

## ANNALS OF THE NEW YORK ACADEMY OF SCIENCES

Issue: *The Year in Cognitive Neuroscience*

REVIEW

**Internally generated hippocampal sequences as a vantage point to probe future-oriented cognition**Giovanni Pezzulo,<sup>1</sup> Caleb Kemere,<sup>2</sup> and Matthijs A.A. van der Meer<sup>3</sup><sup>1</sup>Institute of Cognitive Sciences and Technologies, National Research Council, Rome, Italy. <sup>2</sup>Electrical and Computer Engineering, Rice University, Houston, Texas. <sup>3</sup>Department of Psychological and Brain Sciences, Dartmouth College, Hanover, New Hampshire

Address for correspondence: Giovanni Pezzulo, Institute of Cognitive Sciences and Technologies, National Research Council, Via S. Martino della Battaglia 44, 00185 Rome, Italy. giovanni.pezzulo@istc.cnr.it; Matthijs A.A. van der Meer, Department of Psychological and Brain Sciences, Dartmouth College, Hanover, NH 03755. mvdm@dartmouth.edu

Information processing in the rodent hippocampus is fundamentally shaped by internally generated sequences (IGSs), expressed during two different network states: *theta sequences*, which repeat and reset at the ~8 Hz theta rhythm associated with active behavior, and punctate *sharp wave-ripple (SWR) sequences* associated with wakeful rest or slow-wave sleep. A potpourri of diverse functional roles has been proposed for these IGSs, resulting in a fragmented conceptual landscape. Here, we advance a unitary view of IGSs, proposing that they reflect an inferential process that samples a policy from the animal's generative model, supported by hippocampus-specific priors. The same inference affords different cognitive functions when the animal is in distinct dynamical modes, associated with specific functional networks. Theta sequences arise when inference is coupled to the animal's action–perception cycle, supporting online spatial decisions, predictive processing, and episode encoding. SWR sequences arise when the animal is decoupled from the action–perception cycle and may support offline cognitive processing, such as memory consolidation, the prospective simulation of spatial trajectories, and imagination. We discuss the empirical bases of this proposal in relation to rodent studies and highlight how the proposed computational principles can shed light on the mechanisms of future-oriented cognition in humans.

**Keywords:** future-oriented cognition; internally generated hippocampal sequences; generative model; prediction; prospection

**Introduction**


*It's a poor sort of memory that only works backwards.*  
—Lewis Carroll

A fascinating open question in cognitive neuroscience is *how we escape the present*,<sup>1</sup> or how we temporarily *detach* from the here and now of the current sensorimotor cycle to engage in forms of cognition such as prospection, imagination, and mental time travel. The classic example of reduced sensory awareness is sleep, which enables the internal generation of rich patterns of neural activity only weakly constrained by sensory input. However, awake brains have periods of unconstrained, spontaneous patterns, as activity in the default mode network (DMN) during off-task periods illustrates.<sup>2–5</sup> When quizzed about their thoughts during these

periods, subjects often report thinking about the future.<sup>6</sup> Crucially, the DMN is also engaged during on-task processing; as in the case of sleep, shared brain circuits support both online situated action coupled to the action–perception cycle and detached, offline cognitive processing that escapes the present. To understand cognition, we must understand the principles that govern neural circuits during and between these stimulus-driven or internally generated “dynamical modes.”<sup>7,8</sup>

A model system for exploring how neural circuits switch between stimulus-driven (also stimulus-evoked) and internally generated (spontaneous or self-organized) modes is the hippocampus, and its role in spatial navigation as well as the encoding and retrieval of episodic memories in support of adaptive action. Studies in rodents have provided

doi: 10.1111/nyas.13329



**Figure 1.** Three kinds of sequences in an ensemble of place cells recorded from hippocampal subfield CA1: behavioral sequence (blue), theta sequences (red), and sharp wave-ripple (SWR) sequence (green). (A) Activity of an ensemble of hippocampal place cells as an animal runs a spatial trajectory (in this case, a left turn on a T-maze); a spatiotemporal sequence of place cells is activated, starting with the place cell with its field at the start of the maze (bottom row; each row shows spikes from a single place cell), and ending with the place cell with its field at the end of the maze (top row). Inset: Zooming in on the activity in the gray square shows repeating, compressed theta sequences (red lines). (B) Activity from the same ensemble as the animal rests away from the track (but in the same room). Two synchronous bursts of activity, associated with network events called SWR complexes, punctuate an otherwise quiet (non-theta) activity regime. The second SWR activates cells in an order similar to that seen during behavior, forming an SWR sequence (green line). Data are adapted from Ref. 191.

access to the content of hippocampal dynamics at fast time scales and allow for temporally precise experimental interventions, providing an ideal vantage point to probe detached and future-oriented cognition.<sup>9–19</sup> Internally generated activity occurs in different brain states (awake and attentive, awake rest, different sleep stages), with most experimental studies focusing on one state only. Similarly, individual laboratories typically study such activity in a narrow behavioral setting, different from that studied in other laboratories. These factors have created a fragmented conceptual landscape containing a number of interpretations. Here, we aim to provide a unified theoretical viewpoint on the different forms of internally generated activity in the rodent hippocampus, consider the functions proposed for these phenomena, and finally provide the bones of a unifying account, whose implications may extend to the study of human future-oriented cognition, too.

### Internally generated sequences in the hippocampus and their roles in detached cognition

Studies of information processing in the rodent hippocampus have historically focused on “place cells.” These neurons tend to be active in specific, spa-

tially restricted areas of a given environment, such that the animal’s location can be accurately decoded from an ensemble of simultaneously recorded place cells.<sup>20,21</sup> It is now well established that the activity of place cells encodes other variables beyond location alone;<sup>22,23</sup> moreover, hippocampal cells can tile nonspatial experience with restricted firing fields (“time cells” or “episode cells”;<sup>24,25</sup>) consistent with the view that hippocampal activity encodes a “continuous record of attended experience.”<sup>26</sup> Nevertheless, examining the activity of place cells from a purely spatial viewpoint continues to provide a useful access point into information processing in the hippocampus, as we discuss next.

By moving around in space, an animal will traverse sequentially the firing fields of a number of place cells, resulting in a spatiotemporal sequence of neural activity at the time scale of behavior. This sequence is readily apparent from a raster plot in which the rows, corresponding to neurons, are ordered according to the location of their place field (Fig. 1A). Note that approximately 7 seconds elapse between the activation of the first place cell, at the bottom of the plot, and the last place cell (at the top). Interestingly, there is additional structure in this sequential activation of place cells:

at the time scale of the theta rhythm ( $\sim 8$  Hz in moving rodents), repeating, temporally compressed sequences are formed (Fig. 1A, inset). In these theta sequences, place cells are active in the same order as in the slower time scale behavioral sequence. Theta sequences are compressed because the time between the peak firing rates of a given pair of cells within a theta cycle is shorter than the time between peak firing rates of the same cells at the overall behavioral time scale (estimates of this compression factor vary, but are in the  $5\text{--}10\times$  range;<sup>27–30</sup> note, however, that these sequences do not necessarily unfold at a constant speed, so any compression factor is an approximation). They are repeating because during each theta cycle, a new sequence is initiated and terminated. As a general rule, theta sequences are *local* (i.e., include the animal's current location) and *forward* (i.e., proceed in the same direction the animal is moving). The period of the theta cycle limits the space that can be spanned by a single theta sequence, although it should be noted that there is substantial diversity between theta sequences, and the late phase of the theta cycle, in particular, is associated with a less precise spatial representation (and perhaps even reverse order<sup>31,32</sup>).

Theta sequences are of interest because their structure appears to be internally generated: sensory input as it presents to the animal does not start out with repeating, compressed structure at the theta time scale. Determining the mechanistic basis of theta sequence generation is a topic of active experimental and computational work, discussed in detail elsewhere (including the connection with the single-cell phenomenon of theta phase precession<sup>33,34</sup>); here, we focus on the possible functional benefits of organizing activity in this striking way.

The main proposals in the literature tend to emphasize two quite different ideas: (1) that theta sequences reflect online, short time scale predictions of possible courses of action and their outcomes, useful in decision making,<sup>35</sup> and (2) that theta sequences facilitate the initial storage of episodic-like memories by arranging spike times to be conducive to spike timing-dependent plasticity.<sup>32</sup> Both of these ideas have proven to be difficult to test directly, although there is indirect support for both. Striking experimental observations of theta sequences extending alternately down one arm and then another as animals appear to deliberate at choice points are suggestive of a role in planning,<sup>36</sup>

but in that particular experiment, animals' choices could not be predicted from the content of theta sequences. A subsequent study<sup>11</sup> found a relationship between the length of theta sequences and the goal animals subsequently ran to, although in these and many other recording studies it is unclear if these activity patterns were in fact contributing to behavior. Similarly, the link to plasticity is also indirect. In a limited number of studies in which pharmacological manipulations have disrupted theta sequences,<sup>37,38</sup> memory performance was also disrupted, although it remains unclear to what extent this deficit results from encoding, retrieval, or nonsequence-specific sources. In general, there is a large body of correlative evidence linking the properties of the theta rhythm to both memory encoding and retrieval,<sup>39–41</sup> but relating these to content (i.e., decoded theta sequences) continues to be an experimental challenge that precludes precise interpretations.

A second type of internally generated activity in the hippocampus occurs during sharp wave-ripple (SWR) complexes, punctate events associated with highly synchronous spiking.<sup>42</sup> SWRs generally occur when the animal is at rest but awake, or in slow-wave sleep, and not when moving. However, particularly in novel environments, the boundary between high-gamma oscillations and SWR is blurred, and it has been reported that SWRs occur in animals during exploration.<sup>43–46</sup> During SWRs, place cells can fire in similar order to that seen during behavior, prompting the familiar term “replay.” As with theta sequences, the timing of these SWR sequences is compressed relative to behavior, with estimates of the compression factor ranging  $10\text{--}40\times$ ,<sup>28,47</sup> but note here too that the speed is likely not constant.<sup>19</sup> Unlike theta sequences, however, SWR sequences can be forward and backward,<sup>48–50</sup> even when the animal has not actually run the behavioral sequence in the backward direction. SWR sequences can include the animal's current location, but often do not, instead spanning a trajectory on the opposite side of a large T-maze<sup>51</sup> or a trajectory in a different maze entirely.<sup>52</sup> Multiple SWRs can be chained together, with the next SWR sequence starting where the previous one left off, to span long distances.<sup>47</sup> While early studies emphasized the content of SWR sequences as replay, several studies have shown that SWR sequences do not simply reflect recent experience, but can include trajectories toward an

upcoming goal, be biased away from the currently rewarded trajectory, be novel combinations of previously experienced paths, and be even trajectories that have not been experienced at all<sup>9,19,43,51,53,54</sup> (but see Ref. 55).

Early theories about SWRs proposed that they contribute to systems consolidation (i.e., a way for time-limited memories stored in the hippocampus to be consolidated into neocortical memory structures<sup>56</sup>). This idea is supported by studies interrupting SWRs during offline processing (sleep) that found impairments in the acquisition of hippocampus-dependent spatial tasks<sup>57,58</sup> (see also Ref. 59) and activity in cortical and subcortical structures temporally aligned to the SWR.<sup>60–62</sup> However, this view does not explain why SWRs are frequently observed during pauses in task performance and appear to signal trajectories leading to a goal location (prospective) rather than replaying the recently taken trajectory (retrospective).<sup>12,63</sup> Interrupting SWRs in the awake state, during task performance, led to a performance impairment on the outbound leg of an alternation task (which required the animals to remember the identity of the previous trial) but not the inbound leg (which required no working memory).<sup>14</sup> These observations are more in line with an interpretation of awake SWRs as memory retrieval or behavioral planning.<sup>42,64,65,192</sup>

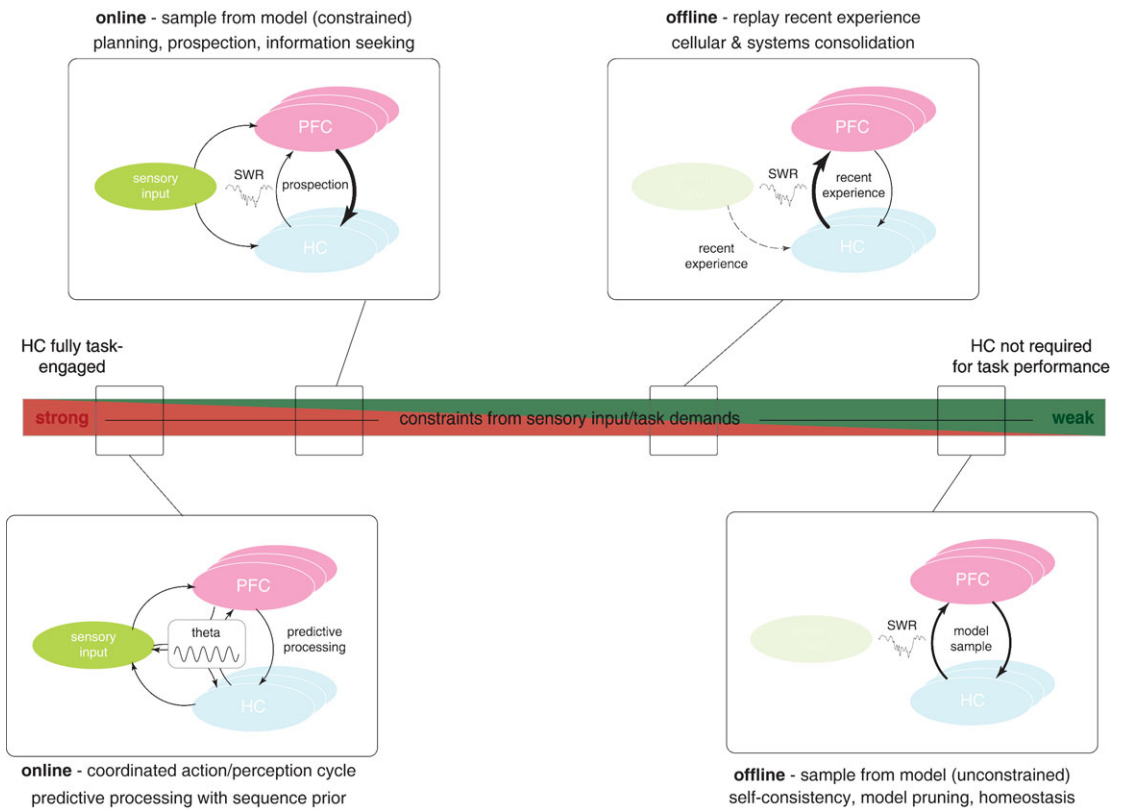
From the above review of the properties of theta sequences and SWR sequences, the main modes of internally generated activity in the hippocampus, a number of open issues are apparent. What is the relationship between theta sequences and SWR sequences? How can we reconcile the differences in content of SWR sequences observed on different tasks? How can the underspecified ideas of consolidation and planning be made more specific as to produce testable predictions and new ideas for experiments and interpretation? To address the above issues, we consider both kinds of internally generated sequences (IGSs) as process components of an overall generative model architecture that supports both adaptive action control and detached cognition, by operating at different dynamical modes (stimulus tied and internally generated) and by forming functional networks that include different brain areas, depending on task demands (see Fig. 2 for a schematic of our overall proposal). This would imply that detached cognition taps the same neurocomputational resources (e.g., those produc-

ing sequential neuronal activity) that afford situated action within action–perception cycles.<sup>171</sup> For this to be possible, the same underlying mechanisms must be flexible enough to operate in different (stimulus-tied and internally generated) modes and at different time scales (i.e., time compression) while also engaging different brain networks in a task-dependent way. To understand how this may be possible, below we discuss computational principles that may underlie sequential neuronal activity and its role across action–perception, memory function, and prospection within our architecture.

### A computational perspective on IGSs

The starting point of this proposal is that the hippocampus forms *internal generative models* jointly with other brain areas, which can be engaged (or not) depending on task demands, including the entorhinal cortex (EC), the ventral striatum (VS), and the prefrontal cortex (PFC).<sup>62</sup> The notion of generative models is central in leading neurophysiological theories of brain structure and function, such as predictive coding, the free energy principle, and the Bayesian brain,<sup>66,67</sup> and is congruent with proposals emphasizing the constructive nature of hippocampal memories and its role in imagination.<sup>5,68</sup> According to these converging perspectives, brains are statistical inference devices that gradually acquire generative models, which encode the statistics of the environment and of agent–environment interactions (i.e., contingencies between actions, sensations, and rewards). The latter are especially important, as ultimately the brain uses generative models to perform inference in the service of adaptive action.

As regularities in the exteroceptive, interoceptive, and proprioceptive domains unfold at different time scales, generative models include various hierarchical levels, which influence one another continuously and reciprocally, in both bottom-up and top-down directions.<sup>69</sup> In predictive coding parlance, higher hierarchical levels propagate predictions (e.g., about an expected sensory stimulus) downward through the hierarchy, while lower hierarchical levels propagate prediction errors upward (e.g., difference between expected and sensed stimulus), and the impact of the “messages” is weighted by their (inverse) uncertainty or precision. This process can be cast with reference to Bayesian inference, if one considers that *prior* beliefs are



**Figure 2.** Schematic illustration of different modes within an overall generative model architecture for prefrontal cortex-hippocampus interactions. The different modes, illustrated in the rounded rectangular panels, are organized along a continuum of constraints from sensory input and task demands. At one extreme, corresponding to full engagement of the hippocampus in the task (bottom left panel), the theta rhythm coordinates an action-perception cycle in which theta sequences create short-term predictions, based on the generative model, that are integrated with sensory input and prediction errors propagated back up to the model. In this theta mode, time-limited encoding of ongoing experience, consisting of episodic-like memory traces, is facilitated by repeating, compressed theta sequences and the associated predictive processing. Top left panel: SWRs shaped by current sensory input are samples from the generative model relevant to the current task and can be used for planning and inference. Their content will reflect current task demands, and will tend to be prospective (e.g., tracing paths to a goal). Top right panel: During offline states, SWR sequences are the vehicles for signaling the content of previously encoded memory episodes (bottom left panel) to cortical regions for model updating (systems consolidation). This consolidation process occurs in the absence of sensory input and is instead driven by short-term plasticity on intrahippocampal synapses. In this mode, SWR content reflects recent experience. Bottom right panel: In the absence of recent short-term plasticity in the hippocampus, SWRs are generated that do not reflect recent experience, but instead serve to improve the generative model by enforcing self-consistency, finding simpler models to prevent overfitting and avoiding catastrophic interference.

iteratively revised as novel information acquired that is (in)compatible with expectations, thus forming *posterior* beliefs that are more informed. The general objective of a predictive coding system is to minimize prediction error (or *free energy*<sup>66</sup>)—in which case, predictions are compatible with the unfolding of sensory events.

The notion of Bayesian inference in (hierarchical) generative models has been used to explain both perception<sup>70–72</sup> and action.<sup>73</sup> In perceptual process-

ing, the higher hierarchical levels encode perceptual hypotheses (e.g., I am looking a dog) that are revised on the basis of bottom-up prediction errors (e.g., hearing a meow, not a bark, as I was expecting) until prediction error (or free energy) is minimized and the “correct” hypothesis is identified (e.g., this is a cat not a dog). However, rather than minimize prediction error by revising my hypothesis, I can also minimize it by acting so as to make the hypothesis true. For example, I can make my original (dog)

hypothesis true by searching for a dog in the visual scene and foveating it. In this case, the predictions propagated by higher hierarchical levels act as goal states, and engaging arc reflexes minimizes the ensuing prediction errors. This scheme can be extended to the planning of *sequences* of actions (or policies) if the agent can predictively compare the (integral of) free energy or surprise conditioned on a series of successive actions (e.g., whether turning twice right or twice left will bring to the expected goal state).<sup>74,75</sup> This example requires engaging the generative model to internally generate and evaluate sequences of predictions. A closely related idea is called planning as inference, in which the current and desired (goal) states are fed to the system (“clamped”) and Bayesian inference is used to “fill gaps” (i.e., to find a trajectory between the current and desired goal location).<sup>76–78</sup> These planning methods that rest on Bayesian inference over generative models (that encode transitions between states and probabilistic relations between states and observations) will become relevant when we next discuss possible neuronal implementations of prospecting.

Various modeling studies have applied the idea of probabilistic (Bayesian) inference in generative models to hippocampal function (and surrounding regions).<sup>79–83</sup> One study highlights that the hippocampus (areas CA1 and CA3), the EC, and prefrontal areas have the right connectivity to encode the agent’s generative model of its environment and to support state estimation, sensory imagery, and other functions, depending on the flow of the information between the areas.<sup>79</sup> In this model, areas CA3–CA1 support state estimation by combining path integration and sensory input from the lateral and medial EC, respectively. The same architecture can also afford sensory imagery when the medial EC receives (virtual) motor commands from the PFC, uses path integration to update the (virtual) position, and communicates it to areas CA3–CA1, which in turn uses its projections to lateral EC to produce predictions of sensory states.

Another proposal posits that a prediction-matching mechanism is implemented within a hierarchical architecture that includes the EC, CA1, and CA3 at increasingly higher hierarchical levels.<sup>27</sup> In this perspective, the most active cell assemblies in CA3 encode predicted positions, whereas EC encodes current inputs. It is when top-down (predictive) and bottom-up (sensory) streams “match”

that CA1 cells fire, thus encoding the updated animal location. This view uses predictive coding to predict current inputs, but it can be expanded to also cover predictions about future stimuli or locations. Accordingly, it has been proposed that the first half of each theta cycle may be stimulus driven and compute the current position (using sensory and path integration information from the lateral and medial EC, respectively), whereas the second half may depend on intrinsic network dynamics involving both grid cells in the medial EC and the hippocampus and compute future spatial positions.<sup>84</sup>

Yet another proposal focuses on a wider brain network formed by the hippocampus and the VS, which may encode two components of a generative model—transitions between locations and contingencies between locations and rewards, respectively—affording Bayesian inference of the best policy or plan.<sup>80</sup> When uncertain, this system sequentially samples from the (combined) internal model to update the “value” of different policies (say, for turning right or left at a junction), which results in the coordinated elicitation of long theta sequences in the hippocampus and of covert reward expectations in the striatum.<sup>64</sup> After sufficient experience, this system can select actions based on “cached” action values without engaging the generative model (hence in a model-free way<sup>85</sup>)—at least while task contingencies remain stable. This latter aspect is consistent with evidence that spatial decisions become hippocampus independent after sufficient experience.<sup>86</sup>

Finally, a series of computational studies explored the idea that various aspects of spatial cognition, including spatial decision making, route planning, model selection, vicarious trial and error (VTE), and the covert evaluation of future spatial trajectories, may be based on probabilistic inference and a common generative model (implemented in the hippocampus and surrounding structures), also discussing various neuronal implementations of (approximate) Bayesian inference.<sup>7,75,79,87,88</sup>

By leveraging and extending these and other models, we discuss below what the notions of statistical inference and generative models can offer to the study of prospective cognition and the relations between stimulus-evoked and internally generated brain activity as applied to theta and SWR sequences.

### Inference


It emerges from the aforementioned models that statistical (Bayesian) inference within generative models affords a homogeneous—yet very flexible—form of computation. We argue that hippocampal IGSs may reflect an inferential process that samples a sequence from the animal's generative model (or sequence generator). During spatial navigation, sequences code for spatial trajectories, which can be evaluated (e.g., to check whether it leads to a goal location) to support use as a plan (or policy) for navigation. In the absence of active navigation, sequences contribute to the construction and maintenance of the generative model (implementing a number of processes in support of memory consolidation). Below, we expand on this idea, discussing how the same generative scheme affords various cognitive functions, including faster prospective processes (e.g., policy selection and control within action–perception cycles) and slower processes (e.g., learning and memory consolidation over longer time scales). The key distinguishing element between these two alternatives is the “dynamical mode” under which the animal operates. The dynamical mode is reflected in the temporal characteristics of brain rhythms: during theta, hippocampal processing is coupled to the action–perception cycle, whereas, during SWRs, there is weaker or even total uncoupling from external events.

**Theta sequences.** How can we represent navigation to a known goal location as an inference problem supported via IGSs? Imagine that the animal knows (i.e., has a model of) the environment and must form and execute a plan to reach a goal location, such as when escaping from a water maze. It can initially select one (e.g., the best *a priori*) among multiple possible policies maintained in its internal model, by “scoring” them according to their usefulness to reach the goal. Then, as it navigates, the animal has to continuously refine this initial plan (e.g., by filling initially incomplete segments and reducing uncertainty about current and goal locations and improving the accuracy of the transitions represented in the internal model) while also remaining flexible enough to exploit novel opportunities or reconsider the situation (e.g., consider other policies) in the light of new evidence. In this perspective, the core process in goal-directed navigation is an inference (e.g., about the spatial tra-

jectory to reach a goal) performed on the basis of the internal model of policies (or plans), and the main role of sensory stimuli is keeping the inference “in register”—not triggering actions directly as in stimulus–response systems. The inference process is coupled to the animal's action–perception cycle, and combines stimulus-tied and internally generated aspects, the latter crucially including prospective aspects, such as the covert evaluation of possible routes.

In Bayesian inference terms, one can think of the policies as hypotheses (on future trajectories, or “where to go next”), which continuously compete and are updated in the light of new evidence, optimized to minimize a measure of distance from the desired goal state and additional constraints, such as energy consumption. In principle, this can be done by testing all policies in parallel before taking any action, but this can be approximated with a more parsimonious scheme in which one or a few candidate policies are considered serially.<sup>73–75</sup> In this scheme, a to-be-tested policy (e.g., the policy having the highest *a priori* “score”) is sampled from the generative model and it generates predictions, which are evaluated and “kept in register” using external inputs (Fig. 3). To formulate a strong hypothesis, *a policy is sampled and evaluated within each theta cycle* and the theta sequence is read out of this process. The first elements of a theta sequence may be those that are kept in register using external stimuli (like in state estimation); this can be done with a predictive-coding style matching of top-down predictions (from the policy) and bottom-up sensations, such as cues or path integration signals in the hippocampal–entorhinal system. The successive elements in the theta sequence represent predictions propagated forward in time using the model, reflecting successive inferential updates of the plan (i.e., predictions about the next states),<sup>89</sup> they participate in the plan evaluation by engaging prefrontal or ventral striatal mechanisms that consider predicted states in relation to goals or reward expectancies.<sup>7,64</sup> According to this hypothesis, policy evaluation would be serial, and different policies may be evaluated within different (e.g., consecutive) theta cycles, as occurs in theta flickering or perhaps also during VTE behavior.<sup>36,90</sup>

This hypothesis assigns theta sequences a more proactive role compared to self-localization and short-term prediction<sup>35,91</sup> by assuming that the span



**Figure 3.** Schematic illustration of theta sequences within an inferential framework that evaluates (and uses) a candidate policy during action–perception cycles. One policy is evaluated for each theta cycle. Potentially, the sequence of states predicted by the policy spans an entire trajectory from start to goal location; however, only a part of it can be expressed within a single theta cycle (e.g., a portion of the plan, from the current location to a goal or another relevant location). Within each theta cycle, it is possible to evaluate the quality of the proximal sensory predictions of a policy (early in the theta cycle) and/or the relation between its distal aspects and current goals (later in the theta cycle). The former part roughly corresponds to state estimation and uses active sensing to temporally coordinate descending predictions and incoming stimuli, such as cues from the entorhinal cortex. The latter part taps goal-related information (or prior preferences) in areas such as the ventral striatum and the prefrontal cortex. This evaluation results in a “score” for the policy that increases (or decreases) the probability that it will be used for action selection or the next evaluation (in successive theta cycles).

of prediction is not limited *a priori* but can “stretch” (to cover, e.g., long distances) when necessary. In this perspective, what is decoded as a theta sequence may be the limited readout of a more sophisticated (and far-looking) inferential process that runs covertly to continuously evaluate and select policies—an implicit trajectory through the plan’s state space, which can be “read” sequentially by downstream regions.<sup>92</sup> The readout is limited, given that inference occurs within theta cycles, which allow for a small number of elements to be encoded (e.g., 5–10 future locations predicted under the current policy<sup>93</sup>). Nevertheless, the mechanism can be flexible and cover different portions of the plan depending on task demands. Theta sequences are usually short, possibly reflecting a current focus of the inference on short-term predictions; in some cases,

they can stretch to cover trajectories toward distal goals<sup>11</sup> or (when the goal is uncertain) to sequentially evaluate and select between alternative spatial plans, as in VTE.<sup>36</sup> Theta sequences may temporarily be centered behind or ahead of the animal,<sup>17</sup> suggesting that different parts of the plan may be inferred or updated (predicted or postdicted) over time depending on task demands—possibly relying on a theta-based segmentation of long streams of experience into meaningful subparts.<sup>17</sup>

Importantly, rhythms and other dynamical phenomena are part and parcel of the inference, in particular, for its temporal aspects and *predictive timing*<sup>94,95</sup> (i.e., predicting *when* something will occur as opposed to *what* will occur). The temporal alignment of top-down predictions and bottom-up sensory streams within each theta cycle may



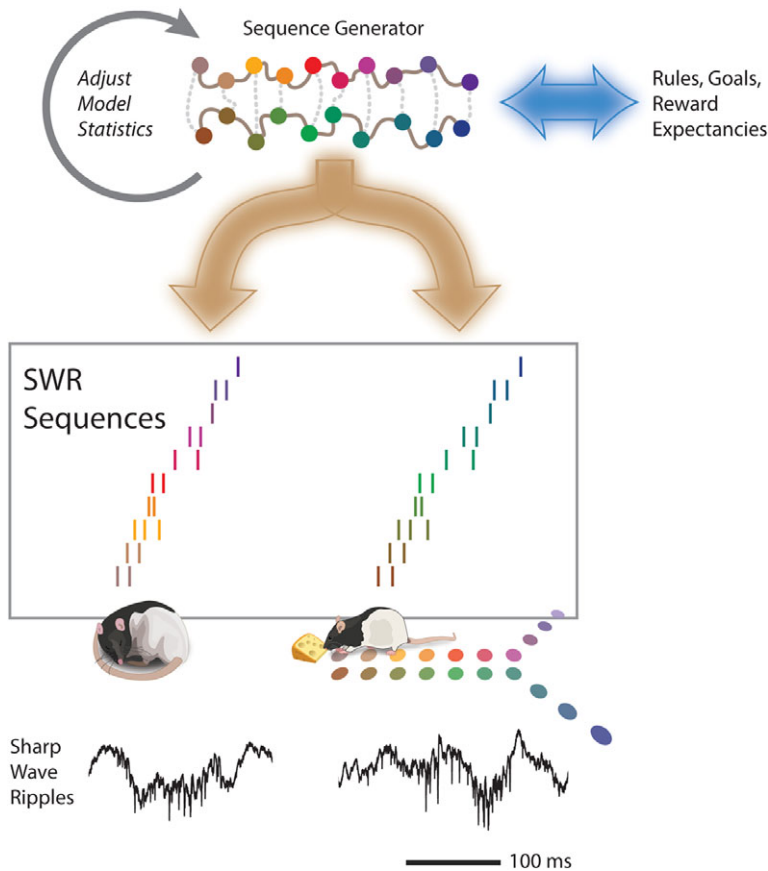
be coordinated by active sensing mechanisms.<sup>96</sup> Internally generated theta oscillations may become coordinated (or even phase locked) with external rhythmic sensory stimuli gathered by active sensing routines, such as whisking and sniffing, which also occur at theta frequency,<sup>97,98</sup> if the onset of these events “phase-resets” hippocampal oscillations<sup>99</sup>—coherent with the idea that action contributes to couple the intrinsic brain rhythms to the dynamics of external events.<sup>100</sup> This implies that theta phase resets may occur to bring specific theta phases in alignment with expected timing of external events<sup>101</sup>—a phenomenon that may be more visible in conditioning experiments (where cues are presented with predictable timing) than in typical spatial navigation tasks. The (rhythmic) entrainment of slow cortical oscillations to the stimulation rate would also enhance perceptual sensitivity,<sup>102</sup> with the theta cycle acting as a filter that increases the gain of signals collected at the right phase and suppresses the other (as noise). Finally, dynamical phenomena, such as synchronization between rhythms in different brain areas, may regulate the routing of information across different brain areas.<sup>92</sup> For example, the temporal alignment of hippocampal activity with VS reward signals during spatial navigation<sup>103,104</sup> (but also during replay,<sup>105</sup> see below) may index communication between these two areas, as if they jointly encode a generative model linking place and reward information.<sup>7,64,106</sup> Coupling between the hippocampus and the PFC may be beneficial instead when memory demands are high.<sup>107</sup> Thus, we suggest that theta sequences reflect ongoing predictions of a generative model synchronized to the action–perception cycle.

**SWR sequences.** Inferential processes may also underlie SWR sequences, the second dynamical mode of the hippocampus (Fig. 4). When the animal pauses during a task or sleeps, inference can be temporarily decoupled from the action–perception cycle and its constraints, such as the necessity to align predictions to external stimuli collected via active sensing routines. The ensuing spontaneous neuronal dynamics—many of which are detected as SWR sequences—can be interpreted as consecutive samples from a (prior) distribution over policies or trajectories encoded in the internal model, which are (relatively) unaffected by external sensory events, as in “mind wandering.”<sup>108</sup> Modeling stud-

ies have shown that, if left free from constraints, the relative frequency of sampled states or trajectories converges on a stationary distribution that is independent from the initial state of the system.<sup>109</sup> This would imply that measuring spontaneous activations should in the long run reveal a statistically optimal model (i.e., a model that is tuned to the statistics of the environment), as has been observed in other brain areas.<sup>110</sup>

However, several studies have shown that the content of SWRs can diverge in interesting ways from uniform distributions, suggesting that SWR content is not completely decoupled from task demands and/or experience. It is important to separate these studies according to the level of task engagement, starting with the distinction between *awake* and *sleep SWRs*. Early studies of sleep SWRs focused on sleep directly following behavior and found a clear contribution of recent experience.<sup>111–113</sup> The rate of SWR emission is increased during sleep that follows novel experience compared with familiar experience,<sup>114</sup> consistent with the generative model being updated with recently acquired information. Although a strict interpretation of the term “replay” implies that SWR content accurately reflects the statistics of recent experience, the efficient learning of generative models instead favors preferential replay of particularly informative or behaviorally relevant experience<sup>115,116</sup> circumventing the “real” statistics of the environment, so that behaviorally relevant events are over-resampled and thus preferentially encoded in the generative model. The dynamics (a large increase followed by a slow decrease) of both the rate of SWRs and what fraction of SWRs contains detectable sequences would then reflect meta-control of the resampling process.<sup>114,117,118</sup> In turn, the coupling of fast-ripple oscillations may regulate what information is transmitted to other brain areas, such as to communicate with the PFC for the formation of declarative memories,<sup>119</sup> to keep episodic memories in register with changing neocortical representations,<sup>120</sup> or to train a behavioral controller (offline) before a choice, thus saving resources during the decision-making process proper.<sup>121</sup> Nevertheless, these views of SWR content and function are compatible with a classical “consolidation” account in which recent experience is used to update a generative model.

Beyond recent experience, the generative model view suggests that other factors should contribute



**Figure 4.** Schematic illustration of SWR sequences within an inferential framework while the animal is decoupled from the action–perception cycle (e.g., while it sleeps). Here, SWR sequences correspond to resamples from the policy distribution, as coded in an internal model. In the absence of biases, the resampling would be from the prior policy distribution and can play roles in structuring the hippocampal internal model (see main text for explanation) or internal models in other brain areas, such as the prefrontal cortex. Other brain areas can bias this process. For example, prefrontal goal information can bias the process to preferentially resample policies that reach a known goal site, thus producing SWR sequences that are relevant for planning.

to sleep SWR content as well. These may include the need to keep the model coherent or self-consistent as learning progresses, ways to prevent overfitting, and preventing destructive interference (well-known ideas in machine learning and connectionist models<sup>56,122</sup>). For example, the risk of overfitting (i.e., the tuning of the generative model to noise or idiosyncratic features of experience that do not generalize) suggests a process to get rid of experiences that complicate the model without improving their predictive power—or model reduction (indeed, there have been suggestions that, under certain conditions, SWRs may elicit long-term depression<sup>123</sup>). Destructive interference refers to the possibility that updating one memory trace may adversely

affect overlapping traces, which may be prevented by active maintenance of relevant traces that have not been experienced for some time. Thus, SWR content during sleep may not be limited to recent experience content for building generative models, but may also include processes for maintaining, tuning, and optimizing such models in a manner independent of specific recent experience.

Awake SWR activity is more constrained by sensory input and task demands. Strong constraints are expected on tasks that require hippocampal function to perform, such as delayed alternation and some place-navigation tasks.<sup>12,14</sup> On such tasks, SWR content can reflect operations like recall and planning that directly support performance

of the task. Although SWR activity is temporarily decoupled from the action–perception cycle, SWR activity can still interact with other brain areas, such as the PFC,<sup>124,125</sup> from which it can receive, for instance, goal information.<sup>126</sup> In turn, this information acts as a constraint to the spontaneous dynamics, which then “drift away” from the (prior) stationary distribution to converge, for example, to a plan to the goal,<sup>12</sup> with the same inferential mechanisms described before. In other words, during performance of hippocampal-dependent tasks, awake SWR sequences may reflect inferential updates for planning in the same way we described for the first dynamical mode, except that it is temporarily decoupled from the action–perception cycle and only uses self-generated information (e.g., about goal locations). If this hypothesis holds, then awake SWR and theta sequences may support planning in a coordinated manner at two different time scales. Awake SWR sequences may be used to select an initial policy (to form a rough plan to a known goal location<sup>12,14,127</sup>) before taking action, and—if needed—theta sequences may evaluate (or revise) the policy during action–perception cycles as the animal approaches the goal site; see Ref. 127 for data consistent with such a cooperative view.

Not all awake SWR content necessarily supports immediate, upcoming behavior. When generative model output is not needed for task performance, such as may occur on tasks or task phases that do not depend on the hippocampus, SWR activity may uncouple more from task demands. In such situations, SWR sequences can include content expected from more offline generative model updating and maintenance, as discussed for sleep above. In addition, SWRs may be important to identify actions that are most informative in updating the generative model.<sup>128</sup> Thus, awake SWR content may emphasize recent experience following rewarding or unexpected outcomes,<sup>49,63,129</sup> upcoming constructive trajectories when the hippocampus is needed to perform the task,<sup>12,14</sup> and trajectories that do not lead to a rewarded goal but to the maximally informative outcome instead<sup>51,192</sup> (see below). In sum, SWR sequences may reflect inferential processes that are temporarily disengaged from the demands of perception–action loops (and hence more spontaneous compared with theta sequences) but nevertheless influenced by different sources of self-generated information,

thus potentially playing multiple roles in planning, memory consolidation, and the generation of novel (imaginary) experiences.

Interestingly, a recent study<sup>46</sup> of the primate hippocampus during visual exploration reports frequent SWRs. This suggests that, just as in rodents exploring novel mazes, there may be task characteristics that can cause the primate hippocampus to enter into an SWR-dominated state. The mechanisms regulating the selection of theta to SWR-dominated modes, and the transitions between them, are incompletely known. While we have emphasized their differences, our computational framework suggests various ways theta and SWR sequences may work synergistically in the service of online behavior. For example, some tasks may require theta and SWR sequences to play complementary functions, such as short-term (theta) and long-term prediction and planning (SWR) or, alternatively, encoding (theta) and simulation (SWR), at fine time scales.<sup>14</sup> These tasks may thus require frequent transitions between dynamical modes and engage a cyclical relationship between theta and SWR sequences, with rapid consecutive cycles of engagement (theta) and detachment (SWR) from the action–perception cycle while the animal is still engaged in the task. Understanding the interplay between theta and SWR sequences remains an exciting avenue for future research.

### *Generative models*

The notion of generative models has been widely used in the context of brain computations and especially cortical processing;<sup>66,67</sup> thus, one may ask what is special about generative models in the hippocampus. An observation that dates back to Tolman<sup>130</sup> is that the hippocampus may support the rapid encoding of a *cognitive map* (a notion that is not limited to spatial maps, but encompasses structured information in nonspatial domains<sup>41</sup>). From a computational perspective, there are at least four (interconnected) aspects of this notion, which we discuss in order.

**Encoding arbitrary statistical relations.** One would expect generative models to be organized differently for, say, perceptual processing (e.g., predicting a ball trajectory) or the encoding of arbitrary associations (e.g., digits of a phone number), because the environmental statistics underlying these phenomena are different. In perceptual

processing, predictions can be often considered to be *extrapolations*, at least in the sense that changes from one visual scene to the next one are small. From a computational perspective, extrapolation may benefit from “averaging” and generalizing from numerous previous experiences, and modern machine learning algorithms excel at this using big data.<sup>131</sup> Rather, the internal model of the hippocampus may excel at a fast form of encoding of more arbitrary, memory-based predictions (e.g., for phone numbers). These arbitrary sequences have a more episodic nature and may require keeping different events separated, not merging them to form (semantic) averages, because generalizing from “average” representations may be ineffective if one has to predict a specific future episode. The coding of place information may not eschew this rule, as manifested by the strong pattern separation of place fields in dentate gyrus and CA3 cell assemblies.<sup>132</sup> Orthogonal representations of this kind may provide a sufficiently large basis to learn separated and specific representations for a huge variety of environments and events, even beyond spatial codes. The idea that hippocampal generative models are specialized for the fast encoding (and recall) of arbitrary sequences would explain the well-recognized importance of this brain structure for episodic memory and one-trial learning,<sup>133</sup> even outside spatial domains. As such, it has been suggested that the spatial and mnemonic functions are manifestations of a more general role of the hippocampus in representing the relationships between objects and events in both space and time.<sup>134</sup>

**Encoding sequences, not just associations.** Given the frequent decoding of sequences (rather than just isolated place cells or pairs) in the hippocampus, one may speculate that the hippocampal internal model may be preconfigured to learn sequences or transitions—or, in other words, that sequences (not just single elements like place cells) may be first-order objects in hippocampal coding. Sequential coding is key across spatial navigation and episodic memory studies. In spatial navigation, sequential organization may stem from the fact that the animal trajectories trace sequences of place cells in space; but the fact that time-compressed sequences are also expressed in internally generated activity suggests that this may be an important dimension of neural coding, or perhaps an adaptation of neural cod-

ing to the sequential nature of spatial navigation. In the episodic memory domain, it has been often argued that episodes may be coded, organized, and recalled as sequences having a limited number of elements.<sup>135,136</sup> This perspective may help in reconciling two main ideas in the literature regarding theta sequences, which emphasize their role in the storage of episodic-like memories<sup>137</sup> or in online prediction.<sup>35</sup> These ideas may be connected if one considers that they focus on two different phases (encoding or storage and retrieval or recall, respectively), but these two processes may operate synergistically. Theta sequences may be primarily configured to develop episodic-like memories; but, once they are developed, these memories may be flexibly used in the service of multiple cognitive processes, such as navigation and decision making.<sup>138</sup>

#### **Sequences as dynamical templates for experience.**

Several aspects of hippocampal processing—fast learning of arbitrary elements, fast remapping across episodes (which implies that previous associations between place cells are no longer valid), and the possibility to express sequences for never-experienced routes (preplay)—all suggest that hippocampal sequences may be largely preconfigured in the model as “dynamical templates”<sup>139</sup> to organize experience, possibly from a limited set of building blocks.<sup>140</sup> Dynamical templates may be preconfigured to represent the temporal structure of multi-item events and be recruited immediately without (at least initially) committing to the “content” of each item in the sequence; but they also permit to quickly “bind” their items to the sensory stimuli experienced one after the other.<sup>141</sup> In this perspective, the coding of a sequence of place cells would not be due to synaptic learning between consecutive place cells but largely to the binding of each item of a preconfigured sequence to a place cell (although stimuli may fine-tune sequential information<sup>139</sup>). This may be possible using computational schemes that form temporal sequences through a state space of attractors or attractor-like states<sup>142–144</sup> or transient trajectories<sup>109,142</sup> possibly corresponding to recurrent computations in CA3/CA1 circuits (see also early models of sequential processing such as Ref. 145). In this perspective, sequences would not be fully represented in the hippocampus, but in the coupling of the hippocampus and other areas, such as the EC. One proposal is that the

**Table 1. Predictions**

<i>Theta sequences</i>
<ul style="list-style-type: none"> <li>In common with other models that stress a probabilistic inference interpretation, we expect that theta sequences contain signals coding for precision (or its inverse, uncertainty). Preliminary support for this idea comes from the observation that passive transport disrupts theta sequences<sup>185</sup> and that experience with a novel environment rapidly improves theta sequence quality.<sup>13</sup> Virtual reality environments, in which the mapping from actions to sensory changes, as well as the transitions in the sensory world itself, can be systematically manipulated and offer a promising avenue for identifying such signals. As pointed out by Penny <i>et al.</i>,<sup>79</sup> such a signal would also allow for optimal multisensory integration.</li> <li>The proposal that theta sequences encode precision implies that the impact of prediction errors will depend on the degree of precision at the time of the error. Episodic memory traces are only encoded in the hippocampus when a sufficiently large prediction error occurs and therefore requires a certain amount of precision in the prediction. Consistent with this idea, animals require a minimum amount of experience in a novel context before they will associate that context with an unpredicted shock.<sup>186</sup></li> <li>Manipulations that disrupt theta sequences affect both “encoding” and “retrieval” functions of hippocampal-dependent memory. For instance, Robbe <i>et al.</i><sup>187</sup> found that systemic cannabinoids reduced animals’ performance on a delayed T-maze alternation task; if such disruption would be applied specifically during a sample (encoding) lap and subsequent retrieval laps (using, for instance, optogenetic manipulation of inputs from the medial septum), impairments would be observed for both.</li> </ul>
<i>SWR sequences</i>
<ul style="list-style-type: none"> <li>The content of SWRs emitted during sleep will reflect experience according to (1) its recency and (2) its informativeness, consistent with a consolidation process in which experience stored in the hippocampus can be resampled to update a cognitive map and cortical knowledge structures. In addition, sleep SWRs contain content that reflects a role in other aspects of generative model learning, such as improving self-consistency, pruning of model features unlikely to generalize, and prevention of catastrophic interference between memory traces. As an example of self-consistency, suppose that an animal has learned that A is connected to B and B is connected to C. Unexpected reward is that received at C after the animal was placed there, causing it to be associated with a reward value. Replay of the B–C connection could be one way in which B could also acquire (discounted) value, as can be shown behaviorally.<sup>188,189</sup></li> <li>The content of SWRs emitted during waking will reflect those actions that lead to maximally informative outcomes. Consistent with Gupta <i>et al.</i>,<sup>51</sup> this means that, when an animal is rewarded for choosing only the left arm on a T-maze, the right arm is the most informative choice and may be more likely to be replayed. In the maze version of spontaneous object recognition tasks, in which animals are biased to seek out an object they have not seen recently,<sup>128,190</sup> SWR content will be biased toward the most informative object. The sensitivity of SWRs to information content can be probed by experiments that manipulate the informativeness or epistemic value<sup>74,75</sup> of places or cues independent of their reward (prediction) value.</li> </ul>
<i>Overall</i>
<ul style="list-style-type: none"> <li>If both theta and SWR sequences stem from the same generative model, manipulations that affect it (e.g., in CA3 or the medial prefrontal cortex) should affect the content of both kinds of sequences in a coherent manner.</li> <li>If theta and SWR sequences are parts of a coordinated planning mechanism, suppressing SWR should enhance theta sequences (and vice versa); see Ref. 127 for preliminary evidence.</li> <li>If IGSs have a truly causal role in behavior and cognition, rather than (for example) being only functional to the readout of cognitive processing in other brain areas such as the PFC, it should be able to target specific cognitive processes (e.g., memory consolidation and planning) by manipulating them.</li> </ul>

coding of place cell sequences can be part of a factorized representation in which the hippocampus supports temporal sequencing (*when* computations) while other areas, such as the EC, code for the content (the *what*) of sequences.<sup>141</sup> In this perspective, the hippocampus would support integrative (what–where–when or what–where–which<sup>146</sup>) functions by acting as a hub (and sequential processor) of information distributed in various brain areas. Interestingly, retrieving a single element (e.g.,

the fifth element) from a sequence requires “replaying” the sequence until the element is tapped, as there are no stand-alone indexes or “pointers” to isolated items of the sequence. This may explain why experience needs to be replayed sequentially for memory retrieval or consolidation.<sup>56,121</sup>

**Structural priors to form maps.** Learning sequences (e.g., for trajectories) does not automatically produce good maps. Rather, one needs also to

discover that, for example, the starting point of two sequences is the same as the end point of another sequence. From a statistical perspective, if the hippocampus has to rapidly form cognitive maps, as in the original notion of Tolman,<sup>130</sup> it would greatly benefit from strong *structural priors*,<sup>147</sup> which reflect (formal) aspects about inputs, such as the fact that sequences of trajectories in the same environment should not have gaps (unless there is a wall or obstacle) or that paths can usually be traversed in both directions. These priors would enable the filling of gaps in experience, such as forming a complete spatial map from a limited set of temporally separated episodes. In this perspective, resampling sequences during SWR activity supports the generation of novel experiences that recombine other ones in constructive ways that are consistent with the priors, which would explain why SWR sequences can reflect novel trajectories or shortcuts.<sup>51</sup> This would resemble the usual inferential process of predictive coding if one has to minimize differences (or prediction errors) between the hypotheses expressed in the priors (e.g., that sequences must be connected) and current evidence (e.g., that recently experienced sequences are not connected). When this is not possible (e.g., when the experienced trajectories cannot be merged), the most parsimonious explanation from a statistical viewpoint would be that these trajectories belong to two (or more) distinct environments, and this in turn may help the animal in building multiple maps. Furthermore, resampling SWR sequences may help restructuring the internal model (e.g., removing redundant parameters, biasing its content,<sup>116,148</sup> and forming state space representations that organize experience within a spatial context<sup>69,149</sup>).

**The role of the hippocampus in information seeking.** An important feature of generative models is that, to the extent that they are adaptive, acquiring information for the purpose of constructing effective models or reducing uncertainty before a choice has behavioral utility. Many organisms, including humans, nonhuman primates, and rodents, are willing to work to obtain information even in the absence of association with reward (“curiosity”<sup>150,151</sup>), thereby following an epistemic drive.<sup>74</sup> Johnson *et al.*<sup>128</sup> pointed out that, in rodents, active information seeking may depend on hippocampal function; specifically, the hip-

poampus is required for inferring which action is maximally informative. Extending their proposal to hippocampal sequences, it is striking to note that the content of awake SWR sequences may reflect such a function, including counterintuitive cases such as that shown in Ref. 51, where animals preferentially replayed the opposite side of a T-maze. Similarly, in the case of an aversive task,<sup>192</sup> SWR sequences travel to the location to be avoided, which may be significantly more relevant to the animal than the alternative, behaviorally neutral path taken. In some cases, the trajectory toward a (rewarded) goal may also be the most informative,<sup>12,14,63</sup> such that the planning and information-seeking accounts overlap, but the two can be explicitly dissociated; see Table 1 for specific predictions.

## Conclusions

Hippocampal IGSs have been extensively studied during goal-directed navigation (in rodents), and they may offer a vantage point to understand how animal brains temporarily detach from the here and now to engage in spontaneous or internally directed actions, which have adaptive value (e.g., for memory consolidation or planning). We proposed that stimulus-tied and internally generated modes of activity may engage the same sequential “inferential machine” (on the basis of generative models that the animal maintains of its environment) to support overt goal-directed action and cover mental activities, including forms of future-oriented and prospective (but also retrospective) cognition that share resemblances with the more sophisticated forms—imagination and mental time travel—studied in humans. Although we have focused on hippocampal IGSs, sequential neuronal activity has been reported also in other brain areas, such as the parietal cortex,<sup>152</sup> the PFC,<sup>153,154</sup> the medial EC,<sup>155</sup> and the striatum<sup>156</sup> of rodents and sensory (olfactory) areas of insects,<sup>157,158</sup> which may speak for the generality of the inferential principles described here.

Several aspects of our proposals remain to be fully tested empirically. One fundamental question regards the validity of the inferential scheme proposed here. The computations we have described can be formalized using probabilistic computations in active inference,<sup>74,75</sup> model-based reinforcement learning,<sup>159</sup> or other methods, such as Monte Carlo tree search,<sup>160</sup> planning as inference,<sup>76</sup>

and generative (deep) networks.<sup>131,161</sup> At the neuronal level, statistical inference has viable biological implementations and can be performed via variational inference<sup>66</sup> or by iteratively sampling individual elements or sequences from the internal model (e.g., using sampling-based probabilistic inference in recurrent networks, similar to Markov chain Monte Carlo methods in statistics.<sup>109,162–164</sup>) For example, theta sequences emerge under the variational scheme of *active inference* if one considers that at each (theta) cycle successive elements of the policy state space are inferred, and the inference proceeds more rapidly for locations that the animal is approaching,<sup>88</sup> which in turn produces a “phase precession” within this inferential scheme. These and other formal schemes may be used to test empirical predictions on how rodents solve challenging navigation tasks or even how humans solve abstract tasks.<sup>165–167</sup> Compared with other computational proposals that also discuss hippocampal function in relation to generative models,<sup>7,75,79,87,88</sup> we have stressed the idea that hippocampal coding and processing may be organized around sequences. Hippocampal generative models may be preconfigured for sequential processing and the rapid encoding of arbitrary sequences of events<sup>139</sup> in ways that current machine learning techniques fail to model. Furthermore, we have stressed the importance of temporal dynamics of probabilistic inference (and predictive timing) to explain various functions that we have associated with theta and SWR sequences. Two examples are active sensing and the formation of declarative memories, which require a fine temporal coordination within the action–perception cycle (active sensing) or during the interaction with other brain areas, such as the PFC (memory consolidation). The validity of the inferential scheme proposed here to address sequential processing and temporal dynamics remains to be fully tested.

Yet another open question regards the specific roles and interplay of theta and SWR sequences during behavior and cognition. We have proposed that they may stem from the same generative process and that their different dynamical signatures (and functions) may depend on the different dynamical modes under which they operate, the interaction with other brain areas, and the input of the generative model. The extent to which the disruption of theta or SWR sequences affects behavior may thus depend on task demands—for example, whether the

animal is facing a decision on the basis of external, episodic input versus prior experiences—or even disengaged from active tasks (e.g., sleeping)—because these tasks would imply different constraints and inputs for the generative model. Furthermore, several lines of evidence reviewed above suggest that theta and SWR sequences may interact (e.g., jointly support planning function or encode-then-simulate function) and also be nested within each other at fine time scales (e.g., SWR sequences occurring between consecutive action–perception cycles), but these interactions and their behavioral consequences remain to be completely charted.

Another crucial aspect of this proposal is assessing the width of the brain networks implied in spatial navigation and planning. Here, we have focused on a restricted brain network that includes only the hippocampus, the VS, and the PFC. Clearly, the brain networks implied in goal-directed navigation are much wider, and some fundamental aspects of their functioning and interactions remain unclear (e.g., what are the specific roles of different brain areas and which aspects of sequential processing are intrinsic to the hippocampus (e.g., through recurrent connections in CA3) and which originate in other areas, such as the PFC<sup>168</sup>). Simultaneous recordings from multiple brain areas and disruption of brain activity (e.g., through optogenetics) may help in shedding light on these and other fundamental questions (see Table 1 for a list of predictions stemming from the proposed framework).

A corollary of our proposal is that, although the cognitive requirements of a specific cognitive task (e.g., goal-directed navigation) are usually compartmentalized into separate processes (e.g., active sensory processing, attention, state and context estimation, working memory of goal location, route planning, and spatial decisions), there is a possibility that a single inferential mechanism simultaneously implements many or all of these processes, as well as others that we have not discussed (e.g., context estimation<sup>79,169</sup>). The variety of functions attributed to SWRs can be reconciled if one considers that they may stem from a common generative model, which can flexibly interact with other brain areas in task-specific ways. Similarly, sequences of neuronal activity found in the rodent parietal cortex may simultaneously play multiple roles in decision making, working memory, and the coordination of sensory processing.<sup>152</sup>

Even more intriguingly, our proposal suggests that the brain may implement distinct cognitive processes by switching its dynamical mode of operation and the coupling with other brain areas. This would imply that the mechanisms supporting detached cognition may not be fundamentally different from those operating during the action–perception cycle; the key differences may lie in the ways the same underlying computations (e.g., inference using a generative model) are triggered and which information they can access. One can also speculate that the latter (detached) dynamics result from the *internalization* of processes that originally supported situated action—or that cognition stems from action.<sup>3,8,170,171</sup>

If this perspective is correct, then evidence on hippocampal IGSs can be used to expose fundamental mechanisms permitting humans to “escape from the present” and afford prospective (future-oriented) or retrospective (past-oriented) forms of cognition, despite the large conceptual and methodological distance between the fields of rodent spatial navigation and human detached cognition. Our proposal implies that memory recall and future-oriented processes, such as planning and imagination, are based on the same inferential mechanisms (e.g., resampling of a policy, covertly evaluating it) and tap into internal generative models in the hippocampus (but also in other brain areas) in the same way. In short, episodic memory (or prediction) is not just recall (or extrapolation) but probabilistic inference using a generative model. The involvement of a common core (inferential) of mechanisms across multiple domains of detached and higher cognition—such as episodic memory, imagination, counterfactual thinking, mind wandering, prospection, and “time travel” into the past and the future—may explain why these abilities recruit shared neuronal circuits.<sup>2,5,172</sup>

In this perspective, mind wandering may be associated with spontaneous thought<sup>173,174</sup> or the sampling from (the prior distribution of) an internal generative model that is minimally constrained from external events, analogous to SWR sequences. This process may have adaptive roles in memory consolidation or model reduction. Furthermore, as for SWRs, the same potentially unbiased (mind wandering) process can be engaged in various other ways. It can become oriented toward future (goal) states, as in the case of planning, when the resam-

pling process is biased by goal information, or toward the past, when a cue engages episodic memory retrieval—as in Proust’s *madeleine* example. In other words, the common core (inferential) mechanisms remain the same across multiple domains of detached cognition, while the orientation (toward the future or past) and the content (e.g., more episodic or semantic) depends on the situation or the ways the resampling is biased.

It is worth noting that the inferential processes described here in relation to theta and SWR sequences—but also possibly human future-oriented cognition—are not (or not necessarily) conscious. The internal inferential updates of generative models that implement planning or deliberation can operate at a time scale below that of cognition, in the same way that the time scale of visual saccades is shorter than the time scale of perception. For example, in the computational scheme for planning introduced in the section on inference, evaluating each single policy requires multiple inferential updates, which then need to operate at a much faster time scale compared with the action–perception cycle.<sup>73–75,193,194</sup> It thus remains to be established under which conditions one can have conscious access to inferential processes that implement planning and mind wandering.

The hippocampus has been consistently reported as an important hub in detached cognitive processes<sup>175–180</sup> and the brain default network<sup>181</sup>—a dynamical mode that epitomizes the “detachment” from action–perception cycles in humans. This has raised questions about the necessity for detached operations to be episodic or autobiographical, requiring projecting oneself into the past or future, rather than more mundanely predicting them.<sup>133</sup> From the computational viewpoint discussed here, the engagement of the hippocampus is not mandatory, but it is beneficial when it is necessary to construct events (in the past or the future) that depend on specific circumstances or when averaging across all evidence (as semantic models do) would not be accurate.<sup>182</sup>

Still, this constructive process would require “meshing” content in semantically coherent ways and hence would require a combination of semantic and episodic processing. We still lack a detailed understanding of such mixed semantic–episodic operations. Perhaps understanding how SWR sequences mesh episodes coherently (e.g.,



adjacent trajectories but not separated ones<sup>51</sup>) may shed light on how human episodic constructive memory works (e.g., meshing memories of oneself at in the grandmother's house and of oneself at a past children's party to imagine how a children's party at the grandmother's house would be). Progress in the field may come from studies that directly probe sequential activity resembling SWR or theta sequences during human cognitive processing,<sup>183,184</sup> possibly designed to test (and disentangle) alternative computational models.

## Acknowledgments

The authors want to thank David Redish, Karl Friston, Paul Verschure, Cyriel Pennartz, Thomas Akam, and the other participants in the workshop on "Internally generated sequences in the hippocampus" held in Ariccia, Rome, on September 27–30, 2016, as well as two anonymous reviewers, for useful comments and suggestions. G.P., M.V.D.M., and C.K. gratefully acknowledge the support of HFSP (Young Investigator Grant RGY0088/2014). CK additionally acknowledges support from the NSF IOS-1550994 and CBET-1351692.

## Competing interests

The authors declare no competing interests.

## References

- Corballis, M.C. 2014. Mental time travel: how the mind escapes from the present. *Cosmology* **18**: 139–145.
- Hassabis, D. & E.A. Maguire. 2009. The construction system of the brain. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **364**: 1263–1271.
- Pezzulo, G. & C. Castelfranchi. 2009. Thinking as the control of imagination: a conceptual framework for goal-directed systems. *Psychol. Res.* **73**: 559–577.
- Pezzulo, G. & P. Cisek. 2016. Navigating the affordance landscape: feedback control as a process model of behavior and cognition. *Trends Cogn. Sci.* **20**: 414–424.
- Schacter, D.L., D.R. Addis, D. Hassabis, *et al.* 2012. The future of memory: remembering, imagining, and the brain. *Neuron* **76**: 677–694.
- Andrews-Hanna, J.R., J. Smallwood & R.N. Spreng. 2014. The default network and self-generated thought: component processes, dynamic control, and clinical relevance. *Ann. N.Y. Acad. Sci.* **1316**: 29–52.
- Pezzulo, G., M.A.A. van der Meer, C.S. Lansink & C.M.A. Pennartz. 2014. Internally generated sequences in learning and executing goal-directed behavior. *Trends Cogn. Sci.* **18**: 647–657.
- Buzsáki, G., A. Peyrache & J. Kubie. 2014. Emergence of cognition from action. *Cold Spring Harb. Symp. Quant. Biol.* **79**: 41–50.
- Dragoi, G. & S. Tonegawa. 2011. Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature* **469**: 397–401.
- Wikenheiser, A.M. & A.D. Redish. 2013. The balance of forward and backward hippocampal sequences shifts across behavioral states. *Hippocampus* **23**: 22–29.
- Wikenheiser, A.M. & A.D. Redish. 2015. Hippocampal theta sequences reflect current goals. *Nat. Neurosci.* **18**: 289–294.
- Pfeiffer, B.E. & D.J. Foster. 2013. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* **497**: 74–79.
- Feng, T., D. Silva & D.J. Foster. 2015. Dissociation between the experience-dependent development of hippocampal theta sequences and single-trial phase precession. *J. Neurosci.* **35**: 4890–4902.
- Jadhav, S.P., C. Kemere, P.W. German & L.M. Frank. 2012. Awake hippocampal sharp-wave ripples support spatial memory. *Science* **336**: 1454–1458.
- Suh, J., D.J. Foster, H. Davoudi, *et al.* 2013. Impaired hippocampal ripple-associated replay in a mouse model of schizophrenia. *Neuron* **80**: 484–493.
- Ólafsdóttir, H.F., F. Carpenter & C. Barry. 2016. Coordinated grid and place cell replay during rest. *Nat. Neurosci.* **19**: 792–794.
- Gupta, A.S., M.A. van der Meer, D.S. Touretzky & A.D. Redish. 2012. Segmentation of spatial experience by hippocampal theta sequences. *Nat. Neurosci.* **15**: 1032–1039.
- Wu, X. & D.J. Foster. 2014. Hippocampal replay captures the unique topological structure of a novel environment. *J. Neurosci.* **34**: 6459–6469.
- Pfeiffer, B.E. & D.J. Foster. 2015. Autoassociative dynamics in the generation of sequences of hippocampal place cells. *Science* **349**: 180–183.
- Wilson, M. & B.L. McNaughton. 1993. Dynamics of the hippocampal ensemble code for space. *Science* **261**: 1055–1058.
- Zhang, K., I. Ginzburg, B.L. McNaughton & T.J. Sejnowski. 1998. Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *J. Neurophysiol.* **79**: 1017–1044.
- Wood, E.R., P.A. Dudchenko & H. Eichenbaum. 1999. The global record of memory in hippocampal neuronal activity. *Nature* **397**: 613–616.
- Allen, K., J.N.P. Rawlins, D.M. Bannerman & J. Csicsvari. 2012. Hippocampal place cells can encode multiple trial-dependent features through rate remapping. *J. Neurosci.* **32**: 14752–14766.
- Pastalkova, E., V. Itskov, A. Amarasingham & G. Buzsáki. 2008. Internally generated cell assembly sequences in the rat hippocampus. *Science* **321**: 1322–1327.
- MacDonald, C.J., K.Q. Lepage, U.T. Eden & H. Eichenbaum. 2011. Hippocampal "time cells" bridge the gap in memory for discontinuous events. *Neuron* **71**: 737–749.
- Morris, R.G.M. 2006. Elements of a neurobiological theory of hippocampal function: the role of synaptic plasticity,

- synaptic tagging and schemas. *Eur. J. Neurosci.* **23**: 2829–2846.
27. Dragoi, G. & G. Buzsáki. 2006. Temporal encoding of place sequences by hippocampal cell assemblies. *Neuron* **50**: 145–157.
  28. Johnson, A., J.C. Jackson & A.D. Redish. 2009. Measuring distributed properties of neural representations beyond the decoding of local variables: implications for cognition. In *Information Processing by Neuronal Populations*. C. Hölscher & M. Munk, Eds.: 95–119. Cambridge University Press.
  29. Maurer, A.P., S.N. Burke, P. Lipa, *et al.* 2012. Greater running speeds result in altered hippocampal phase sequence dynamics. *Hippocampus* **22**: 737–747.
  30. Foster, D.J. & M.A. Wilson. 2007. Hippocampal theta sequences. *Hippocampus* **17**: 1093–1099.
  31. Zheng, C., K.W. Bieri, Y.-T. Hsiao & L.L. Colgin. 2016. Spatial sequence coding differs during slow and fast gamma rhythms in the hippocampus. *Neuron* **89**: 398–408.
  32. Skaggs, W.E. & B.L. McNaughton. 1996. Theta phase precession in hippocampal. *Hippocampus* **6**: 149–172.
  33. Burgess, N. & J. O’Keefe. 2011. Models of place and grid cell firing and theta rhythmicity. *Curr. Opin. Neurobiol.* **21**: 734–744.
  34. Chadwick, A., M. van Rossum & M. Nolan. 2015. Modelling phase precession in the hippocampus. *BMC Neurosci.* **16**(Suppl. 1): O18.
  35. Lisman, J. & A.D. Redish. 2009. Prediction, sequences and the hippocampus. *Philos. Trans. R. Soc. B Biol. Sci.* **364**: 1193–1201.
  36. Johnson, A. & A.D. Redish. 2007. Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J. Neurosci.* **27**: 12176–12189.
  37. Robbe, D. & G. Buzsáki. 2009. Alteration of theta timescale dynamics of hippocampal place cells by a cannabinoid is associated with memory impairment. *J. Neurosci.* **29**: 12597–12605.
  38. Wang, Y., S. Romani, B. Lustig, *et al.* 2015. Theta sequences are essential for internally generated hippocampal firing fields. *Nat. Neurosci.* **18**: 282–288.
  39. Hasselmo, M.E., C. Bodelón & B.P. Wyble. 2002. A proposed function for hippocampal theta rhythm: separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Comput.* **14**: 793–817.
  40. Vertes, R.P. 2005. Hippocampal theta rhythm: a tag for short-term memory. *Hippocampus* **15**: 923–935.
  41. Buzsáki, G. & E.I. Moser. 2013. Memory, navigation and theta rhythm in the hippocampal–entorhinal system. *Nat. Neurosci.* **16**: 130–138.
  42. Buzsáki, G. 2015. Hippocampal sharp wave-ripple: a cognitive biomarker for episodic memory and planning. *Hippocampus* **25**: 1073–1188.
  43. O’Neill, J., T. Senior & J. Csicsvari. 2006. Place-selective firing of CA1 pyramidal cells during sharp wave/ripple network patterns in exploratory behavior. *Neuron* **49**: 143–155.
  44. Kemere, C., M.F. Carr, M.P. Karlsson & L.M. Frank. 2013. Rapid and continuous modulation of hippocampal network state during exploration of new places. *PLoS One* **8**: e73114.
  45. Leonard, T.K. & K.L. Hoffman. 2017. Sharp-wave ripples in primates are enhanced near remembered visual objects. *Curr. Biol.* **27**: 257–262.
  46. Leonard, T.K. *et al.* 2015. Sharp wave ripples during visual exploration in the primate hippocampus. *J. Neurosci.* **35**: 14771–14782.
  47. Davidson, T.J., F. Kloosterman & M.A. Wilson. 2009. Hippocampal replay of extended experience. *Neuron* **63**: 497–507.
  48. Diba, K. & G. Buzsáki. 2007. Forward and reverse hippocampal place-cell sequences during ripples. *Nat. Neurosci.* **10**: 1241–1242.
  49. Ambrose, R.E., B.E. Pfeiffer & D.J. Foster. 2016. Reverse replay of hippocampal place cells is uniquely modulated by changing reward. *Neuron* **91**: 1124–1136.
  50. Foster, D.J. & M.A. Wilson. 2006. Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* **440**: 680–683.
  51. Gupta, A.S., M.A.A. van der Meer, D.S. Touretzky & A.D. Redish. 2010. Hippocampal replay is not a simple function of experience. *Neuron* **65**: 695–705.
  52. Karlsson, M.P. & L.M. Frank. 2009. Awake replay of remote experiences in the hippocampus. *Nat. Neurosci.* **12**: 913–918.
  53. Ólafsdóttir, H.F., C. Barry, A.B. Saleem, *et al.* 2015. Hippocampal place cells construct reward related sequences through unexplored space. *eLife* **4**: e06063.
  54. Grosmark, A.D. & G. Buzsáki. 2016. Diversity in neural firing dynamics supports both rigid and learned hippocampal sequences. *Science* **351**: 1440–1443.
  55. Silva, D., T. Feng & D.J. Foster. 2015. Trajectory events across hippocampal place cells require previous experience. *Nat. Neurosci.* **18**: 1772–1779.
  56. McClelland, J.L., B.L. McNaughton & R.C. O’Reilly. 1995. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol. Rev.* **102**: 419–457.
  57. Girardeau, G., K. Benchenane, S.I. Wiener, *et al.* 2009. Selective suppression of hippocampal ripples impairs spatial memory. *Nat. Neurosci.* **12**: 1222–1223.
  58. Ego-Stengel, V. & M.A. Wilson. 2010. Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus* **20**: 1–10.
  59. Sadowski, J.H.L.P., M.W. Jones & J.R. Mellor. 2016. Sharp-wave ripples orchestrate the induction of synaptic plasticity during reactivation of place cell firing patterns in the hippocampus. *Cell Rep.* **14**: 1916–1929.
  60. Pennartz, C.M., E. Lee, J. Verheul, *et al.* 2004. The ventral striatum in off-line processing: ensemble reactivation during sleep and modulation by hippocampal ripples. *J. Neurosci.* **24**: 6446–6456.
  61. Ji, D. & M.A. Wilson. 2007. Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nat. Neurosci.* **10**: 100–107.
  62. Shin, J.D. & S.P. Jadhav. 2016. Multiple modes of hippocampal–prefrontal interactions in memory-guided behavior. *Curr. Opin. Neurobiol.* **40**: 161–169.

63. Singer, A.C., M.F. Carr, M.P. Karlsson & L.M. Frank. 2013. Hippocampal SWR activity predicts correct decisions during the initial learning of an alternation task. *Neuron* **77**: 1163–1173.
64. van der Meer, M., Z. Kurth-Nelson & A.D. Redish. 2012. Information processing in decision-making systems. *Neuroscientist* **18**: 342–359.
65. Carr, M.F., S.P. Jadhav & L.M. Frank. 2011. Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nat. Neurosci.* **14**: 147–153.
66. Friston, K. 2010. The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.* **11**: 127–138.
67. Doya, K., S. Ishii, A. Pouget & R.P.N. Rao, Eds. 2007. *Bayesian Brain: Probabilistic Approaches to Neural Coding*. 1st ed. The MIT Press.
68. Buckner, R.L. 2010. The role of the hippocampus in prediction and imagination. *Annu. Rev. Psychol.* **61**: 27–48, C1–C8.
69. Friston, K. 2008. Hierarchical models in the brain. *PLoS Comput. Biol.* **4**: e1000211.
70. Lee, T.S. & D. Mumford. 2003. Hierarchical bayesian inference in the visual cortex. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **20**: 1434–1448.
71. Rao, R.P. & D.H. Ballard. 1999. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* **2**: 79–87.
72. Friston, K. 2003. Learning and inference in the brain. *Neural Netw.* **16**: 1325–1352.
73. Friston, K., S. Samothrakis & R. Montague. 2012. Active inference and agency: optimal control without cost functions. *Biol. Cybern.* **106**: 523–541.
74. Friston, K., F. Rigoli, D. Ognibene, *et al.* 2015. Active inference and epistemic value. *Cogn. Neurosci.* **6**: 187–214.
75. Pezzulo, G., E. Cartoni, F. Rigoli, *et al.* 2016. Active inference, epistemic value, and vicarious trial and error. *Learn. Mem.* **23**: 322–338.
76. Botvinick, M. & M. Toussaint. 2012. Planning as inference. *Trends Cogn. Sci.* **16**: 485–488.
77. Solway, A. & M.M. Botvinick. 2012. Goal-directed decision making as probabilistic inference: a computational framework and potential neural correlates. *Psychol. Rev.* **119**: 120–154.
78. Penny, W. 2014. Simultaneous localisation and planning. In *4th International Workshop on Cognitive Information Processing (CIP)* Copenhagen, Denmark. 1–6.
79. Penny, W.D., P. Zeidman & N. Burgess. 2013. Forward and backward inference in spatial cognition. *PLoS Comput. Biol.* **9**: e1003383.
80. Pezzulo, G., F. Rigoli & F. Chersi. 2013. The mixed instrumental controller: using value of information to combine habitual choice and mental simulation. *Front. Cogn.* **4**: 92.
81. Bousquet, O., K. Balakrishnan & V. Honavar. 1998. Is the hippocampus a Kalman filter? *Pac. Symp. Biocomput.* 657–668.
82. Fuhs, M.C. & D.S. Touretzky. 2007. Context learning in the rodent hippocampus. *Neural Comput.* **19**: 3173–3215.
83. Madl, T., S. Franklin, K. Chen, *et al.* 2014. Bayesian integration of information in hippocampal place cells. *PLoS One* **9**: e89762.
84. Sanders, H., C. Rennó-Costa, M. Idiart & J. Lisman. 2015. Grid cells and place cells: an integrated view of their navigational and memory function. *Trends Neurosci.* **38**: 763–775.
85. Sutton, R.S. & A.G. Barto. 1998. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.
86. Packard, M.G. & J.L. McGaugh. 1996. Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. *Neurobiol. Learn. Mem.* **65**: 65–72.
87. Rueckert, E., D. Kappel, D. Tanneberg, *et al.* 2016. Recurrent spiking networks solve planning tasks. *Sci. Rep.* **6**: 21142.
88. Friston, K., T. FitzGerald, F. Rigoli, *et al.* 2016. Active inference: a process theory. *Neural Comput.* **29**: 1–49.
89. Friston, K., T. FitzGerald, F. Rigoli, *et al.* 2016. Active inference and learning. *Neurosci. Biobehav. Rev.* **68**: 862–879.
90. Redish, A.D. 2016. Vicarious trial and error. *Nat. Rev. Neurosci.* **17**: 147–159.
91. Jensen, O. & J.E. Lisman. 1996. Hippocampal CA3 region predicts memory sequences: accounting for the phase precession of place cells. *Learn. Mem.* **3**: 279–287.
92. Buzsáki, G. 2010. Neural syntax: cell assemblies, synapse-sembles, and readers. *Neuron* **68**: 362–385.
93. Lisman, J.E. & O. Jensen. 2013. The theta–gamma neural code. *Neuron* **77**: 1002–1016.
94. Engel, A.K., P. Fries & W. Singer. 2001. Dynamic predictions: oscillations and synchrony in top-down processing. *Nat. Rev. Neurosci.* **2**: 704–716.
95. Arnal, L.H. & A.-L. Giraud. 2012. Cortical oscillations and sensory predictions. *Trends Cogn. Sci.* **16**: 390–398.
96. Zacksenhouse, M. & E. Ahissar. 2006. Temporal decoding by phase-locked loops: unique features of circuit-level implementations and their significance for vibrissal information processing. *Neural Comput.* **18**: 1611–1636.
97. Ranade, S., B. Hangya & A. Kepecs. 2013. Multiple modes of phase locking between sniffing and whisking during active exploration. *J. Neurosci.* **33**: 8250–8256.
98. Grion, N., A. Akrami, Y. Zuo, *et al.* 2016. Coherence between rat sensorimotor system and hippocampus is enhanced during tactile discrimination. *PLoS Biol* **14**: e1002384.
99. Colgin, L.L. 2016. Rhythms of the hippocampal network. *Nat. Rev. Neurosci.* **17**: 239–249.
100. Buzsáki, G. 2006. *Rhythms of the Brain*. Oxford University Press.
101. Lakatos, P., G. Karmos, A.D. Mehta, *et al.* 2008. Entrainment of neuronal oscillations as a mechanism of attentional selection. *Science* **320**: 110–113.
102. Morillon, B., T.A. Hackett, Y. Kajikawa & C.E. Schroeder. 2015. Predictive motor control of sensory dynamics in auditory active sensing. *Curr. Opin. Neurobiol.* **31**: 230–238.
103. van der Meer, M.A.A. & A.D. Redish. 2011. Theta phase precession in rat ventral striatum links place and reward information. *J. Neurosci.* **31**: 2843–2854.

104. van der Meer, M.A.A. & A.D. Redish. 2009. Covert expectation-of-reward in rat ventral striatum at decision points. *Front. Integr. Neurosci* **3**: 1.
105. Lansink, C.S., P.M. Goltstein, J.V. Lankelma, *et al.* 2009. Hippocampus—leads ventral striatum in replay of place-reward information. *PLoS Biol* **7**: e1000173.
106. Lansink, C.S., G.T. Meijer, J.V. Lankelma, *et al.* 2016. Reward expectancy strengthens CA1 theta and beta band synchronization and hippocampal—ventral striatal coupling. *J. Neurosci* **36**: 10598–10610.
107. Jones, M.W. & M.A. Wilson. 2005. Theta rhythms coordinate hippocampal—prefrontal interactions in a spatial memory task. *PLoS Biol* **3**: e402.
108. Smallwood, J. & J.W. Schooler. 2006. The restless mind. *Psychol. Bull.* **132**: 946–958.
109. Buesing, L., J. Bill, B. Nessler & W. Maass. 2011. Neural dynamics as sampling: a model for stochastic computation in recurrent networks of spiking neurons. *PLoS Comput. Biol.* **7**: e1002211.
110. Berkes, P., G. Orbán, M. Lengyel & J. Fiser. 2011. Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science* **331**: 83–87.
111. Wilson, M.A. & B.L. McNaughton. 1994. Reactivation of hippocampal ensemble memories during sleep. *Science* **265**: 676–679.
112. Louie, K. & M.A. Wilson. 2001. Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron* **29**: 145–156.
113. Lee, A.K. & M.A. Wilson. 2002. Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* **36**: 1183–1194.
114. Eschenko, O., W. Ramadan, M. Mölle, *et al.* 2008. Sustained increase in hippocampal sharp-wave ripple activity during slow-wave sleep after learning. *Learn. Mem.* **15**: 222–228.
115. Atherton, L.A., D. Dupret & J.R. Mellor. 2015. Memory trace replay: the shaping of memory consolidation by neuromodulation. *Trends Neurosci.* **38**: 560–570.
116. Kumaran, D., D. Hassabis & J.L. McClelland. 2016. What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends Cogn. Sci.* **20**: 512–534.
117. Buhry, L., A.H. Azizi & S. Cheng. 2011. Reactivation, replay, and preplay: how it might all fit together. *Neural Plast.* **2011**: 203462.
118. Kudrimoti, H.S., C.A. Barnes & B.L. McNaughton. 1999. Reactivation of hippocampal cell assemblies: effects of behavioral state, experience, and EEG dynamics. *J. Neurosci.* **19**: 4090–4101.
119. Simons, J.S. & H.J. Spiers. 2003. Prefrontal and medial temporal lobe interactions in long-term memory. *Nat. Rev. Neurosci.* **4**: 637–648.
120. Káli, S. & P. Dayan. 2004. Off-line replay maintains declarative memories in a model of hippocampal—neocortical interactions. *Nat. Neurosci.* **7**: 286–294.
121. Sutton, R.S. 1990. Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. In *Proceedings of the 7th International Conference on Machine Learning*. San Mateo, CA: 216–224.
122. Bishop, C.M. 2006. *Pattern Recognition and Machine Learning*. Springer.
123. Mehta, M.R. 2007. Cortico-hippocampal interaction during up-down states and memory consolidation. *Nat. Neurosci.* **10**: 13–15.
124. Peyrache, A., M. Khamassi, K. Benchenane, *et al.* 2009. Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nat. Neurosci.* **12**: 919–926.
125. Wierzynski, C.M., E.V. Lubenov, M. Gu & A.G. Siapas. 2009. State-dependent spike—timing relationships between hippocampal and prefrontal circuits during sleep. *Neuron* **61**: 587–596.
126. Hok, V., E. Save, P.P. Lenck-Santini & B. Poucet. 2005. Coding for spatial goals in the prelimbic/infralimbic area of the rat frontal cortex. *Proc. Natl. Acad. Sci. U.S.A.* **102**: 4602–4607.
127. Papale, A.E., M.C. Zielinski, L.M. Frank, *et al.* 2016. Interplay between hippocampal sharp-wave-ripple events and vicarious trial and error behaviors in decision making. *Neuron* **92**: 975–982.
128. Johnson, A., Z. Varberg, J. Benhardus, *et al.* 2012. The hippocampus and exploration: dynamically evolving behavior and neural representations. *Front. Hum. Neurosci.* **6**: 216.
129. Cheng, S. & L.M. Frank. 2008. New experiences enhance coordinated neural activity in the hippocampus. *Neuron* **57**: 303–313.
130. Tolman, E.C. 1948. Cognitive maps in rats and men. *Psychol. Rev.* **55**: 189–208.
131. Hinton, G.E. 2007. Learning multiple layers of representation. *Trends Cogn. Sci* **11**: 428–434.
132. Leutgeb, J.K., S. Leutgeb, M.-B. Moser & E.I. Moser. 2007. Pattern separation in the dentate gyrus and CA3 of the hippocampus. *Science* **315**: 961–966.
133. Tulving, E. 2002. Episodic memory: from mind to brain. *Annu. Rev. Psychol.* **53**: 1–25.
134. Eichenbaum, H. & N.J. Cohen. 2014. Can we reconcile the declarative memory and spatial navigation views on hippocampal function? *Neuron* **83**: 764–770.
135. Atkinson, R.C. & R.M. Shiffrin. 1968. “Human memory: a proposed system and its control processes.” In *The Psychology of Learning and Motivation: Advances in Research and Theory*. 2. K.W. Spence & J.T. Spence, Eds.: 742–775. New York: Academic Press.
136. Miller, G.A. 1956. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychol. Rev.* **63**: 81–97.
137. Lisman, J.E. & M.A. Idiart. 1995. Storage of 7 plus/minus 2 short-term memories in oscillatory subcycles. *Science* **267**: 1512.
138. Moscovitch, M., R. Cabeza, G. Winocur & L. Nadel. 2016. Episodic memory and beyond: the hippocampus and neocortex in transformation. *Annu. Rev. Psychol.* **67**: 105–134.
139. Dragoi, G. 2013. Internal operations in the hippocampus: single cell and ensemble temporal coding. *Front. Syst. Neurosci.* **7**: 46.
140. Malvache, A., S. Reichinnek, V. Villette, *et al.* 2016. Awake hippocampal reactivations project onto orthogonal neuronal assemblies. *Science* **353**: 1280–1283.

141. Friston, K. & G. Buzsáki. 2016. The functional anatomy of time: what and when in the brain. *Trends Cogn. Sci.* **20**: 500–511.
142. Rabinovich, M., A. Volkovskii, P. Lecanda, *et al.* 2001. Dynamical encoding by networks of competing neuron groups: winnerless competition. *Phys. Rev. Lett.* **87**: 068102.
143. Buonomano, D.V. & W. Maass. 2009. State-dependent computations: spatiotemporal processing in cortical networks. *Nat. Rev. Neurosci.* **10**: 113–125.
144. Rolls, E.T. & G. Deco. 2015. Networks for memory, perception, and decision-making, and beyond to how the syntax for language might be implemented in the brain. *Brain Res.* **1621**: 316–334.
145. Grossberg, S. 1978. “A theory of human memory: self-organization and performance of sensory-motor codes, maps, and plans.” In *Progress in Theoretical Biology*. 5. R. Rosen & F. Snell, Eds.: 233–374. New York: Academic Press.
146. Eacott, M.J. & G. Norman. 2004. Integrated memory for object, place, and context in rats: a possible model of episodic-like memory? *J. Neurosci.* **24**: 1948–1953.
147. Kemp, C. & J.B. Tenenbaum. 2008. The discovery of structural form. *Proc. Natl. Acad. Sci. U.S.A.* **105**: 10687–10692.
148. Stickgold, R. & M.P. Walker. 2013. Sleep-dependent memory triage: evolving generalization through selective processing. *Nat. Neurosci.* **16**: 139–145.
149. Maguire, E.A. & S.L. Mullally. 2013. The hippocampus: a manifesto for change. *J. Exp. Psychol. Gen.* **142**: 1180–1189.
150. Bromberg-Martin, E.S., M. Matsumoto & O. Hikosaka. 2010. Dopamine in motivational control: rewarding, aversive, and alerting. *Neuron* **68**: 815–834.
151. Ennaceur, A. & J. Delacour. 1988. A new one-trial test for neurobiological studies of memory in rats. 1: behavioral data. *Behav. Brain Res.* **31**: 47–59.
152. Harvey, C.D., P. Coen & D.W. Tank. 2012. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* **484**: 62–68.
153. Fujisawa, S., A. Amarasingham, M.T. Harrison & G. Buzsáki. 2008. Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat. Neurosci.* **11**: 823–833.
154. Dejean, C. *et al.* 2016. Prefrontal neuronal assemblies temporally control fear behaviour. *Nature* **535**: 420–424.
155. O’Neill, J., C.N. Boccara, F. Stella, *et al.* 2017. Superficial layers of the medial entorhinal cortex replay independently of the hippocampus. *Science* **355**: 184–188.
156. Akhlaghpour, H., J. Wiskerke, J.Y. Choi, *et al.* 2016. Dissociated sequential activity and stimulus encoding in the dorsomedial striatum during spatial working memory. *eLife* **5**: e19507.
157. Laurent, G. 1996. Dynamical representation of odors by oscillating and evolving neural assemblies. *Trends Neurosci.* **19**: 489–496.
158. Mazor, O. & G. Laurent. 2005. Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron* **48**: 661–673.
159. Daw, N.D. & P. Dayan. 2014. The algorithmic anatomy of model-based evaluation. *Philos. Trans. R. Soc. B Biol. Sci.* **369**: 20130478.
160. Gelly, S. & D. Silver. 2011. Monte-Carlo tree search and rapid action value estimation in computer Go. *Artif. Intell.* **175**: 1856–1875.
161. Dayan, P., G.E. Hinton, R.M. Neal & R.S. Zemel. 1995. The helmholtz machine. *Neural Comput.* **7**: 889–904.
162. Fiser, J., P. Berkes, G. Orbán & M. Lengyel. 2010. Statistically optimal perception and learning: from behavior to neural representations. *Trends Cogn. Sci.* **14**: 119–130.
163. Huang, Y. & R.P.N. Rao. 2014. Neurons as Monte Carlo samplers: Bayesian inference and learning in spiking networks. In *Proceedings of the 27th International Conference on Neural Information Processing Systems*, Cambridge, MA: 1943–1951.
164. Pecevski, D. & W. Maass. 2016. Learning probabilistic inference through spike-timing-dependent plasticity. *eNEURO* **3**: pii: ENEURO.0048-15.2016.
165. Donnarumma, F., D. Maisto & G. Pezzulo. 2016. Problem solving as probabilistic inference with subgoal: explaining human successes and pitfalls in the tower of hanoi. *PLoS Comput. Biol.* **12**: e1004864.
166. Balaguer, J., H. Spiers, D. Hassabis & C. Summerfield. 2016. Neural mechanisms of hierarchical planning in a virtual subway network. *Neuron* **90**: 893–903.
167. Stoianov, I., A. Genovesio & G. Pezzulo. 2015. Prefrontal goal codes emerge as latent states in probabilistic value learning. *J. Cogn. Neurosci.* **28**: 140–157.
168. Ito, H.T., S.-J. Zhang, M.P. Witter, *et al.* 2015. A prefrontal-thalamo-hippocampal circuit for goal-directed spatial navigation. *Nature* **522**: 50–55.
169. Jezek, K., E.J. Henriksen, A. Treves, *et al.* 2011. Theta-paced flickering between place-cell maps in the hippocampus. *Nature* **478**: 246–249.
170. Pezzulo, G. & C. Castelfranchi. 2007. The symbol detachment problem. *Cogn. Process.* **8**: 115–131.
171. Pezzulo, G. 2017. Tracing the roots of cognition in predictive processing. In *Philosophy and Predictive Processing*, 20. Frankfurt am Main: MIND Group.
172. Schacter, D.L. & D.R. Addis. 2007. The cognitive neuroscience of constructive memory: remembering the past and imagining the future. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **362**: 773–786.
173. Christoff, K., Z.C. Irving, K.C.R. Fox, *et al.* 2016. Mind-wandering as spontaneous thought: a dynamic framework. *Nat. Rev. Neurosci.* **17**: 718–731.
174. Benedek, M., E. Jauk, R.E. Beaty, *et al.* 2016. Brain mechanisms associated with internally directed attention and self-generated thought. *Sci. Rep.* **6**: 22959.
175. Martin, V.C., D.L. Schacter, M.C. Corballis & D.R. Addis. 2011. A role for the hippocampus in encoding simulations of future events. *Proc. Natl. Acad. Sci. U.S.A.* **108**: 13858–13863.
176. Gaesser, B., R.N. Spreng, V.C. McLelland, *et al.* 2013. Imagining the future: evidence for a hippocampal contribution to constructive processing. *Hippocampus* **23**: 1150–1161.
177. Race, E., M.M. Keane & M. Verfaellie. 2013. Losing sight of the future: impaired semantic prospection following medial temporal lobe lesions. *Hippocampus* **23**: 268–277.
178. Addis, D.R., T. Cheng, R.P. Roberts & D.L. Schacter. 2011. Hippocampal contributions to the episodic simulation of

- specific and general future events. *Hippocampus* **21**: 1045–1052.
179. Brown, T.I. *et al.* 2016. Prospective representation of navigational goals in the human hippocampus. *Science* **352**: 1323–1326.
  180. Hassabis, D., D. Kumaran, S.D. Vann & E.A. Maguire. 2007. Patients with hippocampal amnesia cannot imagine new experiences. *Proc. Natl. Acad. Sci. U.S.A.* **104**: 1726–1731.
  181. Buckner, R.L., J.R. Andrews-Hanna & D.L. Schacter. 2008. The brain's default network: anatomy, function, and relevance to disease. *Ann. N.Y. Acad. Sci.* **1124**: 1–38.
  182. Pezzulo, G. 2016. "The mechanisms and benefits of a future-oriented brain." In *Seeing the Future*. K. Michaelian, S.B. Klein & K.K. Szpunar, Eds.: 267–284. Oxford University Press.
  183. Kurth-Nelson, Z., M. Economides, R.J. Dolan & P. Dayan. 2016. Fast sequences of non-spatial state representations in humans. *Neuron* **91**: 194–204.
  184. Kaplan, R., J. King, R. Koster, *et al.* 2017. The neural representation of prospective choice during spatial planning and decisions. *PLoS Biol.* **15**: e1002588.
  185. Cei, A., G. Girardeau, C. Drieu, *et al.* 2014. Reversed theta sequences of hippocampal cell assemblies during backward travel. *Nat. Neurosci.* **17**: 719–724.
  186. Fanselow, M.S. 1990. Factors governing one-trial contextual conditioning. *Anim. Learn. Behav.* **18**: 264–270.
  187. Robbe, D., S.M. Montgomery, A. Thome, *et al.* 2006. Cannabinoids reveal importance of spike timing coordination in hippocampal function. *Nat. Neurosci.* **9**: 1526–1533.
  188. White, N.M. 1989. Reward or reinforcement: what's the difference?. *Neurosci. Biobehav. Rev.* **13**: 181–186.
  189. Momennejad, I., E.M. Russek, J.H. Cheong, *et al.* 2016. The successor representation in human reinforcement learning. *bioRxiv*. 083824.
  190. Eacott, M.J. & A. Easton. 2007. On familiarity and recall of events by rats. *Hippocampus* **17**: 890–897.
  191. van der Meer, M., A.A. Carey & Y. Tanaka. 2017. Optimizing for generalization in the decoding of internally generated activity in the hippocampus. *Hippocampus*. <https://doi.org/10.1002/10.1002/hipo.22714>.
  192. Wu, C.T., D. Haggerty, C. Kemere & D. Ji. 2017. Hippocampal awake replay in fear memory retrieval. *Nat. Neurosci.* <https://doi.org/10.1002/10.1038/nn.4507>.
  193. Pezzulo, G., F. Rigoli & K. Friston. 2016. Active Inference, homeostatic regulation and adaptive behavioural control. *Prog. Neurobiol.* **134**: 17–35.
  194. Donnarumma, F., M. Costantini, E. Ambrosini, K. Friston & G. Pezzulo. 2017. Action perception as hypothesis testing. *Cortex*. **89**: 45–60.