Trends in
Neurosciences

CellPress
OPEN ACCESS

Opinion

Awake replay: off the clock but on the job

Matthijs A.A. van der Meer[1,3,*] and Daniel Bendor[2,3,*]

Hippocampal replay is widely thought to support two key cognitive functions: online decision-making and offline memory consolidation. In this review, we take a closer look at the hypothesized link between awake replay and online decision-making in rodents, and find only marginal evidence in support of this role. By contrast, the consolidation view is bolstered by new computational ideas and recent data, suggesting that (i) replay performs offline fictive learning for later goal-oriented behavior; and (ii) replay tags memories prior to sleep, prioritizing them for consolidation. Based on these recent advances, we favor an updated and refined role for awake replay – that is, supporting prioritized offline learning and tagging outside the hippocampus – rather than a direct, online role in guiding behavior.

Does hippocampal replay live a double life?

Among the most striking phenomena in the electrical activity of neurons is the hippocampal **sharp wave-ripple (SWR)** (see Glossary): a sudden, spontaneously generated, highly synchronous burst of spikes whose reverberations are felt throughout the brain [1–4]. Why would the brain invest in such a brief but metabolically demanding event? A leading hypothesis is that SWRs serve memory consolidation [5–31]: the fast-learning hippocampus 'trains' a more distributed, slower-learning knowledge network elsewhere, through the **replay** of recent memories ('systems-level consolidation') (Box 1). In support of this idea, early studies showed that the content of replay during sleep reflects recent experience and that disruptions of sleep SWRs interfere with learning [5,6,8,21,22]. It should be noted that while some SWRs contain sequential firing patterns of neuronal ensembles, depicting specific spatio-temporal trajectories (replay), other SWRs do not contain decodable sequences, but may nonetheless contribute to consolidation through other means (e.g., non-sequential reactivation of cell assemblies, or promoting an excitable state in downstream networks) [30,31].

While most prevalent during sleep, SWRs and replay also occur during awake rest [32–36]. Initially, awake replay was hypothesized to support memory consolidation in much the same way as sleep replay [37], but work in rodents showed striking relationships between awake replay content and the animal's immediate next action, including trajectory events to a goal location [38,39]. Moreover, in rats, disruption of awake SWRs were shown to impair performance on the working memory-guided component (but not the procedural-guided component) of a T-maze alternation task [40]. Further studies expanded on these observations by showing correlations between SWR content and subsequent behavior, and improvement of performance by prolonging awake SWRs [41–44]. These findings suggested that SWRs could also play a role in supporting task performance in the moment (i.e., **online** and for immediate use) rather than in the service of learning for an unspecified later purpose. Processes such as the generation and evaluation of future possibilities (i.e., planning) and retrieval of relevant episodic memories are thought to underlie this online role. Several excellent recent reviews have helped canonize this prevailing view that SWRs and replay can serve dual functions: **offline** consolidation and online task performance [45–47].

Highlights

Hippocampal replay is commonly thought to serve a dual purpose: online decision-making (while awake) and offline memory consolidation (during sleep). In this article, we highlight experimental evidence inconsistent with awake replay's postulated role in online decision-making.

One alternative role for awake replay is to improve future goal-oriented behavior by internally simulating past experiences and using them to adjust estimates of future rewards, prioritizing those value updates that best align decisions with newly encountered outcomes.

A second alternative role for awake replay is memory tagging, where salient experiences are reactivated, to create a latent excitable state within hippocampal–cortical circuits. This latent state helps prioritize these memories for replay during sleep, leading to their consolidation.

We hypothesize that the likelihood for a spatial trajectory to reactivate during awake replay increases with prediction error, including novelty and unexpected changes to reward, stimulus, and context.

[1]Department of Psychological & Brain Sciences, Dartmouth College, Hanover, NH 03755, USA
[2]Institute of Behavioural Neuroscience, Dept. of Experimental Psychology, University College London, London, UK
[3]Equal contribution

*Correspondence:
mvdm@dartmouth.edu
(M.A.A. van der Meer) and
d.bendor@ucl.ac.uk (D. Bendor).

Box 1. Why is replay good for memory consolidation?

Hippocampal replay is classically thought of as the sequential reactivation of place cells during sharp wave-ripples (SWRs), providing a memory trace of a previously experienced spatial trajectory [7]. Many of the characteristics of the firing patterns of neurons present during the behavioral episode can be reinstated during a replay event, including which neurons fire, when they fire, and even the relative magnitude of their firing rates [6,8,9]. In rodents, replay can be directly studied by analyzing the sequential patterns of spiking activity from neuronal ensembles. In contrast, human replay studies often need to infer sequential content from indirect measures [19,20,29]. Although direct evidence of replay causing changes in synaptic strength remains to be demonstrated, the key properties of replay believed to be central to its role in memory include the following (see also Figure I):

1) Compression: replay sequences are temporally compressed during SWRs (~10×) compared with the same sequences during behavior, bringing behavioral sequences within the timescale of spike-time dependent plasticity favorable for long-term storage [8,10].
2) Context-specificity: because place cells remap when a rodent explores a new environment, different behavioral episodes get mapped onto different place cell sequences [11–13]. This creates different replay sequences for each context, minimizing interference while enhancing decoding of these memories [14].
3) Content-addressability: awake SWR content is biased to the animal's current location, and cortical activity prior to a SWR is informative about SWR content, demonstrating that SWR content is not random but can be 'queried' by neural activity in cortical (and likely other) brain areas [15,16]. This is an important property because it ensures that SWRs replay information that is relevant to consolidation-related activity in the cortex [17].
4) Compositionality: the retrieval of one SWR can initiate retrieval of a next, related one. For instance, an SWR may start at the location where the previous one ended, a.k.a. 'chaining', thereby creating long sequences or even shortcuts that were not directly experienced [10,18]. This property supports the discovery of indirect or implied associations (e.g., transitive inference) [19,20,29].

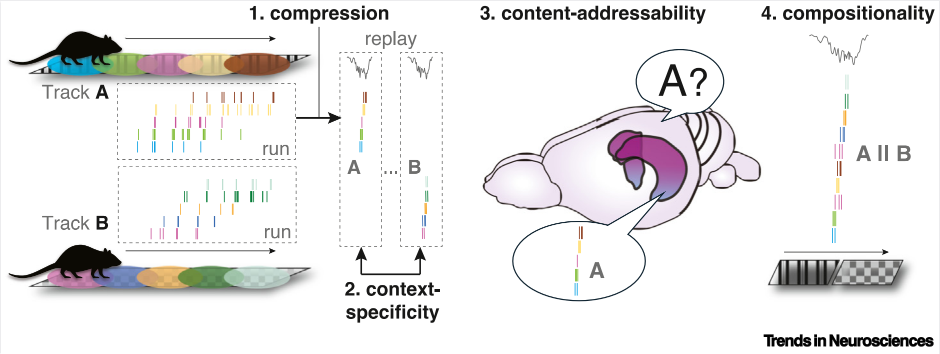

Figure I. Properties of hippocampal replay advantageous for memory consolidation. Different place cell sequences are observed for different spatial trajectories and contexts. During replay, these same sequences are reinstated both in the same temporal order, albeit compressed compared with the timescale of behavior (1), and using the same context-specific ensembles of place cells (2). Hippocampal replay content can be influenced both by cortical networks: content-addressability (3) and by the previously replayed trajectory in the hippocampus: compositionality (4), thought to be important for communication and inference, respectively. Image for panel (3) adapted from doi.org/10.5281/zenodo.3925923, creative commons license (CC-BY).

In this opinion article, we revisit the question of what the role(s) of awake replay are, in the light of new evidence and new theoretical ideas. In particular, we cast a critical eye on previous findings arguing for awake replay's role in task performance (the real-time, online execution of tasks) rather than learning (the offline optimization of decision-making). Based on recent negative results, as well as learning accounts that are hard to distinguish from planning, we highlight a more parsimonious interpretation within the theoretical framework of reinforcement learning. We next review recent studies suggesting that awake replay can also support a previously unexplored intermediate stage of memory consolidation: that of tagging memories prior to sleep, so that they are prioritized for consolidation over other competing memories. Taken together, we believe the weight of evidence favors an updated and refined offline role in learning for awake replay, rather than a direct, online role in guiding behavior.

Glossary

Bellman backup: a core concept in reinforcement learning. Refers to a single value update (of a state, or a state–action pair) as a result of a reward prediction error in algorithms such as temporal-difference reinforcement learning. Bellman backups are used to iteratively improve estimates of a value function by 'backing up' information about neighboring states to refine the value estimate of the current state.

Expected value of backup (EVB): the expected improvement in decision-making performance after performing a particular Bellman backup. It evaluates how beneficial it is to update the value of a given state based on its impact on what the agent will do differently as a result of the updated value function.

Memory tagging: a mechanism by which neurons active during an experience are selectively induced into a latent excitable state. During sleep, this excitable state is unmasked, leading to their increased excitability over other 'untagged' neurons and the prioritization of these neurons to participate in future replay (and the consolidation of these specific memories).

Memory triage: the brain is hypothesized to prioritize salient experiences for consolidation, replaying them first and at a higher rate. Experiences judged to be irrelevant for future behavior are triaged, leading to an insufficient number of replays to be consolidated, eventually resulting in their forgetting.

Offline: occurring outside active task execution (e.g., rest following reward consumption, pauses during inter-trial intervals, or sleep).

Online: occurring during active task execution (e.g., when engaged in real-time behavior such as occurs at a decision point).

Prediction error: a difference (mismatch) between an observed and expected observation.

Replay: the spontaneous reactivation of a sequential pattern of neural activity that has previously occurred during behavior. A common example is the sequential firing of place cells during sleep that mirrors the same activity pattern observed when the rat has previously run a trajectory through a maze.

The farther backward you can look, the farther forward you can see

At first glance, replay appears to be a useful mechanism to support online decisions: episodic retrieval of a relevant previous experience could influence an animal to do the same thing again (if it turned out well last time) [48,49]. Similarly, the imagination of possible upcoming courses of action and their outcomes (i.e., prospective planning or 'rollout' of future states) could bias decisions [38,39,43,50]. However, what is the evidence that awake replay/SWRs support immediate, online decision-making? The literature to date has taken three complementary approaches to investigate this issue: (i) examining replay content and its relationship to behavior (i.e., what the animal chooses to do next); (ii) examining where, when, and how often replay is occurring and how that relates to task performance; and (iii) studying what happens when SWRs are disrupted or enhanced. We review the evidence from these approaches in turn.

As previously noted, several studies have found correspondence between what is replayed and the next trajectory taken by the animal. This effect is particularly clear in open field navigation situations, but also found on some linear maze tasks, in apparent support of a role in online decisions [38,41,51,52]. However, other studies have shown replay content that seems hard to reconcile with a role in online decision-making. In an eight-arm maze where the rewarded arm is changed across blocks of trials, rats tended to replay the previously rewarded arm rather than the currently rewarded arm, echoing earlier results showing replay of the non-chosen arm on two-arm mazes [18,53,54]. In another set of experiments, after conditioning an animal with an aversive stimulus at one end of a linear track, replay was observed to the area actively avoided [55]. One interpretation of these data is that online performance may be supported by replay in different ways: replay of the goal about to be chosen, the alternative to be avoided, or the previously taken option could all be useful for decision-making. Alternatively, awake replay can occur in different behavioral states (e.g., preparation, consumption, resting) that influence its function. How the hippocampus interacts with other brain areas during replay will be critical to understanding this flexibility in function, but it is this very flexibility that complicates our ability to draw clear conclusions about replay function from replay content alone [39,47,49].

Recent results in rats exploring a behavioral apparatus to find rewards indicate that during stopping periods, replay trajectories initially avoid recently visited locations, a consequence of spike-frequency adaptation [56]. This may have a different impact on the content of replay, depending on whether the behavioral task is balanced (e.g., an open field with a random distribution of trajectories), or unbalanced (e.g., a multi-arm maze with a single rewarded arm receiving the majority of trajectories) [57,58]. In a balanced task where the goals of past and future trials are typically in different directions, replay avoiding a past trajectory will be more aligned to trajectories towards the new goal (prospective replay) [38,56]. This may also provide an explanation for the observed bias in prospective replay, during periods of engagement (immediately after stopping and before running), for rats running back and forth along an extended linear track [56]. In contrast, in situations where the rewarded maze arm is fixed across trials (i.e., experience is unbalanced), prospective and retrospective trajectories likely both match the most recently traveled path. This account may help explain 'paradoxical replay', with avoidance of the past trajectory leading to replay of the less visited (or actively avoided) arm of the maze [18,53–55].

If replay supports decision-making, then the presence, number, and/or quality of replays on a given trial should predict task performance. Given the number of studies with pertinent data, it is notable how few positive results there are in the literature. One study in rats found higher quality replay events preceding correct compared with error trials, specifically during the early stages of learning a W-maze alternation task [41]. Additionally, longer duration SWRs can be observed for correct compared with error trials on the same maze; in both cases, effect sizes were modest (e.g.,

Representational drift: the phenomenon in which the neural representation of specific information, such as a stimulus, environment, or task, changes over time despite stable external conditions and consistent behavioral output.
Reward prediction error: the difference between observed and expected reward.
Sharp wave-ripple: transient, high-frequency oscillatory event observed in the hippocampal area CA1 during awake rest or sleep. Consists of a sharp wave (a large depolarization in the local field potential) and a ripple (a fast oscillation). SWRs may be accompanied by particular patterns of spiking activity, which may constitute reactivation or replay of specific stimuli, contexts, or experiences.
Temporal-difference reinforcement learning: a type of machine learning algorithm for predicting future rewards and learning optimal behaviors by updating value functions based on reward prediction errors.
Value function: the expected cumulative reward an agent can obtain from a given state (or state–action pair).

60% chance of predicting an error, when chance is 50%) [44]. However, there are also negative results: using neurofeedback in rats to substantially increase the number of SWRs during the delay period within an eight-arm maze task, no behavioral benefit was reported [59]. Thus, there is not a consistent relationship between awake replay/SWRs and task performance across studies.

Perhaps the most critical test of the hypothesis that replay supports online decision-making is observing whether disrupting (or alternatively prolonging) awake SWRs has a causal effect on task performance. However, the key experiments where animals are first trained to meet a performance criterion on the task before SWR manipulation is applied do not report whether performance was significantly affected during the critical first session [40,44]. Thus, it is possible the main effects found in these studies are due to replay contributing to learning (within sessions and/or across days) rather than to task performance. Indeed, more recent SWR disruption experiments have interpreted these experimental manipulations as affecting within-session learning rather than performance [60]. Furthermore, a recent report found a null effect of SWR disruption on two different tasks that require trial-unique memory, casting further doubt on the role of awake SWRs on decision-making [61].

In summary, while awake replay has properties that in theory would make it helpful for online decision-making, experimental data appears to provide only marginal evidence for such a role. How, then, can it be explained why under some experimental conditions replay still looks strikingly like planning? Could replay closely follow future choices of the animal when in fact it is doing something else? Reinforcement learning, a branch of machine learning concerned with how artificial agents can learn from experience what actions lead to the most reward, offers a possible answer to this mystery [62,63].

How many roads must a map replay, before you've updated your map?

A fundamental challenge faced by reinforcement learning agents is the temporal credit assignment problem: feedback from the environment is typically scarce, and it is not obvious which of the many actions that preceded that feedback deserve how much of the credit (or blame) for the outcome. Early reinforcement learning work noted the value of a replay-like mechanism to help solve this credit assignment problem: Rich Sutton's pioneering Dyna model [64] learned not only from direct experience (online), but could also learn offline from simulated experience, substantially speeding up learning; an idea cited in experimental studies as a potential function of reverse replay in particular [32,65,66]. Many reinforcement learning studies have shown the benefits of replay since, including a key demonstration from a reinforcement learning agent (a 'deep-Q network') learning to play Atari video games [67]. Among the specific functions supported by replay in such models is transitive inference: if you know that A leads to B, and you separately experience B unexpectedly leading to reward, you should update your reward estimate of A even though it has never been paired with reward. By replaying the experience, A→B followed by B→reward, it becomes possible to associate A→reward [19,20]. Thus, offline replay serves as a kind of fictive learning, fixing inconsistencies in expected reward values for different stimuli or states that are connected. A single such update is referred to as a '**Bellman backup**': a single step in the **temporal-difference reinforcement learning** algorithm for iteratively computing the optimal **value function** [68].

Mapping these ideas on to biological systems, it has been previously proposed that replay generated in the hippocampus could speed up learning in a simulated T-maze task, by providing fictive experience for offline learning that happens outside of the hippocampus [69]. This replayed experience was used to update action values stored in the model's 'basal ganglia', which benefited subsequent task performance. Thus, a key idea introduced by this model architecture is that although the function of hippocampal replay is offline learning, stored in network weights elsewhere, these updated weights are available to support task performance the next time the animal returns to the task, independently of the hippocampus or replay [68,70–73]. This idea blurs the distinction between online and offline roles of replay [64,69,71].

In addition to making explicit process-level proposals of how replay could contribute to decision-making, reinforcement learning models can also make normative predictions about what specific content is most useful to replay. As noted earlier, replay is a way to make efficient use of experience, as illustrated by an animal discovering unexpected reward at a new goal location in an open field (Figure 1, left): replay can then be used to increase the future expected reward from all adjacent locations, not just those that the animal took to get to the reward. Recent work has considered the crucial question raised by this scenario: given finite time to replay, what is the most useful [68,73]? The normative assumption that replay will prioritize those learning updates that will be most useful to the animal in the future (quantified as the product of how big the inconsistencies are – 'gain' – and how likely the animal is to encounter them in the future – 'need'; Figure 1) generates a predicted pattern of replay content. This '**expected value of backup**' **(EVB)** pattern matches what is expected from prospective planning even though in this model, the actual role of replay is not in online decision-making but to train reward representations located outside the hippocampus [38]. In addition to accounting for the experimentally observed replay content in many studies, experiments also support the key model requirement that **reward prediction errors** cause replay [71,74,75]. Note that the core idea of using **prediction errors** to prioritize replay is not unique to reward learning, but may also extend to map-learning [76,77]. Thus, what may look like replay content indicative of online decision-making may equally well be predicted by offline learning that prioritizes those updates that are most useful.

To consolidate or not to consolidate, that is the question

During a behavioral task, awake replay trajectories typically originate from the animal's current location and represent possible paths within the local environment. By contrast, when replay

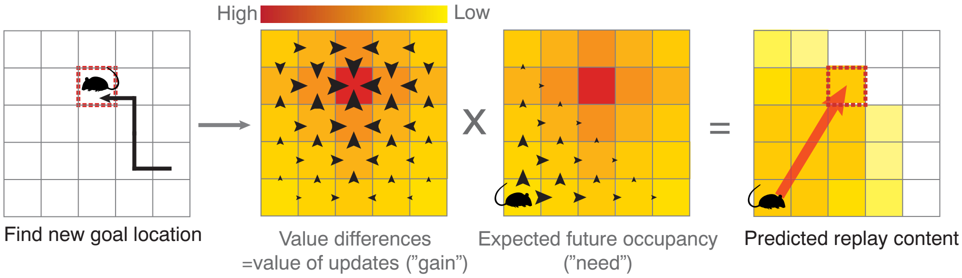

Figure 1. Schematic of the idea that replay implements value function updates from fictive experience, prioritized by how useful those updates are likely to be. After finding a new goal location in an open field (high value, tile delineated in red in left panel), the value of nearby tiles is also increased by virtue of their proximity to the goal, creating a mismatch between the true value function (right) and the animal's current representation of the value function. Replay is a potential mechanism to update these values and make them consistent with each other: in effect, each replay is a fictive experience that the animal can learn from. If replays are prioritized according to the largest value differences needing updating ('gain', arrowheads in second panel) and which locations the animal is likely to encounter ('need'; arrowheads in third panel) then replay trajectories look similar to prospective planning (right panel; red arrow) [68]. However, in this account, replay does not play a role in online decision-making process, but rather updates neural circuits responsible for making decisions. This view also explains why prospective replay is correlated with learning, but may not always reflect the immediate future choice and can still be observed in the absence of choice, such as a simple linear track that only involves running from one side to the other [39,43,56].

occurs later during sleep, replayed trajectories can occur in an unbiased fashion (i.e., any trajectory from any previous context can potentially replay) and, while this may give more flexibility for what the brain can update or consolidate, the challenge of prioritizing is now even more daunting. Recent experiences needing to be replayed reach a bottleneck because there is not enough time during sleep to sufficiently replay all our waking experiences. Thus, the brain must select what it needs to remember and what it can afford to forget, a process referred to as **memory triage** [78]. Analogous to prioritizing more useful updates in reinforcement learning, those memories that are most likely to be relevant in the future also need to be prioritized, replaying first and often during sleep, to increase their likelihood of being consolidated. How does the hippocampus ‘select’ which memories to replay, given that sleep is an internally-driven state and all of the factors that relate to the salience of the memory are no longer present to inform the brain that this memory needs to be prioritized during consolidation?

Awake replay has recently emerged as a potential solution to this problem [79–81]. First, cortico-hippocampal circuits active during a behavioral episode are hypothesized to be actively ‘tagged’ during awake replay (Figure 2). More salient experiences have more awake replay events cumulatively, which strengthens this cortico-hippocampal tag and the associated sleep replay priority. How this tag causes more sleep replay is unknown, but one model suggests that cortico-hippocampal circuits have a latent excitable state, with this excitability proportional to the number of times a memory replays prior to sleep (strength of the tag) [82,83]. During sleep, this latent excitable state is unmasked prior to a replay event and this increases the likelihood that the hippocampus will remap to the context associated with this ‘tagged’ cortico-hippocampal circuit and subsequently replay the memory associated with this context.

Given that behavioral variables increasing a memory’s salience, such as novelty and reward, modulate both awake and sleep replay rates [32,79,84], what is the evidence that awake replay is actually doing something different from sleep replay? When a behavioral episode is repeated, like when a rat runs multiple spatial trajectories on a novel linear track, the rate of awake replay will stay the same or even decrease over time, while the rate of sleep replay will increase proportionally with the number of laps run [79]. However, running more laps will still increase the cumulative number of awake replay events, simply because of more time spent in rest or immobility

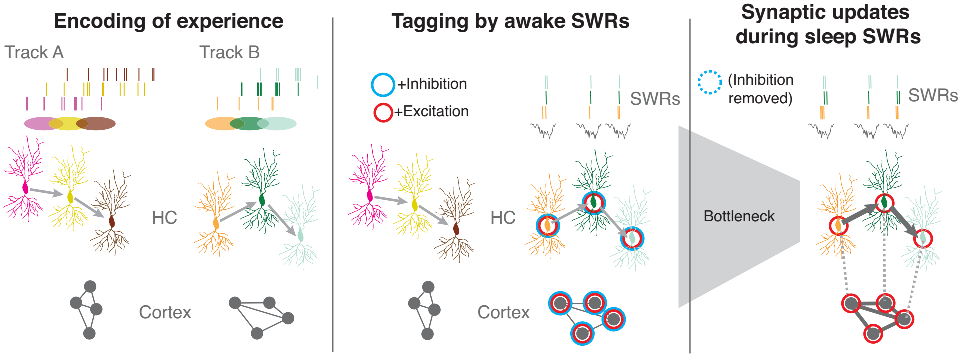

Trends in Neurosciences

Figure 2. Schematic of awake SWRs tagging memories for future consolidation. Different cortical and hippocampal ensembles will encode each experience (left). Behavioral salience (e.g., novelty, reward) will lead to more awake SWRs for some experiences compared with others (here, Track B is prioritized). This creates a memory tag by generating a latent excitable state (middle), by a balanced increase in the strength of excitatory and inhibitory inputs (middle). In sleep, disinhibition during an up state leads to an unmasking of tagged memory circuits, with their increased excitability leading to higher participation in sleep replay, prioritizing these memories for consolidation (right). Abbreviations: HC, hippocampus; SWR, sharp wave-ripple.

when SWRs generally occur [79]. This highlights that while the rate of awake replay may reflect the EVB or salience of a behavioral episode (i.e., higher in a novel vs. familiar track), it is the cumulative number of awake replay events that correlates with the subsequent sleep replay rate (and consolidation priority) of the memory.

There are several behavioral studies that support this process of memory triage and prioritization ([85] mice, [86–90] humans). Importantly, in human subjects, the ability to bias memory processing towards more salient experiences requires sleep within a restricted time window and correlates with hippocampal involvement [91]. Additionally, in human subjects, what makes a memory salient does not need to occur at the same time that the memory is initially encoded (e.g., learning before you go to sleep that you have an exam the next day should be sufficient to prioritize your sleep consolidation towards what the subject you studied earlier in the day matching the exam topic) [92,93]. Given that awake replay can also occur remotely (in a different context to where the memory was originally encoded) allows the brain to use this same mechanism to tag any earlier memory based on new information. While **memory tagging** does not need to be exclusively carried out by hippocampal replay, such a mechanism is particularly advantageous if offline updating is required.

It's tough to make predictions, especially about the future

Which mechanisms could support replay prioritization? Increased prediction error, the result of novelty and unexpected changes in stimulus, reward, and context, leads to higher rates of awake replay [32,71,74,75,79]. Could prediction error be the main causal driver for prioritizing specific trajectories for awake replay (Figure 3)? Unexpected rewards (or the lack of them)

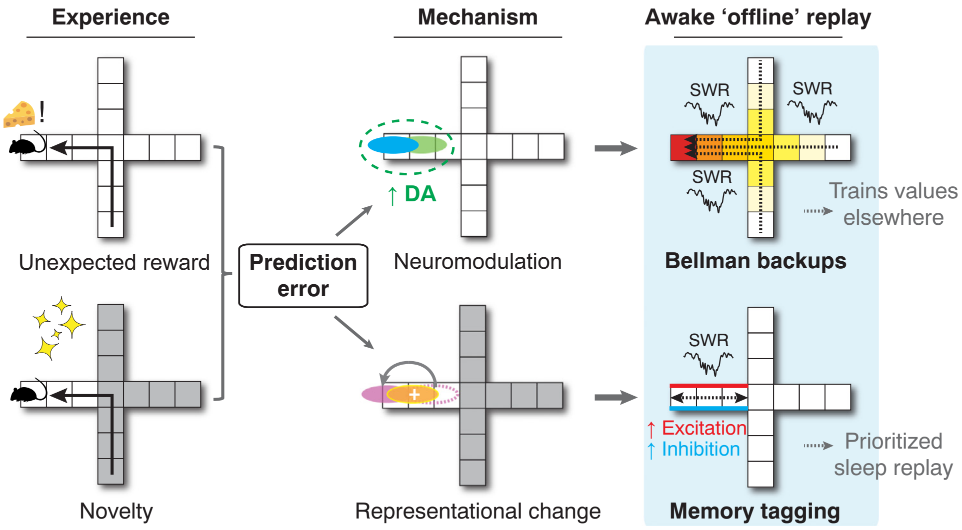

Trends in Neurosciences

Figure 3. Proposed roles for awake replay in reinforcement learning and memory prioritization more generally. (Left) Experience: prediction errors, arising from unexpected reward (top) or novelty (bottom) trigger increases in awake replay through non-exclusive, parallel mechanisms. (Middle) Mechanism: prediction error induces neuromodulation, particularly dopamine (DA) release, which directly promotes excitability in nearby place fields and can also drive the allocation of new place fields (orange) and partial remapping of existing place fields (purple). (Right) Awake replay performs Bellman backups (top) by propagating reward value estimates to preceding states, supporting offline learning outside the hippocampus. Awake replay also can drive memory tagging (bottom), where excitation–inhibition balance modifies cortical–hippocampal dynamics, biasing subsequent sleep replay toward more salient experiences. Together, these processes suggest that awake replay does not serve an online decision-making role but rather facilitates prioritized offline learning and consolidation. Abbreviation: SWR, sharp wave-ripple.

classically cause changes in the firing of midbrain dopamine neurons, which, in addition to strong projections to the striatum, also send projections to the hippocampus (although hippocampal dopamine may also, or even primarily, arise from locus coeruleus neurons) [94–97]. Unexpected state transitions (e.g., A→C instead of the expected A→B) have been shown to similarly cause ventral tegmental area (VTA) activity increases, even in the absence of reward [98]. In turn, hippocampal dopamine release enhances hippocampal replay; putting these ideas together, we speculate that phasic dopamine release increases the latent excitability of recently active place cells, increasing their participation in awake SWRs [99,100]. The use of neuromodulatory signals to prioritize replay may not be unique to dopamine, given that acetylcholine and norepinephrine putatively signal for salience and arousal, respectively [101,102].

An alternative or potentially complementary mechanism to support replay prioritization is representational change. In a completely novel experience (e.g., a new environment, a new internal context inferred by the animal, or a never-before seen stimulus), neurons are ‘allocated’ in the hippocampus to encode that novel experience, manifesting as new ‘emerging’ place fields being formed [103–106]. Place cells can change their place fields in other ways, including partial or global remapping to sudden changes in the environment and smaller changes over longer time scales (i.e., **representational drift**) [107–109]. Here we group together emerging place fields, drift and remapping, as part of the same continuum of representational change and suggest that the underlying cause is prediction error. In turn, this representational change drives replay priority: awake replay rates are observed to increase proportionally to representational change [99,110]. For example, when rats are re-exposed to a linear track, awake replay rates increase with the difference in the population vector correlation between the first and second exposure, with this difference increasing in less familiar environments [79]. Thus, prediction errors during experience can be expressed as representational change in the hippocampus, which in turn could prioritize associated awake replay activity (Figure 3).

Concluding remarks and future perspectives

A widely held view is that hippocampal replay can serve two functional roles: (i) online decision-making (by retrieving relevant memories and/or generating possible courses of action, i.e., planning); and (ii) offline consolidation (i.e., the updating of knowledge structures outside the hippocampus). As discussed in the previous sections, we find only marginal experimental support for (i), with mixed results that, in our interpretation, are equally or more consistent with a role in learning rather than in online performance. We next highlighted two alternative offline roles for awake replay: (i) fictive learning to update value functions, and (ii) the ‘tagging’ of memories to prioritize their later replay during sleep. Note that these hypothesized roles do not need to be exclusively performed by awake replay. For example, theta sequences, the sequential firing of place cells within a theta cycle, which generally occur during locomotion and exploration, could also generate trajectories comparable with replay to support planning and fictive learning. Interestingly, while theta sequences are required for sleep replay to later occur, the number of theta sequences during a behavioral episode does not correlate with the rate of sleep replay (of that experience), suggesting that theta sequences are not involved in memory tagging [30,79,111].

Moving forward, it is worth noting that while unexpected reward and novelty increase awake replay activity, these are also more generally correlated with many other behavioral or neural variables, including arousal levels, motivational state, neuromodulatory signals, prediction errors, and place field remapping and representational drift in the hippocampus and elsewhere. We propose that prediction error is the main driver guiding which awake replay trajectories are prioritized, through the downstream effects of prediction errors on neuromodulatory signals and representational change, that in turn influence which neurons participate in replay. Understanding the

Outstanding questions

The entorhinal cortex provides a major input to the hippocampus, while receiving reciprocal feedback connections. Could the entorhinal cortex help prioritize what the hippocampus replays, based on the difference between the input to the hippocampus (from superficial layers of entorhinal cortex) and the output of the hippocampus (to deep layers of entorhinal cortex), with prediction error causing bigger mismatches?

One speculated role of awake replay is to ‘tag’ memory circuits so that they are prioritized to replay during sleep, leading to their consolidation. Where is this memory prioritization tag found in the cortico-hippocampal circuit? Is this distributed throughout the entire brain, localized to the hippocampus, or restricted to the cortical area driven by the salient cue that has led to the prioritization of the memory?

What plastic changes are caused by awake replay? Learning updates (i.e., Bellman backups) driven by awake replay should cause representational change in cortico-hippocampal circuits. On the contrary, memory tagging should only cause a latent excitability in these circuits. If replay can serve both roles, what controls the balance between replay-for-stability and replay-for-updating? Or, alternatively, are these simply reflections of the same common underlying mechanism?

The predicted replay content from replay-as-online-planning and replay-as-prioritized-offline-value-learning accounts overlaps. What is needed in the design of a strong test of the hypothesis that replay supports online decision-making? At a minimum, animals should be trained to criterion, with SWRs disrupted at the decision point only on a subset of trials, allowing the difference in behavioral performance between disrupted and control trials to be directly compared.

precise role(s) of replay amidst such a tangled web of processes may require thinking more broadly and unifying these many strands of experimental evidence (see Outstanding questions). While it is technically challenging to simultaneously monitor hippocampal replay in combination with other neural or behavioral factors, this may ultimately be required to tease apart the specific contributions of replay, including understanding how replay drives and modifies neural circuits in the brain [112–116]. Similarly, measuring replay content across different types of behavioral tasks will be required to gain insights into the internal logic that helps determine which experiences are selected for consolidation during sleep.

Acknowledgments

This work was supported by funding from NIMH R01 MH123466 and NSF MCA 2422730 (M.A.A.v.d.M.) and BBSRC - BB/Y010345/1 (D.B.).

Declaration of interests

The authors declare no competing interests.

References

1. Logothetis, N.K. *et al.* (2012) Hippocampal–cortical interaction during periods of subcortical silence. *Nature* 491, 547–553
2. Buzsáki, G. (2015) Hippocampal sharp wave-ripple: a cognitive biomarker for episodic memory and planning. *Hippocampus* 25, 1073–1188
3. Liu, A.A. *et al.* (2022) A consensus statement on detection of hippocampal sharp wave ripples and differentiation from other fast oscillations. *Nat. Commun.* 13, 6000
4. Sebastian, E.R. *et al.* (2023) Topological analysis of sharp-wave ripple waveforms reveals input mechanisms behind feature variations. *Nat. Neurosci.* 26, 2171–2181
5. Pavlides, C. and Winson, J. (1989) Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. *J. Neurosci.* 9, 2907–2918
6. Wilson, M.A. and McNaughton, B.L. (1994) Reactivation of hippocampal ensemble memories during sleep. *Science* 265, 676–679
7. Foster, D.J. (2017) Replay comes of age. *Annu. Rev. Neurosci.* 40, 581–602
8. Lee, A.K. and Wilson, M.A. (2002) Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* 36, 1183–1194
9. Tirole, M. *et al.* (2022) Experience-driven rate modulation is reinstated during hippocampal replay. *eLife* 11, e79031
10. Davidson, T.J. *et al.* (2009) Hippocampal replay of extended experience. *Neuron* 63, 497–507
11. Bostock, E. *et al.* (1991) Experience-dependent modifications of hippocampal place cell firing. *Hippocampus* 1, 193–205
12. Wills, T.J. *et al.* (2005) Attractor dynamics in the hippocampal representation of the local environment. *Science* 308, 873–876
13. Leutgeb, S. *et al.* (2005) Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* 309, 619–623
14. Takigawa, M. *et al.* (2024) Evaluating hippocampal replay without a ground truth. *eLife* 13, e85635
15. Bendor, D. and Wilson, M.A. (2012) Biasing the content of hippocampal replay during sleep. *Nat. Neurosci.* 15, 1439–1444
16. Rothschild, G. *et al.* (2017) A cortical–hippocampal–cortical loop of information processing during memory consolidation. *Nat. Neurosci.* 20, 251–259
17. Ji, D. and Wilson, M.A. (2007) Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nat. Neurosci.* 10, 100–107
18. Gupta, A.S. *et al.* (2010) Hippocampal replay is not a simple function of experience. *Neuron* 65, 695–705
19. Barron, H.C. *et al.* (2020) Neuronal computation underlying inferential reasoning in humans and mice. *Cell* 183, 228–243
20. Kurth-Nelson, Z. *et al.* (2023) Replay and compositional computation. *Neuron* 111, 454–469
21. Girardeau, G. *et al.* (2009) Selective suppression of hippocampal ripples impairs spatial memory. *Nat. Neurosci.* 12, 1222–1223
22. Ego-Stengel, V. and Wilson, M.A. (2010) Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus* 20, 1–10
23. Frankland, P.W. and Bontempi, B. (2005) The organization of recent and remote memories. *Nat. Rev. Neurosci.* 6, 119–130
24. Eschenko, O. *et al.* (2008) Sustained increase in hippocampal sharp-wave ripple activity during slow-wave sleep after learning. *Learn. Mem.* 15, 222–228
25. Diekelmann, S. and Born, J. (2010) The memory function of sleep. *Nat. Rev. Neurosci.* 11, 114–126
26. Varela, C. *et al.* (2016) Tracking the time-dependent role of the hippocampus in memory recall using DREADDs. *PLoS One* 11, e0154374
27. Squire, L.R. *et al.* (2015) Memory consolidation. *Cold Spring Harb. Perspect. Biol.* 7, a021766
28. Crowley, R. *et al.* (2019) A review of neurobiological factors underlying the selective enhancement of memory at encoding, consolidation, and retrieval. *Prog. Neurobiol.* 179, 101615
29. Schwartenbeck, P. *et al.* (2023) Generative replay underlies compositional inference in the hippocampal-prefrontal circuit. *Cell* 186, 4885–4897
30. Liu, C. *et al.* (2023) Associative and predictive hippocampal codes support memory-guided behaviors. *Science* 382, eadi8237
31. Genzel, L. *et al.* (2020) A consensus statement: defining terms for reactivation analysis. *Philos. Trans. R. Soc. B* 375, 20200001
32. Foster, D.J. and Wilson, M.A. (2006) Reverse replay of behavioral sequences in hippocampal place cells during the awake state. *Nature* 440, 680–683
33. Kudrimoti, H.S. *et al.* (1999) Reactivation of hippocampal cell assemblies: effects of behavioral state, experience, and EEG dynamics. *J. Neurosci.* 19, 4090–4101
34. O'Neill, J. *et al.* (2006) Place-selective firing of CA1 pyramidal cells during sharp wave/ripple network patterns in exploratory behavior. *Neuron* 49, 143–155
35. Karlsson, M.P. and Frank, L.M. (2009) Awake replay of remote experiences in the hippocampus. *Nat. Neurosci.* 12, 913–918
36. Tang, W. *et al.* (2017) Hippocampal-prefrontal reactivation during learning is stronger in awake compared with sleep states. *J. Neurosci.* 37, 11789–11805
37. Dupret, D. *et al.* (2010) The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nat. Neurosci.* 13, 995–1002
38. Pfeiffer, B.E. and Foster, D.J. (2013) Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* 497, 74–79
39. Ólafsdóttir, H.F. *et al.* (2017) Task demands predict a dynamic switch in the content of awake hippocampal replay. *Neuron* 96, 925–935
40. Jadhav, S.P. *et al.* (2012) Awake hippocampal sharp-wave ripples support spatial memory. *Science* 336, 1454–1458

41. Singer, A.C. *et al.* (2013) Hippocampal SWR activity predicts correct decisions during the initial learning of an alternation task. *Neuron* 77, 1163–1173
42. Ólafsdóttir, H.F. *et al.* (2015) Hippocampal place cells construct reward related sequences through unexplored space. *eLife* 4, e06063
43. Diba, K. and Buzsáki, G. (2007) Forward and reverse hippocampal place-cell sequences during ripples. *Nat. Neurosci.* 10, 1241–1242
44. Fernández-Ruiz, A. *et al.* (2019) Long-duration hippocampal sharp wave ripples improve memory. *Science* 364, 1082–1086
45. Joo, H.R. and Frank, L.M. (2018) The hippocampal sharp wave–ripple in memory retrieval for immediate use and consolidation. *Nat. Rev. Neurosci.* 19, 744–757
46. Ólafsdóttir, H.F. *et al.* (2018) The role of hippocampal replay in memory and planning. *Curr. Biol.* 28, R37–R50
47. Pfeiffer, B.E. (2020) The content of hippocampal "replay". *Hippocampus* 30, 6–18
48. Lengyel, M. and Dayan, P. (2007) Hippocampal contributions to control: the third way. In *Advances in Neural Information Processing Systems 20. Proceedings of the Twenty-First Annual Conference on Neural Information Processing Systems*
49. Carr, M.F. *et al.* (2011) Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nat. Neurosci.* 14, 147–153
50. Redish, A.D. (2016) Vicarious trial and error. *Nat. Rev. Neurosci.* 17, 147–159
51. Widloski, J. and Foster, D.J. (2022) Flexible rerouting of hippocampal replay sequences around changing barriers in the absence of global place field remapping. *Neuron* 110, 1547–1558
52. Xu, H. *et al.* (2019) Assembly responses of hippocampal CA1 place cells predict learned behavior in goal-directed spatial tasks on the radial eight-arm maze. *Neuron* 101, 119–132
53. Gillespie, A.K. *et al.* (2021) Hippocampal replay reflects specific past experiences rather than a plan for subsequent choice. *Neuron* 109, 3149–3163
54. Carey, A.A. *et al.* (2019) Reward revaluation biases hippocampal replay content away from the preferred outcome. *Nat. Neurosci.* 22, 1450–1459
55. Wu, C.T. *et al.* (2017) Hippocampal awake replay in fear memory retrieval. *Nat. Neurosci.* 20, 571–580
56. Mallory, C.S. *et al.* (2025) The time course and organization of hippocampal replay. *Science* 387, 541–548
57. Chen, H.T. and van der Meer, M. (2024) Paradoxical replay can protect contextual task representations from destructive interference when experience is unbalanced. *bioRxiv*, Published online May 9, 2024. https://doi.org/10.1101/2024.05.09.593332
58. Takigawa, M. and Bendor, D. (2025) Replaying off the beaten path. *Science* 387, 476–477
59. Gillespie, A.K. *et al.* (2024) Neurofeedback training can modulate task-relevant memory replay rate in rats. *eLife* 12, RP90944
60. Igata, H. *et al.* (2021) Prioritized experience replays on a hippocampal predictive map for learning. *Proc. Natl Acad. Sci. U. S. A.* 118, e2011266118
61. Deceuninck, L. and Kloosterman, F. (2024) Disruption of awake sharp-wave ripples does not affect memorization of locations in repeated-acquisition spatial memory tasks. *eLife* 13, e84004
62. Roscow, E.L. *et al.* (2021) Learning offline: memory replay in biological and artificial reinforcement learning. *Trends Neurosci.* 44, 808–821
63. Wittkuhn, L. *et al.* (2021) Replay in minds and machines. *Neurosci. Biobehav. Rev.* 129, 367–388
64. Sutton, R.S. (1991) Dyna, an integrated architecture for learning, planning, and reacting. *ACM SIGART Bull.* 2, 160–163
65. Ambrose, R.E. *et al.* (2016) Reverse replay of hippocampal place cells is uniquely modulated by changing reward. *Neuron* 91, 1124–1136
66. Singer, A.C. and Frank, L.M. (2009) Rewarded outcomes enhance reactivation of experience in the hippocampus. *Neuron* 64, 910–921
67. Mnih, V. *et al.* (2015) Human-level control through deep reinforcement learning. *Nature* 518, 529–533
68. Mattar, M.G. and Daw, N.D. (2018) Prioritized memory access explains planning and hippocampal replay. *Nat. Neurosci.* 21, 1609–1617
69. Johnson, A. and Redish, A.D. (2005) Hippocampal replay contributes to within session learning in a temporal difference reinforcement learning model. *Neural Netw.* 18, 1163–1171
70. Stoianov, I.P. *et al.* (2018) Model-based spatial navigation in the hippocampus-ventral striatum circuit: a computational analysis. *PLoS Comput. Biol.* 14, e1006316
71. Momennejad, I. *et al.* (2018) Offline replay supports planning in human reinforcement learning. *eLife* 7, e32548
72. Russek, E.M. *et al.* (2017) Predictive representations can link model-based reinforcement learning to model-free mechanisms. *PLoS Comput. Biol.* 13, e1005768
73. Sagiv, Y. *et al.* (2024) Prioritizing replay when future goals are unknown. *bioRxiv*, Published online March 4, 2024. https://doi.org/10.1101/2024.02.29.582822
74. Roscow, E.L. *et al.* (2019) Behavioural and computational evidence for memory consolidation biased by reward-prediction errors. *bioRxiv*, Published online February 17, 2025. https://doi.org/10.1101/716290
75. Liu, Y. *et al.* (2021) Experience replay is associated with efficient nonlocal learning. *Science* 372, eabf1357
76. Evans, T. and Burgess, N. (2020) Replay as structural inference in the hippocampal-entorhinal system. *bioRxiv*, Published online August 10, 2020. https://doi.org/10.1101/2020.08.07.241547
77. Stoianov, I. *et al.* (2022) The hippocampal formation as a hierarchical generative model supporting generative replay and continual learning. *Prog. Neurobiol.* 217, 102329
78. Stickgold, R. and Walker, M.P. (2013) Sleep-dependent memory triage: evolving generalization through selective processing. *Nat. Neurosci.* 16, 139–145
79. Huelin Gorriz, M. *et al.* (2023) The role of experience in prioritizing hippocampal replay. *Nat. Commun.* 14, 8157
80. Yagi, S. *et al.* (2023) Awake hippocampal synchronous events are incorporated into offline neuronal reactivation. *Cell Rep.* 42, 112871
81. Yang, W. *et al.* (2024) Selection of experience for memory by hippocampal sharp wave ripples. *Science* 383, 1478–1483
82. Barron, H.C. *et al.* (2017) Inhibitory engrams in perception and memory. *Proc. Natl Acad. Sci. U S A* 114, 6666–6674
83. Lewis, P.A. and Bendor, D. (2019) How targeted memory reactivation promotes the selective strengthening of memories in sleep. *Curr. Biol.* 29, R906–R912
84. Cheng, S. and Frank, L.M. (2008) New experiences enhance coordinated neural activity in the hippocampus. *Neuron* 57, 303–313
85. Qiao, Q. *et al.* (2023) Memory capacity and prioritization in female mice. *Sci. Rep.* 13, 14073
86. Baran, B. *et al.* (2012) Processing of emotional reactivity and emotional memory over sleep. *J. Neurosci.* 32, 1035–1042
87. Bennion, K.A. *et al.* (2016) The impact of napping on memory for future-relevant stimuli: prioritization among multiple salience cues. *Behav. Neurosci.* 130, 281
88. Braun, E.K. *et al.* (2018) Retroactive and graded prioritization of memory by reward. *Nat. Commun.* 9, 4886
89. Cowan, E.T. *et al.* (2021) Memory consolidation as an adaptive process. *Psychon. Bull. Rev.* 28, 1796–1810
90. Schapiro, A.C. *et al.* (2018) Human hippocampal replay during rest prioritizes weakly learned information and predicts memory performance. *Nat. Commun.* 9, 3920
91. Rauchs, G. *et al.* (2011) Sleep contributes to the strengthening of some memories over others, depending on hippocampal activity at learning. *J. Neurosci.* 31, 2563–2568
92. Wilhelm, I. *et al.* (2011) Sleep selectively enhances memory expected to be of future relevance. *J. Neurosci.* 31, 1563–1569
93. Zhu, Y. *et al.* (2022) Emotional learning retroactively promotes memory integration through rapid neural reactivation and reorganization. *eLife* 11, e60190
94. Montague, P.R. *et al.* (1996) A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *J. Neurosci.* 16, 1936–1947
95. Gasbarri, A. *et al.* (1997) The dopaminergic mesencephalic projections to the hippocampal formation in the rat. *Prog. Neuro-Psychopharmacol. Biol. Psychiatry* 21, 1–22
96. Takeuchi, T. *et al.* (2016) Locus coeruleus and dopaminergic consolidation of everyday memory. *Nature* 537, 357–362

97. Tsetsenis, T. *et al.* (2023) Dopaminergic regulation of hippocampal plasticity, learning, and memory. *Front. Behav. Neurosci.* 16, 1092420
98. Sharpe, M.J. *et al.* (2017) Dopamine transients are sufficient and necessary for acquisition of model-based associations. *Nat. Neurosci.* 20, 735–742
99. McNamara, C.G. *et al.* (2014) Dopaminergic neurons promote hippocampal reactivation and spatial memory persistence. *Nat. Neurosci.* 17, 1658–1660
100. Duszkiewicz, A.J. *et al.* (2019) Novelty and dopaminergic modulation of memory persistence: a tale of two systems. *Trends Neurosci.* 42, 102–114
101. Hasselmo, M.E. (2006) The role of acetylcholine in learning and memory. *Curr. Opin. Neurobiol.* 16, 710–715
102. Breton-Provencher, V. *et al.* (2021) Locus coeruleus norepinephrine in learned behavior: anatomical modularity and spatiotemporal integration in targets. *Front. Neural Circ.* 15, 638007
103. Josselyn, S.A. and Frankland, P.W. (2018) Memory allocation: mechanisms and function. *Annu. Rev. Neurosci.* 41, 389–413
104. Chen, L. *et al.* (2020) The role of intrinsic excitability in the evolution of memory: significance in memory allocation, consolidation, and updating. *Neurobiol. Learn. Mem.* 173, 107266
105. Lee, D. *et al.* (2012) Hippocampal place fields emerge upon single-cell manipulation of excitability during behavior. *Science* 337, 849–853
106. Priestley, J.B. *et al.* (2022) Signatures of rapid plasticity in hippocampal CA1 representations during novel experiences. *Neuron* 110, 1978–1992
107. Ziv, Y. *et al.* (2013) Long-term dynamics of CA1 hippocampal place codes. *Nat. Neurosci.* 16, 264–266
108. Geva, N. *et al.* (2023) Time and experience differentially affect distinct aspects of hippocampal representational drift. *Neuron* 111, 2357–2366
109. Khatib, D. *et al.* (2023) Active experience, not time, determines within-day representational drift in dorsal CA1. *Neuron* 111, 2348–2356
110. Ormond, J. *et al.* (2023) Enhanced reactivation of remapping place cells during aversive learning. *J. Neurosci.* 43, 2153–2167
111. Drieu, C. *et al.* (2018) Nested sequences of hippocampal assemblies during behavior support subsequent sleep replay. *Science* 362, 675–679
112. Van de Ven, G.M. *et al.* (2016) Hippocampal offline reactivation consolidates recently formed cell assembly patterns during sharp wave-ripples. *Neuron* 92, 968–974
113. Grosmark, A.D. *et al.* (2021) Reactivation predicts the consolidation of unbiased long-term cognitive maps. *Nat. Neurosci.* 24, 1574–1585
114. Roux, L. *et al.* (2017) Sharp wave ripples during learning stabilize the hippocampal spatial map. *Nat. Neurosci.* 20, 845–853
115. Maboudi, K. *et al.* (2024) Retuning of hippocampal representations during sleep. *Nature* 629, 630–638
116. Nguyen, N.D. *et al.* (2024) Cortical reactivations predict future sensory responses. *Nature* 625, 110–118