Neuron Special Issue on *Cognitive Architectures*

The cognitive architecture of spatial navigation: hippocampal and striatal contributions

Fabian Chersi, Neil Burgess

Institute of Cognitive Neuroscience & Institute of Neurology. University College London, 17 Queen Square, London WC1N 3AR, UK
Contact: f.chersi@ucl.ac.uk , n.burgess@ucl.ac.uk

Abstract

Spatial navigation can serve as a model system in cognitive neuroscience, in which specific neural representations, learning rules and control strategies can be inferred from the vast experimental literature that exists across many species, including humans.

Here, we review this literature, focusing on the contributions of hippocampal and striatal systems, and attempt to outline a minimal cognitive architecture that is consistent with the experimental literature and which synthesizes previous related computational modeling. The resulting architecture includes striatal reinforcement learning based on egocentric representations of sensory states and actions, incidental Hebbian association of sensory information with allocentric state representations in the hippocampus, and arbitration of the outputs of both systems based on confidence/uncertainty in medial prefrontal cortex. We discuss the relationship between this architecture and learning in model-free and model-based systems, episodic memory, imagery and planning, including some open questions and directions for further experiments.

1 Introduction

Goal-directed spatial navigation is a good model for general issues in cognitive neuroscience. Like many daily tasks, goal-directed navigation is a complex task that involves a variety of sensory and proprioceptive stimuli, storage and recall of information, and the elaboration of plans. Moreover, there is now an unparalleled literature concerning the neural representations involved (as reflected in the 2014 Nobel Prize, see e.g., Burgess, 2014) and a vast array of experimental data relating behavior to environmental or neurophysiological manipulations.

There are many ways to find a goal location, and the relevant cognitive functions have been categorized in various ways. Here, we follow the nomenclature coming from behavioral and

lesion experiments in animals, following from the suggestion that the hippocampus provides a 'cognitive map' (O'Keefe and Nadel, 1978), see below. Thus, one might navigate by following a sensory cue that directly indicates the goal location ('piloting'), by following a well-learned sequence of actions, each depending on the previous action or a sensory cue ('response learning' using 'route' or 'taxon' strategies), or by following a flexible internal representation of spatial layout ('place learning', using 'cognitive map' or 'locale' strategies). The emphasis on the 'flexibility' of a cognitive map refers to the ability to use it from a new starting location (which would undermine the use of a route) or in the absence of subsets of specific sensory cues (which could undermine the use of piloting or route strategies). This hypothesis that the hippocampus supports a flexible representation of the spatial relationships present in the environment has been extended to include non-spatial information in the relational theory of memory function (Cohen and Eichenbaum, 1993).

Here we examine the cognitive architecture of spatial navigation, with a focus on hippocampal and striatal systems and their interaction. We aim to outline the general principles that can be derived from experimental data and how they constrain the development of formal models of spatial cognition.

2. Brain regions associated with spatial navigation

Like any complex task, spatial navigation involves much of the brain, not least sensory and motor areas. Here we review some of the systems associated with specific roles in navigation. The hippocampus has long been known to be important for episodic memory (Scoville and Milner, 1957), and the discovery of place cells drew attention to its role in spatial memory (O'Keefe and Nadel, 1978), see below. Subsequent experiments using the Morris water maze (Morris et al., 1982), T maze alternation (e.g. Cohen et al., 1971) and the 8 arm maze (e.g. Olton et al., 1977) demonstrated certain aspects of navigation to be particularly sensitive to hippocampal damage. These include navigation to an unmarked location from variable start locations and navigation that requires memory (and potentially avoidance) of previously visited locations. Equally importantly, control conditions in these tasks showed insensitivity to hippocampal damage when piloting or response learning was possible, and emphasized the importance of distal cues in orienting the animal within its environment.

By contrast with the hippocampal formation, cortico-basal ganglia circuits (loops) are thought to store stimulus-response associations and procedural memories (which may underlie route learning or piloting). These loops connect specific neocortical areas unidirectionally to striatal subregions, which project to downstream structures such as the pallidum and the substantia nigra

(SN). These areas connect to the thalamic nuclei that in turn project back to the same neocortical sites of origin. Within this picture, the striatum has been subdivided into several functional regions: the dorsolateral striatum (DLS) associated with stimulus–response learning and habit formation, the dorsomedial striatum (DMS) associated with action–outcome learning, and ventral striatum (VS) associated with motivational and affective processing (Corbit and Balleine, 2003; Packard and McGaugh, 1992; Voorn et al., 2004; Yin and Knowlton, 2006).

An important aspect of theorizing about the function of the striatum concerns its dopaminergic input. In a series of influential experiments Schultz and colleagues recorded the firing of dopamine neurons in the SN/VTA of monkeys performing conditioning experiments. The firing of these neurons was found to be consistent with dopaminergic signaling of the reward prediction error (Schultz et al., 1997) used in theoretical models of reinforcement learning (RL). The strong dopaminergic projections to the striatum strengthen its association with RL, with suggestions that an actor/critic- type functional architecture may exist in dorsal and ventral striatum respectively (e.g. O'Doherty et al., 2004; Bornstein and Daw, 2011; Pennartz et al., 2011), see Box 1 and Figure 1.

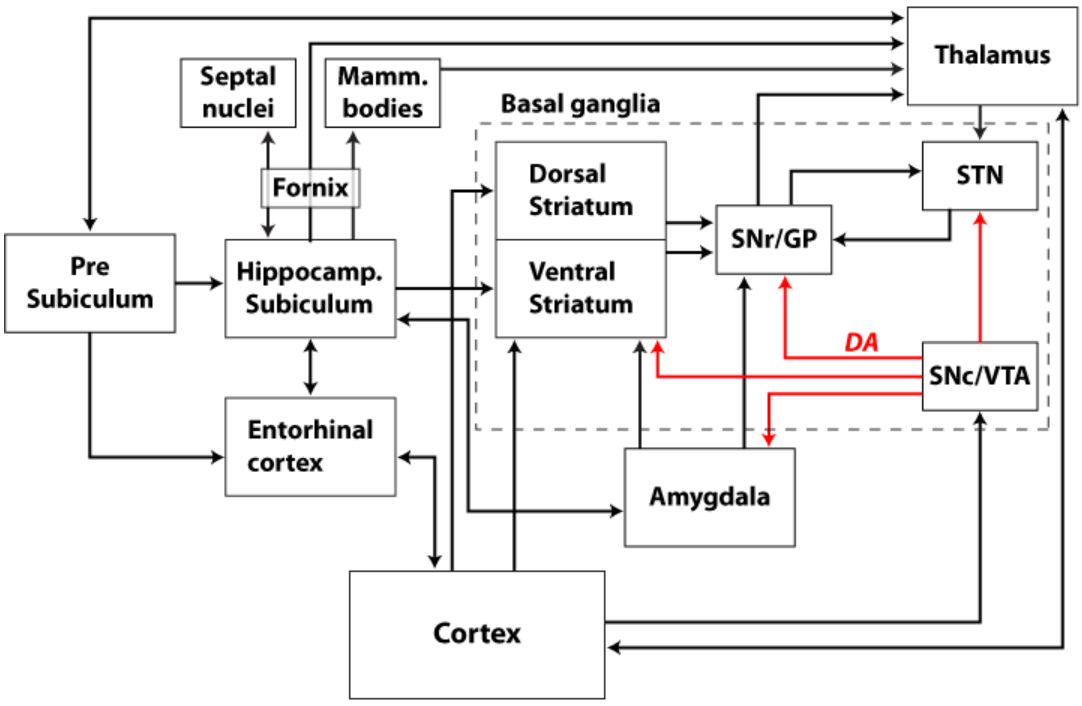

Figure 1. Simplified schematic of the cortical and subcortical connections of the hippocampus and the striatum. Most of the hippocampus's neocortical inputs come from the perirhinal and parahippocampal cortices (not shown), via the Entorhinal cortex. The striatum is part of the basal ganglia (including SNr, GP, SNc, VTA), which we have here represented in a compact form for the sake of clarity. The 'dorsal striatum' usually includes the caudate and

putamen with the nucleus accumbens in the ‘ventral striatum’. Abbreviations: SNr/SNc: substantia nigra pars reticulata/compacta, GP: globus pallidus, STN: sub thalamic nucleus, VTA: ventral tegmental area.

In humans the hippocampus has been specifically implicated in accurate spatial navigation, both in terms of the effects of lesions and metabolic activity during virtual navigation, although showing a greater (right) lateralization of function than in rodents (see Burgess et al., 2002 for a review). Indeed, virtual reality analogues of behavioral tests in rodents show similar associations of hippocampal and striatal activity with place and response learning respectively (Doeller et al., 2008; Hartley et al., 2003; Iaria et al., 2003).

Finally, we note the importance of the parietal cortex in spatial processing. Damage to the posterior parietal cortex often results in optic ataxia – impaired visuospatial coordination of reaching and grasping, consistent with the presence there of neurons tuned to encode the location of visual attention and aspects of reaching and grasping (e.g. Hwang et al., 2014). A common consequence of unilateral damage to the parietal cortex (most often on the right side) is a clinical syndrome known as *unilateral neglect:* an impairment in noticing or paying attention to objects and events in the contralateral hemifield, or the contralateral side of objects. Patients suffering from unilateral neglect can experience *representational neglect*, affecting their spatial imagery and memory performance, as distinct from the more common perceptual neglect (Bisiach and Luzzati, 1978), with neglect in imagery (rather than perception) being more strongly linked to deficits in navigation (Guariglia et al., 2005).

3. Neuronal codes

Single neuron recordings from the hippocampus of freely moving rats showed the existence of ‘place cells’ in regions CA1 and CA3: cells that fire only when the animal is within a limited region of the environment (the 'place field'; O’Keefe and Dostrovsky, 1971; Figure 2A). Place cell firing patterns are established very rapidly when an animal enters an environment for the first time and are stable over the course of several days (Thompson and Best, 1990). These firing patterns ‘remap’ between different environments (Muller and Kubie, 1987), but are robust to smaller changes such as eliminating a subset of environmental cues (O’Keefe and Conway, 1978; Quirk and Muller, 1990), a capacity for “pattern completion” associated with synaptic plasticity in CA3 (Marr, 1971; Nakazawa et al., 2002). Non-visual cues, like olfactory traces or auditory signals, can also contribute to self-location and are sufficient to guide behavior if visual cues are not accessible (Maaswinkel and Whishaw, 1999; Wallace et al., 2002). Place cells have subsequently been identified in a wide range of mammals including bats and humans (Ekstrom et al., 2003; Ulanovsky and Moss, 2007).

A complementary representation is provided by “head direction cells”. These neurons signal the orientation of the animal’s head in the horizontal plane, and are tuned to a narrow range of head directions centered on a preferred firing direction (Taube et al., 1990; Figure 2C). These cells were first reported in the dorsal presubiculum (Ranck, 1985; Taube et al., 1990) and later in a network of structures along the classic Papez circuit, including the thalamic nuclei (Mizumori and Williams, 1993; Taube, 1995), mammillary bodies (Stackman and Taube, 1998), and entorhinal cortex (Sargolini and al., 2006). Interestingly, if two cells share a preferred firing direction in one environment, they will continue to respond in the same way in a second environment even if the absolute firing direction of both cells may have changed (i.e. they remain “in register”).

A third type of spatial cell, “grid cells”, are found in the medial entorhinal cortex (mEC; Hafting et al., 2005) and subicular complex (Boccara et al., 2010). They share some similarities with place cells, but have multiple firing fields arranged on an equilateral triangular grid that covers the environment (Figure 2B). Grid cells appear to be grouped into functional clusters within mEC that share similar characteristics: neighboring cells possess the same grid orientation and scale, having only a different translational offset, while the grid scale increases ventrally along the mEC in discrete steps (Barry et al., 2007; Stensola et al., 2012). Moreover, similarly to head direction cells, their relative position is maintained even after environmental manipulations that change or disrupt the fields of individual cells (Stensola et al., 2012; Yoon et al., 2013).

The sources of information that dictate these spatial responses can be divided into self-motion and environmental sensory inputs. The strong intrinsic organization of the firing patterns of head direction and grid cells, irrespective of the sensory environment, suggests a significant influence of self-motion on their firing patterns. Accordingly, models of head direction and grid cells often rely on continuous attractor dynamics via symmetrical recurrent connectivity (Fuhs and Touretzky, 2006; McNaughton et al., 2006; Zhang, 1996). In these models, the spatial representation is updated by self-motion (a process also known as ‘path integration’) via asymmetric interactions (Zhang, 1996), which can be achieved in an accurate manner by cells with a conjunctive representations of space and movement, such as head direction firing modulated by angular velocity (e.g. Stackman and Taube, 1998) or grid cells modulated by movement velocity (Sargolini and al., 2006), see (Burak and Fiete, 2009; Conklin and Eliasmith, 2005; Skaggs et al., 1994) for computational models.

However, all allocentric (i.e. world-centred) spatial signals relying on self-motion (or ‘path integration’) need resetting relative to the environment to avoid accumulating error.

Correspondingly, spatial firing patterns are strongly influenced by the environment. Distant visual cues, where available, have a controlling influence on head direction and the orientation of other spatial responses (e.g. Taube et al., 1990). The boundaries of an environment also appear to play an important role in determining the firing locations of place cells. Place cell firing patterns across manipulations of environmental shape reflect conjunctions of distances and allocentric directions to environmental boundaries (Hartley et al., 2000; O'Keefe and Burgess, 1996), whereas discrete intra-maze landmarks have relatively little influence (e.g. Cressant et al., 1997). The predicted 'boundary vector cells' mediating this information were subsequently found in subiculum (Barry et al., 2006; Lever et al., 2009) and mEC (Solstad et al., 2008; Figure 2D). Environmental boundaries also affect the firing pattern (Barry et al., 2007) and orientation (Krupic et al., 2015; Stensola et al., 2015) of grid cells, consistent with a role in reducing cumulative error (Hardcastle et al., 2015).

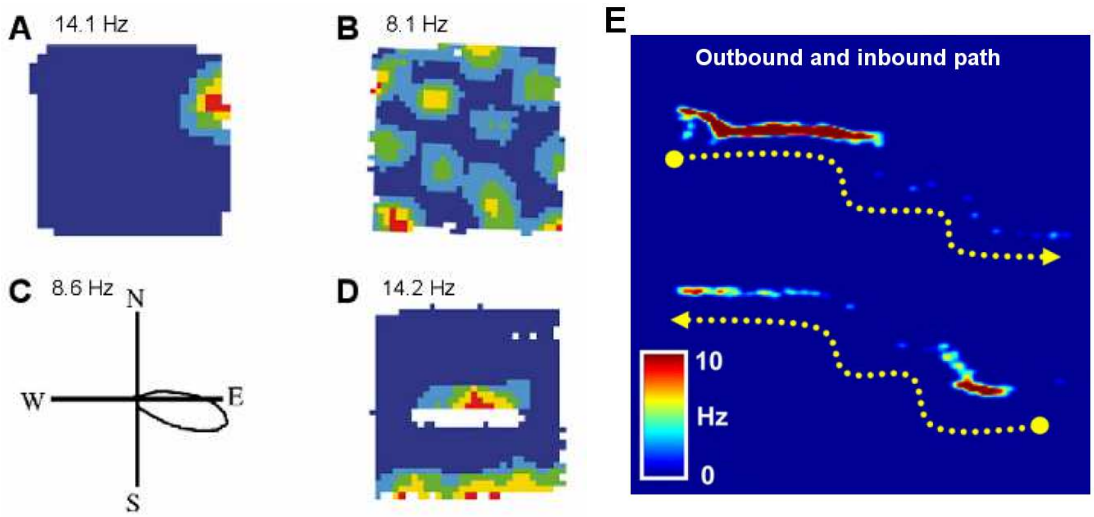

Figure 2. Examples of spatial cells associated with navigation, recorded in freely moving rats, showing firing rate as a function of location or head direction (peak firing rate shown in Hz). (A) Place cells, found in areas CA3 and CA1 of the hippocampus proper, typically fire in a restricted portion of the environment. (B) Grid cells, found in medial entorhinal cortex and pre- and parasubiculum, typically fire in a regular triangular array of locations. Directional grid cells or ''conjunctive'' cells, whose grid-like spatial firing is also modulated by head direction, are also found in these regions. (C) Head-direction cells, found in the presubiculum and deep layers of medial entorhinal cortex, typically fire for a narrow range of allocentric heading directions. (D) Boundary cells, found in subiculum and entorhinal cortex, typically fire at a specific distance from an environmental boundary along a specific allocentric direction. (E) A trajectory neuron from parietal cortex, shown for outbound (upper plots) and inbound (lower plots) traversals of a path (dashed yellow lines). A-C Adapted from (Hartley et al., 2014), E adapted from **(**Nitz, 2006**)**, with permission.

In contrast to the explicit representation of spatial information in the hippocampal formation, neuronal activity in the striatum is more strongly influenced by task stage (Barnes et al., 2005), being modulated by choice points, reward delivery and stereotyped egocentric responses (Schmitzer-Torbert and Redish, 2004; Berke et al., 2009). However, these responses do not

specifically encode route-trajectory information. Evidence for this type of information is found in the posterior parietal cortex (PPC) of navigating rats (McNaughton et al., 1994; Nitz, 2006; Figure 2E). Interestingly, in a very structured environment, grid cell firing patterns become more trajectory dependent (Derdikman et al., 2009) suggesting an influence from parietal cortex combined with a strong contextual modulation (Whitlock et al., 2012). More generally, parietal neurons tend to code all phases of the action sequences used to solve or plan a task (Fogassi et al., 2005; Harvey et al., 2012).

Posterior parietal neurons in monkeys can exhibit conjunctive 'gain field' responses tuned to visual (retinotopic) receptive fields but modulated by eye, head, or body position (Andersen, 1995; Snyder et al., 1998). These responses may allow determination of the location of visual objects relative to the body or in a world-referenced frame (Pouget and Sejnowski, 1997). In particular, area 7a, which contains neurons with world-referenced gain fields (Snyder et al., 1998) projects to the parahippocampal gyrus and presubiculum, and so may allow translation between egocentric parietal representations and allocentric medial temporal representations (Burgess et al., 2001; Byrne et al., 2007).

4. Systems neuroscience of spatial learning

The dependence of place cell firing on environmental boundaries rather than intra-maze landmarks is also reflected in hippocampal-dependent navigation. Pearce et al. (1998) adapted the water maze by adding a local cue at a fixed bearing from the submerged escape platform (Figure 3A). Rats learn relatively direct paths to the goal over the course of a few trials. After four trials (one session), the escape platform and the landmark are moved together to a new location. Rats with and without hippocampal lesions both are able to reach the hidden platform but present distinct performance curves. Hippocampal lesion animals quickly locate the platform on the first trial of a new session, using the intramaze landmark as a cue (i.e. following a 'piloting' strategy), whereas the control animals are slower, continuing to search at the previous location in the maze (i.e. following a cognitive map strategy). On the other hand, control animals learn the new location within each session, and out-perform the lesioned animals by the fourth trial of the session. Thus, the hippocampus appears to support learning of the platform location relative to the maze, rather than the landmark, which can hinder performance when the platform location is moved. Moving the maze relative to the testing room confirms the rats are using the boundary of the maze in combination with distal cues for orientation (e.g. Hamilton et al., 2007).

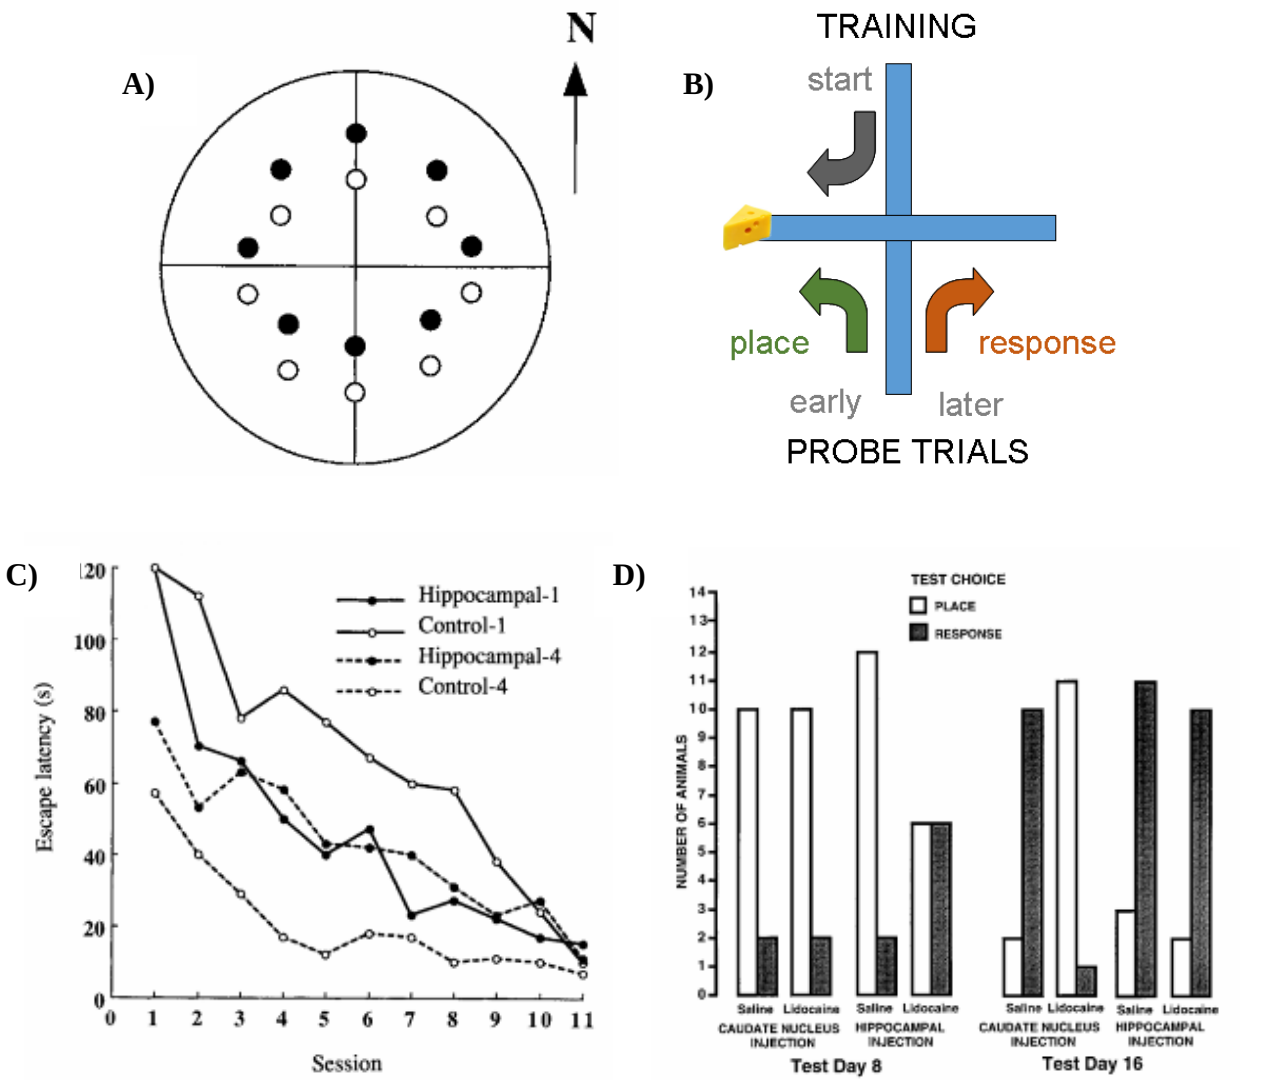

Figure 3. A) Schematic of the water maze used by Pearce et al., (1998). A submerged platform (white circle) and intra-maze landmark (black circle) are placed in the maze at 8 different locations on different sessions, but always with a fixed offset from each other. C) The mean escape latency for the first (continuous line) and fourth (dashed line) trials in each of the 11 sessions for the control and hippocampal lesion rats. Adapted from Pearce et al. (1998). The hippocampal rats perform better at the start of each session, the control rats at the end of each session. B) Schematic of the plus maze used by Packard and McGaugh, (1996). Rats learn to find the food placed at the end of an arm from the start location. In unrewarded probe trials the rat's starting position is moved to the opposite side of the maze. D) Number of rats in the probe trials on days 8 and 16, with either lidocaine or saline injections in the caudate nucleus or hippocampus. White bars indicate those showing a "place" strategy (i.e., going to the rewarded location in the room), while dark bars indicate those showing a "response" strategy (i.e. making the rewarded body-turn). The place strategy is sensitive to hippocampal inactivation on day 8, the response strategy to caudate inactivation on day 16. Adapted from Packard and McGaugh (1996).

The distinct styles of learning supported by hippocampal and striatal systems are further illustrated by experiments using a "plus" maze (Packard and McGaugh, 1996), after (Cohen et al., 1971). Rats were trained to approach a consistently baited arm in the plus maze, starting from the stem (Figure 3B). After several days a single probe trial was given, in which rats were placed

in a start arm opposite that used in training (see Figure 3). Control rats displayed “place learning” (i.e. going to the same allocentric location in the maze) on the Day 8 probe trial and “response learning” (i.e. making the same body turn) on the Day 16 probe trial, indicating that with extended training there is a shift in the systems controlling behavior. Supporting this interpretation, rats with inactivation of the striatum displayed place learning on both Day 8 and Day 16 probe trials, whereas rats with inactivation of the dorsal hippocampus showed no preference for place or response learning on the Day 8 probe trial, but displayed response learning on the Day 16 probe trial. Thus, it seems that response learning, i.e. association of reward with a body-turn, depends on the striatum while place learning, i.e. association of reward with an environmental location, depends on the hippocampus.

4.1 Learning rules in spatial navigation.

There is a long history of debate concerning the nature of spatial learning, spanning from the proponents of stimulus-response associative learning mechanisms driven by trial and error (Hull, 1943; Mackintosh, 1983; Rescorla and Wagner, 1972) to the proponents of incidental learning of internal representations capable of supporting cognition (Tolman, 1948). These arguments are brought to current thinking on spatial navigation in terms of reinforcement learning (RL) based on prediction error (Sutton and Barto, 1981; Foster et al., 2000) and the proposal that the hippocampus is a “cognitive map” (O’Keefe and Nadel, 1978) or relational (Cohen and Eichenbaum, 1993) or episodic (Hirsh, 1974; Scoville and Milner, 1957) memory system.

The chief characteristic of RL is that it relies on prediction error, i.e. the difference between actual reward and expectation based on experience, as opposed to a memory system that relies on incidental one-shot association. There are several consequences of relying on a single prediction error signal. One is that only the amount of future reward associated with a choice can be used to direct behavior, but not the type of reward. To include behavior that can be ‘goal-directed’ (i.e., aimed at a specific type of reward) requires a more elaborate model of the world, perhaps coming closer to the idea of a cognitive map. A second consequence is that learning outcomes based on multiple cues will show “blocking” and “overshadowing” (Kamin, 1969; Pavlov, 1927) between cues. Thus, if a first stimulus already fully predicts reward, learning about a second stimulus that might also predict reward will be “blocked” as there is no prediction error, and partial association of one stimulus to reward reduces the strength of association of a second concurrent cue to reward.

A recent experiment used a virtual reality adaptation of the Pearce et al. (1998) rodent experiment to test whether different types of learning occur within hippocampal and striatal systems in

humans, following (Hirsh, 1974; O'Keefe and Nadel, 1978). Using fMRI, this experiment showed that learning object-locations relative to the environmental boundary correlated with hippocampal activity, whereas learning object-locations relative to an intra-maze landmark correlated with activity in striatum (and also parietal cortex; Doeller et al., 2008). Moreover, parallel behavioral experiments showed that learning relative to an intra-maze landmark is blocked and overshadowed by a second cue (whether a landmark or boundary), whereas learning relative to an environmental boundaries is not – occurring incidentally to other cues (Doeller and Burgess, 2008).

Taken together, these results suggest that the striatum uses RL to associate actions to specific stimuli or landmarks that predict reward (including good performance in the case of conscientious human participants). Whereas the hippocampus forms incidental associations between objects (which might include, but is not restricted to, rewarding objects) and the environmental locations in which they are encountered. The diverse functions and learning rules of these two systems beg the question of how they interact to support a common behavioral outcome. In this context, we note that a large body of research has highlighted the role of the prefrontal cortex in the control and organization of goal-directed behavior (Tremblay and Schultz, 1999; Watanabe, 1996), the monitoring of ongoing voluntary action sequences (Gehring and Knight, 2000), the planning and selection of appropriate actions based on anticipated reward (Petrides, 1995; Rowe et al., 2000), and the ability to learn the contingency between actions and specific outcomes (Balleine and Dickinson, 1998).

The rodent medial prefrontal cortex (mPFC) comprises the ventral infralimbic cortex underneath the more dorsal prelimbic and anterior cingulate regions (Fisk and Wyss, 1999). The former region projects to a variety of limbic and autonomic regions, including the hypothalamus, the amygdala, and the shell region of the nucleus accumbens (Berendse et al., 1992; Hurley et al., 1991; Sesack et al., 1989). In contrast, the more dorsal prelimbic region of PFC projects to core regions of the nucleus accumbens and to dorso-medial regions of the dorsal striatum (Berendse et al., 1992; Gorelova and Yang, 1997), and has reciprocal indirect connections with premotor and motor cortices (Bates and Goldman-Rakic, 1993; Lu et al., 1994). Direct projections to the mPFC stem from the CA1/subiculum of the ventral part of the hippocampus (vHPC) (Jay and Witter, 1991; van Strien et al., 2009). Cells in mPFC can exhibit location specific firing and lesion in vHPC disrupts they goal-related activity (Burton et al., 2009). Anatomical connectivity suggests that the mPFC could be capable of integrating information from brain regions mediating appetitive and emotional motivation, and goal-directed and habitual responses. Consistent with this idea, Doeller et al. (2008) found that, while activity in either hippocampus or striatum

indicated use of the corresponding strategy, increased mPFC activity was seen when both systems were similarly active – suggesting a role in mediating between them.

5. Towards a computational account of spatial navigation

In this section we first briefly review classical computational models of how the hippocampus and the striatum might support spatial navigation, and then describe recent developments that try to reproduce the experimental and neurophysiological data reviewed above.

A number of earlier models of navigation combine the unsupervised learning of place representation with variants of RL (e.g. Brown and Sharp, 1995; Foster et al., 2000). However, this is somewhat at odds with the idea of a distinct rapid incidental learning system in the hippocampus in combination with a slower RL mechanism (see above and Lengyel and Dayan, 2007; Sheynikhovich et al., 2009) and related proposals for complementary hippocampal and neocortical learning systems (Marr, 1971, 1970; McClelland et al., 1995). Viewed as a memory system, learning in the hippocampus need not be driven by reward – its function may be to represent experience (e.g., what is encountered where, for navigation) so that it is available to future planning in which new goals may be specified. In this case currently unrewarding stimuli may become important in future and so also need to be remembered. Nonetheless, encoding may be biased towards stimuli that are novel, for efficiency in encoding, or towards stimuli that precede a rewarding event (see e.g. Tse et al., 2007).

Another family of models (e.g., Blum and Abbott, 1996; Dollé et al., 2010; Martinet et al., 2011; Muller et al., 1996) utilize the hippocampus to build graph-like representations of the environment for use in path planning using activity propagation methods. The connectivity between place cells can in principal support multiple graph-like representations (Samsonovich and McNaughton, 1997), and propagating activity is observed in place cells (e.g. Johnson and Redish, 2007; Pfeiffer and Foster, 2013). However, one issue faced by these models concerns the use of experience-dependent graph learning, which would bias behavior towards well-learned routes: an outcome more associated with the striatum.

Taking inspiration from these previous works, below we outline a minimal cognitive architecture that satisfies the functional and biological constraints reviewed above. For simplicity, we consider four main components: sensory cortex, the hippocampus, the striatum and the prefrontal cortex, see Figure 4.

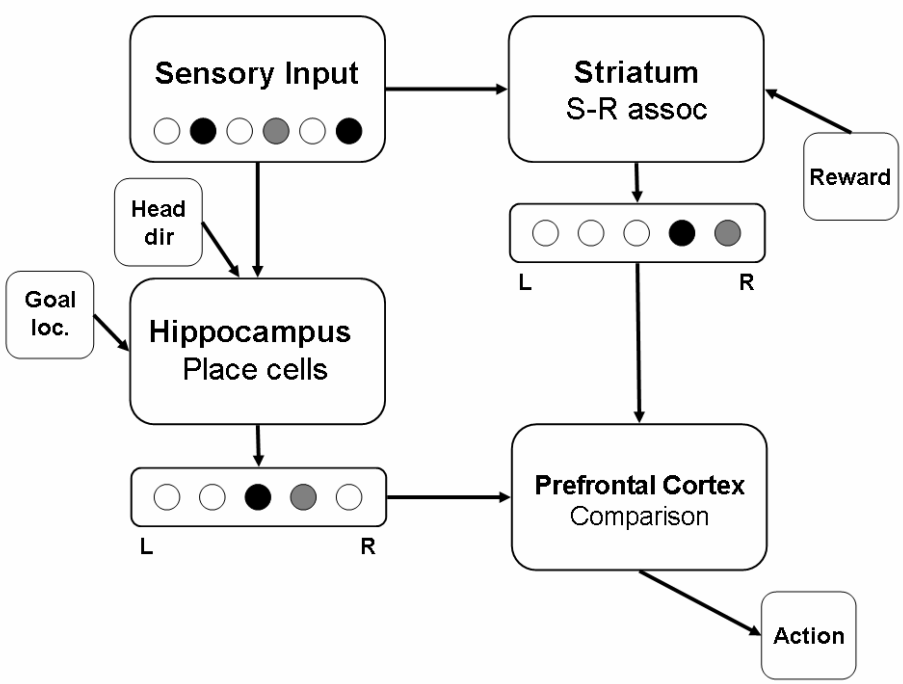

Figure 4. Schematic representation of a minimal circuit for two of the main mechanisms that guide spatial navigation: the hippocampus, providing the "cognitive map" with information about locations for goal-directed decision making, and the striatum that learns stimulus-response associations. Note that the same sensory input is used in different ways by the two systems.

We assume that sensory cortex supports a representation of the quality and distance to objects within the view field relative to the position and direction of the head (i.e. in an egocentric frame of reference). This sensory information, or "sensory snapshot," reaches both hippocampal and striatal systems. Below we describe the different ways in which the two systems would process this information in the tasks described above (Figure 3).

As noted above, the dorsal striatum is thought to be involved in reinforcement learning of stimulus-response associations learning. That is, it may learn sensory motor associations using the mismatch ('prediction-error') between the outcome expected by a ventral striatal 'critic' and the actual outcome in order to produce (statistically) correct associations, see Box 1. Thus, the dorsal and ventral striatum respective may form an "actor-critic" architecture for reinforcement learning (e.g., O'Doherty et al., 2004; Bornstein and Daw, 2011; Pennartz et al., 2011), in which the prediction error signal is encoded by variation in dopamine (Schultz et al., 1997).

Given the egocentric nature of the sensory input, in which distal orientation cues are not specifically salient, the striatum will learn to associate egocentric sensory representations with egocentric actions that lead to reward. The most salient (reward-predictive) feature of the Pearce

et al. (1998) version of the water maze will be the intra-maze landmark due to its proximity to the submerged platform. Thus striatal control of behavior will direct the animal to search near to the landmark. Equally, at the choice point of the plus maze used by Packard and McGaugh (1996), the egocentric sensory input will be similar from both training and probe directions. Thus, under striatal control of behavior, the sensory input associated with a specific action (body turn) during learning will also be driven by the sensory input at the choice point during probe trials.

By contrast, we suppose that the hippocampal system receives head direction information as well as sensory snapshots, so that boundary-vector cell responses can be formed to drive an allocentric place cell representation. Thus, the hippocampus represents states in a way that is relatively insensitive to body-orientation and to discrete intra-maze landmarks compared to extended environmental features. On encountering an object of interest (e.g. a food reward or submerged platform, or a neutral object), its location can be stored by unsupervised Hebbian learning of connections from place cells to "goal cells" coding for that type of object, which might be located in subiculum, ventral striatum or prefrontal cortex (e.g. Hok et al., 2005). The firing of these cells will then provide a value function for navigation to that object – increasing with proximity to it as more of the place cells with potentiated synapses become active (Burgess and O'Keefe, 1996), see Figure 5.

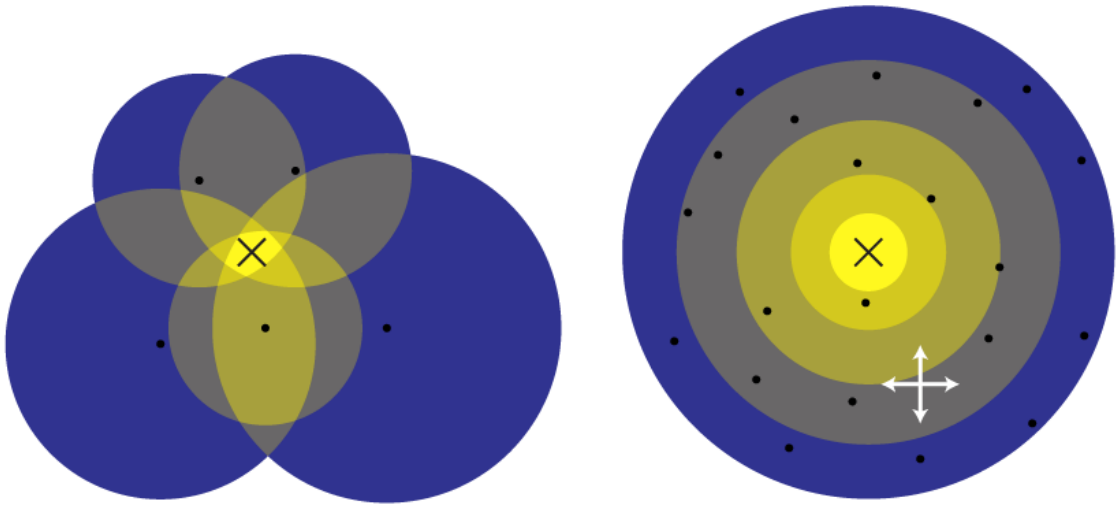

Figure 5. Schematic of how firing fields of different place cells can be combined to obtain a global value function that can be used to reach rewarding locations. The left panel shows an example with few neurons: when the rat reaches the rewarding site (marked with and x) fast Hebbian learning takes place between place cells that are active in that location (represented by black dots surrounded by their blue receptive field) and a "goal cell" downstream (e.g. in subiculum, ventral striatum or prefrontal cortex). In this location the activity of the goal cell is maximal. When the rat moves away from this location, fewer of the place cells with potentiated synapses will be active and therefore the firing rate will be lower. The right panel shows an example with a great number of place cells: when the number of places cells projecting to the value neuron increases, the resulting value function becomes smother and more regular.

Although such goal cells would provide an indication of how close the animal is to the desired location, they would not indicate in which direction to move to reach it. However, the temporal characteristics of place cell firing ("theta phase precession," O'Keefe and Recce, 1993) are such that within each cycle of the theta rhythm, the location represented by the currently active subset of place cells "sweeps forward" from behind to in front of the animal (Burgess et al., 1994; Johnson and Redish, 2007; Skaggs et al., 1996). This potentially allows 'exploration' of the suitability of the current direction of travel. Similar "forward sweeps" are observed during non-locomotor behavior and sleeping (Diba and Buzsáki, 2007; Pfeiffer and Foster, 2013; Wilson and McNaughton, 1994), which may relate to consolidation and planning. Note that the proximity relationships between place cells, as read out by forward sweeps, implicitly encode a model of state-transitions that reflects spatial structure rather than relationships to reward.

In order to produce a motor action or simply a decision, given the potentially different outputs from hippocampal and striatal systems, a mechanism is necessary to compare these outputs and decide which is the most appropriate for the task at hand. One ideal candidate for this function is medial prefrontal cortex. We suppose that this area selects between the possible actions indicated by either system based on the slopes of the associated value functions as a proxy for the confidence or uncertainty of the outputs of either system (Daw et al., 2005; Keramati et al., 2011). Accordingly, the time course of learning in the Packard and McGaugh (1996) experiment indicates slower striatal learning and lower confidence early on, resulting in hippocampal control of behavior, but the striatal system eventually achieves greater confidence and eventually gains control of behaviour. Similarly, in the Pearce et al. (1998) experiment, the intact animals show hippocampal control during the first trial of a new session – responding relative to the boundary of the maze, oriented by distal cues, and so performing worse than hippocampal lesion animals due to the movement of the platform relative to the maze. On later trials within each session however, the lesioned animals have improved little, with the intact animals now out-performing them.

6. Discussion

Taking inspiration from behavioral, lesion/ inactivation, neuroimaging and electrophysiological studies together with existing computational models, we sought to outline a minimal cognitive architecture for spatial navigation. Principal functional components of this architecture include the basal ganglia / striatum and the hippocampal formation, the former using local, incremental, and statistically efficient reinforcement learning rules; and the latter using a one-shot incidental

learning rule. The striatum has previously been proposed to function as an “actor-critic” network supporting classical and instrumental conditioning (O’Doherty et al., 2004; Yin et al., 2005). In our case, it is applied to navigation, using sensory input to provide information about states. By itself, the actor-critic can efficiently learn to solve key experimental tasks, such as the water maze and the plus maze, but its learning is slow and behavior is characterized by the egocentric nature of the sensory state information and the coding of actions.

By contrast, the hippocampus possesses a goal-independent representation of space that is learned rapidly, perhaps reflecting its role in episodic memory. This is obtained by having place cells initially driven in a fixed feed-forward manner by their proximity relations to environmental boundaries (see e.g. Hartley et al., 2000), followed by slower adjustments as an environment becomes familiar (see e.g. Barry and Burgess, 2007; Lever et al., 2002) and the potentially related adjustment of grid cells (Barry et al., 2012). When an object or goal is encountered, “goal” cells can be formed by rapid incidental Hebbian learning in connections from place cells, so that their activity can provide gradients that can be used for goal-directed navigation (Burgess and O’Keefe, 1996, see also Foster et al., 2000). Behaviour in this case is characterized by the allocentric nature of the representations of current and goal locations.

These two systems appear to play complementary roles at different stages of spatial learning. The hippocampus provides an initial rapid associative memory for associations between a goal (or a neutral object) and its context (environmental location in this case) that can guide goal-directed navigation in response to an explicit desire for that goal. However, the contextual association (in this case place cells) provides a similarity gradient that might not be sufficient to support navigation within complex environments. By contrast, the striatal reinforcement learning mechanism is capable of learning the statistics of the task over multiple trials, and thus potentially learning state-action trajectories that cannot be directly inferred from contextual similarity.

The outputs of these two systems must be coordinated by a third component, potentially corresponding to medial prefrontal cortex, see also (Chersi and Pezzulo, 2012; Dollé et al., 2010; Sheynikhovich et al., 2009). The most obvious mechanism for selection is a comparison of the “confidence” in movement directions signaled by either system, possibly utilizing the local gradient of the normalized value function expressed by each system. This proposed architecture relates to more general models of decision making (Daw et al., 2005; Dayan, 2009) which emphasize that goal-directed and habitual mechanisms of choice are linked to model-based and model-free methods of reinforcement learning, respectively.

However, the mapping of habitual and goal-directed (as in e.g. sensitive to devaluation of the type of reward) mechanisms onto hippocampal and striatal learning systems is not straight forward. Both mechanisms may exist outside of the hippocampus, with some authors implicating striatum in habitual learning and prefrontal cortex in goal-directed learning. In this context, the hippocampal system can be seen as a "third way" (Lengyel and Dayan, 2007). It is not just another model-based reinforcement learning system, but appears to be best characterized by one-shot incidental learning of specific examples, as consistent with its well-recognized role in episodic memory (Cohen and Eichenbaum, 1993; Hirsh, 1974; O'Keefe and Nadel, 1978; Scoville and Milner, 1957).

Nonetheless, the hippocampal system can certainly play the role of the 'model' in a model-based learning system. Specifically, the relationship between an explicit memory system and a means of generating imagery for planning is now being recognized. Thus, the hippocampus can be seen as the highest level in a generative model capable of consolidating memory in neocortex (Káli and Dayan, 2004, 2000). It can also be seen as a system for enabling information from long term memory to be rendered in parietal cortex as a coherent egocentric spatial scene consistent with a single viewpoint (Burgess et al., 2001; Byrne et al., 2007). In either case, the hippocampal system can be used to generate information corresponding to upcoming states for use in planning, as consistent with recent experiments relating the human hippocampus to the ability to perform spatially coherent imagery (Hassabis and Maguire, 2007; Schacter and Addis, 2007). The medial prefrontal and striatal areas appears to be involved in representing the reward value of elements of imagined scenarios (Benoit et al., 2014; Lin et al., 2015).

Within the view of the hippocampus as a generative model or means of imagining future states, the imagined movement of the agent could be achieved by iterative interactions with parietal cortex (Byrne et al., 2007; Chersi et al., 2013), or by the generation of "forward sweeps" of place cell representations either during theta states (Burgess et al., 1994; Johnson and Redish, 2007; Skaggs et al., 1996) or off-line "replay" (Pfeiffer and Foster, 2013; Wilson and McNaughton, 1993). These schemes can be seen as architectures for iterative decision making (e.g., Penny et al., 2013), and may enable the striatum to access the expected value associated with specific environmental locations (Lansink et al., 2008; van der Meer and Redish, 2009).

In conclusion, we hope to have shown that spatial navigation can serve as a model system in cognitive neuroscience, in which specific representations, learning rules and control strategies can be inferred from the vast experimental literature that exists across many species, including

humans. In this brief review we have attempted to outline a minimal cognitive architecture consistent with the most obvious of these inferences, both to demonstrate its utility as a model system and to encourage further theoretical and experimental elaboration.

Box 1: What we know (Main facts on the cognitive architecture of spatial navigation)

1. Multiple spatial representations have been identified in neuronal firing, and in behavior.
2. There is a good mapping between representations and brain systems.
3. These systems appear to combine constructively to support spatial memory, which implies that they can be selected between in an appropriate manner, e.g. according to a measure of 'confidence' in each system (e.g. slope of value function).
4. Different systems appear to use different learning rules, potentially reflecting optimization for different aspects of the task (1-shot learning for hippocampal episodic memory, prediction error for striatal action learning)

Box 2: What we need to know (Open questions)

1) Although the hippocampus and the striatum learn by means of two very different mechanisms, what is the influence of the former on the latter during learning? Does the information from the hippocampus directly contribute to the calculation of the prediction error in the striatal system, or is its influence only indirect via behavior (e.g. providing examples of successful routes early in learning). A puzzle here is that if hippocampal information is available to the 'critic', early hippocampal learning would block subsequent learning by the striatum, but if it is not, why does the boundary block learning to the landmark in Doeller and Burgess, (2008)?

2) Are the representations in striatum and parietal areas that could support landmark-related and response learning the same, or are there multiple such representations, and if so what are they like? The Packard and McGaugh (1996) study implies striatal encoding of an egocentric body turn, whereas the Pearce et al., (1998) study implies that hippocampal lesion animals can navigate to an allocentric vector from a landmark, however it is possible that their sub-optimal performance reflects a failure of allocentric representation (e.g. having to circle the landmark).

3) How complex and detailed can place cell "forward sweeps" be, and are they used in planning? It has been shown that in open fields and in simple mazes forward sweeps produce a small number of simple paths, but what happens when environments are complex? Will it be possible to observe branched sweeps? Or do these end at the following decision point?

4) How does the proposed one-shot learning in the hippocampus work? Encounters with non-rewarding objects are encoded in the context of a memory test, but not all coincidences of stimuli can be encoded: what determines which are and which are forgotten? Do they have to be followed by a significant (e.g. rewarding) event as per schema theory (Tse et al., 2007). Novelty

must boost learning (e.g. one-shot learning is presumably restricted to the first experience), how do the role of novelty in the hippocampus compare to that of reward prediction-error in the striatum, and do both involve dopamine (see e.g. Guitart-Masip et al., 2014)?
5) The use of hippocampal place cells for guiding navigation, e.g. via experience-dependent associations between them, requires that distances between places are inferred from moving between them, and so will be distorted by the routes taken rather than reflecting distance per se. A potential solution for large-scale vector navigation could be to make use of the intrinsic metrical regularity of grid cell firing patterns to infer the translation vector between locations (Bush et al., Neuron, submitted).
6) The vast amount of experimental data, both behavioral and neurophysiological, acquired in simple environments allows rather precise hypotheses on the functioning of spatial navigation mechanisms. On the contrary, information from complex mazes is mostly missing, thus it is not known how planning and decision making may work in these environments, and how this kind of knowledge is represented in the brain. It may be that combinations of memory-based and reinforcement-based learning is employed, such as eligibility traces or saliency based learning, or that grid cells play a more important role as they are able to provide global metrical representations of complex environments (Carpenter et al., 2015).

Acknowledgements: The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 604102 (Human Brain Project)" and from the Wellcome Trust.

Supplementary material: Reinforcement learning

The core idea of reinforcement learning is to predict the value of sensory states with respect to the current task and the optimal policy for making actions. When actual outcomes (primary rewards or punishment) fail to match the prediction, the difference (the *prediction error*) is used to modify the connections in order to improve predictions. By providing an internal prediction of value (or *critic*), the animal can learn to improve performance even when actual outcomes occur infrequently (solving the problem of temporal credit assignment).
One of the most common RL algorithms is the temporal difference (TD) learning rule which can be formalized in the following way:

$$\Delta W^{Sens\text{-}Cr} = \alpha [\, r_t + \gamma V(S_t) - V(S_{t-1}) \,] \qquad (1)$$

where $W^{Sens\text{-}Cr}$ are the weights between the representation of the (sensory) state and a critic (perhaps neurons in ventral striatum), $V(S_t)$ is the value of state S at time t (i.e. expected summed discounted future reward) represented by the activity of the critic neurons and direct function of the weights $W^{Sens\text{-}Cr}$, α is the learning rate, r_t is the actual reward at time t (usually 0, except at the rewarding site), and γ is the discount factor (which indicates the factor by which an immediate reward is preferable to a delayed reward). The term in parentheses is the "prediction error", indicating the difference in value at time t compared to that predicted at time *t-1*. In the simplest form of an "actor-critic" architecture, the same learning rule can be used to change connection weights between the neurons representing the state and neurons representing the action taken (an "actor," which might be implemented in dorsal striatum).

One way to handle the problem of sampling combinatorially large spaces of states and actions is to learn the value of all actions compatible with each given state. This method is known as Q-learning, with corresponding learning rule:

$$\Delta W^{Sens\text{-}Cr} = \alpha \cdot [\, r_t + \gamma \cdot max_a\, Q(S_t, a_t) - Q(S_{t-1}, a_{t-1}) \,] \qquad (2)$$

where $Q(S_t, a_t)$ is the value of action a_t when in state S_t , the rest of the parameters being the same as in equation 1. Here, the search for the maximum future Q-value as a function of all possible actions a is indicated by the max_a operation. Similarly as before, the Q-value is encoded by the activity of specific neurons (probably) in the dorsal striatum. As the number of learning trials becomes very high, the decision mechanism tends to optimality (Watkins and Dayan, 1992). We note that, in general, both "actor" and "critic" may need to incorporate more complex 'policies' for generating actions, and that the representations of states and actions may need to be continuous, rather than the discrete representations often used in analyses.

Bibliography

Andersen, R.A., 1995. Encoding of intention and spatial location in the posterior parietal cortex. Cereb. Cortex N. Y. N 1991 5, 457–469.

Balleine, B.W., Dickinson, A., 1998. Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. Neuropharmacology 37, 407–419.

Barnes, T.D., Kubota, Y., Hu, D., Jin, D.Z., Graybiel, A.M., 2005. Activity of striatal neurons reflects dynamic encoding and recoding of procedural memories. Nature 437, 1158–1161.

Barry, C., Burgess, N., 2007. Learning in a geometric model of place cell firing. Hippocampus 17, 786–800.

Barry, C., Bush, D., O'Keefe, J., Burgess, N., 2012. Models of Grid Cells and Theta Oscillations. Nature 488, 103–107.

Barry, C., Hayman, R., Burgess, N., Jeffery, K.J., 2007. Experience-dependent rescaling of entorhinal grids. Nat Neurosci 10, 682–684.
Barry, C., Lever, C., Hayman, R., Hartley, T., Burton, S., O'Keefe, J., Jeffery, K., Burgess, N., 2006. The boundary vector cell model of place cell firing and spatial memory. Rev. Neurosci. 17, 71–97.
Bates, J.F., Goldman-Rakic, P.S., 1993. Prefrontal connections of medial motor areas in the rhesus monkey. J. Comp. Neurol. 336, 211–228.
Benoit, R.G., Szpunar, K.K., Schacter, D.L., 2014. Ventromedial prefrontal cortex supports affective future simulation by integrating distributed knowledge. Proc. Natl. Acad. Sci. 111, 16550–16555.
Berendse, H.W., Galis-de Graaf, Y., Groenewegen, H.J., 1992. Topographical organization and relationship with ventral striatal compartments of prefrontal corticostriatal projections in the rat. J. Comp. Neurol. 316, 314–347.
Berke, J.D., Breck, J.T., Eichenbaum, H., 2009. Striatal Versus Hippocampal Representations During Win-Stay Maze Performance. J. Neurophysiol. 101, 1575–1587.
Bisiach, E., Luzzati, C., 1978. Unilateral neglect of representational space. Cortex 14, 129–133.
Blum, K.I., Abbott, L.F., 1996. A model of spatial map formation in the hippocampus of the rat. Neural Comput. 8, 85–93.
Boccara, C.N., Sargolini, F., Thoresen, V.H., Solstad, T., Witter, M.P., Moser, E.I., Moser, M.-B., 2010. Grid cells in pre- and parasubiculum. Nat. Neurosci. 13, 987–994.
Bornstein, A.M., Daw, N.D., 2011. Multiplicity of control in the basal ganglia: computational roles of striatal subregions. Curr. Opin. Neurobiol. 21, 374–380.
Brown, M.A., Sharp, P.E., 1995. Simulation of spatial learning in the Morris water maze by a neural network model of the hippocampal formation and nucleus accumbens. Hippocampus 5, 171–188.
Burak, Y., Fiete, I.R., 2009. Accurate Path Integration in Continuous Attractor Network Models of Grid Cells. PLoS Comput Biol 5, e1000291.
Burgess, N., 2014. Nobel Prize in physiology or medicine: A spatial model for cognitive neuroscience. Neuron 84, 1120–1125.
Burgess, N., Maguire, E.A., O'Keefe, J., 2002. The human hippocampus and spatial and episodic memory. Neuron 35, 625–641.
Burgess, N., Maguire, E.A., Spiers, H.J., O'Keefe, J., 2001. A temporoparietal and prefrontal network for retrieving the spatial context of lifelike events. NeuroImage 14, 439–453.
Burgess, N., O'Keefe, J., 1996. Neuronal computations underlying the firing of place cells and their role in navigation. Hippocampus 6, 749–762.
Burgess, N., Recce, M., O'Keefe, J., 1994. A model of hippocampal function. Neural Netw 7, 1065–1081.
Burton, B.G., Hok, V., Save, E., Poucet, B., 2009. Lesion of the ventral and intermediate hippocampus abolishes anticipatory activity in the medial prefrontal cortex of the rat. Behav. Brain Res. 199, 222–234.
Byrne, P., Becker, S., Burgess, N., 2007. Remembering the past and imagining the future: a neural model of spatial memory and imagery. Psychol. Rev. 114, 340–375.
Carpenter, F., Manson, D., Jeffery, K., Burgess, N., Barry, C., 2015. Grid cells form a global representation of connected environments. Curr. Biol. CB 25, 1176–1182.
Chersi, F., Donnarumma, F., Pezzulo, G., 2013. Mental imagery in the navigation domain: a computational model of sensory-motor simulation mechanisms. Adapt. Behav.

Chersi, F., Pezzulo, G., 2012. Using hippocampal-striatal loops for spatial navigation and goal-directed decision-making. Cogn. Process. 13 Suppl 1, S125–129.
Cohen, J.S., Laroche, J.P., Beharry, E., 1971. Response perseveration in the hippocampal lesioned rat. Psychon Sci 23, 221–223.
Cohen, N.J., Eichenbaum, H., 1993. Memory, Amnesia, and The Hippocampal System. MIT Press, Cambridge, MA.
Conklin, J., Eliasmith, C., 2005. A Controlled Attractor Network Model of Path Integration in the Rat. J. Comput. Neurosci. 18, 183–203.
Corbit, L.H., Balleine, B.W., 2003. The role of prelimbic cortex in instrumental conditioning. Behav. Brain Res. 146, 145–157.
Cressant, A., Muller, R.U., Poucet, B., 1997. Failure of centrally placed objects to control the firing fields of hippocampal place cells. J Neurosci 17, 2531–2542.
Daw, N.D., Niv, Y., Dayan, P., 2005. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. Nat. Neurosci. 8, 1704–1711.
Dayan, P., 2009. Goal-directed control and its antipodes. Neural Netw 22, 213–219.
Derdikman, D., Whitlock, J.R., Tsao, A., Fyhn, M., Hafting, T., Moser, M.-B., Moser, E.I., 2009. Fragmentation of grid cell maps in a multicompartment environment. Nat. Neurosci. 12, 1325–1332.
Diba, K., Buzsáki, G., 2007. Forward and reverse hippocampal place-cell sequences during ripples. Nat. Neurosci. 10, 1241–1242.
Doeller, C.F., Burgess, N., 2008. Distinct error-correcting and incidental learning of location relative to landmarks and boundaries. Proc. Natl. Acad. Sci. U. S. A. 105, 5909–5914.
Doeller, C.F., King, J.A., Burgess, N., 2008. Parallel striatal and hippocampal systems for landmarks and boundaries in spatial memory. Proc. Natl. Acad. Sci. U. S. A. 105, 5915–5920.
Dollé, L., Sheynikhovich, D., Girard, B., Chavarriaga, R., Guillot, A., 2010. Path planning versus cue responding: a bio-inspired model of switching between navigation strategies. Biol. Cybern. 103, 299–317.
Ekstrom, A.D., Kahana, M.J., Caplan, J.B., Fields, T.A., Isham, E.A., Newman, E.L., Fried, I., 2003. Cellular networks underlying human spatial navigation. Nature 425, 184–188.
Fisk, G.D., Wyss, J.M., 1999. Associational projections of the anterior midline cortex in the rat: intracingulate and retrosplenial connections. Brain Res. 825, 1–13.
Fogassi, L., Ferrari, P.F., Gesierich, B., Rozzi, S., Chersi, F., Rizzolatti, G., 2005. Parietal lobe: from action organization to intention understanding. Science 308, 662–667.
Foster, D.J., Morris, R.G., Dayan, P., 2000. A model of hippocampally dependent navigation, using the temporal difference learning rule. Hippocampus 10, 1–16.
Fuhs, M.C., Touretzky, D.S., 2006. A Spin Glass Model of Path Integration in Rat Medial Entorhinal Cortex. J. Neurosci. 26, 4266–4276.
Gehring, W.J., Knight, R.T., 2000. Prefrontal-cingulate interactions in action monitoring. Nat. Neurosci. 3, 516–520.
Gorelova, N., Yang, C.R., 1997. The course of neural projection from the prefrontal cortex to the nucleus accumbens in the rat. Neuroscience 76, 689–706.
Guariglia, C., Piccardi, L., Iaria, G., Nico, D., Pizzamiglio, L., 2005. Representational neglect and navigation in real space. Neuropsychologia 43, 1138–1143.
Guitart-Masip, M., Duzel, E., Dolan, R., Dayan, P., 2014. Action versus valence in decision making. Trends Cogn. Sci. 18, 194–202.

Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., Moser, E.I.I., 2005. Microstructure of a spatial map in the entorhinal cortex. Nature 801–806.
Hamilton, D.A., Akers, K.G., Weisend, M.P., Sutherland, R.J., 2007. How do room and apparatus cues control navigation in the Morris water task? Evidence for distinct contributions to a movement vector. J. Exp. Psychol. Anim. Behav. Process. 33, 100–114.
Hardcastle, K., Ganguli, S., Giocomo, L.M., 2015. Environmental Boundaries as an Error Correction Mechanism for Grid Cells. Neuron 86, 827–839.
Hartley, T., Burgess, N., Lever, C., Cacucci, F., O'Keefe, J., 2000. Modeling place fields in terms of the cortical inputs to the hippocampus. Hippocampus 10, 369–379.
Hartley, T., Lever, C., Burgess, N., O'Keefe, J., 2014. Space in the Brain: how the hippocampal formation supports spatial cognition. Phil Trans Roy Soc B 369, 1635.
Hartley, T., Maguire, E.A., Spiers, H.J., Burgess, N., 2003. The well-worn route and the path less traveled: distinct neural bases of route following and wayfinding in humans. Neuron 37, 877–888.
Harvey, C.D., Coen, P., Tank, D.W., 2012. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. Nature 484, 62–68.
Hassabis, D., Maguire, E.A., 2007. Deconstructing episodic memory with construction. Trends Cogn. Sci. 11, 299–306.
Hirsh, R., 1974. The hippocampus and contextual retireval of information from memory: a theory. Behav. Biol. 12, 421–444.
Hok, V., Save, E., Lenck-Santini, P.P., Poucet, B., 2005. Coding for spatial goals in the prelimbic/infralimbic area of the rat frontal cortex. Proc. Natl. Acad. Sci. U. S. A. 102, 4602–4607.
Hull, C.L., 1943. Principles of behavior: an introduction to behavior theory. Appleton-Century, Oxford, England.
Hurley, K.M., Herbert, H., Moga, M.M., Saper, C.B., 1991. Efferent projections of the infralimbic cortex of the rat. J. Comp. Neurol. 308, 249–276.
Hwang, E.J., Hauschild, M., Wilke, M., Andersen, R.A., 2014. Spatial and Temporal Eye–Hand Coordination Relies on the Parietal Reach Region. J. Neurosci. 34, 12884–12892.
Iaria, G., Petrides, M., Dagher, A., Pike, B., Bohbot, V.D., 2003. Cognitive strategies dependent on the hippocampus and caudate nucleus in human navigation: variability and change with practice. J. Neurosci. Off. J. Soc. Neurosci. 23, 5945–5952.
Jay, T.M., Witter, M.P., 1991. Distribution of hippocampal CA1 and subicular efferents in the prefrontal cortex of the rat studied by means of anterograde transport of Phaseolus vulgaris-leucoagglutinin. J. Comp. Neurol. 313, 574–586.
Johnson, A., Redish, A.D., 2007. Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. J. Neurosci. Off. J. Soc. Neurosci. 27, 12176–12189.
Káli, S., Dayan, P., 2004. Off-line replay maintains declarative memories in a model of hippocampal-neocortical interactions. Nat. Neurosci. 7, 286–294.
Káli, S., Dayan, P., 2000. Hippocampally-dependent consolidation in a hierarchical model of neocortex, in: NIPS. pp. 24–30.
Kamin, L.J., 1969. Predictability, surprise, attention and conditioning. Punishment and aversive behavior, in: Punishment and Aversive Behavior. Appleton-Century-Crofts, New York, pp. 279–296.
Keramati, M., Dezfouli, A., Piray, P., 2011. Speed/accuracy trade-off between the habitual and the goal-directed processes. PLoS Comput Biol 7, e1002055.

Krupic, J., Bauza, M., Burton, S., Barry, C., O'Keefe, J., 2015. Grid cell symmetry is shaped by environmental geometry. Nature 518, 232–235.
Lansink, C.S., Goltstein, P.M., Lankelma, J.V., Joosten, R.N.J.M.A., McNaughton, B.L., Pennartz, C.M.A., 2008. Preferential reactivation of motivationally relevant information in the ventral striatum. J Neurosci 28, 6372–6382.
Lengyel, M., Dayan, P., 2007. Hippocampal Contributions to Control: The Third Way. NIPS 889–896.
Lever, C., Burgess, N., Cacucci, F., O'Keefe, J., 2002. What can the hippocampal representation of environmental geometry tell us about Hebbian learning. Biol Cybern 87, 356–372.
Lever, C., Burton, S., Jeewajee, A., O'Keefe, J., Burgess, N., 2009. Boundary Vector Cells in the subiculum of the hippocampal formation. J. Neurosci. Off. J. Soc. Neurosci. 29, 9771–9777.
Lin, W.-J., Horner, A.J., Bisby, J.A., Burgess, N., 2015. Medial Prefrontal Cortex: Adding Value to Imagined Scenarios. J. Cogn. Neurosci. 1–11.
Lu, M.T., Preston, J.B., Strick, P.L., 1994. Interconnections between the prefrontal cortex and the premotor areas in the frontal lobe. J. Comp. Neurol. 341, 375–392.
Maaswinkel, H., Whishaw, I.Q., 1999. Homing with locale, taxon, and dead reckoning strategies by foraging rats: sensory hierarchy in spatial navigation. Behav. Brain Res. 99, 143–152.
Mackintosh, N.J., 1983. Conditioning and Associative Learning. Oxford University Press, Oxford.
Marr, D., 1971. Simple memory: a theory for archicortex. Philos Trans R Soc Lond B Biol Sci 262, 23–81.
Marr, D., 1970. A theory for cerebral neocortex. Proc. R. Soc. Lond. B Biol. Sci. 176, 161–234.
Martinet, L.-E., Sheynikhovich, D., Benchenane, K., Arleo, A., 2011. Spatial learning and action planning in a prefrontal cortical network model. PLoS Comput. Biol. 7, e1002045.
McClelland, J.L., McNaughton, B.L., O'Reilly, R.C., 1995. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. Psychol. Rev. 102, 419–457.
McNaughton, B.L., Battaglia, F.P., Jensen, O., Moser, E.I., Moser, M.-B., 2006. Path integration and the neural basis of the "cognitive map." Nat. Rev. Neurosci. 7, 663–678.
McNaughton, B.L., Mizumori, S.J., Barnes, C.A., Leonard, B.J., Marquis, M., Green, E.J., 1994. Cortical representation of motion during unrestrained spatial navigation in the rat. Cereb. Cortex 4, 27–39.
Mizumori, S., Williams, J., 1993. Directionally selective mnemonic properties of neurons in the lateral dorsal nucleus of the thalamus of rats. J Neurosci 13, 4015–4028.
Morris, R.G., Garrud, P., Rawlins, J.N., O'Keefe, J., 1982. Place navigation impaired in rats with hippocampal lesions. Nature 297, 681–683.
Muller, R., Stead, M., Pach, J., 1996. The hippocampus as a cognitive graph. J. Gen. Physiol. 107, 663–694.
Muller, R.U., Kubie, J.L., 1987. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. J Neurosci 7, 1951–1968.
Nakazawa, K., Quirk, M.C., Chitwood, R.A., Watanabe, M., Yeckel, M.F., Sun, L.D., Kato, A., Carr, C.A., Johnston, D., Wilson, M.A., Tonegawa, S., 2002. Requirement for hippocampal CA3 NMDA receptors in associative memory recall. Science 297, 211–218.
Nitz, D.A., 2006. Tracking route progression in the posterior parietal cortex. Neuron 49, 747–756.

O'Doherty, J., Dayan, P., Schultz, J., Deichmann, R., Friston, K., Dolan, R.J., 2004. Dissociable Roles of Ventral and Dorsal Striatum in Instrumental Conditioning. Science 304, 452–454.

O'Keefe, J., Burgess, N., 1996. Geometric determinants of the place fields of hippocampal neurons. Nature 381, 425–428.

O'Keefe, J., Conway, D.H., 1978. Hippocampal place units in the freely moving rat: why they fire where they fire. Exp. Brain Res. 31, 573–590.

O'Keefe, J., Dostrovsky, J., 1971. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. Brain Res. 34, 171–175.

O'Keefe, J., Nadel, L., 1978. The Hippocampus as a Cognitive Map. Oxford University Press.

O'Keefe, J., Recce, M.L., 1993. Phase relationship between hippocampal place units and the EEG theta rhythm. Hippocampus 3, 317–330.

Olton, D.S., Walker, J.A., Gage, F.H., Johnson, C.T., 1977. Choice behavior of rats searching for food. Learn. Motiv. 8, 315–331.

Packard, M.G., McGaugh, J.L., 1996. Inactivation of Hippocampus or Caudate Nucleus with Lidocaine Differentially Affects Expression of Place and Response Learning. Neurobiol. Learn. Mem. 65, 65–72.

Packard, M.G., McGaugh, J.L., 1992. Double dissociation of fornix and caudate nucleus lesions on acquisition of two water maze tasks: further evidence for multiple memory systems. Behav Neurosci 106, 439–446.

Pavlov, I.P., 1927. Conditioned Reflexes. Oxford Univ Press, Oxford.

Pearce, J.M., Roberts, A.D., Good, M., 1998. Hippocampal lesions disrupt navigation based on cognitive maps but not heading vectors. Nature 396, 75–77.

Pennartz, C.M.A., Ito, R., Verschure, P.F.M.J., Battaglia, F.P., Robbins, T.W., 2011. The hippocampal-striatal axis in learning, prediction and goal-directed behavior. Trends Neurosci. 34, 548–559.

Penny, W.D., Zeidman, P., Burgess, N., 2013. Forward and Backward Inference in Spatial Cognition 9, e1003383.

Petrides, M., 1995. Functional organization of the human frontal cortex for mnemonic processing. Evidence from neuroimaging studies. Ann. N. Y. Acad. Sci. 769, 85–96.

Pfeiffer, B.E., Foster, D.J., 2013. Hippocampal place-cell sequences depict future paths to remembered goals. Nature 497, 74–79.

Pouget, A., Sejnowski, T.J., 1997. A new view of hemineglect based on the response properties of parietal neurones. Philos. Trans. R. Soc. B Biol. Sci. 352, 1449–1459.

Quirk, G.J., Muller, R.U., 1990. The firing of hippocampal place cells in the dark depends on the rat's recent experience. J Neurosci 10, 2008–2017.

Ranck, J.B., 1985. Head direction cells in the deep cell layer of dorsal postsubiculum in freely moving rats, in: Electrical Activity of the Archicortex. Akademiai Kiado, Budapest.

Rescorla, R.A., Wagner, A.R., 1972. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement, in: Classical Conditioning II: Current Research and Theory. Appleton Century Crofts, New York, pp. 64–99.

Rowe, J.B., Toni, I., Josephs, O., Frackowiak, R.S., Passingham, R.E., 2000. The prefrontal cortex: response selection or maintenance within working memory? Science 288, 1656–1660.

Samsonovich, A., McNaughton, B.L., 1997. Path Integration and Cognitive Mapping in a Continuous Attractor Neural Network Model. J. Neurosci. 17, 5900–5920.

Sargolini, F., al., 2006. Conjunctive representation of position, direction, and velocity in entorhinal cortex. Science 312.
Schacter, D.L., Addis, D.R., 2007. The cognitive neuroscience of constructive memory: remembering the past and imagining the future. Philos. Trans. R. Soc. Lond. B. Biol. Sci. 362, 773–786.
Schultz, W., Dayan, P., Montague, P.R., 1997. A neural substrate of prediction and reward. Science 275, 1593–1599.
Scoville, W.B., Milner, B., 1957. Loss of recent memory after bilateral hippocampal lesions. J. Neuropsychiatry Clin. Neurosci. 12, 103–113.
Sesack, S.R., Deutch, A.Y., Roth, R.H., Bunney, B.S., 1989. Topographical organization of the efferent projections of the medial prefrontal cortex in the rat: an anterograde tract-tracing study with Phaseolus vulgaris leucoagglutinin. J. Comp. Neurol. 290, 213–242.
Sheynikhovich, D., Chavarriaga, R., Strösslin, T., Arleo, A., Gerstner, W., 2009. Is there a geometric module for spatial orientation? Insights from a rodent navigation model. Psychol. Rev. 116, 540–566.
Skaggs, W.E., Knierim, J.J., Kudrimoti, H., McNaughton, L.B., 1994. A model of the neural basis of the rat's sense of direction. NIPS 173–180.
Skaggs, W.E., McNaughton, B.L., Wilson, M.A., Barnes, C.A., 1996. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. Hippocampus 6, 149–172.
Snyder, L.H., Grieve, K.L., Brotchle, P., Andersen, R.A., 1998. Separate body- and world-referenced representations of visual space in parietal cortex. Nature 394, 887–891.
Solstad, T., Boccara, C.N., Kropff, E., Moser, M.-B., Moser, E.I., 2008. Representation of geometric borders in the entorhinal cortex. Science 322, 1865–1868.
Stackman, R.W., Taube, J.S., 1998. Firing properties of rat lateral mammillary single units: head direction, head pitch, and angular head velocity. J Neurosci 18, 9020–9037.
Stensola, H., Stensola, T., Solstad, T., Froland, K., Moser, M.-B., Moser, E.I., 2012. The entorhinal grid map is discretized. Nature 492.
Stensola, T., Stensola, H., Moser, M.-B., Moser, E.I., 2015. Shearing-induced asymmetry in entorhinal grid cells. Nature 518, 207–212.
Sutton, R.S., Barto, A.G., 1981. Toward a modern theory of adaptive networks: expectation and prediction. Psychol. Rev. 88, 135–170.
Taube, J.S., 1995. Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. J Neurosci 15, 70–86.
Taube, J.S., Muller, R.U., Ranck, J.B., 1990. Head-direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations. J Neurosci 10, 436–447.
Thompson, L.T., Best, P.J., 1990. Long-term stability of place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats. Brain Res 509, 299–308.
Tolman, E.C., 1948. Cognitive maps in rats and men. Psychol. Rev. 55, 189–208.
Tremblay, L., Schultz, W., 1999. Relative reward preference in primate orbitofrontal cortex. Nature 398, 704–708.
Tse, D., Langston, R.F., Kakeyama, M., Bethus, I., Spooner, P.A., Wood, E.R., Witter, M.P., Morris, R.G.M., 2007. Schemas and memory consolidation. Science 316, 76–82.
Ulanovsky, N., Moss, C.F., 2007. Hippocampal cellular network activity in freely moving echolocating bats. Nat Neurosci 10, 222–233.

Van der Meer, M.A.A., Redish, A.D., 2009. Covert Expectation-of-Reward in Rat Ventral Striatum at Decision Points. Front. Integr. Neurosci. 3, 1.
Van Strien, N.M., Cappaert, N.L.M., Witter, M.P., 2009. The anatomy of memory: an interactive overview of the parahippocampal–hippocampal network. Nat. Rev. Neurosci. 10, 272–282.
Voorn, P., Vanderschuren, L.J., Groenewegen, H.J., Robbins, T.W., Pennartz, C.M., 2004. Putting a spin on the dorsal-ventral divide of the striatum. Trends Neurosci 27, 468–474.
Wallace, D.G., Gorny, B., Whishaw, I.Q., 2002. Rats can track odors, other rats, and themselves: implications for the study of spatial behavior. Behav. Brain Res. 131, 185–192.
Watanabe, M., 1996. Reward expectancy in primate prefrontal neurons. Nature 382, 629–632.
Watkins, C.J.C.H., Dayan, P., 1992. Q-learning. Mach. Learn. 8, 279–292.
Whitlock, J.R., Pfuhl, G., Dagslott, N., Moser, M.-B., Moser, E.I., 2012. Functional split between parietal and entorhinal cortices in the rat. Neuron 73, 789–802.
Wilson, M.A., McNaughton, B.L., 1994. Reactivation of hippocampal ensemble memories during sleep. Science 265, 676–679.
Wilson, M.A., McNaughton, B.L., 1993. Dynamics of the hippocampal ensemble code for space. Science 261, 1055–1058.
Yin, H.H., Knowlton, B.J., 2006. The role of the basal ganglia in habit formation. Nat Rev Neurosci 7, 464–476.
Yin, H.H., Ostlund, S.B., Knowlton, B.J., Balleine, B.W., 2005. The role of the dorsomedial striatum in instrumental conditioning. Eur. J. Neurosci. 22, 513–523.
Yoon, K., Buice, M.A., Barry, C., Hayman, R., Burgess, N., Fiete, I.R., 2013. Evidence of low-dimensional continuous attractor dynamics in grid cells. Nat Neurosci 16, 1077–1084.
Zhang, K., 1996. Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. J. Neurosci. Off. J. Soc. Neurosci. 16, 2112–2126.