Replay in minds and machines

–Revised manuscript–

–Revision 1–

Lennart Wittkuhn^{1,2,*}, Samson Chien^{1,2}, Sam Hall-McMaster^{1,2} & Nicolas W. Schuck^{1,2,*}

¹Max Planck Research Group NeuroCode, Max Planck Institute for Human Development

²Max Planck UCL Centre for Computational Psychiatry and Ageing Research

Lentzeallee 94, D-14195 Berlin, Germany

*Corresponding authors: wittkuhn@mpib-berlin.mpg.de, schuck@mpib-berlin.mpg.de

Abstract

Experience-related brain activity patterns reactivate during sleep, wakeful rest, and brief pauses from active behavior. In parallel, machine learning research has found that experience replay can lead to substantial performance improvements in artificial agents. Together, these lines of research suggest replay has a variety of computational benefits for decision-making and learning. Here, we provide an overview of putative computational functions of replay as suggested by machine learning and neuroscientific research. We show that replay can lead to faster learning, less forgetting, reorganization or augmentation of experiences, and support planning and generalization. In addition, we highlight the benefits of reactivating abstracted internal representations rather than veridical memories, and discuss how replay could provide a mechanism to build internal representations that improve learning and decision-making.

Keywords: replay; reinforcement learning; machine learning; representation learning; decision-making

Memory, planning and imagination are important aspects of intelligent behavior. They allow the mind to go beyond merely observing and reacting to its surroundings. How the brain implements these functions, and how they could help to improve artificial intelligent agents, is not yet fully understood. In this review, we will provide an overview of one important candidate mechanism involved in memory, imagination and planning: replay. The term replay is used to refer to a wide variety of mechanisms that relate to the reactivation of past memories. This reactivation has been observed in the brain and is commonly implemented in artificial agents. Replay occurs often, but not always, in sequential form, and mostly while the agent or animal is not interacting with its environment. This article asks *why* the brain and artificial agents might use replay.

1 Replay in the brain

Before we discuss the benefits of replay, we will give a brief historical overview of the subject from a neuroscientific point of view. In neuroscience, much research on memory, planning and imagination has focused on the hippocampus (e.g., [Squire, 1992](#); [Buckner, 2010](#)). Early indications that the hippocampus gives rise to memory functioning came from studies of lesion patients ([Scoville and Milner, 1957](#)), and studies of rodent spatial navigation ([O'Keefe and Nadel, 1978](#)). Of particular importance were rodent recordings from hippocampal pyramidal neurons, known as *place cells*, that demonstrated spatial firing selectivity when animals navigated in a spatial environment ([O'Keefe and Dostrovsky, 1971](#); [O'Keefe and Conway, 1978](#)). These place cells have since been regarded as a core neural substrate for a *cognitive map* ([Tolman, 1948](#)) of physical space that supports spatial navigation ([O'Keefe and Nadel, 1974, 1978](#); [Moser et al., 2008](#)), as well as memory ([Cohen and Eichenbaum, 1993](#); [Redish and Touretzky, 1998](#)).

It soon became clear that hippocampal place cells are also active when an animal is not engaged in a particular task, in line with theoretical proposals that a *reactivation* mechanism could support consolidation of recent memory traces into an aggregated memory store ([Marr, 1971](#)). Following early empirical support for this idea ([Buzsáki, 1989](#); [Pavlides and Winson, 1989](#)), multi-unit recordings in rodents led to the discovery of *replay* – the finding that during periods of rest and sleep, hippocampal cells reactivate sequentially in fast bursts, as if retracing paths the

54 animal had taken during wakefulness ([Wilson and McNaughton, 1994](#); [Skaggs and](#)
55 [McNaughton, 1996](#); [Kudrimoti et al., 1999](#); [Nádasdy et al., 1999](#); [Gerrard et al., 2001](#);
56 [Lee and Wilson, 2002](#); [Louie and Wilson, 2001](#), for reviews of these earlier findings,
57 see [Redish, 1999](#); [Sutherland and McNaughton, 2000](#)). These observations were fol-
58 lowed by a wealth of findings that established the now classic neuroscientific view of
59 replay: replay is sequential, occurs during sleep or rest, reflects previous experience
60 in spatial navigation tasks, and happens on a temporally compressed timescale (for
61 review, see e.g., [Foster, 2017](#)).

62 Over the following decades, much more became known about the biological as-
63 pects of replay, and many findings supported the idea that replay is important for
64 memory. First, replay is commonly detected during brief, high-frequency oscillations
65 called sharp wave-ripples (SWRs) (for review, see e.g., [Buzsáki, 2015](#); [Joo and Frank,](#)
66 [2018](#)), which have also been found in human medial temporal lobe (MTL) ([Bragin](#)
67 [et al., 1999](#); [Staba et al., 2002](#)) and can be linked to memory consolidation during
68 sleep, rest, and awake episodic memory retrieval ([Axmacher et al., 2008](#); [Staresina](#)
69 [et al., 2015](#); [Zhang et al., 2018](#); [Helfrich et al., 2019](#); [Norman et al., 2019](#); [Vaz et al.,](#)
70 [2019, 2020](#)). A link to memory is also supported by findings showing that the se-
71 lective disruption of SWRs during post-task rest slows learning in hippocampus-
72 dependent spatial memory tasks ([Girardeau et al., 2009](#); [Ego-Stengel and Wilson,](#)
73 [2010](#); [Jadhav et al., 2012](#)), and that memory can be influenced by playing sounds
74 during SWR events while an animal sleeps ([Bendor and Wilson, 2012](#); [Rothschild](#)
75 [et al., 2016](#)). Second, replay is much faster than wakeful experience, and this tem-
76 poral compression is believed to induce the conditions that drive learning and the
77 strengthening of memory traces through synaptic plasticity ([Bliss and Collingridge,](#)
78 [1993](#); [Magee, 1997](#); [King et al., 1999](#)). Third, interactions between the hippocampus
79 and prefrontal cortex (PFC) during replay events support the idea of consolidating
80 reactivated memories in the brain (for reviews, see e.g., [Zielinski et al., 2018](#); [Tang](#)
81 [and Jadhav, 2019](#)).

82 Due to the fast and anatomically localized nature of the replay phenomenon,
83 these insights were almost exclusively gained from invasive recordings in rodents and
84 human patient populations. But existing studies focusing on non-invasive detection
85 of replay in humans point to similar conclusions. Memory benefits of non-sequential
86 reactivation during rest or sleep are well documented in humans ([Staresina et al.,](#)

2013; Deuker et al., 2013; Tambini and Davachi, 2013; Tambini et al., 2010; Gruber et al., 2016). Memory consolidation in humans can also be biased by presenting learning-associated sensory cues, a technique known as targeted memory reactivation (TMR), during replay-associated sleep phases in humans (Oudiette and Paller, 2013; Lewis and Bendor, 2019). Recent progress in neuroimaging analyses have also been able to capture the sequentiality of fast replay events using magnetoencephalography (MEG) (Kurth-Nelson et al., 2016; Liu et al., 2021a) and functional magnetic resonance imaging (fMRI) (Schuck and Niv, 2019; Wittkuhn and Schuck, 2021). In combination, these findings have demonstrated that replay exists in a variety of species and support the idea that it reflects a consolidation process that strengthens memory associations (for reviews, see Sutherland and McNaughton, 2000; Rasch and Born, 2007; O'Neill et al., 2010; Diekelmann and Born, 2010; Carr et al., 2011; Zhang et al., 2017; Tambini and Davachi, 2019).

While the above findings have established a foundational knowledge of replay, our understanding of this phenomenon has undergone significant and continued change that sometimes challenged the classic picture of replay (for review, see e.g., Foster, 2017). For instance, replay seems to be significantly more frequent than initially thought, happening not only during sleep or rest but also during brief wakeful pauses from active behavior (Csicsvari et al., 2007; Foster and Wilson, 2006; Diba and Buzsáki, 2007; Davidson et al., 2009; Karlsson and Frank, 2009, for reviews of awake replay, see e.g., Carr et al., 2011; Tambini and Davachi, 2019). Replay-like sequential reactivation patterns occur at various speeds, from highly accelerated to much slower behavioral timescales (Deng et al., 2020; Denovellis et al., 2020; Tang et al., 2021). Recordings outside of the hippocampus have identified replay-like phenomena in a large number of other brain areas, including entorhinal (Ólafsdóttir et al., 2016, 2017; O'Neill et al., 2017; Trettel et al., 2019), prefrontal (Euston et al., 2007; Peyrache et al., 2009; Jadhav et al., 2016; Yu et al., 2018; Shin et al., 2019; Kaefer et al., 2020; Tang et al., 2021), visual and auditory sensory cortices (Ji and Wilson, 2006; Rothschild et al., 2016; Wittkuhn and Schuck, 2021), parietal cortex (Qin et al., 1997; Hoffman and McNaughton, 2002; Harvey et al., 2012), motor cortex (Ramanathan et al., 2015; Gulati et al., 2017), and ventral striatum (Lansink et al., 2009, 2008; Pennartz, 2004; Gomperts et al., 2015). Moreover, replay is not necessarily a faithful replication of previous behavioral sequences but can also reverse

the order of experiences (Csicsvari et al., 2007; Foster and Wilson, 2006; Diba and Buzsáki, 2007; Davidson et al., 2009; Karlsson and Frank, 2009) or change the order of actual experiences according to a learned task rule (Liu et al., 2019). It can also represent remote, non-local and never-experienced locations (Karlsson and Frank, 2009; Gupta et al., 2010; Ólafsdóttir et al., 2015), reflect non-spatial and partially observable task features (Schuck and Niv, 2019), and occur even after tasks without explicit memory requirements (Wittkuhn and Schuck, 2021). Collectively, these results suggest that replay (1) occurs during a variety of behavioral states, including rest, sleep and pausing, (2) occurs on a variety of time scales, (3) occurs in a variety of brain areas, and (4) does not only reflect previous experience, but is involved in a much broader range of cognitive functions than memory consolidation and spatial navigation alone.

Indeed, our understanding of hippocampal place cells, and the neural architecture underlying memory and spatial navigation more generally, has also evolved considerably. The “places” represented by hippocampal neurons are not exclusively determined by location in physical space, but can also incorporate other task-relevant aspects, such as sounds (Aronov et al., 2017) and time (MacDonald et al., 2011), see also O’Keefe and Krupic (2021). Other studies have pointed out that the hippocampus may learn and predict transitions between states in the environment (Gaussier et al., 2002) or encode representations that are predictive of future locations, so called successor representations (SRs), that can be used for reinforcement learning (RL) (Stachenfeld et al., 2017), and that grid-like patterns in the entorhinal cortex and ventromedial PFC may represent coordinates of a non-spatial space (Constantinescu et al., 2016). Thus, today the cognitive map in the hippocampal-entorhinal system is often thought to represent relationships of locations and events beyond physical space, from conceptual knowledge to social cognition (for reviews and perspectives, see Khamassi and Humphries, 2012; Kaplan et al., 2017; Epstein et al., 2017; Schafer and Schiller, 2018; Behrens et al., 2018; Bellmund et al., 2018; Peer et al., 2020; Bottini and Doeller, 2020; Spiers, 2020). Map-like representations also exist beyond the hippocampus, most notably in the medial entorhinal cortex (Hafting et al., 2005; Fyhn et al., 2007; Høydal et al., 2019), and in prefrontal and orbitofrontal cortex (OFC) (Wilson et al., 2014; Schuck et al., 2016; Constantinescu et al., 2016, for a review, see e.g., Schuck et al., 2018). These findings may have important impli-

153 cations for our understanding of the nature of replayed representations, and suggest
154 a mechanism that is much broader than a mere recapitulation of past observations
155 in the hippocampus.

156 How can such a diverse set of findings about replay in the hippocampus and
157 the rest of the brain be integrated? We argue that insights into this question can
158 be gained by considering the machine learning (ML) literature, where “experience
159 replay” was introduced in the early 90s (Lin, 1991). More recently, experience re-
160 play has become particularly popular after its importance for training deep neural
161 networks (DNNs) to play Atari video games became clear (e.g., Mnih et al., 2013;
162 Mnih et al., 2015; Hessel et al., 2018). This led experience replay to rise to promi-
163 nence as a crucial ingredient in building human-level intelligence in artificial agents
164 (Kumaran et al., 2016). Despite the conceptual similarity of biological and artifi-
165 cial replay, research on this subject in neuroscience and ML has progressed largely
166 in parallel. Here, we aim to connect insights from both research fields and review
167 computational perspectives taken in ML on the replay phenomenon. Our goal is to
168 highlight the diversity of possible computational and cognitive functions that might
169 be served by a replay mechanism and attempt to answer the question of *why* agents
170 would replay in the first place. Figure 1 provides a non-exhaustive overview of dif-
171 ferent forms of replay, which differ in which experiences are selected for replay and
172 in how replayed information affects the subsequent behavior of an agent. We will
173 discuss these different instantiations of replay below.

174 2 Computational benefits of replay

175 Replay has become a highly studied aspect of artificial agents. But why do machines
176 need replay, and do animals and machines have the same reasons to employ this pro-
177 cess? In the following sections, we will compare the roles of replay in both biological
178 and artificial agents, and distill the most significant benefits of replay.

179 Before we begin, we would like to point out some significant aspects in which
180 the concept of replay differs between ML and neuroscience. First, neuroscience em-
181 phasizes the sequential and often accelerated nature of replay (Genzel et al., 2020).
182 In ML, in contrast, some methods focus on replaying sets of individual transitions
183 (e.g., Mnih et al., 2015), rather than sequences (but see e.g., Hausknecht and Stone,
184 2015). The issue of replay speed has not been a major consideration in ML, as

Figure 1: Diversity of replay. Each panel outlines one example of one possible form of replay. In each case, we show an agent (black dot / robot) that stores a policy π , a value function Q and a model M , see [Box 1](#). Depicted is also the closest goal / reward (grey square), the relevant episode (blue bar), whether the episode is internally transformed (blue striped bar) and which aspect of the agent is updated through replay (green arrow). **(A)** A case in which encountering a goal triggers reverse replay. Reverse replay is then used to update the agent’s value function, similar to [Lin \(1992\)](#). **(B)** Interleaved replay in which episodes from a previous task are replayed to prevent catastrophic forgetting, see [McClelland et al. \(1995\)](#). **(C)** Replay of uniformly / randomly selected individual transitions ([Mnih et al., 2015](#)). **(D)** An agent can also learn a model through online updating, and replay from the model during offline periods to update its value function ([Sutton, 1991](#)). **(E)** Episodes can be selected for replay based on the magnitude of prediction errors or other reward related signals experienced during task performance ([Schaul et al., 2015](#)). **(F)** Instead of replaying previously experienced episodes, an agent can simulate possible episodes based on its model and policy in order to update the agent’s value function (“offline policy evaluation”, [Sutton and Barto, 2018](#)), or to plan the next actions (the policy π) at a choice point, without updating values. **(G)** Previously experienced episodes can also be abstracted before they are replayed, as done for instance when internal representations instead of observations are reactivated ([Kapturowski et al., 2019](#)). This can be used to e.g., update the agent’s model and / or value function. **(H)** Agents can also insert imagined sub-goals into replayed episodes, in particular in order to leverage information from episodes in which the agent never reached the final goal. This is done in hindsight replay ([Andrychowicz et al., 2017](#)). **(I)** Replay occurring during sleep, i.e., while the agent is not engaging in any task at all. This is commonly observed in animals ([Klinzing et al., 2019](#)), but analogies from the ML literature are lacking because agents do not sleep. Note, that this figure is not meant to be complete and merely illustrates some but not all aspects of the referenced algorithms. © Wittkuhn et al., doi: [10.6084/m9.figshare.14261636.v1](https://doi.org/10.6084/m9.figshare.14261636.v1), CC-BY 4.0 license (<https://creativecommons.org/licenses/by/4.0/>).

185 artificial agents are not bound to physical interaction with the environment and
186 the timescales of biology. A second difference is that understanding the distinction
187 between sleep replay and replay during wakeful pauses from active behavior has a
188 prominence in neuroscience (and is covered extensively in previous reviews; see e.g.,
189 [Findlay et al., 2021](#); [Klinzing et al., 2019](#)) that is not equivalently mirrored in ML
190 research. While the contrast between sleep and wakefulness is a theme that has
191 inspired ML research conceptually (see e.g., [Hinton et al., 1995](#)), the mere fact that
192 artificial agents do not “sleep” in the way that biological agents do, makes it practi-
193 cally impossible to investigate those differences in artificial agents. While biological
194 agents have several “modes” in which they are “off-policy” (sleeping, resting, paus-
195 ing, mind-wandering, etc.), to our knowledge no comparable distinctions have been
196 made for artificial agents. Third, ML researchers often distinguish between experi-
197 ence replay, which corresponds to sampling experiences from a memory buffer, and
198 model-based methods, in which the agent internally generates new experiences from
199 a learned model of the environment. Model-based methods are not universally called
200 replay in the ML literature, and are often referred to as planning instead.

201 Similar to previous work (e.g., [Foster and Knierim, 2012](#); [Cazé et al., 2018](#); [Mo-](#)
202 [mennejad, 2020](#)), our review focuses on the frameworks of RL ([Sutton and Barto,](#)
203 [2018](#)) and neural networks. The formalism of RL allows parallels to be drawn be-
204 tween reactivation of neural patterns in biological agents and replay of task states
205 in artificial agents. The RL framework considers agents that learn from interactions
206 with their environment and thereby gather experiences one at a time. RL techniques
207 are designed to learn from experience gradually, through trial-and-error, using every
208 new experience immediately to adjust the agent’s knowledge about the task. This
209 has the benefit of accruing knowledge without delay, while integrating information
210 over all experiences gained so far, rather than using just the most recent experience
211 to make decisions. Typically, small adjustments are made to the agent’s knowledge
212 with each new experience because large updates risk overwriting the effects of ear-
213 lier learning and can limit generalization. [Box 1](#) describes the fundamental aspects
214 of RL. Next to RL, we will also draw on insights from (supervised) deep learning
215 (for overviews, see e.g., [LeCun et al., 2015](#); [McClelland and Botvinick, 2020](#)) and
216 the successful combination of the two approaches, deep RL ([Mnih et al., 2015](#), see
217 [Tesauro, 1995](#) for an earlier integration).

218 Which benefits can an agent obtain from using replay? In the next sections, we
219 will discuss five potential computational functions of replay: increasing speed and
220 data efficiency of learning, reducing forgetting, reorganizing experiences, planning,
221 and generalization. We do not consider these functions to be entirely separable. We
222 distinguish them because they each offer a unique perspective on what an intelligent
223 agent, biological or artificial, stands to gain from replaying past experiences. This
224 perspective also sheds light on why replay can have different properties in different
225 study contexts, which have found replay to be sometimes backward and sometimes
226 forward, or in some cases to occur immediately and in other cases long after the
227 experience was acquired.

228 In addition to the topics above, we will consider one underexplored aspect of
229 replay: whether replay reflects sensory memories, or past internal representations,
230 and whether replay may also be involved in *shaping* internal representations as well.
231 We hypothesize that the content and function of replay is determined by its interplay
232 with the agent’s current representation of the task and the representational demands
233 of the task at hand, a notion which has recently received some computational ([Russek
234 et al., 2017](#); [Caselles-Dupré et al., 2019](#); [Momennejad, 2020](#)) as well as empirical
235 support ([Schuck and Niv, 2019](#)). In this view, replay can be understood not only as a
236 phenomenon that retrieves relational information stored in a cognitive map, but also
237 as a process that changes relational information and internal state representations
238 of an agent (see [Box 1](#) for a definition of state representations).

239 Each of the sections will be organized as follows. First, we will state a computa-
240 tional problem that any learning agent will be faced with. Then we discuss how this
241 problem has been approached in ML using replay, highlighting both theoretical and
242 empirical results. Finally, we will discuss empirical findings from the neuroscience
243 literature that support a particular ML proposal, or suggest alternative mechanisms.

Box 1 What is reinforcement learning?

Reinforcement learning (RL) theory provides a formal framework to describe how agents learn to optimize their behavior through interactions with an environment that yields rewards or punishments (Sutton and Barto, 2018). The agent-environment interaction is modelled as a Markov decision process (MDP), which consists of (1) an **environment**, described by of a set of states S , (2) a set of **actions** A available to the agent, (3) a **state transition function** $M(s_t, a_t, s_{t+1})$, also called a model, reflecting the probabilities of moving from state s_t to the next state s_{t+1} after taking action a_t and (4) a **reward function** $R(s_t, a_t, s_{t+1})$ that maps each [state, action, next-state]-triplet to a scalar reinforcement signal r . MDPs can have continuous state or action spaces, although most applications consider finite and discrete cases.

In MDPs, the agent-environment interaction is assumed to be Markovian with respect to reward and state, which means that the state and reward at the next time point $t + 1$ depend only on the state and action of the current time point t , but not on any states or actions before. The current state s_t therefore contains all relevant information from the previous history to determine the next state after an action has been performed. In brief, this means that we can think of learning from trial-and-error as the following process: the agent represents the current state of the environment, s_t and then performs an action a_t . The action will affect the environment, changing the agent's state from s_t to s_{t+1} as described by the state transition function M , and potentially yield a reward, as described by the reward function R . The agent's goal is to always perform the actions that maximize return – the expected sum of total rewards over the course of its interaction with the environment. In value-based approaches, the agent learns values that estimate the return, and then implements a **policy** π that maximizes the values. One popular approach is to estimate values with so-called temporal difference (TD) learning, using a Q -learning algorithm (Watkins and Dayan, 1992):

$$Q_{k+1}(s_t, a_t) = Q_k(s_t, a_t) + \alpha \left[r_t + \gamma \max_{a_{t+1}} Q_k(s_{t+1}, a_{t+1}) - Q_k(s_t, a_t) \right] \quad (1)$$

where $\gamma \in [0, 1]$, the discount factor, attenuates the influence of distal rewards, and

273 $\alpha \in [0, 1]$ is a learning rate. Based on the Bellman equation, Q -learning estimates the
274 sum of future rewards in an iterative bootstrapping process that involves the current
275 and future Q -value. Notably, because the algorithm uses the value of the best action
276 on the next step, rather than the value of the action that was actually performed, it
277 is a so-called off-policy approach. Q -learning does not require a transition model of
278 the environment, and hence is a model-free method. One important question that is
279 not addressed by the framework of RL is what information about the environment
280 is encoded in an agent's internal states. It is important to realize that agents mostly
281 do not have a way to simply know the objective, "true" state of the environment,
282 but rather must infer that state from their observations. An agent's internal state
283 representation therefore may not be equivalent to the true state, which, as we will
284 discuss in section 2.6, has many repercussions. The agent could simply take its
285 sensory data to be the states, but this is not sufficient for many tasks; rather the
286 agent needs to supplement its internal state representations with non-observable
287 information (e.g., [Wilson et al., 2014](#); [Schuck et al., 2018](#)).

2.1 Faster learning and data efficiency

The gradual approach to learning in RL has many benefits, but it results in very slow learning that may need thousands of iterations to achieve the optimal policy. Even worse, the slowness of learning grows exponentially with the number of states in the task environment, a phenomenon known as the “curse of dimensionality” (Bellman, 1957). To be a feasible approach to learning in complex and changing environments, gradual methods must therefore be complemented by mechanisms that will speed up learning without sacrificing the benefits of immediate knowledge acquisition and stable long-term memory. In this light, the idea of recapitulating previous experiences seems particularly appealing for machines, because it is easy and cheap for artificial agents to relearn from past experience that is retrieved from a memory buffer.

The brain arguably faces a similar computational challenge. Humans, and other animals, often have to learn directly from the outcomes of their decisions. Yet, repeating errors can pose actual risks, which limits the usefulness of exclusively relying on a slow, trial-and-error-based learning mechanism. More generally, the number of experiences that is acquired with a particular situation in a lifetime is quite limited in relation to the complexity of the environments and the brain, which contains approximately 10^{14} synapses (Tang et al., 2001). In order to make thousands of gradual adjustments to each of these synapses, the ability to reuse experience efficiently is paramount. Replay might be one mechanism to do just that.

2.1.1 Replay can speed-up gradual learning from experience and support credit assignment

In the RL literature, “experience replay” was initially introduced to address the issues of slow learning and data inefficiency (Lin, 1991, 1992, 1993). In his seminal paper, Lin (1992) wrote that “[...] Q-learning algorithms [...] are inefficient in that experiences obtained by trial-and-error are utilized to adjust the networks only once and then thrown away. [...] Experiences should be reused in an effective way.” (Lin, 1992, p. 299). Lin (1992) proposed that experiences can be used to update knowledge in a dual fashion; (1) immediately when experiences are acquired, and (2) at later time points, after experience itself may have long passed. Specifically, Lin (1992) proposed replaying full sequences of experiences, starting from an initial state to a final state, in backward order, and learning from these experiences, as if

320 they were real. [Lin \(1992\)](#) then showed that this is a more efficient use of data that
321 accelerates learning of an RL agent. In line with these ideas, many others have since
322 emphasized the computational benefit of replay for maximizing data efficiency and
323 the speed of learning (for reviews, see e.g., [Kumaran et al., 2016](#); [Hassabis et al.,](#)
324 [2017](#)).

325 There are several reasons why replay can help learning. In the real world, out-
326 comes are often only obtained after a long sequence of events and actions but agents
327 still need to know how to behave at the start of the sequence, as for instance, in a
328 chess game. This problem is known in RL as the *temporal credit assignment prob-*
329 *lem* ([Minsky, 1961](#)) and replay may help to solve it. The early work by [Lin \(1991,](#)
330 [1992, 1993\)](#) pointed out that replay could help an agent to remember the sequence
331 of previous states and actions that led to a given outcome, and assign credit for the
332 reward to the sequence of states and actions that preceded it. This also explains
333 why sequential replay may proceed in backward order ([Lin, 1992](#)). Another aspect is
334 that as the agent’s knowledge of the rewards becomes better with time, outcomes in
335 the past should be re-evaluated in light of this updated knowledge ([van Seijen and](#)
336 [Sutton, 2015](#)). Replay could serve this function by retrieving past rewards which can
337 then be compared to current value estimates.

338 Several neuroscientific studies suggest an important role of replay in speeding up
339 learning in biological agents too. First, studies in rodents reported increases in SWR-
340 associated reactivation following initial learning in novel environments ([Cheng and](#)
341 [Frank, 2008](#); [Eschenko et al., 2008](#); [O’Neill et al., 2008](#); [van de Ven et al., 2016](#); [Tang](#)
342 [et al., 2017](#)), when an acceleration of learning from replay might be most beneficial.
343 Second, several studies reported that it requires only a few experiences in a novel
344 environment for replay to occur, and that it can be detected already during the
345 awake state immediately after behavior ([Foster and Wilson, 2006](#)), but see [Jackson](#)
346 [et al. \(2006\)](#). Third, disrupting replay-related SWRs during awake rest in rodents
347 slows learning in a spatial navigation task ([Jadhav et al., 2012](#)).

348 Previous research has also suggested that backward replay reflects learning through
349 temporal credit assignment in the brain. First, awake backward replay has indeed
350 been frequently observed, where rewarded spatial trajectories of an animal are re-
351 played in reverse order ([Diba and Buzsáki, 2007](#); [Foster and Wilson, 2006](#); [Singer and](#)
352 [Frank, 2009](#)), and the frequency of awake backward (but not forward) replay is mod-

353 ulated by the change in reward magnitude ([Ambrose et al., 2016](#); [Liu et al., 2019](#)).
 354 Second, the rate of backward replay was observed to be more frequent in novel com-
 355 pared to familiar environments ([Foster and Wilson, 2006](#); [Singer and Frank, 2009](#))
 356 and to decrease its bias to reflect previous paths to the goal location as a function
 357 of learning ([Shin et al., 2019](#)). This could suggest that the relevant trajectory has
 358 been learned and does not need to be reinforced through replay anymore ([Foster](#)
 359 [and Knierim, 2012](#)). Interestingly, [Cazé et al. \(2018\)](#) have shown that in particu-
 360 lar model-based replay will also decrease its tendency to replay paths to the goal
 361 with learning, while changes in forward planning ([Johnson and Redish, 2007](#)) might
 362 stem from a model-free process. In a task setting with a stable goal, the replay
 363 buffer of a model-free learner will increasingly accumulate rewarded episodes while
 364 a model-based learner draws on a learned model to sample episodes in a more bal-
 365 anced fashion. The learning-related changes discussed here might therefore reflect a
 366 shift from a model-free to a model-based process with learning – although further
 367 data will be needed, and model-based and model-free replay might be difficult to
 368 disentangle experimentally ([Khamassi and Girard, 2020](#)). The third line of support
 369 comes from computational work that shows how backward replay can strengthen for-
 370 ward synaptic pathways through spike timing dependent plasticity (STDP) ([Haga](#)
 371 [and Fukai, 2018](#)) and thus support forward replay during sleep and active behav-
 372 ior ([Johnson and Redish, 2007](#); [Pfeiffer and Foster, 2013](#); [Wikenheiser and Redish,](#)
 373 [2015b, 2013](#)). Fourth, further evidence for the role of replay in assigning credit is
 374 provided by findings that show replay is coordinated with subcortical activation of
 375 brain areas related to processing reward ([Lansink et al., 2009](#); [Pennartz, 2004](#); [Gom-](#)
 376 [perts et al., 2015](#)), which could convey reward signals to other brain regions like the
 377 hippocampus. Finally, in a recent MEG study in humans, backward replay following
 378 reward receipt was found to be related to non-local learning of task sequences leading
 379 to the reward ([Liu et al., 2021b](#)). In summary, existing empirical studies support the
 380 idea that awake backward replay supports temporal credit assignment by retrieving
 381 states that led to the outcome, accelerating learning for cases in which a long delay
 382 between rewards and actions must be encoded.

Figure 2: Replay speeds up learning to navigate to a goal in a grid world. (A) Square environment (“grid world”) with 20×20 tiles (shown in gray) that contains several walls (black tiles) and one goal location (white tile labelled “G”) that contains a reward. At the beginning of each episode, an RL agent is placed in a random location and can move into one of four cardinal directions (up, down, left, right). The agent receives no reward for moving, a small negative reward for bumping into a wall (-0.1), and a reward of 1 when arriving at the goal location. An episode is terminated once the agent reaches the goal location or a maximum of allowed steps per episode set to 1000. (B) Illustration of the learned value function after the first episode of experience (top left), following replay of the first episode (top right), and after the 50th and 250th episode (bottom). Colors indicate the values of locations from smallest (blue) to highest (red). Values are under the best possible policy, which is assuming that the agent would perform the value-maximizing action in each location. The increasing prevalence of red tiles after 250 episodes therefore reflects that after training the agent has learned a policy for most locations that will avoid any collisions with the wall and reach the goal within the maximum number of allowed steps. Color mapping is scaled for each plot and values smaller than 0.1 are shown as gray tiles. (C) Number of steps (y-axis) needed by the RL agent in each consecutive episode (x-axis) to reach the goal location when using no replay (blue line) between episodes, or when replaying the previous episode in backward order once (brown line) or five times (yellow line). (D) Mean reward (y-axis) achieved by the RL agent in each consecutive episode (x-axis). Colors as in (c). The computer code for the simulations is publicly available at <https://github.com/nschuck/replaysim-wittkuhn-etal2021>. © Wittkuhn et al., doi: 10.6084/m9.figshare.14261636.v1, CC-BY 4.0 license (<https://creativecommons.org/licenses/by/4.0/>).

2.1.2 An example simulation of backward replay

Figure 2 provides an illustration of how backward replay of full sequences works in the context of RL. We consider an RL agent navigating in a square environment with 20×20 tiles that contains several walls and one goal location with a reward (see Figure 2A). The agent can move into one of the four cardinal directions (up, down, left, right). A small negative reward is given for bumping into a wall (-0.1), and a reward of 1 when arriving at the goal location. Otherwise no rewards are provided. The best policy in this case is to navigate to the reward with as little steps as possible, avoiding the wall. This is a well-known “grid world” problem that can be solved using RL, but might be painfully slow without replay. For illustrative purposes, we use the off-policy, model-free Q -learning algorithm described in Equation 1 in Box 1. The learning rate α , temperature τ and discounting factor γ were arbitrarily set to 0.3, 1 and 0.99 for the purposes of this illustration.

Table 1 Q -learning with Replay

```

1: Initialize  $Q(s, a) \leftarrow 0$  for all states  $s$ , actions  $a$  and parameters  $\{\alpha, \gamma, \tau\}$ 
2: for  $episode = 1, 2, \dots$ , number of episodes do
3:    $s \leftarrow$  random start position farther than 10 steps from goal
4:   for  $step = 1, 2, \dots$ , maximum walk length do
5:      $s \leftarrow$  state at current step
6:      $a \leftarrow$  softmax( $s, Q$ ) with temperature  $\tau$ 
7:     Execute  $a$ , obtain next state  $s'$  and reward  $r$ 
8:      $Q(s, a) \leftarrow Q(s, a) + \alpha[r + \gamma \max_{a'} Q(s', a') - Q(s, a)]$ 
9:     if goal reached then
10:      for  $t = \text{step-1, step-2, } \dots$  beginning of episode do
11:         $s, a, r \leftarrow$  state, action, reward at  $t$ 
12:         $s', a' \leftarrow$  state, action at  $t + 1$ 
13:         $Q(s, a) \leftarrow Q(s, a) + \alpha[r + \gamma Q(s', a') - Q(s, a)]$ 
14:      end for
15:    End episode and continue to next
16:   end if
17: end for
18: end for
```

The algorithm is described in Table 1. Briefly, in each episode, the agent starts in a random position and navigates until it has found the reward or the maximum search time has elapsed. The starting locations varied randomly, although start location distances to the reward location were constrained to lie at least 10 tiles away from the goal location (in order to avoid episodes which were too simple). If the agent found the reward, it internally traversed backwards through the sequence of states, actions and rewards until the beginning of the episode, updating its Q -value at each

403 step.

404 The blue lines in Figure 2C–D show the number of steps the agent needs to
405 navigate to the goal location; about 250 episodes are needed before the the agent
406 quickly finds the goal location from a new start position. But the speed of learning
407 increases when we supply the agent with a simple replay mechanism described in
408 Table 1, as can be seen in Figure 2C–D (brown and yellow lines). Adding replay
409 reduced the number of interactions needed to achieve ceiling performance to less
410 than half of what was observed without replay. Note that the choice to replay the
411 full sequence of states, actions and rewards between the start location and the goal
412 location is not without consequence, and a variety of different definitions of what
413 constitutes an episode are common in RL and neuroscience (see Box 2). There are
414 multiple ways to instantiate replay in an RL agent, and the illustration in Figure 2
415 only serves as a basic introduction to computational replay (see Figure 1).

416 **Box 2 What is replayed?**

417

418 Replay is generally thought to represent previous experience. How is this expe-
419 rience stored in artificial and biological agents? In artificial agents, an experience at
420 time t , e_t , is commonly defined as a quadruple consisting of the state s_t , the taken
421 action a_t , the reward r_t received after taking action a_t in state s_t , and the next state
422 s_{t+1} , together $e_t = (s_t, a_t, r_t, s_{t+1})$, effectively describing a single transition between
423 two states as the atomic unit of an artificial replay event. Although in some cases
424 individual transitions are replayed, such as in the Deep Q-Network (DQN) approach
425 by Mnih et al. (2015) where the states S consisted of pre-processed versions of Atari
426 pixel frames, other work uses sequential replay of past states (e.g., the early version of
427 experience replay by Lin (1992), see Figure 2, or replay in recurrent neural networks
428 (RNNs), see e.g., Hausknecht and Stone, 2015; Kapturowski et al., 2019). Interest-
429 ingly, replay techniques in ML increasingly reactivate internal state representations,
430 rather than observations like pixel values (Hayes et al., 2021). We will discuss this
431 aspect in more detail in section 2.6 on representation learning.

432 What constitutes a replayed experience is more difficult to answer for biologi-
433 cal agents. Unlike in artificial agents, replay in biological agents is thought to be
434 sequential (Genzel et al., 2020), and typically involves hippocampal place cells that

435 represent locations in a spatial environment, akin to previously experienced trajec-
 436 tories of locations. However, hippocampal cells appear to be quite flexible in encoding
 437 task-relevant information other than physical space, for instance sounds (Aronov
 438 et al., 2017), trial history (Wood et al., 2000; Sun et al., 2020) or abstract task states
 439 (Schuck and Niv, 2019). Moreover, a prominent theme in neuroscience emphasizes
 440 that the brain segments continuous experience into representations of distinct neural
 441 states that transition at event boundaries or shifts in context (for reviews, see e.g.,
 442 Richmond and Zacks, 2017; Brunec et al., 2018; Shin and DuBrow, 2020; Bird, 2020;
 443 Maurer and Nadel, 2021). To complicate matters, this process might also happen
 444 retroactively, i.e., *after* experiences have been obtained (Clewett et al., 2019). This
 445 formation of segmented memory traces is thought to be driven by various factors,
 446 including inferred changes in the environment (DuBrow et al., 2017), prediction er-
 447 ror signals elicited by reward outcomes (Rouhani et al., 2020) or discontinuities in
 448 the statistical structure of the environment (Gershman et al., 2014). We suggest
 449 that a practical approach for human research therefore seems to be to define events
 450 as “meaningful” units of experience (Bird, 2020) within the current experimental
 451 paradigm, and to potentially formalize them as states in an MDP, as for instance in
 452 Schuck et al. (2016). Finally, in understanding memory as a constructive process, it
 453 is important to note that neural task representations may change from perception
 454 to reactivation (Favila et al., 2020). We argue that this aspect is particularly crucial
 455 for the study of replay in humans, because activity patterns that are expected to
 456 reactivate are commonly determined based on simple localizer tasks that do not in-
 457 volve mnemonic task components (see e.g., Wittkuhn and Schuck, 2021). The brain
 458 might have already transformed its input data to a representation that is different
 459 from what the researcher was hoping to see re-merge from replayed activity patterns.

460 How many experiences are stored and for how long? In DNNs using replay, the
 461 newest experiences are stored at each time step in a memory buffer $D = e_1, \dots, e_t$
 462 with finite size N (e.g., Mnih et al., 2013; Mnih et al., 2015; Zhang and Sutton, 2017).
 463 Since the success of the DNN by Mnih et al. (2015), the memory buffer is typically
 464 set to a size of $N = 10^6$ newest experiences, which continuously replace the oldest ex-
 465 periences (Fedus et al., 2020; Zhang and Sutton, 2017). Recently, Fedus et al. (2020)
 466 investigated the relationship between the number and age of experiences stored in the
 467 memory buffer. First, they found that increased memory capacity improved learn-

ing performance, likely due to a larger coverage of state-action pairs (Fedus et al., 2020). Second, decreasing the age of the oldest experience in the memory buffer also improved performance, likely because of older experiences that resulted from policies that are inconsistent with the current on-policy decision strategy, which is in line with earlier findings noting that experience replay is only beneficial if it is consistent with the current decision policy (Lin, 1991, 1993). An exception to this are certain Atari games that are characterized by sparse rewards and require high levels of exploration. In such tasks, sampling from older off-policy experiences is still beneficial (Fedus et al., 2020). These considerations about the size and age of the memory buffer in artificial agents point to an intriguing trade-off between the utilities of old and new memories: On the one hand, a youthful memory buffer storing only recent experiences can effectively drive the current decision policy and quickly abandon outdated and potentially inefficient behavior. On the other hand, keeping older experiences and integrating them with recent ones may foster generalization and prevent an agent from becoming stuck in a decision policy that is suboptimal.

While the size and content of a memory buffer in artificial agents can be crafted by ML researchers, determining number and nature of memories in brains is topic of ongoing debate for neuroscientists. The human brain is famously known to have a very large storage capacity, owing to the large number of modifiable synapses (Bartol et al., 2015). But forgetting is a common phenomenon. Although decay plays some role in forgetting (Hardt et al., 2013), other factors, such as interference and usage seem to be important as well (Feld and Born, 2017). Indeed forgetting might also be an important aspect of sleep, even while replay processes lead to consolidation (Feld and Born, 2017). Moreover, even if biological agents had an unlimited memory storage, selecting memories for replay from that storage would become challenging with a large amount of experiences, and, from a decision-making perspective, memory representations are only useful in so far they have utility for behavior.

2.2 Less forgetting

Increasing the speed of learning is an important computational benefit of replay, but not the only one. Replay may also help to reduce forgetting. The problem of forgetting arises because many statistical learning mechanisms were built under the assumption that the agent encounters its environment entirely at random, and

500 therefore can learn from examples that are independent and identically distributed
501 (i.i.d.). Yet, experiences in real life are often not “i.i.d.”. First, we typically expe-
502 rience the world as a sequence of related events, chunked in time. Second, some
503 events are much rarer than others, partly because of the way we interact with the
504 environment. These temporal auto-correlations and uneven distributions of events
505 can be an important obstacle for learning. Why does this pose a computational chal-
506 lenge for gradual learning algorithms? Gradual learning mechanisms are designed
507 to integrate experiences over longer periods, but they emphasize the most recent ex-
508 perience. This can cause the agent to forget about important past experiences that
509 were not re-experienced for a long time. This problem is particularly apparent in a
510 supervised learning setting in which neural networks that rely on stochastic gradi-
511 ent descent (SGD) for learning engage in two tasks in a *blocked* manner. RL-based
512 networks also struggle with this phenomenon (Atkinson et al., 2021). If a DNN, for
513 instance, is first trained to perform a task A and subsequently trained with another
514 task B, performance on task A drops dramatically, as if the network forgot how to
515 solve A. In other words: learning task B interfered with what was learned about task
516 A. This problem is known as catastrophic forgetting, or catastrophic interference,
517 and has long been recognized as a major problem in the ML field (McCloskey and
518 Cohen, 1989; Ratcliff, 1990; French, 1999; Hassabis and Maguire, 2007; Kumaran
519 et al., 2016; Parisi et al., 2019). Catastrophic forgetting is one of the main reasons
520 why artificial agents can usually learn a single task quite well but subsequent training
521 on a different task results in poor performance on the previously learned task. This
522 prevents the agent from achieving competencies across multiple tasks, which comes
523 relatively easily to humans. Catastrophic interference can also be understood as an
524 issue threatening the stability of a cognitive map representation (Gupta et al., 2010;
525 McClelland et al., 1995; O’Reilly and McClelland, 1994).

526 2.2.1 Replay can prevent overwriting of previous experiences

527 A potential solution to the computational problem of catastrophic interference is *in-*
528 *terleaved learning* where new experiences are interleaved with existing knowledge to
529 reconcile competing memory representations (McClelland et al., 1995). This influen-
530 tial idea, rooted in the complementary learning systems (CLS) theory (McClelland
531 et al., 1995; O’Reilly et al., 2014; Schapiro et al., 2017), also suggests that replay

may be the mechanism that agents can use to “mentally” interleave past with present experience. While the DNN tunes its connection weights to solve task A, the experienced episodes are stored in a memory buffer. During learning of task B, previous experience with task A is integrated during offline periods via a replay-like mechanism, preventing forgetting, and allowing the agent to perform well on both tasks. [Shin et al. \(2017\)](#), for instance, proposed an approach that learns a generative model based on experience with one classification task A. When switching to an independent classification task B, the system is retrained using a combination of new task data and fictitious sequences from the generative model, resulting in rapid generalization to the new task with little performance loss. Similarly, implementing replay in this way in DNNs can help to overcome performance deficits in incremental task learning scenarios and continuous task environments ([van de Ven and Tolias, 2018](#); [van de Ven et al., 2020](#)). Of note, non-sequential replay has been shown to become necessary in an artificial neural network (ANN) when an internal model of a continuous task environment has to be learned ([Aubin et al., 2018](#)).

Despite its benefits, interleaved replay can result in problems if the agent’s current policy is very different from the behavioral policy when the experiences were collected. To account for this, several authors have argued that replay needs to be corrected for such “off-policyness” using importance sampling ([Meuleau et al., 2010](#)) or other off-policy correction methods such as Retrace ([Munos et al., 2016](#)) or V-trace ([Espeholt et al., 2018](#)). These approaches essentially weight updates that result from replay in proportion to the mismatch between the policy used to generate the replay and the agent’s current policy. This issue is particularly pressing in distributed replay approaches (e.g., [Horgan et al., 2018](#)), where virtual experiences are simulated in parallel and are then used for learning only with some time gap.

Given that humans and other animals do not necessarily seem to suffer from the computational problem of catastrophic interference, the question arises how the brain has apparently solved this issue and, for the purpose of this review, whether replay plays a role in the solution. Humans and animals can solve a wide set of tasks throughout their lifetime, despite temporal autocorrelation of experience and even learn well from blocked experience which troubles DNNs ([Flesch et al., 2018](#)). The idea that this ability might be related to replay ([Antony and Schapiro, 2019](#)) is supported by several studies. [Karlsson and Frank \(2009\)](#) for instance have observed

replay of episodes from a remote spatial context. In humans, reactivation of previously learned events in the hippocampus that overlap with newly encoded memories leads to better retention (Kuhl et al., 2010).

2.2.2 Replay can amplify the influence of rare events on learning

Another challenge arises when learning must occur in environments where some events happen rarely, but are nevertheless of great significance for the agent’s success or well-being. Naive DNNs will, for instance, often forget about dangerous states and revisit them (García et al., 2015). This can be mitigated by replay when separate replay buffers for safe and dangerous states are maintained, such that the model cannot forget, and will frequently be reminded about dangers (Meuleau et al., 2010; Lipton et al., 2016). More generally, importance sampling techniques have been used to ensure that those experiences are sampled which are most important for the current policy of the agent, rather than those that occurred most frequently (Wang et al., 2016).

Evidence that replay might be used to mitigate this problem in animals comes from studies showing that actions which should be avoided will be reactivated, like paths to a shock zone (Wu et al., 2017) or paths to devalued outcomes (Carey et al., 2019). In addition to learning about events that should be avoided, replaying rare events that are only weakly encoded could allow the agent to form a stable representation of the entire environment even if only a smaller subset is experienced frequently. In Gupta et al. (2010) non-local replay was stronger for remote sequences if they were experienced less frequently. Using MEG in humans, Jafarpour et al. (2017) showed that stronger reactivation of one of three previously encoded stimuli was determined by how weakly the stimulus was attended to during encoding. These findings are supported by an fMRI study by Schapiro et al. (2018), who demonstrated that older, less well remembered task stimuli were selectively reactivated during a subsequent rest period resulting in memory improvement, an effect that was particularly strong in participants who slept in the 12-hour interval between test sessions, which likely offered opportunity for additional consolidation through replay. In another study, the benefits of TMR were stronger for weakly learned information (Tambini et al., 2017). Further, replay-associated electroencephalography (EEG) sleep spindles during a nap following difficult (potentially weaker) but not easy (potentially stronger) memory

597 encoding were related to improved subsequent memory performance (Schmidt et al.,
598 2006). Together, we suggest that replay liberates an agent from needing to consider
599 transitions only in proportion to how many times they were experienced. Instead,
600 replay can flexibly increase or decrease the number of opportunities for learning from
601 single episodes.

602 2.3 Re-inventing the past

603 In our introductory example (see Figure 2), replayed content was a close reflection
604 of past experience. Replay occurred immediately after an episode was experienced
605 and reflected past trajectories from start to finish, albeit in reverse order. This setup
606 stands in contrast to the ideas discussed in section 2.2 on forgetting, which imply
607 that replay must not necessarily respect the structure of experiences, but could, for
608 instance, change the order and frequency of events. Beyond dealing with unevenly
609 distributed events, replay could in fact be used to arbitrarily alter the distribution
610 of events upon which memory is built.

611 Such a reorganization of experience also requires a different understanding of
612 what constitutes an episode. In our simulation, we had assumed that the minimal
613 unit of replayed content is one entire sequence of states, actions and rewards that
614 occurred between a random start position and the encounter of a goal. This meant
615 that episodes were often quite long, involving several hundreds of steps particularly
616 early in learning (see Figure 2C), and that the transitions between locations had
617 to be replayed in the order in which they were experienced. But for replay to be
618 able to reorganize experience, an episode could be divided into a much smaller unit
619 of experience, a simple sequence of just one state, one action, one reward and the
620 next state, known as a (s_t, a_t, r_t, s_{t+1}) -tuple. Arguably, replaying such minimal
621 experiences risks losing the benefits of temporal credit assignment, because values
622 will not necessarily propagate along the trajectory to starting positions. But it does
623 offer important advantages, discussed in subsections 2.3.1 to 2.3.4. In consequence,
624 the question of what constitutes an atomic unit of experience from the perspective
625 of replay has important implications and is therefore actively debated (see Box 2).

626 2.3.1 Replay can reactivate experiences randomly

627 Using minimal transitions, there is a large variety of ways in which replay may alter
628 the structure of experiences that have been discussed in the ML and neuroscience
629 literature. One possibility is to reactivate (s_t, a_t, r_t, s_{t+1}) -tuples in a random order,
630 which artificially crafts similar conditions as during supervised learning that allows
631 ANNs trained with stochastic gradient descent (SGD) to excel (Botvinick et al.,
632 2020). Such uniformly sampled (s_t, a_t, r_t, s_{t+1}) -tuples have therefore played an
633 important role in adapting DNNs to RL problems, such as the famous DQN (Mnih
634 et al., 2015). Random replay has also been found to be useful when updates are done
635 incrementally (learning from each example as it arrives), rather than in a batch-wise
636 manner (learning from groups of examples gathered over time), as is common in ML
637 (Chaudhry et al., 2018). Interestingly, some animal studies have also found replay
638 of seemingly random trajectories following exploration of a familiar open-field arena
639 (Stella et al., 2019). Note however, that Stella et al. (2019) still observed replay of
640 sequentially organized transitions that reflected the spatial constraints of the envi-
641 ronment, whereas random replay used in ML can involve sets of single transitions
642 that do not form sequential trajectories. This highlights the different understanding
643 of replay content in ML and neuroscience. Additionally, most animal studies impose
644 the assumption of sequentiality during data analysis, and would discard fully random
645 activation of transitions as noise. In both ML and neuroscience, however, random
646 replay refers to sequential reactivation that is unrelated to previously experienced
647 action sequences, and can be seen at one extreme of a continuum describing how
648 closely replay matches actual behavioral sequences (see Swanson et al., 2020, their
649 Figure 2).

650 2.3.2 Replay can prioritize rewarding experiences

651 Another particularly important idea from the ML literature is to prioritize replay of
652 transitions that led to large surprises, which often prove to be more informative than
653 others and result in more efficient learning (Schaul et al., 2015; Horgan et al., 2018).
654 Such prioritized replay records a prediction error (PE), the difference between the
655 expected and actual reward, for every encountered transition and uses this signal
656 to select experiences for replay later. This method is very similar to, and inspired
657 by, an earlier algorithm in model-based planning known as prioritized sweeping,

658 which selects the state to be updated according to the magnitude of the change
 659 in value upon the execution of the update (Moore and Atkeson, 1993; Peng and
 660 Williams, 1993; Andre et al., 1998). Based on the success of the prioritized replay
 661 approach, more frequent sampling of transitions with a high absolute TD error is
 662 now a common approach to train DNNs (Fedus et al., 2020). Using RL models,
 663 Mattar and Daw (2018) extended previous approaches by focusing prioritization on
 664 behaviorally relevant states that are likely to be encountered again in the future
 665 and those transitions where a policy change would yield the largest net increase
 666 in discounted future reward. Note that although prioritization algorithms assume
 667 selection of replay content on the level of individual transitions, they can, under some
 668 circumstances, still lead to sequential replay. This is true, for example, in the model
 669 from Mattar and Daw (2018), because expectations about increases in future reward
 670 are themselves often auto-correlated.

671 The idea that replay should be influenced by reward and surprise is in line with
 672 several animal studies. Place cell sequences associated with reward are replayed more
 673 often (Ólafsdóttir et al., 2015; Foster and Wilson, 2006; Bhattarai et al., 2019), in
 674 particular those with a high PE (Singer and Frank, 2009; Michon et al., 2019; Roscow
 675 et al., 2019), and the rate of SWRs is also influenced by reward (Ambrose et al., 2016;
 676 Singer and Frank, 2009). These results highlight replay’s role in credit assignment,
 677 as discussed in section 2.1. In human neuroimaging studies, hippocampal activity
 678 is modulated by reward magnitude (Wolosin et al., 2012; Igloi et al., 2015). This
 679 is in line with the link between backward replay and selection of transitions based
 680 on changes in value, proposed by Mattar and Daw (2018) and Cazé et al. (2018).
 681 Replay is also more likely to contain behaviorally significant locations, such as the
 682 current goal (Gupta et al., 2010; Pfeiffer and Foster, 2013; Ólafsdóttir et al., 2015)
 683 and is biased by novelty (Cheng and Frank, 2008; Foster and Wilson, 2006). It has
 684 also been observed that optogenetic manipulation of dopaminergic input neurons,
 685 thought to signal PEs, increase replay during subsequent sleep (McNamara et al.,
 686 2014). Note that reward prediction errors might be accompanied by state prediction
 687 errors, which in the brain might both be conveyed by dopaminergic signals (see e.g.,
 688 Sharpe et al., 2017; Gardner et al., 2018).

689 2.3.3 Replay can connect experiences in novel ways

690 Reactivating reordered sequences can also be used to connect experiences in novel
691 ways or strengthen weakly learned relationships. Replay can for instance correspond
692 more closely to sampling from an internal model of the environment, rather than a
693 veridical recapitulation of past experiences (Sutton, 1991). Among the most early
694 ideas about replay, the Dyna architecture (Sutton, 1991) used an internal model to
695 generate experiences that were then used to train a model-free agent. Indeed, replay
696 can be seen as a way to blur the lines between model-free RL, such as the Q -learning
697 method introduced in Box 1, and model-based RL, during which the agent stores
698 an explicit model of the environment and can use it for planning (van Seijen and
699 Sutton, 2015; Russek et al., 2017; Momennejad et al., 2017).

700 Neuroscientific evidence for reorganized experience has also been reported. Dur-
701 ing behavior, replay events can switch between reflecting immediately preceding,
702 upcoming or more remote episodes, depending on the behavioral state of the animal
703 at the time of replay (Pfeiffer and Foster, 2013; Ólafsdóttir et al., 2017). Even single
704 replay events can depict more than one trajectory, such as the next one and the path
705 the animal will take after reaching the goal location (Pfeiffer and Foster, 2013), as if
706 representing a multi-step planning process (Foster, 2017; Miller and Venditto, 2021).
707 Note, however, that another reason why experiences from a more distant past might
708 be replayed could simply be that the agent is using a period during which it does not
709 have to engage with the environment to optimize memory. This is particularly appar-
710 ent for replay during sleep, when the brain has idle time to process experiences while
711 not being actively engaged with any task. Sleep replay has frequently been observed
712 in animals and humans, and been linked in particular to memory consolidation. Fol-
713 lowing sleep, memory interference is reduced (Baran et al., 2010; McDevitt et al.,
714 2015) and memory integration or differentiation has been found in fMRI patterns
715 after a delay period with sleep (Favila et al., 2016; Tompary and Davachi, 2017).

716 2.3.4 Replay can reuse past experiences to learn about new goals

717 A final aspect of reorganization relates to re-considering the usefulness of past ex-
718 periences in light of one’s knowledge about a goal. The RL framework presented
719 thus far is aimed at the pursuit of a single goal (e.g., the single reward location in
720 our grid world, see Figure 2). However, in many real-world applications, such as the

721 movements of a robotic arm that needs to pick and place objects, an RL strategy
722 incorporating multiple goals would be far more beneficial. Consider again the grid
723 world example in Figure 2, but this time the agent can only move a finite number of
724 steps. Since there is only one goal state that returns a reward, most of the transitions
725 do not land in the goal state and therefore receive no reward. In such a sparse bi-
726 nary reward situation, where success only results from those sequences of transitions
727 ending in the goal state, most sequences of transitions end in uninformative failures,
728 often related to early termination without reward (“giving up”). For instance, when
729 an agent gives up because a goal was not found after a particular amount of time, it
730 can not know how close it was to the goal. Humans, however, can learn from failure
731 as well as success. Inspired by this idea, an ML technique known as hindsight experi-
732 ence replay ([Andrychowicz et al., 2017](#)) is used to relabel the unsuccessful transitions
733 by simply changing the goal state, such that the transitions would now be considered
734 as successful under the new goal, thus contributing to the agent’s learning. To the
735 best of our knowledge, no directly equivalent observation has been made in the brain.

736 2.4 Planning for a better future

737 So far, we have mainly focused on the various ways in which replay serves learning
738 and memory. Yet, psychological, neuroscientific and ML research has pointed out the
739 importance of another mechanism that is crucial for goal-directed behavior: planning.
740 A core aspect of this process is the prospective evaluation during which an agent
741 deliberates which of the available sequences of actions and states leads to the best
742 among several potential outcomes. In most cases, planning requires a mental model,
743 or cognitive map ([Tolman, 1938, 1948](#)), of the environment, that describes the agent’s
744 knowledge about the transition structure of events, including the outcomes at each
745 potential location (e.g., [Moerland et al., 2020](#)). Knowledge about the causal structure
746 of the environment allows an agent to predict and compare the outcomes of sequences
747 of states and actions and to choose the one that yields most reward. Yet, as we will
748 see below, a replay buffer can be used instead of a model in order to perform planning
749 functions too. Figure 3 provides an illustration of how planning differs from the other
750 two aspects of cognition we have considered so far, acting and learning.

751 Within the RL framework, the difference between acting based on learned (cached)
752 values versus acting based on an internal planning process is embodied by the dis-

Figure 3: **Illustration of replay content.** The agent is represented by the robot. The goal location is indicated by the yellow circle. Red arrows indicate behavioral trajectories of the agent in the T-maze environment or internally generated trajectories during replay. During planning (left panel) the agent engages in prospective evaluation of potential behavioral sequences in order to select the one that leads to the goal location, using forward replay. During goal-directed behavior (center panel) the agent instantiates the behavioral trajectory that is immediately relevant to prepare action using forward replay. During learning (right panel), the agent retrospectively evaluates its previous behavior, usually upon reaching a goal location, using backward replay. © Wittkuhn et al., doi: [10.6084/m9.figshare.14261636.v1](https://doi.org/10.6084/m9.figshare.14261636.v1), CC-BY 4.0 license (<https://creativecommons.org/licenses/by/4.0/>).

753 tinction between model-free and model-based systems (Sutton and Barto, 2018; Daw
754 et al., 2005). In model-based RL, the agent uses experience to learn a model of the
755 environment that is described by a function that relates the current state s_t and
756 action a_t to the next state s_{t+1} and the reward r_t (see Box 1). The agent can use
757 this model at decision time or during offline periods to simulate experience (Sutton,
758 1990). Simulated replay can be used to update cached values or to determine which
759 action would be best to execute next, considering the rewards obtained and how
760 the environment would change if a particular action was taken. This deliberation
761 process has two advantages. First, planning allows the agent to remain in the safety
762 of mental imagination and avoid the risk of suffering from potentially harmful conse-
763 quences. Second, planning can be used to decide between never-experienced, entirely
764 hypothetical courses of action (Liu et al., 2019), a feat which would not be possible
765 with purely experience-based replay.

766 Despite these differences between planning and learning, much work in RL has
767 emphasized their similarities (Sutton, 1990, 1991; Sutton et al., 2012; van Seijen
768 and Sutton, 2015). This research points to a function of planning that goes beyond
769 deliberation and has shown that planning functions can be achieved without an
770 explicit model (van Seijen and Sutton, 2015; van Hasselt et al., 2019).

771 As we have seen in our discussion of [Lin \(1992\)](#), the same learning mechanisms
772 can be applied to real or simulated experience. Planning can thus not only be used to
773 determine immediate behavior, but also to shape value functions, a process referred
774 to as background planning (as opposed to decision-time planning and deliberation,
775 see [Pezzulo et al., 2019](#)). This can be illustrated by the Dyna architecture ([Sutton,](#)
776 [1990, 1991](#)). Just as any model-free RL agent, a Dyna agent selects actions according
777 to learned Q -values, and uses experiences to update these Q -values. But it also uses
778 experiences to observe which states and rewards follow the current action, using
779 this information to update its internal model of the world. Importantly, in Dyna the
780 model is then used to train the model-free agent by replaying simulated episodes, and
781 updating the agent's Q -values based on prediction errors, just like real experiences.

782 A second aspect is that replaying experiences stored in a memory buffer in some
783 sense replaces functions that would otherwise be subserved by a model ([van Seijen](#)
784 [and Sutton, 2015](#); [Hessel et al., 2018](#); [van Hasselt et al., 2019](#)), or at least enhance
785 model-based planning functions ([Eysenbach et al., 2019](#)). [van Seijen and Sutton](#)
786 ([2015](#)), for instance, have shown that learning value functions by a model-free method
787 with replay can be equivalent to learning value functions with a model-based method.
788 Empirically, [van Hasselt et al. \(2019\)](#) have shown that state-of-the-art replay meth-
789 ods, involving prioritization based on a Kullback-Leibler (KL) loss, can outperform
790 model-based methods on Atari games ([Kaiser et al., 2019](#)), in part because an in-
791 accurate model can lead to unstable learning. Moreover, [Eysenbach et al. \(2019\)](#)
792 have shown that replay can be used to infer a graph representation of the current
793 task that provides insights into subgoals, which in turn can be used for planning (cf.
794 [Pong et al., 2018](#)). This is reminiscent of hindsight replay, which retroactively inserts
795 rewards into stored replay sequences in order to facilitate learning about hierarchical
796 subgoals ([Andrychowicz et al., 2017](#)). We note, however, that model-based planning
797 methods remain popular in ML ([Pan et al., 2018](#); [Kaiser et al., 2019](#); [Moerland et al.,](#)
798 [2020](#)). Planning methods provide the flexibility needed to generate unseen but pos-
799 sible transitions, and planning over long horizons can be achieved using algorithms
800 such as tree search ([Guo et al., 2014](#); [Silver et al., 2016](#); [Anthony et al., 2017](#)).

801 While the potential benefits of replay for planning have been recognized early
802 on in RL ([Sutton, 1990](#)), consideration of this aspect in neuroscience only appeared
803 later, when studies demonstrated replay events in the awake state, often during short

804 pauses from active behavior (e.g., [Kudrimoti et al., 1999](#); [Foster and Wilson, 2006](#);
805 [Csicsvari et al., 2007](#); [Diba and Buzsáki, 2007](#); [Kurth-Nelson et al., 2016](#); [Eldar et al.,](#)
806 [2020](#)). This allowed researchers to draw closer correspondence between the replayed
807 and the behavioral trajectories, and has resulted in a wealth of findings supporting
808 the idea that replay supports model-based planning in animals as well as humans (for
809 reviews, see e.g., [Yu and Frank, 2015](#); [Pezzulo et al., 2019](#); [Wang et al., 2020](#); [Tambini](#)
810 [and Davachi, 2019](#); [Carr et al., 2011](#); [Ólafsdóttir et al., 2018](#)). Disruption of awake
811 hippocampal SWRs during a spatial alternation task specifically impaired the ability
812 to decide between two trajectories to alternating goal locations, whereas place field
813 representations, reactivation during rest, and other navigation behavior remained
814 intact ([Jadhav et al., 2012](#)). Replay events in the awake state predominantly co-occur
815 with SWRs during short pauses from ongoing exploratory behavior. Forward replay
816 trajectories during awake SWRs often start at the current location of the animal (a
817 well-known “initiation bias”, [Diba and Buzsáki, 2007](#); [Davidson et al., 2009](#); [Karlsson](#)
818 [and Frank, 2009](#); [Pfeiffer and Foster, 2013](#); [Singer et al., 2013](#); [Ambrose et al., 2016](#)),
819 and end at the goal location ([Dupret et al., 2010](#); [Pfeiffer and Foster, 2013](#)), but not
820 always (see e.g., [Johnson and Redish, 2007](#)).

821 A behavioral correlate of deliberation was already described in the 1930s in ro-
822 dents ([Tolman, 1926](#); [Muenzinger and Fletcher, 1936](#)), who tend to pause at a decision
823 point to look back and forth between possible paths, a behavior called vicarious trial
824 and error (VTE) (for review, see [Redish, 2016](#)). Later studies found that during
825 VTE events, hippocampal place cells associated with theta sequences sweep ahead
826 from the animal’s current location ([Johnson and Redish, 2007](#); [Wikenheiser and Re-](#)
827 [dish, 2015b](#); [Amemiya and Redish, 2016](#); [Papale et al., 2016](#)). It was also found
828 that during VTE-like behavior, place cell activity influenced the formation of place
829 fields thought to stabilize the cognitive map ([Monaco et al., 2014](#)). Note that VTE-
830 associated replay is often accompanied by theta sequences, which differ from SWRs
831 in their neurophysiology. Nonetheless, both can be described as sequential activa-
832 tion of hippocampal cell populations, a simplifying assumption that is helpful from
833 a computational perspective ([Foster, 2017](#); [Pezzulo et al., 2019](#)). Recently, theta se-
834 quences have been shown to quickly cycle between possible future trajectories ([Kay](#)
835 [et al., 2020](#)), and increases in theta power in the MTL have been observed in humans
836 in a spatial planning task ([Kaplan et al., 2020](#)). In human fMRI, blood-oxygen-level

837 dependent (BOLD) activity in the hippocampus has been shown to increase with de-
838 liberation time when deciding between two food items with similar value (Bakkour
839 et al., 2019) and hippocampal activity patterns reflect routes to navigational goal
840 locations (Brown et al., 2016). Another study has found that when humans re-learn
841 outcomes associated with choices at lower levels of a decision tree, the extent to
842 which higher levels of the decision tree are reactivated during rest correlates with
843 how much their decisions change, to reach the new downstream reward states (Mo-
844 mennejad et al., 2018).

845 2.4.1 Replay can influence behavior directly or indirectly

846 It should be noted that although awake replay during deliberation of future choices
847 is often related to improved task performance, the replayed trajectories do not neces-
848 sarily correspond to the behavioral trajectory the animal will subsequently take, and
849 sometimes do not end in the goal location (Johnson and Redish, 2007; Singer et al.,
850 2013). In the study by Singer et al. (2013), hippocampal replay during SWRs that
851 preceded correct choices reflected trajectories for the correct *and* incorrect option in a
852 two-alternative W-maze. Once correct performance became stable (at 85% correct),
853 replayed trajectories shifted to represent the correct future choice more frequently
854 than the incorrect one (Singer et al., 2013). One interpretation of these findings is
855 that the hippocampus uses replay to evaluate all potential trajectories and the be-
856 haviorally relevant trajectory is instantiated in a different brain region. Furthermore,
857 backward replay, which backpropagates value information from the goal location, and
858 forward replay, which samples possible trajectories ahead of the animal, might con-
859 nect their trajectories as proposed by models of bidirectional planning (Khamassi and
860 Girard, 2020). Forward replay events have been shown to end at or close to the goal
861 location (Pfeiffer and Foster, 2013) and might efficiently stop in states where value
862 estimates have already been updated by a backward replay mechanism, as could be
863 instantiated by prioritized sweeping (see Khamassi and Girard, 2020). In sum, these
864 findings support the idea that deliberation and learning may interact. Changes in a
865 familiar environment might increase deliberation, while the need for model updating
866 and deliberation could diminish with learning, e.g., because decisions become more
867 habitual and less deliberate (Dolan and Dayan, 2013). Thus, replayed trajectories
868 in the hippocampus that evaluate all potential trajectories might be only predictive

869 of behavior during earlier phases of learning (Singer et al., 2013) or vanish from the
870 hippocampus (e.g., Wimmer and Büchel, 2019) when behavior becomes stereotyped.
871 Findings by Papale et al. (2016) also demonstrate an inverse relationship between
872 SWRs at reward sites and deliberation at choice points.

873 One additional complicating factor regarding the relationship between replay and
874 subsequent behavior concerns the task setting and motivational state of the animal
875 (e.g., Carey et al., 2019; Wu et al., 2017). Take the example of replayed place cell
876 sequences representing the trajectory into a shock zone that is subsequently avoided
877 (Wu et al., 2017). This might serve the purpose of learning strongly from and not
878 forget about significant outcomes, and thus in this circumstance replay is related
879 to avoiding rather than initiating trajectories. In line with this idea, a growing
880 literature on computational psychiatry posits that replay could underlie symptoms
881 like avoidance and rumination that characterize psychiatric disorders like anxiety
882 (Gagne et al., 2018; Heller and Bagot, 2020; Mobbs et al., 2020).

883 If replay is related to planning, but the ultimate determination of behavior also
884 depends on other brain areas, then replayed trajectories might be influenced by
885 concurrent reactivation outside the hippocampus, such as the amygdala in the case
886 of aversive outcomes (Girardeau et al., 2017). A number of studies sheds light on
887 how replay in the hippocampus is coordinated with other brain regions to instantiate
888 behavior. Replay is known to be coordinated with PFC (Jadhav et al., 2016; Pezzulo
889 et al., 2014; Peyrache et al., 2009; Tang et al., 2017), and some work has placed
890 particular focus on the interaction of hippocampal replay and the OFC (Steiner and
891 Redish, 2012; Schuck and Niv, 2019). Indeed, disruption of nearby medial PFC
892 attenuated components of hippocampal theta sequences representing the current
893 location of the animal (Schmidt et al., 2019) and suppression of hippocampal input
894 impaired the integration of task state structure in the OFC (Wikenheiser et al., 2017).
895 Similarly, a recent study in humans has found that hippocampal replay at rest was
896 not directly linked to behavior during a task (Schuck and Niv, 2019). Rather, replay
897 at rest was linked to how well the different task-states were represented in the OFC,
898 which in turn were linked to behavior (Schuck and Niv, 2019). Outside of the PFC,
899 entorhinal grid cells that are thought to enable vector-based spatial navigation likely
900 contribute to planning, as implicated in computational work (see e.g., Erdem and
901 Hasselmo, 2012; Bush et al., 2015).

902 2.4.2 Preplay can help planning in unknown environments

903 While replay is mostly thought to occur *after* experiences have been made, some
904 ideas have been put forth that assume a “preplay” mechanism, in which experiences
905 are mapped out before they are encountered. Few models related to this idea have
906 been proposed in ML, but perhaps the closest concepts are related to attractor
907 dynamics or reservoir computing (for a review, see e.g., [Lukoševičius and Jaeger, 2009](#)).
908 Indeed, recent computational work has suggested a link between preplay and
909 efficient learning, arguing that attractor dynamics can account for replay ([Corneil
910 and Gerstner, 2015](#)) or preexisting internal sequences could be used as a dynamical
911 reservoir ([Leibold, 2020](#)). In work by [Cazin et al. \(2019\)](#), the framework of reservoir
912 computing is used to model the PFC that is shown to integrate replayed sequences
913 into larger sequence assemblies that can be recalled.

914 Preplay has also been observed in neuroscience, with some studies reporting ap-
915 parent “preplay” of place cell sequences before the environment was ever experienced
916 ([Dragoi and Tonegawa, 2011, 2013](#)). While preplay seems reminiscent of a plan-
917 ning process, most findings highlight that apparent sequentiality can also reflect
918 hippocampal cell assemblies that are connected in a way that constrains sequential
919 firing, even prior to experience of a new maze. Nevertheless, other findings indicate
920 that previous experience is required for such spontaneous sequential activation to
921 occur ([Silva et al., 2015](#)). The extent to which the hippocampus is able to seem-
922 ingly preplay novel experiences could depend on the similarity between pre-existing
923 hippocampal representations and new memories about to be formed ([Eichenbaum,
924 2015](#)). Methodologically, this nevertheless highlights the necessity of comparing pre-
925 versus post-task replay (e.g., [Buhry et al., 2011](#)), as shown by recent research that
926 observed pre-vs.-post changes in replay can indeed be explained by cell activation
927 and firing rate correlations during experience ([Farooq et al., 2019](#)).

928 2.5 Inference and generalization

929 Although past information provides a glimpse into what we might expect in the
930 future, every new experience is different from the past in some form or another.
931 In order to use experiences effectively, agents must therefore know how to abstract
932 from their details and store, and replay, information which could generalize best to
933 future challenges. Past experiences should also be used to perform inferences that

934 give novel insights that go beyond what has been observed.

935 **2.5.1 Replay can reflect generalizable information and transition struc-** 936 **ture**

937 Apart from its role in learning and planning, recent developments in ML and neuro-
938 science research suggest that replay also contributes to inference and generalization
939 (for previous reviews and perspectives, see [Kumaran, 2012](#); [Kumaran and McClelland, 2012](#); [Cazé et al., 2018](#); [Herszage and Censor, 2018](#); [Lewis et al., 2018](#); [Momennejad, 2020](#)). One theme in this domain has been to build artificial agents that
942 learn generative models from experience, which can then be used to infer new con-
943 nections based on latent structural rules ([Evans and Burgess, 2019](#)), infer the correct
944 context when given new data ([Stoianov et al., 2020](#)) and generalize information to
945 new tasks to mitigate performance losses ([Shin et al., 2017](#)). In the model proposed
946 by [Stoianov et al. \(2020\)](#), for instance, trajectories through a maze are used to learn
947 a generative model, which can produce new trajectories consistent with the current
948 maze structure during offline periods. As new mazes are learned, novel trajectories
949 continue to be generated offline, but from all the mazes that have been experienced,
950 preventing information about any one maze from being lost (similar to our consid-
951 erations about forgetting in section 2.2). The hierarchical structure of the model
952 results in trajectories being clustered into distinct maze contexts, which allows maze
953 categories to be inferred when presented with new data. Unlike replay used in other
954 contexts, the model by [Stoianov et al. \(2020\)](#) does not suggest that prioritized re-
955 play helps to improve behavioral outcomes. Generative replay that was prioritized
956 based on how surprising observations were under the generative model increased the
957 number of reactivation events that contained important goal locations but did not
958 further improve inference performance ([Stoianov et al., 2020](#)).

959 The clustering of trajectories seen in the model above is related to a broader
960 theoretical view, which has emphasized that separate encoding of transition informa-
961 tion and sensory information during learning will allow knowledge about transitions
962 to be reused across situations with structural similarities but new sensory specifics
963 ([Behrens et al., 2018](#); [Baram et al., 2020](#); [Whittington et al., 2020](#)). Because re-
964 play provides a strong candidate mechanism for learning about transition structure
965 ([Stoianov et al., 2020](#)), replay of abstract (sensory-independent) transition informa-

tion could help to build representations of task structure that can be generalized and used to guide behaviour in new sensory environments (Liu et al., 2019) or combined with sensory observations to make inferences about the current environment (Evans and Burgess, 2019; Stoianov et al., 2020).

Another major computational approach has focused on replay as a mechanism to learn successor representations, a predictive representation that reflects the expected future visitation of states, given the current state (Dayan, 1993). Unlike the one-step transition matrices that are known as models in model-based RL, the successor matrix can reflect non-adjacent dependencies. This allows the agent to understand relationships between a state and multiple successor states, knowledge which can be used to solve inference problems, such as finding the shortest path to a new reward location (Russek et al., 2017; Momennejad et al., 2017). The eigenvectors of a successor matrix can also partition the environment into clusters that help planning (Stachenfeld et al., 2017). Critically, replay of past experiences could be used to update the successor matrix during offline periods (Russek et al., 2017), resonating with the general theme of using replay for model updating (e.g. Aubin et al., 2018). In recent work, Russek et al. (2017) proposed replay of (s_t, a_t, s_{t+1}) -tuples that are prioritized by recency, in which rewards are not needed to update the successor matrix. Using this approach, it was shown that learning SRs with offline replay gives an agent unique benefits compared to agents without replay. In particular, the fast updating of successor states through replay allowed the agent to quickly infer policy updates needed to adapt to changes in the task environment, like a new barrier, that affect the state transition structure.

A potential role of replay in generalization and inference has also been suggested by neuroscientific studies. In a recent study from Barron et al. (2020), hippocampal cells selective to cues and rewarding outcomes that have not been directly experienced together, but whose relationship can be inferred based on sensory pre-conditioning (Brogden, 1939), were found to be co-active during SWR events (Barron et al., 2020). These cells also tended to be reactivated during SWR events in a specific order, with reward selective cells reactivated prior to cue selective cells akin to backward replay. Using MEG, researchers have discovered that visual stimuli are reactivated in a non-experienced order that was based on prior learning of a rule about how items should be reported (Liu et al., 2019). This indicates that sequential reactivation is

999 able to combine prior learning with new sensory inputs to produce behavior relevant
1000 to new environments. Recordings in the hippocampus and PFC have also shown
1001 that hippocampal place cells are reactivated with subsets of prefrontal cells that
1002 encode generalizable task elements (Yu et al., 2018) and that at least some medial
1003 PFC neurons involved in replay have generalized firing fields that cover multiple
1004 starting locations or multiple goal locations within a maze (Kaefer et al., 2020),
1005 suggesting that replay could also contribute to generalization through coordinating
1006 the appropriate reactivation of PFC neurons. Finally, neuroscientific studies have
1007 found supporting evidence for SRs, which have similar properties to place fields,
1008 skewing in the opposite direction of travel and over-representing goal locations, while
1009 the eigenvectors of SRs can account for entorhinal grid cells in some spatial contexts
1010 (Stachenfeld et al., 2017). Hippocampal-entorhinal fMRI signals have also been
1011 shown to reflect relationships between successive non-spatial objects organized in a
1012 graph (Schapiro et al., 2013; Garvert et al., 2017), consistent with the SR (Stachenfeld
1013 et al., 2017).

1014 2.6 Representation learning

1015 A final important aspect for understanding replay in minds and machines concerns
1016 states, the internal representations agents use to describe their environment. In this
1017 last section, we highlight findings indicating that replay is not specific to spatial
1018 locations or sensory observations, but might instead involve task-dependent state
1019 representations. We argue that replaying states has unique benefits as opposed to
1020 replaying only observations. Moreover, we speculate that replay might also have a
1021 role in *learning* the representations that guide behavior. As such, replay could offer
1022 a window into the operations the brain performs to craft useful representations of
1023 the possible task states.

1024 2.6.1 Replay can reflect state representations

1025 The internal states of an agent are a major determinant of its success (see Box 1).
1026 In most environments, sensory input alone will be neither fully necessary nor fully
1027 complete for predicting outcomes. It contains too much task-irrelevant information,
1028 and what is needed to determine the best action can often not directly be observed
1029 (a property called “partial observability”). Even when doing a mundane task such

1030 as crossing the street, there will be many perceived aspects you can safely ignore
1031 (the color of the cars, the behavior of passers-by, etc.), but also factors that are very
1032 important for your decision that might *not* be in your current sensory input, such
1033 as your expectation that cars can appear quickly from behind a sharp bend. Hence,
1034 representing sensory input alone is often insufficient as a state representation. As
1035 [Dayan \(1993\)](#) has put it: “difficult problems can be rendered trivial if looked at in
1036 the correct way” (p. 613).

1037 Moreover, since the agent does not know what the true states of the environment
1038 are, *learning* useful state representations constitutes a major challenge ([Bengio et al.,](#)
1039 [2013](#); [Niv, 2019](#)). This learning involves focusing attention on task-relevant dimen-
1040 sions (e.g., [Niv et al., 2015](#); [Leong et al., 2017](#)), representing non-observable context,
1041 such as past events, in combination with current observations (e.g., [Wilson et al.,](#)
1042 [2014](#); [Schuck et al., 2016, 2018](#)), and leveraging similarity among the states to deter-
1043 mine which experiences might reflect the same hidden causes and which information
1044 can be generalized (e.g., [Gershman and Niv, 2010](#)).

1045 Many replay algorithms for DNNs store past observations, and during replay in-
1046 ternally convert observations into a suitable feature space using a previously learned
1047 transformation, such as a convolutional network (e.g., [Mnih et al., 2015](#)). But the
1048 benefits of directly storing internal representations for replay are increasingly ac-
1049 knowledged ([Kapturowski et al., 2019](#); [Isken et al., 2020](#); [Caccia et al., 2019](#); [Hayes](#)
1050 [et al., 2019](#); [van de Ven et al., 2020](#); [Pellegrini et al., 2019](#); [Hayes et al., 2021](#)).
1051 Amongst others, storing internal representations is often more memory efficient ([Is-](#)
1052 [cen et al., 2020](#); [Hayes et al., 2019](#)), while observations can still be recreated from
1053 compressed internal representations if they are needed ([van de Ven et al., 2020](#)).
1054 Moreover, representational replay can capture unobservable context that was neces-
1055 sary to process a given observation when it was made ([Kapturowski et al., 2019](#)).

1056 However, representational replay, also called state replay, comes with its own set
1057 of challenges. During online processing, the problem of partial observability can be
1058 mitigated by combining long short-term memory (LSTM) ([Hochreiter and Schmid-](#)
1059 [huber, 1997](#)) with DNNs, an architecture known as Deep Recurrent Q-Networks
1060 ([Hausknecht and Stone, 2015](#)). Because the agent’s internal state representations
1061 in these networks depend on the history of previous observations, replay of past
1062 observations risks being out of context, and replay of internal states becomes neces-

sary. Yet, as agents continue to learn from their experience, the way external inputs are mapped onto internal states changes too; in consequence, the way observations were represented internally in the past might be outdated, a phenomenon known as representational drift. It has therefore been suggested that while replaying internal representations, offline learning should be regularized in a way that captures the amount of representational drift since the replay episode (Pomponi et al., 2020; Balaji et al., 2020). The problem of representational drift will be most severe for RNNs, where observational replay will not lead to useful updates if recreated internal states do not match the internal states when the observation was made originally (e.g., Kapturowski et al., 2019). Although initial research has suggested that simply “zeroing” the agents internal state at the start of a replay event is useful in some circumstances (Hausknecht and Stone, 2015), this makes learning longer temporal dependencies more difficult. Accordingly, Kapturowski et al. (2019) have shown that it is beneficial for an agent to store its own past internal states and re-initialize the appropriate state at the start of a replay event. To account for representational drift, a part of the stored state sequences can first be replayed without updating, to reach a more appropriate internal state, and only the remainder of the sequence is then used for offline learning. In sum, replay can be beneficial for learning if it involves not only past observations but also past states.

In animals, several findings indicate that a large variety of representations, including non-spatial sensory as well as state-like representations, might be replayed. First, the firing of hippocampal “place” cells can reflect a number of non-spatial aspects of the environment, if they are task-relevant, such as sounds (Aronov et al., 2017), time (MacDonald et al., 2011), accumulated evidence for a choice (Nieh et al., 2021), or successor representations (Stachenfeld et al., 2017), but see O’Keefe and Krupic (2021). In fact, findings by Cabral et al. (2014) show that hippocampal neurons in mice flexibly switch between representations of spatial or temporal aspects of a task, depending on which strategy was needed to solve it. More directly, one fMRI study by Schuck and Niv (2019) has found that sequential hippocampal replay during post-task rest reflected the non-spatial states of a sequential decision-making task. Importantly, observed transitions between decoded replay events were best explained by replay of states that include non-observable task aspects, such as information from the previous trial, rather than by replay of sensory features of the

task stimuli alone. This study therefore provides direct evidence for the idea that replay involves state representations that are optimized for the operation of RL algorithms. In an MEG-study by [Liu et al. \(2019\)](#), human participants first learned an abstract rule governing how objects should be ordered in a sequence and later replayed a novel set of objects according to the learned rule rather than in order of experience. Replayed sequences consisted of factorized representations of sensory objects, the identity of the sequence they belonged to, as well as the position within that sequence, supporting the notion that replay is not limited to one kind of information. Moreover, [Jadhav et al. \(2012\)](#) showed that disruption of SWRs in a spatial alternation task impaired navigation when it required unobservable knowledge of the previous trial, thus hinting at the activation of state representations rather than observations during replay.

Interestingly, much evidence in neuroscience indicates that replay involves multiple representations which are reactivated in parallel, possibly suggesting that observations might be recreated at the time of replay ([van de Ven et al., 2020](#)). These representations reflect visual ([Ji and Wilson, 2006](#); [Wittkuhn and Schuck, 2021](#)), auditory ([Rothschild et al., 2016](#)) or grid-like ([Ólafsdóttir et al., 2016, 2017](#); [O’Neill et al., 2017](#)) information. These reactivated offline and online representations might interact, as it has been observed for the case of hippocampus and OFC ([Schuck and Niv, 2019](#)). This interaction between the OFC and the hippocampus (for reviews, see [Wikenheiser and Redish, 2015a](#); [Wikenheiser and Schoenbaum, 2016](#)) is particularly interesting given that the OFC might store an agent’s task state representations ([Schuck et al., 2016](#), see also [Kaplan et al., 2017](#)). Disruption of the medial PFC particularly attenuated components of hippocampal theta sequences representing the current location of the animal ([Schmidt et al., 2019](#)) and suppression of hippocampal input to the OFC impaired the integration of task state structure ([Wikenheiser et al., 2017](#)). Conversely, disruption of SWRs during sleep impaired the integrity of hippocampal maps but they re-emerged following re-learning ([Gridchyn et al., 2020](#)) suggesting that relevant maps are stored in brain areas other than the hippocampus ([Niethard and Born, 2020](#)). Despite these interactions, it should also be noted that a number of investigations have shown replay events outside the hippocampus need not be coordinated with hippocampal activity ([O’Neill et al., 2017](#); [Kaefer et al., 2020](#); [Wittkuhn and Schuck, 2021](#)). In sum, hippocampal “place cell” firing can re-

1129 flect a variety of non-spatial but task-relevant aspects (e.g., [Aronov et al., 2017](#)),
1130 replay occurs in a wide variety of interacting brain areas that reflect an animals' un-
1131 derstanding of what is task-relevant, and replay has also be found to directly reflect
1132 partially observable task states ([Schuck and Niv, 2019](#)).

1133 2.6.2 Can replay support learning useful representations?

1134 Simultaneous replay on different levels of representation, including states and sensory
1135 observations, might convey benefits beyond those discussed so far; it might help to
1136 build better state representations. One interesting instance concerns successor rep-
1137 resentations (SRs), a form of state representation that provides an efficient way to
1138 incorporate knowledge about the transitions between states into the state definition.
1139 Computational work by [Russek et al. \(2017\)](#) has shown that SRs can also be learned
1140 and updated through replay. More generally, information about state transitions can
1141 give rise to further graph analytical insights that are known to provide a good basis
1142 for state representations ([Mahadevan and Maggioni, 2007](#)). Sequential replay is a
1143 natural match as a mechanism to learn states that encode transitional information.
1144 Possibly, it could also be used to extract graph properties from experienced transi-
1145 tion structures, such as bottleneck states, which then become integrated into state
1146 representations. A similar approach has been proposed by [Eysenbach et al. \(2019\)](#),
1147 who used replay to infer graph representations that can be used for planning. Other
1148 work has highlighted that representations which predict latent embeddings of future
1149 observations are particularly useful ([Guo et al., 2020](#)). An evaluation of predictive-
1150 ness could therefore be an important contribution of replay to state representation
1151 learning.

1152 More speculatively, coordinated replay across several levels might serve as a mech-
1153 anism to identify which aspects of sensory observations exhibit transitions that are
1154 uncorrelated with state transitions of stored outcomes. Note that RL models bene-
1155 fit from transition information between states, but they could be affected adversely
1156 if transitions of task-irrelevant aspects influenced the agent's internal model. For
1157 example, representing states as specific locations in physical space will result in a
1158 transition matrix that is different from a transition matrix of more abstract task
1159 states, but if spatial position is irrelevant to the task at hand, then transitions be-
1160 tween locations could be harmful for learning and planning. If replay can be used to

1161 find unattended aspects of sensory observations that correlate with reward or rele-
1162 vant transitions, in turn it might be used to determine which dimensions of observed
1163 input are task-irrelevant (see e.g. [Schuck et al., 2015](#), for an example of how recogniz-
1164 ing correlations in the environment could lead to changes in state representations).
1165 To the best of our knowledge, this idea has not been evaluated yet.

1166 Other evidence suggests that the role of replay for state learning could go beyond
1167 information about observed transitions. SRs, for instance, can be extended to deal
1168 with partially observable task environments ([Vértes and Sahani, 2019](#)). Caselles-
1169 Dupré and colleagues proposed another interesting account that involves variational
1170 autoencoders (VAEs) ([Caselles-Dupré et al., 2018, 2019](#)). Building on earlier work
1171 that used generative models to circumvent the memory requirements of observational
1172 replay ([Shin et al., 2017](#)), [Caselles-Dupré et al. \(2019\)](#) proposed storing latent rep-
1173 resentations rather than observations, and using past experiences in this form to
1174 continually train a VAE that acts as a state model. Importantly, only by replaying
1175 past episodes can the VAE learn to form a state representation that allows the agent
1176 to act efficiently across more than one environment.

1177 In the brain, only some evidence so far suggests that replay can change state rep-
1178 resentations. [Schuck and Niv \(2019\)](#) observed that replay in the hippocampus during
1179 rest was related to better decodability of partially observable state representations
1180 from the OFC during the task. Moreover, decoding of state representations in the
1181 OFC increased over time, suggesting that representation learning continued during
1182 the task, and perhaps was related to replay. Although this evidence is correlational,
1183 it hints at a relationship between replay and state representation learning. Yet, much
1184 work is left to do, and uncovering representational changes during or following replay
1185 will require new analytical approaches which for instance do not use localizer tasks.
1186 It is therefore still unclear, whether state learning mechanisms provide a realistic
1187 account for biological replay.

1188 3 Not everything can be solved by replay

1189 In this review, we have outlined the myriad ways in which replay-like mechanisms
1190 can support intelligent behavior. But can the function of replay really be so broad,
1191 or has replay simply become a scientific bandwagon? One part of the problem is
1192 that replay has no definition that is universally agreed on by the whole scientific

community, and given the popularity of the topic, this has led to subsuming a vast variety of phenomena under the same term (but see [Genzel et al., 2020](#), for a consensus statement). Moreover, the search for an inclusive understanding across the ML and neuroscience communities has probably led to further broadening of the concept. Against this backdrop, the concept of replay has occupied a large share of the scientific study of memory, and memory undoubtedly is a very fundamental building block of intelligent behavior in minds and machines alike.

Although our review covers the broad range of functions associated with replay, it is meant as an attempt to differentiate the debate about the topic. We believe that the field needs to be careful with not overburdening a single concept. In this spirit, we hope to have elucidated how, for instance, interleaved replay of past task experiences differs from replay observed during planning, or coordinated replay during offline periods. These differences can be both computational and implementational in nature: replay in these scenarios presumably serves a different function, and it is implemented in the brain, and in computers, in different ways.

In addition, we would like to underline that despite replay’s continued popularity in ML, many state-of-the-art techniques exist which do not use replay. Efficient learning can be achieved without replay, for instance using Asynchronous Advantage Actor-Critic (A3C; [Mnih et al., 2016](#)) or on-policy policy gradient optimization (V-MPO; [Song et al., 2019](#)). Moreover, novel transformer models ([Vaswani et al., 2017](#)), which dispense of the need for convolutional and recurrent computations, have emerged as a powerful framework for solving complex tasks such as language processing ([Dai et al., 2019](#)) or RL ([Parisotto et al., 2019](#)). Some transformer models incorporate replay ([Wu et al., 2020](#)), but many powerful transformers have been proposed which do not require replay, including for RL problems ([Parisotto et al., 2019](#)).

Where replay ends and other forms of memory access start is often unclear as well. Consider, for instance, approaches in which agents rely directly on specific single episodes for behavioral control, such as in the context of *episodic RL* (for recent reviews, see e.g., [Gershman and Daw, 2017](#); [Botvinick et al., 2019](#)). During episodic RL, specific single episodes are stored in memory and retrieved to directly determine behavior when the same or a similar situation is encountered again ([Lengyel and Dayan, 2007](#); [Gershman and Daw, 2017](#); [Botvinick et al., 2019](#)). In humans, the

retrieval of single experiences in decision-making is associated with the hippocampus (Lee et al., 2015; Bornstein and Norman, 2017; Wimmer and Büchel, 2020) and reinstatement of information from past choice trials at decision-time biases present choices towards decisions made previously in the reinstated context (Duncan and Shohamy, 2016; Bornstein et al., 2017; Bornstein and Norman, 2017). To what extent these retrievals of single episodes are supported by sequential replay remains an open question.

Moreover, even if complex memory computations are needed to solve a task, external memory architectures, such as MERLIN (Wayne et al., 2018), can store past experiences in a memory buffer and learn how to read out only relevant experiences when needed. Memory storage in this model can still be efficient as the model can learn to store lower dimensional state representations instead of raw observations, and memory access is targeted to only the currently needed past information. MERLIN has been shown to outperform the LSTM architectures discussed above. But is MERLIN a replay mechanism? In some ways yes, and in others not. But a more important question is which predictions the algorithm makes, and whether they might fit neuroscientific observations. This can be nicely illustrated in the case of the MERLIN algorithm: although this model is computationally distinct from the “traditional” replay-based architectures, MERLIN predicts sophisticated reactivation phenomena. In a task in which the agent had to navigate to a goal location, for instance, the agent’s memory read-out alternated between the subgoals along the way to the goal. In our opinion, asking whether this prediction is true in the brain, and in which environments such a mechanism could be helpful, rather than labeling it as replay or not replay, would be the most fruitful way forward. This also illustrates that replay is not a single testable theory, but rather a framework within which memory, planning and imagination-related functions, as well as their relationships, can be understood.

4 Conclusion and outlook

In this review, we have summarized the literature on replay in neuroscience and ML to showcase which computational benefits biological and artificial agents can gain from replaying previous experience. We have discussed five main computational benefits that, although overlapping, provide useful categories for thinking about what might

1258 motivate an agent to employ replay: faster learning and increased data efficiency, less
1259 forgetting, the reorganization of experience, planning and generalization. In addition
1260 we have argued that replayed content is much richer than a sequence of locations,
1261 and could reflect the agent's current state representation. State representations are
1262 often task- and context-dependent, being influenced by a range of factors, including
1263 the goal-relevant aspects of the agents observations, the transition structure of states,
1264 the location, number and value of goal locations and the motivational and metabolic
1265 state of the animal. We have argued that RL theory provides useful guidance to
1266 understand which form state representations might take in a given task, and which
1267 implications a particular state representation would have for an agent's behavior.
1268 Finally, we have discussed how replay might not only reflect but could help the agent
1269 to learn those states to begin with. While many questions in particular regarding the
1270 latter idea still remain, considering these factors will greatly help to determine what
1271 replayed representations represent and how replay updates decision-making policies
1272 that are used to control behavior.

1273 5 Declaration of competing interest

1274 The authors declare no competing interests.

1275 6 Acknowledgements

1276 This work was supported by an Independent Max Planck Research Group grant
1277 awarded to N.W.S by the Max Planck Society (M.TN.A.BILD0004), a Starting Grant
1278 awarded to N.W.S by the European Union (ERC-2019-StG REPLAY-852669), and a
1279 Humboldt Research Fellowship awarded to S.H.M. by the Alexander von Humboldt
1280 Foundation. We thank members of the Max Planck Research Group NeuroCode for
1281 helpful discussions about the contents of this manuscript. We thank Anika Löwe for
1282 feedback on a previous version of this manuscript. L.W. is a pre-doctoral fellow of the
1283 International Max Planck Research School on Computational Methods in Psychiatry
1284 and Ageing Research (IMPRS COMP2PSYCH). The participating institutions are
1285 the Max Planck Institute for Human Development, Berlin, Germany, and University
1286 College London, London, UK. For more information, see [https://www.mps-ucl-](https://www.mps-ucl-centre.mpg.de/en/comp2psych)
1287 [centre.mpg.de/en/comp2psych](https://www.mps-ucl-centre.mpg.de/en/comp2psych).

1288 References

- 1289 R. Ellen Ambrose, Brad E. Pfeiffer, and David J. Foster. Reverse replay of hippocampal place cells
1290 is uniquely modulated by changing reward. *Neuron*, 91(5):1124–1136, Sep 2016. ISSN 0896-6273.
1291 [10.1016/j.neuron.2016.07.047](https://doi.org/10.1016/j.neuron.2016.07.047). URL <http://dx.doi.org/10.1016/j.neuron.2016.07.047>.
- 1292 Seiichiro Amemiya and Aaron David Redish. Manipulating decisiveness in decision making: Effects
1293 of clonidine on hippocampal search strategies. *Journal of Neuroscience*, 36(3):814–827, 2016. ISSN
1294 0270-6474. [10.1523/JNEUROSCI.2595-15.2016](https://doi.org/10.1523/JNEUROSCI.2595-15.2016). URL [https://www.jneurosci.org/content/36/](https://www.jneurosci.org/content/36/3/814)
1295 [3/814](https://www.jneurosci.org/content/36/3/814).
- 1296 David Andre, Nir Friedman, and Ronald Parr. Generalized prioritized sweeping. In *Advances in Neural*
1297 *Information Processing Systems*. MIT Press, 1998.
- 1298 Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob
1299 McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. Hindsight experience replay.
1300 In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett,
1301 editors, *Advances in Neural Information Processing Systems 30*, pages 5048–5058. Curran Associates,
1302 Inc., 2017. URL <http://papers.nips.cc/paper/7090-hindsight-experience-replay.pdf>.
- 1303 Thomas Anthony, Zheng Tian, and David Barber. Thinking fast and slow with deep learning and tree
1304 search. *arXiv e-prints*, art. arXiv:1705.08439, May 2017.
- 1305 James W. Antony and Anna C. Schapiro. Active and effective replay: systems consolidation reconsid-
1306 ered again. *Nature Reviews Neuroscience*, 20(8):506–507, 2019. [10.1038/s41583-019-0191-8](https://doi.org/10.1038/s41583-019-0191-8). URL
1307 <https://doi.org/10.1038/s41583-019-0191-8>.
- 1308 Dmitriy Aronov, Rhino Nevers, and David W. Tank. Mapping of a non-spatial dimension by the
1309 hippocampal–entorhinal circuit. *Nature*, 543(7647):719–722, Mar 2017. ISSN 1476-4687. [10.1038/na-](https://doi.org/10.1038/nature21692)
1310 [ture21692](https://doi.org/10.1038/nature21692). URL <http://dx.doi.org/10.1038/nature21692>.
- 1311 Craig Atkinson, Brendan McCane, Lech Szymanski, and Anthony Robins. Pseudo-rehearsal:
1312 Achieving deep reinforcement learning without catastrophic forgetting. *Neurocomputing*, 428:
1313 291–307, 2021. ISSN 0925-2312. <https://doi.org/10.1016/j.neucom.2020.11.050>. URL [https:](https://www.sciencedirect.com/science/article/pii/S0925231220318439)
1314 [//www.sciencedirect.com/science/article/pii/S0925231220318439](https://www.sciencedirect.com/science/article/pii/S0925231220318439).
- 1315 Lise Aubin, Mehdi Khamassi, and Benoît Girard. Prioritized sweeping neural dynaq with multiple
1316 predecessors, and hippocampal replays. In Vasiliki Vouloutsi, José Halloy, Anna Mura, Michael
1317 Mangan, Nathan Lepora, Tony J. Prescott, and Paul F.M.J. Verschure, editors, *Biomimetic and*
1318 *Biohybrid Systems*, pages 16–27, Cham, 2018. Springer International Publishing. ISBN 978-3-319-
1319 95972-6.
- 1320 N. Axmacher, C. E. Elger, and J. Fell. Ripples in the medial temporal lobe are relevant for human
1321 memory consolidation. *Brain*, 131(7):1806–1817, May 2008. ISSN 1460-2156. [10.1093/brain/awn103](https://doi.org/10.1093/brain/awn103).
1322 URL <http://dx.doi.org/10.1093/brain/awn103>.

1323 Akram Bakkour, Daniela J Palombo, Ariel Zylberberg, Yul HR Kang, Allison Reid, Mieke Verfaellie,
1324 Michael N Shadlen, and Daphna Shohamy. The hippocampus supports deliberation during value-
1325 based decisions. *eLife*, 8:e46080, jul 2019. ISSN 2050-084X. [10.7554/eLife.46080](https://doi.org/10.7554/eLife.46080). URL <https://doi.org/10.7554/eLife.46080>.
1326

1327 Yogesh Balaji, Mehrdad Farajtabar, Dong Yin, Alex Mott, and Ang Li. The Effectiveness of Memory
1328 Replay in Large Scale Continual Learning. *arXiv e-prints*, art. arXiv:2010.02418, October 2020.

1329 Alon Boaz Baram, Timothy Howard Muller, Hamed Nili, Mona Maria Garvert, and Timo-
1330 thy Edward John Behrens. Entorhinal and ventromedial prefrontal cortices abstract and
1331 generalize the structure of reinforcement learning problems. *Neuron*, 2020. ISSN 0896-
1332 6273. [10.1016/j.neuron.2020.11.024](https://doi.org/10.1016/j.neuron.2020.11.024). URL <http://www.sciencedirect.com/science/article/pii/S0896627320309296>.
1333

1334 Bengi Baran, Jessica Wilson, and Rebecca M. C. Spencer. REM-dependent repair of competitive
1335 memory suppression. *Experimental Brain Research*, 203(2):471–477, 2010. [10.1007/s00221-010-2242-2](https://doi.org/10.1007/s00221-010-2242-2). URL <https://doi.org/10.1007/s00221-010-2242-2>.
1336

1337 Helen C. Barron, Hayley M. Reeve, Renée S. Koolschijn, Pavel V. Perestenko, Anna Shpektor, Hamed
1338 Nili, Roman Rothaermel, Natalia Campo-Urriza, Jill X. O'Reilly, David M. Bannerman, Timo-
1339 thy E.J. Behrens, and David Dupret. Neuronal computation underlying inferential reasoning in
1340 humans and mice. *Cell*, 183(1):228 – 243.e21, 2020. ISSN 0092-8674. [10.1016/j.cell.2020.08.035](https://doi.org/10.1016/j.cell.2020.08.035).
1341 URL <https://doi.org/10.1016/j.cell.2020.08.035>.

1342 Jr Bartol, Thomas M, Cailey Bromer, Justin Kinney, Michael A Chirillo, Jennifer N Bourne, Kristen M
1343 Harris, and Terrence J Sejnowski. Nanoconnectomic upper bound on the variability of synaptic
1344 plasticity. *eLife*, 4:e10778, nov 2015. ISSN 2050-084X. [10.7554/eLife.10778](https://doi.org/10.7554/eLife.10778). URL <https://doi.org/10.7554/eLife.10778>.
1345

1346 Timothy E.J. Behrens, Timothy H. Muller, James C.R. Whittington, Shirley Mark, Alon B. Baram,
1347 Kimberly L. Stachenfeld, and Zeb Kurth-Nelson. What is a cognitive map? organizing knowledge for
1348 flexible behavior. *Neuron*, 100(2):490–509, Oct 2018. ISSN 0896-6273. [10.1016/j.neuron.2018.10.002](https://doi.org/10.1016/j.neuron.2018.10.002).
1349 URL <http://dx.doi.org/10.1016/j.neuron.2018.10.002>.

1350 Richard Bellman. *Dynamic Programming*. Princeton University Press, 1957.

1351 Jacob L. S. Bellmund, Peter Gärdenfors, Edvard I. Moser, and Christian F. Doeller. Navigating cog-
1352 nition: Spatial codes for human thinking. *Science*, 362(6415), 2018. ISSN 0036-8075. [10.1126/sci-ence.aat6766](https://science.sciencemag.org/content/362/6415/eaat6766). URL <https://science.sciencemag.org/content/362/6415/eaat6766>.
1353

1354 Daniel Bendor and Matthew A Wilson. Biasing the content of hippocampal replay during sleep.
1355 *Nature Neuroscience*, 15(10):1439–1444, Sep 2012. ISSN 1546-1726. [10.1038/nn.3203](https://doi.org/10.1038/nn.3203). URL <http://dx.doi.org/10.1038/nn.3203>.
1356

1357 Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives.
1358 *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, Aug 2013. ISSN
1359 2160-9292. [10.1109/tpami.2013.50](https://doi.org/10.1109/TPAMI.2013.50). URL <http://dx.doi.org/10.1109/TPAMI.2013.50>.

1360 Baburam Bhattarai, Jong Won Lee, and Min Whan Jung. Distinct effects of reward and naviga-
1361 tion history on hippocampal forward and reverse replays. *Proceedings of the National Academy*

1362 of Sciences, page 201912533, Dec 2019. ISSN 1091-6490. [10.1073/pnas.1912533117](https://doi.org/10.1073/pnas.1912533117). URL <http://dx.doi.org/10.1073/pnas.1912533117>.

1363

1364 Chris M Bird. How do we remember events? *Current Opinion in Behavioral Sciences*, 32:120–125,
 1365 Apr 2020. ISSN 2352-1546. [10.1016/j.cobeha.2020.01.020](https://doi.org/10.1016/j.cobeha.2020.01.020). URL <http://dx.doi.org/10.1016/j.cobeha.2020.01.020>.

1366

1367 T. V. P. Bliss and G. L. Collingridge. A synaptic model of memory: long-term potentiation in the
 1368 hippocampus. *Nature*, 361(6407):31–39, 1993. [10.1038/361031a0](https://doi.org/10.1038/361031a0). URL <https://doi.org/10.1038/361031a0>.

1369

1370 Aaron M Bornstein and Kenneth A Norman. Reinstated episodic context guides sampling-based
 1371 decisions for reward. *Nature Neuroscience*, 20(7):997–1003, 2017. [10.1038/nn.4573](https://doi.org/10.1038/nn.4573). URL <https://doi.org/10.1038/nn.4573>.

1372

1373 Aaron M. Bornstein, Mel W. Khaw, Daphna Shohamy, and Nathaniel D. Daw. Reminders of
 1374 past choices bias decisions for reward in humans. *Nature Communications*, 8(1):15958, 2017.
 1375 [10.1038/ncomms15958](https://doi.org/10.1038/ncomms15958). URL <https://doi.org/10.1038/ncomms15958>.

1376

1377 Roberto Bottini and Christian F. Doeller. Knowledge across reference frames: Cognitive maps
 1378 and image spaces. *Trends in Cognitive Sciences*, 24(8):606 – 619, 2020. ISSN 1364-
 1379 6613. <https://doi.org/10.1016/j.tics.2020.05.008>. URL <http://www.sciencedirect.com/science/article/pii/S1364661320301327>.

1380

1381 Matthew Botvinick, Sam Ritter, Jane X. Wang, Zeb Kurth-Nelson, Charles Blundell, and Demis Has-
 1382 sabis. Reinforcement learning, fast and slow. *Trends in Cognitive Sciences*, 23(5):408 – 422, 2019.
 1383 ISSN 1364-6613. <https://doi.org/10.1016/j.tics.2019.02.006>. URL <http://www.sciencedirect.com/science/article/pii/S1364661319300610>.

1384

1385 Matthew Botvinick, Jane X. Wang, Will Dabney, Kevin J. Miller, and Zeb Kurth-Nelson. Deep
 1386 reinforcement learning and its neuroscientific implications. *Neuron*, Jul 2020. ISSN 0896-6273.
[10.1016/j.neuron.2020.06.014](https://doi.org/10.1016/j.neuron.2020.06.014). URL <http://dx.doi.org/10.1016/j.neuron.2020.06.014>.

1387

1388 Anatol Bragin, Jerome Engel, Charles L. Wilson, Itzhak Fried, and György Buzsáki. High-
 1389 frequency oscillations in human brain. *Hippocampus*, 9(2):137–142, 1999. ISSN 1098-1063.
 1390 [10.1002/\(sici\)1098-1063\(1999\)9:2<137::aid-hipo5>3.0.co;2-0](https://doi.org/10.1002/(sici)1098-1063(1999)9:2<137::aid-hipo5>3.0.co;2-0). URL [http://dx.doi.org/10.1002/\(SICI\)1098-1063\(1999\)9:2<137::AID-HIP05>3.0.CO;2-0](http://dx.doi.org/10.1002/(SICI)1098-1063(1999)9:2<137::AID-HIP05>3.0.CO;2-0).

1391

1392 Wilfred J Brogden. Sensory pre-conditioning. *Journal of Experimental Psychology*, 25(4):323, 1939.
[10.1037/h0058944](https://doi.org/10.1037/h0058944). URL <https://doi.org/10.1037/h0058944>.

1393

1394 T. I. Brown, V. A. Carr, K. F. LaRocque, S. E. Favila, A. M. Gordon, B. Bowles, J. N. Bailenson,
 1395 and A. D. Wagner. Prospective representation of navigational goals in the human hippocampus.
 1396 *Science*, 352(6291):1323–1326, Jun 2016. ISSN 1095-9203. [10.1126/science.aaf0784](https://doi.org/10.1126/science.aaf0784). URL <http://dx.doi.org/10.1126/science.aaf0784>.

1397

1398 Iva K. Brunec, Morris Moscovitch, and Morgan D. Barense. Boundaries shape cognitive representations
 1399 of spaces and events. *Trends in Cognitive Sciences*, 22(7):637–650, Jul 2018. ISSN 1364-6613.
[10.1016/j.tics.2018.03.013](https://doi.org/10.1016/j.tics.2018.03.013). URL <http://dx.doi.org/10.1016/j.tics.2018.03.013>.

- 1400 Randy Lee Buckner. The role of the hippocampus in prediction and imagination. *Annual Review of*
 1401 *Psychology*, 61(1):27–48, 2010. [10.1146/annurev.psych.60.110707.163508](https://doi.org/10.1146/annurev.psych.60.110707.163508). URL [https://doi.org/](https://doi.org/10.1146/annurev.psych.60.110707.163508)
 1402 [10.1146/annurev.psych.60.110707.163508](https://doi.org/10.1146/annurev.psych.60.110707.163508).
- 1403 Laure Buhry, Amir H. Azizi, and Sen Cheng. Reactivation, replay, and preplay: How it might all
 1404 fit together. *Neural Plasticity*, pages 1–11, 2011. ISSN 1687-5443. [10.1155/2011/203462](https://doi.org/10.1155/2011/203462). URL
 1405 <http://dx.doi.org/10.1155/2011/203462>.
- 1406 Daniel Bush, Caswell Barry, Daniel Manson, and Neil Burgess. Using grid cells for navigation. *Neuron*,
 1407 87(3):507 – 520, 2015. ISSN 0896-6273. <https://doi.org/10.1016/j.neuron.2015.07.006>. URL [http:](http://www.sciencedirect.com/science/article/pii/S0896627315006285)
 1408 [//www.sciencedirect.com/science/article/pii/S0896627315006285](http://www.sciencedirect.com/science/article/pii/S0896627315006285).
- 1409 G. Buzsáki. Two-stage model of memory trace formation: A role for “noisy” brain states. *Neuroscience*,
 1410 31(3):551–570, Jan 1989. ISSN 0306-4522. [10.1016/0306-4522\(89\)90423-5](https://doi.org/10.1016/0306-4522(89)90423-5). URL [http://dx.doi.](http://dx.doi.org/10.1016/0306-4522(89)90423-5)
 1411 [org/10.1016/0306-4522\(89\)90423-5](http://dx.doi.org/10.1016/0306-4522(89)90423-5).
- 1412 György Buzsáki. Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and
 1413 planning. *Hippocampus*, 25(10):1073–1188, Sep 2015. ISSN 1050-9631. [10.1002/hipo.22488](https://doi.org/10.1002/hipo.22488). URL
 1414 <http://dx.doi.org/10.1002/hipo.22488>.
- 1415 Henrique O. Cabral, Martin Vinck, Celine Fouquet, Cyriel M.A. Pennartz, Laure Rondi-Reig, and
 1416 Francesco P. Battaglia. Oscillatory dynamics and place field maps reflect hippocampal ensem-
 1417 ble processing of sequence and place memory under nmda receptor control. *Neuron*, 81(2):402–
 1418 415, 2014. ISSN 0896-6273. <https://doi.org/10.1016/j.neuron.2013.11.010>. URL [https://www.](https://www.sciencedirect.com/science/article/pii/S0896627313010441)
 1419 [sciencedirect.com/science/article/pii/S0896627313010441](https://www.sciencedirect.com/science/article/pii/S0896627313010441).
- 1420 Lucas Caccia, Eugene Belilovsky, Massimo Caccia, and Joelle Pineau. Online learned continual com-
 1421 pression with adaptive quantization modules. *arXiv e-prints*, page arXiv:1911.08019, November
 1422 2019.
- 1423 Alyssa A. Carey, Youki Tanaka, and Matthijs A. A. van der Meer. Reward revaluation biases hippocam-
 1424 pal replay content away from the preferred outcome. *Nature Neuroscience*, 22(9):1450–1459, Aug
 1425 2019. ISSN 1546-1726. [10.1038/s41593-019-0464-6](https://doi.org/10.1038/s41593-019-0464-6). URL [http://dx.doi.org/10.1038/s41593-](http://dx.doi.org/10.1038/s41593-019-0464-6)
 1426 [019-0464-6](http://dx.doi.org/10.1038/s41593-019-0464-6).
- 1427 Margaret F Carr, Shantanu P Jadhav, and Loren M Frank. Hippocampal replay in the awake state:
 1428 a potential substrate for memory consolidation and retrieval. *Nature Neuroscience*, 14(2):147–153,
 1429 Jan 2011. ISSN 1546-1726. [10.1038/nn.2732](https://doi.org/10.1038/nn.2732). URL <http://dx.doi.org/10.1038/nn.2732>.
- 1430 Hugo Caselles-Dupré, Michael Garcia-Ortiz, and David Filliat. Continual state representation learning
 1431 for reinforcement learning using generative replay. *arXiv e-prints*, page arXiv:1810.03880, Oct 2018.
 1432 URL <https://ui.adsabs.harvard.edu/abs/2018arXiv181003880C>.
- 1433 Hugo Caselles-Dupré, Michael Garcia-Ortiz, and David Filliat. S-TRIGGER: Continual State Rep-
 1434 resentation Learning via Self-Triggered Generative Replay. *arXiv e-prints*, art. arXiv:1902.09434,
 1435 February 2019.
- 1436 Romain Cazé, Mehdi Khamassi, Lise Aubin, and Benoît Girard. Hippocampal replays under the
 1437 scrutiny of reinforcement learning models. *Journal of Neurophysiology*, 120(6):2877–2896, Dec 2018.
 1438 ISSN 1522-1598. [10.1152/jn.00145.2018](https://doi.org/10.1152/jn.00145.2018). URL <http://dx.doi.org/10.1152/jn.00145.2018>.

1439 Nicolas Cazin, Martin Llofriu Alonso, Pablo Scleidorovich Chiodi, Tatiana Pelc, Bruce Harland, Alfredo
1440 Weitzenfeld, Jean-Marc Fellous, and Peter Ford Dominey. Reservoir computing model of prefrontal
1441 cortex creates novel combinations of previous navigation sequences from hippocampal place-cell
1442 replay with spatial reward propagation. *PLOS Computational Biology*, 15(7):e1006624, Jul 2019.
1443 ISSN 1553-7358. [10.1371/journal.pcbi.1006624](https://doi.org/10.1371/journal.pcbi.1006624). URL [http://dx.doi.org/10.1371/journal.pcbi.](http://dx.doi.org/10.1371/journal.pcbi.1006624)
1444 [1006624](http://dx.doi.org/10.1371/journal.pcbi.1006624).

1445 Arslan Chaudhry, Puneet K. Dokania, Thalaiyasingam Ajanthan, and Philip H. S. Torr. Riemannian
1446 walk for incremental learning: Understanding forgetting and intransigence. *arXiv e-prints*, page
1447 arXiv:1801.10112, January 2018.

1448 Sen Cheng and Loren M. Frank. New experiences enhance coordinated neural activity in the hip-
1449 pocampus. *Neuron*, 57(2):303 – 313, 2008. ISSN 0896-6273. [10.1016/j.neuron.2007.11.035](https://doi.org/10.1016/j.neuron.2007.11.035). URL
1450 <http://www.sciencedirect.com/science/article/pii/S0896627307010306>.

1451 David Clewett, Sarah DuBrow, and Lila Davachi. Transcending time in the brain: How event mem-
1452 ories are constructed from experience. *Hippocampus*, 29(3):162–183, Feb 2019. ISSN 1050-9631.
1453 [10.1002/hipo.23074](https://doi.org/10.1002/hipo.23074). URL <http://dx.doi.org/10.1002/hipo.23074>.

1454 Neal J. Cohen and Howard Eichenbaum. *Memory, amnesia, and the hippocampal system*. MIT Press,
1455 1993. ISBN 9780262531320.

1456 A. O. Constantinescu, J. X. O’Reilly, and T. E. J. Behrens. Organizing conceptual knowledge in
1457 humans with a gridlike code. *Science*, 352(6292):1464–1468, Jun 2016. ISSN 1095-9203. [10.1126/sci-](https://doi.org/10.1126/science.aaf0941)
1458 [ence.aaf0941](https://doi.org/10.1126/science.aaf0941). URL <http://dx.doi.org/10.1126/science.aaf0941>.

1459 Dane S Corneil and Wulfram Gerstner. Attractor network dynamics enable preplay and rapid
1460 path planning in maze-like environments. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama,
1461 and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28.
1462 Curran Associates, Inc., 2015. URL [https://proceedings.neurips.cc/paper/2015/file/](https://proceedings.neurips.cc/paper/2015/file/e515df0d202ae52fceb14295743063b-Paper.pdf)
1463 [e515df0d202ae52fceb14295743063b-Paper.pdf](https://proceedings.neurips.cc/paper/2015/file/e515df0d202ae52fceb14295743063b-Paper.pdf).

1464 Jozsef Csicsvari, Joseph O’Neill, Kevin Allen, and Timothy Senior. Place-selective firing contributes
1465 to the reverse-order reactivation of CA1 pyramidal cells during sharp waves in open-field explo-
1466 ration. *European Journal of Neuroscience*, 26(3):704–716, Jul 2007. ISSN 0953-816X. [10.1111/j.1460-](https://doi.org/10.1111/j.1460-9568.2007.05684.x)
1467 [9568.2007.05684.x](https://doi.org/10.1111/j.1460-9568.2007.05684.x). URL <http://dx.doi.org/10.1111/j.1460-9568.2007.05684.x>.

1468 Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V. Le, and Ruslan Salakhutdinov.
1469 Transformer-XL: Attentive language models beyond a fixed-length context. *arXiv e-prints*, art.
1470 arXiv:1901.02860, January 2019.

1471 Thomas J. Davidson, Fabian Kloosterman, and Matthew A. Wilson. Hippocampal replay of extended
1472 experience. *Neuron*, 63(4):497–507, Aug 2009. ISSN 0896-6273. [10.1016/j.neuron.2009.07.027](https://doi.org/10.1016/j.neuron.2009.07.027). URL
1473 <http://dx.doi.org/10.1016/j.neuron.2009.07.027>.

1474 Nathaniel D Daw, Yael Niv, and Peter Dayan. Uncertainty-based competition between prefrontal and
1475 dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 8(12):1704–1711, 2005.
1476 [10.1038/nn1560](https://doi.org/10.1038/nn1560). URL <https://doi.org/10.1038/nn1560>.

- 1477 Peter Dayan. Improving generalization for temporal difference learning: The successor representation.
1478 *Neural Computation*, 5(4):613–624, Jul 1993. ISSN 1530-888X. [10.1162/neco.1993.5.4.613](https://doi.org/10.1162/neco.1993.5.4.613). URL
1479 <http://dx.doi.org/10.1162/neco.1993.5.4.613>.
- 1480 Xinyi Deng, Shizhe Chen, Marielena Sosa, Mattias P. Karlsson, Xue-Xin Wei, and Loren M.
1481 Frank. A variable clock underlies internally generated hippocampal sequences. *bioRxiv*, Apr 2020.
1482 [10.1101/2020.04.10.035980](https://doi.org/10.1101/2020.04.10.035980). URL <http://dx.doi.org/10.1101/2020.04.10.035980>.
- 1483 Eric L. Denovellis, Anna K. Gillespie, Michael E. Coulter, Marielena Sosa, Jason E. Chung, Uri T.
1484 Eden, and Loren M. Frank. Hippocampal replay of experience at real-world speeds. *bioRxiv*, 2020.
1485 [10.1101/2020.10.20.347708](https://doi.org/10.1101/2020.10.20.347708). URL [https://www.biorxiv.org/content/early/2020/10/23/2020.](https://www.biorxiv.org/content/early/2020/10/23/2020.10.20.347708)
1486 [10.20.347708](https://www.biorxiv.org/content/early/2020/10/23/2020.10.20.347708).
- 1487 L. Deuker, J. Olligs, J. Fell, T. A. Kranz, F. Mormann, C. Montag, M. Reuter, C. E. Elger, and
1488 N. Axmacher. Memory consolidation by replay of stimulus-specific neural activity. *Journal of*
1489 *Neuroscience*, 33(49):19373–19383, Dec 2013. ISSN 1529-2401. [10.1523/jneurosci.0414-13.2013](https://doi.org/10.1523/JNEUROSCI.0414-13.2013). URL
1490 <http://dx.doi.org/10.1523/JNEUROSCI.0414-13.2013>.
- 1491 Kamran Diba and György Buzsáki. Forward and reverse hippocampal place-cell sequences during
1492 ripples. *Nature Neuroscience*, 10(10):1241–1242, Sep 2007. ISSN 1546-1726. [10.1038/nn1961](https://doi.org/10.1038/nn1961). URL
1493 <http://dx.doi.org/10.1038/nn1961>.
- 1494 Susanne Diekelmann and Jan Born. The memory function of sleep. *Nature Reviews Neuroscience*, 11(2):
1495 114–126, Jan 2010. ISSN 1471-0048. [10.1038/nrn2762](https://doi.org/10.1038/nrn2762). URL <http://dx.doi.org/10.1038/nrn2762>.
- 1496 Ray J. Dolan and Peter Dayan. Goals and habits in the brain. *Neuron*, 80(2):312 – 325, 2013.
1497 ISSN 0896-6273. [10.1016/j.neuron.2013.09.007](https://doi.org/10.1016/j.neuron.2013.09.007). URL [http://www.sciencedirect.com/science/](http://www.sciencedirect.com/science/article/pii/S0896627313008052)
1498 [article/pii/S0896627313008052](http://www.sciencedirect.com/science/article/pii/S0896627313008052).
- 1499 G. Dragoi and S. Tonegawa. Distinct preplay of multiple novel spatial experiences in the rat. *Pro-*
1500 *ceedings of the National Academy of Sciences*, 110(22):9100–9105, May 2013. ISSN 1091-6490.
1501 [10.1073/pnas.1306031110](https://doi.org/10.1073/pnas.1306031110). URL <http://dx.doi.org/10.1073/pnas.1306031110>.
- 1502 George Dragoi and Susumu Tonegawa. Preplay of future place cell sequences by hippocampal cellular
1503 assemblies. *Nature*, 469(7330):397–401, Dec 2011. ISSN 1476-4687. [10.1038/nature09633](https://doi.org/10.1038/nature09633). URL
1504 <http://dx.doi.org/10.1038/nature09633>.
- 1505 Sarah DuBrow, Nina Rouhani, Yael Niv, and Kenneth A Norman. Does mental context
1506 drift or shift? *Current Opinion in Behavioral Sciences*, 17:141–146, 2017. ISSN 2352-
1507 1546. [10.1016/j.cobeha.2017.08.003](https://doi.org/10.1016/j.cobeha.2017.08.003). URL [https://www.sciencedirect.com/science/article/](https://www.sciencedirect.com/science/article/pii/S2352154616302820)
1508 [pii/S2352154616302820](https://www.sciencedirect.com/science/article/pii/S2352154616302820).
- 1509 Katherine D Duncan and Daphna Shohamy. Memory states influence value-based decisions. *Journal*
1510 *of Experimental Psychology: General*, 145(11):1420, 2016.
- 1511 David Dupret, Joseph O’Neill, Barty Pleydell-Bouverie, and Jozsef Csicsvari. The reorganization and
1512 reactivation of hippocampal maps predict spatial memory performance. *Nature Neuroscience*, 13(8):
1513 995–1002, Jul 2010. ISSN 1546-1726. [10.1038/nn.2599](https://doi.org/10.1038/nn.2599). URL <http://dx.doi.org/10.1038/nn.2599>.

1514 Valérie Ego-Stengel and Matthew A. Wilson. Disruption of ripple-associated hippocampal activity
1515 during rest impairs spatial learning in the rat. *Hippocampus*, 20(1):1–10, 2010. ISSN 1098-1063.
1516 [10.1002/hipo.20707](https://doi.org/10.1002/hipo.20707). URL <http://dx.doi.org/10.1002/hipo.20707>.

1517 Howard Eichenbaum. Does the hippocampus preplay memories? *Nature Neuroscience*, 18(12):1701–
1518 1702, 2015. [10.1038/nn.4180](https://doi.org/10.1038/nn.4180). URL <https://doi.org/10.1038/nn.4180>.

1519 Eran Eldar, Gaëlle Lièvre, Peter Dayan, and Raymond J Dolan. The roles of online and offline replay
1520 in planning. *eLife*, 9, Jun 2020. ISSN 2050-084X. [10.7554/elife.56911](https://doi.org/10.7554/elife.56911). URL [http://dx.doi.org/](http://dx.doi.org/10.7554/elife.56911)
1521 [10.7554/elife.56911](https://doi.org/10.7554/elife.56911).

1522 Russell A Epstein, Eva Zita Patai, Joshua B Julian, and Hugo J Spiers. The cognitive map in humans:
1523 spatial navigation and beyond. *Nature Neuroscience*, 20(11):1504–1513, Nov 2017. ISSN 1546-1726.
1524 [10.1038/nn.4656](https://doi.org/10.1038/nn.4656). URL <http://dx.doi.org/10.1038/nn.4656>.

1525 Uğur M. Erdem and Michael Hasselmo. A goal-directed spatial navigation model using forward
1526 trajectory planning based on grid cells. *European Journal of Neuroscience*, 35(6):916–931, 2012.
1527 [10.1111/j.1460-9568.2012.08015.x](https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1460-9568.2012.08015.x). URL [https://onlinelibrary.wiley.com/doi/abs/10.1111/j.](https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1460-9568.2012.08015.x)
1528 [1460-9568.2012.08015.x](https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1460-9568.2012.08015.x).

1529 Oxana Eschenko, Wiâm Ramadan, Matthias Mölle, Jan Born, and Susan J. Sara. Sustained increase in
1530 hippocampal sharp-wave ripple activity during slow-wave sleep after learning. *Learning & Memory*,
1531 15(4):222–228, 2008. [10.1101/lm.726008](https://doi.org/10.1101/lm.726008). URL [http://learnmem.cshlp.org/content/15/4/222.](http://learnmem.cshlp.org/content/15/4/222.abstract)
1532 [abstract](http://learnmem.cshlp.org/content/15/4/222.abstract).

1533 Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Volodymyr Mnih, Tom Ward, Yotam
1534 Doron, Vlad Firoiu, Tim Harley, Iain Dunning, Shane Legg, and Koray Kavukcuoglu. IMPALA:
1535 Scalable Distributed Deep-RL with Importance Weighted Actor-Learner Architectures. *arXiv*
1536 *e-prints*, page arXiv:1802.01561, February 2018. URL [https://ui.adsabs.harvard.edu/abs/](https://ui.adsabs.harvard.edu/abs/2018arXiv180201561E)
1537 [2018arXiv180201561E](https://ui.adsabs.harvard.edu/abs/2018arXiv180201561E).

1538 D. R. Euston, M. Tatsuno, and B. L. McNaughton. Fast-forward playback of recent memory se-
1539 quences in prefrontal cortex during sleep. *Science*, 318(5853):1147–1150, Nov 2007. ISSN 1095-9203.
1540 [10.1126/science.1148979](https://doi.org/10.1126/science.1148979). URL <http://dx.doi.org/10.1126/science.1148979>.

1541 Talfan Evans and Neil Burgess. Coordinated hippocampal-entorhinal replay as structural
1542 inference. *Advances in Neural Information Processing Systems*, pages 1729–1741, 2019.
1543 URL [http://papers.nips.cc/paper/8450-coordinated-hippocampal-entorhinal-replay-as-](http://papers.nips.cc/paper/8450-coordinated-hippocampal-entorhinal-replay-as-structural-inference.pdf)
1544 [structural-inference.pdf](http://papers.nips.cc/paper/8450-coordinated-hippocampal-entorhinal-replay-as-structural-inference.pdf).

1545 Benjamin Eysenbach, Ruslan Salakhutdinov, and Sergey Levine. Search on the replay buffer: Bridging
1546 planning and reinforcement learning. *arXiv e-prints*, page arXiv:1906.05253, June 2019. URL [https:](https://arxiv.org/abs/1906.05253)
1547 [//arxiv.org/abs/1906.05253](https://arxiv.org/abs/1906.05253).

1548 Usman Farooq, Jeremie Sibille, Kefei Liu, and George Dragoi. Strengthened temporal coordination
1549 within pre-existing sequential cell assemblies supports trajectory replay. *Neuron*, 103(4):719–733.e7,
1550 Aug 2019. ISSN 0896-6273. [10.1016/j.neuron.2019.05.040](https://doi.org/10.1016/j.neuron.2019.05.040). URL [http://dx.doi.org/10.1016/j.](http://dx.doi.org/10.1016/j.neuron.2019.05.040)
1551 [neuron.2019.05.040](https://doi.org/10.1016/j.neuron.2019.05.040).

1552 Serra E. Favila, Avi J. H. Chanale, and Brice A. Kuhl. Experience-dependent hippocampal pat-
 1553 tern differentiation prevents interference during subsequent learning. *Nature Communications*, 7(1):
 1554 11066, 2016. [10.1038/ncomms11066](https://doi.org/10.1038/ncomms11066). URL <https://doi.org/10.1038/ncomms11066>.

1555 Serra E. Favila, Hongmi Lee, and Brice A. Kuhl. Transforming the concept of memory reactivation.
 1556 *Trends in Neurosciences*, 43(12):939–950, 2021/01/14 2020. [10.1016/j.tins.2020.09.006](https://doi.org/10.1016/j.tins.2020.09.006). URL <https://doi.org/10.1016/j.tins.2020.09.006>.

1558 William Fedus, Prajit Ramachandran, Rishabh Agarwal, Yoshua Bengio, Hugo Larochelle, Mark
 1559 Rowland, and Will Dabney. Revisiting Fundamentals of Experience Replay. *arXiv e-prints*, art.
 1560 arXiv:2007.06700, July 2020. URL <https://ui.adsabs.harvard.edu/abs/2020arXiv200706700F>.

1561 Gordon B Feld and Jan Born. Sculpting memory during sleep: concurrent consolida-
 1562 tion and forgetting. *Current Opinion in Neurobiology*, 44:20–27, 2017. ISSN 0959-4388.
 1563 <https://doi.org/10.1016/j.conb.2017.02.012>. URL [https://www.sciencedirect.com/science/](https://www.sciencedirect.com/science/article/pii/S0959438816301726)
 1564 [article/pii/S0959438816301726](https://www.sciencedirect.com/science/article/pii/S0959438816301726). Neurobiology of Sleep.

1565 Graham Findlay, Giulio Tononi, and Chiara Cirelli. The evolving view of replay and its functions in
 1566 wake and sleep. *SLEEP Advances*, 01 2021. ISSN 2632-5012. [10.1093/sleepadvances/zpab002](https://doi.org/10.1093/sleepadvances/zpab002). URL <https://doi.org/10.1093/sleepadvances/zpab002>. zpab002.

1568 Timo Flesch, Jan Balaguer, Ronald Dekker, Hamed Nili, and Christopher Summerfield. Comparing
 1569 continual task learning in minds and machines. *Proceedings of the National Academy of Sciences*, 115
 1570 (44):E10313–E10322, 2018. ISSN 0027-8424. [10.1073/pnas.1800755115](https://doi.org/10.1073/pnas.1800755115). URL [https://www.pnas.](https://www.pnas.org/content/115/44/E10313)
 1571 [org/content/115/44/E10313](https://www.pnas.org/content/115/44/E10313).

1572 David J. Foster. Replay comes of age. *Annual Review of Neuroscience*, 40(1):581–602, 2017.
 1573 [10.1146/annurev-neuro-072116-031538](https://doi.org/10.1146/annurev-neuro-072116-031538). URL [https://doi.org/10.1146/annurev-neuro-072116-](https://doi.org/10.1146/annurev-neuro-072116-031538)
 1574 [031538](https://doi.org/10.1146/annurev-neuro-072116-031538).

1575 David J. Foster and James J. Knierim. Sequence learning and the role of the hippocampus in ro-
 1576 dent navigation. *Current Opinion in Neurobiology*, 22(2):294–300, Apr 2012. ISSN 0959-4388.
 1577 [10.1016/j.conb.2011.12.005](https://doi.org/10.1016/j.conb.2011.12.005). URL <http://dx.doi.org/10.1016/j.conb.2011.12.005>.

1578 David J. Foster and Matthew A. Wilson. Reverse replay of behavioural sequences in hippocampal place
 1579 cells during the awake state. *Nature*, 440(7084):680–683, Feb 2006. ISSN 1476-4687. [10.1038/na-](https://doi.org/10.1038/nature04587)
 1580 [ture04587](https://doi.org/10.1038/nature04587). URL <http://dx.doi.org/10.1038/nature04587>.

1581 R French. Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3(4):
 1582 128–135, 1999. [10.1016/S1364-6613\(99\)01294-2](https://doi.org/10.1016/S1364-6613(99)01294-2). URL [https://doi.org/10.1016/S1364-6613\(99\)](https://doi.org/10.1016/S1364-6613(99)01294-2)
 1583 [01294-2](https://doi.org/10.1016/S1364-6613(99)01294-2).

1584 Marianne Fyhn, Torkel Hafting, Alessandro Treves, May-Britt Moser, and Edvard I. Moser. Hip-
 1585 pocampal remapping and grid realignment in entorhinal cortex. *Nature*, 446(7132):190–194, 2007.
 1586 [10.1038/nature05601](https://doi.org/10.1038/nature05601). URL <https://doi.org/10.1038/nature05601>.

1587 Christopher Gagne, Peter Dayan, and Sonia J Bishop. When planning to survive goes wrong: predicting
 1588 the future and replaying the past in anxiety and PTSD. *Current Opinion in Behavioral Sciences*,
 1589 24:89–95, Dec 2018. ISSN 2352-1546. [10.1016/j.cobeha.2018.03.013](https://doi.org/10.1016/j.cobeha.2018.03.013). URL [http://dx.doi.org/10.](http://dx.doi.org/10.1016/j.cobeha.2018.03.013)
 1590 [1016/j.cobeha.2018.03.013](http://dx.doi.org/10.1016/j.cobeha.2018.03.013).

- 1591 Javier García, Fern, and o Fernández. A comprehensive survey on safe reinforcement learning. *Jour-*
1592 *nal of Machine Learning Research*, 16(42):1437–1480, 2015. URL [http://jmlr.org/papers/v16/](http://jmlr.org/papers/v16/garcia15a.html)
1593 [garcia15a.html](http://jmlr.org/papers/v16/garcia15a.html).
- 1594 Matthew P. H. Gardner, Geoffrey Schoenbaum, and Samuel J. Gershman. Rethinking dopamine as
1595 generalized prediction error. *Proceedings of the Royal Society B: Biological Sciences*, 285(1891):
1596 20181645, 2018. [10.1098/rspb.2018.1645](https://doi.org/10.1098/rspb.2018.1645). URL [https://royalsocietypublishing.org/doi/abs/](https://royalsocietypublishing.org/doi/abs/10.1098/rspb.2018.1645)
1597 [10.1098/rspb.2018.1645](https://royalsocietypublishing.org/doi/abs/10.1098/rspb.2018.1645).
- 1598 Mona M Garvert, Raymond J Dolan, and Timothy EJ Behrens. A map of abstract relational
1599 knowledge in the human hippocampal–entorhinal cortex. *eLife*, 6, Apr 2017. ISSN 2050-084X.
1600 [10.7554/eLife.17086](https://doi.org/10.7554/eLife.17086). URL <http://dx.doi.org/10.7554/eLife.17086>.
- 1601 P. Gaussier, A. Revel, J. P. Banquet, and V. Babeau. From view cells and place cells to cognitive map
1602 learning: processing stages of the hippocampal system. *Biological Cybernetics*, 86(1):15–28, 2002.
1603 [10.1007/s004220100269](https://doi.org/10.1007/s004220100269). URL <https://doi.org/10.1007/s004220100269>.
- 1604 Lisa Genzel, George Dragoi, Loren Frank, Karunesh Ganguly, Liset de la Prida, Brad Pfeif-
1605 fer, and Edwin Robertson. A consensus statement: defining terms for reactivation analysis.
1606 *Philosophical Transactions of the Royal Society B: Biological Sciences*, 375(1799):20200001, 2020.
1607 [10.1098/rstb.2020.0001](https://doi.org/10.1098/rstb.2020.0001). URL [https://royalsocietypublishing.org/doi/abs/10.1098/rstb.](https://royalsocietypublishing.org/doi/abs/10.1098/rstb.2020.0001)
1608 [2020.0001](https://royalsocietypublishing.org/doi/abs/10.1098/rstb.2020.0001).
- 1609 Jason L. Gerrard, Hemant Kudrimoti, Bruce L. McNaughton, and Carol A. Barnes. Reactivation of
1610 hippocampal ensemble activity patterns in the aging rat. *Behavioral Neuroscience*, 115(6):1180–1192,
1611 2001. ISSN 0735-7044. [10.1037/0735-7044.115.6.1180](https://doi.org/10.1037/0735-7044.115.6.1180). URL [http://dx.doi.org/10.1037/0735-](http://dx.doi.org/10.1037/0735-7044.115.6.1180)
1612 [7044.115.6.1180](http://dx.doi.org/10.1037/0735-7044.115.6.1180).
- 1613 Samuel J. Gershman and Nathaniel D. Daw. Reinforcement learning and episodic memory in hu-
1614 mans and animals: An integrative framework. *Annual Review of Psychology*, 68(1):101–128, Jan
1615 2017. ISSN 1545-2085. [10.1146/annurev-psych-122414-033625](https://doi.org/10.1146/annurev-psych-122414-033625). URL [http://dx.doi.org/10.1146/](http://dx.doi.org/10.1146/annurev-psych-122414-033625)
1616 [annurev-psych-122414-033625](http://dx.doi.org/10.1146/annurev-psych-122414-033625).
- 1617 Samuel J Gershman and Yael Niv. Learning latent structure: Carving nature at its joints. *Current*
1618 *Opinion in Neurobiology*, 20(2):251–256, 2010. [10.1016/j.conb.2010.02.008](https://doi.org/10.1016/j.conb.2010.02.008). URL [https://www.](https://www.sciencedirect.com/science/article/pii/S0959438810000309)
1619 [sciencedirect.com/science/article/pii/S0959438810000309](https://www.sciencedirect.com/science/article/pii/S0959438810000309).
- 1620 Samuel J. Gershman, Angela Radulescu, Kenneth A. Norman, and Yael Niv. Statistical computations
1621 underlying the dynamics of memory updating. *PLOS Computational Biology*, 10(11):1–13, 11 2014.
1622 [10.1371/journal.pcbi.1003939](https://doi.org/10.1371/journal.pcbi.1003939). URL <https://doi.org/10.1371/journal.pcbi.1003939>.
- 1623 Gabrielle Girardeau, Karim Benchenane, Sidney I Wiener, György Buzsáki, and Michaël B Zugaro.
1624 Selective suppression of hippocampal ripples impairs spatial memory. *Nature Neuroscience*, 12(10):
1625 1222–1223, Sep 2009. ISSN 1546-1726. [10.1038/nn.2384](https://doi.org/10.1038/nn.2384). URL [http://dx.doi.org/10.1038/nn.](http://dx.doi.org/10.1038/nn.2384)
1626 [2384](http://dx.doi.org/10.1038/nn.2384).
- 1627 Gabrielle Girardeau, Ingrid Inema, and György Buzsáki. Reactivations of emotional memory in
1628 the hippocampus–amygdala system during sleep. *Nature Neuroscience*, 20(11):1634–1642, 2017.
1629 [10.1038/nn.4637](https://doi.org/10.1038/nn.4637). URL <https://doi.org/10.1038/nn.4637>.

1630 Stephen N Gomperts, Fabian Kloosterman, and Matthew A Wilson. VTA neurons coordinate
 1631 with the hippocampal reactivation of spatial experience. *eLife*, 4, Oct 2015. ISSN 2050-084X.
 1632 [10.7554/eLife.05360](https://doi.org/10.7554/eLife.05360). URL <http://dx.doi.org/10.7554/eLife.05360>.

1633 Igor Gridchyn, Philipp Schoenenberger, Joseph O'Neill, and Jozsef Csicsvari. Assembly-specific dis-
 1634 ruption of hippocampal replay leads to selective memory deficit. *Neuron*, Feb 2020. ISSN 0896-6273.
 1635 [10.1016/j.neuron.2020.01.021](https://doi.org/10.1016/j.neuron.2020.01.021). URL <http://dx.doi.org/10.1016/j.neuron.2020.01.021>.

1636 Matthias J. Gruber, Maureen Ritchey, Shao-Fang Wang, Manoj K. Doss, and Charan Ranganath.
 1637 Post-learning hippocampal dynamics promote preferential retention of rewarding events. *Neuron*,
 1638 89(5):1110–1120, Mar 2016. ISSN 0896-6273. [10.1016/j.neuron.2016.01.017](https://doi.org/10.1016/j.neuron.2016.01.017). URL <http://dx.doi.org/10.1016/j.neuron.2016.01.017>.

1640 Tanuj Gulati, Ling Guo, Dhakshin S Ramanathan, Anitha Bodepudi, and Karunesh Ganguly. Neural
 1641 reactivations during sleep determine network credit assignment. *Nature Neuroscience*, 20(9):1277–
 1642 1284, Jul 2017. ISSN 1546-1726. [10.1038/nn.4601](https://doi.org/10.1038/nn.4601). URL <http://dx.doi.org/10.1038/nn.4601>.

1643 Xiaoxiao Guo, Satinder Singh, Honglak Lee, Richard L Lewis, and Xiaoshi Wang. Deep learning
 1644 for real-time atari game play using offline monte-carlo tree search planning. In Z. Ghahramani,
 1645 M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Informa-*
 1646 *tion Processing Systems*, volume 27. Curran Associates, Inc., 2014. URL <https://proceedings.neurips.cc/paper/2014/file/8bb88f80d334b1869781beb89f7b73be-Paper.pdf>.

1648 Zhaohan Daniel Guo, Bernardo Avila Pires, Bilal Piot, Jean-Bastien Grill, Florent Alth, Remi Munos,
 1649 and Mohammad Gheshlaghi Azar. Bootstrap latent-predictive representations for multitask rein-
 1650 forcement learning. In Hal Daum III and Aarti Singh, editors, *Proceedings of the 37th International*
 1651 *Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages
 1652 3875–3886. PMLR, 13–18 Jul 2020. URL <http://proceedings.mlr.press/v119/guo20g.html>.

1653 Anoopum S. Gupta, Matthijs A.A. van der Meer, David S. Touretzky, and Aaron David Redish. Hip-
 1654 pocampal replay is not a simple function of experience. *Neuron*, 65(5):695 – 705, 2010. ISSN 0896-
 1655 6273. [10.1016/j.neuron.2010.01.034](https://doi.org/10.1016/j.neuron.2010.01.034). URL <http://www.sciencedirect.com/science/article/pii/S0896627310000607>.

1657 Torkel Hafting, Marianne Fyhn, Sturla Molden, May-Britt Moser, and Edvard I. Moser. Microstructure
 1658 of a spatial map in the entorhinal cortex. *Nature*, 436(7052):801–806, 2005. [10.1038/nature03721](https://doi.org/10.1038/nature03721).
 1659 URL <https://doi.org/10.1038/nature03721>.

1660 Tatsuya Haga and Tomoki Fukai. Recurrent network model for learning goal-directed sequences through
 1661 reverse replay. *eLife*, 7, Jul 2018. ISSN 2050-084X. [10.7554/eLife.34171](https://doi.org/10.7554/eLife.34171). URL <http://dx.doi.org/10.7554/eLife.34171>.

1663 Oliver Hardt, Karim Nader, and Lynn Nadel. Decay happens: the role of active forgetting in memory.
 1664 *Trends in Cognitive Sciences*, 17(3):111–120, 2013. ISSN 1364-6613. [10.1016/j.tics.2013.01.001](https://doi.org/10.1016/j.tics.2013.01.001). URL
 1665 <https://www.sciencedirect.com/science/article/pii/S1364661313000132>.

1666 Christopher D. Harvey, Philip Coen, and David W. Tank. Choice-specific sequences in parietal cortex
 1667 during a virtual-navigation decision task. *Nature*, 484(7392):62–68, Mar 2012. ISSN 1476-4687.
 1668 [10.1038/nature10918](https://doi.org/10.1038/nature10918). URL <http://dx.doi.org/10.1038/nature10918>.

1669 Demis Hassabis and Eleanor A. Maguire. Deconstructing episodic memory with construction. *Trends*
1670 *in Cognitive Sciences*, 11(7):299 – 306, 2007. ISSN 1364-6613. [10.1016/j.tics.2007.05.001](https://doi.org/10.1016/j.tics.2007.05.001). URL
1671 <http://www.sciencedirect.com/science/article/pii/S1364661307001258>.

1672 Demis Hassabis, Dharshan Kumaran, Christopher Summerfield, and Matthew Botvinick.
1673 Neuroscience-inspired artificial intelligence. *Neuron*, 95(2):245–258, Jul 2017. ISSN 0896-6273.
1674 [10.1016/j.neuron.2017.06.011](https://doi.org/10.1016/j.neuron.2017.06.011). URL <http://dx.doi.org/10.1016/j.neuron.2017.06.011>.

1675 Matthew Hausknecht and Peter Stone. Deep Recurrent Q-Learning for Partially Observable MDPs.
1676 *arXiv e-prints*, page arXiv:1507.06527, July 2015. URL [https://ui.adsabs.harvard.edu/abs/](https://ui.adsabs.harvard.edu/abs/2015arXiv150706527H)
1677 [2015arXiv150706527H](https://ui.adsabs.harvard.edu/abs/2015arXiv150706527H).

1678 Tyler L. Hayes, Kushal Kafle, Robik Shrestha, Manoj Acharya, and Christopher Kanan. REMIND your
1679 neural network to prevent catastrophic forgetting. *arXiv e-prints*, page arXiv:1910.02509, October
1680 2019.

1681 Tyler L. Hayes, Giri P. Krishnan, Maxim Bazhenov, Hava T. Siegelmann, Terrence J. Sejnowski, and
1682 Christopher Kanan. Replay in deep learning: Current approaches and missing biological elements.
1683 *arXiv e-prints*, page arXiv:2104.04132, April 2021.

1684 Randolph F. Helfrich, Janna D. Lendner, Bryce A. Mander, Heriberto Guillen, Michelle Paff, Lilit
1685 Mnatsakanyan, Sumeet Vadera, Matthew P. Walker, Jack J. Lin, and Robert T. Knight. Bidirectional
1686 prefrontal-hippocampal dynamics organize information transfer during sleep in humans. *Nature*
1687 *Communications*, 10(1), Aug 2019. ISSN 2041-1723. [10.1038/s41467-019-11444-x](https://doi.org/10.1038/s41467-019-11444-x). URL [http://](http://dx.doi.org/10.1038/s41467-019-11444-x)
1688 dx.doi.org/10.1038/s41467-019-11444-x.

1689 Aaron S. Heller and Rosemary C. Bagot. Is hippocampal replay a mechanism for anxiety and de-
1690 pression? *JAMA Psychiatry*, Feb 2020. ISSN 2168-622X. [10.1001/jamapsychiatry.2019.4788](https://doi.org/10.1001/jamapsychiatry.2019.4788). URL
1691 <http://dx.doi.org/10.1001/jamapsychiatry.2019.4788>.

1692 Jasmine Herszage and Nitzan Censor. Modulation of learning and memory: A shared frame-
1693 work for interference and generalization. *Neuroscience*, 392:270 – 280, 2018. ISSN 0306-4522.
1694 [10.1016/j.neuroscience.2018.08.006](https://doi.org/10.1016/j.neuroscience.2018.08.006). URL [http://www.sciencedirect.com/science/article/](http://www.sciencedirect.com/science/article/pii/S030645221830544X)
1695 [pii/S030645221830544X](http://www.sciencedirect.com/science/article/pii/S030645221830544X).

1696 Matteo Hessel, Joseph Modayil, Hado van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan
1697 Horgan, Bilal Piot, Mohammad Azar, and David Silver. Rainbow: Combining improvements in
1698 deep reinforcement learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1),
1699 2021/02/11 2018. URL <https://ojs.aaai.org/index.php/AAAI/article/view/11796>.

1700 GE Hinton, P Dayan, BJ Frey, and RM Neal. The "wake-sleep" algorithm for unsupervised neural
1701 networks. *Science*, 268(5214):1158–1161, 1995. ISSN 0036-8075. [10.1126/science.7761831](https://doi.org/10.1126/science.7761831). URL
1702 <https://science.sciencemag.org/content/268/5214/1158>.

1703 Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735,
1704 1997. ISSN 0899-7667. [10.1162/neco.1997.9.8.1735](https://doi.org/10.1162/neco.1997.9.8.1735). URL <http://tinyurl.sfx.mpg.de/vpw2>.

1705 K. L. Hoffman and Bruce L. McNaughton. Coordinated reactivation of distributed memory traces
1706 in primate neocortex. *Science*, 297(5589):2070–2073, Sep 2002. ISSN 1095-9203. [10.1126/sci-](https://doi.org/10.1126/science.1073538)
1707 [ence.1073538](https://doi.org/10.1126/science.1073538). URL <http://dx.doi.org/10.1126/science.1073538>.

1708 Dan Horgan, John Quan, David Budden, Gabriel Barth-Maron, Matteo Hessel, Hado van Hasselt, and
 1709 David Silver. Distributed prioritized experience replay. *arXiv e-prints*, page arXiv:1803.00933, Mar
 1710 2018. URL <https://ui.adsabs.harvard.edu/abs/2018arXiv180300933H>.

1711 Øyvind Arne Høydal, Emilie Ranheim Skytøen, Sebastian Ola Andersson, May-Britt Moser, and Ed-
 1712 vard I. Moser. Object-vector coding in the medial entorhinal cortex. *Nature*, 568(7752):400–404,
 1713 2019. [10.1038/s41586-019-1077-7](https://doi.org/10.1038/s41586-019-1077-7). URL <https://doi.org/10.1038/s41586-019-1077-7>.

1714 Kinga Igloi, Giulia Gaggioni, Virginie Sterpenich, and Sophie Schwartz. A nap to recap or how reward
 1715 regulates hippocampal-prefrontal memory networks during daytime sleep in humans. *eLife*, 4:e07903,
 1716 oct 2015. ISSN 2050-084X. [10.7554/eLife.07903](https://doi.org/10.7554/eLife.07903). URL <https://doi.org/10.7554/eLife.07903>.

1717 Ahmet Iscen, Jeffrey Zhang, Svetlana Lazebnik, and Cordelia Schmid. Memory-efficient incremental
 1718 learning through feature adaptation. *arXiv e-prints*, page arXiv:2004.00713, April 2020.

1719 J. C. Jackson, A. Johnson, and Aaron David Redish. Hippocampal sharp waves and reactivation
 1720 during awake states depend on repeated sequential experience. *Journal of Neuroscience*, 26(48):
 1721 12415–12426, Nov 2006. ISSN 1529-2401. [10.1523/jneurosci.4118-06.2006](https://doi.org/10.1523/jneurosci.4118-06.2006). URL <http://dx.doi.org/10.1523/JNEUROSCI.4118-06.2006>.

1723 S. P. Jadhav, C. Kemere, P. W. German, and L. M. Frank. Awake hippocampal sharp-wave ripples
 1724 support spatial memory. *Science*, 336(6087):1454–1458, May 2012. ISSN 1095-9203. [10.1126/sci-](https://doi.org/10.1126/science.1217230)
 1725 [ence.1217230](http://dx.doi.org/10.1126/science.1217230). URL <http://dx.doi.org/10.1126/science.1217230>.

1726 Shantanu P. Jadhav, Gideon Rothschild, Demetris K. Roumis, and Loren M. Frank. Coordinated
 1727 excitation and inhibition of prefrontal ensembles during awake hippocampal sharp-wave ripple events.
 1728 *Neuron*, 90(1):113–127, Apr 2016. ISSN 0896-6273. [10.1016/j.neuron.2016.02.010](https://doi.org/10.1016/j.neuron.2016.02.010). URL <http://dx.doi.org/10.1016/j.neuron.2016.02.010>.

1730 Anna Jafarpour, Will Penny, Gareth Barnes, Robert T. Knight, and Emrah Duzel. Working mem-
 1731 ory replay prioritizes weakly attended events. *eNeuro*, 4(4):1–11, Jul 2017. ISSN 2373-2822.
 1732 [10.1523/eneuro.0171-17.2017](https://doi.org/10.1523/eneuro.0171-17.2017). URL <http://dx.doi.org/10.1523/ENEURO.0171-17.2017>.

1733 Daoyun Ji and Matthew A Wilson. Coordinated memory replay in the visual cortex and hippocampus
 1734 during sleep. *Nature Neuroscience*, 10(1):100–107, Dec 2006. ISSN 1546-1726. [10.1038/nn1825](https://doi.org/10.1038/nn1825). URL
 1735 <http://dx.doi.org/10.1038/nn1825>.

1736 Adam Johnson and Aaron David Redish. Neural ensembles in CA3 transiently encode paths forward
 1737 of the animal at a decision point. *Journal of Neuroscience*, 27(45):12176–12189, Nov 2007. ISSN
 1738 1529-2401. [10.1523/jneurosci.3761-07.2007](https://doi.org/10.1523/jneurosci.3761-07.2007). URL [http://dx.doi.org/10.1523/JNEUROSCI.3761-](http://dx.doi.org/10.1523/JNEUROSCI.3761-07.2007)
 1739 [07.2007](http://dx.doi.org/10.1523/JNEUROSCI.3761-07.2007).

1740 Hannah R. Joo and Loren M. Frank. The hippocampal sharp wave-ripple in memory retrieval for
 1741 immediate use and consolidation. *Nature Reviews Neuroscience*, 19(12):744–757, Oct 2018. ISSN
 1742 1471-0048. [10.1038/s41583-018-0077-1](https://doi.org/10.1038/s41583-018-0077-1). URL <http://dx.doi.org/10.1038/s41583-018-0077-1>.

1743 Karola Kaefer, Michele Nardin, Karel Blahna, and Jozsef Csicsvari. Replay of behavioral sequences
 1744 in the medial prefrontal cortex during rule switching. *Neuron*, Feb 2020. ISSN 0896-6273.
 1745 [10.1016/j.neuron.2020.01.015](https://doi.org/10.1016/j.neuron.2020.01.015). URL <http://dx.doi.org/10.1016/j.neuron.2020.01.015>.

1746 Lukasz Kaiser, Mohammad Babaeizadeh, Piotr Milos, Blazej Osinski, Roy H Campbell, Konrad
 1747 Czechowski, Dumitru Erhan, Chelsea Finn, Piotr Kozakowski, Sergey Levine, Afroz Mohiuddin,
 1748 Ryan Sepassi, George Tucker, and Henryk Michalewski. Model-based reinforcement learning for
 1749 atari. *arXiv e-prints*, page arXiv:1903.00374, March 2019.

1750 Raphael Kaplan, Nicolas W. Schuck, and Christian F. Doeller. The role of mental maps in decision-
 1751 making. *Trends in Neurosciences*, 40(5):256 – 259, 2017. ISSN 0166-2236. [10.1016/j.tins.2017.03.002](https://doi.org/10.1016/j.tins.2017.03.002).
 1752 URL <http://www.sciencedirect.com/science/article/pii/S0166223617300310>.

1753 Raphael Kaplan, Adrià Tauste Campo, Daniel Bush, John King, Alessandro Principe, Raphael Koster,
 1754 Miguel Ley Nacher, Rodrigo Rocamora, and Karl J. Friston. Human hippocampal theta oscilla-
 1755 tions reflect sequential dependencies during spatial planning. *Cognitive Neuroscience*, 11(3):122–131,
 1756 2020. [10.1080/17588928.2019.1676711](https://doi.org/10.1080/17588928.2019.1676711). URL <https://doi.org/10.1080/17588928.2019.1676711>.
 1757 PMID: 31617790.

1758 Steven Kapturowski, Georg Ostrovski, Will Dabney, John Quan, and Remi Munos. Recurrent expe-
 1759 rience replay in distributed reinforcement learning. In *International Conference on Learning Repre-*
 1760 *sentations*, 2019. URL <https://openreview.net/forum?id=r1lyTjAqYX>.

1761 Mattias P Karlsson and Loren M Frank. Awake replay of remote experiences in the hippocampus.
 1762 *Nature Neuroscience*, 12(7):913–918, Jun 2009. ISSN 1546-1726. [10.1038/nn.2344](https://doi.org/10.1038/nn.2344). URL <http://dx.doi.org/10.1038/nn.2344>.
 1763

1764 Kenneth Kay, Jason E. Chung, Marielena Sosa, Jonathan S. Schor, Mattias P. Karlsson, Margaret C.
 1765 Larkin, Daniel F. Liu, and Loren M. Frank. Constant sub-second cycling between representations
 1766 of possible futures in the hippocampus. *Cell*, Jan 2020. ISSN 0092-8674. [10.1016/j.cell.2020.01.014](https://doi.org/10.1016/j.cell.2020.01.014).
 1767 URL <http://dx.doi.org/10.1016/j.cell.2020.01.014>.

1768 Mehdi Khamassi and Benoît Girard. Modeling awake hippocampal reactivations with model-based
 1769 bidirectional search. *Biological Cybernetics*, Feb 2020. ISSN 1432-0770. [10.1007/s00422-020-00817-](https://doi.org/10.1007/s00422-020-00817-x)
 1770 [x](https://doi.org/10.1007/s00422-020-00817-x). URL <http://dx.doi.org/10.1007/s00422-020-00817-x>.

1771 Mehdi Khamassi and Mark D. Humphries. Integrating cortico-limbic-basal ganglia architectures for
 1772 learning model-based and model-free navigation strategies. *Frontiers in Behavioral Neuroscience*,
 1773 6, 2012. ISSN 1662-5153. [10.3389/fnbeh.2012.00079](https://doi.org/10.3389/fnbeh.2012.00079). URL [http://dx.doi.org/10.3389/fnbeh.](http://dx.doi.org/10.3389/fnbeh.2012.00079)
 1774 [2012.00079](http://dx.doi.org/10.3389/fnbeh.2012.00079).

1775 Charles King, Darrell A. Henze, Xavier Leinekugel, and György Buzsáki. Hebbian modification of a
 1776 hippocampal population pattern in the rat. *The Journal of Physiology*, 521(1):159–167, Nov 1999.
 1777 ISSN 0022-3751. [10.1111/j.1469-7793.1999.00159.x](https://doi.org/10.1111/j.1469-7793.1999.00159.x). URL [http://dx.doi.org/10.1111/j.1469-](http://dx.doi.org/10.1111/j.1469-7793.1999.00159.x)
 1778 [7793.1999.00159.x](http://dx.doi.org/10.1111/j.1469-7793.1999.00159.x).

1779 Jens G. Klinzing, Niels Niethard, and Jan Born. Mechanisms of systems memory consolidation during
 1780 sleep. *Nature Neuroscience*, 22(10):1598–1610, Aug 2019. ISSN 1546-1726. [10.1038/s41593-019-](https://doi.org/10.1038/s41593-019-0467-3)
 1781 [0467-3](https://doi.org/10.1038/s41593-019-0467-3). URL <http://dx.doi.org/10.1038/s41593-019-0467-3>.

1782 Hemant S. Kudrimoti, Carol A. Barnes, and Bruce L. McNaughton. Reactivation of hippocampal cell
 1783 assemblies: Effects of behavioral state, experience, and eeg dynamics. *The Journal of Neuroscience*,
 1784 19(10):4090–4101, May 1999. ISSN 1529-2401. [10.1523/JNEUROSCI.19-10-04090.1999](https://doi.org/10.1523/JNEUROSCI.19-10-04090.1999). URL <http://dx.doi.org/10.1523/JNEUROSCI.19-10-04090.1999>.
 1785

1786 Brice A Kuhl, Arpeet T Shah, Sarah DuBrow, and Anthony D Wagner. Resistance to forgetting
1787 associated with hippocampus-mediated reactivation during new learning. *Nature Neuroscience*, 13
1788 (4):501–506, 2010. [10.1038/nn.2498](https://doi.org/10.1038/nn.2498). URL <https://doi.org/10.1038/nn.2498>.

1789 Dharshan Kumaran. What representations and computations underpin the contribution of the hip-
1790 pocampus to generalization and inference? *Frontiers in Human Neuroscience*, 6:157, 2012. ISSN
1791 1662-5161. [10.3389/fnhum.2012.00157](https://www.frontiersin.org/article/10.3389/fnhum.2012.00157). URL [https://www.frontiersin.org/article/10.3389/](https://www.frontiersin.org/article/10.3389/fnhum.2012.00157)
1792 [fnhum.2012.00157](https://www.frontiersin.org/article/10.3389/fnhum.2012.00157).

1793 Dharshan Kumaran and James L. McClelland. Generalization through the recurrent interaction of
1794 episodic memories: A model of the hippocampal system. *Psychological Review*, 119(3):573–616, Jul
1795 2012. ISSN 0033-295X. [10.1037/a0028681](http://dx.doi.org/10.1037/a0028681). URL <http://dx.doi.org/10.1037/a0028681>.

1796 Dharshan Kumaran, Demis Hassabis, and James L. McClelland. What learning systems do intelligent
1797 agents need? Complementary learning systems theory updated. *Trends in Cognitive Sciences*, 20
1798 (7):512–534, Jul 2016. ISSN 1364-6613. [10.1016/j.tics.2016.05.004](http://dx.doi.org/10.1016/j.tics.2016.05.004). URL [http://dx.doi.org/10.](http://dx.doi.org/10.1016/j.tics.2016.05.004)
1799 [1016/j.tics.2016.05.004](http://dx.doi.org/10.1016/j.tics.2016.05.004).

1800 Zeb Kurth-Nelson, Marcos Economides, Raymond J. Dolan, and Peter Dayan. Fast sequences
1801 of non-spatial state representations in humans. *Neuron*, 91(1):194–204, 2016. ISSN 10974199.
1802 [10.1016/j.neuron.2016.05.028](http://dx.doi.org/10.1016/j.neuron.2016.05.028). URL <http://dx.doi.org/10.1016/j.neuron.2016.05.028>.

1803 C. S. Lansink, P. M. Goltstein, J. V. Lankelma, R. N. J. M. A. Joosten, B. L. McNaughton, and C. M. A.
1804 Pennartz. Preferential reactivation of motivationally relevant information in the ventral striatum.
1805 *Journal of Neuroscience*, 28(25):6372–6382, Jun 2008. ISSN 1529-2401. [10.1523/jneurosci.1054-](http://dx.doi.org/10.1523/JNEUROSCI.1054-08.2008)
1806 [08.2008](http://dx.doi.org/10.1523/JNEUROSCI.1054-08.2008). URL <http://dx.doi.org/10.1523/JNEUROSCI.1054-08.2008>.

1807 Carien S. Lansink, Pieter M. Goltstein, Jan V. Lankelma, Bruce L. McNaughton, and Cyriel M. A.
1808 Pennartz. Hippocampus leads ventral striatum in replay of place-reward information. *PLoS Biology*,
1809 7(8):e1000173, Aug 2009. ISSN 1545-7885. [10.1371/journal.pbio.1000173](http://dx.doi.org/10.1371/journal.pbio.1000173). URL [http://dx.doi.](http://dx.doi.org/10.1371/journal.pbio.1000173)
1810 [org/10.1371/journal.pbio.1000173](http://dx.doi.org/10.1371/journal.pbio.1000173).

1811 Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
1812 [10.1038/nature14539](https://doi.org/10.1038/nature14539). URL <https://doi.org/10.1038/nature14539>.

1813 Albert K. Lee and Matthew A. Wilson. Memory of sequential experience in the hippocampus during
1814 slow wave sleep. *Neuron*, 36(6):1183–1194, Dec 2002. ISSN 0896-6273. [10.1016/S0896-6273\(02\)01096-](http://dx.doi.org/10.1016/S0896-6273(02)01096-6)
1815 [6](http://dx.doi.org/10.1016/S0896-6273(02)01096-6). URL [http://dx.doi.org/10.1016/S0896-6273\(02\)01096-6](http://dx.doi.org/10.1016/S0896-6273(02)01096-6).

1816 Sang Wan Lee, John P. O’Doherty, and Shinsuke Shimojo. Neural computations mediating one-shot
1817 learning in the human brain. *PLOS Biology*, 13(4):1–36, 04 2015. [10.1371/journal.pbio.1002137](http://dx.doi.org/10.1371/journal.pbio.1002137).
1818 URL <https://doi.org/10.1371/journal.pbio.1002137>.

1819 Christian Leibold. A model for navigation in unknown environments based on a reservoir
1820 of hippocampal sequences. *Neural Networks*, 124:328–342, Apr 2020. ISSN 0893-6080.
1821 [10.1016/j.neunet.2020.01.014](http://dx.doi.org/10.1016/j.neunet.2020.01.014). URL <http://dx.doi.org/10.1016/j.neunet.2020.01.014>.

1822 Máté Lengyel and Peter Dayan. Hippocampal contributions to control: The third way. In *Proceedings*
1823 *of the 20th International Conference on Neural Information Processing Systems*, NIPS’07, pages
1824 889–896, Red Hook, NY, USA, 2007. Curran Associates Inc. ISBN 9781605603520.

1825 Yuan Chang Leong, Angela Radulescu, Reka Daniel, Vivian DeWoskin, and Yael Niv. Dynamic interac-
1826 tion between reinforcement learning and attention in multidimensional environments. *Neuron*, 93(2):
1827 451 – 463, 2017. ISSN 0896-6273. [10.1016/j.neuron.2016.12.040](https://doi.org/10.1016/j.neuron.2016.12.040). URL <http://www.sciencedirect.com/science/article/pii/S089662731631039X>.
1828

1829 Penelope A. Lewis and Daniel Bendor. How targeted memory reactivation promotes the selective
1830 strengthening of memories in sleep. *Current Biology*, 29(18):R906–R912, Sep 2019. ISSN 0960-9822.
1831 [10.1016/j.cub.2019.08.019](https://doi.org/10.1016/j.cub.2019.08.019). URL <http://dx.doi.org/10.1016/j.cub.2019.08.019>.

1832 Penelope A. Lewis, Günther Knoblich, and Gina Poe. How memory replay in sleep boosts cre-
1833 ative problem-solving. *Trends in Cognitive Sciences*, 22(6):491–503, Jun 2018. ISSN 1364-6613.
1834 [10.1016/j.tics.2018.03.009](https://doi.org/10.1016/j.tics.2018.03.009). URL <http://dx.doi.org/10.1016/j.tics.2018.03.009>.

1835 Long Ji Lin. Programming robots using reinforcement learning and teaching. In *Association for the*
1836 *Advancement of Artificial Intelligence*, pages 781–786, 1991. URL [https://www.aaai.org/Library/](https://www.aaai.org/Library/AAAI/1991/aaai91-122.php)
1837 [AAAI/1991/aaai91-122.php](https://www.aaai.org/Library/AAAI/1991/aaai91-122.php).

1838 Long-Ji Lin. Self-improving reactive agents based on reinforcement learning, planning and teaching.
1839 *Machine Learning*, 8(3-4):293–321, May 1992. ISSN 1573-0565. [10.1007/bf00992699](https://doi.org/10.1007/bf00992699). URL <http://dx.doi.org/10.1007/BF00992699>.
1840

1841 Long-Ji Lin. *Reinforcement Learning for Robots Using Neural Networks*. PhD thesis, School of Com-
1842 puter Science, USA, 1993. URL <https://dl.acm.org/doi/book/10.5555/168871>.

1843 Zachary C. Lipton, Kamyar Azizzadenesheli, Abhishek Kumar, Lihong Li, Jianfeng Gao, and Li Deng.
1844 Combating reinforcement learning’s sisyphian curse with intrinsic fear. *arXiv e-prints*, page
1845 arXiv:1611.01211, November 2016.

1846 B. Liu, X. Ye, Y. Gao, X. Dong, X. Wang, and B. Liu. Forward-looking imaginative planning framework
1847 combined with prioritized-replay double DQN. In *2019 5th International Conference on Control,*
1848 *Automation and Robotics (ICCAR)*, pages 336–341, April 2019. [10.1109/ICCAR.2019.8813352](https://doi.org/10.1109/ICCAR.2019.8813352).

1849 Yunzhe Liu, Raymond J. Dolan, Zeb Kurth-Nelson, and Timothy E.J. Behrens. Human re-
1850 play spontaneously reorganizes experience. *Cell*, 178(3):640–652, Jul 2019. ISSN 0092-8674.
1851 [10.1016/j.cell.2019.06.012](https://doi.org/10.1016/j.cell.2019.06.012). URL <http://dx.doi.org/10.1016/j.cell.2019.06.012>.

1852 Yunzhe Liu, Raymond J Dolan, Cameron Higgins, Hector Penagos, Mark W Woolrich, H Freyja
1853 Ólafsdóttir, Caswell Barry, Zeb Kurth-Nelson, and Timothy E Behrens. Temporally delayed linear
1854 modelling (TDLM) measures replay in both animals and humans. *eLife*, 10:e66917, jun 2021a. ISSN
1855 2050-084X. [10.7554/eLife.66917](https://doi.org/10.7554/eLife.66917). URL <https://doi.org/10.7554/eLife.66917>.

1856 Yunzhe Liu, Marcelo G. Mattar, Timothy E. J. Behrens, Nathaniel D. Daw, and Raymond J. Dolan.
1857 Experience replay is associated with efficient nonlocal learning. *Science*, 372(6544), 2021b. ISSN
1858 0036-8075. [10.1126/science.abf1357](https://doi.org/10.1126/science.abf1357). URL [https://science.sciencemag.org/content/372/6544/](https://science.sciencemag.org/content/372/6544/eabf1357)
1859 [eabf1357](https://science.sciencemag.org/content/372/6544/eabf1357).

1860 Kenway Louie and Matthew A. Wilson. Temporally structured replay of awake hippocampal ensemble
1861 activity during rapid eye movement sleep. *Neuron*, 29(1):145–156, Jan 2001. ISSN 0896-6273.
1862 [10.1016/s0896-6273\(01\)00186-6](https://doi.org/10.1016/s0896-6273(01)00186-6). URL [http://dx.doi.org/10.1016/S0896-6273\(01\)00186-6](http://dx.doi.org/10.1016/S0896-6273(01)00186-6).

- 1863 Mantas Lukoševičius and Herbert Jaeger. Reservoir computing approaches to recurrent neu-
 1864 ral network training. *Computer Science Review*, 3(3):127–149, 2009. ISSN 1574-0137.
 1865 <https://doi.org/10.1016/j.cosrev.2009.03.005>. URL [https://www.sciencedirect.com/science/](https://www.sciencedirect.com/science/article/pii/S1574013709000173)
 1866 [article/pii/S1574013709000173](https://www.sciencedirect.com/science/article/pii/S1574013709000173).
- 1867 Christopher J. MacDonald, Kyle Q. Lepage, Uri T. Eden, and Howard Eichenbaum. Hippocampal
 1868 “time cells” bridge the gap in memory for discontinuous events. *Neuron*, 71(4):737 – 749, 2011.
 1869 ISSN 0896-6273. [10.1016/j.neuron.2011.07.012](https://doi.org/10.1016/j.neuron.2011.07.012). URL [http://www.sciencedirect.com/science/](http://www.sciencedirect.com/science/article/pii/S089662731100609X)
 1870 [article/pii/S089662731100609X](http://www.sciencedirect.com/science/article/pii/S089662731100609X).
- 1871 J. C. Magee. A synaptically controlled, associative signal for hebbian plasticity in hippocampal neurons.
 1872 *Science*, 275(5297):209–213, Jan 1997. ISSN 1095-9203. [10.1126/science.275.5297.209](https://doi.org/10.1126/science.275.5297.209). URL <http://dx.doi.org/10.1126/science.275.5297.209>.
 1873 [//dx.doi.org/10.1126/science.275.5297.209](http://dx.doi.org/10.1126/science.275.5297.209).
- 1874 Sridhar Mahadevan and Mauro Maggioni. Proto-value functions: A laplacian framework for learning
 1875 representation and control in markov decision processes. *Journal of Machine Learning Research*, 8
 1876 (74):2169–2231, 2007. URL <http://jmlr.org/papers/v8/mahadevan07a.html>.
- 1877 D. Marr. Simple memory: A theory for archicortex. *Philosophical Transactions of the Royal Society*
 1878 *B: Biological Sciences*, 262(841):23–81, Jul 1971. ISSN 1471-2970. [10.1098/rstb.1971.0078](https://doi.org/10.1098/rstb.1971.0078). URL
 1879 <http://dx.doi.org/10.1098/rstb.1971.0078>.
- 1880 Marcelo G. Mattar and Nathaniel D. Daw. Prioritized memory access explains planning and hippocam-
 1881 pal replay. *Nature Neuroscience*, 21(11):1609–1617, Oct 2018. ISSN 1546-1726. [10.1038/s41593-018-](https://doi.org/10.1038/s41593-018-0232-z)
 1882 [0232-z](https://doi.org/10.1038/s41593-018-0232-z). URL <http://dx.doi.org/10.1038/s41593-018-0232-z>.
- 1883 Andrew P. Maurer and Lynn Nadel. The continuity of context: A role for the hippocampus. *Trends in*
 1884 *Cognitive Sciences*, 2021/01/15 2021. [10.1016/j.tics.2020.12.007](https://doi.org/10.1016/j.tics.2020.12.007). URL [https://doi.org/10.1016/](https://doi.org/10.1016/j.tics.2020.12.007)
 1885 [j.tics.2020.12.007](https://doi.org/10.1016/j.tics.2020.12.007).
- 1886 James L. McClelland, Bruce L. McNaughton, and Randall C. O’Reilly. Why there are complementary
 1887 learning systems in the hippocampus and neocortex: Insights from the successes and failures of
 1888 connectionist models of learning and memory. *Psychological Review*, 102(3):419–457, Jul 1995. ISSN
 1889 0033-295X. [10.1037/0033-295x.102.3.419](https://doi.org/10.1037/0033-295x.102.3.419). URL [http://dx.doi.org/10.1037/0033-295X.102.3.](http://dx.doi.org/10.1037/0033-295X.102.3.419)
 1890 [419](http://dx.doi.org/10.1037/0033-295X.102.3.419).
- 1891 James Lloyd McClelland and Matthew M. Botvinick. Deep learning: Implications for human learning
 1892 and memory. *PsyArXiv*, nov 2020. [10.31234/osf.io/3m5sb](https://doi.org/10.31234/osf.io/3m5sb). URL [https://doi.org/10.31234/](https://doi.org/10.31234/osf.io/3m5sb)
 1893 [2Fosf.io%2F3m5sb](https://doi.org/10.31234/osf.io/3m5sb).
- 1894 Michael McCloskey and Neal J. Cohen. Catastrophic interference in connectionist networks: The
 1895 sequential learning problem. *Psychology of Learning and Motivation*, 24:109–165, 1989. ISSN 0079-
 1896 7421. [10.1016/S0079-7421\(08\)60536-8](https://doi.org/10.1016/S0079-7421(08)60536-8). URL [http://dx.doi.org/10.1016/S0079-7421\(08\)60536-](http://dx.doi.org/10.1016/S0079-7421(08)60536-8)
 1897 [8](http://dx.doi.org/10.1016/S0079-7421(08)60536-8).
- 1898 Elizabeth A. McDevitt, Katherine A. Duggan, and Sara C. Mednick. REM sleep rescues learning from
 1899 interference. *Neurobiology of Learning and Memory*, 122:51–62, 2015. [10.1016/j.nlm.2014.11.015](https://doi.org/10.1016/j.nlm.2014.11.015).
 1900 URL <http://www.sciencedirect.com/science/article/pii/S1074742714002135>.

Colin G McNamara, Álvaro Tejero-Cantero, Stéphanie Trouche, Natalia Campo-Urriza, and David Dupret. Dopaminergic neurons promote hippocampal reactivation and spatial memory persistence. *Nature Neuroscience*, 17(12):1658–1660, 2014. [10.1038/nn.3843](https://doi.org/10.1038/nn.3843). URL <https://doi.org/10.1038/nn.3843>.

Nicolas Meuleau, Christian Plaunt, D Smith, and Tristan Smith. A POMDP for optimal motion planning with uncertain dynamics. In *ICAPS-10: POMDP Practitioners Workshop*. Citeseer, May 2010.

Frédéric Michon, Jyh-Jang Sun, Chae Young Kim, Davide Ciliberti, and Fabian Kloosterman. Post-learning hippocampal replay selectively reinforces spatial memory for highly rewarded locations. *Current Biology*, 29(9):1436–1444.e5, 2019. ISSN 0960-9822. [10.1016/j.cub.2019.03.048](https://doi.org/10.1016/j.cub.2019.03.048). URL <http://www.sciencedirect.com/science/article/pii/S096098221930346X>.

Kevin J Miller and Sarah Jo C Venditto. Multi-step planning in the brain. *Current Opinion in Behavioral Sciences*, 38:29 – 39, 2021. ISSN 2352-1546. [10.1016/j.cobeha.2020.07.003](https://doi.org/10.1016/j.cobeha.2020.07.003). URL <http://www.sciencedirect.com/science/article/pii/S2352154620301054>.

Marvin Minsky. Steps toward artificial intelligence. *Proceedings of the IRE*, 49(1):8–30, Jan 1961. ISSN 2162-6634. [10.1109/JRPROC.1961.287775](https://doi.org/10.1109/JRPROC.1961.287775). URL <https://ieeexplore.ieee.org/abstract/document/4066245>.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing Atari with Deep Reinforcement Learning. *arXiv e-prints*, art. arXiv:1312.5602, December 2013.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, and et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, Feb 2015. ISSN 1476-4687. [10.1038/nature14236](https://doi.org/10.1038/nature14236). URL <http://dx.doi.org/10.1038/nature14236>.

Volodymyr Mnih, Adrià Puigdomènech Badia, Mehdi Mirza, Alex Graves, Timothy P. Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous Methods for Deep Reinforcement Learning. *arXiv e-prints*, art. arXiv:1602.01783, February 2016.

Dean Mobbs, Drew B. Headley, Weilun Ding, and Peter Dayan. Space, time, and fear: Survival computations along defensive circuits. *Trends in Cognitive Sciences*, Feb 2020. ISSN 1364-6613. [10.1016/j.tics.2019.12.016](https://doi.org/10.1016/j.tics.2019.12.016). URL <http://dx.doi.org/10.1016/j.tics.2019.12.016>.

Thomas M. Moerland, Joost Broekens, and Catholijn M. Jonker. Model-based reinforcement learning: A survey. *arXiv e-prints*, page arXiv:2006.16712, June 2020.

I. Momennejad, E. M. Russek, J. H. Cheong, M. M. Botvinick, N. D. Daw, and S. J. Gershman. The successor representation in human reinforcement learning. *Nature Human Behaviour*, 1(9):680–692, Aug 2017. ISSN 2397-3374. [10.1038/s41562-017-0180-8](https://doi.org/10.1038/s41562-017-0180-8). URL <http://dx.doi.org/10.1038/s41562-017-0180-8>.

Ida Momennejad. Learning structures: Predictive representations, replay, and generalization. *Current Opinion in Behavioral Sciences*, 32:155–166, Apr 2020. ISSN 2352-1546. [10.1016/j.cobeha.2020.02.017](https://doi.org/10.1016/j.cobeha.2020.02.017). URL <http://dx.doi.org/10.1016/j.cobeha.2020.02.017>.

1940 Ida Momennejad, A Ross Otto, Nathaniel D Daw, and Kenneth A Norman. Offline replay supports
 1941 planning in human reinforcement learning. *eLife*, 7:e32548, Dec 2018. [10.7554/eLife.32548](https://doi.org/10.7554/eLife.32548). URL <https://doi.org/10.7554/eLife.32548>.
 1942

1943 Joseph D Monaco, Geeta Rao, Eric D Roth, and James J Knierim. Attentive scanning behavior
 1944 drives one-trial potentiation of hippocampal place fields. *Nature Neuroscience*, 17(5):725–731, 2014.
 1945 [10.1038/nn.3687](https://doi.org/10.1038/nn.3687). URL <https://doi.org/10.1038/nn.3687>.

1946 Andrew W. Moore and Christopher G. Atkeson. Prioritized sweeping: Reinforcement learning
 1947 with less data and less time. *Machine Learning*, 13(1):103–130, Oct 1993. ISSN 1573-0565.
 1948 [10.1007/bf00993104](http://dx.doi.org/10.1007/BF00993104). URL <http://dx.doi.org/10.1007/BF00993104>.

1949 Edvard I. Moser, Emilio Kropff, and May-Britt Moser. Place cells, grid cells, and the brain’s
 1950 spatial representation system. *Annual Review of Neuroscience*, 31(1):69–89, 2021/01/20 2008.
 1951 [10.1146/annurev.neuro.31.061307.090723](https://doi.org/10.1146/annurev.neuro.31.061307.090723). URL <https://doi.org/10.1146/annurev.neuro.31.061307.090723>.
 1952

1953 Karl F Muenzinger and F Milford Fletcher. Motivation in learning. VI. escape from electric shock
 1954 compared with hunger-food tension in the visual discrimination habit. *Journal of Comparative*
 1955 *Psychology*, 22(1):79–91, 1936.

1956 Rémi Munos, Tom Stepleton, Anna Harutyunyan, and Marc G. Bellemare. Safe and efficient off-policy
 1957 reinforcement learning. *arXiv e-prints*, art. arXiv:1606.02647, June 2016.

1958 Zoltán Nádasdy, Hajime Hirase, András Czurkó, Jozsef Csicsvari, and György Buzsáki. Replay and
 1959 time compression of recurring spike sequences in the hippocampus. *The Journal of Neuroscience*,
 1960 19(21):9497–9507, Nov 1999. ISSN 1529-2401. [10.1523/jneurosci.19-21-09497.1999](https://doi.org/10.1523/JNEUROSCI.19-21-09497.1999). URL <http://dx.doi.org/10.1523/JNEUROSCI.19-21-09497.1999>.
 1961

1962 Edward H. Nieh, Manuel Schottdorf, Nicolas W. Freeman, Ryan J. Low, Sam Lewallen, Sue Ann
 1963 Koay, Lucas Pinto, Jeffrey L. Gauthier, Carlos D. Brody, and David W. Tank. Geometry of abstract
 1964 learned knowledge in the hippocampus. *Nature*, 2021. [10.1038/s41586-021-03652-7](https://doi.org/10.1038/s41586-021-03652-7). URL <https://doi.org/10.1038/s41586-021-03652-7>.
 1965

1966 Niels Niethard and Jan Born. A backup of hippocampal spatial code outside the hippocampus?
 1967 new light on systems memory consolidation. *Neuron*, 106(2):204 – 206, 2020. ISSN 0896-6273.
 1968 doi.org/10.1016/j.neuron.2020.03.034. URL <http://www.sciencedirect.com/science/article/pii/S0896627320302415>.
 1969

1970 Yael Niv. Learning task-state representations. *Nature Neuroscience*, 22(10):1544–1553, Sep 2019. ISSN
 1971 1546-1726. [10.1038/s41593-019-0470-8](https://doi.org/10.1038/s41593-019-0470-8). URL <http://dx.doi.org/10.1038/s41593-019-0470-8>.

1972 Yael Niv, Reka Daniel, Andra Geana, Samuel J. Gershman, Yuan Chang Leong, Angela Rad-
 1973 ulescu, and Robert C. Wilson. Reinforcement learning in multidimensional environments relies
 1974 on attention mechanisms. *Journal of Neuroscience*, 35(21):8145–8157, 2015. ISSN 0270-6474.
 1975 [10.1523/JNEUROSCI.2978-14.2015](https://doi.org/10.1523/JNEUROSCI.2978-14.2015). URL <https://www.jneurosci.org/content/35/21/8145>.

1976 Yitzhak Norman, Erin M. Yeagle, Simon Khuvis, Michal Harel, Ashesh D. Mehta, and Rafael Malach.
 1977 Hippocampal sharp-wave ripples linked to visual episodic recollection in humans. *Science*, 365(6454):

- 1978 eaax1030, Aug 2019. ISSN 1095-9203. [10.1126/science.aax1030](https://doi.org/10.1126/science.aax1030). URL <http://dx.doi.org/10.1126/science.aax1030>.
- 1979
- 1980 John O'Keefe and Dulcie H. Conway. Hippocampal place units in the freely moving rat: Why they
1981 fire where they fire. *Experimental Brain Research*, 31(4):573–590, 1978. [10.1007/BF00239813](https://doi.org/10.1007/BF00239813). URL
1982 <https://doi.org/10.1007/BF00239813>.
- 1983 John O'Keefe and Jonathan Dostrovsky. The hippocampus as a spatial map. preliminary evidence from
1984 unit activity in the freely-moving rat. *Brain Research*, 34(1):171–175, Nov 1971. ISSN 0006-8993.
1985 [10.1016/0006-8993\(71\)90358-1](https://doi.org/10.1016/0006-8993(71)90358-1). URL [http://dx.doi.org/10.1016/0006-8993\(71\)90358-1](http://dx.doi.org/10.1016/0006-8993(71)90358-1).
- 1986 John O'Keefe and Julija Krupic. Do hippocampal pyramidal cells respond to non-spatial stimuli? *Phys-*
1987 *iological Reviews*, 2021. [10.1152/physrev.00014.2020](https://doi.org/10.1152/physrev.00014.2020). URL [https://doi.org/10.1152/physrev.](https://doi.org/10.1152/physrev.00014.2020)
1988 [00014.2020](https://doi.org/10.1152/physrev.00014.2020).
- 1989 John O'Keefe and Lynn Nadel. Maps in the brain. *New Scientist*, 62(903):749–751, Jun 1974.
1990 URL [https://books.google.de/books?id=vBRsz7SA0yoC&printsec=frontcover&hl=de&source=](https://books.google.de/books?id=vBRsz7SA0yoC&printsec=frontcover&hl=de&source=gbp_ge_summary_r&cad=0#v=onepage&q&f=false)
1991 [gbp_ge_summary_r&cad=0#v=onepage&q&f=false](https://books.google.de/books?id=vBRsz7SA0yoC&printsec=frontcover&hl=de&source=gbp_ge_summary_r&cad=0#v=onepage&q&f=false).
- 1992 John O'Keefe and Lynn Nadel. *The Hippocampus as a Cognitive Map*. Oxford: Clarendon Press, 1978.
- 1993 H Freyja Ólafsdóttir, Caswell Barry, Aman B Saleem, Demis Hassabis, and Hugo J Spiers. Hippocampal
1994 place cells construct reward related sequences through unexplored space. *eLife*, 4, Jun 2015. ISSN
1995 2050-084X. [10.7554/elife.06063](https://doi.org/10.7554/elife.06063). URL <http://dx.doi.org/10.7554/eLife.06063>.
- 1996 H Freyja Ólafsdóttir, Francis Carpenter, and Caswell Barry. Coordinated grid and place cell replay
1997 during rest. *Nature Neuroscience*, 19(6):792–794, Apr 2016. ISSN 1546-1726. [10.1038/nn.4291](https://doi.org/10.1038/nn.4291). URL
1998 <http://dx.doi.org/10.1038/nn.4291>.
- 1999 H. Freyja Ólafsdóttir, Francis Carpenter, and Caswell Barry. Task demands predict a dynamic switch
2000 in the content of awake hippocampal replay. *Neuron*, 96(4):925–935.e6, Nov 2017. ISSN 0896-6273.
2001 [10.1016/j.neuron.2017.09.035](https://doi.org/10.1016/j.neuron.2017.09.035). URL <http://dx.doi.org/10.1016/j.neuron.2017.09.035>.
- 2002 H. Freyja Ólafsdóttir, Daniel Bush, and Caswell Barry. The role of hippocampal replay in memory and
2003 planning. *Current Biology*, 28(1):R37–R50, Jan 2018. ISSN 0960-9822. [10.1016/j.cub.2017.10.073](https://doi.org/10.1016/j.cub.2017.10.073).
2004 URL <http://dx.doi.org/10.1016/j.cub.2017.10.073>.
- 2005 J. O'Neill, C.N. Boccara, F. Stella, P. Schoenenberger, and J. Csicsvari. Superficial layers of the medial
2006 entorhinal cortex replay independently of the hippocampus. *Science*, 355(6321):184–188, Jan 2017.
2007 ISSN 1095-9203. [10.1126/science.aag2787](https://doi.org/10.1126/science.aag2787). URL <http://dx.doi.org/10.1126/science.aag2787>.
- 2008 Joseph O'Neill, Timothy J Senior, Kevin Allen, John R Huxter, and Jozsef Csicsvari. Reactivation
2009 of experience-dependent cell assembly patterns in the hippocampus. *Nature Neuroscience*, 11(2):
2010 209–215, 2008. [10.1038/nn2037](https://doi.org/10.1038/nn2037). URL <https://doi.org/10.1038/nn2037>.
- 2011 Joseph O'Neill, Barty Pleydell-Bouverie, David Dupret, and Jozsef Csicsvari. Play it again: Reactiva-
2012 tion of waking experience and memory. *Trends in Neurosciences*, 33(5):220–229, May 2010. ISSN
2013 0166-2236. [10.1016/j.tins.2010.01.006](https://doi.org/10.1016/j.tins.2010.01.006). URL <http://dx.doi.org/10.1016/j.tins.2010.01.006>.
- 2014 Randall C. O'Reilly and James L. McClelland. Hippocampal conjunctive encoding, storage, and recall:
2015 Avoiding a trade-off. *Hippocampus*, 4(6):661–682, 1994. [10.1002/hipo.450040605](https://doi.org/10.1002/hipo.450040605).

- 2016 Randall C. O'Reilly, Rajan Bhattacharyya, Michael D. Howard, and Nicholas Ketz. Complementary
2017 learning systems. *Cognitive Science*, 38(6):1229–1248, 2014. [10.1111/j.1551-6709.2011.01214.x](https://doi.org/10.1111/j.1551-6709.2011.01214.x). URL
2018 <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1551-6709.2011.01214.x>.
- 2019 Delphine Oudiette and Ken A. Paller. Upgrading the sleeping brain with targeted memory reactivation.
2020 *Trends in Cognitive Sciences*, 17(3):142 – 149, 2013. ISSN 1364-6613. [10.1016/j.tics.2013.01.006](https://doi.org/10.1016/j.tics.2013.01.006).
2021 URL <http://www.sciencedirect.com/science/article/pii/S136466131300020X>.
- 2022 Yangchen Pan, Muhammad Zaheer, Adam White, Andrew Patterson, and Martha White. Organizing
2023 experience: A deeper look at replay mechanisms for sample-based planning in continuous state
2024 domains. *arXiv e-prints*, page arXiv:1806.04624, June 2018. URL [https://arxiv.org/abs/1806.](https://arxiv.org/abs/1806.04624)
2025 [04624](https://arxiv.org/abs/1806.04624).
- 2026 Andrew E. Papale, Mark C. Zielinski, Loren M. Frank, Shantanu P. Jadhav, and Aaron David Redish.
2027 Interplay between hippocampal sharp-wave-ripple events and vicarious trial and error behaviors in
2028 decision making. *Neuron*, 92(5):975–982, Dec 2016. ISSN 0896-6273. [10.1016/j.neuron.2016.10.028](https://doi.org/10.1016/j.neuron.2016.10.028).
2029 URL <http://dx.doi.org/10.1016/j.neuron.2016.10.028>.
- 2030 German I. Parisi, Ronald Kemker, Jose L. Part, Christopher Kanan, and Stefan Wermter. Con-
2031 tinual lifelong learning with neural networks: A review. *Neural Networks*, 113:54 – 71, 2019.
2032 ISSN 0893-6080. [10.1016/j.neunet.2019.01.012](https://doi.org/10.1016/j.neunet.2019.01.012). URL [http://www.sciencedirect.com/science/](http://www.sciencedirect.com/science/article/pii/S0893608019300231)
2033 [article/pii/S0893608019300231](http://www.sciencedirect.com/science/article/pii/S0893608019300231).
- 2034 Emilio Parisotto, H. Francis Song, Jack W. Rae, Razvan Pascanu, Caglar Gulcehre, Siddhant M.
2035 Jayakumar, Max Jaderberg, Raphael Lopez Kaufman, Aidan Clark, Seb Noury, Matthew M.
2036 Botvinick, Nicolas Heess, and Raia Hadsell. Stabilizing transformers for reinforcement learning.
2037 *arXiv e-prints*, page arXiv:1910.06764, October 2019.
- 2038 C Pavlides and J Winson. Influences of hippocampal place cell firing in the awake state on the
2039 activity of these cells during subsequent sleep episodes. *The Journal of Neuroscience*, 9(8):2907–
2040 2918, 1989. ISSN 1529-2401. [10.1523/jneurosci.09-08-02907.1989](https://doi.org/10.1523/jneurosci.09-08-02907.1989). URL [http://dx.doi.org/10.](http://dx.doi.org/10.1523/JNEUROSCI.09-08-02907.1989)
2041 [1523/JNEUROSCI.09-08-02907.1989](http://dx.doi.org/10.1523/JNEUROSCI.09-08-02907.1989).
- 2042 Michael Peer, Iva K. Brunec, Nora S. Newcombe, and Russell A. Epstein. Structuring knowl-
2043 edge with cognitive maps and cognitive graphs. *Trends in Cognitive Sciences*, 2020/11/26 2020.
2044 [10.1016/j.tics.2020.10.004](https://doi.org/10.1016/j.tics.2020.10.004). URL <https://doi.org/10.1016/j.tics.2020.10.004>.
- 2045 Lorenzo Pellegrini, Gabriele Graffieti, Vincenzo Lomonaco, and Davide Maltoni. Latent replay for
2046 real-time continual learning. *arXiv e-prints*, page arXiv:1912.01100, December 2019. URL [https:](https://arxiv.org/abs/1912.01100)
2047 [//arxiv.org/abs/1912.01100](https://arxiv.org/abs/1912.01100).
- 2048 Jing Peng and Ronald J. Williams. Efficient learning and planning within the dyna framework. *Adaptive*
2049 *Behavior*, 1(4):437–454, 2021/01/19 1993. [10.1177/105971239300100403](https://doi.org/10.1177/105971239300100403). URL [https://doi.org/](https://doi.org/10.1177/105971239300100403)
2050 [10.1177/105971239300100403](https://doi.org/10.1177/105971239300100403).
- 2051 C. M. A. Pennartz. The ventral striatum in off-line processing: Ensemble reactivation during sleep
2052 and modulation by hippocampal ripples. *Journal of Neuroscience*, 24(29):6446–6456, Jul 2004. ISSN
2053 1529-2401. [10.1523/jneurosci.0575-04.2004](https://doi.org/10.1523/jneurosci.0575-04.2004). URL [http://dx.doi.org/10.1523/JNEUROSCI.0575-](http://dx.doi.org/10.1523/JNEUROSCI.0575-04.2004)
2054 [04.2004](http://dx.doi.org/10.1523/JNEUROSCI.0575-04.2004).

2055 Adrien Peyrache, Mehdi Khamassi, Karim Benchenane, Sidney I Wiener, and Francesco P Battaglia.
 2056 Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nature Neuro-*
 2057 *science*, 12(7):919–926, May 2009. ISSN 1546-1726. [10.1038/nn.2337](https://doi.org/10.1038/nn.2337). URL [http://dx.doi.org/](http://dx.doi.org/10.1038/nn.2337)
 2058 [10.1038/nn.2337](http://dx.doi.org/10.1038/nn.2337).

2059 Giovanni Pezzulo, Matthijs A.A. van der Meer, Carien S. Lansink, and Cyriel M.A. Pennartz. Internally
 2060 generated sequences in learning and executing goal-directed behavior. *Trends in Cognitive Sciences*,
 2061 18(12):647–657, Dec 2014. ISSN 1364-6613. [10.1016/j.tics.2014.06.011](https://doi.org/10.1016/j.tics.2014.06.011). URL [http://dx.doi.org/](http://dx.doi.org/10.1016/j.tics.2014.06.011)
 2062 [10.1016/j.tics.2014.06.011](http://dx.doi.org/10.1016/j.tics.2014.06.011).

2063 Giovanni Pezzulo, Francesco Donnarumma, Domenico Maisto, and Ivilin Stoianov. Planning at decision
 2064 time and in the background during spatial navigation. *Current Opinion in Behavioral Sciences*, 29:
 2065 69–76, Oct 2019. ISSN 2352-1546. [10.1016/j.cobeha.2019.04.009](https://doi.org/10.1016/j.cobeha.2019.04.009). URL [http://dx.doi.org/10.](http://dx.doi.org/10.1016/j.cobeha.2019.04.009)
 2066 [1016/j.cobeha.2019.04.009](http://dx.doi.org/10.1016/j.cobeha.2019.04.009).

2067 Brad E. Pfeiffer and David J. Foster. Hippocampal place-cell sequences depict future paths to re-
 2068 membered goals. *Nature*, 497(7447):74–79, Apr 2013. ISSN 1476-4687. [10.1038/nature12112](https://doi.org/10.1038/nature12112). URL <http://dx.doi.org/10.1038/nature12112>.

2070 Jary Pomponi, Simone Scardapane, Vincenzo Lomonaco, and Aurelio Uncini. Efficient continual learn-
 2071 ing in neural networks with embedding regularization. *Neurocomputing*, 397:139–148, 2020. ISSN
 2072 0925-2312. <https://doi.org/10.1016/j.neucom.2020.01.093>. URL [https://www.sciencedirect.com/](https://www.sciencedirect.com/science/article/pii/S092523122030151X)
 2073 [science/article/pii/S092523122030151X](https://www.sciencedirect.com/science/article/pii/S092523122030151X).

2074 Vitchyr Pong, Shixiang Gu, Murtaza Dalal, and Sergey Levine. Temporal difference models: Model-free
 2075 deep RL for model-based control. *arXiv e-prints*, art. arXiv:1802.09081, February 2018.

2076 Yu-Lin Qin, Bruce L. McNaughton, William E. Skaggs, and Carol A. Barnes. Memory reprocessing in
 2077 corticocortical and hippocampocortical neuronal ensembles. *Philosophical Transactions of the Royal*
 2078 *Society of London. Series B: Biological Sciences*, 352(1360):1525–1533, Oct 1997. ISSN 1471-2970.
 2079 [10.1098/rstb.1997.0139](https://doi.org/10.1098/rstb.1997.0139). URL <http://dx.doi.org/10.1098/rstb.1997.0139>.

2080 Dhakshin S. Ramanathan, Tanuj Gulati, and Karunesh Ganguly. Sleep-dependent reactivation of
 2081 ensembles in motor cortex promotes skill consolidation. *PLOS Biology*, 13(9):1–25, Sep 2015.
 2082 ISSN 1545-7885. [10.1371/journal.pbio.1002263](https://doi.org/10.1371/journal.pbio.1002263). URL [http://dx.doi.org/10.1371/journal.pbio.](http://dx.doi.org/10.1371/journal.pbio.1002263)
 2083 [1002263](http://dx.doi.org/10.1371/journal.pbio.1002263).

2084 Björn Rasch and Jan Born. Maintaining memories by reactivation. *Current Opinion in Neurobiology*,
 2085 17(6):698–703, Dec 2007. ISSN 0959-4388. [10.1016/j.conb.2007.11.007](https://doi.org/10.1016/j.conb.2007.11.007). URL [http://dx.doi.org/](http://dx.doi.org/10.1016/j.conb.2007.11.007)
 2086 [10.1016/j.conb.2007.11.007](http://dx.doi.org/10.1016/j.conb.2007.11.007).

2087 Roger Ratcliff. Connectionist models of recognition memory: Constraints imposed by learning and
 2088 forgetting functions. *Psychological Review*, 97(2):285–308, 1990. [10.1037/0033-295x.97.2.285](https://doi.org/10.1037/0033-295x.97.2.285). URL
 2089 <https://doi.org/10.1037/0033-295x.97.2.285>.

2090 Aaron David Redish. *Beyond the cognitive map: From place cells to episodic memory*. MIT Press,
 2091 1999.

2092 Aaron David Redish. Vicarious trial and error. *Nature Reviews Neuroscience*, 17(3):147–159, Feb 2016.
 2093 ISSN 1471-0048. [10.1038/nrn.2015.30](https://doi.org/10.1038/nrn.2015.30). URL <http://dx.doi.org/10.1038/nrn.2015.30>.

2094 Aaron David Redish and David S. Touretzky. The role of the hippocampus in solving the morris water
 2095 maze. *Neural Computation*, 10(1):73–111, Jan 1998. ISSN 1530-888X. [10.1162/089976698300017908](https://doi.org/10.1162/089976698300017908).
 2096 URL <http://dx.doi.org/10.1162/089976698300017908>.

2097 Lauren L. Richmond and Jeffrey M. Zacks. Constructing experience: Event models from per-
 2098 ception to action. *Trends in Cognitive Sciences*, 21(12):962–980, Dec 2017. ISSN 1364-6613.
 2099 [10.1016/j.tics.2017.08.005](https://doi.org/10.1016/j.tics.2017.08.005). URL <http://dx.doi.org/10.1016/j.tics.2017.08.005>.

2100 Emma L. Roscow, Matthew W. Jones, and Nathan F. Lepora. Behavioural and computational evidence
 2101 for memory consolidation biased by reward-prediction errors. *bioRxiv*, 2019. [10.1101/716290](https://doi.org/10.1101/716290). URL
 2102 <https://www.biorxiv.org/content/early/2019/07/26/716290>.

2103 Gideon Rothschild, Elad Eban, and Loren M Frank. A cortical–hippocampal–cortical loop of informa-
 2104 tion processing during memory consolidation. *Nature Neuroscience*, 20(2):251–259, Dec 2016. ISSN
 2105 1546-1726. [10.1038/nn.4457](https://doi.org/10.1038/nn.4457). URL <http://dx.doi.org/10.1038/nn.4457>.

2106 Nina Rouhani, Kenneth A. Norman, Yael Niv, and Aaron M. Bornstein. Reward prediction
 2107 errors create event boundaries in memory. *Cognition*, 203:104269, 2020. ISSN 0010-0277.
 2108 [10.1016/j.cognition.2020.104269](https://doi.org/10.1016/j.cognition.2020.104269). URL [https://www.sciencedirect.com/science/article/pii/](https://www.sciencedirect.com/science/article/pii/S0010027720300883)
 2109 [S0010027720300883](https://www.sciencedirect.com/science/article/pii/S0010027720300883).

2110 Evan M. Russek, Ida Momennejad, Matthew M. Botvinick, Samuel J. Gershman, and Nathaniel D.
 2111 Daw. Predictive representations can link model-based reinforcement learning to model-free mech-
 2112 anisms. *PLoS Computational Biology*, 13(9):e1005768, Sep 2017. ISSN 1553-7358. [10.1371/jour-](https://doi.org/10.1371/journal.pcbi.1005768)
 2113 [nal.pcbi.1005768](https://doi.org/10.1371/journal.pcbi.1005768). URL <http://dx.doi.org/10.1371/journal.pcbi.1005768>.

2114 Matthew Schafer and Daniela Schiller. Navigating social space. *Neuron*, 100(2):476–489, Oct 2018.
 2115 ISSN 0896-6273. [10.1016/j.neuron.2018.10.006](https://doi.org/10.1016/j.neuron.2018.10.006). URL [http://dx.doi.org/10.1016/j.neuron.](http://dx.doi.org/10.1016/j.neuron.2018.10.006)
 2116 [2018.10.006](http://dx.doi.org/10.1016/j.neuron.2018.10.006).

2117 Anna C Schapiro, Timothy T Rogers, Natalia I Cordova, Nicholas B Turk-Browne, and Matthew M
 2118 Botvinick. Neural representations of events arise from temporal community structure. *Nature Neu-*
 2119 *roscience*, 16(4):486–492, Feb 2013. ISSN 1546-1726. [10.1038/nn.3331](https://doi.org/10.1038/nn.3331). URL [http://dx.doi.org/](http://dx.doi.org/10.1038/nn.3331)
 2120 [10.1038/nn.3331](http://dx.doi.org/10.1038/nn.3331).

2121 Anna C. Schapiro, Nicholas B. Turk-Browne, Matthew M. Botvinick, and Kenneth A. Norman. Com-
 2122plementary learning systems within the hippocampus: a neural network modelling approach to
 2123 reconciling episodic memory with statistical learning. *Philosophical Transactions of the Royal Soci-*
 2124 *ety B: Biological Sciences*, 372(1711):20160049, Jan 2017. ISSN 1471-2970. [10.1098/rstb.2016.0049](https://doi.org/10.1098/rstb.2016.0049).
 2125 URL <http://dx.doi.org/10.1098/rstb.2016.0049>.

2126 Anna C. Schapiro, Elizabeth A. McDevitt, Timothy T. Rogers, Sara C. Mednick, and Kenneth A.
 2127 Norman. Human hippocampal replay during rest prioritizes weakly learned information and predicts
 2128 memory performance. *Nature Communications*, 9(1), Sep 2018. ISSN 2041-1723. [10.1038/s41467-](https://doi.org/10.1038/s41467-018-06213-1)
 2129 [018-06213-1](https://doi.org/10.1038/s41467-018-06213-1). URL <http://dx.doi.org/10.1038/s41467-018-06213-1>.

2130 Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay.
 2131 *arXiv e-prints*, page arXiv:1511.05952, Nov 2015. URL [https://ui.adsabs.harvard.edu/abs/](https://ui.adsabs.harvard.edu/abs/2015arXiv151105952S)
 2132 [2015arXiv151105952S](https://ui.adsabs.harvard.edu/abs/2015arXiv151105952S).

2133 Brandy Schmidt, Anneke A. Duin, and Aaron David Redish. Disrupting the medial prefrontal cortex
2134 alters hippocampal sequences during deliberative decision making. *Journal of Neurophysiology*, 121
2135 (6):1981–2000, Jun 2019. ISSN 1522-1598. [10.1152/jn.00793.2018](https://doi.org/10.1152/jn.00793.2018). URL <http://dx.doi.org/10.1152/jn.00793.2018>.

2137 Christina Schmidt, Philippe Peigneux, Vincenzo Muto, Maja Schenkel, Vera Knoblauch, Mirjam
2138 Münch, Dominique J.-F. de Quervain, Anna Wirz-Justice, and Christian Cajochen. Encoding dif-
2139 ficulty promotes postlearning changes in sleep spindle activity during napping. *Journal of Neu-
2140 roscience*, 26(35):8976–8982, 2006. ISSN 0270-6474. [10.1523/JNEUROSCI.2464-06.2006](https://doi.org/10.1523/JNEUROSCI.2464-06.2006). URL
2141 <https://www.jneurosci.org/content/26/35/8976>.

2142 Nicolas W Schuck and Yael Niv. Sequential replay of nonspatial task states in the human hippocampus.
2143 *Science*, 364(6447):eaaw5181, 2019. [10.1126/science.aaw5181](https://doi.org/10.1126/science.aaw5181).

2144 Nicolas W. Schuck, Robert Gaschler, Dorit Wenke, Jakob Heinzle, Peter A. Frensch, John-Dylan
2145 Haynes, and Carlo Reverberi. Medial prefrontal cortex predicts internally driven strategy shifts.
2146 *Neuron*, 86(1):331 – 340, 2015. ISSN 0896-6273. doi.org/10.1016/j.neuron.2015.03.015. URL [http://](http://www.sciencedirect.com/science/article/pii/S0896627315002111)
2147 www.sciencedirect.com/science/article/pii/S0896627315002111.

2148 Nicolas W. Schuck, Ming Bo Cai, Robert C. Wilson, and Yael Niv. Human orbitofrontal cor-
2149 tex represents a cognitive map of state space. *Neuron*, 91(6):1402 – 1412, 2016. ISSN 0896-
2150 6273. [10.1016/j.neuron.2016.08.019](https://doi.org/10.1016/j.neuron.2016.08.019). URL [http://www.sciencedirect.com/science/article/](http://www.sciencedirect.com/science/article/pii/S0896627316305116)
2151 [pii/S0896627316305116](http://www.sciencedirect.com/science/article/pii/S0896627316305116).

2152 Nicolas W. Schuck, Robert Wilson, and Yael Niv. A state representation for reinforcement learning
2153 and decision-making in the orbitofrontal cortex. In Richard Morris, Aaron Bornstein, and Amitai
2154 Shenhav, editors, *Goal-Directed Decision Making*, chapter 12, pages 259–278. Academic Press, 1
2155 edition, 2018. ISBN 978-0-12-812098-9. [10.1016/B978-0-12-812098-9.00012-7](https://doi.org/10.1016/B978-0-12-812098-9.00012-7). URL [https://www.](https://www.sciencedirect.com/science/article/pii/B9780128120989000127)
2156 [sciencedirect.com/science/article/pii/B9780128120989000127](https://www.sciencedirect.com/science/article/pii/B9780128120989000127).

2157 William Beecher Scoville and Brenda Milner. Loss of recent memory after bilateral hippocampal lesions.
2158 *Journal of neurology, neurosurgery, and psychiatry*, 20(1):11–21, 02 1957. [10.1136/jnnp.20.1.11](https://doi.org/10.1136/jnnp.20.1.11). URL
2159 <https://pubmed.ncbi.nlm.nih.gov/13406589>.

2160 Melissa J Sharpe, Chun Yun Chang, Melissa A Liu, Hannah M Batchelor, Lauren E Mueller, Joshua L
2161 Jones, Yael Niv, and Geoffrey Schoenbaum. Dopamine transients are sufficient and necessary for
2162 acquisition of model-based associations. *Nature Neuroscience*, 20(5):735–742, 2017. [10.1038/nn.4538](https://doi.org/10.1038/nn.4538).
2163 URL <https://doi.org/10.1038/nn.4538>.

2164 Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep gener-
2165 ative replay. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan,
2166 and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 2990–2999.
2167 Curran Associates, Inc., 2017. URL [http://papers.nips.cc/paper/6892-continual-learning-](http://papers.nips.cc/paper/6892-continual-learning-with-deep-generative-replay.pdf)
2168 [with-deep-generative-replay.pdf](http://papers.nips.cc/paper/6892-continual-learning-with-deep-generative-replay.pdf).

2169 Justin D. Shin, Wenbo Tang, and Shantanu P. Jadhav. Dynamics of awake hippocampal-prefrontal re-
2170 play for spatial learning and memory-guided decision making. *Neuron*, 104(6):1110–1125, Dec 2019.
2171 ISSN 0896-6273. [10.1016/j.neuron.2019.09.012](https://doi.org/10.1016/j.neuron.2019.09.012). URL [http://dx.doi.org/10.1016/j.neuron.](http://dx.doi.org/10.1016/j.neuron.2019.09.012)
2172 [2019.09.012](http://dx.doi.org/10.1016/j.neuron.2019.09.012).

2173 Yeon Soon Shin and Sarah DuBrow. Structuring memory through inference-based event segmentation.
2174 *Topics in Cognitive Science*, 2021/01/22 2020. [10.1111/tops.12505](https://doi.org/10.1111/tops.12505). URL <https://doi.org/10.1111/tops.12505>.
2175

2176 Delia Silva, Ting Feng, and David J. Foster. Trajectory events across hippocampal place cells re-
2177 quire previous experience. *Nature Neuroscience*, 18(12):1772–1779, Oct 2015. ISSN 1546-1726.
2178 [10.1038/nn.4151](https://doi.org/10.1038/nn.4151). URL <http://dx.doi.org/10.1038/nn.4151>.

2179 David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche,
2180 Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman,
2181 Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach,
2182 Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of go with deep
2183 neural networks and tree search. *Nature*, 529(7587):484–489, 2016. [10.1038/nature16961](https://doi.org/10.1038/nature16961). URL
2184 <https://doi.org/10.1038/nature16961>.

2185 Annabelle C. Singer and Loren M. Frank. Rewarded outcomes enhance reactivation of experience in
2186 the hippocampus. *Neuron*, 64(6):910–921, Dec 2009. ISSN 0896-6273. [10.1016/j.neuron.2009.11.016](https://doi.org/10.1016/j.neuron.2009.11.016).
2187 URL <http://dx.doi.org/10.1016/j.neuron.2009.11.016>.

2188 Annabelle C. Singer, Margaret F. Carr, Mattias P. Karlsson, and Loren M. Frank. Hippocampal SWR
2189 activity predicts correct decisions during the initial learning of an alternation task. *Neuron*, 77(6):
2190 1163–1173, Mar 2013. ISSN 0896-6273. [10.1016/j.neuron.2013.01.027](https://doi.org/10.1016/j.neuron.2013.01.027). URL [http://dx.doi.org/](http://dx.doi.org/10.1016/j.neuron.2013.01.027)
2191 [10.1016/j.neuron.2013.01.027](http://dx.doi.org/10.1016/j.neuron.2013.01.027).

2192 W. E. Skaggs and Bruce L. McNaughton. Replay of neuronal firing sequences in rat hippocampus
2193 during sleep following spatial experience. *Science*, 271(5257):1870–1873, Mar 1996. ISSN 1095-9203.
2194 [10.1126/science.271.5257.1870](https://doi.org/10.1126/science.271.5257.1870). URL <http://dx.doi.org/10.1126/science.271.5257.1870>.

2195 H. Francis Song, Abbas Abdolmaleki, Jost Tobias Springenberg, Aidan Clark, Hubert Soyer, Jack W.
2196 Rae, Seb Noury, Arun Ahuja, Siqi Liu, Dhruva Tirumala, Nicolas Heess, Dan Belov, Martin Ried-
2197 miller, and Matthew M. Botvinick. V-MPO: On-policy maximum a posteriori policy optimization
2198 for discrete and continuous control. *arXiv e-prints*, art. arXiv:1909.12238, September 2019.

2199 Hugo J. Spiers. The hippocampal cognitive map: One space or many? *Trends in Cognitive Sciences*,
2200 Jan 2020. ISSN 1364-6613. [10.1016/j.tics.2019.12.013](https://doi.org/10.1016/j.tics.2019.12.013). URL [http://dx.doi.org/10.1016/j.tics.](http://dx.doi.org/10.1016/j.tics.2019.12.013)
2201 [2019.12.013](http://dx.doi.org/10.1016/j.tics.2019.12.013).

2202 Larry Ryan Squire. Memory and the hippocampus: A synthesis from findings with rats, monkeys,
2203 and humans. *Psychological Review*, 99(2):195–231, 1992. URL [https://doi.org/10.1037/0033-](https://doi.org/10.1037/0033-295x.99.2.195)
2204 [295x.99.2.195](https://doi.org/10.1037/0033-295x.99.2.195).

2205 Richard J. Staba, Charles L. Wilson, Anatol Bragin, Itzhak Fried, and Jerome Engel. Quantita-
2206 tive analysis of high-frequency oscillations (80–500 Hz) recorded in human epileptic hippocampus
2207 and entorhinal cortex. *Journal of Neurophysiology*, 88(4):1743–1752, Oct 2002. ISSN 1522-1598.
2208 [10.1152/jn.2002.88.4.1743](https://doi.org/10.1152/jn.2002.88.4.1743). URL <http://dx.doi.org/10.1152/jn.2002.88.4.1743>.

2209 Kimberly L Stachenfeld, Matthew M Botvinick, and Samuel J Gershman. The hippocampus as a pre-
2210 dictive map. *Nature Neuroscience*, 20(11):1643–1653, Oct 2017. ISSN 1546-1726. [10.1038/nn.4650](https://doi.org/10.1038/nn.4650).
2211 URL <http://dx.doi.org/10.1038/nn.4650>.

2212 B. P. Staresina, A. Alink, N. Kriegeskorte, and R. N. Henson. Awake reactivation predicts memory in
 2213 humans. *Proceedings of the National Academy of Sciences*, 110(52):21159–21164, Dec 2013. ISSN
 2214 1091-6490. [10.1073/pnas.1311989110](https://doi.org/10.1073/pnas.1311989110). URL <http://dx.doi.org/10.1073/pnas.1311989110>.

2215 Bernhard P Staresina, Til Ole Bergmann, Mathilde Bonnefond, Roemer van der Meij, Ole Jensen,
 2216 Lorena Deuker, Christian E Elger, Nikolai Axmacher, and Juergen Fell. Hierarchical nesting of slow
 2217 oscillations, spindles and ripples in the human hippocampus during sleep. *Nature Neuroscience*, 18
 2218 (11):1679–1686, Sep 2015. ISSN 1546-1726. [10.1038/nn.4119](https://doi.org/10.1038/nn.4119). URL [http://dx.doi.org/10.1038/](http://dx.doi.org/10.1038/nn.4119)
 2219 [nn.4119](http://dx.doi.org/10.1038/nn.4119).

2220 Adam Steiner and Aaron David Redish. The road not taken: Neural correlates of decision
 2221 making in orbitofrontal cortex. *Frontiers in Neuroscience*, 6:131, 2012. ISSN 1662-453X.
 2222 [10.3389/fnins.2012.00131](https://doi.org/10.3389/fnins.2012.00131). URL [https://www.frontiersin.org/article/10.3389/fnins.2012.](https://www.frontiersin.org/article/10.3389/fnins.2012.00131)
 2223 [00131](https://www.frontiersin.org/article/10.3389/fnins.2012.00131).

2224 Federico Stella, Peter Baracskey, Joseph O’Neill, and Jozsef Csicsvari. Hippocampal reactivation of
 2225 random trajectories resembling brownian diffusion. *Neuron*, 102(2):450–461, Apr 2019. ISSN 0896-
 2226 6273. [10.1016/j.neuron.2019.01.052](https://doi.org/10.1016/j.neuron.2019.01.052). URL <http://dx.doi.org/10.1016/j.neuron.2019.01.052>.

2227 Ivilin Stoianov, Domenico Maisto, and Giovanni Pezzulo. The hippocampal formation as a hierar-
 2228 chical generative model supporting generative replay and continual learning. *bioRxiv*, Jan 2020.
 2229 [10.1101/2020.01.16.908889](https://doi.org/10.1101/2020.01.16.908889). URL <http://dx.doi.org/10.1101/2020.01.16.908889>.

2230 Chen Sun, Wannan Yang, Jared Martin, and Susumu Tonegawa. Hippocampal neurons represent events
 2231 as transferable units of experience. *Nature Neuroscience*, Apr 2020. ISSN 1546-1726. [10.1038/s41593-](https://doi.org/10.1038/s41593-020-0614-x)
 2232 [020-0614-x](https://doi.org/10.1038/s41593-020-0614-x). URL <http://dx.doi.org/10.1038/s41593-020-0614-x>.

2233 Gary R Sutherland and Bruce McNaughton. Memory trace reactivation in hippocampal and neo-
 2234 cortical neuronal ensembles. *Current Opinion in Neurobiology*, 10(2):180 – 186, 2000. ISSN 0959-
 2235 4388. [10.1016/S0959-4388\(00\)00079-9](https://doi.org/10.1016/S0959-4388(00)00079-9). URL [http://www.sciencedirect.com/science/article/](http://www.sciencedirect.com/science/article/pii/S0959438800000799)
 2236 [pii/S0959438800000799](http://www.sciencedirect.com/science/article/pii/S0959438800000799).

2237 Richard S. Sutton. Integrated architectures for learning, planning, and reacting based on approximating
 2238 dynamic programming. In Bruce Porter and Raymond Mooney, editors, *Machine Learning Proceed-*
 2239 *ings 1990*, pages 216–224. Morgan Kaufmann, San Francisco (CA), 1990. ISBN 978-1-55860-141-3.
 2240 [10.1016/B978-1-55860-141-3.50030-4](https://doi.org/10.1016/B978-1-55860-141-3.50030-4). URL [https://www.sciencedirect.com/science/article/](https://www.sciencedirect.com/science/article/pii/B9781558601413500304)
 2241 [pii/B9781558601413500304](https://www.sciencedirect.com/science/article/pii/B9781558601413500304).

2242 Richard S. Sutton. Dyna, an integrated architecture for learning, planning, and reacting. *ACM*
 2243 *SIGART Bulletin*, 2(4):160–163, Jul 1991. ISSN 0163-5719. [10.1145/122344.122377](https://doi.org/10.1145/122344.122377). URL [http:](http://dx.doi.org/10.1145/122344.122377)
 2244 [//dx.doi.org/10.1145/122344.122377](http://dx.doi.org/10.1145/122344.122377).

2245 Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT Press, 2
 2246 edition, 2018. URL <http://incompleteideas.net/book/the-book-2nd.html>.

2247 Richard S. Sutton, Csaba Szepesvári, Alborz Geramifard, and Michael Bowling. Dyna-style planning
 2248 with linear function approximation and prioritized sweeping. *arXiv e-prints*, abs/1206.3285, 2012.
 2249 URL <http://arxiv.org/abs/1206.3285>.

2250 Rachel A Swanson, Daniel Levenstein, Kathryn McClain, David Tingley, and György Buzsáki. Variable
2251 specificity of memory trace reactivation during hippocampal sharp wave ripples. *Current Opinion in*
2252 *Behavioral Sciences*, 32:126–135, 2020. ISSN 2352-1546. [10.1016/j.cobeha.2020.02.008](https://doi.org/10.1016/j.cobeha.2020.02.008). URL <http://www.sciencedirect.com/science/article/pii/S2352154620300280>. Understanding memory:
2253 Which level of analysis?
2254

2255 A. Tambini and L. Davachi. Persistence of hippocampal multivoxel patterns into postencoding rest is
2256 related to memory. *Proceedings of the National Academy of Sciences*, 110(48):19591–19596, Nov 2013.
2257 ISSN 1091-6490. [10.1073/pnas.1308499110](https://doi.org/10.1073/pnas.1308499110). URL <http://dx.doi.org/10.1073/pnas.1308499110>.

2258 Arielle Tambini and Lila Davachi. Awake reactivation of prior experiences consolidates memories
2259 and biases cognition. *Trends in Cognitive Sciences*, 23(10):876–890, Oct 2019. ISSN 1364-6613.
2260 [10.1016/j.tics.2019.07.008](https://doi.org/10.1016/j.tics.2019.07.008). URL <http://dx.doi.org/10.1016/j.tics.2019.07.008>.

2261 Arielle Tambini, Nicholas Ketz, and Lila Davachi. Enhanced brain correlations during rest are re-
2262 lated to memory for recent experiences. *Neuron*, 65(2):280–290, Jan 2010. ISSN 0896-6273.
2263 [10.1016/j.neuron.2010.01.001](https://doi.org/10.1016/j.neuron.2010.01.001). URL <http://dx.doi.org/10.1016/j.neuron.2010.01.001>.

2264 Arielle Tambini, Alice Berners-Lee, and Lila Davachi. Brief targeted memory reactivation during the
2265 awake state enhances memory stability and benefits the weakest memories. *Scientific Reports*, 7
2266 (1), Nov 2017. ISSN 2045-2322. [10.1038/s41598-017-15608-x](https://doi.org/10.1038/s41598-017-15608-x). URL <http://dx.doi.org/10.1038/s41598-017-15608-x>.
2267

2268 Wenbo Tang and Shantanu P. Jadhav. Sharp-wave ripples as a signature of hippocampal-prefrontal
2269 reactivation for memory during sleep and waking states. *Neurobiology of Learning and Memory*, 160:
2270 11–20, Apr 2019. ISSN 1074-7427. [10.1016/j.nlm.2018.01.002](https://doi.org/10.1016/j.nlm.2018.01.002). URL <http://dx.doi.org/10.1016/j.nlm.2018.01.002>.
2271

2272 Wenbo Tang, Justin D. Shin, Loren M. Frank, and Shantanu P. Jadhav. Hippocampal-prefrontal
2273 reactivation during learning is stronger in awake compared with sleep states. *Journal of Neuroscience*,
2274 37(49):11789–11805, 2017. ISSN 0270-6474. [10.1523/JNEUROSCI.2291-17.2017](https://doi.org/10.1523/JNEUROSCI.2291-17.2017). URL <https://www.jneurosci.org/content/37/49/11789>.
2275

2276 Wenbo Tang, Justin D Shin, and Shantanu P Jadhav. Multiple time-scales of decision making in the hip-
2277 pocampus and prefrontal cortex. *eLife*, 10:e66227, mar 2021. ISSN 2050-084X. [10.7554/eLife.66227](https://doi.org/10.7554/eLife.66227).
2278 URL <https://doi.org/10.7554/eLife.66227>.

2279 Yong Tang, Jens R. Nyengaard, Didima M.G. De Groot, and Hans Jørgen G. Gundersen. Total
2280 regional and global number of synapses in the human brain neocortex. *Synapse*, 41(3):258–273,
2281 2001. [10.1002/syn.1083](https://doi.org/10.1002/syn.1083). URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/syn.1083>.

2282 Gerald Tesauro. Temporal difference learning and TD-gammon. *Communications of the ACM*, 38(3):
2283 58–68, March 1995. ISSN 0001-0782. [10.1145/203330.203343](https://doi.org/10.1145/203330.203343). URL <https://doi.org/10.1145/203330.203343>.
2284

2285 E. C. Tolman. The determiners of behavior at a choice point. *Psychological Review*, 45(1):1–41, 1938.
2286 ISSN 0033-295X. [10.1037/h0062733](https://doi.org/10.1037/h0062733). URL <http://dx.doi.org/10.1037/h0062733>.

2287 Edward C Tolman. A behavioristic theory of ideas. *Psychological Review*, 33(5):352, 1926.
2288 [10.1037/h0070532](https://doi.org/10.1037/h0070532). URL <https://doi.org/10.1037/h0070532>.

2289 Edward C. Tolman. Cognitive maps in rats and men. *Psychological Review*, 55(4):189–208, 1948. ISSN
2290 0033-295X. [10.1037/h0061626](https://doi.org/10.1037/h0061626). URL <http://dx.doi.org/10.1037/h0061626>.

2291 Alexa Tompary and Lila Davachi. Consolidation promotes the emergence of representa-
2292 tional overlap in the hippocampus and medial prefrontal cortex. *Neuron*, 96(1):228–241,
2293 2017. [10.1016/j.neuron.2017.09.005](https://doi.org/10.1016/j.neuron.2017.09.005). URL [http://www.sciencedirect.com/science/article/
2294 pii/S0896627317308401](http://www.sciencedirect.com/science/article/pii/S0896627317308401).

2295 Sean G. Trettel, John B. Trimper, Ernie Hwaun, Ila R. Fiete, and Laura Lee Colgin. Grid cell co-
2296 activity patterns during sleep reflect spatial overlap of grid fields during active behaviors. *Nature*
2297 *Neuroscience*, 22(4):609–617, Mar 2019. ISSN 1546-1726. [10.1038/s41593-019-0359-6](https://doi.org/10.1038/s41593-019-0359-6). URL [http:
2298 //dx.doi.org/10.1038/s41593-019-0359-6](http://dx.doi.org/10.1038/s41593-019-0359-6).

2299 Gido M. van de Ven and Andreas S. Tolias. Generative replay with feedback connections as a general
2300 strategy for continual learning. *arXiv e-prints*, page arXiv:1809.10635, September 2018. URL [https:
2301 //ui.adsabs.harvard.edu/abs/2018arXiv180910635V](https://ui.adsabs.harvard.edu/abs/2018arXiv180910635V).

2302 Gido M. van de Ven, Stéphanie Trouche, Colin G. McNamara, Kevin Allen, and David Dupret. Hip-
2303 pocampal offline reactivation consolidates recently formed cell assembly patterns during sharp wave-
2304 ripples. *Neuron*, 92(5):968–974, Dec 2016. ISSN 0896-6273. [10.1016/j.neuron.2016.10.020](https://doi.org/10.1016/j.neuron.2016.10.020). URL
2305 <http://dx.doi.org/10.1016/j.neuron.2016.10.020>.

2306 Gido M. van de Ven, Hava T. Siegelmann, and Andreas S. Tolias. Brain-inspired replay for continual
2307 learning with artificial neural networks. *Nature Communications*, 11(1):4069, 2020. [10.1038/s41467-
2308 020-17866-2](https://doi.org/10.1038/s41467-020-17866-2). URL <https://doi.org/10.1038/s41467-020-17866-2>.

2309 Hado van Hasselt, Matteo Hessel, and John Aslanides. When to use parametric models in reinforcement
2310 learning? *arXiv e-prints*, art. arXiv:1906.05243, June 2019.

2311 Harm van Seijen and Rich Sutton. A deeper look at planning as learning from replay. In *International
2312 conference on machine learning*, pages 2314–2322, 2015.

2313 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz
2314 Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv e-prints*, page arXiv:1706.03762, June
2315 2017.

2316 Alex P. Vaz, Sara K. Inati, Nicolas Brunel, and Kareem A. Zaghloul. Coupled ripple oscillations
2317 between the medial temporal lobe and neocortex retrieve human memory. *Science*, 363(6430):975–
2318 978, Feb 2019. ISSN 1095-9203. [10.1126/science.aau8956](https://doi.org/10.1126/science.aau8956). URL [http://dx.doi.org/10.1126/
2319 science.aau8956](http://dx.doi.org/10.1126/

2319 science.aau8956).

2320 Alex P. Vaz, John H. Wittig, Sara K. Inati, and Kareem A. Zaghloul. Replay of cortical spiking
2321 sequences during human memory retrieval. *Science*, 367(6482):1131–1134, Mar 2020. ISSN 1095-
2322 9203. [10.1126/science.aba0672](https://doi.org/10.1126/science.aba0672). URL <http://dx.doi.org/10.1126/science.aba0672>.

2323 Eszter Vértés and Maneesh Sahani. A neurally plausible model learns successor represen-
2324 tations in partially observable environments. In H. Wallach, H. Larochelle, A. Beygelz-
2325 imer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural In-
2326 formation Processing Systems 32*, pages 13714–13724. Curran Associates, Inc., 2019.

2327 URL <http://papers.nips.cc/paper/9522-a-neurally-plausible-model-learns-successor->
2328 [representations-in-partially-observable-environments.pdf](http://papers.nips.cc/paper/9522-a-neurally-plausible-model-learns-successor-representations-in-partially-observable-environments.pdf).

2329 Shaoming Wang, Samuel Franklin Feng, and Aaron Bornstein. Mixing memory and desire: How mem-
2330 ory reactivation supports deliberative decision-making. *PsyArXiv*, dec 2020. [10.31234/osf.io/5vksj](https://doi.org/10.31234/osf.io/5vksj).
2331 URL <https://doi.org/10.31234/osf.io/5vksj>.

2332 Ziyu Wang, Victor Bapst, Nicolas Heess, Volodymyr Mnih, Remi Munos, Koray Kavukcuoglu, and
2333 Nando de Freitas. Sample Efficient Actor-Critic with Experience Replay. *arXiv e-prints*, art.
2334 arXiv:1611.01224, November 2016.

2335 Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine Learning*, 8(3-4):279–292, 1992.

2336 Greg Wayne, Chia-Chun Hung, David Amos, Mehdi Mirza, Arun Ahuja, Agnieszka Grabska-Barwinska,
2337 Jack Rae, Piotr Mirowski, Joel Z. Leibo, Adam Santoro, Mevlana Gemici, Malcolm Reynolds, Tim
2338 Harley, Josh Abramson, Shakir Mohamed, Danilo Rezende, David Saxton, Adam Cain, Chloe Hillier,
2339 David Silver, Koray Kavukcuoglu, Matt Botvinick, Demis Hassabis, and Timothy Lillicrap. Unsu-
2340 pervised Predictive Memory in a Goal-Directed Agent. *arXiv e-prints*, page arXiv:1803.10760, March
2341 2018. URL <https://ui.adsabs.harvard.edu/abs/2018arXiv180310760W>.

2342 James C.R. Whittington, Timothy H. Muller, Shirley Mark, Guifen Chen, Caswell Barry, Neil Burgess,
2343 and Timothy E.J. Behrens. The Tolman-Eichenbaum Machine: Unifying space and relational mem-
2344 ory through generalization in the hippocampal formation. *Cell*, 183(5):1249 – 1263.e23, 2020. ISSN
2345 0092-8674. [10.1016/j.cell.2020.10.024](https://doi.org/10.1016/j.cell.2020.10.024). URL <https://doi.org/10.1016/j.cell.2020.10.024>.

2346 Andrew M. Wikenheiser and Aaron David Redish. The balance of forward and backward
2347 hippocampal sequences shifts across behavioral states. *Hippocampus*, 23(1):22–29, 2013.
2348 doi.org/10.1002/hipo.22049. URL <https://doi.org/10.1002/hipo.22049>.

2349 Andrew M Wikenheiser and Aaron David Redish. Decoding the cognitive map: ensemble hippocampal
2350 sequences and decision making. *Current Opinion in Neurobiology*, 32:8–15, Jun 2015a. ISSN 0959-
2351 4388. [10.1016/j.conb.2014.10.002](https://doi.org/10.1016/j.conb.2014.10.002). URL <http://dx.doi.org/10.1016/j.conb.2014.10.002>.

2352 Andrew M. Wikenheiser and Aaron David Redish. Hippocampal theta sequences reflect current goals.
2353 *Nature Neuroscience*, 18(2):289–294, 2015b. [10.1038/nn.3909](https://doi.org/10.1038/nn.3909). URL <https://doi.org/10.1038/nn.3909>.
2354 [3909](https://doi.org/10.1038/nn.3909).

2355 Andrew M. Wikenheiser and Geoffrey Schoenbaum. Over the river, through the woods: cognitive maps
2356 in the hippocampus and orbitofrontal cortex. *Nature Reviews Neuroscience*, 17(8):513–523, 2016.
2357 [10.1038/nrn.2016.56](https://doi.org/10.1038/nrn.2016.56). URL <https://doi.org/10.1038/nrn.2016.56>.

2358 Andrew M. Wikenheiser, Y. Marrero-Garcia, and G. Schoenbaum. Suppression of ventral hippocampal
2359 output impairs integrated orbitofrontal encoding of task structure. *Neuron*, 95(5):1197–1207.e3, 2017.
2360 [10.1016/j.neuron.2017.08.003](https://doi.org/10.1016/j.neuron.2017.08.003). URL <https://doi.org/10.1016/j.neuron.2017.08.003>.

2361 M. Wilson and Bruce L. McNaughton. Reactivation of hippocampal ensemble memories during sleep.
2362 *Science*, 265(5172):676–679, Jul 1994. ISSN 1095-9203. [10.1126/science.8036517](https://doi.org/10.1126/science.8036517). URL <http://dx.doi.org/10.1126/science.8036517>.
2363 doi.org/10.1126/science.8036517.

2364 Robert C. Wilson, Yuji K. Takahashi, Geoffrey Schoenbaum, and Yael Niv. Orbitofrontal
 2365 cortex as a cognitive map of task space. *Neuron*, 81(2):267–279, 2014. ISSN 0896-
 2366 6273. [10.1016/j.neuron.2013.11.005](https://doi.org/10.1016/j.neuron.2013.11.005). URL [http://www.sciencedirect.com/science/article/
 2367 pii/S0896627313010398](http://www.sciencedirect.com/science/article/pii/S0896627313010398).

2368 Elliott G. Wimmer and Christian Büchel. Learning of distant state predictions by the orbitofrontal
 2369 cortex in humans. *Nature Communications*, 10(1), Jun 2019. ISSN 2041-1723. [10.1038/s41467-019-
 2370 10597-z](https://doi.org/10.1038/s41467-019-10597-z). URL <http://dx.doi.org/10.1038/s41467-019-10597-z>.

2371 G. Elliott Wimmer and Christian Büchel. Reactivation of pain-related patterns in the hippocampus
 2372 from single past episodes relates to successful memory-based decision making. *bioRxiv*, May 2020.
 2373 [10.1101/2020.05.29.123893](https://doi.org/10.1101/2020.05.29.123893). URL <http://dx.doi.org/10.1101/2020.05.29.123893>.

2374 Lennart Wittkuhn and Nicolas W. Schuck. Dynamics of fMRI patterns reflect sub-second activation
 2375 sequences and reveal replay in human visual cortex. *Nature Communications*, 12(1):1795, 2021.
 2376 [10.1038/s41467-021-21970-2](https://doi.org/10.1038/s41467-021-21970-2). URL <https://doi.org/10.1038/s41467-021-21970-2>.

2377 Sasha M. Wolosin, Dagmar Zeithamova, and Alison R. Preston. Reward modulation of hippocam-
 2378 pal subfield activation during successful associative encoding and retrieval. *Journal of Cognitive
 2379 Neuroscience*, 24(7):1532–1547, 2020/12/11 2012. [10.1162/jocn{_}a{_}00237](https://doi.org/10.1162/jocn.2012.00237). URL [https:
 2380 //doi.org/10.1162/jocn_a_00237](https://doi.org/10.1162/jocn_a_00237).

2381 Emma R. Wood, Paul A. Dudchenko, R. Jonathan Robitsek, and Howard Eichenbaum. Hippocampal
 2382 neurons encode information about different types of memory episodes occurring in the same location.
 2383 *Neuron*, 27(3):623–633, Sep 2000. ISSN 0896-6273. [10.1016/s0896-6273\(00\)00071-4](https://doi.org/10.1016/S0896-6273(00)00071-4). URL [http:
 2384 //dx.doi.org/10.1016/S0896-6273\(00\)00071-4](http://dx.doi.org/10.1016/S0896-6273(00)00071-4).

2385 Chun-Ting Wu, Daniel Haggerty, Caleb Kemere, and Daoyun Ji. Hippocampal awake replay in fear
 2386 memory retrieval. *Nature Neuroscience*, 20(4):571–580, Feb 2017. ISSN 1546-1726. [10.1038/nm.4507](https://doi.org/10.1038/nn.4507).
 2387 URL <http://dx.doi.org/10.1038/nm.4507>.

2388 Qingyang Wu, Zhenzhong Lan, Jing Gu, and Zhou Yu. Memformer: The memory-augmented trans-
 2389 former. *arXiv e-prints*, page arXiv:2010.06891, October 2020. URL [https://ui.adsabs.harvard.
 2390 edu/abs/2020arXiv201006891W](https://ui.adsabs.harvard.edu/abs/2020arXiv201006891W).

2391 Jai Y. Yu and Loren M. Frank. Hippocampal–cortical interaction in decision making. *Neurobiology
 2392 of Learning and Memory*, 117:34–41, Jan 2015. ISSN 1074-7427. [10.1016/j.nlm.2014.02.002](https://doi.org/10.1016/j.nlm.2014.02.002). URL
 2393 <http://dx.doi.org/10.1016/j.nlm.2014.02.002>.

2394 Jai Y. Yu, Daniel F. Liu, Adrianna Loback, Irene Grossrubatscher, and Loren M. Frank. Specific
 2395 hippocampal representations are linked to generalized cortical representations in memory. *Nature
 2396 Communications*, 9(2209), Jun 2018. ISSN 2041-1723. [10.1038/s41467-018-04498-w](https://doi.org/10.1038/s41467-018-04498-w). URL [http:
 2397 //dx.doi.org/10.1038/s41467-018-04498-w](http://dx.doi.org/10.1038/s41467-018-04498-w).

2398 Hui Zhang, Lorena Deuker, and Nikolai Axmacher. Replay in humans - first evidence and open
 2399 questions. In N. Axmacher and B. Rasch, editors, *Cognitive Neuroscience of Memory Consolidation*,
 2400 pages 251–263. Springer, 2017. [10.1007/978-3-319-45066-7_15](https://doi.org/10.1007/978-3-319-45066-7_15). URL [https://doi.org/10.1007/
 2401 978-3-319-45066-7_15](https://doi.org/10.1007/978-3-319-45066-7_15).

- 2402 Hui Zhang, Juergen Fell, and Nikolai Axmacher. Electrophysiological mechanisms of human memory
2403 consolidation. *Nature Communications*, 9(1), Oct 2018. ISSN 2041-1723. [10.1038/s41467-018-06553-](https://doi.org/10.1038/s41467-018-06553-y)
2404 [y](https://doi.org/10.1038/s41467-018-06553-y). URL <http://dx.doi.org/10.1038/s41467-018-06553-y>.
- 2405 Shangdong Zhang and Richard S. Sutton. A Deeper Look at Experience Replay. *arXiv e-prints*, art.
2406 arXiv:1712.01275, December 2017.
- 2407 Mark C. Zielinski, Wenbo Tang, and Shantanu P. Jadhav. The role of replay and theta sequences in
2408 mediating hippocampal-prefrontal interactions for memory and cognition. *Hippocampus*, Jan 2018.
2409 ISSN 1050-9631. [10.1002/hipo.22821](https://doi.org/10.1002/hipo.22821). URL <http://dx.doi.org/10.1002/hipo.22821>.