
A Unified Framework for Alternating Offline Model Training and Policy Learning

Shentao Yang¹, Shujian Zhang¹, Yihao Feng², Mingyuan Zhou¹

¹The University of Texas at Austin ²Salesforce Research

shentao.yang@mcombs.utexas.edu szhang19@utexas.edu

yihao@salesforce.com mingyuan.zhou@mcombs.utexas.edu

Abstract

In offline model-based reinforcement learning (offline MBRL), we learn a dynamic model from historically collected data, and subsequently utilize the learned model and fixed datasets for policy learning, without further interacting with the environment. Offline MBRL algorithms can improve the efficiency and stability of policy learning over the model-free algorithms. However, in most of the existing offline MBRL algorithms, the learning objectives for the dynamic models and the policies are isolated from each other. Such an *objective mismatch* may lead to inferior performance of the learned agents. In this paper, we address this issue by developing an iterative offline MBRL framework, where we maximize a lower bound of the true expected return, by alternating between dynamic-model training and policy learning. With the proposed unified model-policy learning framework, we achieve competitive performance on a wide range of continuous-control offline reinforcement learning datasets. Source code is released at https://github.com/Shentao-YANG/AMPL_NeurIPS2022.

1 Introduction

Offline reinforcement learning (offline RL) [1, 2], where the agents are trained from static and pre-collected datasets, avoids direct interactions with the underlying real-world environment during the learning process. Unlike traditional online RL, whose success largely depends on simulator-based trial-and-error [e.g., 3, 4], offline RL enables training policies for real-world applications, where it is infeasible or even risky to collect online experimental data, such as robotics, advertisement, or dialog systems [e.g., 5–8]. Though promising, it remains challenging to train agents under the offline setting, due to the discrepancy between the distribution of the offline data and the state-action distribution induced by the current learning policy. With such a discrepancy, directly transferring standard online off-policy RL methods [e.g., 9–11] to the offline setting tends to be problematic [12, 13], especially when the offline data cannot sufficiently cover the state-action space [14]. To tackle this issue in offline RL, recent works [e.g., 15–17] propose to approximate the policy-induced state-action distribution by leveraging a learned dynamic model to draw imaginary rollouts. These additional synthetic rollouts help mitigate the distributional discrepancy and stabilize the policy-learning algorithms under the offline setting.

Most of the prior offline model-based RL (MBRL) methods [e.g., 16, 18–21], however, first *pretrain* a one-step forward dynamic model via maximum likelihood estimation (MLE) on the offline dataset, and then use the learned model to train the policy, without *further improving the dynamic model* during the policy learning process. As a result, the objective function used for model training (e.g., MLE) and the objective of model utilization are *unrelated* with each other. Specifically, the model is trained to be “simply a mimic of the world,” but is used to improve the performance of the learned policy [22–24]. Though such a training paradigm is historically rooted [25, 26], this issue of *objective*

mismatch in the model training and model utilization has been identified as problematic in recent works [23, 24]. In offline MBRL, this issue is exacerbated, since the learned model can hardly be globally accurate, due to the limited amount of offline data and the complexity of the control tasks.

Motivated by the objective-mismatch issue, we develop an iterative offline MBRL method, alternating between training the dynamic model and the policy to maximize a lower bound of the true expected return. This lower bound, leading to a weighted MLE objective for the dynamic-model training, is relaxed to a tractable regularized objective for the policy learning. To train the dynamic model by the proposed objective, we need to estimate the marginal importance weights (MIW) between the offline-data distribution and the stationary state-action distribution of the current policy [27, 28]. This estimation tends to be unstable by standard approaches [e.g., 29, 30], which require saddle-point optimization. Instead, we propose a simple yet stable fixed-point-style method for MIW estimation, which can be directly incorporated into our alternating training framework. With these considerations, our method, offline Alternating Model-Policy Learning (AMPL), performs competitively on a wide range of continuous-control offline RL datasets in the D4RL benchmark [31]. These empirical results and ablation study show the efficacy of our proposed algorithmic designs.

2 Background

Markov decision process and offline RL. A Markov decision process (MDP) is denoted by $\mathcal{M} = (\mathbb{S}, \mathbb{A}, P, r, \gamma, \mu_0)$, where \mathbb{S} is the state space, \mathbb{A} the action space, $P(s' | s, a) : \mathbb{S} \times \mathbb{S} \times \mathbb{A} \rightarrow [0, 1]$ the environmental dynamic, $r(s, a) : \mathbb{S} \times \mathbb{A} \rightarrow [-r_{max}, r_{max}]$ the reward function, $\gamma \in [0, 1)$ the discount factor, and $\mu_0(s) : \mathbb{S} \rightarrow [0, 1]$ the initial state-distribution.

For any policy $\pi(a | s)$, we denote its state-action distribution at timestep $t \geq 0$ as $d_{\pi,t}^P(s, a) \triangleq \Pr(s_t = s, a_t = a | s_0 \sim \mu_0, a_t \sim \pi, s_{t+1} \sim P, \forall t \geq 0)$. The (discounted) stationary state-action distribution of π is denoted as $d_{\pi,\gamma}^P(s, a) \triangleq (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t d_{\pi,t}^P(s, a)$.

Denote $Q_{\pi}^P(s, a) = \mathbb{E}_{\pi,P} [\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | s_0 = s, a_0 = a]$ as the action-value function of policy π under the dynamic P . The goal of RL is to find a policy π maximizing the expected return

$$J(\pi, P) \triangleq (1 - \gamma) \mathbb{E}_{s \sim \mu_0, a \sim \pi(\cdot | s)} [Q_{\pi}^P(s, a)] = \mathbb{E}_{(s,a) \sim d_{\pi,\gamma}^P} [r(s, a)]. \quad (1)$$

In offline RL, the policy π and critic Q_{π}^P are typically approximated by parametric functions π_{ϕ} and Q_{θ} , respectively, with parameters ϕ and θ . The critic Q_{θ} is trained by the Bellman backup

$$\arg \min_{\theta} \mathbb{E}_{(s,a,r,s') \sim \mathcal{D}_{\text{env}}} \left[\left(Q_{\theta}(s, a) - \left(r(s, a) + \gamma \mathbb{E}_{a' \sim \pi_{\phi}(\cdot | s')} [Q_{\theta'}(s', a')] \right) \right)^2 \right], \quad (2)$$

where $Q_{\theta'}$ is the target network [12, 13]. The actor π_{ϕ} is trained in the policy improvement step by

$$\arg \max_{\phi} \mathbb{E}_{s \sim \mathcal{D}_{\text{env}}, a \sim \pi_{\phi}(\cdot | s)} [Q_{\theta}(s, a)], \quad (3)$$

where \mathcal{D}_{env} denotes the offline dataset drawn from $d_{\pi_b,\gamma}^P$ [2, 32], with π_b being the behavior policy.

Offline model-based RL. In offline model-based RL algorithms, the true environmental dynamic P^* is typically approximated by a parametric function $\hat{P}(s' | s, a)$ in some function class \mathcal{P} . With the offline dataset \mathcal{D}_{env} , \hat{P} is trained via the MLE [15, 16, 18] as

$$\arg \max_{\hat{P} \in \mathcal{P}} \mathbb{E}_{(s,a,s') \sim \mathcal{D}_{\text{env}}} \left[\log \hat{P}(s' | s, a) \right]. \quad (4)$$

Similarly, the reward function can be approximated by a parametric model \hat{r} if assumed unknown. With \hat{P} and \hat{r} , the true MDP \mathcal{M} can be approximated by $\hat{\mathcal{M}} = (\mathbb{S}, \mathbb{A}, \hat{P}, \hat{r}, \gamma, \mu_0)$. We further define $d_{\pi,\gamma}^{P^*}(s, a)$ as the stationary state-action distribution induced by π on P^* (or MDP \mathcal{M}), and $d_{\pi,\gamma}^{\hat{P}}(s, a)$ as that on the learned dynamic \hat{P} (or MDP $\hat{\mathcal{M}}$). We approximate $d_{\pi,\gamma}^{P^*}$ by simulating π_{ϕ} on $\hat{\mathcal{M}}$ for a short horizon h starting from state $s \in \mathcal{D}_{\text{env}}$, as in prior work [e.g., 16, 18, 19, 21]. The resulting transitions are stored in a replay buffer $\mathcal{D}_{\text{model}}$, constructed similar to the off-policy RL [33, 9]. To better approximate $d_{\pi,\gamma}^{P^*}$, sampling from \mathcal{D}_{env} in Eqs. (2) and (3) is commonly replaced by sampling from the augmented dataset $\mathcal{D} = f\mathcal{D}_{\text{env}} + (1 - f)\mathcal{D}_{\text{model}}$, $f \in [0, 1]$, denoting sampling from \mathcal{D}_{env} and $\mathcal{D}_{\text{model}}$ with probabilities f and $1 - f$, respectively. We follow Yu et al. [18] to use $f = 0.5$.

3 Offline alternating model-policy learning

Our goal is to derive the objectives for both dynamic-model training and policy learning from a principled perspective. A natural idea is to build a tractable lower bound for $J(\pi, P^*)$, the expected return of the policy π under the true dynamic P^* , and then alternate between training the policy π and the dynamic model \hat{P} to maximize this lower bound. Indeed, we can construct a lower bound as

$$J(\pi, P^*) \geq J(\pi, \hat{P}) - |J(\pi, P^*) - J(\pi, \hat{P})|, \quad (5)$$

where $J(\pi, \hat{P})$ is the expected return of policy π under the learned model \hat{P} . From the right hand side (RHS) of Eq. (5), if the policy evaluation error $|J(\pi, P^*) - J(\pi, \hat{P})|$ is small, $J(\pi, \hat{P})$ will be a good proxy for the true expected return $J(\pi, P^*)$. We can empirically estimate $J(\pi, \hat{P})$ via \hat{P} and π .

Further, if a tractable upper bound for $|J(\pi, P^*) - J(\pi, \hat{P})|$ can be constructed, it can serve as a unified training objective for both dynamic model \hat{P} and policy π . We can then alternate between optimizing the dynamic model \hat{P} and the policy π to maximize the lower bound of $J(\pi, P^*)$, *i.e.*, simultaneously minimizing the upper bound of the evaluation error $|J(\pi, P^*) - J(\pi, \hat{P})|$. This gives us an *iterative, maximization-maximization* algorithm for model and policy learning.

The following theorem indicates a tractable upper bound for $|J(\pi, P^*) - J(\pi, \hat{P})|$, which can be subsequently relaxed for model training and policy learning.

Theorem 1. *Let P^* be the true dynamic and \hat{P} be the approximate dynamic model. Suppose the reward function $|r(s, a)| \leq r_{\max}$, then we have*

$$\left| J(\pi, P^*) - J(\pi, \hat{P}) \right| \leq \frac{\gamma \cdot r_{\max}}{\sqrt{2}(1-\gamma)} \cdot \sqrt{D_{\pi}(P^*, \hat{P})},$$

$$\text{with } D_{\pi}(P^*, \hat{P}) \triangleq \mathbb{E}_{(s,a) \sim d_{\pi_b, \gamma}^{P^*}} \left[\omega(s, a) \text{KL} \left(P^*(s' | s, a) \pi_b(a' | s') \parallel \hat{P}(s' | s, a) \pi(a' | s') \right) \right],$$

where π_b is the behavior policy, $d_{\pi_b, \gamma}^{P^*}$ is the offline-data distribution, and $\omega(s, a) \triangleq \frac{d_{\pi, \gamma}^{P^*}(s, a)}{d_{\pi_b, \gamma}^{P^*}(s, a)}$ is the marginal importance weight (MIW) between the offline-data distribution and the stationary state-action distribution of the policy π [27, 29].

Detailed proof of Theorem 1 can be found in Appendix B.2.

The KL term in $D_{\pi}(P^*, \hat{P})$ indicates the following two principles for model and policy learning:

- ★ For the dynamic model, $s' \sim \hat{P}(\cdot | s, a)$ should be close to the true next state $\tilde{s}' \sim P^*(\cdot | s, a)$, with (s, a) pairs drawn from the stationary state-action distribution of the policy π . Since we cannot directly draw samples from $d_{\pi, \gamma}^{P^*}$, we reweight the offline data with $\omega(s, a)$. This leads to a weighted KL minimization objective for the model training.
- ★ For the policy π , the KL term indicates a regularization term, that the tuple (s', a') from the joint conditional distribution $\hat{P}(s' | s, a) \pi(a' | s')$ should be close to the tuple (\tilde{s}', \tilde{a}') from $P^*(\tilde{s}' | s, a) \pi_b(\tilde{a}' | \tilde{s}')$. (\tilde{s}', \tilde{a}') is simply a sample from the offline dataset.

Based on the above observations, we can fix π and train the dynamic model \hat{P} by minimizing $D_{\pi}(P^*, \hat{P})$ w.r.t. \hat{P} . Similarly, we can fix the dynamic model \hat{P} and learn a policy π to maximize the lower bound of $J(\pi, P^*)$. This alternating training scheme provides a unified approach for model and policy learning. In the following sections, we discuss how to optimize the dynamic model \hat{P} , the policy π , and the MIW ω under our alternating training framework.

3.1 Dynamic model training

Expanding the KL term in $D_{\pi}(P^*, \hat{P})$, we have

$$D_{\pi}(P^*, \hat{P}) = \overbrace{\mathbb{E}_{(s,a,s',a') \sim d_{\pi_b, \gamma}^{P^*}} \left[\omega(s, a) (\log P^*(s' | s, a) + \log \pi_b(a' | s') - \log \pi(a' | s')) \right]}^{\triangleq \textcircled{1}} - \mathbb{E}_{(s,a,s') \sim d_{\pi_b, \gamma}^{P^*}} \left[\omega(s, a) \log \hat{P}(s' | s, a) \right],$$

where the tuple (s, a, s', a') is simply two consecutive state-action pairs in the offline dataset. Further, if the policy π is fixed, the term ① is a constant *w.r.t.* \hat{P} . Thus, given the MIW ω , we can optimize \hat{P} by minimizing the following loss

$$\ell(\hat{P}) \triangleq -\mathbb{E}_{(s,a,s') \sim d_{\pi_b, \gamma}^{P^*}} \left[\omega(s, a) \log \hat{P}(s' | s, a) \right], \quad (6)$$

which is an MLE objective weighted by $\omega(s, a)$. We discuss how to estimate $\omega(s, a)$ in Section 3.3.

3.2 Policy learning

The lower bound for $J(\pi, P^*)$ implied by Theorem 1 is

$$J(\pi, \hat{P}) - \frac{\gamma \cdot r_{\max}}{\sqrt{2}(1-\gamma)} \cdot \sqrt{D_{\pi}(P^*, \hat{P})}, \quad (7)$$

where $J(\pi, \hat{P})$ can be estimated via the action-value function similar to standard offline MBRL algorithms [e.g., 16, 18, 19]. Thus, when the dynamic model \hat{P} is fixed, the main difficulty is to estimate the regularizer $D_{\pi}(P^*, \hat{P})$ for the policy π .

When the policy π is Gaussian, direct estimation of $D_{\pi}(P^*, \hat{P})$ is possible. Empirically, however, it is helpful to learn the policy π in the class of implicit distribution, which is a richer distribution class and can better maximize the action-value function. Specifically, given a noise distribution $p_z(z)$, action $a = \pi_{\phi}(s, z)$ with $z \sim p_z(\cdot)$, where π_{ϕ} is a deterministic network.

Unfortunately, we can not directly estimate the KL term in $D_{\pi}(P^*, \hat{P})$ if π is an implicit policy, since we can only draw samples from π but the density is unknown. A potential solution is to use the dual representation of KL divergence $\text{KL}(p || q) = \sup_T \mathbb{E}_p[T] - \log(\mathbb{E}_q[e^T])$ [34], which can be estimated with samples from the distributions p and q . However, the exponential function therein makes the estimation unstable in practice [30]. We instead use the dual representation of the Jensen–Shannon divergence (JSD) to approximate $D_{\pi}(P^*, \hat{P})$, which can be approximately minimized using the GAN structure [35, 36]. Our framework can thus utilize the many stabilization techniques developed in the GAN community (Appendix E.2.1).

Besides, we remove the MIW $\omega(s, a)$ during the policy training since we do not observe its empirical benefits, which will be discussed in Section 4.2. Further applying the replacement of KL with JSD and ignoring the $\sqrt{\cdot}$ for numerical stability, we get an *approximated new regularization* for policy π :

$$\tilde{D}_{\pi}(P^*, \hat{P}) \triangleq \text{JSD} \left(P^*(s' | s, a) \pi_b(a' | s') d_{\pi_b, \gamma}^{P^*}(s, a) || \hat{P}(s' | s, a) \pi(a' | s') d_{\pi_b, \gamma}^{P^*}(s) \pi(a | s) \right). \quad (8)$$

Informally speaking, Eq. (8) regularizes the imaginary rollouts of π on \hat{P} towards state-action pairs from the offline dataset. Intuitively, $\tilde{D}_{\pi}(P^*, \hat{P})$ is a more effective regularizer for policy training than the original $D_{\pi}(P^*, \hat{P})$, since $\tilde{D}_{\pi}(P^*, \hat{P})$ regularizes action choices at both s and s' . Appendix B.3 discusses how we move from $D_{\pi}(P^*, \hat{P})$ to $\tilde{D}_{\pi}(P^*, \hat{P})$ in detail.

3.3 Marginal importance weight training

A number of methods have been recently proposed to estimate the marginal importance weight ω [27, 29, 30]. These methods typically require solving a complex saddle-point optimization, casting doubts on their training stability especially when combined with policy learning on continuous-control offline MBRL problems. In this section, we mimic the Bellman backup to derive a fixed-point-style method for estimating the MIW.

Denote the true MIW as $\omega^*(s, a) \triangleq \frac{d_{\pi, \gamma}^{P^*}(s, a)}{d_{\pi_b, \gamma}^{P^*}(s, a)}$, we have $d_{\pi_b, \gamma}^{P^*}(s, a) \cdot \omega^*(s, a) = d_{\pi, \gamma}^{P^*}(s, a)$. Expanding the RHS, $\forall s', a'$,

$$d_{\pi_b, \gamma}^{P^*}(s', a') \omega^*(s', a') = \gamma \sum_{s, a} \pi(a' | s') P^*(s' | s, a) \omega^*(s, a) d_{\pi_b, \gamma}^{P^*}(s, a) + (1 - \gamma) \mu_0(s') \pi(a' | s'). \quad (9)$$

The derivation is deferred to Appendix B.4. Therefore, a “Bellman equation” for $\omega(s', a')$ is

$$\begin{aligned} \omega(s', a') &= \mathcal{T} \omega(s', a'), \\ \mathcal{T} \omega(s', a') &\triangleq \frac{\gamma \sum_{s, a} \pi(a' | s') P^*(s' | s, a) \omega(s, a) d_{\pi_b, \gamma}^{P^*}(s, a) + (1 - \gamma) \mu_0(s') \pi(a' | s')}{d_{\pi_b, \gamma}^{P^*}(s', a')}. \end{aligned}$$

Here \mathcal{T} can be viewed as the “Bellman operator” for ω . The update iterate defined by \mathcal{T} has the following convergence property, which is proved in Appendix B.5.

Proposition 2. *On finite state-action space, if the current policy π is close to the behavior policy π_b , then the iterate for ω defined by \mathcal{T} converges geometrically.*

The assumption that π is close to π_b coincides with the regularization term in the policy-learning objective discussed in Section 3.2.

Unfortunately, the RHS of Eq. (9) is not estimable since we do not know the density values therein. We therefore multiply both sides of Eq. (9) by some test function and subsequently sum over (s', a') on both sides to get a tractable objective that only requires samples from the offline dataset and the initial state-distribution μ_0 . It is desired to choose a test function that can better distinguish the difference between the left-hand side (LHS) and the RHS of Eq. (9). A potential choice is the action-value function of the policy π , due to some primal-dual relationship between the stationary state-action density-(ratio) ($d_{\pi, \gamma}^{P^*}(s, a)$ or $\omega^*(s, a)$) and the action-value function [37–39]. A detailed discussion on the choice of the test function is provided in Appendix C.

Practically we use $Q_{\pi}^{\hat{P}}$ as the test function. Note that multiplying both sides of Eq. (9) by the same $Q_{\pi}^{\hat{P}}$ does not undermine the convergence property, under mild conditions on Q . Mimicking the Bellman backup to sum over (s', a') on both sides, with the notation $d_{\pi_b, \gamma}^{P^*}(s, a, s') = d_{\pi_b, \gamma}^{P^*}(s, a)P^*(s' | s, a)$,

$$\overbrace{\mathbb{E}_{(s,a) \sim d_{\pi_b, \gamma}^{P^*}} [\omega^*(s, a) \cdot Q_{\pi}^{\hat{P}}(s, a)]}^{\ell_1(\omega^*)} = \overbrace{\gamma \mathbb{E}_{\substack{(s,a,s') \sim d_{\pi_b, \gamma}^{P^*} \\ a' \sim \pi(\cdot | s')}} [\omega^*(s, a) \cdot Q_{\pi}^{\hat{P}}(s', a')] + (1 - \gamma) \mathbb{E}_{\substack{s \sim \mu_0(\cdot) \\ a \sim \pi(\cdot | s)}} [Q_{\pi}^{\hat{P}}(s, a)]}^{\ell_2(\omega^*)}. \quad (10)$$

Thus for a given ω , we can optimize ω by minimizing the difference between the RHS and the LHS of Eq. (10). For training stability, we use a target network $\omega'(s, a)$ for the RHS, and the final objective for learning ω is

$$(\ell_1(\omega) - \ell_2(\omega'))^2, \quad (11)$$

where the target network $\omega'(s, a)$ is soft-updated after each gradient step, motivated by Q_{θ_j} and $\pi_{\phi'}$.

Our proposed training method is closely related to VPM [40]. By using the MIW $\omega(s, a)$ itself as the test function, VPM leverages the variational power iteration to train MIW iteratively. Instead, our approach uses the current action-value function as the test function, motivated by the primal-dual relationship between the MIW and the action-value function in off-policy evaluation. We compare the empirical performance of several alternative approaches in Section 4.2 and in Table 3 of Appendix A.

3.4 Practical implementation

In this section we briefly discuss some implementation details of our offline Alternating Model-Policy Learning (AMPL) method, whose main steps are in Algorithm 1. Further details are in Appendix E.1.

Dynamic model training. We adopt common practice in offline MBRL [e.g., 16, 18] to use an ensemble of Gaussian probabilistic networks $\hat{P}(\cdot | s, a)$ and $\hat{r}(s, a)$ to parameterize the stochastic transition and reward. We initialize the dynamic model by standard MLE training, and periodically update the model by minimizing Eq. (6).

Critic training. We use the conservative target in the offline RL literature [e.g., 12, 13]:

$$\tilde{Q}(s, a) \triangleq r(s, a) + \gamma \mathbb{E}_{a' \sim \pi_{\phi'}(\cdot | s')} [c \min_{j=1,2} Q_{\theta_j}(s', a') + (1 - c) \max_{j=1,2} Q_{\theta_j}(s', a')], \quad (12)$$

where we set $c = 0.75$. With mini-batch \mathcal{B} sampled from the augmented dataset \mathcal{D} , both critic networks are trained as

$$\forall j = 1, 2, \quad \arg \min_{\theta_j} \frac{1}{|\mathcal{B}|} \sum_{(s,a) \in \mathcal{B}} \text{Huber}(Q_{\theta_j}(s, a), \tilde{Q}(s, a)), \quad (13)$$

where the Huber loss $\text{Huber}(\cdot)$ is used in lieu of the classical MSE for training stability [41].

Estimating $\tilde{D}_{\pi}(P^*, \hat{P})$. In $\tilde{D}_{\pi}(P^*, \hat{P})$, using the notations of GAN, we denote the sample from the left distribution of JSD in Eq. (8) by $\mathcal{B}_{\text{true}}$ (i.e., “true” sample), and the sample from the right distribution of JSD by $\mathcal{B}_{\text{fake}}$ (i.e., “fake” sample).

Algorithm 1 Main steps of AMPL.

Initialize: Dynamic model \hat{P} and \hat{r} , policy π_ϕ , critics Q_{θ_1} and Q_{θ_2} , discriminator D_ψ , MIW ω .
Initialize \hat{P} and \hat{r} via the MLE (Eq. (4)).
for iteration $\in \{1, \dots, \text{total_iterations}\}$ **do**
 if iteration $\% \text{model_retrain_period} == 0$ **then**
 Estimate $\omega(s, a)$ via Eq. (11); train \hat{P} and \hat{r} by weighted MLE (Eq. (6)) with $\omega(s, a)$.
 end if
 Rollout synthetic data with π_ϕ , \hat{P} and \hat{r} , and add the data to $\mathcal{D}_{\text{model}}$.
 Sample mini-batch $\mathcal{B} \sim \mathcal{D} = f\mathcal{D}_{\text{env}} + (1 - f)\mathcal{D}_{\text{model}}$.
 Optimize Q_{θ_1} , Q_{θ_2} via Eqs. (12) – (13).
 Train D_ψ to maximize Eq. (16).
 Optimize π_ϕ by Eq. (17).
end for

The “true” sample $\mathcal{B}_{\text{true}}$ consists of samples from \mathcal{D}_{env} . The “fake” sample $\mathcal{B}_{\text{fake}}$ is formed by first sampling $s \sim \mathcal{D}$, followed by $a \sim \pi_\phi(\cdot | s)$, $s' \sim \hat{P}(\cdot | s, a)$, $a' \sim \pi_\phi(\cdot | s')$. Concretely, the “fake” sample $\mathcal{B}_{\text{fake}}$, generator loss $\mathcal{L}_g(\phi)$, and the discriminator loss $\mathcal{L}_D(\psi)$ can be described as

$$\mathcal{B}_{\text{fake}} \triangleq \left[\begin{matrix} (s, & a) \\ (s', & a') \end{matrix} \right], \quad (14) \quad \mathcal{L}_g(\phi) \triangleq \frac{1}{|\mathcal{B}_{\text{fake}}|} \sum_{(s,a) \in \mathcal{B}_{\text{fake}}} [\log(1 - D_\psi(s, a))], \quad (15)$$

$$\mathcal{L}_D(\psi) \triangleq \frac{1}{|\mathcal{B}_{\text{true}}|} \sum_{(s,a) \sim \mathcal{B}_{\text{true}}} [\log D_\psi(s, a)] + \frac{1}{|\mathcal{B}_{\text{fake}}|} \sum_{(s,a) \sim \mathcal{B}_{\text{fake}}} [\log(1 - D_\psi(s, a))], \quad (16)$$

where $\mathcal{L}_g(\phi)$ is the empirical policy-learning regularizer implied by $\tilde{D}_\pi(P^*, \hat{P})$.

Policy training. Since the constant multiplier in Eq. (7) is unknown, we treat it as a hyperparameter. Adding the regularization term of policy training, our policy optimization objective is

$$\arg \min_\phi -\lambda \cdot \frac{1}{|\mathcal{B}|} \sum_{s \in \mathcal{B}, a \sim \pi_\phi(\cdot | s)} [\min_{j=1,2} Q_{\theta_j}(s, a)] + \mathcal{L}_g(\phi), \quad (17)$$

where the regularization coefficient $\lambda \triangleq \lambda' / Q_{\text{avg}}$, with soft-updated Q_{avg} similar to Fujimoto and Gu [42]. We use $\lambda' = 10$ across all datasets in our experiments (Section 4).

4 Experiments

In this section, we first evaluate our AMPL on continuous-control offline-RL datasets in Section 4.1. Then we conduct ablation study on Section 4.2 to understand the efficacy of some proposed designs.

4.1 Continuous-control results

Our experiments are conducted on a diverse set of datasets in the D4RL benchmark [31], ranging across the Gym-Mojoco, Maze2D, and Adroit domains therein. We use the latest version of the datasets, *i.e.*, “v2” version for the Gym-Mojoco domain and “v1” version for the Maze2D and Adroit domains. Details of our dataset choice are discussed in Appendix E.2.

We compare our AMPL with two DICE methods — AlgaeDICE [43] and OptiDICE [44] — that directly utilize MIW to improve policy learning. We further consider three state-of-the-art (SOTA) offline model-free RL algorithms: CQL [45], FisherBRC [46], and TD3+BC [42]; and three SOTA offline MBRL methods: MOPO [16], COMBO [18], and WMOPO [47]. Experimental details and hyperparameter settings are discussed in Appendix E.2. We run baseline methods using the official implementation under the recommended hyperparameters, except for AlgaeDICE, for which we use the offline version provided by Fu et al. [31]. Table 1 shows the mean and standard deviation of the results of AMPL and baselines over five random seeds.

As shown in Table 1, our AMPL performs comparably well and is relatively stable across the sixteen tested datasets. In particular, AMPL generally performs better than the baselines on the Maze2D and Adroit datasets. The Maze2D and Adroit datasets are considered as more challenging than the MuJoCo datasets [31], since the Maze2D datasets are collected by non-Markovian policies, and the amount of data in the high-dimensional Adroit datasets is limited.

Table 1: Normalized returns for experiments on the D4RL tasks. Mean and standard deviation across five random seeds are reported. High average score and low average rank are desirable. We bold the best result over all methods and underline the best of the model-based methods if different. Here, “hcheetah” denotes “halfcheetah,” “med” denotes “medium,” “rep” denotes “replay,” and “exp” denotes “expert.”

Task Name	AlgaeDICE	OptiDICE	CQL	FisherBRC	TD3+BC	MOPO	COMBO	WMOPO	AMPL
maze2d-large	-2.2 ± 0.6	101.7 ± 50.7	1.5 ± 6.4	0.9 ± 6.4	107.1 ± 45.9	-0.5 ± 2.6	138.5 ± 82.2	1.8 ± 8.6	180.0 ± 39.3
maze2d-med	2.5 ± 14.2	119.4 ± 52.3	6.3 ± 9.1	16.1 ± 30.3	61.4 ± 45.5	12.5 ± 18.3	103.9 ± 42.1	12.7 ± 38.3	<u>107.2</u> ± 45.0
maze2d-umaze	-15.3 ± 0.8	114.0 ± 39.7	37.5 ± 7.2	3.6 ± 16.4	38.6 ± 14.4	-15.4 ± 1.9	<u>112.1</u> ± 56.8	-11.5 ± 3.1	55.8 ± 3.9
hcheetah-med	-0.9 ± 0.7	42.0 ± 3.5	48.6 ± 0.2	47.8 ± 0.3	48.0 ± 0.3	69.2 ± 3.4	73.0 ± 3.6	72.0 ± 4.7	51.7 ± 0.4
walker2d-med	1.0 ± 2.1	55.7 ± 15.7	81.8 ± 1.7	80.7 ± 2.1	83.0 ± 1.4	-0.1 ± 0.0	0.5 ± 0.9	64.8 ± 30.0	83.1 ± 1.8
hopper-med	0.9 ± 0.2	57.5 ± 7.5	68.0 ± 6.0	94.2 ± 4.5	59.1 ± 4.5	44.8 ± 41.5	22.6 ± 38.1	99.7 ± 2.9	58.9 ± 7.9
hcheetah-med-rep	-3.4 ± 2.3	40.5 ± 3.3	38.7 ± 18.6	34.9 ± 18.2	44.5 ± 0.6	62.7 ± 7.5	66.0 ± 1.8	66.4 ± 4.9	44.6 ± 0.7
walker2d-med-rep	0.5 ± 0.6	37.3 ± 21.8	80.8 ± 4.0	84.6 ± 6.7	74.7 ± 8.1	53.8 ± 34.6	60.1 ± 18.3	71.9 ± 15.0	<u>81.5</u> ± 3.0
hopper-med-rep	1.5 ± 0.7	28.1 ± 12.4	96.0 ± 4.7	94.4 ± 1.9	58.9 ± 19.7	84.8 ± 30.0	53.6 ± 29.5	<u>93.8</u> ± 8.5	91.1 ± 9.5
hcheetah-med-exp	-1.8 ± 2.9	66.7 ± 25.8	53.8 ± 13.6	94.5 ± 0.6	91.4 ± 5.2	64.0 ± 20.8	52.5 ± 32.0	68.8 ± 37.5	<u>90.2</u> ± 2.0
walker2d-med-exp	-0.2 ± 0.1	79.4 ± 16.4	110.3 ± 0.6	109.3 ± 0.1	110.4 ± 0.6	-0.2 ± 0.0	1.1 ± 0.6	98.8 ± 16.2	<u>107.7</u> ± 2.1
hopper-med-exp	2.2 ± 1.9	52.6 ± 9.3	79.3 ± 20.5	110.0 ± 4.0	98.7 ± 9.6	15.6 ± 7.4	63.7 ± 51.9	62.6 ± 26.4	<u>76.0</u> ± 15.7
pen-human	-3.0 ± 0.8	-0.9 ± 3.1	-1.5 ± 2.8	-0.1 ± 3.0	1.6 ± 10.2	-1.6 ± 2.6	5.9 ± 8.7	-2.1 ± 1.4	20.6 ± 10.7
pen-cloned	-2.6 ± 1.3	-0.8 ± 3.1	39.5 ± 23.6	-1.2 ± 3.8	6.3 ± 4.6	5.2 ± 10.2	23.2 ± 19.8	-3.0 ± 2.4	57.4 ± 19.2
pen-exp	-1.3 ± 2.4	0.4 ± 7.3	119.3 ± 15.8	2.2 ± 1.1	104.2 ± 40.6	35.4 ± 20.9	57.1 ± 42.6	10.4 ± 20.9	138.3 ± 6.6
door-exp	0.1 ± 0.0	86.1 ± 23.2	94.0 ± 15.8	22.9 ± 28.8	-0.3 ± 0.0	-0.0 ± 0.1	-0.2 ± 0.2	-0.1 ± 0.1	96.3 ± 9.8
Average Score	-1.4	55	59.6	49.7	61.7	26.9	52.1	44.2	83.8
Average Rank	8.5	5.5	4.1	4.3	3.9	6.4	4.8	5.0	2.6

On the Maze2D datasets, the behavioral-cloning (BC) style algorithms, such as FisherBRC, are likely to fail, since these methods use Markovian policy to approximate the non-Markovian behavior policy. Meanwhile, OptiDICE performs relatively well on the Maze2D datasets, mainly because it uses a Gaussian-mixture policy for BC to alleviate this approximation difficulty. This advanced BC policy still offers little help on the higher-dimensional Adroit datasets, due to the challenge of learning the behavior policy. Similarly, TD3+BC performs well on the relatively lower-dimensional Maze2D datasets, but does not work well on the Adroit domain, since the action-space MSE-regularizer in TD3+BC may be insufficient for the high-dimensional tasks. WMOPO also shows a performance drop in these two challenging task domains. The environmental dynamics on these two domains are complex, and thus may not be accurately learned to support WMOPO’s full-trajectory model rollouts, which are used to estimate its MIW. By contrast, our AMPL does not require BC or long model-rollouts, leading to stable performance across task domains. Moreover, AMPL shows generally better performance than MOPO and COMBO, which use fixed pretrained dynamic models without considering the mismatched model objectives.

Compared with the DICE-based methods AlgaeDICE and OptiDICE, our AMPL shows generally better performance. This may indicate that maximizing a lower bound of the true expected return can be a more effective framework than explicit stationary-distribution correction, since maximizing the lower bound is more directly related to RL’s goal of maximizing the policy performance.




Figure 1: Scores of each method in the ablation study (Section 4.2) on each domain of datasets. The faded bars show the scores on each datasets within the stated domain, averaged over five seeds, where each color corresponds to a dataset. The highlighted bar shows the average score on the entire domain. Label “Main” refers to the results of our main method in Section 4.1. Detailed numbers are on Table 2.

4.2 Ablation study

(a): Does the MIW-weighted model (re)training help the performance?

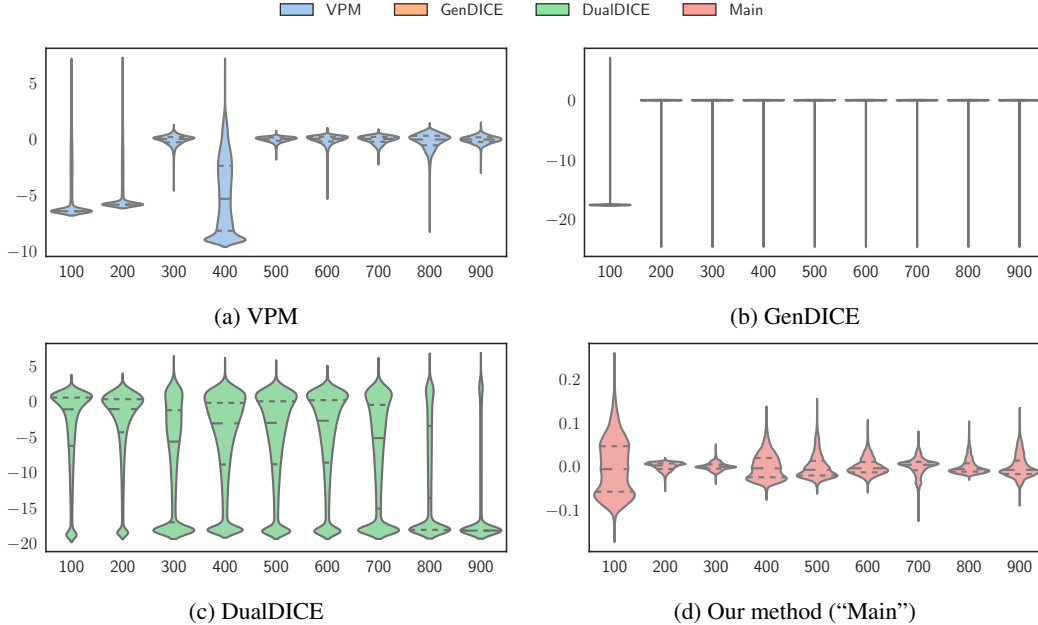


Figure 2: Distribution plots of $\log(\text{MIW})$ of the entire dataset by our method (“Main”) and the three alternatives in Section 4.2 (b), on the “walker2d-medium-replay” dataset during the training process. The dash lines represent the quartiles of the distribution. Recall that the MIW ω is retrained every 100 epochs. The x -axis represents the number of epochs $\{100, 200, \dots, 900\}$, and the y -axis the $\log(\text{MIW})$ values. For clarity, we use the normalized MIWs whose mean on the entire dataset is one. We use random seed two for this plot. Other random seeds have similar patterns. We note the different scales in the y -axis of these four plots.

To verify the effectiveness of our MIW-weighted model (re)training scheme, we compare our AMPL with its variant of training the model only at the beginning using MLE, *i.e.*, No Weights (dubbed as NW). As shown in Figure 1, the performance of the NW variant is generally worse than the main method (dubbed as Main) on all three domains. The performance difference is especially significant on the Maze2D domain. In this domain, the policy is required to “stitch together collected subtrajectories to find the shortest path to the evaluation goal” [31]. As a result, the state-action distribution induced by a good policy should be different from the distribution of the offline-data. Table 2 in Appendix A confirms the benefit of the MIW-weighted model (re)training. Indeed, on the majority of datasets, the NW variant not only shows worse performance, but also has a larger standard deviation relative to the mean score.

(b): *How does the algorithm perform if we change our MIW estimation method to other approaches?*

We compare our AMPL with its three variants where the MIW is instead estimated by the Variational Power Method [40], GENeralized stationary DIstribution Correction Estimation [30], and Dual stationary DIstribution Correction Estimation [29]. Three variants are simply dubbed as VPM, GenDICE, and DualDICE. For numerical stability, the estimated MIW from these three methods is clipped into $(10^{-8}, 500)$. Implementation details are provided in Appendix E.3.

As shown in Figure 1 and Table 2, the variants with these three alternative MIW estimation methods generally perform worse than our approach. A plausible explanation is that these methods can be unstable to provide good MIW estimates for the model training. Figure 2 shows the distribution plots of $\log(\text{MIW})$ of the entire dataset estimated by our method and by the variants using the three alternative MIW estimation methods, on the “walker2d-medium-replay” dataset. We train the MIW every 100 epochs and ignore the MIW initialization from the plots. Notice that the MIWs obtained by these four methods have very different scales, which indicates the instability of the alternative methods compared with ours.

As shown on Figure 2b, the MIWs estimated by the GenDICE variant degenerate, in a sense that most of the MIWs are ≈ 0 at the beginning and are 1 later on, with long tails on the MIW distributions. This may explain the relatively bad performance of GenDICE. Compared with GenDICE, DualDICE provides more diverse MIW estimates, which leads to its relatively better performance. Unfortunately,

Figure 2c shows that, for the DualDICE variant, the distribution of MIWs gradually concentrates on very small and very large values, which indicates the degeneration of the MIWs. The VPM variant performs the best among these three alternatives, mainly because it provides relatively better MIW estimates in the second half of the training process. However, Figure 2a shows that the MIWs from the VPM variant are poorly-distributed at the beginning, though relatively better later on. Our method provides well-behaved MIW estimates in the whole training process, which may explain our better result. As shown by Figure 2d, the MIWs from our method are well-shaped and concentrate around the mean 1^1 over the entire training process. In general, the well-shaping of the estimated MIW is important, since the Cramer-Rao lower bound of the mean-square-error for OPE is related to the square of the density ratio [38, 48].

Besides, on the Maze2D domain, these three variants using the alternative MIW estimation methods perform generally better than the NW variant discussed in Question (a), aligning with the benefit of the MIW-weighted model (re)training scheme. Generally speaking, incorporating the MIW can help model training in offline MBRL.

(c): *What is the performance of a weighted regularizer for policy learning implied by Theorem 1?*

We compare our AMPL with its variant where the policy regularizer is weighted by the MIW $\omega(s, a)$, as suggested by Theorem 1. This Weighted Policy Regularizer variant is dubbed as WPR. As shown in both Figure 1 and Table 2, the WPR variant underperforms our main method. This is because when we estimate the regularization term in WPR, we incorporate weights into the minimax optimization of the policy and the discriminator, which may bring additional instability.

Table 3 of Appendix A provides additional ablation study on the performance of several alternatives.

5 Related work

Offline MBRL. Most of the existing offline MBRL works focus on policy learning under a given model trained by MLE. These works typically constrain the learned policy to avoid visiting regions where the discrepancy between the true and the learned dynamic is large, thus reducing the policy evaluation error of using the learned dynamic [16, 19, 49, 50]. Besides, some recent works [51, 52] also adopt a GAN-style stochastic policy for its flexibility. Rather than using a fixed MLE-trained model, we derive an objective that trains both the policy and the dynamic model toward maximizing a lower bound of true expected return (simultaneously minimizing the policy evaluation error $|J(\pi, P^*) - J(\pi, \hat{P})|$). Several recent works also propose to enhance model training, such as training a reverse dynamic model to encourage conservative model-based imaginations [53], learning a balanced state-action representation for the model to mitigate the distributional discrepancy [54, 55], and using advanced dynamic-model architecture of the GPT-Transformer [56–59] to achieve accurate predictions [60]. These methods are orthogonal to ours and may further improve the performance.

Objective mismatch in MBRL. Our AMPL is related to prior works on online and offline MBRL that mitigate the mismatched model objectives. In online (off-policy) MBRL, Lambert et al. [23] identify the mismatched objectives between the MLE model-training and the model’s usage of improving the control performance. This paper, however, only proposes an impractical weighted learning loss that requires the optimal trajectory. Eysenbach et al. [24] design an objective to jointly optimize the policy and the model in online MBRL. Voelcker et al. [61] propose a per-sample diagonal scaling matrix for training a deterministic model in online RL, which improves on Farahmand et al. [62] and Farahmand [63]. In offline MBRL, D’Oro et al. [64] propose a weighting scheme to learn the transition model. This method requires the behavior policy and estimates the trajectory-wise importance-sampling weight, which is known to suffer from the “curse of horizon” [27]. Thus, this method may not scale to complex high-dimensional offline RL tasks. Voloshin et al. [65] propose a minimax objective for model learning, where the model is optimized over function classes of the MIW and the state-value function. Most similar to our work, Hishinuma and Senda [47] also use a MIW-weighted MLE objective for model training. They use a different version of the MIW estimated via multiple *full-trajectory* rollouts in the learned model. This strategy may suffer from the inaccuracy of the learned model [66–68] and can be time-consuming. In this paper, we develop a fixed-point style method to train the MIW that only requires samples from the offline dataset and the initial state-distribution, which can be more stable and efficient.

¹0 in the log (MIW) plots.

Off-policy evaluation (OPE). OPE [29, 30, 48, 69–74] is a well-studied problem with several lines of research. A natural way is trajectory-wise importance sampling (IS) [75]. This, however, suffers from its variance, which can grow exponentially with the horizon length [27, 76]. Marginal importance sampling methods [27] are developed to address this problem by estimating the IS ratio between two stationary distributions, leading to estimators whose variance are polynomial *w.r.t.* horizon [77]. Though prior works can also be used to estimate the MIW, they may not be directly applicable to our framework, since they either require knowing the behavior policy [27, 37], or need to assume the MIW lying in the space of RKHS [28]. Another approach to estimate the MIW is through saddle-point optimization [29, 30, 39, 78], which can suffer from its complex optimization problem. In this paper, we derive an fixed-point-style and behavior-agnostic MIW estimation method, which is simple to train and does not require additional assumptions.

6 Conclusion

In this paper, we are motivated by the mismatched model objectives in offline MBRL to design an iterative algorithm that alternates between model training and policy optimization. Both the model and the policy are trained to maximize a lower bound of the true expected return. The proposed new algorithm performs competitively with several SOTA baselines.

Limitation. Our current approach requires training two additional networks (discriminator D_ψ and MIW ω), and the model is (re)trained periodically to mitigate the objective mismatch, which brings additional computational cost in order to obtain better empirical results.

Future work. We plan to conduct a theoretical analysis of our proposed framework, and investigate more efficient ways to unify model training and policy learning.

Acknowledgments

S. Yang, S. Zhang, and M. Zhou acknowledge the support of NSF IIS 1812699 and 2212418, and the Texas Advanced Computing Center (TACC) for providing HPC resources that have contributed to the research results reported within this paper.

References

- [1] Sascha Lange, Thomas Gabel, and Martin Riedmiller. *Batch Reinforcement Learning*, pages 45–73. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012. ISBN 978-3-642-27645-3. doi: 10.1007/978-3-642-27645-3_2.
- [2] 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.
- [3] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, February 2015. ISSN 00280836.
- [4] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the Game of Go with Deep Neural Networks and Tree Search. *Nature*, 529(7587):484–489, January 2016. doi: 10.1038/nature16961.
- [5] Lihong Li, Wei Chu, John Langford, and Xuanhui Wang. Unbiased offline evaluation of contextual-bandit-based news article recommendation algorithms. In *Proceedings of the 4th International Conference on Web Search and Data Mining (WSDM)*, pages 297–306, 2011.
- [6] Ekim Yurtsever, Jacob Lambert, Alexander Carballo, and K. Takeda. A Survey of Autonomous Driving: Common Practices and Emerging Technologies. *IEEE Access*, 8:58443–58469, 2020.

- [7] Youngsoo Jang, Jongmin Lee, and Kee-Eung Kim. GPT-critic: Offline reinforcement learning for end-to-end task-oriented dialogue systems. In *International Conference on Learning Representations*, 2022.
- [8] Cenying Yang, Yihao Feng, and Andrew Whinston. Dynamic pricing and information disclosure for fresh produce: An artificial intelligence approach. *Production and Operations Management*, 31(1):155–171, 2022.
- [9] T. Lillicrap, Jonathan J. Hunt, A. Pritzel, N. Heess, T. Erez, Yuval Tassa, D. Silver, and Daan Wierstra. Continuous Control with Deep Reinforcement Learning. *CoRR*, abs/1509.02971, 2016.
- [10] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1861–1870. PMLR, 10–15 Jul 2018.
- [11] Scott Fujimoto, Herke van Hoof, and David Meger. Addressing Function Approximation Error in Actor-Critic Methods. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1587–1596. PMLR, 10–15 Jul 2018.
- [12] Scott Fujimoto, David Meger, and Doina Precup. Off-Policy Deep Reinforcement Learning without Exploration. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2052–2062. PMLR, 09–15 Jun 2019.
- [13] Aviral Kumar, Justin Fu, Matthew Soh, George Tucker, and Sergey Levine. Stabilizing Off-Policy Q-Learning via Bootstrapping Error Reduction. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [14] Rishabh Agarwal, Dale Schuurmans, and Mohammad Norouzi. An optimistic perspective on offline reinforcement learning. In *International Conference on Machine Learning*, pages 104–114. PMLR, 2020.
- [15] Michael Janner, Justin Fu, Marvin Zhang, and Sergey Levine. When to Trust Your Model: Model-Based Policy Optimization. *ArXiv*, abs/1906.08253, 2019.
- [16] Tianhe Yu, Garrett Thomas, Lantao Yu, Stefano Ermon, James Y. Zou, Sergey Levine, Chelsea Finn, and Tengyu Ma. MOPO: Model-based Offline Policy Optimization. *ArXiv*, abs/2005.13239, 2020.
- [17] Aravind Rajeswaran, Igor Mordatch, and Vikash Kumar. A Game Theoretic Framework for Model Based Reinforcement Learning. In *International Conference on Machine Learning*, 2020.
- [18] Tianhe Yu, Aviral Kumar, Rafael Rafailov, Aravind Rajeswaran, Sergey Levine, and Chelsea Finn. COMBO: Conservative Offline Model-Based Policy Optimization. *ArXiv*, abs/2102.08363, 2021.
- [19] Rahul Kidambi, Aravind Rajeswaran, Praneeth Netrapalli, and T. Joachims. MOREL : Model-Based Offline Reinforcement Learning. *ArXiv*, abs/2005.05951, 2020.
- [20] Phillip Swazinna, Steffen Udfluft, and Thomas A. Runkler. Overcoming Model Bias for Robust Offline Deep Reinforcement Learning. *Eng. Appl. Artif. Intell.*, 104:104366, 2021.
- [21] Catherine Cang, Aravind Rajeswaran, P. Abbeel, and M. Laskin. Behavioral Priors and Dynamics Models: Improving Performance and Domain Transfer in Offline RL. *ArXiv*, abs/2106.09119, 2021.
- [22] 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.

- [23] Nathan Lambert, Brandon Amos, Omry Yadan, and Roberto Calandra. Objective mismatch in model-based reinforcement learning. *arXiv preprint arXiv:2002.04523*, 2020.
- [24] Benjamin Eysenbach, Alexander Khazatsky, Sergey Levine, and Ruslan Salakhutdinov. Mismatched No More: Joint Model-Policy Optimization for Model-Based RL. *ArXiv*, abs/2110.02758, 2021.
- [25] Richard S. Sutton. Dyna, an integrated architecture for learning, planning, and reacting. *ACM Sigart Bulletin*, 2:160–163, 1991.
- [26] Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 1. Belmont, MA: Athena Scientific, 1995.
- [27] Q. Liu, Lihong Li, Ziyang Tang, and Dengyong Zhou. Breaking the Curse of Horizon: Infinite-Horizon Off-Policy Estimation. In *Advances in neural information processing systems*, 2018.
- [28] A. Mousavi, Lihong Li, Qiang Liu, and Denny Zhou. Black-box Off-policy Estimation for Infinite-Horizon Reinforcement Learning. *ArXiv*, abs/2003.11126, 2020.
- [29] Ofir Nachum, Yinlam Chow, Bo Dai, and Lihong Li. DualDICE: Behavior-Agnostic Estimation of Discounted Stationary Distribution Corrections. In *Advances in neural information processing systems*, 2019.
- [30] Ruiyi Zhang, Bo Dai, Lihong Li, and D. Schuurmans. GenDICE: Generalized Offline Estimation of Stationary Values. *ArXiv*, abs/2002.09072, 2020.
- [31] 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.
- [32] Justin Fu, Aviral Kumar, Matthew Soh, and Sergey Levine. Diagnosing Bottlenecks in Deep Q-learning Algorithms. In *International Conference on Machine Learning*, 2019.
- [33] Long-Ji Lin. Self-Improving Reactive Agents Based on Reinforcement Learning, Planning and Teaching. *Machine Learning*, 8(3–4):293–321, 1992.
- [34] Monroe D Donsker and SR Srinivasa Varadhan. Asymptotic evaluation of certain markov process expectations for large time. iv. *Communications on pure and applied mathematics*, 36(2):183–212, 1983.
- [35] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.
- [36] Xingchao Liu, Chengyue Gong, Lemeng Wu, Shujian Zhang, Hao Su, and Qiang Liu. Fuse-dream: Training-free text-to-image generation with improved clip+gan space optimization. *arXiv preprint arXiv:2112.01573*, 2021.
- [37] Ziyang Tang, Yihao Feng, Lihong Li, Dengyong Zhou, and Qiang Liu. Doubly robust bias reduction in infinite horizon off-policy estimation. In *International Conference on Learning Representations*, 2020.
- [38] Nathan Kallus and Masatoshi Uehara. Double reinforcement learning for efficient off-policy evaluation in markov decision processes. *J. Mach. Learn. Res.*, 21(167):1–63, 2020.
- [39] Mengjiao Yang, Ofir Nachum, Bo Dai, Lihong Li, and Dale Schuurmans. Off-policy evaluation via the regularized lagrangian. *Advances in Neural Information Processing Systems*, 33:6551–6561, 2020.
- [40] Junfeng Wen, Bo Dai, Lihong Li, and Dale Schuurmans. Batch stationary distribution estimation. *arXiv preprint arXiv:2003.00722*, 2020.
- [41] Peter J. Huber. Robust Estimation of a Location Parameter. *The Annals of Mathematical Statistics*, 35(1):73 – 101, 1964.

- [42] Scott Fujimoto and Shixiang Shane Gu. A Minimalist Approach to Offline Reinforcement Learning. *ArXiv*, abs/2106.06860, 2021.
- [43] Ofir Nachum, Bo Dai, Ilya Kostrikov, Yinlam Chow, Lihong Li, and D. Schuurmans. AlgaeDICE: Policy Gradient from Arbitrary Experience. *ArXiv*, abs/1912.02074, 2019.
- [44] Jongmin Lee, Wonseok Jeon, Byung-Jun Lee, J. Pineau, and Kee-Eung Kim. OptiDICE: Offline Policy Optimization via Stationary Distribution Correction Estimation. *ArXiv*, abs/2106.10783, 2021.
- [45] Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. Conservative Q-Learning for Offline Reinforcement Learning. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1179–1191. Curran Associates, Inc., 2020.
- [46] Ilya Kostrikov, Jonathan Tompson, Rob Fergus, and Ofir Nachum. Offline Reinforcement Learning with Fisher Divergence Critic Regularization. In *International Conference on Machine Learning*, 2021.
- [47] Toru Hishinuma and Kei Senda. Weighted model estimation for offline model-based reinforcement learning. In *Advances in neural information processing systems*, 2021.
- [48] Nan Jiang and Lihong Li. Doubly robust off-policy evaluation for reinforcement learning. In *Proceedings of the 23rd International Conference on Machine Learning*, pages 652–661, 2016.
- [49] Rafael Rafailov, Tianhe Yu, Aravind Rajeswaran, and Chelsea Finn. Offline reinforcement learning from images with latent space models. In *Learning for Dynamics and Control*, pages 1154–1168. PMLR, 2021.
- [50] Cong Lu, Philip J. Ball, Jack Parker-Holder, Michael A. Osborne, and Stephen J. Roberts. Revisiting Design Choices in Model-Based Offline Reinforcement Learning. *ArXiv*, abs/2110.04135, 2021.
- [51] Shentao Yang, Zhendong Wang, Huangjie Zheng, Yihao Feng, and Mingyuan Zhou. A regularized implicit policy for offline reinforcement learning. *arXiv preprint arXiv:2202.09673*, 2022.
- [52] Shentao Yang, Yihao Feng, Shujian Zhang, and Mingyuan Zhou. Regularizing a model-based policy stationary distribution to stabilize offline reinforcement learning. In *International Conference on Machine Learning*, pages 24980–25006. PMLR, 2022.
- [53] Jianhao Wang, Wenzhe Li, Haozhe Jiang, Guangxiang Zhu, Siyuan Li, and Chongjie Zhang. Offline reinforcement learning with reverse model-based imagination. *Advances in Neural Information Processing Systems*, 34, 2021.
- [54] 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. *Advances in Neural Information Processing Systems*, 31, 2018.
- [55] Byung-Jun Lee, Jongmin Lee, and Kee-Eung Kim. Representation Balancing Offline Model-based Reinforcement Learning. In *International Conference on Learning Representations*, 2021.
- [56] Alec Radford and Ilya Sutskever. Improving Language Understanding by Generative Pre-Training. In *arxiv*, 2018.
- [57] Xinjie Fan, Shujian Zhang, Bo Chen, and Mingyuan Zhou. Bayesian attention modules. *Advances in Neural Information Processing Systems*, 33:16362–16376, 2020.
- [58] Shujian Zhang, Xinjie Fan, Bo Chen, and Mingyuan Zhou. Bayesian attention belief networks. In *International Conference on Machine Learning*, pages 12413–12426. PMLR, 2021.

- [59] Shujian Zhang, Chengyue Gong, Xingchao Liu, Pengcheng He, Weizhu Chen, and Mingyuan Zhou. Allsh: Active learning guided by local sensitivity and hardness. *arXiv preprint arXiv:2205.04980*, 2022.
- [60] Michael Janner, Qiyang Li, and Sergey Levine. Reinforcement Learning as One Big Sequence Modeling Problem. *ArXiv*, abs/2106.02039, 2021.
- [61] Claas Voelcker, Victor Liao, Animesh Garg, and Amir-massoud Farahmand. Value Gradient weighted Model-Based Reinforcement Learning. *arXiv preprint arXiv:2204.01464*, 2022.
- [62] 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. PMLR, 2017.
- [63] Amir-massoud Farahmand. Iterative value-aware model learning. *Advances in Neural Information Processing Systems*, 31, 2018.
- [64] Pierluca D’Oro, Alberto Maria Metelli, Andrea Tirinzoni, Matteo Papini, and Marcello Restelli. Gradient-aware model-based policy search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 3801–3808, 2020.
- [65] Cameron Voloshin, Nan Jiang, and Yisong Yue. Minimax model learning. In *International Conference on Artificial Intelligence and Statistics*, pages 1612–1620. PMLR, 2021.
- [66] Stéphane Ross and J. Andrew Bagnell. Agnostic System Identification for Model-Based Reinforcement Learning. *International Conference on Machine Learning*, abs/1203.1007, 2012.
- [67] Ignasi Clavera, Jonas Rothfuss, John Schulman, Yasuhiro Fujita, Tamim Asfour, and P. Abbeel. Model-Based Reinforcement Learning via Meta-Policy Optimization. *ArXiv*, abs/1809.05214, 2018.
- [68] Thanard Kurutach, Ignasi Clavera, Yan Duan, Aviv Tamar, and P. Abbeel. Model-Ensemble Trust-Region Policy Optimization. *International Conference on Learning Representations*, abs/1802.10592, 2018.
- [69] Bo Dai, Ofir Nachum, Yinlam Chow, Lihong Li, Csaba Szepesvári, and Dale Schuurmans. Coindice: Off-policy confidence interval estimation. *Advances in neural information processing systems*, 33:9398–9411, 2020.
- [70] Yihao Feng, Lihong Li, and Qiang Liu. A kernel loss for solving the bellman equation. *Advances in Neural Information Processing Systems*, 32, 2019.
- [71] Yihao Feng, Tongzheng Ren, Ziyang Tang, and Qiang Liu. Accountable off-policy evaluation with kernel bellman statistics. In *International Conference on Machine Learning*, pages 3102–3111. PMLR, 2020.
- [72] Yihao Feng, Ziyang Tang, Na Zhang, and Qiang Liu. Non-asymptotic confidence intervals of off-policy evaluation: Primal and dual bounds. *arXiv preprint arXiv:2103.05741*, 2021.
- [73] Ziyang Tang, Yihao Feng, Na Zhang, Jian Peng, and Qiang Liu. Off-policy interval estimation with lipschitz value iteration. *Advances in Neural Information Processing Systems*, 33:7887–7897, 2020.
- [74] Yihao Feng. *Offline Reinforcement Learning via an Optimization Lens*. PhD thesis, The University of Texas at Austin, 2022.
- [75] Doina Precup, Richard S Sutton, and Sanjoy Dasgupta. Off-policy temporal-difference learning with function approximation. In *International Conference on Machine Learning*, pages 417–424, 2001.
- [76] Lihong Li, Rémi Munos, and Csaba Szepesvári. Toward minimax off-policy value estimation. In *Artificial Intelligence and Statistics*, pages 608–616. PMLR, 2015.

- [77] Tengyang Xie, Yifei Ma, and Yu-Xiang Wang. Towards optimal off-policy evaluation for reinforcement learning with marginalized importance sampling. *Advances in Neural Information Processing Systems*, 32, 2019.
- [78] Ofir Nachum and Bo Dai. Reinforcement learning via fenchel-rockafellar duality. *arXiv preprint arXiv:2001.01866*, 2020.
- [79] Rémi Munos. Error bounds for approximate policy iteration. In *International Conference on Machine Learning*, volume 3, pages 560–567, 2003.
- [80] Jinglin Chen and Nan Jiang. Information-Theoretic Considerations in Batch Reinforcement Learning. *ArXiv*, abs/1905.00360, 2019.
- [81] Masatoshi Uehara, Jiawei Huang, and Nan Jiang. Minimax weight and q-function learning for off-policy evaluation. In *International Conference on Machine Learning*, pages 9659–9668. PMLR, 2020.
- [82] Xi Chen, Ali Ghadirzadeh, Tianhe Yu, Yuan Gao, Jianhao Wang, Wenzhe Li, Bin Liang, Chelsea Finn, and Chongjie Zhang. Latent-variable advantage-weighted policy optimization for offline rl. *arXiv preprint arXiv:2203.08949*, 2022.
- [83] Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep Reinforcement Learning in a Handful of Trials using Probabilistic Dynamics Models. In *Advances in neural information processing systems*, 2018.
- [84] Tatsuya Matsushima, Hiroki Furuta, Yutaka Matsuo, Ofir Nachum, and Shixiang Shane Gu. Deployment-Efficient Reinforcement Learning via Model-Based Offline Optimization. *International Conference on Learning Representations*, abs/2006.03647, 2021.
- [85] H. V. Hasselt. Double Q-learning. In *Advances in neural information processing systems*, 2010.
- [86] Chengyue Gong, Xingchao Liu, and Qiang Liu. Bi-objective Trade-off with Dynamic Barrier Gradient Descent. *Advances in Neural Information Processing Systems*, 34, 2021.
- [87] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*, 2014.
- [88] Yuguang Yue, Zhendong Wang, and Mingyuan Zhou. Implicit Distributional Reinforcement Learning. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [89] Yifan Wu, George Tucker, and Ofir Nachum. Behavior regularized offline reinforcement learning. *arXiv preprint arXiv:1911.11361*, 2019.
- [90] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, M. Laskin, P. Abbeel, A. Srinivas, and Igor Mordatch. Decision Transformer: Reinforcement Learning via Sequence Modeling. *ArXiv*, abs/2106.01345, 2021.
- [91] Ilya Kostrikov, Ashvin Nair, and Sergey Levine. Offline Reinforcement Learning with Implicit Q-Learning. *ArXiv*, abs/2110.06169, 2021.
- [92] Aravind Rajeswaran, Vikash Kumar, Abhishek Gupta, J. Schulman, E. Todorov, and Sergey Levine. Learning Complex Dexterous Manipulation with Deep Reinforcement Learning and Demonstrations. *ArXiv*, abs/1709.10087, 2018.
- [93] Tom White. Sampling generative networks. *arXiv preprint arXiv:1609.04468*, 2016.
- [94] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.

- [95] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [96] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *CoRR*, abs/1511.06434, 2016.
- [97] Tim Salimans, I. Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved Techniques for Training GANs. In *Advances in neural information processing systems*, 2016.
- [98] Ian Goodfellow. Nips 2016 tutorial: Generative adversarial networks. *arXiv preprint arXiv:1701.00160*, 2016.
- [99] Huangjie Zheng and Mingyuan Zhou. Exploiting Chain Rule and Bayes’ Theorem to Compare Probability Distributions. *Advances in Neural Information Processing Systems*, 34, 2021.
- [100] Shujian Zhang, Xinjie Fan, Huangjie Zheng, Korawat Tanwisuth, and Mingyuan Zhou. Alignment attention by matching key and query distributions. *Advances in Neural Information Processing Systems*, 34, 2021.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? **[Yes]** See Sections 3 and 4.
 - (b) Did you describe the limitations of your work? **[Yes]** See Sections 5 and 6.
 - (c) Did you discuss any potential negative societal impacts of your work? **[Yes]** See Appendix F.
 - (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? **[Yes]** See Section 3 and Appendix B.
 - (b) Did you include complete proofs of all theoretical results? **[Yes]** See Appendix B.
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 Section 3 and Appendix E for descriptions. Source code is publicly **released**.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[Yes]** See Appendix E.2 and E.3.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[Yes]** See Tables 1, 2, 3, and 4.
 - (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 Appendix G.
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]**
 - (b) Did you mention the license of the assets? **[Yes]** See the GitHub pages of the assets.
 - (c) Did you include any new assets either in the supplemental material or as a URL? **[No]**
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? **[No]** We use public benchmarks.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[No]** The used benchmarks are widely accepted in the community.
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]