

Bayesian Distributional Policy Gradients

Luchen Li,¹ A. Aldo Faisal^{1,2,3}

¹ Brain & Behaviour Lab, Dept. of Computing, Imperial College London

² Brain & Behaviour Lab, Dept. of Bioengineering, Imperial College London

³ UKRI Centre in AI for Healthcare, Imperial College London

{l.li17, aldo.faisal}@imperial.ac.uk

Abstract

Distributional Reinforcement Learning (RL) maintains the entire probability distribution of the reward-to-go, i.e. the return, providing more learning signals that account for the uncertainty associated with policy performance, which may be beneficial for trading off exploration and exploitation and policy learning in general. Previous works in distributional RL focused mainly on computing the state-action-return distributions, here we model the state-return distributions. This enables us to translate successful conventional RL algorithms that are based on state values into distributional RL. We formulate the distributional Bellman operation as an inference-based auto-encoding process that minimises Wasserstein metrics between target/model return distributions. The proposed algorithm, BDPG (Bayesian Distributional Policy Gradients), uses adversarial training in joint-contrastive learning to estimate a variational posterior from the returns. Moreover, we can now interpret the return prediction uncertainty as an information gain, which allows to obtain a new curiosity measure that helps BDPG steer exploration actively and efficiently. We demonstrate in a suite of Atari 2600 games and MuJoCo tasks, including well known hard-exploration challenges, how BDPG learns generally faster and with higher asymptotic performance than reference distributional RL algorithms.

Introduction

In reinforcement learning (RL), the performance of a policy is evaluated by the (discounted) accumulated future rewards, a random variable known as the reward-to-go or the return. Instead of maintaining the expectations of returns as scalar value functions, distributional RL estimates *return distributions*. Keeping track of the uncertainties around returns has initially been leveraged as a means to raise risk awareness in RL (Morimura et al. 2010; Lattimore and Hutter 2012). Recently, a line of research pioneered by (Bellemare, Dabney, and Munos 2017) applied the distributional Bellman operator for control purposes. Distributional RL is shown to outperform previous successful deep RL methods in Atari-57 when combined with other avant-garde developments in RL (Hessel et al. 2018; Dabney et al. 2018a).

The critical hurdle in distributional RL is to minimise a Wasserstein distance between the distributions of a return and its Bellman target, under which the Bellman operation is a contraction mapping (Bellemare, Dabney, and Munos 2017). A differentiable Wasserstein distance estimator can be obtained in its dual form with constrained Kantorovich potentials (Gulrajani et al. 2017; Arjovsky, Chintala, and Bottou 2017), or approximated by restricting the search for couplings to a set of smoothed joint probabilities with entropic regularisations (Cuturi 2013; Montavon, Müller, and Cuturi 2016; Genevay et al. 2016; Luise et al. 2018). Alternatively, a Bayesian inference perspective redirects the search space to a set of probabilistic encoders that map data in the input space to codes in a latent space (Bousquet et al. 2017; Tolstikhin et al. 2018; Ambrogioni et al. 2018). Bayesian approaches rely on inference to bypass rigid and sub-optimal distributions that are usually entailed otherwise, while retaining differentiability and tractability. Moreover, predictions based on inference, the expectation across a latent space, are more robust to unseen data (Blundell et al. 2015) and thus able to generalise better.

In contrast to previous distributional RL work that focuses on state-action-return distributions, here we investigate state-return distributions and prove that its Bellman operator is also a contraction in Wasserstein metrics. This opens up the possibility of converting state-value algorithms into distributional RL settings. We then formulate the distributional Bellman operation as an inference-based auto-encoding process that minimises Wasserstein metrics between continuous distributions of the Bellman target and estimated return. A second benefit of our inference model is that the learned posterior enables a curiosity bonus in the form of information gain (IG), which is leveraged as internal reward to boost exploration efficiency. We explicitly calculate the entropy reduction in a latent space corresponding to return probability estimation as a KL divergence. In contrast to previous work (Bellemare et al. 2016; Sekar et al. 2020; Ball et al. 2020) in which IG was approximated with ensemble entropy or prediction gains, we obtain analytical results from our variational inference scheme.

To test our fully Bayesian approach and curiosity-driven exploration mechanism against a distributional RL back-

drop, we embed these two innovations into a policy gradient framework. Both innovations would also work for value-based and off-policy policy gradients methods where the state-action-return distribution is modelled instead.

We evaluate and compare our method to other distributional RL approaches on the Arcade Learning Environment Atari 2600 games (Bellemare et al. 2013), including some of the best known hard-exploration cases, and on MuJoCo continuous-control tasks (Todorov, Erez, and Tassa 2012). To conclude we perform ablation experiments, where we investigate our exploration mechanism and the length of bootstrapping in distributional Bellman backup.

Our key contributions in this work are two-fold: we derive first, a fully inference-based generative approach to distributional Bellman operations; and second, a novel curiosity-driven exploration mechanism formulated as posterior information gains attributed to return prediction uncertainty.

Preliminaries

Wasserstein Variational Inference

In this subsection, we discuss how Wasserstein metrics in a data space can be estimated in a Bayesian fashion using adversarial training. Notation-wise we use calligraphic letters for spaces, capital letters for random variables and lower-case letters for values. We denote probability distributions and densities with the same notation, discriminated by the argument being capital or lower-case, respectively.

In optimal transport problems (Villani 2008), divergences between two probability distributions are estimated as the cost required to transport probability mass from one to the other. Consider input spaces $\mathcal{X} \in \mathbb{R}^n$, $\mathcal{Y} \in \mathbb{R}^m$ and a pairwise cost function $c : \mathcal{X} \times \mathcal{Y} \mapsto \mathbb{R}^+$. For two probability measures $\alpha : \mathcal{X} \mapsto \mathcal{P}$, $\beta : \mathcal{Y} \mapsto \mathcal{P}$, an optimal transport divergence is defined as

$$\mathcal{L}_c(\alpha, \beta) := \inf_{\gamma \in \Gamma(\alpha, \beta)} \int_{\mathcal{X} \times \mathcal{Y}} c(x, y) d\gamma(x, y), \quad (1)$$

where $\Gamma(\alpha, \beta)$ is a set of joint distributions or couplings on $\mathcal{X} \times \mathcal{Y}$ with marginals α and β respectively. Particularly, when $\mathcal{Y} = \mathcal{X}$ and the cost function is derived from a metric over \mathcal{X} , $d : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}^+$, via $c(x, y) = d^p(x, y)$, $p \geq 1$, the p -Wasserstein distance on \mathcal{X} is given as

$$W_p(\alpha, \beta) := \mathcal{L}_{d^p}(\alpha, \beta)^{1/p}. \quad (2)$$

Now consider a generative process through a latent variable $Z \in \mathcal{Z} \in \mathbb{R}^l$ with a prior $p_Z(Z)$, a decoder $p_\theta(X|Z)$ and an encoder (amortised inference estimator) $q_\phi(Z|X)$, in which the parameters ϕ , θ are trained to mimic the data distribution $p_X(X)$ implicitly represented by the i.i.d. training samples. The density corresponding to the model distribution can be expressed as $p_G(x) = \mathbb{E}_{z \sim p_Z}[p_\theta(x|z)]$. For a deterministic¹ decoder $X = G_\theta(Z)$, p_G can be thought of

¹For the purpose of generative modelling, the intuition of minimising Wasserstein metrics between target/model distributions (instead of stronger probability density discrepancies such as f -divergences) is to still see meaningful gradients when the model manifold and the true distribution’s support have few intersections without introducing noise to the model distribution (by using a directed continuous mapping) that renders reconstructions blurry.

as the *push-forward* of p_Z through G_θ , i.e. $p_G = G_{\theta\#}p_Z$. Minimising the Wasserstein distance between $p_X(X)$ and $p_G(X)$ is thereby equivalent to finding an optimal transport plan between $p_X(X)$ and $p_Z(Z)$, and matching the aggregated posterior $Q(Z) := \mathbb{E}_{x \sim p_X}[q_\phi(Z|x)]$ to the prior $p_Z(Z)$ (Bousquet et al. 2017; Tolstikhin et al. 2018; Ambrogioni et al. 2018; Rosca, Lakshminarayanan, and Mohamed 2018; He et al. 2019)

$$W_p^p(p_X, p_G) = \inf_{q_\phi: Q=p_Z} \mathbb{E}_{X \sim p_X} \mathbb{E}_{Z \sim q_\phi} [d^p(X, G_\theta(Z))]. \quad (3)$$

Marginal matching in \mathcal{Z} is sometimes preferred for generative models, since it alleviates the posterior collapse problem (Zhao, Song, and Ermon 2017; Hoffman and Johnson 2016) by enabling Z to be more diversely distributed for different x ’s. However, when doing so, Eq. (3) is no longer a proper inference objective, as it enforces neither posterior-contrastive nor joint-contrastive learning. In fact, the encoder needs not to approximate the true posterior $p_\theta(Z|X)$ exactly to satisfy the marginal match. In contrast, our approach maintains a fully Bayesian inference pipeline.

While explicit variational inference requires all probability densities to have analytical expressions, we bypass this by direct density matching through adversarial training (Goodfellow et al. 2014), which requires only that densities can be sampled from for gradient backpropagation, thereby allowing for a degenerate decoder $p_\theta(X|Z) = \delta_{G_\theta(Z)}(X)$.

Lemma 1. (Donahue, Krähenbühl, and Darrell 2017) *Let $p(X, Z)$ and $q(X, Z)$ denote the joint sampling distribution induced by the decoder and encoder respectively, D_ψ a discriminator, and define $F(\psi, \phi, \theta) := \mathbb{E}_{x, z \sim p}[\log D_\psi(x, z)] + \mathbb{E}_{x, z \sim q}[\log(1 - D_\psi(x, z))]$. For any encoder and decoder, deterministic or stochastic, the optimal discriminator $D_{\psi^*} = \operatorname{argmax}_{D_\psi} F(\psi, \phi, \theta)$ is the Radon-Nikodym derivative of measure $p(X, Z)$ w.r.t. $p(X, Z) + q(X, Z)$. The encoder and decoder’s objective for an optimal discriminator $C(\phi, \theta) := F(\psi^*, \phi, \theta)$ can be written in the Jensen-Shannon divergence $C(\phi, \theta) = 2\text{JS}(p, q) - \log 4$, in which the global minimum is achieved if and only if $p(X, Z) = q(X, Z)$.*

We jointly minimise Wasserstein metrics between model/target distributions in \mathcal{X} and conduct variational inference adversarially. p_G will be shown to be modelling the distribution of random returns, leading to a novel approach to accomplishing distributional Bellman operations.

Distributional Reinforcement Learning

We start by laying out RL and policy gradients notation, then explain the distributional perspective of RL, as well as previous solutions to it.

Policy Gradients A standard RL task is framed within a Markov decision process (MDP) $(\mathcal{S}, \mathcal{A}, \mathcal{R}, P, \gamma)$ (Puterman 1994), where \mathcal{S} and \mathcal{A} denote the state and action spaces respectively, $\mathcal{R} : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}^n$ a potentially stochastic reward function, $P : \mathcal{S} \times \mathcal{A} \mapsto \mathcal{P}(\mathcal{S})$ a transition probability

density function, and $\gamma \in (0, 1)$ a temporal discount factor. An RL agent has a policy that maps states to a probability distribution over actions $\pi : \mathcal{S} \mapsto \mathcal{P}(\mathcal{A})$.

The return G^π under the policy π is a random variable that represents the sum of discounted future rewards and the *state-dependent* return is $G^\pi(s) := \sum_{t=0}^{\infty} \gamma^t r_t, s_0 = s$.

A state value function is defined as the expected return $V^\pi(s) := \mathbb{E}[G^\pi(s)]$, a state-action value function the expected state-action return $Q^\pi(s, a) := \mathbb{E}[G^\pi(s, a)]$. The Bellman operator \mathcal{T}^π (Bellman 1957) is defined as

$$\mathcal{T}^\pi V^\pi(s) := \mathbb{E}_{\pi, \mathcal{R}, P} [r + \gamma V^\pi(s')], \quad (4)$$

$$\mathcal{T}^\pi Q^\pi(s, a) := \mathbb{E}_{\mathcal{R}, P, \pi} [r + \gamma Q^\pi(s', a')]. \quad (5)$$

Policy gradient methods (Sutton et al. 1999) optimise a parameterised policy π by directly ascending the gradient of a policy performance objective such as $\mathbb{E}_{s \sim d^\pi, a \sim \pi(\cdot|s)} [\log \pi(a|s) A^\pi(s, a)]$ (Mnih et al. 2016; Schulman et al. 2016) with respect to the parameters of π , where $d^\pi(s)$ is the marginal state density induced by π , and the advantage function A^π can be estimated as $\mathcal{T}^\pi V^\pi(s) - V^\pi(s)$.

Distributional RL In distributional reinforcement learning, the distributions of returns instead of their expectations (i.e. value functions) are maintained. The distributional Bellman equation in terms of the state-action return is (Bellemare, Dabney, and Munos 2017)

$$\mathcal{T}^\pi G^\pi(s, a) \stackrel{D}{=} R(s, a) + \gamma G^\pi(S', A'). \quad (6)$$

The distribution equation $U \stackrel{D}{=} V$ specifies that the random variable U is distributed by the same law as is random variable V . The reward $R(s, a)$, next state-action tuple S', A' and its return $G^\pi(S', A')$ are random variables, with compound randomness stemmed from π, P , and \mathcal{R} .

Eq. (6) is a contraction mapping in the p -th order Wasserstein metrics W_p (Bellemare, Dabney, and Munos 2017). Previously, Eq. (6) is exploited in a Q -learning style value iteration, with the distribution of $G^\pi(s, a)$ represented as a particle set, updated either through cross-entropy loss (Bellemare, Dabney, and Munos 2017), quantile regression (Dabney et al. 2018a,b; Rowland et al. 2019), or Sinkhorn iterations (Martin et al. 2020). Particle-based (ensemble-)critics $G^\pi(s, a)$ are incorporated into conventional off-policy policy gradient methods by (Barth-Maron et al. 2018) and (Kuznetsov et al. 2020). A continuous $G^\pi(s, a)$ distribution can be conferred via Wasserstein-GAN (WGAN) (Arjovsky, Chintala, and Bottou 2017), and has been investigated in both Q -learning (Doan, Mazouze, and Lyle 2018) and policy gradients (Freirich et al. 2019). We remark that these works all estimate return distributions with empirical approximations, e.g. particle set or WGAN.

Methods

We begin with proving that the distributional Bellman operation in terms of *state-return* distributions is also a contraction mapping in Wasserstein metrics. We then show resemblance between distributional Bellman update and a variational Bayesian solution to return distributions, leading to a

novel distributional RL approach. Thereafter, we propose an internal incentive that leverages posterior IG stemmed from return estimation inaccuracy.

Distributional Bellman Operator for State-Return

First, in the same sense that Eq. (6) extends Eq. (5), we extend Eq. (4) and define the distributional Bellman operator regarding the state return $G^\pi(s)$ as

$$\mathcal{T}^\pi G^\pi(s) \stackrel{D}{=} R(s) + \gamma G^\pi(S'). \quad (7)$$

Now we demonstrate that Eq. (7) is also a contraction in p -Wasserstein metrics.

For notional convenience, we write the infimum- p -Wasserstein metric in Eq. (2) in terms of random variables: $d_p(X, Y) := W_p(\alpha, \beta), X \sim \alpha, Y \sim \beta$.

Let $\mathcal{G} \in \mathbb{R}^n$ denote a space of returns valid in the MDP, and $\Omega \in \mathcal{P}(\mathcal{G})(\mathcal{S})$ a space of state-return distributions with bounded moments. Represent as ω the collection of distributions $\{\omega(s) | s \in \mathcal{S}\}$, in which $\omega(s)$ is the distribution of random return $G(s)$. For any two distributions $\omega_1, \omega_2 \in \Omega$, the supremum- p -Wasserstein metric on Ω is defined as (Bellemare, Dabney, and Munos 2017; Rowland et al. 2018)

$$\bar{d}_p(\omega_1, \omega_2) := \sup_{s \in \mathcal{S}} d_p(G_1(s), G_2(s)). \quad (8)$$

Lemma 2. \bar{d}_p is a metric over state-return distributions.

The proof is a straightforward analogue to that of Lemma 2 in (Bellemare, Dabney, and Munos 2017), substituting the state space \mathcal{S} for state-action space $\mathcal{S} \times \mathcal{A}$.

Proposition 1. The distributional Bellman operator for state-return distributions is a γ -contraction in \bar{d}_p .

Proof. The reward $R(s) \in \mathcal{G}$ is a random vector such that $R(s) = \int_{\mathcal{A}} \mathcal{R}(s, a) \tilde{\pi}(a|s) da$, where $\tilde{\pi}$ denotes the normalised policy π .

Represent the marginal state transition kernel under policy π as $P^\pi(s'|s) = \int_{\mathcal{A}} P(s'|s, a) \tilde{\pi}(a|s) da$. Then define a corresponding transition operator $\mathcal{P}^\pi : \mathcal{G} \mapsto \mathcal{G}$

$$\mathcal{P}^\pi G(s) \stackrel{D}{=} G(S'), S' \sim P^\pi(\cdot|s). \quad (9)$$

With the marginal state transition operator substituted for the action-dependent one, the rest of the proof is analogous to that of Lemma 3 presented by (Bellemare, Dabney, and Munos 2017). \square

We therefore conclude that Eq. (7) has a unique fixed point G^π . Proposition 1 vindicates backing up distributions of the state return $G^\pi(s)$ by minimising Wasserstein metrics to a target distribution.

For the policy gradient theorem (Sutton et al. 1999) to hold, one would need at each encountered s_t an unbiased estimator of $\mathbb{E}[\sum_{k=0}^{\infty} \gamma^k r_{t+k}]$ in computing the policy gradient. In distributional RL, such a quantity is obtained by sampling from the approximated return distribution (or averaging across such samples). The Bellman operator being a contraction ensures convergence to a unique true on-policy

Algorithm 1 Bayesian Distributional Policy Gradients

```
1: Initialise encoder  $q_\phi(Z|X, S)$ , generator  $G_\theta(Z, S)$ ,  
   prior  $p_\theta(Z|S)$ , discriminator  $D_\psi(X, Z, S)$  and policy  $\pi$   
2: While not converge:  
   // roll out  
3:   training batch  $\mathcal{D} \leftarrow \emptyset$   
4:   For  $t = 0, \dots, k - 1, \forall$  threads:  
5:     execute  $a_t \sim \pi(\cdot|s_t)$ , get  $r_t, s_{t+1}$   
6:     sample return  $z_t \sim p_\theta(\cdot|s_t)$ ,  $g_t \leftarrow G_\theta(z_t, s_t)$   
7:     update  $\mathcal{D}$   
8:   sample last return  $z_k \sim p_\theta(\cdot|s_k)$ ,  $g_k \leftarrow G_\theta(z_k, s_k)$   
   // estimate advantage for whole batch  
9:   For  $t \in \mathcal{D}$ :  
10:    estimate advantage  $\hat{A}_t$  with  $r_{t:t+k-1}, g_{t:t+k}$   
    using any estimation method  
11:    Bellman target  $x_t \leftarrow \hat{A}_t + g_t$   
12:    get curiosity reward  $r_t^c$  by Eq.(13)-(14)  
13:    get augmented advantage  $\hat{A}_t^c$  by substituting  $r_t$  in  $\hat{A}_t$   
    with  $r_t + r_t^c$   
    // train with mini batch  $B \subset \mathcal{D}$   
14:    For  $t \in B$ :  
15:      encode  $\tilde{z}_t \sim q_\phi(\cdot|x_t, s_t)$   
16:      sample from joint  $p_\theta: z_t \sim p_\theta(\cdot|s_t), \tilde{x}_t \leftarrow G_{\bar{\theta}}(z_t, s_t)$   
    // take gradients  
17:    update  $D_\psi$  by ascending  
     $\frac{1}{|B|} \sum_{t \in B} \log D_\psi(\tilde{x}_t, z_t, s_t) + \log(1 - D_\psi(x_t, \tilde{z}_t, s_t))$   
18:    update encoder, prior by ascending  
     $\frac{1}{|B|} \sum_{t \in B} \log(1 - D_\psi(\tilde{x}_t, z_t, s_t)) + \log D_\psi(x_t, \tilde{z}_t, s_t)$   
19:    update  $G_\theta$  by descending  $\frac{1}{|B|} \sum_{t \in B} \|x_t - G_\theta(\tilde{z}_t, s_t)\|_2^2$   
20:    update  $\pi$  by ascending  $\frac{1}{|B|} \sum_{t \in B} \log \pi(a_t|s_t) \hat{A}_t^c$   
    using any policy gradient method  
21: Return  $\pi$ 
```

return distribution, whose expectation is thereby also unbiased. The same holds also for sample estimates of the state-conditioned reward-to-go and thereby for the advantage function.

Inference in Distributional Bellman Update

We now proceed to show that the distribution of $\mathcal{T}^\pi G^\pi(s)$ can be interpreted as the target distribution p_X , and hence propose a new approach to distributional RL. Specifically, let the data space $\mathcal{X} = \mathcal{G}$ be the space of returns. $\forall s \in \mathcal{S}$, we shorthand as such $x(s) := \mathcal{T}^\pi G^\pi(s)$, $g(s) := G^\pi(s)$, thus $x(s), g(s) \in \mathcal{X}$. We view the Bellman target $x(s)$ as a sample from the empirical data distribution $x(s) \sim p_X(X|s)$, whilst the estimated return $g(s)$ is generated from the model distribution $g(s) \sim p_G(X|s)$. The state s is an observable condition to the generative model: its distribution is of no interest to and not modelled in the Bayesian system.

We factorise the s -conditioned sampling distributions in Lemma 1 such that

$$\begin{aligned} p_\theta(X, Z|s) &:= p_\theta(X|Z, s)p_\theta(Z|s), \\ q_\phi(X, Z|s) &:= p_X(X|s)q_\phi(Z|X, s). \end{aligned} \quad (10)$$

$p_\theta(X|Z, s) = \delta_{G_\theta(Z, s)}(X)$ is a deterministic decoder.

The intuition of a state-conditioned, learned prior for Z instead of a simple, fixed one, is to add stochasticity for the

prior and encoder to meet halfway. Similar to the encoder, we represent the prior also in a variational fashion and sample through re-parameterisation during gradient estimation.

Lemma 1 implies that training D_ψ and the generative model alternately with factorisation in Eq. (10) would suffice to both have the encoder $q_\phi(Z|X, s)$ approximating the true posterior $p_\theta(Z|X, s) := p_\theta(X|Z, s)p_\theta(Z|s)/p_X(X|s)$ and to reconstruct in \mathcal{X} (Dumoulin et al. 2017; Donahue, Krähenbühl, and Darrell 2017). Notice that a globally observable condition s is orthogonal to Lemma 1 and Eq. (3). And so is a learned prior: both the true posterior and the push-forward are *relative to the prior* $p_\theta(Z|s)$.

In our work, in contrast, G_θ is deemed fixed in relation to the minimax game, leaving the *encoder, prior and discriminator* to be trained in the minimax game

$$\begin{aligned} \max_{D_\psi} \min_{q_\phi, p_\theta} \mathbb{E}_{z \sim p_\theta(Z|s)} [\log D_\psi(G_{\bar{\theta}}(z, s), z, s)] \\ + \mathbb{E}_{x \sim p_X(X|s)} \mathbb{E}_{z \sim q_\phi(Z|X, s)} [\log(1 - D_\psi(x, z, s))]. \end{aligned} \quad (11)$$

The overhead bar ($\bar{\cdot}$) denotes that gradient is not back-propagated through the parameter in question. This means $q_\phi(Z|X, s)$ is still trained to approximate the true posterior induced by the current G_θ , irrespective of the capability of the latter for reconstruction. Meanwhile, the reconstruction is achieved by minimising a Wasserstein metric in \mathcal{X}

$$\min_{G_\theta} \mathbb{E}_{x \sim p_X(X|s)} \mathbb{E}_{z \sim q_\phi(Z|X, s)} [d^p(x, G_\theta(z, s))]. \quad (12)$$

Essentially, we are alternating between training the encoder and prior via Eq. (11) and training the generator via Eq. (12). We will use a fixed prior p_Z and omit state dependence in the ensuing discussion, as they do not affect convergence. If the encoder approximates the true posterior everywhere in \mathcal{X} , the aggregated posterior $Q(Z)$ is naturally matched to the prior $p_Z(Z)$, so long as $p_\theta(X|Z)$ is properly normalised, as is indeed the case when it's degenerate. As such, meeting the constraint on the search space in Eq. (3) is a necessary condition to accurate posterior approximation.

Note that in Eq. (3), $\mathbb{E}_{p_X} \mathbb{E}_{q_\phi} [G_\theta(Z)]$ is the push-forward of $Q(Z)$, $G_{\theta\#}Q$. The primal form of $W_p(p_X, p_G)$, where $p_G = G_{\theta\#}p_Z$, is thereby the infimum of $W_p(p_X, G_{\theta\#}Q)$ over q_ϕ s.t. $Q = p_Z$. Therefore, $W_p(p_X, G_{\theta\#}Q)$ is an upper bound to the true objective $W_p(p_X, G_{\theta\#}p_Z)$ upon $Q = p_Z$.

Learning converges as we explain in the following. And to provide intuition, we highlight the resemblances to the Expectation-Maximisation (EM) algorithm. Eq. (11) enforces contrastive learning such that the variational posterior approaches the true posterior, comparable to the E-step in EM. Eq. (11) allows to compute a bound $W_p(p_X, G_{\theta\#}Q)$ in Eq. (12), which is equivalent to the computationally tractable surrogate objective function of the negative free energy in EM, or ELBO in variational Bayes. The expected Wasserstein metric w.r.t. the current q_ϕ is then minimised by updating the parameters of the decoder via Eq. (12). This update is reminiscent of the M-step in EM, which maximises the expected log likelihood while fixing inference for Z .

In our method $W_p(p_X, G_{\theta\#}Q)$ acts as an upper bound when $Q = p_Z$, whereas in EM the surrogate objective is

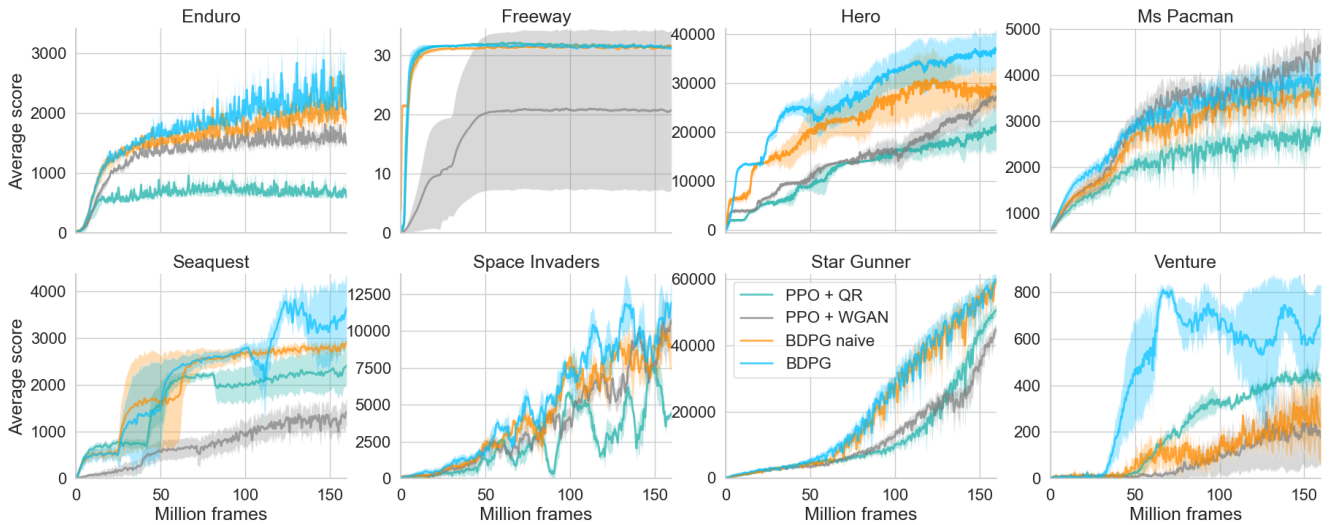


Figure 1: Learning curves on Atari games with the mean (solid line) and standard deviation (shaded area) across 5 runs.

a lower bound. This upper bound decreases in Eq. (11) as it approaches the true objective $W_p(p_X, G_{\theta\#}p_Z)$. Eq. (12) then decreases $W_p(p_X, G_{\theta\#}Q)$ further and consequently also decreases $W_p(p_X, G_{\theta\#}p_Z)$. Note, the condition $Q = p_Z$ does not have to hold on each iteration, but can be amortised over iterations. Assuming infinite model expressiveness, the discrepancy between Q and p_Z shrinks monotonically, as all determinant functions for $Q := \mathbb{E}_{p_X}[q_\phi(\cdot|x)] = p_Z$ in Eq. (11) are fixed irrespective of the value of θ . When q_ϕ converges to the true posterior, $W_p(p_X, G_{\theta\#}Q)$ is more sufficiently an upper bound due to restricted search space in the primal form. While $W_p(p_X, G_{\theta\#}Q)$ functions as an amortised upper bound, $W_p(p_X, G_{\theta\#}p_Z)$ still decreases continually (as opposed to from each iteration) and converges to a local minimum.

The merit of the two-step training is two-fold: 1) with only the distributions over Z under tuning in the minimax game, the adversarial training comes off with a weaker topology and is not relied upon for reconstruction, making its potential instability less of a concern; and 2) an explicit distance loss d in \mathcal{X} minimises W_p to ensure contraction of return distribution backups. If everything was trained adversarially in JS divergence and allowed to reach global optimum, the decoder and encoder would be reversing each other both in density domain. In our setting, Q is matched to p_Z everywhere in \mathcal{Z} , while p_G has minimum W_p distance to p_X .

At each step of environmental interaction, a state return is sampled via the standard two steps $g(s) \sim p_G(X|s) \iff z \sim p_\theta(Z|s), g(s) \leftarrow G_\theta(z, s)$. The one-step Bellman target $x(s)$ is calculated as $r + \gamma g(s')$. Generalisation to k -step bootstrap can be made analogously to the conventional RL.

Exploration through Posterior Information Gain

Curiosity (Schmidhuber 1991; Schaul et al. 2016; Houthoofd et al. 2016; Freirich et al. 2019) produces internal incentives when external reward is sparse. We explore through encouraging visits to regions where the model’s ability to predict

the reward-to-go from current return distribution is weak. However, the Bellman error $x(s) - g(s)$ is not a preferable indicator, as high $x(s) - g(s)$ may well be attributed to high moments of $g(s)$ itself under point estimation (i.e. the *aleatoric* uncertainty), whereas it is the uncertainty in value belief due to estimating parameters with limited data around the state-action tuple (i.e. the *epistemic* uncertainty) that should be driving strengthened visitation.

To measure the true reduction in uncertainty about return prediction, we estimate discrepancies in function space instead of parameter space. Specifically, the insufficiency in data collection can be interpreted as how much a posterior distribution of a statistic or parameter inferred from a condition progresses from a prior distribution with respect to the action execution that changes this condition, i.e., the IG. A large IG means a large amount of data is required to achieve the update. In its simplest form, the condition is implicitly the data trained on. In exact Bayes, the condition itself can be thought of as a variable estimated from data, e.g. the random return X , hence enabling an explicit IG derived from existing posterior model $q_\phi(Z|X)$. Therefore, we define the IG $u(s)$ at s in return estimation as

$$u(s) := \text{KL}(q_\phi(Z|x(s), s) || q_\phi(Z|g(s), s)). \quad (13)$$

Before the transition, the agent’s estimation for return is $g(s)$. The action execution enables the computation of the Bellman target $x(s)$, which would not be viable before the transition, in which $q_\phi(Z|g(s), s)$ acts here as a prior. As a result, $u(s)$ would encourage the agent to make transitions that maximally acquire new information about Z , hence facilitating updating p_G towards p_X . Upon convergence, $g(s)$ and $x(s)$ are indistinguishable and the IG approaches 0. The benefit of our IG is three-fold: it is moments-invariant, makes use of all training data, and increases computation complexity only in forward-passing the posterior model when calculating the KL divergence without even requiring gradient backpropagation.

The curiosity reward $r^c(s)$ is determined by $u(s)$ and a truncation scheme $\mathcal{R} : \mathbb{R}^+ \mapsto \mathbb{R}^{[0, \eta \bar{u}]}$, $\eta, \bar{u} \in \mathbb{R}_*^+$, to prevent radical exploration

$$r^c(s) := \mathcal{R}(u(s)) := \eta_t \cdot \min(u(s), \bar{u}). \quad (14)$$

We exploit relative value by normalising the clipped $u(s)$ by a running mean and standard deviation of previous IGs. The exploration coefficient η_t is logarithmically decayed as $\eta_t = \eta \sqrt{\log t/t}$, by the rate at which the parametric uncertainty decays (Koenker 2005; Mavrin et al. 2019), where t is the global training step, and η an initial value, to assuage exploration getting more sensitive to the value of $u(s)$ as parameters become more accurate.

We use r^c to augment return backup during policy update, as the purpose is for the action to lead to uncertain regions by encouraging curiosity about future steps. When training the generative model for return distributions we use the original reward only.

We investigate a multi-step advantage function. The contraction property of the distributional Bellman operator is propagated from 1-step to k -step scenarios by the same logic as in conventional RL. The benefit of looking into further steps for exploration is intuitive viewed from the long-term goal of RL tasks: the agent should not be complacent about a state just because it is informative to immediate steps.

The pseudocode in Algorithm 1 presents a mini-batch version of our methodology BDPG. We denote state return $g(s_t)$ as g_t for compactness. Other step-dependent values are shorthanded accordingly. We use Euclidean distance for reconstruction, leading to the W_2 metric being minimised. k is the number of unroll steps, and is also the maximum bootstrap length, albeit the two are not necessarily the same.

Related Work

Policy optimisation enables importance sampling based off-policy evaluation for re-sampling weights in experience replay schemes (Wang et al. 2016; Gruslys et al. 2018). In continuous control, where the policy is usually a parametric Gaussian, exploration can be realised by perturbing the Gaussian mean (Lillicrap et al. 2015; Ciosek et al. 2019), or maintaining a mixture of Gaussians (Lim et al. 2018). Alternatively, random actions can be directly incentivised by regularising policy (cross-)entropy (Abdolmaleki et al. 2015; Mnih et al. 2016; Nachum, Norouzi, and Schuurmans 2016; Akrouf et al. 2016; Haarnoja et al. 2018).

A group of works propose to exploit epistemic uncertainty via an approximate posterior distribution of Q values. Stochastic posterior estimates are constructed through training on bootstrapped perturbations of data (Osband et al. 2016, 2019), or overlaying learned posterior variance (Chen et al. 2017; O’Donoghue et al. 2018). While this series of works can be thought of as posterior sampling w.r.t. Q values, (Tang and Agrawal 2018) approximates Bayesian inference by sampling parameters for a distributional RL model. On the other hand, a particle-based distributional RL model itself registers notion of dispersion, inspiring risk-averse and risk-seeking policies (Dabney et al. 2018a) and optimism-in-the-face-of-uncertainty quantified by the variance of the better half of the particle set (Mavrin et al. 2019).




Figure 2: Impact of bootstrap length k and truncation cap \bar{u} for information gain at 10M and 150M steps into training.

There are also approaches exploiting dynamics uncertainty (Houthoofd et al. 2016), pseudo counts (Bellemare et al. 2016; Ostrovski et al. 2017; Tang et al. 2017), gradient of a generative model (Freirich et al. 2019), and good past experiences (Oh et al. 2018), that do not estimate dispersion or model disagreement of value functions.

Evaluation

We evaluate our method on the Arcade Learning Environment Atari 2600 games and continuous control with MuJoCo. We estimate a k -step advantage function using Generalised Advantage Estimation (GAE) (Schulman et al. 2016), and update the policy using Proximal Policy Optimisation (PPO) (Schulman et al. 2017) which maximises a clipped surrogate of the policy gradient objective. For both Atari and MuJoCo environments, we use 16 parallel workers for data collection, and train in mini batches. For Atari, we unroll 128 steps with each worker in each training batch for all algorithms, and average scores every 80 training batches. For MuJoCo, we unroll 256 steps, and average scores every 4 batches. Except for ablation experiments that used rollout length $\max(k, 128)$, the number of unroll steps is also the bootstrap length k .

We compare to other distributional RL baselines on eight of the Atari games, including some of the recognised hard-exploration tasks: Freeway, Hero, Seaquest and Venture.

Direct comparisons to previous works are not meaningful due to compounded discrepancies. To allow for a meaningful comparison, we implement our own versions of baselines, fixing other algorithmic implementation choices such that the tested algorithms vary only in how return distributions are estimated and in the exploration scheme. We modify two previous algorithms (Freirich et al. 2019) and (Dabney et al. 2018b), retaining their return distribution estimators as benchmarks: a generative model Wasserstein-GAN

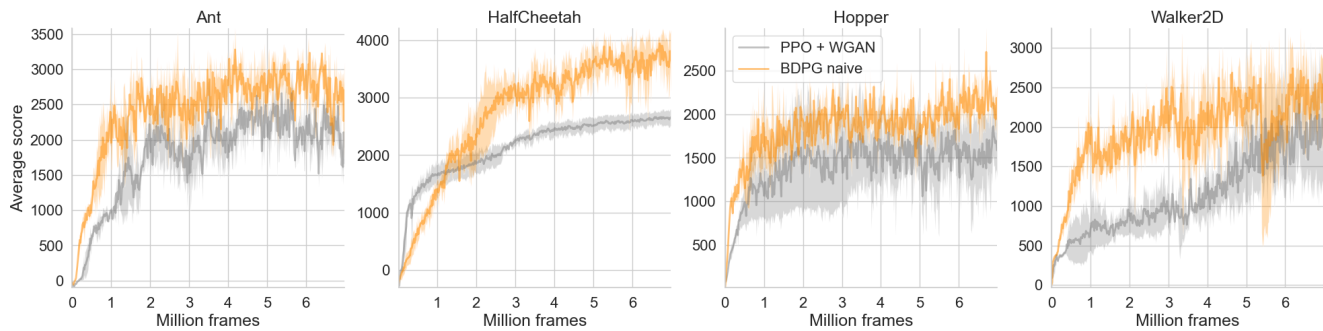


Figure 3: Learning curves on MuJoCo tasks with the mean (solid line) and standard deviation (shaded area) across 5 runs.

(Arjovsky, Chintala, and Bottou 2017) (PPO+WGAN), and a discrete approximation of distribution updated through quantile regression (Koenker 2005) (PPO+QR). Importantly, all of our baselines are *distributional* RL solutions that maintain *state*-return distributions.

Our BDPG is evaluated in two versions: 1) with the naive add-noise-and-argmax (Mnih et al. 2016) exploration mechanism (BDPG naive), and 2) one that explores with the proposed curiosity reward (BDPG). Naive exploration is also used for the PPO+WGAN and PPO+QR baselines. Learning curves in Figure 1 suggest that with exploration mechanism fixed, the proposed Bayesian approach BDPG naive outperforms or is comparable to WGAN and QR in 6 out of 8 games. Moreover, BDPG is always better or equal to BDPG naive, vindicating our exploration scheme, and is able to get the highest score among all tested algorithms in all four of the hard-exploration games tested on.

We conduct ablation and parameter studies to investigate the impact of the bootstrap length k , and of the truncation cap \bar{u} on the IG $u(s)$, on Atari games Breakout and Q*bert. In particular, $k = 1$ and $\bar{u} \rightarrow \infty$ are looked at as ablation cases. Average scores of the batch started at 10M and 150M steps into training are shown in Figure 2. Each coloured pixel corresponds to the best outcome with respect to η value among its selection sweep according to average long-term performance for each combination of k and \bar{u} . We found that as training progresses, short k comes to display a more prohibitive effect, as the model becomes more discriminative about the environment, and lack of learning signals, i.e. fewer rewards to calculate the Bellman target with, becomes increasingly suppressive. Our experiments suggest that although the best bootstrap length depends on the task, longer bootstrapping generally produces better long-term performance. But a long bootstrap length does not work well with a large \bar{u} , a possible explanation is that as k increases, the variance in the Bellman targets multiplies. In this scenario, the agent may encounter states with which it is very unfamiliar. The value of $u(s)$ can grow unbounded and the tendency to explore get out of hand if we do not curb it. Moreover, such extreme values can also jeopardise subsequent curiosity comprehension through the normalisation of $u(s)$. This phenomenon justifies the application of our truncation scheme, especially for larger k . In addition, we found that

choosing too large k does not diminish performance drastically, potentially due to the return distribution already accounting for some degree of reward uncertainty, which is a helpful characteristic when prototyping agents.

In the continuous control tasks with MuJoCo, we focus on the ability of distributional RL algorithms to generalise, and less the challenge of exploration. Therefore, we compare the performance of BDPG naive against the benchmark distributional RL algorithm PPO+WGAN, a generative solution that does not conduct inference. Both are stripped of exploration incentives. Noticeable amounts of variance displayed throughout training for both algorithms may be due to that they both involve adversarial training. As shown in Figure 3, however, our model outperforms the benchmark in all cases with distinct margins. We believe this is because WGAN does not take expectations across an amortised inference space that accounts for better generalisation. This proves to be highly beneficial for reasoning about return distributions in continuous control tasks such as MuJoCo environments, where robustness in the face of unseen data weighs up more in behaviour stability.

Conclusion

We formulate the distributional Bellman operation as an inference-based auto-encoding process. We demonstrated contraction of the Bellman operator for state-return distributions, expanding on the distributional RL body of work that focused on state-action returns, to date. Our tractable solution alternates between minimising Wasserstein metrics to continuous distributions of the Bellman target and adversarial training for joint-contrastive learning to estimate a variational posterior from bootstrapped targets. This allows us to distill the benefits of Bayesian inference into distributional RL, where predictions of returns are based on expectation and thus more accurate in the face of unseen data. As a second innovation we use the availability of a variational posterior to derive a curiosity-driven exploration mechanism, which we show is more efficiently solving hard-exploration tasks. Either of our two contributions can be combined with other building blocks to form new RL algorithms, e.g. in Explainable RL (Beyret, Shafti, and Faisal 2019). We believe that our innovations link and expand the applicability and efficiency of distributional RL methods.

Acknowledgments

We are grateful for our funding support: a Department of Computing PhD Award to LL and a UKRI Turing AI Fellowship (EP/V025449/1) to AAF.

References

- Abdolmaleki, A.; et al. 2015. Model-Based Relative Entropy Stochastic Search. In *Adv. Neural Inform. Proc. Sys. (NIPS)* 28, 3537–3545.
- Akrour, R.; et al. 2016. Model-Free Trajectory Optimization for Reinforcement Learning. In *Proc. the 33rd Intl. Conf. on Machine Learning (ICML)*, volume 48, 2961–2970. New York, NY, USA.
- Ambrogioni, L.; et al. 2018. Wasserstein Variational Inference. In *Adv. Neural Inform. Proc. Sys. (NIPS)* 31, 2473–2482.
- Arjovsky, M.; Chintala, S.; and Bottou, L. 2017. Wasserstein Generative Adversarial Networks. In *Proc. the 34th Intl. Conf. on Machine Learning (ICML)*, volume 70, 214–223. Sydney, Australia.
- Ball, P.; et al. 2020. Ready Policy One: World Building Through Active Learning. In *Proc. the 37th Intl. Conf. on Machine Learning (ICML)*, volume 119, 591–601. Virtual.
- Barth-Maron, G.; et al. 2018. Distributed Distributional Deterministic Policy Gradients. In *Proc. 6th Intl. Conf. on Learning Representations (ICLR)*.
- Bellemare, M.; et al. 2016. Unifying Count-Based Exploration and Intrinsic Motivation. In *Adv. Neural Inform. Proc. Sys. (NIPS)* 29, 1471–1479.
- Bellemare, M. G.; Dabney, W.; and Munos, R. 2017. A Distributional Perspective on Reinforcement Learning. In *Proc. the 34th Intl. Conf. on Machine Learning (ICML)*, volume 70, 449–458. Sydney, Australia.
- Bellemare, M. G.; et al. 2013. The Arcade Learning Environment: An Evaluation Platform for General Agents. In *Journal of Artificial Intelligence Research*.
- Bellman, R. 1957. *Dynamic Programming*. Princeton, NJ, USA: Princeton University Press, 1st edition.
- Beyret, B.; Shafti, A.; and Faisal, A. A. 2019. Dot-to-Dot: Explainable Hierarchical Reinforcement Learning for Robotic Manipulation. In *2019 IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 5014–5019. IEEE.
- Blundell, C.; et al. 2015. Weight Uncertainty in Neural Network. In *Proc. the 32nd Intl. Conf. on Machine Learning (ICML)*, volume 37, 1613–1622. Lille, France.
- Bousquet, O.; et al. 2017. From Optimal Transport to Generative Modeling: the VEGAN cookbook. *arXiv* 1705.07642.
- Chen, R. Y.; et al. 2017. UCB Exploration via Q-Ensembles. *arXiv* 1706.01502.
- Ciosek, K.; et al. 2019. Better Exploration with Optimistic Actor Critic. In *Adv. Neural Inform. Proc. Sys. (NIPS)* 32, 1787–1798.
- Cuturi, M. 2013. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In *Adv. Neural Inform. Proc. Sys. (NIPS)* 26, 2292–2300.
- Dabney, W.; et al. 2018a. Implicit Quantile Networks for Distributional Reinforcement Learning. In *Proc. the 35th Intl. Conf. on Machine Learning (ICML)*, volume 80, 1096–1105. Stockholm, Sweden.
- Dabney, W.; et al. 2018b. Distributional Reinforcement Learning With Quantile Regression. In *Proc. the 32nd AAAI Conf. on Artificial Intelligence*.
- Doan, T.; Mazouze, B.; and Lyle, C. 2018. GAN Q-learning. *arXiv* 1805.04874.
- Donahue, J.; Krähenbühl, P.; and Darrell, T. 2017. Adversarial Feature Learning. In *Proc. the 5th Intl. Conf. Learning Representations (ICLR)*.
- Dumoulin, V.; et al. 2017. Adversarially Learned Inference. In *Proc. the 5th Intl. Conf. Learning Repres. (ICLR)*.
- Freirich, D.; et al. 2019. Distributional Multivariate Policy Evaluation and Exploration with the Bellman GAN. In *Proc. the 36th Intl. Conf. on Machine Learning (ICML)*, volume 97, 1983–1992. Long Beach, CA, USA.
- Genevay, A.; et al. 2016. Stochastic Optimization for Large-scale Optimal Transport. In *Adv. Neural Inform. Proc. Sys. (NIPS)* 29, 3440–3448.
- Goodfellow, I. J.; et al. 2014. Generative Adversarial Networks. In *Adv. Neural Inform. Proc. Sys. (NIPS)* 27, 2672–2680.
- Gruslys, A.; et al. 2018. The Reactor: A Sample-Efficient Actor-Critic Architecture. In *Proc. the 6th Intl. Conf. Learning Representations (ICLR)*.
- Gulrajani, I.; et al. 2017. Improved Training of Wasserstein GANs. In *Adv. Neural Inform. Proc. Sys. (NIPS)* 30, 5769–5779.
- Haarnoja, T.; et al. 2018. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. In *Proc. the 35th Intl. Conf. on Machine Learning (ICML)*, volume 80, 1861–1870.
- He, J.; et al. 2019. Lagging Inference Networks and Posterior Collapse in Variational Autoencoders. In *Proc. the 7th Intl. Conf. Learning Representations (ICLR)*.
- Hessel, M.; et al. 2018. Rainbow: Combining Improvements in Deep Reinforcement Learning. In *Proc. the 32nd AAAI Conf. on Artificial Intelligence*.
- Hoffman, M. D.; and Johnson, M. J. 2016. ELBO surgery: yet another way to carve up the variational evidence lower bound. In *In Workshop of Approximate Bayesian Inference in Neural Information Processing Systems* 29.
- Houthoofd, R.; et al. 2016. VIME: Variational Information Maximizing Exploration. In *Adv. Neural Inform. Proc. Sys. (NIPS)* 29, 1109–1117.
- Koenker, R. 2005. *Quantile Regression*. Cambridge University Press.

- Kuznetsov, A.; et al. 2020. Controlling Overestimation Bias with Truncated Mixture of Continuous Distributional Quantile Critics. In *Proc. 37th Intl. Conf. on Machine Learning (ICML)*. Virtual.
- Lattimore, T.; and Hutter, M. 2012. PAC Bounds for Discounted MDPs. In *Proc. the 23rd Intl. Conf. on Algorithmic Learning Theory*, 320–334.
- Lillicrap, T. P.; et al. 2015. Continuous control with deep reinforcement learning. In *Proc. the 3rd Intl. Conf. Learning Representations (ICLR)*.
- Lim, S.; et al. 2018. Actor-Expert: A Framework for using Q-learning in Continuous Action Spaces. *arXiv* 1810.09103.
- Luise, G.; et al. 2018. Differential Properties of Sinkhorn Approximation for Learning with Wasserstein Distance. In *Adv. Neural Inform. Proc. Sys. (NIPS) 31*, 5859–5870.
- Martin, J.; et al. 2020. Stochastically Dominant Distributional Reinforcement Learning. In *Proc. the 37th Intl. Conf. on Machine Learning (ICML)*. Virtual.
- Mavrin, B.; et al. 2019. Distributional Reinforcement Learning for Efficient Exploration. In *Proc. the 36th Intl. Conf. on Machine Learning (ICML)*, volume 97, 4424–4434. Long Beach, CA, USA.
- Mnih, V.; et al. 2016. Asynchronous Methods for Deep Reinforcement Learning. In *Proc. the 33rd Intl. Conf. on Machine Learning (ICML)*, volume 48, 1928–1937. New York, NY, USA.
- Montavon, G.; Müller, K.-R.; and Cuturi, M. 2016. Wasserstein Training of Restricted Boltzmann Machines. In *Adv. Neural Inform. Proc. Sys. (NIPS) 29*, 3718–3726.
- Morimura, T.; et al. 2010. Parametric Return Density Estimation for Reinforcement Learning. In *Proc. the 26th Conf. on Uncertainty in Artificial Intelligence*.
- Nachum, O.; Norouzi, M.; and Schuurmans, D. 2016. Improving Policy Gradient by Exploring Under-appreciated Rewards. In *Proc. the 4th Intl. Conf. Learning Representations (ICLR)*.
- O’Donoghue, B.; et al. 2018. The Uncertainty Bellman Equation and Exploration. In *Proc. the 35th Intl. Conf. on Machine Learning (ICML)*, volume 80, 3839–3848. Stockholm, Sweden.
- Oh, J.; et al. 2018. Self-Imitation Learning. In *Proc. the 35th Intl. Conf. on Machine Learning (ICML)*, volume 80, 3878–3887. Stockholm, Sweden.
- Osband, I.; et al. 2016. Deep Exploration via Bootstrapped DQN. In *Adv. Neural Inform. Proc. Sys. (NIPS) 29*, 4026–4034.
- Osband, I.; et al. 2019. Deep Exploration via Randomized Value Functions. *Journal of Machine Learning Research* 20: 1–62.
- Ostrovski, G.; et al. 2017. Count-Based Exploration with Neural Density Models. In *Proc. the 34th Intl. Conf. on Machine Learning (ICML)*, volume 70, 2721–2730. International Convention Centre, Sydney, Australia.
- Puterman, M. L. 1994. *Markov Decision Processes : Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc.
- Rosca, M.; Lakshminarayanan, B.; and Mohamed, S. 2018. Distribution Matching in Variational Inference. *arXiv* 1802.06847.
- Rowland, M.; et al. 2018. An Analysis of Categorical Distributional Reinforcement Learning. In *Proc. 21st Intl. Conf. on Artificial Intelligence and Statistics (AISTATS)*, volume 84. Lanzarote, Spain.
- Rowland, M.; et al. 2019. Statistics and Samples in Distributional Reinforcement Learning. In *Proc. the 36th Intl. Conf. on Machine Learning (ICML)*, volume 97, 5528–5536. Long Beach, CA, USA.
- Schaul, T.; et al. 2016. Prioritized Experience Replay. In *Proc. the 4th Intl. Conf. Learning Representations (ICLR)*.
- Schmidhuber, J. 1991. Curious model-building control systems. In *[Proceedings] 1991 IEEE International Joint Conf. on Neural Networks*, volume 2, 1458–1463.
- Schulman, J.; et al. 2016. High-Dimensional Continuous Control Using Generalized Advantage Estimation. In *Proc. the 4th Intl. Conf. Learning Representations (ICLR)*.
- Schulman, J.; et al. 2017. Proximal Policy Optimization Algorithms. *arXiv* 1707.06347.
- Sekar, R.; et al. 2020. Planning to Explore via Self-Supervised World Models. In *Proc. the 37th Intl. Conf. on Machine Learning (ICML)*, volume 119, 8583–8592. Virtual.
- Sutton, R. S.; et al. 1999. Policy Gradient Methods for Reinforcement Learning with Function Approximation. In *Adv. Neural Inform. Proc. Sys. (NIPS) 12*, 1057–1063.
- Tang, H.; et al. 2017. #Exploration: A Study of Count-Based Exploration for Deep Reinforcement Learning. In *Adv. Neural Inform. Proc. Sys. (NIPS) 30*, 2753–2762.
- Tang, Y.; and Agrawal, S. 2018. Exploration by Distributional Reinforcement Learning. In *Proc. the 27th International Joint Conf. on Artificial Intelligence*, 2710–2716.
- Todorov, E.; Erez, T.; and Tassa, Y. 2012. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 5026–5033.
- Tolstikhin, I.; et al. 2018. Wasserstein Auto-Encoders. In *Proc. the 6th Intl. Conf. Learning Representations (ICLR)*.
- Villani, C. 2008. *Optimal Transport: Old and New. Grundlehren der mathematischen Wissenschaften*. Springer Berlin Heidelberg.
- Wang, Z.; et al. 2016. Sample Efficient Actor-Critic with Experience Replay. In *Intl. Conf. Learning Repr. (ICLR)*.
- Zhao, S.; Song, J.; and Ermon, S. 2017. InfoVAE: Information Maximizing Variational Autoencoders. *ArXiv* 1706.02262.