Dealing with the Unknown: Pessimistic Offline Reinforcement Learning

Anonymous Author(s)
Affiliation
Address
email

Abstract: Reinforcement Learning (RL) has been shown effective in domains where the agent can learn policies by actively interacting with its operating environment. However, if we change the RL scheme to offline setting where the agent can only update its policy via static datasets, one of the major issues in offline reinforcement learning emerges, i.e. distribution shift. We propose a Pessimistic Offline Reinforcement Learning (PORL) algorithm to actively lead the agent back to the area where it is familiar by manipulating the value function. We focus on problems caused by out-of-distribution (OOD) states, and deliberately penalize high values at states that are absent in the training dataset, so that the learned pessimistic value function lower bounds the true value anywhere within the state space. We evaluate the PORL algorithm on various benchmark tasks, where we show that our method gains better performance by explicitly handling OOD states, when compared to those methods merely considering OOD actions.

Keywords: Offline Reinforcement Learning, Out-of-Distribution States

1 Introduction

Reinforcement learning (RL), especially with high-capacity models such as deep nets, has shown its power in many domains, e.g. gaming, healthcare, and robotics. However, typical training schemes of RL algorithms rely on active interaction with the environments. It limits their applications in domains where active data collection is expensive or dangerous (e.g., autonomous driving). Recently, offline reinforcement learning (offline RL) has emerged as a promising candidate to overcome this barrier. Different from traditional RL methods, offline-RL learns the policy from a static offline dataset collected without iterative interaction with the environment. Recent works have shown its ability in solving various policy learning tasks [1, 2, 3]. However, offline RL methods suffer from several major problems. One of them is distribution shift. Unlike online RL algorithms, the state and action distributions are different during training and testing. As a result, RL agents may fail dramatically after deployed online.

Many prior works [4, 5, 6, 7, 8, 9] try to mitigate this problem by handling out-of-distribution (OOD) actions. They discourage the policies to visit OOD actions by designing conservative value functions, or estimating the uncertainty of Q-functions. Although constraining the policy can implicitly mitigate the problem of state distribution shift, few works have adopted measures to explicitly handle OOD states during the training stage. In this work, we propose the Pessimistic Offline Reinforcement Learning (PORL) framework to explicitly limit the policy from visiting both unseen states and actions.

Our PORL framework is inspired by the concept of pessimistic MDP in [10], where the reward is significantly small for unseen state-action pairs. We aim to limit the magnitude of the value function at unseen states, so that the agent can avoid or recover from unseen states. It is then crucial to precisely detect OOD states and shape the value function at those states. Since prior methods on OOD actions are derived from a similar concept, we can adapt their approaches to handle OOD states. There are mainly two approaches in the literature. One is to estimate the epistemic uncertainty of Q-function and subtract it from the original Q-function to get a conservative Q-function [5, 6, 7, 8, 9]. The other is to regularize the Q-function during the learning process [4]. The first method is highly

Submitted to the 5th Conference on Robot Learning (CoRL 2021). Do not distribute.

sensitive to the trade-off between the uncertainty estimation and the original Q-function [11, 12] and the quality of uncertainty estimation [13].

Therefore, we follow the second approach, and add a conservative regularization term to the policy evaluation step of PORL to shape the value function. We prove that PORL learns a pessimistic value function that lower bounds the true value function. Moreover, the pessimistic value function is corresponding to a pessimistic Markov Decision Process (MDP), which forces the policy to avoid or recover from out-of-distribution states and actions. We evaluate the PORL algorithm on various benchmark tasks. The performance of our method matches the state-of-the-art offline RL methods. In particular, we show that, by explicitly handling OOD states, we can further improve the policy performance compared to those methods merely considering OOD actions.

2 Related Works

A big challenge for offline reinforcement learning methods is to deal with the problems caused by unvisited states or actions in training data, which is also known as distribution shift. In model-free offline reinforcement learning, some works used importance sampling to fill the gap between the learned policy and the behavior policy in the training dataset [14, 15, 16, 17, 18]. There are also many works constrained the learned policy to be similar to the behavior policy by explicit constraints in the training dataset [5, 19, 20, 21, 22], so that the agent can avoid out-of-distribution actions during test time. The work in [12] proposed a latent space to constrain the policy to avoid deviating from the training data support. One further step to make the agent avoid actions that may cause itself deviate from the training data support is to get a conservative value function and thus a conservative policy. The works in [6, 7, 5, 8, 9] estimate the uncertainty of the learned Q-function, and then directly subtract it from the Q-function to get a conservative Q-function. Another way to get a conservative Q-function is to regularize the Q-function in the optimization problem during the learning process [4]. In model-based reinforcement learning (MBRL), there are also many algorithms that constrain the exploitation in the environment with effective uncertainty estimation methods [23, 24, 25, 26, 27, 10]. It is considered to be mature and reliable to detect OOD actions and states by methods from MBRL. Most of the aforementioned methods focus on OOD actions but not have explicit mechanism to deal with OOD states. In this paper, we focus on OOD states and propose a method to learn a pessimistic value function by adding regularization terms when updating Q-functions, and follow the works in the MBRL domain to establish the module to detect OOD states in our algorithm.

3 Background

3.1 Offline Reinforcement Learning

Given a Markov decision process (MDP), an RL agent aims to maximize the expectation of cumulative rewards. The MDP is represented by a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma)$, where $\mathcal{S}$ is the state space, $\mathcal{A}$ is the action space, $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to [0, 1]$ is the transition function, $r : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ is the reward function, and γ is the discount factor. Typical RL algorithms optimize the policy using experience collected when interacting with the environment. Unlike those online learning paradigms, offline-RL algorithms rely solely on a static offline dataset, denoted by $\mathcal{D} = \left\{ \left(\mathbf{s}_t^i, \mathbf{a}_t^i, \mathbf{s}_{t+1}^i, r_t^i \right) \right\}$.

In this work, we focus on dynamic-programming-based RL algorithms under the offline setting, where we try to extract a policy from a learned value function for the underlying MDP in the training data. Standard Q-learning method estimates an approximate Q-function parametrized by θ, i.e. $\hat{Q}_\theta(\mathbf{s}_t, \mathbf{a}_t)$. In each iteration, the Q-function is updated as follows:

$$\hat{Q}_\theta^{k+1} \leftarrow \arg\min_Q \frac{1}{2} \mathbb{E}_{\mathbf{s},\mathbf{a},\mathbf{s}' \sim \mathcal{D}} \left[\left(Q(\mathbf{s}, \mathbf{a}) - \hat{\mathcal{B}}^\pi \hat{Q}_\theta^k(\mathbf{s}, \mathbf{a}) \right)^2 \right] \text{ (policy evaluation)}, \tag{1}$$

where $\hat{\mathcal{B}}^\pi$ is the empirical Bellman update operator defined as:

$$\hat{\mathcal{B}}^\pi \hat{Q}_\theta^k(\mathbf{s}, \mathbf{a}) = r(\mathbf{s}, \mathbf{a}) + \gamma \mathbb{E}_{\mathbf{a}' \sim \hat{\pi}^k(\mathbf{a}'|\mathbf{s}')} \left[\hat{Q}_\theta^k(\mathbf{s}', \mathbf{a}') \right], \tag{2}$$

where $\hat{\pi}^k(\mathbf{a}|\mathbf{s}) = \delta \left[\mathbf{a} - \arg\max_{\mathbf{a}} \hat{Q}_\theta^k(\mathbf{s}, \mathbf{a}) \right]$. In the offline setting, the empirical Bellman operator $\hat{\mathcal{B}}^\pi$ only back up one sample when updating the Q-function. In the rest of the paper, we denote $\mathcal{E}(Q, \hat{\mathcal{B}}^\pi \hat{Q}_\theta^k) = \frac{1}{2} \mathbb{E}_{\mathbf{s},\mathbf{a},\mathbf{s}' \sim \mathcal{D}} \left[\left(Q(\mathbf{s}, \mathbf{a}) - \hat{\mathcal{B}}^\pi \hat{Q}_\theta^k(\mathbf{s}, \mathbf{a}) \right)^2 \right]$ as the Bellman update error for simplicity.

Different from standard Q-learning, actor-critic algorithms learn a policy network to model $\hat{\pi}$, which is updated after policy evaluation in each iteration:

$$\hat{\pi}^{k+1} \leftarrow \arg\max_{\pi} \mathbb{E}_{\mathbf{s}\sim\mathcal{D},\mathbf{a}\sim\pi(\mathbf{a}|\mathbf{s})}\left[\hat{Q}_{\theta}^{k+1}(\mathbf{s},\mathbf{a})\right] \text{ (policy improvement).} \tag{3}$$

3.2 Uncertainty-Based Methods and Pessimistic Value Functions

By observing Eqn. 1, it is obvious that the Q-function $\hat{Q}_\theta$ is never evaluated or updated at states or actions that never appear in the dataset. The agent may behave unexpectedly or unpredictably at those unseen states or actions during test time. For dynamic-programming-based approaches, one way to address the issue of unseen actions is to estimate the epistemic uncertainty of Q-function and subtract it from the original Q-function [5, 6, 7, 8, 9]. The uncertainty is estimated based on an ensemble of learned Q-functions, and the final conservative Q-function becomes $Q_{\mathrm{c}}(\mathbf{s},\mathbf{a}) = \mathbb{E}_{Q\sim\mathcal{P}_{\mathcal{D}}(Q)}[Q(\mathbf{s},\mathbf{a}) - \alpha\mathrm{Unc}(\mathcal{P}_{\mathcal{D}}(Q))]$, where Unc is defined to be some uncertainty estimation metric, and $\mathcal{P}_{\mathcal{D}}(Q)$ is the distribution over possible Q-functions. Because the uncertainty metric is directly subtracted, uncertainty-based methods is highly sensitive to the quality of uncertainty estimation. Meanwhile, it is difficult to find an ideal α to balance the original Q-function and Unc.

Another way is to regularize the Q-function at the step of policy evaluation. A representative example is Conservative Q-Learning (CQL) [4]. Assuming that the dataset $\mathcal{D}$ is collected with a behavior policy $\pi_\beta(\mathbf{a}|\mathbf{s})$, and μ matches the state-marginal in the dataset, the policy evaluation step in CQL becomes:

$$\hat{Q}^{k+1} \leftarrow \arg\min_{Q} \alpha\left(\mathbb{E}_{\mathbf{s}\sim\mathcal{D},\mathbf{a}\sim\mu(\mathbf{a}|\mathbf{s})}[Q(\mathbf{s},\mathbf{a})] - \mathbb{E}_{\mathbf{s}\sim\mathcal{D},\mathbf{a}\sim\pi_\beta(\mathbf{a}|\mathbf{s})}[Q(\mathbf{s},\mathbf{a})]\right) + \mathcal{E}(Q, \hat{\mathcal{B}}^{\pi}\hat{Q}_{\theta}^{k}). \tag{4}$$

It is worth noting that the aforementioned methods all focus on OOD actions, but they do not have an explicit mechanism to deal with OOD states, which motivates us to develop the PORL framework in this work.

4 Pessimistic Offline Reinforcement Learning Framework

In this section, we introduce the PORL framework to mitigate the issue of state distribution shift. In particular, we propose a novel conservative regularization term in the policy evaluation step. It can then be integrated into Q-learning or actor-critic algorithm, which will be described in Sec. 5.

4.1 How To Deal With OOD States

Assuming the dataset $\mathcal{D}$ is collected with a behavior policy $\pi_\beta(\mathbf{a}|\mathbf{s})$, and the states $\mathbf{s}$ are distributed according to a distribution $d^{\pi_\beta}(\mathbf{s})$ in the dataset, we propose to solve the problem caused by state distribution shift by augmenting the policy evaluation step in CQL [4] with a regularization term scaled by a trade-off factor ε:

$$\min_{Q} \varepsilon\left(\mathbb{E}_{\mathbf{s}\sim d^{\phi}(\mathbf{s}),\mathbf{a}\sim\hat{\pi}^{k}(\mathbf{a}|\mathbf{s})}[Q(\mathbf{s},\mathbf{a})] - \mathbb{E}_{\mathbf{s}\sim d^{\pi_\beta}(\mathbf{s}),\mathbf{a}\sim\hat{\pi}^{k}(\mathbf{a}|\mathbf{s})}[Q(\mathbf{s},\mathbf{a})]\right) + \mathcal{E}(Q, \hat{\mathcal{B}}^{\pi}\hat{Q}_{\theta}^{k}) + \mathcal{C}(Q), \tag{5}$$

where $d^{\phi}(\mathbf{s})$ is a particular state distribution of our choice, and $\mathcal{C}(Q)$ is the cost term adopted from CQL, i.e. $\mathcal{C}(Q) = \alpha\left(\mathbb{E}_{\mathbf{s}\sim\mathcal{D},\mathbf{a}\sim\mu(\mathbf{a}|\mathbf{s})}[Q(\mathbf{s},\mathbf{a})] - \mathbb{E}_{\mathbf{s}\sim\mathcal{D},\mathbf{a}\sim\pi_\beta(\mathbf{a}|\mathbf{s})}[Q(\mathbf{s},\mathbf{a})]\right)$.

The idea is to use the minimization term $\varepsilon\left(\mathbb{E}_{\mathbf{s}\sim d^{\phi}(\mathbf{s}),\mathbf{a}\sim\hat{\pi}^{k}(\mathbf{a}|\mathbf{s})}[Q(\mathbf{s},\mathbf{a})]\right)$ to penalize high values at unseen states in the dataset, and the maximization term $\varepsilon\left(\mathbb{E}_{\mathbf{s}\sim d^{\pi_\beta}(\mathbf{s}),\mathbf{a}\sim\hat{\pi}^{k}(\mathbf{a}|\mathbf{s})}[Q(\mathbf{s},\mathbf{a})]\right)$ to cancel the penalization at in-distribution states. The regularized Q-function could then push the agent towards regions close to the states from the dataset, where the values are higher. To achieve it, we need to find a distribution $d^{\phi}(\mathbf{s})$ that assigns high probabilities to states far away from the dataset, and low probabilities to states near the training dataset.

For example, we can define $d^{\phi}(\mathbf{s})$ as a state distribution that maximizes the KL divergence between itself and the state marginal $d^{\pi_\beta}(\mathbf{s})$ in the training dataset, i.e. $D_{\mathrm{KL}}(d^{\phi}(\mathbf{s}) \,||\, d^{\pi_\beta}(\mathbf{s}))$. However, in general, it is empirically impossible to derive a solution to maximize the KL divergence. Therefore, we need to construct a well-defined distribution $d^{\phi}(\mathbf{s})$. We will instantiate a practical design of $d^{\phi}(\mathbf{s})$ in Sec. 5. For now, we just assume $d^{\phi}(\mathbf{s})$ assigns low probabilities to OOD states.

4.2 Theoretical Analysis

In this section, we analyze the theoretical properties of the proposed policy evaluation step. The proof and more details can be found in Appendix A.

We define $k \in \mathbb{N}$ as the iteration of policy evaluation, i.e. $\hat{Q}^k$ denotes the optimized Q-function in the $k-$th iteration obtained by PORL. Q^π is defined to be the true Q-function under a policy $\pi(\mathbf{a}|\mathbf{s})$ in the underlying MDP without any regularization. The true Q-function can be written in a recursive form via the exact Bellman operator, $\mathcal{B}^\pi$, as $Q^\pi = \mathcal{B}^\pi Q^\pi$. We define $\hat{V}^k$ as the value function under a policy $\pi(\mathbf{a}|\mathbf{s})$, $\hat{V}^k(\mathbf{s}) = \mathbb{E}_{\mathbf{a}\sim\pi(\mathbf{a}|\mathbf{s})}[\hat{Q}^k(\mathbf{s},\mathbf{a})]$. For the true value function V^π in the underlying MDP, we also have $V^\pi = \mathcal{B}^\pi V^\pi$.

We first introduce the theorem that the learned value function is a lower bound of the true one without considering the sampling error defined in the Lemma A.1.

Theorem 4.1 *Assume we can exactly obtain the reward function $r(\mathbf{s},\mathbf{a})$ and the transition function $T(\mathbf{s}'|\mathbf{s},\mathbf{a})$ in the underlying MDP in the training dataset, and the value function $\hat{V}^k$ converges to $\hat{V}^\pi$. Then for any action $\mathbf{a} \in \mathcal{A}$, and any state $\mathbf{s} \in \mathcal{S}$, the learned value function via Eqn. 5 is a lower bound of the true one, i.e., $\hat{V}^\pi(\mathbf{s}) \leq V^\pi(\mathbf{s})$, if the ratio $\frac{\varepsilon}{\alpha}$ satisfies*

$$\frac{\varepsilon}{\alpha} \leq \min_{\mathbf{s},\mathbf{a}} \left(\sum_{\mathbf{a}} \pi(\mathbf{a}|\mathbf{s}) \left[\frac{\pi(\mathbf{a}|\mathbf{s})}{\pi_\beta(\mathbf{a}|\mathbf{s})} - 1 \right] \right) \left(\frac{\left| d^\phi(\mathbf{s}) - d^{\pi_\beta}(\mathbf{s}) \right|}{d^{\pi_\beta}(\mathbf{s})} \sum_{\mathbf{a}} \frac{\pi^2(\mathbf{a}|\mathbf{s})}{\pi_\beta(\mathbf{a}|\mathbf{s})} \right)^{-1}.$$

It is worth noting that the learned value function still lower bounds the true value function for any state and action in the training datasets, i.e. $\mathbf{s},\mathbf{a} \in \mathcal{D}$, even when we consider the sampling error defined in the Lemma A.1. Further details are shown in Corollary A.1. We have no reward or transition pair collected at unseen states or actions outside the training dataset, so it is impossible to bound the error outside the training dataset when consider the sampling error introduced by the reward function and the transition function.

We can now step further and show that the values at OOD states are lower than those at in-distribution states based on the learned value function. The proof can be found in Appendix A.2.

Theorem 4.2 *For any state $\mathbf{s} \in \mathcal{S}$, if $\varepsilon > 0$ is sufficiently large, then the expectation of the learned value function via Eqn. 5 under the state marginal $d^{\pi_\beta}(\mathbf{s})$ in the training data is higher than that under $d^\phi(\mathbf{s})$, i.e., $\mathbb{E}_{\mathbf{s}\sim d^\phi(\mathbf{s})}[\hat{V}^\pi(\mathbf{s})] < \mathbb{E}_{\mathbf{s}\sim d^{\pi_\beta}(\mathbf{s})}[\hat{V}^\pi(\mathbf{s})]$.*

During training time, we can at least evaluate Q-values of OOD actions based on in-distribution states. However, there is actually no information about immediate rewards at OOD states, thus no information about Q-values. Intuitively, under offline settings, the best we can do to mitigate the problem of OOD states is to suppress values at these OOD states, and raising values at in-distribution states, so that the agent can be attracted to the area where it is familiar near the training data. This idea actually aligns with the concept of pessimistic MDP defined in [10]. In a pessimistic MDP, the rewards at state-action pairs far away from the support of the training data are significantly smaller than those at in-distribution states and actions. Therefore, Thm. 4.2 indeed tells us PORL models a value function of a near pessimistic MDP out of the training data. Optimizing the policy under the pessimistic MDP is equivalent to forcing the policy to avoid unknown states and actions.

In summary, PORL can learn a pessimistic value function that lower bounds the true value function. Furthermore, this value function is corresponding to a pessimistic MDP underlying in the training data, and can lead the agent to avoid or even recover from OOD states.

5 Implementing the Algorithm

In this section, we introduce a practical PORL algorithm based on Eqn. 5. This algorithm simply modifies the policy evaluation step of Deep Q-Learning or Soft Actor-Critic algorithms, which is easy to implement.

5.1 Detecting OOD states

In prior to designing the algorithm, we need to choose a proper $d^\phi(\mathbf{s})$, which requires a tool for OOD state detection. Following [10, 13, 28], we use bootstrapping to detect OOD states. In particular, we train a bag of Gaussian dynamics models [28] $\{\hat{P}_1, \hat{P}_2, \ldots, \hat{P}_n\}$ where each model is $\hat{P}_i(\cdot|\mathbf{s}, \mathbf{a}) = \mathcal{N}(\mathbf{s} + \hat{f}_{\phi_i}(\mathbf{s}, \mathbf{a}), \hat{\Sigma}_{\phi_i})$. The function $\hat{f}_{\phi_i}$ outputs the mean difference between the next state and the current state, and Σ_{ϕ_i} models the standard deviation. OOD states are detected by estimating the uncertainty of bootstrap models at a given state $\mathbf{s} \in \mathcal{S}$. Concretely, we define $u_\pi(\mathbf{s}) = \mathbb{E}_{\mathbf{a}\sim\pi(\mathbf{a}|\mathbf{s})}\left[\frac{1}{n}\sum_{i=1}^{n}\left(\hat{f}_{\phi_i}(\mathbf{s}, \mathbf{a}) - \bar{f}_\phi(\mathbf{s}, \mathbf{a})\right)^2\right]$, where $\bar{f}_\phi(\mathbf{s}, \mathbf{a}) = \dfrac{1}{n}\sum_{i=1}^{n}\hat{f}_{\phi_i}(\mathbf{s}, \mathbf{a})$ is the mean of outputs of all $\hat{f}_{\phi_i}$, and the actions are drawn from a policy distribution π. A high $u_\pi(\mathbf{s})$ value indicates the state is more likely to be an unseen state. Given a set of sampled states $\{s_1, s_2, \ldots, s_n\}$, we can define a discrete distribution over it using $u_\pi(\mathbf{s})$: $\zeta(\mathbf{s}_i) = \dfrac{u(\mathbf{s}_i)}{\sum_j u(\mathbf{s}_j)}, i = 1, 2, ..., n$, which assign high probabilities to OOD states. In the following section, we will use it to construct the distribution $d^\phi(\mathbf{s})$.

5.2 Practical Implementation of PORL

We now introduce a practical PORL algorithm. In practice, to obtain a well-defined distribution d^ϕ, we add an additional optimization problem over d^ϕ into the original optimization problem. The resulting optimization problem for the policy evaluation step is:

$$\min_Q \max_{d^\phi} \left[\varepsilon\left(\mathbb{E}_{\mathbf{s}\sim d^\phi(\mathbf{s}), \mathbf{a}\sim\hat{\pi}^k(\mathbf{a}|\mathbf{s})}\left[Q(\mathbf{s}, \mathbf{a})\right] - \mathbb{E}_{\mathbf{s}\sim d^{\pi_\beta}(\mathbf{s}), \mathbf{a}\sim\hat{\pi}^k(\mathbf{a}|\mathbf{s})}[Q(\mathbf{s}, \mathbf{a})]\right) + \mathcal{R}(d^\phi)\right] + \mathcal{E}(Q, \hat{\mathcal{B}}^\pi \hat{Q}^k_\theta) + \mathcal{C}(Q), \quad (6)$$

where $\mathcal{R}(d^\phi)$ is a regularization term inspired by [4] in order to stabilize the training. If we choose $\mathcal{R}(d^\phi) = -D_{\mathrm{KL}}(d^\phi(\mathbf{s}) \,||\, \zeta(\mathbf{s}))$, where $\zeta(\mathbf{s})$ is the distribution we obtained from uncertainty estimations, then $d^\phi(\mathbf{s}) \propto \zeta(\mathbf{s}) \exp\left(V^{\hat{\pi}^k}(\mathbf{s})\right)$, where $V^{\hat{\pi}^k}(\mathbf{s}) = \mathbb{E}_{\mathbf{a}\sim\hat{\pi}^k(\mathbf{a}|\mathbf{s})}\left[Q(\mathbf{s}, \mathbf{a})\right]$. The resulting d^ϕ is intuitively reasonable, because it assigns high probabilities to OOD states with high uncertainty estimations. In particular, d^ϕ assigns higher probabilities to states with high values, because we

Algorithm 1: Pessimistic Offline Reinforcement Learning (PORL)

Initialize: A Q network Q_θ parametrized by θ, A target network $Q_{\bar{\theta}} = Q_\theta$ parametrized by $\bar{\theta}$, a policy network π_φ parametrized by φ, and a bag of dynamics models $\{\hat{P}_1, \hat{P}_2, \ldots, \hat{P}_n\}$ to detect OOD states;
// Dynamics Models Training (Models are used by $u_{\pi_\varphi}(\mathbf{s})$ to detect OOD states in the policy evaluation step)
for *step i in range*$(0, M)$ **do**
  Train dynamics models $\{\hat{P}_1, \hat{P}_2, \ldots, \hat{P}_n\}$ according to the transitions in the dataset $\mathcal{D}$, so that we can later obtain an uncertainty estimation model $u_\pi(\mathbf{s})$ in the policy evaluation step;
end
// Policy Evaluation and Improvement
for *step t in range*$(0, N)$ **do**
  Update Q_θ according to Eqn. 7 with learning rate ϵ_θ and $u_{\pi_\varphi}(\mathbf{s})$:
  $\theta_t \leftarrow \theta_{t-1} + \epsilon_\theta \nabla_\theta J(\theta)$;
  Update π_φ according to the soft actor critic style objective and learning rate ϵ_φ:
  $\varphi_t \leftarrow \varphi_{t-1} + \epsilon_\varphi \mathbb{E}_{\mathbf{s}\sim d^{\pi_\beta}(\mathbf{s}), \mathbf{a}\sim\pi_\varphi(\mathbf{a})}[Q_\theta(\mathbf{s}, \mathbf{a}) - \log \pi_\varphi(\mathbf{a}|\mathbf{s})]$;
  if *t mod target_update* == 0 **then**
    Soft Update the target network $\bar{\theta}_t \leftarrow (1-\tau)\bar{\theta}_{t-1} + \tau\theta_{t-1}$
  end
end

expect to penalize harder on them than those with low values already. With this choice of d^{ϕ} in Eqn. 6, we obtain the following PORL policy evaluation step:

$$\min_Q J(Q) = \min_Q \varepsilon \left(\log \sum_{\mathbf{s}} \zeta(\mathbf{s}) \exp\left(V^{\hat{\pi}^k}(\mathbf{s})\right) - \mathbb{E}_{\mathbf{s}\sim d^{\pi_\beta}(\mathbf{s})}[V^{\hat{\pi}^k}(\mathbf{s})] \right) + \mathcal{E}(Q, \hat{\mathcal{B}}^{\pi}\hat{Q}_{\theta}^{k}) + \mathcal{C}(Q). \quad (7)$$

The first term in Eqn. 7 is very similar to weighted softmax values over the state space. It penalizes the softmax value over the state space, but also considers the distances between sample points and the training data. The two terms following the trade-off factor ε is actually trying to decrease the discrepancy between the softmax value over OOD states and the average value over in-distribution states. Intuitively, it should enforce the learned value function to output higher values at in-distribution states, and lower values at out-of-distribution states. The logsumexp term in Eqn. 7 also mitigates the requirement for an accurate uncertainty estimation $\zeta(\mathbf{s})$ over the entire state space. Only those states with high values contribute to the regularization.

The complete algorithm is shown in Algorithm 1. We include the version for continuous action space which requires a policy network here, and note that if the action space is discrete, then we no longer need a policy network but just an implicit policy based on the learned Q-function. We implement PORL on top of CQL [4], with its default hyperparameters. We also apply Lagrangian dual gradient descent to automatically adjust the trade-off factor ε.

6 Experiments

We compare our algorithm to prior offline algorithms: two state-of-the-art offline RL algorithms BEAR [5] and CQL [4]; two baselines adapted directly from online algorithms, actor-critic algorithm TD3 [29] and DDQN [30]; and behavior cloning (BC). The TD3 baseline is applied when the action space is continuous, whereas DDQN is trained when the action space is discrete. We evaluate each algorithm on a wide range of task domains, including tasks with both continuous and discrete state and action space. All baselines are run with the default code and hyperparameters from the original repositories. In particular, we are interested in the comparison between our algorithm with CQL, because we essentially add an additional state regularization term to the original CQL framework.

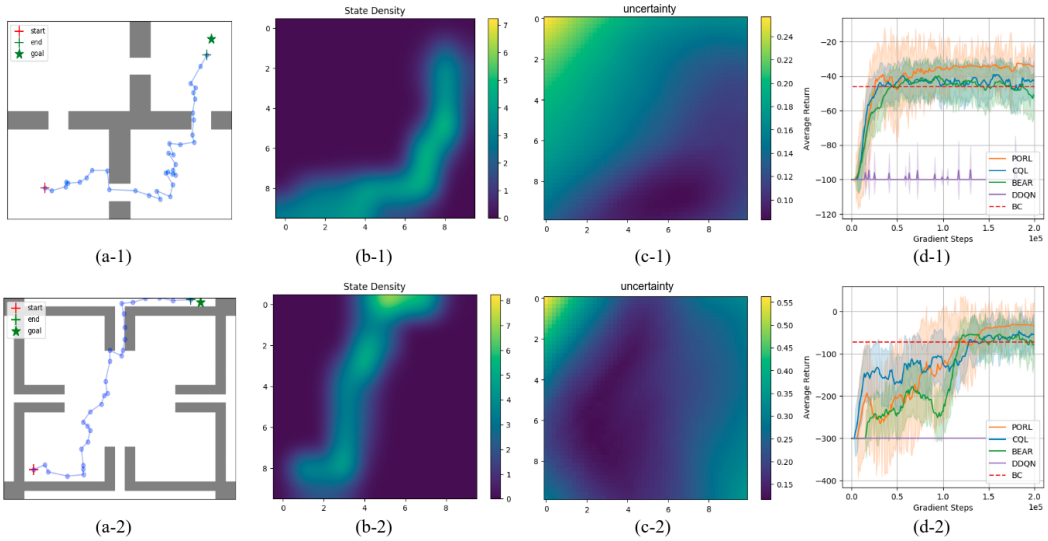

Figure 1: (a) The whole map of the environment; (b) The state density in training dataset; (c) The visualization of uncertainty estimation; (d) The learning curves. The top row (1) and the bottom row (2) are corresponding to PointmassHard-v0 and PointmassSuperHard-v0, respectively. We can see that almost all trajectories in the training datasets are located around the optimal trajectory from the start to the goal in the yellow areas in (b), indicating they are collected by a near-optimal policy.

6.1 Performance on Various Environments

Pointmass Mazes. The task for the agent in this domain is to learn from expert demonstrations to navigate from a random start to a fixed goal. The expert dataset, which contains around 1000 trajectories all from the same start point to the same goal, is collected by online trained RL policy. During the test time of offline RL algorithms, we reset the start to a random point in the state space and the goal to the same fixed point as the dataset. In this way, the performance of the agent at unseen states are evaluated.

Before showing the performance, we first check if the OOD states detection is accurate, and hence, if we can successfully penalize high values at unseen states in training datasets. We evaluate the effectiveness of the OOD states detection method based on the accuracy of uncertainty estimation in the environment Pointmass. Figure 1(b) and (c) are visualizations of the training datasets and estimated uncertainty $u_\pi(\mathbf{s})$, both of which have the coordinate systems the same as that in the map (figure 1(a)). We use different colors in figure 1(b) and (c) to represent different values at each point in the map. The uncertainty estimations tend to be high (yellow areas in Fig. 1(c)) in area where the state density is low (blue areas in Fig 1(b)), and vice versa. This trend empirically shows that our uncertainty estimations are reasonable. We can trust them to detect OOD states when training offline RL algorithms.

We include the learning curves in figure 1(d), in which we evaluate each algorithm based on 3 random seeds, and report the average return. The shaded area represents the standard deviation of each evaluation. As we can see in the figures, PORL outperforms other baselines in both hard and super hard environments. PORL benefits from the augmented policy evaluation step in Eqn. 7. The learned value function produces high values at areas that have low uncertainty estimations, and low values at highly uncertain areas (OOD states). Therefore, the agent can be "attracted" to the high value areas from low value and unfamiliar areas.

Gym Tasks. In this domain, we focus on the locomotion environments from MuJoCo, including Walker2d-v2, Hopper-v2, Halfcheetah-v2, and Ant-v2. Unlike Pointmass environment, we directly adopt the d4rl datasets [31] as our training data in the gym domains. We include four different types of datasets in our experiments, namely, "medium", "medium-expert", "random", and "expert". The

Table 1: Performance on Gym and Adroit Domains

Domain	Task	BC	TD3	BEAR	CQL	PORL
Gym	hopper-medium	29.71	0.80	50.08	63.04	**76.39**
	walker2d-medium	13.06	4.91	34.20	70.67	**75.95**
	halfcheetah-medium	35.81	24.60	41.25	48.66	**49.21**
	ant-medium	83.42	−63.55	−45.78	51.29	**85.31**
	hopper-medium-expert	85.22	12.12	37.67	105.84	**112.80**
	walker2d-medium-expert	16.12	4.77	17.29	75.27	**89.67**
	halfcheetah-medium-expert	**37.04**	−1.41	−2.93	19.13	24.33
	ant-medium-expert	**66.49**	−63.54	31.31	34.44	49.95
	hopper-random	9.44	8.47	9.29	**10.36**	**10.82**
	walker2d-random	2.08	6.04	0.72	1.11	**2.66**
	halfcheetah-random	2.25	27.95	2.16	26.62	**28.58**
	ant-random	26.00	−37.31	24.05	26.65	**27.86**
	hopper-expert	109.82	1.81	0.78	104.41	**110.67**
	walker2d-expert	60.66	−0.95	21.37	105.55	**109.37**
	halfcheetah-expert	**94.22**	−1.40	13.55	80.19	71.33
	ant-expert	**73.57**	55.07	44.65	59.04	60.54
Adroit	pen-human	34.46	−3.83	35.62	54.49	**63.07**
	door-human	1.46	−0.19	−0.33	1.86	**2.28**
	hammer-human	1.35	0.26	0.46	3.91	**4.24**
	relocate-human	0.04	−0.32	−0.30	0.15	**0.33**
	pen-cloned	23.58	−3.91	28.92	35.23	**39.02**
	door-cloned	0.15	−0.33	−0.16	1.72	**1.69**
	hammer-cloned	0.40	0.25	0.21	0.53	**0.95**
	relocate-cloned	−0.24	**−0.14**	−0.23	−0.28	−0.26

"medium", "random", and "expert" dataset are all collected by a single policy, which is an either early-stopping trained, or randomly initialized, or fully trained expert policy. The "medium-expert" dataset is generated by mixing mediocre and expert quality data. We show the normalized scores averaged over 4 random seeds for all methods on gym domain in table 1. We directly ran all baselines from their original repositories with their default parameters, and we only report the average scores we actually obtained. As we can see in the table, PORL outperforms all other offline RL methods on a majority of tasks on gym domains. PORL works especially well with mediocre quality datasets according to the results. In fact, it is one of the advantages of offline RL methods over behavior cloning on medium quality datasets, because offline RL methods take advantage of the information both from the reward and the underlying state and action distributions in training datasets, instead of simply imitating behavior policies as behavior cloning. Medium quality datasets are also considered to be similar to real-world datasets. Therefore, it is important for an offline RL method to perform well in medium quality datasets. We also note that PORL shares some good properties with CQL, such as satisfying performance on mixed quality datasets. PORL and CQL both outperform other offline methods on medium-expert datasets with PORL better between them. The reason is that offline RL methods can "stitch" [31] different trajectories from different policies together according to the information from the reward.

Adroit Tasks The adroit domain [32] provides more challenging tasks than the Pointmass environment and the gym domain. The tasks include controlling a 24-DoF simulated Shadow Hand robot to twirl a pen, open a door, hammer a nail, and relocate a ball. Similar to the datasets in the gym domain, we also directly use the d4rl datasets as the training datasets in our experiments. The performance of PORL and all baselines is shown in table 1. The normalized scores of all methods are average returns on 5 random seeds. We note that PORL has better performance than other baselines on adroit domains. It is a great advantage for PORL to learn useful skills from human demonstrations on these high dimensional and highly realistic robotic simulations.

6.2 Discussions and Limitations

The main contribution of this work is to explicitly limit the values at OOD states, so that the learned policy can act conservatively at OOD states and drives the agent back to the familiar areas near the training data. We are interested to see if our framework can indeed induce a different behavior on OOD states. We use $\Delta_k = \max_{\mathbf{s}\in\mathcal{S}}[V(\mathbf{s})] - \mathbb{E}_{\mathbf{s}\sim\mathcal{D}}[V(\mathbf{s})]$ as a metric to evaluate it at each iteration. If Δ_k is close to zero, then intuitively it indicates the values at OOD states are lower than those at in-distribution states. In Fig. 2, we plot Δ_k at each iteration in hopper-medium-v0. As is shown in the figure, PORL successfully limits Δ_k to be non-positive, which meets our goal in this work and aligns with the statement in theorem 4.2.

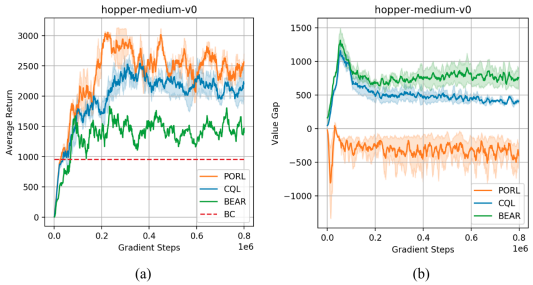

Figure 2: (a) The learning curves in hopper-medium-v0. (b) The discrepancy Δ_k as a function of gradient steps for PORL, CQL, and BEAR.

On the gym domain, we notice that the performance of PORL and CQL on datasets containing expert trajectories is not satisfying, often not as good as BC. We believe it is because of overly conservative value estimation. In fact, it is widely believed that conservative methods suffer from underestimation [13]. The conservative objective function in Eqn. 7 sometimes assign values that are too low to OOD states and actions. Besides, the uncertainty estimation method cannot be guaranteed to be precise on high-dimensional spaces. It is actually a possible future work direction to solve the underestimation and uncertainty estimation problems in conservative methods.

7 Conclusion

We propose a Pessimistic Offline Reinforcement Learning framework to deal with out-of-distribution states. In particular, we add a regularization term in policy evaluation step to shape value function, so that we can improve its extrapolation to OOD states. We also provide theoretical guarantees that the learned pessimistic value function lower bounds the true one and is actually corresponding to a Pessimistic MDP. We evaluate the PORL algorithm on various benchmark tasks, where we show that our method gains better performance by explicitly handling OOD states compared to those methods merely considering OOD actions.

References

[1] N. Jaques, A. Ghandeharioun, J. H. Shen, C. Ferguson, A. Lapedriza, N. Jones, S. Gu, and R. Picard. Way off-policy batch deep reinforcement learning of implicit human preferences in dialog. *arXiv preprint arXiv:1907.00456*, 2019.

[2] D. Kalashnikov, A. Irpan, P. Pastor, J. Ibarz, A. Herzog, E. Jang, D. Quillen, E. Holly, M. Kalakrishnan, V. Vanhoucke, et al. Scalable deep reinforcement learning for vision-based robotic manipulation. In *Conference on Robot Learning*, pages 651–673. PMLR, 2018.

[3] G. Kahn, P. Abbeel, and S. Levine. Badgr: An autonomous self-supervised learning-based navigation system. *IEEE Robotics and Automation Letters*, 6(2):1312–1319, 2021.

[4] A. Kumar, A. Zhou, G. Tucker, and S. Levine. Conservative q-learning for offline reinforcement learning. *arXiv preprint arXiv:2006.04779*, 2020.

[5] A. Kumar, J. Fu, G. Tucker, and S. Levine. Stabilizing off-policy q-learning via bootstrapping error reduction. *arXiv preprint arXiv:1906.00949*, 2019.

[6] I. Osband, C. Blundell, A. Pritzel, and B. Van Roy. Deep exploration via bootstrapped dqn. *arXiv preprint arXiv:1602.04621*, 2016.

[7] B. Eysenbach, S. Gu, J. Ibarz, and S. Levine. Leave no trace: Learning to reset for safe and autonomous reinforcement learning. *arXiv preprint arXiv:1711.06782*, 2017.

[8] R. Agarwal, D. Schuurmans, and M. Norouzi. An optimistic perspective on offline reinforcement learning. In *International Conference on Machine Learning*, pages 104–114. PMLR, 2020.

[9] B. O'Donoghue, I. Osband, R. Munos, and V. Mnih. The uncertainty bellman equation and exploration. In *International Conference on Machine Learning*, pages 3836–3845, 2018.

[10] R. Kidambi, A. Rajeswaran, P. Netrapalli, and T. Joachims. Morel: Model-based offline reinforcement learning. *arXiv preprint arXiv:2005.05951*, 2020.

[11] Y. Wu, G. Tucker, and O. Nachum. Behavior regularized offline reinforcement learning. *arXiv preprint arXiv:1911.11361*, 2019.

[12] W. Zhou, S. Bajracharya, and D. Held. Plas: Latent action space for offline reinforcement learning. *arXiv preprint arXiv:2011.07213*, 2020.

[13] S. Levine, A. Kumar, G. Tucker, and J. Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.

[14] D. Precup. Eligibility traces for off-policy policy evaluation. *Computer Science Department Faculty Publication Series*, page 80, 2000.

[15] S. Gu, E. Holly, T. Lillicrap, and S. Levine. Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In *2017 IEEE international conference on robotics and automation (ICRA)*, pages 3389–3396. IEEE, 2017.

[16] J. Huang and N. Jiang. From importance sampling to doubly robust policy gradient. In *International Conference on Machine Learning*, pages 4434–4443. PMLR, 2020.

[17] S. Pankov. Reward-estimation variance elimination in sequential decision processes. *arXiv preprint arXiv:1811.06225*, 2018.

[18] C.-A. Cheng, X. Yan, and B. Boots. Trajectory-wise control variates for variance reduction in policy gradient methods. In *Conference on Robot Learning*, pages 1379–1394. PMLR, 2020.

[19] S. Fujimoto, D. Meger, and D. Precup. Off-policy deep reinforcement learning without exploration. In *International Conference on Machine Learning*, pages 2052–2062. PMLR, 2019.

[20] N. Y. Siegel, J. T. Springenberg, F. Berkenkamp, A. Abdolmaleki, M. Neunert, T. Lampe, R. Hafner, N. Heess, and M. Riedmiller. Keep doing what worked: Behavioral modelling priors for offline reinforcement learning. *arXiv preprint arXiv:2002.08396*, 2020.

[21] X. B. Peng, A. Kumar, G. Zhang, and S. Levine. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning. *arXiv preprint arXiv:1910.00177*, 2019.

[22] A. Nair, M. Dalal, A. Gupta, and S. Levine. Accelerating online reinforcement learning with offline datasets. *arXiv preprint arXiv:2006.09359*, 2020.

[23] R. S. Sutton. Dyna, an integrated architecture for learning, planning, and reacting. *ACM Sigart Bulletin*, 2(4):160–163, 1991.

[24] M. Watter, J. T. Springenberg, J. Boedecker, and M. Riedmiller. Embed to control: A locally linear latent dynamics model for control from raw images. *arXiv preprint arXiv:1506.07365*, 2015.

[25] M. Zhang, S. Vikram, L. Smith, P. Abbeel, M. Johnson, and S. Levine. Solar: Deep structured representations for model-based reinforcement learning. In *International Conference on Machine Learning*, pages 7444–7453. PMLR, 2019.

[26] D. Hafner, T. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson. Learning latent dynamics for planning from pixels. In *International Conference on Machine Learning*, pages 2555–2565. PMLR, 2019.

[27] M. Janner, J. Fu, M. Zhang, and S. Levine. When to trust your model: Model-based policy optimization. *arXiv preprint arXiv:1906.08253*, 2019.

[28] K. Chua, R. Calandra, R. McAllister, and S. Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper/2018/file/3de568f8597b94bda53149c7d7f5958c-Paper.pdf.

[29] S. Fujimoto, H. Hoof, and D. Meger. Addressing function approximation error in actor-critic methods. In *International Conference on Machine Learning*, pages 1587–1596. PMLR, 2018.

[30] H. Van Hasselt, A. Guez, and D. Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.

[31] J. Fu, A. Kumar, O. Nachum, G. Tucker, and S. Levine. D4rl: Datasets for deep data-driven reinforcement learning, 2020.

[32] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine. Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. *arXiv preprint arXiv:1709.10087*, 2017.