

Sample-Efficient Reinforcement Learning via Conservative Model-Based Actor-Critic

Zhihai Wang¹, Jie Wang^{1,2*}, Qi Zhou¹, Bin Li¹, Houqiang Li^{1,2}

¹CAS Key Laboratory of Technology in GIPAS, University of Science and Technology of China

²Institute of Artificial Intelligence, Hefei Comprehensive National Science Center

{zhwangx, zhouqida}@mail.ustc.edu.cn

{jjewangx, binli, lihq}@ustc.edu.cn

Abstract

Model-based reinforcement learning algorithms, which aim to learn a model of the environment to make decisions, are more sample efficient than their model-free counterparts. The sample efficiency of model-based approaches relies on whether the model can well approximate the environment. However, learning an accurate model is challenging, especially in complex and noisy environments. To tackle this problem, we propose the conservative model-based actor-critic (CMBAC), a novel approach that achieves high sample efficiency without the strong reliance on accurate learned models. Specifically, CMBAC learns multiple estimates of the Q-value function from a set of inaccurate models and uses the average of the bottom-k estimates—a conservative estimate—to optimize the policy. An appealing feature of CMBAC is that the conservative estimates effectively encourage the agent to avoid unreliable “promising actions”—whose values are high in only a small fraction of the models. Experiments demonstrate that CMBAC significantly outperforms state-of-the-art approaches in terms of sample efficiency on several challenging tasks, and the proposed method is more robust than previous methods in noisy environments.

Introduction

Reinforcement learning has achieved great success in decision-making tasks, ranging from playing video games (Mnih et al. 2015; Hessel et al. 2018) to controlling robots in simulators (Lillicrap et al. 2016; Haarnoja et al. 2018). However, many of these results are achieved by model-free algorithms and generally require a massive number of samples, which significantly hinders the applications of model-free methods in real-world tasks (Kurutach et al. 2018; Janner et al. 2019). In contrast, model-based approaches, which build a model of the environment and generate fictitious interactions, are more sample-efficient than their model-free counterparts. Therefore, model-based approaches are promising candidates for dealing with real-world tasks.

The sample efficiency of model-based approaches crucially relies on learning accurate models efficiently, as model errors limit the performance of learned policies, known as the model-bias problem (Deisenroth and Rasmussen 2011; Clavera et al. 2018). Specifically, model errors

can mislead the agent into selecting the unreliable actions—whose values are high in the model but is low in the true environment with a large probability—and thus degrade the performance of the learned policy until learning an accurate model from a large number of interactions (Kurutach et al. 2018; Clavera et al. 2018). Previous methods tackle this problem by improving the expressive power of the models, such as using neural networks and ensemble models (Punjani and Abbeel 2015; Nagabandi et al. 2018; Kurutach et al. 2018; Chua et al. 2018). However, efficiently learning accurate models remains challenging, especially in complex and noisy environments (Janner et al. 2019; Wang et al. 2019; Pan et al. 2020), which creates a hindrance in further improving the sample efficiency of model-based approaches.

To alleviate the strong reliance on accurate models, recent model-based algorithms (Luo et al. 2019; Zhou, Li, and Wang 2020; Yu et al. 2020; Pan et al. 2020) prevent the agent from exploiting model errors by an uncertainty-based penalty. Specifically, they explicitly quantify the uncertainty of the Q-value via the discrepancy of ensemble models and use it as a penalty to learn a lower bound of the true Q-value. However, many works have shown that the uncertainty quantification of existing methods can be unreliable, which acts as a bottleneck to their sample efficiency (Ovadia et al. 2019; Yu et al. 2021). We also empirically show that many existing uncertainty quantification methods can not well approximate the errors of Q-function in Section .

In this paper, we propose the conservative model-based actor-critic (CMBAC), a novel approach that approximates a posterior distribution over Q-values based on the ensemble models and uses the average of the left tail of the distribution approximation to optimize the policy. Specifically, CMBAC alternates between learning multiple estimates of the Q-value from the ensemble models and uses a conservative estimate, i.e., the average of the bottom-k estimates, to optimize the policy. CMBAC has two main advantages: (1) the conservative estimates effectively encourage the agent to avoid unreliable “promising actions”—whose values are high in only a small fraction of the models; (2) the distribution approximation over Q-values can produce reasonable uncertainty estimates of the Q-value. We empirically show that the uncertainty quantification of CMBAC approximates the errors of Q-function more accurately than previous uncertainty quantification methods, which plays a crucial role

*Corresponding author.

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

in the impressive performance of CMBAC (please refer to Figure 3). Experiments show that CMBAC significantly outperforms state-of-the-art methods in terms of sample efficiency on several challenging control tasks (Brockman et al. 2016; Todorov, Erez, and Tassa 2012). Moreover, experiments demonstrate that CMBAC is more robust to model imperfections than previous methods in noisy environments.

Related Work

In this section, we discuss related work, including model-based reinforcement learning, uncertainty in reinforcement learning, and conservatism in reinforcement learning.

Model-Based Reinforcement Learning

Roughly speaking, model-based approaches fall into three categories according to the way of model usage: (1) dyna-style methods (Sutton 1990; Luo et al. 2019; Zhou, Li, and Wang 2020; Janner et al. 2019), which use the model to generate imaginary samples as additional training data; (2) shooting algorithms (de Boer et al. 2005; Chua et al. 2018; Wang and Ba 2020), which use the model to plan to seek the optimal action sequence; (3) policy search with backpropagation (Nguyen and Widrow 1990; Fairbank and Alonso 2012; Heess et al. 2015; Clavera, Fu, and Abbeel 2020; Amos et al. 2021) through time, which exploits the model derivatives and computes the analytic policy gradient. Our work falls into the first category, i.e., the dyna-style algorithm, which has recently shown the potential to achieve high sample efficiency (Janner et al. 2019).

Uncertainty in Reinforcement Learning

Uncertainty estimation plays a crucial role in many reinforcement learning methods (Strehl and Littman 2008; Osband et al. 2016; O’Donoghue et al. 2018; Zhou, Li, and Wang 2020; Yu et al. 2020). In model-based reinforcement learning, recent methods (Luo et al. 2019; Yu et al. 2020; Pan et al. 2020) mainly focus on explicitly quantifying the uncertainty of the Q-value via the discrepancy of the ensemble models to alleviate the strong reliance on accurate models. However, their uncertainty quantification can be unreliable for deep network models (Ovadia et al. 2019; Yu et al. 2021). In contrast, our work approximates a posterior distribution over Q-values based on the ensemble models to capture the uncertainty of the value function. Although the work (Vuong and Tran 2019) also approximates a posterior distribution over Q-values based on the ensemble models, our work differs from it in the way of value estimation. Vuong and Tran (2019) estimates the value in a model via Monte Carlo methods (Sutton and Barto 2018), which can suffer from compounding model errors. In contrast, we use temporal difference methods to estimate the value (Sutton and Barto 2018). In model-free reinforcement learning, many methods (Strehl and Littman 2008; Osband et al. 2016; O’Donoghue et al. 2018) quantify the uncertainty via the true environment samples and use it as a reward bonus to promote exploration, while these uncertainty quantification methods are inappropriate for model-based methods (Zhou, Li, and Wang 2020).

Conservatism in Reinforcement Learning

Previous methods introduce the conservatism into policy optimization by underestimating the true value to improve the robustness (Doyle 1996; Lim, Xu, and Mannor 2013; Rajeswaran et al. 2017) and alleviate the overestimation bias (Fujimoto, van Hoof, and Meger 2018; Kuznetsov et al. 2020; Kumar et al. 2020). In model-based reinforcement learning, some methods (Doyle 1996; Lim, Xu, and Mannor 2013; Rajeswaran et al. 2017) leverage robust policy optimization, which learns a policy that performs well across models. However, the learned policies tend to be over-conservative (Clavera et al. 2018). In contrast, we use the average of the bottom-k estimates instead of the minimum to optimize the policy, controlling the degree of conservatism. In model-free reinforcement learning, many methods incorporate conservatism into policy learning to alleviate the overestimation bias that comes from function approximation errors (Fujimoto, van Hoof, and Meger 2018; Kuznetsov et al. 2020; Kumar et al. 2020). In contrast, CMBAC leverages conservatism to alleviate the overestimation that comes from model errors (please refer to Section).

Background

In this section, we present the notation and provide a brief introduction to the state-of-the-art model-based algorithm, i.e., Model-Based Policy Optimization (Janner et al. 2019).

Preliminaries

We here introduce notation which we will use throughout the paper. We consider an infinite-horizon Markov decision process (MDP) denoted by a tuple $(\mathcal{S}, \mathcal{A}, P^*, r, \gamma, \rho_0)$, where \mathcal{S} and \mathcal{A} are the sets of states and actions, respectively, $P^* : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, \infty)$ is the transition probability density function with $P^*(\cdot | s, a)$ representing the conditional distribution of the next state given the current state s and action a , $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function, $\rho_0 : \mathcal{S} \rightarrow [0, \infty)$ is the starting state distribution, and γ is the discount factor. Let $\pi : \mathcal{S} \rightarrow \mathcal{P}(\mathcal{A})$ be a stationary policy, where $\mathcal{P}(\mathcal{A})$ is a set of probability distribution over \mathcal{A} . Let $\pi(\cdot | s)$ denote the probability distribution over \mathcal{A} at state s . In model-based reinforcement learning, we learn a dynamics model $\hat{P}(\cdot | s, a)$ using data collected from interaction with the true MDP. For simplicity, we assume that the reward function $r(s, a)$ is known throughout the paper, but in practice, we learn a reward function. Let S_0 be the random variable for the initial state. Let $Q^{\pi, P}$ be the state-action value function on the model P and policy π defined by:

$$Q^{\pi, P}(s, a) = \mathbb{E}_{\pi, P} \left[\sum_{t=0}^{\infty} \gamma^t r(S_t, A_t) | S_0 = s, A_0 = a \right].$$

We define $\eta(\pi, P) = \mathbb{E}_{\pi} [Q^{\pi, P}(S_0, A_0)]$ as the expected reward-to-go. Our goal is to maximize the reward-to-go on the true model, that is, $\eta(\pi, P^*)$, over the policy π .

Model-Based Policy Optimization

Model-based policy optimization (MBPO) is a state-of-the-art model-based algorithm that has achieved impressive performance (Janner et al. 2019). MBPO has three ingredients:

Algorithm 1: Pseudo code for CMBAC

```

1: Initialize an ensemble of models  $\{P_{\psi_i}\}_{i=1}^N$ , environ-
   environment dataset  $\mathcal{D}_{\text{env}}$ , and model dataset  $\mathcal{D}_{\text{model}}$ 
2: Initialize policy  $\pi_\phi$ , multi-head Q-network  $\{Q_{\theta_j}\}_{j=1}^K$ 
3: for  $N$  epochs do
4:   Train models  $\{P_{\psi_i}\}_{i=1}^N$  on  $\mathcal{D}_{\text{env}}$ 
5:   for  $E$  steps do
6:     Take action in environment using  $\pi_\phi$ ; add to  $\mathcal{D}_{\text{env}}$ 
7:     for  $M$  model rollouts do
8:       Sample  $s_t$  uniformly from  $\mathcal{D}_{\text{env}}$ 
9:       Perform  $k$ -step model rollouts starting from  $s_t$ 
         using policy  $\pi_\phi$ ; add to  $\mathcal{D}_{\text{model}}$ 
10:    end for
11:    for  $G$  gradient updates do
12:      Train value functions on model data:  $\theta_j \leftarrow \theta_j -$ 
         $\lambda_Q \hat{\nabla} J_Q^j(\theta_j)$  for  $j \in \{1, \dots, K\}$ 
13:      Conservative policy optimization on model data:
         $\phi \leftarrow \phi - \lambda_\pi \hat{\nabla} J_\pi(\phi)$ 
14:    end for
15:  end for
16: end for

```

1. **Ensemble models** MBPO trains a bootstrap ensemble of dynamics models via maximum likelihood technique (Chua et al. 2018) on dataset \mathcal{D}_{env} collected from interactions with the true environment. Each member of the set is modeled as a Gaussian with mean and diagonal covariances given by neural networks. Our work also learns an ensemble of probabilistic models.
2. **Model usage** MBPO selects a model uniformly at random from the ensemble and generates a prediction from the selected model. To reduce compounding model errors introduced by long rollouts, MBPO generates many short rollouts as additional training dataset $\mathcal{D}_{\text{model}}$.
3. **Policy optimization** MBPO uses soft actor-critic (SAC) (Haarnoja et al. 2018)—a state-of-the-art model-free algorithm based on the maximum entropy reinforcement learning framework—to optimize the policy.

Conservative Model-Based Actor-Critic

In this section, we present a detailed description of CMBAC. CMBAC alternates between (1) learning multiple estimates of the Q-value function from the ensemble models and (2) using the average of the bottom-k estimates to optimize the policy. We provide an illustration of CMBAC in Figure 1 and summarize the procedure of CMBAC in Algorithm 1.

Capturing the Uncertainty of the Q-value

To capture the uncertainty of the Q-value, CMBAC directly approximates the posterior distribution over Q-values based on the posterior distribution approximation over models. The uncertainty of the Q-value comes from the unknown true environment, and model-based approaches usually estimate it using a possible set of models.

To approximate the posterior distribution over models, CMBAC learns an ensemble of probabilistic neural

networks $\{p_{\psi_i}\}_{i=1}^N$, which has been shown the potential to capture the uncertainty of models (Lakshminarayanan, Pritzel, and Blundell 2017; Chua et al. 2018). Each probabilistic neural network model the transition probability density as a Gaussian with mean and diagonal covariances given by neural networks. That is, $p_{\psi_i}(s'|s, a) = \mathcal{N}(\mu_{\psi_i}(s, a), \sigma_{\psi_i}(s, a))$. To control the granularity of discrepancy between models, CMBAC constructs a set of models \mathcal{M} using these probabilistic neural networks. Each element in \mathcal{M} , denoted by \mathcal{M}_j , is a set consists of M different networks ($M < N$) and thus the size of \mathcal{M} is $K = \binom{N}{M}$. We view each \mathcal{M}_j as a model, and \mathcal{M}_j generates next state s' given current state-action pair (s, a) under the distribution

$$P_j(s'|s, a) = \sum_{p \in (\{p_{\psi_i}\}_{i=1}^N \cap \mathcal{M}_j)} \frac{1}{M} p(s'|s, a).$$

The set of models \mathcal{M} has the ability to well approximate the posterior distribution of the true environment using non-parametric bootstrap with random initialization (Efron and Tibshirani 1993; Osband et al. 2016; Chua et al. 2018). Given fixed M , CMBAC achieves more fine-grained granularity of discrepancy between models in \mathcal{M} by increasing the total number of networks N . Given fixed N , the discrepancy between models drops with increasing M , i.e., the number of networks in each \mathcal{M}_j . If $M = N$, then CMBAC reduces to MBPO. Moreover, representing each model by an ensemble of neural networks significantly improves the model accuracy (Kurutach et al. 2018; Chua et al. 2018).

To approximate the posterior distribution over Q-values, CMBAC learns multiple estimates via a multi-head Q-network from the distribution approximation over models \mathcal{M} . Similar to bootstrapped DQN (Osband et al. 2016), the multi-head Q-network is a shared neural network architecture with K "heads" branching off independently (please refer to Appendix D for details). Each "head" \hat{Q}_{θ_j} provides an estimate of the Q-value for a policy π , which corresponds to a model $\mathcal{M}_j \in \mathcal{M}$. That is, CMBAC aims to approximate the Q^{π, P_j} via the "head" \hat{Q}_{θ_j} .

The target value of each "head" $\hat{Q}_{\theta_j}(s, a)$ is given by

$$y_j(s, a) = r(s, a) + \gamma(\hat{Q}_{\bar{\theta}_j}(s'_j, a'_j) - \alpha \log(\pi(a'_j|s'_j))),$$

where $j = 1, \dots, K$, s'_j is sampled from $P_j(\cdot|s, a)$, a'_j is sampled from $\pi(\cdot|s'_j)$, α is the temperature parameter, and $\hat{Q}_{\bar{\theta}_j}$ is the target value network with $\bar{\theta}_j$ being an exponentially moving average of the value network weights, which has been shown to stabilize training (Mnih et al. 2015). For each $j \in \{1, \dots, k\}$, the parameter θ_j can be trained to minimize the Bellman residual

$$J_Q(\theta_j) = \mathbb{E}_{(s,a) \sim \mathcal{D}_{\text{model}}} \left[\frac{1}{2} (\hat{Q}_{\theta_j}(s, a) - y_j(s, a))^2 \right].$$

CMBAC naturally approximates the posterior distribution over Q-values, as it learns an estimate from a model sampled from the approximated posterior distribution over models, respectively. We use the clipped double Q-learning as proposed by (Fujimoto, van Hoof, and Meger 2018). In practice, we use two multi-head Q-network and train each "head" using the minimum of the corresponding two "heads".

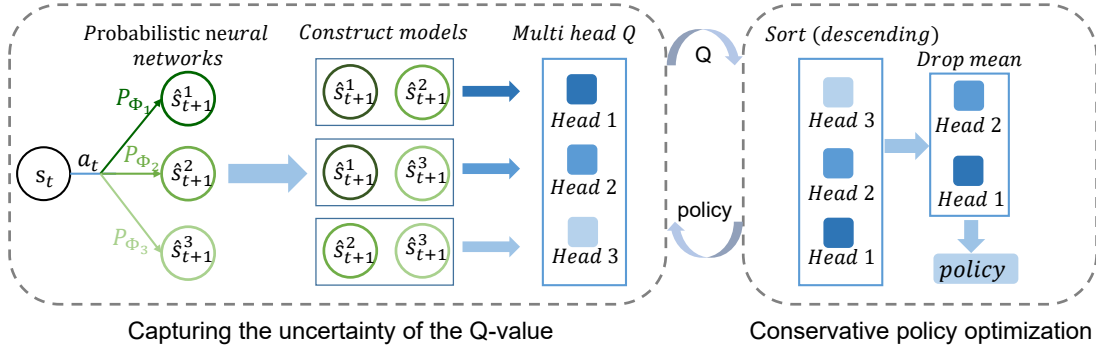


Figure 1: Illustration of CMBAC with $N = 3, M = 2$, and $L = 1$. It first learns three probabilistic neural networks and constructs each model using the arbitrarily two probabilistic neural networks. It then alternates between learning three estimates of the Q-value function from the ensemble models and using the average of the bottom-2 estimates to optimize the policy.

Conservative Policy Optimization

To prevent the agent from exploiting model errors, CMBAC uses a conservative estimate of the Q-value function to optimize the policy, named conservative policy optimization. Previous model-based methods aim to learn a conservative estimate that is an approximate lower bound of the true Q-value Q^{π, P^*} . For example, robust policy optimization (Doyle 1996; Lim, Xu, and Mannor 2013) considers a set of possible models and learns the worst-case Q-value, i.e., $Q^\pi(s, a) = \min_{P \in \mathcal{M}} Q^{\pi, P}(s, a)$. However, these methods tend to learn an over-conservative policy, which severely degrades their sample efficiency and asymptotic performance (Clavera et al. 2018). Unlike these methods, CMBAC introduces conservatism to alleviate the overestimation that comes from the model errors and does not aim to learn a lower bound. We observe that if we learn multiple estimates using the approach in Section , a small fraction of the “heads” severely overestimate the Q-value (see details in Section) while the others provide a relatively reasonable estimation. Based on the observation, we hypothesize that the actions—whose values are high in only a small fraction of the models—are unreliable in the true environment. That is, the true Q-values of these actions are low with a large probability. Therefore, we propose to drop the top-k estimates and use the average of the others for policy optimization to alleviate the effect of such overestimation. Specifically, given a state s and action a , we sort the estimates produced by the multi-head Q-network in ascending order, denoted by $\hat{Q}_{\theta_j}(s, a)$ with $j \in [1..K]$, and then optimize the policy π_ϕ by minimizing the objective as proposed in (Haarnoja et al. 2018)

$$J_\pi(\phi) = \mathbb{E}_{s \sim \mathcal{D}_{\text{model}}} [\text{D}_{\text{KL}}(\pi_\phi(\cdot|s) \parallel \frac{\exp(\frac{1}{\alpha} \hat{Q}(s, \cdot))}{Z(s)})], \quad (1)$$

where $\hat{Q}(s, \cdot) = \frac{1}{K-L} \sum_{j=1}^{K-L} \hat{Q}_{\theta_j}(s, \cdot)$, and $Z(s_t) = \int_a \exp(\frac{1}{\alpha} \hat{Q}(s, a)) da$, which does not contribute to the gradient with respect to the new policy.

In contrast to CMBAC, existing methods use the uncertainty of the Q-value as a penalty. However, we find it may be sensitive to the penalty coefficient (See Appendix C).

Discussion

We discuss some advantages of CMBAC in this subsection.

Capture global uncertainty CMBAC naturally captures the global uncertainty (O’Donoghue et al. 2018; Zhou, Li, and Wang 2020), which considers the compounding model error and its effect on the critic learning. The global uncertainty allows CMBAC to deal with the overestimation that comes from the long-term prediction errors of the model. Otherwise, conservatively optimizing the policy via local uncertainty (i.e., the one-step prediction error) may mislead the agent into exploiting the regions where the one-step prediction is accurate, but multi-step prediction errors are large.

Capture uncertainty with various granularity CMBAC approximates a posterior distribution over Q-values to implicitly capture the uncertainty of Q-value. To achieve granularity in uncertainty capturing, CMBAC controls the granularity of distribution approximation by varying M , i.e., the number of neural networks in each model. The number of estimates K depends on M , and different K correspond to the different granularity of distribution approximation.

Flexibly control the degree of conservatism CMBAC can flexibly control the degree of conservatism by varying M and L . By varying M , CMBAC can provide fine-grained control over the degree of conservatism, as it achieves granularity in capturing uncertainty. By varying L , the degree of conservatism increases with L .

Experiments

Our experiments have four main goals in this section: (1) Test whether CMBAC can significantly outperform state-of-the-art methods. (2) Perform carefully designed ablation study of CMBAC. (3) Perform visualization experiments of CMBAC to explain its effectiveness. (4) Test the robustness of CMBAC in noisy environments.

Comparative Evaluation

For model-based methods, we compare our method to model-based policy optimization (MBPO) (Janner et al.

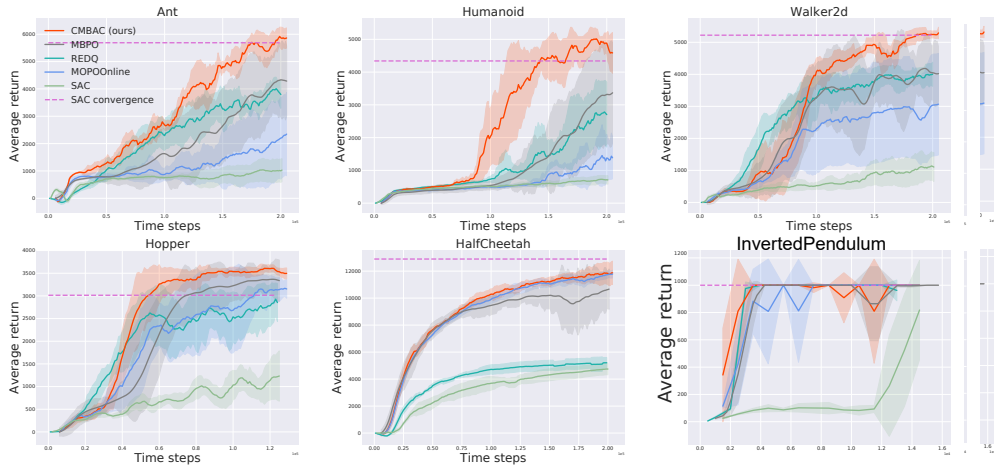


Figure 2: Performance of CMBAC and four baselines on six continuous control tasks. The solid curves correspond to the mean and the shaded region to the standard deviation over five random seeds. For visual clarity, we smooth curves uniformly. The results show that CMBAC significantly outperforms these baselines in terms of sample efficiency on several challenging tasks.




Figure 3: Performance of CMBAC and its ablations on the Ant and Walker2d environments. On the Ant environment, each component of CMBAC is significant for performance improvement. On the Walker2d environment, CMBAC reduces to B-LMEQ-MBPO, as conservative policy optimization does not further improve the performance.

2019), a state-of-the-art algorithm. In addition, we compare to an online variant of model-based offline policy optimization (MOPO) (Yu et al. 2020)—a state-of-the-art offline model-based algorithm—which quantifies the uncertainty of the models and uses the uncertainty as a penalty for policy optimization. For a fair comparison, we implement CMBAC and MOPO-Online, i.e., the online variant of MOPO, both on top of the MBPO. Although masked model-based actor-critic (M2AC) (Pan et al. 2020) outperforms MBPO by using the uncertainty of models, its core component is a masking mechanism, which is orthogonal to our method. Therefore, we do not compare to M2AC. For model-free methods, we compare to soft actor-critic (SAC) (Haarnoja et al. 2018), which is used for policy learning in our method; randomized ensembled double Q-learning (REDQ) (Chen et al. 2021), which has achieved comparable sample efficiency to MBPO.

We evaluate CMBAC and these baselines on MuJoCo (Todorov, Erez, and Tassa 2012) benchmark tasks as used in MBPO. We use two multi-head Q-networks with three hidden layers of 512 neurons each. For all environments except Walker2d, we use the number of dropped estimates $L = 1$.

On Walker2d, we use $L = 0$. For our method, we select the hyperparameter M for each environment independently via grid search. The best hyperparameter for Humanoid, Hopper, Walker2d, and the rest is $M = 1, 3, 4, 2$, respectively. The details of the experimental setup are in Appendix B.

Figure 2 shows that CMBAC significantly outperforms these baselines in terms of sample efficiency on several challenging control tasks. For the most challenging Humanoid environment, CMBAC learns substantially faster than state-of-the-art methods. Specifically, the performance of CMBAC on the Humanoid task at 200 thousand steps matches that of MBPO at 300 thousand steps and SAC at 3 million steps. MOPO-Online achieves poor results on several challenging tasks, which may suggest that its uncertainty-based penalty is inappropriate for the online setting (please refer to Figure 6). The model-free method REDQ has recently achieved comparable sample efficiency to MBPO, which raises the question of whether model-based methods carry the promise of being data efficient. Our results demonstrate that model-based methods can be more sample efficient than their model-free counterparts with careful model usage.

To demonstrate the hyperparameter insensitivity of CMBAC, we replot Figure 2 using unified hyperparameters. Detailed results are in Appendix E. To demonstrate the scalability of CMBAC, we compare CMBAC and the baselines on two additional environments, i.e., Walker2d-NT and Hopper-NT (Wang et al. 2019) (See Appendix E).

Ablation Study

In this subsection, we perform carefully designed ablation study to understand the superiority of CMBAC on Ant, which is one of the most challenging environments, and Walker2d, which has the highest resolution power among all 2D environments. We provide additional results in Appendix C. All results are reported over at least 4 random seeds.

Contribution of each component The path from MBPO (Janner et al. 2019) to CMBAC comprises three modifica-




Figure 4: (1) Left two: varying the number of neural networks in a model M . The results demonstrate that M is essential, and there is an optimal number, i.e., $M = 2$. (2) Right two: varying the number of dropped estimates L . The results show that L is essential, and $L = 1$ provides significant performance improvement on Ant.




Figure 5: Learning curves of CMBAC and its variants on the Ant and Walker2d environments. The results show that increasing ensemble Q-values does not boost the performance, and different implementations of conservative policy optimization significantly impact performance.

tions: Q-network size increase (Big), learning multiple estimates of the Q-value (LMEQ), and conservative policy optimization (CPO). B-MBPO is MBPO with an increased size of Q-networks (Big MBPO). The big Q-network has three layers with 512 neurons versus two layers of 256 neurons in MBPO. Note that the policy network size does not change. B-LMEQ-MBPO is B-MBPO with learning multiple estimates of the Q-value to capture the uncertainty of the Q-value function. B-LMEQ-MBPO uses the average of all the estimates to optimize the policy. CMBAC is B-LMEQ-MBPO with conservative policy optimization, i.e., it uses the average of the bottom-k estimates to optimize the policy.

Figure 3 shows that both LMEQ and CPO are significant for performance improvement (though CPO does not improve the performance on Walker2d). The increased Q-network size does not improve MBPO. The uncertainty estimation method of CMBAC improves its performance in both environments. A possible reason is that it may incorporate more information of the model into the multi-head Q-network and thus improves the value estimation. For the Ant environment, conservative policy optimization provides an additional performance improvement. For the Walker2d environment, we find that our method without conservative policy optimization performs better than that with it (please refer to Figure 4). One possible reason is that the agent may require an optimistic value estimate to promote exploration, and thus avoid suboptimal policies on Walker2d.

Sensitivity to hyperparameters In this part, we analyze the sensitivity of CMBAC to the number of neural networks

in each model M and the number of dropped estimates L . Please refer to Appendix C for additional results.

The number of neural networks in a model M We vary the number $M \in \{1, 2, 3, 4\}$ with $N = 5$ on the Ant and Walker2d environments. Figure 4 shows that (1) the number M is essential as it controls the granularity of uncertainty capturing, and (2) there is an optimal number, i.e., $M = 2$.

The number of dropped estimates L We vary the number of dropped estimates for each ensemble multi-head Q-network $L \in \{0, 1, 2\}$. The total number of estimates dropped is $2L$. Figure 4 suggests that dropping L topmost estimates provides a performance improvement on the Ant environment but degrades the performance on the Walker2d.

Analysis of CMBAC variants To provide further insight into CMBAC, we discuss the performance of some CMBAC variants. These variants implement the two core components of CMBAC, i.e., capturing the uncertainty of the Q-value and conservative policy optimization, in different ways.

LMEQ variants To capture the uncertainty of the Q-value, we propose to increase the number of ensemble Q-value functions, i.e., deep ensembles (Lakshminarayanan, Pritzel, and Blundell 2017). Thus, we propose Model-Based Policy Optimization with Ensemble Q (MBPOEQ), which uses the same multi-head Q-networks as our method. In contrast to CMBAC, MBPOEQ learns all “heads” using the same target, i.e., the average of these “heads”.

CPO variants We additionally propose three possible conservative policy optimization variants, named CMBAC Uncertainty Penalty (CMBACUP), MINimum CMBAC (MINCMBAC), and REDQ-CMBAC, respectively. CMBACUP uses the standard deviation of the Q-value estimates as a penalty. MINCMBAC uses the minimum of all estimates to optimize the policy. Similar to REDQ (Chen et al. 2021), REDQ-CMBAC produces a conservative estimate via each multi-head Q-network and uses the mean of the two conservative estimates to optimize the policy. Due to limited space, we provide the details in Appendix D.

The results in Figure 5 suggest the following conclusions. CMBACUP achieves comparable performance to CMBAC in both environments. However, we find it sensitive to the penalty coefficient (please refer to Appendix C). Increasing the number of ensemble Q-values does not improve performance, indicating that our uncertainty capturing technique is critical to performance improvement. Incorporating con-




Figure 6: We visualize our uncertainty estimation compared to previous methods on the InvertedPendulum environment. Colored points illustrate the uncertainty versus the value estimation error at given state-action pairs. Results show that our uncertainty estimation approximates the errors of the Q-function more accurately than previous methods. (For a fair comparison, we use a fixed policy to generate rollouts in the true environment for computing the real return and use its corresponding Q-value network to estimate the Q-value.)

servatism is significant for performance improvement, and different implementations of conservative policy optimization have a large impact on performance.

Visualization

Uncertainty estimation In this part, we regard the standard deviation of the multiple heads as our estimated uncertainty and compare it to previous uncertainty estimation used in state-of-the-art model-based algorithms. The uncertainty estimation methods include Global, i.e., the cumulative discounted sum of prediction errors similar to that used in model-based offline policy optimization (Yu et al. 2020), M2AC that is used in masked model-based actor-critic (Pan et al. 2020), and SLBO that is used in stochastic lower bounds optimization (Luo et al. 2019). We visualize the prediction error of the Q-value and the estimated uncertainty via scatters. The results in Figure 6 show that our estimated uncertainty can approximate the errors of the Q-function more accurately than previous methods. The superiority of our uncertainty estimation demonstrates that the multi-head Q-networks used in CMBAC can provide a reasonable approximation of the posterior distribution of the true Q-value. Moreover, we find that the cumulative discounted sum of prediction errors can hardly approximate the errors of the Q-function. We provide possible reasons in Appendix D.

Conservative policy optimization We visualize the estimates of the Q-value in each model versus the true return on the 2D point environment. To better understand the superiority of the proposed conservative policy optimization, we use a simplified version of CMBAC (please refer to Appendix D for details). Figure 7 (left) shows that there are a small




Figure 7: We visualize the estimates of the Q-value in each model on the 2D point environment. Colored points illustrate the estimate versus the true return at given state-action pairs. The results demonstrate that dropping several topmost estimates effectively avoid unreliable actions.




Figure 8: Performance of CMBAC, MBPO, and MOPO-Online in noisy environments (100k steps for HalfCheetah and 200k steps for Walker2d). The results show that our method is more robust than state-of-the-art methods.

fraction of models in which the value estimates significantly overestimate due to model errors. Moreover, Figure 7 shows that dropping several topmost estimates effectively encourages the agent to avoid the unreliable “promising actions”—whose values are high in only a small fraction of models.

Robust Analysis

To understand the robustness of CMBAC, we compare it to MOPO-Online and MBPO on noisy Walker2d and HalfCheetah environments. In these noisy environments, we add Gaussian white noises with the standard deviation $\sigma = 0.1$ to the agent’s action at every step. The added noise will decrease the accuracy of the learned models. The results in Figure 8 show that CMBAC performs robustly and significantly outperforms the baselines in terms of sample efficiency in noisy environments.

Conclusion

In this paper, we present conservative model-based actor-critic, a novel approach that approximates a posterior distribution over Q-values based on the ensemble models and uses the average of the left tail of the distribution approximation to optimize the policy. Experiments show that CMBAC significantly outperforms state-of-the-art methods in terms of sample efficiency on several challenging control tasks. Moreover, experiments demonstrate that the proposed method is more robust than previous methods in noisy environments. We believe that our proposed approach will bring new insights into model-based reinforcement learning.

Acknowledgements

We would like to thank all the anonymous reviewers for their insightful comments. This work was supported in part by National Science Foundations of China grants U19B2026, U19B2044, 61822604, 61836006, and 62021001, and the Fundamental Research Funds for the Central Universities grant WK3490000004.

References

- Amos, B.; Stanton, S.; Yarats, D.; and Wilson, A. G. 2021. On the Model-Based Stochastic Value Gradient for Continuous Reinforcement Learning. In *Proceedings of the 3rd Conference on Learning for Dynamics and Control*, volume 144 of *Proceedings of Machine Learning Research*, 6–20. PMLR.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. OpenAI Gym. *CoRR*, abs/1606.01540.
- Chen, X.; Wang, C.; Zhou, Z.; and Ross, K. W. 2021. Randomized Ensembled Double Q-Learning: Learning Fast Without a Model. In *International Conference on Learning Representations*.
- Chua, K.; Calandra, R.; McAllister, R.; and Levine, S. 2018. Deep Reinforcement Learning in a Handful of Trials using Probabilistic Dynamics Models. In *Advances in Neural Information Processing Systems*, volume 31.
- Clavera, I.; Fu, Y.; and Abbeel, P. 2020. Model-Augmented Actor-Critic: Backpropagating through Paths. In *International Conference on Learning Representations*.
- Clavera, I.; Rothfuss, J.; Schulman, J.; Fujita, Y.; Asfour, T.; and Abbeel, P. 2018. Model-Based Reinforcement Learning via Meta-Policy Optimization. In *Proceedings of The 2nd Conference on Robot Learning*, volume 87 of *Proceedings of Machine Learning Research*, 617–629. PMLR.
- de Boer, P.; Kroese, D. P.; Mannor, S.; and Rubinstein, R. Y. 2005. A Tutorial on the Cross-Entropy Method. *Ann. Oper. Res.*, 134(1): 19–67.
- Deisenroth, M.; and Rasmussen, C. E. 2011. PILCO: A Model-Based and Data-Efficient Approach to Policy Search. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, ICML '11, 465–472. ACM. ISBN 978-1-4503-0619-5.
- Doyle, J. 1996. Robust and optimal control. In *Proceedings of 35th IEEE Conference on Decision and Control*, volume 2, 1595–1598 vol.2.
- Efron, B.; and Tibshirani, R. 1993. *An Introduction to the Bootstrap*. Springer. ISBN 978-1-4899-4541-9.
- Fairbank, M.; and Alonso, E. 2012. Value-gradient learning. In *The 2012 International Joint Conference on Neural Networks (IJCNN)*, 1–8.
- Fujimoto, S.; van Hoof, H.; and Meger, D. 2018. Addressing Function Approximation Error in Actor-Critic Methods. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 1587–1596. PMLR.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 1861–1870. PMLR.
- Heess, N.; Wayne, G.; Silver, D.; Lillicrap, T.; Erez, T.; and Tassa, Y. 2015. Learning Continuous Control Policies by Stochastic Value Gradients. In *Advances in Neural Information Processing Systems*, volume 28.
- Hessel, M.; Modayil, J.; Van Hasselt, H.; Schaul, T.; Ostrovski, G.; Dabney, W.; Horgan, D.; Piot, B.; Azar, M.; and Silver, D. 2018. Rainbow: Combining improvements in deep reinforcement learning. In *Thirty-second AAAI conference on artificial intelligence*.
- Janner, M.; Fu, J.; Zhang, M.; and Levine, S. 2019. When to Trust Your Model: Model-Based Policy Optimization. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Kumar, A.; Zhou, A.; Tucker, G.; and Levine, S. 2020. Conservative Q-Learning for Offline Reinforcement Learning. In *Advances in Neural Information Processing Systems*, volume 33, 1179–1191. Curran Associates, Inc.
- Kurutach, T.; Clavera, I.; Duan, Y.; Tamar, A.; and Abbeel, P. 2018. Model-Ensemble Trust-Region Policy Optimization. In *International Conference on Learning Representations*.
- Kuznetsov, A.; Shvechikov, P.; Grishin, A.; and Vetrov, D. 2020. Controlling Overestimation Bias with Truncated Mixture of Continuous Distributional Quantile Critics. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 5556–5566. PMLR.
- Lakshminarayanan, B.; Pritzel, A.; and Blundell, C. 2017. Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2016. Continuous control with deep reinforcement learning. In *4th International Conference on Learning Representations, ICLR 2016*.
- Lim, S. H.; Xu, H.; and Mannor, S. 2013. Reinforcement Learning in Robust Markov Decision Processes. In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.
- Luo, Y.; Xu, H.; Li, Y.; Tian, Y.; Darrell, T.; and Ma, T. 2019. Algorithmic Framework for Model-based Deep Reinforcement Learning with Theoretical Guarantees. In *International Conference on Learning Representations*.
- 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. 2015. Human-level control through deep reinforcement learning. *nature*, 518(7540): 529–533.
- Nagabandi, A.; Kahn, G.; Fearing, R. S.; and Levine, S. 2018. Neural Network Dynamics for Model-Based Deep Reinforcement Learning with Model-Free Fine-Tuning. In

- 2018 *IEEE International Conference on Robotics and Automation (ICRA)*, 7559–7566.
- Nguyen, D.; and Widrow, B. 1990. Neural networks for self-learning control systems. *IEEE Control Systems Magazine*, 10(3): 18–23.
- O’Donoghue, B.; Osband, I.; Munos, R.; and Mnih, V. 2018. The Uncertainty Bellman Equation and Exploration. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 3839–3848. PMLR.
- Osband, I.; Blundell, C.; Pritzel, A.; and Van Roy, B. 2016. Deep Exploration via Bootstrapped DQN. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.
- Ovadia, Y.; Fertig, E.; Ren, J.; Nado, Z.; Sculley, D.; Nowozin, S.; Dillon, J.; Lakshminarayanan, B.; and Snoek, J. 2019. Can you trust your model’s uncertainty? Evaluating predictive uncertainty under dataset shift. In *Advances in Neural Information Processing Systems*, volume 32.
- Pan, F.; He, J.; Tu, D.; and He, Q. 2020. Trust the Model When It Is Confident: Masked Model-based Actor-Critic. In *Advances in Neural Information Processing Systems*, volume 33, 10537–10546.
- Punjani, A.; and Abbeel, P. 2015. Deep learning helicopter dynamics models. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 3223–3230.
- Rajeswaran, A.; Ghotra, S.; Ravindran, B.; and Levine, S. 2017. EPOpt: Learning Robust Neural Network Policies Using Model Ensembles. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Strehl, A. L.; and Littman, M. L. 2008. An analysis of model-based Interval Estimation for Markov Decision Processes. *J. Comput. Syst. Sci.*, 74(8): 1309–1331.
- Sutton, R. S. 1990. Integrated Architectures for Learning, Planning, and Reacting Based on Approximating Dynamic Programming. In *Machine Learning, Proceedings of the Seventh International Conference on Machine Learning, Austin, Texas, USA, June 21-23, 1990*, 216–224. Morgan Kaufmann.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning: An introduction*. MIT press.
- Todorov, E.; Erez, T.; and Tassa, Y. 2012. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 5026–5033.
- Vuong, T.; and Tran, K. 2019. Uncertainty-aware Model-based Policy Optimization. *CoRR*, abs/1906.10717.
- Wang, T.; and Ba, J. 2020. Exploring Model-based Planning with Policy Networks. In *International Conference on Learning Representations*.
- Wang, T.; Bao, X.; Clavera, I.; Hoang, J.; Wen, Y.; Langlois, E.; Zhang, S.; Zhang, G.; Abbeel, P.; and Ba, J. 2019. Benchmarking Model-Based Reinforcement Learning. *CoRR*, abs/1907.02057.
- Yu, T.; Kumar, A.; Rafailov, R.; Rajeswaran, A.; Levine, S.; and Finn, C. 2021. COMBO: Conservative Offline Model-Based Policy Optimization. *CoRR*, abs/2102.08363.
- Yu, T.; Thomas, G.; Yu, L.; Ermon, S.; Zou, J. Y.; Levine, S.; Finn, C.; and Ma, T. 2020. MOPO: Model-based Offline Policy Optimization. In *Advances in Neural Information Processing Systems*, volume 33, 14129–14142. Curran Associates, Inc.
- Zhou, Q.; Li, H.; and Wang, J. 2020. Deep Model-Based Reinforcement Learning via Estimated Uncertainty and Conservative Policy Optimization. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04): 6941–6948.