

Soft Policy Gradient Method for Maximum Entropy Deep Reinforcement Learning

Wenjie Shi* , Shiji Song and Cheng Wu

Department of Automation, Tsinghua University, Beijing, China
 shiwj16@mails.tsinghua.edu.cn, shijis@mail.tsinghua.edu.cn, wuc@tsinghua.edu.cn

Abstract

Maximum entropy deep reinforcement learning (RL) methods have been demonstrated on a range of challenging continuous tasks. However, existing methods either suffer from severe instability when training on large off-policy data or cannot scale to tasks with very high state and action dimensionality such as 3D humanoid locomotion. Besides, the optimality of desired Boltzmann policy set for non-optimal soft value function is not persuasive enough. In this paper, we first derive soft policy gradient based on entropy regularized expected reward objective for RL with continuous actions. Then, we present an off-policy actor-critic, model-free maximum entropy deep RL algorithm called deep soft policy gradient (DSPG) by combining soft policy gradient with soft Bellman equation. To ensure stable learning while eliminating the need of two separate critics for soft value functions, we leverage double sampling approach to making the soft Bellman equation tractable. The experimental results demonstrate that our method outperforms in performance over off-policy prior methods.

1 Introduction

Model-free reinforcement learning (RL) aims to acquire an effective behavior policy through ongoing trial and error interaction with a black box environment. The sole goal of standard RL is to optimize the quality of an agent’s behavior policy by maximizing cumulative discounted rewards. And so far, standard model-free RL has been applied to a range of challenging domains, such as games [Silver *et al.*, 2016], robotics [Levine *et al.*, 2016], finance [Prashanth *et al.*, 2016] and healthcare [Shortreed *et al.*, 2011]. However, some notorious drawbacks of standard model-free RL such as sample complexity, hyperparameter sensitivity or instability, have limited its widespread adoption in real-world domains.

Generally, policy optimization and value iteration are two basic paradigms in standard model-free RL. Policy based approaches, such as TRPO [Schulman *et al.*, 2015], directly

optimize the quantity of interest while remaining stable under function approximation, but policy based RL typically requires on-policy learning that is extravagantly expensive. By contrast, value based approaches, such as DQN [Mnih *et al.*, 2015], can learn from any ”off-policy” transitions sampled from identical environment, making them inherently more sample efficient [Gu *et al.*, 2017]. But unfortunately, off-policy learning does not stably interact with function approximation. Recently, some attempts have been made by combining policy and value based RL under actor-critic framework, such as DDPG [Lillicrap *et al.*, 2015], but there still remain some unsettled issues including how to set choose the type of policy and exploration noise.

Maximum entropy RL, which augments the standard maximum RL objective with an entropy regularization term [Toussaint, 2009], is another popular framework that can handle the tasks with multi-modality. As discussed in prior work, a stochastic policy may emerge as the optimal answer when maximum entropy RL tries to connect the optimal control with probabilistic inference [Todorov, 2008]. Intuitively, framing control as inference produces policies that try to learn all of the ways instead of the best way to perform the task. More importantly, maximum entropy framework shows several amazing advantages when compared to standard RL framework. First, entropy regularization encourages exploration and helps prevent early convergence to sub-optimal policies. Second, the resulting policies can serve as a good initialization for finetuning to a more specific behavior. Third, maximum entropy framework provides a better exploration mechanism for seeking out the best mode in a multimodal reward landscape. Fourth, the resulting policies are more robust in the face of adversarial perturbations as demonstrated in [Haarnoja *et al.*, 2017].

Actually, prior works have proposed various policy and value based RL methods under maximum entropy framework, such as combining policy gradient with Q-learning [O’Donoghue *et al.*, 2016], soft Q-learning [Haarnoja *et al.*, 2017; Schulman *et al.*, 2017]. However, policy based methods requiring on-policy learning suffer from poor sample complexity, while value based methods applying off-policy learning need a complex approximate sampling procedure in continuous action spaces.

In this paper, we explore to design an off-policy and stable model-free deep RL algorithm by combining policy and val-

*Contact Author

ue based methods under maximum entropy RL framework, which we call Deep Soft Policy Gradient (DSPG). First, we propose soft policy gradient under maximum entropy RL framework, and rigorous derivation of this proposition is given. Second, soft policy gradient is combined with soft Bellman equation by employing two deep function approximators to learn the soft Q-function and the stochastic policy, respectively. Finally, we present experimental results that show a significant improvement in performance over prior methods including DDPG and SAC.

2 Related Work

Our deep soft policy gradient (DSPG) algorithm mainly incorporates two ingredients: an off-policy actor-critic architecture with separate stochastic policy and action-value function networks, and an entropy regularization term to enable stability and exploration. We review prior works that draw on some of these ideas in this section.

Off-policy actor-critic has been verified as a feasible method to improve sample efficiency by exploiting off-policy data from other sources, such as past experience. And a particularly popular implementation of off-policy actor-critic architecture is DDPG [Lillicrap *et al.*, 2015], which employs a Q-function estimator to enable off-policy learning, and a deterministic actor to optimize this Q-function by applying deterministic policy gradient [Silver *et al.*, 2014]. However, some issues are still not completely settled. First, for many tasks with continuous action spaces, it is not possible to straightforwardly apply Q-learning which solves the greedy policy, DDPG instead uses deterministic policy to avoid global optimization with respect to the action at every timestep, but which is at the cost of performance loss. Second, additional noise process is required for exploration in the context of deterministic policy, the type and scale of noise must be chosen meticulously for different problem settings to achieve good performance. Furthermore, the interplay between the deterministic actor and the Q-function typically makes DDPG suffer from severe hyperparameter sensitivity [Henderson *et al.*, 2017]. As a consequence, it is difficult to generalize DDPG to very complex, high-dimensional tasks.

Maximum entropy reinforcement learning optimizes policies to maximize the entropy regularized expected reward objective. More recently, several progresses have been made in developing policy and value based methods under the framework of maximum entropy reinforcement learning. Roughly speaking, existing works fall into two categories: softmax temporal consistency [Nachum *et al.*, 2017a; Nachum *et al.*, 2017b] and soft policy iteration [O’Donoghue *et al.*, 2016; Haarnoja *et al.*, 2017; Haarnoja *et al.*, 2018]. Softmax temporal consistency between the optimal policy and soft optimal state value leads to path-wise consistency learning (PCL) methods [Nachum *et al.*, 2017a] and Trust-PCL [Nachum *et al.*, 2017b], which solve jointly the policy and soft state value by minimizing soft consistency error. However, PCL and Trust-PCL depend on complete trajectories and succumb to the instability when training on large off-policy data. Concurrent to PCL and Trust-PCL, some other methods such as PGQ [O’Donoghue *et al.*, 2016] and soft Q-learning [Haarnoja

et al., 2017], instead apply the idea underlying generalized policy iteration to learn maximum entropy policies by alternating policy evaluation and policy improvement. However, PGQ operate on simple tabular representations and are difficult to scale to continuous or high-dimensional domains, while soft Q-learning draws samples from an approximate sampling network. Building on soft Q-learning, soft actor-critic (SAC) [Haarnoja *et al.*, 2018] realizes policy improvement by minimizing Kullback-Leibler divergence between the current policy and the desired policy. However, how to choose the desired policy set for non-optimal value functions is somewhat subjective. Moreover, a separate function approximator for soft state-value is necessary to stabilize the learning, but which results in cumulative approximation errors. In contrast, we derive directly soft policy gradient based on the entropy regularized expected reward objective. And double sampling approach is utilized to stabilize the learning without separate soft value function approximator.

3 Preliminaries

In this section, we will define the reinforcement learning problem addressed in this paper and briefly summarize maximum entropy reinforcement learning framework.

3.1 Notation

We study reinforcement learning and control tasks with continuous action spaces. An agent aims to learn the optimal policies to maximize an entropy regularized expected reward objective by continuing trial-and-error in a stochastic environment. To this end, we define the problem as policy search in an infinite-horizon Markov decision process (MDP), which consists of the tuple $(\mathcal{S}, \mathcal{A}, p, r)$. The state space \mathcal{S} and action space \mathcal{A} are assumed to be continuous, and the state transition probability $p_{\mathbf{s}_t, \mathbf{s}_{t+1}}^{\mathbf{a}_t} : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \rightarrow [0, \infty)$ represents the probability density of next state $\mathbf{s}_{t+1} \in \mathcal{S}$ given current state $\mathbf{s}_t \in \mathcal{S}$ and action $\mathbf{a}_t \in \mathcal{A}$. The environment emits a reward $r : \mathcal{S} \times \mathcal{A} \rightarrow [r_{\min}, r_{\max}]$ on each transition, which we will abbreviate as $r_t \triangleq r(\mathbf{s}_t, \mathbf{a}_t)$ to simplify notation. Throughout this paper, we use $p_k^\pi(\mathbf{s} \rightarrow \mathbf{s}')$ to denote the probability of going from state \mathbf{s} to state \mathbf{s}' in k steps under the policy π . For $\mathbf{s}_0 = \mathbf{s}, \mathbf{s}_k = \mathbf{s}'$, we have

$$p_k^\pi(\mathbf{s} \rightarrow \mathbf{s}') = \sum_{\mathbf{a}_0} \pi(\mathbf{a}_0 | \mathbf{s}_0) \sum_{\mathbf{s}_1} p_{\mathbf{s}_0 \mathbf{s}_1}^{\mathbf{a}_0} \cdots \sum_{\mathbf{a}_{k-1}} \pi(\mathbf{a}_{k-1} | \mathbf{s}_{k-1}) p_{\mathbf{s}_{k-1} \mathbf{s}_k}^{\mathbf{a}_{k-1}}, \quad (1)$$

except that $p_0^\pi(\mathbf{s} \rightarrow \mathbf{s}) = 1$. Besides, we will also use $\rho_\pi(\mathbf{s}_t)$ and $\rho_\pi(\mathbf{s}_t, \mathbf{a}_t)$ to denote the state and state-action marginals of trajectory distribution induced by the policy $\pi(\mathbf{a}_t | \mathbf{s}_t)$.

3.2 Maximum Entropy Reinforcement Learning

Different from standard reinforcement learning, the sole goal of maximum entropy reinforcement learning is to maximize a more general entropy regularized objective, which is defined by augmenting the expected sum of rewards with discounted entropy terms, such that the optimal policy aims to maximize its entropy at each visited state:

$$J(\pi) = \sum_{t=0}^{T-1} \mathbb{E}_{(\mathbf{s}_t, \mathbf{a}_t) \sim \rho_\pi} [r(\mathbf{s}_t, \mathbf{a}_t) + \tau \mathcal{H}^\pi(\cdot | \mathbf{s}_t)], \quad (2)$$

where $\mathcal{H}^\pi(\cdot|\mathbf{s}_t) = -\sum_{\mathbf{a}_t} \pi(\mathbf{a}_t|\mathbf{s}_t) \log \pi(\mathbf{a}_t|\mathbf{s}_t)$ is the entropy of policy π at state \mathbf{s}_t . And $\tau \geq 0$ is a user-specified temperature parameter that controls the degree of entropy regularization. For the rest of this paper, we omit writing the temperature explicitly, as it can always be subsumed into the reward by scaling it by τ^{-1} . Furthermore, this objective can be easily extended to infinite horizon problems by introducing a discount factor γ to ensure that the expected sum of rewards and entropies is finite.

Under the framework of maximum entropy reinforcement learning, soft Q-function $Q^\pi(\mathbf{s}_t, \mathbf{a}_t)$ describes the expected sum of discounted future rewards and entropies except the entropy of state \mathbf{s}_t after taking an action \mathbf{a}_t in state \mathbf{s}_t and thereafter following the policy π :

$$Q^\pi(\mathbf{s}_t, \mathbf{a}_t) = r_t + \sum_{i>t} \gamma^{i-t} \mathbb{E}_{(\mathbf{s}_i, \mathbf{a}_i) \sim \rho_\pi} [r_i + \mathcal{H}^\pi(\cdot|\mathbf{s}_i)]. \quad (3)$$

Similar to standard reinforcement learning, we can derive directly soft Bellman equation from the above definition of soft Q-function as follows.

$$Q^\pi(\mathbf{s}_t, \mathbf{a}_t) = \mathbb{E}_{\mathbf{s}_{t+1} \sim p} [r_t + \gamma \mathbb{E}_{\mathbf{a}_{t+1} \sim \pi} [Q^\pi(\mathbf{s}_{t+1}, \mathbf{a}_{t+1}) + \mathcal{H}^\pi(\cdot|\mathbf{s}_{t+1})]]. \quad (4)$$

Prior works have made great progresses in learning maximum entropy policy by directly solving for optimal soft Q-function [Ziebart *et al.*, 2008; Fox *et al.*, 2015; Haarnoja *et al.*, 2017]. In subsequent sections, we will first propose and prove rigorously soft policy gradient under the framework of maximum entropy reinforcement learning, which corresponds to deterministic policy gradient in standard reinforcement learning. Then, we will further discuss how we can combine soft policy gradient with soft Bellman equation to develop an off-policy and stable deep reinforcement learning algorithm called deep soft policy gradient (DSPG), which corresponds to DDPG in standard reinforcement learning.

4 Soft Policy Gradient

Inspired by the observation that policy gradient methods [Degris *et al.*, 2012] are perhaps the most widely used for tasks with continuous action spaces in standard reinforcement learning, we aim to develop corresponding soft policy gradient methods for maximum entropy reinforcement learning. In this section, we derive directly soft policy gradient based on the entropy regularized expected reward objective and thereafter combine it with function approximation.

4.1 Soft Policy Gradient

The basic idea behind soft policy gradient is to represent the policy by a parametric probability distribution $\pi(\mathbf{a}|\mathbf{s}) = \mathbb{P}[\mathbf{a}|\mathbf{s}; \theta]$ that stochastically selects action \mathbf{a} in state \mathbf{s} according to parameter vector θ . And soft policy gradient typically proceed by sampling this stochastic policy and adjusting the parameter θ of this policy in the direction of greater entropy regularized expected reward objective.

Proposition 1 (Soft Policy Gradient). *Suppose that the MDP satisfies the policy $\pi(\mathbf{a}|\mathbf{s})$ is differentiable with respect to its*

parameter θ , i.e., that $\frac{\partial \pi(\mathbf{a}|\mathbf{s})}{\partial \theta}$ exists. Then,

$$\nabla_\theta J(\pi_\theta) = \mathbb{E}_{(\mathbf{s}, \mathbf{a}) \sim \rho_{\pi_\theta}} [(Q^\pi(\mathbf{s}, \mathbf{a}) - \log \pi(\mathbf{a}|\mathbf{s}) - 1) \nabla_\theta \log \pi(\mathbf{a}|\mathbf{s})]. \quad (5)$$

Proof. According to the definition of soft Q-function (3), we rewrite the entropy regularized expected reward objective (2) as follows.

$$J(\pi_\theta) = \mathbb{E}_{(\mathbf{s}, \mathbf{a}) \sim \rho_{\pi_\theta}} [Q^\pi(\mathbf{s}, \mathbf{a}) + \mathcal{H}^\pi(\cdot|\mathbf{s})]. \quad (6)$$

To get the gradient of (6), we first derive the gradient of the entropy regularized expected reward objective $J(\pi_\theta|\mathbf{s}_0)$ with a designated start state \mathbf{s}_0 .

For the convenience of displaying, we give the following simplified representation.

$$\mathcal{G}(\pi_\theta|\mathbf{s}) = \sum_{\mathbf{a}} \frac{\partial \pi(\mathbf{a}|\mathbf{s})}{\partial \theta} Q^\pi(\mathbf{s}, \mathbf{a}) + \frac{\partial \mathcal{H}^\pi(\cdot|\mathbf{s})}{\partial \theta}. \quad (7)$$

Then, we have

$$\begin{aligned} \frac{\partial J(\pi_\theta|\mathbf{s}_0)}{\partial \theta} &= \frac{\partial}{\partial \theta} \mathbb{E}_{\mathbf{a}_0 \sim \pi_\theta} [Q^\pi(\mathbf{s}_0, \mathbf{a}_0) + \mathcal{H}^\pi(\cdot|\mathbf{s}_0)] \\ &= \frac{\partial}{\partial \theta} \left[\sum_{\mathbf{a}_0} \pi(\mathbf{a}_0|\mathbf{s}_0) Q^\pi(\mathbf{s}_0, \mathbf{a}_0) + \mathcal{H}^\pi(\cdot|\mathbf{s}_0) \right] \\ &= \sum_{\mathbf{a}_0} \pi(\mathbf{a}_0|\mathbf{s}_0) \frac{\partial Q^\pi(\mathbf{s}_0, \mathbf{a}_0)}{\partial \theta} + \mathcal{G}(\pi_\theta|\mathbf{s}_0). \end{aligned} \quad (8)$$

Substituting the soft Bellman equation (4) into (8), we obtain

$$\begin{aligned} \frac{\partial J(\pi_\theta|\mathbf{s}_0)}{\partial \theta} &= \gamma p_1^\pi(\mathbf{s}_0 \rightarrow \mathbf{s}_1) \sum_{\mathbf{a}_1} \pi(\mathbf{a}_1|\mathbf{s}_1) \frac{\partial Q^\pi(\mathbf{s}_1, \mathbf{a}_1)}{\partial \theta} \\ &\quad + \sum_{k \in \{0,1\}} \sum_{\mathbf{s}_k} \gamma^k p_k^\pi(\mathbf{s}_0 \rightarrow \mathbf{s}_k) \mathcal{G}(\pi_\theta|\mathbf{s}_k). \end{aligned} \quad (9)$$

Repeating the above expansion step infinite times, we get

$$\begin{aligned} \frac{\partial J(\pi_\theta|\mathbf{s}_0)}{\partial \theta} &= \sum_{k=0}^{\infty} \sum_{\mathbf{s}_k} \gamma^k p_k^\pi(\mathbf{s}_0 \rightarrow \mathbf{s}_k) \mathcal{G}(\pi_\theta|\mathbf{s}_k) \\ &= \sum_{\mathbf{s}} \sum_{k=0}^{\infty} \gamma^k p_k^\pi(\mathbf{s}_0 \rightarrow \mathbf{s}) \mathcal{G}(\pi_\theta|\mathbf{s}). \end{aligned} \quad (10)$$

Then, we can get the gradient of (6) by calculating the expectation of start state \mathbf{s}_0 in (10) as follows:

$$\begin{aligned} \frac{\partial J(\pi_\theta)}{\partial \theta} &= \sum_{\mathbf{s}_0} \rho_\pi(\mathbf{s}_0) \sum_{\mathbf{s}} \sum_{k=0}^{\infty} \gamma^k p_k^\pi(\mathbf{s}_0 \rightarrow \mathbf{s}) \mathcal{G}(\pi_\theta|\mathbf{s}) \\ &= \sum_{\mathbf{s}} \sum_{\mathbf{s}_0} \rho_\pi(\mathbf{s}_0) \sum_{k=0}^{\infty} \gamma^k p_k^\pi(\mathbf{s}_0 \rightarrow \mathbf{s}) \mathcal{G}(\pi_\theta|\mathbf{s}) \\ &= \sum_{\mathbf{s}} \rho_\pi(\mathbf{s}) \left[\sum_{\mathbf{a}} \frac{\partial \pi(\mathbf{a}|\mathbf{s})}{\partial \theta} Q^\pi(\mathbf{s}, \mathbf{a}) + \frac{\partial \mathcal{H}^\pi(\cdot|\mathbf{s})}{\partial \theta} \right] \\ &= \sum_{\mathbf{s}} \rho_\pi(\mathbf{s}) \sum_{\mathbf{a}} (Q^\pi(\mathbf{s}, \mathbf{a}) - \log \pi(\mathbf{a}|\mathbf{s}) - 1) \frac{\partial \pi(\mathbf{a}|\mathbf{s})}{\partial \theta} \\ &= \mathbb{E}_{(\mathbf{s}, \mathbf{a}) \sim \rho_{\pi_\theta}} [(Q^\pi(\mathbf{s}, \mathbf{a}) - \log \pi(\mathbf{a}|\mathbf{s}) - 1) \nabla_\theta \log \pi(\mathbf{a}|\mathbf{s})]. \end{aligned}$$

□

The above expression of soft policy gradient is surprisingly simple. In particular, there are no terms of the form $\frac{\partial \rho_{\pi_{\theta}}(\mathbf{s})}{\partial \theta}$, which means the effect of policy changes on the distribution of states does not appear.

Moreover, this proposition has important practical value, because it is convenient for approximating the gradient by sampling. Just as policy gradient theorem [Sutton *et al.*, 2000] in standard reinforcement learning, a key issue that soft policy gradient must address is how to estimate the soft Q-function $Q^{\pi}(\mathbf{s}, \mathbf{a})$. Actually, a simple approach is to use a sampled return, which has led to a variant of the REINFORCE algorithm [Williams, 1992].

4.2 Soft Policy Gradient with Approximation

As discussed above, using the true return to approximate soft Q-function $Q^{\pi}(\mathbf{s}, \mathbf{a})$ is a feasible approach, but which is notoriously expensive in terms of sample complexity. Instead we consider the case in which $Q^{\pi}(\mathbf{s}, \mathbf{a})$ is approximated by a learned function approximator. In principle, if the approximation is sufficiently good, we might hope to use it in place of $Q^{\pi}(\mathbf{s}, \mathbf{a})$ in (5) and still point roughly in the direction of the gradient. And [Singh *et al.*, 1994] has proved that for the special case of function approximation arising in a tabular POMDP one could assure positive inner product with the gradient, which is sufficient to ensure improvement for moving in that direction.

However, substituting a function approximator $Q^{\omega}(\mathbf{s}, \mathbf{a})$ for the true soft Q-function $Q^{\pi}(\mathbf{s}, \mathbf{a})$ may introduce bias. Fortunately, there will be no bias if the following two conditions are satisfied according to [Sutton *et al.*, 2000].

- The function approximator $Q^{\omega}(\mathbf{s}, \mathbf{a})$ is compatible in the sense that

$$\frac{\partial Q^{\omega}(\mathbf{s}, \mathbf{a})}{\partial \omega} = \frac{\partial \pi(\mathbf{s}, \mathbf{a})}{\partial \theta} \frac{1}{\pi(\mathbf{s}, \mathbf{a})}. \quad (11)$$

- The parameters ω are chosen to minimise the mean-squared error $\varepsilon^2(\omega)$.

$$\varepsilon^2(\omega) = \mathbb{E}_{(\mathbf{s}, \mathbf{a}) \sim \rho_{\pi_{\theta}}} \left[(Q^{\omega}(\mathbf{s}, \mathbf{a}) - Q^{\pi}(\mathbf{s}, \mathbf{a}))^2 \right]. \quad (12)$$

Then, we have the following soft policy gradient with function approximation.

$$\nabla_{\theta} J(\pi_{\theta}) = \mathbb{E}_{(\mathbf{s}, \mathbf{a}) \sim \rho_{\pi_{\theta}}} \left[(Q^{\omega}(\mathbf{s}, \mathbf{a}) - \log \pi(\mathbf{a}|\mathbf{s}) - 1) \nabla_{\theta} \log \pi(\mathbf{a}|\mathbf{s}) \right]. \quad (13)$$

In other words, compatible function approximators are linear in the same features as the stochastic policy, and the parameters ω are the solution to the linear regression problem that estimates $Q^{\pi}(\mathbf{s}, \mathbf{a})$ from these features. In practice, condition (12) is usually relaxed in favour of policy evaluation algorithms that estimate the value function more efficiently by temporal-difference learning [Bhatnagar *et al.*, 2008]. Based on (13), a form of policy iteration with function approximation can be proved to converge to a locally maximum entropy policy according to prior works [Bertsekas and Tsitsiklis, 1996; Sutton *et al.*, 2000].

5 Combined with Soft Bellman Equation

In this section, we will present our proposed deep soft policy gradient (DSPG) algorithm, which is motivated by combining soft policy gradient with soft Bellman equation. First, We will further discuss how to achieve soft policy evaluation with soft Bellman equation (4) when a function approximator is used for soft Q-function. Then, the details of our DSPG algorithm will be given.

5.1 Soft Bellman Equation with Approximation

To improve the sample efficiency, we employ a function approximator $Q^{\omega}(\mathbf{s}, \mathbf{a})$ to estimate soft Q-function $Q^{\pi}(\mathbf{s}, \mathbf{a})$ in order to make soft policy gradient tractable while avoiding on-policy interaction with environment. Similar to many approaches in standard reinforcement learning, we make use of soft Bellman equation (4) to update the parameters ω of function approximator $Q^{\omega}(\mathbf{s}, \mathbf{a})$.

According to soft Bellman equation (4), we first define soft Bellman backup operator $\mathcal{T}Q^{\omega}$ as follows.

$$\mathcal{T}Q^{\omega} = r_t + \gamma \mathbb{E}_{\mathbf{a}_{t+1} \sim \pi} [Q^{\omega}(\mathbf{s}_{t+1}, \mathbf{a}_{t+1}) + \mathcal{H}^{\pi}(\cdot|\mathbf{s}_{t+1})]. \quad (14)$$

But in practice, we typically use the following sampled soft Bellman backup operator.

$$\tilde{\mathcal{T}}Q^{\omega} = r_t + \frac{\gamma}{M} \sum_j [Q^{\omega}(\mathbf{s}_{t+1}, \mathbf{a}_{(t+1)j}) - \log \pi_{\theta'}(\mathbf{a}_{(t+1)j}|\mathbf{s}_{t+1})]. \quad (15)$$

where the actions $\mathbf{a}_{(t+1)j}$ are sampled from the target policy $\pi_{\theta'}(\cdot|\mathbf{s}_{t+1})$ of current policy. Different from DDPG, which is a commonly used off-policy actor-critic algorithm in standard reinforcement learning and applies the deterministic policy to avoid the inner expectation in Bellman equation, we approximate the expectation of action in (14) by repeatedly sampling with the current policy. Indeed, the deterministic policy may induce catastrophic and unstable update of the policy and value function.

Note that the outer expectation in soft Bellman equation (4) depends only on the environment or state transition probability, which makes it possible to learn $Q^{\omega}(\mathbf{s}, \mathbf{a})$ by using previous off-policy transitions sampled from a replay buffer. Therefore, the parameters ω of $Q^{\omega}(\mathbf{s}, \mathbf{a})$ can be optimized by minimizing the soft loss:

$$\mathcal{L}(\omega) = \frac{1}{N} \sum_i \left[(Q^{\omega}(\mathbf{s}_i, \mathbf{a}_i) - \tilde{\mathcal{T}}Q^{\omega'})^2 \right]. \quad (16)$$

where $Q^{\omega'}$ is the target network used to stabilize the learning of soft Q-function Q^{ω} .

5.2 Deep Soft Policy Gradient Algorithm

It is not possible to straightforwardly apply soft Bellman equation (4) to continuous action spaces, because the inner expectation with respect to the action is intractable when there is not an explicitly feasible way to sample from the current policy. Instead, we used an off-policy actor-critic approach based on soft policy gradient derived in last section. Specifically, as discussed above, soft policy gradient method maintains a parameterized actor function $\pi_{\theta}(\mathbf{a}|\mathbf{s})$, which specifies

the current policy by stochastically mapping states to all possible actions, while the critic $Q^\omega(\mathbf{s}, \mathbf{a})$ is learned to evaluate the current policy by optimizing the soft loss (16) based on soft Bellman equation.

Gradient Clipping. One unique challenge in maximum entropy reinforcement learning when applying soft policy gradient is that there exist the terms $Q^\omega(\mathbf{s}, \mathbf{a})$ explicitly in (13), which means that soft Q-function will directly influence the scale of gradient and thus lead to unstable update of policy and even bad convergence. Motivated by the idea to deal with gradient explosion problem, we leverage gradient clipping approach to address this issue. Gradient clipping approach aims to limit soft policy gradient to an appropriate range by a global norm \mathcal{N} . Intuitively, clipped gradients prevent drastic updates of the current policy and thus stabilize the learning. And an additional benefit from gradient clipping is that constrained soft policy gradient eliminates the impact of different baselines of soft Q-functions in various reinforcement learning tasks. By contrast, it is not necessary for DDPG to use gradient clipping approach, because the deterministic policy gradient depends on $\frac{\partial Q^\omega(\mathbf{s}, \mathbf{a})}{\partial \mathbf{a}}$ rather than $Q^\omega(\mathbf{s}, \mathbf{a})$.

Double Sampling. Another key challenge is that soft policy gradient (13) depends on not only state distribution but also the current policy, which makes it impossible to compute soft policy gradient with transitions sampled from replay buffer. Instead we utilize double sampling approach to estimate the expectation with respect to action. The main idea underlying double sampling approach is to separate the expectation with respect to the state and action by two sampling steps. The first sampling step uses the transitions sampled from replay buffer to approximate the expectation with respect to state distribution induced by the current policy, which is commonly used in policy gradient implementations, such as DPG and DDPG. In the second sampling step, we estimate the expectation with respect to action by sampling with the current policy. Although existing off-policy actor-critic algorithm [Degris *et al.*, 2012] uses a different behaviour policy to generate trajectories, double sampling approach has distinct advantage in terms of performance and sampling complexity.

Based on above gradient clipping and double sampling approaches, soft policy gradient can be rewritten in a tractable form as follows.

$$\nabla_\theta J(\pi_\theta) \approx \frac{1}{NM} \sum_i \sum_j [(Q^\omega(\mathbf{s}_i, \mathbf{a}_{ij}) - \log \pi(\mathbf{a}_{ij}|\mathbf{s}_i) - 1) \nabla_\theta \log \pi(\mathbf{a}_{ij}|\mathbf{s}_i)] \quad (17)$$

$$\nabla_\theta^{CLIP} J(\pi_\theta) = \text{Clip}_{by_norm}(\nabla_\theta J(\pi_\theta), \mathcal{N}) \quad (18)$$

where \mathbf{a}_{ij} is sampled with the current policy $\pi(\cdot|\mathbf{s}_i)$. Clip_{by_norm} denotes the gradient clipping operator.

To summarize, we propose deep soft policy gradient (DSPG, Algorithm 1) algorithm for learning maximum entropy policies in continuous domains. The algorithm proceeds by alternating between evaluating the current policy with soft Bellman equation and improving the current policy with soft policy gradient.

Algorithm 1: DSPG Algorithm

- 1 Randomly initialize the critic $Q^\omega(\mathbf{s}, \mathbf{a})$ and actor $\mu_\theta(\cdot|\mathbf{s})$ with weights ω and θ ;
 - 2 Initialize the target network $Q^{\omega'}$ and $\mu_{\theta'}$ with $\omega' \leftarrow \omega$ and $\theta' \leftarrow \theta$;
 - 3 Initialize replay buffer \mathcal{R} ;
 - 4 **for** each episode **do**
 - 5 Reset initial observation state \mathbf{s}_0 ;
 - 6 **for** each step **do**
 - 7 Sample an action \mathbf{a}_t for \mathbf{s}_t using $\pi_\theta(\cdot|\mathbf{s}_t)$;
 - 8 Execute action \mathbf{a}_t and receive r_t, \mathbf{s}_{t+1} ;
 - 9 Store transition $(\mathbf{s}_t, \mathbf{a}_t, r_t, \mathbf{s}_{t+1})$ in \mathcal{R} ;
 - 10 Sample a random minibatch of N transitions $(\mathbf{s}_i, \mathbf{a}_i, r_i, \mathbf{s}_{i+1})$ from \mathcal{R} ;
 - 11 Sample M actions $\mathbf{a}_{(i+1)j}$ for each state \mathbf{s}_{i+1} using $\pi_{\theta'}(\cdot|\mathbf{s}_{i+1})$;
 - 12 Compute the soft loss $\mathcal{L}(\omega)$ (16);
 - 13 Update the critic by minimizing the soft loss;
 - 14 Sample M actions \mathbf{a}_{ij} for each state \mathbf{s}_i using $\pi_\theta(\cdot|\mathbf{s}_i)$;
 - 15 Compute the clipped soft policy gradient (18);
 - 16 Update the actor by applying the clipped soft policy gradient $\nabla_\theta^{CLIP} J(\pi_\theta)$;
 - 17 Update the target networks:
 $\omega' \leftarrow \alpha\omega + (1 - \alpha)\omega'$
 $\theta' \leftarrow \alpha\theta + (1 - \alpha)\theta'$
 - 18 **end**
 - 19 **end**
-

6 Experiments

We evaluate our proposed algorithm, namely DSPG, across several benchmark continuous control tasks and compare them to standard DDPG and SAC implementations. We find that DSPG can consistently match or beat the performance of these baselines.

6.1 Setup

We chose four well-known benchmark continuous control tasks (Ant, Hopper, HalfCheetah and Walker2d) available from OpenAI Gym and utilizing MuJoCo environment. And two feed-forward neural networks were used to represent the policy and value estimates. In addition, we used uniformly a simple policy represented by a multivariate gaussian for both DSPG and SAC.

For DDPG, identical experimental setup was employed in our DDPG as in [Lillicrap *et al.*, 2015], actually which has to perform a grid search to achieve good performance due to its hyperparameter sensitivity. Besides, to be consistent with DDPG and DSPG, we merely compared to the general form of SAC without applying additional tricks, such as hard target update every fixed interval and using two Q-functions, although these tricks are effective for all three algorithms.

To evaluate DSPG's performance with an appropriate global gradient norm \mathcal{N} , we performed a simple grid search over $\mathcal{N} \in \{1, 3, 5\}$. And the search results show that smaller global gradient norm will slightly slow down the learning but

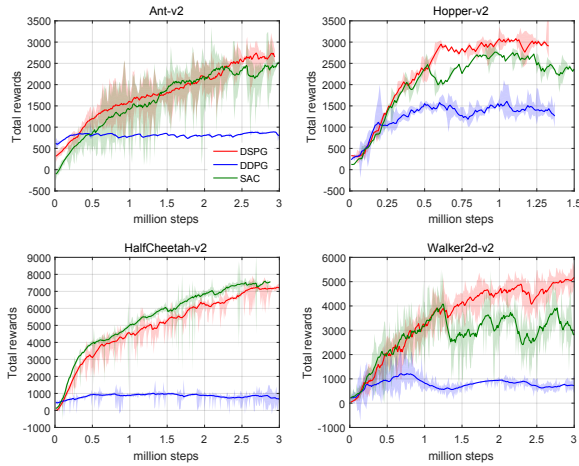


Figure 1: The results of DSPG against DDPG and SAC baselines on continuous control benchmarks. Each plot shows the average total rewards of evaluation rollouts across 7 randomly seeded training runs after choosing best hyperparameters. The x-axis shows millions of environment steps.

make the learning more stable. Based on this conclusion, we chose $\mathcal{N} = 5$ for both Ant and Walker2d, while $\mathcal{N} = 1$ and $\mathcal{N} = 3$ were chosen for Hopper and HalfCheetah, respectively. Moreover, we sampled $M = 64$ times with the current learned policy π_θ to estimate the expectation with respect to the action in each train step, and we performed 4 train steps every interaction step to accelerate the learning for all algorithms. A detailed description of our implementation and experimental setup is available in the Appendix A. The source code of our DSPG implementation will be available online after the paper is accepted.

6.2 Results

We present the average total rewards over training of DSPG, DDPG and SAC in Figure 1. And Table 1 shows the best average total rewards in all steps of training for our implementation and baselines. The results show that, overall, DSPG substantially outperforms DDPG on all of benchmark tasks, in terms of both sample efficiency and convergent performance. The notable gap of performance between DSPG and DDPG is ascribed to the hyperparameter sensitivity of DDPG, but also suggests significant advantage of our DSPG relative to DDPG in challenging tasks with very high state and action dimensionality. Moreover, it can be seen in Figure 1 that DSPG beats the performance of SAC in Hopper and Walker2d, and can match the performance of SAC in Ant and HalfCheetah.

In addition to the average total rewards, the stability of algorithm also plays a crucial role in performance. And a conclusion can be drawn from Figure 1 that DSPG outperforms consistently SAC in terms of the stability in all four benchmark tasks. The good stability of DSPG can be demonstrated from two observations. First, clipped soft policy gradient is applied to guarantee steepest optimization direction and prevent large update step of policy, which is typically considered to be main source of instability. Second, only one critic

Domains	DSPG	SAC	DDPG
Ant-v2	3384.074	3472.346	1012.242
HalfCheetah-v2	7889.747	8103.202	1014.372
Hopper-v2	3674.029	3652.893	3097.082
Walker2d-v2	6060.884	5520.419	2507.199

Table 1: The results for best average total rewards in all training steps for our DSPG, DDPG and SAC implementations. These results are each on different setups with different hyperparameter searches. Thus, although it is not possible to make any definitive claims based on this data, we do conclude that our results are overall competitive with DDPG and SAC baselines.

network is used to estimate soft Q-function with double sampling approach in DSPG, which avoids cumulative approximation errors resulting from a separate soft value function approximator. In practical, we suggest to take into account available computing resources and the cost of sampling when choosing appropriate algorithm for specific task. For tasks with limited computing resources but low cost of sampling, DSPG seems to be better.

7 Conclusion

We presented deep soft policy gradient (DSPG), an off-policy actor-critic and model-free maximum entropy deep reinforcement learning algorithm. Our theoretical results derive soft policy gradient based on entropy regularized expected reward objective. Building on this result, we formulated our DSPG algorithm by combining soft policy gradient with soft Bellman equation. The experimental results suggest that DSPG can perform well on a set of challenging continuous control tasks, improving upon DDPG and SAC in terms of average total rewards and stability. Moreover, soft policy gradient provides a promising avenue for future works.

Acknowledgments

This research was supported by the National Science Foundation of China under Grant 41427806.

A Experimental Setup

Throughout all experiments, we use Adam for learning the neural network parameters with a learning rate 5×10^{-5} and 5×10^{-4} for the actor and critic respectively. For critic we use a discount factor of $\gamma = 0.99$. For the soft target updates we use $\alpha = 0.01$. Both the actor and critic are represented by full-connected feed-forward neural network with two hidden layers of dimensions 512. And all hidden layers use ReLU activation. Specially, we use identity and sigmoid activations for the mean and standard deviation in the output layer respectively. The algorithm use a replay buffer size of three million and train with minibatch sizes of 100 for each train step. Training does not start until the replay buffer has enough samples for a minibatch and does not stop until the global time step equals to the threshold of 3×10^6 . In addition, we scales the reward function by a factor of 5 for all four tasks, as is common in prior works.

References

- [Bertsekas and Tsitsiklis, 1996] D Bertsekas and J Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, 1996.
- [Bhatnagar *et al.*, 2008] Shalabh Bhatnagar, Mohammad Ghavamzadeh, Mark Lee, and Richard S Sutton. Incremental natural actor-critic algorithms. In *Advances in neural information processing systems*, pages 105–112, 2008.
- [Degris *et al.*, 2012] Thomas Degris, Patrick M Pilarski, and Richard S Sutton. Model-free reinforcement learning with continuous action in practice. In *American Control Conference (ACC), 2012*, pages 2177–2182. IEEE, 2012.
- [Fox *et al.*, 2015] Roy Fox, Ari Pakman, and Naftali Tishby. Taming the noise in reinforcement learning via soft updates. *arXiv preprint arXiv:1512.08562*, 2015.
- [Gu *et al.*, 2017] Shixiang Gu, Timothy Lillicrap, Zoubin Ghahramani, Richard E Turner, and Sergey Levine. Q-prop: Sample-efficient policy gradient with an off-policy critic. In *International Conference on Learning Representations*, 2017.
- [Haarnoja *et al.*, 2017] Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. Reinforcement learning with deep energy-based policies. *arXiv preprint arXiv:1702.08165*, 2017.
- [Haarnoja *et al.*, 2018] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018.
- [Henderson *et al.*, 2017] Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. Deep reinforcement learning that matters. *arXiv preprint arXiv:1709.06560*, 2017.
- [Levine *et al.*, 2016] Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *Journal of Machine Learning Research*, 17(1):1334–1373, 2016.
- [Lillicrap *et al.*, 2015] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- [Mnih *et al.*, 2015] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- [Nachum *et al.*, 2017a] Ofir Nachum, Mohammad Norouzi, Kelvin Xu, and Dale Schuurmans. Bridging the gap between value and policy based reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 2775–2785, 2017.
- [Nachum *et al.*, 2017b] Ofir Nachum, Mohammad Norouzi, Kelvin Xu, and Dale Schuurmans. Trust-pcl: An off-policy trust region method for continuous control. *arXiv preprint arXiv:1707.01891*, 2017.
- [O’Donoghue *et al.*, 2016] Brendan O’Donoghue, Remi Munos, Koray Kavukcuoglu, and Volodymyr Mnih. Combining policy gradient and q-learning. *arXiv preprint arXiv:1611.01626*, 2016.
- [Prashanth *et al.*, 2016] LA Prashanth, Cheng Jie, Michael Fu, Steve Marcus, and Csaba Szepesvári. Cumulative prospect theory meets reinforcement learning: Prediction and control. In *International Conference on Machine Learning*, pages 1406–1415, 2016.
- [Schulman *et al.*, 2015] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International Conference on Machine Learning*, pages 1889–1897, 2015.
- [Schulman *et al.*, 2017] John Schulman, Xi Chen, and Pieter Abbeel. Equivalence between policy gradients and soft q-learning. *arXiv preprint arXiv:1704.06440*, 2017.
- [Shortreed *et al.*, 2011] Susan M Shortreed, Eric Laber, Daniel J Lizotte, T Scott Stroup, Joelle Pineau, and Susan A Murphy. Informing sequential clinical decision-making through reinforcement learning: an empirical study. *Machine Learning*, 84(1-2):109–136, 2011.
- [Silver *et al.*, 2014] David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *International Conference on Machine Learning*, 2014.
- [Silver *et al.*, 2016] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.
- [Singh *et al.*, 1994] Satinder P Singh, Tommi Jaakkola, and Michael I Jordan. Learning without state-estimation in partially observable markovian decision processes. In *Machine Learning Proceedings 1994*, pages 284–292. Elsevier, 1994.
- [Sutton *et al.*, 2000] Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pages 1057–1063, 2000.
- [Todorov, 2008] Emanuel Todorov. General duality between optimal control and estimation. In *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*, pages 4286–4292. IEEE, 2008.
- [Toussaint, 2009] Marc Toussaint. Robot trajectory optimization using approximate inference. In *Proceedings of the 26th annual international conference on machine learning*, pages 1049–1056. ACM, 2009.
- [Williams, 1992] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- [Ziebart *et al.*, 2008] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. Maximum entropy inverse reinforcement learning. In *AAAI*, volume 8, pages 1433–1438. Chicago, IL, USA, 2008.