# BooVI: Provably Efficient Bootstrapped Value Iteration 

Boyi Liu*<br>Qi Cai ${ }^{\dagger}$<br>Zhuoran Yang ${ }^{\ddagger}$<br>Zhaoran Wang ${ }^{\S}$


#### Abstract

Despite the tremendous success of reinforcement learning (RL) with function approximation, efficient exploration remains a significant challenge, both practically and theoretically. In particular, existing theoretically grounded RL algorithms based on upper confidence bounds (UCBs), such as optimistic least-squares value iteration (LSVI), are often incompatible with practically powerful function approximators, such as neural networks. In this paper, we develop a variant of bootstrapped LSVI, namely BooVI, which bridges such a gap between practice and theory. Practically, BooVI drives exploration through (re)sampling, making it compatible with general function approximators. Theoretically, BooVI inherits the worst-case $\widetilde{O}\left(\sqrt{d^{3} H^{3} T}\right)$-regret of optimistic LSVI in the episodic linear setting. Here $d$ is the feature dimension, $H$ is the episode horizon, and $T$ is the total number of steps.


## 1 Introduction

Reinforcement learning (RL) with function approximation demonstrates significant empirical success in a broad range of applications [e.g., 16, 38, 40, 45]. However, computationally and statistically efficient exploration of large and intricate state spaces remains a major barrier. Practically, the lack of temporally-extended exploration [31, 28, 30] in existing RL algorithms, e.g., deterministic policy gradient [39] and soft actor-critic [19], hinders them from solving more challenging tasks, e.g., Montezuma's Revenge in Atari Games [42, 17], in a more sample-efficient manner. Theoretically, it remains unclear how to design provably efficient RL algorithms with finite sample complexities or regrets that allow for practically powerful function approximators, e.g., neural networks.

There exist two principled approaches to efficient exploration in RL, namely optimism in the face of uncertainty [e.g., 4, 41, 21, 11, 12, 5, 23, 24] and posterior sampling [e.g., 29, 34, 33, 30, 27, 36].

- The optimism-based approach is often instantiated by incorporating upper confidence bounds (UCBs) into the estimated (action-)value functions as bonuses, which are used to direct exploration. In the tabular setting, the resulting algorithms, e.g., variants of upper confidence bound reinforcement learning (UCRL) [4 21, 11, 5], are known to attain the optimal worstcase regret [32]. Beyond the tabular setting, it remains unclear how to construct closed-form UCBs in a principled manner for general function approximators, e.g., neural networks. The only exception is the linear setting [51, 52, 24, 8], where optimistic least-squares value iteration (LSVI) [24] is known to attain a near-optimal (with respect to $T=K H$ ) worst-case regret. However, the closed-form UCB therein is tailored to linear models [1, 9] rather than fully general-purpose.

[^0]- On the other hand, the posterior-based approach, which originates from Thompson sampling [43, 37], can be instantiated using randomized (action-)value functions [34, 30, 27, 36]. Unlike optimistic LSVI, the resulting algorithm, namely randomized LSVI, straightforwardly allows for general function approximators, as it only requires injecting random noise into the training data of LSVI. Ideally, such injected random noise does not depend on particular function approximators. In the tabular setting, randomized LSVI is known to attain near-optimal worst-case and Bayesian regrets [30, 36]. Meanwhile, in the linear setting, randomized LSVI is very recently shown to attain a near-optimal worst-case regret [54], which is only worse than that of optimistic LSVI by a factor of $\sqrt{d H}$. Here $d$ is the feature dimension and $H$ is the episode horizon. However, achieving such a regret requires a nontrivial modification of the injected random noise, which is tailored to linear models. Such a specialized modification diminishes the supposed practical advantage of randomized LSVI.

To bridge such a gap between practice and theory, we aim to answer the following question: Can we design an RL algorithm that simultaneously achieves the practical advantage of randomized LSVI and the theoretical guarantee of optimistic LSVI?
In this paper, we propose a variant of bootstrapped LSVI, namely BooVI, which combines the advantages of the optimism-based and posterior-based approaches. The key idea of BooVI is to use posterior sampling to implicitly construct an "optimistic version" of the estimated (action-)value functions in a data-driven manner. Unlike randomized LSVI, which samples from the posterior only once, e.g., by injecting random noise into the training data of LSVI, BooVI samples from the posterior multiple times, e.g., via the Langevin dynamics [48 35]. Upon evaluating an action at a state, BooVI ranks the values of the randomized (action-)value functions sampled from the posterior in descending order and returns a top-ranked value, which can be shown to be approximately optimistic. Generally speaking, BooVI corresponds to bootstrapping the noise in the least-squares regression problem of LSVI. As a result, it can be viewed as a parametric bootstrap counterpart of the nonparametric bootstrap technique used in bootstrapped deep Q-networks (DQNs) [31, 30], which demonstrates significant empirical success in terms of exploration.

Compared with existing RL algorithms with function approximation, the advantage of BooVI is twofold:

- Practically, BooVI bypasses the explicit construction of the closed-form UCB in optimistic LSVI. Like randomized LSVI, BooVI straightforwardly allows for general function approximators, as it only requires injecting random noise into the training process of LSVI, e.g., via the Langevin dynamics.
- Theoretically, BooVI inherits the worst-case $\widetilde{O}\left(\sqrt{d^{3} H^{3} T}\right)$-regret of optimistic LSVI in the linear setting. Here $d$ is the feature dimension, $H$ is the episode horizon, and $T$ is the total number of steps. Such a regret is better than the best known worst-case regret of randomized LSVI by a factor of $\sqrt{d H}$. More importantly, unlike the specialized variant of randomized LSVI studied in [54], BooVI can be applied to the linear setting "as is", without tailoring the injected random noise to linear models.

More Related Work: The idea of using posterior sampling to achieve optimism in a data-driven manner is previously studied in the linear bandit setting [26] and the tabular setting of RL [3]. In fact, our linear setting of RL covers the linear bandit setting as a special case, where the episode horizon $H$ is set to one, and the tabular setting of RL as another special case, where the feature mapping is the canonical basis of the state and action spaces. Although BooVI is practically applicable to general function approximators, our theoretical guarantee on its worst-case regret is only applicable to the linear setting. Despite the recent progress [49 22, 13 15], simultaneously achieving provable computational and statistical efficiency in exploration with general function approximators remains challenging. See, e.g., [2] 18] for the recent progress in the contextual bandit setting. Finally, we refer readers to [46, 53, 57, 55] for theories on the generalized models of linear MDPs, which all share similar (approximately) linear structure with the class of linear MDPs considered in the analysis of this paper. We expect our regret bound stays the same for those settings with linear structures, or admits an extra $O(\epsilon T)$ dependency over $T$ for those settings with approximately linear structure, where $\epsilon>0$ describes the level of nonlinearity in the MDP. We also refer readers to [50] for sharper regret bounds in linear MDPs, which are out of the scope of this paper.

## 2 Background

In this section, we introduce the general problem setting.
Episodic Markov Decision Process. We consider the episodic Markov decision process represented by a tuple $(\mathcal{S}, \mathcal{A}, H, \mathcal{P}, r)$, where $\mathcal{S}$ is a compact state space, $\mathcal{A}$ is a finite action space with cardinality $A, H$ is the number of timesteps in each episode, $\mathcal{P}=\left\{\mathcal{P}_{h}\right\}_{h \in[H]}$ with $\mathcal{P}_{h}: \mathcal{S} \times \mathcal{S} \times \mathcal{A} \rightarrow[0,1]$ for all $h \in[H]$ is the set of transition kernels, and $r=\left\{r_{h}\right\}_{h \in[H]}$ with $r_{h}: \mathcal{S} \times \mathcal{A} \rightarrow[0,1]$ for all $h \in[H]$ is the set of reward functions.
At each timestep $h \in[H]$ of episode $k$, an agent at state $s_{h}^{k} \in \mathcal{S}$ with policy $\pi^{k}=\left\{\pi_{h}^{k}\right\}_{h \in[H]}$, where $\pi_{h}^{k}: \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ for all $h \in[H]$, interacts with the environment by first taking an action $a_{h}^{k}$ with probability $\pi_{h}^{k}\left(a_{h}^{k} \mid s_{h}^{k}\right)$ and then receiving the corresponding reward $r_{h}^{k}=r_{h}\left(s_{h}^{k}, a_{h}^{k}\right)$.
We evaluate the performance of policy $\pi=\left\{\pi_{h}\right\}_{h \in[H]}$ starting from timestep $h$, and state-action pair $(s, a)$ using its action-value function $Q_{h}^{\pi}: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, which is defined as

$$
Q_{h}^{\pi}(s, a)=\mathbb{E}_{\pi}\left[\sum_{h^{\prime}=h}^{H} r_{h^{\prime}}\left(s_{h^{\prime}}, a_{h^{\prime}}\right) \mid s_{h}=s, a_{h}=a\right]
$$

for all $h \in[H]$. Correspondingly, the value function $V_{h}^{\pi}: \mathcal{S} \rightarrow \mathbb{R}$ of a policy $\pi$ is defined as

$$
V_{h}^{\pi}(s)=\mathbb{E}_{\pi}\left[\sum_{h^{\prime}=h}^{H} r_{h^{\prime}}\left(s_{h^{\prime}}, a_{h^{\prime}}\right) \mid s_{h}=s\right]=\mathbb{E}_{\pi}\left[Q_{h}^{\pi}\left(s_{h}, a\right) \mid s_{h}=s\right]
$$

for all $h \in[H]$. Here the expectation $\mathbb{E}_{\pi}[\cdot]$ is taken over the trajectory generated by $\pi$. Also, we let $Q_{H+1}^{\pi} \equiv 0$ and thus $V_{H+1}^{\pi} \equiv 0$. Furthermore, we denote by $V_{h}^{*}(s)$ the value function corresponding to the optimal policy $\pi^{*}$. Finally, for notational simplicity, we define $\left[\mathcal{P}_{h} V\right](s, a)=$ $\mathbb{E}_{s^{\prime} \sim \mathcal{P}_{h}(\cdot \mid s, a)}\left[V\left(s^{\prime}\right)\right]$, where $V: \mathcal{S} \rightarrow \mathbb{R}$ can be any function.
For any algorithm that generates a sequence of policies $\left\{\pi^{k}\right\}_{k \in[K]}$, we track its performance via the cumulative regret defined by

$$
\operatorname{Regret}(K)=\sum_{k=1}^{K}\left[V_{1}^{*}\left(s_{1}^{k}\right)-V_{1}^{\pi^{k}}\left(s_{1}^{k}\right)\right]
$$

where $K$ is the number of episodes. Here $s_{1}^{k}$ is the initial state of episode $k$, which is arbitrarily chosen at the start of the episode.

## 3 Bootstrapped Value Iteration

In this section, we introduce Bootstrapped Value Iteration (BooVI in Algorithm 1). For notational simplicity, we write $r_{h}\left(s_{h}^{k}, a_{h}^{k}\right)$ as $r_{h}^{k}$ throughout the rest of this paper. Also, in this paper, we write $\max \{\min \{\cdot, \cdot\}, 0\}$ as $\min \{\cdot, \cdot\}^{+}$, and denote by $\|\cdot\|$ the 2 -norm for vectors.

Least-Squares Value Iteration. At timestep $h \in[H]$ of episode $k$, given the estimated action-value function $\widehat{Q}_{h+1}^{k}\left(s_{h+1}^{\tau}, a\right)$ for all $\tau \in[k-1]$ and $a \in \mathcal{A}$, let $\widehat{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)=\max _{a \in \mathcal{A}} \widehat{Q}_{h+1}^{k}\left(s_{h+1}^{\tau}, a\right)$. Then, Least-Squares Value Iteration (LSVI) updates the parameter $\omega$ of the action-value function via

$$
\begin{equation*}
\widehat{\omega}_{h}^{k} \leftarrow \underset{\omega \in \mathbb{R}^{d}}{\operatorname{argmin}}\left\{\lambda \cdot\|\omega\|^{2}+\sum_{\tau=1}^{k-1}\left(r_{h}^{\tau}+\widehat{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)-Q\left(s_{h}^{\tau}, a_{h}^{\tau} ; \omega\right)\right)^{2}\right\} \tag{3.1}
\end{equation*}
$$

where $Q(\cdot, \cdot ; \cdot): \mathcal{S} \times \mathcal{A} \times \mathbb{R}^{d} \rightarrow \mathbb{R}$ is the parameterization of action-value function with $\omega \in \mathbb{R}^{d}$ being its parameter, and $\lambda \geq 0$ is the regularization parameter. Although the deterministic parameter update by LSVI could exploit the historical data well, it has limited ability to address the exploration need in more challenging tasks.

Bootstrapped Value Iteration. To achieve guided exploration, we introduce BooVI (Algorithm 1), which uses bootstrapping to enforce the optimism of the estimated action-value function.

At the timestep $h$ of episode $k$, given the current data buffer $\mathcal{D}_{h}^{k}=\left\{\left(s_{h}^{\tau}, a_{h}^{\tau}, r_{h}^{\tau}, s_{h+1}^{\tau}\right)\right\}_{\tau \in[k-1]}$, which contains the data collected from the previous $k-1$ episodes, and the current bootstrapped state values $\left\{\widetilde{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)\right\}_{\tau \in[k-1]}$, we define

$$
\begin{equation*}
y_{h}^{\tau}=r_{h}^{\tau}+\widetilde{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right), \quad \text { for all } \tau \in[k-1] . \tag{3.2}
\end{equation*}
$$

With $p_{0}(\omega)$ as prior of $\omega$ and $p\left(y_{h}^{\tau} \mid\left(s_{h}^{\tau}, a_{h}^{\tau}\right), \omega\right)$ as the likelihood of $y_{h}^{\tau}$, the posterior of $\omega$ is given by

$$
\begin{equation*}
p\left(\omega \mid\left\{\widetilde{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)\right\}_{\tau \in[k-1]}, \mathcal{D}_{h}^{k}\right) \propto p_{0}(\omega) \cdot \prod_{\tau=1}^{k-1} p\left(y_{h}^{\tau} \mid\left(s_{h}^{\tau}, a_{h}^{\tau}\right), \omega\right) \tag{3.3}
\end{equation*}
$$

In BooVI (Algorithm(1), we propose to sample repeatedly from such posterior for $N_{k}$ times, which is equivalent to sampling repeatedly from randomized action-value function since each sample $\omega_{h}^{k, i}$ corresponds to one randomized action-value function $Q\left(\cdot, \cdot ; \omega_{h}^{k, i}\right)$. Such a collection of posterior weights is later used to construct the bootstrapped action-value function in Algorithm 2

To independently sample from the posterior in (3.3), we can consider, e.g., using the Langevin dynamics $\omega(t+1) \leftarrow \omega(t)+\Delta \omega(t)$, where

$$
\begin{equation*}
\Delta \omega(t)=\frac{\epsilon_{t}}{2} \cdot\left(\nabla_{\omega} \log p_{0}(\omega(t))+\sum_{\tau=1}^{k-1} \nabla_{\omega} \log p\left(y_{h}^{\tau} \mid\left(s_{h}^{\tau}, a_{h}^{\tau}\right), \omega(t)\right)\right)+\eta_{t} \tag{3.4}
\end{equation*}
$$

where $\epsilon_{t}>0$ is the stepsize, and $\eta_{t} \sim \mathcal{N}\left(0, \epsilon_{t}\right)$. We note here that, after running sufficient many iterations with suitable choices of stepsizes $\epsilon_{t}>0$, the Langevin dynamics gives effectively independent parameter samples from the posterior in (3.3).
To instantiate the connection between the posterior (3.3) with the LSVI, we consider Gaussian prior and likelihood as an example. For the Gaussian prior $\omega \sim \mathcal{N}\left(0, I_{d}\right)$ and the Gaussian likelihood $p\left(y_{h}^{\tau} \mid\left(s_{h}^{\tau}, a_{h}^{\tau}\right), \omega\right) \propto \exp \left\{-\left(y_{h}^{\tau}-Q\left(s_{h}^{\tau}, a_{h}^{\tau} ; \omega\right)\right)^{2} /\left(2 \sigma^{2}\right)\right\}$, the posterior of $\omega$ is given by

$$
\begin{equation*}
p\left(\omega \mid\left\{\widetilde{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)\right\}_{\tau \in[k-1]}, \mathcal{D}_{h}^{k}\right) \propto \exp \left\{-\frac{1}{2} \cdot\|\omega\|^{2}-\frac{1}{2 \sigma^{2}} \sum_{\tau=1}^{k-1}\left(y_{h}^{\tau}-Q\left(s_{h}^{\tau}, a_{h}^{\tau} ; \omega\right)\right)^{2}\right\} \tag{3.5}
\end{equation*}
$$

where $\sigma>0$ is an absolute constant. In this case, the maximum a posteriori (MAP) estimate of $\omega$ coincides with the weight estimate $\widehat{\omega}_{h}^{k}$ obtained by LSVI. Such observation suggests that sampling from the posterior distribution $p\left(\omega \mid\left\{\widetilde{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)\right\}_{\tau \in[k-1]}, \mathcal{D}_{h}^{k}\right)$ would allows us to both achieve exploitation based on the available data, and generate noise for randomized exploration. More specifically, the $\Delta \omega(t)$ in 3.4) takes the form of

$$
\Delta \omega(t)=\frac{\epsilon_{t}}{2} \cdot\left[-\omega(t)+\frac{1}{\sigma^{2}} \sum_{\tau=1}^{k-1}\left(y_{h}^{\tau}-Q\left(s_{h}^{\tau}, a_{h}^{\tau} ; \omega(t)\right)\right) \cdot \nabla_{\omega} Q\left(s_{h}^{\tau}, a_{h}^{\tau} ; \omega(t)\right)\right]+\eta_{t}
$$

which can be viewed as using stochastic gradient descent to solve for the LSVI update (3.1) with $\lambda=\sigma^{2}$ and $\widehat{V}_{h+1}^{k}$ replaced by $\widetilde{V}_{h+1}^{k}$. See also in Lemma E. 1 for the closed form of the posterior in a linear MDP setting with Gaussian prior and likelihood.
We would like to highlight that, although the theoretical guarantee in this paper is built upon the Gaussian prior and likelihood, the choices of prior and likelihood are flexible. For example, we can use uninformative prior for generalized linear model, and Gaussian process prior for kernel of overparameterized neural networks [6]. Moreover, as suggested by theoretical results on Langevin Dynamics [56] and illustrative experiments in Appendix F] the computational overhead caused by replacing the minimization of (3.1) with posterior sampling is mild.

```
Algorithm 1 Bootstrapped Value Iteration (BooVI)
    Require: \(\operatorname{MDP}(\mathcal{S}, \mathcal{A}, H, \mathcal{P}, r)\), action-value function parameterization \(Q(\cdot, \cdot ; \cdot): \mathcal{S} \times \mathcal{A} \times\)
    \(\mathbb{R}^{d} \rightarrow \mathbb{R}\), number of episodes \(K\), number of posterior weights \(\left\{N_{k}\right\}_{k \in[K]}\), lower and upper
    bootstrapping ratios \(\alpha, \beta \in(0,1)\)
    Initialize the data buffer \(\mathcal{D}_{h}^{1} \leftarrow\{ \}\) for \(h \in[H]\)
    For episode \(k=1, \ldots, K\) do
        Set \(N_{k, \alpha} \leftarrow\left\lceil\alpha \cdot N_{k}\right\rceil, N_{k, \beta} \leftarrow\left\lfloor\beta \cdot N_{k}\right\rfloor\), and \(\omega_{H+1}^{k, i} \leftarrow 0\) for all \(i \in\left[N_{k}\right]\)
        Sample \(n_{k}\) uniformly from \(\left\{N_{k, \alpha}, N_{k, \alpha}+1, \ldots, N_{k, \beta}\right\}\)
        For timestep \(h=H, \ldots, 1\) do
            Generate \(\widetilde{Q}_{h+1}^{k}\left(s_{h+1}^{\tau}, a\right)\) using Algorithm 2 with weights \(\left\{\omega_{h+1}^{k, i}\right\}_{i \in\left[N_{k}\right]}\) and parameter \(n_{k}\)
            for all \(a \in \mathcal{A}\) and \(\tau \in[k-1]\)
            \(\widetilde{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right) \leftarrow \max _{a \in \mathcal{A}} \widetilde{Q}_{h+1}^{k}\left(s_{h+1}^{\tau}, a\right)\) for all \(\tau \in[k-1]\)
            Independently sample \(\left\{\omega_{h}^{k, i}\right\}_{i \in\left[N_{k}\right]}\) from the posterior \(p\left(\omega \mid\left\{\widetilde{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)\right\}_{\tau \in[k-1]}, \mathcal{D}_{h}^{k}\right)\)
            defined in (3.3), e.g., using Langevin dynamics in (3.4)
        end
        For timestep \(h=1, \ldots, H\) do
            Generate \(\widetilde{Q}_{h}^{k}\left(s_{h}^{k}, a\right)\) using Algorithm 2 with weights \(\left\{\omega_{h}^{k, i}\right\}_{i \in\left[N_{k}\right]}\) and parameter \(n_{k}\) for
            all \(a \in \mathcal{A}\)
            Take action \(a_{h}^{k} \leftarrow \operatorname{argmax}_{a \in \mathcal{A}} \widetilde{Q}_{h}^{k}\left(s_{h}^{k}, a\right)\), and observe \(r_{h}^{k}\) and \(s_{h+1}^{k}\)
            Update the data buffer \(\mathcal{D}_{h}^{k+1} \leftarrow \mathcal{D}_{h}^{k} \cup\left\{\left(s_{h}^{k}, a_{h}^{k}, r_{h}^{k}, s_{h+1}^{k}\right)\right\}\)
        End
    End
```

The following Algorithm 2 serves as the critical building block of BooVI for enforcing the optimism of the estimated action-value. At episode $k$ in BooVI, for any $(s, a) \in \mathcal{S} \times \mathcal{A}$, Algorithm 2 ranks the estimated action-value functions corresponding to the posterior sample obtained in Line 1 of Algorithm 1 in ascending order. Then, in order to enforce the optimism of the estimated action-value function, Algorithm 2 resamples the $n_{k}$-th top-ranked value from the ordered estimated action-value functions in the manner of bootstrapping. Finally, to ensure sufficient optimism of the obtained bootstrapped action-value function, we extrapolate with a tunable parameter $\nu>1$.

```
Algorithm 2 Bootstrapping Action-Value Function
    Require: Action-value function parameterization \(Q(\cdot, \cdot ; \cdot): \mathcal{S} \times \mathcal{A} \times \mathbb{R}^{d} \rightarrow \mathbb{R}\), posterior sample
    \(\left\{\omega_{h}^{k, i}\right\}_{i \in\left[N_{k}\right]}\), integer \(n_{k} \in\left[N_{k}\right]\), extrapolation parameter \(\nu>1\), and state-action pair \((s, a)\)
    Compute \(Q_{h}^{k, i}(s, a) \leftarrow Q\left(s, a ; \omega_{h}^{k, i}\right)\) for all \(i \in\left[N_{k}\right]\)
    Set \(\widehat{Q}_{h}^{k}(s, a) \leftarrow\left(1 / N_{k}\right) \sum_{i=1}^{N_{k}} Q_{h}^{k, i}(s, a)\)
    Rank \(\left\{Q_{h}^{k, i}(s, a)\right\}_{i \in\left[N_{k}\right]}\) in ascending order to obtain
    \(\left\{Q_{h}^{k,(i)}(s, a)\right\}_{i \in\left[N_{k}\right]}\)
    5: Output: \(\widetilde{Q}_{h}^{k}(s, a) \leftarrow \min \left\{(1-\nu) \cdot \widehat{Q}_{h}^{k}(s, a)+\nu \cdot Q_{h}^{k,\left(n_{k}\right)}(s, a), H-h+1\right\}^{+}\)
```

Remark 3.1 (Sample Efficiency and Computational Efficiency). Here we clarify the differences between sample efficiency and computational efficiency. Throughout this paper, we refer sample efficiency to the learning efficiency with respect to the total number $T=K H$ of interactions with the environment. Such interactions can often time be very resource demanding. Thus, this paper seeks to provide an algorithm that aims to achieve the sample efficiency with respect to the number of such interactions. Although posterior weights are sampled in Algorithm 1, the efficiency of such posterior sampling process is categorized as computational efficiency since it is purely based on the current data buffer $\mathcal{D}_{h}^{k}$ and does not require any extra interaction with the environment.
Remark 3.2 (Additional Computational Cost). Here we discuss the computational cost of BooVI compared to LSVI. (i) Posterior sampling vs. least-square minimization: Consider the Langevin dynamics in (3.4). Since posterior sampling replaces and admits similar steps as least-square minimization, the computational cost should not differ significantly between LSVI and BooVI. (ii) Computing bootstrapped Q-function: To compute bootstrapped Q-function, at each state-action pair,
the extra sorting takes $O\left(N_{k} \log N_{k}\right)$ time complexity, which makes the action selection procedure $O\left(N_{k} \log N_{k}\right)$ times slower. In addition, keeping $N_{k}$ posterior weights takes $O\left(d \cdot N_{k}\right)$ memory. We note that, although the requirement on $N_{k}$ in our subsequent analysis is large, we should be able to use much smaller $N_{k}$ in practice due to the pessimistic nature of the analysis. See Appendix $F$ for illustrative experiment results.

It is worth mentioning that BooVI is applicable to any specific parameterization $Q(\cdot, \cdot ; \cdot): \mathcal{S} \times \mathcal{A} \times$ $\mathbb{R}^{d} \rightarrow \mathbb{R}$ of the action-value function and is fully data-driven. We also note here that the design of BooVI does not rely on the Gaussian prior/posterior or the Langevin dynamics sampling technique.

## 4 Main Results

While BooVI (Algorithm 1) is generally applicable to different parameterizations of the action-value function, we establish its regret for a class of linear MDPs.

Linear Markov Decision Process. We consider a class of MDPs, where the transition kernels and the reward functions are linear in the same feature mapping. Specifically, we have the following definition.

Definition 4.1 (Linear MDP). An $\operatorname{MDP}(\mathcal{S}, \mathcal{A}, H, \mathcal{P}, r)$ is called a linear MDP with feature mapping $\phi: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^{d}$, if for any $h \in[H]$, there exists $d$ unknown signed measures $\mu_{h}=\left(\mu_{h}^{(1)}, \ldots, \mu_{h}^{(d)}\right)$ over $\mathcal{S}$ and an unknown vector $\theta_{h} \in \mathbb{R}^{d}$, such that for any $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have

$$
\begin{equation*}
\mathcal{P}_{h}(\cdot \mid s, a)=\left\langle\phi(s, a), \mu_{h}(\cdot)\right\rangle, \quad r_{h}(s, a)=\left\langle\phi(s, a), \theta_{h}\right\rangle . \tag{4.1}
\end{equation*}
$$

See [51, 52, 24] for examples of such a class of linear MDPs. Specifically, examples include tabular MDPs where $d=S A$ and the feature mapping is the canonical basis $\phi(s, a)=e_{(s, a)}$, and the simplex feature space where the feature space $\{\phi(s, a) \mid(s, a) \in \mathcal{S} \times \mathcal{A}\}$ is a subset of the $d$-dimensional simplex. See also [14, 44, 25] for related discussions on such a linear representation. For notational simplicity, we write the feature $\phi\left(s_{h}^{k}, a_{h}^{k}\right)$ as $\phi_{h}^{k}$ throughout the rest of this paper.
Based on Definition 4.1, without loss of generality, we make the following assumption.
Assumption 4.2. The $\operatorname{MDP}(\mathcal{S}, \mathcal{A}, H, \mathcal{P}, r)$ is a linear $\operatorname{MDP}$ with $\|\phi(s, a)\| \leq 1$ for all $(s, a) \in \mathcal{S} \times \mathcal{A}$ and $\max \left\{\left\|\mu_{h}(\mathcal{S})\right\|,\left\|\theta_{h}\right\|\right\} \leq \sqrt{d}$ for all $h \in[H]$.

To motivate the linear parameterization of the action-value function in the subsequent section, we have the following proposition.
Proposition 4.3 (Linear Action-Value Function, Proposition 2.3 in [24]). For a linear MDP, for any policy $\pi$ and $h \in[H]$, there exists $\omega_{h}^{\pi} \in \mathbb{R}^{d}$ such that for all $(s, a) \in \mathcal{S} \times \mathcal{A}$, we have $Q_{h}^{\pi}(s, a)=\left\langle\phi(s, a), \omega_{h}^{\pi}\right\rangle$.

Regret of BooVI for Linear MDPs. For a linear MDP, by Proposition 4.3 t the parameterization $Q(\cdot, \cdot ; \cdot): \mathcal{S} \times \mathcal{A} \times \mathbb{R}^{d} \rightarrow \mathbb{R}$ in BooVI should take the form of $Q(s, a ; \omega)=\langle\phi(s, a), \omega\rangle$, where $\omega \in \mathbb{R}^{d}$. Furthermore, when the prior and the likelihood are Gaussian, the posterior of $\omega$ is Gaussian distribution with mean being is the weight update by LSVI with chioce of $\lambda=\sigma^{2}$.
For any bootstrapping ratio $q \in(0,1)$, we write $C_{q}=\Phi^{-1}(q)$ throughout the rest of this paper, where $\Phi(\cdot)$ is the cumulative distribution function of the standard Gaussian. Then, we have the following upper bound of regret of BooVI.
Theorem 4.4. Suppose that Assumption 4.2 holds, and that $d \geq 2$. We set $\sigma=1, \nu=d H$, and fix a failure probability $p \in(0,1]$. Then there exists a sequence of posterior sample sizes $\left\{N_{k}\right\}_{k \in[K]}$, and there exist absolute constants $c_{\beta}>c_{\alpha}>0$, such that if $N_{k}=O\left(d^{6} T^{4} k / p^{4}\right)$ for all $k \in[K]$, $C_{\alpha}=c_{\alpha} \cdot \sqrt{\iota}$, and $C_{\beta}=c_{\beta} \cdot \sqrt{\iota}$, BooVI (Algorithm 1) with Gaussian posterior in (3.5) satisfies

$$
\operatorname{Regret}(K)=O\left(\sqrt{d^{3} H^{3} T \iota^{2}}\right)
$$

with probability at least $1-p$, where $T=K H$ and $\iota=\log (3 d T / p)$.
Proof. See Section 5 for a proof sketch.

## 5 Proof Sketch

Notation. At episode $k$, given the number of posterior sample weights $N_{k}$ and given $n_{k}$ sampled in Line 1 of Algorithm 1, we define $C_{k}=\Phi^{-1}\left(n_{k} / N_{k}\right)$. At timestep $h$ of episode $k$, further given posterior weights $\left\{\omega_{h}^{k, i}\right\}_{i \in\left[N_{k}\right]}$, corresponding to the bootstrapped action-value function $\widetilde{Q}_{h}^{k}$ generated by Algorithm 2, we define the bootstrapped value function as

$$
\begin{equation*}
\widetilde{V}_{h}^{k}(s)=\max _{a \in \mathcal{A}} \widetilde{Q}_{h}^{k}(s, a) \tag{5.1}
\end{equation*}
$$

Next, we define the mean bootstrapped action-value functions as

$$
\begin{equation*}
\bar{Q}_{h}^{k}(s, a)=\min \left\{\mathbb{E}_{\omega}\left[(1-\nu) \cdot \widehat{Q}_{h}^{k}(s, a)+\nu \cdot Q_{h}^{k,\left(n_{k}\right)}(s, a)\right], H-h+1\right\}^{+} \tag{5.2}
\end{equation*}
$$

where $\mathbb{E}_{\omega}[\cdot]$ is taken over $\omega$ with respect to the posterior $p\left(\omega \mid\left\{\widetilde{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)\right\}_{\tau \in[k-1]}, \mathcal{D}_{h}^{k}\right)$ defined in (3.3). Correspondingly, we define the mean bootstrapped value function as

$$
\begin{equation*}
\bar{V}_{h}^{k}(s)=\max _{a \in \mathcal{A}} \bar{Q}_{h}^{k}(s, a) . \tag{5.3}
\end{equation*}
$$

Also, we denote by

$$
\begin{equation*}
\bar{\omega}_{h}^{k}=\left(\Lambda_{h}^{k}\right)^{-1} \sum_{\tau=1}^{k-1} y_{h}^{\tau} \cdot \phi_{h}^{\tau} \tag{5.4}
\end{equation*}
$$

the mean of the posterior distribution $p\left(\omega \mid\left\{\widetilde{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)\right\}_{\tau \in[k-1]}, \mathcal{D}_{h}^{k}\right)$, where $y_{h}^{\tau}$ is defined in (3.2). For a fixed failure probability $p \in(0,1]$, we write $\iota=\log (3 d T / p)$, where $T=K H$. Finally, we define matrix $\Lambda_{h}^{k} \in \mathbb{R}^{d \times d}$ by

$$
\begin{equation*}
\Lambda_{h}^{k}=\sum_{\tau=1}^{k-1} \phi_{h}^{\tau}\left(\phi_{h}^{\tau}\right)^{\top}+\sigma^{2} \cdot I_{d} \tag{5.5}
\end{equation*}
$$

Concentration Events. Before proving Theorem 4.4. we first present two lemmas, each characterizing an event that is involved throughout the remaining proofs. The first lemma characterizes the concentration behavior of the bootstrapped action-value function $\widetilde{Q}_{h}^{k}$.
Lemma 5.1. Let $\sigma=1, \nu=d H$, and $C_{\beta}=c_{\beta} \cdot \sqrt{\iota}$ for some $c_{\beta}>0$. For a fixed failure probability $p \in(0,1]$, we define $\mathcal{E}$ as the event that the condition

$$
\left|\widetilde{Q}_{h}^{k}(s, a)-\bar{Q}_{h}^{k}(s, a)\right| \leq\left(c_{\beta} / \sqrt{d}+3\right) \cdot H \sqrt{d \iota / k}
$$

is satisfied for all $(s, a, k, h) \in \mathcal{S} \times \mathcal{A} \times[K] \times[H]$. Then, there exists a sequence of posterior sample sizes $\left\{N_{k}\right\}_{k \in[K]}$ satisfying $N_{k}=O\left(d^{6} T^{4} k / p^{4}\right)$ for all $k \in[K]$, such that $P(\mathcal{E}) \geq 1-p / 3$.

Proof. See Appendix Cfor a detailed proof.

The second lemma characterizes the concentration behavior of the mean bootstrapped state value function $\bar{V}_{h}^{k}$.
Lemma 5.2. Let $\sigma=1, \nu=d H$, and $C_{\beta}=c_{\beta} \cdot \sqrt{\iota}$ for some constant $c_{\beta}>0$. For a fixed failure $p \in(0,1]$, we define $\mathcal{E}^{\prime}$ as the event that the condition

$$
\left\|\sum_{\tau=1}^{k-1} \phi_{h}^{\tau} \cdot\left(\bar{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)-\left[\mathcal{P}_{h} \bar{V}_{h+1}^{k}\right]\left(s_{h}^{\tau}, a_{h}^{\tau}\right)\right)\right\|_{\left(\Lambda_{h}^{k}\right)^{-1}} \leq C \cdot d H \sqrt{\chi}
$$

is satisfied for all $(k, h) \in[K] \times[H]$, where $\left.\chi=\log \left[3\left(1+c_{\beta}\right) d T / p\right)\right]$. Then we have $\mathbb{P}\left(\mathcal{E}^{\prime}\right) \geq 1-p / 3$.
Proof. See Appendix Dfor a detailed proof.

Model Estimation Error. With the events $\mathcal{E}$ and $\mathcal{E}^{\prime}$ ready, we can proceed to characterize the model estimation error $\zeta_{h}^{k}: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ for all $(k, h) \in[K] \times[H]$ defined as

$$
\begin{equation*}
\zeta_{h}^{k}=\bar{Q}_{h}^{k}-\left(r_{h}+\left[\mathcal{P}_{h} \bar{V}_{h+1}^{k}\right]\right), \tag{5.6}
\end{equation*}
$$

which can be comprehended as the estimation error of the model at timestep $h$ of episode $k$ induced by the mean bootstrapped value function $\bar{Q}_{h+1}^{k}$. We have the following lemma charaterizing the model estimation error $\zeta_{h}^{k}$ defined in 5.6.
Lemma 5.3 (Model Estimation Error). Let $d \geq 2, \nu=d H, \sigma=1$, and $p \in(0,1]$. Then there exists a sequence of posterior sample sizes $\left\{N_{k}\right\}_{k \in[K]}$, and there exist absolute constants $c_{\beta}>c_{\alpha}>0$ such that, if $N_{k}=O\left(d^{6} T^{4} k / p^{4}\right)$ for all $k \in[K], C_{\alpha}=c_{\alpha} \cdot \sqrt{\iota}$, and $C_{\beta}=c_{\beta} \cdot \sqrt{\iota}$, for the model estimation error $\zeta_{h}^{k}$ defined in (5.6), we under events $\mathcal{E}$ and $\mathcal{E}^{\prime}$ that

$$
0 \leq \zeta_{h}^{k}(s, a) \leq\left(c_{\alpha}+c_{\beta}\right) \cdot d H \sqrt{\iota} \sqrt{\phi(s, a)^{\top}\left(\Lambda_{h}^{k}\right)^{-1} \phi(s, a)}
$$

for all $(s, a, k, h) \in \mathcal{S} \times \mathcal{A} \times[K] \times[H]$.
Proof. Under Assumption 4.2, using (4.1) we have $\left[\mathcal{P}_{h} \bar{V}_{h+1}^{k}\right]=\int_{\mathcal{S}} \bar{V}_{h+1}^{k}\left(s^{\prime}\right) \cdot\left\langle\phi, \mu_{h}\left(\mathrm{~d} s^{\prime}\right)\right\rangle$. Then, with slight abuse of notation, we denote by $\bar{V}_{h+1}^{k}\left(\mu_{h}\right)=\int_{\mathcal{S}} \bar{V}_{h+1}^{k}(s) \mu_{h}(\mathrm{~d} s) \in \mathbb{R}^{d}$ and write

$$
\begin{aligned}
\mathcal{P}_{h} \bar{V}_{h+1}^{k} & =\left\langle\phi,\left(\Lambda_{h}^{k}\right)^{-1}\left(\Lambda_{h}^{k}\right) \bar{V}_{h+1}^{k}\left(\mu_{h}\right)\right\rangle \\
& =\left\langle\phi,\left(\Lambda_{h}^{k}\right)^{-1}\left(\sum_{\tau=1}^{k-1} \phi_{h}^{\tau}\left(\phi_{h}^{\tau}\right)^{\top} \bar{V}_{h+1}^{k}\left(\mu_{h}\right)\right)\right\rangle+\left\langle\phi,\left(\Lambda_{h}^{k}\right)^{-1} \bar{V}_{h+1}^{k}\left(\mu_{h}\right)\right\rangle \\
& =\left\langle\phi,\left(\Lambda_{h}^{k}\right)^{-1}\left(\sum_{\tau=1}^{k-1} \phi_{h}^{\tau} \cdot\left[\mathcal{P}_{h} \bar{V}_{h+1}^{k}\right]\left(s_{h}^{\tau}, a_{h}^{\tau}\right)\right)\right\rangle+\left\langle\phi,\left(\Lambda_{h}^{k}\right)^{-1} \bar{V}_{h+1}^{k}\left(\mu_{h}\right)\right\rangle,
\end{aligned}
$$

where the second equality follows from (5.5). Thus, we have the following decomposition

$$
\left|\left\langle\phi, \bar{\omega}_{h}^{k}\right\rangle-\left(r_{h}+\mathcal{P}_{h} \bar{V}_{h+1}^{k}\right)\right| \leq u_{1}+u_{2},
$$

where $\bar{\omega}_{h}^{k}$ is defined in (5.4), and

$$
\begin{aligned}
& u_{1}=\left|\left\langle\phi,\left(\Lambda_{h}^{k}\right)^{-1}\left[\sum_{\tau=1}^{k-1} \phi_{h}^{\tau} \cdot\left(\widetilde{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)-\left[\mathcal{P}_{h} \bar{V}_{h+1}^{k}\right]\left(s_{h}^{\tau}, a_{h}^{\tau}\right)\right)\right]\right\rangle\right| \\
& u_{2}=\left|\left\langle\phi,\left(\Lambda_{h}^{k}\right)^{-1}\left(\sum_{\tau=1}^{k-1} r_{h}^{\tau} \cdot \phi_{h}^{\tau}-\bar{V}_{h+1}^{k}\left(\mu_{h}\right)\right)\right\rangle-r_{h}\right|
\end{aligned}
$$

In the sequel, we upper bound $u_{1}$ and $u_{2}$ separately.
Upper bounding $u_{1}$ : First, we further decompose $u_{1}$ as
$\left\langle\phi,\left(\Lambda_{h}^{k}\right)^{-1}\left\{\sum_{\tau=1}^{k-1} \phi_{h}^{\tau} \cdot\left[\left(\widetilde{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)-\bar{V}_{h+1,}^{k}\left(s_{h+1}^{\tau}\right)\right)+\left(\bar{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)-\left[\mathcal{P}_{h} \bar{V}_{h+1}^{k}\right]\left(s_{h}^{\tau}, a_{h}^{\tau}\right)\right)\right]\right\}\right\rangle$,
applying the Cauchy-Schwartz inequality to which we obtain

$$
\begin{align*}
u_{1} \leq\langle\phi & \left.\left(\Lambda_{h}^{k}\right)^{-1}\left[\sum_{\tau=1}^{k-1} \phi_{h}^{\tau} \cdot\left(\widetilde{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)-\bar{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)\right)\right]\right\rangle  \tag{5.7}\\
& +\sqrt{\phi^{\top}\left(\Lambda_{h}^{k}\right)^{-1} \phi} \cdot\left\|\sum_{\tau=1}^{k-1} \phi_{h}^{\tau} \cdot\left(\bar{V}_{h+1}^{k}\left(s_{h}^{\tau}, a_{h}^{\tau}\right)-\left[\mathcal{P}_{h} \bar{V}_{h+1}^{k}\right]\left(s_{h+1}^{\tau}\right)\right)\right\|_{\left(\Lambda_{h}^{k}\right)^{-1}} .
\end{align*}
$$

Now we proceed to upper bound the two terms on the right-hand side of 5.7) in the following. For the first term, since by Lemma 5.1 we have under event $\mathcal{E}$ that $\left|\widetilde{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)-\bar{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)\right| \leq$ $\left(c_{\beta} / \sqrt{d}+3\right) \cdot H \sqrt{d \iota / k}$, we have

$$
\begin{equation*}
\left|\left\langle\phi,\left(\Lambda_{h}^{k}\right)^{-1}\left[\sum_{\tau=1}^{k-1} \phi_{h}^{\tau} \cdot\left(\widetilde{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)-\bar{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)\right)\right]\right\rangle\right| \leq\left(c_{\beta} / \sqrt{d}+3\right) \cdot H \sqrt{d \iota / k} \sum_{\tau=1}^{k-1} \phi^{\top}\left(\Lambda_{h}^{k}\right)^{-1} \phi_{h}^{\tau} . \tag{5.8}
\end{equation*}
$$

By the Cauchy-Schwartz inequality and Lemma D. 1 in [24] (see Appendix E), we have $\sum_{\tau=1}^{k-1} \phi^{\top}\left(\Lambda_{h}^{k}\right)^{-1} \phi_{h}^{\tau} \leq\left[\left(\sum_{\tau=1}^{k-1} \phi^{\top}\left(\Lambda_{h}^{k}\right)^{-1} \phi_{h}^{\tau}\right) \cdot\left(\sum_{\tau=1}^{k-1} \phi^{\top}\left(\Lambda_{h}^{k}\right)^{-1} \phi\right)\right]^{1 / 2} \leq \sqrt{d \iota / k} \cdot \sqrt{d k} \sqrt{\phi^{\top}\left(\Lambda_{h}^{k}\right)^{-1} \phi}$, taking which into (5.8) and gives

$$
\begin{equation*}
\left|\left\langle\phi,\left(\Lambda_{h}^{k}\right)^{-1}\left[\sum_{\tau=1}^{k-1} \phi_{h}^{\tau} \cdot\left(\widetilde{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)-\bar{V}_{h+1}^{k}\left(s_{h+1}^{\tau}\right)\right)\right]\right\rangle\right| \leq\left(c_{\beta} / \sqrt{d}+3\right) \cdot d H \sqrt{\iota} \sqrt{\phi^{\top}\left(\Lambda_{h}^{k}\right)^{-1} \phi} \tag{5.9}
\end{equation*}
$$

Taking (5.9) into (5.7) and applying Lemma 5.2 to the second term on the right-hand side of 5.7, we have under the events $\mathcal{E}$ and $\mathcal{E}^{\prime}$ that

$$
\begin{equation*}
u_{1} \leq\left[\left(c_{\beta} / \sqrt{d}+3\right) \cdot d H \sqrt{\iota}+C \cdot d H \sqrt{\chi}\right] \cdot \sqrt{\phi^{\top}\left(\Lambda_{h}^{k}\right)^{-1} \phi} \tag{5.10}
\end{equation*}
$$

Upper bounding $u_{2}$ : By (4.1), we have

$$
u_{2}=\left|\left\langle\phi,\left(\Lambda_{h}^{k}\right)^{-1}\left(\theta_{h}+\bar{V}_{h+1}^{k}\left(\mu_{h}\right)\right)\right\rangle\right| .
$$

Then, by the Cauchy-Schwartz inequality, we further obtain

$$
\begin{align*}
u_{2} & \leq \sqrt{\phi^{\top}\left(\Lambda_{h}^{k}\right)^{-1} \phi} \cdot\left(\left\|\bar{V}_{h+1}^{k}\left(\mu_{h}\right)\right\|_{\left(\Lambda_{h}^{k}\right)^{-1}}+\left\|\theta_{h}\right\|_{\left(\Lambda_{h}^{k}\right)^{-1}}\right) \\
& \leq \sqrt{\phi^{\top}\left(\Lambda_{h}^{k}\right)^{-1} \phi} \cdot\left(\left\|\bar{V}_{h+1}^{k}\left(\mu_{h}\right)\right\|+\left\|\theta_{h}\right\|\right) \leq 2 H \sqrt{d} \cdot \sqrt{\phi^{\top}\left(\Lambda_{h}^{k}\right)^{-1} \phi} \tag{5.11}
\end{align*}
$$

where the second inequality follows from $\Lambda_{h}^{k} \succeq I_{d}$, and the last inequality follows from $\left\|\theta_{h}\right\| \leq \sqrt{d}$ in Assumption (4.2) as well as the fact that $\bar{V}_{h+1}^{k}(s) \leq H-h \leq H-1$ for all $s \in \mathcal{S}$.
Combining (5.10) and 5.11), we obtain under event $\mathcal{E}$ that

$$
\begin{align*}
& \left|\left\langle\phi, \bar{\omega}_{h}^{k}\right\rangle-\left(r_{h}+\mathcal{P}_{h} \bar{V}_{h+1}^{k}\right)\right| / \sqrt{\phi^{\top}\left(\Lambda_{h}^{k}\right)^{-1} \phi}  \tag{5.12}\\
& \leq\left[\left(c_{\beta} / \sqrt{d}+3\right) \cdot d H \sqrt{\iota}+C \cdot d H \sqrt{\chi}+2 H \sqrt{d}\right] \\
& \leq\left[\left(c_{\beta} / \sqrt{d}+5\right) \cdot d H \sqrt{\iota}+C \cdot d H \sqrt{\chi}\right]=C^{\prime} \cdot d H \sqrt{\iota}
\end{align*}
$$

where $C^{\prime}>0$ is an absolute constant. Next, we need to find an absolute constant $c_{\beta}>0$ such that $C^{\prime} \cdot \sqrt{\iota}=\left(c_{\beta} / \sqrt{d}+5\right) \cdot \sqrt{\iota}+C \cdot \sqrt{\iota+\log \left(1+c_{\beta}\right)}<c_{\beta} \cdot \sqrt{\iota}$. Note that $d \geq 2$ and $\iota \geq \log 2$, it suffices to pick a $c_{\beta}>0$ such that

$$
\begin{equation*}
C \cdot \sqrt{\log 2+\log \left(1+c_{\beta}\right)}<\left[(1-1 / \sqrt{2}) c_{\beta}-5\right] \cdot \sqrt{\log 2} \tag{5.13}
\end{equation*}
$$

which must exist as the left hand side grows logarithmically in $c_{\beta}$ and the right-hand side grows linearly in $c_{\beta}$. For $c_{\beta}>0$ satisfying (5.13), we pick any $c_{\alpha}>0$ such that $C^{\prime} \leq c_{\alpha}<c_{\beta}$ and let $C_{\alpha}=c_{\alpha} \cdot \sqrt{\iota}$. By (5.12) and $r_{h}+\left[\mathcal{P}_{h} \bar{V}_{h+1}^{k}\right] \leq H-h+1$, we obtain under events $\mathcal{E}$ and $\mathcal{E}^{\prime}$ that

$$
\begin{aligned}
\zeta_{h}^{k} & =\min \left\{\left\langle\phi, \bar{\omega}_{h}^{k}\right\rangle+C_{k} \cdot \nu \cdot \sqrt{\phi^{\top}\left(\Lambda_{h}^{k}\right)^{-1} \phi}, H-h+1\right\}^{+}-\left(r_{h}+\mathcal{P}_{h} \bar{V}_{h+1}^{k}\right) \\
& \geq \min \left\{\left(c_{\alpha}-C^{\prime}\right) \cdot d H \sqrt{\iota} \sqrt{\phi^{\top}\left(\Lambda_{h}^{k}\right)^{-1} \phi}, 0\right\} \geq 0
\end{aligned}
$$

On the other hand, by (5.12), we also have under events $\mathcal{E}$ and $\mathcal{E}^{\prime}$ that

$$
\zeta_{h}^{k} \leq\left\langle\phi, \bar{\omega}_{h}^{k}\right\rangle-\left(r_{h}+\mathcal{P}_{h} \bar{V}_{h+1}^{k}\right)+C_{\beta} \cdot \nu \cdot \sqrt{\phi^{\top}\left(\Lambda_{h}^{k}\right)^{-1} \phi} \leq\left(c_{\alpha}+c_{\beta}\right) \cdot d H \sqrt{\iota} \sqrt{\phi^{\top}\left(\Lambda_{h}^{k}\right)^{-1} \phi}
$$

Therefore, we finish the proof of Lemma 5.3
As results of Lemma 5.3. we have the following two lemmas characterizing the optimism and the cumulative estimation error of the mean bootstrapped value function $\bar{V}_{h}^{k}$, respectively.

Lemma 5.4 (Optimistic Random Value Function). Let $\sigma=1$ and $\nu=d H$. There there exists a sequence of posterior sample sizes $\left\{N_{k}\right\}_{k \in[K]}$, and there exist absolute constants $c_{\beta}>c_{\alpha}>0$, such that if $N_{k}=O\left(d^{6} T^{4} k / p^{4}\right)$ for all $k \in[K], C_{\alpha}=c_{\alpha} \cdot \sqrt{\iota}$, and $C_{\beta}=c_{\beta} \cdot \sqrt{\iota}$, we have under the events $\mathcal{E}$ and $\mathcal{E}^{\prime}$ that

$$
\sum_{k=1}^{K}\left[V_{1}^{*}\left(s_{1}^{k}\right)-\bar{V}_{1}^{k}\left(s_{1}^{k}\right)\right] \leq 0
$$

Proof. See Appendix Dfor a detailed proof.
Lemma 5.5 (Cumulative Estimation Error). Let $\sigma=1, \nu=d H$, and $p \in(0,1]$. There exists a sequence of posterior sample sizes $\left\{N_{k}\right\}_{k \in[K]}$, and there exist absolute constants $c_{\beta}>c_{\alpha}>0$, such that if $N_{k}=O\left(d^{6} T^{4} k / p^{4}\right)$ for all $k \in[K], C_{\alpha}=c_{\alpha} \cdot \sqrt{\iota}$, and $C_{\beta}=c_{\beta} \cdot \sqrt{\iota}$, we have under the events $\mathcal{E}$ and $\mathcal{E}^{\prime}$ that

$$
\begin{gathered}
\sum_{k=1}^{K}\left[\bar{V}_{1}^{k}\left(s_{1}^{k}\right)-V_{1}^{\pi^{k}}\left(s_{1}^{k}\right)\right] \leq \sqrt{18 T H^{2} \cdot \log (3 / p)}+4\left(c_{\beta} / \sqrt{d}+3\right) \cdot H^{2} \sqrt{d K \iota} \\
+\left(c_{\alpha}+c_{\beta}\right) \cdot d H^{2} \sqrt{2 d K \iota^{2}}
\end{gathered}
$$

with probability at least $1-p / 3$.
Proof. See Appendix Dfor a detailed proof.
Finally, the regret bound of BooVI for the class of linear MDPs can be established as a consequence of the optimism (Lemmas 5.4) and the upper bound of the cumulative estimation error (Lemma 5.5) of the bootstrapped value function.

Proof of Theorem 4.4 First, recall that we have $\bar{V}_{h}^{k}(s)$ defined in 5.3). We have the regret decomposition

$$
\begin{equation*}
\operatorname{Regret}(K)=\sum_{k=1}^{K}\left[V_{1}^{*}\left(s_{1}^{k}\right)-\bar{V}_{1}^{k}\left(s_{1}^{k}\right)\right]+\sum_{k=1}^{K}\left[\bar{V}_{1}^{k}\left(s_{1}^{k}\right)-V_{1}^{\pi^{k}}\left(s_{1}^{k}\right)\right] \tag{5.14}
\end{equation*}
$$

Applying Lemmas 5.4 and 5.5 to 5.14, we obtain under the events $\mathcal{E}$ and $\mathcal{E}^{\prime}$ that

$$
\begin{align*}
& \operatorname{Regret}(K) \leq 0+ \sqrt{18 T H^{2} \cdot \log (3 / p)}+4\left(c_{\beta} / \sqrt{d}+3\right) \cdot H^{2} \sqrt{d K \iota} \\
&+\left(c_{\alpha}+c_{\beta}\right) \cdot d H^{2} \sqrt{2 d K \iota^{2}} \\
&=O\left(\sqrt{d^{3} H^{3} T \iota^{2}}\right) \tag{5.15}
\end{align*}
$$

with probability at least $1-p / 3$. Finally, by Lemmas 5.1 and 5.2 , we have with at least probability $1-p / 3-p / 3=1-2 p / 3$ that the events $\mathcal{E}$ and $\mathcal{E}^{\prime}$ hold simultaneously. Thus, we have 5.15 hold with probability at least $1-p$, which concludes the proof.

## 6 Conclusion

The cost of collecting data via online experiments is often prohibitive compared to collecting offline data, e.g., via posterior sampling. Thus, designing online learning algorithms that provably explores the environment in a data-driven manner is essential. Moreover, applicability to general environments is critical since reinforcement learning tasks are getting more chanlleging and complex recently. We aspire to motivate the design of novel reinforcement learning algorithms that utilize cheaper data to boost online sample efficiency. Our algorithm and analysis serve as a step towards developing general applicable and provable sample-efficient reinforcement learning. While this paper is mainly on algorithmic and theoretical aspects of the bootstrapping idea in RL, it would be interesting to see the empirical strength of BooVI in more challenging RL environments. We leave this part to our future work.

## Acknowledgement

Zhaoran Wang acknowledges National Science Foundation (Awards 2048075, 2008827, 2015568, 1934931), Simons Institute (Theory of Reinforcement Learning), Amazon, J.P. Morgan, and Two Sigma for their supports.

Zhuoran Yang acknowledges Simons Institute (Theory of Reinforcement Learning).

## References

[1] AbbaSi-Yadkori, Y., PÁL, D. and Szepesvári, C. (2011). Improved algorithms for linear stochastic bandits. In Advances in Neural Information Processing Systems.
[2] Agarwal, A., Hsu, D., Kale, S., Langford, J., Li, L. and Schapire, R. (2014). Taming the monster: A fast and simple algorithm for contextual bandits. In International Conference on Machine Learning.
[3] Agrawal, S. and Jia, R. (2017). Optimistic posterior sampling for reinforcement learning: worst-case regret bounds. In Advances in Neural Information Processing Systems.
[4] Auer, P. and Ortner, R. (2006). Logarithmic online regret bounds for undiscounted reinforcement learning. In Advances in Neural Information Processing Systems.
[5] Azar, M. G., Osband, I. and Munos, R. (2017). Minimax regret bounds for reinforcement learning. In International Conference on Machine Learning.
[6] Bernardo, J., Berger, J., Dawid, A., Smith, A. et al. (1998). Regression and classification using gaussian process priors. Bayesian statistics 6475.
[7] Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J. and Zaremba, W. (2016). Openai gym. arXiv preprint arXiv:1606.01540 .
[8] CAI, Q., Yang, Z., Jin, C. and WANG, Z. (2019). Provably efficient exploration in policy optimization. arXiv preprint arXiv:1912.05830 .
[9] Chu, W., Li, L., Reyzin, L. and Schapire, R. (2011). Contextual bandits with linear payoff functions. In International Conference on Artificial Intelligence and Statistics.
[10] Dani, V., Hayes, T. P. and Kakade, S. M. (2008). Stochastic linear optimization under bandit feedback. Conference on Learning Theory .
[11] DANN, C. and BRUNSKILL, E. (2015). Sample complexity of episodic fixed-horizon reinforcement learning. In Advances in Neural Information Processing Systems.
[12] Dann, C., Lattimore, T. and Brunskill, E. (2017). Unifying PAC and regret: Uniform PAC bounds for episodic reinforcement learning. In Advances in Neural Information Processing Systems.
[13] Dong, K., Peng, J., Wang, Y. and Zhou, Y. (2019). $\sqrt{n}$-regret for learning in Markov decision processes with function approximation and low Bellman rank. arXiv preprint arXiv:1909.02506 .
[14] DU, S. S., Kakade, S. M., Wang, R. and Yang, L. F. (2019). Is a good representation sufficient for sample efficient reinforcement learning? arXiv preprint arXiv:1910.03016 .
[15] Du, S. S., Luo, Y., WANG, R. and Zhang, H. (2019). Provably efficient Q-learning with function approximation via distribution shift error checking oracle. arXiv preprint arXiv:1906.06321.
[16] Duan, Y., Chen, X., Houthooft, R., Schulman, J. and Abbeel, P. (2016). Benchmarking deep reinforcement learning for continuous control. In International Conference on Machine Learning.
[17] Ecoffet, A., Huizinga, J., Lehman, J., Stanley, K. O. and Clune, J. (2019). Go-explore: A new approach for hard-exploration problems. arXiv preprint arXiv:1901.10995.
[18] Foster, D. J., Agarwal, A., Dudík, M., Luo, H. and Schapire, R. E. (2018). Practical contextual bandits with regression oracles. arXiv preprint arXiv:1803.01088.
[19] Hafrnoja, T., Zhou, A., Abbeel, P. and Levine, S. (2018). Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. arXiv preprint arXiv:1801.01290 .
[20] Hsu, D., Kakade, S. and Zhang, T. (2012). A tail inequality for quadratic forms of subgaussian random vectors. Electronic Communications in Probability 17.
[21] JAKSCh, T., ORTNER, R. and AUER, P. (2010). Near-optimal regret bounds for reinforcement learning. Journal of Machine Learning Research 11 1563-1600.
[22] Jiang, N., Krishnamurthy, A., Agarwal, A., Langford, J. and Schapire, R. E. (2017). Contextual decision processes with low Bellman rank are PAC-learnable. In International Conference on Machine Learning.
[23] Jin, C., Allen-Zhu, Z., Bubeck, S. and Jordan, M. I. (2018). Is Q-learning provably efficient? In Advances in Neural Information Processing Systems.
[24] Jin, C., Yang, Z., Wang, Z. and Jordan, M. I. (2019). Provably efficient reinforcement learning with linear function approximation. arXiv preprint arXiv:1907.05388.
[25] Lattimore, T. and Szepesvari, C. (2019). Learning with good feature representations in bandits and in RL with a generative model. arXiv preprint arXiv:1911.07676 .
[26] May, B. C., Korda, N., Lee, A. and Leslie, D. S. (2012). Optimistic bayesian sampling in contextualbandit problems. Journal of Machine Learning Research 13 2069-2106.
[27] Osband, I., Aslanides, J. and Cassirer, A. (2018). Randomized prior functions for deep reinforcement learning. In Advances in Neural Information Processing Systems.
[28] Osband, I., Blundell, C., Pritzel, A. and Van Roy, B. (2016). Deep exploration via bootstrapped DQN. In Advances in Neural Information Processing Systems.
[29] Osband, I., Russo, D. and Van Roy, B. (2013). (more) efficient reinforcement learning via posterior sampling. In Advances in Neural Information Processing Systems.
[30] Osband, I., Russo, D., Wen, Z. and Van Roy, B. (2017). Deep exploration via randomized value functions. Journal of Machine Learning Research .
[31] Osband, I. and Van Roy, B. (2015). Bootstrapped Thompson sampling and deep exploration. arXiv preprint arXiv:1507.00300.
[32] Osband, I. and Van Roy, B. (2016). On lower bounds for regret in reinforcement learning. arXiv preprint arXiv:1608.02732.
[33] Osband, I. and Van Roy, B. (2017). Why is posterior sampling better than optimism for reinforcement learning? In International Conference on Machine Learning.
[34] Osband, I., Van Roy, B. and Wen, Z. (2014). Generalization and exploration via randomized value functions. arXiv preprint arXiv:1402.0635.
[35] Raginsky, M., Rakhlin, A. and Telgarsky, M. (2017). Non-convex learning via stochastic gradient Langevin dynamics: a nonasymptotic analysis. arXiv preprint arXiv:1702.03849 .
[36] RUSSO, D. (2019). Worst-case regret bounds for exploration via randomized value functions. In Advances in Neural Information Processing Systems.
[37] Russo, D. J., Van Roy, B., Kazerouni, A., Osband, I. and Wen, Z. (2018). A tutorial on Thompson sampling. Foundations and Trends® in Machine Learning 11 1-96.
[38] Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M. et al. (2016). Mastering the game of Go with deep neural networks and tree search. Nature 529484.
[39] Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D. and Riedmiller, M. (2014). Deterministic policy gradient algorithms. In International Conference on Machine Learning.
[40] Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., BAKER, L., LAI, M., Bolton, A. ET AL. (2017). Mastering the game of Go without human knowledge. Nature 550354.
[41] Strehl, A. L., Li, L., Wiewiora, E., Langford, J. and Littman, M. L. (2006). PAC model-free reinforcement learning. In International Conference on Machine Learning.
[42] TaïGa, A. A., Fedus, W., Machado, M. C., Courville, A. and Bellemare, M. G. (2019). Benchmarking bonus-based exploration methods on the Arcade Learning Environment. arXiv preprint arXiv:1908.02388.
[43] Thompson, W. R. (1933). On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. Biometrika 25 285-294.
[44] Van Roy, B. and Dong, S. (2019). Comments on the Du-Kakade-Wang-Yang lower bounds. arXiv preprint arXiv:1911.07910.
[45] Wang, W. Y., Li, J. and He, X. (2018). Deep reinforcement learning for NLP. In Association for Computational Linguistics.
[46] Wang, Y., Wang, R., Du, S. S. and Krishnamurthy, A. (2019). Optimism in reinforcement learning with generalized linear function approximation. arXiv preprint arXiv:1912.04136 .
[47] Watkins, C. J. C. H. (1989). Learning from delayed rewards .
[48] Welling, M. and Teh, Y. W. (2011). Bayesian learning via stochastic gradient Langevin dynamics. In International Conference on Machine Learning.
[49] Wen, Z. and Van Roy, B. (2017). Efficient reinforcement learning in deterministic systems with value function generalization. Mathematics of Operations Research 42 762-782.
[50] Xiong, Z., Shen, R. and Du, S. S. (2021). Randomized exploration is near-optimal for tabular mdp. arXiv preprint arXiv:2102.09703.
[51] Yang, L. and Wang, M. (2019). Sample-optimal parametric Q-learning using linearly additive features. In International Conference on Machine Learning.
[52] Yang, L. F. and Wang, M. (2019). Reinforcement leaning in feature space: Matrix bandit, kernels, and regret bound. arXiv preprint arXiv:1905.10389 .
[53] Yang, Z., Jin, C., Wang, Z., Wang, M. and Jordan, M. (2020). Provably efficient reinforcement learning with kernel and neural function approximations. Advances in Neural Information Processing Systems 33.
[54] Zanette, A., Brandfonbrener, D., Pirotta, M. and Lazaric, A. (2019). Frequentist regret bounds for randomized least-squares value iteration. arXiv preprint arXiv:1911.00567.
[55] Zanette, A., Lazaric, A., Kochenderfer, M. and Brunskill, E. (2020). Learning near optimal policies with low inherent bellman error. In International Conference on Machine Learning. PMLR.
[56] Zhang, Y., Liang, P. and Charikar, M. (2017). A hitting time analysis of stochastic gradient langevin dynamics. In Conference on Learning Theory. PMLR.
[57] Zhou, D., Gu, Q. and Szepesvari, C. (2021). Nearly minimax optimal reinforcement learning for linear mixture markov decision processes. In Conference on Learning Theory. PMLR.
[58] Zhou, H., Chen, J., Varshney, L. R. and Jagmohan, A. (2020). Nonstationary reinforcement learning with linear function approximation. arXiv preprint arXiv:2010.04244 .


[^0]:    *Northwestern University; boyiliu2018@u.northwestern.edu
    ${ }^{\dagger}$ Northwestern University; qicai2022@u.northwestern.edu
    ${ }^{\ddagger}$ Princeton University; zy6@princeton.edu
    ${ }^{\text {§ }}$ Northwestern University; zhaoranwang @ gmail.com

