
TEMPORAL DIFFERENCE LEARNING AS GRADIENT SPLITTING

Rui Liu

Division of Systems Engineering
Boston University
Boston, MA, 02215
rliu@bu.edu

Alex Olshevsky

Department of ECE and Division of Systems Engineering
Boston University
Boston, MA, 02215
alexols@bu.edu

October 29, 2020

ABSTRACT

Temporal difference learning with linear function approximation is a popular method to obtain a low-dimensional approximation of the value function of a policy in a Markov Decision Process. We give a new interpretation of this method in terms of a splitting of the gradient of an appropriately chosen function. As a consequence of this interpretation, convergence proofs for gradient descent can be applied almost verbatim to temporal difference learning. Beyond giving a new, fuller explanation of why temporal difference works, our interpretation also yields improved convergence times. We consider the setting with $1/\sqrt{T}$ step-size, where previous comparable finite-time convergence time bounds for temporal difference learning had the multiplicative factor $1/(1-\gamma)$ in front of the bound, with γ being the discount factor. We show that a minor variation on TD learning which estimates the mean of the value function separately has a convergence time where $1/(1-\gamma)$ only multiplies an asymptotically negligible term.

1 Introduction

Reinforcement learning is a basic machine learning paradigm which concerns learning optimal policies in Markov Decision Processes (MDP). It has been applied to many challenging practical problems, such as, autonomous driving (Chen et al., 2015), robotics (Gu et al., 2017), bidding and advertising (Jin et al., 2018), and games (Silver et al., 2016). An important problem in reinforcement learning is to estimate the value function for a given policy, often referred to as the policy evaluation problem. Temporal difference (TD) learning originally proposed by Sutton (1988) is one of the most widely used policy evaluation algorithms. TD uses differences in predictions over successive time steps to drive the learning process, with the prediction at any given time step updated via a carefully chosen step-size to bring it closer to the prediction of the same quantity at the next time step.

Despite its simple implementation, theoretical analysis of TD can be involved. This is particularly true when TD methods are applied to problems with large state-spaces by maintaining an approximation to the value function. Precise conditions for the asymptotic convergence of TD with linear function approximation were established by viewing TD as a stochastic approximation for solving a suitable Bellman equation in (Tsitsiklis and Van Roy, 1997). Before the last few years, there have been few non-asymptotic analyses of TD methods. The first non-asymptotic bounds for TD(0) with linear function approximation were given by Korda and La (2015), obtaining an exponential convergence rate for the centered variant of TD(0) when the underlying Markov chain mixes fast. However, some issues with the proofs of Korda and La (2015) were listed by the subsequent work of Narayanan and Szepesvári (2017).

In Lakshminarayanan and Szepesvári (2018) it was shown that TD algorithms with a problem independent constant step size and iterate averaging, achieve a problem dependent error that decays as $O(1/t)$ with the number of iterations t . Convergence rates in probability with an $O(1/t)$ step-size were provided by Dalal et al. (2018). Both analyses of Dalal et al. (2018) and Lakshminarayanan and Szepesvári (2018) assume samples used by the algorithm are i.i.d. rather than a trajectory in Markov chain. For the Markov chain observation model, Bhandari et al. (2018) provide a $O(1/\sqrt{T})$ convergence rate with step-size that scales as $1/\sqrt{T}$ and $O((\log t)/t)$ convergence rate with step size $O(1/t)$

for projected TD algorithm. The constant factors in the latter bounds depend on $1/(1-\gamma)$, where γ is the discount factor; this scaling is one of the things we will be studying in this paper.

A number of papers also work on algorithms related to and inspired by the classic TD algorithm in the setting with Markovian sampling. [Srikant and Ying \(2019\)](#) give finite-time bounds for the TD algorithms with linear function approximation and a constant step-size. The two time-scale TD with gradient correction algorithm under a Markovian sampling and linear function approximation are discussed by [Xu et al. \(2019\)](#) and shown to converge as fast as $O((\log t)/t^{2/3})$. A method called TD-AMSGrad under linear function approximation is studied by [Xiong et al. \(2020\)](#); with a constant step size, TD-AMSGrad converges to a neighborhood of the global optimum at a rate of $O(1/t)$, and with a diminishing step size, it converges exactly to the global optimum at a rate of $O((\log t)/t)$.

In this paper, we will study the convergence of with linear function approximation under Markov observations. Our main contribution is to provide a new interpretation of temporal difference learning: we show how to view it as a “splitting” (a term we introduce and define later) of an appropriately chosen quadratic form. As a consequence of this interpretation, it is possible to apply convergence proofs for gradient descent almost verbatim to temporal difference learning.

The convergence times bounds we obtain this way improve on existing results. In particular, we study step-sizes of $1/\sqrt{T}$, which are typically recommended because the resulting error bounds do not depend on the inverse eigenvalues of the matrices involved in the linear approximation, which can be quite large; by contrast, methods that achieve faster than $O(1/\sqrt{T})$ decay have performance guarantees that scale with these same eigenvalues. We provide a minor variation on TD(0) for which we obtain a convergence rate that scales as $[(1/(1-\gamma))^2]/T + O(1/\sqrt{T})$, with the constant in the $O(\cdot)$ notation not blowing up as $\gamma \rightarrow 1$. We will also explain why a factor of $1/(1-\gamma)^2$ multiplying the asymptotically negligible $O(1/T)$ term as here is unavoidable.

2 PRELIMINARIES

In this section, we describe the basics of MDPs and TD learning methods. While all this material is standard and available in textbooks (e.g., [Sutton and Barto \(2018\)](#)), it is necessary to standardize notation and make our presentation self-contained.

2.1 Markov Decision Processes

We consider a discounted reward MDP described by a 5-tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma)$, where $\mathcal{S} = [n] = \{1, 2, \dots, n\}$ is the finite state space, \mathcal{A} is the finite action space, $\mathcal{P} = (\mathcal{P}(s'|s, a))_{s, s' \in \mathcal{S}, a \in \mathcal{A}}$ are the transition probabilities, $r = (r(s, a, s'))_{s, s' \in \mathcal{S}, a \in \mathcal{A}}$ are deterministic rewards and $\gamma \in (0, 1)$ is the discount factor. The (stationary) policy to be evaluated is a mapping $\mu: \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$, where $\mu(s, a)$ are the probabilities to select action a when in state s and $\sum_{a \in \mathcal{A}} \mu(s, a) = 1$ for all states $s \in \mathcal{S}$. We adopt the shorthand s_t for the state at step t , a_t for the action taken at step t , and $r_{t+1} = r(s_t, a_t, s_{t+1})$.

The value function of the policy μ , denoted $V^\mu: \mathcal{S} \rightarrow \mathbb{R}$ is defined as

$$V^\mu(s) = E_{\mu, s} \left[\sum_{t=0}^{\infty} \gamma^t r_{t+1} \right],$$

where $E_{\mu, s}[\cdot]$ indicates that s is the initial state and the actions are chosen according to μ .

The immediate reward vector $R^\mu: \mathcal{S} \rightarrow \mathbb{R}$ is defined as

$$R^\mu(s) = E_{\mu, s}(r_1) = \sum_{s' \in \mathcal{S}} \sum_{a \in \mathcal{A}} \mu(s, a) \mathcal{P}(s'|s, a) r(s, a, s').$$

For the remainder of the paper, we will be fixing the policy μ ; consequently, we can talk about the probability transition matrix P^μ defined as

$$P^\mu(s, s') = \sum_{a \in \mathcal{A}} \mu(s, a) \mathcal{P}(s'|s, a).$$

In the following, we will treat V^μ and R^μ as vectors in \mathbb{R}^n , and treat P^μ as a matrix in $\mathbb{R}^{n \times n}$. It is well-known that V^μ satisfies the so-called Bellman equation ([Sutton and Barto, 2018](#)): defining the Bellman operator $T^\mu: \mathbb{R}^n \rightarrow \mathbb{R}^n$ as

$$(T^\mu V^\mu)(s) = \sum_{s'=1}^n P^\mu(s, s') (r(s, s') + \gamma V^\mu(s')) \quad (1)$$

for $s \in [n]$, where $r(s, s') = \sum_{a \in \mathcal{A}} \mu(s, a) r(s, a, s')$, we can then write Bellman equation as $T^\mu V^\mu = V^\mu$ (Sutton and Barto, 2018).

Next, we state some standard assumptions from the literature. The first assumption is on the underlying Markov chain.

Assumption 1. *The Markov chain whose transition matrix is the matrix P^μ is irreducible and aperiodic.*

Following this assumption, the Markov decision process induced by the policy μ is ergodic with a unique stationary distribution $\pi = (\pi_1, \pi_2, \dots, \pi_n)$, a row vector whose entries are non-negative and sum to 1. It also holds that $\pi_{s'} = \lim_{t \rightarrow \infty} (P^\mu)^t(s, s')$ for any two states $s, s' \in [n]$. Note that we are using π to denote the stationary distribution of P^μ , and not the policy (which is denoted by μ).

We will use the notation r_{\max} to denote an upper bound on the rewards; more formally, r_{\max} is a real number such that

$$|r(s, a, s')| \leq r_{\max} \text{ for all } s, s' \in [n], a \in \mathcal{A}.$$

Since the number of actions and states is finite, such an r_{\max} always exists.

We next introduce some notation that will make our analysis more concise. For a symmetric positive definite matrix $A \in \mathbb{R}^{n \times n}$, we define the inner product $\langle x, y \rangle_A = x^T A y$ and the associated norm $\|x\|_A = \sqrt{x^T A x}$. Let $D = \text{diag}(\pi_1, \dots, \pi_n)$ denote the diagonal matrix whose elements are given by the entries of the stationary distribution π . Given value functions V and V' on the state space \mathcal{S} , we have

$$\langle V, V' \rangle_D = V^T D V' = \sum_{s \in \mathcal{S}} \pi_s V(s) V'(s),$$

and the associated norm

$$\|V\|_D^2 = V^T D V = \sum_{s \in \mathcal{S}} \pi_s V(s)^2.$$

Finally, we define the Dirichlet seminorm, which is often called the Dirichlet form in the Markov chain literature (Diaconis et al., 1996); we follow here the notation of Ollivier (2018). The Dirichlet seminorm depends both on the transition matrix P^μ and the invariant measure π :

$$\|V\|_{\text{Dir}}^2 = \frac{1}{2} \sum_{s, s' \in \mathcal{S}} \pi_s P^\mu(s, s') (V(s') - V(s))^2. \quad (2)$$

It is easy to see that, as a consequence of Assumption 1, $\|V\|_{\text{Dir}} = 0$ if and only if V is a multiple of the all-ones vector.

Similarly, we introduce the k -step Dirichlet seminorm, defined as

$$\|V\|_{\text{Dir}, k}^2 = \frac{1}{2} \sum_{s, s' \in \mathcal{S}} \pi_s (P^\mu)^k(s, s') (V(s') - V(s))^2.$$

2.2 Policy Evaluation, Temporal Difference Learning, and Linear Function Approximation

Policy evaluation refers to the problem of estimating the value function V^μ for a given stationary policy μ . If the size of the state space is large, computing $V^\mu(s)$ for all states s may be prohibitively expensive. A standard remedy is to use low dimensional approximation V_θ^μ of V^μ in the classical TD algorithm as in Sutton (1988); Sutton and Barto (2018). For brevity, we omit the superscript μ throughout from now on.

The classical TD(0) algorithm with function approximation V_θ starts with an arbitrary value of the parameters θ_0 ; upon observing the t^{th} transition $s_t \rightarrow s'_t$, it computes the scalar-valued temporal-difference error,

$$\delta_t = r(s_t, a_t, s'_t) + \gamma V_{\theta_t}(s'_t) - V_{\theta_t}(s_t),$$

and updates the parameter vector as

$$\theta_{t+1} = \theta_t + \alpha_t \delta_t \nabla V_{\theta_t}(s_t). \quad (3)$$

Here $\nabla V_{\theta_t}(s_t)$ denotes the gradient of the function $V_{\theta_t}(s_t)$ w.r.t to θ evaluated at $\theta = \theta_t$, and α_t is the step size. Intuitively, updating in the direction $\delta_t \nabla V_{\theta_t}(s_t)$ moves $V_{\theta_t}(s_t)$ closer to the bootstrapped value of $r(s_t, a_t, s'_t) + \gamma V_{\theta_t}(s'_t)$.

We will be considering the TD(0) algorithm with a linear function approximation V_θ defined as

$$V_\theta(s) = \sum_{l=1}^K \theta_l \phi_l(s) \quad \forall s \in \mathcal{S}, \quad (4)$$

for a given set of K feature vectors $\phi_l : \mathcal{S} \rightarrow \mathbb{R}$, $l \in [K]$. For each state s , we will define the vector $\phi(s)$ which stacks up the features of s as $\phi(s) = (\phi_1(s), \phi_1(s), \dots, \phi_K(s))^T \in \mathbb{R}^K$. Finally, $\Phi \in \mathbb{R}^{n \times K}$ is defined to be the matrix $\Phi = [\phi_1, \dots, \phi_K]$.

We thus have that $V_\theta(s) = \theta^T \phi(s)$ and the approximate TD(0) update becomes

$$\theta_{t+1} = \theta_t + \alpha_t (r_t + \gamma \theta_t^T \phi(s'_t) - \theta_t^T \phi(s_t)) \phi(s_t). \quad (5)$$

Next we state a common assumption on the feature vectors, which requires that features used for approximation are linearly independent (Tsitsiklis and Van Roy, 1997; Bhandari et al., 2018).

Assumption 2. *The matrix Φ has full column rank, i.e., the feature vectors $\{\phi_1, \dots, \phi_K\}$ are linearly independent. Additionally, we also assume that $\|\phi(s)\|_2^2 \leq 1$ for $s \in \mathcal{S}$.*

It is always possible to make sure this assumption holds. If the norm bound is unsatisfied, then the standard approach is to normalize the feature vectors so that it does. If the matrix Φ does not have full column rank, one can simply omit enough feature vectors so that it does; since the range of Φ is unaffected, this does not harm the quality of approximation.

It is well-known that under Assumptions 1-2 as well as an additional assumption on the decay of the step-sizes α_t , temporal difference learning converges with probability one; furthermore, its limit is the fixed point of a certain projected Bellman equation (Tsitsiklis and Van Roy, 1997). Henceforth we will use θ^* to denote this fixed point.

It is convenient to introduce the notation

$$g_t(\theta) = (r_t + \gamma \phi(s'_t)^T \theta - \phi(s_t)^T \theta) \phi(s_t). \quad (6)$$

for the direction taken by TD(0) at time t . Note that $g_t(\theta)$ is a scalar multiple of $\phi(s_t)$, the feature vector of the state encountered at time t .

Furthermore, $\bar{g}(\theta)$ will denote the average of $g_t(\theta)$ when the state is sampled according to the stationary distribution:

$$\bar{g}(\theta) = \sum_{s, s' \in \mathcal{S}} \pi(s) P(s, s') (r(s, s') + \gamma \phi(s')^T \theta - \phi(s)^T \theta) \phi(s).$$

Naturally it can be seen (see Tsitsiklis and Van Roy (1997)) that

$$\bar{g}(\theta^*) = 0. \quad (7)$$

2.3 Eligibility traces

We will also study a larger class of algorithms, denoted by TD(λ) and parameterized by $\lambda \in [0, 1]$, that contains as a special case the TD(0) algorithm discussed above. While TD(0) makes parameter updates in the direction of the (scaled) last feature vector $g_t(\theta_t)$, the TD(λ) algorithm maintains the “eligibility trace”

$$z_t = \sum_{k=-\infty}^t (\gamma \lambda)^k \phi(s_{t-k}),$$

which is a geometric weighted average of the feature vectors at all previously visited states, and takes a step in the direction of z_t .

In practice, the sum will start at $k = 0$ (or some other finite time); however, parts of the analysis are done with the sum starting at negative infinity because many of the results are much simpler in this setting, and doing so introduces only an exponentially decaying error term.

It is shown in Tsitsiklis and Van Roy (1997) that, subject to Assumptions 1-2 and appropriate decay of step-sizes, TD(λ) converges with probability one, and its limit is a fixed point of a certain projected & averaged Bellman equation. We will denote this limit by θ_λ^* .

2.4 Markov Chain Observation Model

In this paper, we are interested in TD in the setting where the data is collected from a single sample path of a Markov chain. Our final assumption is that the Markov chain mixes at a uniform geometric rate..

Assumption 3. *There are constants $m > 0$ and $\rho \in (0, 1)$ such that*

$$\sup_{s \in \mathcal{S}} d_{\text{TV}}(P^t(s, \cdot), \pi) \leq m\rho^t \quad t \in \mathbb{N}_0,$$

where $d_{\text{TV}}(P, Q)$ denotes the total-variation distance between probability measures P and Q . In addition, the initial distribution of s_0 is the steady-state distribution π , so (s_0, s_1, \dots) is a stationary sequence.

Under Assumption 1, i.e., for irreducible and aperiodic Markov chains, the uniform mixing assumption always holds (Levin and Peres, 2017). The assumption that the chain begins in steady-state allows us to simplify many mathematical expressions. It is worth noting that the assumption that s_0 is the stationary distribution is primarily done to make the analysis and results tidier: given the uniform mixing assumption, one can apply analysis after the Markov chain is close to its steady-state.

3 Our main result: a new interpretation of temporal difference learning

All existing analysis temporal difference learning proceed by comparing it, either explicitly or implicitly, to the expected update, usually referred to as the mean-path update; for TD(0), this is

$$\theta_{t+1} = \theta_t + \alpha_t \bar{g}(\theta_t). \quad (8)$$

Stochastic approximation Robbins and Monro (1951) is a common tool to make this comparison. Generally, one wants to argue that the mean-path TD update brings θ_t closer to its final value θ^* .

The first theoretical analysis of TD(0) in Tsitsiklis and Van Roy (1997) proceeded based on the observation that $\bar{g}(\theta)$ forms a positive angle with $\theta^* - \theta$, that is

$$\bar{g}(\theta)^T (\theta^* - \theta) > 0. \quad (9)$$

An explicit version of this inequality was used in Bhandari et al. (2018) where it is stated as Lemma 3:

$$\bar{g}(\theta)^T (\theta^* - \theta) \geq (1 - \gamma) \|V_{\theta^*} - V_{\theta}\|_D^2. \quad (10)$$

Our main result is a new interpretation of the quantity $\bar{g}(\theta)$ which explains why such an inequality holds, as well as allows us to derive stronger results. To do this, we first introduce the concept of a “gradient splitting.”

Definition 1. *Let A be a symmetric positive semi-definite matrix. A linear function $h(\theta) = B(\theta - a)$ is called a gradient splitting of the quadratic $f(\theta) = (\theta - a)^T A(\theta - a)$ if*

$$B + B^T = 2A.$$

To the best of our knowledge, the concept is introduced here for the first time. We next explain why it is useful.

3.1 Gradient splitting and gradient descent

Observe first that, as one should expect from the name, $(1/2)\nabla f(\theta)$ is a splitting of the gradient of f since

$$\frac{1}{2}\nabla f(\theta) = A(\theta - a).$$

Of course, it is far from the only splitting, since there are many B that satisfy $B + B^T = 2A$. In particular, B may be non-symmetric. For example, one can take B to be equal to the upper triangular part of $2A$ plus the diagonal of A .

The key property of splittings that make them useful is the following.

Proposition 1. *Suppose $h(\theta)$ is a splitting of $f(\theta)$. Then*

$$(\theta_1 - \theta_2)^T (h(\theta_1) - h(\theta_2)) = \frac{1}{2}(\theta_1 - \theta_2)^T (\nabla f(\theta_1) - \nabla f(\theta_2)).$$

Proof. Indeed,

$$\begin{aligned} (\theta_1 - \theta_2)^T (h(\theta_1) - h(\theta_2)) &= (\theta_1 - \theta_2)^T B(\theta_1 - \theta_2) \\ &= \frac{1}{2}(\theta_1 - \theta_2)^T B(\theta_1 - \theta_2) + \frac{1}{2}(\theta_1 - \theta_2)^T B^T(\theta_1 - \theta_2) \\ &= (\theta_1 - \theta_2)^T A(\theta_1 - \theta_2) \\ &= \frac{1}{2}(\theta_1 - \theta_2)^T (\nabla f(\theta_1) - \nabla f(\theta_2)). \end{aligned}$$

■

Thus, while $h(\theta)$ may be quite different from $\nabla f(\theta)$, the difference disappears once one looks at the inner products considered in Proposition 1.

A particular consequence of Proposition 1 can be obtained by plugging in $\theta_1 = a$, the global minimizer of $f(\theta)$. In that case, $\nabla f(a) = 0$ and $h(a) = 0$ as well, and we obtain that for all θ ,

$$(a - \theta)^T h(\theta) = (a - \theta)^T \nabla f(\theta). \quad (11)$$

Thus the splitting $h(\theta)$ has the exact same angle with the “direction to the optimal solution” $a - \theta$ as the true gradient. This is the punchline of this discussion.

Most analysis of gradient descent on convex functions are ultimately based on the observation that gradient descent “makes progress” towards the optimal solution because it has a positive inner product with the direction to optimality. As a consequence of this discussion, the same argument can be applied to gradient splittings.

3.2 Our contribution

We now come back to temporal difference learning. To analyze TD learning, it is tempting to see if we can write the TD(0) and TD(λ) updates as gradient descent on some appropriately chosen function. Unfortunately, it is well-known (and easy to see) that this cannot work. Indeed, in the TD(0) case, it is possible to express the average direction $\bar{g}(\theta)$ as $\bar{g}(\theta) = B(\theta - \theta^*)$ and in some cases the matrix B is not symmetric; this linear map cannot be the gradient of anything since the non-symmetry of B would contradict equality of partial derivatives (see Maei (2011)).

Our main results show that the temporal difference direction can, however, be viewed as a splitting of the gradient of an appropriately chosen function.

Theorem 1. *Suppose Assumptions 1-2 hold. Then in the TD(0) update, $-\bar{g}(\theta)$ is a splitting of the gradient of the quadratic*

$$f(\theta) = (1 - \gamma)\|V_\theta - V_{\theta^*}\|_D^2 + \gamma\|V_\theta - V_{\theta^*}\|_{\text{Dir}}^2.$$

Theorem 2. *Suppose Assumptions 1-2 hold. Then, in the TD(λ) update, the negative of the expected update $-E[z_t]$ is a splitting of the gradient of the quadratic*

$$f^{(\lambda)}(\theta) = (1 - \gamma\kappa)\|V_\theta - V_{\theta^*}\|_D^2 + (1 - \lambda)\gamma\|V_\theta - V_{\theta^*}\|_{\text{Dir}}^2 + (1 - \lambda)\lambda\gamma^2\|V_\theta - V_{\theta^*}\|_{\text{Dir},2}^2 + (1 - \lambda)\lambda^2\gamma^3\|V_\theta - V_{\theta^*}\|_{\text{Dir},3}^2 + \dots,$$

where $\kappa = (1 - \lambda)/(1 - \gamma\lambda)$.

The proof of these theorems can be found in the supplementary information. Assumptions 1 and 2 are not particularly crucial: they are used only to be able to define the stationary distribution π and the unique fixed point θ^* .

These results provide some new insights into why temporal difference learning works. Indeed, there is no immediate reason why the bootstrapped update of (5) should produce a reasonable answer, and it is well known that the version with nonlinear approximation in (3) can diverge (see Tsitsiklis and Van Roy (1997)). Convergence analyses of TD learning rely on (9), but the proof of this equation from Tsitsiklis and Van Roy (1997) does not yield a conceptual reason for why it should hold.

The previous two theorems provide such a conceptual reason. It turns out that TD(0) and TD(λ) are, on average, attempting to minimize the functions $f(\theta)$ and $f^{(\lambda)}(\theta)$ respectively by moving in direction of a gradient splitting. Moreover, the functions $f(\theta)$ and $f^{(\lambda)}(\theta)$ are plainly convex (they are positive linear combinations of convex quadratics), so that Equation (11) immediately explains why (9) holds. As far as we are aware, these results are new even in the case when no function approximation is used, which can be recovered by simply choosing the vectors ϕ_i to be the unit basis vectors.

These theorems are inspired by the recent preprint (Ollivier, 2018). It is shown there that, if P is reversible, then $-\bar{g}(\theta)$ is exactly the gradient of the function $f(\theta)$. Theorem 1 may be viewed as a way to generalize this observation to the non-reversible case.

4 Consequences

We now discuss several consequences of our new interpretation of temporal difference learning. These will all be along the lines of improved convergence guarantees. Indeed, as we mentioned in the previous section, viewing TD learning as gradient splitting allows us to take existing results for gradient descent and “port” them almost verbatim to the temporal difference setting.

In the main body of the paper, we focus on TD(0); the case of TD(λ) is discussed in the supplementary information. As mentioned earlier, existing analyses of TD(0) rely on (9) as well as its refinement (10). However, as a consequence of Proposition 1, we can actually write out explicitly the inner product between the mean TD(0) direction $\bar{g}(\theta)$ and the direction to optimality $\theta^* - \theta$.

Corollary 1. For any $\theta \in \mathbb{R}^K$,

$$(\theta^* - \theta)^T \bar{g}(\theta) = (1 - \gamma) \|V_{\theta^*} - V_\theta\|_D^2 + \gamma \|V_{\theta^*} - V_\theta\|_{\text{Dir}}^2. \quad (12)$$

Proof. Indeed, we can use (7) to argue that

$$\begin{aligned} (\theta^* - \theta)^T \bar{g}(\theta) &= (\theta^* - \theta)^T (-\bar{g}(\theta^*) - (-\bar{g}(\theta))) \\ &= \frac{1}{2} (\theta^* - \theta)^T (\nabla f(\theta^*) - \nabla f(\theta)), \end{aligned} \quad (13)$$

where the last step follows by Theorem 1 and Proposition 1; here $f(\theta)$ is the function from Theorem 1.

However, for any quadratic function $q(\theta) = (\theta - a)^T P(\theta - a)$ where P is a symmetric matrix, we have that

$$(a - \theta)^T (\nabla q(a) - \nabla q(\theta)) = 2q(\theta).$$

Applying this to the function $f(\theta)$ in (13), we complete the proof. \blacksquare

This corollary should be contrasted to (9) and (10). It is clearly a strengthening of those equations. More importantly, this is an equality, whereas (9) and (10) are inequalities. We thus see that the average TD(0) direction makes more progress in the direction of the optimal solution compared to the previously available bounds.

In the remainder of the paper, we will use Corollary 1 to obtain improved convergence times for TD(0); also, a natural generalization of that Corollary which appeals to Theorem 2 instead of Theorem 1 results in improved convergence times for TD(λ), as explained in the supplementary information. We focus on a particular property which is natural in this context: scaling with the discount factor γ .

Indeed, as we discussed in the introduction, an undesirable feature of some of the existing analyses of temporal difference learning is that they scale multiplicatively with $1/(1 - \gamma)$. It is easy to see why this should be so: existing analyses rely on variations of (10), and as $\gamma \rightarrow 0$, that equation guarantees smaller and smaller progress towards the limit. Unfortunately, it is natural to set the discount factor close to 1 in order to avoid focusing on short-term behavior of the policy.

But now we can instead rely on Corollary 1 and this corollary suggests that as $\gamma \rightarrow 1$, the inner product between the expected TD(0) direction $\bar{g}(\theta)$ and the direction to the optimal solution $\theta^* - \theta$ will be lower bounded by $\gamma \|V_\theta - V_{\theta^*}\|_{\text{Dir}}^2$. A difficulty, however, is that the Dirichlet semi-norm can be zero even when applied to a nonzero vector. We next discuss the results we are able to derive with this approach.

4.1 Improved error bounds

As mentioned earlier, a nice consequence of the gradient splitting interpretation is that we can apply the existing proof for gradient descent almost verbatim to gradient splittings. In particular, temporal difference learning when the states are sampled i.i.d. could be analyzed by simply following existing analyses of noisy gradient descent. However, under our Markov observation model, it is not true that the samples are i.i.d; rather, we proceed by modifying the analysis of so-called *Markov Chain Gradient Descent*, analyzed in the papers Sun et al. (2018); Johansson et al. (2010).

One issue is the choice of step-size. The existing literature on temporal difference learning contains a range of possible step-sizes from $O(1/t)$ to $O(1/\sqrt{t})$ (see Bhandari et al. (2018); Dalal et al. (2018); Lakshminarayanan and Szepesvari (2018)). A step-size that scales as $O(1/\sqrt{t})$ is often preferred because, for faster decaying step-sizes, performance will scale with the smallest eigenvalue of $\Phi^T D \Phi$ or related quantity, and these can be quite small. This is not the case, however, for a step-size that decays like $O(1/\sqrt{t})$.

We will be using the standard notation

$$\bar{\theta}_T = \frac{1}{T} \sum_{t=0}^{T-1} \theta_t,$$

to denote the running average of the iterates θ_t .

We will be considering the projected TD(0) update

$$\theta_{t+1} = \text{Proj}_\Theta(\theta_t + \alpha_t g_t(\theta_t)), \quad (14)$$

where Θ is a convex set containing the optimal solution θ^* . Moreover, we will assume that the norm of every element in Θ is at most R . Setting $G = r_{\max} + 2R$, we have the following error bound.

Corollary 2. *Suppose Assumptions 1-3 hold. Suppose further that $(\theta_t)_{t \geq 0}$ is generated by the Projected TD algorithm of (14) with $\theta^* \in \Theta$ and $\alpha_0 = \dots = \alpha_T = 1/\sqrt{T}$. Then*

$$E \left[(1-\gamma) \|V_{\theta^*} - V_{\bar{\theta}_T}\|_D^2 + \gamma \|V_{\theta^*} - V_{\bar{\theta}_T}\|_{\text{Dir}}^2 \right] \leq \frac{\|\theta^* - \theta_0\|_2^2 + G^2 [9 + 12\tau^{\text{mix}}(1/\sqrt{T})]}{2\sqrt{T}}. \quad (15)$$

The definition of τ^{mix} and proof is available in the supplementary information. We next compare this bound to the existing literature. The closest comparison is Theorem 3(a) in (Bhandari et al., 2018) which shows that

$$E \left[\|V_{\theta^*} - V_{\bar{\theta}_T}\|_D^2 \right] \leq \frac{\|\theta^* - \theta_0\|_2^2}{2(1-\gamma)\sqrt{T}} + \frac{G^2 [9 + 12\tau^{\text{mix}}(1/\sqrt{T})]}{2(1-\gamma)\sqrt{T}}. \quad (16)$$

Corollary 2 is stronger than this, because this bound can be derived from Corollary 2 by ignoring the second term on the left hand side of (15). Moreover, we next argue that Corollary 2 is stronger an interesting way, in that it offers a new insight on the behavior of temporal difference learning.

Observe that the bound of (16) blows up as $\gamma \rightarrow 1$. On the other hand, we can simply ignore the first term on the left-hand side of Corollary 2 to obtain

$$E \left[\|V_{\theta^*} - V_{\bar{\theta}_T}\|_{\text{Dir}}^2 \right] \leq \frac{\|\theta^* - \theta_0\|_2^2}{2\gamma\sqrt{T}} + \frac{G^2 [9 + 12\tau^{\text{mix}}(1/\sqrt{T})]}{2\gamma\sqrt{T}}. \quad (17)$$

In particular, we see that $E \left[\|V_{\theta^*} - V_{\bar{\theta}_T}\|_{\text{Dir}}^2 \right]$ does not blow up as $\gamma \rightarrow 1$. To understand this, recall that the Dirichlet semi-norm is equal to zero if and only if applied to a multiple of the all-ones vector. Consequently, $\|V\|_{\text{Dir}}$ is properly thought of as norm of the projection of V onto $\mathbf{1}^\perp$. We therefore obtain the punchline of this section: the error of (averaged & projected) temporal difference learning projected on $\mathbf{1}^\perp$ does not blow up as the discount factor approaches 1.

There are scenarios where this is already interesting. For example, if TD(0) is a subroutine of policy evaluation, it will be used for a policy improvement step, which is clearly unaffected by adding a multiple of the all-ones vector to the value function. Similarly, Proposition 4 of Ollivier (2018) shows that the bias in the policy gradient computed from an approximation \hat{V} to the true value function V can be bounded solely in terms of $\|V - \hat{V}\|_{\text{Dir}}^2$ (multiplied by a factor that depends on how the policies are parameterized).

It is natural to wonder whether the dependence on $1/(1-\gamma)$ can be removed completely from bounds on the performance of temporal difference learning (not just in terms of projection on $\mathbf{1}^\perp$). We address this next.

4.2 Mean-adjusted temporal difference learning

Unfortunately, it is easy to see that the dependence on $1/(1-\gamma)$ in error bounds for temporal difference learning cannot be entirely removed. We next give an informal sketch explaining why this is so. We consider the case where samples s, s' are i.i.d. with probability $\pi_s P(s, s')$ rather than coming from the Markov chain model, since this only makes the estimation problem easier.

Estimating the mean of the value function. Let us denote by V the true value function; because it is the fixed point of the Bellman operator, $V = R + \gamma PV$, we have that

$$V = (I - \gamma P)^{-1} R = \left(\sum_{m=0}^{\infty} \gamma^m P^m \right) R.$$

Define $\bar{V} = \pi^T V$; then

$$\bar{V} = \pi^T V = \pi^T \left(\sum_{m=0}^{\infty} \gamma^m P^m \right) R = \frac{\pi^T R}{1-\gamma}. \quad (18)$$

Under i.i.d. sampling, what we have are samples from a random variable \tilde{R} which takes the value $r(s, s')$ with probability $\pi_s P(s, s')$. From (18), we have that $(1-\gamma)\bar{V} = E[\tilde{R}]$.

From T samples of the scalar random variable \tilde{R} , the best estimate \hat{R}_T will satisfy $E[(\hat{R}_T - E[\tilde{R}])^2] = \Omega(1/T)$ in the worst-case. This implies that the best estimator \hat{V} of \bar{V} will satisfy $E[(\hat{V} - \bar{V})^2] = \Omega((1/(1-\gamma)^2)/T)$.

To summarize, the squared error in estimating just the mean of the value function will already scale with $1/(1-\gamma)$. If we consider e.g., Φ to be the identity matrix, in which case V_{θ^*} is just equal to the true value function, it easily follows that it is not possible to estimate V_{θ^*} with error that does not scale with $1/(1-\gamma)$.

A better scaling with the discount factor. Note, however, that the previous discussion implied that a term like $(1/(1-\gamma)^2)/T$ in a bound on the squared error is unavoidable. But with $1/\sqrt{T}$ step-size, the error will in general decay more slowly as $1/\sqrt{T}$ as in Corollary 2. Is it possible to derive a bound where the only scaling with $1/(1-\gamma)$ is in the asymptotically negligible $O(1/T)$ term?

As we show next, this is indeed possible. The algorithm is very natural given the discussion in the previous subsection: we run projected and averaged TD(0) and estimate the mean of the value function separately, adjusting the outcome of TD(0) to have the right mean in the end. Building on Corollary 2, the idea is that the mean will have expected square error that scales with $(1/(1-\gamma)^2)/T$ while the temporal difference method will estimate the projection onto $\mathbf{1}^\perp$ without blowing up as $\gamma \rightarrow 1$.

The pseudocode of the algorithm is given next as Algorithm 1 and Corollary 3 bounds its performance. Note that, in contrast to the bounds in the last subsection, Theorem 3 bounds the error to the true value function V directly. This is a more natural bound for this algorithm which tries to directly match the mean of the true value function.

Algorithm 1 Mean-adjusted TD(0)

- 1: Initialize $\bar{A}_{-1} = 0$, $s_0 = \pi$, and some initial condition θ_0 .
 - 2: **for** $t = 0$ to $T - 1$ **do**
 - 3: Projected TD(0) update:
 $\theta_{t+1} = \text{Proj}_\Theta(\theta_t + \alpha_t g_t(\theta_t))$
 - 4: Keep track of the average reward: $\bar{A}_t = \frac{t\bar{A}_{t-1} + r_t}{t+1}$
 - 5: **end for**
 - 6: Set $\hat{V}_T = \frac{\bar{A}_T}{1-\gamma}$
 - 7: Output $V'_T = V_{\hat{\theta}_T} + (\hat{V}_T - \pi^T V_{\hat{\theta}_T}) \mathbf{1}$
-

Corollary 3. Suppose that $(\theta_t)_{t \geq 0}$ and V'_T are generated by Algorithm 1 with step-sizes $\alpha_0 = \dots = \alpha_T = 1/\sqrt{T}$. Suppose further that Θ is a convex set that contains θ^* . Let t_0 be the largest integer which satisfies $t_0 \leq 2\tau^{\text{mix}}\left(\frac{1}{2(t_0+1)}\right)$. Then as long as $T \geq t_0$, we will have

$$E[\|V'_T - V\|_D^2] \leq O\left(\|V_{\theta^*} - V\|_D^2 + \frac{r_{\max}^2 \tau^{\text{mix}}\left(\frac{1}{2(T+1)}\right)}{(1-\gamma)^2 T} + \frac{\|\theta^* - \theta_0\|_2^2 + G^2 [1 + \tau^{\text{mix}}(1/\sqrt{T})]}{\sqrt{T}} \cdot \min\left\{\frac{r(P)}{\gamma}, \frac{1}{1-\gamma}\right\}\right),$$

where $r(P)$ is the inverse spectral gap of the additive reversibilization of the transition matrix P (formally defined in the supplementary information).

The bound of Corollary 3 has the structure mentioned earlier: the blowup with $1/(1-\gamma)$ occurs only in the $\tilde{O}(1/T)$. The last term, which scales like $\tilde{O}(1/\sqrt{T})$ gets multiplied by the minimum of $1/(1-\gamma)$ and a quantity that depends on the matrix P , so that it does not blow up as $\gamma \rightarrow 1$.

Note that, unlike the previous bounds discussed in this paper, this bound depends on an eigenvalue gap associated with the matrix P . However, this dependence is in such a way that it only helps: when $1/(1-\gamma)$ is small, there is no dependence on the eigenvalue gap, and it is only when $\gamma \rightarrow 1$ that performance “saturates” at something depending on an eigenvalue gap.

5 Conclusion

We have given a new interpretation of temporal difference learning in terms of a splitting of gradient descent. As a consequence of this interpretation, analyses of gradient descent can apply to temporal difference learning almost verbatim.

We have exploited this interpretation to observe that temporal difference methods learn the projection of the value function onto $\mathbf{1}^\perp$ without any blowup as $\gamma \rightarrow 1$; by contrast, previous work tended to have error bounds that scaled with

$1/(1 - \gamma)$. While, as we explain, it is not possible to remove the dependence on $O(1/(1 - \gamma))$ in general, we provide an error bound for a simple modification to TD(0) where the only dependence on $1/(1 - \gamma)$ is in the asymptotically negligible term.

An open problem might be to improve the scaling of the bounds we have obtained in this paper with P . Our focus has been on scaling with $1/(1 - \gamma)$ but one could further ask what dependence on the transition matrix P is optimal. It is natural to wonder, for example, whether the $r(P)$ factor in the last term of Corollary 3 measuring how the performance “saturates” as $\gamma \rightarrow 1$, could be improved. Typically, error bounds in this setting scale with the (non-squared) spectral gap of P , which can be much smaller than $r(P)$.

References

- Aldous, D. and Fill, J. (1995). Reversible markov chains and random walks on graphs.
- Bhandari, J., Russo, D., and Singal, R. (2018). A finite time analysis of temporal difference learning with linear function approximation. In *Conference on Learning Theory*, pages 1691–1692.
- Chen, C., Seff, A., Kornhauser, A., and Xiao, J. (2015). Deepdriving: Learning affordance for direct perception in autonomous driving. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2722–2730.
- Dalal, G., Szörényi, B., Thoppe, G., and Mannor, S. (2018). Finite sample analysis for td (0) with linear function approximation.
- Diaconis, P., Saloff-Coste, L., et al. (1996). Logarithmic sobolev inequalities for finite markov chains. *The Annals of Applied Probability*, 6(3):695–750.
- Gu, S., Holly, E., Lillicrap, T., and Levine, S. (2017). Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In *2017 IEEE international conference on robotics and automation (ICRA)*, pages 3389–3396. IEEE.
- Jin, J., Song, C., Li, H., Gai, K., Wang, J., and Zhang, W. (2018). Real-time bidding with multi-agent reinforcement learning in display advertising. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 2193–2201.
- Johansson, B., Rabi, M., and Johansson, M. (2010). A randomized incremental subgradient method for distributed optimization in networked systems. *SIAM Journal on Optimization*, 20(3):1157–1170.
- Korda, N. and La, P. (2015). On td (0) with function approximation: Concentration bounds and a centered variant with exponential convergence. In *International Conference on Machine Learning*, pages 626–634.
- Lakshminarayanan, C. and Szepesvari, C. (2018). Linear stochastic approximation: How far does constant step-size and iterate averaging go? In *International Conference on Artificial Intelligence and Statistics*, pages 1347–1355.
- Levin, D. A. and Peres, Y. (2017). *Markov chains and mixing times*, volume 107. American Mathematical Soc.
- Maei, H. R. (2011). Gradient temporal-difference learning algorithms.
- Narayanan, C. and Szepesvári, C. (2017). Finite time bounds for temporal difference learning with function approximation: Problems with some “state-of-the-art” results. Technical report, Technical Report.
- Ollivier, Y. (2018). Approximate temporal difference learning is a gradient descent for reversible policies. *arXiv preprint arXiv:1805.00869*.
- Robbins, H. and Monroe, S. (1951). A stochastic approximation method. *The Annals of Mathematical Statistics*, pages 400–407.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. (2016). Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489.
- Srikant, R. and Ying, L. (2019). Finite-time error bounds for linear stochastic approximation and td learning. In *Conference on Learning Theory*, pages 2803–2830.
- Sun, T., Sun, Y., and Yin, W. (2018). On markov chain gradient descent. In *Advances in Neural Information Processing Systems*, pages 9896–9905.
- Sutton, R. S. (1988). Learning to predict by the methods of temporal differences. *Machine Learning*, 3(1):9–44.
- Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT Press.
- Tsitsiklis, J. N. and Van Roy, B. (1997). An analysis of temporal-difference learning with function approximation. *IEEE Transactions on Automatic Control*, 42(5):674–690.

- Xiong, H., Xu, T., Liang, Y., and Zhang, W. (2020). Non-asymptotic convergence of adam-type reinforcement learning algorithms under markovian sampling. *arXiv preprint arXiv:2002.06286*.
- Xu, T., Zou, S., and Liang, Y. (2019). Two time-scale off-policy td learning: Non-asymptotic analysis over markovian samples. In *Advances in Neural Information Processing Systems*, pages 10633–10643.

A Proofs of Theorems 1 and 2

In this section, we present the detailed proof of the main results of our paper, i.e., Theorem 1 and 2. We begin with the proof of Theorem 1.

Before the proof, we introduce some necessary notation. Let ϕ be the feature vector of a random state generated according to the stationary distribution π . In other words, $\phi = \phi(s_k)$ with probability π_{s_k} . Let ϕ' be the feature vector of the next state s' and let $r = r(s, s')$. Thus ϕ and ϕ' are random vectors and r is a random variable. As shown in Equation (2) of Bhandari et al. (2018), $\bar{g}(\theta)$ can be written as

$$\bar{g}(\theta) = E[\phi r] + E[\phi(\gamma\phi' - \phi)^T] \theta.$$

With these notations in place, we begin the proof of Theorem 1.

Proof of Theorem 1. Recall that θ^* is the unique vector with $\bar{g}(\theta^*) = 0$ (see Lemma 6 in (Tsitsiklis and Van Roy, 1997)). Consider

$$\bar{g}(\theta) = \bar{g}(\theta) - \bar{g}(\theta^*) = E[\phi(\gamma\phi' - \phi)^T] (\theta - \theta^*). \quad (19)$$

To conclude that $\bar{g}(\theta)$ is a splitting of the gradient for a quadratic form $f(\theta)$, we need to calculate the gradient of $f(\theta)$. Let us begin with the Dirichlet norm and perform the following sequence of manipulations:

$$\begin{aligned} \|V_\theta - V_{\theta^*}\|_{\text{Dir}}^2 &= \frac{1}{2} \sum_{s, s' \in \mathcal{S}} \pi(s) P(s, s') [V_{\theta^*}(s) - V_\theta(s) - V_{\theta^*}(s') + V_\theta(s')]^2 \\ &= \frac{1}{2} \sum_{s, s' \in \mathcal{S}} \pi(s) P(s, s') [(V_{\theta^*}(s) - V_\theta(s))^2 + (V_{\theta^*}(s') - V_\theta(s'))^2] \\ &\quad - \sum_{s, s' \in \mathcal{S}} \pi(s) P(s, s') (V_{\theta^*}(s) - V_\theta(s)) (V_{\theta^*}(s') - V_\theta(s')) \\ &= \frac{1}{2} \sum_{s \in \mathcal{S}} \pi(s) \left(\sum_{s' \in \mathcal{S}} P(s, s') \right) (V_{\theta^*}(s) - V_\theta(s))^2 \\ &\quad + \frac{1}{2} \sum_{s' \in \mathcal{S}} \left(\sum_{s \in \mathcal{S}} \pi(s) P(s, s') \right) (V_{\theta^*}(s') - V_\theta(s'))^2 \\ &\quad - \sum_{s, s' \in \mathcal{S}} \pi(s) P(s, s') (\theta - \theta^*)^T \phi(s) \phi(s')^T (\theta - \theta^*) \\ &= \frac{1}{2} \sum_{s \in \mathcal{S}} \pi(s) (V_\theta(s) - V_{\theta^*}(s))^2 + \frac{1}{2} \sum_{s' \in \mathcal{S}} \phi(s') (V_\theta(s') - V_{\theta^*}(s'))^2 \\ &\quad - (\theta - \theta^*)^T E[\phi \phi'^T] (\theta - \theta^*) \\ &= \|V_\theta - V_{\theta^*}\|_D^2 - (\theta - \theta^*)^T E[\phi \phi'^T] (\theta - \theta^*). \end{aligned} \quad (20)$$

In the above sequence of equations, the first equality is just the definition of Dirichlet semi-norm; the second equality follows by expanding the square; the third equality follows by interchanging sums and the definition of V_θ ; the fourth equality uses that π is a stationary distribution of P , as well as the definition of ϕ and ϕ' ; and the final equality uses the definition of the $\|\cdot\|_D$ norm.

Our next step is to use the identity we have just derived to rearrange the definition of $\|V_\theta - V_{\theta^*}\|_D^2$:

$$\begin{aligned}
\|V_\theta - V_{\theta^*}\|_D^2 &= (V_\theta - V_{\theta^*})^T D (V_\theta - V_{\theta^*}) = (\theta - \theta^*)^T \Phi^T D \Phi (\theta - \theta^*) \\
&= (\theta - \theta^*)^T \sum_{s \in \mathcal{S}} \pi(s) \phi(s) \phi(s)^T (\theta - \theta^*) \\
&= (\theta - \theta^*)^T E[\phi \phi^T] (\theta - \theta^*).
\end{aligned} \tag{21}$$

We now use these identities to write down a new expression for the function $f(\theta)$:

$$\begin{aligned}
f(\theta) &= (1 - \gamma) \|V_\theta - V_{\theta^*}\|_D^2 + \gamma \|V_\theta - V_{\theta^*}\|_{\text{Dir}}^2 \\
&= (1 - \gamma) \|V_\theta - V_{\theta^*}\|_D^2 + \gamma (\|V_\theta - V_{\theta^*}\|_D^2 - (\theta - \theta^*)^T E[\phi \phi^T] (\theta - \theta^*)) \\
&= \|V_\theta - V_{\theta^*}\|_D^2 - \gamma (\theta - \theta^*)^T E[\phi \phi^T] (\theta - \theta^*) \\
&= (\theta - \theta^*)^T E[\phi \phi^T] (\theta - \theta^*) - \gamma (\theta - \theta^*)^T E[\phi \phi^T] (\theta - \theta^*) \\
&= (\theta - \theta^*)^T E[\phi(\phi - \gamma \phi')] (\theta - \theta^*).
\end{aligned}$$

In the above sequence of equations, the first equality is just the definition of the two norms; the second equality is obtained by plugging in (20); the third equality is obtained by cancellation of terms; the fourth equality is obtained by plugging in (21); and the last step follows by merging the two terms together.

As a consequence of writing $f(\theta)$ this way, we can write down a new expression for the gradient of $f(\theta)$ directly:

$$\nabla f(\theta) = (E[\phi(\phi - \gamma \phi')^T] + E[(\phi - \gamma \phi') \phi^T]) (\theta - \theta^*). \tag{22}$$

Combining (19) and (22), it is immediately that $-\bar{g}(\theta)$ is a splitting of $\nabla f(\theta)$. ■

We next turn to the proof of Theorem 2. Before beginning the proof, we introduce some notations.

The operator $T^{(\lambda)}$ is defined as:

$$(T^{(\lambda)} J)(s) = (1 - \lambda) \sum_{m=0}^{\infty} \lambda^m E \left[\sum_{t=0}^m \gamma^t r(s_t, s_{t+1}) + \gamma^{m+1} J(s_{m+1}) \mid s_0 = s \right] \tag{23}$$

for vectors $J \in \mathbb{R}^n$. The expectation is taken over sample paths taken by following actions according to policy μ ; recalling that this results in the transition matrix P , we can write this as

$$T^{(\lambda)} J = (1 - \lambda) \sum_{m=0}^{\infty} \lambda^m \sum_{t=0}^m \gamma^t P^t R + (1 - \lambda) \sum_{m=0}^{\infty} \lambda^m \gamma^{m+1} P^{m+1} J. \tag{24}$$

We next devise new notation that is analogous to the TD(0) case. Let us denote the quantity $\delta_t z_t$ by $x(\theta_t, z_t)$ and its steady-state mean by $\bar{x}(\theta)$. It is known that

$$\bar{x}(\theta) = \Phi^T D (T^{(\lambda)}(\Phi \theta) - \Phi \theta), \tag{25}$$

see Lemma 8 of [Tsitsiklis and Van Roy \(1997\)](#); it also shown there that $\text{TD}(\lambda)$ converges to a unique fixed point of a certain Bellman equation which we'll denote by θ_λ^* , and which satisfies

$$\bar{x}(\theta_\lambda^*) = 0. \tag{26}$$

With these preliminaries in place, we can begin the proof.

Proof of Theorem 2. By the properties of $T^{(\lambda)}$ and $\bar{x}(\theta)$ given in (25) and (23) respectively, we have

$$\begin{aligned}
\bar{x}(\theta) &= \bar{x}(\theta) - \bar{x}(\theta_\lambda^*) \\
&= \Phi^T D (T^{(\lambda)}(\Phi \theta) - \Phi \theta) - \Phi^T D (T^{(\lambda)}(\Phi \theta_\lambda^*) - \Phi \theta_\lambda^*) \\
&= \Phi^T D (T^{(\lambda)}(\Phi \theta) - T^{(\lambda)}(\Phi \theta_\lambda^*) - \Phi(\theta - \theta_\lambda^*)) \\
&= \left[(1 - \lambda) \sum_{m=0}^{\infty} \lambda^m \gamma^{m+1} \Phi^T D P^{m+1} \Phi - \Phi^T D \Phi \right] (\theta - \theta_\lambda^*),
\end{aligned} \tag{27}$$

where the last line used (24).

Our next step is to derive a convenient expression for $f^{(\lambda)}(\theta)$. We begin by finding a clean expression for the Dirichlet form that appears in the definition of $f^{(\lambda)}(\theta)$:

$$\begin{aligned}
\|V_\theta - V_{\theta_\lambda^*}\|_{\text{Dir}, m+1}^2 &= \frac{1}{2} \sum_{s, s' \in \mathcal{S}} \pi_s P^{m+1}(s, s') (V_\theta(s) - V_{\theta_\lambda^*}(s) - V_\theta(s') + V_{\theta_\lambda^*}(s'))^2 \\
&= \frac{1}{2} \sum_{s, s' \in \mathcal{S}} \pi_s P^{m+1}(s, s') \left[(V_\theta(s) - V_{\theta_\lambda^*}(s))^2 + (V_\theta(s') - V_{\theta_\lambda^*}(s'))^2 \right] \\
&\quad - \sum_{s, s' \in \mathcal{S}} \pi_s P^{m+1}(s, s') (V_\theta(s) - V_{\theta_\lambda^*}(s)) (V_\theta(s') - V_{\theta_\lambda^*}(s')) \\
&= \frac{1}{2} \sum_{s \in \mathcal{S}} \pi_s \left(\sum_{s' \in \mathcal{S}} P^{m+1}(s, s') \right) (V_\theta(s) - V_{\theta_\lambda^*}(s))^2 + \frac{1}{2} \sum_{s' \in \mathcal{S}} \left(\sum_{s \in \mathcal{S}} \pi_s P^{m+1}(s, s') \right) (V_\theta(s') - V_{\theta_\lambda^*}(s'))^2 \\
&\quad - \sum_{s, s' \in \mathcal{S}} \pi_s P^{m+1}(s, s') (V_\theta(s) - V_{\theta_\lambda^*}(s)) (V_\theta(s') - V_{\theta_\lambda^*}(s')) \\
&= \frac{1}{2} \sum_{s \in \mathcal{S}} \pi_s (V_\theta(s) - V_{\theta_\lambda^*}(s))^2 + \frac{1}{2} \sum_{s' \in \mathcal{S}} \pi_{s'} (V_\theta(s') - V_{\theta_\lambda^*}(s'))^2 \\
&\quad - \sum_{s, s' \in \mathcal{S}} \pi_s P^{m+1}(s, s') (V_\theta(s) - V_{\theta_\lambda^*}(s)) (V_\theta(s') - V_{\theta_\lambda^*}(s')) \\
&= \sum_{s \in \mathcal{S}} \pi_s (V_\theta(s) - V_{\theta_\lambda^*}(s))^2 - \sum_{s \in \mathcal{S}} \pi_s (V_\theta(s) - V_{\theta_\lambda^*}(s)) \sum_{s' \in \mathcal{S}} P^{m+1}(s, s') (V_\theta(s') - V_{\theta_\lambda^*}(s')) \\
&= (\theta - \theta_\lambda^*)^T (\Phi^T D \Phi - \Phi^T D P^{m+1} \Phi) (\theta - \theta_\lambda^*). \tag{28}
\end{aligned}$$

In the above sequence of equations, the first equality follows by the definition of the $m+1$ -Dirichlet norm; the second equality follows by expanding the square; the third equality follows by interchanging the order of summations; the fourth equality uses that any power of a stochastic matrix is stochastic, and the $\pi P^{m+1} = \pi$; the fifth equality combines terms and rearranges the order of summation; and the last line uses the definition $V_\theta = \Phi \theta$.

We'll also make use of the obvious identity

$$\|V_\theta - V_{\theta_\lambda^*}\|_D^2 = (\theta - \theta_\lambda^*)^T \Phi^T D \Phi (\theta - \theta_\lambda^*). \tag{29}$$

Putting all this together, we can express the function $f^{(\lambda)}(\theta)$ as:

$$\begin{aligned}
f^{(\lambda)}(\theta) &= (1 - \gamma \kappa) \|V_\theta - V_{\theta_\lambda^*}\|_D^2 + (1 - \lambda) \sum_{m=0}^{+\infty} \lambda^m \gamma^{m+1} \|V_\theta - V_{\theta_\lambda^*}\|_{\text{Dir}, m+1}^2 \\
&= (\theta - \theta_\lambda^*)^T \left[(1 - \gamma \kappa) \Phi^T D \Phi + (1 - \lambda) \sum_{m=0}^{\infty} \lambda^m \gamma^{m+1} \Phi^T D (I - P^{m+1}) \Phi \right] (\theta - \theta_\lambda^*) \\
&= (\theta - \theta_\lambda^*)^T \left[\left((1 - \gamma \kappa) + (1 - \lambda) \sum_{m=0}^{\infty} \lambda^m \gamma^{m+1} \right) \Phi^T D \Phi - (1 - \lambda) \sum_{m=0}^{\infty} \lambda^m \gamma^{m+1} \Phi^T D P^{m+1} \Phi \right] (\theta - \theta_\lambda^*) \\
&= (\theta - \theta_\lambda^*)^T \left[\left((1 - \gamma \kappa) + \gamma \frac{1 - \lambda}{1 - \gamma \lambda} \right) \Phi^T D \Phi - (1 - \lambda) \sum_{m=0}^{\infty} \lambda^m \gamma^{m+1} \Phi^T D P^{m+1} \Phi \right] (\theta - \theta_\lambda^*) \\
&= (\theta - \theta_\lambda^*)^T \left[\Phi^T D \Phi - (1 - \lambda) \sum_{m=0}^{\infty} \lambda^m \gamma^{m+1} \Phi^T D P^{m+1} \Phi \right] (\theta - \theta_\lambda^*).
\end{aligned}$$

In the above sequence of equations, the first equality is from the definition of the function $f^{(\lambda)}(\theta)$; the second line comes from plugging in (29) and (28); the third equality from breaking the sum in the second term into two pieces, one of which is then absorbed into the first term; the fourth equality follows by using the sum of a geometric series; and the last equality by the definition of κ from the theorem statement, which, recall, is $\kappa = (1 - \lambda)/(1 - \gamma \lambda)$.

By comparing the expression for $f^{(\lambda)}(\theta)$ we have just derived to (27), it is immediate that $-\bar{x}(\theta)$ is a splitting of the gradient of $f^{(\lambda)}(\theta)$. \blacksquare

B Proof of Corollary 2

We will be using standard notation for the mixing time of the Markov Chain:

$$\tau^{\text{mix}}(\varepsilon) = \min \{t \in \mathbb{N}, t \geq 1 | m\rho^t \leq \varepsilon\}. \quad (30)$$

We will find it convenient to use several observations made in [Bhandari et al. \(2018\)](#). First, Lemma 6 of that paper says that, under the assumptions of Corollary 2, we have that

$$\|g_t(\theta_t)\|_2 \leq G = r_{\max} + 2R. \quad (31)$$

This holds with probability one; note, however, that because the number of states and actions is finite, this just means one takes the maximum over all states and actions to obtain this upper bound.

A second lemma from [Bhandari et al. \(2018\)](#) deals with a measure of “gradient bias,” the quantity $\zeta_t(\theta) = (\bar{g}(\theta) - g_t(\theta))^T(\theta^* - \theta)$. As should be unsurprising, what matters in the analysis is not the natural measure of gradient bias, e.g., $\bar{g}(\theta) - g_t(\theta)$, but rather how the angle with the direction to the optimal solution is affected, which is precisely what is measured by $\zeta_t(\theta)$. We have the following upper bound.

Lemma 1 (Lemma 11 in [Bhandari et al. \(2018\)](#)). *Consider a non-increasing step-size sequence, $\alpha_0 \geq \alpha_1 \geq \dots \geq \alpha_T$. Fix any $t < T$, and set $t^* = \max\{0, t - \tau^{\text{mix}}(\alpha_T)\}$. Then*

$$E[\zeta_t(\theta_t)] \leq G^2 (4 + 6\tau^{\text{mix}}(\alpha_T)) \alpha_{t^*}.$$

With these preliminaries in place, we are now ready to prove the corollary. The proof follows the steps of [Sun et al. \(2018\)](#) to analyze Markov gradient descent, using the fact that the gradient splitting has the same inner product with the direction to the optimal solution as the gradient.

Proof of Corollary 2. From the projected TD(0) recursion, for any t ,

$$\begin{aligned} \|\theta^* - \theta_{t+1}\|_2^2 &= \|\theta^* - \text{Proj}_{\Theta_R}(\theta_t + \alpha_t g_t(\theta_t))\|_2^2 \\ &\leq \|\theta^* - \theta_t - \alpha_t g_t(\theta_t)\|_2^2 \\ &= \|\theta^* - \theta_t\|_2^2 - 2\alpha_t g_t(\theta_t)^T(\theta^* - \theta_t) + \alpha_t^2 \|g_t(\theta_t)\|_2^2 \\ &= \|\theta^* - \theta_t\|_2^2 - 2\alpha_t [\bar{g}(\theta_t)^T - (\bar{g}(\theta) - g_t(\theta))^T](\theta^* - \theta_t) + \alpha_t^2 \|g_t(\theta_t)\|_2^2 \\ &= \|\theta^* - \theta_t\|_2^2 - 2\alpha_t \bar{g}(\theta_t)^T(\theta^* - \theta_t) + 2\alpha_t \zeta_t(\theta_t) + \alpha_t^2 \|g_t(\theta_t)\|_2^2 \\ &\leq \|\theta^* - \theta_t\|_2^2 - 2\alpha_t \bar{g}(\theta_t)^T(\theta^* - \theta_t) + 2\alpha_t \zeta_t(\theta_t) + \alpha_t^2 G^2. \end{aligned}$$

In the above sequence of equations, all the equalities are just rearrangements of terms; whereas the first inequality follows that the projection onto a convex set does not increase distance, while the second inequality follows by (31).

Next we use Corollary 1, rearrange terms, and sum from $t = 0$ to $t = T - 1$:

$$\begin{aligned} &\sum_{t=0}^{T-1} 2\alpha_t E[(1-\gamma)\|V_{\theta^*} - V_{\theta_t}\|_D^2 + \gamma\|V_{\theta^*} - V_{\theta_t}\|_{\text{Dir}}^2] \\ &\leq \sum_{t=0}^{T-1} (E[\|\theta^* - \theta_t\|_2^2] - E[\|\theta^* - \theta_{t+1}\|_2^2]) + \sum_{t=0}^{T-1} 2\alpha_t E[\zeta_t(\theta_t)] + \sum_{t=0}^{T-1} \alpha_t^2 G^2 \\ &= (\|\theta^* - \theta_0\|_2^2 - E[\|\theta^* - \theta_T\|_2^2]) + \sum_{t=0}^{T-1} 2\alpha_t E[\zeta_t(\theta_t)] + \sum_{t=0}^{T-1} \alpha_t^2 G^2 \\ &\leq \|\theta^* - \theta_0\|_2^2 + \sum_{t=0}^{T-1} 2\alpha_t E[\zeta_t(\theta_t)] + \sum_{t=0}^{T-1} \alpha_t^2 G^2. \end{aligned}$$

Now plugging in the step-sizes $\alpha_0 = \dots = \alpha_T = 1/\sqrt{T}$, it is immediate that

$$\sum_{t=0}^{T-1} E[(1-\gamma)\|V_{\theta^*} - V_{\theta_t}\|_D^2 + \gamma\|V_{\theta^*} - V_{\theta_t}\|_{\text{Dir}}^2] \leq \frac{\sqrt{T}}{2} (\|\theta^* - \theta_0\|_2^2 + G^2) + \sum_{t=0}^{T-1} E[\zeta_t(\theta_t)].$$

Using Lemma 1, have that

$$\begin{aligned} \sum_{t=0}^{T-1} E[\zeta_t(\theta_t)] &\leq \sum_{t=0}^{T-1} G^2 \left(4 + 6\tau^{\text{mix}}(\alpha_T)\right) \alpha_{t^*} \\ &= \sqrt{T} G^2 \left(4 + 6\tau^{\text{mix}}\left(1/\sqrt{T}\right)\right). \end{aligned} \quad (32)$$

Putting all this together and using the convexity of the function $f(\theta)$, we can bound the error at the average iterate as:

$$\begin{aligned} E[(1-\gamma)\|V_{\theta^*} - V_{\bar{\theta}_T}\|_D^2 + \gamma\|V_{\theta^*} - V_{\bar{\theta}_T}\|_{\text{Dir}}^2] &\leq \frac{1}{T} \sum_{t=0}^{T-1} E[(1-\gamma)\|V_{\theta^*} - V_{\theta_t}\|_D^2 + \gamma\|V_{\theta^*} - V_{\theta_t}\|_{\text{Dir}}^2] \\ &\leq \frac{\|\theta^* - \theta_0\|_2^2 + G^2}{2\sqrt{T}} + \frac{G^2(4 + 6\tau^{\text{mix}}(1/\sqrt{T}))}{\sqrt{T}} \\ &= \frac{\|\theta^* - \theta_0\|_2^2 + G^2(9 + 12\tau^{\text{mix}}(1/\sqrt{T}))}{2\sqrt{T}}. \end{aligned}$$

■

C Proof of Corollary 3

Before starting the proof, we will need a collection of definitions, observations, and preliminary lemmas. We organize these into subheadings below.

The Dirichlet Laplacian. Let $L = (L(i, j))_{n \times n}$ be a symmetric matrix in $\mathbb{R}^{n \times n}$ defined as

$$L(i, j) = \begin{cases} -(1/2)(\pi_i P(i, j) + \pi_j P(j, i)) & \text{if } i \neq j \\ \sum_{i' \neq i} |L(i, i')| & \text{if } i = j \end{cases}.$$

It is immediate that the diagonal elements of L are positive and its rows sum to zero.

Furthermore, it can be shown that for any vector x , we have that $\|x\|_{\text{Dir}}^2 = x^T L x$. Indeed:

$$\begin{aligned} x^T L x &= \sum_{i=1}^n \left[\sum_{j \neq i} -\frac{1}{2} (\pi_i P(i, j) + \pi_j P(j, i)) x(i) x(j) + \left(\sum_{j \neq i} \frac{1}{2} (\pi_i P(i, j) + \pi_j P(j, i)) \right) x(i)^2 \right] \\ &= \sum_{i < j} \frac{1}{2} (\pi_i P(i, j) + \pi_j P(j, i)) (x(i) - x(j))^2 \\ &= \frac{1}{2} \sum_{i, j \in [n]} \frac{1}{2} (\pi_i P(i, j) + \pi_j P(j, i)) (x(i) - x(j))^2 \\ &= \frac{1}{2} \sum_{i, j \in [n]} \pi_i P(i, j) (x(i) - x(j))^2 = \|x\|_{\text{Dir}}^2. \end{aligned}$$

Connection to the reversed chain. We remark that the matrix L is connected to the so-called “additive reversibilization” of the matrix P , which we explain next. For a stochastic matrix P with stationary distribution π , it is natural to define the matrix P^* as

$$[P^*]_{ij} = \frac{\pi(j)}{\pi(i)} P_{ji}.$$

It is possible to verify that the matrix P^* has the same stationary distribution as the matrix P (see Aldous and Fill (1995)). Intuitively, the equality

$$\pi(i)[P^*]_{ij} = \pi(j)P_{ji},$$

means that it is natural to interpret P^* as the “reversed” chain of P : for all pairs i, j , the link from i to j is traversed as often under the stationary distribution in P^* as the link from j to i in P .

It can then be shown that the matrix $Q = (P + P^*)/2$ is reversible (see Aldous and Fill (1995)); this matrix is called the “additive reversibilization” of the matrix P . It is easy to see that $Q = I - D^{-1}L$; indeed, both the left-hand side and the right-hand side have the same off-diagonal entries and have rows that sum to one. Because Q is reversible, its spectrum is real.

The matrix $D^{-1}L$ is clearly similar to the symmetric matrix $D^{-1/2}LD^{-1/2}$ and thus has a real spectrum, with all the eigenvalues nonnegative. Moreover, $D^{-1}L$ has an eigenvalue of zero as $D^{-1}L\mathbf{1} = 0$. As a consequence of these two observations, if we denote by $r(P)$ the spectral gap of the matrix Q , then we have

$$r(P) = \frac{1}{1 - \lambda_2(Q)} = \frac{1}{\lambda_{n-1}(D^{-1}L)}, \quad (33)$$

where $\lambda_{n-1}(D^{-1}L)$ is the second smallest eigenvalue of $D^{-1}L$.

Equivalence of norms on $\mathbf{1}^\perp$. We will need to pass between the $\|\cdot\|_D$ norm and the $\|\cdot\|_{\text{Dir}}$ norm. To that end, we have the following lemma.

Lemma 2. *For any x with $\langle x, \mathbf{1} \rangle_D = 0$, we have that*

$$\|x\|_D^2 \leq r(P)\|x\|_{\text{Dir}}^2.$$

Proof. Indeed,

$$\min_{\langle x, \mathbf{1} \rangle_D = 0} \frac{\|x\|_{\text{Dir}}^2}{\|x\|_D^2} = \min_{\langle x, \mathbf{1} \rangle_D = 0} \frac{x^T L x}{\langle x, x \rangle_D} = \min_{\langle x, \mathbf{1} \rangle_D = 0} \frac{\langle x, D^{-1} L x \rangle_D}{\langle x, x \rangle_D}.$$

We next observe that the matrix $D^{-1}L$ is self adjoint in the $\langle \cdot, \cdot \rangle_D$ inner product:

$$\langle x, D^{-1} L y \rangle_D = x^T L y = \langle D^{-1} L x, y \rangle_D.$$

Since the smallest eigenvalue of $D^{-1}L$ is zero with associated eigenvector of $\mathbf{1}$, by the Rayleigh-Ritz theorem we have

$$\min_{\langle x, \mathbf{1} \rangle_D = 0} \frac{\langle x, D^{-1} L x \rangle_D}{\langle x, x \rangle_D} = \lambda_{n-1}(D^{-1}L).$$

Putting it all together, we obtain

$$\frac{\|x\|_{\text{Dir}}^2}{\|x\|_D^2} \geq \lambda_{n-1}(D^{-1}L) = r(P)^{-1},$$

where the last step used (33). This completes the proof. ■

Error in mean estimation.

Recall that we set \hat{V}_T be an estimate for the mean of value function in Algorithm 1. Our next lemma upper bounds the error in the estimate \hat{V}_T .

Lemma 3. *Suppose that \hat{V}_T is generated by Algorithm 1 and $\bar{V} = \pi^T V$ denote the mean of value function. Let $t_0 = \max \left\{ t \in \mathbb{N} \mid t_0 \leq 2\tau^{\text{mix}} \left(\frac{1}{2(t_0+1)} \right) \right\}$. Then, for $t > t_0$, we have*

$$E \left[(\hat{V}_t - \bar{V})^2 \right] \leq O \left(\frac{r_{\max}^2 \tau^{\text{mix}} \left(\frac{1}{2(t+1)} \right)}{(1-\gamma)^2 t} \right).$$

Proof. By the definition of \hat{V}_t and \bar{A}_t given in Algorithm 1, we can write the recursion:

$$\hat{V}_t = \frac{\bar{A}_t}{1-\gamma} = \frac{1}{1-\gamma} \left[\bar{A}_{t-1} + \frac{1}{t+1} (r_t - \bar{A}_{t-1}) \right] = \hat{V}_{t-1} + \frac{1}{t+1} \left(\frac{r_t}{1-\gamma} - \hat{V}_{t-1} \right).$$

We next use this recursion to argue:

$$\begin{aligned} E \left[(\hat{V}_t - \bar{V})^2 \right] &= E \left[\left(\hat{V}_{t-1} + \frac{1}{t+1} \left(\frac{r_t}{1-\gamma} - \hat{V}_{t-1} \right) - \bar{V} \right)^2 \right] \\ &= E \left[(\hat{V}_{t-1} - \bar{V})^2 + \frac{1}{(t+1)^2} \left(\frac{r_t}{1-\gamma} - \hat{V}_{t-1} \right)^2 + \frac{2}{t+1} \left(\frac{r_t}{1-\gamma} - \hat{V}_{t-1} \right) (\hat{V}_{t-1} - \bar{V}) \right] \\ &= E \left[(\hat{V}_{t-1} - \bar{V})^2 + \frac{1}{(t+1)^2} \left(\frac{r_t}{1-\gamma} - \hat{V}_{t-1} \right)^2 + \frac{2}{t+1} \left(\frac{r_t}{1-\gamma} - \bar{V} - \hat{V}_{t-1} + \bar{V} \right) (\hat{V}_{t-1} - \bar{V}) \right] \\ &= E \left[\left(1 - \frac{2}{t+1} \right) (\hat{V}_{t-1} - \bar{V})^2 + \frac{1}{(t+1)^2} \left(\frac{r_t}{1-\gamma} - \hat{V}_{t-1} \right)^2 + \frac{2}{t+1} \left(\frac{r_t}{1-\gamma} - \bar{V} \right) (\hat{V}_{t-1} - \bar{V}) \right]. \quad (34) \end{aligned}$$

To bound the second term on the right-hand side of (34), we will use that, since r_{\max} is the upper bound on absolute values of the rewards, we have that

$$\left(\frac{r_t}{1-\gamma} - \hat{V}_{t-1}\right)^2 \leq \left(\frac{r_{\max}}{1-\gamma} + \frac{r_{\max}}{1-\gamma}\right)^2 = \frac{4r_{\max}^2}{(1-\gamma)^2}.$$

We next analyze the third term on the right-hand side of (34). Let $\tau_t = \tau^{\text{mix}}\left(\frac{1}{2(t+1)}\right)$ so that for any state s'' ,

$$\sum_{s=1}^n |P^{\tau_t}(s'', s) - \pi_s| = 2d_{\text{TV}}(P^{\tau_t}(s'', \cdot), \pi) \leq 2m\rho^{\tau_t} \leq \frac{1}{T+1}. \quad (35)$$

We have that

$$\begin{aligned} E\left[\left(\frac{r_t}{1-\gamma} - \bar{V}\right)(\hat{V}_{t-1} - \bar{V})\right] &= E\left[\left(\frac{r_t}{1-\gamma} - \bar{V}\right)(\hat{V}_{t-1} - \hat{V}_{t-1-\tau_t} + \hat{V}_{t-1-\tau_t} - \bar{V})\right] \\ &= E\left[\left(\frac{r_t}{1-\gamma} - \bar{V}\right)(\hat{V}_{t-1} - \hat{V}_{t-1-\tau_t})\right] + E\left[\left(\frac{r_t}{1-\gamma} - \bar{V}\right)(\hat{V}_{t-1-\tau_t} - \bar{V})\right]. \end{aligned}$$

We now bound each of the two terms in the last equation separately. For the first term, we have

$$\begin{aligned} E\left[\left(\frac{r_t}{1-\gamma} - \bar{V}\right)(\hat{V}_{t-1} - \hat{V}_{t-1-\tau_t})\right] &\leq \frac{2r_{\max}}{1-\gamma} \sum_{d=t-\tau_t}^{t-1} E[|\hat{V}_d - \hat{V}_{d-1}|] \\ &= \frac{2r_{\max}}{1-\gamma} \sum_{d=t-\tau_t}^{t-1} \frac{1}{d+1} E\left[\left|\frac{r_d}{1-\gamma} - \hat{V}_{d-1}\right|\right] \\ &\leq \frac{4r_{\max}^2}{(1-\gamma)^2} \sum_{d=t-\tau_t}^{t-1} \frac{1}{d+1} \\ &\leq O\left(\frac{\tau_t r_{\max}^2}{(1-\gamma)^2(t+1)}\right). \end{aligned}$$

where the last inequality follows from $t > 2\tau_t$ (which in turn follows from $t \geq t_0$).

For the second term, we denote the following sigma algebra χ^t denote the sigma algebra generated by the information collected by time t , i.e., by the random variables $s_0, r_0, \theta_0, \dots, s_t, r_t, \theta_t$. We then have that

$$\begin{aligned}
E \left[\left(\frac{r_t}{1-\gamma} - \bar{V} \right) (\hat{V}_{t-1-\tau_t} - \bar{V}) \right] &= E \left[E \left[\left(\frac{r_t}{1-\gamma} - \bar{V} \right) (\hat{V}_{t-1-\tau_t} - \bar{V}) \mid \chi^{t-1-\tau_t} \right] \right] \\
&= E \left[\sum_{s=1}^n \left(\frac{\sum_{s'=1}^n P(s, s') r(s, s')}{1-\gamma} - \bar{V} \right) (\hat{V}_{t-1-\tau_t} - \bar{V}) P^{\tau_t}(s_{t-1-\tau_t}, s) \right] \\
&= E \left[\sum_{s=1}^n \left(\frac{\sum_{s'=1}^n P(s, s') r(s, s')}{1-\gamma} - \bar{V} \right) (\hat{V}_{t-1-\tau_t} - \bar{V}) (P^{\tau_t}(s_{t-1-\tau_t}, s) - \pi_s + \pi_s) \right] \\
&= E \left[\sum_{s=1}^n \left(\frac{\sum_{s'=1}^n P(s, s') r(s, s')}{1-\gamma} - \bar{V} \right) (\hat{V}_{t-1-\tau_t} - \bar{V}) (P^{\tau_t}(s_{t-1-\tau_t}, s) - \pi_s) \right] \\
&\quad + E \left[\sum_{s=1}^n \left(\frac{\sum_{s'=1}^n P(s, s') r(s, s')}{1-\gamma} - \bar{V} \right) (\hat{V}_{t-1-\tau_t} - \bar{V}) \pi_s \right] \\
&= E \left[\sum_{s=1}^n \left(\frac{\sum_{s'=1}^n P(s, s') r(s, s')}{1-\gamma} - \bar{V} \right) (\hat{V}_{t-1-\tau_t} - \bar{V}) (P^{\tau_t}(s_{t-1-\tau_t}, s) - \pi_s) \right] \\
&\quad + E \left[(\hat{V}_{t-1-\tau_t} - \bar{V}) \sum_{s=1}^n \left(\frac{\sum_{s'=1}^n P(s, s') r(s, s')}{1-\gamma} - \bar{V} \right) \pi_s \right] \\
&= E \left[\sum_{s=1}^n \left(\frac{\sum_{s'=1}^n P(s, s') r(s, s')}{1-\gamma} - \bar{V} \right) (\hat{V}_{t-1-\tau_t} - \bar{V}) (P^{\tau_t}(s_{t-1-\tau_t}, s) - \pi_s) \right] \\
&\quad + 0 \\
&\leq \frac{4r_{\max}^2}{(1-\gamma)^2(t+1)} \\
&\leq O \left(\frac{r_{\max}^2}{(1-\gamma)^2(t+1)} \right).
\end{aligned}$$

Here the first equality follows by iterating conditional expectation; the second, third, fourth, and fifth equality is just rearranging terms; the sixth equality follows from (18); and the next inequality follows from (35) as well as the fact that all rewards are upper bounded by r_{\max} in absolute value.

Combining all the inequalities, we can conclude that as long as $t > t_0$, we have that

$$E \left[(\hat{V}_t - \bar{V})^2 \right] \leq \left(1 - \frac{2}{t+1} \right) E \left[(\hat{V}_{t-1} - \bar{V})^2 \right] + O \left(\frac{\tau_t r_{\max}^2}{(1-\gamma)^2(t+1)^2} \right).$$

Let $b_t = O \left(\frac{\tau_t r_{\max}^2}{(1-\gamma)^2} \right)$; then the above equation can be compactly written as

$$E \left[(\hat{V}_t - \bar{V})^2 \right] \leq \left(1 - \frac{2}{t+1} \right) E \left[(\hat{V}_{t-1} - \bar{V})^2 \right] + \frac{b_t}{(t+1)^2}.$$

Let $C_t = \max \left\{ (t_0 + 1) (\hat{V}_{t_0} - \bar{V})^2, b_t \right\}$. We will prove by induction that $t \geq t_0$,

$$E \left[(\hat{V}_t - \bar{V})^2 \right] \leq \frac{C_t}{t+1}.$$

Indeed, the assertion holds for $t = t_0$. Suppose that the assertion holds at time t , i.e., suppose that $E[(\hat{V}_t - \bar{V})^2] \leq C_t/(t+1)$. Then,

$$\begin{aligned} E[(\hat{V}_{t+1} - \bar{V})^2] &\leq \left(1 - \frac{2}{t+2}\right) \frac{C_t}{t+1} + \frac{b_t}{(t+2)^2} \\ &= \frac{C_{t+1}}{t+2} + \left(1 - \frac{2}{t+2}\right) \frac{C_t}{t+1} + \frac{b_t}{(t+2)^2} - \frac{C_{t+1}}{t+2} \\ &= \frac{C_{t+1}}{t+2} + \frac{C_t(t+2)^2 - 2C_t(t+2) + b_t(t+1) - C_{t+1}(t+1)(t+2)}{(t+1)(t+2)^2} \\ &= \frac{C_{t+1}}{t+2} + \frac{(C_t - C_{t+1})(t+1)(t+2) + (b_t - C_t)(t+1) - C_t}{(t+1)(t+2)^2} \\ &\leq \frac{C_{t+1}}{t+2}, \end{aligned}$$

where the last inequality follows because $C_t \leq C_{t+1}$, $b_t \leq C_t$ and $C_t \geq 0$. Therefore, we have that, for $t \geq t_0$,

$$E[(\hat{V}_t - \bar{V})^2] \leq \frac{C_t}{t+1}.$$

Since $(\hat{V}_{t_0} - \bar{V})^2 \leq 4 \frac{r_{\max}^2}{(1-\gamma)^2}$ with probability one, and by definition $t_0 \leq 2\tau^{\max} \left(\frac{1}{2(t_0+1)}\right)$, we have that $C_t = O\left(\frac{r_{\max}^2}{(1-\gamma)^2}\right)$ for $t \geq t_0$; this completes the proof. \blacksquare

With all these preliminary lemmas in place, we can now give the main result of this section, the proof of Corollary 3.

Proof of Corollary 3. By the Pythagorean theorem, we have

$$\|V'_T - V\|_D^2 = \|\pi^T V'_T \mathbf{1} - \pi^T V \mathbf{1}\|_D^2 + \|V'_{T, \mathbf{1}^\perp} - V_{\mathbf{1}^\perp}\|_D^2, \quad (36)$$

where $V'_{T, \mathbf{1}^\perp}, V_{\mathbf{1}^\perp}$ are the projections of V'_T, V onto $\mathbf{1}^\perp$ in the $\langle \cdot, \cdot \rangle_D$ inner product.

Recall, that, in Algorithm 1, we defined

$$V'_T = V_{\hat{\theta}_T} + (\hat{V}_T - \pi^T V_{\hat{\theta}_T}) \mathbf{1}.$$

Therefore,

$$\begin{aligned} \pi^T V'_T \mathbf{1} &= \pi^T V_{\hat{\theta}_T} \mathbf{1} + \pi^T \mathbf{1} (\hat{V}_T - \pi^T V_{\hat{\theta}_T}) \mathbf{1} \\ &= \pi^T V_{\hat{\theta}_T} \mathbf{1} + \hat{V}_T \mathbf{1} - \pi^T V_{\hat{\theta}_T} \mathbf{1} \\ &= \hat{V}_T \mathbf{1}. \end{aligned}$$

Plugging this as well as $\bar{V} = \pi^T V$ into (36) we obtain:

$$\|V'_T - V\|_D^2 = \|\hat{V}_T \mathbf{1} - \bar{V} \mathbf{1}\|_D^2 + \|V'_{T, \mathbf{1}^\perp} - V_{\mathbf{1}^\perp}\|_D^2. \quad (37)$$

For the first term on the right hand side of (37), by the definition of the square norm under π , it is immediate that

$$\|\hat{V}_T \mathbf{1} - \bar{V} \mathbf{1}\|_D^2 = \sum_{i=1}^n \pi_i (\hat{V}_T - \bar{V})^2 = (\hat{V}_T - \bar{V})^2.$$

For the second term on the right hand side of (37), we have

$$\begin{aligned} \|V'_{T, \mathbf{1}^\perp} - V_{\mathbf{1}^\perp}\|_D^2 &= \|V'_{T, \mathbf{1}^\perp} - V_{\theta^*, \mathbf{1}^\perp} + V_{\theta^*, \mathbf{1}^\perp} - V_{\mathbf{1}^\perp}\|_D^2 \\ &\leq 2\|V'_{T, \mathbf{1}^\perp} - V_{\theta^*, \mathbf{1}^\perp}\|_D^2 + 2\|V_{\theta^*, \mathbf{1}^\perp} - V_{\mathbf{1}^\perp}\|_D^2 \\ &\leq 2r(P)\|V'_{T, \mathbf{1}^\perp} - V_{\theta^*, \mathbf{1}^\perp}\|_{\text{Dir}}^2 + 2\|V_{\theta^*} - V\|_D^2 \\ &= 2r(P)\|V_{\hat{\theta}_T} - V_{\theta^*}\|_{\text{Dir}}^2 + 2\|V_{\theta^*} - V\|_D^2, \end{aligned}$$

where the third line follows by the Lemma 2 and the Pythagorean theorem and the fourth line comes from the observation that $\|\cdot\|_{\text{Dir}}$ does not change when we add a multiple of $\mathbf{1}$.

Combining these results and taking expectation of (37), we obtain

$$\begin{aligned} E[\|V'_T - V\|_D^2] &\leq E[(\hat{V}_T - \bar{V})^2] + 2r(P)E[\|V_{\hat{\theta}_T} - V_{\theta^*}\|_{\text{Dir}}^2] + 2E[\|V_{\theta^*} - V\|_D^2] \\ &\leq O\left(\frac{\tau^{\text{mix}}\left(\frac{1}{2(T+1)}\right)r_{\max}^2}{(1-\gamma)^2T}\right) + r(P) \cdot \frac{\|\theta^* - \theta_0\|_2^2 + (9 + 12\tau^{\text{mix}}(1/\sqrt{T}))G^2}{\gamma\sqrt{T}} + 2\|V_{\theta^*} - V\|_D^2, \end{aligned} \quad (38)$$

where the second inequality follows by Lemma 3 and Eq. (17) from the main text.

On the other hand,

$$\begin{aligned} E[\|V'_T - V\|_D^2] &= E[\|V_{\hat{\theta}_T} - V\|_D^2] + E\left[\left\|\left(\hat{V}_T - \pi^T V_{\hat{\theta}_T}\right)\mathbf{1}\right\|_D^2\right] - 2E\left[\left(\pi^T V - \pi^T V_{\hat{\theta}_T}\right)\left(\hat{V}_T - \pi^T V_{\hat{\theta}_T}\right)\right] \\ &= E[\|V_{\hat{\theta}_T} - V\|_D^2] + E\left[\left\|\left(\hat{V}_T - \pi^T V_{\hat{\theta}_T}\right)\mathbf{1}\right\|_D^2\right] - 2E\left[\left(\bar{V} - \pi^T V_{\hat{\theta}_T}\right)\left(\hat{V}_T - \pi^T V_{\hat{\theta}_T}\right)\right] \\ &= E[\|V_{\hat{\theta}_T} - V\|_D^2] + E\left[\left(\hat{V}_T - \pi^T V_{\hat{\theta}_T}\right)^2\right] - 2E\left[\left(\hat{V}_T - \pi^T V_{\hat{\theta}_T} + \bar{V} - \hat{V}_T\right)\left(\hat{V}_T - \pi^T V_{\hat{\theta}_T}\right)\right] \\ &= E[\|V_{\hat{\theta}_T} - V\|_D^2] - E\left[\left(\hat{V}_T - \pi^T V_{\hat{\theta}_T}\right)^2\right] + 2E\left[\left(\hat{V}_T - \bar{V}\right)\left(\hat{V}_T - \pi^T V_{\hat{\theta}_T}\right)\right] \\ &\leq E[\|V_{\hat{\theta}_T} - V\|_D^2] - E\left[\left(\hat{V}_T - \pi^T V_{\hat{\theta}_T}\right)^2\right] + E\left[\left(\hat{V}_T - \bar{V}\right)^2 + \left(\hat{V}_T - \pi^T V_{\hat{\theta}_T}\right)^2\right] \\ &= E[\|V_{\hat{\theta}_T} - V\|_D^2] + E\left[\left(\hat{V}_T - \bar{V}\right)^2\right] \\ &\leq 2E[\|V_{\hat{\theta}_T} - V_{\theta^*}\|_D^2] + 2E[\|V_{\theta^*} - V\|_D^2] + E\left[\left(\hat{V}_T - \bar{V}\right)^2\right] \\ &\leq \frac{2\left[\|\theta^* - \theta_0\|_2^2 + (9 + 12\tau^{\text{mix}}(1/\sqrt{T}))G^2\right]}{(1-\gamma)\sqrt{T}} + 2E[\|V_{\theta^*} - V\|_D^2] + O\left(\frac{r_{\max}^2\tau^{\text{mix}}\left(\frac{1}{2(T+1)}\right)}{(1-\gamma)^2T}\right). \end{aligned} \quad (39)$$

Here the first four equalities come from rearranging; the next inequality comes from the identity $2ab \leq a^2 + b^2$; the next equality comes from cancellation; the next inequality uses $\|u + v\|_D^2 \leq 2\|u\|_D^2 + 2\|v\|_D^2$; and the final inequality uses Corollary 2 and Lemma 3.

We have just derived two bounds on $E[\|V'_T - V\|_D^2]$, one in Eq. (38) and one in Eq. (39). We could, of course, take the minimum of these two bounds. We then obtain:

$$\begin{aligned} E[\|V'_T - V\|_D^2] &\leq 2\|V_{\theta^*} - V\|_D^2 + O\left(\frac{\tau^{\text{mix}}\left(\frac{1}{2(T+1)}\right)r_{\max}^2}{(1-\gamma)^2T}\right) \\ &\quad + \min\left\{r(P) \cdot \frac{\|\theta^* - \theta_0\|_2^2 + (9 + 12\tau^{\text{mix}}(1/\sqrt{T}))G^2}{\gamma\sqrt{T}}, \right. \\ &\quad \left. \frac{2\|\theta^* - \theta_0\|_2^2 + 2(9 + 12\tau^{\text{mix}}(1/\sqrt{T}))G^2}{(1-\gamma)\sqrt{T}}\right\}. \end{aligned}$$

Therefore,

$$E[\|V'_T - V\|_D^2] \leq 2\|V_{\theta^*} - V\|_D^2 + O\left(\frac{\tau^{\text{mix}}\left(\frac{1}{T+1}\right)r_{\max}^2}{(1-\gamma)^2T}\right) + \frac{\|\theta^* - \theta_0\|_2^2 + G^2[1 + \tau^{\text{mix}}(1/\sqrt{T})]}{\sqrt{T}} \cdot \min\left\{\frac{r(P)}{\gamma}, \frac{2}{1-\gamma}\right\},$$

and the proof is complete. \blacksquare

D Error Bound for TD with Eligibility Traces

We now analyze the performance of projected TD(λ) which updates as

$$\theta_{t+1} = \text{Proj}_{\Theta_\lambda}(\theta_t + \alpha_t \delta_t \hat{z}_t), \quad (40)$$

where we now use

$$z_t = \sum_{k=0}^t (\gamma\lambda)^k \phi(s_{t-k}).$$

We remark that this is an abuse of notation, as previously z_t was defined with the sum starting at negative infinity, rather than zero; however, in this section, we will assume that the sum starts at zero. The consequence of this modification of notation is that Theorem 4 does not imply that $-E[z_t]$ is the gradient splitting of an appropriately defined function anymore, as now one needs to account for the error term coming from the beginning of the sum.

We assume Θ_λ is a convex set containing the optimal solution θ_λ^* . We will further assume that the norm of every element in Θ_λ is at most R_λ . Recall that

We begin by introducing some notation. Much of our analysis follows Bhandari et al. (2018) with some deviations where we appeal to Theorem 4, and the notation below is mostly identical to what is used in that paper. First, recall that we denote the quantity δ_{z_t} by $x(\theta_t, z_t)$. We define $\zeta_t(\theta, z_t)$ as a random variable which can be thought of as a measure of the bias that TD(λ) has in estimation of the gradient:

$$\zeta_t(\theta, z_t) = (\bar{x}(\theta) - \delta_{z_t})^T (\theta_\lambda^* - \theta).$$

Analogously to the TD(0) case, what turns out to matter for our analysis is not so much the bias per se, but the inner product of the bias with the direction of the optimal solution as in the definition of $\zeta_t(\theta, z_t)$.

We will next need an upper bound on how big $\|x(\theta, z_t)\|_2$ can get. Since under Assumption 2, we have that $\|\phi(s)\|_2 \leq 1$ for all s , we have that

$$\|z_t\|_2 \leq \frac{1}{1 - \gamma\lambda}.$$

Furthermore, we have that

$$|\delta_t| = |r(s, s') + \gamma\phi(s')^T \theta_t - \phi(s)^T \theta_t| \leq r_{\max} + 2R_\lambda,$$

where we used $|r(s, s')| \leq r_{\max}$ as well as Cauchy-Schwarz. Putting the last two equations together, we obtain

$$\|x(\theta, z_t)\|_2 \leq \frac{r_{\max} + 2R_\lambda}{1 - \gamma\lambda} := G_\lambda. \quad (41)$$

Compared to the result for TD(0), the bound depends on a slightly different definition of the mixing time that takes into account the geometric weighting in the eligibility trace. Define

$$\tau_\lambda^{\text{mix}}(\varepsilon) = \max\{\tau^{\text{MC}}(\varepsilon), \tau^{\text{Algo}}(\varepsilon)\},$$

where

$$\begin{aligned} \tau^{\text{MC}}(\varepsilon) &= \min\{t \in \mathbb{N}_0 | m\rho^t \leq \varepsilon\} \\ \tau^{\text{Algo}}(\varepsilon) &= \min\{t \in \mathbb{N}_0 | (\gamma\lambda)^t \leq \varepsilon\} \end{aligned}$$

The main result of this section is the following corollary of Theorem 2.

Corollary 4. Suppose Assumptions 1-2 hold. Suppose further that $(\theta_t)_{t \geq 0}$ is generated by the Projected TD(λ) algorithm of (40) with θ_λ^* belonging to the convex set Θ_λ and step-sizes $\alpha_0 = \dots = \alpha_T = 1/\sqrt{T}$. Then

$$E[f^{(\lambda)}(\theta)] \leq \frac{\|\theta_\lambda^* - \theta_0\|_2^2 + G_\lambda^2 [14 + 28\tau_\lambda^{\text{mix}}(1/\sqrt{T})]}{2\sqrt{T}},$$

where the function $f^{(\lambda)}(\theta)$ was defined in Theorem 2.

Proof. We begin with the standard recursion for the distance to the limit:

$$\begin{aligned} \|\theta_\lambda^* - \theta_{t+1}\|_2^2 &= \|\theta_\lambda^* - \text{Proj}_{\Theta_\lambda}(\theta_t + \alpha_t \delta_{z_t})\|_2^2 \\ &\leq \|\theta_\lambda^* - \theta_t - \alpha_t \delta_{z_t}\|_2^2 \\ &= \|\theta_\lambda^* - \theta_t\|_2^2 - 2\alpha_t \delta_{z_t}^T (\theta_\lambda^* - \theta_t) + \alpha_t^2 \|\delta_{z_t}\|_2^2 \\ &= \|\theta_\lambda^* - \theta_t\|_2^2 - 2\alpha_t \left(\bar{x}(\theta_t)^T - \left(\bar{x}(\theta_t)^T - \delta_{z_t}^T \right) \right) (\theta_\lambda^* - \theta_t) + \alpha_t^2 \|\delta_{z_t}\|_2^2 \\ &= \|\theta_\lambda^* - \theta_t\|_2^2 - 2\alpha_t \left(\bar{x}(\theta_t) - \bar{x}(\theta_\lambda^*) \right)^T (\theta_\lambda^* - \theta_t) + 2\alpha_t \zeta_t(\theta_t, z_t) + \alpha_t^2 \|x(\theta_t, z_t)\|_2^2 \\ &= \|\theta_\lambda^* - \theta_t\|_2^2 - 2\alpha_t \left[(1 - \gamma\kappa) \|V_\theta - V_{\theta_\lambda^*}\|_D^2 + (1 - \lambda) \sum_{m=0}^{\infty} \lambda^m \gamma^{m+1} \|V_\theta - V_{\theta_\lambda^*}\|_{\text{Dir}, m+1}^2 \right] \\ &\quad + 2\alpha_t \zeta_t(\theta_t, z_t) + \alpha_t^2 \|x(\theta_t, z_t)\|_2^2 \\ &\leq \|\theta_\lambda^* - \theta_t\|_2^2 - 2\alpha_t \left[(1 - \gamma\kappa) \|V_\theta - V_{\theta_\lambda^*}\|_D^2 + (1 - \lambda) \sum_{m=0}^{\infty} \lambda^m \gamma^{m+1} \|V_\theta - V_{\theta_\lambda^*}\|_{\text{Dir}, m+1}^2 \right] \\ &\quad + 2\alpha_t \zeta_t(\theta_t, z_t) + \alpha_t^2 G_\lambda^2. \end{aligned}$$

In the sequence of equations above the first inequality follows that the projection onto a convex set does not increase distance; the remaining equalities are rearrangements, using the quantity $\bar{x}(\theta)$ defined in (25), that $\bar{x}(\theta_\lambda^*) = 0$ from (26), and Proposition 1; and the final inequality used (41).

We next take expectations, rearrange terms, and sum:

$$\begin{aligned}
& \sum_{t=0}^{T-1} 2\alpha_t E \left[(1 - \gamma\kappa) \|V_\theta - V_{\theta_\lambda^*}\|_D^2 + (1 - \lambda) \sum_{m=0}^{\infty} \lambda^m \gamma^{m+1} \|V_\theta - V_{\theta_\lambda^*}\|_{\text{Dir}, m+1}^2 \right] \\
& \leq \sum_{t=0}^{T-1} (E[\|\theta_\lambda^* - \theta_t\|_2^2] - E[\|\theta_\lambda^* - \theta_{t+1}\|_2^2]) + \sum_{t=0}^{T-1} 2\alpha_t E[\zeta_t(\theta_t, z_t)] + \sum_{t=0}^{T-1} \alpha_t^2 G_\lambda^2 \\
& = (\|\theta_\lambda^* - \theta_0\|_2^2 - E[\|\theta_\lambda^* - \theta_T\|_2^2]) + \sum_{t=0}^{T-1} 2\alpha_t E[\zeta_t(\theta_t, z_t)] + \sum_{t=0}^{T-1} \alpha_t^2 G_\lambda^2 \\
& \leq \|\theta_\lambda^* - \theta_0\|_2^2 + \sum_{t=0}^{T-1} 2\alpha_t E[\zeta_t(\theta_t, z_t)] + \sum_{t=0}^{T-1} \alpha_t^2 G_\lambda^2.
\end{aligned}$$

Plugging in the step-sizes $\alpha_0 = \dots = \alpha_T = 1/\sqrt{T}$, we obtain

$$\begin{aligned}
& \sum_{t=0}^{T-1} E \left[(1 - \gamma\kappa) \|V_\theta - V_{\theta_\lambda^*}\|_D^2 + (1 - \lambda) \sum_{m=0}^{\infty} \lambda^m \gamma^{m+1} \|V_\theta - V_{\theta_\lambda^*}\|_{\text{Dir}, m+1}^2 \right] \\
& \leq \frac{\sqrt{T}}{2} (\|\theta_\lambda^* - \theta_0\|_2^2 + G_\lambda^2) + \sum_{t=0}^{T-1} E[\zeta_t(\theta_t, z_t)].
\end{aligned}$$

Using Lemma 20 in (Bhandari et al., 2018), we have that

$$\begin{aligned}
\sum_{t=0}^{T-1} E[\zeta_t(\theta_t, z_t)] & \leq 6\sqrt{T} \left(1 + 2\tau_\lambda^{\text{mix}}(\alpha_T) \right) G_\lambda^2 + \sum_{t=0}^{2\tau_\lambda^{\text{mix}}(\alpha_T)} (\gamma\lambda)^t G_\lambda^2 \\
& \leq 6\sqrt{T} \left(1 + 2\tau_\lambda^{\text{mix}}(\alpha_T) \right) G_\lambda^2 + \left(2\tau_\lambda^{\text{mix}}(\alpha_T) + 1 \right) G_\lambda^2.
\end{aligned}$$

Combining with convexity, we get

$$\begin{aligned}
& E \left[(1 - \gamma\kappa) \|V_{\theta_\lambda^*} - V_{\bar{\theta}_T}\|_D^2 + (1 - \lambda) \sum_{m=0}^{\infty} \lambda^m \gamma^{m+1} \|V_{\theta_\lambda^*} - V_{\bar{\theta}_T}\|_{\text{Dir}, m+1}^2 \right] \\
& \leq E \left[\frac{1}{T} \sum_{t=0}^{T-1} (1 - \gamma\kappa) \|V_\theta - V_{\theta_\lambda^*}\|_D^2 + (1 - \lambda) \sum_{m=0}^{\infty} \lambda^m \gamma^{m+1} \|V_\theta - V_{\theta_\lambda^*}\|_{\text{Dir}, m+1}^2 \right] \\
& \leq \frac{\|\theta_\lambda^* - \theta_0\|_2^2 + G_\lambda^2}{2\sqrt{T}} + \frac{6\sqrt{T} (1 + 2\tau_\lambda^{\text{mix}}(\alpha_T)) G_\lambda^2 + (2\tau_\lambda^{\text{mix}}(\alpha_T) + 1) G_\lambda^2}{T} \\
& \leq \frac{\|\theta_\lambda^* - \theta_0\|_2^2 + G_\lambda^2 (14 + 28\tau^{\text{mix}}(1/\sqrt{T}))}{2\sqrt{T}}.
\end{aligned}$$

■