

Expected Eligibility Traces

Hado van Hasselt¹, Sephora Madjiheurem², Matteo Hessel¹
David Silver¹, André Barreto¹, Diana Borsa¹

¹ DeepMind

² University College London, UK

Abstract

The question of how to determine which states and actions are responsible for a certain outcome is known as the *credit assignment problem* and remains a central research question in reinforcement learning and artificial intelligence. *Eligibility traces* enable efficient credit assignment to the recent sequence of states and actions experienced by the agent, but not to counterfactual sequences that could also have led to the current state. In this work, we introduce *expected eligibility traces*. Expected traces allow, with a single update, to update states and actions that could have preceded the current state, even if they did not do so on this occasion. We discuss when expected traces provide benefits over classic (instantaneous) traces in temporal-difference learning, and show that sometimes substantial improvements can be attained. We provide a way to smoothly interpolate between instantaneous and expected traces by a mechanism similar to bootstrapping, which ensures that the resulting algorithm is a strict generalisation of TD(λ). Finally, we discuss possible extensions and connections to related ideas, such as successor features.

Motivation and Summary

Appropriate credit assignment has long been a major research topic in artificial intelligence (Minsky 1963). To make effective decisions and understand the world, we need to accurately associate events, like rewards or penalties, to relevant earlier decisions or situations. This is important both for learning accurate predictions, and for making good decisions.

Temporal credit assignment can be achieved with repeated temporal-difference (TD) updates (Sutton 1988). One-step TD updates propagate information slowly: when a surprising value is observed, the state immediately preceding it is updated, but no earlier states or decisions are updated. *Multi-step* updates (Sutton 1988; Sutton and Barto 2018) propagate information faster over longer temporal spans, speeding up credit assignment and learning. Multi-step updates can be implemented online using *eligibility traces* (Sutton 1988), without incurring significant additional computational expense, even if the time spans are long; these algorithms have computation that is independent of the temporal span of the predictions (van Hasselt and Sutton 2015).

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

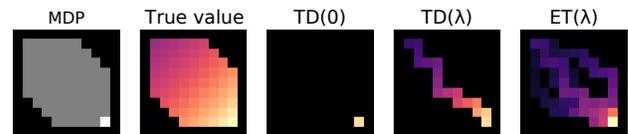


Figure 1: A comparison of TD(0), TD(λ), and the new expected-trace algorithm ET(λ) (with $\lambda = 0.9$). The MDP is illustrated on the left. Each episode, the agent moves randomly down and right from the top left to the bottom right, where any action terminates the episode. Reward on termination are +1 with probability 0.2, and zero otherwise—all other rewards are zero. We plot the value estimates after the first positive reward, which occurred in episode 5. We see a) TD(0) only updated the last state, b) TD(λ) updated the trajectory in this episode, and c) ET(λ) additionally updated trajectories from earlier (unrewarding) episodes.

Traces provide temporal credit assignment, but do not assign credit *counterfactually* to states or actions that *could* have led to the current state, but did not do so this time. Credit will eventually trickle backwards over the course of multiple visits, but this can take many iterations. As an example, suppose we collect a key to open a door, which leads to an unexpected reward. Using standard one-step TD learning, we would update the state in which the door opened. Using eligibility traces, we would also update the preceding trajectory, including the acquisition of the key. But we would not update other sequences that *could* have led to the reward, such as collecting a spare key or finding a different entrance.

The problem of credit assignment to counterfactual states may be addressed by learning a model, and using the model to propagate credit (cf. Sutton 1990; Moore and Atkeson 1993; Chelu, Precup, and van Hasselt 2020); however, it has often proven challenging to construct and use models effectively in complex environments (cf. van Hasselt, Hessel, and Aslanides 2019).

We introduce a new approach to counterfactual credit assignment, based on the concept of *expected eligibility traces*. We present a family of algorithms, which we call ET(λ), that use expected traces to update their predictions. We analyse the nature of these expected traces, and illustrate their benefits empirically in several settings—see Figure 1 for a first

illustration. We introduce a bootstrapping mechanism that provides a spectrum of algorithms between standard eligibility traces and expected eligibility traces, and also discuss ways to apply these ideas with deep neural networks. Finally, we discuss possible extensions and connections to related ideas such as successor features.

Background

Sequential decision problems can be modelled as Markov decision processes¹ (MDP) $(\mathcal{S}, \mathcal{A}, p)$ (Puterman 1994), with state space \mathcal{S} , action space \mathcal{A} , and a joint transition and reward distribution $p(r, s' | s, a)$. An agent selects actions according to its policy π , such that $A_t \sim \pi(\cdot | S_t)$ where $\pi(a | s)$ denotes the probability of selecting a in s , and observes random rewards and states generated according to the MDP, resulting in trajectories $\tau_{t:T} = \{S_t, A_t, R_{t+1}, S_{t+1}, \dots, S_T\}$. A central goal is to predict *returns* of future discounted rewards (Sutton and Barto 2018)

$$G_t \equiv G(\tau_{t:T}) = R_{t+1} + \gamma_{t+1}R_{t+2} + \gamma_{t+1}\gamma_{t+2}R_{t+3} + \dots \\ = \sum_{i=1}^T \gamma_{t+1}^{(i-1)} R_{t+i},$$

where T is for instance the time the current episode terminates or $T = \infty$, and where $\gamma_t \in [0, 1]$ is a (possibly constant) discount factor and $\gamma_t^{(n)} = \prod_{k=0}^{n-1} \gamma_{t+k}$, and $\gamma_t^{(0)} = 1$.

The value $v_\pi(s) = \mathbb{E}[G_t | S_t = s, \pi]$ of state s is the expected return for a policy π . Rather than writing the return as a random variable G_t , it will be convenient to instead write it as an explicit function $G(\tau)$ of the random trajectory τ . Note that $G(\tau_{t:T}) = R_{t+1} + \gamma_{t+1}G(\tau_{t+1:T})$.

We approximate the value with a function $v_w(s) \approx v_\pi(s)$. This can for instance be a table—with a single separate entry $w[s]$ for each state—a linear function of some input features, or a non-linear function such as a neural network with parameters w . The goal is to iteratively update w with

$$w_{t+1} = w_t + \Delta w_t$$

such that v_w approaches the true v_π . Perhaps the simplest algorithm to do so is the Monte Carlo (MC) algorithm

$$\Delta w_t \equiv \alpha(R_{t+1} + \gamma_{t+1}G(\tau_{t+1:T}) - v_w(S_t))\nabla_w v_w(S_t).$$

Monte Carlo is effective, but has high variance, which can lead to slow learning. TD learning (Sutton 1988; Sutton and Barto 2018) instead replaces the return with the current estimate of its expectation $v(S_{t+1}) \approx G(\tau_{t+1:T})$, yielding

$$\Delta w_t \equiv \alpha\delta_t\nabla_w v_w(S_t), \quad (1)$$

$$\text{where } \delta_t \equiv R_{t+1} + \gamma_{t+1}v_w(S_{t+1}) - v_w(S_t),$$

where δ_t is called the temporal-difference (TD) error. We can interpolate between these extremes, for instance with λ -returns which smoothly mix values and sampled returns:

$$G^\lambda(\tau_{t:T}) = R_{t+1} + \gamma_{t+1}((1-\lambda)v_w(S_{t+1}) + \lambda G^\lambda(\tau_{t+1:T})).$$

‘Forward view’ algorithms, like the MC algorithm, use returns that depend on future trajectories and need to wait until the

¹The ideas in this paper extend naturally to POMDPs (cf. ?).

end of an episode to construct their updates, which can take a long time. Conversely, ‘backward view’ algorithms rely only on past experiences and can update their predictions online, during an episode. Such algorithms build an *eligibility trace* (Sutton 1988; Sutton and Barto 2018). An example is TD(λ):

$$\Delta w_t \equiv \alpha\delta_t e_t, \quad \text{with } e_t = \gamma_t\lambda e_{t-1} + \nabla_w v_w(S_t),$$

where e_t is an accumulating eligibility trace. This trace can be viewed as a function $e_t \equiv e(\tau_{0:t})$ of the trajectory of past transitions. The TD update in (1) is known as TD(0), because it corresponds to using $\lambda = 0$. TD($\lambda = 1$) corresponds to an online implementation of the MC algorithm. Other variants exist, using other kinds of traces, and equivalences have been shown between these algorithms and their forward views that use λ -returns: these backward-view algorithms converge to the same solution as the corresponding forward view, and can in some cases yield equivalent weight updates (Sutton 1988; van Seijen and Sutton 2014; van Hasselt and Sutton 2015).

Expected Traces

The main idea of this paper is to use the concept of an *expected eligibility trace*, defined as

$$z(s) \equiv \mathbb{E}[e_t | S_t = s],$$

where the expectation is over the agent’s policy and the MDP dynamics. We introduce a concrete family of algorithms, which we call ET(λ) and ET(λ, η), that learn expected traces and use them in value updates. We analyse these algorithms theoretically, describe specific instances, and discuss computational and algorithmic properties.

ET(λ)

We propose to learn approximations $z_\theta(s) \approx z(s)$, with parameters $\theta \in \mathbb{R}^d$ (e.g., the weights of a neural network). One way to learn z_θ is by updating it toward the instantaneous trace e_t , by minimizing an empirical loss $\mathcal{L}(e_t, z_\theta(S_t))$. For instance, \mathcal{L} could be a component-wise squared loss, optimized with stochastic gradient descent:

$$\theta_{t+1} = \theta_t + \Delta\theta_t, \quad \text{where} \quad (2)$$

$$\Delta\theta_t = -\beta \frac{\partial}{\partial\theta} \frac{1}{2} (e_t - z_\theta(S_t))^\top (e_t - z_\theta(S_t)) \\ = \beta \frac{\partial z_\theta(S_t)}{\partial\theta} (e_t - z_\theta(S_t)), \quad (3)$$

where $\frac{\partial z_\theta(S_t)}{\partial\theta}$ is a $|\theta| \times |e|$ Jacobian² and β is a step size.

The idea is then to use $z_\theta(s) \approx \mathbb{E}[e_t | S_t = s]$ in place of e_t in the value update, which becomes

$$\Delta w_t \equiv \alpha\delta_t z_\theta(S_t). \quad (4)$$

We call this ET(λ). Below, we prove that this update can be unbiased and can have lower variance than TD(λ). Algorithm 1 shows pseudo-code for a concrete instance of ET(λ).

²The Jacobian-vector product can efficiently be computed (e.g., via auto-differentiation) with computational requirements that are comparable to the computation of the loss.

Algorithm 1 ET(λ)

```

1: initialise  $\mathbf{w}, \boldsymbol{\theta}$ 
2: for  $M$  episodes do
3:   initialise  $\mathbf{e} = \mathbf{0}$ 
4:   observe initial state  $S$ 
5:   repeat for each step in episode  $m$ 
6:     generate  $R$  and  $S'$ 
7:      $\delta \leftarrow R + \gamma v_{\mathbf{w}}(S') - v_{\mathbf{w}}(S)$ 
8:      $\mathbf{e} \leftarrow \gamma \lambda \mathbf{e} + \nabla_{\mathbf{w}} v_{\mathbf{w}}(S)$ 
9:      $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} + \beta \frac{\partial z_{\boldsymbol{\theta}}(S)}{\partial \boldsymbol{\theta}} (\mathbf{e} - z_{\boldsymbol{\theta}}(S))$ 
10:     $\mathbf{w} \leftarrow \mathbf{w} + \alpha \delta z_{\boldsymbol{\theta}}(S)$ 
11:   until  $S$  is terminal
12: end for
13: Return  $\mathbf{w}$ 

```

Interpretation and ET(λ, η)

We can interpret TD(0) as taking the MC update and replacing the return from the subsequent state, which is a function of the future trajectory, with a state-based estimate of its expectation: $v_{\mathbf{w}}(S_{t+1}) \approx \mathbb{E}[G(\tau_{t+1:T})|S_{t+1}]$. This becomes most clear when juxtaposing the updates:

$$\alpha(R_{t+1} + \gamma_{t+1}G(\tau_{t+1:T}) - v_{\mathbf{w}}(S_t))\nabla_{\mathbf{w}}v_{\mathbf{w}}(S_t), \quad (\text{MC})$$

$$\alpha(R_{t+1} + \gamma_{t+1}v_{\mathbf{w}}(S_{t+1}) - v_{\mathbf{w}}(S_t))\nabla_{\mathbf{w}}v_{\mathbf{w}}(S_t). \quad (\text{TD})$$

TD(λ) also uses a function of a trajectory: the trace \mathbf{e}_t . We propose replacing this as well with a function of state: the expected trace $\mathbf{z}_{\boldsymbol{\theta}}(S_t) \approx \mathbb{E}[\mathbf{e}(\tau_{0:t})|S_t]$. Again juxtaposing:

$$\Delta \mathbf{w}_t \equiv \alpha \delta_t \mathbf{e}(\tau_{0:t}), \quad (\text{TD}(\lambda))$$

$$\Delta \mathbf{w}_t \equiv \alpha \delta_t \mathbf{z}_{\boldsymbol{\theta}}(S_t). \quad (\text{ET}(\lambda))$$

We can interpolate smoothly between MC and TD(0) via λ . This is often useful to trade off variance of the return with potential bias of the value estimate. For instance, we might not have access to the true state s , and might instead have to rely on features $\mathbf{x}(s)$. Then we cannot always represent or learn the true values $v(s)$ —for instance different states may be aliased (Whitehead and Ballard 1991).

Similarly, when moving from TD(λ) to ET(λ) we replaced a trajectory-based trace with a state-based estimate. This might induce bias and, again, we can smoothly interpolate by using a recursively defined mixture trace \mathbf{y}_t , as defined as³

$$\mathbf{y}_t = (1 - \eta)\mathbf{z}_{\boldsymbol{\theta}}(S_t) + \eta(\gamma_t \lambda \mathbf{y}_{t-1} + \nabla_{\mathbf{w}} v_{\mathbf{w}}(S_t)). \quad (5)$$

This recursive usage of the estimates $\mathbf{z}_{\boldsymbol{\theta}}(s)$ at previous states is analogous to bootstrapping on future state values when using a λ -return, with the important difference that the arrow of time is opposite. This means we do not first have to convert this into a backward view: the quantity can already be computed from past experience directly. We call the algorithm that uses this mixture trace ET(λ, η):

$$\Delta \mathbf{w}_t \equiv \alpha \delta_t \mathbf{y}(S_t). \quad (\text{ET}(\lambda, \eta))$$

³While \mathbf{y}_t depends on both η and λ we leave this dependence implicit, as is conventional for traces.

Note that if $\eta = 1$ then $\mathbf{y}_t = \mathbf{e}_t$ equals the instantaneous trace: ET($\lambda, 1$) is equivalent to TD(λ). If $\eta = 0$ then $\mathbf{y}_t = \mathbf{z}_t$ equals the expected trace; the algorithm introduced earlier as ET(λ) is equivalent to ET($\lambda, 0$). By setting $\eta \in (0, 1)$, we can smoothly interpolate between these extremes.

Theoretical Analysis

We now analyse the new ET algorithms theoretically. First we show that if we use $\mathbf{z}(s)$ directly and s is Markov then the update has the same expectation as TD(λ) (though possibly with lower variance), and therefore also inherits the same fixed point and convergence properties.

Lemma 1. *If s is Markov, then*

$$\mathbb{E}[\delta_t \mathbf{e}_t | S_t = s] = \mathbb{E}[\delta_t | S_t = s] \mathbb{E}[\mathbf{e}_t | S_t = s].$$

Proof. In Appendix . □

Proposition 1. *Let \mathbf{e}_t be any trace vector, updated in any way. Let $\mathbf{z}(s) = \mathbb{E}[\mathbf{e}_t | S_t = s]$. Consider the ET(λ) algorithm $\Delta \mathbf{w}_t = \alpha_t \delta_t \mathbf{z}(S_t)$. For all Markov states s the expectation of this update is equal to the expected update under instantaneous trace \mathbf{e}_t , and its variance is lower or equal:*

$$\mathbb{E}[\alpha_t \delta_t \mathbf{z}(S_t) | S_t = s] = \mathbb{E}[\alpha_t \delta_t \mathbf{e}_t | S_t = s] \quad \text{and}$$

$$\mathbb{V}[\alpha_t \delta_t \mathbf{z}(S_t) | S_t = s] \leq \mathbb{V}[\alpha_t \delta_t \mathbf{e}_t | S_t = s],$$

where the second inequality holds component-wise for the update vector, and is strict when $\mathbb{V}[\mathbf{e}_t | S_t] > 0$.

Proof. We have

$$\begin{aligned} & \mathbb{E}[\alpha_t \delta_t \mathbf{e}_t | S_t = s] \\ &= \mathbb{E}[\alpha_t \delta_t | S_t = s] \mathbb{E}[\mathbf{e}_t | S_t = s] \quad (\text{Lemma 1}) \\ &= \mathbb{E}[\alpha_t \delta_t | S_t = s] \mathbf{z}(s) \\ &= \mathbb{E}[\alpha_t \delta_t \mathbf{z}(S_t) | S_t = s]. \end{aligned} \quad (6)$$

Denote the i -th component of $\mathbf{z}(S_t)$ by $z_{t,i}$ and the i -th component of \mathbf{e}_t by $e_{t,i}$. Then, we also have

$$\begin{aligned} & \mathbb{E}[(\alpha_t \delta_t z_{t,i})^2 | S_t = s] = \mathbb{E}[\alpha_t^2 \delta_t^2 | S_t = s] z_{t,i}^2 \\ &= \mathbb{E}[\alpha_t^2 \delta_t^2 | S_t = s] \mathbb{E}[e_{t,i}^2 | S_t = s]^2 \\ &= \mathbb{E}[\alpha_t^2 \delta_t^2 | S_t = s] (\mathbb{E}[e_{t,i}^2 | S_t = s] - \mathbb{V}[e_{t,i} | S_t = s]) \\ &\leq \mathbb{E}[\alpha_t^2 \delta_t^2 | S_t = s] \mathbb{E}[e_{t,i}^2 | S_t = s] \\ &= \mathbb{E}[(\alpha_t \delta_t e_{t,i})^2 | S_t = s], \end{aligned}$$

where the last step used the fact that s is Markov, and the inequality is strict when $\mathbb{V}[e_{t,i} | S_t] > 0$. Since the expectations are equal, as shown in (6), the conclusion follows. □

Interpretation Proposition 1 is a strong result: it holds for any trace update, including accumulating traces (Sutton 1984, 1988), replacing traces (Singh and Sutton 1996), dutch traces (van Seijen and Sutton 2014; van Hasselt, Mahmood, and Sutton 2014; van Hasselt and Sutton 2015), and future traces that may be discovered. It implies convergence of ET(λ) under the same conditions as TD(λ) (Dayan 1992; Peng 1993;

Tsitsiklis 1994) with lower variance when $\mathbb{V}[e_t|S_t] > 0$, which is the common case.

Next, we consider what happens if we violate the assumptions of Proposition 1. We start by analysing the case of a learned approximation $z_t(s) \approx z(s)$ that relies solely on observed experience.

Proposition 2. *Let e_t an instantaneous trace vector. Then let $z_t(s)$ be the empirical mean $z_t(s) = \frac{1}{n_t(s)} \sum_i^{n_t(s)} e_{t_i^s}$, where t_i^s denotes past times when we have been in state s , that is $S_{t_i^s} = s$, and $n_t(s)$ is the number of visits to s in the first t steps. Consider the expected trace algorithm $\mathbf{w}_{t+1} = \mathbf{w}_t + \alpha_t \delta_t z_t(S_t)$. If S_t is Markov, the expectation of this update is equal to the expected update with instantaneous traces e_t , while attaining a potentially lower variance:*

$$\mathbb{E}[\alpha_t \delta_t z_t(S_t) | S_t] = \mathbb{E}[\alpha_t \delta_t e_t | S_t] \quad \text{and} \\ \mathbb{V}[\alpha_t \delta_t z_t(S_t) | S_t] \leq \mathbb{V}[\alpha_t \delta_t e_t | S_t],$$

where the second inequality holds component-wise. The inequality is strict when $\mathbb{V}[e_t | S_t] > 0$.

Proof. In Appendix. \square

Interpretation Proposition 2 mirrors Proposition 1 but, importantly, covers the case where we estimate the expected traces from data, rather than relying on exact estimates. This means the benefits extend to this pure learning setting. Again, the result holds for any trace update. The inequality is typically strict when the path leading to state $S_t = s$ is stochastic (due to environment or policy).

Next we consider what happens if we do not have Markov states and instead have to rely on, possibly non-Markovian, features $\mathbf{x}(s)$. We then have to pick a function class and for the purpose of this analysis we consider linear expected traces $z_{\Theta}(s) = \Theta \mathbf{x}(s)$ and values $v_{\mathbf{w}}(s) = \mathbf{w}^{\top} \mathbf{x}(s)$, as convergence for non-linear values can not always be assured even for standard TD(λ) (Tsitsiklis and Van Roy 1997), without additional assumptions (e.g., Ollivier 2018; Brandfonbrener and Bruna 2020).

Proposition 3. *When using approximations $z_{\Theta}(s) = \Theta \mathbf{x}(s)$ and $v_{\mathbf{w}}(s) = \mathbf{w}^{\top} \mathbf{x}(s)$ then, if $(1 - \eta)\Theta + \eta\mathbb{I}$ is non-singular, $ET(\lambda, \eta)$ has the same fixed point as $TD(\lambda\eta)$.*

Proof. In Appendix. \square

Interpretation This result implies that linear $ET(\lambda, \eta)$ converges under similar conditions as linear $TD(\lambda')$ for $\lambda' = \lambda \cdot \eta$. In particular, when Θ is non-singular, using the approximation $z_{\Theta}(s) = \Theta \mathbf{x}(s)$ in $ET(\lambda, 0) = ET(\lambda)$ implies convergence to the fixed point of $TD(0)$.

Though $ET(\lambda, \eta)$ and $TD(\lambda\eta)$ have the same fixed point, the algorithms are not equivalent. In general, their updates are not the same. Linear approximations are more general than tabular functions (which are linear functions of a indicator vector for the current state), and we have already seen in Figure 1 that $ET(\lambda)$ behaves quite differently from both $TD(0)$ and $TD(\lambda)$, and we have seen its variance can be lower in Propositions 1 and 2. Interestingly, Θ resembles a preconditioner that speeds up the linear semi-gradient TD update,

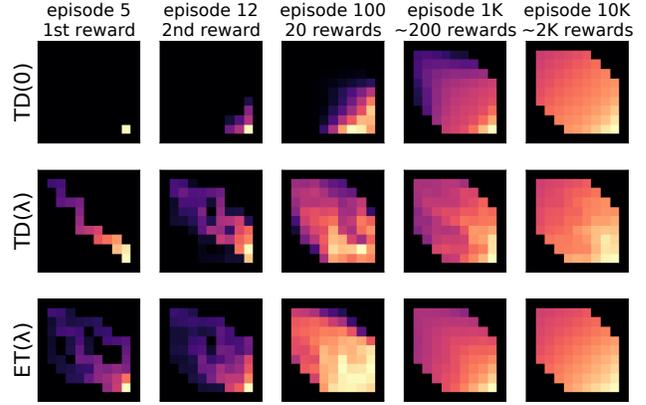


Figure 2: In the same setting as Figure 1, we show later value estimates after more rewards have been observed. $TD(0)$ learns slowly but steadily, $TD(\lambda)$ learns faster but with higher variance, and $ET(\lambda)$ learns both fast and stable.

similar to how second-order optimisation algorithms (Amari 1998; Martens 2016) precondition the gradient updates.

Empirical Analysis

From the insights above, we expect that $ET(\lambda)$ yields lower prediction errors because it has lower variance and aggregates information across episodes better. In this section we empirically investigate expected traces in several experiments. Whenever we refer to $ET(\lambda)$, this is equivalent to $ET(\lambda, 0)$.

An Open World

First consider the grid world depicted in Figure 1. The agent randomly moves right or down (excluding moves that would hit a wall), starting from the top-left corner. Any action in the bottom-right corner terminates the episode with +1 reward with probability 0.2, and 0 otherwise. All other rewards are 0.

Figure 1 shows value estimates after the first positive reward, which occurred in the fifth episode. $TD(0)$ updated a single state, $TD(\lambda)$ updated earlier states in that episode, and $ET(\lambda)$ additionally updated states from previous episodes.

Figure 2 additionally shows value estimates after the second reward (which occurred in episode 12), and after roughly 20, 200, and 2000 rewards (or 100, 1000, and 10,000 episodes, respectively). $ET(\lambda)$ converged faster than $TD(0)$, which propagated information slowly, and faster than $TD(\lambda)$, which exhibited higher variance. All step sizes decayed as $\alpha = \beta = \sqrt{1/k}$, where k is the current episode number.

A Multi-Chain

We now consider the multi-chain shown in Figure 3. We first compare $TD(\lambda)$ and $ET(\lambda)$ with tabular values on various variants of the multi-chain, corresponding to $m \in \{1, 2, 4, 8, \dots, 128\}$ parallel chains of length $n = 4$. The left-most plot in Figure 4 shows the average root mean squared error (RMSE) of the value predictions after 1024 episodes. We ran 10 seeds for each combination of step size $1/t^d$ with $d \in \{0.5, 0.8, 0.9, 1\}$ and $\lambda \in \{0, 0.5, 0.8, 0.9, 0.95, 1\}$.

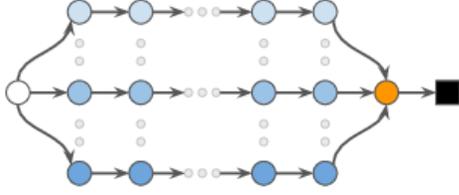


Figure 3: Multi-chain environment. Each episode starts in the left-most (white) state, and randomly transitions to one of m parallel (blue) chains of identical length n . After n steps, the agent always transitions to the same (orange) state, regardless of the chain it was in. The next step the episode terminates. Each reward is $+1$, except on termination when it either is $+1$ with probability 0.9 or -1 with probability 0.1 .

The left plot in Figure 4 shows value errors for different m , minimized over d and λ . The prediction error of $\text{TD}(\lambda)$ (blue) grew quickly with the number of parallel chains. $\text{ET}(\lambda)$ (orange) scaled better, because it updates values in multiple chains (from past episodes) upon receiving a surprising reward (e.g., -1) on termination. The other three plots in Figure 4 show value error as a function of λ for a subset of problems corresponding to $m \in \{8, 32, 128\}$. The dependence on λ differs across algorithms and problem instances, but $\text{ET}(\lambda)$ consistently achieved lower error than $\text{TD}(\lambda)$, especially with high λ . Further analysis, including on step-size sensitivity, is included in the appendix.

Next, we encode each state with a feature vector $\mathbf{x}(s)$ containing a binary indicator vector of the branch, a binary indicator of the progress along the chain, a bias that always equals one, and two binary features indicating when we are in the start (white) or bottleneck (orange) state. We extend the lengths of the chains to $n = 16$. Both $\text{TD}(\lambda)$ and $\text{ET}(\lambda)$ use a linear value function $v_{\mathbf{w}}(s) = \mathbf{w}^\top \mathbf{x}(s)$, and $\text{ET}(\lambda)$ uses a linear expected trace $z_{\Theta}(s) = \Theta \mathbf{x}(s)$. All updates use the same constant step size α . The left plot in Figure 5 shows the average root mean squared value error after 1024 episodes (averaged over 10 seeds). For each point the best constant step size $\alpha \in \{0.01, 0.03, 0.1\}$ (shared across all updates) and $\lambda \in \{0, 0.5, 0.8, 0.9, 0.95, 1\}$ is selected. $\text{ET}(\lambda)$ (orange) attained lower errors across all values of m (left plot), and for all λ (center two plots, for two specific m). The right plot shows results for smooth interpolations via η , for $\lambda = 0.9$ and $m = 16$. The full expected trace ($\eta = 0$) performed well here, we expect in other settings the additional flexibility of η could be beneficial.

Expected Traces in Deep Reinforcement Learning

(Deep) neural networks are a common choice of function class in reinforcement learning (e.g., Werbos 1990; Tesauro 1992, 1994; Bertsekas and Tsitsiklis 1996; Prokhorov and Wunsch 1997; Riedmiller 2005; van Hasselt 2012; Mnih et al. 2015; van Hasselt, Guez, and Silver 2016; Wang et al. 2016; Silver et al. 2016; Duan et al. 2016; Hessel et al. 2018). Eligibility traces are not very commonly combined with deep networks (but see Tesauro 1992; Elfving, Uchibe, and Doya 2018), perhaps in part because of the popularity of experience

replay (Lin 1992; Mnih et al. 2015; Horgan et al. 2018).

Perhaps the simplest way to extend expected traces to deep neural networks is to first separate the value function into a representation $\mathbf{x}(s)$ and a value $v_{(\mathbf{w}, \xi)}(s) = \mathbf{w}^\top \mathbf{x}_\xi(s)$, where \mathbf{x}_ξ is some (non-linear) function of the observations s .⁴ We can then apply the same expected trace algorithm as used in the previous sections by learning a separate linear function $z_{\Theta}(s) = \Theta \mathbf{x}(s)$ using the representation which is learned by backpropagating the value updates:

$$\mathbf{w}_{t+1} = \mathbf{w}_t + \alpha \delta z_{\Theta}(S_t),$$

$$\xi_{t+1} = \xi_t + \alpha \delta e_t^\xi,$$

$$\text{where } e_t^\xi = \gamma_t \lambda e_{t-1}^\xi + \nabla_{\xi} v_{(\mathbf{w}, \xi)}(S_t),$$

$$e_t^{\mathbf{w}} = \gamma_t \lambda e_{t-1}^{\mathbf{w}} + \nabla_{\mathbf{w}} v_{(\mathbf{w}, \xi)}(S_t),$$

and then updating Θ to minimise component-wise squared differences between $e_t^{\mathbf{w}}$ and $z_{\Theta}(S_t)$, as in (2) and (3).

Interesting challenges appear outside the fully linear case. First, the representation will itself be updated and will have its own trace e_t^ξ . Second, in the control case we optimise behaviour: the policy will change. Both these properties of the non-linear control setting imply that the expected traces must track a non-stationary target. We found that being able to track this rather quickly improved performance: the expected trace parameters Θ in the following experiment were updated with a relatively high step size of $\beta = 0.1$.

We tested this idea on two canonical Atari games: Pong and Ms. Pac-Man. The results in Figure 6 show that the expected traces helped speed up learning compared to the baseline which uses accumulating traces, for various step sizes. Unlike most prior work on this domain, which often relies on replay (Mnih et al. 2015; Schaul et al. 2016; Horgan et al. 2018) or parallel streams of experience (Mnih et al. 2016), these algorithms updated the values online from a single stream of experience. Further details on the experimental setup are given in the appendix.

These experiments demonstrate that the idea of expected traces extends to non-linear function approximation, such as deep neural networks. We consider this to be a rich area of further investigations. The results presented here are similar to earlier results (e.g., Mnih et al. 2015) and are not meant to compete with state-of-the-art performance results, which often depend on replay and much larger amounts of experience (e.g., Horgan et al. 2018).

Discussion and Extensions

We now discuss various interesting interpretations and relations, and discuss promising extensions.

Predecessor Features

For linear value functions the expected trace $z(s)$ can be expressed non-recursively as follows:

$$z(s) = \mathbb{E} \left[\sum_{n=0}^{\infty} \lambda_t^{(n)} \gamma_t^{(n)} \mathbf{x}_{t-n} \mid S_t = s \right], \quad (7)$$

⁴Here s denotes observations to the agent, not a full environment state— s is not assumed to be Markovian.

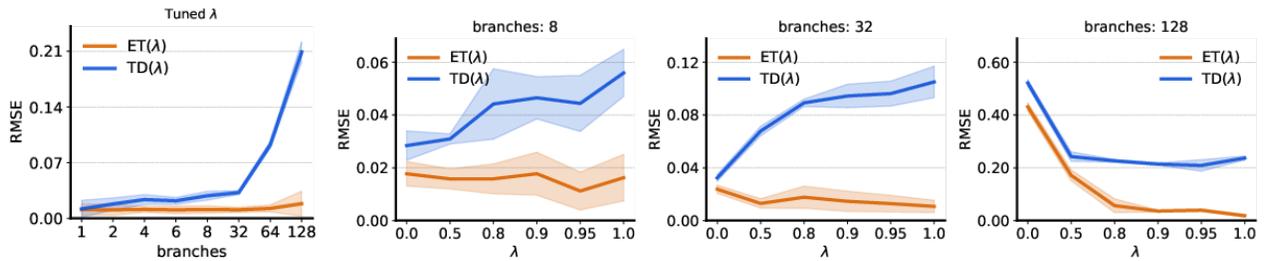


Figure 4: Prediction errors in the multi-chain. $ET(\lambda)$ (orange) consistently outperformed $TD(\lambda)$ (blue). Shaded areas depict standard errors across 10 seeds.

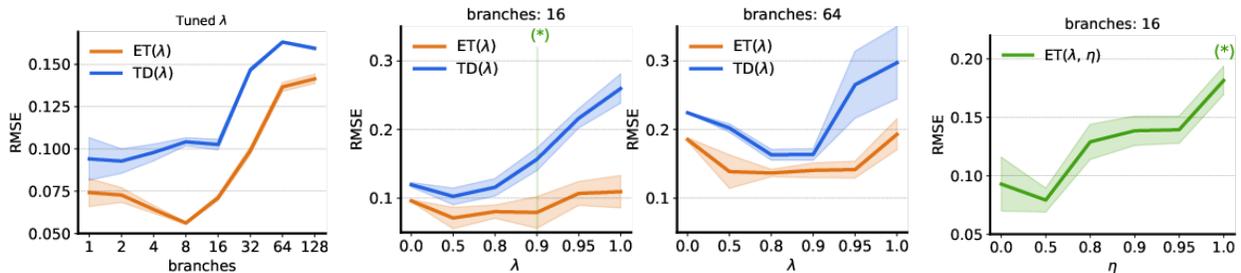


Figure 5: Comparing value error with linear function approximation a) as function of the number of branches (left), b) as function of λ (center two plots) and c) as function of η (right). The left three plots show comparisons of $TD(\lambda)$ (blue) and $ET(\lambda)$ (orange), showing $ET(\lambda)$ attained lower prediction errors. The right plot interpolates between these algorithms via $ET(\lambda, \eta)$, from $ET(\lambda) = ET(\lambda, 0)$ to $ET(\lambda, 1) = TD(\lambda)$, with $\lambda = 0.9$ (corresponding to a vertical slice indicated in the second plot).

where $\gamma_k^{(n)} \equiv \prod_{j=k-n}^k \gamma_j$. This is interestingly similar to the definition of *successor features* (Barreto et al. 2017):

$$\psi(s) = \mathbb{E} \left[\sum_{n=1}^{\infty} \gamma_{t+1}^{(n-1)} \mathbf{x}_{t+n} \mid S_t = s \right]. \quad (8)$$

The summation in (8) is over future features, while in (7) we have a sum over features already observed by the agent. We can thus think of linear expected traces as *predecessor features*. A similar connection was made in the tabular setting by Pitis (2018), relating source traces, which aim to estimate the source matrix $(I - \gamma P)^{-1}$, to successor representations (Dayan 1993). In a sense, the above generalises this insight. In addition to being interesting in its own right, this connection allows for an intriguing interpretation of $\mathbf{z}(s)$ as a multidimensional value function. Like with successor features, the features \mathbf{x}_t play the role of rewards, discounted with $\gamma \cdot \lambda$ rather than γ , and with time flowing backwards.

Although the predecessor interpretation only holds in the linear case, it is also of interest as a means to obtain a practical implementation of expected traces with non-linear function approximation, for instance applied only to the linear ‘head’ of a deep neural network. We used this ‘predecessor feature trick’ in our Atari experiments described earlier.

Relation to Model-Based Reinforcement Learning

Model-based reinforcement learning provides an alternative approach to efficient credit assignment. The general idea is to construct a model that estimates state-transition dynamics, and to update the value function based upon hypothetical

transitions drawn from the model (Sutton 1990), for example by prioritised sweeping (Moore and Atkeson 1993; van Seijen and Sutton 2013). In practice, model-based approaches have proven challenging in environments (such as Atari games) with rich perceptual observations, compared to model-free approaches that more directly update the agent’s policy and predictions (van Hasselt, Hessel, and Aslanides 2019).

In some sense, expected traces also construct a model of the environment—but one that differs in several key regards from standard state-to-state models used in model-based reinforcement learning. First, expected traces estimate *past* quantities rather than *future* quantities. Backward planning (e.g., Chelu, Precup, and van Hasselt 2020) is possible with explicit transition models, but is less common in practice. Second, expected traces estimate the accumulation of *gradients* over a multi-step trajectory, rather than trying to learn the full transition dynamics, thereby focusing only on those aspects that matter for the eventual weight update. Third, expected traces allow credit assignment across these potential past trajectories with a single update, without the iterative computation that is typically required when using a dynamics model. These differences may be important to side-step some of the challenges faced in model-based learning.

Batch Learning and Replay

We have mainly considered the online learning setting in this paper. It is often convenient to learn from batches of data, or replay transitions repeatedly, to enhance data efficiency. A natural extension is replay the experiences sequentially (e.g. Kapturowski et al. 2018), but perhaps alternatives exist. We

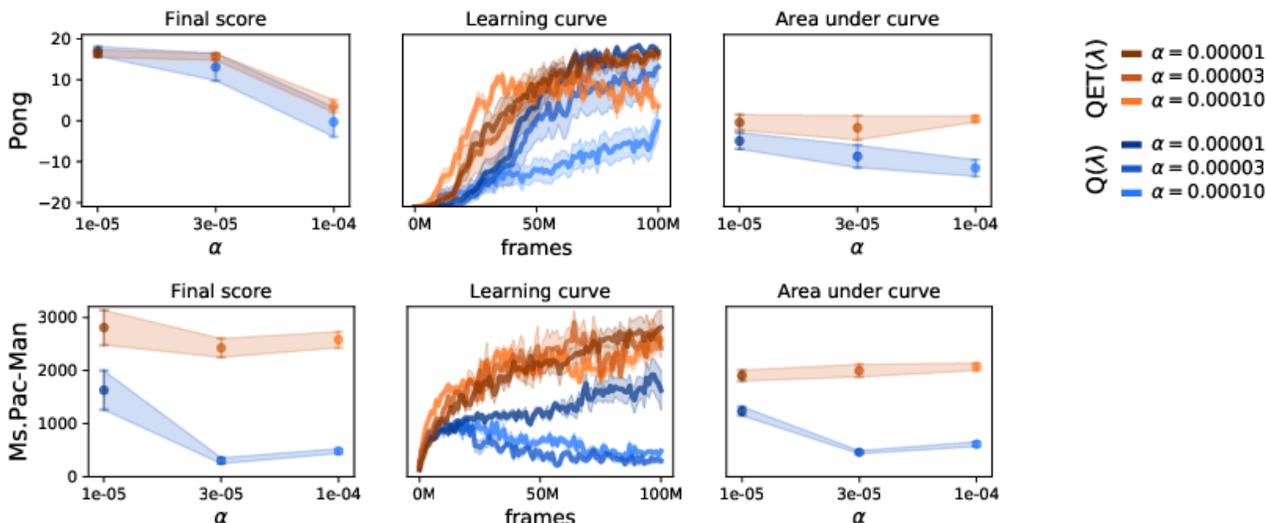


Figure 6: Performance of $Q(\lambda)$ ($\eta = 1$, blue) and $QET(\lambda)$ ($\eta = 0$, orange) on Pong and Ms.Pac-Man for various learning rates. Shaded regions show standard error across 10 random seeds. All results are for $\lambda = 0.95$. Further implementation details and hyper-parameters are in the appendix.

now discuss one potential extension.

We defined a mixed trace \mathbf{y}_t that mixes the instantaneous and expected traces. Optionally the expected trace \mathbf{z}_t can be updated towards the mixed trace \mathbf{y}_t as well, instead of towards the instantaneous trace \mathbf{e}_t . Analogously to TD(λ) we propose to then use at least one real step of data:

$$\Delta\theta_t \equiv \beta (\nabla_t + \gamma_t \lambda_t \mathbf{y}_{t-1} - \mathbf{z}_\theta(S_t))^\top \frac{\partial \mathbf{z}_\theta(S_t)}{\partial \theta}, \quad (9)$$

with $\nabla_t \equiv \nabla_{\mathbf{w}} v_{\mathbf{w}}(S_t)$. This is akin to a forward-view λ -return update, with $\nabla_{\mathbf{w}} v_{\mathbf{w}}(S_t)$ in the role of (vector) reward, and \mathbf{z}_θ of value, and discounted by $\lambda_t \gamma_t$, but reversed in time. In other words, this can be considered a sampled Bellman equation (Bellman 1957) but backward in time.

When we then choose $\eta = 0$, then $\mathbf{y}_{t-1} = \mathbf{z}_\theta(S_{t-1})$, and then the target in (9) only depends on a single transition. Interestingly, that means we can then learn expected traces from *individual* transitions, sampled out of temporal order, for instance in batch settings or when using replay.

Application to Other Traces

We can apply the idea of expected trace to more traces than considered here. We can for instance consider the characteristic eligibility trace used in REINFORCE (Williams 1992) and related policy-gradient algorithms (Sutton et al. 2000).

Another appealing application is to the follow-on trace or *emphasis*, used in emphatic temporal difference learning (Sutton, Mahmood, and White 2016) and related algorithms (e.g., Imani, Graves, and White 2018). Emphatic TD was proposed to correct an important issue with off-policy learning, which can be unstable and lead to diverging learning dynamics. Emphatic TD weights updates according to 1) the inherent interest in having accurate predictions in that state and, 2) the importance of predictions in that state for updating

other predictions. Emphatic TD uses scalar ‘follow-on’ traces to determine the ‘emphasis’ for each update. However, this follow-on trace can have very high, even infinite, variance. Instead, we might estimate and use its expectation instead of the instantaneous emphasis. A related idea was explored by Zhang, Boehmer, and Whiteson (2019) to obtain off-policy actor critic algorithms.

Conclusion

We have proposed a mechanism for efficient credit assignment, using the expectation of an eligibility trace. We have demonstrated this can sometimes speed up credit assignment greatly, and have analyzed concrete algorithms theoretically and empirically to increase understanding of the concept.

Expected traces have several interpretations. First, we can interpret the algorithm as counterfactually updating multiple possible trajectories leading up to the current state. Second, they can be understood as trading off bias and variance, which can be done smoothly via a unifying η parameter, between standard eligibility traces (low bias, high variance) and estimated traces (possibly higher bias, but lower variance). Furthermore, with tabular or linear function approximation we can interpret the resulting expected traces as predecessor states or features—object analogous to successor states or features, but time-reversed. Finally, we can interpret the linear algorithm as preconditioning the standard TD update, thereby potentially speeding up learning. These interpretations suggest that a variety of complementary ways to potentially extend these concepts and algorithms.

We have shown expected traces can already be used to enhance learning in non-linear settings (i.e., deep reinforcement learning), and in the control setting where we update the policy. Further work is needed to determine the full extent of the possibilities of these new algorithms.

References

- Amari, S. I. 1998. Natural gradient works efficiently in learning. *Neural computation* 10(2): 251–276. ISSN 0899-7667.
- Barreto, A.; Dabney, W.; Munos, R.; Hunt, J. J.; Schaul, T.; van Hasselt, H. P.; and Silver, D. 2017. Successor features for transfer in reinforcement learning. In *Advances in neural information processing systems*, 4055–4065.
- Bellemare, M. G.; Naddaf, Y.; Veness, J.; and Bowling, M. 2013. The Arcade Learning Environment: An Evaluation Platform for General Agents. *J. Artif. Intell. Res. (JAIR)* 47: 253–279.
- Bellman, R. 1957. *Dynamic Programming*. Princeton University Press.
- Bertsekas, D. P.; and Tsitsiklis, J. N. 1996. *Neuro-dynamic Programming*. Athena Scientific, Belmont, MA.
- Bradbury, J.; Frostig, R.; Hawkins, P.; Johnson, M. J.; Leary, C.; Maclaurin, D.; and Wanderman-Milne, S. 2018a. JAX: composable transformations of Python+NumPy programs. URL <http://github.com/google/jax>.
- Bradbury, J.; Frostig, R.; Hawkins, P.; Johnson, M. J.; Leary, C.; Maclaurin, D.; and Wanderman-Milne, S. 2018b. JAX: composable transformations of Python+NumPy programs. URL <http://github.com/google/jax>.
- Brandfonbrener, D.; and Bruna, J. 2020. Geometric Insights into the Convergence of Non-linear TD Learning. In *International Conference on Learning Representations*.
- Chelu, V.; Precup, D.; and van Hasselt, H. P. 2020. Forethought and Hindsight in Credit Assignment. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M. F.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 2270–2281.
- Dayan, P. 1992. The convergence of TD(λ) for general lambda. *Machine Learning* 8: 341–362.
- Dayan, P. 1993. Improving generalization for temporal difference learning: The successor representation. *Neural Computation* 5(4): 613–624.
- Duan, Y.; Chen, X.; Houthoofd, R.; Schulman, J.; and Abbeel, P. 2016. Benchmarking deep reinforcement learning for continuous control. In *International Conference on Machine Learning*, 1329–1338.
- Elfving, S.; Uchibe, E.; and Doya, K. 2018. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural Networks* 107: 3–11.
- Hennigan, T.; Cai, T.; Norman, T.; and Babuschkin, I. 2020. Haiku: Sonnet for JAX. URL <http://github.com/deepmind/dm-haiku>.
- Hessel, M.; Modayil, J.; van Hasselt, H. P.; Schaul, T.; Ostrovski, G.; Dabney, W.; Horgan, D.; Piot, B.; Azar, M.; and Silver, D. 2018. Rainbow: Combining Improvements in Deep Reinforcement Learning. *AAAI*.
- Horgan, D.; Quan, J.; Budden, D.; Barth-Maron, G.; Hessel, M.; van Hasselt, H. P.; and Silver, D. 2018. Distributed Prioritized Experience Replay. In *International Conference on Learning Representations*.
- Imani, E.; Graves, E.; and White, M. 2018. An Off-policy Policy Gradient Theorem Using Emphatic Weightings. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 31*, 96–106. Curran Associates, Inc. URL <http://papers.nips.cc/paper/7295-an-off-policy-policy-gradient-theorem-using-emphatic-weightings.pdf>.
- Kaelbling, L. P.; Littman, M. L.; and Cassandra, A. R. 1995. Planning and Acting in Partially Observable Stochastic Domains. Unpublished report.
- Kapturowski, S.; Ostrovski, G.; Quan, J.; Munos, R.; and Dabney, W. 2018. Recurrent experience replay in distributed reinforcement learning. In *International conference on learning representations*.
- Kingma, D. P.; and Adam, J. B. 2015. A method for stochastic optimization. In *International Conference on Learning Representation*.
- Lin, L. 1992. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine Learning* 8(3): 293–321.
- Martens, J. 2016. *Second-order optimization for neural networks*. University of Toronto (Canada).
- Minsky, M. 1963. Steps Toward Artificial Intelligence. In Feigenbaum, E.; and Feldman, J., eds., *Computers and Thought*, 406–450. McGraw-Hill, New York.
- Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T.; Harley, T.; Silver, D.; and Kavukcuoglu, K. 2016. Asynchronous Methods for Deep Reinforcement Learning. In *International Conference on Machine Learning*.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; Petersen, S.; Beattie, C.; Sadik, A.; Antonoglou, I.; King, H.; Kumaran, D.; Wierstra, D.; Legg, S.; and Hassabis, D. 2015. Human-level control through deep reinforcement learning. *Nature* 518(7540): 529–533.
- Moore, A. W.; and Atkeson, C. G. 1993. Prioritized Sweeping: Reinforcement Learning with less Data and less Time. *Machine Learning* 13: 103–130.
- Ollivier, Y. 2018. Approximate Temporal Difference Learning is a Gradient Descent for Reversible Policies. *CoRR* abs/1805.00869.
- Peng, J. 1993. *Efficient dynamic programming-based learning for control*. Ph.D. thesis, Northeastern University.
- Peng, J.; and Williams, R. J. 1996. Incremental Multi-step Q-learning. *Machine Learning* 22: 283–290.
- Pitis, S. 2018. Source Traces for Temporal Difference Learning. In McIlraith, S. A.; and Weinberger, K. Q., eds., *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, 3952–3959. AAAI Press.

- Pohlen, T.; Piot, B.; Hester, T.; Azar, M. G.; Horgan, D.; Budden, D.; Barth-Maron, G.; van Hasselt, H. P.; Quan, J.; Večerík, M.; Hessel, M.; Munos, R.; and Pietquin, O. 2018. Observe and look further: Achieving consistent performance on Atari. *arXiv preprint arXiv:1805.11593* .
- Prokhorov, D. V.; and Wunsch, D. C. 1997. Adaptive critic designs. *IEEE Transactions on Neural Networks* 8(5): 997–1007.
- Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc. New York, NY, USA.
- Riedmiller, M. 2005. Neural Fitted Q Iteration - First Experiences with a Data Efficient Neural Reinforcement Learning Method. In Gama, J.; Camacho, R.; Brazdil, P.; Jorge, A.; and Torgo, L., eds., *Proceedings of the 16th European Conference on Machine Learning (ECML'05)*, 317–328. Springer.
- Schaul, T.; Quan, J.; Antonoglou, I.; and Silver, D. 2016. Prioritized Experience Replay. In *International Conference on Learning Representations*. Puerto Rico.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature* 529(7587): 484–489.
- Singh, S. P.; and Sutton, R. S. 1996. Reinforcement Learning with replacing eligibility traces. *Machine Learning* 22: 123–158.
- Sutton, R. S. 1984. *Temporal Credit Assignment in Reinforcement Learning*. Ph.D. thesis, University of Massachusetts, Dept. of Comp. and Inf. Sci.
- Sutton, R. S. 1988. Learning to predict by the methods of temporal differences. *Machine learning* 3(1): 9–44.
- Sutton, R. S. 1990. Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. In *Proceedings of the seventh international conference on machine learning*, 216–224.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement Learning: An Introduction*. The MIT press, Cambridge MA.
- Sutton, R. S.; Mahmood, A. R.; and White, M. 2016. An Emphatic Approach to the Problem of Off-policy Temporal-Difference Learning. *Journal of Machine Learning Research* 17(73): 1–29.
- Sutton, R. S.; McAllester, D.; Singh, S.; and Mansour, Y. 2000. Policy gradient methods for reinforcement learning with function approximation. *Advances in Neural Information Processing Systems 13 (NIPS-00)* 12: 1057–1063.
- Tesauro, G. 1992. Practical Issues in Temporal Difference Learning. In Lippman, D. S.; Moody, J. E.; and Touretzky, D. S., eds., *Advances in Neural Information Processing Systems 4*, 259–266. San Mateo, CA: Morgan Kaufmann.
- Tesauro, G. J. 1994. TD-Gammon, a self-teaching backgammon program, achieves master-level play. *Neural computation* 6(2): 215–219.
- Tsitsiklis, J. N. 1994. Asynchronous stochastic approximation and Q-learning. *Machine Learning* 16: 185–202.
- Tsitsiklis, J. N.; and Van Roy, B. 1997. An analysis of temporal-difference learning with function approximation. *IEEE Transactions on Automatic Control* 42(5): 674–690.
- van Hasselt, H. P. 2012. Reinforcement Learning in Continuous State and Action Spaces. In Wiering, M. A.; and van Otterlo, M., eds., *Reinforcement Learning: State of the Art*, volume 12 of *Adaptation, Learning, and Optimization*, 207–251. Springer.
- van Hasselt, H. P.; Guez, A.; Hessel, M.; Mnih, V.; and Silver, D. 2016. Learning values across many orders of magnitude. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, 4287–4295.
- van Hasselt, H. P.; Guez, A.; and Silver, D. 2016. Deep reinforcement learning with double Q-Learning. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, 2094–2100.
- van Hasselt, H. P.; Hessel, M.; and Aslanides, J. 2019. When to use parametric models in reinforcement learning? In *Advances in Neural Information Processing Systems*, volume 32, 14322–14333.
- van Hasselt, H. P.; Mahmood, A. R.; and Sutton, R. S. 2014. Off-policy TD(λ) with a true online equivalence. In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence*, 330–339.
- van Hasselt, H. P.; Quan, J.; Hessel, M.; Xu, Z.; Borsa, D.; and Barreto, A. 2019. General non-linear Bellman equations. *arXiv preprint arXiv:1907.03687* .
- van Hasselt, H. P.; and Sutton, R. S. 2015. Learning to predict independent of span. *CoRR* abs/1508.04582.
- van Seijen, H.; and Sutton, R. S. 2013. Planning by Prioritized Sweeping with Small Backups. In *International Conference on Machine Learning*, 361–369.
- van Seijen, H.; and Sutton, R. S. 2014. True online TD(λ). In *International Conference on Machine Learning*, 692–700.
- Wang, Z.; de Freitas, N.; Schaul, T.; Hessel, M.; van Hasselt, H. P.; and Lanctot, M. 2016. Dueling Network Architectures for Deep Reinforcement Learning. In *International Conference on Machine Learning*. New York, NY, USA.
- Werbos, P. J. 1990. A menu of designs for reinforcement learning over time. *Neural networks for control* 67–95.
- Whitehead, S. D.; and Ballard, D. H. 1991. Learning to perceive and act by trial and error. *Machine Learning* 7(1): 45–83.
- Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning* 8: 229–256.
- Zhang, S.; Boehmer, W.; and Whiteson, S. 2019. Generalized off-policy actor-critic. In *Advances in Neural Information Processing Systems*, 2001–2011.