

Replay in minds and machines

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 have been found to 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 have significantly expanded our understanding of the potential computational benefits replay may provide to biological and artificial agents alike. We provide an overview of findings in replay research from neuroscience and machine learning and summarize the computational benefits an agent can gain from replay that cannot be achieved through direct interactions with the world itself. These benefits include faster learning and data efficiency, less forgetting, prioritizing important experiences, as well as improved planning and generalization. In addition to the benefits of replay for improving an agent's decision-making policy, we highlight the less-well studied aspect of replay in representation learning, wherein replay could provide a mechanism to learn the structure and relevant aspects of the environment. Thus, replay might help the agent to build task-appropriate state representations.

Keywords: replay, reinforcement learning, representation learning, decision-making, hippocampus

22 The power of the mind goes far beyond observing and reacting to our surround-
23 ings. Yet, a precise understanding of how the brain can leave the present and enter
24 the realms of memory, planning and imagination is still lacking. Among the most
25 studied aspects is the idea that the hippocampus is central in this process, due to its
26 involvement in retrieving memories when we remember but also when we predict the
27 future (e.g., [Squire, 1992](#); [Buckner, 2010](#)). Early indications that the hippocampus
28 gives rise to memory functioning came from studies of lesion patients, such as patient
29 H.M. ([Scoville & Milner, 1957](#)), and studies of rodent spatial navigation ([O’Keefe &
30 Nadel, 1978](#)). Of particular importance were recordings from hippocampal pyrami-
31 dal neurons in rodents that demonstrated spatial firing selectivity when the animals
32 navigated in a spatial environment like a maze or linear track ([O’Keefe & Dostro-
33 vsky, 1971](#); [O’Keefe & Conway, 1978](#)). These spatially modulated neurons – later
34 also found in humans ([Ekstrom et al., 2003](#)) – became known as *place cells* and have
35 since been regarded as a core neural substrate for a *cognitive map* of physical space
36 that supports spatial navigation ([O’Keefe & Nadel, 1974, 1978](#); [Moser et al., 2008](#)),
37 as well as memory ([Cohen & Eichenbaum, 1993](#); [Redish & Touretzky, 1998](#)). How
38 exactly does the brain use this hippocampal machinery to invoke previous events in
39 the service of remembering the past and predicting the future?

40 1 Replay in the brain

41 Earlier theoretical and computational work, dating back to [Marr \(1971\)](#), proposed
42 that the brain requires a mechanism mediating the consolidation of recent memory
43 traces into an aggregated memory storage via *reactivation*. The idea that reactiva-
44 tion may help memory formation and should occur particularly when the brain
45 would be undisturbed by external input, e.g., during immobility or sleep, has sub-
46 sequently found strong theoretical and empirical support ([Buzsáki, 1989](#); [Pavlides
47 & Winson, 1989](#); [McClelland et al., 1995](#)). A central development was that techno-
48 logical advances in multi-unit recordings in rodents led to the discovery of *replay*.
49 During periods of rest and sleep, hippocampal cells reactivated sequentially in fast
50 bursts, as if retracing paths the animal had taken during wakefulness ([Wilson & Mc-
51 Naughton, 1994](#); [Skaggs & McNaughton, 1996](#); [Kudrimoti et al., 1999](#); [Nádasdy et
52 al., 1999](#); [Gerrard et al., 2001](#); [Lee & Wilson, 2002](#); [Louie & Wilson, 2001](#); for reviews
53 of these earlier findings, see [Redish, 1999](#); [Sutherland & McNaughton, 2000](#)). This

54 observation was followed by a wealth of findings that established the now classic view
55 of replay: the sequential reactivation of hippocampal place cells that occurs during
56 sleep or rest, reflects previous experience in spatial navigation tasks, and happens
57 on a temporally compressed timescale (for review, see e.g., [Foster, 2017](#)).

58 Over the following decades, much more became known about the biological as-
59 pects of replay, and today many studies point to an integral role in memory for-
60 mation as envisioned by [Marr \(1971\)](#). Replay is commonly detected during brief,
61 high-frequency oscillations of approximately 150 to 250 Hz that last about 50 to
62 200 milliseconds, called sharp wave-ripples (SWRs) (for review, see e.g., [Buzsáki,](#)
63 [2015](#); [Joo & Frank, 2018](#)). SWRs have also been found in human medial tempo-
64 ral lobe (MTL) ([Bragin et al., 1999](#); [Staba et al., 2002](#)) and have been linked to
65 memory consolidation during sleep and rest ([Axmacher et al., 2008](#); [Staresina et al.,](#)
66 [2015](#); [Zhang et al., 2018](#); [Helfrich et al., 2019](#)), as well as awake episodic memory
67 retrieval ([Norman et al., 2019](#); [Vaz et al., 2019, 2020](#)). Replay is much faster than
68 wakeful experience, and this temporal compression is believed to induce the condi-
69 tions that drive learning and the strengthening of memory traces through synaptic
70 plasticity ([Bliss & Collingridge, 1993](#); [Magee, 1997](#); [King et al., 1999](#)). A link to
71 memory is also supported by findings showing that the selective disruption of SWRs
72 in rodents during post-task rest slows learning in hippocampus-dependent spatial
73 memory tasks ([Girardeau et al., 2009](#); [Ego-Stengel & Wilson, 2010](#); [Jadhav et al.,](#)
74 [2012](#)). Finally, interactions between the hippocampus and prefrontal cortex (PFC)
75 during replay events support the idea of consolidating reactivated memories in the
76 brain (for reviews, see e.g., [Zielinski et al., 2018](#); [Tang & Jadhav, 2019](#))

77 The study of replay in humans has proven to be more difficult compared to ro-
78 dents because the non-invasive detection of fast and anatomically localized sequential
79 neural events is challenging; but existing findings point to similar conclusions. Mem-
80 ory benefits of non-sequential reactivation during rest or sleep are well documented
81 in humans ([Staresina et al., 2013](#); [Deuker et al., 2013](#); [Tambini & Davachi, 2013](#);
82 [Tambini et al., 2010](#); [Gruber et al., 2016](#)). Memory consolidation can also be bi-
83 ased by presenting learning-associated sensory cues during replay-associated sleep
84 phases in humans ([Oudiette & Paller, 2013](#); [Lewis & Bendor, 2019](#)), as in rodents
85 ([Bendor & Wilson, 2012](#); [Rothschild et al., 2016](#)). Recent progress in neuroimaging
86 analyses have also been able to capture the sequentiality of fast replay events using

87 magnetoencephalography (MEG) (Kurth-Nelson et al., 2016; Liu et al., 2020) and
88 functional magnetic resonance imaging (fMRI) (Schuck & Niv, 2019; Wittkuhn &
89 Schuck, 2021). In combination, these findings have demonstrated that replay exists
90 in a variety of species and support the idea that it reflects a consolidation process
91 that strengthens memory associations (for reviews, see Sutherland & McNaughton,
92 2000; Rasch & Born, 2007; O’Neill et al., 2010; Diekelmann & Born, 2010; Carr et
93 al., 2011; Zhang et al., 2017; Tambini & Davachi, 2019).

94 While the above findings have established a foundational knowledge of replay, our
95 understanding of this phenomenon has undergone significant and continued change
96 that sometimes challenged the classic picture of replay (for review, see e.g., Foster,
97 2017). For instance, replay has been found to be significantly more frequent than
98 initially thought, happening not only during sleep or rest but also during brief wakeful
99 pauses from active behavior (Csicsvari et al., 2007; Foster & Wilson, 2006; Diba &
100 Buzsáki, 2007; Davidson et al., 2009; Karlsson & Frank, 2009; for reviews of awake
101 replay, see e.g., Carr et al., 2011; Tambini & Davachi, 2019). In addition, replay-like
102 sequential reactivation patterns occur at various speeds, from highly accelerated to
103 much slower behavioral timescales (Deng et al., 2020; Denovellis et al., 2020; Tang
104 et al., 2021). Recordings outside of the hippocampus have also made it possible to
105 identify replay-like phenomena in a number of other brain areas, including entorhinal
106 (Ólafsdóttir et al., 2016, 2017; O’Neill et al., 2017; Trettel et al., 2019), prefrontal
107 (Euston et al., 2007; Peyrache et al., 2009; Jadhav et al., 2016; Yu et al., 2018; Shin et
108 al., 2019; Kaefer et al., 2020; Tang et al., 2021), visual and auditory sensory cortices
109 (Ji & Wilson, 2006; Rothschild et al., 2016; Wittkuhn & Schuck, 2021), parietal
110 cortex (Qin et al., 1997; Hoffman & McNaughton, 2002; Harvey et al., 2012), motor
111 cortex (Ramanathan et al., 2015; Gulati et al., 2017), and ventral striatum (Lansink
112 et al., 2009, 2008; Pennartz, 2004; Gomperts et al., 2015). Moreover, replay is
113 not necessarily a faithful replication of previous behavioral sequences but can also
114 reverse the order of experiences (Csicsvari et al., 2007; Foster & Wilson, 2006; Diba
115 & Buzsáki, 2007; Davidson et al., 2009; Karlsson & Frank, 2009) or change the
116 order of actual experiences according to a learned task rule (Liu et al., 2019). It can
117 also represent remote, non-local and never-experienced locations (Karlsson & Frank,
118 2009; Gupta et al., 2010; Ólafsdóttir et al., 2015), reflect non-spatial and partially
119 observable task features (Schuck & Niv, 2019), and occur even after tasks without

120 explicit memory requirements (Wittkuhn & Schuck, 2021). Collectively, these results
121 suggest that replay is involved in a much broader range of cognitive functions than
122 memory consolidation and spatial navigation alone.

123 Indeed, our understanding of hippocampal place cells, and the neural architec-
124 ture underlying memory and spatial navigation more generally, has also evolved
125 considerably. The “places” represented by hippocampal neurons are not exclusively
126 determined by location in physical space, but can also incorporate other task-relevant
127 aspects, such as sounds (Aronov et al., 2017) and time (MacDonald et al., 2011; see
128 also O’Keefe & Krupic, 2021). Next to place cells, many other cells have been de-
129 scribed that contribute to the cognitive map (for review, see e.g., Grieves & Jeffery,
130 2017), for instance grid cells and object vector cells in the medial entorhinal cortex
131 (Hafting et al., 2005; Fyhn et al., 2007; Høydal et al., 2019), and head direction cells
132 in the subiculum (Ranck, 1984; Taube et al., 1990). Other studies have pointed out
133 that the hippocampus may encode representations that are predictive of future loca-
134 tions, so called successor representations (SRs), that can be used for reinforcement
135 learning (RL) (Stachenfeld et al., 2017), and that grid-like patterns in the entorhi-
136 nal cortex and ventromedial PFC may represent coordinates of a non-spatial space
137 (Constantinescu et al., 2016). Thus, today the cognitive map in the hippocampal-
138 entorhinal system is often thought to represent relationships of locations and events
139 beyond physical space, from conceptual knowledge to social cognition (for reviews
140 and perspectives, see Kaplan et al., 2017; Epstein et al., 2017; Schafer & Schiller,
141 2018; Behrens et al., 2018; Bellmund et al., 2018; Peer et al., 2020; Bottini & Doeller,
142 2020; Spiers, 2020). Map-like representations also exist beyond the MTL, most no-
143 tably in the medial prefrontal and orbitofrontal cortex (OFC) (Wilson et al., 2014;
144 Schuck et al., 2016; Constantinescu et al., 2016, for a review, see e.g., Schuck et al.,
145 2018), suggesting that relational information derived from past events is also stored
146 in neocortical brain areas. As we will detail below, these findings may have important
147 implications for our understanding of the nature of replayed representations.

148 How can such a diverse set of findings about replay in the hippocampus and
149 the rest of the brain be integrated? We argue that insights into this question can
150 be gained by considering the machine learning (ML) literature, where “experience
151 replay” was introduced in the early 90s (Lin, 1991). More recently, experience replay
152 has become particularly popular after its importance for training deep neural net-

153 works (DNNs) to play Atari video games became clear (Mnih et al., 2015). This led
154 experience replay to rise to prominence as a crucial ingredient in building human-
155 level intelligence in artificial agents (Kumaran et al., 2016). Despite the conceptual
156 similarity of biological and artificial replay, research on this subject in neuroscience
157 and ML has progressed largely in parallel. Here, we aim to connect insights from
158 both research fields and review computational perspectives taken in ML on the re-
159 play phenomenon. Our goal is to highlight the diversity of possible functions that
160 might be served by a replay mechanism and attempt to answer the question of *why*
161 agents would replay in the first place.

162 2 Computational benefits of replay

163 If replay is not just a phenomenon observed in the brain, but is increasingly im-
164 plemented in artificial agents too, this raises a number of questions: Is replay in
165 machines similar to replay in animals? Why do machines need replay? Do animals
166 and machines have the same reasons to employ this process? Here, we will consider
167 these questions, compare the roles of replay in both biological and artificial agents,
168 and distill the most significant benefits of replay.

169 Before we begin, we would like to point out some of the more significant aspects
170 in which the concept of replay differs between ML and neuroscience. As described
171 above, neuroscience emphasizes the sequential and often accelerated nature of replay
172 (Genzel et al., 2020), but often subsumes reactivation of memories as well as activa-
173 tion of internally generated neural sequences under the same term. ML methods, in
174 contrast, often distinguish between experience replay, which corresponds to sampling
175 from a memory buffer, and model-based methods, often referred to as planning, in
176 which the agent internally generates new experiences from a learned model of the
177 environment. Moreover, several methods in ML focus on replaying sets of individ-
178 ual transitions (akin to only two adjacent place cells), which do not have to form
179 sequences spanning long trajectories through the environment. Finally, the issue of
180 replay speed has not been a major consideration in ML. Bound to physical inter-
181 action with the environment and the timescales of biology, however, the speed of
182 replay takes a more important role in understanding animal cognition than it has in
183 artificial systems.

184 Similar to previous work (e.g., Foster & Knierim, 2012; Cazé et al., 2018; Mo-

185 [mennejad, 2020](#)), our review is guided by the RL framework ([Sutton & Barto, 2018](#)),
186 which allows parallels to be drawn between reactivation of neural patterns in biolog-
187 ical agents and replay of task states in artificial agents. The RL framework considers
188 agents that learn from interactions with their environment and thereby gather expe-
189 riences one at a time. RL techniques are designed to learn from experience gradually,
190 through trial-and-error, using every new experience immediately to adjust the agent’s
191 knowledge about the task. This has the benefit of accruing knowledge without delay,
192 while integrating information over all experiences gained so far, rather than using
193 just the most recent experience to make decisions. Typically, small adjustments
194 are made to the agent’s knowledge with each new experience because large updates
195 risk overwriting the effects of earlier learning and can limit generalization. [Box 1](#)
196 describes the fundamental aspects of RL.

197 Which benefits can an agent obtain from using replay? In the next sections, we
198 will discuss five potential computational functions of replay: increasing the speed and
199 data efficiency of learning, reducing forgetting, reorganizing experiences, planning,
200 and generalization. We do not consider these functions to be entirely separable. We
201 distinguish them because they each offer a unique perspective on what an intelligent
202 agent, biological or artificial, stands to gain from replaying past experiences. This
203 perspective also sheds light on why replay can have different properties in different
204 study contexts, which have found replay to be sometimes backward and sometimes
205 forward, or in some cases to occur immediately and in other cases long after the
206 experience was acquired. In doing so, we try to emphasize how the research fields
207 of neuroscience and ML can learn from each other, identifying commonalities that
208 point to overarching computational principles and delineating differences that could
209 cross-fertilize novel ideas in both fields (for similar perspectives, see e.g., [Hassabis et](#)
210 [al., 2017](#); [Botvinick et al., 2020](#)).

211 Finally, we will consider an underexplored computational aspect of replay. Most
212 research to date has studied how replay may be used to optimize goal-directed behav-
213 ior *given* an already established cognitive map. Here, we explore whether replay may
214 also be involved in shaping the cognitive map itself. We hypothesize that the content
215 and function of replay is determined by its interplay with the current representation
216 of the task and the representational demands of the task at hand, a notion which
217 has recently received some computational ([Russek et al., 2017](#); [Caselles-Dupré et al.,](#)

218 2019; Momennejad, 2020) as well as empirical support (Schuck & Niv, 2019). In this
219 view, replay can be understood not only as a phenomenon that retrieves relational
220 information stored in a cognitive map, but also as a process that changes relational
221 information and internal state representations of an agent (see Box 1 for a definition
222 of state representations).

223 **Box 1 What is reinforcement learning?**

224

225 Reinforcement learning (RL) theory provides a formal framework to describe how
226 agents learn to optimize their behavior through interactions with the environment
227 that yield rewards or punishments and can result in transitions from one situation
228 (or “state”) to the next (Sutton & Barto, 2018). Agents learn which actions to take
229 in which state based on the long-term reward they get. Actions that lead to a net
230 positive sum of future outcomes will be performed more often, whereas actions that
231 lead to a net negative sum of outcomes will be performed less often. At the heart of
232 RL lies the assumption that the agent-environment interaction can be modelled as
233 a Markov decision process (MDP). The MDP consists of the following components:

- 234 • An **environment** that is represented by a discrete set of states S (e.g., locations
235 in a maze or positions in a board game). The agent is assumed to occupy one of
236 those states at each point in time. Note that the states refer to the agent’s internal
237 representation of the environment, which may not be equivalent to the objective,
238 “true” state of the environment. In some environments it can be sufficient to form
239 an internal state representation based on incoming sensory signals alone, but in
240 many cases this approach is not feasible and instead requires to supplement the
241 state definition with non-observable information (e.g. Wilson et al., 2014; Schuck
242 et al., 2018). The question of how an agent can know how to internally represent
243 the state is therefore a topic of great interest in RL research, and will be discussed
244 later in this review. Importantly, by describing the environment as a state space
245 of an MDP, the states can represent any type of task, whether spatial locations in
246 a maze or task states indicated by objects in a non-spatial decision-making task.
247 Note that MDPs can also have continuous state or action spaces, but for simplicity,
248 we only consider finite and discrete state and action spaces here.

- 249 • A discrete set of **actions** A the agent can perform in each state (e.g., movement
250 into a particular direction in a maze).
- 251 • A **state transition function** $T(s_t, a_t, s_{t+1})$ describing the probabilities of moving
252 from the current state $s_t \in S$ to the next state $s_{t+1} \in S$ after taking an action
253 $a_t \in A$. The state transition function can be understood as a cognitive map that
254 entails the connections between states (e.g., route from location A to location B).
- 255 • A **reward function** $R(s_t, a_t, s_{t+1})$ maps each state-action pair to a scalar rein-
256 forcement signal r , received after taking an action a_t in state s_t . Generally, rewards
257 are controlled by the environment, and therefore lie outside of the agent. Still, the
258 value of particular outcomes are inferred by the agent, and learning therefore de-
259 pends on an internal representation of the reward (e.g., [Juechems & Summerfield,](#)
260 [2019](#)). This internal representation can be subject to further transformations such
261 as reward normalization, which in turn may depend on the agent's state.

262 In an MDP, the agent-environment interaction is assumed to be Markovian with
263 respect to reward (and state), which means that the state and reward at the next
264 time point $t + 1$ depend only on the state and action of the current time point t ,
265 but not on any states or actions before. The current state s_t therefore contains all
266 information from the previous history. Formally, the Markov property can be stated
267 as:

$$P(s_{t+1}, r_t | s_t, a_t, s_{t-1}, a_{t-1}, \dots, s_0, a_0) = P(s_{t+1}, r_t | s_t, a_t) \quad (1)$$

268 In brief, this means that we can think of learning from trial-and-error as the
269 following process: the agent represents the current state of the environment, s_t
270 and then performs an action a_t . The action will affect the environment, changing
271 the agent's state from s_t to s_{t+1} (as described by the transition function T), and
272 potentially yield a reward, as described by the reward function R . The agent's goal
273 is to always perform the actions that maximize the sum of total rewards obtained
274 over the course of its interaction with the environment.

275 How can the agent know which actions to perform that will achieve this goal?
276 A first step is to formalize the agent's goal. The goal of the agent is expressed as
277 the expectation over the discounted sum of its future rewards: $E[\sum_{i=0}^{\infty} \gamma^i r_i]$, where
278 i indexes all future time points, and $\gamma \in [0, 1]$ is known as the discount factor which

279 attenuates the influence of distal rewards and ensures that the sum of rewards has a
 280 finite value. In RL, the agent is assumed to gather knowledge about this discounted
 281 sum of rewards and store it as a *value function*. The agents should therefore hope
 282 to eventually have a very accurate value function, which reflects the true sum of
 283 discounted rewards. Future rewards will depend on the agent’s behavior which is
 284 reflected in the agents’ value function, defined as the expected reward after taking
 285 action a_t from state s_t , then following policy π thereafter:

$$Q^\pi(s_t, a_t) = E_\pi \left[\sum_{i=0}^{\infty} \gamma^i r_i \mid s_0 = s_t, a_0 = a_t \right] \quad (2)$$

286 which can be written recursively as:

$$Q_\pi(s_t, a_t) = \sum_{s_{t+1}} T(s_t, a_t, s_{t+1}) [R(s_t, a_t, s_{t+1}) + \gamma Q^\pi(s_{t+1}, \pi(s_{t+1}))] \quad (3)$$

287 Equation 3 is known as the Bellman equation, and its solution is obtained through
 288 an iterative bootstrapping process by first estimating the Q -values, then repeatedly
 289 updating the Q -values using the Bellman equation until they converge to the true
 290 values. This approach requires knowledge of a model of the environment (i.e., state
 291 transition function), and such methods are known as model-based RL (Doya, 1999;
 292 Daw et al., 2005). However, in many situations the agent does not have such a
 293 model, thus a separate category of algorithms known as model-free RL is used. These
 294 model-free methods utilize a combination of sampling and bootstrapping to compute
 295 the value functions directly without the MDP model. The best-known example is
 296 the temporal difference (TD) learning algorithm, which adjusts the value function
 297 incrementally using a prediction error (PE) capturing the discrepancy between the
 298 actual reward and the estimated state value. A widely-used variant of TD learning
 299 is the Q -learning algorithm:

$$Q_{k+1}(s_t, a_t) = Q_k(s_t, a_t) + \alpha [r_t + \gamma \max_{a_{t+1}} Q_k(s_{t+1}, a_{t+1}) - Q_k(s_t, a_t)] \quad (4)$$

300 The Q -learning algorithm assumes that the agent picks the best action on the next
 301 step based on its current knowledge. The resulting Q -value, i.e., $\max_{a_{t+1}} Q_k(s_{t+1}, a_{t+1})$,
 302 is then used as the target to update the current Q -value in iteration k , which circum-
 303 vents the necessity to maintain a transition function as in model-based approaches.
 304 The learning rate $\alpha \in [0, 1]$ determines the extent of adjustment in each update.

305 2.1 Faster learning and data efficiency

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

317 Our brain arguably faces a similar challenge. We, and other animals, often have
318 to learn directly from the outcomes of our decisions. Yet, repeating errors can pose
319 actual risks to us, which limits the usefulness of exclusively relying on a slow, trial-
320 and-error-based learning mechanism. More generally, the number of experiences we
321 acquire with a particular situation in a lifetime is quite limited in relation to the
322 complexity of our environments and our brain, which contains approximately 10^{14}
323 synapses (Tang et al., 2001). In order to make thousands of gradual adjustments to
324 each of these synapses, the ability to reuse experience efficiently is paramount.

325 In the RL literature, “experience replay” was initially introduced to address these
326 issues of slow learning and data inefficiency (Lin, 1991, 1992, 1993). In his seminal
327 paper, Lin (1992) wrote that “[...] Q -learning algorithms [...] are inefficient in
328 that experiences obtained by trial-and-error are utilized to adjust the networks only
329 once and then thrown away. [...] Experiences should be reused in an effective way.”
330 (Lin, 1992, p. 299). Lin (1992) proposed that experiences can be used to update
331 knowledge in a dual fashion; (1) immediately when experiences are acquired, and (2)
332 at later time points, after experience itself may have long passed. Specifically, Lin
333 (1992) proposed replaying full sequences of experiences, starting from an initial state
334 to a final state in backward order, and learning from these experiences, as if they were
335 real. Lin (1992) then showed that this is a more efficient use of data that accelerates
336 learning of an RL agent. In line with these ideas, many others have since emphasized
337 the computational benefit of replay for maximizing data efficiency and the speed of

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

348 There are several reasons why replay can help learning. The core idea of RL is
349 that an agent learns which actions to take through positive and negative outcomes.
350 In the real world, outcomes are often only obtained after a long sequence of events
351 and actions but agents still need to know how to behave at the start of the sequence,
352 as for instance, in a chess game. This problem is known in RL as the *temporal credit*
353 *assignment problem* ([Minsky, 1961](#)) and replay may help to solve it. The early work
354 by [Lin \(1991, 1992, 1993\)](#) pointed out that replay could help an agent to remember
355 the sequence of previous states and actions that led to a given outcome, and assign
356 credit for the reward to the sequence of states and actions that preceded it. This
357 also explains why sequential replay may sometimes proceed in backward order ([Lin,](#)
358 [1992](#)). In the brain, awake backward replay has indeed been frequently observed,
359 where rewarded spatial trajectories of an animal are replayed in reverse order ([Diba](#)
360 [& Buzsáki, 2007](#); [Foster & Wilson, 2006](#); [Singer & Frank, 2009](#)), and the frequency
361 of awake backward (but not forward) replay is modulated by the change in reward
362 magnitude ([Ambrose et al., 2016](#)). Another aspect is that as the agent’s knowledge
363 of the rewards becomes better with time, outcomes in the past should be re-evaluated
364 in light of this updated knowledge ([van Seijen & Sutton, 2015](#)). By comparing past
365 rewards against current value estimates, replay can serve this function.

366 In line with the idea that backward replay reflects learning through temporal
367 credit assignment, the rate of backward replay was also observed to be more fre-
368 quent in novel compared to familiar environments ([Foster & Wilson, 2006](#); [Singer &](#)
369 [Frank, 2009](#)) and to decrease its bias to reflect previous paths to the goal location as
370 a function of learning ([Shin et al., 2019](#)), suggesting that the relevant trajectory has

371 been learned and does not need to be reinforced through replay anymore (Foster &
372 Knierim, 2012). Computational work shows how backward replay can strengthen for-
373 ward synaptic pathways through spike timing dependent plasticity (STDP) (Haga
374 & Fukai, 2018) and thus support forward replay during sleep and active behavior
375 (Johnson & Redish, 2007; Pfeiffer & Foster, 2013; Wikenheiser & Redish, 2015b,
376 2013). Further evidence for the role of replay in assigning credit are findings show-
377 ing that replay is coordinated with subcortical activation of brain areas related to
378 processing reward (Gomperts et al., 2015; Lansink et al., 2009; Pennartz, 2004; Gom-
379 perts et al., 2015), which could convey reward signals to other brain regions like the
380 hippocampus. So far, existing empirical studies support the idea that awake back-
381 ward replay supports temporal credit assignment by retrieving states that led to the
382 outcome, accelerating learning for cases in which a long delay between rewards and
383 actions must be encoded.

384 Figure 1 provides an illustration of how backward replay of full sequences works
385 in the context of RL, and which effect it has on the speed of learning. We consider
386 an RL agent navigating in a square environment with 20×20 tiles that contains
387 several walls and one goal location with a reward (see Figure 1A). The agent can
388 move into one of the four cardinal directions (up, down, left, right). It receives no
389 reward for moving, a small negative reward for bumping into a wall (-0.1), and
390 a reward of 1 when arriving at the goal location. The goal of any RL agent is to
391 maximize the sum of future rewards. Thus, in this case the best policy is to navigate
392 to the reward with as little steps as possible, and without bumping into the wall.
393 This is a well-known “grid world” problem that can be solved using RL because
394 any agent using the gradual updating procedures described above will eventually
395 learn to navigate directly to the reward – but without replay this learning might be
396 painfully slow. For illustrative purposes, we use the off-policy, model-free Q -learning
397 algorithm described in Equation 4 in Box 1. The learning rate α , temperature τ and
398 discounting factor γ were arbitrarily set to 0.3, 1 and 0.99 for the purposes of this
399 illustration.

400 In each episode, the agent starts in a random position and navigates until it has
401 found the reward. After the reward is found, the agent is replaced at a new starting
402 location. The starting locations varied randomly, but in order to avoid excessive
403 fluctuations in performance caused by variability in starting locations, their distances

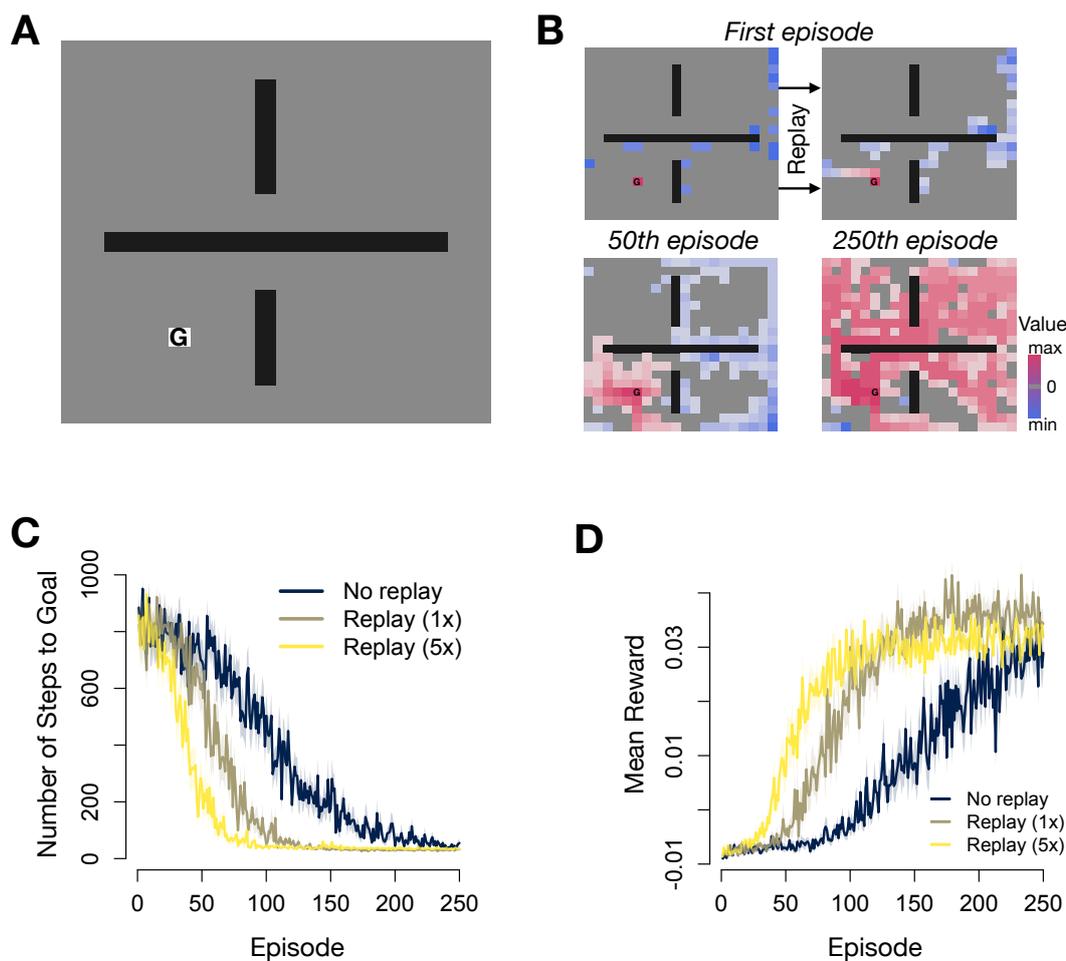


Figure 1: Replay speeds up learning to navigate to a goal in a grid world. (A) Square environment (“grid world”) with 20×20 tiles (shown in gray) that contains several walls (black tiles) and one goal location (white tile labelled “G”) that contains a reward. At the beginning of each episode, an RL agent 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/>).

404 to the reward location were constrained to lie at least 10 tiles away from the goal
405 location (in euclidean distance, i.e., considering direct flight paths). The blue lines
406 in Figure 1C–D show the number of steps the agent needs to navigate to the goal
407 location. As can be seen, the agent needs a considerable amount of experience in
408 order to learn to quickly find the goal location, about 250 episodes in which the agent
409 could explore the states between the start and goal location. Next, we investigate
410 the speed of learning when we supply the agent with a simple replay mechanism. In
411 particular, we used backward replay of entire episodes at the end of each successful
412 episode, similar to the work of Lin (1992), and in line with the considerations about
413 temporal credit assignment discussed above that backward replay can propagate
414 the value back to preceding states. This meant that upon encountering a reward,
415 the agent would not only update the Q -value of the last state. Instead, the agent
416 internally traversed backwards through the sequence of states, actions and rewards
417 until the beginning of the episode $(s_t, a_t, r_t, \dots, s_1, a_1, r_1)$, updating its Q -value at
418 each step just as if it was experiencing the full trajectory episode again. As can be
419 seen in Figure 1C–D (brown and yellow lines), adding replay dramatically speeds up
420 learning of the agent, reducing the number of interactions needed to achieve ceiling
421 performance to less than half of what was observed without replay. Note that the
422 choice to replay the full sequence of states, actions and rewards between the start
423 location and the goal location is not without consequence, and a variety of different
424 definitions of what constitutes an episode are common in RL and neuroscience. We
425 consider them in Box 2 and will discuss their implications further below.

426 One notable aspect of our example is that replayed experiences are used to update
427 Q -values, which slowly integrate knowledge over many experiences to control behav-
428 ior. This is different from relying directly on specific single episodes for behavioral
429 control, an approach often used in the context of *episodic RL* (for recent reviews, see
430 e.g., Gershman & Daw, 2017; Botvinick et al., 2019). During episodic RL, specific
431 single episodes are stored in memory and retrieved to directly determine behavior
432 when the same or a similar situation is encountered again (Lengyel & Dayan, 2007;
433 Gershman & Daw, 2017; Botvinick et al., 2019). Episodic control also has benefits
434 for the speed of learning under particular circumstances. For example, repeating
435 memorized sequences of actions that led to desirable outcomes in the past is a rea-
436 sonable strategy for an agent that prioritizes exploitation over exploration and has

437 only limited experience with a complex and noisy environment that it might rarely
438 encounter again in the future (Lengyel & Dayan, 2007). In humans, the retrieval
439 of single experiences in decision-making is associated with the hippocampus (Lee et
440 al., 2015; Bornstein & Norman, 2017; Wimmer & Büchel, 2020) and reinstatement
441 of information from past choice trials at decision-time biases present choices towards
442 decisions made previously in the reinstated context (Duncan & Shohamy, 2016; Born-
443 stein et al., 2017; Bornstein & Norman, 2017). To what extent these retrievals of
444 single episodes are supported by sequential replay remains an open question. While
445 this direct use of experiences can be very efficient, the brain must strike a balance
446 between such efficiency (i.e., taking the same route to a goal location that was used
447 before) and the flexible updating of behavior in cases when the most efficient be-
448 havioral strategy cannot be used (i.e., the most direct path to a goal location is
449 blocked). The distinction between reusing experiences for episodic RL or updating
450 Q -values highlights the fact that there are multiple ways to instantiate replay in an
451 RL agent, and the illustration in Figure 1 only serves as a basic introduction to com-
452 putational replay. Below we will describe the many different aspects that determine
453 the usefulness of replay for different problems in more detail.

454 **Box 2 What is replayed?**

455
456 Replay is generally thought to represent previous experience. In artificial agents,
457 an experience e_t is commonly defined as a quadruple consisting of the current state
458 s_t , the taken action a_t , the reward r_t received after taking action a_t in state s_t ,
459 and the next state s_{t+1} , together $e_t = (s_t, a_t, r_t, s_{t+1})$, effectively describing a single
460 transition between two states as the atomic unit of an artificial replay event. In earlier
461 RL implementations of experience replay (Lin, 1991) and our own demonstration
462 in Figure 1, replay involves entire episodes (i.e., sequences of chained experiences
463 from the start to the goal location). In contrast, deep neural networks (DNNs) use a
464 single experience as a replayed unit to create independent and identically distributed
465 (i.i.d.) samples required for stochastic gradient descent (SGD). For example, in the
466 DNN by Mnih et al. (2015) the states S provided to the Q -function consist of pre-
467 processed versions of Atari pixel frames. In biological agents, replay is commonly
468 defined as *sequential* reactivation of neural activity patterns relating to previously

469 encoded experience (Genzel et al., 2020). What constitutes a replayed experience is a
470 much more difficult question to answer for biological agents. In non-human animals,
471 this typically involves recordings of hippocampal place cells representing physical
472 locations in a spatial environment, thus equating states to locations and episodes of
473 experiences to trajectories of locations. However, hippocampal cells appear to be
474 quite flexible in encoding task-relevant information other than physical space, for
475 instance sounds (Aronov et al., 2017), trial history (Wood et al., 2000) or run laps
476 (Sun et al., 2020). This notion is even more important in humans, where, depending
477 on the experimental design, replayed memory representations can be multi-modal and
478 thus it is not immediately apparent which elements of an experience should be stored
479 and where in the brain those might be represented. It therefore seems appropriate to
480 consider that replay not only reflects physical location in spatial environments but
481 rather a state of task-relevant information, which, in typical rodent tasks might be
482 majorly, but not exclusively, defined by physical location.

483 Which and how many experiences are stored? How long are they stored? In DNNs
484 using replay, the newest experiences are stored at each time step in a memory buffer
485 $D_t = e_1, \dots, e_t$ with finite size N (e.g., Mnih et al., 2013; Mnih et al., 2015; Zhang
486 & Sutton, 2017). Since the success of the DNN by Mnih et al. (2015), the memory
487 buffer is typically set to a size of $N = 10^6$ newest experiences, which continuously
488 replace the oldest experiences (Fedus et al., 2020; Zhang & Sutton, 2017). Recently,
489 Fedus et al. (2020) investigated the relationship between the number and age of
490 experiences stored in the memory buffer. First, they found that increased memory
491 capacity improved learning performance, likely due to a larger coverage of state-
492 action pairs (Fedus et al., 2020). Second, decreasing the age of the oldest experience
493 in the memory buffer also improved performance, likely because of older experiences
494 that resulted from policies that are inconsistent with the current on-policy decision
495 strategy, and is in line with earlier findings noting that experience replay is only
496 beneficial if it is consistent with the current decision policy (Lin, 1991, 1993). An
497 exception to this are Atari games that are characterized by sparse rewards and require
498 high levels of exploration. In such tasks sampling from older off-policy experiences
499 is still beneficial (Fedus et al., 2020). These considerations about the size of the
500 memory buffer in artificial agents point to an intriguing trade-off between old and
501 new memories: On the one hand, a youthful memory buffer storing only recent

502 experiences can effectively drive the current decision policy and quickly abandon old
503 and potentially inefficient behavior. On the other hand, keeping older experiences
504 and integrating them with recent ones may foster generalization and prevent an agent
505 from becoming stuck in a decision policy that is suboptimal.

506 While the size and content of a memory buffer in artificial agents can be crafted
507 by ML researchers, determining number and nature of memories in brains is topic
508 of ongoing debate for neuroscientists. Humans and other animals are faced with
509 a continuous stream of experiences and, given the gigantic amount of experiences a
510 biological agent acquires in a lifetime, putting a number on the memory size in brains
511 appears virtually impossible. Even if biological agents had an unlimited memory
512 storage, selecting memories for replay from that storage would become challenging
513 with a large amount of experiences. In fact, forgetting irrelevant experiences might
514 be an adaptive biological process (Hardt et al., 2013) and resonates with the idea
515 that the brain should form compact memory representations that have the highest
516 utility for behavior. Tackling the question of what counts as an event or experience, a
517 common theme in neuroscience has been to posit that the brain segments continuous
518 experience into representations of distinct neural states that transition at boundaries
519 between events or shifts in context (for reviews, see e.g., Richmond & Zacks, 2017;
520 Brunec et al., 2018; Shin & DuBrow, 2020; Bird, 2020; Maurer & Nadel, 2021), a
521 process that might also happen retroactively, *after* experiences have been obtained
522 (Clewett et al., 2019). This formation of segmented memory traces is thought to be
523 driven by various factors, including inferred changes in the environment (DuBrow
524 et al., 2017), prediction error signals elicited by reward outcomes (Rouhani et al.,
525 2020) or discontinuities in the statistical structure of the environment (Gershman
526 et al., 2014). A practical approach for human research therefore seems to be to
527 define events as “meaningful” units of experience (Bird, 2020) within the current
528 experimental paradigm, and to potentially formalize them as states in an MDP.
529 Finally, in understanding memory as a constructive process, it is important to note
530 that neural task representations may change from perception to reactivation (Favila
531 et al., 2020). This aspect is particularly crucial for the study of replay in humans,
532 because activity patterns that are expected to reactivate are commonly determined
533 based on simple localizer tasks that do not involve mnemonic task components. The
534 brain might have already transformed its input data to a representation that is

535 different from what the researcher was hoping to see re-merge from replayed activity
536 patterns.

537 2.2 Less forgetting

538 Increasing the speed of learning is an important computational benefit of replay, but
539 not the only one. Many statistical learning mechanisms were built under the assump-
540 tion that the agent encounters its environment entirely at random, and therefore can
541 learn from examples that are independent and identically distributed (i.i.d.). This is
542 often not true for the type of agent-environment interactions encountered in real life,
543 and has been a particular problem for neural networks which use stochastic gradient
544 descent (SGD) for learning, for which i.i.d. samples are one key requirement. Using
545 replay to learn from experience that is not i.i.d. has been an important determinant
546 of the renewed success of neural networks in dealing with RL-type problems (e.g.,
547 [Mnih et al., 2015](#)). Moreover, as we outline below, although RL models acknowledge
548 that successive events are dependent on each other (see [Box 1](#)), they still benefit
549 from replay, especially if an agent cannot freely sample all possible events in a given
550 environment.

551 Temporal auto-correlation in the transition structure of real-world scenarios is an
552 important property of the environment that can be hardly changed. Why does this
553 cause problems for gradual learning algorithms? While gradual learning mechanisms
554 are able to integrate experiences over longer periods, they still emphasize the most
555 recent experience, which can cause the agent to forget about important past expe-
556 riences. This is particularly apparent in a setting in which an agent engages in two
557 tasks in a *blocked* manner. If a DNN is first trained to perform a task A and subse-
558 quently trained with another task B, performance on task A drops dramatically, as
559 if the network forgot how to solve A or that learning task B interferes with what was
560 learned about task A. This problem is known as catastrophic forgetting, or catas-
561 trophic interference, and has long been recognized as a major problem in the ML
562 field ([McCloskey & Cohen, 1989](#); [Ratcliff, 1990](#); [French, 1999](#); [Hassabis & Maguire,](#)
563 [2007](#); [Kumaran et al., 2016](#); [Parisi et al., 2019](#)). Catastrophic forgetting is one of
564 the main reasons why artificial agents can usually learn a single task quite well but
565 subsequent training on a different task results in poor performance on the previously
566 learned task. This prevents the agent from achieving competencies across multiple

567 tasks, which comes relatively easily to humans. Catastrophic interference can also
568 be understood as an issue threatening the stability of a cognitive map representation
569 (Gupta et al., 2010; McClelland et al., 1995; O’Reilly & McClelland, 1994).

570 A potential solution to catastrophic interference is *interleaved learning* where new
571 experiences are interleaved with existing knowledge to reconcile competing memory
572 representations (McClelland et al., 1995). This influential idea, rooted in the com-
573plementary learning systems (CLS) theory (McClelland et al., 1995; O’Reilly et al.,
5742014; Schapiro, Turk-Browne, et al., 2017), also suggests that replay may be the
575mechanism that agents can use to “mentally” interleave past with present experi-
576ence. While the DNN tunes its connection weights to solve task A, the learning
577experience is stored as episodes in a memory buffer. During learning of task B,
578previous experience with task A is integrated during offline periods via a replay-like
579mechanism, allowing the agent to perform well on both tasks. Shin et al. (2017), for
580instance, proposed an approach that learns a generative model based on experience
581with one classification task A. When switching to an independent classification task
582B, the system is retrained using a combination of new task data and fictitious se-
583quences from the generative model, resulting in rapid generalization to the new task
584with little performance loss. Similarly, implementing replay in this way in DNNs
585can help to overcome performance deficits in incremental task learning scenarios and
586continuous task environments (Mnih et al., 2015; van de Ven & Tolias, 2018; van de
587Ven et al., 2020).

588 Humans and other animals do not necessarily seem to suffer from the compu-
589tational problem of catastrophic interference. They can solve a wide set of tasks
590throughout their lifetime, despite temporal autocorrelation of experience and even
591learn well from blocked experience which troubles DNNs (Flesch et al., 2018). The
592idea that this ability might be related to replay (Antony & Schapiro, 2019) is sup-
593ported by several studies. Karlsson and Frank (2009) for instance have observed
594replay of episodes from a remote spatial context. Reactivation studies in humans
595of previously learned events in the hippocampus that overlap with newly encoded
596memories leads to better retention (Kuhl et al., 2010).

597 Another challenge arises when learning must occur in environments where some
598events happen rarely, but are nevertheless of great significance for the agent’s suc-
599cess or well-being. Evidence that replay might be used to mitigate this problem

600 in animals comes from studies showing that actions which should be avoided will
601 be reactivated, like paths to a shock zone (Wu et al., 2017) or paths to devalued
602 outcomes (Carey et al., 2019). In addition to learning about events that should be
603 avoided, replaying rare events that are only weakly encoded could allow the agent to
604 form a stable representation of the entire environment even if only a smaller subset
605 is experienced frequently. In Gupta et al. (2010) non-local replay was stronger for
606 remote sequences if they were experienced less frequently. Using MEG in humans,
607 Jafarpour et al. (2017) showed that stronger reactivation of one of three previously
608 encoded stimuli was determined by how weakly the stimulus was attended to during
609 encoding. These findings are supported by an fMRI study by Schapiro et al. (2018),
610 who demonstrated that older, less well remembered task stimuli were selectively
611 reactivated during a subsequent rest period resulting in memory improvement, an
612 effect that was particularly strong in participants who slept in the 12-hour interval
613 between test sessions. In another study, the benefits of targeted memory reactivation
614 (TMR) were stronger for weakly learned information (Tambini et al., 2017). Further,
615 replay-associated electroencephalography (EEG) sleep spindles during a nap follow-
616 ing difficult but not easy memory encoding were related to improved subsequent
617 memory performance (Schmidt et al., 2006). Overall, the beneficial role of sleep for
618 weakly encoded information is well documented by several studies (Drosopoulos et
619 al., 2007; Peters et al., 2007; Diekelmann et al., 2010; McDevitt et al., 2015; Sio et
620 al., 2013; Kuriyama et al., 2004; Cairney et al., 2016; Talamini et al., 2008; Schapiro,
621 McDevitt, et al., 2017). Together, replay liberates an agent from needing to consider
622 transitions only in proportion to how many times they were experienced. Instead,
623 replay can flexibly increase or decrease the number of opportunities for learning
624 experiences from single episodes.

625 **2.3 Re-inventing the past**

626 In our introductory example (see Figure 1), replayed content was a close reflection
627 of past experience. Replay occurred immediately after an episode was experienced
628 and reflected past trajectories from start to finish, albeit in reverse order. This setup
629 stands in contrast to the ideas discussed in the section on forgetting, which imply
630 that replay must not necessarily respect the structure of experiences, but could, for
631 instance, change the order and frequency of events. Beyond dealing with unevenly

632 distributed events, replay could in fact be used to arbitrarily alter the distribution
633 of events upon which memory is built.

634 Such a reorganization of experience also requires a different understanding of
635 what constitutes an episode. In our simulation, we had assumed that the minimal
636 unit of replayed content is one entire sequence of states, actions and rewards that
637 occurred between a random start position and the encounter of a goal. This meant
638 that episodes were often quite long, involving several hundreds of steps particularly
639 early in learning (see Figure 1C), and that the transitions between locations had to
640 be replayed in the order in which they were experienced. In order for replay to be
641 able to reorganize experience, an episode is often divided into a much smaller unit
642 of experience, a simple sequence of just one state, one action, one reward and the
643 next state, known as a (s_t, a_t, r_t, s_{t+1}) -tuple. Arguably, replaying such minimal
644 experiences risks losing the benefits of temporal credit assignment, which come with
645 backward replay of full trajectories, because values will not necessarily propagate to
646 starting positions. But it does also have benefits, which we will discuss below. In
647 consequence, the question of what constitutes an atomic unit from the perspective
648 of replay has important implications and is therefore actively debated (see Box 2).

649 Using minimal transitions, there is a large variety of ways in which replay may al-
650 ter the structure of experiences that have been discussed in the ML and neuroscience
651 literature. One possibility is to reactivate (s_t, a_t, r_t, s_{t+1}) -tuples in a random order,
652 which artificially crafts similar conditions as during supervised learning that allows
653 artificial neural networks trained with SGD to excel (Botvinick et al., 2020). Such
654 uniformly sampled (s_t, a_t, r_t, s_{t+1}) -tuples have therefore played an important role in
655 adapting DNNs to RL problems, such as the famous Deep Q-Network (DQN) (Mnih
656 et al., 2015). Interestingly, some animal studies have also found replay of seemingly
657 random trajectories following exploration of a familiar open-field arena (Stella et al.,
658 2019). Note however, that Stella et al. (2019) still observed replay of sequentially or-
659 ganized transitions that reflected the spatial constraints of the environment, whereas
660 random replay used in ML can involve sets of single transitions that do not form
661 sequential trajectories. This highlights the different understanding of replay content
662 in ML and neuroscience. Additionally, most animal studies impose the assumption
663 of sequentiality during data analysis, and would discard fully random activation of
664 transitions as noise. In both ML and neuroscience, however, random replay refers to

665 sequential reactivation that is unrelated to previously experienced action sequences,
666 and can be seen at one extreme of a continuum describing how closely replay matches
667 actual behavioral sequences (see [Swanson et al., 2020](#), their Figure 2).

668 Another particularly important idea from the ML literature is to prioritize replay
669 of transitions that led to large surprises, with the idea that some experiences are more
670 informative than others, resulting in more efficient learning ([Schaul et al., 2015](#);
671 [Horgan et al., 2018](#)). Such prioritized replay records a prediction error (PE), the
672 difference between the expected and actual reward, for every encountered transition
673 and uses this signal to select experiences for replay later. This method is very similar
674 to, and inspired by, an earlier algorithm in model-based planning known as prioritized
675 sweeping, which selects the state to be updated according to the magnitude of the
676 change in value upon the execution of the update ([Moore & Atkeson, 1993](#); [Peng &](#)
677 [Williams, 1993](#); [Andre et al., 1998](#)). Based on the success of the prioritized replay
678 approach, more frequent sampling of transitions with a high absolute TD error is
679 now a common approach to train DNNs ([Fedus et al., 2020](#)). Using RL models,
680 [Mattar and Daw \(2018\)](#) extended previous approaches by focusing prioritization on
681 behaviorally relevant states that are likely to be encountered again in the future
682 and those transitions where a policy change would yield the largest net increase
683 in discounted future reward. Note that although prioritization algorithms assume
684 selection of replay content on the level of individual transitions, they can, under some
685 circumstances, still lead to sequential replay. This is true, for example, for the model
686 of [Mattar and Daw \(2018\)](#), because expectations about increases in future reward
687 are themselves often auto-correlated.

688 The core idea that replay should be influenced by reward and surprise is in
689 line with several animal studies. Place cell sequences associated with reward are
690 replayed more often ([Ólafsdóttir et al., 2015](#); [Foster & Wilson, 2006](#); [Bhattarai et](#)
691 [al., 2019](#)), in particular those with a high PE ([Singer & Frank, 2009](#); [Michon et](#)
692 [al., 2019](#); [Roscow et al., 2019](#)), and the rate of SWRs is also influenced by reward
693 ([Ambrose et al., 2016](#); [Singer & Frank, 2009](#)). These results highlight replay's role in
694 credit assignment, as discussed above. In human neuroimaging studies, hippocampal
695 activity is modulated by reward magnitude ([Wolosin et al., 2012](#); [Igloi et al., 2015](#)).
696 This is in line with the link between backward replay and selection of transitions
697 based on changes in value proposed by [Mattar and Daw \(2018\)](#). Replay is also more

698 likely to contain behaviorally significant locations, such as the current goal (Gupta
699 et al., 2010; Pfeiffer & Foster, 2013; Ólafsdóttir et al., 2015) and is biased by novelty
700 (Cheng & Frank, 2008; Foster & Wilson, 2006). It has also been observed that
701 optogenetic manipulation of dopaminergic input neurons, thought to signal PEs,
702 increase replay during subsequent sleep (McNamara et al., 2014).

703 Reactivating smaller sequences in a different order from what was experienced
704 can also be used to connect experiences in novel ways or strengthen weakly learned
705 relationships. Under such conditions replay can correspond more closely to sampling
706 from an internal model of the environment, rather than a veridical recapitulation
707 of past experiences (Sutton, 1991). Among the most early ideas about replay, the
708 Dyna architecture (Sutton, 1991) used an internal model to generate experiences
709 that were then used to train a model-free agent. Indeed, replay can be seen as a way
710 to blur the lines between model-free RL, such as the Q -learning method introduced
711 in Box 1, and model-based RL, during which the agent stores an explicit model of
712 the environment and can use it for planning (van Seijen & Sutton, 2015; Russek et
713 al., 2017; Momennejad et al., 2017).

714 Another dimension of experiential reorganization is to replay experiences that
715 happened long ago, or even refer to imagined experiences in the future. During be-
716 havior, replay events can switch between reflecting immediately preceding, upcoming
717 or more remote episodes, depending on the behavioral state of the animal at the time
718 of replay (Pfeiffer & Foster, 2013; Ólafsdóttir et al., 2017). Moreover, even single
719 replay events can depict more than one trajectory, such as the next one and the
720 path the animal will take after reaching the goal location (Pfeiffer & Foster, 2013),
721 as if representing a multi-step planning process (Foster, 2017; Miller & Venditto,
722 2021). Another reason why experiences from a more distant past might be replayed
723 is simply that the agent is using a period during which it does not have to engage
724 with the environment to optimize memory. This is particularly apparent for replay
725 during sleep, when the brain has idle time to process experiences while not being
726 actively engaged with any task. Sleep replay has frequently been observed in ani-
727 mals and humans, and been linked in particular to memory consolidation. Following
728 sleep, memory interference is reduced (Baran et al., 2010; McDevitt et al., 2015) and
729 memory integration or differentiation has been found in fMRI patterns after a delay
730 period with sleep (Favila et al., 2016; Tompary & Davachi, 2017).

731 A final aspect of re-inventing the past relates to re-considering the usefulness
732 of past experiences in light of one’s knowledge about a goal. The RL framework
733 presented thus far is aimed at the pursuit of a single goal (e.g., the single reward
734 location in our grid world, see Figure 1). However, in many real-world applications,
735 such as the movements of a robotic arm that needs to pick and place objects, an
736 RL strategy incorporating multiple goals would be far more beneficial. Consider
737 again the grid world example in Figure 1, but this time the agent can only move
738 a finite number of steps. Since there is only one goal state that returns a reward,
739 most of the transitions do not land in the goal state and therefore receive no reward.
740 In such a sparse binary reward situation, where success only results from those
741 sequences of transitions ending in the goal state, most sequences of transitions end
742 in uninformative failures, often related to early termination without reward (“giving
743 up”). For instance, when an agent gives up because a goal was not found after a
744 particular amount of time, it can not know how close it was to the goal. Humans,
745 however, can learn from failure as well as success. Inspired by this idea, an ML
746 technique known as hindsight experience replay (Andrychowicz et al., 2017) is used
747 to relabel the unsuccessful transitions by simply changing the goal state, such that
748 the transitions would now be considered as successful under the new goal, thus
749 contributing to the agent’s learning. To the best of our knowledge, no directly
750 equivalent observation has been made in the brain.

751 2.4 Planning for a better future

752 So far, we have mainly focused on the various ways in which replay serves learning
753 and memory. We have shown, for instance, how replay can be used to speed up learn-
754 ing Q -values, which can be used for behavior that maximizes reward (Lin, 1991, see
755 Figure 1). Yet, computational, psychological and neuroscientific research has pointed
756 out the importance of another mechanism that is crucial for goal-directed behavior:
757 planning. A core aspect of this process is the prospective, i.e., future-oriented, eval-
758 uation during which an agent deliberates which of the available sequences of actions
759 and states leads to the best among several potential outcomes. Notably, planning
760 requires a mental model, or cognitive map (Tolman, 1938, 1948), of the environment,
761 that describes the agent’s knowledge about the transition structure of events (e.g.,
762 paths between locations in a maze or contingencies between non-spatial task states)

763 including the outcomes at each potential location. Knowledge about the causal
 764 structure of the environment allows an agent to predict and compare the outcomes
 765 of sequences of states and actions and to choose the one that yields most reward.
 766 Figure 2 provides an illustration of how planning differs from the other two aspects
 767 of cognition we considered so far, acting and learning. While the distinction between
 768 learning and planning is useful, we will see below how agents can learn from planning
 769 processes.

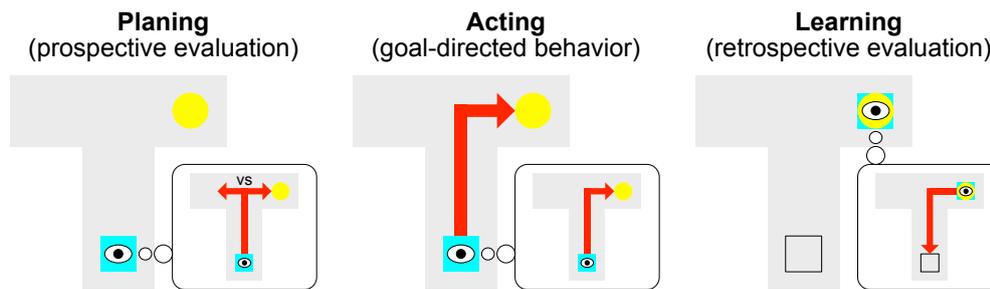


Figure 2: **Illustration of replay content.** The agent is represented by the blue square. The goal location is indicated by the yellow circle. Red arrows indicate behavioral trajectories of the agent in the T-maze environment or internally generated trajectories during replay. During planning (left 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/>).

770 Within the RL framework, the difference between acting based on learned (cached)
 771 values versus acting based on an internal planning process is embodied by the dis-
 772 tinction between model-free and model-based systems (Sutton & Barto, 2018; Daw
 773 et al., 2005). In model-based RL, the agent uses experience to learn a model of the
 774 environment that is described by a function that relates the current state s_t and
 775 action a_t to the next state s_{t+1} and the reward r_t (see Box 1). The agent can use
 776 this model at decision time or during offline periods to simulate experience (Sutton,
 777 1990), i.e., individual (s_t, a_t, r_t, s_{t+1}) -tuples, or longer episodes spanning several
 778 transitions. This allows the agent to update its learned values or to determine which
 779 action would be best to execute next, considering the rewards obtained and how
 780 the environment would change if a particular action was taken. This deliberation

781 process has two advantages. First, planning allows the agent to remain in the safety
782 of mental imagination and avoid the risk of suffering from potentially harmful conse-
783 quences. Second, planning can be used to decide between never-experienced, purely
784 hypothetical courses of action (Liu et al., 2019), a feat which would not be possible
785 with purely experience-based replay.

786 Despite these differences between planning and learning, much work in RL has
787 emphasized their similarities (Sutton, 1990, 1991; Sutton et al., 2012; van Seijen &
788 Sutton, 2015), and points to a function of planning that goes beyond deliberation.
789 As we have seen in our discussion of Lin, 1992, the same learning mechanisms can
790 be applied to real or simulated experience. Planning can thus not only be used to
791 determine immediate behavior, but also to shape value functions, a process referred
792 to as background planning (as opposed to the decision-time planning and deliberation
793 considered above, see Pezzulo et al., 2019). A case in point regarding the similarities
794 between planning and learning is illustrated by the Dyna architecture (Sutton, 1990,
795 1991). Just as any model-free RL agent, a Dyna agent selects actions according to
796 learned Q -values, and uses experiences to update these Q -values. But it also uses
797 experiences to observe rewards and subsequent states that follow each state-action
798 pair, to update an internal model of the environment. Importantly, in Dyna the
799 model is then used to train the model-free agent by producing simulated episodes
800 that will be used to update the agent's Q -values, just like real experiences. Two
801 implications from this perspective should be noted. First, replay might be used
802 not only to update value functions, but also to update a model of the environment
803 (see also, Penagos et al., 2017; Pezzulo et al., 2019, 2017; Momennejad, 2020).
804 Second, such a mechanism also suggests that replay elicits reward prediction error
805 signals, which could explain why striatal replay has been observed before (Lansink
806 et al., 2008, 2009; Pennartz, 2004). Notably, reward prediction errors might be
807 accompanied by state-prediction errors, which in the brain might both be conveyed
808 by dopaminergic signals (see e.g., Sharpe et al., 2017; Gardner et al., 2018).

809 It should be noted that an interaction between deliberation and learning may
810 arise. Entering a novel environment or changes in a familiar environment should
811 increase deliberation because a new model has to be learned or an existing model
812 has to be updated. But if the contingencies in the environment remain unchanged,
813 the need for the agent to update its model and deliberate its choices diminishes

814 with learning. At the same time, decisions become more habitual and less deliber-
815 ate (Dolan & Dayan, 2013). Thus, replayed trajectories in the hippocampus that
816 evaluate all potential trajectories might be only predictive of behavior during ear-
817 lier phases of learning (Singer et al., 2013) or vanish from the hippocampus (e.g.,
818 Wimmer & Büchel, 2019) when behavior becomes stereotyped. Findings by Papale
819 et al. (2016) demonstrate an inverse relationship between SWRs at reward sites and
820 deliberation at choice points.

821 While the potential benefits of replay for planning have been recognized early
822 on in RL (Sutton, 1990), serious consideration of this aspect only appeared in neu-
823 roscience later, where studies demonstrated replay events in the awake state, often
824 during short pauses from active behavior (e.g., Kudrimoti et al., 1999; Foster & Wil-
825 son, 2006; Csicsvari et al., 2007; Diba & Buzsáki, 2007; Kurth-Nelson et al., 2016;
826 Eldar et al., 2020). This allowed researchers to draw closer correspondence between
827 the replayed and the behavioral trajectories, and has resulted in a wealth of findings
828 supporting the idea that replay supports model-based planning in animals as well
829 as humans (for reviews, see e.g., Yu & Frank, 2015; Pezzulo et al., 2019; Wang et
830 al., 2020; Tambini & Davachi, 2019; Carr et al., 2011; Ólafsdóttir et al., 2018). Dis-
831 ruption of awake hippocampal SWRs during a spatial alternation task specifically
832 impaired the ability to decide between two trajectories to alternating goal locations,
833 whereas place field representations, reactivation during rest, and other navigation
834 behavior remained intact (Jadhav et al., 2012). Replay events in the awake state
835 predominantly co-occur with SWRs during short pauses from ongoing exploratory
836 behavior. Replay trajectories during awake SWRs often start at the current location
837 of the animal (Diba & Buzsáki, 2007; Davidson et al., 2009; Karlsson & Frank, 2009;
838 Pfeiffer & Foster, 2013; Singer et al., 2013; Ambrose et al., 2016) and end at the goal
839 location (Dupret et al., 2010; Pfeiffer & Foster, 2013).

840 A behavioral correlate of deliberation was already described in the 1930s in ro-
841 dents (Tolman, 1926; Muenzinger & Fletcher, 1936), who tend to pause at a decision
842 point to look back and forth between possible paths, a behavior called vicarious trial
843 and error (VTE) (for review, see Redish, 2016). Later studies found that during
844 VTE events, hippocampal place cells associated with theta sequences sweep ahead
845 from the animal’s current location (Johnson & Redish, 2007; Wikenheiser & Redish,
846 2015b; Amemiya & Redish, 2016; Papale et al., 2016). It was also found that during

847 VTE-like behavior, place cell activity influenced the formation of place fields thought
848 to stabilize the cognitive map (Monaco et al., 2014). Note that VTE-like replay is
849 often accompanied by theta sequences, which differ from SWR in their neurophysiol-
850 ogy. Nonetheless, both can be described as sequential activation of hippocampal cell
851 populations, a simplifying assumption that is helpful from a computational perspec-
852 tive (Foster, 2017; Pezzulo et al., 2019). Recently, theta sequences have been shown
853 to quickly cycle between possible future trajectories (Kay et al., 2020), and increases
854 in theta power in the MTL have been observed in humans in a spatial planning task
855 (Kaplan et al., 2020). In human fMRI, blood-oxygen-level dependent (BOLD) ac-
856 tivity in the hippocampus has been shown to increase with deliberation time when
857 deciding between two food items with similar value (Bakkour et al., 2019) and hip-
858 pocampal activity patterns reflect routes to navigational goal locations (Brown et
859 al., 2016). Another study has found that when humans re-learn outcomes associated
860 with choices at lower levels of a decision tree, the extent to which higher levels of
861 the decision tree are reactivated during rest correlates with how much their decisions
862 change, to reach the new downstream reward states (Momennejad et al., 2018).

863 It should be noted that although awake replay during deliberation of future
864 choices is often related to improved task performance, the replayed trajectories do
865 not necessarily reveal the behavioral trajectory the animal would subsequently take
866 (Johnson & Redish, 2007; Singer et al., 2013). In the study by Singer et al. (2013),
867 hippocampal replay during SWRs that preceded correct choices reflected trajectories
868 for the correct *and* incorrect option in a two-alternative W-maze. Once correct per-
869 formance became stable (at 85% correct), replayed trajectories shifted to represent
870 the correct future choice more frequently than the incorrect one (Singer et al., 2013).
871 One interpretation of these findings can be that the hippocampus uses replay to eval-
872 uate all potential trajectories and the behaviorally relevant trajectory is instantiated
873 in a different brain region.

874 One additional complicating factor regarding the relation between replay and
875 subsequent behavior concerns the task setting and motivational state of the animal
876 (e.g. Carey et al., 2019; Wu et al., 2017). As discussed above, replayed place
877 cell sequences can represent the trajectory into a shock zone which is subsequently
878 avoided (Wu et al., 2017). This might have the purpose to learn strongly from
879 and not forget about significant outcomes, and thus in this circumstance replay

880 is related to avoiding rather than initiating trajectories. In line with this idea, a
881 growing literature on computational psychiatry considers that replay could underlie
882 symptoms like avoidance and rumination that characterize psychiatric disorders like
883 anxiety (Gagne et al., 2018; Heller & Bagot, 2020; Mobbs et al., 2020).

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

903 Finally, we want to highlight that most of the above interpretations of replay in
904 the neuroscience literature rest on the idea that replay is shaped by previously experi-
905 enced behavioral trajectories. Yet, this account was challenged by a series of findings
906 reporting apparent “preplay” of place cell sequences before the environment was ever
907 experienced (Dragoi & Tonegawa, 2011, 2013). These findings highlight that ap-
908 parent sequentiality can also reflect hippocampal cell assemblies that are connected
909 in a way that constrains sequential firing, even prior to experience of a new maze.
910 Nevertheless, other findings indicate that previous experience is required for such
911 spontaneous sequential activation to occur (Silva et al., 2015). Recent computa-
912 tional work has suggested a link between preplay and efficient learning, arguing that

913 preexisting internal sequences could be used as a dynamical reservoir (Leibold, 2020).
914 The extent to which the hippocampus is able to seemingly preplay novel experiences
915 could depend on the similarity between pre-existing hippocampal representations
916 and new memories about to be formed (Eichenbaum, 2015). Methodologically, this
917 nevertheless highlights the necessity to compare pre- versus post-task replay (e.g.
918 Buhry et al., 2011), as shown by recent research that observed pre-vs.-post changes
919 in replay can indeed be explained by cell activation and firing rate correlations during
920 experience (Farooq et al., 2019).

921 2.5 Inference and generalization

922 Apart from its role in learning and planning, recent developments in ML and neuro-
923 science research suggest that replay also contributes to inference and generalization
924 (for previous reviews and perspectives, see Kumaran, 2012; Kumaran & McClelland,
925 2012; Cazé et al., 2018; Herszage & Censor, 2018; Lewis et al., 2018; Momenne-
926 jad, 2020). A recent theme in this domain has been to build artificial agents that
927 learn generative models from experience, which can then be used to infer new con-
928 nections based on latent structural rules (Evans & Burgess, 2019), infer the correct
929 context when given new data (Stoianov et al., 2020) and generalize information to
930 new tasks to mitigate performance losses (Shin et al., 2017). In the model proposed
931 by Stoianov et al. (2020), for instance, trajectories through a maze are used to learn
932 a generative model, which can produce new trajectories consistent with the current
933 maze structure during offline periods. As new mazes are learned, novel trajectories
934 continue to be generated offline, but from all the mazes that have been experienced,
935 preventing information about any one maze from being lost (similar to our consid-
936 erations about forgetting above). The hierarchical structure of the model results in
937 trajectories being clustered into distinct maze contexts, which allows maze categories
938 to be inferred when presented with new data. Unlike replay used in other contexts,
939 the model by Stoianov et al. (2020) does not suggest that prioritized replay helps
940 to improve behavioral outcomes. Generative replay that was prioritized based on
941 how surprising observations were under the generative model increased the number
942 of reactivation events that contained important goal locations but did not further
943 improve inference performance (Stoianov et al., 2020).

944 The clustering of trajectories seen in the model above is related to a broader

945 theoretical view, which has emphasized that separate encoding of transition infor-
946 mation and sensory information during learning will allow knowledge about tran-
947 sitions to be reused across situations with structural similarities but new sensory
948 specifics (Behrens et al., 2018; Liu et al., 2019; Baram et al., 2020; Whittington et
949 al., 2020). Because replay provides a strong candidate mechanism for learning about
950 transition structure (Stoianov et al., 2020), replay of abstract (sensory-independent)
951 transition information could help to build representations of task structure that can
952 be generalized and used to guide behaviour in new sensory environments (Liu et al.,
953 2019) or combined with sensory observations to make inferences about the current
954 environment (Evans & Burgess, 2019; Stoianov et al., 2020).

955 Another major computational approach has focused on replay as a mechanism
956 to learn successor representations, a predictive representation that reflects the ex-
957 pected future visitation of states, given the current state (Dayan, 1993). Unlike
958 the one-step transition matrices that are known as models in model-based RL, the
959 successor matrix can reflect non-adjacent dependencies. This allows the agent to
960 understand relationships between a state and multiple successor states, knowledge
961 which can be used to solve inference problems, such as finding the shortest path
962 to a new reward location (Russek et al., 2017; Momennejad et al., 2017). Succes-
963 sor representations have similar properties to place fields, skewing in the opposite
964 direction of travel and over-representing goal locations, while the eigenvectors of
965 a successor matrix can partition the environment into clusters that help planning,
966 resembling entorhinal grid cells in some spatial contexts (Stachenfeld et al., 2017).
967 Hippocampal-entorhinal fMRI signals have also been shown to reflect relationships
968 between successive non-spatial objects organized in a graph (Schapiro et al., 2013;
969 Garvert et al., 2017), consistent with the SR (Stachenfeld et al., 2017). Critically,
970 replay of past experiences could be used to update the successor matrix during offline
971 periods (Russek et al., 2017). In recent work, Russek et al. (2017) proposed replay of
972 (s_t, a_t, s_{t+1}) -tuples that are prioritized by recency, in which rewards are not needed
973 to update the successor matrix. Using this approach, it was shown that learning SRs
974 with offline replay gives an agent unique benefits compared to agents without replay.
975 In particular, the fast updating of successor states through replay allowed the agent
976 to quickly infer policy updates needed to adapt to changes in the task environment,
977 like a new barrier, that affect the state transition structure.

978 A potential role of replay in generalization and inference has also been suggested
979 by neuroscientific studies. In a recent study from [Barron et al. \(2020\)](#), hippocampal
980 cells selective to cues and rewarding outcomes that have not been directly experienced
981 together, but whose relationship can be inferred based on sensory pre-conditioning
982 ([Brogden, 1939](#)), were found to be co-active during SWR events ([Barron et al., 2020](#)).
983 These cells also tended to be reactivated during SWR events in a specific order, with
984 reward selective cells reactivated prior to cue selective cells akin to backward replay.
985 Using MEG, researchers have discovered that visual stimuli are reactivated in a non-
986 experienced order that was based on prior learning of a rule about how items should
987 be reported ([Liu et al., 2019](#)). This indicates that sequential reactivation is able to
988 combine prior learning with new sensory inputs to produce behavior relevant to new
989 environments. This aligns with the theoretical proposal that decomposing informa-
990 tion into separate structural and sensory representations can be used to generalize
991 structural relationships to situations with different sensory specifics but similar un-
992 derlying structure ([Behrens et al., 2018](#); [Baram et al., 2020](#); [Whittington et al.,](#)
993 [2020](#)). Nonetheless, this theoretical view is not restricted to replay events and, as a
994 result, there is still limited work testing this idea in the context of replay. Finally,
995 recordings in the hippocampus and PFC have shown that hippocampal place cells are
996 reactivated with subsets of prefrontal cells that encode generalizable task elements
997 ([Yu et al., 2018](#)) and that at least some medial PFC neurons involved in replay
998 have generalized firing fields that cover multiple starting locations or multiple goal
999 locations within a maze ([Kaefer et al., 2020](#)). Together, these results suggest that
1000 replay could also contribute to generalization through coordinating the appropriate
1001 reactivation of PFC neurons.

1002 2.6 Representation learning

1003 The analogy between hippocampal replay and replay in RL models sheds light on
1004 another important aspect: states – the agent’s internal representation of its environ-
1005 ment. State representations are crucial for RL because they can be optimized for
1006 efficient learning and structure the task at hand as an MDP (see [Box 1](#)). In this last
1007 section, we highlight findings indicating that replay is not specific to spatial loca-
1008 tions and could instead involve such task-dependent state representations. We argue
1009 that replaying states has unique benefits as opposed to replaying only observations.

1010 Moreover, we speculate that replay might also have a role in *learning* the representa-
1011 tions that guide behavior. As such, replay could offer a window into the operations
1012 the brain performs to craft useful representations of the possible task states.

1013 The main idea of states is that they meaningfully delineate different situations
1014 in which a specific action has a particular value. In a game of chess, for instance,
1015 the states are the different configurations of the pieces on the board. Given a board
1016 configuration, one can learn how beneficial it is to move a particular piece. Mak-
1017 ing the same move in a different position would have very different consequences,
1018 and differentiating the value of possible moves by board configuration is therefore
1019 paramount. The color of the shirt of your opponent, in contrast, may be a salient
1020 visual observation, but it has no implications for the values of the moves, and it
1021 therefore need not be part of the state. This implies that task state representations
1022 can be shaped by attentional mechanisms that focus on task-relevant dimensions
1023 (e.g., [Niv et al., 2015](#); [Leong et al., 2017](#)). Besides filtering irrelevant information,
1024 good state representations have to overcome the incompleteness of sensory inputs,
1025 which by themselves sometime do not contain all information needed to determine
1026 which action will lead to which outcome. Outcomes, for instance, often depend on
1027 non-observable context, such as past events, but internal representations can be used
1028 to combine current experience with observations of the past to determine the current
1029 task state ([Wilson et al., 2014](#); [Schuck et al., 2018](#)). For example, knowledge about
1030 how your opponent has played in the past can lead you to select a different move in
1031 the same situation as in a previous match. Finally, state representations should be
1032 robust to noise, leverage similarity among the states to determine what information
1033 can be generalized, and identify the hidden causes that may underlie sensory obser-
1034 vations (e.g., [Gershman & Niv, 2010](#)). Different chess positions might for instance
1035 share structural similarities that can be exploited with the same strategy. Defining
1036 the states of a task is therefore crucial for the success of an agent. As [Dayan \(1993\)](#)
1037 has put it: “difficult problems can be rendered trivial if looked at in the correct way”
1038 (p. 613). Importantly, since the agent does not know from the start “how to look
1039 at” the problem in the right way, learning useful state representations constitutes a
1040 major challenge for intelligent agents ([Bengio et al., 2013](#); [Niv, 2019](#)).

1041 In traditional RL models, the input to the algorithm is often engineered by hand
1042 to reflect the true state of the environment. Replay in these models therefore always

1043 reflects states, because no differentiation between sensory observations and states is
1044 made. In DNNs, the model is equipped with a convolutional network that converts
1045 sensory observations into states, over which Q -learning operates (e.g., [Mnih et al.,](#)
1046 [2015](#)). Replay in DNNs involves a memory buffer of observations, which are con-
1047 verted into states inside the network during replay, allowing efficient learning. These
1048 approaches therefore suggest that replay either involves only states, or observations
1049 as well as states.

1050 In animals, several findings indicate that a large variety of representations, in-
1051 cluding non-spatial sensory as well as state-like representations, might be replayed.
1052 First, the firing of hippocampal “place” cells can reflect a number of non-spatial as-
1053 pects of the environment, if they are task-relevant, such as sounds ([Aronov et al.,](#)
1054 [2017](#)), time ([MacDonald et al., 2011](#)) or successor representations ([Stachenfeld et al.,](#)
1055 [2017](#); but see [O’Keefe & Krupic, 2021](#)). More directly, one fMRI study by [Schuck](#)
1056 [and Niv \(2019\)](#) has found that sequential hippocampal replay during post-task rest
1057 reflected the non-spatial states of a sequential decision-making task. Importantly,
1058 observed transitions between decoded replay events were best explained by replay of
1059 states that include non-observable task aspects, such as information from the previ-
1060 ous trial, rather than by replay of sensory features of the task stimuli alone. This
1061 study therefore provides direct evidence for the idea that replay involves state repre-
1062 sentations that are optimized for the operation of RL algorithms. In an MEG-study
1063 by [Liu et al. \(2019\)](#), human participants first learned an abstract rule how objects
1064 should be ordered in a sequence and later replayed a novel set of objects according
1065 to the learned rule rather than in order of experience. Replayed sequences consisted
1066 of factorized representations of sensory objects, the identity of the sequence they
1067 belonged to, as well as the position within that sequence, supporting the notion that
1068 replay is not limited to one kind of information. Moreover, [Jadhav et al. \(2012\)](#)
1069 showed that disruption of SWRs in a spatial alternation task impaired navigation
1070 when it required unobservable knowledge of the previous trial, thus hinting at the
1071 activation of state representations rather than observations during replay.

1072 Other evidence suggests that replay involves multiple representations which are
1073 reactivated in parallel. These representations reflect visual ([Ji & Wilson, 2006](#); [Wit-](#)
1074 [tkuhn & Schuck, 2021](#)), auditory ([Rothschild et al., 2016](#)) or grid-like ([Ólafsdóttir](#)
1075 [et al., 2016, 2017](#); [O’Neill et al., 2017](#)) information and have been observed during

1076 interactions between hippocampus and OFC (for reviews, see [Wikenheiser & Re-](#)
1077 [dish, 2015a](#); [Wikenheiser & Schoenbaum, 2016](#)), which might relate to the agent’s
1078 task state representations ([Kaplan et al., 2017](#)). Disruption of the medial PFC par-
1079 ticularly attenuated components of hippocampal theta sequences representing the
1080 current location of the animal ([Schmidt et al., 2019](#)) and suppression of hippocam-
1081 pal input to the OFC impaired the integration of task state structure ([Wikenheiser](#)
1082 [et al., 2017](#)). Conversely, disruption of SWRs during sleep impaired the integrity of
1083 hippocampal maps but they re-emerged following re-learning ([Gridchyn et al., 2020](#))
1084 suggesting that relevant maps are stored in brain areas other than the hippocampus
1085 ([Niethard & Born, 2020](#)). It should also be noted that a number of investigations
1086 have shown replay events outside the hippocampus need not be coordinated with
1087 hippocampal activity ([O’Neill et al., 2017](#); [Kaefer et al., 2020](#); [Wittkuhn & Schuck,](#)
1088 [2021](#)). In sum, evidence is mounting that hippocampal “place cell” firing can reflect
1089 a variety of non-spatial but task-relevant aspects (e.g., [Aronov et al., 2017](#)), replay
1090 occurs in a wide variety of brain areas that encode different informational aspects,
1091 which at least partially reflect an animals’ understanding of what is task relevant
1092 and even that replay directly reflects partially observable task states ([Schuck & Niv,](#)
1093 [2019](#)).

1094 Replaying states and replaying different levels of representation might convey
1095 unique benefits. RL models benefit from transition information between states, but
1096 they could be affected adversely if transitions of task-irrelevant aspects affect the
1097 agents internal model. For example, representing states as specific locations in phys-
1098 ical space will result in a transition matrix that is different from a transition matrix
1099 of more abstract task states, but if spatial position is irrelevant to the task at hand,
1100 then transitions between locations are not needed for learning and planning. In other
1101 words: the transition matrix, and therefore the foundation of an agent’s ability to
1102 plan and simulate actions, fully depends on its internal state representations ([Guo](#)
1103 [et al., 2020](#)). In addition, observations are, under most circumstances, much more
1104 high dimensional than states, and storing them might require a much larger memory
1105 buffer than an approach in which states are stored (e.g., [Shin et al., 2017](#)). Thus,
1106 replay that has access to and operates on highly adaptive state representations seems
1107 computationally sensible.

1108 If replay indeed operates on the level of states, and learning good state repre-

1109 sentations is a major challenge for any agent, this begs another question: can replay
1110 also help to change the content of the representation itself? Some evidence sug-
1111 gests this might indeed be the case. [Schuck and Niv \(2019\)](#) observed that replay
1112 in the hippocampus during rest was related to better decodability of partially ob-
1113 servable state representations from the OFC during the task. Moreover, decoding
1114 of state representations in the OFC increased over time. Although this evidence
1115 is correlational, it hints at some relation between replay and state representation
1116 learning. How could replay help to form better state representations? One impor-
1117 tant instance concerns successor representations (SRs), which provide an efficient
1118 way to incorporate knowledge about the transitions between states into the state
1119 definition. Computational work by [Russek et al. \(2017\)](#) has shown that SRs are not
1120 only useful state representations, but they can also be learned and updated through
1121 replay. More generally, information about state transitions can give rise to further
1122 graph analytical insights which are known to provide a good basis for state represen-
1123 tations ([Mahadevan & Maggioni, 2007](#)). Sequential replay is a natural match as a
1124 mechanism to learn states encoding transitional information, and possibly might be
1125 involved in extracting graph properties from experienced transition structures, such
1126 as bottleneck states. This knowledge could then be used to form better states. Other
1127 work has highlighted that representations which predict latent embeddings of future
1128 observations are particularly useful ([Guo et al., 2020](#)). An evaluation of predictive-
1129 ness could therefore be an important contribution of replay to state representation
1130 learning.

1131 Some evidence suggests that the role of replay for state learning could go beyond
1132 the focus on information about transitions. SRs, for instance, can be extended to
1133 deal with partially observable task environments ([Vértes & Sahani, 2019](#)). Caselles-
1134 Dupré and colleagues proposed another interesting account that involves variational
1135 autoencoders (VAEs) ([Caselles-Dupré et al., 2018, 2019](#)). Building on earlier work
1136 that used generative models to circumvent the memory requirements of observa-
1137 tional replay ([Shin et al., 2017](#)), [Caselles-Dupré et al. \(2019\)](#) proposed storing latent
1138 representations rather than observations, and using past experiences in this form of
1139 continually training a VAE that acts as a state model. Importantly, only by replay-
1140 ing past episodes can the VAE learn to form a state representation that allows the
1141 agent to act efficiently across more than one environment.

1142 While these computational models represent intriguing proposals, to the best of
1143 our knowledge, not many empirical studies have shed light on whether state learning
1144 mechanisms provide a realistic account for biological replay. As we have shown,
1145 however, some of the above discussed neuroscientific results do suggest that replay
1146 is involved in the construction and maintenance of hidden state representations.

1147 **3 Conclusion and outlook**

1148 In this review, we have summarized the literature on replay in neuroscience and ML to
1149 showcase which computational benefits biological and artificial agents can gain from
1150 replaying previous experience. We have discussed five main computational benefits
1151 that, although overlapping, provide useful categories for thinking about what might
1152 motivate an agent to employ replay: faster learning and increased data efficiency, less
1153 forgetting, the reorganization of experience, planning and generalization. In addition
1154 we have argued that replayed content is much richer than a sequence of locations, and
1155 could reflect the agent's current state representation. State representations are often
1156 task- and context-dependent and are influenced by a range of factors, including the
1157 goal-relevant aspects of the agents observations, the transition structure of states,
1158 the location, number and value of goal locations and the motivational and metabolic
1159 state of the animal. We have argued that RL theory provides useful guidance to
1160 understand which form state representations might take in a given task, and which
1161 implications a particular state representation would have for an agent's behavior.
1162 Finally, we have discussed how replay might not only reflect but could help the agent
1163 to learn those states to begin with. While many questions in particular regarding the
1164 latter idea still remain, considering these factors will greatly help to determine what
1165 replayed representations represent and how replay updates decision-making policies
1166 that are used to control behavior.

1167 4 Declaration of competing interest

1168 The authors declare no competing interests.

1169 5 Acknowledgements

1170 This work was supported by an Independent Max Planck Research Group grant
1171 awarded to N.W.S by the Max Planck Society (M.TN.A.BILD0004), a Starting Grant
1172 awarded to N.W.S by the European Union (ERC-2019-StG REPLAY-852669), and
1173 a Humboldt Research Fellowship awarded to S.H.M. by the Alexander von Hum-
1174 boldt Foundation. We thank members of the Max Planck Research Group Neu-
1175 roCode for helpful discussions about the contents of this manuscript. We thank
1176 Anika Löwe for feedback on a previous version of this manuscript. L.W. is a pre-
1177 doctoral fellow of the International Max Planck Research School on Computational
1178 Methods in Psychiatry and Ageing Research (IMPRS COMP2PSYCH). The partic-
1179 ipating institutions are the Max Planck Institute for Human Development, Berlin,
1180 Germany, and University College London, London, UK. For more information, see
1181 <https://www.mps-ucl-centre.mpg.de/en/comp2psych>.

1182 Glossary

1183 **BOLD** blood-oxygen-level dependent.

1184 **CLS** complementary learning systems.

1185 **DNN** deep neural network.

1186 **DQN** Deep Q-Network.

1187 **EEG** electroencephalography.

1188 **fMRI** functional magnetic resonance imaging.

1189 **i.i.d.** independent and identically distributed.

1190 **MDP** Markov decision process.

1191 **MEG** magnetoencephalography.

1192 **ML** machine learning.

1193 **MTL** medial temporal lobe.

1194 **OFC** orbitofrontal cortex.

1195 **PE** prediction error.

1196 **PFC** prefrontal cortex.

1197 **RL** reinforcement learning.

1198 **SGD** stochastic gradient descent.

1199 **SR** successor representation.

1200 **STDP** spike timing dependent plasticity.

1201 **SWR** sharp wave-ripple.

1202 **TD** temporal difference.

1203 **TMR** targeted memory reactivation.

1204 **VAE** variational autoencoder.

1205 **VTE** vicarious trial and error.

1206 6 References

- 1207 Ambrose, R. E., Pfeiffer, B. E., & Foster, D. J. (2016). Reverse replay of hippocam-
1208 pal place cells is uniquely modulated by changing reward. *Neuron*, *91*(5), 1124–1136.
1209 doi: [10.1016/j.neuron.2016.07.047](https://doi.org/10.1016/j.neuron.2016.07.047)
- 1210 Amemiya, S., & Redish, A. D. (2016). Manipulating decisiveness in decision making: Effects
1211 of clonidine on hippocampal search strategies. *Journal of Neuroscience*, *36*(3), 814–827.
1212 doi: [10.1523/JNEUROSCI.2595-15.2016](https://doi.org/10.1523/JNEUROSCI.2595-15.2016)
- 1213 Andre, D., Friedman, N., & Parr, R. (1998). Generalized prioritized sweeping. In *Advances in neural*
1214 *information processing systems*. MIT Press.
- 1215 Andrychowicz, M., Wolski, F., Ray, A., Schneider, J., Fong, R., Welinder, P., ... Zaremba, W. (2017).
1216 Hindsight experience replay. In I. Guyon et al. (Eds.), *Advances in neural information processing*
1217 *systems 30* (pp. 5048–5058). Curran Associates, Inc.
- 1218 Antony, J. W., & Schapiro, A. C. (2019). Active and effective replay: systems consolidation reconsid-
1219 ered again. *Nature Reviews Neuroscience*, *20*(8), 506–507. doi: [10.1038/s41583-019-0191-8](https://doi.org/10.1038/s41583-019-0191-8)
- 1220 Aronov, D., Nevers, R., & Tank, D. W. (2017). Mapping of a non-spatial dimension by the
1221 hippocampal–entorhinal circuit. *Nature*, *543*(7647), 719–722. doi: [10.1038/nature21692](https://doi.org/10.1038/nature21692)
- 1222 Axmacher, N., Elger, C. E., & Fell, J. (2008). Ripples in the medial temporal lobe are relevant for
1223 human memory consolidation. *Brain*, *131*(7), 1806–1817. doi: [10.1093/brain/awn103](https://doi.org/10.1093/brain/awn103)
- 1224 Bakkour, A., Palombo, D. J., Zylberberg, A., Kang, Y. H., Reid, A., Verfaellie, M., ... Shohamy, D.
1225 (2019). The hippocampus supports deliberation during value-based decisions. *eLife*, *8*, e46080.
1226 doi: [10.7554/eLife.46080](https://doi.org/10.7554/eLife.46080)
- 1227 Baram, A. B., Muller, T. H., Nili, H., Garvert, M. M., & Behrens, T. E. J. (2020). Entorhinal and
1228 ventromedial prefrontal cortices abstract and generalize the structure of reinforcement learning
1229 problems. *Neuron*. doi: [10.1016/j.neuron.2020.11.024](https://doi.org/10.1016/j.neuron.2020.11.024)
- 1230 Baran, B., Wilson, J., & Spencer, R. M. C. (2010). REM-dependent repair of competitive memory
1231 suppression. *Experimental Brain Research*, *203*(2), 471–477. doi: [10.1007/s00221-010-2242-2](https://doi.org/10.1007/s00221-010-2242-2)
- 1232 Barron, H. C., Reeve, H. M., Koolschijn, R. S., Perstenko, P. V., Shpektor, A., Nili, H., ... Dupret,
1233 D. (2020). Neuronal computation underlying inferential reasoning in humans and mice. *Cell*,
1234 *183*(1), 228 - 243.e21. doi: [10.1016/j.cell.2020.08.035](https://doi.org/10.1016/j.cell.2020.08.035)
- 1235 Behrens, T. E., Muller, T. H., Whittington, J. C., Mark, S., Baram, A. B., Stachenfeld, K. L., &
1236 Kurth-Nelson, Z. (2018). What is a cognitive map? organizing knowledge for flexible behavior.
1237 *Neuron*, *100*(2), 490–509. doi: [10.1016/j.neuron.2018.10.002](https://doi.org/10.1016/j.neuron.2018.10.002)
- 1238 Bellman, R. (1957). *Dynamic programming*. Princeton University Press.
- 1239 Bellmund, J. L. S., Gärdenfors, P., Moser, E. I., & Doeller, C. F. (2018). Navigating cognition: Spatial
1240 codes for human thinking. *Science*, *362*(6415). doi: [10.1126/science.aat6766](https://doi.org/10.1126/science.aat6766)
- 1241 Bendor, D., & Wilson, M. A. (2012). Biasing the content of hippocampal replay during sleep. *Nature*
1242 *Neuroscience*, *15*(10), 1439–1444. doi: [10.1038/nn.3203](https://doi.org/10.1038/nn.3203)
- 1243 Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new per-
1244 spectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *35*(8), 1798–1828.
1245 doi: [10.1109/tpami.2013.50](https://doi.org/10.1109/tpami.2013.50)
- 1246 Bhattarai, B., Lee, J. W., & Jung, M. W. (2019). Distinct effects of reward and navigation history
1247 on hippocampal forward and reverse replays. *Proceedings of the National Academy of Sciences*,
1248 201912533. doi: [10.1073/pnas.1912533117](https://doi.org/10.1073/pnas.1912533117)

- 1249 Bird, C. M. (2020). How do we remember events? *Current Opinion in Behavioral Sciences*, *32*,
1250 120–125. doi: [10.1016/j.cobeha.2020.01.020](https://doi.org/10.1016/j.cobeha.2020.01.020)
- 1251 Bliss, T. V. P., & Collingridge, G. L. (1993). A synaptic model of memory: long-term potentiation in
1252 the hippocampus. *Nature*, *361*(6407), 31–39. doi: [10.1038/361031a0](https://doi.org/10.1038/361031a0)
- 1253 Bornstein, A. M., Khaw, M. W., Shohamy, D., & Daw, N. D. (2017). Reminders of past choices bias de-
1254 cisions for reward in humans. *Nature Communications*, *8*(1), 15958. doi: [10.1038/ncomms15958](https://doi.org/10.1038/ncomms15958)
- 1255 Bornstein, A. M., & Norman, K. A. (2017). Reinstated episodic context guides sampling-based decisions
1256 for reward. *Nature Neuroscience*, *20*(7), 997–1003. doi: [10.1038/nn.4573](https://doi.org/10.1038/nn.4573)
- 1257 Bottini, R., & Doeller, C. F. (2020). Knowledge across reference frames: Cogni-
1258 tive maps and image spaces. *Trends in Cognitive Sciences*, *24*(8), 606 - 619.
1259 doi: <https://doi.org/10.1016/j.tics.2020.05.008>
- 1260 Botvinick, M., Ritter, S., Wang, J. X., Kurth-Nelson, Z., Blundell, C., & Hassabis, D. (2019).
1261 Reinforcement learning, fast and slow. *Trends in Cognitive Sciences*, *23*(5), 408 - 422.
1262 doi: <https://doi.org/10.1016/j.tics.2019.02.006>
- 1263 Botvinick, M., Wang, J. X., Dabney, W., Miller, K. J., & Kurth-Nelson, Z. (2020). Deep reinforcement
1264 learning and its neuroscientific implications. *Neuron*. doi: [10.1016/j.neuron.2020.06.014](https://doi.org/10.1016/j.neuron.2020.06.014)
- 1265 Bragin, A., Engel, J., Wilson, C. L., Fried, I., & Buzsáki, G. (1999). High-frequency oscillations in
1266 human brain. *Hippocampus*, *9*(2), 137–142. doi: [10.1002/\(sici\)1098-1063\(1999\)9:2<137::aid-
1267 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)
- 1268 Brogden, W. J. (1939). Sensory pre-conditioning. *Journal of Experimental Psychology*, *25*(4), 323.
1269 doi: [10.1037/h0058944](https://doi.org/10.1037/h0058944)
- 1270 Brown, T. I., Carr, V. A., LaRocque, K. F., Favila, S. E., Gordon, A. M., Bowles, B., . . . Wagner, A. D.
1271 (2016). Prospective representation of navigational goals in the human hippocampus. *Science*,
1272 *352*(6291), 1323–1326. doi: [10.1126/science.aaf0784](https://doi.org/10.1126/science.aaf0784)
- 1273 Brunec, I. K., Moscovitch, M., & Barense, M. D. (2018). Boundaries shape cognitive representations of
1274 spaces and events. *Trends in Cognitive Sciences*, *22*(7), 637–650. doi: [10.1016/j.tics.2018.03.013](https://doi.org/10.1016/j.tics.2018.03.013)
- 1275 Buckner, R. L. (2010). The role of the hippocampus in prediction and imagination. *Annual Review of*
1276 *Psychology*, *61*(1), 27-48. doi: [10.1146/annurev.psych.60.110707.163508](https://doi.org/10.1146/annurev.psych.60.110707.163508)
- 1277 Buhry, L., Azizi, A. H., & Cheng, S. (2011). Reactivation, replay, and preplay: How it might all fit
1278 together. *Neural Plasticity*, 1–11. doi: [10.1155/2011/203462](https://doi.org/10.1155/2011/203462)
- 1279 Bush, D., Barry, C., Manson, D., & Burgess, N. (2015). Using grid cells for navigation. *Neuron*, *87*(3),
1280 507 - 520. doi: <https://doi.org/10.1016/j.neuron.2015.07.006>
- 1281 Buzsáki, G. (1989). Two-stage model of memory trace formation: A role for “noisy” brain states.
1282 *Neuroscience*, *31*(3), 551–570. doi: [10.1016/0306-4522\(89\)90423-5](https://doi.org/10.1016/0306-4522(89)90423-5)
- 1283 Buzsáki, G. (2015). Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and
1284 planning. *Hippocampus*, *25*(10), 1073–1188. doi: [10.1002/hipo.22488](https://doi.org/10.1002/hipo.22488)
- 1285 Cairney, S. A., Lindsay, S., Sobczak, J. M., Paller, K. A., & Gaskell, M. G. (2016). The benefits of
1286 targeted memory reactivation for consolidation in sleep are contingent on memory accuracy and
1287 direct cue-memory associations. *Sleep*, *39*(5), 1139-1150. doi: [10.5665/sleep.5772](https://doi.org/10.5665/sleep.5772)
- 1288 Carey, A. A., Tanaka, Y., & van der Meer, M. A. A. (2019). Reward revaluation biases hippocam-
1289 pal replay content away from the preferred outcome. *Nature Neuroscience*, *22*(9), 1450–1459.
1290 doi: [10.1038/s41593-019-0464-6](https://doi.org/10.1038/s41593-019-0464-6)
- 1291 Carr, M. F., Jadhav, S. P., & Frank, L. M. (2011). Hippocampal replay in the awake state: a
1292 potential substrate for memory consolidation and retrieval. *Nature Neuroscience*, *14*(2), 147–

1293 153. doi: [10.1038/nm.2732](https://doi.org/10.1038/nm.2732)

1294 Caselles-Dupré, H., Garcia-Ortiz, M., & Filliat, D. (2018). Continual state representation learning for
1295 reinforcement learning using generative replay. *arXiv e-prints*, arXiv:1810.03880.

1296 Caselles-Dupré, H., Garcia-Ortiz, M., & Filliat, D. (2019). S-TRIGGER: Continual State Representa-
1297 tion Learning via Self-Triggered Generative Replay. *arXiv e-prints*, arXiv:1902.09434.

1298 Cazé, R., Khamassi, M., Aubin, L., & Girard, B. (2018). Hippocampal replays under the
1299 scrutiny of reinforcement learning models. *Journal of Neurophysiology*, *120*(6), 2877–2896.
1300 doi: [10.1152/jn.00145.2018](https://doi.org/10.1152/jn.00145.2018)

1301 Cheng, S., & Frank, L. M. (2008). New experiences enhance coordinated neural activity in the
1302 hippocampus. *Neuron*, *57*(2), 303 - 313. doi: [10.1016/j.neuron.2007.11.035](https://doi.org/10.1016/j.neuron.2007.11.035)

1303 Clewett, D., DuBrow, S., & Davachi, L. (2019). Transcending time in the brain: How event memories
1304 are constructed from experience. *Hippocampus*, *29*(3), 162–183. doi: [10.1002/hipo.23074](https://doi.org/10.1002/hipo.23074)

1305 Cohen, N. J., & Eichenbaum, H. (1993). *Memory, amnesia, and the hippocampal system*. MIT Press.

1306 Constantinescu, A. O., O'Reilly, J. X., & Behrens, T. E. J. (2016). Organizing conceptual knowledge
1307 in humans with a gridlike code. *Science*, *352*(6292), 1464–1468. doi: [10.1126/science.aaf0941](https://doi.org/10.1126/science.aaf0941)

1308 Csicsvari, J., O'Neill, J., Allen, K., & Senior, T. (2007). Place-selective firing contributes to the
1309 reverse-order reactivation of CA1 pyramidal cells during sharp waves in open-field exploration.
1310 *European Journal of Neuroscience*, *26*(3), 704–716. doi: [10.1111/j.1460-9568.2007.05684.x](https://doi.org/10.1111/j.1460-9568.2007.05684.x)

1311 Davidson, T. J., Kloosterman, F., & Wilson, M. A. (2009). Hippocampal replay of extended experience.
1312 *Neuron*, *63*(4), 497–507. doi: [10.1016/j.neuron.2009.07.027](https://doi.org/10.1016/j.neuron.2009.07.027)

1313 Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and
1314 dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, *8*(12), 1704–1711.
1315 doi: [10.1038/nn1560](https://doi.org/10.1038/nn1560)

1316 Dayan, P. (1993). Improving generalization for temporal difference learning: The successor represen-
1317 tation. *Neural Computation*, *5*(4), 613–624. doi: [10.1162/neco.1993.5.4.613](https://doi.org/10.1162/neco.1993.5.4.613)

1318 Deng, X., Chen, S., Sosa, M., Karlsson, M. P., Wei, X.-X., & Frank, L. M. (2020). A variable clock
1319 underlies internally generated hippocampal sequences. *bioRxiv*. doi: [10.1101/2020.04.10.035980](https://doi.org/10.1101/2020.04.10.035980)

1320 Denovellis, E. L., Gillespie, A. K., Coulter, M. E., Sosa, M., Chung, J. E., Eden, U. T., &
1321 Frank, L. M. (2020). Hippocampal replay of experience at real-world speeds. *bioRxiv*.
1322 doi: [10.1101/2020.10.20.347708](https://doi.org/10.1101/2020.10.20.347708)

1323 Deuker, L., Olligs, J., Fell, J., Kranz, T. A., Mormann, F., Montag, C., ... Axmacher, N. (2013).
1324 Memory consolidation by replay of stimulus-specific neural activity. *Journal of Neuroscience*,
1325 *33*(49), 19373–19383. doi: [10.1523/jneurosci.0414-13.2013](https://doi.org/10.1523/jneurosci.0414-13.2013)

1326 Diba, K., & Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples.
1327 *Nature Neuroscience*, *10*(10), 1241–1242. doi: [10.1038/nn1961](https://doi.org/10.1038/nn1961)

1328 Diekelmann, S., & Born, J. (2010). The memory function of sleep. *Nature Reviews Neuroscience*,
1329 *11*(2), 114–126. doi: [10.1038/nrn2762](https://doi.org/10.1038/nrn2762)

1330 Diekelmann, S., Born, J., & Wagner, U. (2010). Sleep enhances false memories depend-
1331 ing on general memory performance. *Behavioural Brain Research*, *208*(2), 425 - 429.
1332 doi: <https://doi.org/10.1016/j.bbr.2009.12.021>

1333 Dolan, R. J., & Dayan, P. (2013). Goals and habits in the brain. *Neuron*, *80*(2), 312 - 325.
1334 doi: [10.1016/j.neuron.2013.09.007](https://doi.org/10.1016/j.neuron.2013.09.007)

1335 Doya, K. (1999). What are the computations of the cerebellum, the basal ganglia and the cerebral
1336 cortex? *Neural Networks*, *12*(7), 961-974. doi: [10.1016/S0893-6080\(99\)00046-5](https://doi.org/10.1016/S0893-6080(99)00046-5)

- 1337 Dragoi, G., & Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular
1338 assemblies. *Nature*, *469*(7330), 397–401. doi: [10.1038/nature09633](https://doi.org/10.1038/nature09633)
- 1339 Dragoi, G., & Tonegawa, S. (2013). Distinct preplay of multiple novel spatial experi-
1340 ences in the rat. *Proceedings of the National Academy of Sciences*, *110*(22), 9100–9105.
1341 doi: [10.1073/pnas.1306031110](https://doi.org/10.1073/pnas.1306031110)
- 1342 Drosopoulos, S., Schulze, C., Fischer, S., & Born, J. (2007). Sleep’s function in the spontaneous
1343 recovery and consolidation of memories. *Journal of Experimental Psychology: General*, *136*(2),
1344 169. doi: [10.1037/0096-3445.136.2.169](https://doi.org/10.1037/0096-3445.136.2.169)
- 1345 DuBrow, S., Rouhani, N., Niv, Y., & Norman, K. A. (2017). Does mental context drift or shift?
1346 *Current Opinion in Behavioral Sciences*, *17*, 141–146. doi: [10.1016/j.cobeha.2017.08.003](https://doi.org/10.1016/j.cobeha.2017.08.003)
- 1347 Duncan, K. D., & Shohamy, D. (2016). Memory states influence value-based decisions. *Journal of*
1348 *Experimental Psychology: General*, *145*(11), 1420.
- 1349 Dupret, D., O’Neill, J., Pleydell-Bouverie, B., & Csicsvari, J. (2010). The reorganization and reactiva-
1350 tion of hippocampal maps predict spatial memory performance. *Nature Neuroscience*, *13*(8),
1351 995–1002. doi: [10.1038/nn.2599](https://doi.org/10.1038/nn.2599)
- 1352 Ego-Stengel, V., & Wilson, M. A. (2010). Disruption of ripple-associated hippocampal activity during
1353 rest impairs spatial learning in the rat. *Hippocampus*, *20*(1), 1–10. doi: [10.1002/hipo.20707](https://doi.org/10.1002/hipo.20707)
- 1354 Eichenbaum, H. (2015). Does the hippocampus preplay memories? *Nature Neuroscience*, *18*(12),
1355 1701–1702. doi: [10.1038/nn.4180](https://doi.org/10.1038/nn.4180)
- 1356 Ekstrom, A. D., Kahana, M. J., Caplan, J. B., Fields, T. A., Isham, E. A., Newman, E. L., & Fried,
1357 I. (2003). Cellular networks underlying human spatial navigation. *Nature*, *425*(6954), 184–188.
1358 doi: [10.1038/nature01964](https://doi.org/10.1038/nature01964)
- 1359 Eldar, E., Lièvre, G., Dayan, P., & Dolan, R. J. (2020). The roles of online and offline replay in
1360 planning. *eLife*, *9*. doi: [10.7554/elife.56911](https://doi.org/10.7554/elife.56911)
- 1361 Epstein, R. A., Patai, E. Z., Julian, J. B., & Spiers, H. J. (2017). The cognitive map in humans: spatial
1362 navigation and beyond. *Nature Neuroscience*, *20*(11), 1504–1513. doi: [10.1038/nn.4656](https://doi.org/10.1038/nn.4656)
- 1363 Erdem, U. M., & Hasselmo, M. (2012). A goal-directed spatial navigation model using forward
1364 trajectory planning based on grid cells. *European Journal of Neuroscience*, *35*(6), 916–931.
1365 doi: [10.1111/j.1460-9568.2012.08015.x](https://doi.org/10.1111/j.1460-9568.2012.08015.x)
- 1366 Eschenko, O., Ramadan, W., Mölle, M., Born, J., & Sara, S. J. (2008). Sustained increase in hip-
1367 pocampal sharp-wave ripple activity during slow-wave sleep after learning. *Learning & Memory*,
1368 *15*(4), 222–228. doi: [10.1101/lm.726008](https://doi.org/10.1101/lm.726008)
- 1369 Euston, D. R., Tatsuno, M., & McNaughton, B. L. (2007). Fast-forward playback of re-
1370 cent memory sequences in prefrontal cortex during sleep. *Science*, *318*(5853), 1147–1150.
1371 doi: [10.1126/science.1148979](https://doi.org/10.1126/science.1148979)
- 1372 Evans, T., & Burgess, N. (2019). Coordinated hippocampal-entorhinal replay as structural inference.
1373 *Advances in Neural Information Processing Systems*, 1729–1741.
- 1374 Farooq, U., Sibille, J., Liu, K., & Dragoi, G. (2019). Strengthened temporal coordination within
1375 pre-existing sequential cell assemblies supports trajectory replay. *Neuron*, *103*(4), 719–733.e7.
1376 doi: [10.1016/j.neuron.2019.05.040](https://doi.org/10.1016/j.neuron.2019.05.040)
- 1377 Favila, S. E., Chanals, A. J. H., & Kuhl, B. A. (2016). Experience-dependent hippocampal pattern
1378 differentiation prevents interference during subsequent learning. *Nature Communications*, *7*(1),
1379 11066. doi: [10.1038/ncomms11066](https://doi.org/10.1038/ncomms11066)
- 1380 Favila, S. E., Lee, H., & Kuhl, B. A. (2020). Transforming the concept of memory reactivation. *Trends*

1381 *in Neurosciences*, 43(12), 939–950. doi: [10.1016/j.tins.2020.09.006](https://doi.org/10.1016/j.tins.2020.09.006)

1382 Fedus, W., Ramachandran, P., Agarwal, R., Bengio, Y., Larochelle, H., Rowland, M., & Dabney, W.
1383 (2020). Revisiting Fundamentals of Experience Replay. *arXiv e-prints*, arXiv:2007.06700.

1384 Flesch, T., Balaguer, J., Dekker, R., Nili, H., & Summerfield, C. (2018). Comparing continual task
1385 learning in minds and machines. *Proceedings of the National Academy of Sciences*, 115(44),
1386 E10313–E10322. doi: [10.1073/pnas.1800755115](https://doi.org/10.1073/pnas.1800755115)

1387 Foster, D. J. (2017). Replay comes of age. *Annual Review of Neuroscience*, 40(1), 581–602.
1388 doi: [10.1146/annurev-neuro-072116-031538](https://doi.org/10.1146/annurev-neuro-072116-031538)

1389 Foster, D. J., & Knierim, J. J. (2012). Sequence learning and the role of the hippocampus in rodent
1390 navigation. *Current Opinion in Neurobiology*, 22(2), 294–300. doi: [10.1016/j.conb.2011.12.005](https://doi.org/10.1016/j.conb.2011.12.005)

1391 Foster, D. J., & Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place
1392 cells during the awake state. *Nature*, 440(7084), 680–683. doi: [10.1038/nature04587](https://doi.org/10.1038/nature04587)

1393 French, R. (1999). Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*,
1394 3(4), 128–135. doi: [10.1016/s1364-6613\(99\)01294-2](https://doi.org/10.1016/s1364-6613(99)01294-2)

1395 Fyhn, M., Hafting, T., Treves, A., Moser, M.-B., & Moser, E. I. (2007). Hippocampal remapping and
1396 grid realignment in entorhinal cortex. *Nature*, 446(7132), 190–194. doi: [10.1038/nature05601](https://doi.org/10.1038/nature05601)

1397 Gagne, C., Dayan, P., & Bishop, S. J. (2018). When planning to survive goes wrong: predicting the
1398 future and replaying the past in anxiety and PTSD. *Current Opinion in Behavioral Sciences*,
1399 24, 89–95. doi: [10.1016/j.cobeha.2018.03.013](https://doi.org/10.1016/j.cobeha.2018.03.013)

1400 Gardner, M. P. H., Schoenbaum, G., & Gershman, S. J. (2018). Rethinking dopamine as generalized
1401 prediction error. *Proceedings of the Royal Society B: Biological Sciences*, 285(1891), 20181645.
1402 doi: [10.1098/rspb.2018.1645](https://doi.org/10.1098/rspb.2018.1645)

1403 Garvert, M. M., Dolan, R. J., & Behrens, T. E. (2017). A map of abstract relational knowledge in the
1404 human hippocampal–entorhinal cortex. *eLife*, 6. doi: [10.7554/elife.17086](https://doi.org/10.7554/elife.17086)

1405 Genzel, L., Dragoi, G., Frank, L., Ganguly, K., de la Prida, L., Pfeiffer, B., & Robertson, E. (2020). A
1406 consensus statement: defining terms for reactivation analysis. *Philosophical Transactions of the*
1407 *Royal Society B: Biological Sciences*, 375(1799), 20200001. doi: [10.1098/rstb.2020.0001](https://doi.org/10.1098/rstb.2020.0001)

1408 Gerrard, J. L., Kudrimoti, H., McNaughton, B. L., & Barnes, C. A. (2001). Reactivation of hippocam-
1409 pal ensemble activity patterns in the aging rat. *Behavioral Neuroscience*, 115(6), 1180–1192.
1410 doi: [10.1037/0735-7044.115.6.1180](https://doi.org/10.1037/0735-7044.115.6.1180)

1411 Gershman, S. J., & Daw, N. D. (2017). Reinforcement learning and episodic memory in hu-
1412 mans and animals: An integrative framework. *Annual Review of Psychology*, 68(1), 101–128.
1413 doi: [10.1146/annurev-psych-122414-033625](https://doi.org/10.1146/annurev-psych-122414-033625)

1414 Gershman, S. J., & Niv, Y. (2010). Learning latent structure: Carving nature at its joints. *Current*
1415 *Opinion in Neurobiology*, 20(2), 251–256. doi: [10.1016/j.conb.2010.02.008](https://doi.org/10.1016/j.conb.2010.02.008)

1416 Gershman, S. J., Radulescu, A., Norman, K. A., & Niv, Y. (2014). Statistical computations
1417 underlying the dynamics of memory updating. *PLOS Computational Biology*, 10(11), 1–13.
1418 doi: [10.1371/journal.pcbi.1003939](https://doi.org/10.1371/journal.pcbi.1003939)

1419 Girardeau, G., Benchenane, K., Wiener, S. I., Buzsáki, G., & Zugaro, M. B. (2009). Selective suppres-
1420 sion of hippocampal ripples impairs spatial memory. *Nature Neuroscience*, 12(10), 1222–1223.
1421 doi: [10.1038/nn.2384](https://doi.org/10.1038/nn.2384)

1422 Girardeau, G., Inema, I., & Buzsáki, G. (2017). Reactivations of emotional memory in
1423 the hippocampus–amygdala system during sleep. *Nature Neuroscience*, 20(11), 1634–1642.
1424 doi: [10.1038/nn.4637](https://doi.org/10.1038/nn.4637)

- 1425 Gomperts, S. N., Kloosterman, F., & Wilson, M. A. (2015). VTA neurons coordinate with the
1426 hippocampal reactivation of spatial experience. *eLife*, *4*. doi: [10.7554/elife.05360](https://doi.org/10.7554/elife.05360)
- 1427 Gridchyn, I., Schoenenberger, P., O'Neill, J., & Csicsvari, J. (2020). Assembly-specific disruption of hip-
1428 pocampal replay leads to selective memory deficit. *Neuron*. doi: [10.1016/j.neuron.2020.01.021](https://doi.org/10.1016/j.neuron.2020.01.021)
- 1429 Grieves, R. M., & Jeffery, K. J. (2017). The representation of space in the brain. *Behavioural Processes*,
1430 *135*, 113 - 131. doi: [10.1016/j.beproc.2016.12.012](https://doi.org/10.1016/j.beproc.2016.12.012)
- 1431 Gruber, M. J., Ritchey, M., Wang, S.-F., Doss, M. K., & Ranganath, C. (2016). Post-learning
1432 hippocampal dynamics promote preferential retention of rewarding events. *Neuron*, *89*(5), 1110–
1433 1120. doi: [10.1016/j.neuron.2016.01.017](https://doi.org/10.1016/j.neuron.2016.01.017)
- 1434 Gulati, T., Guo, L., Ramanathan, D. S., Bodepudi, A., & Ganguly, K. (2017). Neural reactiva-
1435 tions during sleep determine network credit assignment. *Nature Neuroscience*, *20*(9), 1277–1284.
1436 doi: [10.1038/nn.4601](https://doi.org/10.1038/nn.4601)
- 1437 Guo, Z. D., Pires, B. A., Piot, B., Grill, J.-B., Althé, F., Munos, R., & Azar, M. G. (2020). Bootstrap
1438 latent-predictive representations for multitask reinforcement learning. In H. D. III & A. Singh
1439 (Eds.), *Proceedings of the 37th international conference on machine learning* (Vol. 119, pp. 3875–
1440 3886). PMLR.
- 1441 Gupta, A. S., van der Meer, M. A., Touretzky, D. S., & Redish, A. D. (2010). Hippocampal replay is not
1442 a simple function of experience. *Neuron*, *65*(5), 695 - 705. doi: [10.1016/j.neuron.2010.01.034](https://doi.org/10.1016/j.neuron.2010.01.034)
- 1443 Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., & Moser, E. I. (2005). Microstructure of a spatial
1444 map in the entorhinal cortex. *Nature*, *436*(7052), 801–806. doi: [10.1038/nature03721](https://doi.org/10.1038/nature03721)
- 1445 Haga, T., & Fukai, T. (2018). Recurrent network model for learning goal-directed sequences through
1446 reverse replay. *eLife*, *7*. doi: [10.7554/elife.34171](https://doi.org/10.7554/elife.34171)
- 1447 Hardt, O., Nader, K., & Nadel, L. (2013). Decay happens: the role of active forgetting in memory.
1448 *Trends in Cognitive Sciences*, *17*(3), 111-120. doi: [10.1016/j.tics.2013.01.001](https://doi.org/10.1016/j.tics.2013.01.001)
- 1449 Harvey, C. D., Coen, P., & Tank, D. W. (2012). Choice-specific sequences in parietal cortex during a
1450 virtual-navigation decision task. *Nature*, *484*(7392), 62–68. doi: [10.1038/nature10918](https://doi.org/10.1038/nature10918)
- 1451 Hassabis, D., Kumaran, D., Summerfield, C., & Botvinick, M. (2017). Neuroscience-inspired artificial
1452 intelligence. *Neuron*, *95*(2), 245–258. doi: [10.1016/j.neuron.2017.06.011](https://doi.org/10.1016/j.neuron.2017.06.011)
- 1453 Hassabis, D., & Maguire, E. A. (2007). Deconstructing episodic memory with construction. *Trends in*
1454 *Cognitive Sciences*, *11*(7), 299 - 306. doi: [10.1016/j.tics.2007.05.001](https://doi.org/10.1016/j.tics.2007.05.001)
- 1455 Helfrich, R. F., Lendner, J. D., Mander, B. A., Guillen, H., Paff, M., Mnatsakanyan, L., ... Knight,
1456 R. T. (2019). Bidirectional prefrontal-hippocampal dynamics organize information transfer
1457 during sleep in humans. *Nature Communications*, *10*(1). doi: [10.1038/s41467-019-11444-x](https://doi.org/10.1038/s41467-019-11444-x)
- 1458 Heller, A. S., & Bagot, R. C. (2020). Is hippocampal replay a mechanism for anxiety and depression?
1459 *JAMA Psychiatry*. doi: [10.1001/jamapsychiatry.2019.4788](https://doi.org/10.1001/jamapsychiatry.2019.4788)
- 1460 Herszage, J., & Censor, N. (2018). Modulation of learning and memory: A
1461 shared framework for interference and generalization. *Neuroscience*, *392*, 270 - 280.
1462 doi: [10.1016/j.neuroscience.2018.08.006](https://doi.org/10.1016/j.neuroscience.2018.08.006)
- 1463 Hoffman, K. L., & McNaughton, B. L. (2002). Coordinated reactivation of distributed memory traces
1464 in primate neocortex. *Science*, *297*(5589), 2070–2073. doi: [10.1126/science.1073538](https://doi.org/10.1126/science.1073538)
- 1465 Horgan, D., Quan, J., Budden, D., Barth-Maron, G., Hessel, M., van Hasselt, H., & Silver, D. (2018).
1466 Distributed prioritized experience replay. *arXiv e-prints*, arXiv:1803.00933.
- 1467 Høydal, Ø. A., Skytøen, E. R., Andersson, S. O., Moser, M.-B., & Moser, E. I. (2019). Object-vector
1468 coding in the medial entorhinal cortex. *Nature*, *568*(7752), 400–404. doi: [10.1038/s41586-019-](https://doi.org/10.1038/s41586-019-)

- 1470 Igloi, K., Gaggioni, G., Sterpenich, V., & Schwartz, S. (2015). A nap to recap or how reward regulates
1471 hippocampal-prefrontal memory networks during daytime sleep in humans. *eLife*, *4*, e07903.
1472 doi: [10.7554/eLife.07903](https://doi.org/10.7554/eLife.07903)
- 1473 Jackson, J. C., Johnson, A., & Redish, A. D. (2006). Hippocampal sharp waves and reactivation
1474 during awake states depend on repeated sequential experience. *Journal of Neuroscience*, *26*(48),
1475 12415–12426. doi: [10.1523/jneurosci.4118-06.2006](https://doi.org/10.1523/jneurosci.4118-06.2006)
- 1476 Jadhav, S. P., Kemere, C., German, P. W., & Frank, L. M. (2012). Awake hippocampal sharp-wave
1477 ripples support spatial memory. *Science*, *336*(6087), 1454–1458. doi: [10.1126/science.1217230](https://doi.org/10.1126/science.1217230)
- 1478 Jadhav, S. P., Rothschild, G., Roumis, D. K., & Frank, L. M. (2016). Coordinated excitation and
1479 inhibition of prefrontal ensembles during awake hippocampal sharp-wave ripple events. *Neuron*,
1480 *90*(1), 113–127. doi: [10.1016/j.neuron.2016.02.010](https://doi.org/10.1016/j.neuron.2016.02.010)
- 1481 Jafarpour, A., Penny, W., Barnes, G., Knight, R. T., & Duzel, E. (2017). Working memory replay
1482 prioritizes weakly attended events. *eNeuro*, *4*(4), 1–11. doi: [10.1523/eneuro.0171-17.2017](https://doi.org/10.1523/eneuro.0171-17.2017)
- 1483 Ji, D., & Wilson, M. A. (2006). Coordinated memory replay in the visual cortex and hippocampus
1484 during sleep. *Nature Neuroscience*, *10*(1), 100–107. doi: [10.1038/nn1825](https://doi.org/10.1038/nn1825)
- 1485 Johnson, A., & Redish, A. D. (2007). Neural ensembles in CA3 transiently encode paths for-
1486 ward of the animal at a decision point. *Journal of Neuroscience*, *27*(45), 12176–12189.
1487 doi: [10.1523/jneurosci.3761-07.2007](https://doi.org/10.1523/jneurosci.3761-07.2007)
- 1488 Joo, H. R., & Frank, L. M. (2018). The hippocampal sharp wave–ripple in memory retrieval for immedi-
1489 ate use and consolidation. *Nature Reviews Neuroscience*, *19*(12), 744–757. doi: [10.1038/s41583-
1490 018-0077-1](https://doi.org/10.1038/s41583-018-0077-1)
- 1491 Juechems, K., & Summerfield, C. (2019). Where does value come from? *Trends in Cognitive Sciences*.
1492 doi: [10.1016/j.tics.2019.07.012](https://doi.org/10.1016/j.tics.2019.07.012)
- 1493 Kaefer, K., Nardin, M., Blahna, K., & Csicsvari, J. (2020). Replay of behavioral sequences in the
1494 medial prefrontal cortex during rule switching. *Neuron*. doi: [10.1016/j.neuron.2020.01.015](https://doi.org/10.1016/j.neuron.2020.01.015)
- 1495 Kaplan, R., Campo, A. T., Bush, D., King, J., Principe, A., Koster, R., . . . Friston, K. J. (2020). Hu-
1496 man hippocampal theta oscillations reflect sequential dependencies during spatial planning. *Cog-
1497 nitive Neuroscience*, *11*(3), 122–131. (PMID: 31617790) doi: [10.1080/17588928.2019.1676711](https://doi.org/10.1080/17588928.2019.1676711)
- 1498 Kaplan, R., Schuck, N. W., & Doeller, C. F. (2017). The role of mental maps in decision-making.
1499 *Trends in Neurosciences*, *40*(5), 256 - 259. doi: [10.1016/j.tins.2017.03.002](https://doi.org/10.1016/j.tins.2017.03.002)
- 1500 Karlsson, M. P., & Frank, L. M. (2009). Awake replay of remote experiences in the hippocampus.
1501 *Nature Neuroscience*, *12*(7), 913–918. doi: [10.1038/nn.2344](https://doi.org/10.1038/nn.2344)
- 1502 Kay, K., Chung, J. E., Sosa, M., Schor, J. S., Karlsson, M. P., Larkin, M. C., . . . Frank, L. M. (2020).
1503 Constant sub-second cycling between representations of possible futures in the hippocampus.
1504 *Cell*. doi: [10.1016/j.cell.2020.01.014](https://doi.org/10.1016/j.cell.2020.01.014)
- 1505 King, C., Henze, D. A., Leinekugel, X., & Buzsáki, G. (1999). Hebbian modification of a hippocampal
1506 population pattern in the rat. *The Journal of Physiology*, *521*(1), 159–167. doi: [10.1111/j.1469-
1507 7793.1999.00159.x](https://doi.org/10.1111/j.1469-7793.1999.00159.x)
- 1508 Kudrimoti, H. S., Barnes, C. A., & McNaughton, B. L. (1999). Reactivation of hippocampal cell assem-
1509 blies: Effects of behavioral state, experience, and eeg dynamics. *The Journal of Neuroscience*,
1510 *19*(10), 4090–4101. doi: [10.1523/jneurosci.19-10-04090.1999](https://doi.org/10.1523/jneurosci.19-10-04090.1999)
- 1511 Kuhl, B. A., Shah, A. T., DuBrow, S., & Wagner, A. D. (2010). Resistance to forgetting associated with
1512 hippocampus-mediated reactivation during new learning. *Nature Neuroscience*, *13*(4), 501–506.

1513 doi: [10.1038/nn.2498](https://doi.org/10.1038/nn.2498)

1514 Kumaran, D. (2012). What representations and computations underpin the contribution of the
1515 hippocampus to generalization and inference? *Frontiers in Human Neuroscience*, *6*, 157.
1516 doi: [10.3389/fnhum.2012.00157](https://doi.org/10.3389/fnhum.2012.00157)

1517 Kumaran, D., Hassabis, D., & McClelland, J. L. (2016). What learning systems do intelligent agents
1518 need? Complementary learning systems theory updated. *Trends in Cognitive Sciences*, *20*(7),
1519 512–534. doi: [10.1016/j.tics.2016.05.004](https://doi.org/10.1016/j.tics.2016.05.004)

1520 Kumaran, D., & McClelland, J. L. (2012). Generalization through the recurrent interaction of
1521 episodic memories: A model of the hippocampal system. *Psychological Review*, *119*(3), 573–
1522 616. doi: [10.1037/a0028681](https://doi.org/10.1037/a0028681)

1523 Kuriyama, K., Stickgold, R., & Walker, M. P. (2004). Sleep-dependent learning and motor-skill
1524 complexity. *Learning & Memory*, *11*(6), 705–713. doi: [10.1101/lm.76304](https://doi.org/10.1101/lm.76304)

1525 Kurth-Nelson, Z., Economides, M., Dolan, R. J., & Dayan, P. (2016). Fast sequences of non-spatial
1526 state representations in humans. *Neuron*, *91*(1), 194–204. doi: [10.1016/j.neuron.2016.05.028](https://doi.org/10.1016/j.neuron.2016.05.028)

1527 Lansink, C. S., Goltstein, P. M., Lankelma, J. V., Joosten, R. N. J. M. A., McNaughton, B. L., &
1528 Pennartz, C. M. A. (2008). Preferential reactivation of motivationally relevant information in
1529 the ventral striatum. *Journal of Neuroscience*, *28*(25), 6372–6382. doi: [10.1523/jneurosci.1054-
1530 08.2008](https://doi.org/10.1523/jneurosci.1054-08.2008)

1531 Lansink, C. S., Goltstein, P. M., Lankelma, J. V., McNaughton, B. L., & Pennartz, C. M. A. (2009).
1532 Hippocampus leads ventral striatum in replay of place-reward information. *PLoS Biology*, *7*(8),
1533 e1000173. doi: [10.1371/journal.pbio.1000173](https://doi.org/10.1371/journal.pbio.1000173)

1534 Lee, A. K., & Wilson, M. A. (2002). Memory of sequential experience in the hippocampus during slow
1535 wave sleep. *Neuron*, *36*(6), 1183–1194. doi: [10.1016/s0896-6273\(02\)01096-6](https://doi.org/10.1016/s0896-6273(02)01096-6)

1536 Lee, S. W., O’Doherty, J. P., & Shimojo, S. (2015). Neural computations mediating one-shot learning
1537 in the human brain. *PLOS Biology*, *13*(4), 1–36. doi: [10.1371/journal.pbio.1002137](https://doi.org/10.1371/journal.pbio.1002137)

1538 Leibold, C. (2020). A model for navigation in unknown environments based on a reservoir of hip-
1539 pocampal sequences. *Neural Networks*, *124*, 328–342. doi: [10.1016/j.neunet.2020.01.014](https://doi.org/10.1016/j.neunet.2020.01.014)

1540 Lengyel, M., & Dayan, P. (2007). Hippocampal contributions to control: The third way. In *Proceedings*
1541 *of the 20th international conference on neural information processing systems* (pp. 889–896). Red
1542 Hook, NY, USA: Curran Associates Inc.

1543 Leong, Y. C., Radulescu, A., Daniel, R., DeWoskin, V., & Niv, Y. (2017). Dynamic interaction between
1544 reinforcement learning and attention in multidimensional environments. *Neuron*, *93*(2), 451 -
1545 463. doi: [10.1016/j.neuron.2016.12.040](https://doi.org/10.1016/j.neuron.2016.12.040)

1546 Lewis, P. A., & Bendor, D. (2019). How targeted memory reactivation promotes the
1547 selective strengthening of memories in sleep. *Current Biology*, *29*(18), R906–R912.
1548 doi: [10.1016/j.cub.2019.08.019](https://doi.org/10.1016/j.cub.2019.08.019)

1549 Lewis, P. A., Knoblich, G., & Poe, G. (2018). How memory replay in sleep boosts creative problem-
1550 solving. *Trends in Cognitive Sciences*, *22*(6), 491–503. doi: [10.1016/j.tics.2018.03.009](https://doi.org/10.1016/j.tics.2018.03.009)

1551 Lin, L. J. (1991). Programming robots using reinforcement learning and teaching. In *Association for*
1552 *the advancement of artificial intelligence* (pp. 781–786).

1553 Lin, L.-J. (1992). Self-improving reactive agents based on reinforcement learning, planning and teach-
1554 ing. *Machine Learning*, *8*(3-4), 293–321. doi: [10.1007/bf00992699](https://doi.org/10.1007/bf00992699)

1555 Lin, L.-J. (1993). Reinforcement learning for robots using neural networks. *Carnegie Mellon Univer-*
1556 *sity*.

- 1557 Liu, B., Ye, X., Gao, Y., Dong, X., Wang, X., & Liu, B. (2019). Forward-looking imaginative planning
1558 framework combined with prioritized-replay double DQN. In *2019 5th international conference*
1559 *on control, automation and robotics (iccar)* (p. 336-341). doi: [10.1109/ICCAR.2019.8813352](https://doi.org/10.1109/ICCAR.2019.8813352)
- 1560 Liu, Y., Dolan, R., Kurth-Nelson, Z., & Behrens, T. E. (2020). Measuring sequences of representations
1561 with temporally delayed linear modelling. *bioRxiv*. doi: [10.1101/2020.04.30.066407](https://doi.org/10.1101/2020.04.30.066407)
- 1562 Liu, Y., Dolan, R. J., Kurth-Nelson, Z., & Behrens, T. E. (2019). Human replay spontaneously
1563 reorganizes experience. *Cell*, *178*(3), 640–652. doi: [10.1016/j.cell.2019.06.012](https://doi.org/10.1016/j.cell.2019.06.012)
- 1564 Louie, K., & Wilson, M. A. (2001). Temporally structured replay of awake hippocampal ensemble
1565 activity during rapid eye movement sleep. *Neuron*, *29*(1), 145–156. doi: [10.1016/s0896-
1566 6273\(01\)00186-6](https://doi.org/10.1016/s0896-6273(01)00186-6)
- 1567 MacDonald, C. J., Lepage, K. Q., Eden, U. T., & Eichenbaum, H. (2011). Hippocampal
1568 “time cells” bridge the gap in memory for discontinuous events. *Neuron*, *71*(4), 737 - 749.
1569 doi: [10.1016/j.neuron.2011.07.012](https://doi.org/10.1016/j.neuron.2011.07.012)
- 1570 Magee, J. C. (1997). A synaptically controlled, associative signal for hebbian plasticity in hippocampal
1571 neurons. *Science*, *275*(5297), 209–213. doi: [10.1126/science.275.5297.209](https://doi.org/10.1126/science.275.5297.209)
- 1572 Mahadevan, S., & Maggioni, M. (2007). Proto-value functions: A laplacian framework for learning
1573 representation and control in markov decision processes. *Journal of Machine Learning Research*,
1574 *8*(74), 2169-2231.
- 1575 Marr, D. (1971). Simple memory: A theory for archicortex. *Philosophical Transactions of the Royal*
1576 *Society B: Biological Sciences*, *262*(841), 23–81. doi: [10.1098/rstb.1971.0078](https://doi.org/10.1098/rstb.1971.0078)
- 1577 Mattar, M. G., & Daw, N. D. (2018). Prioritized memory access explains planning and hippocampal
1578 replay. *Nature Neuroscience*, *21*(11), 1609–1617. doi: [10.1038/s41593-018-0232-z](https://doi.org/10.1038/s41593-018-0232-z)
- 1579 Maurer, A. P., & Nadel, L. (2021). The continuity of context: A role for the hippocampus. *Trends in*
1580 *Cognitive Sciences*. doi: [10.1016/j.tics.2020.12.007](https://doi.org/10.1016/j.tics.2020.12.007)
- 1581 McClelland, J. L., McNaughton, B. L., & O’Reilly, R. C. (1995). Why there are complementary learning
1582 systems in the hippocampus and neocortex: Insights from the successes and failures of connection-
1583 ist models of learning and memory. *Psychological Review*, *102*(3), 419–457. doi: [10.1037/0033-
1584 295x.102.3.419](https://doi.org/10.1037/0033-295x.102.3.419)
- 1585 McCloskey, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks:
1586 The sequential learning problem. *Psychology of Learning and Motivation*, *24*, 109–165.
1587 doi: [10.1016/s0079-7421\(08\)60536-8](https://doi.org/10.1016/s0079-7421(08)60536-8)
- 1588 McDevitt, E. A., Duggan, K. A., & Mednick, S. C. (2015). REM sleep rescues learning from interfer-
1589 ence. *Neurobiology of Learning and Memory*, *122*, 51–62. doi: [10.1016/j.nlm.2014.11.015](https://doi.org/10.1016/j.nlm.2014.11.015)
- 1590 McNamara, C. G., Tejero-Cantero, Á., Trouche, S., Campo-Urriza, N., & Dupret, D. (2014). Dopamin-
1591 ergic neurons promote hippocampal reactivation and spatial memory persistence. *Nature Neu-
1592 roscience*, *17*(12), 1658–1660. doi: [10.1038/nn.3843](https://doi.org/10.1038/nn.3843)
- 1593 Michon, F., Sun, J.-J., Kim, C. Y., Ciliberti, D., & Kloosterman, F. (2019). Post-learning hippocampal
1594 replay selectively reinforces spatial memory for highly rewarded locations. *Current Biology*, *29*(9),
1595 1436–1444.e5. doi: [10.1016/j.cub.2019.03.048](https://doi.org/10.1016/j.cub.2019.03.048)
- 1596 Miller, K. J., & Venditto, S. J. C. (2021). Multi-step planning in the brain. *Current Opinion in*
1597 *Behavioral Sciences*, *38*, 29 - 39. doi: [10.1016/j.cobeha.2020.07.003](https://doi.org/10.1016/j.cobeha.2020.07.003)
- 1598 Minsky, M. (1961). Steps toward artificial intelligence. *Proceedings of the IRE*, *49*(1), 8-30.
1599 doi: [10.1109/JRPROC.1961.287775](https://doi.org/10.1109/JRPROC.1961.287775)
- 1600 Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., & Riedmiller, M.

- 1601 (2013). Playing Atari with Deep Reinforcement Learning. *arXiv e-prints*, arXiv:1312.5602.
- 1602 Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... et al.
1603 (2015). Human-level control through deep reinforcement learning. *Nature*, *518*(7540), 529–533.
1604 doi: [10.1038/nature14236](https://doi.org/10.1038/nature14236)
- 1605 Mobbs, D., Headley, D. B., Ding, W., & Dayan, P. (2020). Space, time, and fear: Survival computations
1606 along defensive circuits. *Trends in Cognitive Sciences*. doi: [10.1016/j.tics.2019.12.016](https://doi.org/10.1016/j.tics.2019.12.016)
- 1607 Momennejad, I. (2020). Learning structures: Predictive representations, replay, and generalization.
1608 *Current Opinion in Behavioral Sciences*, *32*, 155–166. doi: [10.1016/j.cobeha.2020.02.017](https://doi.org/10.1016/j.cobeha.2020.02.017)
- 1609 Momennejad, I., Otto, A. R., Daw, N. D., & Norman, K. A. (2018). Offline replay supports planning
1610 in human reinforcement learning. *eLife*, *7*, e32548. doi: [10.7554/eLife.32548](https://doi.org/10.7554/eLife.32548)
- 1611 Momennejad, I., Russek, E. M., Cheong, J. H., Botvinick, M. M., Daw, N. D., & Gershman, S. J. (2017).
1612 The successor representation in human reinforcement learning. *Nature Human Behaviour*, *1*(9),
1613 680–692. doi: [10.1038/s41562-017-0180-8](https://doi.org/10.1038/s41562-017-0180-8)
- 1614 Monaco, J. D., Rao, G., Roth, E. D., & Knierim, J. J. (2014). Attentive scanning behavior
1615 drives one-trial potentiation of hippocampal place fields. *Nature Neuroscience*, *17*(5), 725–731.
1616 doi: [10.1038/nn.3687](https://doi.org/10.1038/nn.3687)
- 1617 Moore, A. W., & Atkeson, C. G. (1993). Prioritized sweeping: Reinforcement learning with less data
1618 and less time. *Machine Learning*, *13*(1), 103–130. doi: [10.1007/bf00993104](https://doi.org/10.1007/bf00993104)
- 1619 Moser, E. I., Kropff, E., & Moser, M.-B. (2008). Place cells, grid cells, and the
1620 brain's spatial representation system. *Annual Review of Neuroscience*, *31*(1), 69–89.
1621 doi: [10.1146/annurev.neuro.31.061307.090723](https://doi.org/10.1146/annurev.neuro.31.061307.090723)
- 1622 Muenzinger, K. F., & Fletcher, F. M. (1936). Motivation in learning. vi. escape from electric shock
1623 compared with hunger-food tension in the visual discrimination habit. *Journal of Comparative*
1624 *Psychology*, *22*(1), 79–91.
- 1625 Nádasdy, Z., Hirase, H., Czurkó, A., Csicsvari, J., & Buzsáki, G. (1999). Replay and time compression
1626 of recurring spike sequences in the hippocampus. *The Journal of Neuroscience*, *19*(21), 9497–
1627 9507. doi: [10.1523/jneurosci.19-21-09497.1999](https://doi.org/10.1523/jneurosci.19-21-09497.1999)
- 1628 Niethard, N., & Born, J. (2020). A backup of hippocampal spatial code outside the hip-
1629 pocampus? new light on systems memory consolidation. *Neuron*, *106*(2), 204 - 206.
1630 doi: doi.org/10.1016/j.neuron.2020.03.034
- 1631 Niv, Y. (2019). Learning task-state representations. *Nature Neuroscience*, *22*(10), 1544–1553.
1632 doi: [10.1038/s41593-019-0470-8](https://doi.org/10.1038/s41593-019-0470-8)
- 1633 Niv, Y., Daniel, R., Geana, A., Gershman, S. J., Leong, Y. C., Radulescu, A., & Wilson, R. C.
1634 (2015). Reinforcement learning in multidimensional environments relies on attention mechanisms.
1635 *Journal of Neuroscience*, *35*(21), 8145–8157. doi: [10.1523/JNEUROSCI.2978-14.2015](https://doi.org/10.1523/JNEUROSCI.2978-14.2015)
- 1636 Norman, Y., Yeagle, E. M., Khuvis, S., Harel, M., Mehta, A. D., & Malach, R. (2019). Hippocampal
1637 sharp-wave ripples linked to visual episodic recollection in humans. *Science*, *365*(6454), eaax1030.
1638 doi: [10.1126/science.aax1030](https://doi.org/10.1126/science.aax1030)
- 1639 O'Keefe, J., & Conway, D. H. (1978). Hippocampal place units in the freely moving rat: Why they
1640 fire where they fire. *Experimental Brain Research*, *31*(4), 573–590. doi: [10.1007/BF00239813](https://doi.org/10.1007/BF00239813)
- 1641 O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. preliminary evidence from
1642 unit activity in the freely-moving rat. *Brain Research*, *34*(1), 171–175. doi: [10.1016/0006-
1643 8993\(71\)90358-1](https://doi.org/10.1016/0006-8993(71)90358-1)
- 1644 O'Keefe, J., & Krupic, J. (2021). Do hippocampal pyramidal cells respond to non-spatial stimuli?

- 1645 *Physiological Reviews*. doi: [10.1152/physrev.00014.2020](https://doi.org/10.1152/physrev.00014.2020)
- 1646 O'Keefe, J., & Nadel, L. (1974). Maps in the brain. *New Scientist*, 62(903), 749–751.
- 1647 O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford: Clarendon Press.
- 1648 Ólafsdóttir, H. F., Barry, C., Saleem, A. B., Hassabis, D., & Spiers, H. J. (2015). Hip-
1649 pocampal place cells construct reward related sequences through unexplored space. *eLife*, 4.
1650 doi: [10.7554/elife.06063](https://doi.org/10.7554/elife.06063)
- 1651 Ólafsdóttir, H. F., Bush, D., & Barry, C. (2018). The role of hippocampal replay in memory and
1652 planning. *Current Biology*, 28(1), R37–R50. doi: [10.1016/j.cub.2017.10.073](https://doi.org/10.1016/j.cub.2017.10.073)
- 1653 Ólafsdóttir, H. F., Carpenter, F., & Barry, C. (2016). Coordinated grid and place cell replay during
1654 rest. *Nature Neuroscience*, 19(6), 792–794. doi: [10.1038/nn.4291](https://doi.org/10.1038/nn.4291)
- 1655 Ólafsdóttir, H. F., Carpenter, F., & Barry, C. (2017). Task demands predict a dy-
1656 namic switch in the content of awake hippocampal replay. *Neuron*, 96(4), 925–935.e6.
1657 doi: [10.1016/j.neuron.2017.09.035](https://doi.org/10.1016/j.neuron.2017.09.035)
- 1658 O'Neill, J., Boccara, C., Stella, F., Schoenenberger, P., & Csicsvari, J. (2017). Superficial layers of the
1659 medial entorhinal cortex replay independently of the hippocampus. *Science*, 355(6321), 184–188.
1660 doi: [10.1126/science.aag2787](https://doi.org/10.1126/science.aag2787)
- 1661 O'Neill, J., Pleydell-Bouverie, B., Dupret, D., & Csicsvari, J. (2010). Play it again: Re-
1662 activation of waking experience and memory. *Trends in Neurosciences*, 33(5), 220–229.
1663 doi: [10.1016/j.tins.2010.01.006](https://doi.org/10.1016/j.tins.2010.01.006)
- 1664 O'Neill, J., Senior, T. J., Allen, K., Huxter, J. R., & Csicsvari, J. (2008). Reactivation of experience-
1665 dependent cell assembly patterns in the hippocampus. *Nature Neuroscience*, 11(2), 209–215.
1666 doi: [10.1038/nn2037](https://doi.org/10.1038/nn2037)
- 1667 O'Reilly, R. C., Bhattacharyya, R., Howard, M. D., & Ketz, N. (2014). Complementary learning
1668 systems. *Cognitive Science*, 38(6), 1229–1248. doi: [10.1111/j.1551-6709.2011.01214.x](https://doi.org/10.1111/j.1551-6709.2011.01214.x)
- 1669 O'Reilly, R. C., & McClelland, J. L. (1994). Hippocampal conjunctive encoding, storage, and recall:
1670 Avoiding a trade-off. *Hippocampus*, 4(6), 661–682. doi: [10.1002/hipo.450040605](https://doi.org/10.1002/hipo.450040605)
- 1671 Oudiette, D., & Paller, K. A. (2013). Upgrading the sleeping brain with targeted memory reactivation.
1672 *Trends in Cognitive Sciences*, 17(3), 142 - 149. doi: [10.1016/j.tics.2013.01.006](https://doi.org/10.1016/j.tics.2013.01.006)
- 1673 Papale, A. E., Zielinski, M. C., Frank, L. M., Jadhav, S. P., & Redish, A. D. (2016). Interplay between
1674 hippocampal sharp-wave-ripple events and vicarious trial and error behaviors in decision making.
1675 *Neuron*, 92(5), 975–982. doi: [10.1016/j.neuron.2016.10.028](https://doi.org/10.1016/j.neuron.2016.10.028)
- 1676 Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., & Wermter, S. (2019). Continual lifelong learning with
1677 neural networks: A review. *Neural Networks*, 113, 54 - 71. doi: [10.1016/j.neunet.2019.01.012](https://doi.org/10.1016/j.neunet.2019.01.012)
- 1678 Pavlides, C., & Winson, J. (1989). Influences of hippocampal place cell firing in the awake state on
1679 the activity of these cells during subsequent sleep episodes. *The Journal of Neuroscience*, 9(8),
1680 2907–2918. doi: [10.1523/jneurosci.09-08-02907.1989](https://doi.org/10.1523/jneurosci.09-08-02907.1989)
- 1681 Peer, M., Brunec, I. K., Newcombe, N. S., & Epstein, R. A. (2020). Structuring knowledge with cogni-
1682 tive maps and cognitive graphs. *Trends in Cognitive Sciences*. doi: [10.1016/j.tics.2020.10.004](https://doi.org/10.1016/j.tics.2020.10.004)
- 1683 Penagos, H., Varela, C., & Wilson, M. A. (2017). Oscillations, neural computations and learning during
1684 wake and sleep. *Current Opinion in Neurobiology*, 44, 193–201. doi: [10.1016/j.conb.2017.05.009](https://doi.org/10.1016/j.conb.2017.05.009)
- 1685 Peng, J., & Williams, R. J. (1993). Efficient learning and planning within the dyna framework.
1686 *Adaptive Behavior*, 1(4), 437–454. doi: [10.1177/105971239300100403](https://doi.org/10.1177/105971239300100403)
- 1687 Pennartz, C. M. A. (2004). The ventral striatum in off-line processing: Ensemble reactivation during
1688 sleep and modulation by hippocampal ripples. *Journal of Neuroscience*, 24(29), 6446–6456.

- doi: [10.1523/jneurosci.0575-04.2004](https://doi.org/10.1523/jneurosci.0575-04.2004)
- 1689
1690 Peters, K. R., Smith, V., & Smith, C. T. (2007). Changes in sleep architecture following mo-
1691 tor learning depend on initial skill level. *Journal of Cognitive Neuroscience*, *19*(5), 817–829.
1692 doi: [10.1162/jocn.2007.19.5.817](https://doi.org/10.1162/jocn.2007.19.5.817)
- 1693 Peyrache, A., Khamassi, M., Benchenane, K., Wiener, S. I., & Battaglia, F. P. (2009). Replay of
1694 rule-learning related neural patterns in the prefrontal cortex during sleep. *Nature Neuroscience*,
1695 *12*(7), 919–926. doi: [10.1038/nn.2337](https://doi.org/10.1038/nn.2337)
- 1696 Pezzulo, G., Donnarumma, F., Maisto, D., & Stoianov, I. (2019). Planning at decision time and in
1697 the background during spatial navigation. *Current Opinion in Behavioral Sciences*, *29*, 69–76.
1698 doi: [10.1016/j.cobeha.2019.04.009](https://doi.org/10.1016/j.cobeha.2019.04.009)
- 1699 Pezzulo, G., Kemere, C., & van der Meer, M. A. (2017). Internally generated hippocampal sequences as
1700 a vantage point to probe future-oriented cognition. *Annals of the New York Academy of Sciences*,
1701 *1396*(1), 144–165. doi: [10.1111/nyas.13329](https://doi.org/10.1111/nyas.13329)
- 1702 Pezzulo, G., van der Meer, M. A., Lansink, C. S., & Pennartz, C. M. (2014). Internally generated
1703 sequences in learning and executing goal-directed behavior. *Trends in Cognitive Sciences*, *18*(12),
1704 647–657. doi: [10.1016/j.tics.2014.06.011](https://doi.org/10.1016/j.tics.2014.06.011)
- 1705 Pfeiffer, B. E., & Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to
1706 remembered goals. *Nature*, *497*(7447), 74–79. doi: [10.1038/nature12112](https://doi.org/10.1038/nature12112)
- 1707 Qin, Y.-L., Menaughton, B. L., Skaggs, W. E., & Barnes, C. A. (1997). Memory reprocess-
1708 ing in corticocortical and hippocampocortical neuronal ensembles. *Philosophical Transac-
1709 tions of the Royal Society of London. Series B: Biological Sciences*, *352*(1360), 1525–1533.
1710 doi: [10.1098/rstb.1997.0139](https://doi.org/10.1098/rstb.1997.0139)
- 1711 Ramanathan, D. S., Gulati, T., & Ganguly, K. (2015). Sleep-dependent reactivation of
1712 ensembles in motor cortex promotes skill consolidation. *PLOS Biology*, *13*(9), 1–25.
1713 doi: [10.1371/journal.pbio.1002263](https://doi.org/10.1371/journal.pbio.1002263)
- 1714 Ranck, J. B. J. (1984). Head direction cells in the deep layer of dorsal presubiculum in freely moving
1715 rats. *Society of Neuroscience*, *10*, 599.
- 1716 Rasch, B., & Born, J. (2007). Maintaining memories by reactivation. *Current Opinion in Neurobiology*,
1717 *17*(6), 698–703. doi: [10.1016/j.comb.2007.11.007](https://doi.org/10.1016/j.comb.2007.11.007)
- 1718 Ratcliff, R. (1990). Connectionist models of recognition memory: Constraints imposed by learning and
1719 forgetting functions. *Psychological Review*, *97*(2), 285–308. doi: [10.1037/0033-295x.97.2.285](https://doi.org/10.1037/0033-295x.97.2.285)
- 1720 Redish, A. D. (1999). *Beyond the cognitive map: From place cells to episodic memory*. MIT Press.
- 1721 Redish, A. D. (2016). Vicarious trial and error. *Nature Reviews Neuroscience*, *17*(3), 147–159.
1722 doi: [10.1038/nrn.2015.30](https://doi.org/10.1038/nrn.2015.30)
- 1723 Redish, A. D., & Touretzky, D. S. (1998). The role of the hippocampus in solving the morris water
1724 maze. *Neural Computation*, *10*(1), 73–111. doi: [10.1162/089976698300017908](https://doi.org/10.1162/089976698300017908)
- 1725 Richmond, L. L., & Zacks, J. M. (2017). Constructing experience: Event models from perception to
1726 action. *Trends in Cognitive Sciences*, *21*(12), 962–980. doi: [10.1016/j.tics.2017.08.005](https://doi.org/10.1016/j.tics.2017.08.005)
- 1727 Roscow, E. L., Jones, M. W., & Lepora, N. F. (2019). Behavioural and computational evidence for
1728 memory consolidation biased by reward-prediction errors. *bioRxiv*. doi: [10.1101/716290](https://doi.org/10.1101/716290)
- 1729 Rothschild, G., Eban, E., & Frank, L. M. (2016). A cortical–hippocampal–cortical loop of in-
1730 formation processing during memory consolidation. *Nature Neuroscience*, *20*(2), 251–259.
1731 doi: [10.1038/nn.4457](https://doi.org/10.1038/nn.4457)
- 1732 Rouhani, N., Norman, K. A., Niv, Y., & Bornstein, A. M. (2020). Reward prediction errors create

- 1733 event boundaries in memory. *Cognition*, 203, 104269. doi: [10.1016/j.cognition.2020.104269](https://doi.org/10.1016/j.cognition.2020.104269)
- 1734 Russek, E. M., Momennejad, I., Botvinick, M. M., Gershman, S. J., & Daw, N. D. (2017). Predictive
1735 representations can link model-based reinforcement learning to model-free mechanisms. *PLoS*
1736 *Computational Biology*, 13(9), e1005768. doi: [10.1371/journal.pcbi.1005768](https://doi.org/10.1371/journal.pcbi.1005768)
- 1737 Schafer, M., & Schiller, D. (2018). Navigating social space. *Neuron*, 100(2), 476–489.
1738 doi: [10.1016/j.neuron.2018.10.006](https://doi.org/10.1016/j.neuron.2018.10.006)
- 1739 Schapiro, A. C., McDevitt, E. A., Chen, L., Norman, K. A., Mednick, S. C., & Rogers, T. T. (2017).
1740 Sleep benefits memory for semantic category structure while preserving exemplar-specific infor-
1741 mation. *Scientific Reports*, 7(1), 14869. doi: [10.1038/s41598-017-12884-5](https://doi.org/10.1038/s41598-017-12884-5)
- 1742 Schapiro, A. C., McDevitt, E. A., Rogers, T. T., Mednick, S. C., & Norman, K. A. (2018). Hu-
1743 man hippocampal replay during rest prioritizes weakly learned information and predicts memory
1744 performance. *Nature Communications*, 9(1). doi: [10.1038/s41467-018-06213-1](https://doi.org/10.1038/s41467-018-06213-1)
- 1745 Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B., & Botvinick, M. M. (2013). Neural
1746 representations of events arise from temporal community structure. *Nature Neuroscience*, 16(4),
1747 486–492. doi: [10.1038/nn.3331](https://doi.org/10.1038/nn.3331)
- 1748 Schapiro, A. C., Turk-Browne, N. B., Botvinick, M. M., & Norman, K. A. (2017). Complementary
1749 learning systems within the hippocampus: a neural network modelling approach to reconciling
1750 episodic memory with statistical learning. *Philosophical Transactions of the Royal Society B:*
1751 *Biological Sciences*, 372(1711), 20160049. doi: [10.1098/rstb.2016.0049](https://doi.org/10.1098/rstb.2016.0049)
- 1752 Schaul, T., Quan, J., Antonoglou, I., & Silver, D. (2015). Prioritized experience replay. *arXiv e-prints*,
1753 arXiv:1511.05952.
- 1754 Schmidt, B., Duin, A. A., & Redish, A. D. (2019). Disrupting the medial prefrontal cortex alters
1755 hippocampal sequences during deliberative decision making. *Journal of Neurophysiology*, 121(6),
1756 1981–2000. doi: [10.1152/jn.00793.2018](https://doi.org/10.1152/jn.00793.2018)
- 1757 Schmidt, C., Peigneux, P., Muto, V., Schenkel, M., Knoblauch, V., Münch, M., ... Cajochen, C.
1758 (2006). Encoding difficulty promotes postlearning changes in sleep spindle activity during nap-
1759 ping. *Journal of Neuroscience*, 26(35), 8976–8982. doi: [10.1523/JNEUROSCI.2464-06.2006](https://doi.org/10.1523/JNEUROSCI.2464-06.2006)
- 1760 Schuck, N. W., Cai, M. B., Wilson, R. C., & Niv, Y. (2016). Human orbitofrontal cortex represents a
1761 cognitive map of state space. *Neuron*, 91(6), 1402 - 1412. doi: [10.1016/j.neuron.2016.08.019](https://doi.org/10.1016/j.neuron.2016.08.019)
- 1762 Schuck, N. W., & Niv, Y. (2019). Sequential replay of nonspatial task states in the human hippocampus.
1763 *Science*, 364(6447), eaaw5181. doi: [10.1126/science.aaw5181](https://doi.org/10.1126/science.aaw5181)
- 1764 Schuck, N. W., Wilson, R., & Niv, Y. (2018). A state representation for reinforcement learning and
1765 decision-making in the orbitofrontal cortex. In R. Morris, A. Bornstein, & A. Shenhav (Eds.),
1766 *Goal-directed decision making* (1st ed., p. 259-278). Academic Press. doi: [10.1016/B978-0-12-812098-9.00012-7](https://doi.org/10.1016/B978-0-12-812098-9.00012-7)
- 1767
- 1768 Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions.
1769 *Journal of neurology, neurosurgery, and psychiatry*, 20(1), 11–21. doi: [10.1136/jnmp.20.1.11](https://doi.org/10.1136/jnmp.20.1.11)
- 1770 Sharpe, M. J., Chang, C. Y., Liu, M. A., Batchelor, H. M., Mueller, L. E., Jones, J. L., ... Schoenbaum,
1771 G. (2017). Dopamine transients are sufficient and necessary for acquisition of model-based
1772 associations. *Nature Neuroscience*, 20(5), 735–742. doi: [10.1038/nn.4538](https://doi.org/10.1038/nn.4538)
- 1773 Shin, H., Lee, J. K., Kim, J., & Kim, J. (2017). Continual learning with deep generative replay. In
1774 I. Guyon et al. (Eds.), *Advances in neural information processing systems 30* (pp. 2990–2999).
1775 Curran Associates, Inc.
- 1776 Shin, J. D., Tang, W., & Jadhav, S. P. (2019). Dynamics of awake hippocampal-prefrontal re-

1777 play for spatial learning and memory-guided decision making. *Neuron*, 104(6), 1110–1125.
1778 doi: [10.1016/j.neuron.2019.09.012](https://doi.org/10.1016/j.neuron.2019.09.012)

1779 Shin, Y. S., & DuBrow, S. (2020). Structuring memory through inference-based event segmentation.
1780 *Topics in Cognitive Science*. doi: [10.1111/tops.12505](https://doi.org/10.1111/tops.12505)

1781 Silva, D., Feng, T., & Foster, D. J. (2015). Trajectory events across hippocampal place cells require
1782 previous experience. *Nature Neuroscience*, 18(12), 1772–1779. doi: [10.1038/nn.4151](https://doi.org/10.1038/nn.4151)

1783 Singer, A. C., Carr, M. F., Karlsson, M. P., & Frank, L. M. (2013). Hippocampal swr activity predicts
1784 correct decisions during the initial learning of an alternation task. *Neuron*, 77(6), 1163–1173.
1785 doi: [10.1016/j.neuron.2013.01.027](https://doi.org/10.1016/j.neuron.2013.01.027)

1786 Singer, A. C., & Frank, L. M. (2009). Rewarded outcomes enhance reactivation of experience in the
1787 hippocampus. *Neuron*, 64(6), 910–921. doi: [10.1016/j.neuron.2009.11.016](https://doi.org/10.1016/j.neuron.2009.11.016)

1788 Sio, U. N., Monaghan, P., & Ormerod, T. (2013). Sleep on it, but only if it is difficult: Effects of sleep
1789 on problem solving. *Memory & Cognition*, 41(2), 159–166. doi: [10.3758/s13421-012-0256-7](https://doi.org/10.3758/s13421-012-0256-7)

1790 Skaggs, W. E., & McNaughton, B. L. (1996). Replay of neuronal firing sequences in rat
1791 hippocampus during sleep following spatial experience. *Science*, 271(5257), 1870–1873.
1792 doi: [10.1126/science.271.5257.1870](https://doi.org/10.1126/science.271.5257.1870)

1793 Spiers, H. J. (2020). The hippocampal cognitive map: One space or many? *Trends in Cognitive*
1794 *Sciences*. doi: [10.1016/j.tics.2019.12.013](https://doi.org/10.1016/j.tics.2019.12.013)

1795 Squire, L. R. (1992). Memory and the hippocampus: A synthesis from findings with rats, monkeys,
1796 and humans. *Psychological Review*, 99(2), 195–231.

1797 Staba, R. J., Wilson, C. L., Bragin, A., Fried, I., & Engel, J. (2002). Quantitative analysis of
1798 high-frequency oscillations (80–500 hz) recorded in human epileptic hippocampus and entorhinal
1799 cortex. *Journal of Neurophysiology*, 88(4), 1743–1752. doi: [10.1152/jn.2002.88.4.1743](https://doi.org/10.1152/jn.2002.88.4.1743)

1800 Stachenfeld, K. L., Botvinick, M. M., & Gershman, S. J. (2017). The hippocampus as a predictive
1801 map. *Nature Neuroscience*, 20(11), 1643–1653. doi: [10.1038/nn.4650](https://doi.org/10.1038/nn.4650)

1802 Staresina, B. P., Alink, A., Kriegeskorte, N., & Henson, R. N. (2013). Awake reactivation predicts
1803 memory in humans. *Proceedings of the National Academy of Sciences*, 110(52), 21159–21164.
1804 doi: [10.1073/pnas.1311989110](https://doi.org/10.1073/pnas.1311989110)

1805 Staresina, B. P., Bergmann, T. O., Bonnefond, M., van der Meij, R., Jensen, O., Deuker, L., . . . Fell, J.
1806 (2015). Hierarchical nesting of slow oscillations, spindles and ripples in the human hippocampus
1807 during sleep. *Nature Neuroscience*, 18(11), 1679–1686. doi: [10.1038/nn.4119](https://doi.org/10.1038/nn.4119)

1808 Steiner, A., & Redish, A. D. (2012). The road not taken: Neural correlates of decision making in
1809 orbitofrontal cortex. *Frontiers in Neuroscience*, 6, 131. doi: [10.3389/fnins.2012.00131](https://doi.org/10.3389/fnins.2012.00131)

1810 Stella, F., Baracska, P., O'Neill, J., & Csicsvari, J. (2019). Hippocampal reactivation
1811 of random trajectories resembling brownian diffusion. *Neuron*, 102(2), 450–461.
1812 doi: [10.1016/j.neuron.2019.01.052](https://doi.org/10.1016/j.neuron.2019.01.052)

1813 Stoianov, I., Maisto, D., & Pezzulo, G. (2020). The hippocampal formation as a hier-
1814 archical generative model supporting generative replay and continual learning. *bioRxiv*.
1815 doi: [10.1101/2020.01.16.908889](https://doi.org/10.1101/2020.01.16.908889)

1816 Sun, C., Yang, W., Martin, J., & Tonegawa, S. (2020). Hippocampal neurons represent events as
1817 transferable units of experience. *Nature Neuroscience*. doi: [10.1038/s41593-020-0614-x](https://doi.org/10.1038/s41593-020-0614-x)

1818 Sutherland, G. R., & McNaughton, B. (2000). Memory trace reactivation in hippocampal and neocorti-
1819 cal neuronal ensembles. *Current Opinion in Neurobiology*, 10(2), 180 - 186. doi: [10.1016/S0959-4388\(00\)00079-9](https://doi.org/10.1016/S0959-4388(00)00079-9)

1820

- 1821 Sutton, R. S. (1990). Integrated architectures for learning, planning, and reacting based on approxi-
1822 mating dynamic programming. In B. Porter & R. Mooney (Eds.), *Machine learning proceedings*
1823 *1990* (pp. 216–224). San Francisco (CA): Morgan Kaufmann. doi: [10.1016/B978-1-55860-141-](https://doi.org/10.1016/B978-1-55860-141-3.50030-4)
1824 [3.50030-4](https://doi.org/10.1016/B978-1-55860-141-3.50030-4)
- 1825 Sutton, R. S. (1991). Dyna, an integrated architecture for learning, planning, and reacting. *ACM*
1826 *SIGART Bulletin*, *2*(4), 160–163. doi: [10.1145/122344.122377](https://doi.org/10.1145/122344.122377)
- 1827 Sutton, R. S., & Barto, A. G. (2018). Reinforcement learning: An introduction. *MIT Press*.
- 1828 Sutton, R. S., Szepesvári, C., Geramifard, A., & Bowling, M. (2012). Dyna-style planning with linear
1829 function approximation and prioritized sweeping. *arXiv e-prints*, *abs/1206.3285*.
- 1830 Swanson, R. A., Levenstein, D., McClain, K., Tingley, D., & Buzsáki, G. (2020). Variable speci-
1831 ficity of memory trace reactivation during hippocampal sharp wave ripples. *Current Opinion*
1832 *in Behavioral Sciences*, *32*, 126–135. (Understanding memory: Which level of analysis?)
1833 doi: [10.1016/j.cobeha.2020.02.008](https://doi.org/10.1016/j.cobeha.2020.02.008)
- 1834 Talamini, L. M., Nieuwenhuis, I. L. C., Takashima, A., & Jensen, O. (2008). Sleep directly following
1835 learning benefits consolidation of spatial associative memory. *Learning & Memory*, *15*(4), 233–
1836 237. doi: [10.1101/lm.771608](https://doi.org/10.1101/lm.771608)
- 1837 Tambini, A., Berners-Lee, A., & Davachi, L. (2017). Brief targeted memory reactivation during the
1838 awake state enhances memory stability and benefits the weakest memories. *Scientific Reports*,
1839 *7*(1). doi: [10.1038/s41598-017-15608-x](https://doi.org/10.1038/s41598-017-15608-x)
- 1840 Tambini, A., & Davachi, L. (2013). Persistence of hippocampal multivoxel patterns into postencoding
1841 rest is related to memory. *Proceedings of the National Academy of Sciences*, *110*(48), 19591–
1842 19596. doi: [10.1073/pnas.1308499110](https://doi.org/10.1073/pnas.1308499110)
- 1843 Tambini, A., & Davachi, L. (2019). Awake reactivation of prior experiences consolidates memories and
1844 biases cognition. *Trends in Cognitive Sciences*, *23*(10), 876–890. doi: [10.1016/j.tics.2019.07.008](https://doi.org/10.1016/j.tics.2019.07.008)
- 1845 Tambini, A., Ketz, N., & Davachi, L. (2010). Enhanced brain correlations during rest are related to
1846 memory for recent experiences. *Neuron*, *65*(2), 280–290. doi: [10.1016/j.neuron.2010.01.001](https://doi.org/10.1016/j.neuron.2010.01.001)
- 1847 Tang, W., & Jadhav, S. P. (2019). Sharp-wave ripples as a signature of hippocampal-prefrontal
1848 reactivation for memory during sleep and waking states. *Neurobiology of Learning and Memory*,
1849 *160*, 11–20. doi: [10.1016/j.nlm.2018.01.002](https://doi.org/10.1016/j.nlm.2018.01.002)
- 1850 Tang, W., Shin, J. D., Frank, L. M., & Jadhav, S. P. (2017). Hippocampal-prefrontal reactivation
1851 during learning is stronger in awake compared with sleep states. *Journal of Neuroscience*, *37*(49),
1852 11789–11805. doi: [10.1523/JNEUROSCI.2291-17.2017](https://doi.org/10.1523/JNEUROSCI.2291-17.2017)
- 1853 Tang, W., Shin, J. D., & Jadhav, S. P. (2021). Multiple time-scales of decision making in the
1854 hippocampus and prefrontal cortex. *eLife*, *10*, e66227. doi: [10.7554/eLife.66227](https://doi.org/10.7554/eLife.66227)
- 1855 Tang, Y., Nyengaard, J. R., De Groot, D. M., & Gundersen, H. J. G. (2001). Total regional
1856 and global number of synapses in the human brain neocortex. *Synapse*, *41*(3), 258–273.
1857 doi: [10.1002/syn.1083](https://doi.org/10.1002/syn.1083)
- 1858 Taube, J., Muller, R., & Ranck, J. (1990). Head-direction cells recorded from the postsubiculum in
1859 freely moving rats. ii. effects of environmental manipulations. *Journal of Neuroscience*, *10*(2),
1860 436–447. doi: [10.1523/JNEUROSCI.10-02-00436.1990](https://doi.org/10.1523/JNEUROSCI.10-02-00436.1990)
- 1861 Tolman, E. C. (1926). A behavioristic theory of ideas. *Psychological Review*, *33*(5), 352.
1862 doi: [10.1037/h0070532](https://doi.org/10.1037/h0070532)
- 1863 Tolman, E. C. (1938). The determiners of behavior at a choice point. *Psychological Review*, *45*(1),
1864 1–41. doi: [10.1037/h0062733](https://doi.org/10.1037/h0062733)

- 1865 Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, *55*(4), 189–208.
1866 doi: [10.1037/h0061626](https://doi.org/10.1037/h0061626)
- 1867 Tompary, A., & Davachi, L. (2017). Consolidation promotes the emergence of representa-
1868 tional overlap in the hippocampus and medial prefrontal cortex. *Neuron*, *96*(1), 228–241.
1869 doi: [10.1016/j.neuron.2017.09.005](https://doi.org/10.1016/j.neuron.2017.09.005)
- 1870 Trettel, S. G., Trimper, J. B., Hwaun, E., Fiete, I. R., & Colgin, L. L. (2019). Grid cell co-activity
1871 patterns during sleep reflect spatial overlap of grid fields during active behaviors. *Nature Neuro-*
1872 *science*, *22*(4), 609–617. doi: [10.1038/s41593-019-0359-6](https://doi.org/10.1038/s41593-019-0359-6)
- 1873 van de Ven, G. M., Siegelmann, H. T., & Tolias, A. S. (2020). Brain-inspired replay for continual learn-
1874 ing with artificial neural networks. *Nature Communications*, *11*(1), 4069. doi: [10.1038/s41467-](https://doi.org/10.1038/s41467-020-17866-2)
1875 [020-17866-2](https://doi.org/10.1038/s41467-020-17866-2)
- 1876 van de Ven, G. M., & Tolias, A. S. (2018). Generative replay with feedback connections as a general
1877 strategy for continual learning. *arXiv e-prints*, arXiv:1809.10635.
- 1878 van de Ven, G. M., Trouche, S., McNamara, C. G., Allen, K., & Dupret, D. (2016). Hippocampal
1879 offline reactivation consolidates recently formed cell assembly patterns during sharp wave-ripples.
1880 *Neuron*, *92*(5), 968–974. doi: [10.1016/j.neuron.2016.10.020](https://doi.org/10.1016/j.neuron.2016.10.020)
- 1881 van Seijen, H., & Sutton, R. (2015). A deeper look at planning as learning from replay. In *International*
1882 *conference on machine learning* (pp. 2314–2322).
- 1883 Vaz, A. P., Inati, S. K., Brunel, N., & Zaghoul, K. A. (2019). Coupled ripple oscillations between
1884 the medial temporal lobe and neocortex retrieve human memory. *Science*, *363*(6430), 975–978.
1885 doi: [10.1126/science.aau8956](https://doi.org/10.1126/science.aau8956)
- 1886 Vaz, A. P., Wittig, J. H., Inati, S. K., & Zaghoul, K. A. (2020). Replay of cortical spiking sequences
1887 during human memory retrieval. *Science*, *367*(6482), 1131–1134. doi: [10.1126/science.aba0672](https://doi.org/10.1126/science.aba0672)
- 1888 Vértés, E., & Sahani, M. (2019). A neurally plausible model learns successor representations in partially
1889 observable environments. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox,
1890 & R. Garnett (Eds.), *Advances in neural information processing systems 32* (pp. 13714–13724).
1891 Curran Associates, Inc.
- 1892 Wang, S., Feng, S. F., & Bornstein, A. (2020). Mixing memory and desire: How memory reactivation
1893 supports deliberative decision-making. *PsyArXiv*. doi: [10.31234/osf.io/5vksj](https://doi.org/10.31234/osf.io/5vksj)
- 1894 Whittington, J. C., Muller, T. H., Mark, S., Chen, G., Barry, C., Burgess, N., & Behrens, T. E. (2020).
1895 The Tolman-Eichenbaum Machine: Unifying space and relational memory through generalization
1896 in the hippocampal formation. *Cell*, *183*(5), 1249 - 1263.e23. doi: [10.1016/j.cell.2020.10.024](https://doi.org/10.1016/j.cell.2020.10.024)
- 1897 Wikenheiser, A. M., Marrero-Garcia, Y., & Schoenbaum, G. (2017). Suppression of ventral hip-
1898 pocampal output impairs integrated orbitofrontal encoding of task structure. *Neuron*, *95*(5),
1899 1197-1207.e3. doi: [10.1016/j.neuron.2017.08.003](https://doi.org/10.1016/j.neuron.2017.08.003)
- 1900 Wikenheiser, A. M., & Redish, A. D. (2013). The balance of forward and back-
1901 ward hippocampal sequences shifts across behavioral states. *Hippocampus*, *23*(1), 22-29.
1902 doi: doi.org/10.1002/hipo.22049
- 1903 Wikenheiser, A. M., & Redish, A. D. (2015a). Decoding the cognitive map: ensemble hip-
1904 pocampal sequences and decision making. *Current Opinion in Neurobiology*, *32*, 8–15.
1905 doi: [10.1016/j.conb.2014.10.002](https://doi.org/10.1016/j.conb.2014.10.002)
- 1906 Wikenheiser, A. M., & Redish, A. D. (2015b). Hippocampal theta sequences reflect current goals.
1907 *Nature Neuroscience*, *18*(2), 289–294. doi: [10.1038/nn.3909](https://doi.org/10.1038/nn.3909)
- 1908 Wikenheiser, A. M., & Schoenbaum, G. (2016). Over the river, through the woods: cognitive maps

1909 in the hippocampus and orbitofrontal cortex. *Nature Reviews Neuroscience*, 17(8), 513–523.
1910 doi: [10.1038/nrn.2016.56](https://doi.org/10.1038/nrn.2016.56)

1911 Wilson, M., & McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during
1912 sleep. *Science*, 265(5172), 676–679. doi: [10.1126/science.8036517](https://doi.org/10.1126/science.8036517)

1913 Wilson, R. C., Takahashi, Y. K., Schoenbaum, G., & Niv, Y. (2014). Orbitofrontal cortex as a cognitive
1914 map of task space. *Neuron*, 81(2), 267–279. doi: [10.1016/j.neuron.2013.11.005](https://doi.org/10.1016/j.neuron.2013.11.005)

1915 Wimmer, E. G., & Büchel, C. (2019). Learning of distant state predictions by the orbitofrontal cortex
1916 in humans. *Nature Communications*, 10(1). doi: [10.1038/s41467-019-10597-z](https://doi.org/10.1038/s41467-019-10597-z)

1917 Wimmer, G. E., & Büchel, C. (2020). Reactivation of pain-related patterns in the hippocam-
1918 pus from single past episodes relates to successful memory-based decision making. *bioRxiv*.
1919 doi: [10.1101/2020.05.29.123893](https://doi.org/10.1101/2020.05.29.123893)

1920 Wittkuhn, L., & Schuck, N. W. (2021). Dynamics of fMRI patterns reflect sub-second activation
1921 sequences and reveal replay in human visual cortex. *Nature Communications*, 12(1), 1795.
1922 doi: [10.1038/s41467-021-21970-2](https://doi.org/10.1038/s41467-021-21970-2)

1923 Wolosin, S. M., Zeithamova, D., & Preston, A. R. (2012). Reward modulation of hippocampal subfield
1924 activation during successful associative encoding and retrieval. *Journal of Cognitive Neuroscience*,
1925 24(7), 1532–1547. doi: [10.1162/jocn.1a.100237](https://doi.org/10.1162/jocn.1a.100237)

1926 Wood, E. R., Dudchenko, P. A., Robitsek, R., & Eichenbaum, H. (2000). Hippocampal neurons encode
1927 information about different types of memory episodes occurring in the same location. *Neuron*,
1928 27(3), 623–633. doi: [10.1016/s0896-6273\(00\)00071-4](https://doi.org/10.1016/s0896-6273(00)00071-4)

1929 Wu, C.-T., Haggerty, D., Kemere, C., & Ji, D. (2017). Hippocampal awake replay in fear memory
1930 retrieval. *Nature Neuroscience*, 20(4), 571–580. doi: [10.1038/nn.4507](https://doi.org/10.1038/nn.4507)

1931 Yu, J. Y., & Frank, L. M. (2015). Hippocampal–cortical interaction in decision making. *Neurobiology
1932 of Learning and Memory*, 117, 34–41. doi: [10.1016/j.nlm.2014.02.002](https://doi.org/10.1016/j.nlm.2014.02.002)

1933 Yu, J. Y., Liu, D. F., Loback, A., Grossrubatscher, I., & Frank, L. M. (2018). Specific hippocampal
1934 representations are linked to generalized cortical representations in memory. *Nature Communi-
1935 cations*, 9(2209). doi: [10.1038/s41467-018-04498-w](https://doi.org/10.1038/s41467-018-04498-w)

1936 Zhang, H., Deuker, L., & Axmacher, N. (2017). Replay in humans - first evidence and open questions.
1937 In N. Axmacher & B. Rasch (Eds.), *Cognitive neuroscience of memory consolidation* (pp. 251–
1938 263). Springer. doi: [10.1007/978-3-319-45066-7_15](https://doi.org/10.1007/978-3-319-45066-7_15)

1939 Zhang, H., Fell, J., & Axmacher, N. (2018). Electrophysiological mechanisms of human memory
1940 consolidation. *Nature Communications*, 9(1). doi: [10.1038/s41467-018-06553-y](https://doi.org/10.1038/s41467-018-06553-y)

1941 Zhang, S., & Sutton, R. S. (2017). A Deeper Look at Experience Replay. *arXiv e-prints*,
1942 arXiv:1712.01275.

1943 Zielinski, M. C., Tang, W., & Jadhav, S. P. (2018). The role of replay and theta sequences
1944 in mediating hippocampal-prefrontal interactions for memory and cognition. *Hippocampus*.
1945 doi: [10.1002/hipo.22821](https://doi.org/10.1002/hipo.22821)