



Published in final edited form as:

Nat Rev Neurosci. 2016 August ; 17(8): 513–523. doi:10.1038/nrn.2016.56.

Over the river, through the woods: cognitive maps in the hippocampus and orbitofrontal cortex

Andrew M. Wikenheiser and

Intramural Research Program, National Institute on Drug Abuse, Baltimore, Maryland 21224, USA

Geoffrey Schoenbaum

Intramural Research Program, National Institute on Drug Abuse, Baltimore, Maryland 21224, USA; the Department of Anatomy and Neurobiology, University of Maryland, Baltimore, Maryland 21201, USA; and the Department of Neuroscience, Johns Hopkins University, Baltimore, Maryland 21205, USA

Abstract

The hippocampus and the orbitofrontal cortex (OFC) both have important roles in cognitive processes such as learning, memory and decision making. Nevertheless, research on the OFC and hippocampus has proceeded largely independently, and little consideration has been given to the importance of interactions between these structures. Here, evidence is reviewed that the hippocampus and OFC encode parallel, but interactive, cognitive ‘maps’ that capture complex relationships between cues, actions, outcomes and other features of the environment. A better understanding of the interactions between the OFC and hippocampus is important for understanding the neural bases of flexible, goal-directed decision making.

Despite possessing distinct neurochemical, anatomical and physiological properties, the hippocampal formation and the orbitofrontal cortex (OFC) have been ascribed broadly similar functional roles. For instance, a core feature of OFC function has been identified as “predicting the specific outcomes that should follow either sensory events or behavioral choices” (REF. 1), whereas others have conceptualized the hippocampal formation as a “system that facilitates predictions about upcoming events” (REF. 2). Although these statements are, of course, meant to provide very general descriptions of how each structure functions, the similarities are striking. Both structures are implicated in forming predictions about the future to support flexible behaviour and in leveraging general knowledge about the world rather than relying exclusively on specific previous experiences^{1,2}.

More formal theories have also arrived at similar functional specifications for these two structures. For decades, the hippocampus has been synonymous with mapping. Indeed, the cognitive map framework of hippocampal function — first proposed more than 40 years ago in the now classic 1978 work by O’Keefe and Nadel³ — is the most enduring theory of hippocampal information processing. Although this mapping function is frequently thought of in terms of spatial mapping, the cognitive map framework revisited the older concept of a

Competing interests statement

The authors declare no competing interests.

cognitive map, as proposed by Tolman⁴. This referred not to a literal map of space, but rather to an abstract map of causal relationships in the world: that is, a set of mental representations that binds external sensory features with internal motivational or emotional factors to form an integrated relational ‘database’ (REFS 4–6; BOX 1). Subsequent work has shown that, beyond the spatial tuning of its principal neurons, many aspects of hippocampal physiology are consistent with the cognitive map function that Tolman envisaged, and it is well established that the hippocampus has a role in encoding information about the world in a way that facilitates flexible and inferential cognitive processes^{7–12}.

Box 1

A modern twist on Tolman’s cognitive map

Tolman’s conception of learning was developed in reaction to the prominent stimulus–response theorizing that dominated in his day. In contrast to other researchers of the time, Tolman framed learning as an active process of extracting information from the world, rather than as a passive accumulation of associations imposed on the animal by the environment^{5,143–145}. Instead of learning individual action–outcome or cue–outcome relationships by storing specific instances of events, Tolman posited that animals track the underlying structure of the world in a map-like representation of causal associations. Tolman named this mental construct the cognitive map, an evocative description of the mental architecture he was proposing⁴. Just as a physical map allows one to plan novel or unique routes to a previously visited destination, a cognitive map would allow one to combine knowledge about causal relationships in the world in a manner that would enable one to derive novel and unique means of achieving outcomes. Importantly, although Tolman tested his ideas using spatial paradigms^{146,147}, he did not intend for cognitive maps to explain spatial planning alone. Instead, he envisioned a much more general system for creating schemas that encapsulate how the world works by tracking latent causal relationships between stimuli, actions, and outcomes.

Although much of Tolman’s thinking is now accepted¹⁴⁸, his ideas met resistance when they were first introduced, partially because of his difficulty in developing a theoretical framework that articulated his rather complex perspective¹⁴³. Although a mathematical description of the cognitive map eluded Tolman, subsequent advances in computational modelling of cognition and behaviour have expanded the range of processes than can be described mathematically. Many of the ideas that Tolman expressed have been subsumed by current models of learning and decision making. In particular, reinforcement-learning models⁶¹ and their progenitors from psychological learning theory^{149–151} have proved to be useful for quantitatively describing different forms of value learning and decision making. Such models distinguish between two fundamental forms of learning and decision making, often labelled ‘model-free’ and ‘model-based’. Model-free algorithms learn the value of actions but do not learn specific information about the sensory properties, identity or other features of outcomes. Model-based systems store a richer set of associations and capitalize on a world model that tracks how different states of the world are linked together and the specific identity of the outcomes those states contain^{152–154}.

Many behaviours that Tolman attributed to cognitive map function can be solved using model-based reasoning^{143,152,154,155}. The cognitive map can thus be defined as an associative structure that facilitates model-based learning and behaviour. This definition remains true to Tolman's thinking but also takes advantage of recent advances in computational modelling of complex behaviours to understand neural function. In line with this definition, we suggest that the cognitive map entails a number of components. First, there must be a mechanism for recognizing and categorizing the world into discrete states based on features that are relevant to current behavioural demands. Second, the cognitive map requires a means of learning and storing the relationships between world states (that is, how states are connected and how they are arranged relative to one another in the broader space of possible states). Third, the map must encode rich representations of outcomes that are associated with states (including their sensory features and identities) that can be used both to predict specific outcomes and, more generally, to estimate how 'good' or 'bad' states are expected to be in a way that reflects fluctuations in motivation, changes in the outcomes themselves and animals' current needs or goals. Finally, there must be a mechanism to use all of this information prospectively to construct novel plans for reaching goals and to predict the outcomes that are likely to follow from combinations of cues or states that have never previously been experienced. The consilience of these modern constructs with Tolman's older ideas suggests that the neural instantiation of the cognitive map might be distributed across the constellation of structures that support model-based, goal-directed behaviour, including the orbitofrontal complex and the hippocampus.

By contrast, the OFC has historically been associated with reward- and value-based behaviour^{13–20}. Recent thinking, however, has suggested that the OFC also has a cognitive-map-like role. Drawing on ideas from computational reinforcement-learning models, it has been proposed that a fundamental function of the OFC is to form and to maintain neural representations of task state: that is, a representation of all the relevant internal and external stimuli or features that define a particular situation in the world^{21,22}. Because this function requires the OFC to encode both features of the environment (including observable sensory properties and unobservable, implicit variables that must be inferred) and how relationships between those features might change in different situations, the OFC has been described as a cognitive map of task state²². Viewed from this perspective, the OFC and hippocampus appear to be involved in very similar cognitive processes.

Here, we examine similarities and differences in OFC and hippocampal contributions to cognitive mapping and flexible behaviour. We propose an updated definition of the cognitive map that is faithful to Tolman's conception but that is grounded in contemporary computational models of reinforcement learning. We review evidence that the hippocampus and OFC have separate roles in cognitive mapping and examine in detail a handful of studies that have directly compared hippocampal and OFC processing in similar behavioural paradigms. Finally, we consider how information might be passed between the OFC and hippocampus and how such a cross-structural dialogue might contribute to behaviour that is dependent on cognitive mapping. Our discussion focuses on studies of the rodent dorsal hippocampus and lateral OFC, although anatomical diversity in both regions^{18,23} and

potentially across species^{24,25} is an important consideration. The overarching message that we hope to convey is that the cognitive map perspective offers a productive unifying direction for future studies of hippocampal and orbitofrontal function.

The road to a cognitive map

Hippocampus

The discovery of place cells (FIG. 1a–c) quickly transformed our understanding of the hippocampus²⁶. In a masterful synthesis, O’Keefe and Nadel³ proposed the hippocampus as the seat of a Tolmanian cognitive map (BOX 1). However, the parallels between hippocampal function and Tolman’s ideas run far deeper than the existence of place cells. Beyond the intuitive correspondence between Tolman’s use of the term ‘map’ and the spatial tuning of place cells, other properties of hippocampal representations fulfil Tolman’s cognitive criteria.

For instance, ensembles of hippocampal neurons quickly and obligatorily adopt place-specific firing patterns without explicit reinforcement, and place-cell activity holistically reflects the topology of the environment, as Tolman’s cognitive map was proposed to do^{3,12}. However, just as Tolman’s cognitive map extended beyond the spatial domain, place-cell activity encodes more than an objective, allocentric map of space. The spatial firing patterns of hippocampal neurons are often modulated by non-spatial factors, such as the presence or absence of objects^{27,28}, attention²⁹, conditioned stimuli³⁰, novelty^{31,32}, perceptual features of the environment³³ and an animal’s internal state³⁴. Place cells sometimes cluster their firing fields around — or show ancillary firing fields near — goal locations³⁵ or places that animals receive reward³⁶, indicating that motivational information can be tied to hippocampal spatial representations.

Moreover, the moment-to-moment dynamics of hippocampal activity suggest a mechanism by which animals might prospectively explore their mental models of the world to aid action selection. For example, in rats carrying out a goal-directed navigation task, hippocampal representations recorded just before an animal initiated a journey encoded paths that led to the next location the animal would visit³⁷. In another study, the extent to which pre-trial hippocampal activity was coordinated at an ensemble level predicted future correct choices³⁸. Even more surprisingly, the hippocampus synthesizes representations that include information beyond the animal’s direct experience. For example, in rats that had been trained on an apparatus with visible but inaccessible corridors, hippocampal representations encoded paths through the environment that crossed into corridors that the animals had never physically entered^{39,40}.

Lest it be thought that the hippocampal cognitive map is solely spatial, we would further emphasize that explicitly non-spatial information is also encoded by hippocampal neurons. This is frequently shown in studies of human hippocampal function. For instance, in a task that challenged participants to estimate the value of never-experienced outcomes composed of unusual combinations of familiar foods, the hippocampus repurposed existing representations of the component familiar foods to construct a representation of the composite snacks⁴¹. The hippocampus was also shown to map social relationships when

participants were forced to use this information in order to perform well on a role-playing video game that required interaction with multiple virtual characters⁴². Hippocampal activity has also been shown to reflect the learning of predictive relationships embedded in sequences of visual cues, and the hippocampus is engaged when participants use these implicit stimulus–stimulus associations to guide their decisions^{43–46}.

Such non-spatial representations are not limited to humans. In rats that had been trained to expect sequences of odour cues in a consistent temporal order, hippocampal neurons showed sequence-dependent selectivity for individual odours: that is, they responded to a given odour only when it was a member of one sequence and not when the odour was presented in a different sequence^{47,48}. This sort of contextual modulation is also observed in directional place cells⁴⁹, which fire only when an animal passes through a place field in one direction of travel. Just as place cells encode sequences of positions in a way that reflects the order in which they are traversed, the activity of these odour-selective neurons preserved the temporal structure that was present in the environment.

In all, these data suggest that the hippocampus provides a supremely flexible system for rapidly encoding complex features of the world (both those that are directly experienced and those that are inferred), allowing important information about experience to be encoded in a way that preserves higher-order spatial and relational information.

Orbitofrontal cortex

Although it has long been known that the OFC is important for selecting behaviour that is suited to the current context, agreeing on a more precise specification of OFC function has been difficult^{1,18,50,51}. Early theories of OFC function centred on its role in suppressing actions that are inappropriate to the current context. However, more recent experimental work has argued against response inhibition as the primary function of the OFC^{50,52}, and contemporary theories have arisen from studies of OFC function in associative learning and decision making. Interestingly, OFC activity is not necessary for simple conditioning or even for certain forms of more complex learning⁵². However, tasks that force subjects to adjust their behaviour in light of new learning generally depend on OFC function. For instance, in outcome devaluation studies, an intact OFC is necessary for animals to respond to previously acquired cue–outcome and response–outcome associations in a manner that accords with recent changes in outcome value^{13,53–56}. Similarly, tasks that hinge on subjects' knowledge of the specific sensory features of outcomes rather than their more general hedonic properties require OFC function^{57–59}. Such diverse findings have proved to be difficult for any single theoretical model to explain⁵⁰. These results might reflect a multiplicity of function within OFC circuits, or it might show that the OFC may have a general underlying function that accounts for its involvement in such a range of behaviours.

A potentially unifying account suggests that OFC represents task state²² (FIG. 1d–f). Reinforcement-learning models hinge on the ability of decision-making agents to parse the complexity of the world into discrete, well-defined states^{60–63}. According to these models, by segmenting the world in this way and tracking information about the collection of possible world states (that is, the 'state space'), individuals can assign values to different states, depending on how good or bad those states themselves are and whether they predict

future reward or punishment. A recent paper²², drawing on previous work that modelled how animals learn deeper world structure when solving conditioning paradigms^{62,64,65}, tested reinforcement learning models on a number of behavioural tasks, including classic benchmarks of OFC function, such as reversal learning and devaluation. The results showed that reinforcement-learning models in which the capacity to represent task states was absent or impoverished behaved remarkably like animals with OFC lesions. This suggested a role for the OFC in representing task state, particularly in cases in which states are abstract or not directly observable and must be derived from experience.

Patterns of activity in the OFC are broadly consistent with task state representations. For example, several results suggest that OFC neurons track variables that are crucial for task performance but that are not directly observable^{50,66,67}. In a task in which correct responses depended on the identity of odour cues presented in both the current and previous trials, OFC neurons encoded cue matches and mismatches⁶⁶. Moreover, recent work suggests that OFC ensembles might even encode counterfactual outcomes^{68–70} — that is, consequences that would have occurred had the animal behaved differently. Animals could use this sort of ‘what if’ representation to improve future behavioural performance. In agreement with this, it was found that when animals are given incomplete information about a change in task contingencies, OFC neurons update their outcome expectancies before experiencing the complete set of new action–outcome pairings⁷¹. The OFC also tracks rule or strategy cues that define correct task performance^{72,73}.

Although the OFC is proposed to be especially important for tracking unobservable states, this does not preclude a role in processing sensory-bound state information. Indeed, evidence for OFC encoding of sensory information related to task performance is abundant. For instance, OFC neurons encode odour identity when odour cues predict a subsequent outcome and also prospectively signal features of impending rewards predicted by cues⁶⁷. However, just as hippocampal place cells show spatial tuning on tasks that do not explicitly tax spatial abilities, the OFC may encode observable state information even when such encoding is not necessary for the behaviour. This proposal is consistent with recent experimental work⁷⁴ showing that rats lacking medial OFC function behaved normally when task state was explicitly signalled but showed deficits when task state was unobservable. This work also highlights a challenge in matching OFC neural activity to animals’ state space representations: even well-controlled behavioural tasks are somewhat under-constrained and could be solved in different ways by individual subjects. Going forward, it will be important to construct paradigms that provide a behavioural readout of the state space that subjects are using⁷⁵ or that bias participants towards using particular state spaces⁷⁶. These approaches — combined with ensemble analyses that measure and interpret high-dimensional state space representations⁷⁷ — will be necessary to rigorously test the idea that OFC maps task state space.

These challenges aside, the state space perspective unifies a broad range of data on OFC function. For instance, economic value representations figure prominently in some models of OFC function^{17,51}. Such value coding might emerge from the representation of task state, particularly when state value most directly drives behaviour. Similarly, behaviours in which other factors dominate task structure would be expected to elicit more diverse OFC

representations. For instance, tasks that depend heavily on action–outcome associations drive strong OFC encoding of action information that is absent in other situations^{78,79}.

It should also be noted that other brain regions, particularly other portions of frontal cortex, may also have an important role in cognitive mapping. Parts of the frontal cortex beyond the OFC have been associated with working memory, representation of task rules and top-down control of behaviour¹⁸. Such processes are clearly related to the cognitive map proposal, although they may, in some situations, be dissociable (this issue is explored in more detail in REF. 22). More work is necessary to determine how each of these brain regions fits within the framework described here.

Cognitive maps in action

Anatomy suggests various pathways by which cognitive maps in the hippocampus and OFC might interact to influence behaviour (BOX 2). Studies that have recorded or manipulated activity in these structures under the same behavioural conditions are of particular value in understanding the nature of this interaction. Here, we discuss four parallel data sets that are particularly amenable to direct comparison of OFC and hippocampal function.

Box 2

Hippocampus and orbitofrontal cortex: pathways for interaction

There are at least three (not mutually exclusive) pathways by which the hippocampus and orbitofrontal cortex (OFC) might interact to influence behaviour: direct projections between the two structures, indirect projections and convergence on a common target (see the figure). We highlight several candidate pathways that could foster cross-structural communication.


Direct projections from the hippocampus reach multiple targets in the frontal cortex, including orbitofrontal regions^{26,156–158}. These projections are densest from the ventral hippocampus and become gradually scarcer towards the dorsal hippocampus. Projections that arise from the subiculum (which receives a strong, direct projection from CA1) largely mirror the CA1 projection to the cortex, contacting a similar constellation of frontal cortical targets, including the OFC¹⁵⁹. The OFC does not appear to return a direct projection to the hippocampus. Instead, there are several indirect channels through which orbitofrontal output might reach the hippocampal area. Recent anatomical work has mapped two polysynaptic pathways that link the frontal cortex with dorsal and ventral regions of the hippocampus by thalamic nuclei¹⁶⁰. In fact, both OFC and hippocampus send projections to and receive projections from the thalamic nucleus reuniens, which appears to have a particularly important role in coordinating bidirectional interactions between frontal cortex and hippocampal nuclei^{161–165}. In addition, OFC projections to parahippocampal structures, including the entorhinal, perirhinal and postrhinal cortices, are another path that links OFC and hippocampal processing streams¹⁶⁶.

Additionally, many brain regions receive broadly overlapping projections from the OFC and hippocampus. Several in particular seem striking for their role in learning and decision making. For instance, the ventral tegmental area (VTA) receives indirect input

from both the hippocampus and the OFC. Dopamine-containing neurons in the VTA are thought to have a crucial role in associative learning¹⁶⁷, signalling reward prediction errors that drive learning. Lesions of the OFC alter the encoding of prediction errors in VTA neurons^{21,168}. Disrupting hippocampal outflow to VTA also has behavioural consequences, preventing context-induced reinstatement of drug seeking¹⁶⁹. In addition, electrophysiological experiments have found that hippocampus, VTA and frontal cortical regions are coupled by coherent, low-frequency local field potential oscillations¹⁷⁰.

The ventral striatum (vStr), which is another structure that has long been implicated in reward and motivational processes^{171–173}, also receives robust inputs from the OFC and the hippocampus. One subset of ventral striatal neurons encodes proximity to reinforcement with graded increases in firing rate^{174–176} ('ramp cells'), and the temporal patterning of ramp cell spiking is organized with respect to oscillations in the hippocampus¹⁷⁷, suggesting that hippocampal information processing may 'clock' reward-related representations in the ventral striatum¹³¹. Interestingly, OFC lesions interfere with the coding of reward magnitude in ventral striatal neurons¹⁷⁸, suggesting that OFC and hippocampal inputs converging in the ventral striatum may support complex, multi-attribute representations of predicted outcomes.

PFC, prefrontal cortex.



Goal-directed spatial decision making

A series of studies have compared hippocampal and OFC neural representations in rats that perform a T-maze decision-making task^{69,80–83} (FIG. 2a,b). In this task, rats made decisions by turning left or right at the choice point of the maze. At the beginning of each session, rats relied on trial and error to determine which of three behavioural patterns (always turn left, always turn right or alternate left and right decisions) would be rewarded. Once animals found the correct strategy, a non-signalled switch in reward contingency occurred, and rats

were again forced to identify which strategy would be reinforced. At the beginning of task sessions, when reward contingencies were unknown, rats often paused at the choice point before committing to a left or right turn. Tolman called this behaviour vicarious trial and error (VTE) and suggested that animals, when confronted with difficult decisions, were mentally simulating the consequences of potential actions before deciding which course of action to undertake^{84–86}. In agreement with this, VTE behaviour in the T-maze was found to follow rats' performance: it decreased as they acquired the correct strategy but then re-emerged following the contingency switch later in the session.

Hippocampal ensemble activity was recorded as rats performed this T-maze task, and representations at the choice point were examined⁸³. During VTE, hippocampal place-cell ensembles encoded paths ahead of the animal, along the potential left- and right-turn options, suggesting that the hippocampus simulated each of the available options to aid decision making. Similar 'look ahead' representations in the hippocampus have been observed in rats navigating towards goal locations^{87,88} and in human participants carrying out virtual navigation tasks^{89,90}. In computational terms, place-cell representations of locations removed from the rats' actual positions could be thought of as a neural mechanism for exploring the environmental state space encoded in the hippocampus. However, for these sorts of state space searches to inform decisions in a useful way, information about the potential outcomes associated with environmental states must also be retrieved. One possibility is that the evaluation of trajectories represented by the hippocampus might take place in structures that receive hippocampal inputs, such as the OFC⁸³.

Subsequent work tested this directly by recording ensembles of OFC neurons as rats performed the T-maze task⁶⁹. OFC cells responded to reward receipt on correct trials, which is consistent with previous work. Surprisingly, however, these same reward-responsive OFC neurons were also active when rats paused at the choice point during VTE. Thus, prospective OFC reward-cell activity occurs approximately when an evaluation of hippocampal VTE representations would be expected to arise. Simultaneous recordings of hippocampal and OFC ensembles could clarify the precise temporal relationship between representations in each structure. Nevertheless, these data suggest that the hippocampus aids action planning by searching through the space of previously learned (in this case, spatial) associations, whereas the OFC evaluates candidate actions that are represented by the hippocampus to determine which is best. Note that such an evaluative role for the OFC is not inconsistent with its proposed function of state representation, because expected state value is likely to be an important feature of the trajectories that are proposed by the hippocampus.

The foregoing studies examined hippocampal and OFC activity under behavioural uncertainty and specifically during punctate moments of deliberation. An interesting counterpoint to these data comes from work that examined OFC and hippocampal coding throughout the course of entire trials^{91–95}. These experiments tested rats on a plus-shaped maze. Animals began each trial in one of the north or south start arms and had to travel to either the east or west goal arm, where food was delivered. In the 'place' version of this task, the reward was reliably located in one of the goal arms, and animals needed to approach this arm regardless of which start arm they began from. In the 'response' version of the task, animals learned to make a single response (turn left, for example), regardless of which arm

they began from. The structure of the task allowed neural representations to be probed by several challenges, such as reversals of reward contingency or changes between response- and place-based task variants.

Hippocampal neurons recorded on the task showed spatial responses that were strongly context- or state-dependent^{91,94,96}. For instance, a place field located on the south start arm might only be active as the animal ran towards the west goal arm, showing no firing at all when the animal passed in the same direction through the identical location en route to the east goal arm. Similarly, a place field located on the west goal arm might be active only when the animal began its journey from the north start arm, showing no response when the animal began from the south. Thus, these prospective and retrospective place cells signalled spatial information only in the context of where the animal began from, or where it was going to, respectively. These data show that the hippocampus differentially parses identical instances of behaviour depending on non-spatial factors.

OFC neurons recorded on the same plus-maze task shared common features with hippocampal responses but also showed telling differences⁹². Rather than forming discrete place fields, OFC neurons tended to fire evenly along entire paths between start and goal arms. These neurons did not encode space per se; instead, path-sensitive OFC neurons reflected the probability that a particular path would lead to reward delivery, and the firing patterns of these cells tracked behavioural performance following contingency reversals and switches between tasks. This suggests that OFC representations integrated information about responses and reward expectation. Like the context-sensitive place cells observed in this task, OFC responses discriminated journeys through the same physical space depending on where the animal was travelling to and where the journey began. Interestingly, because OFC firing was spread over entire journeys rather than concentrated within discrete locations, as is the case for hippocampal responses, single OFC cells could reflect prospective information before reaching the choice point and retrospective information on final approach to the goal arm. This perhaps suggests a more integrated, large-scale representation of the task in single OFC neurons and a more granular, distributed encoding scheme by hippocampal ensembles^{92,97,98}.

Intriguingly, analyses of local field potentials (LFPs) that were recorded simultaneously from the OFC and hippocampus as rats carried out the plus-maze task hinted at interactions between these structures. During stable task performance, LFPs in the OFC and hippocampus oscillated coherently at theta (5–12 Hz) frequency. However, coherence fell following reversals of reward contingency or switches in task type only to rise again slowly as rats acquired the new behaviour⁹². Thus, learning resulted in a transient decoupling of activity in the hippocampus and the OFC, whereas consistent performance on the task was accompanied by stable interactions between these structures.

Taken together, the single-unit recordings suggest that both OFC and hippocampal firing patterns are influenced by contextual information related to task performance, such as the start and end locations of journeys across the maze. However, only in OFC ensembles was this contextual modulation dependent on the presence or absence of a food outcome at the end of trajectories. This suggests that the general spatial context modulation observed in the

hippocampus gives way to a more elaborate representation in the OFC that is coloured by biological meaning — that is, factors that are closely related to the animals' current needs or goals. The field potential data provide intriguing evidence that information flow between the hippocampus and OFC is a dynamic process that varies with learning.

One way to conceptualize OFC and hippocampal responses in the plus-maze task is to assume that both structures encoded somewhat overlapping information but with different emphases. Whereas neurons in the hippocampus encoded context information about rats' journeys while preserving single-cell representations of position, neurons in the OFC seemed to average across large swaths of space, placing a premium on outcome encoding^{91,97}. This suggests that cognitive maps in the hippocampus and OFC might contain similar types of information but that they might format this information in different ways. This is consistent with recent work that assesses the hierarchy of spatial, contextual and outcome representations in the OFC and hippocampus^{99,100}. In these studies, rats were trained to choose between objects presented in the corners of a testing chamber to earn food reward. However, the objects were not consistently presented in the same corner, and the association between objects and reward was context-dependent: an object that was rewarded in one chamber was unrewarded when presented in the context of a second testing chamber. In this way, the task dissociated the location at which objects were presented from the outcomes that they predicted and also dissociated objects' reward contingency from their sensory properties, making it possible to test how these task variables are encoded relative to one another. The authors used representational similarity analysis to determine which combinations of task variables evoked the greatest divergence in ensemble representations and which trial types were coded most similarly.

In hippocampal ensembles, spatial information fractionated neural representations: that is, trials that occurred in different contexts were encoded by anti-correlated patterns of activity, and the location of objects within a context drove the next-greatest divergence in ensemble representations¹⁰⁰. Nevertheless, non-spatial information was also encoded: representations of object–reward associations were more strongly segregated than representations of the individual objects' identities. In OFC ensembles, a different hierarchy of representations emerged. Neurons in the OFC distinguished most sharply between rewarded and non-rewarded objects, suggesting that reward contingency is the major dimension along which OFC representations are constructed in this task⁹⁹. The next greatest separation in ensemble representations reflected the location of items within each context. Finally, the absolute position at which trials occurred was most similarly represented by ensemble activity. These data suggest that, unlike hippocampal representations, which strongly reflected where events occurred, OFC activity principally favoured reward contingency over location. Thus, hippocampal and OFC ensembles both encode a variety of task-relevant information, but each is specialized to emphasize different aspects of the resultant cognitive map.

Inferring implicit value

The OFC and hippocampus are both known to be involved in situations in which previous learning is essential — but in itself insufficient — to support behaviour. A particularly salient example of this is sensory preconditioning¹⁰¹ (FIG. 2c). In this task, subjects are first

exposed to an arbitrary pairing of two neutral stimuli (for example, tone A is followed by tone B). Neither stimulus is paired with reward: subjects are simply presented with the cues in a reliable order. Next, in the conditioning phase of the experiment, one of the previously unrewarded cues is paired with a valuable outcome (for example, tone B is followed by a reward). Finally, in the test session, animals are presented with the preconditioned tone A. Across a range of species^{102–110}, subjects' response to the preconditioned cue is found to be similar to that evoked by the directly conditioned cue (that is, tone B). The interpretation is that subjects inferred the complete causal chain that was made implicit in a piecemeal manner over the different phases of the experiment (tone A is followed by tone B, which is followed by a reward). Model-free cue values are insufficient to explain subjects' responding at test, suggesting that model-based inference must be at work.

Parahippocampal structures have long been implicated in supporting inferential behaviour in studies of sensory preconditioning. Early work suggested that hippocampal lesions prevented value inference from the directly conditioned stimulus to its preconditioned partner^{104,111}. Although subsequent work has reported that the hippocampus is not necessary for sensory preconditioning, when cues are presented simultaneously as compounds¹¹², rather than serially as described above¹⁰⁶, other structures within the hippocampal network have also been implicated. For instance, lesions of the perirhinal¹¹³ or retrosplenial¹¹⁴ cortex abolish sensory preconditioning. Although these studies leave open what function is supported by processing through the circuit, recent work using a chemogenetic approach found that specifically silencing the retrosplenial cortex during the preconditioning phase prevented value inference at test without influencing first-order conditioning¹¹⁵. Thus, hippocampal outflow is potentially crucial to the establishment of the associative scaffolding that is later used to infer the value of the preconditioned cue at test. Consistent with this idea, humans tested on a task that was similar to the sensory preconditioning paradigm used in animals¹⁰⁸ showed enhanced preference for the preconditioned cue at test, and the strength of this preference correlated with hippocampal activity during the learning phase of the task.

Notably, these data dovetail with evidence that the OFC is necessary for this inference process¹¹⁶. Pharmacological inactivation of the rat OFC during the test session abolished responding to the preconditioned cue without affecting responding to the cue that underwent first-order conditioning by being paired directly with reward. These data indicate that the OFC is necessary for using the hippocampus-dependent associative scaffolding acquired in the first phase to predict reward at test.

Overall, these data suggest that both regions are crucial for value inference in sensory preconditioning but that the dynamics of hippocampal and OFC involvement differ in important ways. For example, whereas OFC function was found to be important for inference at the test stage¹¹⁶, hippocampal involvement was confined to the conditioning portion of the experiment in both the human¹⁰⁸ and animal work¹¹⁵. These data suggest that the hippocampus is perhaps crucial for encoding a world model that links preconditioned cues with reward, whereas the OFC is more important for accessing this information in the test session to drive responding.

Putting it all together

The evidence reviewed here suggests that the OFC and hippocampus each contribute to cognitive mapping and the resultant behaviour. We have emphasized experiments that demonstrate the striking similarity in OFC and hippocampal function. This is not to say, however, that the hippocampus and OFC operate synergistically in every situation. For instance, hippocampal lesions selectively alter rats' preferences for delayed outcomes^{117–119} (at least in some testing paradigms; also see REF. 120) without affecting decisions between probabilistic outcomes, whereas OFC lesions produce an inverted pattern of impairment¹¹⁷. Similarly, tests of outcome devaluation paint a somewhat divergent picture in which the OFC appears to be involved in behavioural changes that are induced by specific satiety manipulations^{121–123}, whereas the hippocampus is not^{124–129}. Differences in experimental design often make direct comparisons difficult, but these discrepancies point to potential functional differences.

It also seems clear that OFC and hippocampus, although perhaps contributing to similar types of behaviours, show some level of domain specificity in their information processing. This arises both from the unique anatomical organization within each structure and the pattern of inputs they receive from other brain regions. The hippocampus is remarkably adept at linking information into temporally patterned sequences that span large ensembles of neurons^{130,131} and seems particularly concerned with organizing experience along the axes of space and time. This ability to connect elemental representations and to flexibly produce sequences that reflect learned connections makes the hippocampus well suited to encoding, retrieving and exploring mental models of state spaces^{132–134}. By contrast, representations of similar information in the OFC are more rooted in biological importance^{20,135,136}. Although it is clear that the OFC has an important role in cognitive processes such as learning and decision making, the behaviours that depend on the OFC are motivated by biological needs. This includes, for instance, learning how to respond to obtain food or liquid reward or, in the case of devaluation, updating associations to direct behaviour away from food that had previously been linked with illness, as well as behaving appropriately in social situations^{137–139}. These data argue that, whereas the hippocampus is a flexible and promiscuous processor of abstract associations, OFC information processing may be more grounded in items of immediate biological relevance.

One way to view these distinctions is to consider that there are multiple components to the cognitive map. In other words, just as real-world maps often consist of multiple overlays that describe different aspects of the environment, so too must our global cognitive map be constructed of multiple informational layers that can be turned on and off as necessary. One efficient way of doing this is to task different modules in the brain with representing different informational layers of the global map, thereby distributing cognitive maps across multiple neural structures. If cognitive maps were to be distributed across multiple neural structures, coordination between brain regions would be important for supporting adaptive behaviour. What precise form might this cross-structural dialogue take? One intuitive idea is that OFC input to hippocampus contributes to the extra-spatial modulation of place-selective hippocampal neurons, giving rise to spatial representations that are sensitive to reward, goals or motivational states. The OFC might be thought of as imbuing the hippocampal map with

information about expected outcomes to facilitate goal-dependent navigation. Conversely, spatial or relational information conveyed by the hippocampus to the OFC might allow OFC outcome expectancies to become bound with information about spatial positions or more abstract relationships between potential outcomes and ‘paths’ to obtain those outcomes, whether spatial, as in a maze, or non-spatial, as in sensory preconditioning. Organizing OFC representations in this way could facilitate the development of integrated action–outcome associations by tying abstract outcome predictions (such as a cherry- or banana-flavoured sucrose pellet) to locations that are reached by particular responses (for example, go left at a maze choice point to get the banana pellet, or press the right lever in an operant box to get the cherry pellet). Time-sensitive hippocampal representations^{140–142} might also confer temporal specificity to OFC expectancy representations, either directly or through interactions in some downstream areas, such as the ventral striatum.

Although we are in the early stages of understanding the interplay between brain regions, the increasing sophistication of techniques for measuring and manipulating the activity of neural circuits in projection- and cell-type-specific ways draws the issue of cross-structural interactions to the fore. We suggest that these approaches, when applied to interactions between hippocampus and OFC, might be particularly fruitful in improving our understanding of how cognitive-map-dependent behaviour is learned and deployed.

Acknowledgments

The authors thank members of the Schoenbaum laboratory for helpful discussions on the topics addressed here and for feedback on earlier versions of this manuscript. This work was supported by funding from the US National Institute on Drug Abuse at the Intramural Research Program. The opinions expressed in this article are the authors’ own and do not reflect the view of the US National Institutes of Health, the US Department of Health and Human Services or the US government.

Glossary

Economic value

An integrative measure of how good an outcome is to a decision maker that distils the many multidimensional features of that outcome into a unidimensional measure of worth.

Outcome devaluation

The process of rendering a normally appetitive outcome aversive, typically by pairing it with illness.

Place cells

Pyramidal neurons in the hippocampus that fire action potentials when an animal occupies or passes through particular portions of the environment.

Reinforcement-learning models

A collection of machine-learning models that are inspired by psychological learning theory and that are aimed at solving the problem of using experience of the world to guide future behaviour.

Representational similarity analysis

An analysis approach that quantifies the similarity (or dissimilarity) of neural ensemble representations evoked by different conditions.

Response inhibition

The active suppression of actions that are not adaptive in the current setting.

Specific satiety

A means of devaluing a particular outcome by allowing an animal unrestricted access to it before a test session.

Stimulus–stimulus associations

Associations that are formed between neutral stimuli in the environment in the absence of explicit reinforcement.

Vicarious trial and error (VTE)

A pause and orient pattern of behaviour that decision makers often show when deliberating over potential choices. This is thought to be an overt marker of covert, mental processes that simulate potential outcomes of each course of action.

References

1. Rudebeck PH, Murray EA. The orbitofrontal oracle: cortical mechanisms for the prediction and evaluation of specific behavioral outcomes. *Neuron*. 2014; 84:1143–1156. [PubMed: 25521376]
2. Buckner RL. The role of the hippocampus in prediction and imagination. *Annu Rev Psychol*. 2010; 61:27–48. [PubMed: 19958178]
3. O’Keefe, J., Nadel, L. *The Hippocampus as a Cognitive Map*. Clarendon Press; 1978.
4. Tolman EC. Cognitive maps in rats and men. *Psychol Rev*. 1948; 55:189–208. [PubMed: 18870876]
5. Tolman, EC. *Purposive Behavior in Animals and Men*. Appleton-Century-Crofts; 1932.
6. Tolman EC, Brunswik E. The organism and the causal texture of the environment. *Psychol Rev*. 1935; 42:43.
7. Buzsáki G, Moser EI. Memory, navigation and theta rhythm in the hippocampal–entorhinal system. *Nat Neurosci*. 2013; 16:130–138. [PubMed: 23354386]
8. Eichenbaum H, Cohen Neal J. Can we reconcile the declarative memory and spatial navigation views on hippocampal function? *Neuron*. 2014; 83:764–770. [PubMed: 25144874]
9. Eichenbaum H, Dudchenko P, Wood E, Shapiro M, Tanila H. The hippocampus, memory, and place cells: is it spatial memory or a memory space? *Neuron*. 1999; 23:209–226. [PubMed: 10399928]
10. Wikenheiser AM, Redish AD. Decoding the cognitive map: ensemble hippocampal sequences and decision making. *Curr Opin Neurobiol*. 2015; 32:8–15. [PubMed: 25463559]
11. Dudchenko PA, Wood ER. Place fields and the cognitive map. *Hippocampus*. 2015; 25:709–712. [PubMed: 25788229]
12. Redish, AD. *Beyond the Cognitive Map: From Place Cells to Episodic Memory*. MIT Press; 1999.
13. Gallagher M, McMahan R, Schoenbaum G. Orbitofrontal cortex and representation of incentive value in associative learning. *J Neurosci*. 1999; 19:6610–6614. [PubMed: 10414988]
14. Tremblay L, Schultz W. Relative reward preference in primate orbitofrontal cortex. *Nature*. 1999; 398:704–708. [PubMed: 10227292]
15. O’Doherty J, Kringelbach ML, Rolls ET, Hornak J, Andrews C. Abstract reward and punishment representations in the human orbitofrontal cortex. *Nat Neurosci*. 2001; 4:95–102. [PubMed: 11135651]
16. Gottfried JA, O’Doherty J, Dolan RJ. Encoding predictive reward value in human amygdala and orbitofrontal cortex. *Science*. 2003; 301:1104–1107. [PubMed: 12934011]

17. Padoa-Schioppa C, Assad JA. Neurons in the orbitofrontal cortex encode economic value. *Nature*. 2006; 441:223–226. [PubMed: 16633341]
18. Wallis JD. Orbitofrontal cortex and its contribution to decision-making. *Annu Rev Neurosci*. 2007; 30:31–56. [PubMed: 17417936]
19. McDannald MA, Jones JL, Takahashi YK, Schoenbaum G. Learning theory: a driving force in understanding orbitofrontal function. *Neurobiol Learn Mem*. 2014; 108:22–27. [PubMed: 23770491]
20. Thorpe S, Rolls E, Maddison S. The orbitofrontal cortex: neuronal activity in the behaving monkey. *Exp Brain Res*. 1983; 49:93–115. [PubMed: 6861938]
21. Takahashi YK, et al. Expectancy-related changes in firing of dopamine neurons depend on orbitofrontal cortex. *Nat Neurosci*. 2011; 14:1590–1597. [PubMed: 22037501]
22. Wilson RC, Takahashi YK, Schoenbaum G, Niv Y. Orbitofrontal cortex as a cognitive map of task space. *Neuron*. 2014; 81:267–279. [PubMed: 24462094]
23. Fanselow MS, Dong HW. Are the dorsal and ventral hippocampus functionally distinct structures? *Neuron*. 2010; 65:7–19. [PubMed: 20152109]
24. Wallis JD. Cross-species studies of orbitofrontal cortex and value-based decision-making. *Nat Neurosci*. 2012; 15:13–19.
25. Strange BA, Witter MP, Lein ES, Moser EI. Functional organization of the hippocampal longitudinal axis. *Nat Rev Neurosci*. 2014; 15:655–669. [PubMed: 25234264]
26. Andersen, P., Morris, R., Amaral, D., Bliss, T., O'Keefe, J. *The Hippocampus Book*. Oxford Univ. Press; 2006.
27. Komorowski RW, Manns JR, Eichenbaum H. Robust conjunctive item–place coding by hippocampal neurons parallels learning what happens where. *J Neurosci*. 2009; 29:9918–9929. [PubMed: 19657042]
28. Manns JR, Eichenbaum H. A cognitive map for object memory in the hippocampus. *Learn Mem*. 2009; 16:616–624. [PubMed: 19794187]
29. Fenton AA, et al. Attention-like modulation of hippocampus place cell discharge. *J Neurosci*. 2010; 30:4613–4625. [PubMed: 20357112]
30. Moita MA, Rosis S, Zhou Y, LeDoux JE, Blair HT. Hippocampal place cells acquire location-specific responses to the conditioned stimulus during auditory fear conditioning. *Neuron*. 2003; 37:485–497. [PubMed: 12575955]
31. Larkin MC, Lykken C, Tye LD, Wickelgren JG, Frank LM. Hippocampal output area CA1 broadcasts a generalized novelty signal during an object–place recognition task. *Hippocampus*. 2014; 24:773–783. [PubMed: 24596296]
32. Lever C, et al. Environmental novelty elicits a later theta phase of firing in CA1 but not subiculum. *Hippocampus*. 2010; 20:229–234. [PubMed: 19623610]
33. Quirk GJ, Muller RU, Kubie JL. The firing of hippocampal place cells in the dark depends on the rat's recent experience. *J Neurosci*. 1990; 10:2008–2017. [PubMed: 2355262]
34. Kennedy PJ, Shapiro ML. Motivational states activate distinct hippocampal representations to guide goal-directed behaviors. *Proc Natl Acad Sci USA*. 2009; 106:10805–10810. [PubMed: 19528659]
35. Hollup SA, Molden S, Donnett JG, Moser MB, Moser EI. Accumulation of hippocampal place fields at the goal location in an annular watermaze task. *J Neurosci*. 2001; 21:1635–1644. [PubMed: 11222654]
36. Hok V, et al. Goal-related activity in hippocampal place cells. *J Neurosci*. 2007; 27:472–482. [PubMed: 17234580]
37. Pfeiffer BE, Foster DJ. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*. 2013; 497:74–79. [PubMed: 23594744]
38. Singer AC, Carr MF, Karlsson MP, Frank LM. Hippocampal SWR activity predicts correct decisions during the initial learning of an alternation task. *Neuron*. 2013; 77:1163–1173. [PubMed: 23522050]
39. Dragoi G, Tonegawa S. Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature*. 2011; 469:397–401. [PubMed: 21179088]

40. Ólafsdóttir HF, Barry C, Saleem AB, Hassabis D, Spiers HJ. Hippocampal place cells construct reward related sequences through unexplored space. *eLife*. 2015; 4:e06063. [PubMed: 26112828]
41. Barron HC, Dolan RJ, Behrens TEJ. Online evaluation of novel choices by simultaneous representation of multiple memories. *Nat Neurosci*. 2013; 16:1492–1498. [PubMed: 24013592]
42. Tavares RM, et al. A map for social navigation in the human brain. *Neuron*. 2015; 87:231–243. [PubMed: 26139376]
43. Bornstein AM, Daw ND. Cortical and hippocampal correlates of deliberation during model-based decisions for rewards in humans. *PLoS Comput Biol*. 2013; 9:e1003387. [PubMed: 24339770]
44. Bornstein AM, Daw ND. Dissociating hippocampal and striatal contributions to sequential prediction learning. *Eur J Neurosci*. 2012; 35:1011–1023. [PubMed: 22487032]
45. Schapiro AC, Turk-Browne NB, Norman KA, Botvinick MM. Statistical learning of temporal community structure in the hippocampus. *Hippocampus*. 2015; 26:3–8. [PubMed: 26332666]
46. Shohamy D, Turk-Browne NB. Mechanisms for widespread hippocampal involvement in cognition. *J Exp Psychol Gen*. 2013; 142:1159. [PubMed: 24246058]
47. Ginther MR, Walsh DF, Ramus SJ. Hippocampal neurons encode different episodes in an overlapping sequence of odors task. *J Neurosci*. 2011; 31:2706–2711. [PubMed: 21325539]
48. Allen TA, Salz DM, McKenzie S, Fortin NJ. Nonspatial sequence coding in CA1 neurons. *J Neurosci*. 2016; 36:1547–1563. [PubMed: 26843637]
49. McNaughton BL, Barnes CA, O'Keefe J. The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Exp Brain Res*. 1983; 52:41–49. [PubMed: 6628596]
50. Stalnaker TA, Cooch NK, Schoenbaum G. What the orbitofrontal cortex does not do. *Nat Neurosci*. 2015; 18:620–627. [PubMed: 25919962]
51. Padoa-Schioppa C. Neurobiology of economic choice: a good-based model. *Annu Rev Neurosci*. 2011; 34:333. [PubMed: 21456961]
52. Schoenbaum G, Roesch MR, Stalnaker TA, Takahashi YK. A new perspective on the role of the orbitofrontal cortex in adaptive behaviour. *Nat Rev Neurosci*. 2009; 10:885–892. [PubMed: 19904278]
53. West EA, Forcelli PA, McCue DL, Malkova L. Differential effects of serotonin-specific and excitotoxic lesions of OFC on conditioned reinforcer devaluation and extinction in rats. *Behav Brain Res*. 2013; 246:10–14. [PubMed: 23458741]
54. Rhodes SE, Murray EA. Differential effects of amygdala, orbital prefrontal cortex, and prefrontal cortex lesions on goal-directed behavior in rhesus macaques. *J Neurosci*. 2013; 33:3380–3389. [PubMed: 23426666]
55. Izquierdo AD, Suda RK, Murray EA. Bilateral orbital prefrontal cortex lesions in rhesus monkeys disrupt choices guided by both reward value and reward contingency. *J Neurosci*. 2004; 24:7540–7548. [PubMed: 15329401]
56. Gremel CM, Costa RM. Orbitofrontal and striatal circuits dynamically encode the shift between goal-directed and habitual actions. *Nat Commun*. 2013; 4:2264. [PubMed: 23921250]
57. Ostlund SB, Balleine BW. Orbitofrontal cortex mediates outcome encoding in Pavlovian but not instrumental learning. *J Neurosci*. 2007; 27:4819–4825. [PubMed: 17475789]
58. McDannald MA, Saddoris MP, Gallagher M, Holland PC. Lesions of orbitofrontal cortex impair rats' differential outcome expectancy learning but not conditioned stimulus-potentiated feeding. *J Neurosci*. 2005; 25:4626–4632. [PubMed: 15872110]
59. McDannald MA, Lucantonio F, Burke KA, Niv Y, Schoenbaum G. Ventral striatum and orbitofrontal cortex are both required for model-based, but not model-free, reinforcement learning. *J Neurosci*. 2011; 31:2700–2705. [PubMed: 21325538]
60. Redish AD, Jensen S, Johnson A, Kurth-Nelson Z. Reconciling reinforcement learning models with behavioral extinction and renewal: Implications for addiction, relapse, and problem gambling. *Psychol Rev*. 2007; 114:784–805. [PubMed: 17638506]
61. Sutton, RS., Barto, AG. *Reinforcement Learning: An Introduction*. MIT Press; 1998.
62. Gershman SJ, Niv Y. Learning latent structure: carving nature at its joints. *Curr Opin Neurobiol*. 2010; 20:251–256. [PubMed: 20227271]

63. O'Doherty JP, Lee SW, McNamee D. The structure of reinforcement-learning mechanisms in the human brain. *Curr Opin Behav Sci.* 2015; 1:94–100.
64. Gershman SJ, Blei D, Niv Y. Context, learning and extinction. *Psychol Rev.* 2010; 117:197–209. [PubMed: 20063968]
65. Courville AC, Daw ND, Touretzky DS. Similarity and discrimination in classical conditioning: a latent variable account. *Adv Neural Inform Process Syst.* 2005; 17:313–320.
66. Ramus SJ, Eichenbaum H. Neural correlates of olfactory recognition memory in the rat orbitofrontal cortex. *J Neurosci.* 2000; 20:8199–8208. [PubMed: 11050143]
67. Schoenbaum G, Eichenbaum H. Information coding in the rodent prefrontal cortex. I. Single-neuron activity in orbitofrontal cortex compared with that in piriform cortex. *J Neurophysiol.* 1995; 74:733–750. [PubMed: 7472378]
68. Steiner AP, Redish AD. Behavioral and neurophysiological correlates of regret in rat decision-making on a neuroeconomic task. *Nat Neurosci.* 2014; 17:995–1002. [PubMed: 24908102]
69. Steiner AP, Redish AD. The road not taken: neural correlates of decision making in orbitofrontal cortex. *Front Neurosci.* 2012; 6:131. [PubMed: 22973189]
70. Abe H, Lee D. Distributed coding of actual and hypothetical outcomes in the orbital and dorsolateral prefrontal cortex. *Neuron.* 2011; 70:731–741. [PubMed: 21609828]
71. Stalnaker TA, et al. Orbitofrontal neurons infer the value and identity of predicted outcomes. *Nat Commun.* 2014; 5:3926. [PubMed: 24894805]
72. Tsujimoto S, Genovesio A, Wise SP. Monkey orbitofrontal cortex encodes response choices near feedback time. *J Neurosci.* 2009; 29:2569–2574. [PubMed: 19244532]
73. Tsujimoto S, Sawaguchi T. Neuronal activity representing temporal prediction of reward in the primate prefrontal cortex. *J Neurophysiol.* 2005; 93:3687–3692. [PubMed: 15634707]
74. Bradfield LA, Dezfouli A, van Holstein M, Chieng B, Balleine BW. Medial orbitofrontal cortex mediates outcome retrieval in partially observable task situations. *Neuron.* 2015; 88:1268–1280. [PubMed: 26627312]
75. Gläscher J, Daw N, Dayan P, O'Doherty JP. States versus rewards: dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron.* 2010; 66:585–595. [PubMed: 20510862]
76. Riceberg JS, Shapiro ML. Reward stability determines the contribution of orbitofrontal cortex to adaptive behavior. *J Neurosci.* 2012; 32:16402–16409. [PubMed: 23152622]
77. Doll BB, Duncan KD, Simon DA, Shohamy D, Daw ND. Model-based choices involve prospective neural activity. *Nat Neurosci.* 2015; 18:767–772. [PubMed: 25799041]
78. Feierstein CE, Quirk MC, Uchida N, Sosulski DL, Mainen ZF. Representation of spatial goals in rat orbitofrontal cortex. *Neuron.* 2006; 60:495–507.
79. Roesch MR, Taylor AR, Schoenbaum G. Encoding of time-discounted rewards in orbitofrontal cortex is independent of value representation. *Neuron.* 2006; 51:509–520. [PubMed: 16908415]
80. Gupta AS, van der Meer MA, Touretzky DS, Redish AD. Segmentation of spatial experience by hippocampal theta sequences. *Nat Neurosci.* 2012; 15:1032–1039. [PubMed: 22706269]
81. Blumenthal A, Steiner A, Seeland KD, Redish AD. Effects of pharmacological manipulations of NMDA-receptors on deliberation in the multiple-T task. *Neurobiol Learn Mem.* 2011; 95:376–384. [PubMed: 21296174]
82. Gupta AS, van der Meer MA, Touretzky DS, Redish AD. Hippocampal replay is not a simple function of experience. *Neuron.* 2010; 65:695–705. [PubMed: 20223204]
83. Johnson A, Redish AD. Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J Neurosci.* 2007; 27:12176–12189. [PubMed: 17989284]
84. Tolman EC. Prediction of vicarious trial and error by means of the schematic sowbug. *Psychol Rev.* 1939; 46:318–336.
85. Muenzinger KF. Vicarious trial and error at a point of choice: I. A general survey of its relation to learning efficiency. *Pedagog Semin J Genet Psychol.* 1938; 53:75–86.
86. Redish AD. Vicarious trial and error. *Nat Rev Neurosci.* 2016; 17:147–159. [PubMed: 26891625]
87. Wikenheiser AM, Redish AD. Hippocampal theta sequences reflect current goals. *Nat Neurosci.* 2015; 18:289–294. [PubMed: 25559082]

88. Bieri KW, Bobbitt KN, Colgin LL. Slow and fast gamma rhythms coordinate different spatial coding modes in hippocampal place cells. *Neuron*. 2014; 82:670–681. [PubMed: 24746420]
89. Simon DA, Daw ND. Neural correlates of forward planning in a spatial decision task in humans. *J Neurosci*. 2011; 31:5526–5539. [PubMed: 21471389]
90. Chadwick MJ, Jolly AE, Amos DP, Hassabis D, Spiers HJ. A goal direction signal in the human entorhinal/subicular region. *Curr Biol*. 2015; 25:87–92. [PubMed: 25532898]
91. Ferbinteanu J, Shirvalkar P, Shapiro ML. Memory modulates journey-dependent coding in the rat hippocampus. *J Neurosci*. 2011; 31:9135–9146. [PubMed: 21697365]
92. Young JJ, Shapiro ML. Dynamic coding of goal-directed paths by orbital prefrontal cortex. *J Neurosci*. 2011; 31:5989–6000. [PubMed: 21508224]
93. Rich EL, Shapiro M. Rat prefrontal cortical neurons selectively code strategy switches. *J Neurosci*. 2009; 29:7208–7219. [PubMed: 19494143]
94. Ferbinteanu J, Shapiro ML. Prospective and retrospective memory coding in the hippocampus. *Neuron*. 2003; 40:1227–1239. [PubMed: 14687555]
95. Bahar AS, Shapiro ML. Remembering to learn: independent place and journey coding mechanisms contribute to memory transfer. *J Neurosci*. 2012; 32:2191–2203. [PubMed: 22323731]
96. Shapiro ML, Ferbinteanu J. Relative spike timing in pairs of hippocampal neurons distinguishes the beginning and end of journeys. *Proc Natl Acad Sci USA*. 2006; 103:4287–4292. [PubMed: 16537523]
97. Young JJ, Shapiro ML. The orbitofrontal cortex and response selection. *Ann NY Acad Sci*. 2011; 1239:25–32. [PubMed: 22145872]
98. Shapiro, ML., Riceberg, JS., Seip-Cammack, K., Guise, KG. Space, Time and Memory in the Hippocampal Formation. Derdikman, D., Knierim, JJ., editors. Springer; 2014. p. 517-560.
99. Farovik A, et al. Orbitofrontal cortex encodes memories within value-based schemas and represents contexts that guide memory retrieval. *J Neurosci*. 2015; 35:8333–8344. [PubMed: 26019346]
100. McKenzie S, et al. Hippocampal representation of related and opposing memories develop within distinct, hierarchically organized neural schemas. *Neuron*. 2014; 83:202–215. [PubMed: 24910078]
101. Brogden WJ. Sensory pre-conditioning. *J Exp Psychol*. 1939; 25:323–332.
102. Matsumoto Y, Hirashima D, Mizunami M. Analysis and modeling of neural processes underlying sensory preconditioning. *Neurobiol Learn Mem*. 2013; 101:103–113. [PubMed: 23380289]
103. Muller D, Gerber B, Hellstern F, Hammer M, Menzel R. Sensory preconditioning in honeybees. *J Exp Biol*. 2000; 203:1351–1364. [PubMed: 10729283]
104. Port RL, Beggs AL, Patterson MM. Hippocampal substrate of sensory associations. *Physiol Behav*. 1987; 39:643–647. [PubMed: 3588713]
105. Hall D, Suboski MD. Sensory preconditioning and second-order conditioning of alarm reactions in zebra danio fish (*Brachydanio rerio*). *J Comp Psychol*. 1995; 109:76.
106. Ward-Robinson J, et al. Excitotoxic lesions of the hippocampus leave sensory preconditioning intact: implications for models of hippocampal functioning. *Behav Neurosci*. 2001; 115:1357–1362. [PubMed: 11770066]
107. Yu T, Lang S, Birbaumer N, Kotchoubey B. Neural correlates of sensory preconditioning: a preliminary fMRI investigation. *Hum Brain Mapp*. 2014; 35:1297–1304. [PubMed: 23450811]
108. Wimmer GE, Shohamy D. Preference by association: how memory mechanisms in the hippocampus bias decisions. *Science*. 2012; 338:270–273. [PubMed: 23066083]
109. Karn HW. Sensory pre-conditioning and incidental learning in human subjects. *J Exp Psychol*. 1947; 37:540. [PubMed: 18920627]
110. Kojima S, et al. Sensory preconditioning for feeding response in the pond snail, *Lymnaea stagnalis*. *Brain Res*. 1998; 808:113–115. [PubMed: 9795175]
111. Port RL, Patterson MM. Fimbrial lesions and sensory preconditioning. *Behav Neurosci*. 1984; 98:584. [PubMed: 6466439]
112. Rescorla RA, Cunningham CL. Within-compound flavor associations. *J Exp Psychol Anim Behav Process*. 1978; 4:267–275. [PubMed: 690565]

113. Nicholson DA, Freeman JH Jr. Lesions of the perirhinal cortex impair sensory preconditioning in rats. *Behav Brain Res.* 2000; 112:69–75. [PubMed: 10862937]
114. Robinson S, Keene CS, Iaccarino HF, Duan D, Bucci DJ. Involvement of retrosplenial cortex in forming associations between multiple sensory stimuli. *Behav Neurosci.* 2011; 125:578. [PubMed: 21688884]
115. Robinson S, et al. Chemogenetic silencing of neurons in retrosplenial cortex disrupts sensory preconditioning. *J Neurosci.* 2014; 34:10982–10988. [PubMed: 25122898]
116. Jones JL, et al. Orbitofrontal cortex supports behavior and learning using inferred but not cached values. *Science.* 2012; 338:953–956. [PubMed: 23162000]
117. Abela AR, Chudasama Y. Dissociable contributions of the ventral hippocampus and orbitofrontal cortex to decision-making with a delayed or uncertain outcome. *Eur J Neurosci.* 2013; 37:640–647. [PubMed: 23190048]
118. Mariano TY, et al. Impulsive choice in hippocampal but not orbitofrontal cortex-lesioned rats on a nonspatial decision-making maze task. *Eur J Neurosci.* 2009; 30:472–484. [PubMed: 19656177]
119. Cheung THC, Cardinal RN. Hippocampal lesions facilitate instrumental learning with delayed reinforcement but induce impulsive choice in rats. *BMC Neurosci.* 2005; 6:36. [PubMed: 15892889]
120. Bett D, Murdoch LH, Wood ER, Dudchenko PA. Hippocampus, delay discounting, and vicarious trial-and-error. *Hippocampus.* 2015; 25:643–654. [PubMed: 25483408]
121. O’Doherty J, et al. Sensory-specific satiety-related olfactory activation of the human orbitofrontal cortex. *Neuroreport.* 2000; 11:399–403. [PubMed: 10674494]
122. West EA, DesJardin JT, Gale K, Malkova L. Transient inactivation of orbitofrontal cortex blocks reinforcer devaluation in macaques. *J Neurosci.* 2011; 31:15128–15135. [PubMed: 22016546]
123. Schoenbaum G, Roesch M. Orbitofrontal cortex, associative learning, and expectancies. *Neuron.* 2005; 47:633–636. [PubMed: 16129393]
124. Higgs S, Williamson AC, Rotshtein P, Humphreys GW. Sensory-specific satiety is intact in amnesics who eat multiple meals. *Psychol Sci.* 2008; 19:623–628. [PubMed: 18727773]
125. Corbit LH, Ostlund SB, Balleine BW. Sensitivity to instrumental contingency degradation is mediated by the entorhinal cortex and its efferents via the dorsal hippocampus. *J Neurosci.* 2002; 22:10976–10984. [PubMed: 12486193]
126. Corbit LH, Balleine BW. The role of the hippocampus in instrumental conditioning. *J Neurosci.* 2000; 20:4233–4239. [PubMed: 10818159]
127. Chudasama Y, Wright KS, Murray EA. Hippocampal lesions in rhesus monkeys disrupt emotional responses but not reinforcer devaluation effects. *Biol Psychiatry.* 2008; 63:1084–1091. [PubMed: 18191111]
128. Machado CJ, Bachevalier J. The impact of selective amygdala, orbital frontal cortex, or hippocampal formation lesions on established social relationships in rhesus monkeys (*Macaca mulatta*). *Behav Neurosci.* 2006; 120:761. [PubMed: 16893284]
129. Reichelt AC, Lin TE, Harrison JJ, Honey RC, Good MA. Differential role of the hippocampus in response-outcome and context-outcome learning: evidence from selective satiation procedures. *Neurobiol Learn Mem.* 2011; 96:248–253. [PubMed: 21596147]
130. Carr MF, Jadhav SP, Frank LM. Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nat Neurosci.* 2011; 14:147–153. [PubMed: 21270783]
131. Malhotra S, Cross RW, van der Meer MA. Theta phase precession beyond the hippocampus. *Rev Neurosci.* 2012; 23:39–65. [PubMed: 22718612]
132. van der Meer M, Kurth-Nelson Z, Redish AD. Information processing in decision-making systems. *Neuroscientist.* 2012; 18:342–359. [PubMed: 22492194]
133. Johnson A, van der Meer MA, Redish AD. Integrating hippocampus and striatum in decision-making. *Curr Opin Neurobiol.* 2007; 17:692–697. [PubMed: 18313289]
134. Kurth-Nelson Z, Bickel W, Redish AD. A theoretical account of cognitive effects in delay discounting. *Eur J Neurosci.* 2012; 35:1052–1064. [PubMed: 22487035]

135. Rolls E, Sienkiewicz Z, Yaxley S. Hunger modulates the responses to gustatory stimuli of single neurons in the caudolateral orbitofrontal cortex of the macaque monkey. *Eur J Neurosci.* 1989; 1:53–60. [PubMed: 12106174]
136. Rolls ET. The functions of the orbitofrontal cortex. *Brain Cogn.* 2004; 55:11–29. [PubMed: 15134840]
137. Pearson JM, Watson KK, Platt ML. Decision making: the neuroethological turn. *Neuron.* 2014; 82:950–965. [PubMed: 24908481]
138. Watson KK, Platt ML. Social signals in primate orbitofrontal cortex. *Curr Biol.* 2012; 22:2268–2273. [PubMed: 23122847]
139. Ross R, LoPresti M, Schon K, Stern C. Role of the hippocampus and orbitofrontal cortex during the disambiguation of social cues in working memory. *Cogn Affect Behav Neurosci.* 2013; 13:900–915. [PubMed: 23640112]
140. Eichenbaum H. Memory on time. *Trends Cogn Sci.* 2013; 17:81–88. [PubMed: 23318095]
141. Eichenbaum H. Time cells in the hippocampus: a new dimension for mapping memories. *Nat Rev Neurosci.* 2014; 15:732–744. [PubMed: 25269553]
142. MacDonald CJ, Lepage KQ, Eden UT, Eichenbaum H. Hippocampal “time cells” bridge the gap in memory for discontinuous events. *Neuron.* 2011; 71:737–749. [PubMed: 21867888]
143. Johnson A, Crowe D. Revisiting Tolman: theories and cognitive maps. *Cognitive Critique.* 2009; 1:43–72.
144. Tolman EC. There is more than one kind of learning. *Psychol Rev.* 1949; 56:144. [PubMed: 18128182]
145. Tolman EC. The determiners of behavior at a choice point. *Psychol Rev.* 1938; 45:1–41.
146. Tolman EC, Ritchie BF, Kalish D. Studies in spatial learning. I. Orientation and the short-cut. *J Exp Psychol.* 1946; 36:13–24. [PubMed: 21015338]
147. Tolman EC, Ritchie BF, Kalish D. Studies in spatial learning. II. Place learning versus response learning. *J Exp Psychol.* 1946; 36:221–229. [PubMed: 20985357]
148. Schiller D, et al. Memory and space: towards an understanding of the cognitive map. *J Neurosci.* 2015; 35:13904–13911. [PubMed: 26468191]
149. Rescorla, RA., Wagner, AR. *Classical Conditioning II: Current Research and Theory.* Black, AH., Prokasy, WF., editors. Appleton-Century-Crofts; 1972. p. 64-99.
150. Pearce JM, Hall G. A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychol Rev.* 1980; 87:532–552. [PubMed: 7443916]
151. Mackintosh NJ. A theory of attention: variations in the associability of stimuli with reinforcement. *Psychol Rev.* 1975; 82:276–298.
152. Doll BB, Shohamy D, Daw ND. Multiple memory systems as substrates for multiple decision systems. *Neurobiol Learn Mem.* 2015; 117:4–13. [PubMed: 24846190]
153. Dayan P, Daw ND. Decision theory, reinforcement learning, and the brain. *Cogn Affect Behav Neurosci.* 2008; 8:429–453. [PubMed: 19033240]
154. Balleine, BW., Daw, ND., O’Doherty, JP. *Neuroeconomics: Decision Making and the Brain.* Glimcher, PW., Fehr, E., editors. Academic Press; 2008. p. 367-385.
155. Daw ND, Niv Y, Dayan P. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat Neurosci.* 2005; 8:1704–1711. [PubMed: 16286932]
156. Jay TM, Witter MP. 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.* 1991; 313:574–586. [PubMed: 1783682]
157. Barbas H, Blatt GJ. Topographically specific hippocampal projections target functionally distinct prefrontal areas in the rhesus monkey. *Hippocampus.* 1995; 5:511–533. [PubMed: 8646279]
158. Insausti R, Munoz M. Cortical projections of the non-entorhinal hippocampal formation in the cynomolgus monkey (*Macaca fascicularis*). *Eur J Neurosci.* 2001; 14:435–451. [PubMed: 11553294]
159. Aggleton JP, Christiansen K. The subiculum: the heart of the extended hippocampal system. *Prog Brain Res.* 2015; 219:65–82. [PubMed: 26072234]

160. Prasad JA, Chudasama Y. Viral tracing identifies parallel disynaptic pathways to the hippocampus. *J Neurosci*. 2013; 33:8494–8503. [PubMed: 23658186]
161. McKenna JT, Vertes RP. Afferent projections to nucleus reuniens of the thalamus. *J Comp Neurol*. 2004; 480:115–142. [PubMed: 15514932]
162. Vertes RP, Hoover WB, Szigeti-Buck K, Leranath C. Nucleus reuniens of the midline thalamus: link between the medial prefrontal cortex and the hippocampus. *Brain Res Bull*. 2007; 71:601–609. [PubMed: 17292803]
163. Griffin AL. Role of the thalamic nucleus reuniens in mediating interactions between the hippocampus and medial prefrontal cortex during spatial working memory. *Front Syst Neurosci*. 2015; 9:29. [PubMed: 25805977]
164. Varela C, Kumar S, Yang JY, Wilson MA. Anatomical substrates for direct interactions between hippocampus, medial prefrontal cortex, and the thalamic nucleus reuniens. *Brain Struct Funct*. 2014; 219:911–929. [PubMed: 23571778]
165. Mitchell AS, et al. Advances in understanding mechanisms of thalamic relays in cognition and behavior. *J Neurosci*. 2014; 34:15340–15346. [PubMed: 25392501]
166. Witter MP, et al. Cortico-hippocampal communication by way of parallel parahippocampal-subicular pathways. *Hippocampus*. 2000; 10:398–410. [PubMed: 10985279]
167. Schultz W, Dayan P, Montague R. A neural substrate of prediction and reward. *Science*. 1997; 275:1593–1599. [PubMed: 9054347]
168. Jo, YS., Mizumori, SJY. Prefrontal regulation of neuronal activity in the ventral tegmental area. *Cereb Cortex*. 2015. <http://dx.doi.org/10.1093/cercor/bhv215>
169. Luo AH, Tahsili-Fahadan P, Wise RA, Lupica CR, Aston-Jones G. Linking context with reward: a functional circuit from hippocampal CA3 to ventral tegmental area. *Science*. 2011; 333:353–357. [PubMed: 21764750]
170. Fujisawa S, Buzsáki GA. 4 Hz oscillation adaptively synchronizes prefrontal, VTA, and hippocampal activities. *Neuron*. 2011; 72:153–165. [PubMed: 21982376]
171. Groenewegen, HJ., Wright, CL., Beijer, AVJ. The Emotional Motor System. Holstege, G. Bandler, R., Saper, CB., editors. Elsevier; 1996. p. 485-512.
172. Mogenson GJ, Jones DL, Yim CY. From motivation to action: Functional interface between the limbic system and the motor system. *Prog Neurobiol*. 1980; 14:69–97. [PubMed: 6999537]
173. van der Meer MA, Redish AD. Expectancies in decision making, reinforcement learning, and ventral striatum. *Front Neurosci*. 2010; 4:6. [PubMed: 21221409]
174. Lavoie AM, Mizumori SJY. Spatial-, movement- and reward-sensitive discharge by medial ventral striatum neurons in rats. *Brain Res*. 1994; 638:157–168. [PubMed: 8199856]
175. van der Meer MA, Redish AD. Covert expectation-of-reward in rat ventral striatum at decision points. *Front Integr Neurosci*. 2009; 3:1–15. [PubMed: 19225578]
176. Miyazaki K, Miwazaki KW, Matsumoto G. Different representation of forthcoming reward in nucleus accumbens and medial prefrontal cortex. *Neuroreport*. 2003; 15:721–726.
177. van der Meer MAA, Redish AD. Theta phase precession in rat ventral striatum links place and reward information. *J Neurosci*. 2011; 31:2843–2854. [PubMed: 21414906]
178. Cooch NK, et al. Orbitofrontal lesions eliminate signalling of biological significance in cue-responsive ventral striatal neurons. *Nat Commun*. 2015; 6:7195. [PubMed: 26006060]

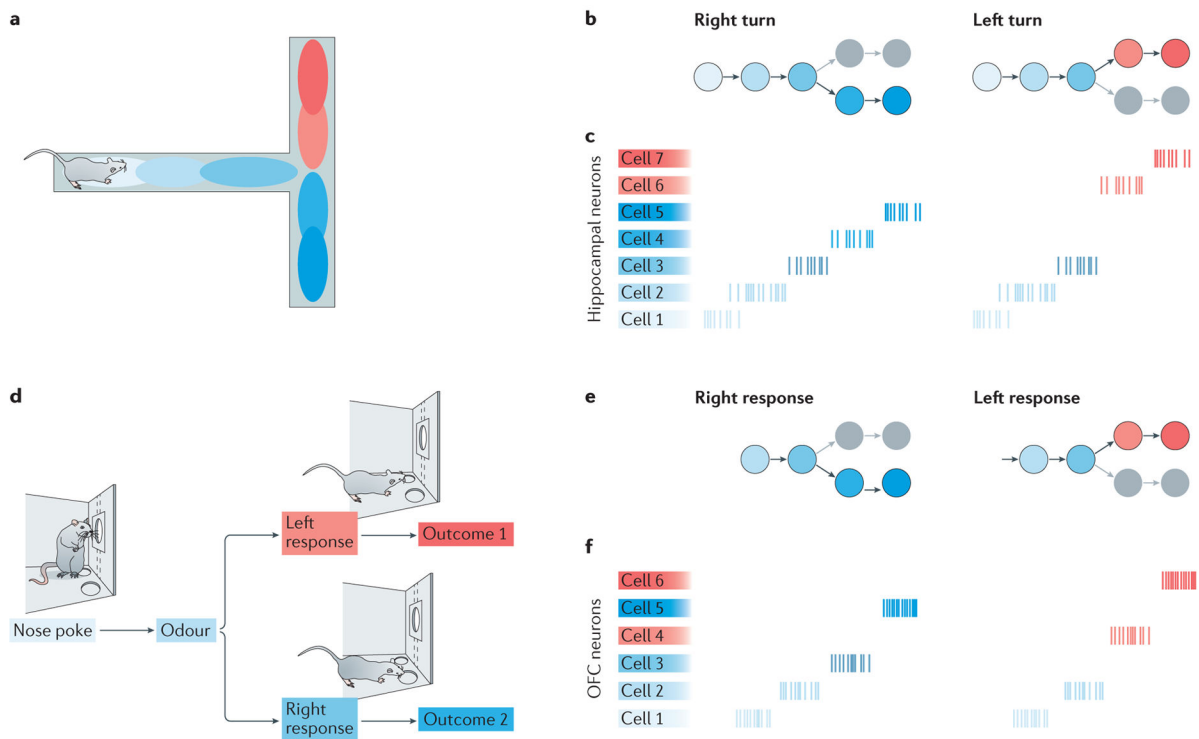


Figure 1. Hippocampal and orbitofrontal cognitive mapping

a | In spatial tasks, many hippocampal neurons exhibit spatially specific firing. The firing fields (place fields, represented in the image by coloured ellipsoids) of these ‘place cells’ tile the environment. At an ensemble level, the firing of place cells encodes the animal’s position in the environment. **b** | In a reinforcement-learning framework, the ensemble firing of spatially tuned place cells could be thought of as encoding an environmental state space: that is, it would represent both individual states (the circumscribed portions of the environment within which each individual place cell is most active) and how they connect to one another. States in this example environmental state space are coloured to correspond to the place fields shown in part **a**. **c** | As animals traverse the environment, the activity of hippocampal neurons could be thought of as representing trajectories through the environmental state space. The figure shows raster plots that illustrate the firing of seven individual hippocampal cells representing two different state space trajectories (coloured to match parts **a**. and **b**.). The trajectories overlap as the animal travels along the central arm of a T-maze but diverge for left and right turns. Different sets of place cells represent positions to the left and right of the choice point. In a similar manner, the activity of non-spatially tuned hippocampal neurons could represent position in a more abstract, non-spatial state space. **d** | Similarly, as animals are engaged in decision-making tasks, orbitofrontal cortex (OFC) neurons that are active during the performance of actions or other task events (such as the presentation of cues and outcomes) could encode the current task state. **e** | In contrast to the example state space shown in part **b**, which was defined entirely by position in the environment, the state space for an operant decision-making task might be structured around important task events, include ing actions (such as making an initial nose poke in an odour sample port). cues (such as the presentation of an odour) and outcomes (such as the delivery of a liquid reward). In

these examples, the states represented might be the ‘odour-sampling state’ or the ‘reward-delivery state’. f | Because different OFC neurons are activated by specific actions and events (for example, they fire in response to particular odour cues rather than to general odour presentation), the trajectories through state space encoded by OFC ensembles vary depending on the animal’s actions (in this case, a left or right response) and on the information received from the environment.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

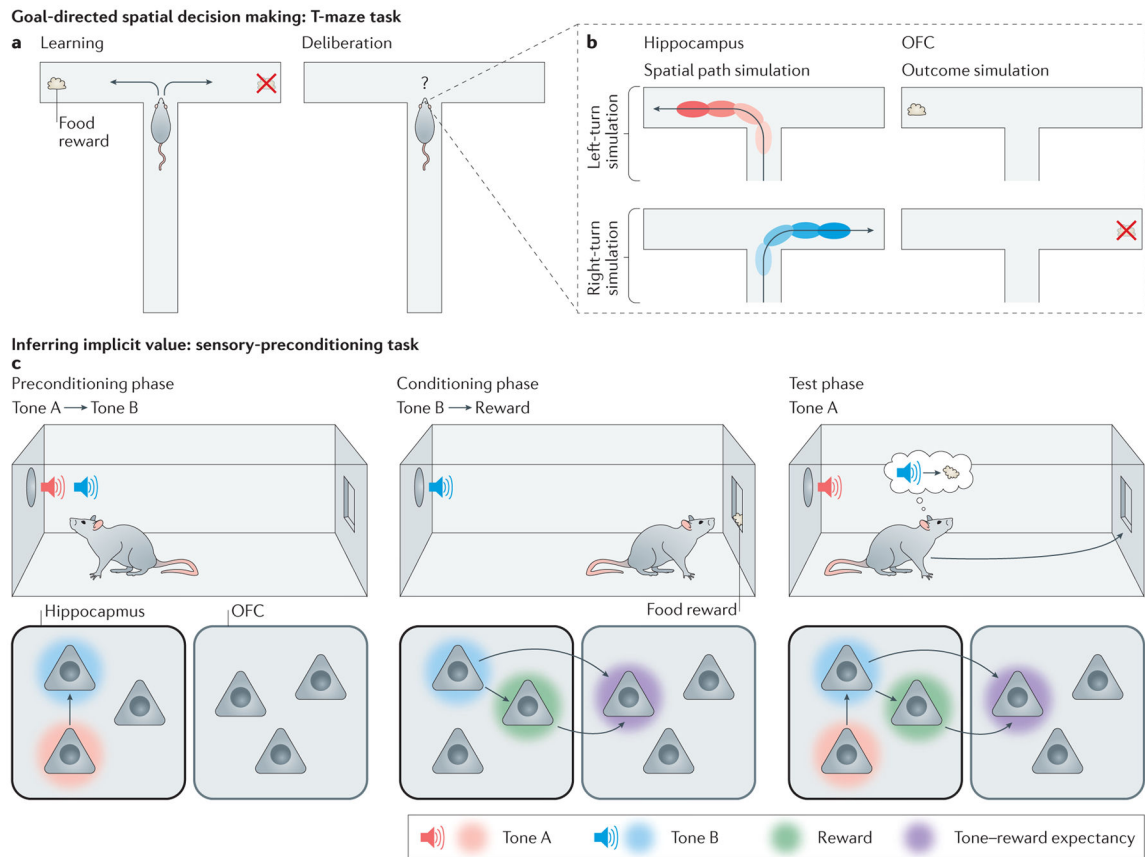


Figure 2. Cognitive maps in action

a | In the T-maze decision-making task, animals learn by trial and error that a particular pattern of choices (for example, always turn left at the choice point) will be rewarded. **b** | When animals deliberate over their options at the choice point, hippocampal ensembles simulate spatial trajectories towards potential reward sites. At the same time, reward-sensitive neurons in the orbitofrontal cortex (OFC) become active. It has been proposed that these neurons are engaged in outcome simulation, potentially providing a substrate for the evaluation of action plans represented by hippocampal ensembles. **c** | In the sensory-preconditioning task (as depicted on the top row), animals learn a predictive relationship between two neutral stimuli, such as tones, during the preconditioning phase. During the conditioning phase of this task, one of these stimuli is paired with reward. In test sessions, animals responded to the preconditioned cue that, although never directly paired with reward, predicts the occurrence of the conditioned cue. This behaviour has been likened to inference, as animals seem to correctly derive the implicit causal structure of the task (that is, tone A is followed by tone B, which is followed by a reward) despite never having directly experienced this arrangement. Lesion and inactivation data suggest a potential neural model of this task (as depicted on the bottom row) in which hippocampal ensembles encode the relationships between elemental stimuli during the preconditioning and conditioning phases of the task, and this information is accessed by the OFC to support responding during the test session.