

Multi-Step Reinforcement Learning: A Unifying Algorithm

Kristopher De Asis,¹ J. Fernando Hernandez-Garcia,¹
G. Zacharias Holland,¹ Richard S. Sutton

Reinforcement Learning and Artificial Intelligence Laboratory, University of Alberta
{kldeasis,jfherman,gholland,rsutton}@ualberta.ca

Abstract

Unifying seemingly disparate algorithmic ideas to produce better performing algorithms has been a longstanding goal in reinforcement learning. As a primary example, $TD(\lambda)$ elegantly unifies one-step TD prediction with Monte Carlo methods through the use of eligibility traces and the trace-decay parameter λ . Currently, there are a multitude of algorithms that can be used to perform TD control, including Sarsa, Q -learning, and Expected Sarsa. These methods are often studied in the one-step case, but they can be extended across multiple time steps to achieve better performance. Each of these algorithms is seemingly distinct, and no one dominates the others for all problems. In this paper, we study a new multi-step action-value algorithm called $Q(\sigma)$ that unifies and generalizes these existing algorithms, while subsuming them as special cases. A new parameter, σ , is introduced to allow the degree of sampling performed by the algorithm at each step during its backup to be continuously varied, with Sarsa existing at one extreme (full sampling), and Expected Sarsa existing at the other (pure expectation). $Q(\sigma)$ is generally applicable to both on- and off-policy learning, but in this work we focus on experiments in the on-policy case. Our results show that an intermediate value of σ , which results in a mixture of the existing algorithms, performs better than either extreme. The mixture can also be varied dynamically which can result in even greater performance.

The Landscape of TD Algorithms

Temporal-difference (TD) methods (Sutton and Barto 1998) are an important concept in *reinforcement learning* (RL) that combines ideas from Monte Carlo and dynamic programming methods. TD methods allow learning to occur directly from raw experience in the absence of a model of the environment's dynamics, like with Monte Carlo methods, while also allowing estimates to be updated based on other learned estimates without waiting for a final result, like with dynamic programming. The core concepts of TD methods provide a flexible framework for creating a variety of powerful algorithms that can be used for both prediction and control.

There are a number of TD control methods that have been proposed. Q -learning (Watkins 1989; Watkins and Dayan

1992) is arguably the most popular, and is considered an *off-policy* method because the policy generating the behaviour (the *behaviour policy*), and the policy that is being learned (the *target policy*) are different. Sarsa (Rummery and Niranjan 1994; Sutton 1996) is the classical *on-policy* control method, where the behaviour and target policies are the same. However, Sarsa can be extended to learn off-policy with the use of importance sampling (Precup, Sutton, and Singh 2000). Expected Sarsa is an extension of Sarsa that, instead of using the action-value of the next state to update the value of the current state, uses the expectation of all the subsequent action-values of the current state with respect to the target policy. Expected Sarsa has been studied as a strictly on-policy method (van Seijen et al. 2009), but in this paper we present a more general version that can be used for both on- and off-policy learning and that also subsumes Q -learning. All of these methods are often described in the simple one-step case, but they can also be extended across multiple time steps.

The $TD(\lambda)$ algorithm unifies one-step TD learning with Monte Carlo methods (Sutton 1988). Through the use of eligibility traces, and the trace-decay parameter, $\lambda \in [0, 1]$, a spectrum of algorithms is created. At one end, $\lambda = 1$, exists Monte Carlo methods, and at the other, $\lambda = 0$, exists one-step TD learning. In the middle of the spectrum are intermediate methods which can perform better than the methods at either extreme (Sutton and Barto 1998). The concept of eligibility traces can also be applied to TD control methods such as Sarsa and Q -learning, which can create more efficient learning and produce better performance (Rummery 1995).

Multi-step TD methods are usually thought of in terms of an average of many multi-step returns of differing lengths and are often associated with eligibility traces, as is the case with $TD(\lambda)$. However, it is also natural to think of them in terms of individual n -step returns with their associated n -step backup (Sutton and Barto 1998). We refer to each of these individual backups as *atomic backups*, whereas the combination of several atomic backups of different lengths creates a *compound backup*.

In the existing literature, it is not clear how best to extend one-step Expected Sarsa to a multi-step algorithm. The *Tree-backup* algorithm was originally presented as a method to perform off-policy evaluation when the behaviour policy

is non-Markov, non-stationary or completely unknown (Precup, Sutton, and Singh 2000). In this paper, we re-present Tree-backup as a natural multi-step extension of Expected Sarsa. Instead of performing the updates with entirely sampled transitions as with multi-step Sarsa, Tree-backup performs the update using the expected values of all the actions at each transition.

$Q(\sigma)$ is an algorithm that was first proposed by Sutton and Barto (2018) which unifies and generalizes the existing multi-step TD control methods. The degree of sampling performed by the algorithm is controlled by the sampling parameter, σ . At one extreme ($\sigma = 1$) exists Sarsa (full sampling), and at the other ($\sigma = 0$) exists Tree-backup (pure expectation). Intermediate values of σ create algorithms with a mixture of sampling and expectation, and σ can be interpreted as a way to control the bias-variance trade-off inherent in multi-step TD algorithms.

In this work, on problems with a tabular representation and a problem requiring function approximation, we show that an intermediate value of σ can outperform the algorithms that exist at either extreme. In addition, we show that σ can be varied dynamically to produce even greater performance. We limit our discussion of $Q(\sigma)$ to the atomic multi-step case without eligibility traces, but a natural extension is to make use of compound backups and is an avenue for future research. Furthermore, $Q(\sigma)$ is generally applicable to both on- and off-policy learning, but for our initial empirical study we examined only on-policy prediction and control problems.

MDPs and One-step Solution Methods

The sequential decision problem encountered in RL is often modeled as a *Markov decision process* (MDP). Under this framework, an *agent* and the environment interact over a sequence of discrete time steps t . At every time step, the agent receives information about the environment's *state*, $S_t \in \mathcal{S}$, where \mathcal{S} is the set of all possible states. The agent uses this information to select an *action*, A_t , from the set of all possible actions \mathcal{A} . Based on the behavior of the agent and the state of the environment, the agent receives a *reward*, $R_{t+1} \in \mathbb{R}$, and moves to another state, $S_{t+1} \in \mathcal{S}$, with a *state-transition probability* $p(s'|s, a) = P(S_{t+1} = s' | S_t = s, A_t = a)$, for $a \in \mathcal{A}$ and $s, s' \in \mathcal{S}$.

The agent behaves according to a *policy* $\pi(a|s)$, which is a probability distribution over the set $\mathcal{S} \times \mathcal{A}$. Through the process of policy iteration (Sutton and Barto 1998), the agent learns the optimal policy, π^* , that maximizes the expected discounted return:

$$G_t = R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + \dots = \sum_{k=0}^{T-t-1} \gamma^k R_{t+1+k}, \quad (1)$$

for a discount factor $\gamma \in [0, 1)$ and $T = \infty$ for continuing tasks, or $\gamma \in [0, 1]$ and T equal to the final time step in episodic tasks.

TD algorithms strive to maximize the expected return by computing value functions that estimate the expected future rewards in terms of the elements of the environment and the

actions of the agent. The *state-value function* is the expected return when the agent is in a state s and follows policy π , defined as $v_\pi(s) = \mathbb{E}_\pi[G_t | S_t = s]$. For control, most of the time we focus on estimating the *action-value function*, which is the expected return when the agent takes an action a , in a state s , while following a policy π , and is defined as:

$$q_\pi(s, a) = \mathbb{E}_\pi[G_t | S_t = s, A_t = a]. \quad (2)$$

Equation 2 can be estimated iteratively by observing new rewards, bootstrapping on old estimates of q_π , and using the update rule:

$$Q(S_t, A_t) \leftarrow Q(S_t, A_t) + \alpha [R_{t+1} + \gamma Q(S_{t+1}, A_{t+1}) - Q(S_t, A_t)], \quad (3)$$

where $\alpha \in (0, 1]$ is the step size parameter. Update rules are also known as *backup* operations because they transfer information back from future states to the current one. A common way to visualize backup operations is by using backup diagrams such as the ones depicted in Figure 1.

For clarity, the algorithmic ideas in this paper are presented initially as *tabular solution methods*, but we also extend them to use function approximation, and thus they also serve as *approximate solution methods*.

The term in brackets in (3):

$$\delta_t^S = R_{t+1} + \gamma Q(S_{t+1}, A_{t+1}) - Q(S_t, A_t), \quad (4)$$

is also known as the *TD error*, denoted δ_t . TD control methods are characterized by their TD error; for example, the TD error in (4) corresponds to the classic on-policy method known as Sarsa.

Because learning requires a certain amount of exploration, behaving greedily with respect to the estimated optimal policy is often infeasible. Therefore, agents are often trained under ϵ -greedy policies for which the agent only chooses the optimal action with a probability $(1 - \epsilon)$ and behaves randomly with probability ϵ , for $\epsilon \in [0, 1]$. Nevertheless, learning the optimal policy is possible if it is done off-policy. When the agent is learning off-policy, it behaves according to a behavior policy, μ , while learning a target policy, π . This can be achieved by using another TD control method, Expected Sarsa. In contrast with Sarsa, Expected Sarsa behaves according to the behavior policy, but updates its estimate by taking an expectation of $Q(S_t, A_t)$ over the actions at time t , according to the target policy (van Seijen et al. 2009). For convenience, let the expected action-value be defined as:

$$V_{t+1} = \sum_a \pi(a|S_{t+1}) Q(S_{t+1}, a). \quad (5)$$

Then, the TD error of Expected Sarsa can be written as:

$$\delta_t^{ES} = R_{t+1} + \gamma V_{t+1} - Q(S_t, A_t). \quad (6)$$

A special case of Expected Sarsa is Q -learning, where the estimate is updated according to the maximum of $Q(S_t, a)$ over the actions (Watkins 1989):

$$\delta_t^{QL} = R_{t+1} + \gamma \max_a Q(S_{t+1}, a) - Q(S_t, A_t). \quad (7)$$

Q -learning is the resulting algorithm when the target policy of Expected Sarsa is the greedy policy with respect to Q .

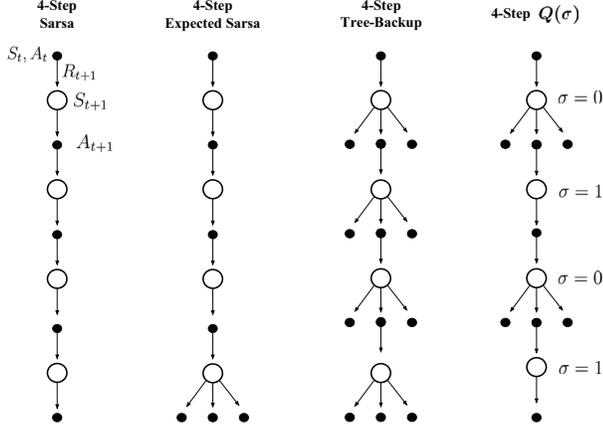


Figure 1: Backup diagrams for atomic 4-step Sarsa, Expected Sarsa, Tree-backup, and $Q(\sigma)$. Here we can see that $Q(\sigma)$ encompasses the other three algorithms based on the setting of σ .

Atomic Multi-Step Algorithms

The TD methods presented in the previous section can be generalized even further by bootstrapping over longer time intervals. This has been shown to decrease the bias of the update at the cost of increasing the variance (Jaakkola, Jordan, and Singh 1994). Nevertheless, in many cases it is possible to achieve better performance by choosing a value for the backup length parameter, n , greater than one (Sutton and Barto 1998). We refer to algorithms which make use of a multi-step atomic backup as *atomic multi-step algorithms*. Just like how one-step methods are defined by their TD error, each atomic multi-step algorithm is characterized by its n -step return. For atomic multi-step Sarsa, the n -step return is:

$$G_{t:t+n} = R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + \dots + \gamma^{n-1} R_{t+n} + \gamma^n Q_{t+n-1}(S_{t+n}, A_{t+n}),$$

$$= \sum_{k=0}^{n-1} \gamma^k R_{t+k+1} + \gamma^n Q_{t+n-1}(S_{t+n}, A_{t+n}), \quad (8)$$

where Q_{t+n-1} is the estimate of q_π at time $t+n-1$, and the subscript range, $t : t+n$, denotes the length of the backup. n -step Sarsa can be adapted for off-policy learning by introducing an importance sampling ratio term (Precup, Sutton, and Singh 2000):

$$\rho_t^{t+n} = \prod_{k=t}^{\tau} \frac{\pi(A_k|S_k)}{\mu(A_k|S_k)}, \quad (9)$$

and multiplying it with the TD error to get the following update rule:

$$Q_{t+n}(S_t, A_t) \leftarrow Q_{t+n-1}(S_t, A_t) + \alpha \rho_{t+1}^{t+n} [G_{t:t+n} - Q_{t+n-1}(S_t, A_t)], \quad (10)$$

where $\tau = \min(t+n-1, T-1)$ is the time step before the end of the update or before the end of the episode. In

the update, the action-values for all other states remain the same – i.e. $Q_{t+n}(s, a) = Q_{t+n-1}(s, a), \forall s \neq S_t$, and $a \neq A_t$. This update rule is not only applicable for off-policy n -step Sarsa, but is a generally useful form for other atomic multi-step algorithms as well. We present the algorithms in this work as general off-policy solution methods, but in the experiments section we evaluate them empirically on-policy only which provides useful insight into their behaviour. We defer the empirical study and comparison of the algorithms in an off-policy setting to future work.

Expected Sarsa can also be generalized to a multi-step method by using the return:

$$G_{t:t+n} = R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + \dots + \gamma^n V_{t+n}. \quad (11)$$

The first $n-1$ states and actions are sampled according to the behaviour policy, as with n -step Sarsa, but the last state is backed up according to the expected action-value under the target policy. To make n -step Expected Sarsa entirely off-policy, an importance sampling ratio term can also be introduced, but it needs to omit the last time step. The resulting update would be the same as in (10), but would use ρ_{t+1}^{t+n-1} and the n -step return for n -step Expected Sarsa from (11).

A drawback to using importance sampling to learn off-policy is that it can create high variance which must be compensated for by using small step sizes; this can slow learning (Precup, Sutton, and Singh 2000). In the next section we present a method that is also a generalization of Expected Sarsa, but that can learn off-policy without importance sampling.

Tree-backup

As shown in (11), the TD return of n -step Expected Sarsa is calculated by taking an expectation over the actions at the last step of the backup. However, it is possible to extend this idea to every time step of the backup by taking an expectation at every step (Precup, Sutton, and Singh 2000). The resulting algorithm is a multi-step generalization of Expected Sarsa that is known as Tree-backup because of its characteristic backup diagram (Figure 1). Moreover, just like Expected Sarsa and Q-learning, this proposed generalization does not require importance sampling to be applied off-policy. Hence, it could be argued that it is a more appropriate generalization of Expected Sarsa to multi-step learning (Sutton and Barto 2018). Because Expected Sarsa subsumes Q-learning, Tree-backup can also be thought of as a multi-step generalization of Q-learning if the target policy is greedy with respect to the action-value function.

Tree-backup has several advantages over n -step Expected Sarsa. Tree-backup has the capacity for learning off-policy without the need for importance sampling, reducing the variance due to the importance sampling ratios. Additionally, because an importance sampling ratio does not need to be computed, the behavior policy does not need to be stationary, Markov, or even known (Precup, Sutton, and Singh 2000).

Each branch of the tree represents an action, while the main branch represents the action taken at time t . The value of each of the branches is the value of $Q_{t+n}(S_t, A_t)$ for the

corresponding t , whereas the value of each segment of the main branch is the reward at the corresponding time step. The n -step return is the sum of the values of each branch weighted by the product of the probabilities of the actions leading to the branch and multiplied by the corresponding power of the discount term. For clarity, it is easier to present the n -step return of the Tree-backup algorithm in terms of the TD error of Expected Sarsa from (6):

$$G_{t:t+n} = Q_{t-1}(S_t, A_t) + \sum_{k=t}^{\tau} \delta_k^{ES} \prod_{i=t+1}^k \gamma \pi(A_i | S_i). \quad (12)$$

This atomic version of multi-step Tree-backup was first presented by Sutton and Barto (2018).

As a result of the product term in (12), in addition to the discount factor γ , future rewards are further discounted by the probabilities of the actions taken. The Tree-backup algorithm therefore assigns less weight to the reward sequence received, and compensates by bootstrapping off of the values of actions not taken. Due to this, Tree-backup is more biased than Sarsa in the multi-step case with a stochastic policy, as Sarsa gives full weight to every reward received prior to bootstrapping. However, this increase in bias (towards the estimates in the value function) is traded off with decreased variance in the reward sequence from taking expectations.

The $Q(\sigma)$ Algorithm

In the previous sections we have incrementally introduced several generalizations for the TD control methods Sarsa and Expected Sarsa, and in this section we present an algorithm that unifies them called $Q(\sigma)$.

Sarsa can be generalized to an atomic multi-step algorithm by using an n -step return, and n -step Sarsa generalizes to an off-policy algorithm through the use of importance sampling. In contrast, Expected Sarsa can learn off-policy without the need for importance sampling, and generalizes to the atomic multi-step algorithms: Tree-backup and n -step Expected Sarsa. All of the algorithms presented so far can be broadly categorized into two families: those that backup their actions as samples, like Sarsa; and those that consider an expectation over all actions in their backup, like Expected Sarsa and Tree-backup. In this section, we introduce a method to unify both families of algorithms by introducing a new parameter, σ . The possibility of unifying Sarsa and Tree-backup was first suggested by Precup et al. (2000), and the first formulation of $Q(\sigma)$ was presented by Sutton and Barto (2018).

The intuition behind $Q(\sigma)$ is based on the idea that we have a choice to update the estimate of q_π based on one action sampled from the set of possible future actions, or based on the expectation over the possible future actions. For example, with n -step Sarsa, a sample is taken at every step of the backup, whereas with the Tree-backup algorithm, an expectation is taken instead. However, the choice of sampling or expectation does not have to remain constant for every step of the backup. Furthermore, the backup at a time step t could be based on a weighted average of both sampling and expectation. In order to implement this, the parameter, $\sigma_t \in [0, 1]$, is introduced to control the degree of sampling

Algorithm 1 Off-policy n -step $Q(\sigma)$ for estimating q_π

Input: a behaviour policy μ and a target policy π
Initialize $S_0 \neq$ terminal; select A_0 according to $\mu(\cdot | S_0)$
Store S_0, A_0 , and $Q(S_0, A_0)$
for $t = 0, 1, 2, \dots, T + n - 1$ **do**
 if $t < T$ **then**
 Take action A_t ; observe R and S_{t+1}
 Store S_{t+1}
 if S_{t+1} is terminal **then**
 Store: $\delta_t^\sigma = R - Q(S_t, A_t)$
 else
 Select A_{t+1} according to $\mu(\cdot | S_{t+1})$ and Store
 Store: $Q(S_{t+1}, A_{t+1}), \sigma_{t+1}, \pi(A_{t+1} | S_{t+1})$
 Store: $\delta_t^\sigma = R + \gamma[\sigma_{t+1}Q(S_{t+1}, A_{t+1}) + (1 - \sigma_{t+1})V_{t+1}] - Q(S_t, A_t)$
 Store: $\rho_{t+1} = \frac{\pi(A_{t+1} | S_{t+1})}{\mu(S_{t+1} | A_{t+1})}$
 end if
 end if
 if $\tau \geq 0$ **then**
 $\tau \leftarrow t - n + 1$
 if $\tau \geq 0$ **then**
 $\rho \leftarrow 1$
 $E \leftarrow 1$
 $G \leftarrow Q(S_\tau, A_\tau)$
 for $k = \tau, \dots, \min(\tau + n - 1, T - 1)$ **do**
 $G \leftarrow G + E \delta_k^\sigma$
 $E \leftarrow \gamma E [(1 - \sigma_k)\pi(A_{k+1} | S_{k+1}) + \sigma_{k+1}]$
 $\rho \leftarrow \rho(1 - \sigma_k + \sigma_k \rho_k)$
 end for
 $Q(S_\tau, A_\tau) \leftarrow Q(S_\tau, A_\tau) + \alpha \rho [G - Q(S_\tau, A_\tau)]$
 end if
 end if
end for

at each step of the backup. Thus, the TD error of $Q(\sigma)$ can be represented in terms of a weighted sum of the TD errors of Sarsa and Expected Sarsa:

$$\begin{aligned} \delta_t^\sigma &= \sigma_{t+1} \delta_t^S + (1 - \sigma_{t+1}) \delta_t^{ES}, \\ &= R_{t+1} + \gamma[\sigma_{t+1} Q_t(S_{t+1}, A_{t+1}) + (1 - \sigma_{t+1}) V_{t+1}] \\ &\quad - Q_{t-1}(S_t, A_t). \end{aligned} \quad (13)$$

The n -step return is then:

$$\begin{aligned} G_{t:t+n} &= Q_{t-1}(S_t, A_t) \\ &\quad + \sum_{k=t}^{\tau} \delta_k^\sigma \prod_{i=t+1}^k \gamma [(1 - \sigma_i)\pi(A_i | S_i) + \sigma_i]. \end{aligned} \quad (14)$$

Moreover, the importance sampling ratio from (9) can be modified to include σ as follows:

$$\rho_{t+1}^{t+n} = \prod_{k=t+1}^{\tau} \left(\sigma_k \frac{\pi(A_k | S_k)}{\mu(A_k | S_k)} + 1 - \sigma_k \right). \quad (15)$$

The update rule for $Q(\sigma)$ can then be obtained by using $G_{t:t+n}$ from (14) and ρ_{t+1}^{t+n} from (15), with the update rule from (10). Algorithm 1 shows the pseudocode for the complete off-policy n -step $Q(\sigma)$ algorithm.

Additionally, a proof for one-step $Q(\sigma)$ is readily available by applying the results from Jakkola et al. (1994), Singh et al. (2000), and van Seijen et al. (2009).

Theorem 1. *The one-step $Q(\sigma)$ estimate defined by*

$$Q_{t+1}(S_t, A_t) = (1 - \alpha_t)Q_t(S_t, A_t) + \alpha_t[R_{t+1} + \gamma(\sigma_{t+1}Q_{t+1}(S_{t+1}, A_{t+1}) + (1 - \sigma_{t+1})V_{t+1})], \quad (16)$$

converges to the optimal policy when the following conditions are satisfied:

1. *The size of the set $\mathcal{S} \times \mathcal{A}$ is finite.*
2. *$\alpha_t = \alpha_t(S_t, A_t) \in [0, 1]$, $\sum_t \alpha_t = \infty$, $\sum_t \alpha_t^2 < \infty$ w.p. 1 and $\forall (s, a) \neq (S_t, A_t) : \alpha_t(S_t, A_t) = 0$.*
3. *The policy is greedy in the limit with infinite exploration.*
4. *The reward function is bounded.*

We defer the full details of the proof to the appendix; however, There are two important results from the proof that are worth emphasizing. First, just as with one-step Q -learning, Sarsa, and Expected Sarsa, one-step $Q(\sigma)$ can be used to learn optimal action-value functions. Second, at each time step t it is possible to choose a σ_t such that the contraction property of the $Q(\sigma)$ update is less than or equal to the contraction induced by the Sarsa or Expected Sarsa updates. This implies that it is possible to choose σ_t at every time step in order to speed up convergence.

It is important to note that every TD control method presented thus far can be obtained with $Q(\sigma)$ by varying the sampling parameter, σ ; when $\sigma = 1$, we obtain Sarsa, when $\sigma = 0$, we obtain Expected Sarsa and Tree-backup, and when $\sigma = 1$ for every step of the backup except for the last, where $\sigma = 0$, we obtain n -step Expected Sarsa. Thus, tuning the hyper-parameter σ is not strictly necessary since it can be set to a fixed value in order to obtain one of the existing TD control algorithms. Nevertheless, intermediate values of σ between 0 and 1 create entirely new algorithms that exist somewhere between full sampling and pure expectation and that could result in better performance. Furthermore, σ does not need to remain constant throughout every episode or even at every time step during an episode or continuing task. σ could be varied dynamically as a function of time, of the current state, or of some measure of the learning progress. In particular, σ could also be varied as a function of the episode number, which we investigate in our experiments. There are potentially a variety of effective schemes for choosing and varying σ , and would be a subject for further research.

Experiments

19-State Random Walk

The *19-state random walk*, shown in Figure 2, is a 1-dimensional environment where an agent randomly transitions to one of two neighboring states. There is a terminal state on each end of the environment, transitioning to one of them gives a reward of -1, and transitioning to the other gives a reward of 1. To compare algorithms that involve taking an expectation based on its policy, the task is formulated such that each state had two actions. Each action deterministically transitions to one of the two neighboring states, and the agent learns on-policy under an equiprobable random behavior policy. This differs from typical random walk setups

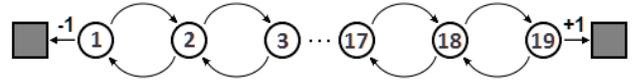


Figure 2: The 19-state random walk MDP. The goal is to accurately estimate the value of each state under equiprobable random behavior.

where each state has one action that will randomly transition to either neighboring state (Sutton and Barto 1998), but the resulting state values are identical.

This environment was treated as a prediction task where a learning algorithm is to estimate a value function under its behavior policy. We conducted an experiment comparing various $Q(\sigma)$ algorithm instances, assessing different multi-step backup lengths, step sizes, and degrees of sampling. The root-mean-square (RMS) error between its estimated value function and the analytically computed values was measured after each episode. Each $Q(\sigma)$ instance and parameter setting ran for 50 episodes and the results are averaged across 100 runs.

Figure 3 shows the results with $n = 3$ and $\alpha = 0.4$, which was found to be representative of the best parameter setting for each instance of $Q(\sigma)$ on this task. Sarsa (full sampling) had better initial performance but poor asymptotic performance, Tree-backup (no sampling) had poor initial performance but better asymptotic performance, and intermediate degrees of sampling traded off between the initial and asymptotic performances. This motivated the idea of dynamically decreasing σ from 1 (full sampling) towards 0 (pure expectation) to take advantage of the initial performance of Sarsa, and the asymptotic performance of Tree-backup. To accomplish this we decreased σ by a factor of 0.95 after each episode. $Q(\sigma)$ with a dynamically varying σ outperformed all of the fixed degrees of sampling.

Stochastic Windy Gridworld

The *windy gridworld* is a tabular navigation task in a standard gridworld which is described by Sutton and Barto (1998). There is a start state and a goal state, and there are four possible moves: right, left, up, and down. When the agent moves into one of the middle columns of the gridworld, it is affected by an upward “wind” which shifts the resultant next state upwards by a number of cells and varies from column to column. If the agent is at the edge of the world and selects a move that would cause it to leave the grid, or would be pushed off the world by the wind, it is simply replaced in the nearest state at the edge of the world. At each time step the agent receives a constant reward of -1 until the goal is reached.

A variation of the windy gridworld, called the *stochastic windy gridworld*, is one where the results of choosing an action are not deterministic. The layout, actions, and wind strengths are the same, but at each time step, with a probability of 10%, the next state that results from picking any action is determined at random from the 8 states currently surrounding the agent.

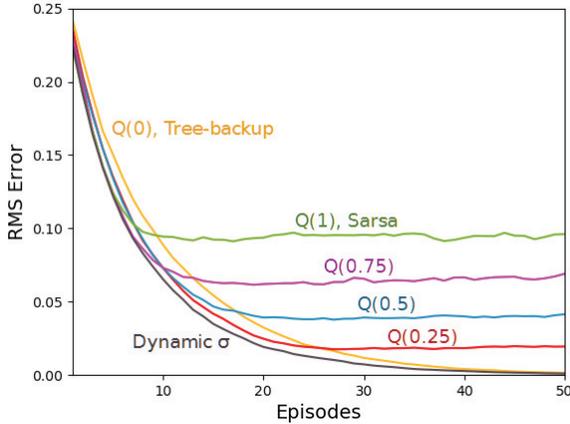


Figure 3: 19-state random walk results. The plot shows the performance of $Q(\sigma)$ in terms of RMS error in the value function. The results are an average of 100 runs, and the standard errors are all less than 0.006. $Q(1)$ had the best initial performance, $Q(0)$ had the best asymptotic performance, and dynamic σ outperformed all fixed values of σ .

We conducted an experiment on the stochastic windy gridworld which consisted of 1000 runs of 100 episodes each to evaluate the performance of various instances of $Q(\sigma)$ with different parameter combinations. All instances of the algorithms behaved and learned according to an ϵ -greedy policy, with $\epsilon = 0.1$. As the performance measure, we compared the average return over the 100 episodes. The results are summarized in Figure 4.

For all the values of σ that we tested, choosing $n = 3$ resulted in the greatest performance; higher and lower values of n decreased the performance. Overall, $Q(\sigma)$ with a dynamic σ performed the best, while $\sigma = 0.5$ was a close second.

Mountain Cliff

We implemented a variant of the classical episodic task, *mountain car*, as described by Sutton and Barto (1998). For this implementation, the rewards, actions and goal remained the same. However, if the agent ever ventured past the top of the leftmost mountain, it would fall off a cliff, be rewarded -100 and returned to a random initial location in the valley between the two hills. We named this environment *mountain cliff*. Both environments were tested and showed the same trend in the results. However, the results obtained in mountain cliff were more pronounced and thus were more suitable for demonstration purposes.

Because the state space is continuous, we approximated q_π using *tile coding* function approximation. Specifically, we used version 3 of Sutton’s tile coding software (n.d.) with 8 tilings, an asymmetric offset by consecutive odd numbers, and each tile taking over 1/8 fraction of the feature space, which gives a resolution of approximately 1.6%.

For each algorithm, we conducted 500 independent runs

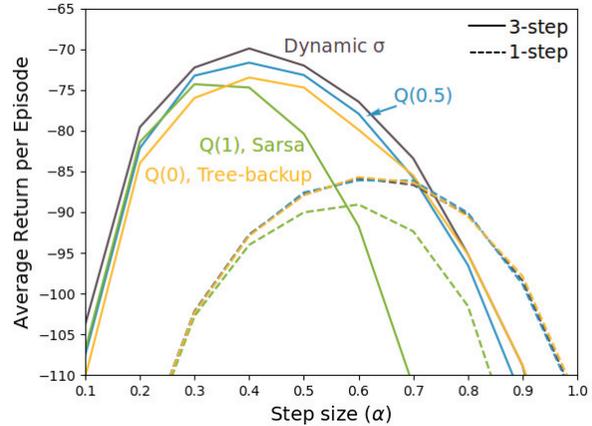


Figure 4: Stochastic windy gridworld results. The plot shows the performance of $Q(\sigma)$ in terms of the average return over 100 episodes as a function of the step size, α , for various values of σ . The results are for selected α values, then are connected by straight lines, and are an average of 1000 runs. The standard errors are all less than 0.3 which is about a line width. 3-step algorithms performed better than their 1-step equivalents, and $Q(\sigma)$ with a dynamic σ performed the best overall.

of 500 episodes each. All training was done on-policy under an ϵ -greedy policy with $\epsilon = 0.1$ and $\gamma = 1$. We optimized for the average return after 500 episodes over different values of the step size parameter, α , and the backup length, n . The results correspond to the best-performing parameter combination for each algorithm: $\alpha = 1/6$ and $n = 4$ for Sarsa; $\alpha = 1/6$ and $n = 8$ for Tree-backup; $\alpha = 1/4$ and $n = 4$ for $Q(0.5)$; and $\alpha = 1/7$ and $n = 8$ for Dynamic σ . We omit n -step Expected Sarsa in the results because its performance was not much different from n -step Sarsa’s performance.

Figure 6 shows the return per episode averaged over 500 runs. To smooth the results, we computed a right-centered moving average with a window of 30 successive episodes. Additionally, we added the average return per episode in a lighter tone to show the variance of each algorithm. As it can be observed, atomic multi-step Sarsa and $Q(0.5)$ had fairly similar performance. Among the atomic multi-step methods with static σ , Tree-backup had the best performance. Nonetheless, $Q(\sigma)$ with dynamic σ outperformed all the algorithms that were using static σ .

In order to gain more insight into the nature of the results, we looked at the average return per episode after 50 (initial performance) and 500 (asymptotic performance) episodes for each algorithm. Additionally, a 95% confidence interval was computed in order to validate the results. After 50 episodes, $Q(0.5)$ had the best performance among the four algorithms with an average return per episode of -398.0; Dynamic σ was a close second with an average return per episode of -406.3. On the other hand, after 500 episodes,

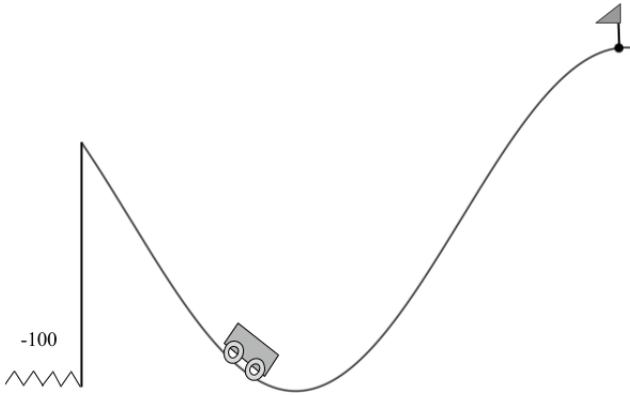


Figure 5: The mountain cliff environment. The goal of the agent is to drive past the flag without falling off the cliff. The agent receives a reward of -1 at every time step, and falling off the cliff returns it to a random initial location in the valley with a reward of -100 .

Dynamic σ managed to outperform all the other algorithms with an average return per episode of -163.7 followed by $Q(0.5)$ with an average return per episode of -167.9 . $Q(1)$ (Sarsa) had the lowest performance with -447.3 average return per episode after 50 episodes and -173.2 after 500 episodes. These results contrast with Figure 6 because the average is taken over all the previous episode instead of the preceding 30 episodes.

Discussion

From our experiments, it is evident that there is merit in unifying the space of algorithms with $Q(\sigma)$. In prediction tasks, such as the 19-state random walk, varying the degree of sampling results in a trade-off between initial and asymptotic performance. In control tasks, such as the stochastic windy gridworld, intermediate degrees of sampling are capable of achieving a higher per-episode average return than either extreme, depending on the number of elapsed episodes.

These findings also extend to tasks with continuous state spaces, such as the mountain cliff. Intermediate values of σ allow for a higher initial performance, whereas small values of σ allow for a better asymptotic performance. As shown in Figure 6, $Q(\sigma)$ with dynamic σ is able to exploit these two benefits by adjusting σ over time.

Moreover, our experiments in the stochastic windy gridworld task demonstrated that it is possible to improve performance by choosing a higher value of the backup length parameter, n . Varying n controls a bias-variance trade-off by adjusting how many rewards are included in the backup before bootstrapping, similar to the parameter λ in the $TD(\lambda)$ algorithm. The parameter σ also has a bias-variance trade-off interpretation, as the Tree-backup algorithm decays the weighting of future rewards based on the stochasticity in the policy (and is therefore more biased). The length parameter n controls the bias-variance trade-off in the direction of

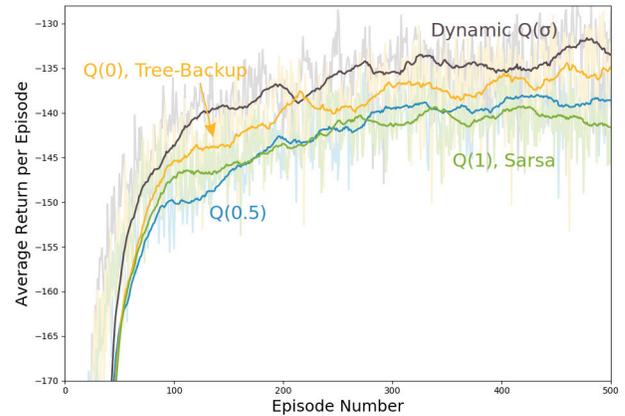


Figure 6: Mountain cliff results. The plot shows the performance of each atomic multi-step algorithm in terms of the average return per episode. The dark lines show the results smoothed using a right-centered moving average with a window of 30 successive episodes, while the light lines show the un-smoothed results. $Q(\sigma)$ with dynamic σ had the best performance among all the algorithms.

the trajectory taken, while the parameter σ manages it by controlling the bootstrapping in the direction of actions not taken. A qualitative result that illustrates the bias-variance trade-off induced by the parameter σ can be observed in the 19-State Random Walk experiment. A large value of σ results in lower bias at the beginning of training and a lower RMS error as a consequence. However, as the bias of the return decreases in the asymptote, the low variance inherent to small values of σ result in more accurate estimates of the action-value function.

Conclusions

In this paper we studied $Q(\sigma)$, which is a unifying algorithm for multi-step TD control methods. $Q(\sigma)$, through the use of the sampling parameter σ , allows for continuous variation between updating based on full sampling and updating based on pure expectation. Our results on prediction and control problems showed that an intermediate fixed degree of sampling can outperform the methods that exist at the extremes (Sarsa and Tree-backup). In addition, we presented simple way of dynamically adjusting σ which outperformed any fixed degree of sampling.

Our presentation of $Q(\sigma)$ was limited to the atomic multi-step case without eligibility traces, we only conducted experiments on on-policy problems, and we only investigated one simple method for dynamically varying σ . This leaves open several avenues for future research. First, $Q(\sigma)$ could be extended to use eligibility traces and compound backups. Second, the performance of $Q(\sigma)$ could be evaluated on off-policy problems. Third, other schemes for dynamically varying σ could be investigated – perhaps as a function of state, the recently observed rewards, or some measure of the learning progress.

Acknowledgments

The authors thank Vincent Zhang, Harm van Seijen, Doina Precup, and Pierre-luc Bacon for insights and discussions contributing to the results presented in this paper, and the entire Reinforcement Learning and Artificial Intelligence research group for providing the environment to nurture and support this research. We gratefully acknowledge funding from Alberta Innovates – Technology Futures, Google Deepmind, and from the Natural Sciences and Engineering Research Council of Canada.

Appendix

Proof of Theorem 1

Let $\mathcal{X} = \mathcal{S} \times \mathcal{A}$, $X_t = (S_t, A_t) \in \mathcal{X}$, $\bar{R}_t = \mathbb{E}\{R_t\}$, and Q^* be the optimal action-value function defined as

$$Q^*(S_t, A_t) = \bar{R}_{t+1} + \gamma \mathbb{E}\{\max_a Q^*(S_{t+1}, a)\}. \quad (17)$$

We define a new stochastic process $(\alpha_t, \Delta_t, F_t)_{t \geq 0}$ by subtracting $Q^*(X_t)$ from both sides of equation (16)

$$\Delta_{t+1}(X_t) = (1 - \alpha_t(X_t))\Delta_t(X_t) - \alpha_t(X_t)F_t(X_t),$$

and letting $\alpha_t \in (0, 1]$, $\Delta_t(X_t) = Q_t(X_t) - Q^*(X_t)$, and $F_t = R_{t+1} + \gamma[\sigma_{t+1}Q_t(X_{t+1}) + (1 - \sigma_{t+1})V_{t+1}] - Q^*(X_t)$. Additionally, let P_t be a sequence of increasing σ -fields representing the history such that α_0 and Δ_0 are P_0 -measurable and α_t , Δ_t , and F_{t-1} are P_t -measurable for $t \geq 1$.

Proving that Δ_t converges to 0 as $t \rightarrow \infty$ is equivalent to showing that Q_t converges to Q^* as $t \rightarrow \infty$. Consequently, the proof is equivalent to showing that the conditions of lemma 1 from Sing et al. (2000) are satisfied for Δ_t .

Conditions one, two, and three of the lemma are satisfied by the corresponding assumptions of the theorem. Hence, we only need to show that $\|\mathbb{E}\{F_t|P_t\}\| \leq k\|\Delta_t\| + C_t$ where $\|\cdot\|$ is the maximum norm, $k \in [0, 1)$, and C_t goes to 0 with probability 1. By adding and subtracting $\max_a Q_t(S_t, a)$, using the definition of Q^* and the triangle inequality, we can show that

$$\begin{aligned} & \|\mathbb{E}\{F_t|P_t\}\| \\ & \leq \|\mathbb{E}\{R_{t+1} + \gamma \max_a Q_t(S_{t+1}, a) - Q^*(S_t, A_t)\}\| \\ & \quad + \gamma \|\mathbb{E}\{\sigma_{t+1}Q_t(s_{t+1}, a_{t+1}) + (1 - \sigma_{t+1})V_{t+1} \\ & \quad \quad - \max_a Q_t(s_{t+1}, a)\}\| \\ & = \gamma \|\mathbb{E}\{\max_a Q_t(S_{t+1}, a) - \max_b Q^*(S_{t+1}, b)\}\| + C_t \\ & \leq \gamma \max_s |\max_a Q_t(s, a) - \max_b Q^*(s, b)| + C_t \\ & \leq \gamma \max_s \max_a |Q_t(s, a) - Q^*(s, a)| + C_t \\ & = \gamma \|\Delta_t\| + C_t. \end{aligned}$$

Note that if the policy is greedy and $\sigma_{t+1} \in [0, 1]$, then $\sigma_{t+1}Q_t(S_{t+1}, A_{t+1}) + (1 - \sigma_{t+1})V_{t+1} = \max_a Q_t(S_{t+1}, a)$. Therefore, C_t goes to 0 as the policy becomes greedy in the limit. Consequently, condition 3 of lemma 1 from Sing et al. (2000) is satisfied. Therefore, Δ_t converges to 0 w.p. 1, which implies that Q_t converges to Q^* w.p. 1. \square

References

- Jaakkola, T.; Jordan, M. I.; and Singh, S. P. 1994. On the convergence of stochastic iterative dynamic programming algorithms. *Neural Computation* 6(6):1185–1201.
- Precup, D.; Sutton, R. S.; and Singh, S. 2000. Eligibility traces for off-policy policy evaluation. In Kaufman, M., ed., *Proceedings of the 17th International Conference on Machine Learning*, 759–766.
- Rummery, G. A., and Niranjan, M. 1994. On-line Q-learning using connectionist systems. Technical report, CUED/F-INFENG/TR 166, Engineering Department, Cambridge University.
- Rummery, G. A. 1995. *Problem Solving with Reinforcement Learning*. Ph.D. Dissertation, Cambridge University.
- Singh, S.; Jaakkola, T.; Littman, M. L.; and Szepesvári, C. 2000. Convergence results for single-step on-policy reinforcement-learning algorithms. *Machine Learning* 38(3):287–308.
- Sutton, R. S., and Barto, A. G. 1998. *Reinforcement Learning: An Introduction*. Cambridge, Massachusetts: MIT Press.
- Sutton, R. S., and Barto, A. G. 2018. *Reinforcement Learning: An Introduction*. 2nd edition. Manuscript in preparation.
- Sutton, R. S. 1988. Learning to predict by the methods of temporal differences. *Machine learning* 3(1):9–44.
- Sutton, R. S. 1996. Generalization in reinforcement learning: Successful examples using sparse coarse coding. In Touretzky, D. S., and Hasselmo, M. E., eds., *Advances in Neural Information Processing Systems* 8, 1038–1044. MIT Press.
- van Seijen, H.; van Hasselt, H.; Whiteson, S.; and Wiering, M. 2009. A theoretical and empirical analysis of expected Sarsa. In *Proceedings of the IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning*, 177–184.
- Watkins, C. J. C. H., and Dayan, P. 1992. Q-learning. *Machine learning* 8(3-4):279–292.
- Watkins, C. J. C. H. 1989. *Learning from Delayed Rewards*. Ph.D. Dissertation, Cambridge University.