

---

# Adversarial Soft Advantage Fitting: Imitation Learning without Policy Optimization

---

**Paul Barde**<sup>\*†</sup>  
Québec AI institute (Mila)  
McGill University  
bardepau@mila.quebec

**Julien Roy**<sup>\*†</sup>  
Québec AI institute (Mila)  
Polytechnique Montréal  
julien.roy@mila.quebec

**Wonseok Jeon**<sup>\*</sup>  
Québec AI institute (Mila)  
McGill University  
jeonwons@mila.quebec

**Joelle Pineau**<sup>‡</sup>  
Québec AI institute (Mila)  
McGill University  
Facebook AI Research

**Derek Nowrouzezahrai**  
Québec AI institute (Mila)  
McGill University

**Christopher Pal**<sup>‡</sup>  
Québec AI institute (Mila)  
Polytechnique Montréal  
Element AI

## Abstract

Adversarial Imitation Learning alternates between learning a discriminator – which tells apart expert’s demonstrations from generated ones – and a generator’s policy to produce trajectories that can fool this discriminator. This alternated optimization is known to be delicate in practice since it compounds unstable adversarial training with brittle and sample-inefficient reinforcement learning. We propose to remove the burden of the policy optimization steps by leveraging a novel discriminator formulation. Specifically, our discriminator is explicitly conditioned on two policies: the one from the previous generator’s iteration and a learnable policy. When optimized, this discriminator directly learns the optimal generator’s policy. Consequently, our discriminator’s update solves the generator’s optimization problem for free: learning a policy that imitates the expert does not require an additional optimization loop. This formulation effectively cuts by half the implementation and computational burden of Adversarial Imitation Learning algorithms by removing the Reinforcement Learning phase altogether. We show on a variety of tasks that our simpler approach is competitive to prevalent Imitation Learning methods.

## 1 Introduction

Imitation Learning (IL) treats the task of learning a policy from a set of expert demonstrations. IL is effective on control problems that are challenging for traditional Reinforcement Learning (RL) methods, either due to reward function design challenges or the inherent difficulty of the task itself [1, 24].

Most IL work can be divided into two branches: Behavioral Cloning and Inverse Reinforcement Learning. Behavioral Cloning casts IL as a supervised learning objective and seeks to imitate the expert’s actions using the provided demonstrations as a fixed dataset [19]. Thus, Behavioral Cloning usually requires a lot of expert data and results in agents that struggle to generalize. As an agent deviates from the demonstrated behaviors – straying outside the state distribution on which it was trained – the risks of making additional errors increase, a problem known as compounding error [24].

---

<sup>\*</sup>Equal contribution.

<sup>†</sup>Work conducted while interning at Ubisoft Montreal’s La Forge R&D laboratory.

<sup>‡</sup>Canada CIFAR AI Chair.

Inverse Reinforcement Learning aims to reduce compounding error by learning a reward function under which the expert policy is optimal [1]. Once learned, an agent can be trained (with any RL algorithm) to learn how to act at any given state of the environment. Early methods were prohibitively expensive on large environments because they required training the RL agent to convergence at each learning step of the reward function [30, 1]. Recent approaches instead apply an adversarial formulation (Adversarial Imitation Learning, AIL) in which a discriminator learns to distinguish between expert and agent behaviors to learn the reward optimized by the expert. AIL methods allow for the use of function approximators and can in practice be used with only a few policy improvement steps for each discriminator update [13, 6, 4].

While these advances have allowed Imitation Learning to tackle bigger and more complex environments [16, 3], they have also significantly complexified the implementation and learning dynamics of Imitation Learning algorithms. It is worth asking how much of this complexity is actually mandated. For example, in recent work, Reddy et al. [20] have shown that competitive performance can be obtained by hard-coding a very simple reward function to incentivize expert-like behaviors and manage to imitate it through off-policy direct RL. Reddy et al. [20] therefore remove the reward learning component of AIL and focus on the RL loop, yielding a regularized version of Behavioral Cloning. Motivated by these results, we also seek to simplify the AIL framework but following the opposite direction: keeping the reward learning module and removing the policy improvement loop.

We propose a simpler yet competitive AIL framework. Motivated by Finn et al. [4] who use the optimal discriminator form, we propose a structured discriminator that estimates the probability of demonstrated and generated behavior using a single parameterized maximum entropy policy. Discriminator learning and policy learning therefore occur simultaneously, rendering seamless generator updates: once the discriminator has been trained for a few epochs, we simply use its policy model to generate new rollouts. We call this approach Adversarial Soft Advantage Fitting (ASAF).

We make the following contributions:

- **Algorithmic:** we present a novel algorithm (ASAF) designed to imitate expert demonstrations without any Reinforcement Learning step.
- **Theoretical:** we show that our method retrieves the expert policy when trained to optimality.
- **Empirical:** we show that ASAF outperforms prevalent IL algorithms on a variety of discrete and continuous control tasks. We also show that, in practice, ASAF can be easily modified to account for different trajectory lengths (from full length to transition-wise).

## 2 Background

**Markov Decision Processes (MDPs)** We use Hazan et al. [12]’s notation and consider the classic  $T$ -horizon  $\gamma$ -discounted MDP  $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{P}_0, \gamma, r, T \rangle$ . For simplicity, we assume that  $\mathcal{S}$  and  $\mathcal{A}$  are finite. Successor states are given by the transition distribution  $\mathcal{P}(s'|s, a) \in [0, 1]$ , and the initial state  $s_0$  is drawn from  $\mathcal{P}_0(s) \in [0, 1]$ . Transitions are rewarded with  $r(s, a) \in \mathbb{R}$  with  $r$  being bounded. The discount factor and the episode horizon are  $\gamma \in [0, 1]$  and  $T \in \mathbb{N} \cup \{\infty\}$ , where  $T < \infty$  for  $\gamma = 1$ . Finally, we consider stationary stochastic policies  $\pi \in \Pi : \mathcal{S} \times \mathcal{A} \rightarrow ]0, 1[$  that produce trajectories  $\tau = (s_0, a_0, s_1, a_1, \dots, s_{T-1}, a_{T-1}, s_T)$  when executed on  $\mathcal{M}$ .

The probability of trajectory  $\tau$  under policy  $\pi$  is  $P_\pi(\tau) \triangleq \mathcal{P}_0(s_0) \prod_{t=0}^{T-1} \pi(a_t|s_t) \mathcal{P}(s_{t+1}|s_t, a_t)$  and the corresponding marginals are defined as  $d_{t,\pi}(s) \triangleq \sum_{\tau:s_t=s} P_\pi(\tau)$  and  $d_{t,\pi}(s, a) \triangleq \sum_{\tau:s_t=s, a_t=a} P_\pi(\tau) = d_{t,\pi}(s) \pi(a|s)$ , respectively. With these marginals, we define the normalized discounted state and state-action occupancy measures as  $d_\pi(s) \triangleq \frac{1}{Z(\gamma, T)} \sum_{t=0}^{T-1} \gamma^t d_{t,\pi}(s)$  and  $d_\pi(s, a) \triangleq \frac{1}{Z(\gamma, T)} \sum_{t=0}^{T-1} \gamma^t d_{t,\pi}(s, a) = d_\pi(s) \pi(a|s)$  where the partition function  $Z(\gamma, T)$  is equal to  $\sum_{t=0}^{T-1} \gamma^t$ . Intuitively, the state (or state-action) occupancy measure can be interpreted as the discounted visitation distribution of the states (or state-action pairs) that the agent encounters when navigating with policy  $\pi$ . The expected sum of discounted rewards can be expressed in term of the occupancy measures as follows:

$$J_\pi[r(s, a)] \triangleq \mathbb{E}_{\tau \sim P_\pi} \left[ \sum_{t=0}^{T-1} \gamma^t r(s_t, a_t) \right] = Z(\gamma, T) \mathbb{E}_{(s,a) \sim d_\pi} [r(s, a)].$$

In the entropy-regularized Reinforcement Learning framework [11], the optimal policy maximizes its entropy at each visited state in addition to the standard RL objective:

$$\pi^* \triangleq \arg \max_{\pi} J_{\pi}[r(s, a) + \alpha \mathcal{H}(\pi(\cdot|s))], \quad \mathcal{H}(\pi(\cdot|s)) = \mathbb{E}_{a \sim \pi(\cdot|s)}[-\log(\pi(a|s))].$$

As shown in [29, 10] the corresponding optimal policy is

$$\pi_{\text{soft}}^*(a|s) = \exp(\alpha^{-1} A_{\text{soft}}^*(s, a)) \quad \text{with} \quad A_{\text{soft}}^*(s, a) \triangleq Q_{\text{soft}}^*(s, a) - V_{\text{soft}}^*(s), \quad (1)$$

$$V_{\text{soft}}^*(s) = \alpha \log \sum_{a \in \mathcal{A}} \exp(\alpha^{-1} Q_{\text{soft}}^*(s, a)), \quad Q_{\text{soft}}^*(s, a) = r(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{P}(\cdot|s, a)} [V_{\text{soft}}^*(s')] \quad (2)$$

**Maximum Causal Entropy Inverse Reinforcement Learning** In the problem of Inverse Reinforcement Learning (IRL), it is assumed that the MDP’s reward function is unknown but that demonstrations from using expert’s policy  $\pi_E$  are provided. Maximum causal entropy IRL [30] proposes to fit a reward function  $r$  from a set  $\mathcal{R}$  of reward functions and retrieve the corresponding optimal policy by solving the optimization problem

$$\min_{r \in \mathcal{R}} \left( \max_{\pi} J_{\pi}[r(s, a) + \mathcal{H}(\pi(\cdot|s))] \right) - J_{\pi_E}[r(s, a)]. \quad (3)$$

In brief, the problem reduces to finding a reward function  $r$  for which the expert policy is optimal. In order to do so, the optimization procedure searches high entropy policies that are optimal with respect to  $r$  and minimizes the difference between their returns and the return of the expert policy, eventually reaching a policy  $\pi$  that approaches  $\pi_E$ . Most of the proposed solutions [1, 29, 13] transpose IRL to the problem of distribution matching; Abbeel and Ng [1] and Ziebart et al. [30] used linear function approximation and proposed to match the feature expectation; Ho and Ermon [13] proposed to cast Eq. (3) with a convex reward function regularizer into the problem of minimizing the Jensen-Shannon divergence between the state-action occupancy measures:

$$\min_{\pi} D_{\text{JS}}(d_{\pi}, d_{\pi_E}) - J_{\pi}[\mathcal{H}(\pi(\cdot|s))] \quad (4)$$

**Connections between Generative Adversarial Networks (GANs) and IRL** For the data distribution  $p_E$  and the generator distribution  $p_G$  defined on the domain  $\mathcal{X}$ , the GAN objective [8] is

$$\min_{p_G} \max_D L(D, p_G), \quad L(D, p_G) \triangleq \mathbb{E}_{x \sim p_E} [\log D(x)] + \mathbb{E}_{x \sim p_G} [\log(1 - D(x))]. \quad (5)$$

In Goodfellow et al. [8], the maximizer of the inner problem in Eq. (5) is shown to be

$$D_{p_G}^* \triangleq \arg \max_D L(D, p_G) = \frac{p_E}{p_E + p_G}, \quad (6)$$

and the optimizer for Eq. (5) is  $\arg \min_{p_G} \max_D L(D, p_G) = \arg \min_{p_G} L(D_{p_G}^*, p_G) = p_E$ . Later, Finn et al. [4] and Ho and Ermon [13] concurrently proposed connections between GANs and IRL. The Generative Adversarial Imitation Learning (GAIL) formulation in Ho and Ermon [13] is based on matching state-action occupancy measures, while Finn et al. [4] considered matching trajectory distributions. Our work is inspired by the discriminator proposed and used by Finn et al. [4],

$$D_{\theta}(\tau) \triangleq \frac{p_{\theta}(\tau)}{p_{\theta}(\tau) + q(\tau)}, \quad (7)$$

where  $p_{\theta}(\tau) \propto \exp r_{\theta}(\tau)$  with reward approximator  $r_{\theta}$  motivated by maximum causal entropy IRL. Note that Eq. (7) matches the form of the optimal discriminator in Eq. (6). Although Finn et al. [4] do not empirically support the effectiveness of their method, the Adversarial IRL approach of Fu et al. [6] (AIRL) successfully used a similar discriminator for state-action occupancy measure matching.

### 3 Imitation Learning without Policy Optimization

In this section, we derive Adversarial Soft Advantage Fitting (ASAF), our novel Adversarial Imitation Learning approach. Specifically, in Section 3.1, we present the theoretical foundations for ASAF to perform Imitation Learning on full-length trajectories. Intuitively, our method is based on the use

of such *structured discriminators* – that match the optimal discriminator form – to fit the trajectory distribution induced by the expert policy. This approach requires being able to evaluate and sample from the learned policy and allows us to learn that policy and train the discriminator simultaneously, thus drastically simplifying the training procedure. We present in Section 3.2 parametrization options that satisfy these requirements. Finally, in Section 3.3, we explain how to implement a practical algorithm that can be used for arbitrary trajectory-lengths, including the transition-wise case.

### 3.1 Adversarial Soft Advantage Fitting – Theoretical setting

Before introducing our method, we derive GAN training with a structured discriminator.

**GAN with structured discriminator** Suppose that we have a generator distribution  $p_G$  and some arbitrary distribution  $\tilde{p}$  and that both can be evaluated efficiently, e.g., categorical distribution or probability density with normalizing flows [22]. We call a *structured discriminator* a function  $D_{\tilde{p}, p_G} : \mathcal{X} \rightarrow [0, 1]$  of the form  $D_{\tilde{p}, p_G}(x) = \tilde{p}(x)/(\tilde{p}(x) + p_G(x))$  which matches the optimal discriminator form for Eq. (6). Considering our new GAN objective, we get:

$$\min_{p_G} \max_{\tilde{p}} L(\tilde{p}, p_G), \quad L(\tilde{p}, p_G) \triangleq \mathbb{E}_{x \sim p_E} [\log D_{\tilde{p}, p_G}(x)] + \mathbb{E}_{x \sim p_G} [\log(1 - D_{\tilde{p}, p_G}(x))]. \quad (8)$$

While the unstructured discriminator  $D$  from Eq. (5) learns a mapping from  $x$  to a Bernoulli distribution, we now learn a mapping from  $x$  to an arbitrary distribution  $\tilde{p}$  from which we can analytically compute  $D_{\tilde{p}, p_G}(x)$ . One can therefore say that  $D_{\tilde{p}, p_G}$  is *parameterized* by  $\tilde{p}$ . For the optimization problem of Eq. (8), we have the following optima:

**Lemma 1.** *The optimal discriminator parameter for any generator  $p_G$  in Eq. (8) is equal to the expert’s distribution,  $\tilde{p}^* \triangleq \arg \max_{\tilde{p}} L(\tilde{p}, p_G) = p_E$ , and the optimal discriminator parameter is also the optimal generator, i.e.,*

$$p_G^* \triangleq \arg \min_{p_G} \max_{\tilde{p}} L(\tilde{p}, p_G) = \arg \min_{p_G} L(p_E, p_G) = p_E = \tilde{p}^*.$$

*Proof.* See Appendix A.1

Intuitively, Lemma 1 shows that the optimal discriminator parameter is also the target data distribution of our optimization problem (i.e., the optimal generator). In other words, solving the inner optimization yields the solution of the outer optimization. In practice, we update  $\tilde{p}$  to minimize the discriminator objective and use it directly as  $p_G$  to sample new data.

**Matching trajectory distributions with structured discriminator** Motivated by the GAN with structured discriminator, we consider the trajectory distribution matching problem in IL. Here, we optimise Eq. (8) with  $x = \tau, \mathcal{X} = \mathcal{T}, p_E = P_{\pi_E}, p_G = P_{\pi_G}$ , which yields the following objective:

$$\min_{\pi_G} \max_{\tilde{\pi}} L(\tilde{\pi}, \pi_G), \quad L(\tilde{\pi}, \pi_G) \triangleq \mathbb{E}_{\tau \sim P_{\pi_E}} [\log D_{\tilde{\pi}, \pi_G}(\tau)] + \mathbb{E}_{\tau \sim P_{\pi_G}} [\log(1 - D_{\tilde{\pi}, \pi_G}(\tau))], \quad (9)$$

with the structured discriminator:

$$D_{\tilde{\pi}, \pi_G}(\tau) = \frac{P_{\tilde{\pi}}(\tau)}{P_{\tilde{\pi}}(\tau) + P_{\pi_G}(\tau)} = \frac{q_{\tilde{\pi}}(\tau)}{q_{\tilde{\pi}}(\tau) + q_{\pi_G}(\tau)}. \quad (10)$$

Here we used the fact that  $P_{\pi}(\tau)$  decomposes into two distinct products:  $q_{\pi}(\tau) \triangleq \prod_{t=0}^{T-1} \pi(a_t | s_t)$  which depends on the stationary policy  $\pi$  and  $\xi(\tau) \triangleq \mathcal{P}_0(s_0) \prod_{t=0}^{T-1} \mathcal{P}(s_{t+1} | s_t, a_t)$  which accounts for the environment dynamics. Crucially,  $\xi(\tau)$  cancels out in the numerator and denominator leaving  $\tilde{\pi}$  as the sole parameter of this structured discriminator. In this way,  $D_{\tilde{\pi}, \pi_G}(\tau)$  can evaluate the probability of a trajectory being generated by the expert policy simply by evaluating products of stationary policy distributions  $\tilde{\pi}$  and  $\pi_G$ . With this form, we can get the following result:

**Theorem 1.** *The optimal discriminator parameter for any generator policy  $\pi_G$  in Eq. (9)  $\tilde{\pi}^* \triangleq \arg \max_{\tilde{\pi}} L(\tilde{\pi}, \pi_G)$  is such that  $q_{\tilde{\pi}^*} = q_{\pi_E}$ , and using generator policy  $\tilde{\pi}^*$  minimizes  $L(\tilde{\pi}^*, \pi_G)$ , i.e.,*

$$\tilde{\pi}^* \in \arg \min_{\pi_G} \max_{\tilde{\pi}} L(\tilde{\pi}, \pi_G) = \arg \min_{\pi_G} L(\tilde{\pi}^*, \pi_G).$$

*Proof.* See Appendix A.2

Theorem 1’s benefits are similar to the ones from Lemma 1: we can use a discriminator of the form of Eq. (10) to fit to the expert demonstrations a policy  $\tilde{\pi}^*$  that simultaneously yields the optimal generator’s policy and produces the same trajectory distribution as the expert policy.

### 3.2 A Specific Policy Class

The derivations of Section 3.1 rely on the use of a learnable policy that can both be evaluated and sampled from in order to fit the expert policy. A number of parameterization options that satisfy these conditions are available.

First of all, we observe that since  $\pi_E$  is independent of  $r$  and  $\pi$ , we can add the entropy of the expert policy  $\mathcal{H}(\pi_E(\cdot|s))$  to the MaxEnt IRL objective of Eq. (3) without modifying the solution to the optimization problem:

$$\min_{r \in \mathcal{R}} \left( \max_{\pi \in \Pi} J_\pi[r(s, a) + \mathcal{H}(\pi(\cdot|s))] \right) - J_{\pi_E}[r(s, a) + \mathcal{H}(\pi_E(\cdot|s))] \quad (11)$$

The max over policies implies that when optimising  $r$ ,  $\pi$  has already been made optimal with respect to the causal entropy augmented reward function  $r'(s, a|\pi) = r(s, a) + \mathcal{H}(\pi(\cdot|s))$  and therefore it must be of the form presented in Eq. (1). Moreover, since  $\pi$  is optimal w.r.t.  $r'$  the difference in performance  $J_\pi[r'(s, a|\pi)] - J_{\pi_E}[r'(s, a|\pi_E)]$  is always non-negative and its minimum of 0 is only reached when  $\pi_E$  is also optimal w.r.t.  $r'$ , in which case  $\pi_E$  must also be of the form of Eq. (1).

With discrete action spaces we propose to parameterize the MaxEnt policy defined in Eq. (1) with the following categorical distribution  $\tilde{\pi}(a|s) = \exp(Q_\theta(s, a) - \log \sum_{a'} \exp Q_\theta(s, a'))$ , where  $Q_\theta$  is a model parameterized by  $\theta$  that approximates  $\frac{1}{\alpha} Q_{\text{soft}}^*$ .

With continuous action spaces, the soft value function involves an intractable integral over the action domain. Therefore, we approximate the MaxEnt distribution with a Normal distribution with diagonal covariance matrix like it is commonly done in the literature [11, 17]. By parameterizing the mean and variance we get a learnable density function that can be easily evaluated and sampled from.

### 3.3 Adversarial Soft Advantage Fitting (ASAF) – practical algorithm

Section 3.1 shows that assuming  $\tilde{\pi}$  can be evaluated and sampled from, we can use the structured discriminator of Eq. (10) to learn a policy  $\tilde{\pi}$  that matches the expert’s trajectory distribution. Section 3.2 proposes parameterizations for discrete and continuous action spaces that satisfy those assumptions.

In practice, as with GANs [8], we do not train the discriminator to convergence as gradient-based optimisation cannot be expected to find the global optimum of non-convex problems. Instead, Adversarial Soft Advantage Fitting (ASAF) alternates between two simple steps: (1) training  $D_{\tilde{\pi}, \pi_G}$  by minimizing the binary cross-entropy loss,

$$\mathcal{L}_{BCE}(D_E, D_G, \tilde{\pi}) \approx -\frac{1}{n_E} \sum_{i=1}^{n_E} \log D_{\tilde{\pi}, \pi_G}(\tau_i^{(E)}) - \frac{1}{n_G} \sum_{i=1}^{n_G} \log \left( 1 - D_{\tilde{\pi}, \pi_G}(\tau_i^{(G)}) \right) \quad (12)$$

where  $\tau_i^{(E)} \sim \mathcal{D}_E$ ,  $\tau_i^{(G)} \sim \mathcal{D}_G$  and  $D_{\tilde{\pi}, \pi_G}(\tau) = \frac{\prod_{t=0}^{T-1} \tilde{\pi}(a_t|s_t)}{\prod_{t=0}^{T-1} \tilde{\pi}(a_t|s_t) + \prod_{t=0}^{T-1} \pi_G(a_t|s_t)}$

with minibatch sizes  $n_E = n_G$ , and (2) updating the generator’s policy as  $\pi_G \leftarrow \tilde{\pi}$  to minimize Eq. (9) (see Algorithm 1).

We derived ASAF considering full trajectories, yet it might be preferable in practice to split full trajectories into smaller chunks. This is particularly true in environments where trajectory length varies a lot or tends to infinity.

To investigate whether the practical benefits of using partial trajectories hurt ASAF’s performance, we also consider a variation, ASAF- $w$ , where we treat trajectory-windows of size  $w$  as if they were full trajectories. Note that considering windows as full trajectories results in approximating that the initial state of these sub-trajectories have equal probability under the expert’s and the generator’s policy (this is easily seen when deriving Eq. (10)).

In the limit, ASAF-1 (window-size of 1) becomes a transition-wise algorithm which can be desirable if one wants to collect roll-outs asynchronously or has only access to unsequential expert data. While ASAF-1 may work well in practice it essentially assumes that the expert’s and the generator’s policies have the same state occupancy measure, which is incorrect until actually recovering the true expert policy.

Finally, to offer a complete family of algorithms based on the structured discriminator approach, we show in Appendix B that this assumption is not mandatory and derive a transition-wise algorithm based on Soft Q-function Fitting (rather than soft advantages) that also gets rid of the RL loop. We call this algorithm ASQF. While theoretically sound, we found that in practice, ASQF is outperformed by ASAF-1 in more complex environments (see Section 5.1).

---

**Algorithm 1: ASAF**

---

**Require:** expert trajectories  $\mathcal{D}_E = \{\tau_i\}_{i=1}^{N_E}$   
 Randomly initialize  $\tilde{\pi}$  and set  $\pi_G \leftarrow \tilde{\pi}$   
**for** steps  $m = 0$  to  $M$  **do**  
   Collect trajectories  $\mathcal{D}_G = \{\tau_i\}_{i=1}^{N_G}$  using  $\pi_G$   
   Update  $\tilde{\pi}$  by minimizing Eq. (12)  
   Set  $\pi_G \leftarrow \tilde{\pi}$   
**end for**

---

## 4 Related works

Ziebart et al. [30] first proposed MaxEnt IRL, the foundation of modern IL. Ziebart [29] further elaborated MaxEnt IRL as well as deriving the optimal form of the MaxEnt policy at the core of our methods. Finn et al. [4] proposed a GAN formulation to IRL that leveraged the energy based models of Ziebart [29]. Finn et al. [5]’s implementation of this method, however, relied on processing full trajectories with Linear Quadratic Regulator and on optimizing with guided policy search, to manage the high variance of trajectory costs. To retrieve robust rewards, Fu et al. [6] proposed a straightforward transposition of [4] to state-action transitions. In doing so, they had to however do away with a GAN objective during policy optimization, consequently minimizing the Kullback–Leibler divergence from the expert occupancy measure to the policy occupancy measure (instead of the Jensen-Shannon divergence) [7].

Later works [25, 15] move away from the Generative Adversarial formulation. To do so, Sasaki et al. [25] directly express the expectation of the Jensen-Shannon divergence between the occupancy measures in term of the agent’s Q-function, which can then be used to optimize the agent’s policy with off-policy Actor-Critic [2]. Similarly, Kostrikov et al. [15] use Dual Stationary Distribution Correction Estimation [18] to approximate the Q-function on the expert’s demonstrations before optimizing the agent’s policy under the initial state distribution using the reparametrization trick [11]. While [25, 15] are related to our methods in their interests in learning directly the value function, they differ in their goal and thus in the resulting algorithmic complexity. Indeed, they aim at improving the sample efficiency in terms of environment interaction and therefore move away from the algorithmically simple Generative Adversarial formulation towards more complicated divergence minimization methods. In doing so, they further complicate the Imitation Learning methods while still requiring to explicitly learn a policy. Yet, simply using the Generative Adversarial formulation with an Experience Replay Buffer can significantly improve the sample efficiency [14]. For these reasons, and since our aim is to propose efficient yet simple methods, we focus on the Generative Adversarial formulation.

While Reddy et al. [20] share our interest for simpler IL methods, they pursue an opposite approach to ours. They propose to eliminate the reward learning steps of IRL by simply hard-coding a reward of 1 for expert’s transitions and of 0 for agent’s transitions. They then use Soft Q-learning [10] to learn a value function by sampling transitions in equal proportion from the expert’s and agent’s buffers. Unfortunately, once the learner accurately mimics the expert, it collects expert-like transitions that are labeled with a reward of 0 since they are generated and not coming from the demonstrations. This effectively causes the reward of expert-like behavior to decay as the agent improves and can severely destabilize learning to a point where early-stopping becomes required [20].

Our work builds on [4], yet its novelty is to explicitly express the probability of a trajectory in terms of the policy in order to directly learn this latter when training the discriminator. In contrast, [6] considers a transition-wise discriminator with un-normalized probabilities which makes it closer to ASQF (Appendix B) than to ASAF-1. Additionally, AIRL [6] minimizes the Kullback-Leiber Divergence [7] between occupancy measures whereas ASAF minimizes the Jensen-Shanon Divergence between trajectory distributions.

Finally, Behavioral Cloning uses the loss function from supervised learning (classification or regression) to match expert’s actions given expert’s states and suffers from compounding error due to co-variate shift [23] since its data is limited to the demonstrated state-action pairs without environment interaction. Contrarily, ASAF-1 uses the binary cross entropy loss in Eq. (12) and does not suffer from compounding error as it learns on both generated and expert’s trajectories.

## 5 Results and discussion

We evaluate our methods on a variety of discrete and continuous control tasks. Our results show that, in addition to drastically simplifying the adversarial IRL framework, our methods perform on par or better than previous approaches on all but one environment. When trajectory length is really long or drastically varies across episodes (see MuJoCo experiments Section 5.3), we find that using sub-trajectories with fixed window-size (ASAF- $w$  or ASAF-1) significantly outperforms its full trajectory counterpart ASAF.

### 5.1 Experimental setup

We compare our algorithms ASAF, ASAF- $w$  and ASAF-1 against GAIL [13], the predominant Adversarial Imitation Learning algorithm in the literature, and AIRL [6], one of its variations that also leverages the access to the generator’s policy distribution. Additionally, we compare against SQIL [20], a recent Reinforcement Learning-only approach to Imitation Learning that proved successful on high-dimensional tasks. Our implementations of GAIL and AIRL use PPO [27] instead of TRPO [26] as it has been shown to improve performance [14]. Finally, to be consistent with [13], we do not use causal entropy regularization.

For all tasks except MuJoCo, we selected the best performing hyperparameters through a random search of equal budget for each algorithm-environment pair (see Appendix D) and the best configuration is retrained on ten random seeds. For the MuJoCo experiments, GAIL required extensive tuning (through random searches) of both its RL and IRL components to achieve satisfactory performances. Our methods, ASAF- $w$  and ASAF-1, on the other hand showed much more stable and robust to hyperparameterization, which is likely due to their simplicity. SQIL used the same SAC[11] implementation and hyperparameters that were used to generate the expert demonstrations.

Finally for each task, all algorithms use the same neural network architectures for their policy and/or discriminator (see full description in Appendix D). Expert demonstrations are either generated by hand (mountaincar), using open-source bots (Pommerman) or from our implementations of SAC and PPO (all remaining). More details are given in Appendix E.

### 5.2 Experiments on classic control and Box2D tasks (discrete and continuous)

Figure 1 shows that ASAF and its approximate variations ASAF-1 and ASAF- $w$  quickly converge to expert’s performance (here  $w$  was tuned to values between 32 to 200, see Appendix D for selected window-sizes). This indicates that the practical benefits of using shorter trajectories or even just transitions does not hinder performance on these simple tasks. Note that for Box2D and classic control environments, we retrain the best configuration of each algorithm for twice as long than was done in the hyperparameter search, which allows to uncover unstable learning behaviors. Figure 1 shows that our methods display much more stable learning: their performance rises until they match the expert’s and does not decrease once it is reached. This is a highly desirable property for an Imitation Learning algorithm since in practice one does not have access to a reward function and thus cannot monitor the performance of the learning algorithm to trigger early-stopping. The baselines on the other hand experience occasional performance drops. For GAIL and AIRL, this is likely due to the concurrent RL and IRL loops, whereas for SQIL, it has been noted that an effective reward decay can occur when accurately mimicking the expert [20]. This instability is particularly severe in the continuous control case. In practice, all three baselines use early stopping to avoid performance decay [20].

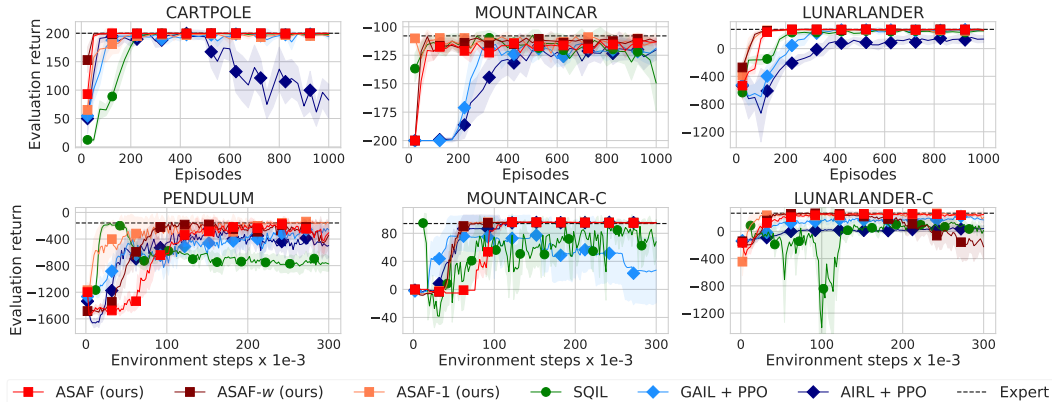


Figure 1: Results on classic control and Box2D tasks for 10 expert demonstrations. First row contains discrete actions environments, second row corresponds to continuous control.

### 5.3 Experiments on MuJoCo (continuous control)

To scale up our evaluations in continuous control we use the popular MuJoCo benchmarks. In this domain, the trajectory length is either fixed at a large value (1000 steps on HalfCheetah) or varies a lot across episodes due to termination when the character falls down (Hopper, Walker2d and Ant). Figure 2 shows that these trajectory characteristics hinder ASAF’s learning as ASAF requires collecting multiple episodes for every update, while ASAF-1 and ASAF-w perform well and are more sample-efficient than ASAF in these scenarios. We focus on GAIL since [6] claim that AIRL performs on par with it on MuJoCo environments. In Figure 5 in Appendix C we evaluate GAIL both with and without gradient penalty (GP) on discriminator updates [9, 14] and while GAIL was originally proposed without GP [13], we empirically found that GP prevents the discriminator to overfit and enables RL to exploit dense rewards, which highly improves its sample efficiency. Despite these ameliorations, GAIL proved to be quite inconsistent across environments despite substantial efforts on hyperparameter tuning. On the other hand, ASAF-1 performs well across all environments. Finally, we see that SQIL’s instability is exacerbated on MuJoCo.

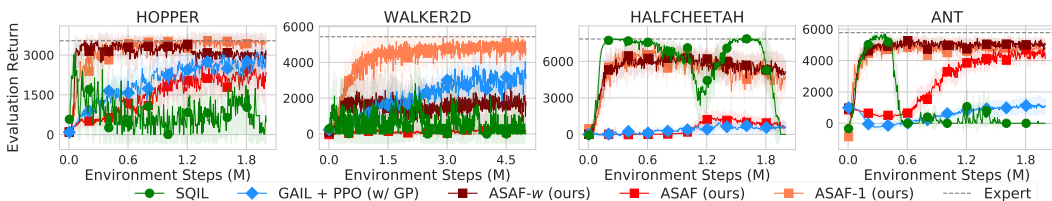


Figure 2: Results on MuJoCo tasks for 25 expert demonstrations.

### 5.4 Experiments on Pommerman (discrete control)

Finally, to scale up our evaluations in discrete control environments, we consider the domain of Pommerman [21], a challenging and very dynamic discrete control environment that uses rich and high-dimensional observation spaces (see Appendix E). We perform evaluations of all of our methods and baselines on a 1 vs 1 task where a learning agent plays against a random agent, the opponent. The goal for the learning agent is to navigate to the opponent and eliminate it using expert demonstrations provided by the champion algorithm of the FFA 2018 competition [28]. We removed the ability of the opponent to lay bombs so that it doesn’t accidentally eliminate itself. Since it can still move around, it is however surprisingly tricky to eliminate: the expert has to navigate across the whole map, lay a bomb next to the opponent and retreat to avoid eliminating itself. This entire routine has then to be repeated several times until finally succeeding since the opponent will often avoid the hit by chance. We refer to this task as *Pommerman Random-Tag*. Note that since we measure success of the imitation task with the win-tie-lose outcome (sparse performance metric), a learning agent has to truly reproduce the expert behavior until the very end of trajectories to achieve higher scores.



Figure 3 shows that all three variations of ASAF as well as Behavioral Cloning (BC) outperform the baselines.

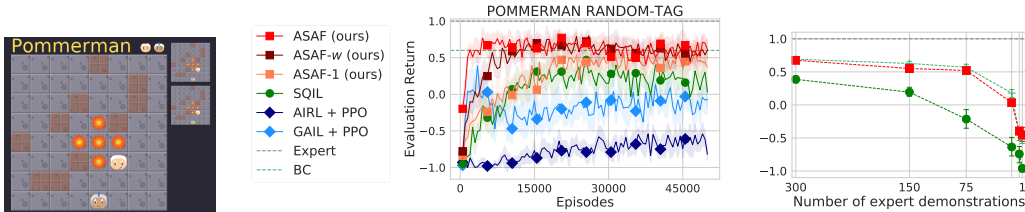


Figure 3: Results on Pommerman Random-Tag: (Left) Snapshot of the environment. (Center) Learning measured as evaluation return over episodes for 150 expert trajectories (Right) Average return on last 20% of training for decreasing number of expert trajectories [300, 150, 75, 15, 5, 1].

## 6 Conclusion

We propose an important simplification to the Adversarial Imitation Learning framework by removing the Reinforcement Learning optimisation loop altogether. We show that, by using a particular form for the discriminator, our method recovers a policy that matches the expert’s trajectory distribution. We evaluate our approach against prior works on many different benchmarking tasks and show that our method (ASAF) compares favorably to the predominant Imitation Learning algorithms. The approximate versions, ASAF-w and ASAF-1, that use sub-trajectories yield a flexible algorithms that work well both on short and long time horizons. Finally, our approach still involves a reward learning module through its discriminator, and it would be interesting in future work to explore how ASAF can be used to learn robust rewards, along the lines of Fu et al. [6].

## Broader Impact

Our contributions are mainly theoretical and aim at simplifying current Imitation Learning methods. We do not propose new applications nor use sensitive data or simulator. Yet our method can ease and promote the use, design and development of Imitation Learning algorithms and may eventually lead to applications outside of simple and controlled simulators. We do not pretend to discuss the ethical implications of the general use of autonomous agents but we rather try to investigate what are some of the differences in using Imitation Learning rather than reward oriented methods in the design of such agents.

Using only a scalar reward function to specify the desired behavior of an autonomous agent is a challenging task as one must weight different desiderata and account for unsuspected behaviors and situations. Indeed, it is well known in practice that Reinforcement Learning agents tend to find bizarre ways of exploiting the reward signal without solving the desired task. The fact that it is difficult to specify and control the behavior of an RL agents is a major flaw that prevent current methods to be applied to risk sensitive situations. On the other hand, Imitation Learning proposes a more natural way of specifying nuanced preferences by demonstrating desirable ways of solving a task. Yet, IL also has its drawbacks. First of all one needs to be able to demonstrate the desired behavior and current methods tend to be only as good as the demonstrator. Second, it is a challenging problem to ensure that the agent will be able to adapt to new situations that do not resemble the demonstrations. For these reasons, it is clear for us that additional safeguards are required in order to apply Imitation Learning (and Reinforcement Learning) methods to any application that could effectively have a real world impact.

## Acknowledgments and Disclosure of Funding

We thank Eloi Alonso, Olivier Delalleau, Félix G. Harvey, Maxim Peter and the entire research team at Ubisoft Montreal’s La Forge R&D laboratory. Their feedback and comments contributed significantly to this work. Christopher Pal and Derek Nowrouzezahrai also acknowledge funding from the Fonds de Recherche Nature et Technologies (FRQNT), Ubisoft Montreal and Mitacs’ Accelerate Program in support of our work, as well as Compute Canada for providing computing resources.

## References

- [1] Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the 21st International Conference on Machine Learning (ICML)*, 2004.
- [2] Thomas Degris, Martha White, and Richard S Sutton. Off-policy actor-critic. In *Proceedings of the 29th International Conference on Machine Learning (ICML)*, pages 179–186, 2012.
- [3] Yiming Ding, Carlos Florensa, Pieter Abbeel, and Mariano Phielipp. Goal-conditioned imitation learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 15298–15309, 2019.
- [4] Chelsea Finn, Paul Christiano, Pieter Abbeel, and Sergey Levine. A connection between generative adversarial networks, inverse reinforcement learning, and energy-based models. *arXiv preprint arXiv:1611.03852*, 2016.
- [5] Chelsea Finn, Sergey Levine, and Pieter Abbeel. Guided cost learning: Deep inverse optimal control via policy optimization. In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, pages 49–58, 2016.
- [6] Justin Fu, Katie Luo, and Sergey Levine. Learning robust rewards with adversarial inverse reinforcement learning. In *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, 2017.
- [7] Seyed Kamyar Seyed Ghasemipour, Richard Zemel, and Shixiang Gu. A divergence minimization perspective on imitation learning methods. In *Proceedings of the 3rd Conference on Robot Learning (CoRL)*, 2019.
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2672–2680, 2014.
- [9] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of Wasserstein GANs. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 5767–5777, 2017.
- [10] Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. Reinforcement learning with deep energy-based policies. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pages 1352–1361, 2017.
- [11] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.
- [12] Elad Hazan, Sham M Kakade, Karan Singh, and Abby Van Soest. Provably efficient maximum entropy exploration. *arXiv preprint arXiv:1812.02690*, 2018.
- [13] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 4565–4573, 2016.
- [14] Ilya Kostrikov, Kumar Krishna Agrawal, Debidatta Dwibedi, Sergey Levine, and Jonathan Tompson. Discriminator-actor-critic: Addressing sample inefficiency and reward bias in adversarial imitation learning. In *Proceedings of the 7th International Conference on Learning Representations (ICLR)*, 2019.
- [15] Ilya Kostrikov, Ofir Nachum, and Jonathan Tompson. Imitation learning via off-policy distribution matching. In *Proceedings of the 8th International Conference on Learning Representations (ICLR)*, 2020.
- [16] Alex Kuefler, Jeremy Morton, Tim Wheeler, and Mykel Kochenderfer. Imitating driver behavior with generative adversarial networks. In *Proceedings of 2017 IEEE Intelligent Vehicles Symposium (IV)*, pages 204–211, 2017.

- [17] Ofir Nachum, Mohammad Norouzi, Kelvin Xu, and Dale Schuurmans. Trust-PCL: An off-policy trust region method for continuous control. In *Proceedings of the 6th International Conference on Learning Representations (ICLR)*, 2018.
- [18] Ofir Nachum, Yinlam Chow, Bo Dai, and Lihong Li. DualDICE: Behavior-agnostic estimation of discounted stationary distribution corrections. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2318–2328, 2019.
- [19] Dean A Pomerleau. Efficient training of artificial neural networks for autonomous navigation. *Neural computation*, 3(1):88–97, 1991.
- [20] Siddharth Reddy, Anca D. Dragan, and Sergey Levine. SQL: Imitation learning via reinforcement learning with sparse rewards, 2019.
- [21] Cinjon Resnick, Wes Eldridge, David Ha, Denny Britz, Jakob Foerster, Julian Togelius, Kyunghyun Cho, and Joan Bruna. Pommerman: A multi-agent playground. *arXiv preprint arXiv:1809.07124*, 2018.
- [22] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, pages 1530–1538, 2015.
- [23] Stéphane Ross and Drew Bagnell. Efficient reductions for imitation learning. In *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 661–668, 2010.
- [24] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 627–635, 2011.
- [25] Fumihiko Sasaki, Tetsuya Yohira, and Atsuo Kawaguchi. Sample efficient imitation learning for continuous control. In *Proceedings of the 6th International Conference on Learning Representations (ICLR)*, 2018.
- [26] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, pages 1889–1897, 2015.
- [27] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [28] Hongwei Zhou, Yichen Gong, Luvneesh Mugrai, Ahmed Khalifa, Andy Nealen, and Julian Togelius. A hybrid search agent in pommerman. In *Proceedings of the 13th International Conference on the Foundations of Digital Games (FDG)*, pages 1–4, 2018.
- [29] Brian D Ziebart. *Modeling purposeful adaptive behavior with the principle of maximum causal entropy*. PhD thesis, 2010.
- [30] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. Maximum entropy inverse reinforcement learning. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence*, pages 1433–1438, 2008.

## Appendix

### A Proofs

#### A.1 Proof of Lemma 1

*Proof.* Lemma 1 states that given  $L(\tilde{p}, p_G)$  defined in Eq. (8):

- (a)  $\tilde{p}^* \triangleq \arg \max_{\tilde{p}} L(\tilde{p}, p_G) = p_E$
- (b)  $\arg \min_{p_G} L(p_E, p_G) = p_E$

Starting with (a), we have:

$$\begin{aligned} \arg \max_{\tilde{p}} L(\tilde{p}, p_G) &= \arg \max_{\tilde{p}} \sum_{x_i} p_E(x_i) \log D_{\tilde{p}, p_G}(x_i) + p_G(x_i) \log(1 - D_{\tilde{p}, p_G}(x_i)) \\ &\triangleq \arg \max_{\tilde{p}} \sum_{x_i} L_i \end{aligned}$$

Assuming infinite discriminator's capacity,  $L_i$  can be made independent for all  $x_i \in \mathcal{X}$  and we can construct our optimal discriminator  $D_{\tilde{p}, p_G}^*$  as a look-up table  $D_{\tilde{p}, p_G}^* : \mathcal{X} \rightarrow ]0, 1[$ ;  $x_i \mapsto D_i^*$  with  $D_i^*$  the optimal discriminator for each  $x_i$  defined as:

$$D_i^* = \arg \max_{D_i} L_i = \arg \max_{D_i} p_{E,i} \log D_i + p_{G,i} \log(1 - D_i), \quad (13)$$

with  $p_{G,i} \triangleq p_G(x_i)$ ,  $p_{E,i} \triangleq p_E(x_i)$  and  $D_i \triangleq D(x_i)$ .

Recall that  $D_i \in ]0, 1[$  and that  $p_{G,i} \in ]0, 1[$ . Therefore the function  $\tilde{p}_i \mapsto D_i = \frac{\tilde{p}_i}{\tilde{p}_i + p_{G,i}}$  is defined for  $\tilde{p}_i \in ]0, +\infty[$ . Since it is strictly monotonic over that domain we have that:

$$D_i^* = \arg \max_{D_i} L_i \Leftrightarrow \tilde{p}_i^* = \arg \max_{\tilde{p}_i} L_i \quad (14)$$

Taking the derivative and setting to zero, we get:

$$\left. \frac{dL_i}{d\tilde{p}_i} \right|_{\tilde{p}_i^*} = 0 \Leftrightarrow \tilde{p}_i^* = p_{E,i} \quad (15)$$

The second derivative test confirms that we have a maximum, i.e.  $\left. \frac{d^2 L_i}{d\tilde{p}_i^2} \right|_{\tilde{p}_i^*} < 0$ . The values of  $L_i$  at the boundaries of the domain of definition of  $\tilde{p}_i$  tend to  $-\infty$ , therefore  $L_i(\tilde{p}_i^* = p_{E,i})$  is the global maximum of  $L_i$  w.r.t.  $\tilde{p}_i$ . Finally, the optimal global discriminator is given by:

$$D_{\tilde{p}, p_G}^*(x) = \frac{p_E(x)}{p_E(x) + p_G(x)} \quad \forall x \in \mathcal{X} \quad (16)$$

This concludes the proof for (a).

The proof for (b) can be found in the work of Goodfellow et al. [8]. We reproduce it here for completion. Since from (a) we know that  $\tilde{p}^*(x) = p_E(x) \forall x \in \mathcal{X}$ , we can write the GAN objective for the optimal discriminator as:

$$\arg \min_{p_G} L(\tilde{p}^*, p_G) = \arg \min_{p_G} L(p_E, p_G) \quad (17)$$

$$= \arg \min_{p_G} \mathbb{E}_{x \sim p_E} \left[ \log \frac{p_E(x)}{p_E(x) + p_G(x)} \right] + \mathbb{E}_{x \sim p_G} \left[ \log \frac{p_G(x)}{p_E(x) + p_G(x)} \right] \quad (18)$$

Note that:

$$\log 4 = \mathbb{E}_{x \sim p_E} [\log 2] + \mathbb{E}_{x \sim p_G} [\log 2] \quad (19)$$

Adding Eq. (19) to Eq. (18) and subtracting  $\log 4$  on both sides:

$$\arg \min_{p_G} L(p_E, p_G) = -\log 4 + \mathbb{E}_{x \sim p_E} \left[ \log \frac{2p_E(x)}{p_E(x) + p_G(x)} \right] + \mathbb{E}_{x \sim p_G} \left[ \log \frac{2p_G(x)}{p_E(x) + p_G(x)} \right] \quad (20)$$

$$= -\log 4 + D_{\text{KL}} \left( p_E \left\| \frac{p_E + p_G}{2} \right. \right) + D_{\text{KL}} \left( p_G \left\| \frac{p_E + p_G}{2} \right. \right) \quad (21)$$

$$= -\log 4 + 2D_{\text{JS}}(p_E \| p_G) \quad (22)$$

Where  $D_{\text{KL}}$  and  $D_{\text{JS}}$  are respectively the Kullback-Leibler and the Jensen-Shannon divergences. Since the Jensen-Shannon divergence between two distributions is always non-negative and zero if and only if the two distributions are equal, we have that  $\arg \min_{p_G} L(p_E, p_G) = p_E$ .

This concludes the proof for (b).  $\square$

## A.2 Proof of Theorem 1

*Proof.* Theorem 1 states that given  $L(\tilde{\pi}, \pi_G)$  defined in Eq. (9):

$$(a) \tilde{\pi}^* \triangleq \arg \max_{\tilde{\pi}} L(\tilde{\pi}, \pi_G) \text{ satisfies } q_{\tilde{\pi}^*} = q_{\pi_E}$$

$$(b) \pi_G^* = \tilde{\pi}^* \in \arg \min_{\pi_G} L(\tilde{\pi}^*, \pi_G)$$

The proof of (a) is very similar to the one from Lemma 1. Starting from Eq. (9) we have:

$$\arg \max_{\tilde{\pi}} L(\tilde{\pi}, \pi_G) = \arg \max_{\tilde{\pi}} \sum_{\tau_i} P_{\pi_E}(\tau_i) \log D_{\tilde{\pi}, \pi_G}(\tau_i) + P_{\pi_G}(\tau_i) \log(1 - D_{\tilde{\pi}, \pi_G}(\tau_i)) \quad (23)$$

$$= \arg \max_{\tilde{\pi}} \sum_{\tau_i} \xi(\tau_i) \left( q_{\pi_E}(\tau_i) \log D_{\tilde{\pi}, \pi_G}(\tau_i) + q_{\pi_G}(\tau_i) \log(1 - D_{\tilde{\pi}, \pi_G}(\tau_i)) \right) \quad (24)$$

$$= \arg \max_{\tilde{\pi}} \sum_{\tau_i} L_i \quad (25)$$

Like for Lemma 1, we can optimise for each  $L_i$  individually. When doing so,  $\xi(\tau_i)$  can be omitted as it is constant w.r.t  $\tilde{\pi}$ . The rest of the proof is identical to the one of but Lemma 1 with  $p_E = q_{\pi_E}$  and  $p_G = q_{\pi_G}$ . It follows that the max of  $L(\tilde{\pi}, \pi_G)$  is reached for  $q_{\tilde{\pi}^*} = q_{\pi_E}$ . From that we obtain that the policy  $\tilde{\pi}^*$  that makes the discriminator  $D_{\tilde{\pi}^*, \pi_G}$  optimal w.r.t  $L(\tilde{\pi}, \pi_G)$  is such that  $q_{\tilde{\pi}^*} = q_{\tilde{\pi}^*} = q_{\pi_E}$  i.e.  $\prod_{t=0}^{T-1} \tilde{\pi}^*(a_t | s_t) = \prod_{t=0}^{T-1} \pi_E(a_t | s_t) \forall \tau$ .

The proof for (b) stems from the observation that choosing  $\pi_G = \tilde{\pi}^*$  (the policy recovered by the optimal discriminator  $D_{\tilde{\pi}^*, \pi_G}$ ) minimizes  $L(\tilde{\pi}^*, \pi_G)$ :

$$\pi_G(a|s) = \tilde{\pi}^*(a|s) \forall (s, a) \in \mathcal{S} \times \mathcal{A} \Rightarrow \prod_{t=0}^{T-1} \pi_G(a_t | s_t) = \prod_{t=0}^{T-1} \tilde{\pi}^*(a_t | s_t) \forall \tau \in \mathcal{T} \quad (26)$$

$$\Rightarrow q_{\pi_G}(\tau) = q_{\pi_E}(\tau) \forall \tau \in \mathcal{T} \quad (27)$$

$$\Rightarrow D_{\tilde{\pi}^*, \tilde{\pi}^*} = \frac{1}{2} \forall \tau \in \mathcal{T} \quad (28)$$

$$\Rightarrow L(\tilde{\pi}^*, \tilde{\pi}^*) = -\log 4 \quad (29)$$

By multiplying the numerator and denominator of  $D_{\tilde{\pi}^*, \tilde{\pi}^*}$  by  $\xi(\tau)$  it can be shown in exactly the same way as in Appendix A.1 that  $-\log 4$  is the global minimum of  $L(\tilde{\pi}^*, \pi_G)$ .  $\square$

## B Adversarial Soft Q-Fitting: transition-wise Imitation Learning without Policy Optimization

In this section we present Adversarial Soft Q-Fitting (ASQF), a principled approach to Imitation Learning without Reinforcement Learning that relies exclusively on transitions. Using transitions rather than trajectories presents several practical benefits such as the possibility to deal with asynchronously collected data or non-sequential experts demonstrations. We first present the theoretical setting for ASQF and then test it on a variety of discrete control tasks. We show that while it is theoretically sound, ASQF is often outperformed by ASAF-1, an approximation to ASAF that also allows to rely on transitions instead of trajectories.

**Theoretical Setting** We consider the GAN objective of Eq. (5) with  $x = (s, a)$ ,  $\mathcal{X} = \mathcal{S} \times \mathcal{A}$ ,  $p_E = d_{\pi_E}$ ,  $p_G = d_{\pi_G}$  and a discriminator  $D_{\tilde{f}, \pi_G}$  of the form of Fu et al. [6]:

$$\begin{aligned} \min_{\pi_G} \max_{\tilde{f}} L(\tilde{f}, \pi_G), \quad L(\tilde{f}, \pi_G) \triangleq \mathbb{E}_{d_{\pi_E}} [\log D_{\tilde{f}, \pi_G}(s, a)] + \mathbb{E}_{d_{\pi_G}} [\log(1 - D_{\tilde{f}, \pi_G}(s, a))], \\ \text{with } D_{\tilde{f}, \pi_G} = \frac{\exp \tilde{f}(s, a)}{\exp \tilde{f}(s, a) + \pi_G(a|s)}, \end{aligned} \quad (30)$$

for which we present the following theorem.

**Theorem 2.** *For any generator policy  $\pi_G$ , the optimal discriminator parameter for Eq. (30) is*

$$\tilde{f}^* \triangleq \arg \max_{\tilde{f}} L(\tilde{f}, \pi_G) = \log \left( \pi_E(a|s) \frac{d_{\pi_E}(s)}{d_{\pi_G}(s)} \right) \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}$$

Using  $\tilde{f}^*$ , the optimal generator policy  $\pi_G^*$  is

$$\arg \min_{\pi_G} \max_{\tilde{f}} L(\tilde{f}, \pi_G) = \arg \min_{\pi_G} L(\tilde{f}^*, \pi_G) = \pi_E(a|s) = \frac{\exp \tilde{f}^*(s, a)}{\sum_{a'} \exp \tilde{f}^*(s, a')} \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}.$$

*Proof.* The beginning of the proof closely follows the proof of Appendix A.1.

$$\begin{aligned} \arg \max_{\tilde{f}} L(\tilde{f}, \pi_G) = \\ \arg \max_{\tilde{f}} \sum_{s_i, a_i} d_{\pi_E}(s_i, a_i) \log D_{\tilde{f}, \pi_G}(s_i, a_i) + d_{\pi_G}(s_i, a_i) \log(1 - D_{\tilde{f}, \pi_G}(s_i, a_i)) \end{aligned} \quad (31)$$

We solve for each individual  $(s_i, a_i)$  pair and note that  $\tilde{f}_i \mapsto D_i = \frac{\exp \tilde{f}_i}{\exp \tilde{f}_i + \pi_{G,i}}$  is strictly monotonic on  $\tilde{f}_i \in \mathbb{R} \forall \pi_{G,i} \in ]0, 1[$  so,

$$D_i^* = \arg \max_{D_i} L_i \Leftrightarrow \tilde{f}_i^* = \arg \max_{\tilde{f}_i} L_i \quad (32)$$

Taking the derivative and setting it to 0, we find that

$$\left. \frac{dL_i}{d\tilde{f}_i} \right|_{\tilde{f}_i} = 0 \quad \Leftrightarrow \quad \tilde{f}_i = \log \left( \pi_{G,i} \frac{d_{\pi_E,i}}{d_{\pi_G,i}} \right) \quad (33)$$

We confirm that we have a global maximum with the second derivative test and the values at the border of the domain i.e.  $\left. \frac{d^2 L_i}{d\tilde{f}_i^2} \right|_{\tilde{f}_i^*} < 0$  and  $L_i$  goes to  $-\infty$  for  $\tilde{f}_i \rightarrow +\infty$  and for  $\tilde{f}_i \rightarrow -\infty$ .

It follows that

$$\tilde{f}^*(s, a) = \log \left( \pi_G(a|s) \frac{d_{\pi_E}(s, a)}{d_{\pi_G}(s, a)} \right) \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A} \quad (34)$$

$$\implies \tilde{f}^*(s, a) = \log \left( \frac{\pi_G(a|s) d_{\pi_E}(s) \pi_E(a|s)}{d_{\pi_G}(s) \pi_G(a|s)} \right) \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A} \quad (35)$$

$$\implies \tilde{f}^*(s, a) = \log \left( \pi_E(a|s) \frac{d_{\pi_E}(s)}{d_{\pi_G}(s)} \right) \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A} \quad (36)$$

This proves the first part of Theorem 2.

To prove the second part notice that

$$\begin{aligned} D_{\tilde{f}^*, \pi_G}(s, a) &= \frac{\pi_E(a|s) \frac{d_{\pi_E}(s)}{d_{\pi_G}(s)}}{\pi_E(a|s) \frac{d_{\pi_E}(s)}{d_{\pi_G}(s)} + \pi_G(a|s)} \\ &= \frac{\pi_E(a|s) d_{\pi_E}(s)}{\pi_E(a|s) d_{\pi_E}(s) + \pi_G(a|s) d_{\pi_G}(s)} \\ &= \frac{d_{\pi_E}(s, a)}{d_{\pi_E}(s, a) + d_{\pi_G}(s, a)} \end{aligned} \quad (37)$$

This is equal to the optimal discriminator of the GAN objective Eq. (16) when  $x = (s, a)$ . For this discriminator we showed in Section A.1 that the optimal generator  $\pi_G^*$  is such that  $d_{\pi_G^*}(s, a) = d_{\pi_E}(s, a) \forall (s, a) \in \mathcal{S} \times \mathcal{A}$ , which is satisfied for  $\pi_G^*(a|s) = \pi_E(a|s) \forall (s, a) \in \mathcal{S} \times \mathcal{A}$ . Using the fact that

$$\sum_{a'} \exp \tilde{f}^*(s, a') = \sum_{a'} \pi_E(a'|s) \frac{d_{\pi_E}(s)}{d_{\pi_G}(s)} = \frac{d_{\pi_E}(s)}{d_{\pi_G}(s)} \sum_{a'} \pi_E(a'|s) = \frac{d_{\pi_E}(s)}{d_{\pi_G}(s)}. \quad (38)$$

we can combine Eq. (36) and Eq. (38) to write the expert's policy  $\pi_E$  as a function of the optimal discriminator parameter  $\tilde{f}^*$ :

$$\pi_E(a|s) = \frac{\exp \tilde{f}^*(s, a)}{\sum_{a'} \exp \tilde{f}^*(s, a')} \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}. \quad (39)$$

This concludes the second part of the proof.  $\square$

**Adversarial Soft-Q Fitting (ASQF) - practical algorithm** In a nutshell, Theorem 2 tells us that training the discriminator in Eq. (30) to distinguish between transitions from the expert and transitions from a generator policy can be seen as retrieving  $\tilde{f}^*$  which plays the role of the expert's soft Q-function (i.e. which matches Eq. (1) for  $\tilde{f}^* = \frac{1}{\alpha} Q_{\text{soft}, E}^*$ ):

$$\pi_E(a|s) = \frac{\exp \tilde{f}^*(s, a)}{\sum_{a'} \exp \tilde{f}^*(s, a')} = \exp \left( \tilde{f}^*(s, a) - \log \sum_{a'} \exp \tilde{f}^*(s, a') \right), \quad (40)$$

Therefore, by training the discriminator, one simultaneously retrieves the optimal generator policy.

There is one caveat though: the summation over actions that is required in Eq. (40) to go from  $\tilde{f}^*$  to the policy is intractable in continuous action spaces and would require an additional step such as a projection to a proper distribution (Haarnoja et al. [11] use a Gaussian) in order to draw samples and evaluate likelihoods. Updating in this way the generator policy to match a softmax over our learned state-action preferences ( $\tilde{f}^*$ ) becomes very similar in requirements and computational load to a policy optimization step, thus defeating the purpose of this work which is to get rid of the policy optimization step. For this reason we only consider ASQF for discrete action spaces.

As explained in Section 3.3, in practice we optimize  $D_{\tilde{f}, \pi_G}$  only for a few steps before updating  $\pi_G$  by normalizing  $\exp \tilde{f}(s, a)$  over the action dimension. See Algorithm 2 for the pseudo-code.

---

**Algorithm 2:** Adversarial Soft-Q Fitting (ASQF)

---

**Require:** expert transitions  $\mathcal{D}_E = \{(s_i, a_i)\}_{i=1}^{N_E}$   
 Randomly initialize  $\tilde{f}$  and get  $\pi_G$  from Eq. (40)  
**for** steps  $m = 0$  to  $M$  **do**  
   Collect transitions  $\mathcal{D}_G = \{(s_i, a_i)\}_{i=1}^{N_G}$  by executing  $\pi_G$   
   Train  $D_{\tilde{f}, \pi_G}$  using binary cross-entropy on minibatches of transitions from  $\mathcal{D}_E$  and  $\mathcal{D}_G$   
   Get  $\pi_G$  from Eq. (40)  
**end for**

---

**Experimental results** Figure 4 shows that ASQF performs well on small scale environments but struggles and eventually fails on more complicated environments. Specifically, it seems that ASQF does not scale well with the observation space size. Indeed mountaincar, cartpole, lunarlander and pommerman have respectively an observation space dimensionality of 2, 4, 8 and 960. This may be due to the fact that the partition function Eq. (38) becomes more difficult to learn. Indeed, for each state, several transitions with different actions are required in order to learn it. Poorly approximating this partition function could lead to assigning too low a probability to expert-like actions and eventually failing to behave appropriately. ASAF on the other hand explicitly learns the probability of an action given the state – in other word it explicitly learns the partition function – and is therefore immune to that problem.

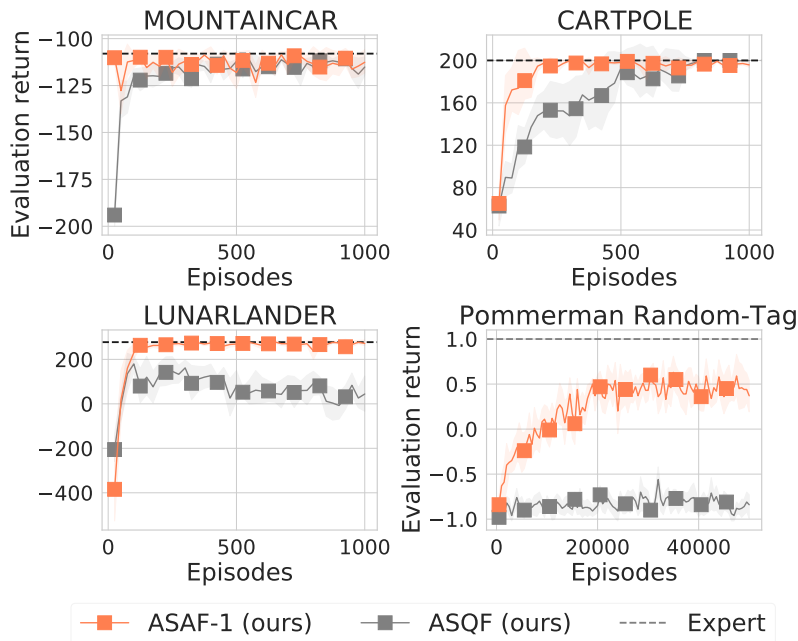


Figure 4: Comparison between ASAF-1 and ASQF, our two transition-wise methods, on environments with increasing observation space dimensionality



## C Additional Experiments

### C.1 GAIL - Importance of Gradient Penalty

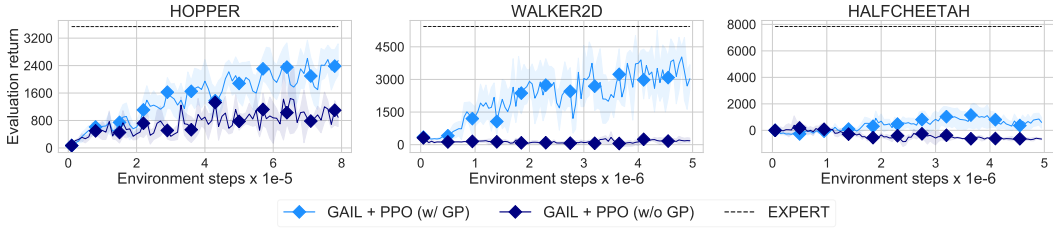


Figure 5: Comparison between original GAIL [13] and GAIL with gradient penalty (GP) [9, 14]

### C.2 Mimicking the expert

To ensure that our method actually mimics the expert and doesn't just learn a policy that collects high rewards when trained with expert demonstrations, we ran ASAF-1 on the Ant-v2 MuJoCo environment using various sets of 25 demonstrations. These demonstrations were generated from a Soft Actor-Critic agent at various levels of performance during its training. Since at low-levels of performance the variance of episode's return is high, we filtered collected demonstrations to lie in the targeted range of performance (e.g. return in [800, 1200] for the 1K set). Results in Figure 6 show that our algorithm succeeds at learning a policy that closely emulates various demonstrators (even when non-optimal).

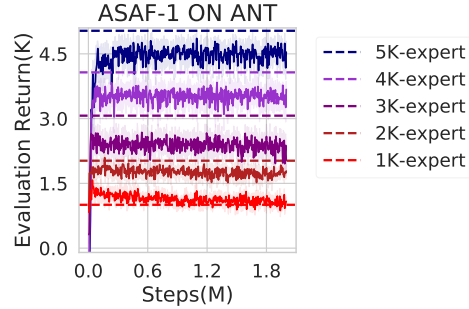


Figure 6: ASAF-1 on Ant-v2. Colors are 1K, 2K, 3K, 4K, 5K expert's performance.

### C.3 Wall Clock Time

We report training times in Figure 7 and observe that ASAF-1 is always fastest to learn. Note however that reports of performance w.r.t wall-clock time should always be taken with a grain of salt as they are greatly influenced by hyper-parameters and implementation details.

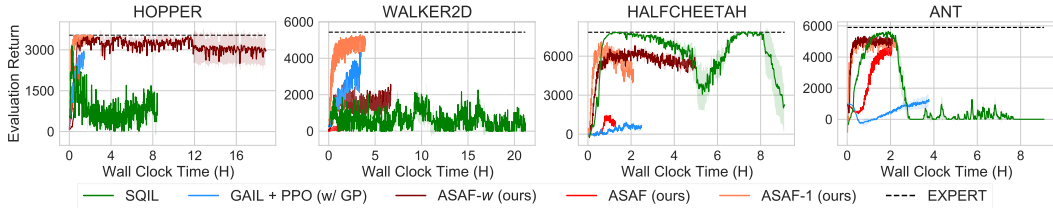


Figure 7: Training times on MuJoCo tasks for 25 expert demonstrations.

## D Hyperparameter tuning and best configurations

### D.1 Classic Control

For this first set of experiments, we use the fixed hyperparameters presented in Table 1.

Table 1: Fixed Hyperparameters for classic control tasks

RL COMPONENT		
HYPER-PARAMETER	DISCRETE CONTROL	CONTINUOUS CONTROL
<b>SAC</b>		
BATCH SIZE (IN TRANSITIONS)	256	256
REPLAY BUFFER LENGTH $ \mathcal{B} $	$10^6$	$10^6$
WARMUP (IN TRANSITIONS)	1280	10240
INITIAL ENTROPY WEIGHT $\alpha$	0.4	0.4
GRADIENT NORM CLIPPING THRESHOLD	0.2	1
TRANSITIONS BETWEEN UPDATE	40	1
TARGET NETWORK WEIGHT $\tau$	0.01	0.01
<b>PPO</b>		
BATCH SIZE (IN TRANSITIONS)	256	256
GAE PARAMETER $\lambda$	0.95	0.95
TRANSITIONS BETWEEN UPDATE	-	2000
EPISODES BETWEEN UPDATES	10	-
EPOCHS PER UPDATE	10	10
UPDATE CLIPPING PARAMETER	0.2	0.2
REWARD LEARNING COMPONENT		
HYPER-PARAMETER	DISCRETE CONTROL	CONTINUOUS CONTROL
<b>AIRL, GAIL, ASAF-1</b>		
BATCH SIZE (IN TRANSITIONS)	256	256
TRANSITIONS BETWEEN UPDATE	-	2000
EPISODES BETWEEN UPDATES	10	-
EPOCHS PER UPDATE	50	50
GRADIENT VALUE CLIPPING THRESHOLD (ASAF-1)	-	1
<b>ASAF, ASAF-<math>w</math></b>		
BATCH SIZE (IN TRAJECTORIES)	10	10
EPISODES BETWEEN UPDATES	10	20
EPOCHS PER UPDATE	50	50
WINDOW SIZE $w$	(SEARCHED)	200
GRADIENT VALUE CLIPPING THRESHOLD	-	1

For the most sensitive hyperparameters, the learning rates for the reinforcement learning and discriminator updates ( $\epsilon_{\text{RL}}$  and  $\epsilon_{\text{D}}$ ), we perform a random search over 50 configurations and 3 seeds each (for each algorithm on each task) for 500 episodes. We consider logarithmic ranges, i.e.  $\epsilon = 10^u$  with  $u \sim \text{Uniform}(-6, -1)$  for  $\epsilon_{\text{D}}$  and  $u \sim \text{Uniform}(-4, -1)$  for  $\epsilon_{\text{RL}}$ . We also include in this search the critic learning rate coefficient  $\kappa$  for PPO also sampled according to a logarithmic scale with  $u \sim \text{Uniform}(-2, 2)$  so that the effective learning rate for PPO’s critic network is  $\kappa \cdot \epsilon_{\text{RL}}$ . For discrete action tasks, the window-size  $w$  for ASAF- $w$  is sampled uniformly within  $\{32, 64, 128\}$ . The best configuration for each algorithm is presented in Tables 2 to 7. Figure 1 uses these configurations retrained on 10 seeds and twice as long.

Finally for all neural networks (policies and discriminators) for these experiments we use a fully-connected MLP with two hidden layers and ReLU activation (except for the last layer). We used hidden sizes of 64 for the discrete tasks and of 256 for the continuous tasks.

Table 2: Best found hyper-parameters for Cartpole

HYPER-PARAMETER	ASAF	ASAF- $w$	ASAF-1	SQIL	AIRL + PPO	GAIL + PPO
DISCRIMINATOR UPDATE LR $\epsilon_D$	0.028	0.039	0.00046	-	$2.5 \cdot 10^{-6}$	0.00036
RL UPDATE LR $\epsilon_{RL}$	-	-	-	0.0067	0.0052	0.012
CRITIC LR COEFFICIENT $\kappa$	-	-	-	-	0.25	0.29
WINDOW SIZE $w$	-	64	1	-	-	-
WINDOW STRIDE	-	64	1	-	-	-

Table 3: Best found hyper-parameters for Mountaincar

HYPER-PARAMETER	ASAF	ASAF- $w$	ASAF-1	SQIL	AIRL + PPO	GAIL + PPO
DISCRIMINATOR UPDATE LR $\epsilon_D$	0.059	0.059	0.0088	-	0.0042	0.00016
RL UPDATE LR $\epsilon_{RL}$	-	-	-	0.062	0.016	0.0022
CRITIC LR COEFFICIENT $\kappa$	-	-	-	-	4.6	0.018
WINDOW SIZE $w$	-	32	1	-	-	-
WINDOW STRIDE	-	32	1	-	-	-

Table 4: Best found hyper-parameters for Lunarlander

HYPER-PARAMETER	ASAF	ASAF- $w$	ASAF-1	SQIL	AIRL + PPO	GAIL + PPO
DISCRIMINATOR UPDATE LR $\epsilon_D$	0.0055	0.0015	0.00045	-	0.0002	0.00019
RL UPDATE LR $\epsilon_{RL}$	-	-	-	0.0036	0.0012	0.0016
CRITIC LR COEFFICIENT $\kappa$	-	-	-	-	0.48	8.5
WINDOW SIZE $w$	-	32	1	-	-	-
WINDOW STRIDE	-	32	1	-	-	-

Table 5: Best found hyper-parameters for Pendulum

HYPER-PARAMETER	ASAF	ASAF- $w$	ASAF-1	SQIL	AIRL + PPO	GAIL + PPO
DISCRIMINATOR UPDATE LR $\epsilon_D$	0.00069	0.00082	0.00046	-	$4.3 \cdot 10^{-6}$	$1.6 \cdot 10^{-5}$
RL UPDATE LR $\epsilon_{RL}$	-	-	-	0.0001	0.00038	0.00028
CRITIC LR COEFFICIENT $\kappa$	-	-	-	-	0.028	84
WINDOW SIZE $w$	-	200	1	-	-	-
WINDOW STRIDE	-	200	1	-	-	-

Table 6: Best found hyper-parameters for Mountaincar-c

HYPER-PARAMETER	ASAF	ASAF- $w$	ASAF-1	SQIL	AIRL + PPO	GAIL + PPO
DISCRIMINATOR UPDATE LR $\epsilon_D$	0.00021	$3.8 \cdot 10^{-5}$	$6.2 \cdot 10^{-6}$	-	$1.7 \cdot 10^{-5}$	$1.5 \cdot 10^{-5}$
RL UPDATE LR $\epsilon_{RL}$	-	-	-	0.0079	0.0012	0.0052
CRITIC LR COEFFICIENT $\kappa$	-	-	-	-	10	12
WINDOW SIZE $w$	-	200	1	-	-	-
WINDOW STRIDE	-	200	1	-	-	-

Table 7: Best found hyper-parameters for Lunarlander-c

HYPER-PARAMETER	ASAF	ASAF- $w$	ASAF-1	SQIL	AIRL + PPO	GAIL + PPO
DISCRIMINATOR UPDATE LR $\epsilon_D$	0.0051	0.0022	0.0003	-	0.0045	0.00014
RL UPDATE LR $\epsilon_{RL}$	-	-	-	0.0027	0.00031	0.00049
CRITIC LR COEFFICIENT $\kappa$	-	-	-	-	14	0.01
WINDOW SIZE $w$	-	200	-	-	-	-
WINDOW STRIDE	-	200	-	-	-	-

## D.2 MuJoCo

For MuJoCo experiments (Hopper-v2, Walker2d-v2, HalfCheetah-v2, Ant-v2), the fixed hyperparameters are presented in Table 8. For all experiments, fully-connected MLPs with two hidden layers and ReLU activation (except for the last layer) were used, where the number of hidden units is equal to 256.

Table 8: Fixed hyperparameters for MuJoCo environments.

<b>RL COMPONENT</b>	
HYPER-PARAMETER	HOPPER, WALKER2D, HALFCHEETAH, ANT
<b>PPO (FOR GAIL)</b>	
GAE PARAMETER $\lambda$	0.98
TRANSITIONS BETWEEN UPDATES	2000
EPOCHS PER UPDATE	5
UPDATE CLIPPING PARAMETER	0.2
CRITIC LR COEFFICIENT $\kappa$	0.25
DISCOUNT FACTOR $\gamma$	0.99
<b>REWARD LEARNING COMPONENT</b>	
HYPER-PARAMETER	HOPPER, WALKER2D, HALFCHEETAH, ANT
<b>GAIL</b>	
TRANSITIONS BETWEEN UPDATES	2000
<b>ASAF</b>	
EPISODES BETWEEN UPDATES	25
<b>ASAF-1 AND ASAF-w</b>	
TRANSITIONS BETWEEN UPDATES	2000

For SQIL we used SAC with the same hyperparameters that were used to generate the expert demonstrations. For ASAF, ASAF-1 and ASAF-w, we set the learning rate for the discriminator at 0.001 and ran random searches over 25 randomly sampled configurations and 2 seeds for each task to select the other hyperparameters for the discriminator training. These hyperparameters included the discriminator batch size sampled from a uniform distribution over  $\{10, 20, 30\}$  for ASAF and ASAF-w (in trajectories) and over  $\{100, 500, 1000, 2000\}$  for ASAF-1 (in transitions), the number of epochs per update sampled from a uniform distribution over  $\{10, 20, 50\}$ , the gradient norm clipping threshold sampled from a uniform distribution over  $\{1, 10\}$ , the window-size (for ASAF-w) sampled from a uniform distribution over  $\{100, 200, 500, 1000\}$  and the window stride (for ASAF-w) sampled from a uniform distribution over  $\{1, 50, w\}$ . For GAIL, we obtained poor results using the original hyperparameters from [13] for a number of tasks so we ran random searches over 100 randomly sampled configurations for each task and 2 seeds to select for the following hyperparameters: the log learning rate of the RL update and the discriminator update separately sampled from uniform distributions over  $[-7, -1]$ , the gradient norm clipping for the RL update and the discriminator update separately sampled from uniform distributions over  $\{None, 1, 10\}$ , the number of epochs per update sampled from a uniform distribution over  $\{5, 10, 30, 50\}$ , the gradient penalty coefficient sampled from a uniform distribution over  $\{1, 10\}$  and the batch size for the RL update and discriminator update separately sampled from uniform distributions over  $\{100, 200, 500, 1000, 2000\}$ .

Table 9: Best found hyper-parameters for the Hopper-v2 environment

HYPER-PARAMETER	ASAF	ASAF- $w$	ASAF-1	SQIL	GAIL + PPO
RL BATCH SIZE (IN TRANSITIONS)	-	-	-	256	200
DISCRIMINATOR BATCH SIZE (IN TRANSITIONS)	-	-	100	-	2000
DISCRIMINATOR BATCH SIZE (IN TRAJECTORIES)	10	10	-	-	-
GRADIENT CLIPPING (RL UPDATE)	-	-	-	-	1.
GRADIENT CLIPPING (DISCRIMINATOR UPDATE)	10.	10.	1.	-	1.
EPOCHS PER UPDATE	50	50	30	-	5
GRADIENT PENALTY (DISCRIMINATOR UPDATE)	-	-	-	-	1.
RL UPDATE LR $\epsilon_{RL}$	-	-	-	$3 * 10^{-4}$	$1.8 * 10^{-5}$
DISCRIMINATOR UPDATE LR $\epsilon_D$	0.001	0.001	0.001	-	0.011
WINDOW SIZE $w$	-	200	1	-	-
WINDOW STRIDE	-	1	1	-	-

Table 10: Best found hyper-parameters for the HalfCheetah-v2 environment

HYPER-PARAMETER	ASAF	ASAF- $w$	ASAF-1	SQIL	GAIL + PPO
RL BATCH SIZE (IN TRANSITIONS)	-	-	-	256	1000
DISCRIMINATOR BATCH SIZE (IN TRANSITIONS)	-	-	100	-	100
DISCRIMINATOR BATCH SIZE (IN TRAJECTORIES)	10	10	-	-	-
GRADIENT CLIPPING (RL UPDATE)	-	-	-	-	-
GRADIENT CLIPPING (DISCRIMINATOR UPDATE)	10.	1	1	-	10
EPOCHS PER UPDATE	50	10	30	-	30
GRADIENT PENALTY (DISCRIMINATOR UPDATE)	-	-	-	-	1.
RL UPDATE LR $\epsilon_{RL}$	-	-	-	$3 * 10^{-4}$	0.0006
DISCRIMINATOR UPDATE LR $\epsilon_D$	0.001	0.001	0.001	-	0.023
WINDOW SIZE $w$	-	200	1	-	-
WINDOW STRIDE	-	1	1	-	-

Table 11: Best found hyper-parameters for the Walker2d-v2 environment

HYPER-PARAMETER	ASAF	ASAF- $w$	ASAF-1	SQIL	GAIL + PPO
RL BATCH SIZE (IN TRANSITIONS)	-	-	-	256	200
DISCRIMINATOR BATCH SIZE (IN TRANSITIONS)	-	-	500	-	2000
DISCRIMINATOR BATCH SIZE (IN TRAJECTORIES)	20	20	-	-	-
GRADIENT CLIPPING (RL UPDATE)	-	-	-	-	-
GRADIENT CLIPPING (DISCRIMINATOR UPDATE)	10.	1.	10.	-	-
EPOCHS PER UPDATE	30	10	50	-	30
GRADIENT PENALTY (DISCRIMINATOR UPDATE)	-	-	-	-	1.
RL UPDATE LR $\epsilon_{RL}$	-	-	-	$3 * 10^{-4}$	0.00039
DISCRIMINATOR UPDATE LR $\epsilon_D$	0.001	0.001	0.001	-	0.00066
WINDOW SIZE $w$	-	100	1	-	-
WINDOW STRIDE	-	1	1	-	-

Table 12: Best found hyper-parameters for the Ant-v2 environment

HYPER-PARAMETER	ASAF	ASAF- $w$	ASAF-1	SQIL	GAIL + PPO
RL BATCH SIZE (IN TRANSITIONS)	-	-	-	256	500
DISCRIMINATOR BATCH SIZE (IN TRANSITIONS)	-	-	100	-	100
DISCRIMINATOR BATCH SIZE (IN TRAJECTORIES)	20	20	-	-	-
GRADIENT CLIPPING (RL UPDATE)	-	-	-	-	-
GRADIENT CLIPPING (DISCRIMINATOR UPDATE)	10.	1.	1.	-	10.
EPOCHS PER UPDATE	50	50	10	-	50
GRADIENT PENALTY (DISCRIMINATOR UPDATE)	-	-	-	-	10
RL UPDATE LR $\epsilon_{RL}$	-	-	-	$3 * 10^{-4}$	$8.5 * 10^{-5}$
DISCRIMINATOR UPDATE LR $\epsilon_D$	0.001	0.001	0.001	-	0.0016
WINDOW SIZE $w$	-	200	1	-	-
WINDOW STRIDE	-	50	1	-	-

### D.3 Pommerman

For this set of experiments, we use a number of fixed hyperparameters for all algorithms either inspired from their original papers for the baselines or selected through preliminary searches. These fixed hyperparameters are presented in Table 13.

Table 13: Fixed Hyperparameters for Pommerman Random-Tag environment.

RL COMPONENT	
HYPER-PARAMETER	POMMERMAN RANDOM-TAG
<b>SAC</b>	
BATCH SIZE (IN TRANSITIONS)	256
REPLAY BUFFER LENGTH $ \mathcal{B} $	$10^5$
WARMUP (IN TRANSITIONS)	1280
INITIAL ENTROPY WEIGHT $\alpha$	0.4
GRADIENT NORM CLIPPING THRESHOLD	0.2
TRANSITIONS BETWEEN UPDATE	10
TARGET NETWORK WEIGHT $\tau$	0.05
<b>PPO</b>	
BATCH SIZE (IN TRANSITIONS)	256
GAE PARAMETER $\lambda$	0.95
EPISODES BETWEEN UPDATES	10
EPOCHS PER UPDATE	10
UPDATE CLIPPING PARAMETER	0.2
CRITIC LR COEFFICIENT $\kappa$	0.5
REWARD LEARNING COMPONENT	
HYPER-PARAMETER	POMMERMAN RANDOM-TAG
<b>AIRL, GAIL, ASAF-1</b>	
BATCH SIZE (IN TRANSITIONS)	256
EPISODES BETWEEN UPDATES	10
EPOCHS PER UPDATE	10
<b>ASAF, ASAF-<math>w</math></b>	
BATCH SIZE (IN TRAJECTORIES)	5
EPISODES BETWEEN UPDATES	10
EPOCHS PER UPDATE	10

For the most sensitive hyperparameters, the learning rates for the reinforcement learning and discriminator updates ( $\epsilon_{\text{RL}}$  and  $\epsilon_{\text{D}}$ ), we perform a random search over 25 configurations and 2 seeds each for all algorithms. We consider logarithmic ranges, i.e.  $\epsilon = 10^u$  with  $u \sim \text{Uniform}(-7, -3)$  for  $\epsilon_{\text{D}}$  and  $u \sim \text{Uniform}(-4, -1)$  for  $\epsilon_{\text{RL}}$ . We also include in this search the window-size  $w$  for ASAF- $w$ , sampled uniformly within  $\{32, 64, 128\}$ . The best configuration for each algorithm is presented in Table 14. Figure 3 uses these configurations retrained on 10 seeds.

Table 14: Best found hyper-parameters for the Pommerman Random-Tag environment

HYPER-PARAMETER	ASAF	ASAF- $w$	ASAF-1	SQIL	AIRL + PPO	GAIL + PPO	BC
DISCRIMINATOR UPDATE LR $\epsilon_{\text{D}}$	0.0007	0.0002	0.0001	-	$3.1 \cdot 10^{-7}$	$9.3 \cdot 10^{-7}$	0.00022
RL UPDATE LR $\epsilon_{\text{RL}}$	-	-	-	0.00019	0.00017	0.00015	-
WINDOW SIZE $w$	-	32	1	-	-	-	-
WINDOW STRIDE	-	32	1	-	-	-	-

Finally for all neural networks (policies and discriminators) we use the same architecture. Specifically, we first process the feature maps (see Section E.3) using a 3-layers convolutional network with number of hidden feature maps of 16, 32 and 64 respectively. Each one of these layers use a kernel size of 3x3 with stride of 1, no padding and a ReLU activation. This module ends with a fully connected layer of hidden size 64 followed by a ReLU activation. The output vector is then concatenated to the unprocessed additional information vector (see Section E.3) and passed through a final MLP with two hidden layers of size 64 and ReLU activations (except for the last layer).

## E Environments and expert data

### E.1 Classic Control

The environments used here are the reference Gym implementations for classic control<sup>2</sup> and for Box2D<sup>3</sup>. We generated the expert trajectories for mountaincar (both discrete and continuous version) by hand using keyboard inputs. For the other tasks, we trained our SAC implementation to get experts on the discrete action tasks and our PPO implementation to get experts on the continuous action tasks.

### E.2 MuJoCo

The experts were trained using our implementation of SAC [11] the state-of-the-art RL algorithm in MuJoCo continuous control tasks. Our implementation basically refactors the SAC implementation from Rlpyt<sup>4</sup>. We trained SAC agent for 1,000,000 steps for Hopper-v2 and 3,000,000 steps for Walker2d-v2 and HalfCheetah-v2 and Ant-v2. We used the default hyper-parameters from Rlpyt.

### E.3 Pommerman

The observation space that we use for Pommerman domain [21] is composed of a set of 15 feature maps as well as an additional information vector. The feature maps whose dimensions are given by the size of the board (8x8 in the case of 1vs1 tasks) are one-hot across the third dimension and represent which element is present at which location. Specifically, these feature maps identify whether a given location is the current player, an ally, an ennemy, a passage, a wall, a wood, a bomb, a flame, fog, a power-up. Other feature maps contain integers indicating bomb blast strength, bomb life, bomb moving direction and flame life for each location. Finally, the additional information vector contains the time-step, number of ammunition, whether the player can kick and blast strength for the current player. The agent has an action space composed of six actions: do-nothing, up, down, left, right and lay bomb.

For these experiments, we generate the expert demonstrations using Agent47Agent, the open-source champion algorithm of the FFA 2018 competition [28] which uses hardcoded heuristics and Monte-Carlo Tree-Search<sup>5</sup>. While this agent occasionally eliminates itself during a match, we only select trajectories leading to a win as being expert demonstrations.

### E.4 Demonstrations summary

Table 15 provides a summary of the expert data used.

Table 15: Expert demonstrations used for Imitation Learning

TASK-NAME	EXPERT MEAN RETURN	NUMBER OF EXPERT TRAJECTORIES
CARTPOLE	200.0	10
MOUNTAINCAR	-108.0	10
LUNARLANDER	277.5	10
PENDULUM	-158.6	10
MOUNTAINCAR-C	93.92	10
LUNARLANDER-C	266.1	10
HOPPER	3537	25
WALKER2D	5434	25
HALFCHEETAH	7841	25
ANT	5776	25
POMMERMAN RANDOM-TAG	1	300, 150, 75, 15, 5, 1

<sup>2</sup>See: [http://gym.openai.com/envs/#classic\\_control](http://gym.openai.com/envs/#classic_control)

<sup>3</sup>See: <http://gym.openai.com/envs/#box2d>

<sup>4</sup>See: <https://github.com/astooke/rlpyt>

<sup>5</sup>See: <https://github.com/YichenGong/Agent47Agent/tree/master/pommerman>

1 We would like to thank the reviewers for their thorough evaluations and for bringing to our attention some missing  
 2 citations and typos, these will be corrected in the updated manuscript. We answer here specific questions raised by the  
 3 reviewers and present requested additional experiments.

4 **Additional MuJoCo experiments (R1, R2, R4).** We focus on GAIL because AIRL claims to perform on par with GAIL  
 5 on MuJoCo. We present in Figure 2 below additional results on the MuJoCo envs (as well as the additional Ant env)  
 6 where GAIL has been re-tuned for further improvement and SQIL and ASAF-1 have been added. We see that even with  
 7 careful tuning GAIL is outperformed by our method and that SQIL’s instability is exacerbated on MuJoCo.

8 We also ran ASAF-1 on the Ant-v2 MuJoCo environment using various sets  
 9 of 25 demonstrations (*as requested by R2*). These demonstrations were  
 10 generated from a Soft Actor-Critic agent at various levels of performance  
 11 during its training. Since at low-levels of performance the variance of  
 12 episode return is high, we filtered collected demonstrations to lie in the  
 13 targeted range of performance (e.g. return in [800, 1200] for the 1K set).  
 14 Results in Figure 1 show that our algorithm succeeds at learning a policy  
 15 that closely emulates various demonstrators (even when non-optimal).

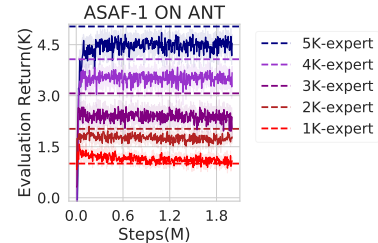


Figure 1: ASAF-1 on Ant-v2. Colors are 1K, 2K, 3K, 4K, 5K expert’s performance.

16 **Comparison of our methods vs Finn et al. [1] and Fu et al. [2] and**  
 17 **Behavioral Cloning (BC) (R1, R3).** Our work builds on [1], yet its novelty  
 18 is to explicitly express the probability of a trajectory in terms of the policy in  
 19 order to directly learn this latter when training the discriminator. In contrast, [2] considers a transition-wise discriminator  
 20 with un-normalized probabilities which makes it closer to ASQF (Appendix B) than to ASAF-1. Additionally, AIRL  
 21 [2] minimizes the Kullback-Leiber Divergence [3] between occupancy measures whereas ASAF minimizes the Jensen-  
 22 Shannon Divergence between trajectories likelihood. Finally, BC uses the loss function from supervised learning  
 23 (classification or regression) to match expert’s actions given expert’s states and suffers from compounding error due  
 24 to co-variate shift [4] since it only learns on the expert state-action visitations (demonstrations) without environment  
 25 interaction. Contrarily, ASAF-1 uses the binary cross entropy loss in Eq. (13) and does not suffer from compounding  
 26 error as it learns on both generated and expert’s trajectories.

27 **Can windowed approach be used for GAIL and AIRL? (R1).** GAIL’s and AIRL’s discriminators are updated based on  
 28 transitions so that a step-wise reward can be learned and used in RL loops. Therefore, the windowed approach, which  
 29 suits trajectory-wise formulations, could not be applied without major modification of their formulations.

30 **Reward Acquisition from ASAF (R3).** Although ASAF does not explicitly acquire reward during training, we can  
 31 retrieve step-wise soft advantages  $\log \pi(a|s)$  from learned agent’s policy  $\pi(a|s)$  which can be used as a reward  
 32 function [2, 5].

33 **Training Time (R4).** Due to lack of room we cannot add here the equivalent of Figure 2 with wall clock time as x-axis  
 34 but we will add it to the updated manuscript. Our observation is that ASAF-1 is always fastest to learn, e.g., 361.2s  
 35 (ASAF-1), 473.1s (SQIL), 1561.0s (ASAF), 1763.6s (GAIL), 2079.1s (ASAF-w) were taken to reach a performance of  
 36 2000 in Hopper environment. Note however that reports of performance w.r.t wall-clock time should always be taken  
 37 with a grain of salt as they are greatly influenced by hyper-parameters and implementation details.

38 **No entropy regularization in loss (R3).** We learn in the softmax policy class of Eq. (2) since it contains the expert’s  
 39 policy (See Section 3.2) but optimize Eq. (13) without entropy regularization as was done in the GAIL paper.

40 **Additional Concerns. (To R3)** We will update our related works with your recommendations. **(To R4)** It seems that our  
 41 supplementary files have been correctly uploaded since R1 was able to read through our code. We are sorry that you  
 42 couldn’t access it. Unfortunately, we are not allowed to put an external link in this rebuttal. We suggest that you contact  
 43 the AC about this issue.

44 **References.** [1] Finn et al., "A connection between generative ...," (2016) [2] Fu et al., "Learning robust rewards ...,"  
 45 (2017) [3] Ghasemipour et al., "A divergence minimization perspective ...," (2019) [4] Ross and Bagnell, "Efficient  
 reductions for imitation learning," (2010) [5] Schulman et al., "High-dimensional continuous control ...," (2015)

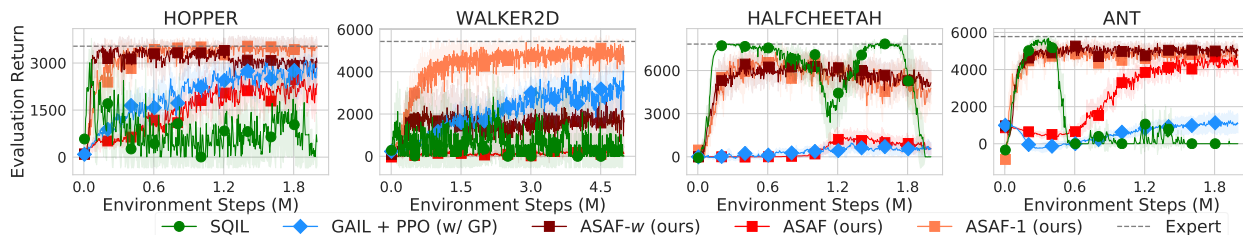


Figure 2: Results on Mujoco environments with improved tuning for GAIL, added SQIL and ASAF-1 and env (Ant)