
Bayesian Optimistic Optimization: Optimistic Exploration for Model-based Reinforcement Learning

Chenyang Wu¹, Tianci Li¹, Zongzhang Zhang^{1*}, Yang Yu^{1, 2}

¹National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing, China

²Pazhou Lab, Guangzhou, China

{wucy, litc}@lamda.nju.edu.cn, {zzzhang, yuy}@nju.edu.cn

Abstract

Reinforcement learning (RL) is a general framework for modeling sequential decision making problems, at the core of which lies the dilemma of exploitation and exploration. An agent failing to explore systematically will inevitably fail to learn efficiently. Optimism in the face of uncertainty (OFU) is a conventionally successful strategy for efficient exploration. An agent following the OFU principle explores actively and efficiently. However, when applied to model-based RL, it involves specifying a confidence set of the underlying model and solving a series of nonlinear constrained optimization, which can be computationally intractable. This paper proposes an algorithm, Bayesian optimistic optimization (BOO), which adopts a dynamic weighting technique for enforcing the constraint rather than explicitly solving a constrained optimization problem. BOO is a general algorithm proved to be sample-efficient for models in a finite-dimensional reproducing kernel Hilbert space. We also develop techniques for effective optimization and show through some simulation experiments that BOO is competitive with the existing algorithms.

1 Introduction

Reinforcement learning (RL) is a sequential decision-making problem in which an agent acts in an unknown environment while maximizing the cumulative rewards it receives [1, 2]. In this paper, we consider the RL in Markov decision processes (MDPs), where the agent observes the state of the environment at each timestep and makes decisions accordingly. Since the environment is unknown, maximizing the cumulative rewards naturally involves a trade-off between exploration and exploitation. Exploitation is to make the best-rewarding decision based on the agent’s information, and exploration means actively gathering information about the environment so that the agent understands better about the environment and thus makes better decisions in the future. An algorithm cannot be sample-efficient without balancing them properly.

Theoretically, the exploration and exploitation dilemma admits a Bayesian optimal solution [3]. That is to consider the RL problem a so-called Bayes-Adaptive MDP (BAMDP), a special case of partially observable Markov decision processes (POMDPs), where the parameter of the dynamics is unobservable. Although this formulation provides useful insights, the POMDP formulation is computationally intractable [4–6]. Therefore, all practical algorithms [7–10] resolve this dilemma by achieving a delicate balance between seeking rewards and gathering information. Optimism in the face of uncertainty (OFU) is one of the conventionally successful approaches for this balance and is established as an efficient learning principle in various cases [11–13].

*Corresponding Author

The OFU principle makes the agent an optimist who is always optimistic about the uncertainty of the environment. If it ever stands a chance that a policy is highly profitable, the agent will try it out. By executing an optimistic policy, the agent either receives high rewards or gains information after observing unexpected outcomes. OFU involves constructing a confidence set of possible MDPs and solving for the most optimistic one within the confidence set. Unfortunately, this is typically a constrained nonlinear non-concave optimization and is impractical to solve [14, 15]. Therefore, a model-based OFU algorithm is previously deemed an inferior choice [15].

In this paper, we propose a model-based OFU algorithm, Bayesian optimistic optimization (BOO), which enforces the constraint via a dynamic weighting technique and therefore makes the optimization more practical. We show that BOO is a general-purpose model-based OFU RL algorithm that is both provably sample-efficient and comparatively computationally tractable. Our contribution is three-fold:

- We extend the OFU algorithms to the Bayesian setting by requiring the model to reside in a Bayesian credible region. BOO is then proposed as the Lagrangian relaxation of this constrained optimization.
- We prove that BOO is efficient in terms of the frequentist regret for a finite-dimensional reproducing kernel Hilbert space (RKHS).
- We derive optimization methods for BOO and demonstrate empirical evidence that BOO is competitive with UCRL2 [16] and PSRL [17].

2 Preliminaries

This section introduces the preliminaries, including RL in MDPs and the principle of optimism in the face of uncertainty.

2.1 Reinforcement Learning in Markov Decision Processes

We consider the problem where a reward-seeking agent repeatedly interacts with a finite-horizon MDPs $M = (\mathcal{S}, \mathcal{A}, P, R, H, s_1)$, where \mathcal{S} and \mathcal{A} are the state and action spaces, respectively. The algorithm adopted by the agent is denoted as \mathfrak{A} . At each episode, the agent is spawned at the initial state $s_1 \in \mathcal{S}$. It takes an action $a_h \in \mathcal{A}$ at each period h within an episode and receives a random reward, $r_h \sim R(s_h, a_h)$. Then, the environmental state transitions to the next state, $s_{h+1} \sim P(s_h, a_h)$. This process repeats until the episode ends at the H -th period, and another episode begins. Note that the fixed initial state is not a stringent condition since starting from an initial distribution ρ_0 is equivalent to starting from a special state s_1 such that $P(s|s_1, a) = \rho_0(s)$ holds for any state s and action a . A finite MDP is an MDP with finite state and action spaces.

A policy π is a function mapping a state and a period to an action distribution. The value function for the MDP M , policy π , and period h is defined recursively as

$$\begin{aligned} V_h^{\pi, M}(s_h) &= \mathbb{E}_{a_h \sim \pi(s_h, h)} \left[Q_h^{\pi, M}(s_h, a_h) \right], \forall h \in [H], \\ Q_h^{\pi, M}(s_h, a_h) &= \bar{R}^M(s_h, a_h) + \mathbb{E}_{s_{h+1} \sim P^M(s_h, a_h)} \left[V_{h+1}^{\pi, M}(s_{h+1}) \right], \forall h \in [H-1], \end{aligned} \quad (1)$$

where $\bar{R}^M(s, a) = \mathbb{E}_{r \sim R^M(s, a)}[r]$, and $Q_H^{\pi, M}(s_H, a_H) = \bar{R}^M(s_H, a_H)$.

We use $X_{k, h}$ to represent the variable X at the h -th period of the k -th episode. For notational convenience, $X_{k, h}$ is sometimes abbreviated as X_{kh} . The history prior to k -th episode is defined as $\mathcal{H}_k = (s_{1,1}, a_{1,1}, r_{1,1}, s_{1,2}, \dots, s_{k-1, H}, a_{k-1, H}, r_{k-1, H})$. The agent is assumed to be capable of memorizing the entire history. In this paper, we consider model-based RL algorithms which produce a model M_k per episode in light of the history \mathcal{H}_k and derive its corresponding optimal policy $\pi_k \in \arg \max_{\pi} V_1^{\pi, M_k}(s_1)$ for execution. We use the terms model and MDP interchangeably. For any history \mathcal{H} , $\mathfrak{A}(\mathcal{H})$ defines a distribution over models and policies. Within the k -th episode, the agent samples an action from $\pi_k(s_{kh}, h)$ at each period h .

In the frequentist viewpoint, there exists an unknown true MDP M^* . We abbreviate $V_h^{\pi^*, M^*}$ as V_h^* , where π^* is an optimal policy of the true MDP M^* . The frequentist performance metric, regret, is

defined in terms of the true MDP:

$$\text{Regret}(T, \mathfrak{A}, M^*) = \mathbb{E}_{\mathcal{H}_{K+1} \sim \mathfrak{A}, M^*} \left[\sum_{k=1}^K \Delta_k \right], \quad (2)$$

where Δ_k is defined as $\mathbb{E}_{\pi_k \sim \mathfrak{A}(\mathcal{H}_k)} \left[V_1^*(s_1) - V_1^{\pi_k, M^*}(s_1) \right]$, K is the total number of episodes, $T = KH$ is the number of the total time steps, and $\mathcal{H}_{K+1} \sim \mathfrak{A}, M^*$ means that the history is sampled by the interaction of the algorithm \mathfrak{A} and the real MDP M^* . In the Bayesian viewpoint, the unknown MDP M^* is treated as a random variable and assigned a prior ρ_M . All MDPs in the support of ρ_M differ only in the transition function P and the reward function R . The Bayesian objective of the agent is to minimize the Bayesian regret up to time T , $\text{BayesRegret}(T, \mathfrak{A}, \rho_M) = \mathbb{E}_{M^* \sim \rho_M} [\text{Regret}(T, \mathfrak{A}, M^*)]$.

2.2 Optimism in the Face of Uncertainty

Optimism in the face of uncertainty is a strategy for information gathering. When the optimal action is not clear given the current information, it is preferable to hazard an optimistic guess. If we make an atrocious guess, we effectively rule that out and pick another next time. Otherwise, we end up finding a competitive solution that incurs little regrets. This idea is mathematically realized as a constrained optimistic optimization, $\max_{\pi_k, M_k} V_1^k(s_1)$ s.t. $M_k \in \mathcal{M}_k$, where \mathcal{M}_k is a confidence set constructed using empirical data \mathcal{H}_k such that $M^* \in \mathcal{M}_k$ with high probability. The pseudocode of the OFU algorithm is shown in Algorithm 1.

Algorithm 1 OFU RL

- 1: **for** episode $k = 1, 2, \dots$ **do**
 - 2: Construct a confidence set \mathcal{M}_k with \mathcal{H}_k
 - 3: Compute $\pi_k \in \arg \max_{\pi} \max_{M_k} V_1^{\pi, M_k}(s_1)$ s.t. $M_k \in \mathcal{M}_k$
 - 4: Execute π_k for an episode
-

3 Related Work

One line of research resolves the exploration and exploitation dilemma by formalizing the RL problem as a planning problem in Bayes-Adaptive MDPs (BAMDPs) [3], which treats the unknown MDP parameter as an additional hidden variable and maintains a belief distribution of the parameter. This line of work shares the scalability problem, which is caused by the exponential increase of possible histories w.r.t. the planning horizon. That is, the planning in BAMDPs is PSPACE-complete [4] and requires exponential time to solve. Intuitively, this is because methods based on BAMDPs deal with the belief distributions of MDPs (or histories) rather than a single MDP. Efficient planning algorithms in BAMDPs do exist. Exploiting the root sampling technique, we can implement an algorithm that gets rid of the posterior distribution and only requires posterior sampling [18]. Nonetheless, the underlying scalability issue remains prominent.

Methods following the OFU principle construct an optimistic estimate of the unknown MDP and execute its optimal policy. Since these methods plan on a single MDP estimate, they are computationally preferable compared to methods based on BAMDPs. The UCRL2 [16] is one such method. However, unlike tabular and linear MDPs for which constructing an optimistic estimate is analytically tractable [11, 13], constructing an optimistic estimate for general MDPs involves a constrained joint optimization of model and policy and is computationally prohibitive. Hence, previous model-based OFU algorithms [19, 20] for general model classes cannot be implemented and rely on posterior sampling for exploration.

The posterior sampling for reinforcement learning (PSRL) works by selecting a random MDP from the posterior distribution and executing its optimal policy [21]. This strategy ensures that a policy is selected according to the probability that it is the optimal policy of the real model. It is shown that PSRL is at least as good as any frequentist OFU algorithm in terms of Bayesian regret [22, 15]. PSRL is argued to be better than optimism since it is more computationally tractable [15]. Nevertheless, as opposed to OFU algorithms, incremental implementation of PSRL is challenging because it requires replanning after each sampling. Practical implementation instead tries to directly sample from the

posterior distribution of the optimal value function [23–25]. However, the theoretical guarantee for these methods is only established for tabular MDPs [25].

H-UCRL was devoted to resolving the intractability of model-based optimistic exploration for general models [26]. It proposes to convert the joint optimization of model and policy into a hallucinated control problem. This approach ignores the correlation between state-action pairs and treats them separately, causing inefficiency as reflected by the extra dependency on the cumulative posterior variance in their regret bound. Our method, BOO, also tries to attack the intractability of optimism, which builds an optimistic model by optimizing both the value and the log-posterior density. It avoids the defect of H-UCRL since maximizing the log-posterior density naturally enforces the correlation between state-action pairs.

We note that, when the prior is uniform, this idea is equivalent to balance value versus log-likelihood, which is first explored in [27] and is named as reward-biased maximum likelihood estimation (RBMLE). They have applied this approach to multi-armed bandits [28], contextual bandits [29], and RL where the model belongs to a known finite set [27, 30]. A constrained version of RBMLE is also successfully applied to RL of linear quadratic control systems [31]. Our algorithm can be considered generalizing RBMLE to a Bayesian perspective and finite dimensional RKHS. In this regard, BOO could also be referred to as reward-biased maximum a posteriori.

Concerning the regret analysis, previous regret analysis for general model-based RL [19, 20] relies on the fact that the constructed model is the most optimistic one in the confidence set. The regret analysis of BOO differs with them significantly since the model constructed by BOO may not belong to a confidence set or a credible region. This difference entails a distinct analysis, where we show that neither the large deviation from the real model nor the possibly pessimistic estimation of the model causes a large regret.

4 Bayesian Optimistic Optimization

In this section, we derive the learning objective of Bayesian optimistic optimization (BOO) as a Lagrangian relaxation of a constrained optimization problem and give an intuitive interpretation of the resulting objective. Assuming the model class resides in a finite-dimensional RKHS, we show that BOO enjoys $\tilde{O}(\sqrt{K})$ regret.

4.1 Constrained BOO

The conventional OFU algorithm, as demonstrated in Algorithm 1, contains a constrained optimistic optimization, where we look for an optimistic MDP $M_k \in \mathcal{M}_k$ and its corresponding optimal policy such that the value is maximized. The constrained BOO is almost the same (see Algorithm 2), except the frequentist confidence set is now replaced with the Bayesian credible region. A credible region with a $1 - \alpha_k$ level of confidence is a set \mathcal{M}_k such that $\Pr(\mathcal{M}_k | \mathcal{H}_k) \geq 1 - \alpha_k$, where $\Pr(\cdot | \mathcal{H}_k)$ is the posterior distribution given history \mathcal{H}_k , and $\Pr(\mathcal{M}_k | \mathcal{H}_k) = \int_{M_k \in \mathcal{M}_k} \Pr(M_k | \mathcal{H}_k) dM_k$.

By the construction of the credible region, we have $M^* \in \mathcal{M}_k$ holds with probability $1 - \alpha_k$ given any history \mathcal{H}_k . Therefore, the per-episode Bayesian regret is bounded with probability $1 - \alpha_k$,

$$\mathbb{E}[\Delta_k] \leq \underbrace{\mathbb{E}[V_1^*(s_1) - V_1^k(s_1) | M^* \in \mathcal{M}_k]}_{\tilde{\Delta}_k^{\text{opt}}} + \underbrace{\mathbb{E}[V_1^k(s_1) - V_1^{\pi_k, M^*}(s_1) | M^* \in \mathcal{M}_k]}_{\tilde{\Delta}_k^{\text{conc}}} \leq \tilde{\Delta}_k^{\text{conc}}, \quad (3)$$

where $V_h^k = V_h^{\pi_k, M_k}$, and the optimism term $\tilde{\Delta}_k^{\text{opt}}$ is less than or equal to 0 by construction. We define a distance metric $d(M_1, M_2) = \max_{\pi} |V_1^{\pi, M_1}(s_1) - V_1^{\pi, M_2}(s_1)|$. It is preferable to have a credible region \mathcal{M} such that the set width $\max_{M_1, M_2 \in \mathcal{M}} d(M_1, M_2)$ is minimized since the set width certifies an upper bound on $\tilde{\Delta}_k^{\text{conc}}$. However, designing a value concentration credible region could be intractable for generic model classes. As a reasonable alternative, we propose the highest density region (HDR) [32] or, in the Bayesian context, the highest posterior density (HPD) region [33], i.e., $\mathcal{M}_k = \{M_k | \Pr(M_k | \mathcal{H}_k) \geq \epsilon_k\}$, where ϵ_k is the largest constant such that $\Pr(\mathcal{M}_k | \mathcal{H}_k) \geq 1 - \alpha_k$. This kind of region features a desirable property that it occupies the smallest volume in the sample space among all credible regions of the same confidence level and has a potentially small set width.

There are two problems preventing the constrained BOO from being a practical algorithm. The first is that the ϵ_k in the definition of the HPD region is unknown. Although we may approximate it with the Monte Carlo approximation [34], an estimator of high/low quantiles has high variance rendering the approximation difficult. The other difficulty is that the constrained joint optimization of model and policy is an NP-hard problem even in bandits with linear reward and quadratic constraints [14].

Algorithm 2 Constrained BOO

- 1: **for** episode $k = 1, 2, \dots$ **do**
 - 2: Construct a **credible region** \mathcal{M}_k with \mathcal{H}_k
 - 3: Compute $\pi_k \in \arg \max_{\pi} \max_{M_k} V_1^{\pi, M_k}(s_1)$ s.t. $M_k \in \mathcal{M}_k$
 - 4: Execute π_k for an episode
-

4.2 BOO as Lagrangian Relaxation of Constrained BOO

By introducing a Lagrange multiplier λ_k , we transform the constrained BOO with HPD regions into an unconstrained optimization problem, $\max_{\pi, M} \left(V_1^{\pi, M}(s_1) + \lambda_k (\log \Pr(M|\mathcal{H}_k) - \log \epsilon_k) \right)$, where $\Pr(M|\mathcal{H}_k) = \frac{\Pr(\mathcal{H}_k|M)\Pr(M)}{\Pr(\mathcal{H}_k)}$. Once the Lagrange multiplier is determined, the constant $\log \epsilon_k$ and the marginal likelihood $\log \Pr(\mathcal{H}_k)$ are irrelevant, and the optimization is equivalent to $\max_{\pi, M} \left(V_1^{\pi, M}(s_1) + \lambda_k (\log \Pr(\mathcal{H}_k|M) + \log \Pr(M)) \right)$. This gives rise to the BOO algorithm as shown in Algorithm 3.

Algorithm 3 BOO

- 1: **for** episode $k = 1, 2, \dots$ **do**
 - 2: Compute $\pi_k \in \arg \max_{\pi} \max_M \left(V_1^{\pi, M}(s_1) + \lambda_k (\log \Pr(\mathcal{H}_k|M) + \log \Pr(M)) \right)$
 - 3: Execute π_k for an episode
-

However, the problem is how to determine the value of λ_k . We note that the following argument by posterior sampling provides an inexact yet insightful viewpoint. A strict discussion is presented in Appendix B. The maximization of the BOO objective can be considered the problem of selecting the best one among multiple posterior samples (see Algorithm 4). When only one sample is taken, Algorithm 4 reduces to the well-known PSRL [21]. As the number of samples j goes to infinity, maximizing among all posterior samples gives approximately the solution of the above optimization.

Algorithm 4 BOO via Posterior Sampling

- 1: **for** episode $k = 1, 2, \dots$ **do**
 - 2: Sample $M_k^1, M_k^2, \dots, M_k^j \sim \Pr(\cdot|\mathcal{H}_k)$
 - 3: Compute $\pi_k \in \arg \max_{\pi} \max_{i \in [j]} \left(V_1^{\pi, M_k^i}(s_1) + \lambda_k (\log \Pr(\mathcal{H}_k|M_k^i) + \log \Pr(M_k^i)) \right)$
 - 4: Execute π_k for an episode
-

Notice that we can regard Algorithm 4 as maximizing over several identically distributed random variables, $\lambda_k \log \Pr(M_k^i|\mathcal{H}_k) + V_1^{\pi_k, M_k^i}(s_1)$. The magnitude of their mean does not matter since subtracting a constant will not change the optimum. The thing that matters is their variation. We need to make sure that the variation is not dominated by either $V_1^{\pi_k, M_k^i}(s_1)$ or $\lambda_k \log \Pr(M_k^i|\mathcal{H}_k)$. If the value dominates the variation, the algorithm will select an unreliable model, which causes inefficiency. If the variation is dominated by the probability, it shows a strong preference for high-probability models and hesitates to explore.

Theorem 4.1 shows that the standard deviation of log probability is at least a constant (proof in Appendix A).

Theorem 4.1 (Variation of the log-posterior density). *Suppose $X^n = (X_1, X_2, \dots, X_n)$ are observations from a stochastic process whose distribution P_{θ} depends on $\theta \in \Theta$, an open subset of \mathbb{R}^m . Assume that the posterior is asymptotic normal, and the log-posterior density is continuous. The variance of the log-posterior density satisfies $\liminf_{n \rightarrow \infty} \text{Var}_{\theta} [\log \Pr(\theta|X^n)] \geq m/2$.*

The process where an algorithm \mathfrak{A} interacts with a random MDP parameterized by θ is a stochastic process where the observation X_t is the state-action pair (s_t, a_t) , and the observation distribution $\Pr(X^n|\theta, \mathfrak{A})$ is a distribution depending on θ . Thus, by Theorem 4.1, we know that if the posterior distribution is asymptotic normal, then, for a random sample θ from the posterior distribution, the variance of the random variable $\log \Pr(\theta|X^n)$ is at least $m/2$ when n is large enough.

Theorem 4.1 relies on the asymptotic normality of the posterior distribution, which is satisfied under some regularity conditions specified in [35]. The theorem only certifies a lower bound on the variance rather than establishing its convergence. Arguably, this is mostly a technical issue arising from the unboundedness of the log-posterior density. It is possible to strengthen this result and derive the convergence of the variance. Indeed, it is empirically observed that this quantity converges rapidly to a constant.

For a proper RL algorithm, it is clear that the standard deviation of $V_1^{\pi_k^i, M_k^i}(s_1)$ should shrink w.r.t. k . Otherwise, the algorithm fails to find out the optimal value and policy. Suppose that the standard deviation of $V_1^{\pi_k^i, M_k^i}$ is $\tilde{O}(k^{-1/\alpha})$, where \tilde{O} is a variant of the big O notation that ignores logarithmic factors. We need to set the scaling parameter λ_k proportionally such that the variation of the log-posterior density term $\log \Pr(M_k^i|\mathcal{H}_k)$ matches that of the value term $V^{\pi_k^i, M_k^i}$. Given a specific model class and noise type, it is possible to derive a worst-case rate for the shrinkage of the value uncertainty and thus determine the proper value of λ_k . Nonetheless, the rate of information revealing could appear in an instance-dependent manner. That is to say, the optimal policy of some MDPs could be inherently easier to determine than the others. Even in a fixed MDP, the rate of uncertainty shrinkage could also change abruptly. For example, consider a ReLU bandit problem, where $B = \{a \in \mathbb{R}^d \mid \|a\|_2 \leq 1\}$, and the agent at each episode selects an action $a \in B$ and receives a Gaussian reward of mean $\max(\theta^\top a, 0)$. If it happens that the agent selects an action a in the inactive region $I = \{a \in B \mid \theta^\top a \leq 0\}$, the resulting observation reveals little information about the optimal action. Suppose the parameter θ is selected such that the positive region $B \setminus I$ is of a maximum width $\epsilon > 0$. Then, in the worst case, the agent needs to explore $\Omega(1/\epsilon^{d-1})$ actions in order to find the positive region, but the uncertainty starts diminishing rapidly whenever the positive region is identified.

The view of matching the log probability variation with the value uncertainty provides rationales for the dynamic adjustment of λ_k . Meanwhile, it points out the potential limitation of decaying λ_k with a fixed rate. In this paper, we focus on methods that scale λ_k with a fixed rate, but as mentioned above, it would be fascinating to adjust it in an instance-dependent manner.

4.3 BOO Regret

In this section, we introduce the regret of BOO. The regret of BOO relies on both the decay rate of the scaling parameter and the complexity of model class. We assume that the model class resides in a d -dimensional RKHS, which roughly means any function in the model class can be represented as a linear function of a potentially unknown finite-dimensional feature map. This assumption is not very restrictive because it places no restriction on the choice of feature map except for the finitude of the dimension.

In Appendix B, we derive the asymptotic regret of the BOO algorithm in Theorem B.1. As suggested by Theorem B.1, the optimal asymptotic regret of BOO is $O(\sqrt{K} \log K)$, i.e., $\tilde{O}(\sqrt{K})$, achieved by setting $\lambda_k^* = c/\sqrt{k}$, where λ_k^* is the optimal scaling parameter and c is a constant associated with the model class. More detailed derivations and conclusions can be found in Appendix B.

We can further interpret $\lambda_k^* = c/\sqrt{k}$ as $\lambda_k^* = \xi_k^V / \xi_k^M$, where ξ_k^M stands for the variation of the log-posterior density for the model M_k , and ξ_k^V represents the value uncertainty. In contrast to the discussion in Section 4, where we consider samples from the posterior distribution and the variation of log-posterior density remains constant, the width of the HPD region measured by the variation of log-posterior density is, in fact, growing at a rate depending on the log covering number. According to Lemma D.7 in Appendix D.1, $\xi_k^M = O(\log k)$. Therefore, we have $\xi_k^V = \lambda_k^* \xi_k^M = O\left(\frac{\log k}{\sqrt{k}}\right)$. Section 5 will highlight the need of manipulating with ξ_k^V and ξ_k^M in optimization.

5 Optimization

In this section, we introduce optimization methods for BOO based on posterior sampling and gradient descent, respectively, and discuss in detail the problems of gradient-based optimization and our proposed solutions.

5.1 Optimization via Posterior Sampling

Algorithm 4 provides a method of optimizing the BOO objective via posterior sampling. As discussed previously, the optimal scaling parameter λ_k for this algorithm is proportional to the value uncertainty, i.e., $\lambda_k = \xi_k^V = \lambda_k^* \xi_k^M$. The potential advantage of Algorithm 4 over PSRL is that, unlike PSRL, it does not require posterior samples to be independent. Even the requirement that samples are from the posterior can be relaxed. These features are essential either when the posterior distribution is approximated or when the posterior samples are produced by Markov chain Monte Carlo methods and are correlated.

5.2 Optimization via Gradient-Based Methods

We can also perform the optimization via gradient-based methods. The objective function of BOO consists of two parts, $V_1^{\pi, M_k^i}(s_1)$ and $\lambda_k(\log \Pr(\mathcal{H}_k | M_k^i) + \log \Pr(M_k^i))$. The gradients of the log-likelihood and the log-prior are easily obtained. For the value part, we provide a value model gradient that admits a similar form as the well-known policy gradient [36, 37] (proof in Appendix E.1). The value model gradient is amenable to Monte Carlo approximation and can be computed exactly for finite MDPs.

Theorem 5.1 (Value model gradient). *Suppose that the transition function P^{M_θ} and reward function R^{M_θ} of model M_θ , the gradient of the value $V_1^{\pi, M_\theta}(s_1)$ w.r.t. the model is*

$$\nabla_\theta V_1^{\pi, M_\theta}(s_1) = \mathbb{E}_{\tau \sim \pi, M_\theta} \left[\sum_{h=1}^H \nabla_\theta \bar{R}^{M_\theta}(s_h, a_h) + \sum_{h=1}^{H-1} V_{h+1}^{\pi, M_\theta}(s_{h+1}) \nabla_\theta \log P^{M_\theta}(s_{h+1} | s_h, a_h) \right], \quad (4)$$

where $\tau = (s_1, a_1, \dots, s_H, a_H)$ is a trajectory, $\tau \sim \pi, M_\theta$ means that the trajectory is formed by the interaction of the policy π and the model M_θ , and $P^{M_\theta}(s_{h+1} | s_h, a_h)$ is the probability of s_{h+1} under distribution $P^{M_\theta}(s_h, a_h)$.

Nevertheless, the value model gradient suffers the same inefficiency just as the policy gradient [38] since the model gradient for a particular state-action pair is 0 whenever it is not visited under the current model and policy.

We propose some techniques to improve the optimization efficiency of BOO, and conduct ablation experiments to verify the effectiveness of our proposed methods in Section G. Two of the most effective techniques are described below, and the rest of the techniques are detailed in Appendix F.

5.2.1 Mean Reward Bonus

In the gradient-based optimization, the model becomes optimistic on state-action pairs it visits frequently, which in turn makes these state-action pairs more appealing. This mutual strengthening phenomenon makes optimization easily stuck at local optima. One way to solve this problem is to increase the rewards of all state-action pairs, which raises the attractiveness of less visited state-action pairs. This can be achieved by adding a bonus term $\xi_k^V H \mathbb{E}_{(s,a) \sim U_{S \times A}} [R(s, a)]$ to the BOO objective. Here, $U_{S \times A}$ is the uniform distribution over the state-action space. The coefficient ξ_k^V ensures that the bonus decays with the value uncertainty. The unvisited state-action pairs will eventually be visited by the policy because they have sufficiently high rewards. Our experiments will show that, in tabular setting, this method is very effective. However, a concern is that it might fail to scale to high dimensional state-action space.

5.2.2 Entropy Regularization

Another intricacy of the optimization is that the optimal solution of the BOO objective could change dramatically across the parameter space from episode to episode, which renders the optimization

extremely hard. A revealing fact is that a small change in the model could change the optimal policy dramatically, making the loss landscape unsmooth. Hence, we introduce an entropy-regularized optimization procedure, which starts with a high initial entropy regularization and gradually annealing. The entropy plays a role in smoothing the policy’s loss landscape such that the optimal policy will not change drastically when the model changes. The smoothing effect of entropy regularization is also discussed previously in [39].

The entropy-regularized learning objective mimics the maximum entropy RL [40]:

$$\begin{aligned} \tilde{J}_k &= V_1^{\pi, M}(s_1) + \lambda_k \log \Pr(M|\mathcal{H}_k) - \xi_k^V \zeta \mathbb{E}_{\tau \sim \pi, M} \left[\sum_{h=1}^H \text{KL}(\pi(s_h, h) \|\hat{\pi}(s_h, h)) \right] \\ &= \tilde{V}_1^{\pi, M}(s_1) + \lambda_k \log \Pr(M|\mathcal{H}_k), \end{aligned} \quad (5)$$

where $\text{KL}(p\|q)$ stands for the relative entropy, $\hat{\pi}$ is a prior policy ensuring π is absolutely continuous w.r.t. $\hat{\pi}$, the hyperparameter ζ controls the amount of entropy, and \tilde{V} is the entropy-regularized value.

The entropy term is downscaled in proportion to the value uncertainty such that the influence of regularization diminishes with time. We denote by \mathcal{M}_θ the set of models parameterized by θ . Let b_ϵ be the smallest number ensuring that, for any $M_\theta \in \mathcal{M}_\theta$, there exists $\pi_\epsilon \in \Pi_\epsilon = \{\pi \mid \text{KL}(\pi(s, h) \|\hat{\pi}(s, h)) \leq b_\epsilon, \forall s \in \mathcal{S}, h \in [H]\}$ such that $\sup_{\pi^*} V_1^{\pi^*, M}(s_1) - V_1^{\pi_\epsilon, M}(s_1) \leq \epsilon$. Then, the value of the maximum entropy policy $\tilde{\pi}^* = \arg \max_{\tilde{\pi}} V_1^{\tilde{\pi}, M}(s_1)$ satisfies that

$$\begin{aligned} V_1^{\tilde{\pi}^*, M}(s_1) &\geq V_1^{\tilde{\pi}^*, M}(s_1) - \xi_k^V \zeta \mathbb{E}_{\tau \sim \tilde{\pi}^*, M} \left[\sum_{h=1}^H \text{KL}(\tilde{\pi}^*(s_h, h) \|\hat{\pi}(s_h, h)) \right] \\ &\geq V_1^{\pi_\epsilon, M}(s_1) - \xi_k^V \zeta \mathbb{E}_{\tau \sim \pi_\epsilon, M} \left[\sum_{h=1}^H \text{KL}(\pi_\epsilon(s_h, h) \|\hat{\pi}(s_h, h)) \right] \\ &\geq \sup_{\pi^*} V_1^{\pi^*, M}(s_1) - \epsilon - \xi_k^V \zeta b_\epsilon. \end{aligned} \quad (6)$$

If ϵ is sufficiently small, then the sub-optimality gap caused by entropy regularization will decrease at the same rate as the decay of value uncertainty, which ensures that the resulting policy covers multiple uncertain actions without detriment to the performance.

The entropy-regularized value is equivalently defined by the following Bellman backup,

$$\begin{aligned} \tilde{V}_h^{\pi, M}(s_h) &= \mathbb{E}_{a_h \sim \pi(s_h, h)} [\tilde{Q}_h^{\pi, M}(s_h, a_h)] - \xi_k^V \zeta \text{KL}(\pi(s_h, h) \|\hat{\pi}(s_h, h)), \forall h \in [H], \\ \tilde{Q}_h^{\pi, M}(s_h, a_h) &= \bar{R}^M(s_h, a_h) + \mathbb{E}_{s_{h+1} \sim P^M(s_h, a_h)} [\tilde{V}_{h+1}^{\pi, M}(s_{h+1})], \forall h \in [H-1], \end{aligned} \quad (7)$$

where $\tilde{Q}_H^{\pi, M}(s_H, a_H) = \bar{R}^M(s_H, a_H)$. The optimization of the policy can be carried out by, for example, the maximum entropy actor-critic algorithm [40] or the soft actor-critic algorithm [41]. In finite MDPs, the optimal value is given by soft value iteration,

$$\tilde{V}_h^{\tilde{\pi}^*, M}(s_h) = \xi_k^V \zeta \log \sum_{a_h \in \mathcal{A}} \hat{\pi}(a_h | s_h, h) \exp \left(\frac{\tilde{Q}_h^{\tilde{\pi}^*, M}(s_h, a_h)}{\xi_k^V \zeta} \right), \quad (8)$$

where $\pi(a_h | s_h, h)$ is the probability of a_h under the distribution $\pi(s_h, h)$. The maximum entropy policy $\tilde{\pi}^*$ is then derived from the optimal value as

$$\tilde{\pi}^*(a_h | s_h, h) = \hat{\pi}(a_h | s_h, h) \exp \left(\frac{\tilde{Q}_h^{\tilde{\pi}^*, M}(s_h, a_h) - \tilde{V}_h^{\tilde{\pi}^*, M}(s_h)}{\xi_k^V \zeta} \right). \quad (9)$$

The gradient of the entropy-regularized value w.r.t. the model is similar to that specified in Theorem 5.1, with the value replaced by the entropy-regularized value.

6 Experiments

This section compares the performance of BOO with PSRL and UCRL2 in RiverSwim, Chain, and Random MDPs. RiverSwim and Chain are hard-exploration MDPs requiring the agent to explore

efficiently, while Random MDPs test the average performance. Two implementations of BOO are presented, namely, FiniteBOO and BPS. FiniteBOO is the BOO with entropy regularization and mean reward bonus mentioned in Section 5.2, which is empirically found to be the best variant in tabular setting according to the ablation study. BPS is an implementation of Algorithm 4 (BOO via posterior sampling).

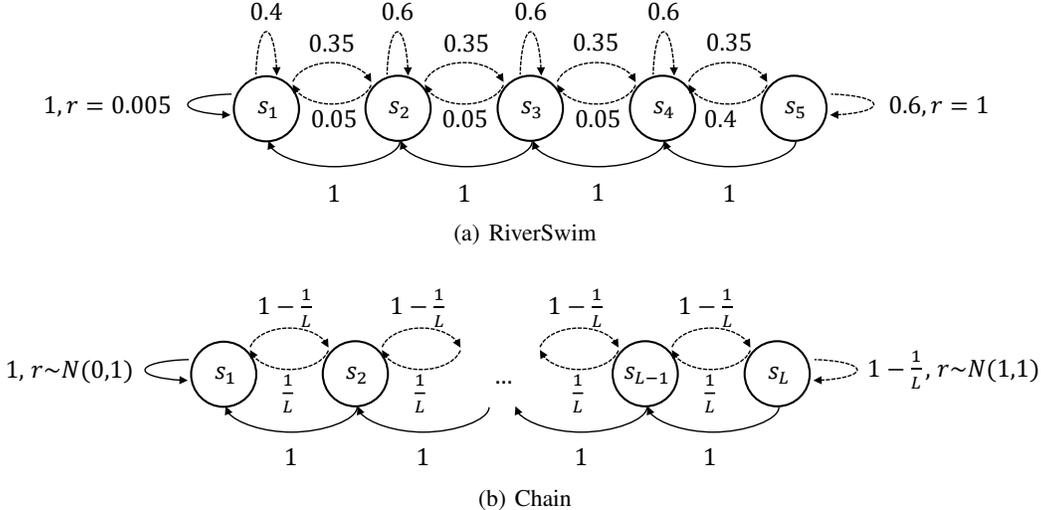


Figure 1: Illustrative diagrams for RiverSwim and Chain. Solid and dotted arrows represent actions, “left” and “right”, respectively, with the transition probability tagged. The action left never fails, but the action right often fails. Except for the rewards shown in the figure, the rewards of other state-action pairs are all zero.

6.1 RiverSwim

The RiverSwim is an MDP where states are organized in chains, and the agent can move left or right, as shown in Figure 1(a). Although the rightmost state has a huge reward, the action of moving right fails with a high probability. Only a policy moves right at each time period has a small chance of success.

We start with the experiment on the RiverSwim to demonstrate the performance of our algorithm in the face of high transition uncertainty. We perform experiments for ten seeds on RiverSwim with $|\mathcal{S}| = H = 5$, $|\mathcal{A}| = 2$ and record the cumulative regret over 100,000 time steps. As shown in Figure 2(a), our algorithm is competitive to PSRL and outperforms UCRL2 significantly.

6.2 Chain

The chain MDP is a variant of the RiverSwim, which has Gaussian rewards and relatively deterministic transitions, as shown in Figure 1(b). Although transitions are relatively certain, the stochastic rewards make the problem difficult to explore. The horizon H and the number of states $|\mathcal{S}|$ are equal to the length of the chain.

We evaluate our algorithms in Chain of a length $L \in \{10, 20, 40\}$ for 100,000 episodes and ten random trials. Figure 2(d~f) illustrates that our algorithm compares favorably with PSRL and UCRL2, which certifies the effectiveness of BOO in problems requiring long-term planning.

6.3 Random MDPs

Random MDPs are tabular MDP models randomly generated from a prior distribution and used to test the general performance of the algorithm.

We randomly generate 100 stochastic MDPs with $|\mathcal{S}| = |\mathcal{A}| = H = 5$ and $|\mathcal{S}| = H = 20$, $|\mathcal{A}| = 5$ from the prior and measure the performance of algorithms over 10,000 timesteps. As shown in

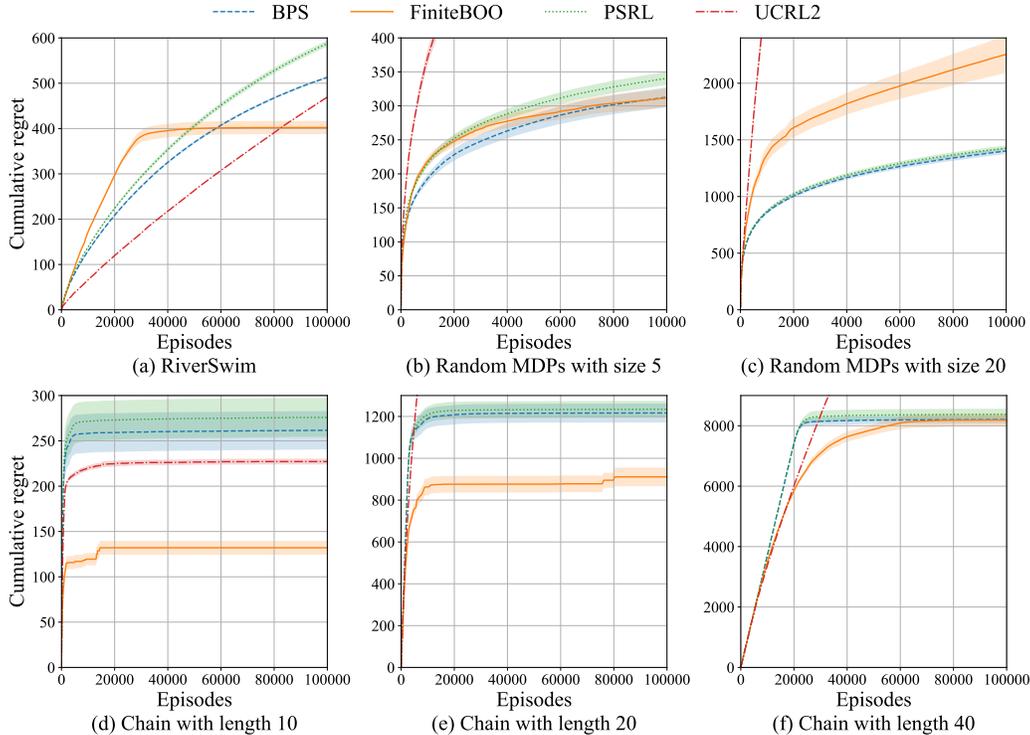


Figure 2: (a~f) show the cumulative regrets of different algorithms on RiverSwim ($|\mathcal{S}| = H = 5$, $|\mathcal{A}| = 2$), Random MDPs ($|\mathcal{S}| = H = 5$, $|\mathcal{A}| = 5$), Random MDPs ($|\mathcal{S}| = H = 20$, $|\mathcal{A}| = 5$), and the chains of different lengths, respectively.

Figure 2(b), the performance of BOO and BPS is close to PSRL, and both algorithms outperform UCRL2 significantly. However, the performance gap between FiniteBOO and PSRL in Figure 2(c) indicates that there remains a challenge in the optimization of large-scale MDPs.

7 Conclusion

This paper proposes BOO as a generic model-based RL algorithm. It is provably sample-efficient and enjoys $\tilde{O}(\sqrt{K})$ regret for models in a finite-dimensional RKHS, where K is the number of episodes. To optimize the BOO objective, we propose the value model gradient and optimization techniques, such as entropy regularization, to improve its efficiency. Through our experiments, we have shown that BOO is competitive with PSRL and outperforms UCRL2 greatly. However, to apply BOO in real-world RL problems, there remains lots of work to be done. Importantly, we need to develop methods that further facilitate the gradient-based optimization of BOO in large-scale problems. It is also an appealing direction for future work to adapt the scaling parameter of BOO on an instance-dependent basis.

Acknowledgements

This work is supported by National Key Research and Development Program of China (2020AAA0107200), the National Science Foundation of China (61921006, 61876119, 62276126), the Natural Science Foundation of Jiangsu (BK20221442), and the Fundamental Research Funds for the Central Universities (022114380010).

References

- [1] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction (Second Edition)*. MIT Press, 2018.

- [2] Q. Liu, J. Zhai, Z. Zhang, S. Zhong, Q. Zhou, P. Zhang, and J. Xu, “A survey on deep reinforcement learning,” *Chinese Journal of Computers*, vol. 41, no. 1, pp. 1–27, 2018.
- [3] M. Ghavamzadeh, S. Mannor, J. Pineau, and A. Tamar, “Bayesian reinforcement learning: A survey,” *Foundations and Trends in Machine Learning*, vol. 8, no. 5-6, pp. 359–483, 2015.
- [4] C. H. Papadimitriou and J. N. Tsitsiklis, “The complexity of Markov decision processes,” *Mathematics of Operations Research*, vol. 12, no. 3, pp. 441–450, 1987.
- [5] O. Madani, S. Hanks, and A. Condon, “On the undecidability of probabilistic planning and related stochastic optimization problems,” *Artificial Intelligence*, vol. 147, no. 1, pp. 5–34, 2003.
- [6] Z. Zhang, M. L. Littman, and X. Chen, “Covering number as a complexity measure for POMDP planning and learning,” in *AAAI Conference on Artificial Intelligence*, pp. 1853–1859, 2012.
- [7] Z. Zhang, D. Hsu, and W. S. Lee, “Covering number for efficient heuristic-based POMDP planning,” in *International Conference on Machine Learning*, pp. 28–36, 2014.
- [8] Z. Zhang, D. Hsu, W. S. Lee, Z. W. Lim, and A. Bai, “PLEASE: Palm leaf search for POMDPs with large observation spaces,” in *International Conference on Automated Planning and Scheduling*, pp. 249–257, 2015.
- [9] C. Wu, R. Kong, G. Yang, X. Kong, Z. Zhang, Y. Yu, D. Li, and W. Liu, “LB-DESPOT: Efficient online POMDP planning considering lower bound in action selection (student abstract),” in *AAAI Conference on Artificial Intelligence*, pp. 15927–15928, 2021.
- [10] C. Wu, G. Yang, Z. Zhang, Y. Yu, D. Li, W. Liu, and J. Hao, “Adaptive online packing-guided search for POMDPs,” in *Advances in Neural Information Processing Systems*, pp. 28419–28430, 2021.
- [11] M. G. Azar, I. Osband, and R. Munos, “Minimax regret bounds for reinforcement learning,” in *International Conference on Machine Learning*, pp. 263–272, 2017.
- [12] S. R. Sinclair, S. Banerjee, and C. L. Yu, “Adaptive discretization for episodic reinforcement learning in metric spaces,” in *ACM on Measurement and Analysis of Computing Systems*, pp. 1–44, 2019.
- [13] C. Jin, Z. Yang, Z. Wang, and M. I. Jordan, “Provably efficient reinforcement learning with linear function approximation,” in *Conference on Learning Theory*, pp. 2137–2143, 2020.
- [14] V. Dani, T. Hayes, and S. Kakade, “Stochastic linear optimization under bandit feedback,” in *Conference on Learning Theory*, pp. 355–366, 2008.
- [15] I. Osband and B. V. Roy, “Why is posterior sampling better than optimism for reinforcement learning?,” in *International Conference on Machine Learning*, pp. 2701–2710, 2017.
- [16] T. Jaksch, R. Ortner, and P. Auer, “Near-optimal regret bounds for reinforcement learning,” *Journal of Machine Learning Research*, vol. 11, no. 51, pp. 1563–1600, 2010.
- [17] I. Osband, D. Russo, and B. Van Roy, “(More) efficient reinforcement learning via posterior sampling,” in *Advances in Neural Information Processing Systems*, pp. 3003–3011, 2013.
- [18] A. Guez, N. Heess, D. Silver, and P. Dayan, “Bayes-adaptive simulation-based search with value function approximation,” in *Advances in Neural Information Processing Systems*, pp. 451–459, 2014.
- [19] S. R. Chowdhury and A. Gopalan, “Online learning in kernelized markov decision processes,” in *International Conference on Artificial Intelligence and Statistics*, pp. 3197–3205, 2019.
- [20] S. Kakade, A. Krishnamurthy, K. Lowrey, M. Ohnishi, and W. Sun, “Information theoretic regret bounds for online nonlinear control,” in *Advances in Neural Information Processing Systems*, pp. 15312–15325, 2020.
- [21] M. J. A. Strens, “A Bayesian framework for reinforcement learning,” in *International Conference on Machine Learning*, pp. 943–950, 2000.

- [22] D. Russo and B. Van Roy, “Learning to optimize via posterior sampling,” *Mathematics of Operations Research*, vol. 39, no. 4, pp. 1221–1243, 2014.
- [23] I. Osband, C. Blundell, A. Pritzel, and B. Van Roy, “Deep Exploration via Bootstrapped DQN,” in *Advances in Neural Information Processing Systems*, vol. 29, pp. 4026–4034, 2016.
- [24] Ian Osband, Benjamin Van Roy, and Zheng Wen, “Generalization and Exploration via Randomized Value Functions,” in *Proceedings of The 33rd International Conference on Machine Learning* (Maria Florina Balcan and Kilian Q. Weinberger, eds.), vol. 48, pp. 2377–2386, PMLR, June 2016.
- [25] I. Osband, B. V. Roy, D. J. Russo, and Z. Wen, “Deep Exploration via Randomized Value Functions,” *Journal of Machine Learning Research*, vol. 20, no. 124, pp. 1–62, 2019.
- [26] S. Curi, F. Berkenkamp, and A. Krause, “Efficient model-based reinforcement learning through optimistic policy search and planning,” in *Advances in Neural Information Processing Systems* (H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, eds.), pp. 14156–14170, 2020.
- [27] P. Kumar and A. Becker, “A new family of optimal adaptive controllers for Markov chains,” *IEEE Transactions on Automatic Control*, vol. 27, no. 1, pp. 137–146, 1982.
- [28] X. Liu, P.-C. Hsieh, Y. H. Hung, A. Bhattacharya, and P. Kumar, “Exploration through reward biasing: Reward-biased maximum likelihood estimation for stochastic multi-armed bandits,” in *International Conference on Machine Learning*, pp. 6248–6258, 2020.
- [29] Y.-H. Hung, P.-C. Hsieh, X. Liu, and P. R. Kumar, “Reward-biased maximum likelihood estimation for linear stochastic bandits,” in *AAAI Conference on Artificial Intelligence*, pp. 7874–7882, 2021.
- [30] A. Mete, R. Singh, X. Liu, and P. R. Kumar, “Reward biased maximum likelihood estimation for reinforcement learning,” in *Annual Learning for Dynamics and Control*, pp. 815–827, 2021.
- [31] A. Mete, R. Singh, and P. R. Kumar, “Augmented RBMLE-UCB approach for adaptive control of linear quadratic systems,” *arXiv preprint arXiv:2201.10542*, 2022.
- [32] R. J. Hyndman, “Computing and graphing highest density regions,” *The American Statistician*, vol. 50, no. 2, pp. 120–126, 1996.
- [33] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian Data Analysis*. Chapman and Hall/CRC, 2013.
- [34] R. J. Hyndman and Y. Fan, “Sample quantiles in statistical packages,” *The American Statistician*, vol. 50, no. 4, pp. 361–365, 1996.
- [35] O. A. Adekola, *Asymptomatic posterior normality for stochastic processes*. PhD thesis, University of Surrey, Guildford, 1987.
- [36] R. J. Williams, “Simple statistical gradient-following algorithms for connectionist reinforcement learning,” *Machine Learning*, vol. 8, no. 3, pp. 229–256, 1992.
- [37] R. S. Sutton, D. McAllester, S. Singh, and Y. Mansour, “Policy gradient methods for reinforcement learning with function approximation,” in *Advances in Neural Information Processing Systems*, pp. 1057–1063, 2000.
- [38] S. Kakade and J. Langford, “Approximately optimal approximate reinforcement learning,” in *International Conference on Machine Learning*, pp. 267–274, 2002.
- [39] Z. Ahmed, N. L. Roux, M. Norouzi, and D. Schuurmans, “Understanding the impact of entropy on policy optimization,” in *International Conference on Machine Learning*, pp. 151–160, 2019.
- [40] S. Levine, “Reinforcement learning and control as probabilistic inference: Tutorial and review,” *arXiv preprint arXiv:1805.00909*, 2018.

- [41] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, “Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor,” in *International Conference on Machine Learning*, pp. 1861–1870, 2018.
- [42] M. Brookes, “The matrix reference manual,” 2020. [online] <http://www.ee.ic.ac.uk/hp/staff/dmb/matrix/intro.html>.
- [43] A. Klenke, *Probability Theory: A Comprehensive Course*. Springer Science & Business Media, 2013.
- [44] C. A. Micchelli and M. Pontil, “On learning vector-valued functions,” *Neural Computation*, vol. 17, pp. 177–204, Jan. 2005.
- [45] D. Russo and B. Van Roy, “Eluder dimension and the sample complexity of optimistic exploration,” in *Advances in Neural Information Processing Systems*, pp. 2256–2264, 2013.
- [46] C. Jin, Q. Liu, and S. Miryoosefi, “Bellman eluder dimension: New rich classes of RL problems, and sample-efficient algorithms,” in *Advances in Neural Information Processing Systems*, pp. 13406–13418, 2021.
- [47] I. Osband and B. Van Roy, “Model-based reinforcement learning and the eluder dimension,” in *Advances in Neural Information Processing Systems*, pp. 1466–1474, 2014.
- [48] A. Maurer, “A bound on the deviation probability for sums of non-negative random variables,” *Journal of Inequalities in Pure and Applied Mathematics*, vol. 4, no. 1, 2003.
- [49] S. M. Kakade, “A natural policy gradient,” in *Advances in Neural Information Processing Systems*, pp. 1531–1538, 2002.
- [50] A. Agarwal, S. M. Kakade, J. D. Lee, and G. Mahajan, “Optimality and approximation with policy gradient methods in Markov decision processes,” in *Conference on Learning Theory*, pp. 64–66, 2020.
- [51] N. N. Schraudolph, “Fast curvature matrix-vector products for second-order gradient descent,” *Neural Computation*, vol. 14, no. 7, pp. 1723–1738, 2002.
- [52] J. Martens, “New insights and perspectives on the natural gradient method,” *Journal of Machine Learning Research*, vol. 21, no. 146, pp. 1–76, 2020.
- [53] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nature Methods*, vol. 17, pp. 261–272, 2020.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes]
 - (c) Did you discuss any potential negative societal impacts of your work? [No] No relevant potential negative impacts is identified.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
 - (b) Did you include complete proofs of all theoretical results? [Yes]

3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [\[Yes\]](#)
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [\[Yes\]](#)
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [\[Yes\]](#)
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [\[Yes\]](#)
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [\[Yes\]](#)
 - (b) Did you mention the license of the assets? [\[Yes\]](#)
 - (c) Did you include any new assets either in the supplemental material or as a URL? [\[Yes\]](#)
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [\[Yes\]](#)
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [\[Yes\]](#)
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [\[N/A\]](#)
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [\[N/A\]](#)
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [\[N/A\]](#)

Appendix

A Variation Analysis

Theorem 4.1 (Variation of the log-posterior density). *Suppose $X^n = (X_1, X_2, \dots, X_n)$ are observations from a stochastic process whose distribution P_θ depends on $\theta \in \Theta$, an open subset of \mathbb{R}^m . Assume that the posterior is asymptotic normal, and the log-posterior density is continuous. The variance of the log-posterior density satisfies $\liminf_{n \rightarrow \infty} \text{Var}_\theta [\log \Pr(\theta|X^n)] \geq m/2$.*

Proof. We first compute the variance of log probability for a m -dimension normal distribution with mean μ and covariance Σ . Let f be the density of the multivariate normal distribution, and $\theta \sim f$. It follows that

$$\begin{aligned} \text{Var}_\theta [\log f(\theta)] &= \text{Var}_\theta \left[-\frac{1}{2} [\log(|\Sigma|) - (\theta - \mu)^\top \Sigma^{-1}(\theta - \mu) + m \log(2\pi)] \right] \\ &= \frac{1}{4} \text{Var}_\theta [(\theta - \mu)^\top \Sigma^{-1}(\theta - \mu)] \\ &= \frac{1}{4} \text{Var}_{\theta'} [\theta'^\top \Sigma^{-1} \theta'], \end{aligned} \quad (10)$$

where the last equality is derived by setting $\theta' = \theta - \mu$. The random variable θ' follows a normal distribution with zero mean and covariance Σ . The expectation of a quadratic form $\theta'^\top A \theta'$ is [42]

$$\mathbb{E} [\theta'^\top A \theta'] = \text{Tr}(A\Sigma), \quad (11)$$

where $A \in \mathbb{R}^{m \times m}$, and $\text{Tr}(\cdot)$ represents the trace of the matrix. Given matrices $A, B \in \mathbb{R}^{m \times m}$, the expectation of a quartic form $\theta'^\top A \theta' \theta'^\top B \theta'$ is [42]

$$\mathbb{E} [\theta'^\top A \theta' \theta'^\top B \theta'] = \text{Tr}(A\Sigma(B + B^\top)\Sigma) + \text{Tr}(A\Sigma) \text{Tr}(B\Sigma). \quad (12)$$

We then derive its variance as follows,

$$\begin{aligned} \text{Var}_{\theta'} [\theta'^\top \Sigma^{-1} \theta'] &= \mathbb{E} \left[(\theta'^\top \Sigma^{-1} \theta' - \mathbb{E} [\theta'^\top \Sigma^{-1} \theta']) (\theta'^\top \Sigma^{-1} \theta' - \mathbb{E} [\theta'^\top \Sigma^{-1} \theta'])^\top \right] \\ &= \mathbb{E} [(\theta'^\top \Sigma^{-1} \theta' - m)(\theta'^\top \Sigma^{-1} \theta' - m)^\top] \\ &= \mathbb{E} [\theta'^\top \Sigma^{-1} \theta' \theta'^\top \Sigma^{-1} \theta'] - 2m \mathbb{E} [\theta'^\top \Sigma^{-1} \theta'] + m^2 \\ &= \text{Tr} (2\Sigma^{-1} \Sigma \Sigma^{-1} \Sigma) + [\text{Tr} (\Sigma^{-1} \Sigma)]^2 - 2m \text{Tr}(\Sigma^{-1} \Sigma) + m^2 \\ &= 2 \text{Tr}(I_m) + \text{Tr}(I_m)^2 - 2m \text{Tr}(I_m) + m^2 \\ &= 2m, \end{aligned} \quad (13)$$

where $I_m \in \mathbb{R}^{m \times m}$ is the identity matrix, and the third-to-last equality makes use of the quadratic and quartic expectations. It follows that the variance of log probability of any normal distribution is $m/2$.

As the posterior distribution is assumed to be asymptotically normal, we complete the proof by applying Portmanteau theorem which states that $\liminf_{n \rightarrow \infty} \mathbb{E}[g(X_n)] \geq \mathbb{E}[g(X)]$ for any lower semi-continuous function g bounded from below if X_n converges in distribution to X [43]. \square

B BOO Regret Analysis

This section analyzes the regret of BOO. The regret of BOO relies on both the decay rate of the scaling parameter and the complexity of the model class. We will first declare our assumptions on the model class. Based on these assumptions, we can measure the complexity of the model class in terms of the covering number and the eluder dimension. Then, we establish the regret bound for different choices of the decay rate of the scaling parameter.

Let the state space \mathcal{S} be a subset of \mathbb{R}^m . Denote the state-action space $\mathcal{S} \otimes \mathcal{A}$ as \mathcal{X} and state-reward space $\mathcal{S} \otimes \mathbb{R}$ as \mathcal{Y} . A model function is a function mapping \mathcal{X} to \mathcal{Y} . For each $y \in \mathcal{Y}$, ϵ_y is a σ -sub-Gaussian noise, i.e. $\log \mathbb{E}[e^{v^\top \epsilon_y}] \leq \|v\|_2^2 \sigma^2 / 2$ for all $v \in \mathbb{R}^m$. We first assume that the real MDP can be fully characterized by the model function and the sub-Gaussian noise.

Assumption 1 (Models with additive sub-Gaussian noises). *Let \mathcal{F} be a class of model functions $f : \mathcal{X} \rightarrow \mathcal{Y}$. Assume that there exists a function $f^* : \mathcal{X} \rightarrow \mathcal{Y}$ in \mathcal{F} such that for any $x \in \mathcal{X}$, $f^*(x) + \epsilon_{f^*(x)}$ is identically distributed as in the real MDP M^* .*

Assumption 1 requires the output distribution to be determined by its mean. This requirement is satisfied when the output distribution is modeled as, for example, Gaussian with known variance. The categorical distribution also satisfies this requirement by representing each category as a one-hot vector, in which case the mean vector equates the probability vector of the distribution.

Since the noise is assumed sub-Gaussian, we adopt the independent Gaussian likelihood as a surrogate for the actual likelihood function [29], i.e.,

$$\log \Pr(y|x, f) = -\frac{n}{2} \log(2\pi\sigma^2) - \frac{1}{2\sigma^2} \|y - f(x)\|_2^2. \quad (14)$$

A linear space \mathcal{H} of functions $f : \mathcal{X} \rightarrow \mathcal{Y}$ on a set \mathcal{X} is said to RKHS if for any $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ the linear functional mapping $f \in \mathcal{H}$ to $f(x)^\top y$ is continuous [44]. We define a d -dimensional RKHS \mathcal{H} via a feature map $\phi : \mathcal{X} \rightarrow \mathbb{R}^d$ as $\{f : \mathcal{X} \rightarrow \mathcal{Y} \mid W \in \mathbb{R}^{(m+1) \times d}, f(x) = W\phi(x), \forall x \in \mathcal{X}\}$. It is easily seen that \mathcal{H} is indeed an RKHS. We assume that the model function class \mathcal{F} is a subset of the d -dimensional RKHS \mathcal{H} .

Assumption 2 (Model functions in d -dimensional RKHS). *Assume that $\mathcal{F} \subseteq \{f \mid W \in \mathbb{R}^{(m+1) \times d}, \|W\|_F \leq r_w, f(x) = W\phi(x), \|\phi(x)\| \leq r_\phi, \forall x \in \mathcal{X}\}$, where r_w and r_ϕ are constants, and $\|\cdot\|_F$ is the Frobenius norm.*

In Appendix C, we prove that the log covering number and the distributional eluder dimension for the function class \mathcal{F} are $O(md \log(r_\phi r_w / \alpha))$ and $O(d \log(r_\phi r_w / \epsilon))$, respectively. Here, the upper bound of the distributional eluder dimension is established for any set of probability measures on \mathcal{X} .

The third assumption is the Lipschitz continuity of the one-step value function, which connects the model difference to the value difference. Let $U_{x,V}(M)$ be the one-step value function with the state-action pair x and the value function V^M , i.e., $U_{x,V}(M) = \bar{R}^M(x) + P_x^M(V)$, where $P_x^M(V) = \int V^M(s') P^M(s'|x) ds'$. We define the Lipschitz constant with respect to the 2-norm as

$$|U_{x,V}(M) - U_{x,V}(M')| \leq L_{x,V} \|f^M(x) - f^{M'}(x)\|_2, \quad (15)$$

where f^M is the model function of M .

Assumption 3 (One-step value Lipschitz continuity). *Let \mathcal{M} be the set of MDPs defined by the class \mathcal{F} of model functions. Assume that the Lipschitz constant $L_{x,V^\pi,M}$ for any $x \in \mathcal{X}$, $M \in \mathcal{M}$, and π is bounded by L , i.e., $L \geq \max_{x \in \mathcal{X}, M \in \mathcal{M}, \pi} L_{x,V^\pi,M}$.*

Finally, the asymptotic regret of BOO is derived in Theorem B.1 based on the assumptions and conclusions we introduced above, see Appendix D.4 for more details of the derivation.

Theorem B.1 (Bayesian optimistic optimization regret). *Let \mathcal{F} be the function class defined in Assumption 2, i.e., $\mathcal{F} \subseteq \{f \mid W \in \mathbb{R}^{(m+1) \times d}, \|W\|_F \leq r_w, f(x) = W\phi(x), \|\phi(x)\| \leq r_\phi, \forall x \in \mathcal{X}\}$. The log-prior probability is assumed to be uniformly bounded by some constant. Setting $\lambda_k = ck^{-v_1}(\log k)^{-v_2}$ for $c > 0$ and $0 < v_1 < 1$ or $v_1 = 0, v_2 \geq 0$, the asymptotic regret of BOO is*

$$\begin{cases} O\left(HLdr_w r_\phi \sigma \sqrt{m(r_c + 1)K} \log K\right) & \text{if } v_1 = \frac{1}{2}, v_2 = 0 \\ O\left(HL^2 d / (c\sigma^2) K^{v_1} (\log K)^{v_2+1}\right) & \text{if } v_1 > \frac{1}{2} \text{ or } v_1 = \frac{1}{2}, v_2 > 0 \\ O\left(Hmdr_w^2 r_\phi^2 \sigma^4 K^{1-v_1} (\log K)^{1-v_2}\right) & \text{if } v_1 < \frac{1}{2} \text{ or } v_1 = \frac{1}{2}, v_2 < 0, \end{cases} \quad (16)$$

where m is the dimension of the state space and r_c is a constant determined by $c = \frac{\sqrt{(r_c+1)}}{r_c} \frac{2\sqrt{2}L}{r_w r_\phi \sigma^3 \sqrt{m}}$.

As suggested by Theorem B.1, the best performance of BOO is achieved by setting $\lambda_k^* = \frac{c}{\sqrt{k}}$.

C Complexity Measurement for the Model Function Class

In this section, we first introduce the definition of the covering number and the eluder dimension and then measure the complexity of the function class \mathcal{F} based on them.

C.1 Definition of the Covering Number and the Eluder Dimension

Our regret analysis uses the notions of the covering number and the eluder dimension [45, 22]. They measure different aspects of function classes.

The covering number is introduced for measuring the size of the function class.

Definition 1 (Covering number). Let $(G, \|\cdot\|)$ be a metric space and $\mathcal{F} \subseteq G$. The covering number $N(\mathcal{F}, \alpha, \|\cdot\|)$ for $\alpha > 0$ is the minimum cardinality of $C \subseteq \mathcal{F}$ such that for all $f \in \mathcal{F}$, there exists $c \in C$ satisfying $\|f - c\| \leq \alpha$.

The eluder dimension is a complexity measure for the exploration difficulty. Similar to [46], we generalize the original eluder dimension [45] to the distributional eluder dimension by measuring the difference on distributions.

Definition 2 (ϵ -dependence between distributions [22, 46]). Let \mathcal{F} be a function class defined on a set \mathcal{X} , and ν, μ_1, \dots, μ_n be probability measures, over \mathcal{X} . We say ν is ϵ -dependent on $\{\mu_1, \mu_2, \dots, \mu_n\}$ w.r.t. \mathcal{F} if any $f_1, f_2 \in \mathcal{F}$ satisfying $\sqrt{\sum_{i=1}^n \mu_i(\|f_1 - f_2\|_2^2)} \leq \epsilon$ also satisfies $\sqrt{\nu(\|f_1 - f_2\|_2^2)} \leq \epsilon$ for ϵ . Here, $\mu(g) = \int g d\mu$. If ν is not ϵ -dependent on $\{\mu_1, \mu_2, \dots, \mu_n\}$, then it is said to be ϵ -independent of $\{\mu_1, \mu_2, \dots, \mu_n\}$ w.r.t. \mathcal{F} .

Let f^* be the real function where observations are generated. Intuitively, a distribution ν is ϵ -independent of $\{\mu_1, \mu_2, \dots, \mu_n\}$ means that a function $f \in \mathcal{F}$ indistinguishable from the real function f^* on historical distributions could still be significantly different on the distribution ν .

Definition 3 (Distributional eluder (DE) dimension [22, 46]). Let \mathcal{F} be a function class defined on a set \mathcal{X} , and Π be a class of probability measures over \mathcal{X} . The distributional eluder dimension $\text{DEdim}(\mathcal{F}, \Pi, \epsilon)$ is the length of the longest sequence in Π such that every element in the sequence is ϵ' -independent of its predecessors for some $\epsilon' \geq \epsilon$.

The distributional eluder dimension states that the bad event that a historically indistinguishable function f fails to match the real function on future data cannot happen indefinitely. In the worst case, after $\text{DEdim}(\mathcal{F}, \Pi, \epsilon)$ such events, all historically indistinguishable function matches the real function with ϵ -precision. Let $\Delta_{\mathcal{X}} = \{\delta_x | \delta_x(y) = \delta(x - y), x, y \in \mathcal{X}\}$ be the set of Dirac delta distribution centered around each $x \in \mathcal{X}$. The distributional eluder dimension equals the original eluder dimension when $\Pi = \Delta_{\mathcal{X}}$. In RL, the distribution set we are concerned with is the distributions induced by the models.

Definition 4 (Probability measures induced by models). Let \mathcal{M} be a set of models and M^* be the real MDP. The probability measures induced by \mathcal{M} and M^* are $\Pi_{\mathcal{M}}^{M^*} = \{\rho^{M, M^*} = \frac{1}{H} \sum_{h=1}^H \rho_h^{M, M^*} | M \in \mathcal{M}\}$, where ρ_h^{M, M^*} denotes the state-action distribution at period h induced by executing an optimal policy π of M in the MDP M^* .

Although $\Pi_{\mathcal{M}}^{M^*}$ is enough for measuring the complexity of the function class, it is not amenable to analytical analysis. Instead, we will establish an upper bound for the distributional eluder dimension which holds for any class of probability measures.

C.2 Derivation of the Covering Number

Let \mathcal{F} be the function class defined in Assumption 2. For all $W_1, W_2 \in \mathcal{W} = \{W \in \mathbb{R}^{(m+1) \times d} | \|W\|_F \leq r_w\}$, we have

$$\begin{aligned} \|f_{W_1} - f_{W_2}\|_{\infty} &= \|(W_1 - W_2)\phi\|_{\infty} \\ &= \text{ess sup}_{x \in \mathcal{X}} \|(W_1 - W_2)\phi(x)\|_2 \\ &\leq \text{ess sup}_{x \in \mathcal{X}} \|W_1 - W_2\|_2 \|\phi(x)\|_2 \\ &\leq \|W_1 - W_2\|_2 r_{\phi} \\ &\leq \|W_1 - W_2\|_F r_{\phi}. \end{aligned} \tag{17}$$

Therefore, an α/r_{ϕ} -covering of \mathcal{W} w.r.t. the Frobenius norm is a α -covering of \mathcal{F} . Since the Frobenius norm of a matrix equals the 2-norm of the corresponding flatten vector, and the α/r_{ϕ} -covering number of $(m+1)d$ -dimensional ball $\{w \in \mathbb{R}^{(m+1) \times d} | \|w\|_2 \leq r_w\}$ w.r.t. 2-norm is

$O((r_\phi r_w/\alpha)^{(m+1)d})$, we have

$$\log N(\mathcal{F}, \alpha, \|\cdot\|_\infty) = O(md \log(r_\phi r_w/\alpha)). \quad (18)$$

C.3 Derivation of Distributional Eluder Dimension

Let Π be a set of probability measures on the measurable space (\mathcal{X}, Σ) . Define $\mathcal{W} = \{W \in \mathbb{R}^{(m+1) \times d} \mid \|W\|_F \leq r_w\}$.

For any $f_{W_1}, f_{W_2} \in \mathcal{F}$ and $\mu \in \Pi$, we have

$$\begin{aligned} \mu(\|f_{W_1} - f_{W_2}\|_2^2) &= \int \|(W_1 - W_2)\phi(x)\|_2^2 d\mu(x) \\ &= \int \text{Tr}((W_1 - W_2)\phi(x)\phi(x)^\top(W_1 - W_2)^\top) d\mu(x) \\ &= \text{Tr}((W_1 - W_2)\mu(\phi\phi^\top)(W_1 - W_2)^\top), \end{aligned} \quad (19)$$

where the third equality follows from the linearity of trace. Then, by definition, the distributional eluder dimension $\text{DEdim}(\mathcal{F}, \Pi, \epsilon)$ is the longest sequence $\mu_1, \dots, \mu_n \in \Pi$ such that

$$w_k = \sup \left\{ \text{Tr}(W\mu_k(\phi\phi^\top)W^\top) \mid \text{Tr}(W\Phi_{k-1}W^\top) \leq \epsilon'^2, W \in \tilde{\mathcal{W}} \right\} > \epsilon'^2 \quad (20)$$

holds for some $\epsilon' \geq \epsilon$ and all $k \in [n]$, where $\Phi_k = \sum_{i=1}^k \mu_i(\phi\phi^\top)$, and $\tilde{\mathcal{W}} = \mathcal{W} - \mathcal{W} = \{W_1 - W_2 \mid W_1, W_2 \in \mathcal{W}\}$. For any $W_1, W_2 \in \mathcal{W}$, it follows by the triangle inequality that $\|W_1 - W_2\|_F \leq \|W_1\|_F + \|W_2\|_F \leq 2r_w$. Let $B_k = \Phi_k + \lambda I$. Notice that

$$\begin{aligned} \epsilon'^2 < w_k &\leq \sup \left\{ \text{Tr}(W\mu_k(\phi\phi^\top)W^\top) \mid \text{Tr}(W\Phi_{k-1}W^\top) \leq \epsilon'^2, \text{Tr}(W^2) \leq 4r_w^2 \right\} \\ &\leq \sup \left\{ \text{Tr}(W\mu_k(\phi\phi^\top)W^\top) \mid \text{Tr}(WB_{k-1}W^\top) \leq \epsilon'^2 + 4\lambda r_w^2 \right\}, \end{aligned} \quad (21)$$

where the second inequality follows because $W^\top B_k W = W^\top \Phi_{k-1} W + \lambda W^\top W \leq \epsilon'^2 + 4\lambda r_w^2$. The solution of this constrained optimization problem is given by the following lemma.

Lemma C.1 (Maximizing trace under trace constraints). Let \mathcal{H} be a Hilbert space. Suppose that $W : \mathcal{H} \rightarrow \mathbb{R}^m$ is a linear operator, $A : \mathcal{H} \rightarrow \mathcal{H}$ is a positive semidefinite linear operator, and $B : \mathcal{H} \rightarrow \mathcal{H}$ is a positive definite linear operator. Let $w = \max_W \text{Tr}(WAW^\top)$ s.t. $\text{Tr}(WBW^\top) \leq \epsilon^2$. Then, $w \leq \epsilon^2 \text{Tr}(B^{-1}A)$.

Proof. The Lagrangian function of the constrained optimization is $\mathcal{L}(W, \lambda) = -\text{Tr}(WAW^\top) + \lambda(\text{Tr}(WBW^\top) - \epsilon^2)$. The KKT condition gives

$$\begin{cases} \lambda \geq 0 \\ \lambda(\text{Tr}(WBW^\top) - \epsilon^2) = 0 \\ \lambda BW^\top = AW^\top. \end{cases} \quad (22)$$

If $\lambda = 0$, then we have $AW^\top = 0$ and, therefore, $\text{Tr}(WAW^\top) = 0$. When $\lambda > 0$, $\text{Tr}(WAW^\top) = \lambda \text{Tr}(WBW^\top) = \lambda \epsilon^2$. Since $\lambda BW^\top = AW^\top$, we conclude that λ is the eigenvalue of $B^{-1}A$ and is at most $\|B^{-1}A\|_2 \geq 0$.

In conclusion, we have

$$\begin{aligned} w &\leq \epsilon^2 \|B^{-1}A\|_2 \\ &\leq \epsilon^2 \|B^{-1}A\|_{\text{HS}} = \epsilon^2 \sqrt{\text{Tr}(B^{-1}A(B^{-1}A)^\top)} = \epsilon^2 \sqrt{\text{Tr}((B^{-1}A)^2)} \\ &\leq \epsilon^2 \text{Tr}(B^{-1}A), \end{aligned} \quad (23)$$

where the second inequality holds because the Hilbert-Schmidt norm of a linear operator is greater than its operator norm, and the third inequality holds for positive semidefinite operators by the definition of trace. \square

According to Lemma C.1, $\epsilon'^2 < w_k \leq (\epsilon'^2 + 4\lambda r_W^2) \text{Tr}(B_{k-1}^{-1} \mu_k(\phi\phi^\top))$. This implies that

$$\text{Tr}(B_{k-1}^{-1} \mu_k(\phi\phi^\top)) > \frac{\epsilon'^2}{\epsilon'^2 + 4\lambda r_W^2}. \quad (24)$$

Since $\mu_k(\phi\phi^\top)$ is positive semidefinite, we can rewrite it by Cholesky decomposition as $\mu_k(\phi\phi^\top) = \phi_k \phi_k^\top$. Consider the case where the feature map ϕ is d -dimensional. By the generalized matrix determinant lemma,

$$\begin{aligned} \det B_n &= \det(I + \phi_n^\top B_{n-1}^{-1} \phi_n) \det B_{n-1} \\ &= \dots = \det(\lambda I) \prod_{i=1}^n \det(I + \phi_i^\top B_{i-1}^{-1} \phi_i) \\ &= \lambda^d \prod_{i=1}^n \det(I + \phi_i^\top B_{i-1}^{-1} \phi_i). \end{aligned} \quad (25)$$

For a positive semidefinite matrix A , the characteristic polynomial is $\det(tI - A) = \prod_{i=1}^d (t - \lambda_i)$, where $\lambda_1, \dots, \lambda_d$ are the eigenvalues of A . Selecting $t = -1$, it follows that

$$\begin{aligned} (-1)^d \det(I + A) &= (-1)^d \prod_{i=1}^d (1 + \lambda_i) \\ \Rightarrow \det(I + A) &\geq 1 + \text{Tr}(A). \end{aligned} \quad (26)$$

Equations (24), (25), and (26) imply that

$$\det B_n > \lambda^d \left(1 + \frac{\epsilon'^2}{\epsilon'^2 + 4\lambda r_W^2}\right)^n. \quad (27)$$

The determinant $\det B_n$ is upper bounded by the AM-GM inequality,

$$\begin{aligned} \det B_n &\leq \left(\frac{\text{Tr}(B_n)}{d}\right)^d = \left(\frac{\text{Tr}(\lambda I) + \sum_{i=1}^n \text{Tr}(\phi_i \phi_i^\top)}{d}\right)^d \\ &\leq \left(\frac{\lambda d + n r_\phi^2}{d}\right)^d. \end{aligned} \quad (28)$$

Combining Equations (27) and (28) and setting $\lambda = \epsilon'^2 / (4r_W^2)$, we get

$$\begin{aligned} \left(1 + \frac{\epsilon'^2}{\epsilon'^2 + 4\lambda r_W^2}\right)^{\frac{n}{d}} &\leq 1 + \frac{r_\phi^2 n}{\lambda d} \\ \Rightarrow \left(1 + \frac{1}{2}\right)^{\frac{n}{d}} &\leq 1 + \frac{4r_W^2 r_\phi^2 n}{\epsilon'^2 d}. \end{aligned} \quad (29)$$

For any $\alpha > 0$, $\beta \geq 0$, we show that if $(1 + \alpha)^k \leq 1 + \beta k$, then $k \leq \frac{\beta+1}{\beta} \frac{2 \log(\beta+1) - \log \log(\alpha+1)}{\log(\alpha+1)}$. Note that

$$k \log(1 + \alpha) \leq \log(1 + \beta k) = \log k + \log(1/k + \beta). \quad (30)$$

If $k \geq 1$, we have

$$k \log(1 + \alpha) \leq \log k + \log(1/k + \beta) \leq \log k + \log(1 + \beta). \quad (31)$$

Since $\log x \leq \frac{x}{e}$ for any $x > 0$, we have

$$k \log(1 + \alpha) \leq \log(k/y) + \log(y) + \log(1 + \beta) \leq k/(ey) + \log y + \log(1 + \beta). \quad (32)$$

holds for any $y > 0$. Selecting $y = \frac{1}{\log(1+\alpha)}$ yields

$$\begin{aligned} k &\leq \frac{\log(1 + \beta) - \log \log(1 + \alpha)}{\log(1 + \alpha) - \frac{\log(1+\alpha)}{e}} \\ &= \frac{e}{e-1} \frac{\log(1 + \beta) - \log \log(1 + \alpha)}{\log(1 + \alpha)} \\ &\leq \frac{e}{e-1} \frac{1 + \alpha}{\alpha} \log \frac{(1 + \beta)(1 + \alpha)}{\alpha}, \end{aligned} \quad (33)$$

where the last inequality holds because $\log(1+x) \geq \frac{x}{x+1}$.

If $k < 1$, the above inequality also holds since according to Equation (30),

$$k \leq \frac{\log(1+\beta)}{\log(1+\alpha)} \leq \frac{1+\alpha}{\alpha} \log(1+\beta) \leq \frac{e}{e-1} \frac{1+\alpha}{\alpha} \log \frac{(1+\beta)(1+\alpha)}{\alpha}. \quad (34)$$

By setting $\alpha = \frac{1}{2}$ and $\beta = \frac{4r_W^2 r_\phi^2}{\epsilon'^2}$, it follows that

$$\begin{aligned} \text{DEdim}(\mathcal{F}, \Pi, \epsilon) = n &\leq d \frac{3e}{e-1} \log \frac{12r_W^2 r_\phi^2}{\epsilon'^2} = d \frac{6e}{e-1} \log \frac{2\sqrt{3}r_W r_\phi}{\epsilon'} \\ &\leq d \frac{6e}{e-1} \log \frac{2\sqrt{3}r_W r_\phi}{\epsilon}. \end{aligned} \quad (35)$$

D Bounding the BOO Regret

In this section, we denote the squared error of K episodes as $L_{2,K}(f) = \sum_{k=1}^K \sum_{h=1}^H \|f(x_{kh}) - y_{kh}\|_2^2$ and empirical 2-norm as $\|g\|_{2,E_K} = \sqrt{\sum_{k=1}^K \sum_{h=1}^H \|g(x_{kh})\|_2^2}$. Our analysis is based on the following regret decomposition:

$$\Delta_k = \underbrace{\mathbb{E}_{\pi_k \sim \mathfrak{A}(\mathcal{H}_k)} [V_1^*(s_1) - V_1^k(s_1) | M^*]}_{\Delta_k^{\text{opt}}} + \underbrace{\mathbb{E}_{\pi_k \sim \mathfrak{A}(\mathcal{H}_k)} [V_1^k(s_1) - V_1^{\pi_k, M^*}(s_1) | M^*]}_{\Delta_k^{\text{conc}}}. \quad (36)$$

The expectation can be eliminated assuming the algorithm deterministically settles on a solution π_k . We will bound the sum of optimism regret $\sum_{k=1}^K \Delta_k^{\text{opt}}$ and concentration regret $\sum_{k=1}^K \Delta_k^{\text{conc}}$ separately.

D.1 Concentration of Squared Error

This subsection establishes the concentration of the model squared error around the expected model difference based on the martingale concentration inequalities in [22].

Lemma D.1 (Martingale concentration [22]). Consider random variables Z_k adapted to the filtration $(\mathcal{H}_k : k \in \mathbb{Z}^+)$. Assume $\mathbb{E}[\exp(\lambda Z_k)]$ is finite for all λ . For all $x \geq 0$ and $\lambda \geq 0$, we have

$$\Pr \left(\left(\sum_{k=1}^K \lambda Z_k \leq \sum_{k=1}^K (\lambda \mu_k + \psi_k(\lambda)) + x \right) \forall K \in \mathbb{Z}^+ \right) \geq 1 - e^{-x}, \quad (37)$$

where $\mu_k = \mathbb{E}[Z_k | \mathcal{H}_k]$ is the conditional mean, and $\psi_k(\lambda) = \log \mathbb{E}[\exp(\lambda(Z_k - \mu_k)) | \mathcal{H}_k]$ is the conditional cumulant generating function of $Z_k - \mu_k$.

Based on Lemma D.1, it is already proved as an intermediate result in [22, 47] that the difference of squared error between the true model function and any other model function will not deviate too much from the squared difference of them.

Lemma D.2 (Concentration of the model squared error around the empirical model difference [22, 47]). Let \mathcal{F} be a set of functions such that $\|f(x)\|_2 \leq r_{\mathcal{F}}$ for all $x \in \mathcal{X}$ and $f \in \mathcal{F}$. Given a set of data $\{(x_{11}, y_{11}), \dots, (x_{KH}, y_{KH})\}$ from K episodes generated by $y_{kh} = f^*(x_{kh}) + \epsilon_{kh}$, $k \in [K]$, $h \in [H]$, where ϵ_{kh} is a σ -sub-Gaussian noise, then

$$\Pr \left(\left(L_{2,K}(f) - L_{2,K}(f^*) \geq \frac{1}{2} \|f - f^*\|_{2,E_K}^2 - \beta_K^{**}(\mathcal{F}, \delta, \alpha) \right) \forall K \in \mathbb{Z}^+, f \in \mathcal{F} \right) \geq 1 - 2\delta, \quad (38)$$

where $\delta > 0$, $\alpha > 0$, and $\beta_K^{**}(\mathcal{F}, \delta, \alpha) = 4\sigma^2 \log(N(\mathcal{F}, \alpha, \|\cdot\|_\infty)/\delta) + \alpha KH(8r_{\mathcal{F}} + \sqrt{8\sigma^2 \log(4K^2 H^2/\delta)})$.

We will further show that the empirical model difference $\|f - f^*\|_{2,E_K}$ is concentrated around the expected model difference. The following lemma is particularly useful for bounding the cumulant generating function of non-negative random variables with bounded second moments.

Lemma D.3 (Cumulant generating function for non-negative random variable [48]). For a random variable $X \geq 0$ with $\mathbb{E}[X^2] < \infty$, we have

$$\log \mathbb{E} [e^{-\lambda X}] \leq -\lambda \mathbb{E}[X] + \frac{\lambda^2}{2} \mathbb{E}[X^2], \quad (39)$$

where $\lambda \geq 0$.

We now establish the concentration result for a single function.

Lemma D.4 (Concentration of the empirical model difference around the expected model difference for a single function). Let f be a function satisfying $\|f(x)\|_2 \leq r$ for all $x \in \mathcal{X}$. For any $\delta > 0$, with probability at least $1 - \delta$,

$$-\|f - f^*\|_{2, E_K}^2 \leq \frac{1}{2} \sum_{k=1}^K \mu_k + 4Hr^2 \log \frac{1}{\delta} \quad (40)$$

simultaneously for all $K \in \mathbb{Z}^+$, where $\mu_k = -\mathbb{E} \left[\sum_{h=1}^H \|f^*(x_{kh}) - f(x_{kh})\|_2^2 \middle| \mathcal{H}_k \right]$.

Proof. Define $Z_k = -\sum_{h=1}^H \|f^*(x_{kh}) - f(x_{kh})\|_2^2$. The conditional mean of Z_k is $\mu_k = \mathbb{E}[Z_k | \mathcal{H}_k] = -\mathbb{E} \left[\sum_{h=1}^H \|f^*(x_{kh}) - f(x_{kh})\|_2^2 \middle| \mathcal{H}_k \right]$. The conditional cumulant generating function of the centered random variable $Z_k - \mu_k$ is

$$\begin{aligned} \psi_k(\lambda) &= \log \mathbb{E}[\exp(\lambda[Z_k - \mu_k]) | \mathcal{H}_k] \\ &= \log \mathbb{E}[\exp(\lambda Z_k) | \mathcal{H}_k] - \lambda \mu_k \\ &\leq \lambda \mathbb{E}[Z_k | \mathcal{H}_k] + \frac{\lambda^2}{2} \mathbb{E}[Z_k^2 | \mathcal{H}_k] - \lambda \mu_k = \frac{\lambda^2}{2} \mathbb{E}[Z_k^2 | \mathcal{H}_k], \end{aligned} \quad (41)$$

where the inequality holds according to Lemma D.3. Since $\sum_{h=1}^H \|f^*(x_{kh}) - f(x_{kh})\|_2^2 \leq \sum_{h=1}^H (\|f^*(x_{kh})\|_2 + \|f(x_{kh})\|_2)^2 \leq \sum_{h=1}^H (2r)^2 = 4Hr^2$, we can bound the conditional second moment of Z_k with its conditional mean:

$$\begin{aligned} \mathbb{E}[Z_k^2 | \mathcal{H}_k] &= \mathbb{E} \left[\left(\sum_{h=1}^H \|f^*(x_{kh}) - f(x_{kh})\|_2^2 \right)^2 \middle| \mathcal{H}_k \right] \\ &\leq 4Hr^2 \mathbb{E} \left[\sum_{h=1}^H \|f^*(x_{kh}) - f(x_{kh})\|_2^2 \middle| \mathcal{H}_k \right] = -4Hr^2 \mu_k. \end{aligned} \quad (42)$$

Selecting $\lambda = \frac{1}{4Hr^2}$ and $x = \log \frac{1}{\delta}$ in Lemma D.1 implies that

$$\begin{aligned} \sum_{k=1}^K Z_k &\leq \sum_{k=1}^K (1 - 2\lambda Hr^2) \mu_k + \frac{x}{\lambda} \\ &\leq \sum_{k=1}^K \frac{1}{2} \mu_k + 4Hr^2 \log \frac{1}{\delta}, \end{aligned} \quad (43)$$

which completes the proof. \square

Lemma D.5 (Concentration of the empirical model difference around the expected model difference). Let \mathcal{F} be a set of functions such that $\|f(x)\|_2 \leq r_{\mathcal{F}}$ for all $x \in \mathcal{X}$ and $f \in \mathcal{F}$. For $\delta > 0$ and $\alpha > 0$,

$$\Pr \left(\left(-\frac{1}{2} \|f - f^*\|_{2, E_K}^2 \leq \frac{1}{4} \sum_{k=1}^K \mu_k + \beta_K^\dagger(\mathcal{F}, \delta, \alpha) \right) \forall K \in \mathbb{Z}^+, f \in \mathcal{F} \right) \geq 1 - \delta, \quad (44)$$

where $\mu_k = -\mathbb{E} \left[\sum_{h=1}^H \|f^*(x_{kh}) - f(x_{kh})\|_2^2 \middle| \mathcal{H}_k \right]$, and $\beta_K^\dagger(\mathcal{F}, \delta, \alpha) = 2Hr_{\mathcal{F}}^2 \log(N(\mathcal{F}, \alpha, \|\cdot\|_\infty)/\delta) + \frac{3}{4}KH(\alpha^2 + 4\alpha r_{\mathcal{F}})$.

Proof. Let $\mathcal{F}^\alpha \subset \mathcal{F}$ be an α -cover of \mathcal{F} w.r.t. the ℓ_2 norm in the sense that for any $f \in \mathcal{F}$ there is an $f^\alpha \in \mathcal{F}^\alpha$ such that $\|f - f^\alpha\|_\infty = \sup_{x \in \mathcal{X}} \|f^\alpha(x) - f(x)\|_2 \leq \alpha$. Applying Lemma D.4 for all $f^\alpha \in \mathcal{F}^\alpha$, it follows by a union bound that

$$-\|f^\alpha - f^*\|_{2,E_K}^2 \leq -\frac{1}{2} \sum_{k=1}^K \mathbb{E} \left[\sum_{h=1}^H \|f^*(x_{kh}) - f^\alpha(x_{kh})\|_2^2 \middle| \mathcal{H}_k \right] + 4Hr_{\mathcal{F}}^2 \log \frac{|\mathcal{F}^\alpha|}{\delta} \quad (45)$$

for all $f^\alpha \in \mathcal{F}^\alpha$ and $t \in \mathbb{Z}^+$ with probability $1 - \delta$. Then, for any $f \in \mathcal{F}$,

$$\begin{aligned} -\|f - f^*\|_{2,E_K}^2 &\leq -\frac{1}{2} \sum_{k=1}^K \mathbb{E} \left[\sum_{h=1}^H \|f^*(x_{kh}) - f(x_{kh})\|_2^2 \middle| \mathcal{H}_k \right] + 4Hr_{\mathcal{F}}^2 \log \frac{|\mathcal{F}^\alpha|}{\delta} \\ &\quad + \text{DiscErr}(\alpha), \end{aligned} \quad (46)$$

where the discretization error is

$$\begin{aligned} \text{DiscErr}(\alpha) &= \min_{f^\alpha \in \mathcal{F}^\alpha} \left\{ \|f^\alpha - f^*\|_{2,E_K}^2 + \frac{1}{2} \sum_{k=1}^K \mathbb{E} \left[\sum_{h=1}^H \|f^*(x_{kh}) - f(x_{kh})\|_2^2 \middle| \mathcal{H}_k \right] \right. \\ &\quad \left. - \|f - f^*\|_{2,E_K}^2 - \frac{1}{2} \sum_{k=1}^K \mathbb{E} \left[\sum_{h=1}^H \|f^*(x_{kh}) - f^\alpha(x_{kh})\|_2^2 \middle| \mathcal{H}_k \right] \right\}. \end{aligned} \quad (47)$$

For an f^α satisfying $\|f^\alpha - f\|_\infty \leq \alpha$,

$$\begin{aligned} &|\|f^\alpha(x) - f^*(x)\|_2^2 - \|f(x) - f^*(x)\|_2^2| \\ &= |\|f^\alpha(x) - f(x)\|_2^2 + 2\langle f^\alpha(x) - f(x), f(x) \rangle + 2\langle f(x) - f^\alpha(x), f^*(x) \rangle| \\ &\leq \alpha^2 + 4\alpha r_{\mathcal{F}}, \end{aligned} \quad (48)$$

via the Cauchy–Schwarz inequality. Summing over all time steps, we have

$$|\|f - f^*\|_{2,E_K}^2 - \|f^\alpha - f^*\|_{2,E_K}^2| \leq KH(\alpha^2 + 4\alpha r_{\mathcal{F}}), \quad (49)$$

and

$$\left| \sum_{k=1}^K \mathbb{E} \left[\sum_{h=1}^H \|f^*(x_{kh}) - f^\alpha(x_{kh})\|_2^2 - \|f^*(x_{kh}) - f(x_{kh})\|_2^2 \middle| \mathcal{H}_k \right] \right| \leq KH(\alpha^2 + 4\alpha r_{\mathcal{F}}). \quad (50)$$

Therefore, the discretization error is bounded by

$$|\text{DiscErr}(\alpha)| \leq \frac{3}{2} KH(\alpha^2 + 4\alpha r_{\mathcal{F}}). \quad (51)$$

The proof is completed by combining Equations (46) and (51). \square

Theorem D.1 (Concentration of the model squared error around the expected model difference). *Let \mathcal{F} be a set of functions such that $\|f(x)\|_2 \leq r_{\mathcal{F}}$ for all $x \in \mathcal{X}$ and $f \in \mathcal{F}$. Given a set of data $\{(x_{11}, y_{11}), \dots, (x_{KH}, y_{KH})\}$ from K episodes generated by $y_{kh} = f^*(x_{kh}) + \epsilon_{kh}$, $k \in [K]$, $h \in [H]$, where ϵ_{kh} is a σ -sub-Gaussian noise, then*

$$\Pr \left(\left(L_{2,K}(f^*) - L_{2,K}(f) \leq \frac{1}{4} \sum_{k=1}^K \mu_k + \beta_K^\ddagger(\mathcal{F}, \delta, \alpha) \right) \forall K \in \mathbb{Z}^+, f \in \mathcal{F} \right) \geq 1 - 3\delta, \quad (52)$$

where $\delta > 0$, $\alpha > 0$, $\mu_k = -\mathbb{E} \left[\sum_{h=1}^H \|f^*(x_{kh}) - f(x_{kh})\|_2^2 \middle| \mathcal{H}_k \right]$, and $\beta_K^\ddagger(\mathcal{F}, \delta, \alpha) = (4\sigma^2 + 2Hr_{\mathcal{F}}^2) \log(N(\mathcal{F}, \alpha, \|\cdot\|_\infty)/\delta) + \alpha KH(\frac{3}{4}\alpha + 11r_{\mathcal{F}} + \sqrt{8\sigma^2 \log(4K^2 H^2/\delta)})$.

Proof. The proof is carried out by combining Lemmas D.2 and D.5. \square

The optimal asymptotic scaling of the confidence set is specified in Lemma D.6.

Lemma D.6 (The asymptotic analysis of confidence set). For any fixed class of functions \mathcal{F} satisfying $\log N(\mathcal{F}, \alpha, \|\cdot\|_\infty) \leq C(1/\alpha)^{c_1}(\log(1/\alpha))^{c_2}$, setting $\alpha_k = (\frac{C\sigma}{H} + \frac{Cr_{\mathcal{F}}^2}{2\sigma})^{1/(c_1+1)}k^{-1/(c_1+1)}(\log k)^{(2c_2-1)/(2c_1+2)}$,

$$\beta_k^\dagger(\mathcal{F}, \delta, \alpha_k) = O\left(4\sigma H \left(\frac{C\sigma}{H} + \frac{Cr_{\mathcal{F}}^2}{2\sigma}\right)^{1/(c_1+1)} k^{c_1/(c_1+1)}(\log k)^{(c_1+2c_2)/(2c_1+2)}\right). \quad (53)$$

Proof. Setting $\alpha = \alpha_k$ in β^\dagger , it is easily checked that

$$\begin{aligned} \beta_k^\dagger(\mathcal{F}, \delta, \alpha_k) &\leq (4\sigma^2 + 2Hr_{\mathcal{F}}^2)C(1/\alpha_k)^{c_1} \log(1/\alpha_k)^{c_2} + (4\sigma^2 + 2Hr_{\mathcal{F}}^2) \log \frac{1}{\delta} \\ &\quad + \alpha_k k H \left(\frac{3}{4}\alpha_k + 11r_{\mathcal{F}} + \sqrt{8\sigma^2 \log(4k^2 H^2/\delta)}\right) \\ &= O\left((4\sigma^2 + 2Hr_{\mathcal{F}}^2)C(1/\alpha_k)^{c_1} \log(1/\alpha_k)^{c_2} + \alpha_k k H \sqrt{8\sigma^2 \log(4k^2 H^2/\delta)}\right) \\ &= O\left(4\sigma H \left(\frac{C\sigma}{H} + \frac{Cr_{\mathcal{F}}^2}{2\sigma}\right)^{1/(c_1+1)} k^{c_1/(c_1+1)}(\log k)^{(c_1+2c_2)/(2c_1+2)}\right), \end{aligned} \quad (54)$$

when $k \rightarrow \infty$. \square

Lemma D.7 (The confidence set of linear function classes). Let \mathcal{F} be the function class defined in Assumption 2. By setting $\alpha_k = md(\frac{\sigma}{H} + \frac{r_w^2 r_\phi^2}{2\sigma})k^{-1}(\log k)^{\frac{1}{2}}$, we have

$$\beta_k^\dagger(\mathcal{F}, \delta, \alpha_k) = O\left(\left(\sigma^2 md + \frac{Hmdr_w^2 r_\phi^2}{2}\right) \log k\right). \quad (55)$$

Proof. According to Assumption 2 and the definition of operator norm, we have

$$\|f(x)\|_2 = \|W\phi(x)\|_2 \leq \|W\|_2 \|\phi(x)\|_2 = r_w r_\phi \quad (56)$$

for all $x \in \mathcal{X}$ and $f \in \mathcal{F}$.

According to Theorem D.1, we have the concentration of the squared error around the expected model difference with probability $1 - 3\delta$ for all $f \in \mathcal{F}$, $K \in \mathbb{Z}^+$,

$$L_{2,K}(f^*) - L_{2,K}(f) \leq \frac{1}{4} \sum_{k=1}^K \mu_k + \beta_K^\dagger(\mathcal{F}, \delta, \alpha), \quad (57)$$

where $\mu_k = -\mathbb{E}\left[\sum_{h=1}^H \|f^*(x_{kh}) - f(x_{kh})\|_2^2 \middle| \mathcal{H}_k\right]$.

According to the Equation (18), the optimal scaling of β_K^\dagger is specified in Lemma D.6 by setting $\alpha_k = md(\frac{\sigma}{H} + \frac{r_w^2 r_\phi^2}{2\sigma})k^{-1}(\log k)^{1/2}$ as follows:

$$\beta_k^\dagger(\mathcal{F}, \delta, \alpha_k) = O\left(\left(\sigma^2 md + \frac{Hmdr_w^2 r_\phi^2}{2}\right) \log k\right). \quad (58)$$

\square

D.2 Concentration Regret

The following Lemma establishes the connection between the value difference and expected model differences by the one-step value Lipschitz continuity.

Lemma D.8 (Connection between value differences and model differences). Given a set of models \mathcal{M} expressed by a set of model functions \mathcal{F} with additive σ -sub-Gaussian noises, the value difference of the policy π_1 on two models $M_1, M_2 \in \mathcal{M}$ is upper bounded by the expected difference of model functions,

$$\left|V_1^{\pi_1, M_1}(s_1) - V_1^{\pi_1, M_2}(s_1)\right| \leq L \mathbb{E}_{\tau \sim \pi_1, M_2} \left[\sum_{h=1}^H \|f^{M_1}(x_{kh}) - f^{M_2}(x_{kh})\|_2 \right], \quad (59)$$

where π_1 is the optimal policy of M_1 , and f^M is the model function of M .

Proof. The value difference can be rewritten as the expected model difference [17],

$$\begin{aligned}
& \left| V_1^{\pi_1, M_1}(s_1) - V_1^{\pi_1, M_2}(s_1) \right| \\
&= \left| \bar{R}^1(x_1) + P_{x_1}^1(V_2^{\pi_1, M_1}) - \bar{R}^2(x_1) - P_{x_1}^2(V_2^{\pi_1, M_2}) \right| \\
&= \left| \bar{R}^1(x_1) - \bar{R}^2(x_1) + (P_{x_1}^1 - P_{x_1}^2)(V_2^{\pi_1, M_1}) + P_{x_1}^2(V_2^{\pi_1, M_1} - V_2^{\pi_1, M_2}) \right| \\
&= \left| \bar{R}^1(x_1) - \bar{R}^2(x_1) + (P_{x_1}^1 - P_{x_1}^2)(V_2^{\pi_1, M_1}) + \mathbb{E}_{s_2 \sim \pi, M_2} \left[V_2^{\pi_1, M_1}(s_2) - V_2^{\pi_1, M_2}(s_2) \right] \right| \quad (60) \\
&= \dots \\
&= \left| \mathbb{E}_{\tau \sim \pi_1, M_2} \left[\sum_{h=1}^H (\bar{R}^1(x_h) - \bar{R}^2(x_h)) + \sum_{h=1}^{H-1} (P_{x_h}^1 - P_{x_h}^2)(V_{h+1}^{\pi_1, M_1}) \right] \right|,
\end{aligned}$$

where $x_h = (s_h, \pi_1(s_h, h))$, $\bar{R}^i(x_h) = \bar{R}^{M_i}(x_h)$, and $P_{x_h}^i = P^{M_i}(\cdot | x_h)$.

The model difference is then upper bounded by one-step value Lipschitz continuity,

$$\begin{aligned}
\left| V_1^{\pi_1, M_1}(s_1) - V_1^{\pi_1, M_2}(s_1) \right| &\leq \mathbb{E}_{\tau \sim \pi_1, M_2} \left[\sum_{h=1}^H L_{x_{kh}, V_{h+1}^{\pi_1, M_1}} \|f^{M_1}(x_{kh}) - f^{M_2}(x_{kh})\|_2 \right] \\
&\leq L \mathbb{E}_{\tau \sim \pi_1, M_2} \left[\sum_{h=1}^H \|f^{M_1}(x_{kh}) - f^{M_2}(x_{kh})\|_2 \right]. \quad (61)
\end{aligned}$$

The proof is completed. \square

Lemma D.9 (Bound on the number of large concentration regret). Consider the set of models \mathcal{M} corresponding to model functions \mathcal{F} . Let $(\beta_K > 0 | K \in \mathbb{Z}^+)$ be a nondecreasing sequence such that $L_{2,K}(f^*) - L_{2,K}(f) \leq \frac{1}{4} \sum_{k=1}^K \mu_k + \beta_K$ for all $f \in \mathcal{F}$ and $K \in \mathbb{Z}^+$, where $\mu_k = -\mathbb{E} \left[\sum_{h=1}^H \|f^*(x_{kh}) - f(x_{kh})\|_2^2 | \mathcal{H}_k \right]$. The log-prior probability is uniformly bounded by Z . For the BOO algorithm, the number of episodes where a large concentration regret is incurred is limited,

$$\sum_{k=1}^K \mathbb{I}(\Delta_k^{\text{conc}} \geq \eta) \leq N(K, \eta / (HL)) \dim_{DE}(\mathcal{F}, \Pi_{\mathcal{M}}^{M^*}, \eta / (HL)), \quad (62)$$

where

$$N(K, \epsilon) = \frac{16Z + \frac{8HL}{\lambda_K} \epsilon + 4\sigma^2 \beta_{K-1}}{H\epsilon^2 \sigma^2}. \quad (63)$$

Proof. Notice that a model function is selected only when it optimizes the BOO objective, which implies that

$$\frac{1}{\lambda_k} V_1^k(s_1) + \log \Pr(M_k) + \log \Pr(\mathcal{H}_k | M_k) \geq \frac{1}{\lambda_k} V_1^*(s_1) + \log \Pr(M^*) + \log \Pr(\mathcal{H}_k | M^*), \quad (64)$$

for any $k \in [K]$. It follows that

$$\begin{aligned}
0 &\leq \frac{1}{\lambda_k} (V_1^k(s_1) - V_1^*(s_1)) + \log \Pr(M_k) - \log \Pr(M^*) + \log \Pr(\mathcal{H}_k | M_k) - \log \Pr(\mathcal{H}_k | M^*) \\
&= \frac{1}{\lambda_k} (V_1^k(s_1) - V_1^*(s_1)) + \log \Pr(M_k) - \log \Pr(M^*) + \frac{\sigma^2}{2} (L_{2,k-1}(f^*) - L_{2,k-1}(f^{M_k})) \\
&\leq \frac{1}{\lambda_k} (V_1^k(s_1) - V_1^*(s_1)) + 2Z + \frac{\sigma^2}{2} (L_{2,k-1}(f^*) - L_{2,k-1}(f^{M_k})) \\
&\leq \frac{1}{\lambda_k} (V_1^k(s_1) - V_1^{\pi^k, M^*}(s_1)) + 2Z + \frac{\sigma^2}{2} (L_{2,k-1}(f^*) - L_{2,k-1}(f^{M_k})) \\
&= \frac{1}{\lambda_k} \Delta_k^{\text{conc}} + 2Z + \frac{\sigma^2}{2} (L_{2,k-1}(f^*) - L_{2,k-1}(f^{M_k})) \\
&\leq \frac{1}{\lambda_k} \Delta_k^{\text{conc}} + 2Z + \frac{\sigma^2}{8} \sum_{i=1}^{k-1} \mu_i + \frac{\sigma^2}{2} \beta_{k-1}, \quad (65)
\end{aligned}$$

where the first equality is obtained by substituting in the Equation (14), and the second equality follows .

Suppose that ρ^{M_k} is $(\Delta_k^{\text{conc}}/(HL))$ -dependent on n_k disjoint subsequences of $(\rho^{M_1}, \dots, \rho^{M_{k-1}})$. Notice that

$$\mathbb{E}_{x \sim \rho^{M_k}} [\|f(x) - f^*(x)\|_2^2] \geq (\mathbb{E}_{x \sim \rho^{M_k}} [\|f(x) - f^*(x)\|_2])^2 \geq \left(\frac{\Delta_k^{\text{conc}}}{HL}\right)^2, \quad (66)$$

where the first inequality follows from the Jensen's inequality, and the second inequality is derived by Lemma D.8.

It follows by the definition of ϵ -dependence that

$$-\sum_{i=1}^{k-1} \mu_i = \sum_{i=1}^{k-1} H \mathbb{E}_{x \sim \rho^{M_i}} [\|f(x) - f^*(x)\|_2^2] > n_k H \left(\frac{\Delta_k^{\text{conc}}}{HL}\right)^2. \quad (67)$$

Equations (65) and (67) collectively suggest that

$$\begin{aligned} 2Z + \frac{1}{\lambda_k} \Delta_k^{\text{conc}} + \frac{\sigma^2}{2} \beta_{k-1} &> \frac{\sigma^2 n_k H}{8} \left(\frac{\Delta_k^{\text{conc}}}{HL}\right)^2 \\ \Rightarrow n_k &< N(k, \Delta_k^{\text{conc}}/(HL)), \end{aligned} \quad (68)$$

where $N(k, \epsilon)$ is defined as $N(k, \epsilon) = \frac{16Z + \frac{8HL}{\lambda_k} \epsilon + 4\sigma^2 \beta_{k-1}}{H\epsilon^2 \sigma^2}$.

The ϵ -errored subsequence of $(1, \dots, K)$ is denoted as $(a_1, \dots, a_{m_\epsilon})$, where $\Delta_{a_i}^{\text{conc}}/(HL) \geq \epsilon$ for $i \in [m_\epsilon]$. Let $(\rho^1, \dots, \rho^{m_\epsilon})$ be the corresponding sequence of probability measures. It follows that each probability measure in the sequence is ϵ -dependent on less than $N(K, \epsilon)$ disjoint subsequences among its predecessors. We show that the length of ϵ -errored subsequence is at most $N(K, \epsilon) \dim_{DE}(\mathcal{F}, \Pi_{\mathcal{M}}^{M^*}, \epsilon)$. To show this, consider a growing sequence (ρ^1, ρ^2, \dots) where each element is ϵ -dependent on less than n disjoint subsequences of its predecessors. Let $B_i = (\rho^i)$ for $i \in [n]$. By definition, the $(n+1)$ -th element ρ^{n+1} should be ϵ -independent of some B_j . Add ρ^{n+1} to B_j , and repeat this process for $i > n+1$. It follows by construction that each element in $B_i, i \in [n]$ is ϵ -independent of all its predecessors. Therefore, such a sequence can contain at most $\dim_{DE}(\mathcal{F}, \Pi_{\mathcal{M}}^{M^*}, \epsilon)$ elements according to the definition of the eluder dimension. Besides, after each sequence is filled, any $x \in \mathcal{X}$ must be ϵ -dependent on all $B_i, i \in [n]$ implying that the sequence (ρ^1, ρ^2, \dots) cannot grow any more. This completes our argument that

$$\sum_{k=1}^K \mathbb{I}(\Delta_k^{\text{conc}} \geq HL\epsilon) = m_\epsilon \leq N(K, \epsilon) \dim_{DE}(\mathcal{F}, \Pi_{\mathcal{M}}^{M^*}, \epsilon). \quad (69)$$

□

Theorem D.2 (Sum of the concentration regret). *Let \mathcal{F} be the function class defined in Assumption 2. The log-prior probability is uniformly bounded by Z . Let $\lambda_k = ck^{-v_1}(\log k)^{-v_2}$ for $0 < v_1 < 1$ or $v_1 = 0, v_2 > 0$, and some $c > 0$. We have*

$$\sum_{k=1}^K \Delta_k^{\text{conc}} = \begin{cases} O\left(HLdr_w r_\phi \sigma \sqrt{m(r_c+1)K} \log K\right) & \text{if } v_1 = \frac{1}{2}, v_2 = 0 \\ O\left(HL^2 d/(c\sigma^2)\right) K^{v_1} (\log K)^{v_2+1} & \text{if } v_1 > v_1^* \text{ or } v_1 = \frac{1}{2}, v_2 > 0 \\ O\left(HLdr_w r_\phi \sigma \sqrt{mK} (\log K)\right) & \text{if } v_1 < \frac{1}{2} \text{ or } v_1 = \frac{1}{2}, v_2 < 0, \end{cases} \quad (70)$$

where r_c is a constant determined by $c = \frac{\sqrt{(r_c+1)}}{r_c} \frac{2\sqrt{2}L}{r_w r_\phi \sigma^3 \sqrt{m}}$.

Proof. The sum of regret can be expressed as a Riemann-Stieltjes integral,

$$\begin{aligned} \sum_{k=1}^K \Delta_k^{\text{conc}} &= \lim_{b \rightarrow \infty} \lim_{n \rightarrow \infty} \sum_{i=1}^n \left(g\left(\frac{(i-1)}{n}b\right) - g\left(\frac{i}{n}b\right) \right) \eta \\ &= - \int_0^\infty \eta dg(\eta), \end{aligned} \quad (71)$$

where $g(\eta) = \sum_{k=1}^K \mathbb{I}(\Delta_k^{\text{conc}} \geq \eta)$.

The optimal scaling of β_K^\ddagger is specified in Lemma D.7, i.e.,

$$\beta_k^\ddagger(\mathcal{F}, \delta, \alpha_k) = O\left(\left(\sigma^2 md + \frac{Hmdr_w^2 r_\phi^2}{2}\right) \log k\right). \quad (72)$$

According to Lemma D.9, Equation (35), and Equation (72), we can get

$$\begin{aligned} \sum_{k=1}^K \mathbb{I}(\Delta_k^{\text{conc}} \geq \eta) &\leq N(K, \eta/(HL)) \dim_{DE}(\mathcal{F}, \Pi_{\mathcal{M}}^{M^*}, \eta/(HL)) \\ &\leq 4HL^2 \frac{4Z + \frac{2\eta}{\lambda_K} + \sigma^4 Hmdr_w^2 r_\phi^2 \log K}{\eta^2 \sigma^2} d \log(HL/\eta). \end{aligned} \quad (73)$$

Along with the fact that $\sum_{k=1}^K \mathbb{I}(\Delta_k^{\text{conc}} \geq \eta) \leq K$, we have, for all η , $g(\eta) \leq h(\eta) = \min\left(K, 4HL^2 \frac{4Z + \frac{2\eta}{\lambda_K} + \sigma^4 Hmdr_w^2 r_\phi^2 \log K}{\eta^2 \sigma^2} d \log(HL/\eta)\right)$.

Since

$$\int_0^\infty \eta d(h(\eta) - g(\eta)) = \eta(h(\eta) - g(\eta))|_0^\infty - \int_0^\infty (h(\eta) - g(\eta)) d\eta \leq 0, \quad (74)$$

we have

$$\begin{aligned} \sum_{k=1}^K \Delta_k^{\text{conc}} &\leq - \int_0^\infty \eta dh(\eta) \\ &= - \int_0^\infty \eta d \min\left(K, 4HL^2 \frac{4Z + \frac{2\eta}{\lambda_K} + \sigma^4 Hmdr_w^2 r_\phi^2 \log K}{\eta^2 \sigma^2} d \log(HL/\eta)\right) \\ &= - \int_{a_K}^\infty \eta d\left(4HL^2 \frac{4Z + \frac{2\eta}{\lambda_K} + \sigma^4 Hmdr_w^2 r_\phi^2 \log K}{\eta^2 \sigma^2} d \log(HL/\eta)\right), \end{aligned} \quad (75)$$

where $a_K > 0$ is the constant satisfying

$$4HL^2 \frac{4Z + \frac{2a_K}{\lambda_K} + \sigma^4 Hmdr_w^2 r_\phi^2 \log K}{a_K^2 \sigma^2} d \log(HL/a_K) = K. \quad (76)$$

This is equivalent to

$$Ka_K^2 = \frac{4HL^2}{\sigma^2} (4Z + 2a_K K^{v_1} (\log K)^{v_2} / c + \sigma^4 Hmdr_w^2 r_\phi^2 \log K) d \log(HL/a_K). \quad (77)$$

Since $K^{v_1} (\log K)^{v_2}$ increases at a sublinear rate w.r.t. K , it must be the case that a_K decreases at a sublinear rate w.r.t. K such that the l.h.s. could match the r.h.s..

If $\lim_{K \rightarrow \infty} \frac{2a_K K^{v_1} (\log K)^{v_2} / c}{\sigma^4 Hmdr_w^2 r_\phi^2 \log K} \rightarrow \infty$, we have for large enough K ,

$$\begin{aligned} Ka_K^2 &\geq \frac{8HdL^2}{c\sigma^2} a_K K^{v_1} (\log K)^{v_2} \log(HL) \\ &\Rightarrow a_K \geq H \left(\frac{8dL^2 \log(HL)}{c\sigma^2}\right) K^{v_1-1} (\log K)^{v_2} \\ &\Rightarrow HL/a_K \leq \left(\frac{8dL \log(HL)}{c\sigma^2}\right)^{-1} K^{v_1-1} (\log K)^{-v_2} \\ &\Rightarrow \log(HL/a_K) \leq \log\left(\frac{c\sigma^2}{8dL \log(HL)}\right) + (1 - v_1) \log K - v_2 \log \log K. \end{aligned} \quad (78)$$

According to the Equation (78), we have for large enough K ,

$$\log(HL/a_K) \leq 3(1 - v_1) \log K. \quad (79)$$

It follows from Equations (77) and (79) that for large enough K ,

$$\begin{aligned}
Ka_K^2 &\leq \frac{4HL^2}{\sigma^2} (6a_K K^{v_1} (\log K)^{v_2} / c) d \log(HL/a_K) \\
&\leq \frac{24HdL^2}{c\sigma^2} a_K K^{v_1} (\log K)^{v_2} (3(1-v_1) \log K) \\
&\Rightarrow a_K \leq H \left(\frac{72dL^2}{c\sigma^2} \right) (1-v_1) K^{(v_1-1)} (\log K)^{(v_2+1)}.
\end{aligned} \tag{80}$$

If $\lim_{K \rightarrow \infty} \frac{2a_K K^{v_1} (\log K)^{v_2} / c}{\sigma^4 H m d r_w^2 r_\phi^2 \log K} \rightarrow 0$, we have for large enough K ,

$$\begin{aligned}
Ka_K^2 &\geq 4(HdL\sigma^2 \sqrt{m r_w r_\phi})^2 \log K \log(HL) \\
&\Rightarrow a_K \geq 2HdL\sigma^2 \sqrt{m r_w r_\phi} \sqrt{\log(HL)} \sqrt{\frac{\log K}{K}} \\
&\Rightarrow HL/a_K \leq (2d\sigma^2 r_w r_\phi \sqrt{m \log(HL)})^{-1} \sqrt{\frac{K}{\log K}} \\
&\Rightarrow \log(HL/a_K) \leq \log \left(\frac{1}{2d\sigma^2 r_w r_\phi \sqrt{m \log(HL)}} \right) + \frac{1}{2} \log K - \frac{1}{2} \log \log K.
\end{aligned} \tag{81}$$

It follows from Equations (77) and (81) that for large enough K ,

$$\begin{aligned}
Ka_K^2 &\leq \frac{4HL^2}{\sigma^2} (3\sigma^4 H m d r_w^2 r_\phi^2 \log K) d \log(HL/a_K) \\
&\leq 12(HdL\sigma r_w r_\phi \sqrt{m})^2 \log K \left(\frac{3}{2} \log K \right) \\
&\Rightarrow a_K \leq (HdL\sigma r_w r_\phi \sqrt{18m}) \frac{\log K}{\sqrt{K}}.
\end{aligned} \tag{82}$$

If $\lim_{K \rightarrow \infty} \frac{2a_K K^{v_1} (\log K)^{v_2} / c}{\sigma^4 H m d r_w^2 r_\phi^2 \log K} \rightarrow r$ for $0 < r < \infty$, then $\frac{c r \sigma^4 H m d r_w^2 r_\phi^2}{2} K^{-v_1} (\log K)^{1-v_2}$ is an equivalent infinitesimal of a_K . Taking the limit in Equation (77),

$$\begin{aligned}
1 &= \frac{4HdL^2}{\sigma^2} \lim_{K \rightarrow \infty} \frac{(4Z + 2a_K K^{v_1} (\log K)^{v_2} / c + \sigma^4 H m d r_w^2 r_\phi^2 \log K) \log(HL/a_K)}{Ka_K^2} \\
\Rightarrow 1 &= \frac{4HdL^2}{\sigma^2} \left(\lim_{K \rightarrow \infty} \frac{2K^{v_1} (\log K)^{v_2} \log(HL/a_K)}{cKa_K} + \lim_{K \rightarrow \infty} \frac{\sigma^4 H m d r_w^2 r_\phi^2 \log K \log(HL/a_K)}{Ka_K^2} \right) \\
\Rightarrow 1 &= \frac{8L^2}{\sigma^6 m r_w^2 r_\phi^2 c^2 r / 2} \left(\lim_{K \rightarrow \infty} \frac{\log(HL/a_K)}{K^{1-2v_1} (\log K)^{1-2v_2}} + \frac{1}{r} \lim_{K \rightarrow \infty} \frac{\log(HL/a_K)}{K^{1-2v_1} (\log K)^{1-2v_2}} \right) \\
\Rightarrow 1 &= \frac{16L^2(r+1)}{\sigma^6 m r_w^2 r_\phi^2 c^2 r^2} \lim_{K \rightarrow \infty} \frac{\log(HL/a_K)}{K^{1-2v_1} (\log K)^{1-2v_2}}.
\end{aligned} \tag{83}$$

Notice that

$$\begin{aligned}
& \lim_{K \rightarrow \infty} \frac{\log(HL/a_K)}{K^{1-2v_1}(\log K)^{1-2v_2}} \\
&= \lim_{K \rightarrow \infty} \frac{-a'_K}{a_K((1-2v_1)K^{-2v_1}(\log K)^{1-2v_2} + (1-2v_2)K^{-2v_1}(\log K)^{-2v_2})} \\
&= \lim_{K \rightarrow \infty} \frac{-2a'_K/\sigma^2}{-cr\sigma^2 Hmdr_w^2 r_\phi^2 (-v_1)K^{-v_1-1}(\log K)^{1-v_2} - cr\sigma^2 Hmdr_w^2 r_\phi^2 (1-v_2)K^{-v_1-1}(\log K)^{-v_2}} \\
&= \lim_{K \rightarrow \infty} \frac{-cr\sigma^2 Hmdr_w^2 r_\phi^2 (-v_1)K^{-v_1-1}(\log K)^{1-v_2} - cr\sigma^2 Hmdr_w^2 r_\phi^2 (1-v_2)K^{-v_1-1}(\log K)^{-v_2}}{(1-2v_1)cr\sigma^2 Hmdr_w^2 r_\phi^2 K^{-3v_1}(\log K)^{2-3v_2} + (1-2v_2)cr\sigma^2 Hmdr_w^2 r_\phi^2 K^{-3v_1}(\log K)^{2-3v_2}} \\
&= \lim_{K \rightarrow \infty} \frac{2a_K}{cr\sigma^4 Hmdr_w^2 r_\phi^2 K^{-v_1}(\log K)^{1-v_2}} \\
&= \lim_{K \rightarrow \infty} \frac{-(-v_1) - (1-v_2)(\log K)^{-1}}{(1-2v_1)K^{1-2v_1}(\log K)^{1-2v_2} + (1-2v_2)K^{1-2v_1}(\log K)^{-2v_2}} \\
&= \lim_{K \rightarrow \infty} \frac{-(-v_1) - (1-v_2)(\log K)^{-1}}{(1-2v_1)K^{1-2v_1}(\log K)^{1-2v_2} + (1-2v_2)K^{1-2v_1}(\log K)^{-2v_2}}, \tag{84}
\end{aligned}$$

where the first and third equalities follow from the l'Hôpital's rule, and the second equality is obtained by substituting in a_K and the multiplication rule for limits. Therefore, we have

$$\lim_{K \rightarrow \infty} \frac{\log(HL/a_K)}{K^{1-2v_1}(\log K)^{1-2v_2}} = \begin{cases} 0 & \text{if } 1-2v_1 > 0 \text{ or } 1-2v_1 = 0, -2v_2 > 0 \\ \frac{1}{2} & \text{if } 1-2v_1 = 0, 2v_2 = 0. \end{cases} \tag{85}$$

Combining Equations (83) and (85), we deduce that $v_1 = \frac{1}{2}$, $v_2 = 0$, and $\frac{8L^2(r+1)}{\sigma^6 m r_w^2 r_\phi^2 c^2 r^2} = 1$. Denote the solution of the system of equations as v_1^*, v_2^*, c^* . We have for some $0 < r < \infty$,

$$\begin{cases} v_1^* = \frac{1}{2} \\ v_2^* = 0 \\ c^* = \frac{\sqrt{(r+1)}}{r} \frac{2\sqrt{2}L}{r_w r_\phi \sigma^3 \sqrt{m}}. \end{cases} \tag{86}$$

It follows from the above discussion that for any $l > 1$ and $K \rightarrow \infty$,

$$a_K \leq \begin{cases} lHLLdr_w r_\phi \sigma \sqrt{(2m(r+1))} K^{-\frac{1}{2}} \log K & \text{if } v_1 = v_1^*, v_2 = v_2^*, c = c^* \\ 72HLL^2 \left(\frac{d}{c\sigma^2}\right) (1-v_1)K^{(v_1-1)}(\log K)^{(v_2+1)} & \text{if } v_1 > v_1^* \text{ or } v_1 = v_1^*, v_2 > v_2^* \\ 3HLLdr_w r_\phi \sigma \sqrt{2m} K^{-\frac{1}{2}} \log K & \text{if } v_1 < v_1^* \text{ or } v_1 = v_1^*, v_2 < v_2^*. \end{cases} \tag{87}$$

It follows from Equation (75) that

$$\begin{aligned}
& \sum_{k=1}^K \Delta_k^{\text{conc}} \\
&= \int_{a_K}^{\infty} \eta d \left(4HLL^2 \frac{4Z + \frac{2\eta}{\lambda_K} + \sigma^4 Hmdr_w^2 r_\phi^2 \log K}{\eta^2 \sigma^2} d \log(HL/\eta) \right) \\
&= \frac{4HdL^2}{\sigma^2} \left(\int_{a_K}^{\infty} \frac{4Z + 2\eta/\lambda_K + \sigma^4 Hmdr_w^2 r_\phi^2 \log K}{\eta^2} d\eta \right. \\
& \quad \left. + \int_{a_K}^{\infty} \left(\frac{2(4Z + \sigma^4 Hmdr_w^2 r_\phi^2 \log K)}{\eta^2} + \frac{2K^{v_1}(\log K)^{v_2}}{c\eta} \right) \left(\log \left(\frac{HL}{\eta} \right) \right) d\eta \right). \tag{88}
\end{aligned}$$

Integration by parts for the second term gives

$$\begin{aligned}
& \int_{a_K}^{\infty} \left(\frac{2(4Z + \sigma^4 H m d r_w^2 r_\phi^2 \log K)}{\eta^2} + \frac{2K^{v_1} (\log K)^{v_2}}{c\eta} \right) (\log(HL/\eta)) \, d\eta \\
&= - \int_{a_K}^{\infty} (\log(HL/\eta)) \, d \left(\frac{2(4Z + \sigma^4 H m d r_w^2 r_\phi^2 \log K)}{\eta} + \frac{2K^{v_1} (\log K)^{v_2} \log \eta}{c} \right) \\
&= (\log(HL/a_K)) \left(\frac{2(4Z + \sigma^4 H m d r_w^2 r_\phi^2 \log K)}{a_k} + \frac{2K^{v_1} (\log K)^{v_2} \log a_k}{c} \right) \\
&+ \int_{a_K}^{\infty} \left(\frac{2(4Z + \sigma^4 H m d r_w^2 r_\phi^2 \log K)}{\eta^2} + \frac{2K^{v_1} (\log K)^{v_2} \log \eta}{c\eta} \right) \, d\eta.
\end{aligned} \tag{89}$$

Combining Equations (88) and (89), we obtain

$$\begin{aligned}
& \sum_{k=1}^K \Delta_k^{\text{conc}} \\
&\leq \frac{4HdL^2}{\sigma^2} \left(\int_{a_K}^{\infty} \frac{4Z + 2\eta/\lambda_K + \sigma^4 H m d r_w^2 r_\phi^2 \log K}{\eta^2} \, d\eta \right. \\
&\quad \left. + (\log(HL/a_K)) \left(\frac{2(4Z + \sigma^4 H m d r_w^2 r_\phi^2 \log K)}{a_k} + \frac{2K^{v_1} (\log K)^{v_2} \log a_k}{c} \right) \right. \\
&\quad \left. + \int_{a_K}^{\infty} \left(\frac{2(4Z + \sigma^4 H m d r_w^2 r_\phi^2 \log K)}{\eta^2} + \frac{2K^{v_1} (\log K)^{v_2} \log \eta}{c\eta} \right) \, d\eta \right) \\
&\leq 2a_K K + \frac{4HdL^2}{\sigma^2} \left(\int_{a_K}^{\infty} \frac{4Z + 2\eta/\lambda_K + \sigma^4 H m d r_w^2 r_\phi^2 \log K}{\eta^2} \, d\eta \right. \\
&\quad \left. + \int_{a_K}^{\infty} \left(\frac{2(4Z + \sigma^4 H m d r_w^2 r_\phi^2 \log K)}{\eta^2} + \frac{2K^{v_1} (\log K)^{v_2} \log \eta}{c\eta} \right) \, d\eta \right) \\
&\leq 2a_K K + o(a_K K),
\end{aligned} \tag{90}$$

where the second inequality is derived from Equation (77) and the third inequality can be verified by l'Hôpital's rule.

By plugging in the discussion of a_K in Equation (87), we have

$$\sum_{k=1}^K \Delta_k^{\text{conc}} = \begin{cases} O \left(HL d r_w r_\phi \sigma \sqrt{m} (r_c + 1) K \log K \right) & \text{if } v_1 = \frac{1}{2}, v_2 = 0 \\ O \left(HL^2 d / (c\sigma^2) \right) K^{v_1} (\log K)^{v_2+1} & \text{if } v_1 > \frac{1}{2} \text{ or } v_1 = \frac{1}{2}, v_2 > 0 \\ O \left(HL d r_w r_\phi \sigma \sqrt{m} K (\log K) \right) & \text{if } v_1 < \frac{1}{2} \text{ or } v_1 = \frac{1}{2}, v_2 < 0, \end{cases} \tag{91}$$

where r_c is a constant determined by $c = \frac{\sqrt{(r_c+1)}}{r_c} \frac{2\sqrt{2}L}{r_w r_\phi \sigma^3 \sqrt{m}}$. The proof is completed. \square

D.3 Optimistic Regret

Theorem D.3. *Let \mathcal{F} be the function class defined in Assumption 2. The log-prior probability is uniformly bounded by Z . Setting $\lambda_k = ck^{-v_1} (\log k)^{-v_2}$ for some $0 < v_1 < 1$ or $v_1 = 0, v_2 > 1$, then*

$$\sum_{k=1}^K \Delta_k^{\text{opt}} = O \left(\sigma^4 c H m d r_w^2 r_\phi^2 K^{-v_1} (\log K)^{1-v_2} \right). \tag{92}$$

Proof. According to Theorem D.1, we have the concentration of the squared error around the expected model difference with probability $1 - 3\delta$ for all $f \in \mathcal{F}, K \in \mathbb{Z}^+$,

$$L_{2,K}(f^*) - L_{2,K}(f) \leq \frac{1}{4} \sum_{k=1}^K \mu_k + \beta_K^\dagger(\mathcal{F}, \delta, \alpha), \tag{93}$$

where $\mu_k = -\mathbb{E} \left[\sum_{h=1}^H \|f^*(x_{kh}) - f(x_{kh})\|_2^2 \middle| \mathcal{H}_k \right]$.

The optimal scaling of β_k^\ddagger is specified in Lemma D.7,

$$\beta_k^\ddagger(\mathcal{F}, \delta, \alpha_k) = O \left(\left(\sigma^2 md + \frac{Hmdr_w^2 r_\phi^2}{2} \right) \log k \right). \quad (94)$$

Notice that a model function is selected only when it optimizes the BOO objective, which implies that

$$\frac{1}{\lambda_k} V_1^k(s_1) + \log \Pr(M_k) + \log \Pr(\mathcal{H}_k | M_k) \geq \frac{1}{\lambda_k} V_1^*(s_1) + \log \Pr(M^*) + \log \Pr(\mathcal{H}_k | M^*), \quad (95)$$

for any $k \in [K]$. It follows that

$$\begin{aligned} \frac{1}{\lambda_k} \Delta_k^{\text{opt}} &\leq \log \Pr(M_k) - \log \Pr(M^*) + \log \Pr(\mathcal{H}_k | M_k) - \log \Pr(\mathcal{H}_k | M^*) \\ &\leq 2Z + \frac{\sigma^2}{2} (L_{2,k-1}(f^*) - L_{2,k-1}(f^{M_k})) \\ &\leq 2Z + \frac{\sigma^2}{8} \sum_{i=1}^{k-1} \mu_i + \frac{\sigma^2}{2} \beta_{k-1}^\ddagger(\mathcal{F}, \delta, \alpha) \leq 2Z + \frac{\sigma^2}{2} \beta_{k-1}^\ddagger(\mathcal{F}, \delta, \alpha), \end{aligned} \quad (96)$$

where the second inequality holds because the log-prior probability is assumed to be uniformly bounded by Z , and the likelihood is defined by Equation (14). The third inequality follows from Equation (93).

For the first episode, we have

$$\Delta_1^{\text{opt}} \leq V_1^*(s_1) - V_1^{\pi^*, M_1}(s_1) \leq LH \mathbb{E}_{x \sim \rho^{M_1, M^*}} \left[\|f^{M^*}(x) - f^{M_1}(x)\|_2 \right] \leq 2LHr_w r_\phi. \quad (97)$$

Therefore, the sum of optimistic regret is at most

$$\begin{aligned} \sum_{k=1}^K \Delta_k^{\text{opt}} &\leq 2LHr_w r_\phi + c \sum_{k=2}^K \frac{4Z + \sigma^2 \beta_k^\ddagger(\mathcal{F}, \delta, \alpha)}{2k^{v_1} (\log k)^{v_2}} \\ &\leq 2LHr_w r_\phi + 2cZ \int_1^K k^{-v_1} (\log k)^{-v_2} dk + \frac{\sigma^4 Hmdcr_w^2 r_\phi^2}{2} \int_1^K k^{-v_1} (\log k)^{1-v_2} dk \\ &\leq 2LHr_w r_\phi + \frac{2cZ}{1-v_1} \left(K^{1-v_1} (\log K)^{-v_2} + v_2 \int_1^K k^{-v_1} (\log k)^{-v_2-1} dk \right) \\ &\quad + \frac{\sigma^4 Hmdcr_w r_\phi}{2(1-v_1)} \left(K^{1-v_1} (\log K)^{1-v_2} + (v_2-1) \int_1^K k^{-v_1} (\log k)^{-v_2} dk \right) \\ &= O \left(\sigma^4 Hmdcr_w^2 r_\phi^2 K^{-v_1} (\log K)^{1-v_2} \right). \end{aligned} \quad (98)$$

The proof is completed. \square

D.4 Proof of BOO Regret

Theorem B.1 (Bayesian optimistic optimization regret). *Let \mathcal{F} be the function class defined in Assumption 2, i.e., $\mathcal{F} \subseteq \{f \mid W \in \mathbb{R}^{(m+1) \times d}, \|W\|_F \leq r_w, f(x) = W\phi(x), \|\phi(x)\| \leq r_\phi, \forall x \in \mathcal{X}\}$. The log-prior probability is assumed to be uniformly bounded by some constant. Setting $\lambda_k = ck^{-v_1} (\log k)^{-v_2}$ for $c > 0$ and $0 < v_1 < 1$ or $v_1 = 0, v_2 \geq 0$, the asymptotic regret of BOO is*

$$\begin{cases} O \left(H L d r_w r_\phi \sigma \sqrt{m(r_c + 1) K} \log K \right) & \text{if } v_1 = \frac{1}{2}, v_2 = 0 \\ O \left(H L^2 d / (c \sigma^2) K^{v_1} (\log K)^{v_2+1} \right) & \text{if } v_1 > \frac{1}{2} \text{ or } v_1 = \frac{1}{2}, v_2 > 0 \\ O \left(H m d r_w^2 r_\phi^2 \sigma^4 K^{1-v_1} (\log K)^{1-v_2} \right) & \text{if } v_1 < \frac{1}{2} \text{ or } v_1 = \frac{1}{2}, v_2 < 0, \end{cases} \quad (16)$$

where m is the dimension of the state space and r_c is a constant determined by $c = \frac{\sqrt{(r_c+1)}}{r_c} \frac{2\sqrt{2}L}{r_w r_\phi \sigma^3 \sqrt{m}}$.

Proof. Theorem B.1 can be obtained by combining Theorem D.2 and Theorem D.3. \square

E Optimization

E.1 Model Gradient of Value Function

Proof. (of Theorem 5.1) We first show that the model gradient satisfies the following recursion,

$$\begin{aligned}
& \sum_{s_h \in \mathcal{S}} \Pr(S_h = s_h | \pi, M_\theta) \nabla_\theta V_h^{\pi, M_\theta}(s_h) \\
&= \sum_{s_h \in \mathcal{S}} \sum_{a_h \in \mathcal{A}} \Pr(S_h = s_h, A_h = a_h | \pi, M_\theta) \sum_{s_{h+1} \in \mathcal{S}} P^{M_\theta}(s_{h+1} | s_h, a_h) \nabla_\theta V_{h+1}^{\pi, M_\theta}(s_{h+1}) \\
&\quad + \sum_{s_h \in \mathcal{S}} \sum_{a_h \in \mathcal{A}} \Pr(S_h = s_h, A_h = a_h | \pi, M_\theta) \left(\nabla_\theta \bar{R}_\theta(s_h, a_h) \right. \\
&\qquad \qquad \qquad \left. + \sum_{s_{h+1} \in \mathcal{S}} V_{h+1}^{\pi, M_\theta}(s_{h+1}) \nabla_\theta P^{M_\theta}(s_{h+1} | s_h, a_h) \right) \tag{99} \\
&= \sum_{s_h \in \mathcal{S}} \sum_{a_h \in \mathcal{A}} \sum_{s_{h+1} \in \mathcal{S}} \Pr(S_{h+1} = s_{h+1}, A_h = a_h | \pi, M_\theta) P^{M_\theta}(s_{h+1} | s_h, a_h) \nabla_\theta V_{h+1}^{\pi, M_\theta}(s_{h+1}) \\
&\quad + \sum_{s_h \in \mathcal{S}} \sum_{a_h \in \mathcal{A}} \Pr(S_h = s_h, A_h = a_h | \pi, M_\theta) \left(\nabla_\theta \bar{R}_\theta(s_h, a_h) \right. \\
&\qquad \qquad \qquad \left. + \sum_{s_{h+1} \in \mathcal{S}} V_{h+1}^{\pi, M_\theta}(s_{h+1}) \nabla_\theta P^{M_\theta}(s_{h+1} | s_h, a_h) \right),
\end{aligned}$$

where $\Pr(S_h = s_h, A_h = a_h | \pi, M_\theta)$ is the marginal probability of s_h and a_h at the period h following policy the π and MDP M_θ . Expanding the recursion, we obtain

$$\begin{aligned}
\nabla_\theta V_1^{\pi, M_\theta}(s_1) &= \sum_{s \in \mathcal{S}} \Pr(S_1 = s) \nabla_\theta V_h^{\pi, M_\theta}(s) \\
&= \sum_{h=1}^H \sum_{s_h \in \mathcal{S}} \sum_{a_h \in \mathcal{A}} \Pr(S_h = s_h, A_h = a_h | \pi, M_\theta) \nabla_\theta \bar{R}_\theta(s_h, a_h) \\
&\quad + \sum_{h=1}^{H-1} \sum_{s_h \in \mathcal{S}} \sum_{a_h \in \mathcal{A}} \Pr(S_h = s_h, A_h = a_h | \pi, M_\theta) \sum_{s_{h+1} \in \mathcal{S}} V_{h+1}^{\pi, M_\theta}(s_{h+1}) \nabla_\theta P^{M_\theta}(s_{h+1} | s_h, a_h). \tag{100}
\end{aligned}$$

It follows from the REINFORCE trick [36] that

$$\begin{aligned}
& \sum_{s_h \in \mathcal{S}} \sum_{a_h \in \mathcal{A}} \Pr(S_h = s_h, A_h = a_h | \pi, M_\theta) \sum_{s_{h+1} \in \mathcal{S}} V_{h+1}^{\pi, M_\theta}(s_{h+1}) \nabla_\theta P^{M_\theta}(s_{h+1} | s_h, a_h) \\
&= \sum_{s_h, a_h, s_{h+1}} \Pr(S_h = s_h, A_h = a_h | \pi, M_\theta) P^{M_\theta}(s_{h+1} | s_h, a_h) V_{h+1}^{\pi, M_\theta}(s_{h+1}) \nabla_\theta \log P^{M_\theta}(s_{h+1} | s_h, a_h) \\
&= \sum_{s_h, a_h, s_{h+1}} \Pr(S_h = s_h, A_h = a_h, S_{h+1} = s_{h+1} | \pi, M_\theta) V_{h+1}^{\pi, M_\theta}(s_{h+1}) \nabla_\theta \log P^{M_\theta}(s_{h+1} | s_h, a_h). \tag{101}
\end{aligned}$$

Here, $\sum_{s_h, a_h, s_{h+1}}$ abbreviates $\sum_{s_h \in \mathcal{S}} \sum_{a_h \in \mathcal{A}} \sum_{s_{h+1} \in \mathcal{S}}$. Combining Equations (100) and (101), we have

$$\begin{aligned}
\nabla_\theta V_1^{\pi, M_\theta}(s_1) &= \sum_{h=1}^H \sum_{s_h \in \mathcal{S}} \sum_{a_h \in \mathcal{A}} \Pr(S_h = s_h, A_h = a_h | \pi, M_\theta) \nabla_\theta \bar{R}_\theta(s_h, a_h) \\
&\quad + \sum_{h=1}^{H-1} \sum_{s_h, a_h, s_{h+1}} \Pr(S_h = s_h, A_h = a_h, S_{h+1} = s_{h+1} | \pi, M_\theta) V_{h+1}^{\pi, M_\theta}(s_{h+1}) \nabla_\theta \log P^{M_\theta}(s_{h+1} | s_h, a_h). \tag{102}
\end{aligned}$$

By replacing the marginal distribution with the joint distribution,

$$\begin{aligned}
\nabla_{\theta} V_1^{\pi, M_{\theta}}(s_1) &= \sum_{\tau} \Pr(\tau|\pi, M_{\theta}) \sum_{h=1}^H \nabla_{\theta} \bar{R}_{\theta}(s_h, a_h) \\
&\quad + \sum_{\tau} \Pr(\tau|\pi, M_{\theta}) \sum_{h=1}^{H-1} V_{h+1}^{\pi, M_{\theta}}(s_{h+1}) \nabla_{\theta} \log P^{M_{\theta}}(s_{h+1}|s_h, a_h) \\
&= \mathbb{E}_{\tau \sim \pi, M_{\theta}} \left[\sum_{h=1}^H \nabla_{\theta} \bar{R}^{M_{\theta}}(s_h, a_h) + \sum_{h=1}^{H-1} V_{h+1}^{\pi, M_{\theta}}(s_{h+1}) \nabla_{\theta} \log P^{M_{\theta}}(s_{h+1}|s_h, a_h) \right].
\end{aligned} \tag{103}$$

The proof is completed. \square

F Implementation Details

In this section, we explain the implementation of BOO in tabular MDPs and detail various tricks for improving the efficiency of optimization.

F.1 Optimization Reformulation

Focusing on the optimization of the model, we directly compute the optimal policy of a given model via dynamic programming according to Equation (9) and differentiate through the Bellman optimality equation to find the model gradient. In other words, we reformulate the joint optimization of policy and model into

$$\max_M \left(V_1^{\pi^*(M), M}(s_1) + \lambda_k (\log \Pr(\mathcal{H}_k|M) + \log \Pr(M)) \right), \tag{104}$$

where $\pi^*(M)$ is the optimal maximum entropy policy of the model M .

F.2 Natural Model Gradient

The natural policy gradient proposed by [49] improves the efficiency of the policy gradient [50] by preconditioning the gradient with the inverse Fisher information matrix. We formalize the natural model gradient in a similar fashion.

Let D_k^R and D_k^P be normalized empirical state-action distribution for rewards and transition:

$$D_k^R(s, a) \propto \sum_{i=1}^{k-1} \sum_{h=1}^H \mathbb{I}(s_{ih} = s, a_{ih} = a) \quad D_k^P(s, a) \propto \sum_{i=1}^{k-1} \sum_{h=1}^{H-1} \mathbb{I}(s_{ih} = s, a_{ih} = a). \tag{105}$$

Denoting the BOO objective in the k -th episode as $J_k = \lambda_k \log \Pr(M_{\theta}|\mathcal{H}_k) + V_1^{\pi, M_{\theta}}(s_1)$, the natural model gradient is $\bar{\nabla}_{\theta} J_k = F_k^{-1} \nabla_{\theta} J_k$, where:

$$F_k = \mathbb{E}_{\tau \sim \pi, M_{\theta}} \left[-\frac{1}{H} \sum_{h=1}^H G_{s_h, a_h}^{\bar{R}} + \frac{1}{H-1} \sum_{h=1}^{H-1} F_{s_h, a_h}^P \right] + \mathbb{E}_{(s,a) \sim D_k^R} [F_{s,a}^R] + \mathbb{E}_{(s,a) \sim D_k^P} [F_{s,a}^P], \tag{106}$$

Here, $F_{s,a}^P$ and $F_{s,a}^R$ are the Fisher information matrices of transition and reward distributions, as follows:

$$\begin{aligned}
F_{s,a}^R &= \mathbb{E}_{r \sim R^{M_{\theta}}(s_h, a_h)} [(\nabla_{\theta} \log R^{M_{\theta}}(r|s_h, a_h))^2], \\
F_{s,a}^P &= \mathbb{E}_{s' \sim P^{M_{\theta}}(s_h, a_h)} [(\nabla_{\theta} \log P^{M_{\theta}}(s'|s_h, a_h))^2].
\end{aligned} \tag{107}$$

The matrix $G_{s,a}^{\bar{R}}$ is the generalized Gauss-Newton matrix [51]:

$$G_{s,a}^{\bar{R}} = \nabla_{\theta} z_R^{\top} \nabla_{z_R}^2 \bar{R}^{M_{\theta}}(s, a) \nabla_{\theta} z_R, \tag{108}$$

where the reward distribution is parameterized by $z_R \subseteq z$, and $z = f_{\theta}(s, a)$ is the output of the model function f_{θ} .

Note that we have $G_{s,a}^{\bar{R}} = \mathbf{0}$ when the mean reward $\bar{R}^{M_\theta}(s, a)$ is linear in z . This is satisfied for most model classes, for example, those specified by Assumption 1. Therefore, commonly, F_k contains only the Fisher information matrix, and the natural gradient can be interpreted geometrically as the steepest descent direction in the information geometry induced by F_k instead of in the Euclidean geometry. Loosely speaking, the distance in the information geometry is measured by the square root of KL divergence between distributions rather than the ℓ_2 distance of parameters [52]. The natural gradient characterizes some global characteristics of the loss surface and therefore facilitates global convergence. We refer interested readers to [52] for a thorough discussion on the interpretation and implementation of the natural gradient.

F.3 Initial State-Action Distribution

It is well-known that the policy gradient suffers from slow convergence since the value gradient of unvisited states are 0. [38] proposes to alleviate this problem by changing the initial state distribution to a distribution with broader coverage. In BOO, this problem is aggravated since the model tends to be optimistic on state-action pairs within its current occupancy measure, causing the model stuck at local optima. Seeing the similarity between the transition function and the policy, we could apply a similar technique to speed up convergence and help escape local optima.

Specifically, we use the state-action distribution induced by the learned policy and model in the past as the initial state-action distribution. Let $0 < \gamma < 1$ be a smoothing factor. We incrementally update the initial state-action distribution φ_{k+1} by $\varphi_{k+1} = \varphi_k * \gamma + \rho_k * (1 - \gamma)$. Here, ρ_k represents the state-action distribution induced by the model and policy in the k -th episode starting from the real initial state s_0 . The initial distribution φ_1 is set to be the uniform distribution over $\mathcal{S} \times \mathcal{A}$.

F.4 Posterior Sampling

In tabular MDPs, the transition distribution is multinomial distribution whose conjugate prior is the Dirichlet distribution. We assume the reward is distributed according to a normal distribution with standard deviation 1 and unknown mean. The conjugate prior of Gaussian reward is Gaussian. For distributions with a conjugate prior, we calculate the posterior distribution analytically and easily sample from the posterior.

F.5 Gradient Ascent

The optimization of BOO is conducted by gradient ascent with line search. One gradient step is performed at each episode. The line search algorithm is implemented in SciPy [53], which enforces the strong Wolfe conditions. For natural gradient, we additionally use the Tikhonov regularization to limit the condition number of the inverse Fisher information matrix.

F.5.1 Entropy Regularization

As discussed in Section 5, the entropy plays a role in smoothing the policy’s loss landscape such that the optimal policy will not change drastically when the model changes. However, the randomness of policy would increase the cumulative regret by selecting sub-optimal actions. Therefore, we derive a policy with reduced entropy for execution in the environment. Specifically, we multiply a coefficient w_ζ to ζ during the derivation of policy according to Equations (8) and (9). In all experiments, w_ζ is set to 0.1.

G Ablation Study

To demonstrate the effectiveness of our proposed method for optimization, we conduct some ablation studies and demonstrate the performance of different algorithm implementations on 100 randomly generated MDPs with $|\mathcal{S}| = |\mathcal{A}| = H = 5$. Hyperparameters used are revealed at the end.

As we can see from Figure 3, when ζ is very small, the optimization failure rate of BOO-E ($\zeta = 0.0001$) is very high, which indicates that the optimization of the BOO algorithm without entropy regularization is very inefficient due to the unsmooth objective and the poor coverage of the policy.

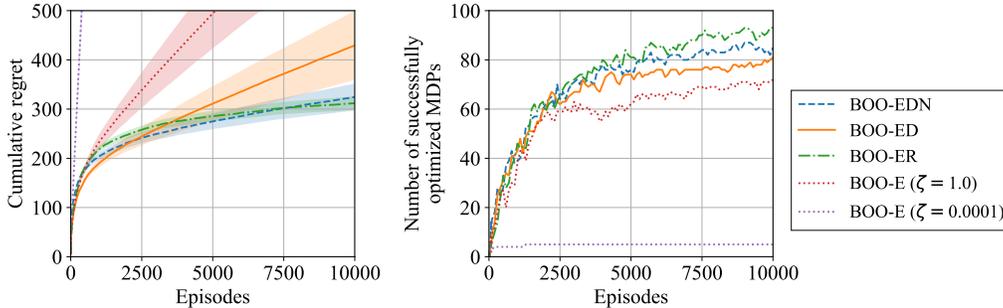


Figure 3: The left and right figures show respectively the cumulative regret and success rate of BOO variants on Random MDPs. An optimization is marked as success when the difference between model value $V_1^k(s_0)$ and the real value $V_1^{\pi_k, M^*}(s_0)$ is within 0.01. E, D, N, and R denote entropy regularization, changing initial distribution, natural gradient and adding mean reward bonus, respectively. For example, BOO-E stands for BOO with entropy regularization.

By using a starting distribution with broader coverage and natural gradients, the optimization success rate increases, and the cumulative regret decreases. It is observed that all proposed techniques could help optimization. Among all variants, BOO-ER achieves the best performance and is used for performance comparison experiments in Section 6.

G.1 Hyperparameters

In this section, we provide the hyperparameters of the algorithms used in Sections 6 and G , as shown in Tables 1 and 2, respectively.

Table 1: Hyperparameters in Section 6

	RiverSwim	Random MDPs	Random MDPs	Chain10	Chain20	Chain40
$ \mathcal{S} $	5	5	20	10	20	40
$ \mathcal{A} $	2	5	5	2	2	2
H	5	5	20	10	20	40
FiniteBOO						
λ_k	$\frac{1}{\sqrt{k}}$	$\frac{1}{\sqrt{k}}$	$\frac{1}{\sqrt{k}}$	$\frac{1}{\sqrt{k}}$	$\frac{1}{\sqrt{k}}$	$\frac{1}{\sqrt{k}}$
ζ	0.1	1	1	1	1	2
w_ζ	0.1	0.1	0.1	0.1	0.1	0.1
γ	-	-	-	-	-	-
BPS						
λ_k	$\frac{\log(k+1)}{\sqrt{k}}$	$\frac{\log(k+1)}{\sqrt{k}}$	$\frac{\log(k+1)}{\sqrt{k}}$	$\frac{\log(k+1)}{\sqrt{k}}$	$\frac{\log(k+1)}{\sqrt{k}}$	$\frac{\log(k+1)}{\sqrt{k}}$
nSamp	10	10	10	10	10	10

‘nSamp’ is a hyperparameter that determines the number of posterior samples.

Table 2: Hyperparameters in Section G

	λ_k	ζ	w_ζ	γ
BOO-E($\zeta = 0.0001$)	$\frac{1}{\sqrt{k}}$	0.0001	0.1	-
BOO-E($\zeta = 1$)	$\frac{1}{\sqrt{k}}$	1	0.1	-
BOO-ED	$\frac{1}{\sqrt{k}}$	1	0.1	0.998
BOO-ER	$\frac{1}{\sqrt{k}}$	1	0.1	-
BOO-EDN	$\frac{1}{\sqrt{k}}$	1	0.1	0.998