
Expected Eligibility Traces

Hado van Hasselt
DeepMind, London, UK

Sephora Madjiheurem
University College London, UK*

Matteo Hessel
DeepMind, London, UK

David Silver
DeepMind, London, UK

André Barreto
DeepMind, London, UK

Diana Borsa
DeepMind, 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.

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 true both when we care about learning accurate predictions, and when we care about making good decisions.

Temporal credit assignment can be achieved with repeated temporal-difference updates (Sutton, 1988). One-step temporal-difference (TD) updates propagate information slowly: when a surprising value is observed, the state immediately preceding it is updated, but none of the earlier states or decisions are updated. *Multi-step* updates (Sutton, 1988; Sutton & Barto, 2018) allow information to be propagated 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 prediction (van Hasselt & Sutton, 2015).

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.

To achieve more efficient learning that includes counterfactual credit assignment, we introduce the concept of *expected eligibility traces*. We present a concrete family of algorithms, which we call

*Work partially conducted during an internship at DeepMind.

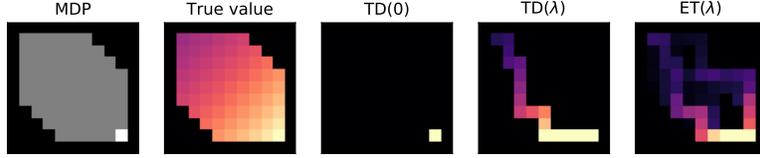


Figure 1: A comparison of TD(0), TD(λ), and ET(λ), which uses expected traces ($\lambda = 0.9$). The MDP is illustrated on the left. In each episode, the agent moves randomly down and right from the top-left corner to the bottom-right. Any action in the bottom right corner terminates the episode. The reward on termination is +1 with probability 0.2, and zero otherwise—all other rewards are zero. For each algorithm we show the value estimates after the first positive reward was observed: in this case in episode 7. We see a) TD(0) only updated the last state, b) TD(λ) updated the trajectory in this episode, and c) ET(λ) additionally updates trajectories from earlier (unrewarding) episodes.

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. 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.

1 Background

Sequential decision problems can be modelled as a Markov decision process² (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 , 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 & Barto, 2018)

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

where T is the end of the current episode or $T = \infty$, and where $\gamma_t \in [0, 1]$ is a (possibly constant) discount factor and $\gamma_t^{(i)} = \prod_{k=1}^i \gamma_{t+k}$. The value $v_\pi(s) = \mathbb{E}[G_t | S_t = s, \pi]$ of state s is the expected return. 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 $v_{\mathbf{w}}(s) \approx v_\pi(s)$. 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 \mathbf{w} . The goal is to iteratively update \mathbf{w} such that $v_{\mathbf{w}}$ approaches the true v_π . Perhaps the simplest algorithm to do so is the Monte Carlo (MC) algorithm

$$\mathbf{w}_{t+1} \leftarrow \mathbf{w}_t + \alpha(R_{t+1} + \gamma_{t+1}G(\tau_{t+1:T}) - v_{\mathbf{w}}(S_t))\nabla_{\mathbf{w}}v_{\mathbf{w}}(S_t).$$

Monte Carlo is effective, but high variance, which can lead to slow learning. A popular alternative is to replace the return with the current estimate of its expectation $v(S_{t+1}) \approx G(\tau_{t+1})$, which yields TD learning (Sutton, 1988; Sutton & Barto, 2018)

$$\mathbf{w}_{t+1} \leftarrow \mathbf{w}_t + \alpha \underbrace{(R_{t+1} + \gamma_{t+1}v_{\mathbf{w}}(S_{t+1}) - v_{\mathbf{w}}(S_t))}_{\equiv \delta_t} \nabla_{\mathbf{w}}v_{\mathbf{w}}(S_t), \quad (1)$$

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

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

‘Forward view’ algorithms, like the MC algorithm, use returns that depend on future trajectories and need to wait until the end of an episode to construct their updates, which can take a long time.

²The ideas extend naturally to POMDPs (cf. Kaelbling et al., 1995).

Conversely, ‘backward view’ algorithms rely only on past experiences and can update their predictions online, during an episode. Such algorithms build a *eligibility trace* (Sutton, 1988; Sutton & Barto, 2018). An example is the TD(λ) update

$$\mathbf{w}_{t+1} \leftarrow \mathbf{w}_t + \alpha \delta_t \mathbf{e}_t, \quad \text{with} \quad \mathbf{e}_t = \gamma_t \lambda \mathbf{e}_{t-1} + \nabla_{\mathbf{w}} v_{\mathbf{w}}(S_t),$$

where \mathbf{e}_t is an accumulating eligibility trace. This trace can be viewed as a function $\mathbf{e}_t \equiv e(\tau_{0:t})$ of the trajectory of past transitions. The TD update (1) is also known as TD(0), because it corresponds to using $\lambda = 0$, whereas $\lambda = 1$ yields a backward-view implementation of the Monte Carlo algorithm. Other variants exist, using other kinds of traces, and equivalences have been shown between these algorithms and their forward view using 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 & Sutton, 2014; van Hasselt & Sutton, 2015).

2 Expected traces

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

$$\mathbf{z}(s) \equiv \mathbb{E}[\mathbf{e}_t \mid S_t = s],$$

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, considerations, and design choices.

2.1 ET(λ)

We propose to learn approximations $\mathbf{z}_{\theta}(S_t) \approx \mathbf{z}(S_t)$, with parameters $\theta \in \mathbb{R}^d$ (e.g., the weights of a neural network). One way to learn \mathbf{z}_{θ} is to keep track of the instantaneous trace \mathbf{e}_t and then to update the expected trace estimates towards this trace, by minimising an empirical loss $\mathcal{L}(\mathbf{e}_t, \mathbf{z}_{\theta}(S_t))$. For instance, \mathcal{L} can be a component-wise squared loss that we optimise with stochastic gradient descent:

$$\theta_{t+1} = \theta_t - \beta \frac{\partial}{\partial \theta} \frac{1}{2} (\mathbf{e}_t - \mathbf{z}_{\theta}(S_t))^{\top} (\mathbf{e}_t - \mathbf{z}_{\theta}(S_t)) = \theta_t + \beta \frac{\partial \mathbf{z}_{\theta}(S_t)}{\partial \theta} (\mathbf{e}_t - \mathbf{z}_{\theta}(S_t)),$$

where $\frac{\partial \mathbf{z}_{\theta}(S_t)}{\partial \theta}$ is a $|\theta| \times |e|$ Jacobian³ and β is a step-size parameter.

The idea is then to use $\mathbf{z}_{\theta}(s) \approx \mathbb{E}[\mathbf{e}_t \mid S_t = s]$ in place of \mathbf{e}_t in the value update, which becomes

$$\mathbf{w}_{t+1} = \mathbf{w}_t + \delta_t \mathbf{z}_{\theta}(S_t). \quad (2)$$

We call this ET(λ). In Section 3 we prove 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).

2.2 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: $v(S_{t+1}) \approx \mathbb{E}[G(\tau_{t+1:T}) \mid S_{t+1}]$. This becomes most clear when juxtaposing the updates

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

$$\mathbf{w}_{t+1} \leftarrow \mathbf{w}_t + \alpha (R_{t+1} + \gamma_{t+1} v_{\mathbf{w}}(S_{t+1}) - v_{\mathbf{w}}(S_t)) \nabla_{\mathbf{w}} v_{\mathbf{w}}(S_t). \quad (\text{TD}(0))$$

³Auto-differentiation can efficiently compute this update with comparable computation to the loss calculation.

Algorithm 1 ET(λ)

```

1: initialise  $\mathbf{w}, \theta$ 
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:      $\theta \leftarrow \theta + \beta \frac{\partial \mathbf{z}_{\theta}(S)}{\partial \theta} (\mathbf{e} - \mathbf{z}_{\theta}(S))$ 
10:     $\mathbf{w} \leftarrow \mathbf{w} + \alpha \delta \mathbf{z}_{\theta}(S)$ 
11:   until  $S$  is terminal
12: end for
13: Return  $\mathbf{w}$ 

```

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

$$\mathbf{w}_{t+1} \leftarrow \mathbf{w}_t + \alpha \delta_t e(\tau_{0:t}), \quad (\text{TD}(\lambda))$$

$$\mathbf{w}_{t+1} \leftarrow \mathbf{w}_t + \alpha \delta_t \mathbf{z}_\theta(S_t). \quad (\text{ET}(\lambda))$$

When switching from MC to TD(0), the dependence on the trajectory was replaced with a state-based value estimate to bootstrap on. 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 have to rely on a small aperture on the world that provides us with features $\mathbf{x}(s)$ instead of the full state s . Then we cannot always represent or learn the true values $v(s)$ —for instance different states may be aliased (Whitehead & Ballard, 1991).

Similarly, when moving from TD(λ) to ET(λ) we replaced a trajectory-based trace with a state-based estimate. Again this might induce bias. Once again, we can smoothly interpolate by using a recursively defined mixture trace we will call \mathbf{y}_t which is defined as⁴

$$\mathbf{y}_t = (1 - \eta) \mathbf{z}_\theta(S_t) + \eta (\gamma_t \lambda \mathbf{y}_{t-1} + \nabla_{\mathbf{w}} v_{\mathbf{w}}(S_t)). \quad (3)$$

This recursive usage of the estimates $\mathbf{z}_\theta(s)$ 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):

$$\mathbf{w}_{t+1} \leftarrow \mathbf{w}_t + \alpha \delta_t \mathbf{y}(S_t). \quad (\text{ET}(\lambda, \eta))$$

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.

3 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.

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 $\mathbf{w}_{t+1} = \mathbf{w}_t + \alpha_t \delta_t \mathbf{z}(S_t)$. For all Markov s the expectation of this update is equal to the expected update with instantaneous trace \mathbf{e}_t , and the variance is lower or equal:*

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

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

$$\begin{aligned} \mathbb{E}[\alpha_t \delta_t \mathbf{e}_t | S_t = s] &= \mathbb{E}[\alpha_t \delta_t | S_t = s] \mathbb{E}[\mathbf{e}_t | S_t = s] && \text{(as } s \text{ is Markov)} \\ &= \mathbb{E}[\alpha_t \delta_t | S_t = s] \mathbf{z}(s) \\ &= \mathbb{E}[\alpha_t \delta_t \mathbf{z}(S_t) | S_t = s]. \end{aligned} \quad (4)$$

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

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

where the last step used the fact that s is Markov, and the inequality is strict when $\mathbb{V}[e_t | S_t] > 0$. Since the expectations are equal, as shown in (4), the conclusion follows. \square

⁴While \mathbf{y}_t depends on both η and λ we leave this dependence implicit, as is conventional for traces.

Interpretation Proposition 1 is a strong result: it holds for any trace update, including accumulating traces (Sutton, 1984, 1988), replacing traces (Singh & Sutton, 1996), dutch traces (van Seijen & Sutton, 2014; van Hasselt et al., 2014; van Hasselt & Sutton, 2015), and future traces that may be discovered. It implies convergence of $ET(\lambda)$ under the same conditions as $TD(\lambda)$ (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 -s denote 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 $\mathbf{w}_{t+1} = \mathbf{w}_t + \alpha_t \delta_t z_t$. 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:*

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

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

Proof. In Appendix A.1. □

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 $TD(\lambda)$ (Tsitsiklis & Van Roy, 1997), without additional assumptions (e.g., Ollivier, 2018; Brandfonbrener & Bruna, 2020). The following property of the mixture trace is used in the proposition below.

Proposition 3. *The mixture trace \mathbf{y}_t defined in (3) can be interpreted as a trace $\mathbf{y}_t = \mu \mathbf{y}_{t-1} + \mathbf{x}_t$ with decay parameter $\mu = \lambda \cdot \eta$ and signal $\mathbf{x}_t = (1 - \eta) z_\Theta(S_t) + \eta \nabla_{\mathbf{w}} v_{\mathbf{w}}(S_t)$, such that*

$$\mathbf{y}_t = \sum_{k=0}^t (\eta \gamma \lambda)^k [(1 - \eta) z_\Theta(S_{t-k}) + \eta \nabla_{\mathbf{w}} v_{\mathbf{w}}(S_{t-k})]. \quad (5)$$

Proof. In Appendix A.2. □

Recall $\mathbf{y}_t = e_t$ when $\eta = 1$, and $\mathbf{y}_t = z_\Theta(S_t)$ when $\eta = 0$, as can be verified by inspecting (5) (and using the convention $0^0 = 1$). We use this proposition to prove the following.

Proposition 4. *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. By Proposition 3 we have that \mathbf{y}_t can be re-written as:

$$\begin{aligned} \mathbf{y}_t &= \sum_{k=0}^t (\eta \gamma \lambda)^k [(1 - \eta) z_\Theta(S_{t-k}) + \eta \mathbf{x}(S_{t-k})] \\ &= \sum_{k=0}^t (\eta \gamma \lambda)^k [(1 - \eta) \Theta \mathbf{x}(S_{t-k}) + \eta \mathbf{x}(S_{t-k})] = [(1 - \eta) \Theta + \eta \mathbb{I}] \underbrace{\sum_{k=0}^t (\eta \gamma \lambda)^k \mathbf{x}(S_{t-k})}_{\text{instantaneous trace } e_t^{\lambda\eta}}. \end{aligned} \quad (6)$$

We examine the fixed point \mathbf{w}_* of the algorithm using this approximation of the expected trace:

$$\mathbb{E}[\delta_t \mathbf{y}_t] = \mathbb{E}[\mathbf{y}_t(R_{t+1} + \gamma \mathbf{x}(S_{t+1}))^\top \mathbf{w}_* - \mathbf{x}(S_t)^\top \mathbf{w}_*] = \mathbf{0}.$$

This implies the fixed point is

$$\mathbf{w}_* = \mathbb{E} \left[\mathbf{y}_t (\gamma \mathbf{x}(S_{t+1}) - \mathbf{x}(S_t))^\top \right]^{-1} \mathbb{E} \left[\mathbf{y}_t R_{t+1} \right].$$

Now, plugging in the relation in (6) above, we get:

$$\begin{aligned} \mathbf{w}_* &= \mathbb{E} \left[\left[(1 - \eta) \Theta + \eta \mathbb{I} \right] e_t^{\lambda \eta} (\gamma \mathbf{x}(S_{t+1}) - \mathbf{x}(S_t))^\top \right]^{-1} \mathbb{E} \left[\left[(1 - \eta) \Theta + \eta \mathbb{I} \right] e_t^{\lambda \eta} R_{t+1} \right] \\ &= \mathbb{E} \left[e_t^{\lambda \eta} (\gamma \mathbf{x}(S_{t+1}) - \mathbf{x}(S_t))^\top \right]^{-1} \left[(1 - \eta) \Theta + \eta \mathbb{I} \right]^{-1} \left[(1 - \eta) \Theta + \eta \mathbb{I} \right] \mathbb{E} \left[e_t^{\lambda \eta} R_{t+1} \right] \\ &= \mathbb{E} \left[e_t^{\lambda \eta} (\gamma \mathbf{x}(S_{t+1}) - \mathbf{x}(S_t))^\top \right]^{-1} \mathbb{E} \left[e_t^{\lambda \eta} R_{t+1} \right]. \end{aligned}$$

This last term is the fixed point for TD($\lambda\eta$). □

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

Though ET(λ, η) 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(λ) behaves quite differently from both TD(0) and TD(λ), 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, similar to how second-order optimisation algorithms (Amari, 1998; Martens, 2016) precondition the gradient updates.

4 Empirical analysis

From the insights above, we expect that ET(λ) will yield 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(λ), this is equivalent to ET($\lambda, 0$).

4.1 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 the value estimates after the first positive reward, which occurred in the seventh episode. TD(0) updated a single state, TD(λ) updated earlier states in that episode, and ET(λ) additionally updated states from previous episodes. Figure 2 shows the values after the second reward,

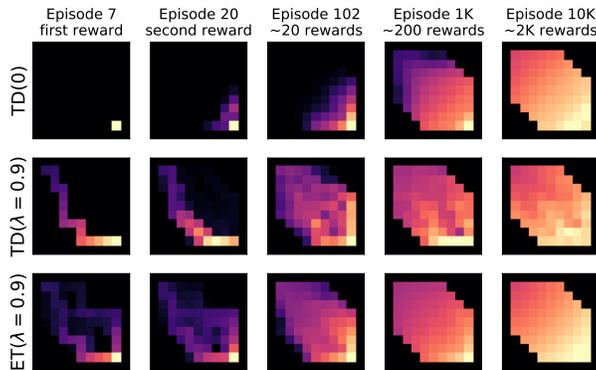


Figure 2: In the same setting as Figure 1, we show later value estimates after more rewards have been observed. TD(0) learns slowly but steadily, TD(λ) learns faster but with higher variance, and ET(λ) learns both fast and stable.

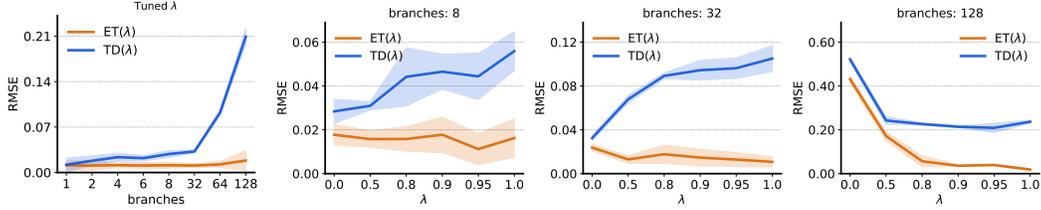


Figure 3: Prediction errors in the multi-chain. $ET(\lambda)$ (orange) consistently outperformed $TD(\lambda)$ (blue). Shaded areas depict standard errors across 10 seeds.

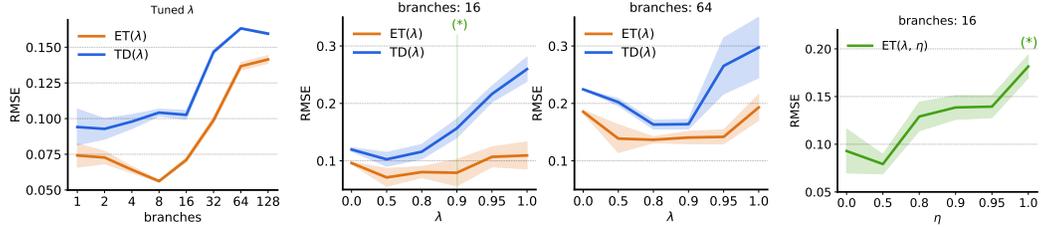
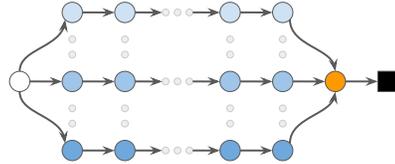


Figure 4: Comparing value error with linear function approximation a) as function of the number of branches (left), b) as function of λ (center two plots) and c) as function of η . The first three plots compare $TD(\lambda)$ (blue) to $ET(\lambda)$ (orange). The right plot examines $ET(\lambda, \eta)$ from $ET(\lambda) = ET(\lambda, 0)$ to $ET(\lambda, 1) = TD(\lambda)$, for $\lambda = 0.9$ (corresponding to a vertical slice indicated in the second plot).

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 than $TD(\lambda)$, which had higher variance. All step sizes decayed as $\alpha = \beta = \sqrt{1/k}$, where k is the current episode number.

4.2 A multi-chain

In the multi-chain shown on the right the agent starts each episode in the (white) state on the left, and then 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 $(1 - p)$ or -1 with probability p .



We first compare $TD(\lambda)$ and $ET(\lambda)$ with tabular values on various variants of the multi-chain, corresponding $n = 4$ and $m \in \{1, 2, 4, 8, \dots, 128\}$. The left-most plot in Figure 3 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 $\lambda \in \{0, 0.5, 0.8, 0.9, 0.95, 1\}$.

The left plot in Figure 3 shows value errors for different m , minimized over d and λ . The prediction error of $TD(\lambda)$ (blue) grew quickly with the number of parallel chains. $ET(\lambda)$ (orange) scaled better, because it can update values in other chains (from past episodes) upon receiving a surprising reward (e.g., -1) on termination. The other three plots in Figure 3 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; in all cases $ET(\lambda)$ achieved lower error than $TD(\lambda)$.

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 $TD(\lambda)$ and $ET(\lambda)$ use a linear value function $v_{\mathbf{w}}(s) = \mathbf{w}^\top \mathbf{x}(s)$, and $ET(\lambda)$ 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 4 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. $ET(\lambda)$ (orange) attained lower errors across all values of m (left plot), and for all λ (center two plots, for two specific m). The right

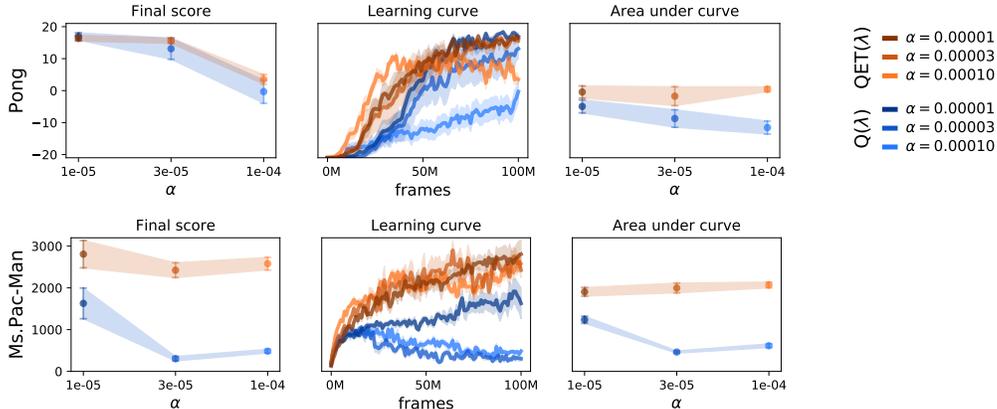


Figure 5: Performance of $Q(\lambda)$ ($\eta = 1$, blue) and $QET(\lambda)$ ($\eta = 0$, orange) on Pong and Ms.Pac-Man for various learning rates. Shaded regions show standard error across 10 random seeds. All results are for $\lambda = 0.95$. Further implementation details and hyper-parameters are in the appendix.

plot shows results for smooth interpolations via η , for $\lambda = 0.9$ and $m = 16$. The full expected trace ($\eta = 0$) performed well here, in other settings the additional flexibility of η could be more beneficial.

4.3 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 & Tsitsiklis, 1996; Prokhorov & Wunsch, 1997; Riedmiller, 2005; van Hasselt, 2012; Mnih et al., 2015; van Hasselt et al., 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 et al., 2018), perhaps 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 $v_{(\mathbf{w}, \boldsymbol{\xi})}(s) = \mathbf{w}^\top \mathbf{x}_\xi(s)$, where \mathbf{x}_ξ is some (non-linear) function of the observations s .⁵ 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:

$$\boldsymbol{\xi}_{t+1} \leftarrow \boldsymbol{\xi}_t + \alpha \delta e_t^\xi \quad \text{and} \quad \mathbf{w}_{t+1} \leftarrow \mathbf{w}_t + \alpha \delta z_{\Theta}(S_t),$$

where $e_t^\xi = \gamma_t \lambda e_{t-1}^\xi + \nabla_{\boldsymbol{\xi}} v_{(\mathbf{w}, \boldsymbol{\xi})}(S_t)$, $e_t^w = \gamma_t \lambda e_{t-1}^w + \nabla_{\mathbf{w}} v_{(\mathbf{w}, \boldsymbol{\xi})}(S_t)$, and Θ is updated by minimising the sum of component-wise squared differences between e_t^w and $z_{\Theta}(S_t)$.

Interesting challenges appear that were not present in the fully linear case. First, the representation will itself be updated. 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 it works best to track rather quickly: the expected trace parameters Θ in the following experiment were updated with a step size of $\beta = 0.1$.

We tested this idea on two canonical Atari games: Pong and Ms.Pac-Man. The results in Figure 5 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 are given in Appendix A.3.

The purpose of these experiments is to show that the idea of expected traces extends to non-linear function approximation, such as deep neural networks. There are other ways to do this than presented here, we consider this a rich area of further investigations. We have no doubt the methods can be

⁵Here s denotes observations to the agent, not a full environment state— s is not assumed to be Markovian.

improved further and generalised to other domains. 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).

5 Discussion and extensions

The work presented here can be interpreted and extended in various interesting ways.

5.1 Predecessor features

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

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

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

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

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*. In addition to being interesting in its own right, this connection allows for an intriguing interpretation of $z(s)$ 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 , though here time flows 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.

5.2 Relationship 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 & Atkeson, 1993; van Seijen & Sutton, 2013). In practice, model-based approaches have proven challenging in environments (such as Atari games) with rich perceptual observations (van Hasselt et al., 2019), compared to model-free approaches that more directly estimate the effects of the agent’s actions and predictions.

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. Second, they estimate the accumulation of gradients over a multi-step trajectory, rather than transition dynamics. Third, they allow credit assignment across these potential past trajectories with a single update, without the iterative computation that is typically required when using a more explicit model. We speculate that these differences may side-step some of the challenges faced in model-based learning.

5.3 Batch learning and replay

We have mainly considered the online learning setting in this paper. Sometimes it is convenient to learn from batches of data, or replay transitions repeatedly, to enhance data efficiency. The most natural extension is to go through the experience sequentially (e.g. Kapturowski et al., 2018), but perhaps alternatives exist. We discuss one such potential extension here.

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

e_t . Analogously to TD(λ) we propose to then use at least one real step of data:

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t + \beta (\nabla_{\mathbf{w}} v(S_t) + \gamma_t \lambda_t \mathbf{y}_{t-1} - \mathbf{z}_{\boldsymbol{\theta}}(S_t))^\top \frac{\partial \mathbf{z}_{\boldsymbol{\theta}}(S_t)}{\partial \boldsymbol{\theta}}. \quad (9)$$

This is akin to a forward-view update with a λ -return, where $\nabla_{\mathbf{w}} v(S_t)$ takes the role of (vector) reward, and $\mathbf{z}_{\boldsymbol{\theta}}$ of value, 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 $\mathbf{y}_{t-1} = \mathbf{z}_{\boldsymbol{\theta}}(S_{t-1})$, 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.

5.4 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 et al., 2016) and related algorithms (e.g., Imani et al., 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 we have in accurate predictions in that state and, 2) the importance of those predictions in the updates to other states. The algorithm uses a scalar ‘follow-on’ trace to determine the ‘emphasis’ for each update. However, this follow-on trace can have very high, even infinite, variance. It is an intriguing idea to instead estimate and use its expectation instead of the instantaneous emphasis.

6 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 analysed concrete instances 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. These interpretations suggest that a variety of complementary ways to potentially extend these concepts and algorithms.

References

- Amari, S. I. (1998). Natural gradient works efficiently in learning. *Neural computation*, 10(2), 251–276.
- Barreto, A., Dabney, W., Munos, R., Hunt, J. J., Schaul, T., van Hasselt, H. P., & Silver, D. (2017). Successor features for transfer in reinforcement learning. In *Advances in neural information processing systems*, (pp. 4055–4065).
- Bellemare, M. G., Naddaf, Y., Veness, J., & 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. & Tsitsiklis, J. N. (1996). *Neuro-dynamic Programming*. Athena Scientific, Belmont, MA.
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., & Wanderman-Milne, S. (2018). JAX: composable transformations of Python+NumPy programs.

- Brandfonbrener, D. & Bruna, J. (2020). Geometric insights into the convergence of non-linear TD learning. In *International Conference on Learning Representations*.
- Dayan, P. (1992). The convergence of TD(λ) for general lambda. *Machine Learning*, 8, 341–362.
- Duan, Y., Chen, X., Houthoofd, R., Schulman, J., & Abbeel, P. (2016). Benchmarking deep reinforcement learning for continuous control. In *International Conference on Machine Learning*, (pp. 1329–1338).
- Elfwing, S., Uchibe, E., & Doya, K. (2018). Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural Networks*, 107, 3–11.
- Hessel, M., Modayil, J., van Hasselt, H., Schaul, T., Ostrovski, G., Dabney, W., Horgan, D., Piot, B., Azar, M., & Silver, D. (2018). Rainbow: Combining improvements in deep reinforcement learning. *AAAI*.
- Horgan, D., Quan, J., Budden, D., Barth-Maroon, G., Hessel, M., van Hasselt, H., & Silver, D. (2018). Distributed prioritized experience replay. In *International Conference on Learning Representations*.
- Imani, E., Graves, E., & White, M. (2018). An off-policy policy gradient theorem using emphatic weightings. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems 31* (pp. 96–106). Curran Associates, Inc.
- Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1995). Planning and acting in partially observable stochastic domains. Unpublished report.
- Kapturowski, S., Ostrovski, G., Quan, J., Munos, R., & Dabney, W. (2018). Recurrent experience replay in distributed reinforcement learning. In *International conference on learning representations*.
- Kingma, D. P. & 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 E. Feigenbaum & J. Feldman (Eds.), *Computers and Thought* (pp. 406–450). McGraw-Hill, New York.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., & 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., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533.
- Moore, A. W. & 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](https://arxiv.org/abs/1805.00869).
- Peng, J. (1993). *Efficient dynamic programming-based learning for control*. PhD thesis, Northeastern University.
- Peng, J. & Williams, R. J. (1996). Incremental multi-step Q-learning. *Machine Learning*, 22, 283–290.

- Pohlen, T., Piot, B., Hester, T., Azar, M. G., Horgan, D., Budden, D., Barth-Maron, G., van Hasselt, H., Quan, J., Večerík, M., Hessel, M., Munos, R., & Pietquin, O. (2018). Observe and look further: Achieving consistent performance on Atari. *arXiv preprint arXiv:1805.11593*.
- Prokhorov, D. V. & 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., & Torgo, L. (Eds.), *Proceedings of the 16th European Conference on Machine Learning (ECML'05)*, (pp. 317–328). Springer.
- Schaul, T., Quan, J., Antonoglou, I., & 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. & 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*. PhD 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*, (pp. 216–224).
- Sutton, R. S. & Barto, A. G. (2018). *Reinforcement Learning: An Introduction*. The MIT press, Cambridge MA.
- Sutton, R. S., Mahmood, A. R., & 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., & 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., & Touretzky, D. S. (Eds.), *Advances in Neural Information Processing Systems 4*, (pp. 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. & 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., Guez, A., Hessel, M., Mnih, V., & 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*, (pp. 4287–4295).
- van Hasselt, H., Guez, A., & Silver, D. (2016). Deep reinforcement learning with double Q-learning. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, (pp. 2094–2100).

- van Hasselt, H., Hessel, M., & Aslanides, J. (2019). When to use parametric models in reinforcement learning? *CoRR*, *abs/1906.05243*.
- van Hasselt, H., Mahmood, A. R., & Sutton, R. S. (2014). Off-policy TD(λ) with a true online equivalence. In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence*, (pp. 330–339).
- van Hasselt, H., Quan, J., Hessel, M., Xu, Z., Borsa, D., & Barreto, A. (2019). General non-linear Bellman equations. *arXiv preprint arXiv:1907.03687*.
- van Hasselt, H. & Sutton, R. S. (2015). Learning to predict independent of span. *CoRR*, *abs/1508.04582*.
- van Hasselt, H. P. (2012). Reinforcement learning in continuous state and action spaces. In M. A. Wiering & M. van Otterlo (Eds.), *Reinforcement Learning: State of the Art*, volume 12 of *Adaptation, Learning, and Optimization* (pp. 207–251). Springer.
- van Seijen, H. & Sutton, R. S. (2013). Planning by prioritized sweeping with small backups. In *International Conference on Machine Learning*, (pp. 361–369).
- van Seijen, H. & Sutton, R. S. (2014). True online TD(λ). In *International Conference on Machine Learning*, (pp. 692–700).
- Wang, Z., de Freitas, N., Schaul, T., Hessel, M., van Hasselt, H., & 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. & 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.

A Appendix

A.1 Proof of Proposition 2

Proposition 2. *Let \mathbf{e}_t an instantaneous trace vector. Then let $\mathbf{z}_t(s)$ be the empirical mean $\mathbf{z}_t(s) = \frac{1}{n_t(s)} \sum_i^{n_t(s)} \mathbf{e}_{t_i^s}$, where t_i^s -s denote 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 $\mathbf{w}_{t+1} = \mathbf{w}_t + \alpha_t \delta_t \mathbf{z}_t$. If S_t is Markov, the expectation of this update is equal to the expected update with instantaneous traces \mathbf{e}_t , while attaining a potentially lower variance:*

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

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

Proof. We have

$$\begin{aligned} \mathbb{E}[\alpha_t \delta_t \mathbf{e}_t | S_t = s] &= \mathbb{E}[\alpha_t \delta_t | S_t = s] \mathbb{E}[\mathbf{e}_t | S_t = s] && \text{(as } s \text{ is Markov)} \\ &= \mathbb{E}[\alpha_t \delta_t | S_t = s] \mathbb{E}[\mathbf{z}_t | S_t = s] && \text{(as } \mathbf{z}_t = \frac{1}{n} \sum_i^n \mathbf{e}_{t_i^s}) \\ &= \mathbb{E}[\alpha_t \delta_t \mathbf{z}_t | S_t = s]. && \square \end{aligned}$$

Now let us look at the conditional variance for each of the dimension of the update vector $\alpha_t \delta_t \mathbf{z}_t$: $\mathbb{V}[\alpha_t \delta_t \mathbf{z}_{t,i} | S_t = s]$, where $\mathbf{z}_{t,i}$ denotes the i -th component of vector \mathbf{z}_t .

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

By a similar argument, we have

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

Now, we also know that $\mathbb{E}[\mathbf{z}_t | S_t = s] = \mathbb{E}[\mathbf{e}_t | S_t = s] = \mu_t$, as \mathbf{z}_t is the empirical mean of \mathbf{e}_t . Thus we also have, component-wise,

$$\mathbb{E}[\mathbf{z}_{t,i} | S_t = s] = \mathbb{E}[\mathbf{e}_{t,i} | S_t = s] = \mu_{t,i}$$

Moreover, from the same reason we have that $\mathbb{V}(\mathbf{z}_{t,i} | S_t = s) = \frac{1}{n_s} \mathbb{V}(\mathbf{e}_{t,i} | S_t = s)$. Thus we obtain:

$$\mathbb{V}[\alpha_t \delta_t \mathbf{z}_{t,i} | S_t = s] = \mathbb{E}[\alpha_t^2 \delta_t^2 | S_t = s] \mathbb{E}[\mathbf{z}_{t,i} (\mathbf{z}_{t,i})^T | S_t = s] - \mathbb{E}[\alpha_t \delta_t | S_t = s]^2 \mu_{t,i}^2$$

Thus:

$$\begin{aligned} &\mathbb{V}[\alpha_t \delta_t \mathbf{z}_{t,i} | S_t = s] - \mathbb{V}[\alpha_t \delta_t \mathbf{e}_{t,i} | S_t = s] \\ &= \mathbb{E}[\alpha_t^2 \delta_t^2 | S_t = s] \underbrace{(\mathbb{E}[\mathbf{z}_{t,i} (\mathbf{z}_{t,i})^T | S_t = s] - \mathbb{E}[\mathbf{e}_{t,i} (\mathbf{e}_{t,i})^T | S_t = s])}_{\leq 0, \text{ from definition of } \mathbf{z}_{t,i}} \leq 0, \end{aligned}$$

with equality holding, if and only if:

- i $\mathbb{E}[(\mathbf{z}_{t,i})^2 | S_t = s] = \mathbb{E}[(\mathbf{e}_{t,i})^2 | S_t = s] \Rightarrow \mathbb{V}(\mathbf{z}_{t,i} | S_t = s) = \mathbb{V}(\mathbf{e}_{t,i} | S_t = s)$, but $\mathbb{V}(\mathbf{z}_{t,i} | S_t = s) = \frac{1}{n_s} \mathbb{V}(\mathbf{e}_{t,i} | S_t = s)$ by definition of $\mathbf{z}_{t,i}$ as the running mean on samples $\mathbf{e}_{t,i}$. This can only happen for $n_s = 1$, or in the absence of stochasticity, for every state s . Thus, in the most general case, this implies $\mathbb{V}(\mathbf{z}_{t,i} | S_t = s) = \mathbb{V}(\mathbf{e}_{t,i} | S_t = s) = 0$; or
- ii $\mathbb{E}[\alpha_t^2 \delta_t^2 | S_t = s] = 0 \Rightarrow \delta_t = 0$

Thus, we have equality only with we have exactly one sample for the average \mathbf{z}_t so far, or only one sample is needed (thus \mathbf{z}_t and \mathbf{e}_t are not actual random variables and there is only one deterministic path to s); or when the TD errors are zero for all transitions following s . \square

A.2 Properties of mixture traces

In this section we explore and proof some of the properties of the proposed mixture trace, defined in Equation (3) in the main text and repeated here:

$$\mathbf{y}_t = (1 - \eta)\mathbf{z}_\theta(S_t) + \eta(\gamma_t\lambda_t\mathbf{y}_{t-1} + \nabla_{\mathbf{w}}v_{\mathbf{w}}(S_t)). \quad (3)$$

The proofs, in this section we will use the notation \mathbf{x}_t to denote the features used in a linear approximation for the value function(s) constructed. Just note that this term can be substituted, in general, by the gradient term $\nabla_{\mathbf{w}}v_{\mathbf{w}}(S_t)$ in the equation above.

Proposition 3. *The mixture trace \mathbf{y}_t defined in (3) can be interpreted as a trace $\mathbf{y}_t = \mu\mathbf{y}_{t-1} + \mathbf{x}_t$ with decay parameter $\mu = \lambda \cdot \eta$ and signal $\mathbf{x}_t = (1 - \eta)\mathbf{z}_\theta(S_t) + \eta \nabla_{\mathbf{w}}v_{\mathbf{w}}(S_t)$, such that*

$$\mathbf{y}_t = \sum_{k=0}^t (\eta\gamma\lambda)^k [(1 - \eta)\mathbf{z}_\theta(S_{t-k}) + \eta \nabla_{\mathbf{w}}v_{\mathbf{w}}(S_{t-k})]. \quad (5)$$

Proof. As mentioned before, under a linear parameterization $\nabla_{\mathbf{w}}v_{\mathbf{w}}(S_t) = \mathbf{x}(S_t) := \mathbf{x}_t$ Let us start with the definition of the mixture trace \mathbf{y}_t :

$$\begin{aligned} \mathbf{y}_t &= (1 - \eta)\mathbf{z}_t + \eta(\gamma_t\lambda_t\mathbf{y}_{t-1} + \mathbf{x}_t) \\ &= [(1 - \eta)\mathbf{z}_t + \eta\mathbf{x}_t] + \eta\gamma_t\lambda_t\mathbf{y}_{t-1} \\ &= [(1 - \eta)\mathbf{z}_t + \eta\mathbf{x}_t] + \eta\gamma_t\lambda_t [(1 - \eta)\mathbf{z}_{t-1} + \eta\mathbf{x}_{t-1}] + \eta^2\gamma_t\lambda_t\gamma_{t-1}\lambda_{t-1}\mathbf{y}_{t-2} \\ &= (1 - \eta) [\mathbf{z}_t + \eta\gamma_t\lambda_t\mathbf{z}_{t-1} + \eta^2\gamma_t\lambda_t\gamma_{t-1}\lambda_{t-1}\mathbf{z}_{t-2} + \dots] + \\ &\quad + \eta [\mathbf{x}_t + \eta\gamma_t\lambda_t\mathbf{x}_{t-1} + \eta^2\gamma_t\lambda_t\gamma_{t-1}\lambda_{t-1}\mathbf{x}_{t-2} + \dots] \\ &= (1 - \eta) \sum_{k=0}^t (\eta\gamma\lambda)^k \mathbf{z}_{t-k} + \eta \sum_{k=0}^t (\eta\gamma\lambda)^k \mathbf{x}_{t-k} \\ &= \sum_{k=0}^t (\eta\gamma\lambda)^k [(1 - \eta)\mathbf{z}_{t-k} + \eta\mathbf{x}_{t-k}] \end{aligned}$$

Substituting \mathbf{x}_t in the above derivation by $\nabla_{\mathbf{w}}v_{\mathbf{w}}(S_t)$ leads to (3). \square

Moreover, it is worth noting that the above equality recovers, for the extreme values of η :

- $\eta = 1 \Rightarrow \mathbf{y}_t = \sum_{k=0}^t (\gamma\lambda)^k \mathbf{x}_{t-k}$ (instantaneous trace for TD(λ))
- $\eta = 0 \Rightarrow \mathbf{y}_t = \sum_{k=0}^t (\eta\gamma\lambda)^k \mathbf{z}_{t-k} = \mathbf{z}_t$ (expected trace for TD(λ))

Moreover, as the extreme values already suggest, the expected update of the mixture traces follows the TD(λ) learning, in expectation, for all the intermediate values $\eta \in (0, 1)$ as well, trading off variance of estimates as η approaches 0.

Proposition 5. *Let \mathbf{e}_t^λ be a λ trace vector. Let $\mathbf{y}_t = (1 - \eta)\mathbf{z}_t + \eta(\gamma\lambda\mathbf{y}_{t-1} + \mathbf{x}_t)$ (as defined in (3)). Consider the ET(λ, η) algorithm $\mathbf{w}_{t+1} = \mathbf{w}_t + \alpha_t\delta_t\mathbf{y}_t$. For all Markov states s the expectation of this update is equal to the expected update with instantaneous traces \mathbf{e}_t^λ :*

$$\mathbb{E} [\alpha_t\delta_t\mathbf{y}(S_t)|S_t = s] = \mathbb{E} [\alpha_t\delta_t\mathbf{e}_t^\lambda|S_t = s],$$

for every $\eta \in [0, 1]$ and any $\lambda \in [0, 1]$.

Proof. Let us revisit Eq. 5 in Proposition 3:

$$\begin{aligned}
\mathbb{E}[\mathbf{y}_t] &= \mathbb{E} \left[\sum_{k=0}^t (\eta\gamma\lambda)^k [(1-\eta)\mathbf{z}_{t-k} + \eta\mathbf{x}_{t-k}] \right] \\
&= \mathbb{E} \left[\sum_{k=0}^t (\eta\gamma\lambda)^k [(1-\eta)\mathbb{E}[(\mathbf{x}_{t-k} + \gamma\lambda\mathbf{z}_{t-k-1})] + \eta\mathbf{x}_{t-k}] \right] \\
&= \mathbb{E} \left[\sum_{k=0}^t (\eta\gamma\lambda)^k \mathbf{x}_{t-k} + (1-\eta)\gamma\lambda \sum_{k=0}^{t-1} (\eta\gamma\lambda)^k \underbrace{\mathbf{z}_{t-k-1}}_{\mathbb{E}[(\mathbf{x}_{t-k-1} + \gamma\lambda\mathbf{z}_{t-k-2})]} \right] \\
&= \mathbb{E} \left[\sum_{k=0}^t (\eta\gamma\lambda)^k \mathbf{x}_{t-k} + (1-\eta)\gamma\lambda \sum_{k=0}^{t-1} (\eta\gamma\lambda)^k \mathbf{x}_{t-k-1} + (1-\eta)(\gamma\lambda)^2 \sum_{k=0}^{t-2} (\eta\gamma\lambda)^k \underbrace{\mathbf{z}_{t-k-2}}_{\mathbb{E}[\mathbf{x}_{t-k-2} + \gamma\lambda\mathbf{z}_{t-k-3}]} \right] \\
&= \mathbb{E} \left[\sum_{k=0}^t (\eta\gamma\lambda)^k \mathbf{x}_{t-k} + (1-\eta) \sum_{i=1}^{t-1} (\gamma\lambda)^i \sum_{k=0}^{t-i} (\eta\gamma\lambda)^k \mathbf{x}_{t-k-i} \right]
\end{aligned}$$

Now, re-writing the sum, gathering all the weighting for each feature \mathbf{x}_{t-k-i} we get:

$$\begin{aligned}
\mathbb{E}[\mathbf{y}_t] &= \mathbb{E} \left[\sum_{k=0}^t (\eta\gamma\lambda)^k \mathbf{x}_{t-k} + (1-\eta) \sum_{i=1}^{t-1} (\gamma\lambda)^i \sum_{k=0}^{t-i} (\eta\gamma\lambda)^k \mathbf{x}_{t-k-i} \right] \\
&= \mathbb{E} \left[\mathbf{x}_t + \sum_{k=1}^t \mathbf{x}_{t-k} \left((\eta\gamma\lambda)^k + (1-\eta) \sum_{i=1}^k (\gamma\lambda)^i \cdot (\gamma\lambda\eta)^{k-i} \right) \right] \\
&= \mathbb{E} \left[\mathbf{x}_t + \sum_{k=1}^t \mathbf{x}_{t-k} (\gamma\lambda)^k \left(\eta^k + (1-\eta) \sum_{i=1}^k \eta^{k-i} \right) \right] \\
&= \mathbb{E} \left[\mathbf{x}_t + \sum_{k=1}^t \mathbf{x}_{t-k} (\gamma\lambda)^k \left(\eta^k + (1-\eta) \frac{1-\eta^k}{(1-\eta)} \right) \right] \\
&= \mathbb{E} \left[\mathbf{x}_t + \sum_{k=1}^t \mathbf{x}_{t-k} (\gamma\lambda)^k \right] \\
&= \mathbb{E} \left[\sum_{k=0}^t (\gamma\lambda)^k \mathbf{x}_{t-k} \right]
\end{aligned}$$

Thus $\mathbb{E}[\mathbf{y}_t] = \mathbb{E}[\sum_{k=0}^t (\gamma\lambda)^k \mathbf{x}_{t-k}] = \mathbb{E}[\mathbf{e}_t^\lambda]$, where \mathbf{e}_t^λ is the instantaneous λ trace on feature space \mathbf{x} . Thus $\mathbb{E}[\mathbf{y}(s)] = z_*^\lambda(s) = \mathbb{E}[\mathbf{e}_t^\lambda]$. Finally we can plug-in this result in the expected update:

$$\begin{aligned}
\mathbb{E}[\alpha_t \delta_t \mathbf{y}(S_t) | S_t = s] &= \mathbb{E}[\alpha_t \delta_t | S_t = s] \mathbb{E}[\mathbf{y}(S_t) | S_t = s] \\
&= \mathbb{E}[\alpha_t \delta_t | S_t = s] z_*^\lambda(s) \\
&= \mathbb{E}[\alpha_t \delta_t | S_t = s] \mathbb{E}[\mathbf{e}_t^\lambda | S_t = s] \\
&= \mathbb{E}[\alpha_t \delta_t \mathbf{e}_t^\lambda | S_t = s].
\end{aligned}$$

□

Finally, please note that in this proposition and its proof we drop the time indices t for λ and γ parameters in the definition of \mathbf{y}_t . This is purely to ease the notation and promote compactness in the derivation.

A.3 Experiment details

For the deep reinforcement learning experiments discussed in Section 4.3 we compare to an implementation of online Q(λ). We first describe this algorithm, and then describe the expected-trace variant. All experiments for Section 4.3 were run with the ALE (Bellemare et al., 2013), exactly as

described in Mnih et al. (2015), including using action repeats (4x), downsampling (to 84×84), and frame stacking. These experiments were conducted using Jax (Bradbury et al., 2018).

In all cases, we used ϵ -greedy exploration (cf. Sutton & Barto, 2018), with an ϵ that quickly decayed from 1 to 0.01 according to $\epsilon_0 = 1$ and $\epsilon_t = \epsilon_{t-1} + 0.01(0.01 - \epsilon_{t-1})$. Unlike Mnih et al. (2015), we did not clip rewards, and we also did not apply any target normalisation (cf. van Hasselt et al., 2016) or non-linear value transformations (Pohlen et al., 2018; van Hasselt et al., 2019). We conjecture that such extensions could be beneficial for performance, but they are orthogonal to the main research questions investigated here and are therefore left for future work.

Algorithm 2 $Q(\lambda)$

```

1: initialise  $\mathbf{w}$ 
2: initialise  $e = \mathbf{0}$ 
3: observe initial state  $S$ 
4: pick action  $A \sim \pi(q_{\mathbf{w}}(S))$ 
5:  $v \leftarrow \max_a q_{\mathbf{w}}(S, a)$ 
6:  $\gamma = 0$ 
7: repeat
8:   take action  $A$ , observe  $R, \gamma'$  and  $S'$            #  $\gamma' = 0$  on a terminating transition
9:    $v' \leftarrow \max_a q_{\mathbf{w}}(S', a)$ 
10:   $\delta \leftarrow R + \gamma v' - v$ 
11:   $e \leftarrow \gamma \lambda e + \nabla_{\mathbf{w}} q_{\mathbf{w}}(S, A)$ 
12:   $\Delta \mathbf{w} \leftarrow \delta e + (v - q_{\mathbf{w}}(S, A)) \nabla_{\mathbf{w}} q_{\mathbf{w}}(S, A)$ 
13:   $\Delta \mathbf{w} \leftarrow \text{transform}(\Delta \mathbf{w})$            # e.g., ADAM-ify
14:   $\mathbf{w} \leftarrow \mathbf{w} + \Delta \mathbf{w}$ 
15: until done

```

A.3.1 Deep $Q(\lambda)$

We assume the typical setting (e.g., Mnih et al., 2015) where we have a neural network $q_{\mathbf{w}}$ that outputs $|A|$ numbers, such that $q(s, a) = q_{\mathbf{w}}(s)[a]$. That is, we forward the observation s through network $q_{\mathbf{w}}$ with weights \mathbf{w} and $|A|$ outputs, and then select the a^{th} output to represent the value of taking action a .

Algorithm 2 then works as follows. For each transition, we first compute a telescoping TD error $\delta = r + \gamma' v' - v$ (line 5), where $\gamma' = 0$ on termination (and then S' is the first observation of the next episode) and, in our experiments, $\gamma' = 0.995$ otherwise. We update the trace e as usual (line 11), using accumulating traces. Note that the weights and, hence, trace will also have elements corresponding to the weights of actions that were not selected. The gradient with respect to those elements is considered to be zero, as is conventional.

Then, we compute a weight update $\Delta \mathbf{w} = \delta e + (v - q_{\mathbf{w}}(S, A)) \nabla_{\mathbf{w}} q_{\mathbf{w}}(S, A)$. The additional term corrects for the fact that our TD error is a telescoping error, and does not have the usual ‘ $-q(s, a)$ ’ term. This is akin to the $Q(\lambda)$ algorithm proposed by Peng & Williams (1996).

Finally, we transform the resulting update, using a transformation exactly like ADAM (Kingma & Adam, 2015), but applied to the update $\Delta \mathbf{w}$ rather than a gradient. The hyper-parameters were $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 0.0001$, and one of the step sizes as given in Figure 5. We then apply the resulting transformed update by adding it to the weights (line 14).

For our experiments in Section 4.3 we used the same architecture as Mnih et al. (2015), but with 128 channels in each convolutional layer, because we ran experiments on TPUs (a single core per experiment) which are most efficient when using tensors where one dimension is a multiple of 128.

A.3.2 Deep QET(λ)

We now describe the expected-trace algorithm, shown in Algorithm 3, which was used for the results in Section 4.3. It is very similar to the $Q(\lambda)$ algorithm described above, and in fact equivalent when we set $\eta = 1$.

The first main change is that we will split the computation of $q(s, a)$ into two separate parts, such that $q_{(\mathbf{w}, \boldsymbol{\xi})}(s, a) = \mathbf{w}_a^\top \mathbf{x}_\xi(s)$. This is equivalent to the previous algorithm: we have just labeled separate subsets of parameters as $(\mathbf{w}, \boldsymbol{\xi})$ rather than merging all of them into a single vector \mathbf{w} , and we have labeled the last hidden layer as $\mathbf{x}(s)$. We keep separate traces for these subset (lines 11 and 12), but this is equivalent to keeping one big trace for the combined set.

Algorithm 3 QET(λ)

```

1: initialise  $\mathbf{w}, \boldsymbol{\xi}, \boldsymbol{\theta}$ 
2: initialise  $\mathbf{e} = \mathbf{0}, \mathbf{y} = \mathbf{0}$ 
3: observe initial state  $S$ 
4: pick action  $A \sim \pi(q(S))$ 
5:  $v \leftarrow \max_a q_{(\mathbf{w}, \boldsymbol{\xi})}(S', a)$            #  $q_{(\mathbf{w}, \boldsymbol{\xi})}(s, a) = \mathbf{w}_a^\top \mathbf{x}_\xi(s)$ , where  $\mathbf{w} = (\mathbf{w}_1, \dots, \mathbf{w}_{|A|})$ 
6:  $\gamma = 0$ 
7: repeat
8:   take action  $A$ , observe  $R, \gamma'$  and  $S'$            #  $\gamma' = 0$  on any terminating transition
9:    $v' \leftarrow \max_a q_{(\mathbf{w}, \boldsymbol{\xi})}(S', a)$ 
10:   $\delta \leftarrow R + \gamma v' - v$ 
11:   $\mathbf{e}^{\mathbf{w}} \leftarrow \gamma \lambda \mathbf{y} + \nabla_{\mathbf{w}} q_{(\mathbf{w}, \boldsymbol{\xi})}(S, A)$ 
12:   $\mathbf{e}^{\boldsymbol{\xi}} \leftarrow \gamma \lambda \mathbf{e}^{\boldsymbol{\xi}} + \nabla_{\boldsymbol{\xi}} q_{(\mathbf{w}, \boldsymbol{\xi})}(S, A)$ 
13:   $\Delta \mathbf{w} \leftarrow \delta \mathbf{e}^{\mathbf{w}} + (v - q_{(\mathbf{w}, \boldsymbol{\xi})}(S, A)) \nabla_{\mathbf{w}} q_{(\mathbf{w}, \boldsymbol{\xi})}(S, A)$ 
14:   $\Delta \boldsymbol{\xi} \leftarrow \delta \mathbf{e}^{\boldsymbol{\xi}} + (v - q_{(\mathbf{w}, \boldsymbol{\xi})}(S, A)) \nabla_{\boldsymbol{\xi}} q_{(\mathbf{w}, \boldsymbol{\xi})}(S, A)$ 
15:   $\Delta \boldsymbol{\theta} \leftarrow \nabla_{\boldsymbol{\theta}} \|\mathbf{e}^{\boldsymbol{\xi}} - \mathbf{z}_{\boldsymbol{\theta}}(S, A)\|_2^2$ 
16:   $\Delta \mathbf{w} \leftarrow \text{transform}(\Delta \mathbf{w})$            # e.g., ADAM-ify
17:   $\Delta \boldsymbol{\xi} \leftarrow \text{transform}(\Delta \boldsymbol{\xi})$ 
18:   $\Delta \boldsymbol{\theta} \leftarrow \text{transform}(\Delta \boldsymbol{\theta})$ 
19:   $\mathbf{w} \leftarrow \mathbf{w} + \Delta \mathbf{w}$ 
20:   $\boldsymbol{\xi} \leftarrow \boldsymbol{\xi} + \Delta \boldsymbol{\xi}$ 
21:   $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} + \Delta \boldsymbol{\theta}$ 
22:   $\mathbf{y} = (1 - \eta) \mathbf{z}_{\boldsymbol{\theta}}(s, a) + \eta \mathbf{e}^{\mathbf{w}}$ 
23: until done

```

This split in parameters helps avoid learning an expected trace for the full trace, which has millions of elements. Instead, we only learn expectations for traces corresponding to the last layer, denoted $\mathbf{e}^{\mathbf{w}}$. Importantly, the function $\mathbf{z}_{\boldsymbol{\theta}}(s, a)$ should condition on both state and action. This was implemented as a tensor $\boldsymbol{\theta} \in \mathbb{R}^{|A| \times |A| \times |x|}$, such that its tensor multiplication with the features $\mathbf{x}(s)$ yields a $|A| \times |A|$ matrix \mathbf{Z} . Then, we interpret the vector $\mathbf{z}_a = [\mathbf{Z}]_a$ as the approximation to the expected trace $\mathbb{E}[\mathbf{e}_t | S_t = s, A_t = a]$, and update it accordingly, using a squared loss (and, again, ADAM-ifying the update before applying it to the parameters). The step size for the expected trace update was always $\beta = 0.1$ in our experiments, and the expected trace loss was not back-propagated into the feature representation. This can be done, but we leave any investigation of this for future work, as it would present a conflating factor for our experiments, because the expected trace update would then serve as an additional learning signal for the features that are also used for the value approximations.