Published in final edited form as:

Curr Opin Behav Sci. 2017 October ; 17: 7–13. doi:10.1016/j.cobeha.2017.05.025.

Predicting the Past, Remembering the Future

Samuel J. Gershman

Department of Psychology and Center for Brain Science, Harvard University

Abstract

Rational analyses of memory suggest that retrievability of past experience depends on its usefulness for predicting the future: memory is adapted to the temporal structure of the environment. Recent research has enriched this view by applying it to semantic memory and reinforcement learning. This paper describes how multiple forms of memory can be linked via common predictive principles, possibly subserved by a shared neural substrate in the hippocampus. Predictive principles offer an explanation for a wide range of behavioral and neural phenomena, including semantic fluency, temporal contiguity effects in episodic memory, and the topological properties of hippocampal place cells.

Keywords

Successor representation; rational analysis; hippocampus; temporal context

Introduction

Why remember the past? George Santayana famously remarked that “Those who cannot remember the past are condemned to repeat it” [1], and William Faulkner expressed a similar sentiment in *Requiem for a Nun*: “The past is never dead. It isn't even the past.” Patterns of recurrence have fascinated historians from Polybius to Arnold Toynbee [2], and some have tried to harvest lessons from these patterns to guide political decision making [3]. More mundanely, the daily lives of most organisms are structured by cycles of sleeping, eating and other routine activities. At shorter timescales, many events can be anticipated using sensory information from the recent past, giving organisms with memory an adaptive advantage. Thus, memory of the past, spanning vastly different timescales, is relevant for predicting and controlling the future [4, 5, 6, 7].

Most research on human memory has focused on understanding, at a descriptive level, how information from the past is stored and retrieved, without contemplating the usefulness of this information for future action. However, researchers have increasingly come to appreciate the fact that memory is organized around predictive design principles [8], evident

Address for correspondence: Samuel Gershman, Department of Psychology, Harvard University, 52 Oxford St., room 295.05, Cambridge, MA 02138, Phone: 773-607-9817, gershman@fas.harvard.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

in multiple forms of memory (semantic, episodic, short-term and procedural). Computational models have formalized these principles mathematically, drawing upon ideas from library science, search engine algorithms, probability theory, and reinforcement learning. Despite this diversity of approaches, we will see that several of them can be unified in terms of a single predictive representation, repurposed (with slight modifications) for the needs of different memory systems. Neural correlates of this predictive representation have been observed in the hippocampus, suggesting a functional explanation for the region's involvement in both retrospective and prospective cognition.

The rational analysis of memory

John Anderson's "rational analysis of memory" was the first theoretical attempt to explain the structure of memory in terms of beliefs about the future [9, 10]. Anderson conceptualized the problem facing memory as one of determining *need probability*: the likelihood that a particular piece of information will be needed in the future. On the assumption that past need predicts future need, memory can be used to produce a forecast. Importantly, because need probability may change over time, this forecast should adapt to the statistics of past experience.

It is worth highlighting here the extent to which this view represents a significant departure from prevailing ideas about episodic memory. The standard view (see [11] for a summary) holds that there is a fundamental tension between memory for specific instances (episodic memory) and memory for statistical regularities (semantic memory). In contrast, Anderson's rational analysis posited that episodic memory is structured by statistical regularities. To formalize this idea, Anderson adapted a model of library borrowing, likening memories to books in a library system. The usage of each book is subject to fluctuations, and the task facing the library system is to track these usage statistics in order to anticipate the probability that a book will be needed in the near future. Anderson showed that this model could capture many basic properties of memory, such as spacing, recency, fan and word frequency effects.

Several authors have explored different assumptions about the statistical regularities governing episodic memory. For example, Gershman and colleagues [12] developed a Bayesian nonparametric model in which the environmental dynamics can switch between an unknown number of "modes." This model makes the prediction (confirmed experimentally) that abrupt changes will result in the inference that a new mode has been encountered, effectively creating an event boundary that reduces interference between items on either side of the boundary. The importance of event boundaries in memory formation and retrieval has been highlighted by a number of other recent studies [13, 14, 15, 16, 17, 18].

Rational analysis has also been applied to semantic memory. In this case, the need probability is dictated by the long-run marginal probability of items, rather than the temporally specific need probability used in the analysis of episodic memory. Griffiths and colleagues [19] showed that the semantic need probability is precisely what is computed by Google's PageRank search algorithm [20]. They conceptualized semantic memory as a directed network (a kind of mental World Wide Web) over which semantic processing flows

according to a random walk. If we let $\bar{L}_{ij} = L_{ij} / \sum_k L_{ik}$ denote the normalized link matrix (where $L_{ij} = 1$ if there is a link from item i to item j , and 0 otherwise), then PageRank computes the vector of item ranks $\mathbf{r}$ according to the recursive definition $\mathbf{r} = \bar{\mathbf{L}}\mathbf{r}$.

We can understand the rank of item i as proportional to its marginal probability $P(i)$ after running the random walk on the semantic network for a long time. Intuitively, this means that the proportion of time a person spends thinking about item i scales with its rank. It can also be interpreted as a probabilistic model of environmental structure: if the need for an item over time can be described by a random walk, then its rank reflects its long-run need probability. Griffiths and colleagues asked participants to generate words beginning with a particular letter, and showed that the number of times a word was retrieved (a measure of semantic fluency) was well-predicted by its rank. This suggests that human semantic memory is structured to make predictions of future need probability over long timescales.

A predictive substrate

Another way to derive the PageRank algorithm is to construct a predictive representation for each item i that encodes the expected discounted number of times item j will be needed on a random walk initiated at i :

$$M_{ij} = \left[\sum_{t=1}^{\infty} \gamma^t \mathbb{I}(s_t=j) | s_0=i \right], \quad (1)$$

where s_t denotes the item needed at time t , $\mathbb{I}(\cdot) = 1$ if its argument is true (0 otherwise), and γ is a discount factor controlling the effective time horizon over which item counts are accumulated. The expectation is taken over randomness in state transitions and rewards. One can interpret $1 - \gamma$ as the constant probability that the random walk will terminate at any given time. An item's rank can be computed from the predictive representation simply by summing the rows of $\mathbf{M}$: $r_j \propto \sum_i M_{ij}$. In other words, the long-run probability of visiting an item is obtained by summing its cumulative expected need across the collection of random walks initiated at different items.

This way of deriving PageRank is intriguing because the matrix $\mathbf{M}$ has another life in the reinforcement learning literature, where it is known as the *successor representation* (SR) [21]. The central problem in reinforcement learning is the computation of *value*, the discounted cumulative reward expected in the future upon entering a state:

$$V_i = \left[\sum_{t=1}^{\infty} \gamma^t R(s_t) | s_0=i \right], \quad (2)$$

where $R(s)$ is the immediate reward expected upon entering state s . The SR renders value computation a linear operation, due to the identity $V_i = \sum_j M_{ij} R(j)$. It therefore offers a

significant computational advantage over traditional model-based planning algorithms that typically scale super-linearly with the number of states [22].

Further insight can be gleaned by considering the long-run average reward. It can be shown that the long-run average reward is a reward-weighted version of the rank (and hence need probability) returned by the PageRank algorithm, revealing a deep and surprising formal connection between reinforcement learning, information retrieval, and the rational analysis of memory. It suggests a way in which reward associations can reweight need probabilities, in accordance with well-known motivational influences on memory [23, 24, 25].

Learning the successor representation with temporal context

The SR is defined as an expectation over infinitely long trajectories, raising the question of how it can be tractably computed. One possibility, borrowing ideas from reinforcement learning, is to directly update an estimate the SR ($\hat{M}$) from state transitions. Specifically, the SR can be updated incrementally using a form of temporal difference learning [26]:

$$\Delta \hat{M}(i, j) = \alpha [(s_{t+1} = j) + \gamma \hat{M}(s_{t+1}, j) - \hat{M}(s_t, j)] e_t(i), \quad (3)$$

where α is a learning rate and $e_t(i)$ is an exponentially decaying memory trace of recent states that specifies the “eligibility” of state i for updating. The term in brackets represents a “prediction error” analogous to the reward prediction error thought to be conveyed by phasic dopamine signaling [27], but in this case the prediction error is defined on state expectations rather than on reward expectations.

The temporal difference update described above is nearly identical to an influential theory of episodic memory, the Temporal Context Model (TCM; [28, 29]). TCM was originally developed to explain the temporal structure of memory retrieval in free recall experiments, but has more recently been applied to many other phenomena, ranging from false memory [30, 31] to spacing and repetition effects [32]. According to TCM, items are bound in memory to a slowly drifting representation of temporal context (a recency-weighted average of previous items), and at test the temporal context acts as a retrieval cue, preferentially drawing items based on their strength of association. It can be shown (see [26] for details) that the temporal context representation corresponds to the eligibility trace e_t , and the matrix of item-context associations corresponds to $\hat{M}$.

The connection between TCM and the SR suggests another way in which episodic memory can be cast in terms of predictive mechanisms. However, the temporal difference learning update is equivalent to the originally proposed Hebbian update rule [28] only when items are presented once per list. In this case, the second term in the temporal difference update is always zero, because no predictions for novel items have been formed yet. When items are repeated, the second term comes into play, altering the theory's predictions. An elegant study [33] drew out these implications, testing the prediction that the temporal difference update will produce a *context repetition effect*: repeating the temporal context of a particular item will strengthen memory for that item, even if the item itself was not repeated. This

prediction follows from the fact that the second term in the temporal difference update will be positive whenever an item's temporal context recurs. The study found support for this prediction across several experiments, as well as a boundary condition, whereby the effect only occurs when the item is strongly predicted by the temporal context. This again is entirely consistent with the theory, since the second term in the update will be smaller when the predictive relationship between context and item is weak. The context repetition effect was also recently demonstrated in rats using an object recognition task [34].

The role of the hippocampus

The multi-faceted role of predictive representations in memory suggests that they may have a dedicated neural substrate. Several lines of evidence suggest that the hippocampus, although traditionally viewed as a repository for episodic memory and spatial representation, may in fact be organized around predictive principles [35, 36]. First, place cells in the hippocampus sweep ahead of an animal's current position when it is at a choice point [37]; these forward sweeps may arise from phase precession, the progressive shift in spike timing relative to the ongoing theta oscillation as an animal moves through a place field [38, 39, 4]. Second, when an animal repeatedly runs on a particular trajectory, place fields tend to expand opposite the direction of travel [40], consistent with the idea that earlier place cells learn to predict upcoming locations. Third, functional MRI experiments in humans have indicated that the hippocampus is sensitive to the predictability of upcoming stimuli [41, 42, 43, 44]. Fourth, the hippocampus is activated when humans engage in episodic future-thinking [45], and damage to the hippocampus appears to severely impair this ability [46] (though this finding is controversial; see [47]).

Recent theoretical work has formalized the predictive function of the hippocampus in terms of the SR [48]. According to this view, place cells do not actually encode spatial location; they encode expectations about future locations (though see [49] for evidence that some place cells have no predictive properties during immobility). In an open field, these future locations will tend to be in the vicinity of the animal's current location, yielding classical radially symmetric place fields. However, predictive fields will become distorted when obstacles are introduced or the environmental topology is altered (Figure 1)—just as observed experimentally [50, 51, 52]. Likewise, replay of place cell sequences appear to respect the environmental topology [53]. The topological sensitivity of the SR is a natural consequence of the fact that an animal's future trajectory is constrained by barriers to movement.

Even when there aren't physical barriers to movement, the SR may still distort along certain paths, for example when rewards occur in reliable locations. This follows from the fact that an animal's policy tends to be reward-seeking. Consistent with this hypothesis, place fields tend to cluster around rewarded locations [54]. Moreover, a recent experiment found that forward sweeps preferentially visit rewarded locations [55].

Early during learning, the SR may resemble a relatively “pure” representation of space, since the predictive relationships between states have not yet been learned. The fact that some place cells show no predictive properties during early exposure to an environment [49]

aligns well with the classic finding that contextual fear conditioning is more effective if the animal is first pre-exposed to the environment (an effect that is hippocampus-dependent [56]). According to the SR theory, predictive relationships between states allow the shock association to propagate to other states (i.e., locations within the environment), leading to a more robust fear memory.

It is important to note that the SR theory does not entirely dispense with some notion of purely spatial representation, as posited by the classical place cell literature. In fact, the SR, when defined over space, is predicated on such a representation, since this is the only way to index spatially distinct states. Likewise, place cells could also index local perceptual features [57], allowing the SR to be defined over this feature space.

Another point of contact between the classical place cell literature and the SR theory concerns the mechanisms of SR updating. The previous section described a temporal difference learning algorithm for updating the SR based on experienced trajectories. However, it is well-known that animals can rapidly change their behavior in response to structural changes in the environment without directly experiencing those changes, such as the detour behavior observed by Tolman [58]. These rapid behavioral changes seem to require an internal model or “cognitive map” that can be updated in the absence of direct experience. One way to accomplish this within the SR theory is to posit a simulation mechanism that can use an internal model to update the SR [59, 60], possibly implemented by forward sweeps of place cells [37].

One of the strengths of the SR theory is that it provides a framework for understanding the origin of non-spatial, relational representations in the hippocampus. For example, the hippocampus is involved in the computation of transitive inferences [61], and damage to the hippocampus impairs transitive inference [62]. Recently, Garvert and colleagues [63] used functional MRI to directly examine the SR account of relational representation. Subjects in their study were exposed to sequences of items generated by a random walk on a graph. Representational similarity in the hippocampus (as measured by adaptation of the neural response across random pairs of items) scaled inversely with distance on the graph, and quantitative analyses revealed that this similarity structure was best described by the SR compared to other graph-theoretic measures.

Predictive spiking neurons

How does the SR arise in biologically realistic neural circuits? Neurons must spike within a few tens of milliseconds of one another for their synapses to be strengthened, much shorter than the temporal horizon of predictive codes like the SR (on the order of seconds). Brea and colleagues [64] have developed a framework for resolving this puzzle, using a variant of traditional spike timing-dependent plasticity [65].

The basic idea is illustrated in Figure 2. When a synapse between two neurons is potentiated, the presynaptic neuron can cause the postsynaptic neuron to fire earlier. This in turn allows the postsynaptic neuron to enter into plasticity with other neurons that are spiking even earlier. Thus, the membrane potential of the postsynaptic neuron comes to reflect the

anticipatory drive from presynaptic inputs; in effect, the postsynaptic neuron is anticipating its own future spiking. This is the same kind of bootstrapping that underlies the temporal difference learning algorithm described above.

Brea and colleagues [64] showed that a form of spike timing-dependent plasticity can be used to implement predictive spiking. In particular, under some fairly general assumptions their plasticity rule enables the dendritic potential to directly encode the expected discounted future spike rate at the soma—i.e., one column of the SR matrix $\mathbf{M}$. Predictive spiking offers an explanation for widespread anticipatory responses in the brain [66, 43, 67].

Conclusions

Starting from the principle that memory is designed to serve a predictive function, a rich web of theoretical insights can be derived. Semantic and episodic memory can be linked via a common predictive representation initially studied in the reinforcement learning literature, allowing us to contemplate the computational properties that these different forms of memory may have in common. These commonalities may derive from a shared neural substrate, as suggested by the widespread involvement of the hippocampus across domains [68].

Despite the appeal of a unifying neural substrate, this viewpoint is in tension with taxonomies of memory that stipulate dissociable systems in the brain [69]. For example, transformational consolidation theories hold that episodic memories stored initially in the hippocampus are gradually transformed into semantic memories stored in neocortex [70, 11]. This viewpoint is supported by evidence that hippocampal amnesics show spared long-term semantic memory [71]. However, more recent experiments have found semantic impairments in amnesia (see [72] for a review), and amnesics exhibit semantically impoverished episodic future-thinking [73]. Beyond semantic memory, there is now considerable evidence that the hippocampus (and episodic memory more generally) plays an important role in reinforcement learning as well [7]. Collectively, these observations lend credence to the idea that multiple forms of memory may draw upon a common predictive representation in the hippocampus.

Another tension derives from dissociations between different forms of memory that support behavioral control in reinforcement learning tasks. For example, stimulus-response strategies depend on sub-regions of the striatum but not on the hippocampus, whereas allocentric “place” strategies depend on the hippocampus but not the striatum [74]. These dissociations indicate that a predictive representation in the hippocampus can only be one part of a larger multi-system architecture, whose components interact both competitively and cooperatively [7].

Acknowledgments

I am grateful to Matthew Botvinick and Nathaniel Daw for helpful discussions, and to Dan Schacter for comments on a previous draft. This work was made possible through grant support from the National Institutes of Health (CRCNS 1207833).

References

1. Santayana, George. *The Life Of Reason*. Vol. 1. Constable And Company Limited; 1905.
2. Trompf, Garry Winston. *The Idea of Historical Recurrence in Western thought: from Antiquity to the Reformation*. University of California Press; 1979.
3. Neustadt, Richard E., Ernest, R. *Thinking in Time: The Uses of History for Decision-Makers*. New York: Free Press; 1986.
4. Shankar, Karthik H., Singh, Inder, Howard, Marc W. Neural mechanism to simulate a scale-invariant future. *Neural Computation*. 2016; 28:2594–2627. [PubMed: 27626961]
5. Heeger, David J. Theory of cortical function. *Proceedings of the National Academy of Sciences*. 2017; 114:1773–1782.
6. Palmer, Stephanie E., Marre, Olivier, Berry, Michael J., Bialek, William. Predictive information in a sensory population. *Proceedings of the National Academy of Sciences*. 2015; 112:6908–6913.
7. Gershman, Samuel J., Daw, Nathaniel D. Reinforcement learning and episodic memory in humans and animals: An integrative framework. *Annual Review of Psychology*. 2017; 68:101–128.
8. Schacter, Daniel L. Adaptive constructive processes and the future of memory. *American Psychologist*. 2012; 67:603–613. [PubMed: 23163437]
9. Anderson, John R., Milson, Robert. Human memory: An adaptive perspective. *Psychological Review*. 1989; 96:703–719.
10. Anderson, John R., Schooler, Lael J. Reflections of the environment in memory. *Psychological Science*. 1991; 2:396–408.
11. Kumaran, Dharshan, Hassabis, Demis, McClelland, James L. What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends in Cognitive Sciences*. 2016; 20:512–534. [PubMed: 27315762]
12. Gershman, Samuel J., Radulescu, Angela, Norman, Kenneth A., Niv, Yael. Statistical computations underlying the dynamics of memory updating. *PLoS Computational Biology*. 2014; 10:e1003939. [PubMed: 25375816]
13. Swallow, Khena M., Zacks, Jeffrey M., Abrams, Richard A. Event boundaries in perception affect memory encoding and updating. *Journal of Experimental Psychology: General*. 2009; 138:236–257. [PubMed: 19397382]
14. Ezzyat, Youssef, Davachi, Lila. What constitutes an episode in episodic memory? *Psychological Science*. 2011; 22:243–252. [PubMed: 21178116]
15. Ezzyat, Youssef, Davachi, Lila. Similarity breeds proximity: pattern similarity within and across contexts is related to later mnemonic judgments of temporal proximity. *Neuron*. 2014; 81:1179–1189. [PubMed: 24607235]
16. DuBrow, Sarah, Davachi, Lila. The influence of context boundaries on memory for the sequential order of events. *Journal of Experimental Psychology: General*. 2013; 142:1277–1286. [PubMed: 23957281]
17. Horner, Aidan J., Bisby, James A., Wang, Aijing, Bogus, Katrina, Burgess, Neil. The role of spatial boundaries in shaping long-term event representations. *Cognition*. 2016; 154:151–164. [PubMed: 27295330]
18. Pettijohn, Kyle A., Thompson, Alexis N., Tamplin, Andrea K., Krawietz, Sabine A., Radvansky, Gabriel A. Event boundaries and memory improvement. *Cognition*. 2016; 148:136–144. [PubMed: 26780472]
19. Griffiths, Thomas L., Steyvers, Mark, Firl, Alana. Google and the mind predicting fluency with pagerank. *Psychological Science*. 2007; 18:1069–1076. [PubMed: 18031414]
20. Brin, Sergey, Page, Lawrence. The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN Systems*. 1998; 30:107–117.
21. Dayan, Peter. Improving generalization for temporal difference learning: The successor representation. *Neural Computation*. 1993; 5(4):613–624.
22. Daw, Nathaniel D., Dayan, Peter. The algorithmic anatomy of model-based evaluation. *Phil Trans R Soc B*. 2014; 369:20130478. [PubMed: 25267820]

23. Wittmann, Bianca C., Schott, Björn H., Guderian, Sebastian, Frey, Julietta U., Heinze, Hans-Jochen, Düzel, Emrah. Reward-related fmri activation of dopaminergic midbrain is associated with enhanced hippocampus-dependent long-term memory formation. *Neuron*. 2005; 45:459–467. [PubMed: 15694331]
24. Adcock, R Alison, Thangavel, Arul, Whitfield-Gabrieli, Susan, Knutson, Brian, Gabrieli, John DE. Reward-motivated learning: mesolimbic activation precedes memory formation. *Neuron*. 2006; 50:507–517. [PubMed: 16675403]
25. Miendlarzewska, Ewa A., Bavelier, Daphne, Schwartz, Sophie. Influence of reward motivation on human declarative memory. *Neuroscience & Biobehavioral Reviews*. 2016; 61:156–176. [PubMed: 26657967]
26. Gershman, Samuel J., Moore, Christopher D., Todd, Michael T., Norman, Kenneth A., Sederberg, Per B. The successor representation and temporal context. *Neural Computation*. 2012; 24:1553–1568. [PubMed: 22364500]
27. Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. *Science*. 1997; 275:1593–1599. [PubMed: 9054347]
28. Howard MW, Kahana MJ. A distributed representation of temporal context. *Journal of Mathematical Psychology*. 2002; 46:269–299.
29. Sederberg PB, Howard MW, Kahana MJ. A context-based theory of recency and contiguity in free recall. *Psychological Review*. 2008; 115(4):893–912. [PubMed: 18954208]
30. Sederberg, Per B., Gershman, Samuel J., Polyn, Sean M., Norman, Kenneth A. Human memory reconsolidation can be explained using the temporal context model. *Psychonomic Bulletin & Review*. 2011; 18:455–468. [PubMed: 21512839]
31. Gershman, Samuel J., Schapiro, Anna C., Hupbach, Almut, Norman, Kenneth A. Neural context reinstatement predicts memory misattribution. *The Journal of Neuroscience*. 2013; 33:8590–8595. [PubMed: 23678104]
32. Siegel, Lynn L., Kahana, Michael J. A retrieved context account of spacing and repetition effects in free recall. *Journal of Experimental Psychology: Learning Memory and Cognition*. 2014; 40:755–764.
33. Smith, Troy A., Hasinski, Adam E., Sederberg, Per B. The context repetition effect: Predicted events are remembered better, even when they don't happen. *Journal of Experimental Psychology: General*. 2013; 142:1298–1308. [PubMed: 23957285]
34. Manns, Joseph R., Galloway, Claire R., Sederberg, Per B. A temporal context repetition effect in rats during a novel object recognition memory task. *Animal Cognition*. 2015; 18:1031–1037. [PubMed: 25917312]
35. Levy, William B., Hocking, Ashlie B., Wu, Xiangbao. Interpreting hippocampal function as recoding and forecasting. *Neural Networks*. 2005; 18:1242–1264. [PubMed: 16269237]
36. Lisman, John, Redish, A David. Prediction, sequences and the hippocampus. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*. 2009; 364:1193–1201. [PubMed: 19528000]
37. Johnson, Adam, Redish, A David. Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *The Journal of Neuroscience*. 2007; 27:12176–12189. [PubMed: 17989284]
38. Jensen, Ole, Lisman, John E. Hippocampal CA3 region predicts memory sequences: accounting for the phase precession of place cells. *Learning & Memory*. 1996; 3:279–287. [PubMed: 10456097]
39. Skaggs, William E., McNaughton, Bruce L. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*. 1996; 6:149–172. [PubMed: 8797016]
40. Mehta, Mayank R., Quirk, Michael C., Wilson, Matthew A. Experience-dependent asymmetric shape of hippocampal receptive fields. *Neuron*. 2000; 25:707–715. [PubMed: 10774737]
41. Strange, Bryan A., Duggins, Andrew, Penny, William, Dolan, Raymond J., Friston, Karl J. Information theory, novelty and hippocampal responses: unpredicted or unpredictable? *Neural Networks*. 2005; 18:225–230. [PubMed: 15896570]

42. Bornstein, Aaron M., Daw, Nathaniel D. Dissociating hippocampal and striatal contributions to sequential prediction learning. *European Journal of Neuroscience*. 2012; 35:1011–1023. [PubMed: 22487032]
43. Schapiro, Anna C., Kustner, Lauren V., Turk-Browne, Nicholas B. Shaping of object representations in the human medial temporal lobe based on temporal regularities. *Current Biology*. 2012; 22:1622–1627. [PubMed: 22885059]
44. Schapiro, Anna C., Turk-Browne, Nicholas B., Norman, Kenneth A., Botvinick, Matthew M. Statistical learning of temporal community structure in the hippocampus. *Hippocampus*. 2016; 26:3–8. [PubMed: 26332666]
45. Schacter, Daniel L., Addis, Donna Rose, Buckner, Randy L. Remembering the past to imagine the future: the prospective brain. *Nature Reviews Neuroscience*. 2007; 8:657–661. [PubMed: 17700624]
46. Hassabis, Demis, Kumaran, Dharshan, Vann, Seralynne D., Maguire, Eleanor A. Patients with hippocampal amnesia cannot imagine new experiences. *Proceedings of the National Academy of Sciences*. 2007; 104:1726–1731.
47. Squire, Larry R., van der Horst, Anna S., McDuff, Susan GR., Frascino, Jennifer C., Hopkins, Ramona O., Mauldin, Kristin N. Role of the hippocampus in remembering the past and imagining the future. *Proceedings of the National Academy of Sciences*. 2010; 107:19044–19048.
48. Stachenfeld, Kimberly L., Botvinick, Matthew, Gershman, Samuel J. Design principles of the hippocampal cognitive map. *Advances in Neural Information Processing Systems*. 2014:2528–2536.
49. Kay, Kenneth, Sosa, Marielena, Chung, Jason E., Karlsson, Mattias P., Larkin, Margaret C., Frank, Loren M. A hippocampal network for spatial coding during immobility and sleep. *Nature*. 2016; 531:185–190. [PubMed: 26934224]
50. Muller, Robert U., Kubie, John L. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *The Journal of Neuroscience*. 1987; 7:1951–1968. [PubMed: 3612226]
51. Skaggs, William E., McNaughton, Bruce L. Spatial firing properties of hippocampal cal populations in an environment containing two visually identical regions. *The Journal of Neuroscience*. 1998; 18:8455–8466. [PubMed: 9763488]
52. Alvernhe, Alice, Save, Etienne, Poucet, Bruno. Local remapping of place cell firing in the tolmán detour task. *European Journal of Neuroscience*. 2011; 33:1696–1705. [PubMed: 21395871]
53. Wu, Xiaojing, Foster, David J. Hippocampal replay captures the unique topological structure of a novel environment. *The Journal of Neuroscience*. 2014; 34:6459–6469. [PubMed: 24806672]
54. Hollup SA, Molden S, Donnett JG, Moser MB, Moser EI. Accumulation of hippocampal place fields at the goal location in an annular watermaze task. *Journal of Neuroscience*. 2001; 21:1635–1644. [PubMed: 11222654]
55. Olafsdottir, H Freyja, Barry, Caswell, Saleem, Aman B., Hassabis, Demis, Spiers, Hugo J. Hippocampal place cells construct reward related sequences through unexplored space. *Elife*. 2015; 4:e06063. [PubMed: 26112828]
56. Rudy, Jerry W., Barrientos, Ruth M., O'reilly, Randall C. Hippocampal formation supports conditioning to memory of a context. *Behavioral Neuroscience*. 2002; 116:530–538. [PubMed: 12148921]
57. Byrne, Patrick, Becker, Suzanna, Burgess, Neil. Remembering the past and imagining the future: a neural model of spatial memory and imagery. *Psychological Review*. 2007; 114:340–375. [PubMed: 17500630]
58. Tolman, Edward C. Cognitive maps in rats and men. 1948
59. Russek, Evan M., Momennejad, Ida, Botvinick, Matthew M., Gershman, Samuel J., Daw, Nathaniel D. Predictive representations can link model-based reinforcement learning to model-free mechanisms. *bioRxiv*. 2016:083857.
60. Momennejad, Ida, Russek, Evan M., Cheong, Jin H., Botvinick, Matthew M., Daw, Nathaniel, Gershman, Samuel J. The successor representation in human reinforcement learning. *bioRxiv*. 2016:083824.

61. Heckers, Stephan, Zalesak, Martin, Weiss, Anthony P., Ditman, Tali, Titone, Debra. Hippocampal activation during transitive inference in humans. *Hippocampus*. 2004; 14:153–162. [PubMed: 15098721]
62. Dusek, Jeffery A., Eichenbaum, Howard. The hippocampus and memory for orderly stimulus relations. *Proceedings of the National Academy of Sciences*. 1997; 94:7109–7114.
63. Garvert, Mona M., Dolan, Raymond J., Behrens, Timothy EJ. A map of abstract relational knowledge in the human hippocampal–entorhinal cortex. *eLife*. 2017; 6:e17086. [PubMed: 28448253]
64. Brea, Johanni, Tamás Gáal, Alexis, Urbanczik, Robert, Senn, Walter. Prospective coding by spiking neurons. *PLoS Computational Biology*. 2016; 12(6):e1005003. [PubMed: 27341100]
65. Caporale, Natalia, Dan, Yang. Spike timing-dependent plasticity: a Hebbian learning rule. *Annual Review Neuroscience*. 2008; 31:25–46.
66. Rainer, Gregor, Rao, S Chenchal, Miller, Earl K. Prospective coding for objects in primate prefrontal cortex. *Journal of Neuroscience*. 1999; 19:5493–5505. [PubMed: 10377358]
67. Schapiro, Anna C., Rogers, Timothy T., Cordova, Natalia I., Turk-Browne, Nicholas B., Botvinick, Matthew M. Neural representations of events arise from temporal community structure. *Nature Neuroscience*. 2013; 16:486–492. [PubMed: 23416451]
68. Shohamy, Daphna, Turk-Browne, Nicholas B. Mechanisms for widespread hippocampal involvement in cognition. *Journal of Experimental Psychology: General*. 2013; 142:1159–1170. [PubMed: 24246058]
69. Squire, Larry R. Memory systems of the brain: a brief history and current perspective. *Neurobiology of Learning and Memory*. 2004; 82:171–177. [PubMed: 15464402]
70. Winocur, Gordon, Moscovitch, Morris, Bontempi, Bruno. Memory formation and long-term retention in humans and animals: Convergence towards a transformation account of hippocampal–neocortical interactions. *Neuropsychologia*. 2010; 48:2339–2356. [PubMed: 20430044]
71. Manns, Joseph R., Hopkins, Ramona O., Squire, Larry R. Semantic memory and the human hippocampus. *Neuron*. 2003; 38:127–133. [PubMed: 12691670]
72. Greenberg, Daniel L., Verfaellie, Mieke. Interdependence of episodic and semantic memory: evidence from neuropsychology. *Journal of the International Neuropsychological Society*. 2010; 16:748–753. [PubMed: 20561378]
73. Race, Elizabeth, Keane, Margaret M., Verfaellie, Mieke. Losing sight of the future: Impaired semantic prospection following medial temporal lobe lesions. *Hippocampus*. 2013; 23:268–277. [PubMed: 23197413]
74. Packard, Mark G., McGaugh, James L. Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. *Neurobiology of Learning and Memory*. 1996; 65:65–72. [PubMed: 8673408]

Of special interest

- [44] A functional MRI study demonstrating that the hippocampus is sensitive to long-range temporal clustering of states.
- [11] A theoretical review linking the neurocognitive mechanisms of episodic and semantic memory to machine learning algorithms.
- [4] A computational model of how hippocampal theta oscillations can be used to predict the future, by translating a representation of stimulus history forward in time.
- [63] Evidence from functional MRI adaptation that the similarity structure of hippocampal representations is sensitive to predictive relationships between stimuli.

Of outstanding interest

- [33] An elegant study of recognition memory that directly supports a temporal difference learning update rule for the successor representation.
- [55] Electrophysiological evidence that place cells are sensitive to reward predictions.
- [64] A spiking neuron implementation of the successor representation.

- An updated rational analysis of memory using concepts from graph theory and reinforcement learning.
- A single predictive representation can unify multiple forms of memory.
- The hippocampus may be responsible for encoding this predictive representation.

Figure 1. The successor representation defined over space

(Left) Neurons tuned to different locations encode prediction about future locations. Shading corresponds to expected number of times a neuron's preferred location will be visited along a trajectory initiated at the animal's current location. (Right) Simulated predictive place fields for different environments.

Figure 2. Prospective coding by spiking neurons

(A) A network of spiking neurons (colored triangles), with spike times indicated by vertical lines. The green neuron drives spiking of the purple neuron, opening a potentiation window (red box) during which other spiking neurons can strengthen their synaptic connections with the purple neuron. (B) As a consequence of plasticity, the purple neuron can now be driven by the blue neuron, shifting the potentiation window earlier and allowing the orange neuron to strengthen its connection with the purple neuron. The activity of the green neuron thus becomes predicted by neurons spiking earlier in time. Adapted from [64].