

Symmetric Q-learning: Reducing Skewness of Bellman Error in Online Reinforcement Learning

Motoki Omura¹, Takayuki Osa^{1,2}, Yusuke Mukuta^{1,2}, Tatsuya Harada^{1,2}

¹The University of Tokyo,

²RIKEN

{omura, osa, mukuta, harada}@mi.t.u-tokyo.ac.jp

Abstract

In deep reinforcement learning, estimating the value function to evaluate the quality of states and actions is essential. The value function is often trained using the least squares method, which implicitly assumes a Gaussian error distribution. However, a recent study suggested that the error distribution for training the value function is often skewed because of the properties of the Bellman operator, and violates the implicit assumption of normal error distribution in the least squares method. To address this, we proposed a method called Symmetric Q-learning, in which the synthetic noise generated from a zero-mean distribution is added to the target values to generate a Gaussian error distribution. We evaluated the proposed method on continuous control benchmark tasks in MuJoCo. It improved the sample efficiency of a state-of-the-art reinforcement learning method by reducing the skewness of the error distribution.

Introduction

Deep reinforcement learning (RL) has shown remarkable performance in control (Haarnoja et al. 2018) and gameplay (Mnih et al. 2013). In RL, it is necessary to estimate a value function to evaluate states and actions. In many cases, the value function is learned through least squares, implicitly assuming that the error distribution is a normal distribution through maximum likelihood estimation. However, during learning, the Bellman operator is used to derive the estimated value for the target. Because of the properties of the Bellman operator, the error distribution can become skewed, which depreciates the performance by not meeting the assumption of a normal error distribution while performing least squares. This issue was addressed in (Garg et al. 2023) by learning the value function using maximum likelihood estimation with a nonnormal error distribution. However, the error distribution can vary depending on the task or learning algorithm, making it unsuitable when the error distribution is unknown.

Therefore, in this study, we propose a method to reduce the skewness of arbitrary error distributions. We corrected the skewed error distributions to be symmetric by adding noise to the target values. As shown in the left panel of Figure 1, we learned the noise η distribution which is the inverse

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Bellman error, negative values of correction noise, and corrected Bellman error from Symmetric REDQ on Hopper-v2. Left: The blue histogram shows the distribution of Bellman errors. The orange histogram represents the distribution of the negative noise added to reduce skewness. It can be observed that the noise distribution fits well with the negative Bellman errors. Right: The green histogram represents the distribution of Bellman errors after adding correction noise. The skewness decreased compared to the blue distribution.

of the error distribution $y - f(x)$. This resulted in a more symmetric error distribution with increased symmetry when noise was added to the target, as shown in the right panel of the plot. By applying least squares to this symmetric error distribution, we can better satisfy the assumption of the error distribution and improve performance. The effectiveness of the proposed method was verified in combination with the Soft Actor-Critic (SAC) (Haarnoja et al. 2018) and Randomized Ensemble Double Q-learning (REDQ) (Chen et al. 2021), which we refer to as Symmetric SAC (SymSAC) and Symmetric REDQ (SymREDQ), respectively. We evaluated our proposed method on five challenging tasks in MuJoCo (Todorov, Erez, and Tassa 2012; Brockman et al. 2016) and achieved comparable or better sample efficiency than that of state-of-the-art online RL methods. We also confirm that our method corrects skewed error distributions to become symmetric by visualizing the error distributions during the learning process. The contributions of this study are as follows.

- We propose a method of adding noise to target variables and reducing the skewness of error distributions, which is a problem when using least squares.
- We evaluated the proposed method in continuous control benchmark tasks and showed that the proposed method can reduce the skewness of error distribution.

- We empirically demonstrated that the proposed method can improve the sample efficiency of REDQ, a state-of-the-art model-free RL algorithm.

Background

Reinforcement Learning

We consider the RL problem under a Markov decision process (MDP) defined by the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma, d)$. \mathcal{S} represents the state space, s represents a state, \mathcal{A} represents the action space, a represents an action, $\mathcal{P}(s_{t+1} | s_t, a_t)$ is the transition probability density, $r(s, a)$ is the reward function, γ is the discount factor, and $d(s_0)$ is the probability density of the initial state. The policy $\pi(a | s)$ is defined as the conditional probability density of a given action. The goal of RL is to identify the policy that maximizes the expected cumulative discounted reward $\mathbb{E}[R_0 | \pi]$, where R_t is the return, given by $R_t = \sum_{k=t}^T \gamma^{k-t} r(s_k, a_k)$.

Bellman Error Distribution

In RL, the Q-function is estimated based on the fact that it satisfies the Bellman equation:

$$Q^\pi(s, a) = \mathbb{E}_{s' \sim \mathcal{P}, a' \sim \pi} [r(s, a) + \gamma Q^\pi(s', a')] \quad (1)$$

Based on the Bellman equation, the Bellman operator is often used to estimate the Q-function and the Bellman operator is known to be contractive. Therefore, minimizing the mean squared error is a common approach for learning the Q-function.

$$\mathcal{L}(\theta) = \mathbb{E}_{s, a \sim \rho} [(y - Q^\pi(s, a; \theta))^2] \quad (2)$$

$\rho(s, a)$ is a probability distribution over states s and actions a , θ denotes the parameters of the Q-function approximator, and the target value y is computed as:

$$y = \mathbb{E}_{s' \sim \mathcal{P}, a' \sim \pi} [r(s, a) + \gamma Q^\pi(s', a'; \theta) | s, a] \quad (3)$$

The prediction error between y and the Q-function is known as the Bellman error.

In the least squares method, it is implicitly assumed that the error distribution follows a normal distribution through maximum likelihood estimation. However, in (Garg et al. 2023), it is explained based on extreme value theory (Fisher and Tippett 1928; Mood 1950) that the distribution of the Bellman error can become skewed when using policies that aim to maximize the Q-value owing to the influence of the maximum operator. They then used a Gumbel regression, which is a maximum likelihood estimation that assumes a Gumbel distribution as the error distribution.

However, that method has several issues. While extreme value theory assumes that the maximum is taken from independently sampled values, in the context of RL, the Q-values in the maximum operator are dependent on state-action pairs. Therefore, it is difficult to satisfy this assumption and obtain an exact extreme value distribution such as the Gumbel distribution. In fact, the Bellman error distribution deviates from the Gumbel distribution during learning,




Figure 2: The pre-corrected Bellman error at three different steps when learning Walker2d with SymREDQ.

depending on the task (Garg et al. 2023). Furthermore, the error distribution is influenced by many factors such as double Q-learning (Hasselt 2010; van Hasselt, Guez, and Silver 2016; Fujimoto, van Hoof, and Meger 2018) and the target network (Mnih et al. 2015), and may change during the learning process as shown in Figure 2, making it difficult to make accurate assumptions regarding the error distribution. Therefore, this study proposes a flexible approach to assumptions regarding error distribution.

Symmetric Q-learning

In this section, we introduce symmetric Q-learning, a technique that corrects the error distribution, making it symmetric and close to a normal distribution by adding noise to the target variable.

Assumption of Error Distribution

Our motivation is to approach the error distribution as a normal distribution to satisfy the assumption of the least squares method. Therefore, we first demonstrate that the error distribution is assumed to follow a normal distribution.

Here, we explain that the least squares method is equivalent to maximum likelihood estimation assuming Gaussian errors. Assuming that we are given a training dataset consisting of pairs of observed values x and the corresponding target values y , we consider the problem of predicting the value of y for a new value of x . We assume that the target variable y is the sum of a function $f(x; \theta)$ with parameters θ and Gaussian noise.

$$y = f(x; \theta) + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma^2) \quad (4)$$

In this case, the probability of y given x , θ , and σ follows a normal distribution with mean $f(x; \theta)$.

$$p(y | x, \theta, \sigma) = \mathcal{N}(y | f(x; \theta), \sigma^2) \quad (5)$$

Therefore, the expected log-likelihood is as follows:

$$\begin{aligned} & \mathbb{E}_{y, x \sim p_{data}} [\log p(y | x, \theta, \sigma)] \\ &= \mathbb{E}_{y, x \sim p_{data}} \left[-\frac{1}{2} \log 2\pi\sigma^2 - \frac{1}{2\sigma^2} (y - f(x; \theta))^2 \right] \end{aligned} \quad (6)$$

If we consider maximizing the expected log-likelihood, the loss function with respect to θ can be expressed as follows:

$$\mathcal{L}(\theta) = \mathbb{E}_{y, x \sim p_{data}} [(y - f(x; \theta))^2] \quad (7)$$

This is equivalent to the loss function used in the least squares method.

Thus, in the least squares method, it is implicitly assumed through a maximum likelihood estimation that the error distribution follows a normal distribution. However, as explained in the previous section, the Bellman error can exhibit a distribution that deviates from the Gaussian distribution. As seen in the left plot of Figure 1, even in REDQ, which has state-of-the-art sample efficiency and uses the least squares method, the pre-corrected Bellman error is also skewed away from a normal distribution. Thus, in online RL, the error distribution varies depending on the task or algorithm. Therefore, we propose an algorithm that can be applied to any error distribution that changes during the learning process.

Normalization of Error Distribution

As seen in the previous section, it is possible for the Bellman error ϵ to not follow either a normal or a Gumbel distribution in practice. Therefore, we consider controlling the error distribution by adding zero-mean noise η to the target variable to ensure that the error distribution follows the normal distribution. We assume that the noise η follows the probabilistic model $q(\eta; \omega)$ with parameter ω . The method used to estimate ω is described in the following section.

$$y + \eta = f(x; \theta) + \epsilon + \eta, \eta \sim q(\eta; \omega) \quad (8)$$

If we assume that the sum of the original error ϵ and the added noise η follows a normal distribution, then the probability of $y + \eta$ follows the normal distribution with mean $f(x; \theta)$.

$$\epsilon + \eta \sim \mathcal{N}(0, \sigma^2) \quad (9)$$

$$p(y + \eta) = \mathcal{N}(f(x; \theta), \sigma^2) \quad (10)$$

Following the same approach as in the previous section, the expected log likelihood and loss function are given as

$$\begin{aligned} & \mathbb{E}_{y, x \sim p_{data}, \eta \sim q} [\log p(y + \eta \mid x, \theta, \sigma, \omega)] \\ &= \mathbb{E}_{y, x \sim p_{data}, \eta \sim q} \left[-\frac{1}{2} \log 2\pi\sigma^2 - \frac{1}{2\sigma^2} (y - f(x; \theta) + \eta)^2 \right] \end{aligned} \quad (11)$$

$$\mathcal{L}(\theta) = \mathbb{E}_{y, x \sim p_{data}, \eta \sim q} [(y - f(x; \theta) + \eta)^2] \quad (12)$$

When the error $y - f(x; \theta) + \eta$ follows the normal distribution, minimizing the mean squared error satisfies the assumption of an error distribution, allowing for more accurate estimates. Therefore, it is necessary to learn the distribution of η which causes $y - f(x; \theta) + \eta$ to approach the normal distribution.

Before explaining the learning of the distribution of η , it should be noted that adding zero-mean noise η does not introduce any bias into the learning of θ . This is easily demonstrated through the following manipulations:

$$\begin{aligned} \mathcal{L}(\theta) &= \mathbb{E}_{y, x \sim p_{data}, \eta \sim q} [(y - f(x; \theta) + \eta)^2] \\ &= \mathbb{E}_{y, x \sim p_{data}, \eta \sim q} [y^2 + f(x; \theta)^2 + \eta^2 - 2yf(x; \theta) \\ &\quad - 2\eta f(x; \theta) + 2y\eta] \\ &= \mathbb{E}_{y, x \sim p_{data}, \eta \sim q} [(y - f(x; \theta))^2 + \eta^2] \end{aligned} \quad (13)$$

Because η^2 does not affect the learning of θ , adding η to the target value does not introduce any bias into the estimator or change the expectation of the mean squared error.

However, η affects the error distribution shape. Let us consider skewness as an indicator of symmetry. By substituting the pre-corrected error $y - f(x)$ with ϵ , the skewness of the post-corrected error distribution can be expressed as follows:

$$\begin{aligned} Skewness(\epsilon + \eta) &= \frac{\mu_3(\epsilon + \eta)}{\sigma^3(\epsilon + \eta)} \\ &= \frac{\mu_3(\epsilon) + \mu_3(\eta)}{(\sigma^2(\epsilon) + \sigma^2(\eta))^{3/2}} \end{aligned} \quad (14)$$

μ_3 and σ represent the third central moment and the variance, respectively. This equation indicates that the skewness, and hence the symmetry, of the error distribution after correction is influenced by the third moment of η . Moreover, correcting the distribution can reduce the skewness and increase the symmetry when the third moment of η has the opposite sign to that of ϵ .

It is expected that learning until convergence while adding zero-mean noise will result in no change in the final return because adding noise does not introduce any bias. However, online RL is a non-stationary learning process, and it is difficult to always achieve sufficient learning during the learning process. Furthermore, bias accumulates through bootstrapping of target values. Therefore, improving the learning efficiency through distribution correction during the learning process may lead to improved performance, including the final return. The following section describes the learning of the distribution of η .

Training of Noise Distribution

Any arbitrary distribution can be used for the probability model $q(\eta; \omega)$, and the parameter ω is estimated in a way that $\epsilon + \eta$ approaches the normal distribution. However, it is difficult to estimate ω accurately so that it becomes a normal distribution. Therefore, we considered making the distribution symmetric.

When η follows a distribution of $-\epsilon$, that means $\eta \sim p_\epsilon(-\eta)$, $\mu_3(\eta) = -\mu_3(\epsilon)$ and the skewness is zero. Therefore, the skewness of $\epsilon + \eta$ decreases and approaches a symmetric distribution by making the distribution of η closer to a distribution of $-\epsilon$. In other words, we bring the probability model $q(\eta; \omega)$ closer to $p_\epsilon(-\eta)$. To do this, we estimate the parameters ω by maximizing the likelihood of $y - f(x; \theta)$ with respect to p_ϵ . The specific steps are shown below.

Assuming that η follows a distribution of $-\epsilon$, then conversely, ϵ follows a distribution of $-\eta$:

$$\epsilon \sim p_\eta(-\epsilon) \quad (15)$$

Then, the probability of ϵ can be expressed as follows:

$$p_\epsilon(\epsilon) = p_\eta(-\epsilon) \simeq q(-\epsilon; \omega) \quad (16)$$

From $\epsilon = y - f(x; \theta)$, the likelihood of the observed sample $y - f(x; \theta)$ becomes $q(-(y - f(x; \theta)); \omega)$, and the loss function for maximizing this likelihood can be expressed as follows:

$$L(\omega) = \mathbb{E}_{y, x \sim p_{data}} [-\log q(-(y - f(x; \theta)); \omega)] \quad (17)$$

By minimizing this loss function, $q(\eta; \omega)$ approaches $p_\epsilon(-\eta)$ and η becomes noise, which reduces the skewness of ϵ .

As a result of the aforementioned learning, it was demonstrated that the skewness of $\epsilon + \eta$ can approach zero. However, a zero skewness does not necessarily imply that the distribution is symmetric. Nevertheless, it is shown from the following formula that the error distribution becomes symmetric:

$$\begin{aligned} p_{\epsilon+\eta}(x) &= p_\epsilon(x) * p_\eta(x) \\ &= \int_{-\infty}^{\infty} p_\epsilon(u) p_\eta(x - u) du \\ &= \int_{-\infty}^{\infty} p_\epsilon(u) p_\epsilon(u - x) du \\ &= \int_{-\infty}^{\infty} p_\epsilon(v + x) p_\epsilon(v) dv, (v = u - x) \\ &= \int_{-\infty}^{\infty} p_\eta(-x - v) p_\epsilon(v) dv \\ &= p_{\epsilon+\eta}(-x) \end{aligned} \quad (18)$$

$p_X(x)$ is the probability density function for the random variable X , and $p_\eta(x) = p_\epsilon(-x)$ when η follows a distribution of $-\epsilon$. From the above formula, it can be seen that the error distribution $p_{\epsilon+\eta}(x)$ is symmetric with respect to $x = 0$.

Practical Algorithms

For the distribution $q(\eta; \omega)$, it is preferable to have a distribution that can represent a skewed distribution with sufficiently wide support. Therefore, as a practical algorithm, we approximate $q(\eta; \omega)$ with the Gaussian mixture models (GMM) by employing variational Bayes expectation maximization (Bishop 2006), which results in a more stable parameter estimation than that of the expectation maximization algorithm based on maximum likelihood estimation. GMM can approximate skewed distributions such as the Gumbel distribution with high accuracy and supports the entire real line. During the training process, we iteratively update the GMM parameters to approach the distribution of negative Bellman errors using the same samples used for Q-function learning. To prevent overfitting the samples, we use a hyperparameter k to update the parameters for every k steps. Because the Bellman error is a one-dimensional data, estimating the GMM has a relatively low computational cost.

Algorithm 1: Symmetric REDQ

Initialize policy parameters θ , N Q-function parameters ϕ_i , $i = 1, \dots, N$, GMM parameters ω , empty replay buffer \mathcal{D} . Set target parameters $\phi_{target, i} \leftarrow \phi_i$, for $i = 1, 2, \dots, N$

repeat

Take one action $a_t \sim \pi_\theta(\cdot | s_t)$. Observe reward r_t , new state s_{t+1} .

Add data to buffer: $\mathcal{D} \leftarrow \mathcal{D} \cup \{(s_t, a_t, r_t, s_{t+1})\}$

for G updates **do**

Sample a mini-batch $B = \{(s, a, r, s')\}$ from \mathcal{D}

Sample a set \mathcal{M} of M distinct indices from $\{1, 2, \dots, N\}$

Compute the Q target y (same for all of the N Q-functions):

$$y = r + \gamma \left(\min_{i \in \mathcal{M}} Q_{\phi_{target, i}}(s', \tilde{a}') - \alpha \log \pi_\theta(\tilde{a}' | s') \right), \tilde{a}' \sim \pi_\theta(\cdot | s')$$

for $i = 1, \dots, N$ **do**

Compute error $\epsilon_i = y - Q_{\phi_i}(s, a)$

Every k steps update ω using $-\epsilon$ by minimizing (17)

for $i = 1, \dots, N$ **do**

Sample η from GMM $q(\eta; \omega)$

Update ϕ_i with gradient descent using

$$\nabla_\phi \frac{1}{|B|} \sum_{(s, a, r, s') \in B} (y - Q_{\phi_i}(s, a) + \eta)^2$$

Update target networks with

$$\phi_{target, i} \leftarrow \rho \phi_{target, i} + (1 - \rho) \phi_i$$



Update policy parameters θ with gradient ascent using

$$\nabla_\phi \frac{1}{|B|} \sum_{s \in B} \left(\frac{1}{N} \sum_{i=1}^N Q_{\phi_i}(s, \tilde{a}_\theta(s)) - \alpha \log \pi_\theta(\tilde{a}_\theta(s) | s) \right), \tilde{a}_\theta(s) \sim \pi_\theta(\cdot | s)$$

until end

To ensure that η becomes unbiased noise, the distribution is shifted such that the mean becomes zero each time the parameters of GMM are updated. One problem with adding noise to the error is that it can increase the error variance, which can be reduced by ensembling. Therefore, when combining this method with algorithms such as REDQ, it is recommended to use a larger ensemble to adjust the variance of the Q-function

The pseudocode for the SymREDQ algorithm, which combines the proposed method with REDQ, is presented in Algorithm 1. When the ensemble size $N = 2$, in-target minimization parameter $M = 2$, and update-to-data (UTD) ratio $G = 1$, it becomes SymSAC. When $N = 1$, $M = 1$, and $G = 1$, it becomes SymSAC without an ensemble. The pseudocode for GMM training is shown in the appendix.

Figure 3: Comparison of SymSAC, SAC and \mathcal{X} -SAC without ensembles for UTD=1Figure 4: Comparison of SymREDQ, REDQ and \mathcal{X} -REDQ for UTD=20.

Experiments

In our experiments, we investigated the following:

- Does distribution correction improve performance?
- Can the error distribution be corrected to a symmetric distribution?

To answer these questions, we implemented the proposed method in addition to prevalent RL methods and evaluated their performance in continuous control tasks.

Setups

We conducted experiments using five challenging tasks in MuJoCo: Humanoid-v2, Walker2d-v2, Hopper-v2, Ant-v2 and HalfCheetah-v2. The average return in the tasks was evaluated by performing independent trials using five random seeds for each task and algorithm. We experimented with two settings: a simple setting without an ensemble and a setting that increased sample efficiency using an ensemble.

In the former, SymSAC, which combined the proposed method with SAC, was compared with SAC and \mathcal{X} -SAC from (Garg et al. 2023). All algorithms used only one critic network to eliminate the effects of clipped double Q-learning or ensemble.

In the latter, SymREDQ, which combines the proposed method with REDQ, was compared with REDQ and \mathcal{X} -REDQ. \mathcal{X} -REDQ is \mathcal{X} -SAC with the addition of ensemble and in-target minimization. To verify sample efficiency, experiments were conducted with the UTD ratio set to 20, as in the REDQ paper. Thus 20 updates were performed per iteration. In SymREDQ, the size of the ensemble was set to 20, instead of the default 10 of REDQ, to reduce variance.

The parameter k , which represents the parameter updating frequency of the distribution of η , was set to the value that showed the best performance between 1, 2, and 3. Regarding \mathcal{X} -SAC and \mathcal{X} -REDQ, the temperature parameter β was

chosen from between 2, 5, 10, and 20 in the same way. Details of the hyperparameters and performance changes based on k , network size and other hyperparameters are provided in the appendix.

Comparative Evaluation

Comparison at UTD=1 without ensemble The average return for the case in which UTD is 1 and there is no ensemble is shown in Figure 3. In Hopper-v2, HalfCheetah-v2 and Ant-v2, our proposed SymSAC method achieved returns similar to those of SAC. For Walker2d-v2 and Humanoid-v2, SymSAC outperformed SAC in terms of returns. \mathcal{X} -SAC outperformed the other methods only on HalfCheetah-v2, a relatively simple task, but fell far short on the other tasks.

Comparison at UTD=20 with ensemble The average return for the case in which UTD is 20 and the ensemble is used is shown in Figure 4. The final average returns were equivalent to or slightly higher than those of the existing methods. The learning speed was also equivalent to or faster, with clear improvements in certain tasks. The numbers of iterations required to achieve a certain return level are listed in Table 1. In Hopper-v2, HalfCheetah-v2, and Ant-v2, the sample efficiency was significantly improved. In addition, Table 2 shows the final average return and the average return for half of the iterations, and the standard error is described in the appendix. In both cases, SymREDQ achieved equivalent or higher returns compared to REDQ and \mathcal{X} -REDQ.

Discussion In SymREDQ, the improvement in learning speed was found to be more significant than that in the final return. This is because the Q-function update amplitude is large in the early stages of learning and the error distribution can change significantly during the learning process. However, using the corrected error distribution with the least

Score	SymREDQ	REDQ	\mathcal{X} -REDQ
Humanoid at 5000	172K	255K	-
Walker2d at 3500	77K	141K	-
Hopper at 3000	34K	46K	-
Ant at 5000	101K	153K	-
HalfCheetah at 8000	131K	196K	-

Table 1: Sample efficiency comparison of SymREDQ, REDQ and \mathcal{X} -REDQ. The numbers indicate the amount of data collected until the specified performance level is reached.

Amount of data	SymREDQ	REDQ	\mathcal{X} -REDQ
Humanoid at 150K	4414	2904	81
Humanoid at 300K	5252	4606	82
Walker2d at 150K	4289	3715	75
Walker2d at 300K	4836	4587	127
Hopper at 62K	3413	3356	1433
Hopper at 125K	3428	3529	1845
Ant at 150K	5241	4910	320
Ant at 300K	5951	5464	1924
HalfCheetah at 150K	8358	7328	5639
HalfCheetah at 300K	9865	9138	7365

Table 2: Performance comparison of SymREDQ, REDQ and \mathcal{X} -REDQ. This shows the final performance and the performance with half the amount of data collected. Mainly SymREDQ showed significant improvement in performance with half the amount of data collected.

squares method leads to stable learning, resulting in an improved learning speed.

We observed that SymSAC and SAC did not show significant improvement, unlike SymREDQ and REDQ. This is likely due to the lack of variance reduction by ensembling in SymREDQ. However, because REDQ is state-of-the-art in terms of sample efficiency and the performance of SymREDQ improved in comparison, the effectiveness of the proposed method is considered sufficient.

In the case of \mathcal{X} -SAC and \mathcal{X} -REDQ, decent learning was only achieved for Hopper-v2 and HalfCheetah-v2, which are relatively simple tasks. To ensure proper learning, more detailed hyperparameter tuning and algorithm improvement are necessary. While extreme Q-learning-based methods were also developed to address the problem of the Bellman error distribution, the experimental results demonstrated that the proposed method was more stable than extreme Q-learning-based methods.

Distribution Plot

Using SymREDQ, we visualized the actual pre-corrected Bellman error, noise, and post-corrected Bellman error during training. In Figure 5, we show how errors were corrected in the steps in which the pre-corrected Bellman error was distorted for each task. During the training process, we ensured that the distribution of noise η had a mean value of zero. However, for visualization purposes, we shifted the mean of the noise distribution to be equal to the mean of

the Bellman error in order to assess the fitting performance to the Bellman error.

In the figure above, the blue histogram shows the distribution of the skewed Bellman errors. The orange histogram represents the distribution of negative noise added to reduce skewness. The distribution of noise η closely approximates the inverse distribution of the pre-corrected Bellman error $y - f(x)$, meaning that the distribution of η approaches the distribution of $-\epsilon$. In the figure below, the green histogram represents the distribution of the Bellman errors after adding the correction noise η . The skewness decreased and the symmetry of the error increased compared with the blue distribution. The Gumbel distribution assumed in extreme Q-learning was a right-skewed distribution. However, the pre-corrected error distribution (blue) in Humanoid and HalfCheetah showed a left-skewed distribution, and not meeting the assumption may be leading to a decrease in performance. The figures for other steps are provided in the appendix.

Related Work

In our study, we addressed the problem of the skewed error distribution in online RL. Therefore, we introduce several representative approaches for online RL and approaches related to the error distribution.

Online Reinforcement Learning

In online RL, the goal is to maximize rewards through iterations with the environment. While table-based RL with methods such as Q-learning (Watkins and Dayan 1992) and SARSA (Rummery and Niranjan 1994) has shown promising results, the combination of deep learning with Q-learning in DQN (Mnih et al. 2013) has achieved performance surpassing human-level capabilities in game playing. For continuous action tasks such as robotics, policy gradient-based methods such as TRPO (Schulman et al. 2015) and PPO (Schulman et al. 2017), as well as actor-critic-based methods such as SAC (Haarnoja et al. 2018) and TD3 (Fujimoto, van Hoof, and Meger 2018), have been used with good performance. All of these existing methods that use function approximation employ least squares to estimate the value function, and can be combined with the proposed method.

When iteration with the environment has high cost, such as in robotics, sample efficiency becomes important in RL. To increase sample efficiency, it is efficient to set $UTD \gg 1$. However, this leads to accumulation of overestimation bias, making learning difficult. To address this issue, REDQ combines ensemble and in-target minimization across a random subset of Q functions to adjust variance and bias suppression, enabling learning even with $UTD \gg 1$. By combining REDQ with SAC, it achieved a sample efficiency equal to or greater than that of state-of-the-art model-based algorithms for MuJoCo benchmark. In DroQ (Hiraoka et al. 2022) and MEPG (He et al. 2022), the ensemble-related computational efficiency was improved using dropout (Srivastava et al. 2014) as a substitute for the ensemble. It can also be used for reducing computational cost of SymREDQ.




Figure 5: The top figure illustrates the density of pre-corrected Bellman error (blue) and negative values of noise used for correction (orange). It shows how closely the distribution of η approaches the distribution of $-\epsilon$. The bottom figure shows the density of the post-corrected error (green), which is the sum of pre-corrected error and noise. This demonstrates the extent to which the distribution approached a symmetric distribution, and the corrected distribution (green) is more symmetric than the pre-corrected distribution (blue).

Error Distribution

In maximum likelihood estimation, a certain distribution is assumed to be the error distribution, and the estimation can be inaccurate when the actual distribution differs from the assumption. The least squares method can be regarded as a maximum likelihood estimation, assuming a normal distribution as the error distribution. It has been experimentally demonstrated that applying the least squares method to errors with nonnormal distributions leads to inaccurate estimation results (Constable 1988; Xu, Chen, and Liang 2018). Therefore, when using the least squares method in machine learning, learning is sometimes performed by transforming the distribution of the data into normal distribution. Many transformations focus on skewness, and it is well-known that logarithmic transformation is used for right-skewed distributions. The Box-Cox transformation (Box and Cox 1964) and Yeo-Johnson (Yeo and Johnson 2000) transformation can adjust the degree of skewness by parameters, and flexible transformations can be performed by estimating the parameters with maximum likelihood estimation, such that the transformed distribution becomes closer to normal distribution. Applying these methods to reinforcement learning is a task for future research.

The least squares method is also used in RL, where normal distribution is required as an error distribution. However, based on extreme value theory, (Garg et al. 2023) showed that the error distribution follows the Gumbel distribution, which is an extreme value distribution because the Bellman equation contains a maximum operator. They addressed this problem using the Gumbel regression, which is a maximum likelihood estimation assuming that the error distribution follows the Gumbel distribution. Gumbel regression can directly estimate optimal soft-value functions without sampling from policies and avoids evaluating out-of-distribution actions, thus significantly improving performance in offline RL settings. However, in online RL, not much improvement is observed. First, Gumbel regression, which assumes a Gumbel distribution as an error distribution, is unstable. It assumes an asymmetric error distribution and has an exponential term in the loss function, making the

gradient prone to being large. Second, the actual error distribution may differ from the Gumbel distribution because the state-action pairs taking the maximum are not independent, violating the premises of extreme value theory. Furthermore, because the properties of algorithms such as double Q-learning and target networks can affect the actual error distribution, it is difficult to assume a single static error distribution for maximum likelihood estimation. Therefore, the error distribution can vary depending on the task and the algorithm, highlighting the need for a method that can be applied to arbitrary error distributions.

The Gauss-Markov theorem states that under the assumptions of homoscedasticity and uncorrelated errors, the linear unbiased estimator estimated by the least squares method has the smallest variance and is the best linear unbiased estimator (BLUE). In RL, a small variance is not necessarily good, and the use of linear parameters is infrequent, but there is room to verify the impact of meeting these assumptions.

Conclusion

This study addresses the issue of the error distribution being far from the assumed normal distribution in online RL methods that use the least squares method. The proposed method adds noise that cancels out the distortion in the error, making it closer to a normal distribution. Experiments conducted on the MuJoCo benchmark demonstrate that the proposed method can symmetrically correct the error distribution, leading to improved sample efficiency and performance equivalent to or better than that of the state-of-the-art RL method. This is because stable learning with an error distribution close to the normal distribution can be achieved in the early stages of learning, resulting in improved efficiency.

In RL using the estimated values as the target, the error distribution is influenced by various factors, such as the Bellman operator and bias-suppression algorithms, which can pose problems during learning. However, there is limited research focusing on error distribution in RL, indicating substantial potential for future advancements in this field.

Acknowledgments

We would like to thank Jen-Yen Chang and Thomas Westfechtel for their valuable proofreading of this manuscript. This work was partially supported by JST Moonshot R&D Grant Number JPMJPS2011, CREST Grant Number JPMJCR2015 and Basic Research Grant (Super AI) of Institute for AI and Beyond of the University of Tokyo. T.O. was partially supported by JSPS KAKENHI Grant Number JP23K18476.

References

- Bishop, C. M. 2006. *Pattern Recognition and Machine Learning*. Springer. ISBN 978-0387-31073-2.
- Box, G. E.; and Cox, D. R. 1964. An analysis of transformations. *Journal of the Royal Statistical Society: Series B (Methodological)*, 26(2): 211–243.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. OpenAI Gym. arXiv:1606.01540.
- Chen, X.; Wang, C.; Zhou, Z.; and Ross, K. W. 2021. Randomized Ensembled Double Q-Learning: Learning Fast Without a Model. In *International Conference on Learning Representations*.
- Constable, C. G. 1988. Parameter estimation in non-Gaussian noise. *Geophysical Journal International*, 94(1): 131–142.
- Fisher, R. A.; and Tippett, L. H. C. 1928. Limiting forms of the frequency distribution of the largest or smallest member of a sample. *Mathematical Proceedings of the Cambridge Philosophical Society*, 24(2): 180–190.
- Fujimoto, S.; van Hoof, H.; and Meger, D. 2018. Addressing Function Approximation Error in Actor-Critic Methods. In Dy, J.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 1587–1596. PMLR.
- Garg, D.; Hejna, J.; Geist, M.; and Ermon, S. 2023. Extreme Q-Learning: MaxEnt RL without Entropy. In *International Conference on Learning Representations*.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. In Dy, J.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 1861–1870. PMLR.
- Hasselt, H. 2010. Double Q-learning. In Lafferty, J.; Williams, C.; Shawe-Taylor, J.; Zemel, R.; and Culotta, A., eds., *Advances in Neural Information Processing Systems*, volume 23. Curran Associates, Inc.
- He, Q.; Su, H.; Gong, C.; and Hou, X. 2022. MEPG: A Minimalist Ensemble Policy Gradient Framework for Deep Reinforcement Learning. arXiv:2109.10552.
- Hiraoka, T.; Imagawa, T.; Hashimoto, T.; Onishi, T.; and Tsuruoka, Y. 2022. Dropout Q-Functions for Doubly Efficient Reinforcement Learning. In *International Conference on Learning Representations*.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; and Riedmiller, M. 2013. Playing Atari with Deep Reinforcement Learning. arXiv:1312.5602.
- 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. 2015. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529–533.
- Mood, A. M. 1950. *Introduction to the theory of statistics*. McGraw-Hill.
- Rummery, G.; and Niranjan, M. 1994. On-Line Q-Learning Using Connectionist Systems. *Technical Report CUED/F-INFENG/TR 166*.
- Schulman, J.; Levine, S.; Abbeel, P.; Jordan, M.; and Moritz, P. 2015. Trust Region Policy Optimization. In Bach, F.; and Blei, D., eds., *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, 1889–1897. Lille, France: PMLR.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. arXiv:1707.06347.
- Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 15(56): 1929–1958.
- Todorov, E.; Erez, T.; and Tassa, Y. 2012. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 5026–5033.
- van Hasselt, H.; Guez, A.; and Silver, D. 2016. Deep Reinforcement Learning with Double Q-Learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 30(1).
- Watkins, C. J.; and Dayan, P. 1992. Q-learning. *Machine learning*, 8: 279–292.
- Xu, W.; Chen, W.; and Liang, Y. 2018. Feasibility study on the least square method for fitting non-Gaussian noise data. *Physica A: Statistical Mechanics and its Applications*, 492: 1917–1930.
- Yeo, I.-K.; and Johnson, R. A. 2000. A new family of power transformations to improve normality or symmetry. *Biometrika*, 87(4): 954–959.