
Robust Reinforcement Learning with Distributional Risk-averse formulation

Pierre Clavier^{1,2} Stéphanie Allasonnière² Erwan Le Pennec¹

Abstract

Robust Reinforcement Learning tries to make predictions more robust to changes in the dynamics or rewards of the system. This problem is particularly important when the dynamics and rewards of the environment are estimated from the data. In this paper, we approximate the Robust Reinforcement Learning constrained with a Φ -divergence using an approximate Risk-Averse formulation. We show that the classical Reinforcement Learning formulation can be robustified using standard deviation penalization of the objective. Two algorithms based on Distributional Reinforcement Learning, one for discrete and one for continuous action spaces are proposed and tested in a classical Gym environment to demonstrate the robustness of the algorithms.

1. Introduction

The classical Reinforcement Learning (RL) problem using Markov Decision Processes (MDPs) modelization gives a practical framework for solving sequential decision problems under uncertainty of the environment. However, for real-world applications, the final chosen policy can be sometimes very sensitive to sampling errors, inaccuracy of the model parameters, and definition of the reward. For discrete state-action space, Robust MDPs is treated by Yang (2017); Petrik & Russel (2019); Grand-Clément & Kroer (2020a;b) or (Behzadian et al., 2021) among others. Here we focus on *more general continuous state space* \mathcal{S} with discrete or continuous action space \mathcal{A} and with constraints defined using Φ -divergence. Robust RL (Morimoto & Doya, 2005) with continuous action space focuses on robustness in the dynamics of the system (changes of P) and has been studied in Abdullah et al. (2019); Singh et al. (2020); Urpí et al. (2021); Eysenbach & Levine (2021) among others.

¹Centre de Mathématiques Appliquées, Ecole polytechnique, France ²INRIA HeKA, INSERM, Université Paris Cité, France. Correspondence to: Pierre Clavier <pierre.clavier@polytechnique.edu>.

Eysenbach & Levine (2021) tackles the problem of both reward and transitions using Max Entropy RL whereas the problem of robustness in action noise perturbation is presented in Tessler et al. (2019). Here we tackle the problem of Robustness *thought dynamics of the system*.

In this paper, we show that it is possible to approximate a Robust Distributional Reinforcement Learning heuristic with Φ -divergence constraints into a Risk-averse formulation, using a formulation based on mean-standard deviation optimization. Moreover, we focus on the idea that generalization, regularization, and robustness are strongly linked together as Husain et al. (2021); Derman & Mannor (2020); Derman et al. (2021); Ying et al. (2021); Brekelmans et al. (2022) noticed in the MDPs or RL framework. The contribution of the work is the following: we motivate the use of standard deviation penalization and derive two algorithms for discrete and continuous action space that are Robust to change in dynamics. These algorithms do not require lots of additional parameter tuning, only the Mean-Standard Deviation trade-off has to be chosen carefully. Moreover, we show that our formulation using Distributional Reinforcement Learning is robust to change dynamics on discrete and continuous action spaces from Mujoco suite.

2. Notations

Considering a Markov Decision Process (MDP) $(\mathcal{S}, \mathcal{A}, P, \gamma)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, $P(s', r | s, a)$ is the reward and transition distribution from state s to s' taking action a and $\gamma \in (0, 1)$ is the discount factor. Stochastic policy are denoted $\pi(a | s) : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ and we consider the cases of action space either discrete or continuous.

A rollout or trajectory using π from state s using initial action a is defined as the random sequence $\tau^{P, \pi | s, a} = ((s_0, a_0, r_0), (s_1, a_1, r_1), \dots)$ with $s_0 = s, a_0 = a, a_t \sim \pi(\cdot | s_t)$ and $(r_t, s_{t+1}) \sim P(\cdot, \cdot | s_t, a_t)$; we denote the distribution over rollouts by $\mathbb{P}(\tau)$ with $\mathbb{P}(\tau) = P_0(s_0) \prod_{t=0}^T P(s_{t+1}, r_t | s_t, a_t) \pi(a_t | s_t) d\tau$ and usually write $\tau \sim \mathbb{P} = (P, \pi)$. Moreover, considering the distribution of discounted cumulative return $Z^{P, \pi}(s, a) = R(\tau^{P, \pi | s, a})$ with $R(\tau) = \sum_{t=0}^{\infty} \gamma^t r_t$, the Q -function $Q^{P, \pi} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ of π is the expected discounted cumulative return of the distribution defined as follows :

$Q^{P,\pi}(s, a) := \mathbb{E}[Z^{P,\pi}(s, a)]$. The classical initial goal of RL also called risk-neutral RL, is to find the optimal policy π^* where $Q^{P,\pi^*}(s, a) \geq Q^{P,\pi}(s, a)$ for all π and $s \in \mathcal{S}, a \in \mathcal{A}$. Finally, the Bellman operator \mathcal{T}^π and Bellman optimal operator \mathcal{T}^* are defined as follow: $\mathcal{T}^\pi Q(s, a) := r(s, a) + \gamma \mathbb{E}_{P,\pi}[Q(s', a')]$ and $\mathcal{T}^* Q(s, a) := r(s, a) + \gamma \mathbb{E}_P[\max_{a'} Q(s', a')]$.

Applying either operator from an initial Q^0 converges to a fixed point Q^π or Q^* at a geometric rate as both operators are contractive. Simplifying the notation with regards to s, a, π and P , we define the set of greedy policies w.r.t. Q called $\mathcal{G}(Q) = \arg \max_{\pi \in \Pi} \langle Q, \pi \rangle$. A classical approach to estimating an optimal policy is known as Approximate Modified Policy Iteration (AMPI) (Scherrer et al., 2015)

$$\begin{cases} \pi_{k+1} \in \mathcal{G}(Q_k) \\ Q_{k+1} = (T^{\pi_{k+1}})^m Q_k + \epsilon_{k+1} \end{cases},$$

which usually reduces to Approximate Value Iteration (AVI, $m = 1$) and Approximate Policy Iteration (API, $m = \infty$) as special cases. The term ϵ_{k+1} accounts for errors made when applying the Bellman Operator in RL algorithms with stochastic approximation.

3. Robust formulation in greedy step of API.

In this section, we would like to find a policy that is robust to change of environment law P as small variations of P should not affect too much the new policy in the greedy step. In our case we are not looking at classical greedy step $\pi' \in \mathcal{G}(Q) = \arg \max_{\pi \in \Pi} \langle Q, \pi \rangle$ rather at the following :

$$\pi' \in \mathcal{G}(Q) = \arg \max_{\pi \in \Pi} \langle \min_P Q^{(P,\pi)}, \pi \rangle.$$

This heuristic in the greedy step can also be justified by trying to avoid an overestimation of the Q functions present in the Deep RL algorithms. Using this formulation, we need to constrain the set of admissible transitions from state-action to the next state P to get a solution to the problem. In general, without constraint, the problem is NP-Hard, so it requires constraining the problem to specific distributions that are not too far from the original one using distance between distributions such as the Wasserstein metric (Abdullah et al., 2019) or other specific distances where the problem can be simplified (Eysenbach & Levine, 2021). If a explicit form of $\min_P Q^{(P,\pi)}$ could be computed exactly for a given divergence, it would lead to a simplification of this max-min optimization problem into a simple maximisation one.

In fact, simplification of the problem is possible using specific Φ -divergence denoted \mathcal{H}_Φ to constrain the problem with Φ a closed convex function such that $\Phi : \mathbb{R} \rightarrow \mathbb{R} \cup \{+\infty\}$ and $\Phi(z) \geq \Phi(1) = 0$ for all $z \in \mathcal{R} : \mathcal{H}_\Phi(\mathbb{Q} | \mathbb{P}) = \sum_{i:p_i>0} p_i \Phi\left(\frac{q_i}{p_i}\right)$ with $\sum_{i:p_i>0} q_i = 1$ and

$q_i \geq 0$. This constraint requires $q_i = 0$ if $p_i = 0$ which makes the measure \mathbb{Q} absolutely continuous with respect to \mathbb{P} . The χ^2 -divergence are a particular case of Φ -divergence with $\Phi(z) = (z - 1)^2$. For trajectories sampled from distribution P_0 and looking at distribution P closed to P_0 with regards to the χ^2 -divergence, the minimisation problem reduces to :

$$\min_{P \in D_{\chi^2}(P || P_0) \leq \alpha} Q^{(P,\pi)} = Q^{(P_0,\pi)} - \alpha \mathbb{V}_{P_0}[Z]^{\frac{1}{2}}. \quad (1)$$

The proof can be found in annex A for α such that $\alpha \leq \mathbb{V}_{P_0}[R(\tau)] / \left\| \tilde{R} \right\|_\infty^2 \leq 1$ with $\tilde{R}(\tau) = R(\tau) - \mathbb{E}_{\tau \sim P_0}[R(\tau)]$ the centered return and $\mathbb{V}_{P_0}[Z]$ the variance of returns. For $\alpha > \mathbb{V}_{P_0}[R(\tau)] / \left\| \tilde{R} \right\|_\infty^2$, the equality becomes an inequality but we still optimize a lower bound of our initial problem. Defining a new greedy step which is penalized by the standard deviation :

$$\pi' \in \mathcal{G}_\alpha(Q) = \arg \max_{\pi \in \Pi} \langle Q^{(P_0,\pi)} - \alpha \mathbb{V}_{P_0}[Z]^{\frac{1}{2}}, \pi \rangle$$

we now look at the current AMPI to improve robustness :

$$\begin{cases} \pi_{k+1} \in \mathcal{G}_\alpha(Q_k) \\ Q_{k+1} = (T^{\pi_{k+1}})^m Q_k + \epsilon_{k+1} \end{cases},$$

This idea is very close to Risk-averse formulation in RL (i.e minimizing risk measure and not only the mean of rewards) but here the idea is to approximate a robustness problem in RL. To do so, the standard deviation of the distribution of the returns must be estimated. Many ways are possible but we favour distributional RL (Bellemare et al., 2017; Dabney et al., 2017; 2018) which achieve very good performances in many RL applications. Estimating quantiles of the distribution of return, we can simply estimate standard deviation using classical estimator of the standard deviation given the quantiles over an uniform grid $\{q_i(s, a)\}_{1 \leq i \leq n}$: $\mathbb{V}[Z(s, a)]^{\frac{1}{2}} = \sigma(s, a) = \sqrt{\sum_{i=1}^n (q_i(s, a) - \bar{q}(s, a))^2}$ where \bar{q} is the classical estimator of the mean. A different interpretation of this formulation could be that by taking actions with less variance, we construct a confidence interval with the standard deviation of the distribution :

$$Z^\pi(s, a) \stackrel{d}{=} \bar{Z}(s, a) - \alpha \sigma(s, a).$$

This idea is present in classical UCB algorithms (Auer, 2002) or pessimism/optimism Deep RL. Here we construct a confidence interval using the distribution of the return and not different estimates of the Q function such as in Moskovitz et al. (2021); Bai et al. (2022). In the next section, we derive two algorithms, one for discrete action space and one for continuous action space using this idea. A very interesting way of doing robust Learning is by doing Max entropy RL such as in the SAC algorithm. In Eysenbach & Levine (2021), a demonstration that SAC is a surrogate of Robust RL is demonstrated formally and numerically and we will compare our algorithm to this method.

4. Distributional RL

Distributional RL aims at approximating the return random variable $Z^\pi(s, a) := \sum_{t=0}^{\infty} \gamma^t \mathcal{R}(s_t, a_t)$ with $s_0 = s, a_0 = a, s_{t+1} \sim \mathcal{P}(\cdot | s_t, a_t), a_t \sim \pi(\cdot | s_t)$ and \mathcal{R} that we explicitly treat as a random variable. The classical RL framework approximate the expectation of the return or the Q-function, $Q^\pi(s, a) := \mathbb{E}[Z^\pi(s, a)]$. Many algorithms and distributional representation of the critic exits (Bellemare et al., 2017; Dabney et al., 2017; 2018) but here we focus on QR-DQN (Dabney et al., 2017) that approximates the distribution of returns $Z^\pi(s, a)$ with $Z_\psi(s, a) := \frac{1}{M} \sum_{m=1}^M \delta(\theta_\psi^m(s, a))$, a mixture of atoms-Dirac delta functions located at $\theta_\psi^1(s, a), \dots, \theta_\psi^M(s, a)$ given by a parametric model $\theta_\psi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^M$. Parameters ψ of a neural network are obtained by minimizing the average over the 1-Wasserstein distance between Z_ψ and the temporal difference target distribution $\mathcal{T}_\pi Z_\psi$, where \mathcal{T}_π is the distributional Bellman operator defined in Bellemare et al. (2017). The control version or optimal operator is denoted $\mathcal{T} Z_\psi$, with $\mathcal{T}^\pi Z(s, a) = \mathcal{R}(s, a) + \gamma Z(s', a')$. According to Dabney et al. (2017), the minimization of the 1-Wasserstein loss can be done by learning quantile locations for fractions $\tau_m = \frac{2m-1}{2M}, m \in [1..M]$ via quantile regression loss, defined for a quantile fraction $\tau \in [0, 1]$ as :

$$\mathcal{L}_{\text{QR}}^\tau(\theta) := \mathbb{E}_{\tilde{Z} \sim Z} [\rho_\tau(\tilde{Z} - \theta)]$$

with $\rho_\tau(u) = u(\tau - \mathbb{I}(u < 0)), \forall u \in \mathbb{R}$. Finally, to obtain better gradients when u is small, Huber quantile loss (or asymmetric Huber loss) can be used:

$$\rho_\tau^H(u) = |\tau - \mathbb{I}(u < 0)| \mathcal{L}_H^1(u),$$

where $\mathcal{L}_H^1(u)$ is a classical Huber loss with parameter 1. The quantile representation has the advantage of not fixing the support of the learned distribution and is used to represent the distribution of return in our algorithm for both discrete and continuous action space.

5. Algorithm

We use a distributional maximum entropy framework for continuous action space which is closed to the TQC algorithm (Kuznetsov et al., 2020). This method uses an actor-critic framework with a distributional truncated critic to avoid overestimation in the estimation with the max operator. This algorithm is based on a soft-policy iteration where we penalize the target $y_i(s, a)$ using the entropy of the distribution. More formally, to compute the target, the principle is to train N approximate estimate $Z_{\psi_1}, \dots, Z_{\psi_C}$ of the distribution of returns Z^π where Z_{ψ_c} maps each (s, a) to $Z_{\psi_c}(s, a) := \frac{1}{M} \sum_{m=1}^M \delta(\theta_{\psi_c}^m(s, a))$, which is supported on atoms $\theta_{\psi_c}^1(s, a), \dots, \theta_{\psi_c}^M(s, a)$. Then

approximations $Z_{\psi_1}, \dots, Z_{\psi_N}$ are trained on the temporal difference target distribution denoted $Y(s, a)$ constructed as follow. First atoms of trained distributions $Z_{\psi_1}(s', a'), \dots, Z_{\psi_C}(s', a')$ are pooled into $\mathcal{Z}(s', a') := \left\{ \theta_{\psi_c}^m(s', a') \mid c \in [1..C], m \in [1..M] \right\}$. We denote elements of $\mathcal{Z}(s', a')$ sorted in ascending order by $z_{(i)}(s', a')$, with $i \in [1..MC]$. Then we only keep the kC smallest elements of $\mathcal{Z}(s', a')$. We remove outliers of distribution to avoid overestimation of the value function. Finally the atoms of the target distribution $Y(s, a) := \frac{1}{kC} \sum_{i=1}^{kC} \delta(y_i(s, a))$ are computed according to a soft policy gradient method where we penalised with the log of the policy :

$$y_i(s, a) := r(s, a) + \gamma [z_{(i)}(s', a') - \eta \log \pi_\phi(a' | s')]. \quad (2)$$

The 1-Wasserstein distance between each of $Z_{\psi_n}(s, a), n \in [1..N]$ and the temporal difference target distribution $Y(s, a)$ is minimized learning the locations for quantile fractions $\tau_m = \frac{2m-1}{2M}, m \in [1..M]$. Similarly, we minimize the loss :

$$J_Z(\psi_c) = \mathbb{E}_{\mathcal{D}, \pi} \left[\frac{1}{MkC} \sum_{j=1}^M \sum_{i=1}^{kC} \rho_{\tau_j}^H(y_i(s, a) - \theta_{\psi_c}^j(s, a)) \right]$$

over the parameters ψ_n , for each critic. With this formulation, the learning of all quantiles $\theta_{\psi_n}^m(s, a)$ is dependent on all atoms of the truncated mixture of target distributions. To optimize the actor, the following loss based on KL-divergence denoted D_{KL} is used for soft policy improvement, :

$$J_{\pi, \alpha}(\phi) = \mathbb{E}_{\mathcal{D}} \left[D_{\text{KL}} \left(\pi_\phi(\cdot | s) \parallel \frac{\exp(\frac{1}{\eta} \xi_\alpha(\theta_\psi(s, \cdot)))}{D} \right) \right]$$

where η can be seen as a temperature and needs to be tuned and D is a constant of normalisation. This expression simplify into :

$$J_{\pi, \alpha}(\phi) = \mathbb{E}_{\mathcal{D}, \pi} \left[\eta \log \pi_\phi(a | s) - \frac{1}{C} \sum_{c=1}^C \xi_\alpha(\theta_{\psi_c}(s, a)) \right]$$

where $s \sim \mathcal{D}, a \sim \pi_\phi(\cdot | s)$. Nontruncated estimate of the Q-value are used for policy optimization to avoid a double truncation, in fact the Z -functions already approximate truncated future distribution. Finally, η is the entropy temperature coefficient and is dynamically adjusted by taking a gradient step with respect to the loss like in Haarnoja et al. (2018) :

$$J(\eta) = \mathbb{E}_{\mathcal{D}, \pi_\phi} [(-\log \pi_\phi(a_t | s_t) - \mathcal{H}_\eta)]$$

at every time the π_ϕ changes. Temperature η decreases if the policy entropy, $-\log \pi_\phi(a_t | s_t)$, is higher than \mathcal{H}_η and increases otherwise. The algorithm is summarized as follows : Our algorithm is based SAC framework but with many distributional critics to improve the estimation of Q -functions while using mean-standard deviation objective in the policy loss to improve robustness.

Algorithm 1 TQC with Standard Deviation penalisation

```

Initialize policy  $\pi_\phi$ , critics  $Z_{\psi_c}, Z_{\bar{\psi}_c}$  for  $c \in [1..C]$ 
for each iteration do
  for each step of the environment do
    collect  $(s_t, a_t, r_t, s_{t+1})$  with policy  $\pi_\phi$ 
     $\mathcal{D} \leftarrow \mathcal{D} \cup \{(s_t, a_t, r_t, s_{t+1})\}$ 
  end for
  for each gradient steps do
    Sample batch  $(s, a, s', r)$  of  $\mathcal{D}$ 
     $y_i(s, a) \leftarrow r(s, a) + \gamma [z_{(i)}(s', a') - \eta \log \pi_\phi(a' | s')]$ 
     $\eta \leftarrow \eta - \lambda_\eta \hat{\nabla}_\eta J(\eta)$ 
     $\phi \leftarrow \phi - \lambda_\pi \hat{\nabla}_\phi J_{\pi, \alpha}(\phi)$ 
     $\psi_c \leftarrow \psi_c - \lambda_Z \hat{\nabla}_{\psi_c} J_Z(\psi_c), c \in [1..C]$ 
     $\bar{\psi}_c \leftarrow \beta \psi_c + (1 - \beta) \bar{\psi}_c, c \in [1..C]$ 
  end for
end for
return policy  $\pi_\phi, Z_{\psi_c}, c \in [1 \dots C]$ .

```

6. Experiments

We try different experiments on continuous and discrete action space (see Annex B.2. for discrete case) to demonstrate the interest of our algorithms for robustness using $\xi : Z \rightarrow \mathbb{E}[Z] - \alpha \mathbb{V}[Z]^{\frac{1}{2}}$ instead of the mean. For continuous action space, we compare our algorithm with SAC which achieves state of the art in robust control (Eysenbach & Levine, 2021) on the Mujoco environment such as Hopper-v3, Walker-v3, or HalfCheetah-v3. We use a version where the entropy coefficient is adjusted during learning for both SAC and our algorithm as it requires less parameter tuning. Moreover, we show the influence of a distributional critic without a mean-standard deviation greedy step using $\alpha = 0$ to demonstrate the advantage of using a distributional critic against the classical SAC algorithm. We also compare our results to TQC algorithm which is in fact very close to SAC algorithm except the distributional critic. Finally, α the penalty is increased to show that for the tested environment, there exists a value of α such as prediction are more robust to change of dynamics. In these simulations, variations of dynamics are carried out by moving the relative mass which is an influential physical parameter in all environments. All algorithms are trained with a relative mass of 1 and then tested on new environment where the mass varies from 0.5 to 2. Two phenomena can be observed for the 3 environments.

In Fig 1, we see that we can find a value of α where the robustness is clearly improved without deteriorating the average performance. Normalized performance using by the maximum of the performance for every curve to highlight robustness and not only mean-performance can be found in annex B.1. If a too strong penalty is applied, the average

performance can be decreased as in the HalfCheetah-v3 environment (see annex B.1). For Hopper-v3, a α calibrated at 5 gives very good robustness performances while for Walker2d-v3, the value is closer to 2. This phenomenon was expected and in agreement with our formulation. Moreover, our algorithm outperforms the SAC algorithm for Robustness tasks in all environments. The tuning of α is discussed in annex B.1.

The second surprising observation is that penalizing our objective also improves performance in terms of stability during training and in terms of average performance, especially for Hopper-v3 and Walker2d-v3 in Fig 1. Similar results are present in the work of (Moskovitz et al., 2021) which gives an interpretation in terms of optimism and pessimism for the environments. This phenomenon is not yet explained, but it is present in some environments that are particularly unstable and have a lot of variance.

We observe similar results for discrete environments such as Cartpole-v1 and Acrobot-v1 (See annex B.2.).

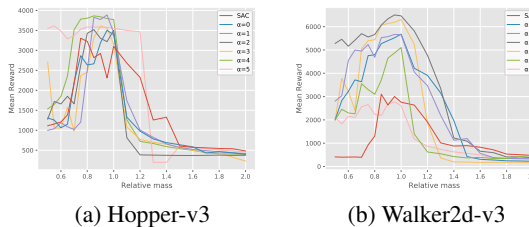


Figure 1: Mean over 20 trajectories varying relative mass of environments.

7. Conclusions

In this paper, we have tried to show that by using a mean-standard deviation formulation to choose our actions pessimistically, we can increase the robustness of our environment for continuous and discrete environments without adding complexity. A single fixed α parameter must be tuned to obtain good performance without penalizing the average performance too much. Moreover, for some environments, it is relevant to penalize to increase the average performance as well when there is lots of variability in the environment.

About the limitations of this work, the convergence of the algorithm to a fixed point is not shown for mean-standard deviation penalization and this formulation is still based on a heuristic even if the link between robustness and penalization is established. Theoretical link with a (Kumar et al., 2022) could be an interesting way of analyzing our algorithm as they use action value to penalize their objective and remains a way to explain this phenomenon. This is left for future work.

References

- Abdullah, M. A., Ren, H., Ammar, H. B., Milenkovic, V., Luo, R., Zhang, M., and Wang, J. Wasserstein robust reinforcement learning. *arXiv preprint arXiv:1907.13196*, 2019.
- Auer, P. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002.
- Bai, C., Wang, L., Yang, Z., Deng, Z., Garg, A., Liu, P., and Wang, Z. Pessimistic bootstrapping for uncertainty-driven offline reinforcement learning. *arXiv preprint arXiv:2202.11566*, 2022.
- Behzadian, B., Petrik, M., and Ho, C. P. Fast algorithms for l_∞ -constrained s-rectangular robust mdps. *Advances in Neural Information Processing Systems*, 34, 2021.
- Bellemare, M. G., Dabney, W., and Munos, R. A distributional perspective on reinforcement learning. *34th International Conference on Machine Learning, ICML 2017*, 1:693–711, 7 2017. URL <https://arxiv.org/abs/1707.06887v1>.
- Brekelmans, R., Genewein, T., Grau-Moya, J., Delétang, G., Kunesch, M., Legg, S., and Ortega, P. Your policy regularizer is secretly an adversary. *arXiv preprint arXiv:2203.12592*, 2022.
- Dabney, W., Rowland, M., Bellemare, M. G., and Munos, R. Distributional reinforcement learning with quantile regression. *32nd AAAI Conference on Artificial Intelligence, AAAI 2018*, pp. 2892–2901, 10 2017. URL <https://arxiv.org/abs/1710.10044v1>.
- Dabney, W., Ostrovski, G., Silver, D., and Munos, R. Implicit quantile networks for distributional reinforcement learning. *35th International Conference on Machine Learning, ICML 2018*, 3:1774–1787, 6 2018. URL <https://arxiv.org/abs/1806.06923v1>.
- Derman, E. and Mannor, S. Distributional robustness and regularization in reinforcement learning. *arXiv preprint arXiv:2003.02894*, 2020.
- Derman, E., Geist, M., and Mannor, S. Twice regularized mdps and the equivalence between robustness and regularization. *Advances in Neural Information Processing Systems*, 34, 2021.
- Duchi, J. and Namkoong, H. Learning models with uniform performance via distributionally robust optimization. *arXiv preprint arXiv:1810.08750*, 2018.
- Duchi, J., Glynn, P., and Namkoong, H. Statistics of robust optimization: A generalized empirical likelihood approach. *arXiv preprint arXiv:1610.03425*, 2016.
- Eysenbach, B. and Levine, S. Maximum entropy rl (provably) solves some robust rl problems. *arXiv preprint arXiv:2103.06257*, 2021.
- Geist, M., Scherrer, B., and Pietquin, O. A theory of regularized markov decision processes. In *International Conference on Machine Learning*, pp. 2160–2169. PMLR, 2019.
- Gotoh, J.-y., Kim, M. J., and Lim, A. E. Robust empirical optimization is almost the same as mean-variance optimization. *Operations research letters*, 46(4):448–452, 2018.
- Grand-Clément, J. and Kroer, C. First-order methods for wasserstein distributionally robust mdp. *arXiv preprint arXiv:2009.06790*, 2020a.
- Grand-Clément, J. and Kroer, C. Scalable first-order methods for robust mdps. *arXiv preprint arXiv:2005.05434*, 2020b.
- Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *35th International Conference on Machine Learning, ICML 2018*, 5:2976–2989, 1 2018. URL <http://arxiv.org/abs/1801.01290>.
- Husain, H., Ciosek, K., and Tomioka, R. Regularized policies are reward robust. In *International Conference on Artificial Intelligence and Statistics*, pp. 64–72. PMLR, 2021.
- Jaimungal, S., Pesenti, S. M., Wang, Y. S., and Tatsat, H. Robust risk-aware reinforcement learning. *SSRN Electronic Journal*, 8 2021. doi: 10.2139/SSRN.3910498. URL <https://papers.ssrn.com/abstract=3910498>.
- Jain, A., Patil, G., Jain, A., Khetarpal, K., and Precup, D. Variance penalized on-policy and off-policy actor-critic. 2021a. URL www.aaai.org.
- Jain, A., Patil, G., Jain, A., Khetarpal, K., and Precup, D. Variance penalized on-policy and off-policy actor-critic. *arXiv preprint arXiv:2102.01985*, 2021b.
- Kumar, N., Levy, K., Wang, K., and Mannor, S. Efficient policy iteration for robust markov decision processes via regularization. *arXiv preprint arXiv:2205.14327*, 2022.
- Kuznetsov, A., Shvechikov, P., Grishin, A., and Vetrov, D. Controlling overestimation bias with truncated mixture of continuous distributional quantile critics. In *International Conference on Machine Learning*, pp. 5556–5566. PMLR, 2020.

- Ma, Y. J., Jayaraman, D., and Bastani, O. Conservative offline distributional reinforcement learning. 7 2021. URL <https://arxiv.org/abs/2107.06106v2>.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. Playing atari with deep reinforcement learning. 12 2013. URL <https://arxiv.org/abs/1312.5602v1>.
- Morimoto, J. and Doya, K. Robust reinforcement learning. *Neural computation*, 17(2):335–359, 2005.
- Moskovitz, T., Parker-Holder, J., Pacchiano, A., Arbel, M., and Jordan, M. Tactical optimism and pessimism for deep reinforcement learning. *Advances in Neural Information Processing Systems*, 34, 2021.
- Nam, D. W., Kim, Y., and Park, C. Y. Gmac: A distributional perspective on actor-critic framework. In *International Conference on Machine Learning*, pp. 7927–7936. PMLR, 2021.
- Petrik, M. and Russel, R. H. Beyond confidence regions: Tight bayesian ambiguity sets for robust mdps. *Advances in neural information processing systems*, 32, 2019.
- Raffin, A., Hill, A., Ernestus, M., Gleave, A., Kanervisto, A., and Dormann, N. Stable baselines3, 2019.
- Scherrer, B., Ghavamzadeh, M., Gabillon, V., Lesner, B., and Geist, M. Approximate modified policy iteration and its application to the game of tetris. *J. Mach. Learn. Res.*, 16(49):1629–1676, 2015.
- Schulman, J., Levine, S., Moritz, P., Jordan, M. I., and Abbeel, P. Trust region policy optimization. *32nd International Conference on Machine Learning, ICML 2015*, 3:1889–1897, 2 2015. URL <http://arxiv.org/abs/1502.05477>.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv*, 7 2017a. ISSN 23318422. URL <https://arxiv.org/abs/1707.06347v2>. PPO algorithm premier papier
Important à citer

.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017b.
- Singh, R., Zhang, Q., and Chen, Y. Improving robustness via risk averse distributional reinforcement learning. In *Learning for Dynamics and Control*, pp. 958–968. PMLR, 2020.
- Smirnova, E., Dohmatob, E., and Mary, J. Distributionally robust reinforcement learning. 2 2019. URL <https://arxiv.org/abs/1902.08708v2>.
- Tessler, C., Efroni, Y., and Mannor, S. Action robust reinforcement learning and applications in continuous control. In *International Conference on Machine Learning*, pp. 6215–6224. PMLR, 2019.
- Urpí, N. A., Curi, S., and Krause, A. Risk-averse offline reinforcement learning. *arXiv preprint arXiv:2102.05371*, 2021.
- Vieillard, N., Kozuno, T., Scherrer, B., Pietquin, O., Munos, R., and Geist, M. Leverage the average: an analysis of kl regularization in reinforcement learning. *Advances in Neural Information Processing Systems*, 33:12163–12174, 2020.
- Wang, H. and Zhou, X. Y. Continuous-time mean–variance portfolio selection: A reinforcement learning framework. *Mathematical Finance*, 30(4):1273–1308, 2020.
- Wiesemann, W., Kuhn, D., and Rustem, B. Robust markov decision processes. *Mathematics of Operations Research*, 38(1):153–183, 2013.
- Yang, I. A convex optimization approach to distributionally robust markov decision processes with wasserstein distance. *IEEE Control Systems Letters*, 1:164–169, 2017. ISSN 24751456. doi: 10.1109/LCSYS.2017.2711553.
- Ying, C., Zhou, X., Su, H., Yan, D., and Zhu, J. Towards safe reinforcement learning via constraining conditional value-at-risk. 2021.
- Zhang, J. and Weng, P. Safe distributional reinforcement learning.
- Zhang, J. and Weng, P. Safe distributional reinforcement learning. In *International Conference on Distributed Artificial Intelligence*, pp. 107–128. Springer, 2021.

A. Proof of (1)

We consider the following equality :

$$\min_{P \in D_{\chi^2}(P||P_0) \leq \alpha} Q^{(P,\pi)} = Q^{(P_0,\pi)} - \alpha \mathbb{V}_{P_0}[Z]^{\frac{1}{2}}. \quad (3)$$

Consider that trajectories τ is drawn from \mathbb{P} but here we will write P the transition of the environment as the policy π is fixed and it is the only part which differ.

Writing $\tilde{R}(\tau) = R(\tau) - \mathbb{E}_{\tau \sim P_0}[R(\tau)]$ we get :

$$\begin{aligned} \|\mathbb{E}_{\tau \sim P}[R(\tau)] - \mathbb{E}_{\tau \sim P_0}[R(\tau)]\| &= \left\| \int_{\tau} \tilde{R}(\tau) (p(\tau) - p_0(\tau)) d\tau \right\| \\ &= \left\| \int_{\tau} \tilde{R}(\tau) \sqrt{p_0(\tau)} \frac{(p(\tau) - p_0(\tau))}{\sqrt{p_0(\tau)}} d\tau \right\| \\ &\leq \left\| \int_{\tau} \tilde{R}(\tau)^2 p_0(\tau) d\tau \right\|^{\frac{1}{2}} \left\| \int_{\tau} \frac{(p(\tau) - p_0(\tau))^2}{p_0(\tau)} d\tau \right\|^{\frac{1}{2}} \\ &= \mathbb{V}_{P_0}[R(\tau)]^{\frac{1}{2}} D_{\chi^2}(P||P_0)^{\frac{1}{2}}, \end{aligned}$$

because of the positivity of divergence and of the variance, norms are removed. This inequality comes from Cauchy-Schwarz inequality and becomes an equality if for $\lambda \in \mathbb{R}$:

$$\tilde{R}(\tau) p_0(\tau) = \lambda (p(\tau) - p_0(\tau)) \iff p(\tau) = p_0(\tau) \left(1 + \frac{1}{\lambda} \tilde{R}(\tau)\right). \quad (4)$$

However, $p(\tau)$ needs to be non-negative and sum to one as it is a measure. Normalisation condition is respected by construction however to ensure that the measure is non-negative, this requires $\left\| \tilde{R}(\tau)/\lambda \right\| \leq 1$ in the case where $\lambda \leq 0$. In this case of equality, we obtain from 4 that $D_{\chi^2}(P||P_0) = \frac{\mathbb{V}_{P_0}[R(\tau)]}{\lambda^2}$. Replacing the divergence in the inequality, the following result holds :

$$\|\mathbb{E}_{\tau \sim P}[R(\tau)] - \mathbb{E}_{\tau \sim P_0}[R(\tau)]\| \leq \frac{\mathbb{V}_{P_0}[R(\tau)]}{\lambda}.$$

For proving (3) we are interested in the case where $D_{\chi^2}(P||P_0) \leq \alpha$, from the initial inequality we obtain :

$$\min_{P \in D_{\chi^2}(P||P_0) \leq \alpha} Q^{(P,\pi)} \geq \min_{P \in D_{\chi^2}(P||P_0) \leq \alpha} Q^{(P_0,\pi)} - D_{\chi^2}(P||P_0) \mathbb{V}_{P_0}[Z]^{\frac{1}{2}} = Q^{(P_0,\pi)} - \alpha \mathbb{V}_{P_0}[Z]^{\frac{1}{2}}$$

with the maximum value of α equals to $D_{\chi^2}(P||P_0) = \frac{\mathbb{V}_{P_0}[R(\tau)]}{\lambda^2} \leq \frac{\mathbb{V}_{P_0}[R(\tau)]}{\|\tilde{R}\|_{\infty}^2} = \frac{\|\tilde{R}\|_2^2}{\|\tilde{R}\|_{\infty}^2} \leq 1$, where the first inequality comes from the conditions $\left\| \tilde{R}(\tau)/\lambda \right\| \leq 1$ and the last one comes from that the L_2 norm is smaller than ∞ -norm.

If our problem is constrained, assuming $\alpha \leq \frac{\mathbb{V}_{P_0}[R(\tau)]}{\|\tilde{R}\|_{\infty}^2} \leq 1$, we obtain the following results with the maximum attained for $D_{\chi^2}(P||P_0) = \alpha$:

$$\min_{P \in D_{\chi^2}(P||P_0) \leq \alpha} Q^{(P,\pi)} = Q^{(P_0,\pi)} - \alpha \mathbb{V}[Z]^{\frac{1}{2}}. \quad (5)$$

For $\alpha > 1$, we still optimize a lower bound of the quantity of interest. The formulation of our algorithm becomes:

$$\begin{cases} \pi_{k+1} \in \mathcal{G}_\alpha(Z_k) = \mathcal{G}(\xi_\alpha(Z_k) = \arg \max_{\pi \in \Pi} \langle \mathbb{E}[Z_k] - \alpha \sqrt{\mathbb{V}[Z_k]}, \pi \rangle) \\ Z_{k+1} = (T^{\pi_{k+1}})^m Z_k \end{cases}$$

B. Further Experimental Details

All experiments were run on a cluster containing an Intel Xeon CPU Gold 6230, 20 cores and every single experiment was performed on a single CPU between 3 and 6 hours for continuous control and less than 1 hour for discrete control environment.

Pre-trained models will be available for all algorithms and environments on a GitHub link.

The Mujoco OpenAI Gym tasks licensing information is given at <https://github.com/openai/gym/blob/master/LICENSE.md>. Baseline implementation of PPO, SAC, TQC and QRDQN can be found in [Raffin et al. \(2019\)](#). Moreover, hyperparameters across all experiments used are displayed in Table 2, 1 and 3.

B.1. Results for continuous action spaces

Tuning of α must be chosen carefully, for example, α is chosen in $\{0, 1, \dots, 5\}$ for Hopper-v3 and Walker2d-v3 whereas values of α are chosen smaller in $\{0, 0.1, 0.5, 1, 1.5, 2\}$ and not in a bigger interval. As a rule of thumb for choosing α , we can look at the empirical mean and variance at the end of the trajectories to see if the environment has rewards that fluctuate a lot. The smaller the mean/variance ratio, the more likely we are to penalise our environment. For HalfCheetah-v3, the mean/variance ratio is about approximately 100, so we will favour smaller penalties than for Walker2d where the mean/variance ratio is about 50 or 10 for Hopper-v3.

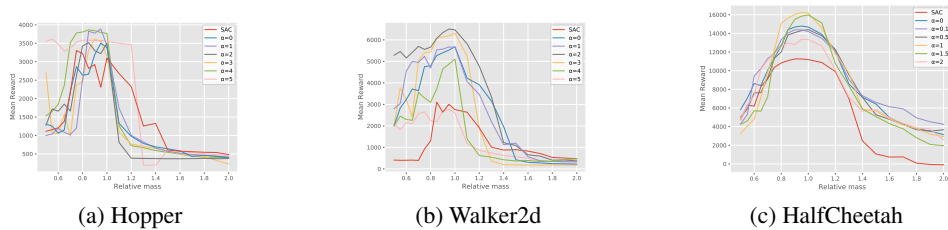


Figure 2: Mean results on 20 trajectories varying relative mass.

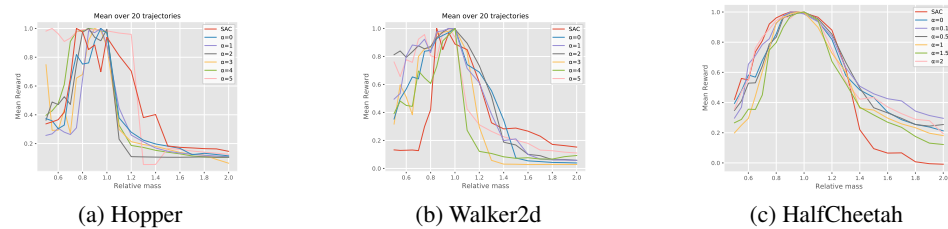


Figure 3: Normalised mean results on 20 trajectories varying relative mass.

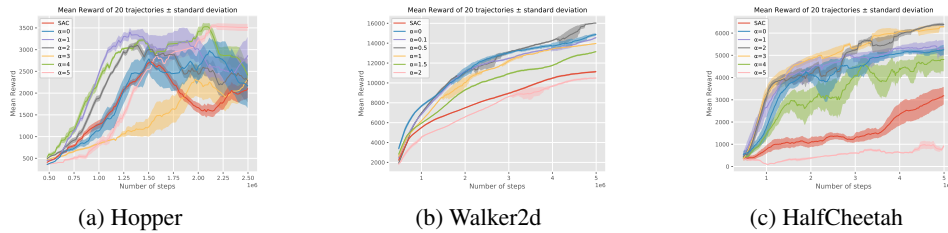


Figure 4: Mean performances over 20 trajectories \pm variance.

B.2. Results on discrete action spaces

We test our algorithm QRDQN with standard deviation penalization on discrete action space, varying the length of the pole in Cartpole-v1 and Acrobot-v1 environments. We observe similar results for the discrete environment in terms of robustness. Training is done for a length of the pole equal to the x-axis of the black star on the graph, then for testing, the length of the pole is increased or decreased. We show that robustness is increased when we penalised our distributional critic. We have compared our algorithm to PPO which has shown relatively good results in terms of robustness for discrete action space in (Abdullah et al., 2019) as SAC does not apply to discrete action space. The same phenomenon is observed in terms of robustness as for the continuous environments. However, the more surprising observations on Hopper and Walker2d with an improvement in performance on average is to be qualified. This is partly explained by the fact that the maximum reward is reached in Cartpole and Acrobot quickly. An ablation study can be found in annex C where we study the impact of penalization on our behavior policy during testing and on the policy used during learning. It is shown that both are needed in the algorithm.

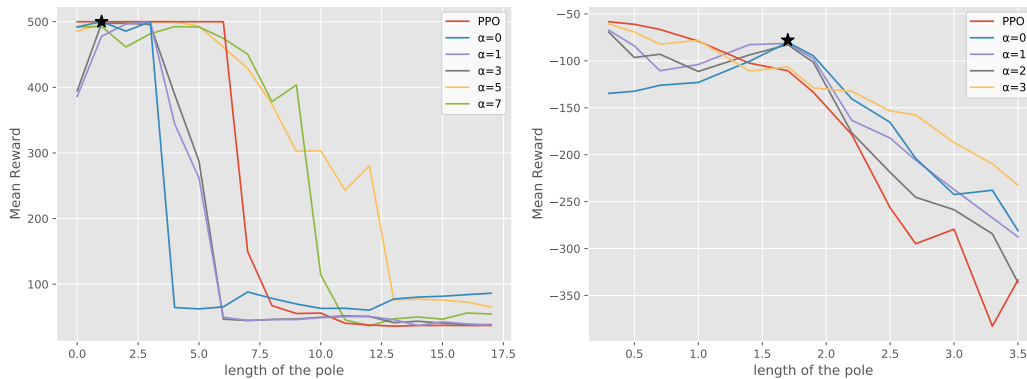


Figure 5: Mean over 20 trajectories varying length of the pole trained on the x-axis of the black star for Cartpole-v1 and Acrobot-v1 environments

B.3. Ablation study

The purpose of this ablation study is to look at the influence of penalization in the discrete action space with QRDQN. In the figures below, we look at the influence of penalizing only during training, which will have the effect of choosing less risky actions during training in order to increase robustness. This curve is denoted *Train penalized*.

Then we look at the influence of penalizing only once the policy has been learned using classic QRDQN without penalization. Only mean-var actions are selected here during testing and not during training. This experience is denoted *Train Penalization*.

Finally, we compare its variants with our algorithm called *Full penalization*. The results of the ablation are: to achieve optimal performance, both phases are necessary.

When penalties are applied only during training. Good performance is obtained in general close to the length 1 where we train our algorithm. However, the performance is difficult to generalize when the pole length is increased as we do not

penalize during testing.

When we penalize only during testing: even if the performances deteriorate, we see that it tends to add robustness because the curves have less tendency to decrease when we increase the length of the pole. The performances are not very high as we play different acts than those taken during the learning.

So both phases are therefore necessary for our algorithm. Penalizing during training allows for safer exploration and penalizing during testing allows for better generalization.

The ablation study for the continuous case is more difficult to do. Indeed, the fact that the penalty occurs only in the gradient descent phase makes it difficult to penalize only in the test phase.

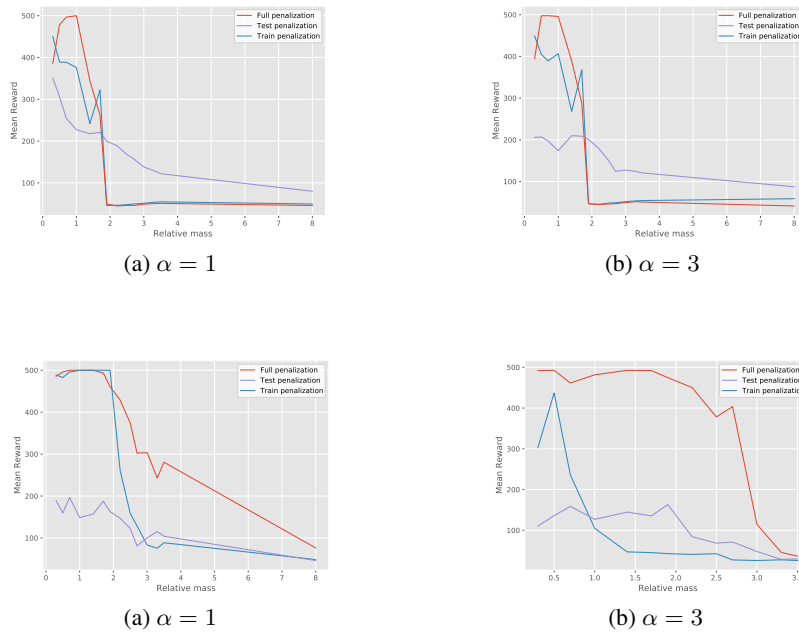


Figure 6: Ablation study for different values of α

C. Hyperparameters

For HalfCheetah-v3 , penalisation is chosen in $[0, 2]$ and not $[0, 5]$ like in Walker-v3 and Hopper-v3.

Table 1: Table of best hyperparameter for Cartpole-v1

Hyperparameter	QRDQN with standard deviation penalisation	PPO
Learning Rate	2.3e-3	3e-4
Optimizer	Adam	Adam
Replay Buffer Size	10e5	N/A
Number of Quantiles	10	N/A
Huber parameter κ	1	N/A
Penalisation α	{0,1,3,5,7 }	N/A
Network Hidden Layers for Policy	N/A	256:256
Network Hidden Layers for Critic	256:256	256:256
Number of samples per Minibatch	64	256
Discount factor γ	0.99	0.99
Target smoothing coefficient β	.0005	N/A
Non-linearity	ReLu	ReLu
Target update interval	10	N/A
Gradient steps per iteration	1	1
Entropy coefficient	N/A	0
GAE λ	0.95	0.8

Table 2: Table of best hyperparameter for Acrobot-v1

Hyperparameter	QRDQN with standard deviation penalisation	PPO
Learning Rate	6.3e-4	3e-4
Optimizer	Adam	Adam
Replay Buffer Size	50 000	N/A
Number of Quantiles	25	N/A
Huber parameter κ	1	N/A
Penalisation α	{0, 0.5, 1, 2, 3}	N/A
Network Hidden Layers for Critic	256:256	256:256
Network Hidden Layers for Policy	N/A	256:256
Number of samples per Minibatch	128	64
Discount factor γ	0.99	0.99
Target smoothing coefficient β	.0005	N/A
Non-linearity	ReLu	ReLu
Target update interval	250	N/A
Gradient steps per iteration	4	1
Entropy coefficient	N/A	0
GAE λ	0.95	0.95

Table 3: Table of best hyperparameter for all continuous environments

Hyperparameter	TQC with standard deviation penalisation	SAC
Learning Rate	linear decay from $7.3e-4$	linear decay from $7.3e-4$
Optimizer	Adam	Adam
Replay Buffer Size	10^6	10^6
Expected Entropy Target	$-\dim \mathcal{A}$	$-\dim \mathcal{A}$
Number of Quantiles	25	N/A
Huber parameter κ	1	N/A
Penalisation α	$\{0, 1, \dots, 5\}$	N/A
Network Hidden Layers for Policy	256:256	256:256
Network Hidden Layers for Critic	512:512:512	256:256
Number of dropped atoms	2	N/A
Number of samples per Minibatch	256	256
Discount factor γ	0.99	0.99
Target smoothing coefficient β	.0005	0.005
Non-linearity	ReLu	ReLu
Target update interval	1	1
Gradient steps per iteration	1	1