

Reward Biased Maximum Likelihood Estimation for Reinforcement Learning

Akshay Mete*

Texas A&M University, TX, USA.

AKSHAYMETE@TAMU.EDU

Rahul Singh*

Indian Institute of Science, Bangalore, India.

RAHULSINGH@IISC.AC.IN

Xi Liu

Texas A&M University, TX, USA.

XILIU.TAMU@GMAIL.COM

P. R. Kumar

Texas A&M University, TX, USA.

PRK@TAMU.EDU

Abstract

The Reward-Biased Maximum Likelihood Estimate (RBMLE) for adaptive control of Markov chains was proposed in (Kumar and Becker, 1982) to overcome the central obstacle of what is variously called the fundamental “closed-identifiability problem” of adaptive control (Borkar and Varaiya, 1979), the “dual control problem” by Feldbaum (Feldbaum, 1960a,b), or, contemporaneously, the “exploration vs. exploitation problem”. It exploited the key observation that since the maximum likelihood parameter estimator can asymptotically identify the closed-transition probabilities under a certainty equivalent approach (Borkar and Varaiya, 1979), the limiting parameter estimates must necessarily have an optimal reward that is less than the optimal reward attainable for the true but unknown system. Hence it proposed a counteracting reverse bias in favor of parameters with larger optimal rewards, providing a carefully structured solution to the fundamental problem alluded to above. It thereby proposed an optimistic approach of favoring parameters with larger optimal rewards, now known as “optimism in the face of uncertainty.” The RBMLE approach has been proved to be long-term average reward optimal in a variety of contexts including controlled Markov chains, linear quadratic Gaussian (LQG) systems, some nonlinear systems, and diffusions. However, modern attention is focused on the much finer notion of “regret,” or finite-time performance for *all* time, espoused by (Lai and Robbins, 1985). Recent analysis of RBMLE for multi-armed stochastic bandits (Liu et al., 2020) and linear contextual bandits (Hung et al., 2020) has shown that it not only has state-of-the-art regret, but it also exhibits empirical performance comparable to or better than the best current contenders, and leads to several new and strikingly simple index policies for these classical problems. Motivated by this, we examine the finite-time performance of RBMLE for reinforcement learning tasks that involve the general problem of optimal control of unknown Markov Decision Processes. We show that it has a regret of $O(\log T)$ over a time horizon of T steps, similar to state-of-the-art algorithms. Simulation studies show that RBMLE outperforms other algorithms such as UCRL2 (Auer et al., 2009) and Thompson Sampling (Ouyang et al., 2017; Gopalan and Mannor, 2015; Abbasi-Yadkori and Szepesvári, 2015).

Keywords: Reinforcement Learning; Markov Decision Process; Adaptive Control

* denotes equal contribution

1. Introduction

Consider a controlled Markov chain with finite state space X , finite action set U , and controlled transition probabilities $\mathbb{P}(x(t+1) = y \mid x(t) = x, u(t) = u) = p(x, y, u)$, where $x(t) \in X$ denotes the state at time t , and $u(t) \in U$ denotes the action taken at time t . A reward $r(x, u)$ is received when action u is taken in state x . Let $J^*(p)$ denote the maximal long-term average reward $\liminf_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T r(x(t), u(t))$ obtainable. We consider the case where the transition probabilities p are only known to belong a set Θ , but otherwise unknown. We address the adaptive control problem of minimizing the expected “regret”

$$TJ^*(p) - \mathbb{E} \sum_{t=1}^T r(x(t), u(t)) \quad (1)$$

as a function of T .

This broad problem has a long history. Let $J(\theta, \pi)$ denote the long-term average reward accrued by a stationary deterministic policy $\pi : X \rightarrow U$ when the transition probabilities are given by $\theta = \{\theta(x, y, u), x \in X, y \in X, u \in U\}$, let $J^*(\theta) := \max_{\pi} J(\theta, \pi)$ denote the optimal long-term average reward attainable under θ , and let $\pi^\theta \in \arg \max_{\pi} J(\theta, \pi)$ be an optimal policy for θ . In early work, (Mandl, 1974) studied the problem of using a “certainty equivalent” approach, where a maximum likelihood estimate (MLE)

$$\hat{\theta}(t) \in \arg \max_{\theta} \prod_{s=1}^{t-1} \theta(x(s), x(s+1), u(s)) \quad (2)$$

of the unknown transition probabilities is made at each time t , and an action $u(t) = \pi^{\hat{\theta}(t)}(x(t))$ is taken that is optimal in state $x(t)$ for the transition probabilities $\hat{\theta}(t)$. Mandl showed that if an “identifiability condition”,

$$\theta \neq \theta' \text{ with } \theta, \theta' \in \Theta \implies \theta(x, \cdot, u) \neq \theta'(x, \cdot, u) \quad \forall (x, u) \in X \times U \quad (3)$$

holds, then the maximum likelihood estimates $\hat{\theta}(t)$ converge to the true transition probabilities p as $t \rightarrow \infty$, and the corresponding long-term average reward obtained by the adaptive controller is the optimal reward $J^*(p)$. This identifiability condition is however restrictive, e.g., it is not satisfied for the two-armed bandit problem or any problem with a fundamental exploration vs. exploitation dilemma.

In general, in the absence of the identifiability condition, (Borkar and Varaiya, 1979) showed that one only obtains “closed-loop identification”: the maximum likelihood estimates converge to a θ^* for which

$$\theta^*(x, y, \pi^{\theta^*}(x)) = p(x, y, \pi^{\theta^*}(x)) \quad \forall (x, y). \quad (4)$$

However, the limiting policy π^{θ^*} is generally not an optimal long-term average policy for the true transition probabilities p . Indeed this is the central challenge of the exploration vs. exploitation problem: As the parameter estimates begin to converge exploration ceases, and one ends up only identifying the behavior of the system under the limited actions being applied to the system. One misses out on other potentially valuable policies.

This central difficulty was overcome in (Kumar and Becker, 1982). They first noted that (4) implies that $J(\theta^*, \pi^{\theta^*}) = J(p, \pi^{\theta^*})$. As a consequence of this, since π^{θ^*} is optimal for θ^* , i.e., $J(\theta^*, \pi^{\theta^*}) = J^*(\theta^*)$, but not for p , i.e., $J(p, \pi^{\theta^*}) \leq J^*(p)$, they made the critical observation that the optimal long-term average reward accruable for the limiting estimate must necessarily be lower than the optimal long-term average reward accruable for the true parameter:

$$J^*(\theta^*) \leq J^*(p). \quad (5)$$

Therefore, the maximum likelihood estimator is inherently *biased* in favor of θ 's with lower optimal rewards than p . Therefore to extricate oneself from this bind, one must necessarily tilt the balance toward exploring parameters with larger optimal rewards. Motivated by this, they proposed a certainty equivalent approach using a *Reward Biased MLE (RBMLE)* that attempts to counteract this with a bias in the reverse direction, favoring parameters θ with a larger optimal reward:

$$\hat{\theta}(t) \in \arg \max_{\theta} f(J^*(\theta))^{\alpha(t)} \prod_{s=1}^{t-1} \theta(x(s), x(s+1), u(s)), \quad (6)$$

where f is any strictly monotone increasing function. This biasing however has to be delicate in that $\alpha(t)$ has to be large enough so that it asymptotically does choose parameters with larger optimal reward than under p , but has to be small enough in that it does not lose the consistency property (4) of the MLE. They showed that the choice $\alpha(t) = o(t)$ with $\lim_{t \rightarrow +\infty} \alpha(t) = +\infty$ suffices in ensuring (4) for every Cesaro-limit point θ^* of the RBMLE estimator (6), but also satisfies

$$J^*(\theta^*) \geq J^*(p). \quad (7)$$

From this it follows that

$$J^*(p) \leq J^*(\theta^*) = J(\theta^*, \pi^{\theta^*}) = J(p, \pi^{\theta^*}) \leq J^*(p), \quad (8)$$

resulting in π^{θ^*} being an optimal long-term average reward policy for p .

The RBMLE policy therefore proposed the optimistic philosophy of favoring parameters with larger rewards, now known as “optimism in the face of uncertainty” (OFU). The inequality (5) indicates why this is fundamentally necessary, since otherwise there is a one-sided exploration bias. RBMLE was the first long-term average reward optimal (also called “asymptotically optimal”) learning algorithm in the frequentist setting (Berry and Fristedt, 1985) that does not resort to forced explorations. In the special case of Bernoulli bandits it was shown to yield particularly simple index policies (Becker and Kumar, 1981). The RBMLE approach has since been applied to a wide range of sequential decision-making, learning, and adaptive control problems. The approach was extended to more general MDPs in (Kumar, 1982; Kumar and Lin, 1982; Borkar, 1990), to LQG systems in (Kumar, 1983a; Campi and Kumar, 1998; Prandini and Campi, 2000), to linear time-invariant systems in (Bittanti et al., 2006), to adaptive control of nonlinear systems in (Kumar, 1983b), to more general ergodic problems in (Stettner, 1993), and to controlled diffusions (Borkar, 1991; Duncan et al., 1994), where its long-term average optimality was established.

A finer notion of optimality than long-term average reward optimality is “regret” (1), which was proposed in (Lai and Robbins, 1985) in the context of multi-armed bandits (MABs). Long-term average optimality corresponds to a regret of $o(T)$, but (Lai and Robbins, 1985) asked the much

more delicate question of how small exactly can regret be made. They were able to sharply characterize the optimal regret as $c \log T + o(\log T)$ for MABs. The performance criterion of regret has now become central to the broader field of Reinforcement learning (RL) (Sutton et al., 1998), which involves an agent repeatedly interacting with an unknown environment that is modeled as a Markov decision process (MDP) (Puterman, 2014) to maximize a total reward. Many algorithms such as UCRL (Auer and Ortner, 2007), UCRL2 (Auer et al., 2009), R-Max (Brafman and Tennenholtz, 2002), REGAL (Bartlett and Tewari, 2012), Posterior Sampling (Strens, 2000), (Osband et al., 2013) and TSMDP (Gopalan and Mannor, 2015) have been studied in great detail, and their learning regret analyzed.

Lai and Robbins also proposed an “Upper Confidence Bound” (UCB) policy which plays the bandit whose upper confidence bound is highest, and showed that it attains the optimal order of regret. The UCB policy also employs the OFU principle by trying arms with larger potential rewards, but in a different way from RBMLE. It has been extended to a wide variety of learning problems: (Brafman and Tennenholtz, 2002; Auer et al., 2002, 2009; Bartlett and Tewari, 2012; Singh et al., 2020).

While the original work analyzed its long-term average optimality, the finite-time regret analysis of RBMLE based algorithms in various settings is an overdue topic of topical interest. An initial step in this direction was taken in (Liu et al., 2020) by analyzing RBMLE for the special case of stochastic multi-armed bandits. The index policy for Bernoulli bandits suggested in (Becker and Kumar, 1981) was generalized to the exponential family of bandits. They analyzed RBMLE’s performance for the exponential family of multi-armed bandits (MABs) and showed that the regret scales as $O(\log T)$. Moreover, numerical experiments in (Liu et al., 2020) clearly exhibited that RBMLE outperforms the UCB in terms of empirical regret, and in fact RBMLE is competitive or slightly better than current state-of-art contenders. Moreover, RBMLE does so with low computational cost in view of its simple indices. Recently (Hung et al., 2020) have proposed an extension of RBMLE for linear contextual bandits which achieves an $O(\sqrt{T} \log T)$ regret, better than the existing state of the art policies like LinTS (Agrawal and Goyal, 2013) and GPUCB (Srinivas et al., 2010). They also show that RBMLE has a competitive regret performance in simulations, and is computationally efficient in comparison with current state-of the art policies such as in (Agrawal and Goyal, 2013), (Srinivas et al., 2010).

Another recent effort (Abbasi-Yadkori and Szepesvári, 2011), motivated by the RBMLE approach of (Campi and Kumar, 1998; Bittanti et al., 2006), addressed the performance of regret for linear quadratic Gaussian (LQG) systems, and established a regret of $\tilde{O}(\sqrt{T})$ ¹.

Due to these developments showing optimal regret performance of RBMLE in these two contexts, it is of interest to examine the regret performance of RBMLE in more general settings. This paper takes the first step in finite-time regret analysis and empirical analysis of the RBMLE algorithm for reinforcement learning (RL) tasks that involve the general problem of optimal control of unknown Markov Decision Processes. Its key contributions are:

1. We propose a new RL algorithm for maximizing rewards for unknown MDPs, that utilizes the RBMLE principle while making control decisions.
2. We analyze the finite-time performance, i.e., the learning regret, of the proposed learning algorithm. We show that the regret is $O(\log T)$.

1. \tilde{O} hides factors that are logarithmic in T .

3. We provide simulation results to show that RBMLE outperforms UCRL2 and TSDE.

With these results, together with the positive results in the context of stochastic MABs (Liu et al., 2020) and linear contextual bandits (Hung et al., 2020), RBMLE provides a second tool for reinforcement learning, complementing the UCB approach.

2. System Model

We consider the MDP described in Section 1, assuming, without loss of generality, that $r(x, u) \in [0, 1]$ for all $(x, u) \in X \times U$. We denote by Π_{sd} the set of all stationary deterministic policies that map X into U , by Π_s the set of all stationary possibly randomized policies, and by $\Pi^*(\theta) := \arg \max_{\pi \in \Pi_{sd}} J(\theta, \pi)$ the set of all optimal stationary deterministic policies for the parameter θ .

Definition 1 (*Unichain MDP*) Under a stationary policy π , let $\tau_{x,y}^\pi$ denote the time taken to hit the state y when started in state x . The MDP is called unichain if $\mathbb{E}[\tau_{x,y}^\pi]$ is finite for all (x, y, π) .

Definition 2 (*Mixing Time*) For a unichain MDP with parameter p , its mixing time T_p is defined as

$$T_p := \max_{\pi \in \Pi_s} \max_{x, y \in X} \mathbb{E}[\tau_{x,y}^\pi].$$

Its ‘‘conductivity’’ κ_p is

$$\kappa_p := \max_{\pi \in \Pi_s} \max_{x \in X} \frac{\max_{y \neq x} \mathbb{E}[\tau_{x,y}^\pi]}{2\mathbb{E}[\tau_{x,x}^\pi]}.$$

Definition 3 (*Gap*) For a stationary policy π , let $\Delta(\pi)$ denote the difference between the optimal average reward and the average reward accrued by π under parameter p , and by Δ_{\min} the gap between the rewards of the best and second best policies,

$$\Delta(\pi) := J^*(p) - J(p, \pi) \text{ and } \Delta_{\min} := \min_{\pi \notin \Pi^*(p)} \Delta(\pi).$$

Definition 4 (*Kullback-Leibler divergence*) For $p_2 = \{p_2(x)\}$ absolutely continuous with respect to $p_1 = \{p_1(x)\}$ the KL-divergence between them is

$$KL(p_1, p_2) := \sum_{x \in X} p_1(x) \log \frac{p_1(x)}{p_2(x)}. \quad (9)$$

For two integers x, y , we use $[x, y]$ to denote the set $\{x, x + 1, \dots, y\}$ and for $a, b \in \mathbb{R}$ we let $a \vee b := \max\{a, b\}$.

Assumption 1 We assume that the following information is known about the unknown transition probabilities $p(x, y, u)$:

- the set of tuples (x, y, u) for which $p(x, y, u) = 0$,
- a lower bound p_{\min} on the non-zero transition probabilities,

$$p_{\min} := \min_{(x,y,u):p(x,y,u)>0} p(x, y, u). \quad (10)$$

We let Θ denote the set

$$\left\{ \theta \in [0, 1]^{|X| \times |X| \times |U|} : \theta(x, y, u) = 0 \text{ if } p(x, y, u) = 0, \sum_{y \in X} \theta(x, y, u) = 1 \forall (x, u), \theta(x, y, u) \geq 0 \right\}. \quad (11)$$

We occasionally refer to $\theta \in \Theta$ as a ‘‘parameter’’ describing the model or transition probabilities.

3. The RBMLE-Based Learning Algorithm

For an MDP parameter $\theta \in \Theta$, denote by $\theta(x, u)$ the vector $\{\theta(x, y, u)\}_{y \in X}$. Let $n(x, u; t)$ be the number of times an action u has been applied in state x until time t , and by $n(x, y, u; t)$ the number of $x \rightarrow y$ one-step transitions under the application of action u . Let $\hat{p}(t) = \{\hat{p}(x, y, u; t)\}$ be the empirical estimate of p at time t , with $\hat{p}(x, y, u; t)$ the MLE of $p(x, y, u)$ at time t ,

$$\hat{p}(x, y, u; t) := \frac{n(x, y, u; t)}{n(x, u; t) \vee 1} \forall x, y \in X \text{ and } u \in U. \quad (12)$$

The RBMLE algorithm: The algorithm evolves in an episodic manner. For episode k , we let τ_k denote its start time and $\mathcal{E}_k := [\tau_k, \tau_{k+1} - 1]$ the set of time-slots that comprise it. The episode durations increase exponentially with episode index, with $|\mathcal{E}_k| = 2^k$. Clearly $\tau_k = \sum_{\ell=1}^{k-1} |\mathcal{E}_\ell|$. Throughout we abbreviate $\hat{p}(\tau_k)$ as \hat{p}_k , $n(x, y; \tau_k)$ as $n_k(x, y)$ and $n(x, y, u; \tau_k)$ as $n_k(x, y, u)$. At the beginning of each episode \mathcal{E}_k , the RBMLE determines²:

(i) A ‘‘reward-biased MLE’’ θ_k :

$$\theta_k \in \arg \max_{\theta \in \Theta} \left\{ \max_{\pi \in \Pi_{sd}} \left\{ \alpha(\tau_k) J(\theta, \pi) - \sum_{(x, u)} n_k(x, u) KL(\hat{p}_k(x, u), \theta(x, u)) \right\} \right\}, \quad (13)$$

$$\text{where } \alpha(t) := a \log(t^b |X|^2 |U|), \text{ with } b > 2, \text{ and } a > \frac{|X|^3 |U|}{2p_{\min} \Delta_{\min}}. \quad (14)$$

(ii) A stationary deterministic policy $\pi_k \in \arg \max_{\pi \in \Pi_{sd}} J(\theta_k, \pi)$ that is optimal for θ_k .

(iii) The action applied for $t \in \mathcal{E}_k$ is $u(t) = \pi_k(x(t))$.

An equivalent Index description of the RBMLE learning algorithm: At the beginning of each episode \mathcal{E}_k , RBMLE attaches an index $I_k(\pi)$ to each $\pi \in \Pi_{sd}$,

$$I_k(\pi) := \max_{\theta \in \Theta} \left\{ \alpha(\tau_k) J(\theta, \pi) - \sum_{(x, u)} n_k(x, u) KL(\hat{p}_k(x, u), \theta(x, u)) \right\}. \quad (15)$$

Within \mathcal{E}_k it implements the policy π_k that has the largest index, i.e.,

$$\pi_k \in \arg \max_{\pi \in \Pi_{sd}} I_k(\pi). \quad (16)$$

For each $\pi \in \Pi_{sd}$, define $\theta_{k, \pi}$ as

$$\theta_{k, \pi} \in \arg \max_{\theta \in \Theta} \left\{ \alpha(\tau_k) J(\theta, \pi) - \sum_{(x, u)} n_k(x, u) KL(\hat{p}_k(x, u), \theta(x, u)) \right\}. \quad (17)$$

2. Throughout the paper, a pre-specified priority order is used to choose a particular maximizer in $\arg \max$ if needed.

4. Preliminary Results

Define the following ‘‘confidence interval’’ $\mathcal{C}(t)$ at time t associated with the empirical estimate $\hat{p}(t)$,

$$\mathcal{C}(t) := \left\{ \theta \in \Theta : |\theta(x, y, u) - \hat{p}(x, y, u; t)| \leq d_1(x, u; t), \forall (x, y, u) \in X \times X \times U \right\}, \quad (18)$$

where

$$d_1(x, u; t) := \sqrt{\frac{\log(t^b |X|^2 |U|)}{n(x, u; t)}}, \quad (19)$$

and $b > 2$. Define also the set \mathcal{G}_1 ,

$$\mathcal{G}_1 := \{\omega : p \in \mathcal{C}(t), \forall t \in \mathbb{N}\}. \quad (20)$$

Lemma 5 *The probability that p lies in $\mathcal{C}(t)$ is bounded as follows:*

$$\mathbb{P}(p \in \mathcal{C}(t)) > 1 - \frac{2}{t^{2b-1} |X|^2 |U|}, \quad \forall t \in \mathbb{N}.$$

Lemma 6

$$|\theta_{k,\pi}(x, y, u) - \hat{p}_k(x, y, u)| \leq d_2(x, u; \tau_k), \quad \forall (x, y, u) \in X \times X \times U,$$

where

$$d_2(x, u; t) := \sqrt{\frac{\alpha(t)}{2n(x, u; t)}}, \quad \forall (x, u) \in X \times U. \quad (21)$$

We now derive a lower bound on the index of any optimal stationary policy $\pi^* \in \Pi^*(p)$ that holds with high probability.

Lemma 7 *On the set \mathcal{G}_1 , the index $I_k(\pi^*)$ of any optimal policy $\pi^* \in \Pi^*(p)$ is lower bounded as follows:*

$$I_k(\pi^*) \geq \alpha(\tau_k)(1 - \gamma)J^*(p), \quad \forall k = 1, 2, \dots, K,$$

where $\gamma := \frac{|X|^3 |U|}{2a \cdot p_{\min} J^*(p)}$.

Next, we show that if the state-action pairs corresponding to a sub-optimal policy have been visited for a sufficiently large number of times, then its index $I_k(\pi)$ is lower than the index of any optimal policy.

Lemma 8 *Let $\pi \notin \Pi^*(p)$ be any sub-optimal stationary deterministic policy. Suppose that the number of visits to each $(x, \pi(x))$ until τ_k is lower bounded as follows,*

$$n(x, \pi(x); \tau_k) > \frac{\alpha(\tau_k)}{c^2} \quad \forall x \in X, \quad (22)$$

where $c := \frac{\beta \Delta_{\min}}{\kappa_p |X|^2 \left(\frac{1}{\sqrt{2}} + \frac{1}{\sqrt{a}}\right)}$, $\beta \in \left(0, 1 - \gamma \frac{J^*(p)}{\Delta_{\min}}\right)$ and $\gamma = \frac{|X|^3 |U|}{2a p_{\min} J^*(p)}$. Then, the index of the sub-optimal policy π at the beginning of the episode k is strictly lower than the index of any optimal policy π^* , i.e., $I_k(\pi) < I_k(\pi^*)$.

5. Regret Analysis

We begin by decomposing the cumulative regret of the learning rule ϕ , $R(\phi, p, T)$ into the sum of episodic regrets as $R(\phi, p, T) := \sum_k \left(J^*(p) |\mathcal{E}_k| - \sum_{t \in \mathcal{E}_k} r(x(t), u(t)) \right)$. Since the RBMLE algorithm implements stationary policy π_k during \mathcal{E}_k , we obtain the following bound on the expected $\mathbb{E}R(\phi, p, T)$ (Lemma 12, Mete et al.),

$$\mathbb{E}R(\phi, p, T) \leq \sum_{k=1}^{K(T)} \mathbb{E} \left(J^*(p) |\mathcal{E}_k| - |\mathcal{E}_k| J(\pi_k, p) \right) + T_p K(T), \quad (23)$$

where $K(T)$ is the number of episodes till T . The first summation can be regarded as the sum of the regrets arising from the policies chosen in the episodes $k = 1, 2, \dots, K(T)$, assuming that each episode is started with a steady-state distribution for the state corresponding to the policy chosen in that episode. The last term $T_p K(T)$ can be regarded as the additional regret due to not starting in a steady state in each episode. We now present the main result of this paper which shows that the expected regret of the RBMLE algorithm is bounded by $c' \log T + c''$ for all T :

Theorem 9 *The regret of the RBMLE based-policy is upper-bounded as*

$$\mathbb{E}R(\phi, p, T) \leq c_1 \kappa_p^2 |X|^5 |U| \left(\frac{\sqrt{\frac{a}{2}} + 1}{\beta \Delta_{\min}} \right)^2 \log T + (\kappa_p |X| |U| + 1) \log_2 T + C \text{ for all } T,$$

where $\beta \in \left(0, 1 - \gamma \frac{J^*(p)}{\Delta_{\min}}\right)$, $\gamma = \frac{|X|^3 |U|}{2ap_{\min} J^*(p)}$, $c_1 \in \left(0, \frac{11}{\kappa_p^2}\right)$, and

$$C = c_1 \kappa_p^2 |X|^5 |U| \left(\frac{\sqrt{\frac{a}{2}} + 1}{\beta \Delta_{\min}} \right)^2 \log(|X|^2 |U|) + (\kappa_p |X| |U| + 1) + |X| |U| + \frac{8}{|X|^2 |U|}.$$

Proof The decomposition (23) shows that: (a) Episodic regret is 0 in those episodes in which π_k is optimal, i.e., $\pi_k \in \Pi^*(p)$. (b) if π_k is not optimal then the episodic regret is bounded by the length of the episode $|\mathcal{E}_k|$, since the magnitude of rewards is less than 1.

Let \mathcal{G}_2 denote the set with $\mathbb{P}(\mathcal{G}_2) \geq 1 - \frac{|X| |U|}{T}$ (Lemma 11, Mete et al.) such that

$$n(x, u; T) \geq \frac{y_{x,u}}{2} - \sqrt{y_{x,u} \log T} \text{ for all } (x, u) \text{ on } \mathcal{G}_2, \quad (24)$$

where $y_{x,u} := \sum_{k \in \mathcal{K}_{x,u}} \left\lfloor \frac{|\mathcal{E}_k|}{2T_p} \right\rfloor$, and $\mathcal{K}_{x,u}$ denotes the set of indices of those episodes up to time T in which action u is taken when state is equal to x . Define the ‘‘good set’’ $\mathcal{G} := \mathcal{G}_1 \cap \mathcal{G}_2$. We first consider the regret on \mathcal{G} .

(i) *Regret due to suboptimal episodes on \mathcal{G}* : Define $n_c := \alpha(T) \left(\frac{\kappa_p |X|^2}{\beta \Delta_{\min}} \left(\frac{1}{\sqrt{2}} + \frac{1}{\sqrt{a}} \right) \right)^2$.

On the ‘‘good set’’ \mathcal{G} , confidence intervals $\mathcal{C}(t)$, defined in (18), hold true for all episode starting times τ_k , $k \in [1, K(T)]$, and also the conclusions of (24) are true. Hence it follows from Lemma 8 that if $n(x, u; \tau_k) > n_c$ for all (x, u) then the regret in \mathcal{E}_k is 0. Otherwise, there exists at least one state, action pair (x, u) with $n(x, u; \tau_k) \leq n_c$. We now upper bound the number of time-steps in such ‘‘sub-optimal’’ episodes $\mathcal{K}_{x,u}$ in which control u is applied in state x . Since $n(x, u; T) \leq n_c$, we have $n_c \geq \frac{y_{x,u}}{2} - \sqrt{y_{x,u} \log T}$. Note that $n_c \geq \kappa_p^2 \log T$. Then, there exists $c_1 < \frac{11}{\kappa_p^2}$ such

that $y_{x,u} \leq c_1 n_c$ (Lemma 13, Mete et al.). So $\sum_{k \in \mathcal{K}_{x,u}} |\mathcal{E}_k| \leq 2T_p c_1 n_c + 2|\mathcal{K}_{x,u}|T_p$. Note that $\mathcal{K}_{x,u} < K(T)$, where $K(T)$ is the total number of episodes till T . We let \mathcal{R}_1 be the total regret until T due to suboptimal episodes on the good set \mathcal{G} . Then

$$\mathcal{R}_1 \leq \sum_{(x,u)} \sum_{k \in \mathcal{K}_{x,u}} |\mathcal{E}_k| \leq 2|X||U|T_p(c_1 n_c + K(T)), \quad (25)$$

where \mathcal{R}_1 is the expected regret on \mathcal{G} .

(ii) *Regret on \mathcal{G}_1^c* : For any episode $k \in [0, K(T)]$, the probability of failure of the confidence interval at the beginning of the episode $C(\tau_k)$ is upper bounded by $\frac{2}{|X|^2|U|\tau_k^{2b-1}}$ as shown in Lemma 5. The expected regret in each such episode is bounded by the length of the episode \mathcal{E}_k . We let \mathcal{R}_3 be the total expected regret till T due to the failure of confidence intervals. It can be upper-bounded as:

$$\mathcal{R}_3 \leq \sum_{k=1}^{K(T)} \frac{2(\tau_{k+1} - \tau_k)}{|X|^2|U|\tau_k^{2b-1}} \leq \sum_{k=1}^{\infty} \frac{2}{|X|^2|U|\tau_k^{2b-2}} \left(\frac{\tau_{k+1}}{\tau_k} \right) \leq \sum_{k=1}^{\infty} \frac{4}{|X|^2|U|\tau_k^{2b-2}} \leq \frac{8}{|X|^2|U|}.$$

(iii) *Regret on \mathcal{G}_2^c* : The probability of the set where the conclusions of (24) do not hold true for a state-action pair (x, u) is upper bounded by $\frac{|X||U|}{T}$. Since the sample-path regret can be trivially upper bounded by T , it follows that $\mathcal{R}_2 \leq |X||U|$.

(iv) *Additional regret due to not starting in a steady state in each episode*: The RBMLE algorithm implements the policy π_k at the beginning of episodes k . The total expected reward in the episode depends on the state at the beginning of the episode $x(\tau_k)$ (Lemma 12, Mete et al.). The policy incurs an additional loss if it starts in an unfavorable state at τ_k which is upper bounded by T_p in each episode. Let \mathcal{R}_4 be the total expected regret due to not starting in a steady state in each episode. Then $\mathcal{R}_4 \leq K(T)T_p$ where, $K(T) = \lceil \log_2 T \rceil$ is the total number of episodes till T .

The proof is completed by adding the bounds on $\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3$ and \mathcal{R}_4 . ■

6. Simulation Experiments

We evaluate the performance of the RBMLE algorithm by empirical comparison with UCRL2 (Auer et al., 2009) and Thompson Sampling. Among the many Thompson Sampling variants, we use TSDE (Ouyang et al., 2017) since the simulation results in (Ouyang et al., 2017) show that TSDE has lower empirical regret than Lazy PSRL (Abbasi-Yadkori and Szepesvári, 2015) and TSMDP (Gopalan and Mannor, 2015). We maintain the fairness of the comparison by following approach: For every state-action pair, we generate a sample path of transitions at the beginning of each experiment. We use these same samples for all three algorithms. For the sake of uniformity with UCRL2 and TSDE, the length of an episode of RBMLE is dynamically determined as follows: An episode is terminated if the number of visits to any state-action pair in the episode exceeds the total visits till the beginning of the episode. For all experiments, the bias term for RBMLE is $10 \log t$ and the confidence parameter for UCRL2 is 0.01. We compare the cumulative regret for three different MDPs in Figure 1. In all experiments, RBMLE outperforms UCRL2 and TSDE.




Figure 1: Empirical Comparison of RBMLE with UCRL2 and TSDE

7. Concluding Remarks

A fundamental challenge in online learning is what is prosaically called the closed-loop identifiability problem (Borkar and Varaiya, 1979). It is the same as the “dual control” problem raised by (Feldbaum, 1960a,b), or the more contemporaneous “exploration vs. exploitation” problem: When a learning algorithm begins to converge, it ceases to learn or explore. It was noticed by (Kumar and Becker, 1982) that in a certainty equivalence context this means that the learnt model will automatically have a one-sided bias of having a smaller optimal reward than the true model. RBMLE was proposed to overcome this fundamental problem by incorporating a counteracting bias in favor of parameters with larger optimal rewards. It provides a general purpose reinforcement learning algorithm for dynamic stochastic systems. Most of the work on RBMLE has been focused on the problem of long-term average optimality, the context in which it was originally proposed. However, current applications emphasize the much finer performance of regret, which captures the growth of the total reward as a function of the horizon T . Recent work examining RBMLE for stochastic bandits (Liu et al., 2020), linear contextual bandits (Hung et al., 2020) has shown that not only does RBMLE have optimal order of regret but it also has excellent empirical performance competitive or better than state of the art algorithms, and it also achieves this with low computational complexity. For the LQG context, recent work motivated by RBMLE also establishes near optimal regret performance (Abbasi-Yadkori and Szepesvári, 2011). With the present paper establishing optimal order of regret for reinforcement learning problems modeled as Markov Decision Processes, the RBMLE complements the UCB approach and provides a second tool for reinforcement learning.

Acknowledgments

This research has been partially supported by NSF under CCF-1934904, Science & Technology Center CCF-0939370, and CMMI-2038625, the USARO under W911NF-18-10331 and W911NF-2-120064, USARL W911NF-19-2-0243, and USONR N00014-18-1-2048. The views and conclusions here do not represent the official policies, either expressed or implied, of NSF, USARO, USONR, USARL, or U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

References

- Yasin Abbasi-Yadkori and Csaba Szepesvári. Regret bounds for the adaptive control of linear quadratic systems. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 1–26, 2011.
- Yasin Abbasi-Yadkori and Csaba Szepesvári. Bayesian optimal control of smoothly parameterized systems. In *UAI*, pages 1–11. Citeseer, 2015.
- Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *International Conference on Machine Learning*, pages 127–135, 2013.
- Peter Auer and Ronald Ortner. Logarithmic online regret bounds for undiscounted reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 49–56, 2007.
- Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- Peter Auer, Thomas Jaksch, and Ronald Ortner. Near-optimal regret bounds for reinforcement learning. In *Advances in neural information processing systems*, pages 89–96, 2009.
- Peter L. Bartlett and Ambuj Tewari. REGAL: A regularization based algorithm for reinforcement learning in weakly communicating mdps. *CoRR*, abs/1205.2661, 2012. URL <http://arxiv.org/abs/1205.2661>.
- A. Becker and P. R. Kumar. Optimal strategies for the n-armed bandit problem. *Univ. Maryland. Baltimore County, Math. Res. Rep.*, pages 81–1, 1981.
- Donald A. Berry and Bert Fristedt. Bandit problems: sequential allocation of experiments (monographs on statistics and applied probability). *London: Chapman and Hall*, 5(71-87):7–7, 1985.
- Sergio Bittanti, Marco C. Campi, et al. Adaptive control of linear time invariant systems: the “bet on the best” principle. *Communications in Information & Systems*, 6(4):299–320, 2006.
- V. Borkar and P. Varaiya. Adaptive control of Markov chains, I: Finite parameter set. *IEEE Transactions on Automatic Control*, 24(6):953–957, 1979.
- V. S. Borkar. The Kumar-Becker-Lin scheme revisited. *Journal of Optimization Theory and Applications*, 66(2):289–309, 1990.
- V. S. Borkar. Self-tuning control of diffusions without the identifiability condition. *Journal of optimization theory and applications*, 68(1):117–138, 1991.
- Ronen I Brafman and Moshe Tennenholtz. R-max-a general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research*, 3(Oct):213–231, 2002.
- Marco C. Campi and P. R. Kumar. Adaptive linear quadratic Gaussian control: the cost-biased approach revisited. *SIAM Journal on Control and Optimization*, 36(6):1890–1907, 1998.
- Grace E. Cho and Carl D. Meyer. Markov chain sensitivity measured by mean first passage times. *Linear Algebra and its Applications*, 316(1-3):21–28, 2000.

- Thomas M. Cover. *Elements of information theory*. John Wiley & Sons, 1999.
- T. E. Duncan, B. Pasik-Duncan, and L. Stettner. Almost self-optimizing strategies for the adaptive control of diffusion processes. *Journal of optimization theory and applications*, 81(3):479–507, 1994.
- A. A. Feldbaum. Dual control theory. i. *Avtomatika i Telemekhanika*, 21(9):1240–1249, 1960a.
- A. A. Feldbaum. Dual control theory. ii. *Avtomatika i Telemekhanika*, 21(11):1453–1464, 1960b.
- Aditya Gopalan and Shie Mannor. Thompson sampling for learning parameterized Markov decision processes. In *Conference on Learning Theory*, pages 861–898, 2015.
- Y-H. Hung et al. Reward-biased maximum likelihood estimation for linear stochastic bandits. *arXiv preprint arXiv:2010.04091*, 2020.
- P. R. Kumar. Adaptive control with a compact parameter set. *SIAM Journal on Control and Optimization*, 20(1):9–13, 1982.
- P. R. Kumar. Optimal adaptive control of linear-quadratic-Gaussian systems. *SIAM Journal on Control and Optimization*, 21(2):163–178, 1983a.
- P. R. Kumar. Simultaneous identification and adaptive control of unknown systems over finite parameter sets. *IEEE Transactions on Automatic Control*, 28(1):68–76, 1983b.
- P. R. Kumar and A. Becker. A new family of optimal adaptive controllers for Markov chains. *IEEE Transactions on Automatic Control*, 27(1):137–146, 1982.
- P. R. Kumar and W. Lin. Optimal adaptive controllers for unknown Markov chains. *IEEE Transactions on Automatic Control*, 27(4):765–774, 1982.
- Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- Xi Liu, Ping-Chun Hsieh, Yu Heng Hung, Anirban Bhattacharya, and P Kumar. Exploration through reward biasing: Reward-biased maximum likelihood estimation for stochastic multi-armed bandits. In *International Conference on Machine Learning*, pages 6248–6258. PMLR, 2020.
- P. Mandl. Estimation and control in Markov chains. *Advances in Applied Probability*, pages 40–60, 1974.
- Akshay Mete, Rahul Singh, and P. R. Kumar. Reward Biased Maximum Likelihood Estimation for Reinforcement Learning, 2020. URL <https://arxiv.org/abs/2011.07738>.
- Michael Mitzenmacher and Eli Upfal. *Probability and computing: Randomization and probabilistic techniques in algorithms and data analysis*. Cambridge university press, 2017.
- Ian Osband, Daniel Russo, and Benjamin Van Roy. (More) efficient reinforcement learning via posterior sampling. In *Advances in Neural Information Processing Systems*, pages 3003–3011, 2013.

- Yi Ouyang, Mukul Gagrani, Ashutosh Nayyar, and Rahul Jain. Learning unknown Markov decision processes: A thompson sampling approach. *arXiv preprint arXiv:1709.04570*, 2017.
- Maria Prandini and Marco C. Campi. Adaptive lqg control of input-output systems—a cost-biased approach. *SIAM Journal on Control and Optimization*, 39(5):1499–1519, 2000.
- Martin L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- Rahul Singh, Abhishek Gupta, and Ness B. Shroff. Learning in Markov decision processes under constraints. *arXiv preprint arXiv:2002.12435*, 2020.
- Niranjan Srinivas, Andreas Krause, Sham M. Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *27-th International Conference on Machine Learning*, 2010.
- Lukasz Stettner. On nearly self-optimizing strategies for a discrete-time uniformly ergodic adaptive model. *Applied Mathematics and Optimization*, 27(2):161–177, 1993.
- Malcolm Strens. A Bayesian framework for reinforcement learning. In *ICML*, volume 2000, pages 943–950, 2000.
- Richard S. Sutton, Andrew G. Barto, et al. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge, 1998.