
Off-Policy Adversarial Inverse Reinforcement Learning

Samin Yeasar Arnob¹

Abstract

Adversarial Imitation Learning (AIL) is a class of algorithms in Reinforcement learning (RL), which tries to imitate an expert without taking any reward from the environment and does not provide expert behavior directly to the policy training. Rather, an agent learns a policy distribution that minimizes the difference from expert behavior in adversarial setting. Adversarial Inverse Reinforcement Learning (AIRL) leverages the idea of AIL, integrates a reward function approximation along with learning the policy and shows the utility of IRL in the transfer learning setting. But the reward function approximator that enables transfer learning does not perform well in imitation tasks. We propose an Off-Policy Adversarial Inverse Reinforcement Learning (Off-policy-AIRL) algorithm which is sample efficient as well as gives good imitation performance compared to the state-of-the-art AIL algorithm in the continuous control tasks. For the same reward function approximator, we show the utility of learning our algorithm over AIL by using the learned reward function to retrain the policy over a task under significant variation where expert demonstrations are absent.

1. Introduction

Reinforcement learning (RL) is a very useful framework for learning complex behavior in control, where an agent interacts with an environment and tries to learn the expected behavior by optimizing an objective function that ensures the highest cumulative reward over specific time steps. A reward is often human-engineered function that tries to ensure higher rewards for expected behavior within the environment. Designing a reward function for a complex environment can be tricky. Poorly designed reward function may

lead to a fatal mistake when applied in real-world applications.

Generative Adversarial Imitation Learning (GAIL) (Ho & Ermon, 2016) introduces AIL where a generator (policy) is used to compute trajectory and a binary classifier called discriminator classifies generated behavior from the expert. As imitation learning algorithms only learn policy, they miserably fail when there is a change in environmental dynamics.

Inverse reinforcement learning (IRL) (Ng et al., 1999a; Abbeel & Ng, 2004) addresses the importance of learning reward function and thus learn reward function approximator along with policy. Adversarial Inverse Reinforcement Learning (AIRL) (Fu et al., 2018) proposes an IRL algorithm in adversarial learning, which shows a promising result when there is considerable variability in the environment from the demonstration setting. But compared to AIL algorithms, AIRL does not perform well in imitation performance (Kostrikov et al., 2019; Fu et al., 2018). It is also important to note the best performing reward function approximator configuration for imitation and transfer learning tasks in AIRL algorithm are different. In our work we perform transfer learning using the same configuration that works best in imitation performance.

The paper makes the following contributions: (i) We propose an Off-Policy Adversarial Inverse Reinforcement Learning algorithm (off-policy-AIRL) which is sample efficient and improves imitation performance of prior IRL algorithms through comparing imitation performance with state-of-the-art AIL algorithm on continuous control tasks. (ii) We show utility of learning IRL algorithm over imitation by reusing the learned reward function to train policy under certain changes in dynamics where expert demonstrations are absent. We use *Soft Actor-Critic* (Haarnoja et al., 2018) (SAC), which makes our algorithm more sample efficient. (iii) We also show using *Multiplicative Compositional Policies* (Peng et al., 2019) (MCP) allows more flexibility in retraining policy in transfer task and at the same time improves imitation performance.

^{*}Equal contribution ¹Department of Electrical and Computer Engineering, McGill University, Montreal, Canada. Correspondence to: samin.arnob <@mail.mcgill.ca>.

2. Background

We define a Markov decision process (MDP) as a tuple $(S, A, T, r, \gamma, \rho_0)$, where S, A are the state and action space respectively, $\gamma \in [0, 1)$ is the discount factor. We consider the dynamic or transition distribution $T(s^j|s, a)$, initial state distribution $\rho_0(s)$ and reward function $r(s, a)$ are unknown in IRL setup and can be learned through interaction with the MDP.

AIL evolved from maximum casual entropy IRL framework (Ziebart et al., 2010), which considers an entropy regularized MDP with the goal to find the optimal policy π that maximizes the expected entropy-regularized discounted reward under π, T and ρ_0 :

$$\pi = \arg \max_{\pi} \mathbf{E}_{\tau} \left[\sum_{t=0}^T \gamma^t r(s_t, a_t) + H(\pi(\cdot|s_t)) \right]. \quad (1)$$

Here $\tau = (s_0, a_0 \dots s_T, a_T)$ defines the sequence of states and actions induced by policy and $H(\pi(\cdot|s_t))$ is the discounted causal entropy of policy π . In imitation task, expert demonstrations $D = \tau_1, \tau_2 \dots \tau_n$ are given and an agent is expected to produce trajectory which are similar to the expert and infer reward function. Under certain constrain of feature matching (Ziebart et al., 2008) derived following objective function:

$$P_{\psi}(\tau) / z e^{R(\tau)} = \frac{1}{z} e^{R(\tau)}, \quad (2)$$

$$/ p(s_0) \prod p(s_{t+1}|s_t, a_t) e^{\gamma^t r_{\psi}(s_t, a_t)}, \quad (3)$$

where $z = \int_{\tau} e^{R(\tau)}$ is the partition function and can be expressed as probability of trajectories with equal reward/cost are equal likely and probability of trajectory with higher rewards are exponentially more likely. Equation (2) can be interpreted as solving a maximum likelihood problem:

$$\max_{\psi} \mathbf{E}_{\tau} [D[\log p_{\psi}(\tau)]]. \quad (4)$$

But the assumptions that are made in (Ziebart et al., 2010) are model / transition function $T(s^j|s, a)$ is known and applied on small state-action space.

GAIL (Ho & Ermon, 2016), a model-free algorithm, is the first framework to draw connection between generative adversarial network (GAN) and imitation learning which works in continuous state-action space. GAIL uses a generative model or generator, G that acts as the policy. Purpose of the generator is to learn policy π_{θ} such that the state-action visitation frequency $\rho_{\pi_{\theta}}$ is similar to expert's ρ_{π_E} without directly being provided with the expert demonstrations. A discriminator D is a binary classifier which tries to

distinguish the data distribution of the generator ρ_{π} from the expert's ρ_{π_E} . Objective for a generator is to minimize the difference with expert visitation frequency, while a discriminator wants to maximize the difference and thus derived the entropy regularized objective as following:

$$\begin{aligned} & \text{distance}_{\min_{\pi_{\theta}} \max_{D_{\psi}} (\rho_{\pi}, \rho_{\pi_E})} = \\ & \min_{\pi_{\theta}} \max_{D_{\psi}} \left[\mathbf{E}_{\pi_{\theta}} [\log D_{\psi}(s, a)] + \right. \\ & \left. \mathbf{E}_{\pi_E} [\log(1 - D_{\psi}(s, a))] - \lambda H(\pi_{\theta}) \right], \quad (5) \end{aligned}$$

where entropy term H is policy regularizer controlled by $\lambda \geq 0$. GAIL uses two separate neural-networks to represent generator and discriminator. For any given state generator tries to take expert like action. Discriminator takes state-action pairs as input and computes the likelihood of the input coming from an expert. Generator uses a reward function $R(s, a) = -\log D_{\psi}(s, a)$ in order to train it's network parameters.

GAIL is sample efficient in terms of number of expert demonstrations required but not quite sample efficient in terms of required interaction with the environment to learn the policy. Thus GAIL is not suitable for many real-world applications. *Discriminator Actor critic* (DAC) (Kostrikov et al., 2019) uses Twin-delayed deep deterministic actor-critic (TD3) (Fujimoto et al., 2018), an off-policy algorithm as generator to improve upon the sample efficiency of existing methods, and it extends the learning environment with absorbing states. DAC criticises AIL algorithms (i.e GAIL) for inducing reward biases by using strict positive or negative reward function and proposes a reward function $r(s, a) = \log(D_{\theta}(s, a)) - \log(1 - D_{\theta}(s, a))$ in order to achieve unbiased rewards. Combined, these changes remove the survival bias and solves sample inefficiency problem of GAIL.

For being an imitation learning algorithm, instead of learning cost/reward function, both GAIL and DAC only recover expert policy and thus will fail miserably in dynamic environments. A discriminator function, that learns policy as well as the cost/reward function, is proposed by (Finn et al., 2016) by optimizing following objective for discriminator network:

$$D_{\psi} = \frac{p_{\psi}(\tau)}{p_{\psi}(\tau) + q(\tau)} \stackrel{(a)}{=} \frac{1/z e^{R_{\psi}}}{1/z e^{R_{\psi}} + q(\tau)}, \quad (6)$$

$$\begin{aligned} \text{thus } L_{dis} = & \max_{D_{\psi}} \left[\mathbf{E}_{\tau} [p[\log D_{\psi}(\tau)]] + \right. \\ & \left. \mathbf{E}_{\tau} [q[\log(1 - D_{\psi}(\tau))]] \right], \quad (7) \end{aligned}$$

where $p(\tau)$ is the data distribution and $q(\tau)$ is the policy distribution and (a) consider derived objective function (2) from (Ziebart et al., 2010). But (Finn et al., 2016) contains

no experimental demonstration of imitation performance. Later on AIRL (Fu et al., 2018) shows, when we consider a whole trajectory to compute discriminator update it suffers from high variance issue and that leads to a poor performing generator. Instead AIRL computes discriminator update using (s, a, s') tuple:

$$D_{\psi, \omega} = \frac{p_{\psi, \omega}(s, a, s^{\theta})}{p_{\psi, \omega}(s, a, s^{\theta}) + q(s, a)}, \quad (8)$$

$$= \frac{e^{f_{\psi, \omega}(s, a, s^{\theta})}}{e^{f_{\psi, \omega}(s, a, s^{\theta})} + q(s, a)}, \quad (9)$$

$$= \frac{e^{f_{\psi, \omega}(s, a, s^{\theta})}}{e^{f_{\psi, \omega}(s, a, s^{\theta})} + \pi_{\theta}(a|s)}, \quad (10)$$

where $f_{\psi, \omega}(s, a, s^{\theta}) = r_{\psi}(s, a) + \gamma V_{\omega}(s^{\theta}) - V_{\omega}(s)$ is a potential-based reward function, $r_{\psi}(s, a)$ is a reward function approximator and V_{ω} is reward shaping term controlled by the dynamics. In (Fu et al., 2018) $f_{\psi, \omega}(s, a, s^{\theta})$ is referred as disentangled reward function (Ng et al., 1999b) due to $r_{\psi}(s)$ being indifferent to the changes in dynamics T .

3. Off-policy Adversarial Inverse Reinforcement Learning

Off-policy-AIRL algorithm is inspired from DAC (Kostrikov et al., 2019), which is an adversarial imitation learning algorithm. DAC addresses two existing problem with AIL algorithms, which are reward function bias and sample inefficiency. Reward functions used in these algorithms induce an implicit bias by using strict positive or negative reward functions, which may work for some environments but in many cases become a reason for sub-optimal behavior.

Furthermore, (Kostrikov et al., 2019) shows adversarial methods improperly handle terminal states. This introduces implicit reward priors that can affect the policy’s performance. In particular, many imitation learning implementations (Kim & Park, 2018; Ho & Ermon, 2016; Fu et al., 2018) and MDPs omit absorbing states s_a . Thus they implicitly assign 0 reward to terminal/absorbing states and adds a bias to the reward learning process. Thus we use the same approach to learn the reward for absorbing states. Return of the terminal state is defined as $R_T = r(s_T, a_T) + \sum_{t=T+1}^{\infty} \gamma^t r(s_a, \cdot)$, where $r(s_a, 0)$ is learned reward, instead of just $R_T = r(s_T, a_T)$. To make the implementation more stable, the terminal reward is analytically derived using the following equation:

$$R_T = r(s_T, a_T) + \gamma \frac{r(s_a, \cdot)}{1 - \gamma} \quad (11)$$

Like any other imitation learning algorithm DAC only learns the policy. We want to use the best of both worlds and thus

replace the binary discriminator of DAC and use formulated discriminator function from AIRL (Fu et al., 2018) to learn a reward function approximator along with the policy. From (Fu et al., 2018) it is evident that AIRL provides a poor imitation performance compared to GAIL. It is important to note that the best performing policy in (Fu et al., 2018) for imitation task is trained using state-action dependent reward function $r_{\psi}(s, a)$, which fails to re-train policy in the transfer learning task. For state dependent reward function $r_{\psi}(s)$ (Fu et al., 2018) successfully completes transfer learning task but it comes with a cost of poor imitation performance. In this work, we perform transfer learning using the same configuration that works best in imitation performance. We improve upon the prior IRL algorithm performance and while still being sample efficient by using Soft-Actor-Critic (SAC) (Haarnoja et al., 2018) as the generator.

We store our experience in a replay buffer \mathcal{R} to utilize them during the off-policy update of the generator. After each time the environment reaches a terminal condition and resets, we update our discriminator and generator for the same episodic timesteps. We use an expert buffer \mathcal{R}_E to store the expert trajectory and sample the (s_E, a_E, s_E^{θ}) pair in batch while update our discriminator.

3.1. Discriminator update rule

During discriminator update we randomly sample mini-batch of state-action-next-state pairs $(s, a, \cdot, s^{\theta})_b$ from both expert buffer \mathcal{R}_E and replay buffer \mathcal{R} . We use following equation derived in (Fu et al., 2018) to compute the discriminator output,

$$D_{\psi, \omega} = \frac{e^{f_{\psi, \omega}(s, a, s^{\theta})}}{e^{f_{\psi, \omega}(s, a, s^{\theta})} + \pi_{\theta}(a|s)}. \quad (12)$$

where

$$f_{\psi, \omega} = r_{\psi}(s) + \gamma V_{\omega}(s^{\theta}) - V_{\omega}(s), \quad (13)$$

$$= r_{\psi}(s) + \gamma V_{\omega}(s^{\theta}) - V_{\omega}(s). \quad (14)$$

Here reward approximator r_{ψ} can be a function of state, state-action or state-action-next-state and V_{ω} can be any function that gives a measure of being at any state s . Similar to (Fu et al., 2018) we use value function V_{ω} as reward shaping terms. We find the best imitation performance using state dependent reward function.

Output of the discriminator will predict likelihood of being an expert and thus objective of the discriminator is to minimize following binary cross entropy loss:

$$\min L_{\psi, \omega} = \min \sum_{b=1}^B [\log D_{\psi, \omega}(s_b, a_b, s_b^{\theta}) + \log(1 - D_{\psi, \omega}(s_{b_E}, a_{b_E}, s_{b_E}^{\theta}))]. \quad (15)$$

With in few iterations of the discriminator update, it easily can classify the expert from the generator data. Thus to make the generator learning more stable we use gradient penalty (Gulrajani et al., 2017; Kostrikov et al., 2019).

3.2. Generator update rule

In IRL setting, we consider the actual reward from the environment is non-existent and rather formulate a reward approximator to train policy. AIRL (Fu et al., 2018) algorithm uses a on-policy policy-gradient algorithm as its generator and proposes following reward approximator $\hat{r}(s, a, s^\theta)$ to update the policy:

$$\hat{r}_{\psi,\omega}(s, a, s^\theta) = \log D_{\psi,\omega}(s, a, s^\theta) - \log(1 - D_{\psi,\omega}(s, a, s^\theta)). \quad (16)$$

Using equation (12) it is also referred as entropy regularized reward function:

$$\begin{aligned} \hat{r}_{\psi,\omega}(s, a, s^\theta) &= \log \frac{e^{f_{\psi,\omega}(s,a,s^\theta)}}{e^{f_{\psi,\omega}(s,a,s^\theta)} + \pi(a|s)} \\ &\quad \log \frac{\pi_\theta(a|s)}{e^{f_{\psi,\omega}(s,a,s^\theta)} + \pi_\theta(a|s)}, \\ &= f_{\psi,\omega}(s, a, s^\theta) - \log \pi_\theta(a|s). \end{aligned} \quad (17)$$

For improving imitation performance and sample efficiency we use SAC as our generator. SAC being an off-policy algorithm, samples $(s, a, s^\theta, done)$ tuple from buffer when trains the network. For our algorithm we experiment different reward approximator which is discussed in following section and use trained reward approximator to update critic or Q_ϕ function in SAC with entropy term through following gradient step:

$$\begin{aligned} \nabla_\phi J_Q(\phi) &= \nabla_\theta Q_\phi(s_t, a_t) \left[Q_\phi(s_t, a_t) - \hat{r}_{\psi,\omega}(s, a, s^\theta) + \right. \\ &\quad \left. \gamma Q_{\phi^\theta}(s_{t+1}, a_{t+1}) - \alpha \log \pi_\theta(a_{t+1}|s_{t+1}) \right]. \end{aligned} \quad (18)$$

Similar to (Fujimoto et al., 2018; Haarnoja et al., 2018) we update actor or policy π_θ for every second update of our Q_ϕ function using following gradient step:

$$\nabla_\theta [E_{a_t \sim \pi_\theta} [\alpha \log \pi_\theta(a_t|s_t) - Q_\phi(s_t, a_t)]]. \quad (19)$$

3.3. Reward function selection

Implementation of AIRL discriminator (Fu et al., 2018) does not work off the shelf as now we have to train an off-policy generator. We experiment on different way to compute the reward function for the generator update.

In AIRL, the author uses the value function of the states $V_\omega(s)$ to measure additional shaping terms. We experiment

on different variations of reward approximator $\hat{r}_{\psi,\omega}$ to update policy and disentangled reward function $f_{\psi,\omega}$ to compute discriminator output so that we find the combination of these two approximators that gives the best performance in off-policy-AIRL.

Implementation 1:

Policy update using :

$$\begin{aligned} \hat{r}_{\psi,\omega}(s, a, s^\theta) &= \log D_{\psi,\omega}(s, a, s^\theta) \\ &\quad \log(1 - D_{\psi,\omega}(s, a, s^\theta)) \end{aligned}$$

Disentangled reward function :

$$f_{\psi,\omega} = r_\psi(s, a) + \gamma V_\omega(s^\theta) - V_\omega(s)$$

Implementation 2

Policy update using :

$$\begin{aligned} \hat{r}_{\psi,\omega}(s, a, s^\theta) &= \log D_{\psi,\omega}(s, a, s^\theta) \\ &\quad \log(1 - D_{\psi,\omega}(s, a, s^\theta)) \end{aligned}$$

Disentangled reward function :

$$f_{\psi,\omega} = r_\psi(s) + \gamma V_\omega(s^\theta) - V_\omega(s)$$

Implementation 3

Policy update using :

$$\hat{r}_{\psi,\omega}(s, a, s^\theta) = r_\psi(s, a)$$

Disentangled reward function :

$$f_{\psi,\omega} = r_\psi(s, a) + \gamma V_\omega(s^\theta) - V_\omega(s)$$

Implementation 4

Policy update using :

$$\hat{r}_{\psi,\omega}(s, a, s^\theta) = r_\psi(s)$$

Disentangled reward function :

$$f_{\psi,\omega} = r_\psi(s) + \gamma V_\omega(s^\theta) - V_\omega(s)$$

For implementation (1) and (2), we use the exact policy update rule which is used in AIRL (Fu et al., 2018). In implementation (1) we consider a state dependent reward function $r_\psi(s)$ and for implementation (2) we consider reward to be a function of state-action pair $r_\psi(s, a)$. As we see from equation (16), reward computed to update policy is equivalent to entropy regularized reward function. We evaluate the disentangled reward function for both state dependent $r_\psi(s)$ and state-action dependent $r_\psi(s, a)$. For implementation (3) and (4) we follow similar comparative study where policy update discard entropy regularization and use direct output from our reward approximator r_ψ . As discussed in DAC (Kostrikov et al., 2019), strict negative or positive reward approximator induces bias. Similar to DAC, our reward approximator r_ψ gives both positive and negative rewards thus does not suffer reward biasing.

4. Transfer learning task

Advantage of IRL over imitation learning is that we can leverage the learned reward function to retrain a new policy. If the reward function captures the underlying objective of an agent, then it is possible to use the reward function under robust changes of the dynamics. Imitation learning successfully learn policy in the training domain but it is not possible to use the policy when there is a significant domain shift (Fu et al., 2018). We use the learned reward function from off-policy-AIRL algorithm to re-train our policy in transfer learning setting. To the best of our knowledge AIRL (Fu et al., 2018) is the only IRL framework that shows utility of reward function in transfer learning for continuous control tasks and thus we demonstrate our experiments using same environments from (Fu et al., 2018) (see Figure 1) and we also consider expert demonstrations to be absent for the transfer tasks.

For transfer learning experiments we consider dynamic changes in following two criterion, where either (1) dynamics of the agent or (2) dynamics of the environment changes while the goal/objective of the agent and action dimension remain the same. Through these experiments we show the learned reward function r_{ψ} indeed can help the policy to behave under dynamic changes even when expert behavior is completely absent. For transfer learning experiments we use the same hyper-parameter and update rules described in prior.

4.1. Criterion:1

By changing dynamics of an agent, we refer to a significant structural change of an agent. To evaluate transfer learning performance under this criterion we use *Customized Ant* (Fu et al., 2018) environment. In both imitation and transfer task the quadrupedal Ant requires to run forward, while dynamics of the agent during test session is changed by disabling and shrinking two legs. In MuJoCo setting, (Fu et al., 2018) sets the gear ratios to 1, whereas the other joints were set to 150. The gear ratio is a parameter which translates actions to torques on the joints, so the disabled joints are 150x more difficult to move. This significantly changes its gait. In a MDP this puts a restriction in action space A , which results in a new transition dynamics $T^{\theta}(s|s, a)$. Thus optimal policy π is changed even though the objective and action-dimension remains the same.

4.2. Criterion:2

By changing dynamics of the environment, we consider notable variation in the trained environment such that it directly results in a different transition dynamics $T^{\theta}(s|s, a)$ without putting restriction in action-space A . Objective of *Shifting Maze* (Fu et al., 2018) task is to reach to a specific

goal position, where during imitation learning, agent is trained to go around the wall on the left side, but during transfer task environment setup is changed and must go around the wall on the right. Thus optimal policy π is changed even though the objective and the dynamics of the agent remains same.

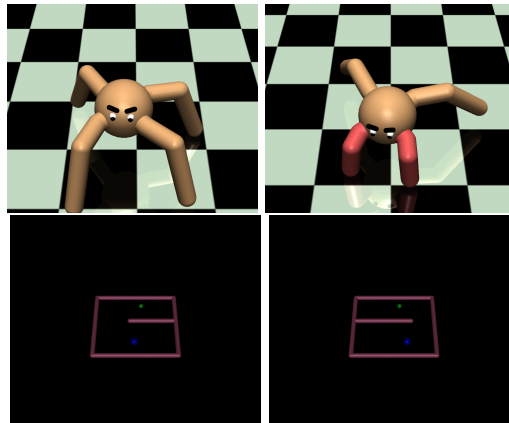


Figure 1. Top row: Ant environment is used to perform experiment under criterion 1 Bottom row: For shifting maze task, the agent (blue) is required to reach at the goal (green) and is being used for experiment under criterion 2

5. Multiplicative Compositional Policies

Policy often gets overfitted over one task and thus become harder to re-train during transfer learning task. Different techniques such as early stopping or regularization can be helpful in this case. An alternative solution is provided by (Peng et al., 2018) through using *variational bottleneck* technique to control over the information that is fed to the discriminator using mutual information theory, which allows more controlled update in the generator. But this technique is highly dependent on hyper-parameter of the bottleneck itself, thus performance can be varied vastly if not tuned properly for each task. We want to learn a policy that gives a good performance during transfer learning without being dependent on exact hyper-parameter tuning. We propose using *Multiplicative Compositional Policy* (MCP) (Peng et al., 2019) to improve policy performance in transfer learning task. MCP was introduced to combine multiple primitive skills in RL setting and these primitive skills are controlled using a gating function. We propose using multiple primitive networks to learn single skill in IRL setting. Using a weighted multiplicative composition of Gaussian primitives improves imitation performance and allows more flexibility in composing into new policy in transfer learning.

MCP (Peng et al., 2019) tries to learn different skills with different primitive policies and then reuse those learned skills by combining them to do a more sophisticated task.

In hierarchical learning (Faloutsos et al., 2001; Sutton et al., 1999; Jordan & Jacobs, 1992; Peng et al., 2016) it is common to learn primitives and then find a composite policy $\pi(a|s)$ by weighted sum of distribution from primitives $\pi_i(a|s)$. A gating function computes the weight w_i , that determines the probability of activating each primitive for a given state, s . Composite policy can be written as:

$$\pi(a|s) = \sum_{i=1}^k w_i(s) \pi_i(a|s), \sum_{i=1}^k w_i(s) = 1, w_i(s) > 0. \quad (20)$$

Here k denotes the number of primitives and this can be referred as *additive model*. Standard hierarchical learning models can sequentially select particular skill over time. In other words, can activate single primitive at each timestep. *MCP* (Peng et al., 2019) proposes an *multiplicative model*, where multiple primitives are activated at a given time-step. This allows an agent to learn a complex task requiring to perform more than one sub-task at a time.

MCP (Peng et al., 2019) decomposes agent’s policy into primitives, train those primitive policies on different sub-tasks and then obtain a composite-policy by multiplicative composition of these primitives. Here each primitive is considered as distribution over actions and the composite policy is a multiplicative composition of these distribution. Thus,

$$\pi(a|s, g) = \frac{1}{Z(s, g)} \prod_{i=1}^k \pi_i(a|s, g)^{w_i(s, g)}, w_i(s, g) > 0. \quad (21)$$

MCP enables an agent to activate multiple primitives simultaneously with each primitive specializing in different behaviors that can be composed to produce a continuous spectrum of skills, whereas, standard hierarchical algorithm only activate single primitive at a time.

6. MCP in off-policy-AIRL

MCP learns primitive policies π_i by training them into different sub-tasks and then try to learn more complex task leveraging the learned skills. But in transfer learning experiments, we use MCP to learn single task. The underlying idea is to observe whether it can relearn the composition of different action parameters when we put them in dynamic scenarios. For example, during transfer learning experiments (Fu et al., 2018) disables two legs and thus it is important for the agent to re-optimize its policy to walk in this transfer learning setup. Our initial objective is to see if MCP can decompose motor skills for a single task and re-optimize the learned policy by re-training the gating function. We refer this new policy as SAC-MCP.

In MCP (Peng et al., 2019), experiments were conducted

using on-policy algorithms in RL setting and the policy parameters were considered to be state and goal dependent. We explore this concept to decompose motor skill over single task in IRL setting using off-policy algorithm.




Figure 2. Generator architecture for SAC-MCP.

7. Experimental Setup

Similar to (Ho & Ermon, 2016; Kostrikov et al., 2019) discriminator is 2 layer MLP of 100 hidden units with tanh activation. Our generator consists of separate Actor and Critic neural network and follows the architecture used in (Kostrikov et al., 2019; Fujimoto et al., 2018), where both of these networks have 2 layer MLP of 400 and 300 hidden units with ReLU activation. To implement SAC-MCP, we modify Actor network (see Figure 2) by adding multiple primitive networks with same configuration described in (Peng et al., 2019).

We have trained all networks with the Adam optimizer (Kingma & Ba, 2014) from *PyTorch*, which uses default learning rate of $1e^{-3}$. For our experiment we trained our algorithms on *MuJoCo* (Todorov et al., 2012) continuous control task environments. For transfer learning experiments we use Custom-Ant and Shifting Maze environments from (Fu et al., 2018). Performance curve is obtained using the mean over 10 experiments for 0-9 seeds and evaluated after each 5000 interaction with the environment. During each evaluation we have stored the average performance of 10 runs.

8. Results

8.1. Reward function selection

We conduct experiment on *CustomAnt-v0* (Fu et al., 2018) to select reward function for policy update. We train SAC policy and collected 50 trajectory from the expert. We keep our seed fixed to 0. Figure 3(a) gives the performance curve using actual reward from the environment and it is

evident that reward signal directly from the reward function approximator (Implementation 4) gives better policy performance in imitation task. We also see the cumulative reward achieved from reward approximator r_θ (see Figure 3(b)) replicates the actual performance curve. As our experiment shows drastic difference in performance we have not done comparative study for multiple seeds or other environments.




Figure 3. Performance comparison of Off-policy AIRL for different update rules. Here figure (a) gives the performance curve using actual reward from the environment (b) shows the cumulative reward achieved from reward approximator r_θ

As demonstrated in Figure 4 using implementation-4 off-policy AIRL gives a better imitation performance than DAC in Hopper-v2 and Walker2d-v2 environments within $1e^6$ iterations, while provides comparable performance for HalfCheetah-v2 and Ant-v2 environments after $2e^6$ iterations.

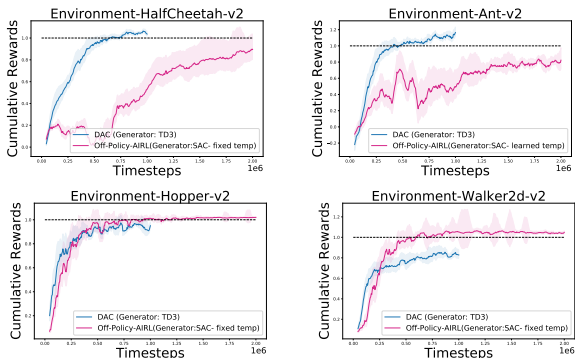


Figure 4. Performance Comparison of DAC and Off-policy AIRL over continuous control task

8.2. Performance in transfer learning

We test trained agent in transfer learning task under criterion (1) and (2). We demonstrate the utility of learning reward function which gives us an advantage of learning IRL over imitation learning. In IRL setting, Off-policy AIRL learns expert like behavior in Custom Ant environment but suffers from variance in Shifting maze task. We do not tune separate hyper-parameter for individual environment but doing so may reduce performance variance for Shifting maze environment.

As illustrated in Figure 1 during transfer learning under criterion-1, agent’s gait changes significantly and as can be




Figure 5. Top row: Performance Off-policy AIRL over continuous control task Bottom row: Performance Off-policy AIRL during transfer learning task.

seen in Figure 4 off-policy-AIRL fails to re-learn the policy using reward function. But we do not conclude that it is not possible to relearn policy under this criterion using off-policy-AIRL. It is demonstrated in (Ha, 2018) that changing a better structure of agent’s body not only is better suited for the task but also facilitates policy learning. The environment that we use in criterion-1 was introduced in (Fu et al., 2018), where the structure of the body of the quadrupedal Ant was changed (Fu et al., 2018) for transfer learning by shrinking two legs and making other two larger. But the specific increment or decrement of the legs may have favored for (Fu et al., 2018) to re-learn policy when two legs are paralyzed in transfer learning. A hyper-parameter tuning of the structural change may work for our proposed algorithm as well. Experiments with more structural flexibility are required to come to a solid conclusion under this criterion.

On the other hand for Shifting maze environment (see Figure 4) it successfully completes the task but suffer from high variance in performance.

8.3. Performance comparison for using MCP

Performance of our MCP implementation in IRL setting is shown in Figure 6. Using SAC-MCP improves the imitation performance in multiple (HalfCheetah-v2 and Ant-v2) MuJoCo control tasks.

We again conduct our transfer learning experiments on Shifting-Maze tasks. Here we do not load prior gating function. We compare performance with and without re-training the actor. Our initial hypothesis is that having multiplicative composition of Gaussian primitives should allow us with diverse behavior by only training the gating function during transfer task. But Figure 7 shows we are not able to learn shifting maze task without retraining the actor networks. We compare performance of SAC-MCP for $k = 74, 8g$ primitives (see Figure 7). For $K = 8$ it reduces the performance

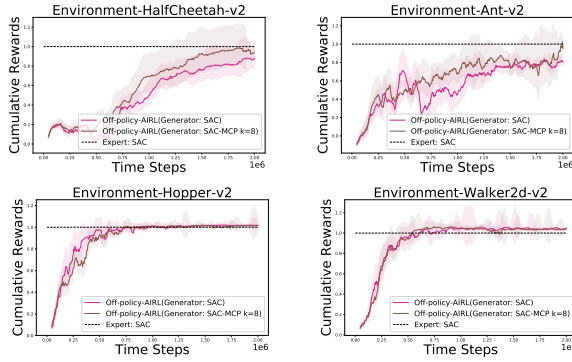


Figure 6. Performance Comparison of SAC and SAC-MCP as generator in off-policy-AIRL over continuous control task.

variance suffered by SAC.




Figure 7. Performance of off-policy-AIRL (Generator: SAC-MCP) during transfer learning. *Left*:Train only gating function. *Right*:Train both policy and gating function

8.4. Random State Initialization

We initialize the environment from a random state. As we retrain already learned policy over one task, it does not give much exploratory action. Thus initializing environment at different state will help to explore the new task. We use expert agent to run for 0-100 steps and then stop at random, thus enabling the learning agent to start at random state every time the environment resets. It allows an agent to explore from a state that is expected to be visited by an expert. At the same time, using learned policy (from prior imitation task) to take actions will result into better quality of sequential observations. We see in Figure 8 random initialization improves performance and also reduces performance variance for both SAC and SAC-MCP on transfer learning task for shifting maze environment.

9. Discussion and Conclusion

We propose an off-policy IRL algorithm in adversarial learning which learns policy using its approximated reward function. We use SAC as generator to improve imitation performance while still being sample efficient. We compare different way to compute reward functions for off-policy generator update and experimentally find simple state depen-

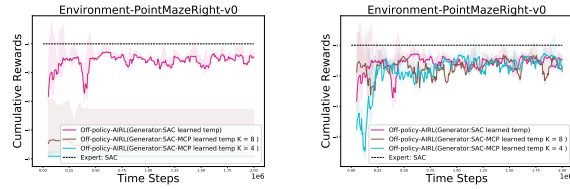


Figure 8. Performance Comparison of SAC and SAC-MCP as generator during transfer learning with random state initialization. *Left*:Train only gating function of SAC-MCP. *Right*:Train both policy and gating function of SAC-MCP

dent reward network without entropy regularization works best for our algorithm. We compare our imitation performance with state-of-the-art imitation learning algorithm to show our persistent expert like performance over continuous control environments.

We experimentally demonstrate the utility of learned reward function using off-policy-AIRL when the expert demonstrations are non-existent in transfer learning. We compare performance over two transfer learning criterion, where either (1) the dynamics of the agent or (2) the dynamics of the environment changes while the goal/objective and action dimension remain the same.

Furthermore, we propose using MCP model to learn single primitive to achieve more flexibility in retraining policy in transfer learning. We observe an improved performance for using SAC-MCP on both imitation performance and transfer learning. We also conduct experiments under random state initialization, which helps the agent to explore from an state that is expected to be visited in the new environment and reduces performance variance over prior experiments.

But unfortunately, our algorithm (using both SAC and SAC-MCP as generator) fails when the (criterion 1) dynamics of the agent is changed in transfer learning. The quadrupedal Ant has a very specific increment and decrement of the legs, which may have favored for (Fu et al., 2018) to re-learn policy when two legs are paralyzed in transfer learning. As specific body structure of an agent can facilitate learning policy over a task (Ha, 2018), experiments using environments with more structural flexibility is required to come to a conclusion under this criterion.

References

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

Faloutsos, P., Van de Panne, M., and Terzopoulos, D. Composable controllers for physics-based character animation.

- In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pp. 251–260. ACM, 2001.
- Finn, C., Christiano, P. F., Abbeel, P., and Levine, S. A connection between generative adversarial networks, inverse reinforcement learning, and energy-based models. *CoRR*, abs/1611.03852, 2016. URL <http://arxiv.org/abs/1611.03852>.
- Fu, J., Luo, K., and Levine, S. Learning robust rewards with adversarial inverse reinforcement learning. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=rkHywl-A->.
- Fujimoto, S., van Hoof, H., and Meger, D. Addressing function approximation error in actor-critic methods. *CoRR*, abs/1802.09477, 2018. URL <http://arxiv.org/abs/1802.09477>.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. C. Improved training of wasserstein gans. In *Advances in neural information processing systems*, pp. 5767–5777, 2017.
- Ha, D. Reinforcement learning for improving agent design. 2018. URL <https://designrl.github.io>. <https://designrl.github.io>.
- Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., Abbeel, P., and Levine, S. Soft actor-critic algorithms and applications. *CoRR*, abs/1812.05905, 2018. URL <http://arxiv.org/abs/1812.05905>.
- Ho, J. and Ermon, S. Generative adversarial imitation learning. *CoRR*, abs/1606.03476, 2016. URL <http://arxiv.org/abs/1606.03476>.
- Jordan, M. I. and Jacobs, R. A. Hierarchies of adaptive experts. In *Advances in neural information processing systems*, pp. 985–992, 1992.
- Kim, K.-E. and Park, H. S. Imitation learning via kernel mean embedding. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kostrikov, I., Agrawal, K. K., Dwibedi, D., Levine, S., and Tompson, J. Discriminator-actor-critic: Addressing sample inefficiency and reward bias in adversarial imitation learning. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=Hk4fpoA5Km>.
- Ng, A. Y., Harada, D., and Russell, S. Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*, volume 99, pp. 278–287, 1999a.
- Ng, A. Y., Harada, D., and Russell, S. J. Policy invariance under reward transformations: Theory and application to reward shaping. In *Proceedings of the Sixteenth International Conference on Machine Learning, ICML ’99*, pp. 278–287, San Francisco, CA, USA, 1999b. Morgan Kaufmann Publishers Inc. ISBN 1-55860-612-2. URL <http://dl.acm.org/citation.cfm?id=645528.657613>.
- Peng, X. B., Berseth, G., and Van de Panne, M. Terrain-adaptive locomotion skills using deep reinforcement learning. *ACM Transactions on Graphics (TOG)*, 35(4):81, 2016.
- Peng, X. B., Kanazawa, A., Toyer, S., Abbeel, P., and Levine, S. Variational discriminator bottleneck: Improving imitation learning, inverse rl, and gans by constraining information flow. *arXiv preprint arXiv:1810.00821*, 2018.
- Peng, X. B., Chang, M., Zhang, G., Abbeel, P., and Levine, S. MCP: learning composable hierarchical control with multiplicative compositional policies. *CoRR*, abs/1905.09808, 2019. URL <http://arxiv.org/abs/1905.09808>.
- Sutton, R. S., Precup, D., and Singh, S. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2): 181–211, 1999.
- Todorov, E., Erez, T., and Tassa, Y. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033. IEEE, 2012.
- Ziebart, B. D., Maas, A., Bagnell, J. A., and Dey, A. K. Maximum entropy inverse reinforcement learning. 2008.
- Ziebart, B. D., Bagnell, J. A., and Dey, A. K. Modeling interaction via the principle of maximum causal entropy. 2010.