
ORIGINAL ARTICLE

An introduction to reinforcement learning for neuroscience

Kristopher T. Jensen^{1,2}

¹Sainsbury Wellcome Centre, University College London

²Computational and Biological Learning Lab, University of Cambridge

Correspondence

Kristopher Jensen, Sainsbury Wellcome Centre, University College London, London W1T 4JG, UK

Email: kris.torp.jensen@gmail.com

Reinforcement learning (RL) has a rich history in neuroscience, from early work on dopamine as a reward prediction error signal (Schultz et al., 1997) to recent work proposing that the brain could implement a form of ‘distributional reinforcement learning’ popularized in machine learning (Dabney et al., 2020). There has been a close link between theoretical advances in reinforcement learning and neuroscience experiments throughout this literature, and the theories describing the experimental data have therefore become increasingly complex. Here, we provide an introduction and mathematical background to many of the methods that have been used in systems neuroscience. We start with an overview of the RL problem and classical temporal difference algorithms, followed by a discussion of ‘model-free’, ‘model-based’, and intermediate RL algorithms. We then introduce deep reinforcement learning and discuss how this framework has led to new insights in neuroscience. This includes a particular focus on meta-reinforcement learning (Wang et al., 2018) and distributional RL (Dabney et al., 2020). Finally, we discuss potential shortcomings of the RL formalism for neuroscience and highlight open questions in the field. Code that implements the methods discussed and generates the figures is also [provided](#).

1 | INTRODUCTION

Humans and other animals learn from their experiences. Sometimes, this takes the form of explicit demonstration, as is often the case during formal education. However, we also often have to learn from trial and error together with feedback received from the world around us – sometimes implicit and sometimes explicit. This is well illustrated by the classical case study of Pavlov's dogs, who learned to associate a so-called 'conditioned stimulus' (CS; e.g. the ringing of a bell) with the availability of food shortly after (the 'unconditioned stimulus'; US). Following a brief period of learning, the dogs would start to salivate in response to the CS in advance of any food actually being served. This suggests that the dogs had learned to associate the CS with the availability of 'reward' in the form of food, and that they produced an appropriate physiological response to take advantage of this food availability. Importantly, this occurred without any explicit instruction or description of the sequence of events preceding food being served. Instead, the dogs learned from experience with their environment and the presence of a salient, rewarding stimulus.

Such passive stimulus-response predictions are also called 'Pavlovian learning' and have been commonly used in neuroscience to study learning from external rewards (Niv, 2009). This forms a specific instantiation of the concept of 'reinforcement learning', which is a general term for settings where an agent learns from reward signals in the environment rather than explicit demonstration, as is the case in 'supervised learning'. Importantly, the past decades have shown that principles of reinforcement learning can be used to explain not just behaviour, but also neural activity in biological circuits (Niv, 2009; Botvinick et al., 2020). An explicit neural basis of RL was initially demonstrated in foundational work by Schultz et al. (1997), which showed that the firing rates of dopaminergic neurons in the Ventral Tegmental Area (VTA) reflected the difference between expected and actual 'value' when animals received a juice reward following a CS consisting of a lever-press in response to a small light turning on. This provided a potential neural substrate of the classical 'temporal difference' learning algorithm (Schultz et al., 1997; Sutton, 1988), which has since been expanded to a wealth of evidence for reinforcement learning in neural dynamics (Niv, 2009; Dabney et al., 2020; Watabe-Uchida et al., 2017). However, these classical algorithms are generally restricted to simple problem settings, while humans and other animals are capable of solving complex high-dimensional problems involving extended planning and fine motor control. The field of 'deep reinforcement learning' has recently emerged to tackle such problems in a machine learning setting, which has led to impressive results across a range of tasks (Mnih et al., 2013; Schrittwieser et al., 2020; Wurman et al., 2022; Vinyals et al., 2019). Intriguingly, recent research has demonstrated that these deep RL algorithms also have parallels in both behaviour and neural dynamics (Botvinick et al., 2020; Wang et al., 2018; Dabney et al., 2020; Jensen et al., 2023; Aldarondo et al., 2024), suggesting that neuroscience can continue to learn from advances in reinforcement learning.

In this review, we provide an overview of the reinforcement learning problem and popular algorithms, with a particular focus on parallels and uses of these algorithms in neuroscience. This overview starts from classical tabular TD learning and Q-learning algorithms, which have guided neuroscience research for decades. We then consider the important distinction between model-based and model-free reinforcement learning, as well as methods that fall somewhere in the gray area between these extremes, and discuss their neural correlates. Finally, we generalize the tabular methods to the non-linear function approximation setting and the resulting deep RL methods, which have revolutionized machine learning in recent years. We do this with a focus on methods that have had a strong influence on neuroscience to give the reader a better idea of the mathematical and computational background of recent neuroscientific findings. These include the 'meta-reinforcement learning' model of PFC by Wang et al. (2018) and the 'distributional reinforcement learning' model of VTA dopaminergic neurons by Dabney et al. (2020) in particular. We hope this review will be useful both for those who are interested in the theory underlying reinforcement learning in neuroscience and for those who want an overview of how the neuroscience literature builds on principles from reinforcement learning theory. Throughout the paper, the focus will be on an intuitive understanding of the relevant RL methods, and explicit

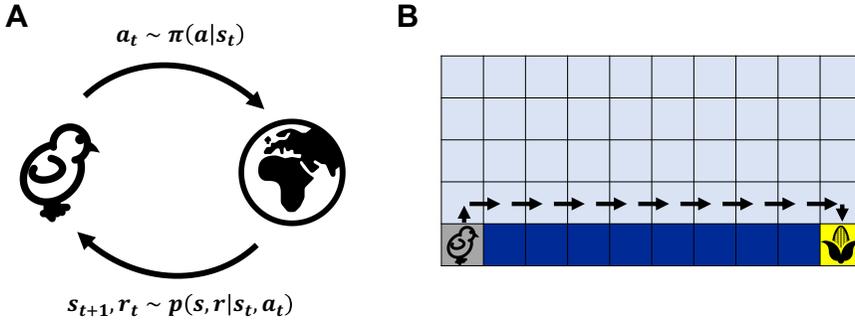


FIGURE 1 The reinforcement learning problem and cliffworld environment. **(A)** An agent (here the bird) interacts with the world to maximize reward. This involves a balance between exploring potentially interesting new states (e.g. searching for food in a new field) while also exploiting states known to yield high reward (e.g. the field that had many worms yesterday). At a given point in time, the bird is in some state s_t from which it can take an action a_t , with the probability of different actions determined by the ‘policy’ $\pi(a|s_t)$, which is controlled by the agent. a_t then leads to a change in the environment according to the non-controllable environment dynamics $s_{t+1}, r_t \sim p(s, r|s_t, a_t)$. Here, r_t is the empirical ‘reward’ received by the agent, and its objective is to collect as much cumulative reward as possible. Often, reinforcement learning problems are divided into ‘episodes’, with the agent learning over the course of multiple repeated exposures to the environment. This could for example consist of the bird learning over multiple days which fields are likely to be rich in food, while minimizing the distance travelled and exposure to predators. **(B)** The ‘cliffworld’ environment, which will be used to demonstrate the performance and behaviour of a range of reinforcement learning algorithms in this work. The agent starts in the lower left corner (location $[0, 0]$), and the episode finishes when it encounters either the ‘cliff’ (dark blue) or the goal (yellow; location $[9, 0]$). If the agent walks off the cliff, it receives a reward of -100 . If it finds the goal, it receives a reward of $+50$. In any other state, it receives a reward of -1 . Such negative rewards for ‘neutral’ actions are commonly used to encourage the agent to achieve its goal as fast as possible. The arrows indicate the ‘optimal’ policy, which takes the agent to the goal via the shortest possible route that avoids the cliff.

derivations are included only where we consider them conducive to such understanding. We refer to [Sutton and Barto \(2018\)](#) for a more in-depth treatment of the underlying theory.

2 | PROBLEM SETTING

Here we provide a short introduction to the reinforcement learning problem in a discrete state and action space with a finite time horizon – a common setting for neuroscience experiments consisting of repeated trials or episodes in a controlled environment. In this setting, the environment consists of states from a discrete set $s \in \mathcal{S} = \{s_1, s_2, \dots\}_1^{|\mathcal{S}|}$, and the agent can take actions $a \in \mathcal{A} = \{a_1, a_2, \dots\}_1^{|\mathcal{A}|}$. The environment is characterized by transition and reward probabilities $p(s_{t+1}, r_t|s_t, a_t)$, where r_t is the reward at time t . We will use $r(s, a)$ to denote either the reward when it depends deterministically on the state and action, or its expectation otherwise. We will further make the *Markov assumption* that the next state and reward only depend on the current state and action, $p(s_{t+1}, r_t|s_t, a_t, s_{t-1}, a_{t-1}, \dots, s_0, a_0) = p(s_{t+1}, r_t|s_t, a_t)$.

We can now define a *trajectory* $\tau = \{s_t, a_t, r_t\}_{t=0}^T$. The probability of a trajectory occurring is

$$p_\pi(\tau) = p(s_0) \prod_{t=0}^T p(s_{t+1}, r_t|s_t, a_t) \pi(a_t|s_t). \quad (1)$$

$p_\pi(\tau)$ depends on the *policy* of the agent, $\pi(a|s)$, which specifies the probability of taking action a in state s (Figure 1A).

The objective is to learn a policy that maximizes the expected total discounted reward

$$J(\pi) = \mathbb{E}_{\tau \sim p_{\pi}(\tau)} [R_{\tau}] = \mathbb{E}_{\tau \sim p_{\pi}(\tau)} \left[\sum_{t=0}^T \gamma^t r_t | \tau \right], \quad (2)$$

where $R_{\tau} := \sum_{t=0}^T \gamma^t r_t | \tau$, and we have written $J(\pi)$ since the policy uniquely specifies J in a stationary environment. In Equation 2, γ is a ‘discount factor’, which stipulates that we should care more about immediate rewards than rewards far in the future. We can provide three interpretations for this discount factor. One is that agents intrinsically care more about immediate reward than distant reward. A second is that there is a fixed non-zero probability $(1 - \gamma)$ of the current ‘episode’ or environment terminating or changing at each timestep, in which case we should weight putative future reward by the probability that we are still engaged in the task at that time. The third view is that γ simply provides a tool for reducing the variance of our learning methods, especially in temporally extended tasks. This third view is most compatible with the fact that *evaluation* of RL agents after training is generally done without discounting.

Since $J(\pi)$ depends on the policy of the agent, it is possible to search in the space of policies for one that maximizes J , which is the topic of reinforcement learning. It is often assumed that the experience $\{\tau\}$ is generated by the agent acting according to its policy, and the resulting experience is then used to update the policy in a way that increases $J(\pi)$. However, ‘off-policy’ and ‘offline’ reinforcement learning methods also exist, where the agent learns on the basis of experience generated by a policy different from π (Levine et al., 2020; Section 11.2). This can be either an ‘old’ version of the agent itself, when it acted according to a different policy, or data generated by an entirely different agent. Off-policy learning is important for biological organisms, where learning can happen ‘offline’ during sleep after initial data collection during wake, or from observing other individuals (also the topic of ‘imitation learning’; Section 11.3).

3 | TEMPORAL DIFFERENCE LEARNING

A simple way to maximize reward in an environment is to learn the ‘value’ of different states, and then move towards states with high value. The potential importance of such an algorithm for neuroscience is evident from the value-seeking behaviour of many organisms, and the widespread findings of neural ‘codes’ for value across the brain (Schultz et al., 1992; Padoa-Schioppa and Assad, 2006; Rushworth et al., 2011). This leaves the question of how such value codes can be learned in a biologically plausible setting.

One answer to this question takes the form of the classical ‘temporal difference learning’ algorithm (Sutton, 1988; Sutton and Barto, 2018). This involves defining a *value function* for a given state s and policy π , which quantifies the expected future reward when following π starting from s :

$$V^{\pi}(s) = \mathbb{E}_{\tau \sim p_{\pi}(\tau)} \left[\sum_{t' \geq t} \gamma^{t'-t} r_{t'} | s_t = s \right]. \quad (3)$$

Here, $\mathbb{E}_{\tau \sim p_{\pi}(\tau)}[\cdot]$ indicates an expectation taken over trajectories τ resulting from the agent following the policy π . For the true value function, we can expand this as a self-consistency equation

$$V^{\pi}(s) = r_{\pi}(s) + \sum_{s'} p_{\pi}(s_{t+1} = s' | s_t = s) \mathbb{E}_{\tau \sim p_{\pi}(\tau)} \left[\sum_{t'=t+1} \gamma^{t'-t} r_{t'} | s_{t+1} = s' \right] \quad (4)$$

$$= r_{\pi}(s) + \gamma \sum_{s'} p_{\pi}(s_{t+1} = s' | s_t = s) V^{\pi}(s'), \quad (5)$$

where $p_\pi(s_{t+1} = s' | s_t = s) = \sum_a \pi(a|s)p(s_{t+1} = s' | s_t = s, a_t = a)$ is the probability of transitioning from s to s' under π , and $r_\pi(s) = \mathbb{E}_{a \sim \pi}[r(s, a)]$ is the expected reward in state s , averaged over actions. Importantly, Equation 5 would not hold if $V^\pi(s)$ was not the true value function (Sutton and Barto, 2018). When learning an approximate value function $V(s)$, we can therefore use this bootstrapped self-consistency relation as an objective function:

$$\mathcal{L} \propto (V(s) - (r_\pi(s) + \gamma \mathbb{E}_{p_\pi(s'|s)} [V(s')]))^2, \quad (6)$$

Gradient descent w.r.t $V(s)$ gives us an update rule

$$\Delta V(s) \propto -\frac{\partial \mathcal{L}}{\partial V(s)} \quad (7)$$

$$\propto -V(s) + r_\pi(s) + \gamma \mathbb{E}_{p_\pi(s'|s)} [V(s')] \quad (8)$$

$$\approx -V(s_t) + r_t + \gamma V(s_{t+1}). \quad (9)$$

The last line approximates the expected update with a single sample corresponding to the states and reward actually experienced. As more experience is collected and many small gradient steps are taken according Equation 9, these single-sample estimates average out to the expectation in Equation 8 (Figure 2A). Variants of this algorithm can also learn about multiple past states at once using the notion of *eligibility traces* (Sutton and Barto, 2018). However, Equation 9 is the canonical temporal difference learning rule (Sutton, 1988), and it leads to learning dynamics where the temporal difference error $\delta_t := -V(s_t) + r_t + \gamma V(s_{t+1})$ gradually propagates from the rewarding state to prior states that predict the upcoming reward (Figure 2C).

This gradual propagation of prediction errors from the reward state to its predecessors has been of great interest in neuroscience. In particular, classical work by Schultz et al. (1997) demonstrated a similar pattern of neural activity in dopaminergic VTA neurons, which formed the foundation of a now well-established theory that dopamine provides a biological reward prediction error signal that drives learning (Niv, 2009; Watabe-Uchida et al., 2017). At a behavioural level, this is supported by experiments showing that artificial stimulation of dopamine neurons can be a strong driver of learning (Olds and Milner, 1954; Tsai et al., 2009; Steinberg et al., 2013). The simple narrative of dopamine as a reward prediction error has also been challenged in recent years (Coddington and Dudman, 2019; Howe et al., 2013; Horvitz, 2000), which has led to theories of dopamine as a more general prediction error for both value but also other quantities like salience and movement (Kakade and Dayan, 2002; Gershman et al., 2024).

Much classical work in neuroscience has focused on value learning in Pavlovian conditioning tasks, but animals in natural environments also have to take actions on the basis of this information. However, after learning a value function, it can be used for optimal action selection if we can estimate $r(s, a)$ and $p(s'|s, a)$. We can then compute a 'state-action value' $Q^\pi(s, a)$, defined as the expected discounted future reward associated with taking action a in state s and then following policy π :

$$Q^\pi(s, a) := \mathbb{E} \left[\sum_{t'=t} \gamma^{t'-t} r_{t'} | s_t = s, a_t = a \right] = r(s, a) + \gamma \sum_{s'} p(s'|s, a) V^\pi(s'). \quad (10)$$

Approximating the true value function V^π by the learned value function V in Equation 10 yields an approximate state-action value Q , which can be used to choose the action with the highest expected reward,

$$a^*(s) = \operatorname{argmax}_a Q(s, a). \quad (11)$$

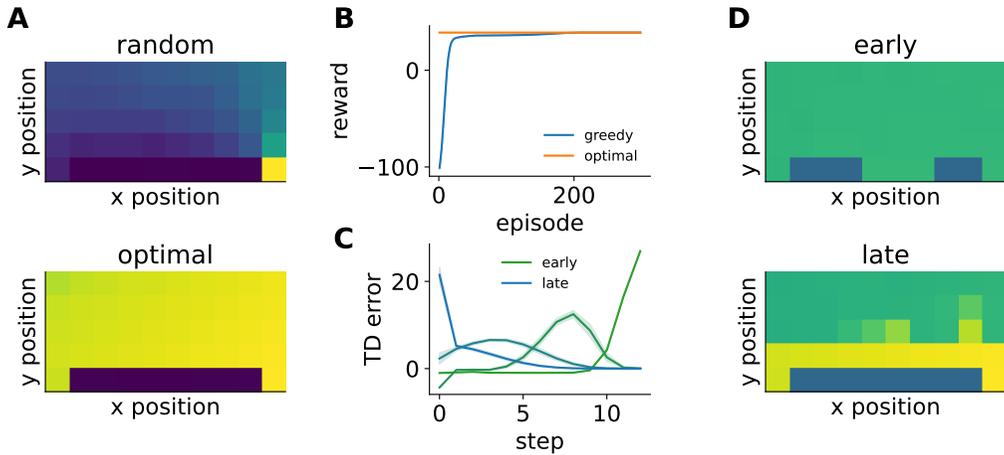


FIGURE 2 Temporal difference learning. (A) Value functions acquired through temporal difference learning (Equation 9) while acting according to either a random (top) or an optimal (bottom) policy. These simulations were performed with a random start state in the cliffworld environment to ensure full coverage of the space. Dark blue indicates negative expected reward (-100) and yellow indicates positive expected reward (+50). These simulations used a learning rate of $\alpha = 0.05$ and no temporal discounting ($\gamma = 1$). Under the random policy, states near the cliff have low value even if they are close to the goal, since the agent often falls off the cliff from there. Under the optimal policy, all states have high expected reward, since the agent always reaches the goal. States nearer the goal have slightly higher value than those further away. (B) Empirical reward as a function of episode number for a TD-learning agent that acts according to Equation 11 while updating its value estimates according to Equation 9. For this agent, action selection assumes access to a ‘one-step’ world model in order to evaluate the consequence of each putative action. The agent gradually converges to an optimal policy. Parameters for the agent are as in (A), except that the start state is always the lower left corner. (C) TD error (Equation 9) as a function of the step number along the optimal path for the agent in (B) at different stages of learning (green to blue). This TD signal gradually propagates backwards from the reward to preceding states, mirroring biological recordings of dopamine activity (Schultz et al., 1997). (D) Value function learned by a greedy TD agent as in (B), plotted either early (top) or late (bottom) in training. Early in training, the agent has learned that the cliff is bad but doesn’t know where the goal is or how to get there. Late in training, the agent has learned a value function that locally resembles the optimal value function from (A), while it has not learned the value of distant states that are rarely or never visited from the start state. This is a potential shortcoming of ‘greedy’ agents that can easily converge to a sub-optimal local maximum in more complicated environments. For this analysis, we used a high learning rate of $\alpha = 0.5$ to make the early TD updates larger and therefore more visible.

Updating the value function according to Equation 9 while acting in the environment according to Equation 11 leads to an agent that gradually learns to take better actions as it learns a better value function (Figure 2B-D). This provides a biologically plausible algorithm for reward-driven learning in agents with access to a one-step predictive model.

4 | Q-LEARNING

In some cases, we may not know the transition function or it could be expensive to simulate. Additionally, it has been found that dopamine activity can reflect learning signals not just for state values but also for *action* values (Roesch et al., 2007; Morris et al., 2006). This suggests an alternative model of biological learning, where animals directly learn the state-action values defined in Equation 10. This is in contrast to the algorithm in Section 3, where the agent only learned the *state* values, and then computed the action values at decision-time using a one-step world model. Q-learning is the most prominent model for learning state-action values, and it has commonly been used to explain animal behaviour and neural activity (Niv, 2009; Mattar and Daw, 2018).

To learn the Q-values necessary for action selection directly, we start by expanding Equation 10,

$$Q^\pi(s, a) = r(s, a) + \gamma \sum_{s'} p(s_{t+1} = s' | s_t = s, a_t = a) \sum_{a'} \pi(a' | s') Q^\pi(s', a'). \quad (12)$$

For the greedy policy $\pi^g(a|s) := I_{a=a^*(s)}$ (where the indicator function $I_{a=b} = 1$ for $a = b$ and 0 otherwise), this gives rise to a self-consistency expression for the state-action values:

$$Q^{\pi^g}(s, a) = r(s, a) + \gamma \mathbb{E}_{s' \sim p(s'|s, a)} \left[\max_{a'} Q^{\pi^g}(s', a') \right]. \quad (13)$$

Importantly, this self-consistency expression only holds when the Q-values have converged to the true expected rewards, and the associated greedy policy is therefore optimal (Sutton and Barto, 2018). We can now use Equation 13 as an objective by defining

$$\mathcal{L} \propto \left(Q(s, a) - \left(r(s, a) + \gamma \mathbb{E}_{s' \sim p(s'|s, a)} \left[\max_{a'} Q(s', a') \right] \right) \right)^2, \quad (14)$$

Gradient descent w.r.t $Q(s, a)$ gives us an update rule

$$\Delta Q(s, a) \propto -Q(s, a) + r(s, a) + \gamma \mathbb{E}_{s' \sim p(s'|s, a)} \left[\max_{a'} Q(s', a') \right] \quad (15)$$

$$\approx -Q(s_t, a_t) + r_t + \gamma \max_{a'} Q(s_{t+1}, a'). \quad (16)$$

This is the so-called Q-learning update rule (Watkins, 1989; Figure 3A), where we have estimated the expectation with the single sample actually seen by the agent in the last line.

Q-learning is guaranteed to converge to the optimal policy in the limit of infinitesimal learning rates and infinite sampling of the state-action space (Watkins and Dayan, 1992; Sutton and Barto, 2018). However, following the greedy policy $a^*(s) = \operatorname{argmax}_a Q(s, a)$ before convergence of the Q-values can lead to undersampling of the state-action space and poor performance. It is therefore common to either use an 'ε-greedy' policy, $\pi(a|s) = \epsilon/|\mathcal{A}| + (1-\epsilon)I_{a=a^*(s)}$, or a softmax-policy, $\pi(a|s) \propto \exp(\beta Q(s, a))$, to collect the experience used to update the Q-values (Figure 3B). Such exploration strategies and their biological correlates are discussed in more detail in Section 9.

These approaches make Q-learning an 'off-policy' algorithm, since the policy used in the learning update (the greedy policy) is different from the policy used for action selection (the stochastic policy). An on-policy alternative known as 'SARSA' (state-action-reward-state-action) is also commonly used, where the update rule uses the Q-value corresponding to the action a_{t+1} sampled at the next timestep instead of the greedy action (Figure 3C):

$$\Delta Q(s_t, a_t) \propto -Q(s_t, a_t) + r_t + \gamma Q(s_{t+1}, a_{t+1}). \quad (17)$$

This will converge to the true Q-values for a given policy π , similar to how the TD learning rule in Equation 9 converges to the true value function for a given policy, again under assumptions of infinite sampling of the space.

When animals have to choose between actions with different values, studies have found evidence for midbrain dopamine neurons encoding the prediction error used for either Q-learning (Roesch et al., 2007; Niv, 2009) or SARSA (Morris et al., 2006; Niv, 2009). They may therefore not just be abstract learning algorithms, but instead have plausible implementations in biological neural circuits. However, the methods considered so far also have notable shortcomings. For example, the amount of data needed to learn state(-action) values and the assumption of a stationary environment can be prohibitive for animals needing to act in a rapidly changing world, where bad decisions have fatal consequences.

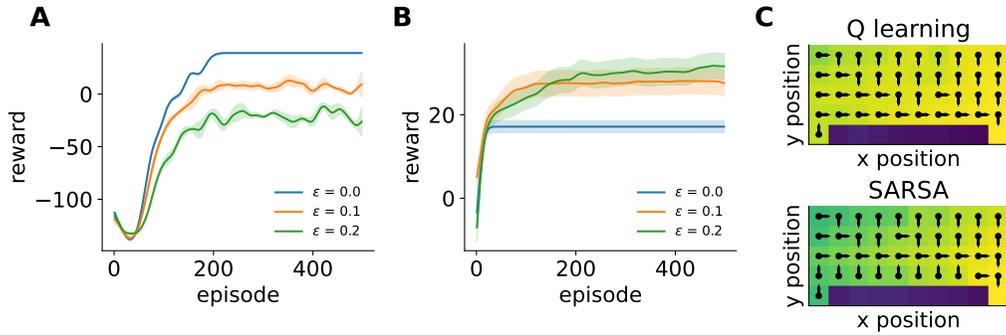


FIGURE 3 Q-learning. (A) Empirical reward as a function of episode number for Q-learners with different levels of stochasticity in their policy ($\epsilon \in \{0, 0.1, 0.2\}$; legend). For these simulations, we used a learning rate of $\alpha = 0.05$ for all agents and no temporal discounting ($\gamma = 1$). The agent with $\epsilon = 0$ converges to an optimal policy, similar to the TD agent in Figure 2A. However, convergence is in this case slower despite using the same learning rate, because the Q-learner has to learn about each action independently, while the TD agent used its one-step world model to aggregate learning across actions reaching the same state. In this cliffworld environment, increasing epsilon leads to worse performance since it increases the probability of falling off the cliff. Additionally, there is no risk of getting stuck in a local minimum since there is only one rewarding state, which decreases the value of exploration. Lines and shading indicate mean and standard error across 10 simulations. (B) As in (A), now for a non-cliffworld grid environment with two goals: one with a reward of +20 at location (0, 4), and one with a reward of +50 at location (5,0). In this case, having non-zero epsilon can increase the probability of discovering the 'high reward' goal rather than getting stuck with a locally optimal policy of moving to the 'low reward' goal. In these simulations, we used a learning rate of $\alpha = 1$, since this effect is less robust with lower learning rates that lead to more exploration of the environment across all agents. (C) Cliffworld policy learned by a Q-learning (top) or SARSA (bottom) agent with $\epsilon = 0.3$. Colours indicate the maximum value of any action in a state from blue (-100) to yellow (+50), and arrows indicate which action has the highest value. The Q-learning agent learns to move right above the cliff, because this is the optimal thing to do under the assumption that subsequent actions are also optimal. This is because it is an 'off-policy' algorithm that does not take into account the actual policy of the agent. In contrast, the SARSA agent learns to move a 'safe distance' away from the cliff, since it is an 'on-policy' algorithm that takes into account the finite probability of the agent choosing to move off the cliff from upcoming states. Q-learning agents are also frequently trained using a stochastic ϵ -greedy policy and then evaluated with the greedy policy corresponding to $\epsilon = 0$, or they can be trained while 'annealing' ϵ from some finite value to 0 over several episodes to allow for initial exploration.

5 | MODEL-FREE AND MODEL-BASED REINFORCEMENT LEARNING

We have so far considered what is known as 'model-free' reinforcement learning algorithms. These involve learning a stimulus-response function that says 'when in state s , take action a '. Such algorithms do not require much computation at decision time, where they rely on cached state or action values. However, it can require a lot of experience with the environment to learn these model-free policies, and they can be inflexible in changing environments. This is incompatible with many aspect of animal behaviour, which we know is adaptive and can benefit from 'latent learning' in an environment before a reward-driven task is ever encountered (Blodgett, 1929; Tolman, 1948).

On the other hand, 'model-based' reinforcement learning uses a model of the world to simulate the consequences of different actions at decision time. This can be much more data efficient, since learning a world model is often easier than learning a full policy (Figure 4A). In machine learning settings, model-based RL has exhibited impressive performance across a range of tasks with large state spaces, including Atari, chess, shogi, and Go (Silver et al., 2018; Schrittwieser et al., 2020; Deisenroth and Rasmussen, 2011). In a biological context, the idea of first learning a model of the environment, and then using it to guide reward-driven behaviour also provides one plausible explanation for latent learning and other types of rapid adaptation. However, model-based decision making can be computationally

intensive at decision-time, which is a challenge for animals that rely on rapid decision making for survival (Figure 4B).

In model-based RL, an approximate transition-and-reward function $\tilde{p}(s', r|s, a)$ is learned from past experience. Once this model has been learned, it can be used for planning at decision time. This can be done for example by expanding the Q-value relation from Equation 13:

$$Q(s_t, a_t) \approx r(s_t, a_t) + \gamma \mathbb{E}_{\tilde{p}(s_{t+1}|s_t, a_t)} [\text{argmax}_{a_{t+1}} Q(s_{t+1}, a_{t+1})] \quad (18)$$

$$\approx r(s_t, a_t) + \gamma \mathbb{E}_{\tilde{p}(s_{t+1}|s_t, a_t)} [\text{argmax}_{a_{t+1}} [r(s_{t+1}, a_{t+1}) + \gamma \mathbb{E}_{\tilde{p}(s_{t+2}|s_{t+1}, a_{t+1})} [\text{argmax}_{a_{t+2}} Q(s_{t+2}, a_{t+2})]]]] \quad (19)$$

$$= \dots \quad (20)$$

If the environment is deterministic, $p(s'|s, a)$ is a delta function, and otherwise the next-state expectations may need to be approximated with multiple samples. Unfortunately, optimizing over all possible action sequences in Equation 18 is in general an exponentially large search problem in the planning depth, which makes it infeasible for any reasonably sized problem. It is therefore common to either use 'depth-first search' with limited breadth, or 'breadth-first search' with limited depth. In breadth-first search, we consider all possible actions at each level of the search tree but terminate the search at a finite depth, instead using cached 'model-free' state-values to estimate the reward-to-go from the terminal states. Such 'plan-until-habit' has also been proposed as a model of human behaviour (Keramati et al., 2016). In depth-first search, we instead sample a series of paths from s_t to termination (or some upper bound), using a heuristic to prioritize actions expected to be good, and then pick an action with high expected reward (Huys et al., 2012).

For both of these strategies, it is necessary to trade off the temporal opportunity cost of planning with the increase in expected reward (Botvinick and Cohen, 2014; Agrawal et al., 2022). This has been a popular research area in cognitive science, where a wealth of literature on 'resource-rational' decision making has emerged in recent years (Griffiths et al., 2019; Callaway et al., 2022). However, this literature has often focused on the behaviour of optimal agents, with less focus on the learning process and neural mechanisms that might implement the necessary computations. Bridging this gap, recent work has suggested that frontal cortex and striatum might initially store a 'model-free' policy in its network state, which is gradually updated with model-based information from the hippocampal formation until the policy improvement is outweighed by the temporal opportunity cost of planning (Jensen et al., 2023).

While several model-based and model-free reinforcement learning methods have thus been developed and used to model animal learning and behaviour, it remains an open question when and whether these different strategies drive animal behaviour. A different line of research has therefore explicitly investigated the balance between model-based and model-free RL in biological agents (Daw et al., 2005; Geerts et al., 2020; Lengyel and Dayan, 2007), where the choice between the two approaches is thought to be guided by some notion of optimality on the basis of available resources and uncertainty about the environment. A popular paradigm for these studies is the so-called 'two-step' task developed by Daw and colleagues (Daw et al., 2011; Momennejad et al., 2017; Wang et al., 2018; although note Akam et al., 2015).

Such work has shown that animals can use both model-free and model-based decision making, with the dorsolateral striatum being particularly important for model-free reinforcement learning (Yin et al., 2004, 2005), and the dorsomedial striatum, prefrontal cortex, and hippocampal formation being important for model-based decision making (Vikbladh et al., 2019; Geerts et al., 2020; Miller et al., 2017; Niv, 2009; Killcross and Coutureau, 2003). This also has interesting parallels to recent work in motor learning, where the basal ganglia were found to be sufficient for 'habitual' motor sequences even in the absence of motor cortex, while motor cortex was necessary for more flexible motor behaviours that are likely to require a high-level 'schema' of the task structure (Mizes et al., 2023b,a). In Section 9 we will see how combining these model-based and model-free ideas with deep learning can lead to human-level performance in tasks such as chess and Go that require long-term planning.

6 | THE SUCCESSOR REPRESENTATION

As we saw in the previous section, an important distinction can be made between model-free reinforcement learning methods, which cache stimulus-response mappings based on prior experience, and model-based reinforcement learning methods, which compute a policy by simulating possible futures using a world model at decision-time. However, we have also noted how animals both need the flexibility of model-based methods as well as the rapid decision making afforded by model-free methods. It has therefore been suggested that animals use intermediate methods that combine some model-free and some model-based features. A particularly prevalent theory has been that of the ‘successor representation’ (SR), which has been proposed to explain both human behaviour (Momennejad et al., 2017) and features of neural activity (Stachenfeld et al., 2017). In particular, the SR allows for flexible adaptation to changing reward functions without having to perform expensive simulations at decision time.

The SR (Dayan, 1993) rewrites the expected reward starting from state s as:

$$V^\pi(s) = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t r_t \mid s_0 = s \right] \quad (21)$$

$$= \sum_{t=0}^{\infty} \gamma^t \sum_{s'} p_\pi(s_t = s' \mid s_0 = s) r(s') \quad (22)$$

$$= \sum_{s'} r(s') \sum_t \gamma^t p_\pi(s_t = s' \mid s_0 = s) \quad (23)$$

$$= \mathbf{r}^T \mathbf{m}_s^\pi. \quad (24)$$

Here, \mathbf{r} is a vector of the average reward associated with each state, and \mathbf{m}_s^π is a vector of the expected discounted future occupancy of state s' if the agent starts in state s and follows the policy π :

$$M_{ss'}^\pi = \sum_{t=0}^{\infty} \gamma^t p_\pi(s_t = s' \mid s_0 = s). \quad (25)$$

The full matrix M^π , constructed from stacking the \mathbf{m}_s^π corresponding to all states s , is denoted the ‘successor matrix’, and it allows us to write down a vector of expected rewards from all states as

$$\mathbf{v}^\pi = M^\pi \mathbf{r}. \quad (26)$$

Here, we have retained the superscript π to indicate that the successor matrix and value function depend on the policy of the agent, which affects the expected occupancy of different states. Having computed the value of each state, we can perform action selection using Equation 11.

The flexibility of the successor representation arises when the reward structure of the environment changes, $\mathbf{r} \rightarrow \mathbf{r}'$. We can now compute the expected reward associated with each state under the new reward function and old policy,

$$\mathbf{v}'^\pi = M^\pi \mathbf{r}'. \quad (27)$$

This provides a better starting point than the old policy *and* reward function (Figure 4C), but the SR does not generalize perfectly since the new value function can lead to policy changes and therefore having to update M (Figure 4D). The successor matrix can be learned by temporal-difference learning when transitioning from s_t , analogous to Equation 9:

$$\Delta M_{s_t s'}^\pi \propto -M_{s_t s'}^\pi + I_{s_t=s'} + \gamma M_{s_{t+1} s'}^\pi. \quad (28)$$

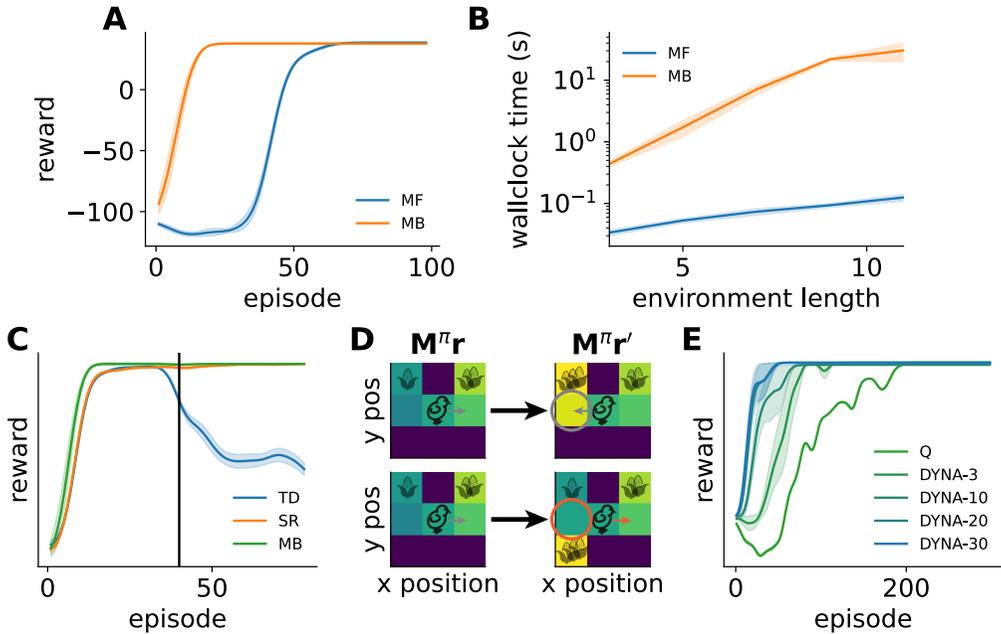


FIGURE 4 Model-based reinforcement learning. (A) Learning curves for model-free (MF) and model-based (MB) RL agents. The MB agent used depth-first search to compute an optimal path at each decision point, gradually learning the reward and transition functions while exploring the environment. The MF agent was a Q-learning agent with $\epsilon = 0$ and learning rate $\alpha = 1$. (B) Wallclock time needed to run 100 episodes of cliffworld with either the MB or MF agents from (A), as a function of the length of the environment. While the MB agent required less experience to learn a good policy, the wallclock time per episode was much larger than for the MF agent. This illustrates an important balance between model-based and model-free reinforcement learning, where MF methods usually require more experience but MB methods require more compute at decision time. (C) Learning curve for an agent using the successor representation (SR) together with learning curves for the model-based agent in (A) and the greedy TD-agent from Figure 2. The goal was moved from location (9, 0) to location (0, 4) at episode 40 (vertical black line), and location (9, 0) was instead given a reward of -5. The MB and SR agents had their reward functions updated to reflect this change and rapidly adapted their policies, while the TD agent had no such mechanism for robustness to changing reward functions. Reward curves were convolved with a Gaussian kernel ($\sigma = 3$ episodes), which is why performance appears to decrease slightly before episode 40. The TD and SR agents were assumed to have access to a 1-step world model at initialization, while the MB agent learned the transition structure from experience. (D) SR agents cannot always adapt to new reward functions if the newly rewarded states have low probability under the old policy. Left column: Value function for an agent that learned an initial policy in an environment with a small reward in the upper left corner and intermediate reward in the upper right corner. The middle top and bottom states are ‘cliffs’. The agent learned to make an initial rightward choice (grey arrows). Right column: A large reward was introduced in either the top left corner (top row) or bottom left corner (bottom row) and the value function recomputed (Equation 27). The agent was unable to adapt to a large reward in the bottom left corner, since the old policy had low probability of reaching this state, even after initially going to the left. This results in a low expected value for going left from the start state (red circle), and a suboptimal policy that continues to go right (red arrow). (E) Learning curve for a standard Q-learning agent (blue) or Dyna agents that perform different numbers of Q-value updates after each physical action (legend). These Dyna updates used cached experience rather than data from a learned world model. Dyna agents make better use of limited experience at the cost of increased compute (proportional to the number of updates).

s'_t is any state, and s_{t+1} is the next state actually observed. Intuitively, transitioning from s_t to s_{t+1} means that (i) we have just been in state s_t , and (ii) we should increase the expected occupancy of all states commonly reached from s_{t+1} (including s_{t+1} itself). Alternatively, if the policy-dependent transition matrix T^π is known, where $T^\pi_{ss'} = p_\pi(s_{t+1} =$

$s' | s_t = s$), the successor matrix can be computed as the geometric series $M^\pi = \mathbf{I} + \gamma \mathbf{T}^\pi + \gamma^2 (\mathbf{T}^\pi)^2 + \dots = (\mathbf{I} - \gamma \mathbf{T}^\pi)^{-1}$.

The SR has been proposed as a model of how humans and other animals learn and generalize (Momennejad et al., 2017; Stachenfeld et al., 2017; Geerts et al., 2020; Gershman, 2018). For example, humans adapt more readily to changes in reward functions than changes in transition functions in a simple sequential binary decision-making task (Momennejad et al., 2017), consistent with the SR facilitating rapid adaptation to changes in r but only slow learning of M . Additionally, hippocampal ‘place cells’ have been proposed to encode a predictive map, with each cell corresponding to a column of M (Stachenfeld et al., 2017). In this model, the firing of a place cell in a given location s reflects the expected future occupancy of its ‘preferred’ location s' conditioned on currently being at s . Stachenfeld et al. (2017) showed that the SR model explains a range of findings in the hippocampal literature. For example, this model explains the asymmetry of directional place fields on a one-dimensional track (Mehta et al., 2000), where s' is more likely to be reached from other states s that precede s' than equidistant states that follow s' . The SR also explains why place fields change near a newly inserted barrier, since two states on either side of the barrier can no longer be visited in quick succession (Alvernhe et al., 2011).

While the SR is perhaps the most prominent model in systems neuroscience that combines features of model-free and model-based RL, it is not the only one. Another interesting algorithm is the ‘Dyna’ architecture of Sutton (1991). In this framework, a model of the world is learned from experience and used to train a model-free policy offline by bootstrapping imagined experience sampled from the model. This allows for more data-efficient learning of model-free policies at the cost of additional compute during ‘rest’, but without needing more compute at decision time (Figure 4E). The model used to simulate data for offline training can either be an explicit learned world model, or it can simply be a memory buffer of past experiences in the form of $(s_t, a_t, r_t, s_{t+1}, a_{t+1})$ tuples. Such experience replay has also proven crucial to the success of modern deep reinforcement learning agents by allowing for higher data efficiency and reducing the instability arising from online experience being autocorrelated (Mnih et al., 2013; Schaul et al., 2015). A prominent theory in neuroscience posits that hippocampal replays could be implementing such a Dyna-like algorithm by generating imagined experience that is used to train the model-free RL systems of the brain (Mattar and Daw, 2018). This theory is supported by the finding that patterns of rodent replay in multiple navigation tasks are consistent with the optimal replays of a Q-learning agent with Dyna, and it has recently been extended to explain not just the content of replays but also their timing (Agrawal et al., 2022).

7 | DEEP REINFORCEMENT LEARNING

We have so far considered small state and action spaces, where tabular policies are tractable. Unfortunately, most ethologically relevant state spaces are large enough that we cannot enumerate all possible states and actions. However, novel situations often resemble previously encountered states, allowing agents to generalize shared structure to these new but related settings (Botvinick et al., 2020). In these cases, we can use *function approximation* (Sutton and Barto, 2018) instead of tabular policies. This involves an assumption that similar states will have similar state-action values and should therefore have similar policies. By making this assumption, we can generalize to unseen states based on previous experience. The use of deep or recurrent neural networks as powerful function approximators for reinforcement learning has driven impressive progress in this setting – the domain of ‘deep reinforcement learning’ (deep RL). Deep RL has seen increasing interest not just in machine learning, but also as a model of neural dynamics and behaviour in humans and other animals (Wang et al., 2018; Jensen et al., 2023; Makino, 2023; Merel et al., 2019; Banino et al., 2018; Aldarondo et al., 2024; Botvinick et al., 2020). Popular approaches in (model-free) deep RL can largely be divided into two categories: value-based methods, which compute state-action values that can be used for action selection; and policy gradient methods, which train a neural network to output a policy directly.

Value-based methods

The deep RL approaches most similar to the tabular methods considered in Section 3 and Section 4 use neural networks to compute state-action values, which can be used for action selection as we saw in Equation 11. However, by using function approximation instead of the tabular values considered previously, these networks can generalize to unseen states in large state spaces. This gives rise to the family of ‘deep Q-learning’ methods, which closely mirror the tabular Q-learning considered previously, but now with function approximation.

The simplest approach involves defining a state-action value function $Q_\theta(s, a)$, where the parameters θ of the neural network defining our agent are learned as follows:

- Collect experience (s_t, a_t, r_t, s_{t+1}) .
- Define a loss $\mathcal{L} = 0.5[Q_\theta(s_t, a_t) - (r_t + \gamma \max_a Q_\theta(s_{t+1}, a))]^2$.
- Update parameters $\Delta\theta \propto -\frac{\partial \mathcal{L}}{\partial \theta}$ [often treating the ‘target value’ $y_t = (r_t + \gamma \max_a Q_\theta(s_{t+1}, a))$ as constant w.r.t. θ].

When acting according to our policy, we simply pick the action predicted to have the highest value, usually using some variant of ϵ -greedy or softmax to increase exploration.

On the surface, this looks like a straightforward generalization of tabular Q-learning, and it may seem surprising that deep Q-learning did not see significant use or success until the foundational work of Mnih et al. (2013). However, a major difficulty arises from the autocorrelation of the states observed by the agent, which destabilizes training. This can be mitigated by the use of ‘experience replay’, where the experience generated by the agent is added to a global replay buffer \mathcal{B} . One or more experiences are then sampled randomly from the buffer at each iteration and used to update the network parameters – reminiscent of the ‘Dyna’ architecture described previously. Additionally, the target value $y_t = (r_t + \gamma \max_a Q_\theta(s_{t+1}, a))$ itself depends on θ and therefore changes when any Q-value is updated (in contrast to tabular Q-learning, where there is no parameter sharing). It is therefore common to use a ‘target network’ $Q_{\theta'}$ that remains fixed for multiple rounds of data collection and parameter updates. This reduces fluctuations in the target values, and the resulting parameter updates are gradients of a well-defined objective function. Together, these two approaches give rise to the ‘deep Q network’ (DQN) developed by Mnih et al. (2013), which is trained as follows:

- Collect experience (s_t, a_t, r_t, s_{t+1}) and add to \mathcal{B} [optionally many iterations and optionally removing stale experiences].
- Randomly sample an experience $(s'_t, a'_t, r'_t, s'_{t+1}) \sim \mathcal{B}$ [optionally a full batch].
- Define a loss $\mathcal{L}(\theta) = 0.5[Q_\theta(s'_t, a'_t) - (r'_t + \gamma \max_a Q_{\theta'}(s'_{t+1}, a))]^2$ [optionally averaged over the full batch]. Note the ‘student network’ has parameters θ and the target network inside the max has parameters θ' .
- Update the network parameters $\Delta\theta \propto -\frac{\partial \mathcal{L}(\theta)}{\partial \theta}$.
- At regular intervals, set our target network to the student network, $\theta' \leftarrow \theta$.

This algorithm is effectively off-policy, since most of the data in \mathcal{B} is collected by a policy defined by an old set of parameters – and the data in \mathcal{B} can in fact be generated completely independently of the agent being trained. Even though the DQN is more stable than naive deep Q-learning, an additional instability arises from the fact that $Q_{\max}(s'_{t+1}) = \max_a Q_{\theta'}(s'_{t+1}, a)$ uses the same Q values both to estimate which action is best and what the value of that action is, which leads to a positively biased estimate. This can be mitigated by ‘double Q-learning’ (Van Hasselt et al., 2016), where the student network selects the best action and the target network evaluates its value, $Q_{\max}(s'_{t+1}) \leftarrow Q_{\theta'}(s'_{t+1}, \operatorname{argmax}_a(Q_\theta(s'_{t+1}, a)))$.

While modern deep Q-learning has reached impressive performance across a range of machine learning settings (Mnih et al., 2013; Lillicrap et al., 2015; Schaul et al., 2015; Kalashnikov et al., 2018), it is unclear whether the various modifications needed to stabilize the algorithm could be implemented in biological circuits. This is perhaps the reason why

neuroscience research using deep Q-learning has been relatively scarce, despite the prevalence of tabular Q-learning in theoretical neuroscience. An interesting exception is recent work by Makino (2023), which shows parallels between the values learned by a DQN and neural representations in mammalian cortex during a compositional behavioural task. Additionally, the importance of experience replay in DQNs (Mnih et al., 2013; Schaul et al., 2015) has close parallels to the proposal that hippocampal replay constitutes a form of experience replay (Matar and Daw, 2018).

Policy gradient methods

A conceptually simpler approach for deep reinforcement learning uses policy gradient methods (Sutton and Barto, 2018), where a neural network with parameters θ takes as input the (observable) state of the environment and directly outputs a policy π_θ . This has also found more support and use in the neuroscience literature, where policy gradient methods have recently been used as models of learning and neural dynamics in the biological brain (Wang et al., 2018; Jensen et al., 2023; Merel et al., 2019; Song et al., 2017).

The objective in a policy gradient network is to find the setting of θ that maximizes expected reward. A naive way to achieve this would be to define $R_\tau := \sum_{t=0}^T \gamma^t r_t$ and compute gradients given by

$$\nabla_\theta J(\theta) = \nabla_\theta \mathbb{E}_{\tau \sim p_{\pi_\theta}(\tau)} [R_\tau] \quad (29)$$

$$= \sum_{\tau} R_\tau \nabla_\theta p_{\pi_\theta}(\tau). \quad (30)$$

Here, $\tau \sim p_{\pi_\theta}(\tau)$ indicates trajectories sampled from the distribution induced by the policy π_θ , and $J(\theta)$ indicates the expectation of R_τ under $p_{\pi_\theta}(\tau)$ (c.f. Equation 2). However, evaluating Equation 30 requires us to know how the environment will respond to our actions, which in general may not be the case. Instead, we use the ‘log-derivative trick’, which takes advantage of the linearity of the expectation and the identity $\nabla_\theta \log f(\theta) = f(\theta)^{-1} \nabla_\theta f(\theta)$ to write

$$\nabla_\theta J(\theta) = \sum_{\tau} R_\tau \nabla_\theta p_{\pi_\theta}(\tau) \quad (31)$$

$$= \sum_{\tau} R_\tau p_{\pi_\theta}(\tau) \nabla_\theta \log p_{\pi_\theta}(\tau) \quad (32)$$

$$= \mathbb{E}_{\tau \sim p_{\pi_\theta}(\tau)} [R_\tau \nabla_\theta \log p_{\pi_\theta}(\tau)], \quad (33)$$

Since the environment does not depend on θ , we can simplify the calculation of $\nabla_\theta \log p_{\pi_\theta}(\tau)$:

$$\nabla_\theta \log p_{\pi_\theta}(\tau) = \nabla_\theta \left[\log p(s_0) + \sum_{t=0}^T \log p(s_{t+1} | s_t, a_t) + \log \pi_\theta(a_t | s_t) \right] \quad (34)$$

$$= \sum_{t=0}^T \nabla_\theta \log \pi_\theta(a_t | s_t). \quad (35)$$

Inserting Equation 34 in Equation 31, we arrive at the REINFORCE algorithm (Williams, 1992):

$$\nabla_\theta J(\theta) = \mathbb{E}_{\tau \sim p_{\pi_\theta}(\tau)} \left[R_\tau \sum_{t=0}^T \nabla_\theta \log \pi_\theta(a_t | s_t) \right] \quad (36)$$

$$\approx \frac{1}{N} \sum_{\tau \sim p_{\pi_\theta}(\tau)} \left(\sum_{t=0}^T \gamma^t r_t \right) \left(\sum_{t=0}^T \nabla_\theta \log \pi_\theta(a_t | s_t) \right), \quad (37)$$

where the second line approximates the expectation with N empirical rollouts of the policy in the environment. In-

tuitively, Equation 37 says that we should preferentially upregulate the probability of trajectories with high reward. Importantly, it no longer differentiates through the environment – only the policy.

While the REINFORCE algorithm is unbiased, it also has high variance, which can make learning slow and unstable. It is therefore common to introduce modifications that reduce the variance. The first of these comes from noting that an action taken at time t cannot affect the reward at times $t' < t$. We therefore define $R_t := \sum_{t'=t}^T \gamma^{t'-t} r_{t'}$ and write a new update rule as

$$\hat{\nabla}_{\theta} J(\theta) \approx \frac{1}{N} \sum_{\tau \sim p_{\pi_{\theta}}(\tau)} \sum_{t=0}^T R_t \nabla_{\theta} \log \pi_{\theta}(a_t | s_t). \quad (38)$$

This is the formulation most commonly used in the literature, but it is not actually the same as Equation 37, which would use $R_t = \sum_{t'=t}^T \gamma^{t'-0} r_{t'}$. As briefly discussed in Section 2, this is because the discount factor γ is generally used as a variance reduction method rather than because we intrinsically care less about rewards later in the task. In fact, Equation 38 is not strictly speaking a gradient (Nota and Thomas, 2019), which is why we denote it $\hat{\nabla}$.

It can also be shown that subtracting an action-independent baseline from R_t does not change the expectation in Equation 38, while potentially reducing its variance. A common choice is the expected future reward $V(s_t)$:

$$\hat{\nabla}_{\theta} J(\theta) \approx \frac{1}{N} \sum_{\tau \sim p_{\pi_{\theta}}(\tau)} \sum_{t=0}^T (R_t - V(s_t)) \nabla_{\theta} \log \pi_{\theta}(a_t | s_t). \quad (39)$$

Intuitively, Equation 39 upregulates the probability of actions that lead to higher-than-expected reward and downregulates the probability of actions that lead to lower-than-expected reward.

Finally, it is common to reduce the variance of the gradient estimate further through an approach known as ‘bootstrapping’, which approximates $R_t \approx r_t + \gamma V(s_{t+1})$. This is useful because $r_t + \gamma V(s_{t+1})$ has lower variance than $\sum_{t'=t}^T \gamma^{t'-t} r_{t'}$. We therefore replace $(R_t - V(s_t))$ in Equation 39 with the ‘advantage function’ $A(s_t, a_t) = Q(s_t, a_t) - V(s_t) \approx r_t + \gamma V(s_{t+1}) - V(s_t)$. In between these two extreme cases of a full Monte Carlo estimate of R_t and a ‘one-step’ bootstrap, the sum in R_t can be truncated to any order, with $R_{t'}$ replaced by $V(s_{t'})$ (Sutton and Barto, 2018). In theory, this gradient estimate remains unbiased if the value function is correct. In practice, the learned estimate of $V(s_{t'})$ will be inexact, which biases the parameter updates. Bootstrapping therefore leads to a tradeoff between the bias and variance of parameter updates.

These variance reduction approaches give rise to the so-called ‘actor-critic’ algorithm, where an agent both computes a policy π (the actor) and an ‘evaluation’ of the policy in the form of state(-action) values (the critic). A rich neuroscience literature suggests that the basal ganglia of biological agents implement an actor-critic-like algorithm. Here, dorsal striatum is proposed to implement the ‘actor’ and ventral striatum the ‘critic’ (Takahashi et al., 2008; Sutton and Barto, 2018; O’Doherty et al., 2004).

For these actor-critic algorithms, it is common to parameterize both the policy $\pi_{\theta}(a|s)$ and value function $V_{\theta}(s)$ with neural networks. To optimize these parameters using out-of-the-box automatic differentiation, we need to write down an ‘objective function’ with the correct gradients – but we saw in Equation 30 that this cannot simply be the expected reward. Instead, we define an auxiliary utility (i.e. negative loss)

$$\tilde{J}(\theta) = \frac{1}{N} \sum_{\tau \sim p_{\pi_{\theta}}(\tau)} \sum_{t=0}^T (R_t - V(s_t)) \log \pi_{\theta}(a_t | s_t), \quad (40)$$

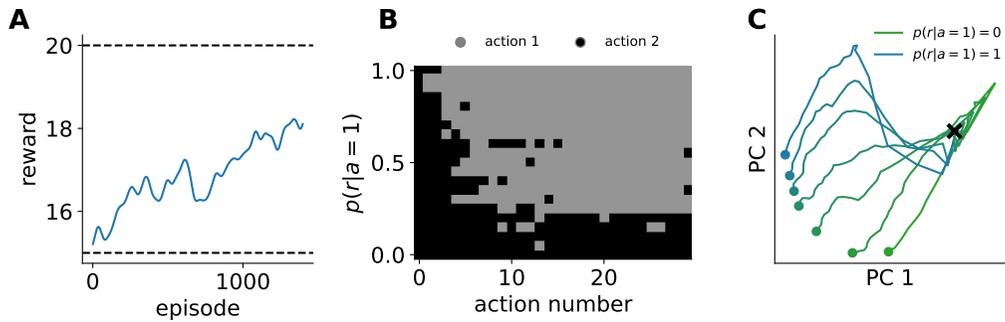


FIGURE 5 Meta-reinforcement learning. The results in this figure reproduce some of the analyses in Figure 1 of Wang et al. (2018). **(A)** We trained a recurrent meta-reinforcement learning agent in a two-armed bandit task, where the reward probabilities of each arm were sampled independently from $\mathcal{U}(0, 1)$ at the beginning of each episode and remained fixed throughout the episode. A recurrent neural network was trained across many episodes with *different reward probabilities* using an actor-critic algorithm. The input to the agent consisted of the previous action, the previous reward, and the time-within-trial. The average reward per episode is plotted against the episode number, showing that the agent gradually learns to adapt within each episode to the particular instantiation of the bandit task. Importantly, the parameters of the network are fixed within an episode, meaning that this adaptation occurs through the recurrent dynamics. Dashed horizontal lines indicate the reward of an agent selecting random actions and an ‘oracle’ agent that always chooses the best arm. **(B)** Heatmap showing example behaviour of the agent in episodes with different reward probabilities for the first arm, $p(r|a = 1)$. For the analysis here and in (C), we set $p(r|a = 2) = 1 - p(r|a = 1)$. Across episodes, the agent experiments with different actions and eventually converges on the optimal action. For episodes with more similar reward probabilities (near the middle), it takes longer to identify the optimal action. This balance between exploration and exploitation is mediated by the recurrent network dynamics, which are learned over many episodes using deep reinforcement learning. **(C)** We averaged the hidden state of the RNN over 100 episodes for each of several different reward probabilities, ranging from low (green) to high (blue) $p(r|a = 1)$. We then performed PCA on the resulting matrix of average hidden states to compute a low-dimensional trajectory over the course of an episode for each reward probability. This two-dimensional embedding of neural activity converges to different regions of state space during the episode for different reward probabilities. Black cross indicates the hidden state at the beginning of an episode, and coloured points indicate the final hidden state in an episode for the different reward probabilities.

where R_t can optionally be approximated by $r_t + \gamma V(s_{t+1})$. While $\tilde{J}(\theta)$ has no intrinsic interpretation, it is chosen such that $\nabla_{\theta} \tilde{J}(\theta) = \tilde{\nabla}_{\theta} J(\theta)$ when treating $\delta_t := (R_t - V(s_t))$ as constant w.r.t θ , and the gradients can be computed using standard automatic differentiation. The gradient of the value function loss is then given by $\nabla_{\theta} \sum_t \frac{1}{2} (R_t - V_{\theta}(s_t))^2 = \sum_t [-\delta_t \nabla_{\theta} V_{\theta}(s_t)]$.

While these policy gradient methods may seem far removed from neuroscience, it has been found that neural networks trained with policy gradients often learn representations and behaviours reminiscent of biological organisms (Wang et al., 2018; Jensen et al., 2023; Merel et al., 2019; Li et al., 2022; Song et al., 2017).

Meta-reinforcement learning

A prominent example of deep reinforcement learning providing insights into biological circuits is the ‘recurrent meta-reinforcement learning’ model developed by Wang et al. (2018). The authors trained a *recurrent* deep RL agent using policy gradients, where the RNN parameters are configured by learning from rewards over long periods of time from many tasks that have a shared underlying structure. Importantly, this ‘slow’ model-free learning process gives rise to an agent that can rapidly learn from experience with *fixed parameters* when exposed to a new task from the same task distribution. This is achieved by the agent learning to effectively implement a fast RL-like algorithm in the *dynamics* of the network (Figure 5). This process, whereby an agent trained slowly on a large distribution of tasks can rapidly

adapt to a new task, is known as ‘meta-reinforcement learning’ (Finn et al., 2017; Ritter et al., 2018; Duan et al., 2016; Wang et al., 2016). Wang et al. (2018) suggested that prefrontal cortex resembles such a recurrent meta-RL system, and their model explained a range of neuroscientific findings. This included

- Dynamic adaptation of the effective learning rate of an agent to the volatility of the environment (Behrens et al., 2007).
- The emergence of ‘model-based’ behaviour in the ‘two-step’ task commonly used to distinguish between model-free and model-based RL (Miller et al., 2017; Daw et al., 2011).
- The ability of animals to get progressively faster at learning when exposed to multiple tasks with a consistent abstract task structure (Harlow, 1949).

Further experimental evidence for this meta-RL framework comes from Hattori et al. (2023), who showed that across-session learning in a reversal learning task relied on synaptic plasticity in orbitofrontal cortex (OFC; a subregion of PFC), while within-session learning relied on recurrent dynamics in OFC. Recently, Jensen et al. (2023) also extended the work of Wang et al. (2018) to allow the meta-RL network dynamics to update the policy from *imagined* experience using a learned model of the environment – reminiscent of Dyna, but now implemented in RNN dynamics instead of parameter updates.

8 | DISTRIBUTIONAL REINFORCEMENT LEARNING

In Equation 3 and Equation 10, we defined the expected future reward for a given state or state-action pair. The methods considered so far have only used such expectations as a learning signal. However, recent research suggests that animals may in fact estimate entire future reward distributions (Dabney et al., 2020; Sousa et al., 2023). These studies were inspired by findings that such *distributional RL* can improve the performance of artificial agents (Bellemare et al., 2017, 2023; Dabney et al., 2018). To formalize this, we use $Z^\pi(s, a)$ to denote a single sample from the distribution over possible cumulative discounted future rewards resulting from following policy π after taking action a in state s :

$$Z^\pi(s, a) \sim p_{Z^\pi} \left(\sum_{t'=t}^{\infty} \gamma^{t'-t} r_{t'} \mid s_t = s, a_t = a \right) \quad (41)$$

The stochasticity of Z^π can both be due to stochasticity in environment dynamics and reward, and it can be due to stochasticity in the policy of the agent itself. Clearly, the expectation of $Z^\pi(s, a)$ equals the corresponding Q value:

$$\mathbb{E}_{p_{Z^\pi}} [Z^\pi(s, a)] = Q^\pi(s, a). \quad (42)$$

Instead of only estimating this expectation, we now want to learn the full distribution of returns, $p_{Z^\pi}(Z^\pi(s, a))$. One normative reason to learn this distribution is to develop methods that are risk averse (Morimura et al., 2010, 2012) or explicitly take into account uncertainty (Dearden et al., 1998). However, recent work has suggested that such a distributional approach can also increase expected reward by improving representation learning in the deep RL setting (Bellemare et al., 2017; Dabney et al., 2018; Rowland et al., 2019; Bellemare et al., 2023). This is because traditional deep RL only distinguishes states that have different expected value, while distributional RL learns to distinguish states that have different value distributions (Figure 6A).

To implement distributional RL, we consider the τ^{th} *expectile* of $p_{Z^\pi}(Z^\pi(s, a))$, e_τ , which is defined for a random variable Z as the solution to the equation

$$\tau \mathbb{E}[\max(0, Z - e_\tau)] = (1 - \tau) \mathbb{E}[\max(0, e_\tau - Z)]. \quad (43)$$

This is a generalization of the mean, which is recovered for $\tau = 0.5$, similar to how the quantile generalizes the notion of a median. A distribution is uniquely specified by its expectiles, and we can therefore represent $p_{Z^\pi}(\mathcal{Z}^\pi(s, a))$ in terms of $\{\epsilon_\tau\}$. Translating this to an RL algorithm involves training a network (or tabular values) to predict a set of expectiles for a given state (and action). The parameters of the agent are then updated by propagating the distribution implied by the predicted expectiles through the Bellman equation and minimizing the difference between the initial and propagated distributions (Bellemare et al., 2023).

In the tabular value learning case (c.f. Section 3), this can be achieved using a modified TD-update rule (c.f. Equation 9; Lowet et al., 2020; Figure 6B). In particular, we consider a set of units $\{V_{\tau_i}(s)\}$, each with a target expectile $\tau_i := \frac{\alpha_i^+}{\alpha_i^+ + \alpha_i^-}$ of the return distribution $p_{Z^\pi}(\mathcal{Z}^\pi(s))$ (where $p_{Z^\pi}(\mathcal{Z}^\pi(s)) = \sum_a [\pi(a|s)p_{Z^\pi}(\mathcal{Z}^\pi(s))]$). These expectiles can be learned by sampling experience from the environment under the policy and defining a TD error for each unit and state transition as

$$\delta_i := r_t + \gamma \tilde{Z}_j^\pi(s_{t+1}) - V_{\tau_i}(s_t). \quad (44)$$

Here, $\tilde{Z}_j^\pi(s_{t+1})$ is a *random sample* from the learned approximate distribution of cumulative future returns from state s_{t+1} , $p_{\{V_{\tau_i}\}}(\tilde{Z}^\pi(s_{t+1}))$ (Lowet et al., 2020; Dabney et al., 2020). We then apply the following update rule to all units:

$$\Delta V_{\tau_i}(s_t) = \alpha_i^+ \max(\delta_i, 0) + \alpha_i^- \min(\delta_i, 0). \quad (45)$$

In other words, we apply the TD update rule to each unit with learning rate α_i^+ for positive TD errors and learning rate α_i^- for negative TD errors. When running this algorithm to convergence, $V_{\tau_i}(s)$ will approach the τ_i^{th} expectile (ϵ_{τ_i}) of $p_{Z^\pi}(\mathcal{Z}^\pi(s, a))$. In the deep RL setting, we would compute gradients of the form $\nabla_\theta \mathcal{L} = \sum_i \Delta V_{\tau_i}(s_t) \partial V_{\tau_i}(s_t) / \partial \theta$ to learn a model with parameters θ that predicts the full set of expectiles. We refer to Bellemare et al. (2017); Dabney et al. (2018); Rowland et al. (2019); Bellemare et al. (2023); and Dabney et al. (2020) for additional mathematical details, alternative parameterizations of the return distribution, and extensions to the control setting.

Intriguingly, recent work in neuroscience suggests that distributional RL could underlie value learning in biological neural circuits (Dabney et al., 2020; Lowet et al., 2020; Sousa et al., 2023). In particular, Dabney et al. (2020) recorded the activity of dopaminergic VTA neurons during a task with stochastic rewards and showed that the neurons appeared to represent the full distribution of possible outcomes using an expectile representation. More concretely, they showed that:

- The VTA neurons exhibited a range of different ‘reversal points’ – defined as the reward magnitude at which the firing rate of a neuron did not change from its baseline firing rate. This is consistent with a distributional RL theory, where the changes in neural firing rates from baseline correspond to the expectile TD updates considered above. In this case, the reversal point of a neuron i should be $V_{\tau_i} \approx \epsilon_{\tau_i}$ (Figure 6C).
- Neurons had different slopes describing the relationship between expected reward and firing rate in the regimes where expected reward was above and below the reversal point (V_{τ_i}) of each neuron. This is consistent with the algorithm described in Equation 45 (Figure 6C).
- When independently fitting a slope to the data above (α_i^+) and below (α_i^-) the reversal point of neuron i , the reversal point of the neuron was positively correlated with $\tau_i = \frac{\alpha_i^+}{\alpha_i^- + \alpha_i^+}$. This is consistent with the expectile distributional RL setting, where the reversal point is $V_{\tau_i} \approx \epsilon_{\tau_i}$, which is monotonic in τ_i (Figure 6C).
- When ‘imputing’ the distribution implied by the VTA neurons when interpreted as expectiles (Figure 6D), the resulting fit resembled the true distribution of rewards in the experiment.

These findings generalize the canonical RPE view of Schultz et al. (1997), which can be seen as the averaged version

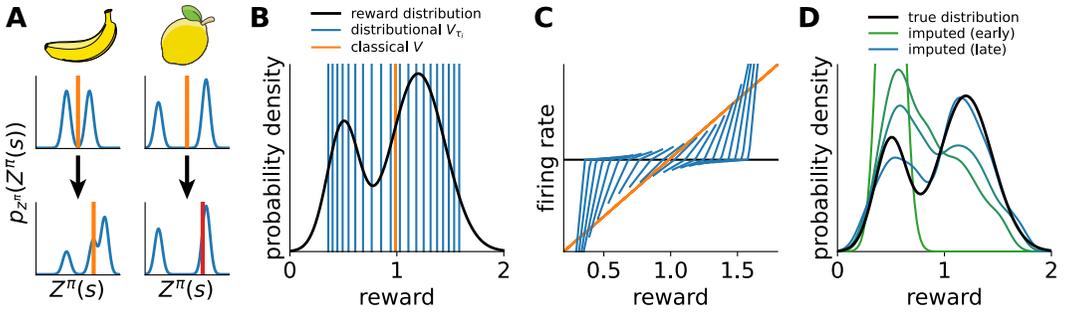


FIGURE 6 Distributional reinforcement learning. **(A)** Example of distributional RL improving representation learning. Two states (e.g. ‘banana’ and ‘lemon’) may have the same expected future reward but different reward *distributions* (top row). A standard RL agent only has to predict the mean (orange vertical lines) and may learn a simple predictive feature like ‘yellow’ that combines both states. If the expected future reward now changes for one state (bottom row), e.g. because the agent learned to make a banana smoothie, it may erroneously generalize to *all* yellow fruits (bottom right; red line). However, a distributional agent is forced to learn an initial representation that distinguishes the states, which can improve downstream learning and prevent overgeneralization (blue curves). **(B)** Distributional RL simulations on a value estimation task with a single state and no actions, reproducing key ideas from [Dabney et al. \(2020\)](#). Distribution of rewards (black), plotted together with the learned values V_{τ_i} across 20 units learning through standard TD learning (orange; all $\tau_i = 0.5$), or through distributional RL with different $\tau_i = \frac{\alpha_i^+}{\alpha_i^+ + \alpha_i^-}$ (blue). The TD units all converge to the mean reward, while the expectile units end up tiling the distribution. **(C)** For both the TD units and distributional units from (B), we plot the temporal difference updates performed in response to different rewards from the reward distribution. These updates have been proposed to be represented in the firing rates of dopaminergic VTA neurons relative to baseline ([Schultz et al., 1997](#); [Dabney et al., 2020](#)). The TD units show a constant linear scaling across positive and negative rewards, while the distributional units show an asymmetric scaling of firing rate with reward magnitude above and below their reversal point (black horizontal line). The ratio of slopes above and below the reversal point scales positively with the value of the reversal point. These features of dopaminergic VTA neurons were used by [Dabney et al. \(2020\)](#) to argue that the brain implements a form of distributional RL. **(D)** True reward distribution (black) reproduced from (B), now plotted together with the reward distribution $p_{\{V_{\tau_i}\}}(\tilde{Z}^\pi(s))$ implied by the distributional units from (B) and (C) at different stages of learning (green to blue). These imputed distributions were computed by *assuming* that the value V_{τ_i} learned by unit i corresponds to expectile $\tau_i = \frac{\alpha_i^+}{\alpha_i^+ + \alpha_i^-}$ of the reward distribution. We then infer the distribution implied by these expectiles under the assumption that it consists of a set of 20 delta functions ([Rowland et al., 2019](#)). Finally, we convolved the resulting delta functions with a Gaussian kernel ($\sigma = 0.1$) for visualization. This process was repeated using $\{V_{\tau_i}\}$ at different stages of learning. The units were all initialized at $V_{\tau_i} = 0.5$, so the initial distribution is a delta function at $r = 0.5$ (green). At the end of learning, the population faithfully represents the true reward distribution, capturing key features including bimodality and the relative magnitude of the two modes (blue). [Dabney et al. \(2020\)](#) used a similar approach to infer the distribution encoded by dopaminergic VTA neurons at the end of animal training and found a close match to the true reward distribution.

of the theory by [Dabney et al. \(2020\)](#). The expectile regression algorithm investigated by [Dabney et al. \(2020\)](#) relies on non-local TD updates and non-linear ‘imputation’ of the return distribution $p_{\{V_{\tau_i}\}}(\tilde{Z}^\pi(s))$ induced by the learned expectiles $\{V_{\tau_i}(s)\}$. However, recent work has suggested more biologically plausible distributional RL updates to address these challenges ([Tano et al., 2020](#); [Wenliang et al., 2023](#)).

9 | LEARNING FROM SCALAR REWARDS?

The methods discussed so far have all relied on some scalar ‘reward’ available from the environment. Distributional RL predicted the full distribution of possible rewards but still assumed that an external reward was eventually observed. This is perhaps reasonable in many experimentally controlled settings, where we deliver juice or water to animals as a

function of their actions. It is also often appropriate in machine learning, where we want to optimize some externally imposed objective that can be quantified. However, in natural environments, there is no such 'global' reward – instead, different actions can be rewarded in different ways that must be converted to an internal learning signal (Juechems and Summerfield, 2019). For example, foraging for food might reduce hunger, while collecting water might reduce thirst. Starting a family might have an initial energy cost, but it could improve survival later in life and propagate the gene pool. An RL purist would convert all of these different gains and losses into a single common unit and balance them appropriately to decide what actions to take (Silver et al., 2021). However, it has also been suggested that biological agents adaptively change their instantaneous objective via a higher-order controller in e.g. prefrontal cortex, which determines what the current lower-level objective is (Juechems and Summerfield, 2019; Miller and Cohen, 2001; Botvinick, 2008). This resembles ideas in hierarchical reinforcement learning (Pateria et al., 2021; Botvinick, 2008), but here the higher-order policies are not necessarily learned through reinforcement learning within a lifetime, instead emerging during evolution where survival depends on the discovery of 'useful' objectives.

A related challenge is the technical difficulty of learning from scalar rewards, which are generally sparse. While the methods in Section 7 can train a neural network to maximize reward in theory, in practice they are often noisy, unstable, or find local minima. As a topical example, consider the challenge of training large language models to interact with human users. Training such a model from scratch using reinforcement learning would be near-impossible, but 'pretraining' the model on a large unsupervised dataset followed by 'reinforcement learning from human feedback' has proven hugely successful and revolutionized language models for human interactions (Team et al., 2023). This works because the large unlabelled dataset provides a way to learn good *representations* that distinguish different concepts, since this is necessary to solve the base task of predicting the next word. Once such representations are learned, the subsequent finetuning for human interaction is an easier problem that can be solved with reinforcement learning.

Of more relevance to neuroscience, it is also common in deep RL to introduce additional 'auxiliary costs' to the utility function that are jointly optimized together with reward maximization, and which can use a richer data source to drive representation learning in the network (Jaderberg et al., 2016). A popular approach is to include losses that require the agent to predict the next observation from the current state and action (Jaderberg et al., 2016; Zintgraf et al., 2019). Such predictive losses have close parallels in neuroscience, where it has been suggested that predictive learning could drive many of the representations observed in biological circuits (Rao and Ballard, 1999; Stachenfeld et al., 2017; Whittington et al., 2020; Blanco-Pozo et al., 2021) and serve as a foundation for model-based planning (Jensen et al., 2023). Fang and Stachenfeld (2023) also recently showed that augmenting an RL agent with an auxiliary predictive objective leads to neural representations that resemble biological circuits more closely. This suggests a potentially important interaction between self-supervised representation learning and reward-driven reinforcement learning in biological circuits.

If a good predictive model of the world is learned, this model can also be used together with a search algorithm to turn the reinforcement learning problem into a supervised learning problem. As an example, the MuZero model developed by Schrittwieser et al. (2020) to learn Atari, chess, shogi, and Go, was trained to predict future values and actions by unrolling the environment using a learned latent representation. It then used Monte Carlo tree search (MCTS; a model based-search algorithm c.f. Section 5) and the predicted value function to improve the provisional actions predicted by the base policy network. These final actions, which had been optimized using MCTS, were treated as a supervised learning target for the base policy network. Through many iterations of such predictive learning, MCTS-based policy improvement, and training of the base policy, the network learned both the transition structure of the task and to select good actions. Such semi-supervised approaches to reinforcement learning are useful in the common setting where learning the transition structure of the world is easier than learning a policy (Jensen et al., 2023). These methods also have interesting parallels to learning in biological networks, where interactions between model-based and model-

free systems are similarly thought to drive action selection based on learned latent representations (Botvinick et al., 2020).

Finally, algorithms that seek to maximize scalar rewards often run into challenges related to exploration. In particular, once an above-average policy has been identified, exploration is disfavored because it has lower expected reward, even if there is a finite probability of discovering new and better policies. This is why imposing stochasticity in the form of e.g. ϵ -greedy policies is common in the tabular RL literature. In deep RL, it is instead popular to include various forms of ‘exploration bonuses’ in the objective function. Policy gradient algorithms for example often add an auxiliary entropy loss of the form $\mathcal{L}_E = \sum_{\tau \sim p_{\pi_{\theta}}(\tau)} \sum_{t=0}^T \sum_a \pi_{\theta}(a_t | s_t) \log \pi_{\theta}(a_t | s_t)$, and indeed the bandit example in Figure 5 does not work without such an entropy loss. Other approaches to improve exploration include hierarchical reinforcement learning, which introduces autocorrelations in the policy to improve coverage of the state space. This is reminiscent of the biological ‘Lévy flight’ hypothesis, which suggests that animals explore an environment using a heavy-tailed distribution of ‘step sizes’ to maximize the probability of finding sparse rewards (Viswanathan et al., 1999). Biological agents are also thought to engage in periods of ‘random exploration’ that can be interpreted as a biological parallel to ϵ -greedy-like algorithms. Random exploration appears to be under noradrenergic control (Tervo et al., 2014; Dubois et al., 2021), while the zona incerta can drive more directed exploration (Ahmadlou et al., 2021). Such neural control of exploration is particularly well-characterized in the zebra finch song circuit (Ölveczky et al., 2005), and it is consistent with a view of top-down imposition of flexible reward functions – where one possible objective could be to increase variability or reduce uncertainty – rather than optimization of a global scalar reward.

10 | DISCUSSION

In this review, we have provided a mathematical overview of some of the many reinforcement learning methods that are commonly used in systems and computational neuroscience. We have also highlighted a range of explicit parallels between these methods and experimental results in neuroscience and cognitive science to illustrate the utility of reinforcement learning as a framework for understanding biological learning and decision making. This has ranged from classical work on reward prediction errors (Schultz et al., 1997) to recent findings of distributional reinforcement learning in biological circuits (Dabney et al., 2020).

While RL has already had a profound influence on systems neuroscience, several open questions remain. In particular, much work has focused on simple stimulus-response or binary decision making tasks. This is a far reach from ethologically relevant problems that involve processing multimodal stimuli, decision making with long-lasting consequences, and high-dimensional motor control. Some recent work building on deep RL has started to bridge this gap. For example, Banino et al. (2018) demonstrated the emergence of grid cells in agents navigating complex environments, Aldarondo et al. (2024) showed similarities between an RL agent trained to control a virtual rodent and corresponding biological motor representations, and Jensen et al. (2023) showed parallels between a recurrent meta-RL agent and human behaviour in a navigation task requiring temporally extended thinking. However, much work remains to extend our neuroscientific understanding to ethologically relevant settings, both experimental and computational.

A related challenge will be to combine different components of existing models to capture the generalist nature of biological circuits. This is in contrast to past work, which has often focused on a single neural circuit or function, such as motor control or navigation. Such a generalist approach will involve explicit modeling of the roles of different brain regions, and more importantly it will require us to capture how they interact with one another during learning and decision making. Clearly, such models will need to be constrained by experimental data, both at the level of behaviour and at the level of neural activity. This is becoming increasingly feasible with recent advances in recording technologies, both for high-resolution behavioural tracking (Mathis et al., 2018; Dunn et al., 2021) and for simultaneous and long-term recording of neural activity (Steinmetz et al., 2021; Pachitariu et al., 2016; Dhawale et al., 2017).

Finally, most work on reinforcement learning in neuroscience has considered short-term decision making tasks, where planning and decision making in primitive state and action spaces are feasible. This is in stark contrast to most human decision making, which occurs over extended timescales and often involves hierarchies of decisions or priorities that change over time. For example, we may decide to pursue an undergraduate degree at Cambridge University, which then requires us to (i) write an application, (ii) prepare for an interview, and (iii) arrange our travel. Each of these processes in turn require us to plan increasingly low-level decisions, such as booking a flight or deciding which bus to take to the airport. This is the topic of hierarchical reinforcement learning, which has already been highlighted as a potentially useful model of human behaviour (Eckstein and Collins, 2020; Botvinick, 2008; Botvinick et al., 2009) and is becoming an increasingly important area of research in machine learning (Pateria et al., 2021). While reinforcement learning will undoubtedly remain important for understanding such flexible animal behaviour, we should also keep in mind the shortcomings and challenges of learning from scalar rewards (Section 9). It will be important to investigate the regimes in which reinforcement learning provides a good model of behaviour and neural dynamics, while also exploring other frameworks with potentially richer learning substrates. This will also enable the study of interactions between such different learning algorithms, which will likely be necessary to understand biological learning in rich environments.

Acknowledgements

Guillaume Hennequin, Will Dorrell, and the NBDT reviewers and editor provided valuable feedback that improved and clarified the manuscript. Most of the content in this review was originally prepared for the 2023 Janelia Theoretical Neuroscience Workshop.

References

- Agrawal, M., Mattar, M. G., Cohen, J. D. and Daw, N. D. (2022) The temporal dynamics of opportunity costs: A normative account of cognitive fatigue and boredom. *Psychological review*, **129**, 564.
- Ahmadlou, M., Houba, J. H., van Vierbergen, J. F., Giannouli, M., Gimenez, G.-A., van Weeghel, C., Darbanfouladi, M., Shirazi, M. Y., Dziubek, J., Kacem, M. et al. (2021) A cell type-specific cortico-subcortical brain circuit for investigatory and novelty-seeking behavior. *Science*, **372**, eabe9681.
- Akam, T., Costa, R. and Dayan, P. (2015) Simple plans or sophisticated habits? state, transition and learning interactions in the two-step task. *PLoS computational biology*, **11**, e1004648.
- Aldarondo, D., Merel, J., Marshall, J. D., Hasenclever, L., Klibaite, U., Gellis, A., Tassa, Y., Wayne, G., Botvinick, M. and Ölveczky, B. P. (2024) A virtual rodent predicts the structure of neural activity across behaviors. *Nature*, 1–3.
- Alvernhe, A., Save, E. and Poucet, B. (2011) Local remapping of place cell firing in the Tolman detour task. *European Journal of Neuroscience*, **33**, 1696–1705.
- Banino, A., Barry, C., Uria, B., Blundell, C., Lillicrap, T., Mirowski, P., Pritzel, A., Chadwick, M. J., Degris, T., Modayil, J. et al. (2018) Vector-based navigation using grid-like representations in artificial agents. *Nature*, **557**, 429–433.
- Barreto, A., Dabney, W., Munos, R., Hunt, J. J., Schaul, T., van Hasselt, H. P. and Silver, D. (2017) Successor features for transfer in reinforcement learning. *Advances in neural information processing systems*, **30**.
- Behrens, T. E., Woolrich, M. W., Walton, M. E. and Rushworth, M. F. (2007) Learning the value of information in an uncertain world. *Nature neuroscience*, **10**, 1214–1221.

- Bellemare, M. G., Dabney, W. and Munos, R. (2017) A distributional perspective on reinforcement learning. In *International conference on machine learning*, 449–458. PMLR.
- Bellemare, M. G., Dabney, W. and Rowland, M. (2023) *Distributional reinforcement learning*. MIT Press.
- Blanco-Pozo, M., Akam, T. and Walton, M. (2021) Dopamine reports reward prediction errors, but does not update policy, during inference-guided choice. *bioRxiv*.
- Blodgett, H. C. (1929) The effect of the introduction of reward upon the maze performance of rats. *University of California publications in psychology*.
- Botvinick, M. and Toussaint, M. (2012) Planning as inference. *Trends in cognitive sciences*, **16**, 485–488.
- Botvinick, M., Wang, J. X., Dabney, W., Miller, K. J. and Kurth-Nelson, Z. (2020) Deep reinforcement learning and its neuroscientific implications. *Neuron*, **107**, 603–616.
- Botvinick, M. M. (2008) Hierarchical models of behavior and prefrontal function. *Trends in cognitive sciences*, **12**, 201–208.
- Botvinick, M. M. and Cohen, J. D. (2014) The computational and neural basis of cognitive control: charted territory and new frontiers. *Cognitive science*, **38**, 1249–1285.
- Botvinick, M. M., Niv, Y. and Barto, A. G. (2009) Hierarchically organized behavior and its neural foundations: A reinforcement learning perspective. *cognition*, **113**, 262–280.
- Callaway, F., van Opheusden, B., Gul, S., Das, P., Krueger, P. M., Griffiths, T. L. and Lieder, F. (2022) Rational use of cognitive resources in human planning. *Nature Human Behaviour*, **6**, 1112–1125.
- Coddington, L. T. and Dudman, J. T. (2019) Learning from action: reconsidering movement signaling in midbrain dopamine neuron activity. *Neuron*, **104**, 63–77.
- Dabney, W., Kurth-Nelson, Z., Uchida, N., Starkweather, C. K., Hassabis, D., Munos, R. and Botvinick, M. (2020) A distributional code for value in dopamine-based reinforcement learning. *Nature*, **577**, 671–675.
- Dabney, W., Rowland, M., Bellemare, M. and Munos, R. (2018) Distributional reinforcement learning with quantile regression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32.
- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P. and Dolan, R. J. (2011) Model-based influences on humans' choices and striatal prediction errors. *Neuron*, **69**, 1204–1215.
- Daw, N. D., Niv, Y. and Dayan, P. (2005) Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature neuroscience*, **8**, 1704–1711.
- Dayan, P. (1993) Improving generalization for temporal difference learning: The successor representation. *Neural computation*, **5**, 613–624.
- Dearden, R., Friedman, N. and Russell, S. (1998) Bayesian Q-learning. *Aaai/iaai*, **1998**, 761–768.
- Deisenroth, M. and Rasmussen, C. E. (2011) Pilco: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning (ICML-11)*, 465–472.
- Dhawale, A. K., Poddar, R., Wolff, S. B., Normand, V. A., Kopelowitz, E. and Ölveczky, B. P. (2017) Automated long-term recording and analysis of neural activity in behaving animals. *Elife*, **6**, e27702.
- Duan, Y., Schulman, J., Chen, X., Bartlett, P. L., Sutskever, I. and Abbeel, P. (2016) RL2: Fast reinforcement learning via slow reinforcement learning. *arXiv preprint arXiv:1611.02779*.
- Dubois, M., Habicht, J., Michely, J., Moran, R., Dolan, R. J. and Hauser, T. U. (2021) Human complex exploration strategies are enriched by noradrenaline-modulated heuristics. *Elife*, **10**, e59907.

- Dunn, T. W., Marshall, J. D., Severson, K. S., Aldarondo, D. E., Hildebrand, D. G., Chettih, S. N., Wang, W. L., Gellis, A. J., Carlson, D. E., Aronov, D. et al. (2021) Geometric deep learning enables 3D kinematic profiling across species and environments. *Nature methods*, **18**, 564–573.
- Eckstein, M. K. and Collins, A. G. (2020) Computational evidence for hierarchically structured reinforcement learning in humans. *Proceedings of the National Academy of Sciences*, **117**, 29381–29389.
- Espeholt, L., Soyer, H., Munos, R., Simonyan, K., Mnih, V., Ward, T., Doron, Y., Firoiu, V., Harley, T., Dunning, I. et al. (2018) IMPALA: Scalable distributed deep-RL with importance weighted actor-learner architectures. In *International conference on machine learning*, 1407–1416. PMLR.
- Fang, C. and Stachenfeld, K. L. (2023) Predictive auxiliary objectives in deep RL mimic learning in the brain. *arXiv preprint arXiv:2310.06089*.
- Finn, C., Abbeel, P. and Levine, S. (2017) Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, 1126–1135. PMLR.
- Geerts, J. P., Chersi, F., Stachenfeld, K. L. and Burgess, N. (2020) A general model of hippocampal and dorsal striatal learning and decision making. *Proceedings of the National Academy of Sciences*, **117**, 31427–31437.
- Gershman, S. J. (2018) The successor representation: its computational logic and neural substrates. *Journal of Neuroscience*, **38**, 7193–7200.
- Gershman, S. J., Assad, J. A., Datta, S. R., Linderman, S. W., Sabatini, B. L., Uchida, N. and Wilbrecht, L. (2024) Explaining dopamine through prediction errors and beyond. *Nature Neuroscience*, 1–11.
- Griffiths, T. L., Callaway, F., Chang, M. B., Grant, E., Krueger, P. M. and Lieder, F. (2019) Doing more with less: meta-reasoning and meta-learning in humans and machines. *Current Opinion in Behavioral Sciences*, **29**, 24–30.
- Gronauer, S. and Diepold, K. (2022) Multi-agent deep reinforcement learning: a survey. *Artificial Intelligence Review*, 1–49.
- Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., Abbeel, P. et al. (2018) Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*.
- Harlow, H. F. (1949) The formation of learning sets. *Psychological review*, **56**, 51.
- Hattori, R., Hedrick, N. G., Jain, A., Chen, S., You, H., Hattori, M., Choi, J.-H., Lim, B. K., Yasuda, R. and Komiyama, T. (2023) Meta-reinforcement learning via orbitofrontal cortex. *Nature Neuroscience*, 1–10.
- Horvitz, J. C. (2000) Mesolimbocortical and nigrostriatal dopamine responses to salient non-reward events. *Neuroscience*, **96**, 651–656.
- Howe, M. W., Tierney, P. L., Sandberg, S. G., Phillips, P. E. and Graybiel, A. M. (2013) Prolonged dopamine signalling in striatum signals proximity and value of distant rewards. *nature*, **500**, 575–579.
- Huys, Q. J., Eshel, N., O’Nions, E., Sheridan, L., Dayan, P. and Roiser, J. P. (2012) Bonsai trees in your head: how the Pavlovian system sculpts goal-directed choices by pruning decision trees. *PLoS computational biology*, **8**, e1002410.
- Jaderberg, M., Mnih, V., Czarnnecki, W. M., Schaul, T., Leibo, J. Z., Silver, D. and Kavukcuoglu, K. (2016) Reinforcement learning with unsupervised auxiliary tasks. *arXiv preprint arXiv:1611.05397*.
- Jensen, K. T., Hennequin, G. and Mattar, M. G. (2023) A recurrent network model of planning explains hippocampal replay and human behavior. *bioRxiv*, 2023–01.
- Jie, T. and Abbeel, P. (2010) On a connection between importance sampling and the likelihood ratio policy gradient. *Advances in Neural Information Processing Systems*, **23**.
- Juechems, K. and Summerfield, C. (2019) Where does value come from? *Trends in cognitive sciences*, **23**, 836–850.

- Kakade, S. and Dayan, P. (2002) Dopamine: generalization and bonuses. *Neural Networks*, **15**, 549–559.
- Kalashnikov, D., Irpan, A., Pastor, P., Ibarz, J., Herzog, A., Jang, E., Quillen, D., Holly, E., Kalakrishnan, M., Vanhoucke, V. et al. (2018) QT-Opt: Scalable deep reinforcement learning for vision-based robotic manipulation. *arXiv preprint arXiv:1806.10293*.
- Keramati, M., Smittenaar, P., Dolan, R. J. and Dayan, P. (2016) Adaptive integration of habits into depth-limited planning defines a habitual-goal-directed spectrum. *Proceedings of the National Academy of Sciences*, **113**, 12868–12873.
- Killcross, S. and Coutureau, E. (2003) Coordination of actions and habits in the medial prefrontal cortex of rats. *Cerebral cortex*, **13**, 400–408.
- Lai, L. and Gershman, S. J. (2021) Policy compression: An information bottleneck in action selection. In *Psychology of Learning and Motivation*, vol. 74, 195–232. Elsevier.
- Lengyel, M. and Dayan, P. (2007) Hippocampal contributions to control: the third way. *Advances in neural information processing systems*, **20**.
- Levine, S. (2018) Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909*.
- Levine, S., Kumar, A., Tucker, G. and Fu, J. (2020) Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*.
- Li, C., Kreiman, G. and Ramanathan, S. (2022) Integrating artificial and biological neural networks to improve animal task performance using deep reinforcement learning. *bioRxiv*, 2022–09.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D. and Wierstra, D. (2015) Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*.
- Loukola, O. J., Solvi, C., Coscos, L. and Chittka, L. (2017) Bumblebees show cognitive flexibility by improving on an observed complex behavior. *Science*, **355**, 833–836.
- Lowet, A. S., Zheng, Q., Matias, S., Drugowitsch, J. and Uchida, N. (2020) Distributional reinforcement learning in the brain. *Trends in neurosciences*, **43**, 980–997.
- Makino, H. (2023) Arithmetic value representation for hierarchical behavior composition. *Nature Neuroscience*, **26**, 140–149.
- Mathis, A., Mamidanna, P., Cury, K. M., Abe, T., Murthy, V. N., Mathis, M. W. and Bethge, M. (2018) DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nature neuroscience*, **21**, 1281–1289.
- Mattar, M. G. and Daw, N. D. (2018) Prioritized memory access explains planning and hippocampal replay. *Nature neuroscience*, **21**, 1609–1617.
- Mehta, M. R., Quirk, M. C. and Wilson, M. A. (2000) Experience-dependent asymmetric shape of hippocampal receptive fields. *Neuron*, **25**, 707–715.
- Merel, J., Aldarondo, D., Marshall, J., Tassa, Y., Wayne, G. and Ölveczky, B. (2019) Deep neuroethology of a virtual rodent. *arXiv preprint arXiv:1911.09451*.
- Miller, E. K. and Cohen, J. D. (2001) An integrative theory of prefrontal cortex function. *Annual review of neuroscience*, **24**, 167–202.
- Miller, K. J., Botvinick, M. M. and Brody, C. D. (2017) Dorsal hippocampus contributes to model-based planning. *Nature neuroscience*, **20**, 1269–1276.
- Mizes, K. G., Lindsey, J., Escola, G. S. and Ölveczky, B. P. (2023a) Dissociating the contributions of sensorimotor striatum to automatic and visually guided motor sequences. *Nature Neuroscience*, 1–14.

- Mizes, K. G., Lindsey, J., Escola, G. S. and Olveczky, B. P. (2023b) Motor cortex is required for flexible but not automatic motor sequences. *bioRxiv*, 2023–09.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D. and Riedmiller, M. (2013) Playing Atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*.
- Momennejad, I., Russek, E. M., Cheong, J. H., Botvinick, M. M., Daw, N. D. and Gershman, S. J. (2017) The successor representation in human reinforcement learning. *Nature human behaviour*, **1**, 680–692.
- Morimura, T., Sugiyama, M., Kashima, H., Hachiya, H. and Tanaka (2012) Parametric return density estimation for reinforcement learning. *arXiv preprint arXiv:1203.3497*.
- Morimura, T., Sugiyama, M., Kashima, H., Hachiya, H. and Tanaka, T. (2010) Nonparametric return distribution approximation for reinforcement learning. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, 799–806.
- Morris, G., Nevet, A., Arkadir, D., Vaadia, E. and Bergman, H. (2006) Midbrain dopamine neurons encode decisions for future action. *Nature neuroscience*, **9**, 1057–1063.
- Niv, Y. (2009) Reinforcement learning in the brain. *Journal of Mathematical Psychology*, **53**, 139–154.
- Nota, C. and Thomas, P. S. (2019) Is the policy gradient a gradient? *arXiv preprint arXiv:1906.07073*.
- Nowé, A., Vrancx, P. and De Hauwere, Y.-M. (2012) Game theory and multi-agent reinforcement learning. *Reinforcement Learning: State-of-the-Art*, 441–470.
- O’Doherty, J., Dayan, P., Schultz, J., Deichmann, R., Friston, K. and Dolan, R. J. (2004) Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *science*, **304**, 452–454.
- Olds, J. and Milner, P. (1954) Positive reinforcement produced by electrical stimulation of septal area and other regions of rat brain. *Journal of comparative and physiological psychology*, **47**, 419.
- Ölveczky, B. P., Andalman, A. S. and Fee, M. S. (2005) Vocal experimentation in the juvenile songbird requires a basal ganglia circuit. *PLoS biology*, **3**, e153.
- Pachitariu, M., Stringer, C., Schröder, S., Dipoppa, M., Rossi, L. F., Carandini, M. and Harris, K. D. (2016) Suite2p: beyond 10,000 neurons with standard two-photon microscopy. *BioRxiv*, 061507.
- Padoa-Schioppa, C. and Assad, J. A. (2006) Neurons in the orbitofrontal cortex encode economic value. *Nature*, **441**, 223–226.
- Pan, Y., Cheng, C.-A., Saigol, K., Lee, K., Yan, X., Theodorou, E. and Boots, B. (2017) Agile autonomous driving using end-to-end deep imitation learning. *arXiv preprint arXiv:1709.07174*.
- Pateria, S., Subagdja, B., Tan, A.-h. and Quek, C. (2021) Hierarchical reinforcement learning: A comprehensive survey. *ACM Computing Surveys (CSUR)*, **54**, 1–35.
- Peshkin, L. and Shelton, C. R. (2002) Learning from scarce experience. *arXiv preprint cs/0204043*.
- Piray, P. and Daw, N. (2024) Reconciling flexibility and efficiency: Medial entorhinal cortex represents a compositional cognitive map. *bioRxiv*, 2024–05.
- Piray, P. and Daw, N. D. (2021) Linear reinforcement learning in planning, grid fields, and cognitive control. *Nature communications*, **12**, 1–20.
- Rao, R. P. and Ballard, D. H. (1999) Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, **2**, 79–87.
- Ritter, S., Wang, J., Kurth-Nelson, Z., Jayakumar, S., Blundell, C., Pascanu, R. and Botvinick, M. (2018) Been there, done that: Meta-learning with episodic recall. In *International conference on machine learning*, 4354–4363. PMLR.

- Roesch, M. R., Calu, D. J. and Schoenbaum, G. (2007) Dopamine neurons encode the better option in rats deciding between differently delayed or sized rewards. *Nature neuroscience*, **10**, 1615–1624.
- Rowland, M., Dadashi, R., Kumar, S., Munos, R., Bellemare, M. G. and Dabney, W. (2019) Statistics and samples in distributional reinforcement learning. In *International Conference on Machine Learning*, 5528–5536. PMLR.
- Rushworth, M. F., Noonan, M. P., Boorman, E. D., Walton, M. E. and Behrens, T. E. (2011) Frontal cortex and reward-guided learning and decision-making. *Neuron*, **70**, 1054–1069.
- Schaul, T., Quan, J., Antonoglou, I. and Silver, D. (2015) Prioritized experience replay. *arXiv preprint arXiv:1511.05952*.
- Schrittwieser, J., Antonoglou, I., Hubert, T., Simonyan, K., Sifre, L., Schmitt, S., Guez, A., Lockhart, E., Hassabis, D., Graepel, T. et al. (2020) Mastering Atari, Go, chess and shogi by planning with a learned model. *Nature*, **588**, 604–609.
- Schultz, W., Apicella, P., Scarnati, E. and Ljungberg, T. (1992) Neuronal activity in monkey ventral striatum related to the expectation of reward. *Journal of neuroscience*, **12**, 4595–4610.
- Schultz, W., Dayan, P. and Montague, P. R. (1997) A neural substrate of prediction and reward. *Science*, **275**, 1593–1599.
- Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T. et al. (2018) A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, **362**, 1140–1144.
- Silver, D., Singh, S., Precup, D. and Sutton, R. S. (2021) Reward is enough. *Artificial Intelligence*, **299**, 103535.
- Solway, A. and Botvinick, M. M. (2012) Goal-directed decision making as probabilistic inference: a computational framework and potential neural correlates. *Psychological review*, **119**, 120.
- Song, H. F., Yang, G. R. and Wang, X.-J. (2017) Reward-based training of recurrent neural networks for cognitive and value-based tasks. *Elife*, **6**, e21492.
- Sousa, M., Bujalski, P., Cruz, B. F., Louie, K., McNamee, D. and Paton, J. J. (2023) Dopamine neurons encode a multidimensional probabilistic map of future reward. *bioRxiv*, 2023–11.
- Stachenfeld, K. L., Botvinick, M. M. and Gershman, S. J. (2017) The hippocampus as a predictive map. *Nature neuroscience*, **20**, 1643–1653.
- Steinberg, E. E., Keiflin, R., Boivin, J. R., Witten, I. B., Deisseroth, K. and Janak, P. H. (2013) A causal link between prediction errors, dopamine neurons and learning. *Nature neuroscience*, **16**, 966–973.
- Steinmetz, N. A., Aydin, C., Lebedeva, A., Okun, M., Pachitariu, M., Bauza, M., Beau, M., Bhagat, J., Böhm, C., Broux, M. et al. (2021) Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science*, **372**, eabf4588.
- Sutton, R. (1991) Dyna, an integrated architecture for learning, planning, and reacting. *ACM Sigart Bulletin*, **2**, 160–163.
- Sutton, R. S. (1988) Learning to predict by the methods of temporal differences. *Machine learning*, **3**, 9–44.
- Sutton, R. S. and Barto, A. G. (2018) *Reinforcement learning: An introduction*. MIT press.
- Sutton, R. S., Precup, D. and Singh, S. (1999) Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, **112**, 181–211.
- Takahashi, Y., Schoenbaum, G. and Niv, Y. (2008) Silencing the critics: understanding the effects of cocaine sensitization on dorsolateral and ventral striatum in the context of an actor/critic model. *Frontiers in neuroscience*, **2**, 282.
- Tano, P., Dayan, P. and Pouget, A. (2020) A local temporal difference code for distributional reinforcement learning. *Advances in neural information processing systems*, **33**, 13662–13673.
- Team, G., Anil, R., Borgeaud, S., Wu, Y., Alayrac, J.-B., Yu, J., Soricut, R., Schalkwyk, J., Dai, A. M., Hauth, A. et al. (2023) Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.

- Tervo, D. G., Proskurin, M., Manakov, M., Kabra, M., Vollmer, A., Branson, K. and Karpova, A. Y. (2014) Behavioral variability through stochastic choice and its gating by anterior cingulate cortex. *Cell*, **159**, 21–32.
- Todorov (2009) Efficient computation of optimal actions. *Proceedings of the national academy of sciences*, **106**, 11478–11483.
- Todorov, E. (2006) Linearly-solvable Markov decision problems. *Advances in neural information processing systems*, **19**.
- Tolman, E. C. (1948) Cognitive maps in rats and men. *Psychological review*, **55**, 189.
- Tsai, H.-C., Zhang, F., Adamantidis, A., Stuber, G. D., Bonci, A., De Lecea, L. and Deisseroth, K. (2009) Phasic firing in dopaminergic neurons is sufficient for behavioral conditioning. *Science*, **324**, 1080–1084.
- Van Hasselt, H., Guez, A. and Silver, D. (2016) Deep reinforcement learning with double Q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, vol. 30.
- Vértes, E. and Sahani, M. (2019) A neurally plausible model learns successor representations in partially observable environments. *Advances in Neural Information Processing Systems*, **32**.
- Vikbladh, O. M., Meager, M. R., King, J., Blackmon, K., Devinsky, O., Shohamy, D., Burgess, N. and Daw, N. D. (2019) Hippocampal contributions to model-based planning and spatial memory. *Neuron*, **102**, 683–693.
- Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., Choi, D. H., Powell, R., Ewalds, T., Georgiev, P. et al. (2019) Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, **575**, 350–354.
- Viswanathan, G. M., Buldyrev, S. V., Havlin, S., da Luz, M. G., Raposo, E. P. and Stanley, H. E. (1999) Optimizing the success of random searches. *nature*, **401**, 911–914.
- Wang, J. X., Kurth-Nelson, Z., Kumaran, D., Tirumala, D., Soyer, H., Leibo, J. Z., Hassabis, D. and Botvinick, M. (2018) Prefrontal cortex as a meta-reinforcement learning system. *Nature neuroscience*, **21**, 860–868.
- Wang, J. X., Kurth-Nelson, Z., Tirumala, D., Soyer, H., Leibo, J. Z., Munos, R., Blundell, C., Kumaran, D. and Botvinick, M. (2016) Learning to reinforcement learn. *arXiv preprint arXiv:1611.05763*.
- Watabe-Uchida, M., Eshel, N. and Uchida, N. (2017) Neural circuitry of reward prediction error. *Annual review of neuroscience*, **40**, 373–394.
- Watkins, C. J. and Dayan, P. (1992) Q-learning. *Machine learning*, **8**, 279–292.
- Watkins, C. J. C. H. (1989) Learning from delayed rewards.
- Wenliang, L. K., Délatang, G., Aitchison, M., Hutter, M., Ruoss, A., Gretton, A. and Rowland, M. (2023) Distributional Bellman operators over mean embeddings. *arXiv preprint arXiv:2312.07358*.
- Whittington, J. C., Muller, T. H., Mark, S., Chen, G., Barry, C., Burgess, N. and Behrens, T. E. (2020) The Tolman-Eichenbaum machine: unifying space and relational memory through generalization in the hippocampal formation. *Cell*, **183**, 1249–1263.
- Williams, R. J. (1992) Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, **8**, 229–256.
- Wurman, P. R., Barrett, S., Kawamoto, K., MacGlashan, J., Subramanian, K., Walsh, T. J., Capobianco, R., Devlic, A., Eckert, F., Fuchs, F. et al. (2022) Outracing champion Gran Turismo drivers with deep reinforcement learning. *Nature*, **602**, 223–228.
- Yin, H. H., Knowlton, B. J. and Balleine, B. W. (2004) Lesions of dorsolateral striatum preserve outcome expectancy but disrupt habit formation in instrumental learning. *European journal of neuroscience*, **19**, 181–189.
- Yin, H. H., Ostlund, S. B., Knowlton, B. J. and Balleine, B. W. (2005) The role of the dorsomedial striatum in instrumental conditioning. *European Journal of Neuroscience*, **22**, 513–523.
- Zintgraf, L., Shiarlis, K., Igl, M., Schulze, S., Gal, Y., Hofmann, K. and Whiteson, S. (2019) VariBAD: A very good method for bayes-adaptive deep RL via meta-learning. *arXiv preprint arXiv:1910.08348*.

11 | ADDITIONAL TOPICS OF INTEREST

While we have tried to provide a fairly comprehensive overview of topics in reinforcement learning of interest to neuroscience, there are naturally many interesting areas that we have had to omit. Here we provide a brief description of some of these together with pointers to relevant literature for those who are interested in exploring them further.

11.1 | Hierarchical reinforcement learning

So far, we have considered a simple environment consisting of discrete states and actions, and all planning and decision making has taken place in the space of action primitives. However, when planning over longer horizons, it can be necessary to break down the overall policy into a series of sub-goals, sub-policies, or 'skills' (Sutton et al., 1999; Pateria et al., 2021). This is the topic of hierarchical reinforcement learning (HRL) and 'options', where an agent learns a high-level policy over policies that can themselves be specified in terms of primitive actions or even lower-level policies. Such HRL has been found to explain features of human behaviour (Eckstein and Collins, 2020; Botvinick, 2008; Botvinick et al., 2009) and remains an area of substantial interest in the neuroscience literature.

11.2 | Off-policy & offline reinforcement learning

In most of the work considered in this paper, the experience used to train the RL agents has been sampled from the policy of the agent itself. Indeed this is required for the gradients to be unbiased in the vanilla policy gradient setting. However, an area of substantial interest is that of offline reinforcement learning, where the agent is trained from scratch on the basis of pre-collected experience (Levine et al., 2020). This is particularly important in cases where online data collection is expensive or too risky but large-scale datasets exist, such as in many healthcare settings. Off-policy reinforcement learning is the related problem of learning from a combination of online data and pre-generated data, possibly from a 'stale' version of the current agent. The off-policy setting is especially relevant to biology, where data collection is expensive and we therefore wish to make maximum use of existing data. This can e.g. be achieved through experience replay, which can be prioritized (instead of sampled at random) to maximize future reward and minimize temporal opportunity costs (Mattar and Daw, 2018; Agrawal et al., 2022; Schaul et al., 2015). A variety of 'off-policy' policy gradient methods have also been developed to improve sample efficiency, which de-bias the gradients e.g. through the use of importance sampling (Espeholt et al., 2018; Jie and Abbeel, 2010; Peshkin and Shelton, 2002; Haarnoja et al., 2018).

11.3 | Imitation learning

Related to the problem of offline reinforcement learning is that of *imitation learning*, where we also learn from pre-collected data. However, in contrast to offline RL where we make no assumption about the quality of the policy used to collect the data, imitation learning assumes that the data has been collected by an 'expert' we wish to imitate (Levine et al., 2020). This is useful in cases where a large amount of expert data is available, such as the case of autonomous driving (Pan et al., 2017). Imitation learning is clearly important during early development in biological organisms, where we learn from observing the individuals around us. Indeed, such imitation learning is a hallmark not just of humans but has also been demonstrated in organisms as 'simple' as the bumblebee (Loukola et al., 2017). Imitation learning has also recently been used to learn models of biological neural circuits from high-resolution behavioural data (Aldarondo et al., 2024).

11.4 | Linear reinforcement learning

As we have seen in most of this tutorial, reinforcement learning is generally difficult and requires iterative algorithms that often scale poorly with the problem size. However, there are settings where we can simplify the problem to the

point where it becomes analytically tractable in an approach known as ‘linear reinforcement learning’ (Todorov, 2006, 2009). This is similar to the SR approach, where we saw that the value function reduces to a linear function of the reward-per-state. Similar to how the SR matrix can be seen as describing the dynamics of some ‘base policy’, we also define a base policy in linear RL and compute a ‘control cost’ as the KL divergence between transition dynamics with and without our controller:

$$\mathcal{L}_{ctrl}(s) = KL[u(s'|s)||p(s'|s)], \quad (46)$$

where $p(s'|s)$ are the prior transition dynamics and $u(s'|s)$ are the controlled transition dynamics marginalized over the policy. For $\mathcal{L}_{ctrl}(s)$ to be well-defined, we require $u(s'|s) = 0$ whenever $p(s'|s) = 0$, which prevents impossible transitions even under our flexible controller. When subtracting this loss from the RL objective, the resulting utility turns out to be convex in u and can therefore be solved efficiently for the controller, which implicitly specifies the policy. This approach has recently been used as an explicit model of biological decision making (Piray and Daw, 2021, 2024). It also has close parallels to learning and planning as inference (Levine, 2018; Solway and Botvinick, 2012; Botvinick and Toussaint, 2012) and to RL with information bottlenecks (Lai and Gershman, 2021). Both of these families of approaches involve reinforcement learning with a KL-regularized reward function, and they have also been used as models of biological decision making.

11.5 | Successor features

In Section 6, we saw that the successor representation can be used for decision making with flexible adaptation in environments with changing reward structures. However, we developed this framework only in the tabular setting despite extending TD-learning and Q-learning to the ‘deep RL’ setting with function approximation. This leaves open the question of whether a similar generalization of the SR exists. This turns out to be the case and is known as ‘successor features’ (SF; Barreto et al., 2017), where the expected future observation of learned features of the environment are used in place of the expected future state occupancy. Successor features have also been shown to have a biologically plausible implementation that facilitates learning and generalization in noisy and partially observable environments (Vértes and Sahani, 2019).

11.6 | Multi-agent reinforcement learning

We have only considered the case of single agents interacting with a black-box environment. However, in many cases, multiple agents are simultaneously interacting with each other and the environment around them (Gronauer and Diepold, 2022). This means that, from the point of view of a single agent, the other agents are part of its environment. In such settings, there are interesting learning dynamics beyond the scope of the present tutorial, but which are covered in detail by e.g. Gronauer and Diepold (2022), and which are also of substantial interest in game theory (Nowé et al., 2012). In some cases, a whole group of agents may be working together to maximize a single joint reward function – as is the case for members of a single sports team. Interestingly, the learning of many individual neurons in the brain from a single common reinforcing signal (such as dopamine) can be modelled as such a multi-agent reinforcement learning problem (Sutton and Barto, 2018). If the ‘agents’ (or neurons) are assumed to have Bernoulli-logistic outputs, Williams (1992) shows that the independent learning of individual agents from the global reward signal leads to the implementation of a policy gradient algorithm at the population level (Sutton and Barto, 2018).