Weighted model estimation for offline model-based reinforcement learning

Anonymous Author(s)
Affiliation
Address
email

Abstract

This paper discusses model estimation in offline model-based reinforcement learning (MBRL), which is important for subsequent policy improvement using an estimated model. From the viewpoint of covariate shift, a natural idea is model estimation weighted by the ratio of the state-action distributions of offline data and real future data. However, estimating such a natural weight is one of the main challenges for off-policy evaluation, which is not easy to use. As an artificial alternative, this paper considers weighting with the state-action distribution ratio of offline data and simulated future data, which can be estimated relatively easily by standard density ratio estimation techniques for supervised learning. Based on the artificial weight, this paper defines a loss function for offline MBRL and presents an algorithm to optimize it. Weighting with the artificial weight is justified as evaluating an upper bound of the policy evaluation error. Numerical experiments demonstrate the effectiveness of weighting with the artificial weight.

1 Introduction

Reinforcement learning (RL) is a framework to learn a policy in an unknown environment [1]. Model-based RL (MBRL) is an approach that explicitly estimates a transition model and utilizes it to improve a policy. Compared to model-free RL, one major advantage of MBRL is the data efficiency [2–4], which is important for applications where the data collection is expensive, such as robotics or healthcare. This paper focuses on offline MBRL (also known as batch MBRL).

One of the main challenges in offline MBRL is distribution shift [5]. The distribution of offline data is different from the distribution of data obtained in the future when applying an improved policy. This is a situation called covariate shift, where a model estimated using standard supervised learning techniques such as empirical risk minimization (ERM) may not generalize to test data. A major idea in MBRL is to estimate a model using such a standard method and to improve a policy using the estimated model in regions where the predictions are accurate [6–8]. A point to be improved here is model estimation without considering covariate shift. In supervised learning under covariate shift, importance-weighted ERM can obtain good predictive performance [9]. If a model in offline MBRL can be estimated in a similar way, then the estimated model will make more accurate predictions, resulting in increased policy improvement. Motivated by this, this paper tackles the issue of weighted model estimation considering covariate shift in offline MBRL.

From the viewpoint of covariate shift, a natural idea is weighting with the ratio of the state-action distributions of offline data and real future data, as the agent uses a model to predict real future data. For convenience, this paper calls the ratio "the natural weight." However, estimating the natural weight is one of the main challenges for off-policy evaluation [10, 5]. Compared to density ratio estimation for supervised learning [11], the difficulty of off-policy evaluation is the inaccessibility of samples from the distribution of real future data.

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

As an alternative, this paper considers weighting with the state-action distribution ratio of offline data and "simulated" future data. This paper calls the ratio "the artificial weight," in contrast to the natural weight. Simulated future data can be generated in offline simulation, unlike real future data. Since samples from the both distributions are available, the artificial weight can be relatively easily obtained by standard density ratio estimation techniques for supervised learning. Based on the artificial weight, this paper defines a loss function for offline MBRL in Section 4. This paper presents an algorithm to optimize it in Section 5.

The question here is the validity of weighting with the artificial weight, as it may not be natural for covariate shift. This paper justifies it as evaluating an upper bound of the policy evaluation error in Section 4. This paper demonstrates that it works in practice in numerical experiments in Section 6. The above is the contribution of this paper.

2 Related Works

Most of offline MBRL researches [7, 8, 12] discuss how to learn a policy using a learned model. In contrast, this paper studies how to estimate a model considering covariate shift in MBRL, which depends on a policy to be applied. In principle, these can be combined as an algorithm that repeats between estimating a model and learning a policy, like Algorithm 3.

Importance-sampling-based methods estimate the natural weight, namely, the distribution ratio of offline data and real future data on a trajectory or state(-action) space [13–15, 10]. The advantage is unbiasedness, while the disadvantage is high variance. Model-based methods, including this paper, evaluate policies utilizing models trained using data. Model-based methods have the opposite advantage and disadvantage of importance-sampling-based methods. Doubly robust methods combine importance-sampling-based and model-based methods [16–18]. In principle, incorporating the idea of this paper as a model-based component is possible.

VAML [19, 20] defines a state-action-wise loss function based on the expected error of Bellman operator and learns a model to optimize a weighted sum. Similar ideas are found in [21–23]. The main focus is the state-action-wise loss function, while how to weight it is not discussed sufficiently. In contrast, this paper simply uses the negative log likelihood loss function for a state-action pair and discusses how to weight it. Combining these two ideas is an interesting direction.

This paper is similar to VMBPO [24] in that an EM-style algorithm updates a model and a policy using a universal loss function. The differences is that this paper considers covariate shift but does not use the idea of RL as inference [25]. Formulation that explicitly incorporates RL as inference is also an interesting topic.

Another approach is learning a model robust against covariate shift by reducing the difference between the distributions of offline data and future data on a feature space [22, 26, 27]. Since covariate shift measured in a feature space is encouraged to decrease, policy improvement increasing covariate shift is also discouraged. That is, such an idea implicitly limits policy improvement. This paper differs in that it is trying to utilize covariate shift.

3 Preliminary

This paper discusses a discounted infinite horizon MDP [28]. Let s be a state. Let a be a action. Let $\rho(s)$ be an initial state distribution. Let $\mathcal{P}(s'|s,a)$ be a transition probability function. Let $r(s,a)$ be a reward function. Let γ be a discount factor. Let π be a policy. The value function is

$$V_{\mathcal{P}}^{\pi}(s^0) = \mathbb{E}_{a^t \sim \pi(\cdot|s^t), s_{t+1} \sim \mathcal{P}(\cdot|s^t,a^t)} \left[\sum_{t=0}^{\infty} \gamma^t r(s^t, a^t) \middle| s^0 \right]. \tag{1}$$

The expected return is $\eta_{\mathcal{P}}^{\pi} = \mathbb{E}_{s^0 \sim \rho}\left[V_{\mathcal{P}}^{\pi}(s)\right]$. The state distribution is

$$d_{\mathcal{P}}^{\pi}(s) = (1-\gamma) \sum_{t=0}^{\infty} \gamma^t \mathrm{Pr}\left(s^t = s \middle| \rho, \pi, \mathcal{P}\right). \tag{2}$$

The state-action distribution is $d_{\mathcal{P}}^{\pi}(s,a) = d_{\mathcal{P}}^{\pi}(s)\pi(a|s)$. Note that $\eta_{\mathcal{P}}^{\pi} = \mathbb{E}_{(s,a)\sim d_{\mathcal{P}}^{\pi}}\left[\frac{r(s,a)}{1-\gamma}\right]$.

Let $\mathcal{P}_*$ denote the true transition probability, which is unknown to the agent. Let $\mathcal{P}_\theta$ denote the transition probability model parameterized by θ. Let $\mathcal{D} = \{(s_n, a_n, s'_n)\}_{n=1}^N$ be offline data, which is obtained from $\mathcal{P}_*$. Let $d^{\mathcal{D}}_{\mathcal{P}_*}$ be the distribution of the input part of $\mathcal{D}$, $\{(s_n, a_n)\}_{n=1}^N$. Let $\hat{\mathcal{D}}^\pi_\theta$ be simulation data generated given $(\mathcal{P}_\theta, \pi)$ with state initialization with probability $1-\gamma$, which is equal to sampling (s, a) from $d^\pi_{\mathcal{P}_\theta}$. Let $c_\theta(s, a) = \mathbb{E}_{s' \sim \mathcal{P}_*(\cdot|s,a)} [-\ln \mathcal{P}_\theta(s'|s, a)]$ be the cross entropy between $\mathcal{P}_*(\cdot|s, a)$ and $\mathcal{P}_\theta(\cdot|s, a)$. Let $h(s, a) = \mathbb{E}_{s' \sim \mathcal{P}_*(\cdot|s,a)} [-\ln \mathcal{P}_*(s'|s, a)]$ be the self-entropy of $\mathcal{P}_*(\cdot|s, a)$.

4 Loss Function

This paper discusses the theory for the case where the reward function is known. For the unknown reward function, if the prediction of a reward model is bounded, the same loss functions can be derived, see Appendix.

4.1 Policy Evaluation

The goal of policy evaluation is estimating $\eta^\pi_{\mathcal{P}_*}$ when given π. For a model-based setting, the goal reduces to finding θ decreasing $|\eta^\pi_{\mathcal{P}_*} - \eta^\pi_{\mathcal{P}_\theta}|$, which is bounded as

$$\begin{aligned}
|\eta^\pi_{\mathcal{P}_*} - \eta^\pi_{\mathcal{P}_\theta}| &\leq \frac{\gamma \mathbb{E}_{(s,a)\sim d^\pi_{\mathcal{P}_\theta}} \left[\sum_{s'} \left(\mathcal{P}_*(s'|s, a) - \mathcal{P}_\theta(s'|s, a)\right) V^\pi_{\mathcal{P}_*}(s')\right]}{1-\gamma} \\
&\leq \frac{B \mathbb{E}_{(s,a)\sim d^\pi_{\mathcal{P}_\theta}} \left[||\mathcal{P}_*(\cdot|s, a) - \mathcal{P}_\theta(\cdot|s, a)||_1\right]}{\sqrt{2}} \\
&\leq B\sqrt{\mathbb{E}_{(s,a)\sim d^\pi_{\mathcal{P}_\theta}} \left[c_\theta(s, a) - h(s, a)\right]} \\
&= B\sqrt{\mathbb{E}_{(s,a)\sim d^{\mathcal{D}}_{\mathcal{P}_*}} \left[w^\pi_\theta(s, a)(c_\theta(s, a) - h(s, a))\right]} \qquad (3) \\
&\leq B\sqrt{\mathbb{E}_{(s,a)\sim d^{\mathcal{D}}_{\mathcal{P}_*}} \left[w^\pi_\theta(s, a) c_\theta(s, a)\right] - h_{\min}}, \qquad (4)
\end{aligned}$$

where $B = \frac{\gamma \max_{s,a} |r(s,a)|\sqrt{2}}{(1-\gamma)^2}$, $w^\pi_\theta(s, a) = \frac{d^\pi_{\mathcal{P}_\theta}(s,a)}{d^{\mathcal{D}}_{\mathcal{P}_*}(s,a)}$, and $h_{\min} = \min_{s,a} h(s, a)$. The first inequality is the Telescoping lemma [21]. The second inequality holds from Holder's inequality and $||V^\pi_{\mathcal{P}}||_\infty \leq \frac{\max_{s,a} |r(s,a)|}{1-\gamma}$. The third inequality holds from Pinsker's inequality and Jensen's inequality. When trying to optimize Equation (3), $w^\pi_\theta(s, a)h(s, a)$ is difficult to handle, as $h(s, a)$ is unknown. The fourth inequality (4) excludes the difficult part from the optimization.

This paper considers decreasing the policy evaluation error by minimizing Equation (4). What is to be minimized is

$$L^\pi(\theta) = \mathbb{E}_{(s,a)\sim d^{\mathcal{D}}_{\mathcal{P}_*}} \left[w^\pi_\theta(s, a) c_\theta(s, a)\right] = \mathbb{E}_{(s,a)\sim d^{\mathcal{D}}_{\mathcal{P}_*}, s' \sim \mathcal{P}_*(\cdot|s,a)} \left[-w^\pi_\theta(s, a) \ln \mathcal{P}_\theta(s'|s, a)\right]. \quad (5)$$

The empirical average of Equation (5) is

$$\hat{L}^\pi(\theta) = -\frac{1}{N} \sum_{n=1}^N w^\pi_\theta(s_n, a_n) \ln \mathcal{P}_\theta(s'_n|s_n, a_n), \quad (6)$$

Note that Equation(6) is a loss function weighted with $w^\pi_\theta(s, a)$, which is the artificial weight, namely, the distribution ratio of offline data $\mathcal{D}$ and simulated future data $\hat{\mathcal{D}}^\pi_\theta$. Thus, weighted model estimation weighting with the artificial weight is justified as evaluating an upper bound of the policy evaluation error.

4.2 Policy optimization

The goal of policy optimization is finding π increasing $\eta^\pi_{\mathcal{P}_*}$. In offline MBRL, the agent utilizes $\eta^\pi_{\mathcal{P}_\theta}$ for the optimization. A lower bound of $\eta^\pi_{\mathcal{P}_*}$ is

$$\eta^\pi_{\mathcal{P}_*} \geq \eta^\pi_{\mathcal{P}_\theta} - |\eta^\pi_{\mathcal{P}_*} - \eta^\pi_{\mathcal{P}_\theta}| \geq \eta^\pi_{\mathcal{P}_\theta} - B\sqrt{L^\pi(\theta) - h_{\min}}. \quad (7)$$

The first term evaluates the expected return by model $\mathcal{P}_\theta$, and the second penalizes the policy evaluation error.

Inspired from increasing $\eta^\pi_{\mathcal{P}_*}$ by maximizing Equation (7), this paper defines a generalization loss function for offline MBRL by

$$J(\theta, \pi) = -\eta^\pi_{\mathcal{P}_\theta} + B'\sqrt{L^\pi(\theta) - h_{\min}} \tag{8}$$

where $B' \in [0, B]$ is a tunable coefficient to scale the penalty for the policy evaluation error.

The agent optimizes an estimate of Equation (8), where the artificial weight is used to estimate $\eta^\pi_{\mathcal{P}_\theta}$ and $L^\pi(\theta)$. That is, the agent optimizes a loss function weighted with w^π_θ, similarly to Section 4.1. For the same reason as Section 4.1, weighting with the artificial ratio is justified as evaluating the expected return with the penalty for the policy evaluation error.

5 Algorithm

5.1 Policy Evaluation

5.1.1 Full Version

This paper considers a gradient-based optimization algorithm. The gradient of Equation (6) is

$$\nabla_\theta \hat{L}^\pi(\theta) = -\frac{1}{N}\sum_{n=1}^{N} w^\pi_\theta(s_n, a_n)\left\{\nabla_\theta \ln \mathcal{P}_\theta(s'_n|s_n, a_n) + \ln \mathcal{P}_\theta(s'_n|s_n, a_n) z^\pi_\theta(s_n)\right\}, \tag{9}$$

where $z^\pi_\theta(s) = \nabla_\theta \ln d^\pi_{\mathcal{P}_\theta}(s)$, which can be estimated using an extension of LSDG [29], described later. The artificial weight, w^π_θ, can be estimated using a density ratio estimation method given $(D, \hat{D}^\pi_{\mathcal{P}_\theta})$. The gradient of Equation (6) can be estimated using them.

In practice, the computational efforts to estimate (w^π_θ, z_θ) are non-negligible. This paper presents a method that iterates between estimating (w^π_θ, z_θ) and updating θ locally. To locally update from θ_0, this paper uses a local approximate loss function defined by

$$\hat{L}^\pi(\theta; \theta_0) = -\frac{1}{N}\sum_{n=1}^{N} w^\pi_\theta(s_n, a_n)\left\{\ln \mathcal{P}_\theta(s'_n|s_n, a_n) + (\theta - \theta_0)\cdot\left(\ln \mathcal{P}_{\theta_0}(s'_n|s_n, a_n) z^\pi_{\theta_0}(s_n)\right)\right\}. \tag{10}$$

During updating θ, since $(w^\pi_{\theta_0}, z^\pi_{\theta_0})$ is fixed, the gradient of Equation (10) can be obtained immediately. For $\theta = \theta_0$, Equation (6) and Equation (10) are equal, and their gradients are also equal. That is, for $\theta = \theta_0$, the first-order necessary conditions for optimality are the same. In theory, if assuming that $(w^\pi_{\theta_0}, z^\pi_{\theta_0})$ is accurately estimated, the iterative method described above converges.

Importance-weighted ERM is consistent but can be unstable in practice [9], even when the importance weight is exactly known. Importance-weighted ERM weighted with w^π_θ can be more unstable due to the estimation and approximation errors of density ratio estimation of w^π_θ. To stabilize the estimation of θ, this paper replaces w^π_θ in Equation (10) with $[w^\pi_{\theta_0}(s_n, a_n)]^\alpha$, where $\alpha = 0$ corresponds to ERM, and $\alpha = 1$ corresponds to importance-weighted ERM weighted with w^π_θ. That is, α is a hyperparameter controlling the trade-off between consistency and stability.

Extending LSDG. LSDG [29] uses the gradient of $\ln d^\pi_{\mathcal{P}}(s)$ with respect to policy parameters in model-free RL. This paper extends this idea to the gradient with respect to model parameters in MBRL. In a discounted MDP, $d^\pi_{\mathcal{P}_\theta}$ satisfies

$$d^\pi_{\mathcal{P}_\theta}(s') = (1-\gamma)\rho(s') + \gamma\sum_{s,a} d^\pi_{\mathcal{P}_\theta}(s)\pi(a|s)\mathcal{P}_\theta(s'|s, a). \tag{11}$$

The gradient with respect to the i-th parameter, θ_i, is

$$d^\pi_{\mathcal{P}_\theta}(s') z^\pi_{\theta_i}(s') = \sum_{s,a} d^\pi_{\mathcal{P}_\theta}(s)\pi(a|s)\mathcal{P}_\theta(s'|s, a)\left\{\gamma\nabla_{\theta_i}\ln \mathcal{P}_\theta(s'|s, a) + \gamma z^\pi_{\theta_i}(s)\right\}, \tag{12}$$

Algorithm 1 Policy evaluation based on weighted model estimation (full version).

Input: $(\mathcal{D}, \pi)$ and α.
Initialize $\mathcal{P}_{\theta_0}$.
for $i = 1, 2, \cdots$ **do**
  Generate $\hat{\mathcal{D}}^{\pi}_{\theta_0}$, using simulation given $(\mathcal{P}_{\theta_0}, \pi)$.
  Estimate $w^{\pi}_{\theta_0}$, using density ratio estimation given $(\mathcal{D}, \hat{\mathcal{D}}^{\pi}_{\theta_0})$.
  Estimate $z^{\pi}_{\theta_0}$, using extended-LSDG given $(\mathcal{P}_{\theta_0}, \pi)$.
  Update $\mathcal{P}_{\theta_0}$, maximizing Equation (10) given $(\mathcal{D}, w^{\pi}_{\theta_0}, z^{\pi}_{\theta_0}, \alpha)$.
end for

Algorithm 2 Policy evaluation based on weighted model estimation (simplified version).

Input: $(\mathcal{D}, \pi)$ and α.
Initialize $\mathcal{P}_{\theta_0}$.
for $i = 1, 2, \cdots$ **do**
  Generate $\hat{\mathcal{D}}^{\pi}_{\theta_0}$, using simulation given $(\mathcal{P}_{\theta_0}, \pi)$.
  Estimate $w^{\pi}_{\theta_0}$, using density ratio estimation given $(\mathcal{D}, \hat{\mathcal{D}}^{\pi}_{\theta_0})$.
  Update $\mathcal{P}_{\theta_0}$, maximizing Equation (13) given $(\mathcal{D}, w^{\pi}_{\theta_0}, \alpha)$.
end for

where $z^{\pi}_{\theta_i}(s) = \nabla_{\theta_i} \ln d^{\pi}_{\mathcal{P}_\theta}(s)$. Equation (12) implies that $z^{\pi}_{\theta_i}$ is the value function satisfying a "forward" Bellman equation, where the transition probability function is $\mathcal{P}_\theta$, and the reward function is $\gamma \nabla_{\theta_i} \ln \mathcal{P}_\theta(s'|s,a)$. To estimate $z^{\pi}_{\theta_i}$, this paper extends LSDG by replacing the gradient of $\ln \pi(a|s)$ with $\gamma \nabla_{\theta_i} \ln \mathcal{P}_\theta(s'|s,a)$. The extended-LSDG estimates $z^{\pi}_{\theta_i}$ using simulation data generated using $(\mathcal{P}_\theta, \pi)$. Since some off-policy evaluation methods estimate the discounted state distribution ratio as a solution to another forward Bellman equation [14, 15, 5], their extensions could also be used instead of LSDG.

Algorithm 1 summarizes how to optimize θ for policy evaluation, described above.

5.1.2 Simplified version

The computationally expensive part in Algorithm 1 is estimating $z^{\pi}_{\theta_0}$. This requires estimating the same number of the value functions satisfying Equation (12) as the number of model parameters. This is unrealistic for neural networks with many parameters. As a practical simplification, this paper also considers replacing Equation (10) with

$$\hat{L}^{\pi}(\theta; \theta_0) = -\frac{1}{N} \sum_{n=1}^{N} w^{\pi}_{\theta_0}(s_n, a_n) \ln \mathcal{P}_\theta(s'_n|s_n, a_n). \tag{13}$$

Algorithm 2 is the resulting simplified algorithm. The price of the simplification is the loss of convergence in theory, as the first-order necessary conditions are no longer the same. For this reason, the number of iterations should be specified. In practice, Algorithm 2 also works well, see Section 6.

5.2 Policy Optimization

This paper considers an EM-style optimization algorithm, where the E-step fixes π and optimizes θ, and the M-step fixes θ and optimizes π. To make each optimization step easier, this paper uses the majorization minimization method [30]. As the surrogate function for Equation (8), this paper uses

$$J_{\text{surr}}(\theta, \pi; \theta_0, \pi_0) = -\eta^{\pi}_{\theta} + b(L^{\pi}(\theta) - h_{\min}) + C = -\mathbb{E}_{(s,a)\sim d^{\pi}_{\mathcal{P}_\theta}} \left[\frac{r(s,a)}{1-\gamma} - b(c_\theta(s,a) - h_{\min}) \right] + C \tag{14}$$

where $b = \frac{B'}{2\sqrt{L^{\pi_0}(\theta_0) - h_{\min}}}$ and $C = b(L^{\pi_0}(\theta_0) - h_{\min})$. Let $\hat{b} = \frac{B'}{2\sqrt{\hat{L}^{\pi_0}(\theta_0) - \hat{h}_{\min}}}$ be an estimate of b, where $\hat{h}_{\min} = \min_{(s,a,s') \in \mathcal{D}} [-\ln \mathcal{P}_{\theta_0}(s'|s,a)]$. Note that $\hat{L}^{\pi_0}(\theta_0)$ is computed using Equation (6) given $(\mathcal{D}, \mathcal{P}_{\theta_0}, w^{\pi_0}_{\theta_0})$, where $w^{\pi_0}_{\theta_0}$ can be estimated using density ratio estimation given $(\mathcal{D}, \hat{\mathcal{D}}^{\pi_0}_{\theta_0})$, and

Algorithm 3 Policy optimization based on weighted model estimation.

1: **Input:** D and (α, B').
2: Initialize $(\mathcal{P}_{\theta_0}, \pi_0)$.
3: **for** $i = 1, 2, \cdots$ **do**
4: **if** $i = 1$ **then**
5: Estimate $\mathcal{P}_{\theta_0}$, using ERM given $\mathcal{D}$.
6: **else**
7: Compute $\hat{b}$, given $(B', \mathcal{D}, \mathcal{P}_{\theta_0}, \pi_0)$.
8: Update $\mathcal{P}_{\theta_0}$, using E-step given $(\mathcal{D}, \hat{b}, \pi_0, \alpha)$.
9: **end if**
10: Estimate $\hat{c}_{\theta_0}$, using ERM given $(\mathcal{D}, \mathcal{P}_{\theta_0})$.
11: Compute $\hat{b}$, given $(B', \mathcal{D}, \mathcal{P}_{\theta_0}, \pi_0)$.
12: Update π_0, using M-step given $(\mathcal{P}_{\theta_0}, \hat{b}, \hat{c}_{\theta_0})$.
13: **end for**

$\hat{\mathcal{D}}_{\theta_0}^{\pi_0}$ can be generated using simulation given $(\mathcal{P}_{\theta_0}, \pi_0)$. That is, $\hat{b}$ can be computed when given $(\mathcal{D}, \mathcal{P}_{\theta_0}, \pi_0)$.

E-step. To update θ while fixing $\pi = \pi_0$, the E-step estimates Equation (14) by

$$-\frac{1}{N}\sum_{n=1}^{N} w_\theta^{\pi_0}(s_n, a_n)\left[\frac{r(s_n, a_n)}{(1-\gamma)} + \hat{b}\ln \mathcal{P}_\theta(s'_n|s_n, a_n)\right], \tag{15}$$

where the constant terms are ignored. Equation (15) can be minimized using Algorithm 1 replacing $\ln \mathcal{P}_\theta(s'_n|s_n, a_n)$ with $\ln \mathcal{P}_\theta(s'_n|s_n, a_n) + \frac{r(s_n, a_n)}{\hat{b}(1-\gamma)}$, or Algorithm 2.

M-step. To update π while fixing $\theta = \theta_0$, the M-step estimates Equation (14) by

$$-\mathbb{E}_{(s,a)\sim d_{\mathcal{P}_{\theta_0}}^{\pi}}\left[\frac{r(s,a)}{1-\gamma} - \hat{b}(\hat{c}_{\theta_0}(s,a) - \hat{h}_{\min})\right], \tag{16}$$

where $\hat{c}_{\theta_0}$ is a model of c_{θ_0}, and C is ignored. This paper trains $\hat{c}_{\theta_0}$ using input data $\{(s_n, a_n)\}_{n=1}^N$ and output data $\{-\ln \mathcal{P}_{\theta_0}(s'_n|s_n, a_n)\}_{n=1}^N$. Note that Equation (16) is the expected return in a MDP, where the transition probability is $\mathcal{P}_{\theta_0}$, and the reward function is

$$r(s,a) - (1-\gamma)\hat{b}(\hat{c}_{\theta_0}(s,a) - \hat{h}_{\min}). \tag{17}$$

Here, $\hat{c}_{\theta_0}(s,a) - \hat{h}_{\min}$ can be seen as an estimate of the policy evaluation error, and $(1-\gamma)\hat{b}$ can be seen as the coefficient to scale the penalty for the error. Since this MDP is known to the agent, Equation (16) can be optimized using any MDP optimization approaches, including RL algorithms.

Algorithm 3 summarizes the procedure described above.

Remark. If $\alpha = 0$, then the E-step reduces to ERM, and Lines 7-8 in Algorithm 3 are skipped. This case can be seen as an instantiation of MOPO [8], where $\hat{c}_{\theta_0} - \hat{h}_{\min}$ and $\hat{b}$ are the admissible error estimator and the penalty coefficient, respectively. That is, the M-step is essentially the same as MOPO, and the E-step is different from MOPO. In other words, MOPO can be seen applying only M-step to optimize the penalized return. Thus, Algorithm 3 improves MOPO in that it optimizes the penalized return with respect to both a model and a policy by the E-step and M-step.

6 Numerical Experiment

6.1 Pendulum

Firstly, to simply illustrate the effectiveness of weighting with the artificial weight, this paper studies policy evaluation on a pendulum environment. State s is the pair of angle and angular velocity, and action a is the torque. Offline data $\mathcal{D}$ is collected using a random policy. Target policy π is a

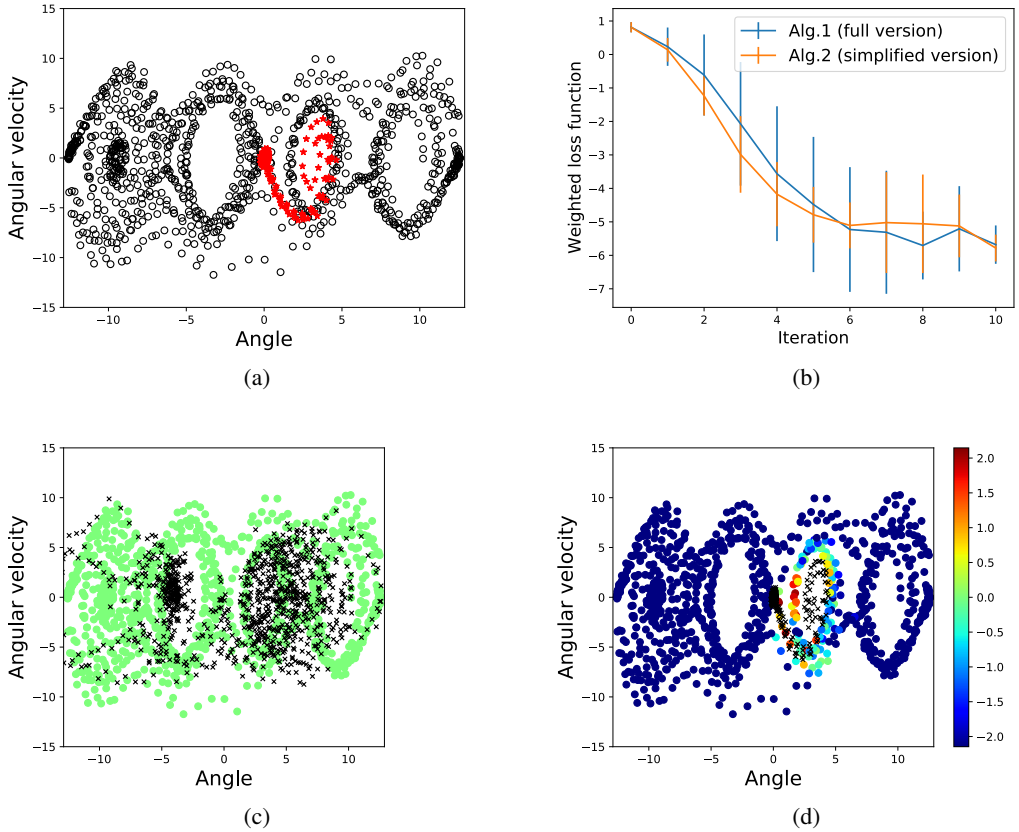

Figure 1: Pendulum task. (a) state part of offline data and real future data, denoted by black circle red star. (b) learning curves of weighted loss function. (c) simulated future states when using ERM. (d) simulated future states when using Algorithm 1. In (c) and (d), black crosses are simulated future states, and colormap show $\log w_\theta^\pi(s_n, a_n)$ for offline data.

controller that stabilizes the pendulum at $(0, 0)$ in the real environment. That is, real future data $\mathcal{D}^\pi_{\mathcal{P}_*}$ is the stabilization behavior towards $(0, 0)$. Figure 1 shows the state part of $\mathcal{D}$ and $\mathcal{D}^\pi_{\mathcal{P}_*}$.

The outline of the implementation is as follows. The agent uses $\mathcal{P}_\theta$ represented by two-layer neural networks with 8 units with tanh activation. The agent estimates w_θ^π using density ratio estimation via logistic regression [31, 11], with normalizing it to satisfy $\frac{1}{N}\sum_n w_\theta^\pi(s_n, a_n) = 1$. The agent estimates z_θ^π using an extension of LSDG [29]. For more details, see Appendix.

Unweighted vs weighted ERM. Figure 1 (c) shows simulated future data $\hat{D}^\pi_{\mathcal{P}_\theta}$ when using the unweighted maximum likelihood method, a typical ERM method. The pendulum is predicted not to stabilize at $(0, 0)$, unlike $D^\pi_{\mathcal{P}_*}$. This result can be described as follows. Based on the unweighted loss function, $\mathcal{P}_\theta$ tries to represent $\mathcal{D}$ as much as possible. However, its expressiveness is insufficient for global prediction. As a result, $\mathcal{P}_\theta$ does not predict the stabilization behavior.

Figure 1 (d) shows $\hat{D}^\pi_{\mathcal{P}_\theta}$ when using Algorithm 1 with $\alpha = 1$. The state part of $\mathcal{D}$ is colored based on the common logarithm of $w_\theta^\pi(s_n, a_n)$ when Algorithm 1 converges. The pendulum is predicted to stabilize at $(0, 0)$, like $D^\pi_{\mathcal{P}_*}$. This situation can be described as follows. When Algorithm 1 converges, Line 3 in Algorithm 1 predicts the stabilization behavior shown in Figure 1. Using the stabilization behavior as $\hat{D}^\pi_{\mathcal{P}_\theta}$, Line 4 estimates $w_\theta^\pi(s_n, a_n)$ shown in Figure 1. After estimating z_θ^π at Line 5, Line 6 estimates $\mathcal{P}_\theta$ to predict more accurately where the weight is larger. Since the expressive capability of $\mathcal{P}_\theta$ is collected near the larger weights, $\mathcal{P}_\theta$ can predict the same stabilization behavior

Table 1: D4RL MuJoCo Benchmark: normalized scores.

dataset	$\alpha = 0$	$\alpha = 0.2$
HalfCheetah-random	50.0 $\pm$ 5.6	45.5 $\pm$ 4.6
HalfCheetah-medium	73.4 $\pm$ 3.0	71.1 $\pm$ 8.6
HalfCheetah-medium-replay	69.7 $\pm$ 2.0	69.8 $\pm$ 1.9
HalfCheetah-medium-expert	88.8 $\pm$ 11.4	85.8 $\pm$ 17.6
Hopper-random	24.4 $\pm$ 9.8	27.6 $\pm$ 10.4
Hopper-medium	100.8 $\pm$ 4.4	105.3 $\pm$ 0.9
Hopper-medium-replay	100.1 $\pm$ 6.7	102.5 $\pm$ 5.9
Hopper-medium-expert	110.5 $\pm$ 1.2	109.8 $\pm$ 0.5
Walker2d-random	17.3 $\pm$ 9.0	9.6 $\pm$ 3.1
Walker2d-medium	81.1 $\pm$ 3.7	43.1 $\pm$ 48.6
Walker2d-medium-replay	81.0 $\pm$ 2.0	84.7 $\pm$ 3.7
Walker2d-medium-expert	79.1 $\pm$ 52.47	108.5 $\pm$ 3.4

again at Line 4 in the next iteration, despite the lack of global expressiveness as mentioned in the paragraph above. This is what happens when Algorithm 1 converges.

The discounted return of π averaged over 100 episodes in the real environment is -15.7. This paper regards this value as the expected return, namely, the goal of policy evaluation. When using the ERM method, the estimated value is -80.7. When using Algorithm 1, the estimated value is -17.8. Thus, Algorithm 1 estimates the expected return more accurately, as it can captures the stabilization behavior.

Simplified version. Figure 1 (b) shows the lerning curves of Algorithm 1 and Algorithm 2, where θ is initialized by the learning result of the ERM method. The results are averaged over 5 runs, and the errorbar is the standard deviation. The simplified version also reduces the weighted loss function, compared to the ERM method, i.e. the start of the iteration. This result suggests that, although the simplified version has the limitation of not converging, it can be used as better than ERM. The simplified version can also predict the stabilization behavior, similarly to the full version, see Appendix.

6.2 D4RL MuJoCo Benchmark

Secondly, to demonstrate weighting with the artificial weight on large-scale environments, this paper studies policy optimization on the D4RL Benchmark [32] based on the MuJoCo simulator [33]. This paper shows the results of Algorithm 3 with $\alpha = 0$ and $\alpha = 0.2$. This paper regards the case of $\alpha = 0$ as a MBRL baseline, as it can be seen as an instantiation of MOPO [8].

The outline of the implementation is as follows. The agent iterates the main loop of Algorithm 3 twice. For the E-step, the agent estimates $\Pr(s', r|s, a)$ instead of $\mathcal{P}(s'|s, a)$, to handle the unknown reward function, as in [8]. the agent represents $\Pr(s', r|s, a)$ using an ensemble of models based on [4, 6, 8], where each model is parametrized as 7-layer neural networks with 256 hidden units. The agent estimates w_θ^π in the same way as Section 6.1. The agent estimates c_θ using L2 norm minimization. For the M-step, the agent improves π using SAC [34], where samples are drawn from $\mathcal{D} \bigcup \hat{\mathcal{D}}_\theta^\pi$. This paper implements SAC by modifying the implementation code by [35]. To stabilize the M-step, the agent terminates a simulation episode when a state or action variable diverges. In addition, the agent clips a penalized reward so that the absolute value is not more than 10 times the maximum absolute value of the reward data. For more details, see Appendix.

Table 1 shows the normalized scores of Algorithm 3, where the results are averaged averaged over 4 runs, and $\pm$ means the stanrard deviations. The major differences are found for Walker2d-medium-expert and Walker2d-medium datasets. The normalized scores of $\alpha = 0.2$ is better for the former and worse for the latter. Below, this paper describes the reasons.

Walker2d-medium-expert dataset. The case of $\alpha = 0.2$ moves forward without falling for all runs, while the case of $\alpha = 0$ falls down for one run. Figure 2 (a) shows the learning curves of real and simulated returns in the M-step at the second iteration for the seed for that run. The case

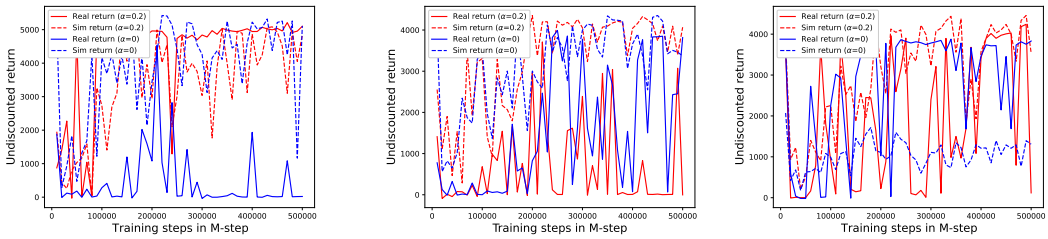

(a) (b) (c)

Figure 2: D4RL MuJoCo Benchmark: (a) learning curves for walker2d-medium-expert. (b,c) learning curves for walker2d-medium.

of $\alpha = 0.2$ obtains similar curves for real and simulation, while the case of $\alpha = 0$ does not. This difference comes from the weighted model estimation. In the case of $\alpha = 0$, the model for the first move is not sufficiently accurate, and this lead to falling at the first move. In the case of $\alpha = 0.2$, simulated future data includes predictions of the first move, and similar parts of offline data are weighted with large weights. Based on the weights, the estimated model more accurately predicts the first move. As a result, the case of $\alpha = 0.2$ can learn to move forward without falling, based on the more accurate prediction. That is, Algorithm 3 improves the performance based on weighted model estimation. For the other runs, although the case of $\alpha = 0$ moves forward without falling in a real environment, it obtains real and simulation returns that are relatively different, see Appendix.

Walker2d-medium dataset. The case of $\alpha = 0.2$ falls down for two runs, while the case of $\alpha = 0$ moves forward without falling. Figure 2 (b,c) shows the learning curves fo those runs. In Figure 2 (b), for both $\alpha = 0$ and $\alpha = 0.2$, the simulation curves do not seem to capture the real returns. The difference of the score seems due to the high variance of the real returns. In Figure 2 (c), the case of $\alpha = 0.2$ suffers from the high variance, though the simulation return captures the real return to some extent. For $\alpha = 0$, the real return is good, but it cannot be predicted by the simulation return. Note that the real return is unknown to the agent before applying a learned policy. Thus, this paper does not find advantages for $\alpha = 0$ from these results.

7 Limitations

This paper has limitations common to standard offline MBRL approaches, such as Markov assumption and reward design. The limitations specific to this paper are as follows. This paper derives Algorithm 1 based on the theory, but it cannot be applied to large-scale problems due to the large amount of computations. This paper also presents Algorithm 2 as a practical simplified version, but it has no convergence with respect to Equation (8), see Section 5.1.

8 Conclusion

This paper discusses weighted model estimation considering covariate shift in offline MBRL. The main idea of this paper is weighting with the artificial weight, which is relatively easy to estimate. Based on this idea, this paper defines a loss function for offline MBRL and presents algorithms to optimize it. This paper justifies this idea as evaluating the upper bound of the policy evaluation error. This paper demonstrates the effectiveness of this idea in numerical experiments. The main contribution of this paper is showing the validity of weighting with the artificial weight.

An interesting issue is extending to Bayesian MBRL [36, 37]. In general, the posterior distribution under covariate shift can be defined using a weighted likelihood function [9]. In offline MBRL, the likelihood function weighted with the artificial weight is effective, as shown in this paper. Combining these ideas will introduce a weighted posterior considering covariate shift into Bayesian MBRL.

References

[1] Richard S Sutton, Andrew G Barto, et al. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge, 1998.

[2] Christopher G Atkeson and Juan Carlos Santamaria. A comparison of direct and model-based reinforcement learning. In *International Conference on Robotics and Automation*, volume 4, pages 3557–3564. IEEE, 1997.

[3] Marc Peter Deisenroth, Dieter Fox, and Carl Edward Rasmussen. Gaussian processes for data-efficient learning in robotics and control. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2):408–423, 2013.

[4] Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Neural Information Processing Systems*, pages 4759–4770, 2018.

[5] Sergey Levine, Aviral Kumar, George Tucker, and Justin Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.

[6] Michael Janner, Justin Fu, Marvin Zhang, and Sergey Levine. When to trust your model: Model-based policy optimization. In *Neural Information Processing Systems*, volume 32, 2019.

[7] Rahul Kidambi, Aravind Rajeswaran, Praneeth Netrapalli, and Thorsten Joachims. Morel: Model-based offline reinforcement learning. In *Neural Information Processing Systems*, volume 33, pages 21810–21823, 2020.

[8] Tianhe Yu, Garrett Thomas, Lantao Yu, Stefano Ermon, James Y Zou, Sergey Levine, Chelsea Finn, and Tengyu Ma. Mopo: Model-based offline policy optimization. In *Neural Information Processing Systems*, volume 33, pages 14129–14142, 2020.

[9] Hidetoshi Shimodaira. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of Statistical Planning and Inference*, 90(2):227–244, 2000.

[10] Ofir Nachum, Yinlam Chow, Bo Dai, and Lihong Li. Dualdice: Behavior-agnostic estimation of discounted stationary distribution corrections. In *Neural Information Processing Systems*, volume 32, 2019.

[11] Masashi Sugiyama, Taiji Suzuki, and Takafumi Kanamori. *Density ratio estimation in machine learning*. Cambridge University Press, 2012.

[12] Arthur Argenson and Gabriel Dulac-Arnold. Model-based offline planning. In *International Conference on Learning Representations*, 2021.

[13] Doina Precup, Richard S Sutton, and Satinder Singh. Eligibility traces for off-policy policy evaluation. In *International Conference on Machine Learning*, pages 759–766, 2000.

[14] Assaf Hallak and Shie Mannor. Consistent on-line off-policy evaluation. In *International Conference on Machine Learning*, pages 1372–1383, 2017.

[15] Carles Gelada and Marc G Bellemare. Off-policy deep reinforcement learning by bootstrapping the covariate shift. In *AAAI Conference on Artificial Intelligence*, volume 33, pages 3647–3655, 2019.

[16] Nan Jiang and Lihong Li. Doubly robust off-policy value evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 652–661. PMLR, 2016.

[17] Philip Thomas and Emma Brunskill. Data-efficient off-policy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 2139–2148, 2016.

[18] Mehrdad Farajtabar, Yinlam Chow, and Mohammad Ghavamzadeh. More robust doubly robust off-policy evaluation. In *International Conference on Machine Learning*, pages 1447–1456, 2018.

[19] Amir-Massoud Farahmand, Andre Barreto, and Daniel Nikovski. Value-aware loss function for model-based reinforcement learning. In *Artificial Intelligence and Statistics*, pages 1486–1494, 2017.

[20] Amir-Massoud Farahmand. Iterative value-aware model learning. In *Neural Information Processing Systems*, pages 9072–9083, 2018.

[21] Yuping Luo, Huazhe Xu, Yuanzhi Li, Yuandong Tian, Trevor Darrell, and Tengyu Ma. Algorithmic framework for model-based deep reinforcement learning with theoretical guarantees. In *International Conference on Learning Representations*, 2019.

[22] Yao Liu, Omer Gottesman, Aniruddh Raghu, Matthieu Komorowski, Aldo A Faisal, Finale Doshi-Velez, and Emma Brunskill. Representation balancing MDPs for off-policy policy evaluation. In *Neural Information Processing Systems*, pages 2644–2653, 2018.

[23] Romina Abachi, Mohammad Ghavamzadeh, and Amir-massoud Farahmand. Policy-aware model learning for policy gradient methods. *arXiv preprint arXiv:2003.00030*, 2020.

[24] Yinlam Chow, Brandon Cui, MoonKyung Ryu, and Mohammad Ghavamzadeh. Variational model-based policy optimization. *arXiv preprint arXiv:2006.05443*, 2020.

[25] Sergey Levine. Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909*, 2018.

[26] Jian Shen, Han Zhao, Weinan Zhang, and Yong Yu. Model-based policy optimization with unsupervised model adaptation. In *Neural Information Processing Systems*, 2020.

[27] Byung-Jun Lee, Jongmin Lee, and Kee-Eung Kim. Representation balancing offline model-based reinforcement learning. In *International Conference on Learning Representations*, 2021.

[28] Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 1994.

[29] Tetsuro Morimura, Eiji Uchibe, Junichiro Yoshimoto, Jan Peters, and Kenji Doya. Derivatives of logarithmic stationary distributions for policy gradient reinforcement learning. *Neural Computation*, 22(2):342–376, 2010.

[30] Kenneth Lange, David R Hunter, and Ilsoon Yang. Optimization transfer using surrogate objective functions. *Journal of Computational and Graphical Statistics*, 9(1):1–20, 2000.

[31] Jing Qin. Inferences for case-control and semiparametric two-sample density ratio models. *Biometrika*, 85(3):619–630, 1998.

[32] Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. D4RL: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*, 2020.

[33] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012.

[34] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning*, pages 1861–1870. PMLR, 2018.

[35] Pranjal Tandon. pytorch-soft-actor-critic. `https://github.com/pranz24/pytorch-soft-actor-critic`, 2018.

[36] Michael O'Gordon Duff. *Optimal Learning: Computational procedures for Bayes-adaptive Markov decision processes*. PhD thesis, University of Massachusetts Amherst, 2002.

[37] Mohammad Ghavamzadeh, Shie Mannor, Joelle Pineau, and Aviv Tamar. Bayesian reinforcement learning: A survey. *Foundations and Trends® in Machine Learning*, 8(5-6):359–483, 2015.

[38] Tom Le Paine, Cosmin Paduraru, Andrea Michi, Caglar Gulcehre, Konrad Zolna, Alexander Novikov, Ziyu Wang, and Nando de Freitas. Hyperparameter selection for offline reinforcement learning. *arXiv preprint arXiv:2007.09055*, 2020.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to [Yes] , [No] , or [N/A] . You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? [Yes] See Section **??**.
- Did you include the license to the code and datasets? [No] The code and the data are proprietary.
- Did you include the license to the code and datasets? [N/A]

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] See Section 7.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] The possibility of negative societal impacts is essentially the same with existing offline MBRL approaches [8], or other automation research. That is, there is still a risk of deploying a learned policy that may behave unpredictably in high-risk domains. This paper discusses improving predictive performance, but it does not guarantee guarantee safety. Much additional work is needed.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See the supplemental material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See the supplemental material.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] The $\pm$ in the Table means standard deviation, averaged over 5 random seeds.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See the supplemental material.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] We use D4RL dataset [32]. We use a modification of pytorch SAC implementation [35].
   (b) Did you mention the license of the assets? [Yes] D4RL datasets are licensed under the Creative Commons Attribution 4.0 License (CC BY), and D4RL code is licensed under the Apache 2.0 License [32]. Pytorch SAC implementation [35] is under MIT License.
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] See the zip file of the code.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]