

# MICRO: Model-Based Offline Reinforcement Learning with a Conservative Bellman Operator

Xiao-Yin Liu<sup>1,2</sup>, Xiao-Hu Zhou<sup>1,2</sup>\*, Guotao Li<sup>1,2</sup>, Hao Li<sup>1,2</sup>, Mei-Jiang Gui<sup>1,2</sup>, Tian-Yu Xiang<sup>1,2</sup>, De-Xing Huang<sup>1,2</sup> and Zeng-Guang Hou<sup>1,2</sup>\*

<sup>1</sup>State Key Laboratory of Multimodal Artificial Intelligence Systems Institute of Automation, Chinese Academy of Sciences, Beijing 100190, China

<sup>2</sup>The School of Artificial Intelligence University of Chinese Academy of Sciences, Beijing 100049, China  
liuxiaoyin2023@ia.ac.cn, xiaohu.zhou@ia.ac.cn, zengguang.hou@ia.ac.cn

## Abstract

Offline reinforcement learning (RL) faces a significant challenge of distribution shift. Model-free offline RL penalizes the  $Q$  value for out-of-distribution (OOD) data or constrains the policy closed to the behavior policy to tackle this problem, but this inhibits the exploration of the OOD region. Model-based offline RL, which uses the trained environment model to generate more OOD data and performs conservative policy optimization within that model, has become an effective method for this problem. However, the current model-based algorithms rarely consider agent robustness when incorporating conservatism into policy. Therefore, the new **model-based offline** algorithm with a **conservative Bellman operator** (MICRO) is proposed. This method trades off performance and robustness via introducing the robust Bellman operator into the algorithm. Compared with previous model-based algorithms with robust adversarial models, MICRO can significantly reduce the computation cost by only choosing the minimal  $Q$  value in the state uncertainty set. Extensive experiments demonstrate that MICRO outperforms prior RL algorithms in offline RL benchmark and is considerably robust to adversarial perturbations.

## 1 Introduction

Reinforcement learning (RL) has achieved state-of-the-art performance on numerous control decision-making tasks [Singh *et al.*, 2022; Yu *et al.*, 2021a]. However, the trial-and-error procedure in which the agent interacts with the environment directly can be costly and unsafe in some scenarios. Offline RL effectively overcomes this problem by learning policy from static dataset [Levine *et al.*, 2020].

Offline RL suffers from the distribution shift between the learned policy and the behavior policy in the offline dataset due to limited coverage of state-action pairs. To tackle this issue, model-free offline RL algorithms either penalize value function for out-of-distribution (OOD) state-action pairs [Kumar *et al.*, 2020; Kostrikov *et al.*, 2021] or constrain the policy

closed to the behavior policy directly [Kumar *et al.*, 2019; Fujimoto *et al.*, 2019]. However, inhibition of exploration for OOD region hinders the improvement of agent performance.

Model-based offline RL (MBORL) trains the environment model using offline dataset and optimizes the policy based on the constructed environment model. Compared with model-free methods, MBORL can perform better to state not occur in offline dataset via generating more model data [Yu *et al.*, 2020]. However, due to the gap between the estimated and true environment, the generated model data is unreliable, which degrades agent performance. Therefore, conservatism is incorporated into the algorithm through different means.

One type is applying reward penalties for OOD data by uncertainty quantification directly [Yu *et al.*, 2020; Kim and Oh, 2023; Sun *et al.*, 2023] or value function regularization indirectly [Yu *et al.*, 2021b; Liu *et al.*, 2023]. Another type employs a robust adversarial training framework to enforce conservatism by training an adversarial dynamics model [Rigter *et al.*, 2022; Bhardwaj *et al.*, 2023]. However, the dynamics model of the second type is constantly updated during the policy learning process, which demands more computation costs.

Due to the gap between simulation to reality and disturbances in the real world, a robust policy is a necessity in many real-world RL applications. However, the current MBORL algorithms rarely consider agent robustness when improving agent performance. To tackle the gap between the simulation (source) and actual (target) environment, the robust RL based on Robust Markov Decision Process (RMDP) [Nilim and El Ghaoui, 2005] optimizes policy under the worst possible environment [Pinto *et al.*, 2017; Kamalaruban *et al.*, 2020; Panaganti *et al.*, 2022]. Notice that, this is consistent with the core issue of MBORL: the gap between the estimated and the actual environment. Therefore, *is it possible to introduce robust RL into MBORL as a way to improve agent performance and guarantee robustness?*

Motivated by the above question, a new MBORL algorithm is proposed, named MICRO, where a new conservative Bellman operator is designed to incorporate conservatism and guarantee robustness. Compared with [Rigter *et al.*, 2022; Bhardwaj *et al.*, 2023], MICRO only needs to choose the minimal  $Q$  value in the state uncertainty set, which significantly reduces computation cost. Fig. 1 illustrates the framework of MICRO conceptually. The main contributions are summarized

\*Corresponding Authors.

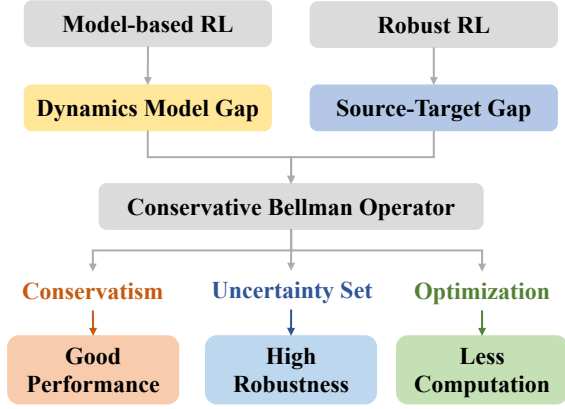


Figure 1: Conceptual illustration of model-based offline RL with a conservative Bellman operator.

as follows:

- MICRO, a novel and theoretically grounded method, is the first MBORL algorithm that trades off performance and robustness by introducing robust Bellman operator.
- A conservative Bellman operator, which combines standard and robust Bellman operators, is designed to impose conservatism and guarantee agent robustness.
- Compared with previous MBORL algorithms, MICRO can guarantee agent robustness with less computation cost while improving performance.

In this paper, the theoretical and experimental results show that MICRO has the guarantee of a robust policy improvement and outperforms current state-of-the-art algorithms on the D4RL dataset benchmark. Furthermore, MICRO achieves better robustness on different adversarial attacks.

## 2 Related Work

**Model-free offline RL:** Online RL algorithms perform poorly in offline RL due to value overestimation and distribution shift. Existing model-free algorithms typically adopt two methods: policy constraint and value regularization. The goal of policy constraint is to restrict the learned policy close to the behavior policy [Kumar *et al.*, 2019; Fujimoto *et al.*, 2019]. In contrast, value regularization methods underestimate the value for OOD data by conservatively estimating  $Q$  values [Kumar *et al.*, 2020; Cheng *et al.*, 2022] or assigning penalization based on the  $Q$  function uncertainty [An *et al.*, 2021; Bai *et al.*, 2021]. However, the root of challenge for offline RL stems from limited data coverage. Model-free offline RL algorithms learn policy only from offline dataset, which hinders agent exploration.

**Model-based offline RL:** Model-based offline RL algorithms train a dynamics model from the offline dataset and utilize the model to broaden data coverage. However, there is a gap between the learned and actual dynamics model due to the inadequate offline dataset. The conservatism should be incorporated into algorithms to prevent agent from areas where

the prediction of the learned dynamic model is unreliable. One method is to penalize the reward by uncertainty quantification [Yu *et al.*, 2020; Kim and Oh, 2023; Sun *et al.*, 2023]. Another method is to enforce small  $Q$ -value for imprecise model data generated by the dynamics model [Yu *et al.*, 2021b; Li *et al.*, 2023; Liu *et al.*, 2023]. Different from the above methods, RAMBO [Rigter *et al.*, 2022] and ARMOR [Bhardwaj *et al.*, 2023] incorporated conservatism by training the robust adversarial model. However, the above two methods demand high computation cost and perform poorly compared with current SOTA methods.

**Robust RL:** The research line of robust RL mainly includes two parts. One line of works focus on the adversarial learning. They applied adversarial attacks to states or actions during policy learning [Pinto *et al.*, 2017; Pattanaik *et al.*, 2018; Tessler *et al.*, 2019]. Another line of works are based on RMDP, which optimize policy under the worst-case environment. The related studies are centered around how to solve minimization problem under uncertainty set [Wang and Zou, 2021; Badrinath and Kalathil, 2021; Panaganti *et al.*, 2022]. Similar to [Kuang *et al.*, 2022], the robust RL in MICRO ultimately only needs to consider the state uncertainty set, which reduces computation complexity.

**This work:** The works that are closest to our work MICRO are by RAMBO [Rigter *et al.*, 2022] and ARMOR [Bhardwaj *et al.*, 2023] that introduces the robust adversarial idea into algorithm. MICRO is significantly different from these in the following ways: 1) MICRO incorporates conservatism into policy by introducing robust Bellman operator, instead of training the adversarial model like RAMBO and ARMOR. 2) MICRO directly selects minimal  $Q$  value in state uncertainty set rather than constantly updating environment model during policy learning, which consumes less computation time. 3) MICRO learns the robust policy against environment perturbations and adversarial attacks, which RAMBO and ARMOR fail to consider. *To the best of our knowledge, MICRO is the first work that incorporates the robust Bellman operator into model-based offline RL, and guarantees the performance and robustness of the learned policy with less computation.*

## 3 Preliminaries

**Markov Decision Process and Offline RL:** An Markov Decision Process (MDP) is defined by the tuple,  $\mathcal{M} = (\mathcal{S}, \mathcal{A}, r, T_{\mathcal{M}}, \mu_0, \gamma)$ , where  $\mathcal{S}$  is the state space,  $\mathcal{A}$  is the action space,  $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$  is the reward function in  $\mathcal{M}$ ,  $T_{\mathcal{M}} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$  is the transition dynamics under  $\mathcal{M}$ ,  $\Delta(\mathcal{S})$  denotes the set of probability distribution over  $\mathcal{S}$ ,  $\mu_0$  is the initial state distribution, and  $\gamma \in (0, 1]$  is the discount factor. A policy  $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$  is a conditional distribution over action given a state. The state-action value function  $Q_{\mathcal{M}}^{\pi}(s, a)$  (i.e.,  $Q$ -function) represents the expected discounted return from taking action  $a$  at state  $s$  and then executing  $\pi$  thereafter in  $\mathcal{M}$ :  $Q_{\mathcal{M}}^{\pi}(s, a) := \mathbb{E}_{\pi, T_{\mathcal{M}}}[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \mid s_0 = s, a_0 = a]$ . The value function is defined as  $V_{\mathcal{M}}^{\pi}(s) := \sum_{a \in \mathcal{A}} \pi(a \mid s) Q_{\mathcal{M}}^{\pi}(s, a)$ .

The goal of RL is to optimize the policy  $\pi$  that maximizes the expected discounted return  $J(\mathcal{M}, \pi) := \sum_{s \in \mathcal{S}} \mu_0(s) V_{\mathcal{M}}^{\pi}(s)$ . Such policy  $\pi$  can be derived from  $Q$ -

learning, which learns the  $Q$ -function that satisfies the following Bellman optimal operator:

$$\mathcal{T}_{\mathcal{M}}Q = r(s, a) + \gamma \mathbb{E}_{s' \sim T_{\mathcal{M}}(s|s, a)} \left[ \max_{a' \in \mathcal{A}} Q(s', a') \right]. \quad (1)$$

where the operator  $\mathcal{T}_{\mathcal{M}}$  is a contraction on  $\mathbb{R}^{|\mathcal{S} \times \mathcal{A}|}$  and  $Q_{\mathcal{M}}^*$  is the fixed point of the above operator. Therefore, the  $Q$ -iteration can be defined as  $Q_{k+1} = \mathcal{T}_{\mathcal{M}}Q_k$ , and it satisfies that  $Q_k \rightarrow Q_{\mathcal{M}}^*$ . The  $Q$ -iteration uses a function class  $\mathcal{Q}$ , defined as  $\mathcal{Q} = \{Q : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1/(1 - \gamma)]\}$ , to approximate  $Q_{\mathcal{M}}^*$ . The sample-based counterpart of Bellman operator  $\mathcal{T}_{\mathcal{M}}$  can be expressed as  $\hat{\mathcal{T}}_{\mathcal{M}}Q(s, a) := r(s, a) + \gamma [\max_{a' \in \mathcal{A}} Q(s', a')]$ . The  $Q$ -function is updated by  $Q_{k+1} = \arg \min_{Q \in \mathcal{Q}} \sum_{i=1}^N [\hat{\mathcal{T}}_{\mathcal{M}}Q_k(s_i, a_i) - Q(s_i, a_i)]^2$ . The dataset  $\{(s_i, a_i, r_i, s'_i)\}_{i=1}^N$  can be a fixed offline dataset or replay buffer generated by the current policy interacting with the environment.

The goal of offline RL is to learn a policy which maximizes  $J(\mathcal{M}, \pi)$  from the fixed dataset  $\mathcal{D}_{\text{offline}} = \{(s_i, a_i, r_i, s'_i)\}_{i=1}^N$  collected by a behavior policy  $\mu$ . The discounted state visitation distribution obtained by rolling out policy  $\pi$  on  $\mathcal{M}$ , denoted as  $d_{\mathcal{M}}^{\pi}(s) := (1 - \gamma) \mathbb{E}_{\pi, T_{\mathcal{M}}} [\sum_{t=0}^{\infty} \gamma^t \mathcal{P}(s_t = s | \pi, \mathcal{M})]$ , where  $\mathcal{P}(s_t = s | \pi, \mathcal{M})$  is the probability of reaching state  $s$  at time  $t$ . The state-action visitation distribution is  $d_{\mathcal{M}}^{\pi}(s, a) = d_{\mathcal{M}}^{\pi}(s) \pi(a | s)$ . We denote the ground truth MDP as  $\mathcal{M}^*$ , so the state-action visitation distribution of a fixed dataset can be written as  $d_{\mathcal{M}^*}^{\mu}(s, a)$ . However, since some  $(s, a)$  barely occurs in  $\mathcal{D}_{\text{offline}}$ , implying  $d_{\mathcal{M}^*}^{\mu}(s, a)$  approaches zero, offline RL tends to suffer from extrapolation error.

**Robust MDP and Robust RL:** The Robust Markov Decision Process (RMDP) considers the discrepancy between simulation and real environments, and gives an *uncertainty set*  $\mathcal{M}_{\varepsilon}$  of possible environments. Following [Abdullah *et al.*, 2019; Kuang *et al.*, 2022], *Wasserstein Distance* of order  $p$  is used to describe uncertainty set  $\mathcal{M}_{\varepsilon}$ , which is defined as

$$\mathcal{M}_{\varepsilon} = \left\{ \hat{\mathcal{M}} : \mathbb{W}_p(T_{\mathcal{M}^*}(\cdot | s, a), T_{\hat{\mathcal{M}}}(\cdot | s, a)) < \varepsilon \right\}. \quad (2)$$

where  $T_{\hat{\mathcal{M}}}(\cdot | s, a)$  and  $T_{\mathcal{M}^*}(\cdot | s, a)$  are the transition dynamics under simulation and real environment, respectively. If the reward function is unknown, since the analysis is similar to the above, we can concatenate the reward onto the state.  $\mathbb{W}_p$  is  $p^{\text{th}}$ -order Wasserstein distance given by

$$\mathbb{W}_p(T_{\hat{\mathcal{M}}}, T_{\mathcal{M}^*}) = \inf_{\kappa \in \mathbf{K}(T_{\hat{\mathcal{M}}}, T_{\mathcal{M}^*})} \left\{ \int_{\mathcal{S} \times \mathcal{S}} d(x, y)^p d\kappa \right\}^{\frac{1}{p}}. \quad (3)$$

where  $\mathbf{K}(T_{\hat{\mathcal{M}}}, T_{\mathcal{M}^*})$  denotes the set of couplings between  $T_{\hat{\mathcal{M}}}$  and  $T_{\mathcal{M}^*}$ ,  $d(\cdot, \cdot) : \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}_+$  is a distance metric for  $\mathcal{S}$  [Hou *et al.*, 2020].

The objective of robust RL is to find the optimal robust policy that maximizes the expected discounted return against the worst possible environment in uncertainty set  $\mathcal{M}_{\varepsilon}$ . The *robust value function*  $V^{\pi}$  and the *robust  $Q$ -function*  $Q^{\pi}$  are defined as  $V^{\pi}(s) = \inf_{\mathcal{M} \in \mathcal{M}_{\varepsilon}} V_{\mathcal{M}}^{\pi}(s)$ ,  $Q^{\pi}(s, a) = \inf_{\mathcal{M} \in \mathcal{M}_{\varepsilon}} Q_{\mathcal{M}}^{\pi}(s, a)$ , respectively [Nilim and El Ghaoui, 2005]. By incorporating uncertainty set into Bellman operator, the *robust Bellman operator* can be denoted as [Iyengar, 2005]

$$\mathcal{T}_{\mathcal{M}_{\varepsilon}}Q = r + \gamma \inf_{\mathcal{M} \in \mathcal{M}_{\varepsilon}} \mathbb{E}_{s' \sim T_{\mathcal{M}}} \left[ \max_{a' \in \mathcal{A}} Q(s', a') \right]. \quad (4)$$

Since the above formulation involves a minimization over the uncertainty set  $\mathcal{M}_{\varepsilon}$ , RMDP is much more computationally intensive compared to standard MDP, which poses a great computational challenge.

**Model-Based Offline RL:** MBORL uses the offline dataset  $\mathcal{D}_{\text{offline}}$  to estimate transition dynamics model  $T_{\hat{\mathcal{M}}}(s' | s, a)$ , which is typically trained via maximum likelihood estimation (MLE) :  $\min_{T_{\hat{\mathcal{M}}}} \mathcal{L}_{\hat{\mathcal{M}}} := \mathbb{E}_{(s, a, s') \sim \mathcal{D}_{\text{offline}}} [-\log T_{\hat{\mathcal{M}}}(s' | s, a)]$ . In the following part, we assume that the reward function  $r$  is known, although the  $r$  can be learned together with dynamics model if unknown. Then, the estimated MDP  $\hat{\mathcal{M}} = (\mathcal{S}, \mathcal{A}, r, T_{\hat{\mathcal{M}}}, \mu_0, \gamma)$  can be constructed to approximate the true MDP  $\mathcal{M}^*$ . The MBORL methods use the estimated MDP  $\hat{\mathcal{M}}$  to generate model data  $\mathcal{D}_{\text{model}}$  through  $h$ -step rollouts, which broadens the data coverage and mitigates the problem of extrapolation errors. The current MBORL algorithms use an augmented dataset  $\mathcal{D}_{\text{offline}} \cup \mathcal{D}_{\text{model}}$  to train the policy, where each data is sampled from the model data  $\mathcal{D}_{\text{model}}$  with the probability  $f$ ,  $f \in [0, 1]$ , and from the offline data  $\mathcal{D}_{\text{offline}}$  with probability  $1 - f$ . However, due to the gap between the estimated and true MDP, some generated model data are inaccurate, leading to poor performance of agent. Therefore, how to strike a balance between accurate offline data and inaccurate model data becomes the key to improving the performance of MBORL algorithms.

## 4 Problem Formulation

To address the MBORL problem proposed in Section 3, conservative policy optimization is performed within that model [Kim and Oh, 2023; Yu *et al.*, 2021b; Yu *et al.*, 2020]. Different from the above methods, [Rigter *et al.*, 2022] and [Bhardwaj *et al.*, 2023] introduced the robust adversarial idea into MBORL. They both incorporate the policy learning information into the dynamics model learning to train the robust dynamics model, which effectively improves agent performance.

However, there are the following drawbacks for the above two methods: 1) The dynamics model is constantly updated during the policy learning process, which needs much more computation cost. 2) They only consider the agent policy improvement and ignore the agent robustness. Therefore, we want to design an MBORL algorithm with good performance, strong robustness and low computation cost. Before this, we first make Assumption 1 that is widely used in the offline RL literature [Chen and Jiang, 2019; Xie *et al.*, 2021; Wang *et al.*, 2021]. This assumption states that there must exist a function  $Q \in \mathcal{Q}$  which can well-approximate  $Q_{\mathcal{M}}^*$  under the operator  $\mathcal{T}_{\mathcal{M}}$ .

**Assumption 1** (Approximate realizability). *Let  $d^{\pi} \in \Delta(\mathcal{S} \times \mathcal{A})$  be the arbitrary data distribution. Then, for any policy  $\pi \in \Pi$ ,  $\inf_{Q \in \mathcal{Q}} \mathbb{E}_{(s, a) \sim d^{\pi}} [(Q - \mathcal{T}_{\mathcal{M}}Q)^2] < \varepsilon_r$  holds.*

Inspired by the goal of robust RL to train a robust policy that can be against the gap between the simulation and the actual environment, we introduce robust RL to MBORL as a way to cope with the gap between the trained and the actual environment in MBORL. [Rigter *et al.*, 2022] gave the following proposition, proven by [Uehara and Sun, 2021], which

shows that the performance of any policy in the true MDP is at least as good as in the worst-case MDP with high probability.

**Proposition 1** (Pessimistic value function). *Let  $\mathcal{M}^*$  be the true MDP and  $\mathcal{M}_\varepsilon$  be the uncertainty set of MDP. Then with high probability, for any policy  $\pi$ ,  $\inf_{\mathcal{M} \in \mathcal{M}_\varepsilon} V_{\mathcal{M}}^\pi \leq V_{\mathcal{M}^*}^\pi$  holds.*

Proposition 1 states that the value under the worst MDP is a lower bound on that in the real MDP, which guarantees the conservatism of value function in uncertainty set  $\mathcal{M}_\varepsilon$ . Therefore, we combine robust RL and MBORL directly and propose a new conservative Bellman operator  $\mathcal{T}Q(s, a)$ .

**Definition 1.** *The conservative Bellman operator  $\mathcal{T}$  is defined as*

$$\mathcal{T}Q = \begin{cases} r + \gamma \mathbb{E}_{s' \sim T_{\mathcal{M}}} \left[ \max_{a' \in \mathcal{A}} Q(s', a') \right], & \mathcal{M} = \mathcal{M}^*. \\ r + \gamma \inf_{\mathcal{M} \in \mathcal{M}_\varepsilon} \mathbb{E}_{s' \sim T_{\mathcal{M}}} \left[ \max_{a' \in \mathcal{A}} Q(s', a') \right], & \mathcal{M} = \widehat{\mathcal{M}}. \end{cases}$$

The above operator expresses: 1) When the data are accurate offline data collected from the true MDP  $\mathcal{M}^*$ , the standard Bellman operator is applied for the computation. 2) When the data are inaccurate model data collected from the estimated MDP  $\widehat{\mathcal{M}}$ , the robust Bellman operator is used to perform conservative computation. The conservative Bellman operator incorporates conservatism into the algorithm, which reduces training risk from model errors. Meanwhile, this operator can guarantee robustness by optimizing policy under worst-case MDP. Compared with [Rigter *et al.*, 2022] and [Bhardwaj *et al.*, 2023], this method only needs to consider the worst MDP and does not require updating the dynamics model constantly, which significantly reduces computation cost. Next, we theoretically analyze dynamic programming properties of  $\mathcal{T}$  in the tabular MDP setting and give the following proposition.

**Proposition 2** ( $\gamma$ -contraction mapping operator). *Let  $\|\cdot\|_\infty$  be the  $\mathcal{L}_\infty$  norm and  $(\mathbb{R}^{|\mathcal{S} \times \mathcal{A}|}, \|\cdot\|_\infty)$  is complete space. Then, the conservative Bellman operator  $\mathcal{T}$  is a  $\gamma$ -contraction mapping operator; i.e.  $\|\mathcal{T}Q_1 - \mathcal{T}Q_2\|_\infty \leq \gamma\|Q_1 - Q_2\|_\infty$ .*

The proof of Proposition 2 is deferred to Appendix A.1. According to the fixed point theorem, any initial  $Q$  function can converge to a unique fixed point by repeatedly applying the conservative Bellman operator  $\mathcal{T}$ . Note that the ground truth MDP  $\mathcal{M}^*$  is unknown in Eq. (2). Here, considering  $\mathcal{M}^* \in \mathcal{M}_\varepsilon$  holds with high probability, the MDP  $\widehat{\mathcal{M}}_{\text{MLE}}$  trained by MLE given offline dataset  $\mathcal{D}_{\text{offline}}$  is employed to approximate the true MDP  $\mathcal{M}^*$ .

However, the computation of operator  $\mathcal{T}$  also faces the following crucial challenge. It is infeasible to consider all MDPs in the uncertainty set  $\mathcal{M}_\varepsilon$  for minimization in Eq. (4). *How to solve the optimization over uncertainty set  $\mathcal{M}_\varepsilon$ ?* For this challenge, [Panaganti *et al.*, 2022] gave the dual reformulation of RMDP and solved the dual problem using empirical risk minimization. However, they introduced a new network that is updated constantly in the solving process, which also demands much computation cost. Therefore, the following section gives a solution for this challenge and derives how to further incorporate robust RL into MBORL.

## 5 Methodology

### 5.1 Reformulation for Conservative Bellman Operator

Following the previous works [Kuang *et al.*, 2022; Blanchet and Murthy, 2019], the dual reformulation of optimization,  $\inf_{\mathcal{M} \in \mathcal{M}_\varepsilon} \mathbb{E}_{s' \sim T_{\mathcal{M}}(s'|s, a)} [\max_{a' \in \mathcal{A}} Q(s', a')]$ , is given below.

**Proposition 3.** *Let  $\mathcal{M}_\varepsilon$  be the uncertainty set defined by (2). Then, for any  $Q : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1/(1 - \gamma)]$ , the term  $\inf_{\mathcal{M} \in \mathcal{M}_\varepsilon} \mathbb{E}_{s' \sim T_{\mathcal{M}}(s'|s, a)} [\max_{a' \in \mathcal{A}} Q(s', a')]$  can be equivalently written as*

$$\sup_{\lambda \geq 0} \left\{ \mathbb{E}_{s' \sim T_{\mathcal{M}^*}} \left[ \inf_{\bar{s} \in \mathcal{S}} \left( \max_{a' \in \mathcal{A}} Q(\bar{s}, a') + \lambda d(s', \bar{s})^p \right) - \lambda \varepsilon \right] \right\}.$$

The proof can be found in Appendix A.2. Note that, the expectation only relies on the true MDP  $\mathcal{M}^*$ . This avoids the demand to consider all MDP in  $\mathcal{M}_\varepsilon$ . However, the above equation involves minimum and maximum optimization, which is also difficult to solve. The above equation can be reduced to the following optimization problem:  $\max_{\lambda} \inf_{x \in \mathcal{S}} \{q(x) + \lambda(d(x) - \varepsilon)\}, \lambda \geq 0$ . The above optimization is essentially a Lagrange dual problem for the following optimization:  $\min_x q(x), d(x) \leq \varepsilon$ . Therefore, the above equation is equivalent to the following form

$$\min_{\bar{s}} \mathbb{E}_{s' \sim T_{\mathcal{M}^*}} \left[ \max_{a' \in \mathcal{A}} Q(\bar{s}, a') \right] \quad \text{s.t.} \quad d(\bar{s}, s')^p \leq \varepsilon. \quad (5)$$

Let  $U_\varepsilon(s')$  be the uncertainty set of  $s'$ , denoted as  $U_\varepsilon(s') = \{\bar{s} \in \mathcal{S} : d(\bar{s}, s')^p \leq \varepsilon\}$ . Then, the Eq. (5) can be written as  $\mathbb{E}_{s' \sim T_{\mathcal{M}^*}} \left\{ \inf_{\bar{s} \in U_\varepsilon(s')} [\max_{a' \in \mathcal{A}} Q(\bar{s}, a')] \right\}$ . Since the  $\widehat{\mathcal{M}}_{\text{MLE}}$  is used to approximate the true  $\mathcal{M}^*$ , the conservative Bellman operator  $\mathcal{T}$  can be expressed as

$$\mathcal{T}Q = \begin{cases} r + \gamma \mathbb{E}_{s' \sim T_{\mathcal{M}}} \left[ \max_{a' \in \mathcal{A}} Q(s', a') \right], & \mathcal{M} = \mathcal{M}^*. \\ r + \gamma \mathbb{E}_{s' \sim T_{\mathcal{M}}} \left\{ \inf_{\bar{s} \in U_\varepsilon(s')} \left[ \max_{a' \in \mathcal{A}} Q(\bar{s}, a') \right] \right\}, & \widehat{\mathcal{M}}_{\text{MLE}}. \end{cases}$$

The operator  $\mathcal{T}$  performs conservative estimation for model data by choosing the worst-case state  $\bar{s}$  from  $U_\varepsilon(s')$ , which also improves the agent robustness. Although the above operator still faces the optimization dilemma, it is easier to analyze the state compared with MDP. Next, the best function mapping  $f(s, a)$  is chosen from the function class  $\mathcal{F} = \{f : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}\}$  that satisfies the following form

$$\mathcal{T}Q = r + \gamma \mathbb{E}_{s' \sim T_{\mathcal{M}}} \left[ \max_{a' \in \mathcal{A}} Q(s', a') - \beta f(s, a) \right]. \quad (6)$$

where  $f(s, a)$  is the penalty term,  $\beta$  is a tuning coefficient. Since  $f(s, a) = 0$  when  $\mathcal{M} = \mathcal{M}^*$ , we only need to consider the solution for  $f \in \mathcal{F}$  under the estimated MDP  $\widehat{\mathcal{M}}_{\text{MLE}}$ . Theoretically, the empirical risk minimization approach is applied to find  $f$ , denoted as

$$f = \arg \min_{f \in \mathcal{F}} \mathbb{E}_{(s, a, s') \sim \mathcal{D}_{\text{model}}} \left\{ \left( \max_{a' \in \mathcal{A}} Q(s', a') - \inf_{\bar{s} \in U_\varepsilon(s')} \left[ \max_{a' \in \mathcal{A}} Q(\bar{s}, a') \right] \right) - \beta f(s, a) \right\}^2. \quad (7)$$

where  $\mathcal{D}_{\text{model}}$  is the dataset collected from  $\widehat{\mathcal{M}}_{\text{MLE}}$ . Due to uncertainty in the model parameters, the ensemble of  $N$  dynamics models,  $\{T_\varphi^i = \mathcal{N}(\mu_\varphi^i, \sigma_\varphi^i)\}_{i=1}^N$ , is trained similar to [Yu *et al.*, 2020; Janner *et al.*, 2019], each of which is a neural network that outputs a Gaussian distribution. Notice that there is a bias in the states predicted by each dynamics model, so the set  $\mathcal{X}(s, a) = \{s' | s' \sim T_\varphi^i(s' | s, a), i = 1, 2, \dots, N\}$  is applied to approximate  $U_\varepsilon(s')$ . Further, the following equation is adopted to approximate  $f(s, a)$ :

$$f(s, a) = \max_{a' \in \mathcal{A}} Q(s', a') - \inf_{\bar{s} \in \mathcal{X}(s, a)} \left[ \max_{a' \in \mathcal{A}} Q(\bar{s}, a') \right], \quad (8)$$

where  $s' \sim \sum_{i=1}^N T_\varphi^i(s' | s, a)/N$ . Since  $\mathcal{X}(s, a)$  still fails to cover all region of the state uncertain set, the coefficient  $\beta$  is introduced in Eq. (6) to adjust the value. From this, the  $f(s, a)$  also reflects the gap of  $Q$ -function value between the worst-case and average level. Obviously, Eq. (6) also satisfies  $\gamma$ -contraction mapping, which guarantees the iteration convergence.

Note that  $f(s, a)$  can be regarded as a type of uncertainty quantification like MOPO [Yu *et al.*, 2020]. However,  $f(s, a)$  contains more robustness information. This term trade off performance and robustness. For example, if  $\varepsilon_1 \leq \varepsilon_2$ , then  $U_{\varepsilon_1}(s') \subseteq U_{\varepsilon_2}(s')$  holds, thereby  $\inf_{\bar{s} \in U_{\varepsilon_1}(s')} [F(\bar{s})] \geq \inf_{\bar{s} \in U_{\varepsilon_2}(s')} [F(\bar{s})]$  holds. Therefore, as  $\varepsilon$  increases, the penalty  $f(s, a)$  for model data increases, which means robustness increases while performance is difficult to improve. The  $\varepsilon$  for different data  $(s, a)$  is adjusted dynamically by model uncertainty indirectly in Eq. (8), instead of the fixed value. This is beneficial to trading off the performance and robustness.

## 5.2 Theoretical Analysis

The theoretical guarantees on *Robust Policy Improvement* property of MICRO are given in this section. First, the policy concentrability coefficient is introduced below [Bhardwaj *et al.*, 2023; Uehara and Sun, 2021].

**Definition 2** (Policy concentrability coefficient). *We define  $\mathcal{C}_{\mathcal{M}_\varepsilon}(\pi)$  to measure how well model errors transfer between the visitation distribution  $d_{\widehat{\mathcal{M}}}^\pi$  under estimated MDP  $\widehat{\mathcal{M}}$  and policy  $\pi$  and offline data distribution  $d_{\mathcal{M}^*}^\mu$ ,*

$$\mathcal{C}_{\mathcal{M}_\varepsilon}(\pi) = \sup_{\widehat{\mathcal{M}} \in \mathcal{M}_\varepsilon} \frac{\mathbb{E}_{(s,a) \sim d_{\widehat{\mathcal{M}}}^\pi} [\text{TV}(T_{\mathcal{M}^*}, T_{\widehat{\mathcal{M}}})]}{\mathbb{E}_{(s,a) \sim d_{\mathcal{M}^*}^\mu} [\text{TV}(T_{\mathcal{M}^*}, T_{\widehat{\mathcal{M}}})]}.$$

where TV is total variation distance, and the  $\mathcal{C}_{\mathcal{M}_\varepsilon}(\pi)$  is upper bounded by  $\|d_{\widehat{\mathcal{M}}}^\pi/d_{\mathcal{M}^*}^\mu\|_\infty$  [Uehara and Sun, 2021]. It represents the discrepancy between the offline data distribution compared to the visitation distribution under policy  $\pi$ .

The optimal policy  $\pi^*$  of standard RL can be obtained by maximizing  $J(\mathcal{M}, \pi) = \mathbb{E}_{(s,a) \sim d_{\mathcal{M}}^\pi} [r(s, a)] / (1 - \gamma)$ . Note that in Eq. (6), the conservative Bellman operator  $\mathcal{T}$  is equivalent to penalize the reward  $r(s, a)$ , i.e.  $\widehat{r}(s, a) = r(s, a) - \gamma\beta f(s, a)$ . Therefore, the output policy  $\pi^\dagger$  of MICRO is expressed as

$$\pi^\dagger = \arg \max_{\pi \in \Pi} \left\{ J(\mathcal{M}_f, \pi) - \frac{\gamma\beta}{1-\gamma} \mathbb{E}_{d_{\mathcal{M}_f}^\pi} [f(s, a)] \right\}. \quad (9)$$

**Algorithm 1** Model-based offline reinforcement learning with a conservative Bellman operator (MICRO)

**Input:** offline dataset  $\mathcal{D}_{\text{offline}}$ , critic  $\{Q_{\omega_i}\}_{i=1}^K$ , policy  $\pi_\theta$ , environment model  $T_\varphi$ , and rollout length  $h$ .

**Output:** the final policy  $\pi^\dagger$  learned by the algorithm MICRO.

- 1: Initialize policy  $\pi_\theta$ , target network  $Q_{\bar{\omega}_i} \leftarrow Q_{\omega_i}$ , and model dataset  $\mathcal{D}_{\text{model}} = \emptyset$ .
- 2: Train an ensemble dynamics model  $\{T_\varphi^i\}_{i=1}^N$  on  $\mathcal{D}_{\text{offline}}$ .
- 3: **for**  $t = 1, 2, \dots, n_{\text{iter}}$  **do**
- 4:   Generate synthetic  $h$ -step rollouts by  $T_\varphi^i$ . Add transition data into  $\mathcal{D}_{\text{model}}$ .
- 5:   Draw samples  $B = (s, a, s', r)$  from  $\mathcal{D}_{\text{offline}} \cup \mathcal{D}_{\text{model}}$ .
- 6:   Compute  $f(s, a)$  according to Eq. (8).
- 7:   Update policy  $\pi_\theta$  and critic network  $Q_{\omega_i}$  using  $B$  according to Eqs. (11) and (12).
- 8:   **if**  $t \% \text{update period} = 0$  **then**
- 9:     Soft update periodically:  $\bar{\omega}_i \leftarrow \tau\omega_i + (1 - \tau)\bar{\omega}_i$ .
- 10:   **end if**
- 11: **end for**

where the  $\mathcal{M}_f = f\widehat{\mathcal{M}} + (1-f)\bar{\mathcal{M}}$  denotes  $f$ -interpolant MDP, which combines sampling error and model error. The  $\bar{\mathcal{M}}$  is empirical MDP that is based on sample data. The  $d_{\mathcal{M}_f}^\pi = fd_{\widehat{\mathcal{M}}}^\pi + (1-f)d_{\bar{\mathcal{M}}}^\pi$  is the mixed distribution of offline and model dataset. Since  $f(s, a) = 0$  when  $\mathcal{M} = \mathcal{M}^*$  in Eq. (9), the final learned policy  $\pi^\dagger$  can be denoted as

$$\pi^\dagger = \arg \max_{\pi \in \Pi} \left\{ J(\mathcal{M}_f, \pi) - \frac{f\gamma\beta}{1-\gamma} \eta(\pi) \right\}. \quad (10)$$

where  $\eta(\pi) = \mathbb{E}_{d_{\widehat{\mathcal{M}}}^\pi(s,a)} [f(s, a)]$  denotes the expectation of the discrepancy of  $\widehat{Q}$  value between the worst-case and average level under estimated MDP  $\widehat{\mathcal{M}}$ . The Bellman operator on empirical MDP  $\bar{\mathcal{M}}$  suffers from sampling error but is unbiased in expectation, whereas Bellman operator on estimated MDP  $\widehat{\mathcal{M}}$  induces bias but no sampling error [Yu *et al.*, 2021b; Liu *et al.*, 2023]. The performance between the true MDP  $\mathcal{M}^*$  and the estimated MDP  $\widehat{\mathcal{M}}$  and the empirical MDP  $\bar{\mathcal{M}}$  under policy  $\pi$  can be found in Lemma 2 and Lemma 3, respectively.

**Theorem 1** (Robust policy improvement). *Under Assumption 1, for any policy  $\pi \in \Pi$ , the policy  $\pi^\dagger$  learned by MICRO in Eq.(10), with high probability  $1 - \delta$ , satisfies that  $J(\mathcal{M}^*, \pi) - J(\mathcal{M}^*, \pi^\dagger)$  is upper bound by*

$$\underbrace{\mathcal{O}\left(\frac{f\gamma\beta}{1-\gamma}\right) \left[ \eta(\pi) - \eta(\pi^\dagger) \right]}_{:=\Delta_1} + \underbrace{\mathcal{O}\left(\frac{1-f}{1-\gamma}\right) \left[ \omega(\pi) + \omega(\pi^\dagger) \right]}_{:=\Delta_2} + \underbrace{\mathcal{O}\left(\frac{fR_{\max}}{(1-\gamma)^2} \sqrt{\frac{2\log(|\mathcal{M}|/\delta)}{n}}\right) \left[ \mathcal{C}_{\mathcal{M}_\varepsilon}(\pi) + \mathcal{C}_{\mathcal{M}_\varepsilon}(\pi^\dagger) \right]}_{:=\Delta_3}.$$

where  $\omega(\pi) = \mathbb{E}_{d_{\widehat{\mathcal{M}}}^\pi(s,a)} [(C_{r,\delta} + R_{\max}C_{T,\delta})/\sqrt{|\mathcal{D}(s, a)|}]$  reflects the sampling error under policy  $\pi$ . The proof of Theorem 1 can be found in Appendix A.3. Theorem 1 proves

Type	Environment	Ours	Model-based Offline RL baseline				Model-free Offline RL baseline			
		MICRO	ARMOR	RAMBO	COUNT	MOBILE	EDAC	ATAC	PBRL	IQL
Medium	Halfcheetah	73.9 $\pm$ 1.7	54.2	<b>77.6</b>	76.5	74.6	65.9	54.3	57.9	47.4
	Hopper	106.5 $\pm$ 0.8	101.4	92.8	103.6	<b>106.6</b>	101.6	102.8	75.3	66.3
	Walker2d	86.6 $\pm$ 1.5	90.7	86.9	87.6	87.7	<b>92.5</b>	<b>91.0</b>	89.6	78.3
Medium Replay	Halfcheetah	71.5 $\pm$ 1.9	50.5	68.9	71.5	<b>71.7</b>	61.3	49.5	45.1	44.2
	Hopper	<b>104.6 <math>\pm</math> 0.9</b>	97.1	96.6	101.7	<b>103.9</b>	101.0	<b>102.8</b>	100.6	94.7
	Walker2d	<b>94.2 <math>\pm</math> 0.9</b>	85.6	85.0	87.7	89.9	87.1	<b>94.1</b>	77.7	73.9
Medium Expert	Halfcheetah	106.8 $\pm$ 1.4	93.5	93.7	100.0	<b>108.2</b>	<b>106.3</b>	95.5	92.3	86.7
	Hopper	<b>113.1 <math>\pm</math> 0.4</b>	103.4	83.3	<b>111.4</b>	<b>112.6</b>	110.7	<b>112.6</b>	110.8	91.5
	Walker2d	<b>116.3 <math>\pm</math> 0.7</b>	112.2	68.3	112.3	<b>115.2</b>	<b>114.7</b>	<b>116.3</b>	110.1	109.6
Average Score		<b>97.1</b>	87.6	83.7	94.7	96.7	93.5	91.0	84.5	76.9

Table 1: Results for the D4RL tasks during the last 5 iterations of training averaged over 4 seeds.  $\pm$  captures the standard deviation over seeds. Bold indicates the highest score. Highlighted numbers indicate results within 2% of the most performant algorithm.

Method	Environment			Average
	Halfcheetah	Hopper	Walker2d	
MICRO	48031.4	45647.6	49271.5	47650.2
RAMBO	133438.0	118020.2	132339.5	127932.7

Table 2: Comparison results of computation time (unit: seconds).

that the policy  $\pi^\dagger$  learned by MICRO can compete with any arbitrary policy  $\pi$  with a large enough dataset [Cheng *et al.*, 2022]. The term  $\Delta_1$  captures the penalty discrepancy of model uncertainty between the policy  $\pi$  and  $\pi^\dagger$ . The terms  $\Delta_2$  and  $\Delta_3$  reflect the sampling error and model error, respectively. When the policy  $\pi$  is the behavior policy  $\mu$ , the generated model data are more accurate under  $\widehat{\mathcal{M}}$  and  $\mu$ , i.e.  $f(s, a)$  approaching zero, which leads to  $\eta(\mu) - \eta(\pi^\dagger) < 0$ . Therefore, the upper bound of  $J(\mathcal{M}^*, \mu) - J(\mathcal{M}^*, \pi^\dagger)$  is more tight.

### 5.3 Practical Implementation

In practical implementation, the value function  $Q$  and the policy  $\pi$  are approximated via neural networks, denoted as  $Q_\omega$  and  $\pi_\theta$ , respectively. The algorithm MICRO, presented in algorithm 1, is built on model-based policy optimization (MBPO) [Janner *et al.*, 2019] and soft actor-critic (SAC) [Haarnoja *et al.*, 2018]. The MLE method is used to train  $N$  ensemble dynamics models  $\{T_\varphi^i = \mathcal{N}(\mu_\varphi^i, \sigma_\varphi^i)\}_{i=1}^N$ . Before agent updating, the model data is collected through  $h$ -step rollouts and is added into model dataset  $\mathcal{D}_{\text{model}}$ . Then, higher coverage dataset,  $\mathcal{D} = \mathcal{D}_{\text{offline}} \cup \mathcal{D}_{\text{model}}$ , are utilized to update agent. The principle of agent updates is described below

$$\begin{aligned}
 Q_{\omega_i} &\leftarrow \arg \min_{\omega_i} \mathbb{E}_{s,a,s' \sim \mathcal{D}} \left[ \left( Q_{\omega_i}(s, a) - \widehat{\mathcal{T}} Q_{\omega_i}(s, a) \right)^2 \right] \\
 \pi_\theta &\leftarrow \arg \max_{\theta} \mathbb{E}_{s,a \sim \mathcal{D}} \left[ \min_{i=1, \dots, K} Q_{\omega_i}(s, a) - \alpha \mathcal{H}(s, a) \right].
 \end{aligned} \tag{11}$$

where  $\widehat{\mathcal{T}} Q_{\omega_i}(s, a)$  is the target value,  $\alpha$  is the entropy regularization coefficient,  $Q_{\omega_i}$  denotes the  $i$ -th  $Q$ -function, and  $K$  is the number of critic. The entropy regularization term

$\mathcal{H}(s, a) := \log \pi_\theta(a | s)$  is applied to prevent policy degradation [Haarnoja *et al.*, 2018]. The minimization operation in above equation is to prevent overestimation of agent training [Van Hasselt *et al.*, 2016]. For any pairs  $(s, a, r, s') \in \mathcal{D}$ , the target value can be defined as

$$\begin{aligned}
 \widehat{\mathcal{T}} Q_{\omega_i}(s, a) &= r(s, a) + \gamma \left[ \min_{i=1, \dots, K} Q_{\omega_i}(s', a') \right. \\
 &\quad \left. - \alpha \mathcal{H}(s', a') - \beta f(s, a) \right].
 \end{aligned} \tag{12}$$

where  $Q_{\omega_i}$  represents the  $i$ -th target  $Q$ -function, and  $\widehat{\mathcal{T}}$  denotes the practical conservative Bellman operator.

## 6 Experiments

This section focuses on the following questions. **Q<sub>1</sub>**: How does MICRO compare with state-of-the-art RL algorithms in standard offline RL benchmarks? **Q<sub>2</sub>**: How does MICRO of computation cost compare with RAMBO and ARMOR? **Q<sub>3</sub>**: Does the method perform better robustness than other RL algorithms? We answer the above questions using D4RL benchmark [Fu *et al.*, 2020] with several control tasks and datasets. More hyperparameters and implementation details are provided in Appendix B. The code for MICRO is available at [github.com/xiaoyinliu0714/MICRO](https://github.com/xiaoyinliu0714/MICRO).

### 6.1 Benchmark Results (Q<sub>1</sub> and Q<sub>2</sub>)

This section compares recent model-free offline RL algorithms, such as ATAC [Cheng *et al.*, 2022], EDAC [An *et al.*, 2021], IQL [Kostrikov *et al.*, 2021], and PBRL [Bai *et al.*, 2021], as well as model-based offline RL algorithms, including ARMOR [Bhardwaj *et al.*, 2023], RAMBO [Rigter *et al.*, 2022], MOBILE [Sun *et al.*, 2023], and Count-MORL [Kim and Oh, 2023] (denoted as COUNT here). These methods are evaluated on Gym domain involving three environments (HalfCheetah, Hopper, and Walker2d) with three types of datasets (medium, medium-replay, medium-expert) for each environment.

Table 1 gives the results of the average normalized score with standard deviation. The results for the above methods



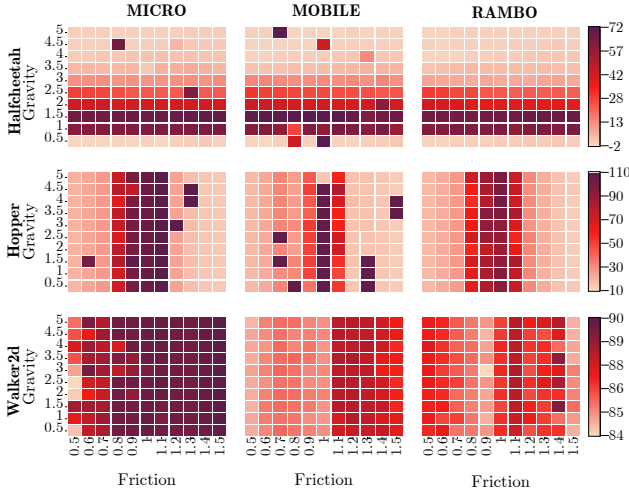


Figure 2: The performance of MICRO, RAMBO and MOBILE under environment parameters perturbation in Halfcheetah, Hopper and Walker2d environments. The gravity and friction vary from 0.5 to 5 times and 0.5 to 1.5 times the value of the simulation environment, respectively.

are obtained from the original papers. The table shows that MICRO outperforms prior RL methods on 4 of total 9 tasks and achieves comparable results in the remaining settings. Overall, MICRO achieves best average performance compared to previous RL algorithms. Table 2 compares computation time between MICRO and RAMBO under three environments (ARMOR is not compared, as code is not open-source). It shows MICRO reduces computation cost by nearly three times.

## 6.2 Environment Parameter Perturbations ( $Q_3$ )

In most cases, it is difficult to achieve high-fidelity simulation environment, which inevitably results in sim-to-real gaps in policy learning. The simulation environment of D4RL tasks is based on the MuJoCo physics simulator. To test the performance of method in high dynamics-gap environment, we modify the parameters of MuJoCo environment, including friction and gravity, to construct high sim-to-real gap environment following [Niu *et al.*, 2022].

Fig. 2 gives the results under various test conditions, where different gravity and friction coefficients are set for the three environments and darker colour means higher performance. The robustness of MICRO is significantly superior to MOBILE and RAMBO in sim-to-real gaps, and MICRO achieves the best generalization in Walker2d environment.

## 6.3 External Adversarial Attacks ( $Q_3$ )

Due to the observation errors and external disturbances within real-world scenarios, three noise perturbation methods, namely **RA**, **AD**, **MQ** following [Yang *et al.*, 2022; Zhang *et al.*, 2020], are used to test the robustness of policy. The observation states are attacked through the above perturbation in evaluation phase. Details about three perturbation types are as follows.

- **RA** randomly samples perturbed states from the ball  $U_\varepsilon(s) = \{s' \in \mathcal{S} : d(s, s')^p \leq \varepsilon\}$ , where  $\varepsilon$  denotes the scale of perturbation.

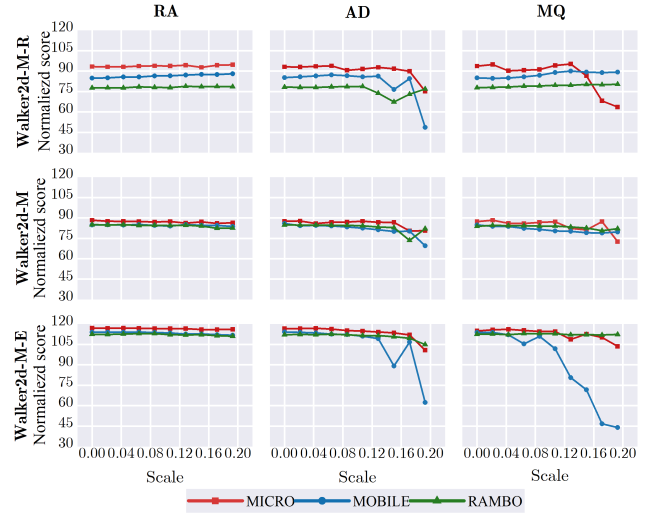


Figure 3: The performance of MICRO, MOBILE and RAMBO under attack scales range [0, 0.2] of different attack types in Walker2d environment. M, M-R and M-E are the abbreviations of Medium, Medium-Replay and Medium-Expert, respectively.

- **AD** is an effective perturbation based on agent policy. It selects the state  $s'$  from  $U_\varepsilon(s)$  that maximizes the gap between action distribution  $\pi_\theta(\cdot|s)$  and  $\pi_\theta(\cdot|s')$ , i.e.  $\min_{s' \in U_\varepsilon(s)} -D_J(\pi_\theta(\cdot|s) || \pi_\theta(\cdot|s'))$ .
- **MQ** performs the perturbation based policy and  $Q$ -function. It finds the state  $s'$  from  $U_\varepsilon(s)$  that minimizes the  $Q$ -function, i.e.  $\min_{s' \in U_\varepsilon(s)} Q(s, \pi_\theta(s'))$ .

Fig. 3 gives the result of different methods with varying perturbation scale  $\varepsilon$  under three perturbation types in Walker2d environment. This figure shows that MICRO achieves better robustness than other methods under **RA** and **AD** attacks. However, the performance of MICRO is inferior to RAMBO in large **MQ** attack. This is mainly because this attack selects the minimum  $Q$ -value, which is more conservative for MICRO. The results in Halfcheetah and Hopper environments can be found in Appendix C. Overall, MICRO greatly outperforms MOBILE in external attacks and achieves superior robustness with less time compared with RAMBO.

## 7 Conclusion

MICRO is a promising novel model-based offline RL algorithm that introduces robust Bellman operator into method, where a conservative Bellman operator is proposed to improve performance and guarantee robustness with less computation. Extensive experiments demonstrate that MICRO can achieve state-of-the-art performance on standard benchmarks and better robustness under different adversarial attacks and perturbations. The main limitation of this work is the robustness slightly performs poorly under strong attacks, which should be improved in future works. MICRO is cutting-edge work in how to trade off performance and robustness in model-based offline RL field. We hope this can provide a promising direction for future research.

## Acknowledgements

This work was supported in part by the National Key Research and Development Program of China under 2023YFC2415100; in part by the National Natural Science Foundation of China under Grant 62222316, Grant 62373351; in part by the Youth Innovation Promotion Association of Chinese Academy of Sciences (CAS) under Grant 2020140; and in part by the CIE-Tencent Robotics X Rhino-Bird Focused Research Program.

## References

- [Abdullah *et al.*, 2019] Mohammed Amin Abdullah, Hang Ren, Haitham Bou Ammar, Vladimir Milenkovic, Rui Luo, Mingtian Zhang, and Jun Wang. Wasserstein robust reinforcement learning. *arXiv preprint arXiv:1907.13196*, 2019.
- [An *et al.*, 2021] 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.
- [Badrinath and Kalathil, 2021] Kishan Panaganti Badrinath and Dileep Kalathil. Robust reinforcement learning using least squares policy iteration with provable performance guarantees. In *International Conference on Machine Learning*, pages 511–520. PMLR, 2021.
- [Bai *et al.*, 2021] Chenjia Bai, Lingxiao Wang, Zhuoran Yang, Zhi-Hong Deng, Animesh Garg, Peng Liu, and Zhao-ran Wang. Pessimistic bootstrapping for uncertainty-driven offline reinforcement learning. In *International Conference on Learning Representations*, 2021.
- [Bhardwaj *et al.*, 2023] Mohak Bhardwaj, Tengyang Xie, Byron Boots, Nan Jiang, and Ching-An Cheng. Adversarial model for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 2023.
- [Blanchet and Murthy, 2019] Jose Blanchet and Karthyek Murthy. Quantifying distributional model risk via optimal transport. *Mathematics of Operations Research*, 44(2):565–600, 2019.
- [Chen and Jiang, 2019] Jinglin Chen and Nan Jiang. Information-theoretic considerations in batch reinforcement learning. In *International Conference on Machine Learning*, pages 1042–1051. PMLR, 2019.
- [Cheng *et al.*, 2022] Ching-An Cheng, Tengyang Xie, Nan Jiang, and Alekh Agarwal. Adversarially trained actor critic for offline reinforcement learning. In *International Conference on Machine Learning*, pages 3852–3878. PMLR, 2022.
- [Fu *et al.*, 2020] 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.
- [Fujimoto *et al.*, 2019] 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.
- [Haarnoja *et al.*, 2018] 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.
- [Hou *et al.*, 2020] Linfang Hou, Liang Pang, Xin Hong, Yanyan Lan, Zhiming Ma, and Dawei Yin. Robust reinforcement learning with wasserstein constraint. *arXiv preprint arXiv:2006.00945*, 2020.
- [Iyengar, 2005] Garud N Iyengar. Robust dynamic programming. *Mathematics of Operations Research*, 30(2):257–280, 2005.
- [Janner *et al.*, 2019] 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.
- [Kamalaruban *et al.*, 2020] Parameswaran Kamalaruban, Yu-Ting Huang, Ya-Ping Hsieh, Paul Rolland, Cheng Shi, and Volkan Cevher. Robust reinforcement learning via adversarial training with langevin dynamics. *Advances in Neural Information Processing Systems*, 33:8127–8138, 2020.
- [Kim and Oh, 2023] Byeongchan Kim and Min-hwan Oh. Model-based offline reinforcement learning with count-based conservatism. In *International Conference on Machine Learning*, 2023.
- [Kostrikov *et al.*, 2021] 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.
- [Kuang *et al.*, 2022] Yufei Kuang, Miao Lu, Jie Wang, Qi Zhou, Bin Li, and Houqiang Li. Learning robust policy against disturbance in transition dynamics via state-conservative policy optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 7247–7254, 2022.
- [Kumar *et al.*, 2019] 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.
- [Kumar *et al.*, 2020] 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.
- [Levine *et al.*, 2020] 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.
- [Li *et al.*, 2023] Hao Li, Xiao-Hu Zhou, Xiao-Liang Xie, Shi-Qi Liu, Zhen-Qiu Feng, Xiao-Yin Liu, Mei-Jiang Gui, Tian-Yu Xiang, De-Xing Huang, Bo-Xian Yao, et al. Crop: Conservative reward for model-based offline policy optimization. *arXiv preprint arXiv:2310.17245*, 2023.
- [Liu *et al.*, 2023] Xiao-Yin Liu, Xiao-Hu Zhou, Xiao-Liang Xie, Shi-Qi Liu, Zhen-Qiu Feng, Hao Li, Mei-Jiang Gui,



- Tian-Yu Xiang, De-Xing Huang, and Zeng-Guang Hou. Domain: Mildly conservative model-based offline reinforcement learning. *arXiv preprint arXiv:2309.08925*, 2023.
- [Nilim and El Ghaoui, 2005] Arnab Nilim and Laurent El Ghaoui. Robust control of markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798, 2005.
- [Niu et al., 2022] Haoyi Niu, Yiwen Qiu, Ming Li, Guyue Zhou, Jianming HU, Xianyuan Zhan, et al. When to trust your simulator: Dynamics-aware hybrid offline-and-online reinforcement learning. *Advances in Neural Information Processing Systems*, 35:36599–36612, 2022.
- [Panaganti et al., 2022] Kishan Panaganti, Zaiyan Xu, Dileep Kalathil, and Mohammad Ghavamzadeh. Robust reinforcement learning using offline data. *Advances in Neural Information Processing Systems*, 35:32211–32224, 2022.
- [Pattanaik et al., 2018] Anay Pattanaik, Zhenyi Tang, Shuijing Liu, Gautham Bommannan, and Girish Chowdhary. Robust deep reinforcement learning with adversarial attacks. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 2040–2042, 2018.
- [Pinto et al., 2017] Lerrel Pinto, James Davidson, Rahul Sukthankar, and Abhinav Gupta. Robust adversarial reinforcement learning. In *International Conference on Machine Learning*, pages 2817–2826. PMLR, 2017.
- [Rigter et al., 2022] Marc Rigter, Bruno Lacerda, and Nick Hawes. Rambo-rl: Robust adversarial model-based offline reinforcement learning. *Advances in Neural Information Processing Systems*, 35:16082–16097, 2022.
- [Singh et al., 2022] Bharat Singh, Rajesh Kumar, and Vinay Pratap Singh. Reinforcement learning in robotic applications: a comprehensive survey. *Artificial Intelligence Review*, pages 1–46, 2022.
- [Sun et al., 2023] Yihao Sun, Jiaji Zhang, Chengxing Jia, Haoxin Lin, Junyin Ye, and Yang Yu. Model-bellman inconsistency for model-based offline reinforcement learning. In *International Conference on Machine Learning*, 2023.
- [Tessler et al., 2019] 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.
- [Uehara and Sun, 2021] Masatoshi Uehara and Wen Sun. Pessimistic model-based offline reinforcement learning under partial coverage. *arXiv preprint arXiv:2107.06226*, 2021.
- [Van Hasselt et al., 2016] Hado Van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.
- [Wang and Zou, 2021] Yue Wang and Shaofeng Zou. Online robust reinforcement learning with model uncertainty. *Advances in Neural Information Processing Systems*, 34:7193–7206, 2021.
- [Wang et al., 2021] Ruosong Wang, Dean Foster, and Sham M Kakade. What are the statistical limits of offline rl with linear function approximation? In *International Conference on Learning Representations*, 2021.
- [Xie et al., 2021] 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*, 34:6683–6694, 2021.
- [Yang et al., 2022] Rui Yang, Chenjia Bai, Xiaoteng Ma, Zhaoran Wang, Chongjie Zhang, and Lei Han. Rorl: Robust offline reinforcement learning via conservative smoothing. *Advances in Neural Information Processing Systems*, 35:23851–23866, 2022.
- [Yu et al., 2020] 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.
- [Yu et al., 2021a] Chao Yu, Jiming Liu, Shamim Nemati, and Guosheng Yin. Reinforcement learning in healthcare: A survey. *ACM Computing Surveys (CSUR)*, 55(1):1–36, 2021.
- [Yu et al., 2021b] 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*, 34:28954–28967, 2021.
- [Zhang et al., 2020] Huan Zhang, Hongge Chen, Chaowei Xiao, Bo Li, Mingyan Liu, Duane Boning, and Cho-Jui Hsieh. Robust deep reinforcement learning against adversarial perturbations on state observations. *Advances in Neural Information Processing Systems*, 33:21024–21037, 2020.