arXiv:2210.05918v3 [cs.LG] 19 Sep 2024

Finite time analysis of temporal difference learning with linear function approximation: Tail averaging and regularisation[1]

Gandharv Patil[1], **Prashanth L.A.**[2], **Dheeraj Nagaraj**[3], **Doina Precup**[1]

1: McGill University & Mila, 2: Indian Institute of Technology Madras, 3: Google Research

gandharv.patil@mail.mcgill.ca, prashla@cse.iitm.ac.in, dheerajnagaraj@google.com, dprecup@cs.mcgill.ca

Abstract: We study the finite-time behaviour of the popular temporal difference (TD) learning algorithm when combined with tail-averaging. We derive finite time bounds on the parameter error of the tail-averaged TD iterate under a step-size choice that does not require information about the eigenvalues of the matrix underlying the projected TD fixed point. Our analysis shows that tail-averaged TD converges at the optimal $O(1/t)$ rate, both in expectation and with high probability. In addition, our bounds exhibit a sharper rate of decay for the initial error (bias), which is an improvement over averaging all iterates. We also propose and analyse a variant of TD that incorporates regularisation. From analysis, we conclude that the regularised version of TD is useful for problems with ill-conditioned features.

1. Introduction

Temporal difference (TD) [21] learning is an efficient and easy to implement stochastic approximation algorithm used for evaluating the long-term performance of a decision policy. The algorithm predicts the value function using a single sample path obtained by simulating the Markov decision process (MDP) with a given policy. Analysis of TD algorithms is challenging, and researchers have devoted significant effort in studying its asymptotic properties [10, 15, 19, 23]. In recent years, there has been an interest in characterising the finite-time behaviour of TD, and several papers [2, 3, 6, 12, 17] have tackled this problem under various assumptions.

For t iterations/updates, most existing works either provide a $O\left(\frac{1}{t^\alpha}\right)$ (with universal step-size) [2, 6] or a $O\left(\frac{1}{t}\right)$ (with constant step-size) [2, 12, 17] convergence rate to the TD-fixed point $\theta^\star$ defined as $\theta^\star \triangleq A^{-1}b$, where A and b are quantities which depend on the MDP and the policy (see Section 2 for the notational information). To obtain a $O\left(\frac{1}{t}\right)$ rate with a constant step-size, [2, 17] assume that the minimum eigenvalue of the matrix A is known apriori. However, in a typical RL setting, such eigenvalue information is not available. Estimating the matrix A and its lowest eigenvalue accurately might require a large number of additional samples, which makes the algorithm more complicated. Therefore, obtaining a $O\left(\frac{1}{t}\right)$ rate for TD with a *universal* step-size is an important open problem.

In this paper, we provide a solution to this problem by establishing a $O\left(\frac{1}{t}\right)$ bound on the convergence rate for a variant of TD that incorporates tail-averaging, and uses a constant 'universal' step-size. In [2, 12, 17] the authors study an alternate version called iterate averaging which was introduced independently by Polyak and Juditsky [16] and Ruppert [18] for general stochastic-approximation algorithms. A shortcoming of iterate averaging is that the initialisation error (i.e., distance between θ_0 and $\theta^\star$) is forgotten at a slower rate than the non-averaged case, and in practical implementations, one usually performs averaging after a sufficient number of iterations have been performed. This type of delayed averaging, called 'tail-averaging', has been explored in the context of ordinary least squares by Jain et al. in [11].

Inspired by the analysis of TD learning, we propose a variant of TD that incorporates regularisation, wherein we introduce a parameter λ and solve for the regularised TD fixed point given by $\theta^\star_{\text{reg}} = (A+\lambda \mathbf{I})^{-1}b$. The update rule for this algorithm is similar to vanilla TD except that it involves an additional factor with λ. Through our analysis we observe that using regularisation can be helpful in obtaining better non-asymptotic bounds for many problems, where the discount factor is close to 1.

Concretely, the contributions of this paper are as follows: First, we establish a $O(1/t)$ finite time bounds on the convergence rate of tail-averaged TD and tail-averaged TD with regularisation. Similar to [2, 6], the analysis assumes that the data is sampled in an i.i.d. fashion from a fixed distribution. The resulting bounds are valid under a universal step-size and hold in expectation as well as high probability. We also show that Markov sampling can be handled with simple mixing arguments. The salient features of the bounds for each variant are as follows:

Tail averaged TD: In this variant, the step-size is a function of the discount factor and a bound on the norm of the state features. The expectation bound provides a $O(1/t)$ convergence rate for tail-averaged TD iterate, while

[1] This is longer version of the paper presented at AISTATS 2023 [14]

the high-probability bound establishes an exponential concentration of tail-averaged TD around the projected TD fixed point.

Tail-averaged TD with regularisation: For this variant, the step-size is a function of the discount factor, regularisation parameter λ, and a bound on the norm of the state features. Although this variant converges to the regularised TD fixed-point $\theta^\star_{\text{reg}}$, we show that the worse-case bound on the difference between TD fixed point $\theta^\star$ and $\theta^\star_{\text{reg}}$ is $O(\lambda)$ in the ℓ_2 norm. Moreover, our analysis makes a case for using the regularised TD algorithm for problems with ill-conditioned features.

Next, we show that under mixing assumptions, we can extend our results to Markov sampling instead of i.i.d. sampling. These error bounds contain an extra $\tilde{O}(\tau_{\text{mix}})$, where τ_{mix} is the underlying Markov chain's mixing time. This is no better than making the samples appear approximately i.i.d. by considering one out of every $\tilde{O}(\tau_{\text{mix}})$ samples, and then dropping the rest. In fact, as per Nagaraj et al. [13, Theorem 2], even with the discount factor $\beta = 0$, it is information-theoretically impossible to do any better without further assumptions on the nature of the linear approximation. Recently Agarwal et al. [1] showed that for linear MDPs, one can use reverse experience replay with function approximation to obtain finite time bounds which are independent of the mixing time constant. We leave the study of TD with different experience replay strategies as an interesting future direction, and for the sake of completeness, present the bounds for Markov sampling in Remark 8, and provide a proof sketch in Section 7.

In Table 1 we compare our expectation bounds with existing bounds in the literature. In addition, we also derive high-probability bounds for tail-averaged TD with/without regularisation, and we provide a summary of these bounds in a tabular form in Table 2.

Table 1. Summary of the bounds in expectation of the form $\mathbb{E}[\|\theta_{\text{Alg},t} - \theta^\star\|_2^2]$, where $\theta^\star$ is the TD fixed point, and $\theta_{\text{Alg},t}$ is the parameter picked by an algorithm after t iterations of TD.

Reference	Algorithm	Step-size	Rate
Bhandari et al. [2]	Last iterate	c/t[1]	$O(1/t)$
	Averaged iterate	$\frac{1}{\sqrt{t}}$	$O(1/\sqrt{t})$
Dalal et al. [6]	Last iterate	$1/t^\alpha$	$O(1/t^\alpha)$
Lakshminarayanan and Szepesvari [12]	Constant step-size with averaging	c	$O(1/t)$
Prashanth et al. [17]	Last iterate	c/n, $c \propto 1/\mu$	$O(1/t)$
	Averaged iterate	c/t^α, $c > 0$	$O(1/t^\alpha)$
Our work	Tail-averaged TD	$c > 0$	$O(1/t)$
	Regularised TD[2]	$c > 0$	$O(1/t)$

[1] Step-size requires information about eigenvalue of the feature covariance matrix Σ.

[2] The convergence here is to the regularised TD solution.

Related work. Over the past few years, there has been significant interest in understanding the finite-time behaviour of TD learning. Several researchers have proposed interesting frameworks establishing bounds on TD's convergence rate under different assumptions. In [7, 9, 20, 24, 25] the authors analyse the finite time behaviour of TD using Lyapunov drift-conditions and establish finite time bounds that hold under expectation. The advantage of this framework is that it can be used directly for analysing TD with Markov noise. However, to provide an $O(1/t)$ bound, these analyses use a step-size which depends on the eigenvalue of A. For *eg.*, in [20, Theorem 7], we have $\epsilon = O(\frac{\log T}{\gamma_{\max} T})$ where $\gamma_{\max}$ is essentially the smallest eigenvalue of A. Similar conditions can also be found in [7, Eq. (88)], [24, Proposition 2], and [9, Eq. (18)].

The analysis presented in this work is closely related to bounds established in [2, 12, 17], where the authors

Table 2. Summary of the high-probability bounds of the form $\mathbb{P}\left[\left\|\theta_{\mathrm{Alg},t}-\theta^\star\right\|_2^2 \le h(t)\right]$, where $\theta^\star$ is the TD fixed point, $\theta_{\mathrm{Alg},t}$ is the parameter picked by an algorithm after t iterations of TD, and $h(t)$ is a function of t that depends on Alg.

Reference	Algorithm	Step-size	$h(t)$
Dalal et al. [6]	Last iterate	$1/t^\alpha$	$O(1/t^\alpha)$
Prashanth et al. [17]	Last iterate	$c/n,\ c \propto 1/\mu$	$O(1/t)$
	Averaged iterate	$c/t^\alpha,\ c>0$	$O(1/t^\alpha)$
Our work	Tail-averaged TD	$c>0$	$O(1/t)$
	Regularised TD[2]	$c>0$	$O(1/t)$

[2]The convergence here is to the regularised TD solution.

provide an $O(1/t)$ bound in expectation on the mean square error of the parameters. Our bounds match the overall order of these bounds under comparable assumptions. The principal advantage with our bounds is that they hold for a 'universal' step-size choice, while the aforementioned references required the knowledge of μ. Another advantage with our bounds, owing to tail averaging, is that the initial error is forgotten exponentially fast, while the corresponding term in the aforementioned references exhibit a power law decay. In another related work, for a universal step-size the authors in [6] provide a $O(1/t^\alpha)$ bound in expectation, where $\alpha \in (0,1)$, while we obtain a $O(1/t)$ bound under similar assumptions.

Finally, high-probability bounds for TD have been derived in [6, 17]. In comparison to these works, the high-probability bound that we derive is easy to interpret and exhibits better concentration properties. The related Q-learning algorithm and modifications have also been considered in the finite-time regime with linear function approximation (cf. [4, 5] and the references therein). However, these results too require the knowledge of the condition number to set the step-size.

The rest of the paper is organised as follows: In Section 2, we present the main model of TD with function approximation used for our analysis. In Section 3, we describe the tail-averaged TD algorithm, and also present the finite time bounds for this algorithm. In Section 4, we combine tail-averaging with regularisation in a TD algorithm, and provide finite time bounds for this algorithm. In Section 5, we present a sketch of the proofs of our main results, and the detailed proofs are available in section 6. In Section 7, we discuss the extension of our results to address the case of Markov sampling. Finally, in Section 8, we provide the concluding remarks.

2. TD with linear function approximation

Consider an MDP $\langle \mathcal{S}, \mathcal{A}, P, r, \beta\rangle$, where $\mathcal{S}$ is the state space, $\mathcal{A}$ is the action space, $P(s'|s,a)$ is the probability of transitioning to the state s' from the state s on choosing action a, $r : \mathcal{S}\times\mathcal{A} \to \mathbb{R}$ is the per step reward, and $\beta \in [0,1)$ is the discount factor. We assume that the state and action spaces are both finite. A stationary randomised policy π maps every state s to a distribution over actions. For a given policy π, we define the value function V^π as follows:

$$V^\pi(s) = \mathbb{E}^{\pi,P}\left[\sum_{t=0}^{\infty} \beta^t r(s_t,a_t) | S_0 = s\right], \tag{1}$$

where the action a_t in state s_t is chosen using policy π, i.e., $a_t \sim \pi(s_t)$. The value function V^π obeys the Bellman equation $\mathcal{T}^\pi V^\pi = V^\pi$, where the Bellman operator $\mathcal{T}^\pi$ is defined by $(\mathcal{T}^\pi V)(s) \triangleq \mathbb{E}^{\pi,P}\left[r(s,a) + \beta V(s')\right]$, where the action a is chosen using π, i.e., $a \sim \pi(s)$ and the next state s' is drawn from $P(\cdot|s)$.

2.1. *Value function approximation*

Most practical applications have high-dimensional state-spaces making exact computation of the value function infeasible. One solution to overcome this problem is to use a parametric approximation of the value function. In

this work, we consider the linear function approximation architecture [22], where the value function $V^\pi(s)$, for any $s \in \mathcal{S}$, is approximated as follows:

$$V^\pi(s) \approx \widetilde{V}(s;\theta) := \phi(s)^\top \theta. \tag{2}$$

In the above, $\phi(s) \in \mathbb{R}^d$ is a fixed feature vector for state s, and $\theta \in \mathbb{R}^d$ is a parameter vector that is shared across states. When the state space is a finite set, say $\mathcal{S} = \{1, 2, \ldots, n\}$, the n-vector $\widetilde{V}(\theta)$ with components $\widetilde{V}(s;\theta)$ can be expressed as follows:

$$\widetilde{V}(\cdot;\theta) = \underbrace{\begin{bmatrix} \phi_1(1) & \phi_2(1) & \ldots & \phi_d(1) \\ \vdots & \vdots & \ddots & \vdots \\ \phi_1(n) & \phi_2(n) & \ldots & \phi_d(n) \end{bmatrix}}_{\Phi} \underbrace{\begin{bmatrix} \theta_1 \\ \vdots \\ \theta_d \end{bmatrix}}_{\theta}, \tag{3}$$

where $\Phi \in \mathbb{R}^{n \times d}$, and $\theta \in \mathbb{R}^d$.
The objective is to learn the best parameter for approximating V^π within the following linear space:

$$\mathcal{B} := \{\Phi\theta \mid \theta \in \mathbb{R}^d\}. \tag{4}$$

Naturally, with a linear function approximation, it is not possible to find the fixed point $V^\pi = \mathcal{T}^\pi V^\pi$. Instead, one can approximate V^π within $\mathcal{B}$ by solving a projected system of equations. The system of equations, which is also referred to as the projected Bellman equation, is given by

$$\Phi\theta^\star = \Pi\mathcal{T}^\pi(\Phi\theta^\star), \tag{5}$$

where Π is the orthogonal projection operator onto the set $\mathcal{B}$ using a weighted ℓ_2-norm. More precisely, let $D = \mathrm{diag}(\rho(1), \ldots, \rho(n)) \in \mathbb{R}^{n \times n}$ denote a diagonal matrix, whose elements are given by the stationary distribution ρ of the Markov chain underlying the policy π. We assume that the stationary distribution exists (see Assumption 1). Let $\|V\|_D = \sqrt{V^\top D V}$ denote the weighted norm of a n-vector V, and assume that the matrix Φ has full column rank. Then, the operator Π projects orthogonally onto $\mathcal{B}$ using the $\|\cdot\|_D$ norm, and it can be shown that $\Pi = \Phi(\Phi^\top D \Phi)^{-1}\Phi^\top D$.
Next, the projected TD fixed point θ^* for (5) is given by:

$$A\theta^\star = b, \text{ where } A \triangleq \Phi^\top D(\mathbf{I} - \beta P)\Phi, \quad b \triangleq \Phi^\top D\mathcal{R}, \tag{6}$$

and $\mathcal{R} = \sum_{a \in \mathcal{A}} \pi(s,a) r(s,a)$.

2.2. *Temporal Difference (TD) Learning*

Temporal difference (TD) [22] algorithms are a class of stochastic approximation methods used for solving the projected linear system given in (5). These algorithms start with a initial guess for the θ_0, and at every time-step t and update them using samples from the Markov chain induced by a policy π.
The update rule is given as follows:

$$\begin{aligned} \theta_t &= \theta_{t-1} + \gamma f_t(\theta_{t-1}), \text{ where} \\ f_t(\theta) &\triangleq (r_t + \beta\theta^\top\phi(s_t') - \theta^\top\phi(s_t))\phi(s_t). \end{aligned} \tag{7}$$

In the above, γ is the step-size parameter.
An alternate version of the algorithm (which we consider for deriving the high probability bounds) uses the projection Γ as follows:

$$\theta_t = \Gamma(\theta_{t-1} + \gamma f_t(\theta_{t-1})). \tag{8}$$

In (8), operator Γ projects the iterate θ_t onto the nearest point in a closed ball $C \in \mathbb{R}^d$ with a radius H, which is large enough to include $\theta^\star$.
An interesting result by [23] tells us that for any $\theta \in \mathbb{R}^d$, the function
$f(\theta) \triangleq (r(s,a) + \beta\theta^\top\phi(s') - \theta^\top\phi(s))\phi(s)$ has a well defined steady-state expectation given by

$$\mathbb{E}^{\rho,P}[f(\theta)] = \sum_{s,s' \in \mathcal{S}, a \in \mathcal{A}} \rho(s)\pi(s,a)\Big((r(s,a) + P(s'|s,a)\beta\theta^\top\phi(s') - \theta^\top\phi(s))\phi(s)\Big). \tag{9}$$

We can rearrange (9) as $\sum_{s,s'\in\mathcal{S},a\in\mathcal{A}} P(s'|s,a)(r(s,s')+\beta\theta^\top\phi(s'_t)) = (\mathcal{T}^\pi\Phi\theta)(s)$, and use [23, Lemma 8] to get the following:

$$\mathbb{E}^{\rho,P}[f(\theta)] = \Phi^\top D(\mathcal{T}^\pi(\Phi\theta) - \Phi\theta) \tag{10}$$

$$= -A\theta + b, \tag{11}$$

where A and b are as defined in (6). We can then characterise the mean behaviour of TD algorithm using the following update rule:

$$\theta_t = \theta_{t-1} + \gamma\Big(\Phi^\top D(\mathcal{T}^\pi(\Phi\theta_{t-1}) - \Phi\theta_{t-1})\Big)$$

$$= \theta_{t-1} + \gamma\mathbb{E}^{\rho,P}[f(\theta_{t-1})]. \tag{12}$$

The characterisation of TD's behaviour in (12) is of particular importance as it forms the basis of our analysis.

3. Tail-averaged TD

3.1. Basic algorithm

Tail averaging or suffix averaging refers to returning the average of the final few iterates of the optimisation process, to improve its variance properties. Specifically, for any t, the tail-averaged iterate $\theta_{k+1,N}$ is the average of $\{\theta_{k+1},\dots,\theta_t\}$, computed as follows:

$$\theta_{k+1,N} = \frac{1}{N}\sum_{i=k+1}^{k+N}\theta_i, \tag{13}$$

where $N = t-k$.

Note that Polyak and Juditsky [16], showed that averaging all the iterates produces best asymptotic convergence rate, but from a non-asymptotic analysis viewpoint, it is usually observed that the initial error (the rate at which the initial point is forgotten) is forgotten slower with iterate averaging as compared to the non-averaged case, see [8]. Tail averaging retains the advantages of iterate averaging, while ensuring that the initial error is forgotten exponentially fast – a conclusion that can be inferred from the finite time bounds that we derive for the TD algorithm.

Algorithm 1 presents the pseudocode of the tail-averaged TD algorithm.

Algorithm 1: Tail-averaged TD(0)

Input : Initial parameter θ_0, step-size γ, initial state distribution ζ_0, tail-average index k.

Sample an initial state $s_0 \sim \zeta_0$;

for $t = 0, 1, \dots$ **do**

  Choose an action $a_t \sim \pi(s_t)$;

  Observe r_t, and next state s'_t;

  Update parameters: $\theta_t = \theta_{t-1} + \gamma f(\theta_{t-1})$;

  Average the final N iterates: $\theta_{k+1,N} = \frac{1}{N}\sum_{i=k+1}^{k+N}\theta_i$, where $N = t-k$.

end

3.2. Finite time bounds

Before presenting our results, we list the assumptions under which we conduct our analysis.

Assumption 1. The Markov chain underlying the policy π is irreducible.

Assumption 2. The samples $\{s_t, r_t, s'_t\}_{t\in\mathbb{N}}$ are independently and identically drawn from: $\rho(s)P(s'|s)$ where, at time t the state s_t, where ρ stationary distribution underlying policy π, and $P(\cdot|s_t)$ is the transition probability matrix of the MDP.

Assumption 3. For all $s \in \mathcal{S}$, $\|\phi(s)\|_2 \le \Phi_{\max} < \infty$.

Assumption 4. For all $s \in \mathcal{S}$, and $a \in \mathcal{A}$, $|r(s,a)| \le R_{\max} < \infty$.

Assumption 5. The matrix Φ has full column rank.

Assumption 6. The set $C \triangleq \{\theta \in \mathbb{R}^d | \, \|\theta\|_2 \le H\}$ used for projection through Γ satisfies $H > \frac{\|b\|_2}{\mu}$.

We now discuss the assumptions listed above. Assumption 1 ensures the existence of the stationary distribution for the Markov chain underlying policy π. We study the non-asymptotic behaviour of the tail-averaged TD algorithm under the i.i.d observation model as specified in Assumption 2, and later show that our results can be extended to handle Markov sampling. Next, Assumptions 3 and 4 are boundedness requirements on the underlying features and rewards, and are common in the finite time analysis of TD algorithm, see [2, 17]. Assumption 5 requires the columns of the feature matrix Φ to be linearly independent, in turn ensuring the uniqueness of the TD solution θ^*. Moreover, this assumption ensures that the minimum eigenvalue, say μ' of $B = \mathbb{E}^{\rho,P}[\Phi\Phi^\top]$ is strictly positive, in turn implying that the minimum eigenvalue μ of the matrix A defined in (6) is strictly positive. Assumption 6 is required for the high-probability bounds, while the bounds in expectation do not require projection.

The first result we state below is a bound in expectation on the parameter error $\left\|\theta_{k+1,N} - \theta^\star\right\|_2^2$.

Theorem 1 (Bound in expectation). *Suppose Assumptions 1 to 5 hold. Choose a step size γ satisfying*

$$\gamma \leq \gamma_{\max} = \frac{1-\beta}{(1+\beta)^2\Phi_{\max}^2}, \tag{14}$$

where β is the discount factor and $\Phi_{\max}$ is a bound on the features (see Assumption 3).
Then the expected error of the tail-averaged iterate $\theta_{k+1,N}$ when using Algorithm 1 satisfies

$$\mathbb{E}\left[\left\|\theta_{k+1,N} - \theta^\star\right\|_2^2\right] \leq \frac{10e^{(-k\gamma(1-\beta)\mu')}}{\gamma^2(1-\beta)^2\mu'^2N^2}\mathbb{E}\left[\left\|\theta_0 - \theta^\star\right\|_2^2\right] + \frac{10\sigma^2}{(1-\beta)^2\mu'^2N}, \tag{15}$$

where $N = t-k$, θ_0 is the initial point, $\sigma^2 = (R_{\max} + (1+\beta)\Phi_{\max}^2\left\|\theta^\star\right\|_2^2)$, with $\theta^\star$ denoting the TD fixed point specified in (6), and μ' is the minimum eigenvalue of $B = \mathbb{E}^{\rho,P}[\Phi\Phi^\top]$.

Proof. See Section 5.1 for a sketch and Subsection 6.2.4 for a detailed proof. □

A few remarks are in order.

Remark 1. It is apparent that the bound presented above scales inversely with the square of $(1-\beta)\mu'$. More importantly, the bound presented above is for a step-size choice that does not require information about the eigenvalues of matrices A or B. To the best of our knowledge, this is the first bound of $O(1/t)$ for a 'universal' step-size. Previous results, such as those by [2, 17] provide a comparable bound, albeit for a diminishing step-size of the form c/k, where setting c requires knowledge of μ. On the other hand, [6, 17] provide a $O(1/t^\alpha)$ bound for larger step-sizes of the form c/t^α, where c is a universal constant.

Remark 2. The first term on the RHS of (15) relates to the rate at which the initial parameter θ_0 is forgotten, while the second term arises from a martingale difference noise term associated with the i.i.d. sampling model. Setting $k = t/2$, we observe that the first term is forgotten at an exponential rate, while the noise term is $O(1/t)$.

Remark 3. In [12], the authors consider iterate averaging in a linear stochastic approximation setting. Comparing their Theorem 1 to the result we have presented above, we note that the first term on the RHS of (15) exhibits an exponential decay, while the corresponding decay is of order $O(1/t)$ in [12]. The second term in their result as well as in (15) is of order $O(1/t)$. While the second dominates the rate, the first term, which relates to the rate at which the initial parameter is forgotten, decays much faster with tail averaging. Intuitively, it makes sense to average after sufficient iterations have passed, instead of averaging from the beginning, and our bounds confirm this viewpoint.

Remark 4. A closely related result under comparable assumptions is Theorem 2 of [2]. This result provides two bounds corresponding to constant and diminishing step-sizes, respectively, while assuming the knowledge of μ. The bound there corresponding to the constant step-size for the last iterate of TD is the sum of an exponentially decaying 'initial error' term and a constant offset with the noise variance. The second bound in the aforementioned work is $O(1/t)$ for both initial error and noise terms. The bound we derived in (15) combines the best of these two bounds through tail averaging, i.e., an exponentially decaying initial error, and a $O(1/t)$ noise term. As an aside, our bound is for the projection-free variant of TD, while the bounds in [2] requires projection, with an assumption similar to Assumption 6.

Remark 5. Another closely related result is Theorem 4.4 of [17], where the authors analyse TD with linear function approximation, with input data from a batch of samples. The analysis there can be easily extended to cover our i.i.d. sampling model. As in the remark above, while the overall rate is $O(1/t)$ in their result as well as (15), the initial error in our bound is forgotten much faster. A similar observation also holds w.r.t. the bound in the recent work [3], but the authors do not state their bound explicitly.

Remark 6. It is possible to extend our analysis to cover the Markov noise observation model, as specified in Section 8 of [2]. In this model, we assume that the underlying Markov chain is uniformly ergodic, which intuitively translates to a fast mixing rate. For finite Markov chains irreducibility and aperiodicity are sufficient to

establish this assumption. The fast mixing assumption allows us to translate the i.i.d. sampling resutls to Markov sampling. We provide the details of such an extension in Section 7.

Next, we turn to providing a bound that holds with high probability the parameter error $\left\|\theta_{k+1,N}-\theta^\star\right\|_2^2$ of the projected TD algorithm. For this result, we require the TD update parameter to stay within a bounded region that houses θ^*, which is formalized in Assumption 6.

Theorem 2 (High-probability bound). *Suppose Assumptions 1 to 6 hold. Choose the step size such that $\gamma \le \gamma_{\max}$, where $\gamma_{\max}$ is defined in (14). Then, for any $\delta \in (0,1]$, we have the following bound for the projected tail-averaged iterate $\theta_{k+1,N}$:*

$$P\left(\left\|\theta_{k+1,N}-\theta^\star\right\|_2 \le \frac{2\sigma}{(1-\beta)\mu'\sqrt{N}}\sqrt{\log\left(\frac{1}{\delta}\right)} + \frac{4e^{(-k\gamma(1-\beta)^2\mu')}}{\gamma(1-\beta)\mu' N}\mathbb{E}\left[\left\|\theta_0-\theta^\star\right\|_2\right] + \frac{4\sigma}{(1-\beta)\mu'\sqrt{N}}\right) \ge 1-\delta,$$

where $N, \sigma, \mu, \theta_0, \theta^\star$ are as specified in Theorem 1.

Proof. See Subsection 5.2 for a sketch and Subsection 6.3.2 for a detailed proof. □

Remark 7. High-probability bounds for TD algorithm have been derived earlier in [6, 17]. In comparison to Theorem 4.2 of [17], we note that our bound is an improvement since the sampling error (the first and third terms in $\mathcal{K}(n)$ defined above) decays at a much faster rate for tail-averaged TD. Next, unlike [6], we note that our bound requires projection. However it does exhibit a $O(1/t)$ rate. The result by [6] (Theorem 3.6) is of the form $O(1/t^\lambda)$ where λ is related to the minimum eigenvalue μ of matrix A, and hence cannot be guaranteed to be of order $O(1/t)$.

Remark 8. Consider the case when Assumption 2 does not hold, but we sample (s_t, r_t, s_t') from a trajectory corresponding the policy π. We assume exponential ergodicity for the total variation distance as used in [2, 20], with mixing time τ_{mix}. For this case, we consider a variant of TD which uses one sample in every $\tilde{O}(\tau_{\text{mix}})$ consecutive samples for the update iteration. The guarantees for the resulting TD algorithm with tail averaging with N data points in the trajectory corresponds to the guarantees for $\theta_{k+1,N'}$ in Theorem 2 where $N' = \tilde{O}\left(\frac{N}{\tau_{\text{mix}}}\right)$, and $(1-\delta)$ is replaced by $(1-2\delta)$. This gives an error of the order $\tilde{O}\left(\sqrt{\frac{\tau_{\text{mix}}}{N}}\right)$, which is similar to the bounds in [2, Theorem 3]. These follow from standard mixing arguments and we refer to Section 7 for further details. As an aside, we remark that one cannot get a better bound from an information-theoretic viewpoint without further assumptions on the nature of the linear approximation (cf. Theorem 2 in [13]).

4. Regularized TD Learning

In this section, we analyse the regularised TD algorithm. From the results in Theorems 1 and 2 one can observe that although tail-averaged TD achieves a $O\left(\frac{1}{t}\right)$ rate of convergence, the bounds depend inversely on $(1-\beta)\mu'$, where μ' is the minimum eigenvalue of $B = \mathbb{E}^{\rho,P}[\Phi\Phi^\top]$. In the following results we will show that the non-asymptotic bounds for regularised TD scale inversely with μ (minimum eigenvalue of matrix A). Such a dependence may be preferable over vanilla TD, as there are problem instances where $(1-\beta)\mu' \ll \mu$. To make this intuition more concrete, consider the following problem instance.

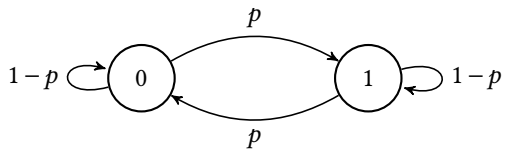

Fig. 1. A two state Markov chain

Example 1. Consider a two state MDP with the transition dynamics as depicted in Figure 1, for a given policy, say π. The one-dimensional state features are given as follows: $\phi(1) = 1$, and $\phi(2) = \frac{1}{2}$. For the case of $p = \frac{1}{2}$, we have

$$A = (1/2)\left(\phi(1)^2 + \phi(2)^2\right) - \beta/4\left(\phi(1)^2 + \phi(2)^2 + \phi(1)\phi(2) + \phi(2)\phi(1)\right) = \frac{5}{8} - \frac{9\beta}{16}.$$

Further, $B = \frac{5}{8}$. Thus, for any $\beta \in [0,1]$, $(1-\beta)B \le A$. Further, as β approaches 1, $(1-\beta B) \to 0$, while $A \to \frac{1}{16}$. Since convergence rate of tail-averaged TD depends inversely on $(1-\beta)\mu'$ (see Theorems 1 and 2), there is a concrete case for an algorithm whose convergence rate depends on μ instead of $(1-\beta)\mu'$. The regularised TD variant that we present next achieves this objective.

4.1. Basic algorithm

Instead of the TD solution (28), we solve the following regularised problem for a given regularisation parameter $\lambda > 0$:

$$\theta^\star_{\text{reg}} = (A + \lambda \mathbf{I})^{-1} b, \tag{16}$$

The update iteration of the TD analogue for the regularised case is as follows:

$$\hat{\theta}_t = (\mathbf{I} - \gamma\lambda)\hat{\theta}_{t-1} + \gamma(r_t + \beta\hat{\theta}_{t-1}^\top \phi(s_t') - \hat{\theta}_{t-1}^\top \phi(s_t))\phi(s_t). \tag{17}$$

Similarly, the projected regularised TD update (which we consider for deriving the high probability bounds) uses the projection Γ as follows:

$$\hat{\theta}_t = \Gamma((\mathbf{I} - \gamma\lambda)\hat{\theta}_{t-1} + \gamma(r_t + \beta\hat{\theta}_{t-1}^\top \phi(s_t') - \hat{\theta}_{t-1}^\top \phi(s_t))\phi(s_t)). \tag{18}$$

In (18), operator Γ projects the iterate θ_t onto the nearest point in a closed ball $C \in \mathbb{R}^d$ with a radius H which is large enough to include $\theta^\star_{\text{reg}}$.

Using arguments similar to vanilla TD, it is easy to see that the iterate $\hat{\theta}_t$ converges (16) under Assumptions 1 to 5, and a standard stochastic approximation condition on the step-size.

The overall flow of the regularised TD algorithm would be similar to Algorithm 1, except that the iterate is updated according to (17), and an additional regularisation parameter is involved.

4.2. Finite time bounds

Using a technique similar to that used in establishing the bound for tail-averaged TD in Theorem 1, we arrive at the following bound in expectation for regularised TD.

Theorem 3 (Bound in expectation). *Suppose Assumptions 1 to 4 hold. Choose a step size γ satisfying*

$$\gamma \leq \gamma_{max} = \frac{\lambda}{\lambda^2 + 2\lambda(1+\beta)\Phi^2_{\text{max}} + (1+\beta)^2\Phi^4_{\text{max}}}. \tag{19}$$

Then the expected error of the tail-averaged regularised TD iterate $\hat{\theta}_{k+1,N}$ satisfies

$$\mathbb{E}\left[\left\|\hat{\theta}_{k+1,N} - \theta^\star_{\text{reg}}\right\|_2^2\right] \leq \frac{10e^{(-k\gamma(\mu+\lambda))}}{\gamma^2(\mu+\lambda)^2N^2}\mathbb{E}\left[\left\|\hat{\theta}_0 - \theta^\star_{\text{reg}}\right\|_2^2\right] + \frac{10\sigma^2}{(\mu+\lambda)^2N}, \tag{20}$$

where $N = t - k$, and $\sigma^2 = (R_{\text{max}} + (1+\beta)\Phi^2_{\text{max}}\left\|\theta^\star_{\text{reg}}\right\|_2^2)$.

Proof. See Section 5.3 for a sketch and Subsection 6.2.4 for a detailed proof. □

While the result above bounded the distance to the regularised TD solution, the next result shows that the distance between regularised TD iterate and vanilla projected TD fixed point is of $O(\lambda)$.

Corollary 1. Under conditions of Theorem 3, we have

$$\mathbb{E}\left[\left\|\hat{\theta}_{k+1,N} - \theta^\star\right\|_2^2\right] \leq \frac{20e^{(-k\gamma(\mu+\lambda))}}{\gamma^2(\mu+\lambda)^2N^2}\mathbb{E}\left[\left\|\hat{\theta}_0 - \theta^\star_{\text{reg}}\right\|_2^2\right] + \frac{20\sigma^2}{(\mu+\lambda)^2N} + \frac{2\lambda^2\Phi^2_{\text{max}}R^2_{\text{max}}}{\sigma_{\text{min}}(A)^2(\mu+\lambda)^2}, \tag{21}$$

where $\sigma_{\text{min}}(A)$ is A's minimum singular value

Proof. See Subsection 6.4.4 □

With a suitable choice of λ, the next result shows that regularised TD obtains a $O(1/t)$ rate (e.g., with $k = t/2$), and the bound scales inversely with the eigenvalue μ of the matrix A. From the discussion earlier, recall that there are problem instances where $\mu \gg (1-\beta)\mu'$, and the bound for tail-averaged TD sans regularisation depended inversely on $(1-\beta)\mu'$.

Corollary 2. Under conditions of Theorem 3, and with $\lambda = \frac{1}{\sqrt{N}}$, we obtain

$$\mathbb{E}\left[\left\|\hat{\theta}_{k+1,N} - \theta^\star\right\|_2^2\right] \leq \frac{20(1+(1+\beta)\Phi^2_{\text{max}}\sqrt{N})^4 e^{\frac{(-k\mu)}{(1+\beta)^2\Phi^4_{\text{max}}\sqrt{N}}}}{\mu^2N^3}\mathbb{E}\left[\left\|\hat{\theta}_0 - \theta^\star_{\text{reg}}\right\|_2^2\right] + \frac{20\sigma^2}{\mu^2N} + \frac{2\Phi^2_{\text{max}}R^2_{\text{max}}}{\mu^2N} \tag{22}$$

A few remarks are in order.

Remark 9. The choice of step-size in the bound of Theorem 3 is universal, i.e., does not require the knowledge of μ, and the rate of convergence is $O(1/t)$, if we set $k = t/2$, or any constant multiple of t. However, the regularised TD iterate converges to (16), which is different from the vanilla TD fixed point. But, the distance is between the regularised and vanilla TD solutions is $O(\lambda)$, implying that for small value of λ, the regularised TD solution is a good proxy, which in turn implies that the regularised TD iterate can be used in place of vanilla TD iterate, to obtain a good approximation to the TD fixed point. Corollary 1 makes this intuition precise.

Remark 10. The initial and sampling errors in (20) are as in the tail-averaged TD (see Theorem 1), i.e., initial error is forgotten at an exponential rate, while the sampling error is $O(1/t)$.

Remark 11. In [17], the authors analyse the iterate-average variant of TD, and derive a $O(1/t^{\alpha})$ bound for a step-size $\Theta(1/k^{\alpha})$, where $1/2 < \alpha < 1$. Further, their step-size choice is universal as is the case of tail-averaged TD in comparison to [17]. Our bound for regularised TD exhibits a better rate, though the bound measures the distance to the regularised TD solution.

Next, we present a high-probability bound for regularised TD in the spirit of Theorem 2.

Theorem 4 (High-probability bound). *Suppose Assumptions 1 to 4, and 6 hold. Choose the step size such that $\gamma \leq \gamma_{\max}$, where $\gamma_{\max}$ is defined in (19). Then, for any $\delta \in (0,1]$, we have the following bound for the projected tail-averaged regularised TD iterate $\hat{\theta}_{k+1,N}$:*

$$P\left(\left\|\hat{\theta}_{k+1,N} - \theta^{\star}_{\text{reg}}\right\|_2 \leq \frac{2\sigma}{(\mu+\lambda)\sqrt{N}}\sqrt{\log\left(\frac{1}{\delta}\right)} + \frac{4e^{(-k\gamma(\mu+\lambda))}}{\gamma(\mu+\lambda)N}\mathbb{E}\left[\left\|\hat{\theta}_0 - \theta^{\star}_{\text{reg}}\right\|_2\right] + \frac{4\sigma}{(\mu+\lambda)\sqrt{N}}\right) \geq 1-\delta,$$

where $N, \sigma, \mu, \hat{\theta}_0, \theta^{\star}_{\text{reg}}$ are as specified in Theorem 3.

Proof. The proof follows the same template as Theorem 2, and is given in Subsection 6.5.2. □

As discussed in Remark 8 and Section 7 for tail-averaged TD sans regularisation, it is straightforward to extend the results in Theorems 3 and 4 to cover the case of Markov sampling.

5. Proof Ideas

5.1. Proof of Theorem 1 (Sketch)

Proof. We present here the framework for obtaining the results obtained in the paper; the framework has been introduced in the work of [2, 6, 17]. Towards that end, we begin by introducing some notation. First, we define the centered error $z_t \triangleq \theta_t - \theta^{\star}$. Using the TD update (7), the centered error can be seen to satisfy the following recursive relation:

$$z_t = (\mathbf{I} - \gamma a_t) z_{t-1} + \gamma f_t(\theta^{\star}), \tag{23}$$

where $f(\cdot)$ defined as in (7), and $a_t \triangleq \phi(s_t)\phi(s_t)^{\top} - \beta\phi(s_t)\phi(s_t')^{\top}$.

The centered error is decomposed into a bias and variance term as follows:

$$\begin{aligned}\mathbb{E}\left[\|z_t\|_2^2\right] &= 2\mathbb{E}\left[\left\|C^{t:1}z_0\right]\right\|_2^2\Big] + 2\gamma^2\sum_{k=0}^{t}\mathbb{E}\left[\left\|\sum_{k=0}^{t} C^{t:k+1} f_k(\theta^{\star})\right\|_2^2\right] \\ &= 2z_t^{\text{bias}} + 2\gamma^2 z_t^{\text{variance}},\end{aligned}$$

where

$$C^{i:j} = \begin{cases}(\mathbf{I} - \gamma a_i)(\mathbf{I} - \gamma a_{i-1})\ldots(\mathbf{I} - \gamma a_j), & \text{if } i >= j \\ \mathbf{I}, & \text{otherwise}\end{cases} \tag{24}$$

The bias term is then bounded as $z_t^{\text{bias}} \leq \exp(-\gamma(1-\beta)\mu' t)\mathbb{E}[\|z_0\|_2^2]$, while the variance term is bounded as $z_t^{\text{variance}} \leq \frac{\sigma^2}{(1-\beta)\mu'}$.

The centered error corresponding to the tail-averaged iterate $\theta_{k+1,N}$ is given by $z_{k+1,N} = \frac{1}{N}\sum_{i=k+1}^{k+N} z_i$. The analysis proceeds by bounding the expectation of the norm $\mathbb{E}[\left\|z_{k+1,N}\right\|_2^2]$ using the following decomposition:

$$\mathbb{E}\left[\left\|z_{k+1,N}\right\|_2^2\right] \leq \frac{1}{N^2}\left(\sum_{i=k+1}^{k+N}\mathbb{E}\left[\|z_i\|_2^2\right] + 2\sum_{i=k+1}^{k+N-1}\sum_{j=i+1}^{k+N}\mathbb{E}\left[z_i^{\top}z_j\right]\right).$$

Using the definitions of z_t^{bias} and z_t^{variance}, we simplify the RHS above as follows:

$$\mathbb{E}\left[\left\|z_{k+1,N}\right\|_2^2\right] \leq \underbrace{\frac{2}{N^2}\left(1+\frac{4}{\gamma(1-\beta)\mu'}\right)\sum_{i=k+1}^{k+N} z_i^{\text{bias}}}_{z_{k+1,N}^{\text{bias}}} + \underbrace{\frac{2}{N^2}\left(1+\frac{4}{\gamma(1-\beta)\mu'}\right)\gamma^2\sum_{i=k+1}^{k+N} z_i^{\text{variance}}}_{z_{k+1,N}^{\text{variance}}}, \tag{25}$$

where, $z_i^{\text{bias}} = \mathbb{E}\left[\left\|C^{t:1}z_0\right\|_2^2\right]$ and $z_i^{\text{variance}} = \sum_{k=0}^{t}\mathbb{E}\left[\left\|\sum_{k=0}^{t} C^{t:k+1}f_k(\theta^\star)\right\|_2^2\right]$, and $f(\cdot)$ defined as in (7).
The main result follows by substituting the bounds on z_i^{bias} and z_i^{variance} followed by some algebraic manipulations. □

5.2. *Proof of Theorem 2 (Sketch)*

Proof. To obtain the high-probability bound we use the proof technique by Prashanth et al. [17], where we consider separately the deviation of the centered error from its mean, *i.e.*, $\left\|z_{k+1,N}\right\|_2^2 - \mathbb{E}\left[\left\|z_{k+1,N}\right\|_2^2\right]$. We decompose this quantity as a sum of martingale differences, establish a Lipschitz property followed by a sub-Gaussian concentration bound to infer

$$P\left(\left\|z_{k+1,N}\right\|_2 - \mathbb{E}\left[\left\|z_{k+1,N}\right\|_2\right] > \epsilon\right) \leq \exp\left(-\frac{\epsilon^2}{(R_{\max}+(1+\beta)H\Phi_{\max}^2)^2\sum_{i=k+1}^{k+N}L_i^2}\right), \tag{26}$$

where $L_i \triangleq \frac{\gamma}{N}\sum_{j=i+1}^{i+N}\left(1-\frac{\gamma(1-\beta)\mu'}{2}\right)^{j-i+1}$ Next, under the choice of step-size γ specified in the theorem statement, we establish that

$$\sum_{i=k+1}^{k+N} L_i^2 \leq \frac{4}{N(1-\beta)\mu'^2}.$$

The main claim follows by (i) substituting the bound obtained above in (26); (ii) using the bound on $\mathbb{E}\left[\left\|z_{k+1,N}\right\|_2\right]$ specified in Theorem 1; and (iii) converting the tail bound resulting from (i) and (ii) into a high-probability bound. The detailed proof is given in Subsection 6.3. □

5.3. *Proof of Theorem 3 (Sketch)*

The proof of Theorem 3 follows the same template as Theorem 1. The following lemma captures the interplay of the step size and regularization parameters and its subsequent effect on the constants and decay rates in Theorem 3's result.

Lemma 1. *With $\gamma \leq \gamma_{max}$ as given in (19), the following bound holds*

$$\left\|\left(\mathbf{I}-\gamma(A+\lambda\mathbf{I})\right)^\top\left(\mathbf{I}-\gamma(A+\lambda\mathbf{I})\right)\right\|_2 \leq 1-\gamma(\mu+\lambda), \quad \textit{and} \quad \|(\mathbf{I}-\gamma A)\|_2 \leq 1-\frac{\gamma(\mu+\lambda)}{2}.$$

6. Convergence Analysis

6.1. *Preliminaries*

Let $\mathcal{F}_t$ denote the sigma-field generated by $\{\theta_0 \ldots \theta_t\}$, $t \geq 0$, and let

$$f_t(\theta) \triangleq (r_t + \beta\theta^\top\phi(s_t') - \theta^\top\phi(s_t))\phi(s_t). \tag{27}$$

Recall that n denotes the number of states in the underlying MDP and the feature matrix $(|\mathcal{S}|\times d)$-matrix, where $\Phi \triangleq (\phi_1(s_1), \ldots, \phi_d(s_n))$.
According to the characterisation of TD's steady-state behaviour in (12), it's final solution can be written as follows:

$$\theta^\star = A^{-1}b, \tag{28}$$

Using (12), we rewrite the TD update in eq. (7) as follows:

$$\begin{aligned}\theta_t &= \theta_{t-1} + \gamma f_t(\theta_{t-1}) \\ &= \theta_{t-1} + \gamma\Big(-A\theta_{t-1} + b + \Delta M_t\Big),\end{aligned} \tag{29}$$

where $\Delta M_t = f_t(\theta_{t-1}) - \mathbb{E}^{\rho,P}[f_t(\theta_{t-1})|\mathcal{F}_{t-1}]$ is a martingale difference sequence with f defined as in (27).

We shall first establish a few useful results in Lemma 3 to 6.

Lemma 2. *For any $a, b \in \mathbb{R}^d$, we have*

$$-\frac{\theta^\top(aa^\top + bb^\top)\theta}{2} \le \theta^\top(ab^\top)\theta \le \frac{\theta^\top(aa^\top + bb^\top)\theta}{2}.$$

Proof.

$$\begin{aligned}\theta^\top ab^\top \theta &\overset{(a)}{\le} \left(\left(\theta^\top aa^\top\theta\right)\left(\theta^\top bb^\top\theta\right)\right)^{1/2} \\ &\overset{(b)}{\le} \frac{(\theta^\top aa^\top\theta) + (\theta^\top bb^\top\theta)}{2} \\ &= \frac{\theta^\top(aa^\top + bb^\top)\theta}{2},\end{aligned}$$

where (a) follows from Cauchy-Schwarz inequality and (b) follows from AM-GM inequality. □

Lemma 3. *The matrix A defined in (28) satisfies*

$$\|A\|_2 \le (1+\beta)\Phi_{\max}^2,$$

where $\Phi_{\max}$ is specified in Assumption 3, and β is the discount factor.

Proof.

$$\begin{aligned}\|A\|_2 &= \|\mathbb{E}[\phi(s_t)\phi(s_t)^\top - \beta\phi(s_t)\phi(s_t')^\top]\|_2 \\ &\overset{(a)}{\le} \mathbb{E}[\|\phi(s_t)\phi(s_t)^\top\|_2] + \beta\mathbb{E}[\|\phi(s_t)\phi(s_t')^\top\|_2] \\ &\overset{(b)}{\le} (1+\beta)\Phi_{\max}^2,\end{aligned}$$

where (a) is due to Jensen's inequality, and (b) uses Assumption 3. □

Lemma 4. *Recall that $B \triangleq \mathbb{E}^\rho[\phi(s_t)\phi(s_t)^\top]$, For any $\theta \in \mathbb{R}^d$ the following inequality holds*

$$2(1-\beta)\theta^\top B\theta \le \theta^\top(A + A^\top)\theta \le 2(1+\beta)\theta^\top B\theta.$$

Proof. We will first provide a proof for the lower bound

$$\begin{aligned}\theta^\top(A+A^\top)\theta &= \theta^\top\left(\mathbb{E}[\phi(s_t)\phi(s_t)^\top - \beta\phi(s_t)\phi(s_t')^\top] + \mathbb{E}[(\phi(s_t)\phi(s_t)^\top)^\top - \beta(\phi(s_t)\phi(s_t')^\top)^\top]\right)\theta \\ &= \theta^\top(2B - \beta\mathbb{E}[\phi(s_t)\phi(s_t')^\top + \phi(s_t')\phi(s_t)^\top])\theta \\ &\overset{(a)}{\ge} \theta^\top(2B - \beta(\mathbb{E}[\phi(s_t)\phi(s_t)^\top] + \mathbb{E}[\phi(s_t')\phi(s_t')^\top])\theta \\ &\overset{(b)}{\ge} 2(1-\beta)\theta^\top B\theta,\end{aligned} \tag{30}$$

where (a) is uses Lemma 2, and (b) follows from the fact that $\mathbb{E}^\rho[\phi(s_t)\phi(s_t)^\top] = \mathbb{E}^\rho[\phi(s_t')\phi(s_t')^\top]$ which holds because $\phi(s_t)$ and $\phi(s_t')$ are both sampled from the stationary distribution ρ.
For the upper bound, we use Lemma 2 (30) to obtain

$$\begin{aligned}\theta^\top(A+A^\top)\theta &\le \theta^\top(2B + \beta\mathbb{E}[\phi(s_t)\phi(s_t)^\top + \phi(s_t')\phi(s_t')^\top])\theta \\ &= 2(1+\beta)\theta^\top B\theta.\end{aligned}$$

□

Lemma 5. *Let $a_j \triangleq [\phi(s_j)\phi(s_j)^\top - \beta\phi(s_j)\phi(s_j')^\top]$, and $B \triangleq \mathbb{E}^\rho[\phi(s_t)\phi(s_t)^\top]$. Then for any $\theta \in \mathbb{R}^d$, the following inequality holds*

$$\theta^\top(\mathbb{E}[a_j^\top a_j])\theta \le \Phi_{\max}^2(1+\beta)^2\theta^\top B\theta.$$

Proof. Note that

$$a_j^\top a_j = \left\|\phi(s_j)\right\|_2^2 \left[\phi(s_j)\phi(s_j)^\top - \beta\{\phi(s_j)\phi(s_j')^\top + \phi(s_j')\phi(s_j)^\top\} + \beta^2\phi(s_j')\phi(s_j')^\top\right].$$

Therefore,

$$\begin{aligned}
\theta^\top \mathbb{E}[a_j^\top a_j]\theta &= \theta^\top \left(\left\|\phi(s_j)\right\|_2^2 \mathbb{E}\left[\phi(s_j)\phi(s_j)^\top - \beta\{\phi(s_j)\phi(s_j')^\top + \phi(s_j')\phi(s_j)^\top\} + \beta^2\phi(s_j')\phi(s_j')^\top\right]\right)\theta \\
&\overset{(a)}{\leq} \theta^\top \left(\Phi_{\max}^2 \mathbb{E}\left[\phi(s_j)\phi(s_j)^\top - \beta\{\phi(s_j)\phi(s_j')^\top + \phi(s_j')\phi(s_j)^\top\} + \beta^2\phi(s_j')\phi(s_j')^\top\right)^\top\right]\theta \\
&\overset{(b)}{\leq} \theta^\top \left(\Phi_{\max}^2 \mathbb{E}\left[\phi(s_j)\phi(s_j)^\top - \beta\{\phi(s_j)\phi(s_j)^\top + \phi(s_j')\phi(s_j')^\top\} + \beta^2\phi(s_j')\phi(s_j')^\top\right]\right)\theta \\
&\leq \theta^\top \left(\Phi_{\max}^2 \left(\mathbb{E}[\phi(s_j)\phi(s_j)^\top + \beta^2\phi(s_j')\phi(s_j')^\top] - \beta\,\mathbb{E}[\phi(s_j)\phi(s_j)^\top + \phi(s_j')\phi(s_j')^\top]\right)\right)\theta \\
&\overset{(c)}{\leq} \Phi_{\max}^2(1+2\beta+\beta^2)\theta^\top B\theta \\
&= \Phi_{\max}^2(1+\beta)^2\theta^\top B\theta,
\end{aligned}$$

where (a) follows from Assumption 3, (b) follows from Lemma 2, and (c) holds because $\mathbb{E}^\rho[\phi(s_t)\phi(s_t)^\top] = \mathbb{E}^\rho[\phi(s_t')\phi(s_t')^\top]$.

□

Lemma 6. *Let $a_j \triangleq [\phi(s_j)\phi(s_j)^\top - \beta\phi(s_j)\phi(s_j')^\top]$. With $\gamma \leq \gamma_{\max} = \frac{1-\beta}{(1+\beta)^2\Phi_{\max}^2}$, the following bounds hold for any random variable $\theta \in \mathbb{R}^d$ that is $\mathcal{F}_j$ measurable:*

$$\mathbb{E}\left[\theta^\top \left(\mathbf{I} - \gamma a_j\right)^\top \left(\mathbf{I} - \gamma a_j\right)\theta | \mathcal{F}_j\right] \leq (1 - \gamma(1-\beta)\mu')\left\|\theta\right\|_2^2,$$

and

$$\mathbb{E}\left[\left\|\left(\mathbf{I} - \gamma a_j\right)\theta\right\|_2 | \mathcal{F}_j\right] \leq \left(1 - \frac{\gamma(1-\beta)\mu'}{2}\right)\left\|\theta\right\|_2.$$

Proof. Note that

$$\begin{aligned}
\mathbb{E}[\theta^\top(\mathbf{I}-\gamma a_j)^\top(\mathbf{I}-\gamma a_j)\theta|\mathcal{F}_j] &= \mathbb{E}[\theta^\top(\mathbf{I} - \gamma(a_j^\top + a_j) + \gamma^2 a_j^\top a_j)\theta|\mathcal{F}_j] \\
&= \left\|\theta\right\|_2^2 - \underbrace{\gamma\,\theta^\top\mathbb{E}[a_j^\top + a_j|\mathcal{F}_j]\theta}_{\text{Term 1}} + \underbrace{\gamma^2\,\theta^\top\mathbb{E}[a_j^\top a_j|\mathcal{F}_j]\theta}_{\text{Term 2}}.
\end{aligned} \tag{31}$$

We now bound "Term 1" on the RHS above as follows:

$$\begin{aligned}
\theta^\top\mathbb{E}[a_j^\top + a_j|\mathcal{F}_j]\theta &= \theta^\top(A^\top + A)\theta \\
&\overset{(a)}{\geq} 2(1-\beta)\theta^\top B\theta,
\end{aligned} \tag{32}$$

where (a) follows from Lemma 4, with $B \triangleq \mathbb{E}^\rho[\phi(s_t)\phi(s_t)^\top]$.
Next, we bound "Term 2" on the RHS of eq. (31) as follows:

$$\theta^\top\mathbb{E}[a_j^\top a_j|\mathcal{F}_j]\theta \overset{(a)}{\leq} \Phi_{\max}^2(1+\beta)^2\theta^\top B\theta, \tag{33}$$

where (a) is due to Lemma 5.
Compiling (32) and (33), we get

$$\begin{aligned}
\mathbb{E}[\theta^\top(\mathbf{I}-a_j)^\top(\mathbf{I}-a_j)\theta|\mathcal{F}_j] &\leq \left\|\theta\right\|_2^2 - 2\gamma(1-\beta)\theta^\top B\theta + \gamma^2\Phi_{\max}^2(1+\beta)^2\theta^\top B\theta \\
&= \left\|\theta\right\|_2^2 - \gamma(1-\beta)\left(2 - \gamma\frac{(1+\beta)^2}{(1-\beta)}\Phi_{\max}^2\right)\theta^\top B\theta \\
&\overset{(a)}{\leq} \left\|\theta\right\|_2^2 - \gamma(1-\beta)\theta^\top B\theta \\
&\overset{(b)}{\leq} \left\|\theta\right\|_2^2 - \gamma(1-\beta)\mu'\left\|\theta\right\|_2^2
\end{aligned}$$

$$= (1-\gamma\mu'(1-\beta))\,\|\theta\|_2^2, \tag{34}$$

where (a) uses $\gamma \le \gamma_{\max}$, and (b) follows from using Cauchy-Schwarz inequality along with the fact that μ' as the minimum eigenvalue of B.
Applying Jensen's inequality in conjunction with (34), we obtain

$$\begin{aligned}\mathbb{E}[\left\|(\mathbf{I}-a_j)\theta\right\|_2|\mathcal{F}_j] &\le (1-\gamma\mu'(1-\beta))^{1/2}\,\|\theta\|_2 \\ &\overset{(a)}{\le} \left(1-\frac{\gamma\mu'(1-\beta)}{2}\right)\|\theta\|_2,\end{aligned}$$

where (a) uses $1-\gamma\mu'(1-\beta) \ge 0 \implies (1-\gamma\mu'(1-\beta))^{\frac{1}{2}} \le 1-\frac{\gamma\mu'(1-\beta)}{2}$. □

6.2. *Expectation bound for Tail-averaged TD (Proof of Theorem 1)*

6.2.1. Bias-variance decomposition of the non-asymptotic error

Let,

$$C^{i:j} = \begin{cases}(\mathbf{I}-\gamma a_i)(\mathbf{I}-\gamma a_{i-1})\dots(\mathbf{I}-\gamma a_j), & \text{if } i >= j \\ \mathbf{I}, & \text{otherwise}\end{cases} \tag{35}$$

where a_t is defined in Lemma 5
Next let,

$$z_t^{\text{bias}} \triangleq \mathbb{E}\left[\left\|C^{t:1}z_0\right]\right\|_2^2\Big], \tag{36}$$

and (37)

$$z_t^{\text{variance}} \triangleq \mathbb{E}\left[\left\|\sum_{k=1}^{t} C^{t:k+1}f_k(\theta^\star)\right\|_2^2\right]. \tag{38}$$

Let the centered error be $z_t = \theta_t - \theta^\star$. Then, using eq. (7), we obtain

$$\begin{aligned}z_t &= \theta_{t-1} - \theta^\star + \gamma(r_t + \beta\theta_{t-1}^\top\phi(s_t') - \theta_{t-1}^\top\phi(s_t))\phi(s_t) \\ &= \theta_{t-1} - \theta^\star + \gamma(r_t\phi(s_t) - a_t\theta_{t-1}) + \gamma a_t\theta^\star - \gamma a_t\theta^\star \\ &= (\mathbf{I}-\gamma a_t)z_{t-1} + \gamma f_t(\theta^\star) \\ &\overset{(a)}{=} C^{t:1}z_0 + \gamma\sum_{k=1}^{t} C^{t:k+1}f_k(\theta^\star),\end{aligned} \tag{39}$$

where (a) follows from unrolling the update rule.
Taking expectations on both sides of (39), we obtain

$$\begin{aligned}\mathbb{E}[\|z_t\|_2^2] &= \mathbb{E}\left[\left\|C^{t:1}z_0 + \gamma\sum_{k=0}^{t} C^{t:k+1}f_k(\theta^\star)\right\|_2^2\right] \\ &\overset{(a)}{\le} 2z_t^{\text{bias}} + 2\gamma^2 z_t^{\text{variance}},\end{aligned} \tag{40}$$

where (a) is obtained by using the following inequalities, i. $\|(a+b)\|_2^2 \le 2\|a\|_2^2 + 2\|b\|_2^2$, ii. $\left\|\sum_{i=0}^n \mathbf{x}_i\right\| \le \sum_{i=0}^n \|\mathbf{x}_i\|$. z_t^{bias} and z_t^{variance} are given by (36) and (38).
Therefore, a bound on $\mathbb{E}\,\|z_t\|_2^2$ can obtained by bounding individual terms in eq. (40).

6.2.2. Bounding z_t^{bias}

Lemma 7. *Under conditions of Theorem 1, we have*

$$z_t^{\text{bias}} \le \exp(-\gamma(1-\beta)\mu' t)\mathbb{E}\left[\|z_0\|_2^2\right],\ \forall t \ge 1. \tag{41}$$

Proof. Notice that

$$\begin{aligned}z_t^{\text{bias}} &= \mathbb{E}[\left\|C^{t:1}z_0\right]\|_2^2] \\ &= \mathbb{E}\left[(C^{t-1:1}z_0)^\top(\mathbf{I}-\gamma a_t)^\top(\mathbf{I}-\gamma a_t)C^{t-1:1}z_0\right]\end{aligned}$$

$$\begin{aligned}
&\overset{(a)}{\leq} (1-\gamma(1-\beta)\mu')\mathbb{E}\left[\left\|C^{t-1:1}z_0\right\|_2^2\right] \\
&\overset{(b)}{\leq} (1-\gamma(1-\beta)\mu')^t\mathbb{E}\left[\|z_0\|_2^2\right] \\
&\leq \exp(-\gamma(1-\beta)\mu' t)\mathbb{E}\left[\|z_0\|_2^2\right],
\end{aligned} \tag{42}$$

where (a) follows from Lemma 6, and (b) follows from using the argument in (a) repeatedly. □

6.2.3. Bounding z_t^{variance}

Lemma 8. *For any $\mathcal{F}_i$ - measurable random vector $\mathbf{x} \in \mathbb{R}^d$, and $t \geq i$, we have*

$$\sum_{i=0}^{t}\mathbb{E}\left[\left\|C^{t:i+1}\mathbf{x}\right\|_2^2\right] \leq \frac{\mathbb{E}\left[\|\mathbf{x}\|_2^2\right]}{\gamma(1-\beta)\mu'}. \tag{43}$$

Proof.

$$\begin{aligned}
\mathbb{E}\left[\sum_{i=0}^{t}\left\|C^{t:i+1}\mathbf{x}\right\|_2^2\right] &= \mathbb{E}\left[\sum_{i=0}^{t}\left(C^{t:i+1}\mathbf{x}\right)^\top\left(C^{t:i+1}\mathbf{x}\right)\right] \\
&= \mathbb{E}\left[\sum_{i=0}^{t}(C^{t-1:i+1}\mathbf{x})^\top(\mathbf{I}-\gamma a_i)^\top(\mathbf{I}-\gamma a_i)C^{t-1:i+1}\mathbf{x}\right] \\
&\overset{(a)}{=} \sum_{i=0}^{t}\mathbb{E}\left[(C^{t-1:i+1}\mathbf{x})^\top(\mathbf{I}-\gamma a_i)^\top(\mathbf{I}-\gamma a_i)C^{t-1:i+1}\mathbf{x}\right] \\
&\overset{(b)}{\leq} \sum_{i=0}^{t}(1-\gamma(1-\beta)\mu')\mathbb{E}\left[\left\|C^{t-1:i+1}\mathbf{x}\right\|_2^2\right] \\
&\overset{(c)}{\leq} \sum_{i=0}^{t}(1-\gamma(1-\beta)\mu')^{t-i}\mathbb{E}\left[\|\mathbf{x}\|_2^2\right] \\
&\overset{(d)}{\leq} \frac{\mathbb{E}[\|\mathbf{x}\|_2^2]}{\gamma(1-\beta)\mu'},
\end{aligned}$$

where (a) follows from linearity of the expectation operator, (b) follows from Lemma 6, (c) follows from unrolling the recursion, and (d) follows from the fact that $\sum_{i=1}^{t}(1-\gamma(1-\beta)\mu')^i \leq \frac{1}{\gamma(1-\beta)\mu'}$.

□

Lemma 9. *Under conditions of Theorem 1, we have*

$$z_t^{\text{variance}} \leq \frac{\sigma^2}{\gamma(1-\beta)\mu'}. \tag{44}$$

Proof.

$$z_t^{\text{variance}} = \sum_{i=1}^{t}\mathbb{E}\left[\left(C^{t:i+1}f_i(\theta^\star)\right)^\top\left(C^{t:i+1}f_i(\theta^\star)\right)\right] \tag{45}$$

$$\overset{(a)}{\leq} \frac{1}{\gamma(1-\beta)\mu'}\mathbb{E}\left[\left\|f_t(\theta^\star)\right\|_2^2|\mathcal{F}_{t-1}\right] \tag{46}$$

$$\overset{(b)}{\leq} \frac{\sigma^2}{\gamma(1-\beta)\mu'}, \tag{47}$$

where (a) follows from Lemma 8, and (b) follows from Assumption 4, and the fact that $\mathbb{E}\left[\left\|f_t(\theta^\star)\right\|_2^2|\mathcal{F}_{t-1}\right] \leq \sigma^2$, where $\sigma^2 = (R_{\max} + (1+\beta)\Phi_{\max}^2\left\|\theta^\star\right\|_2^2)$. □

6.2.4. Proof of Theorem 1

Recall that

$$z_{k+1,N} = \frac{1}{N}\sum_{i=k+1}^{k+N} z_i.$$

Now,

$$\mathbb{E}\left[\left\|z_{k+1,N}\right\|_2^2\right] = \frac{1}{N^2}\sum_{i,j=k+1}^{k+N}\mathbb{E}\left[z_i^\top z_j\right]$$
$$\overset{(a)}{\leq} \frac{1}{N^2}\left(\sum_{i=k+1}^{k+N}\mathbb{E}\left[\|z_i\|_2^2\right] + 2\sum_{i=k+1}^{k+N-1}\sum_{j=i+1}^{k+N}\mathbb{E}\left[z_i^\top z_j\right]\right), \tag{48}$$

where (a) follows from separating out the diagonal and off-diagonal terms.
For the bound in expectation for tail-averaged TD in Theorem 1, we establish a useful result that bounds the second term in the RHS of (48).

Lemma 10. *For all $i \geq 1$, we have*

$$\sum_{i=k+1}^{k+N-1}\sum_{j=i+1}^{k+N}\mathbb{E}\left[z_i^\top z_j\right] \leq \frac{2}{\gamma(1-\beta)\mu'}\sum_{i=k+1}^{k+N}\mathbb{E}\left[\|z_i\|_2^2\right]. \tag{49}$$

Proof.

$$\sum_{i=k+1}^{k+N-1}\sum_{j=i+1}^{k+N}\mathbb{E}\left[z_i^\top z_j\right] = \sum_{i=k+1}^{k+N-1}\sum_{j=i+1}^{k+N}\mathbb{E}\left[z_i^\top\left(C^{(j:i+1)}z_i + \gamma\sum_{l=0}^{j-i-1}C^{j:l+1}f_l(\theta^\star)\right)\right]$$
$$\overset{(a)}{=} \sum_{i=k+1}^{k+N-1}\sum_{j=i+1}^{k+N}\mathbb{E}\left[z_i^\top C^{(j:i+1)}z_i\right]$$
$$= \sum_{i=k+1}^{k+N-1}\sum_{j=i+1}^{k+N}\mathbb{E}\left[z_i^\top(\mathbf{I}-\gamma a_j)(C^{(j-1:i+1)}z_i)\right]$$
$$\overset{(b)}{\leq} \sum_{i=k+1}^{k+N-1}\sum_{j=i+1}^{k+N}\mathbb{E}\left[\|z_i\|_2^2\left\|C^{j:i+1}\right\|_2\right]$$
$$\overset{(c)}{\leq} \sum_{i=k+1}^{k+N-1}\sum_{j=i+1}^{k+N}\left(1-\frac{\gamma(1-\beta)\mu'}{2}\right)^{j-i}\mathbb{E}\left[\|z_i\|_2^2\right]$$
$$\leq \sum_{i=k+1}^{k+N}\mathbb{E}\left[\|z_i\|_2^2\right]\sum_{j=i+1}^{\infty}\left(1-\frac{\gamma(1-\beta)\mu'}{2}\right)^{j-i}$$
$$\overset{(d)}{\leq} \frac{2}{\gamma(1-\beta)\mu'}\sum_{i=k+1}^{k+N}\mathbb{E}\left[\|z_i\|_2^2\right],$$

where (a) follows from the fact that $\mathbb{E}[f_t(\theta^\star)|\mathcal{F}_{t-1}] = 0$, (b) follows from the fact that $z^\top C^{j:i+1}z \leq \|z\|_2^2\left\|C^{j:i+1}\right\|_2$ due to Cauchy-Schwarz inequality, (c) is because $\left\|C^{j:i+1}\right\|_2 \leq \left(\|I-\gamma a\|_2\right)^{j-i} \leq \left(1-\frac{\gamma(1-\beta)\mu'}{2}\right)^{j-i}$ (due to Cauchy-Schwarz inequality and Lemma 6), and (d) follows from the summation of the geometric series. □

For the sake of readability, we will restate Theorem 1 below.

Theorem 5. *Suppose Assumptions 1 to 5 hold. Choose a step size γ satisfying*

$$\gamma \leq \gamma_{\max} = \frac{1-\beta}{(1+\beta)^2\Phi_{\max}^2}, \tag{50}$$

where β is the discount factor and $\Phi_{\max}$ is a bound on the features (see Assumption 3).
Then the expected error of the tail-averaged iterate $\theta_{k+1,N}$ formed using Algorithm 1 satisfies

$$\mathbb{E}\left[\left\|\theta_{k+1,N}-\theta^\star\right\|_2^2\right] \leq \frac{10e^{(-k\gamma(1-\beta)\mu')}}{\gamma^2(1-\beta)^2\mu'^2N^2}\mathbb{E}\left[\left\|\theta_0-\theta^\star\right\|_2^2\right] + \frac{10\sigma^2}{(1-\beta)^2\mu'^2N}, \tag{51}$$

where $N = t-k$, $\sigma^2 = (R_{\max}+(1+\beta)\Phi_{\max}^2\left\|\theta^\star\right\|_2^2)$, θ_0 is the initial point, and $\theta^\star$ is the TD fixed point specified in (6).

Proof. Substituting the result of Lemma 10 in eq. (48), we obtain

$$\begin{aligned}
\mathbb{E}\left[\left\|z_{k+1,N}\right\|_2^2\right] &\leq \frac{1}{N^2}\left(\sum_{i=k+1}^{k+N}\mathbb{E}\left[\|z_i\|_2^2\right]+\frac{4}{\gamma(1-\beta)\mu'}\sum_{i=k+1}^{k+N}\mathbb{E}\left[\|z_i\|_2^2\right]\right)\\
&= \frac{1}{N^2}\left(1+\frac{4}{\gamma(1-\beta)\mu'}\right)\sum_{i=k+1}^{k+N}\mathbb{E}\left[\|z_i\|_2^2\right]\\
&\overset{(a)}{\leq} \underbrace{\frac{2}{N^2}\left(1+\frac{4}{\gamma(1-\beta)\mu'}\right)\sum_{i=k+1}^{k+N} z_i^{\text{bias}}}_{z_{k+1,N}^{\text{bias}}} + \underbrace{\frac{2}{N^2}\left(1+\frac{4}{\gamma(1-\beta)\mu'}\right)\gamma^2\sum_{i=k+1}^{k+N} z_i^{\text{variance}}}_{z_{k+1,N}^{\text{variance}}}. \qquad (52)
\end{aligned}$$

where (a) follows from eq. (40).
$z_{k+1,N}^{\text{bias}}$ in eq. (52) is bounded as follows:

$$\begin{aligned}
z_{k+1,N}^{\text{bias}} &\leq \frac{2}{N^2}\left(1+\frac{4}{\gamma(1-\beta)\mu'}\right)\sum_{i=k+1}^{\infty} z_i^{\text{bias}}\\
&\overset{(a)}{\leq} \frac{2}{N^2}\left(1+\frac{4}{\gamma(1-\beta)\mu'}\right)\sum_{i=k+1}^{\infty}(1-\gamma(1-\beta)\mu')^i\mathbb{E}\left[\|z_0\|_2^2\right]\\
&\overset{(b)}{=} \frac{2}{\gamma(1-\beta)\mu' N^2}(1-\gamma(1-\beta)\mu')^{k+1}\left(1+\frac{4}{\gamma(1-\beta)\mu'}\right)\mathbb{E}\left[\|z_0\|_2^2\right],
\end{aligned}$$

where (a) follows from eq. (42) in proof of Lemma 7, and (b) follows from the bound on summation of a geometric series.
$z_{k+1,N}^{\text{variance}}$ in eq. (52) is bounded as follows:

$$\begin{aligned}
z_{k+1,N}^{\text{variance}} &\overset{(a)}{\leq} \frac{2\gamma^2}{N^2}\left(1+\frac{4}{\gamma(1-\beta)\mu'}\right)\sum_{i=k+1}^{k+N}\frac{\sigma^2}{\gamma(1-\beta)\mu'}\\
&\leq \frac{2\gamma^2}{N^2}\left(1+\frac{4}{\gamma(1-\beta)\mu'}\right)\sum_{i=0}^{N}\frac{\sigma^2}{\gamma(1-\beta)\mu'}\\
&= \left(1+\frac{4}{\gamma(1-\beta)\mu'}\right)\frac{2\gamma\sigma^2}{(1-\beta)\mu' N},
\end{aligned}$$

where (a) follows from Lemma 9.
Finally substituting the bounds on $z_{k+1,N}^{\text{bias}}$ and $z_{k+1,N}^{\text{variance}}$ in (52), we get

$$\begin{aligned}
\mathbb{E}[\left\|z_{k+1,N}\right\|_2^2] &\leq \left(1+\frac{4}{\gamma(1-\beta)\mu'}\right)\left(\frac{2}{\gamma(1-\beta)\mu' N^2}(1-\gamma(1-\beta)\mu')^{k+1}\mathbb{E}[\|z_0\|_2^2]+\frac{2\gamma\sigma^2}{(1-\beta)\mu' N}\right),\\
&\overset{(a)}{\leq} \left(1+\frac{4}{\gamma(1-\beta)\mu'}\right)\left(\frac{2\exp(-k\gamma(1-\beta)\mu')}{\gamma(1-\beta)\mu' N^2}\mathbb{E}[\|z_0\|_2^2]+\frac{2\gamma\sigma^2}{(1-\beta)\mu' N}\right)\\
&\overset{(b)}{\leq} \frac{10\exp(-k\gamma(1-\beta)\mu')}{\gamma^2(1-\beta)^2\mu'^2 N^2}\mathbb{E}\left[\|z_0\|_2^2\right]+\frac{10\sigma^2}{(1-\beta)^2\mu'^2 N},
\end{aligned}$$

where (a) follows from the fact that $(1+x)^y = \exp(y\log(1+x)) \leq \exp(xy)$, and (b) uses $\gamma(1-\beta)\mu' < 1$ as $\mu' < \Phi_{\max}^2$, which implies that $1+\frac{4}{\gamma(1-\beta)\mu'} \leq \frac{5}{\gamma(1-\beta)\mu'}$. □

6.3. *High probability bound for tail-averaged TD (Proof of Theorem 2)*

Proposition 1. *Suppose Assumptions 3 and 4 hold, then for all $\epsilon \geq 0$, and $t \geq 1$,*

$$P\left(\left\|z_{k+1,N}\right\|_2 - \mathbb{E}\left[\left\|z_{k+1,N}\right\|_2 > \epsilon\right]\right) \leq \exp\left(-\frac{\epsilon^2}{(R_{\max}+(1+\beta)H\Phi_{\max}^2)^2\sum_{i=k+1}^{k+N}L_i^2}\right),$$

where $L_i \triangleq \frac{\gamma}{N}\sum_{j=i+1}^{i+N}\left(1-\frac{\gamma(1-\beta)\mu'}{2}\right)^{j-i+1}$.

Proof. To derive the result we follow the technique from [17].
Step 1. We decompose the centered error $\left\|z_{k+1,N}\right\|_2 - \mathbb{E}[\left\|z_{k+1,N}\right\|_2]$ as follows:

$$\left\|z_{k+1,N}\right\|_2 - \mathbb{E}\left[\left\|z_{k+1,N}\right\|_2\right] = \sum_{i=k+1}^{k+N} D_i, \tag{53}$$

where $D_i \triangleq g_i - \mathbb{E}[g_i|\mathcal{F}_{i-1}]$, and $g_i \triangleq \mathbb{E}[\left\|z_{k+1,N}\right\|_2|\mathcal{F}_{i-1}]$.
Step 2.
We prove that functions g_i are Lipschitz continuous in the random innovation f_i (given by (27)) at time i with constant L_i.
First, let $\Theta^i_{k+1,N}(\theta)$ to be the value of the tail-averaged iterate at time t that evolves according to eq. (7) beginning from θ at time i. Next, let $\bar{\Theta}^i_{k+1,N}(\bar{\theta},\theta)$ be defined as follows:

$$\bar{\Theta}^i_{k+1,N}(\bar{\theta},\theta) \triangleq \frac{(i-k)\bar{\theta}}{N} + \frac{1}{N}\sum_{j=i+1}^{i+N}\Theta^i_j(\theta),$$

where $\bar{\theta}$ is the value of the tail averaged iterate at time i.
Now let f and f' denote two possible values of the random innovation at time i, and set $\theta = \theta_{i-1} + \gamma f$ and $\theta' = \theta_{i-1} + \gamma f'$. Therefore,

$$\mathbb{E}\left[\left\|\bar{\Theta}^i_{i+1,N}(\bar{\theta},\theta) - \bar{\Theta}^i_{k+1,N}(\bar{\theta},\theta')\right\|_2\right] = \mathbb{E}\left[\frac{1}{N}\sum_{j=i+1}^{i+N}\left\|(\Theta^i_j(\theta) - \Theta^i_j(\theta'))\right\|_2\right]. \tag{54}$$

We will now bound the term $\Theta^i_j(\theta) - \Theta^i_j(\theta')$ inside the summation of (54).
We will first , note that as the projection Γ is non-expansive, we have the following

$$\mathbb{E}\left[\left\|\Theta^i_j(\theta) - \Theta^i_j(\theta')\right\|_2|\mathcal{F}_{j-1}\right] \le \mathbb{E}\left[\left\|\Theta^i_{j-1}(\theta) - \Theta^i_{j-1}(\theta') - \gamma[f_i(\Theta^i_{j-1}(\theta)) - f_i(\Theta^i_{j-1}(\theta'))]\right\|_2|\mathcal{F}_{j-1}\right]. \tag{55}$$

Expanding f_i and using the definition a_j from Lemma 5, we have

$$\begin{aligned}
&\Theta^i_{j-1}(\theta) - \Theta^i_j(\theta') - \gamma[f_i(\Theta^i_{j-1}(\theta)) - f_i(\Theta^i_{j-1}(\theta')] \\
&\quad = \Theta^i_{j-1}(\theta) - \Theta^i_{j-1}(\theta') - \gamma[\phi(s_j)\phi(s_j)^\top - \beta\phi(s_j)\phi(s_j')^\top][(\Theta^j_{i-1}(\theta)) - (\Theta^j_{i-1}(\theta')] \\
&\quad = [\mathbf{I} - \gamma a_j](\Theta^i_{j-1}(\theta) - \Theta^i_{j-1}(\theta')).
\end{aligned}$$

Using the tower property of conditional expectations, it follows that

$$\begin{aligned}
\mathbb{E}[\left\|\Theta^i_j(\theta) - \Theta^i_j(\theta')\right\|_2] &= \mathbb{E}\left[\mathbb{E}\left[\|\Theta^i_j(\theta) - \Theta^i_j(\theta')\|_2\middle|\mathcal{F}_{j-1}\right]\right] \\
&= \mathbb{E}\left[\mathbb{E}\left[\left\|(\mathbf{I} - \gamma a_j)(\Theta^i_{j-1}(\theta) - \Theta^i_{j-1}(\theta'))\right\|_2\middle|\mathcal{F}_{j-1}\right]\right] \\
&\overset{(a)}{\le} \left(1 - \frac{\gamma(1-\beta)\mu'}{2}\right)\mathbb{E}[\left\|\Theta^i_{j-1}(\theta) - \Theta^i_{j-1}(\theta')\right\|_2] \\
&\overset{(b)}{=} \left(1 - \frac{\gamma(1-\beta)\mu'}{2}\right)^{j-i+1}\|\theta - \theta'\|_2,
\end{aligned}$$

where (a) follows from Lemma 6, and (b) follows from repeated application of argument used in arriving at (a).
Now, using Jensen's inequality, we get

$$\begin{aligned}
|\mathbb{E}[\left\|\theta_j - \theta^\star\right\|_2|\theta_i = \theta] - \mathbb{E}[\left\|\theta_j - \theta^\star\right\|_2|\theta_i = \theta']| &\le \mathbb{E}\left[\left\|(\Theta^i_j(\theta) - \Theta^i_j(\theta'))\right\|_2\right], \\
&\le \gamma\left(1 - \frac{\gamma(1-\beta)\mu'}{2}\right)^{j-i+1}\|f - f'\|_2.
\end{aligned} \tag{56}$$

Substituting (56) in (54), we obtain

$$\mathbb{E}[\left\|\bar{\Theta}^i_{k+1}(\bar{\theta},\theta)-\bar{\Theta}^i_{k+1}(\bar{\theta},\theta')\right\|_2] \leq \frac{\gamma}{N}\sum_{j=i+1}^{i+N}\left(1-\frac{\gamma(1-\beta)\mu'}{2}\right)^{j-i+1}\|f-f'\|_2. \tag{57}$$

From (56), it is clear that g_i is L_i- Lipschtiz in the random innovation at time i, which implies that D_i is Lipschitz with the constant L_i.

Step 3.

Next, we derive a standard martingale concentration bound for the tail-averaged iterate $z_{k+1,N}$. For any $\eta > 0$,

$$\begin{aligned}
P(\left\|z_{k+1,N}\right\|_2^2 - \mathbb{E}[\left\|z_{k+1,N}\right\|_2^2] \geq \epsilon) &= P\left(\sum_{i=k+1}^{k+N} D_i \geq \epsilon\right) \\
&\overset{(a)}{\leq} \exp(-\eta\epsilon)\mathbb{E}\left[\exp(\eta\sum_{i=k+1}^{k+N} D_i)\right] \\
&\overset{(b)}{=} \exp(-\eta\epsilon)\mathbb{E}\left[\exp(\eta\sum_{i=k+1}^{k+N} D_i)\mathbb{E}\left[\exp(\eta D_{k+1,N})|\mathcal{F}_{t-1}\right]\right],
\end{aligned}$$

where are (a) follows from Markov inequality and (b) follows from eq. (53).

From [17, Proof of proposition 1 part 3, page 585], we have the following bound for a zero-mean r.v. Z with $|Z| \leq B$ w.p 1, and a L-Lipschitz function g:

$$\mathbb{E}[\exp(\eta g(Z))] \leq \exp\left(\frac{\eta^2 B^2 L^2}{2}\right).$$

Next, from assumption 4, and the projection step of the algorithm we have that $f_i(\theta_{i-1}) < (R_{\max} + (1+\beta)H\Phi^2_{\max})$ is a bounded random variable, and conditioned on $\mathcal{F}_{i-1}$, D_i is Lipshitz in $f_i(\theta_{i-1})$ with constant L_i. Hence,

$$\mathbb{E}[\exp(\eta D_t)|\mathcal{F}_{t-1}] \leq \exp\left(\frac{\eta^2(R_{\max}+(1+\beta)H\Phi^2_{\max})^2 L_t^2}{2}\right).$$

Using these facts we get,

$$P\left(\left\|z_{k+1,N}\right\|_2 - \mathbb{E}\left[\left\|z_{k+1,N}\right\|_2\right] > \epsilon\right) \leq \exp(-\eta\epsilon)\exp\left(\frac{\eta^2(R_{\max}+(1+\beta)H\Phi^2_{\max})^2\sum_{i=k+1}^{k+N}L_i^2}{2}\right).$$

Optimising over η in the above inequality leads to

$$P\left(\left\|z_{k+1,N}\right\|_2 - \mathbb{E}\left[\left\|z_{k+1,N}\right\|_2\right] > \epsilon\right) \leq \exp\left(-\frac{\epsilon^2}{(R_{\max}+(1+\beta)H\Phi^2_{\max})^2\sum_{i=k+1}^{k+N}L_i^2}\right).$$

□

6.3.1. Bounding the Lipschitz constant

Lemma 11. *With L_i as defined in Proposition 1, we have*

$$\sum_{i=k+1}^{k+N} L_i^2 \leq \frac{4}{N(1-\beta)^2\mu'^2}.$$

Proof. Notice that

$$\begin{aligned}
\sum_{i=k+1}^{k+N} L_i^2 &= \frac{\gamma^2}{N^2}\sum_{i=k+1}^{k+N}\left(\sum_{j=i+1}^{i+N}\left(1-\frac{\gamma(1-\beta)\mu'}{2}\right)^{j-i+1}\right)^2 \\
&\leq \frac{\gamma^2}{N^2}\sum_{i=k+1}^{k+N}\left(\sum_{j=i+1}^{\infty}\left(1-\frac{\gamma(1-\beta)\mu'}{2}\right)^{j-i+1}\right)^2 \\
&= \frac{\gamma^2}{N^2}\sum_{i=k+1}^{k+N}\left(\frac{4}{\gamma^2(1-\beta)^2\mu'^2}\right) \\
&= \frac{4}{N(1-\beta)^2\mu'^2}.
\end{aligned}$$

□

6.3.2. Proof of Theorem 2

For the sake of convenience, we restate the high probability bound for tail-averaged TD below.

Theorem 6 (High-probability bound). *Suppose Assumptions 1 to 6 hold. Choose the step size such that $\gamma \leq \gamma_{\max}$, where $\gamma_{\max}$ is defined in (14). Then, for any $\delta \in (0,1]$, we have the following bound for the projected tail-averaged iterate $\theta_{k+1,N}$:*

$$P\left(\left\|\theta_{k+1,N}-\theta^\star\right\|_2 \leq \frac{2\sigma}{(1-\beta)\mu'\sqrt{N}}\sqrt{\log\left(\frac{1}{\delta}\right)}+\frac{\sqrt{10}e^{(-k\gamma(1-\beta)\mu')}}{\gamma(1-\beta)\mu' N}\mathbb{E}\left[\left\|\theta_0-\theta^\star\right\|_2\right]+\frac{\sqrt{10}\sigma}{(1-\beta)\mu'\sqrt{N}}\right)\geq 1-\delta,$$

where $N, \sigma, \mu, \theta_0, \theta^\star$ are as specified in Theorem 1.

Proof.

$$\begin{aligned}P(\left\|z_{k+1,N}\right\|_2^2-\mathbb{E}[\left\|z_{k+1,N}\right\|_2^2]>\epsilon) &\overset{(a)}{\leq} \exp\left(-\frac{\epsilon^2}{(R_{\max}+(1+\beta)H\Phi_{\max}^2)^2\sum_{i=k+1}^{k+N}L_i^2}\right)\\ &\overset{(b)}{\leq} \exp\left(-\frac{N(1-\beta)\mu'^2\epsilon^2}{4\sigma^2}\right),\end{aligned} \tag{58}$$

where (a) follows from Proposition 1, and (b) follows from Lemma 11.
The inequality in eq. (58) can be re-written in high-confidence form as follows: For any $\delta \in (0,1]$,

$$P\left(\left\|z_{k+1,N}\right\|_2^2-\mathbb{E}[\left\|z_{k+1,N}\right\|_2^2]\leq \frac{2\sigma}{(1-\beta)\mu'\sqrt{N}}\sqrt{\log\left(\frac{1}{\delta}\right)}\right)\geq 1-\delta.$$

The final bound follows by substituting the bound on $\mathbb{E}\left[\left\|z_{k+1,N}\right\|_2^2\right]$ from Theorem 1 in the inequality above. □

6.4. Expectation bound for Tail-averaged TD with regularisation

We will first begin by rewriting the regularised TD iteration as follows:

$$\hat{\theta}_{t+1} = (\mathbf{I}-\gamma\lambda)\hat{\theta}_t+\gamma(-A\hat{\theta}_t+b\Delta M_t),$$

where $\Delta M_t = \hat{f}_t(\hat{\theta}_{t-1})-\mathbb{E}^{\rho,P}\left[\hat{f}_t(\hat{\theta}_{t-1})|\mathcal{F}_{t-1}\right]$.
Recall that the solution for the regularised TD iteration is given by

$$\theta^\star_{\text{reg}} = (A+\lambda\mathbf{I})^{-1}b,$$

where A and b are defined in eq. (6).

Lemma 12. *With $\gamma \leq \gamma_{\max} = \frac{\lambda}{\lambda^2+2\lambda(1+\beta)\Phi_{\max}^2+(1+\beta)^2\Phi_{\max}^4}$, we have*

$$\left\|\left(\mathbf{I}-\gamma(A+\lambda\mathbf{I})\right)^\top\left(\mathbf{I}-\gamma(A+\lambda\mathbf{I})\right)\right\|_2 \leq 1-\gamma(\mu+\lambda),$$

and

$$\|(\mathbf{I}-\gamma A)\|_2 \leq 1-\frac{\gamma(\mu+\lambda)}{2}.$$

Proof.

$$\begin{aligned}\left\|(\mathbf{I}-\gamma(A+\lambda\mathbf{I}))^\top(\mathbf{I}-\gamma(A+\lambda\mathbf{I}))\right\|_2 &= \left\|\mathbf{I}-\gamma(A+\lambda\mathbf{I})-\gamma(A+\lambda\mathbf{I})^\top+\gamma^2(A+\lambda\mathbf{I})^\top(A+\lambda\mathbf{I})\right\|_2\\ &= \left\|\mathbf{I}-\gamma(A+A^\top+2\lambda\mathbf{I})+\gamma^2(A+\lambda\mathbf{I})^\top(A+\lambda\mathbf{I})\right\|_2\\ &\leq \left\|\mathbf{I}-\gamma(A+A^\top+2\lambda\mathbf{I})\right\|_2+\left\|\gamma^2(A+\lambda\mathbf{I})^\top(A+\lambda\mathbf{I})\right\|_2\\ &= \left\|\mathbf{I}-\gamma(A+A^\top+2\lambda\mathbf{I})\right\|_2+\gamma^2\left\|A^\top A+\lambda(A^\top+A)+\lambda^2\mathbf{I}\right\|_2\end{aligned}$$

$$\overset{(a)}{\leq} 1-\gamma\big(2\mu+2\lambda-\gamma(\lambda^2+2\lambda(1+\beta)\Phi_{\max}^2+(1+\beta)^2\Phi_{\max}^4)\big)$$
$$\overset{(b)}{\leq} 1-(2\mu+\lambda)\gamma$$
$$\leq 1-\gamma(\mu+\lambda),$$

where (a) follows from Lemma 3, and $\|(A+A^\top)\|_2^2 \geq 2\mu$ and (b) follows as per definition of γ.
For the second claim, note that $1-\gamma(\mu+\lambda) \geq 0 \implies (1-\gamma(\mu+\lambda))^{\frac{1}{2}} \leq 1-\frac{\gamma(\mu+\lambda)}{2}$. Therefore,

$$\|(\mathbf{I}-\gamma A)\|_2 \leq 1-\frac{\gamma(\mu+\lambda)}{2}.$$

□

Lemma 13. *Let $\gamma \leq \gamma_{\max}$, where $\gamma_{\max}$ is set as per (19), and a_j be as defined in Lemma 5. Then for any $\mathcal{F}_{i-1}$ measurable $\hat{\theta} \in \mathbb{R}^d$ we have*

$$\mathbb{E}\left[\hat{\theta}^\top(\mathbf{I}-\gamma(\lambda\mathbf{I}+a_j))^\top(\mathbf{I}-\gamma(\lambda\mathbf{I}+a_j))\hat{\theta}|\mathcal{F}_j\right] \leq (1-\gamma(\mu+\lambda))\left\|\hat{\theta}\right\|_2^2$$
and,
$$\mathbb{E}\left[\left\|(\mathbf{I}-\gamma(\lambda\mathbf{I}+a_j))\hat{\theta}\right\|_2|\mathcal{F}_j\right] \leq \left(1-\frac{\gamma(\mu+\lambda)}{2}\right)\left\|\hat{\theta}\right\|_2.$$

Proof.

$$\begin{aligned}
&\mathbb{E}[\hat{\theta}^\top(\mathbf{I}-\gamma(\lambda\mathbf{I}+a_j))^\top(\mathbf{I}-\gamma(\lambda\mathbf{I}+a_j))\hat{\theta}|\mathcal{F}_j] \\
&\quad= \mathbb{E}[\hat{\theta}^\top(\mathbf{I}-2\gamma\lambda\mathbf{I}-\gamma(a_j+a_j^\top)+\gamma^2\{\lambda^2\mathbf{I}+\lambda(a_j+a_j^\top)+a_j^\top a_j\})\hat{\theta}|\mathcal{F}_j] \\
&\quad= \mathbb{E}[\hat{\theta}^\top\hat{\theta}|\mathcal{F}_j]-\gamma\mathbb{E}[\hat{\theta}^\top 2\lambda\mathbf{I}\hat{\theta}|\mathcal{F}_j]-\underbrace{\gamma\,\hat{\theta}^\top\mathbb{E}[a_j^\top+a_j|\mathcal{F}_j]\hat{\theta}}_{\text{Term 1}}+ \\
&\gamma^2\underbrace{\hat{\theta}^\top\mathbb{E}[a_j^\top a_j|\mathcal{F}_j]\hat{\theta}}_{\text{Term 2}}+\gamma^2\underbrace{\hat{\theta}^\top\mathbb{E}[a_j+a_j^\top|\mathcal{F}_j]\hat{\theta}}_{\text{Term 3}}+\gamma^2\mathbb{E}[\hat{\theta}^\top\lambda^2\mathbf{I}\hat{\theta}|\mathcal{F}_j].
\end{aligned}$$

We proceed by bounding the Terms 1, 2, and 3 individually.
Term 1:

$$\begin{aligned}
\theta^\top\mathbb{E}[a_j^\top+a_j|\mathcal{F}_j]\theta &= \theta^\top(A^\top+A)\theta \\
&\overset{(a)}{\geq} 2\mu\left\|\hat{\theta}\right\|_2^2,
\end{aligned} \tag{59}$$

where, μ is the minimum eigenvalue of matrix A.

Term 2:
From Lemma 5 for any $\hat{\theta} \in \mathbb{R}^d$, we have

$$\begin{aligned}
\theta^\top(\mathbb{E}[a_j^\top a_j])\theta &\leq \Phi_{\max}^2(1+\beta)^2\hat{\theta}^\top B\hat{\theta} \\
&\overset{(a)}{\leq} \Phi_{\max}^2(1+\beta)^2\left\|\hat{\theta}\right\|_2^2\|B\|_2 \\
&\overset{(b)}{\leq} (1+\beta)^2\Phi_{\max}^4\left\|\hat{\theta}\right\|_2^2,
\end{aligned}$$

where (a) follows from the Cauchy-Schwarz inequality, and (b) follows from Assumption 3.

Term 3:

$$\begin{aligned}
\hat{\theta}^\top\mathbb{E}[a_j+a_j^\top|\mathcal{F}_j]\hat{\theta} &= \hat{\theta}^\top(A+A^\top)\hat{\theta} \\
&\overset{(a)}{\leq} \left\|\hat{\theta}\right\|_2^2\left\|A+A^\top\right\|_2 \\
&\overset{(b)}{\leq} 2\left\|\hat{\theta}\right\|_2^2(1+\beta)\Phi_{\max}^2,
\end{aligned}$$

where (a) follows from Cauchy-Schwarz inequality and (b) follows from Lemma 3.
Combining the bounds on Term 1,2, and 3, we obtain

$$\begin{aligned}
&\mathbb{E}[\hat{\theta}^\top(\mathbf{I}-\gamma(\lambda\mathbf{I}+a_j))^\top(\mathbf{I}-\gamma(\lambda\mathbf{I}+a_j))\hat{\theta}|\mathcal{F}_j] \\
&\qquad \leq (1-\gamma(2\mu+2\lambda-\gamma(\lambda^2+2(1+\beta)\Phi_{\max}^2+(1+\beta)^2\Phi_{\max}^4)))\left\|\hat{\theta}\right\|_2^2 \\
&\qquad \overset{(a)}{\leq} (1-\gamma(2\mu+\lambda))\left\|\hat{\theta}\right\|_2^2, \qquad (60)
\end{aligned}$$

where (a) follows from using the using the value of $\gamma_{\max}$ given in (19).
Using Jensen's inequality along with (60), we have

$$\begin{aligned}
\mathbb{E}\left[\left\|(\mathbf{I}-\gamma(\lambda\mathbf{I}+a_j))\hat{\theta}\right\|_2|\mathcal{F}_j\right] &\leq (1-\gamma(\mu+\lambda))^{1/2}\left\|\hat{\theta}\right\|_2 \\
&\overset{(a)}{\leq} \left(1-\frac{\gamma(\mu+\lambda)}{2}\right)\left\|\hat{\theta}\right\|_2,
\end{aligned}$$

where (a) uses the following fact: $1-\gamma(\mu+\lambda)\geq 0 \implies (1-\gamma(\mu+\lambda))^{\frac{1}{2}} \leq 1-\frac{\gamma(\mu+\lambda)}{2}$. □

6.4.1. Bias-variance decomposition of the non-asymptotic error

Let

$$\hat{z}_t^{\text{bias}} \triangleq \mathbb{E}\left[\left\|C^t\hat{z}_0\right\|_2^2\right], \qquad (61)$$

and

$$\hat{z}_t^{\text{variance}} \triangleq \sum_{k=0}^{t}\mathbb{E}\left[\left\|C^k\Delta M_{t-k}\right\|_2^2\right]. \qquad (62)$$

Now define the centered error rule as $\hat{z}_t = \hat{\theta}_t - \theta^\star_{\text{reg}}$. Using eq. (17), we obtain

$$\begin{aligned}
\hat{z}_t &= (\mathbf{I}-\gamma\lambda)\hat{\theta}_{t-1}-\theta^\star_{\text{reg}}+\gamma\left(-A\hat{\theta}_{t-1}+b+\Delta M_t\right) \\
&= (\mathbf{I}-\gamma\lambda)\hat{\theta}_{t-1}-\theta^\star_{\text{reg}}+\gamma\left(-A\hat{\theta}_{t-1}+(A+\lambda\mathbf{I})(A+\lambda\mathbf{I})^{-1}b+\Delta M_t\right) \\
&= (\mathbf{I}-\gamma(A+\lambda\mathbf{I}))\hat{z}_{t-1}+\gamma\Delta M_t \\
&\overset{(a)}{=} C^t\hat{z}_0+\gamma\sum_{k=0}^{t}C^k\Delta M_{t-k}, \qquad (63)
\end{aligned}$$

where $C=(\mathbf{I}-\gamma(A+\lambda\mathbf{I}))$, and (a) follows from unrolling the update rule.
Taking expectation on both sides of eq. (63) we obtain

$$\begin{aligned}
\mathbb{E}\left[\|\hat{z}_t\|_2^2\right] &= \mathbb{E}\left[\left\|C^t\hat{z}_0+\gamma\sum_{k=0}^{t}C^k\Delta M_{t-k}\right\|_2^2\right] \\
&\overset{(a)}{\leq} 2\hat{z}_t^{\text{bias}}+2\gamma^2\hat{z}_t^{\text{variance}}, \qquad (64)
\end{aligned}$$

where (a) is obtained by using the following inequalities i. $\|a+b\|_2^2 \leq 2\|a\|_2^2+2\|b\|_2^2$, ii. $\left\|\sum_{i=0}^n \mathbf{x}_i\right\| \leq \sum_{i=0}^n\|\mathbf{x}_i\|$, and $\hat{z}_t^{\text{bias}}$ & $\hat{z}_t^{\text{variance}}$ are defined in eqs. (61) and (62)
Therefore, a bound on $\mathbb{E}\|\hat{z}_t\|_2^2$ can obtained by bounding individual terms in eq. (64).

6.4.2. Bounding $\hat{z}_t^{\text{bias}}$

Lemma 14. *For any step size $\gamma \leq \gamma_{\max}$, the bias or the initial error of the TD update is upper bounded as*

$$\hat{z}_t^{\text{bias}} \leq \exp(-\gamma(\mu+\lambda)t)\mathbb{E}[\|\hat{z}_0\|_2^2].$$

Proof.

$$\hat{z}_t^{\text{bias}} = \mathbb{E}[\left\|C^{t:1}\hat{z}_0\right\|_2^2]$$

$$
\begin{aligned}
&= \mathbb{E}\left[(C^{t-1:1}\hat{z}_0)^\top(\mathbf{I}-\gamma(A+\lambda\mathbf{I}))^\top(\mathbf{I}-\gamma(A+\lambda\mathbf{I}))C^{t-1:1}\hat{z}_0\right] \\
&\overset{(a)}{\leq} \left\|(\mathbf{I}-\gamma(\mathbf{I}-\gamma(A+\lambda\mathbf{I})))^\top(\mathbf{I}-\gamma(A+\lambda\mathbf{I}))\right\|_2 \mathbb{E}[\left\|C^{t-1:1}\hat{z}_0\right\|_2^2], \\
&\overset{(b)}{\leq} (1-\gamma(\mu+\lambda))\mathbb{E}[\left\|C^{t-1:1}\hat{z}_0\right\|_2^2], \qquad (65) \\
&\overset{(c)}{=} (1-\gamma(\mu+\lambda))^t\mathbb{E}[\|\hat{z}_0\|_2^2], \qquad (66) \\
&\leq \exp(-\gamma(\mu+\lambda)t)\mathbb{E}[\|\hat{z}_0\|_2^2],
\end{aligned}
$$

where (a) follows Cauchy-Schwarz inequality, (b) follows from Lemma 6, and (c) is the recursive application of the bound in (b). □

6.4.3. Bounding $\hat{z}_t^{\text{variance}}$

Lemma 15. *For any random vector* $\mathbf{x} \in \mathbb{R}^d$,

$$
\sum_{i=0}^{t}\mathbb{E}[\left\|C^i\mathbf{x}\right\|_2^2] \leq \frac{\mathbb{E}[\|\mathbf{x}\|_2^2]}{\gamma(\mu+\lambda)}.
$$

Proof.

$$
\begin{aligned}
\mathbb{E}\left[\sum_{i=0}^{t}\left\|C^i\mathbf{x}\right\|_2^2\right] &= \mathbb{E}\left[\sum_{i=0}^{t}\left(C^i\mathbf{x}\right)^\top\left(C^i\mathbf{x}\right)\right] \\
&= \sum_{i=0}^{t}\mathbb{E}\left[\left(\mathbf{x}^\top(C^i)^\top C^i\mathbf{x}\right)\right] \\
&\overset{(a)}{\leq} \sum_{i=0}^{t}\left\|(\mathbf{I}-\gamma(\mathbf{I}-\gamma(A+\lambda\mathbf{I})))^\top(\mathbf{I}-\gamma(\mathbf{I}-\gamma(A+\lambda\mathbf{I})))\right\|_2\mathbb{E}\left[\mathbf{x}^\top(C^{i-1})^\top C^{i-1}\mathbf{x}\right] \\
&\overset{(b)}{\leq} \sum_{i=0}^{t}(1-\gamma(\mu+\lambda))\mathbb{E}\left[\mathbf{x}^\top(C^{i-1})^\top C^{i-1}\mathbf{x}\right] \\
&\overset{(c)}{\leq} \sum_{i=0}^{t}(1-\gamma(\mu+\lambda))^i\mathbb{E}\left[\|\mathbf{x}\|_2^2\right] \\
&\overset{(d)}{\leq} \frac{\mathbb{E}\left[\|\mathbf{x}\|_2^2\right]}{\gamma(\mu+\lambda)},
\end{aligned}
$$

Where (a) follows from Cauchy-Schwarz inequality, (b) follows from Lemma 6 and definition of C^k, (c) follows from unrolling the recursion, and (d) follows from the fact that $\sum_{i=1}^{t}(1-\gamma(\mu+\lambda))^i \leq \frac{1}{\gamma(\mu+\lambda)}$.

□

Lemma 16.

$$
\hat{z}_t^{\text{variance}} \leq \frac{\sigma^2}{\gamma(\mu+\lambda)}.
$$

Proof.

$$
\begin{aligned}
\hat{z}_t^{\text{variance}} &= \sum_{i=1}^{t}\mathbb{E}\left[\left(C^i\Delta M_{i-t}\right)^\top\left(C^i\Delta M_{i-t}\right)\right] \\
&\overset{(a)}{\leq} \frac{1}{(\mu+\lambda)}\mathbb{E}[\|\Delta M_t\|_2^2|\mathcal{F}_{t-1}] \\
&\overset{(b)}{\leq} \frac{\sigma^2}{\gamma(\mu+\lambda)},
\end{aligned}
$$

where (a) follows from Lemma 15, and (b) follows from Assumption 4, and the fact that $\mathbb{E}[\|\Delta M_t\|_2^2|\mathcal{F}_{t-1}] \leq \sigma^2$, where $\sigma^2 = (R_{\max} + (1+\beta)\Phi_{\max}^2\left\|\theta^\star\right\|_2^2)$. □

6.4.4. Proof of Theorem 3

We will first proceed by getting the bias variance decomposition of the tail-averaged regularised TD update.

$$\mathbb{E}[\left\|\hat{z}_{k+1,N}\right\|_2^2] = \frac{1}{N^2}\sum_{i,j=k+1}^{k+N}\mathbb{E}[\hat{z}_i^\top \hat{z}_j]$$
$$\overset{(a)}{\leq} \frac{1}{N^2}\left(\sum_{i=k+1}^{k+N}\mathbb{E}[\|\hat{z}_i\|_2^2] + 2\sum_{i=k+1}^{k+N-1}\sum_{j=i+1}^{k+N}\mathbb{E}[\hat{z}_i^\top \hat{z}_j]\right), \tag{67}$$

where (a) follows from separating out the diagonal and off-diagonal terms and using Cauchy-Schwarz inequality.

Lemma 17. *For all $i \geq 1$, we have*

$$\sum_{i=k+1}^{k+N-1}\sum_{j=i+1}^{k+N}\mathbb{E}[\hat{z}_i^\top \hat{z}_j] \leq \frac{2}{\gamma(\mu+\lambda)}\sum_{i=k+1}^{k+N}\mathbb{E}[\|\hat{z}_i\|_2^2].$$

Proof.

$$\sum_{i=k+1}^{k+N-1}\sum_{j=i+1}^{k+N}\mathbb{E}[\hat{z}_i^\top \hat{z}_j] = \sum_{i=k+1}^{k+N-1}\sum_{j=i+1}^{k+N}\mathbb{E}[\hat{z}_i^\top (C^{(j-i)}\hat{z}_i + \gamma\sum_{l=0}^{j-i-1} C^l \Delta M_{j-l}))]$$
$$\overset{(a)}{=} \sum_{i=k+1}^{k+N-1}\sum_{j=i+1}^{k+N}\mathbb{E}[\hat{z}_i^\top C^{(j-i)}\hat{z}_i]$$
$$\leq \sum_{i=k+1}^{k+N-1}\sum_{j=i+1}^{k+N}\mathbb{E}[\|\hat{z}_i\|_2^2\left\|C^{(j-i)}\right\|_2]$$
$$\overset{(b)}{\leq} \sum_{i=k+1}^{k+N-1}\sum_{j=i+1}^{k+N}(1-\gamma(\mu+\lambda))^{j-i}\mathbb{E}[\|\hat{z}_i\|_2^2]$$
$$\leq \sum_{i=k+1}^{k+N}\mathbb{E}[\|\hat{z}_i\|_2^2]\sum_{j=i+1}^{\infty}(1-\gamma(\mu+\lambda))^{j-i}$$
$$\overset{(c)}{\leq} \frac{2}{\gamma(\mu+\lambda)}\sum_{i=k+1}^{k+N}\mathbb{E}[\|\hat{z}_i\|_2^2],$$

where (a) follows from the fact that $\mathbb{E}[\Delta M_t] = 0$, (b) follows from Lemma 12, (c) follows from the summation of the geometric series. □

For the sake of convenience, we restate the bound in expectation for regularised TD below.

Theorem 7 (Bound in expectation). *Suppose Assumptions 1 to 4 hold. Choose a step size γ satisfying*

$$\gamma \leq \gamma_{\max} = \frac{\lambda}{\lambda^2 + 2\lambda(1+\beta)\Phi_{\max}^2 + (1+\beta)^2\Phi_{\max}^4}. \tag{68}$$

Then the expected error of the tail-averaged regularised TD iterate $\hat{\theta}_{k+1,N}$ satisfies

$$\mathbb{E}[\left\|\hat{\theta}_{k+1,N} - \theta^\star_{\text{reg}}\right\|_2^2] \leq \frac{10e^{(-k\gamma(\mu+\lambda))}}{\gamma^2(\mu+\lambda)^2N^2}\mathbb{E}\left[\left\|\hat{\theta}_0 - \theta^\star_{\text{reg}}\right\|_2^2\right] + \frac{10\sigma^2}{(\mu+\lambda)^2N}, \tag{69}$$

where $N = t-k$, and $\sigma^2 = (R_{\max} + (1+\beta)\Phi_{\max}^2\left\|\hat{\theta}^\star\right\|_2^2)$.

Proof. Substituting the result of Lemma 17 in eq. (67) we get

$$\mathbb{E}[\left\|\hat{z}_{k+1,N}\right\|_2^2] \leq \frac{1}{N^2}\left(\sum_{i=k+1}^{k+N}\mathbb{E}[\|\hat{z}_i\|_2^2] + \frac{4}{\gamma(\mu+\lambda)}\sum_{i=k+1}^{k+N}\mathbb{E}[\|\hat{z}_i\|_2^2]\right)$$
$$= \frac{1}{N^2}\left(1 + \frac{4}{\gamma(\mu+\lambda)}\right)\sum_{i=k+1}^{k+N}\mathbb{E}[\|\hat{z}_i\|_2^2]$$

$$\overset{(a)}{\leq} \underbrace{\frac{2}{N^2}\left(1+\frac{4}{\gamma(\mu+\lambda)}\right)\sum_{i=k+1}^{k+N}\hat{z}_i^{\text{bias}}}_{\hat{z}_{k+1,N}^{\text{bias}}} + \underbrace{\frac{2}{N^2}\left(1+\frac{4}{\gamma(\mu+\lambda)}\right)\gamma^2\sum_{i=k+1}^{k+N}\hat{z}_i^{\text{variance}}}_{\hat{z}_{k+1,N}^{\text{variance}}}. \tag{70}$$

where (a) follows from eq. (64), and $N = t-k$.
$\hat{z}_{k+1,N}^{\text{bias}}$ in eq. (70) is bounded as follows

$$\begin{aligned}
\hat{z}_{k+1,N}^{\text{bias}} &\leq \frac{2}{N^2}\left(1+\frac{4}{\gamma(\mu+\lambda)}\right)\sum_{i=k+1}^{\infty}\hat{z}_i^{\text{bias}} \\
&\overset{(a)}{\leq} \frac{2}{N^2}\left(1+\frac{4}{\gamma(\mu+\lambda)}\right)\sum_{i=k+1}^{\infty}(1-\gamma(\mu+\lambda))^i\mathbb{E}[\|\hat{z}_0\|_2^2] \\
&\overset{(b)}{=} \frac{2}{\gamma(\mu+\lambda)N^2}(1-\gamma(\mu+\lambda))^{k+1}\left(1+\frac{4}{\gamma(\mu+\lambda)}\right)\mathbb{E}[\|\hat{z}_0\|_2^2],
\end{aligned}$$

where (a) follows from eq. (66) in proof of Lemma 14, and (b) follows from the summation of the geometric series.
$\hat{z}_{k+1,N}^{\text{variance}}$ in eq. (70) is bounded as follows

$$\begin{aligned}
\hat{z}_{k+1,N}^{\text{variance}} &\overset{(a)}{\leq} \frac{2\gamma^2}{N^2}\left(1+\frac{4}{\gamma(\mu+\lambda)}\right)\sum_{i=k+1}^{k+N}\frac{\sigma^2}{\gamma(\mu+\lambda)} \\
&\leq \frac{2\gamma}{N^2}\left(1+\frac{4}{\gamma(\mu+\lambda)}\right)\sum_{i=0}^{N}\frac{\sigma^2}{(\mu+\lambda)} \\
&= \left(1+\frac{4}{\gamma(\mu+\lambda)}\right)\frac{2\gamma\sigma^2}{(\mu+\lambda)N},
\end{aligned}$$

where (a) follows from Lemma 16.
Finally substituting the bounds on $\hat{z}_{k+1,N}^{\text{bias}}$ and $\hat{z}_{k+1,N}^{\text{variance}}$ in (70), we get

$$\begin{aligned}
\mathbb{E}\left[\left\|\hat{z}_{k+1,N}\right\|_2^2\right] &\leq \left(1+\frac{4}{\gamma(\mu+\lambda)}\right)\left(\frac{2}{\gamma(\mu+\lambda)N^2}(1-\gamma(\mu+\lambda))^{k+1}\mathbb{E}\left[\|\hat{z}_0\|_2^2\right]+\frac{2\gamma\sigma^2}{(\mu+\lambda)N}\right), \\
&\overset{(a)}{\leq} \left(1+\frac{4}{\gamma(\mu+\lambda)}\right)\left(\frac{2\exp(-k\gamma(\mu+\lambda))}{\gamma(\mu+\lambda)N^2}\mathbb{E}[\|\hat{z}_0\|_2^2]+\frac{2\gamma\sigma^2}{(\mu+\lambda)N}\right) \\
&\overset{(b)}{\leq} \frac{10\exp(-k\gamma(\mu+\lambda))}{\gamma^2(\mu+\lambda)^2N^2}\mathbb{E}\left[\|\hat{z}_0\|_2^2\right]+\frac{10\sigma^2}{(\mu+\lambda)^2N},
\end{aligned}$$

where (a) is because $(1+x)^y = \exp(y\log(1+x)) \leq \exp(xy)$, and (b) is because $1+\frac{4}{\gamma(\mu+\lambda)} \leq \frac{5}{\gamma(\mu+\lambda)}$ □

Corollary 3. Under conditions of Theorem 3, we have

$$\mathbb{E}\left[\left\|\hat{\theta}_{k+1,N}-\theta^\star\right\|_2^2\right] \leq \frac{20e^{(-k\gamma(\mu+\lambda))}}{\gamma^2(\mu+\lambda)^2N^2}\mathbb{E}\left[\left\|\hat{\theta}_0-\theta_{\text{reg}}^\star\right\|_2^2\right]+\frac{20\sigma^2}{(\mu+\lambda)^2N}+\frac{2\lambda^2\Phi_{\max}^2R_{\max}^2}{\sigma_{\min}(A)^2(\lambda+\mu)^2},$$

where $\sigma_{\min}(A)$ as A's minimum singular value.

Proof.

$$\mathbb{E}\left[\left\|\hat{\theta}_{k+1,N}-\theta^\star\right\|_2^2\right] \overset{(a)}{\leq} \underbrace{2\mathbb{E}\left[\left\|\theta_{\text{reg}}^\star-\theta^\star\right\|_2^2\right]}_{\text{Term 1}} + \underbrace{2\mathbb{E}\left[\left\|\hat{\theta}_{k+1,N}-\theta_{\text{reg}}^\star\right\|_2^2\right]}_{\text{Term 2}}), \tag{71}$$

where (a) is because $\|a+b\|_2^2 \leq 2\|a\|_2^2+2\|b\|_2^2$.
Bounding Term 1

$$\mathbb{E}\left[\left\|\theta^\star-\theta_{\text{reg}}^\star\right\|_2^2\right] = \left\|A^{-1}b-(A+\lambda\mathbf{I})^{-1}b\right\|_2^2$$

$$
\begin{aligned}
&\overset{(a)}{\leq} \left\|A^{-1} - (A+\lambda \mathbf{I})^{-1}\right\|_2^2 \|b\|_2^2 \\
&= \left\|A^{-1}(A+\lambda \mathbf{I} - A)(A+\lambda \mathbf{I})^{-1}\right\|_2^2 \|b\|_2^2 \\
&\overset{(a)}{\leq} \left\|A^{-1}\right\|_2^2 \lambda^2 \left\|(A+\lambda \mathbf{I})^{-1}\right\|_2^2 \|b\|_2^2 \\
&\overset{(b)}{\leq} \frac{\lambda^2 \Phi_{\max}^2}{\sigma_{\min}(A)^2(\lambda+\mu)^2},
\end{aligned} \tag{72}
$$

where (a) follows from Cauchy-Schwarz inequality, and (b) follows from the fact that $\left\|A^{-1}\right\| = 1/\sigma_{\min}(A)$, with $\sigma_{\min}(A)$ as A's minimum singular value.

□

6.5. *High probability bound for Tail-averaged TD with regularisation*

Proposition 2. *Under assumptions 3, 4, for all $\epsilon \geq 0$, and $t \geq 1$,*

$$
P\left(\left\|\hat{z}_{k+1,N}\right\|_2 - \mathbb{E}\left[\left\|\hat{z}_{k+1,N}\right\|_2 > \epsilon\right]\right) \leq \exp\left(-\frac{\epsilon^2}{(R_{\max} + (1+\beta)H\Phi_{\max}^2)^2 \sum_{i=k+1}^{k+N} L_i^2}\right),
$$

where $L_i = \frac{\gamma}{N}\sum_{j=i+1}^{i+N}\left(1 - \frac{\gamma(\mu+\lambda)}{2}\right)^{j-i+1}$.

For this proof we can use Step 1 and Step 3 of Proposition 1. However, as the update rule is differs from the usual TD update, we will derive the Lipschitz constant (Step 2 in proof of Proposition 1) for the regularised TD algorithm separately. Towards that end,

We need to prove that that functions g_i are Lipschitz continuous in f_i at time i with new constant L_i.

Towards that end, first, define $\hat{\Theta}^i_{t,k+1}(\theta)$ to be the value of the regularised-tail-averaged iterate at time t that evolves according to eq. (7) beginning from θ at time i, and next, define $\bar{\hat{\Theta}}^i_{k+1,N}(\bar{\hat{\theta}}, \hat{\theta})$ as follows:

$$
\bar{\hat{\Theta}}^i_{k+1,N}(\bar{\hat{\theta}}, \hat{\theta}) = \frac{(i-k)\bar{\hat{\theta}}}{N} + \frac{1}{N}\sum_{j=i+1}^{i+N} \hat{\Theta}^i_j(\hat{\theta}),
$$

where $\hat{\theta}$ is the value of the tail-averaged regularised TD iterate at time i.

Let f and f' denote two different values of random innovations at time i. Then we know that, $\hat{\theta} = \hat{\theta}_{i-1} + \gamma f$, and $\hat{\theta}' = \hat{\theta}_{i-1} + \gamma f'$ are the parameter values corresponding to each f. Therefore,

$$
\mathbb{E}\left[\left\|\bar{\hat{\Theta}}^i_{i+1,N}(\bar{\hat{\theta}}, \hat{\theta}) - \bar{\hat{\Theta}}^i_{i+1,N}(\bar{\hat{\theta}}, \hat{\theta}')\right\|_2\right] = \mathbb{E}\left[\frac{1}{N}\sum_{j=i+1}^{i+N}\left\|\hat{\Theta}^i_j(\hat{\theta}) - \hat{\Theta}^i_j(\hat{\theta}')\right\|_2\right]. \tag{73}
$$

We will now bound the term $\hat{\Theta}^i_j(\hat{\theta}) - \hat{\Theta}^i_j(\hat{\theta}')$ inside the summation of (73).

Note that as the projection Γ is non-expansive, we have the following

$$
\mathbb{E}\left[\left\|\hat{\Theta}^i_j(\hat{\theta}) - \hat{\Theta}^i_j(\hat{\theta}')\right\|_2 \middle| \mathcal{F}_{j-1}\right] \leq \mathbb{E}\left[\left\|\left(\hat{\Theta}^i_{j-1}(\hat{\theta}) - \hat{\Theta}^i_{j-1}(\hat{\theta}')\right) - \gamma\left[f_i(\hat{\Theta}^i_{j-1}(\hat{\theta})) - f_i(\hat{\Theta}^i_{j-1}(\hat{\theta}'))\right]\right\|_2 \middle| \mathcal{F}_{j-1}\right].
$$

Expanding on f_i and using$a_j \triangleq [\phi(s_j)\phi(s_j)^\top - \beta\phi(s_j)\phi(s'_j)^\top]$, we have

$$
\begin{aligned}
&\hat{\Theta}^i_{j-1}(\hat{\theta}) - \hat{\Theta}^i_j(\hat{\theta}') - \gamma[f_i(\hat{\Theta}^i_{j-1}(\hat{\theta})) - f_i(\hat{\Theta}^i_{j-1}(\hat{\theta}')] = \\
&(1-\gamma\lambda\mathbf{I})\left(\hat{\Theta}^i_{j-1}(\hat{\theta}) - \hat{\Theta}^i_{j-1}(\hat{\theta}')\right) - \gamma[\phi(s_j)\phi(s_j)^\top - \beta\phi(s_j)\phi(s'_j)^\top][(\hat{\Theta}^j_{i-1}(\hat{\theta})) - (\hat{\Theta}^j_{i-1}(\hat{\theta}')] \\
&= [\mathbf{I} - \gamma(a_j + \lambda\mathbf{I})](\hat{\Theta}^i_{j-1}(\hat{\theta}) - \hat{\Theta}^i_{j-1}(\hat{\theta}')).
\end{aligned} \tag{74}
$$

Using the tower property of conditional expectations, it follows that:

$$
\begin{aligned}
\mathbb{E}\left[\left\|\hat{\Theta}^i_j(\hat{\theta}) - \hat{\Theta}^i_j(\hat{\theta}')\right\|_2\right] &= \mathbb{E}\left[\mathbb{E}\left[\left\|\hat{\Theta}^i_j(\hat{\theta}) - \hat{\Theta}^i_j(\hat{\theta}')\right\|_2 \middle| \mathcal{F}_{j-1}\right]\right] \\
&= \mathbb{E}\left[\mathbb{E}\left[\left\|(\mathbf{I} - (\lambda\mathbf{I} + a_j))\hat{\Theta}^i_{j-1}(\hat{\theta}) - \hat{\Theta}^i_{j-1}(\hat{\theta}')\right\|_2 \middle| \mathcal{F}_{j-1}\right]\right]
\end{aligned}
$$

$$\overset{(a)}{\leq} \left(1-\frac{\gamma(\mu+\lambda)}{2}\right)\mathbb{E}\left[\left\|\hat{\Theta}^i_{j-1}(\hat{\theta})-\hat{\Theta}^i_{j-1}(\hat{\theta}')\right\|_2\middle|\mathcal{F}_{j-1}\right]$$
$$\overset{(b)}{\leq} \left(1-\frac{\gamma(\mu+\lambda)}{2}\right)^{j-i+1}\left\|\hat{\theta}-\hat{\theta}'\right\|_2,$$

where (a) follows from Lemma 13, and (b) follows from repeated application of argument that helps us arrive at (a).
Now using Jensen's inequality we get,

$$\left|\mathbb{E}\left[\left\|\hat{\theta}_j-\theta^\star_{\text{reg}}\right\|_2|\hat{\theta}_j=\hat{\theta}\right]-\mathbb{E}\left[\left\|\hat{\theta}_j-\theta^\star_{\text{reg}}\right\|_2|\hat{\theta}_j=\hat{\theta}'\right]\right|\leq\mathbb{E}\left[\left\|\hat{\Theta}^i_j(\hat{\theta})-\hat{\Theta}^i_j(\hat{\theta}')\right\|_2\right]\leq\gamma\left(1-\frac{\gamma(\mu+\lambda)}{2}\right)^{j-i+1}\|f-f'\|_2. \tag{75}$$

Substituting (75) in (73), and using Lemma 10, we get the following

$$\mathbb{E}\left[\left\|\bar{\hat{\Theta}}^i_{k+1}(\bar{\hat{\theta}}_{i-1},\hat{\theta})-\bar{\hat{\Theta}}^i_{k+1}(\bar{\hat{\theta}}_{i-1},\hat{\theta}')\right\|_2\right]\leq\frac{\gamma}{N}\sum_{j=i+1}^{i+N}\left(1-\frac{\gamma(\mu+\lambda)}{2}\right)^{j-i+1}\|f-f'\|_2. \tag{76}$$

From (76) it is clear that g_i is L_i-Lipschtiz in f_i at time i, which implies that D_i is Lipschitz with Lipschitz constant $L_i=\frac{\gamma}{N}\sum_{j=i+1}^{i+N}\left(1-\frac{\gamma(\mu+\lambda)}{2}\right)^{j-i+1}$.

6.5.1. Bounding the Lipschitz constant

Lemma 18. *For the tail-averaged TD, the Lipschitz constant L_i in Proposition 1 is upper bounded as follows:*

$$\sum_{i=k+1}^{k+N}L_i^2\leq\frac{4}{(\mu+\lambda)^2N}.$$

Proof.

$$\begin{aligned}\sum_{i=k+1}^{k+N}L_i^2&=\frac{\gamma^2}{N^2}\sum_{i=k+1}^{k+N}\left(\sum_{j=i+1}^{i+N}\left(1-\frac{\gamma(\mu+\lambda)}{2}\right)^{j-i+1}\right)^2\\&\leq\frac{\gamma^2}{N^2}\sum_{i=k+1}^{k+N}\left(\sum_{j=i+1}^{\infty}\left(1-\frac{\gamma(\mu+\lambda)}{2}\right)^{j-i+1}\right)^2\\&\overset{(a)}{=}\frac{\gamma^2}{N^2}\sum_{i=k+1}^{k+N}\frac{4}{\gamma^2(\mu+\lambda)^2}\\&\leq\frac{\gamma^2}{N^2}\sum_{i=k+1}^{\infty}\frac{4}{\gamma^2(\mu+\lambda)^2}\\&=\frac{4}{(\mu+\lambda)^2N},\end{aligned}$$

where (a) is due to summing the geometric series. □

6.5.2. Proof of Theorem 4

For the sake of convenience, we restate the high probability bound for regularised TD below.
Theorem 8 (High-probability bound). *Assume 1 to 4, and 6. Choose the step size such that $\gamma\leq\gamma_{\max}$, where $\gamma_{\max}$ is defined in (19). Then, for any $\delta\in(0,1]$, we have the following bound for the projected tail-averaged regularised TD iterate $\hat{\theta}_{k+1,N}$:*

$$P\left(\left\|\hat{\theta}_{k+1,N}-\theta^\star_{\text{reg}}\right\|_2\leq\frac{2\sigma}{(\mu+\lambda)\sqrt{N}}\sqrt{\log\left(\frac{1}{\delta}\right)}+\frac{4e^{(-k\gamma(\mu+\lambda))}}{\gamma(\mu+\lambda)N}\mathbb{E}\left[\left\|\hat{\theta}_0-\theta^\star_{\text{reg}}\right\|_2\right]+\frac{4\sigma}{(\mu+\lambda)\sqrt{N}}\right)\geq1-\delta,$$

where $N,\sigma,\mu,\hat{\theta}_0,\theta^\star_{\text{reg}}$ are as specified in Theorem 3.

Proof.

$$P(\|\hat{z}_{k+1,N}\|_2^2 - \mathbb{E}[\|\hat{z}_{k+1,N}\|_2^2] > \epsilon) \overset{(a)}{\leq} \exp\left(-\frac{\epsilon^2}{(R_{\max}+(1+\beta)H\Phi_{\max}^2)^2 \sum_{i=k+1}^{k+N} L_i^2}\right),$$
$$\overset{(b)}{\leq} \exp\left(-\frac{N(\mu+\lambda)^2\epsilon^2}{4\sigma^2}\right), \tag{77}$$

where (a) follows from Proposition 2, and (b) follows from Lemma 18.
The main claim follows by converting the bound above to a high-confidence form. In particular,

$$\exp\left(-\frac{N(\mu+\lambda)^2\epsilon^2}{4\sigma^2}\right) = \delta, \text{ leads to } \quad \epsilon = \frac{2\sigma}{(\mu+\lambda)}\sqrt{\frac{\log\left(\frac{1}{\delta}\right)}{N}}. \tag{78}$$

The final bound follows by substituting the value for ϵ from (78) in (77), and using the result from Theorem 3. □

7. Bounds for Under Mixing Assumptions

Instead of Assumption 2, we now consider the case when $(s_t)_{t\in\mathbb{N}}$ are drawn from a single stationary trajectory of the Markov chain with policy π. We assume exponential ergodicity for the Markov chain, which holds when any finite Markov chain is irreducible. Let ρ be the corresponding stationary distribution.
Assumption 7. $s_1 \sim \rho$ and there exist constants C and τ_{mix} such that for every $t, \tau \in \mathbb{N}$

$$D(\tau) := \sup_{s\in\mathcal{S}} \text{TV}(s_{t+\tau}|s_t = s, \rho) \leq C\exp(-\tfrac{\tau}{\tau_{\text{mix}}}),$$

where TV denotes the total variation distance between probability measures.
This is a standard assumption in the literature [2, 20]. We now adapt Lemma 3 from [13] to our present setting.
Lemma 19 (Adaptation of Lemma 3 in [13])**.** *For any $K \in \mathbb{N}$, define the random variable*

$$S_{K,n} := ((s_1, s_2), (s_{K+1}, s_{K+2}), (s_{2K+1}, s_{2K+2}), \ldots, (s_{nK+1}, s_{nK+2})).$$

Let P denote the transition kernel for the Markov chain under policy π. By $\rho^{(2)}$ denote the joint distribution of (s_1, s_2). UnderAssumption 7, we have

$$\text{TV}(S_{K,n}, (\rho^{(2)})^{\otimes n}) \leq nD(K-1) \leq nC\exp(-\tfrac{K-1}{\tau_{\text{mix}}}).$$

Now, let $R_{K,n} = (r_1, r_{K+1}, \ldots, r_{nK+1})$ be the random rewards corresponding to $S_{K,n}$ and consider i.i.d random variables $\tilde{S}_{K,n} = ((\tilde{s}_1, \tilde{s}_2), (\tilde{s}_{K+1}, \tilde{s}_{K+2}), (\tilde{s}_{2K+1}, \tilde{s}_{2K+2}), \ldots, (\tilde{s}_{nK+1}, \tilde{s}_{nK+2})) \sim (\rho^{(2)})^{\otimes n}$ along with the corresponding rewards $\tilde{R}_{K,n}$. We can define these random variables on a common probability space such that

$$\mathbb{P}((S_{K,n}, R_{K,n}) \neq (\tilde{S}_{K,n}, \tilde{R}_{K,n})) \leq nD(K-1) \leq nC\exp(-\tfrac{K-1}{\tau_{\text{mix}}}).$$

We will use the mixing technique used in SGD-DD in [13]. Here, when we obtain samples (s_t, r_t, s_{t+1}) from a trajectory instead of from the i.i.d distribution as considered before. we modify the Algorithm 1 in the following ways to account for mixing. We fix $K \in \mathbb{N}$.

Modification 1: Run Algorithm 1 with data $S_{K,n}, R_{K,n}$ - i.e, we input $(s_{tK+1}, r_{tK+1}, s_{tK+2})$ at step t.

Modification 2: Run Algorithm 1 with data $\tilde{S}_{K,n}, \tilde{R}_{K,n}$.
Note that the Modification 2 is exactly same as running the algorithm under Assumption 2 for n steps and therefore the results of Theorem 2 apply to this case if we replace N with n. By the results in Lemma 19, we conclude that the trajectories (θ_t) generated by modification 1 and $(\tilde{\theta}_t)$ generated by modification 2 can be coupled such that

$$\mathbb{P}\left[(\theta_t)_{t=1}^{n+1} \neq (\tilde{\theta}_t)_{t=1}^{n+1}\right] \leq nD(K-1).$$

This is based on the fact that whenever the algorithm is fed with the same input, we obtain the same output. Setting $K = \tau_{\text{mix}}\log(\frac{Cn}{\delta})$, we conclude that under Assumption 7, we have:

$$\mathbb{P}\left[(\theta_t)_{t=1}^{n+1} \neq (\tilde{\theta}_t)_{t=1}^{n+1}\right] \leq \delta.$$

Therefore, we conclude the bounds in Remark 8.

8. Conclusions

We presented a finite time analysis of tail-averaged TD algorithm with/without regularisation. Our bounds are easy to interpret, and improve the previously known results. To the best of our knowledge, this is the first result that establishes a $O\left(\frac{1}{t}\right)$ convergence rate for a TD algorithm with a universal step size. Finally, we also analyse TD with regularisation and show how it can be useful in certain problem instances with ill-conditioned features.

References

[1] N. Agarwal, S. Chaudhuri, P. Jain, D. M. Nagaraj, and P. Netrapalli. Online Target Q-learning with Reverse Experience Replay: Efficiently finding the Optimal Policy for Linear MDPs. In *The Tenth International Conference on Learning Representations, ICLR*, 2022.

[2] J. Bhandari, D. Russo, and R. Singal. A finite time analysis of temporal difference learning with linear function approximation. In *Conference On Learning Theory (COLT)*, volume 75 of *Proceedings of Machine Learning Research*, pages 1691–1692. PMLR, 2018.

[3] S. Chen, A. Devraj, A. Busic, and S. Meyn. Explicit Mean-Square Error Bounds for Monte-Carlo and Linear Stochastic Approximation. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 108 of *Proceedings of Machine Learning Research*, pages 4173–4183. PMLR, 2020.

[4] Z. Chen, J. P. Clarke, and S. T. Maguluri. Target Network and Truncation Overcome The Deadly triad in Q-Learning. *arXiv preprint arXiv:2203.02628*, 2022.

[5] Z. Chen, S. Zhang, T. T. Doan, J. P. Clarke, and S. T. Maguluri. Finite-sample analysis of nonlinear stochastic approximation with applications in reinforcement learning. In *Automatica*, volume 146, 2022.

[6] G. Dalal, B. Szörényi, G. Thoppe, and S. Mannor. Finite sample analyses for TD(0) with function approximation. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 6144–6160, 2018.

[7] A. Durmus, É. Moulines, A. Naumov, S. Samsonov, and H.-T. Wai. On the stability of random matrix product with markovian noise: Application to linear stochastic approximation and td learning. In *COLT*, 2021.

[8] M. Fathi and N. Frikha. Transport-entropy inequalities and deviation estimates for stochastic approximation schemes. *Electronic Journal of Probability*, 18:1–36, 2013.

[9] B. Hu and U. A. Syed. Characterizing the exact behaviors of temporal difference learning algorithms using markov jump linear system theory. In *NeurIPS*, 2019.

[10] T. S. Jaakkola, M. I. Jordan, and S. P. Singh. On the convergence of stochastic iterative dynamic programming algorithms. *Neural Comput.*, 6(6):1185–1201, 1994.

[11] P. Jain, S. Kakade, R. Kidambi, P. Netrapalli, and A. Sidford. Parallelizing stochastic gradient descent for least squares regression: mini-batching, averaging, and model misspecification. *Journal of Machine Learning Research*, 18, 2018.

[12] C. Lakshminarayanan and C. Szepesvari. Linear stochastic approximation: How far does constant step-size and iterate averaging go? In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 84 of *Proceedings of Machine Learning Research*, pages 1347–1355. PMLR, Apr 2018.

[13] D. Nagaraj, X. Wu, G. Bresler, P. Jain, and P. Netrapalli. Least squares regression with markovian data: Fundamental limits and algorithms. *Advances in neural information processing systems*, 33:16666–16676, 2020.

[14] G. Patil, P. L.A., D. Nagaraj, and D. Precup. Finite time analysis of temporal difference learning with linear function approximation: Tail averaging and regularisation. In F. Ruiz, J. Dy, and J.-W. van de Meent, editors, *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pages 5438–5448. PMLR, 25–27 Apr 2023. URL https://proceedings.mlr.press/v206/patil23a.html.

[15] F. J. Pineda. Mean-field theory for batched-td(l). *Neural Comput.*, 9(7):1403–1419, 1997.

[16] B. T. Polyak and A. B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992.

[17] L. A. Prashanth, N. Korda, and R. Munos. Concentration bounds for temporal difference learning with linear function approximation: the case of batch data and uniform sampling. *Mach. Learn.*, 110(3):559–618, 2021.

[18] D. Ruppert. Stochastic approximation. *Handbook of Sequential Analysis*, pages 503–529, 1991.

[19] R. Schapire and M. K. Warmuth. On the worst-case analysis of temporal-difference learning algorithms. *Machine Learning*, 22:95–121, 2004.

[20] R. Srikant and L. Ying. Finite-Time Error Bounds For Linear Stochastic Approximation and TD Learning. In *Conference on Learning Theory (COLT)*, volume 99 of *Proceedings of Machine Learning Research*, pages 2803–2830. PMLR, 2019.

[21] R. S. Sutton. Learning to predict by the methods of temporal differences. *Mach. Learn.*, 3:9–44, 1988.

[22] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2018. ISBN 0-262-19398-1.

[23] J. N. Tsitsiklis and B. Van Roy. An analysis of temporal-difference learning with function approximation. *IEEE Transactions on Automatic Control*, 42(5):674–690, 1997.

[24] G. Wang and G. B. Giannakis. Finite-Time Error Bounds for Biased Stochastic Approximation with Applications to Q-Learning. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 108 of *Proceedings of Machine Learning Research*, pages 3015–3024. PMLR, 2020.

[25] S. Zhang, Z. Zhang, and S. T. Maguluri. Finite sample analysis of average-reward td learning and q-learning. In *NeurIPS*, 2021.