SIMULTANEOUSLY LEARNING STOCHASTIC AND ADVERSARIAL MARKOV DECISION PROCESS WITH LINEAR FUNCTION APPROXIMATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Reinforcement learning (RL) has been commonly used in practice. To deal with the numerous states and actions in real applications, the function approximation method has been widely employed to improve the learning efficiency, among which the linear function approximation has attracted great interest both theoretically and empirically. Previous works on the linear Markov Decision Process (MDP) mainly study two settings, the stochastic setting where the reward is generated in a stochastic way and the adversarial setting where the reward can be chosen arbitrarily by an adversary. All these works treat these two environments separately. However, the learning agents often have no idea of how rewards are generated and a wrong reward type can severely disrupt the performance of those specially designed algorithms. So a natural question is whether an algorithm can be derived that can efficiently learn in both environments but without knowing the reward type. In this paper, we first consider such best-of-both-worlds problem for linear MDP with the known transition. We propose an algorithm and prove it can simultaneously achieve $O(\text{poly}\log K)$ regret in the stochastic setting and $O(\sqrt{K})$ regret in the adversarial setting where K is the horizon. To the best of our knowledge, it is the first such result for linear MDP.

1 INTRODUCTION

Reinforcement learning (RL) studies the problem where a learning agent interacts with the environment over time and aims to maximize its cumulative rewards in a given horizon. It has a wide range of real applications including robotics (Kober et al., 2013), games (Mnih et al., 2013; Silver et al., 2016), etc. The environment dynamics are usually modeled by the Markov Decision Process (MDP) with a fixed transition function. We consider the general episodic MDP setting where the interactions last for several episodes and the length of each episode is fixed (Jin et al., 2018; 2020b; Luo et al., 2021; Yang et al., 2021). In each episode, the agent first observes its current state and would decide which action to take. After making the decision, it receives an instant reward and the environment will then transfer to the next state. The cumulative reward in an episode is called the value and the objective of the agent is equivalent to minimizing the regret defined as the cumulative difference between the optimal value and its received values over episodes.

Many previous works focus on the tabular MDP setting where the state and action space are finite and the values can be represented by a table (Azar et al., 2017; Jin et al., 2018; Chen et al., 2021; Luo et al., 2021). Most of them study the stochastic setting with the stationary reward in which the reward of a state-action pair is generated from a fixed distribution (Azar et al., 2017; Jin et al., 2018; Simchowitz & Jamieson, 2019; Yang et al., 2021). Since the reward may change over time in applications, some works consider the adversarial MDP where the reward can be arbitrarily generated among different episodes (Yu et al., 2009; Rosenberg & Mansour, 2019; Jin et al., 2020a; Chen et al., 2021; Luo et al., 2021). All of these works pay efforts to learn the value function table to find the optimal policy and the computation complexity highly depends on the state and action space size.

However, in real applications such as the Go game, there are numerous states and the value function table is huge, which brings a great challenge to the computation complexity for traditional algorithms in the tabular case. To cope with the dimensionality curse, a rich line of works employ the function

approximation methods, such as the linear function and deep neural networks, to approximate the value functions or the policies to improve learning efficiency. These methods also achieve great success in practical scenarios such as the Atari and Go games (Mnih et al., 2013; Silver et al., 2016). Despite their great empirical performances, it also brings a series of challenges in deriving theoretical analysis. To build a better theoretical understanding of these approximation methods, lots of works start from deriving regret guarantees for linear function classes.

The linear MDP is a popular model which assumes both the transition and reward at a state-action pair are linear in the corresponding d-dimensional feature (Jin et al., 2020b; He et al., 2021; Hu et al., 2022). There are also mainly two types of the reward. For the stochastic setting, Jin et al. (2020b) provides the first efficient algorithm named Least-Square Value Iteration UCB (LSVI-UCB) and show that the its regret over K episodes can be upper bounded by $O(\sqrt{K})$. To seek for a tighter result with respect to the specific problem structure, He et al. (2021) provide a new analysis for LSVI-UCB and show it achieves an $O(\text{poly} \log K)$ instance-dependent regret upper bound. The adversarial setting is much harder than the stochastic one since the reward can change arbitrarily but the agent can only observe the rewards on the experienced trajectory. For this more challenging case, a regret upper bound of order $O(\sqrt{K})$ is only obtained in the case with known transition by Neu & Olkhovskaya (2021). All these works separately treat two environment types.

However, the learning agent usually has no idea of how the reward is generated. And once the reward type is wrong, the specially designed algorithm for a separate setting may suffer great loss. Thus deriving an algorithm that can adapt to different environment types becomes a natural solution for this problem. This direction has attracted great research interest in simpler bandit (Bubeck & Slivkins, 2012; Zimmert et al., 2019; Lee et al., 2021; Kong et al., 2022) and tabular MDP settings (Jin & Luo, 2020; Jin et al., 2021b) but still remains open in linear MDP.

In this paper, we try to answer the question of deriving best-of-both-worlds (BoBW) guarantees for linear MDP. Due to the challenge of learning in the adversarial setting, we also consider the known transition case. We propose an algorithm that continuously detects the real environment type and adjusts its strategy. It has been shown that our algorithm can simultaneously achieve $O(\text{poly} \log K)$ regret in the stochastic setting and $O(\sqrt{K})$ regret in the adversarial setting. To the best of our knowledge, these are the first BoBW results for linear MDP. It is also worth noting that our BoBW algorithm relies on an algorithm that can achieve a high-probability guarantee for the adversarial setting, which previous works fail to provide. And we propose the first analysis for a high-probability regret bound in the adversarial linear MDP.

2 Related Work

Linear MDP. Recently, deriving theoretically guaranteed algorithms for RL with linear function approximation has attracted great interests. The linear MDP model is one of the most popular one. Jin et al. (2020b) develop the first efficient algorithm LSVI-UCB both in sample and computation complexity for this setting. They show that the algorithm achieves $O(\sqrt{d^3H^3K})$ regret where d is the feature dimension and H is the length of each episode. This result is recently improved to the optimal order $O(dH\sqrt{K})$ by Hu et al. (2022) with a tighter concentration analysis. Apart from UCB, the TS-type algorithm has also been proposed for this setting (Zanette et al., 2020a). All these results do not consider the specific problem structure. In the stochastic setting, deriving an instance-dependent regret is more attractive to show the tighter performances of algorithms in a specific problem. This type of regret has been widely studied under the tabular MDP setting (Simchowitz & Jamieson, 2019; Yang et al., 2021). He et al. (2021) is the first to provide this type of regret in linear MDP. Using a different proof framework, they show that the LSVI-UCB algorithm can achieve $O(d^3H^5 \log K/\Delta)$ where Δ is the minimum value gap in the episodic MDP.

All these works consider the stochastic setting with stationary rewards. Neu & Olkhovskaya (2021) first attempts to analyze the more challenging adversarial environment. They consider a simplier setting with known transition and provide an $O(\sqrt{dHK})$ regret upper bound. For unknown transition case, Luo et al. (2021) provide an $O(d^{2/3}H^2K^{2/3})$ upper bound with the help of a simulator and $O(d^2H^4K^{14/15})$ guarantee for the general case. Above all, even in the separate adversarial setting, $O(\sqrt{K})$ regret is only derived for known transition case. We also study the known transition

setting and try to provide $O(\sqrt{K})$ regret in the adversarial setting while simultaneously achieving $O(\text{poly}\log K)$ regret if the environment is truly stochastic.

Best-of-both-worlds. The question of reaching best-of-both-worlds results is first proposed by Bubeck & Slivkins (2012) for bandit setting, a special case of episodic MDP with $H = 1$. Their proposed algorithm assumes the setting is stochastic and continuously detects whether the assumption is satisfied. Such a detection-based method is shown to achieve $O(\text{poly}\log K)$ regret in the stochastic setting and $O(\sqrt{K})$ in the adversarial setting, which is later improved by Auer & Chiang (2016). Similar detection-based techniques have also been adopted in more general linear bandit (Lee et al., 2021) and graph feedback (Kong et al., 2022) settings to achieve BoBW guarantees. Another line of works consider using Follow-the-Regularized-Leader (FTRL) to adapt to different environment types (Zimmert & Seldin, 2019; 2021). This type of algorithm is shown to be tighter than Bubeck & Slivkins (2012); Auer & Chiang (2016) in the bandit setting and also attracts lots of interest in more complex problems such as combinatorial bandits (Zimmert et al., 2019; Chen et al., 2022).

The first BoBW result in the MDP setting is provided by Jin & Luo (2020) in the tabular case. Due to the challenge of the problem, they first study the known transition setting. Their approach to achieving BoBW is the FTRL algorithm with a newly designed regularizer, which result is later improved by Jin et al. (2021b) and also generalized to the unknown transition case. To the best of our knowledge, we are the first to consider the BoBW problem in linear MDP. We also start from the known transition setting and our algorithm is based on detection.

RL with general function approximation The linear mixture MDP is another popular RL model with linear function approximation. It assumes the transition function can be approximated by a weighted average over several transition kernels. In the stochastic setting, both instance-independent (Ayoub et al., 2020; Zhou et al., 2021) and dependent regret bound (He et al., 2021) have been derived. And in the adversarial setting, only full information case has been studied, where the agent has access to the rewards of all state-action pairs (Cai et al., 2020; He et al., 2022). Apart from linear function approximation, there is also a rich line of works considering general function classes, such as the setting with low Bellman rank (Jiang et al., 2017; Zanette et al., 2020b), low Eluder dimension (Wang et al., 2020; Kong et al., 2021) and low Bellman Eluder dimension (Jin et al., 2021a).

3 Setting

We consider the episodic Markov Decision Process (MDP) where the agent interacts with the environment for K episodes. The episodic MDP can be represented by $\mathcal{M}(\mathcal{S}, \mathcal{A}, H, \{r_k\}_{k=1}^K, P)$, where $\mathcal{S}$ is the state space, $\mathcal{A}$ is the action space, H is the length of each episode, $r_k = \{r_{k,h}\}_{h=1}^H$ is the reward function and $P = \{P_h\}_{h=1}^H$ is the transition probability function. Specifically, at each episode k and step $h \in [H]$, $r_{k,h}(s,a) \in [0,1]$ and $P_h(\cdot \mid s,a) \in [0,1]^{|\mathcal{S}|}$ are the reward and the transition probability at state $s \in \mathcal{S}$ by taking action $a \in \mathcal{A}$, respectively.

We focus on stationary policies. Denote $\pi = \{\pi_h\}_{h=1}^H$ as a policy mapping from the state space to an action distribution where $\pi_h : \mathcal{S} \to \Delta_{\mathcal{A}}$. For each episode $k \in [K]$, the agent would start from the first state $s_{k,1} := s_1$[1] and determine the policy π_k. Then at each step $h \in [H]$ of episode k, it first observes the current state $s_{k,h}$ and then perform the action $a_{k,h} \sim \pi_{k,h}(\cdot \mid s_{k,h})$. The agent would receive a random reward $y_{k,h} := r_{k,h}(s_{k,h}, a_{k,h}) + \epsilon_{k,h}$, where $\epsilon_{k,h}$ is an independent zero-mean noise. The environment then transfers to the next state $s_{k,h+1}$ based on the transition function $P(\cdot \mid s_{k,h}, a_{k,h})$. The episode ends when the last state $s_{k,H+1}$ is reached.

We focus on the linear MDP setting where both the transition and the reward functions are linear in a given feature mapping (Jin et al., 2020b; He et al., 2021; Hu et al., 2022). And we first consider the known transition case. The formal definition is as follows.

Assumption 1. *(Linear MDP with known transition) $\mathcal{M}(\mathcal{S}, \mathcal{A}, H, \{r_k\}_{k=1}^K, P)$ is a linear MDP with a known feature mapping $\phi : \mathcal{S} \times \mathcal{A} \to \mathbb{R}^d$ such that for each step $h \in [H]$, there exists a known*

[1]The deterministic starting state is only for expositional convenience. Our algorithm and analysis can directly handle random starting states with a distribution.

vector μ_h and unknown vector θ_h where for each $(s,a) \in \mathcal{S} \times \mathcal{A}$ and state $s' \in \mathcal{S}$,

$$P_h(s' \mid s,a) = \left\langle \phi(s,a), \mu_h(s') \right\rangle \ , r_{k,h}(s,a) = \left\langle \phi(s,a), \theta_{k,h} \right\rangle \ .$$

In the stochastic setting, the reward function $\{r_k\}_{k=1}^K$, or namely the reward parameter $\{\theta_k\}_{k=1}^K$, is fixed over different episodes $k \in [K]$. And in the adversarial setting, the reward parameter $\{\theta_k\}_{k=1}^K$ can be chosen arbitrarily by an adversary (which may be possibly dependent on previous policies).

We evaluate the performance of a policy π by its value functions. Specifically, for any episode k and step h, denote the Q-value function $Q_{k,h}^{\pi}(s,a)$ as the expected reward that will be obtained by the agent starting from $(s_{k,h}, a_{k,h}) = (s,a)$ with policy π, which is formally defined as

$$Q_{k,h}^{\pi}(s,a) = \mathbb{E}\left[\sum_{h'=h+1}^{H} y_{k,h'} \mid \pi, s_{k,h} = s, a_{k,h} = a \right] .$$

Similarly, the value function $V_{k,h}^{\pi}(s)$ of any state s is defined as

$$V_{k,h}^{\pi}(s) = \mathbb{E}\left[\sum_{h'=h+1}^{H} y_{k,h'} \mid \pi, s_{k,h} = s \right] .$$

In the following paper, we abuse a bit notation by using $V_k^{\pi} := V_{k,1}^{\pi}(s_1)$ to represent the value of policy π at the starting state s_1 and episode k. Define $\phi_{\pi,h} = \mathbb{E}\left[\phi(s_h, a_h) \mid \pi\right]$ as the expected feature vector that the policy π visits at step h. It is worth noting that this recovers the state-action visitation probability in the tabular setting. And according to the definition, the value function of policy π at episode k can be represented as $V_k^{\pi} = \sum_{h=1}^{H} \left\langle \phi_{\pi,h}, \theta_{k,h} \right\rangle$.

In this paper, we consider optimizing both the stochastic and adversarial environments within a finite policy set Π. Given a policy set Π, the learning agent would determine the policy $\pi_k \in \Pi$ in each episode $k \in [K]$. Let $\pi^* \in \arg\max_{\pi} \sum_{k=1}^{K} V_k^{\pi}$ be one optimal policy in Π that maximizes the cumulative value functions over K episodes, which is assumed to be unique in the stochastic setting similar to previous works in tabular MDP (Jin & Luo, 2020; Jin et al., 2021b) and bandit setting (Lee et al., 2021; Zimmert & Seldin, 2019; 2021). Denote the cumulative regret compared with $\pi^* \in \Pi$ over K episodes as

$$Reg(K;\Pi) = \sum_{k=1}^{K} \left(V_k^{\pi^*} - V_k^{\pi_k} \right) . \tag{1}$$

4 ALGORITHM

In this section, we propose a detection-based algorithm to optimize both stochastic and adversarial environments for linear MDP with a given policy set Π. Our algorithm is mainly inspired by the detection technique of Lee et al. (2021) for BoBW in linear bandits. At a high level, the algorithm first assumes the environment is truly adversarial and continuously detect whether it could be a stochastic one. Its design relies on a new linear MDP algorithm that can return well-concentrated estimators for values of policies and also achieve sub-linear regret in the adversarial setting with high probability.

Previous works on adversarial linear MDP fail to provide a high-probability guarantee and thus no existing algorithms satisfy this property. In Appendix D, we propose a variant of Geometric Hedge (Algorithm 4), which is initially designed for the simple bandit case (Bartlett et al., 2008), and provide a new analysis for it in the linear MDP setting. We show that this algorithm satisfies the following properties and can be used to derive the BoBW results. It is also worth noting that this algorithm is the first to achieve a high-probability regret guarantee for adversarial linear MDP.

Theorem 1. *Given a policy set Π, the regret of our proposed Algorithm 4 in the adversarial setting can be upper bounded by*

$$\mathrm{Reg}(K;\Pi) \leq \mathcal{O}\left(\sqrt{dH^3 K \log\left(|\Pi|/\delta\right)} \right) \tag{2}$$

with probability at least $1-\delta$.

Further, at each episode k*, Algorithm 4 returns a value estimator* $\hat{V}_k^\pi$ *for each policy* $\pi \in \Pi$*. Choosing constant* $L_0 = 4dH \log\left(|\Pi|/\delta\right)$*,* $C_1 \geq 2^{15} dH^3 \log\left(K|\Pi|/\delta\right)$ *and* $C_2 \geq 20$*, it holds that for any* $k_0 \geq L_0$ *and policy* $\pi \in \Pi$*,*

$$\sum_{k=1}^{k_0} \left(V_k^\pi - V_k^{\pi_k}\right) \leq \sqrt{C_1 k_0} - C_2 \left| \sum_{k=1}^{k_0} \left(\hat{V}_k^\pi - V_k^\pi\right) \right|$$

with probability larger than $1-\delta$.

Our main BoBW algorithm is a phased algorithm and is presented in Algorithm 1. It takes Algorithm 4 satisfying Theorem 1 with parameter L_0 and C_1 as input. The first epoch is of length L_0 and the length of the following epochs would grow exponentially as Line 4. During each epoch, it executes Algorithm 2, which we refer to as the BoBW main body (Line 3).

Algorithm 1 BoBW for linear MDP

1: Input: Algorithm 4 with parameter C_1 and L_0; Set $L := L_0$
2: **while** true **do**
3: Execute Algorithm 2 (BoBW main body) with parameter L and receive output k_0
4: Set $L = 2k_0$
5: **end while**

Algorithm 2 (BoBW main body) takes the Algorithm 4 with parameter C_1 and L as input. Here L can just be regarded as the minimum number of episodes that Algorithm 4 needs to run to collect enough observations. The algorithm first assumes the environment is adversarial and executes Algorithm 4 in at least L episodes (which we refer to as the first phase). As shown in Theorem 1, Algorithm 4 guarantees that when running for more than L episodes, the regret compared with any policy π and the distance between its estimated V value and the real V value would be no larger. Based on these concentration properties, if a policy $\hat{\pi}$ shows consistent better performance than all of the other policies (Line 5), we have the reason to believe that the environment is truly stochastic. Being aware of this, as shown in Line 6, Algorithm 2 would transfer to the stochastic phase (Line 9-18, which we refer to as the second phase) with the estimated V values returned by Algorithm 4.

Since the estimated values by Algorithm 4 can well approximate the real values of policies, the exploration in the stochastic setting can be conducted by the estimated value gaps to obtain a problem-dependent regret bound. The objective is to identify the optimal policy and maximize the collected rewards, which can be implemented by an optimization problem (Algorithm 3). Taking the estimated value gap $\hat{\Delta}$ as input, Algorithm 3 would return a probability distribution p^* over the policy set Π. Intuitively, p^* maximizes the expected values of policies while ensuring the uncertainty of all policies to be smaller than its sub-optimality gap. In Appendix A, we show that when $\hat{\Delta}$ is estimated accurately, selecting policies based on p^* can reach a problem-dependent regret upper bound.

Back to the main body of Algorithm 2, after computing the distribution p_k based on the current estimated $\hat{\Delta}$ (Line 10), it also mixes this policy distribution with a one-hot vector $e_{\hat{\pi}}$ to ensure that the estimated optimal policy $\hat{\pi}$ can be observed for enough times and the variance of its following estimators can thus be low (Line 11). The algorithm then samples a policy π_k according to this mixed distribution and executes it in this episode. Then based on the received rewards $y_{k,h}$ at each step h and the total reward $Y_k = \sum_{h=1}^H y_{k,h}$, the value estimations of policies can be further updated. Due to the technical reason, here for the estimated optimal policy $\hat{\pi}$ and other policies π, we use different estimators. Specifically, we use the importance weighted estimator to approximate the value of $\hat{\pi}$. The reason for using mixed policy distribution $\tilde{p}$ is just to ensure the low variance of this estimator. And for other policies, we use the standard linear square estimators (Line 13). Based on these newly estimated values of policies, the algorithm can then update the estimation for their sub-optimality gaps as Line 14. To get a tighter analysis, when computing the value gap of π, we use the traditional estimator of it by Algorithm 4 in the first k_0 episodes and the Catoni estimator for the recent $k-k_0$ episodes. Formally speaking, the Catoni estimator $\textbf{Catoni}_\alpha\left(\{X_1, X_2, \cdots X_n\}\right)$ is defined as the unique root of $f(z) = \sum_{i=1}^n \Phi\left(\alpha\left(X_i - z\right)\right)$, where $\Phi(y) = \log\left(1 + y + y^2/2\right)$ if

Algorithm 2 BOBW main body

Input: A new instance of Algorithm 4 with parameter C_1, parameter L
Define: $f_K = \log K$
for each episode $k = 1, 2, \cdots$ **do**
  Execute and update Algorithm 4, receive value estimators $\hat{V}_k^\pi$ for each $\pi \in \Pi$
  if $k \geq L$ and there exists a policy $\hat{\pi} \in \Pi$ such that

$$\sum_{s=1}^{k} Y_s \leq \sum_{s=1}^{k} \hat{V}_s^{\hat{\pi}} + 5\sqrt{f_K C_1 k}\,, \tag{3}$$

$$\sum_{s=1}^{k} Y_s \geq \sum_{s=1}^{k} \hat{V}_s^{\pi} + 25\sqrt{f_K C_1 k}\,, \quad \forall \pi \neq \hat{\pi}\,. \tag{4}$$

  then
    $k_0 = k$, $\hat{\Delta}_\pi = \frac{1}{k_0}\left(\sum_{s=1}^{k} \hat{V}_s^{\hat{\pi}} - \hat{V}_s^{\pi}\right)$, $\hat{\Delta} = \left\{\left(\hat{\Delta}_\pi\right)_{\pi\in\Pi}\right\}$; break
  end if
end for
for episode $k = k_0 + 1, k_0 + 2, \cdots$ **do**
  Compute $p_k = \mathbf{OP}(k, \hat{\Delta})$
  Compute $\tilde{p}_k(\pi) = \frac{1}{2}e_{\hat{\pi}} + \frac{1}{2}p_k(\pi)$, where $e_{\hat{\pi}}$ is a one-hot vector with 1 only at policy $\hat{\pi}$
  Sample $\pi_k \sim \tilde{p}_k$ and execute π_k
  Receive rewards $Y_k = \sum_{h=1}^{H} y_{k,h}$ and calculate $\hat{V}_k^\pi$ for each $\pi \in \Pi$ as follows

$$\forall \pi \neq \hat{\pi} \;: \hat{V}_k^\pi = \sum_{h=1}^{H} \phi_{\pi,h}^\top \Sigma_{k,h}^{-1} \phi_{\pi_k,h} y_{k,h}\,, \text{ where } \Sigma_{k,h} = \sum_{\pi} \tilde{p}_k(\pi)\phi_{\pi,h}\phi_{\pi,h}^\top\,; \tag{5}$$

$$\hat{V}_k^{\hat{\pi}} = \frac{Y_k}{\tilde{p}_k(\hat{\pi})}\mathbb{1}\{\pi_k = \hat{\pi}\}\,. \tag{6}$$

  For each $\pi \neq \hat{\pi}$, compute $\hat{\Delta}_\pi$ as

$$\hat{\Delta}_k^\pi = \frac{1}{k}\left(\sum_{s=1}^{k_0} \hat{V}_s^\pi + (k - k_0)\mathrm{Rob}_{k,\pi} - \sum_{s=1}^{k} \hat{V}_s^{\hat{\pi}}\right)\,, \text{ with } \mathrm{Rob}_{k,\pi} = \mathbf{Catoni}_{\alpha_k^\pi}\left(\{\hat{V}_s^\pi\}_{s=k_0+1}^{k}\right) \tag{7}$$

  if

$$\exists \pi \neq \hat{\pi}\,, \hat{\Delta}_k^\pi \notin \left[0.39\hat{\Delta}_\pi, 1.81\hat{\Delta}_\pi\right] \text{ or} \tag{8}$$

$$\sum_{s=k_0+1}^{k} \left(\hat{V}_s^{\hat{\pi}} - Y_s\right) \geq 20\sqrt{f_K C_1 k_0} \tag{9}$$

  then
    Return k_0
  end if
end for

$y \geq 0$ and $\Phi(y) = -\log\left(1 - y + y^2/2\right)$ otherwise. The hyper-parameter α_k^π is set as

$$\alpha_k^\pi = \sqrt{\frac{4\log\left(k|\Pi|/\delta\right)}{\sum_{\kappa=k_0+1}^{k}\left(2\kappa\hat{\Delta}_\pi^2/\beta_\kappa + 9dH\right)}}\,, \tag{12}$$

which mainly follows Lee et al. (2021) but with the careful consideration in MDP setting.

Algorithm 3 Optimization problem (**OP**)$\left(k, \hat{\Delta}\right)$

1: Define $\beta_k = 2^{15} H \log\left(|\Pi| k/\delta\right)$
2: Return the minimizer p^* of the following constrained optimization problem

$$\min_{p} \sum_{\pi \in \Pi} p_\pi \hat{\Delta}_\pi \tag{10}$$

$$\text{s.t.} \sum_{h=1}^{H} \left\|\phi_{\pi,h}\right\|^2_{\Sigma_h^{-1}(p)} \leq \frac{k\hat{\Delta}_\pi}{\beta_k} + 4dH \text{ and } \sum_{\pi \in \Pi} p_\pi = 1\,, \tag{11}$$

where $\Sigma_h(p) = \sum_\pi p_\pi \phi_{\pi,h} \phi_{\pi,h}^\top$.

Is is worth noting that an adversarial setting may be disguised as stochastic scenarios and foolish the algorithm to enter in the stochastic phase. Thus the agent still needs to be vigilant about the possible change of the environment. The detection conditions (Line 15) are set for this objective. To be specific, if the estimated sub-optimality gap of a policy changes obviously compared with the original estimation by Algorithm 4 in the adversarial phase or the regret compared with $\hat{\pi}$ is large, the algorithm can determine that the environment may not be stochastic and would terminiate the current epoch and enter in the next epoch with parameter k_0 (Line 16).

5 Theoretical Analysis

In this section, we provide the theoretical guarantee and the analysis of Algorithm 1 in both stochastic and adversarial settings

The first is about the stochastic setting. Since the value function remains the same for different episodes, we simplify the notation and use V^π to represent the real value of policy $\pi \in \Pi$. Before presenting the main results, we first introduce the sub-optimality gaps that will be used.

Definition 1. *For each policy $\pi \in \Pi$, define $\Delta_\pi = V^{\pi^*} - V^\pi$ as the sub-optimality gap of π compared with the optimal policy $\pi^* \in \arg\max_{\pi \in \Pi} V^\pi$. Further let*

$$\Delta_{\min} = \min_{\pi : \Delta_\pi > 0} \Delta_\pi \tag{13}$$

be the minimum non-negative value gap.

Theorem 2 provides a regret upper bound for Algorithm 1 in the stochastic setting.

Theorem 2. *(Regret bound in the stochastic setting) With probability at least $1-\delta$, Algorithm 1 guarantees that*

$$\mathrm{Reg}(K;\Pi) \leq \mathcal{O}\left(\frac{dH^2 \log(K) \log\left(|\Pi| K/\delta\right)}{\Delta_{\min}}\right)\,. \tag{14}$$

And if the environment is adversarial, the regret of Algorithm 1 can be upper bounded as the following Theorem 3.

Theorem 3. *(Regret bound in the adversarial setting) With probability at least $1-\delta$, Algorithm 1 guarantees that*

$$\mathrm{Reg}(K;\Pi) \leq \mathcal{O}\left(\sqrt{dH^3 K \log(K) \log\left(|\Pi| K/\delta\right)}\right)\,. \tag{15}$$

Due to the space limit, the full proof of these two theorems are deferred to Appendix B and C. We will provide a proof sketch for them in later sections.

Discussions To the best of our knowledge, Algorithm 1 is the first that can simultaneously achieve $O(\mathrm{poly}\log K)$ regret in the stochastic setting and $O(\sqrt{K})$ regret in the adversarial setting for linear MDP problem. Our approach relies on a new observation that the value of each policy can be written as the inner product between the expected state-action visitation feature ϕ_π and the unknown reward

feature θ. Thus the problem can be regarded as an online linear optimization problem and we can optimize the value functions based on policies. It is also worth noting that previous works on the separate adversarial setting only provide the upper bound for the expected regret (Neu & Olkhovskaya, 2021), and we are the first to provide a high-probability guarantee.

If the real global optimal policy is just contained in the policy set Π, our regret is equivalent to the traditional definition in previous works (Jin et al., 2020b; He et al., 2021; Neu & Olkhovskaya, 2021). And if the policy space is too large (or infinite) in some applications, one can construct a policy covering such that for the global optimal value $\sup_\pi \sum_{k=1}^K V_k^\pi$, there exists a policy representative π^* in the covering which satisfies $\sup_\pi \sum_{k=1}^K V_k^\pi - \sum_{k=1}^K V_k^{\pi^*} \leq K\epsilon$. In this case, the regret compared with the global optimal policy can be upper bounded by

$$Reg(K) := \sup_\pi \sum_{k=1}^K \left(V_k^\pi - V_k^{\pi_k}\right) \leq \sum_{k=1}^K \left(V_k^{\pi^*} - V_k^{\pi_k}\right) + K\epsilon = Reg(K;\Pi) + K\epsilon\,.$$

By choosing $\epsilon = O(1/K)$, we can get the same upper bound for $Reg(K)$. Since the regret upper bound in Theorem 2 and 3 only depends on $\log|\Pi|$, using the covering does not influence the main order of the regret.

5.1 Regret Analysis in the Stochastic Setting

In the stochastic setting, we consider the regret in the two phases of each epoch separately. We first show that, with high probability, the first phase (which we call as the adversarial phase, Line 3-8 in Algorithm 2) will terminate after k_0 episodes where $k_0 \in \left[64 f_K C_1/\Delta_{\min}^2, 900 f_K C_1/\Delta_{\min}^2\right]$, and we can identify the optimal policy $\pi^* \in \Pi$. The formal claims are summarized in the following Lemma 1.

Lemma 1. *In the stochastic setting, in each epoch, the following* 4 *claims hold.*

1. *With probability at least* $1-4\delta$, $k_0 \leq \max\left\{\frac{900 f_K C_1}{\Delta_{\min}^2}, L\right\}$.
2. *With probability at least* $1-\delta$, $\hat{\pi} = \pi^*$.
3. *With probability at least* $1-2\delta$, $k_0 \geq \frac{64 f_K C_1}{\Delta_{\min}^2}$.
4. *With probability at least* $1-3\delta$, $\hat{\Delta}_\pi \in [0.7\Delta_\pi, 1.3\Delta_\pi]\,, \forall \pi \neq \pi^*$.

The detailed proof of Lemma 1 is deferred to Appendix B. We next will give a proof sketch of Theorem 2 based on the results of Lemma 1. According to claim 1 in Lemma 1, we know that $k_0 = \mathcal{O}\left(f_K C_1/\Delta_{\min}^2\right)$, so we can bound the regret in the first phase using the guarantees of Algorithm 4 in Theorem 1 by $\sqrt{C_1 k_0} = \mathcal{O}\left(C_1\sqrt{\log(K)}/\Delta_{\min}\right)$.

And after k_0 episodes, the algorithm would transfer to the second phase, which we call as the stochastic phase (Line 9-18 in Algorithm 2). If the environment is truly stochastic, the values of all policies would remain stationary and the detection condition (Line 15 in Algorithm 2) would never be satisfied. Thus the stochastic phase will not end (proved in Lemma 8 in Appendix B). As for the regret suffered in this phase, we can analyze it using the properties of **OP**. According to claim 4 in Lemma 1, the estimated sub-optimality gap is close to the real sub-optimality gap as $\hat{\Delta}_\pi \in \left[\Delta_\pi/\sqrt{3}, \sqrt{3}\Delta_\pi\right]$. Thus performing policies based on the solution of **OP** with $\hat{\Delta}$ can reach the real instance optimality.

Specifically, we can first bridge the gap between the expected regret and the regret that occurred using Freedman inequality (Lemma 12).

$$\sum_{s=k_0+1}^{k} \Delta_{\pi_s} \leq 2\sum_{s=k_0+1}^{k}\sum_{\pi} \tilde{p}_s(\pi)\Delta_\pi + 2H\log\left(\frac{1}{\delta}\right)\,.$$

Recall that $\tilde{p}$ is computed based on **OP** under $\hat{\Delta}$, which is close to the real sub-optimality gap Δ. The regret occurred in phase 2 in episodes k larger than a problem dependent constant $M =$

$\mathcal{O}\left(dH\beta_K/\Delta_{\min}^2\right)$, which is the dominating part in the expected regret, can be bounded using Lemma 7 with $r = 3$:

$$\sum_{s=M}^{k}\sum_{\pi}\tilde{p}_s(\pi)\Delta_\pi \le \sum_{s=M}^{k}\frac{72dH\beta_s}{\Delta_{\min}s} = \mathcal{O}\left(\frac{dH\beta_k\log(k)}{\Delta_{\min}}\right).$$

Above all, we can conclude that with probability at least $1-\delta$, the regret can be upper bounde as

$$\mathrm{Reg}\left(K;\Pi\right) = \mathcal{O}\left(\frac{dH\beta_K\log(K)}{\Delta_{\min}}\right) = \mathcal{O}\left(\frac{dH^2\log(K)\log\left(|\Pi|K/\delta\right)}{\Delta_{\min}}\right).$$

5.2 Regret Analysis in the Adversarial Setting

In the adversarial setting, the regret in the first phase can be guaranteed with the property of Algorithm 4 in Theorem 1. Here we will mainly analyze the second phase. We first show that in the second phase of each epoch, the returned policy $\hat{\pi}$ is actually the optimal policy in Π.

Lemma 2. *For any episode k in the second phase, the policy $\hat{\pi}$ has the most accumulated value in Π during episodes 1 to k. That is, $\hat{\pi} \in \arg\max_{\pi\in\Pi}\sum_{\kappa=1}^{k}V_\kappa^\pi$.*

Since $\hat{\pi}$ is the optimal policy in Π, the regret can be written as the sum of the deviation between the value of the selected policy π_s and $V_s^{\hat{\pi}}$

$$\begin{aligned}
&\sum_{s=k_0+1}^{k}\left(V_s^{\hat{\pi}} - V_s^{\pi_s}\right)\\
=&\sum_{s=k_0+1}^{k}\Bigg[\left(\hat{V}_s^{\hat{\pi}} - Y_s\right) + \left(Y_s\mathbb{1}\{\pi_s=\hat{\pi}\} + V_s^{\hat{\pi}}\mathbb{1}\{\pi_s\neq\hat{\pi}\} - \hat{V}_s^{\hat{\pi}}\right)\\
&+\left(\left(V_s^{\hat{\pi}}\mathbb{1}\{\pi_s=\hat{\pi}\} - Y_s\mathbb{1}\{\pi_s=\hat{\pi}\}\right) + (Y_s - V_s^{\pi_s})\right)\Bigg].
\end{aligned}$$

According to the detection condition (equation 9) in Algorithm 2, the first term can be upper bounded by $\mathcal{O}\left(\sqrt{f_K C_1 k_0}\right)$. As for the second and last term, Freedman inequality (Lemma 12) also provides an upper bound $\mathcal{O}\left(C_1 k_0\right)$ for them. Above all, the regret in a single epoch can be upper bounded by $\mathcal{O}\left(\sqrt{f_K C_1 k_0}\right)$. According to the choice of the minimal duration L in each epoch, the length k_0 of the first phase in an epoch is at most half of that in the next epoch. Thus the final regret can be bounded as

$$\mathrm{Reg}\left(K;\Pi\right) = \mathcal{O}\left(\sqrt{C_1 K\log(K)}\right) = \mathcal{O}\left(\sqrt{dH^2K\beta_K\log(K)}\right).$$

6 Conclusion

In this paper, we propose the first BoBW algorithm for linear MDP that can simultaneously achieve $O(\mathrm{poly}\log K)$ regret in the stochastic setting and $O(\sqrt{K})$ regret in the adversarial setting. Our approach relies on the new observation that the value function of a policy can be written as the sum of the inner products between the expected state-action visitation feature $\phi_{\pi,h}$ and the unknown reward parameter θ_h at different steps $h \in [H]$. And the problem can thus be regarded as an online linear optimization problem. Apart from these BoBW results, we also propose a new analysis that can reach a high-probability regret guarantee for adversarial linear MDP, which is also the first such result in the literature.

An important future direction is to remove the assumption of unique optimal policy in the stochastic setting. This assumption also appears in previous BoBW works for tabular MDP (Jin & Luo, 2020; Jin et al., 2021b) and linear bandits (Lee et al., 2021) which destroys the generality of the results but is challenging to be removed due to the hardness of the BoBW objective and the complex structure of MDP. Extending the current results to the unknown transition setting is also prospective. This is hoped to be solved by some new techniques since the current approach highly depends on the known state-action visitation feature. Deriving an FTRL-type algorithm for this objective is also an interesting future direction. It still remains open even in the simpler linear bandit setting without the transition between different states.

REFERENCES

Peter Auer and Chao-Kai Chiang. An algorithm with nearly optimal pseudo-regret for both stochastic and adversarial bandits. In *Conference on Learning Theory*, pp. 116–120. PMLR, 2016.

Alex Ayoub, Zeyu Jia, Csaba Szepesvari, Mengdi Wang, and Lin Yang. Model-based reinforcement learning with value-targeted regression. In *International Conference on Machine Learning*, pp. 463–474. PMLR, 2020.

Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *International Conference on Machine Learning*, pp. 263–272. PMLR, 2017.

Peter Bartlett, Varsha Dani, Thomas Hayes, Sham Kakade, Alexander Rakhlin, and Ambuj Tewari. High-probability regret bounds for bandit online linear optimization. In *Proceedings of the 21st Annual Conference on Learning Theory-COLT 2008*, pp. 335–342. Omnipress, 2008.

Sébastien Bubeck and Aleksandrs Slivkins. The best of both worlds: Stochastic and adversarial bandits. In *Conference on Learning Theory*, pp. 42–1. JMLR Workshop and Conference Proceedings, 2012.

Qi Cai, Zhuoran Yang, Chi Jin, and Zhaoran Wang. Provably efficient exploration in policy optimization. In *International Conference on Machine Learning*, pp. 1283–1294. PMLR, 2020.

Cheng Chen, Canzhe Zhao, and Shuai Li. Simultaneously learning stochastic and adversarial bandits under the position-based model. In *AAAI*, 2022.

Liyu Chen, Haipeng Luo, and Chen-Yu Wei. Minimax regret for stochastic shortest path with adversarial costs and known transition. In *Conference on Learning Theory*, pp. 1180–1215. PMLR, 2021.

David Freedman. On tail probabilities for martingales. 1975.

Jiafan He, Dongruo Zhou, and Quanquan Gu. Logarithmic regret for reinforcement learning with linear function approximation. In *International Conference on Machine Learning*, pp. 4171–4180. PMLR, 2021.

Jiafan He, Dongruo Zhou, and Quanquan Gu. Near-optimal policy optimization algorithms for learning adversarial linear mixture mdps. In *International Conference on Artificial Intelligence and Statistics*, pp. 4259–4280. PMLR, 2022.

Pihe Hu, Yu Chen, and Longbo Huang. Nearly minimax optimal reinforcement learning with linear function approximation. In *International Conference on Machine Learning*, pp. 8971–9019. PMLR, 2022.

Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, and Robert E Schapire. Contextual decision processes with low bellman rank are pac-learnable. In *International Conference on Machine Learning*, pp. 1704–1713. PMLR, 2017.

Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is q-learning provably efficient? *Advances in neural information processing systems*, 31, 2018.

Chi Jin, Tiancheng Jin, Haipeng Luo, Suvrit Sra, and Tiancheng Yu. Learning adversarial markov decision processes with bandit feedback and unknown transition. In *International Conference on Machine Learning*, pp. 4860–4869. PMLR, 2020a.

Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pp. 2137–2143. PMLR, 2020b.

Chi Jin, Qinghua Liu, and Sobhan Miryoosefi. Bellman eluder dimension: New rich classes of rl problems, and sample-efficient algorithms. *Advances in neural information processing systems*, 34:13406–13418, 2021a.

Tiancheng Jin and Haipeng Luo. Simultaneously learning stochastic and adversarial episodic mdps with known transition. *Advances in neural information processing systems*, 33:16557–16566, 2020.

Tiancheng Jin, Longbo Huang, and Haipeng Luo. The best of both worlds: stochastic and adversarial episodic mdps with unknown transition. *Advances in Neural Information Processing Systems*, 34: 20491–20502, 2021b.

Jens Kober, J Andrew Bagnell, and Jan Peters. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274, 2013.

Dingwen Kong, Ruslan Salakhutdinov, Ruosong Wang, and Lin F Yang. Online sub-sampling for reinforcement learning with general function approximation. *arXiv preprint arXiv:2106.07203*, 2021.

Fang Kong, Yichi Zhou, and Shuai Li. Simultaneously learning stochastic and adversarial bandits with general graph feedback. In *International Conference on Machine Learning*, pp. 11473–11482. PMLR, 2022.

Chung-Wei Lee, Haipeng Luo, Chen-Yu Wei, Mengxiao Zhang, and Xiaojin Zhang. Achieving near instance-optimality and minimax-optimality in stochastic and adversarial linear bandits simultaneously. In *International Conference on Machine Learning*, 2021.

Haipeng Luo, Chen-Yu Wei, and Chung-Wei Lee. Policy optimization in adversarial mdps: Improved exploration via dilated bonuses. *Advances in Neural Information Processing Systems*, 34:22931–22942, 2021.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.

Gergely Neu and Julia Olkhovskaya. Online learning in mdps with linear function approximation and bandit feedback. *Advances in Neural Information Processing Systems*, 34:10407–10417, 2021.

Aviv Rosenberg and Yishay Mansour. Online convex optimization in adversarial markov decision processes. In *International Conference on Machine Learning*, pp. 5478–5486. PMLR, 2019.

David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.

Max Simchowitz and Kevin G Jamieson. Non-asymptotic gap-dependent regret bounds for tabular mdps. *Advances in Neural Information Processing Systems*, 32, 2019.

Ruosong Wang, Russ R Salakhutdinov, and Lin Yang. Reinforcement learning with general value function approximation: Provably efficient approach via bounded eluder dimension. *Advances in Neural Information Processing Systems*, 33:6123–6135, 2020.

Chen-Yu Wei and Haipeng Luo. More adaptive algorithms for adversarial bandits. *CoRR*, abs/1801.03265, 2018. URL http://arxiv.org/abs/1801.03265.

Kunhe Yang, Lin Yang, and Simon Du. Q-learning with logarithmic regret. In *International Conference on Artificial Intelligence and Statistics*, pp. 1576–1584. PMLR, 2021.

Jia Yuan Yu, Shie Mannor, and Nahum Shimkin. Markov decision processes with arbitrary reward processes. *Mathematics of Operations Research*, 34(3):737–757, 2009.

Andrea Zanette, David Brandfonbrener, Emma Brunskill, Matteo Pirotta, and Alessandro Lazaric. Frequentist regret bounds for randomized least-squares value iteration. In *International Conference on Artificial Intelligence and Statistics*, pp. 1954–1964. PMLR, 2020a.

Andrea Zanette, Alessandro Lazaric, Mykel Kochenderfer, and Emma Brunskill. Learning near optimal policies with low inherent bellman error. In *International Conference on Machine Learning*, pp. 10978–10989. PMLR, 2020b.

Dongruo Zhou, Quanquan Gu, and Csaba Szepesvari. Nearly minimax optimal reinforcement learning for linear mixture markov decision processes. In *Conference on Learning Theory*, pp. 4532–4576. PMLR, 2021.

Julian Zimmert and Yevgeny Seldin. An optimal algorithm for stochastic and adversarial bandits. In *International Conference on Artificial Intelligence and Statistics*, pp. 467–475. PMLR, 2019.

Julian Zimmert and Yevgeny Seldin. Tsallis-inf: An optimal algorithm for stochastic and adversarial bandits. *Journal of Machine Learning Research*, 22:1–49, 2021.

Julian Zimmert, Haipeng Luo, and Chen-Yu Wei. Beating stochastic and adversarial semi-bandits optimally and simultaneously. In *International Conference on Machine Learning*, pp. 7683–7692. PMLR, 2019.