
Improved Regret Bound and Experience Replay in Regularized Policy Iteration

Nevena Lazić¹ Dong Yin¹ Yasin Abbasi-Yadkori¹ Csaba Szepesvári^{1,2}

Abstract

In this work, we study algorithms for learning in infinite-horizon undiscounted Markov decision processes (MDPs) with function approximation. We first show that the regret analysis of the POLITEX algorithm (a version of regularized policy iteration) can be sharpened from $O(T^{3/4})$ to $O(\sqrt{T})$ under nearly identical assumptions, and instantiate the bound with linear function approximation. Our result provides the first high-probability $O(\sqrt{T})$ regret bound for a computationally efficient algorithm in this setting. The exact implementation of POLITEX with neural network function approximation is inefficient in terms of memory and computation. Since our analysis suggests that we need to approximate the average of the action-value functions of past policies well, we propose a simple efficient implementation where we train a single Q-function on a replay buffer with past data. We show that this often leads to superior performance over other implementation choices, especially in terms of wall-clock time. Our work also provides a novel theoretical justification for using experience replay within policy iteration algorithms.

1. Introduction

Model-free reinforcement learning (RL) algorithms combined with powerful function approximation have achieved impressive performance in a variety of application domains over the last decade. Unfortunately, the theoretical understanding of such methods is still quite limited. In this work, we study single-trajectory learning in infinite-horizon undiscounted Markov decision processes (MDPs), also known as average-reward MDPs, which capture tasks such as routing and the control of physical systems.

One line of works with performance guarantees for the average-reward setting follows the “online MDP” approach

¹DeepMind ²University of Alberta. Correspondence to: Nevena Lazić <nevena@google.com>.

proposed by Even-Dar et al. (2009), where the agent selects policies by running an online learning algorithm in each state, typically mirror descent. The resulting algorithm is a version of approximate policy iteration (API), which alternates between (1) estimating the state-action value function (or Q-function) of the current policy and (2) setting the next policy to be optimal w.r.t. the sum of all previous Q-functions plus a regularizer. Note that, by contrast, standard API sets the next policy only based on the most recent Q-function. The policy update can also be written as maximizing the most recent Q-function minus KL-divergence to the previous policy, which is somewhat similar to recently popular versions of API (Schulman et al., 2015; 2017; Achiam et al., 2017; Abdolmaleki et al., 2018; Song et al., 2019a).

The original work of Even-Dar et al. (2009) studied this scheme with known dynamics, tabular representation, and adversarial reward functions. More recent works (Abbasi-Yadkori et al., 2019; Hao et al., 2020; Wei et al., 2020a) have adapted the approach to the case of unknown dynamics, stochastic rewards, and value function approximation. With linear value functions, the POLITEX algorithm of Abbasi-Yadkori et al. (2019) achieves $O(T^{3/4})$ high-probability regret in ergodic MDPs, and the results only scale in the number of features rather than states. Wei et al. (2020a) later show an $O(\sqrt{T})$ bound on *expected* regret for a similar algorithm named MDP-EXP2. In this work, we revisit the analysis of POLITEX and show that it can be sharpened to $O(\sqrt{T})$ under nearly identical assumptions, resulting in the first $O(\sqrt{T})$ *high-probability* regret bound for a computationally efficient algorithm in this setting.

In addition to improved analysis, our work also addresses practical implementation of POLITEX with neural networks. The policies produced by POLITEX in each iteration require access to the sum of all previous Q-function estimates. With neural network function approximation, exact implementation requires us to keep all past networks in memory and evaluate them at each step, which is inefficient in terms of memory and computation. Some practical implementation choices include subsampling Q-functions and/or optimizing a KL-divergence regularized objective w.r.t. a parametric policy at each iteration. We propose an alternative approach, where we approximate the average of all past Q-functions by training a single network on a replay buffer with past

data. We demonstrate that this choice often outperforms other approximate implementations, especially in terms of run-time. When available memory is constrained, we propose to subsample transitions using the notion of coresets (Bachem et al., 2017).

Our work also provides a novel perspective on the benefits of experience replay. Experience replay is a standard tool for stabilizing learning in modern deep RL, and typically used in *off-policy* methods like Deep Q-Networks (Mnih et al., 2013), as well as “value gradient” methods such as DDPG (Lillicrap et al., 2016) and SVG (Heess et al., 2015). A different line of *on-policy* methods typically does not rely on experience replay; instead, learning is stabilized by constraining consecutive policies to be close in terms of KL divergence (Schulman et al., 2015; Song et al., 2019a; Degraeve et al., 2019). We observe that both experience replay and KL-divergence regularization can be viewed as approximate implementations of POLITEX. Thus, we provide a theoretical justification for using experience replay in API, as an approximate implementation of online learning in each state. Note that this online-learning view differs from the commonly used justifications for experience replay, namely that it “breaks temporal correlations” (Schaul et al., 2016; Mnih et al., 2013). Our analysis also suggests a new objective for subsampling or priority-sampling transitions in the replay buffer, which differs priority-sampling objectives of previous work (Schaul et al., 2016).

In summary, our main contributions are (1) an improved analysis of POLITEX, showing an $O(\sqrt{T})$ regret bound under the same assumptions as the original work, and (2) an efficient implementation that also offers a new perspective on the benefits of experience replay.

2. Setting and Notation

We consider learning in infinite-horizon undiscounted ergodic MDPs $(\mathcal{X}, \mathcal{A}, r, P)$, where $\mathcal{X}$ is the state space, $\mathcal{A}$ is a finite action space, $r : \mathcal{X} \times \mathcal{A} \rightarrow [0, 1]$ is an unknown reward function, and $P : \mathcal{X} \times \mathcal{A} \rightarrow \Delta_{\mathcal{X}}$ is the unknown probability transition function. A policy $\pi : \mathcal{X} \rightarrow \Delta_{\mathcal{A}}$ is a mapping from a state to a distribution over actions. Let $\{(x_t^\pi, a_t^\pi)\}_{t=1}^\infty$ denote the state-action sequence obtained by following policy π . The expected average reward of policy π is defined as

$$J_\pi := \lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T r(x_t^\pi, a_t^\pi) \right]. \quad (2.1)$$

Let μ_π denote the stationary state distribution of a policy π , satisfying $\mu_\pi = \mathbb{E}_{x \sim \mu_\pi, a \sim \pi}[P(\cdot|x, a)]$. We will sometimes write μ_π as a vector, and use $\nu_\pi = \mu_\pi \otimes \pi$ to denote the stationary state-action distribution. In ergodic MDPs, J_π and μ_π are well-defined and independent of the initial state. The optimal policy π_* is a policy that maximizes the

expected average reward. We will denote by J_* and μ_* the expected average reward and stationary state distribution of π_* , respectively.

The value function of a policy π is defined as:

$$V_\pi(x) := \mathbb{E} \left[\sum_{t=1}^{\infty} (r(x_t^\pi, a_t^\pi) - J_\pi) | x_1^\pi = x \right]. \quad (2.2)$$

The state-action value function $Q_\pi(x, a)$ is defined as

$$Q_\pi(x, a) := r(x, a) - J_\pi + \sum_{x'} P(x'|x, a) V_\pi(x'). \quad (2.3)$$

Notice that

$$V_\pi(x) = \sum_a \pi(a|x) Q_\pi(x, a). \quad (2.4)$$

Equations (2.3) and (2.4) are known as the Bellman equation. If we do not use the definition of $V_\pi(x)$ in Eq. (2.2) and instead solve for $V_\pi(x)$ and $Q_\pi(x, a)$ using the Bellman equation, we can see that the solutions are unique up to an additive constant. Therefore, in the following, we may use $V_\pi(x)$ and $Q_\pi(x, a)$ to denote the same function up to a constant. We will use the shorthand notation $Q_\pi(x, \pi') = \mathbb{E}_{a \sim \pi'(\cdot|x)}[Q_\pi(x, a)]$; note that $Q_\pi(x, \pi) = V_\pi(x)$. The advantage function of a policy π is defined as $A_\pi(x, a) = Q_\pi(x, a) - V_\pi(x)$.

The agent interacts with the environment as follows: at each round t , the agent observes a state $x_t \in \mathcal{X}$, chooses an action $a_t \sim \pi_t(\cdot|x_t)$, and receives a reward $r_t := r(x_t, a_t)$. The environment then transitions to the next state x_{t+1} with probability $P(x_{t+1}|x_t, a_t)$. Recall that π_* is the optimal policy and J_* is its expected average reward. The regret of an algorithm with respect to this fixed policy is defined as

$$R_T := \sum_{t=1}^T (J_* - r(x_t, a_t)). \quad (2.5)$$

The learning goal is to find an algorithm that minimizes the long-term regret R_T .

Our analysis will require the following assumption on the mixing rate of policies.

Assumption 2.1 (Uniform mixing). Let H_π be the state-action transition matrix of a policy π . Let $\gamma(H_\pi)$ be the corresponding *ergodicity coefficient* (Seneta, 1979), defined as $\gamma(H_\pi) := \max_{z: z^\top \mathbf{1}=0, \|z\|_1=1} \|z^\top H_\pi\|_1$. We assume that there exists a scalar $\gamma < 1$ such that for any policy π , $\gamma(H_\pi) \leq \gamma < 1$.

Assumption 2.1 implies that for any pair of distributions ν, ν' , $\|(\nu - \nu')^\top H_\pi\|_1 \leq \gamma \|\nu - \nu'\|_1$; see Lemma A.1 in Appendix A for a proof.

3. Related Work

Regret bounds for average-reward MDPs. Most no-regret algorithms for infinite-horizon undiscounted MDPs are only applicable to tabular representations and model-based (Bartlett, 2009; Jaksch et al., 2010; Ouyang et al., 2017; Fruit et al., 2018; Jian et al., 2019; Talebi & Maillard, 2018). For weakly-communicating MDPs with diameter D , these algorithms nearly achieve the minimax lower bound $\Omega(\sqrt{D|\mathcal{X}||\mathcal{A}|T})$ (Jaksch et al., 2010) with high probability. Wei et al. (2020b) provide model-free algorithms with regret bounds in the tabular setting. In the model-free setting with function approximation, the POLITEX algorithm (Abbasi-Yadkori et al., 2019) achieve $O(d^{1/2}T^{3/4})$ regret in uniformly ergodic MDPs, where d is the size of the compressed state-action space. Hao et al. (2020) improve these results to $O(T^{2/3})$. More recently, Wei et al. (2020a) present three algorithms for average-reward MDPs with linear function approximation. Among these, FOPO achieves $O(\sqrt{T})$ regret but is computationally inefficient, and OLSVLFH is efficient but obtains $O(T^{3/4})$ regret. The MDP-EXP2 algorithm is computationally efficient, and under similar assumptions as in Abbasi-Yadkori et al. (2019) it obtains a $O(\sqrt{T})$ bound on *expected regret* (a weaker guarantee than the high-probability bounds in other works). Our analysis shows a high-probability $O(\sqrt{T})$ regret bound under the same assumptions.

KL-regularized approximate policy iteration. Our work is also related to approximate policy iteration algorithms which constrain each policy to be close to the previous policy in the sense of KL divergence. This approach was popularized by TRPO (Schulman et al., 2015), where it was motivated as an approximate implementation of conservative policy iteration (Kakade & Langford, 2002). Some of the related subsequent works include PPO (Schulman et al., 2017), MPO (Abdolmaleki et al., 2018), V-MPO (Song et al., 2019a), and CPO (Achiam et al., 2017). While these algorithms place a constraint on consecutive policies and are mostly heuristic, another line of research shows that using KL divergence as a regularizer has a theoretical justification in terms of either regret guarantees (Abbasi-Yadkori et al., 2019; Hao et al., 2020; Wei et al., 2020b;a) or error propagation (Vieillard et al., 2020a;b).

Experience replay. Experience replay (Lin, 1992) is one of the central tools for achieving good performance in deep reinforcement learning. While it is mostly used in off-policy methods such as deep Q-learning (Mnih et al., 2013; 2015), it has also shown benefits in value-gradient methods (Heess et al., 2015; Lillicrap et al., 2016), and has been used in some variants of KL-regularized policy iteration (Abdolmaleki et al., 2018; Tomar et al., 2020). Its success has been attributed to removing some temporal correlations from data fed to standard gradient-based optimization algorithms.

Schaul et al. (2016) have shown that non-uniform replay sampling based on the Bellman error can improve performance. Unlike these works, we motivate experience replay from the perspective of online learning in MDPs (Even-Dar et al., 2009) with the goal of approximating the average of past value functions well.

Continual learning. Continual learning (CL) is the paradigm of learning a classifier or regressor that performs well on a set of tasks, where each task corresponds to a different data distribution. The tasks are observed sequentially, and the learning goal is to avoid forgetting past tasks without storing all the data in memory. This is quite similar to our goal of approximating the average of sequentially-observed Q-functions, where the data for approximating each Q-function has different distribution. In general, approaches to CL can be categorized as regularization-based (Kirkpatrick et al., 2017; Zenke et al., 2017; Farajtabar et al., 2020; Yin et al., 2020), expansion-based (Rusu et al., 2016), and replay-based (Lopez-Paz & Ranzato, 2017; Chaudhry et al., 2018; Borsos et al., 2020), with the approaches based on experience replay typically having superior performance over other methods.

4. Algorithm

Our algorithm is similar to the POLITEX schema (Abbasi-Yadkori et al., 2019). In each phase k , POLITEX obtains an estimate $\hat{Q}_{\pi_k}$ of the action-value function Q_{π_k} of the current policy π_k , and then sets the next policy using the mirror descent update rule:

$$\begin{aligned} \pi_{k+1}(\cdot|x) &= \operatorname{argmax}_{\pi} \hat{Q}_{\pi_k}(x, \pi) - \eta^{-1} D_{KL}(\pi \parallel \pi_k(\cdot|x)) \\ &= \operatorname{argmax}_{\pi} \sum_{i=1}^k \hat{Q}_{\pi_i}(x, \pi) + \eta^{-1} \mathcal{H}(\pi) \\ &\propto \exp \left(\eta \sum_{i=1}^k \hat{Q}_{\pi_i}(x, \cdot) \right), \end{aligned} \quad (4.1)$$

where $\mathcal{H}(\cdot)$ is the entropy function. When the functions $\{\hat{Q}_{\pi_i}\}_{i=1}^k$ are tabular or linear, the above update can be implemented efficiently by simply summing all the table entries or weight vectors. However, with neural network function approximation, we need to keep all networks in memory and evaluate them at each step, which quickly becomes inefficient in terms of storage and computation. Some of the efficient implementations proposed in literature include keeping a subset of the action-value functions (Abbasi-Yadkori et al., 2019), and using a parametric policy and optimizing the KL-regularized objective w.r.t. the parameters over the available data (Tomar et al., 2020).

Our proposed method, presented in Algorithm 1, attempts to directly approximate the average of all previous action-value

Algorithm 1 Schema for policy iteration with replay

- 1: **Input:** phase length τ , num. phases K , parameter η
- 2: **Initialize:** $\pi_1(a|x) = 1/|\mathcal{A}|$, empty replay buffer $\mathcal{R}$
- 3: **for** $k = 1, \dots, K$ **do**
- 4: Execute π_k for τ time steps and collect data $\mathcal{D}_k$
- 5: Compute $\hat{Q}_k$, an estimate of $Q_k = \frac{1}{k} \sum_{i=1}^k Q_{\pi_i}$, from data $\mathcal{D}_k$ and replay $\mathcal{R}$
- 6: Set $\pi_{k+1}(a|x) \propto \exp(\eta k \hat{Q}_k(x, a))$
- 7: Update replay $\mathcal{R}$ with $\mathcal{D}_k$
- 8: **end for**
- 9: **Output:** π_{K+1}

functions

$$Q_k(x, a) := \frac{1}{k} \sum_{i=1}^k Q_{\pi_i}(x, a).$$

To do so, we only use a single network, and continually train it to approximate Q_k . At each iteration k , we obtain a dataset of tuples $\mathcal{D}_k = \{(x_t, a_t, R_t)\}$, where R_t is the empirical return from the state-action pair (x_t, a_t) . We initialize $\hat{Q}_k$ to $\hat{Q}_{k-1}$ and update it by minimizing the squared error over the union of $\mathcal{D}_k$ and the replay buffer $\mathcal{R}$. We then update the replay buffer with all or a subset of data in $\mathcal{D}_k$.

In the sequel, in Section 5, we first show that by focusing on estimating the average Q-function, we can improve the regret bound of POLITEX under nearly identical assumption; in Section 6, we instantiate this bound for linear value functions, where we can estimate the average Q-function simply by weight averaging; in Section 7, we focus on practical implementations, in particular, we discuss the limitations of weight averaging, and provide details on how to leverage replay data when using non-linear function approximation; in Section 8 we present our experimental results; and in Section 9 we make final remarks.

5. Regret Analysis of POLITEX

In this section, we revisit the regret analysis of POLITEX (Abbasi-Yadkori et al., 2019) in ergodic average-reward MDPs, and show that it can be improved from $O(T^{3/4})$ to $O(\sqrt{T})$ under similar assumptions. Our analysis relies in part on a simple modification of the regret decomposition. Namely, instead of including the estimation error of each value-function Q_{π_k} in the regret, we consider the error in the running average Q_k . When this error scales as $O(1/\sqrt{k})$ in a particular weighted norm, the regret of POLITEX is $O(\sqrt{T})$. As we show in the following section, this bound can be instantiated for linear value functions under the same assumptions as Abbasi-Yadkori et al. (2019).

Assumption 5.1 (Boundedness). Let $\hat{Q}_{\pi_k} := k\hat{Q}_k - (k-1)\hat{Q}_{k-1}$. We assume that there exists a constant $Q_{\max}$ such

that for all $k = 1, \dots, K$ and for all $x \in \mathcal{X}$,

$$\max_a \hat{Q}_{\pi_k}(x, a) - \min_a \hat{Q}_{\pi_k}(x, a) \leq Q_{\max}.$$

For the purpose of our algorithm, functions $\hat{Q}_{\pi_k}$ are unique up to a constant. Thus we can equivalently assume that $\|\hat{Q}_{\pi_k}(x, \cdot)\|_{\infty} \leq Q_{\max}$ for all x .

Define $\hat{V}_{\pi_k}(x) := \hat{Q}_{\pi_k}(x, \pi_k)$. Let $V_k := \frac{1}{k} \sum_{i=1}^k V_{\pi_i}(x)$ and $\hat{V}_k(x) := \sum_{i=1}^k \hat{V}_{\pi_i}(x)$ be the average of the state-value functions and its estimate. We will require the estimation error of the running average to scale as in the following assumption.

Assumption 5.2 (Estimation error). Let μ_* be the stationary state distribution of the optimal policy π_* . With probability at least $1 - \delta$, for a problem-dependent constant C , the errors in $\hat{Q}_K$ and $\hat{V}_K$ are bounded as

$$\begin{aligned} \mathbb{E}_{x \sim \mu_*} [\hat{V}_K(x) - V_K(x)] &\leq C \sqrt{\log(1/\delta)/K} \\ \mathbb{E}_{x \sim \mu_*} [Q_K(x, \pi_*) - \hat{Q}_K(x, \pi_*)] &\leq C \sqrt{\log(1/\delta)/K}. \end{aligned}$$

Define $S_{\delta}(|\mathcal{A}|, \mu_*)$ as in Abbasi-Yadkori et al. (2019):

$$S_{\delta}(|\mathcal{A}|, \mu_*) := \sqrt{\frac{\log |\mathcal{A}|}{2}} + \langle \mu_*, \sqrt{\frac{1}{2} \log \frac{1}{\delta \mu_*}} \rangle.$$

We bound the regret of POLITEX in the following theorem. Here, recall that τ is the length of each phase, and γ and η are defined in Assumption 2.1 and Eq. (4.1), respectively.

Theorem 5.3 (Regret of POLITEX). Let Assumptions 2.1, 5.2, and 5.1 hold. For $\tau \geq \frac{\log T}{2 \log(1/\gamma)}$ and $\eta = \frac{\sqrt{8 \log |\mathcal{A}|}}{Q_{\max} \sqrt{K}}$, for a constant C_1 , with probability at least $1 - 4\delta$, the regret of POLITEX in ergodic average-reward MDPs is bounded as

$$R_T \leq \frac{C_1(1 + Q_{\max})S_{\delta}(|\mathcal{A}|, \mu_*)\sqrt{\tau}}{(1 - \gamma)^2} \sqrt{T}.$$

Proof. We start by decomposing the cumulative regret, following similar steps as Abbasi-Yadkori et al. (2019):

$$R_T = \sum_{k=1}^K \sum_{t=(k-1)\tau+1}^{k\tau} (J_* - J_{\pi_k}) + (J_{\pi_k} - r_t). \quad (5.1)$$

The second term $V_T = \sum_{k=1}^K \sum_{t=(k-1)\tau+1}^{k\tau} (J_{\pi_k} - r_t)$ captures the sum of differences between observed rewards and their long term averages. In previous work, this term was shown to scale as $O(K\sqrt{\tau})$. We show that the analysis can in fact be tightened to $O(\sqrt{T})$ using improved concentration bounds and the slow-changing nature of the policies. See Lemma B.1 in Appendix B for precise details.

The first term, which is also called *pseudo-regret* in literature, measures the difference between the expected reward

of the reference policy and the policies produced by the algorithm. Applying the performance difference lemma (Cao, 1999), we can write each pseudo-regret term as

$$J_* - J_{\pi_k} = \mathbb{E}_{x \sim \mu_*} [Q_{\pi_k}(x, \pi_*) - Q_{\pi_k}(x, \pi_k)].$$

Now, notice that POLITEX is running exact mirror descent for loss functions $\hat{Q}_{\pi_k}$. We bridge the pseudo-regret by the $\hat{Q}_{\pi_k}$ terms:

$$R_{T1a} = \tau \sum_{k=1}^K \mathbb{E}_{x \sim \mu_*} [Q_{\pi_k}(x, \pi_*) - \hat{Q}_{\pi_k}(x, \pi_*)] \quad (5.2)$$

$$R_{T1b} = \tau \sum_{k=1}^K \mathbb{E}_{x \sim \mu_*} [\hat{Q}_{\pi_k}(x, \pi_k) - Q_{\pi_k}(x, \pi_k)] \quad (5.3)$$

$$R_{T2} = \tau \sum_{k=1}^K \mathbb{E}_{x \sim \mu_*} [\hat{Q}_{\pi_k}(x, \pi_*) - \hat{Q}_{\pi_k}(x, \pi_k)]. \quad (5.4)$$

R_{T2} can be bounded using the regret of mirror descent as in previous work (Abbasi-Yadkori et al., 2019). Setting $\eta = \frac{\sqrt{\log |A|}}{Q_{\max} \tau \sqrt{2K}}$, and using a union bound over all states, with probability at least $1 - \delta$, R_{T2} is bounded as

$$R_{T2} \leq \tau Q_{\max} S_\delta(|A|, \mu_*) \sqrt{K}.$$

Bounding regret due to estimation error. We now focus on bounding $R_{T1} = R_{T1a} + R_{T1b}$ under Assumption 5.2. We have the following:

$$\begin{aligned} R_{T1a} &= \tau \sum_{k=1}^K \mathbb{E}_{x \sim \mu_*} [Q_{\pi_k}(x, \pi_*) - \hat{Q}_{\pi_k}(x, \pi_*)] \\ &= (\tau K) \mathbb{E}_{x \sim \mu_*, a \sim \pi_*} [Q_K(x, a) - \hat{Q}_K(x, a)] \\ &\leq C\tau \sqrt{K \log(1/\delta)}. \\ R_{T1b} &= \tau \sum_{k=1}^K \mathbb{E}_{x \sim \mu_*} [\hat{V}_{\pi_k}(x) - V_{\pi_k}(x)] \\ &= \tau K \mathbb{E}_{x \sim \mu_*} [\hat{V}_K(x) - V_K(x)] \\ &\leq C\tau \sqrt{K \log(1/\delta)}. \end{aligned}$$

We can then obtain the final result by combining the bounds on R_{T1a} , R_{T1b} , R_{T2} , V_T from Appendix B, and using union bound as well as the fact that $K = T/\tau$. $\square$

6. Linear Value Functions

In this section, we show that the estimation error condition in Assumption 5.2 (and thus $O(\sqrt{T})$ regret) can be achieved under similar assumptions as in Abbasi-Yadkori et al. (2019) and Wei et al. (2020a), which we state next.

Assumption 6.1 (Linear value functions). The action-value function Q_π of any policy π is linear: $Q_\pi(x, a) = w_\pi^\top \phi(x, a)$, where $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$ is a known feature function such that $\max_{x,a} \|\phi(x, a)\| \leq C_\Phi$.

Assumption 6.2 (Feature excitation). There exists a constant σ^2 such that for any policy π , $\lambda_{\min}(\mathbb{E}_{(x,a) \sim \mu_\pi \otimes \pi} [\phi(x, a) \phi(x, a)^\top]) \geq \sigma^2 > 0$.

We now describe a simple procedure for estimating the average action-value functions $Q_k(x, a) = \phi(x, a)^\top w_k$, where $w_k = \frac{1}{k} \sum_{i=1}^k w_{\pi_i}$, such that the conditions of Assumption 5.2 are satisfied. Essentially, we estimate each Q_{π_i} using least-squares Monte Carlo and then average the weights. We will use the shorthand notation $\phi_t = \phi(x_t, a_t)$. Let $\mathcal{H}_i$ and $\mathcal{T}_i$ be subsets of time indices in phase i (defined later). We estimate Q_{π_i} as follows:

$$\hat{w}_{\pi_i} = \left(\sum_{t \in \mathcal{H}_i} \phi_t \phi_t^\top + \alpha I \right)^{-1} \sum_{t \in \mathcal{H}_i} \phi_t R_t \quad (6.1)$$

where R_t are the empirical b -step returns (b is specified later), computed as

$$R_t = \sum_{j=t}^{t+b} (r_j - \hat{J}_{\pi_i}), \quad \hat{J}_{\pi_i} = \frac{1}{|\mathcal{T}_i|} \sum_{t \in \mathcal{T}_i} r_t. \quad (6.2)$$

We then estimate w_k as $\hat{w}_k = \frac{1}{k} \sum_{i=1}^k \hat{w}_{\pi_i}$. Note that for this special case of linear value functions, we do not need to use the replay buffer in Algorithm 1. For analysis purposes, we divide each phase of length τ into $2m$ blocks of size b and let $\mathcal{H}_i$ ($\mathcal{T}_i$) be the starting indices of odd (even) blocks in phase i . Due to the gaps between indices and fast mixing, this makes the data almost independent (we make this precise in Appendix C) and the error easier to analyze. In practice, one may simply want to use all data.

For a distribution μ , let $\|x\|_\mu$ denote the distribution-weighted norm such that $\|x\|_\mu^2 = \sum_i \mu_i x_i^2$. Using Jensen's inequality, we have that $(\mathbb{E}_{x \sim \mu} [q(x)])^2 \leq \mathbb{E}_{x \sim \mu} [q(x)^2]$. Thus, it suffices to bound the Q -function error in the distribution-weighted norm, $\|Q_K - \hat{Q}_K\|_{\mu_* \otimes \pi_*}$. Furthermore, given bounded features,

$$\|\hat{Q}_K - Q_K\|_{\mu_* \otimes \pi_*} \leq C_\Phi \|\hat{w}_K - w_K\|_2,$$

so it suffices to bound the error in the weights. We bound this error in the following Lemma, proven in Appendix C.

Lemma 6.3 (Estimation error for linear functions). Suppose that Assumptions 2.1, 6.1 and 6.2 hold and that true action-value weights are bounded as $\|w_{\pi_i}\|_2 \leq C_w$ for all $i = 1, \dots, K$. Then for any policy π , for $\alpha = \sqrt{\tau/K}$, $m \geq 72C_\Phi^4 \sigma^{-2} (1 - \gamma)^{-2} \log(d/\delta)$, and $b \geq \frac{\log(T\delta^{-1}(1-\gamma)^{-1})}{\log(1/\gamma)}$,

there exists an absolute constant c such that with probability at least $1 - \delta$,

$$\|\hat{w}_K - w_K\|_2 \leq c\sigma^{-2}(C_w + C_\Phi)b\sqrt{\frac{\log(2d/\delta)}{Km}}.$$

Furthermore, in Appendix D, we show that the error in the average state-value function satisfies the following:

$$\mathbb{E}_{\mu_*}[\hat{V}_K(x) - V_K(x)] \leq cC_\Phi|\mathcal{A}|(C_w + C_\Phi)\frac{b}{\sigma^2}\sqrt{\frac{\log(2d/\delta)}{Km}}.$$

We have demonstrated that Assumption 5.2 can be satisfied with linear value functions. For Assumption 5.1, it suffices for the weight estimates $\{\hat{w}_{\pi_i}\}$ to be bounded. This will be true, since we assume that the true weights are bounded and we can bound the error in the weight space. Thus POLITEX has an $O(\sqrt{T})$ regret in this setting (though note that we incur an extra $|\mathcal{A}|$ factor coming from the $\hat{V}_K$ error bound).

7. Practical Implementation

As mentioned earlier, the key idea in our policy update is to obtain an estimate $\hat{Q}_k$ of the average of all the Q-functions in previous phases. We have seen that when we use linear functions to approximate the Q-functions, we can simply average the weights in order to get an estimate of the average Q-function. However, in practice, we often need to use non-linear functions, especially neural networks, to approximate complex Q-functions. In this section, we discuss how to efficiently implement our algorithm with non-linear function approximation.

7.1. Weight Averaging

The simplest idea may be averaging the weights of neural networks. However, a crucial difference from the linear setting is that averaging the weights of the neural networks is not equivalent to averaging the functions that they represent: the function that a neural network represent is invariant to the permutation of the hidden units, and thus two networks with very different weights can represent similar functions. Therefore, this implementation may only succeed when all the Q-function approximations are around the same local region in the weight space. Thus, when we learn the new Q-function $\hat{Q}_{\pi_k}$ in phase k , we should initialize it with $\hat{Q}_{k-1}$, run SGD with new data, and then average with $\hat{Q}_{k-1}$.

7.2. Experience Replay

Another natural idea is to leverage the data from replay buffer to obtain an estimate of the average Q-function. We elaborate the details below. We use simple b -step Monte Carlo estimate for the Q value of each state-action pair. For any $(i-1)\tau + 1 \leq t \leq i\tau - b$, we can estimate the

state-action value of (x_t, a_t) by b -step cumulative reward¹

$$R_t = \sum_{j=t}^{t+b} (r_j - \hat{J}_{\pi_i}), \quad \hat{J}_{\pi_i} = \frac{1}{\tau} \sum_{j=(i-1)\tau+1}^{i\tau} r_j.$$

In the following, we denote by $\tau' := \tau - b$ the maximum number of data that we can collect from every phase. At the end of each phase, we store all or a subset of the (x_t, a_t, R_t) tuples in our replay buffer $\mathcal{R}$. We extract feature $\phi(x, a) \in \mathbb{R}^d$ for the state-action pair (x, a) and let $\mathcal{F} \subseteq \{f : \mathbb{R}^d \mapsto \mathbb{R}\}$ be a class of functions that we use to approximate the Q-functions. For phase i , we propose the following method to estimate Q_{π_i} : $\hat{Q}_{\pi_i}(x, a) = \hat{f}(\phi(x, a))$, where $\hat{f} \in \arg \min_{f \in \mathcal{F}} \ell_i(f)$ and

$$\ell_i(f) := \frac{1}{\tau'} \sum_{t=(i-1)\tau+1}^{i\tau-b} (f(\phi(x_t, a_t)) - R_t)^2. \quad (7.1)$$

Suppose that we store all the data from the previous phases in the replay buffer, then in order to estimate the average of the Q-functions of the first k phases, i.e., $\hat{Q}_k$, we propose to use the heuristic that minimizes the average of the k squared loss functions defined in Eq. (7.1), i.e., $\frac{1}{k} \sum_{i=1}^k \ell_i(f)$.

Subsampling and coreset. In practice, due to the high memory cost, it may be hard to store all the data from the previous phases in the replay buffer. We found that a simple strategy to resolve this issue is to begin with storing all the data from every phase, and add a limit on size of the replay buffer. When the buffer size exceeds the limit, we eliminate a subset of the data uniformly at random.

Another approach is to sample a subset of size s from the data collected in each phase. Denote this subset by $\mathcal{R}_i$ for phase i . Thus in the k -th phase, we have τ' data $\mathcal{D}_k$ from the current phase as well as $s(k-1)$ data from the replay buffer $\mathcal{R}$. First, suppose that the s data points are sampled uniformly at random. We can then minimize the following objective:

$$\min_{f \in \mathcal{F}} \frac{1}{k} \left(\ell_i(f) + \sum_{i=1}^{k-1} \hat{\ell}_i(f) \right), \quad (7.2)$$

where $\hat{\ell}_i(f) := \frac{1}{s} \sum_{(x_t, a_t, R_t) \in \mathcal{R}_i} (f(\phi(x_t, a_t)) - R_t)^2$ is an unbiased estimate of $\ell_i(f)$. Further, uniform sampling is not the only way to construct an unbiased estimate. In fact, for any discrete distribution, with PMF $q = \{q_t\}$, over the τ' data in $\mathcal{D}_i$, we can sample s data points according to q and construct

$$\hat{\ell}_i(f) := \frac{1}{\tau'} \sum_{(x_t, a_t, R_t) \in \mathcal{R}_i} \frac{1}{q_t} (f(\phi(x_t, a_t)) - R_t)^2, \quad (7.3)$$

¹This is a practical implementation of Eq. (6.2), i.e., we do not split the data in each phase into blocks.

in order to obtain an unbiased estimate of $\ell_i(f)$. As shown by Bachem et al. (2017), by choosing $q_t \propto (f(\phi(x_t, a_t)) - R_t)^2$, we can minimize the variance of $\widehat{\ell}_i(f)$ for any fixed f . In the following, we call the subset of data sampled according to this distribution a *coreset* of the data. In the experiments in Section 8, we show that thanks to the variance reduction effect, sampling a coreset often produces better performance than sampling a subset uniformly at random, especially when the rewards are sparse.

Comparison to Abbasi-Yadkori et al. (2019). Our algorithm can be considered as an efficient implementation of POLITEX (Abbasi-Yadkori et al., 2019) via experience replay. In the original POLITEX algorithm, the Q-functions are estimated using Eq. (7.1) for each phase, and all the functions are stored in memory. When an agent interacts with the environment, it needs to evaluate all the Q-functions in order to obtain the probability of each action. This implies that the time complexity of computing action probabilities increases with the number of phases, and as a result the algorithm is hard to scale to a large number of phases. Although we can choose to evaluate a random subset of the Q-functions to estimate the action probabilities in POLITEX, our implementation via experience replay can still be faster since we only need to evaluate a single function, trained with replay data, to take actions.

8. Experiments

In this section, we evaluate our implementations empirically. We make comparisons with several baselines in two control environments.

Environments. We use two control environments with the simulators described in Tassa et al. (2018). Both environments are episodic with episode length 1000. The environments we evaluate are:

- *Cart-pole* (Barto et al., 1983): The goal of this environment is to balance an unactuated pole by applying forces to a cart at its base. We discretize the continuous force to 5 values: $\{-2, -1, 0, 1, 2\}$. The reward at each time step is a real number in $[0, 1]$. We end episodes early when the pole falls (we use rewards less than 0.5 as an indicator for falling), and assume zero reward for the remaining steps when reporting results.
- *Ball-in-cup*: Here, an actuated planar receptacle can translate in the vertical plane in order to swing and catch a ball attached to its bottom. This task has a sparse reward: 1 when the ball is in the cup, and 0 otherwise. We discretize the two-dimensional continuous action space to a 3×3 grid, i.e., 9 actions in total.

Algorithms. We compare the following algorithms: our proposed implementation of POLITEX using experience re-

play, POLITEX using weight averaging, the original POLITEX algorithm, the variant of POLITEX that averages 10 randomly selected Q-functions, the mirror descent policy optimization (MDPO) algorithm (Tomar et al., 2020), the constrained policy optimization (CPO) (Achiam et al., 2017), and the V-MPO algorithm (Song et al., 2019b).

For all the algorithms, we extract Fourier basis features (Konidaris et al., 2011) from the raw observations for both environments: for Cart-pole, we use 4 bases and for Ball-in-cup we use 2 bases. For variants of POLITEX we construct the state-action features $\phi(x, a)$ by block one-hot encoding, i.e., we partition the $\phi(x, a)$ vector into $|\mathcal{A}|$ blocks, and set the a -th block to be the Fourier features of x and other blocks to be zero. We approximate Q-functions using neural networks with one hidden layer and ReLU activation: the width of the hidden layer is 50 for Cart-pole and 250 for Ball-in-cup. For MDPO, CPO and V-MPO, the algorithms use a policy network whose input and output are the state feature and action probability, respectively. We use one hidden layer networks with width 50 for Cart-pole and 250 for Ball-in-cup. These three algorithms also need a value network, which takes the state feature as input and outputs its value for the current policy. For both environments, we use one hidden layer network with width 50.

For Cart-pole, we choose phase length $\tau = 10^4$ and for Ball-in-cup, we choose $\tau = 2 \times 10^4$. Since the environments are episodic, we have multiple episodes in each phase. For Cart-pole, since the training usually does not need a large number of phases, we do not set limit on the size of the replay buffer, whereas for Ball-in-cup, we set the limit as 2×10^6 . For our algorithm and the variants of POLITEX, we choose parameter η in $\{5, 10, 20, 40, 80, 160\}$ and report the result of each algorithm using the value of η that produces the best average reward. For MDPO, CPO and V-MPO, we note that according to Eq. (4.1), the KL regularization coefficient between policies is the reciprocal of the η parameter in our algorithm, we also run these baseline algorithms in the same range of η and report the best result. In addition, in CPO, we set the limit in KL divergence between adjacent policies 0.001η , which, in our experiments, leads to good performance. We treat the length of Monte Carlo estimate b as a hyper parameter, so we choose it in $\{100, 300\}$ and report the best performance.

Results. We run each algorithm in each environment 20 times and report the average results as well as the standard deviation of the 20 runs (shaded area). We report the average reward in each phase before training on the data collected during the phase. The results are presented in Figure 1. Every run is conducted on a single P100 GPU.

From Fig. 1(a), we can see that the iteration complexity (average reward vs number of phases) and final performance of the experience replay implementation using all the data from




Figure 1. Experiments on the Cart-pole environment (top) and Ball-in-cup (bottom).

past phases (blue curve) are similar to many state-of-the-art algorithms, such as POLITEX and V-MPO. Meanwhile, according to Fig. 1(b), the experience replay implementation achieves strong performance in training speed, i.e., best performance at 30 minute in Cart-pole and second best performance at 200 minute in Ball-in-cup. We note here that the training time includes both the time for data collection, i.e., interacting with the environment simulator and the training of a new policy. The observation that the experience replay based implementation is faster than the original POLITEX and the implementation that uses 10 Q-functions can be explained by the fact that the replay-based algorithm only uses a single Q-function and thus is faster at data collection, as discussed in Section 7.2. In addition, that the replay-based algorithm runs faster than MDPO, CPO, and V-MPO can be explained by its simplicity: in variants of POLITEX, we only need to approximate the Q-function, whereas in MDPO, CPO, and V-MPO, we need to train both the policy and value networks.

The POLITEX weight averaging scheme achieves the best performance on Ball-in-cup in both iteration complexity and training speed, but does not converge to a solution that is sufficiently close to the optimum on Cart-pole. Notice that for weight averaging, we used the initialization technique mentioned in Section 7.1. Considering the consistent strong performance of experience replay in both environments, we still recommend applying experience replay when using non-linear function approximation.

In Fig. 1(c), we can see that when we only sample a small subset (1% or 4%) of the data to store in the replay buffer,

the coreset technique described in Section 7.2 achieves better performance due to the variance reduction effect. This improvement is quite significant in the Ball-in-cup environment, which has sparse rewards. Notice that sampling coreset only adds negligible computational overhead during training, and thus we recommend the coreset technique when there exists a strict memory constraint.

9. Discussion

The main contributions of our work are an improved analysis and practical implementation of POLITEX. On the theoretical side, we show that POLITEX obtains an $O(\sqrt{T})$ high-probability regret bound in uniformly mixing average-reward MDPs with linear function approximation, which is the first such bound for a computationally efficient algorithm. The main limitation of these result is that, similarly to previous works, they hold under somewhat strong assumptions which circumvent the need for exploration. An interesting future work direction would be to relax these assumptions and incorporate explicit exploration instead.

On the practical side, we propose an efficient implementation with neural networks that relies on experience replay, a standard tool in modern deep RL. Our work shows that experience replay and KL regularization can both be viewed as approximately implementing mirror descent policy updates within a policy iteration scheme. This provides an online-learning justification for using replay buffers in policy iteration, which is different than the standard explanation for their success. Our work also suggests a new objective

for storing and prioritizing replay samples, with the goal of approximating the average of value functions well. This goal has some similarities with continual learning, where experience replay has also lead to empirical successes. One interesting direction for future work would be to explore other continual learning techniques in the approximate implementation of mirror descent policy updates.

References

- Abbasi-Yadkori, Y., Bartlett, P., Bhatia, K., Lazic, N., Szepesvari, C., and Weisz, G. POLITEX: Regret bounds for policy iteration using expert prediction. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 3692–3702. PMLR, 09–15 Jun 2019.
- Abdolmaleki, A., Springenberg, J. T., Tassa, Y., Munos, R., Heess, N., and Riedmiller, M. Maximum a posteriori policy optimisation. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=SlANxQW0b>.
- Achiam, J., Held, D., Tamar, A., and Abbeel, P. Constrained policy optimization. In *International Conference on Machine Learning*, pp. 22–31, 2017.
- Bachem, O., Lucic, M., and Krause, A. Practical core-set constructions for machine learning. *arXiv preprint arXiv:1703.06476*, 2017.
- Bartlett, P. L. Regal: A regularization based algorithm for reinforcement learning in weakly communicating mdps. In *In Proceedings of the 25th Annual Conference on Uncertainty in Artificial Intelligence*, 2009.
- Barto, A. G., Sutton, R. S., and Anderson, C. W. Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE Transactions on Systems, Man, and Cybernetics*, (5):834–846, 1983.
- Borsos, Z., Mutn , M., and Krause, A. Coresets via bilevel optimization for continual learning and streaming. *arXiv preprint arXiv:2006.03875*, 2020.
- Cao, X.-R. Single sample path-based optimization of Markov chains. *Journal of optimization theory and applications*, 100(3):527–548, 1999.
- Cesa-Bianchi, N. and Lugosi, G. *Prediction, learning, and games*. Cambridge university press, 2006.
- Chaudhry, A., Ranzato, M., Rohrbach, M., and Elhoseiny, M. Efficient lifelong learning with a-gem. *arXiv preprint arXiv:1812.00420*, 2018.
- Cho, G. E. and Meyer, C. D. Comparison of perturbation bounds for the stationary distribution of a Markov chain. *Linear Algebra and its Applications*, 335(1-3):137–150, 2001.
- Degrave, J., Abdolmaleki, A., Springenberg, J. T., Heess, N., and Riedmiller, M. Quinoa: a Q-function you infer normalized over actions. *arXiv preprint arXiv:1911.01831*, 2019.
- Even-Dar, E., Kakade, S. M., and Mansour, Y. Online Markov decision processes. *Mathematics of Operations Research*, 34(3):726–736, 2009.
- Farajtabar, M., Azizan, N., Mott, A., and Li, A. Orthogonal gradient descent for continual learning. In *International Conference on Artificial Intelligence and Statistics*, pp. 3762–3773, 2020.
- Fruit, R., Pirotta, M., Lazaric, A., and Ortner, R. Efficient bias-span-constrained exploration-exploitation in reinforcement learning. In *International Conference on Machine Learning*, 2018.
- Hao, B., Lazic, N., Abbasi-Yadkori, Y., Joulani, P., and Szepesvari, C. Provably efficient adaptive approximate policy iteration. *arXiv preprint arXiv:2002.03069*, 2020.
- Heess, N., Wayne, G., Silver, D., Lillicrap, T., Erez, T., and Tassa, Y. Learning continuous control policies by stochastic value gradients. *Advances in Neural Information Processing Systems*, 28:2944–2952, 2015.
- Jaksch, T., Ortner, R., and Auer, P. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(Apr):1563–1600, 2010.
- Jian, Q., Fruit, R., Pirotta, M., and Lazaric, A. Exploration bonus for regret minimization in discrete and continuous average reward mdps. In *Advances in Neural Information Processing Systems*, pp. 4891–4900, 2019.
- Jin, C., Netrapalli, P., Ge, R., Kakade, S. M., and Jordan, M. I. A short note on concentration inequalities for random vectors with subgaussian norm. *arXiv preprint arXiv:1902.03736*, 2019.
- Kakade, S. and Langford, J. Approximately optimal approximate reinforcement learning. In *ICML*, volume 2, pp. 267–274, 2002.
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
- Konidaris, G., Osentoski, S., and Thomas, P. Value function approximation in reinforcement learning using the fourier

- basis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 25, 2011.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. In *ICLR (Poster)*, 2016.
- Lin, L.-J. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine learning*, 8(3-4):293–321, 1992.
- Lopez-Paz, D. and Ranzato, M. Gradient episodic memory for continual learning. In *Advances in neural information processing systems*, pp. 6467–6476, 2017.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *nature*, 518(7540): 529–533, 2015.
- Neu, G., György, A., Szepesvári, C., and Antos, A. Online Markov decision processes under bandit feedback. *IEEE Transactions on Automatic Control*, 59(3):676–691, December 2014.
- Ouyang, Y., Gagrani, M., Nayyar, A., and Jain, R. Learning unknown Markov decision processes: A Thompson sampling approach. In *Advances in Neural Information Processing Systems*, pp. 1333–1342, 2017.
- Rusu, A. A., Rabinowitz, N. C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., Pascanu, R., and Hadsell, R. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016.
- Schaul, T., Quan, J., Antonoglou, I., and Silver, D. Prioritized experience replay. In *International Conference on Learning Representations (ICLR)*, 2016.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. Trust region policy optimization. In *International conference on machine learning*, pp. 1889–1897, 2015.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Seneta, E. Coefficients of ergodicity: structure and applications. *Advances in applied probability*, pp. 576–590, 1979.
- Seneta, E. Perturbation of the stationary distribution measured by ergodicity coefficients. *Advances in Applied Probability*, 20(1):228–230, 1988.
- Song, H. F., Abdolmaleki, A., Springenberg, J. T., Clark, A., Soyer, H., Rae, J. W., Noury, S., Ahuja, A., Liu, S., Tirumala, D., Heess, N., Belov, D., Riedmiller, M. A., and Botvinick, M. M. V-MPO: on-policy maximum a posteriori policy optimization for discrete and continuous control. *CoRR*, abs/1909.12238, 2019a. URL <http://arxiv.org/abs/1909.12238>.
- Song, H. F., Abdolmaleki, A., Springenberg, J. T., Clark, A., Soyer, H., Rae, J. W., Noury, S., Ahuja, A., Liu, S., Tirumala, D., et al. V-mpo: on-policy maximum a posteriori policy optimization for discrete and continuous control. *arXiv preprint arXiv:1909.12238*, 2019b.
- Talebi, M. S. and Maillard, O.-A. Variance-aware regret bounds for undiscounted reinforcement learning in mdps. In *Algorithmic Learning Theory*, 2018.
- Tassa, Y., Doron, Y., Muldal, A., Erez, T., Li, Y., Casas, D. d. L., Budden, D., Abdolmaleki, A., Merel, J., Lefrancq, A., et al. DeepMind control suite. *arXiv preprint arXiv:1801.00690*, 2018.
- Tomar, M., Shani, L., Efroni, Y., and Ghavamzadeh, M. Mirror descent policy optimization. *arXiv preprint arXiv:2005.09814*, 2020.
- Tropp, J. A. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12(4):389–434, 2012.
- Vieillard, N., Kozuno, T., Scherrer, B., Pietquin, O., Munos, R., and Geist, M. Leverage the average: an analysis of regularization in rl. *arXiv preprint arXiv:2003.14089*, 2020a.
- Vieillard, N., Scherrer, B., Pietquin, O., and Geist, M. Momentum in reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, pp. 2529–2538, 2020b.
- Wei, C.-Y., Jafarnia-Jahromi, M., Luo, H., and Jain, R. Learning infinite-horizon average-reward mdps with linear function approximation. *arXiv preprint arXiv:2007.11849*, 2020a.
- Wei, C.-Y., Jahromi, M. J., Luo, H., Sharma, H., and Jain, R. Model-free reinforcement learning in infinite-horizon average-reward Markov decision processes. In *International Conference on Machine Learning*, pp. 10170–10180. PMLR, 2020b.
- Yin, D., Farajtabar, M., and Li, A. SOLA: Continual learning with second-order loss approximation. *arXiv preprint arXiv:2006.10974*, 2020.

Yu, B. Rates of convergence for empirical processes of stationary mixing sequences. *The Annals of Probability*, 22(1):94–116, 1994.

Zenke, F., Poole, B., and Ganguli, S. Continual learning through synaptic intelligence. *Proceedings of machine learning research*, 70:3987, 2017.

Appendix

A. Preliminaries

We state some useful definitions and lemmas in this section.

Lemma A.1. Let X_1 and X_2 be a pair of distribution vectors. Let H be the transition matrix of an ergodic Markov chain with a stationary distribution ν , and ergodicity coefficient (defined in Assumption 2.1) upper-bounded by $\gamma < 1$. Then

$$\|(H^m)^\top (X_1 - X_2)\|_1 \leq \gamma^m \|X_1 - X_2\|_1.$$

Proof. Let $\{v_1, \dots, v_n\}$ be the normalized left eigenvectors of H corresponding to ordered eigenvalues $\{\lambda_1, \dots, \lambda_n\}$. Then $v_1 = \nu$, $\lambda_1 = 1$, and for all $i \geq 2$, we have that $\lambda_i < 1$ (since the chain is ergodic) and $v_i^\top \mathbf{1} = 0$. Write X_1 in terms of the eigenvector basis as:

$$X_1 = \alpha_1 \nu + \sum_{i=2}^n \alpha_i v_i \quad \text{and} \quad X_2 = \beta_1 \nu + \sum_{i=2}^n \beta_i v_i.$$

Since $X_1^\top \mathbf{1} = 1$ and $X_2^\top \mathbf{1} = 1$, it is easy to see that $\alpha_1 = \beta_1 = 1$. Thus we have

$$\|H^\top (X_1 - X_2)\|_1 = \|H^\top \sum_{i=2}^n (\alpha_i - \beta_i) v_i\|_1 \leq \gamma \left\| \sum_{i=2}^n (\alpha_i - \beta_i) v_i \right\|_1 = \gamma \|X_1 - X_2\|_1$$

where the inequality follows from the definition of the ergodicity coefficient and the fact that $\mathbf{1}^\top v_i = 0$ for all $i \geq 2$. Since

$$\mathbf{1}^\top H^\top \sum_{i=2}^n (\alpha_i - \beta_i) v_i = \mathbf{1}^\top \sum_{i=2}^n \lambda_i (\alpha_i - \beta_i) v_i = 0,$$

the inequality also holds for powers of H . $\square$

Lemma A.2 (Doob martingale). Let Assumption 2.1 hold, and let $\{(x_t, a_t)\}_{t=1}^T$ be the state-action sequence obtained when following policies $\pi_1, \dots, \pi_k$ for τ steps each from an initial distribution ν_0 . For $t \in [T]$, let X_t be a binary indicator vector with a non-zero element at the linear index of the state-action pair (x_t, a_t) . Define for $i \in [T]$,

$$B_i = \mathbb{E} \left[\sum_{t=1}^T X_t | X_1, \dots, X_i \right], \quad \text{and} \quad B_0 = \mathbb{E} \left[\sum_{t=1}^T X_t \right].$$

Then, $\{B_i\}_{i=0}^T$ is a vector-valued martingale: $\mathbb{E}[B_i - B_{i-1} | B_0, \dots, B_{i-1}] = 0$ for $i = 1, \dots, T$, and $\|B_i - B_{i-1}\|_1 \leq 2(1 - \gamma)^{-1}$ holds for $i \in [T]$.

The constructed martingale is known as the Doob martingale underlying the sum $\sum_{t=1}^T X_t$.

Proof. That $\{B_i\}_{i=0}^T$ is a martingale follows from the definition. We now bound its difference sequence. Let H_t be the state-action transition matrix at time t , and let $H_{i:t} = \prod_{j=i}^{t-1} H_j$, and define $H_{i:i} = I$. Then, for $t = 0, \dots, T-1$, $\mathbb{E}[X_{t+1} | X_t] = H_t^\top X_t$ and by the Markov property, for any $i \in [T]$,

$$B_i = \sum_{t=1}^i X_t + \sum_{t=i+1}^T \mathbb{E}[X_t | X_i] = \sum_{t=1}^i X_t + \sum_{t=i+1}^T H_{i:t}^\top X_t, \quad \text{and} \quad B_0 = \sum_{t=1}^T H_{0:t}^\top X_0.$$

For any $i \in [T]$,

$$\begin{aligned} B_i - B_{i-1} &= \sum_{t=1}^i X_t - \sum_{t=1}^{i-1} X_t + \sum_{t=i+1}^T H_{i:t}^\top X_t - \sum_{t=i}^T H_{i-1:t}^\top X_{i-1} \\ &= \sum_{t=i}^T H_{i:t}^\top (X_t - H_{i-1:t}^\top X_{i-1}). \end{aligned} \tag{A.1}$$

Since X_i and $H_{i-1}^\top X_{i-1}$ are distribution vectors, under Assumption 2.1 and using Lemma A.1,

$$\|B_i - B_{i-1}\|_1 \leq \sum_{t=i}^T \|H_{i:t}^\top (X_i - H_{i-1}^\top X_{i-1})\|_1 \leq 2 \sum_{j=0}^{T-i} \gamma^j \leq 2(1-\gamma)^{-1}.$$

□

Let $(\mathcal{F}_k)_k$ be a filtration and define $\mathbb{E}_k[\cdot] := \mathbb{E}[\cdot | \mathcal{F}_k]$. We will make use of the following concentration results for the sum of random matrices and vectors.

Theorem A.3 (Matrix Azuma, [Tropp \(2012\)](#) Thm 7.1). Consider a finite $(\mathcal{F})_k$ -adapted sequence $\{X_k\}$ of Hermitian matrices of dimension m , and a fixed sequence $\{A_k\}$ of Hermitian matrices that satisfy $\mathbb{E}_{k-1} X_k = 0$ and $X_k^2 \preceq A_k^2$ almost surely. Let $v = \|\sum_k A_k^2\|$. Then with probability at least $1 - \delta$, $\|\sum_k X_k\|_2 \leq 2\sqrt{2v \ln(m/\delta)}$.

A version of Theorem A.3 for non-Hermitian matrices of dimension $m_1 \times m_2$ can be obtained by applying the theorem to a Hermitian dilation of X , $\mathcal{D}(X) = \begin{bmatrix} 0 & X \\ X^* & 0 \end{bmatrix}$, which satisfies $\lambda_{\max}(\mathcal{D}(X)) = \|X\|$ and $\mathcal{D}(X)^2 = \begin{bmatrix} XX^* & 0 \\ 0 & X^*X \end{bmatrix}$. In this case, we have that $v = \max(\|\sum_k X_k X_k^*\|, \|\sum_k X_k^* X_k\|)$.

Lemma A.4 (Hoeffding-type inequality for norm-subGaussian random vectors, [Jin et al. \(2019\)](#)). Consider random vectors $X_1, \dots, X_n \in \mathbb{R}^d$ and corresponding filtrations $\mathcal{F}_i = \sigma(X_1, \dots, X_i)$ $i \in [n]$, such that $X_i | \mathcal{F}_{i-1}$ is zero-mean norm-subGaussian with $\sigma_i \in \mathcal{F}_{i-1}$. That is:

$$\mathbb{E}[X_i | \mathcal{F}_i] = 0, \quad P(\|X_i\| \geq t | \mathcal{F}_{i-1}) \leq 2 \exp(-t^2/2\sigma_i^2) \quad \forall t \in \mathbb{R}, \forall i \in [n].$$

If the condition is satisfied for fixed $\{\sigma_i\}$, there exists a constant c such that for any $\delta > 0$, with probability at least $1 - \delta$,

$$\left\| \sum_{i=1}^n X_i \right\| \leq c \sqrt{\sum_{i=1}^n \sigma_i^2 \log(2d/\delta)}.$$

B. Bounding the Difference Between Empirical and Average Rewards

In this section, we bound the second term in Equation 5.1, corresponding to the difference between empirical and average rewards.

Lemma B.1. Let Assumption 2.1 hold, and assume that $\tau \geq \frac{\log T}{2 \log(1/\gamma)}$ and that $r(x, a) \in [0, 1]$ for all x, a . Then, by choosing $\eta = \frac{\sqrt{8 \log |\mathcal{A}|}}{Q_{\max} \sqrt{K}}$, we have with probability at least $1 - \delta$,

$$\sum_{k=1}^K \sum_{t=(k-1)\tau+1}^{k\tau} (r_t - J_{\pi_k}) \leq 2(1-\gamma)^{-1} \sqrt{2T \log(2/\delta)} + 2\sqrt{T} + (1-\gamma)^{-2} \sqrt{8K \log |\mathcal{A}|}.$$

Proof. Let r denote the vector of rewards, and recall that $J_\pi = \nu_\pi^\top r$. Let X_t be the indicator vector for the state-action pair at time t , as in Lemma A.2, and let $\nu_t = \mathbb{E}[X_t]$. We have the following:

$$V_T := \sum_{k=1}^K \sum_{t=(k-1)\tau+1}^{k\tau} (r_t - J_{\pi_k}) = \sum_{k=1}^K \sum_{t=(k-1)\tau+1}^{k\tau} r^\top (X_t - \nu_t + \nu_t - \nu_{\pi_k})$$

We slightly abuse the notation above by letting ν_t denote the state-action distribution at time t , and ν_π the stationary distribution of policy π . Let $\{B_i\}_{i=0}^T$ be the Doob martingale in Lemma A.2. Then $B_0 = \sum_{t=1}^T \nu_t$ and $B_T = \sum_{t=1}^T X_t$, and the first term can be expressed as

$$V_{T1} := \sum_{t=1}^T r^\top (X_t - \nu_t) = r^\top (B_T - B_0).$$

By Lemma A.2, $|\langle B_i - B_{i-1}, r \rangle| \leq \|B_i - B_{i-1}\|_1 \|r\|_\infty \leq 2(1 - \gamma)^{-1}$. Hence by Azuma's inequality, with probability at least $1 - \delta$,

$$V_{T1} \leq 2(1 - \gamma)^{-1} \sqrt{2T \log(2/\delta)}. \quad (\text{B.1})$$

For the second term we have

$$\begin{aligned} V_{T2} &:= \sum_{k=1}^K \sum_{t=(k-1)\tau+1}^{k\tau} r^\top (\nu_t - \nu_{\pi_k}) \\ &= \sum_{k=1}^K r^\top \left(\sum_{i=1}^{\tau} (H_{\pi_k}^i)^\top \nu_{(k-1)\tau} - \nu_{\pi_k} \right) \\ &\leq \sum_{k=1}^K \|r\|_\infty \sum_{i=1}^{\tau} \left\| (H_{\pi_k}^i)^\top (\nu_{(k-1)\tau} - \nu_{\pi_{k-1}} + \nu_{\pi_{k-1}}) - \nu_{\pi_k} \right\|_1 \\ &\leq \sum_{k=1}^K \sum_{i=1}^{\tau} \left\| \nu_{(k-1)\tau} - \nu_{\pi_{k-1}} \right\|_1 + \left\| (H_{\pi_k}^i)^\top \nu_{\pi_{k-1}} - \nu_{\pi_k} \right\|_1 \\ &\leq \sum_{k=1}^K \sum_{i=1}^{\tau} \left\| (H_{\pi_{(k-1)}}^\tau)^\top \nu_{(k-2)\tau} - \nu_{\pi_{k-1}} \right\|_1 + \gamma^i \left\| \nu_{\pi_{k-1}} - \nu_{\pi_k} \right\|_1 \\ &\leq 2T\gamma^\tau + \frac{1}{1 - \gamma} \sum_{k=1}^K \left\| \nu_{\pi_k} - \nu_{\pi_{k-1}} \right\|_1. \end{aligned}$$

For $\tau \geq \frac{\log T}{2 \log(1/\gamma)}$, the first term is upper-bounded by $2\sqrt{T}$.

Using results on perturbations of Markov chains (Seneta, 1988; Cho & Meyer, 2001), we have that

$$\left\| \nu_{\pi_k} - \nu_{\pi_{k-1}} \right\|_1 \leq \frac{1}{1 - \gamma} \|H_{\pi_k} - H_{\pi_{k-1}}\|_\infty \leq \frac{1}{1 - \gamma} \max_x \left\| \pi_k(\cdot|x) - \pi_{k-1}(\cdot|x) \right\|_1$$

Note that the policies $\pi_k(\cdot|x)$ are generated by running mirror descent on reward functions $\hat{Q}_{\pi_k}(x, \cdot)$. A well-known property of mirror descent updates with entropy regularization (or equivalently, the exponentially-weighted-average algorithm) is that the difference between consecutive policies is bounded as

$$\left\| \pi_{k+1}(\cdot|x) - \pi_k(\cdot|x) \right\|_1 \leq \eta \left\| \hat{Q}_{\pi_k}(x, \cdot) \right\|_\infty.$$

See e.g. Neu et al. (2014) Section V.A for a proof, which involves applying Pinsker's inequality and Hoeffding's lemma (Cesa-Bianchi & Lugosi (2006) Section A.2 and Lemma A.6). Since we assume that $\left\| \hat{Q}_{\pi_k} \right\|_\infty \leq Q_{\max}$, we can obtain

$$V_{T2} \leq 2\sqrt{T} + (1 - \gamma)^{-2} K \eta Q_{\max}.$$

By choosing $\eta = \frac{\sqrt{8 \log |\mathcal{A}|}}{Q_{\max} \sqrt{K}}$, we can bound the second term as

$$V_{T2} \leq 2\sqrt{T} + (1 - \gamma)^{-2} \sqrt{8K \log |\mathcal{A}|}. \quad (\text{B.2})$$

Putting Eq. (B.1) and (B.2) together, we obtain that with probability at least $1 - \delta$,

$$V_T \leq 2(1 - \gamma)^{-1} \sqrt{2T \log(2/\delta)} + 2\sqrt{T} + (1 - \gamma)^{-2} \sqrt{8K \log |\mathcal{A}|}.$$

□

C. Proof of Lemma 6.3

Proof. Recall that we split each phase into $2m$ blocks of size b and let $\mathcal{H}_i$ and $\mathcal{T}_i$ denote the starting indices of odd and even blocks, respectively. We let R_t denote the empirical b -step returns from the state action pair (x_t, a_t) in phase i :

$$R_t = \sum_{i=t}^{t+b} (r_i - \hat{J}_{\pi_i}), \quad \hat{J}_{\pi_i} = \frac{1}{|\mathcal{T}_i|} \sum_{t \in \mathcal{T}_i} r_t.$$

We start by bounding the error in R_t . Let X be a binary indicator vector for a state-action pair (x, a) . Let H_π be the state-action transition kernel for policy π , and let ν_π be the corresponding stationary state-action distribution. We can write the action-value function at (x, a) as

$$\begin{aligned} Q_\pi(x, a) &= r(x, a) - J_\pi + X^\top H_\pi Q_\pi \\ &= (X - \nu_\pi)^\top r + X^\top H_\pi (r - J_\pi \mathbf{1} + H_\pi Q_\pi) \\ &= \sum_{i=0}^{\infty} (X - \nu_\pi)^\top H_\pi^i r. \end{aligned}$$

Let $Q_\pi^b(x, a) = \sum_{i=0}^b (X - \nu_\pi)^\top H_\pi^i r$ be a version of Q_π truncated to b steps. Under uniform mixing, the difference to the true Q_π is bounded as

$$|Q_\pi(x, a) - Q_\pi^b(x, a)| \leq \sum_{i=1}^{\infty} |(X - \nu_\pi)^\top H_\pi^{i+b} r| \leq \frac{2\gamma^{b+1}}{1-\gamma}. \quad (\text{C.1})$$

Let $b_t = Q_{\pi_i}^b(x_t, a_t) - Q_{\pi_i}(x_t, a_t)$ denote the truncation bias at time t , and let $z_t = \sum_{i=t}^{t+b} r_i - X_t^\top H_{\pi_i}^{(i-t)} r$ denote the reward noise. We will write

$$R_t = Q_{\pi_i}(x_t, a_t) + b(J_{\pi_i} - \hat{J}_{\pi_i}) + z_t + b_t.$$

Note that $m = |\mathcal{H}_i|$ and let

$$\widehat{M}_i = \frac{1}{m} \sum_{t \in \mathcal{H}_i} \phi_t \phi_t^\top + \frac{\alpha}{m} I.$$

We estimate the value function of each policy π_i using data from phase i as

$$\begin{aligned} \widehat{w}_{\pi_i} &= \widehat{M}_i^{-1} m^{-1} \sum_{t \in \mathcal{H}_i} \phi_t R_t \\ &= \widehat{M}_i^{-1} m^{-1} \sum_{t \in \mathcal{H}_i} \phi_t (\phi_t^\top w_{\pi_i} + b_t + z_t + b(J_{\pi_i} - \hat{J}_{\pi_i})) + \widehat{M}_i^{-1} \frac{\alpha}{m} (w_{\pi_i} - w_{\pi_i}) \\ &= w_{\pi_i} + \widehat{M}_i^{-1} m^{-1} \sum_{t \in \mathcal{H}_i} \phi_t (z_t + b_t + b(J_{\pi_i} - \hat{J}_{\pi_i})) - \widehat{M}_i^{-1} m^{-1} \alpha w_{\pi_i} \end{aligned}$$

Our estimate $\widehat{w}_k$ of $w_k = \frac{1}{k} \sum_{i=1}^k w_{\pi_i}$ can thus be written as follows:

$$\widehat{w}_k - w_k = \frac{1}{km} \sum_{i=1}^k \sum_{t \in \mathcal{H}_i} \widehat{M}_i^{-1} \phi_t (z_t + b_t + b(J_{\pi_i} - \hat{J}_{\pi_i})) - \frac{\alpha}{km} \sum_{i=1}^k \widehat{M}_i^{-1} w_{\pi_i}.$$

We proceed to upper-bound the norm of the RHS.

Set $\alpha = \sqrt{m/k}$. Let C_w be an upper-bound on the norm of the true value-function weights $\|w_{\pi_i}\|_2$ for $i = 1, \dots, K$. In Appendix C.3, we show that with probability at least $1 - \delta$, for $m \geq 72C_\Phi^4 \sigma^{-2} (1 - \gamma)^{-2} \log(d/\delta)$, $\|\widehat{M}_i^{-1}\|_2 \leq 2\sigma^{-2}$. Thus with probability at least $1 - \delta$, the last error term is upper-bounded as

$$\frac{\alpha}{km} \left\| \sum_{i=1}^k \widehat{M}_i^{-1} w_{\pi_i} \right\|_2 \leq 2\sigma^{-2} C_w (km)^{-1/2}. \quad (\text{C.2})$$

Similarly, for

$$b \geq \frac{\log((1-\gamma)^{-1}\sqrt{km})}{\log(1/\gamma)}, \quad (\text{C.3})$$

the norm of the truncation bias term is upper-bounded as

$$\frac{1}{km} \sum_{i=1}^k \sum_{t \in \mathcal{H}_i} \|\widehat{M}_i^{-1} \phi_t b_t\|_2 \leq \frac{2\gamma^b}{km(1-\gamma)} \sum_{i=1}^k \sum_{t \in \mathcal{H}_i} \|\widehat{M}_i^{-1} \phi_t\|_2 \leq 2\sigma^{-2} C_\Phi (km)^{-1/2}. \quad (\text{C.4})$$

To bound the error terms corresponding to reward noise z_t and average-error noise $J_{\pi_i} - \widehat{J}_{\pi_i}$, we rely on the independent blocks techniques of Yu (1994). We show in Sections C.1 and C.2 that with probability $1 - 2\delta$, for constants c_1 and c_2 , each of these terms can be bounded as:

$$\begin{aligned} \frac{1}{km} \left\| \sum_{i=1}^k \sum_{t \in \mathcal{H}_i} \widehat{M}_i^{-1} \phi_t z_t \right\|_2 &\leq 2c_1 C_\Phi \sigma^{-2} \sqrt{\frac{b \log(2d/\delta)}{km}} \\ \frac{b}{km} \left\| \sum_{i=1}^k (J_{\pi_i} - \widehat{J}_{\pi_i}) \sum_{t \in \mathcal{H}_i} \widehat{M}_i^{-1} \phi_t \right\|_2 &\leq 2c_2 C_\Phi \sigma^{-2} b \sqrt{\frac{\log(2d/\delta)}{km}}. \end{aligned}$$

Thus, putting terms together, we have for an absolute constant c , with probability at least $1 - \delta$,

$$\|\widehat{w}_k - w_k\|_2 \leq c\sigma^{-2} (C_w + C_\Phi) b \sqrt{\frac{\log(2d/\delta)}{km}}.$$

Note that this result holds for every $k \in [K]$ and thus also holds for $k = K$. $\square$

C.1. Bounding $\sum_{i=1}^k \widehat{M}_i^{-1} \sum_{t \in \mathcal{H}_i} \phi_t z_t$

Let $\|\cdot\|_{\text{tv}}$ denote the total variation norm.

Definition C.1 (β -mixing). Let $\{Z_t\}_{t=1,2,\dots}$ be a stochastic process. Denote by $Z_{1:t}$ the collection $(Z_1, \dots, Z_t)$, where we allow $t = \infty$. Let $\sigma(Z_{i:j})$ denote the sigma-algebra generated by $Z_{i:j}$ ($i \leq j$). The k^{th} β -mixing coefficient of $\{Z_t\}$, β_k , is defined by

$$\begin{aligned} \beta_k &= \sup_{t \geq 1} \mathbb{E} \sup_{B \in \sigma(Z_{t+k:\infty})} |P(B|Z_{1:t}) - P(B)| \\ &= \sup_{t \geq 1} \mathbb{E} \|P_{Z_{t+k:\infty}|Z_{1:t}}(\cdot|Z_{1:t}) - P_{Z_{t+k:\infty}}(\cdot)\|_{\text{tv}}. \end{aligned}$$

$\{Z_t\}$ is said to be β -mixing if $\beta_k \rightarrow 0$ as $k \rightarrow \infty$. In particular, we say that a β -mixing process mixes at an *exponential* rate with parameters $\bar{\beta}, \alpha, \gamma > 0$ if $\beta_k \leq \bar{\beta} \exp(-\alpha k^\gamma)$ holds for all $k \geq 0$.

Let X_t be the indicator vector for the state-action pair (x_t, a_t) as in Lemma A.2. Note that the distribution of (x_{t+1}, a_{t+1}) given (x_t, a_t) can be written as $\mathbb{E}[X_{t+1}|X_t]$. Let H_t be the state-action transition matrix at time t , let $H_{i:t} = \prod_{j=i}^{t-1} H_j$, and define $H_{i:i} = I$. Then we have that $\mathbb{E}[X_{t+k}|X_{1:t}] = H_{t:t+k}^\top X_t$ and $\mathbb{E}[X_{t+k}] = H_{1:t+k}^\top \nu_0$, where ν_0 is the initial state distribution. Thus, under the uniform mixing Assumption 2.1, the k^{th} β -mixing coefficient is bounded as:

$$\beta_k \leq \sup_{t \geq 1} \mathbb{E} \sum_{j=k}^{\infty} \|H_{t:t+j}^\top X_t - H_{1:t+j}^\top \nu_0\|_1 \leq \sup_{t \geq 1} \mathbb{E} \sum_{j=k}^{\infty} \gamma^j \|X_t - H_{1:t}^\top \nu_0\|_1 \leq \frac{2\gamma^k}{1-\gamma}.$$

We bound the noise terms using the independent blocks technique of Yu (1994). Recall that we partition each phase into $2m$ blocks of size b . Thus, after k phases we have a total of $2km$ blocks. Let $\mathbb{P}$ denote the joint distribution of state-action pairs in *odd* blocks. Let $\mathcal{I}_i$ denote the set of indices in the i^{th} block, and let $x_{\mathcal{I}_i}, a_{\mathcal{I}_i}$ denote the corresponding states and actions. We factorize the joint distribution according to blocks:

$$\begin{aligned} \mathbb{P}(x_{\mathcal{I}_1}, a_{\mathcal{I}_1}, x_{\mathcal{I}_3}, a_{\mathcal{I}_3}, \dots, x_{\mathcal{I}_{2km-1}}, a_{\mathcal{I}_{2km-1}}) &= \mathbb{P}_1(x_{\mathcal{I}_1}, a_{\mathcal{I}_1}) \times \mathbb{P}_3(x_{\mathcal{I}_3}, a_{\mathcal{I}_3} | x_{\mathcal{I}_1}, a_{\mathcal{I}_1}) \times \dots \\ &\quad \times \mathbb{P}_{2km-1}(x_{\mathcal{I}_{2km-1}}, a_{\mathcal{I}_{2km-1}} | x_{\mathcal{I}_{2km-3}}, a_{\mathcal{I}_{2km-3}}). \end{aligned}$$

Let $\tilde{\mathbb{P}}_i$ be the marginal distribution over the variables in block i , and let $\tilde{\mathbb{P}}$ be the product of marginals of odd blocks.

Corollary 2.7 of [Yu \(1994\)](#) implies that for any Borel-measurable set E ,

$$|\mathbb{P}(E) - \tilde{\mathbb{P}}(E)| \leq (km - 1)\beta_b \quad (\text{C.5})$$

where β_b is the b^{th} β -mixing coefficient of the process. The result follows since the size of the “gap” between successive blocks is b ; see Appendix E for more details.

Recall that our estimates $\hat{w}_{\pi_i}$ are based only on data in odd blocks in each phase. Let $\tilde{\mathbb{E}}$ denote the expectation w.r.t. the product-of-marginals distribution $\tilde{\mathbb{P}}$. Then $\tilde{\mathbb{E}}[\hat{M}_i^{-1} \sum_{t \in \mathcal{H}_i} \phi_t z_t] = 0$ because for $t \in \mathcal{H}_i$ and under $\tilde{\mathbb{P}}$, z_t is zero-mean given ϕ_t and is independent of other feature vectors outside of the block. Furthermore, by Hoeffding’s inequality $\tilde{\mathbb{P}}(|z_t|/b \geq a) \leq 2 \exp(-2ba^2)$. Since $\|\phi_t\|_2 \leq C_\Phi$ and $\|\hat{M}_i^{-1}\|_2 \leq 2\sigma^{-2}$ for large enough m , we have that

$$\tilde{\mathbb{P}}(\|\hat{M}_i^{-1} \phi_t z_t\|_2 \geq 2b\sigma^{-2}C_\Phi a) \leq 2 \exp(-2ba^2).$$

Since $\hat{M}_i^{-1} \phi_t z_t$ are norm-subGaussian vectors, using Lemma A.4, there exists a constant c_1 such that for any $\delta \geq 0$

$$\tilde{\mathbb{P}} \left(\left\| \sum_{i=1}^k \hat{M}_i^{-1} \sum_{t \in \mathcal{H}_i} \phi_t z_t \right\|_2 \geq 2c_1 C_\Phi \sigma^{-2} \sqrt{bkm \log(2d/\delta)} \right) \leq \delta.$$

Thus, using (C.5),

$$\mathbb{P} \left(\left\| \sum_{i=1}^k \hat{M}_i^{-1} \sum_{t \in \mathcal{H}_i} \phi_t z_t \right\|_2 \geq 2c_1 C_\Phi \sigma^{-2} \sqrt{bkm \log(2d/\delta)} \right) \leq \delta + (km - 1)\beta_b.$$

Under Assumption 2.1, we have that $\beta_b \leq 2\gamma^b(1 - \gamma)^{-1}$. Setting $\delta = 2km\gamma^b(1 - \gamma)^{-1}$ and solving for b we get

$$b = \frac{\log(2km\delta^{-1}(1 - \gamma)^{-1})}{\log(1/\gamma)}. \quad (\text{C.6})$$

Notice that when b is chosen as in Eq. (C.6), the condition (C.3) is also satisfied. Plugging this into the previous display gives that with probability at least $1 - 2\delta$,

$$\left\| \sum_{i=1}^k \hat{M}_i^{-1} \sum_{t \in \mathcal{H}_i} \phi_t z_t \right\|_2 \leq 2c_1 C_\Phi \sigma^{-2} \sqrt{bkm \log(2d/\delta)}.$$

C.2. Bounding $\left\| \sum_{i=1}^k \hat{M}_i^{-1} \sum_{t \in \mathcal{H}_i} \phi_t (J_{\pi_i} - \hat{J}_{\pi_i}) \right\|_2$

Recall that the average-reward estimates $\hat{J}_{\pi_i}$ are computed using time indices corresponding to the starts of even blocks, $\mathcal{T}_i$. Thus this error term is only a function of the indices corresponding to block starts. Now let $\mathbb{P}$ denote the distribution over state-action pairs (x_t, a_t) for indices t corresponding to block starts, i.e. $t \in \{1, b+1, 2b+1, \dots, (2km-1)b+1\}$. We again factorize the distribution over blocks as $\mathbb{P} = \mathbb{P}_1 \otimes \mathbb{P}_2 \otimes \dots \otimes \mathbb{P}_{2km}$. Let $\tilde{\mathbb{P}} = \tilde{\mathbb{P}}_1 \otimes \tilde{\mathbb{P}}_2 \otimes \dots \otimes \tilde{\mathbb{P}}_{2km}$ be a product-of-marginals distribution defined as follows. For odd j , let $\tilde{\mathbb{P}}_j$ be the marginal of $\mathbb{P}$ over (x_{jb+1}, a_{jb+1}) . For even j in phase i , let $\tilde{\mathbb{P}}_j = \nu_{\pi_i}$ correspond to the stationary distribution of the corresponding policy π_i . Using arguments similar to independent blocks, we show in Appendix E that

$$\|\mathbb{P} - \tilde{\mathbb{P}}\|_1 \leq 2(2km - 1)\gamma^{b-1}.$$

Let $\tilde{\mathbb{E}}$ denote expectation w.r.t. the product-of-marginals distribution $\tilde{\mathbb{P}}$. Then $\tilde{\mathbb{E}}[\hat{M}_i^{-1} \sum_{t \in \mathcal{H}_i} \phi_t (J_{\pi_i} - \hat{J}_{\pi_i})] = 0$, since under $\tilde{\mathbb{P}}$, $\hat{J}_{\pi_i}$ is the sum of rewards for state-action pairs distributed according to ν_{π_i} , and these state-action pairs are independent of other data. Using a similar argument as in the previous section, for $b = 1 + \frac{\log(4km/\delta)}{\log(1/\gamma)}$, there exists a constant c_2 such that with probability at least $1 - 2\delta$,

$$\left\| \sum_{i=1}^k \hat{M}_i^{-1} \sum_{t \in \mathcal{H}_i} \phi_t (J_{\pi_i} - \hat{J}_{\pi_i}) \right\|_2 \leq 2c_2 C_\Phi \sigma^{-2} \sqrt{km \log(2d/\delta)}.$$

C.3. Bounding $\|\widehat{M}_i^{-1}\|_2$

In this subsection, we show that with probability at least $1 - \delta$, for $m \geq 72C_\Phi^4\sigma^{-2}(1 - \gamma)^{-2}\log(d/\delta)$, $\|M_i^{-1}\|_2 \leq 2\sigma^{-2}$.

Let Φ be a $|\mathcal{X}||\mathcal{A}| \times d$ matrix of all features. Let $D_i = \text{diag}(\nu_{\pi_i})$, and let $\widehat{D}_i = \text{diag}(\sum_{t \in \mathcal{H}_i} X_t)$, where X_t is a state-action indicator as in Lemma A.2. Let $M_i = \Phi^\top D_i \Phi + \alpha m^{-1} I$. We can write $\widehat{M}_i^{-1}$ as

$$\begin{aligned}\widehat{M}_i^{-1} &= (\Phi^\top \widehat{D}_i \Phi + \alpha \tau^{-1} I + \Phi^\top (D_i - \widehat{D}_i) \Phi)^{-1} \\ &= (M_i + \Phi^\top (D_i - \widehat{D}_i) \Phi)^{-1} \\ &= (I + M_i^{-1} \Phi^\top (D_i - \widehat{D}_i) \Phi)^{-1} M_i^{-1}\end{aligned}$$

By Assumption 6.2 and 6.1, $\|M_i^{-1}\|_2 \leq \sigma^{-2}$. In Appendix C.4, we show that w.p. at least $1 - \delta$,

$$\|\Phi^\top (\widehat{D}_i - D_i) \Phi\|_2 \leq 6m^{-1/2} C_\Phi^2 (1 - \gamma)^{-1} \sqrt{2\log(d/\delta)}$$

Thus

$$\|\widehat{M}_i^{-1}\|_2 \leq \sigma^{-2} (1 - \sigma^{-2} 6m^{-1/2} C_\Phi^2 (1 - \gamma)^{-1} \sqrt{2\log(d/\delta)})^{-1}$$

For $m \geq 72C_\Phi^4\sigma^{-2}(1 - \gamma)^{-2}\log(d/\delta)$, the above norm is upper-bounded by $\|\widehat{M}_i^{-1}\|_2 \leq 2\sigma^{-2}$.

C.4. Bounding $\|\Phi^\top (\widehat{D}_i - D_i) \Phi\|_2$

For any matrix A ,

$$\|\Phi^\top A \Phi\|_2 = \left\| \sum_{i,j} A_{ij} \phi_i \phi_j^\top \right\|_2 \leq \sum_{i,j} |A_{ij}| \|\phi_i \phi_j^\top\|_2 \leq C_\Phi^2 \sum_{i,j} |A_{ij}| = C_\Phi^2 \|A\|_{1,1}. \quad (\text{C.7})$$

where $\|A\|_{1,1}$ denotes the sum of absolute entries of A . Using the same notation for X_t as in Lemma A.2,

$$\begin{aligned}\|\Phi^\top (\widehat{D}_i - D_i) \Phi\|_2 &= \frac{1}{m} \sum_{t \in \mathcal{H}_i} \Phi^\top \text{diag}(X_t - \nu_t + \nu_t - \nu_{\pi_i}) \Phi \\ &\leq \frac{1}{m} \left\| \sum_{t \in \mathcal{H}_i} \Phi^\top \text{diag}(X_t - \nu_t) \Phi \right\|_2 + \frac{C_\Psi^2}{m} \sum_{t \in \mathcal{H}_i} \|\nu_t - \nu_{\pi_i}\|_1.\end{aligned}$$

Under the fast-mixing assumption 2.1, the second term is bounded by $2C_\Psi^2 m^{-1} (1 - \gamma)^{-1}$.

For the first term, we can define a martingale $(B_i)_{i=0}^m$ similar to the Doob martingale in Lemma A.2, but defined only on the m indices $\mathcal{H}_i$. Note that $\sum_{t \in \mathcal{H}_i} \Phi^\top \text{diag}(X_t - \nu_t) \Phi = \Phi^\top \text{diag}(B_m - B_0) \Phi$. Thus we can use matrix-Azuma to bound the difference sequence. Given that

$$\|(\Phi^\top (B_i - B_{i-1}) \Phi)^2\|_2 \leq 4C_\Phi^4 (1 - \gamma)^{-2},$$

combining the two terms, we have that with probability at least $1 - \delta$,

$$\begin{aligned}\|\Phi^\top (\widehat{D}_i - D_i) \Phi\|_2 &\leq 4m^{-1/2} C_\Phi^2 (1 - \gamma)^{-1} \sqrt{2\log(d/\delta)} + 2m^{-1} C_\Phi^2 (1 - \gamma)^{-1} \\ &\leq 6m^{-1/2} C_\Phi^2 (1 - \gamma)^{-1} \sqrt{2\log(d/\delta)}.\end{aligned}$$

D. Bounding $\|V_K - \widehat{V}_K\|_{\mu_*}$

We write the value function error as follows:

$$\begin{aligned}\mathbb{E}_{x \sim \mu_*} [\widehat{V}_K(x) - V_K(x)] &= \sum_x \mu_*(x) \sum_a \phi(x, a)^\top \frac{1}{K} \sum_{i=1}^K \pi_i(a|x) (\widehat{w}_{\pi_i} - w_{\pi_i}) \\ &\leq \frac{1}{K} \sum_x \mu_*(x) \sum_a \|\phi(x, a)\|_2 \left\| \sum_{i=1}^K \pi_i(a|x) (\widehat{w}_{\pi_i} - w_{\pi_i}) \right\|_2\end{aligned}$$

Note that for any set of scalars $\{p_i\}_{i=1}^K$ with $p_i \in [0, 1]$, the term $\left\| \sum_{i=1}^K p_i (\hat{w}_{\pi_i} - w_{\pi_i}) \right\|_2$ has the same upper bound as $\left\| \sum_{i=1}^K (\hat{w}_{\pi_i} - w_{\pi_i}) \right\|_2$. The reason is as follows. One part of the error includes bias terms (C.2) and (C.4), whose upper bounds are only smaller when reweighted by scalars in $[0, 1]$. Thus we can simply upper-bound the bias by setting all $\{p_i\}_{i=1}^K$ to 1. Another part of the error, analyzed in Appendices C.1 and C.2 involves sums of norm-subGaussian vectors. In this case, applying the weights only results in these vectors potentially having smaller norm bounds. We keep the same bounds for simplicity, again corresponding to all $\{p_i\}_{i=1}^K$ equal to 1. Thus, reusing the results of the previous section, we have

$$\mathbb{E}_{x \sim \mu^*} [\hat{V}_K(x) - V_K(x)] \leq C_\Phi |\mathcal{A}| c \sigma^{-2} (C_w + C_\Phi) b \sqrt{\frac{\log(2d/\delta)}{Km}}.$$

E. Independent Blocks

Blocks. Recall that we partition each phase into $2m$ blocks of size b . Thus, after k phases we have a total of $2km$ blocks. Let $\mathbb{P}$ denote the joint distribution of state-action pairs in odd blocks. Let $\mathcal{I}_i$ denote the set of indices in the i^{th} block, and let $x_{\mathcal{I}_i}, a_{\mathcal{I}_i}$ denote the corresponding states and actions. We factorize the joint distribution according to blocks:

$$\begin{aligned} \mathbb{P}(x_{\mathcal{I}_1}, a_{\mathcal{I}_1}, x_{\mathcal{I}_3}, a_{\mathcal{I}_3}, \dots, x_{\mathcal{I}_{2km-1}}, a_{\mathcal{I}_{2km-1}}) &= \mathbb{P}_1(x_{\mathcal{I}_1}, a_{\mathcal{I}_1}) \times \mathbb{P}_3(x_{\mathcal{I}_3}, a_{\mathcal{I}_3} | x_{\mathcal{I}_1}, a_{\mathcal{I}_1}) \times \dots \\ &\quad \times \mathbb{P}_{2km-1}(x_{\mathcal{I}_{2km-1}}, a_{\mathcal{I}_{2km-1}} | x_{\mathcal{I}_{2km-3}}, a_{\mathcal{I}_{2km-3}}). \end{aligned}$$

Let $\tilde{\mathbb{P}}_i$ be the marginal distribution over the variables in block i , and let $\tilde{\mathbb{P}}$ be the product of marginals. Then the difference between the distributions $\tilde{\mathbb{P}}$ and $\mathbb{P}$ can be written as

$$\begin{aligned} \mathbb{P} - \tilde{\mathbb{P}} &= \mathbb{P}_1 \otimes \mathbb{P}_3 \otimes \dots \otimes \mathbb{P}_{2km-1} - \mathbb{P}_1 \otimes \tilde{\mathbb{P}}_3 \otimes \dots \otimes \tilde{\mathbb{P}}_{2km-1} \\ &= \mathbb{P}_1 \otimes (\mathbb{P}_3 - \tilde{\mathbb{P}}_3) \otimes \mathbb{P}_5 \otimes \dots \otimes \mathbb{P}_{2km-1} \\ &\quad + \mathbb{P}_1 \otimes \tilde{\mathbb{P}}_3 \otimes (\mathbb{P}_5 - \tilde{\mathbb{P}}_5) \otimes \mathbb{P}_7 \otimes \dots \otimes \mathbb{P}_{2km-1} \\ &\quad + \dots \\ &\quad + \mathbb{P}_1 \otimes \tilde{\mathbb{P}}_3 \otimes \tilde{\mathbb{P}}_5 \otimes \dots \otimes \tilde{\mathbb{P}}_{2km-3} \otimes (\mathbb{P}_{2km-1} - \tilde{\mathbb{P}}_{2km-1}). \end{aligned}$$

Under β -mixing, since the gap between the blocks is of size b , we have that

$$\|\mathbb{P}_i(x_{\mathcal{I}_i}, a_{\mathcal{I}_i} | x_{\mathcal{I}_{i-2}}, a_{\mathcal{I}_{i-2}}) - \tilde{\mathbb{P}}_i(x_{\mathcal{I}_i}, a_{\mathcal{I}_i})\|_1 \leq \beta_b = \frac{2\gamma^b}{1-\gamma}.$$

Thus the difference between the joint distribution and the product of marginals is bounded as

$$\|\mathbb{P} - \tilde{\mathbb{P}}\|_1 \leq (km - 1)\beta_b.$$

Block starts. Now let $\mathbb{P}$ denote the distribution over state-action pairs (x_t, a_t) for indices t corresponding to block starts, i.e. $t \in \{1, b+1, 2b+1, \dots, (2km-1)b+1\}$. We again factorize the distribution over blocks:

$$\mathbb{P}(x_1, a_1, x_{b+1}, a_{b+1}, \dots, x_{(2km-1)b+1}, a_{(2km-1)b+1}) = \mathbb{P}_1(x_1, a_1) \prod_{j=2}^{2km} \mathbb{P}_j(x_{jb+1}, a_{jb+1} | x_{(j-1)b+1}, a_{(j-1)b+1}).$$

Define a product-of-marginals distribution $\tilde{\mathbb{P}} = \tilde{\mathbb{P}}_1 \otimes \tilde{\mathbb{P}}_2 \otimes \dots \otimes \tilde{\mathbb{P}}_{2km}$ over the block-start variables as follows. For odd j , let $\tilde{\mathbb{P}}_j$ be the marginal of $\mathbb{P}$ over (x_{jb+1}, a_{jb+1}) . For even j in phase i , let $\tilde{\mathbb{P}}_j = \nu_{\pi_i}$ correspond to the stationary distribution of the policy π_i . Using the same notation as in Appendix A, let X_t be the indicator vector for (x_t, a_t) and let $H_{i:j}$ be the product of state-action transition matrices at times $i+1, \dots, j$. For odd blocks j , we have

$$\|\mathbb{P}_j(\cdot | x_{(j-1)b+1}, a_{(j-1)b+1}) - \tilde{\mathbb{P}}_j(\cdot)\|_1 = \|H_{(j-1)b+1:j}^\top (X_{(j-1)b+1} - \tilde{\mathbb{P}}_{j-1})\|_1 \leq 2\gamma^{b-1}.$$

Slightly abusing notation, let H_{π_i} be the state-action transition matrix under policy π_i . For even blocks j in phase i , since they always follow an odd block in the same phase,

$$\|\mathbb{P}_j(\cdot | x_{(j-1)b+1}, a_{(j-1)b+1}) - \tilde{\mathbb{P}}_j(\cdot)\|_1 = \|(H_{\pi_i}^{b-1})^\top (X_{(j-b)+1} - \nu_{\pi_i})\|_1 \leq 2\gamma^{b-1}.$$

Thus, using a similar distribution decomposition as before, we have that $\|\mathbb{P} - \tilde{\mathbb{P}}\|_1 \leq 2(2km - 1)\gamma^{b-1}$.