
Implicit Policy for Reinforcement Learning

Yunhao Tang
Columbia University
yt2541@columbia.edu

Shipra Agrawal *
Columbia University
sa3305@columbia.edu

Abstract

We introduce *Implicit Policy*, a general class of expressive policies that can flexibly represent complex action distributions in reinforcement learning, with efficient algorithms to compute entropy regularized policy gradients. We empirically show that, despite its simplicity in implementation, entropy regularization combined with a rich policy class can attain desirable properties displayed under maximum entropy reinforcement learning framework, such as robustness and multi-modality.

1 Introduction

Reinforcement Learning (RL) combined with deep neural networks have led to a wide range of successful applications, including the game of Go, robotics control and video game playing [32, 30, 24]. During the training of deep RL agent, the injection of noise into the learning procedure can usually prevent the agent from premature convergence to bad locally optimal solutions, for example, by entropy regularization [30, 23] or by explicitly optimizing a maximum entropy objective [13, 25].

Though entropy regularization is much simpler to implement in practice, it greedily optimizes the policy entropy at each time step, without accounting for future effects. On the other hand, maximum entropy objective considers the entropy of the distribution over entire trajectories, and is more conducive to theoretical analysis [2]. Recently, [13, 14] also shows that optimizing the maximum entropy objective can lead to desirable properties such as robustness and multi-modal policy.

Can we preserve the simplicity of entropy regularization while attaining desirable properties under maximum entropy framework? To achieve this, a necessary condition is an expressive representation of policy. Though various flexible probabilistic models have been proposed in generative modeling [10, 37], such models are under-explored in policy based RL. To address such issues, we propose flexible policy classes and efficient algorithms to compute entropy regularized policy gradients.

In Section 3, we introduce *Implicit Policy*, a generic policy representation from which we derive two expressive policy classes, Normalizing Flows Policy (*NFP*) and more generally, Non-invertible Blackbox Policy (*NBP*). *NFP* provides a novel architecture that embeds state information into Normalizing Flows; *NBP* assumes little about policy architecture, yet we propose algorithms to efficiently compute entropy regularized policy gradients when the policy density is not accessible. In Section 4, we show that entropy regularization optimizes a lower bound of maximum entropy objective. In Section 5, we show that when combined with entropy regularization, expressive policies achieve competitive performance on benchmarks and leads to robust and multi-modal policies.

2 Preliminaries

2.1 Background

We consider the standard RL formalism consisting of an agent interacting with the environment. At time step $t \geq 0$, the agent is in state $s_t \in \mathcal{S}$, takes action $a_t \in \mathcal{A}$, receives instant reward $r_t \in \mathbb{R}$ and transitions to next state $s_{t+1} \sim p(s_{t+1}|s_t, a_t)$. Let $\pi : \mathcal{S} \mapsto \mathcal{A}$ be a policy. The objective of

* This research was supported by an Amazon Research Award (2017) and AWS cloud credits.

RL is to search for a policy which maximizes cumulative expected reward $J(\pi) = \mathbb{E}_\pi [\sum_{t=0}^{\infty} r_t \gamma^t]$, where $\gamma \in (0, 1]$ is a discount factor. The action value function of policy π is defined as $Q^\pi(s, a) = \mathbb{E}_\pi [\sum_{t=0}^{\infty} r_t \gamma^t | s_0 = s, a_0 = a]$. In policy based RL, a policy is explicitly parameterized as π_θ with parameter θ , and the policy can be updated by policy gradients $\theta \leftarrow \theta + \alpha \nabla_\theta J(\pi_\theta)$, where α is the learning rate. So far, there are in general two ways to compute policy gradients for either on-policy or off-policy updates.

Score function gradient & Pathwise gradient. Given a stochastic policy $a_t \sim \pi_\theta(\cdot | s_t)$, the score function gradient for on-policy update is computed as $\nabla_\theta J(\pi_\theta) = \mathbb{E}_{\pi_\theta} [\sum_{t=0}^{\infty} Q^{\pi_\theta}(s_t, a_t) \nabla_\theta \log \pi_\theta(a_t | s_t)]$ as in [31, 30, 23, 35]. For off-policy update, it is necessary to introduce importance sampling weights to adjust the distribution difference between the behavior policy and current policy. Given a deterministic policy $a_t = \pi_\theta(s_t)$, the pathwise gradient for on-policy update is computed as $\nabla_\theta J(\pi_\theta) = \mathbb{E}_{\pi_\theta} [\sum_{t=0}^{\infty} \nabla_a Q^{\pi_\theta}(s_t, a)|_{a=\pi_\theta(s_t)} \nabla_\theta \pi_\theta(s_t)]$. In practice, this gradient is often computed off-policy [33, 32], where the exact derivation comes from a modified off-policy objective [3].

Entropy Regularization. For on-policy update, it is common to apply entropy regularization [38, 26, 23, 31]. Let $\mathbb{H}[\pi(\cdot | s)]$ be the entropy of policy π at state s . The entropy regularized update is

$$\theta \leftarrow \theta + \alpha \{ \nabla_\theta J(\pi_\theta) + \beta \mathbb{E}_{\pi_\theta} [\nabla_\theta \sum_{t=0}^{\infty} \mathbb{H}[\pi_\theta(\cdot | s_t)] \gamma^t] \}, \quad (1)$$

where $\beta > 0$ is a regularization constant. By boosting policy entropy, this update can potentially prevent the policy from premature convergence to bad locally optimal solutions. In Section 3, we will introduce expressive policies that leverage both on-policy/off-policy updates, and algorithms to efficiently compute entropy regularized policy gradients.

Maximum Entropy RL. In maximum entropy RL formulation, the objective is to maximize the cumulative reward and the policy entropy $J_{\text{MaxEnt}}(\pi_\theta) = \mathbb{E}_{\pi_\theta} [\sum_{t=0}^{\infty} r_t \gamma^t + \beta \sum_{t=0}^{\infty} \mathbb{H}[\pi_\theta(\cdot | s_t)] \gamma^t]$, where $\beta > 0$ is a tradeoff constant. Note that $\nabla_\theta J_{\text{MaxEnt}}(\pi_\theta)$ differs from the update in (1) by an exchange of expectation and gradient. The intuition of $J_{\text{MaxEnt}}(\pi_\theta)$ is to achieve high reward while being as random as possible over trajectories. Since there is no simple low variance gradient estimate for $J_{\text{MaxEnt}}(\pi_\theta)$, several previous works [31, 13, 25] have proposed to optimize $J_{\text{MaxEnt}}(\pi_\theta)$ primarily using off-policy value based algorithms.

2.2 Related Work

A large number of prior works have implemented policy gradient algorithms with entropy regularization [30, 31, 23, 26], which boost exploration by greedily maximizing policy entropy at each time step. In contrast to such greedy procedure, maximum entropy objective considers entropy over the entire policy trajectories [13, 25, 29]. Though entropy regularization is simpler to implement in practice, [12, 13] argues in favor of maximum entropy objective by showing that trained policies can be robust to noise, which is desirable for real life robotics tasks; and multi-modal, a potentially desired property for exploration and fine-tuning for downstream tasks. However, their training procedure is fairly complex, which consists of training a soft Q function by fixed point iteration and a neural sampler by Stein variational gradient [21]. We argue that properties as robustness and multi-modality are attainable through simple entropy regularized policy gradient algorithms combined with expressive policy representations.

Prior works have studied the property of maximum entropy objective [25, 39], entropy regularization [26] and their connections with variants of operators [2]. It is commonly believed that entropy regularization greedily maximizes local policy entropy and does not account for how a policy update impacts future states. In Section 4, we show that entropy regularized policy gradient update maximizes a lower bound of maximum entropy objective, given constraints on the differences between consecutive policy iterates. This partially justifies why simple entropy regularization combined with expressive policy classes can achieve competitive empirical performance in practice.

There is a number of prior works that discuss different policy architectures. The most common policy for continuous control is unimodal Gaussian [30, 31, 23]. [14] discusses mixtures of Gaussian, which

can represent multi-modal policies but it is necessary to specify the number of modes in advance. [13] also represents a policy using implicit model, but the policy is trained to sample from the soft Q function instead of being trained directly. Recently, we find [11] also uses Normalizing Flows to represent policies, but their focus is learning an hierarchy and involves layers of pre-training. Contrary to early works, we propose to represent flexible policies using implicit models/Normalizing Flows and efficient algorithms to train the policy end-to-end.

Implicit models have been extensively studied in probabilistic inference and generative modeling [10, 17, 19, 37]. Implicit models define distributions by transforming source noise via a forward pass of neural networks, which in general sacrifice tractable probability density for more expressive representation. Normalizing Flows are a special case of implicit models [27, 5, 6], where transformations from source noise to output are invertible and allow for maximum likelihood inference. Borrowing inspirations from prior works, we introduce implicit models into policy representation and empirically show that such rich policy class entails multi-modal behavior during training. In [37], GAN [10] is used as an optimal density estimator for likelihood free inference. In our work, we apply similar idea to compute entropy regularization when policy density is not available.

3 Implicit Policy for Reinforcement Learning

We assume the action space \mathcal{A} to be a compact subset of \mathbb{R}^m . Any sufficiently smooth stochastic policy can be represented as a blackbox $f_\theta(\cdot)$ with parameter θ that incorporates state information s and independent source noise ϵ sampled from a simple distribution $\rho_0(\cdot)$. In state s , the action a is sampled by a forward pass in the blackbox.

$$a = f_\theta(s, \epsilon), \epsilon \sim \rho_0(\cdot). \quad (2)$$

For example, Gaussian policy is reduced to $a = \sigma_\theta(s) \cdot \epsilon + \mu_\theta(s)$ where ρ_0 is standard Gaussian [30]. In general, the distribution of a_t is implicitly defined: for any set A of \mathcal{A} , $\mathbb{P}(a \in A|s) = \int_{\epsilon: f_\theta(s, \epsilon) \in A} \rho_0(\epsilon) d\epsilon$. Let $\pi_\theta(\cdot|s)$ be the density of this distribution². We call such policy *Implicit Policy* as similar ideas have been previous explored in implicit generative modeling literature [10, 19, 37]. In the following, we derive two expressive stochastic policy classes following this blackbox formulation, and propose algorithms to efficiently compute entropy regularized policy gradients.

3.1 Normalizing Flows Policy (NFP)

We first construct a stochastic policy with Normalizing Flows. Normalizing Flows [27, 6] have been applied in variational inference and probabilistic modeling to represent complex distributions. In general, consider transforming a source noise $\epsilon \sim \rho_0(\cdot)$ by a series of invertible nonlinear function $g_{\theta_i}(\cdot)$, $1 \leq i \leq K$ each with parameter θ_i , to output a target sample x ,

$$x = g_{\theta_K} \circ g_{\theta_{K-1}} \circ \dots \circ g_{\theta_2} \circ g_{\theta_1}(\epsilon). \quad (3)$$

Let Σ_i be the Jacobian matrix of $g_{\theta_i}(\cdot)$, then the density of x is computed by chain rule,

$$\log p(x) = \log p(\epsilon) + \sum_{i=1}^K \log \det(\Sigma_i). \quad (4)$$

For a general invertible transformation $g_{\theta_i}(\cdot)$, computing $\det(\Sigma_i)$ is expensive. We follow the architecture of [5] to ensure that $\det(\Sigma_i)$ is computed in linear time. To combine state information, we embed state s by another neural network $L_{\theta_s}(\cdot)$ with parameter θ_s and output a state vector $L_{\theta_s}(s)$ with the same dimension as ϵ . We can then insert the state vector between any two layers of (3) to make the distribution conditional on state s . In our implementation, we insert the state vector after the first transformation (we detail our architecture design in Appendix C).

$$a = g_{\theta_K} \circ g_{\theta_{K-1}} \circ \dots \circ g_{\theta_2} \circ (L_{\theta_s}(s) + g_{\theta_1}(\epsilon)). \quad (5)$$

Though the additive form of $L_{\theta_s}(s)$ and $g_{\theta_1}(\epsilon)$ may in theory limit the capacity of the model, in practice we find the resulting policy still very expressive. For simplicity, we denote the above

²In future notations, when the context is clear, we use $\pi_\theta(\cdot|s)$ to denote both the density of the policy as well as the policy itself: for example, $a \sim \pi_\theta(\cdot|s)$ means sampling a from the policy; $\log \pi_\theta(a|s)$ means the log density of policy at a in state s .

transformation (5) as $a = f_\theta(s, \epsilon)$ with parameter $\theta = \{\theta_s, \theta_i, 1 \leq i \leq K\}$. It is obvious that $\epsilon \leftrightarrow a = f_\theta(s, \epsilon)$ is still invertible between a and ϵ , which is critical for computing $\log \pi_\theta(a|s)$ according to (4). Such representations build complex policy distributions with explicit probability density $\pi_\theta(\cdot|s)$, and hence entail training using score function gradient estimators.

Since there is no analytic form for entropy, we use samples to estimate entropy by re-parameterization, $\mathbb{H}[\pi_\theta(\cdot|s)] = \mathbb{E}_{a \sim \pi_\theta(\cdot|s)}[-\log \pi_\theta(a|s)] = \mathbb{E}_{\epsilon \sim \rho_0(\cdot)}[-\log \pi_\theta(f_\theta(s, \epsilon)|s)]$. The gradient of entropy can be easily computed by a pathwise gradient and easily implemented using back-propagation $\nabla_\theta \mathbb{H}[\pi_\theta(\cdot|s)] = \mathbb{E}_{\epsilon \sim \rho_0(\cdot)}[-\nabla_\theta \log \pi_\theta(f_\theta(s, \epsilon)|s)]$.

On-policy algorithm for NFP. Any on-policy policy optimization algorithms can be easily combined with NFP. Since NFP has explicit access to policy density, it allows for training using score function gradient estimators with efficient entropy regularization.

3.2 Non-invertible Blackbox Policy (NBP)

The forward pass in (2) transforms the simple noise distribution $\epsilon \sim \rho_0(\cdot)$ to complex action distribution $a_t \sim \pi_\theta(\cdot|s_t)$ through the blackbox $f_\theta(\cdot)$. However, the mapping $\epsilon \mapsto a_t$ is in general non-invertible and we do not have access to the density $\pi_\theta(\cdot|s_t)$. We derive a pathwise gradient for such cases and leave all the proof in Appendix A.

Theorem 3.1 (Stochastic Pathwise Gradient). *Given an implicit stochastic policy $a_t = f_\theta(s_t, \epsilon)$, $\epsilon \sim \rho_0(\cdot)$. Let π_θ be the implicitly defined policy. Then the pathwise policy gradient for the stochastic policy is*

$$\nabla_\theta J(\pi_\theta) = \mathbb{E}_{\pi_\theta} [\mathbb{E}_{\epsilon \sim \rho_0(\cdot)} [\nabla_\theta f_\theta(s, \epsilon) \nabla_a Q^{\pi_\theta}(s, a)|_{a=f_\theta(s, \epsilon)}]]. \quad (6)$$

To compute the gradient of policy entropy for such general implicit policy, we propose to train an additional classifier $c_\psi : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$ with parameter ψ along with policy π_θ . The classifier c_ψ is trained to minimize the following objective given a policy π_θ

$$\min_{\psi} \mathbb{E}_{a \sim \pi_\theta(\cdot|s)} [-\log \sigma(c_\psi(a, s))] + \mathbb{E}_{a \sim \mathcal{U}(\mathcal{A})} [-\log(1 - \sigma(c_\psi(a, s)))], \quad (7)$$

where $\mathcal{U}(\mathcal{A})$ is a uniform distribution over \mathcal{A} and $\sigma(\cdot)$ is the sigmoid function. We have lemma A.1 in Appendix A.2 to guarantee that the optimal solution ψ^* of (7) provides an estimate of policy density, $c_{\psi^*}(s, a) = \log \frac{\pi_\theta(a|s)}{|\mathcal{A}|-1}$. As a result, we could evaluate the entropy by simple re-parametrization $\mathbb{H}[\pi_\theta(\cdot|s)] = \mathbb{E}_{\epsilon \sim \rho_0(\cdot)} [-\log \pi(f_\theta(s, \epsilon)|s)] \approx \mathbb{E}_{\epsilon \sim \rho_0(\cdot)} [-c_{\psi^*}(f_\theta(s, \epsilon), s)]$. Further, we can compute gradients of the policy entropy through the density estimate as shown by the following theorem.

Theorem 3.2 (Unbiased Entropy Gradient). *Let ψ^* be the optimal solution from (7), where the policy $\pi_\theta(\cdot|s)$ is given by implicit policy $a = f_\theta(s, \epsilon)$, $\epsilon \sim \rho_0(\cdot)$. The gradient of entropy $\nabla_\theta \mathbb{H}[\pi_\theta(\cdot|s)]$ can be computed as*

$$\nabla_\theta \mathbb{H}[\pi_\theta(\cdot|s)] = -\mathbb{E}_{\epsilon \sim \rho_0(\cdot)} [\nabla_\theta c_{\psi^*}(f_\theta(s, \epsilon), s)] \quad (8)$$

It is worth noting that to compute $\nabla_\theta \mathbb{H}[\pi_\theta(\cdot|s)]$, simply plugging in $c_{\psi^*}(a, s)$ to replace $\log \pi_\theta(a|s)$ in the entropy definition does not work in general, since the optimal solution ψ^* of (7) implicitly depends on θ . However, fortunately in this case the additional term vanishes. The above theorem guarantees that we could apply entropy regularization even when the policy density is not accessible.

Off-policy algorithm for NBP. We develop an off-policy algorithm for NBP. The agent contains an implicit $f_\theta(s, \epsilon)$ with parameter θ , a critic $Q_\phi(s, a)$ with parameter ϕ and a classifier $c_\psi(s, a)$ with parameter ψ . At each time step t , we sample action $a_t = f_\theta(s_t, \epsilon)$, $\epsilon \sim \rho_0(\cdot)$ and save experience tuple $\{s_t, a_t, r_t, s_{t+1}\}$ to a replay buffer B . During training, we sample a mini-batch of tuples from B , update critic $Q_\phi(s, a)$ using TD learning, update policy $f_\theta(s, \epsilon)$ using pathwise gradient (6) and update classifier $c_\psi(s, a)$ by gradient descent on (7). We also maintain target networks $f_{\theta^-}(s, \epsilon)$, $Q_{\phi^-}(s, a)$ with parameter θ^- , ϕ^- to stabilize learning [24, 32]. The pseudocode is listed in Appendix D.

4 Entropy Regularization and Maximum Entropy RL

Though policy gradient algorithms with entropy regularization are easy to implement in practice, they are harder to analyze due to the lack of a global objective. Now we show that entropy regularization maximizes a lower bound of maximum entropy objective when consecutive policy iterates are close.

At each iteration of entropy regularized policy gradient algorithm, the policy parameter is updated as in (1). Following similar ideas in [15, 30], we now interpret such update as maximizing a linearized surrogate objective in the neighborhood of the previous policy iterate $\pi_{\theta_{\text{old}}}$. The surrogate objective is

$$J_{\text{surr}}(\pi_{\theta}) = J(\pi_{\theta}) + \beta \mathbb{E}_{\pi_{\theta_{\text{old}}}} \left[\sum_{t=0}^{\infty} \mathbb{H}[\pi_{\theta}(\cdot|s_t)] \gamma^t \right]. \quad (9)$$

The first-order Taylor expansion of (9) centering at θ_{old} gives a linearized surrogate objective $J_{\text{surr}}(\pi_{\theta}) \approx J_{\text{surr}}(\pi_{\theta_{\text{old}}}) + \nabla_{\theta} J_{\text{surr}}(\pi_{\theta})|_{\theta=\theta_{\text{old}}} (\theta - \theta_{\text{old}})$. Let $\delta\theta = \theta - \theta_{\text{old}}$, the entropy regularized update (1) is equivalent to solving the following optimization problem then update according to $\theta \leftarrow \theta_{\text{old}} + \delta\theta$,

$$\begin{aligned} \min_{\delta\theta} & \left[\nabla_{\theta} J_{\text{surr}}(\pi_{\theta})|_{\theta=\theta_{\text{old}}} \right]^T \delta\theta \\ \text{s.t.} & \|\delta\theta\|_2 \leq C(\alpha, \theta_{\text{old}}), \end{aligned}$$

where $C(\alpha, \theta_{\text{old}})$ is a positive constant depending on both the learning rate α and the previous iterate θ_{old} , and can be recovered from (1). The next theorem shows that by constraining the KL divergence of consecutive policy iterates, the surrogate objective (9) forms a non-trivial lower bound of maximum entropy objective,

Theorem 4.1 (Lower Bound). *If $\mathbb{KL}[\pi_{\theta}||\pi_{\theta_{\text{old}}}] \leq \alpha$, then*

$$J_{\text{MaxEnt}}(\pi) \geq J_{\text{surr}}(\pi) - \frac{\beta\gamma\sqrt{\alpha}\epsilon}{(1-\gamma)^2}, \quad \text{where } \epsilon = \max_s |\mathbb{H}[\pi(\cdot|s)]|. \quad (10)$$

By optimizing $J_{\text{surr}}(\pi_{\theta})$ at each iteration, entropy regularized policy gradient algorithms maximize a lower bound of $J_{\text{MaxEnt}}(\pi_{\theta})$. This implies that though entropy regularization is a greedier procedure than optimizing maximum entropy objective, it accounts for certain effects that the maximum entropy objective is designed to capture. Nevertheless, the optimal solutions of both optimization procedures are different. Previous works [26, 13] have shown that the optimal solutions of both procedures are energy based policies, with energy functions being fixed points of Boltzmann operator and Mellowmax operator respectively [2]. In Appendix B, we show that Boltzmann operator interpolates between Bellman operator and Mellowmax operator, which asserts that entropy regularization is greedier than optimizing $J_{\text{MaxEnt}}(\pi_{\theta})$, yet it still maintains uncertainties in the policy updates.

Though maximum entropy objective accounts for long term effects of policy entropy updates and is more conducive to analysis [2], it is hard to implement a simple yet scalable procedure to optimize the objective [13, 14, 2]. Entropy regularization, on the other hand, is simple to implement in both on-policy and off-policy setting. In experiments, we will show that entropy regularized policy gradients combined with expressive policies achieve competitive performance in multiple aspects.

5 Experiments

Our experiments aim to answer the following questions: (1) Will expressive policy be hard to train, does implicit policy provide competitive performance on benchmark tasks? (2) Are implicit policies robust to noises on locomotion tasks? (3) Does implicit policy + entropy regularization entail multi-modal policies as displayed under maximum entropy framework [13]?

To answer (1), we evaluate both NFP and NBP agent on benchmark continuous control tasks in MuJoCo [36] and compare with baselines. To answer (2), we compare NFP with unimodal Gaussian policy on locomotion tasks with additive observational noises. To answer (3), we illustrate the multi-modal capacity of both policy representations on specially designed tasks illustrated below, and compare with baselines. In all experiments, for NFP, we implement with standard PPO for on-policy update to approximately enforce the KL constraint (10) as in [31]; for NBP, we implement the off-policy algorithm developed in Section 3. In Appendix C and F, we detail hyper-parameter settings in the experiments and provide a small ablation study.

5.1 Locomotion Tasks

Benchmark tasks. One potential disadvantage of expressive policies compared to simple policies (like unimodal Gaussian) is that they pose a more serious statistical challenge due to a larger number of parameters. To see if implicit policy suffers from such problems, we evaluate NFP and NBP on MuJoCo benchmark tasks. For each task, we train for a prescribed number of time steps, then report the results averaged over 5 random seeds. We compare the results with baseline algorithms, such as DDPG [32], SQL [13], TRPO [30] and PPO [31], where baseline TRPO and PPO use unimodal Gaussian policies. As can be seen from Table 1, both NFP and NBP achieve competitive performances on benchmark tasks: they outperform DDPG, SQL and TRPO on most tasks. However, baseline PPO tends to come on top on most tasks. Interestingly on HalfCheetah, baseline PPO gets stuck on a locally optimal gait, which NFP improves upon by a large margin.

Tasks	Timesteps	DDPG	SQL	TRPO	PPO	NFP	NBP
Hopper	$2.00 \cdot 10^6$	≈ 1100	≈ 1500	≈ 1250	$\approx \mathbf{2130}$	≈ 1640	$\approx \mathbf{1880}$
HalfCheetah	$1.00 \cdot 10^7$	≈ 6500	$\approx \mathbf{8000}$	≈ 1800	≈ 1590	≈ 4000	$\approx \mathbf{6560}$
Walker2d	$5.00 \cdot 10^6$	≈ 1600	≈ 2100	≈ 800	$\approx \mathbf{3800}$	$\approx \mathbf{3000}$	≈ 2450
Ant	$1.00 \cdot 10^7$	≈ 200	≈ 2000	≈ 0	$\approx \mathbf{4440}$	$\approx \mathbf{2500}$	≈ 2070

Table 1: A comparison of implicit policy optimization with baseline algorithms on MuJoCo benchmark tasks. For each task, we show the average rewards achieved after training the agent for a fixed number of time steps. The results for NFP and NBP are averaged over 5 random seeds. The results for DDPG, SQL and TRPO are approximated based on the figures in [14], PPO is from OpenAI baseline implementation [4]. We highlight the top two algorithms for each task in bold font. Both TRPO and PPO use unimodal Gaussian policies.

Robustness to Noisy Observations. We add independent Gaussian noise $\mathcal{N}(0, 0.1^2)$ to each component of the observations to make the original tasks partially observable. Since PPO with unimodal Gaussian achieves leading performance on noise-free locomotion tasks across on-policy baselines (A2C [23], TRPO [30]) as shown in [31] and Appendix E.1, we compare NFP only with PPO with unimodal Gaussian on such noisy locomotion tasks. In Figure 1, we show the learning curves of both agents, where on many tasks NFP learns significantly faster than unimodal Gaussian. Why complex policies may add to robustness? We propose that since these control tasks are known to be solved by multiple separate modes of policy [22], observational noises potentially blur these modes and make it harder for a unimodal Gaussian policy to learn any single mode (e.g. unimodal Gaussian puts probability mass between two neighboring modes [18]). On the contrary, NFP can still navigate a more complex reward landscape thanks to a potentially multi-modal policy distribution and learn effectively. We leave a more detailed study of robustness, multi-modality and complex reward landscape as interesting future work.

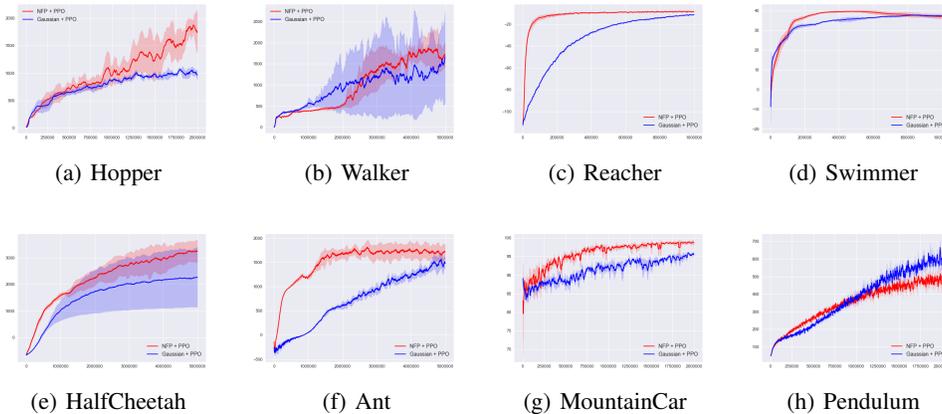


Figure 1: Noisy Observations: learning curves on noisy locomotion tasks. For each task, the observation is added a Gaussian noise $\mathcal{N}(0, 0.1^2)$ component-wise. Each curve is averaged over 4 random seeds. Red is NFP and blue is unimodal Gaussian, both implemented with PPO. NFP beats Gaussian on most tasks.

5.2 Multi-modal policy

Gaussian Bandits. Though factorized unimodal policies suffice for most benchmark tasks, below we motivate the importance of a flexible policy by a simple example: Gaussian bandits. Consider a two dimensional bandit $\mathcal{A} = [-1, 1]^2$. The reward of action a is $-a^T \Sigma^{-1} a$ for a positive definite matrix Σ . The optimal policy for maximum entropy objective is $\pi^*(a) \propto \exp(-\frac{a^T \Sigma^{-1} a}{\beta})$, i.e. a Gaussian policy with covariance matrix Σ . We compare NFP with PPO with factorized Gaussian. As illustrated in Figure 2(a), NFP can approximate the optimal Gaussian policy pretty closely while the factorized Gaussian cannot capture the high correlation between the two action components.

Navigating 2D Multi-goal. We motivate the strength of implicit policy to represent multi-modal policy by Multi-goal environment [13]. The agent has 2D coordinates as states $\mathcal{S} \subset \mathbb{R}^2$ and 2D forces as actions $\mathcal{A} \subset \mathbb{R}^2$. A ball is randomly initialized near the origin and the goal is to push the ball to reach one of the four goal positions plotted as red dots in Figure 2(b). While a unimodal policy can only deterministically commit the agent to one of the four goals, a multi-modal policy obtained by NBP can stochastically commit the agent to multiple goals. On the right of Figure 2(b) we also show sampled actions and contours of Q value functions at various states: NBP learns a very flexible policy with different number of modes in different states.

Learning a Bimodal Reacher. For a more realistic example, consider learning a bimodal policy for reaching one of two targets (Figure 3(a)). The agent has the physical coordinates of the reaching arms as states $\mathcal{S} \subset \mathbb{R}^9$ and applies torques to the joints as actions $\mathcal{A} \subset \mathbb{R}^2$. The objective is to move the reacher head to be close to one of the targets. As illustrated by trajectories in Figure 2(c), while a unimodal Gaussian policy can only deterministically reach one target (red curves), a NFP agent can capture both modes by stochastically reaching one of the two targets (blue curves).

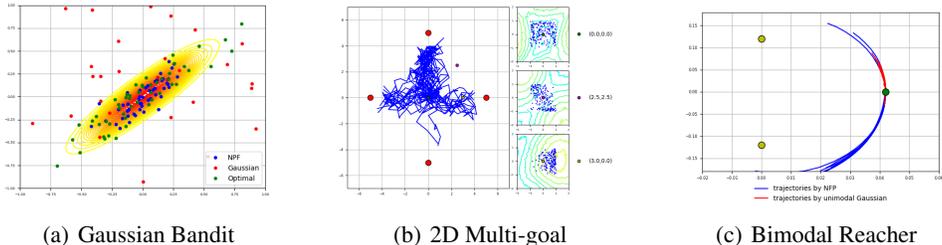


Figure 2: (a): Illustration of Gaussian bandits. The x and y axes are actions. Green dots are actions from the optimal policy, a Gaussian distribution with covariance structure illustrated by the contours. Red dots and blue dots are actions sampled from a learned factorized Gaussian and NFP. NFP captures the covariance of the optimal policy while factorized Gaussian cannot. (b): Illustration of 2D multi-goal environment. Left: trajectories generated by trained NBP agent (solid blue curves). The x and y axes are coordinates of the agent (state). The agent is initialized randomly near the origin. The goals are red dots, and instant rewards are proportional to the agent’s minimum distance to one of the four goals. Right: predicted Q value contours by the critic (light blue): low value, light green: high value and actions sampled from the policy (blue dots) at three selected states. The NFP policy has different number of modes at different states. (c): Trajectories of the reacher head by NFP (blue curves) and unimodal Gaussian policies (red curves) for the bimodal reacher. Yellow dots are locations of the two targets, and the green dot is the starting location of the reacher.

Fine-tuning for downstream tasks. A recent paradigm for RL is to pre-train an agent to perform a conceptually high-level task, which may accelerate fine-tuning the agent to perform more specific tasks [13]. We consider pre-training a quadrupedal robot (Figure 3(b)) to run fast, then fine-tune the robot to run fast in a particular direction [13] as illustrated in Figure 3(c), where we set walls to limit the directions in which to run. Wide and Narrow Hallways tasks differ by the distance of the opposing walls. If an algorithm does not inject enough diversity during pre-training, it will commit the agent to prematurely run in a particular direction, which is bad for fine-tuning. We compare the pre-training capacity of DDPG [20], SQL [13] and NBP. As shown in Figure 3(d), after pre-training, NBP agent manages to run in multiple directions, while DDPG agent runs in a single direction due to a deterministic policy (Appendix E.2). In Table 2, we compare the cumulative rewards of agents after fine-tuning on downstream tasks with different pre-training as initializations. In both tasks, we find NBP to outperform DDPG, SQL and random initialization (no pre-training) by statistically significant margins,

potentially because NBP agent learns a high-level running gait that is more conducive to fine-tuning. Interestingly, in Narrow Hallway, randomly initialized agent performs better than DDPG pre-training, which is probably because running fast in Narrow Hallway requires running in a very narrow direction, and DDPG pre-trained agent needs to first unlearn the overtly specialized running gait acquired from pre-training. In Wide Hallway, randomly initialized agent easily gets stuck in a locally optimal gait (running between two opposing walls) while pre-training in general helps avoid such problem.

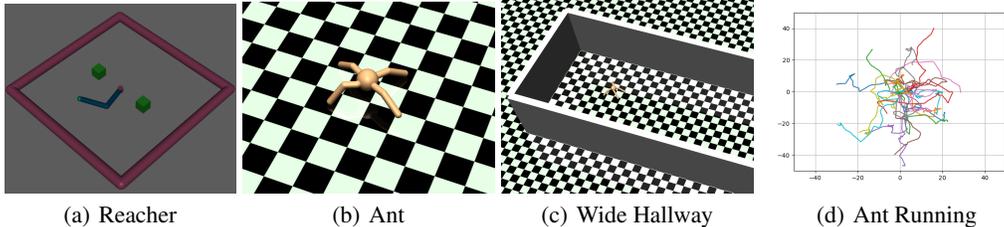


Figure 3: Illustration of locomotion tasks: (a) Bimodal Reacher. Train a reacher to reach one of two targets. Green boxes are targets. (b) Ant-Running. Train a quadrupedal robot to run fast. The instant reward is the robot’s center of mass velocity; (c) Ant-Hallway. Train a quadrupedal robot to run fast under the physical constraints of walls, the instant reward is the same as in (b). Narrow and Wide Hallway tasks differ by the distance between the opposing walls; (d) Trajectories by NBP agent in Ant-Running. The agent learns to run in multiple directions.

Tasks	Random init	DDPG init	SQL init	NBP init
Wide Hallway	522 ± 111	3677 ± 472	3624 ± 312	4306 ± 571
Narrow Hallway	3023 ± 280	2866 ± 248	3026 ± 352	3752 ± 408

Table 2: A comparison of downstream fine-tuning under different initializations. For each task, we show the cumulative rewards after pre-training for $2 \cdot 10^6$ steps and fine-tuning for 10^6 steps. The rewards are shown in the form (mean ± std), all results are averaged over 5 seeds. Random init means the agent is trained from scratch.

Combining multiple modes by Imitation Learning. We propose another paradigm that can be of practical interest. In general, learning a multi-modal policy from scratch is hard for complex tasks since it requires good exploration and an algorithm to learn multi-modal distributions [13], which is itself a hard inference problem [10]. A big advantage of policy based algorithm over value based algorithm [13] is that the policy can be easily combined with imitation learning. We could decompose a complex task into several simpler tasks, each representing a simple mode of behavior easily learned by a RL agent, then combine them into a single agent using imitation learning or inverse RL [1, 8, 28].

We illustrate with a stochastic Swimmer example (see Appendix E.3). Consider training a Swimmer to move fast either forward or backward. The aggregate behavior has two modes and it is easy to solve each single mode. We train two separate Swimmers to move forward/backward and generate expert trajectories using the trained agents. We then train a NBP / NFP agent using GAN [10] / maximum likelihood estimation to combine both modes. Training with the same algorithms, a unimodal policy either commits to only one mode or learns a policy that puts large probability mass between the two modes [18, 10], which greatly deviates from the expert policy. On the contrary, expressive policies can more flexibly incorporate multiple modes into a single agent.

6 Conclusion

We have proposed *Implicit Policy*, a rich class of policy that can represent complex action distributions. We have derived efficient algorithms to compute entropy regularized policy gradients for generic implicit policies. Importantly, we have also showed that entropy regularization maximizes a lower bound of maximum entropy objective, which implies that in practice entropy regularization + rich policy class can lead to desired properties of maximum entropy RL. We have empirically showed that implicit policy achieves competitive performance on benchmark tasks, is more robust to observational noise, and can flexibly represent multi-modal distributions.

Acknowledgements. This research was supported by an Amazon Research Award (2017) and AWS cloud credits. The authors would like to thank Jalaj Bhandari for helpful discussions, and Sergey Levine for helpful comments on early stage experiments of the paper.

References

- [1] Abbeel, P. and Ng, A. Y. (2004). Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, page 1. ACM.
- [2] Asadi, K. and Littman, M. L. (2017). An alternative softmax operator for reinforcement learning. In *International Conference on Machine Learning*, pages 243–252.
- [3] Degris, T., White, M., and Sutton, R. S. (2012). Off-policy actor-critic. *arXiv preprint arXiv:1205.4839*.
- [4] Dhariwal, P., Hesse, C., Klimov, O., Nichol, A., Plappert, M., Radford, A., Schulman, J., Sidor, S., and Wu, Y. (2017). Openai baselines. <https://github.com/openai/baselines>.
- [5] Dinh, L., Krueger, D., and Bengio, Y. (2014). Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*.
- [6] Dinh, L., Sohl-Dickstein, J., and Bengio, S. (2016). Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*.
- [7] Duan, Y., Chen, X., Houthoofd, R., Schulman, J., and Abbeel, P. (2016). Benchmarking deep reinforcement learning for continuous control. In *International Conference on Machine Learning*, pages 1329–1338.
- [8] Finn, C., Christiano, P., Abbeel, P., and Levine, S. (2016). A connection between generative adversarial networks, inverse reinforcement learning, and energy-based models. *arXiv preprint arXiv:1611.03852*.
- [9] Fortunato, M., Azar, M. G., Piot, B., Menick, J., Osband, I., Graves, A., Mnih, V., Munos, R., Hassabis, D., Pietquin, O., et al. (2017). Noisy networks for exploration. *arXiv preprint arXiv:1706.10295*.
- [10] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680.
- [11] Haarnoja, T., Hartikainen, K., Abbeel, P., and Levine, S. (2018a). Latent space policies for hierarchical reinforcement learning. *arXiv preprint arXiv:1804.02808*.
- [12] Haarnoja, T., Pong, V., Zhou, A., Dalal, M., Abbeel, P., and Levine, S. (2018b). Composable deep reinforcement learning for robotic manipulation. *arXiv preprint arXiv:1803.06773*.
- [13] Haarnoja, T., Tang, H., Abbeel, P., and Levine, S. (2017). Reinforcement learning with deep energy-based policies. *arXiv preprint arXiv:1702.08165*.
- [14] Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. (2018c). Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*.
- [15] Kakade, S. and Langford, J. (2002). Approximately optimal approximate reinforcement learning. In *ICML*, volume 2, pages 267–274.
- [16] Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- [17] Kingma, D. P. and Welling, M. (2013). Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- [18] Levine, S. (2018). Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909*.
- [19] Li, Y. and Turner, R. E. (2017). Gradient estimators for implicit models. *arXiv preprint arXiv:1705.07107*.
- [20] Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. (2015). Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*.

- [21] Liu, Q. and Wang, D. (2016). Stein variational gradient descent: A general purpose bayesian inference algorithm. In *Advances In Neural Information Processing Systems*, pages 2378–2386.
- [22] Mania, H., Guy, A., and Recht, B. (2018). Simple random search provides a competitive approach to reinforcement learning. *arXiv preprint arXiv:1803.07055*.
- [23] Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., and Kavukcuoglu, K. (2016). Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, pages 1928–1937.
- [24] Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. (2013). Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.
- [25] Nachum, O., Norouzi, M., Xu, K., and Schuurmans, D. (2017). Bridging the gap between value and policy based reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 2772–2782.
- [26] O’Donoghue, B., Munos, R., Kavukcuoglu, K., and Mnih, V. (2016). Pqg: Combining policy gradient and q-learning. *arXiv preprint arXiv:1611.01626*.
- [27] Rezende, D. J. and Mohamed, S. (2015). Variational inference with normalizing flows. *arXiv preprint arXiv:1505.05770*.
- [28] Ross, S., Gordon, G., and Bagnell, D. (2011). A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635.
- [29] Schulman, J., Chen, X., and Abbeel, P. (2017a). Equivalence between policy gradients and soft q-learning. *arXiv preprint arXiv:1704.06440*.
- [30] Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. (2015). Trust region policy optimization. In *International Conference on Machine Learning*, pages 1889–1897.
- [31] Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017b). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- [32] Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. (2016). Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489.
- [33] Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., and Riedmiller, M. (2014). Deterministic policy gradient algorithms. In *ICML*.
- [34] Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958.
- [35] Sutton, R. S., McAllester, D. A., Singh, S. P., and Mansour, Y. (2000). Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pages 1057–1063.
- [36] Todorov, E., Erez, T., and Tassa, Y. (2012). Mujoco: A physics engine for model-based control. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pages 5026–5033. IEEE.
- [37] Tran, D., Ranganath, R., and Blei, D. M. (2017). Hierarchical implicit models and likelihood-free variational inference. *arXiv preprint arXiv:1702.08896*.
- [38] Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. In *Reinforcement Learning*, pages 5–32. Springer.
- [39] Ziebart, B. D. (2010). *Modeling purposeful adaptive behavior with the principle of maximum causal entropy*. Carnegie Mellon University.

A Proof of Theorems

A.1 Stochastic Pathwise Gradient

Theorem 3.1 (Stochastic Pathwise Gradient). *Given an implicit stochastic policy $a_t = f_\theta(s_t, \epsilon)$, $\epsilon \sim \rho_0(\cdot)$. Let π_θ be the implicitly defined policy. Then the pathwise policy gradient for the stochastic policy is*

$$\nabla_\theta J(\pi_\theta) = \mathbb{E}_{\pi_\theta} \left[\mathbb{E}_{\epsilon \sim \rho_0(\cdot)} [\nabla_\theta f_\theta(s, \epsilon) \nabla_a Q^{\pi_\theta}(s, a)|_{a=f_\theta(s, \epsilon)}] \right]. \quad (6)$$

Proof. We follow closely the derivation of deterministic policy gradient [33]. We assume that all conditions are satisfied to exchange expectations and gradients when necessary. Let $\pi = \pi_\theta$ denote the implicit policy $a_t = f_\theta(s_t, \epsilon)$, $\epsilon \sim \rho_0(\cdot)$. Let V^π, Q^π be the value function and action value function under such stochastic policy. We introduce $p(s \rightarrow s', k, \pi)$ as the probability of transitioning from s to s' in k steps under policy π . Overloading the notation a bit, $p(s \rightarrow s', 1, a)$ is the probability of $s \rightarrow s'$ in one step by taking action a (i.e., $p(s \rightarrow s', 1, a) = p(s'|s, a)$). We have

$$\begin{aligned} & \nabla_\theta V^\pi(s) \\ &= \nabla_\theta \mathbb{E}_{a \sim \pi(\cdot|s)} [Q^\pi(s, a)] = \nabla_\theta \mathbb{E}_{\epsilon \sim \rho_0(\cdot)} [Q^\pi(s, f_\theta(s, \epsilon))] \\ &= \nabla_\theta \mathbb{E}_{\epsilon \sim \rho_0(\cdot)} [r(s, f_\theta(s, \epsilon)) + \int_{\mathcal{S}} \gamma p(s'|s, f_\theta(s, \epsilon)) V^\pi(s') ds'] \\ &= \mathbb{E}_{\epsilon \sim \rho_0(\cdot)} [\nabla_\theta r(s, f_\theta(s, \epsilon)) + \nabla_\theta \int_{\mathcal{S}} \gamma p(s'|s, f_\theta(s, \epsilon)) V^\pi(s') ds'] \\ &= \mathbb{E}_{\epsilon \sim \rho_0(\cdot)} [\nabla_\theta r(s, f_\theta(s, \epsilon)) + \int_{\mathcal{S}} \gamma V^\pi(s') \nabla_\theta p(s'|s, f_\theta(s, \epsilon)) ds' + \int_{\mathcal{S}} \gamma p(s'|s, f_\theta(s, \epsilon)) \nabla_\theta V^\pi(s') ds'] \\ &= \mathbb{E}_{\epsilon \sim \rho_0(\cdot)} [\nabla_\theta f_\theta(s, \epsilon) \nabla_a [r(s, a) + \gamma \int_{\mathcal{S}} \gamma p(s'|s, a) V^\pi(s') ds']|_{a=f_\theta(s, \epsilon)} \\ &\quad + \int_{\mathcal{S}} \gamma p(s'|s, f_\theta(s, \epsilon)) \nabla_\theta V^\pi(s') ds'] \\ &= \mathbb{E}_{\epsilon \sim \rho_0(\cdot)} [\nabla_\theta f_\theta(s, \epsilon) \nabla_a Q^\pi(s, a)|_{a=f_\theta(s, \epsilon)}] + \mathbb{E}_{\epsilon \sim \rho_0(\cdot)} \left[\int_{\mathcal{S}} \gamma p(s \rightarrow s', 1, f_\theta(s, \epsilon)) \nabla_\theta V^\pi(s') ds' \right]. \end{aligned}$$

In the above derivation, we have used the Fubini theorem to interchange integral (expectation) and gradients. We can iterate the above derivation and have the following

$$\begin{aligned} & \nabla_\theta V^\pi(s) \\ &= \nabla_\theta \mathbb{E}_{a \sim \pi(\cdot|s)} [Q^\pi(s, a)] \\ &= \mathbb{E}_{\epsilon \sim \rho_0(\cdot)} [\nabla_\theta f_\theta(s, \epsilon) \nabla_a Q^\pi(s, a)|_{a=f_\theta(s, \epsilon)}] + \mathbb{E}_{\epsilon \sim \rho_0(\cdot)} [\gamma p(s \rightarrow s', 1, f_\theta(s, \epsilon)) \nabla_\theta V^\pi(s') ds'] \\ &= \mathbb{E}_{\epsilon \sim \rho_0(\cdot)} [\nabla_\theta f_\theta(s, \epsilon) \nabla_a Q^\pi(s, a)|_{a=f_\theta(s, \epsilon)}] + \\ &\quad \mathbb{E}_{\epsilon \sim \rho_0(\cdot)} \left[\int_{\mathcal{S}} \gamma p(s \rightarrow s', 1, f_\theta(s, \epsilon)) \mathbb{E}_{\epsilon' \sim \rho_0(\cdot)} [\nabla_\theta f_\theta(s', \epsilon') \nabla_a Q^\pi(s', a')|_{a'=f_\theta(s', \epsilon')}] ds' \right] + \\ &\quad \mathbb{E}_{\epsilon \sim \rho_0(\cdot)} \left[\int_{\mathcal{S}} \gamma p(s \rightarrow s', 1, f_\theta(s, \epsilon')) \mathbb{E}_{\epsilon' \sim \rho_0(\cdot)} \left[\int_{\mathcal{S}} \gamma p(s' \rightarrow s'', 1, f_\theta(s', \epsilon')) \nabla_\theta V^\pi(s'') ds'' \right] ds' \right] \\ &= \dots \\ &= \int_{\mathcal{S}} \sum_{t=0}^{\infty} \gamma^t p(s \rightarrow s', t, \pi) \mathbb{E}_{\epsilon' \sim \rho_0(\cdot)} [\nabla_\theta f_\theta(s', \epsilon') \nabla_a Q^\pi(s', a')|_{a'=f_\theta(s', \epsilon')}] ds'. \end{aligned}$$

With the above, we derive the pathwise policy gradient as follows

$$\begin{aligned}
\nabla_{\theta} J(\pi_{\theta}) &= \nabla_{\theta} \int_{\mathcal{S}} p_1(s) V^{\pi}(s) ds \\
&= \int_{\mathcal{S}} p_1(s) \nabla_{\theta} V^{\pi}(s) ds \\
&= \int_{\mathcal{S}} \int_{\mathcal{S}} \sum_{t=0}^{\infty} \gamma^t p_1(s) p(s \rightarrow s', t, \pi) ds \mathbb{E}_{\epsilon' \sim \rho_0(\cdot)} [\nabla_{\theta} f_{\theta}(s', \epsilon') \nabla_a Q^{\pi}(s', a') |_{a'=f_{\theta}(s, \epsilon')}] ds' \\
&= \int_{\mathcal{S}} \rho_{\pi}(s') \mathbb{E}_{\epsilon' \sim \rho_0(\cdot)} [\nabla_{\theta} f_{\theta}(s', \epsilon') \nabla_a Q^{\pi}(s', a') |_{a'=f_{\theta}(s, \epsilon')}] ds',
\end{aligned}$$

where $\rho_{\pi}(s') = \int_{\mathcal{S}} \sum_{t=0}^{\infty} \gamma^t p_1(s) p(s \rightarrow s', t, \pi) ds$ is the discounted state visitation probability under policy π . Writing the whole integral as an expectation over states, the policy gradient is

$$\nabla_{\theta} J(\pi_{\theta}) = \mathbb{E}_{s \sim \rho_{\pi}(s)} [\mathbb{E}_{\epsilon \sim \rho_0(\cdot)} [\nabla_{\theta} f_{\theta}(s, \epsilon) \nabla_a Q^{\pi}(s, a) |_{a=f_{\theta}(s, \epsilon)}]].$$

which is equivalent to in (6) in theorem 3.1. \square

We can recover the result for deterministic policy gradient by using a degenerate functional form $f_{\theta}(s, \epsilon) = f_{\theta}(s)$, i.e. with a deterministic function to compute actions.

A.2 Unbiased Entropy Gradient

Lemma A.1 (Optimal Classifier as Density Estimator). *Assume c_{ψ} is expressive enough to represent any classifier (for example c_{ψ} is a deep neural net). Assume \mathcal{A} to be bounded and let $U(\mathcal{A})$ be uniform distribution over \mathcal{A} . Let ψ^* be the optimizer to the optimization problem in (7). Then $c_{\psi^*}(s, a) = \log \frac{\pi_{\theta}(a|s)}{|\mathcal{A}|^{-1}}$ and $|\mathcal{A}|$ is the volume of \mathcal{A} .*

Proof. Observe that (7) is a binary classification problem with data from $a \sim \pi_{\theta}(\cdot|s)$ against $a \sim U(\mathcal{A})$. The optimal classifier of the problem produces the density ratio of these two distributions. See for example [10] for a detailed proof. \square

Theorem 3.2 (Unbiased Entropy Gradient). *Let ψ^* be the optimal solution from (7), where the policy $\pi_{\theta}(\cdot|s)$ is given by implicit policy $a = f_{\theta}(s, \epsilon), \epsilon \sim \rho_0(\cdot)$. The gradient of entropy $\nabla_{\theta} \mathbb{H}[\pi_{\theta}(\cdot|s)]$ can be computed as*

$$\nabla_{\theta} \mathbb{H}[\pi_{\theta}(\cdot|s)] = -\mathbb{E}_{\epsilon \sim \rho_0(\cdot)} [\nabla_{\theta} c_{\psi^*}(f(\theta, \epsilon), s)] \quad (8)$$

Proof. Let $\pi_{\theta}(\cdot|s)$ be the density of implicit policy $a = f_{\theta}(\epsilon, s), \epsilon \sim \rho_0(\cdot)$. The entropy is computed as follows

$$\mathbb{H}[\pi_{\theta}(\cdot|s)] = -\mathbb{E}_{a \sim \pi_{\theta}(\cdot|s)} [\log \pi_{\theta}(a|s)].$$

Computing its gradient

$$\begin{aligned}
\nabla_{\theta} \mathbb{H}[\pi_{\theta}(\cdot|s)] &= -\nabla_{\theta} \mathbb{E}_{\epsilon \sim \rho_0(\cdot)} [\log \pi_{\theta}(f_{\theta}(s, \epsilon)|s)] \\
&= -\mathbb{E}_{a \sim \pi_{\theta}(\cdot|s)} [\nabla_{\theta} \log \pi_{\theta}(a|s)] - \mathbb{E}_{\epsilon \sim \rho_0(\cdot)} [\nabla_a \log \pi_{\theta}(a|s) |_{a=f_{\theta}(s, \epsilon)} \nabla_{\theta} f_{\theta}(s, \epsilon)] \\
&= -\mathbb{E}_{\epsilon \sim \rho_0(\cdot)} [\nabla_a \log \pi_{\theta}(a|s) |_{a=f_{\theta}(s, \epsilon)} \nabla_{\theta} f_{\theta}(s, \epsilon)] \\
&= -\mathbb{E}_{\epsilon \sim \rho_0(\cdot)} [\nabla_a c_{\psi^*}(f(\theta, \epsilon), s) \nabla_{\theta} f_{\theta}(s, \epsilon)] \quad (11)
\end{aligned}$$

In the second line we highlight the fact that the expectation depends on parameter θ both implicitly through the density π_{θ} and through the sample $f_{\theta}(s, \epsilon)$. After decomposing the gradient using chain rule, we find that the first term vanishes, leaving the result shown in the theorem. \square

A.3 Lower Bound

We recall that given a policy π , the standard RL objective is $J(\pi) = \mathbb{E}_\pi [\sum_{t=0}^{\infty} \gamma^t r_t]$. In maximum entropy formulation, the maximum entropy objective is

$$J_{\text{MaxEnt}}(\pi) = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t (r_t + \beta \mathbb{H}[\pi(\cdot|s_t)]) \right], \quad (12)$$

where $\beta > 0$ is a regularization constant and $\mathbb{H}(\pi(\cdot|s_t))$ is the entropy of policy π at s_t . We construct a surrogate objective based on another policy $\tilde{\pi}$ as follows

$$J_{\text{surr}}(\pi, \tilde{\pi}) = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t r_t \right] + \beta \mathbb{E}_{\tilde{\pi}} \left[\sum_{t=0}^{\infty} \gamma^t \mathbb{H}[\pi(\cdot|s_t)] \right]. \quad (13)$$

The following proof highly mimics the proof in [30]. We have the following definition for coupling two policies

Definition A.1 (α -coupled). *Two policies $\pi, \tilde{\pi}$ are α -coupled if $\mathbb{P}(\pi(\cdot|s) \neq \tilde{\pi}(\cdot|s)) \leq \alpha$ for any $s \in \mathcal{S}$.*

Lemma A.2. *Given $\pi, \tilde{\pi}$ are α -coupled, then*

$$|\mathbb{E}_{s_t \sim \pi} [\mathbb{H}[\pi(\cdot|s_t)]] - \mathbb{E}_{s_t \sim \tilde{\pi}} [\mathbb{H}[\tilde{\pi}(\cdot|s_t)]]| \leq 2(1 - (1 - \alpha)^t) \max_s |\mathbb{H}[\pi(\cdot|s)]|.$$

Proof. Let n_t denote the number of times that $a_i \neq \tilde{a}_i$ for $i < t$, i.e. the number of times that $\pi, \tilde{\pi}$ disagree before time t . We can decompose the expectations as follows

$$\begin{aligned} \mathbb{E}_{s_t \sim \pi} [\mathbb{H}[\pi(\cdot|s_t)]] &= \mathbb{E}_{s_t \sim \pi | n_t=0} [\mathbb{H}[\pi(\cdot|s_t)]] \mathbb{P}(n_t = 0) + \mathbb{E}_{s_t \sim \pi | n_t>0} [\mathbb{H}[\pi(\cdot|s_t)]] \mathbb{P}(n_t > 0), \\ \mathbb{E}_{s_t \sim \tilde{\pi}} [\mathbb{H}[\tilde{\pi}(\cdot|s_t)]] &= \mathbb{E}_{s_t \sim \tilde{\pi} | n_t=0} [\mathbb{H}[\tilde{\pi}(\cdot|s_t)]] \mathbb{P}(n_t = 0) + \mathbb{E}_{s_t \sim \tilde{\pi} | n_t>0} [\mathbb{H}[\tilde{\pi}(\cdot|s_t)]] \mathbb{P}(n_t > 0). \end{aligned}$$

Note that $n_t = 0$ implies $a_i = \tilde{a}_i$ for all $i < t$ hence

$$\mathbb{E}_{s_t \sim \tilde{\pi} | n_t=0} [\mathbb{H}[\pi(\cdot|s_t)]] = \mathbb{E}_{s_t \sim \pi | n_t=0} [\mathbb{H}[\pi(\cdot|s_t)]].$$

The definition of α -coupling implies $\mathbb{P}(n_t = 0) \geq (1 - \alpha)^t$, and so $\mathbb{P}(n_t > 0) \leq 1 - (1 - \alpha)^t$. Now we note that

$$|\mathbb{E}_{s_t \sim \pi | n_t>0} [\mathbb{H}[\pi(\cdot|s_t)]] \mathbb{P}(n_t > 0) - \mathbb{E}_{s_t \sim \tilde{\pi} | n_t>0} [\mathbb{H}[\tilde{\pi}(\cdot|s_t)]] \mathbb{P}(n_t > 0)| \leq 2(1 - (1 - \alpha)^t) \max_s |\mathbb{H}[\pi(\cdot|s)]|.$$

Combining previous observations, we have proved the lemma. \square

Note that if we take $\pi = \pi_\theta, \tilde{\pi} = \pi_{\theta_{\text{old}}}$, then the surrogate objective $J_{\text{surr}}(\pi_\theta)$ in (9) is equivalent to $J_{\text{surr}}(\pi, \tilde{\pi})$ defined in (13). With lemma A.2, we prove the following theorem.

Theorem 4.1 (Lower Bound). *If $\mathbb{KL}[\pi_\theta || \pi_{\theta_{\text{old}}}] \leq \alpha$, then*

$$J_{\text{MaxEnt}}(\pi) \geq J_{\text{surr}}(\pi) - \frac{\beta \gamma \sqrt{\alpha} \epsilon}{(1 - \gamma)^2}, \quad \text{where } \epsilon = \max_s |\mathbb{H}[\pi(\cdot|s)]|. \quad (10)$$

Proof. We first show the result for general policies π and $\tilde{\pi}$ with $\mathbb{KL}[\pi || \tilde{\pi}] \leq \alpha$. As a result, [30] shows that $\pi, \tilde{\pi}$ are $\sqrt{\alpha}$ -coupled. Recall the maximum entropy objective $J_{\text{MaxEnt}}(\pi)$ defined in (12) and surrogate objective in (13), take the difference of two objectives

$$\begin{aligned} |J_{\text{MaxEnt}}(\pi) - J_{\text{surr}}(\pi, \tilde{\pi})| &= \beta \left| \sum_{t=0}^{\infty} \gamma^t \mathbb{E}_{s_t \sim \pi} [\mathbb{H}[\pi(\cdot|s_t)]] - \mathbb{E}_{s_t \sim \tilde{\pi}} [\mathbb{H}[\tilde{\pi}(\cdot|s_t)]] \right| \\ &\leq \beta \sum_{t=0}^{\infty} \gamma^t |\mathbb{E}_{s_t \sim \pi} [\mathbb{H}[\pi(\cdot|s_t)]] - \mathbb{E}_{s_t \sim \tilde{\pi}} [\mathbb{H}[\tilde{\pi}(\cdot|s_t)]]| \\ &\leq \beta \sum_{t=0}^{\infty} \gamma^t (1 - (1 - \sqrt{\alpha})^t) \max_s |\mathbb{H}[\pi(\cdot|s)]| \\ &= \left(\frac{1}{1 - \gamma} - \frac{1}{1 - \gamma(1 - \sqrt{\alpha})} \right) \beta \sqrt{\alpha} \max_s |\mathbb{H}[\pi(\cdot|s)]| \\ &= \frac{\beta \gamma \sqrt{\alpha}}{(1 - \gamma)^2} \max_s |\mathbb{H}[\pi(\cdot|s)]| \end{aligned}$$

Now observe that by taking $\pi = \pi_\theta$, $\tilde{\pi} = \pi_{\theta_{\text{old}}}$, the above inequality implies the theorem. \square

In practice, α -coupling enforced by KL divergence is often relaxed [30, 31]. The theorem implies that, by constraining the KL divergence between consecutive policy iterates, the surrogate objective of entropy regularization maximizes a lower bound of maximum entropy objective.

B Operator view of Entropy Regularization and Maximum Entropy RL

Recall in standard RL formulation, the agent is in state s , takes action a , receives reward r and transitions to s' . Let the discount factor $\gamma < 1$. Assume that the reward r is deterministic and the transitions $s' \sim p(\cdot|s, a)$ are deterministic, i.e. $s' = f(s, a)$, it is straightforward to extend the following to general stochastic transitions. For a given policy π , define linear Bellman operator as

$$\mathcal{T}^\pi Q(s, a) = r + \gamma \mathbb{E}_{a' \sim \pi(\cdot|s')} [Q(s', a')].$$

Any policy π satisfies the linear Bellman equation $\mathcal{T}^\pi Q^\pi = Q^\pi$. Define Bellman optimality operator (we will call it Bellman operator) as

$$\mathcal{T}^* Q(s, a) = r + \gamma \max_{a'} Q(s', a').$$

Now we define Mellowmax operator [2, 13] with parameter $\beta > 0$ as follows,

$$\mathcal{T}_s Q(s, a) = r + \gamma \beta \log \int_{a' \in \mathcal{A}} \exp\left(\frac{Q(s', a')}{\beta}\right) da'.$$

It can be shown that both \mathcal{T}^* and \mathcal{T} are contractive operator when $\gamma < 1$. Let Q^* be the unique fixed point of $\mathcal{T}^* Q = Q$, then Q^* is the action value function of the optimal policy $\pi^* = \arg \max_\pi J(\pi)$. Let Q_s^* be the unique fixed point of $\mathcal{T}_s Q = Q$, then Q_s^* is the soft action value function of $\pi_s^* = \arg \max_\pi J_{\text{MaxEnt}}(\pi)$. In addition, the optimal policy $\pi^*(\cdot|s) = \arg \max_a Q^*(s, a)$ and $\pi_s^*(a|s) \propto \exp\left(\frac{Q_s^*(s, a)}{\beta}\right)$.

Define Boltzmann operator with parameter β as follows

$$\mathcal{T}_B Q(s, a) = r + \gamma \mathbb{E}_{a' \sim p_B(\cdot|s')} [Q(s', a')],$$

where $p_B(a'|s') \propto \exp\left(\frac{Q(s', a')}{\beta}\right)$ is the Boltzmann distribution defined by $Q(s', a')$. [26] shows that the stationary points of entropy regularization procedure are policies of the form $\pi(a|s) \propto \exp\left(\frac{Q^\pi(s, a)}{\beta}\right)$. We illustrate the connection between such stationary points and fixed points of Boltzmann operator as follows.

Theorem B.1 (Fixed points of Boltzmann Operators). *Any fixed point $Q(s, a)$ of Boltzmann operator $\mathcal{T}_B Q = Q$, defines a stationary point for entropy regularized policy gradient algorithm by $\pi(a|s) \propto \exp\left(\frac{Q(s, a)}{\beta}\right)$; reversely, any stationary point of entropy regularized policy gradient algorithm π , has its action value function $Q^\pi(s, a)$ as a fixed point to Boltzmann operator.*

Proof. Take any fixed point Q of Boltzmann operator, $\mathcal{T}_B Q = Q$, define a policy $\pi(a|s) \propto \exp\left(\frac{Q(s, a)}{\beta}\right)$. From the definition of Boltzmann operator, we can easily check that π 's entropy regularized policy gradient is exactly zero, hence it is a stationary point for entropy regularized gradient algorithm.

Take any policy π such that its entropy regularized gradient is zero, from [26] we know for such policy $\pi(a|s) \propto \exp\left(\frac{Q^\pi(s, a)}{\beta}\right)$. The linear Bellman equation for such a policy $\mathcal{T}^\pi Q^\pi = Q^\pi$ translates directly into the Boltzmann equation $\mathcal{T}_B Q^\pi = Q^\pi$. Hence Q^π is indeed a fixed point of Boltzmann operator. \square

The above theorem allows us to associate the policies trained by entropy regularized policy gradient with the Boltzmann operator. Unfortunately, [2] shows that unlike MellowMax operator, Boltzmann operator does not have unique fixed point and is not a contractive operator in general, though this does not necessarily prevent policy gradient algorithms from converging. We make a final observation that

shows that Boltzmann operator \mathcal{T}_B interpolates Bellman operator \mathcal{T}^* and MellowMax operator \mathcal{T}_s : for any Q and fixed $\beta > 0$ (see theorem B.2),

$$\mathcal{T}^*Q \geq \mathcal{T}_BQ \geq \mathcal{T}_sQ. \quad (14)$$

If we view all operators as picking out the largest value among $Q(s', a')$, $a' \in \mathcal{A}$ in next state s' , then \mathcal{T}^* is the greediest and \mathcal{T}_s is the most conservative, as it incorporates trajectory entropy as part of the objective. \mathcal{T}_B is between these two operators, since it looks ahead for only one step. The first inequality in (14) is trivial, now we show the second inequality.

Theorem B.2 (Boltzmann Operator is greedier than Mellowmax Operator). *For any $Q \in \mathbb{R}^n$, we have $\mathcal{T}_BQ \geq \mathcal{T}_sQ$, and the equality is tight if and only if $Q_i = Q_j$ for $\forall i, j$.*

Proof. Recall the definition of both operators, we essentially need to show the following inequality

$$\mathbb{E}_{a' \sim p_B(\cdot|s')} [Q(s', a')] \geq \beta \log \int_{a' \in \mathcal{A}} \exp\left(\frac{Q(s', a')}{\beta}\right).$$

Without loss of generality, assume there are n actions in total and let $x_i = \exp(\frac{Q(s', a_i)}{\beta})$ for the i th action. The above inequality reduces to

$$\frac{\sum_{i=1}^n x_i \log x_i}{\sum_{i=1}^n x_i} \geq \log \frac{1}{n} \sum_{i=1}^n x_i.$$

We did not find any reference for the above inequality so we provide a proof below. Notice that $x_i > 0, \forall i$. Introduce the objective $J(x) = \frac{\sum_{i=1}^n x_i \log x_i}{\sum_{i=1}^n x_i} - \log \frac{1}{n} \sum_{i=1}^n x_i$. Compute the gradient of $J(x)$,

$$\frac{\partial J(x)}{\partial x_j} = \frac{\sum_{i=1}^n (\log x_j - \log x_i) x_i}{(\sum_{i=1}^n x_i)^2}, \forall j.$$

The stationary point at which $\frac{\partial J(x)}{\partial x} = 0$ is of the form $x_i = x_j, \forall i, j$. At such point, let $x_i = x_0, \forall i$ for some generic x_0 . Then we compute the Hessian of $J(x)$ at such stationary point

$$\begin{aligned} \frac{\partial^2 J(x)}{\partial x_j^2} \Big|_{x_i=x_0, \forall i} &= \frac{n-1}{n^2 x_0^2}, \forall j. \\ \frac{\partial^2 J(x)}{\partial x_k \partial x_j} \Big|_{x_i=x_0, \forall i} &= \frac{1}{n^2 x_0^2}, \forall j \neq k. \end{aligned}$$

Let $H(x_0)$ be the Hessian at this stationary point. Let $t \in \mathbb{R}^n$ be any vector and we can show

$$t^T H(x_0) t = \sum_{i \neq j} \frac{1}{n^2 x_0^2} (t_i - t_j)^2 \geq 0,$$

which implies that $H(x_0)$ is positive semi-definite. It is then implied that at such x_0 we will achieve local minimum. Let $x_i = x_0, \forall i$ we find $J(x) = 0$, which implies that $J(x) \geq 0, \forall x$. Hence the proof is concluded. \square

C Implicit Policy Architecture

C.1 Normalizing Flows Policy Architecture

We design the neural network architectures following the idea of [5, 6]. Recall that Normalizing Flows [27] consist of layers of transformations as follows ,

$$x = g_{\theta_K} \circ g_{\theta_{K-1}} \circ \dots \circ g_{\theta_2} \circ g_{\theta_1}(\epsilon),$$

where each $g_{\theta_i}(\cdot)$ is an invertible transformation. We focus on how to design each atomic transformation $g_{\theta_i}(\cdot)$. We overload the notations and let x, y be the input/output of a generic layer $g_{\theta}(\cdot)$,

$$y = g_{\theta}(x).$$

We design a generic transformation $g_\theta(\cdot)$ as follows. Let x_I be the components of x corresponding to subset indices $I \subset \{1, 2, \dots, m\}$. Then we propose as in [6],

$$\begin{aligned} y_{1:d} &= x_{1:d} \\ y_{d+1:m} &= x_{d+1:m} \odot \exp(s(x_{1:d})) + t(x_{1:d}), \end{aligned} \quad (15)$$

where $t(\cdot), s(\cdot)$ are two arbitrary functions $t, s : \mathbb{R}^d \mapsto \mathbb{R}^{m-d}$. It can be shown that such transformation entails a simple Jacobian matrix $|\frac{\partial y}{\partial x^T}| = \exp(\sum_{j=1}^{m-d} [s(x_{1:d})]_j)$ where $[s(x_{1:d})]_j$ refers to the j th component of $s(x_{1:d})$ for $1 \leq j \leq m-d$. For each layer, we can permute the input x before apply the simple transformation (15) so as to couple different components across layers. Such coupling entails a complex transformation when we stack multiple layers of (15). To define a policy, we need to incorporate state information. We propose to preprocess the state $s \in \mathbb{R}^n$ by a neural network $L_{\theta_s}(\cdot)$ with parameter θ_s , to get a state vector $L_{\theta_s}(s) \in \mathbb{R}^m$. Then combine the state vector into (15) as follows,

$$\begin{aligned} z_{1:d} &= x_{1:d} \\ z_{d+1:m} &= x_{d+1:m} \odot \exp(s(x_{1:d})) + t(x_{1:d}) \\ y &= z + L_{\theta_s}(s). \end{aligned} \quad (16)$$

It is obvious that $x \leftrightarrow y$ is still bijective regardless of the form of $L_{\theta_s}(\cdot)$ and the Jacobian matrix is easy to compute accordingly.

In our experiments, we implement s, t both as 4-layers neural networks with $k = 3$ or $k = 6$ units per hidden layer. We stack $K = 4$ transformations: we implement (16) to inject state information only after the first transformation, and the rest is conventional coupling as in (15). $L_{\theta_s}(s)$ is implemented as a feedforward neural network with 2 hidden layers each with 64 hidden units. Value function critic is implemented as a feedforward neural network with 2 hidden layers each with 64 hidden units with rectified-linear between hidden layers.

C.2 Non-Invertible Blackbox Policy Architecture

Any implicit model architecture as in [10, 37] can represent a Non-invertible Blackbox Policy (NBP). On MuJoCo control tasks, consider a task with state space $\mathcal{S} \subset \mathbb{R}^n$ and action space $\mathcal{A} \subset \mathbb{R}^m$. Consider a feedforward neural network with n input units and m output units. The intermediate layers have parameters θ and the output is a deterministic mapping from the input $a = f_\theta(s)$. We choose an architecture similar to NoisyNet [9]: introduce a distribution over θ . In our case, we choose factorized Gaussian $\theta = \mu_\theta + \sigma_\theta \cdot \epsilon$. The implicit policy is generated as

$$a = f_\theta(s), \theta = \mu_\theta + \sigma_\theta \cdot \epsilon, \epsilon \sim \mathcal{N}(0, 1),$$

which induces an implicit distribution over output a . In practice, we find randomizing parameters θ to generate implicit policy works well and is easy to implement, we leave other approaches for future research.

In all experiments, we implement the network $f_\theta(\cdot)$ as a feedforward neural network with 2 hidden layers each with 64 hidden units. Between layers we use rectified-linear for non-linear activation, layer normalization to standardize inputs, and dropout before the last output. Both value function critic and classifier critic are implemented as feedforward neural networks with 2 hidden layers each with 64 hidden units with rectified-linear between hidden layers. Note that $\mu_\theta, \sigma_\theta$ are the actual parameters of the model: we initialize μ_θ using standard initialization method and initialize $\sigma_\theta = \log(\exp(\rho_\theta) + 1)$ with $\rho_\theta \in [-9.0, -1.0]$. For simplicity, we set all ρ_θ to be the same and let $\rho = \rho_\theta$. We show below that dropout is an efficient technique to represent multi-modal policy.

Dropout for multi-modal distributions. Dropout [34] is an efficient technique to regularize neural networks in supervised learning. However, in reinforcement learning where overfitting is not a big issue, the application of dropout seems limited. Under the framework of implicit policy, we want to highlight that dropout serves as a natural method to parameterize multi-modal distributions. Consider a feed-forward neural network with output $y \in \mathbb{R}^m$. Assume that the last layer is a fully-connected network with h inputs. Let $x \in \mathbb{R}^n$ be an input to the original neural network and $\phi(x) \in \mathbb{R}^h$ be the input to the last layer (we get $\phi(x)$ by computing forward pass of x through the network until the last layer), where $\phi(x)$ can be interpreted as a representation learned by previous layers.

Let $W \in \mathbb{R}^{m \times h}$, $b \in \mathbb{R}^m$ be the weight matrix and bias vector of the last layer, then the output is computed as (we ignore the non-linear activation at the output)

$$y_i = \sum_{j=1}^h \phi_j(x) W_{ij} + b_i, \forall 1 \leq i \leq m. \quad (17)$$

If dropout is applied to the last layer, let z be the Bernoulli mask i.e. $z_i \sim \text{Bernoulli}(p)$, $1 \leq i \leq h$ where p is the probability for dropping an input to the layer. Then

$$y_i = \sum_{j=1}^h (\phi_j(x) \cdot z_j) W_{ij} + b_i, \forall 1 \leq i \leq m \quad (18)$$

Given an i , if each $\phi_j(x) W_{ij}$ has a different value, their stochastic sum $\sum_{j=1}^h \phi_j(x) \cdot z_j W_{ij}$ in (18) can take up to about 2^h values. Despite some redundancy in these 2^h values, in general y_i in (18) has a multi-modal distribution supported on multiple values. We have hence moved from a unimodal distribution (17) to a multi-modal distribution (18) by adding a simple dropout.

D Algorithm Pseudocode

Below we present the pseudocode for an off-policy algorithm to train NBP. On the other hand, for NFP we can apply any on-policy optimization algorithms [31, 30] and we omit the pseudocode here.

Algorithm 1 Non-invertible Blackbox Policy (NBF) Off-policy update

- 1: INPUT: target parameter update period τ ; learning rate $\alpha_\theta, \alpha_\phi, \alpha_\psi$; entropy regularization constant β .
- 2: INITIALIZE: parameters θ, ϕ, ψ and target network parameters θ^-, ϕ^- ; replay buffer $B \leftarrow \{\}$; step counter $counter \leftarrow 0$.
- 3: **for** $e = 1, 2, 3 \dots E$ **do**
- 4: **while** episode not terminated **do**
- 5: **// Control**
- 6: $counter \leftarrow counter + 1$.
- 7: In state s , sample noise $\epsilon \sim \rho_0(\cdot)$, compute action $a = f_\theta(s, \epsilon)$, transition to s' and receive instant reward r .
- 8: Save experience tuple $\{s, a, r, s'\}$ to buffer B .
- 9: Sample N tuples $D = \{s_j, a_j, r_j, s'_j\}$ from B .
- 10: **// Update Critic**
- 11: Compute TD error as in [24, 32] as follows, where $a'_j = f_{\theta^-}(s'_j, \epsilon_j)$, $\epsilon_j \sim \rho_0(\cdot)$.

$$J_\phi = \frac{1}{N} \sum_{j=1}^N (Q_\phi(s_j, a_j) - r_j - \gamma Q_{\phi^-}(s'_j, a'_j))^2.$$

- 12: Update $\phi \leftarrow \phi - \alpha_\phi \nabla_\phi J_\phi$.
 - 13: **// Update classifier**
 - 14: Sample N actions uniformly from action space $a_j^{(u)} \sim \mathcal{U}(\mathcal{A})$. Compute classification objective C_ψ (7) using data $\{s_j, a_j\}_{j=1}^N$ against $\{s_j, a_j^{(u)}\}_{j=1}^N$.
 - 15: Update classifier $\psi \leftarrow \psi - \alpha_\psi \nabla_\psi C_\psi$.
 - 16: **// Update policy with entropy regularization.**
 - 17: Compute pathway gradient $\nabla_\theta J(\pi_\theta)$ (6) with $Q^\pi(s, a)$ replaced by critic $Q_\phi(s, a)$ and states replaced by sampled states s_j .
 - 18: Compute entropy gradient $\nabla_\theta \mathbb{H}[\pi_\theta(\cdot|s)]$ using (8) on sampled data.
 - 19: Update $\theta \leftarrow \theta + \alpha_\theta (\nabla_\theta J(\pi_\theta) + \beta \nabla_\theta \mathbb{H}[\pi_\theta(\cdot|s)])$
 - 20: **if** $counter \bmod \tau = 0$ **then**
 - 21: Update target parameter $\phi^- \leftarrow \phi, \theta^- \leftarrow \theta$.
 - 22: **end if**
 - 23: **end while**
 - 24: **end for**
-

E Additional Experiment Results

E.1 Locomotion tasks

As has been shown in previous works [31], PPO is almost the most competitive on-policy optimization baseline on locomotion control tasks. We provide a table of comparison among on-policy baselines below. On each task we train for a specified number of time steps and report the average results over 5 random seeds. Though NFP remains a competitive algorithm, PPO with unimodal Gaussian generally achieves better performance.

Tasks	Timesteps	PPO	A2C	CEM	TRPO	NFP
Hopper	10^6	≈ 2300	≈ 900	≈ 500	≈ 2000	≈ 1880
HalfCheetah	10^6	≈ 1900	≈ 1000	≈ 500	≈ 0	≈ 2200
Walker2d	10^6	≈ 3500	≈ 900	≈ 800	≈ 1000	≈ 1980
InvertedDoublePendulum	10^6	≈ 8000	≈ 6500	≈ 0	≈ 0	≈ 8000

Table 1: A comparison of NFP with (on-policy) baseline algorithms on MuJoCo benchmark tasks. For each task, we show the average rewards achieved after training the agent for a fixed number of time steps. The results for NFP are averaged over 5 random seeds. The results for A2C, CEM [7] and TRPO are approximated based on the figures in [31]. PPO is from OpenAI baseline implementation [4]. We highlight the top two algorithms for each task in bold font. PPO, A2C and TRPO all use unimodal Gaussians. PPO is the most competitive.

E.2 Multi-modal policy: Fine-tuning for downstream tasks

In Figure 4, we compare trajectories generated by agents pre-trained by DDPG and NBP on the running task. Since DDPG uses a deterministic policy, starting from a fixed position, the agent can only run in a single direction. On the other hand, NBP agent manages to run in multiple directions. This comparison partially illustrates that a NBP agent can learn the concept of *general running*, instead of *specialized running* – running in a particular direction.

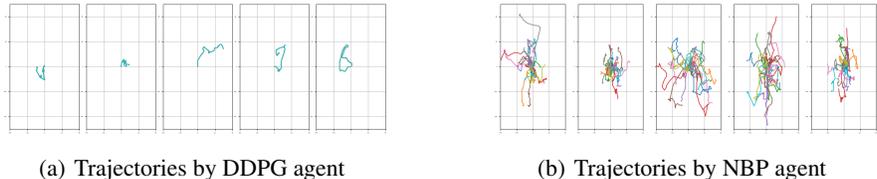


Figure 4: (a)(b): Trajectories generated by DDPG pre-trained agents and NBP pre-trained agent on Ant-Running task under different random seeds. Starting from the initial position, DDPG agent can only produce a single deterministic trajectory due to the deterministic policy; NBP agent produces trajectories that are more diverse, illustrating that NBP agent learns to run in multiple directions.

E.3 Multi-modal policy: Combining multiple modes by Imitation Learning.

Didactic Example. We motivate combining multiple modes of imitation learning with a simple example: imitating an expert with two modes of behavior. Consider a simple MDP on an axis with state space $\mathcal{S} = [-10, 10]$, action space $\mathcal{A} = [-1, 1]$. The agent chooses which direction to move and transitions according to the equation $s_{t+1} = s_t + a_t$. We design an expert that commits itself randomly to one of the two endpoints of the state space $s = -10$ or $s = 10$ by a bimodal stochastic policy. We generate 10000 trajectories from the expert and use them as training data for direct behavior cloning.

We train a NBP agent using GAN training [10]: given the expert trajectories, train a NBP as a generator that produces similar trajectories and train a separate classifier to distinguish true expert/generated trajectories. Unlike maximum likelihood, GAN training tends to capture modes of the expert trajectories. If we train a unimodal Gaussian policy using GAN training, the agent may commit to a single mode; below we show that trained NBP policy captures both modes.

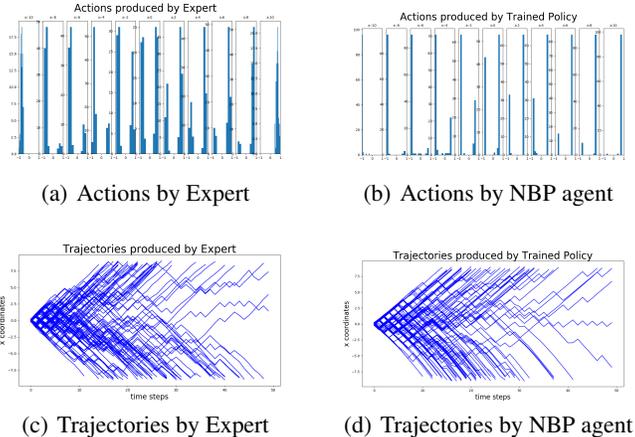


Figure 5: Imitating a bimodal expert: (a)(b) compare the actions produced by the expert and the trained NBP agent at different states s . The expert policy has a bimodal policy across different states and becomes increasingly unimodal when $s \approx \pm 10$; the trained policy captures such bimodal behavior. (c)(d) compare the trajectories of expert/trained agent. The vertical axis indicates the states s and horizontal axis is the time steps in an episode, each trajectory is terminated at $s = \pm 10$. Trajectories of both expert and trained policy are very similar.

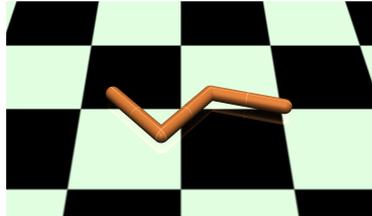
Stochastic Swimmer. The goal is to train a Swimmer robot that moves either forward or backward. It is not easy to specify a reward function that directly translates into such bimodal behavior and it is not easy to train a bimodal agent under such complex dynamics even if the reward is available. Instead, we train two Swimmers using RL objectives corresponding to two deterministic modes: swimming forward and swimming backward. Combining the trajectories generated by these two modes provides a policy that stochastically commits to either swimming forward or backward. We train a NFP agent (with maximum likelihood behavior cloning [8]) and NBP agent (with GAN training [10]) to imitate expert policy. The trajectories generated by trained policies show that the trained policies have fused these two modes of movement.

F Hyper-parameters and Ablation Study

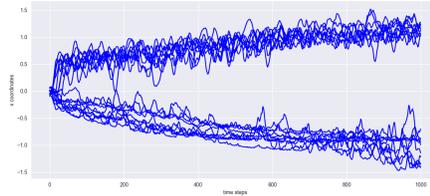
Hyper-parameters. Refer to Appendix C for a detailed description of architectures of NFP and NBP and hyper-parameters used in the experiments. For NFP, critical hyper-parameters are entropy coefficient β , number of transformation layers K and number of hidden units per layer k for transformation function s, t . For NBP, critical hyper-parameters are entropy coefficient β and the initialized variance parameter for factorized Gaussian ρ . In all conventional locomotion tasks, we set $\beta = 0.0$; for multi-modal policy tasks, we set $\beta \in \{0.1, 0.01, 0.001\}$. We use Adam [16] for optimization with learning rate $\alpha \in \{3 \cdot 10^{-5}, 3 \cdot 10^{-4}\}$.

Ablation Study. For NBP, the default baseline is $\beta = 0.0, \rho = -4.0$. We fix other hyper-parameters and change only one set of hyper-parameter to observe its effect on the performance. Intuitively, large ρ encourages a widespread distribution over parameter θ and consequently a more uniform initial distribution over actions. From Figure 7 we see that the performance is not monotonic in ρ, β . We find the model is relatively sensitive to hyper-parameter ρ and a general good choice is $\rho \in \{-4.0, -5.0, -6.0\}$.

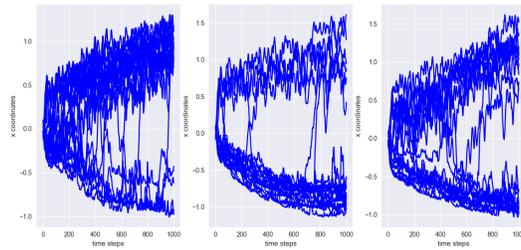
For NFP, the default baseline is $\beta = 0.0, K = 4, k = 3$. We fix other hyper-parameters and change only one set of hyper-parameter to observe its effect on the performance. In general, we find the model's performance is fairly robust to hyper-parameters (see Figure 8): large K, k will increase the complexity of the policy but does not necessarily benefit performance on benchmark tasks; strictly positive entropy coefficient $\beta > 0.0$ does not make much difference on benchmark tasks, though for learning multi-modal policies, adding positive entropy regularization is more likely to lead to multi-modality.



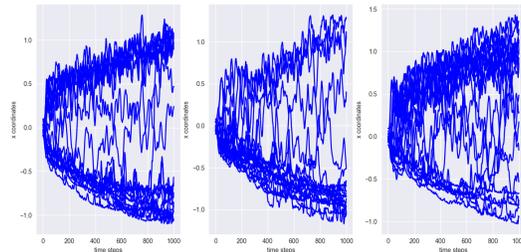
(a) Swimmer



(b) Trajectories by Expert

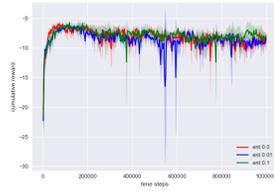


(c) Trajectories by NBP agent

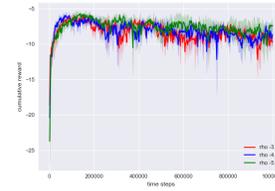


(d) Trajectories by NFP agent

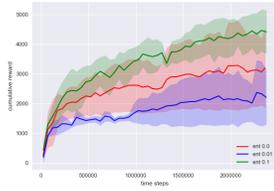
Figure 6: Combining multiple modes by Imitation Learning: stochastic Swimmer. (a) Illustration of Swimmer; (b) Expert trajectories produced by two Swimmers moving in two opposite directions (forward and backward). Vertical axis is the x coordinate of the Swimmer, horizontal axis is the time steps; (c) Trajectories produced by NBP agent trained using GAN under different seeds; (d) Trajectories produced by NFP agent trained using maximum likelihood under different seeds. Implicit policy agents have incorporated two modes of behavior into a single policy, yet a unimodal Gaussian policy can at most commit to one mode.



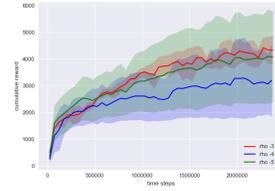
(a) Reacher: entropy



(b) Reacher: initial Gaussian variance parameter



(c) HalfCheetah: entropy

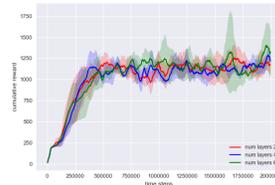


(d) HalfCheetah: initial Gaussian variance parameter

Figure 7: Ablation study: NBP



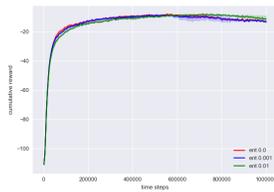
(a) Hopper: entropy



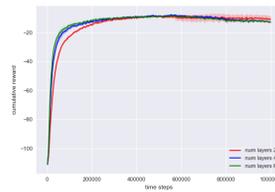
(b) Hopper: # of layers



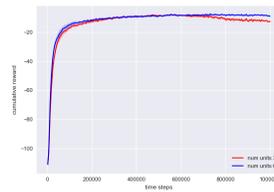
(c) Hopper: # of units



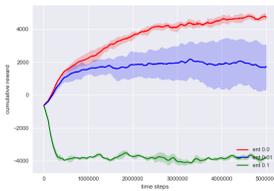
(d) Reacher: entropy



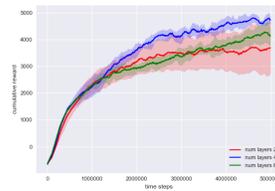
(e) Reacher: # of layers



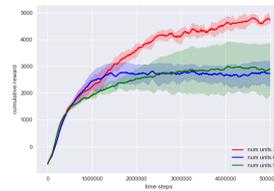
(f) Reacher: # of units



(g) HalfCheetah: entropy



(h) HalfCheetah: # of layers



(i) HalfCheetah: # of units

Figure 8: Ablation study: NFP