

## A PROOF FOR THEOREM 4.1

Theorem 4.1 shows that the LIL problem (as a bilevel optimization problem), *i.e.*,

$$\max_{\pi} \underbrace{\mathbb{E}_{(s,a) \in \mathcal{D}_E} [\log \pi(a|s)]}_{\textcircled{1}} + \underbrace{\mathbb{E}_{s \in \mathcal{D}_E} [\log P_{\pi}(s)]}_{\textcircled{2}}, \quad \text{s.t. } P_{\pi} = \arg \max_P \mathbb{E}_{s \in \mathcal{D}_{\pi}} [\log P(s)]. \quad (2)$$

can be relaxed to a two-stage optimization problem, *i.e.*,

$$\max_{\pi} \mathbb{E}_{(s,a) \in \mathcal{D}_E} [\log \pi(a|s)] + \mathbb{E}_{s \in \mathcal{D}_{\pi}} [\log P_E(s)], \quad \text{s.t. } P_E = \arg \max_P \mathbb{E}_{s \in \mathcal{D}_E} [\log P(s)]. \quad (3)$$

which estimates the expert state distribution  $P_E$  (*in Stage #1*), and trains the learner policy  $\pi$  based on  $P_E$  (*in Stage #2*).

*Proof.* Given sufficient expert data  $\mathcal{D}_E$ , the expert state distribution  $P_E(s)$  can be estimated by maximizing the likelihood of  $\mathcal{D}_E$  as below, which is naturally an upper-bound of component  $\textcircled{2}$  in eq.(2).

$$P_E = \arg \max_P \mathbb{E}_{s \in \mathcal{D}_E} [\log P(s)], \text{ with } \mathbb{E}_{s \in \mathcal{D}_E} [\log P_E(s)] \geq \max_{\pi} \mathbb{E}_{s \in \mathcal{D}_E} [\log P_{\pi}(s)].$$

This indicates that ideally the LIL problem in eq.(2) aims to learn an optimal policy  $\pi^*$  with a state distribution  $P_{\pi^*}$  matching the expert distribution  $P_E$ , say,  $P_{\pi^*} = P_E$ . Then, by importance sampling (Neal, 2001), we have,

$$\mathbb{E}_{s \sim P_E} [\log P_{\pi}(s)] = \mathbb{E}_{s \sim P_{\pi}} \left[ \frac{P_E(s)}{P_{\pi}(s)} \log P_{\pi}(s) \right].$$

With sufficient expert data  $\mathcal{D}_E$ ,  $\textcircled{2}$  in eq.(2) can be expressed with  $\mathbb{E}_{s \in \mathcal{D}_{\pi}} \left[ \frac{P_E(s)}{P_{\pi}(s)} \log P_{\pi}(s) \right]$ . For ease of notation, we define  $\textcircled{2} = g(\pi, P_{\pi}) \triangleq \mathbb{E}_{s \in \mathcal{D}_{\pi}} \left[ \frac{P_E(s)}{P_{\pi}(s)} \log P_{\pi}(s) \right]$  which is a function of policy  $\pi$  and its state distribution  $P_{\pi}$ . We denote the optimal policy as  $\pi^*$ . Since the optimal state distribution  $P_{\pi^*} = P_E$ , we have the following inequality, *i.e.*,

$$g(\pi^*, P_{\pi^*}) = g(\pi^*, P_E) = \mathbb{E}_{s \in \mathcal{D}_{\pi^*}} [\log P_{\pi^*}(s)] \geq \max_{\pi} g(\pi, P_E) = \max_{\pi} \mathbb{E}_{s \in \mathcal{D}_{\pi}} [\log P_E(s)],$$

where equality holds when  $\pi = \pi^*$ , at least. Therefore, eq.(2) can be rewritten as

$$\begin{aligned} & \max_{\pi} \left( \mathbb{E}_{(s,a) \in \mathcal{D}_E} [\log \pi(a|s)] + \mathbb{E}_{s \in \mathcal{D}_E} [\log P_{\pi}(s)] \right) \\ & \stackrel{\pi = \pi^*}{=} \mathbb{E}_{(s,a) \in \mathcal{D}_E} [\log \pi^*(a|s)] + \mathbb{E}_{s \in \mathcal{D}_E} [\log P_{\pi^*}(s)] \\ & \geq \max_{\pi} \mathbb{E}_{(s,a) \in \mathcal{D}_E} [\log \pi(a|s)] + \max_{\pi} \mathbb{E}_{s \in \mathcal{D}_{\pi}} [\log P_E(s)] \\ & \geq \max_{\pi} \left( \mathbb{E}_{(s,a) \in \mathcal{D}_E} [\log \pi(a|s)] + \mathbb{E}_{s \in \mathcal{D}_{\pi}} [\log P_E(s)] \right), \end{aligned}$$

as the log likelihood terms are all non-positive. This completes the proof.  $\square$

## B EXPERIMENT SETUPS

**Resource usage and running time.** All experiments are run on GeForce RTX2080. The training of  $P_{E\omega^*}$  takes roughly 6 hours for CartPole and Reacher, and 12 hours for Hopper and Walker for 200 iterations.

**Evaluation setup.** We use the same amount of environment interactions and expert demonstrations in GAIL, GPRIL, DRIL, and SLIL, where information of each task is shown in Tab. 2. For all approaches, we evaluate their learner policy performances in every iteration. The experiments with multiple mode task settings all use 18 expert trajectories. We do not subsample expert trajectories for any of the experiment tasks. The evaluation score is achieved via evaluating the mean and std of 50 trajectories generated with the learner policy.

**Model details.** All baselines and SLIL shares the same policy network structure in all experiments – tanh nonlinearities sandwiched with two hidden layers of 100 units. Consistent with GAIL (Ho &

Table 2: Parameters for baselines SLIL.

Task	Training iterations	Number of $(s, a)$ per iteration	Expert performance	Random policy performance
CartPole-v0	200	200	200±0	17± 4
Hopper-v2	1000	1000	3624±19	8± 6
Reacher-v2	200	1000	-4.5±1.7	-93.7 ±4.8
Walker-v2	1000	1000	7002±33	-2±3

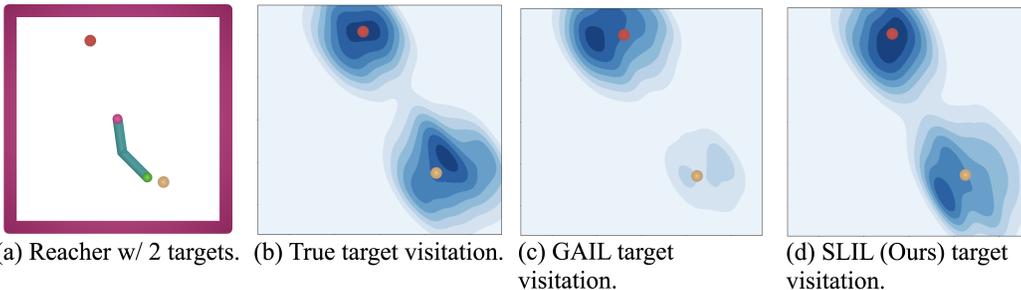


Figure 15: Results obtained by SLIL (Ours) and baselines on mode coverage. (a): A Reacher task, with two targets in different colors. (b)-(d) show the mode coverage (*i.e.*, state distribution) with expert policy (b), GAIL policy (c), and our SLIL policy (d). All the distributions are visualized using KDE (Sheather & Jones, 1991).

Ermon (2016), value functions in SLIL has the the same neural network architecture as the policy networks and employs generalized advantage estimation Schulman et al. (2015b) with  $\lambda = 0.99$  and  $\gamma = 0.95$  to decrease the gradient variance. The DCNF model features a hypernetwork-Ha et al. (2016) with hidden dimension of 32 and width of 64. To train it, we use Dopri5 ODE integrator.

**Hyperparameter details.** For training  $P_{E,\omega^*}$ , we schedule a linear noise level decay with  $\sigma_0 = 1$  and  $\sigma_{200} = 0$ . The learning rate is  $1e - 4$  for all tasks. The training of  $\pi$  also has a learning rate at  $1e - 4$  with gradient clip 0.1.

## C MORE EXPERIMENT RESULTS

We further show the learning curves of DCNF and SoftFlow to better understand the efficacy of the denoising mechanism in Fig. 14. The figure unfolds that in each training epoch the training loss of DCNF lies below SoftFlow. This likely indicates that by decreasing the noise level in each training iteration, the denoising mechanism expose different levels of regularization for optimization. Such types of regularization enable the DCNF model to find a better local minimum than SoftFlow. We therefore hypothesize that the denoising mechanism is able to make the loss landscape more smooth than SoftFlow, and reaches a better minimum entailing better expert state distribution learning. We omit the learning curves in other tasks as similar observations are made.

Fig. 15 shows the mode coverage results of SLIL (ours) and baselines over the Reacher task with two target modes. It shows two mode targets from expert demonstrations Fig. 15(b), where GAIL only visited the red one and SLIL visited both targets. Tab. 3 shows detailed results of SLIL and baselines in tasks with single expert demonstration mode. Fig. 16 shows the learning curves of the test tasks with a single mode when demonstration number is 4. In all tasks, SLIL has more stable training curves (with less mean return perturbation) with higher convergence speed.

Tab. 4 shows the EMD (Ling & Okada, 2007), KL, and reverse KL (RKL) results between expert and learner final state distributions in the Reacher2, Reacher4 and HalfCheetah2 tasks. It reflects that SLIL is better at recovering expert modes than baseline methods.

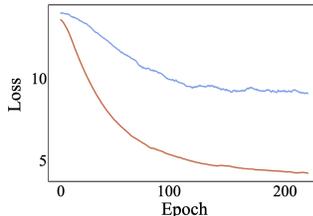


Figure 14: Learning curves of DCNF and SoftFlow in the Hopper task with 18 demonstration trajectories.

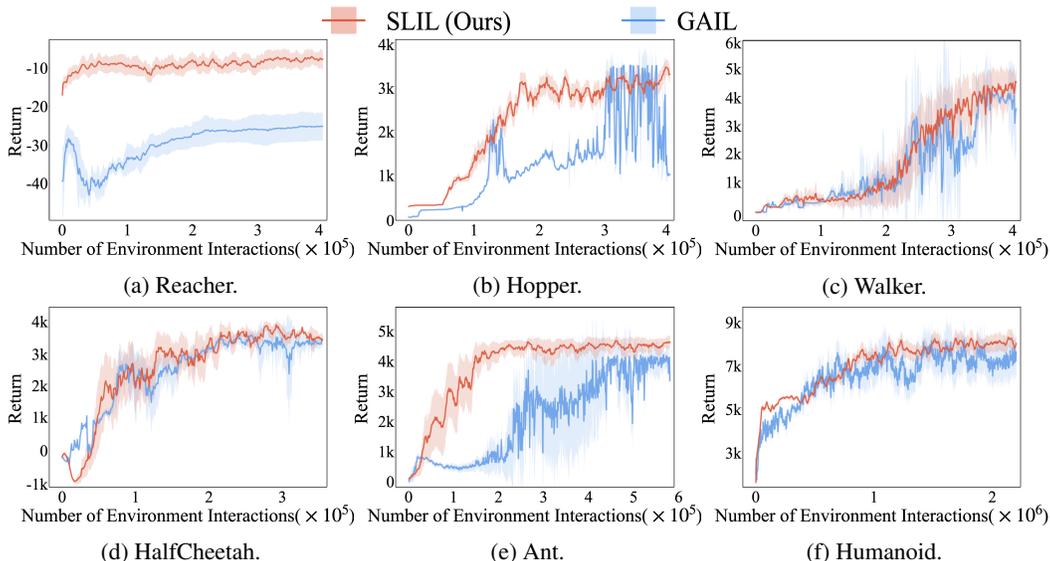


Figure 16: Learning curve comparison between GAIL and SLIL (Ours). All tasks are shown 4 demonstrations. The y-axis is the obtained return (*i.e.*, total reward).

Table 3: Learned policy performance.

Task	Datasize	BC	GAIL	GPRIL	DRIL	SLIL (Ours)
CartPole	1	59 $\pm$ 27	200 $\pm$ 0	53 $\pm$ 16	200 $\pm$ 0	<b>200<math>\pm</math>0</b>
	4	81 $\pm$ 31	200 $\pm$ 0	187 $\pm$ 8	200 $\pm$ 0	<b>200<math>\pm</math>0</b>
	7	137 $\pm$ 27	200 $\pm$ 0	200 $\pm$ 0	200 $\pm$ 0	<b>200<math>\pm</math>0</b>
	10	167 $\pm$ 30	200 $\pm$ 0	200 $\pm$ 0	200 $\pm$ 0	<b>200<math>\pm</math>0</b>
Reacher	4	-10.27 $\pm$ 2.14	-26.90 $\pm$ 7.48	-12.55 $\pm$ 3.54	<b>-9.13<math>\pm</math>.83</b>	-9.44 $\pm$ 3.16
	11	-9.49 $\pm$ 3.66	-12.77 $\pm$ 8.90	-10.45 $\pm$ 5.21	-7.11 $\pm$ 2.23	<b>-6.34<math>\pm</math>3.26</b>
	18	-8.89 $\pm$ 3.83	-7.34 $\pm$ 2.63	-9.96 $\pm$ 5.01	-6.93 $\pm$ 2.37	<b>-6.33<math>\pm</math>2.52</b>
	25	-9.63 $\pm$ 3.84	-6.64 $\pm$ 2.47	-11.87 $\pm$ 4.71	-6.90 $\pm$ 2.65	<b>-5.78<math>\pm</math>2.50</b>
Hopper	4	2352 $\pm$ 894	<b>3394<math>\pm</math>37</b>	22 $\pm$ 1	898 $\pm$ 132	3381 $\pm$ 183
	11	2589 $\pm$ 635	<b>3599<math>\pm</math>4</b>	407 $\pm$ 202	3150 $\pm$ 184	3500 $\pm$ 31
	18	3331 $\pm$ 66	3631 $\pm$ 3	1339 $\pm$ 1390	3611 $\pm$ 3	<b>3686<math>\pm</math>8</b>
	25	3589 $\pm$ 56	3476 $\pm$ 5	1406 $\pm$ 844	3580 $\pm$ 8	<b>3595<math>\pm</math>8</b>
Walker2d	4	1233 $\pm$ 969	4070 $\pm$ 1010	557 $\pm$ 357	555 $\pm$ 148	<b>4230<math>\pm</math>1108</b>
	11	3456 $\pm$ 863	5108 $\pm$ 410	1042 $\pm$ 75	4567 $\pm$ 1231	<b>6537<math>\pm</math>560</b>
	18	4477 $\pm$ 1329	6671 $\pm$ 39	1464 $\pm$ 637	<b>6886<math>\pm</math>202</b>	6873 $\pm$ 55
	25	5294 $\pm$ 1860	6815 $\pm$ 20	2254 $\pm$ 1006	6693 $\pm$ 130	<b>6923<math>\pm</math>58</b>
Ant	4	4204 $\pm$ 289	4218 $\pm$ 240	2722 $\pm$ 36	3837 $\pm$ 259	<b>4613<math>\pm</math>161</b>
	11	<b>4577<math>\pm</math>145</b>	4105 $\pm$ 223	2510 $\pm$ 27	4515 $\pm$ 239	4540 $\pm$ 169
	18	4736 $\pm$ 75	4690 $\pm$ 102	2755 $\pm$ 183	4703 $\pm$ 40	<b>4752<math>\pm</math>91</b>
	25	4682 $\pm$ 89	4735 $\pm$ 54	2656 $\pm$ 85	4690 $\pm$ 75	<b>4825<math>\pm</math>45</b>
HalfCheetah	4	2070 $\pm$ 528	3254 $\pm$ 133	558 $\pm$ 148	359 $\pm$ 266	<b>3687<math>\pm</math>416</b>
	11	3979 $\pm$ 61	4015 $\pm$ 344	2655 $\pm$ 253	<b>4063<math>\pm</math>50</b>	4052 $\pm$ 236
	18	3911 $\pm$ 416	4393 $\pm$ 212	2666 $\pm$ 186	4185 $\pm$ 30	<b>4531<math>\pm</math>65</b>
	25	4027 $\pm$ 91	<b>4423<math>\pm</math>104</b>	3619 $\pm$ 257	4227 $\pm$ 26	4416 $\pm$ 77
Humanoid	80	6145 $\pm$ 1918	8268 $\pm$ 1401	2048 $\pm$ 1140	<b>8800<math>\pm</math>639</b>	8404 $\pm$ 571
	160	6722 $\pm$ 1126	<b>9994<math>\pm</math>1053</b>	6023 $\pm$ 1006	9507 $\pm$ 832	9771 $\pm$ 835
	240	8834 $\pm$ 998	<b>9430<math>\pm</math>906</b>	8091 $\pm$ 878	9185 $\pm$ 492	9294 $\pm$ 385

Task	Approach	Measure		
		EMD	KL	RKL
Reacher2	BC	1.01	2.60	4.38
	GAIL	0.84	2.47	4.51
	DRIL	1.00	2.51	4.25
	SLIL	<b>0.47</b>	<b>2.39</b>	<b>3.25</b>
Reacher4	BC	0.81	4.23	6.33
	GAIL	0.41	4.49	6.22
	DRIL	0.55	5.16	6.16
	SLIL	<b>0.33</b>	<b>3.94</b>	<b>4.53</b>
HalfCheetah2	BC	1.81	6.74	12.78
	GAIL	1.75	4.52	18.13
	DRIL	1.83	6.90	12.76
	SLIL	<b>1.01</b>	<b>3.89</b>	<b>11.54</b>

Table 4: The EMD (Ling & Okada, 2007), KL and RKL between expert and learned policy state distribution.