

# A Convergent Off-Policy Temporal Difference Algorithm

Raghuram Bharadwaj Diddigi\* Chandramouli Kamanchi\* Shalabh Bhatnagar

**Abstract**—Learning the value function of a given policy (target policy) from the data samples obtained from a different policy (behavior policy) is an important problem in Reinforcement Learning (RL). This problem is studied under the setting of off-policy prediction. Temporal Difference (TD) learning algorithms are a popular class of algorithms for solving the prediction problem. TD algorithms with linear function approximation are shown to be convergent when the samples are generated from the target policy (known as on-policy prediction). However, it has been well established in the literature that off-policy TD algorithms under linear function approximation diverge. In this work, we propose a convergent on-line off-policy TD algorithm under linear function approximation. The main idea is to penalize the updates of the algorithm in a way as to ensure convergence of the iterates. We provide a convergence analysis of our algorithm. Through numerical evaluations, we further demonstrate the effectiveness of our algorithm.

## I. INTRODUCTION

The two important problems in Reinforcement Learning (RL) are Prediction and Control [1]. The prediction problem deals with computing the value function of a given policy. In a discounted reward setting, value function refers to the total expected discounted reward obtained by following the given policy. The control problem refers to computing the optimal policy, i.e., the policy that maximizes the total expected discount reward. When the model information (probability transition matrix and single-stage reward function) is fully known, techniques like value iteration and policy iteration are used to solve the control problem. Policy iteration is a two-step iterative algorithm where the task of prediction is performed in the first step for a given policy followed by the policy improvement task in the second step. However, in most of the practical scenarios, the model information is not known and instead (state, action, reward and next-state) samples are only available. Under such a model-free setting, popular RL algorithms for prediction are Temporal Difference (TD) and for control are Q-Learning and Actor-Critic algorithms [2]. Actor-Critic algorithms can be seen as model-free analogs of the policy iteration algorithm and involve a model-free prediction step. Therefore, it is clear that model-free prediction

is an important problem for which optimal and convergent solutions are desired.

TD algorithms under the tabular approach (where there is no approximation of the value function) are a very popular class of algorithms for computing the exact value function of a given policy (henceforth referred to as target policy) from samples. In many of the real-life problems though, we encounter situations where the number of states is large or even infinite. In such cases, it is not possible to use tabular approaches and one has to resort to approximation based methods. TD algorithms are shown to be stable and convergent under linear function approximation, albeit under the setting of on-policy. On-policy refers to the setting where state and action samples are obtained using the target policy itself. As we approach practical scenarios, it can be noted that such samples are not always available to the practitioner. For example in games, say a practitioner would like to evaluate a (target) strategy. However, the data available to her might be from a player following a different strategy. The question that arises in this scenario is whether she can make use of this data and still evaluate the target strategy. These problems are studied under the setting of off-policy prediction where the goal is to evaluate the value function of the target policy from the data generated from a different policy (commonly referred to as behavior policy). The recent empirical success of the Deep Q-Learning algorithm (model-free control algorithm) motivates us to understand its convergence behavior, which is a very difficult problem. It has been noted in (Section 11.3 of [2]) that convergence and stability issues arise when we combine three components - function approximation, bootstrapping (TD algorithms) and off-policy learning, what they refer to as the “deadly triad”.

In our work, we propose an online off-policy stable TD algorithm for a prediction problem under linear function approximation. The idea is to penalize the parameters of the TD update to mitigate the divergence problem. We note here that the recent work [3] provides a comprehensive and excellent survey of algorithms for off-policy prediction problems and performs a comparative study. However, for the sake of completeness, we now discuss some of the important and relevant works on the off-policy prediction problem. In [4], Least-Squares TD algorithms (LSTD) with linear function approximation have been proposed that are shown to be convergent under both on-policy and off-policy settings. However, the per-step complexity of LSTD algorithms is quadratic in the number of parameters. In [5], off-policy TD algorithms are proposed that make use of an importance sampling idea to convert the expected value of total discounted reward under behavior

\* Equal Contribution.

This work was supported by the Robert Bosch Centre for Cyber-Physical Systems, Indian Institute of Science and a grant from the Department of Science and Technology, Government of India.

R. B. Diddigi and C. Kamanchi are with the Department of Computer Science and Automation, Indian Institute of Science, Bengaluru 560012, India (e-mail: raghub@iisc.ac.in; chandramouli@iisc.ac.in.).

S. Bhatnagar is with the Department of Computer Science and Automation and the Department of Robert Bosch Centre for Cyber-Physical Systems, Indian Institute of Science, Bengaluru 560012, India (e-mail: shalabh@iisc.ac.in).

policy to expected value under target policy. However, the variance of such algorithms is very high and in some cases tends to be infinite. In [6], the Gradient TD (GTD) algorithm has been proposed that is stable under off-policy learning and linear approximation and has linear (in the number of parameters) complexity. Since then, there have been a lot of improvements on the GTD algorithm under various settings like prediction, control, and non-linear function approximation [7]–[10]. The idea of adding the penalty in the form of a regularization term has been considered in [11] where Regularized off-policy TD (RO-TD) algorithm has been proposed based on GTD algorithms and convex-concave saddle point formulations. Emphatic TD algorithms (ETD) [12]–[15] are another popular class of off-policy TD algorithms that achieve stability by emphasizing or de-emphasizing updates of the algorithm. These updates also have linear-time complexity. Moreover, these algorithms learn only one set of parameters, unlike GTD algorithms which are two-time scale stochastic approximation algorithms that learn two sets of parameters. Recently in [16, 17], a co-variance off-policy TD (COP-TD) algorithm has been proposed that includes a co-variance shift term in the TD update. This shift term is also learned along with the parameter of the algorithm.

Our algorithm, like the Emphatic TD algorithm, trains only one set of parameters and like ETD and GTD algorithms, has per-update complexity that is linear in the number of parameters. The contributions of our paper are as follows:

- We derive an online off-policy TD learning algorithm with linear function approximation. Our algorithm has linear per-iteration computational complexity in the number of parameters.
- We prove the convergence of our algorithm utilizing the techniques of [12, 18].
- We show the empirical performance of our algorithm on standard benchmark RL environments.

The rest of the paper is organized as follows. In Section II, we introduce the background and preliminaries. We propose our algorithm in Section III. Sections IV and V describe the analysis of our algorithm. Section VI presents the results of our numerical experiments. Finally, Section VII presents concluding remarks and future research directions.

## II. BACKGROUND AND PRELIMINARIES

We consider a Markov Decision Process (MDP) of form  $(S, U, p, r, \gamma)$  where  $S$  denotes the state space.  $U$  is the set of actions,  $p$  is a probability transition matrix where  $p(s'|s, a)$  denotes the probability of system transition to state  $s'$  when action  $a$  is chosen in state  $s$ .  $r$  is the single-stage reward where  $r(s, a)$  denotes the reward obtained by taking action  $a$  in state  $s$ . Finally,  $\gamma$  denotes the discount factor. Let  $\pi : S \rightarrow \Delta(U)$  be the target policy where  $\Delta(U)$  denotes the set of probability distributions over actions. The objective of the MDP prediction problem is to estimate the value function ( $V^\pi$ ) of the target policy  $\pi$ , where the value function of a state  $s \in S$  denoted by  $V^\pi(s)$  is given by:

$$V^\pi(s) = \mathbb{E} \left[ \sum_{i=0}^{\infty} \gamma^i r(s_i, a_i) \middle| s_0 = s, \pi \right], \quad (1)$$

where the state-action trajectory  $(s_0, a_0, s_1, \dots)$  is obtained following the policy  $\pi$  and  $\mathbb{E}[\cdot]$  denotes the expectation.

As the number of states of the MDP can be very large, we resort to approximation techniques to compute the value function. In our work, we consider the linear function approximation architecture where

$$\widehat{V}(s) = \theta^T \phi(s), \quad (2)$$

where  $\widehat{V}(s)$  denotes the approximate value function associated with state  $s$  (that we desire to be very close to the exact value function),  $\phi(s)$  is a  $d \times 1$  feature vector associated with state  $s$  and  $\theta$  is a  $d \times 1$  weight vector. Note that the exact value function  $V^\pi$  may not be representable by (2). Therefore, the objective is to estimate the weight vector  $\theta$  so that the approximate value function denoted by (2) is as close as possible to the exact value function.

The on-policy TD(0) [2] is a popular on-line algorithm for computing the weight vector  $\theta$ . The update equation is given by:

$$\theta_{n+1} = \theta_n + \alpha_n (r_n + \gamma \theta_n^T \phi(s_{n+1}) - \theta_n^T \phi(s_n)) \phi(s_n), \quad (3)$$

where  $(s_n, r_n, s_{n+1})$  is the state, reward and next state samples obtained at time  $n$ ,  $\alpha_n$ ,  $n \geq 0$  is the step-size sequence and  $\theta_0$  denotes the initial parameter vector.

The stability of the on-policy TD(0) algorithm is well established in the literature [12]. We outline the proof of the convergence of this algorithm. Following the notation of [12], please note that the update rule (3) can be re-written as:

$$\theta_{n+1} = \theta_n + \alpha_n (b_n - A_n \theta_n), \quad (4)$$

where  $A_n = \phi(s_n)(\phi(s_n) - \gamma \phi(s_{n+1}))^T$  and  $b_n = r_{n+1} \phi(s_n)$ .

It is shown in [2] that the algorithm with update rule (4) is stable if the matrix  $A$  given by:

$$A = \lim_{n \rightarrow \infty} A_n = \Phi^T D_\pi (I - \gamma P_\pi) \Phi \quad (5)$$

is positive definite. In (5),  $\Phi$  is a  $|S| \times d$  matrix with the feature vector  $\phi(s)$  in row  $s$ .  $D_\pi$  is the  $|S| \times |S|$  diagonal matrix with the diagonal being the stationary distribution of state  $i$  obtained under policy  $\pi$ . Finally,  $P_\pi$  is a  $|S| \times |S|$  matrix with  $[P_\pi]_{ij} = \sum_a \pi(i, a) p(j|i, a)$ . For the on-policy TD(0) algorithm,  $A$  is shown to be positive definite [12] proving the stability of the algorithm.

In the off-policy prediction problem, the data samples are obtained from a behavior policy  $\mu$  instead of the target policy  $\pi$ . In this case, the off-policy TD(0) update [12] is given by:

$$\theta_{n+1} = \theta_n + \alpha_n \rho_n (r_n + \gamma \theta_n^T \phi(s_{n+1}) - \theta_n^T \phi(s_n)) \phi(s_n), \quad (6)$$

where  $r_n$  is the reward obtained by taking action  $a_n$  in state  $s_n$  and  $\rho_n$  is the importance sampling ratio given by  $\frac{\pi(s_n, a_n)}{\mu(s_n, a_n)}$ . The corresponding matrix  $A$  for this algorithm is given by:

$$A = \Phi^T D_\mu (I - \gamma P_\pi) \Phi, \quad (7)$$

where  $D_\mu$  is a diagonal matrix with diagonal of  $D_\mu$  being the stationary distribution obtained under policy  $\mu$ .

The matrix  $A$  defined in (7) need not be positive definite [12]. Therefore stability and convergence of the off-policy TD(0) are not guaranteed.

---

**Algorithm 1** Perturbed Off-Policy Prediction Algorithm

---

**Input:**

$\mu, \pi$ : behaviour and target policies respectively  
 $(s_n, a_n, r_n)_{n=0}^\infty$ : data from behaviour policy  
 $\theta_0$ : initial parameter vector  
 $\gamma$ : discount factor  
 $\phi(s)$ : feature vector of state  $s$   
 $\eta$ : non-negative regularization parameter  
 $\{\alpha_n\}$ : step-size sequence  
Iter: total number of iterations

**Output:**  $\theta_{\text{Iter}}$ 

```
1: procedure OFF-POLICY PREDICTION:  
2:   while  $n < \text{Iter}$  do  
3:      $\rho_n = \frac{\pi(s_n, a_n)}{\mu(s_n, a_n)}$   
4:      $\delta_n = r_n + \gamma \phi(s_{n+1})^T \theta_n - (1 + \eta) \phi(s_n)^T \theta_n$   
5:      $\theta_{n+1} = \theta_n + \alpha_n \rho_n \delta_n \phi(s_n)$   
6:   return  $\theta_{\text{Iter}}$ 
```

---

The off-policy TD(0) algorithm, if it converges, may perform comparably to some of the off-policy convergent algorithms in the literature. For example, in Figure 5 of [3], it has been shown that the performance of off-policy TD(0) is comparable to the GTD(0) algorithm. However, as the algorithm is not stable, off-policy TD(0) can diverge. In this paper, we propose a simple and stable off-policy TD algorithm. In the next section, we propose our algorithm and in Section IV, we provide its convergence analysis.

### III. THE PROPOSED ALGORITHM

The input to our algorithm is the target policy, whose value function we want to estimate and the behavior policy, using which the samples are generated. Also, provided as an input to our algorithm is the regularization parameter ( $\eta$ ). The algorithm works as follows. At each time step  $n$ , we obtain a sample  $(s_n, a_n, r_n, s_{n+1})$  using which importance sampling is computed as shown in the Step 3. We then compute our modified temporal difference term as show in Step 4. Finally, the parameters of the algorithm are updated as shown in Step 5.

**Remark 1.** *It is clear from the Algorithm 1 that, the per-step complexity is  $\mathcal{O}(d)$ , where  $d$  is the number of parameters.*

**Remark 2.** *The choice of  $\eta$  is critical in our algorithm. Larger values of  $\eta$  ensure convergence (see Lemma 2) and smaller values of  $\eta$  ensure more accurate solution (see Lemma 3)*

In the next section, we provide the convergence analysis of our proposed algorithm.

### IV. CONVERGENCE ANALYSIS

The update rule of Algorithm 1 can be rewritten as follows.

$$\begin{aligned} \theta_{n+1} &= \theta_n + \alpha_n \rho_n \delta_n \phi(s_n) \\ &= \theta_n + \alpha_n (b_n - A_n \theta_n), \end{aligned}$$

where  $A_n$  and  $b_n$  are given by

$$\begin{aligned} A_n &= -\rho_n (\gamma \phi(s_n) \phi(s_{n+1})^T - (1 + \eta) \phi(s_n) \phi(s_n)^T), \quad (8) \\ b_n &= \rho_n r_n \phi(s_n). \quad (9) \end{aligned}$$

We state and invoke Theorem 2 (also see Th. 17, p. 239 of [19]) of [18] to show the convergence of our algorithm.

**Theorem 1.** *Consider an iterative algorithm of the form*

$$\theta_{n+1} = \theta_n + a_n (b(X_n) - A(X_n) \theta_n)$$

where

- 1) *the step-size sequence satisfies  $\sum a_n = \infty$ ,  $\sum a_n^2 < \infty$ .*
- 2)  *$X_n$  is a Markov process with a unique invariant distribution.*
- 3)  *$A = \mathbb{E}_0[A(X_n)]$  and  $b = \mathbb{E}_0[b(X_n)]$ . Here  $\mathbb{E}_0$  is the expectation with respect to the stationary distribution of the Markov chain given by the behaviour policy  $\mu$ .*
- 4) *The matrix  $A$  is positive definite.*
- 5) *There exist positive constants  $C, q$  and a positive real valued function  $h$  from the states of the Markov chain  $\{X_n\}$  such that  $\sum_{n=0}^\infty \|\mathbb{E}[A(X_n)|X_0 = X] - A\| \leq C(1 + h^q(X))$  and  $\sum_{n=0}^\infty \|\mathbb{E}[b(X_n)|X_0 = X] - b\| \leq C(1 + h^q(X))$ .*
- 6) *For any  $q > 1$  there exists a constant  $\kappa_q$  such that for all  $X$  and  $n$ ,  $\mathbb{E}[h^q(X_n)|X_0 = X] \leq \kappa_q(1 + h^q(X))$ .*

*Under these assumptions, i.e. 1-6 above, it is known that  $\theta_n$  converges to the solution of  $b - A\theta = 0$ .*

To begin we define the process  $X_n$  as follows. Let  $X_n = (s_n, a_n, s_{n+1})$ . Observe that  $X_n$  is a Markov chain. In particular,  $s_{n+1}$  is a deterministic function of  $X_n$  and the distribution of  $a_{n+1}$  and  $s_{n+2}$  depends only on  $s_{n+1}$ . Also note that, in our algorithm,  $A(X_n) = A_n$  and  $b(X_n) = b_n$  given by equations (8) and (9) respectively with  $X_n = (s_n, a_n, s_{n+1})$ .

Assumptions 1 and 2 are fairly general. The assumptions 5 and 6 can be shown to hold with  $h$  as a constant function for finite state-action MDPs and the arguments are similar to those in theorem 1 of [18]. Therefore, the most important assumption to verify is that the matrix  $A$  is positive definite. In this section, we prove that  $A$  is positive definite, thereby proving the convergence of our proposed algorithm.

We begin by proving some important lemmas that are used in our main theorem.

**Lemma 1.** *Let  $\Phi$  be the  $|S| \times d$  matrix where the  $i^{\text{th}}$  row of  $\Phi$  is given by  $\phi(i)$ , the feature vector of state  $i$  and  $r_\pi$  be the  $|S| \times 1$  vector where the  $i^{\text{th}}$  component is given by  $r_\pi(i) = \sum_{a \in A} r(i, a) \pi(i, a)$ . Let  $\mathbb{E}_0$  be the expectation with respect to the stationary distribution of the Markov chain realized by  $\mu$ . Then  $A = \mathbb{E}_0[A_n]$  and  $b = \mathbb{E}_0[b_n]$  are given by*

$$\begin{aligned} A &= \Phi^T D_\mu ((1 + \eta)I - \gamma P_\pi) \Phi, \\ b &= \Phi^T D_\mu r_\pi, \end{aligned}$$

where  $D_\mu$  is a diagonal matrix with the  $i^{\text{th}}$  diagonal element being  $d_\mu(i)$ .

*Proof.*

$$\begin{aligned}
\mathbb{E}_0[A_n] &= -\mathbb{E}_0 \left[ \rho_n (\gamma \phi(s_n) \phi(s_{n+1})^T - (1 + \eta) \phi(s_n) \phi(s_n)^T) \right] \\
&= - \sum_{i,j \in S, a \in U} \mu(i, a) \left[ \frac{\pi(i, a)}{\mu(i, a)} (\gamma \phi(i) \phi(j)^T d_\mu(i) p(j|i, a) \right. \\
&\quad \left. - (1 + \eta) d_\mu(i) \phi(i) \phi(i)^T) \right] \\
&= - \sum_{i,j} d_\mu(i) (\gamma \phi(i) \phi(j)^T p_\pi(j|i) - (1 + \eta) \phi(i) \phi(i)^T) \\
&= \Phi^T D_\mu ((1 + \eta) I - \gamma P_\pi) \Phi.
\end{aligned}$$

Similarly

$$b = \mathbb{E}_0[b_n] = \mathbb{E}_0[\rho_n r_n \phi(s_n)] = \Phi^T D_\mu r_\pi.$$

**Definition 1.** A  $d \times d$  matrix  $M$  is positive definite if for all  $0 \neq y \in \mathbb{R}^d$ ,  $y^T M y > 0$ .

**Lemma 2.** Given a  $d \times d$  matrix  $M$ ,  $M$  is positive definite iff the symmetric matrix  $S = M + M^T$  is positive definite.

*Proof.* For  $0 \neq y \in \mathbb{R}^d$  observe that

$$y^T S y = y^T M y + y^T M^T y = 2y^T M y,$$

since  $(y^T M y)^T = y^T M y$  as both are scalars and  $y^T M^T y = (y^T M y)^T$ . Hence  $S$  is positive definite if and only if  $M$  is positive definite  $\square$

**Theorem 2.** Suppose  $M = D((1 + \eta)I - \gamma P)$  where  $D$  is a diagonal matrix with positive diagonal entries,  $P$  is a Markov matrix and  $\eta \geq \max_i \left( \frac{\gamma d^T p(\cdot|i)}{d_i} - 1, 0 \right)$  and  $0 < \gamma < 1$  are positive constants. Then  $M = [m_{ij}]$  is positive definite.

*Proof.* Consider the symmetric matrix  $S = M + M^T$ . From Lemma 2, it is enough to show that  $S$  is positive definite. Since  $S$  is symmetric it is diagonalizable. Therefore it is enough to show that the eigen-values of  $S$  are positive. From the Gershgorin circle theorem (see [20]) for any eigen value  $\lambda$  of  $S$  there exists  $i$  such that

$$\begin{aligned}
|\lambda - 2m_{ii}| &\leq \sum_{j \neq i} |m_{ij}| + \sum_{j \neq i} |m_{ji}| \\
\implies \lambda &\geq 2m_{ii} - \sum_{j \neq i} |m_{ij}| - \sum_{j \neq i} |m_{ji}|.
\end{aligned}$$

Now  $m_{ii} = d_i((1 + \eta) - \gamma p(i|i))$  and for  $i \neq j$  we have  $m_{ij} = -d_i \gamma p(j|i)$ . Therefore  $m_{ii} - \sum_{j \neq i} |m_{ij}| = (1 + \eta - \gamma) d_i$  and  $m_{ii} - \sum_{j \neq i} |m_{ji}| = ((1 + \eta) d_i - \gamma d^T p(i|\cdot))$ ,

$$\implies \lambda \geq (1 + \eta - \gamma) d_i + ((1 + \eta) d_i - \gamma d^T p(i|\cdot)) > 0$$

from the hypothesis  $\eta > \max_i \frac{\gamma d^T p(i|\cdot)}{d_i} - 1$ . We see that every eigen-value of  $S > 0$  i.e.,  $S$  is positive definite. Hence  $M$  is positive definite. In particular, given the behaviour policy  $\mu$  and the target policy  $\pi$ , there exists  $\eta > 0$  such that  $A = \Phi^T D_\mu ((1 + \eta)I - \gamma P_\pi) \Phi$  is positive definite.  $\square$

To describe the point of convergence of our algorithm consider for a given policy  $\mu$  and a parameter  $\eta$ ,  $T_\mu^\eta : \mathbb{R}^{|S|} \rightarrow \mathbb{R}^{|S|}$

as  $T_\mu^\eta = \frac{1}{1+\eta} T_\mu$ . We state and prove the following properties about  $T_\mu^\eta$ .

**Lemma 3.**  $T_\mu^\eta$  is a  $\|\cdot\|_\infty$ -contraction and converges point-wise to  $T_\mu$  as  $\eta \rightarrow 0$ .

*Proof.* From the definition  $T_\mu^\eta = \frac{1}{1+\eta} T_\mu$ . For any  $V \in \mathbb{R}^{|S|}$ ,

$$T_\mu^\eta V = \frac{1}{1+\eta} T_\mu V \rightarrow T_\mu V \text{ as } \eta \rightarrow 0.$$

It is easy to see that for any  $V, W \in \mathbb{R}^{|S|}$ ,

$$\begin{aligned}
\|T_\mu^\eta V - T_\mu^\eta W\|_\infty &= \frac{\gamma}{1+\eta} \|P_\mu(V - W)\|_\infty \\
&\leq \frac{\gamma}{1+\eta} \|V - W\|_\infty.
\end{aligned}$$

$\square$  Hence  $T_\mu^\eta$  is  $\|\cdot\|_\infty$ -contraction.  $\square$

## V. ABOUT THE POINT OF CONVERGENCE

The algorithm converges to the point  $\theta^*$  such that  $b - A\theta^* = 0$ . Now

$$\begin{aligned}
b - A\theta^* &= 0 \\
\implies \Phi^T D_\mu ((1 + \eta)I - \gamma P_\pi) \Phi \theta^* &= \Phi^T D_\mu r_\pi \\
\implies \Phi^T D_\mu \left( I - \frac{\gamma}{1+\eta} P_\pi \right) \Phi \theta^* &= \Phi^T D_\mu \frac{r_\pi}{1+\eta} \\
\implies \Phi^T D_\mu \Phi \theta^* &= \Phi^T D_\mu \left( \frac{r_\pi}{1+\eta} + \frac{\gamma}{1+\eta} P_\pi \Phi \theta^* \right) \\
\implies \Phi \theta^* &= \Phi (\Phi^T D_\mu \Phi)^{-1} \Phi^T D_\mu \left( \frac{r_\pi}{1+\eta} + \frac{\gamma}{1+\eta} P_\pi \Phi \theta^* \right) \\
\implies \Phi \theta^* &= \Pi_{D_\mu} T_\pi^\eta \Phi \theta^*,
\end{aligned}$$

where  $\Pi_{D_\mu} = \Phi (\Phi^T D_\mu \Phi)^{-1} \Phi^T D_\mu$  is the projection operator that projects any  $V \in \mathbb{R}^{|S|}$  to the subspace  $\{\Phi r | r \in \mathbb{R}^d\}$  with respect to the norm  $\|\cdot\|_{D_\mu}$ . Hence we observe that, similar to online on-policy TD, our online off-policy TD is a projected stochastic fixed point iteration with respect to the perturbed Bellman operator  $T_\pi^\eta$ .

**Remark 3.** Note that the bound derived for  $\eta$  in Theorem 2 is a sufficient but not a necessary condition. If the value of  $\eta$  is large, the algorithm converges but to a poorly approximated solution. Therefore, in experiments, we select the value of  $\eta$  that is large enough to ensure convergence and small enough to ensure that approximation is reasonable.

## VI. EXPERIMENTS AND RESULTS

In this section, we describe the performance of our proposed algorithm on three tasks. We first perform experiments on two benchmark counter-examples for off-policy divergence. Finally, we perform our experiment on a 3-state MDP example and analyze various properties of our proposed algorithm<sup>1</sup>. The evaluation metric considered is Root Mean Square Error (RMSE) defined as:

$$RMSE(\theta) = \sqrt{d_\mu(s)(V_\pi(s) - \widehat{V}_\theta(s))^2}, \quad (10)$$

<sup>1</sup>The implementation codes for our experiments is available at: <https://github.com/raghuddigi/Off-Policy-Convergent-Algorithm>

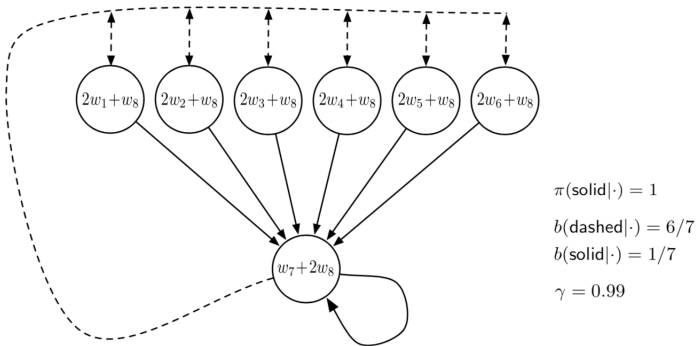


Figure 1: Baird’s Counterexample. Figure taken from [21]

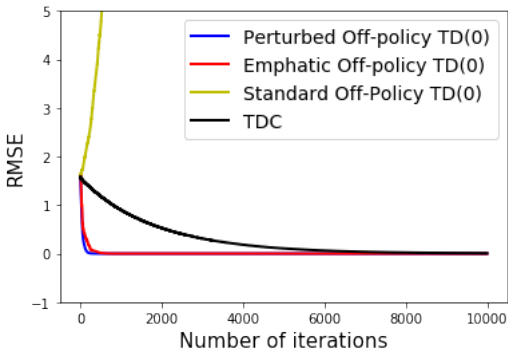


Figure 2: Performance of algorithms on “ $\theta \rightarrow 2\theta$ ”. RMSE is the value averaged across 10 independent runs

where  $\theta$  is the parameter that is used to approximate the value function,  $d_\mu$  is the stationary distribution associated with the behavior policy  $\mu$ ,  $V_\pi$  is the exact value function of the target policy  $\pi$  and  $\hat{V}_\theta$  is the approximate value function that is estimated. We perform 10 independent runs and present the average of RMSE obtained on all the three experiments. For comparison purposes, we also implement Emphatic TD (ETD(0)) algorithm [12] and a gradient-family algorithm, linear TD with gradient correction (TDC) [7].

First, we consider the “ $\theta \rightarrow 2\theta$ ” example ([22], Section 3 of [12]). In this example, there are two states - 1 and 2 and two actions - ‘left’ and ‘right’. Left action in state 1 results in state 1, while right action results in state 2. Similarly, right action in state 2 results in state 2 and left action results in state 1. The target policy is to take right in both the states, whereas behavior policy is to take left and right actions with equal probability in both the states. The value function is linearly approximated with one feature. The feature of state 1 is 1 and that of state 2 is 2. The discount factor is taken to be 0.9. The update parameter  $\theta$  is initialized to 1 and the  $\eta$  for our algorithm is taken to be 1. The step-size for the algorithms is held constant at 0.01. In Figure 2, we show the performance of algorithms over 10000 iterations. We can see that the standard off-policy TD(0) diverges whereas the other three algorithms including our proposed perturbed off-policy TD(0) converges to a point where the RMSE is zero.

Next, we consider the “7-star” example, first proposed in [23]. This is completely described in Figure 1 [21]. There are




Figure 3: Performance of algorithms on “Baird’s Counterexample”. RMSE is the value averaged across 10 runs

7 states represented as circles. The expression inside the circle  $i$  represents the linear approximation of the state  $i$ . The policy  $\pi$  in Figure 1 represents the target policy and  $b$  represents the behavior policy. We run all the algorithms, i.e., standard off-policy TD(0), Emphatic off-policy TD(0), TDC and our algorithm, Perturbed off-policy TD(0) for 1000000 iterations. The step-size for the algorithms is set to  $0.0001^2$ . From Figure 3, we can see that our perturbed off-policy TD converges to the exact solution while the Emphatic TD(0) appears to oscillate. On the other hand, the TDC algorithm appears to converge slowly. Moreover, it is known that standard off-policy diverges for this example, which can also be observed from Figure 3.

Finally, we construct an MDP as follows. There are 3 states and 2 actions - ‘left’ and ‘right’ possible in each state. The ‘Left’ action in states 1 and 2 leads to state 1. And the ‘right’ action in states 2 and 3 leads to state 3. Finally ‘left’ action in state 3 leads to state 2. The single-stage rewards in all transitions is taken to be 1 and the discount factor is 0.9. The target policy  $\pi = [[0, 1], [0.5, 0.5], [1, 0]]$  and the behavior policy  $\mu = [[0.9, 0.1], [0.5, 0.5], [0.1, 0.9]]$  (where the first component represents the probability to take ‘left’ and the second component represents the probability to take ‘right’). The feature vectors of the three states are  $[1, 0], [1, 1], [0, 1]$  respectively. The step-size for the algorithms is set to 0.0001. We run all the algorithms for 1000000 iterations. From Figure 4, we can see that perturbed off-policy TD(0) converges. For this experiment, the best possible RMSE is 2.548 and our proposed algorithm achieves 2.97.

In the experimental setting above, the value of  $\eta$  is set to 0.5. In Figure 5, we run our algorithm with two other values of  $\eta = 0.4$  and 0.6 respectively. We observe that, for  $\eta = 0.4$ , convergence is not guaranteed as this  $\eta$  correction is not enough. On the other hand, for  $\eta = 0.6$ , convergence is ensured. However, the converged solution is not close due to the over-correction. Hence, it is to be noted that an optimal value of  $\eta$  is desired for ensuring the convergence and near-optimal solution at the same time (recall that a higher value of  $\eta$  is enough to ensure the convergence alone).

<sup>2</sup>We have run experiments with three other step-sizes and included it in our supplementary material. Please find them at: <https://github.com/raghudiddigi/Off-Policy-Convergent-Algorithm/blob/master/Supplementary.pdf>

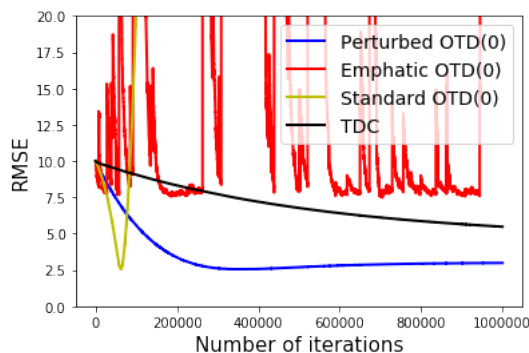


Figure 4: Performance of algorithms on a 3-state MDP. RMSE is the value averaged across 10 runs. The best possible RMSE for this MDP is 2.548




Figure 5: Performance of our proposed algorithm with three different  $\eta$  values

**Remark 4.** It has to be noted that the objective of the experiments is to show that our proposed algorithm mitigates the divergence problem of the off-policy TD algorithm. Moreover, if we choose a good value of  $\eta$ , it ensures that the algorithm converges to a solution closer to the optimal solution. At this point, we do not make any claims about the quality of the converged solution compared to the Emphatic TD(0) and TDC algorithms. We have seen that our proposed algorithm performed better than Emphatic TD and TDC in the last two examples. Further empirical analysis is needed to compare the quality of the converged solution with Emphatic TD(0), TDC as well as other off-policy algorithms in the literature.

## VII. CONCLUSIONS AND FUTURE WORK

In this work, we have proposed an off-policy TD algorithm for mitigating the divergence problem of the standard off-policy TD algorithm. Our proposed algorithm makes use of a penalty parameter to ensure the stability of the iterates. We have then proved that this addition of penalty parameter makes the matrix  $A$  positive definite, which in turn ensures the convergence. Finally, we empirically show the convergence on benchmark counter-examples for off-policy divergence.

As seen from the experiments, the choice of  $\eta$  is critical for our algorithm. The lower-bound that we have provided in our analysis is not tight and coming up with a tight bound is an interesting future direction. Also, in future, we would like to

extend our algorithm to include eligibility traces and study its applications on real world problems.

## REFERENCES

- [1] D. P. Bertsekas and J. N. Tsitsiklis, *Neuro-dynamic programming*. Athena Scientific Belmont, MA, 1996, vol. 5.
- [2] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [3] S. Ghiassian, A. Patterson, M. White, R. S. Sutton, and A. White, "Online off-policy prediction," *arXiv preprint arXiv:1811.02597*, 2018.
- [4] S. J. Bradtko and A. G. Barto, "Linear least-squares algorithms for temporal difference learning," *Machine learning*, vol. 22, no. 1-3, pp. 33–57, 1996.
- [5] D. Precup, R. S. Sutton, and S. Dasgupta, "Off-policy temporal-difference learning with function approximation," in *ICML*, 2001, pp. 417–424.
- [6] R. S. Sutton, C. Szepesvári, and H. R. Maei, "A convergent  $O(n)$  algorithm for off-policy temporal-difference learning with linear function approximation," *Advances in Neural Information Processing Systems*, vol. 21, no. 21, pp. 1609–1616, 2008.
- [7] 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 *Proceedings of the 26th Annual International Conference on Machine Learning*. ACM, 2009, pp. 993–1000.
- [8] H. R. Maei, C. Szepesvári, S. Bhatnagar, D. Precup, D. Silver, and R. S. Sutton, "Convergent temporal-difference learning with arbitrary smooth function approximation," in *Advances in Neural Information Processing Systems*, 2009, pp. 1204–1212.
- [9] H. R. Maei and R. S. Sutton, "GQ( $\lambda$ ): A general gradient algorithm for temporal-difference prediction learning with eligibility traces," in *3d Conference on Artificial General Intelligence (AGI-2010)*. Atlantis Press, 2010.
- [10] H. R. Maei, C. Szepesvári, S. Bhatnagar, and R. S. Sutton, "Toward off-policy learning control with function approximation." in *ICML*, 2010, pp. 719–726.
- [11] B. Liu, S. Mahadevan, and J. Liu, "Regularized off-policy TD-learning," in *Advances in Neural Information Processing Systems*, 2012, pp. 836–844.
- [12] R. S. Sutton, A. R. Mahmood, and M. White, "An emphatic approach to the problem of off-policy temporal-difference learning," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 2603–2631, 2016.
- [13] H. Yu, "On convergence of emphatic temporal-difference learning," in *Conference on Learning Theory*, 2015, pp. 1724–1751.
- [14] A. Hallak, A. Tamar, R. Munos, and S. Mannor, "Generalized emphatic temporal difference learning: Bias-variance analysis," in *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [15] S. Ghiassian, B. Rafiee, and R. S. Sutton, "A first empirical study of emphatic temporal difference learning," *arXiv preprint arXiv:1705.04185*, 2017.
- [16] A. Hallak and S. Mannor, "Consistent on-line off-policy evaluation," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 1372–1383.
- [17] C. Gelada and M. G. Bellemare, "Off-policy deep reinforcement learning by bootstrapping the covariate shift," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 3647–3655.
- [18] J. N. Tsitsiklis and B. Van Roy, "An analysis of temporal-difference learning with function approximation," *IEEE Transactions on Automatic Control*, vol. 42, no. 5, pp. 674–690, 1997.
- [19] A. Benveniste, M. Métivier, and P. Priouret, *Adaptive algorithms and stochastic approximations*. Springer Science & Business Media, 2012, vol. 22.
- [20] G. Golub and C. Van Loan, "Matrix computations, (Johns Hopkins University Press, Baltimore, 1996)."
- [21] J. Zhang, "Bairdexample," Nov. 2019. [Online]. Available: <https://github.com/MJeremy2017/Reinforcement-Learning-Implementation/tree/master/BairdExample>
- [22] J. N. Tsitsiklis and B. Van Roy, "Feature-based methods for large scale dynamic programming," *Machine Learning*, vol. 22, no. 1-3, pp. 59–94, 1996.
- [23] L. Baird, "Residual algorithms: Reinforcement learning with function approximation," in *Machine Learning Proceedings 1995*. Elsevier, 1995, pp. 30–37.