

GMAC: A Distributional Perspective on Actor-Critic Framework

Daniel Wontae Nam¹ Younghoon Kim¹ Chan Y. Park¹

Abstract

In this paper, we devise a distributional framework on actor-critic as a solution to distributional instability, action type restriction, and conflation between samples and statistics. We propose a new method that minimizes the Cramér distance with the multi-step Bellman target distribution generated from a novel Sample-Replacement algorithm denoted $SR(\lambda)$, which learns the correct value distribution under multiple Bellman operations. Parameterizing a value distribution with Gaussian Mixture Model further improves the efficiency and the performance of the method, which we name GMAC. We empirically show that GMAC captures the correct representation of value distributions and improves the performance of a conventional actor-critic method with low computational cost, in both discrete and continuous action spaces using Arcade Learning Environment (ALE) and PyBullet environment.

1. Introduction

The ability to learn complex representations via neural networks has enjoyed success in various applications of reinforcement learning (RL), such as pixel-based video games (Mnih et al., 2015), the game of Go (Silver et al., 2016), robotics (Levine et al., 2016), and high dimensional controls like humanoid robots (Lillicrap et al., 2016; Schulman et al., 2015). Starting from the seminal work of Deep Q-Network (DQN) (Mnih et al., 2015), the advance in value prediction network, in particular, has been one of the main driving forces for the breakthrough.

Among the milestones of the advances in value function approximation, distributional reinforcement learning (DRL) further develops the scalar value function to a distributional representation. The distributional perspective offers various benefits by providing more information on the characteristics and the behavior of the value. One such benefit is the

¹KC Machine Learning Lab, Seoul, Korea. Correspondence to: Daniel <dwtam@kc-ml2.com>.

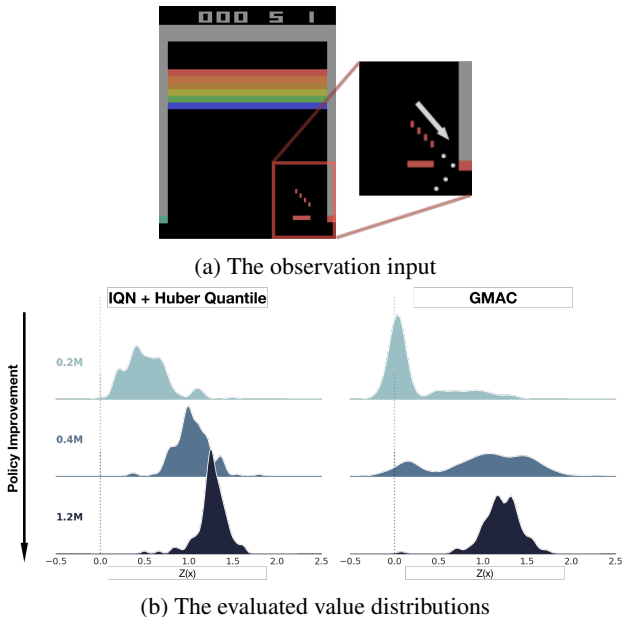


Figure 1: Modality of value distribution during the learning process of Breakout-v4. (a) An arrow is added in the inset to indicate the ball’s direction of travel. The episode reaches a terminal state if the paddle misses the ball. (b) Learned probability density functions of the value distributions of IQN + Huber-quantile (left) vs. GMAC (Gaussian mixture model + energy distance) (right) for a same policy when trained on $\{0.2, 0.4, 1.2\}$ M frames. As the policy improves, the probability of losing a turn ($V = 0$) should decrease while the probability of earning scores ($V > 0$) increases. Note that the modality transition from $V = 0$ is clearly captured by GMAC.

preservation of multimodality in value distributions, which leads to more stable learning of the value function (Bellemare et al., 2017a).

Despite the development, several issues remain, hindering DRL from becoming a robust framework. First, a theoretical instability exists in the control setting of value-based DRL methods (Bellemare et al., 2017a). Second, previous DRL algorithms are limited to a single type of action space, either discrete (Bellemare et al., 2017a; Dabney et al., 2018b;a) or continuous (Barth-Maron et al., 2018; Singh et al., 2020). Third, a common choice of loss function for DRL is the Huber quantile regression loss, which is vulnerable to conflation between samples and statistics without an imputation

strategy (Rowland et al., 2019).

While the instability and action space issue can be avoided simply by applying a general actor-critic framework (Williams, 1988; 1992; Sutton et al., 1999), practical methods critical to actor-critic framework such as TD(λ) have not been established in the distributional perspective. Therefore we suggest a novel sample-replacement algorithm denoted by SR(λ) to generate multi-step Bellman target distribution with high efficiency. Furthermore, we avoid the conflation problem by directly learning samples through minimizing the Cramér distance between distributions.

As proven in (Rowland et al., 2019), using an imputation strategy can help DRL methods to learn a more accurate representation of a distribution. However, many actor-critic methods are designed to use multi-step returns such as the λ -return (Watkins, 1989) for which the imputation strategy can be a computational burden. Therefore, we instead construct the multi-step returns from samples and parameters to avoid the necessity of imputation. We propose to parameterize the value distribution as a Gaussian mixture model (GMM), and minimize the Cramér distance between the distributions. When combining GMM with the energy distance, a specific case of the Cramér distance, we can derive an analytic solution and obtain unbiased sample gradients at a much lower computational cost compared to the method using the Huber quantile loss. We call our framework GMAC (Gaussian mixture actor-critic).

We present experimental results to demonstrate how GMAC can successfully solve the three problems of DRL. Firstly, we illustrate that GMAC is a competitive actor-critic framework by showing that the framework outperforms its baseline algorithms in the Atari games (Bellemare et al., 2013). Secondly, the experiments on the continuous control tasks in PyBullet environments (Coumans & Bai, 2016–2020) show that the same framework can be used for both tasks with discrete and continuous action spaces. Lastly, we share the FLOP measurement results to show that the accurate representation of value distributions can be learned with less computational cost.

2. Related Works

Bellemare et al. (2017a) has shown that the distributional Bellman operator derived from the distributional Bellman equation is a contraction in a maximal form of the Wasserstein distance. Based on this point, Bellemare et al. (2017a) proposed a categorical distributional model, C51, which is later discussed to be minimizing the Cramér distance in the projected distributional space (Rowland et al., 2018; Bellemare et al., 2019; Qu et al., 2019). Dabney et al. (2018b) proposed quantile regression-based models, QR-DQN, which parameterizes the distribution with a uniform mixture of

Diracs and uses sample-based Huber quantile loss (Huber, 1964). Dabney et al. (2018a) later expanded it further so that a full continuous quantile function can be learned through the implicit quantile network (IQN). Yang et al. (2019) then further improved the approximation of the distribution by adjusting the set of quantiles. Choi et al. (2019) suggested parameterizing the value distribution using Gaussian mixture and minimizing the Tsallis-Jensen divergence as the loss function on a value-based method. Outside of RL, Bellemare et al. (2017b) proposed to use Cramér distance in place of Wasserstein distance used in WGAN due to its unbiasedness in sample gradients (Arjovsky et al., 2017).

There have been many applications of the distributional perspective, which exploit the additional information from value distribution. Dearden et al. (1998) modeled parametric uncertainty and Morimura et al. (2010a;b) designed a risk-sensitive algorithm using a distributional perspective, which can be seen as the earliest concept of distributional RL. Mavrin et al. (2019) utilized the idea of the uncertainty captured from the variance of value distribution. Nikolov et al. (2019) has also utilized the distributional representation of the value function by using information-directed-sampling for better exploration of the value-based method. While multi-step Bellman target was considered (Hessel et al., 2018), the sample-efficiency was directly addressed by combining multi-step off-policy algorithms like Retrace(λ) (Gruslys et al., 2017).

Just as C51 has been expanded deep RL to distributional perspective, Barth-Maron et al. (2018) studied a distributional perspective on DDPG (Lillicrap et al., 2016), an actor-critic method, by parameterizing a distributional critic as categorical distribution and Gaussian mixture model. Singh et al. (2020) has further expanded the work by using an implicit quantile network for the critic. Several works (Duan et al., 2020; Kuznetsov et al., 2020; Ma et al., 2020) have proposed a distributional version of the soft-actor-critic (SAC) framework to address the error from over-estimating the value. These works mainly focused on combining a successful distributional method with a specific RL algorithm. To this end, this paper aims to suggest a more general method that can extend any actor-critic to the distributional perspective.

3. Distributional Reinforcement Learning

We consider a conventional RL setting, where an agent’s interaction with its environment is described by a Markov Decision Process (MDP) $(\mathcal{X}, \mathcal{A}, R, P, \gamma)$, where \mathcal{X} and \mathcal{A} are state and action spaces, $R(x, a)$ is the stochastic reward function for a pair of state x and action a , $P(x'|x, a)$ is the transition probability of observing x' given the pair (x, a) , and $\gamma \in (0, 1)$ is a time discount factor. A policy $\pi(\cdot|x)$ maps a state $x \in \mathcal{X}$ to a probability distribution over actions $a \in \mathcal{A}$.

The objective of RL is to maximize the expected return, $\mathbb{E}[G_t]$ where $G_t = \sum_{t=0}^{\infty} \gamma^t R(x_t, a_t)$ is the sum of discounted rewards from state x_t given a policy π at time t . Then for any state x_t , the value V and state-action value Q under the given policy π can be defined as

$$V(x_t) = \mathbb{E}[G_t \mid X = x_t], \quad (1)$$

$$Q(x_t, a_t) = \mathbb{E}[G_t \mid X = x_t, A = a_t]. \quad (2)$$

A recursive relationship in the value in terms of the reward and the transition probability is described by the Bellman equation (Bellman, 1957) given by

$$Q(x, a) = \mathbb{E}[R(x, a)] + \gamma \mathbb{E}_{a' \sim \pi, x' \sim P} [Q(x', a')], \quad (3)$$

where the first expectation is calculated over a given state-action pair (x, a) and the second expectation is taken over the next possible states $x' \sim P(\cdot | x, a)$ and actions $a' \sim \pi(\cdot | x)$.

DRL extends the Bellman equation to an analogous recursive equation, termed the distributional Bellman equation (Morimura et al., 2010a;b; Bellemare et al., 2017a), using a distribution of the possible sum of discounted rewards $Z(x, a)$:

$$Z^\pi(x, a) \stackrel{D}{=} R(x, a) + \gamma Z^\pi(X', A'), \quad (4)$$

where $\stackrel{D}{=}$ denotes having equal distributions and $Q(x, a) = \mathbb{E}[Z(x, a)]$. Then Z is learned through distributional Bellman operator \mathcal{T}^π defined as

$$\mathcal{T}^\pi Z(x, a) \stackrel{D}{=} R(x, a) + \gamma P^\pi Z(x, a) \quad (5)$$

where $P^\pi : \mathcal{Z} \rightarrow \mathcal{Z}$ is a state transition operator under policy π , $P^\pi Z(x, a) \stackrel{D}{=} Z(X', A')$, where $X' \sim P(\cdot | x, a)$ and $A' \sim \pi(\cdot | X')$. Analogously, the distributional Bellman optimality operator \mathcal{T} can be defined as

$$\mathcal{T}Z(x, a) \stackrel{D}{=} R(x, a) + \gamma Z(X', \arg \max_{a'} \mathbb{E}[Z(X', a')]). \quad (6)$$

The distributional Bellman operator has been proven to be a γ -contraction in a maximal form of Wasserstein distance (Bellemare et al., 2017a), which has a practical definition given by

$$d_p(U, V) = \left(\int_0^1 |F_U^{-1}(\omega) - F_V^{-1}(\omega)|^p d\omega \right)^{1/p}, \quad (7)$$

where U, V are random variables and F_U, F_V are their cumulative distribution functions (cdf).

However, unlike the distributional Bellman operator, the distributional Bellman optimality operator is not a contraction

in any metric (Bellemare et al., 2017a), causing an instability where the distance $d_p(\mathcal{T}Z_1, \mathcal{T}Z_2)$ between some random variables Z_1, Z_2 may not converge to a unique solution. This issue has been discussed in Bellemare et al. (2017a), with an example of oscillating value distribution caused by a specific tie-breaker design of the *argmax* operator.

The instability can be removed simply by learning the value distribution under the evaluation setting of the Bellman operation described in (5). This, on the other hand, poses a new problem on the RHS of (5): the next state-action value distribution $Z(X', A')$ becomes a mixture distribution of all possible state-action value distributions, the computation of which can be infeasible for value-based methods in continuous action space. We avoid this issue by directly approximating the state value distribution $Z(X')$ instead of the state-action value distribution $Z(X', A')$. This lets us to use the general actor-critic policy gradient,

$$\nabla J(\theta) = \mathbb{E}[A_t \nabla \ln(\pi_\theta(x_t, a_t))], \quad (8)$$

where the advantage A_t may be estimated using the temporal-difference (TD) error between expectations of the value distributions, or using generalized advantage estimation (Schulman et al., 2016) in a similar manner. At this point, we are left with estimating multi-step distributional Bellman target while avoiding the data conflation problem (Rowland et al., 2019).

4. Algorithm

4.1. SR(λ): Sample-Replacement for λ -return Distribution

Here we introduce SR(λ), a novel method for estimating the multi-step distributional Bellman target, analogous to the TD(λ) in the case of scalar value functions. The actor-critic method is a temporal-difference (TD) learning method in which the value function, the critic, is learned through the TD error defined by the difference between the TD target given by n -step return, $G_t^{(n)} = \sum_{i=1}^n \gamma^{i-1} r_{t+i} + \gamma^n V(x_{t+n})$, and the current value estimate $V(x_t)$. A special case of TD method, called TD(λ) (Sutton, 1988), generates a weighted average of n -step returns for the TD target, also known as the λ -return,

$$G_t^{(\lambda)} = (1 - \lambda) \sum_{n=1}^{\infty} \lambda^{n-1} G_t^{(n)}, \quad \lambda \in [0, 1], \quad (9)$$

to mitigate the variance and bias trade-off between Monte Carlo and the TD(0) return to enhance data efficiency.

An important piece of SR(λ) is the use of a random variable $Z_t^{(\lambda)}$ that is the distributional analogue of the λ -return, which we propose in the following. First let us define a random variable \tilde{G} whose sample space is the set of all n -step




Figure 2: A visual representation of the SR(λ) algorithm for generating N distributional Bellman targets from given a trajectory of length N . Set of dirac samples ($m = 4$) are represented by the colored shapes. Shapes of same colors denote that they are transformed from the same $Z(x_t)$ and different shapes denote numbers of Bellman operation applied to a specific sample. Each row of samples in the rightmost figure represent the samples the λ -return Bellman target $Z_{t-1}^{(\lambda)}$ for each time step.

returns, $\{G_t^{(1)}, \dots, G_t^{(\infty)}\}$ with the probability distribution given by

$$\Pr[\tilde{G} = G_t^{(n)}] = (1 - \lambda)\lambda^{n-1}. \quad (10)$$

Then, (9) is same as the expectation of the random variable \tilde{G} . Similar to $G_t^{(n)}$, we define n -step approximation of the value distribution as

$$Z_t^{(n)} := \sum_{i=0}^{n-1} \gamma^i R(x_{t+i}, a_{t+i}) + \gamma^n Z(x_{t+n}), \quad (11)$$

where $\mathbb{E}[Z_t^{(n)}] = G_t^{(n)}$. Then we can imagine a random variable \tilde{Z} whose sample space is a set of all n -step approximations, $\{Z_t^{(1)}, \dots, Z_t^{(\infty)}\}$, which are random variables as well. However, unlike G whose expectation is a scalar value, i.e. the weighted mean of its supports, the expectation of \tilde{Z} is a random variable that has a distribution equal to the mixture of distributions of $Z_t^{(n)}$. To avoid the ambiguity of “an expectation of a random variable of random variables”, we define the distributional analogue of (10) in terms of cdfs:

$$\Pr[\tilde{F} = F_{Z_t^{(n)}}] = (1 - \lambda)\lambda^{n-1}. \quad (12)$$

$F_{Z_t^{(n)}}$ denotes the cdf of the n -step return $Z_t^{(n)}$, and \tilde{F} is a random variable over the set of $F_{Z_t^{(n)}} \in \{F_{Z_t^{(0)}}, \dots, F_{Z_t^{(\infty)}}\}$. Then using (12), we can successfully define the expectation of \tilde{F} as a linear combination of $F_{Z_t^{(n)}}$

$$\mathbb{E}[\tilde{F}] = (1 - \lambda) \sum_{n=1}^{\infty} \lambda^{n-1} F_{Z_t^{(n)}}. \quad (13)$$

Let us define a random variable $Z_t^{(\lambda)}$ that has $\mathbb{E}[\tilde{F}]$ as its cdf, i.e. the probability distribution of $Z_t^{(\lambda)}$ is a mixture distribution of the probability distributions of $Z_t^{(n)}$'s. Then

the expectation of $Z_t^{(\lambda)}$ and the expectation of $Z_t^{(n)}$ have an analogous relationship to (9) (see Appendix B), meaning that the expectation of $Z_t^{(\lambda)}$ is equal to the λ -return.

Note that, in practice, collecting infinite horizon trajectory is infeasible and thus the truncated sum is often used (Cichosz, 1995; van Seijen et al., 2011):

$$F_{Z_t^{(\lambda)}} = (1 - \lambda) \sum_{n=1}^N \lambda^{n-1} F_{Z_t^{(n)}} + \lambda^N F_{Z_t^{(N)}}. \quad (14)$$

Given a trajectory of length N , naively speaking, finding $Z_t^{(\lambda)}$ for each time step requires finding N different $Z_t^{(n)}$. As a result, we need to find total of $O(N^2)$ different distributions to find $Z_t^{(\lambda)}$ for all states in the given trajectory. But the number of distributions to find reduces to $O(N)$ when we create approximations of $Z_t^{(n)}$ beforehand and reuse them for calculating $Z_t^{(\lambda)}$ for each time step.

One choice among such approximations is to use a mixture of diracs from the sample values, as described in (Dabney et al., 2018b):

$$F_{Z_t^{(n)}} \approx F_{Z_\theta(x_t)} := \frac{1}{m} \sum_{i=1}^m \delta_{\theta_i(x_t)}, \quad (15)$$

where $\theta : \mathcal{X} \rightarrow \mathbb{R}^m$ is some parametric model. In this case, we can approximate the distribution of $Z_t^{(\lambda)}$ by aggregating the samples from each $Z_\theta(x_t)$ with probability λ^{n-1} . Since the total set of samples remains unchanged during the calculation, we can create the N different $Z_t^{(\lambda)}$ in a single sweep by replacing a portion of samples for each time step, which leads to the name of our method Sample-Replacement, or SR(λ).

Figure 2 describes SR(λ) schematically. The approximated distribution of the λ -returns, $Z_t^{(\lambda)}$, for the last state in a trajectory is simply given by the samples of the last value distribution $Z_\theta(x_t)$. Then traversing the trajectory in a reversed time order, we replace each of the samples with a

Algorithm 1 SR(λ)

Input: Trajectory of states and value distributions $\{(x_1, Z_1), \dots, (x_N, Z_N)\}$ of length N , discount factor γ , weight parameter λ

Output: Set of λ -returns $\{Z_0^{(\lambda)}, \dots, Z_{N-1}^{(\lambda)}\}$

$\mathbf{Z} \leftarrow$ Collect m samples $\{z_1, \dots, z_m\}$ from Z_N

for $t = N - 1$ **to** 0 **do**

$\mathbf{Z} \leftarrow r_t + \gamma \mathbf{Z}$ {Bellman operation}

$Z_t^{(\lambda)} \leftarrow \sum_{i=1}^m \delta_{z_i}$
{empirical distribution using m Diracs}

$\mathbf{Z}' \leftarrow$ Collect m samples $\{z'_1, \dots, z'_m\}$ from $Z_t^{(\lambda)}$

for $i = 1$ **to** m **do**

$z_i \leftarrow z'_i$ with probability $1 - \lambda$

end for

end for

new sample from the earlier time step with a probability of $1 - \lambda$. The replaced collection of samples is used as the approximation for $Z_t^{(\lambda)}$ in that time step. We iterate this process until the beginning of the trajectory to create a total of N approximations. A more detailed description of the algorithm can be found in Algorithm 1.

We further propose to apply SR(λ) to the parameters of GMM instead of dirac samples in the following sections. There exists a closed-form solution for minimizing the Cramér distance between Gaussian mixtures, which enables us to create unbiased gradients at a lower computational cost compared to when using samples or statistics. In the case of statistics, Rowland et al. (2019) has shown that one should use an imputation strategy on the statistics to acquire samples of the distribution, which may add significant computational overhead.

4.2. Cramér Distance

Let P and Q be probability distributions over \mathbb{R} . If we define the cdf of P, Q as F_P, F_Q respectively, the l_p family of divergence between P and Q is

$$l_p(P, Q) := \left(\int_{-\infty}^{\infty} |F_P(x) - F_Q(x)|^p dx \right)^{1/p}. \quad (16)$$

When $p = 2$, it is termed the Cramér distance. The distributional Bellman operator in the evaluation setting is a $|\gamma|^{1/p}$ -contraction mapping in the Cramér metric space (Rowland et al., 2019; Qu et al., 2019), whose worked out proof can also be found in Appendix C.

A notable characteristic of the Cramér distance is the unbiasedness of the sample gradient,

$$\mathbb{E}_{X \sim Q} \nabla_{\theta} l_2^2(\hat{P}_m, Q_{\theta}) = \nabla_{\theta} l_2^2(P, Q_{\theta}) \quad (17)$$

where $\hat{P}_m := \frac{1}{m} \sum_{i=1}^m \delta_{X_i}$ is the empirical distribution, and Q_{θ} is a parametric approximation of a distribution. The unbiased sample gradient makes it suitable to use Cramér distance with stochastic gradient descent method and empirical distributions for updating the value distribution.

Székely (2002) showed that, in the univariate case, the squared Cramér distance is equivalent to one half of *energy distance* ($l_2^2(P, Q) = \frac{1}{2} \mathcal{E}(P, Q)$) defined as

$$\begin{aligned} \mathcal{E}(P, Q) &:= \mathcal{E}(U, V) \\ &= 2 \mathbb{E} \|U - V\|_2 - \mathbb{E} \|U - U'\|_2 - \mathbb{E} \|V - V'\|_2, \end{aligned} \quad (18)$$

where U, U' and V, V' are random variables that follow P, Q , respectively. Then, energy distance can be approximated using the random samples of U and V .

4.3. Energy Distance between Gaussian Mixture Models

We take a step further to enhance the approximation accuracy and computational efficiency by considering the parameterized model of the value distribution as a GMM (Choi et al., 2019; Barth-Maron et al., 2018). Following the same assumption used for (15), the approximation using GMM is given using parametric models $\mu, \sigma, w : \mathcal{X} \rightarrow \mathbb{R}^K$

$$\begin{aligned} Z_{\theta}(x_t) &\sim \mathcal{N}(\mu_k(x_t), \sigma_k^2(x_t)), \\ k &\sim \text{Categorical}(w(x_t)). \end{aligned} \quad (19)$$

If random variables U, V follow the distributions P, Q parameterized as GMMs, the energy distance has the following closed-form

$$\begin{aligned} \mathcal{E}(U, V) &= 2\delta(U, V) - \delta(U, U') - \delta(V, V'), \\ \text{where} \\ \delta(U, V) &= \sum_{i,j} w_{ui} w_{vj} \mathbb{E} [|X_{ij}|], \\ X_{ij} &\sim \mathcal{N}(\mu_{ui} - \mu_{vj}, \sigma_{ui}^2 + \sigma_{vj}^2). \end{aligned} \quad (20)$$

Here, μ_{ui} refers to the mean of the i^{th} component for random variable U and same applies for σ and w for both U and V . The closed-form solution of the energy distance defined in (20) has a computational advantage over sample-based approximations like the Huber quantile loss. When using the GMM, the analytic approximation of (15) can be derived as

$$\begin{aligned} Z_t^{(\lambda)} &\sim \mathcal{N}(\mu_{nk}(x_t), \sigma_{nk}^2(x_t)), \\ n &\sim \text{Geo}(1 - \lambda), \\ k &\sim \text{Categorical}(w(x_t)), \end{aligned} \quad (21)$$

$$F_{Z_t^{(\lambda)}} \approx \frac{1}{m} \sum_{i=1}^m \Phi(z | \mu_{nk}, \sigma_{nk}^2),$$

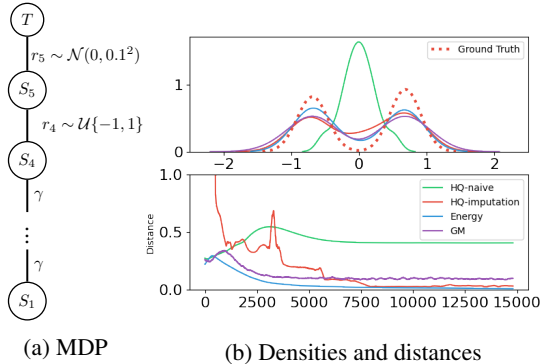


Figure 3: (a) An environment with five states and stochastic rewards with expected value of zero. (b) The probability density functions (top) and the energy distance between the ground truth and the estimated distributions (bottom) for tabular methods with Huber-quantile, Huber-quantile with imputation, energy distance, and GMM with energy distance.

where μ_{nk} refers to the k^{th} component of μ_n , and $\Phi(z|\mu_{nk}, \sigma_{nk}^2)$ denotes the cdf of a normal distribution given its mean and standard deviation, μ_{nk} and σ_{nk} , respectively. This can be empirically approximate as a mixture of m Gaussians and we can simply perform sample replacement on the parameters (μ, σ^2) , instead of realizations of the random variables as in (15). Then, the distance function described in (20) can easily be applied.

When bringing all the components together, we have a distributional actor-critic framework with $SR(\lambda)$ that minimizes the energy distance between Gaussian mixture value distributions. Comprehensively, we call this method GMAC. A brief sketch of the algorithm is shown in Appendix E.

5. Experiments

In this section, we present experimental results for three different distributional versions of Proximal Policy Optimization (PPO) with $SR(\lambda)$: IQAC (IQN + Huber quantile), IQAC-E (IQN + energy distance), and GMAC (GMM + energy distance), in the order of the progression of our suggested approach. The performance of the scalar version of PPO with value clipping (Schulman et al., 2016) is used as the baseline for comparison. Details about the loss function of each method can be found in Appendix D. For a fair comparison, we keep all common hyperparameters consistent across the algorithms except for the value heads and their respective hyperparameters (see Appendix F).

The results demonstrate three contributions of our proposed DRL framework: 1) the ability to correctly capture the multimodality of value distributions, 2) generalization to both discrete and continuous action spaces, and 3) significantly reduced computational cost.

Representing Multimodality As discussed throughout Section 4, we expect minimizing the Cramér distance to produce a correct depiction of a distribution without using an imputation strategy. First, we demonstrate this with a simple value regression problem for an MDP of five sequential states, as shown in Figure 3 (a). The reward function r_i of last two state S_i is stochastic, with r_4 from a uniform discrete distribution and r_5 from a normal distribution. Then the value distribution of S_1 should be bimodal with expectation of zero (Figure 3 (b)). In this example, minimizing the Huber-quantile loss (labeled as HQ-naive) of dirac mixture underestimates the variance of S_1 due to conflation and does not capture the locations of the modes. By applying an imputation strategy as suggested in Rowland et al. (2019), a slight improvement on the underestimation of variance can be seen. On the other hand, both dirac mixture and GMM, labeled as Energy and GM respectively in the figure, show that minimizing the energy distance converges to correct mode locations. For a fair comparison, GMM uses one-third of the parameters used in the dirac mixtures as its number of mixtures. More details about the experimental setup and further results can be found in Appendix G. The comparison is extended to complex tasks such as the Atari games, of which an example result is shown in Figure 1, and additional visualizations of the value distribution during the learning process from different games can be found in Appendix G.

Discrete and Continuous Action Spaces The human normalized score for 57 Atari games in ALE (Bellemare et al., 2013) is presented in Figure 4. The results show that GMAC outperforms its scalar baseline PPO and other known distributional methods IQN and QR-DQN in mean scores. On the other hand, in the median scores, GMAC places between IQN and QR-DQN. The results tell us that GMAC significantly outperforms the value-based distributional methods in some of the Atari games while its overall performance is competitive. Another clear distinction is that there is a significant decrease in performance when implicit quantile network is used with Huber-quantile loss for the critic with same architecture with same hyperparameters. In contrast, using energy distance as the loss function ensures non-degenerative performance. The learning curves for each of 61 Atari games, including ones that did not have human scores, can be found in Appendix G.

The same exact algorithm is taken to continuous control task of PyBullet environments (Coumans & Bai, 2016–2020), with the changes only made in the hyperparameters and policy parameters, from softmax logits to mean and variance of normal distribution. Without any continuous-control specific modifications made, our methods produce competitive performance compared to the scalar version PPO with slight improvements in the hard tasks such as HumanoidBulletEnv-v0. More results can be found in Appendix G.

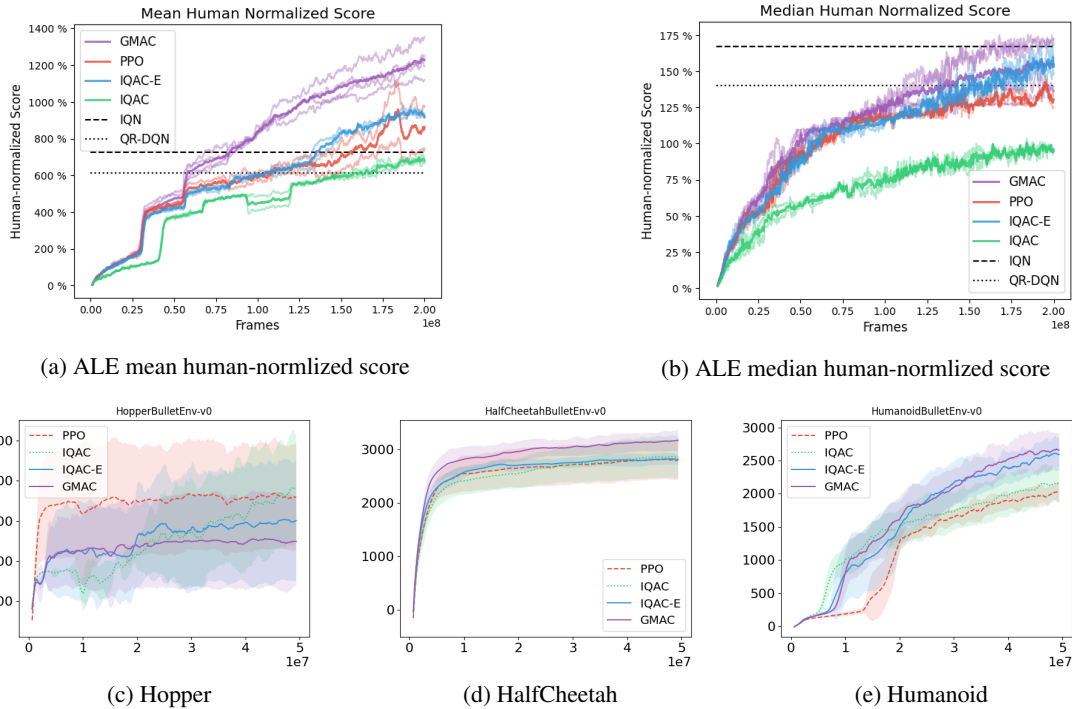


Figure 4: Learning curves for Atari games from ALE and 3 continuous control tasks from PyBullet during training. The scores for IQN and QR-DQN are taken from Dabney et al. (2018a). For Atari, GMAC was run for 3 seeds, PPO and other baselines were run for 2 seeds. Results of selected Atari games run for 5 seeds can be found in Appendix G. For PyBullet, all algorithms were run for 5 seeds.

Computational Cost Table 1 shows the number of parameters and the number of floating-point operations (FLOPs) required for a single inference and update step of each agent. We emphasize three points here. Firstly, the implicit quantile network requires more parameters due to the intermediate embeddings of random quantiles. Secondly, the difference between the FLOPs for a single update in IQAC and IQAC-E indicates that the proposed energy distance requires less computation than the Huber quantile regression. Lastly, the results for GMAC show that using GMM can greatly reduce the cost even to match the numbers of PPO while having improved performance.

Table 1: FLOP measurement results for a single process in Breakout-v4

| Algorithm | Params (M) | FLOPs (G) | |
|-----------|-------------|-------------|-------------|
| | | Inference | Update |
| PPO | 0.44 | 1.73 | 5.19 |
| IQAC | 0.52 | 2.98 | 12.98 |
| IQAC-E | 0.52 | 2.98 | 8.98 |
| GMAC | 0.44 | 1.73 | 5.27 |

Using Distributions By capturing the correct modes of a value distribution, an additional degree of freedom on top of




Figure 5: Learning curve of Montezuma’s Revenge using modality information as intrinsic reward. For a fair comparison, both TD error and energy distance used for intrinsic rewards are normalized in a similar manner to that of RND (Burda et al., 2019) to keep the scales at a similar level.

the expected value can be accurately obtained, from which richer information can be derived to distinguish states by their value distributions. In particular, the extra information may be utilized as an intrinsic motivation in sparse-reward exploration tasks. To demonstrate the plausibility of such application, we compare using Cramér distance between value distributions as intrinsic reward to using TD error between scalar value estimates in a sparse reward environment of Montezuma’s Revenge in Figure 5, which shows a clear improvement in performance.

6. Conclusion

In this paper, we have developed the distributional perspective of the actor-critic framework which integrates the $SR(\lambda)$ method, Cramér distance, and Gaussian mixture models for improved performance in both discrete and continuous action spaces at a lower computational cost. Furthermore, we show that our proposed method can capture the correct modality in the value distribution, while the extension of the conventional method with the stochastic policy fails to do so.

Capturing the correct modality of value distributions can improve the performance of various policy-based RL applications that exploit statistics from the value distribution. Such applications may include training risk-sensitive policies and learning control tasks with sparse rewards that require heavy exploration, where transient information from the value distribution can give benefit to the learning process. We leave further development of these ideas as future works.

References

- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein gan. In Proceedings of the 34th International Conference on Machine Learning (ICML), 2017.
- Barth-Maron, G., Hoffman, M. W., Budden, D., Dabney, W., Horgan, D., TB, D., Muldal, A., Heess, N., and Lillicrap, T. Distributed distributional deterministic policy gradients. In International Conference on Learning Representations (ICLR), 2018.
- Bellemare, M. G., Naddaf, Y., Veness, J., and Bowling, M. The arcade learning environment: An evaluation platform for general agents. Journal of Artificial Intelligence Research, 47:253–279, 2013.
- Bellemare, M. G., Dabney, W., and Munos, R. A distributional perspective on reinforcement learning. In Proceedings of the 34th International Conference on Machine Learning (ICML), 2017a.
- Bellemare, M. G., Danihelka, I., Dabney, W., Mohamed, S., Lakshminarayanan, B., Hoyer, S., and Munos, R. The cramer distance as a solution to biased wasserstein gradients. arXiv preprint arXiv:1705.10743, 2017b.
- Bellemare, M. G., Roux, N. L., Castro, P. S., and Moitra, S. Distributional reinforcement learning with linear function approximation. In Artificial Intelligence and Statistics, volume 89 of Proceedings of Machine Learning Research, pp. 2203–2211, 2019.
- Bellman, R. Dynamic Programming. Dover Publications, 1957.
- Bertsekas, D. P. and Tsitsiklis, J. N. Neuro-Dynamic Programming. Athena Scientific, 1st edition, 1996.
- Burda, Y., Edwards, H., Storkey, A. J., and Klimov, O. Exploration by random network distillation. In 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019. OpenReview.net, 2019.
- Choi, Y., Lee, K., and Oh, S. Distributional deep reinforcement learning with a mixture of gaussians. In 2019 International Conference on Robotics and Automation (ICRA), pp. 9791–9797, 2019.
- Cichosz, P. Truncating temporal differences: On the efficient implementation of $TD(\lambda)$ for reinforcement learning. Journal on Artificial Intelligence, 2:287–318, 1995.
- Coumans, E. and Bai, Y. Pybullet, a python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>, 2016–2020.
- Dabney, W., Ostrovski, G., Silver, D., and Munos, R. Implicit quantile networks for distributional reinforcement learning. In Proceedings of the 35th International Conference on Machine Learning (ICML), 2018a.
- Dabney, W., Rowland, M., Bellemare, M. G., and Munos, R. Distributional reinforcement learning with quantile regression. In AAAI, pp. 2892–2901, 2018b.
- Dearden, R., Friedman, N., and Russell, S. Bayesian q-learning. Proceedings of the fifteenth national/tenth conference on Artificial intelligence/Innovative applications of artificial intelligence, pp. 761 – 768, 1998.
- Duan, J., Guan, Y., Li, S. E., Ren, Y., and Cheng, B. Distributional soft actor-critic: Off-policy reinforcement learning for addressing value estimation errors. arXiv preprint arXiv:2001.02811, 2020.
- Gruslys, A., Dabney, W., Azar, M. G., Piot, B., Bellemare, M., and Munos, R. The reactor: A fast and sample-efficient actor-critic agent for reinforcement learning. In International Conference on Learning Representation (ICLR), 2017.
- Hessel, M., Modayil, J., Hasselt, H. V., Schaul, T., Ostrovski, G., Dabney, W., Horgan, D., Piot, B., Azar, M. G., and Silver, D. Rainbow: Combining improvements in deep reinforcement learning. In AAAI, 2018.
- Huber, P. J. Robust estimation of a location parameter. The Annals of Mathematical Statistics, 35(1):73–101, 1964.
- Kakade, S. and Langford, J. Approximately optimal approximate reinforcement learning. In Proceedings of the 19th

- International Conference on Machine Learning (ICML), 2002.
- Kuznetsov, A., Shvechikov, P., Grishin, A., and Vetrov, D. P. Controlling overestimation bias with truncated mixture of continuous distributional quantile critics. arXiv preprint arXiv:2005.04269, 2020.
- Levine, S., Finn, C., Darrell, T., and Abbeel, P. End-to-end training of deep visuomotor policies. Journal of Machine Learning Research, 17:39:1–39:40, 2016.
- 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. In Proceedings of the 33rd International Conference on Learning Representations (ICML), 2016.
- Ma, X., Zhang, Q., Xia, L., Zhou, Z., Yang, J., and Zhao, Q. Distributional soft actor critic for risk sensitive learning. arXiv preprint arXiv:2004.14547, 2020.
- Mavrin, B., Yao, H., Kong, L., Wu, K., and Yu, Y. Distributional reinforcement learning for efficient exploration. In Proceedings of the 36th International Conference on Machine Learning (ICML), 2019.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., and Hassabis, D. Human-level control through deep reinforcement learning. Nature, 518(7540): 529–533, 2015.
- Morimura, T., Sugiyama, M., Kashima, H., Hachiya, H., and Tanaka, T. Parametric return density estimation for reinforcement learning. In UAI, pp. 368–375. AUAI Press, 2010a.
- Morimura, T., Sugiyama, M., Kashima, H., Hachiya, H., and Tanaka, T. Nonparametric return distribution approximation for reinforcement learning. In Proceedings of the 27th International Conference on Machine Learning (ICML), 2010b.
- Nikolov, N., Kirschner, J., Berkenkamp, F., and Krause, A. Information-directed exploration for deep reinforcement learning. In International Conference on Learning Representations (ICLR), 2019.
- Puterman, M. L. Markov Decision Processes: Discrete Stochastic Dynamic Programming. John Wiley & Sons, Inc., 1st edition, 1994.
- Qu, C., Mannor, S., and Xu, H. Nonlinear distributional gradient temporal-difference learning. In Proceedings of the 36th International Conference on Machine Learning, volume 97 of Proceedings of Machine Learning Research, pp. 5251–5260. PMLR, 09–15 Jun 2019.
- Rowland, M., Bellemare, M. G., Dabney, W., Munos, R., and Teh, Y. W. An analysis of categorical distributional reinforcement learning. In Storkey, A. J. and Pérez-Cruz, F. (eds.), Artificial Intelligence and Statistics (AISTATS), volume 84 of Proceedings of Machine Learning Research, pp. 29–37. PMLR, 2018.
- Rowland, M., Dadashi, R., Kumar, S., Munos, R., Bellemare, M. G., and Dabney, W. Statistics and samples in distributional reinforcement learning. In Chaudhuri, K. and Salakhutdinov, R. (eds.), Proceedings of the 36th International Conference on Machine Learning (ICML), 2019.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. Trust region policy optimization. In Proceedings of the 32nd International Conference on Machine Learning (ICML), 2015.
- Schulman, J., Moritz, P., Levine, S., Jordan, M. I., and Abbeel, P. High-dimensional continuous control using generalized advantage estimation. In International Conference on Learning Representations (ICLR), 2016.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347, 2017.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. Mastering the game of go with deep neural networks and tree search. Nature, 529(7587):484–489, 2016.
- Singh, R., Lee, K., and Chen, Y. Sample-based distributional policy gradient. arXiv preprint arXiv:2001.02652, 2020.
- Sutton, R. S. Learning to predict by the methods of temporal differences. Machine Learning, 3(1):9–44, aug 1988. ISSN 0885-6125. doi: 10.1023/A:1022633531479.
- Sutton, R. S. and Barto, A. G. Introduction to Reinforcement Learning. MIT Press, Cambridge, MA, USA, 1st edition, 1998. ISBN 0262193981.
- Sutton, R. S., McAllester, D., Singh, S., and Mansour, Y. Policy gradient methods for reinforcement learning with function approximation. In Proceedings of the 12th International Conference on Neural Information Processing Systems (NeurIPS), Cambridge, MA, USA, 1999.
- Székely, G. E-statistics: The energy of statistical samples. 10 2002. doi: 10.13140/RG.2.1.5063.9761.

- van Seijen, H., Whiteson, S., van Hasselt, H., and Wiering, M. Exploiting best-match equations for efficient reinforcement learning. Journal of Machine Learning Research, 12:2045–2094, 2011.
- Watkins, C. J. C. H. Learning from Delayed Rewards. PhD thesis, King's College, Oxford, 1989.
- Williams, R. J. Toward a theory of reinforcement-learning connectionist systems. Technical Report NU-CCS-88-3, College of Comp. Sci., Northeastern University, Boston, MA, 1988.
- Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. Machine Learning, 8:229–256, 1992.
- Yang, D., Zhao, L., Lin, Z., Qin, T., Bian, J., and Liu, T.-Y. Fully parameterized quantile function for distributional reinforcement learning. In 33rd Annual Conference on Neural Information Processing Systems (NeurIPS), pp. 6190–6199, 2019.