
Bidirectional Model-based Policy Optimization

Hang Lai¹ Jian Shen¹ Weinan Zhang¹ Yong Yu¹

Abstract

Model-based reinforcement learning approaches leverage a forward dynamics model to support planning and decision making, which, however, may fail catastrophically if the model is inaccurate. Although there are several existing methods dedicated to combating the model error, the potential of the single forward model is still limited. In this paper, we propose to additionally construct a backward dynamics model to reduce the reliance on accuracy in forward model predictions. We develop a novel method, called Bidirectional Model-based Policy Optimization (BMPO) to utilize both the forward model and backward model to generate short branched rollouts for policy optimization. Furthermore, we theoretically derive a tighter bound of return discrepancy, which shows the superiority of BMPO against the one using merely the forward model. Extensive experiments demonstrate that BMPO outperforms state-of-the-art model-based methods in terms of sample efficiency and asymptotic performance.

1. Introduction

Reinforcement learning (RL) methods are commonly divided into two categories: (1) model-free RL (MFRL) which directly learns a policy or value function from observation data, and (2) model-based RL (MBRL) which builds a predictive model of the environment and generates samples from it to derive a policy or a controller. While MFRL algorithms have achieved remarkable success in different ranges of areas (Mnih et al., 2015; Lillicrap et al., 2015; Schulman et al., 2017), they need a large number of samples, which limit their applications mostly to simulators. Model-based methods, in contrast, have shown great potential in reducing the sample complexity (Deisenroth et al., 2013). However, the asymptotic performance of MBRL methods often lags

¹Shanghai Jiao Tong University, Shanghai, China. Correspondence to: Hang Lai <laih@apex.sjtu.edu.cn>, Weinan Zhang <wnzhang@sjtu.edu.cn>.

behind their model-free counterparts due to model error, which is especially severe for multi-step rollout because of the compounding error (Asadi et al., 2018b).

Several previous works have been proposed to alleviate compounding model error in different ways. For example, Whitney & Fergus (2018) and Asadi et al. (2019) introduced multi-step models which directly predict the consequence of executing a sequence of actions; Mishra et al. (2017) tried to divide trajectories into temporal segments and make predictions over segments instead of one timestep; Talvitie (2017), Kaiser et al. (2019) and Freeman et al. (2019) proposed to train the dynamics model on its own outputs, hoping the model is capable of giving accurate predictions in the region of its outputs. Wu et al. (2019) attributed long horizon rollouts mismatching to the one-step supervised learning and proposed to learn the model via model imitation. Besides the above innovations in model learning, Nguyen et al. and Xiao et al. (2019) developed methods of using adaptive rollout horizon according to the estimated compounding error, and Janner et al. (2019) proposed to use truncated short rollouts branched from real states. However, even with the compounding model error mitigated, it is almost impossible to obtain perfect multi-step predictions in practice, and thus the potential of the single forward model is still limited.

In this paper, we investigate how to maximally take advantage of the underlying dynamics of the environment to reduce the reliance on the accuracy of the forward model. When making decisions, human beings can not only predict future consequences of current behaviors forward but also imagine possible traces leading to a target goal backward. Based on this consideration, inspired by Edwards et al. (2018) and Goyal et al. (2018), we propose to learn a backward model in addition to the forward dynamics model to generate simulated rollouts backwards. To be specific, given a state and a previous action, the backward model predicts the preceding state, which means it can generate trajectories that terminate at one certain state.

Now, intuitively, if we utilize both the forward model and the backward model to sample bidirectionally from some encountering states for policy optimization, the model error will compound separately in different directions. As Figure 1 shows, when we generate trajectories with the same length using a forward model and bidirectional models re-




Figure 1. Comparison of compounding error of trajectories generated using only a forward model (top) and bidirectional models (bottom). States in real trajectories are shown in red, and states predicted by the forward (backward) model are shown in blue (green). The orange lines between real states and simulated states roughly represent the error, and when the rollout length increases, the error will compound. For bidirectional models, model error compounds k_1 steps backward and k_2 steps forward while for the forward model, model error compounds $k_1 + k_2$ steps.

spectively, the compounding error of the latter will be less than that of the former. More theoretical analysis is given in Section 5.

With this insight, we combine bidirectional models with recent MBPO method (Janner et al., 2019) and propose a practical MBRL algorithm called Bidirectional Model-based Policy Optimization (BMPO). Besides, we develop a novel state sampling strategy to sample states for model rollouts and incorporate model predictive control (MPC) (Camacho & Alba, 2013) into our algorithm for further performance improvement. Furthermore, we theoretically prove that BMPO has a tighter bound of return discrepancy compared against MBPO, which verifies the advantage of bidirectional models in theory. We evaluate our BMPO and previous state-of-the-art algorithms (Haarnoja et al., 2018; Janner et al., 2019) on a range of continuous control benchmark tasks. Experiments demonstrate that BMPO achieves higher sample efficiency and better asymptotic performance compared with prior model-based methods which only use forward models.

2. Related Work

Model-based reinforcement learning methods are expected to reduce sample complexity by learning a model as a simulator of the environment (Sutton & Barto, 2018). However,

model error tends to cripple the performance of model-based approaches, which is also known as model-bias (Deisenroth & Rasmussen, 2011). As is discovered in previous works (Venkatraman et al., 2015; Talvitie, 2017; Asadi et al., 2018b), even small model error can severely degrade multi-step rollouts since the model error will compound, and the predicted states will move out of the region where the model has high accuracy after a few steps (Asadi et al., 2018a).

To mitigate the compounding error problem, multi-step models (Whitney & Fergus, 2018; Asadi et al., 2019) were proposed to predict the outcome of executing a sequence of actions directly. Segments-based models have also been considered to make stable and accurate predictions over temporal segments (Mishra et al., 2017). Alternatively, a model may also be trained on its own outputs, hoping that the model can perform reasonable prediction in the region of its outputs (Talvitie, 2017; Kaiser et al., 2019; Freeman et al., 2019). Besides, model imitation tried to learn the model by matching the multi-step rollouts distributions via WGAN (Wu et al., 2019). Furthermore, adaptive rollout horizon techniques were investigated to stop rolling out based on compound model error estimates (Nguyen et al.; Xiao et al., 2019). The MBPO algorithm (Janner et al., 2019) avoided the compounding error by generating short branched rollouts from real states. This paper is mainly based on MBPO framework backbone and allows for extended rollouts with less compounding error by generating rollouts bidirectionally.

There are many model architecture choices, such as linear models (Parr et al., 2008; Sutton et al., 2012; Levine & Koltun, 2013; Levine & Abbeel, 2014; Kumar et al., 2016), nonparametric Gaussian processes (Kuss & Rasmussen, 2004; Ko et al., 2007; Deisenroth & Rasmussen, 2011), and neural networks (Draeger et al., 1995; Gal et al., 2016; Nagabandi et al., 2018). Model ensembles have shown to be effective in preventing a policy or a controller from exploiting the inaccuracies of any single model (Rajeswaran et al., 2016; Chua et al., 2018; Kurutach et al., 2018; Janner et al., 2019). In this paper, we adopt the ensemble of probabilistic networks for both forward and backward models.

Model predictive control (MPC) (Camacho & Alba, 2013) is considered as an efficient and robust way in model-based planning, which utilizes the dynamic model to look forward several steps and optimizes the actions sequence over a finite horizon. Nagabandi et al. (2018) combined deterministic neural networks model with a simple MPC of random shooting and acquired stable and plausible locomotion gaits in high-dimensional tasks. Chua et al. (2018) improved this method by capturing two kinds of uncertainty in modeling and replacing random shooting with the cross-entropy method (CEM). Wang & Ba (2019) further expanded MPC algorithms by applying policy networks to generate action

sequence proposals. In this paper, we also incorporate MPC into our algorithm to refine the actions sampled from policy when interacting with the environment.

Theoretical analysis for tabular- and linear-setting model-based methods has been conducted in prior works (Szita & Szepesvári, 2010; Dean et al., 2017). As for non-linear settings, Sun et al. (2018) provided a convergence analysis for the model-based optimal control framework. Luo et al. (2018) built a lower bound of the expected reward while enforced a trust region on the policy. Janner et al. (2019), instead, derived a bound of discrepancy between returns in real environment and those under the branched rollout scheme of MBPO in terms of rollout length, model error, and policy shift divergence. In this paper, we also provide a theoretical analysis of leveraging bidirectional models in MBPO and obtain a tighter bound.

To facilitate research in MBRL, Langlois et al. (2019) gathered a wide collection of MBRL algorithms and benchmarked them on a series of environments specially designed for MBRL. Beyond comparing different methods, they also raised several crucial challenges for future MBRL research.

This work is closely related to Goyal et al. (2018), which built a backtracking model of the environment and used it to generate traces leading to high value states. These traces are then used for imitation learning to improve the model-free learner. Our approach differs from theirs in the following two aspects: (1) their method is more like model-free RL as they mainly train the policy on real observation data and use the model generated traces to fine-tuning the policy, while our method is purely model-based RL and optimize the policy with model-generated data. (2) they only build a backward model of the environment, whereas we build a forward model and a backward model at the same time.

3. Preliminaries

3.1. Reinforcement Learning

A discrete-time Markov decision process (MDP) with infinite horizon is defined by the tuple $(\mathcal{S}, \mathcal{A}, M^*, r, \gamma, \rho_0)$. Here, $\mathcal{S} \in \mathbb{R}^{d_s}$ and $\mathcal{A} \in \mathbb{R}^{d_a}$ denote the state and action spaces, respectively. $\gamma \in [0, 1)$ is the discount factor. $M^*(s'|s, a)$ denotes the transition distribution function, and $r(s, a)$ denotes the reward function. ρ_0 represents the initial state distribution. Let η denote the expected return, *i.e.*, the sum of the discounted rewards along a trajectory $\tau := (s_0, a_0, \dots, s_t, a_t \dots)$. The goal of reinforcement learning is to find the optimal policy π^* that maximizes the expected return:

$$\pi^* = \arg \max_{\pi} \eta[\pi] = \arg \max_{\pi} \mathbb{E}_{\pi} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \right]. \quad (1)$$

where $s_0 \sim \rho_0$, $a_t \sim \pi(a_t|s_t)$, $s_{t+1} \sim M^*(s_{t+1}|s_t, a_t)$.

Algorithm 1: Bidirectional Model-based Policy Optimization (BMPO)

```

1 Initialize policy  $\pi_{\phi}$ , backward policy  $\tilde{\pi}_{\phi'}$ , forward
  model  $p_{\theta}$ , backward model  $q_{\theta'}$ , environment replay
  buffer  $\mathcal{D}_{\text{env}}$ , model replay buffer  $\mathcal{D}_{\text{model}}$ 
2 for  $N$  epochs do
3   Train model  $p_{\theta}$  and  $q_{\theta'}$  using  $\mathcal{D}_{\text{env}}$  via applying
  gradient descent on Equation 2 and 3
4   for  $E$  steps do
5     Take an action in environment according to
  MPC; add to  $\mathcal{D}_{\text{env}}$ 
6     for  $M$  model rollouts do
7       Sample state  $s_t$  from  $\mathcal{D}_{\text{env}}$  with
  probability computed via Equation 6
8       Perform  $k_1$  steps backward rollouts and
   $k_2$  steps forward rollouts starting from
   $s_t$ ; add to  $\mathcal{D}_{\text{model}}$ 
9     for  $G$  gradient updates do
10      Update policy parameters on model data:
   $\phi \leftarrow \phi - \lambda_{\pi} \hat{\nabla}_{\phi} J_{\pi}(\phi, \mathcal{D}_{\text{model}})$ 
11    for  $G'$  gradient updates do
12      Update backward policy on  $\mathcal{D}_{\text{env}}$  by
  optimizing Equation 4 or 5

```

3.2. Dyna-style Algorithm

Dyna-style algorithm (Sutton, 1991; Langlois et al., 2019) is one kind of MBRL method that uses a model for policy optimization. To be specific, in Dyna, a forward dynamic model $p_{\theta}(s'|s, a)$ is learned from agent interactions with the environment by executing the policy $\pi_{\phi}(a|s)$. The policy π_{ϕ} is then optimized using model-free algorithms on real data and model-generated data. To perform the Dyna-style algorithm using bidirectional models, we further learn a backward model $q_{\theta'}(s|s', a)$ and a backward policy $\tilde{\pi}_{\phi'}(a|s')$. As such, forward rollouts and backward rollouts can be generated simultaneously for policy optimization.

4. BMPO Framework

In this section, we introduce how BMPO leverages bidirectional models to generate more plentiful simulated data for policy optimization in detail. Although bidirectional models can be incorporated into almost any Dyna-style model-based algorithms (Sutton, 1991), we choose the Model-based Policy Optimization (MBPO) (Janner et al., 2019) algorithm as the framework backbone since it is the state-of-the-art MBRL method and is sufficiently general. The overall algorithm is demonstrated in Algorithm 1, and an overview of the algorithm architecture is shown in Figure 2.




Figure 2. Overview of the BMPO algorithm. When interacting with the environment, the agent uses the model to predict the future states over a finite horizon H and then take the first action of the sequence with the highest simulated estimated return. The transitions are stored in \mathcal{D}_{env} , where the state value is represented by the shade of color (value increases from light red to dark red). High value states are then sampled from \mathcal{D}_{env} to perform bidirectional model rollouts, which are stored in \mathcal{D}_{model} for policy optimization afterward.

4.1. Bidirectional Models Learning in BMPO

4.1.1. FORWARD MODEL

For the forward model, we use an ensemble of bootstrapped probabilistic dynamics models, which were first introduced in PETS (Chua et al., 2018) to capture two kinds of uncertainty. To be specific, individual probabilistic models can capture the aleatoric uncertainty aroused from the inherent stochasticity of a system, and the bootstrapped ensemble aims to capture the epistemic uncertainty due to the lack of sufficient training data. Prior works (Chua et al., 2018; Janner et al., 2019) have demonstrated that the ensemble of probabilistic models is quite effective in MBRL, even when the ground truth dynamics are deterministic.

In detail, each in the forward model ensembles $\{p_\theta^i\}_{i=1}^B$ is parameterized by a multi-layer neural network, and we denote the parameters in the forward model ensembles as θ . Given a state s and an action a , each probabilistic neural network outputs a Gaussian distribution with diagonal covariances of the next state: $p_\theta^i(s'|s, a) = \mathcal{N}(\mu_\theta^i(s, a), \Sigma_\theta^i(s, a))$.¹ We train the model ensembles with different initializations and bootstrapped samples of the real environment data via maximum likelihood, and the corresponding loss of the forward model is

$$\mathcal{L}_f(\theta) = \sum_{t=1}^N [\mu_\theta(s_t, a_t) - s_{t+1}]^\top \Sigma_\theta^{-1}(s_t, a_t) [\mu_\theta(s_t, a_t) - s_{t+1}] + \log \det \Sigma_\theta(s_t, a_t), \quad (2)$$

where μ and Σ are the mean and covariance respectively, and N denotes the total number of transition data.

4.1.2. BACKWARD MODEL

Besides the traditional forward model learning, we additionally learn a backward model for extended simulated data generation. In this way, we can mitigate the error caused by excessively exploiting the forward model. In other words,

¹The model network will output a similar Gaussian distribution of the reward as well, which is omitted here for simplicity.

the backward model is used to reduce the burden of the forward model on generating data.

Due to the powerful capabilities of the probabilistic network ensemble, similarly to the forward model, we adopt the same parameterization for the backward model. In detail, another multi-layer neural network with parameters θ' is used to output a Gaussian prediction, i.e., $q_{\theta'}^i(s|s', a) = \mathcal{N}(\mu_{\theta'}^i(s', a), \Sigma_{\theta'}^i(s', a))$, and the loss function of the backward model is

$$\mathcal{L}_b(\theta') = \sum_{t=1}^N [\mu_{\theta'}(s_{t+1}, a_t) - s_t]^\top \Sigma_{\theta'}^{-1}(s_{t+1}, a_t) [\mu_{\theta'}(s_{t+1}, a_t) - s_t] + \log \det \Sigma_{\theta'}(s_{t+1}, a_t). \quad (3)$$

We note that although we can directly train the backward model with data sampled from the environment, it remains a problem how to choose the actions when we use the backward model to generate trajectories backwards. Recall that in the forward model situation, the actions are usually taken by the current policy $\pi_\phi(a|s)$ or an exploration policy. Thus we need an additional backward policy $\tilde{\pi}_{\phi'}(a|s')$ to generate actions in the backward rollouts.

4.2. Policy Optimization in BMPO

4.2.1. BACKWARD POLICY

To sample trajectories backwards, we need to design a backward policy $\tilde{\pi}_{\phi'}(a|s')$ to take action given the next state. There are many alternatives for backward policy design. In our experiments, we use two heuristic methods for comparison.

The first way is to train the backward policy by maximum likelihood estimation according to the (a, s') data sampled in the environment. The corresponding loss function is

$$\mathcal{L}_{MLE}(\phi') = - \sum_{t=0}^N \log \tilde{\pi}_{\phi'}(a_t | s_{t+1}). \quad (4)$$

The second way is to use a conditional generative adversarial

network (CGAN) (Mirza & Osindero, 2014) to generate an action a conditioned on the next state s' , and the adversarial loss can be written as

$$\min_{\tilde{\pi}} \max_D V(D, \tilde{\pi}) = \mathbb{E}_{(a, s') \sim \pi} [\log D(a, s')] + \mathbb{E}_{s' \sim \pi} [\log (1 - D(\tilde{\pi}(\cdot | s'), s'))], \quad (5)$$

where D is an additional discriminator and the generator here is the backward policy $\tilde{\pi}$.

Our goal is to make the backward rollouts resemble the real trajectory sampled by the current forward policy. Thus when training the backward policy, we only use the recent trajectories sampled by the agent in the real environment.

4.2.2. MBPO WITH BIDIRECTIONAL MODELS

The original MBPO (Janner et al., 2019) algorithm iterates between three stages: data collection, model learning, and policy optimization. In the data collection stage, data is collected by executing the latest policy in real environment and added to replay buffer \mathcal{D}_{env} . In the model learning stage, an ensemble of bootstrapped forward models are trained using all the data in \mathcal{D}_{env} through maximum likelihood. Then, in the policy optimization stage, short-length rollouts starting from states randomly chosen from \mathcal{D}_{env} are generated using the forward model. Model-generated data is then used to train a policy π (see line 10 in Algorithm 1) through Soft Actor-Critic (SAC) (Haarnoja et al., 2018) by minimizing the expected KL-divergence $J_{\pi}(\phi, \mathcal{D}) = \mathbb{E}_{s_t \sim \mathcal{D}} [D_{KL}(\pi \| \exp\{Q^{\pi} - V^{\pi}\})]$.

Bidirectional models can be incorporated into MBPO naturally. Specifically, we iteratively collect data, train bidirectional models and backward policy, and then use them to generate short-length rollouts bidirectionally starting from some real states to optimize our policy. Besides, in Section 4.3 we design other components specifically for BMPO, *i.e.*, state sampling and MPC, which are crucial for the improvement of performance.

4.3. Design Decisions

4.3.1. STATE SAMPLING STRATEGY

To better exploit the bidirectional models, inspired by Goyal et al. (2018), we begin simulated branched rollouts bidirectionally from high value states instead of randomly chosen states in the environment replay buffer. In such a way, the agent could learn to reach high value states through backward rollouts, and also learn to act better after these states through forward rollouts.

An intuitive idea is simply choosing the first K highest-value states to begin rollouts, which, however, will cause the agent not knowing how to behave in low value states due to the absence of low value states data. For better

generalization and stability, we choose starting states from the environment replay buffer \mathcal{D}_{env} according to Boltzmann distribution based on the value $V(s)$ estimated by SAC. Let $p(s)$ be the probability of a state s being chosen as a starting state, then we have

$$p(s) \propto e^{\beta V(s)}, \quad (6)$$

where β is the hyperparameter to control the ratio of high value states.

4.3.2. INCORPORATING MPC

In BMPO, when taking actions in the real environment, we adopt model predictive control (MPC) to exploit the forward model further. MPC is a common model-based planning method using the learned model to look forward and optimize actions sequence. In detail, at each timestep, N candidate action sequences $\mathbf{A}_t^{1:N} = (a_t^{1:N}, \dots, a_{t+H-1}^{1:N})$ are generated, where H is the planning horizon, and the corresponding trajectories are simulated in the learned model. Then the first action of the sequence that yields the highest accumulated rewards is selected. Here, we use a variant of traditional MPC: (1) we generate action sequences from the current policy like Wang & Ba (2019) instead of uniform distribution; (2) we add the value estimate of the last state in the simulated trajectory to the planning objective, *i.e.*,

$$a_t = \operatorname{argmax}_{a_t^{1:N} \sim \pi} \sum_{t'=t}^{t+H-1} \gamma^{t'-t} r(s_{t'}, a_{t'}) + \gamma^H V(s_{t+H}). \quad (7)$$

It is worth noting that we only use MPC in the training phase but not in the evaluation phase. In the training phase, MPC helps select actions to take in real environments, and the obtained high value states will be used to generate simulated rollouts for policy optimization. In the evaluation phase, the trained policy is evaluated directly.

5. Theoretical Analysis

In this section, we theoretically analyze the discrepancy between the expected returns in the real environment and those under the bidirectional branched rollout scheme of BMPO. The proofs can be found in the appendix as provided in supplementary materials. When we use the bidirectional models to generate simulated rollouts from some encountered state, we assume that the length of backward rollouts is k_1 , and the length of forward rollouts is k_2 . Under such a scheme, we first consider a more general return discrepancy of two arbitrary bidirectional branched rollouts.

Lemma 5.1. (*Bidirectional Branched Rollout Returns Bound*). *Let η_1, η_2 be the expected returns of two bidirectional branched rollouts. Out of the branch, we assume that the expected total variation distance between*

these two dynamics at each timestep t is bounded as $\max_t E_{(s,a) \sim p_1^t(s,a)} D_{TV}(p_1^{\text{pre}}(s'|s,a) \| p_2^{\text{pre}}(s'|s,a)) \leq \epsilon_m^{\text{pre}}$, similarly, the forward branch dynamic bounded as $\max_t E_{(s,a) \sim p_1^t(s,a)} D_{TV}(p_1^{\text{for}}(s'|s,a) \| p_2^{\text{for}}(s'|s,a)) \leq \epsilon_m^{\text{for}}$, and the backward branch dynamic bounded as $\max_t E_{(s',a) \sim p_1^t(s',a)} D_{TV}(p_1^{\text{back}}(s|s',a) \| p_2^{\text{back}}(s|s',a)) \leq \epsilon_m^{\text{back}}$. Likewise, the total variation distance of policy is bounded by $\epsilon_\pi^{\text{pre}}$, $\epsilon_\pi^{\text{for}}$ and $\epsilon_\pi^{\text{back}}$, respectively. Then the returns are bounded as:

$$|\eta_1 - \eta_2| \leq 2r_{\max} \left[\frac{\gamma^{k_1+k_2+1}}{(1-\gamma)^2} (\epsilon_m^{\text{pre}} + \epsilon_\pi^{\text{pre}}) + \frac{\gamma^{k_1+k_2}}{1-\gamma} \epsilon_\pi^{\text{pre}} \right. \\ \left. + \frac{1-\gamma^{k_1}}{1-\gamma} (k_1 (\epsilon_m^{\text{back}} + \epsilon_\pi^{\text{back}}) + \epsilon_\pi^{\text{back}}) \right. \\ \left. + \frac{\gamma^{k_1}}{1-\gamma} (k_2 (\epsilon_m^{\text{for}} + \epsilon_\pi^{\text{for}}) + \epsilon_\pi^{\text{for}}) \right]. \quad (8)$$

Proof. See Appendix A, Lemma A.1. \square

Now, we can bound the discrepancy between the returns in the environment and in the branched rollouts of BMPO. Let $\eta[\pi]$ be the expected return of executing current policy in the true dynamics, and $\eta^{\text{branch}}[\pi]$ be the expected return of executing current policy in the model generated branch and executing old policy π_D out of the branch. Then we can derive the discrepancy between them as follows.

Theorem 5.1. (BMPO Return Discrepancy Upper Bound) Assume that the expected total variation distance between the learned forward model \hat{p} and the true dynamics p at each timestep t is bounded as $\max_t E_{(s,a) \sim \pi_t} [D_{TV}(p(s'|s,a) \| \hat{p}(s'|s,a))] \leq \epsilon_m^{\text{for}}$. Similarly, the error of backward model \hat{q} is bounded as $\max_t E_{(s',a) \sim \pi_t} [D_{TV}(q(s|s',a) \| \hat{q}(s|s',a))] \leq \epsilon_m^{\text{back}}$ and the variation between current policy and the behavioral policy is bounded as $\max_s D_{TV}(\pi_D(a|s) \| \pi(a|s)) \leq \epsilon_\pi$. Assume $\epsilon_m^{\text{for}} \approx \epsilon_m^{\text{back}} = \epsilon_m$, then under a branched rollouts scheme with a backward branch length of k_1 and a forward branch length of k_2 , we have

$$|\eta[\pi] - \eta^{\text{branch}}[\pi]| \leq 2r_{\max} \left[\frac{\gamma^{k_1+k_2+1} \epsilon_\pi}{(1-\gamma)^2} \right. \\ \left. + \frac{\gamma^{k_1+k_2} \epsilon_\pi}{(1-\gamma)} + \frac{\max(k_1, k_2) \epsilon_m}{1-\gamma} \right]. \quad (9)$$

Proof. See Appendix A, Theorem A.1. \square

We notice that in MBPO (Janner et al., 2019), the authors derived a similar return discrepancy bound (refer to Theorem 4.3 therein) with only one forward dynamics model. Setting the forward rollout length as $k_1 + k_2$, the bound is

$$|\eta[\pi] - \eta^{\text{branch}}[\pi]| \leq 2r_{\max} \left[\frac{\gamma^{k_1+k_2+1} \epsilon_\pi}{(1-\gamma)^2} \right. \\ \left. + \frac{\gamma^{k_1+k_2} \epsilon_\pi}{(1-\gamma)} + \frac{(k_1 + k_2) \epsilon_m}{1-\gamma} \right]. \quad (10)$$

By comparing the two bounds, it is evident that BMPO obtains a tighter upper bound of the return discrepancy by employing bidirectional models. The main difference is that in BMPO the coefficient of the model error is $\max(k_1, k_2)$, while in MBPO the coefficient is $k_1 + k_2$. It is easy to understand: as we generate the branched rollouts bidirectionally, the model error will compound only k_1 steps in the backward model or k_2 steps in the forward model and thus at most $\max(k_1, k_2)$ steps, instead of $k_1 + k_2$ steps using forward model only.

6. Experiments

Our experiments aim to answer the following three questions: 1) How does BMPO perform compared with model-free RL methods and previous state-of-the-art model-based RL methods using only one forward model? 2) Does using the bidirectional models reduce compounding error compared with using one forward model? 3) What are the critical components of our overall algorithm?

6.1. Comparison with State-of-the-Arts

In this section, we compare our method with previous state-of-the-art baselines. Specifically, for model-based methods, we compare against MBPO (Janner et al., 2019), as our method builds on top of it; SLBO (Luo et al., 2018) and PETS (Chua et al., 2018), both performing well in the model-based benchmarking test (Langlois et al., 2019). For model-free methods, we compare to Soft Actor-Critic (SAC) (Haarnoja et al., 2018), which is proved to be effective in continuous control tasks. For the MBPO baseline, we only generate k_2 -steps forward rollout, where k_2 is set as the default value used in the original MBPO paper. We do not report the result of MBPO using rollouts of $k_1 + k_2$ steps since using only k_2 steps achieves better performance. Notice that we do not include the backtracking model method (Goyal et al., 2018) for comparison since it is not a purely model-based method, and its performance improvement is limited compared with SAC.

We evaluate all the algorithms in six environments in total using OpenAI Gym (Brockman et al., 2016). Among them, Pendulum is one traditional control task, and Hopper, Walker2D, Ant are three complex MuJoCo tasks (Todorov et al., 2012). We additionally add two variants of MuJoCo tasks without early termination states, denoted as Hopper-NT and Walker2d-NT, which have been released as benchmarking environments for MBRL (Langlois et al., 2019). More details about the environments can be found in Appendix C.

The comparison results are shown in Figure 3. In different locomotion tasks, our method BMPO learns faster and has better asymptotic performance than previous model-based

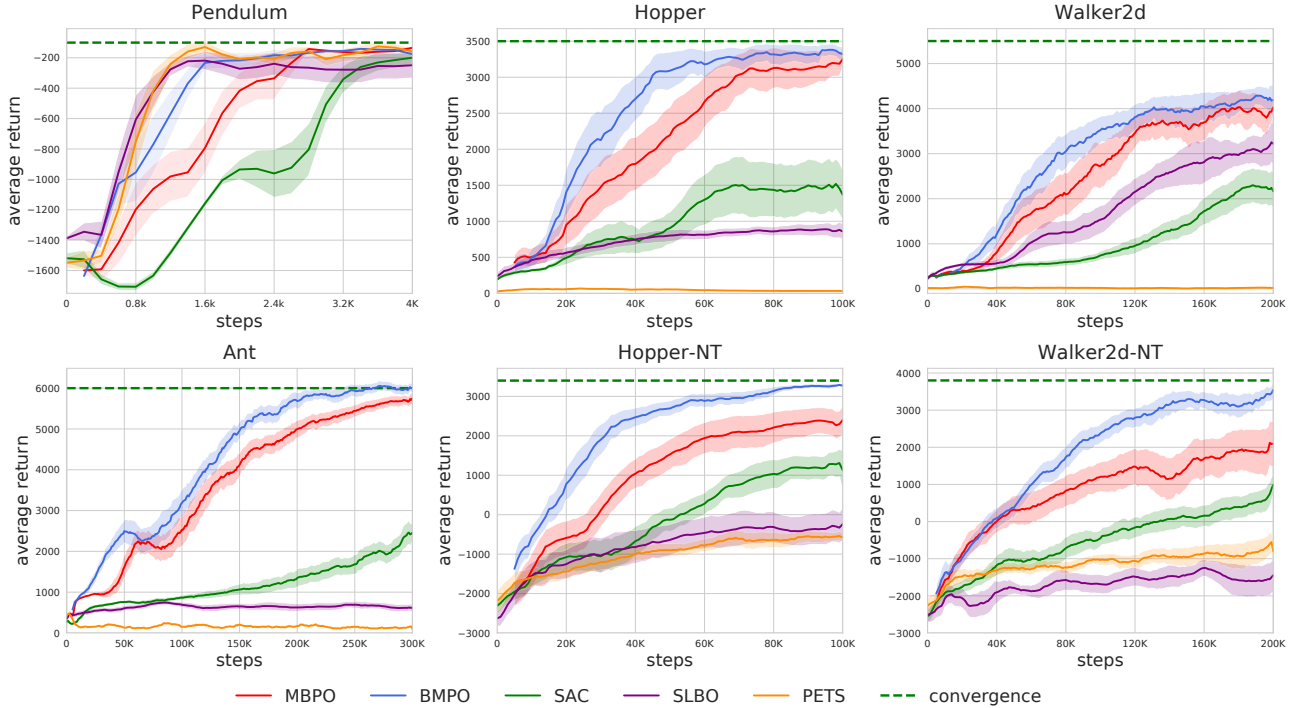


Figure 3. Learning curves of BMPO (ours) and four baselines on different continuous control environments. The solid lines indicate the mean and shaded areas indicate the standard error of six trails over different random seeds. Each trial is evaluated every 1000 environment steps (200 steps for Pendulum), where each evaluation reports the average return over ten episodes. The dashed reference lines are the asymptotic performance of SAC.

algorithms using only the forward model, which empirically demonstrates the advantage of bidirectional models. This gap is even more significant in benchmarking environments Hopper-NT and Walker2d-NT. One possible reason is that in the environments without early termination, the space of encountered states is larger, and thus the introduced state sampling strategy is more effective.

6.2. Model Error

In this section, we first plot the validation loss of the models in Figure 4a, which can roughly represent the single-step prediction error. As the figure shows, the difficulty of learning forward/backward models is task-dependent. Then, we investigate the compounding error of traditional forward model and our bidirectional models when generating the same length simulated trajectories. We calculate the multi-step prediction error for evaluation. A similar validation error is also used in Nagabandi et al. (2018). More specifically, assume a real trajectory of length $2h$ is denoted as $(s_0, a_0, s_1, \dots, s_{2h})$. For the forward model, we sample from s_0 and generate forward rollouts $(\hat{s}_0, a_0, \hat{s}_1, \dots, \hat{s}_{2h})$ where $\hat{s}_0 = s_0$ and for $i \geq 0$, $\hat{s}_{i+1} = p_\theta(\hat{s}_i, a_i)$. Then the corresponding compounding error is defined as

$$\text{Error}_{\text{for}} = \frac{1}{2h} \sum_{i=1}^{2h} \|\hat{s}_i - s_i\|_2^2.$$

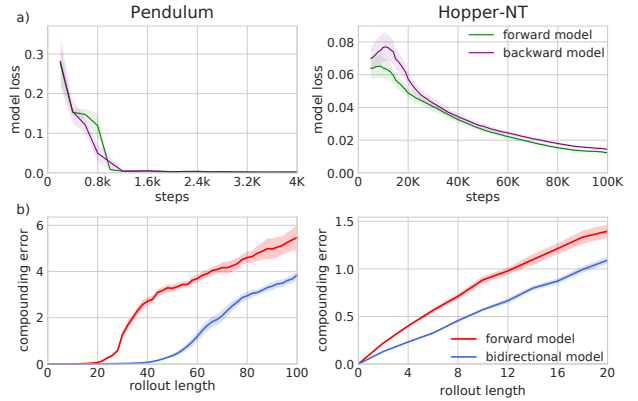


Figure 4. (a) Validation loss of forward/backward models. (b) Compounding error of the forward model and the bidirectional models in environment Pendulum (left) and Hopper-NT (right).

Similarly, for the bidirectional models, suppose we sample from s_h and generate both forward and backward rollouts of h steps, compounding error is defined as

$$\text{Error}_{\text{bi}} = \frac{1}{2h} \sum_{i=1}^h \left(\|\hat{s}_{h+i} - s_{h+i}\|_2^2 + \|\hat{s}_{h-i} - s_{h-i}\|_2^2 \right),$$

where $\hat{s}_h = s_h$ and for $i \geq 0$, $\hat{s}_{h+i+1} = p_\theta(\hat{s}_{h+i}, a_{h+i})$ and $\hat{s}_{h-i-1} = q_{\theta'}(\hat{s}_{h-i}, a_{h-i-1})$.

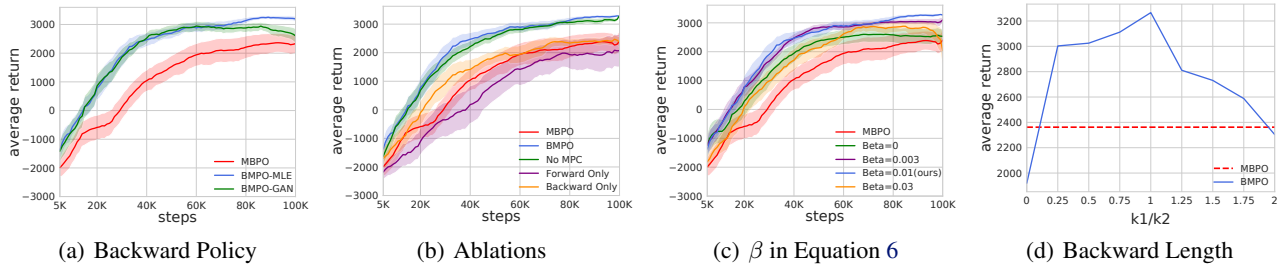


Figure 5. Design evaluation of our algorithm BMPO on task Hopper-NT. (a) Comparison of two heuristic design choices for the backward policy loss: MLE loss and GAN loss. (b) Ablation study of three crucial components: forward model, backward model, and MPC. (c) We study the sensitivity of our algorithm to the hyperparameter β in Equation 6. (d) Average return of the last 10 epochs over six trials with different backward rollout lengths k_1 and fixed forward length k_2 .

We conduct experiments with different rollout length $2h$ and plot the results in Figure 4b. We observe that employing bidirectional models can significantly reduce model compounding error, which is consistent with our intuition and theoretical analysis.

6.3. Design Evaluation

In this section, we evaluate the importance of each design decision for our overall algorithm. The results are demonstrated in Figure 5.

6.3.1. BACKWARD POLICY DESIGN

In Figure 5(a), we compare two heuristic design choices for the backward policy loss: MLE loss and GAN loss, which are described in detail in Section 4.2.1. From the comparison, we notice a slight performance degradation in the late iterations of training when adopting GAN loss. One possible reason may be the instability and model collapse of GAN (Goodfellow, 2016). More advanced methods to stabilize GAN’s training can be incorporated into backward policy training, which remains as future work. As for other comparisons in this paper, we use the MLE loss for the backward policy as default.

6.3.2. ABLATION STUDY

We further carry out an ablation study to characterize the importance of three main components of our algorithm: 1) no backward model to sample trajectory (Forward Only); 2) no forward model to sample trajectory (Backward Only); 3) no MPC when interacting with the environment (No MPC). The results are shown in Figure 5(b). We find that ablating MPC decreases the performance slightly, and using only one forward model or backward model causes a more severe performance dropping. This further reveals the superiority of using bidirectional models. Notice that the performance of not using backward model is even worse than the vanilla MBPO, which means that sampling from high value states

does not benefit the traditional forward model since the agent may only focus on how to act at high value states, but not care how to reach these states.

6.3.3. HYPERPARAMETER STUDY

In this section, we investigate the sensitivity of BMPO to the hyperparameters. First, we test BMPO with different hyperparameter β used in Equation 6. We vary β from 0 to a relatively large number, where $\beta = 0$ means random sampling, and larger β means focusing more on high value states. The results are shown in Figure 5(c). We observe that up to a certain level, increasing β yields better performance while too large β can degrade the performance. This may be due to the fact that with a too large β , low value states are almost impossible to be chosen, and the value estimation at the beginning is difficult. Nevertheless, our algorithm with different β all outperform the baseline, which indicates the robustness to β .

Though it has been discovered that linearly increasing rollout length achieves excellent performance (Janner et al., 2019), it remains a problem that how to choose the backward rollout length k_1 according to the forward length k_2 . We fix k_2 to be the same as Janner et al. (2019) and vary k_1 from 0 to $2k_2$. As is shown in Figure 5(d), setting $k_1 = k_2$ provides the best result, while too short and too long backward length k_1 are both detrimental. In practice, we use $k_1 = k_2$ in most environments except Ant, where the backward model error is too large compared with forward model error. All hyperparameters settings are provided in Appendix D.

7. Conclusion

In this work, we present a novel model-based reinforcement learning method, namely bidirectional model-based policy optimization (BMPO), using the newly introduced bidirectional models. We theoretically prove the advantage of bidirectional models by deriving a tighter return discrepancy upper bound of RL objective compared with only one

forward model. Experimental results show that BMPO achieves better asymptotic performance and higher sample efficiency than previous state-of-the-art model-based methods on several benchmark continuous control tasks. For future work, we will investigate the usage of bidirectional models in other model-based RL frameworks and study how to leverage the bidirectional models better.

Acknowledgments

The corresponding author Weinan Zhang thanks the support of "New Generation of AI 2030" Major Project 2018AAA0100900 and NSFC (61702327, 61772333, 61632017).

References

- Asadi, K., Cater, E., Misra, D., and Littman, M. L. Towards a simple approach to multi-step model-based reinforcement learning. *arXiv preprint arXiv:1811.00128*, 2018a.
- Asadi, K., Misra, D., and Littman, M. L. Lipschitz continuity in model-based reinforcement learning. *arXiv preprint arXiv:1804.07193*, 2018b.
- Asadi, K., Misra, D., Kim, S., and Littman, M. L. Combating the compounding-error problem with a multi-step model. *arXiv preprint arXiv:1905.13320*, 2019.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- Camacho, E. F. and Alba, C. B. *Model predictive control*. Springer Science & Business Media, 2013.
- Chua, K., Calandra, R., McAllister, R., and Levine, S. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems*, pp. 4754–4765, 2018.
- Dean, S., Mania, H., Matni, N., Recht, B., and Tu, S. On the sample complexity of the linear quadratic regulator. *arXiv preprint arXiv:1710.01688*, 2017.
- Deisenroth, M. and Rasmussen, C. E. Pilco: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning (ICML-11)*, pp. 465–472, 2011.
- Deisenroth, M. P., Neumann, G., Peters, J., et al. A survey on policy search for robotics. *Foundations and Trends® in Robotics*, 2(1–2):1–142, 2013.
- Draeger, A., Engell, S., and Ranke, H. Model predictive control using neural networks. *IEEE Control Systems Magazine*, 15(5):61–66, 1995.
- Edwards, A. D., Downs, L., and Davidson, J. C. Forward-backward reinforcement learning. *arXiv preprint arXiv:1803.10227*, 2018.
- Freeman, C. D., Metz, L., and Ha, D. Learning to predict without looking ahead: World models without forward prediction. *ArXiv*, abs/1910.13038, 2019.
- Gal, Y., McAllister, R., and Rasmussen, C. E. Improving pilco with bayesian neural network dynamics models. In *Data-Efficient Machine Learning workshop, ICML*, volume 4, 2016.
- Goodfellow, I. Nips 2016 tutorial: Generative adversarial networks. *arXiv preprint arXiv:1701.00160*, 2016.
- Goyal, A., Brakel, P., Fedus, W., Singhal, S., Lillicrap, T., Levine, S., Larochelle, H., and Bengio, Y. Recall traces: Backtracking models for efficient reinforcement learning. *arXiv preprint arXiv:1804.00379*, 2018.
- Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018.
- Janner, M., Fu, J., Zhang, M., and Levine, S. When to trust your model: Model-based policy optimization. *arXiv preprint arXiv:1906.08253*, 2019.
- Kaiser, L., Babaeizadeh, M., Milos, P., Osinski, B., Campbell, R. H., Czechowski, K., Erhan, D., Finn, C., Koza-kowski, P., Levine, S., et al. Model-based reinforcement learning for atari. *arXiv preprint arXiv:1903.00374*, 2019.
- Ko, J., Klein, D. J., Fox, D., and Haehnel, D. Gaussian processes and reinforcement learning for identification and control of an autonomous blimp. In *Proceedings 2007 IEEE international conference on robotics and automation*, pp. 742–747. IEEE, 2007.
- Kumar, V., Todorov, E., and Levine, S. Optimal control with learned local models: Application to dexterous manipulation. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 378–383. IEEE, 2016.
- Kurutach, T., Clavera, I., Duan, Y., Tamar, A., and Abbeel, P. Model-ensemble trust-region policy optimization. *arXiv preprint arXiv:1802.10592*, 2018.
- Kuss, M. and Rasmussen, C. E. Gaussian processes in reinforcement learning. In *Advances in neural information processing systems*, pp. 751–758, 2004.
- Langlois, E., Zhang, S., Zhang, G., Abbeel, P., and Ba, J. Benchmarking model-based reinforcement learning. *arXiv preprint arXiv:1907.02057*, 2019.

- Levine, S. and Abbeel, P. Learning neural network policies with guided policy search under unknown dynamics. In *Advances in Neural Information Processing Systems*, pp. 1071–1079, 2014.
- Levine, S. and Koltun, V. Guided policy search. In *International Conference on Machine Learning*, pp. 1–9, 2013.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- Luo, Y., Xu, H., Li, Y., Tian, Y., Darrell, T., and Ma, T. Algorithmic framework for model-based deep reinforcement learning with theoretical guarantees. *arXiv preprint arXiv:1807.03858*, 2018.
- Mirza, M. and Osindero, S. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- Mishra, N., Abbeel, P., and Mordatch, I. Prediction and control with temporal segment models. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 2459–2468. JMLR. org, 2017.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529, 2015.
- Nagabandi, A., Kahn, G., Fearing, R. S., and Levine, S. Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 7559–7566. IEEE, 2018.
- Nguyen, N. M., Singh, A., and Tran, K. Improving model-based rl with adaptive rollout using uncertainty estimation.
- Parr, R., Li, L., Taylor, G., Painter-Wakefield, C., and Littman, M. L. An analysis of linear models, linear value-function approximation, and feature selection for reinforcement learning. In *Proceedings of the 25th international conference on Machine learning*, pp. 752–759. ACM, 2008.
- Rajeswaran, A., Ghotra, S., Ravindran, B., and Levine, S. Epopt: Learning robust neural network policies using model ensembles. *arXiv preprint arXiv:1610.01283*, 2016.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Sun, W., Gordon, G. J., Boots, B., and Bagnell, J. Dual policy iteration. In *Advances in Neural Information Processing Systems*, pp. 7059–7069, 2018.
- Sutton, R. S. Dyna, an integrated architecture for learning, planning, and reacting. *ACM Sigart Bulletin*, 2(4):160–163, 1991.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 2018.
- Sutton, R. S., Szepesvári, C., Geramifard, A., and Bowling, M. P. Dyna-style planning with linear function approximation and prioritized sweeping. *arXiv preprint arXiv:1206.3285*, 2012.
- Szita, I. and Szepesvári, C. Model-based reinforcement learning with nearly tight exploration complexity bounds. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pp. 1031–1038, 2010.
- Talvitie, E. Self-correcting models for model-based reinforcement learning. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- Todorov, E., Erez, T., and Tassa, Y. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033. IEEE, 2012.
- Venkatraman, A., Hebert, M., and Bagnell, J. A. Improving multi-step prediction of learned time series models. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- Wang, T. and Ba, J. Exploring model-based planning with policy networks. *arXiv preprint arXiv:1906.08649*, 2019.
- Whitney, W. and Fergus, R. Understanding the asymptotic performance of model-based rl methods. 2018.
- Wu, Y.-H., Fan, T.-H., Ramadge, P. J., and Su, H. Model imitation for model-based reinforcement learning. *arXiv preprint arXiv:1909.11821*, 2019.
- Xiao, C., Wu, Y., Ma, C., Schuurmans, D., and Müller, M. Learning to combat compounding-error in model-based reinforcement learning. *arXiv preprint arXiv:1912.11206*, 2019.