

Reinforcement learning approaches to hippocampus-dependent flexible spatial navigation

Charline Tessereau^{1,2,3} , Reuben O'Dea^{1,3}, Stephen Coombes^{1,3}
and Tobias Bast^{2,3}

Brain and Neuroscience Advances
Volume 5: 1–18
© The Author(s) 2021
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/2398212820975634
journals.sagepub.com/home/bna



Abstract

Humans and non-human animals show great flexibility in spatial navigation, including the ability to return to specific locations based on as few as one single experience. To study spatial navigation in the laboratory, watermaze tasks, in which rats have to find a hidden platform in a pool of cloudy water surrounded by spatial cues, have long been used. Analogous tasks have been developed for human participants using virtual environments. Spatial learning in the watermaze is facilitated by the hippocampus. In particular, rapid, one-trial, allocentric place learning, as measured in the delayed-matching-to-place variant of the watermaze task, which requires rodents to learn repeatedly new locations in a familiar environment, is hippocampal dependent. In this article, we review some computational principles, embedded within a reinforcement learning framework, that utilise hippocampal spatial representations for navigation in watermaze tasks. We consider which key elements underlie their efficacy, and discuss their limitations in accounting for hippocampus-dependent navigation, both in terms of behavioural performance (i.e. how well do they reproduce behavioural measures of rapid place learning) and neurobiological realism (i.e. how well do they map to neurobiological substrates involved in rapid place learning). We discuss how an actor–critic architecture, enabling simultaneous assessment of the value of the current location and of the optimal direction to follow, can reproduce one-trial place learning performance as shown on watermaze and virtual delayed-matching-to-place tasks by rats and humans, respectively, if complemented with map-like place representations. The contribution of actor–critic mechanisms to delayed-matching-to-place performance is consistent with neurobiological findings implicating the striatum and hippocampus–striatal interaction in delayed-matching-to-place performance, given that the striatum has been associated with actor–critic mechanisms. Moreover, we illustrate that hierarchical computations embedded within an actor–critic architecture may help to account for aspects of flexible spatial navigation. The hierarchical reinforcement learning approach separates trajectory control via a temporal-difference error from goal selection via a goal prediction error and may account for flexible, trial-specific, navigation to familiar goal locations, as required in some arm-maze place memory tasks, although it does not capture one-trial learning of new goal locations, as observed in open field, including watermaze and virtual, delayed-matching-to-place tasks. Future models of one-shot learning of new goal locations, as observed on delayed-matching-to-place tasks, should incorporate hippocampal plasticity mechanisms that integrate new goal information with allocentric place representation, as such mechanisms are supported by substantial empirical evidence.

Keywords

Reinforcement learning, computational modelling, spatial navigation, Morris watermaze, one-shot learning, place learning and memory, hierarchical agent

Received: 31 July 2020; accepted: 21 October 2020

Introduction

Successful spatial navigation is required for many everyday tasks: animals have to find food and shelter and remember where and how to find these, humans need to navigate to work, home or to the supermarket. Since natural environments are inherently heterogeneous and subject to continuous change, animal brains have evolved robust and flexible solutions to solve challenges in spatial navigation.

In the mammalian brain, and to some extent also the avian brain (Bingman and Sharp, 2006; Colombo and Broadbent, 2000), the hippocampus has long been recognised to play a central role in place learning, memory and spatial navigation, based on the behavioural effects of lesions and other manipulations of the hippocampus (Morris et al., 1982, 1986, 1990) and based on the spatial tuning of certain hippocampal neurons, so-called place cells (Jeffery, 2018; Moser et al., 2017; O'Keefe, 2014; O'Keefe and Dostrovsky, 1971). Studies combining hippocampal manipulations with behavioural testing in rodents have revealed that the hippocampus is particularly important for

¹School of Mathematical Sciences, University of Nottingham, Nottingham, UK

²School of Psychology, University of Nottingham, Nottingham, UK

³Neuroscience@Nottingham

Corresponding authors:

Charline Tessereau, School of Mathematical Sciences, University of Nottingham, University Park, Nottingham NG7 2RD, UK.
Email: tessereau.charline@gmail.com

Reuben O'Dea, School of Mathematical Sciences, University of Nottingham, University Park, Nottingham NG7 2RD, UK.
Email: Reuben.O'Dea@nottingham.ac.uk

Stephen Coombes, School of Mathematical Sciences, University of Nottingham, University Park, Nottingham NG7 2RD, UK.
Email: Stephen.Coombes@nottingham.ac.uk

Tobias Bast, School of Psychology, University of Nottingham, University Park, Nottingham NG7 2RD, UK.
Email: Tobias.Bast@nottingham.ac.uk



flexible spatial navigation based on rapid allocentric place learning, in which places are learned based on their relationship to environmental cues (Bast et al., 2009; Eichenbaum, 1990; Morris et al., 1990; Steele and Morris, 1999).

Animal experiments on spatial navigation have been complemented by tools from theoretical and computational neuroscience. Many theoreticians have targeted spatial navigation problems, either trying to reproduce behaviours (Banino et al., 2018; Dayan, 1991) or to explain properties of neurons that show spatial tuning, including hippocampal place cells (Banino et al., 2018; Fuhs and Touretzky, 2006; O'Keefe and Burgess, 1996; Samsonovich and McNaughton, 1997; Widloski and Fiete, 2014). Some have also studied the functional properties of network representations of neural computations. For example, exploring how static attributes such as place coding can lead to or be integrated within network dynamics (Kanitscheider and Fiete, 2017), or have probed the storage capacity of spatial representations (Battaglia and Treves, 1998).

In a spatial navigation context, most real-world situations involve choosing a behavioural response that leads to a goal location associated with a reward. These could be direct rewards, such as food, or escape from an unpleasant situation, for example, escape from water in the watermaze. Successful navigation requires animals to distinguish diverse cues in their environment, to encode their own current position, to access a memory of where the goal is, to choose an appropriate trajectory, and to recognise the goal and its vicinity. Learning how to reach a goal location is a problem that, in principle, fits very well within a reinforcement learning (RL) context. RL commonly refers to a computational framework that studies how intelligent systems learn to associate situations with actions in order to maximise the rewards within an environment (Sutton and Barto, 2018). When applied to spatial navigation, a RL model can be used to infer how neuronal representations of space, as revealed by electrophysiological recordings, may serve to maximise reward (Banino et al., 2018; Corneil and Gerstner, 2015; Dollé et al., 2018; Foster et al., 2000; Gerstner and Abbott, 1997; Russek et al., 2017). RL models have led to numerous successes in understanding how biological networks could produce observed behaviours (Frankenhuis et al., 2019; Haferlach et al., 2007), yet there are still substantial challenges in using RL approaches to account for animal behaviour, such as hippocampus-dependent flexible navigation based on rapid place learning.

The aim of this article is to review RL models that may account for hippocampus-dependent rapid place learning, especially as seen in the watermaze delayed-matching-to-place (DMP) task. We will especially focus on an exemplar approach to this problem proposed by Foster et al. (2000). In the section 'Flexible hippocampal spatial navigation', we review briefly some key experimental findings on the involvement of the hippocampus in spatial navigation tasks in the watermaze, highlighting its particular importance in rapid place learning. The section 'RL for spatial navigation' contains an overview of the key concepts in RL. The section 'A model-free agent using an actor-critic architecture' describes the first part of the model by Foster et al. (2000), an RL architecture that provides a computational approach to how a place can become associated with actions to pursue a reward. We present a detailed description of the computations underlying the behaviour of the model and

their possible biological substrates, which we hope may make the model more accessible to neuroscientists without a strong neurocomputational background. Then, we focus on two minimal extensions to this architecture that enable adaptation to a changing reward location. The first one involves a map-like representation of location that enables vector-based navigation and was proposed by Foster et al. (2000); we show that this extension can reproduce key measures of rapid place learning performance on the DMP task, including sharp latency reductions from trial 1 to 2 (Steele and Morris, 1999), and also the more recent finding that rats show search preference for the correct location within one trial (Bast et al. (2009), see section 'Map-like representation of locations for goal-directed trajectories'). The second uses ideas drawn from hierarchical RL (Botvinick et al., 2009; Schweighofer and Doya, 2003), in which adding layers of control allows the agent more flexible behaviours. We discuss details of these computations, their correspondence to neurobiological findings, and the plausibility of their implementation, in particular, to account for rapid place learning within an artificial watermaze set-up (see section 'Hierarchical control to flexibly adjust to task requirements'). In the 'Conclusion' section, we emphasise some of the computational principles that we propose hold particular promise for neuropsychologically realistic models of rapid place learning in the watermaze.

Flexible hippocampal spatial navigation

Humans and other animals show remarkable *flexibility* in spatial navigation. In this context, flexibility refers to the ability to adjust to a changing environment, such as the variation in the goal or start location (Tolman, 1948). Watermaze tasks, in which rodents learn to find a hidden escape platform in a circular pool of water surrounded by spatial cues (Morris, 2008), have been important tools to study the neuropsychological mechanisms of such flexible spatial navigation in rodents. In the original task, the platform location remains the same over many trials and days of training. The animals can incrementally learn the place of the hidden platform using distal cues surrounding the watermaze and then navigate to it from different start positions (Morris, 1981). Learning is reflected by a reduction in the time taken to reach the platform location ('escape latencies') across trials and a search preference for the vicinity of the goal location when the platform is removed in probe trials.

Rapid place learning can be assessed in the watermaze through the DMP task, where the location of the platform remains constant during trials within a day (typically four trials per day), but is changed every day (Figure 1(a); Bast et al., 2009; Steele and Morris, 1999). A key observation from the behaviour of rats on the DMP task is that a single trial to a new goal location is sufficient for the animal to learn this location and subsequently to navigate to it efficiently (Steele and Morris, 1999). This phenomenon is therefore commonly referred to as 'one-shot' or 'one-trial' place learning. Such one-trial place learning is reflected by a marked latency reduction between the first and second trials to a new goal location (Figure 1(b)), with little further improvement on subsequent trials, and by a marked search preference for the vicinity of the correct location when trial 2 is run as probe (Figure

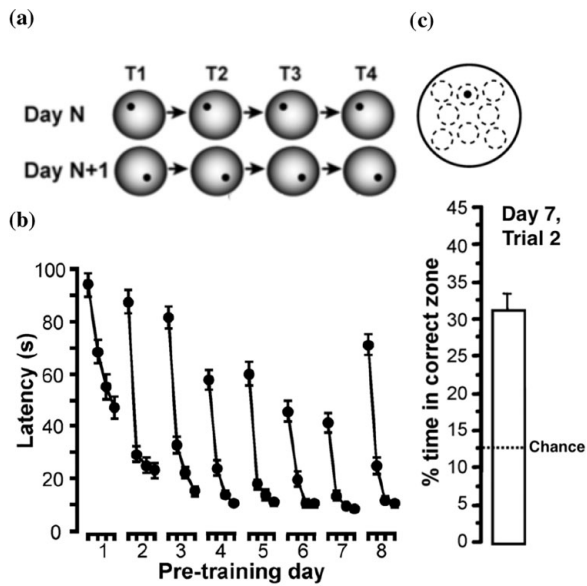


Figure 1. One-shot place learning by rats in the delayed-matching-to-place (DMP) watermaze task. (a) Rats have to learn a new goal location (location of escape platform) every day and complete four navigation trials to the new location on each day. (b) The time taken to find the new location reduces markedly from trial 1 to 2, with little further improvements on trials 3 and 4, and minimal interference between days. (c) When trial 2 is run as a probe trial, during which the platform is unavailable, rats show marked search preference for the vicinity of the goal location. To measure search preference, the watermaze surface is divided into eight equivalent symmetrically arranged zones (stippled lines in sketch), including the ‘correct zone’ centred on the goal location (black dot). The search preference corresponds to time spent searching in the ‘correct zone’, expressed as a percentage of time spent in all eight zones together. The chance level corresponds to 12.5%, corresponding to the rat spending the same time in each of the eight zones depicted in the sketch. These behavioural measures highlight successful one-shot place learning. Figure adapted from Figure 2 in Bast et al. (2009).

1(c) with the platform removed (Bast et al., 2009). Buckley and Bast (2018) reverse-translated the watermaze DMP task into a task for human participants, using a virtual environment presented on a computer screen, and have shown that human participants exhibit similar one-trial place learning to rats.

On the incremental place learning task in the watermaze, hippocampal lesions are known to disrupt rats’ performance (Morris et al., 1982), slowing down learning (Morris et al., 1990) and severely limiting rats’ ability to navigate to the goal from variable start positions (Eichenbaum, 1990). However, rats with partial hippocampal lesions sparing less than half of the hippocampus can show relatively intact performance on the incremental place learning task (De Hoz et al., 2003; Moser et al., 1995), and even rats with complete hippocampal lesions can show intact place memory following extended incremental training (Bast et al., 2009; Morris et al., 1990). Rats can also show intact incremental place learning on the watermaze with blockade of hippocampal synaptic plasticity if they received pretraining (Bannerman et al., 1995; Inglis et al., 2013). These findings suggest that incremental

place learning, although normally facilitated by hippocampal mechanisms, can partly be sustained by extra-hippocampal mechanisms.

In contrast to incremental place learning, rapid place learning, based on one or a few experiences, may absolutely require the hippocampus, with extra-hippocampal mechanisms unable to sustain such learning (Bast, 2007). Studies in rodents have shown that spatial navigation based on one-trial place learning on the DMP watermaze task is highly sensitive to hippocampal dysfunction that may leave incremental place learning performance in the watermaze relatively intact.

Specifically, one-trial place learning performance on the watermaze DMP test is severely impaired, and often virtually abolished, by complete and partial hippocampal lesions (Bast et al., 2009; De Hoz et al., 2005; Morris et al., 1990), as well as by disruption of hippocampal plasticity mechanisms (Inglis et al., 2013; Nakazawa et al., 2003; O’Carroll et al., 2006; Pezze and Bast, 2012; Steele and Morris, 1999; also compare similar findings by Bast et al. (2005) in a dry-land food-reinforced DMP task) or by aberrant hippocampal firing patterns (McGarrity et al., 2017). Rats with hippocampal lesions and NMDA (*n*-methyl-*d*-aspartate) receptor blockade show similar swim paths on trial 1 and trial 2 to the same goal location, swimming in circles over large areas of the watermaze surface (Redish and Touretzky, 1998; Steele and Morris, 1999), suggesting that they do not have, or cannot access, information about the recent goal location and/or the history of their positions.

Consistent with findings in rats that watermaze DMP performance is highly hippocampus-dependent, human participants’ one-trial place learning performance on the virtual DMP task is strongly associated with theta oscillations in the medial temporal lobe (including the hippocampus, Bauer et al., 2020). Overall, the findings reviewed above suggest that the DMP paradigm is a more sensitive assay of hippocampus-dependent navigation than incremental place learning paradigms, as good performance on the DMP task may absolutely require the hippocampus, with extra-hippocampal mechanisms unable to sustain such learning (Bast, 2007). In the following sections, we will review some RL approaches that link spatial representations in the hippocampus with successful navigation. We will focus on the performance and limitations of these methods in accounting for navigation based on rapid, one-trial, hippocampal place learning, especially as assessed by the watermaze and virtual DMP tasks in rodents and humans, respectively.

RL for spatial navigation

Typical RL problems involve four components: states, actions, values and policy (Sutton and Barto, 2018), as shown in Figure 2(a). In spatial navigation, states usually represent the agent location, but can be extended to describe more abstract concepts such as contexts or stimuli (Sutton and Barto, 2018). Actions are decisions to transition between states (i.e. a decision to move from one location to another). Values quantify the mean expected reward to be obtained under a given state or action. Rewards are scalars usually given at certain spatial locations, mimicking goal locations in navigation tasks. The value function can be a function of (i.e. dependent upon) the state alone, in which case it refers to the discounted total amount of reward that an agent can expect to receive in the future from a current state s at time t .

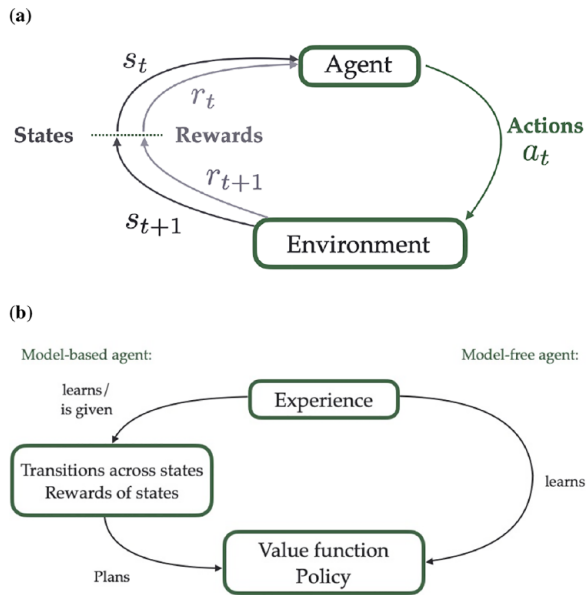


Figure 2. Basic principles of reinforcement learning (RL). (a) Key components of RL models. An agent in the state s_t (which in spatial context often corresponds to a specific location in the environment) associated with the reward r_t takes the action a_t to move from one state to another within its environment. Depending on the available routes and on the rewards in the environment, this action leads to the reception of a potential reward r_{t+1} in the subsequent state (or location) s_{t+1} . (b) Model-free versus model-based approaches in RL. In model-free RL (right), an agent learns the values of the states on the fly, that is, by trial and error, and adjusts its behaviour accordingly (in order to maximise its expected rewards). In model-based RL (left), the agent learns or is given the transition probabilities between states within an environment and the rewards associated with the states (a ‘model’ of the environment), and uses this information to plan ahead and select the most successful trajectory.

Alternatively, the value can also refer to the state–action pair, in which case it refers to the value of taking a particular action at a certain state. The value function is given by

$$V(s) = E[r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \dots | s_t = s] \quad (1)$$

In equation (1), the value of state s , $V(s)$, is computed by summing all future rewards r_j that will be received at time j , $j \geq 0$ discounted by a factor γ , which quantifies the extent to which immediate rewards are favoured compared to delayed ones of the same magnitude. $E[\cdot]$ refers to the expectation, which sums the possible rewards depending on their associated probability. A policy is a probability distribution over the set of possible actions. This defines which actions are more likely to be chosen in a certain location and has to be learned in order to maximise the *value function* (i.e. to maximise the expected amount of reward).

Mathematically, the problem can be represented as a Markov decision process, equipped with transition probabilities between states that shapes the way actions change states and a reward function that maps states to reward (see, for more details, Howard, 1960). In a spatial navigation context, the transition probabilities typically depend on the spatial structure of the

environment, and the latter can also vary temporally in certain contexts, as routes might open or close on particular occasions, for example. The probabilities of transition between states/locations in an environment and the rewards available at every location form a *model* of the environment.

In model-free RL (Figure 2(b), right), the *model* is unknown and the agent must discover its environment, and associated rewards, and learn how to optimise behaviour on the fly, through trial and error. Conversely, in model-based RL approaches (Figure 2(b), left), the agent has access to the *model*, from which a tree of possible chains of actions and states can be built and used for planning. In this way, the best possible chain of actions can be defined, for example, using Dynamic Programming, which selects at every location the optimal action, using one-step transition probabilities (Sutton and Barto, 2018).

We can assess whether humans and animals use model-free or model-based strategies by comparing their performance to both types of agent on a two-step decision task (Da Silva and Hare, 2019; Daw et al., 2011; Miller et al., 2017). In this task, participants first choose between two states, both of which afterwards lead to final states with different probabilities, making one transition ‘rare’ and the other ‘common’. The final states have unbalanced reward probability distributions (see, for example, Figure 1(a) in Miller et al., 2017, for a diagram of the task). Investigating how subjects adjust to rare transition outcomes indicates whether they have access to the *model* or not. A model-free agent will adjust its behaviour only based on the outcome, whereas a model-based agent will adjust also according to the probability of this transition. Both humans and animals show behavioural correlates of model-based and model-free agents (Daw et al., 2011; Gershman, 2017; Keramati et al., 2011; Miller et al., 2017; Yin and Knowlton, 2006).

Model-based approaches require the calculation and storage of the transition probability matrix and tree-search computations (Huys et al., 2013). As the number of states can be very high, depending on the complexity of the problem and the precision required, model-based methods are usually computationally costly (Huys et al., 2013). However, as they contain exhaustive information about the available routes between states, they are more flexible towards changing goal locations than model-free approaches (Keramati et al., 2011). A study of spatial navigation in human participants showed that, although paths to the goal were shorter, choice times were higher in trials when the behaviour matches that of a model-based agent compared to trials where it matches that of a model-free agent (Angraini et al., 2018). In studies involving rats in a T-maze, vicarious trial and error (VTE) behaviour, short pauses that rats make at decision points, tend to get shorter with repetitive exposure to the same goal location (Redish, 2016). Experimental studies suggest that VTE behaviour reflects simulations of scenarios of future trajectories in order to make a decision (Redish, 2016), which would correspond to a model-based approach of task solving (Penny et al., 2013; Pezzulo et al., 2017). This suggests that model-based strategies require more processing time than model-free strategies, which is thought to represent ‘planning’ time (Keramati et al., 2011).

The control of behaviour could be coordinated between model-free and model-based systems, either depending on uncertainty (Daw et al., 2005), depending on a trade-off between the cost of engaging in complex computations and the associated improvement in the value of decisions (Pezzulo et al., 2013), or depending on how well the different systems perform on a task

(Dollé et al., 2018). Moreover, the state and action representations that enable solution of a task seem to dynamically evolve on the timescale of learning in order to adjust to the task requirements (Dezfouli and Balleine, 2019). When the task gives an illusion of determinism – for example, when the task is overtrained – the neural representations and the behaviour shift from model-based, purposeful, behaviour, to habitual behaviour, which is faster but less flexible to any change (Smith and Graybiel, 2013). When the situation is inherently stochastic, for example, when the task evolves to incorporate more steps and complexity, the neural representations and behaviour evolve to incorporate the multi-step dependencies, and simultaneously prune the tree of possible outcomes depending on the most likely scenarios (Tomov et al., 2020). The adaption of neural representations to task demand suggests that a continuum of state and action representations for behavioural control, between the two extremes of model-based and model-free systems, exists in the brain (Dezfouli and Balleine, 2019).

The rapidity with which rats adjust to a changing goal location in the DMP task (one exposure only, as discussed in section ‘Flexible hippocampal spatial navigation’) indicates that some kind of ‘model’ is being used to enable adaptive route selection. In fact, a model-based approach has recently been proposed to solve a range of spatial learning tasks including the watermaze DMP task (Dollé et al., 2018). Dollé et al. (2018) investigated the interaction between model-free and model-based strategies using a model-free controller to gate interactions between the two. The controller learns to select one of the two strategies depending on the reward in the current task. The model reproduced the flexibility towards new goal locations in the watermaze DMP task, through the gating mechanism, which switched to the model-based strategy for this particular task.

Generally, current model-based approaches, including the one proposed by Dollé et al. (2018), have several limitations in accounting for watermaze DMP task performance in a neuropsychologically realistic way. First, unlike what would be expected based on model-based mechanisms, rats do not reach optimal trajectories in the DMP task, as reflected by the observation that escape latencies on trials 2 to 4 remain higher than would typically be observed following incremental place learning in the watermaze (compare Morris et al. (1986), Steele and Morris (1999) and Bast et al. (2009)) and than would be expected from a model-based agent (Sutton and Barto, 2018). Findings in humans by Anggraini et al. (2018) suggest that participants that used more model-based approaches more often took the shortest path to goal locations. Second, classical model-based approaches are currently mostly implemented in discrete state space, although they can be approximated to continuous spaces (Jong and Stone, 2007). The size of the graph to model the environment, requiring a high number of states for fine discretisation to mimic continuity, and the width and depth of the trees to search at possible decision points (e.g. at the start location) that would be required to account for such behaviours (every possible trajectory) suggest that a long planning time would be required at a start of every trial (Keramati et al., 2011), which does not fit with the behaviour of rats and human participants, who take off virtually immediately at the start of the second trial to the new goal location on the DMP task in the watermaze and virtual maze, respectively (Buckley and Bast, 2018; Steele and Morris, 1999). Dollé et al. (2018) overcome this problem using an approximation: the chosen trajectory between the current position and the goal

location is in fact the trajectory between their respective closest nodes within the graph. Overall, this suggests that the control in the watermaze DMP task cannot be explained by a model-free RL mechanism alone, but also is unlikely to be fully model-based.

Lying in between model-free and model-based approaches, the successor representation (SR; Dayan, 1993) enables more flexibility than model-free computational approaches, but without the heavy computational requirement of a model-based agent (Ducarouge and Sigaud, 2017). In the SR, the link between two states depends on how many times the agent can expect to visit one state when starting from the other in the future. It is therefore a predictive representation of space occupancy. Properties of place cell firing, such as shaping of the activity profile by changes in the environment, have been linked to key features of the SR (Gershman, 2018; Stachenfeld et al., 2017). The SR can be computed not only from the transition probability matrix, but also by online learning, usually through counting the occupancy of states (Dayan, 1993; Russek et al., 2017). Connecting place cells, using this representation, within an attractor network allows one to generate trajectories from any starting location to any goal location within a maze (Corneil and Gerstner, 2015). The SR can be adapted to a continuous state representation (Barreto et al., 2017; Jong and Stone, 2007). However, the SR still represents a complex state representation, since the size of the SR matrix is similar to the size of the model-based representation. In the following section, we present two minimal extensions to a model-free architecture that enable flexibility. We will provide an in-depth discussion of the underlying computations and of their possible neurobiological substrates.

A model-free agent using an actor–critic architecture

An actor–critic network model for incremental learning

Learning through temporal-difference error. Temporal-difference (TD) learning refers to a class of model-free RL methods that improve the estimate of the value function using successive comparison of its value, called the TD error (this approach is commonly referred to as ‘bootstrapping’). In traditional TD learning, the agent follows a fixed policy and discovers how good this policy is through the computation of the value function (Sutton and Barto, 2018). Conversely, in an actor–critic learning architecture, an agent explores the environment and progressively forms a policy that maximises expected reward using a TD error. Interactions with the environment allow simultaneous sampling of the rewards (obtained at certain locations) and of the effect of the policy (when there is no reward, the effect of a policy can be judged depending on the difference in value between two consecutive locations), thereby appraising predictions of values and actions, so that both can be updated accordingly. The ‘actor’ refers to the part of the architecture that learns and executes the policy, and the ‘critic’ to the part that estimates the value function (Figure 3(a)).

Foster et al. (2000) proposed an actor–critic framework to perform spatial learning in a watermaze equivalent (Figure 3(a)). The agent’s location is represented through a network of units mimicking hippocampal place cells, which have a Gaussian type of activity around their preferred location (the further the agent is

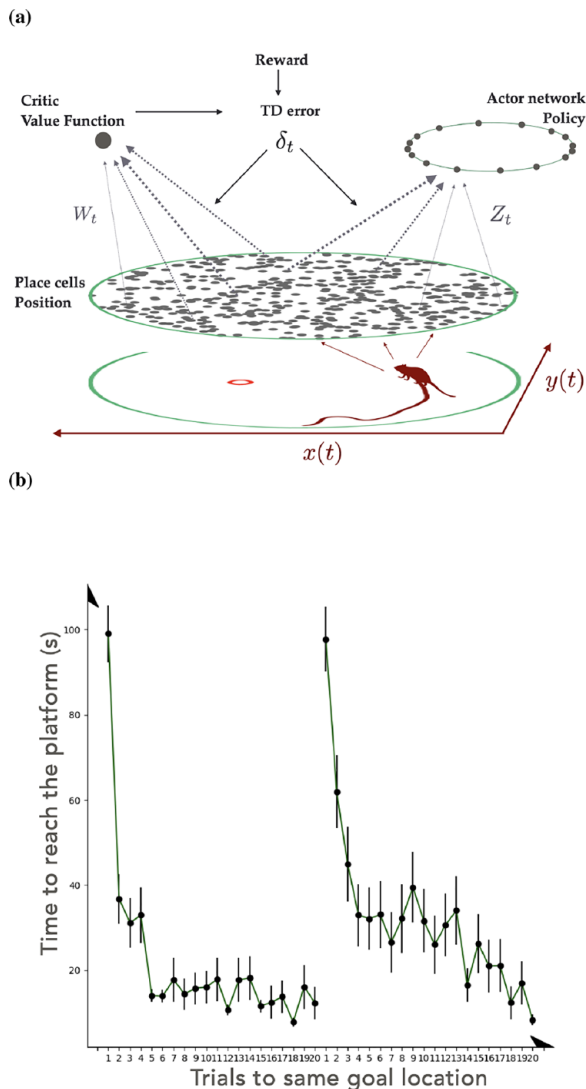


Figure 3. (a) Classical actor–critic architecture for a temporal-difference (TD) agent learning to solve a spatial navigation problem in the watermaze, as proposed by Foster et al. (2000). The state (location of the agent $(x(t), y(t))$) is encoded within a neural network (in this case, the units mimic place cells in the hippocampus). State information is fed into an action network, which computes the best direction to go next, and to a critic network that computes the value of the states encountered. The difference in critic activity along with the reception or not of the reward (given at the goal location) are used to compute the TD error δ_t , such that successful moves (that lead to a positive TD error) are more likely to be taken again in the future and less likely otherwise. Simultaneously, the critic’s estimation of the value function is adjusted in order to be more accurate. These updates occur through changing the critic and actor weights, respectively W_t and Z_t . The goal location, marked as a circle within the maze, is the only location in which a reward is given. (b) Performance of the agent, obtained by implementing the model from Foster et al. (2000). The time that the agent requires to get to the goal (‘Latencies’, vertical axis) reduces with trials (horizontal axis) and reaches almost a minimum (after trial 5). When the goal changes (on trial 20), the agent takes a very long time to adapt to this new goal location.

from the place cell’s preferred location, the less active the unit will be). Each place cell projects to an actor network and a critic unit, or network in other related models, for example, in Frémaux et al. (2013), through plastic connections (respectively denoted Z_t and W_t in Figure 3(a)).

The actor network activity defines which action is optimal, with each unit in the network coding for motion in a particular direction, together covering a 360° angle. In Foster et al. (2000), this angular direction is quantised, whereas in Frémaux et al. (2013), using a more detailed spiking neuron model, the action network codes for continuous movement directions (although finely quantised in numerical simulations). The action is chosen according to the activity of the actor, so that directions corresponding to more active cells are prioritised. In Foster et al. (2000), the action corresponding to the maximum of a softmax probability distribution of the activities is selected. In Frémaux et al. (2013), the activities of the units of the network are used as weights to sum all possible directions, giving rise to a mean vector that defines the chosen direction. If the motion leads to the goal location, a reward is obtained. In the case of computational models of the watermaze task, the reward is only delivered when the agent is within a small circle representing the platform (see Figure 3(a)). This reward information, encoded within the environment is used along with the difference in the successive critic activities to compute the TD error δ_t , via

$$\delta_t = r(s_{t+1}) + \gamma V(s_{t+1}) - V(s_t) \quad (2)$$

In equation (2), the reward $r(s_{t+1})$ along with the discounted value of the new state $\gamma V(s_{t+1})$ is compared to the prior prediction $V(s_t)$. The TD error contains two pieces of information: it reflects both how good is the decision that was just taken and the accuracy of the critic at estimating the value of the state.

To reduce the TD error, the model simultaneously updates the connections weights Z_t and W_t via defined learning rules in order to improve the actor and the critic. The learning rule updates the connection weights according to the current error and the place cell activity, such that the probability of taking the same decision in the current state/location increases if it leads to a new position that has a higher value, and decreases otherwise (see Doya, (2000), for learning rules). Using the model proposed by Foster et al. (2000), we can reproduce their finding that within a few trials, the agent reaches the goal using an almost optimal path, reflected by low latencies to reach the goal, in a watermaze-like environment (Figure 3(b)). The full model can be found in Supplemental Appendix A (available at: <https://journals.sagepub.com/doi/suppl/10.1177/2398212820975634>), ‘Actor–critic equations, based on Foster et al. (2000)’.

Important variables for learning: discount factor for value propagation and spatial scale of place cells representations for experience generalisation. The previous section illustrated how place cells can be integrated into a network for spatial navigation through the association of values and actions within an RL framework. A benefit of this approach is that (1) it allows for relatively fast learning, within only a few trials, agents reach ‘short’ average escape latencies and (2) the agent obtains information about which action could lead to the goal from variable and distal start positions. These two properties rely on two major

components that enable learning and influence the learning speed (i.e. how fast latencies reduce).

First, the TD error allows to ‘back-propagate’ value information from successive states, with the speed of this back propagation modulated by the discount factor γ . The update of the state value $V(s_t)$ and of the policy at state s_t , after having moved to the new state s_{t+1} , depends on the TD error. Let us consider the latter, defined by equation (2): the TD error is the difference between the received reward and the discounted value at the next state (given by $r(s_{t+1}) + \gamma V(s_{t+1})$, the first two terms in equation (2)), and the value at the current state ($V(s_t)$, the last term in equation (2)). Note that the update of the state value is performed ‘in the future’ for ‘the past’: the value underlying the decision taken at time t will only be updated at time $t+1$. Moreover, the extent to which the future is taken into account is modulated by the parameter γ .

Let us consider the extreme cases. If $\gamma = 0$, the only update takes place when the reward is found, and the location updated is the one immediately prior to the reward. All the other locations, which do not precede the reward reception, will never be associated with a non-zero value. If $\gamma = 1$, the state value $V(s_t)$ will be updated until it is equal to $V(s_{t+1})$. This leads to a constant value function over the maze (i.e. all locations have the same value). In both cases, the value function being uniform, the actor computes all actions as equally good (except very near the goal in the case where $\gamma = 0$), as it moves through the environment. Therefore, only intermediate values of γ allow the model to learn, and its value defines how fast it learns. Ideally, one wants to adjust the discount factor γ in order to maximise the slope of the value function, so that the policy is ‘concentrated’ on the optimal choice, and to obtain a uniform slope across the space, as this guarantees that the agent has good information on which to base its decision from any location within the environment.

Second, the spatial scale of the place cell representation, determined by the width of these neurons’ place fields, strongly affects speed and precision of place learning. The state representation through place cells enables the generalisation of learning from a single experience across states, that is, to update information on many locations based on the experience within one particular location. Every update amends the value and policy for all states depending on the current place cell activity, such that more distal locations are less concerned by the update than proximal ones. The spatial reach of a particular update increases with the width of the place cell activity profiles. This process speeds up learning, because when the agent encounters a location with similar place cell representation to those already encountered, the prior experiences have already shaped the current policy and value function and can be used to inform subsequent actions.

Let us consider the extreme cases. If the width is very small, the agent cannot generalise enough from experience, and this considerably slows down learning, as the agent must comprehensively search the environment in order to learn. At the opposite extreme, if the activity profile is very wide, generalisation occurs where it is not appropriate: for example, at opposite ends of the goal, when the best actions to choose would be opposite to each other, as at the North end it would be best to go South, whereas at the South end, it would be best to go North. The optimal width, therefore, lies in a trade-off between speed of learning and precision of knowledge: it should be scaled to the size of the environment in order to speed up learning and is constrained by the size

of the goal. Optimising these parameters allows one to reduce the number of trials to obtain good performance. Supplemental Appendix B (available at: <https://journals.sagepub.com/doi/suppl/10.1177/2398212820975634>), ‘Actor–critic component: effect of changing the place cell activity width and the discount factor on incremental learning towards the goal location’, shows how changing these parameters affects learning.

Along the hippocampal longitudinal axis, places are represented over a continuous range of spatial scales, with the width of place cell activity profiles gradually increasing from the dorsal (also known as septal) towards the ventral (also known as temporal) end of the hippocampus in rats (Kjelstrup et al., 2008). A recent RL model suggests that smaller scales of representation would support the generation of optimal path length, whereas larger scales would enable faster learning, defining a trade-off between path optimality and speed of learning (Sclaidorovich et al., 2020). In Figure 6 of Supplemental Appendix B (available at: <https://journals.sagepub.com/doi/suppl/10.1177/2398212820975634>), ‘Actor–critic component: effect of changing the place cell activity width and the discount factor on incremental learning towards the goal location’, for the actor–critic model, we also see that a wide activity profile of place cells leads to suboptimal routes, characterised by escape latencies that stay high.

Bast et al. (2009) found that the intermediate hippocampus is critical to maintain DMP performance in the watermaze, particularly search preference. Moreover, the trajectories used by rats in the watermaze DMP task are suboptimal, that is, path lengths are higher, compared to the incremental learning task (compare results in Morris et al. (1990), Steele and Morris (1999) and Bast et al. (2009)). These findings may partly reflect that place neurons in the intermediate hippocampus, which show place fields of an intermediate width and, thereby, may deliver a trade-off between fast and precise learning, are particularly important for navigation performance during the first few trials of learning a new place, as on the DMP task. Another potential explanation for the importance of the intermediate hippocampus is that this region combines sufficiently accurate place representations, provided by place cells with intermediate-width place fields, with strong connectivity to prefrontal and subcortical sites that support use of these place representations for navigation (Bast, 2011; Bast et al., 2009), including striatal RL mechanisms (Humphries and Prescott, 2010). With incremental learning of a goal location, spatial navigation can become more precise, with path lengths getting shorter and search preference values increasing (Bast et al., 2009). Interestingly, the model by Sclaidorovich et al. (2020) suggests that such precise incremental place learning performance may be particularly dependent on narrow place fields, which are shown by place cells in the dorsal hippocampus (Kjelstrup et al., 2008). This may help to understand why incremental place learning performance has been found to be particularly dependent on the dorsal hippocampus (Moser et al., 1995).

Using eligibility traces allows to update past decisions according to the current experience. The particular actor–critic implementation proposed by Foster et al. (2000) and described above involves a one-step update: only the value and policy of the state that the agent just left are updated. The weight associated to the value of the following state for any update depends on γ , as described in the previous section. However, past decisions sometimes affect current situations, and one-step

updates only improve the last choice and estimate. This can be addressed by incorporating past decisions when performing the current update, weighted according to an *eligibility trace* (Sutton and Barto, 2018). Eligibility traces keep a record of how much past decisions influence the current situation. This makes possible to update value functions and policies from previous states of the same trajectory according to the current outcome.

The most commonly known example of the use of eligibility traces is TD(λ) learning, which updates the value and policy of previous states within the same trajectory according to the outcome observed after a certain subsequent period, weighted by a decay rate λ . λ refers to how far in the past the current situation affects previous states' value and policy. One extreme is TD(0) in which the only step updated is the state the agent just left, as described in the section 'An actor-critic network model for incremental learning'. When λ increases towards 1, more events within the trajectory are taken into account, and the method becomes more reminiscent of a Monte-Carlo approach (Sutton and Barto, 2018), where all the states and actions encountered during the trial are updated at every step. In Scleidorovich et al. (2020), the authors show that the optimal value of λ depends on the width of the place cell activity distributions: for wide place fields, adding eligibility traces does not speed up learning much (i.e. reducing the number of trials needed to reach asymptotic performance), but it does for narrower place fields.

Striatal and dopaminergic mechanisms as candidate substrates for the actor and critic components. In the RL literature, the ventral striatum is often considered to play the role of the 'critic' (Humphries and Prescott, 2010; Khamassi and Humphries, 2012; O'Doherty et al., 2004; Van Der Meer and Redish, 2011). The firing of neurons in the ventral striatum ramps up when rats approach a goal location in a T-maze (Van Der Meer and Redish, 2011), consistent with the critic activity representing the value function in actor-critic models of spatial navigation (Foster et al., 2000; Frémaux et al., 2013). Striatal activity also correlates with action selection (Kimchi and Laubach, 2009) and with action-specific reward values (Roesch et al., 2009; Samejima et al., 2005).

In line with the architecture proposed in the model by Foster et al. (2000) (Figure 3(a)), there are hippocampal projections to the ventral and medial dorsal striatum (Groenewegen et al., 1987; Humphries and Prescott, 2010). Studies combining watermaze testing with manipulations of ventral and medial dorsal striatum support the notion that these regions are required for spatial navigation. Lesions of the ventral striatum (Annett et al., 1989) and of the medial dorsal striatum (Devan and White, 1999) have been reported to impair spatial navigation on the incremental place learning task. In addition, crossed unilateral lesions disconnecting hippocampus and medial dorsal striatum also impair incremental place learning performance, suggesting hippocampo-striatal interactions are required (Devan and White, 1999). To our knowledge, it has not been tested experimentally if there is a dichotomy between 'actor' and 'critic'. The experimental evidence outlined above is consistent with both actor and critic roles of the striatum (Van Der Meer and Redish, 2011), but whether distinct or the same striatal neurons or regions act as actor and critic needs to be addressed.

Li and Daw (2011) address a related dichotomy in a study on human participants who have to choose between two arms associated with reward probabilities on a bandit task. The participants are

given the outcomes of their decision after their choice: namely, how much they win and how much they would have won if they had selected the other arm. Li and Daw (2011) compare two ways of updating the weights which determine which arm to choose: one compares the reward to the predicted value ('value update'), and the other one compares the reward to the forgone reward ('policy update'). They show that striatal BOLD activity correlates more with 'policy' than 'value' update, and correlates positively with the chosen reward and negatively with the reward that was not chosen. They also show correlation with a value-based decision variable, the difference between the action value of the chosen and the not chosen arm. The translation of their analysis to spatial navigation in the watermaze is not straightforward. First, in the two armed bandit task, there are no states, but only actions. Although the design of the analysis by Li and Daw (2011) allows one to disentangle the rewards from the predicted values, it does not allow one to separate action from state value in a spatial navigation context, if it makes sense at all to separate the two. In spatial navigation, as states can be passed through to reach any goal, it seems to be more efficient not to separate actions and values. However, it is interesting to see an experimental set-up aimed at testing such a dichotomy. Perhaps, an architecture such as SARSA (State-Action-Reward-State-Action, Sutton and Barto, 2018), in which the values of a state-action pair are learned instead of the values of states only, could be considered, as it unites the actor and critic computation within the same network.

Phasic dopamine release from dopaminergic midbrain projections to the striatum has long been suggested to reflect reward prediction errors (Glimcher, 2011; Schultz et al., 1997), which correspond to the TD error in the model in Figure 3(a), and dopamine release in the striatum shapes action selection (Gerfen and Surmeier, 2011; Humphries et al., 2012; Morris et al., 2010). Direct optogenetic manipulation of striatal neurons expressing dopamine receptors modified decisions (Tai et al., 2012), consistent with the actor activity in actor-critic models of spatial navigation (Foster et al., 2000; Frémaux et al., 2013). Moreover, 6-hydroxydopamine lesions to the striatum, depleting striatal dopamine (and, although to a lesser extent, also dopamine in other regions, including hippocampus) impaired spatial navigation on the incremental place learning task in the watermaze (Braun et al., 2012). These findings suggest that aspects of the dopaminergic influence on striatal activity could be consistent with the modulation of action selection by the TD errors in an actor-critic architecture.

However, although there is long-term potentiation (LTP) like synaptic plasticity at hippocampo-ventral striatal connections, consistent with the plastic connections between the place cell network and the critic and actor in the model by Foster et al. (2000), a recent study failed to provide evidence that this plasticity depends on dopamine (LeGates et al., 2018). Absence of dopamine modulation of hippocampo-striatal plasticity contrasts with the suggested modulation of connections between place cell representations and the critic and actor components by the TD error signal in the RL model. Thus, currently available evidence fails to support one key feature of the architecture described in section 'An actor-critic network model for incremental learning' (Foster et al., 2000).

Requirement of hippocampal plasticity. In the implementation of the model described above, plasticity takes place within

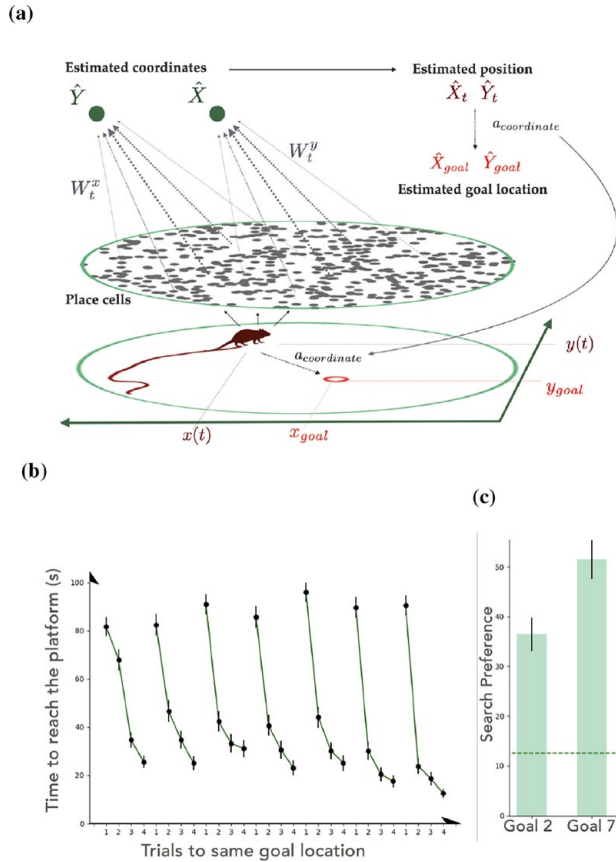


Figure 4. (a) Architecture of the coordinate-based navigation system, which was added to the actor-critic system shown in Figure 3(a) to reproduce accurate spatial navigation based on one-trial place learning, as observed in the watermaze DMP task (Foster et al., 2000). Place cells are linked to coordinate estimators through plastic connections W_t^x, W_t^y . The estimated coordinates \hat{X}, \hat{Y} are used to compare the estimated location of the goal $\hat{X}_{goal}, \hat{Y}_{goal}$ to the agent estimated location \hat{X}_t, \hat{Y}_t in order to form a vector towards the goal, that is being followed when choosing the ‘coordinate action’ $a_{coordinate}$. The new action a_{good} is integrated into the actor network described in Figure 3(a). (b, c) Performance of the extended model using coordinate-based navigation. (b) Escape latencies of the agent when the goal location is changed every four trials, mimicking the watermaze DMP task. (c) ‘Search preference’ for the area surrounding the goal location, as reflected by the percentage of time the agent spends in an area centred on the goal location when the second trial to a new goal location is run as probe trial, with the goal removed (stippled line indicates percentage of time spent in the correct zone by chance, that is, if the agent had no preference for any particular area), computed for the second and the seventh goal locations. One-trial learning of the new goal location is reflected by the marked latency reduction from trial 1 to trial 2 to a new goal location (without interference between successive goal locations) and by the marked search preference for the new goal location when trial 2 is run as probe. The data in (b) were obtained by computing the model in (Foster et al., 2000) and the data in (c) by adapting the model in order to reproduce search preference measures when trial 2 was run as a probe trial. The increase in search preference observed between the second and seventh goal location is addressed in the section ‘Limitations of the model in reproducing DMP behaviour in rats and humans’.

the feedforward connections from the place cell network modelling the hippocampus and the actor and critic networks that, as discussed above, could correspond to parts of the striatum. The model does not capture the finding that hippocampal NMDA receptor-dependent plasticity is required for incremental place learning in the watermaze if rats have not been pretrained on the task (Morris et al., 1986, 1989; Nakazawa et al., 2004).

The agent is less flexible than animals in adapting to changing goal locations. When the goal changes (on trial 20 in Figure 3(b)), the agent takes many trials to adapt and takes more than 10 trials to reach asymptotic performance levels (also see Figure 4(a) in Foster et al., 2000). The high accuracy, but limited flexibility with overtraining, are well-known features of TD RL methods (e.g. discussed in Botvinick et al., (2019); Gershman, (2017); Gershman et al., (2014); Sutton and Barto, (2018)). These cached methods have been proposed to account for the progressive development of habitual behaviours (Balleine, 2019). TD learning is essentially an implementation of Thorndike’s law of effect (Thorndike, 1927), which increases the probability of reproducing an action if it is positively rewarded. In the RL model discussed above (Figure 3(a)), a particular location, represented by activities of place cells with overlapping place fields, is associated with only one ‘preferred action’, due to the unique weights that need to be fully relearned when the goal changes. Therefore, the way actions and states are linked only allows navigation to one specific goal location.

The model produces a general control mechanism, that, in this example, makes it possible to generate trajectories to a particular goal location. This control mechanism could be integrated within an architecture that allows more flexibility, for example, the goal location may be represented by different means than a unique value function computed via slow and incremental steps from visits to single goal locations. The next section considers how the RL model of Figure 3(a) can be used, along with a uniform representation of both the agent and the goal location, to reproduce the flexibility shown by rats and humans towards changing goal locations in the watermaze DMP task.

Map-like representation of locations for goal-directed trajectories

The RL architecture discussed above (Figure 3(a)) cannot reproduce rapid learning of a new location as observed in the DMP watermaze task, but instead there is substantial interference between successive goal locations, with latencies increasing across goal locations and only a gradual small decrease in latencies across the four trials to the same goal location (see Foster et al., (2000), Figure 4(b)). To reproduce flexible spatial navigation based on one-trial place learning as observed on the DMP task, Foster et al. (2000) proposed to incorporate a coordinate system into their original actor-critic architecture (Figure 4(a)). This coordinate system is composed of two additional cells \hat{X} and \hat{Y} that learn to estimate the real coordinates x and y throughout the maze. These cells receive input from the place cell network through plastic connections W_t^x, W_t^y . The connections evolve dependent on a TD error that represents the difference between the displacement estimated from the coordinate cells and the real displacement of the agent. The weights between

place cells and the coordinate cells are reduced if the estimated displacement is higher than the actual displacement and increased if it is lower, so that the estimated coordinates progressively become consistent with the real coordinates (see Figure 8 in Supplemental Appendix D (available at: <https://journals.sagepub.com/doi/suppl/10.1177/2398212820975634>), ‘Hierarchical model – equations’).

Foster et al. (2000) added an additional action to the set of actions already available. Instead of defining movement in specific allocentric directions, as the other action cells do, going North, East and so on, that we will refer to as ‘allocentric direction cells’, the coordinate action cell a_{coord} points the agent in the direction of the estimated goal location. The estimates of x and y are used to compare the agent’s estimated position \hat{X}_t, \hat{Y}_t to the estimated goal location $\hat{X}_{goal}, \hat{Y}_{goal}$ (which is stored after the first trial of every day) in order to form a vector leading to the estimated goal location (Figure 4(a)).

The agent very quickly adapts to new goal locations, reproducing performance similar to rats and humans on the watermaze (Figure 1) and virtual (Buckley and Bast, 2018) DMP task, respectively. Using the model of Foster et al. (2000), we can replicate their finding that the model reproduces the characteristic pattern of latencies shown by rats and human on the DMP task, that is, a marked reduction from trial 1 to 2 to a new goal location and no interference between successive goal location (Figure 4(b)). Moreover, extending the findings of Foster et al. (2000), we find that the model also reproduces markedly above-chance search preference for the vicinity of the goal location when trial 2 to a new goal location is run as probe trial where the platform is removed (Figure 4(c)).

Limitations of the model in reproducing DMP behaviour in rats and humans. The ‘coordinate’ approach relies on computational ‘tricks’ that are required to make the approach work, but for which plausible neurobiological substrates remain to be identified. Early in training, movement of the agent is based on activity of the ‘allocentric direction cells’, which are used to lead the exploration of the environment. This exploratory phase allows learning of the coordinates. As the estimated coordinates \hat{X}, \hat{Y} become more consistent with the real coordinates, the coordinate action a_{coord} becomes more reliable, as it will always lead the agent in the direction of the goal. During the first trial to a new goal location, the coordinate action cell encodes random displacement, until the goal is found and its estimated location is stored. During this trial, the coordinate action is not reinforced, a trick that prevents its devaluation. On the subsequent trials, the coordinate action encodes the displacement towards the stored estimated goal location (as described before) and is reinforced. Therefore, the probability of choosing the coordinate action gradually becomes one, and it comes to be the only action followed.

One consequence of this is that, unlike in rats and people, the agent’s performance both in terms of latency reduction and in terms of search preference gradually improves across successive new goal locations (see Figure 4(b) and (c)). The gradual improvement of latency reductions and search preferences contrasts with behaviour shown by rats (Figure 1 and see Figure 3(b) and (c) in Bast et al. (2009) for search preference across days) and human participants (Buckley and Bast, 2018). More specifically, in the model, latency reductions from trial 1 to 2 gradually increase across successive new goal locations, and the latencies on trial 2 to 4 to a new goal location gradually decrease (Figure 4(b)). In contrast, rats basically reach asymptotic performance levels with

no systematic increases in latency reductions from trial 1 to 2 or decreases in latency values on trial 2 to 4 after a few successive goal locations; in the example shown in Figure 1(b), asymptotic performance levels are reached from about day 4. It should also be noted that the overall latency reductions across the first few locations in rats is likely to mainly reflect procedural learning, with rats learning that they cannot escape by climbing the wall of the pool or by diving. Human participants on the virtual DMP task, who do not need to learn the task requirements because they receive task instructions, show virtually asymptotic latency and path lengths values from the first new goal locations, with hardly any improvements across successive goal locations (Buckley and Bast, 2018). Moreover, search preference for the correct location substantially increases across successive probe trials in the model (Figure 4(c)), whereas in rats and humans, search preference remains stable across successive new goal locations on the DMP task (Bast et al., 2009; Buckley and Bast, 2018).

In addition, the random search during trial 1 in the model is inconsistent with the finding that rats on the DMP task (but not human participants (Buckley and Bast, 2018)) tend to go towards the previous goal location on trial 1 to a new goal location (Pearce et al., 1998; Steele and Morris, 1999; and our own unpublished observations); in addition, both rats and human participants show systematic search patterns on trial 1 (Buckley and Bast, 2018; Gehring et al., 2015). The random search in the model during trial 1 leads to consistently and similarly high trial 1 latencies (Figure 4(b)). In contrast, in rats, trial 1 latencies are more variable (Figure 1(b)); this partly reflects procedural learning across the first few new goal locations, which results in reductions of trial 1 latencies, and the different spatial relationship between the start location and the previous and current goal location affecting trial 1 latencies (e.g. if the current goal location lies on the path from the start location to the previous goal location, rats are more likely to ‘bump’ into the current goal location on trial 1, leading to lower trial 1 latencies). The adjustment of the policy when the predicted goal is not encountered is not addressed in the current approach, a point that section ‘Hierarchical control to flexibly adjust to task requirements’ will address.

The model’s actor–critic component and striatal contributions to DMP performance. Given the association of actor–critic mechanisms with the striatum (Joel et al., 2002; Khamassi and Humphries, 2012; O’Doherty et al., 2004; Van Der Meer and Redish, 2011), the actor–critic component in the model is consistent with our recent findings that the striatum is associated with rapid place learning performance on the DMP task. More specifically, using functional inhibition of the ventral striatum in rats, we have shown that the ventral striatum is required for one-trial place learning performance on the watermaze DMP task (Seaton, 2019); moreover, using high-density electroencephalogram (EEG) recordings with source reconstruction in human participants, we found that theta oscillations in a circuit including both temporal lobe and striatum are associated with one-trial place learning performance on the virtual DMP task (Bauer et al., 2020).

The model suggests that, after a few trials, once the action probability for the coordinate action has reached the value 1, the movement is predefined by following a vector pointing to the goal location. The critic becomes inconsistent, as the action now does not follow the gradient of the value function anymore, and, therefore, there is no control over the behaviour by the TD error. The continued association of the striatum with DMP performance, beyond the first few trials, is consistent with the role of the

striatum as the ‘actor’ (Van Der Meer and Redish, 2011), and the model would suggest that the striatum reads out estimated locations and computes a vector towards the estimated goal location.

Neural substrates of the goal representation and hippocampal plasticity required for rapid learning of new goal locations. Notwithstanding some limitations, findings with this model support the important idea that, embedded within a model-free RL framework, a map-like representation of locations within an environment may allow computations by the agent to produce efficient navigation to new goal locations within as little as one trial. This idea is also present in recently proposed agents that are capable of flexible spatial navigation based on a RL system complemented by path integration mechanisms to derive a grid-like map of the environment (which resembles entorhinal grid cell representations) that can be used to compute trajectories from the agent’s location to the goal (Banino et al., 2018) and has also led the watermaze DMP task to be solved using a graph search algorithm in Dollé et al. (2018). The findings of goal-vector cells in the bat hippocampus (Sarel et al., 2017) and of ‘predictive reward place cells’ in mouse hippocampus (Gauthier and Tank, 2018) support the idea implemented in the model that consistency of representations – unified representations of goals and locations across tasks and environment – could help goal-directed navigation. In particular, egocentric boundary encoding neurons have been found in the striatum of rats, although in the dorso-medial striatum (Hinman et al., 2019). As rats navigate in the watermaze using surrounding cues, these cells could inform striatal navigation in the DMP task (Bicanski and Burgess, 2020).

In the extension to the classical TD architecture (Foster et al., 2000), the encounter with a new goal location does not involve a change in place cell representation, and the formation of the memory of the new goal location is not addressed. Experimental evidence suggests that a goal representation could lie within hippocampal representations themselves (Gauthier and Tank, 2018; Hok et al., 2007; McKenzie et al., 2013; Poucet and Hok, 2017). McKenzie et al. (2013) studied hippocampal CA1 representations during learning of new goal locations in an environment where previous places were already associated with goals. They showed that neurons coding for existing goals would also encode new goal locations and that these representations progressively separate with repetitive learning of the new goal location, but maintain an overlap of representations between all goal locations. Moreover, Hok et al. (2007) observed an increase in firing rate around goal locations outside of place cells’ main firing field, and Dupret et al. (2010) showed that learning of new goal locations by rats in a food-reinforced dry-land DMP task is associated with an increase in the number of CA1 neurons that have a place field around the goal location. Furthermore, Dupret et al. (2010) showed that both this accumulation of place fields around the goal location and rapid learning of new goal locations is disrupted by systemic NMDA receptor blockade. These findings suggest that goal representation can be embedded within the hippocampus, that new goal locations are represented within similar networks as previous goal locations, and that the hippocampal remapping emerging from new goal locations is linked to behavioural performance and may depend on NMDA receptor-mediated synaptic plasticity.

In line with this suggestion, studies, combining intra-hippocampal infusion of an NMDA receptor antagonist with behavioural testing and electrophysiological measurements of

hippocampal LTP, showed that hippocampal NMDA receptor-dependent LTP-like synaptic plasticity is required during trial 1 for rats to learn a new goal location in the watermaze DMP task (Steele and Morris, 1999) and in a dry-land DMP task (Bast et al., 2005). LTP-like synaptic plasticity may give rise to changes in place cell representations (Dragoi et al., 2003), which could contribute to changes in hippocampal place cell networks associated with the learning of new goal locations (Dupret et al., 2010).

Map-like representations of locations, integrated within a RL architecture, may be part of neural mechanisms that enable flexibility to changing goal locations in the watermaze DMP task. Cartesian coordinates are convenient here because the task is implemented within an open-field arena, however they do not seem to provide a biologically realistic implementation of spatial navigation problems in general. For example, they do not allow navigation in an environment with walls, for which geodesic coordinates would be more appropriate (Gustafson and Daw, 2011). Moreover, the approach described here does not address how the goal representation comes about, and the model does not specify how the policy adjusts when the agent does not encounter the predicted goal. The next section describes how a hierarchical architecture can provide a solution to this problem.

Hierarchical control to flexibly adjust to task requirements

The actor–critic approach described in section ‘An actor–critic network model for incremental learning’ requires many trials to adjust to changes in goal locations partly because there is only one possible association between location and action, which depends upon a particular goal (section ‘The agent is less flexible than animals in adapting to changing goal locations’). However, brains are able to perform multiple tasks in the same environment. Those tasks often involve sequential behaviour at multiple timescales (Bouchacourt et al., 2019). Pursuing goals sometimes requires following a sequence of subroutines, with short-term/interim objectives, themselves divided into elemental skills (Botvinick et al., 2019). Hierarchically organised goal-directed behaviours allow computational RL agents to be more flexible (Botvinick et al., 2009; Dayan and Hinton, 1993).

In the watermaze DMP task, rats tend to navigate to the previous goal location on trial 1 with a new goal location (Pearce et al., 1998; Steele and Morris, 1999, and our own unpublished observations) and then find out that this remembered goal location is not the current goal. This suggests that preexisting goal networks can flexibly adjust to errors and are linked to control mechanisms over shorter timescales that allow movement realisation in order to navigate to the new goal location. The critic controls the selection of direction at every time step, finely chosen to mimic the generation of a smooth trajectory. The critic sits at an intermediary level of control. It does not perform the lower level of the control, the motor mechanisms responsible for the generation of the limb movements, but also does not control the choice or retrieval process of the goal that is being pursued. The critic allows progressive decisions in order to reach one particular goal (Van Der Meer and Redish, 2011).

We hypothesised that the computation of a goal prediction error within a hierarchical architecture could enable flexibility towards changing goal locations. We implemented a hierarchical agent, but the architecture itself does not perform hierarchical learning. In

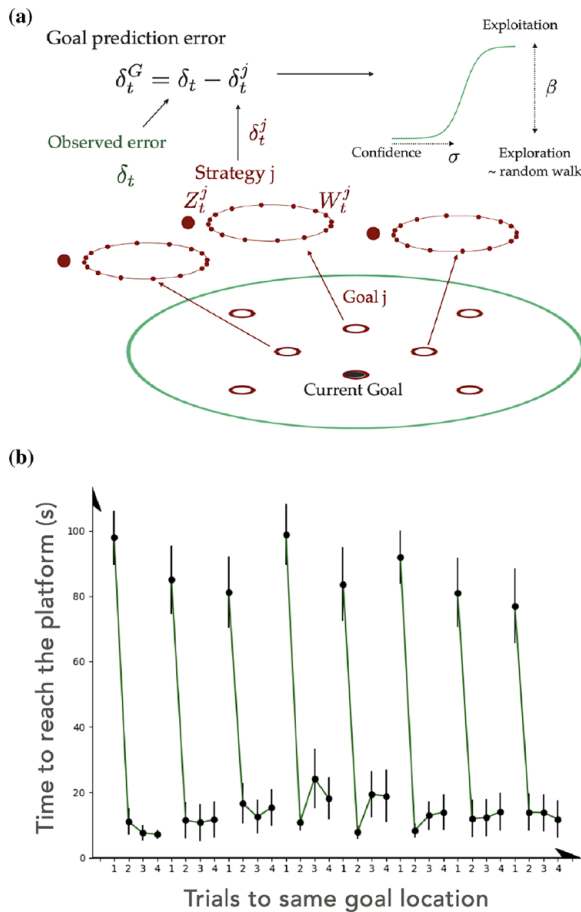


Figure 5. (a) Hierarchical RL model. The agent has learned the critic and action connection weights (Z^j and W^j , respectively) for each goal j (red circles around the maze). The actor and critic networks together, as represented in Figure 3(a), form the strategy j . A goal prediction error δ_t^G is used to compute a confidence parameter σ , which measures how good the current strategy is in reaching the current goal location. The confidence level shapes the degree of exploitation of the current strategy β through a sigmoid function of confidence. When the confidence level is very high, the strategy chosen is closely followed, as shown by a high exploitation parameter β . On the contrary, a low confidence level leads to more exploration of the environment. (b) Performance of the hierarchical agent. The model is able to adapt to changing goal locations, as seen in the reduction of latencies to reach the goal.

Botvinick (2012), agents are trained to find subroutines, for example, through looking for bottleneck states in a graph. In our implementation, we simply added a layer, which selects which one of the subroutines is most suited for the current situation.

In our implementation of a hierarchical RL architecture to allow for flexible one-trial place learning, we consider familiar goal locations, that is, although the goal location changes every four trials, the goal locations are always chosen from a set of eight locations where the agent has learned to navigate to the goal during a pre-training period. This contrasts with the most commonly used watermaze DMP procedure where the goal locations are novel (e.g.

Bast et al., 2009; Steele and Morris, 1999), although there are also DMP variations where the platform location changes daily, but is always chosen from a limited number of platform locations (Whishaw, 1985). Pretrained, familiar, goal locations are also a feature of delayed-non-matching-to-place (DNMP) tasks in the radial maze (e.g. Floresco et al., 1997; Lee and Kesner, 2002), where rats are first pretrained to learn that food can be found in any of eight arms (i.e. these are familiar goal locations); after this, the rats are required to use a ‘non-matching-to-place’ rule to choose between several open arms during daily test trials, based on whether they found food in the arms during a daily study or sample trial: arms that contained food during the sample trial will not contain food during the test trial and vice versa. Based on work by Schweighofer and Doya (2003), we propose that the agent’s behaviour could be shaped by the chosen policy depending on how confident the agent is about the policy leading to the goal.

The agent learns different policies and value functions using the model described in section ‘An actor–critic network model for incremental learning’, each of them associated with one of eight possible goal locations presented in the DMP task (which would be at the centre of the eight zones shown in Figure 1(c)). After multiple trials necessary to learn the actor and critic weights (as presented in section ‘An actor–critic network model for incremental learning’) for each goal location, the policies and value functions associated with each one of them are stored. We refer to the value and associated policy as a ‘strategy’. The choice of a strategy depends on a goal prediction error $\delta_t^G = \delta_t - \delta_t^j$ (see Figure 5(a)). The goal prediction error is used to compute a level of confidence that the agent has in the strategy it follows. When the strategy followed does not lead to the goal, the confidence level decreases, leading to more exploration of the environment until the goal is discovered. Once the goal is discovered, the strategy that minimises the prediction error is selected. Figure 5(b) represents the latencies of the agent. The agent can quickly adapt to changing goal locations, as reflected by the steep reduction in latencies between trial 1 and 2 of the new goal location.

Prefrontal areas have been proposed to carry out meta-learning computations, integrating information over multiple trials to perform computations related to a rule or a specific task (Wang et al., 2018). Neurons in prefrontal areas seem to carry goal information (Hok et al., 2005; Poucet and Hok, 2017), and their population activity dynamic correlates with the adoption of new behavioural strategies (Maggi et al., 2018). In previous work, prefrontal areas have been modelled as defining the state–action–outcomes contingencies according to the rule requirement (Daw et al., 2005; Rusu and Pennartz, 2020). Moreover, prefrontal dopaminergic activity affects flexibility towards changing rules (Ellwood et al., 2017; Goto and Grace, 2008), and frontal dopamine concentration increases during reversal learning and rule changes (Van Der Meulen et al., 2007). Therefore, the goal prediction error that shapes which goal location is pursued according to our hierarchical RL model could be computed by frontal areas from dopaminergic signals.

Limitations in accounting for open-field DMP performance. We present this approach as an illustration of how a hierarchical agent could be more flexible by separating the computation of the choice of the goal from the computation of the choice of the actions to reach it. However, the model has several

features that limit its use to provide a neuropsychologically plausible explanation of the computations underlying DMP performance in the watermaze and related open-field environments. First, the agent has to learn beforehand the connections between place cells and action and critic cells that lead to successful navigation towards every possible goal location of the maze. This would involve pretraining with the possible goal locations, whereas the agent would fail to learn a completely new goal location within one trial (i.e. return to a location that contained the goal for the very first time). Hence, the model can be considered as a model of one-shot recall, rather than one-shot learning. This cannot account for the one-trial place learning performance shown by rats and human participants on DMP tasks towards new goal locations, rather than familiar ones (Bast et al., 2009; Buckley and Bast, 2018). One discrepancy between the model's behaviour and rats can be seen during the first few trials, in which the agent automatically shows adaptation to the new goal location (as reflected by sharp latency reductions from trial 1 to 2; Figure 5(b)), whereas rats need a few trials to learn the task (Figure 1(b)). However, the hierarchical RL model may account for one-trial place learning performance on the DMP task when the changing goal locations are familiar goal locations, that is, always chosen from a limited number of locations (Whishaw, 1985, see the third point below).

Second, if the agent does not find the goal in the location to which its current strategy leads, it starts exploring the maze randomly until it finds the current goal location and selects a strategy that predicted it the best. This results in consistently high latencies during the first trial of every new goal location (Figure 5(b)). In contrast, the trial 1 latencies of rats are more variable (Figure 1(b)), for reasons considered in the section 'Map-like representation of locations for goal-directed trajectories' (last paragraph). On probe trials, removing the goal location would lead the agent to start exploring, therefore failing to reproduce the search preference as shown in open-field DMP tasks in the watermaze (Bast et al., 2009), virtual maze (Buckley and Bast, 2018) and in a dry-land arena (Bast et al., 2005). Interestingly, this suggests that rats may not show search preference during probe trials when they are tested on a DMP task variant that uses familiar goal location (Whishaw, 1985) and, therefore, may be solved by a hierarchical RL mechanism.

Third, a lesion study (Jo et al., 2007), as well as our own inactivation studies (McGarrity et al., 2015), in rats indicate that prefrontal areas are not required for successful one-shot learning of new goal locations, or the expression of such learning, in the watermaze DMP task, and frontal areas were also not among the brain areas where EEG oscillations were associated with virtual DMP performance in our recent study in human participants (Bauer et al., 2020). This contrasts with the hierarchical RL model, which implicates 'meta-control' processes that may be associated with the prefrontal cortex. However, on a DMP task variant that uses familiar goal location (Whishaw, 1985) and, therefore, may be solved by this hierarchical agent, prefrontal contributions may become more important, a hypothesis that remains to be tested. This suggests that the two DMP variants may rely on different neuro-behavioural mechanisms. Interestingly, the prefrontal cortex and hippocampo-prefrontal interactions are required for one-trial place learning performance on radial-maze (Floresco et al., 1997; Seamans et al., 1995) and T-maze (Spellman et al., 2015) DNMP tasks, which involve daily changing familiar

goal locations and, hence, may be supported by hierarchical RL mechanisms similar to our model (see section 'Hierarchical control to flexibly adjust to task requirements'). Moreover, on the T-maze DNMP task, Spellman et al. (2015) found that hippocampal projections to the medial prefrontal cortex are especially important during encoding of the reward-place association, but less so during retrieval and expression of this association. This is partly in line with the behaviour of the model, as the goal prediction error is important to select the appropriate strategy when the agent finds the correct goal location during the sample trial. However, hippocampal-prefrontal interactions are not yet considered in the model.

Fourth, hippocampal plasticity is required in open-field DMP tasks (Bast et al., 2005; Steele and Morris, 1999). The current approach suggests that the adaptation necessary during trial 1 gives rise to the selection of a set of actor and critic weights that lead to the goal through the computation of a goal prediction error. The model does not explain how the computation of the goal prediction error would be linked to hippocampal mechanisms. It is possible that a positive prediction error would make the current event (being in the right goal location) salient enough to affect its neural representation to stay in memory, for example. In a spatial navigation task in which rats had to remember reward locations chosen according to different rules, McKenzie et al. (2014) have shown that hippocampal representations are hierarchical depending on the task requirement: if the context was determining the reward location, the context would be the most discriminant factor within hippocampal representations. Recent work by Sanders et al. (2020) suggests that hierarchical inference could be used to explain remapping processes. It may be that a hierarchical representation of the task within the hippocampus could help adaptation to new goal locations through remapping processes. Hok et al. (2013) found that prefrontal lesions decreased variability of hippocampal place cell firing and hypothesised that this was linked to flexibility mechanisms and rule-based object associations (Navawongse and Eichenbaum, 2013) within hippocampal firing patterns. This finding shows that the prefrontal cortex can modulate hippocampal place cell activity. If the goal prediction error is coded by the prefrontal cortex, these findings imply that the goal prediction error could act on hippocampal representations in order to incorporate new task requirements (e.g. information about the new goal location) and modify expectations.

A potential account of arm-maze DNMP performance?

Although the hierarchical RL approach may be limited in accounting for key features of performance on DMP tasks using novel goal locations, it may have more potential in accounting for flexible trial-dependent behaviour displayed by rats on DNMP tasks in the radial arm maze, which involve trial-dependent choices between familiar goal locations. DNMP performance in radial-maze tasks requires NMDA receptors, including in the hippocampus, during pretraining, although after pretraining, and contrary to the watermaze (Steele and Morris, 1999) and event arena DMP tasks (Bast et al., 2005), rats can acquire and maintain trial-specific place information independent of hippocampal NMDA receptor-mediated plasticity, even though the hippocampus is still required (Caramanos and Shapiro, 1994; Lee and Kesner, 2002; Shapiro and O'Connor, 1992). The hierarchical RL architecture may account for this phase of acquisition of

arm-reward association via pretraining to the eight possible goal locations, via the formation of actor and critic weights of every strategy. However, the plasticity considered in the hierarchical model is more consistent with changes in hippocampal-striatal connections, whereas the model does not address the role of plasticity within the hippocampus during this phase. Moreover, the hierarchical RL model also fits with the requirement of the prefrontal cortex for flexible spatial behaviour on arm-maze tasks, as described in the previous section.

However, to test if a hierarchical RL architecture can reproduce behaviour on DNMP arm-maze task, our implementation of a hierarchical RL model outlined above (see Figure 5(a)) would need to be adapted to the arm-maze environments, to the DNMP rule and to an error measure of performance that is typically used in arm-maze tasks (Floresco et al., 1997; Seamans et al., 1995). The goal prediction error would drive exploration to other arms and provide a long-term control allowing to carry memories of previously visited goals.

Conclusion

We presented an actor–critic architecture, which leads to action selection based on the difference of estimated rewards to be received. The model (Foster et al., 2000) uses an estimate of the value function over the maze to drive behaviour, through a critic network that receives place cell activities as input. It can successfully learn to select which action is best through an actor network, which also receives place cell input and is trained to follow the gradient of the value function from the critic difference in successive activities. This agent can follow trajectories towards a particular, fixed, goal location, that corresponds to the maximum of the value function. However, when the goal location changes, the model needs many trials to adjust in order to accurately navigate to the new goal, which is in marked contrast with real DMP performance of rats (Figure 1(b)) and humans (Buckley and Bast, 2018).

To account for one-trial place learning performance on the DMP task, a possible extension to the actor–critic approach is to learn map-like representation of locations throughout the maze that facilitate the direct comparison between the goal location and the agent's location. This enables the computation of a goal-directed displacement towards any new goal location throughout the maze and reproduces flexibility shown by humans and animals towards new goal location, as reflected by sharp latency reductions from trial 1 to 2 to a new goal location and marked search preference for the new goal location when trial 2 is run as probe (Figure 4(b) and (c)).

Given that the striatum has been associated with actor–critic mechanisms (Joel et al., 2002; Khamassi and Humphries, 2012; O'Doherty et al., 2004; Van Der Meer and Redish, 2011), using an actor–critic agent for flexible spatial navigation is consistent with empirical evidence associating striatal regions with place learning performance on both incremental (Annett et al., 1989; Braun et al., 2010; Devan and White, 1999) and DMP (Bauer et al., 2020; Seaton, 2019) tasks. However, contrasting with the coordinate extension to the actor–critic architecture, experimental evidence suggests that goal location memory may lie within hippocampal place cell representations (Dupret et al., 2010; McKenzie et al.,

2013) and that one-trial place learning performance on DMP tasks in rats requires NMDA receptor-dependent LTP-like hippocampal synaptic plasticity (Bast et al., 2005; Steele and Morris, 1999).

Finally, we illustrated how flexibility may be generated through hierarchical organisation of task control (Balleine et al., 2015; Botvinick et al., 2009; Dayan and Hinton, 1993). In an extension of Foster et al. (2000), we separated the selection of the goal from the control of the displacement towards it by means of a different read-out of the TD error by the different control systems. The agent first learns different 'strategies', each of which correspond to a different critic-actor component that leads to one of the possible goal locations. The critic and actor are used to perform the displacement to the goal location. An additional hierarchical layer is added to compute a goal prediction error which compares location of the goal predicted from the strategies to the real goal location and selects a strategy accordingly. The agent follows its strategy depending on a confidence parameter that integrates goal prediction errors information over multiple trials. This hierarchical RL agent can adapt to changing goal locations, although these goal locations need to be familiar, whereas in open-field DMP tasks, the changing goal locations are new (Bast et al., 2005, 2009). However, the hierarchical RL approach may be more suitable to account for situation when trial-specific memories of familiar goal location need to be formed, as on arm-maze DNMP tasks (section 'A potential account of arm-maze DNMP performance?')

To conclude, elements of an actor–critic architecture may account for some important aspects of rapid place learning performance in the DMP watermaze task. Together with a map-like representation of location, an actor–critic architecture can support the efficient, fast, goal-directed computations required for such performance, and a hierarchical structure is useful for efficient, distributed, control. Future models of hippocampus-dependent flexible spatial navigation should involve LTP-like plasticity mechanisms and goal location representation within the hippocampus, which have been implicated in trial-specific place memory by substantial empirical evidence.

Acknowledgements

We thank Georg Raiser for fruitful discussions and comments on the manuscript.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship and/or publication of this article: C.T. was funded by a PhD studentship supported by the School of Mathematical Sciences and the School of Psychology and the University of Nottingham.

ORCID iD

Charline Tessereau  <https://orcid.org/0000-0002-0385-2802>

Supplemental material

Supplemental material for this article is available at: <https://journals.sagepub.com/doi/suppl/10.1177/2398212820975634>.

References

- Anggraini D, Glasauer S and Wunderlich K (2018) Neural signatures of reinforcement learning correlate with strategy adoption during spatial navigation. *Scientific Reports* 8(1): 10110.
- Annett L, McGregor A and Robbins T (1989) The effects of ibotenic acid lesions of the nucleus accumbens on spatial learning and extinction in the rat. *Behavioural Brain Research* 31(3): 231–242.
- Balleine BW (2019) The meaning of behavior: Discriminating reflex and volition in the brain. *Neuron* 104(1): 47–62.
- Balleine BW, Dezfouli A, Ito M, et al. (2015) Hierarchical control of goal-directed action in the cortical–basal ganglia network. *Current Opinion in Behavioral Sciences* 5: 1–7.
- Banino A, Barry C, Uria B, et al. (2018) Vector-based navigation using grid-like representations in artificial agents. *Nature* 557(7705): 429–433.
- Bannerman D, Good M, Butcher S, et al. (1995) Distinct components of spatial learning revealed by prior training and NMDA receptor blockade. *Nature* 378(6553): 182–186.
- Barreto A, Dabney W, Munos R, et al. (2017) Successor features for transfer in reinforcement learning. In: MC Mozer, MI Jordan and T Petsche (eds) *Advances in Neural Information Processing Systems*. Cambridge, MA: The MIT Press, pp. 4055–4065.
- Bast T (2007) Toward an integrative perspective on hippocampal function: From the rapid encoding of experience to adaptive behavior. *Reviews in the Neurosciences* 18(3–4): 253–282.
- Bast T (2011) The hippocampal learning-behavior translation and the functional significance of hippocampal dysfunction in schizophrenia. *Current Opinion in Neurobiology* 21(3): 492–501.
- Bast T, da Silva BM and Morris RG (2005) Distinct contributions of hippocampal NMDA and AMPA receptors to encoding and retrieval of one-trial place memory. *Journal of Neuroscience* 25(25): 5845–5856.
- Bast T, Wilson IA, Witter MP, et al. (2009) From rapid place learning to behavioral performance: A key role for the intermediate hippocampus. *PLOS Biology* 7(4): e1000089.
- Battaglia FP and Treves A (1998) Attractor neural networks storing multiple space representations: A model for hippocampal place fields. *Physical Review E* 58(6): 7738–7753.
- Bauer M, Buckley MG and Bast T (2020) Individual differences in theta-band oscillations in a spatial memory network revealed by EEG predict rapid place learning. *bioRxiv*. Epub ahead of print 5 June. DOI: 10.1101/2020.06.05.134346.
- Bicanski A and Burgess N (2020) Neuronal vector coding in spatial cognition. *Nature Reviews Neuroscience* 21(9): 453–447.
- Bingman VP and Sharp PE (2006) Neuronal implementation of hippocampal-mediated spatial behavior: A comparative evolutionary perspective. *Behavioral and Cognitive Neuroscience Reviews* 5(2): 80–91.
- Botvinick M, Ritter S, Wang JX, et al. (2019) Reinforcement learning, fast and slow. *Trends in Cognitive Sciences* 23(5): 408–422.
- Botvinick MM (2012) Hierarchical reinforcement learning and decision making. *Current Opinion in Neurobiology* 22(6): 956–962.
- Botvinick MM, Niv Y and Barto AC (2009) Hierarchically organized behavior and its neural foundations: A reinforcement learning perspective. *Cognition* 113(3): 262–280.
- Bouchacourt F, Palminteri S, Koechlin E, et al. (2019) Temporal chunking as a mechanism for unsupervised learning of task-sets. *bioRxiv*. Epub ahead of print 24 July. DOI: 10.1101/713156.
- Braun AA, Graham DL, Schaefer TL, et al. (2012) Dorsal striatal dopamine depletion impairs both allocentric and egocentric navigation in rats. *Neurobiology of Learning and Memory* 97(4): 402–408.
- Braun DA, Mehring C and Wolpert DM (2010) Structure learning in action. *Behavioural Brain Research* 206(2): 157–165.
- Buckley MG and Bast T (2018) A new human delayed-matching-to-place test in a virtual environment reverse-translated from the rodent watermaze paradigm: Characterization of performance measures and sex differences. *Hippocampus* 28(11): 796–812.
- Caramanos Z and Shapiro ML (1994) Spatial memory and N-methyl-D-aspartate receptor antagonists APV and MK-801: Memory impairments depend on familiarity with the environment, drug dose, and training duration. *Behavioral Neuroscience* 108(1): 30–43.
- Colombo M and Broadbent N (2000) Is the avian hippocampus a functional homologue of the mammalian hippocampus? *Neuroscience & Biobehavioral Reviews* 24(4): 465–484.
- Corneil DS and Gerstner W (2015) Attractor network dynamics enable preplay and rapid path planning in maze-like environments. In: C Cortes, N Lawrence, D Lee, et al. (eds) *Advances in Neural Information Processing Systems*. New York: Curran Associates, Inc., pp. 1684–1692.
- Da Silva CF and Hare TA (2019) Humans are primarily model-based and not model-free learners in the two-stage task. *bioRxiv*. Epub ahead of print 25 July. DOI: 10.1101/682922.
- Daw ND, Gershman SJ, Seymour B, et al. (2011) Model-based influences on humans' choices and striatal prediction errors. *Neuron* 69(6): 1204–1215.
- Daw ND, Niv Y and Dayan P (2005) Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience* 8(12): 1704–1711.
- Dayan P (1991) Navigating through temporal difference. In: RP Lippmann, JE Moody and DS Touretzky (eds) *Advances in Neural Information Processing Systems*. San Mateo, CA: Morgan Kaufmann, pp. 464–470.
- Dayan P (1993) Improving generalization for temporal difference learning: The successor representation. *Neural Computation* 5(4): 613–624.
- Dayan P and Hinton GE (1993) Feudal reinforcement learning. In: SJ Hanson, JD Cowan and CL Giles (eds) *Advances in Neural Information Processing Systems*. San Mateo, CA: Morgan Kaufmann, pp. 271–278.
- De Hoz L, Knox J and Morris RG (2003) Longitudinal axis of the hippocampus: Both septal and temporal poles of the hippocampus support water maze spatial learning depending on the training protocol. *Hippocampus* 13(5): 587–603.
- De Hoz L, Moser EI and Morris RG (2005) Spatial learning with unilateral and bilateral hippocampal networks. *European Journal of Neuroscience* 22(3): 745–754.
- Devan BD and White NM (1999) Parallel information processing in the dorsal striatum: Relation to hippocampal function. *Journal of Neuroscience* 19(7): 2789–2798.
- Dezfouli A and Balleine BW (2019) Learning the structure of the world: The adaptive nature of state-space and action representations in multi-stage decision-making. *PLOS Computational Biology* 15(9): e1007334.
- Dollé L, Chavarriaga R, Guillot A, et al. (2018) Interactions of spatial strategies producing generalization gradient and blocking: A computational approach. *PLOS Computational Biology* 14(4): e1006092.
- Doya K (2000) Reinforcement learning in continuous time and space. *Neural Computation* 12(1): 219–245.
- Dragoi G, Harris KD and Buzsáki G (2003) Place representation within hippocampal networks is modified by long-term potentiation. *Neuron* 39(5): 843–853.
- Ducarouge A and Sigaud O (2017) The successor representation as a model of behavioural flexibility. Available at: <https://hal.archives-ouvertes.fr/hal-01576352>
- Dupret D, O'Neill J, Pleydell-Bouverie B, et al. (2010) The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nature Neuroscience* 13(8): 995–1002.
- Eichenbaum H (1990) Hippocampal representation in spatial learning. *Journal of Neuroscience* 10(11): 331–339.
- Ellwood IT, Patel T, Wadia V, et al. (2017) Tonic or phasic stimulation of dopaminergic projections to prefrontal cortex causes mice to maintain

- or deviate from previously learned behavioral strategies. *Journal of Neuroscience* 37(35): 8315–8329.
- Floresco SB, Seamans JK and Phillips AG (1997) Selective roles for hippocampal, prefrontal cortical, and ventral striatal circuits in radial-arm maze tasks with or without a delay. *Journal of Neuroscience* 17(5): 1880–1890.
- Foster D, Morris R and Dayan P (2000) A model of hippocampally dependent navigation, using the temporal difference learning rule. *Hippocampus* 10(1): 1–16.
- Frankenhuis WE, Panchanathan K and Barto AG (2019) Enriching behavioral ecology with reinforcement learning methods. *Behavioural Processes* 161(2019): 94–100.
- Frémaux N, Sprekeler H and Gerstner W (2013) Reinforcement learning using a continuous time actor-critic framework with spiking neurons. *PLoS Computational Biology* 9(4): e1003024.
- Fuhs MC and Touretzky DS (2006) A spin glass model of path integration in rat medial entorhinal cortex. *Journal of Neuroscience* 26(16): 4266–4276.
- Gauthier JL and Tank DW (2018) A dedicated population for reward coding in the hippocampus. *Neuron* 99(1): 179–193.
- Gehring TV, Luksys G, Sandi C, et al. (2015) Detailed classification of swimming paths in the Morris Water Maze: Multiple strategies within one trial. *Scientific Reports* 5(1): 14562.
- Gerfen CR and Surmeier DJ (2011) Modulation of striatal projection systems by dopamine. *Annual Review of Neuroscience* 34: 441–466.
- Gershman SJ (2017) Reinforcement learning and causal models. In: M Waldmann (ed.) *The Oxford Handbook of Causal Reasoning*. New York: Oxford University Press, pp. 295–306.
- Gershman SJ (2018) The successor representation: Its computational logic and neural substrates. *Journal of Neuroscience* 38(33): 7193–7200.
- Gershman SJ, Markman AB and Otto AR (2014) Retrospective reevaluation in sequential decision making: A tale of two systems. *Journal of Experimental Psychology: General* 143(1): 182–194.
- Gerstner W and Abbott L (1997) Learning navigational maps through potentiation and modulation of hippocampal place cells. *Journal of Computational Neuroscience* 4(1): 79–94.
- Glimcher PW (2011) Understanding dopamine and reinforcement learning: The dopamine reward prediction error hypothesis. *Proceedings of the National Academy of Sciences of the United States of America* 108(Supplement 3): 15647–15654.
- Goto Y and Grace AA (2008) Dopamine modulation of hippocampal–prefrontal cortical interaction drives memory-guided behavior. *Cerebral Cortex* 18(6): 1407–1414.
- Groenewegen H, Vermeulen-Van der Zee E, Te Kortschot A, et al. (1987) Organization of the projections from the subiculum to the ventral striatum in the rat: A study using anterograde transport of phaseolus vulgaris leucoagglutinin. *Neuroscience* 23(1): 103–120.
- Gustafson NJ and Daw ND (2011) Grid cells, place cells, and geodesic generalization for spatial reinforcement learning. *PLoS Computational Biology* 7(10): e1002235.
- Haferlach T, Wessnitzer J, Mangan M, et al. (2007) Evolving a neural model of insect path integration. *Adaptive Behavior* 15(3): 273–287.
- Hinman JR, Chapman GW and Hasselmo ME (2019) Neuronal representation of environmental boundaries in egocentric coordinates. *Nature Communications* 10(1): 1–8.
- Hok V, Chah E, Save E, et al. (2013) Prefrontal cortex focally modulates hippocampal place cell firing patterns. *Journal of Neuroscience* 33(8): 3443–3451.
- Hok V, Lenck-Santini PP, Roux S, et al. (2007) Goal-related activity in hippocampal place cells. *Journal of Neuroscience* 27(3): 472–482.
- Hok V, Save E, Lenck-Santini P, et al. (2005) Coding for spatial goals in the prelimbic/infralimbic area of the rat frontal cortex. *Proceedings of the National Academy of Sciences of the United States of America* 102(12): 4602–4607.
- Howard RA (1960) *Dynamic Programming and Markov Processes*. Cambridge, MA: MIT Press.
- Humphries MD and Prescott TJ (2010) The ventral basal ganglia, a selection mechanism at the crossroads of space, strategy, and reward. *Progress in Neurobiology* 90(4): 385–417.
- Humphries MD, Khamassi M and Gurney K (2012) Dopaminergic control of the exploration-exploitation trade-off via the basal ganglia. *Frontiers in Neuroscience* 6(9): 9.
- Huys QJM, Cruickshank A and Series P (2013) *Reward-Based Learning, Model-Based and Model-Free*. New York: Springer.
- Inglis J, Martin SJ and Morris RG (2013) Upstairs/downstairs revisited: Spatial pretraining-induced rescue of normal spatial learning during selective blockade of hippocampal N-methyl-D-aspartate receptors. *European Journal of Neuroscience* 37(5): 718–727.
- Jeffery KJ (2018) Cognitive representations of spatial location. *Brain and Neuroscience Advances* 2: 2398212818810686.
- Jo YS, Park EH, Kim IH, et al. (2007) The medial prefrontal cortex is involved in spatial memory retrieval under partial-cue conditions. *Journal of Neuroscience* 27(49): 13567–13578.
- Joel D, Niv Y and Ruppin E (2002) Actor–critic models of the basal ganglia: New anatomical and computational perspectives. *Neural Networks* 15(4–6): 535–547.
- Jong NK and Stone P (2007) Model-based exploration in continuous state spaces. In: I Miguel and W Ruml (eds) *7th International Symposium on Abstraction, Reformulation, and Approximation (SARA 2007), Volume 4612 of Lecture Notes in Computer Science*. Whistler, BC, Canada: Springer, pp. 258–272.
- Kanitscheider I and Fiete I (2017) Making our way through the world: Towards a functional understanding of the brain’s spatial circuits. *Current Opinion in Systems Biology* 3: 186–194.
- Keramati M, Dezfouli A and Piray P (2011) Speed/accuracy trade-off between the habitual and the goal-directed processes. *PLoS Computational Biology* 7(5): e1002055.
- Khamassi M and Humphries MD (2012) Integrating corticolimbic-basal ganglia architectures for learning model-based and model-free navigation strategies. *Frontiers in Behavioral Neuroscience* 6: 79.
- Kimchi EY and Laubach M (2009) Dynamic encoding of action selection by the medial striatum. *Journal of Neuroscience* 29(10): 3148–3159.
- Kjelstrup KB, Solstad T, Brun VH, et al. (2008) Finite scale of spatial representation in the hippocampus. *Science* 321(5885): 140–143.
- Lee I and Kesner RP (2002) Differential contribution of NMDA receptors in hippocampal subregions to spatial working memory. *Nature Neuroscience* 5(2): 162–168.
- LeGates TA, Kvarta MD, Tooley JR, et al. (2018) Reward behaviour is regulated by the strength of hippocampus–nucleus accumbens synapses. *Nature* 564(7735): 258–262.
- Li J and Daw ND (2011) Signals in human striatum are appropriate for policy update rather than value prediction. *Journal of Neuroscience* 31(14): 5504–5511.
- Maggi S, Peyrache A and Humphries MD (2018) An ensemble code in medial prefrontal cortex links prior events to outcomes during learning. *Nature Communications* 9(1): 1–12.
- McGarrity S, Mason R, Fone KC, et al. (2017) Hippocampal neural disinhibition causes attentional and memory deficits. *Cerebral Cortex* 27(9): 4447–4462.
- McGarrity S, Somerled S, Eaton C, et al. (2015) Medial prefrontal cortex is not required for, but can modulate, hippocampus-dependent behaviour based on rapid learning of changing goal locations on the Water Maze delayed-matching-to-place test. *British Neuroscience Association Abstracts* 23: P118.
- McKenzie S, Frank AJ, Kinsky NR, et al. (2014) Hippocampal representation of related and opposing memories develop within distinct, hierarchically organized neural schemas. *Neuron* 83(1): 202–215.
- McKenzie S, Robinson NT, Herrera L, et al. (2013) Learning causes reorganization of neuronal firing patterns to represent related experiences

- within a hippocampal schema. *Journal of Neuroscience* 33(25): 10243–10256.
- Miller KJ, Botvinick MM and Brody CD (2017) Dorsal hippocampus contributes to model-based planning. *Nature Neuroscience* 20(9): 1269–1276.
- Morris G, Schmidt R and Bergman H (2010) Striatal action-learning based on dopamine concentration. *Experimental Brain Research* 200(3–4): 307–317.
- Morris R, Anderson E, Lynch Ga and Baudry M (1986) Selective impairment of learning and blockade of long-term potentiation by an N-methyl-D-aspartate receptor antagonist, AP5. *Nature* 319(6056): 774–776.
- Morris R, Halliwell R and Bowery N (1989) Synaptic plasticity and learning II: Do different kinds of plasticity underlie different kinds of learning? *Neuropsychologia* 27(1): 41–59.
- Morris R, Schenk F, Tweedie F, et al. (1990) Ibotenate lesions of hippocampus and/or subiculum: Dissociating components of allocentric spatial learning. *European Journal of Neuroscience* 2(12): 1016–1028.
- Morris RG (1981) Spatial localization does not require the presence of local cues. *Learning and Motivation* 12(2): 239–260.
- Morris RG (2008) Morris water maze. *Scholarpedia* 3(8): 6315–6317.
- Morris RG, Garrud P, Rawlins Ja and O’Keefe J (1982) Place navigation impaired in rats with hippocampal lesions. *Nature* 297(5868): 681–683.
- Moser EI, Moser MB and McNaughton BL (2017) Spatial representation in the hippocampal formation: A history. *Nature Neuroscience* 20(11): 1448–1464.
- Moser MB, Moser EI, Forrest E, et al. (1995) Spatial learning with a minislab in the dorsal hippocampus. *Proceedings of the National Academy of Sciences of the United States of America* 92(21): 9697–9701.
- Nakazawa K, McHugh TJ, Wilson MA, et al. (2004) NMDA receptors, place cells and hippocampal spatial memory. *Nature Reviews Neuroscience* 5(5): 361–372.
- Nakazawa K, Sun LD, Quirk MC, et al. (2003) Hippocampal CA3 NMDA receptors are crucial for memory acquisition of one-time experience. *Neuron* 38(2): 305–315.
- Navawongse R and Eichenbaum H (2013) Distinct pathways for rule-based retrieval and spatial mapping of memory representations in hippocampal neurons. *Journal of Neuroscience* 33(3): 1002–1013.
- O’Carroll CM, Martin SJ, Sandin J, et al. (2006) Dopaminergic modulation of the persistence of one-trial hippocampus-dependent memory. *Learning & Memory* 13(6): 760–769.
- O’Doherty J, Dayan P, Schultz J, et al. (2004) Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science* 304(5669): 452–454.
- O’Keefe J (2014) *Spatial cells in the hippocampal formation*. Nobel Lecture on 7 December 2014 at Aula Medica, Karolinska Institutet, Stockholm.
- O’Keefe J and Burgess N (1996) Geometric determinants of the place fields of hippocampal neurons. *Nature* 381(6581): 425–428.
- O’Keefe J and Dostrovsky J (1971) The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely-moving rat. *Brain Research* 34(1): 171–175.
- Pearce JM, Roberts AD and Good M (1998) Hippocampal lesions disrupt navigation based on cognitive maps but not heading vectors. *Nature* 396(6706): 75–77.
- Penny WD, Zeidman P and Burgess N (2013) Forward and backward inference in spatial cognition. *PLOS Computational Biology* 9(12): e1003383.
- Pezze M and Bast T (2012) Dopaminergic modulation of hippocampus-dependent learning: Blockade of hippocampal D1-class receptors during learning impairs 1-trial place memory at a 30-min retention delay. *Neuropharmacology* 63(4): 710–718.
- Pezzulo G, Kemere C and Van Der Meer MA (2017) Internally generated hippocampal sequences as a vantage point to probe future-oriented cognition. *Annals of the New York Academy of Sciences* 1396(1): 144–165.
- Pezzulo G, Rigoli F and Chersi F (2013) The mixed instrumental controller: Using value of information to combine habitual choice and mental simulation. *Frontiers in Psychology* 4(92): 92.
- Poucet B and Hok V (2017) Remembering goal locations. *Current Opinion in Behavioral Sciences* 17: 51–56.
- Redish AD (2016) Vicarious trial and error. *Nature Reviews Neuroscience* 17(3): 147–159.
- Redish AD and Touretzky DS (1998) The role of the hippocampus in solving the Morris water maze. *Neural Computation* 10(1): 73–111.
- Roesch MR, Singh T, Brown PL, et al. (2009) Ventral striatal neurons encode the value of the chosen action in rats deciding between differently delayed or sized rewards. *Journal of Neuroscience* 29(42): 13365–13376.
- Russek EM, Momennejad I, Botvinick MM, et al. (2017) Predictive representations can link model-based reinforcement learning to model-free mechanisms. *PLOS Computational Biology* 13(9): e1005768.
- Rusu SI and Pennartz CM (2020) Learning, memory and consolidation mechanisms for behavioral control in hierarchically organized cortico-basal ganglia systems. *Hippocampus* 30(1): 73–98.
- Samejima K, Ueda Y, Doya K, et al. (2005) Representation of action-specific reward values in the striatum. *Science* 310(5752): 1337–1340.
- Samsonovich A and McNaughton BL (1997) Path integration and cognitive mapping in a continuous attractor neural network model. *Journal of Neuroscience* 17(15): 5900–5920.
- Sanders H, Wilson MA and Gershman SJ (2020) Hippocampal remapping as hidden state inference. *Elife* 9: e51140.
- Sarel A, Finkelstein A, Las L, et al. (2017) Vectorial representation of spatial goals in the hippocampus of bats. *Science* 355(6321): 176–180.
- Schultz W, Dayan P and Montague PR (1997) A neural substrate of prediction and reward. *Science* 275(5306): 1593–1599.
- Schweighofer N and Doya K (2003) Meta-learning in reinforcement learning. *Neural Networks* 16(1): 5–9.
- Scleidorovich P, Llofriu M, Fellous JM, et al. (2020) A computational model for spatial cognition combining dorsal and ventral hippocampal place field maps: Multiscale navigation. *Biological Cybernetics* 114: 187–207.
- Seamans JK, Floresco SB and Phillips AG (1995) Functional differences between the prelimbic and anterior cingulate regions of the rat prefrontal cortex. *Behavioral Neuroscience* 109(6): 1063–1073.
- Seaton A (2019) *An investigation of the role of the nucleus accumbens in the hippocampal learning-behaviour translation*. PhD Thesis, University of Nottingham, Nottingham.
- Shapiro ML and O’Connor C (1992) N-methyl-D-aspartate receptor antagonist MK-801 and spatial memory representation: Working memory is impaired in an unfamiliar but not in a familiar environment. *Behavioral Neuroscience* 106(4): 604–612.
- Smith KS and Graybiel AM (2013) A dual operator view of habitual behavior reflecting cortical and striatal dynamics. *Neuron* 79(2): 361–374.
- Spellman T, Rigotti M, Ahmari SE, et al. (2015) Hippocampal-prefrontal input supports spatial encoding in working memory. *Nature* 522(7556): 309–314.
- Stachenfeld KL, Botvinick MM and Gershman SJ (2017) The hippocampus as a predictive map. *Nature Neuroscience* 20(11): 1643–1653.
- Steele R and Morris R (1999) Delay-dependent impairment of a matching-to-place task with chronic and intrahippocampal infusion of the NMDA-antagonist D-AP5. *Hippocampus* 9(2): 118–136.
- Sutton RS and Barto AG (2018) *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.
- Tai LH, Lee AM, Benavidez N, et al. (2012) Transient stimulation of distinct subpopulations of striatal neurons mimics changes in action value. *Nature Neuroscience* 15(9): 1281–1289.
- Thorndike EL (1927) The law of effect. *The American Journal of Psychology* 39(1–4): 212–222.

- Tolman EC (1948) Cognitive maps in rats and men. *Psychological Review* 55(4): 189–208.
- Tomov MS, Truong VQ, Hundia RA, et al. (2020) Dissociable neural correlates of uncertainty underlie different exploration strategies. *Nature Communications* 11(1): 1–12.
- Van Der Meer MA and Redish AD (2011) Ventral striatum: A critical look at models of learning and evaluation. *Current Opinion in Neurobiology* 21(3): 387–392.
- Van Der Meulen JA, Joosten RN, de Bruin JP, et al. (2007) Dopamine and noradrenaline efflux in the medial prefrontal cortex during serial reversals and extinction of instrumental goal-directed behavior. *Cerebral Cortex* 17(6): 1444–1453.
- Wang JX, Kurth-Nelson Z, Kumaran D, et al. (2018) Prefrontal cortex as a meta-reinforcement learning system. *Nature Neuroscience* 21(6): 860–868.
- Whishaw IQ (1985) Formation of a place learning-set by the rat: A new paradigm for neurobehavioral studies. *Physiology & Behavior* 35(1): 139–143.
- Widloski J and Fiete IR (2014) A model of grid cell development through spatial exploration and spike time-dependent plasticity. *Neuron* 83(2): 481–495.
- Yin HH and Knowlton BJ (2006) The role of the basal ganglia in habit formation. *Nature Reviews Neuroscience* 7(6): 464–476.