Memory consolidation from a reinforcement learning perspective

Jong Won Lee[1]* and Min Whan Jung[1,2]*

[1]Center for Synaptic Brain Dysfunctions, Institute for Basic Science, Daejeon, Republic of Korea,
[2]Department of Biological Sciences, Korea Advanced Institute of Science and Technology, Daejeon,
Republic of Korea

Memory consolidation refers to the process of converting temporary memories into long-lasting ones. It is widely accepted that new experiences are initially stored in the hippocampus as rapid associative memories, which then undergo a consolidation process to establish more permanent traces in other regions of the brain. Over the past two decades, studies in humans and animals have demonstrated that the hippocampus is crucial not only for memory but also for imagination and future planning, with the CA3 region playing a pivotal role in generating novel activity patterns. Additionally, a growing body of evidence indicates the involvement of the hippocampus, especially the CA1 region, in valuation processes. Based on these findings, we propose that the CA3 region of the hippocampus generates diverse activity patterns, while the CA1 region evaluates and reinforces those patterns most likely to maximize rewards. This framework closely parallels Dyna, a reinforcement learning algorithm introduced by Sutton in 1991. In Dyna, an agent performs offline simulations to supplement trial-and-error value learning, greatly accelerating the learning process. We suggest that memory consolidation might be viewed as a process of deriving optimal strategies based on simulations derived from limited experiences, rather than merely strengthening incidental memories. From this perspective, memory consolidation functions as a form of offline reinforcement learning, aimed at enhancing adaptive decision-making.

KEYWORDS

simulation-selection model, offline learning, value, dyna, imagination, CA3, CA1

Introduction

Müller and Pilzecker (1900) introduced the term "consolidation", proposing that after successful encoding, a physiological process known as "perseveration" stabilizes memory representations, gradually reducing their susceptibility to interference from new learning (Lechner et al., 1999). Despite over a century of research, the underlying reasons for consolidation and the transformations memories undergo during this process remain incompletely understood. According to the most influential theory, the standard systems consolidation theory, memories, initially encoded and temporarily stored in the hippocampus, are gradually reorganized and distributed across the neocortex. This process ensures that memories become less dependent on the hippocampus over time and are integrated into broader cortical networks for stable, long-term storage (Squire and Alvarez, 1995).

Although substantial evidence supports the systems consolidation theory, there are also findings that challenge its main premises. For instance, recent studies have observed impaired episodic memory without a temporal gradient following hippocampal damage in both humans and animals (Sekeres et al., 2018; Sutherland and Lehmann, 2011; Sutherland et al., 2010). Similarly, brain imaging studies have revealed hippocampal activation associated with the

recollection of vivid, detailed, context-specific memories spanning the entire lifespan (Moscovitch et al., 2016; Moscovitch et al., 2005; Sekeres et al., 2018). These observations have prompted the development of alternative theories.

One such alternative is the multiple trace theory, which posits that memories are not merely transferred from the hippocampus to the neocortex but remain reliant on the hippocampus indefinitely for detailed episodic recollection. According to multiple trace theory, each retrieval of an episodic memory generates a new trace or representation within the hippocampus, reinforcing the memory and enhancing its accessibility over time (Nadel and Moscovitch, 1997). In contrast to systems consolidation theory, multiple trace theory suggests that while semantic memories may become independent of the hippocampus, episodic memories remain hippocampus-dependent for vivid, context-rich retrieval.

Another proposal, the trace transformation theory, emphasizes dynamic, bidirectional interactions between the hippocampus and neocortex. Rather than depicting memory consolidation as a simple handover of responsibility from the hippocampus to the neocortex, trace transformation theory envisions a lifelong process of hippocampal-neocortical collaboration. This theory emphasizes the coexistence and interaction of different forms of memory, suggesting that consolidation involves ongoing reorganization and expression of memories based on hippocampal-neocortical dynamics (Sekeres et al., 2018; Winocur and Moscovitch, 2011; Winocur et al., 2010; Winocur et al., 2007).

Despite extensive debate and research on these theories, one critical aspect has received less attention: the selective nature of memory consolidation. After encoding, not all memories share the same fate—some are forgotten, while others persist for a lifetime, often in a transformed or reconstructed form. Systems consolidation is widely understood to involve extracting general information from specific experiences, resulting in the formation of gists, schemas, and semantic knowledge (Cheng, 2017; McClelland et al., 1995; Moscovitch and Gilboa, 2024). However, indiscriminate generalization of individual experiences may be maladaptive. For survival, prioritizing memories of critical events—such as visiting a location associated with a significant reward or encountering a predator at a specific time and place—provides a clear advantage. Understanding the mechanisms that drive such selective memory consolidation remains a key challenge in memory research.

It has been proposed that the slow consolidation of memories serves an adaptive function by allowing endogenous processes triggered by an experience to influence memory strength (Gold and McGaugh, 1975). Research indicates that adrenal stress hormones, such as epinephrine and cortisol, released during emotional arousal, play a crucial role in modulating memory strength based on the significance of the experience, with the amygdala mediating the effects of these hormones on memory consolidation (Paré and Headley, 2023; Roozendaal et al., 2007; Mcgaugh and Roozendaal, 2002). These findings suggest that experiences with behavioral significance are more likely to be consolidated due to their activation of the emotional arousal system. However, the specificity issue remains unresolved, as emotional arousal modulates memory over a broad time scale. Various stimulant drugs that act on epinephrine and cortisol pathways enhance memory consolidation when administered within minutes or even hours after training (Mcgaugh, 2000; McIntyre et al., 2012). Nonetheless, it remains unclear whether and how emotional arousal selectively impacts memories of the numerous sensory experiences preceding the arousal. This highlights the need for further research to uncover the mechanisms that govern the selection and prioritization of memories for consolidation.

In this article, we discuss the memory selection issue from a different perspective. This perspective is based on two relatively recent findings: that the hippocampus is involved not only in remembering the past but also in imagining the future, and that it encodes robust value information. We begin by summarizing evidence supporting the role of the hippocampus in imagination and then review evidence that the hippocampus, particularly the CA1 region, encodes robust value information. Next, we introduce the simulation-selection model, which we propose as a framework to explain the functions of the CA3-CA1 neural network. Finally, we draw parallels between our model and the Dyna reinforcement learning algorithm (Sutton, 1991), suggesting that hippocampal processes underlying memory consolidation may be conceptualized as a form of offline reinforcement learning—a mechanism for reinforcing valuable future strategies by recombining past experiences through simulation.

Hippocampus and imagination

The hippocampus, a critical brain structure traditionally recognized for its role in memory encoding and retrieval, is increasingly understood as essential for imagination and simulating future events. In humans, research has shown that the hippocampus is fundamental for generating detailed and coherent imagined scenarios. Patients with bilateral damage to the medial temporal lobes, which include the hippocampus, exhibit significant impairments in imagining hypothetical episodes (Hassabis et al., 2007). Furthermore, as a key component of the default mode network, the hippocampus is not only activated during the recall of autobiographical memories but also while envisioning future scenarios (Addis et al., 2007; Szpunar et al., 2007). These findings indicate the hippocampus is essential for synthesizing elements of memory into cohesive hypothetical episodes.

Animal studies complement and extend these findings by revealing the underlying neural mechanisms that allow the hippocampus to support imagination and predictive thought. Research on hippocampal replay, a phenomenon where patterns of neuronal activity representing past experiences are reactivated, has been particularly informative. In rats, hippocampal place cells go through rapid sequential discharges during periods of quiet rest and sleep, mirroring the order of activity observed during active navigation (Lee and Wilson, 2002; Foster and Wilson, 2006; Diba and Buzsáki, 2007). Initially, these replays were thought to support the recall and consolidation of prior navigation experiences. However, subsequent studies revealed that hippocampal replays also involve novel recombinations of previously learned trajectories (Gupta et al., 2010). Further research demonstrated that the hippocampus engages in preplay, where neuronal sequences representing paths in a novel environment are activated before the animal encounters them (Dragoi and Tonegawa, 2011). These findings suggest that the hippocampus constructs forward-looking models of the world, enabling prediction and preparation for future scenarios. Thus, findings from human and animal studies converge on the idea that the hippocampus is not merely a repository for memories but a flexible, predictive system

capable of constructing mental representations of the past, present, and future.

Value representation in the hippocampus

Early efforts to identify value-related neural activity primarily focused on brain regions outside the hippocampus, such as the parietal cortex, frontal cortex, and basal ganglia (O'doherty, 2004; Lee et al., 2012a; Glimcher, 2014). However, an early human imaging study detected value-related BOLD signals in the hippocampus alongside well-established value-related regions like the orbitofrontal cortex and striatum (Tanaka et al., 2004). In rats, value-related neuronal activity was identified in the CA1 region of the hippocampus (Lee et al., 2012b). The strength and characteristics of these CA1 value signals were comparable to those observed in traditional value-related regions such as the orbitofrontal cortex and striatum (Shin et al., 2021). Moreover, CA1 value signals temporally overlapped with choice and reward signals, indicating that CA1 integrates the necessary information for computing reward prediction errors and updating reward values (Lee et al., 2012b).

Additional studies have corroborated these findings across species. Human imaging studies further showed value-related BOLD signals in the hippocampus (Bornstein and Daw, 2013; Dombrovski et al., 2020), and physiological recordings in monkeys demonstrated value-dependent neuronal activity in this region (Knudsen and Wallis, 2021). In mice, calcium imaging studies revealed robust value signals in both dorsal and ventral CA1, with population activity in dorsal CA1 neurons increasing as a function of value (Biane et al., 2023; Yun et al., 2023). Collectively, these findings establish that value-related hippocampal neural processes are conserved across species, including rats, mice, monkeys, and humans.

Comparisons along the hippocampal transverse axis revealed that CA1 exhibits significantly stronger value signals than CA3 or the subiculum, its main input and output structures (Lee et al., 2012b; Lee et al., 2017). This observation aligns with evidence that CA1 neurons, unlike CA3, remap their place fields in response to changes in reward locations (Dupret et al., 2010). Furthermore, chemogenetic inactivation of CA1, but not CA3, impaired value learning without affecting value-dependent action selection, indicating CA1's critical role in valuation (Jeong et al., 2018). These findings suggest that value processing sets CA1 apart from other hippocampal subregions, establishing it as a key region for integrating valuation with other hippocampal mnemonic processes.

Simulation-selection model

Building on discoveries indicating the hippocampus's role in imagination and the CA1 region's unique function in value processing, we propose a new framework for the CA3-CA1 neural network: the simulation-selection model (Jung et al., 2018). The model's core concept is straightforward—CA3 acts as a simulator, generating diverse activity patterns, while CA1 functions as a selector, prioritizing and reinforcing patterns associated with high value (Figure 1). This selective reinforcement ensures that neural representations of high-value events and actions are strengthened, making them more likely to influence future decisions in similar contexts.

Anatomical and physiological evidence strongly supports CA3's role in generating diverse activity patterns. A key anatomical distinction between CA3 and CA1 is the presence of extensive recurrent collateral projections in CA3. In rats, each CA3 pyramidal neuron receives approximately 12,000 Schaffer collateral synapses from other CA3 neurons, which comprise 75% of its excitatory inputs (Amaral et al., 1990). Because CA3 neurons are heavily interconnected, the activity of some neurons often triggers the activation of others (self-excitation), especially when inhibitory neuronal activity is low, such as during slow-wave sleep and quiet rest (Ranck, 1973). During these states, the hippocampus shows slow, irregular rhythmic activity interspersed with occasional sharp-wave ripples (SWRs)—synchronized neuronal discharges accompanied by 140–200 Hz oscillations (Buzsáki et al., 1992; Buzsáki and Vanderwolf, 1983; O'keefe, 1976). CA3 serves as the primary initiator of SWRs, during which the majority of hippocampal replays are observed (Buzsáki, 2015).

Several features of CA3's architecture and dynamics facilitate the generation of novel activity patterns during SWRs (Jung et al., 2018). The CA3 network is characterized by many weak, recurrent synapses rather than a few strong ones (Debanne et al., 1999; Miles and Wong, 1986), making it more effective for generating variable sequences compared to networks dominated by fewer strong connections. Also, recurrent collateral synapses in CA3 support symmetric, rather than asymmetric, spike timing-dependent plasticity within a relatively broad time window (~150 ms) (Mishra et al., 2016). This promotes extensive associations among CA3 neurons with overlapping place fields, regardless of navigation trajectory. These features are consistent with the view that CA3 operates as a simulator, capable of generating activity patterns related to both previously experienced and unexperienced events.

In contrast, CA1 lacks the strong recurrent projections found in CA3, possessing only weak, short, longitudinally directed connections (Yang et al., 2014). Consequently, CA1 does not independently generate SWR-associated replays but instead processes activity sequences received from CA3. What distinguishes CA1 is its robust encoding of value, enabling it to process CA3-generated activity patterns differently based on their associated values. Although the value dependence of CA1 activity during SWRs remains incompletely characterized, existing findings strongly support this hypothesis. For example, CA1 place cells with firing fields near rewarding locations are preferentially reactivated during SWRs, whereas CA3 place cells do not exhibit such reward dependence (Dupret et al., 2010). Furthermore, CA1 replay preferentially encodes trajectories leading to reward locations (Foster and Wilson, 2006; Gupta et al., 2010; Ólafsdóttir et al., 2015; Pfeiffer and Foster, 2013). Reward also enhances the rate and fidelity of awake replays in CA1 (Ambrose et al., 2016; Bhattarai et al., 2020), which facilitates the consolidation of memories associated with these replays (Yang et al., 2024). In humans, rewards have been shown to enhance the imagination of episodic future events (Bulganin and Wittmann, 2015) and to preferentially reactivate high-reward contexts during post-learning rest, improving memory retention (Gruber et al., 2016; Sterpenich et al., 2021).

Collectively, these findings support the core premise of the simulation-selection model: CA3 generates diverse activity patterns, while CA1 selectively reinforces those associated with high value. The functional outcome of this interplay is the prioritization of high-value activity patterns, strengthening their neural representations and enhancing the likelihood of optimal future decision-making. The simulation-selection model generates numerous testable predictions.
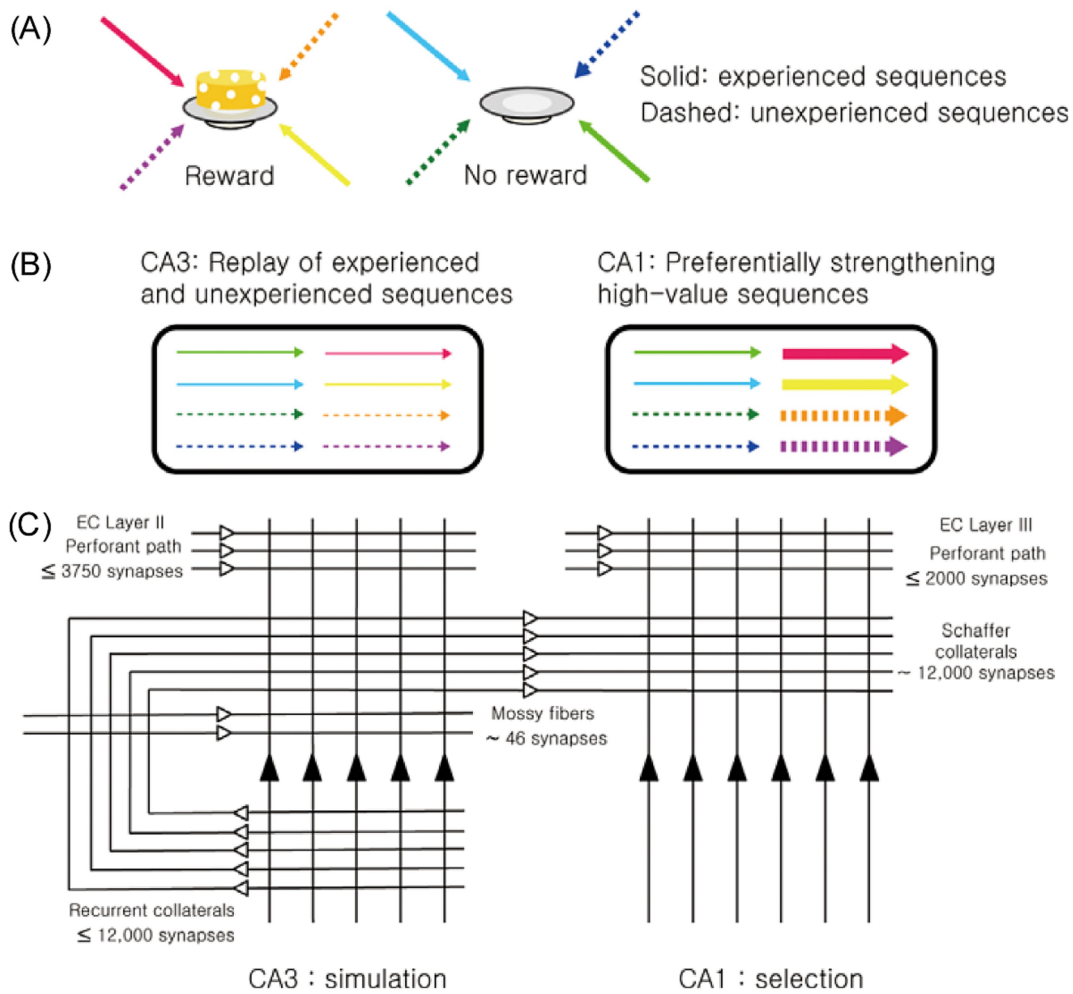
FIGURE 1
Overview of the simulation-selection model. **(A)** Navigation sequences to two locations—one where a reward was obtained (high-value sequence) and one where it was not (low-value sequence)—are represented using different colors. Solid arrows indicate experienced sequences, while dashed arrows represent unexperienced (novel) sequences. **(B)** CA3 generates both experienced and novel (unexperienced) navigation sequences, independent of their value. Among these, CA1 selectively reinforces high-value sequences, whether experienced or novel. **(C)** The schematic diagram illustrates the basic circuit organization of CA3 and CA1. The numbers denote the average number of synapses for each projection pathway in a single CA3 or CA1 pyramidal neuron (Amaral et al., 1990). The extensive but individually weak recurrent collaterals in CA3 enable the generation of both remembered (experienced) and novel (unexperienced) sequences. In contrast, CA1, which lacks recurrent collateral projections but conveys strong value signals, selectively reinforces high-value sequences. Figure adapted from Jung et al., 2018, licensed under CC-BY 4.0.

For instance, it predicts that CA1 replays will be more value-dependent than CA3 replays and that blocking CA3–CA1 synaptic plasticity during exploration will diminish the value dependence of CA1 replays (Jung et al., 2018). Additionally, the model suggests differential effects of CA3 and CA1 modulation on the diversity and value dependence of hippocampal replays, although the serial organization of the CA3–CA1 circuit poses challenges for interpreting such results. While further research is needed to test these predictions and address unresolved questions, such as differentiating the processing of positive and negative values, current findings align with this model and suggest its potential relevance in adaptive decision-making.

Offline reinforcement learning

Reinforcement learning is a branch of artificial intelligence focused on discovering optimal action strategies in dynamic and uncertain environments. A central concept in reinforcement learning is the value function, which estimates the expected cumulative reward an agent can achieve from a specific state (or state-action pair) by following a particular strategy (policy). Agents use value functions to select actions and continuously update these functions based on the outcomes of their decisions. Through this iterative process, agents approximate true value functions and adapt their choices accordingly (Sutton and Barto, 1998). However, this trial-and-error approach can be highly inefficient, often requiring a vast number of trials to converge to accurate value estimates. This inefficiency becomes especially problematic in scenarios where achieving a goal involves long action sequences or where the environment changes rapidly.

To address these challenges, the Dyna algorithm, introduced by Richard Sutton in 1991, provides an integrated approach that significantly enhances learning efficiency. Dyna combines direct interactions with the environment and simulated experiences to accelerate the reinforcement learning process. During direct interaction, the agent collects data by exploring the environment and updating its value functions and policies based on observed outcomes.

In parallel, the agent builds an internal model of the environment, capturing relationships between states, actions, and rewards. Using this model, the agent performs offline simulations to generate additional experiences, which are then used to further refine value functions and policies (Sutton and Barto, 1998; Sutton, 1991). This dual approach leverages the strengths of both model-free and model-based learning. The use of direct interaction ensures robustness and adaptability to environmental variations, while simulations enable faster learning and more efficient exploration of the state space by allowing the agent to explore hypothetical scenarios without real-world trials.

The similarity between the simulation-selection model and the Dyna algorithm is striking. The Dyna algorithm overcomes the inefficiencies of pure trial-and-error learning by combining real-world experiences with simulated planning, enabling agents to learn and adapt efficiently in complex and dynamic environments. Similarly, the simulation-selection model allows agents to navigate their environments more efficiently by complementing actual experiences with the simulation and evaluation of diverse scenarios, such as spatial trajectories, during idle states. Both frameworks accelerate value learning by leveraging simulations derived from limited experiences. In this analogy, the key components of the Dyna algorithm map seamlessly onto the simulation-selection model: the recurrent network architecture of CA3 supports the generation of simulated trajectories, functioning analogously to model-based planning, while the strong value signals in CA1 facilitate the selection and reinforcement of these simulations, akin to refining value functions and policies in Dyna.

At the outset, we noted that the ultimate goal of memory consolidation remains unclear. The simulation-selection theory offers a novel perspective, suggesting that memory consolidation is not merely the transformation or reorganization of temporary experiences stored in the hippocampus into long-term memory. Instead, it may function as a process of finding optimal strategies for navigating an environment through simulation, using limited experiences as a foundation. In this process, high-value behavioral strategies are selectively reinforced, facilitating better decision-making in the future. From this viewpoint, memory consolidation can be understood as a form of offline reinforcement learning. This idea also aligns with the constructive episodic simulation hypothesis, which posits that individuals flexibly extract and recombine elements of past experiences to simulate potential future scenarios (Schacter and Addis, 2007). Such flexibility allows past information to be effectively repurposed for simulating alternative future possibilities, reducing the reliance on actual trial-and-error behavior.

Conclusion

Recent studies increasingly highlight the hippocampus's role in predictive coding, emphasizing its capacity to prepare the brain for future scenarios (Barron et al., 2020; Buckner, 2010; Pezzulo et al., 2014). Theoretical advancements further support this view, with reinforcement learning frameworks emerging as powerful tools to explain hippocampal functions (Ambrogioni and Ólafsdóttir, 2023; Gershman and Daw, 2017; Stachenfeld et al., 2017; Tessereau et al., 2021). Integrating the simulation-selection model into this framework, we propose that hippocampal neural processes underlying memory consolidation might be understood as a form of offline reinforcement learning. From this perspective, memory consolidation is not a passive

process of fortifying memories based on initial encoding strength or arousal level but an active process of selecting and reinforcing valuable options for the future by recombining past experiences through imagination (Cowan et al., 2021; Jung, 2023; Jung et al., 2018).

This perspective does not negate other proposed roles of memory consolidation, such as schema formation, emotional memory modulation, or semantic abstraction. Rather, it broadens current thinking on memory consolidation by positioning reinforcement learning as a valuable theoretical framework for understanding hippocampal processes. It emphasizes the hippocampus's dual role in retaining past experiences and actively transforming them into actionable strategies for navigating future challenges. By bridging neural, behavioral, and computational frameworks, this approach provides new insights into the mechanisms and biological functions of hippocampal memory consolidation.

Data availability statement

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding authors.

Author contributions

JL: Writing – original draft, Writing – review & editing. MJ: Writing – original draft, Writing – review & editing.

Funding

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declare that no Gen AI was used in the creation of this manuscript.

Publisher's note

References

Addis, D. R., Wong, A. T., and Schacter, D. L. (2007). Remembering the past and imagining the future: common and distinct neural substrates during event construction and elaboration. *Neuropsychologia* 45, 1363–1377. doi: 10.1016/j.neuropsychologia.2006.10.016

Amaral, D. G., Ishizuka, N., and Claiborne, B. (1990). Chapter neurons, numbers and the hippocampal network. *Prog. Brain Res.* 83, 1–11. doi: 10.1016/S0079-6123(08)61237-6

Ambrogioni, L., and Ólafsdóttir, H. F. (2023). Rethinking the hippocampal cognitive map as a meta-learning computational module. *Trends Cogn. Sci.* 27, 702–712. doi: 10.1016/j.tics.2023.05.011

Ambrose, R. E., Pfeiffer, B. E., and Foster, D. J. (2016). Reverse replay of hippocampal place cells is uniquely modulated by changing reward. *Neuron* 91, 1124–1136. doi: 10.1016/j.neuron.2016.07.047

Barron, H. C., Auksztulewicz, R., and Friston, K. (2020). Prediction and memory: a predictive coding account. *Prog. Neurobiol.* 192:101821. doi: 10.1016/j.pneurobiol.2020.101821

Bhattarai, B., Lee, J. W., and Jung, M. W. (2020). Distinct effects of reward and navigation history on hippocampal forward and reverse replays. *Proc. Natl. Acad. Sci.* 117, 689–697. doi: 10.1073/pnas.1912533117

Biane, J. S., Ladow, M. A., Stefanini, F., Boddu, S. P., Fan, A., Hassan, S., et al. (2023). Neural dynamics underlying associative learning in the dorsal and ventral hippocampus. *Nat. Neurosci.* 26, 798–809. doi: 10.1038/s41593-023-01296-6

Bornstein, A. M., and Daw, N. D. (2013). Cortical and hippocampal correlates of deliberation during model-based decisions for rewards in humans. *PLoS Comput. Biol.* 9:e1003387. doi: 10.1371/journal.pcbi.1003387

Buckner, R. L. (2010). The role of the hippocampus in prediction and imagination. *Annu. Rev. Psychol.* 61, 27–48. doi: 10.1146/annurev.psych.60.110707.163508

Bulganin, L., and Wittmann, B. C. (2015). Reward and novelty enhance imagination of future events in a motivational-episodic network. *PLoS One* 10:e0143477. doi: 10.1371/journal.pone.0143477

Buzsáki, G. (2015). Hippocampal sharp wave-ripple: a cognitive biomarker for episodic memory and planning. *Hippocampus* 25, 1073–1188. doi: 10.1002/hipo.22488

Buzsáki, G., Horvath, Z., Urioste, R., Hetke, J., and Wise, K. (1992). High-frequency network oscillation in the hippocampus. *Science* 256, 1025–1027. doi: 10.1126/science.1589772

Buzsáki, G., and Vanderwolf, C. H. (1983). Cellular bases of hippocampal Eeg in the behaving rat. *Brain Res. Rev.* 6, 139–171. doi: 10.1016/0165-0173(83)90037-1

Cheng, S. (2017). Consolidation of episodic memory: an epiphenomenon of semantic learning. In N. Axmacher and B. Rasch (eds), *Cogn. Neurosci. Memory consolid.*, 57–72. Cham, Switzerland: Springer International Publishing. doi: 10.1007/978-3-319-45066-7_4

Cowan, E. T., Schapiro, A. C., Dunsmoor, J. E., and Murty, V. P. (2021). Memory consolidation as an adaptive process. *Psychon. Bull. Rev.* 28, 1796–1810. doi: 10.3758/s13423-021-01978-x

Debanne, D., Gähwiler, B. H., and Thompson, S. M. (1999). Heterogeneity of synaptic plasticity at unitary Ca3–Ca1 and Ca3–Ca3 connections in rat hippocampal slice cultures. *J. Neurosci.* 19, 10664–10671. doi: 10.1523/JNEUROSCI.19-24-10664.1999

Diba, K., and Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nat. Neurosci.* 10, 1241–1242. doi: 10.1038/nn1961

Dombrovski, A. Y., Luna, B., and Hallquist, M. N. (2020). Differential reinforcement encoding along the hippocampal long axis helps resolve the explore–exploit dilemma. *Nat. Commun.* 11:5407. doi: 10.1038/s41467-020-18864-0

Dragoi, G., and Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature* 469, 397–401. doi: 10.1038/nature09633

Dupret, D., O'neill, J., Pleydell-Bouverie, B., and Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nat. Neurosci.* 13, 995–1002. doi: 10.1038/nn.2599

Foster, D. J., and Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* 440, 680–683. doi: 10.1038/nature04587

Gershman, S. J., and Daw, N. D. (2017). Reinforcement learning and episodic memory in humans and animals: an integrative framework. *Annu. Rev. Psychol.* 68, 101–128. doi: 10.1146/annurev-psych-122414-033625

Glimcher, P. W. (2014). Value-based decision making. In P. W. Axmacher and E. Fehr (eds), *Neuroeconomics: Decision making and the brain*, (2nd ed., 373–391). London: Academic Press. doi: 10.1016/B978-0-12-416008-8.00020-6

Gold, P., and McGaugh, J. L. (1975). A single-trace, two-process view of memory strage processes. In D. Deutsch and J. A. Deutsch (eds), *Short-term Memory*, 355–378. New York: Academic Press.

Gruber, M. J., Ritchey, M., Wang, S.-F., Doss, M. K., and Ranganath, C. (2016). Post-learning hippocampal dynamics promote preferential retention of rewarding events. *Neuron* 89, 1110–1120. doi: 10.1016/j.neuron.2016.01.017

Gupta, A. S., Van Der Meer, M. A., Touretzky, D. S., and Redish, A. D. (2010). Hippocampal replay is not a simple function of experience. *Neuron* 65, 695–705. doi: 10.1016/j.neuron.2010.01.034

Hassabis, D., Kumaran, D., Vann, S. D., and Maguire, E. A. (2007). Patients with hippocampal amnesia cannot imagine new experiences. *Proc. Natl. Acad. Sci.* 104, 1726–1731. doi: 10.1073/pnas.0610561104

Jeong, Y., Huh, N., Lee, J., Yun, I., Lee, J. W., Lee, I., et al. (2018). Role of the hippocampal Ca1 region in incremental value learning. *Sci. Rep.* 8:9870. doi: 10.1038/s41598-018-28176-5

Jung, M. W. (2023). A brain for innovation: The neuroscience of imagination and abstract thinking. New York: Columbia University Press.

Jung, M. W., Lee, H., Jeong, Y., Lee, J. W., and Lee, I. (2018). Remembering rewarding futures: a simulation-selection model of the hippocampus. *Hippocampus* 28, 913–930. doi: 10.1002/hipo.23023

Knudsen, E. B., and Wallis, J. D. (2021). Hippocampal neurons construct a map of an abstract value space. *Cell* 184:e10.

Lechner, H. A., Squire, L. R., and Byrne, J. H. (1999). 100 years of consolidation—remembering Müller and Pilzecker. *Learn. Mem.* 6, 77–87. doi: 10.1101/lm.6.2.77

Lee, H., Ghim, J.-W., Kim, H., Lee, D., and Jung, M. (2012b). Hippocampal neural correlates for values of experienced events. *J. Neurosci.* 32, 15053–15065. doi: 10.1523/JNEUROSCI.2806-12.2012

Lee, S.-H., Huh, N., Lee, J. W., Ghim, J.-W., Lee, I., and Jung, M. W. (2017). Neural signals related to outcome evaluation are stronger in Ca1 than Ca3. *Front Neural Circuits* 11:40. doi: 10.3389/fncir.2017.00040

Lee, D., Seo, H., and Jung, M. W. (2012a). Neural basis of reinforcement learning and decision making. *Annu. Rev. Neurosci.* 35, 287–308. doi: 10.1146/annurev-neuro-062111-150512

Lee, A. K., and Wilson, M. A. (2002). Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* 36, 1183–1194. doi: 10.1016/S0896-6273(02)01096-6

Mcclelland, J. L., Mcnaughton, B. L., and O'reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol. Rev.* 102, 419–457. doi: 10.1037/0033-295X.102.3.419

Mcgaugh, J. L. (2000). Memory--a century of consolidation. *Science* 287, 248–251. doi: 10.1126/science.287.5451.248

Mcgaugh, J. L., and Roozendaal, B. (2002). Role of adrenal stress hormones in forming lasting memories in the brain. *Curr. Opin. Neurobiol.* 12, 205–210. doi: 10.1016/S0959-4388(02)00306-9

Mcintyre, C. K., Mcgaugh, J. L., and Williams, C. L. (2012). Interacting brain systems modulate memory consolidation. *Neurosci. Biobehav. Rev.* 36, 1750–1762. doi: 10.1016/j.neubiorev.2011.11.001

Miles, R., and Wong, R. (1986). Excitatory synaptic interactions between Ca3 neurones in the guinea-pig hippocampus. *J. Physiol.* 373, 397–418. doi: 10.1113/jphysiol.1986.sp016055

Mishra, R., Kim, S., Guzman, S., and Jonas, P. (2016). Symmetric spike timing-dependent plasticity at Ca3-Ca3 synapses optimizes storage and recall in autoassociative networks. *Nat. Commun.* 7:11552. doi: 10.1038/ncomms11552

Moscovitch, M., Cabeza, R., Winocur, G., and Nadel, L. (2016). Episodic memory and beyond: the hippocampus and neocortex in transformation. *Annu. Rev. Psychol.* 67, 105–134. doi: 10.1146/annurev-psych-113011-143733

Moscovitch, M., and Gilboa, A. (2024). Systems consolidation, transformation and reorganization: Multiple trace theory, trace transformation theory and their competitors. Oxford: Oxford Academic. doi: 10.1093/oxfordhb/9780190917982.013.43

Moscovitch, M., Rosenbaum, R. S., Gilboa, A., Addis, D. R., Westmacott, R., Grady, C., et al. (2005). Functional neuroanatomy of remote episodic, semantic and spatial memory: a unified account based on multiple trace theory. *J. Anat.* 207, 35–66. doi: 10.1111/j.1469-7580.2005.00421.x

Müller, G. E., and Pilzecker, A. (1900). Experimentelle beiträge zur lehre vom gedächtniss. Leipzig, Germany: Ja Barth.

Nadel, L., and Moscovitch, M. (1997). Memory consolidation, retrograde amnesia and the hippocampal complex. *Curr. Opin. Neurobiol.* 7, 217–227. doi: 10.1016/S0959-4388(97)80010-4

O'doherty, J. P. (2004). Reward representations and reward-related learning in the human brain: insights from neuroimaging. *Curr. Opin. Neurobiol.* 14, 769–776. doi: 10.1016/j.conb.2004.10.016

O'keefe, J. (1976). Place units in the hippocampus of the freely moving rat. *Exp. Neurol.* 51, 78–109. doi: 10.1016/0014-4886(76)90055-8

Ólafsdóttir, H. F., Barry, C., Saleem, A. B., Hassabis, D., and Spiers, H. J. (2015). Hippocampal place cells construct reward related sequences through unexplored space. *eLife* 4:e06063. doi: 10.7554/eLife.06063

Paré, D., and Headley, D. B. (2023). The amygdala mediates the facilitating influence of emotions on memory through multiple interacting mechanisms. *Neurobiol. Stress* 24:100529. doi: 10.1016/j.ynstr.2023.100529

Pezzulo, G., Van Der Meer, M. A., Lansink, C. S., and Pennartz, C. M. (2014). Internally generated sequences in learning and executing goal-directed behavior. *Trends Cogn. Sci.* 18, 647–657. doi: 10.1016/j.tics.2014.06.011

Pfeiffer, B. E., and Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* 497, 74–79. doi: 10.1038/nature12112

Ranck, J. B. Jr. (1973). Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats: part I. Behavioral correlates and firing repertoires. *Exp. Neurol.* 41, 462–531. doi: 10.1016/0014-4886(73)90290-2

Roozendaal, B., Barsegyan, A., and Lee, S. (2007). Adrenal stress hormones, amygdala activation, and memory for emotionally arousing experiences. *Prog. Brain Res.* 167, 79–97. doi: 10.1016/S0079-6123(07)67006-X

Schacter, D. L., and Addis, D. R. (2007). The cognitive neuroscience of constructive memory: remembering the past and imagining the future. *Phil. Trans. Royal Society B* 362, 773–786. doi: 10.1098/rstb.2007.2087

Sekeres, M. J., Winocur, G., and Moscovitch, M. (2018). The hippocampus and related neocortical structures in memory transformation. *Neurosci. Lett.* 680, 39–53. doi: 10.1016/j.neulet.2018.05.006

Shin, E. J., Jang, Y., Kim, S., Kim, H., Cai, X., Lee, H., et al. (2021). Robust and distributed neural representation of action values. *eLife* 10:e53045. doi: 10.7554/eLife.53045

Squire, L. R., and Alvarez, P. (1995). Retrograde amnesia and memory consolidation: a neurobiological perspective. *Curr. Opin. Neurobiol.* 5, 169–177. doi: 10.1016/0959-4388(95)80023-9

Stachenfeld, K. L., Botvinick, M. M., and Gershman, S. J. (2017). The hippocampus as a predictive map. *Nat. Neurosci.* 20, 1643–1653. doi: 10.1038/nn.4650

Sterpenich, V., Van Schie, M. K., Catsiyannis, M., Ramyead, A., Perrig, S., Yang, H.-D., et al. (2021). Reward biases spontaneous neural reactivation during sleep. *Nat. Commun.* 12:4162. doi: 10.1038/s41467-021-24357-5

Sutherland, R., and Lehmann, H. (2011). Alternative conceptions of memory consolidation and the role of the hippocampus at the systems level in rodents. *Curr. Opin. Neurobiol.* 21, 446–451. doi: 10.1016/j.conb.2011.04.007

Sutherland, R. J., Sparks, F. T., and Lehmann, H. (2010). Hippocampus and retrograde amnesia in the rat model: a modest proposal for the situation of systems consolidation. *Neuropsychologia* 48, 2357–2369. doi: 10.1016/j.neuropsychologia.2010.04.015

Sutton, R. S. (1991). Dyna, an integrated architecture for learning, planning, and reacting. *ACM SIGART Bull.* 2, 160–163. doi: 10.1145/122344.122377

Sutton, R., and Barto, A. (1998). Reinforcement learning. Cambridge, MA: MIT Press.

Szpunar, K. K., Watson, J. M., and Mcdermott, K. B. (2007). Neural substrates of envisioning the future. *Proc. Natl. Acad. Sci.* 104, 642–647. doi: 10.1073/pnas.0610082104

Tanaka, S. C., Doya, K., Okada, G., Ueda, K., Okamoto, Y., and Yamawaki, S. (2004). Prediction of immediate and future rewards differentially recruits cortico-basal ganglia loops. *Nat. Neurosci.* 7, 887–893. doi: 10.1038/nn1279

Tessereau, C., O'dea, R., Coombes, S., and Bast, T. (2021). Reinforcement learning approaches to hippocampus-dependent flexible spatial navigation. *Brain Neurosci. Adv.* 5:2398212820975634. doi: 10.1177/2398212820975634

Winocur, G., and Moscovitch, M. (2011). Memory transformation and systems consolidation. *J. Int. Neuropsychol. Soc.* 17, 766–780. doi: 10.1017/S1355617711000683

Winocur, G., Moscovitch, M., and Bontempi, B. (2010). Memory formation and long-term retention in humans and animals: convergence towards a transformation account of hippocampal–neocortical interactions. *Neuropsychologia* 48, 2339–2356. doi: 10.1016/j.neuropsychologia.2010.04.016

Winocur, G., Moscovitch, M., and Sekeres, M. (2007). Memory consolidation or transformation: context manipulation and hippocampal representations of memory. *Nat. Neurosci.* 10, 555–557. doi: 10.1038/nn1880

Yang, W., Sun, C., Huszár, R., Hainmueller, T., Kiselev, K., and Buzsáki, G. (2024). Selection of experience for memory by hippocampal sharp wave ripples. *Science* 383, 1478–1483. doi: 10.1126/science.adk8261

Yang, S., Yang, S., Moreira, T., Hoffman, G., Carlson, G. C., Bender, K. J., et al. (2014). Interlamellar Ca1 network in the hippocampus. *Proc. Natl. Acad. Sci.* 111, 12919–12924. doi: 10.1073/pnas.1405468111

Yun, M., Hwang, J. Y., and Jung, M. W. (2023). Septotemporal variations in hippocampal value and outcome processing. *Cell Rep.* 42:112094. doi: 10.1016/j.celrep.2023.112094