RAMBO-RL: Robust Adversarial Model-Based Offline Reinforcement Learning

Anonymous Author(s)
Affiliation
Address
email

Abstract

Offline reinforcement learning (RL) aims to find performant policies from logged data without further environment interaction. Model-based algorithms, which learn a model of the environment from the dataset and perform conservative policy optimisation within that model, have emerged as a promising approach to this problem. In this work, we present Robust Adversarial Model-Based Offline RL (RAMBO), a novel approach to model-based offline RL. To enforce conservatism, we formulate the problem as a two-player zero sum game against an adversarial environment model. The model is trained to minimise the value function while still accurately predicting the transitions in the dataset, forcing the policy to act conservatively in areas not covered by the dataset. To approximately solve the two-player game, we alternate between optimising the policy and adversarially optimising the model. The problem formulation that we address is theoretically grounded, resulting in a probably approximately correct (PAC) performance guarantee and a pessimistic value function which lower bounds the value function in the true environment. We evaluate our approach on widely studied offline RL benchmarks, and demonstrate that it outperforms existing state-of-the-art baselines.

1 Introduction

Reinforcement learning (RL) [61] has achieved state-of-the-art performance on many sequential decision-making problems [40, 58, 44]. However, the need for extensive exploration prohibits the application of RL to many real world domains where such exploration is costly or dangerous. Offline RL [32, 34] overcomes this limitation by learning policies from static, pre-recorded datasets.

Online RL algorithms perform poorly in the offline setting due to the distribution shift between the state-action pairs in the dataset and those taken by the learnt policy. Thus, an important aspect of offline RL is to introduce *conservatism* to prevent the learnt policy from executing state-action pairs which are out of distribution. Model-free offline RL algorithms [28, 30, 72, 22, 26, 14] train a policy from only the data present in the fixed dataset, and incorporate conservatism either into the value function or by directly constraining the policy.

On the other hand, model-based offline RL algorithms [76, 75, 24, 63, 39] use the dataset to learn a model of the environment, and train a policy using additional synthetic data generated from that model. By training on additional synthetic data, model-based algorithms can potentially generalise better to states not present in the dataset, or to solving new tasks. Previous approaches to model-based offline RL incorporate conservatism by estimating the uncertainty in the model and applying reward penalties for state-action pairs that have high uncertainty [76, 24]. However, uncertainty estimation can be unreliable for neural network models [75, 36]. Like recent work [75], we propose an approach for offline model-based RL which *does not require uncertainty estimation*.

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

In this work we present Robust Adversarial Model-Based Offline (RAMBO) RL, a new algorithm for model-based offline RL. RAMBO incorporates conservatism by modifying the *transition dynamics* of the learnt Markov decision process (MDP) model in an adversarial manner. We formulate the problem of offline RL as a zero-sum game against an adversarial environment. To solve the resulting maximin optimisation problem, we alternate between optimising the agent and optimising the adversary in the style of Robust Adversarial RL (RARL) [49]. Unlike existing RARL approaches, our model-based approach forgoes the need to define and train an adversary policy, and instead only learns an adversarial model of the MDP. We train the agent policy with an actor-critic algorithm using synthetic data generated from the model in addition to data sampled from the dataset, similar to Dyna [60] and a number of recent methods [21, 76, 24, 75]. We update the environment model so that it reduces the value function for the agent policy, whilst still accurately predicting the transitions in the dataset. As a result, our approach introduces conservatism by generating *pessimistic synthetic transitions* for state-action pairs which are out-of-distribution. The theoretical formulation of offline RL that our algorithm addresses yields a PAC bound for the performance gap with respect to any policy covered by the dataset, and a pessimistic value function that lower bounds the value function in the true environment.

In summary, the main contributions of this work are:

- RAMBO, a novel and theoretically-grounded model-based offline RL algorithm which enforces conservatism by training an adversarial dynamics model.
- Adapting the Robust Adversarial RL approach to model-based offline RL by proposing a new formulation of RARL, where instead of defining and training an adversary policy, we directly train the model adversarially.

In our experiments we demonstrate that RAMBO outperforms current state-of-the-art algorithms on the D4RL benchmarks [13], and we provide ablation results which show that if the model is not trained adversarially, the performance of RAMBO degrades significantly.

2 Related Work

Offline RL: Offline RL addresses the problem of learning policies from fixed datasets, and has been applied to domains such as healthcare [42, 56], natural language processing [23, 22], and robotics [29, 37, 50]. Model-free offline RL algorithms do not require a learnt model. Approaches for model-free offline RL include importance sampling algorithms [35, 41], constraining the learnt policy to be similar to the behaviour policy [15, 27, 72, 22, 57, 14], incorporating conservatism into the value function during training [9, 26, 30, 73], using uncertainty quantification to generate more robust value estimates [1, 2, 28], or applying only a single iteration of policy iteration [7, 48]. In contrast, model-based approaches learn a model of the environment and generate synthetic data from that model [60] to optimise a policy using either planning [4] or RL algorithms [24, 76, 75]. By training a policy on additional synthetic data, model-based approaches have the potential for broader generalisation and for solving new tasks [6, 76]. A simple approach to ensuring conservatism is to constrain the policy to be similar to the behaviour policy in the same fashion as some model-free approaches [8, 39, 63]. Another approach is to apply reward penalties for executing state-action pairs with high uncertainty in the environment model [24, 74, 76]. However, this requires explicit uncertainty estimates which may be unreliable for neural network models [16, 36, 45, 75]. COMBO [75] obviates the need for uncertainty estimation in model-based offline RL by adapting model-free techniques [30] to regularise the value function for out-of-distribution samples. Like COMBO, our approach does not require uncertainty estimation.

Most approaches to model-based offline RL use maximum likelihood estimates (MLE) of the MDP trained using standard supervised learning [4, 39, 63, 76, 75]. However, other methods have been proposed to learn models which are more suitable for offline policy optimisation. One approach is to reweight the loss function to ensure the model is accurate under the state-action distribution generated by the policy [33, 52, 19]. A *reverse* model is used by [67] to generate rollouts to goal states which are in-distribution. In contrast, our approach produces pessimistic synthetic transitions when out-of-distribution.

Most related to our work is a recent paper [66] which introduces the maximin formulation of offline RL that we address. This existing work motivates our approach theoretically by showing that the problem formulation obtains probably approximately correct (PAC) guarantees. However, [66] only addresses the theoretical aspects of the problem formulation and does not propose a practical

algorithm. In this work, we propose a practical RL algorithm to solve the maximin formulation of model-based offline RL.

Robust RL: Algorithms for Robust MDPs [5, 20, 43, 53, 64, 70] find the policy with the best worst-case performance over a set of possible MDPs. Typically, it is assumed that the uncertainty set of MDPs is specified a priori. To eliminate the need to specify the set of possible MDPs, model-free approaches to Robust MDPs [68, 54] instead assume that samples can be drawn from a misspecified MDP which is similar to the true MDP. As our work addresses offline RL, we assume that we have a fixed dataset from the true MDP.

Our approach is conceptually similar to Robust Adversarial RL (RARL) [49], a method proposed to improve the robustness of RL policies in the *online* setting. RARL is posed as a two-player zero-sum game where the agent plays against an adversary which perturbs the environment. Formulations of model-free RARL differ in how they define the action space of the adversary. Options include allowing the adversary to apply perturbation forces to the simulator [49], add noise to the agent's actions [65], or periodically take over control [47]. A model-based approach to RARL is proposed in [12], which learns an optimistic and pessimistic model to encourage online exploration. However, this existing approach requires uncertainty estimation as well as an adversarial policy to be learnt in *addition to the model*. Our work follows the paradigm of RARL and alternates between agent and adversarial updates in a maximin formulation. We adapt model-based RARL to the *offline* setting and propose an alternative formulation: we eliminate the need to learn an adversary policy and instead *directly modify the MDP model adversarially*. Related to our work is [38] which adversarially modifies trajectories while training a policy. However, [38] assumes that the dynamics are known in the form of a differential equation, whereas we assume that we only have access to sampled data.

3 Preliminaries

MDPs and Offline RL: An MDP is defined by the tuple, $M = (S, A, T, R, \mu_0, \gamma)$. S and A denote the state and action spaces respectively, $R(s, a)$ is the reward function, $T(s'|s, a)$ is the transition function, μ_0 is the initial state distribution, and $\gamma \in (0, 1)$ is the discount factor. In this work we consider Markovian policies, $\pi \in \Pi$, which map from each state to a distribution over actions. We denote the (improper) discounted state visitation distribution of a policy by $d^\pi_M(s) := \sum_{t=0}^{\infty} \gamma^t \Pr(s_t = s|\pi, M)$, where $\Pr(s_t = s|\pi, M)$ is the probability of reaching state s at time t by executing policy π in M. The improper state-action visitation distribution is $d^\pi_M(s, a) = \pi(a|s) \cdot d^\pi_M(s)$. We also denote the normalised state-action visitation distribution by $\tilde{d}^\pi_M(s, a) = (1 - \gamma) \cdot d^\pi_M(s, a)$.

The value function, $V^\pi_M(s)$, represents the expected discounted return from executing π from state s in M: $V^\pi_M(s) = \mathbb{E}_{\pi,M}\big[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t)\big]$. We write V^π_M to indicate the value function under the initial state distribution, i.e. $V^\pi_M = \sum_{s \in S} \mu_0(s) V^\pi_M(s)$. The standard objective for MDPs is to find the policy which maximises V^π_M. The state-action value function, $Q^\pi_M(s, a)$, is the expected discounted cumulative reward from taking action a at state s and then executing π thereafter.

In offline RL we only have access to a fixed dataset of transitions from the MDP, $\mathcal{D} = \{(s_i, a_i, r_i, s'_i)\}_{i=1}^{|\mathcal{D}|}$. The goal of offline RL is to find the best possible policy using the fixed dataset.

Model-Based Offline RL Algorithms: Model-based approaches to offline RL use a model of the MDP to help train a policy. The dataset is used to learn a dynamics model, $\widehat{T}$, which is typically trained via maximum likelihood estimation: $\min_{\widehat{T}} \mathbb{E}_{(s,a,s') \sim \mathcal{D}}\big[- \log \widehat{T}(s'|s, a)\big]$. A model of the reward function, $\widehat{R}(s, a)$, can also be learnt if it is unknown. The estimated MDP, $\widehat{M} = (S, A, \widehat{T}, \widehat{R}, \mu_0, \gamma)$, has the same state and action space as the true MDP but uses the learnt transition and reward functions. Thereafter, any planning or RL algorithm can be used to recover optimal policy in the learnt model, $\widehat{\pi} = \arg\max_{\pi \in \Pi} V^\pi_{\widehat{M}}$.

Unfortunately, directly applying this approach to the offline RL setting does not perform well due to distribution shift. In particular, if the dataset does not cover the entire state-action space, the model will inevitably be inaccurate for some state-action pairs. Thus, naive policy optimisation on a learnt model in the offline setting can result in *model exploitation* [21, 31, 52]. To mitigate this issue, we propose the novel approach of enforcing conservatism by adversarially modifying the transition dynamics of $\widehat{M}$.

In line with existing works [76, 75, 8], we use model-based policy optimisation (MBPO) [21] to learn the optimal policy for $\widehat{M}$. MBPO utilises a standard actor-critic RL algorithm. However, the value function is trained using an augmented dataset $\mathcal{D} \cup \mathcal{D}_{\widehat{M}}$, where $\mathcal{D}_{\widehat{M}}$ is synthetic data generated by simulating rollouts in the learnt model. To generate the synthetic data, MBPO performs k-step rollouts in $\widehat{M}$ starting from states $s \in \mathcal{D}$, and adds this data to $\mathcal{D}_{\widehat{M}}$. To train the policy, minibatches of data are drawn from $\mathcal{D} \cup \mathcal{D}_{\widehat{M}}$, where each datapoint is sampled from the real data, $\mathcal{D}$, with probability f, and from $\mathcal{D}_{\widehat{M}}$ with probability $1 - f$.

Robust Adversarial Reinforcement Learning: RARL addresses the problem of finding a robust agent policy, π, in the online RL setting by posing the problem as a two-player zero sum game against adversary policy, $\bar{\pi}$:

$$\pi = \arg\max_{\pi \in \Pi} \min_{\bar{\pi} \in \bar{\Pi}} V_M^{\pi, \bar{\pi}} \tag{1}$$

where $V_M^{\pi,\bar{\pi}}$ is the expected value from executing π and $\bar{\pi}$ in environment M. Different approaches define the action space for $\bar{\pi}$ in different ways, as discussed in Section 2. For a scalable approximation to the optimisation problem in Equation 1, algorithms for RARL alternate between applying steps of stochastic gradient ascent to the agent's policy to increase the expected value, and stochastic gradient descent to the adversary's policy to decrease the expected value. In our work, we follow the RARL paradigm of alternating between agent and adversarial updates and adapt it to the model-based offline setting. Instead of defining a separate adversary policy, we treat the *model itself* as the policy to be adversarially trained.

4 Problem Formulation

For the sake of generality, we assume that both the transition function and reward function are unknown. Hereafter, we write $\widehat{T}$ to denote both the learnt dynamics and reward function, where $\widehat{T}(s', r|s, a)$ is the probability of receiving reward r and transitioning to s' after executing (s, a). We address the maximin formulation of offline RL recently proposed by [66]:

Problem 1. *For some dataset, $\mathcal{D}$, and some fixed constant $\xi > 0$, find the policy π defined by*

$$\pi = \arg\max_{\pi \in \Pi} \min_{\widehat{T} \in \mathcal{M}_{\mathcal{D}}} V_{\widehat{T}}^{\pi}, \text{ where} \tag{2}$$

$$\mathcal{M}_{\mathcal{D}} = \Big\{ \widehat{T} \mid \mathbb{E}_{\mathcal{D}} \big[\text{TV}(\widehat{T}_{\text{MLE}}(\cdot|s,a), \widehat{T}(\cdot|s,a))^2 \big] \leq \xi \Big\}, \tag{3}$$

where $\text{TV}(P_1, P_2)$ is the total variation distance between distributions P_1 and P_2, and $\widehat{T}_{\text{MLE}}$ denotes the maximum likelihood estimate of the MDP given the offline dataset, $\mathcal{D}$.

Thus, the set defined in Equation 3 contains MDPs which are similar to the maximum likelihood estimate under state-action pairs in $\mathcal{D}$. However, because the expectation in Equation 3 is taken under $\mathcal{D}$, there is no restriction on $\widehat{T}$ for regions of the state-action space not covered by $\mathcal{D}$. We present a brief overview of the theoretical guarantees from [66] in the following subsection.

Remark 1. Note that Problem 1 differs from the pessimistic MDP formulations introduced by MOPO [76] and MOReL [24]. Problem 1 considers the worst-case transition dynamics, while the pessimistic MDPs constructed by MOPO and MOReL only modify the reward function by applying reward penalties for state-action pairs with high uncertainty.

4.1 Theoretical Motivation

The theoretical analysis from [66] shows that solving Problem 1 outputs a policy that with high probability is approximately as good as any policy with a state-action distribution that is covered by the dataset. This is formally stated in the following theorem.

Theorem 1 (PAC guarantee from [66], Theorem 5)**.** *Denote the true MDP transition function by T, and let $\mathcal{M}$ denote a hypothesis class of MDP models such that $T \in \mathcal{M}$. Let π denote the solution to Problem 1 for dataset $\mathcal{D}$. Then with probability $1 - \delta$ for any policy, $\pi^* \in \Pi$, we have*

$$V_T^{\pi^*} - V_T^{\pi} \leq (1-\gamma)^{-2} c_1 \sqrt{C_{\pi^*}} \sqrt{G_{\mathcal{M}_1} + G_{\mathcal{M}_2} + \xi_n^2 + \frac{\ln(c/\delta)}{|\mathcal{D}|}}, \text{ where}$$

$$C_{\pi^*} = \max_{T' \in \mathcal{M}} \frac{\mathbb{E}_{(s,a)\sim \tilde{d}_T^{\pi^*}} \left[\mathrm{TV}(T'(\cdot|s,a), T(\cdot|s,a))^2 \right]}{\mathbb{E}_{(s,a)\sim\rho} \left[\mathrm{TV}(T'(\cdot|s,a), T(\cdot|s,a))^2 \right]},$$

where ρ is the state-action distribution from which $\mathcal{D}$ was sampled, and c and c_1 are universal constants. We refer the reader to Appendix A of [66] for the definitions of $G_{\mathcal{M}_1}$, $G_{\mathcal{M}_2}$, and ξ_n.

The quantity C_{π^*} is upper bounded by the maximum density ratio between the comparator policy, π^*, and the offline distribution, i.e. $C_{\pi^*} \leq \max_{(s,a)} \tilde{d}_T^{\pi^*}(s,a)/\rho(s,a)$. It represents the discrepancy between the distribution of data in the dataset compared to the visitation distribution of policy π^*. Theorem 1 shows that if we find a policy by solving Problem 1, the performance gap of that policy is bounded with respect to any other policy π^* which has a state-action distribution which is covered by the dataset.

Furthermore, the value function under the worst-case model in the set defined by Problem 1 is a lower bound on the value function in the true environment, as stated by Proposition 1.

Proposition 1 (Pessimistic value function)**.** *Let T denote the true transition function for some MDP, and let $\mathcal{M}_\mathcal{D}$ be the set of MDP models defined in Equation 3. Then for any policy π, with probability $1-\delta$ we have that*

$$\min_{\widehat{T} \in \mathcal{M}_\mathcal{D}} V_{\widehat{T}}^\pi \leq V_T^\pi.$$

Proposition 1 follows from the fact that $T \in \mathcal{M}_\mathcal{D}$ with high probability, which is proven in [66] (Appendix E.2). Proposition 1 shows that we can expect the performance of any policy in the true MDP to be at least as good as the value in the worst-case model defined in Problem 1.

While [66] provides the theoretical motivation for solving Problem 1, it does not propose a practical algorithm. In this work, we focus on developing a practical approach to solving Problem 1.

5 RAMBO-RL

In this section, we present Robust Adversarial Model-Based Offline RL (RAMBO), our algorithm for solving Problem 1. The main difficulty with solving Problem 1 is that it is unclear how to find the worst-case MDP in the set defined in Equation 3. To arrive at a scalable solution, we propose a novel approach which is in the spirit of RARL. We alternate between optimising the agent policy to increase the expected value, and adversarially optimising the model to decrease the expected value. In this section, we first describe how we compute the gradient to adversarially train the model. Then, we discuss how to ensure that the model remains approximately within the constraint set defined in Problem 1. Finally, we present our overall algorithm.

5.1 Model Gradient

We propose a policy gradient-inspired approach to adversarially optimise the model. Typically, policy gradient algorithms are used to modify the distribution over actions taken by a policy at each state [61, 62, 71]. In contrast, the update that we propose modifies the likelihood of the successor states and rewards in the MDP model.

We assume that the MDP model is defined by parameters, ϕ, and we write $\widehat{T}_\phi$ to indicate this. We denote by V_ϕ^π the value function for policy π in model $\widehat{T}_\phi$. To approximately find $\min_{\widehat{T}_\phi \in \mathcal{M}_\mathcal{D}} V_\phi^\pi$ as required by Problem 1 via gradient descent, we wish to compute the gradient of the model parameters which reduces the value of the policy within the model, i.e. $\nabla_\phi V_\phi^\pi$.

Proposition 2 (Model Gradient)**.** *Let ϕ denote the parameters of a parametric MDP model $\widehat{T}_\phi$, and let V_ϕ^π denote the value function for policy π in $\widehat{T}_\phi$. Then:*

$$\nabla_\phi V_\phi^\pi = \mathbb{E}_{s\sim d_\phi^\pi, a\sim\pi, (s',r)\sim\widehat{T}_\phi} \left[(r + \gamma V_\phi^\pi(s')) \cdot \nabla_\phi \log \widehat{T}_\phi(s', r|s, a) \right] \tag{4}$$

The proof of Proposition 2 is given in Appendix A. We can subtract the baseline $Q_\phi^\pi(s,a)$ without biasing the gradient estimate (see Appendix A.1 for details):

$$\nabla_\phi V_\phi^\pi = \mathbb{E}_{s\sim d_\phi^\pi, a\sim\pi, (s',r)\sim\widehat{T}_\phi} \left[(r + \gamma V_\phi^\pi(s') - Q_\phi^\pi(s,a)) \cdot \nabla_\phi \log \widehat{T}_\phi(s', r|s, a) \right] \tag{5}$$

The Model Gradient differs from the standard policy gradient in that a) it is used to update the likelihood of successor states in the model, rather than actions in a policy, and b) the "advantage" term $r + \gamma V^\pi_\phi(s') - Q^\pi_\phi(s,a)$ compares the utility of receiving reward r and transitioning to s' to the expected value of state action pair (s,a). In contrast, the standard advantage term, $Q(s,a) - V(s)$ [55], compares the utility of executing state-action pair (s,a) to the expected value at s. To estimate V^π_ϕ and Q^π_ϕ, we use the critic learnt by the actor-critic algorithm used for policy optimisation. Thus, we use the critic *both* for training the policy and adversarially training the model.

Remark 2. The Model Gradient can be thought of as a specific instantiation of the policy gradient if we view $\widehat{T}_\phi$ as a adversarial "policy" on an augmented MDP, M^+, i.e. $\widehat{T}_\phi : S^+ \to \text{Dist}(A^+)$. The augmented state space, $S^+ = S \times A$ includes the state in the original MDP augmented by the action taken by the agent. The augmented action space consists of the reward applied and the successor state, $A^+ = S \times [R_{\text{min}}, R_{\text{max}}]$. Thus, we can think of this approach as an instantiation of RARL in which the adversary policy ($\bar{\pi}$ in Equation 1) that we train is the *model itself*.

Remark 3. In principle, we could have instead adapted backpropagation through time (BPTT) [11, 69] to adversarially train the model. However, BPTT is known to have issues with numerical instability [11].

5.2 Adversarial Model Training

If we were to update the model using Equation 5 alone, this would allow the model to be modified arbitrarily such that the value function in the model is reduced. However, the set of plausible MDPs given by Equation 3 states that over the dataset, $\mathcal{D}$, the model $\widehat{T}_\phi$ should be close to the maximum likelihood estimate, $\widehat{T}_{\text{MLE}}$. Specifically, in the inner optimisation of Problem 1 we wish to find a solution to the constrained optimisation problem

$$\min_{\widehat{T}_\phi} V^\pi_\phi, \quad s.t. \;\; \mathbb{E}_\mathcal{D}\big[\text{TV}(\widehat{T}_{\text{MLE}}(\cdot|s,a), \widehat{T}_\phi(\cdot|s,a))^2\big] \leq \xi. \tag{6}$$

The Lagrangian relaxation leads to the unconstrained problem

$$\max_{\lambda \geq 0} \min_{\widehat{T}_\phi} \Big(L(\hat{T}, \lambda) := V^\pi_\phi + \lambda\big(\mathbb{E}_\mathcal{D}\big[\text{TV}(\widehat{T}_{\text{MLE}}(\cdot|s,a), \widehat{T}_\phi(\cdot|s,a))^2\big] - \xi\big)\Big), \tag{7}$$

where λ is the Lagrange multiplier. Rather than optimising the Lagrange multiplier, we find that in practice fixing λ to apply a constant weighting between the two terms works well with minimal tuning. To facilitate easier tuning of the learning rate, in our implementation we apply the weighting constant to the value function term rather than the model term, which is equivalent up to a scaling factor. This leads to

$$\min_{\widehat{T}_\phi} \Big(\lambda V^\pi_\phi + \mathbb{E}_\mathcal{D}\big[\text{TV}(\widehat{T}_{\text{MLE}}(\cdot|s,a), \widehat{T}_\phi(\cdot|s,a))^2\big]\Big). \tag{8}$$

Rather than explicitly minimising the TV distance between the model and MLE model, we apply the standard MLE loss for the model term. This leads to the final loss function:

$$\mathcal{L}_\phi = \lambda V^\pi_\phi - \mathbb{E}_{(s,a,s') \sim \mathcal{D}}\big[\log \widehat{T}_\phi(s'|s,a)\big]. \tag{9}$$

By minimising the loss function in Equation 9, the model is trained to a) predict the transitions within the dataset, and b) reduce the value function of the policy, with λ determining the tradeoff between these two objectives. Choosing λ to be small ensures that the MLE term dominates for transitions within $\mathcal{D}$, ensuring that the model fits the dataset accurately. Because the MLE term is only computed over $\mathcal{D}$, the value function term dominates outside of the dataset meaning that the model is modified adversarially for transitions outside of the dataset.

To estimate the gradient of the loss function in Equation 9 for stochastic gradient descent, we sample a minibatch of transitions from $\mathcal{D}$ to estimate the MLE term. The gradient for the value function term is computed using the Model Gradient. The transitions used to compute the Model Gradient term must be sampled under the current policy and model (Equation 5). Therefore, to estimate the Model Gradient term, we generate a minibatch of transitions by simulating the current policy in $\widehat{T}_\phi$.

Like previous works [10, 76, 75, 8], we represent the dynamics model using an ensemble of neural networks. Each neural network produces a Gaussian distribution over the next state and reward:

$\widehat{T}_\phi(s', r|s, a) = \mathcal{N}(\mu_\phi(s, a), \Sigma_\phi(s, a))$. A visualisation of the result of adversarially training the dynamics model can be found in Appendix C.4.

Normalisation The composite loss function in Equation 9 comprises two terms which may have different magnitudes across domains depending on the scale of the states and rewards. To enable easier tuning of the adversarial loss weighting, λ, across different domains we perform the following normalisation procedure. Prior to training, we normalise the states in $\mathcal{D}$ in the manner proposed by [14], by subtracting the mean and dividing by the standard deviation of each state dimension in the dataset. Additionally, when computing the gradient in Equation 5 we normalise the advantage terms, $r + \gamma V^\pi_\phi(s') - Q^\pi_\phi(s, a)$, according to the mean and standard deviation across each minibatch. Advantage normalisation is already common practice in policy gradient RL implementations [3, 51].

Algorithm

We are now ready to present our overall approach in Algorithm 1. The first step of RAMBO is to pretrain the environment dynamics model using standard MLE (Line 1). Thereafter, the algorithm follows the format of RARL. At each iteration, we apply gradient updates to the agent to increase the expected value, followed by gradient updates to the model to decrease the expected value.

Prior to each agent update, we generate synthetic k-step rollouts starting from states in $\mathcal{D}$ by simulating rollouts in the current MDP model $\widehat{T}_\phi$. This data is added to the synthetic dataset $\mathcal{D}_{\widehat{T}_\phi}$ (Line 3). Following previous approaches [76, 75, 21] we store data in $\mathcal{D}_{\widehat{T}_\phi}$ in a first in, first out manner so that only data generated from recent iterations is stored in $\mathcal{D}_{\widehat{T}_\phi}$. The agent's policy and value functions are trained using an off-policy actor critic algorithm using samples from $\mathcal{D} \cup \mathcal{D}_{\widehat{T}_\phi}$ (Line 4). In our implementation, we use soft actor-critic (SAC) [18] for agent training.

To update the model to minimise the loss in Equation 9 in Line 5, we sample data from $\mathcal{D}$ to estimate the gradient for the MLE component. To compute the adversarial component, we generate samples by simulating the current policy and model and utilise the value function learnt by the agent to compute the gradient according to Equation 5.

Algorithm 1 RAMBO-RL

Require: Normalised dataset, $\mathcal{D}$;

1: $\widehat{T}_\phi \leftarrow$ MLE dynamics model.
2: **for** $i = 1, 2, \ldots, n_{\text{iter}}$ **do**
3: Generate synthetic k-step rollouts. Add transition data to $\mathcal{D}_{\widehat{T}_\phi}$.
4: *Agent update*: Update π and Q^π_ϕ with an actor critic algorithm, using samples from $\mathcal{D} \cup \mathcal{D}_{\widehat{T}_\phi}$.
5: *Adversarial model update*: Update $\widehat{T}_\phi$ according to Eq. 9, using samples from $\mathcal{D}$ for the MLE component and the current critic Q^π_ϕ and synthetic data sampled from π and $\widehat{T}_\phi$ for the adversarial component.

6 Experiments

In our experiments, we examine how well RAMBO performs compared to state-of-the-art baselines, and explore the importance of adversarial training to the performance of the algorithm. The code for our experiments is included in the supplementary material. We evaluate our approach on the following domains from the D4RL benchmarks [13].

MuJoCo There are three different environments representing different robots (*HalfCheetah, Hopper, Walker2D*), each with 4 datasets (*Random, Medium, Medium-Replay, Medium-Expert*). *Random* contains transitions collected by a random policy. *Medium* contains transitions collected by an early-stopped SAC policy. *Medium-Replay* consists of the replay buffer generated while training the *Medium* policy. The *Medium-Expert* dataset contains a mixture of suboptimal and expert data.

AntMaze The agent controls a robot and navigates to reach a goal, receiving a sparse reward only if the goal is reached. There are three different layouts of maze (*Umaze, Medium, Large*), and different dataset types (*Fixed, Play, Diverse*) which differ in terms of the variety of start and goal locations used to collect the dataset.

Single Transition Example This domain has a one-dimensional state and action space, and several distinct regions of the action space are covered by the dataset. We use this domain to illustrate the difference between RAMBO and the most similar prior algorithm, COMBO. See Appendix C.1 for a more detailed description.

Hyperparameter Selection and Evaluation The base hyperparameters that we use for RAMBO mostly follow those used in SAC [18] and COMBO [75]. We find that the performance of RAMBO is sensitive to the choice of rollout length, k, consistent with findings in previous works [21, 36]. The other critical parameter for RAMBO is the choice of the adversarial weighting, λ.

For each dataset, we choose the rollout length and the adversarial weighting from one of three possible configurations: $(k, \lambda) \in \{(2, 3e\text{-}4), (5, 3e\text{-}4), (5, 0)\}$. We included $(k, \lambda) = (5, 0)$ as we found that an adversarial weighting of 0 worked well for some datasets. For all MuJoCo datasets we performed model rollouts using the current policy, but for some AntMaze datasets we used a random rollout policy as we found that this performed better. Further details about the hyperparameters are in Appendix B.

To evaluate RAMBO, we ran each of the three hyperparameter configurations for five seeds each, and report the best performance across the three configurations. Thus, our evaluation of RAMBO utilises limited online hyperparameter tuning to choose hyperparameters for each dataset. Note that this is the most common practice among existing model-based offline RL algorithms [8, 24, 36, 39, 76]. The performance obtained for each of the hyperparameter configurations is included in Appendix C.2.

Offline hyperparameter selection is an important topic in offline RL [46, 77]. This is because in some applications selecting hyperparameters based on online performance may be infeasible. Therefore, we present additional results in Appendix C.3 where we select between the three choices of hyperparameters offline using a simple heuristic based on the magnitude and stability of the Q-values during training.

Computational Resources During our experiments, each run of RAMBO has access to 2 CPUs of an Intel Xeon Platinum 8259CL processor at 3.1GHz, and half of an Nvidia T4 GPU. With this hardware, each run of RAMBO takes 24-30 hours.

Baselines We compare RAMBO against state-of-the-art model-based (COMBO [75], MOReL [24], and MOPO [76]) and model-free (CQL [30], IQL [27], and TD3+BC [14]) offline RL algorithms. To facilitate a fair comparison, we provide results for all algorithms for the MuJoCo-v2 D4RL datasets which contain more performant data than v0 [6] (details in Appendix B.6). We report results for the AntMaze-v0 datasets.

Results The results in Table 1 show that RAMBO outperforms existing state-of-the-art methods on the MuJoCo locomotion domains. Relative to existing algorithms, RAMBO obtains the best performance on the Medium-Replay datasets, and performs less well on the Medium-Expert datasets. For the high-quality data in the Medium-Expert datasets, simpler approaches such as performing behaviour cloning on the best 10% of trajectories can be used to achieve stronger performance than offline RL methods [27]. Therefore, the suboptimal performance of RAMBO on the Medium-Expert datasets is less of a concern, as applying offline RL algorithms may not be the most suitable approach for these datasets.

For AntMaze, RAMBO performs considerably less well than the model-free baselines. This echoes the findings of [67] that model-based approaches that use forward rollouts struggle to perform well in the AntMaze domains, potentially because model-based algorithms are too aggressive and collide with walls. The authors of [67] show that *reverse* rollouts can lead to stronger performance for model-based methods in these domains. In future work, we wish to investigate whether combining RAMBO with reverse rollouts can lead to stronger performance on these domains.

In Table 2 we present ablations comparing the performance of RAMBO to the same approach without any adversarial updates to the model. These results demonstrate that overall performance degrades if the adversarial training is removed. This parallels previous findings that conservatism is crucial to obtaining strong performance in offline RL. Interestingly however, for some specific datasets we obtain the best performance with no adversarial training (Appendix C.2). This suggests that a potential direction for future work could be trying to identify which types of problems do not require conservatism for a successful policy to be trained offline with model-based RL.

Table 1: Results for the D4RL benchmark using the normalisation procedure proposed by [13]. We report the normalised performance during the last 10 iterations of training averaged over 5 seeds. $\pm$ captures the standard deviation over seeds. Highlighted numbers indicate results within 1% of the most performant algorithm. * indicates the total without random datasets.

		Ours	Model-based baselines			Model-free baselines			
		RAMBO	**COMBO**	**MOPO**	**MOReL**	**CQL**	**IQL**	**TD3+BC**	**BC**
Random	HalfCheetah	39.5 ± 3.5	38.8	35.4	25.6	19.6	-	11.0	2.1
	Hopper	25.4 ± 7.5	17.9	4.1	53.6	6.7	-	8.5	9.8
	Walker2D	0.0 ± 0.3	7.0	4.2	37.3	2.4	-	1.6	1.6
Medium	HalfCheetah	77.9 ± 4.0	54.2	69.5	42.1	49.0	47.4	48.3	36.1
	Hopper	87.0 ± 15.4	94.9	48.0	95.4	66.6	66.3	59.3	29.0
	Walker2D	84.9 ± 2.6	75.5	-0.2	77.8	83.8	78.3	83.7	6.6
Medium Replay	HalfCheetah	68.7 ± 5.3	55.1	68.2	40.2	47.1	44.2	44.6	38.4
	Hopper	99.5 ± 4.8	73.1	39.1	93.6	97.0	94.7	60.9	11.8
	Walker2D	89.2 ± 6.7	56.0	69.4	49.8	88.2	73.9	81.8	11.3
Medium Expert	HalfCheetah	95.4 ± 5.4	90.0	72.7	53.3	90.8	86.7	90.7	35.8
	Hopper	88.2 ± 20.5	111.1	3.3	108.7	106.8	91.5	98.0	111.9
	Walker2D	56.7 ± 39.0	96.1	-0.3	95.6	109.4	109.6	110.1	6.4
MuJoCo-v2 Total:		812.4 ± 47.8	769.7	413.4	773.0	767.4	692.6*	698.5	300.8
AntMaze	Umaze	25.0 ± 12.0	80.3	0.0	0.0	74.0	87.5	78.6	65.0
	Medium-Play	16.4 ± 17.9	0.0	0.0	0.0	61.2	71.2	3.0	0.0
	Large-Play	0.0 ± 0.0	0.0	0.0	0.0	15.8	39.6	0.0	0.0
	Umaze-Diverse	0.0 ± 0.0	57.3	0.0	0.0	84.0	62.2	71.4	55.0
	Medium-Diverse	23.2 ± 14.2	0.0	0.0	0.0	53.7	70.0	10.6	0.0
	Large-Diverse	2.4 ± 3.3	0.0	0.0	0.0	14.9	47.5	0.2	0.0
AntMaze-v0 Total:		67.0 ± 14.9	137.6	0.0	0.0	303.6	378.0	163.8	120.0

Table 2: Ablation of the adversarial updates for RAMBO. These results use the same rollout length for each dataset as RAMBO but with no adversarial updates (i.e. $\lambda = 0$). The scores are averaged over 5 seeds. To produce these results, we gave runs where the value function diverged a score of 0.

RAMBO (No Adversarial Training)			
MuJoCo-v2 Total:	671.5 ± 43.5	**AntMaze-v0 Total:**	45.8 ± 21.8

Table 3: Comparison between RAMBO and COMBO for the Single Transition Example. The results are averaged over 10 seeds. $\pm$ captures the standard deviation over seeds.

RAMBO:	$\mathbf{1.49 \pm 0.04}$	**COMBO:**	1.34 ± 0.14

Table 3 compares the performance of RAMBO and COMBO on the Single Transition Example. Further analysis in Appendix C.1 shows that for this problem, the pessimistic value function updates used by COMBO create local maxima in the Q-function throughout training. Policy optimisation can become stuck in these local maxima. On the other hand, the value function produced by RAMBO is initially optimistic, and pessimism is introduced into the value function *gradually* as the transition function is modified adversarially. As a result, the likelihood of the policy becoming stuck in poor local maxima is reduced. This may explain why RAMBO is able to achieve strong performance relative to existing algorithms in the MuJoCo domains. Gradually increasing the level of pessimism could be a useful modification for existing offline RL algorithms.

7 Conclusion and Future Directions

RAMBO is a promising new approach to offline RL which imposes conservatism by adversarially modifying the transition dynamics of a learnt model. Our approach is theoretically justified, and achieves state-of-the-art performance on standard benchmarks.

There are a number of possible extensions to RAMBO, some of which we have already discussed. In addition, we would like to apply RAMBO to image-space domains by using deep latent variable models to compress the state space [17, 50] and adversarially perturbing the transition dynamics in the latent space representation. Furthermore, we wish to improve the computational efficiency of RAMBO. RAMBO is more computationally demanding than many prior offline RL methods due to the need to simultaneously train the model as well as the policy. Potential directions to improve the efficiency include warmstarting the policy with behaviour cloning [24, 59] or reducing the frequency of model updates. Finally, another potential research direction is to apply the idea of adversarially trained dynamics models to improve robustness in the online setting.

References

[1] 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.

[2] Gaon An, Seungyong Moon, Jang-Hyun Kim, and Hyun Oh Song. Uncertainty-based offline reinforcement learning with diversified q-ensemble. *Advances in neural information processing systems*, 34:7436–7447, 2021.

[3] Marcin Andrychowicz, Anton Raichuk, Piotr Stańczyk, Manu Orsini, Sertan Girgin, Raphaël Marinier, Leonard Hussenot, Matthieu Geist, Olivier Pietquin, Marcin Michalski, Sylvain Gelly, and Olivier Bachem. What matters for on-policy deep actor-critic methods? A large-scale study. In *International Conference on Learning Representations*, 2021.

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

[5] J Andrew Bagnell, Andrew Y Ng, and Jeff G Schneider. Solving uncertain Markov decision processes. Technical report, 2001.

[6] Philip J Ball, Cong Lu, Jack Parker-Holder, and Stephen Roberts. Augmented world models facilitate zero-shot dynamics generalization from a single offline environment. In *International Conference on Machine Learning*, pages 619–629. PMLR, 2021.

[7] David Brandfonbrener, Will Whitney, Rajesh Ranganath, and Joan Bruna. Offline RL without off-policy evaluation. *Advances in Neural Information Processing Systems*, 34:4933–4946, 2021.

[8] Catherine Cang, Aravind Rajeswaran, Pieter Abbeel, and Michael Laskin. Behavioral priors and dynamics models: Improving performance and domain transfer in offline RL. In *Deep RL Workshop NeurIPS 2021*, 2021.

[9] Ching-An Cheng, Tengyang Xie, Nan Jiang, and Alekh Agarwal. Adversarially trained actor critic for offline reinforcement learning. *arXiv preprint arXiv:2202.02446*, 2022.

[10] Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. *Advances in neural information processing systems*, 31, 2018.

[11] Ignasi Clavera, Violet Fu, and Pieter Abbeel. Model-augmented actor-critic: Backpropagating through paths. *International Conference on Learning Representations*, 2020.

[12] Sebastian Curi, Ilija Bogunovic, and Andreas Krause. Combining pessimism with optimism for robust and efficient model-based deep reinforcement learning. In *International Conference on Machine Learning*, pages 2254–2264. PMLR, 2021.

[13] 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.

[14] Scott Fujimoto and Shixiang Shane Gu. A minimalist approach to offline reinforcement learning. *Advances in Neural Information Processing Systems*, 34, 2021.

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

[16] Jakob Gawlikowski, Cedrique Rovile Njieutcheu Tassi, Mohsin Ali, Jongseok Lee, Matthias Humt, Jianxiang Feng, Anna Kruspe, Rudolph Triebel, Peter Jung, Ribana Roscher, et al. A survey of uncertainty in deep neural networks. *arXiv preprint arXiv:2107.03342*, 2021.

[17] David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. *Advances in Neural Information Processing Systems*, 31, 2018.

[18] 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.

[19] Toru Hishinuma and Kei Senda. Weighted model estimation for offline model-based reinforcement learning. *Advances in Neural Information Processing Systems*, 34, 2021.

[20] Garud N Iyengar. Robust dynamic programming. *Mathematics of Operations Research*, 30(2):257–280, 2005.

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

[22] Natasha Jaques, Asma Ghandeharioun, Judy Hanwen Shen, Craig Ferguson, Agata Lapedriza, Noah Jones, Shixiang Gu, and Rosalind Picard. Way off-policy batch deep reinforcement learning of implicit human preferences in dialog. *arXiv preprint arXiv:1907.00456*, 2019.

[23] Kirthevasan Kandasamy, Yoram Bachrach, Ryota Tomioka, Daniel Tarlow, and David Carter. Batch policy gradient methods for improving neural conversation models. *International Conference on Learning Representations*, 2017.

[24] Rahul Kidambi, Aravind Rajeswaran, Praneeth Netrapalli, and Thorsten Joachims. MOReL: Model-based offline reinforcement learning. *Advances in neural information processing systems*, 33:21810–21823, 2020.

[25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 2015.

[26] Ilya Kostrikov, Rob Fergus, Jonathan Tompson, and Ofir Nachum. Offline reinforcement learning with Fisher divergence critic regularization. In *International Conference on Machine Learning*, pages 5774–5783. PMLR, 2021.

[27] Ilya Kostrikov, Ashvin Nair, and Sergey Levine. Offline reinforcement learning with implicit Q-learning. In *International Conference on Learning Representations*, 2022.

[28] Aviral Kumar, Justin Fu, Matthew Soh, George Tucker, and Sergey Levine. Stabilizing off-policy Q-learning via bootstrapping error reduction. *Advances in Neural Information Processing Systems*, 32, 2019.

[29] Aviral Kumar, Anikait Singh, Stephen Tian, Chelsea Finn, and Sergey Levine. A workflow for offline model-free robotic reinforcement learning. *Conference on Robot Learning*, 2021.

[30] Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. Conservative Q-learning for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 33:1179–1191, 2020.

[31] Thanard Kurutach, Ignasi Clavera, Yan Duan, Aviv Tamar, and Pieter Abbeel. Model-ensemble trust-region policy optimization. In *International Conference on Learning Representations*, 2018.

[32] Sascha Lange, Thomas Gabel, and Martin Riedmiller. Batch reinforcement learning. In *Reinforcement learning*, pages 45–73. Springer, 2012.

[33] Byung-Jun Lee, Jongmin Lee, and Kee-Eung Kim. Representation balancing offline model-based reinforcement learning. In *International Conference on Learning Representations*, 2020.

[34] 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.

[35] Yao Liu, Adith Swaminathan, Alekh Agarwal, and Emma Brunskill. Off-policy policy gradient with state distribution correction. *International Conference on Machine Learning RL4RealLife Workshop*, 2019.

[36] Cong Lu, Philip Ball, Jack Parker-Holder, Michael Osborne, and Stephen J Roberts. Revisiting design choices in offline model based reinforcement learning. In *International Conference on Learning Representations*, 2022.

[37] Ajay Mandlekar, Fabio Ramos, Byron Boots, Silvio Savarese, Li Fei-Fei, Animesh Garg, and Dieter Fox. Iris: Implicit reinforcement without interaction at scale for learning control from offline robot manipulation data. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4414–4420. IEEE, 2020.

[38] Ajay Mandlekar, Yuke Zhu, Animesh Garg, Li Fei-Fei, and Silvio Savarese. Adversarially robust policy learning: Active construction of physically-plausible perturbations. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3932–3939. IEEE, 2017.

[39] Tatsuya Matsushima, Hiroki Furuta, Yutaka Matsuo, Ofir Nachum, and Shixiang Gu. Deployment-efficient reinforcement learning via model-based offline optimization. In *International Conference on Learning Representations*, 2021.

[40] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

[41] Ofir Nachum, Bo Dai, Ilya Kostrikov, Yinlam Chow, Lihong Li, and Dale Schuurmans. AlgaeDICE: Policy gradient from arbitrary experience. *arXiv preprint arXiv:1912.02074*, 2019.

[42] Xinkun Nie, Emma Brunskill, and Stefan Wager. Learning when-to-treat policies. *Journal of the American Statistical Association*, 116(533):392–409, 2021.

[43] Arnab Nilim and Laurent El Ghaoui. Robust control of Markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798, 2005.

[44] OpenAI, Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, Jonas Schneider, Nikolas Tezak, Jerry Tworek, Peter Welinder, Lilian Weng, Qiming Yuan, Wojciech Zaremba, and Lei Zhang. Solving rubik's cube with a robot hand. *arXiv preprint*, 2019.

[45] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model's uncertainty? Evaluating predictive uncertainty under dataset shift. *Advances in Neural Information Processing Systems*, 32, 2019.

[46] 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.

[47] Xinlei Pan, Daniel Seita, Yang Gao, and John Canny. Risk averse robust adversarial reinforcement learning. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 8522–8528. IEEE, 2019.

[48] Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning. *arXiv preprint arXiv:1910.00177*, 2019.

[49] Lerrel Pinto, James Davidson, Rahul Sukthankar, and Abhinav Gupta. Robust adversarial reinforcement learning. In *International Conference on Machine Learning*, pages 2817–2826. PMLR, 2017.

[50] 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.

[51] Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann. Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 2021.

[52] Aravind Rajeswaran, Igor Mordatch, and Vikash Kumar. A game theoretic framework for model based reinforcement learning. In *International Conference on Machine Learning*, pages 7953–7963. PMLR, 2020.

[53] Marc Rigter, Bruno Lacerda, and Nick Hawes. Minimax regret optimisation for robust planning in uncertain Markov decision processes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 11930–11938. Association for the Advancement of Artificial Intelligence, 2021.

[54] Aurko Roy, Huan Xu, and Sebastian Pokutta. Reinforcement learning under model mismatch. *Advances in Neural Information Processing Systems*, 30, 2017.

[55] John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. *International Conference on Learning Representations*, 2016.

[56] Susan M Shortreed, Eric Laber, Daniel J Lizotte, T Scott Stroup, Joelle Pineau, and Susan A Murphy. Informing sequential clinical decision-making through reinforcement learning: an empirical study. *Machine learning*, 84(1):109–136, 2011.

[57] Noah Siegel, Jost Tobias Springenberg, Felix Berkenkamp, Abbas Abdolmaleki, Michael Neunert, Thomas Lampe, Roland Hafner, Nicolas Heess, and Martin Riedmiller. Keep doing what worked: Behavior modelling priors for offline reinforcement learning. In *International Conference on Learning Representations*, 2020.

[58] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354–359, 2017.

[59] Avi Singh, Albert Yu, Jonathan Yang, Jesse Zhang, Aviral Kumar, and Sergey Levine. Cog: Connecting new skills to past experience with offline reinforcement learning. *Conference on Robot Learning*, 2020.

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

[61] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

[62] Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. *Advances in Neural Information Processing Systems*, 12, 1999.

[63] Phillip Swazinna, Steffen Udluft, and Thomas Runkler. Overcoming model bias for robust offline deep reinforcement learning. *Engineering Applications of Artificial Intelligence*, 104:104366, 2021.

[64] Aviv Tamar, Shie Mannor, and Huan Xu. Scaling up robust MDPs using function approximation. In *International Conference on Machine Learning*, pages 181–189. PMLR, 2014.

[65] Chen Tessler, Yonathan Efroni, and Shie Mannor. Action robust reinforcement learning and applications in continuous control. In *International Conference on Machine Learning*, pages 6215–6224. PMLR, 2019.

[66] Masatoshi Uehara and Wen Sun. Pessimistic model-based offline reinforcement learning under partial coverage. *International Conference on Learning Representations*, 2022.

[67] 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.

[68] Yue Wang and Shaofeng Zou. Online robust reinforcement learning with model uncertainty. *Advances in Neural Information Processing Systems*, 34, 2021.

[69] Paul J Werbos. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560, 1990.

[70] Wolfram Wiesemann, Daniel Kuhn, and Berç Rustem. Robust Markov decision processes. *Mathematics of Operations Research*, 38(1):153–183, 2013.

[71] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256, 1992.

[72] Yifan Wu, George Tucker, and Ofir Nachum. Behavior regularized offline reinforcement learning. *arXiv preprint arXiv:1911.11361*, 2019.

[73] Tengyang Xie, Ching-An Cheng, Nan Jiang, Paul Mineiro, and Alekh Agarwal. Bellman-consistent pessimism for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 34, 2021.

[74] Yijun Yang, Jing Jiang, Tianyi Zhou, Jie Ma, and Yuhui Shi. Pareto policy pool for model-based offline reinforcement learning. In *International Conference on Learning Representations*, 2021.

[75] Tianhe Yu, Aviral Kumar, Rafael Rafailov, Aravind Rajeswaran, Sergey Levine, and Chelsea Finn. COMBO: Conservative offline model-based policy optimization. *Advances in Neural Information Processing Systems*, 34, 2021.

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

[77] Siyuan Zhang and Nan Jiang. Towards hyperparameter-free policy selection for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 34, 2021.

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] We described two main limitations: a) our approach performs less well on the AntMaze datasets (discussed in experiments), and b) requires more computation than many other approaches (discussed in conclusion).
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (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]
   (b) Did you include complete proofs of all theoretical results? [Yes] We include a complete proof of Proposition 2. For the theoretical results from [66], we provide references to the appropriate parts of that paper.
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 supplementary material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Detailed information about implementation details and hyperparameters is in the appendix.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] Standard deviation with respect to seeds is reported for all results.
   (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] Details of compute used per run is in the experiments section. Details of total compute used for full evaluation is in the appendix.
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 cite the D4RL benchmark datasets [13].
   (b) Did you mention the license of the assets? [Yes] In Appendix B.6.
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] Code for our work is in the supplementary material.
   (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]