Machine Learning, -, 1–15 (1999)
© 1999 Kluwer Academic Publishers, Boston. Manufactured in The Netherlands.

Technical Update: Least-Squares Temporal Difference Learning

JUSTIN A. BOYAN jboyan@arc.nasa.gov
NASA Ames Research Center, Moffett Field, CA 94035

Received —, 1999

Editor: Satinder Singh

Abstract. TD(λ) is a popular family of algorithms for approximate policy evaluation in large MDPs. TD(λ) works by incrementally updating the value function after each observed transition. It has two major drawbacks: it may make inefficient use of data, and it requires the user to manually tune a stepsize schedule for good performance. For the case of linear value function approximations and $\lambda = 0$, the Least-Squares TD (LSTD) algorithm of Bradtke and Barto [5] eliminates all stepsize parameters and improves data efficiency.

This paper updates Bradtke and Barto's work in three significant ways. First, it presents a simpler derivation of the LSTD algorithm. Second, it generalizes from $\lambda = 0$ to arbitrary values of λ; at the extreme of $\lambda = 1$, the resulting new algorithm is shown to be a practical, incremental formulation of supervised linear regression. Third, it presents a novel and intuitive interpretation of LSTD as a *model-based* reinforcement learning technique.

Keywords: reinforcement learning, temporal difference learning, value function approximation, linear least-squares methods

1. Background

We address the problem of approximating the value function V^π of a fixed policy π in a large Markov decision process [2, 10]. This is an important subproblem of several algorithms for sequential decision making, including optimistic policy iteration [2] and STAGE [3]. $V^\pi(x)$ simply predicts the expected long-term sum of future rewards obtained when the process starts in state x and follows policy π until termination. For simplicity we will assume that π is proper (guaranteed to terminate), which implies that V^π is well-defined without the use of a discount factor; however, the algorithms presented here extend straightforwardly to the discounted case.

For small Markov chains whose transition probabilities are all explicitly known, computing V^π is a trivial matter of solving a system of linear equations. However, in many practical applications, the transition probabilities of the chain are available only implicitly—either in the form of a simulation model or in the form of an agent's actual experience executing π in its environment. In either case, we must compute V^π or an approximation thereof (denoted $\tilde{V}^\pi$) solely from a collection of trajectories sampled from the chain. This is where the TD(λ) family of algorithms applies.

TD(λ) was introduced in [11]; excellent summaries may now be found in several books [2, 10]. For each state on each observed trajectory, TD(λ) incrementally adjusts the coefficients of $\tilde{V}^\pi$ toward new target values. The target values depend on

the parameter $\lambda \in [0,1]$. At $\lambda = 1$, the target at each visited state x_t is the *Monte-Carlo return*, i.e., the actual observed sum of future rewards $R_t + R_{t+1} + \cdots + R_{\text{END}}$. This is an unbiased sample of $V^\pi(x_t)$, but may have significant variance since it depends on a long stochastic sequence of rewards. At the other extreme, $\lambda = 0$, the target value is set by a sampled one-step lookahead: $R_t + \tilde{V}^\pi(x_{t+1})$. This value has lower variance—the only random component is a single state transition—but is biased by the potential inaccuracy of the lookahead estimate of V^π. The parameter λ trades off between bias and variance. Empirically, several studies have found intermediate values of λ to perform best [11, 10].

TD(λ) has been shown to converge to a good approximation of V^π when *linear architectures* are used, assuming a suitable decreasing schedule of stepsizes for the incremental weight updates [16]. Linear architectures—which include lookup tables, state aggregation methods, CMACs, radial basis function networks with fixed bases, and multi-dimensional polynomial regression—approximate $V^\pi(x)$ by first mapping the state x to a feature vector $\boldsymbol{\phi}(x) \in \Re^K$, and then computing a linear combination of those features, $\boldsymbol{\phi}(x)^\mathsf{T}\boldsymbol{\beta}$. Figure 1 gives a convenient form of TD(λ) that exploits this representation. Table 1 summarizes the notation used throughout this paper.

On each transition, the algorithm computes the scalar one-step TD error $R_t + (\boldsymbol{\phi}(x_{t+1}) - \boldsymbol{\phi}(x_t))^\mathsf{T}\boldsymbol{\beta}$, and apportions that error among all state features according to their respective *eligibilities* $\mathbf{z}_t$. The eligibility vector may be seen as an algebraic trick by which TD(λ) propagates rewards backward over the current trajectory without having to remember the trajectory explicitly. Each feature's eligibility at time t depends on the trajectory's history and on λ: $\mathbf{z}_t = \sum_{i=t_0}^{t} \lambda^{t-i}\boldsymbol{\phi}(x_i)$, where t_0 is the time at which the current trajectory started. In the case of TD(0), only the current state's features are eligible to be updated, so $\mathbf{z}_t = \boldsymbol{\phi}(x_t)$; whereas in TD(1), the features of all states seen so far on the current trajectory are eligible, so $\mathbf{z}_t = \sum_{i=t_0}^{t} \boldsymbol{\phi}(x_i)$.

Table 1. A glossary of notation used in this paper.

Symbol	Type	Meaning
$\mathbf{X}$	set	state space for a Markov Decision Problem (MDP)
R_t	scalar	the one-step reward received at time t
V^π	$: X \to \Re$	the true value function of a given policy π in the MDP
$\tilde{V}^\pi$	$: X \to \Re$	a linear approximation to V^π given by $\tilde{V}^\pi(x) = \boldsymbol{\phi}(x)^\mathsf{T}\boldsymbol{\beta}$
K	$\in \mathbb{N}$	dimensionality of the feature vector representing each state
$\boldsymbol{\phi}$	$: X \to \Re^K$	the mapping from state to real-valued feature vector
$\boldsymbol{\beta}$	$\in \Re^K$	the vector of coefficients that specifies $\tilde{V}^\pi$
λ	scalar $\in [0,1]$	bias-variance parameter for temporal-difference learning
α_i	scalar $\in (0,1)$	stepsize for incremental coefficient updating in TD(λ)
$\mathbf{z}_t$	$\in \Re^K$	"eligibility vectors" used to trace a trajectory's history
$\boldsymbol{\beta}_\lambda$	$\in \Re^K$	the coefficients to which TD(λ) converges (dependent on λ)
$\mathbf{d}$	$\in \Re^K$	vector used in convergence analysis of TD(λ) (see Eq. 1)
$\mathbf{C}$	$K \times K$ matrix	matrix used in convergence analysis of TD(λ) (see Eq. 1)
$\mathbf{b}$	$\in \Re^K$	vector produced by LSTD(λ) for direct solution of $\boldsymbol{\beta}$ (see Eq. 2)
$\mathbf{A}$	$K \times K$ matrix	matrix produced by LSTD(λ) for direct solution of $\boldsymbol{\beta}$ (see Eq. 2)

TD($\boldsymbol{\lambda}$) for approximate policy evaluation:

Given: • a *simulation model* for a proper policy π in MDP X;
• a *featurizer* $\boldsymbol{\phi} : X \to \Re^K$ mapping states to feature vectors, $\boldsymbol{\phi}(\text{END}) \stackrel{\text{def}}{=} \mathbf{0}$;
• a parameter $\lambda \in [0, 1]$; and
• a sequence of *stepsizes* $\alpha_1, \alpha_2, \ldots$ for incremental coefficient updating.

Output: a coefficient vector $\boldsymbol{\beta}$ for which $V^\pi(x) \approx \boldsymbol{\beta} \cdot \boldsymbol{\phi}(x)$.

Set $\boldsymbol{\beta} := \mathbf{0}$ *(or an arbitrary initial estimate)*, $t := 0$.
for $n := 1, 2, \ldots$ **do**: {
  Set $\boldsymbol{\delta} := 0$.
  Choose a start state $x_t \in X$.
  Set $\mathbf{z}_t := \boldsymbol{\phi}(x_t)$.
  while $x_t \neq$ END, **do**: {
    Simulate one step of the process, producing a reward R_t and next state x_{t+1}.
    Set $\boldsymbol{\delta} := \boldsymbol{\delta} + \mathbf{z}_t\big(R_t + (\boldsymbol{\phi}(x_{t+1}) - \boldsymbol{\phi}(x_t))^\mathsf{T}\boldsymbol{\beta}\big)$.
    Set $\mathbf{z}_{t+1} := \lambda\mathbf{z}_t + \boldsymbol{\phi}(x_{t+1})$.
    Set $t := t + 1$.
  }
  Set $\boldsymbol{\beta} := \boldsymbol{\beta} + \alpha_n\boldsymbol{\delta}$.
}

Figure 1. Ordinary TD(λ) for linearly approximating the undiscounted value function of a fixed proper policy.

To what weights does TD(λ) converge? Examining the update rule for $\boldsymbol{\delta}$ in Figure 1, it is not difficult to see that the coefficient changes made by TD(λ) after an observed trajectory $(x_0, x_1, \ldots, x_L, \text{END})$ have the form $\boldsymbol{\beta} := \boldsymbol{\beta} + \alpha_n(\mathbf{d} + \mathbf{C}\boldsymbol{\beta} + \boldsymbol{\omega})$, where

$$\mathbf{d} = \mathrm{E}\Big\{\sum_{i=0}^{L} \mathbf{z}_i R_i\Big\}; \qquad \mathbf{C} = \mathrm{E}\Big\{\sum_{i=0}^{L} \mathbf{z}_i\big(\boldsymbol{\phi}(x_{i+1}) - \boldsymbol{\phi}(x_i)\big)^\mathsf{T}\Big\}; \tag{1}$$

and $\boldsymbol{\omega}$ = zero-mean noise. The expectations are taken with respect to the distribution of trajectories through the Markov chain. It is shown in [2, §6.3.4] that $\mathbf{C}$ is negative definite and that the noise $\boldsymbol{\omega}$ has sufficiently small variance, which together with the stepsize conditions mentioned above, imply that $\boldsymbol{\beta}$ converges to a fixed point $\boldsymbol{\beta}_\lambda$ satisfying $\mathbf{d} + \mathbf{C}\boldsymbol{\beta}_\lambda = \mathbf{0}$. In effect, TD($\lambda$) solves this system of equations by performing stochastic gradient descent on a potential function $\|\boldsymbol{\beta} - \boldsymbol{\beta}_\lambda\|^2$. It never explicitly represents $\mathbf{d}$ or $\mathbf{C}$. The changes to $\boldsymbol{\beta}$ depend only on the most recent trajectory, and after those changes are made, the trajectory and its rewards are simply forgotten. This approach, while requiring little computation per iteration, wastes data and may require sampling many trajectories to reach convergence.

One technique for using data more efficiently is "experience replay" [6]: explicitly remember all trajectories ever seen, and whenever asked to produce an updated set of coefficients, perform repeated passes of TD(λ) over all the saved trajectories until convergence. This technique is similar to the batch training methods

commonly used to train neural networks. However, in the case of linear function approximators, there is another way.

2. The Least-Squares TD(λ) Algorithm

The Least-Squares TD(λ) algorithm, or LSTD(λ), converges to the same coefficients $\boldsymbol{\beta}_\lambda$ that TD(λ) does. However, instead of performing gradient descent, LSTD(λ) builds explicit estimates of the $\mathbf{C}$ matrix and $\mathbf{d}$ vector (actually, estimates of a constant multiple of $\mathbf{C}$ and $\mathbf{d}$), and then solves $\mathbf{d} + \mathbf{C}\boldsymbol{\beta}_\lambda = \mathbf{0}$ directly. The actual data structures that LSTD(λ) builds from experience are the matrix $\mathbf{A}$ (of dimension $K \times K$, where K is the number of features) and the vector $\mathbf{b}$ (of dimension K):

$$\mathbf{b} = \sum_{i=0}^{t} \mathbf{z}_i R_i \qquad\qquad \mathbf{A} = \sum_{i=0}^{t} \mathbf{z}_i \big(\boldsymbol{\phi}(x_i) - \boldsymbol{\phi}(x_{i+1})\big)^{\mathsf{T}} \tag{2}$$

After n independent trajectories have been observed, $\mathbf{b}$ is an unbiased estimate of $n\mathbf{d}$, and $\mathbf{A}$ is an unbiased estimate of $-n\mathbf{C}$. Thus, $\boldsymbol{\beta}_\lambda$ can be estimated as $\mathbf{A}^{-1}\mathbf{b}$. As is standard in least-squares algorithms, Singular Value Decomposition is used to invert $\mathbf{A}$ robustly [8]. The complete LSTD(λ) algorithm is specified in Figure 2.

LSTD($\boldsymbol{\lambda}$) for approximate policy evaluation:

Given: a *simulation model*, *featurizer*, and λ as in ordinary TD(λ).
(No stepsize schedules or initial estimates of $\boldsymbol{\beta}$ are necessary.)

Output: a coefficient vector $\boldsymbol{\beta}$ for which $V^\pi(x) \approx \boldsymbol{\beta} \cdot \boldsymbol{\phi}(x)$.

Set $\mathbf{A} := \mathbf{0}$, $\mathbf{b} := \mathbf{0}$, $t := 0$.
for $n := 1, 2, \ldots$ **do**: {
  Choose a start state $x_t \in X$.
  Set $\mathbf{z}_t := \boldsymbol{\phi}(x_t)$.
  while $x_t \neq$ END, **do**: {
    Simulate one step of the chain, producing a reward R_t and next state x_{t+1}.
    Set $\mathbf{A} := \mathbf{A} + \mathbf{z}_t(\boldsymbol{\phi}(x_t) - \boldsymbol{\phi}(x_{t+1}))^{\mathsf{T}}$. */* outer product */*
    Set $\mathbf{b} := \mathbf{b} + \mathbf{z}_t R_t$.
    Set $\mathbf{z}_{t+1} := \lambda \mathbf{z}_t + \boldsymbol{\phi}(x_{t+1})$.
    Set $t := t + 1$.
  }
  Whenever updated coefficients are desired: Set $\boldsymbol{\beta} := \mathbf{A}^{-1}\mathbf{b}$. */* Use SVD. */*
}

Figure 2. A least-squares version of TD(λ) (compare Figure 1). Note that $\mathbf{A}$ has dimension $K \times K$, and $\mathbf{b}$, $\boldsymbol{\beta}$, $\mathbf{z}$, and $\boldsymbol{\phi}(x)$ all have dimension $K \times 1$.

When $\lambda = 0$, LSTD(0) reduces precisely to Bradtke and Barto's LSTD algorithm, which they derived using a more complex approach based on regression with instrumental variables [5]. At the other extreme, when $\lambda = 1$, LSTD(1) produces the

same $\mathbf{A}$ and $\mathbf{b}$ that would be produced by supervised linear regression on training pairs of {state features $\mapsto$ observed Monte-Carlo returns}; the proof is given in the appendix. Thanks to the algebraic trick of the eligibility vectors, LSTD(1) builds the regression matrices *fully incrementally*—without having to store the trajectory while waiting to observe the eventual outcome. When trajectories through the chain are long, this provides significant memory savings over linear regression.

The computation per timestep required to update $\mathbf{A}$ and $\mathbf{b}$ is the same as least-squares linear regression: $O(K^2)$, where K is the number of features. LSTD(λ) must also perform a matrix inversion at a cost of $O(K^3)$ whenever $\boldsymbol{\beta}$'s coefficients are needed—typically, once per complete trajectory. (If updated coefficients are required more frequently, then the $O(K^3)$ cost can be avoided by recursive-least-squares [5] or Kalman-filtering techniques [2, §3.2.2], which update $\boldsymbol{\beta}$ on each timestep at a cost of only $O(K^2)$.) LSTD(λ) performs more computation per observation than incremental TD(λ), which updates the coefficients using only $O(K)$ computation per timestep. However, as pointed out in [5], LSTD(λ) offers several significant advantages:

- Least-squares algorithms "extract more information from each additional observation" [5] and would thus be expected to converge with fewer training samples.
- TD(λ)'s convergence can be slowed dramatically by a poor choice of the stepsize parameters α_n. LSTD(λ) eliminates these parameters.
- TD(λ)'s performance is sensitive to the initial estimate for $\boldsymbol{\beta}_\lambda$. LSTD($\lambda$) does not rely on an arbitrary initial estimate.
- TD(λ) is also sensitive to the ranges of the individual features. LSTD(λ) is not.

Section 4 below presents experimental results comparing the data efficiency of gradient-based and least-squares-based TD learning.

3. LSTD(λ) as Model-Based Learning

Surprisingly, the move from an incremental, gradient-based update rule to a direct, least-squares-based update rule for TD(λ) turns out to be mathematically equivalent to a move from a model-free to a model-based reinforcement learning algorithm. This equivalence provides interesting new intuitions about the space of temporal difference learning algorithms.

To begin, let us restrict our attention to the case of a small discrete state space X, over which V^π can be represented and learned exactly by a lookup table. A classical model-based algorithm for learning V^π from simulated trajectory data would proceed as follows:

1. From the state transitions and rewards observed so far, build in memory an *empirical model* of the Markov chain. The sufficient statistics of this model are
   - a vector $\mathbf{n}$ recording the number of times each state has been visited;

- a matrix $\mathbf{T}$ recording the observed state-transition counts: $\mathbf{T}_{ij}$ = how many times x_j was seen to directly follow x_i; and
- a vector $\mathbf{s}$ recording, for each state, the sum of all one-step rewards observed on transitions leaving that state.

2. Whenever a new estimate of the value function V^π is desired, solve the linear system of Bellman equations corresponding to the current empirical model. Writing $\mathbf{N} = \mathbf{diag}(\mathbf{n})$, the solution vector of V^π values is given by

$$\mathbf{v} = (\mathbf{N} - \mathbf{T})^{-1}\mathbf{s}. \tag{3}$$

This model-based technique contrasts with TD(λ), a model-free approach to the same problem. TD(λ) does not maintain any statistics on observed transitions and rewards; it simply updates the components of $\mathbf{v}$ directly. In the limit, assuming a lookup-table representation, both converge to the optimal V^π. The advantage of TD(λ) is its low computational burden per step; the advantage of the classical model-based method is that it makes the most of the available training data. The relative advantages of model-based and model-free reinforcement learning methods have been investigated in, e.g., [12, 7, 1].

Where does LSTD(λ) fit in? In fact, for the case of $\lambda = 0$, it precisely duplicates the classical model-based method sketched above. The assumed lookup-table representation for $\tilde{V}^\pi$ means that we have one independent feature per state: the feature vector $\boldsymbol{\phi}$ corresponding to state 1 is $(1, 0, 0, \ldots, 0)$; corresponding to state 2 is $(0, 1, 0, \ldots, 0)$; etc. Referring to the algorithm of Figure 2, we see that LSTD(0) performs the following operations upon each observed transition:

$$\mathbf{b} := \mathbf{b} + \boldsymbol{\phi}(x_t)R_t \qquad\qquad \mathbf{A} := \mathbf{A} + \boldsymbol{\phi}(x_t)(\boldsymbol{\phi}(x_t) - \boldsymbol{\phi}(x_{t+1}))^\mathsf{T} \tag{4}$$

Clearly, the role of $\mathbf{b}$ is to sum all the rewards observed at each state, exactly as the vector $\mathbf{s}$ does in the classical technique. $\mathbf{A}$, meanwhile, accumulates the statistics $(\mathbf{N} - \mathbf{T})$. To see this, note that the outer product in Eq. 4 is a matrix consisting of an entry of $+1$ on the single diagonal element corresponding to state x_t; an entry of -1 on the element in row x_t, column x_{t+1}; and all the rest zeroes. Summing one such sparse matrix for each observed transition gives $\mathbf{A} \equiv \mathbf{N} - \mathbf{T}$. Finally, LSTD(0) performs the inversion $\boldsymbol{\beta} := \mathbf{A}^{-1}\mathbf{b} = (\mathbf{N} - \mathbf{T})^{-1}\mathbf{s}$, giving the same solution as in Equation 3.

Thus, when $\lambda = 0$, the $\mathbf{A}$ and $\mathbf{b}$ matrices built by LSTD(λ) effectively record a model of all the observed transitions. What about when $\lambda > 0$? Again, $\mathbf{A}$ and $\mathbf{b}$ record the sufficient statistics of an empirical Markov model—but in this case, the model being captured is one whose single-step transition probabilities directly encode the multi-step TD(λ) backup operations. That is, the model links each state x to all the downstream states that follow x on any trajectory, and records how much influence each has on estimating $\tilde{V}^\pi(x)$ according to TD(λ). In the case of $\lambda = 0$, the TD(λ) backups correspond to the one-step transitions, resulting in the equivalence described above. The opposite extreme, the case of $\lambda = 1$, is also interesting: the empirical Markov model corresponding to TD(1)'s backups is

the chain in which each state x leads directly to absorption, and $\boldsymbol{\beta}$ then simply computes the average Monte-Carlo return at each state. For general λ, we produce the statistics of the "simple beta-model" of multi-scale reinforcement learning [14]. In short, if we assume a lookup-table representation for the function $\tilde{V}^\pi$, we can view the LSTD(λ) algorithm as performing these two steps:

1. It implicitly uses the observed simulation data to build a Markov chain. This chain compactly models all the backups that TD(λ) would perform on the data.

2. It solves the chain by performing a matrix inversion.

The lookup-table representation for $\tilde{V}^\pi$ is intractable in practical problems; in practice, LSTD(λ) operates on states only via their (linearly dependent) feature representations $\boldsymbol{\phi}(x)$. In this case, LSTD(λ) implicitly builds a *compressed* version of the empirical model's transition matrix $\mathbf{N} - \mathbf{T}$ and summed-reward vector $\mathbf{s}$:

$$\mathbf{b} = \boldsymbol{\Phi}^\mathsf{T}\mathbf{s} \qquad\qquad \mathbf{A} = \boldsymbol{\Phi}^\mathsf{T}(\mathbf{N} - \mathbf{T})\boldsymbol{\Phi} \tag{5}$$

where $\boldsymbol{\Phi}$ is the $|X| \times K$ matrix representation of the function $\boldsymbol{\phi} : X \to \Re^K$. From the compressed empirical model, LSTD(λ) computes these coefficients for $\tilde{V}^\pi$:

$$\boldsymbol{\beta}_\lambda = \mathbf{A}^{-1}\mathbf{b} = (\boldsymbol{\Phi}^\mathsf{T}(\mathbf{N} - \mathbf{T})\boldsymbol{\Phi})^{-1}(\boldsymbol{\Phi}^\mathsf{T}\mathbf{s}). \tag{6}$$

Ideally, these coefficients $\boldsymbol{\beta}_\lambda$ would be equivalent to the *empirical optimal* coefficients $\boldsymbol{\beta}^*_\lambda$. The empirical optimal coefficients are those that would be found by building the full uncompressed empirical model (represented by $\mathbf{N} - \mathbf{T}$ and $\mathbf{s}$), using a lookup table to solve for that model's value function ($\mathbf{v} = (\mathbf{N} - \mathbf{T})^{-1}\mathbf{s}$), and then performing a least-squares linear fit from the state features $\boldsymbol{\Phi}$ to the lookup-table value function:

$$\boldsymbol{\beta}^*_\lambda \stackrel{\text{def}}{=} (\boldsymbol{\Phi}^\mathsf{T}\boldsymbol{\Phi})^{-1}(\boldsymbol{\Phi}^\mathsf{T}\mathbf{v}) = (\boldsymbol{\Phi}^\mathsf{T}\boldsymbol{\Phi})^{-1}\boldsymbol{\Phi}^\mathsf{T}(\mathbf{N} - \mathbf{T})^{-1}\mathbf{s}. \tag{7}$$

It can be shown that Equations 6 and 7 are indeed equivalent for the case of $\lambda = 1$, because that setting of λ implies that $\mathbf{T} = \mathbf{0}$, thus $(\mathbf{N} - \mathbf{T})^{-1}$ is diagonal and commutes. However, for the case of $\lambda < 1$, solving the compressed empirical model does not in general produce the optimal least-squares fit to the solution of the uncompressed model [16].

4. Experimental Comparison of TD(λ) and LSTD(λ)

This section reports experimental results comparing TD(λ) and LSTD(λ) on two domains: a simple illustrative Markov chain and the large-scale game of backgammon. We begin with the simple chain, pictured in Figure 3. It consists of 13 states, each of which is represented by four state features as shown. We seek to represent the value function compactly as a linear function of these four state features. In fact, this domain's optimal V^π function is exactly linear in these features: the optimal coefficients $\boldsymbol{\beta}^*_\lambda$ are $(-24, -16, -8, 0)$.

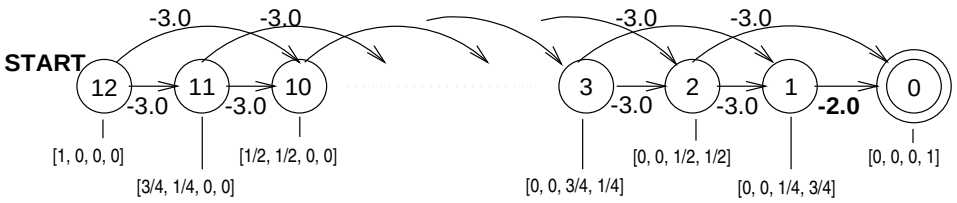

Figure 3. A 13-state Markov chain. In states 2–12, each outgoing arc is taken with probability 0.5. For value function approximation, each state is represented by four features, as follows: the representations for states 12, 8, 4, and 0 are, respectively, [1,0,0,0], [0,1,0,0], [0,0,1,0], and [0,0,0,1]; and the representations for the other states are obtained by linearly interpolating between these.

Because V^π is exactly linear in this domain, LSTD(λ) is guaranteed to converge with probability 1 to the optimal $\boldsymbol{\beta}_\lambda^*$ for any λ. TD(λ) is also guaranteed to converge to $\boldsymbol{\beta}_\lambda^*$, under the additional condition that an appropriate schedule of stepsizes is chosen. The following three criteria on the schedule (α_n) are sufficient: $\alpha_n \geq 0 \ \ \forall n$; $\sum_{n=1}^{\infty} \alpha_n = \infty$; and $\sum_{n=1}^{\infty} \alpha_n^2 < \infty$. Our experiments use schedules that satisfy these criteria, having the form:

$$\alpha_n \stackrel{\text{def}}{=} a_0 \frac{n_0 + 1}{n_0 + n} \qquad n = 1, 2, \ldots \tag{8}$$

The parameter a_0 determines the initial stepsize, and n_0 determines how gradually the stepsize decreases over time. Each TD(λ) experiment was run with six different stepsize schedules, corresponding to the six combinations of $a_0 \in \{0.1, 0.01\}$ and $n_0 \in \{10^2, 10^3, 10^6\}$. These settings produce a range of learning rates that is typical for applications of gradient descent. TD(λ) also requires an initial setting for the coefficients $\boldsymbol{\beta}$; we initialized $\boldsymbol{\beta}$ to $\mathbf{0}$ on each experiment.

Comparative results are given in Figures 4 and 5. Figure 4 focuses on the case of $\lambda = 0.4$, comparing the learning curve for LSTD(λ) against those of all six schedules of TD(λ). Each point plotted represents the average of 10 trials. The plot shows clearly that for $\lambda = 0.4$, LSTD(λ) learns a good approximation to V^π in fewer trials than any of the TD(λ) experiments, and performs better asymptotically as well.

Figure 5 graphically summarizes six learning-curve plots similar to Figure 4, corresponding to varying λ over the settings $\{0, 0.2, 0.4, 0.6, 0.8, 1.0\}$. The results may be summarized as follows:

- Across all values of λ, LSTD(λ) learns a good approximation to V^π in fewer trials than any of the TD(λ) experiments, and performs better asymptotically as well.

- The performance of TD(λ) depends critically on the stepsize schedule chosen. LSTD(λ) has no tunable parameters other than λ itself.

- Varying λ has a relatively small effect on LSTD(λ)'s performance.

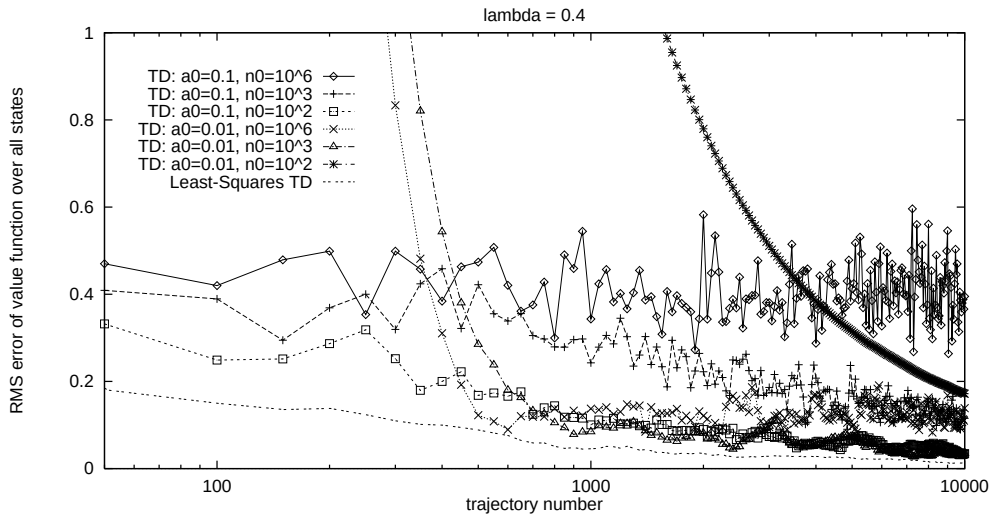

Figure 4. Performance of TD(0.4) and LSTD(0.4) on the sample domain. Note the log scale on the x-axis. All points plotted represent the average of 10 trials.

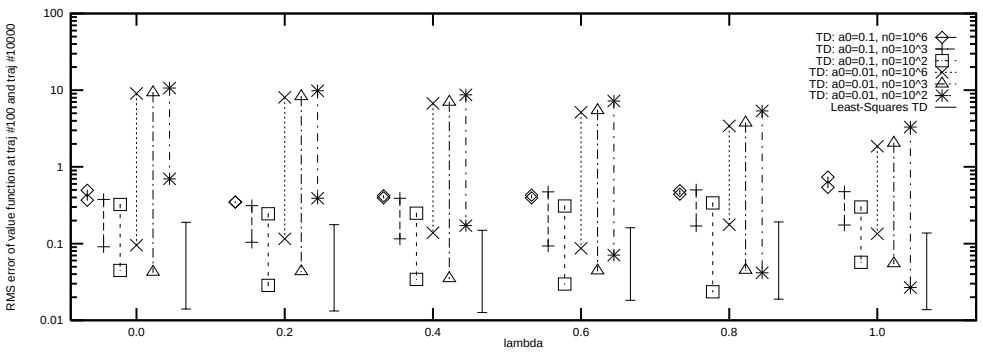

Figure 5. Summary of results at six settings of λ. At each setting, seven algorithms are compared: TD(λ) (with six different stepsize schedules) and LSTD(λ). The plotted segment shows the mean RMS value function approximation error after 100 trajectories (top of segment) and 10,000 trajectories (bottom of segment). Note the log scale on the y-axis. LSTD(λ) is best in all cases.

Large-Scale Example: Backgammon Policy Evaluation

We now turn to a larger-scale test domain, familiar in the reinforcement-learning literature: the game of backgammon. In this larger domain, we compare TD(λ) with LSTD(λ) in terms of overall computational efficiency as well as data efficiency.

We generate trajectories by opposing two fixed, pre-trained backgammon policies: "Fiona," a neural network patterned after TD-Gammon [15], and "pubeval," a benchmark backgammon evaluator.[1] We sidestep the two-player aspect of backgammon by defining a Markov chain consisting of only those states in which it is Fiona's turn to play. That is, each single transition covers two backgammon moves, one by Fiona and one by pubeval. The value function we seek to learn specifies, for any backgammon position, the expected winnings of Fiona over pubeval, starting from that position with Fiona to roll the dice and move first.

In the small-scale domain of the previous section, both TD(λ) and LSTD(λ) were known to converge to the optimal value function V^π, because V^π was exactly linear in that domain's features. In backgammon, by contrast, we cannot guarantee convergence to the optimal value function—nor do we even have that function available for measuring learning performance. Instead, we measure performance against a "gold standard" value function approximation $\tilde{G}(x)$, built by training LSTD(1) on a set of 100,000 sample games (amounting to nearly 3 million positions). The learning architecture for $\tilde{G}(x)$ was quadratic regression over 15 backgammon-specific high-level features.[2] To measure the performance of a new value function approximation $\tilde{V}^\pi$, then, we compute the sampled RMS difference between $\tilde{V}^\pi$ and $\tilde{G}$:

$$\text{Error}(\tilde{V}^\pi) \stackrel{\text{def}}{=} \left(\frac{1}{|S|}\sum_{x\in S}\left(\tilde{V}^\pi(x) - \tilde{G}(x)\right)^2\right)^{1/2}$$

over an independent test set S of 10,000 backgammon positions.

Our experiments compared TD(λ) with LSTD(λ) using both linear and quadratic approximations over the 15 backgammon features noted above. The coefficient vectors corresponding to these approximations comprise $K = 16$ and $K = 136$ parameters, respectively. The parameter λ was set arbitrarily to 0.9. For TD(λ), four different learning rate schedules were used: $a_0 \in \{10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}\}$, $n_0 = 1000$. The coefficients $\boldsymbol{\beta}$ were always initialized to $\mathbf{0}$. Performance measurements were taken after every 50 games. Each algorithm was trained on the same 9 sets of 10,000 games each, and the average error curves over these 9 runs are plotted.

Results are shown in Figure 6. The upper row of the figure shows the average error curves for the linear case (at left) and quadratic case (at right) as a function of the number of sample trajectories observed. As in the smaller-scale results presented earlier, we can see that LSTD(λ) brings down the approximation error using far less sample data than TD(λ). We also see again that TD(λ) is quite sensitive to the learning rate schedule chosen; in fact, in the quadratic case, the highest learning rate of $a_0 = 0.01$ produced error curves greater than the bounds of the plot.

The lower row of Figure 6 plots the same error curves, but with respect to elapsed time rather than the number of trajectories observed. Experiments were run on a 296 MHz Ultrasparc II. The linear regression results (at left) show that when $\boldsymbol{\beta}$

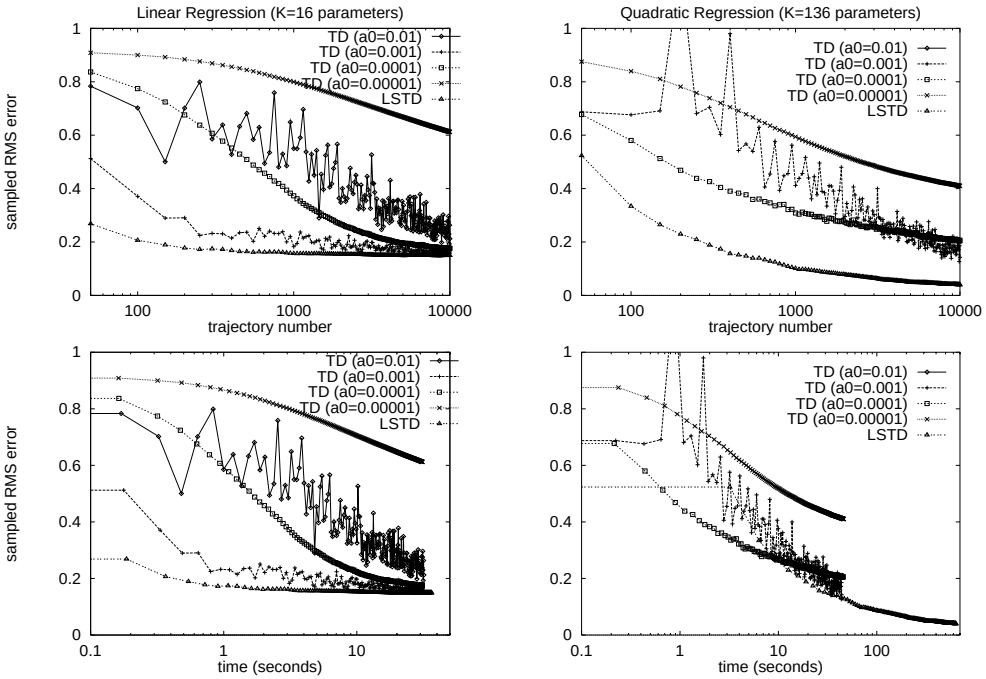

Figure 6. Results of TD(0.9) and LSTD(0.9) on backgammon policy evaluation. The upper two graphs plot learning curves with respect to the number of trajectories sampled, for linear and quadratic models, respectively. The lower two graphs plot the identical learning curves, but with respect to the amount of elapsed wall-clock time used by the learning algorithms (excluding time spent generating moves). All points plotted represent the average of 9 trials.

has relatively few parameters, the time penalty for using LSTD(λ) is slight. With quadratic regression (at right), by contrast, the large matrix inversions performed by LSTD(λ) cause it to run markedly more slowly than TD(λ) (note the semilog scale). Nevertheless, because of its better data efficiency, the RMS error of LSTD(λ) surpasses that of the best TD(λ) run after only about 10 seconds of computation.

In interpreting these timing results, two important points should be considered. First, LSTD(λ) only performed its coefficient update $\boldsymbol{\beta} := \mathbf{A}^{-1}\mathbf{b}$ once every 50 trajectories, when needed to collect performance statistics. Practical applications may require updating the $\boldsymbol{\beta}$ coefficients more or less frequently, with a corresponding increase or decrease in computational requirements. Second, the timings reported here measure only the time spent in the algorithmic calculations of TD(λ) or LSTD(λ); the time spent in trajectory generation is excluded. In many practical domains, including backgammon, the time spent in generating sample state transitions in the Markov chain may exceed the time required for temporal-difference learning. When data is expensive to generate, maximizing the data efficiency of

learning as LSTD(λ) does will reduce overall running times despite the greater computation spent in learning.

5. Conclusions and Future Work

We have argued, both from the experimental results above and from the deep connection to model-based reinforcement learning presented in Section 3, that the least-squares formulation of TD learning makes better use of simulation data than TD(λ). As a practical matter, the decision of when to prefer one algorithm over the other will depend on the application domain. If a domain has many features and simulation data is available cheaply, then incremental methods such as TD(λ) may have better real-time performance than least-squares methods [13]. On the other hand, some reinforcement learning applications have been successful with small numbers of features (e.g., [9, 3]), and in these situations LSTD(λ) should be superior.

LSTD(λ) has been applied successfully in the context of STAGE, a reinforcement learning algorithm for combinatorial optimization [4]. An exciting possibility for future work is to apply LSTD(λ) in the context of approximation algorithms for general Markov decision problems. LSTD(λ) provides an alternative to TD(λ) for the inner loop of optimistic policy iteration [2], and should enable good control policies to be discovered with fewer trial simulations.

Acknowledgments

Thanks to Andrew Moore, Jeff Schneider, Geoff Gordon, Jeremy Frank, Hamid Berenji, and the anonymous reviewers for many helpful comments. Thanks to Marc Ringuette for helping with the design of the backgammon evaluator. This work was supported by a NASA GSRP fellowship while the author was at Carnegie Mellon University.

Appendix

Equivalence of LSTD(1) and Linear Regression

We show here that when $\lambda = 1$, the LSTD(λ) algorithm of Figure 2 produces an approximate value function which is equivalent to that which would be produced by standard, non-incremental, least-squares linear regression. To be precise, assume we are given a sample trajectory $(x_0, x_1, \ldots, x_L, \text{END})$ of a Markov chain, with corresponding feature vectors $\boldsymbol{\phi}_t$ and one-step rewards R_t on each step. From this trajectory, a supervised-learning system learning from Monte-Carlo returns would

generate the following training pairs:

$$\begin{aligned}\boldsymbol{\phi}_0 &\mapsto R_0 + R_1 + \cdots + R_L \\ \boldsymbol{\phi}_1 &\mapsto R_1 + \cdots + R_L \\ &\vdots \quad \vdots \\ \boldsymbol{\phi}_L &\mapsto R_L\end{aligned}$$

Performing standard least-squares linear regression on the above training set would produce the coefficients $\boldsymbol{\beta} = \mathbf{A}^{-1}\mathbf{b}$, with $\mathbf{A}$ and $\mathbf{b}$ computed as follows:

$$\mathbf{A}_{\mathrm{LR}} = \sum_{i=0}^{L} \boldsymbol{\phi}_i \boldsymbol{\phi}_i^{\mathsf{T}} \qquad \mathbf{b}_{\mathrm{LR}} = \sum_{i=0}^{L} \boldsymbol{\phi}_i y_i$$

$$\text{where } \; y_i = \sum_{j=i}^{L} R_j$$

We now show that, thanks to the algebraic trick of the eligibility vectors $\mathbf{z}_t$, LSTD(1) builds the equivalent $\mathbf{A}$ and $\mathbf{b}$ fully incrementally—without having to store the trajectory while waiting to observe the eventual outcome y_i.

Proof: With simple algebraic manipulations, the sums built by LSTD(1)'s $\mathbf{A}$ and $\mathbf{b}$ telescope neatly into $\mathbf{A}_{\mathrm{LR}}$ and $\mathbf{b}_{\mathrm{LR}}$, as follows:

$$\begin{aligned}
\mathbf{A} &= \sum_{i=0}^{L} \mathbf{z}_t \big(\boldsymbol{\phi}_i - \boldsymbol{\phi}_{i+1}\big)^{\mathsf{T}} \\
&= \sum_{i=0}^{L} \big(\sum_{j=0}^{i} \boldsymbol{\phi}_j\big)(\boldsymbol{\phi}_i - \boldsymbol{\phi}_{i+1})^{\mathsf{T}} && \textit{(by definition of } \mathbf{z}_t\textit{)} \\
&= \sum_{i=0}^{L} \sum_{j=0}^{i} \boldsymbol{\phi}_j \boldsymbol{\phi}_i^{\mathsf{T}} - \sum_{i=0}^{L} \sum_{j=0}^{i} \boldsymbol{\phi}_j \boldsymbol{\phi}_{i+1}^{\mathsf{T}} \\
&= \big(\boldsymbol{\phi}_0 \boldsymbol{\phi}_0^{\mathsf{T}} + \sum_{i=1}^{L} \sum_{j=0}^{i} \boldsymbol{\phi}_j \boldsymbol{\phi}_i^{\mathsf{T}}\big) - \big(\sum_{k=1}^{L+1} \sum_{j=0}^{k-1} \boldsymbol{\phi}_j \boldsymbol{\phi}_k^{\mathsf{T}}\big) && \textit{(substituting } k \equiv i+1\textit{)} \\
&= \boldsymbol{\phi}_0 \boldsymbol{\phi}_0^{\mathsf{T}} + \sum_{i=1}^{L} \sum_{j=0}^{i} \boldsymbol{\phi}_j \boldsymbol{\phi}_i^{\mathsf{T}} - \sum_{k=1}^{L} \sum_{j=0}^{k-1} \boldsymbol{\phi}_j \boldsymbol{\phi}_k^{\mathsf{T}} && \textit{(since } \boldsymbol{\phi}_{L+1} \overset{\text{def}}{=} \mathbf{0}\textit{)} \\
&= \boldsymbol{\phi}_0 \boldsymbol{\phi}_0^{\mathsf{T}} + \sum_{i=1}^{L} \big(\sum_{j=0}^{i} \boldsymbol{\phi}_j \boldsymbol{\phi}_i^{\mathsf{T}} - \sum_{j=0}^{i-1} \boldsymbol{\phi}_j \boldsymbol{\phi}_i^{\mathsf{T}}\big) && \textit{(substituting } i \equiv k\textit{)} \\
&= \sum_{i=0}^{L} \boldsymbol{\phi}_i \boldsymbol{\phi}_i^{\mathsf{T}} \\
&= \mathbf{A}_{\mathrm{LR}}, \qquad \text{as desired;}
\end{aligned}$$

and

$$
\begin{aligned}
\mathbf{b} &= \sum_{i=0}^{L} \mathbf{z}_i R_i \\
&= \sum_{i=0}^{L} (\sum_{j=0}^{i} \boldsymbol{\phi}_j) R_i && \textit{(by definition of } \mathbf{z}_t\textit{)} \\
&= \sum_{i=0}^{L} \sum_{j=0}^{L} 1(j \le i) \boldsymbol{\phi}_j R_i && \textit{(where } 1(\mathit{True}) \stackrel{\text{def}}{=} 1,\ 1(\mathit{False}) \stackrel{\text{def}}{=} 0\textit{)} \\
&= \sum_{j=0}^{L} \sum_{i=0}^{L} 1(j \le i) \boldsymbol{\phi}_j R_i \\
&= \sum_{j=0}^{L} \boldsymbol{\phi}_j \sum_{i=j}^{L} R_i \\
&= \mathbf{b}_{\mathrm{LR}}, \qquad \text{as desired.}
\end{aligned}
$$

These reductions prove that the contributions to $\mathbf{A}$ and $\mathbf{b}$ by any single trajectory are identical in LSTD(1) and least-squares linear regression. In both algorithms, contributions from multiple trajectories are simply summed into the matrices. Thus, LSTD(1) and linear regression compute the same statistics and, ultimately, the same coefficients for the approximated value function. ■

Notes

1. The Fiona network, designed by Marc Ringuette and myself, consists of 200 input units, 20 sigmoidal hidden units, and 3 linear output units. Pubeval is a linear network designed by Gerry Tesauro, available from `http://www.cs.cmu.edu/afs/cs/project/connect/code/tesauro/`.
2. The 15 backgammon features used for training were measures of the race status, amount of contact, prime strength, inner board strength, checkers on bar, and immediate hit probability. Please contact the author for details.

References

1. C. G. Atkeson and J. C. Santamaria. A comparison of direct and model-based reinforcement learning. In *International Conference on Robotics and Automation*, 1997.
2. D. Bertsekas and J. Tsitsiklis. *Neuro-Dynamic Programming.* Athena Scientific, Belmont, MA, 1996.
3. J. A. Boyan and A. W. Moore. Learning evaluation functions for global optimization and Boolean satisfiability. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI)*, 1998.
4. J. A. Boyan. *Learning Evaluation Functions for Global Optimization.* PhD thesis, Carnegie Mellon University, 1998.
5. S. J. Bradtke and A. G. Barto. Linear least-squares algorithms for temporal difference learning. *Machine Learning*, 22(1/2/3):33–57, 1996.
6. L.-J. Lin. *Reinforcement Learning for Robots Using Neural Networks.* PhD thesis, Carnegie Mellon University, 1993.

7. A. W. Moore and C. G. Atkeson. Prioritized sweeping: Reinforcement learning with less data and less time. *Machine Learning*, 13:103–130, 1993.
8. W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, second edition, 1992.
9. S. Singh and D. Bertsekas. Reinforcement learning for dynamic channel allocation in cellular telephone systems. In M. C. Mozer, M. I. Jordan, and T. Petsche, editors, *NIPS-9*, page 974. The MIT Press, 1997.
10. R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 1998.
11. R. S. Sutton. Learning to predict by the methods of temporal differences. *Machine Learning*, 3, 1988.
12. R. S. Sutton. Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. In *Proceedings of the Seventh International Conference on Machine Learning*. Morgan Kaufmann, 1990.
13. R. S. Sutton. Gain adaptation beats least squares. In *Proceedings of the 7^{th} Yale Workshop on Adaptive and Learning Systems*, pages 161–166, 1992.
14. R. S. Sutton. TD models: Modeling the world at a mixture of time scales. In *Machine Learning: Proceedings of the 12th International Conference*, pages 531–539. Morgan Kaufmann, 1995.
15. G. Tesauro. TD-Gammon, a self-teaching backgammon program, achieves master-level play. *Neural Computation*, 6(2):215–219, 1994.
16. J. N. Tsitsiklis and B. Van Roy. An analysis of temporal-difference learning with function approximation. *IEEE Trans. Auto. Control*, 42(5):674–690, May 1997.