
Dueling RL: Reinforcement Learning with Trajectory Preferences

Aldo Pacchiano*
Microsoft Research,
New York City

Aadirupa Saha*#
TTI, Chicago

Jonathan Lee
Stanford University

Abstract

We consider the problem of preference-based reinforcement learning (PbRL), where, unlike traditional reinforcement learning (RL), an agent receives feedback only in terms of 1 bit (0/1) preferences over a trajectory pair instead of absolute rewards for it. The success of the traditional reward-based RL framework crucially depends on how accurately a system designer can express an appropriate reward function, which is often a non-trivial task. The main novelty of our framework is the ability to learn from preference-based trajectory feedback that eliminates the need to hand-craft numeric reward models. This paper sets up a formal framework for the PbRL problem with non-Markovian rewards, where the trajectory preferences are encoded by a generalized linear model of dimension d . Assuming the transition model is known, we propose an algorithm with a regret guarantee of $\tilde{O}\left(SHd\log(T/\delta)\sqrt{T}\right)$. We further extend the above algorithm to the case of unknown transition dynamics and provide an algorithm with regret $\tilde{O}((\sqrt{d}+H^2+|\mathcal{S}|)\sqrt{dT}+\sqrt{|\mathcal{S}||\mathcal{A}|TH})$. To the best of our knowledge, our work is one of the first to give tight regret guarantees for preference-based RL problem with trajectory preferences.

* Equal contribution (alphabetically).
#Author is currently with Apple ML Research, US.
Corresponding email: aadirupa@ttic.edu

1 INTRODUCTION

Classical reinforcement learning (RL) with absolute reward feedback is a well-studied framework which is a sequential experience-driven learning process to optimize an accumulated long-term reward (Sutton and Barto, 2018; Auer et al., 2009; Singh et al., 2002). Over the years, several works have addressed RL in terms of both the optimal sample complexity for finding the best policy (Azar et al., 2013; Dann and Brunskill, 2015; Dann et al., 2017; Domingues et al., 2020a; Lattimore and Hutter, 2012) and minimizing regret via balancing exploration and exploitation (Zhang and Ji, 2019; Azar et al., 2017; Ortner, 2020; Talebi and Maillard, 2018; Efroni et al., 2020; Domingues et al., 2020b).

However, a major limitation of the standard RL setting is that its success crucially depends on the prior knowledge encoded into the definition of the reward function. The learned policy can often be sensitive to small changes of the reward, possibly yielding very different behaviors depending on the relative values of the rewards. The choice of reward function in applications such as robotics consequently entails a high amount of non-trivial effort in reward engineering, leading to challenges such as reward shaping, reward hacking, infinite rewards, and multi-objective outcomes (Wirth and Furnkranz, 2013; Wirth et al., 2017).

The framework of *Preference-based Reinforcement Learning* (PbRL) (Busa-Fekete et al., 2014; Wirth et al., 2016, 2017) has been proposed as a fix to this problem, to enforce learning from non-numerical, relative feedback which need not suffer from issues due to the inaccuracy of reward modeling or engineering. This framework widely applies to multiple areas including robot training, stock-prediction, recommender systems, clinical trials, etc. (Novoseller et al., 2019; Sadigh et al., 2017; Christiano et al., 2017; Kupcsik et al., 2018; Jain et al., 2013; Wirth et al., 2017).

While the problem of PbRL was introduced almost a decade ago, most work in it has been primarily ap-

plied or experimental in nature (Riad et al., 2012; Jain et al., 2013; Busa-Fekete et al., 2014; Christiano et al., 2017; Wirth and Fürnkranz, 2013; Wirth et al., 2016, 2017; Ibarz et al., 2018; Kupcsik et al., 2018). There have also been attempts to design suitable algorithms based on varying preference models and problem objectives (Novoseller et al., 2019; Xu et al.), but, to the best of our knowledge, existing theoretical guarantees on PbRL literature are sparse. The performance guarantees of most of the proposed algorithms are not well-understood (Wirth et al., 2017; Xu et al.) except for some very recent attempts (Novoseller et al., 2019; Xu et al.) as discussed below in the section on related work. We consider the problem of provably finding the best finite-horizon policy (i.e., one with highest expected reward) for an unknown Markov decision process (MDP), but with only relative preference feedback on H -length trajectories.

Problem Setup (informal). Consider a T -round, H -horizon MDP $(\mathbb{P}, \mathcal{S}, \mathcal{A}, H)$, with \mathcal{S} and \mathcal{A} being the finite sets of states and actions, respectively, and \mathbb{P} representing the transition dynamics of the MDP. We consider a real-valued score function $s(\tau)$, which is neither known nor queryable, that scores a given trajectory τ . We assume preference of any two trajectories τ_1 and τ_2 is determined by their underlying score difference, i.e. $P(\tau_1 \succ \tau_2) = \sigma(s(\tau_1) - s(\tau_2))$, $\sigma : \mathbb{R} \mapsto [0, 1]$ being a suitable link-function. In particular, we assume $s(\cdot)$ is an (unknown) linear function of the trajectory-feature $\phi(\tau) \in \mathbb{R}^d$, and the link-function σ is the sigmoid Li et al. (2017). The goal is to minimize the regret with respect to the optimal policy.

An important thing to note is that in our setting the trajectory features $\phi(\tau) \in \mathbb{R}^d$ are not necessarily sum-decomposable (over individual state-action features of the trajectory) and the underlying reward function is non-Markovian. In this case, the optimal policy may be history dependent. This is more general than assuming the trajectory reward is a linear function of the sum of per-state features, e.g. in Novoseller et al. (2019). Under the latter more limiting assumption, the traditional linear bandit techniques can be easily used to derive regret guarantees. Since the number of history dependent policies is super exponential, to deal with this more general setting we first show the log-covering number of the history dependent policies that are an optimal policy of an MDP with a trajectory score specified by our form of trajectory feedback is upper bounded by a polynomial quantity. Our specific contributions are as follows:

1. To the best of our knowledge, we are the first to formulate and analyze the finite time regret guarantee for preference-based linear bandits problem

with non-Markovian reward models (Sec. 2).

2. We propose an algorithm for known transition dynamics which is shown to yield a regret guarantee of $\tilde{O}\left(SHd \log(T/\delta) \sqrt{T} \right)$ (Sec. 3).¹
3. We further generalize our algorithm to the case of unknown models and propose an algorithm with regret guarantee $\tilde{O}((\sqrt{d} + H^2 + |\mathcal{S}|) \sqrt{dT} + \sqrt{|\mathcal{S}||\mathcal{A}|TH})$ (Sec. 4).

Related Work. Over the last two decades the problem of learning from preference feedback in bandits, known as *dueling bandits*, has gained much attention (Yue et al., 2012; Zoghi et al., 2014b, 2015). Dueling bandits generalizes the standard multi-armed bandit (MAB) (Auer et al., 2002). The goal is to identify a set of ‘good’ arms from a larger fixed set of arms by querying preference feedback for pairs of actively chosen arms. Yue and Joachims (2009, 2011); Saha and Krishnamurthy (2022); Ghoshal and Saha (2022); Saha and Gopalan (2018a) The setting is relevant in various real-world systems which aim to collect information from user preferences, including recommender systems, retail management, search engine optimization, job scheduling, etc. Towards these goals, several algorithms have been proposed (Ailon et al., 2014; Zoghi et al., 2014a; Komiyama et al., 2015; Gajane et al., 2015; Saha and Gopalan, 2018b, 2019).

Though there has been a fair amount of research for preference-based bandits (no state information), few works consider incorporating preference feedback in the reinforcement learning (RL) framework, which considers the problem of long-term objectives over a markov decision process (Singh et al., 2002; Ng et al., 2006; Talebi and Maillard, 2018; Ortner, 2020; Zhang and Ji, 2019; Zanette and Brunskill, 2019). However the classical RL setup assumes access to reward feedback for each state-action pair which might be impractical in many real world scenarios. Few very recent works considered training RL agents based on general trajectory-based reward which are available only at the end of each trajectory (Efroni et al., 2021; Chatterji et al., 2021), but their setting still assumes access to absolute reward feedback, unlike the case in PbRL. Some initial works consider the applied PbRL problem inspired by the problems of reward hacking, reward shaping, difficulty to model infinite rewards or multi-objective trade-offs (Busa-Fekete et al., 2014; Wirth et al., 2016, 2017; Christiano et al., 2017) etc.

Novoseller et al. (2019) made the first attempt to analyze the finite T -round regret guarantee for the PbRL

¹The notation $\tilde{O}(\cdot)$ hides logarithmic factors in $T, H, |\mathcal{A}|, |\mathcal{S}|$.

problem with trajectory preference feedback, where the learner is allowed to run two independent trajectories in parallel and receive 0/1 preference feedback after every such h -length roll out. Assuming an underlying MDP model, the preference between two h -length trajectories is modeled as being proportional to the accumulated reward of the corresponding trajectories. The authors propose a Double Posterior Sampling (DPS) technique with asymptotically sublinear regret.

Xu et al. models reward-free trajectory preferences and analyses the sample complexity of finding the ϵ -best-policy. Their proposed algorithm crucially depends on an underlying dueling bandit black box whose performance guarantee is restricted to preference structures like Strong Stochastic Transitivity and Stochastic Triangle Inequality. Furthermore, the algorithms proposed in this work are not shown to enjoy provably optimal sample complexity, and, moreover, the fundamental performance limit of sample complexity is also not explicitly analyzed.

The literature of multi-agent reinforcement learning in Markov games closely relates to the setup of PbRL which attempts the problem of reaching Nash equilibrium of a simultaneous move markov game based on per-state win-loss feedback of the two (or multiple) players. Bai and Jin (2020); Bai et al. (2020); Liu et al. (2021) address the problems from finite action two player markov games, while Xie et al. (2020) extended this setting to zero sum games with linear function approximation. However, all these works analyzed the episodic sample complexity of the learning algorithm towards finding an ϵ -approximate Nash equilibrium which is fairly unrelated to the regret objective of PbRL problem we considered in this paper.

Another closely related sub-field of RL, *imitation learning*, addresses the objective of learning optimal behavior from trajectories suggested by an expert. In Ng et al. (2000); Boularias et al. (2011); Neu and Szepesvári (2012); Wulfmeier et al. (2015), inverse reinforcement learning problems have been considered, where the objective is to extract (unknown) reward function from the trajectories given by an oracle or expert. Once the reward functions are computed, any RL algorithm could, in principle, be applied to compute the optimal policy. Ho and Ermon (2016) propose a generative adversarial network based imitation learning algorithm that computes the optimal policy directly from the trajectories of expert. Our work is fundamentally different in the sense that we do not receive trajectories or optimal actions from an expert. Instead, we get preferences over sample trajectories that are posed as queries to a system expert for preference feedback.

2 PROBLEM SETUP

Notation. Let $[n]$ denote the set $\{1, 2, \dots, n\}$. Given a set A , for any two items $x, y \in A$, we denote that i is preferred over j by $x \succ y$. By $\mathcal{B}_r(d)$ we denote the ℓ_2 -norm ball of radius r in dimension d . Lower case bold letters denote vectors, upper case bold letters denote matrices.

RL Model. Consider a T -episode, H -horizon RL setup $(\mathbb{P}, \mathcal{S}, \mathcal{A}, H, \rho)$, \mathcal{S} is a finite set of states, \mathcal{A} is a set of actions, $\mathbb{P}(\cdot | s, a)$ is the MDP transition dynamics given a state and action pair $(s; a)$, $H \in \mathbb{N}$ is the length of an episode, ρ denotes the initial distribution over states.

We denote a trajectory by concatenation of all states and actions visited during H steps $\tau := (s_1, a_1, \dots, s_H, a_H)$. In general let $\tau_{h:H} := (s_h, a_h, \dots, s_H, a_H)$ denote the states and action from step h until the end of the episode. We denote by τ_h to be all the states and actions taken up to step h and define $\tau_0 = \emptyset$. Let Γ be the set of all possible trajectories of length H , similarly Γ_h denotes the set of all sub-trajectories up to step h . We use the superscript t as in τ^t to denote a trajectory sampled during the t -th episode. At the start of each episode, we assume the initial state s_1 is drawn from a fixed distribution ρ known to the learner apriori (for example concentrated on an initial state s_0).

Trajectory embedding. For any trajectory τ we assume the existence of a trajectory embedding function $\phi : \Gamma \rightarrow \mathbb{R}^d$. We denote by $\phi(\tau)$ to the d -dimensional embedding of trajectory τ . The map ϕ is known to the learner. One special case of such a trajectory-dependent feature map is a decomposed embedding, where $\phi(\tau) = \sum_{h=1}^H \phi(s_h, a_h)$ and $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$ is a mapping from state-actions pairs to \mathbb{R}^d . Examples of these trajectory embeddings can be borrowed from the Behavior Guided class of algorithms for policy optimization found in Pacchiano et al. (2020). Many practically relevant trajectory or state-action embeddings can be found defined in Pacchiano et al. (2020); Parker-Holder et al. (2020). It is conceivable the preference model may be based on one of these embedding maps.

Policy embedding. The above feature embedding also leads to a natural mean embedding of any policy $\pi : \mathcal{S} \mapsto \mathcal{A}$ given by $\phi(\pi) := \mathbf{E}_{\tau \sim \pi}[\phi(\tau)]$.

Preference modeling. Assuming $\mathbf{w}^* \in \mathbb{R}^d$ to be an unknown vector, we define the pairwise-preference of trajectory τ_1 over τ_2 as:

$$\begin{aligned} \mathbb{P}(\tau_1 \succ \tau_2) &= \sigma(\langle \phi(\tau_1) - \phi(\tau_2), \mathbf{w}^* \rangle) \\ &= \frac{\exp(\phi(\tau_1)^\top \mathbf{w}^*)}{\exp(\phi(\tau_1)^\top \mathbf{w}^*) + \exp(\phi(\tau_2)^\top \mathbf{w}^*)} \end{aligned} \quad (1)$$

where $\sigma : \mathbb{R} \mapsto [0, 1]$ is the logistic link function, i.e. $\sigma(x) = (1 + e^{-x})^{-1}$. We can ‘lift’ the definition of a comparison from trajectories to policies by setting,

$$\mathbb{P}(\pi_1 \succ \pi_2) = \sigma(\langle \phi(\pi_1) - \phi(\pi_2), \mathbf{w}^* \rangle) \quad (2)$$

Equation 1 says the probability of any trajectory τ_1 being preferred over τ_2 is essentially proportional to the score difference of the individual trajectories, assuming the score for any trajectory τ is defined as the function

$$s(\tau) := \langle \phi(\tau), \mathbf{w}^* \rangle.$$

The linear score of any policy π (expectation over trajectories) can be similarly defined as $s(\pi) := \mathbf{E}_{\tau \sim \pi}[\langle \phi(\tau), \mathbf{w}^* \rangle]$ and therefore $\mathbb{P}(\pi_1 \succ \pi_2) = \sigma(s(\pi_1) - s(\pi_2))$.

Non markovian policy class. The performance of all our algorithms will be measured against the policy that maximizes $s(\pi)$. Since $s(\tau)$ may be a non-markovian function of the trajectory, the policy optimizing this objective need not be markovian. We therefore set Π as the set of all history dependent policies. In contrast with standard markovian RL works, this is one of the main sources of technical complexity of our setting.

Assumption 1. [Bounded parameter] We assume that $\|\mathbf{w}^*\| \leq W$ for some known $W > 0$.

Assumption 2. [Bounded feature maps] For all trajectories τ we assume that $\|\phi(\tau)\| \leq B$ for some known $B > 0$.²

Definition 1. The degree of non-linearity of the sigmoid σ over the parameter space (denoting the first derivative of σ by σ') is given by

$$\kappa := \sup_{\mathbf{x} \in \mathcal{B}_B(d), \mathbf{w} \in \mathcal{B}_S(d)} \frac{1}{\sigma'(\mathbf{w}^\top \mathbf{x})}.$$

Objective: Alternative: The objective of the learner is to minimize regret by finding policies to maximize the sum of their expected scores over T rounds. At each round t , the learner proposes two policies, π_t^1 and π_t^2 , which are executed in the MDP generating trajectories τ_t^1 and τ_t^2 . The learner then receives feedback in the form of the Bernoulli variable $o_t \in \{0, 1\}$ which specifies whether τ_t^1 is preferred ($o_t = 1$) or τ_t^2 is preferred ($o_t = 0$). The preference feedback o_t is distributed according to $P(\tau_t^1 \succ \tau_t^2)$. We measure the learner’s performance via its pseudo-regret w.r.t. policy class Π , which we define as:

$$\begin{aligned} R_T^{\text{scr}} &:= \max_{\pi \in \Pi} \sum_{t=1}^T \frac{[(2\phi(\pi) - \phi(\pi_t^1) - \phi(\pi_t^2))^\top \mathbf{w}^*]}{2} \\ &= \sum_{t=1}^T \frac{2s(\pi^*) - (s(\pi_t^1) + s(\pi_t^2))}{2}, \end{aligned} \quad (3)$$

where $\pi^* := \max_{\pi \in \Pi} s(\pi)$. This essentially measures the performance of the learner at round t in terms of average score of the played policies π_t^1, π_t^2 w.r.t. the score maximizing policy π^* .

Remark 1. An important thing to note is that representing any trajectory pair (τ_1, τ_2) by the feature $(\phi(\tau_1) - \phi(\tau_2)) \in \mathbb{R}^d$, our underlying preference model is similar to reward model of Chatterji et al. (2021) (see Assumption 2.1). The fundamental difference between our setting and that of Chatterji et al. (2021) is the nature of the dueling feedback. In our work, the only way to gather any information when interacting with the world is by comparing the trajectories of two different policies. This adds a layer of complexity not present in the per trajectory feedback model from Chatterji et al. (2021), that makes their algorithms not immediately applicable to our setting.

One may think of using our preference model (Equation 2) to define an alternative notion of regret:

$$R_T^{\text{pref}} := \max_{\pi \in \Pi} \sum_{t=1}^T \frac{\mathbb{P}(\pi \succ \pi_t^1) + \mathbb{P}(\pi \succ \pi_t^2) - 1}{2}. \quad (4)$$

Fortunately, these two notions of regret can be shown to be ‘equivalent’ in the following sense,

Claim 1. Let $\pi^* \in \arg \max_{\pi \in \Pi} s(\pi)$. Then π^* also achieves the max in Eqn. 4.

The logistic link function is increasing w.r.t. its argument, thus for any π, π' we have $\mathbb{P}(\pi^* \succ \pi') = \sigma(s(\pi^*) - s(\pi')) \geq \sigma(s(\pi) - s(\pi')) = \mathbb{P}(\pi \succ \pi')$. It follows that for all t and all π :

$$\mathbb{P}(\pi^* \succ \pi_t^1) + \mathbb{P}(\pi^* \succ \pi_t^2) \geq \mathbb{P}(\pi \succ \pi_t^1) + \mathbb{P}(\pi \succ \pi_t^2)$$

thus establishing the claim.

This argument can also be used to show R_T^{scr} and R_T^{pref} are equivalent up to constant factors when $B, W \leq 1$. The proof is given in Appendix A.

Claim 2. $\frac{R_T^{\text{scr}}}{2(e+1)} \leq R_T^{\text{pref}} \leq \frac{R_T^{\text{scr}}}{2}$.

We conclude that a strategy that attains sublinear R_T^{scr} regret also has sublinear R_T^{perf} regret.

²Note B could essentially depend on the trajectory-length H .

3 KNOWN MODEL: ALGORITHM AND ANALYSIS

In this section, we introduce and analyze an algorithm for solving the preference-based RL problem when the transition model, \mathbb{P} , that governs the probability of transitioning to a next state is known to the learner. In this case, it becomes possible to directly compute expected features induced by policies; however, the *difficulty of learning based only on preference feedback as opposed to rewards remains*. This is because we have access to feedback only through relative preferences on the trajectories rather than an assumed known reward function. Before stating the algorithm, we first detail a method of estimating the underlying parameter \mathbf{w}_* in the logistic model. This procedure serves as a basis for the algorithm.

3.1 Maximum Likelihood Estimation

In the logistic model, a natural way of computing an estimator \mathbf{w}_t of \mathbf{w}_* given trajectory pairs $\{(\tau_\ell^1, \tau_\ell^2)\}_{\ell=1}^{t-1}$ and preference feedback values $\{o_\ell\}_{\ell=1}^{t-1}$ is via maximum likelihood estimation. At time t the regularized log-likelihood (or negative cross-entropy loss) of a parameter \mathbf{w} can be written as:

$$\mathcal{L}_t^\lambda(\mathbf{w}) = \sum_{\ell=1}^{t-1} (o_\ell \log(\sigma(\langle \phi(\tau_\ell^1) - \phi(\tau_\ell^2), \mathbf{w} \rangle)) - \frac{\lambda}{2} \|\mathbf{w}\|_2^2 + (1 - o_\ell) \log(1 - \mu(\langle \phi(\tau_\ell^1) - \phi(\tau_\ell^2), \mathbf{w} \rangle)),$$

where $\lambda > 0$ is a regularization parameter. The function \mathcal{L}_t^λ is strictly concave for $\lambda > 0$. The maximum likelihood estimator $\widehat{\mathbf{w}}_t^{\text{MLE}}$ can be written as $\widehat{\mathbf{w}}_t^{\text{MLE}} = \arg \max_{\mathbf{w} \in \mathbb{R}^d} \mathcal{L}_t^\lambda(\mathbf{w})$. Unfortunately, $\widehat{\mathbf{w}}_t^{\text{MLE}}$ may not satisfy the boundedness Assumption 1, so we instead make use of a projected version of $\widehat{\mathbf{w}}_t^{\text{MLE}}$. Following Faury et al. (2020), and recalling Assumption 1, we define a data matrix and a transformation of $\widehat{\mathbf{w}}_t^{\text{MLE}}$ given by

$$\mathbf{V}_t = \kappa \lambda \mathbb{I}_d + \sum_{\ell=1}^{t-1} (\phi(\tau_\ell^1) - \phi(\tau_\ell^2)) (\phi(\tau_\ell^1) - \phi(\tau_\ell^2))^\top$$

$$g_t(\mathbf{w}) = \sum_{\ell=1}^{t-1} \sigma(\langle \phi(\tau_\ell^1) - \phi(\tau_\ell^2), \mathbf{w} \rangle) (\phi(\tau_\ell^1) - \phi(\tau_\ell^2)) + \lambda \mathbf{w}$$

Then, the projected parameter, along with its confidence set, is given by

$$\mathbf{w}_t^L = \arg \min_{\mathbf{w} \text{ s.t. } \|\mathbf{w}\| \leq W} \|g_t(\mathbf{w}) - g_t(\widehat{\mathbf{w}}_t^{\text{MLE}})\|_{\mathbf{V}_t^{-1}} \quad (5)$$

$$\mathcal{C}_t(\delta) = \{\mathbf{w} \text{ s.t. } \|\mathbf{w} - \mathbf{w}_t^L\|_{\mathbf{V}_t} \leq 2\kappa\beta_t(\delta)\} \quad (6)$$

where $\beta_t(\delta) = \sqrt{\lambda W} + \sqrt{\log(1/\delta) + 2d \log(1 + \frac{tB}{\kappa\lambda d})}$. We restate a bound by Faury et al. (2020) that shows

the probability of \mathbf{w}_* being in $\mathcal{C}_t(\delta)$ for all $t \geq 1$ can be lower bounded.

Lemma 1. [Lemma 1 from Faury et al. (2020)³] Let $\delta \in (0, 1]$ and define the event that \mathbf{w}_* is in the confidence interval $\mathcal{C}_t(\delta)$ for all $t \in \mathbb{N}$:

$$\mathcal{E}_\delta = \{\forall t \geq 1, \mathbf{w}_* \in \mathcal{C}_t(\delta)\}.$$

Then $\mathbb{P}(\mathcal{E}_\delta) \geq 1 - \delta$.

3.2 Algorithm and Analysis

We are now ready to state the Logistic Preference based Reinforcement Learning (LPbRL) algorithm with known model, shown in Algorithm 1. Before any interaction or feedback, we initialize identical data matrices $\mathbf{V}_1 = \bar{\mathbf{V}}_1 = \kappa \lambda \mathbb{I}_d$, $\lambda > 0$ being a regularization parameter. \mathbf{V}_t , as defined before, is designed to track the exact covariates used in the maximum likelihood estimation. $\bar{\mathbf{V}}_t$ (Line 10) on the other hand tracks a similar quantity, but instead uses the expected features under a given policy.

At each round t , we then compute an estimate \mathbf{w}_t^L and determine a set of candidate policies Π_t for which no other policy π significantly outperforms a member of Π_t . The threshold for what constitutes “significant” is determined by the uncertainty in the estimate of \mathbf{w}_t^L . We then search over this set to identify two policies, π_t^1 and π_t^2 , with expected features that maximize the uncertainty determined by $\bar{\mathbf{V}}_t$, precisely by choosing $(\pi_t^1, \pi_t^2) = \arg \max_{\pi^1, \pi^2 \in \Pi_t} \|\phi(\pi^1) - \phi(\pi^2)\|_{\bar{\mathbf{V}}_t^{-1}}$. Both policies are deployed, inducing trajectories τ_t^1 and τ_t^2 and feedback o_t is received. We then update the data matrices \mathbf{V}_t and $\bar{\mathbf{V}}_t$ with the trajectory features $\phi(\tau_t^1) - \phi(\tau_t^2)$ and expected features $\phi(\pi_t^1) - \phi(\pi_t^2)$, respectively. The procedure is repeated for each round $t \in [T]$.

Theorem 1. Let $\delta \leq 1/e$ and $\lambda \geq \frac{B}{\kappa}$. Then, with probability at least $1 - \delta$, the expected regret of Algorithm 1 can be bounded by

$$R_t \leq (4\kappa\beta_t(\delta) + 2\alpha_{d,T}(\delta)) \sqrt{2Td \log\left(1 + \frac{TB}{\kappa d}\right)}.$$

Note there is no dependence on the size of the state or action spaces on account of the model being known in this setting. Furthermore, we note that any dependence on the horizon H is effectively accounted for in the size

³A slight modification in the expression of $\beta_t(\delta)$ is needed to incorporate the fact that we assume $\|\phi(\tau)\| \leq B$ for any τ (Assumption 2) while $B = 1$ in Faury et al. (2020). But this can be easily incorporated using Thm. 1 and Lem. 10 of Abbasi-Yadkori et al. (2011) in the final step of the proof of Lem. 12 of Faury et al. (2020)

of the constant B that bounds the norm of the trajectory features $\phi(\tau)$. For example, if $\phi(\tau)$ decomposes in a per-timestep fashion as $\phi(\tau) = \sum_{h \in [H]} \phi(s_h, a_h)$ where each h satisfies $\|\phi(s_h, a_h)\| \leq B'$, then a trivial bound would give $B \leq B'H$. However, Assumption 2 allows for greater generality.

Algorithm 1 LPbRL: Regret minimization (Known Model)

- 1: **input:** Regularization parameter λ , Learning rate $\eta_t > 0$, exploration length $t_0 > 0$
- 2: Define $\alpha_{d,T}(\delta) = 20BW\sqrt{d \log(T(1+2T)/\delta)}$ and $\gamma_t(\delta) = 2\kappa\beta_t(\delta) + \alpha_{d,T}(\delta)$.
- 3: Initialize $\bar{\mathbf{V}}_t = \kappa\lambda\mathbb{I}_d$
- 4: **for** $t = 1, 2, \dots, T$ **do**
- 5: Compute \mathbf{w}_t^L (see Eqn. Eq. (5))
- 6: Set $\Pi_t = \{\pi^1 | (\phi(\pi^1) - \phi(\pi))^T \mathbf{w}_t^L +$

$$\gamma_t(\delta) \|\phi(\pi^1) - \phi(\pi)\|_{\bar{\mathbf{V}}_t^{-1}} \geq 0 \ \forall \pi\}$$

- 7: Compute

$$(\pi_t^1, \pi_t^2) = \arg \max_{\pi^1, \pi^2 \in \Pi_t} \|\phi(\pi^1) - \phi(\pi^2)\|_{\bar{\mathbf{V}}_t^{-1}}.$$

- 8: Sample $\tau_t^1 \sim \pi_t^1$ and $\tau_t^2 \sim \pi_t^2$.
- 9: Play the duel (τ_t^1, τ_t^2) and receive $o_t = \mathbb{1}(\tau_t^1 \text{ beats } \tau_t^2)$
- 10: Update

$$\bar{\mathbf{V}}_{t+1} = \bar{\mathbf{V}}_t + (\phi(\pi_t^1) - \phi(\pi_t^2))(\phi(\pi_t^1) - \phi(\pi_t^2))^T$$

- 11: **end for**
-

Remark 2. *Theorem 1 shows that for a sufficiently large choice of the regularization parameter λ , the pseudo-regret of Algorithm 1 is at most $R_t = \mathcal{O}\left(\left(W\sqrt{\kappa B} + WB\right) d \log(TB/\kappa\delta)\sqrt{T}\right)$. Importantly, the regret scales nearly optimally with $d\sqrt{T}$ dependency given existing $\Omega(d\sqrt{T})$ lower bounds for linear bandits (Lattimore and Szepesvári, 2020) and known reductions between the standard and preference regret Saha (2021). Assuming κ to be constant, we pay the additional factors in B and W due to non-Markovian rewards which are only indirectly revealed to the learner in terms of preferences.*

3.3 Regret Analysis: Proof Sketch of Thm. 1

We now sketch the proof of Theorem 1. Details and proofs of supporting results can be found in Appendix B.1. The *main idea* of the proof is to ensure that Π_t contains only candidate policies that are predicted to be “sufficiently good” under the learned model \mathbf{w}_t^L using the size of the confidence set $\mathcal{C}_t(\delta)$. We must also verify that Π_t always contains the optimal policy π^* .

Thus, as long as the set $\mathcal{C}_t(\delta)$ shrinks at a sufficiently fast rate, our algorithm will have sublinear regret.

However, in order to judge the uncertainty in predictions of the expected value $\phi(\pi)^T \mathbf{w}_t^L$ of a policy π , we must relate the data matrix \mathbf{V}_t that controls the accuracy of the learned parameter \mathbf{w}_t^L (see Lemma 1), and its expected counterpart $\bar{\mathbf{V}}_t$ (used to define Π_t). The set Π_t is characterized via $\bar{\mathbf{V}}_t$ because this way it allows us to relate it to the algorithm’s regret, a quantity that depends on the expected features of the played policies. Corollary 1 establishes that distances $\|\mathbf{w}_t^L - \mathbf{w}^*\|_{\bar{\mathbf{V}}_t}$ weighted by $\bar{\mathbf{V}}_t$ are not too far from the same distances $\|\mathbf{w}_t^L - \mathbf{w}^*\|_{\mathbf{V}_t}$ weighted by \mathbf{V}_t . Let

$$\mathcal{E}_{prec} = \{\bar{\mathbf{V}}_T \preceq 2\mathbf{V}_T + 84B^2 d \log((1+2T)/\delta)\mathbb{I}_d\}.$$

Corollary 1. *Under Assumption 1, conditioned on event $\mathcal{E}_\delta \cap \mathcal{E}_{prec}$, for any $t \in [T]$*

$$\|\mathbf{w}^* - \mathbf{w}_t^L\|_{\bar{\mathbf{V}}_t} \leq 4\kappa\beta_t(\delta) + \alpha_{d,T}(\delta),$$

where $\alpha_{d,T}(\delta) = 20BW\sqrt{d \log(T(1+2T)/\delta)}$. Furthermore, if $\delta \leq 1/e$, then $\mathbb{P}(\mathcal{E}_\delta \cap \mathcal{E}_{prec}) \geq 1 - \delta - \delta \log_2 T$.

The proof of above is given in Appendix B. Leveraging this relationship, we can establish that the confidence set of policies Π_t defined in line 5 of Algorithm 1 will contain the optimal policy.

Lemma 2. *Conditioned on event $\mathcal{E}_\delta \cap \mathcal{E}_{prec}$, $\pi^* \in \Pi_t$,*

The remainder of the proof now consists of showing the instantaneous regret can be bounded in terms of the size of the confidence sets and the uncertainty values $\|\phi(\pi_t^1) - \phi(\pi_t^2)\|_{\bar{\mathbf{V}}_t}$. We defer the final details to Appendix B.3.

4 UNKNOWN MODEL: ALGORITHM AND ANALYSIS

Algorithm description. The LPbRL algorithm for unknown dynamics models works in a similar way to Algorithm 1. The main differences lay in the definition of the set Π_t . Whereas in Algorithm 1 this set of policies can be defined without taking into account the model uncertainty, in this case the set of policies to optimize over needs to be carefully constructed in such a way that it can be shown to contain π_* (see Lemma 4). With this in mind we start by introducing the necessary technical tools that will be used throughout this section to deal with model uncertainty.

4.1 Analysis of instantaneous regret:

For any policy π and any MDP model \mathbb{P} , we denote by $\phi^{\mathbb{P}}(\pi)$ to the mean feature of policy π in model \mathbb{P} .

We use the notation $N_t(s, a)$ to denote the number of samples of action a at state s the learner has collected up to time t . We use the notation $\widehat{\mathbb{P}}_t$ to denote the empirical model at time t . We use an 'empirical' version of $\widehat{\mathbf{V}}_t$ defined using the average features computed using the model available at time t :

$$\widehat{\mathbf{V}}_t = \kappa \lambda \mathbb{I}_d + \sum_{\ell=1}^{t-1} \left(\widehat{\phi}^{\widehat{\mathbb{P}}_t}(\pi_\ell^1) - \widehat{\phi}^{\widehat{\mathbb{P}}_t}(\pi_\ell^2) \right) \left(\widehat{\phi}^{\widehat{\mathbb{P}}_t}(\pi_\ell^1) - \widehat{\phi}^{\widehat{\mathbb{P}}_t}(\pi_\ell^2) \right)^\top \quad (7)$$

Our confidence intervals will use a Mahalanobis norm defined by this covariance matrix. Throughout this section we will make heavy use of some of the results from Chatterji et al. (2021). With that in mind we will define a variety of bonus terms. Given any $\eta > 0$ define,

$$\xi_{s,a}^{(t)}(\eta, \delta) = \min \left(2\eta, 4\eta \sqrt{\frac{U}{N_t(s, a)}} \right) \\ \text{s.t. } U = H \log(|\mathcal{S}||\mathcal{A}|) + \log \left(\frac{6 \log(N_t(s, a))}{\delta} \right).$$

We define the following 'bonus' function corresponding to the expectation of these bonus terms summed over a trajectory sampled from a policy π in the model $\widehat{\mathbb{P}}_t$,

$$\widehat{B}_t(\pi, \eta, \delta) = \mathbb{E}_{s_1 \sim \rho, \tau \sim \widehat{\mathbb{P}}_t^\pi(\cdot|s_1)} \left[\sum_{h=1}^{H-1} \xi_{s_h, a_h}^{(t)}(\eta, \delta) \right].$$

Similar to the previous theorem, we must relate $\|\mathbf{w}^* - \mathbf{w}_t^L\|_{\widehat{\mathbf{V}}_t}$ and $\|\mathbf{w}^* - \mathbf{w}_t^L\|_{\widehat{\mathbf{V}}_t}$. We do this via a series of Lemmas.

Lemma 3. *Let $\bar{\mathcal{E}}_0$ be the event that for all $t \in \mathbb{N}$,*

$$\|\mathbf{w}_t^L - \mathbf{w}^*\|_{\widehat{\mathbf{V}}_t} \leq \sqrt{2} \|\mathbf{w}_t^L - \mathbf{w}^*\|_{\widehat{\mathbf{V}}_t} + \sqrt{\sum_{\ell=1}^{t-1} 4 \left(\widehat{B}_t \left(\pi, 2WB, \frac{\delta'}{8\ell^3|\mathcal{A}|\mathcal{S}} \right) \right)^2} + \frac{1}{t}.$$

where $\delta' = \frac{\delta}{(1+4W)^\alpha}$ and $\epsilon = \frac{1}{t^2 \kappa \lambda + 4B^2 t^3}$. Then $\mathbb{P}(\bar{\mathcal{E}}_0) \geq 1 - \delta$.

The proof of Lemma 3 is in Appendix C.2. We now proceed to define the set of plausible policies Π_t . To do so, it will be useful to introduce the following confidence radius multiplier

$$\gamma_t = \sqrt{2} (4\kappa\beta_t(\delta) + \alpha_{d,T}(\delta)) + \frac{1}{t} + 2 \sqrt{\sum_{\ell=1}^{t-1} \widehat{B}_t^2 \left(\pi_\ell^1, 2WB, \frac{\delta'}{8\ell^3|\mathcal{A}|\mathcal{S}} \right) + \widehat{B}_t^2 \left(\pi_\ell^2, 2WB, \frac{\delta'}{8\ell^3|\mathcal{A}|\mathcal{S}} \right)}.$$

Finally,

$$\Pi_t = \left\{ \pi^1 \mid \left(\widehat{\phi}^{\widehat{\mathbb{P}}_t}(\pi^1) - \widehat{\phi}^{\widehat{\mathbb{P}}_t}(\pi) \right)^\top \mathbf{w}_t^L + \gamma_t \|\widehat{\phi}^{\widehat{\mathbb{P}}_t}(\pi^1) - \widehat{\phi}^{\widehat{\mathbb{P}}_t}(\pi)\|_{\widehat{\mathbf{V}}_t^{-1}} + \widehat{B}_t \left(\pi^1, 2SB, \frac{\delta}{2|\mathcal{A}|\mathcal{S}} \right) + \widehat{B}_t \left(\pi, 2SB, \frac{\delta}{2|\mathcal{A}|\mathcal{S}} \right) \geq 0, \forall \pi \right\}.$$

Algorithm 2 LPbRL: Regret minimization (Unknown Model)

- 1: **input:** Learning rate $\eta_t > 0$, exploration length $t_0 > 0$
 - 2: Initialize empirical model $\widehat{\mathbb{P}}_1$.
 - 3: **for** $t = 1, 2, \dots, T$ **do**
 - 4: Compute \mathbf{w}_t^L and Π_t .
 - 5: Compute $(\pi_t^1, \pi_t^2) = \arg \max_{\pi^1, \pi^2 \in \Pi_t} \gamma_t \|\widehat{\phi}^{\widehat{\mathbb{P}}_t}(\pi^1) - \widehat{\phi}^{\widehat{\mathbb{P}}_t}(\pi^2)\|_{\widehat{\mathbf{V}}_t^{-1}} + 2\widehat{B}_t(\pi^1, 2WB, \delta) + 2\widehat{B}_t(\pi^2, 2WB, \delta)$
 - 6: Sample $\tau_t^1 \sim \pi_t^1$ and $\tau_t^2 \sim \pi_t^2$.
 - 7: Play the duel (τ_t^1, τ_t^2) and receive $o_t = \mathbb{1}(\tau_t^1 \text{ beats } \tau_t^2)$
 - 8: Update $\widehat{\mathbf{V}}_{t+1} = \widehat{\mathbf{V}}_t + \left(\widehat{\phi}^{\widehat{\mathbb{P}}_t}(\pi_\ell^1) - \widehat{\phi}^{\widehat{\mathbb{P}}_t}(\pi_\ell^2) \right) \left(\widehat{\phi}^{\widehat{\mathbb{P}}_t}(\pi_\ell^1) - \widehat{\phi}^{\widehat{\mathbb{P}}_t}(\pi_\ell^2) \right)^\top$
 - 9: Update empirical model and build $\widehat{\mathbb{P}}_{t+1}$.
 - 10: **end for**
-

Algorithm 2 shares the structure of Algorithm 1. The main difference lies in the definition of Π_t and in the optimization problem to find (π_t^1, π_t^2) . We can prove a result similar to Lemma 2 and show that $\pi^* \in \Pi_t$.

Lemma 4. *Let $\bar{\mathcal{E}}_{-1}$ be the event that $\pi^* \in \Pi_t$ for all $t \in \mathbb{N}$. Then $\mathbb{P}(\bar{\mathcal{E}}_{-1}) \geq 1 - 5\delta$.*

The proof of Lemma 4 can be found in Appendix C.3. The next step in the proof is to exhibit a bound on the instantaneous regret,

Lemma 5. *Let $\bar{\mathcal{E}}_2$ be the event that for all $t \in \mathbb{N}$,*

$$2r_t \leq \left(\widehat{\phi}^{\widehat{\mathbb{P}}_t}(\pi^*) - \widehat{\phi}^{\widehat{\mathbb{P}}_t}(\pi_t^1) \right)^\top \mathbf{w}^* + \left(\widehat{\phi}^{\widehat{\mathbb{P}}_t}(\pi^*) - \widehat{\phi}^{\widehat{\mathbb{P}}_t}(\pi_t^2) \right)^\top \mathbf{w}^* + \widehat{B}_t(\pi^*, 4WB, \delta) + \widehat{B}_t(\pi_t^1, 2WB, \delta) + \widehat{B}_t(\pi_t^2, 2WB, \delta).$$

Then $\mathbb{P}(\bar{\mathcal{E}}_2) \geq 1 - 2\delta$.

Proof. Note that we can write:

$$\begin{aligned} 2r_t &= (\phi(\pi^*) - \phi(\pi_t^1))^\top \mathbf{w}^* + (\phi(\pi^*) - \phi(\pi_t^2))^\top \mathbf{w}^* \\ &= (\widehat{\phi}^{\mathbb{P}_t}(\pi^*) - \widehat{\phi}^{\mathbb{P}_t}(\pi_t^1))^\top \mathbf{w}^* + (\widehat{\phi}^{\mathbb{P}_t}(\pi^*) - \widehat{\phi}^{\mathbb{P}_t}(\pi_t^2))^\top \mathbf{w}^* + \\ &\quad 2(\phi(\pi^*) - \widehat{\phi}^{\mathbb{P}_t}(\pi^*))^\top \mathbf{w}^* + \\ &\quad (\phi^{\mathbb{P}_t}(\pi_t^1) - \phi(\pi_t^1))^\top \mathbf{w}^* + (\phi^{\mathbb{P}_t}(\pi_t^2) - \phi(\pi_t^2))^\top \mathbf{w}^* \end{aligned}$$

By Lemma 12 (Lemma B.1 in Chatterji et al. (2021)), we conclude that with probability at least $1 - \delta$, for all $t \in \mathbb{N}$, setting $\eta = 4WB$,

$$2(\phi(\pi^*) - \widehat{\phi}^{\mathbb{P}_t}(\pi^*))^\top \mathbf{w}^* \leq \widehat{B}_t(\pi^*, 4WB, \delta)$$

Similarly, as a consequence of Lemma 12 and a union bound, setting $\eta = 2WB$, with probability at least $1 - 2\delta$

$$\begin{aligned} (\phi^{\mathbb{P}_t}(\pi_t^1) - \phi(\pi_t^1))^\top \mathbf{w}^* &\leq \widehat{B}_t(\pi_t^1, 2WB, \delta) \\ (\phi^{\mathbb{P}_t}(\pi_t^2) - \phi(\pi_t^2))^\top \mathbf{w}^* &\leq \widehat{B}_t(\pi_t^2, 2WB, \delta) \end{aligned}$$

The result follows. \square

Armed with the results of Lemma 4 we can show the following bound for the regret.

Lemma 6. *With probability at least $1 - 15\delta$ the regret is bounded by,*

$$\begin{aligned} R_T &\leq 2\gamma_T \sqrt{2Td \log\left(1 + \frac{TB}{d}\right)} + \\ &\quad \sum_{t \in [T]} 4\widehat{B}_t(\pi_t^1, 4WB, \delta) + 4\widehat{B}_t(\pi_t^2, 4WB, \delta) \end{aligned}$$

The proof of Lemma 6 can be found in Appendix C.4. The derivation follows from a repeated use of the instantaneous regret upper bound derived from Lemma 5.

The rest of the proof is dedicated to bound the $\widehat{B}_t(\cdot)$ terms. The general idea is to relate these bonus expectations under the empirical model with an expected sum of bonus terms under the true model and sampled according to policies π_t^1 and π_t^2 . Once this is achieved we have reduced the problem to bound a sum of vanishing markovian errors under the sampling distribution defined by the policies that were selected during optimization. This can be done via a similar argument as many existing RL works. Finally, we also show the γ_T term can be bounded by a term of the form $\widetilde{O}(\kappa\beta_t(\delta) + \alpha_{d,T}(\delta) + \text{poly}(H, |\mathcal{S}|, |\mathcal{A}|))$, hides logarithmic factors in $\delta, |\mathcal{S}|$ and $|\mathcal{A}|$. A detailed discussion of these arguments can be found in Appendix C. Our final main result (simplified) is thus,

Theorem 2. *The regret of LPbRL satisfies,*

$$R_T \leq \widetilde{O}(\kappa d\sqrt{T} + H^{3/2} \sqrt{|\mathcal{S}||\mathcal{A}|dT} + H|\mathcal{S}| \sqrt{|\mathcal{A}|dT}).$$

For all $T \in \mathbb{N}$ simultaneously with probability at least $1 - 15\delta$. Where \widetilde{O} hides logarithmic factors in $\delta, |\mathcal{S}|$ and $|\mathcal{A}|$.

The complete version of Theorem 2 can be found in Appendix C. Similar to Theorem 1, the leading term in the regret scales as $\widetilde{O}(d\sqrt{T})$ due to estimation based on the preferences. In addition to this, we now have dependence on $|\mathcal{S}|$ and $|\mathcal{A}|$ unlike before. These arise due to the tabular nature of the problem since the transition dynamics are unknown in this case.

5 DISCUSSIONS AND FUTURE SCOPES

In this work we addressed the problem of reinforcement learning from relative preference feedback where the agent does not get to see the absolute reward of actions taken at each state but instead observes the relative preferences between trajectories. We modeled the preference feedback in terms of the underlying non-Markovian linear reward model and proposed algorithms for both known as well as unknown MDP transition models. Precisely the regret guarantees of our proposed algorithms are analyzed to be respectively $\widetilde{O}(d \log(T/\delta)\sqrt{T})$ and $\widetilde{O}((\sqrt{d} + H^2 + |\mathcal{S}|)\sqrt{dT} + \sqrt{|\mathcal{S}||\mathcal{A}|TH})$ for the case of known and unknown transition models.

As discussed in the introduction, preference-based reinforcement learning has applications in several fields including training robots, stock market, recommender systems, two player games, chatbot interactions, etc. Thus there are plenty of scopes to extend the above setup to incorporate the corresponding system requirements, e.g. generalizing dueling trajectory preferences to subsets, considering alternative preference feedback without assuming an underlying reward model, extending to infinite horizon settings with more complex state-actions spaces, etc. Analyzing the fundamental performance limits of the PbRL regret minimization problem and designing algorithms with tighter performance guarantees would also be another interesting direction to investigate.

References

- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24:2312–2320, 2011.
- Nir Ailon, Zohar Shay Karnin, and Thorsten Joachims. Reducing dueling bandits to cardinal bandits. In *ICML*, volume 32, pages 856–864, 2014.

- 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.
- Mohammad Gheshlaghi Azar, Rémi Munos, and Hilbert J Kappen. Minimax pac bounds on the sample complexity of reinforcement learning with a generative model. *Machine learning*, 91(3):325–349, 2013.
- Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. *arXiv preprint arXiv:1703.05449*, 2017.
- Yu Bai and Chi Jin. Provable self-play algorithms for competitive reinforcement learning. In *International Conference on Machine Learning*, pages 551–560. PMLR, 2020.
- Yu Bai, Chi Jin, and Tiancheng Yu. Near-optimal reinforcement learning with self-play. In *Advances in Neural Information Processing Systems*, 2020.
- 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*, pages 335–342. Omnipress, 2008.
- Abdeslam Boularias, Jens Kober, and Jan Peters. Relative entropy inverse reinforcement learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pages 182–189, 2011.
- Róbert Busa-Fekete, Balázs Szörényi, Paul Weng, Weiwei Cheng, and Eyke Hüllermeier. Preference-based reinforcement learning: evolutionary direct policy search using a preference-based racing algorithm. *Machine Learning*, 97(3):327–351, 2014.
- Niladri S Chatterji, Aldo Pacchiano, Peter L Bartlett, and Michael I Jordan. On the theory of reinforcement learning with once-per-episode feedback. *arXiv preprint arXiv:2105.14363*, 2021.
- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems*, pages 4299–4307, 2017.
- Christoph Dann and Emma Brunskill. Sample complexity of episodic fixed-horizon reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 2818–2826, 2015.
- Christoph Dann, Tor Lattimore, and Emma Brunskill. Unifying pac and regret: Uniform pac bounds for episodic reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 5713–5723, 2017.
- Omar Darwiche Domingues, Pierre Ménard, Emilie Kaufmann, and Michal Valko. Episodic reinforcement learning in finite mdps: Minimax lower bounds revisited. *arXiv preprint arXiv:2010.03531*, 2020a.
- Omar Darwiche Domingues, Pierre Ménard, Matteo Pirodda, Emilie Kaufmann, and Michal Valko. Regret bounds for kernel-based reinforcement learning. *arXiv preprint arXiv:2004.05599*, 2020b.
- Yonathan Efroni, Nadav Merlis, and Shie Mannor. Reinforcement learning with trajectory feedback. *arXiv preprint arXiv:2008.06036*, 2020.
- Yonathan Efroni, Nadav Merlis, and Shie Mannor. Reinforcement learning with trajectory feedback. In *AAAI Conference on Artificial Intelligence, AAAI*, 2021.
- Louis Fauray, Marc Abeille, Clément Calauzènes, and Olivier Fercoq. Improved optimistic algorithms for logistic bandits. *arXiv preprint arXiv:2002.07530*, 2020.
- Pratik Gajane, Tanguy Urvoy, and Fabrice Clérot. A relative exponential weighing algorithm for adversarial utility-based dueling bandits. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 218–227, 2015.
- Suprovat Ghoshal and Aadirupa Saha. Exploiting correlation to achieve faster learning rates in low-rank preference bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 456–482. PMLR, 2022.
- Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *Advances in neural information processing systems*, pages 4565–4573, 2016.
- Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform chernoff bounds via nonnegative supermartingales. *Probability Surveys*, 17:257–317, 2020.
- Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-Uniform, Nonparametric, Nonasymptotic Confidence Sequences. *The Annals of Statistics*, 49(2):1055–1080, 2021.
- Borja Ibarz, Jan Leike, Tobias Pohlen, Geoffrey Irving, Shane Legg, and Dario Amodei. Reward learning from human preferences and demonstrations in atari. *Advances in neural information processing systems*, 31, 2018.
- Ashesh Jain, Brian Wojcik, Thorsten Joachims, and Ashutosh Saxena. Learning trajectory preferences for

- manipulators via iterative improvement. In *Advances in neural information processing systems*, pages 575–583, 2013.
- Junpei Komiyama, Junya Honda, Hisashi Kashima, and Hiroshi Nakagawa. Regret lower bound and optimal algorithm in dueling bandit problem. In *COLT*, pages 1141–1154, 2015.
- Andras Kupcsik, David Hsu, and Wee Sun Lee. Learning dynamic robot-to-human object handover from human feedback. In *Robotics research*, pages 161–176. Springer, 2018.
- Tor Lattimore and Marcus Hutter. Pac bounds for discounted mdps. In *International Conference on Algorithmic Learning Theory*, pages 320–334. Springer, 2012.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Lihong Li, Yu Lu, and Dengyong Zhou. Provably optimal algorithms for generalized linear contextual bandits. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2071–2080. JMLR. org, 2017.
- Qinghua Liu, Tiancheng Yu, Yu Bai, and Chi Jin. A sharp analysis of model-based reinforcement learning with self-play. In *International Conference on Machine Learning*, pages 7001–7010. PMLR, 2021.
- Gergely Neu and Csaba Szepesvári. Apprenticeship learning using inverse reinforcement learning and gradient methods. *arXiv preprint arXiv:1206.5264*, 2012.
- Andrew Y Ng, Stuart J Russell, et al. Algorithms for inverse reinforcement learning. In *Icml*, volume 1, page 2, 2000.
- Andrew Y Ng, Adam Coates, Mark Diel, Varun Ganapathi, Jamie Schulte, Ben Tse, Eric Berger, and Eric Liang. Autonomous inverted helicopter flight via reinforcement learning. In *Experimental robotics IX*, pages 363–372. Springer, 2006.
- Ellen R Novoseller, Yanan Sui, Yisong Yue, and Joel W Burdick. Dueling posterior sampling for preference-based reinforcement learning. *arXiv preprint arXiv:1908.01289*, 2019.
- Ronald Ortner. Regret bounds for reinforcement learning via markov chain concentration. *Journal of Artificial Intelligence Research*, 67:115–128, 2020.
- Aldo Pacchiano, Jack Parker-Holder, Yunhao Tang, Krzysztof Choromanski, Anna Choromanska, and Michael Jordan. Learning to score behaviors for guided policy optimization. In *International Conference on Machine Learning*, pages 7445–7454. PMLR, 2020.
- Jack Parker-Holder, Aldo Pacchiano, Krzysztof M Choromanski, and Stephen J Roberts. Effective diversity in population based reinforcement learning. *Advances in Neural Information Processing Systems*, 33:18050–18062, 2020.
- Akrour Riad, Marc Schoenauer, and Michèle Sebag. April: Active preference learning-based reinforcement learning. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD*. Springer Berlin Heidelberg, 2012.
- Dorsa Sadigh, Anca D Dragan, Shankar Sastry, and Sanjit A Seshia. Active preference-based learning of reward functions. In *Robotics: Science and Systems*, 2017.
- Aadirupa Saha. Optimal algorithms for stochastic contextual preference bandits. *Advances in Neural Information Processing Systems*, 34:30050–30062, 2021.
- Aadirupa Saha and Aditya Gopalan. Battle of bandits. In *Uncertainty in Artificial Intelligence*, 2018a.
- Aadirupa Saha and Aditya Gopalan. Active ranking with subset-wise preferences. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2018b.
- Aadirupa Saha and Aditya Gopalan. PAC Battling Bandits in the Plackett-Luce Model. In *Algorithmic Learning Theory*, pages 700–737, 2019.
- Aadirupa Saha and Akshay Krishnamurthy. Efficient and optimal algorithms for contextual dueling bandits under realizability. In *International Conference on Algorithmic Learning Theory*, pages 968–994. PMLR, 2022.
- Satinder Singh, Diane Litman, Michael Kearns, and Marilyn Walker. Optimizing dialogue management with reinforcement learning: Experiments with the nfun system. *Journal of Artificial Intelligence Research*, 16:105–133, 2002.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- Mohammad Sadegh Talebi and Odalric-Ambrym Maillard. Variance-aware regret bounds for undiscounted reinforcement learning in mdps. *arXiv preprint arXiv:1803.01626*, 2018.
- Christian Wirth and Johannes Fürnkranz. Preference-based reinforcement learning: A preliminary survey. In *Proceedings of the ECML/PKDD-13 Workshop on Reinforcement Learning from Generalized Feedback: Beyond Numeric Rewards*, 2013.
- Christian Wirth, Johannes Fürnkranz, Gerhard Neumann, et al. Model-free preference-based reinforcement learning. In *30th AAAI Conference on Artificial Intelligence, AAAI 2016*, pages 2222–2228, 2016.

- Christian Wirth, Riad Akrou, Gerhard Neumann, and Johannes Fürnkranz. A survey of preference-based reinforcement learning methods. *The Journal of Machine Learning Research*, 18(1):4945–4990, 2017.
- Markus Wulfmeier, Peter Ondruska, and Ingmar Posner. Maximum entropy deep inverse reinforcement learning. *arXiv preprint arXiv:1507.04888*, 2015.
- Qiaomin Xie, Yudong Chen, Zhaoran Wang, and Zhuoran Yang. Learning zero-sum simultaneous-move markov games using function approximation and correlated equilibrium. In *Conference on Learning Theory*, 2020.
- Yichong Xu, Ruosong Wang, Lin Yang, Aarti Singh, and Artur Dubrawski. Preference-based reinforcement learning with finite-time guarantees. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, pages 18784–18794. Curran Associates, Inc.
- Yisong Yue and Thorsten Joachims. Interactively optimizing information retrieval systems as a dueling bandits problem. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 1201–1208. ACM, 2009.
- Yisong Yue and Thorsten Joachims. Beat the mean bandit. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 241–248, 2011.
- Yisong Yue, Josef Broder, Robert Kleinberg, and Thorsten Joachims. The k -armed dueling bandits problem. *Journal of Computer and System Sciences*, 78(5):1538–1556, 2012.
- Andrea Zanette and Emma Brunskill. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. *arXiv preprint arXiv:1901.00210*, 2019.
- Zihan Zhang and Xiangyang Ji. Regret minimization for reinforcement learning by evaluating the optimal bias function. In *Advances in Neural Information Processing Systems*, pages 2827–2836, 2019.
- Masrour Zoghi, Shimon Whiteson, Remi Munos, Maarten de Rijke, et al. Relative upper confidence bound for the k -armed dueling bandit problem. In *JMLR Workshop and Conference Proceedings*, number 32, pages 10–18. JMLR, 2014a.
- Masrour Zoghi, Shimon A Whiteson, Maarten De Rijke, and Remi Munos. Relative confidence sampling for efficient on-line ranker evaluation. In *Proceedings of the 7th ACM international conference on Web search and data mining*, pages 73–82. ACM, 2014b.
- Masrour Zoghi, Zohar S Karnin, Shimon Whiteson, and Maarten De Rijke. Copeland dueling bandits. In *Advances in Neural Information Processing Systems*, pages 307–315, 2015.

Contents

1	INTRODUCTION	1
2	PROBLEM SETUP	3
3	KNOWN MODEL: ALGORITHM AND ANALYSIS	5
3.1	Maximum Likelihood Estimation	5
3.2	Algorithm and Analysis	5
3.3	Regret Analysis: Proof Sketch of Thm. 1	6
4	UNKNOWN MODEL: ALGORITHM AND ANALYSIS	6
4.1	Analysis of instantaneous regret:	6
5	DISCUSSIONS AND FUTURE SCOPES	8
A	Appendix for Section 2	12
A.1	Proof of Claim 2	13
B	Appendix for Section 3	13
B.1	Proof of Corollary 1	13
B.2	Proof of Lemma 2	15
B.3	Proof of Theorem 1	15
C	Appendix for Section 4	16
C.1	Supporting Related Work Lemmas	22
C.2	Proof of Lemma 3	22
C.3	Proof of Lemma 4	24
C.4	Proof of Lemma 6	25
D	Miscellaneous Technical Lemmas	26

**Supplementary for Dueling RL: Reinforcement Learning with
Trajectory Preferences**

A Appendix for Section 2

Claim 2. $\frac{R_T^{\text{scr}}}{2(e+1)} \leq R_T^{\text{pref}} \leq \frac{R_T^{\text{scr}}}{2}$.

A.1 Proof of Claim 2

Proof. Recall by Eq. (3), Eq. (4) and Claim 1.

$$R_T^{\text{scr}} = \sum_{t=1}^T \frac{2s(\pi^*) - (s(\pi_t^1) + s(\pi_t^2))}{2},$$

$$R_T^{\text{pref}} := \sum_{t=1}^T \frac{P(\pi^* \succ \pi_t^1) + P(\pi^* \succ \pi_t^2) - 1}{2}.$$

Now assume $S, B < 1$. Then for any two policies π_1 and $\pi_2 \in \Pi$, such that $s(\pi_1) \geq s(\pi_2)$, we have:

$$\begin{aligned} P(\pi_1, \pi_2) - 1/2 &= \frac{e^{s(\pi_1)} - e^{s(\pi_2)}}{2(e^{s(\pi_1)} + e^{s(\pi_2)})} \\ &= \frac{e^{s(\pi_1) - s(\pi_2)} - 1}{2(e^{s(\pi_1) - s(\pi_2)} + 1)} \\ &> \frac{s(\pi_1) - s(\pi_2)}{2(e + 1)} \end{aligned}$$

On the other hand denoting $x = s(\pi_1) - s(\pi_2) \in (0, 1)$ we get:

$$\begin{aligned} P(\pi_1, \pi_2) - 1/2 &= \frac{e^x - 1}{2(e^x + 1)} < \frac{(x + x^2/2! + x^3/3! + \dots)}{4} \\ &< \frac{x(1 + x/2 + x^2/2^2 + \dots)}{4} < x/2 \end{aligned}$$

The claim now follows combining the above two inequalities and noting that by definition $\pi^* := \arg \max_{\pi \in \Pi} s(\pi)$. \square

B Appendix for Section 3

By first-order optimality conditions, $\widehat{\mathbf{w}}_t^{\text{MLE}}$ is the point in \mathbb{R}^d satisfying:

$$\begin{aligned} \nabla_{\mathbf{w}} \mathcal{L}_t^\lambda(\widehat{\mathbf{w}}_t^{\text{MLE}}) &= \sum_{\ell=1}^{t-1} o_\ell (\phi(\tau_\ell^1) - \phi(\tau_\ell^2)) \\ &\quad - \left(\sum_{\ell=1}^{t-1} \sigma(\langle \phi(\tau_\ell^1) - \phi(\tau_\ell^2), \mathbf{w} \rangle (\phi(\tau_\ell^1) - \phi(\tau_\ell^2)) + \lambda \mathbf{w}) \right). \end{aligned}$$

B.1 Proof of Corollary 1

The primary mechanism behind Corollary 1 is the following lemma for matrix concentration.

Lemma 7. *Let $\delta \leq e^{-1}$. Then, with probability $1 - \delta \log_2 T$, for all $t \in [T]$, it holds that*

$$\|\mathbf{w}^* - \mathbf{w}_t^L\|_{\mathbf{V}_t}^2 \leq 2\|\mathbf{w}^* - \mathbf{w}_t^L\|_{\mathbf{V}_t}^2 + 84B^2 d \log(T(1 + 2T)/\delta) \|\mathbf{w}^* - \mathbf{w}_t^L\|_{\mathbf{V}_t}^2 \quad (8)$$

Proof. Fix $\mathbf{v} \in \mathbb{R}^d$ such that $\|\mathbf{v}\|_2 = 1$. For $\ell \in [T]$, let $X_\ell = \mathbf{v}^\top (\phi(\tau_\ell^1) - \phi(\tau_\ell^2)) (\phi(\tau_\ell^1) - \phi(\tau_\ell^2))^\top \mathbf{v}$. Furthermore define $X_0 = \kappa\lambda$. Observe that $X_\ell - \mathbf{E}_{\ell-1}X_\ell$ for $\ell = 0, \dots, T$ is an $\{\mathcal{F}_\ell\}$ -adapted martingale difference sequence where $\mathbf{E}_\ell[\cdot]$ denotes the conditional expectation $\mathbf{E}[\cdot | \mathcal{F}_\ell]$.

Note that the conditional variance of the individual terms may be bounded above by

$$\text{var}(X_\ell) = \mathbf{E}_{\ell-1} \left[X_\ell^2 - \mathbf{E}_{\ell-1} [X_\ell]^2 \right] \quad (9)$$

$$\leq \mathbf{E}_{\ell-1} [X_\ell^2] \quad (10)$$

$$\leq 4B^2 \mathbf{E}_{\ell-1} [X_\ell] \quad (11)$$

where we have used the fact that X_ℓ is non-negative and $\|\phi(\tau)\|_2 \leq B$. Let $\widehat{\mathbf{V}}_t = \kappa\lambda\mathbb{I} + \sum_{s \in [t]} \mathbf{E}_{s-1} (\phi(\tau_s^1) - \phi(\tau_s^2)) (\phi(\tau_s^1) - \phi(\tau_s^2))^\top$.

By (Bartlett et al., 2008, Lemma 2), we have, with probability at least $1 - \delta \log_2 T$,

$$\mathbf{v}^\top \widehat{\mathbf{V}}_T \mathbf{v} = \sum_{\ell=0}^T \mathbf{E}_{\ell-1} X_\ell \leq \sum_{\ell=0}^T X_\ell + \sqrt{16 \log(1/\delta) \sum_{\ell=0}^T \text{var}_{\ell-1}(X_\ell) + 2B^2 \log(1/\delta)} \quad (12)$$

$$\leq \sum_{\ell=0}^T X_\ell + \sqrt{64B^2 \log(1/\delta) \sum_{\ell=0}^T \mathbf{E}_{\ell-1} X_\ell + 2B^2 \log(1/\delta)} \quad (13)$$

$$\leq \sum_{\ell=0}^T X_\ell + \frac{1}{2} \sum_{\ell=0}^T \mathbf{E}_{\ell-1} X_\ell + 34B^2 \log(1/\delta) \quad (14)$$

$$= \mathbf{v}^\top \mathbf{V}_T \mathbf{v} + \frac{1}{2} \mathbf{v}^\top \widehat{\mathbf{V}}_T \mathbf{v} + 34B^2 \log(1/\delta) \quad (15)$$

where the third line applied the AM-GM inequality. Rearranging shows that

$$\frac{1}{2} \mathbf{v}^\top \widehat{\mathbf{V}}_T \mathbf{v} \leq \mathbf{v}^\top \mathbf{V}_T \mathbf{v} + 34B^2 \log(1/\delta) \quad (16)$$

This holds for a fixed \mathbf{v} . We now show that it approximately holds for all \mathbf{v} such that $\|\mathbf{v}\|_2 = 1$ via a covering argument.

Let \mathcal{C} be a minimal ϵ -cover of $\Pi^{d-1} = \{\mathbf{v} \in \mathbb{R}^d : \|\mathbf{v}\|_2 = 1\}$. A standard result states that $|\mathcal{C}_\epsilon| \leq (1 + 2/\epsilon)^d$. Then, by the union bound, with probability $1 - \delta \log_2 T$, for all $\mathbf{v} \in \mathcal{C}_\epsilon$,

$$\frac{1}{2} \mathbf{v}^\top \widehat{\mathbf{V}}_T \mathbf{v} \leq \mathbf{v}^\top \mathbf{V}_T \mathbf{v} + 34B^2 d \log((1 + 2/\epsilon)/\delta) \quad (17)$$

Let $A = \frac{1}{2} \mathbf{V}_T - \widehat{\mathbf{V}}_T$. Note that $\|A\| \leq 4B^2 T$ by definition. Let $\mathbf{v} \in \Pi^{d-1}$ be arbitrary and let $\mathbf{u}_v \in \mathcal{C}_\epsilon$ be the closest vector in the cover so that $\|\mathbf{v} - \mathbf{u}_v\|_2 \leq \epsilon$. Then,

$$\mathbf{v}^\top A \mathbf{v} = \mathbf{v}^\top A \mathbf{v} + \mathbf{u}^\top A \mathbf{u} - \mathbf{u}^\top A \mathbf{u} \quad (18)$$

$$\leq \mathbf{u}^\top A \mathbf{u} + 8\epsilon B^2 T \quad (19)$$

$$\leq 34B^2 d \log((1 + 2T)/\delta) + 8B^2 \quad (20)$$

$$\leq 42B^2 d \log((1 + 2T)/\delta) \quad (21)$$

under the good event and choosing $\epsilon = 1/T$. Since this holds for all $\mathbf{v} \in \Pi^{d-1}$, we conclude that

$$\widehat{\mathbf{V}}_T \preceq 2\mathbf{V}_T + 84B^2 d \log((1 + 2T)/\delta) \mathbb{I}_d \quad (22)$$

with probability at least $1 - \delta \log_2 T$. Finally, by Jensen's inequality we have $\overline{\mathbf{V}}_T \preceq \widehat{\mathbf{V}}_T$. Then, we apply the union bound over $t \in [T]$, which gives the result. \square

The proof of the corollary now follows immediately as a consequence.

Corollary 1. Under Assumption 1, conditioned on event $\mathcal{E}_\delta \cap \mathcal{E}_{prec}$, for any $t \in [T]$

$$\|\mathbf{w}^* - \mathbf{w}_t^L\|_{\bar{\mathbf{V}}_t} \leq 4\kappa\beta_t(\delta) + \alpha_{d,T}(\delta),$$

where $\alpha_{d,T}(\delta) = 20BW\sqrt{d\log(T(1+2T)/\delta)}$. Furthermore, if $\delta \leq 1/e$, then $\mathbb{P}(\mathcal{E}_\delta \cap \mathcal{E}_{prec}) \geq 1 - \delta - \delta \log_2 T$.

Proof of Corollary 1. Assuming that \mathcal{E}_δ holds, we have that $\|\mathbf{w}^* - \mathbf{w}_t^L\|_{\mathbf{V}_t} \leq 2\kappa\beta(\delta)$. Furthermore, Lemma 7 gives

$$\begin{aligned} \|\mathbf{w}^* - \mathbf{w}_t^L\|_{\bar{\mathbf{V}}_t} &\leq \sqrt{2}\|\mathbf{w}^* - \mathbf{w}_t^L\|_{\mathbf{V}_t} + 10B\sqrt{d\log(T(1+2T)/\delta)}\|\mathbf{w}^* - \mathbf{w}_t^L\|_2 \\ &\leq 4\kappa\beta_t(\delta) + 20BS\sqrt{d\log(T(1+2T)/\delta)} \end{aligned}$$

Also, $\mathbb{P}(\mathcal{E}_\delta \cap \mathcal{E}_2) \geq 1 - \delta - \delta \log_2 T$ follows from above combined with the claim of Lemma 1. \square

B.2 Proof of Lemma 2

Lemma 2. Conditioned on event $\mathcal{E}_\delta \cap \mathcal{E}_{prec}$, $\pi^* \in \Pi_t$,

Proof. Condition on $\mathcal{E}_\delta \cap \mathcal{E}_{prec}$. By definition of π^* , we have $(\phi(\pi^*) - \phi(\pi))^\top \mathbf{w}^* \geq 0$ for any arbitrary π . This implies

$$\begin{aligned} 0 &\leq (\phi(\pi^*) - \phi(\pi))^\top \mathbf{w}_t^L + \|\phi(\pi^*) - \phi(\pi)\|_{\bar{\mathbf{V}}_t^{-1}} \cdot \|\mathbf{w}^* - \mathbf{w}_t^L\|_{\bar{\mathbf{V}}_t} \\ &\leq (\phi(\pi^*) - \phi(\pi))^\top \mathbf{w}_t^L + (4\kappa\beta_t(\delta) + \alpha_{d,T}(\delta)) \cdot \|\phi(\pi^*) - \phi(\pi)\|_{\bar{\mathbf{V}}_t^{-1}} \end{aligned}$$

where the second line follows from Corollary 1. \square

B.3 Proof of Theorem 1

We require a standard determinant bound to complete the proof.

Lemma 8. Let $\lambda \geq B$. Consider the sequence $\mathbf{v}_1, \dots, \mathbf{v}_T \in \mathbb{R}^d$ such that $\|\mathbf{v}_i\| \leq B$ and define $V_t = \lambda I + \sum_{s \in [t-1]} \mathbf{v}_s \mathbf{v}_s^\top$. Then,

$$\sum_{t \in [T]} \|\mathbf{v}_t\|_{V_t^{-1}}^2 \leq 2d \log \left(1 + \frac{TB}{d} \right)$$

Proof. Since $\lambda \geq B$, we have that $\|\mathbf{v}_t\|_{V_t^{-1}} \leq 1$. Therefore, from Lemma 19.4 of Lattimore and Szepesvári (2020)

$$\begin{aligned} \sum_{t \in [T]} \|\mathbf{v}_t\|_{V_t^{-1}}^2 &\leq \sum_{t \in [T]} \log \left(1 + \|\mathbf{v}_t\|_{V_t^{-1}}^2 \right) \\ &\leq 2d \log \left(\frac{d\lambda + TB^2}{d\lambda} \right) \\ &= 2d \log \left(1 + \frac{TB}{d} \right) \end{aligned}$$

\square

Theorem 1. Let $\delta \leq 1/e$ and $\lambda \geq \frac{B}{\kappa}$. Then, with probability at least $1 - \delta$, the expected regret of Algorithm 1 can be bounded by

$$R_t \leq (4\kappa\beta_t(\delta) + 2\alpha_{d,T}(\delta)) \sqrt{2Td \log \left(1 + \frac{TB}{\kappa d} \right)}.$$

Proof of Theorem 1. Armed with the supporting results, we now focus on completing the proof of Theorem 1. The result may be shown by bounding the instantaneous regret. Condition on the event \mathcal{E}_δ . Then,

$$\begin{aligned}
 2r_t &:= (\phi(\pi^*) - \phi(\pi_t^1))^\top \mathbf{w}^* + (\phi(\pi^*) - \phi(\pi_t^2))^\top \mathbf{w}^* \\
 &= (\phi(\pi^*) - \phi(\pi_t^1))^\top \mathbf{w}_t^L + (\phi(\pi^*) - \phi(\pi_t^1))^\top (\mathbf{w}^* - \mathbf{w}_t^L) + (\phi(\pi^*) - \phi(\pi_t^2))^\top \mathbf{w}_t^L \\
 &\quad + (\phi(\pi^*) - \phi(\pi_t^2))^\top (\mathbf{w}^* - \mathbf{w}_t^L) \\
 &\leq (\phi(\pi^*) - \phi(\pi_t^1))^\top \mathbf{w}_t^L + (\phi(\pi^*) - \phi(\pi_t^2))^\top \mathbf{w}_t^L \\
 &\quad + \|\mathbf{w}^* - \mathbf{w}_t^L\|_{\bar{\mathbf{V}}_t} \cdot \|\phi(\pi^*) - \phi(\pi_t^1)\|_{\bar{\mathbf{V}}_t^{-1}} + \|\mathbf{w}^* - \mathbf{w}_t^L\|_{\bar{\mathbf{V}}_t} \cdot \|\phi(\pi^*) - \phi(\pi_t^2)\|_{\bar{\mathbf{V}}_t^{-1}}
 \end{aligned}$$

The last two terms in the above sum can be bounded using Corollary 1 as follows:

$$\begin{aligned}
 &\|\mathbf{w}^* - \mathbf{w}_t^L\|_{\bar{\mathbf{V}}_t} \cdot \|\phi(\pi^*) - \phi(\pi_t^1)\|_{\bar{\mathbf{V}}_t^{-1}} + \|\mathbf{w}^* - \mathbf{w}_t^L\|_{\bar{\mathbf{V}}_t} \cdot \|\phi(\pi^*) - \phi(\pi_t^2)\|_{\bar{\mathbf{V}}_t^{-1}} \\
 &\leq (2\kappa\beta_t(\delta) + \alpha_{T,d}(\delta)) \cdot \left(\|\phi(\pi^*) - \phi(\pi_t^1)\|_{\bar{\mathbf{V}}_t^{-1}} + \|\phi(\pi^*) - \phi(\pi_t^2)\|_{\bar{\mathbf{V}}_t^{-1}} \right)
 \end{aligned}$$

The first two terms leverage the optimistic bonus, using the fact that $\pi_t^1, \pi_t^2 \in \mathcal{S}_t$:

$$(\phi(\pi^*) - \phi(\pi_t^1))^\top \mathbf{w}_t^L + (\phi(\pi^*) - \phi(\pi_t^2))^\top \mathbf{w}_t^L \leq (2\kappa\beta_t(\delta) + \alpha_{T,d}(\delta)) \cdot \left(\|\phi(\pi^*) - \phi(\pi_t^1)\|_{\bar{\mathbf{V}}_t^{-1}} + \|\phi(\pi^*) - \phi(\pi_t^2)\|_{\bar{\mathbf{V}}_t^{-1}} \right)$$

In summary, we have that the instantaneous regret is upper bounded as

$$\begin{aligned}
 2r_t &\leq 2(2\kappa\beta_t(\delta) + \alpha_{T,d}(\delta)) \cdot \left(\|\phi(\pi^*) - \phi(\pi_t^1)\|_{\bar{\mathbf{V}}_t^{-1}} + \|\phi(\pi^*) - \phi(\pi_t^2)\|_{\bar{\mathbf{V}}_t^{-1}} \right) \\
 &\leq 4(2\kappa\beta_t(\delta) + \alpha_{T,d}(\delta)) \cdot \|\phi(\pi_t^1) - \phi(\pi_t^2)\|_{\bar{\mathbf{V}}_t^{-1}}
 \end{aligned}$$

where the last inequality follows from the fact that $\pi^* \in \mathcal{S}_t$ by Lemma 2 and since π_t^1 and π_t^2 were chosen the maximizer of the weighted difference $\|\phi(\pi_t^1) - \phi(\pi_t^2)\|_{\bar{\mathbf{V}}_t}$. The regret is therefore

$$\begin{aligned}
 R_T &= \sum_{t \in [T]} r_t \\
 &\leq 2(2\kappa\beta_T(\delta) + \alpha_{T,d}(\delta)) \cdot \sum_{t \in [T]} \|\phi(\pi_t^1) - \phi(\pi_t^2)\|_{\bar{\mathbf{V}}_t^{-1}} \\
 &\leq 2(2\kappa\beta_T(\delta) + \alpha_{T,d}(\delta)) \cdot \sqrt{T \sum_{t \in [T]} \|\phi(\pi_t^1) - \phi(\pi_t^2)\|_{\bar{\mathbf{V}}_t^{-1}}^2} \\
 &\leq 2(2\kappa\beta_T(\delta) + \alpha_{T,d}(\delta)) \cdot \sqrt{2Td \log \left(1 + \frac{TB}{d} \right)}
 \end{aligned}$$

where the second inequality follows from Cauchy-Schwarz and the last inequality applies Lemma 8. \square

C Appendix for Section 4

In this section we will use the notation $N_t(s, a)$ to denote the number of times action a was executed at state s up to time $t - 1$. Recall the bonus terms,

Given any $\eta > 0$ define,

$$\begin{aligned}
 \xi_{s,a}^{(t)}(\eta, \delta) &= \min \left(2\eta, 4\eta \sqrt{\frac{U}{N_t(s, a)}} \right) \\
 \text{s.t. } U &= H \log(|\mathcal{S}||\mathcal{A}|) + \log \left(\frac{6 \log(N_t(s, a))}{\delta} \right).
 \end{aligned}$$

and the empirical average of $\xi_{s,a}^{(t)}(\eta, \delta)$ bonuses,

$$\widehat{B}_t(\pi, \eta, \delta) = \mathbb{E}_{s_1 \sim \rho, \tau \sim \widehat{\mathbb{P}}_t^\pi(\cdot | s_1)} \left[\sum_{h=1}^{H-1} \xi_{s_h, a_h}^{(t)}(\eta, \delta) \right].$$

Additionally we also define the error terms

$$\begin{aligned} \xi_{s,a}^{(t)}(\epsilon, \eta, \delta) &= \min \left(2\eta, 4\eta \sqrt{\frac{U}{N_t(s, a)}} \right) \\ \text{s.t. } U &= H \log(|\mathcal{S}||\mathcal{A}|H) + |\mathcal{S}| \log \left(\left\lceil \frac{4\eta H}{\epsilon} \right\rceil \right) + \log \left(\frac{6 \log(N_t(s, a))}{\delta} \right). \end{aligned}$$

In contrast with the definition of bonus $\xi_{s,a}^{(t)}(\eta, \delta)$ this quantity depends on an extra parameter ϵ . These error terms induce the the following ‘bonus’ function,

$$B_t(\pi, \eta, \delta, \epsilon) = \mathbb{E}_{s_1 \sim \rho, \tau \sim \mathbb{P}^\pi(\cdot | s_1)} \left[\sum_{h=1}^{H-1} \xi_{s_h, a_h}^{(t)}(\epsilon, \eta, \delta) \right].$$

Here the expectation is under the true MDP dynamics.

Once we have established the validity of Lemma 6, and therefore that with probability at least $1 - 15\delta$,

$$R_T \leq 2\gamma_T \sqrt{2Td \log \left(1 + \frac{TB}{d} \right)} + \sum_{t \in [T]} 4\widehat{B}_t(\pi_t^1, 4WB, \delta) + 4\widehat{B}_t(\pi_t^2, 4WB, \delta)$$

it remains to show the $\widehat{B}_t(\cdot)$ terms are small. We’ll do so by showing that for any $\eta > 0$ and $\delta \in (0, 1)$ and for all policies π simultaneously we can bound the empirical expected bonuses $\widehat{B}_t(\pi, \eta, \delta)$ in terms of the population quantities $B_t(\pi, \eta H, \delta)$,

Lemma 9. *Let $\eta, \epsilon > 0$. For all π simultaneously and for all $t \in \mathbb{N}$, with probability $1 - \delta$,*

$$\widehat{B}_t(\pi, \eta, \delta) \leq 2B_t(\pi, 2H\eta, \delta, \epsilon) + \epsilon$$

Proof. Recall that,

$$\widehat{B}_t(\pi, \eta, \delta) = \mathbb{E}_{s_1 \sim \rho, \tau \sim \widehat{\mathbb{P}}_t^\pi(\cdot | s_1)} \left[\sum_{h=1}^{H-1} \xi_{s_h, a_h}^{(t)}(\eta, \delta) \right].$$

Let $f : \Gamma \rightarrow \mathbb{R}$ be defined as,

$$f(\tau) = \sum_{h=1}^{H-1} \xi_{s_h, a_h}^{(t)}(\eta).$$

It is easy to see that $f(\tau) \in (0, 2\eta H]$ for all $\tau \in \Gamma$. Therefore, a direct application of Lemma 13 implies that with probability at least $1 - \delta$ and simultaneously for all π , and $t \in \mathbb{N}$,

$$\widehat{B}_t(\pi, \eta, \delta) \leq \mathbb{E}_{s_1 \sim \rho, \tau \sim \mathbb{P}^\pi(\cdot | s_1)} \left[\sum_{h=1}^{H-1} \xi_{s_h, a_h}^{(t)}(\eta, \delta) \right] + B_t(\pi, 2H\eta, \delta, \epsilon) + \epsilon$$

Since $\xi_{s,a}^{(t)}(\epsilon, \eta, \delta) \geq \xi_{s,a}^{(t)}(\eta, \delta)$ for all $\epsilon > 0$, $s, a \in \mathcal{S} \times \mathcal{A}$ and $\xi_{s,a}^{(t)}(\epsilon, \eta, \delta)$ is monotonic in η we conclude that,

$$\begin{aligned} \mathbb{E}_{s_1 \sim \rho, \tau \sim \mathbb{P}^\pi(\cdot | s_1)} \left[\sum_{h=1}^{H-1} \xi_{s_h, a_h}^{(t)}(\eta, \delta) \right] &\leq \mathbb{E}_{s_1 \sim \rho, \tau \sim \mathbb{P}^\pi(\cdot | s_1)} \left[\sum_{h=1}^{H-1} \xi_{s_h, a_h}^{(t)}(\epsilon, \eta, \delta) \right] \\ &\leq \mathbb{E}_{s_1 \sim \rho, \tau \sim \mathbb{P}^\pi(\cdot | s_1)} \left[\sum_{h=1}^{H-1} \xi_{s_h, a_h}^{(t)}(\epsilon, 2H\eta, \delta) \right] \\ &= B_t(\pi, 2H\eta, \delta, \epsilon) \end{aligned}$$

Combining these inequalities the result follows. \square

Let $\bar{\mathcal{E}}_3$ such that for all $T \in \mathbb{N}$ be the event that,

$$\sum_{t \in [T]} 4\widehat{B}_t(\pi_t^1, 4SB, \delta) + 4\widehat{B}_t(\pi_t^2, 4SB, \delta) \leq \epsilon T + \sum_{t \in [T]} 8B_t(\pi_t^1, 8HSB, \delta, \epsilon) + 8B_t(\pi_t^2, 8HSB, \delta, \epsilon)$$

Invoking Lemmas 6 and 9 we can show $\bar{\mathcal{E}}_3$ occurs with probability at least $1 - 2\delta$. Let's bound the sum $\sum_{t \in [T]} 8B_t(\pi_t^1, 8HSB, \delta, \epsilon) + 8B_t(\pi_t^2, 8HSB, \delta, \epsilon)$. Consider the martingale difference sequences $\{B_t(\pi_t^1, 8HSB, \delta, \epsilon) - \sum_{h=1}^{H-1} \xi_{s_{t,h}^1, a_{t,h}^1}(\epsilon, 8HSB, \delta)\}_{t=1}^\infty$ and $\{B_t(\pi_t^2, 8HSB, \delta, \epsilon) - \sum_{h=1}^{H-1} \xi_{s_{t,h}^2, a_{t,h}^2}(\epsilon, 8HSB, \delta)\}_{t=1}^\infty$ each with norm upper bound $32H^2SB$. By an anytime Hoeffding inequality (see Lemma 16) (since $\xi_{s,a}(\epsilon, \eta, \delta) \leq 2\eta$ and therefore $\sum_h \xi_{s_h, a_h}(\epsilon, \eta, \delta) \leq 2H\eta$) applied to with probability at least $1 - 2\delta$ for all $T \in \mathbb{N}$ simultaneously

$$\sum_{t \in [T]} 8B_t(\pi_t^1, 4SB, \delta, \epsilon) + 8B_t(\pi_t^2, 4SB, \delta, \epsilon) \leq 8 \sum_{t \in [T]} \left(\sum_{h=1}^{H-1} \xi_{s_{t,h}^1, a_{t,h}^1}(\epsilon, 8HSB, \delta) + \sum_{h=1}^{H-1} \xi_{s_{t,h}^2, a_{t,h}^2}(\epsilon, 8HSB, \delta) \right) + \mathbf{I}$$

Where $\mathbf{I} = 128HSB \sqrt{TH \log \left(\frac{6 \log(TH)}{\delta} \right)}$. In order to bound the remaining empirical error terms, we to the following standard result,

Lemma 10. For $i \in \{1, 2\}$ the empirical sum of errors satisfies the following bound

$$\begin{aligned} & \sum_{t \in [T]} \sum_{h=1}^{H-1} \xi_{s_{t,h}^i, a_{t,h}^i}(\epsilon, 8HSB, \delta) \leq \\ & 64HSB \sqrt{\left(H \log(|\mathcal{S}||\mathcal{A}|H) + |\mathcal{S}| \log \left(\left\lceil \frac{32H^2SB}{\epsilon} \right\rceil \right) + \log \left(\frac{6 \log(HTs)}{\delta} \right) \right) |\mathcal{S}||\mathcal{A}|TH}. \end{aligned}$$

Proof. Let's rewrite this sum by instead summing over states and actions,

$$\begin{aligned} & \sum_{t \in [T]} \sum_{h=1}^{H-1} \xi_{s_{t,h}^i, a_{t,h}^i}(\epsilon, 8HSB, \delta) = \\ & \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{t=1}^{N_{T+1}(s,a)} \min \left(16HSB, 32HSB \sqrt{\frac{H \log(|\mathcal{S}||\mathcal{A}|H) + |\mathcal{S}| \log \left(\left\lceil \frac{32H^2SB}{\epsilon} \right\rceil \right) + \log \left(\frac{6 \log(t)}{\delta} \right)}{t}} \right) \\ & \leq 32HSB \sqrt{H \log(|\mathcal{S}||\mathcal{A}|H) + |\mathcal{S}| \log \left(\left\lceil \frac{32H^2SB}{\epsilon} \right\rceil \right) + \log \left(\frac{6 \log(HT)}{\delta} \right)} \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{t=1}^{N_{T+1}(s,a)} \frac{1}{\sqrt{t}} \\ & \leq 32HSB \sqrt{H \log(|\mathcal{S}||\mathcal{A}|H) + |\mathcal{S}| \log \left(\left\lceil \frac{32H^2SB}{\epsilon} \right\rceil \right) + \log \left(\frac{6 \log(HT)}{\delta} \right)} \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} 2\sqrt{N_{T+1}(s,a)} \\ & \leq 64HSB \sqrt{\left(H \log(|\mathcal{S}||\mathcal{A}|H) + |\mathcal{S}| \log \left(\left\lceil \frac{32H^2SB}{\epsilon} \right\rceil \right) + \log \left(\frac{6 \log(HT)}{\delta} \right) \right) |\mathcal{S}||\mathcal{A}|TH}. \end{aligned}$$

The result follows. \square

We will use Lemma 10 with $\epsilon = 1/T$. As a consequence of Lemma 10 and Lemma 6 we see that when $\mathcal{E}_\delta \cap \mathcal{E}_2 \cap \bar{\mathcal{E}}_0 \cap \bar{\mathcal{E}}_{-1} \cap \bar{\mathcal{E}}_2 \cap \bar{\mathcal{E}}_3$ holds

$$R_T \leq 2\gamma_T \sqrt{2Td \log \left(1 + \frac{TB}{d}\right)} + \tilde{\mathcal{O}} \left(H^{3/2} \sqrt{|\mathcal{A}||\mathcal{S}|TH} + H|\mathcal{S}| \sqrt{|\mathcal{A}|TH} + H\sqrt{TH} \right).$$

Now it remains to bound term γ_T . From now on let's set $\epsilon = \min(1/T, 8HSB)$ and let $\bar{\mathcal{E}}_4$ be the event that for all $t \in \mathbb{N}$ and all $i \in \{1, 2\}$,

$$\widehat{B}_t \left(\pi_\ell^i, 2SB, \frac{\delta'}{8\ell^3 \mathcal{A} \mathcal{S}} \right) \leq 2B_t \left(\pi_\ell^i, 4HSB, \frac{\delta'}{8\ell^3 \mathcal{A} \mathcal{S}}, \epsilon \right) + \epsilon$$

for all $t \in \mathbb{N}$ and all $i \in \{1, 2\}$. As a consequence of Lemma 9 we can bound $\mathbb{P}(\bar{\mathcal{E}}_4) \geq 1 - 2\delta$. Squaring both sides,

$$\begin{aligned} \left(\widehat{B}_t(\pi_\ell^i, 2SB, \delta'_\ell) \right)^2 &\leq \left(2B_t(\pi_\ell^i, 4HSB, \delta'_\ell, \epsilon) + \epsilon \right)^2 \\ &\stackrel{(i)}{\leq} 4 \left(B_t(\pi_\ell^i, 4HSB, \delta'_\ell, \epsilon) \right)^2 + 16\epsilon HSB + \epsilon^2 \\ &\leq 4 \left(B_t(\pi_\ell^i, 4HSB, \delta'_\ell, \epsilon) \right)^2 + 24\epsilon HSB \end{aligned}$$

Where $\delta'_\ell = \frac{\delta'}{8\ell^3 \mathcal{A} \mathcal{S}}$. Inequality (i) used that $B(\pi, \eta, \delta, \epsilon) \leq 2H\eta$. The last inequality holds because $\epsilon \leq 8HSB$. Therefore if $\bar{\mathcal{E}}_4$ holds for all $t \in \mathbb{N}$,

$$\gamma_t \leq \sqrt{2} (4\kappa\beta_t(\delta) + \alpha_{d,T}(\delta)) + \frac{1}{t} + 4\sqrt{\sum_{\ell=1}^{t-1} B_t^2(\pi_\ell^1, 4HSB, \delta'_\ell, \epsilon) + B_t^2(\pi_\ell^2, 4HSB, \delta'_\ell, \epsilon) + \mathbf{I}}.$$

Where $\mathbf{I} = 96(t-1)\epsilon HSB$. We are just left with bounding the sum of squares $\sum_{\ell=1}^{t-1} \left(B_t(\pi_\ell^1, 4HSB, \frac{\delta'}{8\ell^3 \mathcal{A} \mathcal{S}}, \epsilon) \right)^2 + \left(B_t(\pi_\ell^2, 4HSB, \frac{\delta'}{8\ell^3 \mathcal{A} \mathcal{S}}, \epsilon) \right)^2$.

Lemma 11. Let $\eta, \epsilon > 0$ and $\delta, \delta' \in (0, 1)$ and define $\bar{\mathcal{E}}_5(\delta')$ be the event that for all $t \in \mathbb{N}$ and $i \in \{1, 2\}$

$$\begin{aligned} \sum_{i \in \{1, 2\}} \sum_{\ell=1}^{t-1} (B_\ell(\pi_\ell^i, \eta, \delta/\ell^3))^2 &\leq 12\eta^2 H^2 \left(1.4 \ln \ln (2 (\max(4\eta^2 Ht, 1))) + \ln \frac{5.2}{\delta'} + 1 \right) + \\ &64\eta^2 H \left(H \log(|\mathcal{S}||\mathcal{A}|H) + |\mathcal{S}| \log \left(\left\lceil \frac{4\eta H}{\epsilon} \right\rceil \right) + \log \left(\frac{6 \log(HT)}{\delta} \right) \right) |\mathcal{S}||\mathcal{A}| \log(TH + |\mathcal{S}||\mathcal{A}|) \end{aligned}$$

Then $\mathbb{P}(\bar{\mathcal{E}}_5(\delta')) \geq 1 - 2\delta'$.

Proof. Observe that,

$$\begin{aligned} \left(B_\ell(\pi_\ell^1, \eta, \frac{\delta}{\ell^3}, \epsilon) \right)^2 + \left(B_\ell(\pi_\ell^2, \eta, \frac{\delta}{\ell^3}, \epsilon) \right)^2 &= \left(\mathbb{E}_{s_1^1 \sim \rho, \tau \sim \mathbb{P}^{\pi_\ell^1}(\cdot | s_1^1)} \left[\sum_{h=1}^{H-1} \xi_{s_h^1, a_h^1}^{(t)}(\epsilon, \eta, \frac{\delta}{\ell^3}) \right] \right)^2 + \\ &\left(\mathbb{E}_{s_1^2 \sim \rho, \tau \sim \mathbb{P}^{\pi_\ell^2}(\cdot | s_1^2)} \left[\sum_{h=1}^{H-1} \xi_{s_h^2, a_h^2}^{(t)}(\epsilon, \eta, \frac{\delta}{\ell^3}) \right] \right)^2 \\ &\stackrel{(i)}{\leq} H \mathbb{E}_{s_1^1 \sim \rho, \tau \sim \mathbb{P}^{\pi_\ell^1}(\cdot | s_1^1)} \left[\sum_{h=1}^{H-1} \left(\xi_{s_h^1, a_h^1}^{(t)}(\epsilon, \eta, \frac{\delta}{\ell^3}) \right)^2 \right] + \\ &H \mathbb{E}_{s_1^2 \sim \rho, \tau \sim \mathbb{P}^{\pi_\ell^2}(\cdot | s_1^2)} \left[\sum_{h=1}^{H-1} \left(\xi_{s_h^2, a_h^2}^{(t)}(\epsilon, \eta, \frac{\delta}{\ell^3}) \right)^2 \right] \end{aligned}$$

Where inequality (i) is a consequence of $\left(\mathbb{E} \left[\sum_{h=1}^H a_h \right] \right)^2 \leq H \mathbb{E} \left[\sum_{h=1}^H a_h^2 \right]$. Define the martingale-difference sequences for $i \in \{1, 2\}$,

$$D_\ell^{(i)} = \mathbb{E}_{s_1^i \sim \rho, \tau \sim \mathbb{P}^{\pi_\ell^i}(\cdot | s_1^i)} \left[\sum_{h=1}^{H-1} \left(\xi_{s_h^i, a_h^i}^{(\ell)}(\epsilon, \eta, \delta) \right)^2 \right] - \sum_{h=1}^{H-1} \left(\xi_{s_h^i, a_h^i}^{(\ell)}(\epsilon, \eta, \delta) \right)^2$$

Since $\xi_{s,a}^{(\ell)}(\epsilon, \eta, \delta) \leq 2\eta$, we see that $|D_\ell^{(i)}| \leq 8\eta^2 H$. Observe that for $i \in \{1, 2\}$,

$$\begin{aligned} \text{Var}_\ell^{(i)} \left(\sum_{h=1}^{H-1} \left(\xi_{s_h^i, a_h^i}^{(\ell)}(\epsilon, \eta, \delta) \right)^2 \right) &\leq \mathbb{E}_{s_1^i \sim \rho, \tau \sim \mathbb{P}^{\pi_\ell^i(\cdot|s_1^i)}} \left[\left\{ \sum_{h=1}^{H-1} \left(\xi_{s_h^i, a_h^i}^{(\ell)}(\epsilon, \eta, \delta) \right)^2 \right\}^2 \right] \\ &\stackrel{(i)}{\leq} 4\eta^2 H \mathbb{E}_{s_1^i \sim \rho, \tau \sim \mathbb{P}^{\pi_\ell^i(\cdot|s_1^i)}} \left[\sum_{h=1}^{H-1} \left(\xi_{s_h^i, a_h^i}^{(\ell)}(\epsilon, \eta, \delta) \right)^2 \right] \\ &\leq 16\eta^4 H^2 \end{aligned}$$

Where (i) follows because for $\xi_{s,a}^{(\ell)}(\epsilon, \eta, \delta) \leq 2\eta$.

Since the variance can be bounded by the mean, we can make use of a Uniform Empirical Bernstein Bound from Lemma 17. Let $S_t^{(i)} = \sum_{\ell=1}^t D_\ell^{(i)}$ for $i \in \{1, 2\}$ and $W_t^{(i)} = \sum_{\ell=1}^t \text{Var}_\ell^{(i)} \left(\sum_{h=1}^{H-1} \left(\xi_{s_h, a_h}^{(\ell, i)}(\epsilon, \eta, \delta) \right)^2 \right)$. Let $c = 8\eta^2 H$ and $m = 4\eta^2 H$. With probability $1 - \delta'$ for all $t \in \mathbb{N}$,

$$\begin{aligned} \sum_{\ell=1}^{t-1} D_\ell^{(i)} &\leq \sqrt{\max \left(4\eta^2 H \sum_{\ell=1}^{t-1} \mathbb{E}_{s_1 \sim \rho, \tau \sim \mathbb{P}^{\pi_\ell^i(\cdot|s_1)}} \left[\sum_{h=1}^{H-1} \left(\xi_{s_h, a_h}^{(\ell, i)}(\epsilon, \eta, \delta) \right)^2 \right], 4\eta^2 H \right) \left(1.4 \ln \ln (2 (\max (4\eta^2 H t, 1))) + \ln \frac{5.2}{\delta'} \right)} \\ &\quad + 3.28\eta^2 H \eta \left(1.4 \ln \ln (2 (\max (4\eta^2 H t, 1))) + \ln \frac{5.2}{\delta'} \right) \\ &\leq \sqrt{\left(4\eta^2 H \sum_{\ell=1}^{t-1} \mathbb{E}_{s_1 \sim \rho, \tau \sim \mathbb{P}^{\pi_\ell^i(\cdot|s_1)}} \left[\sum_{h=1}^{H-1} \left(\xi_{s_h, a_h}^{(\ell, i)}(\epsilon, \eta, \delta) \right)^2 \right] + 4\eta^2 H \right) \left(1.4 \ln \ln (2 (\max (4\eta^2 H t, 1))) + \ln \frac{5.2}{\delta'} \right)} \\ &\quad + 3.28\eta^2 H \left(1.4 \ln \ln (2 (\max (4\eta^2 H t, 1))) + \ln \frac{5.2}{\delta'} \right) \end{aligned}$$

Since $\sqrt{ab} \leq \frac{a+b}{2}$,

$$\begin{aligned} \sum_{\ell=1}^{t-1} D_\ell^{(i)} &\leq \frac{1}{2} \mathbb{E}_{s_1 \sim \rho, \tau \sim \mathbb{P}^{\pi_\ell^i(\cdot|s_1)}} \left[\sum_{h=1}^{H-1} \left(\xi_{s_h, a_h}^{(\ell, i)}(\epsilon, \eta, \delta) \right)^2 \right] + \\ &\quad 2\eta^2 H + \underbrace{(3.28\eta^2 H + 2\eta^2 H)}_{\leq 6\eta^2 H} \left(1.4 \ln \ln (2 (\max (4\eta^2 H t, 1))) + \ln \frac{5.2}{\delta'} \right). \end{aligned}$$

Therefore with high probability for $i \in \{1, 2\}$,

$$\begin{aligned} \mathbb{E}_{s_1 \sim \rho, \tau \sim \mathbb{P}^{\pi_\ell^i(\cdot|s_1)}} \left[\sum_{h=1}^{H-1} \left(\xi_{s_h, a_h}^{(\ell, i)}(\epsilon, \eta, \delta) \right)^2 \right] &\leq 2 \sum_{\ell=1}^{t-1} \sum_{h=1}^{H-1} \left(\xi_{s_h, a_h}^{(\ell, i)}(\epsilon, \eta, \delta) \right)^2 + 4\eta^2 H + \\ &\quad 6\eta^2 H \left(1.4 \ln \ln (2 (\max (4\eta^2 H t, 1))) + \ln \frac{5.2}{\delta'} \right). \end{aligned}$$

Therefore with probability $1 - 2\delta'$,

$$\begin{aligned} \sum_{i \in \{1,2\}} \sum_{\ell=1}^{t-1} (B_\ell(\pi_\ell^i, \eta, \delta/\ell^3))^2 &\leq 2H \sum_{i \in \{1,2\}} \sum_{\ell=1}^{t-1} \sum_{h=1}^{H-1} \left(\xi_{s_h, a_h}^{(\ell, i)}(\epsilon, \eta, \delta) \right)^2 + \\ &12\eta^2 H^2 \left(1.4 \ln \ln (2 (\max (4\eta^2 Ht, 1))) + \ln \frac{5.2}{\delta'} + 1 \right). \end{aligned}$$

We are left with the task of bounding the terms $\sum_{\ell=1}^{t-1} \sum_{h=1}^{H-1} \left(\xi_{s_h, a_h}^{(\ell, i)}(\epsilon, \eta, \delta) \right)^2$.

Let's rewrite this sum by instead summing over states and actions,

$$\begin{aligned} \sum_{t \in [T]} \sum_{h=1}^{H-1} \left(\xi_{s_h^i, a_h^i}^{(t)}(\epsilon, \eta, \delta) \right)^2 &= \\ \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{t=1}^{N_{T+1}(s, a)} \min \left(4\eta^2, 16\eta^2 \frac{H \log(|\mathcal{S}||\mathcal{A}|H) + |\mathcal{S}| \log \left(\left\lceil \frac{4\eta H}{\epsilon} \right\rceil \right) + \log \left(\frac{6 \log(t)}{\delta} \right)}{t} \right) & \\ = \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{t=1}^{N_{T+1}(s, a)} 16\eta^2 \frac{H \log(|\mathcal{S}||\mathcal{A}|H) + |\mathcal{S}| \log \left(\left\lceil \frac{4\eta H}{\epsilon} \right\rceil \right) + \log \left(\frac{6 \log(t)}{\delta} \right)}{t} & \\ = 16\eta^2 \left(H \log(|\mathcal{S}||\mathcal{A}|H) + |\mathcal{S}| \log \left(\left\lceil \frac{4\eta H}{\epsilon} \right\rceil \right) + \log \left(\frac{6 \log(HT)}{\delta'} \right) \right) \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \sum_{t=1}^{N_{T+1}(s, a)} \frac{1}{t} & \\ 32\eta^2 \left(H \log(|\mathcal{S}||\mathcal{A}|H) + |\mathcal{S}| \log \left(\left\lceil \frac{4\eta H}{\epsilon} \right\rceil \right) + \log \left(\frac{6 \log(HT)}{\delta'} \right) \right) \sum_{s \in \mathcal{S}} \sum_{a \in \mathcal{A}} \log(N_{T+1}(s, a) + 1) & \\ \leq 32\eta^2 \left(H \log(|\mathcal{S}||\mathcal{A}|H) + |\mathcal{S}| \log \left(\left\lceil \frac{4\eta H}{\epsilon} \right\rceil \right) + \log \left(\frac{6 \log(HT)}{\delta'} \right) \right) |\mathcal{S}||\mathcal{A}| \log(TH + |\mathcal{S}||\mathcal{A}|) & \end{aligned}$$

Therefore with probability $1 - 2\delta'$,

$$\begin{aligned} \sum_{i \in \{1,2\}} \sum_{\ell=1}^{t-1} (B_\ell(\pi_\ell^i, \eta, \delta/\ell^3))^2 &\leq 12\eta^2 H^2 \left(1.4 \ln \ln (2 (\max (4\eta^2 Ht, 1))) + \ln \frac{5.2}{\delta'} + 1 \right) + \\ 64\eta^2 H \left(H \log(|\mathcal{S}||\mathcal{A}|H) + |\mathcal{S}| \log \left(\left\lceil \frac{4\eta H}{\epsilon} \right\rceil \right) + \log \left(\frac{6 \log(HT)}{\delta'} \right) \right) |\mathcal{S}||\mathcal{A}| \log(TH + |\mathcal{S}||\mathcal{A}|) & \end{aligned}$$

The result follows. \square

The main takeaway from this lemma is that the sum of the square errors grows only logarithmically in T . Applying this bound to γ_T and setting $\epsilon = O(1/T)$ we obtain,

$$\gamma_T \leq \sqrt{2} (4\kappa\beta_T(\delta) + \alpha_{d,T}(\delta)) + 2\sqrt{\omega_T(\delta)} + \frac{1}{T} = \tilde{\mathcal{O}} \left(\kappa\sqrt{d} + H^2\sqrt{|\mathcal{S}||\mathcal{A}|} + H^{3/2}|\mathcal{S}|\sqrt{|\mathcal{A}|} \right)$$

Where

$$\begin{aligned} \omega_T(\delta) &= 192H^4 S62B^2 \left(1.4 \ln \ln (2 (\max (64H^3 S^2 B^2 t, 1))) + \ln \frac{5.2}{\delta'} + 1 \right) + \\ 1024H^3 S^2 B^2 \left(H \log(|\mathcal{S}||\mathcal{A}|H) + |\mathcal{S}| \log \left(\left\lceil \frac{64H^3 S^2 B^2}{\epsilon} \right\rceil \right) + \log \left(\frac{6 \log(HT)}{\delta'} \right) \right) |\mathcal{S}||\mathcal{A}| \log(TH + |\mathcal{S}||\mathcal{A}|) & \end{aligned}$$

Applying this bound to γ_T and setting $\epsilon = 1/96T$ we obtain,

$$\gamma_T \leq \sqrt{2} (4\kappa\beta_T(\delta) + \alpha_{d,T}(\delta)) + 2\sqrt{\omega_T(\delta) + HSB} + \frac{1}{T}$$

Combining these observations we can derive our main result,

Theorem 3 (Formal version of Theorem 2). *If $\mathcal{E}_\delta \cap \mathcal{E}_2 \cap \bar{\mathcal{E}}_0 \cap \bar{\mathcal{E}}_{-1} \cap \bar{\mathcal{E}}_2 \cap \bar{\mathcal{E}}_3 \cap \bar{\mathcal{E}}_4 \cap \bar{\mathcal{E}}_5(\frac{\delta}{2})$ holds then the regret of LPbRL satisfies,*

$$\begin{aligned} R_T \leq & 2 \left(4\kappa\beta_t(\delta) + \alpha_{d,T}(\delta) + 2\sqrt{\omega_T(\delta) + 96(t-1)\epsilon HSB} + \frac{1}{T} \right) \sqrt{2Td \log \left(1 + \frac{TB}{d} \right)} + \\ & 128H^{3/2}SB \sqrt{TH \log \left(\frac{6 \log(T)}{\delta} \right)} + \\ & 1024HSB \sqrt{\left(H \log(|\mathcal{S}||\mathcal{A}|H) + |\mathcal{S}| \log \left(\left\lceil \frac{32H^2SB}{\epsilon} \right\rceil \right) + \log \left(\frac{6 \log(HT)}{\delta} \right) \right)} SATH. \end{aligned}$$

For all $T \in \mathbb{N}$ simultaneously. Where $\mathbb{P}(\mathcal{E}_\delta \cap \mathcal{E}_2 \cap \bar{\mathcal{E}}_0 \cap \bar{\mathcal{E}}_{-1} \cap \bar{\mathcal{E}}_2 \cap \bar{\mathcal{E}}_3 \cap \bar{\mathcal{E}}_4 \cap \bar{\mathcal{E}}_5(\frac{\delta}{2})) \geq 1 - 15\delta$.

C.1 Supporting Related Work Lemmas

We will make use of the following Lemma (see Lemma B.1 in Chatterji et al. (2021)),

Lemma 12. *For any fixed policy π , and any function $f : \Gamma \rightarrow \mathbb{R}$ satisfying $\max_{\tau \in \Gamma} |f(\tau)| \leq \eta$, with probability at least $1 - \delta$ for all $t \in \mathbb{N}$,*

$$\mathbb{E}_{s_1 \sim \rho, \tau \sim \mathbb{P}^\pi(\cdot|s_1)} [f(\tau)] - \mathbb{E}_{s_1 \sim \rho, \tau \sim \hat{\mathbb{P}}_t^\pi(\cdot|s_1)} [f(\tau)] \leq \hat{B}_t(\pi, \eta, \delta)$$

We will also make use of the following Lemma (see Lemma B.2 from Chatterji et al. (2021)) corresponding to the uniform version of lemma 12.

Lemma 13 (Uniform version of Lemma 12). *Let $\epsilon > 0$. For any function $f : \Gamma \rightarrow \mathbb{R}$ satisfying $\max_{\tau \in \Gamma} |f(\tau)| \leq \eta$, for all policies π simultaneously and all $t \in \mathbb{N}$,*

$$\mathbb{E}_{s_1 \sim \rho, \tau \sim \hat{\mathbb{P}}^\pi(\cdot|s_1)} [f(\tau)] - \mathbb{E}_{s_1 \sim \rho, \tau \sim \mathbb{P}_t^\pi(\cdot|s_1)} [f(\tau)] \leq B_t(\pi, \eta, \delta, \epsilon) + \epsilon.$$

We will make use of the following standard bound on the covering number of the l_2 ball.

Lemma 14. *For any $\epsilon \in (0, 1]$ the ϵ -covering number of the Euclidean ball in \mathbb{R}^d with radius $r > 0$ i.e. $\{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\|_2 \leq r\}$ is upper bounded by $(\frac{1+2r}{\epsilon})^d$.*

C.2 Proof of Lemma 3

Lemma 3. *Let $\bar{\mathcal{E}}_0$ be the event that for all $t \in \mathbb{N}$,*

$$\begin{aligned} \|\mathbf{w}_t^L - \mathbf{w}_*\|_{\tilde{\mathbf{v}}_t} &\leq \sqrt{2} \|\mathbf{w}_t^L - \mathbf{w}_*\|_{\bar{\mathbf{v}}_t} + \\ &\sqrt{\sum_{\ell=1}^{t-1} 4 \left(\hat{B}_t \left(\pi, 2WB, \frac{\delta'}{8\ell^3|\mathcal{A}|\mathcal{S}} \right) \right)^2} + \frac{1}{t}. \end{aligned}$$

where $\delta' = \frac{\delta}{(\frac{1+4W}{\epsilon})^d}$ and $\epsilon = \frac{1}{t^{2\kappa\lambda+4B^2t^3}}$. Then $\mathbb{P}(\bar{\mathcal{E}}_0) \geq 1 - \delta$.

Proof. Recall that as a result of assumption 1 and the definition of \mathbf{w}_t^L we can bound $\|\mathbf{w}_t^L - \mathbf{w}_*\| \leq 2S$. Let \mathbf{v} be such that $\|\mathbf{v}\| \leq 2S$.

Let's consider $\mathbf{v}^\top \tilde{\mathbf{V}}_t \mathbf{v}$,

$$\mathbf{v}^\top \tilde{\mathbf{V}}_t \mathbf{v} = \kappa\lambda \|\mathbf{v}\|^2 + \sum_{\ell=1}^{t-1} \left(\left\langle \mathbf{v}, \phi^{\widehat{\mathbb{P}}_\ell}(\pi_\ell^1) - \phi^{\widehat{\mathbb{P}}_\ell}(\pi_\ell^2) \right\rangle \right)^2$$

Let's focus on a single summand $\left\langle \mathbf{v}, \phi^{\widehat{\mathbb{P}}_\ell}(\pi_\ell^1) - \phi^{\widehat{\mathbb{P}}_\ell}(\pi_\ell^2) \right\rangle^2$ with $\ell \in [t-1]$.

By Lemma 12, with probability at least $1 - \frac{\delta'}{4\ell^3}$ for all $\pi \in \Pi$ simultaneously,

$$\langle \phi^{\widehat{\mathbb{P}}_\ell}(\pi), \mathbf{v} \rangle \leq \langle \phi(\pi), \mathbf{v} \rangle + \widehat{B}_\ell \left(\pi, 2SB, \frac{\delta'}{8\ell^3 \mathcal{A}S} \right).$$

And

$$-\langle \phi^{\widehat{\mathbb{P}}_\ell}(\pi), \mathbf{v} \rangle \leq -\langle \phi(\pi), \mathbf{v} \rangle + \widehat{B}_\ell \left(\pi, 2SB, \frac{\delta'}{8\ell^3 \mathcal{A}S} \right).$$

Then

$$\left| \langle \phi^{\widehat{\mathbb{P}}_\ell}(\pi_\ell^1) - \phi^{\widehat{\mathbb{P}}_\ell}(\pi_\ell^2), \mathbf{v} \rangle \right| \leq \left| \langle \phi(\pi_\ell^1) - \phi(\pi_\ell^2), \mathbf{v} \rangle \right| + \widehat{B}_\ell \left(\pi_\ell^1, 2SB, \frac{\delta'}{8\ell^3 \mathcal{A}S} \right) + \widehat{B}_\ell \left(\pi_\ell^2, 2SB, \frac{\delta'}{8\ell^3 \mathcal{A}S} \right).$$

And therefore,

$$\langle \phi^{\widehat{\mathbb{P}}_\ell}(\pi_\ell^1) - \phi^{\widehat{\mathbb{P}}_\ell}(\pi_\ell^2), \mathbf{v} \rangle^2 \leq 2\langle \phi(\pi_\ell^1) - \phi(\pi_\ell^2), \mathbf{v} \rangle^2 + 4 \left(\widehat{B}_\ell \left(\pi_\ell^1, 2SB, \frac{\delta'}{8\ell^3 \mathcal{A}S} \right) \right)^2 + 4 \left(\widehat{B}_\ell \left(\pi_\ell^2, 2SB, \frac{\delta'}{8\ell^3 \mathcal{A}S} \right) \right)^2.$$

Therefore with probability at least $1 - \frac{\delta}{2t^2}$, for all $t \in \mathbb{N}$ simultaneously,

$$\begin{aligned} \mathbf{v}^\top \tilde{\mathbf{V}}_t \mathbf{v} &= \kappa\lambda \|\mathbf{v}\|^2 + \sum_{\ell=1}^{t-1} \left\langle \mathbf{v}, \phi^{\widehat{\mathbb{P}}_\ell}(\pi_\ell^1) - \phi^{\widehat{\mathbb{P}}_\ell}(\pi_\ell^2) \right\rangle^2 \\ &\leq 2\kappa\lambda \|\mathbf{v}\|^2 + 2 \sum_{\ell=1}^{t-1} \langle \mathbf{v}, \phi(\pi_\ell^1) - \phi(\pi_\ell^2) \rangle^2 + 4 \left(\widehat{B}_t \left(\pi_\ell^1, 2SB, \frac{\delta'}{8\ell^3 \mathcal{A}S} \right) \right)^2 + 4 \left(\widehat{B}_t \left(\pi_\ell^2, 2SB, \frac{\delta'}{8\ell^3 \mathcal{A}S} \right) \right)^2 \\ &= 2\mathbf{v}^\top \tilde{\mathbf{V}}_t \mathbf{v} + \sum_{\ell=1}^{t-1} 4 \left(\widehat{B}_t \left(\pi_\ell^1, 2SB, \frac{\delta'}{8\ell^3 \mathcal{A}S} \right) \right)^2 + 4 \left(\widehat{B}_t \left(\pi_\ell^2, 2SB, \frac{\delta'}{8\ell^3 \mathcal{A}S} \right) \right)^2 \end{aligned}$$

Consider an ϵ -cover of the $2S$ ball in \mathbb{R}^d and let's see that for any other \mathbf{v} in the $2S$ ball, the closest vector in the covering $\tilde{\mathbf{v}}$ satisfies, $\|\mathbf{v} - \tilde{\mathbf{v}}\| \leq \epsilon$ and therefore,

$$\begin{aligned} \left| \mathbf{v}^\top \tilde{\mathbf{V}}_t \mathbf{v} - \tilde{\mathbf{v}}^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{v}} \right| &= \left| \mathbf{v}^\top \tilde{\mathbf{V}}_t \mathbf{v} - \mathbf{v}^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{v}} + \mathbf{v}^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{v}} - \tilde{\mathbf{v}}^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{v}} \right| \\ &\leq \left| \mathbf{v}^\top \tilde{\mathbf{V}}_t \mathbf{v} - \mathbf{v}^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{v}} \right| + \left| \mathbf{v}^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{v}} - \tilde{\mathbf{v}}^\top \tilde{\mathbf{V}}_t \tilde{\mathbf{v}} \right| \\ &\leq \|\tilde{\mathbf{V}}_t \mathbf{v}\| \|\mathbf{v} - \tilde{\mathbf{v}}\| + \|\tilde{\mathbf{V}}_t \tilde{\mathbf{v}}\| \|\mathbf{v} - \tilde{\mathbf{v}}\| \\ &\leq \epsilon (\kappa\lambda + 4B^2(t-1)) \end{aligned}$$

Similarly

$$|\mathbf{v}^\top \bar{\mathbf{V}}_t \mathbf{v} - \tilde{\mathbf{v}}^\top \bar{\mathbf{V}}_t \tilde{\mathbf{v}}| \leq \epsilon (\kappa\lambda + 4B^2(t-1))$$

Invoking Lemma 14 and setting $\delta' = \frac{\delta}{\left(\frac{1+4S}{\epsilon}\right)^d}$ applying it to $\mathbf{w}_t^L - \mathbf{w}_*$

$$\begin{aligned} \|\mathbf{w}_t^L - \mathbf{w}_*\|_{\tilde{\mathbf{V}}_t}^2 &\leq 2\|\mathbf{w}_t^L - \mathbf{w}_*\|_{\bar{\mathbf{V}}_t}^2 + \\ &\quad \sum_{\ell=1}^{t-1} 4 \left(\hat{B}_t \left(\pi_\ell^1, 2SB, \frac{\delta'}{8\ell^3 \mathcal{A}S} \right) \right)^2 + 4 \left(\hat{B}_t \left(\pi_\ell^2, 2SB, \frac{\delta'}{8\ell^3 \mathcal{A}S} \right) \right)^2 + 2\epsilon (\kappa\lambda + 4B^2(t-1)). \end{aligned}$$

Setting $\epsilon = \frac{1}{t^2 \kappa\lambda + 4B^2 t^3}$ and using the fact that all $a, b, c \geq 0$ we have $\sqrt{a^2 + b^2 + c^2} \leq a + b + c$,

$$\|\mathbf{w}_t^L - \mathbf{w}_*\|_{\tilde{\mathbf{V}}_t} \leq \sqrt{2}\|\mathbf{w}_t^L - \mathbf{w}_*\|_{\bar{\mathbf{V}}_t} + \sqrt{\sum_{\ell=1}^{t-1} 4 \left(\hat{B}_t \left(\pi_\ell^1, 2SB, \frac{\delta'}{8\ell^3 \mathcal{A}S} \right) \right)^2 + 4 \left(\hat{B}_t \left(\pi_\ell^2, 2SB, \frac{\delta'}{8\ell^3 \mathcal{A}S} \right) \right)^2} + \frac{1}{t}.$$

Since the inequality holds for any $t \in \mathbb{N}$ with probability at least $1 - \frac{\delta}{2t^2}$, by the union bound, the inequality holds for all $t \in \mathbb{N}$ simultaneously with probability at least $1 - \delta$. \square

C.3 Proof of Lemma 4

Lemma 4. *Let $\bar{\mathcal{E}}_{-1}$ be the event that $\pi^* \in \Pi_t$ for all $t \in \mathbb{N}$. Then $\mathbb{P}(\bar{\mathcal{E}}_{-1}) \geq 1 - 5\delta$.*

Proof. Let's start by conditioning on \mathcal{E}_δ , $\bar{\mathcal{E}}_0$ and \mathcal{E}_2 (see Corollary 1 for a definition of \mathcal{E}_2). By Lemmas 1 and 3, $\mathbb{P}(\mathcal{E}_\delta \cap \bar{\mathcal{E}}_0 \cap \mathcal{E}_2) \geq 1 - 4\delta$

By definition of π_* , $(\phi(\pi^*) - \phi(\pi))^\top \mathbf{w}^* \geq 0$ for any arbitrary π . Therefore,

$$0 \leq (\phi(\pi^*) - \phi(\pi))^\top \mathbf{w}^*$$

By Lemma 12, with probability at least $1 - \delta$ for all $\pi_1, \pi_2 \in \Pi$ simultaneously and all $t \in \mathbb{N}$

$$(\phi(\pi_1) - \phi(\pi_2))^\top \mathbf{w}^* \leq (\hat{\phi}^{\hat{\mathbb{P}}_t}(\pi_1) - \hat{\phi}^{\hat{\mathbb{P}}_t}(\pi_2))^\top \mathbf{w}^* + \hat{B}_t \left(\pi_1, 2SB, \frac{\delta}{2\mathcal{A}S} \right) + \hat{B}_t \left(\pi_2, 2SB, \frac{\delta}{\mathcal{A}^2 S} \right), \quad (23)$$

In particular this implies that with probability at least $1 - \delta$ for π^* and any π ,

$$0 \leq (\phi(\pi^*) - \phi(\pi))^\top \mathbf{w}^* \leq (\hat{\phi}^{\hat{\mathbb{P}}_t}(\pi^*) - \hat{\phi}^{\hat{\mathbb{P}}_t}(\pi))^\top \mathbf{w}^* + \hat{B}_t \left(\pi^*, 2SB, \frac{\delta}{2\mathcal{A}S} \right) + \hat{B}_t \left(\pi, 2SB, \frac{\delta}{\mathcal{A}^2 S} \right), \quad (24)$$

Let's bound the term $(\hat{\phi}^{\hat{\mathbb{P}}_t}(\pi^*) - \hat{\phi}^{\hat{\mathbb{P}}_t}(\pi))^\top \mathbf{w}^*$

$$\begin{aligned} (\hat{\phi}^{\hat{\mathbb{P}}_t}(\pi^*) - \hat{\phi}^{\hat{\mathbb{P}}_t}(\pi))^\top \mathbf{w}^* &= (\hat{\phi}^{\hat{\mathbb{P}}_t}(\pi^*) - \hat{\phi}^{\hat{\mathbb{P}}_t}(\pi))^\top \mathbf{w}_t^L + (\hat{\phi}^{\hat{\mathbb{P}}_t}(\pi^*) - \hat{\phi}^{\hat{\mathbb{P}}_t}(\pi))^\top (\mathbf{w}^* - \mathbf{w}_t^L) \\ &\leq (\hat{\phi}^{\hat{\mathbb{P}}_t}(\pi^*) - \hat{\phi}^{\hat{\mathbb{P}}_t}(\pi))^\top \mathbf{w}_t^L + \|\hat{\phi}^{\hat{\mathbb{P}}_t}(\pi^*) - \hat{\phi}^{\hat{\mathbb{P}}_t}(\pi)\|_{\tilde{\mathbf{V}}_t} \|\mathbf{w}^* - \mathbf{w}_t^L\|_{\tilde{\mathbf{V}}_t} \end{aligned} \quad (25)$$

Since $\bar{\mathcal{E}}_0$ holds, by Lemma 3

$$\|\mathbf{w}_t^L - \mathbf{w}_*\|_{\tilde{\mathbf{V}}_t} \leq \sqrt{2}\|\mathbf{w}_t^L - \mathbf{w}_*\|_{\tilde{\mathbf{V}}_t} + \sqrt{\sum_{\ell=1}^{t-1} 4 \left(\hat{B}_t \left(\pi, 2SB, \frac{\delta'}{8\ell^3 \mathcal{A}\mathcal{S}} \right) \right)^2} + \frac{1}{t}.$$

Since $\mathcal{E}_\delta \cap \mathcal{E}_2$ is assumed to hold Corollary 1 implies that $\|\mathbf{w}_t^L - \mathbf{w}_*\|_{\tilde{\mathbf{V}}_t} \leq 4\kappa\beta_t(\delta) + \alpha_{d,T}(\delta)$ and therefore,

$$\|\mathbf{w}_t^L - \mathbf{w}_*\|_{\tilde{\mathbf{V}}_t} \leq \sqrt{2} (4\kappa\beta_t(\delta) + \alpha_{d,T}(\delta)) + \sqrt{\sum_{\ell=1}^{t-1} 4 \left(\hat{B}_t \left(\pi, 2SB, \frac{\delta'}{8\ell^3 \mathcal{A}\mathcal{S}} \right) \right)^2} + \frac{1}{t}.$$

Since $\gamma_t = \sqrt{2} (4\kappa\beta_t(\delta) + \alpha_{d,T}(\delta)) + \sqrt{\sum_{\ell=1}^{t-1} 4 \left(\hat{B}_t \left(\pi, 2SB, \frac{\delta'}{8\ell^3 \mathcal{A}\mathcal{S}} \right) \right)^2} + \frac{1}{t}$, combining these results with Equations 25 and 24 yields,

$$0 \leq (\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi))^\top \mathbf{w}_t^L + \gamma_t \|\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi)\|_{\tilde{\mathbf{V}}_{t-1}} + \hat{B}_t \left(\pi^*, 2SB, \frac{\delta}{2\mathcal{A}\mathcal{S}} \right) + \hat{B}_t \left(\pi, 2SB, \frac{\delta}{\mathcal{A}^2\mathcal{S}} \right),$$

Thus implying $\pi^* \in \Pi_t$. Taking a union bound between $\mathcal{E}_\delta \cap \bar{\mathcal{E}}_0 \cap \mathcal{E}_2$ and the $1 - \delta$ probability event from Equation 25 yields the result. □

C.4 Proof of Lemma 6

Full version of Lemma 6,

Lemma 15. *If $\mathcal{E}_\delta \cap \mathcal{E}_2 \cap \bar{\mathcal{E}}_0 \cap \bar{\mathcal{E}}_{-1} \cap \bar{\mathcal{E}}_2$ the regret is bounded by,*

$$R_T \leq 2\gamma_T \sqrt{2Td \log \left(1 + \frac{TB}{d} \right)} + \sum_{t \in [T]} 4\hat{B}_t(\pi_t^1, 4WB, \delta) + 4\hat{B}_t(\pi_t^2, 4WB, \delta)$$

Proof. We first condition on $\mathcal{E}_\delta \cap \mathcal{E}_2 \cap \bar{\mathcal{E}}_0 \cap \bar{\mathcal{E}}_{-1}$. Let's start by showing the following bound on the instantaneous regret,

$$2r_t \leq 2\gamma_t \|\phi^{\hat{\mathbb{P}}_t}(\pi_t^1) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^2)\|_{\tilde{\mathbf{V}}_{t-1}} + 4\hat{B}_t(\pi_t^1, 4SB, \delta) + 4\hat{B}_t(\pi_t^2, 4SB, \delta)$$

Since we are conditioning on $\bar{\mathcal{E}}_2$, by Lemma 5 follows that for all t ,

$$2r_t \leq (\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^1))^\top \mathbf{w}^* + (\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^2))^\top \mathbf{w}^* + \hat{B}_t(\pi^*, 4SB, \delta) + \hat{B}_t(\pi_t^1, 2SB, \delta) + \hat{B}_t(\pi_t^2, 2SB, \delta).$$

Let's focus on bounding the term $(\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^1))^\top \mathbf{w}^* + (\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^2))^\top \mathbf{w}^*$.

$$\begin{aligned} (\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^1))^\top \mathbf{w}^* + (\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^2))^\top \mathbf{w}^* &= (\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^1))^\top \mathbf{w}_t^L + (\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^1))^\top (\mathbf{w}^* - \mathbf{w}_t^L) \\ &\quad + (\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^2))^\top \mathbf{w}_t^L + (\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^2))^\top (\mathbf{w}^* - \mathbf{w}_t^L) \\ &\leq (\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^1))^\top \mathbf{w}_t^L + (\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^2))^\top \mathbf{w}_t^L \\ &\quad + \|\mathbf{w}^* - \mathbf{w}_t^L\|_{\tilde{\mathbf{V}}_t} \cdot \|\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^1)\|_{\tilde{\mathbf{V}}_{t-1}} + \\ &\quad \|\mathbf{w}^* - \mathbf{w}_t^L\|_{\tilde{\mathbf{V}}_t} \cdot \|\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^2)\|_{\tilde{\mathbf{V}}_{t-1}} \end{aligned} \tag{26}$$

Since $\mathcal{E}_\delta \cap \mathcal{E}_2 \cap \bar{\mathcal{E}}_0$ holds, the last two terms in the sum above can be bounded using Lemma 3 and Corollary 1 by

$$\begin{aligned} & \|\mathbf{w}^* - \mathbf{w}_t^L\|_{\tilde{\mathbf{V}}_t} \cdot \|\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^1)\|_{\tilde{\mathbf{V}}_{t-1}} + \|\mathbf{w}^* - \mathbf{w}_t^L\|_{\tilde{\mathbf{V}}_t} \cdot \|\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^2)\|_{\tilde{\mathbf{V}}_{t-1}} \\ & \leq \gamma_t \left(\|\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^1)\|_{\tilde{\mathbf{V}}_{t-1}} + \|\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^2)\|_{\tilde{\mathbf{V}}_{t-1}} \right) \end{aligned}$$

Where

$$\gamma_t = \sqrt{2} (4\kappa\beta_t(\delta) + \alpha_{d,T}(\delta)) + \sqrt{\sum_{\ell=1}^{t-1} 4 \left(\hat{B}_t \left(\pi_\ell^1, 2SB, \frac{\delta'}{8\ell^3 \mathcal{A}\mathcal{S}} \right) \right)^2 + 4 \left(\hat{B}_t \left(\pi_\ell^2, 2SB, \frac{\delta'}{8\ell^3 \mathcal{A}\mathcal{S}} \right) \right)^2} + \frac{1}{t}.$$

The first two terms on the right hand side of inequality 26 leverage the optimistic bonus, using the fact that $\pi_t^1, \pi_t^2 \in \Pi_t$ and therefore,

$$\begin{aligned} & (\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^1))^\top \mathbf{w}_t^L + (\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^2))^\top \mathbf{w}_t^L + \hat{B}_t(\pi^*, 4SB, \delta) + \hat{B}_t(\pi_t^1, 2SB, \delta) + \hat{B}_t(\pi_t^2, 2SB, \delta) \\ & \leq \gamma_t \left(\|\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^1)\|_{\tilde{\mathbf{V}}_{t-1}} + \|\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^2)\|_{\tilde{\mathbf{V}}_{t-1}} \right) + 3\hat{B}_t(\pi^*, 4SB, \delta) + \hat{B}_t(\pi_t^1, 4SB, \delta) + \hat{B}_t(\pi_t^2, 4SB, \delta) \\ & \leq \gamma_t \left(\|\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^1)\|_{\tilde{\mathbf{V}}_{t-1}} + \|\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^2)\|_{\tilde{\mathbf{V}}_{t-1}} \right) + 4\hat{B}_t(\pi^*, 4SB, \delta) + 2\hat{B}_t(\pi_t^1, 4SB, \delta) + 2\hat{B}_t(\pi_t^2, 4SB, \delta) \end{aligned}$$

Putting these together we can conclude that,

$$\begin{aligned} 2r_t & \leq 2\gamma_t \left(\|\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^1)\|_{\tilde{\mathbf{V}}_{t-1}} + \|\phi^{\hat{\mathbb{P}}_t}(\pi^*) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^2)\|_{\tilde{\mathbf{V}}_{t-1}} \right) + \\ & \quad 4\hat{B}_t(\pi^*, 4SB, \delta) + 2\hat{B}_t(\pi_t^1, 4SB, \delta) + 2\hat{B}_t(\pi_t^2, 4SB, \delta) \end{aligned}$$

Recall that whenever $\bar{\mathcal{E}}_{-1}$ holds, $\pi^* \in \Pi_t$ and that as a result of how π_t^1, π_t^2 are chosen (see Algorithm 2)

$$2r_t \leq 2\gamma_t \|\phi^{\hat{\mathbb{P}}_t}(\pi_t^1) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^2)\|_{\tilde{\mathbf{V}}_{t-1}} + 4\hat{B}_t(\pi_t^1, 4SB, \delta) + 4\hat{B}_t(\pi_t^2, 4SB, \delta)$$

The regret is therefore upper bounded by,

$$\begin{aligned} R_T & = \sum_{t \in [T]} 2r_t \\ & \leq \sum_{t \in [T]} 2\gamma_t \|\phi^{\hat{\mathbb{P}}_t}(\pi_t^1) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^2)\|_{\tilde{\mathbf{V}}_{t-1}} + 4\hat{B}_t(\pi_t^1, 4SB, \delta) + 4\hat{B}_t(\pi_t^2, 4SB, \delta) \\ & \leq 2\gamma_T \sqrt{T \sum_{t \in [T]} \|\phi^{\hat{\mathbb{P}}_t}(\pi_t^1) - \phi^{\hat{\mathbb{P}}_t}(\pi_t^2)\|_{\tilde{\mathbf{V}}_{t-1}}^2} + 4\hat{B}_t(\pi_t^1, 4SB, \delta) + 4\hat{B}_t(\pi_t^2, 4SB, \delta) \\ & \leq 2\gamma_T \sqrt{2Td \log \left(1 + \frac{TB}{d} \right)} + \sum_{t \in [T]} 4\hat{B}_t(\pi_t^1, 4SB, \delta) + 4\hat{B}_t(\pi_t^2, 4SB, \delta) \end{aligned}$$

Where the last inequality follows from Lemma 8. □

D Miscellaneous Technical Lemmas

We will make use of the following Lemmas

Lemma 16 (Hoeffding Inequality). *Let $\{x_t\}_{t=1}^\infty$ be a martingale difference sequence with $|x_t| \leq \zeta$ and let $\delta \in (0, 1]$. Then with probability $1 - \delta$ for all $T \in \mathbb{N}$*

$$\sum_{t=1}^T x_t \leq 2\zeta \sqrt{T \ln \left(\frac{6 \ln T}{\delta} \right)}.$$

Proof. Observe that $\frac{|x_t|}{\zeta} \leq 1$. By invoking a time-uniform Hoeffding-style concentration inequality (Howard et al., 2020, Equation (11)) we find that

$$\Pr \left[\forall t \in \mathbb{N} : \sum_{i=1}^t \frac{x_i}{\zeta} \leq 1.7 \sqrt{T \left(\log \log(T) + 0.72 \log \left(\frac{5.2}{\delta} \right) \right)} \right] \geq 1 - \delta.$$

Rounding up the constants for the sake of simplicity we get

$$\Pr \left[\forall t \in \mathbb{N} : \sum_{i=1}^t x_i \leq 2\zeta \sqrt{T \left(\log \left(\frac{6 \log(T)}{\delta} \right) \right)} \right] \geq 1 - \delta,$$

which establishes our claim. □

Lemma 17 (Uniform empirical Bernstein bound). *In the terminology of Howard et al. (2021), let $S_t = \sum_{i=1}^t Y_i$ be a sub- ψ_P process with parameter $c > 0$ and variance process W_t . Then with probability at least $1 - \delta$ for all $t \in \mathbb{N}$*

$$\begin{aligned} S_t \leq & 1.44 \sqrt{\max(W_t, m) \left(1.4 \ln \ln \left(2 \left(\max \left(\frac{W_t}{m}, 1 \right) \right) \right) + \ln \frac{5.2}{\delta} \right)} \\ & + 0.41c \left(1.4 \ln \ln \left(2 \left(\max \left(\frac{W_t}{m}, 1 \right) \right) \right) + \ln \frac{5.2}{\delta} \right) \end{aligned}$$

where $m > 0$ is arbitrary but fixed.