
Offline Reinforcement Learning with Reverse Model-based Imagination

Jianhao Wang^{*1}, Wenzhe Li^{*1}, Haozhe Jiang¹, Guangxiang Zhu^{2†}, Siyuan Li¹,
Chongjie Zhang¹

¹Institute for Interdisciplinary Information Sciences, Tsinghua University, China

²Baidu Inc., China

{wjh19, lwz21, jianghz20}@mails.tsinghua.edu.cn

guangxiangzhu@outlook.com

sy-li17@mails.tsinghua.edu.cn

chongjie@tsinghua.edu.cn

Abstract

In offline reinforcement learning (offline RL), one of the main challenges is to deal with the distributional shift between the learning policy and the given dataset. To address this problem, recent offline RL methods attempt to introduce conservatism bias to encourage learning in high-confidence areas. Model-free approaches directly encode such bias into policy or value function learning using conservative regularizations or special network structures, but their constrained policy search limits the generalization beyond the offline dataset. Model-based approaches learn forward dynamics models with conservatism quantifications and then generate imaginary trajectories to extend the offline datasets. However, due to limited samples in offline datasets, conservatism quantifications often suffer from overgeneralization in out-of-support regions. The unreliable conservative measures will mislead forward model-based imaginations to undesired areas, leading to overaggressive behaviors. To encourage more conservatism, we propose a novel model-based offline RL framework, called *Reverse Offline Model-based Imagination* (ROMI). We learn a reverse dynamics model in conjunction with a novel reverse policy, which can generate rollouts leading to the target goal states within the offline dataset. These reverse imaginations provide informed data augmentation for model-free policy learning and enable conservative generalization beyond the offline dataset. ROMI can effectively combine with off-the-shelf model-free algorithms to enable model-based generalization with proper conservatism. Empirical results show that our method can generate more conservative behaviors and achieve state-of-the-art performance on offline RL benchmark tasks.

1 Introduction

Deep reinforcement learning (RL) has achieved tremendous successes in a range of domains [1–3] by utilizing a large number of interactions with the environment. However, in many real-world applications, collecting sufficient exploratory interactions is usually impractical, because online data collection can be costly or even dangerous, such as in healthcare [4] and autonomous driving [5]. To address this challenge, offline RL [6, 7] develops a new learning paradigm that trains RL agents only with pre-collected offline datasets and thus can abstract away from the cost of online exploration [8–17]. For such offline settings, recent studies demonstrate that directly applying

^{*}Equal contribution.

[†]Work done while Guangxiang was a Ph.D. student at Tsinghua University.

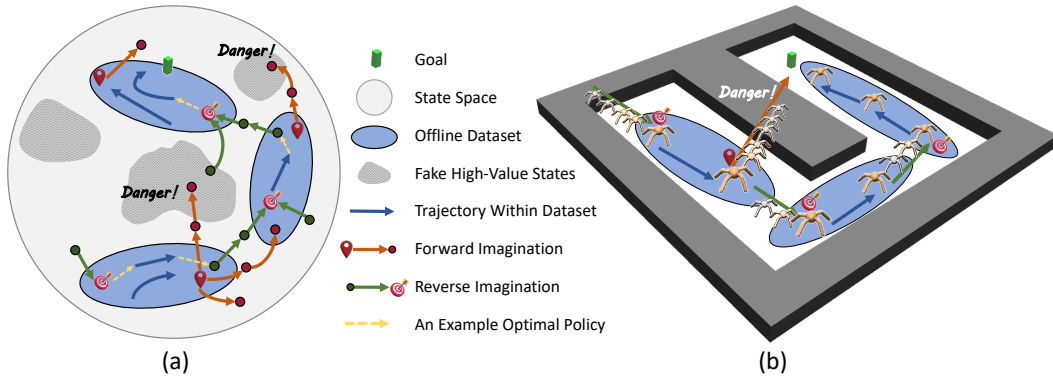


Figure 1: (a) Basic idea of ROMI. (b) A concrete RL task to demonstrate the superiority of ROMI.

the online RL algorithms can lead to poor performance [8, 9, 16]. This phenomenon is primarily attributed to *distributional shift* [7] between the learning policy and the behavior policy induced by the given dataset. With function approximation, the learning policy can overgeneralize the offline dataset and result in unexpected or dangerous behaviors. Thus developing techniques to handle distributional shift is becoming an active topic in the community of offline RL.

Recently, a variety of advanced offline RL algorithms have been proposed, which introduce conservatism bias and constrain the policy search in high-confidence regions induced by the offline dataset. Model-free offline RL methods [8–11] explicitly encode such bias into policy or value functions by using conservative regularizations or specially designed network structures. These methods often effectively address distributional shift issues, but their constrained policy search can limit the generalization beyond the offline dataset. In contrast, model-based offline RL [12–15] adopts more aggressive approaches. They first learn a forward dynamics model from the offline dataset with conservatism quantifications and then generate imaginary trajectories on high confidence regions to extend the offline dataset. Specifically, these methods may use a model-uncertainty quantification [12, 13], representation learning of a robust model [14], or a conservative estimation of value functions [15] to ensure the confidence of model-based rollouts. However, because samples in offline datasets are limited, conservatism quantifications often suffer from overgeneralization, especially in out-of-support regions. Thus, these unreliable measures can overestimate some unknown states and mislead forward model-based imaginations to undesired areas, leading to radicalism. In this paper, we will investigate a new direction in the context of model-based offline RL and propose reverse model imagination that enables effective conservative generalization.

We use Figure 1a to illustrate our basic ideas. When an offline dataset contains expert or nearly optimal trajectories, model-free offline RL methods [8–11] show promising performance. In other situations, such as in Figure 1a, when the optimal policy requires a composition of multiple trajectories (\rightarrow) in the offline dataset, model-free offline RL usually fails because it may get stuck in isolated regions of the offline dataset (\circ). In contrast, model-based approaches [12–15] have advantages of connecting trajectories in the offline dataset by generating bridging rollouts. When using forward dynamics models, model-based methods can generate aggressive rollouts from the dataset to outside areas (\rightarrow). Such *forward imaginations* potentially discover a better policy outside the offline dataset, but may also lead to undesirable regions consisting of fake high-value states (\bullet) due to overgeneralization errors. Things will be different if we reverse the model imagination. *Reverse imaginations* (\rightarrow) generate possible traces leading to target goal states (\odot) inside the offline dataset, which provides a conservative way of augmenting the offline dataset. We assume that the reverse dynamics model has similar accuracy to the forward model. In this way, reverse models not only maintain the generalization ability to interpolate between given trajectories and potential better policies but also avoid radical model imagination by enabling bidirectional search in conjunction with the existing forward real trajectories (\rightarrow) in the offline dataset. The scenario illustrated in Figure 1a is common in RL tasks [18], where the agent may encounter new obstacles or unexpected dangers at any time (e.g., bumping into walls like in Figure 1b or falling). As many unexpected data are not recorded by an offline dataset, there are likely large generalization errors outside the support of the dataset. Thus, reverse models that can generate conservative imaginations are more appealing for real-world offline RL tasks.

Based on the above observation, we present a novel model-based offline RL framework, called *Reverse Offline Model-based Imagination* (ROMI). ROMI trains a reverse dynamics model and generates backward imaginary trajectories by a backtracking rollout policy. This rollout policy is learned by a conditional generative model to produce diverse reverse actions for each state and lead to the unknown data space with high probabilities. ROMI ensures start states of the reverse rollout trajectories (i.e., target goals of forward trajectories) are within the offline dataset, and thus naturally imposes conservative constraints on imaginations. With this conservative data augmentation, ROMI has the advantage of effectively combining with off-the-shelf model-free algorithms (e.g., BCQ [8] and CQL [11]) to further strengthen its generalization with proper conservatism. Taken together, the whole ROMI framework enables extensive interpolation of the dataset and potentially better performance (contributed by diverse policy) in a safe manner (contributed by reverse imagination). To our best knowledge, ROMI is the first offline RL approach that utilizes reverse model-based imaginations to induce conservatism bias with data augmentation. It provides a novel bidirectional learning paradigm for offline RL, which connects reverse imaginary trajectories with pre-collected forward trajectories in the offline dataset. Such a bidirectional learning paradigm shares similar motivations with humans’ bidirectional reasoning. Studies from Psychology [19] show that, during the decision-making process, humans not only consider the consequences of possible future actions from a forward view but also imagine possible traces leading to the ideal goal through backward inference.

We conduct extensive evaluations on the D4RL offline benchmark suite [18]. Empirical results show that ROMI significantly outperforms state-of-the-art model-free and model-based baselines. Our method achieves the best or comparable performance on 16 out of 24 tasks among all algorithms. Ablation studies verify that the reverse model imagination can effectively generate more conservative behaviors than forward model imagination. Videos of the experiments are available online³.

2 Preliminaries

We consider a Markov decision process (MDP) defined by a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, T, r, \mu_0, \gamma)$, where \mathcal{S} and \mathcal{A} denote the state space and the action space, respectively. $T(s'|s, a) : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ denotes the transition distribution function, $r(s, a) : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ denotes the reward function, $\mu_0 : \mathcal{S} \rightarrow [0, 1]$ is the initial state distribution, and $\gamma \in (0, 1)$ is the discount factor. Moreover, we denote the reverse transition distribution by $T_r(s|s', a) = T^{-1} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$. The goal of an RL agent is to optimize a policy $\pi(a|s) : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ that maximizes the expected cumulative reward, i.e., $\mathcal{J}(\pi) = \mathbb{E}_{s_0 \sim \mu_0, s_{t+1} \sim T(\cdot|s_t, \pi(s_t))} [\sum_{t=0}^{\infty} \gamma^t r(s_t, \pi(s_t))]$.

In the offline RL setting, the agent only has access to a static dataset $\mathcal{D}_{\text{env}} = \{(s, a, r, s')\}$ and is not allowed to interact with the environment for additional online explorations. The data can be collected through multi-source logging policies and we denote the empirical distribution of behavior policy in a given dataset \mathcal{D}_{env} collected by π_D . Logging policies are not accessible in our setting.

Model-based RL methods aim at performing planning or policy searches based on a learned model of the environment. They usually learn a dynamics model \hat{T} and a reward model $\hat{r}(s, a)$ from a collection of environmental data in a self-supervised manner. Most of the existing approaches use the forward model $\hat{T}_f(s'|s, a)$ for dynamics learning, but we will show in Section 3 that the reverse dynamics model $\hat{T}_r(s|s', a)$ can induce more conservatism and is critical for the offline RL tasks.

3 Reverse Offline Model-based Imagination

In the offline RL setting, the agent can only access a given dataset without additional online exploration. Model-based offline RL algorithms usually face three challenges in this setting. (i) The offline dataset has limited samples and generalization over the given data is required. (ii) There are a lot of uncertainties that are difficult to estimate out of the support of the dataset. (iii) The model-based rollouts cannot receive feedback from the real environment. Thus it is important to augment diverse data and keep conservative generalization at the same time. In this section, we will introduce the *Reverse Offline Model-based Imagination* (ROMI) framework to combine model-based imagina-

³<https://sites.google.com/view/romi-offlinerl/>.

tion with model-free offline policy learning. This framework encourages diverse augmentation of model-based rollouts and enables conservative generalization of the generated imaginations.

We use a concrete example to show the superiority of our framework in Figure 1b. The agent navigates in a u-maze to reach the goal (■). We gather an offline dataset consisting of three data blocks (●). This dataset only contains the trajectories that do not hit the wall, thus the agent will be unaware of the walls during the offline learning process. In this situation, the conventional forward model may generate a greedy and radical imagination (→) that takes a shortcut to go straight to the goal (whose value function may be overestimated) and hit the wall. On the contrary, ROMI adopts backward imaginations and generates traces that can lead to the target goal states in the support of the dataset (↔). Such imaginations can be further connected with the given forward trajectories (→) in the dataset to form an optimal policy. In this example, we assume that the model error is mainly due to the prediction of the OOD states (outside of the support of the data), then we can expect reverse models to help combat overestimation since they prevent rollout trajectories that end in the OOD states. A more detailed illustrative example is deferred to Appendix A.

Specifically, our framework consists of two main components: (i) a reverse model learned from the offline dataset, (ii) a diverse rollout policy to generate reverse actions that are close to the dataset. ROMI pre-generates the model-based trajectories under the rollout policy based on the reverse model. Then we use the mixture of the imaginary data and the original dataset to train a model-free offline agent. The whole algorithm is shown in Algorithm 1.

Training the reverse model. To backtrace the dynamics and reward function of the environment, we introduce a reverse model to estimate the reverse dynamics model $\widehat{T}_r(s|s', a)$ and reward model $\widehat{r}(s, a)$ from the offline dataset simultaneously. For simplicity, we unify the dynamics and reward function into our reverse model $p(s, r|s', a)$, i.e.,

$$p(s, r|s', a) = p(s|s', a)p(r|s', a, s) = T_r(s|s', a)p(r|s, a), \quad (1)$$

where we assume that the reward function only depends on the current state and action. This unified reverse model represents the probability of the current state and immediate reward conditioned on the next state and current action. We parameterize it by ϕ and optimize it by minimizing the loss function $\mathcal{L}_M(\phi)$ (which is equivalent to maximizing the log-likelihood):

$$\mathcal{L}_M(\phi) = \mathbb{E}_{(s, a, r, s') \sim \mathcal{D}_{\text{env}}} [-\log \widehat{p}_\phi(s, r|s', a)], \quad (2)$$

where \mathcal{D}_{env} is the offline dataset.

Training the reverse rollout policy. To encourage diversity for the reverse model-based imagination near the dataset, we train a generative model $\widehat{G}_\theta(a|s')$, which samples diverse reverse actions from the offline dataset using stochastic inference. Specifically, we use a conditional variational auto-encoder (CVAE) [20, 8] to represent the diverse rollout policy $\widehat{G}_\theta(a|s')$, which is parameterized by θ and depends on the next state. The rollout policy $\widehat{G}_\theta(a|s')$ contains two modules: (i) an action encoder $\widehat{E}_\omega(s', a)$ that outputs a latent vector z under the gaussian distribution $z \sim \widehat{E}_\omega(s', a)$, and (ii) an action decoder $\widehat{D}_\xi(s', z)$ whose input is the latent vector z and reconstructs the given action $\widehat{a} = \widehat{D}_\xi(s', z)$. The action encoder and decoder are parameterized by ω and ξ , respectively. For simplicity, denote the parameters of the rollout policy $\widehat{G}_\theta(a|s')$ by $\theta = \{\omega, \xi\}$. We defer the detailed discussion of CVAE to Appendix B.

We train the rollout policy $\widehat{G}_\theta(a|s')$ by maximizing the variational lower bound $\mathcal{L}_p(\theta)$,

$$\mathcal{L}_p(\theta) = \mathbb{E}_{(s, a, r, s') \sim \mathcal{D}_{\text{env}}, z \sim \widehat{E}_\omega(s', a)} \left[\left(a - \widehat{D}_\xi(s', z) \right)^2 + D_{\text{KL}} \left(\widehat{E}_\omega(s', a) \parallel \mathcal{N}(0, \mathbf{I}) \right) \right], \quad (3)$$

where \mathbf{I} is an identity matrix. To optimize such loss function, we adopt similar optimization techniques as those used in the context of image generalization and video prediction [20, 8].

After the rollout policy $\widehat{G}_\theta(a|s')$ is well trained, we can sample reverse actions \widehat{a} based on the policy. We first draw a latent vector from the multivariate normal distribution, $\widehat{z} \sim \mathcal{N}(0, \mathbf{I})$, and then utilize the action decoder to sample actions conditioned on the next state, $\widehat{a} = \widehat{D}_\xi(s', \widehat{z})$. To explore more possibilities, the rollout policy $\widehat{G}_\theta(a|s')$ uses stochastic layers to generate a variety of reverse actions

for multiple times. In addition, if we need an extremely diverse policy in easy tasks, we can replace the generative model based policy with a uniformly random rollout policy.

Combination with model-free algorithms. Based on the learned reverse dynamics model and the reverse rollout policy, ROMI can generate reverse imaginations. We collect these rollouts to form a model-based buffer $\mathcal{D}_{\text{model}}$ and further compose the total dataset with the offline dataset, $\mathcal{D}_{\text{total}} = \mathcal{D}_{\text{env}} \cup \mathcal{D}_{\text{model}}$. Since our rollout policy is agnostic to policy learning, such buffer can be obtained before the policy learning stage, i.e., ROMI can be combined with any model-free offline RL algorithm (e.g., BCQ [8] or CQL [11]). Specifically, during the model-based imagination, we sample the target state s_{t+1} from the dataset \mathcal{D}_{env} , and generate reverse imaginary trajectory $\hat{\tau}$ with the rollout horizon h by the reverse model \hat{p}_ϕ and rollout policy \hat{G}_θ :

$$\hat{\tau} = \left\{ (s_{t-i}, a_{t-i}, r_{t-i}, s_{t+1-i}) \mid a_{t-i} \sim \hat{G}_\theta(\cdot | s_{t+1-i}) \text{ and } s_{t-i}, r_{t-i} \sim \hat{p}_\phi(\cdot | s_{t+1-i}, a_{t-i}) \right\}_{i=0}^{h-1}.$$

We gather trajectories $\hat{\tau}$ to form the buffer $\mathcal{D}_{\text{model}}$ and further combine it with \mathcal{D}_{env} to obtain $\mathcal{D}_{\text{total}}$. Then we will run the model-free offline policy learning algorithm on the total buffer to derive the final policy π_{out} . Compared to existing model-based approaches [12–15], ROMI provides informed data augmentation to extend the offline dataset. It is agnostic to policy optimization and thus can be regarded as an effective and flexible plug-in component to induce conservative model-based imaginations for offline RL.

Algorithm 1 ROMI: Reverse Offline Model-based Imagination

- 1: **Require:** Offline dataset \mathcal{D}_{env} , rollout horizon h , the number of iterations C_ϕ, C_θ, T , learning rates $\alpha_\phi, \alpha_\theta$, model-free offline RL algorithm (i.e., BCQ or CQL)
 - 2: Randomly initialize reverse model parameters ϕ
 - 3: **for** $i = 0 \dots C_\phi - 1$ **do** ▷ Learning a reverse dynamics model \hat{p}_ϕ
 - 4: Compute \mathcal{L}_M using the dataset \mathcal{D}_{env}
 - 5: Update $\phi \leftarrow \phi - \alpha_\phi \nabla_\phi \mathcal{L}_M$
 - 6: Randomly initialize rollout policy parameters θ
 - 7: **for** $i = 0 \dots C_\theta - 1$ **do** ▷ Learning a diverse rollout policy \hat{G}_θ
 - 8: Compute \mathcal{L}_p using the dataset \mathcal{D}_{env}
 - 9: Update $\theta \leftarrow \theta - \alpha_\theta \nabla_\theta \mathcal{L}_p$
 - 10: Initialize the replay buffer $\mathcal{D}_{\text{model}} \leftarrow \emptyset$
 - 11: **for** $i = 0 \dots T - 1$ **do** ▷ Collecting the replay buffer $\mathcal{D}_{\text{model}}$
 - 12: Sample target state s_{t+1} from the dataset \mathcal{D}_{env}
 - 13: Generate reverse model rollout $\hat{\tau} = \{(s_{t-i}, a_{t-i}, r_{t-i}, s_{t+1-i})\}_{i=0}^{h-1}$ from s_{t+1} by drawing samples from the dynamics model \hat{p}_ϕ and rollout policy \hat{G}_θ
 - 14: Add model rollouts to replay buffer, $\mathcal{D}_{\text{model}} \leftarrow \mathcal{D}_{\text{model}} \cup \{(s_{t-i}, a_{t-i}, r_{t-i}, s_{t+1-i})\}_{i=0}^{h-1}$
 - 15: Compose the final dataset $\mathcal{D}_{\text{total}} \leftarrow \mathcal{D}_{\text{env}} \cup \mathcal{D}_{\text{model}}$
 - 16: Combine model-free offline RL algorithms to derive the final policy π_{out} using the dataset $\mathcal{D}_{\text{total}}$
 - 17: **Return:** π_{out}
-

4 Experiments

In this section, we conduct a bunch of experiments in the offline RL benchmark [18] to answer the following questions: (i) Does ROMI outperform the state-of-the-art offline RL baselines (see Table 1 and 2)? (ii) Does ROMI achieve excellent performance because of the reverse model-based imagination (see Section 4.3)? (iii) Is CVAE-based rollout policy critical for ROMI (see Table 4)? (iv) Compared with the forward imagination, does ROMI trigger more conservative and effective behaviors (see Figure 3)?

4.1 Evaluation Environments

We evaluate ROMI on a wide range of domains in the D4RL benchmark [18], including the Maze2D domain, the Gym-MuJoCo tasks, and the AntMaze domain. Figure 2 shows the snapshots of nine environments used in our experiments. We defer the quantification of ROMI’s model accuracy in

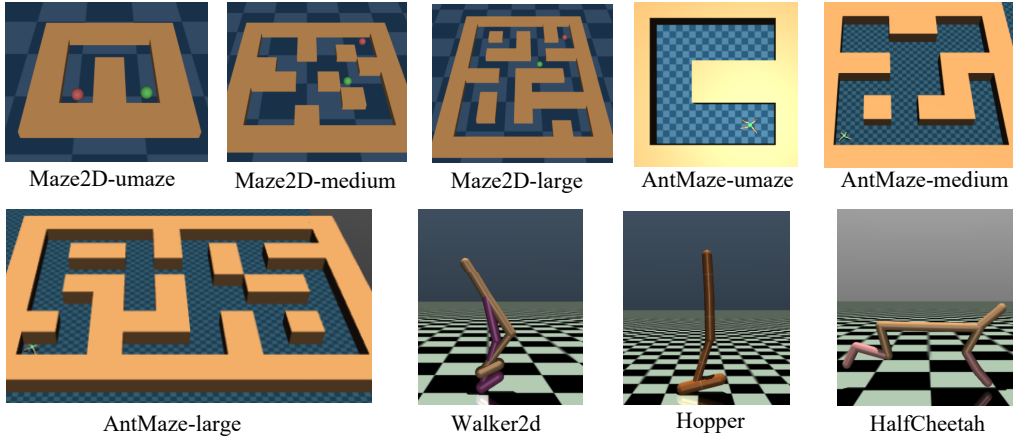


Figure 2: Experimental environments.

these domains to Appendix C and empirical evaluations show that reverse models have comparable accuracy, if not worse, than forward models.

Maze2D. The *maze2d* domain requires a 2D agent to learn to navigate in the maze to reach a fixed target goal and stay there. As shown in Figure 2, there are three maze layouts (i.e., *umaze*, *medium*, and *large*) and two dataset types (i.e., *sparse* and *dense* reward signal) in this domain. The dataset of each layout is generated by a planner moving between randomly sampled waypoints. From the detailed discussion of the D4RL benchmark [18], we found that the agents in the *maze2d* dataset are always moving on the clearing and will not stay in place. We will visualize the dataset of *mazed2d-umaze* in Section 4.5.

Gym-MuJoCo. The Gym-MuJoCo tasks consist of three different environments (i.e., *walker2d*, *hopper*, and *halfcheetah*), and four types of datasets (i.e., *random*, *medium*, *medium-replay*, and *medium-expert*). *Random* dataset contains experiences selected by a random policy. *Medium* dataset contains experiences from an early-stopped SAC policy. *Medium-replay* dataset records the samples in the replay buffer during the training of the "medium" SAC policy. *Medium-expert* dataset is mixed with suboptimal samples and samples generated from an expert policy.

AntMaze. The *antmaze* domain combines challenges of the previous two domains. The policy needs to learn to control the robot and navigate to the goal simultaneously. This domain also contains three different layouts (i.e., *umaze*, *medium*, and *large*) shown in Figure 2. D4RL benchmark [18] introduces three flavors of datasets (i.e., *fixed*, *diverse*, and *play*) in this setting, which commands the ant from different types of start locations to various types of goals.

4.2 Overall Results

In this subsection, the experimental results are presented in Table 1 and 2, which are evaluated in the D4RL benchmark tasks [18] illustrated in Figure 2. We compare ROMI with 11 state-of-the-art baselines: *MF* denotes the best performance from offline model-free algorithms, including BCQ [8], BEAR [10], BRAC-v, BRAC-p [9], BAIL [17], and CQL [11]; *MB* denotes the best performance from offline model-based algorithms, including MOPO [13], MOREL [12], Repb-SDE [14], and COMBO [15]; *BC* denotes the popular behavior cloning from the dataset in the imitation learning. The implementation details of ROMI and these algorithms are deferred to Appendix D.2. Towards fair evaluation, all experimental results are illustrated with the averaged performance with \pm standard deviation over three random seeds.

We evaluate ROMI in nine D4RL benchmark domains with 24 tasks. Our experiments show that ROMI significantly outperforms baselines and achieves the best or comparable performance on 16 out of 24 continuous control tasks. Table 1 and 2 only contain the best performance of *MF* and *MB* categories⁴. We defer the pairwise comparison of ROMI and each baseline to Appendix E, which

⁴In Table 2, each score of COMBO or MOREL is the better one between the score our reproduction and their reported score.

Table 1: Performance of ROMI and best performance of prior methods on the *maze* and *antmaze* domains, on the normalized return metric proposed by D4RL benchmark [18]. Scores roughly range from 0 to 100, where 0 corresponds to a random policy performance and 100 corresponds to an expert policy performance. *med* is short for *medium*.

Environment	BC	ROMI-BCQ	MF	MB
sparse-maze2d-umaze	-3.2	139.5 \pm 3.6	65.7 \pm 6.9 ^{BEAR}	76.4 \pm 19.2 ^{COMBO}
sparse-maze2d-med	-0.5	82.4 \pm 15.2	70.6 \pm 34.3 ^{BRAC-v}	68.5 \pm 83.6 ^{COMBO}
sparse-maze2d-large	-1.7	83.1 \pm 22.1	81.0 \pm 65.3 ^{BEAR}	14.1 \pm 10.7 ^{COMBO}
dense-maze2d-umaze	-6.9	98.3 \pm 2.5	51.5 \pm 8.2 ^{BRAC-p}	94.3 \pm 13.6 ^{Repb-SDE}
dense-maze2d-med	2.7	102.6 \pm 32.4	41.7 \pm 2.0 ^{BAIL}	84.2 \pm 9.5 ^{COMBO}
dense-maze2d-large	-0.3	124 \pm 1.3	133.0 \pm 25.5 ^{BEAR}	36.8 \pm 12.4 ^{MOPO}
fixed-antmaze-umaze	82.0	68.7 \pm 2.7	75.3 \pm 13.7 ^{BCQ}	80.3 \pm 18.5 ^{COMBO}
play-antmaze-med	0.0	35.3 \pm 1.3	1.7 \pm 1.0 ^{BAIL}	0.0
play-antmaze-large	0.0	20.2 \pm 14.8	2.2 \pm 1.3 ^{BAIL}	0.0
diverse-antmaze-umaze	47.0	61.2 \pm 3.3	54.0 \pm 15.0 ^{BAIL}	57.3 \pm 33.6 ^{COMBO}
diverse-antmaze-med	0.0	27.3 \pm 3.9	61.5 \pm 10.0 ^{CQL}	0.0
diverse-antmaze-large	0.0	41.2 \pm 4.2	1.0 \pm 0.9 ^{BAIL}	0.0

Table 2: Performance of ROMI and best performance of prior methods on Gym-MuJoCo tasks.

Environment	BC	ROMI-CQL	MF	MB
random-walker2d	0.0	7.5 \pm 20.0	11.1 \pm 8.8 ^{BEAR}	7.0 ^{COMBO}
random-hopper	0.9	30.2 \pm 4.4	31.4 \pm 0.1 ^{CQL}	31.7 \pm 0.1 ^{Repb-SDE}
random-halfcheetah	-0.1	24.5 \pm 0.7	19.6 \pm 1.2 ^{CQL}	38.8 ^{COMBO}
medium-walker2d	41.7	84.3 \pm 1.1	83.8 \pm 0.2 ^{CQL}	85.3 \pm 2.2 ^{Repb-SDE}
medium-hopper	40.0	72.3 \pm 17.5	66.6 \pm 4.1 ^{CQL}	95.4 ^{MOReL}
medium-halfcheetah	39.2	49.1 \pm 0.8	49.0 \pm 0.4 ^{CQL}	69.5 \pm 0.0 ^{MOPO}
medium-replay-walker2d	2.2	109.7 \pm 9.8	88.4 \pm 1.1 ^{CQL}	83.8 \pm 7.6 ^{Repb-SDE}
medium-replay-hopper	8.1	98.1 \pm 2.6	97.0 \pm 0.8 ^{CQL}	93.6 ^{MOReL}
medium-replay-halfcheetah	25.6	47.0 \pm 0.7	46.4 \pm 0.3 ^{CQL}	68.2 \pm 3.2 ^{MOPO}
medium-expert-walker2d	73.4	109.7 \pm 5.3	109.5 \pm 0.1 ^{CQL}	111.2 \pm 0.2 ^{Repb-SDE}
medium-expert-hopper	36.0	111.4 \pm 5.6	106.8 \pm 2.9 ^{CQL}	111.1 ^{COMBO}
medium-expert-halfcheetah	39.7	86.8 \pm 19.7	90.8 \pm 5.6 ^{CQL}	95.6 ^{MOReL}

can demonstrate that ROMI in conjunction with off-the-shelf model-free methods (i.e., BCQ [8] and CQL [11]) can outperform all offline RL baselines. Specifically, we denote the methods with ROMI as ROMI-BCQ, and ROMI-CQL, respectively. The suffix *-BCQ* or *-CQL* indicates that ROMI adopts BCQ [8] or CQL [11] as base learning algorithms for policy optimization.

Table 1 shows that ROMI-BCQ is the best performer on 10 out of 12 tasks in the *maze2d* and *antmaze* domains. In these tasks, model-free methods can achieve reasonable performance, while current model-based algorithms with forward imagination cannot perform well, especially in *antmaze*. This may be because that messy walls (see Figure 2) are unknown from the given datasets, and the forward imagination may lead RL agent to bump into walls and lose the game. In contrast, reverse model-based imagination can avoid directing agents towards unconfident area, which makes a safe way to imagine in the complex domains. Table 2 shows that ROMI-CQL achieves the best performer on six out of 12 tasks in *gym* domain. CQL is the best performer among all model-free methods, and ROMI-CQL can further outperform CQL on four challenging tasks. In this domain, current model-based methods perform pretty well. We argue that similar to forward imagination, reverse direction can also generalize beyond offline datasets for better performance in the relatively safe tasks, i.e., there is no obstacle around, as shown in Figure 2.

4.3 Ablation Study with Model-based Imagination

In this subsection, we conduct an ablation study to investigate whether ROMI works due to the reverse model-based imagination. Specifically, we replace the reverse imagination with the forward direction in ROMI, which is denoted as *Forward rROMI (FOMI)*. In this subsection, we study the performance of ROMI and FOMI in *maze2d* and *antmaze* domains and defer the ablation study in *gym* domain to Appendix F. Towards fair comparison, we integrate FOMI with BCQ [8] in these settings, called *FOMI-BCQ*. Table 3 shows that ROMI-BCQ significantly outperforms FOMI-BCQ and the base model-free method BCQ [8], which implies that reverse model-based imagination is critical for ROMI in the offline RL settings.

In the *maze2d* domain, compared with the base model-free algorithm BCQ, ROMI-BCQ outperforms all settings, while FOMI-BCQ achieves the superior performance in *maze2d-medium* but performs poorly in the *umaze* and *large* layouts. As illustrated in Figure 2, *maze2d-medium* enjoys less obstacles on the diagonal to the goal than *maze2d-umaze* and *maze2d-large*. In this case, the conservative reverse model-based imagination can enable safe generalization in all layouts. A detailed case study of *maze2d-umaze* will be provided in Section 4.5. Moreover, in the *antmaze* domain, ROMI-BCQ achieve the best performance in the *medium* and *large* layouts, while FOMI-BCQ and BCQ perform well in *antmaze-umaze*. From Figure 2, we find that the *medium* and *large* layouts of *antmaze* have larger mazes with narrower passages, which may frustrate forward imagination and make reverse imagination more effective.

Table 3: Ablation study about ROMI with model-based imagination. Delta equals the improvement of ROMI-BCQ over BCQ on the normalized return metric.

Dataset type	Environment	ROMI-BCQ (ours)	FOMI-BCQ	BCQ (base)	Delta
sparse	maze2d-umaze	139.5 ± 3.6	8.1 ± 15.5	41.1 ± 7.6	98.4
sparse	maze2d-medium	82.4 ± 15.2	93.6 ± 41.3	9.7 ± 14.2	72.7
sparse	maze2d-large	83.1 ± 22.1	-2.5 ± 0.0	38.3 ± 10.4	44.8
dense	maze2d-umaze	98.3 ± 2.5	30.7 ± 0.9	37.0 ± 5.3	61.3
dense	maze2d-medium	102.6 ± 32.4	64.7 ± 37.0	37.9 ± 4.5	64.7
dense	maze2d-large	124.0 ± 1.3	-0.7 ± 7.1	79.8 ± 12.2	44.2
fixed	antmaze-umaze	68.7 ± 2.7	79.5 ± 2.5	75.3 ± 13.7	-6.6
play	antmaze-medium	35.3 ± 1.3	26.2 ± 5.5	0.0	35.3
play	antmaze-large	20.2 ± 14.8	12.0 ± 3.3	0.0	20.2
diverse	antmaze-umaze	61.2 ± 3.3	66.8 ± 3.5	49.3 ± 9.9	11.9
diverse	antmaze-medium	27.3 ± 3.9	12.3 ± 2.1	0.0	27.3
diverse	antmaze-large	41.2 ± 4.2	17.8 ± 2.1	0.0	41.2

4.4 Ablation Study with Different Rollout Policies

In this subsection, we conduct an ablation study to investigate the effect of ROMI’s different rollout policies. To compare with a CVAE-based policy, we propose a new rollout policy (i.e., reverse behavior cloning) for ROMI-BCQ, denoted by *ROMI-RBC-BCQ*. To realize the reverse behavior cloning method, we train a stochastic policy $\hat{\pi}_\varphi(a|s')$, which is parameterized by φ and can sample current action depended on the next state. During training $\hat{\pi}_\varphi(a|s')$, the objective $\mathcal{L}_{rbc}(\varphi)$ is formalized by

$$\mathcal{L}_{rbc}(\varphi) = \mathbb{E}_{(s,a,r,s') \sim \mathcal{D}_{env}} [-\log \hat{\pi}_\varphi(a|s')], \tag{4}$$

where \mathcal{D}_{env} is the given offline dataset, and minimizing the loss function $\mathcal{L}_{rbc}(\varphi)$ is equivalent to maximizing the log-likelihood of probability.

We illustrate the performance of ROMI with different rollout policies in Table 4, where ROMI-BCQ achieves the best performance and ROMI-RBC-BCQ also outperforms BCQ. As suggested by prior generative models [21, 20], in comparison to the deterministic layers (e.g., fully connected layers), CVAE-based methods can generate more diverse and realistic structured samples using stochastic inference. We argue that, the rollout policy implemented by CVAE is critical for ROMI to rollout more diverse trajectories for proposed generalization, and BC-based implementation is also effective in reverse model-based imagination.

Table 4: Ablation study about ROMI with different rollout policies.

Dataset type	Environment	ROMI-BCQ (ours)	ROMI-RBC-BCQ	BCQ (base)
fixed	antmaze-umaze	68.7 \pm 2.7	62.2 \pm 5.6	75.3 \pm 13.7
play	antmaze-medium	35.3 \pm 1.3	33.8 \pm 6.2	0.0
play	antmaze-large	20.2 \pm 14.8	13.3 \pm 16.1	0.0
diverse	antmaze-umaze	61.2 \pm 3.3	43.8 \pm 13.3	49.3 \pm 9.9
diverse	antmaze-medium	27.3 \pm 3.9	20.8 \pm 15.5	0.0
diverse	antmaze-large	41.2 \pm 4.2	14.2 \pm 9.8	0.0

4.5 A Case Study in *maze2d-umaze*

To dive deeper into how ROMI triggers more conservative and effective behaviors, we provide a detailed visual demonstration of one particular task in D4RL benchmark [18]: *maze2d-umaze-sparse*.

As mentioned in Section 4.1, experiences in *maze2d* domain are generated by a planner moving between randomly sampled waypoints on the clearing. Figure 3a shows the movement of the agent from randomly sampled trajectories. To earn high returns, the agent not only needs to learn how to direct to the goal, but also how to stay in the high reward region — the latter behavior is not in the dataset yet. Model-free offline RL algorithms constrain their policy "close" to the dataset, thus it is hard to learn such behaviors out of the support. To see this, Figure 3d shows the behavior of a trained BCQ policy during execution. After reaching the goal, the agent will still oscillate between the high-reward and low-reward regions. This motivates us to use model imagination to generalize beyond the dataset.

As shown in Table 1, ROMI solves this task but previous model-based methods have poor performance, sometimes even worse than model-free algorithms. To better understand this counter-intuitive phenomenon, we compare the rollout trajectories and the learned policy of ROMI-BCQ, FOMI-BCQ (mentioned in Section 4.3), and MOPO [13]. Figure 3(g-i) shows the imagined trajectories in the learning process of the three methods. Figure 3(b,c,e) shows the learned policy behavior at the execution phase. While all model-based imagination will leave the dataset for better generalization, forward model rollout naturally takes some risks as it directs the agent to unknown areas. Undesired forward model imagination will ruin the policy learning (e.g., FOMI-BCQ in Figure 3e and Table 3) or mislead the policy optimization to the suboptimal solution (e.g., MOPO in Figure 3c). Moreover, as shown in Figure 3f, the regularization penalty based on model uncertainty also failed, which is also pointed out in the literature [15]. On the other hand, reverse model imagination inherits the conservatism of the dataset, as it is always ended in the real experience points. Figure 3(b,g) shows that ROMI induces the conservative and optimal behaviors: ROMI will stay around the goal point, and will stick to the data points for higher expected returns. We aim to quantify the aggressiveness of each learned policy and thus define the following trajectory-based discrepancy to measure the distance between learning policy and dataset:

Definition 1 (Average Trajectory Discrepancy). *Given a dataset $\mathcal{D} = \{(s, a, r, s')\}$ and a trajectory $\tau = (s_0, a_0, \dots, s_H, a_H)$, the discrepancy between \mathcal{D} and τ is defined as:*

$$\mathbb{D}(\mathcal{D}, \tau) = \frac{1}{H+1} \sum_{t=0}^H \min_{(s, a, r, s') \in \mathcal{D}} \|s_t - s\|_2. \quad (5)$$

We report the average trajectory discrepancy for policies during execution in Figure 3j and defer results for other maze environments in Appendix G. The results implicate that ROMI is as conservative as model-free method BCQ, but forward model-based offline RL policy (i.e., ROMI and MOPO) tends to deviate from the dataset and will touch the undesirable walls.

5 Related Work

Offline RL. In the literature, offline RL methods can be primarily divided into model-free and model-based algorithms. Prior model-free offline RL methods constrain their policy search in the offline dataset. They can realize their bias by either estimating uncertainty quantification to the value function [7, 9, 10, 16], using importance sampling based algorithms [22–28], explicitly constraining




Figure 3: A case study in *maze2d-umaze*. Note that in figure (f), the penalty is reflected by the color, where blue represents large penalty while white represents small penalty.

the learning policy to be close to the dataset [8, 29], learning the conservative value function [11], and using KL divergence [30, 9, 31] or MMD [10]. On the other hand, prior model-based offline RL methods [12–14, 32–38] have studied in model-uncertainty quantification [12, 13, 39], representation learning [14], constraining the policy to imitate the behavioral policy [35], and using conservative estimation of value function [15]. Different from these works, we propose ROMI to investigate a new direction in the model-based offline RL, which will provide natural conservatism bias with maintaining superior model-based generalization benefits.

Reverse model-based RL. The idea of learning a reverse dynamics model to imagine reversal samples from goal states first emerges in the literature of online RL algorithms [19, 40–43], and shows the potential of speeding up learning and improving sample efficiency by aimed exploration. Similar backward model shows benefit in planning for credit assignment [44, 45] and robustness [46]. Lai et al. (2020) [42] utilizes a backward model to reduce the reliance on accuracy in forward model predictions. In contrast to the backward model in online RL, we show that reverse imagination inherently incorporates the conservatism into rollout trajectories in the offline RL, which can lead rollouts towards target states in dataset. To our best knowledge, ROMI is the first offline RL method with reverse model-based imagination to induce conservatism bias with data augmentation.

6 Conclusion

This paper introduces ROMI, a novel model-based framework for offline RL. To enable conservative generalization, it adopts reverse dynamics model to imagine possible trajectories that can lead to the states within the offline dataset. We demonstrate ROMI can effectively combine with off-the-shelf model-free algorithms and achieve state-of-the-art performance on offline RL benchmark tasks. ROMI leverages good trajectories (e.g., reaching goal states or with high rewards) in an offline dataset to generate effective augmentation data. Although most practical offline datasets contain such trajectories, it is challenging for ROMI to work in other rare cases. One potential future direction is to extend ROMI with forward models, generating bolder imaginations and reaching goal states outside a given dataset. Theoretical formalization of the reverse imagination in offline RL is also an interesting future direction.

Acknowledgments and Disclosure of Funding

The authors would like to thank the anonymous reviewers, Zhizhou Ren, Kun Xu, and Hang Su for valuable and insightful discussions and helpful suggestions. This work is supported in part by Science and Technology Innovation 2030 – “New Generation Artificial Intelligence” Major Project (No. 2018AAA0100904), a grant from the Institute of Guo Qiang, Tsinghua University, and a grant from Turing AI Institute of Nanjing.

References

- [1] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- [2] 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.
- [3] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- [4] Omer Gottesman, Fredrik Johansson, Matthieu Komorowski, Aldo Faisal, David Sontag, Finale Doshi-Velez, and Leo Anthony Celi. Guidelines for reinforcement learning in healthcare. *Nature medicine*, 25(1):16–18, 2019.
- [5] Fisher Yu, Wenqi Xian, Yingying Chen, Fangchen Liu, Mike Liao, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving video database with scalable annotation tooling. *arXiv preprint arXiv:1805.04687*, 2(5):6, 2018.
- [6] Sascha Lange, Thomas Gabel, and Martin Riedmiller. Batch reinforcement learning. In *Reinforcement learning*, pages 45–73. Springer, 2012.
- [7] 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.
- [8] 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.
- [9] Yifan Wu, George Tucker, and Ofir Nachum. Behavior regularized offline reinforcement learning. *arXiv preprint arXiv:1911.11361*, 2019.
- [10] Aviral Kumar, Justin Fu, George Tucker, and Sergey Levine. Stabilizing off-policy q-learning via bootstrapping error reduction. *arXiv preprint arXiv:1906.00949*, 2019.
- [11] Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. Conservative q-learning for offline reinforcement learning. *arXiv preprint arXiv:2006.04779*, 2020.
- [12] Rahul Kidambi and Aravind Rajeswaran. Morel: Model-based offline reinforcement learning. *Neurips*, 2020.
- [13] Tianhe Yu, Garrett Thomas, Lantao Yu, Stefano Ermon, James Zou, Sergey Levine, Chelsea Finn, and Tengyu Ma. Mopo: Model-based offline policy optimization. *arXiv preprint arXiv:2005.13239*, 2020.
- [14] Byung-Jun Lee, Jongmin Lee, and Kee-Eung Kim. Representation balancing offline model-based reinforcement learning. In *International Conference on Learning Representations*, 2021.
- [15] Tianhe Yu, Aviral Kumar, Rafael Rafailov, Aravind Rajeswaran, Sergey Levine, and Chelsea Finn. Combo: Conservative offline model-based policy optimization. *arXiv preprint arXiv:2102.08363*, 2021.

- [16] 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.
- [17] Xinyue Chen, Zijian Zhou, Zheng Wang, Che Wang, Yanqiu Wu, and Keith Ross. Bail: Best-action imitation learning for batch deep reinforcement learning. *Advances in Neural Information Processing Systems*, 33, 2020.
- [18] 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.
- [19] Keith J Holyoak and Dan Simon. Bidirectional reasoning in decision making by constraint satisfaction. *Journal of Experimental Psychology: General*, 128(1):3, 1999.
- [20] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems*, 28:3483–3491, 2015.
- [21] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [22] Doina Precup, Richard S Sutton, and Sanjoy Dasgupta. Off-policy temporal-difference learning with function approximation. In *ICML*, pages 417–424, 2001.
- [23] Richard S Sutton, A Rupam Mahmood, and Martha White. An emphatic approach to the problem of off-policy temporal-difference learning. *The Journal of Machine Learning Research*, 17(1):2603–2631, 2016.
- [24] Yao Liu, Adith Swaminathan, Alekh Agarwal, and Emma Brunskill. Off-policy policy gradient with stationary distribution correction. In *Uncertainty in Artificial Intelligence*, pages 1180–1190. PMLR, 2020.
- [25] Adith Swaminathan and Thorsten Joachims. Batch learning from logged bandit feedback through counterfactual risk minimization. *The Journal of Machine Learning Research*, 16(1):1731–1755, 2015.
- [26] 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.
- [27] Avi Singh, Albert Yu, Jonathan Yang, Jesse Zhang, Aviral Kumar, and Sergey Levine. Cog: Connecting new skills to past experience with offline reinforcement learning. *arXiv preprint arXiv:2010.14500*, 2020.
- [28] 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.
- [29] Yao Liu, Adith Swaminathan, Alekh Agarwal, and Emma Brunskill. Provably good batch reinforcement learning without great exploration. *arXiv preprint arXiv:2007.08202*, 2020.
- [30] 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.
- [31] Wenxuan Zhou, Sujay Bajracharya, and David Held. Plas: Latent action space for offline reinforcement learning. *arXiv preprint arXiv:2011.07213*, 2020.
- [32] Chelsea Finn and Sergey Levine. Deep visual foresight for planning robot motion. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2786–2793. IEEE, 2017.
- [33] Frederik Ebert, Chelsea Finn, Sudeep Dasari, Annie Xie, Alex Lee, and Sergey Levine. Visual foresight: Model-based deep reinforcement learning for vision-based robotic control. *arXiv preprint arXiv:1812.00568*, 2018.

- [34] Gregory Kahn, Adam Villaflor, Pieter Abbeel, and Sergey Levine. Composable action-conditioned predictors: Flexible off-policy learning for robot navigation. In *Conference on Robot Learning*, pages 806–816. PMLR, 2018.
- [35] Tatsuya Matsushima, Hiroki Furuta, Yutaka Matsuo, Ofir Nachum, and Shixiang Gu. Deployment-efficient reinforcement learning via model-based offline optimization. *arXiv preprint arXiv:2006.03647*, 2020.
- [36] Arthur Argenson and Gabriel Dulac-Arnold. Model-based offline planning. *arXiv preprint arXiv:2008.05556*, 2020.
- [37] Phillip Swazinna, Steffen Udluft, and Thomas Runkler. Overcoming model bias for robust offline deep reinforcement learning. *arXiv preprint arXiv:2008.05533*, 2020.
- [38] Rafael Rafailov, Tianhe Yu, Aravind Rajeswaran, and Chelsea Finn. Offline reinforcement learning from images with latent space models. *arXiv preprint arXiv:2012.11547*, 2020.
- [39] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua V Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. *arXiv preprint arXiv:1906.02530*, 2019.
- [40] Anirudh Goyal, Philemon Brakel, William Fedus, Soumye Singhal, Timothy Lillicrap, Sergey Levine, Hugo Larochelle, and Yoshua Bengio. Recall traces: Backtracking models for efficient reinforcement learning. *arXiv preprint arXiv:1804.00379*, 2018.
- [41] Ashley D Edwards, Laura Downs, and James C Davidson. Forward-backward reinforcement learning. *arXiv preprint arXiv:1803.10227*, 2018.
- [42] Hang Lai, Jian Shen, Weinan Zhang, and Yong Yu. Bidirectional model-based policy optimization. In *International Conference on Machine Learning*, pages 5618–5627. PMLR, 2020.
- [43] Kimin Lee, Younggyo Seo, Seunghyun Lee, Honglak Lee, and Jinwoo Shin. Context-aware dynamics model for generalization in model-based reinforcement learning. In *International Conference on Machine Learning*, pages 5757–5766. PMLR, 2020.
- [44] Hado P van Hasselt, Matteo Hessel, and John Aslanides. When to use parametric models in reinforcement learning? *Advances in Neural Information Processing Systems*, 32:14322–14333, 2019.
- [45] Veronica Chelu, Doina Precup, and Hado van Hasselt. Forethought and hindsight in credit assignment. *arXiv preprint arXiv:2010.13685*, 2020.
- [46] Taher Jafferjee, Ehsan Imani, Erin Talvitie, Martha White, and Micheal Bowling. Hallucinating value: A pitfall of dyna-style planning with imperfect environment models. *arXiv preprint arXiv:2006.04363*, 2020.
- [47] Michael Janner, Justin Fu, Marvin Zhang, and Sergey Levine. When to trust your model: Model-based policy optimization. *arXiv preprint arXiv:1906.08253*, 2019.