

# Discerning Temporal Difference Learning

Jianfei Ma

Northwestern Polytechnical University  
 School of Mathematics and Statistics  
 matrixfeeney@gmail.com

## Abstract

Temporal difference learning (TD) is a foundational concept in reinforcement learning (RL), aimed at efficiently assessing a policy's value function. TD( $\lambda$ ), a potent variant, incorporates a memory trace to distribute the prediction error into the historical context. However, this approach often neglects the significance of historical states and the relative importance of propagating the TD error, influenced by challenges such as visitation imbalance or outcome noise. To address this, we propose a novel TD algorithm named discerning TD learning (DTD), which allows flexible emphasis functions—predetermined or adapted during training—to allocate efforts effectively across states. We establish the convergence properties of our method within a specific class of emphasis functions and showcase its promising potential for adaptation to deep RL contexts. Empirical results underscore that employing a judicious emphasis function not only improves value estimation but also expedites learning across diverse scenarios.

## Introduction

In reinforcement learning, efficiently predicting future rewards based on past experiences is a fundamental challenge. TD(0) assigns credit using the difference between successive predictions (Sutton 1988), offering online and recursive capabilities, but its scope is limited to the current observed state. On the other hand, TD( $\lambda$ ) (Sutton and Barto 2018), when combined with the eligibility trace, assigns credit to all historical states, using a recency heuristic to propagate credit. However, it lacks consideration for the relative importance of each state, that is, how much emphasis should be given to each historical state to make better predictions? Both approaches have found success in modern RL algorithms (Mnih et al. 2013) (Schulman et al. 2017), but they uniformly weigh states, overlooking the potential benefits of emphasis-aware credit assignment.

In various circumstances, the emphasis-awareness becomes crucial. The accurate estimation of a value function often faces inherent challenges, including visitation imbalances and noisy outcomes, particularly in reward-seeking tasks. Due to variations in initial state distributions or transition models, agents tend to update high-frequency states

more frequently, while less attention is given to less frequent or near-terminating states, such as goal states. This phenomenon is especially prevalent in episodic tasks with eligibility traces, where more frequent states are updated each time a new state is encountered while states close to termination have shorter trace positions, resulting in fewer updates. In such cases, diverging update frequencies can lead to imbalanced value estimations. Furthermore, the added complexity arising from noisy observations (Chelu et al. 2022) or rewards (Wang, Liu, and Li 2020) exacerbates the estimation challenge. Injected noise can lead to erroneous predictions, propagating inaccuracies to other states. Moreover, in reality, the most rewarding states are often rare occurrences, particularly in situations where most states have sparse rewards. The fundamental idea is to emphasize less frequent or more valuable states by increasing individual update magnitudes or reducing attention on states with adverse factors, using a nonnegative emphasis function. While existing approaches (Sutton, Mahmood, and White 2016) (Anand and Precup 2021) (Chelu et al. 2022) provide some insights, they either lack a thorough investigation of emphasis functions in diverse scenarios or overlook the mutual influence of the emphasis between states.

In this paper, we introduce a novel class of TD learning methods based on a fundamental identity. This identity, when viewed forward, directly incorporates an emphasis function that prioritizes various multi-step return combinations, offering enhanced flexibility. We establish a connection to a computationally efficient backward view that updates online, with its structure revealing the emphasis function's role. We provide theoretical analysis demonstrating that for a specific class of emphasis functions, our method converges to the optimal solution in the sense of an emphasized objective. Illustrative examples are presented to investigate emphasis choices in diverse scenarios. These examples reveal that our proposed approach, DTD, can enhance value estimation and expedite learning, with particularly noticeable benefits from more compact choices like the absolute expected TD error. Moreover, the newly developed type of return function holds promise for adaptation to DRL scenarios, especially where accurate advantage estimation is crucial (Schulman et al. 2016). Additionally, we establish a connection to prioritized sampling (Schaul et al. 2016) in cases where the data is Markovian. Lastly, we initiate a dis-

cussion on the design of the emphasis function from various perspectives.

## Preliminaries

### Notation

Let's denote  $\|\cdot\|_{\Lambda}$  the vector norm induced by a positive definite matrix  $\Lambda$ , i.e.  $\|x\|_{\Lambda} = \sqrt{x^T \Lambda x}$ . And the corresponding induced matrix norm is  $\|\Lambda\|_{\Lambda} = \max_{\|x\|_{\Lambda}=1} \|\Lambda x\|_{\Lambda}$ . With  $\Lambda = \mathbf{I}$ , it comes to the Euclidean-induced norm, for which we drop the subscript as  $\|\Lambda\|$ . For simplicity,  $\mathbf{1}$  denotes the all-one vector. We indicate random variables by capital letters (e.g.  $S_t, A_t$ ), realization by lowercase letters (e.g.  $s_t, a_t$ ).

### Problem Setting

Consider an infinite-horizon discounted MDP, defined by a tuple  $(\mathcal{S}, \mathcal{A}, P, r, \rho_0, \gamma)$ , with a finite state space  $\mathcal{S}$ , a finite action space  $\mathcal{A}$ , a transition kernel  $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ , a reward function  $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ , an initial state distribution  $\rho_0 : \mathcal{S} \rightarrow \mathbb{R}$ , and a discount factor  $\gamma \in [0, 1]$ . Being at a state  $s_t \in \mathcal{S}$ , the agent takes an action  $a_t \in \mathcal{A}$  according to some policy  $\pi$ , which assigns a probability  $\pi(a_t|s_t)$  to the choice. After the environment receives  $a_t$ , it emits a reward  $r_t$ , and sends the agent to a new state  $s_{t+1} \sim P(s_{t+1}|s_t, a_t)$ . Repeating this procedure, the discounted return can be fulfilled as  $G_t = \sum_{t=0}^{\infty} \gamma^t R_t$ . We denote  $\mathbf{P}_{\pi} \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}|}$  the state transition matrix and  $\mathbf{r}_{\pi} \in \mathbb{R}^{|\mathcal{S}|}$  the expected immediate reward vector. And the steady-state distribution is denoted as  $\mathbf{d}_{\pi}(s)$ , which we assume exists and is positive at all states. Let  $\mathbf{D}$  denote the diagonal matrix with  $\mathbf{d}_{\pi}$  on its diagonal. The prediction problem we are interested in is to estimate the value function:

$$\hat{v}_{\pi}(s) = \mathbb{E}_{\pi}[G_t | S_t = s]. \quad (1)$$

When the state space is large or even continuous, it is beneficial to use function approximation  $\hat{v}(s, \theta)$  to represent  $v$  to generalize across states. In particular, if the feature is expressive, it is convenient to use linear function approximation:

$$\hat{v}(s, \theta) = \phi(s)^T \theta, \quad (2)$$

where  $\phi(s)$  is the feature vector at state  $s$ . With each feature vector of length  $K$  being at the row of the matrix  $\Phi$ , we can compactly represent the value function vector as  $\mathbf{V}_{\theta} = \Phi \theta$ . For any value function  $\mathbf{V}$ , the most representable solution in the span of  $\Phi$  corresponds to (Sutton et al. 2009) (Yu and Bertsekas 2009):

$$\Pi \mathbf{V} = \Phi \theta^* \text{ where } \theta^* = \arg \min_{\theta} \|\Phi \theta - \mathbf{V}\|_{\mathbf{D}}, \quad (3)$$

where  $\Pi$  is the projection matrix in the form of:

$$\Pi = \Phi (\Phi^T \mathbf{D} \Phi)^{-1} \Phi^T \mathbf{D}. \quad (4)$$

To solve Eq. (3), simulation-based approaches are often utilized. With  $\mathbf{V}$  equal to the one-step TD target, TD(0) performs stochastic gradient descent to minimize the TD error:

$$\theta_{t+1} = \theta_t + \alpha_t \delta_t \phi_t, \quad (5)$$

where

$$\delta_t = R_{t+1} + \gamma \theta_t^T \phi_{t+1} - \theta_t^T \phi_t \quad (6)$$

is the TD error, and  $\alpha_t$  is the learning rate. The advantage of this approach is that it incrementally updates the weight vector at every time step, without requiring waiting until the end of an episode. However, it only takes effect on the current observed state. TD( $\lambda$ ), on the other hand, while sustains the same benefit of the online update, it is able to influence past experiences. Those past experiences can be viewed as eligible experiences that receive credit from the latest experience. This results in a more efficient update:

$$\begin{aligned} \mathbf{e}_t &= \gamma \lambda \mathbf{e}_{t-1} + \phi_t \\ \theta_{t+1} &= \theta_t + \alpha_t \mathbf{e}_t \delta_t, \end{aligned} \quad (7)$$

where  $\mathbf{e}_t$  is called *eligibility trace*, with  $\mathbf{e}_{-1} = 0$ . It is this temporally extended memory that allows the TD error at the current time step to be propagated to the states along the path that leads to the current state.

While the additional parameter  $\lambda \in [0, 1]$  is seamlessly integrated into the trace, it originates from the conventional forward view that directly interpolates  $n$ -step return exponentially forming the  $\lambda$ -return:

$$G_t^{\lambda} = (1 - \lambda) \sum_{n=1}^{\infty} \lambda^{n-1} G_t^{(n)}, \quad (8)$$

where

$$G_t^{(n)} = \sum_{k=1}^n \gamma^{k-1} R_{t+k} + \gamma^n \hat{v}(S_{t+n}, \theta), \quad (9)$$

is the  $n$ -step return. The method based on the target  $G_t^{\lambda}$  is called the  $\lambda$ -return algorithm, which has been proven to achieve the same weight updates as offline TD( $\lambda$ ) (Sutton 1988) (Sutton and Barto 2018).

### Discerning Temporal Difference Learning

The limitation of TD( $\lambda$ ) is that it fails to account for the importance of each historical state or to consider the relative significance of propagating the TD error. To address this issue, we derive our new return function that directly incorporates emphasis start based on an important identity. Consider any function  $f : \mathcal{S} \rightarrow \mathbb{R}$ , the following identity holds:

$$\sum_{n=0}^{\infty} \lambda^n (f_{t+n} - f_{t+n+1} \lambda) = f_t, \quad (10)$$

which generalizes the multiplier  $1 - \lambda$  in the  $\lambda$ -return. An interesting property is that the above holds for any real-valued function. However, such a function class would be too large to accommodate our purpose. We, therefore, constrain it into the bounded positive real-valued function as the emphasis function, measuring the significance of each state, with which we can derive a new return function as follows:

**Proposition 1.** For any  $f : \mathcal{S} \rightarrow \mathbb{R}^+$ , it holds that:

$$\bar{G}_t^{\lambda, f} \doteq \sum_{n=1}^{\infty} \lambda^{n-1} (f_{t+n-1} - f_{t+n} \lambda) G_t^{(n)} \quad (11)$$

$$= f_t \hat{v}(S_{t+n}, \theta) + \sum_{n=0}^{\infty} (\gamma \lambda)^n \delta_{t+n} f_{t+n}. \quad (12)$$

We defer the precise proof to the Appendix.

Intuitively, as Eq. 12 indicates, each TD error term is reweighted by the emphasis function so as to control the relative strength of each future state.  $\bar{G}_t^{\lambda, f}$  here nonetheless is unnormalized unless with a scalar multiplier  $\frac{1}{f_t}$ . Henceforth, we will reload the notation as  $G_t^{\lambda, f} \doteq \frac{1}{f_t} \bar{G}_t^{\lambda, f}$ , named as discerning  $\lambda$ -return.

Next, we will formally deliver DTD with a composition of an emphasized objective and the discerning  $\lambda$ -return as mentioned above. Denote  $\mathbf{F}$  as a diagonal matrix with  $f$  on its diagonal, furthermore  $\mathbf{\Lambda} = \mathbf{F}\mathbf{D}\mathbf{F}$ , we therefore minimize an emphasized objective analogous to Eq. 3:

$$\boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta}} \|\boldsymbol{\Phi}\boldsymbol{\theta} - \mathbf{V}\|_{\mathbf{\Lambda}}, \quad (13)$$

which modulates the steady state probability with the square of the emphasis function. The projection matrix can be expressed as:

$$\boldsymbol{\Pi}^f = \boldsymbol{\Phi}(\boldsymbol{\Phi}^\top \mathbf{\Lambda} \boldsymbol{\Phi})^{-1} \boldsymbol{\Phi}^\top \mathbf{\Lambda}. \quad (14)$$

Any solution to Eq. 13 will have an orthogonal difference to the emphasized basis such that:

$$\mathbf{V} - \boldsymbol{\Pi}^f \mathbf{V} \perp \mathbf{\Lambda} \boldsymbol{\Phi}. \quad (15)$$

By combining the discerning  $\lambda$ -return as the target  $\mathbf{V}$ , and manipulating the equivalence to the backward view similar to the deduction of the TD( $\lambda$ ), we can derive the DTD( $\lambda$ ) update:

$$\begin{aligned} \mathbf{e}_t &= \gamma \lambda \mathbf{e}_{t-1} + f_t \phi_t \\ \boldsymbol{\theta}_{t+1} &= \boldsymbol{\theta}_t + \alpha_t \mathbf{e}_t \delta_t f_t \\ \mathbf{e}_{-1} &= 0, \end{aligned} \quad (16)$$

which distinguishes the historical state as well as regulates the relative importance of propagating the TD error. The complete algorithm is outlined in Alg. 1 with a general function approximator.

---

#### Algorithm 1: DTD( $\lambda$ )

---

**Input:**  $\pi, v_{\boldsymbol{\theta}}, f, \gamma, \lambda$   
**Initialize:**  $\boldsymbol{\theta}$  arbitrarily

1: **for** each episode **do**  
2:   Initialize  $S$   
3:   Initialize  $\mathbf{e}$   
4:   **repeat**  
5:     Take action  $A \sim \pi(\cdot | S)$ , observe  $R, S'$   
6:      $\mathbf{e} \leftarrow \gamma \lambda \mathbf{e} + f(S) \nabla \hat{v}(S, \boldsymbol{\theta})$   
7:      $\delta \leftarrow R + \gamma \hat{v}(S', \boldsymbol{\theta}) - \hat{v}(S, \boldsymbol{\theta})$   
8:      $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} + \alpha_t \delta \mathbf{e} f(S)$   
9:      $S \leftarrow S'$   
10:   **until**  $S$  is terminal  
11: **end for**  
12: **return**  $\boldsymbol{\theta}$

---

## Theoretical Analysis

We embark on a theoretical exploration of the algorithm's convergence behavior concerning the emphasis function, offering analyses for both parameter-independent and

parameter-dependent scenarios. To establish the foundation, we introduce several necessary assumptions:

**Assumption 1.** *The Markov chain  $\{S\}$  is irreducible and aperiodic.*

**Assumption 2.**  *$\boldsymbol{\Phi}$  has linearly independent columns.*

**Assumption 3.** *The learning rate  $\{\alpha_t\}$  is non-increasing, and satisfies Robbins-Monro conditions:*

$$\sum_{t=0}^{\infty} \alpha_t = \infty \quad \text{and} \quad \sum_{t=0}^{\infty} \alpha_t^2 < \infty. \quad (17)$$

**Assumption 4.**  *$f_{t+n-1} - f_{t+n} \lambda$  is independent of  $G_t^{(n)}$  for  $n \in \mathbb{N}^+$ .*

Assumptions 1–3 adhere to the standard framework for analyzing linear TD methods (see, for instance, (Tsitsiklis and Roy 1997) (Yu 2010)). Assumption 4 is introduced to facilitate analytical operator analysis.

To characterize the discerning  $\lambda$ -return in its expected behavior, we introduce a notion of the DTD( $\lambda$ ) operator, which encapsulates the essence of the forward-view DTD( $\lambda$ ):

**Definition 1.** *Discerning  $\lambda$ -return operator:*

$$\begin{aligned} \mathcal{T}^{\lambda, f}(\mathbf{V}_{\boldsymbol{\theta}}) &= \mathbf{F}^{-1} \sum_{n=0}^{\infty} \lambda^n (\mathbf{P}_{\pi}^n (\mathbf{I} - \lambda \mathbf{P}_{\pi}) \mathbf{F}) \mathbf{1} \circ \\ &\quad \left( \sum_{t=0}^n (\gamma \mathbf{P}_{\pi})^t \mathbf{r}_{\pi} + (\gamma \mathbf{P}_{\pi})^{n+1} \mathbf{V}_{\boldsymbol{\theta}} \right), \end{aligned} \quad (18)$$

where  $\circ$  is the Hadamard product between matrices.

Next, we examine the contraction condition about  $\mathcal{T}^{\lambda, f}$ .

**Theorem 1.** *Let  $\sigma_{\min}(\mathbf{F})$  represent the smallest singular value of matrix  $\mathbf{F}$ . The mapping  $\mathcal{T}^{\lambda, f}$  is a contraction for the parameter-independent case if it satisfies:*

- i)  $\|\mathbf{F}\|_{\mathbf{\Lambda}} < \frac{\sigma_{\min}(\mathbf{F})(1-\gamma\lambda)}{\gamma \|\mathbf{1}\|_{\mathbf{\Lambda}} \|\mathbf{I} - \lambda \mathbf{P}_{\pi}\|_{\mathbf{\Lambda}}}$ .
- ii) *For the parameter-dependent case, if further there exists a Lipschitz constant  $\kappa \in [0, \frac{(1-\lambda)(1-\gamma)\sigma_{\min}(\mathbf{F})}{r_{\max} \|\mathbf{1}\|_{\mathbf{\Lambda}} \|\mathbf{I} - \lambda \mathbf{P}_{\pi}\|_{\mathbf{\Lambda}}})$  such that for any  $\mathbf{F}_1(\mathbf{V}_{\boldsymbol{\theta}_1}), \mathbf{F}_2(\mathbf{V}_{\boldsymbol{\theta}_2})$ :*

$$\|\mathbf{F}_1 - \mathbf{F}_2\|_{\mathbf{\Lambda}} \leq \frac{\kappa}{2} \|\mathbf{V}_{\boldsymbol{\theta}_1} - \mathbf{V}_{\boldsymbol{\theta}_2}\|_{\mathbf{\Lambda}},$$

*then it is a contraction mapping provided that:*

$$\|\mathbf{F}\|_{\mathbf{\Lambda}} < \frac{\sigma_{\min}(\mathbf{F})(1-\gamma\lambda)}{\gamma \|\mathbf{1}\|_{\mathbf{\Lambda}} \|\mathbf{I} - \lambda \mathbf{P}_{\pi}\|_{\mathbf{\Lambda}}} - \frac{(1-\gamma\lambda)r_{\max}\kappa}{\gamma(1-\lambda)(1-\gamma)}, \forall \boldsymbol{\theta} \in \Theta,$$

where  $\Theta$  is the parameter space that can be a suitable subset of  $\mathbb{R}^K$ .

This property guarantees the uniqueness of the fixed point. To avoid repetition, we define the function class  $\Xi$  as a set that satisfies either condition i) or ii).

**Remark 1.** *Note  $\|\mathbf{1}\|_{\mathbf{\Lambda}}$  is simply the expected value of the squared emphasis function under the steady-state distribution, i.e.  $\mathbb{E}_{\mathbf{d}_{\pi}}[f^2(S)]$ . In practice, if we can scale the emphasis function into a considerably small range (i.e.  $[0, 1]$ ), we will have a broader spectrum that enhances the contraction.*

**Corollary 1.**  *$\boldsymbol{\Pi}^f \mathcal{T}^{\lambda, f}$  is a contraction mapping for any  $f \in \Xi$ .*

*Proof.* From Eq. 15 we know that the difference between  $\mathbf{V}$  and  $\mathbf{\Pi}^f \mathbf{V}$  is orthogonal to the  $\Phi$  in the sense of the  $\|\cdot\|_\Lambda$ , whereas  $\mathbf{\Pi}^f \mathbf{V}$  is a linear combination of  $\Phi$ , therefore  $\mathbf{V} - \mathbf{\Pi}^f \mathbf{V} \perp \Lambda \mathbf{\Pi}^f \mathbf{V}$ . By Pythagorean theorem, it follows that  $\mathbf{\Pi}^f$  is non-expansive. Since  $\mathcal{T}^{\lambda, f}$  is a contraction mapping, thus the composition is also a contraction mapping.  $\square$

Consider a process  $X_t = \{S_t, S_{t+1}, \mathbf{e}_t\}$ , which is a finite Markov process as  $\mathbf{e}_t$  is only dependent up to  $S_t$ . Thereby, the update in Eq. 16 can be simplified as follows:

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t + \alpha_t (A(X_t)\boldsymbol{\theta}_t + b(X_t)), \quad (19)$$

where  $A(X_t) = \mathbf{e}_t(\gamma\phi(S_{t+1}) - \phi(S_t))^\top f(S_t)$  and  $b(X_t) = \mathbf{e}_t R_t f(S_t)$ . It was shown that this update exhibits asymptotic behavior akin to a deterministic variant in the sense of the steady state distribution (Benveniste, Métivier, and Priouret 1990) (Tsitsiklis and Roy 1997). As a result, we now delve into the essential quantities required to represent such a variant. Denoting  $\mathbf{A} = \mathbb{E}_{\mathbf{d}_\pi}[A(X_t)]$  and  $\mathbf{b} = \mathbb{E}_{\mathbf{d}_\pi}[b(X_t)]$ , they can be succinctly expressed in the matrix form:

**Lemma 1.** Denote  $\bar{\mathbf{F}} = \lim_{t \rightarrow \infty} \mathbf{F}_t$ , which is assumed to exist, then

$$\begin{aligned} \mathbf{A} &= \Phi^\top \bar{\mathbf{F}} \mathbf{D} (\mathbf{I} - \gamma\lambda \mathbf{P}_\pi)^{-1} \bar{\mathbf{F}} (\gamma \mathbf{P}_\pi - \mathbf{I}) \Phi \\ \mathbf{b} &= \Phi^\top \bar{\mathbf{F}} \mathbf{D} (\mathbf{I} - \gamma\lambda \mathbf{P}_\pi)^{-1} \bar{\mathbf{F}} \mathbf{r}_\pi. \end{aligned} \quad (20)$$

In the parameter-independent case,  $\mathbf{F}_t \equiv \mathbf{F}$ , while in the parameter-dependent case,  $\mathbf{F}_t = \mathbf{F}(\mathbf{V}_{\boldsymbol{\theta}_t})$ .

Using those quantities, we can express the deterministic variant as follows:

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t + \alpha_t (\mathbf{A}\boldsymbol{\theta}_t + \mathbf{b}). \quad (21)$$

To establish a connection with the earlier contraction results, it can be shown that:

$$\mathbf{A}\boldsymbol{\theta} + \mathbf{b} = \Phi^\top \Lambda (\mathcal{T}^{\lambda, f}(\Phi\boldsymbol{\theta}) - \Phi\boldsymbol{\theta}). \quad (22)$$

Building upon this result and the established contraction condition, we can derive a fundamental component for the convergence:

**Lemma 2.**  $\mathbf{A}$  is negative definite for any  $f \in \Xi$ .

With the above results, we can now demonstrate the convergence result:

**Theorem 2.** The updates induced by  $DTD(\lambda)$  converge to a unique fixed point  $\boldsymbol{\theta}^*$  satisfying  $\mathbf{A}\boldsymbol{\theta}^* + \mathbf{b} = 0$  for any  $f \in \Xi$ .

## Experiments

In this section, we delve into the impact of  $DTD(\lambda)$ 's emphasizing effect, whether the emphasis function is predetermined or adapted during training. We examine scenarios involving visitation imbalance or noisy outcomes to determine if  $DTD(\lambda)$  can address these challenges and enhance overall performance using predetermined emphasis. Regarding adaptive emphasis, we explore a more compact form, namely the absolute expected TD error, to assess the influence of non-stationary emphasis on the prediction tasks. Our findings demonstrate that, firstly, the update-rebalancing and noise-averse effects effectively handle inherent prediction difficulties; secondly, the promising adaptive emphasis surpasses numerous baselines across diverse tasks.

## Evaluation

We choose the mean-square projected Bellman error (MSPBE) (Sutton et al. 2009) as the performance metric, as it quantifies the deviation from the most representative functions attainable with the given features of  $\mathcal{T}\mathbf{V}_\theta$ , where  $\mathcal{T}$  is the Bellman operator. The MSPBE is defined as follows:

$$MSPBE(\boldsymbol{\theta}) = \|\mathbf{V}_\theta - \mathbf{\Pi} \mathcal{T} \mathbf{V}_\theta\|_\mathbf{D}. \quad (23)$$

The experiments are carried out over 50 independent runs, each spanning 5000 environment steps. The depicted curves report the best performance with extensive parameter sweeping. They present the aggregated mean along with error bars representing the standard deviation. All of the problems are episodic, undiscounted, and involve only a fixed target policy.

## More or Less

To illustrate the impact of the emphasis in scenarios with the visitation imbalance, we examine a 5-state random-walk problem featuring three distinct initial state distributions. In each episode, the agent starts deterministically from either the leftmost, middle, or rightmost state. This selection of the initial state leads to varying visitation frequencies among the neighboring states. States near the initial choice are more frequently visited, while those farther away are less likely to be encountered. Consequently, these three initial state distributions result in overall state visitation frequencies that exhibit left-skewed, center-elevated, or right-skewed patterns. The chain includes two terminal states located at opposite ends, with all transitions uniformly distributed across states. Rewards are uniformly zero, except when transitioning into the right terminal state. Tabular features represent the state characteristics. The challenge with  $TD(\lambda)$  arises from the tendency to update more frequently visited states heavily, while paying less attention to states that are visited infrequently. This discrepancy emerges due to the higher occurrence of more frequent states in the eligibility trace, resulting in more updates from ahead-time steps. This effect can be even more amplified if these states persist in the trace for an extended duration. Conversely, states that are infrequent visitors or have shorter durations in the trace undergo fewer updates. To rectify these imbalanced updates, our approach,  $DTD(\lambda)$ , addresses the shortage of total updates by increasing the update magnitude for infrequent states. The emphasis function we tailor is based on the inverse of the normalized empirical state visitation counts, which are then scaled to lie within the range of  $[0, 1]$  as recommended by Remark 1. Finally, we take the square root to restore the original quantity.

To demonstrate the efficacy of our method for combatting the noisy outcome in scenarios with the perturbed reward, we consider a larger problem with 10 states where transitions have a uniform reward level with added noises. In order to isolate the influence of visitation imbalance, the initial state is selected uniformly from all available states, and all transitions are executed uniformly. The noise is symmetric for the transition from a state but varies across states. We consider the noises as  $\sigma = [0, 1, 0.2, \dots, 1]$  for states  $s_1, s_2, \dots, s_{10}$ . Three different reward levels  $r = [-1, 0, 1]$

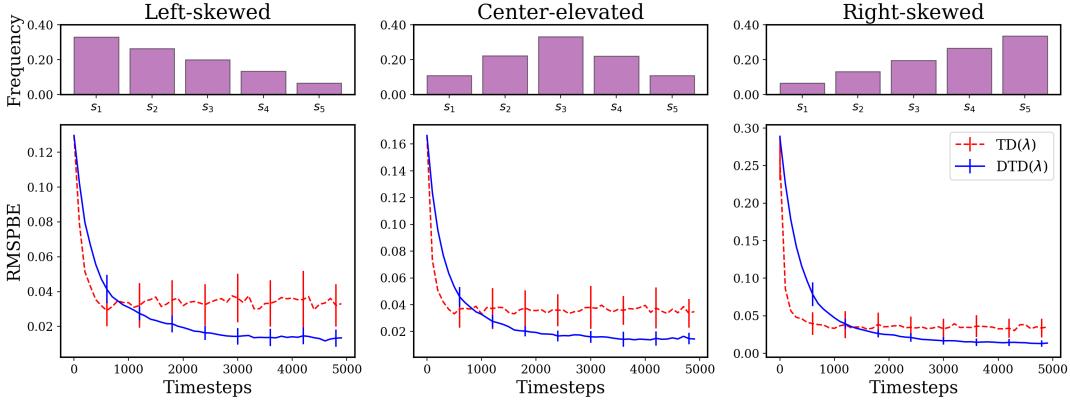


Figure 1: Top: State visitation frequency for different states; Bottom: Learning curve of MSPBE for algorithmic comparison. The three tasks are based on three different initial distributions.

are tested with varying difficulties. The reward that the agent actually receives is the reward level with an added Gaussian noise  $\mathcal{N}(0, \sigma(s_i))$  for transition from  $s_i$ . Incorrect predictions can occur when the agent is unaware of underlying noises, leading it astray from the true value. Moreover, this can lead to even more severe consequences, as the erroneous TD error may propagate to other states. DTD( $\lambda$ ) offers greater flexibility in addressing this situation through the emphasis function, which places resistance on states with high noise levels while prioritizing those with low noise levels. To that end, we introduce a prior into the design of the emphasis function, specifically the negative exponential of the noise levels, denoted as  $\exp(-\sigma(s))$ , to mitigate the influence of unpredictable outcomes. It is also scaled to lie within the range of  $[0, 1]$  and applied the square root.

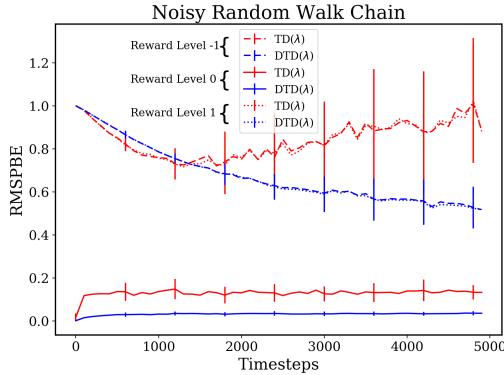


Figure 2: Learning curve of MSPBE of different reward levels with added noises.

The results depicted in Fig. 1 indicate that regardless of the skewness of the state visitation frequency, DTD( $\lambda$ ) effectively rebalances updates to achieve improved overall predictions. While TD( $\lambda$ ) shows faster progress in the early stages, the aliasing effect of TD error and the lack of updates for infrequent states become more pronounced, leading to its struggle in further reducing the error. In contrast, DTD( $\lambda$ )

allocates more attention to those infrequent states, resulting in a more balanced update process. In the case illustrated in Fig. 2, even at a zero reward level, TD( $\lambda$ ) results in a larger prediction error with higher variation, while DTD( $\lambda$ ) consistently maintains a relatively small prediction error with less variability. As the reward level becomes non-zero, the increased complexity involved in predicting the true value causes TD( $\lambda$ ) to progressively deviate from its initial prediction. In contrast, DTD( $\lambda$ ) effectively discerns different noise levels, leading to a reduction in prediction errors. From the learning curve, We hypothesize that with an increased computational budget, DTD( $\lambda$ ) can yield a much lower prediction error.

The idea of allocating attention selectively can be enlightening. In reality, valuable states are often infrequently encountered, and achieving a goal can require substantial effort. By focusing more attention on these crucial outcomes, we can enhance the influence of pathways leading to them.

### Adaptive Emphasis

the predetermined emphasis can vary depending on the specific problems, making manual crafting challenging. Is it possible to devise a compact emphasis that directly aligns with the nature of the prediction task? In this part, we examine the parameter-dependent emphasis, namely the absolute expected TD error, evaluated using the true dynamics, to showcase the effectiveness of the prediction-oriented emphasis for accelerating learning. We investigate four additional tasks, three of which share the same setup as the 5-state problem discussed earlier, but the initial state distribution is set to the middle state by default. There involves three representations as introduced in (Sutton et al. 2009): tabular, inverted (inappropriate state generalization), and dependent (insufficient representation), posing aliasing and representation challenges that standard methods are difficult to solve. Due to space limits, we refer the reader to (Sutton et al. 2009) for more detailed descriptions. The last task is a 13-state Boyan chain with 4 features (Boyan 2002), which serves as a standard benchmark for evaluating TD-style algorithms. In addition to comparing DTD( $\lambda$ ) with

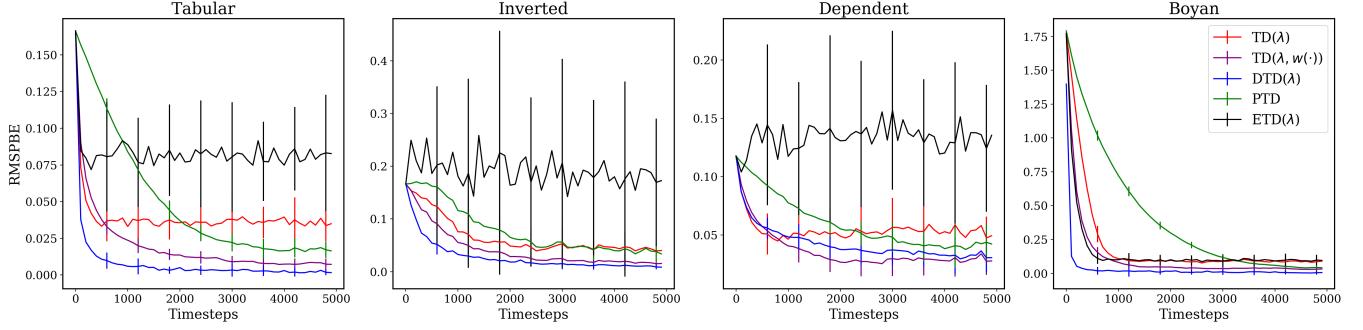


Figure 3: Learning curve of MSPBE on 5-state random walk chain with tabular, inverted, and dependent feature representation and the 13-state Boyan chain. Baselines are chosen to be emphatic and with selective updating.

TD( $\lambda$ ), we assess its performance against several baselines that incorporate varying levels of emphasis. These baselines include an on-policy emphatic variant of ETD( $\lambda$ ) (Sutton, Mahmood, and White 2016), the preferential approach PTD (Anand and Precup 2021), and the selective updating TD( $\lambda, w(\cdot)$ ) (Chelu et al. 2022).

The results presented in Fig. 3 demonstrate that DTD( $\lambda$ ) outperforms the other methods across the majority of tasks, exhibiting both rapid initial learning and minimal variability. Even under challenging representations, it can not only mitigate incorrect state aliasing where the update from one state only changes the parameters of other states, but also manifest the benefit of adaptive updating for limited capacity where the span of the feature space is insufficient to solve the task exactly. It is worth noting that ETD( $\lambda$ ) may suffer from the high variance issue of the follow-on trace, which could explain its poor performance. On the other hand, PTD appears to interpolate between TD(0) update and no update, causing its updates to be centered around TD(0) and possibly missing out on the advantages of combining different  $n$ -step returns. The approach most closely related to ours is TD( $\lambda, w(\cdot)$ ), which employs a similar eligibility trace. However, it does not take into account the importance of propagating the TD error. Comparing our approach to TD( $\lambda, w(\cdot)$ ) is essentially a direct test of the significance of our emphasis factor multiplied by the TD error. The results clearly show that removing this emphasis factor significantly degrades the performance, underscoring its crucial role in amplifying the propagation of TD error and its relative influence to historical states when combined with the eligibility trace.

## Extensibility for DRL

In this section, we delve deeper into the aspects of DTD, particularly its connection to advantage estimation and its relevance to prioritized sampling. The former is closely linked to the concept of discerning  $\lambda$ -returns, which holds the potential for further enhancing variance reduction. The latter aspect establishes a relationship with non-uniform sampling, wherein DTD(0) can yield a similar prioritization effect.

## Discerning Advantage Estimator

In the realm of DRL algorithms, the variance of policy gradients often becomes a bottleneck for overall performance,

particularly in on-policy algorithms (Schulman et al. 2015) (Schulman et al. 2017). Generalized Advantage Estimator (GAE) (Schulman et al. 2016) offers an effective approach to mitigate the high variance stemming from lengthy trajectory estimates. Notably, (Peng et al. 2018) and (Ma 2023) establish a close connection between GAE and the  $\lambda$ -return, albeit with a baseline function integrated to reduce the variance. Similarly, we can derive the Discerning Advantage Estimator (DAE), in the context of the discerning  $\lambda$ -return:

**Definition 2.** *Discerning Advantage Estimator:*

$$\hat{A}_t^{DAE(\lambda, \gamma, f)} = \frac{1}{f_t} \sum_{n=0}^{\infty} (\gamma \lambda)^n \delta_{t+n} f_{t+n}. \quad (24)$$

Aside from reducing variance with the baseline, DAE incorporates emphasis to reweight the TD error terms. We hypothesize that a well-chosen emphasis function can additionally lower the variance of the advantage estimate, specifically by quantifying the variance of the  $n$ -step return.

## Connection to Prioritized Sampling

Prioritized experience replay (PER) (Schaul et al. 2016) highlights that the “uniformness” of experience replay (Mnih et al. 2013) adheres to the same frequency as the original experiences, yet it fails to account for the significance of individual samples. This method updates the value function by assigning a priority to each sample, thus giving precedence to samples with higher significance, specifically those proportional to the sampled absolute TD error. This approach, referred to as the “frequency-centric” approach, emphasizes rolling out these significant samples more frequently, leading to more updates. On the other hand, DTD(0) follows a “magnitude-centric” approach, increasing the magnitude of each update to directly address the imbalanced frequency. The following proposition demonstrates the connection between DTD(0) and PER:

**Proposition 2.** *For a Markovian dataset  $\mathcal{D}$  generated from  $\pi$ , of a size  $N$ , then:*

$$\begin{aligned} & \mathbb{E}_{\text{uniform}} \left[ f^2(s) (v^{\text{target}} - \hat{v}(s, \theta))^2 \right] \\ &= c \cdot \mathbb{E}_q \left[ (v^{\text{target}} - \hat{v}(s, \theta))^2 \right], \end{aligned} \quad (25)$$

where

$$q(s) = \frac{f^2(s)}{\sum_{s' \in \mathcal{D}} f^2(s')} \quad c = \frac{\sum_{s' \in \mathcal{D}} f^2(s')}{N}. \quad (26)$$

What this conveys is that sampling from a priority distribution  $q(s)$ , scaled by a constant factor  $c$ , is analogous to uniform sampling with an emphasized objective. It is worth noting that any priority distribution of interest can be obtained by taking the square root of the corresponding emphasis function. This equivalence between DTD(0) and PER is compelling, as it directly integrates the emphasis into the objective to achieve the same prioritization effect.

However, we refrain from specifying a fixed form for the emphasis function, as it should ideally be tailored to the specifics of each problem, considering factors like problem complexity and size. In practical applications, employing function approximations to extend the emphasis across similar states could prove useful. However, delving into the details of the estimation method for such function approximations lies beyond the scope of this study and presents an intriguing avenue for future research.

## Discussions on Possible Forms

In this section, we open a dialogue on designing emphasis functions that would maximize the effectiveness of our approach.

The emphasis can be future-predicting, such as selecting the conditional entropy  $f_{t+n} = H(G|A_{\leq t+n-1}, S_{\leq t+n})$ . This choice will prioritize the  $n$ -step return with maximal information contained in  $(A_{t+n-1}, S_{t+n})$  about the return  $G$ . To provide a more intuitive understanding, for  $\lambda = 1$ , the quantity  $f_{t+n-1} - f_{t+n}$  equates to the mutual information  $I(G|A_{\leq t+n-2}, S_{\leq t+n-1}; A_{t+n-1}, S_{t+n})$ .

The emphasis can also be history-summarizing, such as choosing the negative exponential of variance of  $n$ -step return  $f_{t+n} = \exp(-\mathbb{V}[G_t^{(n)}])$ . This approach would resemble the experiments involving perturbed rewards, allowing the distinction of various return functions based on their noise levels. Such an approach could be beneficial for model-based planning, as the accumulation of errors in the model can render predictions less reliable (Janner et al. 2019).

It is also intriguing to assess the expected immediate reward for each state, with which it becomes possible to categorize the state space based on higher rewards. This enables the allocation of more resources towards predicting the value function of these valuable states, enhancing their utility in control tasks.

## Related Work

In the context of TD learning, ETD (Sutton, Mahmood, and White 2016) employs a follow-on trace coupled with an interest function to address the stability issue in the off-policy TD learning. PTD (Anand and Precup 2021) introduces a preference function that is reversely related to  $\lambda$ , enabling interpolation between TD(0) update and no update to handle partial observation challenges. Additionally,  $\text{TD}(\lambda, w(\cdot))$

(Chelu et al. 2022) proposes a selective eligibility trace for reweighting historical states, similar to our approach. However, it disregards the consideration of the relative influence of propagating the TD error.

Emphasizing the significance of certain states is a recurring concept in various domains. (McLeod et al. 2021) addresses the multi-prediction problem with a GVF (Sutton et al. 2011) by focusing on learning a subset of states for each prediction, facilitated by an underlying interest function. (Imani, Graves, and White 2018) introduces an extension of emphatic weighting into the domain of control, resulting in an off-policy emphatic policy gradient that incorporates a state-dependent interest function. However, the process of adapting or selecting an appropriate interest function can be challenging. In response, (Klissarov et al. 2022) proposes a meta-gradient to dynamically adjust the interest, highlighting the advantages of identifying crucial states, thereby enhancing the efficacy of transfer learning across RL tasks. It is possible to combine our method with these techniques. The intriguing question, however, is how our approach can be most suitable for control problems. The notion of selective updating also finds application in option learning, manifesting either through initiation sets (Sutton, Precup, and Singh 1999), or via the utilization of an interest function (Khetarpal et al. 2020). In model-based RL, (Abbas et al. 2020) combines the learned variance to adjust the weighting of targets derived from model planning to account for the limited model capacity, and (Buckman et al. 2018) leverages a bias-variance trade-off to determine a weighting scheme.

Regarding the prioritized sampling, PER (Schaul et al. 2016) addresses the initial step of considering the significance of different samples. In terms of the expected gradient perspective, shows the correlation between an  $l^1$  loss employing a prioritized sampling scheme and the uniformly sampled MSE loss. Additionally, (Pan et al. 2022) establishes an alternative equivalence between the uniformly sampled cubic loss and the prioritized MSE loss.

## Conclusions

In this paper, we introduced an emphasis-aware TD learning approach that takes into account the importance of historical states and the relative significance of propagating TD errors. Our method offers enhanced flexibility in selecting the emphasis. From various angles, we demonstrated its efficacy in challenging scenarios involving visitation imbalance and outcome noise. It not only restores balance to updates but also distinguishes between different noise levels, leading to improved predictions. We explored adaptive emphasis and confirmed its effectiveness in accelerating learning. Theoretical analysis established a contraction condition for algorithm convergence, offering practical insights into selecting the emphasis function. We also presented insights into extensions for DRL, including the proposed DAE and an equivalence between DTD(0) and PER. Additionally, we discussed potential forms of emphasis, which could be valuable when integrating with function approximations for future work.

## References

Abbas, Z.; Sokota, S.; Talvitie, E.; and White, M. 2020. Selective Dyna-Style Planning Under Limited Model Capacity. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, 1–10.

Anand, N. V.; and Precup, D. 2021. Preferential Temporal Difference Learning. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139, 286–296.

Benveniste, A.; Métivier, M.; and Priouret, P. 1990. *Adaptive Algorithms and Stochastic Approximations*, volume 22. Springer.

Boyan, J. A. 2002. Technical Update: Least-Squares Temporal Difference Learning. *Mach. Learn.*, 49(2-3): 233–246.

Buckman, J.; Hafner, D.; Tucker, G.; Brevdo, E.; and Lee, H. 2018. Sample-Efficient Reinforcement Learning with Stochastic Ensemble Value Expansion. In *Advances in Neural Information Processing Systems 31*, 8234–8244.

Chelu, V.; Borsa, D.; Precup, D.; and van Hasselt, H. 2022. Selective Credit Assignment. *CoRR*, abs/2202.09699.

Fujimoto, S.; Meger, D.; and Precup, D. 2020. An Equivalence between Loss Functions and Non-Uniform Sampling in Experience Replay. In *Advances in Neural Information Processing Systems 33*.

Imani, E.; Graves, E.; and White, M. 2018. An Off-policy Policy Gradient Theorem Using Emphatic Weightings. In *Advances in Neural Information Processing Systems 31*, 96–106.

Janner, M.; Fu, J.; Zhang, M.; and Levine, S. 2019. When to Trust Your Model: Model-Based Policy Optimization. In *Advances in Neural Information Processing Systems 32*, 12498–12509.

Khetarpal, K.; Klissarov, M.; Chevalier-Boisvert, M.; Bacon, P.; and Precup, D. 2020. Options of Interest: Temporal Abstraction with Interest Functions. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence*, 4444–4451.

Klissarov, M.; Fakoor, R.; Mueller, J. W.; Asadi, K.; Kim, T.; and Smola, A. J. 2022. Adaptive Interest for Emphatic Reinforcement Learning. In *NeurIPS*.

Ma, J. 2023. Distillation Policy Optimization. *CoRR*, abs/2302.00533.

McLeod, M.; Lo, C.; Schlegel, M.; Jacobsen, A.; Kumaraswamy, R.; White, M.; and White, A. 2021. Continual Auxiliary Task Learning. In *Advances in Neural Information Processing Systems 34*, 12549–12562.

Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; and Riedmiller, M. A. 2013. Playing Atari with Deep Reinforcement Learning. *CoRR*, abs/1312.5602.

Pan, Y.; Mei, J.; Farahmand, A.; White, M.; Yao, H.; Rohani, M.; and Luo, J. 2022. Understanding and mitigating the limitations of prioritized experience replay. In *Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence*, volume 180, 1561–1571.

Peng, X. B.; Abbeel, P.; Levine, S.; and van de Panne, M. 2018. DeepMimic: example-guided deep reinforcement learning of physics-based character skills. *ACM Trans. Graph.*, 37(4): 143.

Schaul, T.; Quan, J.; Antonoglou, I.; and Silver, D. 2016. Prioritized Experience Replay. In *4th International Conference on Learning Representations*.

Schulman, J.; Levine, S.; Abbeel, P.; Jordan, M. I.; and Moritz, P. 2015. Trust Region Policy Optimization. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37, 1889–1897.

Schulman, J.; Moritz, P.; Levine, S.; Jordan, M. I.; and Abbeel, P. 2016. High-Dimensional Continuous Control Using Generalized Advantage Estimation. In *4th International Conference on Learning Representations*.

Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. *CoRR*, abs/1707.06347.

Sutton, R. S. 1988. Learning to Predict by the Methods of Temporal Differences. *Mach. Learn.*, 3: 9–44.

Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement Learning: An Introduction*. The MIT Press, second edition.

Sutton, R. S.; Maei, H. R.; Precup, D.; Bhatnagar, S.; Silver, D.; Szepesvári, C.; and Wiewiora, E. 2009. Fast gradient-descent methods for temporal-difference learning with linear function approximation. In *Proceedings of the 26th Annual International Conference on Machine Learning*, volume 382, 993–1000.

Sutton, R. S.; Mahmood, A. R.; and White, M. 2016. An Emphatic Approach to the Problem of Off-policy Temporal-Difference Learning. *J. Mach. Learn. Res.*, 17: 73:1–73:29.

Sutton, R. S.; Modayil, J.; Delp, M.; Degris, T.; Pilarski, P. M.; White, A.; and Precup, D. 2011. Horde: a scalable real-time architecture for learning knowledge from unsupervised sensorimotor interaction. In *10th International Conference on Autonomous Agents and Multiagent Systems*, 761–768.

Sutton, R. S.; Precup, D.; and Singh, S. 1999. Between MDPs and Semi-MDPs: A Framework for Temporal Abstraction in Reinforcement Learning. *Artif. Intell.*, 112(1-2): 181–211.

Tsitsiklis, J. N.; and Roy, B. V. 1997. An analysis of temporal-difference learning with function approximation. *IEEE Trans. Autom. Control*, 42(5): 674–690.

Wang, J.; Liu, Y.; and Li, B. 2020. Reinforcement Learning with Perturbed Rewards. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence*, 6202–6209.

Yu, H. 2010. Convergence of Least Squares Temporal Difference Methods Under General Conditions. In *Proceedings of the 27th International Conference on Machine Learning*, 1207–1214.

Yu, H.; and Bertsekas, D. P. 2009. Convergence Results for Some Temporal Difference Methods Based on Least Squares. *IEEE Trans. Autom. Control*, 54(7): 1515–1531.