2

What Is a Memory That It Can Be Changed?

Lynn Nadel

Introduction

The concern of this book, the workshop on which it is based, and the theory that propelled it is the nature of memory and how it can be altered in the course of effective therapy. The purpose of this chapter is to set the stage for what is to come in the rest of the book by providing a selective overview of current thinking about memory. We will discuss some features of memory that are particularly important in understanding what is going on when one changes one's mind—or at least the part of it that is responsible for memories, which we will shortly see is pretty much the entire brain.

Already in the late 19th century, Ebbinghaus (1885) had initiated the experimental study of learning, forgetting, and memory, and James (1890) wrote presciently about several forms of memory. James highlighted the distinction between what he called primary and secondary memory, by which he was referring to what is now called short-term and long-term memory. James noted that memory results from the association of ideas—a definition that had major consequences for the study of memory, including in the relatively short term the banishment of memory altogether from the emerging science of psychology as strictly behaviorist approaches dominated from 1915 through the 1940s and early 1950s. There were a few exceptions outside of North America, most notably Bartlett (1932) in England, who presciently described memory as a momentary construction in his book on *Remembering*.

Telescoping history, we jump over the behaviorist perspective that had no need for the notion of memory, straight to the postwar era. Among other things, the war brought into existence the computational devices that have slowly come to dominate our lives and the computational attitude that gave rise to a renewed interest in memory. Even the earliest computers needed memory to function, and the distinction between short-term and long-term memory made by James

Lynn Nadel, *What Is a Memory That It Can Be Changed?* In: *Neuroscience of Enduring Change*. Edited by: Richard D. Lane and Lynn Nadel, Oxford University Press (2020).
DOI: 10.1093/oso/9780190881511.003.0002

applied readily to the information-processing approach suggested by thinking of human brains as just another kind of computer.

One can see in Figure 2.1 all of the components of memory as it was conceived at the time—the important thing to note about this picture is that the various memory boxes differ from one another in terms of how long the memories in that box last. Another thing worth noting is that in the paper first promulgating this model Atkinson and Shiffrin (1968) made direct reference to the famous case of the amnesic patient HM (who had portions of both temporal lobes including the hippocampus removed surgically due to intractable epilepsy) and how his memory impairment supported the distinction between short-term memory, which was intact in HM, and long-term memory, which HM could no longer form.

The picture presented by HM at the time his case was first reported (Scoville & Milner, 1957) is rather different than what he is now taken to show, but in understanding his impact on the emerging field of cognitive neuropsychology in general and the study of memory disorders in particular, we need to consider how memory was conceptualized at that time, the mid 1950s. Early reports on HM were taken as showing that the medial temporal areas resected in his surgery were critically important in a *consolidation* process that transformed short-term into long-term memories. That is, HM was assumed to have intact sensory memory, short-term memory, and also long-term memory in that he seemed capable of retrieving autobiographical memories from his life before surgery. What was clearly impaired, in a devastating way, was HM's ability to transform new experiences into long-term memories. We now know that HM's remote memory was nowhere near as good as the early investigations mistakenly concluded, which has implications for thinking about the role of the hippocampus in long-term memory.

Shortly after the first reports on HM, a team of neuropsychologists, including Brenda Milner, who led the study of HM, joined forces to use what was known about HM's surgery to try to create an animal model of amnesia. Orbach, Milner and Rasmussen (1960) recreated HM's lesion in a group of rhesus monkeys and tested these animals on a variety of tasks they imagined would be comparable to the memory tasks HM failed completely. Unfortunately, monkeys with what

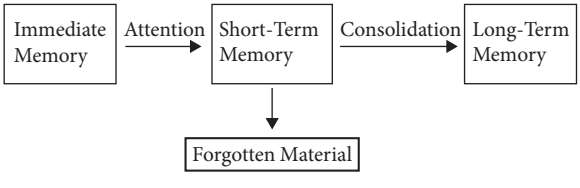

Figure 2.1 Memory in the 1960s.

seemed like comparable lesions did not evince any memory loss. Some years later the same result emerged from studies with rats—lesions in the hippocampus did not seem to cause anything like the memory defect observed in HM and other amnesics (e.g., Kimble, 1963). This led to a decade of relative chaos in the world of hippocampal research, with some entertaining the possibility that this core temporal lobe structure might have a rather different function in humans than it has in other animals.

This quite unsatisfying state of affairs was jolted by the discovery of "place cells" in the hippocampus of freely moving rats (O'Keefe & Dostrovsky, 1971), an experiment that changed the field and ultimately earned John O'Keefe a share of the Nobel Prize in 2014. This discovery served as the foundation for the idea that the hippocampus served as the neural core of "cognitive maps" and that this gave it a critical function in a certain kind of spatial cognition, as well as a particular kind of memory, the kind that depends upon spatiotemporal context (Nadel & O'Keefe, 1974; O'Keefe & Nadel, 1978). This selectivity helps us understand, in retrospect, why the human and animal data did not go together. In brief, the tasks that were used in these different species were not fully equivalent—the tasks devised for animals had not paid attention to the core spatial/contextual function of this structure. Animals with damage to the hippocampus could solve these tasks using other memory systems since spatial learning wasn't typically part of the experiment.

The cognitive map theory of hippocampal function had a number of major implications for understanding how the brain organizes memories, including the idea that there are multiple forms of memory that differ as a function of the nature of the information being stored. According to this view, these different kinds of memory are undergirded by distinct neural networks with unique operating characteristics, likely have different developmental trajectories, are differentially affected by such things as stress and emotion, and have different plasticity functions. The latter highlights the fact that different kinds of memory can be learned, and unlearned, at different rates. In some systems learning will be fast, and unlearning as well. In others, learning and unlearning might be slow.

In addition to providing a way of thinking about the apparent discrepancy between human and animal studies of the hippocampal role in memory, the notion of multiple memory systems led, in the early 1980s, to advances in thinking about consolidation—the memory "fixing" process attributed to the hippocampus after the study of HM. What has come to be known as the standard model of memory consolidation (SMC) proposed that control over the expression of both episodic and semantic memories passes from hippocampus to neocortex during consolidation—initially the hippocampus is needed to retrieve from neocortex the various parts of a memory, but over consolidation time, the neocortex

acquires the ability to do this without hippocampal intervention (Squire et al., 1984; Alvarez & Squire, 1994).

Problems with this view, however, centered around the question of what happens to remote memories in amnesia and whether semantic and episodic memories suffer the same fate in amnesia. Resolution of the former question emerged only when the initial observations of HM were replaced by more refined assays of autobiographical memory in HM and other amnesics. Such studies, beginning to emerge in the 1990s, suggested that the retrograde amnesia seen in various clinical cases went further back in time than one might have expected based on SMC and, further, that any remote memories amnesics could retrieve were atypically impoverished in characteristic ways, lacking episodic details.

Concerns about the extent of remote memory deficits, and the relative impact on episodic and semantic memory, led to the development of multiple trace theory (MTT) in the late 1990s (Nadel & Moscovitch, 1997). MTT proposed that the hippocampus was *always* involved in the retrieval of suitably vivid and detailed episodic memories, in contrast to what SMC had proposed. The advent of neuroimaging methods made it possible to test this idea directly, and from the outset the data have uniformly supported the MTT perspective—the recall of remote episodic memories activates the hippocampus (Nadel et al., 2000; Ryan et al., 2001), in particular if they are vivid (Robin & Moscovitch, 2017).

MTT made another prediction that is particularly important for the science of memory change—namely, that when a memory is reactivated it is likely to be re-encoded in a changed form. Some years after the promulgation of MTT memory "reconsolidation" burst upon the field—evidence showed that re-activating a memory rendered it labile and open to change. This was the reawakening of an old idea, perhaps first noted by Freud (1895/1966) in 1896, who used the term "retranscription" to refer to a process during therapy that sounds a lot like reconsolidation in modern terms. Whether or not changes actually occur, the reactivated and now labile memory has to be stabilized again, and this "reconsolidation" process takes time, although likely less time than original consolidation. The research that brought reconsolidation back to the fore included work with both mazes (Przybyslawski & Sara, 1997), and fear learning (Nader et al., 2000). The latter in particular stoked interest in the psychiatric community, as the potential for erasing fearful or other negative memories was clear.

The phenomenon of reconsolidation calls into question the original thinking behind the notion of memory consolidation. The time-limited nature of consolidation, and the notion that after that consolidation period a memory was permanently "fixed," were fundamental facts about memory for proponents of this view. If memory consolidation is about something other than making memories permanently stable, what might that be? Before answering that question, it is useful to consider the adaptive advantage of having a malleable memory

system. If memory is only about reliving the past, then malleability is a bad idea. However, if memory is to be used to guide behavior in the future, malleability is a very good idea. Being able to flexibly manipulate our memories helps us plan for alternate futures based on past outcomes. Being able to update memories allows us to take advantage of the most current information about the environment when consulting the past to inform action in the present and future.

What then is consolidation for? In addition to modifying neural connections representing the specifics of a particular episode, consolidation allows for the integration of information gleaned during such episodes into pre-existing knowledge networks. These networks represent both the things we've experienced in the world and the concepts, categories, schemas and scripts that have been abstracted from experience over our entire life. Adaptive behavior is generated by flexibly utilizing information in both the episodic system, informing us about specific instances in our life, and in the other systems, informing us about regularities that govern both the world and how we should behave in it if we are to thrive. This is consistent with the notion that consolidation serves the purposes of integrating episodic memories with semantic memories or schemas. The kinds of change sought in therapy will likely engage all of these systems, as a goal of psychotherapy is to create momentary experiences that will lead to lasting changes in semantic memories or schemas. This means we must understand how different forms of memory are stored in the brain, so that we might better understand how best to induce enduring change.

Forming and Re-Forming Memories: A Memory Systems Perspective

Forming memories in the brain involves changing effective connectivity between neurons in an ensemble; the nature of these ensembles likely differs as a function of the kind of memory in question. Without going into detail, it is important to point out the statistical nature of the connections we are talking about. A memory cannot be merely a series of direct links connecting a chain of specific neurons. Instead, it is a pattern of activity across multiple neurons, organized in re-entrant circuits, each cycle of which could in principle be executed by a different subset of neurons. The informational content of these neural ensembles remains constant across minor variations in which neurons are sustaining the pattern at any particular time, and that is what matters for present purposes. As we have noted, successful therapy is likely to involve changes in multiple systems, so we must now turn to a discussion of several types of memory critical to the therapeutic enterprise.

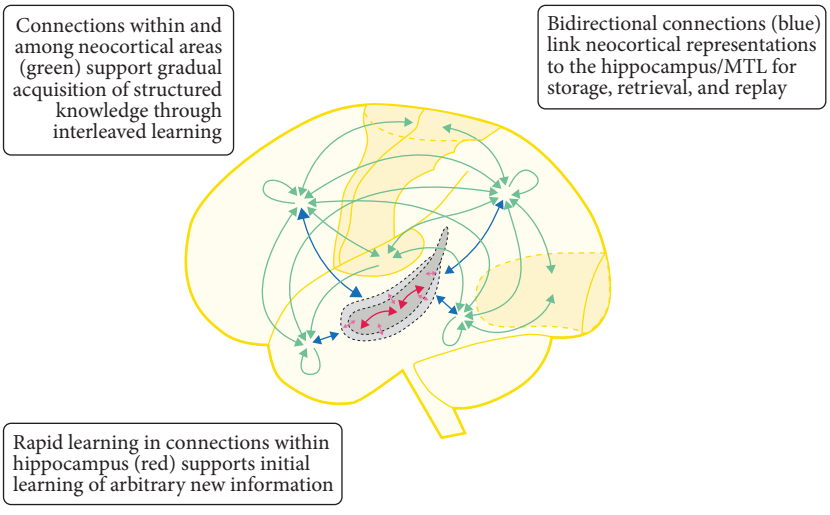

Figure 2.2 Episodic memory circuits.

Episodic Memory

Episodic memory is about singular events, occurring in specific spatio-temporal contexts, involving the things, people, and actions comprising the event. Events are multidimensional and engage widely dispersed neural ensembles. Such widely distributed neurons cannot all be directly connected—so how can an event involving each of them be memorized? The hippocampus plays a critical role here. Elements (clusters of neurons one assumes) in disparate cortical zones activate elements in the hippocampus, which rapidly link to each other, forming a template that indexes the distributed bits and pieces comprising the details of the memory.

Thus, an episodic memory trace incorporates linkages among cortical neurons, among hippocampal neurons, and between hippocampal and cortical neurons (in both directions) as captured in a recent review by Kumaran et al. (2016), from whom Figure 2.2 was adapted. Changing such a memory could, in principle, involve changing any or all of the linkages that collectively comprise that memory.

Episodic Memory Reactivation

Episodic memories can be reactivated, triggered by some aspect of the current situation that serves to "cue" that older memory, by activating a subset of the

cortical elements that are part of that memory. These cortical elements activate the hippocampal elements to which they connect. Activation of these elements initiates a process of pattern completion within the hippocampus, activating the remaining elements of the hippocampal "template" representing that episode. This template then "retrieves" the cortical elements that were not part of the "cue" but that were part of the old memory. This retrieved episode may, or may not, reach conscious awareness.

Forming this kind of memory requires plasticity between widely separated elements, a challenge because such plasticity in the connections between two elements is only possible if they are active *together* within a very narrow time window—perhaps as short as 5 to 10 msec. This coordination among dispersed neurons is thought to be accomplished by neurons firing in phase with each other, which creates a brief slice of time when plasticity between them is enabled (see Çalışkan & Stork, 2018, for a review). Within this very brief window, neurons that "fire together" can "wire together," as Hebb (1949) suggested. In sum, reactivating a memory seems to involve a multistep process, both bottom–up and top–down, that ultimately results in a fleshed out recalled memory, of which one could become conscious.

Changing An Episodic Memory

At first blush, it seems that updating an episodic memory should involve adding (and/or subtracting) elements from the ensemble of connected elements that comprise the reactivated memory. **It is an important fact that this only happens to a reactivated memory**. Episodic memories that are not reactivated are not changed or changeable. But is reactivation alone enough to change a memory representation? Apparently not, and that is where prediction, and prediction error (PE), are thought to come in.

Predictions (Expectations) and Prediction Error

Reactivated memories, fleshed out through pattern completion, generate "predictions"—for example, what I might see if I turn the corner in a familiar house I've lived in for years. Such predictions are "compared" with what actually happens, generating varying amounts of "mismatch" or what is called "prediction error." There is what appears to be an inverted U-shaped function regarding PE in episodic memory—updating only occurs at a particular level. Too little PE likely results in no change, or a strengthening of existing linkages. Too much PE and pattern separation processes take hold, forcing the creation of a distinct

memory, engaging a different ensemble of neurons. This system-level "decision" is critical in the therapeutic context. It is essential, as we will see, to stay out of the part of the PE curve where pattern separation dominates. This would make it less likely that anything learned in the therapist's office would generalize to situations in the real world.

My colleagues and I developed a human episodic memory paradigm to study reconsolidation under conditions where updating might be possible. Participants learn a list of objects (Set 1) in a particular context (Room A). Two days later they come back to either Room A or a new place, Room B. Those brought back to Room A are further reminded of their previous learning of Set 1 by having the same experimenter and being asked a leading question. Those brought to Room B are met by a new experimenter and no questions at all. We expected that the participants brought back to the same context would reactivate their memory of Set 1 and its objects. Both groups of participants learned a new list of objects (Set 2). Our thought was that when the memory of Set 1 was reactivated just as Set 2 was being learned, some of the new objects in Set 2 might be incorrectly added to the activated memory of Set 1. That is, the reactivated memory would be changed by current experience. And this is precisely what happened.

As Figure 2.3 shows at the top, participants in the Reminder group reported a significant number of Set 2 items (red bar) as having belonged to Set 1. This was not observed in the participants who did not have their memory of Set 1 learning reactivated. Figure 2.3 shows that it was the spatial context, and not the experimenter or leading question, that produced the memory reactivation essential to this memory updating effect (Hupbach et al., 2007).

In a series of further studies (Hupbach et al., 2008, 2009), we explored the various conditions under which memory updating occurs and does not occur. We showed, for example, that context played a critical role only if it was relatively novel. In familiar environments other cues take on the critical reactivating role.

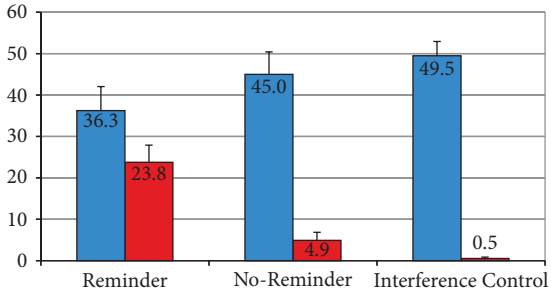

Figure 2.3 Extent of updating observed in Hupbach et al. (2007) reconsolidation study.

Another early (unpublished) finding showed that a full retrieval was not followed by updating of List 1 with List 2 items, a result that might reflect the role of prediction error in enabling memory updating. Other investigators have studied the presence, or absence, of memory reconsolidation in humans, with somewhat mixed results, as another contribution to this volume make clear (chapter 10).

Most recently, we looked at memory reconsolidation in a variant of our paradigm that we could bring into the scanner (Simon et al., 2017). We showed that whether or not a memory was updated and which specific items from Set 2 were recalled incorrectly as having been on Set 1 were both related to the extent of activity in a particular brain region, the temporo-parietal junction (TPJ). The more activity in this region the less likely was memory updating to occur. We interpreted this in terms of the well-established role of the TPJ in monitoring the extent of retrieved details of a memory. Our results best fit within the "predictive" framework sketched out earlier. Extensive prediction error, which is most likely when memory reactivation is highly detailed, leads to a pattern separation outcome, wherein the reactivated memory and the newly forming memory for the current state are kept separate and distinct. A modest amount of prediction error might lead to a different outcome—the system deciding that the reactivated memory and the current state are somehow part of the same "event," leading to an updating of the reactivated memory with some of the contents of the current situation. That is, some Set 2 items are attached to the reactivated Set 1 memory instead of being included in the emerging Set 2 memory.

The work from our lab and others allows us to summarize some of the key features of episodic memory reconsolidation and updating as it is understood now:

1. Reactivation is necessary for updating to occur. This can be triggered by something in the current situation that is similar to a previously stored memory, perhaps a person, or some salient object/event shared by the prior experience and the present circumstance. It can also be triggered internally, during bouts of mind-wandering (see Chapter 5 of this volume), or while asleep (see Chapter 7 of this volume). The emotional state you are in can contribute to which memories you will reactivate or retrieve.
2. Reactivated memories can vary along several dimensions. They can vary in detail—some memories are highly detailed, others less so. They can vary in strength—some strong, some weak. They can vary in terms of emotional salience. Each of these reflects activation in different networks. All of these will affect the possibility of a reactivated memory being altered by what is happening in the present.
3. Reactivation can be followed by any of several outcomes. At issue first is the question of whether reactivation is necessarily followed by memory

destabilization, which appears to be determined by the ratio of two subunits of the NMDA receptor (see reviews by Zhang et al., 2018; Vigil & Giese, 2018). Subunit GluN2B regulates memory destabilization—when this subunit is inhibited, reactivation-induced destabilization and the memory-erasing effects of protein-synthesis inhibitors are not seen. GluN2A regulates memory restabilization. The ratio between N2A and N2B determines what will happen when a "memory" is reactivated. One variable that might influence this ratio is amount of "training" (Lee et al., 2017). There is evidence that strong memories are associated with decreased N2B receptors and hence are less likely to destabilize and be updated, even in the presence of substantial PE. Should this be the case, one way of facilitating change when quite strong memories are involved might be to pharmacologically manipulate N2B receptor levels. Such an approach is complicated by the apparent fact that control over destabilization might rest with different modulators in different brain systems. There is some evidence, for example, that cannabinoid receptors control destabilization in some hippocampal circuits (Nasehi et al., 2017), while serotonin receptors play such a role in the perirhinal cortex (Morici et al., 2018).

4. Assuming a memory is destabilized at the cellular level, updating it demands prediction error—in the absence of PE there is nothing to update. In the absence of PE reactivated memories would return to their pre-reactivation state eventually, although during the time they are destabilized they are open to being strengthened or weakened, or even erased.

Semantic, Conceptual and Action-Based Memory Systems

What was just described applies to explicit memories that engage widely dispersed neural elements—namely, specific episodes. Matters are different when one considers other kinds of memory, for example, semantic memories. Some are explicit, such as facts about the world or about one's personal history; others are implicit, such as concepts and categories, and action schemas that tie these to behaviors through scripts.

Most agree that semantic memories, facts, scripts, and schemas engage neural systems outside the hippocampal system, in cortical and subcortical regions. Many also subscribe to the idea that learning is very rapid in hippocampal circuits and less so in extra-hippocampal circuits, an idea suggested by Nadel & O'Keefe (1974; O'Keefe & Nadel, 1978), and based to some extent on the discovery in the 1970s of long-term potentiation in rabbit hippocampus (Bliss & Lomo, 1973). This idea was formalized by McClelland et al. (1995) in the complementary learning systems theory, which asserted that because of the way

knowledge is represented in cortical circuits—in overlapping codes, rather than in sparse, orthogonalized codes—effective change must be made slowly. Rapid change would alter the delicate balance of synaptic weights that is required to store multiple items in overlapping sets of neurons, with the consequence of catastrophically interfering with previously stored knowledge.

While it is now clear that rapid plasticity can occur outside the hippocampus, for the most part the notion that slow changes are most effective in the systems storing conceptual and schematic knowledge has been retained. This means that scripts and schemas, including well-established action routines, might not be as easily changed and updated as is the case for episodic memories. This makes sense, intuitively—episodic knowledge is based on a single experience, by definition. Semantic knowledge is the aggregate of multiple experiences and hence more likely to provide valid predictions on which to base current and future actions. It should be harder to change, and the slow change in synaptic weights that can occur each time the updating circuit is activated accomplishes this goal. But, there is a caveat in that during sleep, and the replay observed in slow-wave sleep, it might be possible for the hippocampal episode system to send hundreds, or even thousands, of training signals to the cortex, making it possible that slow change on a trial-by-trial basis need not translate into slow change in the real world. Replay during sleep can translate slow physiological alterations into rapid psychological change. Recent work using targeted memory reactivation (e.g., Simon et al., 2018) raises the possibility that this replay process can be influenced in ways that can both strengthen or weaken memory traces.

As noted earlier, memory consolidation includes processes that integrate new experiences into established knowledge networks. Work in both animal models (Tse et al., 2007) and humans (Fernandez & Morris, 2018) shows that new experience is *more rapidly* integrated into extra-hippocampal circuits when it fits with the knowledge one has previously acquired. These may be the conditions under which fast acquisition of new information can occur in neocortical circuits. Research that focuses on whether new experience is schema-congruent or schema-incongruent shows that the brain reacts differently to these two kinds of experiences. Schema-congruent inputs predominantly activate the medial PFC, while schema-incongruent inputs predominantly activate the medial temporal lobe, including the hippocampus. The latter case is relevant to the therapeutic situation—new information is introduced that is schema-incongruent, with the hope that it will lead to a reformulation of the schema and hence changes in behavior predicated on it. The fact that this situation activates the hippocampal system suggests several things. First, substantial change in schemas is either orchestrated, or strongly impacted, by the fast-learning episodic system. Second, preferential engagement of the hippocampal system could have an impact on the perspective from which one views one's experience, shifting between

first- and third-person views. Which perspective a memory is recalled from clearly matters (see Chapter 4 of this volume). Finally, to the extent to which this episodic driving of schema-change is critical to enduring change, it is essential that the therapeutic situation be one in which stress is managed at levels below those that would downregulate hippocampal function, (e.g., Kim & Diamond, 2002). And it is also likely important that sleep is well managed, as previously hinted at.

The literature on change in semantic systems also suggests that one way such networks can react to the addition of incongruent information is to split off new subschemas (or subcategories). My schema for "apple" included "has red skin" until I discovered Granny Smiths, at which point my schema expanded to include red and green types, both nested under the enlarged category of APPLE. Understanding how such splitting and reconfiguring of schemas plays out with complex and often abstract, beliefs about oneself and the world is a task for the future, because many of the critical properties of schema-based knowledge systems remain to be carefully studied. It is an open question whether the splitting of semantic networks is at all akin to the way in which memories are kept separate in the episode system.

We return to some of these issues in our final chapter of the book, where we focus on a future research agenda.

In sum, a few points about memory should be emphasized: (a) there are many forms of memory, and the rules governing the storage, retrieval, and updating of these various forms likely differ—some forms of memory may be relatively easy to change; others (like habits), less so; (2) memory reactivation and updating, sometimes called "reconsolidation" can apparently result in changed memories, but the exact conditions under which this happens are unclear at present; (3) interactions between memory and emotion are critical, affecting the likelihood of a memory being made in the first place, or remade after reactivation; and (4) predictive coding is at the core of what memory exists to do; therefore, computational approaches to prediction (e.g., Bayesian modeling) could inform how memories are remade and hence be of use in therapy.

References

Alvarez, P., & Squire, L. R. (1994). Memory consolidation and the medial temporal lobe: A simple network model. *Proceedings of the National Academy of Sciences*, *91*, 7041–7045.

Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation: II* (pp. 89–195). Oxford, England: Academic Press.

Bartlett, F. C. (1932). *Remembering*, Cambridge, England: Cambridge University Press.

Bliss, T., & Lømo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *Journal of Physiology.* (*Lond.*), *232*, 331–356.

Çalışkan, G., & Stork, O. (2018). Hippocampal network oscillations as mediators of behavioural metaplasticity: Insights from emotional learning. *Neurobiology of Learning and Memory*, *154*, 37–53.

Ebbinghaus, H. (1885). *Memory: A contribution to experimental psychology*. New York, NY, Dover Press.

Fernández, G., & Morris, R. G. (2018). Memory, novelty and prior knowledge. *Trends in Neurosciences*, *41*, 654–659.

Freud, S. (1966). Letter 52 from extracts from the Fliess papers. In J. Strachey (Ed.), *The standard edition of the complete psychological works of Sigmund Freud: Vol. 1, 1886–1899: Pre-psycho-analytic publications and unpublished drafts* (pp. 233–239). London, England: Hogarth Press. (Work originally published 1895)

Hebb, D. (1949). *The organization of behavior*. New York, NY: Wiley.

Hupbach, A., Gomez, R., & Nadel, L. (2009). Episodic memory reconsolidation: Updating or source confusion? *Memory*, *17*, 502–510.

Hupbach, A., Gomez, R., Nadel, L. & Hardt, O. (2007). Reconsolidation of episodic memories: A subtle reminder triggers integration of new information. *Learning and Memory*, *14*, 47–53.

Hupbach, A., Hardt, O., Gomez, R. & Nadel, L. (2008). The dynamics of memory: Context-dependent updating, *Learning and Memory*, *15*, 574–579.

James, W. (1890). *The principles of psychology*. New York, NY: Henry Holt.

Kim, J. J., & Diamond, D. M. (2002). The stressed hippocampus, synaptic plasticity and lost memories. *Nature Reviews Neuroscience*, *3*, 453–462.

Kimble, D. P. (1963). The effects of bilateral hippocampal lesions in rats. *Journal of Comparative and Physiological Psychology*, *56*, 273–283.

Kumaran, D., Hassabis, D., & McClelland, J. L. (2016). What learning systems do intelligent agents need? Complementary learning systems theory updated. *Trends in Cognitive Sciences*, *20*, 512–534.

Lee, J. L., Nader, K., & Schiller, D. (2017). An update on memory reconsolidation updating. *Trends in Cognitive Sciences*, *21*, 531–545.

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*(3), 419–457.

Morici, J. F., Miranda, M., Gallo, F. T., Zanoni, B., Bekinschtein, P., & Weisstaub, N. V. (2018). 5-HT2a receptor in mPFC influences context-guided reconsolidation of object memory in perirhinal cortex. *eLife*, *7*, e33746.

Nadel, L., & Moscovitch, M. (1997). Consolidation, retrograde amnesia and the hippocampal formation. *Current Opinion in Neurobiology*, *7*, 217–227.

Nadel, L., & O'Keefe, J. (1974). The hippocampus in pieces and patches: an essay on modes of explanation in physiological psychology. In R. Bellairs & E. G. Gray (Eds.), *Essays on the nervous system*: *A festschrift for J. Z. Young* (pp. 367–390). Oxford, England: Clarendon Press.

Nadel, L., Samsonovitch, A., Ryan, L., & Moscovitch, M. (2000). Multiple trace theory of human memory: Computational, neuroimaging and neuropsychological results. *Hippocampus*, *10*, 352–368.

Nader, K., Schafe, G. E., & LeDoux, J. E. (2000). Fear memories require protein synthesis in the amygdala for reconsolidation after retrieval. *Nature*, *406*, 722–726.

Nasehi, M., Hajikhani, M., Ebrahimi-Ghiri, M., & Zarrindast, M. R. (2017). Interaction between NMDA and CB2 function in the dorsal hippocampus on memory consolidation impairment: An isobologram analysis. *Psychopharmacology*, *234*, 507–514.

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, *34*, 171–175.

O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford: The Clarendon Press.

Orbach, J., Milner, B., & Rasmussen, T. (1960). Learning and retention in monkeys after amygdala-hippocampus resection. *Archives of Neuro*logy, *3*, 230–251.

Przybyslawski, J., & Sara, S. J. (1997). Reconsolidation of memory after its reactivation. *Behavioural Brain Research*, *84*, 241–246.

Robin, J., & Moscovitch, M. (2017). Details, gist and schema: Hippocampal–neocortical interactions underlying recent and remote episodic and spatial memory. *Current Opinion in Behavioral Sciences*, *17*, 114–123.

Ryan, L., Nadel, L., Keil, K., Putnam, K., Schnyer, D., Trouard, T., & Moscovitch, M. (2001). The hippocampal complex is equally involved in retrieving recent and very remote autobiographical memories: Evidence from functional magnetic resonance imaging in neurologically intact people. *Hippocampus*, *11*, 707–714.

Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesion. *Journal of Neurology, Neurosurgery and Psychiat*ry, *20*, 11–21.

Simon, K. C., Gómez, R. L., & Nadel, L. (2018). Losing memories during sleep after targeted memory reactivation. *Neurobiology of Learning and Memory*, *15*(1), 10–17.

Simon, K. C., Gómez, R. L., Nadel, L. & Scalf, P. E. (2017). Brain correlates of memory reconsolidation: a role for the TPJ. *Neurobiology of Learning and Memory*, *142*, 154–161.

Squire, L. R., Cohen, N. J., & Nadel, L. (1984). The medial temporal region and memory consolidation: a new hypothesis. In H. Weingartner & E. Parker (Eds.), *Memory consolidation* (pp. 185–210). Hillsdale, NJ: Erlbaum, 1984.

Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. A., Wood, E. R., . . . Morris, R. G. (2007). Schemas and memory consolidation. *Science*, *316*, 76–82.

Vigil, F. A., & Giese, K. P. (2018). Calcium/calmodulin-dependent kinase II and memory destabilization: a new role in memory maintenance. *Journal of Neurochemistry*, 147(1), 12–23.

Zhang, J. J., Haubrich, J., Bernabo, M., Finnie, P. S. B., & Nader, K. (2018). Limits on lability: Boundaries of reconsolidation and the relationship to metaplasticity. *Neurobiology of Learning and Memory*, *154*, 78–86.