
Diversity Actor-Critic: Sample-Aware Entropy Regularization for Sample-Efficient Exploration

Seungyul Han¹ Youngchul Sung¹

Abstract

In this paper, sample-aware policy entropy regularization is proposed to enhance the conventional policy entropy regularization for better exploration. Exploiting the sample distribution obtainable from the replay buffer, the proposed sample-aware entropy regularization maximizes the entropy of the weighted sum of the policy action distribution and the sample action distribution from the replay buffer for sample-efficient exploration. A practical algorithm named diversity actor-critic (DAC) is developed by applying policy iteration to the objective function with the proposed sample-aware entropy regularization. Numerical results show that DAC significantly outperforms existing recent algorithms for reinforcement learning.

1. Introduction

Reinforcement learning (RL) aims to maximize the expected return under Markov decision process (MDP) (Sutton & Barto, 1998). When the given task is complex, e.g., the environment has high action-dimensions or sparse rewards, it is important to explore state-action pairs well for high performance (Agre & Rosenschein, 1996). For better exploration, recent RL considers various methods: maximizing the policy entropy to take actions more uniformly (Ziebart et al., 2008; Fox et al., 2015; Haarnoja et al., 2017), maximizing diversity gain that yields intrinsic reward to explore rare states by counting the number of visiting states (Strehl & Littman, 2008; Lopes et al., 2012), maximizing information gain (Houthoofd et al., 2016; Hong et al., 2018), maximizing model prediction error (Achiam & Sastry, 2017; Pathak et al., 2017). In particular, based on policy iteration for soft Q-learning, Haarnoja et al. (2018a) extended maximum entropy RL and proposed an off-policy actor-critic algorithm, soft actor-critic (SAC), which has competitive performance

for challenging continuous control tasks.

In this paper, we consider the problem of policy entropy regularization in off-policy learning and propose a generalized approach to policy entropy regularization for sample-efficient exploration. In off-policy learning, we store samples in the replay buffer and reuse old samples to update the current policy (Mnih et al., 2015). Thus, the sample buffer has information about the old samples. However, the simple policy entropy regularization tries to maximize the entropy of the current policy irrespective of the distribution of the previous samples in the replay buffer. In order to exploit the sample information in the replay buffer and enhance performance, we propose sample-aware entropy regularization, which tries to maximize the entropy of the weighted sum of the current policy action distribution and the sample action distribution from the replay buffer. We develop a practical and efficient algorithm for return maximization based on the proposed sample-aware entropy regularization without explicitly computing the replay-buffer sample distribution, and demonstrate that the proposed algorithm yields significant enhancement in exploration and final performance on various difficult environments such as tasks with sparse reward or high action dimensions.

2. Related Works

Entropy regularization: Entropy regularized RL maximizes the sum of the expected return and the policy action entropy. It encourages the agent to visit the action space uniformly for each given state, and can provide more accurate model prediction (Ziebart, 2010). Entropy regularization is considered in various domains: inverse reinforcement learning (Ziebart et al., 2008), stochastic optimal control problems (Todorov, 2008; Toussaint, 2009; Rawlik et al., 2013), and off-policy reinforcement learning (Fox et al., 2015; Haarnoja et al., 2017). Nachum et al. (2017a) showed that there exists a connection between value-based and policy-based RL under entropy regularization. O’Donoghue et al. (2016) proposed an algorithm combining value-based and policy-based RL, and Schulman et al. (2017a) proved that they are equivalent. Hazan et al. (2019) maximized the entropy of state distribution induced by the current policy by using state mixture distribution for better pure exploration.

¹Department of Electrical Engineering, Korea Advanced Institute of Science and Technology, Daejeon, South Korea. Correspondence to: Youngchul Sung <yicsung@kaist.ac.kr>.

Diversity gain: Diversity gain is used to provide a guidance for exploration to the agent. To achieve diversity gain, many intrinsically-motivated approaches and intrinsic reward design methods have been considered, e.g., intrinsic reward based on curiosity (Chentanez et al., 2005; Baldassarre & Mirolli, 2013), model prediction error (Achiam & Sastry, 2017; Pathak et al., 2017; Burda et al., 2018), divergence/information gain (Houthoofd et al., 2016; Hong et al., 2018), counting (Strehl & Littman, 2008; Lopes et al., 2012; Tang et al., 2017; Martin et al., 2017), and unification of them (Bellemare et al., 2016). Eysenbach et al. (2018) explicitly maximized diversity based on mutual information.

Off-policy learning: Off-policy learning reuses samples generated from behaviour policies for policy update (Sutton & Barto, 1998; Degris et al., 2012), so it is sample-efficient compared to on-policy learning. In order to reuse old samples, a replay buffer that stores trajectories generated by previous policies is used for Q-learning (Mnih et al., 2015; Lillicrap et al., 2015; Fujimoto et al., 2018; Haarnoja et al., 2018a). To further enhance both stability and sample efficiency, several methods were considered, e.g., combining on-policy and off-policy (Wang et al., 2016; Gu et al., 2016; 2017), and generalization from on-policy to off-policy (Nachum et al., 2017b; Han & Sung, 2019).

3. Background

Setup: We assume a basic RL setup composed of an environment and an agent. The environment follows an infinite horizon Markov decision process $(\mathcal{S}, \mathcal{A}, P, \gamma, r)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, P is the transition probability, γ is the discount factor, and $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function. In this paper, we consider continuous state and action spaces. The agent has policy distribution $\pi \in \Pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, \infty)$ which selects an action a_t for given state s_t at time step t , where Π is the policy space. Then, the agent receives reward $r_t := r(s_t, a_t)$ from the environment and the state changes to s_{t+1} . Standard RL aims to maximize the discounted return $\mathbb{E}_{s_0 \sim p_0, \tau_0 \sim \pi} [\sum_{t=0}^{\infty} \gamma^t r_t]$, where $\tau_t = (s_t, a_t, s_{t+1}, a_{t+1}, \dots)$ is an episode trajectory.

Soft Actor-Critic: Soft actor-critic (SAC) includes a policy entropy regularization term in the objective function with the aim of performing more diverse actions for each given state and visiting states with higher entropy for better exploration (Haarnoja et al., 2018a). The entropy-augmented policy objective function of SAC is given by

$$J_{SAC}(\pi) = \mathbb{E}_{\tau_0 \sim \pi} \left[\sum_{t=0}^{\infty} \gamma^t (r_t + \beta \mathcal{H}(\pi(\cdot|s_t))) \right], \quad (1)$$

where \mathcal{H} is the entropy function and $\beta \in (0, \infty)$ is the entropy coefficient. SAC is a practical off-policy actor-critic algorithm based on soft policy iteration (SPI) that alternates

soft policy evaluation to estimate the true soft Q -function and soft policy improvement to find the optimal policy that maximizes (1). SPI theoretically guarantees convergence to the optimal policy that maximizes (1) for finite MDPs.

4. The Diversity Actor-Critic Algorithm

4.1. Motivation of Sample-Aware Entropy

In order to guarantee the convergence of Q-learning, there is a key assumption: *Each state-action pair must be visited infinitely often* (Watkins & Dayan, 1992). Without proper exploration, policy can converge to local optima and task performance can be degraded severely (Plappert et al., 2017). Therefore, exploration for visiting diverse state-action pairs is important for RL. There has been extensive research for better exploration in RL. One important line of recent methods is to use intrinsic reward based on prediction model (Chentanez et al., 2005; Baldassarre & Mirolli, 2013; Achiam & Sastry, 2017; Pathak et al., 2017; Burda et al., 2018). In this approach, we have a prediction model for a target value or distribution, and the prediction model is learned with samples. Then, the prediction error is used as the intrinsic reward added to the actual reward from the environment, and the discounted sum of the actual and intrinsic rewards is maximized. The fundamental rationale behind this approach is that the prediction model is well learned for the frequently-observed state-action pairs in the sample history and hence the prediction error is small. On the other hand, for unobserved or less-observed state-action pairs in the sample history, the prediction model training is not enough and the prediction error is large. In this way, un- or less-explored state-action pairs are favored.

Another successful method for exploration is policy entropy regularization with the representative method shown in (1). In (1), one can view the policy action entropy $\mathcal{H}(\pi(\cdot|s_t))$ as an intrinsic reward added to the actual reward r_t . This method relies on the fact that the entropy attains maximum when the distribution is uniform (Cover & Thomas, 2006). Thus, maximizing the discounted sum of the actual reward r_t and the policy action entropy as in (1) yields a policy that tries not only to maximize the actual reward but also to visit states with high action entropy and to take more uniform actions for better exploration. Furthermore, in this method the weighting factor β in (1) can be learned adaptively based on the Lagrangian method to maintain a certain level of entropy (Haarnoja et al., 2018b). However, on the contrary to the prediction model-based method, the entropy regularization method does not exploit the previously-observed sample information to construct the intrinsic reward at current time t since the intrinsic reward $\mathcal{H}(\pi(\cdot|s_t))$ depends only on the policy $\pi(\cdot|s_t)$, and $\pi(\cdot|s_t)$ for given s_t does not directly capture the sample distribution information from the replay buffer.

In this paper, we consider the maximum entropy framework in off-policy learning and extend this framework by devising an efficient way to exploit the sample information in the replay buffer so as to harness the merits of the two aforementioned approaches: taking more uniform actions and promoting un- or less-performed actions in the past.

4.2. Proposed Policy Objective Function

In order to use the previous sample information in entropy-based exploration, we first define the mixture distribution

$$q_{mix}^{\pi, \alpha}(\cdot|s) := \alpha\pi(\cdot|s) + (1 - \alpha)q(\cdot|s), \quad (2)$$

where $\alpha \in [0, 1]$ is the weighting factor, $\pi(\cdot|s)$ is the policy (action) distribution, and $q(\cdot|s)$ is the sample action distribution of the replay buffer \mathcal{D} which stores previous samples. Then, we propose maximizing the following objective function

$$J(\pi) = \mathbb{E}_{\tau_0 \sim \pi} \left[\sum_{t=0}^{\infty} \gamma^t (r_t + \beta \mathcal{H}(q_{mix}^{\pi, \alpha}(\cdot|s_t))) \right], \quad (3)$$

where we refer to $\mathcal{H}(q_{mix}^{\pi, \alpha}(\cdot|s_t))$ as the *sample-aware entropy*. Note that maximizing the sample-aware entropy enhances sample-efficient exploration because in this case the learning guides the policy to choose actions so that the mixture distribution $q_{mix}^{\pi, \alpha}(\cdot|s_t)$ becomes uniform. That is, $\pi(\cdot|s_t)$ will choose actions rare in the replay buffer (i.e., the density $q(\cdot|s_t)$ is low) with high probability and choose actions stored many times in the replay buffer (i.e., the density $q(\cdot|s_t)$ is high) with low probability so as to make the mixture distribution uniform. Indeed, we can decompose the sample-aware entropy $\mathcal{H}(q_{mix}^{\pi, \alpha})$ for given s_t as

$$\mathcal{H}(q_{mix}^{\pi, \alpha}) = - \int_{a \in \mathcal{A}} (\alpha\pi + (1 - \alpha)q) \log \underbrace{(\alpha\pi + (1 - \alpha)q)}_{=q_{mix}^{\pi, \alpha}} \quad (4)$$

$$= \int \alpha\pi \log \frac{\alpha\pi}{\alpha\pi + (1 - \alpha)q} + \int (1 - \alpha)q \log \frac{(1 - \alpha)q}{\alpha\pi + (1 - \alpha)q} - \int \alpha\pi \log(\alpha\pi) - \int (1 - \alpha)q \log((1 - \alpha)q) \quad (5)$$

$$= D_{JS}^{\alpha}(\pi||q) + \alpha\mathcal{H}(\pi) + (1 - \alpha)\mathcal{H}(q) + \text{constant}, \quad (6)$$

where $D_{JS}^{\alpha}(\pi||q) := \alpha \int \pi \log \frac{\pi}{\alpha\pi + (1 - \alpha)q} + (1 - \alpha) \int q \log \frac{q}{\alpha\pi + (1 - \alpha)q}$ is the α -skew Jensen-Shannon (JS)-symmetrization of KL divergence (Nielsen, 2019). D_{JS}^{α} reduces to the standard JS divergence for $\alpha = \frac{1}{2}$ and to zero for $\alpha = 0$ or 1. When $\alpha = 1$, $\mathcal{H}(q_{mix}^{\pi, \alpha})$ reduces to the simple entropy and the problem reduces to (1). When $\alpha \in (0, 1)$, on the other hand, all the first three terms in the right-hand side (RHS) of (6) remain. Thus, the added regularized term in (3) will guide the policy to have more uniform actions due to $\mathcal{H}(\pi)$ and simultaneously to promote actions away from q due to $D_{JS}^{\alpha}(\pi||q)$. Thus, the proposed policy objective function (3) has the desired

properties. Note that the sample-aware entropy is included as reward not as an external regularization term added to the discounted return, and this targets optimization for high total sample-aware entropy of the entire trajectory. (An analytic toy example showing the efficiency of the sample-aware entropy regularization is provided in Appendix A.) The main challenge to realize policy design with (3) is how to compute the sample distribution q , which is necessary to compute the objective function. Explicit computation of q requires a method such as discretization and counting for continuous state and action spaces. This should be done for each environment and can be a difficult and tedious job for high dimensional environments. Even if such empirical q is obtained by discretization and counting, generalization of q to arbitrary state-action pairs is typically required to actually implement an algorithm based on function approximation and this makes the problem difficult further. In the remainder of this paper, circumventing this difficulty, we develop a practical and efficient algorithm to realize (3) without explicitly computing q .

4.3. Algorithm Construction

Our algorithm construction for the objective function (3) is based on *diverse policy iteration*, which is a modification of the soft policy iteration of Haarnoja et al. (2018a). Diverse policy iteration is composed of diverse policy evaluation and diverse policy improvement. Note that the sample action distribution q is updated as iteration goes on. However, it changes very slowly since the buffer size is much larger than the time steps of one iteration. Hence, for the purpose of algorithm derivation, we regard the action distribution q as a fixed distribution in this section.

As in typical policy iteration, for diverse policy iteration, we first define the true diverse Q -function Q^{π} as

$$Q^{\pi}(s_t, a_t) := \frac{1}{\beta} r_t + \mathbb{E}_{\tau_{t+1} \sim \pi} \left[\sum_{l=t+1}^{\infty} \gamma^{l-t-1} \left(\frac{1}{\beta} r_l + \mathcal{H}(q_{mix}^{\pi, \alpha}(\cdot|s_l)) \right) \right], \quad (8)$$

by including the term $\mathcal{H}(q_{mix}^{\pi, \alpha}(\cdot|s_l))$. Since $Q^{\pi}(s_t, a_t)$ includes $\mathcal{H}(q_{mix}^{\pi, \alpha}(\cdot|s_l))$, it seemingly requires computation of q , as seen in (4)-(6). In order to circumvent this difficulty, we define the following ratio function:

$$R^{\pi, \alpha}(s_t, a_t) = \frac{\alpha\pi(a_t|s_t)}{\alpha\pi(a_t|s_t) + (1 - \alpha)q(a_t|s_t)}, \quad (9)$$

and express the objective and all required loss functions in terms of the ratio function not q . For this, based on (4), we rewrite $\mathcal{H}(q_{mix}^{\pi, \alpha}(\cdot|s_t))$ as follows:

$$\mathcal{H}(q_{mix}^{\pi, \alpha}) = \alpha \mathbb{E}_{a_t \sim \pi(\cdot|s_t)} [\log R^{\pi, \alpha}(s_t, a_t) - \log \alpha\pi(a_t|s_t)] + (1 - \alpha) \mathbb{E}_{a_t \sim q(\cdot|s_t)} [\log R^{\pi, \alpha}(s_t, a_t) - \log \alpha\pi(a_t|s_t)]. \quad (10)$$

$$J_{\pi_{old}}(\pi(\cdot|s_t)) := \beta \{ \mathbb{E}_{a_t \sim \pi} [Q^{\pi_{old}}(s_t, a_t) + \alpha(\log R^{\pi, \alpha}(s_t, a_t) - \log \alpha \pi(a_t|s_t))] + (1 - \alpha) \mathbb{E}_{a_t \sim q} [\log R^{\pi, \alpha}(s_t, a_t) - \log \alpha \pi(a_t|s_t)] \}. \quad (13)$$

$$\tilde{J}_{\pi_{old}}(\pi(\cdot|s_t)) := \beta \mathbb{E}_{a_t \sim \pi} [Q^{\pi_{old}}(s_t, a_t) + \alpha(\log R^{\pi_{old}, \alpha}(s_t, a_t) - \log \pi(a_t|s_t))], \quad (14)$$

The above equation is obtained by exploiting the entropy definition (4). As seen in (4), $\mathcal{H}(q_{mix}^{\pi, \alpha})$ is the sum of $-\alpha \mathbb{E}_{a_t \sim \pi(\cdot|s_t)} \log q_{mix}^{\pi, \alpha}$ and $-(1 - \alpha) \mathbb{E}_{a_t \sim q(\cdot|s_t)} \log q_{mix}^{\pi, \alpha}$, but $\log q_{mix}^{\pi, \alpha}$ in both terms can be expressed as $\log \frac{\alpha \pi(a_t|s_t)}{R^{\pi, \alpha}(s_t, a_t)}$ by the definition of the ratio function (9). So, we obtain (10). Now, the $\mathcal{H}(q_{mix}^{\pi, \alpha})$ expression in the RHS of (10) contains only the ratio function $R^{\pi, \alpha}$ and the policy π , and thus fits our purpose. Note that the expectation $\mathbb{E}_{a_t \sim q(\cdot|s_t)}$ will eventually be replaced by empirical expectation based on the samples in the replay buffer in a practical algorithm. So, it does not cause a problem. Thus, the added term $\mathcal{H}(q_{mix}^{\pi, \alpha})$ in (3) and (8) is fully expressed in terms of the ratio function $R^{\pi, \alpha}$ and the policy π . Now, we present the diverse policy iteration composed of diverse policy evaluation and diverse policy improvement.

Diverse policy evaluation: We first define a diverse action value function estimate $Q : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. Then, we define a modified Bellman backup operator acting on Q to estimate Q^π as

$$\mathcal{T}^\pi Q(s_t, a_t) := \frac{1}{\beta} r_t + \gamma \mathbb{E}_{s_{t+1} \sim P} [V(s_{t+1})], \quad (11)$$

where $V(s_t)$ is the estimated diverse state value function given by the sum of $\mathbb{E}_{a_t \sim \pi} [Q(s_t, a_t)]$ and $\mathcal{H}(q_{mix}^{\pi, \alpha})$ for given s_t , i.e.,

$$V(s_t) = \mathbb{E}_{a_t \sim \pi} [Q(s_t, a_t) + \alpha \log R^{\pi, \alpha}(s_t, a_t) - \alpha \log \alpha \pi(a_t|s_t)] + (1 - \alpha) \mathbb{E}_{a_t \sim q} [\log R^{\pi, \alpha}(s_t, a_t) - \log \alpha \pi(a_t|s_t)], \quad (12)$$

where we used the expression (10) for $\mathcal{H}(q_{mix}^{\pi, \alpha})$. Note that for the diverse policy evaluation, the policy π under evaluation is given. Hence, the ratio function $R^{\pi, \alpha}(s_t, a_t)$ is given for given π by its definition (9). Hence, $V(s_t)$ in (12) is well defined and thus the mapping $\mathcal{T}^\pi Q(s_t, a_t)$ on the current estimate $Q(s_t, a_t)$ in (11) is well defined. By repeating the mapping \mathcal{T}^π on Q , the resulting sequence converges to Q^π . Proof is given in Lemma 1 in Appendix B.

Diverse policy improvement: Now, consider diverse policy improvement. Suppose that we are given $Q^{\pi_{old}}(\cdot, \cdot)$ for the current policy π_{old} . (In this diverse policy improvement step, we use the notation π_{old} for the given current policy to distinguish from the notation π as the optimization argument.) Then, we construct the diverse policy objective function $J_{\pi_{old}}(\pi(\cdot|s_t))$ as shown in (13), where the notation π in (13) represents the optimization argument. $J_{\pi_{old}}(\pi)$ is the policy objective function estimated under $Q^{\pi_{old}}$. If we replace π_{old} in $J_{\pi_{old}}(\pi)$ with π and view state s_t as the initial state, then (13) reduces to $J(\pi)$ in (3). (This can be

checked with (3), (8), (10) and (13).) Note that π in the $R^{\pi, \alpha}$ and $\log(\alpha \pi)$ terms inside the expectations in (13) is the optimization argument π . We update the policy from π_{old} to π_{new} as

$$\pi_{new} = \arg \max_{\pi \in \Pi} J_{\pi_{old}}(\pi). \quad (15)$$

Then, π_{new} satisfies $Q^{\pi_{new}}(s_t, a_t) \geq Q^{\pi_{old}}(s_t, a_t)$, $\forall (s_t, a_t) \in \mathcal{S} \times \mathcal{A}$. Proof is given in Lemma 2 in Appendix B.

Then, in a similar way to the proof of the convergence of the soft policy iteration (Haarnoja et al., 2018a), we can prove the convergence of the diverse policy iteration, stated in the following theorem.

Theorem 1 (Diverse Policy Iteration) *By repeating iteration of the diverse policy evaluation applying the Bellman operator (11) and the diverse policy improvement (15), any initial policy converges to the optimal policy π^* s.t. $Q^{\pi^*}(s_t, a_t) \geq Q^{\pi'}(s_t, a_t)$, $\forall \pi' \in \Pi$, $\forall (s_t, a_t) \in \mathcal{S} \times \mathcal{A}$. Furthermore, such π^* achieves maximum J in (3).*

Proof. See Appendix B.1.

For proof of Theorem 1, we need the assumption of finite MDPs as in the proof of usual policy iteration or SPI. Later, we consider function approximation for the policy and the value functions to implement the diverse policy iteration in continuous state and action spaces, based on the convergence proof in finite MDPs.

Although Theorem 1 proves convergence of the diverse policy iteration for finite MDPs and provides a basis for implementation with function approximation for continuous MDPs, actually finding the optimal policy by using Theorem 1 is difficult due to the step (15) used in Theorem 1. The reason is as follows. In order to facilitate proof of monotone improvement by the step (15), we set π in the $R^{\pi, \alpha}$ term in (13) as the optimization argument, as seen in Appendix B.1. Otherwise, proof of monotone improvement is not tractable. However, this setup causes a problem in practical implementation. For practical implementation with function approximation, we will eventually use parameterized estimates for the required functions, as we do shortly. For the policy π we will use π_θ with parameter θ . Under this situation, let us consider the ratio function again. The ratio function $R^{\pi, \alpha}$ for a given π is a mapping from $(\mathcal{S}, \mathcal{A})$ to $[0, 1)$, as seen in (9). In the case that π in $R^{\pi, \alpha}$ is the optimization argument policy with parameter θ , we need to

define a mapping from $(\mathcal{S}, \mathcal{A})$ to $[0, 1]$ for each of all possible values of θ . That is, the output value of $R^{\pi_{\theta}, \alpha}(s_t, a_t)$ depends not only on (s_t, a_t) but also on θ . To capture this situation, we need to set the input to the ratio function $R^{\pi_{\theta}, \alpha}$ as (s_t, a_t, θ) . However, the dimension of the policy (neural network) parameter θ is huge and thus implementation of $R^{\pi_{\theta}, \alpha}(s_t, a_t)$ as a function of (s_t, a_t, θ) is not simple. To circumvent this difficulty, we need to modify the policy objective function so that it involves a much simpler form for the ratio function for easy implementation. For this, instead of using $R^{\pi, \alpha}$ with π being the optimization argument, we use the ratio function $R^{\pi_{old}, \alpha}$ for the given current policy π_{old} so that π in the $R^{\pi, \alpha}$ term in the policy objective function is not the optimization argument anymore but fixed as the given current policy π_{old} . With this replacement, we manage to show the following result:

Theorem 2 Consider the new objective function for policy improvement $\tilde{J}_{\pi_{old}}(\pi(\cdot|s_t))$ in (14), where the ratio function inside the expectation in (14) is the ratio function for the given current policy π_{old} . Suppose that the policy is parameterized with parameter θ . Then, for parameterized policy π_{θ} , the two objective functions $J_{\pi_{\theta_{old}}}(\pi_{\theta}(\cdot|s_t))$ in (13) and $\tilde{J}_{\pi_{\theta_{old}}}(\pi_{\theta}(\cdot|s_t))$ in (14) have the same gradient direction for θ at $\theta = \theta_{old}$ for all $s_t \in \mathcal{S}$, where θ_{old} is the parameter of the given current policy π_{old} .

Proof. See Appendix B.2.

Note that maximizing the new policy objective function (14) is equivalent to minimizing $D_{KL}(\pi(\cdot|s_t) || \exp(Q^{\pi_{old}}(s_t, \cdot)/\alpha + R^{\pi_{old}, \alpha}(s_t, \cdot))) = -\tilde{J}_{\pi_{old}}(\pi(\cdot|s_t))/\beta$, and the improved policy obtained by maximizing (14) can be expressed as $\pi_{new}(\cdot|s_t) \propto \exp(Q^{\pi_{old}}(s_t, \cdot)/\alpha + R^{\pi_{old}, \alpha}(s_t, \cdot))$. Note also that the new policy objective function $\tilde{J}_{\pi_{old}}(\pi(\cdot|s_t))$ is a simplified version in two aspects. First, we require the ratio function only for the given current policy. Second, the $\mathbb{E}_{a_t \sim q}$ term in (13) disappeared. Note that π in the $\log \pi$ inside the expectation of $\tilde{J}_{\pi_{old}}(\pi(\cdot|s_t))$ in (14) is still the optimization argument. However, this is not a problem since we have the parameter θ for the policy in implementation and this parameter will be updated. Now, the ratio function in the policy objective function $\tilde{J}_{\pi_{old}}(\pi(\cdot|s_t))$ in (14) in the diverse policy improvement step is the ratio function for the given current policy. Furthermore, the ratio function in (12) in the diverse policy evaluation step is also for the given current policy. Hence, we need to implement and track only the ratio function for the current policy. Now, by Theorems 1 and 2, we can find the optimal policy maximizing (3) by iterating the diverse policy evaluation (11) and the diverse policy improvement maximizing $\tilde{J}_{\pi_{old}}(\pi(\cdot|s_t))$ in (14).

The final step to complete the proposed diverse policy iteration is learning of the ratio function for the current policy.

For this, we define an estimate function $R^{\alpha} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ for the true ratio function $R^{\pi, \alpha}$ of the current policy π and adopt the learning method proposed in the works of Sugiyama et al. (2012); Goodfellow et al. (2014). That is, we first define the objective function for R^{α} as

$$J_{ratio}(R^{\alpha}(s_t, \cdot)) = \alpha \mathbb{E}_{a_t \sim \pi(\cdot|s_t)} [\log R^{\alpha}(s_t, a_t)] + (1 - \alpha) \mathbb{E}_{a_t \sim q(\cdot|s_t)} [\log(1 - R^{\alpha}(s_t, a_t))]. \quad (16)$$

Then, we learn R^{α} by maximizing the objective $J_{ratio}(R^{\alpha})$. Note that for given s , $J_{ratio}(R^{\alpha}(s, \cdot)) = \int_a [c_1 \log R^{\alpha}(s, a) + c_2 \log(1 - R^{\alpha}(s, a))] da$, where $c_1 = \alpha\pi$, and $c_2 = (1 - \alpha)q$. The integral is maximized when the integrand for each a is maximized. The integrand $f(R^{\alpha}(s, a)) = c_1 \log R^{\alpha}(s, a) + c_2 \log(1 - R^{\alpha}(s, a))$ is a concave function of $R^{\alpha}(s, a)$, and its maximum is achieved when $R^{\alpha}(s, a) = c_1/(c_1 + c_2) = \alpha\pi/(\alpha\pi + (1 - \alpha)q)$. Hence, $J_{ratio}(R^{\alpha})$ is maximized when $R^{\alpha}(s_t, a_t) = \alpha\pi/(\alpha\pi + (1 - \alpha)q)$, which is exactly the desired ratio function shown in (9). Therefore, the ratio function for the current policy can be estimated by maximizing the objective function $J_{ratio}(R^{\alpha})$.

4.4. Diversity Actor Critic Implementation

We use deep neural networks to implement the policy π , the ratio function R^{α} , and the diverse value functions Q and V , and their network parameters are θ , η , ϕ , and ψ , respectively. Based on $\tilde{J}_{\pi_{old}}(\pi)$ in (14) and $J_{ratio}(R^{\alpha})$ in (16), we provide the practical objective functions: $\hat{J}_{\pi}(\theta)$ for the parameterized policy π_{θ} , and $\hat{J}_{R^{\alpha}}(\eta)$ for the parameterized ratio function estimator R_{η}^{α} , given by

$$\hat{J}_{\pi}(\theta) = \mathbb{E}_{s_t \sim \mathcal{D}, a_t \sim \pi_{\theta}} [Q_{\phi}(s_t, a_t) + \alpha \log R_{\eta}^{\alpha}(s_t, a_t) - \alpha \log \pi_{\theta}(a_t | s_t)], \quad (17)$$

$$\hat{J}_{R^{\alpha}}(\eta) = \mathbb{E}_{s_t \sim \mathcal{D}} [\alpha \mathbb{E}_{a_t \sim \pi_{\theta}} [\log R_{\eta}^{\alpha}(s_t, a_t)] + (1 - \alpha) \mathbb{E}_{a_t \sim \mathcal{D}} [\log(1 - R_{\eta}^{\alpha}(s_t, a_t))]], \quad (18)$$

where \mathcal{D} denotes the replay buffer. Based on the Bellman operator \mathcal{T}^{π} in (11), we provide the loss functions: $\hat{L}_Q(\phi)$ and $\hat{L}_V(\psi)$ for the parameterized value functions Q_{ϕ} and V_{ψ} , respectively, given by

$$\hat{L}_Q(\phi) = \mathbb{E}_{(s_t, a_t) \sim \mathcal{D}} \left[\frac{1}{2} (Q_{\phi}(s_t, a_t) - \hat{Q}(s_t, a_t))^2 \right], \quad (19)$$

$$\hat{L}_V(\psi) = \mathbb{E}_{s_t \sim \mathcal{D}} \left[\frac{1}{2} (V_{\psi}(s_t) - \hat{V}(s_t))^2 \right], \quad (20)$$

Algorithm 1 Diversity Actor Critic

```

Initialize parameter  $\theta, \eta, \psi, \bar{\psi}, \xi, \phi_i, i = 1, 2$ 
for each iteration do
    Sample a trajectory  $\tau$  of length  $N$  by using  $\pi_\theta$ 
    Store the trajectory  $\tau$  in the buffer  $\mathcal{D}$ 
    for each gradient step do
        Sample random minibatch of size  $M$  from  $\mathcal{D}$ 
        Compute  $\hat{J}_\pi(\theta), \hat{J}_{R^\alpha}(\eta), \hat{L}_Q(\phi_i), \hat{L}_V(\psi)$  from the minibatch
         $\theta \leftarrow \theta + \delta \nabla_\theta \hat{J}_\pi(\theta)$ 
         $\eta \leftarrow \eta + \delta \nabla_\eta \hat{J}_{R^\alpha}(\eta)$ 
         $\phi_i \leftarrow \phi_i - \delta \nabla_{\phi_i} \hat{L}_Q(\phi_i), i = 1, 2$ 
         $\psi \leftarrow \psi - \delta \nabla_\psi \hat{L}_V(\psi)$ 
        Update  $\bar{\psi}$  by EMA from  $\psi$ 
        if  $\alpha$ -adaptation is applied then
            Compute  $\hat{L}_\alpha(\xi)$  from the minibatch
             $\xi \leftarrow \xi - \delta \nabla_\xi \hat{L}_\alpha(\xi)$ 
        end if
    end for
end for
    
```

where the target values \hat{Q} and \hat{V} are defined as

$$\hat{Q}(s_t, a_t) = \frac{1}{\beta} r_t + \gamma \mathbb{E}_{s_{t+1} \sim P} [V_{\bar{\psi}}(s_{t+1})] \quad (21)$$

$$\begin{aligned} \hat{V}(s_t) &= \mathbb{E}_{a_t \sim \pi_\theta} [Q_\phi(s_t, a_t) \\ &\quad + \alpha \log R_\eta^\alpha(s_t, a_t) - \alpha \log \alpha \pi_\theta(a_t | s_t)] \\ &\quad + (1 - \alpha) \mathbb{E}_{a_t \sim \mathcal{D}} [\log R_\eta^\alpha(s_t, a_t) - \log \alpha \pi_\theta(a_t | s_t)]. \end{aligned} \quad (22)$$

Here, $\bar{\psi}$ is the network parameter of the target value $V_{\bar{\psi}}$ updated by exponential moving average (EMA) of ψ for stable learning (Mnih et al., 2015). In addition, we use two Q -functions $Q_{\phi_i}, i = 1, 2$ to reduce overestimation bias as proposed in (Fujimoto et al., 2018) and applied in SAC, and each Q -function is updated to minimize their loss function $\hat{L}_Q(\phi_i)$. For the policy and the value function update, the minimum of two Q -functions is used for the policy update. Combining all up to now, we propose the diversity actor-critic (DAC) algorithm summarized as Algorithm 1. Here, note that DAC becomes SAC when $\alpha = 1$, and becomes standard off-policy RL without entropy regularization when $\alpha = 0$. When $0 < \alpha < 1$, we accomplish sample-aware entropy regularization. Detailed implementation of DAC is provided in Appendix C. For DAC, we can consider the technique of SAC proposed in (Haarnoja et al., 2018b) using Q -function only for reducing complexity or automatic tuning of β for balancing the entropy and the return. However, in the case of DAC, both α and β affect the entropy term, so both should be tuned simultaneously.

5. α -Adaptation

In the proposed sample-aware entropy regularization, the weighting factor α between the policy and the sample distribution plays an important role in controlling the ratio between the policy distribution π and the sample action distribution q . However, it is difficult to find optimal α for each environment. To circumvent this α search, we propose an automatic adaptation method for α based on max-min principle widely considered in game theory, robust learning, and decision making problems (Chinchuluun et al., 2008). That is, since we do not know optimal α , an alternative formulation is that we maximize the return while maximizing the worst-case sample-aware entropy, i.e., $\min_\alpha \mathcal{H}(q_{mix}^{\pi, \alpha})$. Then, the max-min approach can be formulated as follows:

$$\max_\pi \mathbb{E}_{\tau_0 \sim \pi} \left[\sum_t \gamma^t (r_t + \beta \min_\alpha [\mathcal{H}(q_{mix}^{\pi, \alpha}) - \alpha c]) \right] \quad (23)$$


where c is a control hyperparameter for α adaptation. Note that we learn α to minimize the sample-aware entropy so that the entropy is maintained above a certain level to explore the state and action spaces well. So, the α learning objective is given by a Lagrangian form. Thus, when the α -learning model is parameterized with parameter ξ , the α learning objective is given by $\hat{L}_\alpha(\xi) = \mathbb{E}_{s_t \sim \mathcal{D}} [\mathcal{H}(q_{mix}^{\pi_\theta, \alpha \xi}) - \alpha \xi c]$. Detailed implementation of α -adaptation is given in Appendix C.1.

6. Experiments


In this section, we evaluate the proposed DAC algorithm on various continuous-action control tasks and provide ablation study. We first consider the pure exploration performance and then the performance on challenging sparse-reward or delayed Mujoco tasks. The source code of DAC based on Python Tensorflow is available at <http://github.com/seungyulhan/dac/>.

6.1. Pure Exploration Performance

For comparison baselines, we first considered the state-of-the-art entropy regularization methods: SAC and SAC-Div. SAC-Div is SAC combined with the exploration method in (Hong et al., 2018) that diversifies the policy from the buffer distribution by simply maximizing $J_{SAC}(\pi) + \alpha_d D(\pi || q)$ for some divergence D , where $J_{SAC}(\pi)$ is given in (1). For SAC-Div, we considered the KL divergence and the adaptive scale α_d with $\delta_d = 0.2$, as suggested in (Hong et al., 2018). The case of JS divergence used for SAC-Div is provided in Appendix F. Note that both SAC-Div and DAC contain a divergence term in their objective functions (DAC contains D_{JS}^α , as seen in (6) and (3)), but there is an important difference. SAC-Div adds a single divergence term on the reward sum $J_{SAC}(\pi)$. So, SAC-Div keeps π



(a) Continuous 4-room maze (b) Number of state visitation



(c) State visit histogram at 5k (left) 50k (middle) 300k (right) steps

Figure 1: Pure exploration task: Continuous 4-room maze

away from q , but does not guide learning of the policy to visit states on which the divergence between π and q is large. On the contrary, DAC contains the divergence term D_{JS}^{α} as an intrinsic reward at each time step inside the reward sum of $J(\pi)$, as seen in (3) and (6). Hence, DAC not only keeps π away from q but also learns a policy to visit states on which the divergence between π and q is large, to have large $J(\pi)$, so that more actions different from q are possible. This situation is analogous to the situation of SAC in which the entropy is included as an intrinsic reward inside the sum of $J_{SAC}(\pi)$, as seen in (1), and hence for large $J_{SAC}(\pi)$, SAC learns a policy to visit states on which the policy action entropy is large. In addition to SAC and SAC-Div, we considered the recent high-performance state-based exploration methods: random network distillation

(RND) (Burda et al., 2018) and MaxEnt(State) (Hazan et al., 2019). RND explores rare states by adding an intrinsic reward based on model prediction error, and MaxEnt(State) explores rare states by using a reward functional based on the entropy of state mixture distribution. Detailed simulation setup is provided in Appendix D.

In order to see the pure exploration performance of DAC ($\alpha = 0.5$ is used), we considered state visitation on a 100×100 continuous 4-room maze. The maze environment was designed by modifying a continuous grid map available at <https://github.com/huyaoyu/GridMap>, and it is shown in Fig. 1(a). State is the (x, y) position of the agent in the maze, action is (dx, dy) bounded by $[-1, 1] \times [-1, 1]$, and the agent location after action becomes $(x + dx, y + dy)$. The agent starts from the left lower corner $(0.5, 0.5)$ and explores the maze without any reward. Fig. 1(b) shows the number of total different state visitations averaged over 30 seeds, where the number of state visitations is obtained based on quantized 1×1 squares. Here, the shaded region in the figure represents one standard deviation (1σ) from the mean. As seen in Fig. 1(b), DAC visited much more states than the other methods, which shows the superior exploration performance of DAC. Fig. 1(c) shows the corresponding state visit histogram of all seeds with 1×1 square quantization. Here, as the color of a state becomes brighter, the state is visited more times. It is seen that SAC/SAC-Div rarely visit the right upper room even at 300k time steps, RND and MaxEnt(State) visit the right upper room more than SAC/SAC-Div, and DAC visits the right upper room far earlier and more than the other methods.

6.2. Performance on Sparse-Reward Mujoco Tasks

Then, we evaluated the performance on sparse-reward Mujoco tasks, which have been widely used as difficult environments for RL in many previous studies (Hong et al., 2018; Mazoure et al., 2019; Burda et al., 2018). We considered two versions. One was SparseMujoco, which is a sparse version of Mujoco (Todorov et al., 2012) in OpenAI Gym (Brockman et al., 2016), and the reward is 1 if the agent satisfies a certain condition, otherwise 0 (Hong et al., 2018; Mazoure et al., 2019). The other was DelayedMujoco (Zheng et al., 2018; Guo et al., 2018), which has the same state-action spaces as original Mujoco tasks but reward is sparsified. That is, rewards for D time steps are accumulated and the accumulated reward sum is delivered to the agent once every D time steps, so the agent receives no reward during the accumulation time.

First, we fixed $\alpha = 0.5$ for DAC and tested DAC on SparseMujoco. The result is shown in Fig. 2, which shows the performance averaged over 10 random seeds. For all performance plots, we used deterministic evaluation which

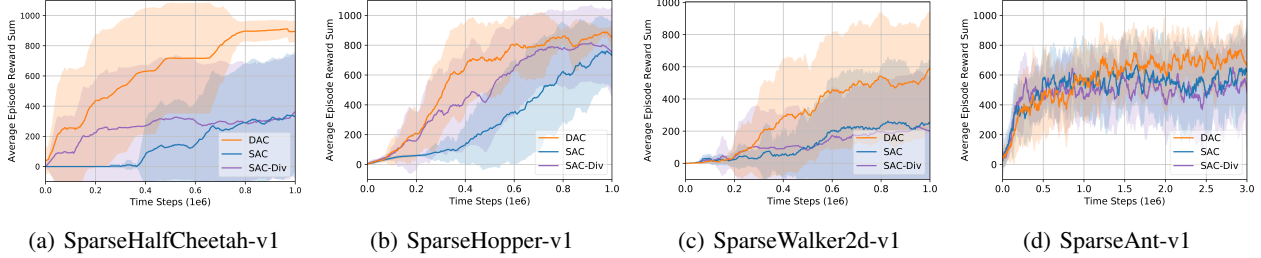



Figure 2: Performance comparison on Sparse Mujoco tasks

generated an episode by deterministic policy for each iteration, and the shaded region in the figure represents one standard deviation (1σ) from the mean. It is seen that DAC has significant performance gain over the competitive SAC and SAC-Div baselines. Figs. 3 (a) and (b) show the divergence $D_{JS}^\alpha(\pi||q)$ curve and the corresponding number of discretized state visitation curve, respectively, on SparseHalfCheetah shown in Fig. 2(a). (The curves for the other tasks are provided in Appendix E.1. See Appendix E.1 about the discretization.) It is seen in Fig. 3(a) that the divergence of DAC is much higher than those of SAC/SAC-Div throughout the learning time. This implies that the policy of DAC choose more diverse actions from the policy distribution far away from the sample action distribution q , so DAC visits more diverse states than the baselines, as seen in Fig. 3(b).

Next, we tested DAC on DelayedMujoco and HumanoidStandup. Note that HumanoidStandup is one of the difficult high-action dimensional Mujoco tasks, so its reward is not sparsified for test. We considered three cases for α of DAC: $\alpha = 0.5, 0.8$, and α -adaptation. Fig. 4 shows the result. Again, we can observe significant performance improvement by DAC over the SAC baselines. We can also observe that the best α depends on tasks. For example, $\alpha = 0.8$ is the best for DelayedHalfCheetah, but $\alpha = 0.5$ is the best for DelayedAnt. Thus, the result shows that α -adaptation method is necessary in order to adapt α properly for each


 Figure 3: (a) α -skewed JS symmetrization of KLD $D_{JS}^\alpha(\pi||q)$ with $\alpha = 0.5$ and (b) the corresponding mean number of state visitation

task. Although the proposed α -adaptation in Section 5 is sub-optimal, DAC with our α -adaptation method has top-level performance across all the considered tasks and further enhances the performance in some cases such as DelayedHalfCheetah and DelayedHopper tasks.

We studied more on the α -adaptation proposed in Section 5 and the behavior of sample-awareness over the learning phase. Fig. 5 shows the learning curve of α , $D_{JS}(\pi||q)$ and the policy entropy $\mathcal{H}(\pi)$, which are intertwined in the DAC objective function as seen in (6). In the case of DelayedHalfCheetah, α increases to one as time step goes on, and the initially nonzero JS divergence term $D_{JS}(\pi||q)$ diminishes to zero as time goes. This means that the sample action distribution is exploited in the early phase of learning, and DAC operates like SAC as time goes. On the other hand, in the case of DelayedHopper, the learned α gradually settle down around 0.5, and the JS divergence term $D_{JS}(\pi||q)$ is non-zero throughout the learning phase. Thus, it is seen that the proposed α -adaptation learns the weighting factor α with a completely different strategy depending on the task, and this leads to better overall performance for each task as seen in Fig. 4.

6.3. Analysis on the Change of q

We assumed that the sample action distribution q of the replay buffer \mathcal{D} is fixed for theoretical development and proof of diverse policy iteration in Section 4.3. However, q changes as iteration goes in real situation, so there exist a gap between the assumption and the real situation for DAC. Changing distribution was considered in some previous work. Hazan et al. (2019) considered the change of previous distributions to guarantee convergence, but they still have a common objective function (i.e., the entropy of state distribution d^π induced by policy π) to maximize. In our case, on the other hand, the objective function (3) itself changes over time as q changes, so it is difficult to show convergence with incorporation of the change of q . Thus, we assumed locally fixed q because q changes slowly when the buffer size is large. In order to investigate the impact of the lapse in the assumption and check the robustness of




Figure 4: Performance comparison on HumanoidStandup and Delayed Mujoco tasks (A zoomed version of the figure is available at Figure E.3 in Appendix E.)




Figure 5: Averaged learning curve for α -adaptation




Figure 6: Robustness against the change speed of q

DAC with respect to the change speed of q , we performed an additional study. In the study, we maintained the buffer size of the replay buffer \mathcal{D} as $N = 1000k$. Then, instead of using original q , i.e., the sample action distribution of whole \mathcal{D} , we now used q' , which is the sample action distribution of the latest N' samples (we call \mathcal{D}') stored in \mathcal{D} with $N' \leq 1000k$. The smaller N' is, the faster q' changes. Then, with others remaining the same, the ratio objective function $\hat{J}_{R^\alpha}(\eta)$ and the target value $\hat{V}(s_t)$ in DAC were changed to incorporate q' as

$$\begin{aligned} \hat{J}_{R^\alpha}(\eta) &= \mathbb{E}_{s_t \sim \mathcal{D}} [\alpha \mathbb{E}_{a_t \sim \pi_\theta} [\log R_\eta^\alpha(s_t, a_t)] \\ &\quad + (1 - \alpha) \mathbb{E}_{a_t \sim q'} [\log(1 - R_\eta^\alpha(s_t, a_t))]], \\ \hat{V}(s_t) &= \mathbb{E}_{a_t \sim \pi_\theta} [Q_\phi(s_t, a_t) \\ &\quad + \alpha \log R_\eta^\alpha(s_t, a_t) - \alpha \log \alpha \pi_\theta(a_t | s_t)] \\ &\quad + (1 - \alpha) \mathbb{E}_{a_t \sim q'} [\log R_\eta^\alpha(s_t, a_t) - \log \alpha \pi_\theta(a_t | s_t)], \end{aligned}$$

where s_t is still drawn from the original buffer \mathcal{D} and only q' considers samples distribution of \mathcal{D}' . Hence, s_t drawn from \mathcal{D} may not belong to \mathcal{D}' used to compute q' . So, we used generalization: To sample actions from q' for arbitrary states in \mathcal{D} , we learned q' by using variational auto-encoder. Fig. 6 shows the corresponding performance. As shown, the performance degrades as q' changes faster by decreasing N' from 1000k to 1k (Note that the original DAC is the case when $N' = 1000k$), but the performance is quite robust against the q change speed. Note that the performance is better than SAC even for $N' = 1k$.

We provided more results including the max average return table, more ablation study (control coefficient c , entropy coefficient β , and the effect of JS divergence) and the per-

formance comparison with various state-of-the-art RL algorithms in Appendix E. The results there also show that DAC yields top level performance.

7. Conclusion

In this paper, we have proposed a sample-aware entropy framework for off-policy RL to overcome the limitation of simple policy entropy regularization. With the sample-aware entropy regularization, we can achieve diversity gain by exploiting sample history in the replay buffer in addition to policy entropy for sample-efficient exploration. For practical implementation of sample-aware entropy regularized RL, we have used the ratio function to make computation of the sample action distribution from the replay buffer unnecessary, and have proposed the DAC algorithm with convergence proof. We have also provided an adaptation method for DAC to automatically control the ratio of the sample action distribution to the policy distribution. Numerical results show that the proposed DAC algorithm significantly outperforms other state-of-the-art RL algorithms.

8. Acknowledgement

This work was supported in part by the ICT R&D program of MSIP/IITP(2016-0-00563, Research on Adaptive Machine Learning Technology Development for Intelligent Autonomous Digital Companion) and in part by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Science, ICT & Future Planning(NRF2017R1E1A1A03070788).

References

- Achiam, J. and Sastry, S. Surprise-based intrinsic motivation for deep reinforcement learning. *arXiv preprint arXiv:1703.01732*, 2017.
- Agre, P. and Rosenschein, S. J. *Computational theories of interaction and agency*. Mit Press, 1996.
- Baldassarre, G. and Mirolli, M. *Intrinsically motivated learning in natural and artificial systems*. Springer, 2013.
- Bellemare, M., Srinivasan, S., Ostrovski, G., Schaul, T., Saxton, D., and Munos, R. Unifying count-based exploration and intrinsic motivation. In *Advances in Neural Information Processing Systems*, pp. 1471–1479, 2016.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- Burda, Y., Edwards, H., Storkey, A., and Klimov, O. Exploration by random network distillation. *arXiv preprint arXiv:1810.12894*, 2018.
- Chentanez, N., Barto, A. G., and Singh, S. P. Intrinsically motivated reinforcement learning. In *Advances in neural information processing systems*, pp. 1281–1288, 2005.
- Chinchuluun, A., Pardalos, P. M., Migdalas, A., and Pitsoulis, L. *Pareto optimality, game theory and equilibria*. Springer, 2008.
- Cover, T. M. and Thomas, J. A. *Elements of Information Theory*. Wiley, 2006.
- Degrís, T., White, M., and Sutton, R. S. Off-policy actor-critic. *arXiv preprint arXiv:1205.4839*, 2012.
- Dhariwal, P., Hesse, C., Klimov, O., Nichol, A., Plappert, M., Radford, A., Schulman, J., Sidor, S., Wu, Y., and Zhokhov, P. Openai baselines. <https://github.com/openai/baselines>, 2017.
- Eysenbach, B., Gupta, A., Ibarz, J., and Levine, S. Diversity is all you need: Learning skills without a reward function. *arXiv preprint arXiv:1802.06070*, 2018.
- Fox, R., Pakman, A., and Tishby, N. Taming the noise in reinforcement learning via soft updates. *arXiv preprint arXiv:1512.08562*, 2015.
- Fujimoto, S., van Hoof, H., and Meger, D. Addressing function approximation error in actor-critic methods. *arXiv preprint arXiv:1802.09477*, 2018.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.
- Gu, S., Lillicrap, T., Ghahramani, Z., Turner, R. E., and Levine, S. Q-prop: Sample-efficient policy gradient with an off-policy critic. *arXiv preprint arXiv:1611.02247*, 2016.
- Gu, S. S., Lillicrap, T., Turner, R. E., Ghahramani, Z., Schölkopf, B., and Levine, S. Interpolated policy gradient: Merging on-policy and off-policy gradient estimation for deep reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 3846–3855, 2017.
- Guo, Y., Oh, J., Singh, S., and Lee, H. Generative adversarial self-imitation learning. *arXiv preprint arXiv:1812.00950*, 2018.
- Haarnoja, T., Tang, H., Abbeel, P., and Levine, S. Reinforcement learning with deep energy-based policies. *arXiv preprint arXiv:1702.08165*, 2017.
- Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018a.
- Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., Abbeel, P., et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018b.
- Han, S. and Sung, Y. Dimension-wise importance sampling weight clipping for sample-efficient reinforcement learning. *International Conference on Machine Learning*, 2019.
- Hazan, E., Kakade, S., Singh, K., and Van Soest, A. Provably efficient maximum entropy exploration. In *International Conference on Machine Learning*, pp. 2681–2691. PMLR, 2019.
- Hong, Z.-W., Shann, T.-Y., Su, S.-Y., Chang, Y.-H., Fu, T.-J., and Lee, C.-Y. Diversity-driven exploration strategy for deep reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 10489–10500, 2018.
- Houthoofd, R., Chen, X., Duan, Y., Schulman, J., De Turck, F., and Abbeel, P. Vime: Variational information maximizing exploration. In *Advances in Neural Information Processing Systems*, pp. 1109–1117, 2016.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.

- Lopes, M., Lang, T., Toussaint, M., and Oudeyer, P.-Y. Exploration in model-based reinforcement learning by empirically estimating learning progress. In *Advances in neural information processing systems*, pp. 206–214, 2012.
- Martin, J., Sasikumar, S. N., Everitt, T., and Hutter, M. Count-based exploration in feature space for reinforcement learning. *arXiv preprint arXiv:1706.08090*, 2017.
- Mazouze, B., Doan, T., Durand, A., Hjelm, R. D., and Pineau, J. Leveraging exploration in off-policy algorithms via normalizing flows. *arXiv preprint arXiv:1905.06893*, 2019.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529–533, 2015.
- Nachum, O., Norouzi, M., Xu, K., and Schuurmans, D. Bridging the gap between value and policy based reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 2775–2785, 2017a.
- Nachum, O., Norouzi, M., Xu, K., and Schuurmans, D. Trust-pcl: An off-policy trust region method for continuous control. *arXiv preprint arXiv:1707.01891*, 2017b.
- Nielsen, F. On the jensen–shannon symmetrization of distances relying on abstract means. *Entropy*, 21(5):485, 2019.
- O’Donoghue, B., Munos, R., Kavukcuoglu, K., and Mnih, V. Combining policy gradient and q-learning. *arXiv preprint arXiv:1611.01626*, 2016.
- Pathak, D., Agrawal, P., Efros, A. A., and Darrell, T. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 16–17, 2017.
- Plappert, M., Houthoofd, R., Dhariwal, P., Sidor, S., Chen, R. Y., Chen, X., Asfour, T., Abbeel, P., and Andrychowicz, M. Parameter space noise for exploration. *arXiv preprint arXiv:1706.01905*, 2017.
- Puterman, M. L. and Brumelle, S. L. On the convergence of policy iteration in stationary dynamic programming. *Mathematics of Operations Research*, 4(1):60–69, 1979.
- Rawlik, K., Toussaint, M., and Vijayakumar, S. On stochastic optimal control and reinforcement learning by approximate inference. In *Twenty-Third International Joint Conference on Artificial Intelligence*, 2013.
- Santos, M. S. and Rust, J. Convergence properties of policy iteration. *SIAM Journal on Control and Optimization*, 42(6):2094–2115, 2004.
- Schulman, J., Chen, X., and Abbeel, P. Equivalence between policy gradients and soft q-learning. *arXiv preprint arXiv:1704.06440*, 2017a.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017b.
- Strehl, A. L. and Littman, M. L. An analysis of model-based interval estimation for markov decision processes. *Journal of Computer and System Sciences*, 74(8):1309–1331, 2008.
- Sugiyama, M., Suzuki, T., and Kanamori, T. *Density ratio estimation in machine learning*. Cambridge University Press, 2012.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. The MIT Press, Cambridge, MA, 1998.
- Tang, H., Houthoofd, R., Foote, D., Stooke, A., Chen, O. X., Duan, Y., Schulman, J., DeTurck, F., and Abbeel, P. # exploration: A study of count-based exploration for deep reinforcement learning. In *Advances in neural information processing systems*, pp. 2753–2762, 2017.
- Todorov, E. General duality between optimal control and estimation. In *2008 47th IEEE Conference on Decision and Control*, pp. 4286–4292. IEEE, 2008.
- Todorov, E., Erez, T., and Tassa, Y. Mujoco: A physics engine for model-based control. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pp. 5026–5033. IEEE, 2012.
- Toussaint, M. Robot trajectory optimization using approximate inference. In *Proceedings of the 26th annual international conference on machine learning*, pp. 1049–1056. ACM, 2009.
- Wang, Z., Bapst, V., Heess, N., Mnih, V., Munos, R., Kavukcuoglu, K., and de Freitas, N. Sample efficient actor-critic with experience replay. *arXiv preprint arXiv:1611.01224*, 2016.
- Watkins, C. J. and Dayan, P. Q-learning. *Machine Learning*, 8(3):279–292, 1992.
- Wu, Y., Mansimov, E., Grosse, R. B., Liao, S., and Ba, J. Scalable trust-region method for deep reinforcement learning using kronecker-factored approximation. In *Advances in neural information processing systems*, pp. 5279–5288, 2017.

Zheng, Z., Oh, J., and Singh, S. On learning intrinsic rewards for policy gradient methods. In *Advances in Neural Information Processing Systems*, pp. 4644–4654, 2018.

Ziebart, B. D. *Modeling purposeful adaptive behavior with the principle of maximum causal entropy*. PhD thesis, figshare, 2010.

Ziebart, B. D., Maas, A., Bagnell, J. A., and Dey, A. K. Maximum entropy inverse reinforcement learning. 2008.