MUSBO: Model-based Uncertainty Regularized and Sample Efficient Batch Optimization for Deployment Constrained Reinforcement Learning

DiJia(Andy) Su [1]   **Jason D. Lee** [1]   **John M. Mulvey** [2]   **H. Vincent Poor** [1]

Abstract

In many contemporary applications such as healthcare, finance, robotics, and recommendation systems, continuous deployment of new policies for data collection and online learning is either cost ineffective or impractical. We consider a setting that lies between pure offline reinforcement learning (RL) and pure online RL called deployment constrained RL in which the number of policy deployments for data sampling is limited. To solve this challenging task, we propose a new algorithmic learning framework called Model-based Uncertainty regularized and Sample Efficient Batch Optimization (MUSBO). Our framework discovers novel and high quality samples for each deployment to enable efficient data collection. During each offline training session, we bootstrap the policy update by quantifying the amount of uncertainty within our collected data. In the high support region (low uncertainty), we encourage our policy by taking an aggressive update. In the low support region (high uncertainty) when the policy bootstraps into the out-of-distribution region, we downweight it by our estimated uncertainty quantification. Experimental results show that MUSBO achieves state-of-the-art performance in the deployment constrained RL setting.

1. Introduction

Recent advances in deep learning have enabled reinforcement learning (RL) to achieve remarkable success in various applications (Silver et al., 2017)(OpenAI et al., 2019b)(Vinyals et al., 2019)(OpenAI et al., 2019a). However, despite RL's success, it suffers many problems. In particular, traditional RL algorithms require the agent to interact with the real world to collect large amounts of on-

[1]Department of Electrical and Computer Engineering, Princeton University [2]Department of Operation Research and Financial Engineering, Princeton University. Correspondence to: DiJia(Andy) Su <andy2008su@gmail.com>.

line data with the latest learned policy. However, the online deployment of the agent to the real world might be impossible in many real world applications (Matsushima et al., 2020). In the field of robotics or self-driving cars, for example, the cost of deploying the agent to the field may be too risky for the agent itself as well as its surrounding environment. In quantitative finance, trading strategy is usually carefully back-tested and calibrated offline with historical data and simulation. Since the market data has a low signal-to-noise ratio (Chen et al., 2020), online training can easily lead the policy fits to the noise, which is dangerous and can potentially trigger unexpected large monetary loss. In the recommendation system setting (Peska & Vojtas, 2020), after the policy has been trained offline, it will be deployed across different servers for serving many users. In such setting, online training of policy is difficult because the resulting policy might become unstable and/or cause bad user experience. Thus, large chunk of data is collected for the deployed policy. Then, the data is used for offline training. During the next scheduled deployment, the production level (data collection) policy is then updated.
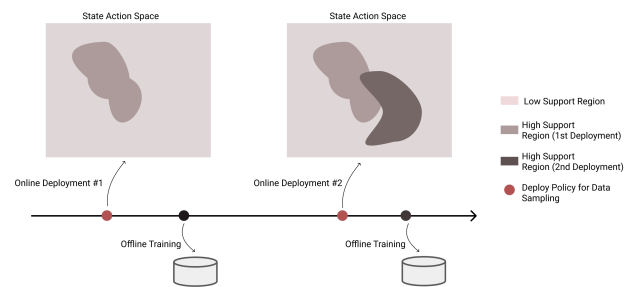

Figure 1. Illustration showing the deployment constrained RL setting. Large batch of data collection only occurs at the deployment point (showing two of them in red). The deployment and offline training occurs in an interleaved fashion. As the number of deployment increases, the state-action space will be explored more and more. Here, thicker-brown colored regions show the area of high support whereas the lighter color shows an under-explored (low support) region.

To address this challenge, training an RL agent in an offline fashion seems to offer solution. In offline RL (Levine et al.,

2020), a static dataset is collected by a behavioral (or data collection) policy. Since the RL agent has no access to the online environment, training the RL agent faces many challenges. First, because the learning policy and the behavioral policy have different state visitation frequencies, evaluation of the offline policy is difficult. Second, during the training process, this discrepancy of the distribution might increase over the training time, a phenomenon known as distributional shift. Third, there might be a large extrapolation error when the value function is bootstrapped into out-of-distribution actions (Kumar et al., 2019), a phenomenon that can lead to learning divergence and instability. There is a large body of research work on pure offline RL. Despite their empirical success, since the pure offline RL environment is fundamentally different than our setting, we suspect developing novel deployment constrained RL can possibly achieve a large headroom of improvement, since offline RL algorithms assume zero interaction with the environment.

In this work, we present Model-based Uncertainty regularized and Sample Efficient Batch Optimization (MUSBO), a general framework for doing batch policy optimization in the deployment constrained setting. At a high level, our method encourages the data collection policy to *discover high quality and novel data batches during each online deployment*. During the offline training, MUSBO performs policy update weighted with *uncertainty regularized coefficient*, a term that quantifies the amount of uncertainty with respect to each state-action pair within the data. Our optimization process regularizes the update toward the high confidence regions. In areas of low confidence or low data support, MUSBO takes a pessimistic point of view and discounts the update by the uncertainty regularized coefficient when the policy bootstraps into out-of-distribution states and actions regions. We empirically compare our MUSBO to other strong baselines such as the state-of-the-art deployment constrained RL algorithms (BREMEN) (Matsushima et al., 2020), and we show that MUSBO is capable of achieving significant policy improvement while using smaller amounts of data and fewer deployment.

2. Related Works

The related research literature can be broadly categorized into three categories: deployment-constrained RL, offline RL, and model-based RL.

Deployment-constrained RL To the best of our knowledge, (Matsushima et al., 2020) is the first work that proposed the concept of deployment-constrained efficiency. In their paper, the authors proposed the algorithm Behavior-Regularized Model-ENsemble (BREMEN) that enforces KL-divergence between the learning policy and the behavioral policy for learning update. They compared their approach across numerous baseline such as Soft-Actor-

Critic (SAC)(Haarnoja et al., 2018), Model-based-TRPO (METRO)(Kurutach et al., 2018a), and show that their methodology provides the strongest empirical performance. Here in our paper, we compare our method MUSBO directly with a state-of-the-art method (BREMEN). Different from BREMEN, we first provide a theoretical analysis characterizing the improvement in terms of value function and optimization lower bounds. Second, our method is different because we emphasize at 1) propose to regularize the learning policy update and the behavioral policy by weighting each state-action pair with uncertainty quantification, 2) propose a method of measuring state action uncertainty by utilizing a new set of dynamics models with next state estimation error, 3) propose to make use of uncertainty quantification to maximize the discovery of novel data transitions during each deployment.

On the other hand, there also some related research works such as (Bai et al., 2020) and in semi-batch RL such as (Ernst et al., 2005; Lange et al., 2012; Jaakkola et al., 1999; Chu & Kitani, 2020).

Offline RL. Unlike deployment-constrained setting, Offline RL assumes no interaction with the environment. Thus, the learning policy needs to reason about the behavioral policy and make policy update base on that. The two most related literatures are Yu et al. (2020) and Kidambi et al. (2020). In Yu et al. (2020), the authors proposed an uncertainty penalized MDPs in which the reward function was used to explicitly penalize the uncertainty. On the other hand, the authors from Kidambi et al. (2020) proposed a pessimistic MDP that divides the environment into two regions: known or unknown. When the agent is entering into the unknown region, the MDP undergoes a halt state (or absorbing state) in which a large negative reward will be assigned to penalize this action. In both cases, the proposed uncertainty penalization comes from the reward function, which is a totally different setup than what is being proposed in our paper. Here, our method enables a regularization approach. Our method regularizes our policy update toward high confidence region while down-weights or regularized it away when the policy bootstraps into the unfamiliar *state and actions* regions.

Besides these two, there is also a large body of offline RL research. In model-free offline RL, the common techniques are either 1) enforcing the learning policy to stay close with the behavioral policy as in Fujimoto et al. (2019); Kumar et al. (2019); Wu et al. (2019); Nachum et al. (2019b); Zhang et al. (2020); Nachum et al. (2019a) or ensembles of Q values for stabilizing the learning and behaviors as in Ghasemipour et al. (2021); Wu et al. (2019); Nair et al. (2020).

Model-Based RL. In model-based approach, the world representation is learned first and then is used for generating

imaginary rollouts. Related research works are Chua et al. (2018); Janner et al. (2019); Luo et al. (2019); Munos & Szepesvári (2008). However, direct application of MBRL methods into the offline setting can be challenging due to distribution shifts. Nevertheless, the closely related research is Kurutach et al. (2018a), in which the authors proposed to use an ensemble of estimated model dynamics for generating imaginary rollous for stabilizing effects. Similar approaches have also been investigated by Zhang et al. (2019); Kaiser et al. (2020); Veerapaneni et al. (2020); Feinberg et al. (2018).

2.1. Background

We consider a discounted, infinite horizon *Markov decision process (MDP)*, and denote $\Omega = (S, A, M, r, \mu_0, \gamma)$, where S is the state space, A is the action space, $M(s'|s,a)$ is the state transition kernel of transiting to a next state s' from state s while taking action a, $r(s,a)$ is the reward function, $\gamma \in (0,1)$ is the discounting factor, and μ_0 is the initial state distribution.

In RL, the goal is to optimize a policy $\pi(a|s)$ such that the expected discounted return $J_M(\pi)$ is maximized: $J_M(\pi) = \mathop{\mathbf{E}}_{\pi,M,\mu_0} \sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)$. The value function is defined as $V_M^\pi(s) = \mathop{\mathbf{E}}_{\pi,s_{t+1}\sim M(s_t,a_t)} \{\sum_{t=0}^{\infty} \gamma^t r(s_t,a_t)|s_0 = s\}$. We denote the optimal policy $\pi^* = \operatorname{argmax}_\pi J_M(\pi)$. We further define ρ_M^π to be the discounted state visited by π on M such that $\rho_M^\pi = (1-\gamma)\sum_{t=0} \gamma^t p_{S_t^\pi}$, where $p_{S_t^\pi}$ is the distribution of state at time t.

In the **deployment constrained** setting, there is a limitation on the number of times that the policy can be deployed online for data collection. In real world applications such as recommender system or robotic control, due to the associated cost with each deployment, it is desirable to minimize the number of deployments to as small as possible. Let I be the number of times of deployment, and let B be the size of a large batch of training data collected when a policy has been taken online (or deployed), then, the total size of data collected throughout the entire process is $I \times B$. Further, let $B^{(i)}$ be the batch of data collected during the i_{th} deployment. In each batch B, we store the data transition $\{(s, a, r, s')\}$. Note the difference between the deployment constrained setting and the pure offline RL setting. In the pure offline RL setting, we consider it as a single batch of data with $I = 1$, whereas in the deployment constrained setting, the agent will have the opportunity to interact with the environment by deploying its learning or other behavioral policy for data collection purposes.

In **Model-Based RL** (MBRL), the world dynamics are learned from the sampled data. We estimate the ground truth environment transition dynamics M^* by learning a forward next state transition dynamics model $\hat{M} : S \times A \to S$.

Let $\hat{M}^{(i)}$ be the estimated dynamics bootstrapped from data batch $B^{(1)} \bigcup B^{(2)} \bigcup ... \bigcup B^{(i)}$, in which we assume the model is updated with the aggregated data batches up to the i_{th} deployment. We denote by $\hat{M}_\theta^{(i)}$ the estimated dynamics learned by neural networks parameterized by θ.

3. Algorithmic Framework

For the purpose of exposition, we start by presenting an idealized version of MUSBO algorithmic framework. Then, we describe a practical version that we have implemented and perform quite well. The detailed algorithm is summarized in Algorithm 1.

The heart of the algorithm relies on uncertainty estimation. To get the most out of each deployment, our algorithm tries to explore the under-represented region (the low support area of our collected data) so as to minimize the amount of uncertainty or maximize the information gain. We postulate that by having more diversified samples, we can learn a more accurate representation of the environment. After collecting high quality data (batch of transitions), the idea is that we want to weight each policy update by taking the uncertainty quantification on each state action pair by **uncertainty regularized coefficient** ($U_{\hat{M}^{(i)}, \hat{M}^*}(s,a)$) which characterize the level of uncertainty results from model estimation error (eq. (6)). For low support region, we discount the policy update by $U_{\hat{M}^{(i)}, \hat{M}^*}(s,a)$ and discourage the policy from bootstrapping into unknown regions. In doing this, we are *regularizing* the policy update into high support/confidence regions.

3.1. Theoretical Results

All the proofs in this section are deferred to the appendix. In the deployment constrained setting, we are interested in iterative improvement after each deployment. Here, we characterize the iterative improvement by developing a lower bound for V_M^π with $D(\cdot, \cdot)$ as the discrepancy between the M^* and $\hat{M}$:

$$V_M^\pi \geq V_{\hat{M}}^\pi - D(\hat{M}, \pi) \tag{1}$$

Let $u(\hat{M}, \pi)$ be a scalar such that

$$u(\hat{M}^{(i)}, \pi) \triangleq \frac{V_{\hat{M}}^\pi - D(\hat{M}^{(i)}, \pi)}{V_{\hat{M}}^\pi} \tag{2}$$

Let π_{ref} to be a reference policy that we will deploy to collect the data. Then, we make the followings three assumptions:

$$V_{M^*}^\pi \geq V_{\hat{M}}^\pi u_{\pi_{\text{ref}},\delta}(\hat{M}, \pi), \text{ s.t.} d(\pi, \pi_{\text{ref}}) \leq \delta \tag{A1}$$

$$\hat{M} = M^* \implies D_{\pi_{\text{ref}}}(\hat{M}, \pi) = 0$$
$$\implies u_{\pi_{\text{ref}},\delta}(\hat{M}, \pi) = 1 \tag{A2}$$

$D_{\pi_{\text{ref}}}(\hat{M}, \pi)$ is given by the form of $E_{\tau \sim \pi_{\text{ref}}, M^*}\{f(\hat{M}, \pi, \tau)\}$ (A3)

Starting from the monotone improvement result of Luo et al. (2019); Schulman et al. (2017), we have:

Theorem 1. *Let* $\pi_k, \hat{M}_k = \underset{\pi \in \Pi, M}{argmax} \, u_{\pi_k, \delta}(M, \pi)V_M^\pi$ *be the policy and model optimized offline after each deployment k such that $d(\pi, \pi_k) \leq \delta$. Then, the sequence of $\pi_0, \pi_1, \pi_2, ..\pi_k$ yields monotonically increasing values:*

$$V_{M^*}^{\pi_0} \leq V_{M^*}^{\pi_1} \leq .. \leq V_{M^*}^{\pi_k} \tag{3}$$

As a result of Theorem 1, our MUSBO formulation (line 4 and line 2) which iteratively builds a model $\hat{M}_\theta$ and π will have its resulting value function getting better and better.

Algorithm 1 MUSBO: Model-based Uncertainty Regularized and Sample Efficient Batch Optimization

---

Given: Size of data batches B, Number of deployments I
$D_{all} \leftarrow \{\}$
for $i = 1, 2, 3, ..., I$ *of deployments* **do**
    deploy π online for data collection
1    $D_{all} \leftarrow D_{all} \bigcup$ Collect-Data-Low-Support-Region(π)
2    $\hat{M}_\theta \leftarrow$ Learn approximate transition dynamics model
3   Train Uncertainty-Labeler(D_{all})
4    $\pi \leftarrow$ Train with MBRL(Uncertainty-Labeler, $\hat{M}_\theta, D_{all}$)
        as in section 3.3
end
return π

Now, we define the following function $d_{\hat{M}^{(i)}, \hat{M}^{(i+1)}}(s, a)$ such that it bounds the estimation error between of value function of the approximated $\hat{M}^{(i)}$ and next deployed estimated $\hat{M}^{(i+1)}$ from above:

$d_{\hat{M}^{(i)}, \hat{M}^{(i+1)}}(s, a)$
$\geq \Big| \underset{\hat{s}' \sim \hat{M}^{(i)}(\cdot|s,a)}{\mathrm{E}}\{V_{\hat{M}^{(i)}}^\pi(\hat{s}')\} - \underset{\hat{s}' \sim \hat{M}^{(i+1)}(\cdot|s,a)}{\mathrm{E}}\{V_{\hat{M}^{(i)}}^\pi(\hat{s}')\}\Big|$ (4)

Proposition 1. *Assume there exists a function $d_{\hat{M}^{(i)}, \hat{M}^{(i+1)}}(s, a)$ such that it upper bounds the error induced by the model estimation between the approximated $\hat{M}$ and the ground truth M^* (as in eq. (4)). Then:*

$$V_{\hat{M}^{(i+1)}}^\pi \geq V_{\hat{M}^{(i)}}^\pi - \kappa \underset{(s,a) \sim \rho_{\hat{M}^{(i+1)}}^\pi}{\mathrm{E}}\{d_{\hat{M}^{(i)}, \hat{M}^{(i+1)}}(s, a)\}, \tag{5}$$

where $\kappa \triangleq (1 - \gamma)^{-1}\gamma$. We further define the following function $U_{\hat{M}^{(i)}, \hat{M}^{(i+1)}}(s, a): S \times A \to \mathbb{R}$ as an **uncertainty regularized coefficient**:

$$\underset{(s,a) \sim \rho_{\hat{M}^{(i+1)}}^\pi}{\mathrm{E}}\{U_{\hat{M}^{(i)}, \hat{M}^{(i+1)}}(s, a)\} \triangleq$$
$$\Big(V_{\hat{M}^{(i)}}^\pi - \kappa \underset{(s,a) \sim \rho_{\hat{M}^{(i+1)}}^\pi}{\mathrm{E}}\{d_{\hat{M}^{(i)}, \hat{M}^{(i+1)}}(s, a)\}\Big)(V_{\hat{M}^{(i)}}^\pi)^{-1} \tag{6}$$

As an immediate corollary, we combine eq. (6) and Proposition 2 to establish a similar lower bound between $V_{\hat{M}^{(i+1)}}^\pi$ and the ground-truth $V_{\hat{M}^*}^\pi$:

Corollary 1. *Assume the existence of the uncertainty regularized coefficient $U_{\hat{M}^{(i+1)}, M^*}(s, a)$ as defined by eq. (6), then the performance of π on the ground-truth M^* is lower bounded by that of the estimated $\hat{M}^{(i+1)}$, weighted by $U_{\hat{M}^{(i+1)}, M^*}(s, a)$*

$$V_{M^*}^\pi \geq V_{\hat{M}^{(i+1)}}^\pi \underset{(s,a) \sim \rho_{M^*}^\pi}{\mathrm{E}}\{U_{\hat{M}^{(i+1)}, M^*}(s, a)\}. \tag{7}$$

Corollary.1 says that in optimizing the lower bound (the RHS of eq. (7)) with $V_{\hat{M}^{(i+1)}}^\pi$, we can maximize the overall performance of π on the real dynamics M^*.

Interpretation of $U_{\hat{M}^{(i+1)}, M^*}(s, a)$ (uncertainty regularized coefficient). Here, the term can be interpreted as an uncertainty quantification measure as a result of the model discrepancy between $\hat{M}^{(i+1)}$ and M^*, with $\hat{M}^{(i+1)}$ being learned from the collected (and limited amount of) data. On the regions where we have high data support, $\hat{M}^{(i+1)}(s, a)$ is close to $M^*(s, a)$, thus their discrepancy decreases and $U_{\hat{M}^{(i+1)}, M^*}(s, a)$ approaches to unity. On the regions that are less certain, the discrepancy increases and thus $U_{\hat{M}^{(i+1)}, M^*}(s, a)$ decreases. *Thus to optimize $V_{M^*}^\pi$, we instead optimize its lower bound, which is the uncertainty regularized version of the value function at the estimated model.* We assign a higher weight at the regions of high confidence, and a lower weight at the regions of low confidence. Next, we utilize our result from eq. (7) for establishing our MUSBO lower bound optimization algorithm.

3.2. Practical Implementation

In this section, we explain the practical implementation of algorithm 1 in detail. For readers convenience, we have labeled the line number in the algorithm, and we will refer to the line number as we explain below.

Learning the transition dynamics (line 2): For estimating the transition dynamics $\hat{M}_\theta$, we use an ensemble of N deterministic dynamics models with multi-layer fully-connected perceptrons as in Mishra et al. (2017); Kurutach et al. (2018b) to reduce model bias. Here, each of these models is parameterized a neural network which different weights initialization denoted as $\hat{f}_\theta$. They are trained to predict the next state with L_2 loss:

$$\min_\theta \frac{1}{|D_{all}|} \sum_{(s_t, a_t, s_{t+1}) \in D_{all}} \|s_{t+1} - \hat{f}_\theta(s_t, a_t)\|_2^2. \tag{8}$$

Uncertainty-Labeler (line 3). This module is responsible for characterizing the levels of uncertainty in the collected data and approximated $U_{\hat{M}^{(i+1)}, M^*}(S, A)$. Specifically, we

want to identify the regions in our data that have high support or low support with respect to uncertainty. Since the oracle is unavailable to us, here we can only approximate the uncertainty by state-actions visitation frequency. Shall the state-actions visited frequently, more pairs of them will show up in the data. If we were to train models on the batches of data to predict the next state transition dynamics, then, our prediction will be more accurate on the states that we have seen (in the collected data), and less accurate on the unfamiliar state. Thus, we have established that the uncertainty quantification measure as the next state prediction error. If we train a set of neural networks that learns the dynamics models $\hat{P}_\phi : S \times A \to S$ utilizing previous batch of data $(B^{(1)}, ..B^{(i-1)})$, and for the latest collected batch $(B^{(i)})$, we can quantify the uncertainty by prediction error. The lower the prediction error in relation to the ground truth, the higher the support.

Formally, our uncertainty-labeler approximates the uncertainty regularized coefficient $U_{\hat{M}^{(i)}, M^*}(s, a)$ by $\hat{U}(a, s)$:

$$\hat{U}(a, s) = \max_{i \in \{\hat{P}_\phi\}_i^K} 1 - \text{clip}(\|s_{t+1} - \hat{s}_{t+1}\|_2^2, 0, 1) \quad (9)$$

where $\hat{s}_{t+1}$ is given by $\hat{P}_\phi(s, a)$, and we take the maximum prediction error out of the K ensembles $\hat{P}_\phi$, and clip the error between 0 and 1.

For practical implementation, we use K ensembles of Probabilistic Neural Networks (PNN) (Chua et al., 2018). In PNN, the network has its output neurons parameterized by a Gaussian distribution in the effort of capturing the aleatoric uncertainty. This module is trained to minimize the following loss:

$$\text{loss}_{PNN}(\phi) = \sum_{n=1}^N \{\mu_\phi(s_n, a_n) - s_{n+1}\}^M \Sigma_\phi^{-1}(s_n, a_n)$$
$$(\mu_\phi(s_n, a_n) - s_{n+1}) + \log \det\Sigma_\phi(s_n, a_n) \quad (10)$$

where ϕ is the neural network learning parameters, μ_ϕ and Σ_ϕ are the mean and variance of the Gaussian distribution.

Note that here, we use different sets (types) of networks for estimating the model dynamics $\hat{M}_\theta$ and the uncertainty $\hat{P}_\phi$. We separate out these two networks intentionally to avoid possible error propagation between the fictitious trajectory generating model dynamics($\hat{M}_\theta$) and the uncertainty labelling ($\hat{P}_\phi$).

Collect Data from Low Support Region (line 1). To maximize the benefit out of each deployment, we emphasize that we want to achieve the maximal data coverage as well as exploring the under-explored region (the low support area).

Exploration in the deployment constrained setting is nontrivial because it is hard to evaluate which action-state pairs

will fill the low support region and lead to high data coverage. During online data sampling, for each state that our agent encounters, we identify the amount of uncertainty by taking the actions which lead to maximal *prediction error* between our predicted next state ($\hat{s}'$) and the ground truth next state as by 9. Following a trajectory of maximal prediction error leads to novel experience discovery and reduce the number of unknown regions within the data. (For detailed implementation, please refer to Appendix E).

3.3. Offline Training with MBRL Method

Next, we define our offline model-based training method (line 4).

Function   MUSBO MBRL Training($\hat{M}_\theta$, D_{all}, *Uncertainty-Labeler*) :

5   Initialize π_{init} with behavior cloning as given by eqn. 11

6   **for** *training iterations* **do**

7   Randomly sample dynamics $\hat{M}_\theta$ from K ensembles

     for *optimization steps* **do**

8        $\hat{\tau} \leftarrow$ sample fictitious trajectory from $\hat{M}_\theta$

9        $\hat{\tau}_{labeled} \leftarrow$ label fictitious trajectory with Uncertainty-Labeler ($\hat{\tau}$) as in eq.(12)

10   train with TRPO with $\hat{\tau}_{labeled}$, π_{init} as in eq.(13)

     end

   end

return

In the function above, we provide a practical instantiation of MUSBO offline model-based method for learning a policy.

Our method utilizes trust region policy optimization (TRPO) (Schulman et al., 2017) with fictitious trajectory generated with learned ensembles of transition dynamics $\hat{M}_\theta$ weighted by Uncertainty-Labeler (or uncertainty regularized coefficient).

Behavioral Cloning (line 6). Here, we start by initializing the policy with π_{init} with behavioral cloning weighted by uncertainty regularized coefficient. Specifically, our initial policy $\hat{\pi}_{\text{init}} = \hat{\pi}_\beta$ is parameterized by β will be:

$$\min_\beta \frac{1}{D_{\text{all}}} \frac{1}{2} \sum_{(s_t, a_t) \in D_{\text{all}}} \hat{U}(a_t, s_t) \|a_t - \hat{\pi}_\beta\|_2^2 \quad (11)$$

Here eq. (11) signifies the fact that we want our policy to be cloned from the high confidence region (high data support) by assigning higher weights, leaving our policy to be away from the region of low confidence (higher uncertainty).

Fictitious trajectory generated from learned dynamics (line 7 to line 8). After each deployment, the environment dynamics transitions are estimated by an ensemble

of $\{f_\theta\}_1^M$ trained using D_{all}. To generate a fictitious trajectory, we first randomly sample a model $i \in (1, 2, .., M)$, and then we roll-out the trajectory $\hat{\tau}$ by running the learning policy π with the next state as $\hat{s}_{t+1} = \hat{f}_{\theta_i}(\hat{s}_t, a_t = \pi(\hat{s}_t))$, for $t \in 1, ..., M$.

TRPO training with Uncertainty (line 9 to line 10). Next, we train the policy with uncertainty regularized TRPO. We weight each TRPO update with an uncertainty regularized coefficient from the Uncertainty-Labeler. This time, since the ground-truth next state is unavailable to us (because we are using fictitious trajectory), instead, we compute the uncertainty regularized coefficient as the intra-discrepancy of the next-state prediction error within the models of the Uncertainty-Labeler ($\{\hat{P}_\phi\}_1^K$). To do this, we first randomly sample two models (i, j) from the K Uncertainty-Labeler ensembles. For the generated fictitious trajectory $\hat{\tau}$, we label each $(s_t, a_t) \in \hat{\tau}$ with the approximated uncertainty regularized coefficient as:

$$\widehat{\widehat{U}}(s_t, a_t) = \exp(-\alpha \|\hat{s}_{t+1}^{(i)} - \hat{s}_{t+1}^{(j)}\|_2^2),$$
$$\text{for } (s_t, a_t) \in \hat{\tau} \tag{12}$$

where $s_{t+1}^{(i)} = \hat{P}_\theta^{(i)}(\hat{s}_t, a_t)$ is the next state predicted by the sampled i_{th} model from the Uncertainty-Labeler (and similarly for j index), and $\alpha > 0$ is a temperature parameter.

Utilizing Corollary 1, we use TRPO to optimize the $V_{M^*}^\pi$ by improving its lower-bound (the RHS). Here, we replace $V_{\hat{M}^{(i)}}^\pi(s)$ by $A^{\pi_{\vartheta_k}}(s, a)$, and we approximated $U_{\hat{M}^{(i)}, M^*}(s, a)$ by $\widehat{\widehat{U}}(s, a)$, with TRPO:

$$\underset{\vartheta}{\text{argmax}} \, E_{\pi_\vartheta(s), s, \hat{P}_\theta^{(i,j)}, \hat{f}_\phi^{(i)}} \left\{ \frac{\pi_\vartheta(a|s)}{\pi_{\vartheta_k}(a|s)} A^{\pi_{\vartheta_k}}(s, a) \widehat{\widehat{U}}(a, s) \right\}$$
$$\text{s.t. } E_{a \sim \pi_\vartheta(s), s, \hat{P}_\theta^{(i,j)}, \hat{f}_\phi^{(i)}} \{ D_{KL}(\pi_\vartheta(\cdot|s) \| \pi_{\vartheta_k}(\cdot|s)) \} \leq \delta \tag{13}$$

where we set $\pi_{\vartheta_0} = \pi_{\text{init}}$ as the initial policy of the TRPO at the first iteration. Here, $A^{\pi_{\vartheta_k}}(s, a)$ is the advantage function of policy π_{ϑ_k} following a fictitious trajectory generated by $\hat{f}_\phi$. Here, our optimization process regularizes the update toward the high confidence regions. When the policy bootstraps into out-of-distribution states and actions regions (low support area), we discount the update.

4. Empirical Results

We empirically evaluate our proposed MUSBO algorithm with five continuous control benchmarks using the MuJOCO[1] physics simulator: Walker2d, Hopper, Half-Cheetah, Ant, and Cheetah-Run.

Baseline. We compare our MUSBO with BREMEN, a state-of-the-art algorithm[2] that is designed specifically for

---

[1]http://www.mujoco.org/

[2]BREMEN has been compared against to SAC, Model-

deployment constrained setting. For BREMEN, we used the exact same hyper-parameters as what was originally proposed in Matsushima et al. (2020). As for comparison, we also adapted MOPO (Yu et al., 2020), a pure-offline algorithm, to deployment-constrained setting. We used the latest learned policy for data collection.

Evaluation Setup. Following Matsushima et al. (2020), our evaluation set up consisted of 5 deployments for Half-Cheetah, Ant, Cheetah-Run and 10 deployments for Walker2D and Hopper. To see the trade-off between the number of deployment versus data sample size, we perform our empirical tests on two separate cumulative data sizes of *500k* and *250k*.

For the settings of *500k / (250k)* sample size experiments, the environments Half-Cheetah, Ant and Cheetah-Run have a per deployment data collection batch size $B = 100k/(50k)$ (with 5 deployments in total). For the environments of Walker2d and Hopper, they have a per deployment data collection batch size of $B = 50k/(25k)$ (with 10 deployments in total).

Deployment Setup. In the deployment constrained setting, the data collection only happens during the deployment. No training happens during this period. Only until a data batch size of $|B|$ has been collected, the agent will be taken offline for training and policy update. At the first (initial) deployment, a random policy is used for data collection. After that, the data collection will be replaced by the updated policy and will be launched to deploy again.

We plot our 500k data size results in Fig.2 and Fig.3, and we plot our 250k data size results on Fig.4 and Fig.5. For all experimental results, we averaged over 5 random seeds.

Empirical Details: 500k data size. The top(first) figure of Fig.2 shows the total sample size of each deployment on the y-axis, and on the x-axis, we show the number of deployments. We align along the x-axis for all figures. Our method works the best on the Walker2D and Hopper, in which our MUSBO approach achieves significant cumulative rewards especially in the longer deployments (after 6_{th}). In the Hopper environment, we see that the baselines (BREMEN and MOPO) show incapable of learning a meaningful policy while our MUSBO significantly outperforms. On the other hand, in the Ant, Cheetah-Run and Half-Cheetah environment (as shown in Fig.3), our MUSBO is capable of achieving a higher performance using a smaller amount of data. For instance, in the Cheetah-Run environment, MUSBO already achieved 550 points in the second deployment but the baselines take four to five deployments to achieve the same score.

Ensembles-TRPO, BCQ, and BRAC (all adapted to deployment-constrained setting) and shows state-of-the-art performance.
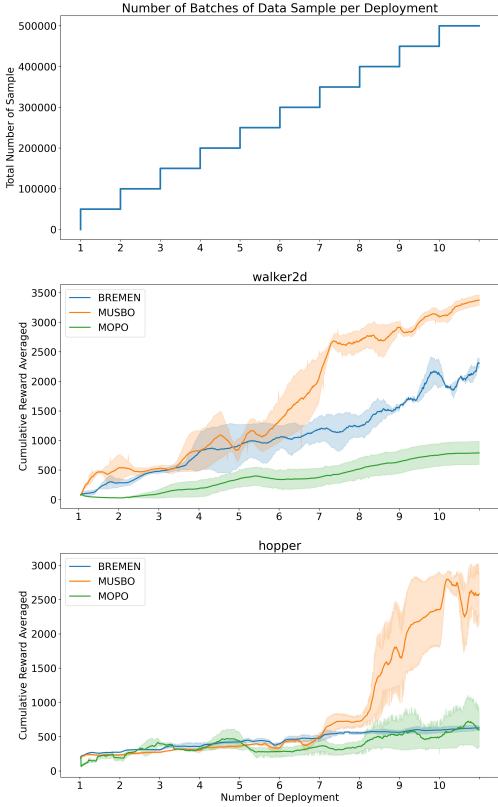
Figure 2. Empirical Evaluations of Walker2D and Hopper for $I \times B = 500k$ settings with 10 deployments and $B = 50k$. The x-axis is aligned showing the number of deployment.

Figure 3. Empirical Evaluations of Ant, Cheetah-Run and Half-Cheetah for $I \times B = 500k$ settings with 5 deployments and $B = 100k$. The x-axis are aligned showing the number of deployment.

Empirical Details: 250k data size. Different than the previous set of experiments, here, we reduce the total data size by half to 250k. In general, when we reduce the data size, the performances of all algorithms will be reduced. Despite this, our MUSBO algorithm performs quite well even in this setting. Comparing to the 500k data size experiment, we also observe similar trends. In Fig.4, we plot the results for Walker2d and Hopper. For the former, we see a significant winning margin in Walker2d whereas in Hopper, the winning only happens in the $8_{th} \sim 9_{th}$ deployments. In Fig.5, we plot the results for Ant, Cheetah-Run and Half-Cheetah. We observe that our MUSBO algorithm achieves faster learning and stronger performance.

5. Ablation Study

In this section, we examine the MUSBO algorithm in terms of the following two aspects: 1) discovery of high quality and novel data batches during each deployment; 2) whether this will lead to more accurate learning of model dynamics. Averaged over three random seeds, we perform two ablation studies as blow. To assess the novelty of data-batches during each deployment, we calculate the cosine similarity score.
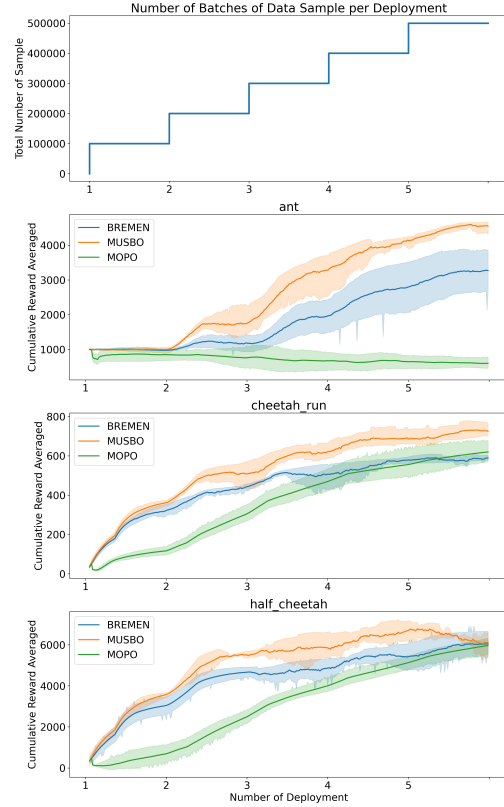
For each deployment, we calculate the cosine score between the current batch ($B^{(i)}$) versus the aggregation of all of the previous batches ($B^1 \bigcup B^2 \bigcup ... B^{(i-1)}$). Our procedure is that for each state observation in the current batch $B^{(i)}$, we measure the aggregated cosine score against each state observation in the previous aggregated batches, averaged over the number of transitions. We repeat this calculation for all state observation for each deployment, and then show the result in Table.1. Note that deployment 1 is not shown because the initial data collection policy is a random policy. Our result shows that our MUSBO algorithm leads to much higher novel transitions discovery. In Fig.6, we compare the performance of the fictitious rollouts from the learned model dynamics versus the rollouts from the real environment. Here, we take the comparison between MUSBO and baseline (BREMEN). We first plot the energy distance in the left hand plot of Fig.6. On the right hand plot, we show the mean square error (MSE) of trajectory-wise rollouts in Fig.6. As is shown in the figure, in the beginning, the performance of both methods is very close to each (due to initial random exploration). As the number of deployment increases, our method is capable of discovering high quality transitions which result in a better estimated model.
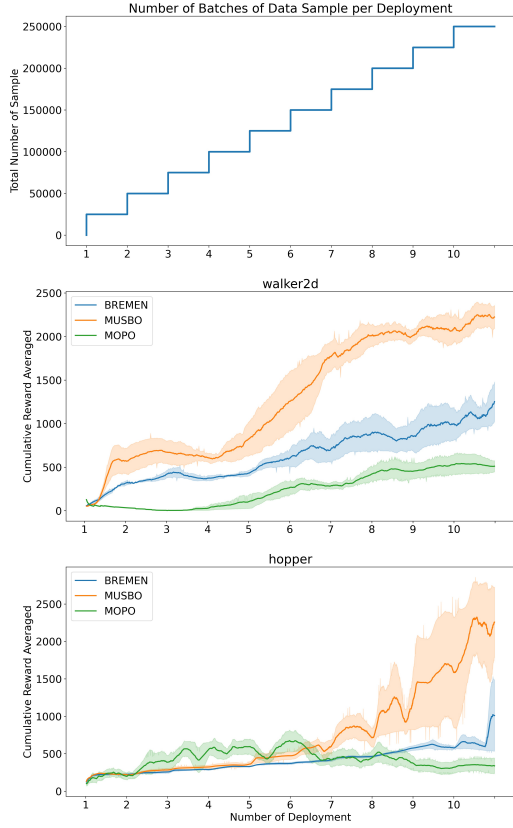
Figure 4. Empirical Evaluations of Walker2D and Hopper for $I \times B = 250k$ settings with 10 deployments and $B = 25k$. The x-axis are aligned showing the number of deployment.

Deployment	2nd	3rd	4th	5th
MUSBO	**0.994**	**0.963**	**0.945**	**0.944**
Baseline(BREMEN)	0.974	0.909	0.894	0.865

Table 1. Table showing the amount of novelty of each data batches collected between each deployment. Novelty is measured as cosine similarity between each observation (state) in $B^{(i)}$ versus the previous aggregated batches $(B^{(1)}.. \bigcup B^{(i-1)})$, averaged over the number of transitions.

6. Conclusion

In this paper, we have proposed the algorithmic framework Model-based Uncertainty Regularized Sample efficient Batch Optimization (MUSBO) for optimizing the policy learning under the deployment-constrained setting. Through theoretical analysis, we first build a lower bound that utilizes uncertainty quantification of state-action pairs in the dataset. Our proposed strategy encourages novel data transition discovery during each deployment. Then, during offline training, we weight each state-action pair with uncertainty measures and discourage the policy update from bootstrapping into unknown (or low support) regions. Our empirical analysis shows that MUSBO outperforms the
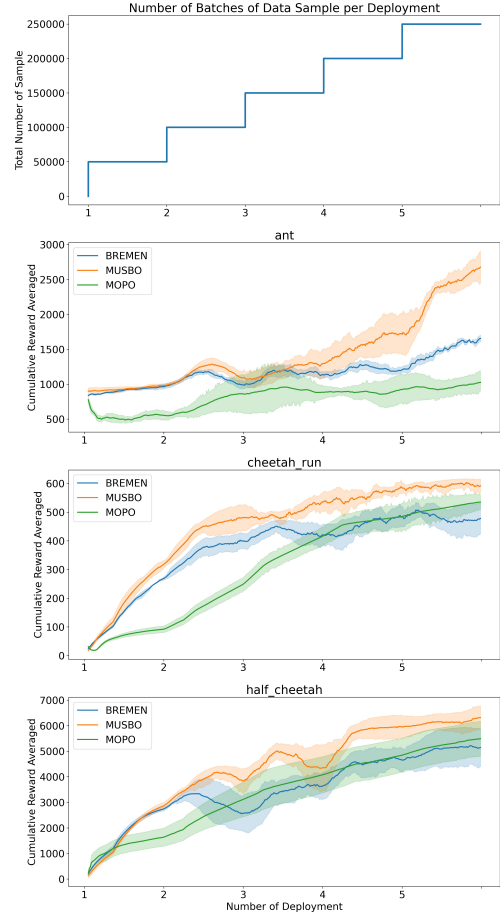

Figure 5. Empirical Evaluations of Ant, Cheetah-Run and Half-Cheetah for $I \times B = 250k$ settings with 5 deployments and $B = 50k$. The x-axis are aligned showing the number of deployment.
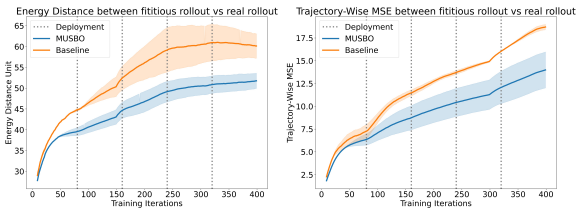

Figure 6. Fictitious rollouts accuracy of our learned model dynamics for the Cheetah-Run environment, comparing between MUSBO and baseline (BREMEN). The vertical dotted lines show the point at which the deployments take place. The left hand plot shows the energy distance between the fictitious rollouts generated from the learned model versus the real rollouts coming from the real environment (the **lower** the curve the better). The plot on the right shows the Trajectory-wise MSE between the fictitious rollouts versus the real fictitious(the **lower** the curve the better). Both plots show that our method MUSBO achieves learning a more accurate model.

state-of-the-art algorithm in the MuJuco benchmarks.

References

Bai, Y., Xie, T., Jiang, N., and Wang, Y.-X. Provably efficient q-learning with low switching cost, 2020.

Chen, L., Pelger, M., and Zhu, J. Deep learning in asset pricing, 2020.

Chu, W.-H. and Kitani, K. M. Neural batch sampling with reinforcement learning for semi-supervised anomaly detection. In Vedaldi, A., Bischof, H., Brox, T., and Frahm, J.-M. (eds.), *Computer Vision – ECCV 2020*, pp. 751–766, Cham, 2020. Springer International Publishing. ISBN 978-3-030-58574-7.

Chua, K., Calandra, R., McAllister, R., and Levine, S. Deep reinforcement learning in a handful of trials using probabilistic dynamics models, 2018.

Ernst, D., Geurts, P., and Wehenkel, L. Tree-based batch mode reinforcement learning. *Journal of Machine Learning Research*, 6:503–556, 04 2005.

Feinberg, V., Wan, A., Stoica, I., Jordan, M. I., Gonzalez, J. E., and Levine, S. Model-based value estimation for efficient model-free reinforcement learning, 2018.

Fujimoto, S., Meger, D., and Precup, D. Off-policy deep reinforcement learning without exploration, 2019.

Ghasemipour, S. K. S., Schuurmans, D., and Gu, S. S. Emaq: Expected-max q-learning operator for simple yet effective offline and online rl, 2021.

Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor, 2018.

Jaakkola, T., Singh, S., and Jordan, M. Reinforcement learning algorithm for partially observable markov decision problems. *Advances in Neural Information Processing Systems*, 7, 11 1999.

Janner, M., Fu, J., Zhang, M., and Levine, S. When to trust your model: Model-based policy optimization, 2019.

Kaiser, L., Babaeizadeh, M., Milos, P., Osinski, B., Campbell, R. H., Czechowski, K., Erhan, D., Finn, C., Kozakowski, P., Levine, S., Mohiuddin, A., Sepassi, R., Tucker, G., and Michalewski, H. Model-based reinforcement learning for atari, 2020.

Kidambi, R., Rajeswaran, A., Netrapalli, P., and Joachims, T. Morel : Model-based offline reinforcement learning, 2020.

Kumar, A., Fu, J., Tucker, G., and Levine, S. Stabilizing off-policy q-learning via bootstrapping error reduction, 2019.

Kurutach, T., Clavera, I., Duan, Y., Tamar, A., and Abbeel, P. Model-ensemble trust-region policy optimization, 2018a.

Kurutach, T., Clavera, I., Duan, Y., Tamar, A., and Abbeel, P. Model-ensemble trust-region policy optimization, 2018b.

Lange, S., Gabel, T., and Riedmiller, M. Batch reinforcement learning. *Reinforcement Learning: State of the Art*, 01 2012. doi: 10.1007/978-3-642-27645-3_2.

Langley, P. Crafting papers on machine learning. In Langley, P. (ed.), *Proceedings of the 17th International Conference on Machine Learning (ICML 2000)*, pp. 1207–1216, Stanford, CA, 2000. Morgan Kaufmann.

Levine, S., Kumar, A., Tucker, G., and Fu, J. Offline reinforcement learning: Tutorial, review, and perspectives on open problems, 2020.

Luo, Y., Xu, H., Li, Y., Tian, Y., Darrell, T., and Ma, T. Algorithmic framework for model-based deep reinforcement learning with theoretical guarantees, 2019.

Matsushima, T., Furuta, H., Matsuo, Y., Nachum, O., and Gu, S. Deployment-efficient reinforcement learning via model-based offline optimization, 2020.

Mishra, N., Abbeel, P., and Mordatch, I. Prediction and control with temporal segment models, 2017.

Munos, R. and Szepesvári, C. Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research*, 9(27):815–857, 2008. URL http://jmlr.org/papers/v9/munos08a.html.

Nachum, O., Chow, Y., Dai, B., and Li, L. Dualdice: Behavior-agnostic estimation of discounted stationary distribution corrections, 2019a.

Nachum, O., Dai, B., Kostrikov, I., Chow, Y., Li, L., and Schuurmans, D. Algaedice: Policy gradient from arbitrary experience, 2019b.

Nair, A., Dalal, M., Gupta, A., and Levine, S. Accelerating online reinforcement learning with offline datasets, 2020.

OpenAI, :, Berner, C., Brockman, G., Chan, B., Cheung, V., Debiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., Jozefowicz, R., Gray, S., Olsson, C., Pachocki, J., Petrov, M., de Oliveira Pinto, H. P., Raiman, J., Salimans, T., Schlatter, J., Schneider, J., Sidor, S., Sutskever, I., Tang, J., Wolski, F., and Zhang, S. Dota 2 with large scale deep reinforcement learning, 2019a.

OpenAI, Akkaya, I., Andrychowicz, M., Chociej, M., Litwin, M., McGrew, B., Petron, A., Paino, A., Plappert, M., Powell, G., Ribas, R., Schneider, J., Tezak, N., Tworek, J., Welinder, P., Weng, L., Yuan, Q., Zaremba, W., and Zhang, L. Solving rubik's cube with a robot hand, 2019b.

Peska, L. and Vojtas, P. Off-line vs. on-line evaluation of recommender systems in small e-commerce. In *Proceedings of the 31st ACM Conference on Hypertext and Social Media*, HT '20, pp. 291–300, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450370981. doi: 10.1145/3372923.3404781. URL https://doi.org/10.1145/3372923.3404781.

Schulman, J., Levine, S., Moritz, P., Jordan, M. I., and Abbeel, P. Trust region policy optimization, 2017.

Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T., Lillicrap, T., Simonyan, K., and Hassabis, D. Mastering chess and shogi by self-play with a general reinforcement learning algorithm, 2017.

Veerapaneni, R., Co-Reyes, J. D., Chang, M., Janner, M., Finn, C., Wu, J., Tenenbaum, J. B., and Levine, S. Entity abstraction in visual model-based reinforcement learning, 2020.

Vinyals, O., Babuschkin, I., Czarnecki, W., Mathieu, M., Dudzik, A., Chung, J., Choi, D., Powell, R., Ewalds, T., Georgiev, P., Oh, J., Horgan, D., Kroiss, M., Danihelka, I., Huang, A., Sifre, L., Cai, T., Agapiou, J. P., Jaderberg, M., Vezhnevets, A. S., Leblond, R., Pohlen, T., Dalibard, V., Budden, D., Sulsky, Y., Molloy, J., Paine, T. L., Gulcehre, C., Wang, Z., Pfaff, T., Wu, Y., Ring, R., Yogatama, D., Wünsch, D., McKinney, K., Smith, O., Schaul, T., Lillicrap, T., Kavukcuoglu, K., Hassabis, D., Apps, C., and Silver, D. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, pp. 1–5, 2019.

Wang, T., Bao, X., Clavera, I., Hoang, J., Wen, Y., Langlois, E., Zhang, S., Zhang, G., Abbeel, P., and Ba, J. Benchmarking model-based reinforcement learning, 2019.

Wu, Y., Tucker, G., and Nachum, O. Behavior regularized offline reinforcement learning, 2019.

Yu, T., Thomas, G., Yu, L., Ermon, S., Zou, J., Levine, S., Finn, C., and Ma, T. Mopo: Model-based offline policy optimization, 2020.

Zhang, M., Vikram, S., Smith, L., Abbeel, P., Johnson, M. J., and Levine, S. Solar: Deep structured representations for model-based reinforcement learning, 2019.

Zhang, R., Dai, B., Li, L., and Schuurmans, D. Gendice: Generalized offline estimation of stationary values, 2020.

A. Proof of Theorem 1

Proof. Following Theorem 3.1 of (Luo et al., 2019), we see: that for each deployment k:

$$V_{M^*}^{\pi_{k+1}} \geq V_{M_{k+1}}^{\pi_{k+1}} u_{\pi_k,\delta}(M_{k+1}, \pi_{k+1}) \text{(by Assumption A1)}$$

$$\begin{aligned} V_{M^*}^{\pi_k} &= V_{M^*}^{\pi_k} u_{\pi_k,\delta}(M^*, \pi_k) \\ &\leq V_{M_{k+1}}^{\pi_{k+1}} u_{\pi_k,\delta}(M_{k+1}, \pi_{k+1}) \text{(by Assumption A2)} \end{aligned}$$

Thus,

$$V_{M^*}^{\pi_{k+1}} \geq V_{M^*}^{\pi_k} \geq ...$$

$\square$

B. Proof of Proposition 1

Following Lemma 4.2 of (Luo et al., 2019), we let W_j be the expected return when executing $\hat{M}^{(i+1)}$ for the first j steps, and then switch to $\hat{M}^{(i)}$ for the remaining steps.

Proof.

$$W_j = \mathop{\mathbf{E}}_{\substack{a_t \sim \pi(s_t) \\ \forall j > t \geq 0, s_{t+1} \sim \hat{M}^{(i+1)}(\cdot|s_t,a_t) \\ \forall t \geq j, s_{t+1} \sim \hat{M}^{(i)}(\cdot|s_t,a_t)}} \left\{ \sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | s_0 = s \right\} \tag{14}$$

Thus, we have $W_0 = V_{\hat{M}^{(i)}}^{\pi}$, and $W_{\infty} = V_{\hat{M}^{(i+1)}}^{\pi}$. Next, we write:

$$V_{\hat{M}^{(i+1)}}^{\pi} - V_{\hat{M}^{(i)}}^{\pi} = \sum_{j=0}^{\infty} (W_{j+1} - W_j) \tag{15}$$

We expand W_j and W_{j+1} so that we can cancel the shared terms:

$$W_j = R_j + \mathop{\mathbf{E}}_{s_j,a_j \sim \pi, \hat{M}^{(i+1)}} \left\{ \mathop{\mathbf{E}}_{s_{j+1} \sim \hat{M}^{(i)}(\cdot|s_t,a_t)} \left\{ \gamma^{j+1} V_{\hat{M}^{(i)}}^{\pi}(s_{j+1}) \right\} \right\}$$

$$W_{j+1} = R_j + \mathop{\mathbf{E}}_{s_j,a_j \sim \pi, \hat{M}^{(i+1)}} \left\{ \mathop{\mathbf{E}}_{s_{j+1} \sim \hat{M}^{(i+1)}(\cdot|s_t,a_t)} \left\{ \gamma^{j+1} V_{\hat{M}^{(i)}}^{\pi}(s_{j+1}) \right\} \right\}$$

where R_j is the expected return of the first j time step. Next, we cancel the share terms so that:

$$W_{j+1} - W_j = \gamma^{j+1} \mathop{\mathbf{E}}_{s_j,a_j \sim \pi, \hat{M}^{(i+1)}} \Big\{ \\ \mathop{\mathbf{E}}_{s' \sim \hat{M}^{(i+1)}(\cdot|s_j,a_j)} \{V_{\hat{M}^{(i)}}^{\pi}(s')\} - \mathop{\mathbf{E}}_{s' \sim \hat{M}^{(i)}(\cdot|s_j,a_j)} \{V_{\hat{M}^{(i)}}^{\pi}(s')\} \Big\} \tag{16}$$

Thus, based on eq. (16), we have:

$$V_{\hat{M}^{(i+1)}}^{\pi} - V_{\hat{M}^{(i)}}^{\pi} = \kappa \mathop{\mathbf{E}}_{s_j,a_j \sim \pi, \hat{M}^{(i+1)}} \Big\{ \\ \mathop{\mathbf{E}}_{s' \sim \hat{M}^{(i+1)}(\cdot|s_j,a_j)} \{V_{\hat{M}^{(i)}}^{\pi}(s')\} - \mathop{\mathbf{E}}_{s' \sim \hat{M}^{(i)}(\cdot|s_j,a_j)} \{V_{\hat{M}^{(i)}}^{\pi}(s')\} \Big\} \tag{17}$$

Multiply eq. (17) by -1 on both side and substitute in eq. (4), thus we have

$$V_{\hat{M}^{(i+1)}}^{\pi} \geq V_{\hat{M}^{(i)}}^{\pi} - \\ \kappa \mathop{\mathbf{E}}_{(s,a) \sim \rho_{\hat{M}^{(i+1)}}^{\pi}} \{d_{\hat{M}^{(i)},\hat{M}^{(i+1)}}(s,a)\}. \tag{18}$$

$\square$

C. Proof of Corollary 1

Proof. We apply the definition of $U_{\hat{M}^{(i+1)},M^*}(s,a)$ from eq. (6) into Proposition 1, and this completes the proof. $\square$

D. Characterizing the improvement gap between each deployment

Adapted from Luo et al. (2019), next, we provide a bound on the difference of the value functions at each deployment,

Proposition 2. *Assume* $V_{\hat{M}}^{\pi}$ *is L-Lipschitz such that* $|V_{\hat{M}}^{\pi}(s) - V_{\hat{M}}^{\pi}(s')| \leq L\|s - s'\|, \forall s, s'$, *then:*

$$|V_{\hat{M}^{(i+1)}}^{\pi} - V_{\hat{M}^{(i)}}^{\pi}| \leq \kappa L \\ \mathop{\mathbf{E}}_{(s,a) \sim \rho_{\hat{M}^{(i+1)}}^{\pi}} (\|\hat{M}^{(i+1)}(s,a) - \hat{M}^{(i)}(s,a)\|) \tag{19}$$

Here, Proposition 2 says that the improvement between each value function is upper bounded by the improvement of estimated model dynamics. As the number of deployment increases, we get more training data, and thus our model estimation is getting better and better. For the proof, it is basically the same step as Luo et al. (2019).

Proof. Since we have assumed the Lipschitzeness of $V_{\hat{M}}^{\pi}$, we can rewrite the RHS of eq. (17) that:

$$\mathop{\mathbf{E}}_{s' \sim \hat{M}^{(i+1)}(\cdot|s,a)} \{V_{\hat{M}^{(i)}}^{\pi}(s')\} - \mathop{\mathbf{E}}_{s' \sim \hat{M}^{(i)}(\cdot|s,a)} \{V_{\hat{M}^{(i)}}^{\pi}(s')\} \\ \leq L|\hat{M}^{(i+1)}(s,a) - \hat{M}^{(i)}(s,a)| \tag{20}$$

We then combine eq. (20) and eq. (17) with triangle inequality, we get:

$$|V_{\hat{M}^{(i+1)}}^{\pi} - V_{\hat{M}^{(i)}}^{\pi}| \leq \kappa L \\ \mathop{\mathbf{E}}_{(s,a) \sim \rho_{\hat{M}^{(i+1)}}^{\pi}} (\|\hat{M}^{(i+1)}(s,a) - \hat{M}^{(i)}(s,a)\|) \tag{21}$$

$\square$

E. Exploration to Collect Data from Low Support Region

During each deployment, we want to collect the data from the low support (or un-visited) regions. we make use of the uncertainty labeler to guide exploration to the un-visited regions for novel data discovery. This is achieved by injecting the $\hat{U}(a, s)$ as an exploration noise with the zero-mean normal distribution: $\mathcal{N}(0, \sigma = \hat{U}(a, s))$. In the Ablation Study (Section.5), we show that this exploration strategy leads to novel data discovery, and also contribute to better learned model dynamics.

Specifically, the action will be parameterized by a stochastic Gaussian policy (with parameter μ_ϑ) as:

$$a_t = \tanh(\mu_\vartheta(s_t)) + \epsilon_{\text{const}} + \epsilon_{\hat{U}} \qquad (22)$$

where ϵ_{const} and $\epsilon_{\hat{U}}$ are:

$$\epsilon_{\text{const}} \sim \mathcal{N}(\mu = 0, \sigma = 0.01)$$

$$\epsilon_{\hat{U}} \sim \mathcal{N}\left(\mu = 0, \sigma = \hat{U}\left(a = \tanh(\mu_\vartheta(s_t)) + \epsilon_{\text{const}}, s = s_t\right)\right)$$

Here, ϵ_{const} is an additive noise with a constant variance of 0.01, and on top of this, we also added another Gaussian noise with variance equal to $\hat{U}(a, s)$ to guide exploration.

F. Detailed Implementation

We used ADAM as the optimizer with a learning rate of 1e-3 for the model dynamics $\hat{T}$ with ensembles size of $N = 5$. For the uncertainty-labeler, we use ensembles of PNN with $K = 3$ and a learning rate of 1e-3. For the behavioral cloning, we used a learning rate of 5e-4. For all the collected data, we divided them into 85% for training, and 15% for validation (for model validation). The $\hat{T}, \hat{P}$ are trained with early stopping when their performance on the validation set no longer improves after consecutive 3 episodes. We used this same setting for the behavioral cloning as well.

Model Architecture For our policy network, we parameterized it by two layers of fully-connected neural network with hidden units of 200. For the $\hat{T}$ model dynamics, we parameterized it by two layers of fully-connected neural network with hidden units of 1024. We used this same configuration with $\hat{P}$ uncertainty-labeler and implemented with PNN.

Training Time The overall training time differs for each environment. We train all models on Nvidia T4 GPU. For the $500k$ data size experiment, the entire training duration (1 run) for Walker2D and Hopper environments are 18 hours and 26 hours respectively. For the Half-Cheetah, Ant, and Cheetah-Run environment, it is a a lot more faster. It takes 7 hours, 12 hours, and 8 hours per run, respectively. For the $250k$ data size experiments, the training time is about 25 minutes faster than the $500k$ experiments.

	L	Rollouts Length	TRPO's δ
Ant	2,000	250	0.05
HalfCheetah	2,000	250	0.1
Hopper	6,000	1,000	0.05
Walker2d	2,000	1,000	0.05
CheetahRun	2,000	250	0.05

Table 2. Hyper-parameters for MUSBO Algorithm

Since we are applying Dyna-style update with neural network based dynamics models, following Wang et al. (2019) Matsushima et al. (2020), we used the following reward functions for our dynamics models (for model-based training only) as:

- Walker2d, Hopper: $\dot{x}_t - 0.001\|a_t\|_2^2 + 1$

- CheetahRun: $\max(0, \min(\dot{x}_t/10, 1))$

- Ant: $\dot{x}_t - 0.1\|a_t\|_2^2 - 3.0(z_t - 0.57)^2 + 1$

- HalfCheetah: $\dot{x}_t - 0.1\|a_t\|_2^2$

We enabled termination in rollouts for the Hopper and Walker2D environments, and disabled that of the Ant, HalfCheetah, and CheetahRun environments (with a maximum step of 1000 for each episode). For the CheetahRun task, we adopt it from the DM control suit[3].

Other Hyper-parameters We used the same $\alpha = 0.028$ (the temperature parameter for eq.(12)) for all environments. Similarly, for the action parameterization eq. (22), we used the same constant $\sigma = 0.01$ (variance term of ϵ_{const}) for all environments. The σ of ϵ_{const} is searched over the set of $\{0.01, 0.05, 0.1\}$.

We searched the rollout length on $\{250, 1000\}$, and the δ on $\{0.01, 0.05, 0.1\}$. We summarized these three parameters (L, Rollouts Length, δ) as above in table 2.

For discount factor γ and GAE λ, we used the same set of hyperparmaters as in Wang et al. (2019). Specifically, we used the same $\gamma = 0.99$ for all environment. Also, we used GAE $\lambda = 0.95$ for all environment except for Ant which has a GAE $\lambda = 0.97$.

Hyper-parameters for Baseline. For the BREMEN, we used the exact parameters as (Matsushima et al., 2020). For MOPO, we adapted it to the deployment setting. We used the latest learned policy for deployment, and then launched it to collect data batch of size $|B|$. For the CheetahRun task, we used the same set of parameters of HalfCheetah. On all environments, we tried hyper-parameters search on the ranges as originally proposed by the paper Yu et al.

---

[3]https://github.com/deepmind/dm_control

(2020), we didn't find any improvement over the same set of parameters as originally proposed. Thus, we used the same set of parameters as originally proposed.