jonathan Booher, Max Spero, Albert Tung, Julian Gao, John Emmons, Anchit Gupta, Emre Orbay, Silvio Savarese, and Li Fei-Fei. ROBOTURK: A crowdsourcing platform for robotic skill learning through imitation. In *Conference on Robot Learning*, 2018.
- Kevin P. Murphy. *Machine learning : a probabilistic perspective*. 2013.
- Nagarajan Natarajan, Inderjit S Dhillon, Pradeep K Ravikumar, and Ambuj Tewari. Learning with noisy labels. In *Advances in Neural Information Processing Systems*, 2013.
- Andrew Y. Ng and Stuart J. Russell. Algorithms for inverse reinforcement learning. In *International Conference on Machine Learning*, 2000.
- Takayuki Osa, Joni Pajarinen, Gerhard Neumann, J. Andrew Bagnell, Pieter Abbeel, and Jan Peters. An algorithmic perspective on imitation learning. *Foundations and Trends in Robotics*, 2018.
- Xue Bin Peng, Angjoo Kanazawa, Sam Toyer, Pieter Abbeel, and Sergey Levine. Variational discriminator bottleneck: Improving imitation learning, inverse RL, and GANs by constraining information flow. In *International Conference on Learning Representations*, 2019.
- Dean Pomerleau. ALVINN: an autonomous land vehicle in a neural network. In *Advances in Neural Information Processing Systems*, 1988.
- Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. 1994.
- Rajesh Ranganath, Sean Gerrish, and David M. Blei. Black box variational inference. In *International Conference on Artificial Intelligence and Statistics*, 2014.
- Vikas C. Raykar, Shipeng Yu, Linda H. Zhao, Gerardo Hermosillo Valadez, Charles Florin, Luca Bogoni, and Linda Moy. Learning from crowds. *Journal of Machine Learning Research*, 2010.
- Christian P. Robert and George Casella. *Monte Carlo Statistical Methods*. 2005.
- Stephane Ross and Drew Bagnell. Efficient reductions for imitation learning. In *International Conference on Artificial Intelligence and Statistics*, 2010.
- Stefan Schaal. Is imitation learning the route to humanoid robots? *Trends in Cognitive Sciences*, 1999.
- Yannick Schroecker, Mel Vecerik, and Jon Scholz. Generative predecessor models for sample-efficient imitation learning. In *International Conference on Learning Representations*, 2019.
- John Schulman, Sergey Levine, Philipp Moritz, Michael Jordan, and Pieter Abbeel. Trust region policy optimization. In *International Conference on Machine Learning*, 2015.
- Kyriacos Shiarlis, João V. Messias, and Shimon Whiteson. Inverse reinforcement learning from failure. In *International Conference on Autonomous Agents & Multiagent Systems*, 2016.
- David Silver, J. Andrew Bagnell, and Anthony Stentz. Learning autonomous driving styles and maneuvers from expert demonstration. In *International Symposium on Experimental Robotics*, 2012.
- 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*, 2017.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning - an Introduction*. MIT Press, 1998.
- G. E. Uhlenbeck and L. S. Ornstein. On the theory of the brownian motion. *Physical Review*, 1930.

Robert J. van Beers, Patrick Haggard, and Daniel M. Wolpert. The role of execution noise in movement variability. *Journal of Neurophysiology*, 2004.

Ziyu Wang, Josh Merel, Scott E. Reed, Nando de Freitas, Gregory Wayne, and Nicolas Heess. Robust imitation of diverse behaviors. In *Advances in Neural Information Processing Systems*, 2017.

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

Brian D. Ziebart, J. Andrew Bagnell, and Anind K. Dey. Modeling interaction via the principle of maximum causal entropy. In *International Conference on Machine Learning*, 2010.

A DERIVATIONS

This section derives the lower-bounds of $f(\phi, \omega)$ and $g(\phi, \omega)$ presented in the paper. We also derive the objective function $\mathcal{H}(\phi, \omega, \psi, \theta)$ of VILD.

A.1 LOWER-BOUND OF f

Let $l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k) = r_{\phi}(\mathbf{s}_t, \mathbf{a}_t) + \log p_{\omega}(\mathbf{u}_t | \mathbf{s}_t, \mathbf{a}_t, k)$, we have that $f(\phi, \omega) = \mathbb{E}_{p_d(\mathbf{s}_{1:T}, \mathbf{u}_{1:T} | k) p(k)} \left[\sum_{t=1}^T f_t(\phi, \omega) \right]$, where $f_t(\phi, \omega) = \log \int_{\mathcal{A}} \exp(l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k)) d\mathbf{a}_t$. By using a variational distribution $q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k)$ with parameter ψ , we can bound $f_t(\phi, \omega)$ from below by using the Jensen inequality as follows:

$$\begin{aligned} f_t(\phi, \omega) &= \log \left(\int_{\mathcal{A}} \exp(l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k)) \frac{q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k)}{q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k)} d\mathbf{a}_t \right) \\ &\geq \int_{\mathcal{A}} q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k) \log \left(\exp(l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k)) \frac{1}{q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k)} \right) d\mathbf{a}_t \\ &= \mathbb{E}_{q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k)} [l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k) - \log q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k)] \\ &= \mathcal{F}_t(\phi, \omega, \psi). \end{aligned} \quad (6)$$

Then, by using the linearity of expectation, we obtain the lower-bound of $f(\phi, \omega)$ as follows:

$$\begin{aligned} f(\phi, \omega) &\geq \mathbb{E}_{p_d(\mathbf{s}_{1:T}, \mathbf{u}_{1:T} | k) p(k)} \left[\sum_{t=1}^T \mathcal{F}_t(\phi, \omega, \psi) \right] \\ &= \mathbb{E}_{p_d(\mathbf{s}_{1:T}, \mathbf{u}_{1:T} | k) p(k)} \left[\sum_{t=1}^T \mathbb{E}_{q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k)} [l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k) - \log q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k)] \right] \\ &= \mathcal{F}(\phi, \omega, \psi). \end{aligned} \quad (7)$$

To verify that $f(\phi, \omega) = \max_{\psi} \mathcal{F}(\phi, \omega, \psi)$, we maximize $\mathcal{F}_t(\phi, \omega, \psi)$ w.r.t. q_{ψ} under the constraint that q_{ψ} is a valid probability density, i.e., $q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k) > 0$ and $\int_{\mathcal{A}} q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k) d\mathbf{a}_t = 1$. By setting the derivative of $\mathcal{F}_t(\phi, \omega, \psi)$ w.r.t. q_{ψ} to zero, we obtain

$$\begin{aligned} q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k) &= \exp(l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k) - 1) \\ &= \frac{\exp(l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k))}{\int_{\mathcal{A}} \exp(l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k)) d\mathbf{a}_t}, \end{aligned}$$

where the last line follows from the constraint $\int_{\mathcal{A}} q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k) d\mathbf{a}_t = 1$. To show that this is indeed the maximizer, we substitute $q_{\psi^*}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k) = \frac{\exp(l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k))}{\int_{\mathcal{A}} \exp(l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k)) d\mathbf{a}_t}$ into $\mathcal{F}_t(\phi, \omega, \psi)$:

$$\begin{aligned} \mathcal{F}_t(\phi, \omega, \psi^*) &= \mathbb{E}_{q_{\psi^*}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k)} [l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k) - \log q_{\psi^*}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k)] \\ &= \log \left(\int_{\mathcal{A}} \exp(l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k)) d\mathbf{a}_t \right). \end{aligned}$$

This equality verifies that $f_t(\phi, \omega) = \max_{\psi} \mathcal{F}_t(\phi, \omega, \psi)$. Finally, by using the linearity of expectation, we have that $f(\phi, \omega) = \max_{\psi} \mathcal{F}(\phi, \omega, \psi)$.

A.2 LOWER-BOUND OF g

Next, we derive the lower-bound of $g(\phi, \omega)$ presented in the paper. We first derive a trivial lower-bound using a general variational distribution over trajectories and reveal its issues. Then, we derive a lower-bound presented in the paper by using a structured variational distribution. Recall that the function $g(\phi, \omega) = \log Z_{\phi, \omega}$ is

$$g(\phi, \omega) = \log \left(\sum_{k=1}^K p(k) \int \cdots \int_{(\mathcal{S} \times \mathcal{A} \times \mathcal{A})^T} p_1(\mathbf{s}_1) \prod_{t=1}^T p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{u}_t) \exp(l(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k)) d\mathbf{s}_{1:T} d\mathbf{u}_{1:T} d\mathbf{a}_{1:T} \right).$$

Lower-bound via a variational distribution A lower-bound of g can be obtained by using a variational distribution $\bar{q}_{\beta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)$ with parameter β . We note that this variational distribution allows any dependency between the random variables $\mathbf{s}_{1:T}$, $\mathbf{u}_{1:T}$, $\mathbf{a}_{1:T}$, and k . By using this distribution, we have a lower-bound

$$\begin{aligned} g(\phi, \omega) &= \log \left(\sum_{k=1}^K p(k) \int \cdots \int_{(\mathcal{S} \times \mathcal{A} \times \mathcal{A})^T} p_1(\mathbf{s}_1) \prod_{t=1}^T p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{u}_t) \exp(l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k)) \right. \\ &\quad \left. \times \frac{\bar{q}_{\beta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)}{\bar{q}_{\beta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)} d\mathbf{s}_{1:T} d\mathbf{u}_{1:T} d\mathbf{a}_{1:T} \right) \\ &\geq \mathbb{E}_{\bar{q}_{\beta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)} \left[\log p(k) p_1(\mathbf{s}_1) + \sum_{t=1}^T \{ \log p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{u}_t) + l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k) \} \right. \\ &\quad \left. - \log \bar{q}_{\beta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k) \right] \\ &= \bar{\mathcal{G}}(\phi, \omega, \beta). \end{aligned} \tag{8}$$

The main issue of using this lower-bound is that, $\bar{\mathcal{G}}(\phi, \omega, \beta)$ can be computed or approximated only when we have an access to the transition probability $p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{u}_t)$. In many practical tasks, the transition probability is unknown and needs to be estimated. However, estimating the transition probability for large state and action spaces is known to be highly challenging (Sutton & Barto, 1998). For these reasons, this lower-bound is not suitable for our method.

Lower-bound via a structured variational distribution To avoid the above issue, we use the structure variational approach (Hoffman & Blei, 2015), where the key idea is to pre-define conditional dependency to ease computation. Specifically, we use a variational distribution $q_{\theta}(\mathbf{a}_t, \mathbf{u}_t | \mathbf{s}_t, k)$ with parameter θ and define dependencies between states according to the transition probability of an MDP. With this variational distribution, we lower-bound g as follows:

$$\begin{aligned} g(\phi, \omega) &= \log \left(\sum_{k=1}^K p(k) \int \cdots \int_{(\mathcal{S} \times \mathcal{A} \times \mathcal{A})^T} p_1(\mathbf{s}_1) \prod_{t=1}^T p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{u}_t) \exp(l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k)) \right. \\ &\quad \left. \times \frac{q_{\theta}(\mathbf{a}_t, \mathbf{u}_t | \mathbf{s}_t, k)}{q_{\theta}(\mathbf{a}_t, \mathbf{u}_t | \mathbf{s}_t, k)} d\mathbf{s}_{1:T} d\mathbf{u}_{1:T} d\mathbf{a}_{1:T} \right) \\ &\geq \mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)} \left[\sum_{t=1}^T l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k) - \log q_{\theta}(\mathbf{a}_t, \mathbf{u}_t | \mathbf{s}_t, k) \right] \\ &= \mathcal{G}(\phi, \omega, \theta), \end{aligned} \tag{9}$$

where $\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k) = p(k) p_1(\mathbf{s}_1) \prod_{t=1}^T p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{u}_t) q_{\theta}(\mathbf{a}_t, \mathbf{u}_t | \mathbf{s}_t, k)$. The optimal variational distribution $q_{\theta^*}(\mathbf{a}_t, \mathbf{u}_t | \mathbf{s}_t, k)$ can be founded by maximizing $\mathcal{G}(\phi, \omega, \theta)$ w.r.t. q_{θ} . Solving this maximization problem is identical to solving a maximum entropy RL (MaxEnt-RL) problem (Ziebart et al., 2010) for an MDP defined by a tuple $\mathcal{M} = (\mathcal{S} \times \mathbb{N}_+, \mathcal{A} \times \mathcal{A}, p(\mathbf{s}' | \mathbf{s}, \mathbf{u}) \mathbb{1}_{k=k'}, p_1(\mathbf{s}_1) p(k_1), l_{\phi, \omega}(\mathbf{s}, \mathbf{a}, \mathbf{u}, k))$. Specifically, this MDP is defined with a state variable $(\mathbf{s}_t, k_t) \in \mathcal{S} \times \mathbb{N}$, an action variable $(\mathbf{a}_t, \mathbf{u}_t) \in \mathcal{A} \times \mathcal{A}$, a transition probability density $p(\mathbf{s}_{t+1}, k_{t+1} | \mathbf{s}_t, k_t, \mathbf{a}_t, \mathbf{u}_t) \mathbb{1}_{k_t=k_{t+1}}$, an initial state density $p_1(\mathbf{s}_1) p(k_1)$, and a reward function $l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k)$. Here, $\mathbb{1}_{a=b}$ is the indicator function which equals to 1 if $a = b$ and 0 otherwise. By adopting the optimality results of MaxEnt-RL (Ziebart et al., 2010; Haarnoja et al., 2018), we have $g(\phi, \omega) = \max_{\theta} \mathcal{G}(\phi, \omega, \theta)$, where the optimal variational distribution is

$$q_{\theta^*}(\mathbf{a}_t, \mathbf{u}_t | \mathbf{s}_t, k) = \exp(Q(\mathbf{s}_t, k, \mathbf{a}_t, \mathbf{u}_t) - V(\mathbf{s}_t, k)). \tag{10}$$

The functions Q and V are soft-value functions defined as

$$Q(\mathbf{s}_t, k, \mathbf{a}_t, \mathbf{u}_t) = l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k) + \mathbb{E}_{p(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{u}_t)} [V(\mathbf{s}_{t+1}, k)], \quad (11)$$

$$V(\mathbf{s}_t, k) = \log \iint_{\mathcal{A} \times \mathcal{A}} \exp(Q(\mathbf{s}_t, k, \mathbf{a}_t, \mathbf{u}_t)) d\mathbf{a}_t d\mathbf{u}_t. \quad (12)$$

A.3 OBJECTIVE FUNCTION \mathcal{H} OF VILD

This section derives the objective function $\mathcal{H}(\phi, \omega, \psi, \theta)$ from $\mathcal{F}(\phi, \omega, \psi) - \mathcal{G}(\phi, \omega, \theta)$. Specifically, we substitute the models $p_{\omega}(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k) = \mathcal{N}(\mathbf{u}_t|\mathbf{a}_t, \mathbf{C}_{\omega}(k))$ and $q_{\theta}(\mathbf{a}_t, \mathbf{u}_t|\mathbf{s}_t, k) = q_{\theta}(\mathbf{a}_t|\mathbf{s}_t)\mathcal{N}(\mathbf{u}_t|\mathbf{a}_t, \Sigma)$. We also give an example when using a Laplace distribution for $p_{\omega}(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k)$ instead of the Gaussian distribution.

First, we substitute $q_{\theta}(\mathbf{a}_t, \mathbf{u}_t|\mathbf{s}_t, k) = q_{\theta}(\mathbf{a}_t|\mathbf{s}_t)\mathcal{N}(\mathbf{u}_t|\mathbf{a}_t, \Sigma)$ into \mathcal{G} :

$$\begin{aligned} \mathcal{G}(\phi, \omega, \theta) &= \mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)} \left[\sum_{t=1}^T l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k) - \log \mathcal{N}(\mathbf{u}_t|\mathbf{a}_t, \Sigma) - \log q_{\theta}(\mathbf{a}_t|\mathbf{s}_t) \right] \\ &= \mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)} \left[\sum_{t=1}^T l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k) + \frac{1}{2} \|\mathbf{u}_t - \mathbf{a}_t\|_{\Sigma^{-1}}^2 - \log q_{\theta}(\mathbf{a}_t|\mathbf{s}_t) \right] + c_1, \end{aligned}$$

where c_1 is a constant corresponding to the log-normalization term of the Gaussian distribution. Next, by using the re-parameterization trick, we rewrite $\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)$ as

$$\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k) = p(k)p_1(\mathbf{s}_1) \prod_{t=1}^T p_1(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t + \Sigma^{1/2}\epsilon_t)\mathcal{N}(\epsilon_t|0, \mathbf{I})q_{\theta}(\mathbf{a}_t|\mathbf{s}_t),$$

where we use $\mathbf{u}_t = \mathbf{a}_t + \Sigma^{1/2}\epsilon_t$ with $\epsilon_t \sim \mathcal{N}(\epsilon_t|0, \mathbf{I})$. With this, the expectation of $\sum_{t=1}^T \|\mathbf{u}_t - \mathbf{a}_t\|_{\Sigma^{-1}}^2$ over $\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)$ can be written as

$$\begin{aligned} \mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)} \left[\sum_{t=1}^T \|\mathbf{u}_t - \mathbf{a}_t\|_{\Sigma^{-1}}^2 \right] &= \mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)} \left[\sum_{t=1}^T \|\mathbf{a}_t + \Sigma^{1/2}\epsilon_t - \mathbf{a}_t\|_{\Sigma^{-1}}^2 \right] \\ &= \mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)} \left[\sum_{t=1}^T \|\Sigma^{1/2}\epsilon_t\|_{\Sigma^{-1}}^2 \right] \\ &= Td_{\mathbf{a}}, \end{aligned}$$

which is a constant. Then, the quantity \mathcal{G} can be expressed as

$$\mathcal{G}(\phi, \omega, \theta) = \mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)} \left[\sum_{t=1}^T l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k) - \log q_{\theta}(\mathbf{a}_t|\mathbf{s}_t) \right] + c_1 + Td_{\mathbf{a}}.$$

By ignoring the constant, the optimization problem $\max_{\phi, \omega, \psi} \min_{\theta} \mathcal{F}(\phi, \omega, \psi) - \mathcal{G}(\phi, \omega, \theta)$ is equivalent to

$$\begin{aligned} \max_{\phi, \omega, \psi} \min_{\theta} \mathbb{E}_{p_d(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}|k)p(k)} \left[\sum_{t=1}^T \mathbb{E}_{q_{\psi}(\mathbf{a}_t|\mathbf{s}_t, \mathbf{u}_t, k)} [l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k) - \log q_{\psi}(\mathbf{a}_t|\mathbf{s}_t, \mathbf{u}_t, k)] \right] \\ - \mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)} \left[\sum_{t=1}^T l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k) - \log q_{\theta}(\mathbf{a}_t|\mathbf{s}_t) \right]. \quad (13) \end{aligned}$$

Our next step is to substitute $p_{\omega}(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k)$ by our choice of model. First, let us consider a Gaussian distribution $p_{\omega}(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k) = \mathcal{N}(\mathbf{u}_t|\mathbf{a}_t, \mathbf{C}_{\omega}(\mathbf{s}_t, k))$, where the covariance depends on state. With this model, the second term in Eq. (13) is given by

$$\begin{aligned} \mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)} \left[\sum_{t=1}^T l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k) - \log q_{\theta}(\mathbf{a}_t|\mathbf{s}_t) \right] \\ = \mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)} \left[\sum_{t=1}^T r_{\phi}(\mathbf{s}_t, \mathbf{a}_t) + \log \mathcal{N}(\mathbf{u}_t|\mathbf{a}_t, \mathbf{C}_{\omega}(\mathbf{s}_t, k)) - \log q_{\theta}(\mathbf{a}_t|\mathbf{s}_t) \right] \\ = \mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)} \left[\sum_{t=1}^T r_{\phi}(\mathbf{s}_t, \mathbf{a}_t) - \frac{1}{2} \|\mathbf{u}_t - \mathbf{a}_t\|_{\mathbf{C}_{\omega}^{-1}(\mathbf{s}_t, k)}^2 - \frac{1}{2} \log |\mathbf{C}_{\omega}(\mathbf{s}_t, k)| - \log q_{\theta}(\mathbf{a}_t|\mathbf{s}_t) \right] + c_2, \end{aligned}$$

where $c_2 = -\frac{d_{\mathbf{a}}}{2} \log 2\pi$ is a constant. By using the reparameterization trick, we write the expectation of $\sum_{t=1}^T \|\mathbf{u}_t - \mathbf{a}_t\|_{\mathbf{C}_{\omega}^{-1}(\mathbf{s}_t, k)}^2$ as follows:

$$\begin{aligned} \mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)} \left[\sum_{t=1}^T \|\mathbf{u}_t - \mathbf{a}_t\|_{\mathbf{C}_{\omega}^{-1}(\mathbf{s}_t, k)}^2 \right] &= \mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)} \left[\sum_{t=1}^T \|\mathbf{a}_t + \boldsymbol{\Sigma}^{1/2} \boldsymbol{\epsilon}_t - \mathbf{a}_t\|_{\mathbf{C}_{\omega}^{-1}(\mathbf{s}_t, k)}^2 \right] \\ &= \mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)} \left[\sum_{t=1}^T \|\boldsymbol{\Sigma}^{1/2} \boldsymbol{\epsilon}_t\|_{\mathbf{C}_{\omega}^{-1}(\mathbf{s}_t, k)}^2 \right]. \end{aligned}$$

Using this equality, the second term in Eq. (13) is given by

$$\mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)} \left[\sum_{t=1}^T r_{\phi}(\mathbf{s}_t, \mathbf{a}_t) - \log q_{\theta}(\mathbf{a}_t | \mathbf{s}_t) - \frac{1}{2} \left(\|\boldsymbol{\Sigma}^{1/2} \boldsymbol{\epsilon}_t\|_{\mathbf{C}_{\omega}^{-1}(\mathbf{s}_t, k)}^2 + \log |\mathbf{C}_{\omega}(\mathbf{s}_t, k)| \right) \right]. \quad (14)$$

Maximizing this quantity w.r.t. θ has an implication as follows: $q_{\theta}(\mathbf{a}_t | \mathbf{s}_t)$ maximizes the expected cumulative reward while avoiding states that are difficult for demonstrators. Specifically, a large value of $\mathbb{E}_{p(k)} [\log |\mathbf{C}_{\omega}(\mathbf{s}_t, k)|]$ indicates that demonstrators have a low level of expertise for state \mathbf{s}_t on average, given by our estimated covariance. In other words, this state is difficult to accurately execute optimal actions for all demonstrators on averages. Since the policy $q_{\theta}(\mathbf{a}_t | \mathbf{s}_t)$ should minimize $\mathbb{E}_{p(k)} [\log |\mathbf{C}_{\omega}(\mathbf{s}_t, k)|]$, the policy should avoid states that are difficult for demonstrators. We expect that this property may improve exploration-exploitation trade-off in IL. Nonetheless, we leave an investigation of this property for future work, since this is not in the scope of the paper.

In this paper, we specify that the covariance does not depend on state: $\mathbf{C}_{\omega}(\mathbf{s}_t, k) = \mathbf{C}_{\omega}(k)$. This model specification enables us to simplify Eq. (14) as follows:

$$\begin{aligned} &\mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)} \left[\sum_{t=1}^T r_{\phi}(\mathbf{s}_t, \mathbf{a}_t) - \log q_{\theta}(\mathbf{a}_t | \mathbf{s}_t) - \frac{1}{2} \left(\|\boldsymbol{\Sigma}^{1/2} \boldsymbol{\epsilon}_t\|_{\mathbf{C}_{\omega}^{-1}(k)}^2 + \log |\mathbf{C}_{\omega}(k)| \right) \right] \\ &= \mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, \mathbf{a}_{1:T}, k)} \left[\sum_{t=1}^T r_{\phi}(\mathbf{s}_t, \mathbf{a}_t) - \log q_{\theta}(\mathbf{a}_t | \mathbf{s}_t) \right] - \frac{T}{2} \mathbb{E}_{p(k) \mathcal{N}(\boldsymbol{\epsilon}_t | 0, \mathbf{I})} \left[\|\boldsymbol{\Sigma}^{1/2} \boldsymbol{\epsilon}_t\|_{\mathbf{C}_{\omega}^{-1}(k)}^2 + \log |\mathbf{C}_{\omega}(k)| \right] \\ &= \mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})} \left[\sum_{t=1}^T r_{\phi}(\mathbf{s}_t, \mathbf{a}_t) - \log q_{\theta}(\mathbf{a}_t | \mathbf{s}_t) \right] - \frac{T}{2} \mathbb{E}_{p(k)} [\text{Tr}(\mathbf{C}_{\omega}^{-1}(k) \boldsymbol{\Sigma}) + \log |\mathbf{C}_{\omega}(k)|], \end{aligned}$$

where $\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{a}_{1:T}) = p_1(\mathbf{s}_1) \prod_{t=1}^T \int_{\mathcal{A}} p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{u}_t) \mathcal{N}(\mathbf{u}_t | \mathbf{a}_t, \boldsymbol{\Sigma}) d\mathbf{u}_t q_{\theta}(\mathbf{a}_t | \mathbf{s}_t)$. The last line follows from the quadratic form identity: $\mathbb{E}_{\mathcal{N}(\boldsymbol{\epsilon}_t | 0, \mathbf{I})} [\|\boldsymbol{\Sigma}^{1/2} \boldsymbol{\epsilon}_t\|_{\mathbf{C}_{\omega}^{-1}(k)}^2] = \text{Tr}(\mathbf{C}_{\omega}^{-1}(k) \boldsymbol{\Sigma})$. Next, we substitute $p_{\omega}(\mathbf{u}_t | \mathbf{s}_t, \mathbf{a}_t, k) = \mathcal{N}(\mathbf{u}_t | \mathbf{a}_t, \mathbf{C}_{\omega}(k))$ into the first term of Eq. (13).

$$\begin{aligned} &\mathbb{E}_{p_d(\mathbf{s}_{1:T}, \mathbf{u}_{1:T} | k) p(k)} \left[\sum_{t=1}^T \mathbb{E}_{q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k)} [l_{\phi, \omega}(\mathbf{s}_t, \mathbf{a}_t, \mathbf{u}_t, k) - \log q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k)] \right] \\ &= \mathbb{E}_{p_d(\mathbf{s}_{1:T}, \mathbf{u}_{1:T} | k) p(k)} \left[\sum_{t=1}^T \mathbb{E}_{q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k)} \left[r_{\phi}(\mathbf{s}_t, \mathbf{a}_t) - \frac{1}{2} \|\mathbf{u}_t - \mathbf{a}_t\|_{\mathbf{C}_{\omega}^{-1}(k)}^2 - \frac{1}{2} \log |\mathbf{C}_{\omega}(k)| \right. \right. \\ &\quad \left. \left. - \log q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k) \right] \right] - T d_{\mathbf{a}} \log 2\pi / 2. \end{aligned} \quad (15)$$

Lastly, by ignoring constants, Eq. (13) is equivalent to $\max_{\phi, \omega, \psi} \min_{\theta} \mathcal{H}(\phi, \omega, \psi, \theta)$, where

$$\begin{aligned} \mathcal{H}(\phi, \omega, \psi, \theta) &= \mathbb{E}_{p_d(\mathbf{s}_{1:T}, \mathbf{u}_{1:T} | k) p(k)} \left[\sum_{t=1}^T \mathbb{E}_{q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k)} \left[r_{\phi}(\mathbf{s}_t, \mathbf{a}_t) - \frac{1}{2} \|\mathbf{u}_t - \mathbf{a}_t\|_{\mathbf{C}_{\omega}^{-1}(k)}^2 - \log q_{\psi}(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k) \right] \right] \\ &\quad - \mathbb{E}_{\tilde{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})} \left[\sum_{t=1}^T r_{\phi}(\mathbf{s}_t, \mathbf{a}_t) - \log q_{\theta}(\mathbf{a}_t | \mathbf{s}_t) \right] + \frac{T}{2} \mathbb{E}_{p(k)} [\text{Tr}(\mathbf{C}_{\omega}^{-1}(k) \boldsymbol{\Sigma})]. \end{aligned}$$

This concludes the derivation of VILD.

As mentioned, other distributions beside the Gaussian distribution can be used for p_{ω} . For instance, let us consider a multivariate-independent Laplace distribution: $p_{\omega}(\mathbf{u}_t | \mathbf{s}_t, \mathbf{a}_t, k) = \prod_{d=1}^{d_{\mathbf{a}}} \frac{1}{2c_k^{(d)}} \exp(-\|\frac{\mathbf{u}_t - \mathbf{a}_t}{\mathbf{c}_k}\|_1)$, where a division of vector by vector denotes element-wise division. The Laplace distribution has heavier tails when compared to the Gaussian distribution, which makes the Laplace distribution more suitable for modeling

demonstrators who tend to execute outlier actions. By using the Laplace distribution for $p_\omega(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k)$, we obtain an objective

$$\begin{aligned} \mathcal{H}_{\text{Lap}} = & \mathbb{E}_{p_d(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, k)} \left[\sum_{t=1}^T \mathbb{E}_{q_\psi(\mathbf{a}_t|\mathbf{s}_t, \mathbf{u}_t, k)} \left[r_\phi(\mathbf{s}_t, \mathbf{a}_t) - \left\| \frac{\mathbf{u}_t - \mathbf{a}_t}{\mathbf{c}_k} \right\|_1 - \log q_\psi(\mathbf{a}_t|\mathbf{s}_t, \mathbf{u}_t, k) \right] \right] \\ & - \mathbb{E}_{\tilde{q}_\theta(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})} \left[\sum_{t=1}^T r_\phi(\mathbf{s}_t, \mathbf{a}_t) - \log q_\theta(\mathbf{a}_t|\mathbf{s}_t) \right] + \frac{T\sqrt{2}}{\sqrt{\pi}} \mathbb{E}_{p(k)} \left[\text{Tr}(\mathbf{C}_\omega^{-1}(k)\boldsymbol{\Sigma}^{1/2}) \right]. \end{aligned}$$

We can see that differences between \mathcal{H}_{Lap} and \mathcal{H} are the absolute error and scaling of the trace term.

B IMPLEMENTATION DETAILS

We implement VILD using the PyTorch deep learning framework. For all function approximators, we use neural networks with 2 hidden-layers of 100 tanh units, except for the Humanoid task where we use neural networks with 2 hidden-layers of 100 relu units. We optimize parameters ϕ , ω , and ψ by Adam with step-size 3×10^{-4} , $\beta_1 = 0.9$, $\beta_2 = 0.999$ and mini-batch size 256. To optimize the policy parameter θ , we use trust region policy optimization (TRPO) (Schulman et al., 2015) with batch size 1000, except on the Humanoid task where we use soft actor-critic (SAC) (Haarnoja et al., 2018) with mini-batch size 256. Note that TRPO is an on-policy RL method that uses only trajectories collected by the current policy, while SAC is an off-policy RL method that use trajectories collected by previous policies. On-policy methods are generally more stable than off-policy methods, while off-policy methods are generally more data-efficient (Gu et al., 2017). We use SAC for Humanoid mainly due to its high data-efficiency. When SAC is used, we also use trajectories collected by previous policies to approximate the expectation over the trajectory density $\tilde{q}_\theta(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})$.

For the distribution $p_\omega(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k) = \mathcal{N}(\mathbf{u}_t|\mathbf{a}_t, \mathbf{C}_\omega(k))$, we use diagonal covariances $\mathbf{C}_\omega(k) = \text{diag}(\mathbf{c}_k)$, where $\omega = \{\mathbf{c}_k\}_{k=1}^K$ and $\mathbf{c}_k \in \mathbb{R}^{d_a}$ are parameter vectors to be learned. For the distribution $q_\psi(\mathbf{a}_t|\mathbf{s}_t, \mathbf{u}_t, k)$, we use a Gaussian distribution with diagonal covariance, where the mean and logarithm of the standard deviation are the outputs of neural networks. Since k is a discrete variable, we represent $q_\psi(\mathbf{a}_t|\mathbf{s}_t, \mathbf{u}_t, k)$ by neural networks that have K output heads and take input vectors $(\mathbf{s}_t, \mathbf{u}_t)$; The k -th output head corresponds to (the mean and log-standard-deviation of) $q_\psi(\mathbf{a}_t|\mathbf{s}_t, \mathbf{u}_t, k)$. We also pre-train the mean function of $q_\psi(\mathbf{a}_t|\mathbf{s}_t, \mathbf{u}_t, k)$, by performing least-squares regression for 1000 gradient steps with target value \mathbf{u}_t . This pre-training is done to obtain reasonable initial predictions. For the policy $q_\theta(\mathbf{a}_t|\mathbf{s}_t)$, we use a Gaussian policy with diagonal covariance, where the mean and logarithm of the standard deviation are outputs of neural networks. We use $\boldsymbol{\Sigma} = 10^{-8}\mathbf{I}$ in experiments.

To control exploration-exploitation trade-off, we use an entropy coefficient $\alpha = 0.0001$ in TRPO. In SAC, the value of α is optimized so that the policy has a certain value of entropy, as described by Haarnoja et al. (2018). Note that including α in VILD is equivalent to rescaling quantities in the model by α , i.e., $\exp(r_\phi(\mathbf{s}_t, \mathbf{a}_t)/\alpha)$ and $(p_\omega(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k))^{1/\alpha}$. A discount factor $0 < \gamma < 1$ may be included similarly, and we use $\gamma = 0.99$ in experiments.

For all methods, we regularize the reward/discriminator function by the gradient penalty (Gulrajani et al., 2017) with coefficient 10, since it was previously shown to improve performance of generative adversarial learning methods. For methods that learn a reward function, namely VILD, AIRL, and MaxEnt-IRL, we apply a sigmoid function to the output of a reward network to bound reward values. We found that without the bounds, reward values of the agent can be highly negative in the early stage of learning, which makes RL methods prematurely converge to poor policies. An explanation of this phenomenon is that, in MDPs with large state and action spaces, distribution of demonstrations and distribution of agent’s trajectories are not overlapped in the early stage of learning. In such a scenario, it is trivial to learn a reward function which tends to positive-infinity values for demonstrations and negative-infinity values for agent’s trajectories. While the gradient penalty regularizer slightly remedies this issue, we found that the regularizer alone is insufficient to prevent this scenario. Moreover, for VILD, it is beneficial to bound the reward function to control a trade-off between the immediate reward and the squared error when optimizing ψ .

A pseudo-code of VILD with IS is given in Algorithm 1, where the reward parameter is updated by IS gradient in line 8. For VILD without IS, the reward parameter is instead updated by an estimate of $\nabla_\phi \mathcal{H}(\phi, \omega, \psi, \theta)$. The regularizer $L(\omega) = T\mathbb{E}_{p(k)}[\log|\mathbf{C}_\omega^{-1}(k)|]/2$ penalizes large value of $\mathbf{C}_\omega(k)$. A source-code of our implementation will be publicly available.

C EXPERIMENT DETAILS

In this section, we describe experimental settings and data generation. We also give brief reviews of methods compared against VILD in the experiments.

Algorithm 1 VILD: Variational Imitation Learning with Diverse-quality demonstrations

-
- 1: **Input:** Diverse-quality demonstrations $\mathcal{D}_d = \{(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, k)\}_{n=1}^N$ and a replay buffer $\mathcal{B} = \emptyset$.
 - 2: **while** Not converge **do**
 - 3: **while** $|\mathcal{B}| < B$ with batch size B **do** \triangleright Collect samples from $\tilde{q}_\theta(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})$
 - 4: Sample $\mathbf{a}_t \sim q_\theta(\mathbf{a}_t|\mathbf{s}_t)$ and $\epsilon_t \sim \mathcal{N}(\epsilon_t|\mathbf{0}, \Sigma)$.
 - 5: Execute $\mathbf{a}_t + \epsilon_t$, observe $\mathbf{s}'_t \sim p(\mathbf{s}'_t|\mathbf{s}_t, \mathbf{a}_t + \epsilon_t)$, and include $(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}'_t)$ into \mathcal{B}
 - 6: Update q_ψ by an estimate of $\nabla_\psi \mathcal{H}(\phi, \omega, \psi, \theta)$.
 - 7: Update p_ω by an estimate of $\nabla_\omega \mathcal{H}(\phi, \omega, \psi, \theta) + \nabla_\omega L(\omega)$.
 - 8: Update r_ϕ by an estimate of $\nabla_\phi \mathcal{H}_{\text{IS}}(\phi, \omega, \psi, \theta)$.
 - 9: Update q_θ by an RL method (e.g., TRPO or SAC) with reward function r_ϕ .
-

C.1 SETTINGS AND DATA GENERATION

We evaluate VILD on four continuous-control benchmark tasks from OpenAI gym platform (Brockman et al., 2016) with the Mujoco physics simulator: HalfCheetah, Ant, Walker2d, and Humanoid. To obtain the optimal policy for generating demonstrations, we use the ground-truth reward function of each task to pre-train π^* with TRPO. We generate diverse-quality demonstrations by using $K = 10$ demonstrators according to the graphical model in Figure 1(b). We consider two types of the noisy policy $p(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k)$: a Gaussian noisy policy and a time-signal-dependent (TSD) noisy policy.

Gaussian noisy policy. We use a Gaussian noisy policy $\mathcal{N}(\mathbf{u}_t|\mathbf{a}_t, \sigma_k^2 \mathbf{I})$ with a constant covariance. The value of σ_k for each of the 10 demonstrators is 0.01, 0.05, 0.1, 0.25, 0.4, 0.6, 0.7, 0.8, 0.9 and 1.0, respectively. Note that our model assumption on p_ω corresponds to this Gaussian noisy policy. Table 2 shows the performance of demonstrators (in terms of cumulative ground-truth rewards) with this Gaussian noisy policy. A random policy π_0 is an initial policy neural network for learning; The network weights are initialized such that the magnitude of actions is small. Note that this initialization scheme is a common practice in deep RL (Gu et al., 2017).

TSD noisy policy. To make learning more challenging, we generate demonstrations according to a noise characteristic of human motor control, where a magnitude of noises is proportion to a magnitude of actions and increases with execution time (van Beers et al., 2004). Specifically, we generate demonstrations using a Gaussian distribution $\mathcal{N}(\mathbf{u}_t|\mathbf{a}_t, \text{diag}(\mathbf{b}_k(t) \times \|\mathbf{a}_t\|_1/d_a))$, where the covariance is proportion to the magnitude of action and depends on time steps and \times denotes an element-wise product. We call this policy time-signal-dependent (TSD) noisy policy. Here, $\mathbf{b}_k(t)$ is a sample of a noise process whose noise variance increases over time, as shown in Figure 3. We obtain this noise process for the k -th demonstrator by reversing Ornstein–Uhlenbeck (OU) processes with parameters $\theta = 0.15$ and $\sigma = \sigma_k$ (Uhlenbeck & Ornstein, 1930)⁴. The value of σ_k for each demonstrator is 0.01, 0.05, 0.1, 0.25, 0.4, 0.6, 0.7, 0.8, 0.9, and 1.0, respectively. Table 3 shows the performance of demonstrators with this TSD noisy policy. Learning from demonstrations generated by TSD is challenging; The Gaussian model of p_ω cannot perfectly model the TSD noisy policy, since the ground-truth variance is a function of actions and time steps.

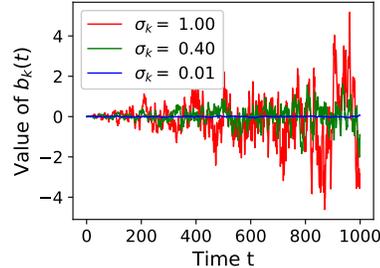


Figure 3: Samples $\mathbf{b}_k(t)$ drawn from noise processes used for the TSD noisy policy.

C.2 COMPARISON METHODS

Here, we briefly review methods compared against VILD in our experiments. We firstly review online IL methods, which learn a policy by RL and require additional transition samples from MDPs.

MaxEnt-IRL. Maximum (causal) entropy IRL (MaxEnt-IRL) (Ziebart et al., 2010) is a well-known IRL method. The original derivation of the method is based on the maximum entropy principle (Jaynes, 1957) but for causal interactions, and uses a linear-in-parameter reward function: $r_\phi(\mathbf{s}_t, \mathbf{a}_t) = \phi^\top b(\mathbf{s}_t, \mathbf{a}_t)$ with a basis function b . Here, we consider an alternative derivation which is applicable to non-linear reward function (Finn et al., 2016). Briefly speaking, MaxEnt-IRL learns a reward parameter by minimizing a KL divergence from a data distribution $p^*(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})$ to a model $p_\phi(\mathbf{s}_{1:T}, \mathbf{a}_{1:T}) =$

⁴OU process is commonly used to generate time-correlated noises where the noise variance decays towards zero. We reserve this process along the time axis, so that the noise variance grows over time.

Table 2: Performance of a random policy π_0 , the optimal policy π^* , and demonstrators with the Gaussian noisy policy.

σ_k	Cheetah	Ant	Walker	Humanoid
(π_0)	-0.58	995	131	222
(π^*)	4624	4349	4963	5093
0.01	4311	3985	4434	4315
0.05	3978	3861	3486	5140
0.01	4019	3514	4651	5189
0.25	1853	536	4362	3628
0.40	1090	227	467	5220
0.6	567	-73	523	2593
0.7	267	-208	332	1744
0.8	-45	-979	283	735
0.9	-399	-328	255	538
1.0	-177	-203	249	361

Table 3: Performance of a random policy π_0 , the optimal policy π^* , and demonstrators with the TSD noisy policy.

σ_k	Cheetah	Ant	Walker	Humanoid
(π_0)	-0.58	995	131	222
(π^*)	4624	4349	4963	5093
0.01	4362	3758	4695	5130
0.05	4015	3623	4528	5099
0.01	3741	3368	2362	5195
0.25	1301	873	644	1675
0.40	-203	231	302	610
0.6	-230	-51	29	249
0.7	-249	-37	24	221
0.8	-416	-567	14	191
0.9	-389	-751	7	178
1.0	-424	-269	4	169

$\frac{1}{Z_\phi} p_1(\mathbf{s}_1) \prod_{t=1}^T p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t) \exp(r_\phi(\mathbf{s}_t, \mathbf{a}_t) / \alpha)$, where Z_ϕ is the normalization term. Minimizing this KL divergence is equivalent to solving $\max_{\phi} \mathbb{E}_{p^*(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})} [\sum_{t=1}^T r_\phi(\mathbf{s}_t, \mathbf{a}_t)] - \log Z_\phi$. To compute $\log Z_\phi$, we can use the importance sampling approach (Finn et al., 2016) or the variational approach as done in VILD. The latter leads to a max-min problem

$$\max_{\phi} \min_{\theta} \mathbb{E}_{p^*(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})} \left[\sum_{t=1}^T r_\phi(\mathbf{s}_t, \mathbf{a}_t) \right] - \mathbb{E}_{q_\theta(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})} \left[\sum_{t=1}^T r_\phi(\mathbf{s}_t, \mathbf{a}_t) - \alpha \log q_\theta(\mathbf{a}_t | \mathbf{s}_t) \right],$$

where $q_\theta(\mathbf{s}_{1:T}, \mathbf{a}_{1:T}) = p_1(\mathbf{s}_1) \prod_{t=1}^T p(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t) q_\theta(\mathbf{a}_t | \mathbf{s}_t)$. The policy $q_\theta(\mathbf{a}_t | \mathbf{s}_t)$ maximizes the learned reward function and is the solution of IL.

As we mentioned, the proposed model in VILD is based on the model of MaxEnt-IRL. By comparing the max-min problem of MaxEnt-IRL and the max-min problem of VILD, we can see that the main difference are the variational distribution q_ψ and the noisy policy model p_ω . If we assume that q_ψ and p_ω are Dirac delta functions: $q_\psi(\mathbf{a}_t | \mathbf{s}_t, \mathbf{u}_t, k) = \delta_{\mathbf{a}_t = \mathbf{u}_t}$ and $p_\omega(\mathbf{u}_t | \mathbf{a}_t, \mathbf{s}_t, k) = \delta_{\mathbf{u}_t = \mathbf{a}_t}$, then the max-min problem of VILD reduces to the max-min problem of MaxEnt-IRL. In other words, if we assume that all demonstrators execute the same optimal policy and have an equal level of expertise, then VILD reduces to MaxEnt-IRL.

GAIL. Generative adversarial IL (GAIL) (Ho & Ermon, 2016) performs occupancy measure matching via generative adversarial networks (GAN) to learn the optimal policy from expert demonstrations. Specifically, GAIL finds a parameterized policy π_θ such that the occupancy measure $\rho_{\pi_\theta}(\mathbf{s}, \mathbf{a})$ of π_θ is similar to the occupancy measure $\rho_{\pi^*}(\mathbf{s}, \mathbf{a})$ of π^* . Here, $\rho_\pi(\mathbf{s}, \mathbf{a}) = \mathbb{E}_{p_\pi(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})} [\sum_{t=0}^T \delta(\mathbf{s}_t - \mathbf{s}, \mathbf{a}_t - \mathbf{a})]$ is the state-action occupancy measure of π and satisfies the equality $\mathbb{E}_{p_\pi(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})} [\sum_{t=1}^T r(\mathbf{s}_t, \mathbf{a}_t)] = \iint_{\mathcal{S} \times \mathcal{A}} \rho_\pi(\mathbf{s}, \mathbf{a}) r(\mathbf{s}, \mathbf{a}) d\mathbf{s} d\mathbf{a} = \mathbb{E}_\pi [r(\mathbf{s}, \mathbf{a})]$. To measure the similarity, GAIL uses the Jensen-Shannon divergence, which is estimated and minimized by the following generative-adversarial training objective:

$$\min_{\theta} \max_{\phi} \mathbb{E}_{\rho_{\pi^*}} [\log D_\phi(\mathbf{s}, \mathbf{a})] + \mathbb{E}_{\rho_{\pi_\theta}} [\log(1 - D_\phi(\mathbf{s}, \mathbf{a})) + \alpha \log \pi_\theta(\mathbf{a}_t | \mathbf{s}_t)],$$

where $D_\phi(\mathbf{s}, \mathbf{a}) = \frac{d_\phi(\mathbf{s}, \mathbf{a})}{d_\phi(\mathbf{s}, \mathbf{a}) + 1}$ is called a discriminator. The minimization problem w.r.t. θ is achieved using RL with a reward function $-\log(1 - D_\phi(\mathbf{s}, \mathbf{a}))$.

AIRL. Adversarial IRL (AIRL) (Fu et al., 2018) was proposed to overcome a limitation of GAIL regarding reward function: GAIL does not learn the expert reward function, since GAIL has $D_\phi(\mathbf{s}, \mathbf{a}) = 0.5$ at the saddle point for every states and actions. To overcome this limitation while taking advantage of generative-adversarial training, AIRL learns a reward function by solving

$$\max_{\phi} \mathbb{E}_{p^*(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})} \left[\sum_{t=1}^T \log D_\phi(\mathbf{s}, \mathbf{a}) \right] + \mathbb{E}_{q_\theta(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})} \left[\sum_{t=1}^T \log(1 - D_\phi(\mathbf{s}, \mathbf{a})) \right],$$

where $D_\phi(\mathbf{s}, \mathbf{a}) = \frac{r_\phi(\mathbf{s}, \mathbf{a})}{r_\phi(\mathbf{s}, \mathbf{a}) + q_\theta(\mathbf{a} | \mathbf{s})}$. The policy $q_\theta(\mathbf{a}_t | \mathbf{s}_t)$ is learned by RL with a reward function $r_\phi(\mathbf{s}_t, \mathbf{a}_t)$. Fu et al. (2018) showed that the gradient of this objective w.r.t. ϕ is equivalent to the gradient of MaxEnt-IRL w.r.t. ϕ . The authors also proposed an approach to disentangle reward function, which leads to a better performance in transfer learning settings. Nonetheless, this disentangle approach is general and can be applied to other IRL methods, including MaxEnt-IRL and VILD. We do not evaluate AIRL with disentangle reward function.

We note that, based on the relation between MaxEnt-IRL and VILD, we can extend VILD to use a training procedure of AIRL. Specifically, by applying the same derivation from MaxEnt-IRL to AIRL

by Fu et al. (2018), we can derive a variant of VILD which learns a reward parameter by solving $\max_{\phi} \mathbb{E}_{p_d(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}|k)p(k)} [\sum_{t=1}^T \mathbb{E}_{q_{\phi}(\mathbf{a}_t|\mathbf{s}_t, \mathbf{u}_t, k)} [\log D_{\phi}(\mathbf{s}, \mathbf{a})]] + \mathbb{E}_{\hat{q}_{\theta}(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})} [\sum_{t=1}^T \log(1 - D_{\phi}(\mathbf{s}, \mathbf{a}))]$. We do not evaluate this variant of VILD in our experiment.

VAIL. Variational adversarial IL (VAIL) (Peng et al., 2019) improves upon GAIL by using variational information bottleneck (VIB) (Alemi et al., 2017). VIB aims to compress information flow by minimizing a variational bound of mutual information. This compression filters irrelevant signals, which leads to less over-fitting. To achieve this in GAIL, VAIL learns the discriminator D_{ϕ} by an optimization problem

$$\min_{\phi, E} \max_{\beta \geq 0} \mathbb{E}_{\rho_{\pi^*}} [\mathbb{E}_{E(\mathbf{z}|\mathbf{s}, \mathbf{a})} [-\log D_{\phi}(\mathbf{z})]] + \mathbb{E}_{\rho_{\pi_{\theta}}} [\mathbb{E}_{E(\mathbf{z}|\mathbf{s}, \mathbf{a})} [-\log(1 - D_{\phi}(\mathbf{z}))]] \\ + \beta \mathbb{E}_{(\rho_{\pi^*} + \rho_{\pi_{\theta}})/2} [\text{KL}(E(\mathbf{z}|\mathbf{s}, \mathbf{a})|p(\mathbf{z})) - I_c],$$

where \mathbf{z} is an encode vector, $E(\mathbf{z}|\mathbf{s}, \mathbf{a})$ is an encoder, $p(\mathbf{z})$ is a prior distribution of \mathbf{z} , I_c is the target value of mutual information, and $\beta > 0$ is a Lagrange multiplier. With this discriminator, the policy $\pi_{\theta}(\mathbf{a}_t|\mathbf{s}_t)$ is learned by RL with a reward function $-\log(1 - D_{\phi}(\mathbb{E}_{E(\mathbf{z}|\mathbf{s}, \mathbf{a})} [\mathbf{z}]))$.

It might be expected that the compression may make VAIL robust against diverse-quality demonstrations, since irrelevant signals in low-quality demonstrations are filtered out via the encoder. However, we find that this is not the case, and VAIL does not improve much upon GAIL in our experiments. This is perhaps because VAIL compress information from both demonstrators and agent’s trajectories. Meanwhile in our setting, irrelevant signals are generated only by demonstrators. Therefore, the information bottleneck may also filter out relevant signals in agent’s trajectories, which lead to poor performances.

InfoGAIL. Information maximizing GAIL (InfoGAIL) (Li et al., 2017) is an extension of GAIL for learning a multi-modal policy in MM-IL. The key idea of InfoGAIL is to introduce a context variable z to the GAIL formulation and learn a context-dependent policy $\pi_{\theta}(\mathbf{a}|\mathbf{s}, z)$, where each context represents each mode of the multi-modal policy. To ensure that the context is not ignored during learning, InfoGAIL regularizes GAIL’s objective so that a mutual information between contexts and state-action variables is maximized. This mutual information is indirectly maximized via maximizing a variational lower-bound of mutual information. By doing so, InfoGAIL solves a min-max problem

$$\min_{\theta, Q} \max_{\phi} \mathbb{E}_{\rho_{\pi^*}} [\log D_{\phi}(\mathbf{s}, \mathbf{a})] + \mathbb{E}_{\rho_{\pi_{\theta}}} [\log(1 - D_{\phi}(\mathbf{s}, \mathbf{a})) + \alpha \log \pi_{\theta}(\mathbf{a}|\mathbf{s}, z)] + \lambda L(\pi_{\theta}, Q),$$

where $L(\pi_{\theta}, Q) = \mathbb{E}_{p(z)\pi_{\theta}(\mathbf{a}|\mathbf{s}, z)} [\log Q(z|\mathbf{s}, \mathbf{a}) - \log p(z)]$ is a lower-bound of mutual information, $Q(z|\mathbf{s}, \mathbf{a})$ is an encoder neural network, and $p(z)$ is a prior distribution of contexts. In our experiment, the number of context z is set to be the number of demonstrators K . As discussed in Section 1, when knowing the level of demonstrators’ expertise, we may choose contexts that correspond to high-expertise demonstrator. In other words, we may hand-craft the prior distribution $p(z)$ so that a probability of contexts is proportion to the level of demonstrators’ expertise. Nonetheless, for fair comparison, we do not use the oracle knowledge about the level of demonstrators’ expertise, and set $p(z)$ to be a uniform distribution. For the Humanoid task in our experiment, we use the Wasserstein-distance variant of InfoGAIL, since the Jensen-Shannon-divergence variant of InfoGAIL does not perform well in this task.

Next, we review offline IL methods. These methods learn a policy based on supervised learning and do not require additional transition samples from MDPs.

BC. Behavior cloning (BC) (Pomerleau, 1988) is perhaps the simplest IL method. BC treats an IL problem as a standard supervised learning problem and ignores dependency between states distributions and policy. For continuous action space, BC solves a least-square regression problem to learn a parameter θ of a deterministic policy $\pi_{\theta}(\mathbf{s}_t)$:

$$\min_{\theta} \mathbb{E}_{p^*(\mathbf{s}_{1:T}, \mathbf{a}_{1:T})} \left[\sum_{t=1}^T \|\mathbf{a}_t - \pi_{\theta}(\mathbf{s}_t)\|_2^2 \right].$$

BC-D. BC with Diverse-quality demonstrations (BC-D) is a simple extension of BC for handling diverse-quality demonstrations. This method is based on the naive model in Section 3.1, and we consider it mainly for evaluation purpose. BC-D uses supervised learning to learn a policy parameter θ and expertise parameter ω of a model $p_{\theta, \omega}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, k) = p(k)p(\mathbf{s}_1) \prod_{t=1}^T p(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{u}_t) \int_{\mathcal{A}} \pi_{\theta}(\mathbf{a}_t|\mathbf{s}_t) p_{\omega}(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k) d\mathbf{a}_t$. To learn the parameters, we minimize the KL divergence from data distribution to the model. By using the variational approach to handle integration over the action space, BC-D solves an optimization problem

$$\max_{\theta, \omega, \nu} \mathbb{E}_{p_d(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}|k)p(k)} \left[\sum_{t=1}^T \mathbb{E}_{q_{\nu}(\mathbf{a}_t|\mathbf{s}_t, \mathbf{u}_t, k)} \left[\log \frac{\pi_{\theta}(\mathbf{a}_t|\mathbf{s}_t) p_{\omega}(\mathbf{u}_t|\mathbf{s}_t, \mathbf{a}_t, k)}{q_{\nu}(\mathbf{a}_t|\mathbf{s}_t, \mathbf{u}_t, k)} \right] \right],$$

where $q_{\nu}(\mathbf{a}_t|\mathbf{s}_t, \mathbf{u}_t, k)$ is a variational distribution with parameters ν . We note that the model $p_{\theta, \omega}(\mathbf{s}_{1:T}, \mathbf{u}_{1:T}, k)$ of BC-D can be regarded as a regression-extension of the two-coin model proposed by Raykar et al. (2010) for classification with noisy labels.

Co-teaching. Co-teaching (Han et al., 2018) is the state-of-the-art method to perform classification with noisy labels. This method trains two neural networks such that mini-batch samples are exchanged under a small loss criteria. We extend this method to learn a policy by least-square regression. Specifically, let $\pi_{\theta_1}(\mathbf{s}_t)$ and $\pi_{\theta_2}(\mathbf{s}_t)$ be two neural networks representing policies, and $\nabla_{\theta} L(\theta, \mathcal{B}) = \nabla_{\theta} \sum_{(\mathbf{s}, \mathbf{a}) \in \mathcal{B}} \|\mathbf{a} - \pi_{\theta}(\mathbf{s})\|_2^2$ be gradients of a least-square loss estimated by using a mini-batch \mathcal{B} . The parameters θ_1 and θ_2 are updated by iterates:

$$\theta_1 \leftarrow \theta_1 - \eta \nabla_{\theta_1} L(\theta_1, \mathcal{B}_{\theta_2}), \quad \theta_2 \leftarrow \theta_2 - \eta \nabla_{\theta_2} L(\theta_2, \mathcal{B}_{\theta_1}).$$

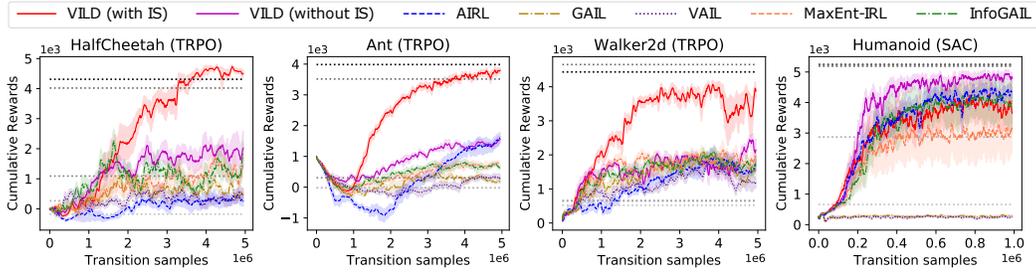
The mini-batch \mathcal{B}_{θ_2} for updating θ_1 is obtained such that \mathcal{B}_{θ_2} incurs small loss when using prediction from π_{θ_2} , i.e., $\mathcal{B}_{\theta_2} = \operatorname{argmin}_{\mathcal{B}'} L(\theta_2, \mathcal{B}')$. Similarly, the mini-batch \mathcal{B}_{θ_1} for updating θ_2 is obtained such that \mathcal{B}_{θ_1} incurs small loss when using prediction from π_{θ_1} . For evaluating the performance, we use the policy network π_{θ_1} .

D ADDITIONAL EXPERIMENTAL RESULTS

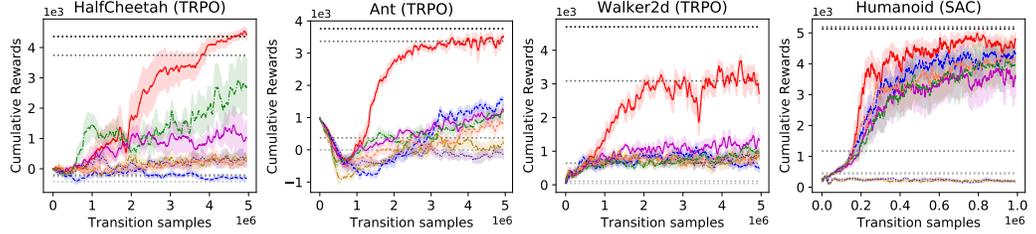
Results against online IL methods. Figure 4 shows the learning curves of VILD and existing online IL methods against the number of transition samples. It can be seen that for both types of noisy policy, VILD with and without IS outperform existing methods overall, except on the Humanoid tasks where most methods achieve comparable performance.

Results against offline IL methods. Figure 5 shows learning curves of offline IL methods, namely BC, BC-D, and Co-teaching. For comparison, the figure also shows the final performance of VILD with and without IS, according to Table 1. We can see that these offline methods do not perform well, especially on the high-dimensional Humanoid task. The poor performance of these methods is due to the issues of compounding error and low-quality demonstrations. Specifically, BC performs the worst, since it suffers from both issues. Still, BC may learn well in the early stage of learning, but its performance sharply degrades, as seen in Ant and Walker2d. This phenomena can be explained as an empirical effect of memorization in deep neural networks (Arpit et al., 2017). Namely, deep neural networks learn to remember samples with simple patterns first (i.e., high-quality demonstrations from experts), but as learning progresses the networks overfit to samples with difficult patterns (i.e., low-quality demonstrations from amateurs). Co-teaching is the-state-of-the-art method to avoid this effect, and we can see that it performs significantly better than BC. Meanwhile, BC-D, which learns the policy and level of demonstrators’ expertise, also performs better than BC and is comparable to Co-teaching. Nonetheless, the performance of Co-teaching and BC-D is still much worse than VILD with IS.

Accuracy of estimated expertise parameter. Figure 6 shows the estimated parameters $\omega = \{\mathbf{c}_k\}_{k=1}^K$ of $\mathcal{N}(\mathbf{u}_t | \mathbf{a}_t, \operatorname{diag}(\mathbf{c}_k))$ and the ground-truth variance $\{\sigma_k^2\}_{k=1}^K$ of the Gaussian noisy policy $\mathcal{N}(\mathbf{u}_t | \mathbf{a}_t, \sigma_k^2 \mathbf{I})$. The results show that VILD learns an accurate ranking of the variance compared to the ground-truth. The values of these parameters are also quite accurate compared to the ground truth, except for demonstrators with low-levels of expertise. A possible reason for this phenomena is that low-quality demonstrations are highly dissimilar, which makes learning the expertise more challenging. We can also see that the difference between expertise parameters of VILD with IS and VILD without IS is small and negligible.

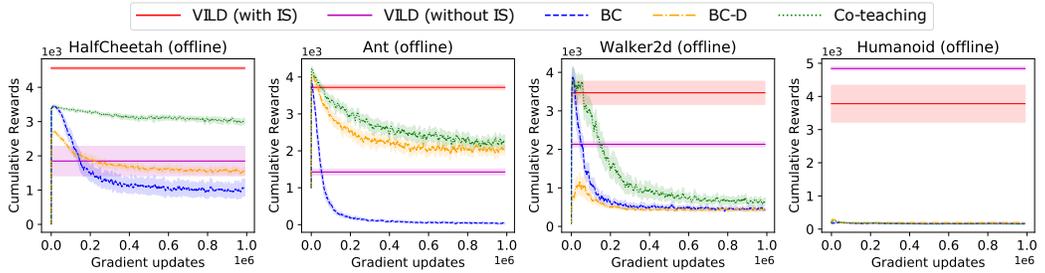


(a) Performance of online IL methods when demonstrations are generated using Gaussian noisy policy.

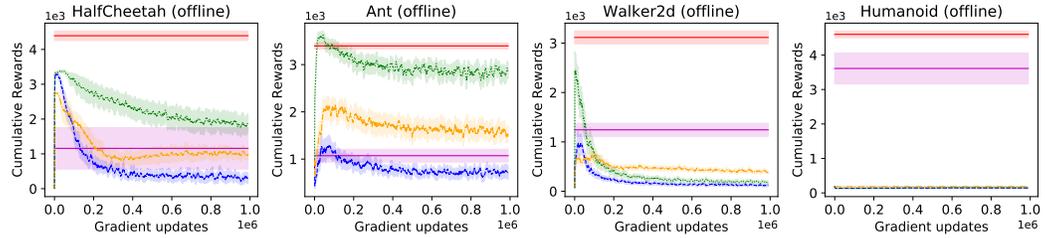


(b) Performance of online IL methods when demonstrations are generated using TSD noisy policy.

Figure 4: Performance averaged over 5 trials of online IL methods against the number of transition samples. Horizontal dotted lines indicate performance of $k = 1, 3, 5, 7, 10$ demonstrators.



(a) Performance of offline IL methods when demonstrations are generated using Gaussian noisy policy.



(b) Performance of offline IL methods when demonstrations are generated using TSD noisy policy.

Figure 5: Performance averaged over 5 trials of offline IL methods against the number of gradient update steps. For VILD with and without IS, we report the final performance in Table 1.

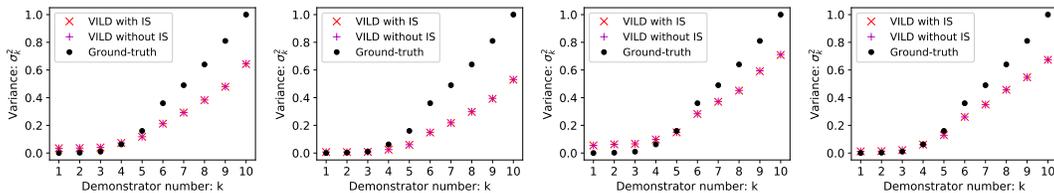


Figure 6: Expertise parameters $\omega = \{c_k\}_{k=1}^K$ learned by VILD and the ground-truth $\{\sigma_k^2\}_{k=1}^K$ for the Gaussian noisy policy. For VILD, we report the value of $\|c_k\|_1/d_a$.