

THE POINT TO WHICH SOFT ACTOR-CRITIC CONVERGES

Jianfei Ma

School of Mathematics and Statistics
Northwestern Polytechnical University
matrixfeeney@gmail.com

ABSTRACT

Soft actor-critic is a successful successor over soft Q-learning. While lived under maximum entropy framework, their relationship is still unclear. In this paper, we prove that in the limit they converge to the same solution. This is appealing since it translates the optimization from an arduous to an easier way. The same justification can also be applied to other regularizers such as KL divergence.

1 PRELIMINARIES

Consider a regularized infinite-horizon discounted MDP, defined by a tuple $(\mathcal{S}, \mathcal{A}, P, r, \rho_0, \gamma, \Delta)$, where \mathcal{S} is the state space, \mathcal{A} is the action space with finite cardinality $|\mathcal{A}|$, $p : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ is the transition probability distribution, $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function, assumed to be bounded $\rho_0 : \mathcal{S} \rightarrow \mathbb{R}$ is the distribution of the initial state s_0 , and $\gamma \in [0, 1)$ is the discount factor. We denote $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ as a stochastic policy. We restrict our attention on the regularizer $\Delta : \pi \rightarrow \mathbb{R}$, that is, being a function of the policy. Whenever noticed, $\mathcal{H}(s)$ abbreviates the entropy of $\pi(\cdot|s)$.

In SAC Haarnoja et al. (2018), the soft Bellman operator \mathcal{T}^π is defined as follows

$$\mathcal{T}^\pi Q(s_t, a_t) = r(s_t, a_t) + \gamma \mathbb{E}_{s_{t+1}} [V(s_{t+1})], \quad (1)$$

where

$$V(s_t) = \mathbb{E}_{a_t \sim \pi} [Q(s_t, a_t) - \eta \log \pi(a_t|s_t)] \quad (2)$$

It is not difficult to see that \mathcal{T}^π is a contraction by modifying the reward as $r(s, a) + \gamma \mathbb{E}_{s' \sim p} [\mathcal{H}(s')]$.

With repeatedly applying this operator over an arbitrary starting action-value function Q , it approaches to the soft value function.

2 CONVERGENCE ANALYSIS

In this section, we firstly study the regularized proxies from an optimization perspective, then state the soft policy iteration, and finally point out the convergence result.

2.1 OPTIMIZING WITH REGULARIZATION

Define the regularized state-value function as

$$\tilde{V}^\pi(s) = \mathbb{E} \left[\sum_{l=0}^{\infty} \gamma^l (r_{t+l} + \eta \Delta_{t+l}) | s_0 = s \right] \quad (3)$$

where η is the temperature parameter, usually positive, determining the relative importance of the regularization term against the reward.

The optimal regularized value function $\tilde{V}^*(s)$ should satisfy the corresponding optimal Bellman equation¹ for all $s \in \mathcal{S}$

$$\tilde{V}^*(s) = \sup_{\pi} \sum_{a \in \mathcal{A}} \pi(a|s) [r(s, a) + \eta \Delta(s) + \gamma \mathbb{E}_{s' \sim p} [\tilde{V}^*(s')]] \quad (4)$$

¹Though we use the summation for simplicity, it can be readily replaced by the integral.

For $\Delta(s) = \mathcal{H}(\pi(\cdot|s))$, we have

Lemma 1. For all $(s, a) \in \mathcal{S} \times \mathcal{A}$, the optimal value function $\tilde{V}^*(s)$ and the optimal policy $\tilde{\pi}^*(a|s)$, satisfy

$$\begin{aligned} \tilde{V}^*(s) &= \eta \log \sum_{a \in \mathcal{A}} \exp \frac{1}{\eta} (r(s, a) + \gamma \mathbb{E}_{s' \sim p} [\tilde{V}^*(s')]) \\ \tilde{\pi}^*(a|s) &= \frac{\exp \frac{1}{\eta} (r(s, a) + \gamma \mathbb{E}_{s' \sim p} [\tilde{V}^*(s')])}{\sum_{a \in \mathcal{A}} \exp \frac{1}{\eta} (r(s, a) + \gamma \mathbb{E}_{s' \sim p} [\tilde{V}^*(s')])} \end{aligned} \quad (5)$$

from which we define an auxiliary optimal action-value function (not the true one)

$$\tilde{Q}^*(s, a) = r(s, a) + \gamma \mathbb{E}_{s' \sim p} [\tilde{V}^*(s')] \quad (6)$$

Proposition 1. For any $V : \mathcal{S} \rightarrow \mathbb{R}$ that satisfies $V(s) \leq \tilde{V}^*(s)$ for all $s \in \mathcal{S}$, then

$$Q(s, a) \triangleq r(s, a) + \gamma \mathbb{E}_{s' \sim p} [V(s')] \leq \tilde{Q}^*(s, a) \quad (7)$$

2.2 SOFT POLICY ITERATION

Consider the softmax policy class Π

Lemma 2. (Soft Policy Iteration). Repeatedly application of soft policy evaluation (Haarnoja et al., 2018, Lemma 1) and soft policy improvement (Haarnoja et al., 2018, Lemma 2) to any $\pi \in \Pi$ converges to a policy π^* such that $Q^{\pi^*}(s, a) \geq Q^\pi(s, a)$ for all $\pi \in \Pi$ and $(s, a) \in \mathcal{S} \times \mathcal{A}$.

Combining all the aforementioned statements, we formally arrive at

Theorem 2. For any initial policy π_0 and corresponding action-value function Q^{π_0} , the convergent points induced by SPI 2 satisfy $Q^{\pi^*}(s, a) = \tilde{Q}^*(s, a)$ and $\pi^* = \tilde{\pi}^*$.

Proof. The backward direction is obvious as shown in (Haarnoja et al., 2018, proof of Theorem 1), that is, $Q^{\pi^*} \geq \tilde{Q}^*$. We only need show the other direction. Since Q^{π^*} is the fixed point of the soft Bellman operator \mathcal{T}^{π^*} , thus it must satisfy the Bellman equation as defined in Equation 7 with a value function V^{π^*} . And since V^* is the regularized value function that at most can be obtained, it must have $V^{\pi^*} \leq V^*$. By Proposition 1, it follows that $Q^{\pi^*} \leq \tilde{Q}^*$. And since $\pi^* \in \Pi$, it immediately follows that $\pi^* = \tilde{\pi}^*$. \square

This theorem connects LogSumExp optimization to policy evaluation and improvement, providing an alternative approach. It links SQL Haarnoja et al. (2017) and SAC, with SAC being superior in optimization. It allows for optimizing regularizers like KL divergence using a different procedure. With a prior $\tilde{\pi}$ on the policy, setting $\Delta(s) = -D_{\text{KL}}(\pi|\tilde{\pi})$ allows us to derive conservative optimal points and define the conservative Bellman operator, using similar justifications

$$\begin{aligned} V^{\pi^*}(s) &= \eta \log \sum_{a \in \mathcal{A}} \tilde{\pi}(a|s) \exp \frac{1}{\eta} (r(s, a) + \gamma \mathbb{E}_{s' \sim p} [V^{\pi^*}(s')]) \\ \pi^*(a|s) &= \frac{\tilde{\pi}(a|s) \exp \frac{1}{\eta} (r(s, a) + \gamma \mathbb{E}_{s' \sim p} [V^{\pi^*}(s')])}{\sum_{a \in \mathcal{A}} \tilde{\pi}(a|s) \exp \frac{1}{\eta} (r(s, a) + \gamma \mathbb{E}_{s' \sim p} [V^{\pi^*}(s')])} \end{aligned} \quad (8)$$

$$\begin{aligned} \mathcal{T}^\pi Q(s_t, a_t) &= r(s_t, a_t) + \gamma \mathbb{E}_{s_{t+1}} [V(s_{t+1})], \\ V(s_t) &= \mathbb{E}_{a_t \sim \pi} [Q(s_t, a_t) - \eta \log \frac{\pi(a_t|s_t)}{\tilde{\pi}(a_t|s_t)}] \end{aligned} \quad (9)$$

Intervening between policy evaluation based on the conservative Bellman operator, and policy improvement with the softmax policy of the conservative action-value function, we are guaranteed to converge to the optimal policy.

URM STATEMENT

The authors acknowledge that at least one key author of this work meets the URM criteria of ICLR 2023 Tiny Papers Track.

REFERENCES

- Mohammad Gheshlaghi Azar, Vicenç Gómez, and Hilbert J. Kappen. Dynamic policy programming. *J. Mach. Learn. Res.*, 13:3207–3245, 2012. doi: 10.5555/2503308.2503344. URL <https://dl.acm.org/doi/10.5555/2503308.2503344>.
- Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. Reinforcement learning with deep energy-based policies. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1352–1361. PMLR, 2017. URL <http://proceedings.mlr.press/v70/haarnoja17a.html>.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In Jennifer G. Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1856–1865. PMLR, 2018. URL <http://proceedings.mlr.press/v80/haarnoja18b.html>.

A PROOF OF LEMMA 1

Following the sketch of Azar et al. (2012), we define the Lagrangian function $\mathcal{L}(s; \lambda) : \mathcal{S} \rightarrow \mathbb{R}$

$$\mathcal{L}(s; \lambda) = \sum_{a \in \mathcal{A}} \pi(a|s) [r(s, a) + \gamma \mathbb{E}_{s' \sim p} [\tilde{V}^*(s')]] + \eta \mathcal{H}(s) - \lambda \left(\sum_{a \in \mathcal{A}} \pi(a|s) - 1 \right) \quad (10)$$

Since the objective is linear and \mathcal{H} is strictly-concave in π , and the probability simplex is at least non-empty, thus slater condition is satisfied, which implies the optimum by solving

$$0 = \frac{\partial \mathcal{L}(s; \lambda)}{\partial \pi(a|s)} = r(s, a) + \gamma \mathbb{E}_{s' \sim p} [\tilde{V}^*(s')] - \eta \log \pi(a|s) - \eta - \lambda \quad (11)$$

The solution is

$$\pi^* = \exp \left(-\frac{\lambda}{\eta} - 1 \right) \exp \frac{1}{\eta} (r(s, a) + \gamma \mathbb{E}_{s' \sim p} [\tilde{V}^*(s')]) \quad (12)$$

With the equality constraint

$$\sum_{a \in \mathcal{A}} \pi^*(a|s) = 1 \quad (13)$$

by applying log transformation on both sides, we can solve for the multiplier as

$$\lambda = \eta \log \sum_{a \in \mathcal{A}} \exp \frac{1}{\eta} (r(s, a) + \gamma \mathbb{E}_{s' \sim p} [\tilde{V}^*(s')]) - \eta \quad (14)$$

inserting which into Equation 12, we get

$$\tilde{\pi}^*(a|s) = \frac{\exp \frac{1}{\eta} (r(s, a) + \gamma \mathbb{E}_{s' \sim p} [\tilde{V}^*(s')])}{\sum_{s \in \mathcal{A}} \exp \frac{1}{\eta} (r(s, a) + \gamma \mathbb{E}_{s' \sim p} [\tilde{V}^*(s')])} \quad (15)$$

And finally plug this result into Equation 4, we get

$$\tilde{V}^*(s) = \eta \log \sum_{a \in \mathcal{A}} \exp \frac{1}{\eta} (r(s, a) + \gamma \mathbb{E}_{s' \sim p} [\tilde{V}^*(s')]) \quad (16)$$