

Imitation learning with Sinkhorn Distances

Georgios Papagiannis^{1,*}✉ and Yunpeng Li²

¹Imperial College London

²University of Surrey

g.papagiannis21@imperial.ac.uk, yunpeng.li@surrey.ac.uk

Abstract. Imitation learning algorithms have been interpreted as variants of divergence minimization problems. The ability to compare occupancy measures between experts and learners is crucial in their effectiveness in learning from demonstrations. In this paper, we present tractable solutions by formulating imitation learning as minimization of the Sinkhorn distance between occupancy measures. The formulation combines the valuable properties of optimal transport metrics in comparing non-overlapping distributions with a cosine distance cost defined in an adversarially learned feature space. This leads to a highly discriminative critic network and optimal transport plan that subsequently guide imitation learning. We evaluate the proposed approach using both the reward metric and the Sinkhorn distance metric on a number of MuJoCo experiments. For the implementation and reproducing results please refer to the following repository <https://github.com/gpapagiannis/sinkhorn-imitation>.

1 Introduction

Recent developments in reinforcement learning (RL) have allowed agents to achieve state-of-the-art performance on complex tasks from learning to play games [33,38,18] to dexterous manipulation [24], provided with well defined reward functions. However, crafting such a reward function in practical scenarios to encapsulate the desired objective is often non-trivial. Imitation learning (IL) [20] aims to address this issue by formulating the problem of learning behavior through expert demonstration and has shown promises on various application domains including autonomous driving and surgical task automation [2,21,43,23,13].

The main approaches to imitation learning include that of behavioral cloning (BC) and inverse reinforcement learning (IRL). BC mimics the expert’s behavior by converting the task into a supervised regression problem [23,30]. While simple to implement, it is known to suffer from low sample efficiency and poor generalization performance due to covariate shift and high sample correlations in the expert’s trajectory [27,26]. Algorithms such as Dataset Aggregation (DAgger) [26] and Disturbances for Augmenting Robot Trajectories (DART) [17] alleviate

*Work done as a student at the University of Surrey

this issue. However, they require constantly querying an expert for the correct actions.

Inverse reinforcement learning instead aims to recover a reward function which is subsequently used to train the learner’s policy [42,19]. IRL approaches have shown significantly better results [2,5,40,22,1,16] including being sample efficient in terms of expert demonstration. However, IRL itself is an ill-posed problem - multiple reward functions can characterize a specific expert behavior, therefore additional constraints need to be imposed to recover a unique solution [19,42,41]. In addition, the alternating optimization procedure between reward recovery and policy training leads to increased computational cost.

Adversarial imitation learning, on the other hand, bypasses the step of explicit reward inference as in IRL and directly learns a policy that matches that of an expert. Generative adversarial imitation learning (GAIL) [11] minimizes the Jensen-Shannon (JS) divergence between the learner’s and expert’s occupancy measures through a generative adversarial networks (GANs)-based training process. GAIL was developed as a variant of the reward regularized maximum entropy IRL framework [42], where different reward regularizers lead to different IL methods. GAIL has been extended by various other methods aiming to improve its sample efficiency in regard to environment interaction through off-policy RL [31,4,36,15,14]. Recent development [9] provides a unified probabilistic perspective to interpret different imitation learning methods as f -divergence minimization problems and showed that the state-marginal matching objective of IRL approaches is what contributes the most to their superior performance compared to BC. While these methods have shown empirical success, they inherit the same issues from f -divergence and adversarial training, such as training instability in GAN-based training [10] and mode-covering behavior in the JS and Kullback-Leibler (KL) divergences [9,12].

An alternative approach is to utilize optimal transport-based metrics to formulate the imitation learning problem. The optimal transport (OT) theory [37] provides a flexible and powerful tool to compare probability distributions through coupling of distributions based on the metric in the underlying spaces. The Wasserstein adversarial imitation learning (WAIL) [39] was proposed to minimize the dual form of the Wasserstein distance between the learner’s and expert’s occupancy measures, similar to the training of the Wasserstein GAN [3]. The geometric property of the Wasserstein distance leads to numerical stability in training and robustness to disjoint measures. However, the solution to the dual formulation is intractable; approximations are needed in the implementation of neural networks to impose the required Lipschitz condition [28]. [7] introduced Primal Wasserstein imitation learning (PWIL), that uses a reward proxy derived based on an upper bound to the Wasserstein distance between the state-action distributions of the learner and the expert. While PWIL leads to successful imitation, it is unclear how it inherits the theoretical properties of OT, since the transport map between occupancy measures is suboptimal, based on a greedy coupling strategy whose approximation error is difficult to quantify.

In this paper we present Sinkhorn imitation learning (SIL), a *tractable* solution to optimal transport-based imitation learning by leveraging the coupling of occupancy measures and the computational efficiency of the Sinkhorn distance [6], that inherits the theoretical properties of OT. Our main contributions include: (i) We propose and justify an imitation learning training pipeline that minimizes the Sinkhorn distance between occupancy measures of the expert and the learner; (ii) We derive a reward proxy using a set of trainable and highly discriminative optimal transport ground metrics; (iii) We demonstrate through experiments on the MuJoCo simulator [35] that SIL obtains comparable results with the state-of-the-art, outperforming the baselines on a number of experiment settings in regard to both the commonly used reward metric and the Sinkhorn distance.

The rest of this paper is organized as follows. In Section 2 we provide the necessary background for this work. Section 3 introduces the proposed Sinkhorn Imitation Learning (SIL) framework. Section 4 provides details of experiments to evaluate the performance of SIL on a number of MuJoCo environments. We conclude the paper and discuss future research directions in Section 5.

2 Background

2.1 Imitation Learning

Notation. We consider a Markov Decision Process (MDP) which is defined as a tuple $\{\mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma\}$, where \mathcal{S} is a set of states, \mathcal{A} is a set of possible actions an agent can take on the environment, $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is a transition probability matrix, $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is a reward function and $\gamma \in (0, 1)$ is a discount factor. The agent’s behavior is defined by a stochastic policy $\pi : \mathcal{S} \rightarrow \text{Prob}(\mathcal{A})$ and Π is the set of all such policies. We use $\pi_E, \pi \in \Pi$ to refer to the expert and learner policy respectively. The performance measure of policy π is defined as $\mathcal{J} = \mathbb{E}_\pi[r(s, a)] = \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | \mathcal{P}, \pi]$ where $s_t \in \mathcal{S}$ is a state observed by the agent at time step t . With a slight abuse of notations, we also use $r((s, a)_\pi)$ to denote explicitly that $(s, a)_\pi \sim \pi$. τ_E and τ_π denote the set of state-action pairs sampled by an expert and a learner policy respectively during interaction with the environment, also referred to as trajectories. The distribution of state-action pairs generated by policy π through environment interaction, also known as the occupancy measure $\rho_\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, is defined as $\rho_\pi(s, a) = (1 - \gamma)\pi(a|s) \sum_{t=0}^{\infty} \gamma^t P_\pi[s_t = s]$ where $P_\pi[s_t = s]$ denotes the probability of a state being s at time step t following policy π .

Generative Adversarial Imitation Learning. Ho and Ermon [11] extended the framework of MaxEnt IRL by introducing a reward regularizer $\psi(r) : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$:

$$\text{IRL}_\psi(\pi_E) := \arg \max_r -\psi(r) + \min_{\pi \in \Pi} \left(-\mathcal{H}^{\text{causal}}(\pi) - \mathbb{E}_\pi[r(s, a)] \right) + \mathbb{E}_{\pi_E}[r(s, a)] , \quad (1)$$

where $\mathcal{H}^{causal}(\pi) := \mathbb{E}_{\rho_\pi}[-\log \pi(a|s)]/(1 - \gamma)$ [41]. The process of RL following IRL can be formulated as that of occupancy measure matching [11]:

$$\text{RL} \circ \text{IRL}_\psi(\pi_E) := \arg \min_{\pi \in \Pi} -\mathcal{H}^{causal}(\pi) + \psi^*(\rho_\pi - \rho_E), \quad (2)$$

where ψ^* corresponds to the convex conjugate of the reward regularizer $\psi(r)$. The regularized MaxEnt IRL framework bypasses the expensive step of reward inference and learns how to imitate an expert by matching its occupancy measure. Different realizations of the reward regularizer lead to different IL frameworks. A specific choice of the regularizer leads to the Generative Adversarial Imitation Learning (GAIL) framework that minimizes the Jensen-Shannon divergence between the learner’s and expert’s occupancy measures [11].

f-Divergence MaxEnt IRL. Recently, Ghasemipour et al. [9] showed that training a learner policy π to minimize the distance between two occupancy measures can be generalised to minimize any f -divergence between ρ_E and ρ_π denoted as $D_f(\rho_E \parallel \rho_\pi)$. Different choices of f yield different divergence minimization IL algorithms [9] and can be computed as:

$$\max_{T_\omega} \mathbb{E}_{(s,a) \sim \rho_E} [T_\omega(s,a)] - \mathbb{E}_{(s,a) \sim \rho_\pi} [f^*(T_\omega(s,a))], \quad (3)$$

where $T_\omega : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ and f^* is the convex conjugate of the selected f -divergence. The learner’s policy is optimized with respect to the reward proxy $f^*(T_\omega(s,a))$.

2.2 Optimal Transport

While divergence minimization methods have enjoyed empirical success, they are still difficult to evaluate in high dimensions [34], due to the sensitivity to different hyperparameters and difficulty in training depending on the distributions that are evaluated [28]. The optimal transport (OT) theory [37] provides effective methods to compare degenerate distributions by accounting for the underlying metric space. Consider $P_k(\Gamma)$ to be the set of Borel probability measures on a Polish metric space (Γ, d) with finite k -th moment. Given two probability measures $p, q \in P_k(\Gamma)$, the k -Wasserstein metric is defined as [37]:

$$\mathcal{W}_k(p, q)_c = \left(\inf_{\zeta \in \Omega(p, q)} \int_\Gamma c(x, y)^k d\zeta(x, y) \right)^{\frac{1}{k}}, \quad (4)$$

where $\Omega(p, q)$ denotes the set of joint probability distributions whose marginals are p and q , respectively. $c(x, y)$ denotes the cost of transporting sample $x \sim p$ to $y \sim q$. The joint distribution ζ that minimizes the total transportation cost is referred to as the optimal transport plan.

Sinkhorn Distances. The solution to Equation (4) is generally intractable for high dimensional distributions in practice. A regularized form of the optimal transport formulation was proposed by Cuturi [6] that can efficiently compute

the Wasserstein metric. The Sinkhorn distance $\mathcal{W}_s^\beta(p, q)_c$ between p and q is defined as:

$$\mathcal{W}_s^\beta(p, q)_c = \inf_{\zeta_\beta \in \Omega_\beta(p, q)} \mathbb{E}_{x, y \sim \zeta_\beta} [c(x, y)] , \quad (5)$$

where $\Omega_\beta(p, q)$ denotes the set of all joint distributions in $\Omega(p, q)$ with entropy of at least $\mathcal{H}(p) + \mathcal{H}(q) - \beta$ and $\mathcal{H}(\cdot)$ computes the entropy of a distribution. The distance is evaluated on two distributions p and q where in the context of adversarial IL correspond to the state-action distributions of the learner and the expert policies.

3 SIL: Sinkhorn Imitation Learning

We consider the problem of training a learner policy π to imitate an expert, by matching its state-action distribution ρ_E in terms of minimizing their Sinkhorn distance. To facilitate the development of the learning pipeline, we begin by discussing how the Sinkhorn distance is used to evaluate similarity between occupancy measures.

Consider the case of a learner π interacting with an environment and generating a trajectory of state-action pairs $\tau_\pi \sim \pi$ that characterizes its occupancy measure. A trajectory of expert demonstrations $\tau_E \sim \pi_E$ is also available as the expert trajectories. The optimal transport plan ζ_β between the samples of τ_π and τ_E can be obtained via the Sinkhorn algorithm [6]. Following Equation (5) we can evaluate the Sinkhorn distance of τ_π and τ_E as follows:

$$\mathcal{W}_s^\beta(\tau_\pi, \tau_E)_c = \sum_{(s, a)_\pi \in \tau_\pi} \sum_{(s, a)_{\pi_E} \in \tau_E} c((s, a)_\pi, (s, a)_{\pi_E}) \zeta_\beta((s, a)_\pi, (s, a)_{\pi_E}) . \quad (6)$$

Reward Proxy. We now introduce a reward proxy suitable for training a learner policy that minimizes $\mathcal{W}_s^\beta(\tau_\pi, \tau_E)_c$ in order to match the expert's occupancy measure.

The reward function $v_c((s, a)_\pi)$ for each sample $(s, a)_\pi$ in the learner's trajectory is defined as:

$$v_c((s, a)_\pi) := - \sum_{(s, a)_{\pi_E} \in \tau_E} c((s, a)_\pi, (s, a)_{\pi_E}) \zeta_\beta((s, a)_\pi, (s, a)_{\pi_E}) . \quad (7)$$

The optimization objective of the learner policy $\mathcal{J} = \mathbb{E}_\pi[r((s, a)_\pi)]$ under $r((s, a)_\pi) := v_c((s, a)_\pi)$ corresponds to minimizing the Sinkhorn distance between the learner's and expert's trajectories defined in Equation (6). Hence, by maximizing the optimization objective \mathcal{J} with reward $v_c((s, a)_\pi)$, a learner is trained to minimize the Sinkhorn distance between the occupancy measures of the learner and the expert demonstrator.

Adversarial reward proxy. The reward specified in Equation (7) can only be obtained after the learner has generated a complete trajectory. The optimal

transport plan $\zeta_\beta((s, a)_\pi, (s, a)_{\pi_E})$ then weighs the transport cost of each sample $(s, a)_\pi \in \tau_\pi$ according to the samples present in τ_π and τ_E . The dependence of $v_c((s, a)_\pi)$ to all state-action pairs in τ_π and τ_E can potentially result in the same state-action pair being assigned significantly different rewards depending on the trajectory that it is sampled from. Such dependence can lead to difficulty in maximizing the optimization objective \mathcal{J} (and equivalently in minimizing the Sinkhorn distance between the occupancy measures from the learner and the expert). Empirical evidence is provided in the ablation study in Section 4.

In order to provide a discriminative signal to the learner’s policy and aid the optimization process, we consider adversarially training a critic to penalize non-expert state-action pairs by increasing their transport cost to the expert’s distribution, drawing inspiration from the adversarially trained transport ground metric in the OT-GAN framework [29]. The critic $c_w((s, a)_\pi, (s, a)_{\pi_E})$ parameterized by w is defined as follows:

$$c_w((s, a)_\pi, (s, a)_{\pi_E}) = 1 - \frac{f_w((s, a)_\pi) \cdot f_w((s, a)_{\pi_E})}{\|f_w((s, a)_\pi)\|_2 \|f_w((s, a)_{\pi_E})\|_2}, \quad (8)$$

where \cdot denotes the inner product between two vectors. $f_w(\cdot) : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$ maps the environment’s observation space $\mathcal{S} \times \mathcal{A}$ to an adversarially learned feature space \mathbb{R}^d where d is the feature dimension. The adversarial reward proxy $v_{c_w}((s, a)_\pi)$ is obtained by substituting the transport cost $c(\cdot, \cdot)$ in Equation (7) with $c_w(\cdot, \cdot)$ defined by Equation (8). SIL learns π by solving the following minimax optimization problem:

$$\arg \min_{\pi} \max_w \mathcal{W}_s^\beta(\rho_\pi, \rho_E)_{c_w}. \quad (9)$$

Remark 1. For SIL, the adversarial training part of the transport cost is not part of the approximation procedure of the distance metric, as in GAIL [11] and WAIL [39]. The Sinkhorn distance is computed directly via the Sinkhorn iterative procedure [6] with the transport cost defined in Equation (8).

Algorithm. The pseudocode for the proposed Sinkhorn imitation learning (SIL) framework is presented in Algorithm 1. In each iteration we randomly match each of the learner’s generated trajectories to one of the expert’s and obtain their Sinkhorn distance. The reason behind this implementation choice is to maintain a constant computational complexity with respect to a potentially increasing number of demonstrations. We then alternate between one step of updating a critic network c_w to maximize the Sinkhorn distance between the learner’s and expert’s trajectories and a policy update step to minimize the distance between occupancy measures with the learned reward proxy. As SIL depends on complete environment trajectories to compute the Sinkhorn distance, it is inherently an on-policy method. Hence, to train our imitator policy we use Trust Region Policy Optimization (TRPO) [32] for our experiments.

Algorithm 1: Sinkhorn imitation learning (SIL)

- Input:** Set of expert trajectories $\{\tau_E\} \sim \pi_E$, Sinkhorn regularization parameter β , initial learner’s policy parameters θ_0 , initial critic network parameters w_0 , number of training iterations K
- 1: **for** iteration $k = 0$ to K **do**
 - 2: Sample a set of trajectories $\{\tau_{\pi_{\theta_k}}\}_k \sim \pi_{\theta_k}$.
 - 3: Create a set of trajectory pairs $\{(\tau_{\pi_{\theta_k}}, \tau_E)\}_k$ by randomly matching trajectories from the learner’s set to the expert’s.
 - 4: For each pair in $\{(\tau_{\pi_{\theta_k}}, \tau_E)\}_k$, calculate $\mathcal{W}_s^\beta(\tau_{\pi_{\theta_k}}, \tau_E)_{c_w}$ using the Sinkhorn algorithm (Equation (5)) and transport cost as in Equation (8), in order to update the reward proxy $v_{c_{w_k}}((s, a)_{\pi_{\theta_k}})$ for each state action pair.
 - 5: Update w_k to maximize $\mathcal{W}_s^\beta(\tau_{\pi_{\theta_k}}, \tau_E)_{c_w}$ using gradient ascent with the gradient:

$$\nabla_{w_k} \frac{1}{m} \sum_{\{(\tau_{\pi_{\theta_k}}, \tau_E)\}_k} \mathcal{W}_s^\beta(\tau_{\pi_{\theta_k}}, \tau_E)_{c_w} , \quad (10)$$

where m is the number of trajectory pairs.

- 6: Update policy parameter θ_k using TRPO and reward $v_{c_{w_k}}((s, a)_{\pi_{\theta_k}})$ updated in Step 4.
- 7: **end for**

Output: Learned policy π_{θ_k} .

3.1 Connection to regularized MaxEnt IRL.

We now show how SIL can be interpreted as a variant of the regularized MaxEnt IRL framework [11] given a specific choice of $\psi(r)$.

Definition 1. Consider a learner’s policy and expert’s demonstrations, as well as their induced occupancy measures ρ_π and ρ_E . We define the following reward regularizer:

$$\psi_{\mathcal{W}}(r) := -\mathcal{W}_s^\beta(\rho_\pi, \rho_E)_{c_w} + \mathbb{E}_{\rho_\pi}[r(s, a)] - \mathbb{E}_{\rho_E}[r(s, a)] . \quad (11)$$

Proposition 1. The reward regularizer $\psi_{\mathcal{W}}(r)$ defined in Equation (11) leads to an entropy regularized MaxEnt IRL algorithm. When $r((s, a)_\pi) = v_{c_w}((s, a)_\pi)$,

$$RL \circ IRL_{\psi_{\mathcal{W}}}(\pi_E) = \arg \min_{\pi \in \Pi} -\mathcal{H}^{causal}(\pi) + \sup_w \mathcal{W}_s^\beta(\rho_\pi, \rho_E)_{c_w} . \quad (12)$$

Equation (12) corresponds to the process of updating a critic network to maximize the Sinkhorn distance between the learner’s and expert’s occupancy measures, followed by the process of finding a policy π to minimize it. The added term $\mathcal{H}^{causal}(\pi)$ is treated as a regularization parameter.

Proof. Consider the set of possible rewards $\mathcal{R} := \{r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}\}$ in finite state-action space as in [11] and [9]. The joint state-action distributions ρ_π and ρ_E are represented as vectors in $[0, 1]^{\mathcal{S} \times \mathcal{A}}$.

Define $\psi_{\mathcal{W}}(r) := -\mathcal{W}_s^\beta(\rho_\pi, \rho_E) + \mathbb{E}_{\rho_\pi}[r(s, a)] - \mathbb{E}_{\rho_E}[r(s, a)]$, where $\mathcal{W}_s^\beta(\rho_\pi, \rho_E)$ is obtained with the transport cost c_w defined in Equation (8). Given $r(s, a) = v_{c_w}(s, a)$ and recall that the convex conjugate of a function g is $g^*(y) = \sup_{x \in \text{dom}(g)} (y^T x - g(x))$, we obtain

$$\begin{aligned} \psi_{\mathcal{W}}^*(\rho_\pi - \rho_E) &= \sup_{r \in \mathcal{R}} [(\rho_\pi - \rho_E)^T r - \psi_{\mathcal{W}}(r)] = \sup_{r \in \mathcal{R}} \left[\sum_{\mathcal{S} \times \mathcal{A}} (\rho_\pi(s, a) - \rho_E(s, a)) \cdot r(s, a) \right. \\ &\quad \left. + \mathcal{W}_s^\beta(\rho_\pi, \rho_E) - \sum_{\mathcal{S} \times \mathcal{A}} (\rho_\pi(s, a) - \rho_E(s, a)) \cdot r(s, a) \right] = \\ &= \sup_{r \in \mathcal{R}} \mathcal{W}_s^\beta(\rho_\pi, \rho_E) = \sup_{v_{c_w} \in \mathcal{R}} \mathcal{W}_s^\beta(\rho_\pi, \rho_E) = \sup_w \mathcal{W}_s^\beta(\rho_\pi, \rho_E). \end{aligned} \quad (13)$$

From Equation (2),

$$\begin{aligned} \text{RL} \circ \text{IRL}_\psi(\pi_E) &= \arg \min_{\pi \in \Pi} -\mathcal{H}^{\text{causal}}(\pi) + \psi_{\mathcal{W}}^*(\rho_\pi - \rho_E) \\ &= \arg \min_{\pi \in \Pi} -\mathcal{H}^{\text{causal}}(\pi) + \sup_w \mathcal{W}_s^\beta(\rho_\pi, \rho_E). \end{aligned} \quad (14)$$

4 Experiments

To empirically evaluate the Sinkhorn imitation learning (SIL) algorithm, we benchmark SIL against BC in the four MuJoCo [35] environments studied in [9], namely Hopper-v2, Walker2d-v2, Ant-v2 and HalfCheetah-v2, as well as the Humanoid-v2 environment. Given that SIL is an on-policy method due to the requirement of complete trajectories, two on-policy adversarial IL algorithms, namely GAIL [11] and AIRL [8], are also included as baselines. All algorithms are evaluated against the true reward metric obtained through environment interaction, in addition to the Sinkhorn distance between the samples from the learned policy and the expert demonstrations.

Initially we train policies using TRPO [32] to obtain expert performance. The expert policies are used to generate sets of expert demonstrations. The performance of the obtained expert policies can be found in Table 1. To study the robustness of SIL in learning from various lengths of trajectory sets we train the algorithms on sets of $\{2, 4, 8, 16, 32\}$ and for Humanoid-v2 for $\{8, 16, 32\}$ sets. All trajectories are subsampled by a factor of 20 starting from a random offset, a common practice found in [11, 9, 8]. SIL, GAIL and AIRL are trained for 250 iterations allowing approximately 50,000 environment interactions per iteration. For Humanoid-v2 we train the algorithms for 350 iterations. All reported results correspond to performance metrics obtained after testing the learner policies on 50 episodes.

Table 1: Performance of expert policies providing the demonstrations trained using TRPO.

Environments	Expert Performance
Hopper-v2	3354.74 ± 1.87
HalfCheetah-v2	4726.53 ± 133.12
Walker2d-v2	3496.44 ± 8.79
Ant-v2	5063.11 ± 337.50
Humanoid-v2	6303.36 ± 97.71

4.1 Implementation Details

Adversarial Critic. The critic network consists of a 2-layer MLP architecture with 128 units each with ReLU activations. For each experiment we report the best performing result after training the critic with the following learning rates $\{0.0004, 0.0005, 0.0006, 0.0007, 0.0008, 0.0009\}$ and output dimensions $\{5, 10, 30\}$. Although different choices of the critic network output dimension may yield better results for the proposed SIL algorithm in different environments, no further attempt was made to fine-tune the output for the critic. We note that for most experiment settings a critic output dimension of 30 and learning rate of 0.0005 among the pool of candidate values yield the best results.

Reward Proxy. After obtaining the value of v_{c_w} as defined in Equations (7) and (8), we add a value of $\frac{2}{L}$ where L is the trajectory length and scale the reward by 2. By doing so we set the range of v_{c_w} to be $0 \leq v_{c_w} \leq 4$ which proved to be effective for environments requiring a survival bonus. We keep track of a running standard deviation to normalize rewards.

Policy Architecture & Training. For both the expert and learner policies, we use the same architecture comprised of a 2-layer MLP architecture each with 128 units with ReLU activations. The same architecture is used amongst all imitation learning algorithms. For all adversarial IL algorithms, as well as obtaining expert performance, we train the policies using Trust Region Policy Optimization [32]. Finally, we normalize environment observations by keeping track of the running mean and standard deviation.

GAIL & AIRL. To aid the performance of the benchmarks algorithms GAIL and AIRL in the HalfCheetah-v2 environment, we initialize the policies with that from behavioural cloning.

Computational Resource. The experiments were run on a computer with an Intel (R) Xeon (R) Gold 5218 CPU 2.3 GHz and 16GB of RAM, and a RTX 6000 graphic card with 22GB memories.




Fig. 1: Mean and standard deviation of the Sinkhorn distance evaluated during training of SIL using a fixed cosine transport cost by stochastically sampling an action from the learner’s policy.

4.2 Results

Sinkhorn metric. We begin by evaluating performance amongst IL methods using the Sinkhorn metric. Since our goal is to assess how well imitation learning algorithms match the expert’s occupancy measure, the Sinkhorn distance offers a valid metric of similarity between learner’s and expert’s trajectories compared to the reward metric which is also often unavailable in practical scenarios. We report the Sinkhorn distance between occupancy measures computed with a fixed cosine distance-based transport cost during testing and evaluation:

$$c((s, a)_\pi, (s, a)_{\pi_E}) = 1 - \frac{[s, a]_\pi \cdot [s, a]_{\pi_E}}{\|[s, a]_\pi\|_2 \|[s, a]_{\pi_E}\|_2}, \quad (15)$$

where $[s, a]_\pi$ denotes the concatenated vector of state-action of policy π and $\|\cdot\|_2$ computes the L2 norm. Table 2 reports the Sinkhorn metric evaluated between the trajectories generated by the learned policies with the demonstrations provided by the expert. A smaller Sinkhorn distance corresponds to higher similarity between the learner’s and expert’s generated trajectories. SIL, AIRL and GAIL obtain comparable performance in most of the environments. The proposed SIL algorithm outperforms the baselines in almost all experiments on the environments of HalfCheetah-v2 and Ant-v2, while AIRL achieves superior performance on the environments of Hopper-v2 and Walker2d-v2. GAIL on the other hand obtains relatively poor performance with regard to the Sinkhorn distance when provided with only 2 expert trajectories on the environments of Hopper-v2, HalfCheetah-v2 and Ant-v2. As expected, behavioral cloning fails to obtain competitive performance in almost all experiment settings especially when provided with a small number of expert demonstrations.

In addition, SIL outperforms GAIL and AIRL on the Humanoid-v2 environment when provided with 8 and 16 trajectories, where SIL demonstrates significantly improved sample efficiency in terms of both expert demonstrations and environment interactions. GAIL outperforms the rest when trained with 32 trajectories on the Humanoid-v2 environment. Interestingly, BC obtains superior performance with regard to the Sinkhorn distance on the Humanoid-v2 environment when provided with 8 trajectories, but low performance regarding the reward metric as shown in Table 3.

Table 2: Mean and standard deviation of the Sinkhorn distance between the expert demonstrations and samples from imitator policies for BC, GAIL, AIRL and SIL. A fixed cosine transport cost is used only for evaluation (Smaller distance denotes better performance).

Environments	Trajectories	BC	GAIL	AIRL	SIL
Hopper-v2	2	0.467 ± 0.009	0.098 ± 0.003	0.069 ± 0.001	0.073 ± 0.001
	4	0.408 ± 0.080	0.120 ± 0.010	0.066 ± 0.009	0.082 ± 0.010
	8	0.300 ± 0.029	0.074 ± 0.004	0.068 ± 0.006	0.071 ± 0.005
	16	0.182 ± 0.042	0.106 ± 0.008	0.074 ± 0.010	0.078 ± 0.012
	32	0.157 ± 0.084	0.071 ± 0.008	0.072 ± 0.009	0.089 ± 0.008
HalfCheetah-v2	2	1.043 ± 0.058	0.940 ± 0.181	0.577 ± 0.157	0.546 ± 0.138
	4	0.791 ± 0.096	0.633 ± 0.095	0.630 ± 0.091	0.620 ± 0.101
	8	0.841 ± 0.071	0.702 ± 0.095	0.708 ± 0.054	0.700 ± 0.052
	16	0.764 ± 0.166	0.670 ± 0.128	0.671 ± 0.112	0.688 ± 0.131
	32	0.717 ± 0.129	0.695 ± 0.113	0.699 ± 0.091	0.685 ± 0.083
Walker2d-v2	2	0.474 ± 0.023	0.067 ± 0.008	0.034 ± 0.005	0.080 ± 0.004
	4	0.694 ± 0.011	0.067 ± 0.006	0.036 ± 0.002	0.079 ± 0.005
	8	0.335 ± 0.004	0.069 ± 0.005	0.036 ± 0.003	0.063 ± 0.003
	16	0.199 ± 0.013	0.061 ± 0.004	0.037 ± 0.005	0.102 ± 0.007
	32	0.196 ± 0.098	0.052 ± 0.003	0.042 ± 0.004	0.147 ± 0.003
Ant-v2	2	0.843 ± 0.033	0.344 ± 0.068	0.164 ± 0.006	0.158 ± 0.008
	4	0.684 ± 0.159	0.165 ± 0.119	0.163 ± 0.008	0.157 ± 0.014
	8	0.996 ± 0.029	0.159 ± 0.016	0.164 ± 0.019	0.155 ± 0.012
	16	0.724 ± 0.149	0.225 ± 0.106	0.173 ± 0.062	0.165 ± 0.022
	32	0.452 ± 0.094	0.176 ± 0.029	0.172 ± 0.020	0.173 ± 0.018
Humanoid-v2	8	0.336 ± 0.089	0.386 ± 0.011	1.015 ± 0.015	0.379 ± 0.296
	16	0.290 ± 0.086	0.428 ± 0.027	1.034 ± 0.017	0.182 ± 0.011
	32	0.182 ± 0.028	0.162 ± 0.144	1.026 ± 0.015	0.250 ± 0.180

Reward metric. To better understand how performance changes in terms of the Sinkhorn distance metric translates to the true reward, Table 3 shows the reward obtained with the learned policies in the same experiments reported in Table 2. While all adversarial imitation learning algorithms exhibit similar reward values compared to the expert policies, we observe that SIL generally obtains lower reward compared to AIRL on Ant-v2. In addition, AIRL obtains lower reward compared to SIL and GAIL on Walker2d-v2. However, both SIL and AIRL yield superior performance in these environments when evaluated using the Sinkhorn distance as shown in Table 2. The result suggests that evaluating the

Table 3: Mean and standard deviation of the reward metric performance of imitator policies for BC, GAIL, AIRL and SIL.

Environments	Trajectories	BC	GAIL	AIRL	SIL
Hopper-v2	2	391.38 \pm 42.98	3341.27 \pm 38.96	3353.33 \pm 2.05	3376.70 \pm 2.45
	4	659.51 \pm 166.32	3206.85 \pm 1.56	3353.75 \pm 1.67	3325.66 \pm 4.24
	8	1094.39 \pm 145.93	3216.93 \pm 3.08	3369.17 \pm 3.04	3335.31 \pm 2.66
	16	2003.71 \pm 655.85	3380.97 \pm 2.16	3338.07 \pm 2.14	3376.55 \pm 2.65
	32	2330.82 \pm 1013.71	3333.93 \pm 1.47	3361.56 \pm 1.93	3326.52 \pm 3.62
HalfCheetah-v2	2	-60.80 \pm 23.12	764.91 \pm 546.47	4467.83 \pm 61.13	4664.65 \pm 91.73
	4	1018.68 \pm 236.13	5183.67 \pm 118.74	4578.84 \pm 102.92	4505.88 \pm 130.50
	8	1590.73 \pm 279.05	4902.46 \pm 721.43	4686.22 \pm 147.89	4818.82 \pm 251.27
	16	2434.30 \pm 733.29	4519.49 \pm 157.99	4783.79 \pm 197.27	4492.37 \pm 134.35
	32	3598.98 \pm 558.70	4661.17 \pm 147.21	4633.48 \pm 116.89	4795.68 \pm 191.90
Walker2d-v2	2	591.92 \pm 32.77	3509.37 \pm 8.08	3497.80 \pm 9.64	3566.32 \pm 16.11
	4	314.77 \pm 9.21	3537.63 \pm 4.14	3496.61 \pm 10.94	3523.73 \pm 21.91
	8	808.37 \pm 5.28	3394.15 \pm 4.74	3488.68 \pm 10.67	3420.13 \pm 16.38
	16	1281.80 \pm 81.11	3444.96 \pm 23.99	3459.84 \pm 8.25	3557.51 \pm 11.67
	32	1804.74 \pm 1154.36	3427.61 \pm 9.79	3495.04 \pm 17.18	3203.32 \pm 23.65
Ant-v2	2	845.14 \pm 172.37	3443.87 \pm 716.61	5190.89 \pm 67.94	4981.70 \pm 50.89
	4	897.54 \pm 2.14	4912.92 \pm 606.99	5182.42 \pm 65.70	5020.71 \pm 89.74
	8	991.92 \pm 2.92	5112.21 \pm 102.23	5083.30 \pm 77.48	5112.55 \pm 62.87
	16	1014.14 \pm 447.66	4854.87 \pm 895.63	5034.80 \pm 331.64	4935.33 \pm 87.15
	32	2197.20 \pm 487.00	5009.60 \pm 247.43	5013.36 \pm 119.12	4581.27 \pm 123.75
Humanoid-v2	8	1462.47 \pm 1139.19	1249.26 \pm 187.71	3897.47 \pm 1047.03	4456.09 \pm 2707.92
	16	2100.93 \pm 1116.79	496.11 \pm 113.28	4396.01 \pm 433.63	6380.37 \pm 40.35
	32	4807.86 \pm 1903.08	6252.73 \pm 570.72	1884.92 \pm 764.89	5593.19 \pm 1967.86

performance of imitation learning algorithms with a true similarity metric, such as the Sinkhorn distance, can be more reliable since our objective is to match state-action distributions.

Training Stability. Table 2 showcases that SIL consistently minimizes the Sinkhorn distance while being robust to varying lengths of expert demonstrations. Figure 1 depicts the evolution of the Sinkhorn distance between occupancy measures of the learner and the expert in the training process of SIL. In spite of the training instability observed on Walker2d-v2 with 2 or 32 expert trajectories on the Humanoid-v2 environment, SIL still successfully learns to imitate the expert demonstrator. We speculate that training stability could be improved in these settings with further hyperparameter tuning as discussed in Section 5 which we leave for future work. Training stability of SIL is evident on the Hopper-v2,




Fig. 2: *Ablation Study.* Mean and standard deviation of the Sinkhorn distance during training of SIL for three sets of varying number of trajectories. The critic network update has been replaced with a fixed cosine transport cost defined in Equation (15).

Table 4: Mean and standard deviation of the reward and Sinkhorn metric performance after re-training SIL with a *fixed* cosine transport cost defined in Equation (15).

Environments	Metric	2	8	32
Hopper-v2	Reward	264.72 ± 1.28	520.88 ± 29.83	9.44 ± 0.31
	Sinkhorn	0.036 ± 0.007	0.552 ± 0.008	0.777 ± 0.007
HalfCheetah-v2	Reward	-1643.98 ± 198.31	-844.52 ± 267.42	-1220.92 ± 217.86
	Sinkhorn	0.670 ± 0.141	0.841 ± 0.035	0.424 ± 0.031
Walker2d-v2	Reward	60.64 ± 7.92	-2.39 ± 14.05	-11.38 ± 1.22
	Sinkhorn	0.538 ± 0.006	0.487 ± 0.005	0.466 ± 0.009
Ant-v2	Reward	1482.03 ± 480.99	607.87 ± 87.09	114.22 ± 123.24
	Sinkhorn	0.398 ± 0.090	0.419 ± 0.025	0.424 ± 0.031
		8	16	32
Humanoid-v2	Reward	447.87 ± 31.26	505.47 ± 67.62	335.48 ± 65.14
	Sinkhorn	0.760 ± 0.011	0.789 ± 0.013	0.835 ± 0.012

Ant-v2 and HalfCheetah-v2 environments.

Ablation study. To study the effect of minimizing the Sinkhorn distance between occupancy measures using a fixed transport cost, we repeat our experiments on the environments Hopper-v2, HalfCheetah-v2, Walker2d-v2 and Ant-v2 with $\{2, 8, 32\}$ trajectory sets. For Humanoid-v2 we conduct the experiments on sets of $\{8, 16, 32\}$. In this ablation study, instead of training a critic network in an adversarially learned feature space, we assign a reward proxy defined by Equation (7) with a *fixed* cosine transport cost introduced in Equation (15).

Figure 2 depicts the evolution of the Sinkhorn distance between occupancy measures during training of SIL, after replacing the adversarial objective of the critic network with a fixed transport cost. While the training process is more stable, it fails to achieve good performance in terms both of the Sinkhorn distance metric (Figures 1 and 2) and reward metric (see Table 4). The result suggests

that the training objective of the critic network has been a crucial part of the proposed algorithm in providing sufficiently strong signals to the learner policy to match the expert’s state-action distribution.

5 Conclusion

In this work we presented Sinkhorn imitation learning (SIL), a solution to optimal transport based imitation learning, by formulating the problem of matching an expert’s state-action distribution as minimization of their Sinkhorn distance. We utilized an adversarially trained critic that maps the state-action observations to an adversarially learned feature space. The use of the critic provides a discriminative signal to the learner policy to facilitate the imitation of an expert demonstrator’s behavior. Experiments on 5 MuJoCo environments demonstrate that SIL exhibits competitive performance compared to the baselines.

The Sinkhorn imitation learning framework can be extended in several directions to address current limitations which we aim to study in future work. Currently, SIL’s formulation makes it compatible with only on-policy RL methods as computing the Sinkhorn distance necessitates complete trajectories. While SIL is efficient compared to other on-policy adversarial IL benchmarks, it still requires more environment interactions to learn compared to off-policy adversarial IL methods. Hence, it is an interesting future direction to extend SIL to be compatible with off-policy RL algorithms, in line with previous work [7,14,15,25] to yield a method that both inherits the theoretical benefits of OT while being sample efficient. Additionally, performance of SIL was reported with a fixed critic network structure in all studied experiments. Hence, it is unclear what is the effect of the network architecture in guiding imitation learning. It will be of practical significance to investigate the impact of different critic network architectures on training stability and computational efficiency, as well as its relationship to the dimension of state-action space. Another interesting research area is to extend the current framework to incorporate the temporal dependence of the trajectory in the construction of the optimal transport coupling and subsequently the reward proxy. We anticipate that this will be a promising direction for improving the sample efficiency and generalization performance of the optimal transport-based adversarial imitation learning framework.

References

1. Abbeel, P., Dolgov, D., Ng, A.Y., Thrun, S.: Apprenticeship learning for motion planning with application to parking lot navigation. In: Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 1083–1090 (2008)
2. Abbeel, P., Coates, A., Ng, A.Y.: Autonomous helicopter aerobatics through apprenticeship learning. *The International Journal of Robotics Research* **29**(13), 1608–1639 (2010)
3. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein generative adversarial networks. In: Proc. International Conference on Machine Learning (ICML). pp. 214–223. Sydney, Australia (2017)

4. Blondé, L., Kalousis, A.: Sample-efficient imitation learning via generative adversarial nets. In: Proc. International Conference on Artificial Intelligence and Statistics (AISTATS). Okinawa, Japan (2019)
5. Coates, A., Abbeel, P., Ng, A.Y.: Learning for control from multiple demonstrations. In: Proc. International Conference on Machine Learning (ICML). p. 144–151. Helsinki, Finland (2008)
6. Cuturi, M.: Sinkhorn distances: Lightspeed computation of optimal transport. In: Proc. Advances in Neural Information Processing Systems (NeurIPS). pp. 2292–2300. Lake Tahoe, Nevada, USA (2013)
7. Dadashi, R., Hussenot, L., Geist, M., Pietquin, O.: Primal wasserstein imitation learning. In: Proc. International Conference on Learning Representations (ICLR) (2021)
8. Fu, J., Luo, K., Levine, S.: Learning robust rewards with adversarial inverse reinforcement learning. In: Proc. International Conference on Learning Representations (ICLR). Vancouver, Canada (2018)
9. Ghasemipour, S.K.S., Zemel, R., Gu, S.: A divergence minimization perspective on imitation learning methods. In: Proc. Conference on Robot Learning (CoRL). Osaka, Japan (2019)
10. Goodfellow, I., et al.: Generative adversarial nets. In: Proc. Advances in Neural Information Processing Systems (NeurIPS). pp. 2672–2680. Montréal, Canada (2014)
11. Ho, J., Ermon, S.: Generative adversarial imitation learning. In: Proc. Advances in Neural Information Processing Systems (NeurIPS). pp. 4565–4573. Barcelona, Spain (2016)
12. Ke, L., Barnes, M., Sun, W., Lee, G., Choudhury, S., Srinivasa, S.S.: Imitation learning as f-divergence minimization. arXiv:1905.12888 (2019)
13. Kober, J., Peters, J.: Learning motor primitives for robotics. In: Proc. IEEE International Conference on Robotics and Automation. pp. 2112–2118. Kobe, Japan (2009)
14. Kostrikov, I., Agrawal, K.K., Levine, S., Tompson, J.: Discriminator-actor-critic: Addressing sample inefficiency and reward bias in adversarial imitation learning. In: Proc. International Conference on Learning Representations (ICLR). New Orleans, USA (2019)
15. Kostrikov, I., Nachum, O., Tompson, J.: Imitation learning via off-policy distribution matching. In: Proc. International Conference on Learning Representations (ICLR) (2020)
16. Kuderer, M., Kretschmar, H., Burgard, W.: Teaching mobile robots to cooperatively navigate in populated environments. In: Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 3138–3143 (2013)
17. Laskey, M., Lee, J., Fox, R., Dragan, A., Goldberg, K.: Dart: Noise injection for robust imitation learning. In: Proc. Conference on Robot Learning (CoRL). Mountain View, USA (2017)
18. Mnih, V., Kavukcuoglu, K., Silver, D., et al.: Human-level control through deep reinforcement learning. *Nature* **518**, 529–533 (2015)
19. Ng, A.Y., Russell, S.: Algorithms for inverse reinforcement learning. In: Proc. International Conference on Machine Learning (ICML). pp. 663–670. Stanford, CA, USA (2000)
20. Osa, T., Pajarinen, J., Neumann, G., Bagnell, J.A., Abbeel, P., Peters, J.: An Algorithmic Perspective on Imitation Learning. Now Foundations and Trends (2018)

21. Osa, T., Sugita, N., Mitsuishi, M.: Online trajectory planning in dynamic environments for surgical task automation. In: Proc. Robotics: Science and Systems. Berkley, CA, USA (2014)
22. Park, D., Noseworthy, M., Paul, R., Roy, S., Roy, N.: Inferring task goals and constraints using bayesian nonparametric inverse reinforcement learning. In: Proc. Conference on Robot Learning (CoRL). Osaka, Japan (2019)
23. Pomerleau, D.A.: Alvin: An autonomous land vehicle in a neural network. In: Proc. Advances in Neural Information Processing Systems (NeurIPS). pp. 305–313 (1989)
24. Rajeswaran, A., Kumar, V., Gupta, A., Vezzani, G., Schulman, J., Todorov, E., Levine, S.: Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. In: Proc. Robotics: Science and Systems (RSS). Pittsburgh, Pennsylvania (2018)
25. Reddy, S., Dragan, A.D., Levine, S.: SQIL: imitation learning via reinforcement learning with sparse rewards. In: Proc. International Conference on Learning Representations (ICLR) (2020)
26. Ross, S., Gordon, G., Bagnell, J.: A reduction of imitation learning and structured prediction to no-regret online learning. In: Proc. Conference on Artificial Intelligence and Statistics. pp. 627–635. Fort Lauderdale, FL, USA (2011)
27. Ross, S., Bagnell, A.: Efficient reductions for imitation learning. In: Proc. International Conference on Artificial Intelligence and Statistics. pp. 661–668. Chia Laguna Resort, Sardinia, Italy (13–15 May 2010)
28. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X., Chen, X.: Improved techniques for training GANs. In: Proc. Advances in Neural Information Processing Systems (NeurIPS). pp. 2234–2242. Barcelona, Spain (2016)
29. Salimans, T., Zhang, H., Radford, A., Metaxas, D.N.: Improving GANs using optimal transport. In: Proc. International Conference on Learning Representation (ICLR). Vancouver, Canada (2018)
30. Sammut, C., Hurst, S., Kedzier, D., Michie, D.: Learning to fly. In: Proc. International Conference on Machine Learning (ICML). pp. 385–393. Aberdeen, Scotland, United Kingdom (1992)
31. Sasaki, F., Yohira, T., Kawaguchi, A.: Sample efficient imitation learning for continuous control. In: Proc. International Conference on Learning Representations (ICLR) (2019)
32. Schulman, J., Levine, S., Abbeel, P., Jordan, M., Moritz, P.: Trust region policy optimization. In: Proc. International Conference on Machine Learning (ICML). pp. 1889–1897. Lille, France (2015)
33. Silver, D., et al.: Mastering the game of go with deep neural networks and tree search. *Nature* **529**, 484–503 (2016)
34. Sriperumbudur, B.K., Gretton, A., Fukumizu, K., Lanckriet, G.R.G., Schölkopf, B.: A note on integral probability metrics and ϕ -divergences. arXiv:0901.2698 (2009)
35. Todorov, E., Erez, T., Tassa, Y.: Mujoco: A physics engine for model-based control. In: Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 5026–5033 (2012)
36. Torabi, F., Geiger, S., Warnell, G., Stone, P.: Sample-efficient adversarial imitation learning from observation. arXiv:1906.07374 (2019)
37. Villani, C.: Optimal Transport: old and new, vol. 338. Berlin, Germany: Springer Science & Business Media (2008)
38. Vinyals, O., et al.: Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* **575**, 350–354 (2019)

- 39. Xiao, H., Herman, M., Wagner, J., Ziesche, S., Etesami, J., Linh, T.H.: Wasserstein adversarial imitation learning. arXiv:1906.08113 (2019)
- 40. Zhu, Y., Wang, Z., Merel, J., Rusu, A.A., Erez, T., Cabi, S., Tunyasuvunakool, S., Kramár, J., Hadsell, R., de Freitas, N., Heess, N.: Reinforcement and imitation learning for diverse visuomotor skills. arXiv:1802.09564 (2018)
- 41. Ziebart, B., Bagnell, A.J.: Modeling interaction via the principle of maximum causal entropy. In: Proc. International Conference on Machine Learning (ICML). pp. 1255–1262. Haifa, Israel (2010)
- 42. Ziebart, B., Mass, A., Bagnell, A.J., Dey, A.K.: Maximum entropy inverse reinforcement learning. In: Proc. AAAI Conference on Artificial intelligence. pp. 1433–1438 (2008)
- 43. Zucker, M., Ratliff, N., Stolle, M., Chestnutt, J., Bagnell, J.A., Atkeson, C.G., Kuffner, J.: Optimization and learning for rough terrain legged locomotion. The International Journal of Robotics Research **30**(2), 175–191 (2011)