

On Generalized Bellman Equations and Temporal-Difference Learning

Huizhen Yu

JANEY.HZYU@GMAIL.COM

*Reinforcement Learning and Artificial Intelligence Group
Department of Computing Science, University of Alberta
Edmonton, AB, T6G 2E8, Canada*

A. Rupam Mahmood

RUPAM@KINDRED.AI

*Kindred Inc.
243 College St
Toronto, ON M5T 1R5, Canada*

Richard S. Sutton

RSUTTON@UALBERTA.CA

*Reinforcement Learning and Artificial Intelligence Group
Department of Computing Science, University of Alberta
Edmonton, AB, T6G 2E8, Canada*

Editor: Csaba Szepesvári

Abstract

We consider off-policy temporal-difference (TD) learning in discounted Markov decision processes, where the goal is to evaluate a policy in a model-free way by using observations of a state process generated without executing the policy. To curb the high variance issue in off-policy TD learning, we propose a new scheme of setting the λ -parameters of TD, based on generalized Bellman equations. Our scheme is to set λ according to the eligibility trace iterates calculated in TD, thereby easily keeping these traces in a desired bounded range. Compared with prior work, this scheme is more direct and flexible, and allows much larger λ values for off-policy TD learning with bounded traces. As to its soundness, using Markov chain theory, we prove the ergodicity of the joint state-trace process under nonrestrictive conditions, and we show that associated with our scheme is a generalized Bellman equation (for the policy to be evaluated) that depends on both the evolution of λ and the unique invariant probability measure of the state-trace process. These results not only lead immediately to a characterization of the convergence behavior of least-squares based implementation of our scheme, but also prepare the ground for further analysis of gradient-based implementations.

Keywords: Markov decision process, approximate policy evaluation, generalized Bellman equation, reinforcement learning, temporal-difference method, Markov chain, randomized stopping time

1. Introduction

We consider discounted Markov decision processes (MDPs) and off-policy temporal-difference (TD) learning methods for approximate policy evaluation with linear function approxima-

tion. The goal is to evaluate a policy in a model-free way by using observations of a state process generated without executing the policy. Off-policy learning is an important part of the reinforcement learning methodology (Sutton and Barto, 1998) and has been studied in the areas of operations research and machine learning. (For an incomplete list of references, see e.g., Glynn and Iglehart, 1989; Precup et al., 2000, 2001; Randhawa and Juneja, 2004; Sutton et al., 2008, 2009; Maei, 2011; Yu, 2012; Dann et al., 2014; Geist and Scherrer, 2014; Mahadevan et al., 2014; Mahmood et al., 2014; Liu et al., 2015; Sutton et al., 2016; Dai et al., 2018.) Available TD algorithms, however, tend to have very high variances due to the use of importance sampling, an issue that limits their applicability in practice. The purpose of this paper is to introduce a new TD learning scheme that can help address this problem.

Our work is motivated by the recently proposed Retrace algorithm (Munos et al., 2016) and ABQ algorithm (Mahmood et al., 2017), and by the Tree-Backup algorithm (Precup et al., 2000) that existed earlier. These algorithms, as explained by Mahmood et al. (2017), all try to use the λ -parameters of TD to curb the high variance issue in off-policy learning. In particular, they all choose the values of λ according to the current state or state-action pair in such a way that guarantees the boundedness of the eligibility traces in TD learning, which can help reduce significantly the variance of the TD iterates. A limitation of these algorithms, however, is that they tend to be over-conservative and restrict λ to small values, whereas small λ can result in large approximation bias in TD solutions.

In this paper, we propose a new scheme of setting the λ -parameters of TD, based on generalized Bellman equations. Our scheme is to set λ according to the eligibility trace iterates calculated in TD, thereby easily keeping those traces in a desired bounded range. Compared with the schemes used in the previous work just mentioned, this is a direct way to bound the traces in TD, and it is also more flexible and allows much larger λ values for off-policy learning.

Regarding generalized Bellman equations, in our context, they will correspond to a family of dynamic programming equations for the policy to be evaluated. These equations all have the true value function as their unique solution, and their associated operators have contraction properties, like the standard Bellman operator. We will refer to the associated operators as generalized Bellman operators or Bellman operators for short. Some authors have considered, at least conceptually, the use of an even broader class of equations for policy evaluation. For example, Ueno et al. (2011) have considered treating the policy evaluation problem as a parameter estimation problem in the statistical framework of estimating equations, and in their framework, any equation that has the true value function as the unique solution can be used to estimate the value function. The family of generalized Bellman equations we consider has a more specific structure. They generalize multistep Bellman equations, and they are associated with randomized stopping times and arise from the strong Markov property (see Section 3.1 for details).

Generalized Bellman equations and operators are powerful tools. In classic MDP theory they have been used in some intricate optimality analyses (e.g., Schäl and Sudderth, 1987). Their computational use, however, seems to emerge primarily in the field of reinforcement learning. Through the λ -parameters and eligibility traces, TD learning is naturally connected with, not a single Bellman operator, but a family of Bellman operators, with different choices of λ or different rules of calculating the eligibility trace iterates corresponding

to different Bellman operators. Early efforts that use this aspect to broaden the scope of TD algorithms and to analyze such algorithms include Sutton’s work (1995) on learning at multiple timescales and Tsitsiklis’ work on generalized TD algorithms in the tabular case (see the book by Bertsekas and Tsitsiklis, 1996, Chap. 5.3). In the context of off-policy learning, there are more recent approaches that try to utilize this connection of TD with generalized Bellman operators to make TD learning more efficient (Precup et al., 2000; Yu and Bertsekas, 2012; Munos et al., 2016; Mahmood et al., 2017). This is also our aim, in proposing the new scheme of setting the λ -parameters.

Our analyses of the new TD learning scheme will focus on its theoretical side. Using Markov chain theory, we prove the ergodicity of the joint state and trace process under nonrestrictive conditions (see Theorem 2.1), and we show that associated with our scheme is a generalized Bellman equation (for the policy to be evaluated) that depends on both the evolution of λ and the unique invariant probability measure of the state-trace process (see Theorem 3.2 and Corollary 3.1). These results not only lead immediately to a characterization of the convergence behavior of least-squares based implementation of our scheme (see Corollary 2.1 and Remark 3.2), but also prepare the ground for further analysis of gradient-based implementations. (The latter analysis has been carried out recently by Yu (2017); see Remark 3.3.)

In addition to the theoretical study, we also present the results from a preliminary numerical study that compares several ways of setting λ for the least-squares based off-policy algorithm. The results demonstrate the advantages of the proposed new scheme with its greater flexibility.

We remark that although we shall focus exclusively on policy evaluation in this paper, approximate policy evaluation methods are highly pertinent to finding near-optimal policies in MDPs. They can be applied in approximate policy iteration, in policy-gradient algorithms for gradient estimation or in direct policy search (see e.g., Konda, 2002; Mannor et al., 2003). In addition to solving MDPs, they can also be used in artificial intelligence and robotics applications as a means to generate experience-based world models (see e.g., Sutton, 2009). It is, however, beyond the scope of this paper to discuss these applications of our results.

The rest of the paper is organized as follows. In Section 2, after a brief background introduction, we present our scheme of TD learning with bounded traces, and we establish the ergodicity of the joint state-trace process. In Section 3, we first discuss generalized Bellman operators associated with randomized stopping times, and we then derive the generalized Bellman equation associated with our scheme. In Section 4, we present the experimental results on the least-squares based implementation of our scheme. Appendices A-B include a proof for generalized Bellman operators and materials about approximation properties of TD solutions that are too long to include in the main text.

2. Off-Policy TD Learning with Bounded Traces

We describe the off-policy policy evaluation problem and the algorithmic form of TD learning in Section 2.1. We then present our scheme of history-dependent λ in Section 2.2, and analyze the properties of the resulting eligibility trace iterates and the convergence of the corresponding least-squares based algorithm in Section 2.3.

2.1 Preliminaries

The off-policy learning problem we consider in this paper concerns two Markov chains on a finite state space $\mathcal{S} = \{1, \dots, N\}$. The first chain has transition matrix P , and the second P^o . Whatever physical mechanisms that induce the two chains shall be denoted by π and π^o , and referred to as the target policy and behavior policy, respectively. The second Markov chain we can observe; however, it is the system performance of the first Markov chain that we want to evaluate.

Specifically, we consider a one-stage reward function $r_\pi : \mathcal{S} \rightarrow \mathbb{R}$ and an associated discounted total reward criterion with state-dependent discount factors $\gamma(s) \in [0, 1]$, $s \in \mathcal{S}$. Let Γ denote the $N \times N$ diagonal matrix with diagonal entries $\gamma(s)$. We assume that P and P^o satisfy the following conditions:

Condition 2.1 (Conditions on the target and behavior policies)

- (i) P is such that the inverse $(I - P\Gamma)^{-1}$ exists, and
- (ii) P^o is such that for all $s, s' \in \mathcal{S}$, $P_{ss'}^o = 0 \Rightarrow P_{ss'} = 0$, and moreover, P^o is irreducible.

The performance of π is defined as the expected discounted total rewards for each initial state $s \in \mathcal{S}$:

$$v_\pi(s) := \mathbb{E}_s^\pi [r_\pi(S_0) + \sum_{t=1}^{\infty} \gamma(S_1) \gamma(S_2) \cdots \gamma(S_t) \cdot r_\pi(S_t)], \quad (2.1)$$

where the notation \mathbb{E}_s^π means that the expectation is taken with respect to (w.r.t.) the Markov chain $\{S_t\}$ starting from $S_0 = s$ and induced by π (i.e., with transition matrix P). The function v_π is well-defined under Condition 2.1(i). It is called the *value function* of π , and by standard MDP theory (see e.g., Puterman, 1994), we can write it in matrix/vector notation as

$$v_\pi = r_\pi + P\Gamma v_\pi, \quad \text{i.e.,} \quad v_\pi = (I - P\Gamma)^{-1} r_\pi.$$

The first equation above is known as the Bellman equation (or dynamic programming equation) for a stationary policy (cf. Footnote 2).

We compute an approximation of v_π of the form $v(s) = \phi(s)^\top \theta$, $s \in \mathcal{S}$, where $\theta \in \mathbb{R}^n$ is a parameter vector and $\phi(s)$ is an n -dimensional feature representation for each state s (here $\phi(s), \theta$ are column vectors and the symbol $^\top$ stands for transpose). Data available for this computation are:

- (i) a realization of the Markov chain $\{S_t\}$ with transition matrix P^o generated by π^o , and
- (ii) rewards $R_t = r(S_t, S_{t+1})$ associated with state transitions, where the function r relates to $r_\pi(s)$ as $r_\pi(s) = \mathbb{E}_s^\pi[r(s, S_1)]$ for all $s \in \mathcal{S}$.¹

To find a suitable parameter θ for the approximation $\phi(s)^\top \theta$, we use the off-policy TD learning scheme. Define $\rho(s, s') = P_{ss'}/P_{ss'}^o$ (the importance sampling ratio),² and write

$$\rho_t = \rho(S_t, S_{t+1}), \quad \gamma_t = \gamma(S_t).$$

1. One can add to R_t a zero-mean finite-variance noise term. This makes little difference to our analyses, so we have left it out for notational simplicity.

2. Our problem formulation entails both value function and state-action value function estimation for a stationary policy in the standard MDP context. In these applications, it is the state-action space of the MDP that corresponds to the state space \mathcal{S} here. In particular, for value function estimation, S_t here

Given an initial $e_0 \in \mathfrak{R}^n$, for each $t \geq 1$, the eligibility trace vector $e_t \in \mathfrak{R}^n$ and the scalar temporal-difference term $\delta_t(v)$ for any approximate value function $v : \mathcal{S} \rightarrow \mathfrak{R}$ are calculated according to

$$e_t = \lambda_t \gamma_t \rho_{t-1} e_{t-1} + \phi(S_t), \quad (2.2)$$

$$\delta_t(v) = \rho_t (R_t + \gamma_{t+1} v(S_{t+1}) - v(S_t)). \quad (2.3)$$

Here $\lambda_t \in [0, 1], t \geq 1$, are important parameters in TD learning, the choice of which we shall elaborate on shortly.

There exist a number of TD algorithms that use e_t and δ_t to generate a sequence of parameters θ_t for approximate value functions. One such algorithm is LSTD (Boyan, 1999; Yu, 2012), which obtains θ_t by solving the linear equation for $\theta \in \mathfrak{R}^n$,

$$\frac{1}{t} \sum_{k=0}^{t-1} e_k \delta_k(v) = 0, \quad v = \Phi \theta \quad (2.4)$$

(if it admits a solution), where Φ is a matrix with row vectors $\phi(s)^\top, s \in \mathcal{S}$. LSTD updates the equation (2.4) iteratively by incorporating one by one the observation of (S_t, S_{t+1}, R_t) at each state transition. We will discuss primarily this algorithm in the paper, as its behavior can be characterized directly using our subsequent analyses of the joint state-trace process.

As mentioned earlier, our analyses will also provide bases for analyzing other gradient-based TD algorithms (e.g., Sutton et al., 2008, 2009; Maei, 2011; Mahadevan et al., 2014) by using stochastic approximation theory (Kushner and Yin, 2003; Borkar, 2008; Karmakar and Bhatnagar, 2018). Because of the complexity of this subject, however, we will not delve into it in the present paper, and we refer the reader to the recent work (Yu, 2017) for details.

2.2 Our Scheme of History-dependent λ

We now come to the choices of λ_t in the trace iterates (2.2). For TD with function approximation, one often lets λ_t be a constant or a function of S_t (Sutton, 1988; Tsitsiklis and Van Roy, 1997; Sutton and Barto, 1998). If neither the behavior policy nor the λ_t 's are further constrained, $\{e_t\}$ can have unbounded variances and is also unbounded in many natural situations (see e.g., Yu, 2012, Section 3.1), and this makes off-policy TD learning challenging.³ If we let the behavior policy to be close enough to the target policy so that $P^o \approx P$, then variance can be reduced, but it is not a satisfactory solution, for the applicability of off-policy learning would be seriously limited.

Without restricting the behavior policy, as mentioned earlier, the two recent papers (Munos et al., 2016; Mahmood et al., 2017), as well as the closely related early work by

corresponds to the pair of previous action and current state in the MDP, whereas for state-value function estimation, S_t here corresponds to the current state-action pair in the MDP. The ratio $\rho(s, s') = P_{ss'}^o / P_{ss'}$ then comes out as the ratio of action probabilities under π and π^o , the same as what appears in most of the off-policy learning literature. For the details of these correspondences, see (Yu, 2012, Examples 2.1, 2.2). The third application is in a simulation context where P^o corresponds to a simulated system and both P^o, P are known so that the ratio $\rho(s, s')$ is available. Such simulations are useful, for example, in studying system performance under perturbations, and in speeding up the computation when assessing the impacts of events that are rare under the dynamics P .

3. However, asymptotic convergence can still be ensured for several algorithms (Yu, 2012, 2015, 2016b), thanks partly to a powerful law of large numbers for stationary processes.

Precup et al. (2000), exploit state-dependent λ 's to control variance. Their choices of λ_t are such that $\lambda_t \gamma_t \rho_{t-1} < 1$ for all t , so that the trace iterates e_t are made bounded, which can help reduce the variance of the iterates.

Motivated by this prior work, our proposal is to set λ_t according to e_{t-1} directly, so that we can keep e_t in a desired range straightforwardly and at the same time, allow a much larger range of values for the λ -parameters. As a simple example, if we use λ_t to scale the vector $\gamma_t \rho_{t-1} e_{t-1}$ to be within a ball with some given radius, then we keep e_t always bounded.

In the rest of this paper, we shall focus on analyzing the iteration (2.2) with a particular choice of λ_t of the kind just mentioned. We want to be more general than the preceding simple example. However, since the dependence on the trace e_{t-1} would make λ_t dependent on the entire past history (S_0, \dots, S_{t-1}) , we also want to retain certain Markovian properties that are very useful for convergence analysis. This leads us to consider λ_t *being a certain function of the previous trace and past states*. More specifically, we will let λ_t be a function of the previous trace e_{t-1} and a certain memory state that is a summary of the states observed so far. The formulation is as follows.

2.2.1 FORMULATION AND EXAMPLES

We denote the memory state at time t by y_t . For simplicity, we assume that y_t can only take values from a finite set \mathcal{M} , and its evolution is Markovian: $y_t = g(y_{t-1}, S_t)$ for some given function g . The joint process $\{(S_t, y_t)\}$ is then a simple finite-state Markov chain. Each y_t is a function of the history (S_0, \dots, S_t) and y_0 . We further require, besides the irreducibility of $\{S_t\}$ (cf. Condition 2.1(ii)), that

Condition 2.2 (Evolution of memory states) *Under the behavior policy π^o , the Markov chain $\{(S_t, y_t)\}$ on $\mathcal{S} \times \mathcal{M}$ has a single recurrent class.*

This recurrence condition is nonrestrictive: If the Markov chain has multiple recurrent classes, each recurrent class can be treated separately by using the same arguments we present in this paper. However, we remark that the finiteness assumption on \mathcal{M} is a simplification. We choose to work with finite \mathcal{M} mainly for the reason that with the traces lying in a continuous space, to study the joint state and trace process, we need to resort to properties of Markov chains on infinite spaces. With an infinite \mathcal{M} , we would need to introduce more technical conditions that are not essential to our analysis and can obscure our main arguments.

We thus let y_t and λ_t evolve as

$$y_t = g(y_{t-1}, S_t), \quad \lambda_t = \lambda(y_t, e_{t-1}) \quad (2.5)$$

where $\lambda : \mathcal{M} \times \mathbb{R}^n \rightarrow [0, 1]$. We require the function λ to satisfy two conditions.

Condition 2.3 (Conditions for $\lambda(\cdot)$) *For some norm $\|\cdot\|$ on \mathbb{R}^n , the following hold for each memory state $y \in \mathcal{M}$:*

- (i) *For any $e, e' \in \mathbb{R}^n$, $\|\lambda(y, e) e - \lambda(y, e') e'\| \leq \|e - e'\|$.*
- (ii) *For some constant C_y , $\|\gamma(s') \rho(s, s') \cdot \lambda(y, e) e\| \leq C_y$ for all $e \in \mathbb{R}^n$ and all possible state transitions (s, s') that can lead to the memory state y .*

In the above, the second condition is to restrict $\{e_t\}$ in a desired range (as it makes $\|e_t\| \leq \max_{y \in \mathcal{M}} C_y + \max_{s \in \mathcal{S}} \|\phi(s)\|$). The first condition is about the continuity of the function $\lambda(y, e)e$ in the trace variable e for each memory state y , and it plays a key role in the subsequent analysis, where we will use this condition to ensure that the traces e_t and the states (S_t, y_t) jointly form a Markov chain with appealing properties. We shall defer a further discussion on the technical roles of these conditions to the end of Section 2.3 (cf. Remark 2.2).

Let us give a few simple examples of choosing λ that satisfy Condition 2.3. We will later use these examples in our experimental study (Section 4).

Example 2.1 We consider again the simple scaling example mentioned earlier and describe it using the terminologies just introduced. In this example, we let $y_t = (S_{t-1}, S_t)$. For each $y = (s, s')$, we define the function $\lambda(y, \cdot)$ so that when multiplied with $\lambda(y, e)$, the vector $\gamma(s')\rho(s, s')e$ is scaled down whenever its length exceeds a given threshold $C_{ss'}$:

$$\lambda(y, e) = \begin{cases} 1 & \text{if } \gamma(s')\rho(s, s')\|e\|_2 \leq C_{ss'}; \\ \frac{C_{ss'}}{\gamma(s')\rho(s, s')\|e\|_2} & \text{otherwise.} \end{cases} \quad (2.6)$$

Condition 2.3(i) is satisfied because for $y = (s, s')$ with $\gamma(s')\rho(s, s') = 0$, $\lambda(y, e)e = e$, whereas for $y = (s, s')$ with $\gamma(s')\rho(s, s') \neq 0$, $\lambda(y, e)e$ is simply the Euclidean projection of e onto the ball (centered at the origin) with radius $C_{ss'}/(\gamma(s')\rho(s, s'))$ and is therefore Lipschitz continuous in e with modulus 1 w.r.t. $\|\cdot\|_2$. Corresponding to (2.6), the update rule (2.2) of e_t becomes

$$e_t = \begin{cases} \gamma_t \rho_{t-1} e_{t-1} + \phi(S_t) & \text{if } \gamma_t \rho_{t-1} \|e_{t-1}\|_2 \leq C_{S_{t-1}S_t}; \\ C_{S_{t-1}S_t} \cdot \frac{e_{t-1}}{\|e_{t-1}\|_2} + \phi(S_t) & \text{otherwise.} \end{cases} \quad (2.7)$$

Note that this scheme of setting λ encourages the use of large λ_t : $\lambda_t = 1$ will be chosen whenever possible. A variation of the scheme is to multiply the right-hand side (r.h.s.) of (2.6) by another factor $\beta_{ss'} \in [0, 1]$, so that λ_t can be at most $\beta_{S_{t-1}S_t}$. In particular, one such variation is to simply multiply the r.h.s. of (2.6) by a constant $\beta \in (0, 1)$ so that $\lambda_t \leq \beta < 1$ for all t . \blacksquare

Example 2.2 The Retrace algorithm (Munos et al., 2016) modifies the trace updates in off-policy TD learning by truncating the importance sampling ratios by 1. In particular, for the off-policy TD(λ) algorithm with a constant $\lambda = \beta \in (0, 1]$, Retrace modifies the trace updates to be

$$e_t = \beta \gamma_t \cdot \min\{1, \rho_{t-1}\} \cdot e_{t-1} + \phi(S_t). \quad (2.8)$$

As pointed out by Mahmood et al. (2017), to retain the original interpretation of λ as a bootstrapping parameter in TD learning, we can rewrite the above update rule of Retrace equivalently as

$$e_t = \lambda_t \gamma_t \rho_{t-1} e_{t-1} + \phi(S_t) \quad \text{for } \lambda_t = \beta \cdot \frac{\min\{1, \rho_{t-1}\}}{\rho_{t-1}} \quad (\text{with } 0/0 = 0). \quad (2.9)$$

Each λ_t here is a function of (S_{t-1}, S_t) only and does not depend on e_{t-1} , so this choice of λ -parameters automatically satisfies Condition 2.3(i) with the memory states being $y_t =$

(S_{t-1}, S_t) . When the discount factors $\gamma(s)$ are all strictly less than 1, $\|e_t\|$ for all t are bounded by a deterministic constant that depends on the initial e_0 . Then for each initial e_0 , Retrace's choice of λ coincides with a choice in our framework, since the C -parameters in Condition 2.3(ii) can be made vacuously large so that the condition is satisfied by all the traces e_t that could be encountered by Retrace. Thus in this case our framework for choosing λ effectively encompasses the particular choice used by Retrace.

One can make variations on Retrace's trace update rule. For example, instead of truncating each importance sampling ratio $\rho(s, s')$ by 1, one can truncate it by a constant $K_{ss'} \geq 1$, and then use a scaling scheme similar to Example 2.1 to bound the traces. The simplest such variation is to choose two memory-independent positive constants K and C , and replace the definition of λ_t in (2.9) by the following: with $\tilde{\lambda}_t = \frac{\min\{K, \rho_{t-1}\}}{\rho_{t-1}}$ (where we treat $0/0 = 0$),

$$\lambda_t = \begin{cases} \beta \tilde{\lambda}_t & \text{if } \tilde{\lambda}_t \gamma_t \rho_{t-1} \|e_{t-1}\|_2 \leq C; \\ \beta \tilde{\lambda}_t \cdot \frac{C}{\tilde{\lambda}_t \gamma_t \rho_{t-1} \|e_{t-1}\|_2} & \text{otherwise.} \end{cases} \quad (2.10)$$

Correspondingly, instead of (2.8), the update rule of e_t becomes

$$e_t = \begin{cases} \beta \gamma_t \cdot \min\{K, \rho_{t-1}\} \cdot e_{t-1} + \phi(S_t) & \text{if } \gamma_t \cdot \min\{K, \rho_{t-1}\} \cdot \|e_{t-1}\|_2 \leq C; \\ \beta C \cdot \frac{e_{t-1}}{\|e_{t-1}\|_2} + \phi(S_t) & \text{otherwise.} \end{cases} \quad (2.11)$$

These variations of Retrace are similar to Example 2.1 and satisfy Condition 2.3.⁴ ■

2.2.2 COMPARISON WITH PREVIOUS WORK

For policy evaluation, the Retrace algorithm (Munos et al., 2016) and the ABQ algorithm (Mahmood et al., 2017) are very similar (ABQ was actually developed independently of Retrace before the Munos et al. (2016) paper was published, although the ABQ paper itself was released much later). Both Retrace and ABQ include the Tree-Backup algorithm (Precup et al., 2000) as a special case. They can use additional parameters to select λ from a range of values, whereas Tree-Backup specifies λ , implicitly, in a particular way (which has the advantage of requiring no knowledge of the behavior policy) and does not have the freedom in choosing λ . Because of the relations between these algorithms, when comparing our method to them, we will compare it with Retrace only. In the experimental study given later in Section 4 on the performance of LSTD for various ways of setting λ , we will compare our scheme of choosing λ with that of Retrace for $\beta = 1$, which lets Retrace use the largest λ that it can take.

We see in Example 2.2 that the eligibility trace update rule of Retrace can be written in two equivalent forms, (2.8) and (2.9). The second form (2.9) has the advantage that the λ -parameters involved are shown explicitly. In TD learning, the λ -parameters directly affect the associated Bellman operators and can be meaningfully interpreted as stopping probabilities (see Section 3), whereas the importance sampling ratio terms in the eligibility

4. To see this, let the memory states be $y_t = (S_{t-1}, S_t)$. For each $y = (s, s')$, let $\lambda(y, e)$ be defined according to (2.10), and let C_y in Condition 2.3(ii) be $C_{ss'} = \frac{\rho(s, s')}{\min\{K, \rho(s, s')\}} \cdot C$ (treat $0/0 = 0$). Then note that $\|\lambda(y, e)e - \lambda(y, e')e'\|_2 \leq \beta \min\left\{\frac{K}{\rho(s, s')}, 1\right\} \cdot \|e - e'\|_2 \leq \beta \|e - e'\|_2$.

trace iterates are essentially unchanged, for they have to be there in order to correct for the discrepancy between the behavior and target policies. For this reason, we prefer (2.9) to (2.8) and prefer thinking in terms of the selection of λ -parameters to that of what occurs *apparently* to those importance sampling ratio terms in the trace updates.

As mentioned in Example 2.2, the Munos et al. (2016) paper does not make the connection between (2.8) and (2.9). Mahmood et al. (2017) recognized the role of the λ -parameters and made explicit use of it to derive the ABQ algorithm. However, in the ABQ paper, the discussion and the presentation of the algorithm still emphasize the apparent changes in those importance sampling ratio terms in the trace iterates. This is an unsatisfactory point in that paper that we hope we have clarified with our present work.

We mentioned in the introduction that Retrace, ABQ and Tree-Backup are too conservative and tend to use too small λ values. Let us now make this statement more precise and also explain the reason behind.

These algorithms tend to behave effectively like $TD(\lambda)$ with small constant λ , despite that they can have $\lambda_t = 1$ at some time steps t . This is due to the nature of TD learning with time-varying λ , which is very different from that of TD with constant λ . For time-varying λ , a large λ_t at one time step need not mean that we are using the information of the cumulative rewards over a long time horizon to estimate the value at the state S_t encountered at time t . Because the next λ_{t+1} could be very small or even zero, forcing a TD algorithm to “bootstrap” immediately. When large λ_t ’s are interleaved with small ones, we are effectively in the situation of TD with small λ . This could occur to our proposed scheme as well if, for example, in Example 2.1 the thresholds $C_{ss'}$ are set too small. When we use larger thresholds, we allow larger λ . By comparison, Retrace, ABQ, and Tree-Backup constrain the state-dependent λ -parameters to be small enough so that all the products $\lambda_t \gamma_t \rho_{t-1} < 1$, and this makes them prone to the small- λ issue just mentioned. (See the experiments in Section 4.2 for demonstrations.)

While we consider Retrace for approximate policy evaluation, the Munos et al. (2016) paper actually focuses primarily on finding an optimal policy for an MDP, in the tabular case, and it has demonstrated good empirical performance of Retrace and Tree-Backup for that purpose. Despite this, its results are not adequate yet to establish asymptotic optimality of these algorithms in the online optimistic policy iteration setting (personal communication with Munos), and it is still an open theoretical question whether online TD algorithms can solve an MDP like the Q-learning algorithm (Watkins, 1989; Tsitsiklis, 1994), when positive λ (small or not) and rapidly changing target policies are involved.

We also mention that for policy evaluation, Munos et al. (2016, Section 3.1) have also conceived the use of generalized Bellman operators, although they did not relate these operators explicitly to history-dependent λ ’s and did not study corresponding algorithms in this general case.

2.3 Ergodicity Result

The properties of the joint state-trace process $\{(S_t, y_t, e_t)\}$ are important for understanding and characterizing the behavior of our proposed TD learning scheme. We study them in this subsection. Most importantly, we shall establish the ergodicity of the state-trace process. The result will be useful in convergence analysis of several associated TD algorithms (Yu,

2017), although in this paper we discuss only the LSTD algorithm. In the next section we will also use the ergodicity result when we relate the LSTD equation (2.4) to a generalized Bellman equation for the target policy in order to interpret the LSTD solutions.

We note that to obtain the results in this subsection, we will follow similar lines of argument used in (Yu, 2012) for analyzing off-policy LSTD with constant λ . However, because λ is now history-dependent, some proof steps in (Yu, 2012) no longer apply. We shall explain this in more detail after we prove the main result of this subsection.

As another side note, one can introduce nonnegative coefficients $i(y)$ for memory states y to weight the state features (similarly to the use of “interest” weights in the ETD algorithm (Sutton et al., 2016)) and update e_t according to

$$e_t = \lambda_t \gamma_t \rho_{t-1} e_{t-1} + i(y_t) \phi(S_t). \quad (2.12)$$

The results given below apply to this update rule as well.

Let us start with two basic properties of $\{(S_t, y_t, e_t)\}$ that follow directly from our choice of the λ function:

- (i) By Condition 2.3(i), for each y , $\lambda(y, e)e$ is a continuous function of e , and thus e_t depends continuously on e_{t-1} . This, together with the finiteness of $\mathcal{S} \times \mathcal{M}$, ensures that $\{(S_t, y_t, e_t)\}$ is a weak Feller Markov chain.⁵
- (ii) Then, by a property of weak Feller Markov chains (Meyn and Tweedie, 2009, Theorem 12.1.2(ii)), the boundedness of $\{e_t\}$ ensured by Condition 2.3(ii) implies that $\{(S_t, y_t, e_t)\}$ has at least one invariant probability measure.

The third property, given in the lemma below, concerns the behavior of $\{e_t\}$ for different initial e_0 . It is an important implication of Condition 2.3(i); actually, it is our purpose of introducing the condition 2.3(i) in the first place. In the lemma, $\xrightarrow{a.s.}$ stands for “converges almost surely to.”

Lemma 2.1 *Let $\{e_t\}$ and $\{\hat{e}_t\}$ be generated by the iteration (2.2) and (2.5), using the same trajectory of states $\{S_t\}$ and initial y_0 , but with different initial e_0 and \hat{e}_0 , respectively. Then under Conditions 2.1(i) and 2.3(i), $e_t - \hat{e}_t \xrightarrow{a.s.} 0$.*

Proof The proof is similar to that of (Yu, 2012, Lemma 3.2). Let $\Delta_t = \|e_t - \hat{e}_t\|$, and let \mathcal{F}_t denote the σ -algebra generated by $S_k, k \leq t$. Note that under our assumption, in the generation of the two trace sequences $\{e_t\}$ and $\{\hat{e}_t\}$, the states $\{S_t\}$ and the memory states $\{y_t\}$ are the same, but the λ -parameters are different. Let us denote them by $\{\lambda_t\}$ and $\{\hat{\lambda}_t\}$ for the two trace sequences, respectively. Then by (2.2), $e_t - \hat{e}_t = \gamma_t \rho_{t-1} (\lambda_t e_{t-1} - \hat{\lambda}_t \hat{e}_{t-1})$, and by Condition 2.3(i), $\|\lambda_t e_{t-1} - \hat{\lambda}_t \hat{e}_{t-1}\| \leq \|e_{t-1} - \hat{e}_{t-1}\|$. Hence $\|e_t - \hat{e}_t\| \leq \gamma_t \rho_{t-1} \|e_{t-1} - \hat{e}_{t-1}\|$, so $\mathbb{E}[\Delta_t | \mathcal{F}_{t-1}] \leq \mathbb{E}[\gamma_t \rho_{t-1} | \mathcal{F}_{t-1}] \cdot \Delta_{t-1} \leq \Delta_{t-1}$. This shows $\{(\Delta_t, \mathcal{F}_t)\}$ is a nonnegative supermartingale. By the supermartingale convergence theorem (Dudley, 2002, Theorem 10.5.7 and Lemma 4.3.3), $\{\Delta_t\}$ converges a.s. to a nonnegative random variable Δ_∞ with $\mathbb{E}[\Delta_\infty] \leq \liminf_{t \rightarrow \infty} \mathbb{E}[\Delta_t]$. From the inequality $\|e_t - \hat{e}_t\| \leq \gamma_t \rho_{t-1} \|e_{t-1} - \hat{e}_{t-1}\|$ for all t , we have $\Delta_t \leq \Delta_0 \cdot \prod_{k=1}^t \gamma_k \rho_{k-1}$, from which a direct calculation shows $\mathbb{E}[\Delta_t] \leq \Delta_0 \cdot \mathbf{1}^\top (P\Gamma)^t \mathbf{1}$

5. This means that for any bounded continuous function f on $\mathcal{S} \times \mathcal{M} \times \mathbb{R}^n$ (endowed with the usual topology), with $X_t = (S_t, y_t, e_t)$, $\mathbb{E}[f(X_1) | X_0 = x]$ is a continuous function of x (Meyn and Tweedie, 2009, Prop. 6.1.1).

where $\mathbf{1}$ denotes the n -dimensional vector of all 1's. As $t \rightarrow \infty$, $(P\Gamma)^t$ converges to the zero matrix under Condition 2.1(i). Therefore, $\liminf_{t \rightarrow \infty} \mathbb{E}[\Delta_t] = 0$ and consequently, we must have $\Delta_\infty = 0$ a.s., i.e., $\Delta_t \xrightarrow{a.s.} 0$. \blacksquare

We use Lemma 2.1 and ergodicity properties of weak Feller Markov chains (Meyn, 1989) to prove the ergodicity theorem below. A direct application to LSTD will be discussed immediately after the theorem, before we give its proof.

To state the result, we need some terminology and notation. For $\{(S_t, y_t, e_t)\}$ starting from the initial condition $x = (s, y, e)$, we write \mathbf{P}_x for its probability distribution, and we write “ \mathbf{P}_x -a.s.” for “almost surely with respect to \mathbf{P}_x .” The *occupation probability measures* are denoted by $\{\mu_{x,t}\}$, and they are random probability measures on $\mathcal{S} \times \mathcal{M} \times \mathbb{R}^n$ given by

$$\mu_{x,t}(D) := \frac{1}{t} \sum_{k=0}^{t-1} \mathbb{1}((S_k, y_k, e_k) \in D) \quad \forall \text{ Borel sets } D \subset \mathcal{S} \times \mathcal{M} \times \mathbb{R}^n,$$

where $\mathbb{1}(\cdot)$ is the indicator function. We are interested in the asymptotic convergence of these occupation probability measures in the sense of *weak convergence*: for probability measures $\{\mu_t\}$ and μ on a metric space, $\{\mu_t\}$ converges weakly to μ if $\int f d\mu_t \rightarrow \int f d\mu$ as $t \rightarrow \infty$, for every bounded continuous function f .

We shall also consider the Markov chain $\{(S_t, S_{t+1}, y_t, e_t)\}$, whose occupation probability measures are defined likewise. This Markov chain is essentially the same as $\{(S_t, y_t, e_t)\}$, but it is more convenient for applying our ergodicity result to TD algorithms because the temporal-difference term $\delta_t(v)$ involves (S_t, S_{t+1}, e_t) . Regarding invariant probability measures of the two Markov chains, obviously, if ζ is an invariant probability measure of $\{(S_t, y_t, e_t)\}$, then an invariant probability measure of $\{(S_t, S_{t+1}, y_t, e_t)\}$ is the probability measure ζ_1 composed from the marginal ζ and the conditional distribution of S_1 given (S_0, y_0, e_0) specified by P^o ; i.e.,

$$\zeta_1(D) = \int \sum_{s' \in \mathcal{S}} P_{ss'}^o \mathbb{1}((s, s', y, e) \in D) \zeta(d(s, y, e)) \quad \forall \text{ Borel sets } D \subset \mathcal{S}^2 \times \mathcal{M} \times \mathbb{R}^n. \quad (2.13)$$

(In the above, we used the notation $\int f(x) \zeta(dx)$ to write the integral of f w.r.t. ζ , and the notation $\zeta(d(s, y, e))$ is the same as $\zeta(dx)$ with $x = (s, y, e)$.)

Theorem 2.1 *Let Conditions 2.1-2.3 hold. Then $\{(S_t, y_t, e_t)\}$ is a weak Feller Markov chain and has a unique invariant probability measure ζ . For each initial condition $x := (s, y, e)$ of (S_0, y_0, e_0) , the occupation probability measures $\{\mu_{x,t}\}$ converge weakly to ζ , \mathbf{P}_x -a.s.*

Likewise, the same holds for $\{(S_t, S_{t+1}, y_t, e_t)\}$, whose unique invariant probability measure is as given in (2.13).

If the initial distribution of (S_0, y_0, e_0) is ζ , the state-trace process $\{(S_t, y_t, e_t)\}$ is stationary. Let \mathbb{E}_ζ denote expectation w.r.t. this stationary process. We now state a corollary of the above theorem for LSTD, before we prove the theorem.

Consider the sequence of equations in v , $\frac{1}{t} \sum_{k=0}^{t-1} e_k \delta_k(v) = 0$, appeared in (2.4) for LSTD. From the definition (2.3) of $\delta_t(v)$,

$$\delta_t(v) = \rho_t (R_t + \gamma_{t+1} v(S_{t+1}) - v(S_t)),$$

we see that for fixed v , every $e_k \delta_k(v)$ can be expressed as $f(S_k, S_{k+1}, e_k)$ for a continuous function f . Since the traces and hence the entire process lie in a bounded set under Condition 2.3(ii), the weak convergence of the occupation probabilities measures of $\{(S_t, S_{t+1}, y_t, e_t)\}$ shown by Theorem 2.1 implies that this sequence of equations has an asymptotic limit that can be expressed in terms of the stationary state-trace process as follows.

Corollary 2.1 *Let Conditions 2.1-2.3 hold. Then for each initial condition of (S_0, y_0, e_0) , almost surely, the sequence of linear equations in v , $\frac{1}{t} \sum_{k=0}^{t-1} e_k \delta_k(v) = 0$, tends asymptotically to $\mathbb{E}_\zeta[e_0 \delta_0(v)] = 0$ (also a linear equation in v), in the sense that the random coefficients in the former equations converge to the corresponding coefficients in the latter equation as $t \rightarrow \infty$.*

In the rest of this section we prove Theorem 2.1. Broadly speaking, the line of argument is as follows: We first prove the weak convergence of occupation probability measures to the same invariant probability measure, for each initial condition. This will in turn imply the uniqueness of the invariant probability measure.

After the proof we will first comment in Remark 2.1 on the differences between our proof and that of a similar result in the previous work (Yu, 2012). We will then comment in Remark 2.2 about the technical roles of Condition 2.3 (which concerns the choice of the function $\lambda(\cdot)$) and whether some part of that condition can be relaxed.

Proof of Theorem 2.1 As we discussed before Lemma 2.1, under Conditions 2.3, $\{(S_t, y_t, e_t)\}$ is weak Feller and has at least one invariant probability measure ζ . Then, by (Meyn, 1989, Prop. 4.1), there exists a set $D \subset \mathcal{S} \times \mathcal{M} \times \mathbb{R}^n$ with ζ -measure 1 such that for each initial condition $x = (s, y, e) \in D$, the occupation probability measures $\{\mu_{x,t}\}$ converge weakly, \mathbf{P}_x -a.s., to an invariant probability measure μ_x that depends only on the initial condition x . To prove the theorem using this result, we need to show that (i) all these $\{\mu_x \mid x \in D\}$ are the same invariant probability measure, and (ii) for all $x \notin D$, $\{\mu_{x,t}\}$ has the same weak convergence property.

To this end, we first consider an arbitrary pair (s, y_s) in the recurrent class of $\{(S_t, y_t)\}$ (cf. Condition 2.2). Let us show that for all initial conditions $x \in \{(s, y_s, e) \mid e \in \mathbb{R}^n\}$, $\{\mu_{x,t}\}$ converges weakly to the same invariant probability measure, almost surely.

Since the finite-state Markov chain $\{(S_t, y_t)\}$ has a single recurrent class (Condition 2.2) and its evolution is not affected by $\{e_t\}$, the marginal of ζ on $\mathcal{S} \times \mathcal{M}$ coincides with the unique invariant probability distribution of $\{(S_t, y_t)\}$. So the fact that $\zeta(D) = 1$ and (s, y_s) is a recurrent state of $\{(S_t, y_t)\}$ implies that there exists some \hat{e} with $(s, y_s, \hat{e}) \in D$. For the initial condition $\hat{x} = (s, y_s, \hat{e})$, by the result of (Meyn, 1989) mentioned earlier, $\{\mu_{\hat{x},t}\}$ converges weakly to $\mu_{\hat{x}}$, almost surely.

Now consider $x = (s, y_s, e)$ for an arbitrary $e \in \mathbb{R}^n$. Generate iterates $\{\hat{e}_t\}$ and $\{e_t\}$ according to (2.2), using the same trajectory $\{(S_t, y_t)\}$ with $(S_0, y_0) = (s, y_s)$, but with $\hat{e}_0 = \hat{e}$ and $e_0 = e$. By Lemma 2.1, $\hat{e}_t - e_t \xrightarrow{a.s.} 0$. Therefore, except on a null set of sample

paths, it holds for all bounded Lipschitz continuous functions f on $\mathcal{S} \times \mathcal{M} \times \mathbb{R}^n$ that⁶

$$\left| \int f d\mu_{\hat{x},t} - \int f d\mu_{x,t} \right| = \left| \frac{1}{t} \sum_{k=0}^{t-1} f(S_k, y_k, \hat{e}_k) - \frac{1}{t} \sum_{k=0}^{t-1} f(S_k, y_k, e_k) \right| \rightarrow 0. \quad (2.14)$$

By the a.s. weak convergence of $\mu_{\hat{x},t}$ to $\mu_{\hat{x}}$ proved earlier, except on a null set, $\int f d\mu_{\hat{x},t} \rightarrow \int f d\mu_{\hat{x}}$ for all such functions f . Combining this with (2.14) yields that almost surely, $\int f d\mu_{x,t} \rightarrow \int f d\mu_{\hat{x}}$ for all such f . By (Dudley, 2002, Theorem 11.3.3), this implies that almost surely, $\mu_{x,t} \rightarrow \mu_{\hat{x}}$ weakly.

Thus we have proved that for all initial conditions $x = (s, y_s, e)$, $e \in \mathbb{R}^n$, $\{\mu_{x,t}\}$ converges weakly, almost surely, to the same invariant probability measure $\mu_{\hat{x}}$. Denote $\mu = \mu_{\hat{x}}$. Let us now show that for any initial condition x , $\{\mu_{x,t}\}$ also converges to μ , \mathbf{P}_x -a.s.

Consider $\{(S_t, y_t, e_t)\}$ with an arbitrary initial condition $\bar{x} = (\bar{s}, \bar{y}, \bar{e})$. Let $\tau = \min\{t \mid (S_t, y_t) = (s, y_s)\}$ (the pair (s, y_s) is as in the proof above). Note that $\tau < \infty$ a.s., because (s, y_s) is a recurrent state of $\{(S_t, y_t)\}$. Define $(\tilde{S}_k, \tilde{y}_k) = (S_{\tau+k}, y_{\tau+k})$, $\tilde{e}_k = e_{\tau+k}$ for $k \geq 0$.

By the strong Markov property (see e.g. Nummelin, 1984, Theorem 3.3), $\{(\tilde{S}_k, \tilde{y}_k)\}_{k \geq 0}$ has the same probability distribution as the Markov chain $\{(S_t, y_t)\}$ that starts from $(S_0, y_0) = (s, y_s)$. Therefore, by the preceding proof, $\mathbf{P}_{\bar{x}}$ -almost surely, for all bounded continuous functions f on $\mathcal{S} \times \mathcal{M} \times \mathbb{R}^n$,

$$\lim_{m \rightarrow \infty} \frac{1}{m} \sum_{k=0}^{m-1} f(\tilde{S}_k, \tilde{y}_k, \tilde{e}_k) = \int f d\mu. \quad (2.15)$$

Denote $a \wedge b = \min\{a, b\}$. Using (2.15) and the fact $\tau < \infty$ a.s., we have that $\mathbf{P}_{\bar{x}}$ -almost surely,

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{k=0}^{t-1} f(S_k, y_k, e_k) &= \lim_{t \rightarrow \infty} \left(\frac{1}{t} \sum_{k=0}^{t \wedge (\tau-1)} f(S_k, y_k, e_k) + \frac{1}{t} \sum_{k=\tau}^{t-1} f(S_k, y_k, e_k) \right) \\ &= \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{k=0}^{t-\tau-1} f(S_{\tau+k}, y_{\tau+k}, e_{\tau+k}) \\ &= \lim_{m \rightarrow \infty} \frac{1}{m} \sum_{k=0}^{m-1} f(\tilde{S}_k, \tilde{y}_k, \tilde{e}_k) = \int f d\mu. \end{aligned}$$

This proves that $\{\mu_{x,t}\}$ converges weakly to μ almost surely, for each initial condition x .

It now follows that μ must be the unique invariant probability measure of $\{(S_t, y_t, e_t)\}$. To see this, suppose ζ is another invariant probability measure. For any bounded continuous function f , by stationarity, $\mathbb{E}_{\zeta}[\frac{1}{t} \sum_{k=0}^{t-1} f(S_k, y_k, e_k)] = \int f d\zeta$ for all $t \geq 1$. On the other hand, the preceding proof has established that for all initial conditions x ,

$$\frac{1}{t} \sum_{k=0}^{t-1} f(S_k, y_k, e_k) = \int f d\mu_{x,t} \rightarrow \int f d\mu, \quad \mathbf{P}_x\text{-a.s.},$$

which implies that if ζ is the initial distribution of (S_0, y_0, e_0) , then $\frac{1}{t} \sum_{k=0}^{t-1} f(S_k, y_k, e_k) \rightarrow \int f d\mu$, \mathbf{P}_{ζ} -a.s. We thus have

$$\begin{aligned} \int f d\zeta &= \mathbb{E}_{\zeta} \left[\frac{1}{t} \sum_{k=0}^{t-1} f(S_k, y_k, e_k) \right] = \lim_{t \rightarrow \infty} \mathbb{E}_{\zeta} \left[\frac{1}{t} \sum_{k=0}^{t-1} f(S_k, y_k, e_k) \right] \\ &= \mathbb{E}_{\zeta} \left[\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{k=0}^{t-1} f(S_k, y_k, e_k) \right] = \int f d\mu, \end{aligned}$$

6. Here we are using the same (S_k, y_k) , $k \leq t$ in the occupation probability measures $\mu_{\hat{x},t}$ and $\mu_{x,t}$. This is valid because the e_t 's do not affect the evolution of $\{(S_t, y_t)\}$ and are functions of these states and the given initial e_0 . If we call the $\mu_{x,t}$ here $\tilde{\mu}_{x,t}$ instead and define $\mu_{x,t}$ using another independent copy of $\{(S_t, y_t)\}$, then since the two sequences of occupation probability measures will have the same probability distribution, $\{\mu_{x,t}\}$ will have the same weak convergence property as $\{\tilde{\mu}_{x,t}\}$.

where the third equality follows from the bounded convergence theorem. This shows $\int f d\zeta = \int f d\mu$ for all bounded continuous functions f , and hence $\zeta = \mu$ by (Dudley, 2002, Prop. 11.3.2), proving the uniqueness of the invariant probability measure.

The conclusions for the Markov chain $\{(S_t, S_{t+1}, y_t, e_t)\}$ follow from the same arguments given above, if we replace S_t with (S_t, S_{t+1}) and replace the set \mathcal{S} with the set of possible state transitions. (We could have proved the assertions for $\{(S_t, S_{t+1}, y_t, e_t)\}$ first and then deduced as their implications the assertions for $\{(S_t, y_t, e_t)\}$. We treated the latter first, as it makes the notation in the proof simpler.) ■

Remark 2.1 (About the proof) Theorem 2.1 is similar to (Yu, 2012, Theorem 3.2) for off-policy LSTD with constant λ (the analysis given in (Yu, 2012) also applies to state-dependent λ). Some of the techniques used to prove the two theorems are also similar. The main difference to (Yu, 2012) is that in the proof here we used an argument based on the strong Markov property to extend the weak convergence property of $\{\mu_{x,t}\}$ for a subset of initial conditions $x \in \{(s, y_s, e) \mid e \in \mathbb{R}^n\}$ to all initial conditions, whereas in (Yu, 2012) this step was proved using a result on the convergence-in-mean of LSTD iterates established first. The latter approach would not work here due to the dependence of λ_t on the history. Indeed, due to this dependence, the proof of the convergence-in-mean of LSTD given in (Yu, 2012) does not carry over to our case, even though that convergence does hold as a consequence of Theorem 2.1, in view of the boundedness of traces by construction. Compared with the proof of the ergodicity result in (Yu, 2012), the proof we gave here is more direct and therefore better.

Regarding possible alternative proofs of Theorem 2.1, let us also mention that if we prove first the uniqueness of the invariant probability measure, then, since $\{(S_t, y_t, e_t)\}_{t \geq 1}$ lie in a bounded set, the weak convergence of occupation probability measures will follow immediately from (Meyn, 1989, Prop. 4.2). However, because the evolution of the λ_t 's depends on both states and traces, it does not seem easy to us to prove directly the uniqueness part first. ■

Remark 2.2 (About the conditions on the function $\lambda(\cdot)$) Our proof of Theorem 2.1 relied on Lemma 2.1 and the two properties discussed preceding that lemma, namely, that $\{(S_t, y_t, e_t)\}$ is a weak Feller Markov chain and has at least one invariant probability measure. As long as these hold when we weaken or change the conditions on the function $\lambda(\cdot)$, the proof and the conclusions of the theorem will remain applicable.

We introduced Condition 2.3(ii) to bound the traces for algorithmic concerns. For the ergodicity of the state-trace process, Condition 2.3(ii) is unimportant—in fact, it can be removed from the conditions of Theorem 2.1. The reason is that we used this condition before Lemma 2.1 to quickly infer that $\{(S_t, y_t, e_t)\}$ has at least one invariant probability measure, but this is still true without Condition 2.3(ii), in view of (Meyn and Tweedie, 2009, Theorem 12.1.2(ii)) and the fact that under Condition 2.1(i), $\{e_t\}$ is bounded in probability (the proof of this fact is straightforward and similar to the proof of (Yu, 2012, Lemma 3.1) or (Yu, 2015, Prop. A.1)).

Condition 2.3(i) is actually two conditions combined into one. The first is the continuity of $\lambda(y, e)e$ in e for each y , which was used to ensure that the state-trace process is a weak

Feller Markov chain. To be more general, instead of letting the evolutions of the traces and memory states be governed by the functions λ and g , one may consider letting them be governed by stochastic kernels. Then by placing a suitable continuity condition on the stochastic kernel λ , one can ensure that the state-trace process has the desired weak Feller Markov property.

The second condition packed into Condition 2.3(i) is that for each y , $\lambda(y, e)e$ is a Lipschitz continuous function of e with modulus 1. This condition is somewhat restrictive, and one may consider instead allowing the function to have Lipschitz modulus greater than 1. However, additional conditions are then needed to ensure that Lemma 2.1 holds. (If this lemma does not hold, then the state-trace process may not be ergodic and one will need a different approach than the one we took to characterize the sample path properties of the state-trace process.)

From an algorithmic perspective, if it is desirable to choose even larger λ_t 's or to have greater flexibility in choosing these λ -parameters, some of the generalizations just mentioned can be considered. For example, Condition 2.3(ii) can be replaced and stochastic kernels can be introduced to allow for occasionally large traces e_t , so that instead of having the traces bounded, one only make their variances bounded in a desired range. ■

3. Generalized Bellman Equations

In this section, we continue the analysis started in Section 2.3. Recall that Corollary 2.1 established that the asymptotic limit of the linear equations (2.4) for LSTD is the linear equation (in v):

$$\mathbb{E}_\zeta[e_0 \delta_0(v)] = 0.$$

Our goal now is to relate this equation to a generalized Bellman equation for the target policy π . This will then allow us to interpret solutions of (2.4) computed by LSTD as solutions of approximate versions of that generalized Bellman equation.

To this end, we will first give a general description of randomized stopping times and associated Bellman operators (Section 3.1). We will then use these notions to derive the particular Bellman operators that correspond to our choices of the λ -parameters and appear in the linear equations for LSTD (Section 3.2). We will also discuss a composite scheme of choosing the λ -parameters as a direct application and extension of our results.

To simplify notation in subsequent derivations, we shall use the following shorthand notation: For $k \leq m$, denote $S_k^m = (S_k, S_{k+1}, \dots, S_m)$,

$$\rho_k^m = \prod_{i=k}^m \rho_i, \quad \lambda_k^m = \prod_{i=k}^m \lambda_i, \quad \gamma_k^m = \prod_{i=k}^m \gamma_i. \quad (3.1)$$

Also, we shall treat $\rho_k^m = \lambda_k^m = \gamma_k^m = 1$ if $k > m$.

3.1 Randomized Stopping Times and Associated Bellman Operators

Consider the Markov chain $\{S_t\}$ induced by the target policy π . Let Condition 2.1(i) hold. Recall that for the value function v_π , we have that for each state $s \in \mathcal{S}$,

$$v_\pi(s) = \mathbb{E}_s^\pi \left[\sum_{t=0}^{\infty} \gamma_1^t r_\pi(S_t) \right] \quad (\text{by definition})$$

and

$$v_\pi(s) = r_\pi(s) + \mathbb{E}_s^\pi[\gamma_1 v_\pi(S_1)].$$

The second equation is the standard one-step Bellman equation.

To write generalized Bellman equations for π , we shall make use of *randomized stopping times* for $\{S_t\}$, a notion that generalizes naturally stopping times for $\{S_t\}$ in that whether to stop at time t depends not only on the past states S_0^t but also on certain random outcomes. A simple example is to toss a coin at each time and stop as soon as the coin lands on heads, regardless of the history S_0^t . (The corresponding Bellman equation is the one associated with TD(λ) for a constant λ ; cf. Example 3.1.) Of interest here is the general case where the stopping decision does depend on the entire history.

To define a randomized stopping time formally, first, the probability space of $\{S_t\}$ is enlarged to take into account whatever randomization scheme that is used to make the stopping decision. (The enlargement will be problem-dependent, as the next subsection will demonstrate.) Then, on the enlarged space, a randomized stopping time τ for $\{S_t\}$ is a stopping time⁷ relative to some increasing sequence of σ -algebras $\mathcal{F}_0 \subset \mathcal{F}_1 \subset \dots$, where the sequence $\{\mathcal{F}_t\}$ is such that

- (i) for all $t \geq 0$, $\mathcal{F}_t \supset \sigma(S_0^t)$ (the σ -algebra generated by S_0^t), and
- (ii) relative to $\{\mathcal{F}_t\}$, $\{S_t\}$ remains to be a Markov chain with transition probability P , i.e., for all $s \in \mathcal{S}$, $\text{Prob}(S_{t+1} = s \mid \mathcal{F}_t) = P_{S_t s}$.

See (Nummelin, 1984, Chap. 3.3); in particular, see Prop. 3.6 in p. 31-32 therein for several equivalent definitions of randomized stopping times.

Note that if $\mathcal{F}_t = \sigma(S_0^t)$ for all t , then the history of states S_0^t fully determines whether $\tau \leq t$ and τ reduces to a stopping time for the Markov chain $\{S_t\}$. The properties (i)-(ii) in the above definition encapsulate our earlier intuitive discussion about making stopping decisions, namely, stopping decisions are made based on the history S_0^t and additional random outcomes that do not affect the evolution of the Markov chain.

Like stopping times, the strong Markov property also holds for randomized stopping times for a Markov chain. This is an important basic property. It says that in the event $\tau < \infty$, conditioned on the σ -algebra \mathcal{F}_τ associated with the stopping time τ relative to $\{\mathcal{F}_t\}$ (which is the σ -algebra generated by the events that “happen before τ ”), the conditional distribution of $(S_\tau, S_{\tau+1}, \dots)$ is the same as the probability distribution of a Markov chain (S_0, S_1, \dots) with initial state $S_0 = S_\tau$ (Nummelin, 1984, Theorem 3.3).

The above abstract definition of a randomized stopping time allows us to write Bellman equations in general forms without worrying about the details of the enlarged space, which are not important at this point. For notational simplicity, when there is no confusion, we shall still write \mathbf{P}^π for the probability measure on the enlarged probability space and use \mathbb{E}^π and \mathbb{E}_s^π to denote the expectation and conditional expectation given $S_0 = s$, respectively, w.r.t. \mathbf{P}^π .

If τ is a randomized stopping time for $\{S_t\}$, the strong Markov property (Nummelin, 1984, Theorem 3.3) allows us to express v_π in terms of $v_\pi(S_\tau)$ and the total discounted

7. A random time τ is called a stopping time relative to a sequence $\{\mathcal{F}_t\}$ of increasing σ -algebras if the event $\{\tau = t\} \in \mathcal{F}_t$ for every t .

rewards R^τ prior to stopping:

$$\begin{aligned} v_\pi(s) &= \mathbb{E}_s^\pi \left[\sum_{t=0}^{\tau-1} \gamma_1^t r_\pi(S_t) + \sum_{t=\tau}^\infty \gamma_1^\tau \cdot \gamma_{\tau+1}^t r_\pi(S_t) \right] \\ &= \mathbb{E}_s^\pi [R^\tau + \gamma_1^\tau v_\pi(S_\tau)], \end{aligned} \quad (3.2)$$

where $R^\tau = \sum_{t=0}^{\tau-1} \gamma_1^t r_\pi(S_t)$ for $\tau \in \{0, 1, 2, \dots\} \cup \{+\infty\}$.⁸ We can also write the Bellman equation (3.2) in terms of $\{S_t\}$ only, by taking expectation over τ :

$$\begin{aligned} v_\pi(s) &= \mathbb{E}_s^\pi \left[\sum_{t=0}^\infty \left(\mathbb{1}(\tau > t) \cdot \gamma_1^t r_\pi(S_t) + \mathbb{1}(\tau = t) \cdot \gamma_1^t v_\pi(S_t) \right) \right], \\ &= \mathbb{E}_s^\pi \left[\sum_{t=0}^\infty \left(q_t^+(S_0^t) \cdot \gamma_1^t r_\pi(S_t) + q_t(S_0^t) \cdot \gamma_1^t v_\pi(S_t) \right) \right], \end{aligned} \quad (3.3)$$

where

$$q_t^+(S_0^t) = \mathbf{P}^\pi(\tau > t \mid S_0^t), \quad q_t(S_0^t) = \mathbf{P}^\pi(\tau = t \mid S_0^t). \quad (3.4)$$

The r.h.s. of (3.2) or (3.3) defines a generalized Bellman operator $T : \mathfrak{R}^N \rightarrow \mathfrak{R}^N$ associated with τ , which has several equivalent expressions; e.g.,

$$(Tv)(s) = \mathbb{E}_s^\pi [R^\tau + \gamma_1^\tau v(S_\tau)] = \mathbb{E}_s^\pi \left[\sum_{t=0}^\infty \left(q_t^+(S_0^t) \cdot \gamma_1^t r_\pi(S_t) + q_t(S_0^t) \cdot \gamma_1^t v(S_t) \right) \right], \quad s \in \mathcal{S}.$$

Depending on the context, one expression can be more convenient to use than the other. For example, the first expression is convenient for defining T through the associated τ and for deducing the contraction property of T , whereas expressions like the second will be of interest when we want to know more explicitly the particular T for our TD learning scheme and its dependence on the λ -parameters.

In common with one-step Bellman operator, the generalized Bellman operator T is affine and involves a substochastic matrix. If $\tau \geq 1$ a.s., then the value function v_π is the unique fixed point of T , i.e., the unique solution of $v = Tv$, and T is a sup-norm contraction. In fact, this can be shown for slightly more general τ :

Theorem 3.1 *Let Condition 2.1(i) hold, and let the randomized stopping time τ be such that $\mathbf{P}^\pi(\tau \geq 1 \mid S_0 = s) > 0$ for all states $s \in \mathcal{S}$. Then v_π is the unique fixed point of the generalized Bellman operator T associated with τ , and T is a contraction w.r.t. a weighted sup-norm on \mathfrak{R}^N .*

8. We explain the derivations in this footnote. In the case $\tau = 0$, $R^0 = 0$. In the case $\tau = \infty$, by Condition 2.1(i), $R^\infty = \sum_{t=0}^\infty \gamma_1^t r_\pi(S_t)$ is almost surely well-defined, while the second term $\gamma_1^\tau v_\pi(S_\tau)$ in (3.2) is 0 because $\gamma_1^\infty := \prod_{k=1}^\infty \gamma_k = 0$ a.s., under Condition 2.1(i). Equation (3.2) is derived as follows: By the strong Markov property (Nummelin, 1984, Theorem 3.3), on $\{\tau < \infty\}$,

$$\mathbb{E}^\pi \left[\sum_{t=\tau}^\infty \gamma_1^\tau \cdot \gamma_{\tau+1}^t r_\pi(S_t) \mid \mathcal{F}_\tau \right] = \gamma_1^\tau \cdot \mathbb{E}_{S_\tau}^\pi \left[\sum_{t=0}^\infty \gamma_1^t r_\pi(S_t) \right] = \gamma_1^\tau v_\pi(S_\tau).$$

Then, since the term $\mathbb{E}_s^\pi \left[\sum_{t=\tau}^\infty \gamma_1^\tau \cdot \gamma_{\tau+1}^t r_\pi(S_t) \right] = \mathbb{E}_s^\pi \left[\mathbb{1}(\tau < \infty) \cdot \sum_{t=\tau}^\infty \gamma_1^\tau \cdot \gamma_{\tau+1}^t r_\pi(S_t) \right]$, we use the property of the conditional expectation given \mathcal{F}_τ and the fact $\mathcal{F}_\tau \supset \sigma(S_0)$ to rewrite this term as

$$\mathbb{E}_s^\pi \left[\mathbb{1}(\tau < \infty) \cdot \mathbb{E}^\pi \left[\sum_{t=\tau}^\infty \gamma_1^\tau \cdot \gamma_{\tau+1}^t r_\pi(S_t) \mid \mathcal{F}_\tau \right] \right] = \mathbb{E}_s^\pi \left[\mathbb{1}(\tau < \infty) \cdot \gamma_1^\tau v_\pi(S_\tau) \right] = \mathbb{E}_s^\pi \left[\gamma_1^\tau v_\pi(S_\tau) \right],$$

where in the last equality we also used the fact $\gamma_1^\infty = 0$ a.s. This gives (3.2).

We prove this theorem in Appendix A. The proof amounts to showing that if a state process evolves according to the substochastic matrix \tilde{P} involved in the affine operator T , then all the states in \mathcal{S} are transient (equivalently, the spectral radius of \tilde{P} is less than 1 and $I - \tilde{P}$ is invertible (Puterman, 1994, Appendix A.4)). From this the conclusions of the theorem follow as a basic fact from nonnegative matrix theory (Seneta, 2006, Theorem 1.1), and one specific choice of the weights of the sup-norm in the theorem is simply the expected time for the process to leave \mathcal{S} from each initial state (see e.g., the proof of (Bertsekas and Tsitsiklis, 1996, Prop. 2.2)).

For TD algorithms that do not use history-dependent λ , the random times τ and the corresponding Bellman operators T have simple descriptions:

Example 3.1 (TD with constant or state-dependent λ) Depending on the choice of λ , TD(λ) algorithms are associated with different randomized stopping times τ . In the case of constant λ , starting from time 1, we stop the system with probability $1 - \lambda$ if it has not stopped yet; i.e.,

$$\tau \geq 1 \quad \text{and} \quad \mathbf{P}^\pi(\tau = t \mid \tau > t-1, S_0^t) = 1 - \lambda, \quad \forall t \geq 1.$$

In particular, we always stop at $t = 1$ if $\lambda = 0$, and we never stop if $\lambda = 1$. Similarly, for state-dependent λ where $\lambda_t = \lambda(S_t)$, a function of the current state, the preceding stopping probability is replaced by $1 - \lambda(S_t)$: $\mathbf{P}^\pi(\tau = t \mid \tau > t-1, S_0^t) = 1 - \lambda(S_t)$ for $t \geq 1$. In these cases, by taking expectations over τ , the corresponding Bellman operators can be expressed solely in terms of λ and the model parameters for the target policy. ■

3.2 Bellman Equation for the Proposed TD Learning Scheme

With the terminology of randomized stopping times, we are now ready to write down the generalized Bellman equation associated with the TD learning scheme proposed in Section 2.2. It corresponds to a particular randomized stopping time. We shall first describe this random time, from which a generalized Bellman equation follows as seen in the preceding subsection. That this is indeed the Bellman equation for our TD learning scheme will then be proved.

Consider the Markov chain $\{S_t\}$ under the target policy π . We define a randomized stopping time τ for $\{S_t\}$:

- Let $y_t, \lambda_t, e_t, t \geq 1$, evolve according to (2.5) and (2.2):

$$y_t = g(y_{t-1}, S_t), \quad \lambda_t = \lambda(y_t, e_{t-1}), \quad e_t = \lambda_t \gamma_t \rho_{t-1} e_{t-1} + \phi(S_t), \quad t \geq 1.$$

- Let the initial (S_0, y_0, e_0) be distributed according to ζ , the unique invariant probability measure in Theorem 2.1 for the state-trace process induced by the behavior policy.
- At time $t \geq 1$, we stop the system with probability $1 - \lambda_t$ if it has not yet been stopped. Let τ be the time when the system stops ($\tau = \infty$ if the system never stops).

To make the dependence on the initial distribution ζ explicit, we write \mathbf{P}_ζ^π for the probability measure of this process.

Note that by definition λ_t and $\lambda_1^t = \prod_{k=1}^t \lambda_k$ are functions of the initial (y_0, e_0) and states S_0^t . From how the random time τ is defined, we have for all $t \geq 1$,

$$\mathbf{P}_\zeta^\pi(\tau > t \mid S_0^t, y_0, e_0) = \lambda_1^t =: h_t^+(y_0, e_0, S_0^t), \quad (3.5)$$

$$\mathbf{P}_\zeta^\pi(\tau = t \mid S_0^t, y_0, e_0) = \lambda_1^{t-1}(1 - \lambda_t) =: h_t(y_0, e_0, S_0^t), \quad (3.6)$$

and hence

$$q_t^+(S_0^t) := \mathbf{P}_\zeta^\pi(\tau > t \mid S_0^t) = \int h_t^+(y, e, S_0^t) \zeta(d(y, e) \mid S_0), \quad (3.7)$$

$$q_t(S_0^t) := \mathbf{P}_\zeta^\pi(\tau = t \mid S_0^t) = \int h_t(y, e, S_0^t) \zeta(d(y, e) \mid S_0), \quad (3.8)$$

where $\zeta(d(y, e) \mid s)$ is the conditional distribution of (y_0, e_0) given $S_0 = s$, w.r.t. the initial distribution ζ . As before, we can write the generalized Bellman operator T associated with τ in several equivalent forms. Let \mathbb{E}_ζ^π denote expectation under \mathbf{P}_ζ^π . Similarly to the derivation of (3.3), we can rewrite (3.2) in this case by taking expectation over τ conditioned on (S_0^t, y_0, e_0) to derive that for all $v : \mathcal{S} \rightarrow \mathbb{R}, s \in \mathcal{S}$,

$$(Tv)(s) = \mathbb{E}_\zeta^\pi \left[\sum_{t=0}^\infty \lambda_1^t \gamma_1^t r_\pi(S_t) + \sum_{t=1}^\infty \lambda_1^{t-1} (1 - \lambda_t) \gamma_1^t v(S_t) \mid S_0 = s \right]. \quad (3.9)$$

Or express T in the form of (3.3) by further integrating over (y_0, e_0) and using (3.7)-(3.8):

$$(Tv)(s) = \mathbb{E}_\zeta^\pi \left[\sum_{t=0}^\infty \left(q_t^+(S_0^t) \cdot \gamma_1^t r_\pi(S_t) + q_t(S_0^t) \cdot \gamma_1^t v(S_t) \right) \mid S_0 = s \right], \quad (3.10)$$

for all $v : \mathcal{S} \rightarrow \mathbb{R}, s \in \mathcal{S}$, where in the case $t = 0$, $q_0^+(S_0) = 1$ and $q_0(S_0) = 0$ since $\tau > 0$ by construction.

It will be useful later to express $TV - V$ in terms of temporal differences. From (3.9), by writing $\lambda_1^{t-1}(1 - \lambda_t)\gamma_1^t v(S_t) = \lambda_1^{t-1}\gamma_1^t v(S_t) - \lambda_1^t\gamma_1^t v(S_t)$ and rearranging terms, we have for all $v : \mathcal{S} \rightarrow \mathbb{R}, s \in \mathcal{S}$,

$$\begin{aligned} (Tv)(s) - v(s) &= \mathbb{E}_\zeta^\pi \left[\sum_{t=0}^\infty \lambda_1^t \gamma_1^t r_\pi(S_t) + \sum_{t=0}^\infty \lambda_1^t \gamma_1^{t+1} v(S_{t+1}) - \sum_{t=0}^\infty \lambda_1^t \gamma_1^t v(S_t) \mid S_0 = s \right] \\ &= \mathbb{E}_\zeta^\pi \left[\sum_{t=0}^\infty \lambda_1^t \gamma_1^t \cdot \left(r_\pi(S_t) + \gamma_{t+1} v(S_{t+1}) - v(S_t) \right) \mid S_0 = s \right]. \end{aligned} \quad (3.11)$$

In a similar way, from (3.10), we can write⁹

$$(Tv)(s) - v(s) = \mathbb{E}_\zeta^\pi \left[\sum_{t=0}^\infty q_t^+(S_0^t) \cdot \gamma_1^t \cdot \left(r_\pi(S_t) + \gamma_{t+1} v(S_{t+1}) - v(S_t) \right) \mid S_0 = s \right].$$

Remark 3.1 Comparing the two expressions (3.9) and (3.10) of T , we remark that the expression (3.9) reflects the role of the λ_t 's in determining the stopping time, whereas the expression (3.10), which has eliminated the auxiliary variables y_t and e_t , shows more clearly the dependence of the stopping time on the entire history S_0^t . It can also be seen, from the

9. Since τ is a randomized stopping time for the Markov chain $\{S_t\}$, we have $\mathbf{P}_\zeta^\pi(\tau > t \mid S_0^{t+1}) = \mathbf{P}_\zeta^\pi(\tau > t \mid S_0^t)$, so $\mathbf{P}_\zeta^\pi(\tau > t \mid S_0^t) - \mathbf{P}_\zeta^\pi(\tau = t+1 \mid S_0^{t+1}) = \mathbf{P}_\zeta^\pi(\tau > t+1 \mid S_0^{t+1})$, i.e., $q_t^+(S_0^t) - q_{t+1}^+(S_0^{t+1}) = q_{t+1}^+(S_0^{t+1})$. Thus we can write the term $q_t(S_0^t)$ in (3.10) for $t \geq 1$ as $q_{t-1}^+(S_0^{t-1}) - q_t^+(S_0^t)$, and the expression for $(Tv - v)(s)$ then follows by rearranging terms.

initial distribution ζ , the dependence of λ_t on the traces and the dependence of the traces on the function $\rho(\cdot)$ (which describes importance sampling ratios), that both the behavior policy and the choice of the feature representation assert a significant role in determining the Bellman operator T for the target policy. This is in contrast with off-policy TD learning that uses a constant λ , where the behavior policy and the approximation subspace affect only how one approximates the Bellman equation underlying TD, not the Bellman equation itself, which is solely determined by λ (cf. Example 3.1).

Furthermore, note that as the invariant distribution of the state-trace process, ζ is associated with the dynamic behavior of the states *and traces* under the behavior policy. Generally, there is no explicit expression of ζ in terms of P^o and the parameters in the λ function. As a result, in general we cannot express the operator T in terms of these parameters in the learning scheme. This is different from the case of TD(λ) where λ is a function of the present state only. ■

We now proceed to show how the Bellman equation $v = Tv$ given above relates to the off-policy TD learning scheme in Section 2.2. Some notation is needed. Denote by $\zeta_{\mathcal{S}}$ the invariant probability measure of the Markov chain $\{S_t\}$ *induced by the behavior policy*; note that it coincides with the marginal of ζ on \mathcal{S} . For two functions v_1, v_2 on \mathcal{S} , we write $v_1 \perp_{\zeta_{\mathcal{S}}} v_2$ if $\sum_{s \in \mathcal{S}} \zeta_{\mathcal{S}}(s) v_1(s) v_2(s) = 0$. If \mathcal{L} is a linear subspace of functions on \mathcal{S} and $v \perp_{\zeta_{\mathcal{S}}} v'$ for all $v' \in \mathcal{L}$, we write $v \perp_{\zeta_{\mathcal{S}}} \mathcal{L}$. Recall that ϕ is a function that maps each state s to an n -dimensional feature vector. Denote by \mathcal{L}_{ϕ} the subspace spanned by the n component functions of ϕ , which is the space of approximate value functions for our TD learning scheme. Recall also that \mathbb{E}_{ζ} denotes expectation w.r.t. the *stationary* state-trace process $\{(S_t, y_t, e_t)\}$ under the behavior policy (cf. Theorem 2.1).

Theorem 3.2 *Let Conditions 2.1-2.3 hold. Then as a linear equation in v , $\mathbb{E}_{\zeta}[e_0 \delta_0(v)] = 0$ is equivalently $Tv - v \perp_{\zeta_{\mathcal{S}}} \mathcal{L}_{\phi}$, where T is the generalized Bellman operator for π given in (3.9) or (3.10).*

Remark 3.2 (On LSTD) Note that

$$Tv - v \perp_{\zeta_{\mathcal{S}}} \mathcal{L}_{\phi}, \quad v \in \mathcal{L}_{\phi}$$

is a projected version of the generalized Bellman equation $Tv - v = 0$ (projecting the left-hand side onto the approximation subspace \mathcal{L}_{ϕ} w.r.t. the $\zeta_{\mathcal{S}}$ -weighted Euclidean norm). Theorem 3.2 and Corollary 2.1 together show that this is what LSTD solves in the limit.

Note also that although the generalized Bellman operator T is a contraction (Theorem 3.1), the composition of projection with T is in general not a contraction (cf. Example B.1 in Appendix B). Thus we cannot use contraction-based arguments to analyze approximation properties. For that purpose, we use the oblique projection viewpoint of Scherrer (2010). Specifically, if the preceding projected Bellman equation admits a unique solution \bar{v} , then \bar{v} can be viewed as an oblique projection of v_{π} (Scherrer, 2010) and the approximation error $\bar{v} - v_{\pi}$ can be characterized as in (Yu and Bertsekas, 2010) by using the oblique projection viewpoint. The details of these are given in Appendix B. ■

Remark 3.3 (On gradient-based TD) While Theorem 3.2 is about the LSTD algorithm, it also helps prepare the ground for analyzing gradient-based algorithms similar to

those discussed in (Maei, 2011; Mahadevan et al., 2014). Like LSTD, these algorithms aim to solve the same projected generalized Bellman equation as characterized by Theorem 3.2 (cf. Remark 3.2). Their average dynamics, which is important for analyzing their convergence using the mean ODE approach from stochastic approximation theory (Kushner and Yin, 2003), can be studied based on the ergodicity result of Theorem 2.1, in essentially the same way as we did in Section 2.3 for the LSTD algorithm. For details of the convergence analysis of these gradient-based TD algorithms, see the recent work (Yu, 2017). ■

In the rest of this subsection, we give a corollary to Theorem 3.2, deferring the proofs of both the theorem and the corollary to the next subsection. The corollary concerns a composite scheme of setting λ , which is slightly more general than what Section 2.2 described. It results in a Bellman operator that is a composition of the components of other Bellman operators, and it can be useful in practice for variance control. Let us describe the scheme first, before explaining our motivation for it.

Partition the state space into m nonempty disjoint sets: $\mathcal{S} = \cup_{i=1}^m \mathcal{S}_i$. Associate each set \mathcal{S}_i with a possibly different scheme of setting λ that is of the type described in Section 2.2, and denote its memory states by $y_t^{(i)}$ and λ -function by $\lambda^{(i)}(\cdot, \cdot)$. Keep m trace vectors $e_t^{(1)}, \dots, e_t^{(m)}$, one for each set, and update them according to

$$e_t^{(i)} = \lambda_t^{(i)} \gamma_t \rho_{t-1} e_{t-1}^{(i)} + \phi(S_t) \mathbb{1}(S_t \in \mathcal{S}_i), \quad 1 \leq i \leq m, \quad (3.12)$$

where $\lambda_t^{(i)} = \lambda^{(i)}(y_t^{(i)}, e_{t-1}^{(i)})$. We then have m ergodic state-trace processes that share the same state variables, $\{(S_t, y_t^{(i)}, e_t^{(i)})\}$, $i = 1, 2, \dots, m$. Each process has a unique invariant probability measure $\zeta^{(i)}$ (Theorem 2.1) and an associated randomized stopping time $\tau^{(i)}$ and generalized Bellman operator $T^{(i)}$, as discussed in this subsection. Define now an operator T by concatenating the component mappings of $T^{(i)}$ for \mathcal{S}_i as follows: for all $v \in \mathbb{R}^N$ and $s \in \mathcal{S}$,

$$(Tv)(s) := (T^{(i)}v)(s) \quad \text{if } s \in \mathcal{S}_i. \quad (3.13)$$

Consider an LSTD algorithm that defines the trace e_t to be the sum of the m trace vectors,

$$e_t = \sum_{i=1}^m e_t^{(i)}, \quad (3.14)$$

and uses the traces to form the linear equation as before,

$$\frac{1}{t} \sum_{k=0}^{t-1} e_k \delta_k(v) = 0, \quad v = \Phi\theta.$$

Note that $\frac{1}{t} \sum_{k=0}^{t-1} e_k \delta_k(v) = 0$ is the same as $\sum_{i=1}^m \frac{1}{t} \sum_{k=0}^{t-1} e_k^{(i)} \delta_k(v) = 0$. By Corollary 2.1, as a linear equation in v , it tends asymptotically (as $t \rightarrow \infty$) to the linear equation $\sum_{i=1}^m \mathbb{E}_{\zeta^{(i)}}[e_0^{(i)} \delta_0(v)] = 0$.

Corollary 3.1 *Let Condition 2.1 hold. Consider the composite scheme of setting λ discussed above, and let Conditions 2.2-2.3 hold for each of the m schemes involved. Let LSTD calculate traces according to (3.12) and (3.14). Then the limiting linear equation (in v) associated with LSTD, $\sum_{i=1}^m \mathbb{E}_{\zeta^{(i)}}[e_0^{(i)} \delta_0(v)] = 0$, is equivalently $Tv - v \perp_{\zeta_S} \mathcal{L}_\phi$, where T is the generalized Bellman operator for π given by (3.13) and has the same fixed point and contraction properties as stated in Theorem 3.1.*

The use of composite schemes will be demonstrated by experiments in Section 4.2.2. Here let us explain informally our motivation for such schemes.

Remark 3.4 (About composite schemes of setting λ) Our motivation for using the composite schemes is revealed by the equation (3.13). Typically each $T^{(i)}$ is designed to be simple to implement in TD learning. For example, if we ignore for now the bounding of traces introduced in Section 2.2 and just consider TD(λ) with constant λ , $T^{(i)}$ can be the Bellman operator $T^{(\lambda)}$ for TD(λ) with some constant λ . A simple, extreme example is to partition the state space into two sets, and associate one with $T^{(\lambda)}, \lambda = 1$, and the other with $T^{(\lambda)}, \lambda = 0$. Using the combination (3.13) of the two operators in TD then means that for the first set of states whose $\lambda = 1$, we want to estimate their values by using the information about the total rewards received when starting from those states, whereas for the second set of states whose $\lambda = 0$, we only use the information about their one-stage rewards and how these states relate to the “neighboring” states in the transition graph. While this way of using different kinds of information for different states is natural and useful for TD-based policy evaluation, it cannot be realized by keeping a single trace sequence as before and only letting λ_t evolve with states or histories. Indeed, in that case, as discussed in Section 2.2.2, interleaving large and small λ_t ’s would make the algorithm behave effectively like TD with small λ over the entire state space.

In the context of the more complex scheme of setting λ discussed in this paper, our motivation and reasons for considering composite schemes are the same. Each $T^{(i)}$ can be designed to be simple to implement, such as in the simple scaling example in Section 2.2. The parameters in the i th scheme can be chosen so that they encourage the use of large λ_t ’s throughout time or dictate the use of only small λ_t ’s. By combining component mappings of $T^{(i)}$ through (3.13), composite schemes allow us to use cumulative rewards and transition structures at different timescales for different states. This provides additional flexibility in managing the bias-variance trade-off when estimating the value function (see Figure 9 and Figure 11 in Section 4.2.2 for a demonstration).

Finally, we mention that for off-policy LSTD(λ) with constant λ , composite schemes were proposed in (Yu and Bertsekas, 2012) and analyzed in (Yu, 2012, Proposition 4.5, Section 4.3). Our Corollary 3.1 extends that result. The convergence analysis of the gradient-based algorithms for the composite schemes is given in (Yu, 2017). ■

3.3 Proofs of Theorem 3.2 and Corollary 3.1

We divide the proof of Theorem 3.2 into two steps. The first step deals with an expression for the trace vector, given in the following lemma. It is more subtle than the other step in the proof, which involves mostly calculations.

We start by extending the stationary state-trace process $\{(S_t, y_t, e_t)\}_{t \geq 0}$ to $t = -1, -2, \dots$, and work with a double-ended stationary process $\{(S_t, y_t, e_t)\}_{-\infty < t < \infty}$ (by Kolmogorov’s existence theorem (Dudley, 2002, Theorem 12.1.2), such a process exists). Note that as before this is a Markov chain whose transition probability is defined by the behavior policy π° together with the update rules (2.2) and (2.5) for e_t , y_t and λ_t , and the marginal distribution of each (S_t, y_t, e_t) is ζ . We keep using the notation \mathbf{P}_ζ and \mathbb{E}_ζ for this double-ended stationary Markov chain.

Recall the shorthand notation (3.1) introduced at the beginning of Section 3: For $k \leq m$, $\rho_k^m = \prod_{i=k}^m \rho_i$, $\lambda_k^m = \prod_{i=k}^m \lambda_i$, $\gamma_k^m = \prod_{i=k}^m \gamma_i$, and in addition, $\lambda_1^0 = \gamma_1^0 = \rho_0^{-1} = 1$ by convention.

Lemma 3.1 *\mathbf{P}_ζ -almost surely, $\sum_{t=1}^\infty \lambda_{1-t}^0 \gamma_{1-t}^0 \rho_{-t}^{-1} \phi(S_{-t})$ is well-defined and finite, and*

$$e_0 = \phi(S_0) + \sum_{t=1}^\infty \lambda_{1-t}^0 \gamma_{1-t}^0 \rho_{-t}^{-1} \phi(S_{-t}). \quad (3.15)$$

Proof First, we show $\mathbb{E}_\zeta[\sum_{t=1}^\infty \gamma_{1-t}^0 \rho_{-t}^{-1}] < \infty$. Indeed,

$$\mathbb{E}_\zeta[\sum_{t=1}^\infty \gamma_{1-t}^0 \rho_{-t}^{-1}] = \sum_{t=1}^\infty \mathbb{E}_\zeta[\gamma_{1-t}^0 \rho_{-t}^{-1}] = \sum_{t=1}^\infty \zeta_S^\top (P\Gamma)^t \mathbf{1} < \infty,$$

where the first equality follows from the monotone convergence theorem, the second equality from Condition 2.1(ii) and a direct calculation, and the last inequality follows from Condition 2.1(i) (since $(I - P\Gamma)^{-1} = \sum_{t=0}^\infty (P\Gamma)^t$). This implies $\sum_{t=1}^\infty \gamma_{1-t}^0 \rho_{-t}^{-1} < \infty$, \mathbf{P}_ζ -a.s., so $\gamma_{1-t}^0 \rho_{-t}^{-1} \rightarrow 0$ as $t \rightarrow \infty$, \mathbf{P}_ζ -a.s. Since $\lambda_{1-t}^0 \leq 1$ for all t , it also implies that

$$\mathbb{E}_\zeta[\sum_{t=1}^\infty \lambda_{1-t}^0 \gamma_{1-t}^0 \rho_{-t}^{-1} \|\phi(S_{-t})\|] \leq \max_{s \in \mathcal{S}} \|\phi(s)\| \cdot \mathbb{E}_\zeta[\sum_{t=1}^\infty \gamma_{1-t}^0 \rho_{-t}^{-1}] < \infty. \quad (3.16)$$

It then follows from a theorem on integration (Rudin, 1966, Theorem 1.38, p. 28-29) that \mathbf{P}_ζ -almost surely, the infinite series $\sum_{t=1}^\infty \lambda_{1-t}^0 \gamma_{1-t}^0 \rho_{-t}^{-1} \phi(S_{-t})$ converges to a finite limit.

We now prove the expression for e_0 . By unfolding the iteration (2.2) for e_t backwards in time, we have for all $m \geq 1$,

$$e_0 = \phi(S_0) + \sum_{t=1}^{m-1} \lambda_{1-t}^0 \gamma_{1-t}^0 \rho_{-t}^{-1} \phi(S_{-t}) + \lambda_{1-m}^0 \gamma_{1-m}^0 \rho_{-m}^{-1} e_{-m}. \quad (3.17)$$

Let $m \rightarrow \infty$ in the r.h.s. of (3.17). For the last term, the trace e_{-m} lies in a bounded set by Condition 2.3(ii), $\lambda_{1-m}^0 \leq 1$, and as we just showed, $\gamma_{1-m}^0 \rho_{-m}^{-1} \rightarrow 0$, \mathbf{P}_ζ -a.s. So the last term converges to zero \mathbf{P}_ζ -a.s. Also as we just showed, the second term converges \mathbf{P}_ζ -almost surely to $\sum_{t=1}^\infty \lambda_{1-t}^0 \gamma_{1-t}^0 \rho_{-t}^{-1} \phi(S_{-t})$. The expression (3.15) for e_0 then follows. \blacksquare

Proof of Theorem 3.2 Treating $\lambda_1^0 = \gamma_1^0 = \rho_0^{-1} = 1$, we write the expression of e_0 given in Lemma 3.1 as $e_0 = \sum_{t=0}^\infty \lambda_{1-t}^0 \gamma_{1-t}^0 \rho_{-t}^{-1} \phi(S_{-t})$, \mathbf{P}_ζ -a.s. We use this expression to calculate first $\mathbb{E}_\zeta[e_0 \cdot \rho_0 f(S_0^1)]$ for an arbitrary function f on $\mathcal{S} \times \mathcal{S}$. (Note that f is bounded and measurable, since \mathcal{S} is finite.) We have

$$\begin{aligned} \mathbb{E}_\zeta[e_0 \cdot \rho_0 f(S_0^1)] &= \sum_{t=0}^\infty \mathbb{E}_\zeta[\lambda_{1-t}^0 \gamma_{1-t}^0 \rho_{-t}^{-1} \phi(S_{-t}) \cdot \rho_0 f(S_0^1)] \\ &= \sum_{t=0}^\infty \mathbb{E}_\zeta[\lambda_1^t \gamma_1^t \rho_0^{t-1} \phi(S_0) \cdot \rho_t f(S_t^{t+1})] \\ &= \sum_{t=0}^\infty \mathbb{E}_\zeta[\phi(S_0) \cdot \mathbb{E}_\zeta[\lambda_1^t \gamma_1^t \rho_0^t f(S_t^{t+1}) \mid S_0, y_0, e_0]] \end{aligned} \quad (3.18)$$

where we used the stationarity of the double-ended state-trace process to derive the second equality, and we changed the order of expectation and summation in the first equality. This change is justified by the dominated convergence theorem (cf. (3.16)), and so are similar interchanges of expectation and summation that will appear in the rest of this proof.

To proceed with the calculation, we relate the expectations in the summation in (3.18) to expectations w.r.t. the process with probability measure \mathbf{P}_ζ^π introduced in Section 3.2

(which we recall is induced by the target policy π and involves the randomized stopping time τ). Let $\tilde{\mathbb{E}}_\zeta^\pi$ denote expectation w.r.t. the marginal of \mathbf{P}_ζ^π on the space of $\{(S_t, y_t, e_t)\}_{t \geq 0}$. From the change of measure performed through ρ_0^t , we have

$$\mathbb{E}_\zeta [\lambda_1^t \gamma_1^t \rho_0^t f(S_t^{t+1}) \mid S_0, y_0, e_0] = \tilde{\mathbb{E}}_\zeta^\pi [\lambda_1^t \gamma_1^t f(S_t^{t+1}) \mid S_0, y_0, e_0], \quad t \geq 0. \quad (3.19)$$

Combining this with (3.18) and using the fact that ζ is the marginal distribution of (S_0, y_0, e_0) in both processes, we obtain

$$\begin{aligned} \mathbb{E}_\zeta [e_0 \cdot \rho_0 f(S_0^1)] &= \sum_{t=0}^\infty \tilde{\mathbb{E}}_\zeta^\pi [\phi(S_0) \cdot \tilde{\mathbb{E}}_\zeta^\pi [\lambda_1^t \gamma_1^t f(S_t^{t+1}) \mid S_0, y_0, e_0]] \\ &= \tilde{\mathbb{E}}_\zeta^\pi [\phi(S_0) \cdot \sum_{t=0}^\infty \tilde{\mathbb{E}}_\zeta^\pi [\lambda_1^t \gamma_1^t f(S_t^{t+1}) \mid S_0]]. \end{aligned} \quad (3.20)$$

We now use (3.20) to calculate $\mathbb{E}_\zeta [e_0 \delta_0(v)]$ for a given function v . Recall from (2.3) that $\delta_0(v) = \rho_0 \cdot (r(S_0^1) + \gamma_1 v(S_1) - v(S_0))$, so we let $f(S_t^{t+1}) = r(S_t^{t+1}) + \gamma_{t+1} v(S_{t+1}) - v(S_t)$ in (3.20). Since $\mathbb{E}_\zeta^\pi [r(S_t^{t+1}) \mid S_0^t] = r_\pi(S_t)$, we have

$$\begin{aligned} \sum_{t=0}^\infty \tilde{\mathbb{E}}_\zeta^\pi [\lambda_1^t \gamma_1^t f(S_t^{t+1}) \mid S_0] &= \sum_{t=0}^\infty \tilde{\mathbb{E}}_\zeta^\pi \left[\lambda_1^t \gamma_1^t \left(r_\pi(S_t) + \gamma_{t+1} v(S_{t+1}) - v(S_t) \right) \mid S_0 \right] \\ &= (Tv - v)(S_0), \end{aligned}$$

where the last equality follows from the expression (3.11) for $Tv - v$. Therefore, by (3.20),

$$\mathbb{E}_\zeta [e_0 \delta_0(v)] = \sum_{s \in \mathcal{S}} \zeta_S(s) \phi(s) \cdot (Tv - v)(s), \quad (3.21)$$

and this shows that $\mathbb{E}_\zeta [e_0 \delta_0(v)] = 0$ is equivalent to $Tv - v \perp_{\zeta_S} \mathcal{L}_\phi$. \blacksquare

We now prove Corollary 3.1.

Proof of Corollary 3.1 We apply Theorem 3.2 to each state-trace process $\{(S_t, y_t^{(i)}, e_t^{(i)})\}$ for $i = 1, 2, \dots, m$. Specifically, by (3.21) and the definition (3.12) of $e_t^{(i)}$,

$$\mathbb{E}_{\zeta^{(i)}} [e_0^{(i)} \delta_0(v)] = \sum_{s \in \mathcal{S}} \zeta_S(s) \cdot \phi(s) \mathbb{1}(s \in \mathcal{S}_i) \cdot (T^{(i)}v - v)(s).$$

Hence

$$\begin{aligned} \sum_{i=1}^m \mathbb{E}_{\zeta^{(i)}} [e_0^{(i)} \delta_0(v)] &= \sum_{s \in \mathcal{S}} \zeta_S(s) \phi(s) \cdot \left[\sum_{i=1}^m \mathbb{1}(s \in \mathcal{S}_i) \cdot (T^{(i)}v - v)(s) \right] \\ &= \sum_{s \in \mathcal{S}} \zeta_S(s) \phi(s) \cdot (Tv - v)(s), \end{aligned}$$

where the last equality follows from the definition (3.13) of the operator T . This shows that the linear equation in v , $\sum_{i=1}^m \mathbb{E}_{\zeta^{(i)}} [e_0^{(i)} \delta_0(v)] = 0$, is equivalently $Tv - v \perp_{\zeta_S} \mathcal{L}_\phi$.

We now prove that T has v_π as its unique fixed point and is a contraction with respect to a weighted sup-norm—in other words, Theorem 3.1 applies to T . For this, it suffices to show that T satisfies the conditions of Theorem 3.1, namely, T is a generalized Bellman operator associated with a randomized stopping time τ that satisfies $\mathbf{P}^\pi(\tau \geq 1 \mid S_0 = s) > 0$ for all states $s \in \mathcal{S}$. We can define such a random time τ from the randomized stopping times $\tau^{(i)}$ associated with the Bellman operators $T^{(i)}$. In particular, by enlarging

the probability space if necessary, we can regard $\tau^{(i)}$, $i = 1, 2, \dots, m$, as being defined on the same probability space.¹⁰ We then let $\tau = \tau^{(i)}$ if $S_0 \in \mathcal{S}_i$. With this definition, we have $\mathbf{P}^\pi(\tau \geq 1 \mid S_0 = s) > 0$ for all states $s \in \mathcal{S}$ (since $\tau^{(i)} \geq 1$ a.s. for all i). For each set \mathcal{S}_i , by (3.2), the component mappings of the generalized Bellman operator T_τ associated with τ are given by

$$(T_\tau v)(s) = \mathbb{E}_s^\pi [R^\tau + \gamma_1^\tau v(S_\tau)] = \mathbb{E}_s^\pi [R^{\tau^{(i)}} + \gamma_1^{\tau^{(i)}} v(S_{\tau^{(i)}})] = (T^{(i)}v)(s), \quad s \in \mathcal{S}_i.$$

So $T_\tau = T$ by the definition (3.13) of T ; i.e., T is the Bellman operator associated with the randomized stopping time τ . \blacksquare

4. Numerical Study

In this section, we first use a toy problem to illustrate the behavior of traces calculated by off-policy LSTD(λ) for constant λ and for λ that evolves according to a simple special case of our proposed scheme described in Example 2.1. We then compare the behavior of LSTD for various choices of λ , on the toy problem and on the Mountain Car problem.

4.1 Behavior of Traces

The toy problem we use in this study has 21 states, arranged as shown in Figure 1 (left). One state is located at the centre, and the rest of the states split evenly into four groups, indicated by the four loops in the figure. The topology of the transition graph is the same for the target and behavior policies. We have drawn the transition graph only for the northeast group in Figure 1 (left); the states in each of the other three groups are arranged in the same manner and have the same transition structure. Given this symmetry, to specify the transition matrices P and P^o for the target and behavior policies π and π^o respectively, it suffices to specify the submatrices for the central state and one of the groups. If we label the central state as state 1 and the states in the northeast group clockwise as states 2-6, the submatrices of P and P^o for these states are given, respectively, by

$$\pi : \begin{pmatrix} 0 & 0.25 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0.2 & 0 & 0.8 & 0 & 0 \\ 0 & 0 & 0.2 & 0 & 0.8 & 0 \\ 0 & 0 & 0 & 0.2 & 0 & 0.8 \\ 0.8 & 0 & 0 & 0 & 0.2 & 0 \end{pmatrix}, \quad \pi^o : \begin{pmatrix} 0 & 0.25 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0.5 & 0 & 0.5 & 0 & 0 \\ 0 & 0 & 0.5 & 0 & 0.5 & 0 \\ 0 & 0 & 0 & 0.5 & 0 & 0.5 \\ 0.5 & 0 & 0 & 0 & 0.5 & 0 \end{pmatrix}.$$

Intuitively speaking, from the central state, the system enters one group of states by moving diagonally in one of the four directions with equal probability, and after spending some time in that group, eventually returns to the central state and the process repeats. The behavior policy on average spends more time wandering inside each group than the target policy, while the target policy tends to traverse clockwise through the group more quickly.

10. That we can do so is clear from the definition of each $\tau^{(i)}$ as described at the beginning of Section 3.2 and from the fact that for each i , the initial distribution $\zeta^{(i)}$ on $(S_0, y_0^{(i)}, e_0^{(i)})$ has the same marginal on \mathcal{S} , which is ζ_S .




Figure 1: The transition graph of a toy problem (left) and a cycle pattern in it (right). The numbers appearing in the right graph indicate the importance sampling ratios for the state transitions represented by each directed edge. From such cycle patterns one can infer whether the trace sequence is unbounded almost surely.

All the rewards are zero except for the middle state in each group—for the northeast group, this is the shaded state in Figure 1 (left). For the two northern (southern) groups, their middle states have reward 1 (−1). The discount factor is $\gamma = 0.9$ for all states. As to features, we aggregate states into 5 groups: the 4 groups mentioned earlier and the central state forming its own group, and we let each state have 5 binary features indicating its membership.

We now discuss and illustrate the behavior of traces in this toy problem. For comparison, we first do this for the off-policy TD(λ) with a constant λ . It can help explain the challenges in off-policy TD learning and our motivation for proposing the new scheme of setting λ .

4.1.1 TRACES FOR CONSTANT λ

In this experiment we let $\lambda = 1$ and consider the trace iterates $\{e_t\}$ calculated by TD(1). In general, by identifying certain cycle patterns in the transition graph, one can infer whether $\{e_t\}$ will be unbounded over time almost surely (Yu, 2012, Section 3.1). Figure 1 (right) shows such a cycle of states in the transition graph of the toy problem. It consists of the central state and the northeast group of states. Labeled on each edge of the cycle is the importance sampling ratio for that state transition. Traversing through the cycle once from any starting state, and multiplying together the importance sampling ratios of each edge and the discount factors of the destination states, we get $(\frac{0.8}{0.5})^4 \cdot \gamma^6 = (\frac{0.8}{0.5})^4 \cdot 0.9^6 > 1$. From this one can infer that $\{e_t\}$ calculated by off-policy TD(1) will be unbounded in this problem (cf. Yu, 2012, Prop. 3.1).

We plotted in the upper left graph of Figure 2 the Euclidean norm $\|e_t\|$ of the traces over 8×10^5 iterations for TD(1). One can see the recurring spikes and the exceptionally large values of some of these spikes in the plot. This is consistent with the unboundedness of $\{e_t\}$ just discussed.

The unboundedness of $\{e_t\}$ tells us that the invariant probability measure ζ of the state-trace process $\{(S_t, e_t)\}$ has an unbounded support. Despite this unboundedness, $\{e_t\}$ is bounded in probability (Yu, 2012, Lemma 3.4) and under the invariant distribution ζ , $\mathbb{E}_\zeta[\|e_0\|] < \infty$ (Yu, 2012, Prop. 3.2). The latter property implies that under the invariant distribution, the probability of $\|e_0\| > x$ decreases as $o(1/x)$ for large x . Since the empirical

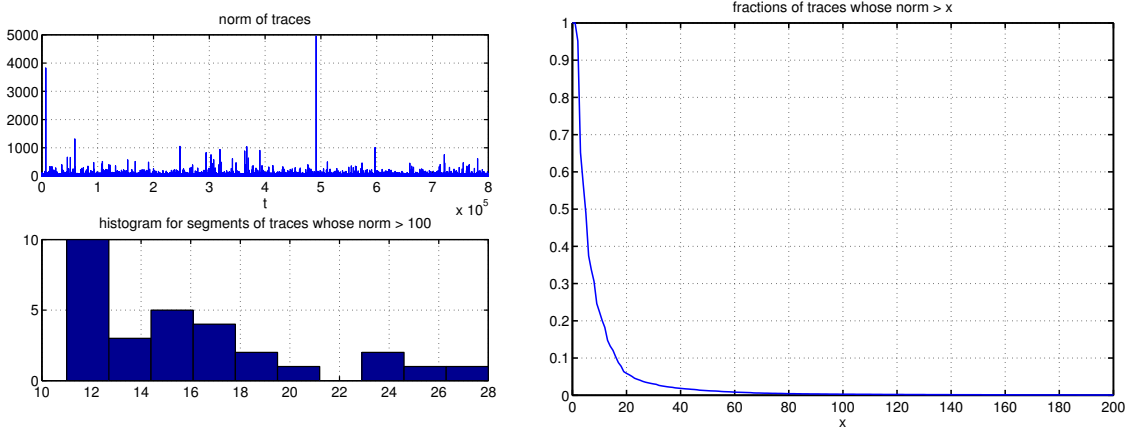


Figure 2: Statistics of traces for TD(1) in a toy problem (see the text in Section 4.1.1 for detailed explanations).

distribution of the state-trace process converges to ζ almost surely (Yu, 2012, Theorem 3.2), during a run of many iterations, we expect to see the fraction of traces with $\|e_t\| > x$ drop in a similar way as x increases.

The simulation result shown in the right part of Figure 2 agrees with the preceding discussion. Plotted in the graph are fractions of traces with $\|e_t\| > x$ during 8×10^5 iterations (the vertical axis indicates the fraction, and the horizontal axis indicates x). It can be seen that despite the recurring spikes in $\|e_t\|$ during the entire run, the fraction of traces with large magnitude x drops sharply with the increase in x .

While only a small fraction of traces have exceptionally large magnitude, they can occur in consecutive iterations. This is illustrated by the histogram in the lower left part of Figure 2. The histogram concerns the excursions of the trajectory $\{e_t\}$ outside of the ball $\{e \in \mathbb{R}^n \mid \|e\| \leq 100\}$. The horizontal axis indicates the lengths of the excursions (where the length is the number of iterations an excursion contains), and the vertical axis indicates how many excursions of length x occurred during the 8×10^5 iterations of the experimental run. We plotted the histogram for lengths $x > 10$. It can be seen that one can have large traces during many consecutive iterations. Such behavior, although tolerable by LSTD, is especially detrimental to TD algorithms and can disrupt their learning. This is our main motivation for suggesting the use of λ -parameters to bound the traces directly.

4.1.2 TRACES WITH EVOLVING λ

We now proceed to illustrate the behavior of traces and LSTD for λ that evolves according to our proposed scheme. Specifically, for this demonstration, we will use the simple scaling example given by (2.6)-(2.7) in Example 2.1, with all the thresholds $C_{ss'}$ being the same constant C . That is, the update rule for e_t used in this experiment is

$$e_t = \begin{cases} \gamma_t \rho_{t-1} e_{t-1} + \phi(S_t) & \text{if } \gamma_t \rho_{t-1} \|e_{t-1}\|_2 \leq C; \\ C \cdot \frac{e_{t-1}}{\|e_{t-1}\|_2} + \phi(S_t) & \text{otherwise.} \end{cases} \quad (4.1)$$

We first simulate the state-trace process to illustrate the ergodicity of this process stated by Theorem 2.1. We will shortly study the performance of LSTD for different values of C in Section 4.2.1. The results of these two experiments are shown in Figure 3 and Figures 4-5, respectively, and the details are as follows.

According to the ergodicity result of Theorem 2.1, no matter from which initial state and trace pair (S_0, e_0) we generate a trajectory $\{(S_t, e_t)\}_{0 \leq t \leq \bar{t}}$ according to the behavior policy, the empirical distribution of state and trace pairs in this trajectory should converge, as $\bar{t} \rightarrow \infty$, to the same distribution on $\mathcal{S} \times \mathbb{R}^n$, which is the marginal of the invariant probability measure ζ on that space. Since \mathcal{S} is discrete, we can verify this fact by examining the empirical conditional distribution of the trace given the state. In other words, for each state s , we examine the empirical distribution of the trace for the sub-trajectory (S_{t_k}, e_{t_k}) , $k = 1, 2, \dots$, where $S_{t_k} = s$ and it is the k th visit to state s by the trajectory. We check if this empirical distribution converges to the same one as we increase the length \bar{t} of the trajectory and as we vary the initial condition (S_0, e_0) .




Figure 3: Demonstration of convergence of empirical conditional distributions on the trace space. (See the text in Section 4.1.2 for detailed explanations.)

To give a sense of what these limiting distributions over the trace space look like, we set the parameter $C = 50$ and generated a long trajectory with 8×10^5 iterations. In the top row of Figure 3, we plotted three normalized histograms of the first trace component for the sub-trajectories associated with three states of the toy problem, respectively: the central state (left), the middle state of the northeast group (middle), and the first state of the southeast group (right).

To check whether the empirical conditional distributions on the trace space converge along the sub-trajectory $\{(S_{t_k}, e_{t_k})\}$ for a given state, we compare the characteristic functions f_k of these distributions with the characteristic function f of the empirical conditional distribution obtained at the end of the sub-trajectory during the experimental run. In par-

ticular, we evaluate all these (complex-valued) characteristic functions at 500 points, which are chosen randomly according to the multivariate normal distribution on \mathbb{R}^5 with mean 0 and covariance matrix $200^2 I$. We take the maximal difference between f_k and f at these 500 points as an indicator of the deviation between the two corresponding distributions.¹¹ In the bottom row of Figure 3, the first three plots show the difference curves obtained in the way just described, for three different states, respectively. These three states are the same ones mentioned earlier in the description of the top row of Figure 3. The horizontal axis of these plots indicates k , the number of visits to the corresponding state. As can be seen, the difference curves all tend to zero as k increases, which is consistent with the predicted convergence of the empirical conditional distributions on the trace space.

So far we compared the empirical distributions along the same trajectory. Next we compare them against the one obtained at the end of another trajectory that starts from a different initial condition. The difference curve for one state (the first state in the southeast group) is plotted in the last graph in the bottom row of Figure 3, and it is the lower curve in that graph. As can be seen, the curve tends to zero, suggesting that the limiting distribution of these empirical distributions does not change if we vary the initial condition, which is consistent with Theorem 2.1.

For comparison, we also plotted the difference curve when these same empirical conditional distributions are compared against the empirical conditional distribution obtained from the same trajectory but for a different state (specifically, the middle state of the north-east group). This is the upper curve in the last graph in the bottom row of Figure 3. It clearly indicates that for the two states, the associated limiting conditional distributions on the trace space are different. It also shows that the characteristic function approach we adopted in this experiment can effectively distinguish between two different distributions (cf. Footnote 11).

4.2 LSTD with Evolving λ

We now present experiments on the LSTD algorithm.

4.2.1 A TOY PROBLEM

Let us first continue with the toy problem of the previous subsection and show how the LSTD algorithm performs in this problem as we vary the parameter C in the λ function for bounding the traces. We ran LSTD for $C = 10, 20, \dots, 100$, using the same trajectory, for 3×10^5 iterations, and we computed the (Euclidean) distance of these LSTD solutions to the asymptotic TD(1) solution (in the space of the θ -parameters), normalized by the norm of the latter. We then repeat this calculation 10 times, each time with an independently generated trajectory. Plotted in Figure 4 (left) against the values of C are the means and standard deviations of the normalized distances of LSTD solutions thus obtained.

11. Recall that a tight sequence $\{p_k\}$ of probability distributions on \mathbb{R}^m converges to a probability distribution p if and only if the characteristic functions of p_k converge pointwise to the characteristic function of p (Dudley, 2002, Lemma 9.5.5). Recall also that we are dealing with convergence in distribution here, which is much weaker than convergence in total variation, so we cannot use total variation as a metric on the distribution space in this case.




Figure 4: Compare LSTD solutions with evolving λ (left) and with constant λ (right). For constant λ , the red dot-dash curve in the right plot shows the quality of the asymptotic TD(λ) solutions, and LSTD(λ) would approach this curve in the limit, but due to variance issues, it can require an impractically large number of iterations to exhibit this convergent behavior. LSTD with evolving λ outperforms LSTD with constant λ in this case and effectively archives the quality of TD(λ) solutions for large constant λ .




Figure 5: Approximation quality of asymptotic TD(λ) solutions with constant λ for the toy problem. Larger λ yields better approximations.

For comparison, we did the same for LSTD with a large constant λ : $\lambda = 0.9, 0.92, \dots, 1$. Figure 4 (right) shows the result, where the dash-dot curve indicates the normalized distance of the asymptotic TD(λ) solution—the solution that LSTD(λ) would obtain in the limit. It can be seen that the performance of LSTD deteriorates as λ gets close to 1. We think that this is because due to the high variance issue, the convergence of LSTD with a large constant λ is too slow and requires far more iterations than the 3×10^5 iterations performed. In comparison, LSTD with evolving λ behaves better: it works effectively for $C \geq 20$, and

the approximation quality it achieved with such C is comparable to that of asymptotic $\text{TD}(\lambda)$ solutions for a large constant λ around 0.96.

Figure 5 shows the quality of the asymptotic solutions of $\text{TD}(\lambda)$ with constant λ , for the full range of λ values. Plotted in the left graph is the normalized distance of the $\text{TD}(\lambda)$ solution to the $\text{TD}(1)$ solution. Plotted in the right graph is the normalized approximation error of the corresponding approximate value function, where the error is measured by the weighted Euclidean norm with weights specified by $\zeta_{\mathcal{S}}$ (the invariant distribution on \mathcal{S} under the behavior policy), and the normalization is over the weighted norm $\|v_{\pi}\|_{\zeta_{\mathcal{S}}}$ of the true value function v_{π} of the target policy. It can be seen that using $\lambda > 0.9$ provides considerably better approximations for this problem than using small λ .

We also found that for this problem, if we set λ according to the Retrace algorithm with $\beta = 1$ (cf. (2.8) in Example 2.2), then the performance of LSTD is comparable to $\text{TD}(\lambda)$ with a small λ around 0.5. (Specifically, for Retrace, the normalized distance to the $\text{TD}(1)$ solution is 0.37, and the normalized approximation error is 0.54, which are comparable to the numbers for $\text{TD}(0.5)$, as Figure 5 shows.) This is not surprising, because, as we discussed earlier in Section 2.2.2, in keeping $\lambda_t \rho_{t-1} \leq 1$ always, Retrace and ABQ can be too “conservative,” resulting in an overall effect that is like using a small λ , even though λ_t may appear to be large at times. Recall also that this can happen to our proposed scheme too. In the present experiment, for instance, this can happen when C is small; in particular, the case $C = 0$ reduces to $\text{LSTD}(0)$.

4.2.2 MOUNTAIN CAR PROBLEM

In this subsection we demonstrate LSTD with evolving λ on a problem adapted from the well-known Mountain Car problem (Sutton and Barto, 1998). The details of this adaptation, including the target and behavior policies involved, can be found in the report (Yu, 2016a, Section 5.1, p. 23-26); most of these details are not crucial for our experiments, so to avoid distraction, we only describe briefly the experimental setup here.

In Mountain Car, the goal is to drive an underpowered car to reach the top of a steep hill, from the bottom of a valley. A state consists of the position and velocity of the car, whose values lie in the intervals $[-1.2, 0.5]$, $[-0.07, 0.07]$, respectively. The position 0.5 corresponds to the desired hill top destination, while the position $-\pi/6$ (≈ -0.52) lies at the bottom of a valley that is between the destination and a second hill peaked at -1.2 in the opposite direction (see the illustration in Figure 6). Except for the destination state, each state has three available actions: `{back, coast, forward}`, and the rewards depend only on the action taken and are -1.5 , 0 and -1 for the three actions, respectively. The dynamics is as given in (Sutton and Barto, 1998). We consider undiscounted expected total rewards, so the discount factor is 1 except at the destination state, where the discount factor is 0 and from where the car enters a rewardless termination state permanently.

The target policy π is a simple but reasonably well-behaved policy. On either slopes between the two hills, it tries to increase its energy (kinetic plus gravitational potential energy) by accelerating in the direction of its current motion. If this brings it up to the opposite hill (position < -1), it coasts; otherwise, if its velocity drops to near zero, it goes forward or backward with equal probability. Figure 6 visualizes the total costs $-v_{\pi}$ of the




Figure 6: Left: Illustration of the Mountain Car problem. Right: Costs $-v_\pi$ for this problem are estimated and visualized as a color image, with the coloring scheme indicated by the colorbar. (The horizontal and vertical axes of the image correspond to position and velocity, respectively, for the 2-dimensional state space of this problem.)




Figure 7: Left: Visualization of weights on states induced by the behavior policy. Right: The color image visualizes the approximation of $-v_\pi$ obtained from a discretized model for the Mountain Car problem, where the coloring scheme is the same as that shown in Figure 6 (right). The quality of this approximation is close to that of TD(0).

target policy,¹² where the horizontal (vertical) axis indicates position (velocity) and the

12. The values of v_π shown in Figure 6 are estimated by simulating the target policy for each starting state in a set of 171×141 points evenly spaced in the position-velocity space. In particular, the position (velocity) interval is evenly divided into subintervals of length 0.01 (0.001), and for each starting state, the target policy is simulated 600 times.

colorbar on the right shows the value corresponding to each color. The discontinuity of the function in certain regions can be seen in this figure.

The behavior policy π^o is an artificial policy that takes a random action (chosen with equal probability from the three actions) 90% of the time, and explores the state space by jumping to some random state 10% of the time. It also restarts when it is at the destination: with equal probability, it either restarts near the bottom of the valley or restarts from a random point sampled uniformly from the state space.¹³

We now explain how we will measure approximation qualities. Since the Mountain Car problem has a continuous state space, it is actually not covered by our analysis, which is for finite-state problems. Although we can treat it as essentially a finite-state problem (since the simulation is done with finite precision in computers), the number of states would still be too large to calculate the weights ζ_S . So, to measure weighted approximation errors $\|v - v_\pi\|$ for approximate value functions v produced by various LSTD algorithms in the subsequent experiments, we will compare v and v_π at a grid of points in the state space and calculate a weighted Euclidean distance between the function values at these grid points, using a set of weights precalculated by simulating the behavior policy.¹⁴ The image in Figure 7 (left) visualizes these weights.

We use tile-coding (Sutton and Barto, 1998) to generate 145 binary features for our experiments.¹⁵ The approximate value functions obtained with LSTD algorithms are thus piecewise constant. For comparison, we also build a discrete approximate model by state aggregation. The discretization is done at a resolution comparable to our tile-coding scheme, and the dynamics and rewards of this model are calculated based on data collected under the behavior policy. The solution of the discrete approximate model is shown in Figure 7 (right) (the coloring scheme for this and the subsequent images are the same as shown in Figure 6). It is similar to the approximate value function calculated by LSTD(0), which is shown in Figure 8 (first image, top row). As will be seen shortly, with positive λ , the approximation quality of LSTD improves. Thus the discrete model approximation approach is not as effective as the TD method in this case.

We now report the results of our experiments on the Mountain Car problem.

First Experiment: In this experiment, we compare three ways of setting λ : (i) Retrace with $\beta = 1$ (cf. Example 2.2); (ii) our simple scaling scheme with parameter C used in the previous experiments (cf. (4.1) and Example 2.1); and (iii) a composite scheme of the type discussed at the end of Section 3.2, which partitions the state space into two sets¹⁶

-
13. The behavior policy is exactly the same as described in (Yu, 2016a, p. 24-25) except for the possibility of restarting near the bottom of the valley whenever the destination is reached.
 14. Specifically, we chose a grid of 171×141 points evenly spaced in the position-velocity space. We ran the behavior policy for 8×10^5 *effective* iterations, where an iteration is considered to be *ineffective* if the behavior policy takes an action, e.g., the restart action, that is impossible for the target policy. A visit to a state at an effective iteration was counted as a visit to the nearest grid point. At the end of the run, visits to a boundary point $(-1.2, 0)$ were disregarded as they were due to boundary effects in the dynamics of this problem, and the final counts were normalized to produce a set of weights on the grid points that sum to 1.
 15. Two tilings are used: the first (second) comprises of 64 (81) uneven-sized rectangles that cover the state space. Together they produce a total of 145 binary features. The details of the coding scheme are as described in (Yu, 2016a, p. 28).
 16. The first set consists of those states (position, velocity) with either position ≤ -0.9 or velocity ≥ 0.04 . The rest of the states belong to the second set.





Figure 8: Visualized in the color images are approximations of $-v_\pi$ obtained by LSTD with different schemes of setting λ , where the coloring scheme is as that shown in Figure 6 (right). The choices of λ for each image, from left to right and top to bottom, are as follows. Top row: $C = 0$ (equivalent to LSTD(0)), $C = 5$, Retrace. Middle row: $C = 25$, $C = 125$, $C: (125, 25)$. Bottom row: $C = 200$, $C = 300$, $C: (300, 50)$.



and applies the simple scaling scheme with parameters C_1, C_2 for the first and second set, respectively. When referring to this composite scheme in the figures, we will use the designation $C: (C_1, C_2)$.

We ran LSTD with different ways of setting λ just mentioned, on the same state trajectory generated by the behavior policy, for 6×10^5 effective iterations (cf. Footnote 14). Some of the approximate value functions obtained at the end of the run are visualized as images in Figure 8. It can be seen that the result of Retrace is similar to that of the simple scaling scheme with a small C , and as we increase C , the approximation from the scaling scheme improves.

To compare more precisely the approximation errors and see how they change over time for each algorithm, we did 10 independent runs, each of which consists of 6×10^5 effective


 Figure 9: Compare the approximation error of LSTD for different schemes of setting λ .

iterations. The results are shown in Figures 9-11. Plotted in Figure 9 for each algorithm are the mean and standard deviation of the approximation errors for the 10 approximate value functions obtained by that algorithm at the end of the 10 runs. We can see from this figure the improvement in approximation quality as C increases. We can also see that the result of Retrace is in between those of $C = 0$ and $C = 5$, which is consistent with what the images in the top row of Figure 8 tell us.


 Figure 10: Compare the temporal behavior of LSTD for different schemes of setting λ .

 Figure 11: Compare the temporal behavior of LSTD for different schemes of setting λ .

Plotted in Figures 10-11 are the approximation errors calculated per 2000 effective iterations for each algorithm, during one of the 10 experimental runs. Figure 10 (left) shows how Retrace compares with the simple scaling with $C = 5$ and $C = 25$. It can be seen that the latter two achieved better approximation quality than Retrace without increases in variance. Figure 10 (right) shows that for larger values of C , variances also became larger initially; however, after about 5×10^4 iterations, these schemes overtook Retrace, yielding better approximations.

Figure 11 shows how the composite scheme of setting λ performs. Comparing the plots in this figure with the right plot in Figure 10, it can be seen that the composite schemes helped in reducing variances, and in about 2×10^4 iterations the schemes $C: (125, 0)$ and $C: (125, 25)$ overtook Retrace and yielded better approximations. Together with Figure 9, Figure 11 shows clearly the bias-variance trade-off of using composite schemes in this problem.

Second Experiment: Similarly to the previous experiment, we now compare our proposed method with several other ways of setting λ for the LSTD algorithm as well as with a constrained variant of LSTD: (i) Retrace with $\beta = 1$ as before; (ii) constant λ ; (iii) constrained LSTD with constant λ ; and (iv) the simple scaling scheme with parameter C . For constant $\lambda \in [0, 1]$, the constrained LSTD(λ) used in this experiment evolves the trace vectors as off-policy LSTD(λ) does, but it forms and solves the linear equation $\frac{1}{t} \sum_{k=0}^{t-1} [e_k]_{50} \cdot \delta_k(v) = 0, v = \Phi\theta$ instead, where the function $[\cdot]_{50}$ truncates each component of the trace vector to be within the interval $[-50, 50]$. (Such an algorithm follows naturally from the ergodicity of the state-trace process and the approximation of an unbounded integrable function by a bounded one. For a detailed discussion, see (Yu, 2016b, Section 3.2) or (Yu, 2017, Section 3.3).)




Figure 12: Compare the approximation errors of several LSTD algorithms.

Figure 12 is similar to Figure 9 and shows, for each algorithm, the mean and standard deviation of the approximation errors of 10 approximate value functions obtained at the end of 10 independent experimental runs (each of which consists of 6×10^5 effective iterations). The horizontal axis indicates the algorithms and their parameters. As can be seen from this figure, for constant small λ , LSTD(λ) performed well in this problem, and LSTD(0.3) and Retrace are comparable. For constant $\lambda > 0.7$, LSTD(λ) failed to give sensible results, and LSTD(0.7) started to show this unreliable behavior. This behavior of LSTD(λ) is related to what we observed in Figure 4(right) in the small toy problem, and it is, we think, due to

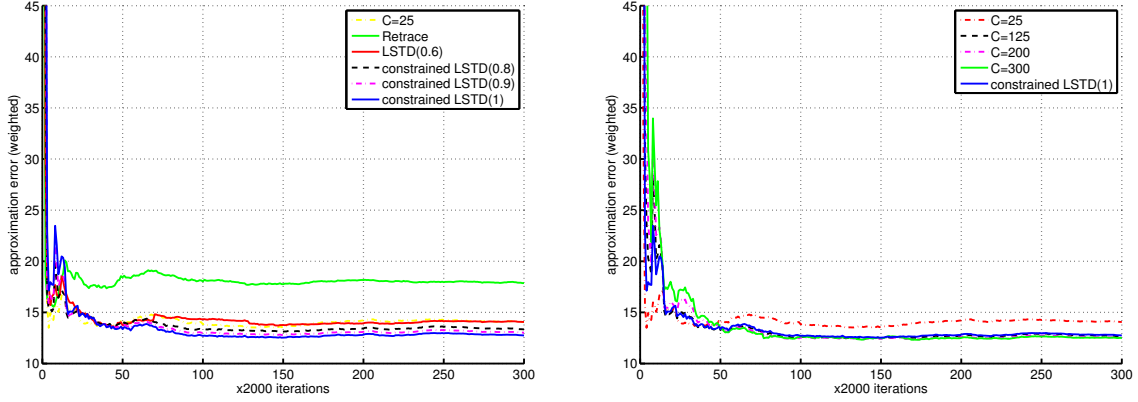


Figure 13: Compare the temporal behavior of several LSTD algorithms.

the high variance issue, which becomes more severe as λ gets larger. Constrained LSTD(λ) is much more reliable, and it did consistently well for all values of λ tested. LSTD with evolving λ also did well, and with $C > 125$, it achieved a slightly better approximation quality than constrained LSTD(1).

Figure 13 compares the temporal behavior of several algorithms during one experimental run. Plotted are the approximation errors calculated per 2000 iterations for each algorithm in the comparison. It can be seen that in this problem constrained LSTD(λ) did not suffer from large variances even with large λ values, and compared with constrained LSTD(1), the behavior of LSTD with evolving λ was also reasonable for large C .

Third Experiment: In this experiment we first compare the simple scaling scheme and Retrace, where both schemes now use an additional parameter $\beta \in [0, 1]$, as discussed in Examples 2.1-2.2. Recall that when $\beta = 1$, they reduce to the schemes that we already compared in the previous experiments. For both schemes, as β becomes smaller, we expect the approximation quality to drop but the variance to get smaller.

Plotted in Figure 14 for $C = 125$, $C = 200$ and Retrace are the results obtained from a single experimental run consisting of 3×10^5 effective iterations. As before the approximation errors were calculated per 2000 iterations. The results do show the expected bias-variance trade-off effects of the parameter β , during the initial period of the experimental run, although the effects on Retrace turned out to be smaller and hard to discern at the scale of the plot.

Next we test some of the variations on Retrace discussed in Example 2.2. Specifically, we consider (2.10)-(2.11) with parameters $\beta = 0.9$, $K \in \{1.5, 2.0, 2.5, 3.0\}$ and $C \in \{50, 125\}$. Plotted in Figure 15 are the results from one experimental run of 3×10^5 effective iterations. For comparison, the plots also show the behavior of Retrace and the simple scaling scheme with the same values of C during that run (these algorithms used the same $\beta = 0.9$). As expected and can be seen from the figure, the approximation quality improves with K and C . While the variances also tend to increase during the initial part of the run, the variants that truncate the importance sampling ratios by $K = 1.5$ performed comparably to Retrace initially, soon overtook Retrace and achieved better approximation quality.




Figure 14: Compare the approximation error and temporal behavior of LSTD for different schemes of setting λ . From left to right: $C = 125$, $C = 200$, Retrace.




Figure 15: Compare the approximation error and temporal behavior for some variations on Retrace.

5. Conclusion

We developed in this paper a new scheme of setting the λ -parameters for off-policy TD learning, using the ideas of randomized stopping times and generalized Bellman equations for MDPs. Like the two recently proposed algorithms Retrace (Munos et al., 2016) and ABQ (Mahmood et al., 2017), our scheme keeps the traces bounded to reduce variances, but it is much more general and flexible. To study its theoretical properties, we analyzed the resulting state-trace process and established convergence and solution properties for the associated LSTD algorithm, and these results have prepared the ground for convergence analysis of the gradient-based implementation of our proposed scheme (Yu, 2017). In addition we did a preliminary numerical study. It showed that with the proposed scheme LSTD can outperform several existing off-policy LSTD algorithms. It also demonstrated that in order to achieve better bias-variance trade-offs in off-policy learning, it is helpful to have more flexibility in choosing the λ -parameters and to allow for large λ values. Future research is to conduct a more extensive numerical study of both least-squares based and gradient-based algorithms, with more versatile ways of using the memory states and λ -parameters, in off-policy learning applications.

Acknowledgments

An abridged version of this article appeared first at *The 30th Canadian Conference on Artificial Intelligence (CAI)*, May 2017. We thank several anonymous reviewers for CAI and JMLR for their helpful feedback. We also thank Dr. Ajin Joseph and Dr. Martha Steensrup for reading parts of our manuscript and providing comments that helped improve our presentation. This research was supported by a grant from Alberta Innovates—Technology Futures.

Appendix A. Proof of Theorem 3.1

We prove Theorem 3.1 in this appendix. Recall from Section 3.1 that the generalized Bellman operator T associated with a randomized stopping time τ satisfies that $v_\pi = Tv_\pi$. By (3.2), the substochastic matrix \tilde{P} in this affine operator T is given by

$$\tilde{P}_{ss'} = \mathbb{E}_s^\pi[\gamma_1^\tau \mathbb{1}(S_\tau = s')], \quad s, s' \in \mathcal{S}. \quad (\text{A.1})$$

We can extend \tilde{P} to a transition matrix \tilde{P}^e by adding an additional absorbing state Δ to the system, so that $\tilde{P}_{\Delta\Delta}^e = 1$ and

$$\tilde{P}_{s\Delta}^e = 1 - \sum_{s' \in \mathcal{S}} \tilde{P}_{ss'} = 1 - \mathbb{E}_s^\pi[\gamma_1^\tau], \quad s \in \mathcal{S}.$$

Both conclusions of Theorem 3.1 will follow immediately if we show that $I - \tilde{P}$ is invertible. Indeed, if $I - \tilde{P}$ is invertible, then $v = Tv$ has a unique solution, which must be v_π since v_π is always a solution of this equation. In addition, if $I - \tilde{P}$ is invertible, then since \tilde{P} is a substochastic matrix, the spectral radius of \tilde{P} must be less than 1. Consider adding to \tilde{P} a small enough perturbation ϵM , where M is the matrix of all ones and ϵ is a sufficiently small positive number so that the spectral radius of $\tilde{P} + \epsilon M$ is less than 1. Applying (Seneta, 2006, Theorem 1.1) to the nonnegative primitive matrix $\tilde{P} + \epsilon M$, we have that $(\tilde{P} + \epsilon M)w < w$ where w is a positive eigenvector of the matrix $\tilde{P} + \epsilon M$ corresponding to a positive eigenvalue that is strictly less than 1. Consequently $\tilde{P}w \leq (\tilde{P} + \epsilon M)w < w$, implying that \tilde{P} is a linear contraction w.r.t. a weighted sup-norm (with weights w), which is the second conclusion of the theorem.

Hence, to prove Theorem 3.1, it suffices to show that the inverse $(I - \tilde{P})^{-1}$ exists, which is equivalent to that for the Markov chain on $\mathcal{S} \cup \{\Delta\}$ with transition matrix \tilde{P}^e , all the states in \mathcal{S} are transient (see e.g., Puterman 1994, Appendix A.4).

We prove this by contradiction. Suppose it is not true, and let $\tilde{\mathcal{S}} \subset \mathcal{S}$ be a recurrent class of the Markov chain. Then for all $s \in \tilde{\mathcal{S}}$, $\tilde{P}_{ss'}^e = 0$ for $s' \notin \tilde{\mathcal{S}}$ (i.e., the submatrix of \tilde{P} corresponding to $\tilde{\mathcal{S}}$ is a transition matrix). In particular, $\tilde{P}_{s\Delta}^e = 0$, so

$$\mathbb{E}_s^\pi[\gamma_1^\tau] = 1 - \tilde{P}_{s\Delta}^e = 1, \quad \forall s \in \tilde{\mathcal{S}}, \quad (\text{A.2})$$

implying that given $S_0 \in \tilde{\mathcal{S}}$, $\gamma_1^\tau = 1$ a.s. (since $\gamma_1^\tau \in [0, 1]$). Then by (A.1),

$$\tilde{P}_{ss'} = \mathbb{E}_s^\pi[\mathbb{1}(S_\tau = s')], \quad \forall s, s' \in \tilde{\mathcal{S}}. \quad (\text{A.3})$$

Observe from (A.2) that given $S_0 \in \tilde{\mathcal{S}}$, the event $\{\tau = \infty\}$ has zero probability. This is because under Condition 2.1(i), $\gamma_1^t \xrightarrow{a.s.} 0$ (as $t \rightarrow \infty$) and consequently $\gamma_1^\infty := \prod_{t=1}^\infty \gamma_t = 0$

a.s. Indeed, the existence of the inverse $(I - P\Gamma)^{-1} = \sum_{t=0}^{\infty} (P\Gamma)^t$ under Condition 2.1(i) implies $(P\Gamma)^t \rightarrow 0$ as $t \rightarrow \infty$. Since the s th entry of $(P\Gamma)^t \mathbf{1}$ equals $\mathbb{E}_s^\pi[\gamma_1^t]$, we have $\mathbb{E}_s^\pi[\gamma_1^t] \rightarrow 0$ as $t \rightarrow \infty$. As the nonnegative sequence $\{\gamma_1^t\}_{t \geq 1}$ is nonincreasing (since each $\gamma_t \in [0, 1]$), this implies, by Fatou's lemma (Dudley, 2002, Lemma 4.3.3), that $\gamma_1^t \xrightarrow{a.s.} 0$. Thus if $S_0 \in \tilde{\mathcal{S}}$, τ is almost surely finite.

We now consider a Markov chain $\{S_t\}$ with transition matrix P and $S_0 \in \tilde{\mathcal{S}}$. We will extract from it a Markov chain $\{\tilde{S}_k\}_{k \geq 0}$ with transition matrix \tilde{P}^e on the recurrent class $\tilde{\mathcal{S}}$, by employing multiple stopping times. We will show $\gamma_1^\infty = 1$ a.s., contrary to the fact $\gamma_1^\infty = 0$ a.s. just discussed.

For ease of explanation, let us imagine that there is a device that if we give it a sequence of states S_0, S_1, \dots generated according to P , it will output the stopping decision at a random time τ that is exactly the randomized stopping time τ associated with the operator T . (Because τ is a randomized stopping time for $\{S_t\}$, a device that correctly implements τ does not affect the evolution of $\{S_t\}$, so we can give $\{S_t\}$ to the device as inputs.)

Let $\{S_t\}$ start from some state $S_0 = s \in \tilde{\mathcal{S}}$. We use the device just mentioned to generate the first randomized stopping time τ_1 . We then reset the device so that it now “sees” the time-shifted process $\{S_{\tau_1+t'} \mid t' \geq 0\}$ with the initial state being S_{τ_1} . We wait till the device makes another stopping decision, and we designate that time by τ_2 . We repeat this procedure as soon as the device makes yet another stopping decision. This gives us a nondecreasing sequence $0 \leq \tau_1 \leq \tau_2 \leq \dots$.

Since $S_0 = s \in \tilde{\mathcal{S}}$, by what we proved earlier, τ_1 is almost surely finite. Let $\tilde{S}_1 = S_{\tau_1}$ and $\tilde{S}_0 = S_0$. The transition from \tilde{S}_0 to \tilde{S}_1 is according to the transition matrix \tilde{P}^e by construction (cf. (A.3)). Since $\tilde{\mathcal{S}}$ is a recurrent class for \tilde{P}^e , we must have $\tilde{S}_1 \in \tilde{\mathcal{S}}$ almost surely. Then, repeating the same argument and using induction, we have that almost surely, for all $k \geq 1$, τ_k is defined and finite and $\tilde{S}_k := S_{\tau_k} \in \tilde{\mathcal{S}}$. Thus we obtain an infinite sequence $\{\tilde{S}_k\}$, which is a recurrent Markov chain on $\tilde{\mathcal{S}}$ with its transition matrix given by the corresponding submatrix of \tilde{P}^e .

Now in view of (A.2), almost surely,

$$\gamma_1^{\tau_1} = 1, \quad \gamma_{\tau_1+1}^{\tau_2} = 1, \quad \dots, \quad \gamma_{\tau_k+1}^{\tau_{k+1}} = 1, \quad \dots \quad (\text{A.4})$$

(recall that if $\tau_k + 1 > \tau_{k+1}$, $\gamma_{\tau_k+1}^{\tau_{k+1}} = 1$ by definition). Consider first a simpler case of the randomized stopping time τ that defines T : for any initial state S_0 , $\tau \geq 1$ a.s. Then, $\{\tau_k\}$ is strictly increasing, and by multiplying the variables in (A.4) together, we have $\gamma_1^\infty = 1$ a.s. For the general case of τ assumed in the theorem, $\tau = 0$ is possible, but $\mathbf{P}^\pi(\tau \geq 1 \mid S_0 = s) > 0$ for all states $s \in \mathcal{S}$. This means that the event of τ_k being the same for all k greater than some (random) \bar{k} has probability zero. So $\{\tau_k\}$ must converge to $+\infty$ almost surely, and we again obtain, by multiplying the variables in (A.4) together, that $\gamma_1^\infty = 1$ a.s. This contradicts the fact proved earlier, namely, that for any initial state S_0 , $\gamma_1^\infty = 0$ a.s. So the assumption of a recurrent class $\tilde{\mathcal{S}} \subset \mathcal{S}$ for \tilde{P}^e must be false. This proves Theorem 3.1.

Appendix B. Oblique Projection Viewpoint and Error Bound for TD

In this appendix we first explain Scherrer's interpretation of TD solutions as oblique projections (Scherrer, 2010), and we then give approximation error bounds for TD similar to those

given by Yu and Bertsekas (2010), which do not rely on contraction properties. We will explain these properties of TD in the context of generalized Bellman operators discussed in this paper. Although this was not the framework used in (Scherrer, 2010; Yu and Bertsekas, 2010) and our setup here is more general than the one discussed in those previous papers, the arguments and reasoning are essentially the same.

B.1 Solutions of TD as Oblique Projections of the Value Function

Let us start with the projected Bellman equation associated with TD/LSTD that we noted in Remark 3.2:

$$Tv - v \perp_{\zeta_S} \mathcal{L}_\phi, \quad v \in \mathcal{L}_\phi,$$

where T is a generalized Bellman operator with v_π as its unique fixed point, and \mathcal{L}_ϕ is the approximation subspace. We can write this equation equivalently as

$$v = \Pi_{\zeta_S} T v, \tag{B.1}$$

where Π_{ζ_S} denotes the projection onto the approximation subspace \mathcal{L}_ϕ with respect to the ζ_S -weighted Euclidean norm. In the subsequent derivations, we will not use the fact that ζ_S is the invariant probability measure induced by the behavior policy on \mathcal{S} , so the analyses we give in this appendix apply to any weighted Euclidean norm $\|\cdot\|_{\zeta_S}$.

Scherrer (2010) first realized that the solution of the projected Bellman equation (B.1) can be viewed as an *oblique projection of the value function* v_π on the approximation subspace \mathcal{L}_ϕ . This viewpoint provides an intuitive geometric interpretation of the TD solution and explains conceptually the source of its approximation bias. Analytically, this view also gives tight bounds on the approximation bias, as we will elaborate later in Section B.2.

An oblique projection is defined by two nonorthogonal subspaces of equal dimensions: it is the projection onto the first subspace orthogonally to the second (Saad, 2003), as illustrated in Figure 16. More precisely, for any two n -dimensional subspaces $\mathcal{L}_1, \mathcal{L}_2$ of \mathbb{R}^N such that no vector in \mathcal{L}_2 is orthogonal to \mathcal{L}_1 , there is an associated *oblique projection operator* $\Pi_{\mathcal{L}_1 \mathcal{L}_2} : \mathbb{R}^N \rightarrow \mathcal{L}_1$ defined by

$$\Pi_{\mathcal{L}_1 \mathcal{L}_2} x \in \mathcal{L}_1, \quad x - \Pi_{\mathcal{L}_1 \mathcal{L}_2} x \perp \mathcal{L}_2, \quad \forall x \in \mathbb{R}^N. \tag{B.2}$$

If the two subspaces are the same: $\mathcal{L}_1 = \mathcal{L}_2$ or if x lies in \mathcal{L}_1 , then the oblique projection $\Pi_{\mathcal{L}_1 \mathcal{L}_2} x$ is the same as $\Pi_{\mathcal{L}_1} x$, the orthogonal projection of x onto \mathcal{L}_1 . In general this need not be the case and $\Pi_{\mathcal{L}_1 \mathcal{L}_2} x \neq \Pi_{\mathcal{L}_1} x$ typically (cf. Figure 16). A matrix representation of the projection operator $\Pi_{\mathcal{L}_1 \mathcal{L}_2}$ is given by

$$\Pi_{\mathcal{L}_1 \mathcal{L}_2} = \Phi_1 (\Phi_2^\top \Phi_1)^{-1} \Phi_2^\top, \tag{B.3}$$

where Φ_1 and Φ_2 are $N \times n$ matrices whose columns form a basis of \mathcal{L}_1 and \mathcal{L}_2 , respectively (see Saad 2003, Chap. 1.12). For comparison, a matrix representation of the orthogonal projection operator $\Pi_{\mathcal{L}_1}$ is $\Pi_{\mathcal{L}_1} = \Phi_1 (\Phi_1^\top \Phi_1)^{-1} \Phi_1^\top$. (Below we will use the same notation for a projection operator and its matrix representations.)

Back to the projected Bellman equation (B.1), let us assume it has a unique solution v_{TD} and express v_{TD} in terms of v_π . We have $v_{\text{TD}} = \Pi_{\zeta_S} T v_{\text{TD}}$. By Theorem 3.1, v_π is the




Figure 16: Oblique projection of x onto the subspace \mathcal{L}_1 orthogonally to the subspace \mathcal{L}_2 .

unique solution of $v = Tv$. Recall that the generalized Bellman operator T is affine and can be expressed as $Tv = \tilde{r}_\pi + \tilde{P}_\pi v$ for some vector \tilde{r}_π and substochastic matrix \tilde{P}_π . It follows that $\tilde{r}_\pi = (I - \tilde{P}_\pi) v_\pi$ and

$$(I - \Pi_{\zeta_S} \tilde{P}_\pi) v_{\text{TD}} = \Pi_{\zeta_S} \tilde{r}_\pi = \Pi_{\zeta_S} (I - \tilde{P}_\pi) v_\pi. \quad (\text{B.4})$$

Let the columns of Φ form a basis of the approximation subspace \mathcal{L}_ϕ , and let D be the diagonal matrix with ζ_S as its diagonal elements. Then a matrix representation of the projection operator Π_{ζ_S} is $\Pi_{\zeta_S} = \Phi (\Phi^\top D \Phi)^{-1} \Phi^\top D$. Using this representation of Π_{ζ_S} and the fact $v_{\text{TD}} \in \mathcal{L}_\phi$, we obtain from (B.4) an expression of v_{TD} in terms of v_π :

$$v_{\text{TD}} = \Phi (\Phi^\top D (I - \tilde{P}_\pi) \Phi)^{-1} \Phi^\top D (I - \tilde{P}_\pi) v_\pi. \quad (\text{B.5})$$

(Here the invertibility of the matrix $\Phi^\top D (I - \tilde{P}_\pi) \Phi$ is equivalent to our assumption that v_{TD} is the unique solution of (B.1).)

Let us compare the expression (B.5) with (B.3). We see that if the geometry on \mathfrak{R}^N is determined by the usual Euclidean norm $\|\cdot\|_2$, then v_{TD} is an oblique projection of v_π for the two subspaces $\mathcal{L}_1 = \mathcal{L}_\phi$ and $\mathcal{L}_2 = \text{column-space}((I - \tilde{P}_\pi)^\top D \Phi)$.

Alternatively, consider the case where the geometry on \mathfrak{R}^N is determined by some weighted Euclidean norm $\|\cdot\|_\xi$ with weights ξ (for example, $\xi = \zeta_S$, which is one of the cases of interest in off-policy learning). In this case we can first scale each coordinate by the square root $\sqrt{\xi(s)}$ of its weight to reduce the case to that of the norm $\|\cdot\|_2$. Specifically, let \perp_ξ denote orthogonality with respect to $\|\cdot\|_\xi$, and let Ξ denote the diagonal matrix that has ξ as its diagonal elements. For the linear mapping $h : v \mapsto \Xi^{1/2} v$, we have

$$\|v\|_\xi = \|h(v)\|_2, \quad \text{and} \quad v_1 \perp_\xi v_2 \Leftrightarrow h(v_1) \perp h(v_2). \quad (\text{B.6})$$

The second relation means that \bar{v} is an oblique projection of v for two subspaces $\mathcal{L}_1, \mathcal{L}_2$ with respect to $\|\cdot\|_\xi$ (i.e., $\bar{v} \in \mathcal{L}_1$ and $v - \bar{v} \perp_\xi \mathcal{L}_2$), if and only if $h(\bar{v})$ is an oblique projection of $h(v)$ for the two subspaces $h(\mathcal{L}_1), h(\mathcal{L}_2)$ with respect to $\|\cdot\|_2$:

$$h(\bar{v}) \in h(\mathcal{L}_1), \quad h(v) - h(\bar{v}) \perp h(\mathcal{L}_2).$$

With these facts in mind, we rewrite (B.5) equivalently as follows:

$$\begin{aligned} \Xi^{1/2} v_{\text{TD}} &= \Xi^{1/2} \Phi \cdot (\Phi^\top D(I - \tilde{P}_\pi) \Xi^{-1} \Xi^{1/2} \cdot \Xi^{1/2} \Phi)^{-1} \cdot \Phi^\top D(I - \tilde{P}_\pi) \Xi^{-1} \Xi^{1/2} \cdot \Xi^{1/2} v_\pi \\ \implies h(v_{\text{TD}}) &= h(\Phi) \cdot (h(\Phi_2)^\top \cdot h(\Phi))^{-1} h(\Phi_2)^\top \cdot h(v_\pi), \end{aligned} \quad (\text{B.7})$$

where

$$\Phi_2 = \Xi^{-1}(I - \tilde{P}_\pi)^\top D\Phi \quad (\text{B.8})$$

and h applied to a matrix denotes the result of applying h to each column of that matrix. Comparing (B.7) with (B.3), we see that $h(v_{\text{TD}})$ is an oblique projection of $h(v_\pi)$ with respect to $\|\cdot\|_2$ for the two subspaces $h(\mathcal{L}_1)$, $h(\mathcal{L}_2)$, where

$$\mathcal{L}_1 = \mathcal{L}_\phi = \text{column-space}(\Phi), \quad \mathcal{L}_2 = \text{column-space}(\Phi_2). \quad (\text{B.9})$$

So based on the discussion earlier, with respect to the weighted Euclidean norm $\|\cdot\|_\xi$ on \mathbb{R}^N , v_{TD} is an oblique projection of v_π for the two subspaces $\mathcal{L}_1, \mathcal{L}_2$.

Note that by (B.8), the second subspace \mathcal{L}_2 defining the above oblique projection is the image of the approximation subspace \mathcal{L}_ϕ under the linear transformation $\Xi^{-1}(I - \tilde{P}_\pi)^\top D$. Thus \mathcal{L}_2 depends on the dynamics induced by the target policy as well as the generalized Bellman operator T that we choose. Relating the oblique projection interpretation of v_{TD} to Figure 16, we can see where the approximation bias of TD, $v_{\text{TD}} - \Pi_\xi v_\pi$, comes from.

B.2 Approximation Error Bound

We now consider the approximation error of v_{TD} and use the oblique projection viewpoint to derive a sharp bound on the approximation bias $v_{\text{TD}} - \Pi_\xi v_\pi$. Before proceeding, however, let us first remind the reader that unless the norm $\|\cdot\|_{\zeta_S}$ for the projection operator Π_{ζ_S} is purposefully chosen, the composition of Π_{ζ_S} with a generalized Bellman operator T is usually not a contraction, and thus error bounds for projected generalized Bellman equations usually cannot be obtained with contraction-based arguments. This is the case even for on-policy learning, as the following example shows.

Example B.1 (Non-contractive $\Pi_{\zeta_S} T$) If the target policy π induces an irreducible Markov chain with invariant probability measure ζ_S , and if T is the Bellman operator for TD(λ) with a constant λ , then, as Tsitsiklis and Van Roy (1997) showed, $\Pi_{\zeta_S} T$ is a contraction operator w.r.t. the weighted Euclidean norm $\|\cdot\|_{\zeta_S}$. Consequently, the matrix $\Phi^\top D(\tilde{P}_\pi - I) \Phi$ associated with the TD(λ) algorithm is negative definite (Tsitsiklis and Van Roy, 1997). The derivation of the contraction property of $\Pi_{\zeta_S} T$ in this case relies critically on the inequality $\zeta_S^\top \tilde{P}_\pi < \zeta_S^\top$. This inequality generally does not hold for the substochastic matrix \tilde{P}_π in the generalized Bellman operator T , when λ is not constant. So, for non-constant λ , we can no longer expect $\Pi_{\zeta_S} T$ to be a contraction or the matrix $\Phi^\top D(\tilde{P}_\pi - I) \Phi$ to be negative definite.

As an example, consider a simple two-state problem in which the system under the target policy π moves from one state to another in a cycle. Let $\pi^o = \pi$, let the discount factor γ be a constant, and let λ be a function of states with $\lambda(1) = 0$, $\lambda(2) = 1$. Then $\zeta_S^\top = (0.5, 0.5)$ and $\tilde{P}_\pi = \begin{pmatrix} \gamma^2 & 0 \\ \gamma & 0 \end{pmatrix}$. For γ near 1, e.g., $\gamma = 0.95$, and for Φ as given below,

we can calculate the ζ_S -weighted norm of $\Pi_{\zeta_S} \tilde{P}_\pi$ and the matrix associated with $\text{TD}(\lambda)$:

$$\Phi = \begin{pmatrix} 3 & 1 \\ 1 & 1 \end{pmatrix}, \quad \|\Pi_{\zeta_S} \tilde{P}_\pi\|_{\zeta_S} = \|\tilde{P}_\pi\|_{\zeta_S} \approx 1.31 > 1,$$

$$\Phi^\top D(\tilde{P}_\pi - I) \Phi = \begin{pmatrix} 0.4862 & -0.1713 \\ 0.7787 & -0.0738 \end{pmatrix}.$$

The latter matrix is not negative definite; in fact, its eigenvalues have positive real parts, so it is not even a Hurwitz matrix ($\text{TD}(\lambda)$ can diverge in this case). In the above, we have $\Pi_{\zeta_S} \tilde{P}_\pi = \tilde{P}_\pi$, and while $\Pi_{\zeta_S} \tilde{P}_\pi$ is not a contraction w.r.t. $\|\cdot\|_{\zeta_S}$, it is still a contraction w.r.t. some matrix norm. If we now let $\Phi = (3, 1)^\top$ instead, then the spectral radius of the matrix $\Pi_{\zeta_S} \tilde{P}_\pi$ comes out as $\sigma(\Pi_{\zeta_S} \tilde{P}_\pi) \approx 1.10 > 1$, so $\Pi_{\zeta_S} \tilde{P}_\pi$ (and hence $\Pi_{\zeta_S} T$) cannot be a contraction w.r.t. any matrix norm. \blacksquare

We now proceed to bound the bias term $v_{\text{TD}} - \Pi_\xi v_\pi$ relative to $\|v_\pi - \Pi_\xi v_\pi\|_\xi$, the distance between v_π and the approximation subspace measured with respect to $\|\cdot\|_\xi$. It is more transparent to derive the bound for the case of a general oblique projection operator $\Pi_{\mathcal{L}_1 \mathcal{L}_2}$ with respect to the usual Euclidean norm $\|\cdot\|_2$, so let us do that first and then use the linear transformation $h(\cdot)$ to translate the result to TD, as we did earlier in the preceding subsection.

We bound the bias $\|\Pi_{\mathcal{L}_1 \mathcal{L}_2} x - \Pi_{\mathcal{L}_1} x\|_2$ relative to $\|x - \Pi_{\mathcal{L}_1} x\|_2$, by calculating

$$\kappa := \sup_{x \in \mathbb{R}^N} \frac{\|\Pi_{\mathcal{L}_1 \mathcal{L}_2} x - \Pi_{\mathcal{L}_1} x\|_2}{\|x - \Pi_{\mathcal{L}_1} x\|_2} = \sup_{x \in \mathbb{R}^N} \frac{\|\Pi_{\mathcal{L}_1 \mathcal{L}_2}(x - \Pi_{\mathcal{L}_1} x)\|_2}{\|x - \Pi_{\mathcal{L}_1} x\|_2} \quad (\text{B.10})$$

(where we treat $0/0 = 0$). This constant κ depends on the two subspaces \mathcal{L}_1 , \mathcal{L}_2 , and reflects the “angle” between them. It has several equivalent expressions, e.g.,

$$\kappa = \sup_{x \perp \mathcal{L}_1, \|x\|_2=1} \|\Pi_{\mathcal{L}_1 \mathcal{L}_2} x\|_2 = \|\Pi_{\mathcal{L}_1 \mathcal{L}_2}(I - \Pi_{\mathcal{L}_1})\|_2, \quad (\text{B.11})$$

or with $\sigma(F)$ denoting the spectral radius of a square matrix F ,

$$\kappa = \sqrt{\sigma\left(\Pi_{\mathcal{L}_1 \mathcal{L}_2}(I - \Pi_{\mathcal{L}_1}) \cdot (I - \Pi_{\mathcal{L}_1})^\top \Pi_{\mathcal{L}_1 \mathcal{L}_2}^\top\right)} = \sqrt{\sigma(\Pi_{\mathcal{L}_1 \mathcal{L}_2} \Pi_{\mathcal{L}_2 \mathcal{L}_1} - \Pi_{\mathcal{L}_1})}. \quad (\text{B.12})$$

(After the definition of κ , each expression of κ in (B.10)-(B.12) follows from the preceding one; in particular, for the last expression in (B.12), we used the fact that $\Pi_{\mathcal{L}_1}^\top = \Pi_{\mathcal{L}_1}$, $\Pi_{\mathcal{L}_1 \mathcal{L}_2}^\top = \Pi_{\mathcal{L}_2 \mathcal{L}_1}$, and $\Pi_{\mathcal{L}_1} \Pi_{\mathcal{L}_2 \mathcal{L}_1} = \Pi_{\mathcal{L}_1}$.)

We can express κ in terms of the spectral radius of an $n \times n$ matrix, similarly to what was done in (Yu and Bertsekas, 2010). In particular, we take the last expression of κ in (B.12) and rewrite the symmetric matrix in that expression using the matrix representations of the projection operators as follows:

$$\Pi_{\mathcal{L}_1 \mathcal{L}_2} \Pi_{\mathcal{L}_2 \mathcal{L}_1} - \Pi_{\mathcal{L}_1} = \Phi_1(\Phi_2^\top \Phi_1)^{-1} \Phi_2^\top \cdot \Phi_2(\Phi_1^\top \Phi_2)^{-1} \Phi_1^\top - \Phi_1(\Phi_1^\top \Phi_1)^{-1} \Phi_1^\top.$$

By a result in matrix theory (Horn and Johnson, 1985, Theorem 1.3.20), for any $N \times n$ matrix F_1 and $n \times N$ matrix F_2 , $\sigma(F_1 F_2) = \sigma(F_2 F_1)$. Applying this result to the preceding

expression with $F_1 = \Phi_1$, we have that $\sigma(\Pi_{\mathcal{L}_1\mathcal{L}_2}\Pi_{\mathcal{L}_2\mathcal{L}_1} - \Pi_{\mathcal{L}_1})$ is equal to the spectral radius of the matrix

$$(\Phi_2^\top \Phi_1)^{-1}(\Phi_2^\top \Phi_2)(\Phi_1^\top \Phi_2)^{-1}(\Phi_1^\top \Phi_1) - I.$$

Combing this with (B.12), we obtain that

$$\kappa^2 = \sigma(F) - 1, \quad \text{where } F = (\Phi_2^\top \Phi_1)^{-1}(\Phi_2^\top \Phi_2)(\Phi_1^\top \Phi_2)^{-1}(\Phi_1^\top \Phi_1). \quad (\text{B.13})$$

Thus, for the above κ , the bound below holds for all $x \in \mathbb{R}^N$ and with equality attained at some x :

$$\|\Pi_{\mathcal{L}_1\mathcal{L}_2}x - \Pi_{\mathcal{L}_1}x\|_2 \leq \kappa \|x - \Pi_{\mathcal{L}_1}x\|_2. \quad (\text{B.14})$$

We now translate the result (B.13)-(B.14) to our TD context. We want to bound the relative bias $\|v_{\text{TD}} - \Pi_\xi v_\pi\|_\xi / \|v_\pi - \Pi_\xi v_\pi\|_\xi$. As discussed earlier in Section B.1, we can replace $\|\cdot\|_\xi$ with $\|\cdot\|_2$ by using the linear transformation $h(\cdot)$ to scale the coordinates. In particular, by the two relations given in (B.6),

$$\frac{\|v_{\text{TD}} - \Pi_\xi v_\pi\|_\xi}{\|v_\pi - \Pi_\xi v_\pi\|_\xi} = \frac{\|h(v_{\text{TD}}) - \Pi h(v_\pi)\|_2}{\|h(v_\pi) - \Pi h(v_\pi)\|_2}, \quad (\text{B.15})$$

where Π on the r.h.s. stands for the orthogonal projection onto $h(\mathcal{L}_\phi)$ with respect to $\|\cdot\|_2$. As shown by (B.7), $h(v_{\text{TD}})$ is an oblique projection of $h(v_\pi)$ for the two subspaces $h(\mathcal{L}_\phi)$ and $h(\mathcal{L}_2)$, where \mathcal{L}_2 is given by (B.8)-(B.9). According to (B.13), the constant κ for this oblique projection is $\sqrt{\sigma(F) - 1}$ where, if we take $\Phi_1 = \Phi$ and Φ_2 as defined by (B.8), F is now given by the expression in (B.13) with $h(\Phi_1), h(\Phi_2)$ in place of Φ_1, Φ_2 , respectively. Thus, by (B.14) and (B.15) we obtain that

$$\|v_{\text{TD}} - \Pi_\xi v_\pi\|_\xi \leq \kappa \|v_\pi - \Pi_\xi v_\pi\|_\xi \quad \text{for } \kappa = \sqrt{\sigma(F) - 1}, \quad (\text{B.16})$$

where F is an $n \times n$ matrix given by

$$F = [h(\Phi_2)^\top h(\Phi_1)]^{-1} \cdot [h(\Phi_2)^\top h(\Phi_2)] \cdot [h(\Phi_1)^\top h(\Phi_2)]^{-1} \cdot [h(\Phi_1)^\top h(\Phi_1)] \quad (\text{B.17})$$

for

$$\Phi_1 = \Phi, \quad \Phi_2 = \Xi^{-1}(I - \tilde{P}_\pi)^\top D\Phi, \quad (\text{B.18})$$

or more explicitly, after substituting the expressions of $h(\Phi_1)$ and $h(\Phi_2)$ in the formula of F and removing h , we have

$$F = (\Psi^\top \Phi)^{-1}(\Psi^\top \Xi^{-1} \Psi)(\Phi^\top \Psi)^{-1}(\Phi^\top \Xi \Phi), \quad \text{where } \Psi = (I - \tilde{P}_\pi)^\top D\Phi.$$

Note that by the definition of κ , the bound (B.16) is a worst-case bound that depends only on the two subspaces involved in the oblique projection operator. In other words, given the approximation subspace \mathcal{L}_ϕ , the dynamics described by \tilde{P}_π , the projection norm $\|\cdot\|_{\zeta_S}$ and the norm $\|\cdot\|_\xi$ for measuring the approximation quality, the bound (B.16) is attained by a worst-case choice of the rewards \tilde{r}_π for the target policy. In this sense, the bound (B.16) is tight.

References

- D. P. Bertsekas and J. N. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, Belmont, MA, 1996.
- V. S. Borkar. *Stochastic Approximation: A Dynamic Viewpoint*. Cambridge University Press, Cambridge, 2008.
- J. A. Boyan. Least-squares temporal difference learning. In *The 16th Int. Conf. Machine Learning (ICML)*, 1999.
- B. Dai, A. Shaw, L. Li, L. Xiao, N. He, Z. Liu, J. Chen, and L. Song. SBEED: Convergent reinforcement learning with nonlinear function approximation. In *The 35th Int. Conf. Machine Learning (ICML)*, 2018.
- C. Dann, G. Neumann, and J. Peters. Policy evaluation with temporal differences: A survey and comparison. *Journal of Machine Learning Res.*, 15:809–883, 2014.
- R. M. Dudley. *Real Analysis and Probability*. Cambridge University Press, Cambridge, 2002.
- M. Geist and B. Scherrer. Off-policy learning with eligibility traces: A survey. *Journal of Machine Learning Res.*, 15:289–333, 2014.
- P. W. Glynn and D. L. Iglehart. Importance sampling for stochastic simulations. *Management Science*, 35:1367–1392, 1989.
- R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, Cambridge, UK, 1985.
- P. Karmakar and S. Bhatnagar. Two timescale stochastic approximation with controlled Markov noise and off-policy temporal difference learning. *Math. Oper. Res.*, 43(1):130–151, 2018.
- V. R. Konda. *Actor-Critic Algorithms*. PhD thesis, MIT, 2002.
- H. J. Kushner and G. G. Yin. *Stochastic Approximation and Recursive Algorithms and Applications*. Springer-Verlag, New York, 2nd edition, 2003.
- B. Liu, J. Liu, M. Ghavamzadeh, S. Mahadevan, and M. Petrik. Finite-sample analysis of proximal gradient TD algorithms. In *The 31st Conf. Uncertainty in Artificial Intelligence (UAI)*, 2015.
- H. R. Maei. *Gradient Temporal-Difference Learning Algorithms*. PhD thesis, University of Alberta, 2011.
- S. Mahadevan, B. Liu, P. Thomas, W. Dabney, S. Giguere, N. Jacek, I. Gemp, and J. Liu. Proximal reinforcement learning: A new theory of sequential decision making in primal-dual spaces, 2014. arXiv:1405.6757.

- A. R. Mahmood, H. van Hasselt, and R. S. Sutton. Weighted importance sampling for off-policy learning with linear function approximation. In *Advances in Neural Information Processing Systems (NIPS)* 27, 2014.
- A. R. Mahmood, H. Yu, and R. S. Sutton. Multi-step off-policy learning without importance-sampling ratios, 2017. arXiv:1702.03006.
- S. Mannor, R. Rubinstein, and Y. Gat. The cross entropy method for fast policy search. In *The 20th Int. Conf. Machine Learning (ICML)*, 2003.
- S. Meyn. Ergodic theorems for discrete time stochastic systems using a stochastic Lyapunov function. *SIAM J. Control Optim.*, 27:1409–1439, 1989.
- S. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability*. Cambridge University Press, Cambridge, 2nd edition, 2009.
- R. Munos, T. Stepleton, A. Harutyunyan, and M. G. Bellemare. Safe and efficient off-policy reinforcement learning. In *Advances in Neural Information Processing Systems (NIPS)* 29, 2016.
- E. Nummelin. *General Irreducible Markov Chains and Non-Negative Operators*. Cambridge University Press, Cambridge, 1984.
- D. Precup, R. S. Sutton, and S. Singh. Eligibility traces for off-policy policy evaluation. In *The 17th Int. Conf. Machine Learning (ICML)*, 2000.
- D. Precup, R. S. Sutton, and S. Dasgupta. Off-policy temporal-difference learning with function approximation. In *The 18th Int. Conf. Machine Learning (ICML)*, 2001.
- M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, New York, 1994.
- R. S. Randhawa and S. Juneja. Combining importance sampling and temporal difference control variates to simulate Markov chains. *ACM Trans. Modeling and Computer Simulation*, 14(1):1–30, 2004.
- W. Rudin. *Real and Complex Analysis*. McGraw-Hill, New York, 1966.
- Y. Saad. *Iterative Methods for Sparse Linear Systems*. SIAM, Philadelphia, 2nd edition, 2003.
- M. Schäl and W. Sudderth. Stationary policies and Markov policies in Borel dynamic programming. *Probability Theory and Related Fields*, 74:91–111, 1987.
- B. Scherrer. Should one compute the temporal difference fix point or minimize the Bellman residual? The unified oblique projection view. In *The 27th Int. Conf. Machine Learning (ICML)*, 2010.
- E. Seneta. *Non-negative Matrices and Markov Chains*. Springer, New York, 2006.

- R. S. Sutton. Learning to predict by the methods of temporal differences. *Machine Learning*, 3:9–44, 1988.
- R. S. Sutton. TD models: Modeling the world at a mixture of time scales. In *The 12th Int. Conf. Machine Learning (ICML)*, 1995.
- R. S. Sutton. The grand challenge of predictive empirical abstract knowledge. In *IJCAI Workshop on Grand Challenges for Reasoning from Experiences*, 2009.
- R. S. Sutton and A. G. Barto. *Reinforcement Learning*. MIT Press, Cambridge, MA, 1998.
- R. S. Sutton, C. Szepesvári, and H. Maei. A convergent $O(n)$ algorithm for off-policy temporal-difference learning with linear function approximation. In *Advances in Neural Information Processing Systems (NIPS) 21*, 2008.
- R. S. Sutton, H. R. Maei, D. Precup, S. Bhatnagar, D. Silver, C. Szepesvári, and E. Wiewiora. Fast gradient-descent methods for temporal-difference learning with linear function approximation. In *The 26th Int. Conf. Machine Learning (ICML)*, 2009.
- R. S. Sutton, A. R. Mahmood, and M. White. An emphatic approach to the problem of off-policy temporal-difference learning. *Journal of Machine Learning Res.*, 17(73):1–29, 2016.
- J. N. Tsitsiklis. Asynchronous stochastic approximation and Q-learning. *Machine Learning*, 16:185–202, 1994.
- J. N. Tsitsiklis and B. Van Roy. An analysis of temporal-difference learning with function approximation. *IEEE Trans. Autom. Control*, 42(5):674–690, 1997.
- T. Ueno, S. Maeda, M. Kawanabe, and S. Ishii. Generalized TD learning. *Journal of Machine Learning Res.*, 12:1977–2020, 2011.
- C. J. C. H. Watkins. *Learning from Delayed Rewards*. PhD thesis, Cambridge Univ., England, 1989.
- H. Yu. Least squares temporal difference methods: An analysis under general conditions. *SIAM J. Control Optim.*, 50:3310–3343, 2012.
- H. Yu. On convergence of emphatic temporal-difference learning. In *The 28th Ann. Conf. Learning Theory (COLT)*, 2015. A longer version at arXiv:1506.02582.
- H. Yu. Some simulation results for emphatic temporal-difference learning algorithms, 2016a. arxiv:1605.02099.
- H. Yu. Weak convergence properties of constrained emphatic temporal-difference learning with constant and slowly diminishing stepsize. *Journal of Machine Learning Res.*, 17(220):1–58, 2016b.
- H. Yu. On convergence of some gradient-based temporal-differences algorithms for off-policy learning, 2017. arxiv:1712.09652.

- H. Yu and D. P. Bertsekas. Error bounds for approximations from projected linear equations. *Math. Oper. Res.*, 35(2):306–329, 2010.
- H. Yu and D. P. Bertsekas. Weighted Bellman equations and their applications in approximate dynamic programming. LIDS Technical Report 2876, MIT, 2012.