2014

Electrophysiological evidence for memory schemas in the rat hippocampus

https://hdl.handle.net/2144/15122

BOSTON UNIVERSITY

GRADUATE SCHOOL OF ARTS AND SCIENCES

Dissertation

ELECTROPHYSIOLOGICAL EVIDENCE FOR MEMORY SCHEMAS IN THE RAT HIPPOCAMPUS

by

SAMUEL McKENZIE

B.Sc., McGill University, 2007

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

2014

Approved by

First Reader ____
Howard Eichenbaum, Ph.D.
Professor of Psychology

Second Reader ____
Michael Hasselmo, Ph.D.
Professor of Psychology

Third Reader ____
Chantal Stern, Ph.D.
Professor of Psychology

Dedicated to Ricki and Stewart McKenzie

ACKNOWLEDGMENTS

First I would like to acknowledge Howard Eichenbaum. Throughout my time at Boston University, Howard has been extremely supportive providing a combination of freedom and insight that has transformed me from someone with ideas about how the brain works into someone who can test those ideas. His door has always been open and the conversations that have followed his friendly, “What’s cookin’?” will continue to shape how I relate neurobiological observations to cognitive theory.

I would also like to acknowledge the contributions of my committee members. Mike Hasselmo introduced me to the concept of neural networks in 2008 in a fantastic course on models of memory. I will carry his emphasis on model guided research throughout the next stages of my scientific career. I would also like to thank Chantal Stern for allowing me to guest lecture in her Master’s level neuroscience of memory course. I gained a broader understanding of memory systems in that course than in any other period of my life. Marc Howard has offered almost clairvoyant suggestions as to the right type of analysis to address the neural patterns that I would bring to his office. His insights and skepticisms have been extremely useful. Finally, I would like to thank Adam Johnson for being excited about my experiments; I look forward testing whether my data can be described by his models of schema development.

I have had many other fantastic mentors along the way. I highly value the emphasis on multiple memory systems that I learned while working with Veronique

Bohbot in Montreal. I am grateful for the opportunity that Paul Lipton provided me in running the lab portion of NE203. I would also like to acknowledge the Cold Spring Harbor memory systems course. My current understanding of memory can be traced back to those two weeks in the summer of 2009.

All of my technical skills and many of my theoretical ones I owe to Howard's lab members. In particular I would like to thank: Norbert Fortin, Ben Kraus, Rob Komorowski, Chris MacDonald, Dan Salz, Nick Robinson, Jon Rueckemann, Nat Kinsky, Lara Rangel, Anja Farovik, and Dan Sheehan for teaching me what I know, for thoughtful conversations, and for their contributions to the experiments described in this thesis. Andrea Frank collected that data recorded in reported in Chapter 4 and I greatly appreciate the amount of work she contributed to the project. I would also like to thank my students, Lauren Herrera, Jordana Churchill and Pam Riviere who have helped immensely over the years.

I have had great teachers throughout my life starting with my parents, Ricki and Stewart. They taught me the power of creativity and observation in problem solving. Aside from being great teachers, they have also given me so many opportunities to have the best possible education and a happy life. Thank you both.

ELECTROPHYSIOLOGICAL EVIDENCE FOR MEMORY SCHEMAS IN THE RAT HIPPOCAMPUS

(Order No.)

SAMUEL McKENZIE

Boston University Graduate School of Arts and Science, 2014

Major Professor: Howard Eichenbaum, Professor of Psychology

ABSTRACT

According to Piaget and Bartlett, learning involves both assimilation of new memories into networks of preexisting knowledge and alteration of existing networks to accommodate new information into existing schemas. Recent evidence suggests that the hippocampus integrates related memories into schemas that link representations of separately acquired experiences. In this thesis, I first review models for how memories of individual experiences become consolidated into the structure of world knowledge. Disruption of consolidated memories can occur during related learning, which suggests that consolidation of new information is the reconsolidation of related memories. The accepted role of the hippocampus during memory consolidation and reconsolidation suggests that it is also involved in modifying appropriate schemas during learning. To study schema development, I trained rats to retrieve rewards at different loci on a maze while recording hippocampal calls. About a quarter of cells were active at multiple goal sites, though the ensemble as a whole distinguished goal loci from one another. When new goals were introduced, cells that had been active at old goal locations began firing at the new locations. This initial generalization decreased in the days after learning. Learning also caused changes in firing patterns at well-learned goal locations. These results suggest that learning was supported by modification of an active schema of spatially related reward loci. In another experiment, I extended these findings to explore a schema of object and place associations. Ensemble activity was influenced by a hierarchy of task dimensions which included: experimental context, rat's spatial location, the reward potential and the identity of sampled objects. As rats learned about new objects, the cells that had previously fired for particular object-place conjunctions generalized their firing patterns to new conjunctions that similarly predicted reward. In both experiments, I observed highly structured representations for a set of related experiences. This organization of hippocampal activity counters key assumptions in standard models of hippocampal function that predict relative independence between memory traces. Instead, these findings reveal neural mechanisms for how the hippocampus develops a relational organization of memories that could support novel, inferential judgments between indirectly related events.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

2DG	2-deoxyglucose
AMPAR	α-Amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid receptor
ANOVA	Analysis of Variance
AP	Anterior/posterior
BIC	Binding of Items and Contexts
BOLD	Blood Oxygen Level Dependent
C	Celsius
CA1	Cornu Ammonis 1
CA3	Cornu Ammonis 3
cm	Centimeter
CPA	Cross-place Analysis
CPVSI	Context Position Valence Set Item
CS	Conditioned Stimulus
DG	Dentate Gyrus
DMS	Delayed Match to Sample
DNMS	Delayed Non-match to Sample
DV	Dorsal/Ventral
EC	Entorhinal Cortex
ECII	Layer II of Entorhinal Cortex
fMRI	function Magnetic Resonance Imaging
Hz	Hertz
IEG	Immediate Early Gene
IQ	Intelligence Quotient
Kg	Kilogram
kHz	Kilohertz
LEC	Lateral Entorhinal Cortex
LED	Light Emitting Diode
LFP	Local Field Potential
μA	Microamp
μL	Microliter
μm	Micrometer
MEC	Medial Entorhinal Cortex
Mg	Milligram
ML	Medial/Lateral
mm	Millimeter
mPFC	medial Prefrontal Cortex
ms	millisecond
MVPA	Multi Voxel Pattern Analysis
NiCr	Nichrome
NMDAR	N-Methyl-D-aspartic Acid Receptor
PCA	Principal Components Analysis
PETH	Peri-event Time Histogram
ROC	Receiver Operator Characteristic
s	Seconds
S.C.	Subcutaneously
SI	Selectivity Index
US	Unconditioned Stimulus
VMPFC	Ventromedial Prefrontal Cortex

Reappearance presumes disappearance, it may not be nice
Or proper or easily analysed not to be static
But none of your slide snide rules can catch what is sliding so fast
And, all you advisers on this by the time it is that,
I just do not want your advice
Nor need you be troubled to pin me down in my room
Since the room and I will escape for I tell you flat:
One cannot live in the same room twice.

-Louis MacNeice

1 General Introduction

The fundamental challenge for a memory system is to recognize past experiences that have utility for the present. While no two situations are identical – we can never live in the same room twice – clearly we can use acquired knowledge to make inferences and devise rules to apply in novel situations. Though little is known about the neural computations that allow for the interaction of memories formed at different times, several lines of research indicate an important role for the hippocampus in this process.

Past experience can shape future behavior in many ways and in the standard taxonomy of memory (Squire and Zola, 1996), different brain systems serve distinctive computational roles. Typically, the hippocampus is thought to be responsible for the formation and retrieval of declarative memory, or the information which is available for conscious expression. In particular, the field is unanimous that the hippocampus is required for the conscious recollection of life events, known as episodic memories. Models of how the hippocampus contributes to episodic memory emphasize its role in the creation of distinct memory traces that do not interfere or interact with one another, at least within the intra-hippocampal circuitry.

In addition to episodic memory, the classic memory taxonomy of declarative memory also includes our knowledge of facts and rules, our semantic memory. This is the information that has lost its spatiotemporal source and is divorced from a feeling of personal involvement, or an autonoetic awareness. While it is clear that episodic memory and semantic memory are mutually influential, it is unknown how episodic memories combine to form semantic memory, and accordingly what is the role of the hippocampus in the formation and retrieval of semantic memory.

Recently, the hippocampus has been implicated in the formation of memory schema, the organization of related knowledge that contains semantic knowledge *as well as* episodic details. The hippocampal role in forming memory schema is to build relational networks of episodic memories. Under this model, memory for the common elements of different memories, often defined as semantic memories, can be used to associatively bridge the memory for the unique aspects of related experiences (Eichenbaum, 1999; Cohen and Eichenbaum, 1993; McKenzie and Eichenbaum, 2011).

I will focus on two conflicting roles of the hippocampus: one account that emphasizes the storage of episodes with minimal interference and the other that emphasizes the storage of episodes to allow for their associative interaction. In the remainder of this chapter, I will review evidence for both accounts of hippocampal function and describe models that describe each theoretical role. In Chapter 2, I will discuss the findings from studies on systems consolidation and memory reconsolidation and argue for a critical role of the hippocampus in schema formation. In Chapter 3, I will present findings from an electrophysiological study of the hippocampus during the formation of a spatial schema, providing evidence to support key predictions of a schema account of memory. In Chapter 4, I will present neural evidence for a hippocampal schema for item and context memories. Finally, in Chapter 5, I will discuss my experimental results within the context of the two divergent hypotheses of the hippocampus. I propose that the primary goal of the hippocampus is to realize novel relationships between related memories by linking experiences that share common features, rather than storing minimally-interfering memory traces, as current dogma holds.

1.1 Episodic memory is dependent upon hippocampal function

1.1.1 Hippocampal damage in amnestic patients causes episodic memory deficits

While there were hints of a relationship between the hippocampus and episodic memory in the late 19th and early 20th century (e.g. von Bechterew, 1900; Glees and Griffith, 1952), it was Brenda Milner's seminal work on patient Henry Molaison (H.M.) that ignited decades of research as to how the hippocampus contributes to the conscious recollection of prior experience. As an experimental treatment for intractable epilepsy, H.M. received bilateral resectomy of the hippocampus, amygdala and surrounding medial temporal lobe cortices including the entorhinal cortex, parahippocampal cortex and piriform cortex. After the operation, Milner noted that H.M. could "no longer recognize the hospital staff nor find his way to the bathroom, and he seemed to recall nothing of the day-to-day events of his hospital life" (Milner and Scoville 1957).

Over 50 years of research on H.M. revealed that his anterograde declarative memory for stimuli of all modalities was severely impaired (Milner et al., 1968; Corkin, 2002) along with his ability to recall information up to 11 years prior to his surgery, including day-to-day occurrences, tunes, news events, etc (Corkin, 1984; Sagar et al., 1985). Subsequent research on subjects with hippocampal-specific damage has confirmed combined retrograde and anterograde memory declarative memory deficits supporting a role for the hippocampus in storing old memories and making new ones (Bayley et al., 2006; Manns et al., 2003).

An impressive dissociation between episodic and semantic memory has been reported in a case study of three people with childhood damage limited to the hippocampus (Vargha-Khadem et al., 1997). Remarkably, these patients attended mainstream school, were competent in reading and writing, had normal vocabulary, and showed average performance on a test of verbal IQ that explicitly tests for common-knowledge semantic information. In stark contrast, the same patients showed profound amnesia for everyday events and were unable to navigate familiar environments or orient in time; the hippocampal damage caused a selective episodic memory deficit. Other studies that have specifically tested for episodic memory in subjects with damage to the hippocampus have also shown impairments (e.g. Bayley et al., 2005, 2006; Steinvorth et al., 2005; Gilboa et al., 2006; Zola-Morgan et al., 1986; Rempel-Clower et al., 1996; Victor and Agamanolis, 1990; Manns and Squire, 1999), though there is debate as to whether semantic memory is also affected (Manns et al., 2003). There are also divergent findings as to whether episodic memories remain dependent on hippocampal function for all times or just within a consolidation window after learning (cf. Steinvorth et al., 2005; Kirwan et al., 2008), though there is a consensus that retrieval of recently formed episodic memories depends upon the hippocampus.

It is impossible to conclude from patients with permanent damage to the hippocampus that both encoding and retrieval processes are disrupted, since intact encoding could be masked by deficits in retrieval. However, in support of the conclusion that hippocampal damage affects both encoding and retrieval of episodic details, lesions of the CA1 region of the hippocampus are associated with transient global amnesia with retrograde episodic memory impairments lasting from hours to years (Bartsch et al., 2006; Guillery-Girard et al., 2004). Unlike in patients with permanent brain damage, those with transient global amnesia regain their ability to form new memories with intact general retrieval of episodic information learned before and after the acute episode, while events that took place during the amnestic period are forgotten (Sander and Sander, 2005). Furthermore, electrical stimulation of the medial temporal lobe of epileptic patients impairs encoding and subsequent retrieval of complex scenes (Halgren et al., 1985). These studies, in addition to work done in animals, show a clear role for the hippocampus in encoding as well as in retrieval of episodic information.

It is important to emphasize that not all mnemonic faculties are impaired after hippocampal damage or even medial temporal lobe damage. H.M. had intact motor learning (Corkin, 1965) as well as visual priming for at least some types of material

(Gabrieli et al., 1990). As discussed, the patients of Vargha-Khadem, with damage limited to the hippocampus, could learn a great deal of new information with a specific deficit for autobiographic information or information associated with a time or a place.

The critical hippocampal contribution may be allowing one to explicitly ‘remember’ the moment when some information was acquired rather than simply ‘knowing’ that the information has been encountered before (Aggleton and Brown, 1999; Eichenbaum et al., 1994; Yonelinas et al., 2002; Tulving and Markowitsch, 1998). This distinction between remembering and knowing was key in developing the concept of episodic memory, the expression of which requires remembering or recollecting the original experience (Tulving, 1972).

Central to the notion of recollection is the association between elements of an experience (Diana et al., 2008). For example, seeing an old friend can cue a memory of the last encounter with that person which in turn can bring to mind another friend who happened to be there at the time. Recognition of your friend could be accomplished by successful recollection of this previous episode or simply by recognizing that your friend is not a stranger. In the latter case, employing a vague sense of familiarity can be successful in the absence of any information about when or where or in what context that person entered your life.

Dual process models of memory posit that the hippocampus and associated diencephalic regions enable recollection/knowing/episodic memory while the perirhinal cortex and other rhinal cortices support familiarity. Several methods have been used to disentangle the contributions of familiarity and recollection in accurate recognition of studied material. The observation of poor free recall with intact recognition is often considered evidence for impaired recollection (Mayes et al., 2002) and this combined phenotype is accompanied with self reports of ‘knowing’ study items rather than ‘remembering’ them (Aggleton et al., 2005). Though recognition can be supported by familiarity alone, recall of paired associates has been reported to be particularly sensitive to hippocampal damage (Vargha-Khadem, 1997; Mayes et al., 2004). Other model based methods such as analysis of Receiver Operator Characteristics (ROC) have also been used to argue for impaired recollection and spared familiarity in patients with damage primarily limited to the hippocampus (Quamme et al., 2004, Yonelinas et al., 2002; Aggleton et al., 2005; Diana et al., 2010).

The dual process model is controversial and others find equivalent damage to both recollection and familiarity in patients with damage limited to the hippocampus (Wais et al., 2006; Dede et al., 2013). Whether or not the recollection is the sole and proprietary function of the hippocampus, the hippocampus appears to have a privileged role in subjective experience of recollection while other brain structures can support familiarity judgments in the absence of hippocampal function.

1.1.2 Neural recordings show hippocampal correlates of episodic memory

As suggested by examination of amnestic patients, brain imaging studies have shown hippocampal activation associated with processing of episodic information. As episodic memories require associating multiple aspects of an experience, experiments have tried to capture this feature by using stimuli such as images of complex scenes or explicit pairing of an item (e.g. words, pictures of objects or faces) with information present during item presentation, such as: other items, a colored background, the item's spatial location, instructed mental imagery, etc. During learning, these kinds of stimuli consistently evoke activation of the hippocampus that correlates with subsequent successful retrieval of episodic ('remembered') or associative information (Stern et al., 1996; Kirchhoff et al., 2000; Fernández et al., 1998; Eldridge et al., 2005; Song et al., 2011; Davachi and Wagner, 2002; Ranganath and Rainer., 2003; Kumaran and Maguire, 2006; Davachi et al., 2003; Staresina and Davachi, 2006). Greater hippocampal activation is also observed during accurate retrieval of associative memories made within the lab (Eldridge et al., 2000, 2005; Montaldi et al., 2006; Diana et al., 2010; Yonelinas et al., 2005; Yamashita et al., 2009; Takashima et al., 2009) as well as autobiographic episodic memories formed in everyday settings (Smith and Squire, 2009; Haist et al., 2001; Douville et al., 2005; Bayley et al., 2006).

Interestingly, hippocampal activation has also been reported for confident retrieval of the contextual source of a studied word irrespective of accuracy of the source memory (Wais et al., 2010). Indeed, the hippocampus is active during the imagination of events that have never occurred (Hassabis et al., 2007; Addis et al., 2007) and together these findings conform to the notion that episodic memory is a constructive process as opposed to a filing cabinet of personal experiences (Schacter, 1999).

These studies done in humans, and others done in animal models, inspired the binding of item in context model (BIC) which states that the hippocampus creates episodic memories by uniting the anatomical input stream carrying item information with that carrying contextual (source) information (Eichenbaum et al., 2007; Diana et al., 2007; Ranganath, 2010). In the BIC model context is defined broadly as either: spatial, temporal, semantic, social, affective, etc. The association of items with their context is hypothesized to create ‘hippocampal firing patterns [that] reflect unique conjunctions of stimuli with their significance, the animal's specific behaviors, and the places and contexts in which the stimuli occur’ (Eichenbaum et al., 2007). Reactivation of these cells is thought to drive cortical areas active during the original episode and thereby enable the conscious experience of recollection.

Due to the poor spatial resolution of fMRI, it is impossible to determine the hippocampal firing patterns that give rise to a blood-oxygen-level dependent (BOLD) signal that correlates with good memory for item and context. However, to partially address the issue of information content embedded within BOLD activity patterns, recent studies have employed multi-voxel pattern analysis (MVPA) and classification algorithms to determine whether a stimulus evokes reliable patterns related to recall. In several such studies, hippocampal activity recorded after a cue to recall recent (minutes to two week-old) or remote (ten year-old) autobiographic memory could be accurately attributed to the appropriate individual memory (Chadwick et al., 2010; Bonnici et al., 2012, 2013). The observation of activity patterns that are predictive of memory for individual experiences assuage concerns that hippocampal activation observed during retrieval reflects general novelty detection or a mismatch signal inherited from another brain region (Brewer et al., 1998) and lends credence to the theory that hippocampal computations allow for retrieval of episodic details.

Several studies have also examined episodic encoding and retrieval in epileptic patients with electrodes implanted into their hippocampi as well as other regions. These experiments serve as a crucial bridge between human studies linking neural activity with conscious episodic experience and the animal models discussed in the following section. In the first such study to link activity of the human hippocampus with memory, researchers found cells that responded to conjunctions of facial expressions and facial identity. Averaging across all stimuli presentations, about 20% cells fired differently (both higher and lower rates observed) for novel images as compared to a second presentation 10 hours later (Fried et al., 1997). The same percentage of cells discriminated old versus new pictures for both correct and incorrect assessments of familiarity indicating that the veracity of the hippocampal record of experience was superior to the subjects' conscious awareness of the image memory.

Even with the small number of neurons sampled in the study, the authors observed 37% of the cells fired robustly for any one of the specific conjunctions of task parameters, however, during the presentation of any single image only 1% of cells were active. Other studies showed that most hippocampal cells are inhibited after repeated presentation of the same stimulus, while a small minority show increased firing with familiarity, a combination of firing patterns that supports a model in which learning potentiates a few cells that in turn inhibit their neighbors (Viskontas et al., 2006). Since the stimuli in these studies were not specifically chosen to activate the recorded cells, these results indicate a sparse network of cells that each are broadly tuned to a large number of stimuli, a point that I will return to in Chapter 1.3 and in Chapter 5.2.

Follow up studies on single cell firing patterns revealed both the subsequent memory effect and important correlates of retrieval, thereby providing a mechanism for the fMRI studies discussed earlier. In paired associate learning, a significant number of cells (~20%) in the hippocampus but not the entorhinal cortex or amygdala fired differentially to words for which the paired associate would be correctly recalled versus words for which memory would subsequently fail (Cameron et al., 2001). Since analysis in this study was done across stimuli, these results suggest a global hippocampal process that enables subsequent memory. Similarly, in a study in which images of different categories were presented only once during study, neurons were identified that detected either novelty or familiarity of multiple stimuli during a test viewing (Rutishauser et al., 2006). The firing rates of these cells during the subsequent item presentation predicted accurate recollection of the original location of the study item, with novelty detecting cells firing less and familiarity detecting cells firing more during test viewing for accurate subsequent recollection (Rutishauser et al., 2008).

These studies by Rutishauser *et. al* are often used to argue in favor of a highly distributed global memory strength signal in the hippocampus (e.g. Wixted and Squire, 2010). Conversely, other studies have identified hippocampal cells that only detect novelty for specific stimulus types (Viskontas et al., 2006; Fried et al., 1997), though these results were not linked to memory performance.

Free recall studies have provided this important link between cells tuned to specific stimuli and memory. For example, cells that were active while a patient viewed a movie fire again prior to voluntarily recall of the same movie clip (Gelbard-Sagiv et al., 2008). In another study, cells that fire when a patient occupies a virtual location in a first-person video game are preferentially reactivated when an item that was located in that ‘location’ is remembered (Miller et al., 2013). Finally, the pattern of activity observed during the repeated presentation of a stimulus reflects encoding of the perceptual details of the stimulus itself, as well as elements of the temporal context of the item during its first presentation: ensemble activity recorded during repeated presentations of a stimulus correlates significantly with activity recorded before, during and after the item’s original presentation (Howard et al., 2012). Together, these studies in humans link episodic encoding and retrieval with specific hippocampal firing patterns and suggest that reactivation of particular hippocampal cells active during learning elicits conscious memory recall of the associated episode.

In summary, the hippocampus is at least temporarily required for the encoding and retrieval of events that occurred once. This is in contrast to the ability to recognize that single items have been encountered, a faculty which remains relatively intact after hippocampal damage and likely depends upon the perirhinal cortex (Aggleton and Brown, 1999; Eichenbaum et al., 1994). These findings suggest an important role for the hippocampus in associating items with their spatiotemporal context (Aggleton and Brown, 1999) or source (Diana et al., 2008) to create unique patterns of neural activity that, when reactivated, define the conscious experience of an episodic memory.

1.2 Episodic memory depends upon the hippocampus in animals

1.2.1 Hippocampal lesions cause episodic-like memory deficits

Animal models offer two critical features that are not possible in studies on human patients. First, with the modern pharmacological and genetic techniques it is possible to precisely control the extent and even the duration of a neural intervention such that the issues of ‘hidden pathology’ are far less of a concern (but see Rudy and Sutherland, 2008). Second, in animals it is possible to dictate the timing between learning, hippocampal disruption, and test.

Endel Tulving has perhaps been the most vocal proponent of the view that animals do not possess episodic memory due to their inability for autonoetic awareness, or conscious awareness of self (e.g. Tulving, 2005; see Cheke and Clayton, 2010 for counter-examples to Tulving’s ‘spoon test’). However, due to the impressive cross-species homology in hippocampal anatomy, physiology, and function (Manns and Eichenbaum, 2006; Clark and Squire, 2013) and the inability to definitely query the conscious state of another being (Nagel, 1974), parsimony suggests that the same neural processes that give rise to episodic memory in the human likely give rise to episodic memory animals as well. However, the major limitation of animal research is that animals cannot verbally report their recollections, and therefore operational definitions of ‘episodic-like’ memories are required (e.g. Clayton et al., 1998). While early tests (e.g. the delayed non-match to sample task) of episodic memory could be solved with familiarity judgments alone, contemporary animal studies recognize the importance of testing memory for an item's context at the time of study and ideally employ tests of source memory or what happened where and when.

The delayed matching to sample (DMS) and the delayed non-matching to sample (DNMS) tests were thought to be 'benchmark' tests for hippocampal-dependent memory in the non-human primate (Gaffan, 1974; Mishkin, 1978; Clark and Squire, 2010). These tests require studying an item and then either selecting the same item (DMS) or a different item (DNMS) rather than a foil during a delayed memory test. The key innovation in these tasks was to introduce a large number of 'trial-unique' stimuli, presumably tapping into the need to use episodic memory, rather than habit memory, for recognition judgments.

Medial temporal lobe damage causes impairments of a monkey's ability to recognize familiar stimuli for both the DNMS and DMS task. The impairment depends upon the amount of time between study and test and the number of items studied before test (list length) as well as whether there are distractions between study and test (Mishkin, 1978; Gaffan, 1974; Zola-Morgan and Squire, 1985). Spared performance at short durations and with short lists controls for perception, motivation and attention during study and mirrors the observation that the medial temporal lobe is more related to long term memory than working memory or perception (Scoville and Milner, 1957).

Decades of work done on the DNMS task have shown that the surrounding cortical regions, in particular the perirhinal cortex, can be solely responsible for successful performance on DNMS task (Meunier et al., 1993, 1996; Murray and Mishkin, 1998; Baxter and Murray, 2001; Nemanic et al., 2004). However, other labs have shown memory impairments on the DNMS task after selective hippocampal lesions (Beason-Held et al., 1999; Zola et al., 2000) and performance on a spontaneous item recognition task is severely impaired in monkeys with selective hippocampal damage (Nemanic et al., 2004). These results in the monkey parallel human studies that show that memory for single items can be supported either by the hippocampus or the rhinal cortices and highlight the need to test item and context associations in a model of episodic memory.

Manipulations of either the rodent and monkey hippocampus have confirmed a dissociation between deficits on item and location associations versus item recognition. Since rats and monkeys, as well as human infants (Fagan, 1970), have a natural tendency to spontaneously attend to novelty, tasks can be designed to compare the time spent ‘exploring’ novel versus familiar items. As mentioned, lesions to the hippocampus in monkeys in some tasks can disrupt normal visual exploration of novel items (Nemanic et al., 2004) though in other tasks, attention to novel items is intact. However, in versions of the task that test for the conjunction of item and place, hippocampal lesions reliably cause monkeys to fail to recognize spatial rearrangements of familiar items (Bachevalier and Nemanic, 2008). Similarly, rats with hippocampal lesions fail to spend more time exploring an item that previously was presented within a different context (Mumby et al., 2002) or within a different position of the same context (Ennaceur et al., 1997; Mumby et al., 2002; Eacott and Norman; 2004), showing a hippocampal role for item and place memory. Pharmacological studies have shown the inactivation of the mouse hippocampus with lidocaine or disruption of plasticity with scopolamine, NMDAR antagonism, or anisomycin disrupts novelty preference for familiar items in a new positions (Assini et al., 2009) or for familiar items placed in a new context (Balderas et al., 2008) while sparing item recognition. Therefore, under some circumstances item recognition does not depend upon the hippocampus proper, though disruption of hippocampal processing severely impairs recall of the spatial context in which an item was studied.

Instrumental learning of item and place associations also depends upon the hippocampus. When rats are cued with a flavored pellet, they can readily learn to retrieve pellets of the same flavor from a previously paired location (Day et al., 2003; Tse et al., 2007, 2011). The ability to encode these flavor and place associations depends upon glutamatergic transmission in the hippocampus, since blocking AMPAR or NMDAR during initial flavor and position pairing disrupts subsequent ability to find the location associated with the flavor cue (Day et al., 2003; Tse et al., 2007). Retrieval also depends upon AMPAR function in the hippocampus (Day et al., 2003), though only for recently formed item and place associations (Tse et al., 2007). Severe memory deficits for object and place associations with spared object recognition have also been identified in monkeys with fornix lesions (Gaffan et al., 1994), supporting a critical role for the hippocampus in association items with their spatial context.

Even item and context associations can eventually be supported by extra hippocampal regions as shown in a different task, and one that is formally identical to that used in the experiments described in Chapter 4. In this task rats are trained that reward is available when a particular CS is present in one of two contexts, while presentation of a different CS signals reward in the other context (Good et al., 1998). In one set of experiments, rats with pre-training lesions acquired the conditioned discrimination task for a first set of two CSs, revealing successful item and context association. Training on the second set of CSs was done in stages, with single CSs initially presented in the context paired with reward. Following this simple conditioning, CSs were then presented in both the conditioned context and, for the first time, in the non-conditioned, non-rewarded context. Normal rats immediately used the change in context to respond less to the CS in the non-conditioned context while lesioned rats responded equally to the CS in both contexts. Lesioned rats did not spontaneously bind the context with the CS in the absence of instrumental training (Good et al., 1998) despite being able to use contextual information to guide behavior as evidenced by intact biconditional discrimination of the first set of CSs and subsequent learning for the second set. Other studies have confirmed that pre-training hippocampal lesions do not affect overall performance on this task (Coutureau et al., 2002). In a similar task, mice with NMDAR ‘knocked-out’ in CA3 show retarded learning when both CSs and contexts were experienced after the gene deletion but intact learning when either paired associate was familiar (Rajji et al., 2006). Together these results suggest that multiple brain regions can associate specific stimuli with background context and the hippocampal involvement is especially evident for rapid binding of novel items with their context, as predicted for an episodic memory system.

In support of a role of the hippocampus during item and context association, post-training lesions of the entire hippocampus cause retrograde memory impairments in a similar task (Komorowski et al., 2013). These same rats rapidly relearned the pre-lesion associations and anterograde memory for new item and context pairings was intact (Robert Komorowski, personal communication). These studies show that the hippocampus can be involved in the normal association of items and context, though other brain regions can compensate when the hippocampus is lesioned.

Although item and space are important to episodic memory, time may be the key feature that distinguishes episodic and semantic memory. Clayton and Dickinson (1998) were the first to show episodic-like memory in scrub jays that could remember where and when they buried particular food items for later recovery. To test “what”, “where”, and “when” memories, rodent studies have employed clever adaptations of the spontaneous object recognition (Dere et al., 2006; DeVito et al., 2010a), tests of fear generalization to contexts linked with a time of day (Iordanova et al., 2009), or explicit instrumental training to recognize sequential odor presentations (Ergorul and Eichenbaum, 2004; Fortin et al., 2002; Kesner et al., 2002). These rodent experiments all show a critical role of the hippocampus in the association of what happened where and when – the association of an item with its spatiotemporal context.

Many other studies have shown that hippocampal lesions cause contextual learning deficits (e.g. Bouton, 1993; Kim and Fanselow, 1992; Matus-Amat et al., 2004) and some have argued that the context must be incidentally encoded as a perceptual ‘background’ to identify the unique hippocampal-dependent contribution (Wang and Morris, 2010; Kim and Fanselow, 1992). Accordingly, it has been suggested that the hippocampus is a passive, ‘automatic’ recorder of the ongoing activity patterns in the neocortex (Wang and Morris, 2010), and it is this dragnet approach to binding disparate cortical regions that makes the region so important for episodic memory.

1.2.2 Recording studies reveal the 'raw ingredients' for episodic memory formation are present in the hippocampus

The contemporary study of hippocampal single cell physiology began with the discovery of the 'place cell' (O'Keefe and Dostrovsky, 1971). In their first characterization of the spatial correlate of hippocampal cells, O'Keefe and Dostrovsky (1971) mention a cell in dorsal CA1 that fired preferentially in a particular position and only when the rat was restrained. This observation that hippocampal receptive fields are dependent upon the conjunction of spatial and transient non-spatial stimuli was overlooked for many years with the advent and success of the random foraging paradigm that equalizes behavior across all positions and reliably produces the spatial correlate, or place field (Muller and Kubie, 1987; Muller et al, 1987). However, the importance of non-spatial information to hippocampal to the spatial receptive field has since been well-characterized in rodents, monkeys and humans (see section 1.1.2) and I will review the data from primarily experiments on rodents that are relevant to models of episodic memory storage.

Decades of work on the place fields of hippocampal cells, and CA1 cells in particular, have shown preserved spatial correlates after a staggering number of neural manipulations. Place fields in CA1 are preserved after lesions or disruptions to each of its inputs, including: the prefrontal cortex (Kyd et al., 2005; Hok et al., 2013), anterodorsal thalamus and post-subiculum (‘head direction system’; Calton et al., 2003), the Schaffer Collaterals (Brun et al., 2002), CA3 (Brun et al., 2002; Nakashiba et al., 2008); large entorhinal lesions of both LEC and MEC (Van Cauter et al., 2008; Miller and Best, 1980), Layer III of the EC (Brun et al., 2008; Suh et al., 2011), LEC (Lu et al., 2013), perirhinal cortex (Lee and Park, 2013), the mammillary bodies (Sharp and Koester, 2008), the fornix (Miller and Best, 1980), or the medial septum (Leutgeb and Mizumori, 1999). Combined, these studies show that no one region is necessary for formation of a place field, at least in area CA1.

Many models of hippocampal place cells hypothesize that grid cells in the MEC combine to define the spatial firing field of hippocampal cells (e.g. Hafting et al., 2005; O’Keefe and Burgess, 2005; McNaughton et al., 2006; Moser et al., 2008; Rolls et al., 2006; Hasselmo, 2009). However, place fields are preserved after large lesions to the MEC (Van Cauter et al., 2008; Miller and Best, 1980). Grid field organization is also severely disrupted after medial septal inactivation (Brandon et. al., 2011; Koenig et al., 2011), which, when lesioned, spares the spatial firing fields in the hippocampus (Leutgeb and Mizumori, 1999). On the other hand, disruption of the hippocampus causes grid cell firing fields to become disorganized leaving only some residual spatial information (Bonnevie et al., 2013; Fyhn et al., 2004). Together, these results suggest that grid cells are not necessary for place cells, though they may be sufficient.

Place fields largely persist in a variety of situations in which critical sensory information is impoverished, for example: in the darkness (Quirk et al., 1990), during involuntary motion (Song et al., 2005), in zero gravity (Knierim et al., 2000), in virtual reality (Harvey et al., 2009), after surgical deafness and visual masking (Hill and Best, 1981), and in pups that have not yet opened their eyes (Save et al., 1998). As suggested by the lesion studies, these results show that multiple streams of input can produce a spatial correlate.

A large number of experimental factors influence changes in the position of a place field (global remapping) and the firing rate (rate remapping) within a place field (Leutgeb et al., 2005; Muller et al., 1987). A cell’s firing field is displaced with discordant movements of local and distal cues (Shapiro et al., 1997), changes in either color or odor of the environment (Anderson and Jeffery, 2003), changes in the geometry of the testing chamber (Wills et al., 2005; Lever et al., 2002), estimates of path integration (Gothard et al., 1996a; Skaggs and McNaughton, 1998; Kraus et al., 2013; Ravassard et al., 2013), and other features (e.g. recording room) which typically are thought of as defining spatial context (Leutgeb et al., 2005).

In contrast, the firing rate of a cell can be dictated by changes in the geometry or color of an open arena (Leutgeb et al., 2005; Lu et al., 2013), the identity of sampled item (e.g. scented sand) (Komorowski et al., 2009; Tort et al., 2011; Wood et al., 1999; Manns et al., 2009; Rolls et al., 1989), the item sampled moments earlier (MacDonald et al., 2011), the removal of landmarks (Hetherington and Shapiro, 1997), the interval of time after a salient event (Kraus et al., 2013; MacDonald et al., 2011; Pastalkova et al., 2008), conditioned stimuli (Moita et al., 2003), reward magnitude (Lee et al., 2012; Kobayashi et al., 1997) and the direction with which an animal is running (Navratilova et al., 2012;

McNaughton et al., 1983), was running (Ferbinteanu et al., 2003) or will run (Wood et al., 2000; Ferbinteanu et al., 2003).

Combined these observations led to the hypothesis that the identity of the active ensemble defines the spatial reference frame, the location of the firing fields dictates position within that reference frame, and the firing rate of the cells the non-spatial events that occur within that position (Leutgeb et al., 2005; Lu et al., 2013). This view conforms well to the notion that the hippocampus is a cognitive map within which events can take place (O'Keefe and Nadel, 1978).

Manipulations that do not intuitively redefine the current spatial map (e.g. removing the rat and returning it to the same environment to do the same task; McNaughton et al., 1996) can also change the location in which a cell fires (e.g. Markus et al., 1995; Bostock et al., 1991; Bahar et al., 2011; Eschenko and Mizumori, 2007). Regardless, this global remapping has been considered evidence for a change in spatial reference frame (Samsonovich and McNaughton, 1997, McNaughton et al., 1996). Due to the circularity in this line of reasoning and since the extreme low end of rate remapping is equivalent to global remapping, the distinction between the two processes seems unsatisfactory and may simply reflect the *degree* of difference between the two conditions rather than the *kind* of difference.

Remapping experiments often emphasize that the location in which single cells fire and the distance between firing fields of pairs of cells are independent in two unrelated contexts (Kubie and Ranck, 1987; Leutgeb et al., 2004; Vazdarjanova Guzowski, 2004). Extrapolation from these kinds of data led to the theory that the hippocampus performs a ‘patter separation’ and ‘orthogonalization’ function on cortical traces to reduce the representational similarity of experiences with overlapping features (Norman and O’Reilly, 2003; Norman, 2010; McNaughton and Morris, 1987; Rolls, 1989, 1990; O’Reilly and McClelland, 1994; Hasselmo and Wyble, 1997). I will return to this critical point in Chapter 1.3, Chapter 4.4 and Chapter 5.2.

These studies all show that the inputs required to form an episodic memory are available to the hippocampus, and show experience dependent changes that are suggestive of a memory process, though few studies have identified changes in hippocampal firing patterns that correlate with learning. For example, cells remap after fear conditioning but no correlation between the degree of training-induced remapping and the subsequent fidelity of recall has been reported (Moita et al., 2004; Wang et al., 2012); remapping was less in control conditions, however.

Accumulation of firing fields at goal areas has been observed as rats learn to swim to escape platforms (Hollup et al., 2001) or to find hidden food locations (Dupret et al., 2010, 2013; Pfeiffer and Foster, 2013). These results suggest that firing at goals reflects the long-term memory for the events that happened at those locations. However, the decrease in the over-representation with subsequent training does not reflect the elimination of the goal memory (Fyhn et al., 2002). Pharmacological blockade of NMDAR blocks both the overrepresentation of goal sites in CA1 and tendency to dwell at goal locations on non-rewarded probe trials; however, systemic delivery of NMDAR antagonists in the reported dose causes drastic changes in wait behavior (personal observations). Finally, global remapping does not alter the ability of rats to retrieve food rewards within the equivalent position of an altered environment (Jeffery et al., 2003), an observation that highlights the ambiguous relationship between remapping and memory.

Of the recording studies that have identified memory correlates, the majority report that learning parallels encoding similar experiences with different patterns of hippocampal activity, a process interpreted in the pattern separation framework described above. For example, as monkeys learn to associate images of scenes with subsequent saccadic eye movements, hippocampal cells begin to discriminate scenes with changes in firing rate (Wirth et al., 2003). Unlike other experiments that have reported experience dependent pattern separation (Bostock et al., 1991; Lever et al., 2002), in this study the degree of scene selectivity was correlated with an estimated learning curve. Other studies have found that when rodents traverse opposite directions over the same space, cells fire differentially, and the degree to which CA1 cells, but not CA3 cells, differentiated the opposite trajectories correlated with rats' ability to remember which of two overlapping journeys predicted reward (Bahar et al., 2011). Cell firing rates also differentiate the same traversal through a position based off of the start location and the goal of a rat (i.e. journey-dependent coding). This journey-dependent coding decreases when rats make errors on a matching-to-place task, suggesting a link between memory for specific routes and hippocampal differentiation of positions that are common to overlapping journeys (Bahar et al., 2012). Finally, Komorowski et al. (2009) reported that cells become tuned to item and place conjunctions as rats learn to use context to predict which of two items contain reward. Due to the striking correlation between item tuning and learning curve, this training protocol was adopted for the experiment reported in Chapter 4. Therefore in each of these studies, encoding was accompanied by increases in the degree to which the network discriminates experiences, as expected in a system that supports episodic memory via pattern separation.

Studies on retrieval have also linked hippocampal activity to memory performance. It is possible that the hippocampus activity during rhythmic bursts of activity known as sharp-wave ripples constitute the physiological basis for recollection (Buzsaki, 1986; Chrobak et al., 2000; Pfeiffer and Foster, 2013; Carr et al., 2011). Analysis of cell activity as rats run around linear tracks or open fields and subsequently stop or sleep have shown that cells that were active during an experience are reactivated during ripples at compressed timescales during ‘replays’ (Wilson and McNaughton, 1994; reviewed in Carr et al., 2011). To demonstrate a link between replay activity and recall, hippocampal activity was recorded as rats alternated between a random and stationary goal. Replays recorded while the rat was still, transiently encoded future trajectories to the stationary goal location (Pfeiffer and Foster, 2013). These results suggest that reactivation of the cells active during prior journeys elicited a memory that could be used in navigational decisions (see also Johnson and Redish, 2007).

Information required for episodic memory (what, where and when) is encoded by hippocampal cells, hippocampal cells change firing properties as animals learn, and patterns of activity during recall contain information that directly relates to successful retrieval. Overwhelmingly, the changes in firing patterns that correlate with learning lead to an increase in the discriminability of learned stimuli, findings that fit with the standard ‘Hebb-Marr’ model of hippocampal function discussed in the following section.

1.3 The ‘Hebb-Marr’ model of episodic memory

Many theories of memory have been proposed that link hippocampal physiology to specific computations that could underlie learning and memory. Perhaps the most influential of these theories is the ‘Hebb-Marr’ model (Hebb, 1949, Marr, 1971) along with its variants (Hasselmo et al., 1996, McNaughton and Morris 1987, McNaughton and Nadel 1990, Treves and Rolls 1994; O’Reilly and McClelland, 1994). The common thread in these models is that the role of the hippocampus is to store representations of correlated inputs in independent networks which can be queried with partial input to reinstate the pattern of activity that was present during initial learning. This reinstatement constitutes the episodic memory.

Nearly all computational theories of memory model learning as changes in the connection strength between pre- and post-synaptic neurons that are coactive during learning. This concept of memory was famously proposed by Donald Hebb (1949) and commonly summarized simply as: cells that fire together wire together. Anti-correlation between pre- and post-synaptic cell firing can also lead to decreases in the connection strength between cells (Markram et al., 1997), thus bidirectional changes in synaptic strength reflect patterns of coactivity.

Studies on activity-dependent changes in synaptic strength led to the hypothesis that associative learning can be accomplished by finding optimal connection strengths between every pre-synaptic cell and every post-synaptic cell. This set of optimal connections strengths constitutes a weight matrix that defines the desired transform of the input patterns to the associated output patterns (Rumelhart and McClelland, 1986; McCulloch and Pitts, 1943; Oja, 1989; Hopfield, 1982; Anderson, 1977, Hinton 1989, Kohonen, 1984). Many learning rules for synaptic weight modification have been proposed, though regardless of the precise strategy, weight matrices in general can define a transform between a set of inputs and a set of outputs.

However, when multiple items are sequentially stored within the same weight space, the adjustments to the synaptic weights driven by new items disrupts the ability of artificial neuron networks to recall old items. In linear systems this interference is proportional to the correlation of the stored input patterns (Frean and Robbin, 1999), however, in non-linear training regimes (e.g. feedforward backpropogation), new learning can severely disrupt stored memories in what has been termed ‘catastrophic interference’ (McCloskey and Cohen, 1989). The reason for this is that the solution to ideal weight matrix to associate two patterns of activity is distributed across all of the individual weights in the matrix. Big differences in retrieval accuracy can occur with two similar conformations of the synaptic weight matrix, highlighting the existence of ‘weight cliffs’ in the solution space for an ideal transform function between inputs and outputs (French, 1999).

Since interference is particularly detrimental for correlated patterns, one solution to the problem is to decorrelate the input patterns prior to storage (French, 1992). Decorrelation can occur by mapping a smaller representational space onto a larger one (Marr, 1971) and also by representing inputs with a sparse network of active cells with the majority silent (French, 1999). The large number of mostly silent cells in the dentate gyrus suggested that the first stop in the trisynaptic loop involves pattern separation of correlated inputs (McNaughton and Morris, 1987; O’Reilly and McClelland, 1994; Hasselmo and Wyble, 1997). Adult neurogenesis in the dentate offers another mechanism by which to map inputs onto virgin weight space, thereby ensuring minimal interference (Wiskott et al., 2006). The sparse dentate code is hypothesized to be the result of a ‘winner-take-all’ computation in which a few cells fire *only* when specific conjunctions of stimuli (e.g. item in location) are present and once active inhibit their neighbors (Hasselmo and Wyble, 1997; Sloviter and Brisman et al., 1995). The connections between the EC to the DG and from DG to CA3 are considered random (O’Reilly and McClelland, 1994) and therefore the EC input is mapped onto a few random dentate cells. The combination of a random graph and the DG pattern separation strategies ensures that correlated inputs will be orthogonalized within the hippocampus. The pattern separation function of the dentate is thought to be the critical feature that allows storage of correlated episodic memories without catastrophic interference.

After the dentate, information is propagated to CA3, a region that is known to have axons which synapse widely back over CA3 to form an autoassociative or recurrent network. Once a pattern of activity is stored within the weight matrix of a recurrent network, subsequent presentations of degraded inputs can associatively recall the originally stored pattern, a process known as pattern completion (Hopfield, 1982; Kohonen, 1972; Anderson, 1972). Therefore CA3 is hypothesized to be important for pattern completion, or recall, of familiar memories with degraded inputs (O’Reilly and McClelland, 1994) either from DG of from direction projections from EC.

Finally, CA3 inputs onto CA1 and finally back to the ECII. Hippocampal signals must be inverted back into a form that can be read by the EC. The CA1 to EC 'point-to-point' topography is considered evidence for this translator role and therefore CA1 patterns should reflect some of the correlations of the input structure, but, due to pattern separation in the DG and CA, represent correlated inputs with relatively distinct patterns (O'Reilly and McClelland, 1994; Hasselmo and Wyble, 1997).

In Hebb-Marr models, the pattern completion function of CA3 is thought to be limited to recreating a pattern of activity during encoding. Novel conjunctions of familiar elements would not associatively activate prior traces but instead create a new representation that is ideally independent of those representing the previously experienced conjunctions (Kumaran and McClelland, 2012; O'Reilly and McClelland, 1994). These transforms on the signal within the hippocampus, in the words and spirit of McClelland, "by nature tend to support episodic memory at the expense of capturing the higher order structure of a set of experiences" (Kumaran and McClelland, 2012). It is this final conclusion that I will argue against throughout the rest of this thesis.

1.4 The hippocampus supports memory schemas that define the relationships between episodic memories

The important role of schemas in memory formation was introduced to cognitive science in classic studies of Piaget (1926) and Bartlett (1932), who proposed that new memories that are consistent with pre-existing knowledge are readily assimilated within established memory networks, often requiring updating or modification of the existing schema to accommodate the new information. The relational theory of hippocampal function states that hippocampus encodes the relationships that compare and contrast items and that allow for "weaving new items into the existing organization of memory" (Cohen and Eichenbaum, 1993).

The 'Hebb-Marr' account of episodic memory described above states that different hippocampal cells encode unique conjunctions of stimuli. In contrast, the relational memory theory states that the hippocampus codes individual elements of a memory *as well as* their relationships to one another which in many cases may be evidenced by cells responding to both individual stimuli and their conjunctions. Due to this difference in models, the relational theory also hypothesizes that experiences with overlapping features would generate correlated hippocampal patterns with the correlation structure dictated by the features that are functionality equivalent or distinct as defined by motivational needs and/or predictiveness. This theoretical shift in hippocampal functions allows for the linking of the distinct elements of experiences that overlap in other respects. This transitive association is possible due to cells representing distinct features of multiple episodes being associated with a subset of cells that encode the features common across episodes (Eichenbaum et al., 2004).

McClelland et al. (1995) introduced the idea that the hippocampus supports schema assimilation and accommodation in an influential model of hippocampal-cortical interactions during consolidation. In this model, the role of the hippocampus was to present patterns of activity for the cortex to discover the underlying correlation structure.

Subsequent modeling explicitly adopts “a theoretical ideal of perfect pattern separation in the hippocampus” (Kumaran and McClelland, 2012) which is “optimal for an unstructured episodic memory system in which each memory must be kept distinct from other memories” (Rolls, 2013). I will argue against the notion that the hippocampus dismantles the correlation structure between memories in the following sections on hippocampal schemas.

1.4.1 Disrupting hippocampal processing in animals and humans impairs relational memory schemas

Much of the evidence reviewed above for the hippocampal role in episodic memory can be reinterpreted within the relational model of memory. For example, the ‘Hebb-Marr’ account hypothesizes that the failure of people, monkeys, and rats with hippocampal damage to recognize the novelty of different spatial reconfigurations of familiar items can be explained to the absence of conjunctive cells that respond only to the unique combination of an item in its position – the activity snapshot that defined the episodic experience. Alternatively, the same deficit can be reinterpreted under the relational memory framework as a lack of a code for how the individual items were spatially related to one another (Eichenbaum, 1999). The episode was defined by a representation of the items, their spatial relationship to one another and their relationship to other aspects of the study context. The code for the relationships rather than the unique conjunction is hypothesized to be the critical hippocampal contribution.

More formal tests of the relational memory theory are required that make differential predictions from the ‘Hebb-Marr’ models of episodic memory. Typically, these tests go beyond asking subjects to report whether a precise conformation of stimuli was present at study and instead test for whether subjects learned relationships acquired over multiple trials to make inferences when faced with novel recombinations of familiar material.

Cognitive mapping experiments offer support for the relational memory theory. For example, the ability of rats to find a hidden platform in the Morris Water maze is known to depend upon the hippocampus when rats must begin from a position that differed from the start position during prior training (Eichenbaum et al., 1990; Morris et al., 1982). In contrast, ‘navigation’ towards the platform can remain intact despite hippocampal lesions if the start location is held constant over trials (Eichenbaum et al., 1990). Therefore, the hippocampus is selectively required for using the acquired knowledge about spatial relationships to make novel inferences, as has been confirmed in other studies testing relational memory in mice on a radial arm maze (Etchamendy et al., 2003). These experiments are consistent with the hippocampus as a cognitive map (O’Keefe and Nadel, 1978); however, other experiments have demonstrated that non-spatial schemas also require hippocampal function.

By examining the non-spatial organization of related memories, Bunsey et al. (1996) showed that normal rats link overlapping paired associates and make new inferences about indirectly related elements. Rats learned that odor A predicted reward with odor B which in turn predicted reward with odor C. Rats were then tested on whether odor A can be used to inferentially predict reward association with odor C (transitive probe) and whether rats could use odor C to predict reward with odor B (symmetry probe). Both tests of flexibility were impaired with large hippocampal lesions. Interestingly, in humans the ability to do this kind of transitive inference emerges around the same time that children learn to use maps (Piaget, 1928) which in turn correlates well with postnatal maturations in the hippocampal circuit (Lavenex et al., 2013).

In another test of transitive inference, subjects are first given overlapping ‘premise pairs’ of items in which they must choose A>B, B>C, C>D, D>E and are then probed with a novel pairings (e.g. B>D) to determine whether relationships learned in separate experiences can integrated to make novel inferences. Lesions of the rodent hippocampus prior to any training cause memory deficits specifically for the inferential probes leaving intact performance on the explicitly trained ‘premise pairs’ (Dusek et al., 1997). Studies in mice have shown that lesions between initial training on the ‘premise pairs’ and the inferential probe also cause memory deficits revealing an important role for the hippocampus in retrieval during inferential performance (DeVito et al., 2010b). Another lab showed a critical role for encoding but not retrieval (Van der Jeugd et al., 2009) and the difference may be the number of days of training (mean number of days was 10 in DeVito *et al.* and 20 in the Van de Jeugd *et al.* study) on the premise pairs prior to hippocampal surgery (see Chapter 2.1). Finally, inferential probe performance is selectively disrupted in human patients with damage limited to the hippocampus (Smith and Squire, 2005) and sleep turns out to be critical for probe performance suggesting an ‘offline’ role of the hippocampus in integrating memories into a unified schema (Ellenbogen et al., 2007) (see Chapter 2.1).

In a separate test of relational processing, Tse et al. (2007) demonstrated that rats develop a schema of locations where different foods are buried by showing that once several food/location associations had been formed, new ones could be added within a single trial; however, in a different environment, the learning of new associations was much more gradual. Moreover, when new associations could be integrated within a pre-existing schema, hippocampal lesions after 3 hours, but not 48 hours, impaired subsequent performance, revealing a consolidation gradient considerably steeper than those reported in studies in which learning did not benefit from an existing schema. Follow up studies revealed the mPFC as the important brain region for memory recall for remote memories and, alongside the hippocampus, for learning new food and position combinations once a prior schema had been formed (Tse et al., 2011; Wang et al., 2012).

Combined, these studies show that the hippocampus does more than allow successful memory for the contents of an episode, but it is also important for the integration of related episodes in order to make new inferences or to enhance the speed of learning. These results are not obvious predictions from a model in which episodic memories are stored independently of one another, though updates to the ‘Hebb-Marr’ model have been suggested (Kumaran and McClelland, 2012; McClelland, 2013).

1.4.2 Imaging studies show a hippocampal role in activating and integrating related memories

In humans, hippocampal activity has also been related to performance on a variety of transitive inference tasks. Activity during 'premise pair' learning (Nagode and Pardo, 2002; Greene et al., 2006; Shohamy and Wagner, 2008; Wimmer and Shohamy, 2012; Preston and Zeithamova, 2010, 2012; Kumaran et al., 2012) correlates with subjects' future ability to make novel inferences based on information integrated over initial training. Furthermore, there is higher hippocampal activity during inferential probe tests than during controls tests for recognition of study material (Heckers et al., 2004; Preston et al., 2004; Zalesak and Heckers, 2009) or during probe tests for which subjects failed to show correct inferential judgments (Preston and Zeithamova, 2010) showing hippocampal involvement during learning and test. Hippocampal activity also scaled linearly with subjects' hierarchical judgment of the value of individual stimuli (Kumaran et al., 2012) which were presented separated as overlapping 'premise pairs'.

Other tasks have also revealed hippocampal activity that relates to subjects' discovery (implicit or explicit) of patterns that can only be achieved by integrating information across trials. For example, hippocampal activation was seen in a motor task which contained embedded patterns that could only be learned by realizing the higher order structures of a set of sequences of in finger movements (Schendan et al., 2003). In another demonstration that the hippocampus is required in combining episodes to extract commonalities, Kumaran et al. (2009) adapted the weather prediction task such that it was possible to predict an outcome in two ways: the episodic strategy of linking each individual pattern with the outcome and a schema strategy where only a subset of the presented information was required for perfect performance, though realization of which subset required integrating knowledge over trials. Hippocampal and VMPFC connectivity during learning strongly correlated with subjects' ability to discover these hidden regularities and only hippocampal activity predicted subjects' ability to use the conceptual information in novel circumstances. These results highlight the dichotomous role of the hippocampus in forming unique episodic memories, but also in realizing the correlation structure between individual experiences.

Clever task designs have enabled testing whether memories for individual episodes become integrated during encoding or only during inferential probes, a time in which hippocampal activation does not always correlate with performance (Shohamy and Wagner, 2008). Since images of houses and faces activate different brain regions, it is possible to teach subjects an association between a house and a face, then in a second stage of encoding, present the originally paired face or house and look for covert reactivation of the non-presented stimuli. Two studies using this type of design showed reactivation of the associated, though non-present stimulus category during this second stage of encoding (Zeithamova and Preston, 2012; Shohamy and Wagner, 2008). Changes in hippocampal activity correlated with the degree of associative reactivation (Zeithamova and Preston, 2012) which in turn correlated with inferential probe performance showing a compelling role for the hippocampus and relational recall during learning. These results led to the integrative encoding hypothesis which states that stored trace is dependent upon both the present stimuli as well as other stimuli evoked by the current situation (Gluck and Myers, 1993; Zeithamova and Preston, 2010, 2012).

Few single unit recording studies have explicitly tested predictions from the relational memory theory and most supporting evidence is indirect. For example, in studies in which a rat runs through the same location with different start and end positions, some cells differentiate the journeys (‘splitter cells’) while others do not (Wood et al., 2000). It has been suggested that the cells that do not differentiate journeys “captured a set of regularities in stimulus relations that were common across both conditions—the nodal features” and that these cells could activate cells that do differentiate events to build a relational network of experience (Eichenbaum, 1999). However, another interpretation is that non-splitter cells were simply not tuned to journey differences and could no sooner help to link experiences than cells in V1 that fire for oriented lines irrespective of where the viewer was moments earlier.

A more compelling case for these nodal cells comes from studies of functional equivalence in which the perceptual features of stimuli are distinct though the meaning is similar. Striking examples of ‘nodal cells’ that show functional equivalence come from studies in human epileptic patients in which cells increase firing rate in response to pictures all from the same category of stimuli such as houses or faces (Kreiman et al., 2000) or most remarkably from cells that respond only to symbols that represent an individual, such as a cell that responded to colored photographs, line drawing, and the written name of the actor Halle Barry (Quiroga et al., 2005). Similar category cells have been reported in the monkey during a delayed matching task (Hampson et al., 2004).

Finally, as rats learn to run down the arms of a W-maze to retrieve reward, cells begin to fire at equivalent positions on each arm containing reward (Singer et al., 2010).

It is unclear how these cells become tuned to diverse, equivalent stimuli. The relational theory predicts that the tuning to equivalent stimuli is a function of associative properties of the hippocampus (Eichenbaum, 1999); however, some types of category learning are independent of hippocampal function (Knowlton and Squire, 1993; but see Zeithamova et al., 2008). Therefore, equivalent firing in the hippocampus may be a reflection of category coding in afferent networks.

One study suggests that the hippocampus reactivates experienced patterns in order to form novel combinations that are so critical in the inferential probes that test for ‘flexible’ memory expression. Gupta et al. (2010) reported replays of spatial representations that composed overlapping spatial trajectories that occasionally linked to form representations of novel routes that would be consistent with a navigational inference of related previous experiences.

Few studies have linked learning with changes in single cell firing patterns that would be uniquely predicted by the relational theory of hippocampal function. However, the existence of cells that reflect perceptually distinct though functionally equivalent stimuli show evidence that the correlation structure of memories has the potential to be stored within the hippocampus, not just in the cortex (Kumaran and McClelland, 2012; McClelland et al., 1995).

1.5 Conclusions and motivation for current experiments

The episodic memory theory and relational theory of the hippocampus have not been satisfactorily wed. The relational theory can account for a large body of data relating the hippocampus and episodic memory, such as the important role of context in dictating item representations. However, the role of retrograde, catastrophic interference has not been addressed and computational studies show that preservation of the correlation structure of a set of memories within the same set of synapses imperils the stability of existing memories. On the other hand, the ‘Hebb-Marr’ model of episodic memory solves the issue of interference through the use of a dedicated hippocampal system that is distinct, though coupled, with a cortical network in which schema is realized. Unfortunately, the data on transitive inference and the observation of categorical coding schemes within the hippocampus challenge the notion of representational independence, a fundamental feature of the model that protects episodic memories from each other.

Therefore, there is a need for further recording experiments to be conducted in which related memories are sequentially learned to understand the contribution of prior knowledge on learning and the influence of learning on the stability of existing networks. Data of this type are presented in Chapters 3 and 4. My results support the relational theory and strongly argue against a model that strives for independent hippocampal codes for related experiences.

Before turning to the electrophysiological evidence for hippocampal schema, I will argue in the next chapter that a large body of work done on consolidation and reconsolidation has already shown a role of the hippocampus in updating and potentially creating related memory networks.

2 Reconsolidation and systems consolidation experiments support a schema model for hippocampal function

Most studies on memory consolidation consider the new information as if it were imposed on a tabula rasa, but considerable evidence indicates that new memories must be interleaved within a large network of relevant pre-existing knowledge. Early studies on reconsolidation highlighted that a newly consolidated memory could be erased after reactivation, but new evidence has shown that an effective reactivation experience must also involve memory reorganization to incorporate new learning. The combination of these observations on consolidation and reconsolidation highlights the fundamental similarities of both phenomena as the integration of new information and old, and it suggests reconsolidation = consolidation as a never ending process of schema modification.

I will present evidence that show a critical role of the hippocampus in schema modification, suggesting that the correlations between experiences are used by the hippocampus to update relevant schema via reconsolidation.

2.1 Systems Consolidation and the hippocampus

Since the pioneering studies on retrograde amnesia, it has been accepted that memories undergo a process of consolidation (Ribot, 1882; Müller and Pilzecker, 1900; Burnham, 1903). Immediately after learning, memories are labile, that is, subject to interference and trauma, but later they are stabilized, such that they are not disrupted by the same interfering events. It is well recognized that memory consolidation involves a relatively brief cascade of molecular and cellular events that alter synaptic efficacy as well as a prolonged systems level interaction between the hippocampus and cerebral cortex (McGaugh, 2000; Dudai, 2004); I will focus on the latter. As discussed in the Chapter 1, amnesic patients show a retrograde memory deficit with recall of information learned long before hippocampal damage spared relative to recently acquired memories.

A major limitation on studies of retrograde amnesia in humans is that there is no control over the extent of exposure to events during acquisition, as well as no control over how often the memories for those events are re-experienced or remembered. This problem has been addressed in several ‘prospective’ studies on amnesia in animals, where hippocampal damage occurs at different time points after learning and temporally graded amnesia emerges across multiple species and memory tasks (reviewed in Milner et al., 1998; but see Sutherland and Lehmann 2011; Broadbent and Clark, 2013). The duration of the systems consolidation period is highly variable across species and tasks. The evidence for temporally limited hippocampal involvement is compelling; however, this observation does not provide direct evidence on what brain areas support memory when the hippocampus is no longer necessary.

Insights about the relative engagement of other brain areas over the course of consolidation have come from recent experiments that have measured brain activation during memory retrieval at different times after learning in humans and animals. In humans, activation of the hippocampus during accurate memory retrieval in normal subjects was maximal for recent news stories and declined over approximately 9 years, parallel with the course of retrograde amnesia (Smith and Squire, 2009). Conversely, activation of widespread cortical areas was lowest for the most recent accurately remembered events and increased for more remote memories (see also Haist et al., 2001; Douville et al., 2005; Bayley et al., 2006). Recent ‘prospective’ studies using functional imaging have identified greater activation of the hippocampus during recall of recently over remotely studied paired associations and the opposite temporal gradient in cortical areas (Yamashita et al., 2009; Takashima et al., 2009). Furthermore, in the latter study, functional connectivity between the hippocampus and cortical areas decreased, whereas connectivity within the cortical network increased, over time following learning.

Studies on animals have employed 2-deoxyglucose (2DG) uptake and immediate early gene (IEG) activation as measures of neural activity in brain areas during memory retrieval for recently and remotely acquired memories. Bontempi et al. (1999) reported greater 2DG uptake in hippocampal area for recently acquired spatial discriminations, and conversely greater activation of frontal and temporal cortical areas for remotely acquired spatial memories. Similar patterns of decreasing hippocampal activation and increasing widespread cortical activation were observed using IEG expression as a reflection of neural activation in a contextual fear paradigm (Frankland et al., 2004; Maviel et al., 2004). Whereas the latter studies have focused on the medial prefrontal area as a key site for post-consolidation spatial memory, other studies have localized greater activation of olfactory recipient cortical areas for remote social-olfactory memories (Ross and Eichenbaum 2006), and greater activation of a higher-order auditory cortical area for remotely acquired tone-cued fear conditioning (Sacco and Sacchetti, 2010). The overall findings on cellular imaging studies in rodents impressively parallels the findings from functional imaging in humans, providing compelling evidence of systems consolidation characterized by early greater involvement of the hippocampus and later greater involvement of the task-relevant cortical areas.

Additional evidence for cortical-hippocampal interactions during consolidation comes from studies on hippocampal “replay” of memories during sleep and other off-line states, suggesting that the strengthening of cortical linkages depends on inputs from the hippocampus (e.g., Wilson and McNaughton 1994; reviewed in Carr et al., 2011). These interactions were highlighted in a study where, during sleep following maze running, populations of simultaneously recorded hippocampal and visual cortical cells fired in coordinated replays of the sequences of activity observed during awake behavior (Ji and Wilson, 2007). Additional support for the idea that hippocampal replay drives memory consolidation came from a report that replay following new spatial learning predicts subsequent memory performance (Dupret et al., 2010), and findings that stimulation-produced suppression of hippocampal sharp waves, when most replay events occur, impairs subsequent spatial memory, whereas stimulation at other times has no effect (Girardeau et al., 2009; Ego-Stengel and Wilson, 2010; see also Nakashiba et al., 2009). Other studies have focused on the cerebral cortex and shown development of a coordinated cortical neural network activation following learning (Alvarez and Eichenbaum, 2002; Takehara-Nishiuchi and McNaughton, 2008; Sakai and Miyashita, 1991), and cortical reorganization that depends upon an early tag within the regions that subsequently support the memory (Lesburguères et al., 2011; Tse et al., 2011).

These recent findings support the classic idea that a newly acquired memory depends initially on the hippocampus and eventually on widespread areas of the cerebral cortex. The classic and new observations have generated three current hypotheses about different aspects of the consolidation process (see Figure 1).

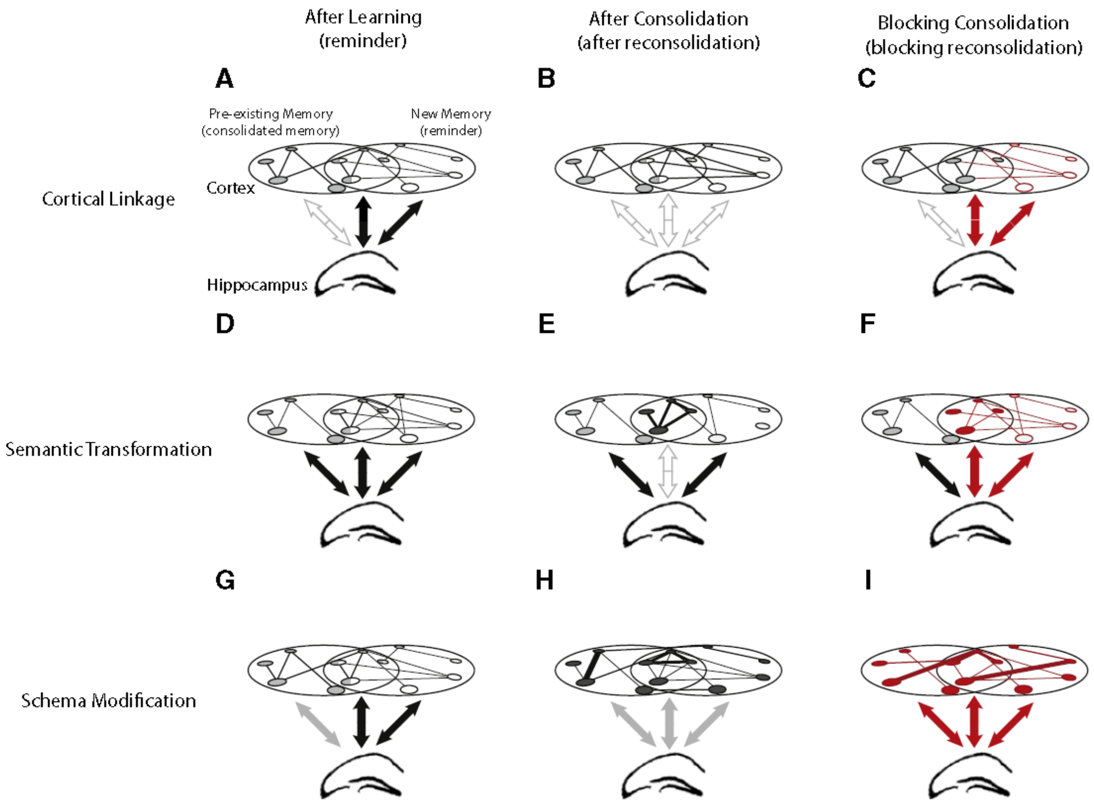

Figure 1 Three models of systems consolidation, each of which also supports aspects of reconsolidation (parentheses – see text). The hippocampus and an idealized network of connections among nodes (neurons) in the cortex represent a pre-existing memory (grey nodes), a newly acquired memory (white nodes), and shared components (grey/white nodes). The thickness of lines between cortical nodes reflects the strengths of intracortical connections. The dependence of cortical networks on the hippocampus is indicated by the arrows (dark = strong; gray = weak; open = no dependence). Throughout, red indicates a network disrupted by amnestic agents. A-C. The cortical linkage model. (A) Immediately after learning, the pre-existing memory has strong connections (thick lines) between nodes, whereas connections in the newly acquired memory network are weak (thin) and hippocampal dependent (dark arrows). (B) Consolidation is accomplished by strengthening intracortical projections (thick lines) and eventually the hippocampal connections are not required (open arrows). (C) Blocking consolidation prevents cortical strengthening of the new information (red, thin lines), leaving those cortical regions still dependent on the hippocampus (red arrow). Pre-existing memories (thick lines) that are supported by cortical networks are spared. D-F. The

semantic transformation model. (D) Immediately after learning, hippocampal connections are critical (dark arrows) for pre-existing episodic memories as well as for the newly acquired memory. (E) Consolidation involves strengthening (very thick lines) of intracortical connections shared by overlapping nodes to create a system that differs from either the first or second network (black nodes) and no longer depends upon the hippocampus (open arrow). Network elements that represent unique (episodic) portions of those memories remain weakly connected (thin lines) and dependent on the hippocampus (dark arrows). (F) Blocking consolidation prevents the formation of a new episodic memory stored within the hippocampal-cortical networks (red thin lines) and also the strengthening of the overlapping portion of the cortical network, leaving the newly formed network (red arrow), as well as the pre-existing cortical networks, hippocampal-dependent (dark arrows). G-I. The schema modification model. (G) Immediately after learning, the hippocampus is essential for supporting connections within the newly acquired memory, and important for the full elaboration of details of episodic memories (gray arrow). (H) During consolidation, the hippocampus supports both increases (thick lines) and decreases (thin lines) in connection strengths within the networks for pre-existing and new memories, resulting in an interleaving of the memories into a composite network (black nodes) that remains dependent on the hippocampus for the full range of detailed information (grey arrows). (I) Blocking consolidation after learning disrupts not only the consolidation of the newly acquired memory, but also corrupts the synaptic weight changes that support the interleaving of networks (reconsolidation), resulting in an altered set of cortical networks (red) still partially dependent on the hippocampus (red arrows).

2.1.1 Consolidation as linking cortical representations.

There are several variations of the hypothesis that the hippocampus rapidly stores critical information for linking cortical representations and that during multiple iterations of cortical-hippocampal interaction, connections within the cortex are strengthened and eventually support these associations in the absence of hippocampal function (Marr 1971; Squire et al., 1984; Teyler and DiScenna, 1986; Damasio, 1989; Squire, 1992). Each of these models proposes that, during learning, information from cortical areas that are activated in perceptual processing and working memory send inputs to the hippocampus, which encodes a “sketch” or “conjunction” of that information or “index” of loci within the cortex that contain the detailed information. During the consolidation period, memory cues that replicate partial information from the learning experience reach the hippocampus, activating the hippocampal representation or index, which, via back projections to the cortex reactivates the complete pattern of activations in cortical networks that were generated during learning (Figure 1A). Each time this reactivation occurs, intracortical connections between the disparate active cortical networks are gradually strengthened. After many such reactivations the intracortical connections are sufficiently strong to support reactivation of the entire set of cortical networks without assistance from the hippocampus (Figure 1B). Under this model, blocking consolidation prevents the strengthening of the intracortical connections for a newly acquired memory, but leaves pre-existing, memories intact (Figure 1C).

With regard to the functional imaging studies described above, it is notable that these models do not explicitly predict that the hippocampus should be less activated during effortful recall of remote memories. Indeed, a recent experiment showed increased c-fos expression in the hippocampus for older memories for the escape location on the Morris water maze (Lopez et al., 2012) and *greater* CA3 overlap in the cells that were active during encoding and retrieval for remote memories (Tayler et al., 2013). Furthermore, these models predict that the relevant cortical networks should be activated for both recent and remote memories, albeit those activations might be generated differentially through the hippocampus for recent memories and directly for remote memories. These issues remain to be resolved for models of the hippocampus as temporarily linking cortical representations.

2.1.2 Consolidation as semantic transformation.

The multiple trace theory, frequently opposed with the cortical linkage view, proposes that memories are qualitatively transformed from episodic memories into semantic memories during the consolidation period (Nadel and Moscovitch, 1997; Winocur et al., 2010). On this view, memories that are initially stored in cortical-hippocampal circuitry are episodic, defined as context-specific, and repeated “off line” reactivations create multiple distinct traces (Figure 1D, ‘New Memory’) from which the common information is extracted and integrated within pre-existing semantic networks in the cortex. Eventually the cortical representations that are common among memories, i.e., semantic memories free of episodic/contextual detail (Figure 1E, thick lines), do not depend on the hippocampus (empty arrow), but retrieval of episodic details continues to depend upon cortical-hippocampal connections (black arrows). In this model, blocking consolidation prevents the strengthening of intracortical connections that support semantic transformation, leaving new as well as remotely acquired episodic memories dependent on the hippocampus (Figure 1F, red).

In support of this view are reports that amnesic patients show temporally ungraded retrograde impairment for episodic memories (e.g., Rosenbaum et al., 2001; Steinvorth et al., 2005). However, contrary to the view that episodic and contextual memories always depend on the hippocampus, there are also findings of spared remote autobiographical memories in patients with medial temporal lobe damage (Bayley et al., 2003; reviewed in Squire and Bayley, 2007) and it is argued that flat retrograde gradients for episodic memory occur only following damage extending beyond the hippocampus into cortical areas (Reed and Squire, 1998). However, functional imaging studies have consistently reported that the hippocampus is activated for both recently and remotely acquired episodic and autobiographical memories (Ryan et al., 2001; Maguire et al., 2001; Piolino et al., 2004; Addis et al., 2004; Gilboa et al., 2004; Viard et al., 2007). These findings contrast with the above described observations of declining hippocampal activation during retrieval of famous faces and names, news events, i.e. semantic memories (Smith and Squire 2009; Haist et al., 2001; Douville et al., 2005).

Rodent studies also support the view that consolidation involves the semantic transformation of memories. In these studies, a memory that generalizes to testing conditions that differ from original training is typically considered an animal model of semantic memory. Parallel to the human literature, several experiments have shown that remote contextual memories become more generalized and independent of the hippocampus (Wiltgen and Silva, 2007; Wiltgen et al., 2010; Winocur et al, 2007; see also Lehmann et al., 2010; but also Weinberger et al., 2009). Conversely, hippocampal damage results in ungraded retrograde amnesia for spatial memories (Clark et al., 2005; Martin et al., 2005; Winocur et al., 2005a), except under circumstances of extensive and varied experience in environments wherein remote spatial memories are spared following hippocampal damage in both humans (Teng and Squire, 1999) and rats (Winocur et al., 2005b). Notably, these findings are also consistent with a simpler view that details of memories and information not repeated or contradicted across repeated experiences are most likely to be forgotten or overwritten, which also would be expected to result in a residual and strengthened semantic memory.

2.1.3 Consolidation as schema modification.

A distinct idea on memory transformation argues that newly acquired memories are not stored in isolation, but rather new memories are gradually incorporated into a "schema", an organization of related knowledge that contains semantic knowledge *as well as* episodic details. Unlike the semantic transformation view, schemas do not distinguish episodic and semantic memories, but rather interleave all memories via common elements and, unlike the focus on semantic transformation of multiple hippocampal traces, schemas involve the interleaving of new learning initially with previously acquired memories and subsequently with future memories. The schema idea, as originally proposed by Bartlett in 1932, was extended from the perspective of consolidation theory by McClelland et al. (1995), who contrasted rapid synaptic modification in the hippocampus with slowly modified connections within the cortex, and suggested that that the hippocampus supports memory for a brief period after learning, during which system reactivations integrate the new information via modifications of a pre-existing schema that connects related memories (Figure 1G,H). In this model, blocking consolidation disrupts the reorganization of pre-existing cortical representations and leaves newly acquired memories corrupted and dependent on the hippocampus (Figure 1I).

An important set of studies supporting the role of the hippocampus and the prefrontal cortex showed that both the hippocampus and the prefrontal cortex are required for rapid learning of novel locations of a flavored pellet, but only when those memories were learned within the same context as prior associations between flavor and location

(Tse et al., 2007, 2011). Learning in these experiments drove IEG expression in both the hippocampus and prefrontal areas, though in the mPFC expression was only observed when new flavor/place associations were learned within the same context where old associations had been made, suggesting a reorganization of mPFC traces that depends upon hippocampal function.

Brain imaging studies in humans have shown schema-dependent interactions between the hippocampus and ventromedial prefrontal cortex (VMPFC). Functional connectivity between these regions is low during repeat viewing of movies and high when subjects previously viewed a temporally scrambled version of the movie (van Kesteren et al., 2010), suggesting that prior exposure to the movie in the correct order established a schema that involved connectivity between the two regions. Furthermore, hippocampal activity during learning was lowest and VMPFC activity highest when item and context could be plausibly linked according to a common sense world schema (van Kesteren et al., 2013), suggesting that the hippocampus becomes active with schema incongruent learning. Other studies have found the converse result, that congruent information increases the correlation between hippocampal activity and subsequent memory (Staresina et al., 2009) and considering the hypothesis that lower activity may reflect efficient coding (Gonsalves et al., 2005; Demb et al., 1995), these results suggest an important role for the interactions between the hippocampus and prefrontal cortex in the formation and update of memory schema.

The three models of consolidation described above are not mutually exclusive. The hippocampus plays a key role in linking elements of memories processed in the cortex, including links that compose representations of discrete events and representations of episodes composed of sequences of events (Eichenbaum, 2004). Memories interact through ‘nodal’ representations of features common to multiple experiences. Importantly, these common nodal elements characterize information that is not bound to a particular event or episode, and is consistent across experiences, and in that sense underlie a “semantic transformation”. Also, it is precisely via the nodal elements that memories are connected and therefore underlie the structure of schemas. The evidence presented above suggests a critical role for the hippocampus in the establishment of the cortical nodes that link and relate disparate experiences. As illustrated in Figure 1, the different models of consolidation may best be viewed as focusing on different aspects of the larger process by which memories are interleaved during consolidation.

2.2 Reconsolidation studies suggest that learning causes modification of cortico-hippocampal networks that store related memories

Standard consolidation theories described above characterize consolidation as a one-time event, after which a memory is impermeable to subsequent disruption. However, this view was challenged in the late 1960’s by studies reporting that presentation of a “reminder” cue made a completely consolidated memory again labile to the same agents that would block consolidation (Misanin et al. 1968, Schneider and Sherman, 1968). In 2000, Nader and colleagues raised this challenge again in experiments that targeted the known role of the amygdala in synaptic consolidation of the Pavlovian association of a tone (CS) with a shock (US; Falls et al., 1992; Duvarci et al., 2006), showing that a CS alone reminder delivered long after consolidation was complete re-engaged the temporary susceptibility of the memory. They hypothesized that the reminder reactivated the original memory trace, making it necessary to "reconsolidate" the memory, or else suffer erasure of the memory. Over the last decade, many experiments have supported the observation of memory susceptibility following reminders and these findings have been reviewed extensively in recent papers (Nader and Hardt, 2009; Dudai and Eisenberg, 2004; Lee, 2010; Alberini, 2011; Sara, 2010).

Results supporting the existence of reconsolidation have been reported in several species across a broad range of learning tests, and using a variety of manipulations to block memory (e.g. Rose and Rankin, 2006; Pedreira et al., 2002, Eisenberg et al., 2003; Frankland et al., 2006; Lee et al., 2005; Hupbach et al., 2007; Monfils et al., 2009; Schiller et al., 2010). Despite broad support for the generality of reconsolidation (Nader 2009), several studies have failed to find that amnestic agents block memory in the reconsolidation paradigm (Biedenkapp and Rudy, 2004) or have observed that the memory deficits are temporary (Lattal and Abel, 2004; Power et al., 2006), leading to the idea that the reconsolidation phenomenon has 'boundary conditions' (Eisenberg et al., 2003; Milekic and Alberini, 2002; Morris et al., 2006). Several experimental parameters have been shown to be important in determining whether reconsolidation occurs, including how memories are reactivated (Debiec et al., 2006; Tronel et al., 2005), whether novelty is introduced during memory reactivation (Pedreira et al., 2004), and the age and strength of a memory (Eisenberg et al., 2003; Milekic and Alberini, 2002). I consider two main categories of boundary conditions: which memory is active at the time of amnestic treatment and whether the reminder generates new learning.

Early in the recent series of studies on reconsolidation there were conflicting reports on whether reminders reinstated lability of memories for classical aversive conditioning. Several studies (Berman et al., 2003; Vianna et al., 2001; Koh and Bernstein, 2003; Pedreira and Maldonado, 2003) noted that, in the Pavlovian conditioning studies, presentation of the CS alone can have two opposing effects: it can act as a reminder to engage the original memory trace, and it can generate extinction, which involves development of a new and competing memory trace. Based on the dual roles of CS alone presentation, Eisenberg et al. (2003) suggested that the effects of amnestic agents differ depending on whether the original memory trace or the newly developed memory for extinction was dominant at the time of amnesic treatment. To test the trace dominance theory, subjects were given either more initial CS/US training or more CS alone trials after initial conditioning, with the assumption that more initial training would cause the fear memory to dominate during the reminder, while extinction memory would dominate after more sessions with the CS alone. Consistent with the trace dominance hypothesis, more CS/US pairings resulted in disrupted reconsolidation of the original aversive memory whereas more CS alone presentations resulted in subsequent loss of extinction and preserved fear memory, in different species and different memory tests (Eisenberg et al., 2003). These findings can also explain why extensive training and/or specific time periods between initial training and reminder could result in strong, original memory traces that are reactivated as dominant following a reminder (Suzuki et al., 2004; Wang et al., 2009; Milekic and Alberini, 2002; Eisenberg and Dudai, 2004; Robinson et al., 2010, but see Duvarci et al., 2006), and why effective reminders must be presented for reconsolidation of the original memory (Bozon et al., 2003).

The other major factor in determining the efficacy of amnestic agents in the reconsolidation protocol is whether the reminder event involves new learning in addition to recovery of the initial memory trace. One study reported that whereas original memories are blocked by an amnestic agent following a CS alone reminder, there was no loss of the original memory following reminder presentations that involve a combination of CS and US presentations, suggesting that CS alone reminder constituted a new learning experience (Pedreira et al., 2004). However, there are several examples of successful disruption of reconsolidation following presentation of both a CS and US (Duvarci et al., 2004; Rodriguez-Ortiz et al., 2008; Valjent et al., 2006). In these studies, it is not clear that performance was at asymptote, leaving open the possibility that new learning still occurred during the reminder event, a factor that proved critical in another study (Rodriguez-Ortiz et al., 2005). Also, Morris et al. (2006) directly compared reconsolidation following reminder trials in rats trained to asymptotic performance in standard (“reference memory”) water maze task versus a (“working memory”) variant of the task where new escape locations were learned daily, and found that anisomycin was effective after reminders only in the condition of new learning each day. Also, in other studies on human declarative and motor memory, providing subjects with a reminder that involves new learning is key to alteration of existing memories (Walker et al., 2003;

Hupbach et al., 2007, 2008, 2009; Forcato et al., 2009). These findings and several other studies indicate that learning during the reminder session is a critical boundary condition for reconsolidation (Winters et al., 2009; Robinson et al., 2010; Lee, 2010).

The combination of requirements for dominance of the original memory and new learning suggest that the key conditions for blockade of reconsolidation involve a re-activated memory trace that is susceptible to modification and new related learning that occurs during the interfering event. Thus, the encoding of new information occurs within the context of retrieval, and the circuits that are modulated by new information are the ones that are activated by the reminder. At the same time, blockade of reconsolidation is only observed in conditions that favor new learning related to the re-activated memory (e.g. additional training, extinction), suggesting that reconsolidation involves some kind of reconciliation or integration of a vulnerable memory trace and new relevant information (Eichenbaum, 2006).

2.3 Reconsolidation and Systems Consolidation: the hippocampus can become re-engaged when new information must be incorporated within existing networks

The three models of systems consolidation introduced earlier differ in the nature of interactions between pre-existing and new memory networks and their dependence on the hippocampus. In the cortical linkage model, consolidated memories are independent of the hippocampus (Figure 1A,B). Therefore, in a reconsolidation protocol, amnestic agents delivered to the hippocampus could only affect the newly acquired network, that is, the reminder, but leave intact the previously consolidated memories (Figure 1C). This outcome is not consistent with the findings that even consolidated memories are affected by reminders and damage to the hippocampus (Debiec et al., 2002; Winocur et al., 2009). Theories that hypothesize elimination of hippocampal connectivity to cortical networks during systems consolidation must be somehow updated to incorporate the findings that even consolidated memories can regain hippocampal dependence after a reminder (systems reconsolidation).

In the semantic transformation model, newly acquired memories are overlaid with pre-existing semantic memory networks, such that the common elements and connections become hippocampal-independent and semantic (Figure 1D, E). Reconsolidation has been suggested as having two roles: to potentiate intracortical connections to form semantic memories; and, to strengthen episodic memories when new learning, or a reminder, re-engages the hippocampal networks active during original learning (Figure 1E; Hupbach et al., 2007; Winocur et al., 2009). In this scheme, hippocampal amnestic treatments after a reminder should block the retention of any new episodic memory, prevent new semantic memory formation, and disrupt reconsolidation of other, similar episodic memories (Figure 1F, red); pre-existing, semantic memories would be left intact. Initial support for these claims came from a study that found systemic, but not hippocampal, delivery of amnesic agents disrupted reconsolidation of remote fear memories, findings that were interpreted to be evidence of corrupted semantic memory (Frankland et al., 2006). However, other studies showed that hippocampal reconsolidation is necessary for consolidated memories (Debiec et al., 2002; Winocur et al., 2009), and a recent experiment on remote memories showed that the generalized, ‘semantic’, fear responding that normally occurs in non-conditioned contexts was also dependent on hippocampal reconsolidation (Winocur et al., 2009). Therefore, somehow pre-existing semantic networks must become hippocampus-dependent, a condition that counters predictions of the original theory (Nadel and Moscovitch, 1997). Indeed a recent optogenetic study has demonstrated that remote contextual fear memories thought to be independent of hippocampal function using standard pharmacological methods are revealed to actually depend upon the hippocampus using very rapid (ms) inhibition upon entry to a conditioned context (Goshen et al., 2011).

In the schema modification model, consolidation occurs by integrating the new memory into active, pre-existing memories via reorganization of common elements within the hippocampus and the cortex (Figure 1G, H). In reconsolidation experiments, the reminder determines which memories will be active during encoding and therefore which synapses will be affected by the new learning (Figure 1G). In this model, systemic amnestic treatment after a reminder would result in a partial integration of the newly learned information into the hippocampal and cortical networks, resulting in a corruption of the reorganizing network (Figure 1I, red). Manipulations limited to the hippocampus could cause disruption of cortical reconsolidation due to interrupted replay of the newly acquired learning (Eichenbaum, 2006) or errant discharges from a damaged hippocampus driving molecular changes in reorganizing cortical circuits (Rudy and Sutherland, 2008) or perhaps another mechanism that would affect cortical circuits undergoing plastic remodeling. While each of the models described here captures some of the phenomenology of reconsolidation experiments, none has compelling support, and this is likely to remain the case until we better understand the nature of neural representations in the hippocampus and cortex and how they change during consolidation and its breakdown.

2.4 Conclusion

Taken together, the findings on blockade of reconsolidation following molecular interventions, hippocampal lesions, and inference has led several to suggest that reconsolidation normally involves an “updating” of memories (Lewis, 1979; Sara, 2010; Morris et al., 2006; Lee 2009, 2010; Dudai and Eisenberg, 2004). It has been suggested that updating can occur via two mechanisms, a destabilization of existing memory traces and modification of the contents of the original memory to add new related material (Lee 2008, 2010). Common among these views is the idea that reconsolidation is the mechanism by which initially consolidated memories are changed with new learning.

I take a different view and propose that even initial consolidation occurs through a reorganization of pre-existing memories. Thus, while there is still much to be discovered about the mechanisms of consolidation and reconsolidation, I suggest that it would be valuable to consider that reconsolidation = consolidation. Dudai and Eisenberg (2004) adopted a very similar hypothesis, suggesting that reconsolidation is a manifestation of a "lingering" consolidation process. Here I take this idea one step further and suggest that reconsolidation is the never-ending consolidation process. When we refer to consolidation, we cannot consider new learning to occur in a *tabula rasa*. Rather, the consolidation of new learning, the first life of a memory, is a reorganization (and therefore a “re”-consolidation) of the existing schema. Correspondingly, after the new learning has been consolidated into the existing schema, reminders and new related experiences normally constitute memories that must be consolidated by further reorganization of the current relevant schema. Counter to models that state independence of hippocampal traces (see section 1.3), I suggest that reconsolidation will induce plasticity preferentially in pre-existing hippocampal networks that are relevant to learning. The fundamental conclusion here is that new information is continually being integrated and thereby repeatedly consolidated in a never-ending reorganization of memory networks.

3 Learning Causes Reorganization of Neuronal Firing Patterns to Represent Related Experiences within a Hippocampal Schema

3.1 Introduction

Despite evidence that the hippocampus plays a central role in schemas, our knowledge about how neuronal networks in the hippocampus integrate memories into schemas has received little attention. Indirect evidence has come from studies that observe a “partial remapping” of hippocampal spatial firing patterns following alterations in salient spatial cues, which could reflect both assimilation of new information and accommodation of existing network representations (reviewed in Eichenbaum et al., 1999). Several other studies have reported similarities in the firing patterns of hippocampal neurons associated with traversing similar routes through different mazes (Singer et al., 2010), with objects in different locations (Wood et al., 1999) or in different examples of the stimuli with the same meaning (Quiroga et al., 2005; Hampson et al., 2004). The common coding of events that are shared among memories could act as “nodes” that connect different memories within schematic networks. Yet, no studies have examined how common representations develop as new memories are added to a pre-existing framework. In the present study I extend previous findings on hippocampal neurons that develop firing patterns associated with goal locations when goals are moved in water mazes (Hollup et al, 2001) and open fields (Dupret et al., 2010), and I explore how the hippocampus represents the addition of new goals within an environment for which the spatial representation of pre-existing goals has been characterized. Pursuing the neural representations that underlie characteristics of schemas described by Piaget

(1929), I explored how new goals are assimilated into a pre-existing hippocampal schema and how the original schema is modified to accommodate the new related information. The results show that a pre-existing schema assimilates new goal-memories while distinguishing pre-existing memories, and accommodates its structure to integrate the new knowledge within the pre-existing network representation.

3.2 Methods

3.2.1 Subjects. Five male Long Evans rats (Charles River Laboratories, Wilmington, MA) were housed within the Laboratory Animal Care Facility on the Boston University Charles River Campus. Subjects were individually housed and kept on a 12 hour light/dark cycle. During behavioral testing, subjects were food and water deprived ensuring a minimum of 85% free feeding weight (300-500g) with access to 15 minutes of water per day. All animal protocols were approved by the Boston University Charles River Campus Institutional Animal Care and Use Committee.

*3.2.2 Surgery***.** Rats were anesthetized using a mixture of 3% isoflurane in oxygen and were injected with atropine (0.03mg/kg SC) and buprenorphine (0.1mg/kg SC) prior to surgery and meloxicam (1mg/kg, SC) immediately following. An ophthalmic ointment was applied to the eyes and temperature was maintained between 37°C and 38°C. Ringer's solution was injected to maintain hydration. Stainless steel bone screws, including two ground screws, were inserted into the skull. A 2.0mm diameter hole was drilled into the skull using a dental drill and the dura was removed. The base of the microdrive was aimed at the CA1 region of the left dorsal hippocampus (−3.5 AP, +1.9

ML). The craniotomy was sealed with Kwik Sil (World Precision Instruments, Inc., Sarasota, FL), the microdrive and ground wires were secured in place using dental acrylic (Henry Schein, Melville, NY) and C&B Metabond (Parkell, Inc, Edgewood, NY), and the surgical site was sutured shut.

*3.2.3 Electrophysiological recordings***.** Individually movable 24-tetrode microdrives were built in-house. Tetrodes were spun from 12.5µm NiCr Kanthal fine wire (Sandvik, Palm Coast, FL) and gold plated for an impedance of 200kΩ at 1kHz. At the end of surgery, each tetrode was lowered ~850 µm into the brain. After 5 days of recovery, the tetrodes were lowered over 7–14 days toward the CA1 layer. After the experiments, 25 µA of current was passed through each tetrode for 30 s before perfusion and histological confirmation of tetrode placement.

As previously reported (MacDonald et al., 2011), the electrical signal was referenced to a common skull screw and differentially filtered for single unit activity (154 Hz–8.8 kHz) and LFPs (1.5–400 Hz). The amplified spikes from each wire were digitized at 40 kHz while the field potential was digitized at 1 kHz and monitored with the Multineuron Acquisition Processor (Plexon Inc., Dallas). Individual pyramidal neurons were isolated using Offline Sorter (Plexon Inc.) by visualizing combinations of waveform features (peak-valley, valley, peak, principal components, and time-stamps) extracted from wires making up a single tetrode (i.e., “manual cluster cutting”). Single-neuron selectivity was verified by the interspike interval histograms that contained no successive spikes within a 2 ms refractory period. Single-neuron stability was verified by comparing cluster stability across a recording session.

3.2.4 Behavioral paradigm. Prior to surgery, water-deprived rats were trained to retrieve water at several locations on a symmetrical wooden circular track (107cm in diameter) that contained 20 potential reward sites. The reward sites were small water dishes with adjacent LEDs connected by an angle bracket. The location of the goal for each trial was chosen pseudo-randomly among 4-5 evenly spaced locations, with the goal never repeating more than three consecutive times. Each trial began with a blinking LED (300ms on, 700ms off) at the selected location and, when the rat was within 15cm of an LED-cued goal location, the light was held on for 5 seconds, after which 75μL of water was delivered from the adjacent water port. If the rat did not wait at the reward site within two minutes, there was a 15s ITI and then the task progressed to the next trial.

Following surgery, rats were retrained on the LED-cued version of the task. When tetrodes were positioned in dorsal CA1, three new spatially defined goal locations (SP1, SP2, SP3) were introduced over the course of a week; these goal locations were not cued by an LED and could only be identified by memory of their location relative to extra-maze cues (overhead lights and large white cue card). On Day 1, 30 LED-cued trials were presented and then SP1 was introduced with water delivered to that site prior to the rats' arrival. Rats initially found the new site by first searching unsuccessfully for water at the well-learned cued sites and then searching each well around the maze. After finding the new spatial goal, SP1 and LED-cued trials alternated such that, after completion of each LED-cued trial, the rat could run to the SP1 location and await a reward delivery; over the course of the first seven SP1 trials, the interval between arrival to the SP1 site and reward delivery ramped from 0.3 - 5s. Data collected at the new goal sites during ‘ramped-training’ trials were not included in neural or behavioral analysis. On Day 2, rats were trained to alternate between SP1 and LED-cued locations for the shorter of 100 trials or 60 minutes. On Day 3, the rats were given a block of 30 SP1 and LED-cued trials before SP2 was introduced. Between LED-cued trials, rats were required to first visit SP2 then SP1, waiting at each location for 5s before a reward was presented. On Day 4 the rat continued to be trained on SP2-SP1 and LED-cued trials. Day 5 began with a block of 30 SP2-SP1 and LED-cued trials, before SP3 was introduced. Between LED-cued trials, rats were required to visit three spatial locations in the order SP3-SP2-SP1, waiting for 5 sec to receive water at each place. Day 6 involved continued training on the three spatial locations along with the original LED-cued locations. Following this sequence, rats were overtrained on the sequence of three positions for 1-12 days and then given at least two days on LED-cued trials only before being trained on a second sequence of spatial locations that never overlapped with LED-cued locations. One rat was also trained on a third sequence.

To assess the development of a schema as reflected in facilitation of new spatial sequence learning following initial sequence acquisition, I compared the latency to arrive at SP1 leading to a reward for trials on the first sequence to that for the second sequence over initial blocks of training trials (Figure 2E). Analyses of learning focused on the trials after ‘ramped-training’ when rats were required to wait five seconds at the appropriate reward site.

Neural data with only SP1 and the cued sites are reported from 11 sessions from five rats. Data with SP2 and SP1 are reported from nine sessions with four rats and neural data with SP1, SP2 and SP3 are reported from seven sessions from three rats. Data are reported from 15 learning sessions (either SP1, SP2 or SP3) from five rats, seven sessions in which the new goal location is one day old (from four rats), three sessions when the new goal is two days old (from three rats), and 27 sessions in which the new goal is 8.34±3.61 days old (from three rats).

3.2.5 Analyses. The timestamps of unit action potentials, rat position, and reward delivery times were imported into MATLAB R2012b for all analyses.

Behavioral events were distinguished as periods of waiting (WAIT) when the animal arrived at a goal and remained within 15 cm of the goal site for 5 seconds prior to, and not including, reward delivery, and epochs of running (RUN) at velocity over 10cm/s within 4.2cm of a goal site, excluding events 7s before through 3s after reward delivery.

Three spatial firing rate maps were constructed: one for WAIT events, one for clockwise RUN events, and one for counter-clockwise RUN events. Occupancy on the circle track was linearized and binned at 8.4cm and not smoothed. Analyses were performed separately for rate maps in each running direction and then averaged for a single RUN metric.

Temporal analysis of WAIT epochs was based on peri-event time histograms (PETHs) centered on reward delivery and binned at 250ms. I also defined cells as activated during WAIT events if a cell had an average firing rate above 2Hz and 30% of the maximum firing rate around reward delivery in any 250ms bin during the WAIT epoch (-5s to 0s to reward delivery). The maximal firing rate was calculated ±20s around reward delivery.

The local field potential (LFP) was taken from a channel with spiking units. Theta power was the power of the LFP band-pass filtered between 6-10Hz. Sharp-wave ripples were events that were seven standard deviations higher than the average LFP filtered at 140-200Hz with minimum duration of 20ms and maximum duration of 150ms. The time of the ripple was the timestamp of the maximum waveform.

To probe the organization of the schema I assumed that similar neuronal activity patterns observed when the rat is at different locations encode the common features of those events, whereas divergent firing patterns represent the distinct features of those events. The commonality and divergence of representations at multiple locations were assessed in six analyses that emphasize different aspects of similarity in neuronal firing patterns:

3.2.5.1 Cross-place activity. The similarity of temporal patterns of firing at multiple locations was measured by pair-wise multiplication of the firing rate histograms (-10s to +5s centered on reward delivery and binned at 250ms) at each location (mean goals/session = 5.78, range 4-8), normalized by the maximal firing rate (Figure 3D;

Figure 7 B,C). The average cross-place activity ($\overline{\mathbf{CPA}}$) was computed between every pair of locations for each cell during rewarded epochs.

$$\overline{\mathrm{CPA(t)}} = \langle \frac{\lambda_i(t) * \lambda_j(t)}{\lambda_{max} * \lambda_{max}} \rangle$$

$\lambda_i(t)$ is the trial averaged firing rate at position i at time t relative to reward onset. $\lambda_j(t)$ is the firing rate at the equivalent time t relative to reward onset at position j. λ_{max} is the maximum average firing rate observed for that cell in any time bin at any location. Firing rates below a sliding threshold were then set to zero, driving the product to zero, and resulting in a metric of cross-place activity over time that captures when high activity is present in two locations. With a threshold up to 5Hz, a significant increase in CPA was observed upon arrival to the goal location.

3.2.5.2 Firing pattern selectivity index. The degree to which cells had different firing patterns across locations was measured with a selectivity index (SI; Komorowski et al., 2009; Wirth et al., 2003):

$$\mathrm{SI(t)} = \frac{n - \sum_{i=1}^{n} \left(\lambda_i(t)/\lambda_{pref}(t)\right)}{n-1}$$

Where n is the number of goal locations, $\lambda_i(t)$ is the trial averaged firing rate at position i at time t relative to reward onset, and $\lambda_{pref}(t)$ is the maximal trial averaged firing rate at any position at time *t*. SI values range between 1, indicating activity at only one location, to 0, indicating equivalent activity at all locations (Figure 4 A,B).

3.2.5.3 Spatial selectivity index. Spatial selectivity was calculated based on each of the three firing rate maps. Unlike the firing pattern SI, λ_i and λ_{pref} did not vary over time and were based upon the average firing rate recorded at the goal locations (8.4 cm bin) during either WAIT map (Spatial SI WAIT) or the RUN maps (Spatial SI RUN) (Figure 4E).

3.2.5.4 Temporal pattern ensemble correlation. Population vectors were composed of the mean firing rates for simultaneously recorded cells at each reward location at each time bin centered on reward delivery. The average similarity over time, $\overline{\boldsymbol{COR}(\boldsymbol{t})}$, was measured by calculating the Pearson correlation coefficient of the population vectors, $\boldsymbol{X_i}(\boldsymbol{t})$ **and** $\boldsymbol{X_j}(\boldsymbol{t})$, at every time point, *t*, at pairs of locations, i and *j*. Only ensembles of cells with a more than eight units with a baseline rate less than 4Hz were considered (22 ensembles, avg. number of cells 14.9, STD 5.63, range 8-30). Low correlations indicate distinctive coding (Figure 4F).

$$\overline{COR(t)} = \langle corr\big(X_i(t), X_j(t)\big)\rangle$$

3.2.5.5 Spatial ensemble correlation. As with the selectivity index, ensemble correlations were also calculated using the same three spatial firing rate maps: the WAIT map and the two RUN maps. The spatial ensembles were constructed by average firing rates recorded at the goal locations (8.4cm bin) during either WAIT or RUN. The spatial ensemble correlation was calculated by taking the average Pearson correlation coefficient of the spatial ensembles recorded at each goal site (Figure 7 D,E,F).

3.2.5.6 Bayesian decoding. The extent to which firing patterns across locations differ should be reflected in accurate decoding. The spatial posterior distribution was calculated using the firing rates from both RUN and WAIT epochs using every other 500ms block of spiking and positional data. Time was binned at 250ms and the spatial firing rate map was binned as described above. Decoding accuracy was defined as the conditional probability of occupying the actual location given the observed spike counts for the time blocks not used to create the posterior distribution (Zhang et al., 1998) (Figure 5). Decoding was only applied to sessions with ensembles of cells of more than eight units with baseline rates less than 4Hz.

3.2.5.7 Changes in firing patterns on the day of learning. To measure the extent to which neuronal firing patterns changed with learning, spatial firing rate maps were made for all periods when the rat's velocity was over 10cm/s before and after introduction of the new goal or for the first and second half of a session without new learning. Remapping was measured by Pearson’s correlation of each cell's before and after rate maps.

To measure learning-related changes in firing patterns during WAIT events, ensemble vectors of the average firing rates recorded at each WAIT site before and after learning were correlated (Figure 7F).

3.2.5.8 Changes in firing patterns over days. Two methods were used to study how cells encoded WAIT events across days. First, the spatial ensemble correlation was used to compare the new WAIT site to each of the well-learned, LED-cued sites. The average of these pair-wise correlations was used as a metric for ensemble discrimination of the new WAIT event (Figure 7E). Second, I used a single cell metric to analyze cells that were active at the new WAIT site on every day after learning. For these active cells, a cell's average rate at the new site (8.4cm bin) was subtracted from its rate at the other LED-cued WAIT sites. The minimum firing rate difference between the new and original sites was then calculated to measure the degree to which new WAIT site was differentiated from any other WAIT site.

3.2.5.9 Principal Component Analysis. For visualization purposes, PCA was used to reduce the N-dimensional (N = number of cells in the ensemble) representation of the goals. PETHs recorded on each trial were converted into z-scores based on overall firing rate means and standard deviations. The firing rates were then transformed into principal component space using the covariance of firing rates between neurons recorded ±20s around goal delivery. For each goal location, the trial-averaged scores of the first two components are plotted as a function of time to reward (Figure 8B).

To show how the spatial representation of new and old WAIT events changed over days, firing rate maps for each cell were computed as described and then z-score normalized. Ensemble vectors were created using the average rates for each cell at each goal location. These vectors were then transformed into principal component space based off of the covariance of the firing rates recorded across all locations. For each goal location, the first two components of the spatial ensemble representation are plotted (Figure 8C).

3.2.5.10 Significance testing. Metrics that were normally distributed were compared using ANOVAs and student's t-test. However, most metrics were not normally distributed and therefore significance testing of means was done using Monte Carlo sampling with replacement with 10000 repetitions (bootstrap). Changes from baseline were measured as the difference between the averages at every time point against a baseline calculated from the average -20 to +20s around reward delivery.

3.3 Results

3.3.1 Rats rapidly new spatial goal locations.

Rats were trained on a task that required them to learn new goal sites in an environment with several well-learned goal locations (see Figure 2). Rats first learned to wait five seconds at one of several LED-cued locations to retrieve a water reward. Over the course of six days, 1-3 new spatially defined goal sites were progressively introduced (SP1, SP2, SP3). Rats learned these new goal locations as reflected in significantly increased percentage of time spent in the rewarded area (not including reward consumption) after introduction of the reward (Figure 2 C,D; *2-way ANOVA main effect of time* $F(1,24) = 91.55$, $p = 11.6^{-9}$).

Schemas are commonly characterized as facilitating new learning that is consistent with prior accumulated knowledge (e.g., Tse et al., 2007). Here, to examine whether initial acquisition of a spatial goal sequence facilitated subsequent learning of second spatial goal sequence, I compared the latency to arrive at the first spatial goal (SP1) of the first learned spatial sequence to that for the second spatial goal sequence. As shown in Figure 2E, during initial spatial sequence acquisition, rats arrived in substantially longer latencies in first trials after the response requirement reached 5 sec (see Methods), and thereafter latencies shortened. During learning of the second spatial sequence, unlike in the first sequence, latencies were initially low (*2-way ANOVA: main effect of sequence, $F(1,3) = 19.93$, $p = 0.0004$; time X sequence interaction, $F(2,3) =0 .0038$, $p = 0.0038$; post-hoc paired t-test for trials 1-4, $t(3) = 5.11$, $p = 0.014$)* and remained low throughout testing. Savings in learning, as reflected by the significantly lower latencies during the second sequence, suggest that rats developed a schema for learning goal sequences that occurred within the familiar context.

3.3.2 Hippocampal neurons activate during multiple well-learned WAIT events

I first characterized how learned WAIT events are organized within a well-established schema of goal locations for which the rat has had at least two days of training. Within a well-established schema, it is expected that related memories will be both differentiated and linked via “nodes” that are characterized by similar representations of the common features between events. To investigate whether hippocampal networks contain a schema of WAIT events at multiple goal locations, I first asked whether the same behaviors at different locations were represented similarly by firing patterns of hippocampal neurons. Then I also asked whether these WAIT representations also differentiated distinct memories of goal locations.

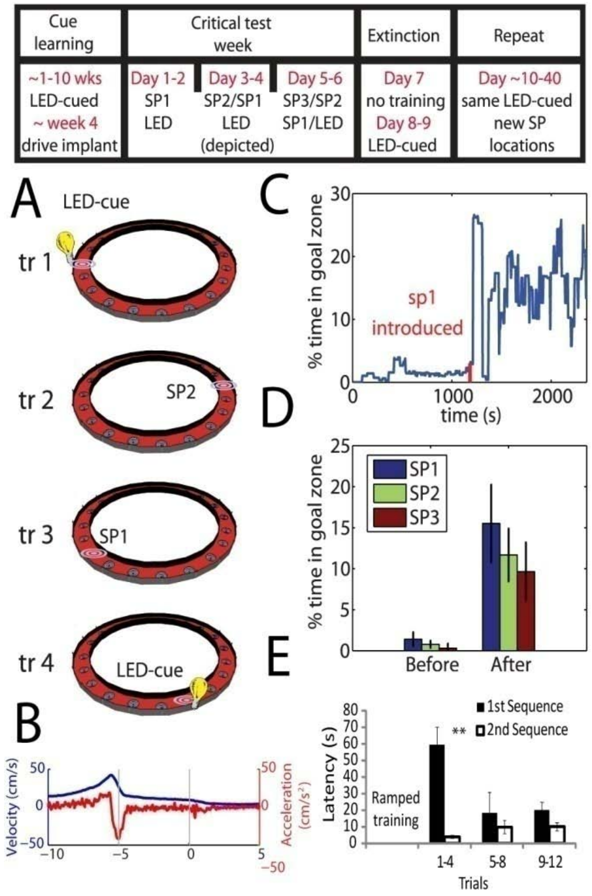

Figure 2 Summary of behavior at goal sites. (A) The rat must visit the two spatial positions (SP1 and SP2) in order, waiting at each for 5s for water reward. The next sequence begins with a LED-cued trial in one of four positions chosen pseudo-randomly. (B) The average velocity and acceleration profile as the rat approached and waited at the goal sites. (C) The percentage of time a rat spent in the new spatial position not including reward consumption in a sliding 2 minute window throughout an example session. Red tick marks indicated the time in which the new spatial location was first introduced. (D) The average percentage of time rats spent in the new spatial location before and after its introduction for all rats and for one, two or three spatial goal locations, SP1, SP2, and SP3, respectively. (E) The latencies to arrive at SP1 for the first and second spatial sequence significantly differed in the early trials (trials 1-4) at the full 5 sec wait period thus showing that learning one spatial sequence facilitated the learning of another. ** p = .01.

Several single neuron and ensemble analyses were adopted to characterize firing patterns of putative CA1 pyramidal cells (*n = 410, avg. firing rate = 0.40 ± 0.021 Hz*) through time and space. These analyses focused on WAIT events when rats stopped at the goal loci prior to reward delivery and RUN epochs when rats ran through the locations and no reward was available. WAIT epochs allowed me to assess firing associated with anticipatory behavior at the goal location, whereas RUN epochs provide a baseline of spatial firing without the anticipatory behavior.

Examination of firing patterns before and during WAIT events revealed that many cells fired at multiple goal locations (Figure 3A). For example, Cell 1 in Figure 3A had two firing fields around the spatially separate goals sites 1 and 12 and did not fire at other goal locations, while Cell 2 fired at goals 1, 17 and 12. An analysis of the number of locations where activation occurred during WAIT events revealed that 40% of neurons were active during WAIT events at one location, 15.6% at two locations, and 9.8% at more than two goal locations, indicating that 39% of the neurons that were active during at least one WAIT event fired during more than one WAIT event.

Firing associated with WAIT events was not uniform in time such that, upon arrival to a subset of goal areas, neurons often released a burst of spikes followed either by a period of decaying activity or silence (Figure 3 B,C). This firing pattern can be observed clearly in the firing behavior of Cell 6 in Figure 3B, which increased firing rate upon arrival to locations 1 and 12 and decreased firing at the other two locations. Comparing the average response from all cells that activated during WAIT events versus those that did not, activation during WAIT events was characterized by an increase in firing rate upon arrival (5s before water delivery) to the goal site followed by a slowly tapering level of firing during WAIT events. In contrast, cells that did not exceed the activation threshold at WAIT locations decreased firing rates from baseline (Figure 3C).

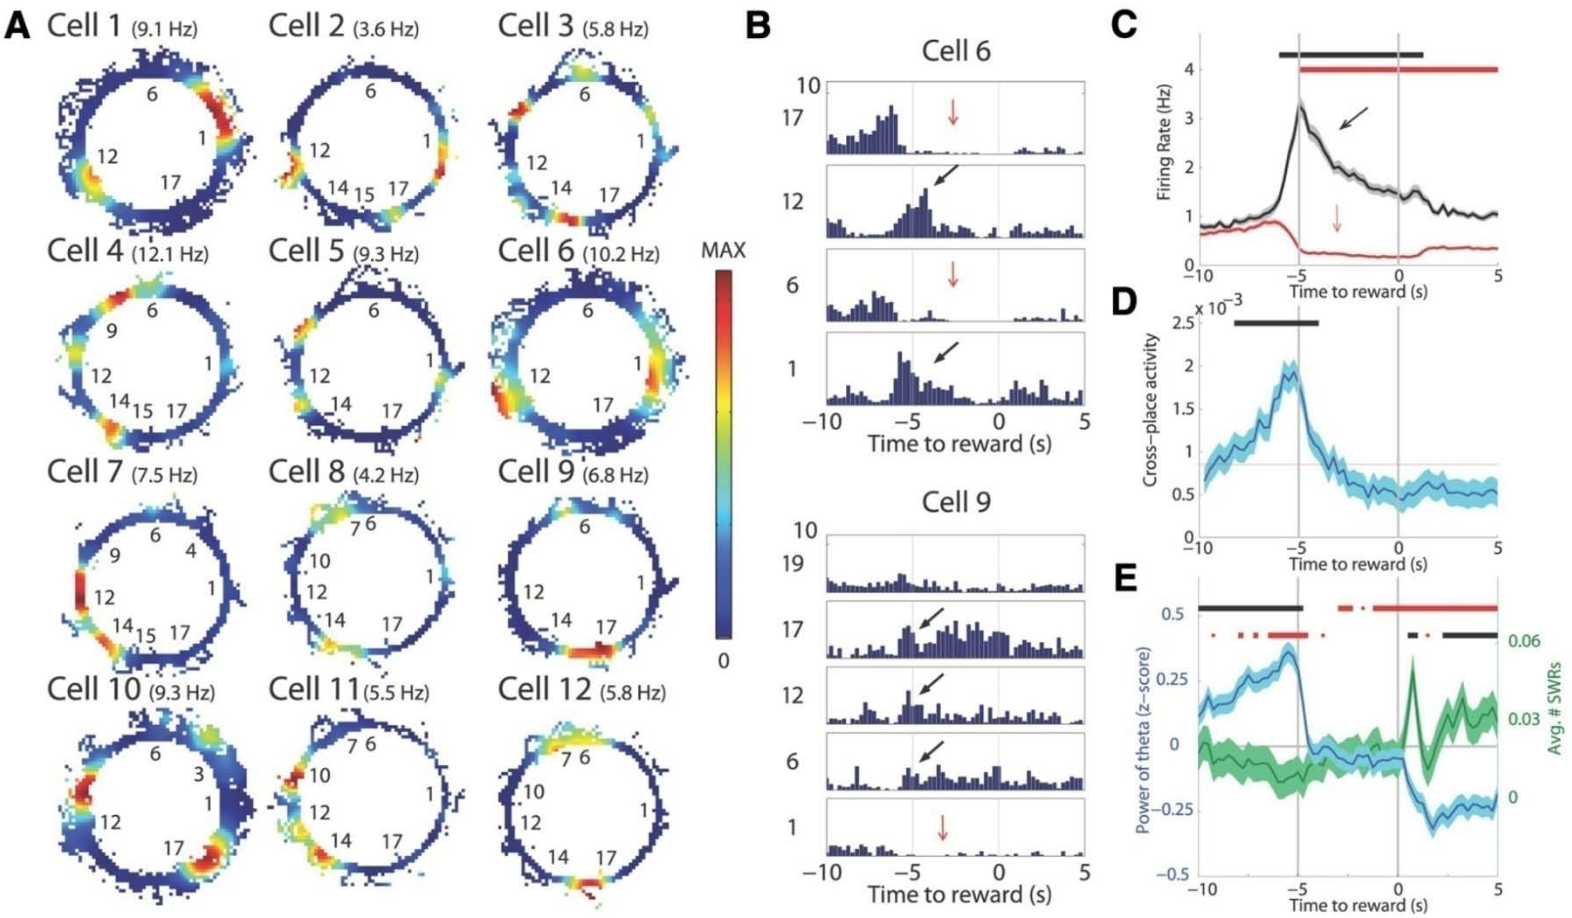
A
Cell 1 (9.1 Hz)
Cell 2 (3.6 Hz)
Cell 3 (5.8 Hz)
Cell 4 (12.1 Hz)
Cell 5 (9.3 Hz)
Cell 6 (10.2 Hz)
Cell 7 (7.5 Hz)
Cell 8 (4.2 Hz)
Cell 9 (6.8 Hz)
Cell 10 (9.3 Hz)
Cell 11 (5.5 Hz)
Cell 12 (5.8 Hz)
1
3
4
6
7
9
10
12
14
15
17
MAX
0
B
Cell 6
10
17
12
6
1
−10
−5
0
5
Time to reward (s)
Cell 9
10
19
17
12
6
1
−10
−5
0
5
Time to reward (s)
C
Firing Rate (Hz)
0
1
2
3
4
−10
−5
0
5
Time to reward (s)
D
x 10^−3
Cross−place activity
0
0.5
1
1.5
2
2.5
−10
−5
0
5
Time to reward (s)
E
Power of theta (z−score)
0.5
0.25
0
−0.25
−0.5
Avg. # SWRs
0.06
0.03
0
−10
−5
0
5
Time to reward (s)

Figure 3 Cells fire at multiple goal sites (A) Example rate maps (both RUN and WAIT events) for 12 cells that fired at multiple locations goal locations (numbered). The color for each rate map is scaled to each cell's maximum firing rate reported in parentheses by the cell label. (B) Histograms from two cells centered on reward delivery at t=0s. Each block is a different reward site. The grey bars indicate the 5s the rat waited for reward. Black arrows show burst of firing upon arrival, red arrows show slight decrease in baseline firing rate as rats wait for reward. (C) Mean cell firing rate during active WAIT event (black) and during non-active WAIT event (red). Black bar is significant firing above baseline ($p < .05$) for the active cells, and red bar is significant firing below baseline for the non-active cells. (D) The cross-place activity was calculated as the product of a cell's normalized firing rate during approach to two goal locations. On average, cells were more likely to be active in multiple locations during the beginning of the WAIT period. (E) The power of theta decreases upon arrival to the goal site and further still after reward consumption. In contrast, the number of sharp-wave ripples increased only after reward consumption, indication that the burst of activity upon arrival to the goal is not due to a ripple event. Black bars show data above baseline, red bars represent data below. The top error bars show significance testing for theta power and the bars underneath correspond to analysis of sharp-wave ripples.

To examine the similarity of activation patterns across locations, I measured cross-place activity (CPA) between all pair-wise combinations of WAIT events at well-learned sites for all cells. This analysis revealed a common activity peak at the beginning of WAIT events, followed by diminished cross-place activity during the course of the ensuing WAIT event (Figure 3D). The CPA increase at the outset of WAIT events was observed even after setting rates below a sliding threshold to zero (see methods) to measure only time points of high activity in both goal locations.

The activation of cells at the goal site coincided with the rat decelerating (Figure 2B), thus possibly confounding a change neural activity due to reduction in running speed with firing that encodes the WAIT event *per se*. The firing rate of hippocampal cells generally is positively correlated with running speed (McNaughton et al., 1983; Geisler et al., 2007), except for the brief bursts of activity during sharp-wave ripples when the animal is typically not moving (Buzsaki et al., 1983, O'Neill et al., 2006). Here I also observed a positive correlation between running speed and population firing rate in the 20 seconds prior to arrival to goal sites (r = 0.24, p = 4.9^{-167}) and a weak but significant positive correlation between acceleration and population rate during the same epoch (r = 0.02, p = 0.01). In contrast, during WAIT epochs population firing rate was negatively correlated with acceleration (r = -0.11, p=5.7^{-15}) though the positive correlation with velocity was maintained (r = 0.25, p = 4.35^{-70}). Most important, even though the velocity and acceleration profiles were very similar at all goal sites, cells were never active at all the goal locations. Also, the probability of observing sharp-wave ripples was lowest upon arrival to the goal sites (Figure 3E). Combined, these observations show that the burst of activity upon arrival to the goals cannot be explained simply by changes in running speed, acceleration, or spiking in a sharp-wave burst, suggesting they likely reflect neural processing related to behavior during the WAIT events.

3.3.3 Hippocampal activity distinguishes well-learned WAIT events at different locations

In addition to relating memories by their common features, schemas must distinguish memories by differences in their features. There are potentially many ways by which hippocampal ensembles could discriminate memories. A simple mechanism for generating a high degree of memory separation would be for cells to activate during a single event and remain silent for others. Another possibility is that cells distinguish events based on firing rate and previous studies have shown that rate differences could allow for comparable discrimination as the simpler, binary coding scheme (Leutgeb et al., 2005). Finally, it is possible that while single cells may discriminate a subset of events

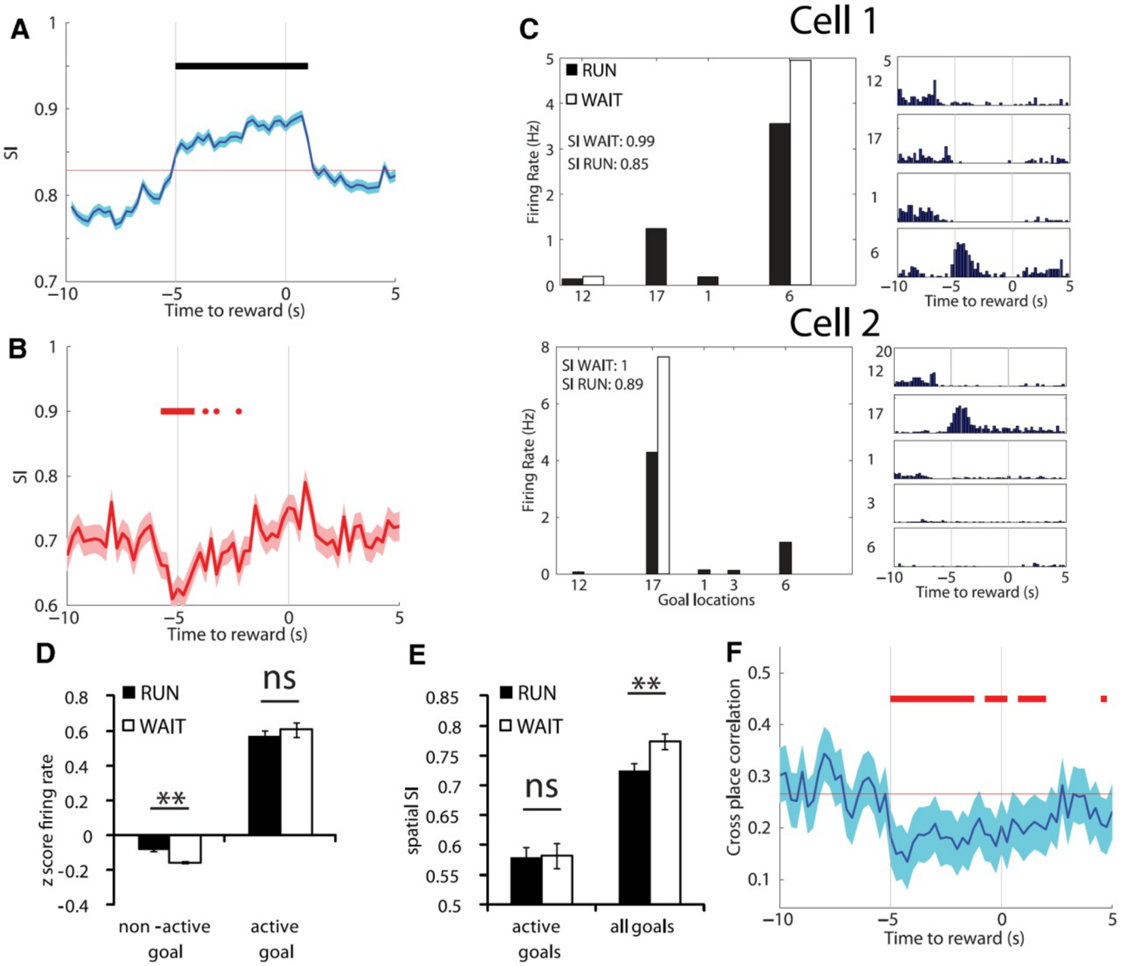

Figure 4 Ensembles differentiate well-learned goals (A) The average of cells' selectivity index (SI) as measured by the PETH recorded at each WAIT location. The SI increased upon arrival to the goal site (t = -5s) and remained high until just after reward delivery. (B) The same selectivity metric but calculated only for cells that were active at multiple sites. The SI for only the sites in which a cell was active decreased upon arrival to the goal sites during the same period of high cross-place activity. (C) The firing rates of two cells at each WAIT site as calculated from the firing rate map (left) and the PETH (right). During WAIT epochs the firing rate is more selective to the preferred goal location. (D) The average z-score firing rate as calculated from the firing rate map is higher during RUN epochs at sites in which the cell does not activate during WAIT epochs, while there is no difference in rate at the sites for which a cell is active during the WAIT epoch. (E) For cells that are active for multiple WAIT epochs, the selectivity across the active sites is equally poor during WAIT and RUN epochs, though across all goal locations, selectivity is higher during WAIT epochs, indicating that cells are selective during WAIT epochs due to decreases in firing rate in locations where they were active during RUNs. (F) The ensemble correlation decreased upon arrival to the goal location indicating more distinct neural representations. Red bars show significant (bootstrap, p < .05) epochs below baseline and black lines shows epochs above baseline.

poorly, on the ensemble level those events are well discriminated due to the contribution of other cells that are selective to that subset. I explored whether any of these three coding scenarios could describe how cells distinguish goal sites from one another.

Though many neurons were active at the outset of multiple WAIT events, these events could be potentially discriminated by firing rate and differing temporal patterns of activity. For example, Cell 6 depicted in Figure 3B fired during WAIT events in two locations with different firing rates at different times and decreased firing rate at the other two locations. To evaluate the extent to which WAIT events at different locations were associated with distinct temporal firing patterns, I calculated the SI for each cell ($n = 410$) during WAIT events at all well learned goal locations. There was a prominent increase in SI that began upon arrival at the goal and persisted throughout the 5s WAIT period, indicating that single neurons highly discriminated those events at different locations (Figure 4A). When these analyses were restricted to cells that were active in multiple locations ($n = 164$), there was a drop in selectivity for the locations in which a cell showed cross-place activity (Figure 4B). This drop in SI can be seen in Figure 4B upon arrival to the goal site, the same time in which cells showed high CPA. While many cells discriminated WAIT events by being active only at a single site, other cells encoded a subset of the WAIT events similarly by increasing firing rate upon the outset of the events.

While the SI analysis can capture the temporal dynamics as rats remain within the 15cm of the goal site, this time series analysis is potentially confounded by differences in occupancy of specific locations; so, firing rates at the goal locations were additionally analyzed based on the average firing rates, as calculated from the spatial rate map, at goal locations during WAIT and RUN events. As can be seen in the examples in Figure 4C, cells showed different firing rates during WAIT and RUN epochs through the same goal locations. In the example shown in Figure 4C, Cell 1 fired in locations 6 and 17 during RUN epochs, but during WAIT epochs only fired in location 6 and was virtually silent in the other locations, as demonstrated by the low firing rates in locations 1, 12 and 17 during the WAIT event in the adjacent histogram. In both examples, the cells fired in more goal locations during RUN events. The difference in firing rate during RUN epochs was not a general increase in drive to the cells since there was no difference in firing rate during WAIT epochs (*mean rate = 2.30 ± 0.13 Hz*) as compared to the RUN epochs (*mean rate = 2.20 ± 0.11 Hz*) through the goal location in which a cell was active (*bootstrap, p = 0.56*) (Figure 4D). In goal locations in which a cell was not active during WAIT events, firing rates increased during RUN events (*mean rate WAIT = 0.30 ± 0.03 Hz; mean rate RUN = 0.49 ± 0.03 Hz; bootstrap, p < 0.0001*) (Figure 4D). It is important to note that cells never fired at all goal locations during WAIT events despite identical behavioral demands. By definition, locations in which a cell was active were associated with significantly higher average firing rates than locations in which a cell was not active (*bootstrap, p < 0.0001*)

As suggested by these rate differences, hippocampal neurons discriminated locations during WAIT events (*spatial SI = 0.83 ± 0.01*) better than during RUN events (*spatial SI = 0.79 ± 0.01*) through the same locations (*bootstrap, p = 0.039*). This difference in spatial discrimination during WAIT and RUN epochs was more pronounced in the cells that were active in multiple locations (*spatial SI WAIT = 0.77 ± 0.01, spatial SI RUN = 0.72 ± 0.01; bootstrap, p =0.009*) (Figure 4E). However, for the same cells, there was no difference in spatial SI when the analysis was restricted to only the subset of goals in which a cell was active (*spatial SI WAIT = 0.58 ± 0.02, spatial SI RUN = 0.58 ± 0.02; bootstrap, p =0 .93*) (Figure 4E). Therefore, single cells better discriminated goal locations during WAIT events by limiting the number of events for which they were active and not by increasing rate differences in the subset of active locations.

Given that many cells fire at multiple locations, I tested the extent to which the pattern of population activity distinguished WAIT events at different locations. As rats approached the reward site and waited for the water reward, the pair-wise correlation of ensemble activity of all putative pyramidal cells recorded at well-learned goal locations decreased upon arrival and remained low throughout the 5s WAIT epochs (Figure 4F). Therefore, as the rat arrived at different goal locations, there was a rapid change in hippocampal activity that resulted in the associated neural ensembles becoming less similar. Spatial ensemble correlation also indicated that goal locations were less correlated during WAIT epochs (*spatial ensemble correlation = 0.176 ± 0.02*) than RUN events (*spatial ensemble correlation = 0.22 ± 0.02*) through the same location (*bootstrap, p = 0.044*). These findings show that while a subset of cells showed common activity in a subset of WAIT locations, the ensemble of neurons discriminated WAIT events throughout the entire 5s WAIT period.

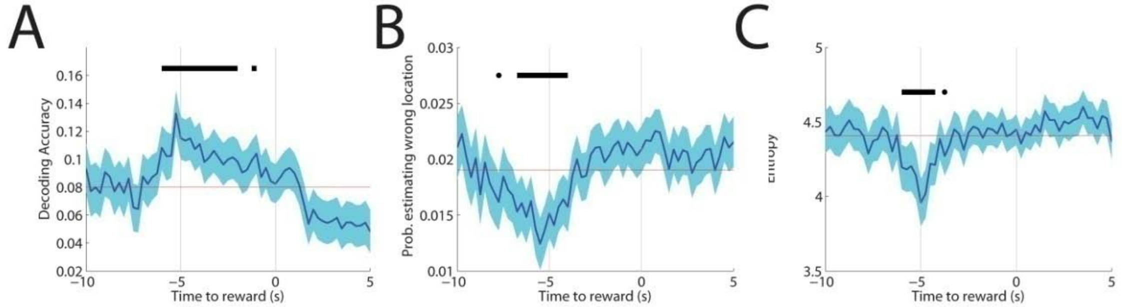

Figure 5 Decoding accuracy is maximal at goal sites (A) Bayesian decoding of hippocampal spiking data revealed the highest probability of the rat occupying its actual location upon arrival to the WAIT site. (B) Also during arrival, the probability of decoding the rats' position to the wrong location and (C) the overall uncertainty of the rats' location (Shannon's entropy) were less than baseline (red line). Black bar indicated significant ($p < .05$) deviations from baseline.

The analyses above relied upon averaging firing rates over space or over time and only considered how distinct goal locations were from one another rather than being globally distinguished from other maze locations. To estimate the overall discrimination of WAIT events at each moment in time, I examined Bayesian decoding of neural activity for the locations of different WAIT events. The best decoding accuracy was observed at the beginning of the WAIT epochs when there was the highest probability of the rat occupying its actual location based on ensemble activity (Figure 5A). Consistent with my results on high spatial selectivity and low ensemble correlation during WAIT epochs, the beginning of WAIT events was also associated with low probability of the rat occupying the other goal locations spatially separate from the rat's actual location (Figure 5B). In addition, the overall uncertainty of the rats' location (Shannon's entropy) was lowest during the beginning of the WAIT epochs (Figure 5C), showing that the rat was least likely to be in any other location other than its actual location during WAIT event. Therefore, despite the subset of cells that fired in multiple goal locations, on the ensemble level WAIT events were well discriminated from one another and from locations elsewhere on the track.

3.3.4 Hippocampal coding of new and previously-learned WAIT events changes over days

I also explored characteristics of firing patterns prior to new learning that predict the assimilation of representations of new WAIT events with learning. Based on the property of schemas by which new learning is accelerated for material that is most consistent with existing schema structure (Bartlett 1932; Tse et al., 2007), I hypothesized that cells that begin to fire during WAIT events with learning would be the same cells that already fired during WAIT events prior to new learning.

To examine this prediction, I introduced a new, spatially-defined goal location mid-way into a recording session to determine how cells change firing at the new and old goal locations with learning. I distinguished cells (n = 44/192) that were activated (see methods) during the WAIT event at the new goal site from those that fell below this activation threshold. These cells showed clear activity at the new WAIT location, but also at previously learned WAIT locations (e.g. Figure 6, Cell 1 in locations 1,3, and 6). The average response during the WAIT event of all cells that were active at the new site was a strong initial burst of activity upon arrival to the site followed by a decrease in average firing rate as cells that were active stop firing (Figure 7A). Also, cells that were not active during new WAIT events showed a slight decrease in firing rate during the

WAIT event as can be seen by the lower than baseline firing rate starting 5s before reward delivery (Figure 7A).

Before the new goal site was introduced, 91.3% of cells that would become active at the new goal site were active at the previously learned WAIT locations. These active cells also fired at more of the previously learned WAIT locations (*2.26 ± 0.22 sites)* than those cells that were not active at the new site (*1.39 ± 0.07 sites; bootstrap,* $p < 0.0001$)(Figure 7B, inset). CPA analysis showed that this activity common to multiple locations occurred as rats arrived to the goal locations (Figure 7B). These results suggest that the networks that encoded the new goal locations were those that were active during other WAIT events.

After introduction of the new WAIT site, cells that were active at the new WAIT site also showed activity at previously-learned WAIT sites (*mean number of sites = 2.29 ± 0.20)* as reflected in a high cross-place activity between the new goal location and the other goal locations at arrival (Figure 7C). To determine whether the high CPA between new and original locations was associated with similar ensemble activity at new and original WAIT events, spatial ensemble correlations (10 ensembles with 16.0 ± 1.25 cells) were calculated in a pair-wise fashion between the new WAIT site and original WAIT sites before and after learning (52 pairs total). These analyses showed that before learning, activity during RUNs through the to-be-rewarded site was poorly correlated ($r = 0.10 \pm 0.04$) with activity during RUNs through the other goal sites (Figure 7D). However, after the introduction of the new WAIT site, the ensemble activity at the new

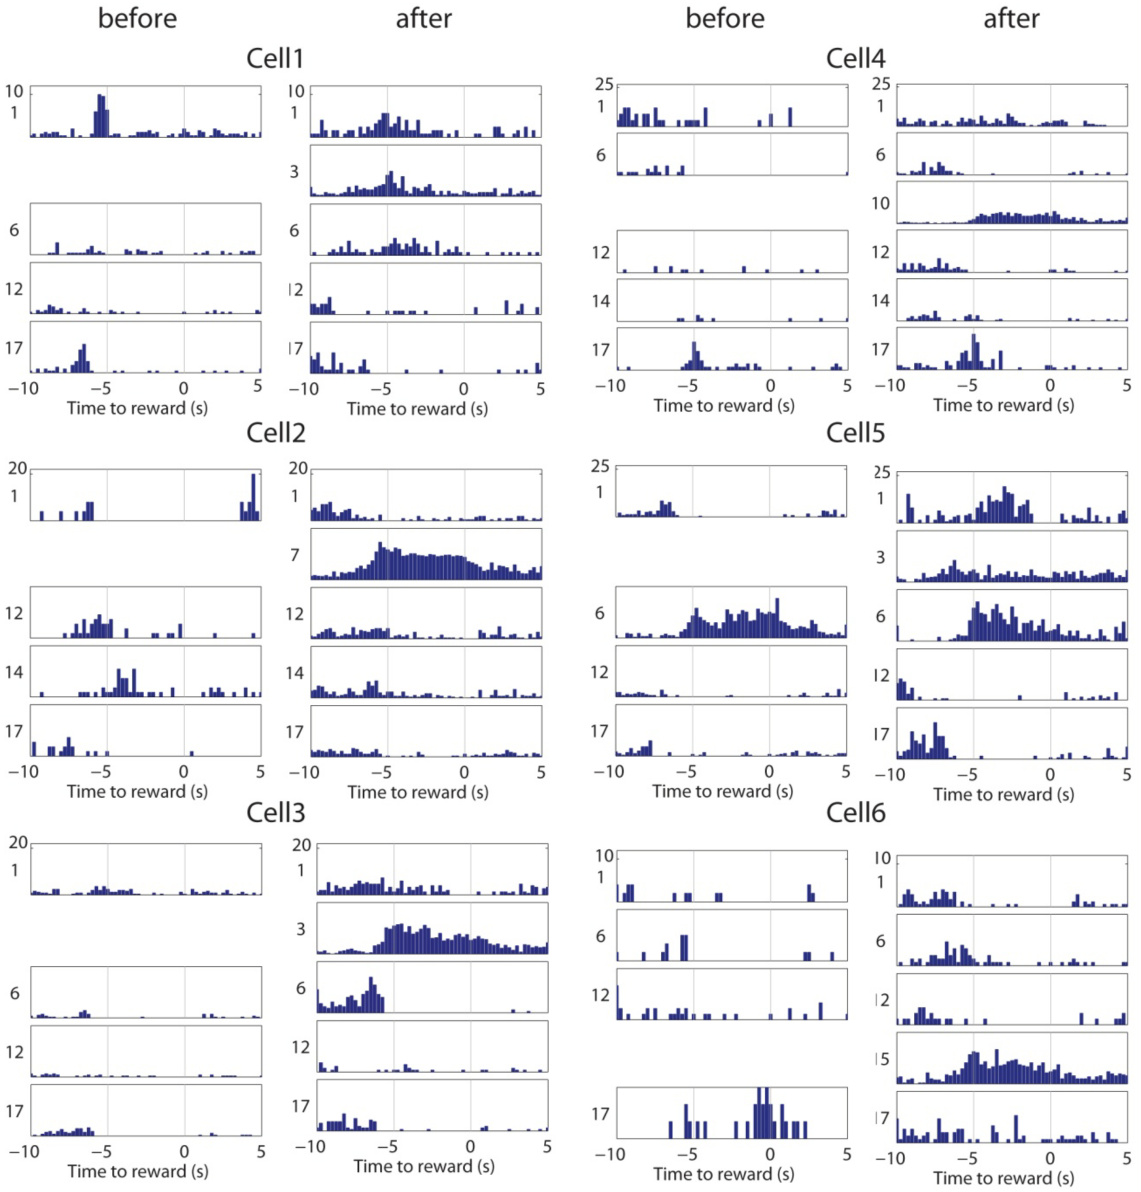

Figure 6 Six cells that fired during the new WAIT events. Cells that fired at the new WAIT locations also fired at original WAIT locations before learning (e.g. cell 2 in location 14) and/or after learning (e.g. cell 1 in locations 1 and 6). Cells also increased firing rate at goal locations where they were not active prior to the new WAIT event (e.g cell 5 begins to fire in location 1).

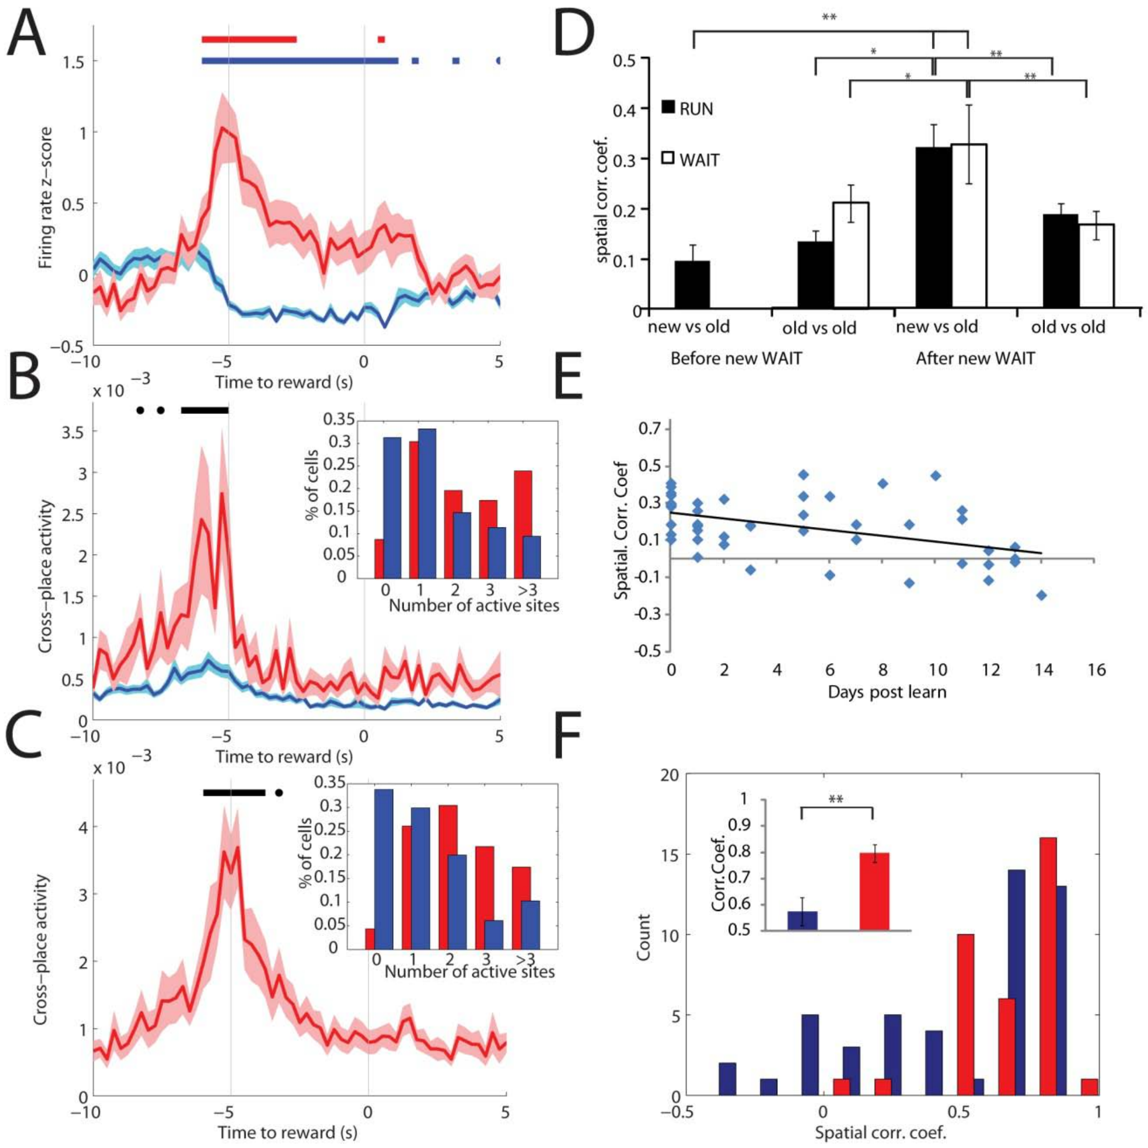

Figure 7 New WAIT events are represented by cells that fired at old WAIT events. (A) The average firing rate at the new WAIT site for cells that were active and those that were not (blue). Red bar shows epochs of firing rate significantly above baseline for the active cells, the blue bar is epochs below baseline for non-active cells. (B) Cells that were active at the new goal site (red) fired upon arrival to multiple WAIT events prior to the introduction of the new goal site, as can be seen by the high CPA. Black bar shows epochs were the CPA prior to learning is above baseline and higher for cells active at the new goal site as compared to the CPA for cells not active at the new goal site (blue). *Inset* Distribution of the number of sites in which a cells (red=active at new site, blue=not active at new site) were active before the introduction of the new site. (C) Cells that fired at the new WAIT location fired at original WAIT locations after learning as seen by a high CPA between original and new sites. *Inset* Distribution of the number of sites in which a cells (red=active at new site, blue=not active at new site) were active after the introduction of the new site. (D) Before introduction of the new WAIT site, ensemble activity at the goal sites was poorly correlated (new vs

old). After the new WAIT site was introduced, ensemble activity at the new site became more correlated with activity at the previously learned goal sites (new vs old). In contrast, ensemble correlation at the original goal locations remained low throughout learning (old vs old). (E) In the days following the introduction of a new WAIT site, the average ensemble correlation between that location and the well learned LED-cued location decreased, indicating increased distinctiveness of the WAIT event representation over time. (F) The distribution of ensemble correlations of the over-trained LED-cued sites for the 1st and 2nd half of a recording session on learning days (blue) and days with only LED locations (red). The distributions significantly differed and the representation of the well learned locations changed more on learning days as shown by lower average ensemble correlations (*inset*).

site became more correlated with activity at the previously learned goals during both RUN ($r = 0.32 \pm 0.04, p < 0.0001$) and WAIT ($r = 0.33 \pm 0.08, p = 0.003$) events (Figure 7D). This high correlation between the new WAIT site and previously learned WAIT sites was in contrast to the low ensemble correlations restricted ($r = 0.14 \pm 0.04$) to the previously learned WAIT events (*bootstrap*, $p = 0.02$) (Figure 7D).

To address at what point activity at the new WAIT site became decorrelated with activity at other WAIT sites, I tracked activity during WAIT events in the days following learning. Activity at the new WAIT site was only compared to the over-trained LED-cued sites in order to see how new information was incorporated into the well-established schema. On the first day of learning, the ensemble activity at the new WAIT site showed the highest correlation with activity at the LED-cued sites (Figure 7E). Over the course of days, the spatial ensemble correlation between the new WAIT site and the over-trained LED-cued sites decreased, such that there was a significant correlation between time since learning and the average spatial ensemble correlation between new and original WAIT sites ($r = -0.45, p = 0.009$).

High average ensemble correlations between new and original WAIT sites were thought to be due to single cells that fired at a similar rate at the new WAIT and at one or more of the original WAIT sites. As suggested by the change in ensemble correlation, the smallest rate difference between the new and original sites was observed on the day of learning (rate = 1.73±.49 Hz) and increased over the course of days giving rise to a significant correlation between minimum rate difference and days post-learning (see methods) ($r = 0.42, p = 0.02$).

The pair of locations that produced the minimum rate difference was not necessarily the closest two locations in actual space as there was no correlation between the distance between the goal sites and minimum rate difference ($r = 0.12, p = 0.53$). Furthermore, there was no correlation in the total amount of distance traveled during the WAIT epoch and the number of days since learning ($r = -0.14, p = 0.36$), suggesting that there was no systematic change in behavior at the new WAIT site in the days after learning.

These data suggest that integrating the new goal location into the pre-existing schema involved cells initially generalizing the original WAIT events to the new WAIT event, but over extended time, the full set of goal locations became well distinguished by hippocampal cell activity.

3.3.5 Learning induces modification of pre-existing schemas to accommodate new memories

After characterizing the dynamics of the schema for previously acquired goal locations, I next analyzed how an established schema changes to assimilate new WAIT events. I observed that learning a new goal location caused a remapping of spatial firing patterns as shown by a lower correlation of the spatial firing rate maps before and after learning ($r = 0.44$ ± 0.02) than the correlation of rate maps of the first and second half of a session without new learning ($r = 0.53$ ± 0.02*; bootstrap,* $p = 0.002$). Rate map correlations did not differ between the learning sessions for the first ($r = 0.45$ ± 0.04), second ($r = 0.45$ ± 0.04) and third goal locations ($r = 0.43$ ± 0.04)(*all* $p > 0.19$) and the rate map correlations for each of these three types of learning session were lower than the rate map correlations on days without learning (*bootstrap, all* $p < 0.05$). These results indicate that each new learning event caused an equivalent degree of reorganization in hippocampal networks.

More specific to theories of schema modification, it is expected that the assimilation of new information causes changes in the pre-existing schema to accommodate new memories, and this should be reflected in substantial changes in firing rates during the WAIT events. To calculate how the representation of a goal site changed over the course of a session, the spatial ensemble correlation was calculated between the 1st and 2nd half of the recording sessions. Low correlations indicate a change in the representation of the WAIT sites. On days with only LED-cued locations, the majority of WAIT events had a stable representation as indicated by the high average spatial correlation between the first and second half of the session (Figure 7F). In contrast, on days with learning, some sites maintained a stable representation throughout learning (e.g. in Figure 6, Cell 5 in location 6), while other locations were associated with cells that changed their firing rates as the new goal location was learned (e.g. in Figure 6, Cell

5 in location 1; Cell 1 in location 6). The learning of the new goal site was accommodated by changes in firing rate in at least a subset of the original goal locations as seen by a lower average spatial correlation between the first and second half of sessions with learning ($r = 0.57 \pm 0.05$) versus days without ($r = 0.79 \pm 0.03$; *bootstrap*, $p = 0.002$) and a significantly different distribution of correlations for days with and without learning (*Kolmogorov-Smirnov statistic* $= 0.38$, $p = 0.004$) (Figure 7F).

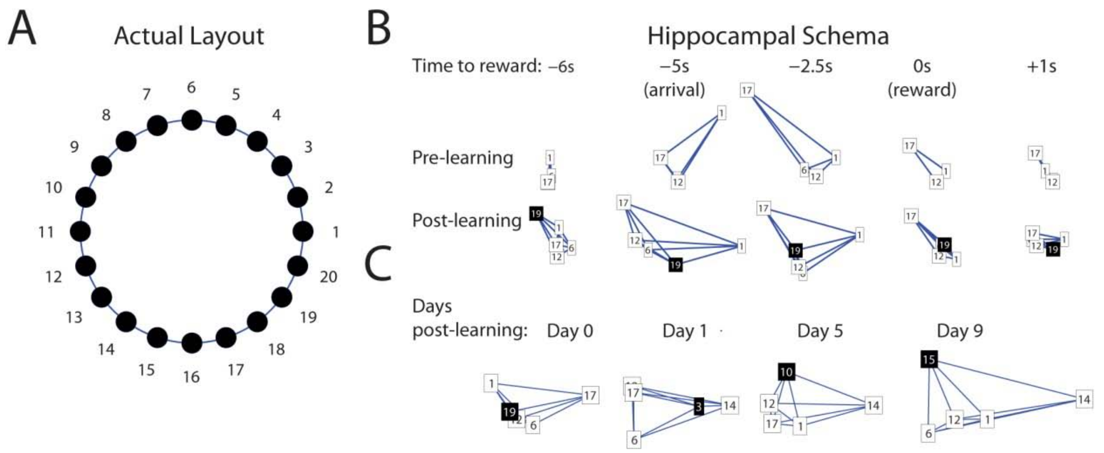

Figure 8 Schemas are dynamic at multiple time scale. (A) The circle track consisted of 20 possible water ports 22cm apart of which 4-8 ports could be rewarded during a single session. (B) The number boxes are the identity of a rewarded water ports and the new goal site is depicted in black with white text. The ensemble representations of the WAIT events are plotted by the first two principal components at each time relative to reward. The top row is the pre-learning schema and the bottom row is the post-learning schema with the addition of the new goal location plotted in the same PC space. At arrival (-5s to reward), the representations of the goal locations became less correlated due to the fact that the ensemble representation of the goal locations occupy different regions of hippocampal state space. In the post-learning schema, pre-existing locations are represented differently by the same ensemble, but original WAIT events are still well discriminated by the hippocampal network. (C) Immediately after learning, the new WAIT site is represented similarly to pre-existing WAIT sites, as visualized by the new and original WAIT sites occupying a similar region of PC space. Over the course of days the new WAIT site is represented differently from the original WAIT sites as seen by an increase in the distance between the new WAIT site to the old. The data are plotted from different sessions and different rats to clearly depict the trend described in Figure 7E.

3.3.6 Schemas are dynamic during task performance and across days

I examined the structure of neural ensemble representations of the WAIT events at different locations (Figure 8A) in terms of projections onto the first two principal components of the ensemble (see methods) (Figure 8B, C). This analysis revealed that hippocampal schemas of WAIT events were dynamic in two ways. First, during stable performance, as the rat traversed the maze and arrived at the goal there was a rapid distortion of the schema (Figure 8B). Upon arrival and as the rat awaited reward, the WAIT events were represented in different regions in hippocampal state space, thus increasing cells' spatial selectivity and decreasing ensemble correlation between locations. Thereafter the representations of WAIT events converged.

Second, upon learning a new WAIT location, the existing schema accommodated during learning to assimilate the new WAIT events. The WAIT events in the updated schema were represented in a different part of hippocampal state space, however, as before the pre-existing goals were well distinguished (Figure 8B). The new WAIT events in the updated schema were initially represented similarly to the pre-existing WAIT events. Over the course of days, the new WAIT location became well distinguished from the old, as visualized by an increase in the distance between the new and original goal locations in the principal component space (Figure 8C).

3.4 Discussion

I observed that incorporating a new goal location into a pre-existing schema of goal locations changed the hippocampal network in several ways. First, cells that fired at original goals began to fire at the new goal site in addition to firing at the previously learned goal sites. The common activity across the new and original goal locations diminished the discriminability of the new goal location by hippocampal cell firing. Only in the days after learning did the representation of the new goal location become decorrelated from those recorded at the original goal locations. I also observed that firing at the original goal sites changed as the new goal location was learned. Despite the plasticity that occurred as new goal information was accommodated, the ensemble continued to discriminate the well learned goals during reward anticipation.

Several previous studies have reported that changing a goal location does not influence place cell firing (Speakman and O'Keefe, 1990; Trullier et al., 1999) while others report shifts in the mapping of goal locations (Gothard et al 1996b; Kobayashi et al., 1997) and more recent studies have shown that there may be an over-representation of goal areas in a water maze task (Hollup et al., 2001) and in paradigm that require association of reward and location (Dupret et al., 2010). However, in all of these studies, goal locations were moved, rather than added as in the current study, and were aimed at examining the influence of the goal event on spatial mapping. By exploring the addition of goal locations to a pre-existing and continuing representation of goal events I am able to observe the evolution of a schema for the same behavior at multiple goal sites within an environment.

My findings that cells change their firing patterns at both original and new sites parallel several studies that have focused on the spatial firing patterns of hippocampal place cells as animals explore multiple environments with overlapping features. These studies have shown that alterations in salient spatial or non-spatial cues often result in a “partial remapping” of the spatial firing patterns of hippocampal neurons, reflected in a combination of subsets of place cells that maintain their spatial firing patterns and other cells that cease firing or develop a new spatial firing pattern. Partial remapping has been observed when spatial cues are eliminated or reorganized (Shapiro et al., 1997), when rats move between environments with identical local spatial cues (Skaggs and McNaughton, 1998), and in the same environment when task demands are changed (Markus et al., 1995). Partial remapping could both distinguish experiences that share overlapping features and bind related experiences via coding elements that are common across representations (Eichenbaum, 1999).

Computational models (see Chapter 1.3) highlight the deleterious effects of learning on retrograde interference which would manifest as remapping. Indeed, remapping at old sites occurred which suggests that new information was integrated within the existing related network. However, despite this remapping, critical features of the code for the goal locations was preserved, such as their discriminability, suggesting that the interference was not ‘catastrophic’. The learning-induced partial remapping at the over-trained goal locations may be essential in maintaining the structure of original and new location memories and blocking this plasticity could result in the reconsolidation deficits observed in many studies (see Chapter 2)(Debiec et al., 2002; Morris et al., 2006).

Other studies that have also included common stimuli or behavioral events in different places, have observed hippocampal neurons that fire in response to the common events. Wood et al. (1999) reported similar firing patterns of hippocampal neurons when rats sampled the same odors in different places while performing a delayed non-match to sample task. Similarly, Singer et al. (2010) reported “path equivalence” of spatial firing patterns of hippocampal neurons as rats traversed separate but parallel segments of routes through a maze. Notably, in these experiments, and in several others where rats traversed a common path segment while pursuing different trajectories (Wood et al., 2000; Frank et al., 2000; Ferbinteanu and Shapiro, 2003; Smith and Mizumori, 2006; Bahar et al., 2012), some neurons have similar activity associated with the common events, whereas others fire distinctly, thus disambiguating overlapping memories. In monkeys (Hampson et al., 2004) and humans (Quiroga et al., 2005) neurons have been observed that fire similarly in response to stimuli that are visually quite various but are similar in meaning (e.g., a cell that fires to various views of the same famous person, and even the name of that person). Such “categorical” (Lin et al., 2005; Hampson et al;, 2004) and “invariant” (Quiroga et al., 2005) firing patterns are consistent with the findings on rodents indicating that hippocampal neurons can link between related memories via common coding elements (Eichenbaum 2004).

In the well-learned schema, the common coding of reward sites was not at the expense of pattern separation, though this was not the case for newly learned goal locations. The cells that fired in the new goal learning were those that fired at other goal sites suggesting the incorporation of new cells into related networks. The stronger correlation between the new and original goal locations may share a common mechanism with the over-representation of the escape platform in a water maze paradigm (Hollup 2001; Fyhn et al., 2002). In those studies, when the goal was moved, cells that fired at the original goal locations tended to fire in the new goal location, though only during the early phase of learning. Furthermore, just as I have reported, several of the cells had multiple fields on the annular water maze. However, since the goal location was moved – and not added – in the water maze experiments, it is unclear whether cells in those studies came to relate the new and old goal locations through common firing.

I observed that while cells fire at multiple locations during learning, the disambiguation of those locations takes days, potentially involving processes of consolidation. These results are reminiscent earlier studies that have shown that multiple exposures can be necessary for CA1 cells to distinguish arenas that differ only by the color of a cue card (Bostock et al., 1991) or how rotations of proximal and distal cues affect place cell firing (Shapiro et al., 1997; Brown and Skaggs et al., 2002). When a rat must choose between one of several well-learned trajectories through the same space, the different trajectories are encoded by anti-correlated representations in CA1 (Bahar et al., 2011), suggesting that experience may increase the degree to which one representation suppresses competitors.

Common firing across locations was especially prominent during running through goals as compared to waiting in the same locations. In mice (Kentros et al., 2004) and rats (Fenton et al., 2010), goal directed behavior increases the spatial information (spatial discrimination) of hippocampal activity, and it has been suggested that decreases in spatial information during foraging may be due to multiple representations that are transiently expressed (Harris et al., 2003; Jackson and Redish, 2007, Kelemen and Fenton, 2010, Jezek et al., 2011; Dupret et al., 2013). Hok et al. (2007b) observed that a large proportion of the cells are active when rats run to an unmarked goal zone in an open arena, and proposed that the activity reflected a transient reward prediction signal (Hok 2007a,b). I also observed cell activity upon arrival to goal sites, though this activity clearly signaled more than the possibility of reward since decoding of this activity led to the best estimate of a rat’s location. It is unknown whether cell assemblies are composed of cells that fire in response to multiple related events or whether the separate goals are encoded by cells that can be part of distinct cell assemblies that are not simultaneously active. In either scenario, cell activity at multiple reward sites during running, and potentially upon arrival to the goals, may reflect a decision making process that requires transiently relating spatially separate events.

I hypothesize that learning in a familiar environment adds new hippocampal neurons to a pre-existing neural network to support the memory of the relationship to the new memory with the old. The addition of new information to an existing schema causes reorganization of related networks, in my case those that represent events at other goal locations. The reorganization of original memories to accommodate new information establishes new meaningful relationships and also maintains the pre-existing similarities and differences of the original network. I propose that consolidation of new memories within pre-existing networks is the reconsolidation of those existing networks, a process that ultimately stores new and original relationships within a common schema (McClelland et al., 1995; Tse et al., 2007; McKenzie and Eichenbaum, 2011).

4 Hippocampal representation of related and opposing memories develop within distinct, hierarchically organized neural schemas

4.1 Introduction

Recording studies investigating hippocampal activity during memory formation have found that cells fire similarly in locations that involve the same behaviors (Singer et al., 2010; McKenzie et al., 2013). Conversely, when task demands require different behaviors in the same location, firing patterns of place cells can become anti-correlated (Markus et al., 1995; Bahar et al., 2011). These findings suggest that related memories within a spatial schema can be supported by common firing patterns by place cells, and that behavioral context can drive the generation of distinct schemas.

To explore how hippocampal neural networks represent distinct experiences in the same and different behavioral and spatial contexts, I recorded from hippocampal neuronal ensembles while rats learned to use the spatial context they occupy to determine which of two objects to select in order to obtain rewards. As rats learned about an initial set of stimuli (Context 1: choose stimulus A not B; Context 2: choose B not A), hippocampal neurons came to fire preferentially during stimulus sampling at one or more locations (Komorowski et al., 2009 ; Tort et al., 2011). Furthermore, when rats learned a second stimulus set (Context 1: choose C not D; Context 2: choose D not C), hippocampal ensembles represented a hierarchy of stimulus relationships during stimulus sampling. This representational schema embodied relevant relations among the two problems via their overlapping spatial and non-spatial features.

4.2 Methods

4.2.1 Surgery

Rats were anesthetized using a mixture of 3% isoflurane in oxygen and were injected with atropine (0.03 mg/kg, S.C.) and buprenorphine (0.1 mg/kg, S.C.) before surgery and meloxicam (1 mg/kg, S.C.) immediately after surgery. An ophthalmic ointment was applied to the eyes, and temperature was maintained between 36°C and 38°C. Ringer's solution was injected to maintain hydration. Stainless steel bone screws, including two ground screws, were inserted into the skull over the cerebellum. A 2.0-mm-diameter hole was drilled into the skull using a dental drill, and the dura was removed. The base of the microdrive was aimed at the CA1 and CA3 region of the right dorsal hippocampus (−3.6 mm anteroposterior, +2.9 mm mediolateral). The craniotomy was sealed with Kwik Sil (World Precision Instruments), the microdrive and ground wires were secured in place using dental acrylic (Henry Schein), and the surgical site was sutured shut.

4.2.2 Electrophysiological recordings

Individually movable 24-tetrode microdrives were built in-house. Tetrodes were spun from four nichrome wires (12.5 μm diameter; California Fine Wire, Grover Beach, CA, USA) and gold plated for an impedance of 200 kΩ at 1 kHz. At the end of surgery, each tetrode was lowered ~850 μm into the brain. After 5 d of recovery, the tetrodes were lowered further over 7–14 d toward the CA1 layer and CA3 layers. After the experiments, 25 µA of current was passed through each tetrode for 30 s before perfusion and histological confirmation of tetrode placement.

As reported previously (McKenzie et al., 2013) the electrical signal was referenced to a common skull screw and differentially filtered for single-unit activity (154 Hz to 8.8 kHz) and local field potentials (LFPs) (1.5–400 Hz). The amplified spikes from each wire were digitized at 40 kHz, whereas the field potential was digitized at 1 kHz and monitored with the Multineuron Acquisition Processor (Plexon). Individual pyramidal neurons were isolated using Offline Sorter (Plexon) by visualizing combinations of waveform features (peak valley, valley, peak, principal components, nonlinear energy and timestamps) extracted from wires making up a single tetrode (i.e., “manual cluster cutting”). Single-neuron selectivity was verified by the interspike interval histograms that contained no successive spikes within a 2 ms refractory period.

4.2.3 Apparatus.

The apparatus consisted of two wooden chambers connected by an alleyway. Each chamber was 40.6 x 40.6 x 43cm. The central alley connecting the chambers was 40.6cm long. Scented terra cotta pots could be positioned in two locations 30cm apart within each context at the corners of the context on the opposite side of the central alleyway. Four contexts were used in the study, two in each chamber, that were textured and colored inserts on the walls and floor. For shaping and pre-training, Context 0 was Clear Plastic taped to the wood chamber and Context 00 was Green Foam Strips taped to the wood. During the training, two new contexts were used in the same chambers: Context 1 was black anti-static cloth and Context 2 was white sand paper.

Items were terracotta pots (diameter:10.2cm, height:10.2cm) weighted with sand, capped with melted wax, filled with different digging media (3cm deep), and scented with essential oils mixed with Crisco around the rim. Scent concentration was calibrated for human just-noticeable-detection. Two items were used for pre-training: Item X (medium: purple plastic beads, scent: banana) and Item Y (medium: shredded paper, scent: hazelnut). Four different items were used for training: Item A (scent: geranium, medium: multicolored triangle foam pieces), Item B (scent: maple, medium: gravel), Item C (scent: raspberry, medium: 1cm cut red coffee stirrers), and Item D (scent: cinnamon, medium: blue rubber mulch).

4.2.4 Behavioral Training.

<u>4.2.4.1 Habituation and Shaping</u>: Five male long-Evans rats were maintained at 85% of their free feed body weight throughout training and testing. Rats were handled for approximately one week for 20 minutes daily during which point they were given Froot Loops which would later serve as motivation learning.

On the first two habituation days, rats freely explored the apparatus with Context 0 and 00 inserts for 20 minutes with Froot Loops placed randomly throughout the environment. On the third day, Froot Loops were only placed in the corners of the two contexts where the pots would later be present (Positions 1 - 4). On the fourth day, Froot

Loops were placed inside of a pot that was filled 2/3 with wax and no other digging media. The single pot could appear in any of the four locations. On subsequent days, 1/4 Froot Loop pieces were used as reward buried within the pot that was to be associated with the context during pre-training (Item X in Context 0, Item Y in Context 00).

4.2.4.2 Pre-training. Once rats reliably dug in the pots in either context, they were introduced to the full version task. At the start of a trial, the rat was blocked into one chamber and allowed to explore for 10 seconds. A divider was then put down and the two pre-training items (Item X and Item Y) were placed in the corners of the chamber. The divider was lifted and the rat was free to dig in either pot. In the Context 0, Item X contained a buried reward and in Context00, Item Y contained the reward. If the rat dug (snout or paws touching the media) in the non-rewarded pot, no reward was given, both pots were quickly removed, and the trial was over. Rats were permitted to sample each pot multiple times in the absence of digging. The position of the pot (two positions per context) was pseudo-randomized for each trial such that the same item would not occur in the same location on more than two consecutive occasions. Once the rat consumed the Froot Loop, the trial ended, the alley dividers were lifted and the rat shuttled to the other context for the next trial. Every 10 trials, the rat was kept in the same context for two trials in a row to ensure that the rats did not adopt an alternating strategy. Every 10 trials, there was one trial in which no pot contained a reward to ensure behavior was not guided by the Froot Loop odor. Rats were rewarded on these trials after digging in the correct pot. At the beginning of pre-training, rats were given 40 trials/day which increased to 80 trials/day over 1-2 weeks. Training was carried out for approximately 3-4 weeks until the rat reached a criterion of 85% correct within a session. Once rats reached criterion, they were scheduled for micro-drive implant surgery. After at least one week of recovery, rats were retrained on the pre-training items while tetrodes were lowered into position (up to one month of daily training).

4.2.4.3 Training. Once tetrodes were in place, rats were introduced to a novel set of contexts (Context 1 and Context 2) and a new set of items (Item A and Item B). Item A was rewarded in Context 1 and Item B was rewarded in Context 2. Otherwise the task was identical to the pre-training protocol. One 75-100 trials were given per day, though more sampling events could be recorded if rats correctly rejected the non-rewarded pot (or incorrectly rejected the rewarded pot). Rats learned the reward contingencies for Item A and Item B over three days (AB1, AB2, and AB3). The day after AB3, rats were given a new item set (Item C and Item D) to be learned within the same contexts. Item C was rewarded in Context 1, Item D was rewarded in Context 2. Rats were given 75-100 trials per day for three testing days (CD1, CD2, CD3). Rats were permitted to self correct on the first 10 trials of AB1 and CD1, though neural data collected on these trials was not analyzed and only the first dig response was considered for assessment of learning. The day after CD3, rats were given a block of 20 trials with only Items A and B as a reminder followed by a block of 124 trials (ABCD1) in which the rat could either receive the first item set (Item A and Item B) *or* the second item set (Item C and Item D) from one trial to the next. Items belonging to different sets were never presented on the same trial. The order in which item sets were presented was pseudo-randomized with no more than three item set repetitions in a row. The two item sets were presented inter-mixed in the same fashion for the following two days (ABCD2 and ABCD3). Data from AB1 from one rat was excluded due to experimenter error during behavioral training.

I studied how hippocampal cells encoded several task dimensions (*Item, set, valence, position and context*). *Items* (A,B,C, and D) that were presented together were part of the same *set* (AB and CD). Items rewarded in the same context were of equivalent *valence* (AC and BD). During ABCD sessions, any of the four items could appear in any of the four *positions* (1-4) so my analyses focused on how cells modulated firing for different items within the same position and how cells fired for the same item (or related items) in different positions and in different *contexts* (Context 1 and Context 2).

Three of the five rats also completed an open field recording session after each recording session. The arena was 61x71cm and rats foraged for randomly distributed 1/4 Froot Loops for 20-30 minutes to quantify place fields in a geometrically and visually distinct environment.

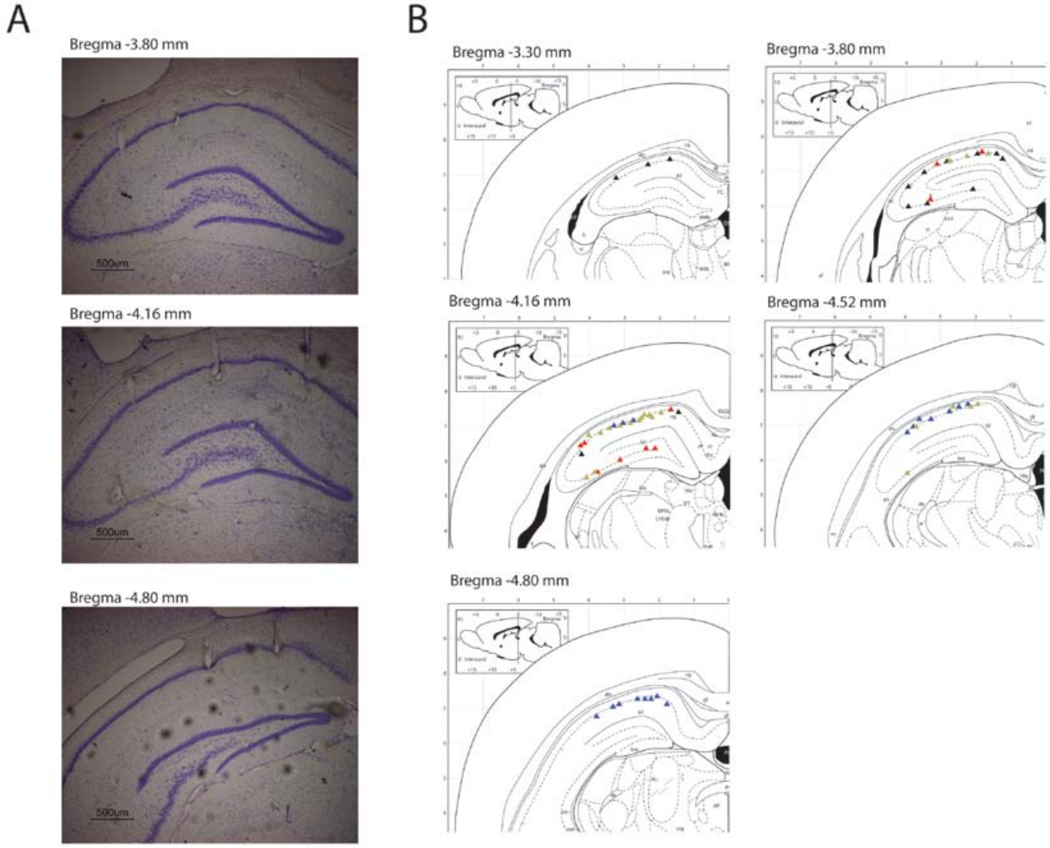

Figure 9 Histological confirmation of tetrode placement. (A) Histology from 3 rats showing tetrode lesion marks in CA1 and CA3. (B) Color coded distribution of lesion marks for 4 rats. Histology from one rat was damaged and unit assignment was done based off of local field potential and turn count.

4.2.5 Analysis

<u>4.2.5.1 Single neuron analysis.</u> I recorded from 571 units from five rats: 292 units from CA1 and another 279 cells from CA3 (see Figure 9). Since there were no obvious differences between cells recorded from CA1 or CA3. I did not find any difference between the regions in my analyses (Figure 14). Instead my analyses focused on ensemble representations sampled in both hippocampal subfields.

Rats could sample the two items presented on each trial multiple times, and all item sampling events in which the rat dug in the rewarded pot and refrained from digging in the non-rewarded pot were considered in my analysis. To estimate the hippocampal representation of every item sampling event, for each cell the number of spikes fired was counted for up to the first 2s of item sampling and this count was divided by the sampling duration to give the average firing rate for each cell on each item sampling event.

4.2.5.2 Population analyses. The goal of the experiment was to determine how the hippocampus codes related events. In this task, events were the item sampling epochs that could be related by the following dimensions: context, position, item, valence or set. When neural activity for events that were within a condition for a particular dimension (e.g. all item sampling events in the same position) was more similar than activity recorded during events that were between conditions (e.g. events in different positions), I interpreted this as evidence for coding of that dimension within the hippocampus. The within and between conditions for each dimension are given in Table 1.

Based on firing rates from these units, two different ensemble metrics were used to measure event similarity: ensemble correlation and a naive Bayesian classifier. These methods are described below.

Task Dimension	Within condition	Between condition
Context	Same context, different position, any item	Different context, different position, any item
Position (Pos)	Same position, same context, any item	Different position, same context, any item
Item in position (Item/Pos)	Same item, same position, same context	Different item, same valence, same position, same context
Valence in position (Val/Pos)	Same valence, difference item, same position, same context	Different valence, different item, same position, same context
Set in position (Set/Pos)	Same set, different item, different valence, same position, same context	Different set, different item, different valence, same position, same context
Item in other position (Item/xPos)	Same item, different position, same context	Different item, same valence, different position, same context
Valence in other position (Val/xPos)	Same valence, difference item, different position, same context	Different valence, different item, different position, same context
Set in other position (Set/xPos)	Same set, different item, different valence, different position, same context	Different set, different item, different valence, different position, same context

Table 1 Definition of the within and between filters for analysis.

4.2.5.2.1 Ensemble Correlation. Firing rates for individual neurons were z-score normalized using the mean and standard deviation among all item sampling events to create a population vector of normalized rates for each event. The Pearson's correlation coefficient of these population vectors was calculated for every pair of events.

To calculate the strength of a coding dimension (e.g. item in position) using the ensemble correlation, I compared the average correlation for events within a condition (e.g. all item sampling events in the same position with Item A) for a task dimension (see table 1) versus events that were between conditions for that dimension (e.g. events in the same position with Item A versus Item B) (Figure 13C, Figure 15A,C). For each coding dimensions, a single d' was generated either for individual sessions or for the experiment as a whole by pooling the correlations recorded in different rats. The d' was calculated as follows:

$$d' = \frac{\mu_W - \mu_B}{\sqrt{\frac{1}{2}(\sigma_W^2 + \sigma_B^2)}}$$

where, μ_W is the mean correlation coefficient for within condition events for that dimension with variance, σ^2_W, and μ_B is the mean correlation coefficient for between condition events with variance, σ^2_B. The observed *d'* was compared to bootstrap data in which I randomly shuffled event identities 10000 times and then recomputed the correlation analysis and d' metric for each bootstrap sample (Figure 15B). When the observed d' was greater than 95% of the 10000 shuffled d' metrics, the dimension captured by the d' was considered to have been significantly coded by the hippocampus (Figure 15D).

To assess when different dimensions emerged during item sampling, I ran a similar analysis though firing rates were taken at different 250ms time bins centered ±3s around pot sampling (Figure 17). Only events for which the rat's head remained over the pot for over 1.5s were included. Significance testing was done using the same bootstrap analysis at each time point and comparing whether the observed d' was significant at $p < 0.002$ (Bonferroni correction for 24 time points).

4.2.5.2.2 Bayesian Classifier. As a separate test of event similarity that had a different set of assumptions, a naive Bayesian classifier was used to determine the probability that a pattern or neural activity was recorded for each item and place combination (four items in four positions). Due to uneven sampling, rats often preferred a particular item and position combination and would sample those more often. Therefore, I only considered the last six sampling events for each item and place combination. When there were fewer than six events, that category of item and position trial was eliminated. Next, the z-scored population vector was calculated for each event, as described above. The dimensionality of the ensemble representation for each event was reduced via principal component analysis (PCA) and only the first four components were used to categorize item/position combinations (Figure 16A,B). Then, the mean, variance and covariance of each item/position four-dimensional ensemble representation was estimated with one event missing from each item and position combination. Next, a multidimensional normal distribution with the estimated means, variance and covariance matrices were fit to the each cluster of item/position ensemble representations (maximum 16). Finally, the probability of the missing events being any of the possible item/position combinations was calculated based on the probability of that item/position combination given the ensemble representation as estimated by the normal distributions above. Since there is a probability associated with each item/position combination, I could determine whether there is hierarchy of coding probability with the correct item and place conjunction generating the highest probability, followed by correct valence, correct set, correct position, and finally correct context (Figure 16C,D).

The degree of dimensional coding was calculated in a similar fashion as that for the correlation coefficients, though the d' was calculated based on the probability of classifying within a condition or between conditions for each dimension. To test significance, I ran a bootstrap analysis in which event identities were scrambled 10000 times and if the observed d' for the difference in probabilities was greater than the 95% confidence interval, I concluded that the hippocampus coded that dimension.

4.2.5.2.3 Dendrogram. The hierarchical nature of the task schema was visualized using the MATLAB functions linkage and dendrogram. On days in which Items A, B, C and D were presented, the median z-score firing rates were calculated for each of the 16 item and position combinations. For each item and position combination, a large firing rate vector was created composed of the rates of every cell recorded from a session in which correct behavior was observed for every item in every position (N=560). The agglomerative hierarchical cluster tree was then created using the unweighted average distance between pairs of vectors and the Pearson's correlation coefficient as the distance metric. The nearest vector pairs were merged to form clusters. Then the nearest clusters were merged iteratively until all 16 vectors were within a single cluster (Figure 13E).

4.2.5.2.4 Cross-day analysis. Due to the low number of cells recorded in each rat, using trial by trial ensemble analyses to compare activity across testing days was not possible. Therefore, to assess the similarity of unit activity for one item set (e.g. AB) versus that for the other item set (e.g. CD), the median z-score firing rates were calculated for each of the 16 item and position combinations from both days. The z-score was calculated using the mean and standard deviation recorded for that day. As described for the dendrogram, each item and position combination (n=16) was associated with a firing rate vector composed of the rates of every cell (N=38). The pair-wise correlating of these vectors across days generated an 8x8 similarity matrix which reflects the overlap in neural activity for each item and place combination. Statistical testing was done by shuffling the item/position identity of each median rate vector and recalculating the correlations on the randomized data set.

4.2.5.3 Interneuron identification. Three criteria of cell firing were considered to determine whether a unit was an interneuron, A) mean firing rate greater than 8 Hz, B) spike length at 25% of spike height less than 0.4ms, and C) time to first moment of autocorrelogram greater than 9ms (Csicsvari et al., 1999). A unit was classified as an interneuron if conditions A and (B or C) were met.

4.2.5.4 Cell stability criteria. To determine whether the neural networks that coded the first item set (AB) were re-engaged to code related items of the second item set (CD), it was necessary to merge data across 24 hours therefore requiring assessment of recording stability. To increase the chances of stable recordings, from two days before the beginning of the experiment until the final recording day, no tetrodes were moved. Thereafter, cell stability was assessed in four ways. First, during manual cluster cutting, waveforms were viewed across time to visually confirm stability in the recording across days. Second, to quantitatively ensure high waveform stability, the average waveforms on each tetrode (four concatenated waveforms) recorded on the first day were correlated

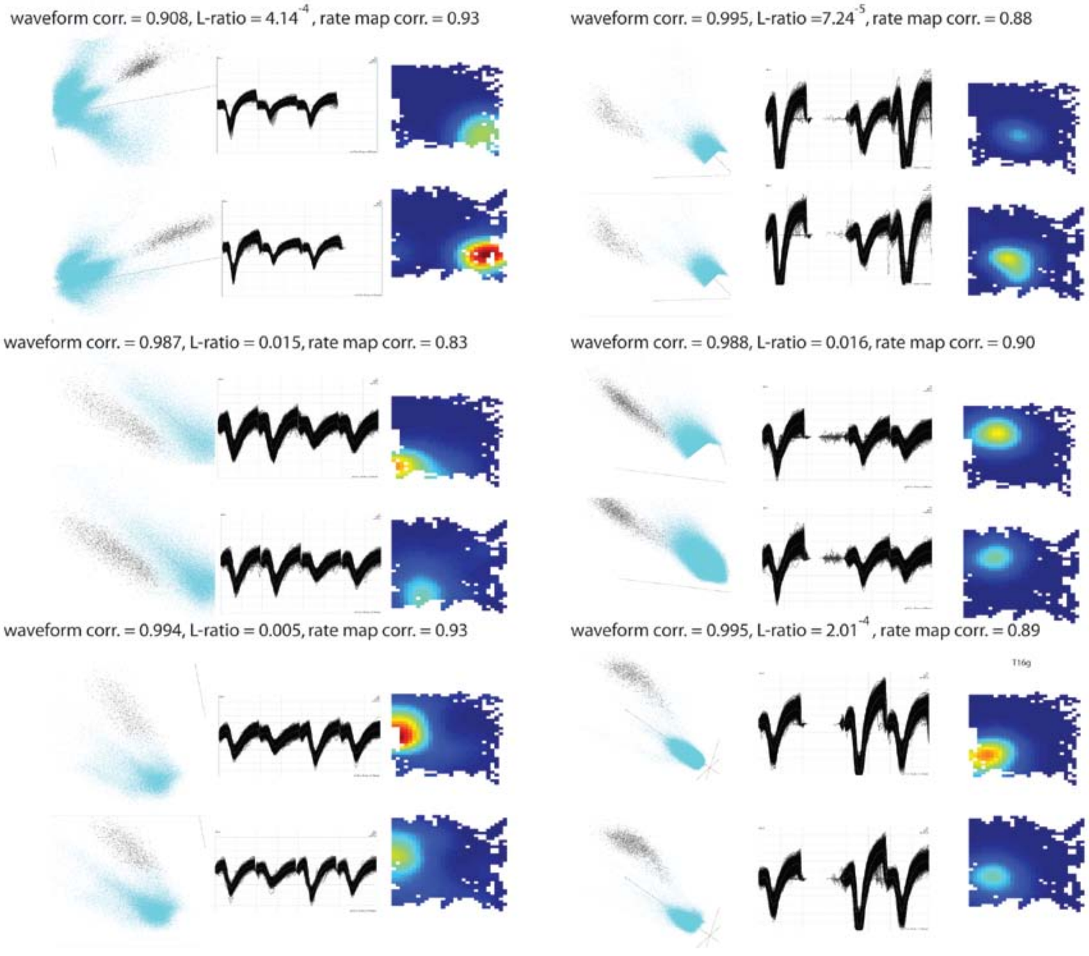

Figure 10 Stability analysis of six cells recorded over 24hrs. For each cell each row is a day. *Left* waveforms plotted by peak-valley for three electrodes, *Middle* Waveforms on each electrode, *Right* Rate map on open field scaled to max firing over both days.

with those recorded on the second and only units with a waveform correlation greater than 0.90 were included. Next, I computed the L-ratio (Schmitzer-Torbert et al., 2005) for the merged recording sessions using the energy of each electrode and the first principal component of the energy normalized waveform from each electrode and only accepted units with an L-ratio less than 0.08. Finally, after completing each test session three of the five rats ran for 20-30 minutes on an open field retrieving scattered food reward. Place fields were identified by calculating mean firing rate in 5x5cm bins and smoothing the binned rate map with a Gaussian kernel with σ = 15cm. Rate maps recorded across the two days were correlated, and only cells with a Pearson's correlation coefficients greater than 0.50 were retained. For the merging of AB3 with CD1, these inclusion criteria yielded 38 cells in four rats (Rat1 = 9 cells, Rat2 = 20 cells, Rat3 = 1 cell, Rat4 = 0 cells, Rat5 = 8 cells) (Figure 10). Only four cells were excluded due to the rate map stability criterion alone, and the inclusion of these cells results did not change the conclusion of any of my analyses.

4.3 Results

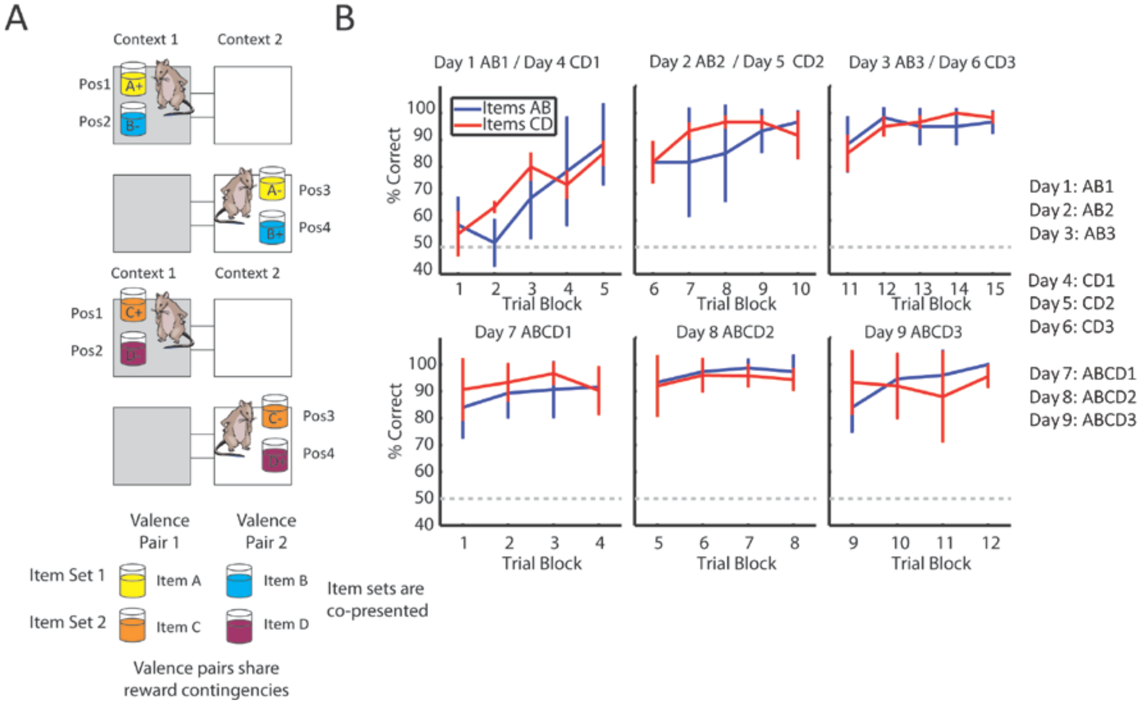

Figure 11 Summary of behavior. (A) For three days, rats learned to associate Item A with Context 1 and Item B with Context 2. For the next three days, rats learned to association Item C with Context 1 and Item D with Context 2. The items could be in one of four positions with the reward potential of an item is defined by the context (two positions per context). Items A & C are associated by common valence since they are rewarded (and not rewarded) in the same contexts; likewise for Items B & D. After the first six

days of training all items are presented within the same session for three days, though Items A and B always co-occur as a set as do Items C and D (B) Mean performance across all nine days of training. Rats performed higher than chance on Items A and B by the fifth trial block on the first day of training (day 1: AB1). In contrast, rats performed higher than chance on Items C and D on the second trial block on the first day in which those items were introduced (day 4: CD1). There were 15 trials per block. Error bars are standard deviation.

4.3.1 Rats rapidly acquired the association of items in context

Over the course of three days, rats learned that Item A was rewarded in Context 1 and Item B in Context 2 (Figure 11B). For the next three days, rats learned that Item C was rewarded in Context 1 and Item D in Context 2. Rats learned each item set within a single day (AB1 and CD1), as shown by the increase in the number of correct responses from chance performance during the first 15 trials (AB t(3) = 1.67, p =0.19; CD t(3) = 0.63, p =0.57) to performance significantly above chance by the final block of trials (AB t(3) =5.19, p =0.01; CD t(3) =8.34, p = 0.004; two-way repeated measures ANOVA item set X trial block, main effect of trial block, F(4,3) = 8.39, p = 0.002). Though there was no significant interaction between trial block and item set (F(4,3) = 0.39, p =0.77), rats reached above chance performance on the first item set only on the fifth training block (tr: 60 - 75), while on the second item set rats reached above chance performance by the second training block (tr: 16-30) (t(3) = 9.0, p = 0.003), at which point performance was marginally better than that observed during training on the first item set during the equivalent block of trials (t(3) = 2.83, p = 0.066). These results show rapid acquisition of the second item set, with rats performing above chance after one or two trials on each of the 8 item-position combinations. For both item sets, performance remained above chance on the subsequent two days of testing and throughout the three days in which all items were inter-mixed (all trial blocks, Ps <= 0.05). These results show that rats

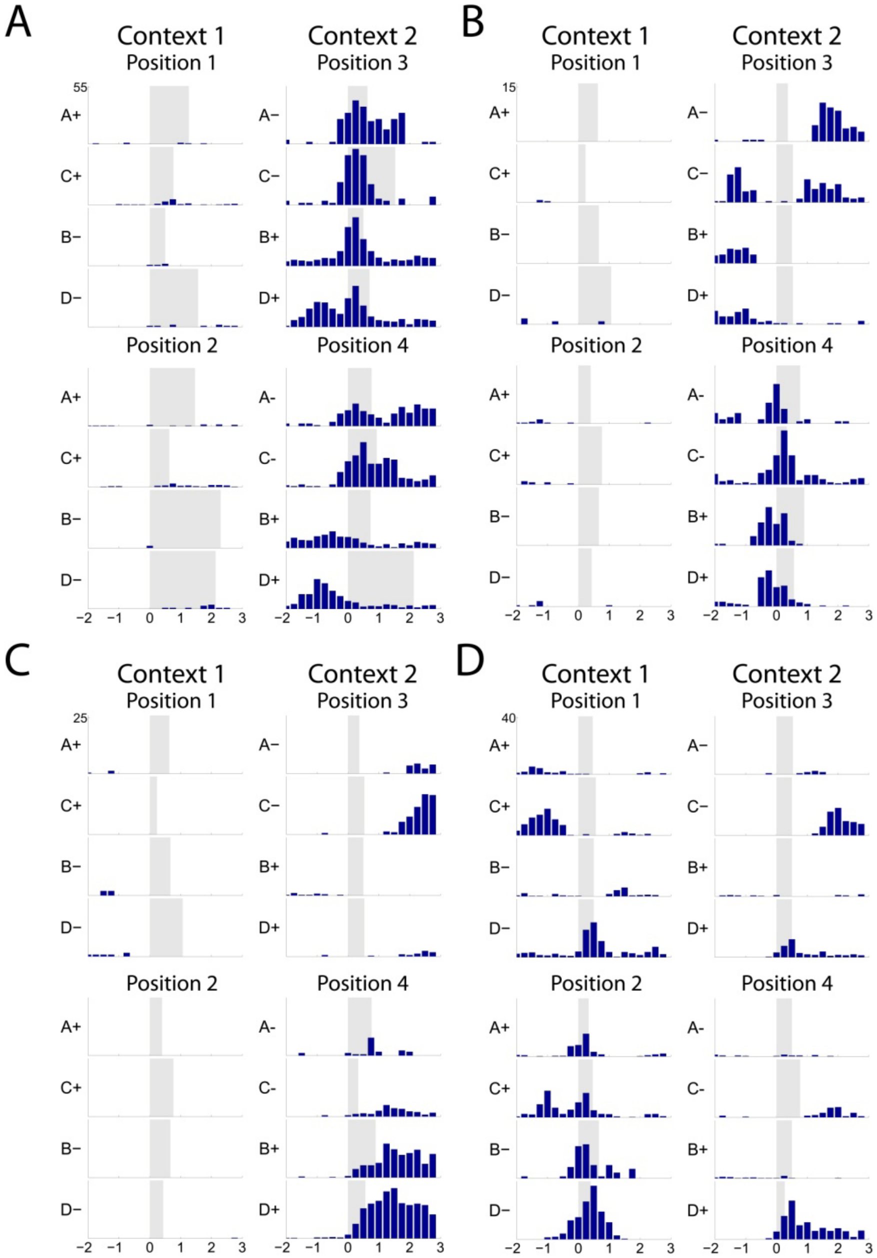

Figure 12 Example peri-event time histograms (PETHs). A-E. PETHs recorded for the four items presented within a location are grouped together with activity centered on item sampling. The y-axis trial average firing rate and the grey shading indicated the minimal trial length for that average. (A) A context cell fires in both positions in Context 1. (B) A position cell fires for any item in position 4, (C) A valence by position cell fires for rewarded (+) items (D & F) in position 4, (D) This conjunctive cell fired more for Item D in all positions but most of all in position 4.

associated items and context within a single recording session and suggest that learning of the second item set occurred faster than the first, similar to previous reports of rapid assimilation of new information into schemas (Tse et al., 2007).

4.3.2 The hippocampus encodes context, place, valence and item

I reasoned that rats could organize a schema of memories for the stimulus items by multiple dimensions: similarity in item reward assignment (valence), by the position in a context in which they were experienced, by the set of items that co-occur on each trial (Set AB, and Set CD), and by the context in which valence is defined. My initial analyses focused on identifying the extent to which item sampling events were related by common neural activity along these dimensions.

On the overtraining days in which all items were inter-mixed, I observed many single neurons that fired during stimulus sampling associated with one or more of these dimensions. A 4-way ANOVA was conducted to compare how firing rates were modulated by: context, position nested within context, set, and valence. The interaction of set and valence reflects item coding and interactions with those dimensions and position reflects the item and place conjunctive firing previously described (Komorowski et al., 2009). I found that firing rates of many of the cells (n = 571, CA1: 292, CA3: 279) were significantly modulated by task dimensions. I found that 40.8% of cells showed a main effect of context (CA1: 124, CA3: 109) , 34.2% of cells showed a main effect of position (CA1: 105, CA3: 90), 28.7% of cells showed a main effect for valence (CA1: 84, CA3: 80), 18.9% of cells showed a main effect for item set (CA1: 65, CA3: 43), and

13.3% of cells showed a significant interaction of valence and set reflecting item coding (CA1: 36, CA3: 40). An equal number of CA1 and CA3 cells showed a significant main effect for each task dimension.

Of the neurons that showed a main effect, 77.5% (n = 237/306) also showed an interaction between spatial (i.e. context and position) and non-spatial (i.e. set and valence) task dimensions. For example, Figure 12A shows a neuron with a main effect of context (greater activity in Context 2) but not position and an interaction with context and valence (greater activity for rewarded items). There were other neurons that distinguished positions within a context; for example, the cell in Figure 12B showed a main effect of position (firing more in Position 4). Some neurons fired during item sampling for either rewarded (or unrewarded) items; for example, the cell shown in Figure 12C that showed a position by valence interaction as well as a position by valence by set interaction (most firing for Item D in Position 4). I also observed neurons that fired for an item in all positions, though at different rates, like the neuron in Figure 12D that showed significant interaction of valence and set reflecting an item code and an interaction between valence, set and position, reflecting the conjunctive item in place code. A summary of the average firing rates for each cell for each item and place combination is shown in Figure 13A. Most cells had a preferred position, as indicated by high firing rate in only one position, and a preferred item and/or valence within that position, reflected in very few cells displaying equally high rates for all four items in the preferred position. Thus, the activity of individual hippocampal neurons is most appropriately characterized as multi-dimensional associated with relevant task parameters.

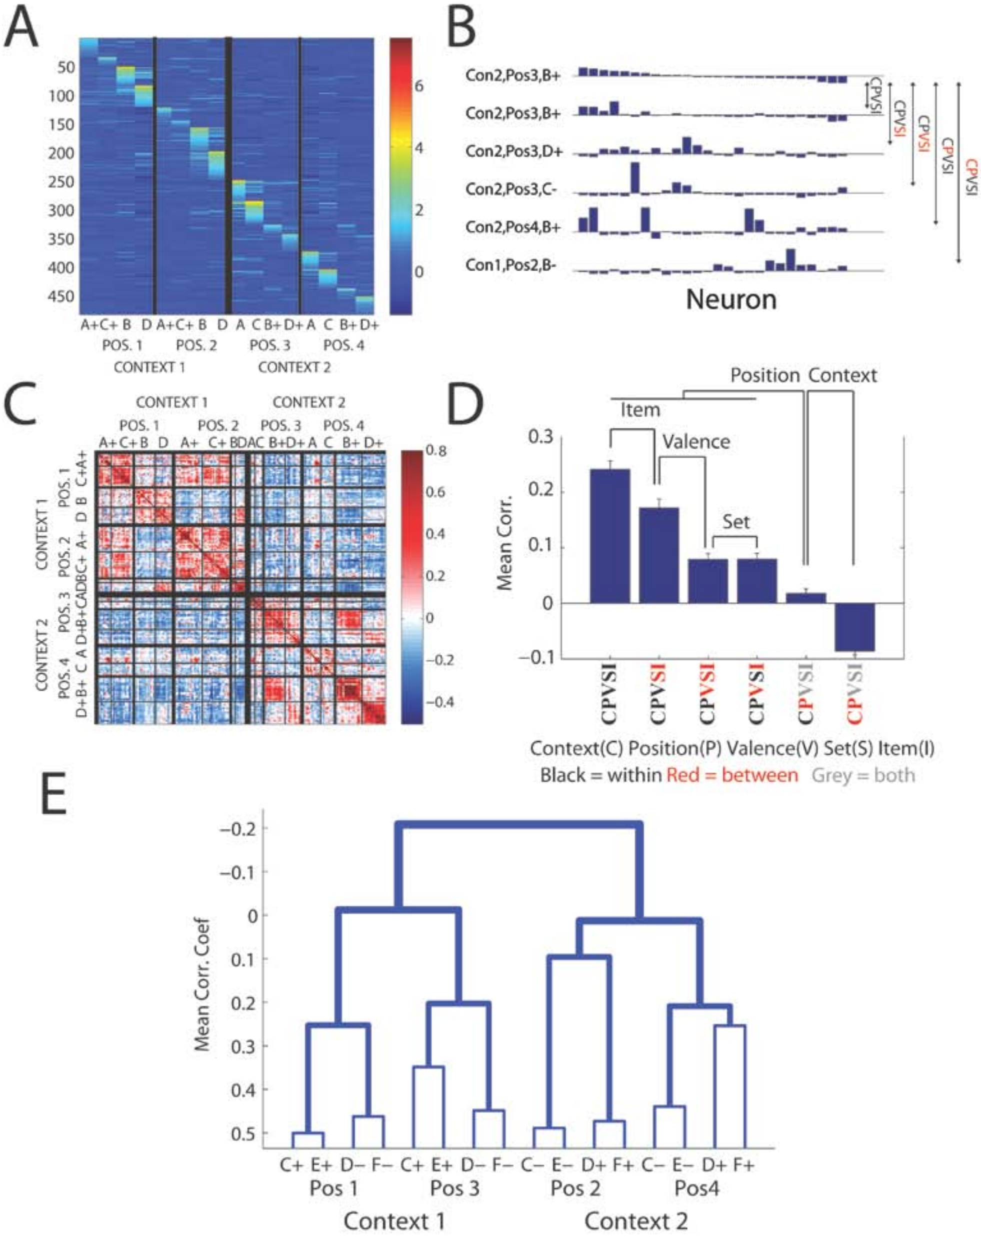

Figure 13 Related experiences are represented by correlated hippocampal ensembles (A) Summary of cell firing rate for all cells recorded during ABCD sessions. The color of the z-axis is the trial averaged z-normalized firing rate. The x-axis is each of the trial types sorted first by context, then by position, then by

valence, then by item. The strong item coding is evident in that very few cells show equivalence firing for all items present in a location. (B) Simultaneously recorded ensembles for seven trials. Each histogram is the firing rates (z-score) for each cell during one trial identified by Context (Con), Position (Pos), Item (A,B,C,D) and reward potential of the item in that pos (+,-). The range is the same for each trial (-1.5 to 7.5 st. devs). CPVSI is an acronym for each task dimension: Context Position Valence Set Item. Black lettering indicates that recordings were from trials of the same condition for that dimension (e.g. all black test indicates repetition of the same item in the same positions). Red lettering indicates that recordings were taken from trials between conditions (e.g. the red P shows ensembles from trials in Position 4 instead of Position 2. (C) An example correlation matrix from one ABCD session showing the ensemble correlation for each trial with each other trial with the trials sorted as above. The trials were presented in pseudo-random order, but were sorted them according to: context, place, valence, and then item. Lines divide trial types and from thickest to thinnest: context, place, valence, item. Red is positive correlation, blue is negative and white is zero. High ensemble correlation of trials with the same items in the same location can be seen by the high correlation in the square trial blocks along the diagonal. (D) The average correlation for within and between condition trials is shown for each task dimension (see table 1). For all dimensions except set, the correlations are higher for within condition trials than between. The error bars are SEM for the 15 sessions. CPVSI as defined above, with black indicating within condition correlations, red indicating between condition correlations and grey indicating both. (E) A dendrogram showing that ensembles of the same valence in the same position are most similar, followed by ensembles associated with items of opposing valence within the same position. Positions within the same context were also coded more similarly than positions in the opposing context.

To measure the similarity of ensemble representations between item sampling events, I calculated the average z-normalized firing rate for each neuron during each item sampling epoch, and constructed a population vector for every trial based on these normalized rates. Examples of activity patterns of a simultaneously recorded ensemble taken from one recording session are shown in Figure 13B. The overall ensemble pattern was similar for repetitions of items sampling events with the same item in the same position (top two ensemble events). Ensemble patterns for item sampling events with different items of the same valence were partially overlapping; several new cells joined the ensemble while others dropped out (top vs third events). Ensemble patterns for events with different valence in the same position showed more divergent firing patterns (top vs fourth event), and the overlap in the ensemble pattern decreased further still for events in different positions and in different contexts (top vs fifth and sixth events, respectively).

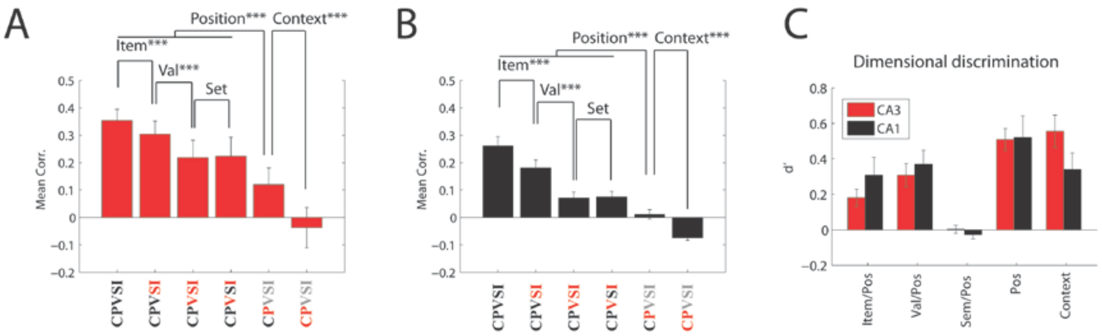

Figure 14 CA1 and CA3 similarly code task dimensions. The average correlation for within and between condition trials is shown for each task dimension for (A) CA3 and (B) and CA1. (C) The d' value estimating the degree of dimensional discrimination. All dimensions were significantly encoded except for set and both regions discriminated the dimensions equally after controlling for sample size (data not shown).

I created similarity matrixes to quantify the degree of ensemble similarity across all item sampling events for each recording session. Figure 13C shows an example similarity matrix for one session in which item sampling events have been sorted by four task dimensions: context, position, valence, and item; valence was chosen over set because, as will be presented below, set is not encoded by hippocampal ensembles. The similarity matrix shows that correlation of hippocampal ensembles recorded on different trials reflects the identified task dimensions. For example, in this similarity matrix, in the upper left corner it can be seen that ensemble patterns of Items A+ and C+ are strongly correlated, showing representational similarity of items that have the same valence in the same position. Looking down the left row, it can also be seen that ensemble patterns for A+ and C+ are strongly correlated between Positions 1 and 2, showing similarity in representations of these items across positions in the same context. By contrast, looking further down the left row, ensemble patterns for A+ and C+ in Position 1 are inversely correlated with those all of the items in positions 3 and 4, showing anti-correlations with the same items and items of the same valence when sampled in the alternate context.

To compare the similarity of ensemble representations associated with these dimensions, ensemble correlations were computed using the similarity matrices taken from all subjects and each session (n=15, 3 ABCD sessions from 5 rats). Five different task dimensions were considered: context, position, valence, set, and item (Table 1). To measure how distinctly hippocampal ensembles encode each dimension, I contrasted ensemble correlations for trials within the same condition of each dimension to those between conditions within a dimension while holding other dimensions constant. For example, to measure the extent to which ensembles encoded items, I compared ensemble correlations on events that involved the same item to those that involved a different item of the same valence at the same position. Ensemble patterns among events with the same item in the same location were most similar (mean $r = 0.25 \pm 0.02$), and were larger than that for events in the same position with different items of the same valence (mean $r = 0.18 \pm 0.02$; $p < 0.0001$), providing strong evidence that hippocampal ensembles differentiate items at specific locations. To measure coding of valence I compared ensemble correlations among events that involved different items of the same valence to those that involved different items of different valences at the same position. The similarities of ensemble patterns for events involving different items with the same valence were greater than those with different valence in the same position ($r = 0.08 \pm 0.01$, $p < 0.0001$), indicating strong coding evidence that hippocampal ensembles differentiate items by valence at each location. To measure coding of sets (AB and CD) I compared ensemble correlations among events that involved different items from the same set versus different items across sets at the same position. The similarities of ensemble patterns for events involving the item from the same and different sets were equivalent (within set mean $r = 0.08 \pm 0.01$; between set mean corr. coef. = 0.08 ± 0.01), indicating that hippocampal ensembles do not represent co-occurrence of items within sets ($p = 0.15$). To measure ensemble coding of positions, I compared the ensemble correlation of trials that occurred within the same position to the correlation of ensembles that occurred in different positions within the same context. Ensemble patterns of events occurring within the same position were more similar (within position mean corr. coef. = 0.16 ± 0.01) than those for events across positions (mean corr. coef. = 0.02 ± 0.01; $p < 0.0001$), reflecting the well known place code. Finally, hippocampal patterns were anti-correlated between the two contexts (between context mean corr. coef = -0.09 ± 0.01, probability of the observed percentage of negative correlations < 0.0001) , and the ensemble correlations for events that occurred in different contexts were significantly lower than the correlation of ensembles recorded in different positions within the same context ($p < 0.0001$), indicating that contexts that differentially define item valence are associated with distinct patterns of neural activity. No differences between CA1 and CA3 were observed (Figure 14 A-C). The same overall conclusions were found when significance testing for each dimension was done for individual sessions rather than for data pooled across rats (Figure 15D).

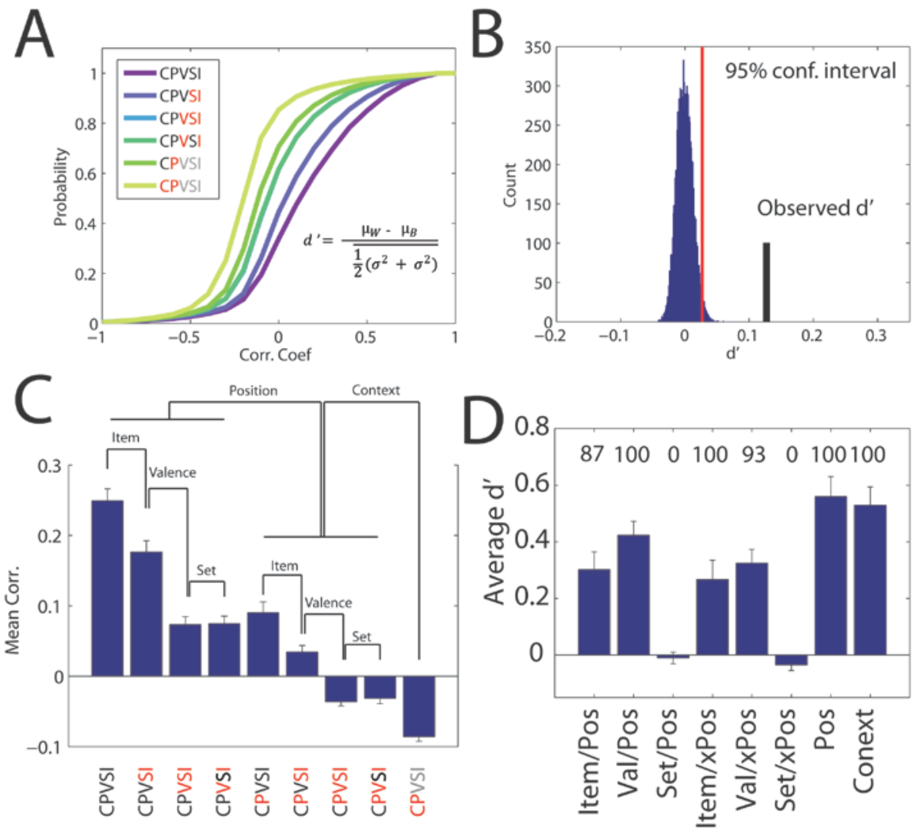

Figure 15 Control analyses confirm dimensional discrimination (A) The cumulative distribution function of ensemble correlations recorded during trials with varying degrees of overlap in the identity and position of the sampled item. CPVSI is an acronym for each task dimension: Context Position Valence Set Item. Black lettering indicates that trials were correlated within the same condition for that context (e.g. correlation of trials within Position 1). Red lettering indicates that correlations were taken from trials between conditions (e.g. correlation of ensembles from trials in Position 1 with those in Position 3). Grey lettering indicates correlation both within and between conditions for that dimension. Distribution were compared with a d' metric. (B) The item and position labels for each trial were randomly shuffled 10000 times and the d' metric specifying the degree of dimensional coding was for each of shuffled data set. The observed d' was then compared to the distribution calculated from the bootstrapped data. (C) Expansion of Figure 13D showing within (/Pos) and across position (/xPos) comparisons. Item and valence were coded within and between positions. Set was not encoded. (D) The d' values for each dimension. The number above each bar shows the percentage of sessions for which that dimension was statistically significant at $p < 0.05$.

This combination of results indicates a hierarchy of ensemble similarity during item sampling for events. To illustrate this hierarchy, I constructed a dendrogram in which each item and position combination (n = 16) was associated with a firing rate vector from neurons across all sessions and all rats. These rate vectors were correlated and rate vectors that produced the largest correlations were grouped into clusters (see methods). This analysis shows that the highest average similarity of ensemble patterns for item sampling events of the same valence in the same position, followed that for events of opposite valence in the same position, followed by that for the same items in the other position in the same context, and finally, the anti-correlation for events that occurred in the alternate context where items had opposite reward contingencies (Figure 13E). There are a larger number of possible dendrograms (C_{15} = 9,694,845) and therefore the likelihood of observing this binary tree by chance is low.

I confirmed these results by estimating the probability that a pattern of ensemble firing rates was recorded in each of the 16 item in position combinations using a Bayesian decoding algorithm. The decoding algorithm generated the same hierarchy of

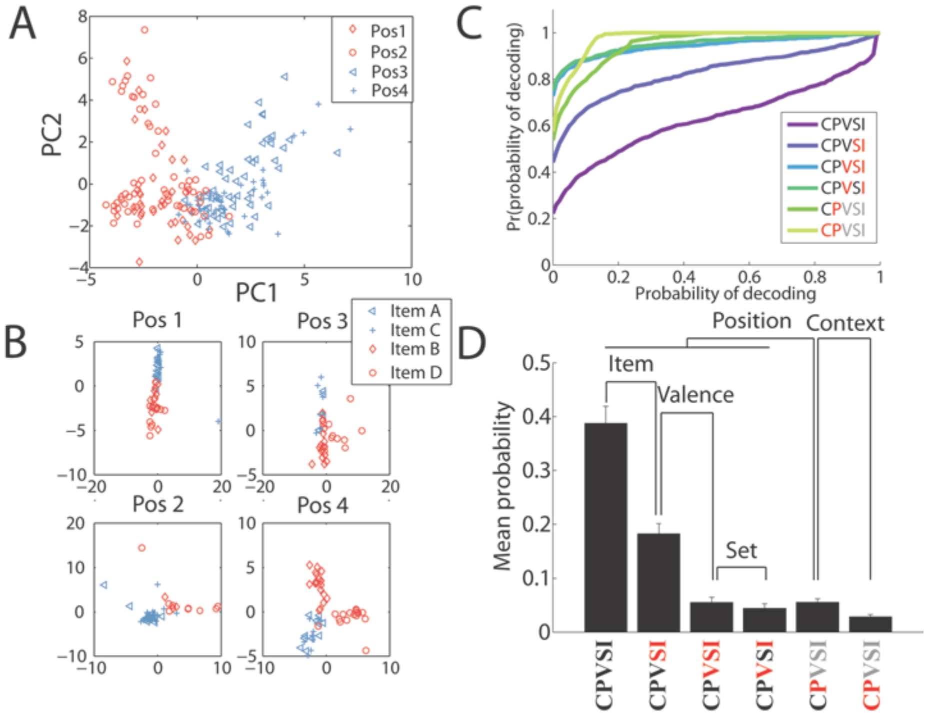

Figure 16 Bayesian classification confirms dimensional discrimination. (A) The principal component representation of every trial from one session (ABCD2) for one rat projected onto the first two PC axes. Context 1 = red, Context 2 = blue. Pos1 = ◊ , Pos2 = o, Pos 3 = Δ , Pos 4 =+. Note how ensemble representation from the same positions and contexts cluster. (B) Each subplot shows the trials in one location plotted onto the first two PC axes for that position. Color shows different valence, marker shows different items. Valence difference is evident by color separation in Positions 1 and 4. Item difference is shown by clustering of circles and diamonds in Position 4, for example. (C) The cumulative probability function of a Bayesian classifier estimating the probability that an ensemble was recorded in response to the same or different conditions from which the data was actually recorded. (D) The mean probability associated with decoding a trial to different item and place conjunctions that were the same and/or different as that in which the data was collected. The probability of classification to trials of the same condition was greater than classification between conditions within that dimension for: context, position, valence and item but not set. These results are qualitatively similar as those found using correlation coefficients. Color coded CPVSI as in Figure 13, black = decoded from same condition, red = decoded from different dimensions, grey = average of same and different.

ensemble similarity as the correlational techniques described above (Figure 16). Ensembles were most likely to have been recorded from the correct item and position combination (mean probability = 0.38 ± 0.03) which was greater than the probability of the unit activity originating from trials with a different item of the same valence in the same position (mean probability = 0.19 ± 0.02, $p < 0.0001$). This significant difference in probability reflects the strong item coding. The next most likely origin of the recorded ensemble was from trials occurring in the same position though containing an item of opposing valence (mean probability = 0.05 ± 0.01). This significant difference in probability reflects the valence code. Ensembles were equally likely to originate from trials of the same set (mean probability = 0.05 ± 0.01) as from trials of the other set (mean probability = 0.04 ± 0.01), confirming the lack of a distinct code for item pairing ($p = 0.11$). The position code was reflected by the higher probability that ensemble activity was recorded in the correct position (mean probability = 0.20 ± 0.02) than the probability of being recorded in the incorrect position within the correct context (mean probability = 0.05 ± 0.01, $p < 0.0001$). Finally, ensembles were least likely to have been recorded in the opposing context (mean probability = 0.03 ± 0.004), with a mean probability of decoding to the wrong context lower than decoding to the wrong position within the same context ($p < 0.0001$) and lower than decoding to the opposing context by chance ($p = 0.0013$).

Based on these analyses, I conclude that items, their valence, and the location where they appear within a context are encoded by hippocampal ensembles during item sampling. Notably, set (item pairing) is not encoded by hippocampal ensembles in the current task. Finally, hippocampal ensembles separate representations of events in different contexts, suggesting opposing schemas are created for events in contexts that are meaningfully distinct.

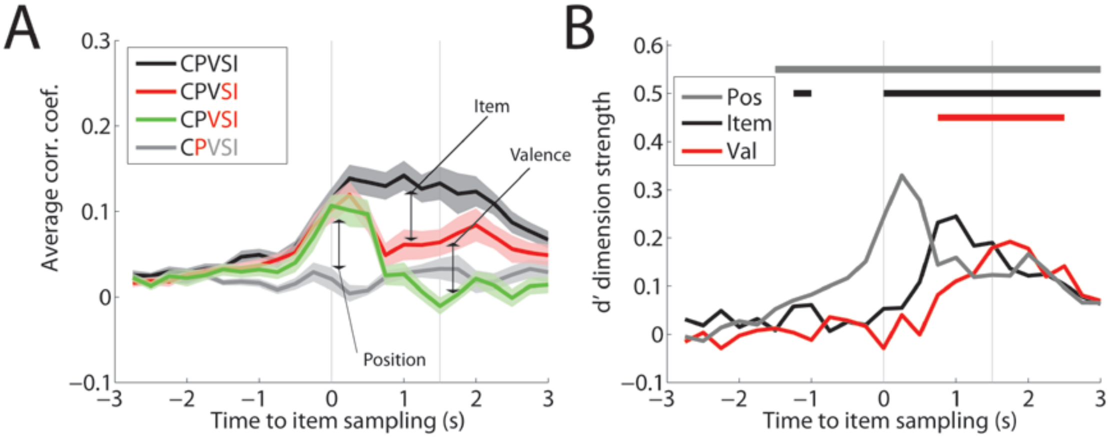

Figure 17 During item sampling, position is coded first, followed by item and finally valence (A) The ensemble correlation analyses was done using the rates taken from 250ms bins centered around sampling for trials in the same location. Sampling epochs shorter that 1.5s were excluded. Color coded CPVSI as in Figure 13. Error bar is SEM (n = 15). (B) The average d' for item (black) and valence (red) and position (green) for trials in the same position. Color coded bar above graph shows time points in which that dimension was significantly coded.

4.3.3 Task dimensions are expressed at different times during item sampling

I next explored whether different dimensions of an event are encoded as a unified representation of the event, or whether each dimension is encoded distinctly. I reasoned that, if an item's identity, valence and position are integrated within a unified representation, they should all be expressed simultaneously. On the other hand, if these dimensions are distinctly encoded, then their representations might be expected to appear sequentially during the item sampling period. To address this question, firing rates for each item sampling event were calculated for 250ms bins centered around the onset of item sampling and ensembles recorded on different trials were correlated at each time bin.

These analyses indicated that information about item quality, valence, and position are expressed sequentially rather than simultaneously. Position information was expressed first. For all types of events that occurred within the same position, at the outset of item sampling there was an initial increase in the mean ensemble correlation irrespective of whether the trials contained the same item (Figure 17A, black), different items of the same valence (Figure 17A, red) or different items of opposing valence (Figure 17A, green) indicating a reliable position code at the onset of sampling. In contrast, ensemble correlations between positions were, on average, low throughout item sampling (Figure 17A, grey). The difference between the within position versus between position correlations was assessed with a d' metric that reflects the degree of position coding (Figure 17B, green). Position d' was significant for approximately 1000ms before item sampling and peaked 250ms after item sampling before significantly decreasing (max d' - min d' = 0.19, probability observed d' range in bootstrap data < 0.001) (Figure 17B, grey).

I hypothesized that this decrease in position coding was caused by the onset of item and valence coding that increased the variability of firing within a position. For trials with the same item in the same position, the average ensemble correlation increased upon arrival to the item and remained high throughout sampling (Figure 17A, black), indicating reliable coding throughout the sampling epoch. Trials with different items of the same valence (Figure 17A, red) showed the same initial increase in ensemble correlation, which subsequently decreased throughout item sampling, reflecting the divergence of neural firing patterns in response to different stimuli – the item code. The d' metric for item coding (Figure 17B, black) was statistically significant from the onset of sampling until the reward was retrieved, and peaked 1000 ms after item sampling. Therefore, item information increased at the same time points when position coding decreased.

Valence information influenced hippocampal firing rates last during item sampling. When comparing trials with different items and different reward valence (Figure 17A, green), neural activity was initially well correlated, reflecting the position code. However, 750ms after item sampling, ensembles recorded during trials with different items of the same valence were significantly more correlated than ensembles recorded during trials of items with different valence, as shown by the increase in the valence d' metric at these times (Figure 17B, red). After 1000ms of item sampling, the correlation of neural activity recorded on trials with opposite valence became statistically indistinguishable from the randomly shuffled data ($Ps > 0.2$ at all time points 1000-1500 ms), despite the rat occupying the same position throughout the 1500 ms sampling epoch. Therefore, neural patterns did not distinguish between items of opposing valence until late into item sampling. However, once valence coding emerged, ensembles recorded in the same position for trials of opposing valence were completely uncorrelated, suggesting that item and valence inputs dominated over position input by the end of item sampling.

Firing rates were modulated by position early during item sampling followed by item identity and finally item valence. The item identity and valence coding that emerged during sampling decreased position coding by increasing the variability of firing rates within a fixed location. Moreover, similar activity in response to the same item presented in different positions (Figure 15C) further diminished the position code. These results show that different types of information dominate hippocampal computations at different times, suggesting that these task dimensions are encoded as distinct elements rather than integrated conjunctions.

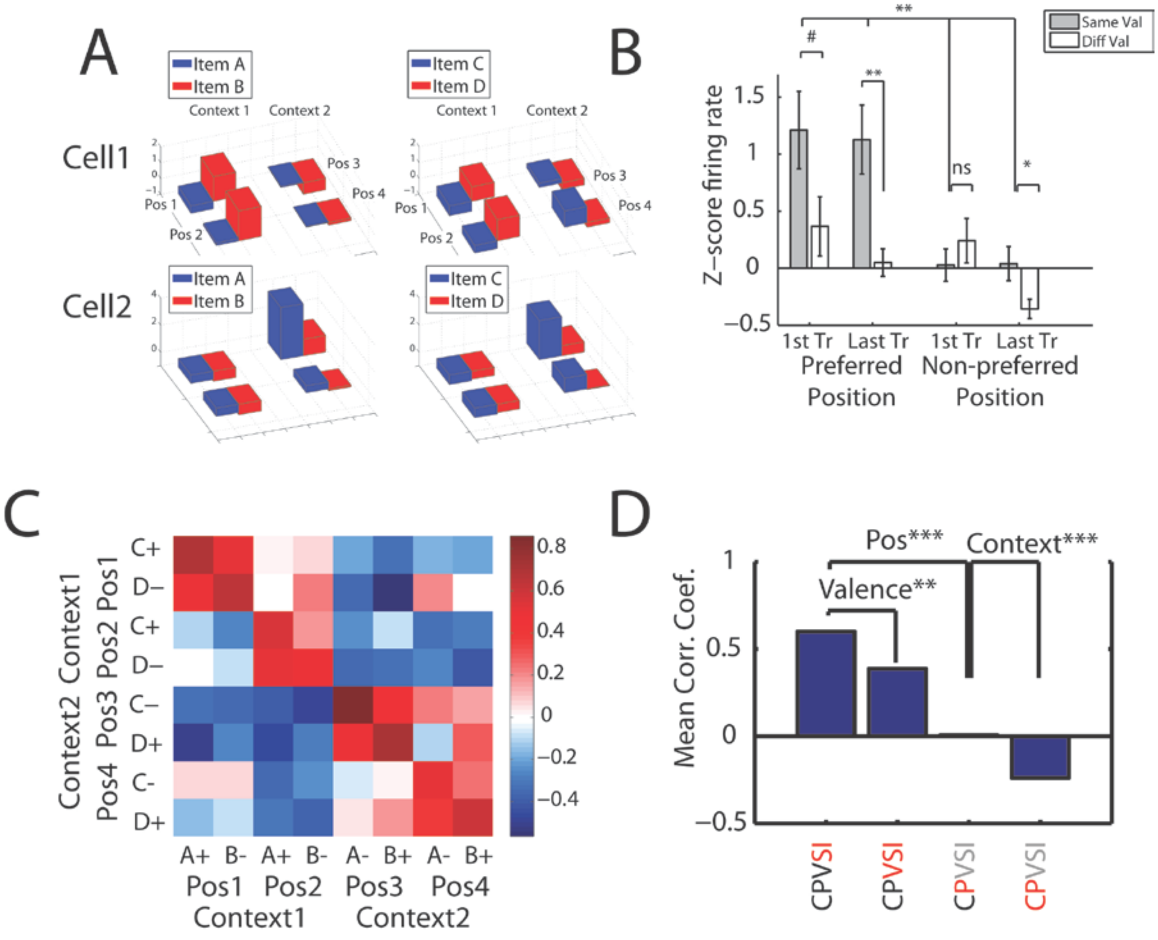

Figure 18. New items are encoded within established schema (A) The firing rate for two neurons recorded on the last day of training on the first item set (AB3) and the first day of training on the second item set (CD1). Cell 1 fired for Item B in both positions in Context 1 and the next day for the equivalent item in the same two positions. Cell 2 fired for equivalent Items A and C, but only in Position 4. (B) The mean z-score firing rate recorded on the first and last encounter with each item, as identified on AB3 training. Initially, there was a non-significant trend for cells to fire at higher rates for the valence equivalent item in the preferred position and no difference between firing for the two items in the non-preferred position within the same context. However, by the final trial for each item, cells fired significantly more for the valence equivalent item in both positions within the preferred context. (C) Data from AB3 and CD1 were merged and the median ensemble rates for each item in place conjuction were calculated and ensemble vectors were correlated from AB3 to CD1 to create the 8x8 similarity matrix. (D) The valence, position and context coding that developed over CD training was preserved and extended to training with the new item set. Color coded CPVSI as in Figure 13

4.3.4 New associations are encoded within related hippocampal networks

I next explored how new information is assimilated within an existing representation of related memories. Views on schemas suggest that new memories are stored within the structure of related existing memories (McKenzie et al., 2011; McClelland et al., 2013, 1995). Therefore, I tested the hypothesis that neural activity recorded during training of the second item set (CD) would be similar to activity already established during training on the first item set (AB). I expected a stable position and context code, because the meaning of these task dimensions was not altered by adding new items within the same spatial organizations. Also, based on the view that schemas link closely related events, I also expected similarity in the ensemble representations of items that similarly predict reward or non-reward within a context.

To determine whether firing patterns reflected the similarity of the two items sets, I first identified the item and position conjunction that generated the highest firing rate during the last day of training on the first item set (AB3). I then used this preferred item and position for each neuron to compare firing rates in responses to items of the same and different valence during the first day of training on the second item set (CD1). This required comparing the activity patterns of the same neurons between two successive daily sessions, and so this analysis is focused on 38 neurons that were carefully selected as having identical characteristics across days (see Methods). As predicted, I observed that in 52.6% of the cells (n = 20/38), the median firing rate recorded during the second item set was higher for the same-valence item in the preferred position than in any of the other seven item and place combinations ($\chi^2 = 52.4$, $p = 4.65^{-13}$). For example, Cell 1 in Figure 18A fired for unrewarded Item B- in both positions in Context 1 and fired equivalently for unrewarded Item D- in both positions in Context 1 (main effect valence F(1,213) = 16.96, p = 0.0001, Item B vs Item D in Positions 1 and 2, Ps > 0.05). Similarly, Cell 2 showed equivalent firing for unrewarded Items A- and C- in only one of the positions in Context 2 (main effect valence F(1,213) = 171,1. $p = 4.41^{-29}$; Item A vs Item C in Position 3, t(8) = 0.33, p = 0.75).

To examine when common firing across item sets emerged, firing rates were compared on the first and last encounter with each new item (Items C & D) for which animals made the appropriate behavioral response. On both the first and last trials, firing rates were higher in the preferred positions (mixed-model repeated measured ANOVA, main effect position $F_{1,37} = 10.68$, p = 0.002, post-hoc t-tests at each time point, Ps < 0.05) revealing a stable position code across item sets. Even on the first encounter with each item in the preferred position (actual trial number 1-48, mean trial number = 11.25, s.d. = 11.86), there was a trend that cells fired more for the same-valence item (mean z-score rate = 1.24 ± 0.35) than the different-valence item (mean z-score rate = 0.38 ± 0.28;paired t-test, t(37) = 1.99, p = 0.053) (Figure 18B, first trial) though an equal number of cells fired maximally to the same and different valence items (n=20/38; χ^2 = 0.02, p = 0.87) . By the end of training, cells clearly showed a firing rate preference for the equivalent item. On the last trial for each item in the preferred position, cells fired at a significantly higher rate for same-valence items (mean z-score rate = 1.17 ±0.31) as compared to items of opposing valence (mean z-score rate 0.11±0.13, t(37) = 3.52, p = 0.001) (Figure 18B, last trial) and far more cells (n = 33/38) fired maximally to the equivalent item (χ^2 = 19.2, p = 1.19^{-5}).

I then tested for differences in firing to the same- and different-valence items in the non-preferred location. On the first trial with each item, there were no differences in firing rates (t(35) = 1.47, p = 0.15), though by the final trial there was greater firing to same-valence items (t(35) = 2.81, p = 0.008) (Figure 18B, non-preferred position). These results show that during learning, cells rapidly developed a preference for the item of the same valence in both the preferred and non-preferred positions.

I next asked whether the overall neural representation developed during training of the first item set was re-invoked during training of the second item set. I adopted a similar correlation analysis as that previously described for analyzing days when all items were inter-mixed, though instead of comparing simultaneously recorded ensembles on individual trials, I calculated the median firing rate for each item and place combination (four items in four positions) for each cell (n = 38) and concatenated these rates across rats for subsequent analysis. The ensemble similarity in response to different item and position combinations was established by examining the correlations of cell activity recorded at every item and position combination for the first item set (two items in four positions) with firing rates recorded during initial training of the second item set (two new items in the same four positions). The resulting 8x8 similarity matrix describes which item and position combinations result in overlapping hippocampal ensembles (high correlations), and therefore operationally defines the neural network organization of the task dimensions (Figure 18C).

As suggested by the single cell findings, when analyses were limited to trials occurring in the same position, trials with different items of the same valence were more correlated (mean $r = 0.60 \pm 0.04$) than trials with different items of opposing valence (mean $r = 0.38 \pm 0.04$, $p = 0.001$), suggesting a subset of cells that fired during the first item set fired similarly for the equivalent item of the second item set.

I also found evidence for a stable position and context code. Trials recorded in the same position across days were more correlated (mean $r = 0.49 \pm 0.04$) than trials recorded in different positions (mean $r = 0.02 \pm 0.04$, $p < 0.0001$), which were in turn more correlated than trials recorded within the other context (mean $r = -0.24 \pm 0.03$, $p < 0.0001$) (Figure 18D). Together these data show that neural correlates of the entire task structure that developed during training on the first item set were adopted and extended to encode related items that occurred within the same context.

Together these data suggest that neural activity during learning of new item associations reflected binding of information acquired over multiple days into a unified task schema that represents equivalent item and position associations within overlapping hippocampal networks. The item and place conjunctions that generated correlated neural activity at the end of learning both item sets were qualitatively the same as those observed during overtraining days in which all items were presented, suggesting that the schema that developed during learning was maintained until the final day of behavioral testing.

4.4 Discussion

The present findings show that hippocampal cells are driven by a hierarchy of multimodal inputs that influence their firing rates associated with relevant task dimensions. For both spatial and non-spatial stimuli, related features developed an overlapping hippocampal code, while features that required divergent behavior and differential reward expectation developed distinct hippocampal codes. These data show that hippocampal functions go beyond unifying spatial and non-spatial information in the service of creating unique memory traces for places where particular events occur. Instead, the hierarchical nature of the single trial correlations suggests that the hippocampus encodes the relationships between items, their meaning, and the places they occur. Moreover, the anti-correlation of the hippocampal representation of events that occur in meaningfully different contexts suggests the existence two schemas, one for each context.

Firing rates were modulated by spatial context, position within a context, item identity and item valence. These results add to a confluence of evidence supporting the hypothesis that the hippocampus encodes spatial and non-spatial features of an experience. Several studies on hippocampal neural activity in rodents have shown that cells respond to conditioned items and tones only within the place field for each cell (Tort et al., 2011; Moita et al., 2003), while others have reported responses to common stimuli across multiple locations (Wood et al., 1999; Manns et al., 2009; McKenzie et al., 2013; Singer et al., 2010; Eichenbaum et al., 1987). I observed both types of responses as shown by the valence and item coding within and across positions.

I observed that most cells were not spatial or non-spatial, but rather responded conjunctly to spatial and non-spatial dimensions (Komorowski et al., 2009; Manns et al., 2009; Anderson and Jeffery, 2003; Wiebe et al., 1999; Deshmukh and Knierim, 2013). The convergence of the ‘what’ and ‘where’ pathways at the level of the hippocampus provides an anatomical basis for the spatial and non-spatial firing correlates (Witter et al., 2000a). The mnemonic importance of this conjunctive firing is suggested by studies in animals (Balderas et al., 2008; Komorowski et al., 2013; Tse et al., 2007; Day et al., 2003; Parkinson et al., 1988) and humans (Holdstock et al., 2002; Vargha-Khadem et al., 1997) that have reported profound deficits in the recollection of an item with its context after hippocampal damage. Together, these results, and others (Davachi et al., 2003; Diana et al., 2010) have generated the hypothesis that the hippocampal contribution to memory is binding of items within a contextual framework (Eichenbaum et al., 2007; Diana et al., 2007; Ranganath 2010).

Many have argued that binding item in context increases the uniqueness of memory traces allowing the storage and recollection of distinct episodic memories (Gilbert et al., 1998; Yassa et al., 2011; Norman, 2010; Norman and O'Reilly 2003; Xu and Südhof, 2013; Hasselmo and Wyble, 1997). This hypothesis is supported by reports that damage to the hippocampus (Wiltgen and Silva 2007, 2010; Frankland et al., 1998; Winocur et al. 2007, 2009) or the passage of time (Winocur 2007; Biedenkapp and Rudy, 2007) causes learned behaviors to generalize to novel spatial contexts. The observation that place cells generate independent 'maps' in different contexts (Paz-Villagrán et al., 2004; Spiers et al., 2013; Leutgeb et al., 2004; Kubie and Ranck 1983; Hayman et al., 2003) has also been interpreted as evidence that the hippocampus generates distinct memory traces.

Here I found strong evidence for a hippocampal code that distinguishes between different item and place conjunctions, but I also found evidence that related items and positions were associated with correlated hippocampal firing patterns. Other recording studies have also found a common hippocampal code for related events. When animals traverse visually similar parallel arms of a W-shaped maze to retrieve reward, cells begin to fire at the equivalent positions on each arm as learning progresses (Singer et al., 2010). Another study in which rats retrieved water at different location within a circular track found that the cells that fire at new goal locations are those that fire for other goal locations that required the same behavioral demands to retrieve reward (McKenzie et al., 2013). Studies on human hippocampal cells have found that cells respond to categories of items (Kreiman et al., 2000) and also to a broad range of stimuli that gain functional equivalence, for example images of people from multiple angles and text strings of the person's name (Quiroga et al., 2005). With these coding properties, the hippocampus has the potential to record of experience by coding for both the unique and common features of events within overlapping networks (Eichenbaum, 2004).

The overlapping code for related events hints at the critical computation performed by the hippocampus in tasks in which the correct behavioral response can only be realized when information learned across episodes is integrated and generalized. In a task similar to the one employed here, fear memory generalized between items that similarly predicted reward within a context (Honey and Watt, 1998), thus necessitating a mechanism by which items that shared a common contextual association become associated with one another. In other tasks, hippocampal damage in rodents (Dusek and Eichenbaum, 1997; Bunsey and Eichenbaum 1996; DeVito et al., 2010b; Van der Jeugd 2010) and people (Myers et al., 2003) causes memory deficits limited to probe tests that measure generalization to a novel condition. In a direct parallel to my results, an fMRI study using multi-voxel pattern classification of the BOLD response has shown that when two images come to predict a third common image, the pattern of activity in the medial temporal lobe reflects this transitive association by responding similarly to the two distinct images (Zeithamova et al., 2012). I suggest that similar firing patterns associated with related events can potentially link those memories to allow accurate generalization.

I found that events that took place in different contexts were associated with anti-correlated hippocampal activity. This is in contrast to the independent place code that is often reported for unrelated contexts and also distinct from observation of anti-correlated activity for different behaviors executed within the same space (Bahar et al., 2012; Markus et al., 1995). In my experiment, rats executed the same behaviors in both contexts, though in response to different items. Therefore, the anti-correlation reflects either the context and item associations or the behavior and item associations, but cannot simply reflect behavior or arousal alone (O'Keefe, 1999).

I conclude that related item and context associations are encoded within overlapping hippocampal networks, while item and context associations that lead to opposite outcomes are stored within distinct networks. In tasks for which the same behavior is required, this overlapping contextual code can lead to enhanced transfer of knowledge through generalization (Eich, 1985; Tse et al., 2007), though in other circumstances, such as reversal training, a common code can lead to interference (McDonald et al., 2002; Spear et al., 1980; Corcoran and Maren, 2001).

Previous studies have established that place cells maintain a somewhat stable place code across days (Ziv et al., 2013; Thompson and Best, 1990; Mankin et al., 2012). I extend these findings and show that cells also maintain a stable valence code across days and across different items. After characterizing the items that cells prefer, 79% of those cells went on to fire more in response to other items of equivalent valence. These results suggest that, at the time of learning, new information is encoded within extant networks that stored related information. Reports of hippocampal ‘preplay’, where neural patterns recorded during behavior can be observed before the subject explores a well learned (Louie and Wilson, 2001) or novel (Dragoi and Tonegawa, 2011,2013) environment, suggest a potential mechanism by which retrieval at the time of learning can link past experience with present. My findings suggest that this linkage could occur selectively for meaningfully related experiences. I hypothesize that this overlapping code at the time of learning allows for a transitive association of the two items by virtue of their common association with a context.

Central to the relational theory for hippocampal schema development is the notion that information is never stored in isolation, but rather assimilated within networks of related memory traces. Here I show that a neural correlate of memory, item-in-place coding, generalized to new related experiences as predicted by models that require novel experiences to be integrated within hippocampal networks of related experiences (Eichenbaum, 2004). Similarities in hippocampal activity between familiar and novel conditions likely reflect the integration and generalization of related memories, arguably a primary purpose of memory systems.

5 Discussion

In this thesis, I have argued that the key function of the hippocampus is to incorporate new information within relevant existing memory networks to store memory schemas (Chapter 2). In Chapters 3 and 4, I begin to explore the neural substrate of a hippocampal schema as reflected by network activity that links and distinguishes related associations, integrates new memories to pre-existing schemas, and modifies the pre-existing schema to accommodate additional memories. I have identified neural correlates that are consistent with the relational schema hypothesis and that challenge key assumptions that episodic memories are stored as unrelated traces within the hippocampus (Chapter 1.3).

5.1 Novel findings that add to our understanding of schema

I have taken advantage of two known correlates of memory – firing at goal sites (Dupret et al., 2010; Chapter 3) and conjunctive firing fields for items in positions (Komorowski et al., 2009, 2013; Chapter 4) – to investigate how these signatures of memory change to reflect the overall structure of a set of experiences. The results from the two recording experiments (Chapter 3 and Chapter 4) revealed several properties of hippocampal cells in CA1 and CA3 that are important ‘raw ingredients’ for a schema-based memory system stored within the hippocampal circuitry.

First, I have found evidence that related experiences are represented by correlated firing patterns while situations that require divergent behaviors are associated with anti-correlated activity. Data points from both studies speak to this conclusion. From

Experiment 1, I observed that the cells that fired at novel goal locations were those that fired at previously learned goal locations. This generalization of firing observed in a subset of cells caused an increase in correlation of firing patterns observed at old and new locations. The related behavior or common reward expectancy in different positions was encoded by an overlapping set of cells. From Experiment 2, I observed that activity recorded in locations that similarly predicted the reward potential of an item (i.e. locations within the same context) was more correlated than activity recorded in locations that differentially predicted reward potential (i.e. locations in opposite contexts). Similarly, activity recorded in response to items that predicted reward within the same locations was more correlated than activity recorded in response to items that differentially predicted reward. These results show that information about spatial and non-spatial commonalities between experiences is present in the firing rates of hippocampal ensembles. Furthermore, the cells that fired in response to new items also fired in response to previously learned items that similarly predicted reward in the same position. In both experiments, cells that fired in response to familiar stimuli generalized this firing to novel, related stimuli during learning. These results suggest that related schemas are recalled during learning. The cognitive effect of this neural generalization may correspond to known memory benefits of learning information that is consistent with prior schema (Chase and Simon, 1973).

I also observed that cells responded to single elements and to conjunctions of elements. From Experiment 1, there were cells that fired at a subset of locations in which rats engaged in the same behavior to retrieve reward. From Experiment 2, I observed cells that: fired for the same items in multiple positions, fired for multiple items of the same reward potential, fired for multiple items in the same position, and fired in multiple positions within the same context. Therefore the majority of cells were tuned to related item and place conjunctions rather than to a unique item and place conjunction. During item sampling, information about position, item identity and reward potential was expressed sequentially, further suggesting that the elements of the experience could separately *or* in conjunction drive cells to fire. The addition of the 'OR' gate allows for events with overlapping features to be encoded by similar ensembles of neurons. Future studies will be needed to determine whether 'nodal cells' that fire for memories with common elements (e.g. cells that fire for rewarded items) actually associate cells that fire to distinct aspects of an experience (e.g. cells that fire for specific item and place conjunctions).

Finally, data from Experiment 1 shows that learning in one position causes firing patterns at other well-learned goal sites to change, though not to the detriment of the amount of spatial information carried by cells that fired at the familiar goal locations. These data suggest that new information was stored via the reorganization and reconsolidation of related memory networks (Chapter 2). Models of memory explicitly warn of the dangers of this kind of retrograde interference (Section 1.3), however, in my experiments reorganization did not disrupt the amount of spatial information nor the degree of correlation between ensembles recorded at well-learned goal locations.

I observed three basic ‘raw ingredients’ that are key requirements of a relational schemas model and that run counter to predictions of a model for episodic memory in which traces are stored independently of one another. I found that, 1.) the correlation in activity patterns between situations reflects mnemonic relatedness, 2.) single elements of an experience can drive cell activity, not just conjunctions, and 3.) learning causes reorganization in related memories networks without causing ‘catastrophic inference’.

The ‘Hebb-Marr’ model emphasizes that unique combinations of entorhinal inputs drive a random, small number of dentate cells. This process of pattern separation produces decorrelated hippocampal traces which are stored in recurrent collaterals of CA3. The goal of these transforms is to protect the episodic traces from subsequent interference and modeling work has shown that storing orthogonalized traces can achieve this goal. Despite the attractive theoretical advantage of orthogonalizing the code for a set of experiences, my results show that hippocampal firing patterns reflect both spatial and non-spatial relatedness.

Abandoning pattern separation as the ideal computational goal of the hippocampus will ultimately require simultaneous recordings of the hippocampal inputs and the hippocampus to determine whether activity patterns recorded during related situations are less correlated within the hippocampus that within its inputs. A set of studies that have rotated local and distal cues in opposite directions have indeed found that rotational discordance causes independent patterns of activity to emerge in the dentate gyrus while CA3 maintains stable and relatively cohesive representations that rotate mostly with local cues (Neunuebel et al., 2014). In a related study, cells were recorded in the rhinal cortices as animals performed the same task and were analyzed with the same analyses revealed that MEC rotates with distal cues while LEC rotates more with local cues. Analysis of pattern separation within each entorhinal region reveals nearly orthogonal representations with the largest angular displacement of local and distal cues (Neunuebel et al., 2013). Therefore, it is not at all clear whether the dentate gyrus orthogonalizes correlated patterns of activity or whether the independent codes are inherited from the LEC and MEC. Furthermore, even if the dentate does create independent representations, this code was not inherited in area CA3. Imaging studies that lack resolution to distinguish dentate activity from CA3 activity also confirm that pattern separation of activity patterns originating in DG/CA3 is not inherited in CA1 or the subiculum (Yassa and Stark, 2011; Bakker et al., 2008). The Hebb-Marr model strongly predicts that the pattern separation function of the dentate must be preserved in its post-synaptic targets if catastrophic interference is to be avoided.

My data shows that a rich hierarchical correlation structure exists and therefore is likely used by the hippocampus as evidenced by the hippocampal involvement in inference learning and rapid consolidation. However, the pattern separation function of the dentate and CA3 was critical in computational models of episodic memory to prevent ‘catastrophic interference’. Below I discuss features of the hippocampal circuit that may allow the hippocampus to minimize the trade-off between long-term memory storage and rapid encoding while promoting the integration of memories across time.

5.2 The hippocampus as a plastic, recurrent hub: dimensional expansion realizes memory overlap and decreases interference

Several features of the hippocampus make it well suited as a relational memory structure. First, the hippocampus has one of the highest densities of NMDAR in the brain (Monaghan and Cotman, 1985) endowing its synapses with rapid plasticity that is contingent upon coincident pre- and post-synaptic activity (Bliss and Collingridge 1993). Second, the hippocampus has a large degree of intra- and inter-region recurrent connectivity (Kjonigsen et al., 2011). Third, as reviewed in Chapter 1, there is a high degree of convergence of many brain regions upon the hippocampus; a large percentage of the brain is two synapses away from the hippocampus and therefore the hippocampus allows these regions to interact. In the parlance of small world networks, the hippocampus is a hub. These properties as well as several of those highlighted by the Hebb-Marr model potentially enable the hippocampus to store unique episodes and associate multiple related episodes within a common network or schema.

An important aspect of the ‘Hebb-Marr’ model is the recognition of sparse firing in the hippocampus. Sparse activity was theoretically accomplished in two ways: by cells only firing to unique feature conjunctions, and via lateral inhibition in which the most active neurons silence competitors. I have argued against the notion that each conjunction is stored in independent traces; however, the observed sparse firing patterns need to be reconciled with correlated activity for related circumstances.

As suggested in the early models, it is useful and biologically feasible, for the hippocampal network to have a homeostatic ideal for a given number of cells to be active (Shew et al., 2011). Many recording and imaging studies have confirmed that in general there are 20-40% of cells active in any given environment (Ziv et al., 2013; Wilson and McNaughton, 1994; Guzowski et al., 1999) though only 1-2% active in a given moment (Viskontas et al., 2006). My data suggests that active cells are not a random subset, though the question remains: which cells are they?

Like others, I propose that the recurrent circuitry is important for pattern completion, though not simply for recall of patterns originally elicited by a pure feedforward perceptual drive. If the current input pattern is permitted to drive recall within the recurrent collaterals, then representations reflecting the present and past may be coactive and bound together (Wallenstein and Hasselmo, 1997; Hasselmo and Eichenbaum, 2005; Levy, 1996; Lisman, 1999). Associating the past and the present within a recurrent network allows for some interesting properties to emerge and dramatically alters the observable representations to include not only what is present but what is related to the present. Accordingly, hippocampal representations are sparse but not random. Cells respond to a small number of stimuli, but the stimuli in which they respond to are related: hippocampal cells are fundamentally category cells.

Several computational models (Gluck and Myers, 1993; Howard and Kahana, 2002; Levy, 1996; Lisman, 1999; Wallenstein and Hasselmo, 1997) predict that the representational code may reflect future predictions or past states in addition to the present stimuli. For example, in a model put forth by Gluck and Myers (1993) hippocampal representations are defined both by the stimulus itself and the learned associations between current and future states. Therefore stimuli with a common fate can transitively associate one another via the ‘nodes’ representing the same future prediction. The advantage of this coding scheme is a ‘compression’ of information such that multiple stimuli that lead to a common outcome can be coded by the same hippocampal cells.

While Gluck and Myers do not model how network activity persists from one moment to the next, other models have proposed explicit solutions. For example, the persistent activity observed in the entorhinal cortex (Egorov et al., 2002; Fransen et al., 2006) or the hippocampus (Knauer et al., 2013) may constitute a ‘buffer’ that maintains representations from the past that can be associated with representations of the current state (Koene and Hasselmo, 2008; Jensen and Lisman, 2005). Others have used recurrent connectivity in CA3 to allow previous states to influence future ones (Levy, 1996; Wallenstein and Hasselmo, 1997). Chaining one moment to the next allows for associative retrieval of entire sequences (Levy, 1996; Wallenstein and Hasselmo, 1997) and if this chaining occurs rapidly, rate estimates of a stimulus representation would include a code that partially reflects the future (Itskov et al., 2008). It is possible that my observation of cells that responded to multiple items or behaviors that similarly predict reward arose due predictions of a common future.

The temporal context model states that the current state is in part defined by a decaying trace of the past (Howard and Kahana, 2002). If the entire temporal context is bound together, then repeated presentations of an item can cue retrieval of its original temporal context through pattern completion. If the retrieved temporal context can also be used to cue the recurrent network, then items with a common history or future can be associated (Rao and Howard, 2008) providing another mechanism by which correlations between related stimuli could be established.

Several models of memory explicitly separate learning and retrieval phases to prevent recall of past associations during learning (Hasselmo, 1995; Hasselmo et al., 2002; Hopfield, 1984; Amit, 1988). There are two related problems that motivated this separation: runaway synaptic modification and higher-order interference (Hasselmo, 1994). Runaway synaptic modification occurs when multiple correlated patterns are stored within the same network. In this scenario, presentation of a new pattern causes retrieval of old related patterns. This retrieval causes association of units representing the new pattern but also units representing the old related pattern; synapses are modified between neurons representing the current state *and* past states. If this hybrid representation can also cue recall, then further activation of other related patterns is possible thus driving higher-order interference. Clearly some strategy is required to stop this associative expansion from consuming the entire brain. Several solutions have been proposed such as ensuring a constant total strength of synaptic weights to a single cell (Oja, 1989) or using lateral inhibition to limit the total number of active cells (Grossberg, 1976; Kohonen, 1993).

Associative networks where the influence of recurrent recall is minimized are often used to model formation of episodic memories within the hippocampus (e.g. Hasselmo 1994; Hasselmo et al., 1996). According to these models, episodic memories are formed by creating unique ‘snapshots’ of the constellation of the current stimuli divorced from the associative influence of recurrent networks (Treves and Rolls, 1994; O’Reilly, 2010; Hasselmo et al., 1996). Instead, I propose that the hippocampus stores a representation that is a composite of the current state and related past states. This description of hippocampal function is consistent with a class of models in which information is stored within self-organizing networks (Grossberg and Merrill, 1996).

In contrast to associative networks that only bind feedforward representations of present stimuli, self-organizing networks develop units that respond to categories of related inputs (Kohonen, 1984). If categorical representations are stored within the hippocampus, as they appear to be (e.g. Quiroga et al., 2005; Chapter 4), then why does damage to hippocampal function cause such severe episodic memory deficits? I propose that episodic memory reflects the unique association of a diverse set of categorical representations. If episodic memory is conceptualized as a movie, then the movie is populated by categories of information encountered at the time of learning rather than imprints of specific states of neocortical activity. Reactivating the unique set of categorical representations that were active during learning corresponds to the conscious experience of episodic memory.

As pointed out by others (Cohen and Eichenbaum, 1993; Eichenbaum, 1999) place cells represent the category of being within a particular location, which may be a particularly important category for demarcating one episodic memory from the next. Cohen and Eichenbaum (1993) suggested that episodic memory is encoded by a subset of cells being driven by complex conjunctions of stimuli and internal states whose unique combination defines a unique moment in life. I take a different approach and suggest instead that the episode is defined by the population of active cells, none of which is uniquely tuned and all of which respond to different sets of related input patterns. According to this view, ensembles pattern separate similar stimuli during some learning tasks because in those studies reward depended upon discriminating previously neutral stimuli (e.g. Komorowski et al., 2009). The increased pattern separation may reflect the emergence of a new category of information rather than an episodic memory code. For example, no two learning trials are identical, so a subject must learn the necessary and sufficient features of a trial that constitute a category of circumstances that predict a reliable outcome. Category boundaries may be supported by cells that discriminate the dimensions of an experience that differentially predict reward, or other commonalities that have significance. I have shown that when animals learn that different stimuli predict a common outcome, cells show a pattern completion function instead (Chapter 4), which also could reflect the category of circumstances with common meaning.

In some situations, category learning involves hippocampal processing (e.g. Kumaran et al., 2009; Zeithamova et al., 2008), while in others (Knowlton and Squire,

2000) category boundaries can be formed in the absence of hippocampal function. The boundary conditions that define which categories depend on hippocampal function are poorly understood. The hippocampal ability to reactivate and bind synaptically distant brain regions could allow for the discovery of relationships between stimuli that otherwise would not have been possible. The hippocampus sits atop the 'dorsal' and 'ventral' visual processing streams (Felleman and Van Essen, 1991) and these inputs are integrated along with monosynaptic and disynaptic inputs from the prefrontal cortex (Ohara et al., 2013), amygdala (Pitkänen et al., 2000), thalamus (Aggleton and Brown, 1999), and topological projections to the hypothalamus amongst others (Wyss et al., 1979; Witter et al., 2000a,b). The role of the hippocampus as a hub connecting these diverse regions has long been recognized (Teyler and DiScenna, 1986; Battaglia et al., 2011) and through this capacity it is thought that stimuli of different modalities can be associated. Indeed the combination of the recurrent collaterals and its action as a hub suggests that the hippocampus allows for an associative expansion of a memory trace to include the stimuli present as well as the states associated with those stimuli. I predict hippocampal involvement in category decisions that depend upon conjunctions of properties of stimulus that are represented in brain regions that do not have strong direct connections, such as combinations of color and spatial location.

Network hubs may also have another important role in defining the dimensionality in which a memory 'lives' and the role of hubs on memory interference has barely been explored, if at all. The role of hubs in small world networks is to allow short paths between different nodes of the network without requiring full network connectivity (Watts and Strogatz, 1998; Barabasi and Albert, 1999). In the brain, this could be particularly important since neural activity that occurs within a short period of time can be integrated and associated via spike-timing dependent plasticity (Markram et al., 1997). Therefore, the hippocampal role as a hub may increase the amount of connected brain – the dimensionality of the memory space – that can be mutually influential within critically brief time windows. Computational work has shown that Hopfield’s fully connected graphs have the theoretical maximal memory capacity (Treves and Rolls, 1994), but the brain is not fully connected. Small world networks with hubs fair better than other sparsely connected networks both in terms of the stability of established attractors when tested with a training pattern and the ability for corrupted training patterns to settle on learned inputs (McGraw and Menzinger, 2003). Formal modeling efforts are required to determine whether patterns stored in network with a hub are more resilient to interference than a patterns stored in the same set of neurons that communicate solely via local connectivity. I hypothesize that the small world created by a hippocampal hub allows for episodic memories that are less subject to ‘catastrophic interference’ by virtue of its role in increasing the functional dimensionality of the memory space.

Finally, the issue of interference and network capacity has largely focused upon static signals in the brain presumably after reaching equilibrium – point attractors in neural state space (e.g. Treves and Rolls, 1994). However, recurrent connections offer the opportunity for dynamic attractors of sequences of activity (Aihara and Adachi, 1997). Sequential activity is observed within theta periods (Dragoi and Buzsaki, 2006; Skaggs et al 1996; Gupta et al., 2012; see also Johnson and Redish, 2007) and during sharp-wave ripples (Wilson and McNaughton 1994; Davidson et al., 2009), suggesting that sequences may be a fundamental computational unit of the hippocampus. Furthermore, hippocampal networks can differentially fire in response to common inputs active in different orders (Hyde and Strowbridge, 2012; but see Liu et al., 2013 for evidence of a population code). Therefore, the issue of network capacity and interference highly depends upon how downstream 'readers' integrate and differentially respond to ordered sequences of the same and different neurons (Buzsaki, 2010). Current estimates of the number of stable attractors (memories) that can exist within the hippocampal network must be far too low since extant analyses have, for the sake of tractability, ignored the temporal dimension. Additional consideration of the permutations of network patterns over time vastly increases the dimensionality of the memory space. I observed snapshots of hippocampal activity whose correlations reflected meaningful relationships between experimental conditions. It is likely that each condition and potentially each moment was encoded by unique sequences of activity, the storage of which may not suffer from the catastrophic interference that has motivated most models of hippocampal function.

5.3 Conclusions

Hippocampal research has focused on how to build a system that stores isolated, unique episodic memories. Given the growing literature on the important role of the hippocampus in integrating information into a schema of related world knowledge, the time has come to reconcile these two potentially conflicting computational goals. I found that hippocampal ensembles capture the higher-order structure of a set of experiences and that neural correlates of well-learned memories show plasticity during learning. These data suggest that new information is assimilated via the accommodation of cortico-hippocampal networks that store related information. How the hippocampus avoids the trade-off between encoding new information and long term memory stability is not clear. The hippocampus may solve this 'plasticity-stability dilemma' by increasing the dimensionality of the memory space by binding synaptically distant brain regions or by increasing the dimensionality in the temporal domain with neuronal sequences as the fundamental representational unit. I propose that relational memory is particularly dependent upon the hippocampus because this region is a plastic, recurrent hub that changes quickly with experience and allows associative reactivation across modalities.

BIBLIOGRAPHY

Addis, D.R., Moscovitch, M., Crawley, A.P., and McAndrews, M.P. (2004). Recollective qualities modulate hippocampal activation during autobiographical memory retrieval. Hippocampus *14*, 752–762.

Addis, D.R., Wong, A.T., and Schacter, D.L. (2007). Remembering the past and imagining the future: common and distinct neural substrates during event construction and elaboration. Neuropsychologia *45*, 1363–1377.

Aggleton, J.P., and Brown, M.W. (1999). Episodic memory, amnesia, and the hippocampal-anterior thalamic axis. The Behavioral and Brain Sciences *22*, 425–489.

Aggleton, J.P., Vann, S.D., Denby, C., Dix, S., Mayes, A.R., Roberts, N., and Yonelinas, A.P. (2005). Sparing of the familiarity component of recognition memory in a patient with hippocampal pathology. Neuropsychologia *43*, 1810–1823.

Aihara, K., and Adachi, M. (1997). Associative Dynamics in a Chaotic Neural Network. Neural Networks *10*, 83–98.

Alberini, C.M. (2011). The role of reconsolidation and the dynamic process of long-term memory formation and storage. Frontiers in Behavioral Neuroscience *5*.

Alvarez, P., and Eichenbaum, H. (2002). Representations of odors in the rat orbitofrontal cortex change during and after learning. Behavioral Neuroscience *116*, 421–433.

Amit, D.J. (1988). Modeling brain function: The world of attractor neural networks. (Cambridge, UK: Cambridge University Press).

Anderson, J.A. (1972). A simple neural network generating an interactive memory. Mathematical Biosciences *14*, 197–220.

Anderson, M.I., and Jeffery, K.J. (2003). Heterogeneous modulation of place cell firing by changes in context. The Journal of Neuroscience *23*, 8827–8835.

Anderson, J.A., Silverstein, J.W., Ritz, S.A., and Jones, R.S. (1977). Distinctive features, categorical perception, and probability learning: some applications of a neural model. Psychological Review *84*, 413.

Assini, F.L., Duzzioni, M., and Takahashi, R.N. (2009). Object location memory in mice: pharmacological validation and further evidence of hippocampal CA1 participation. Behavioural Brain Research *204*, 206–211.

Bachevalier, J., and Nemanic, S. (2008). Memory for spatial location and object-place associations are differently processed by the hippocampal formation, parahippocampal areas TH/TF and perirhinal cortex. Hippocampus *18*, 64–80.

Bahar, A.S., and Shapiro, M.L. (2012). Remembering to learn: independent place and journey coding mechanisms contribute to memory transfer. The Journal of Neuroscience *32*, 2191–2203.

Bahar, A.S., Shirvalkar, P.R., and Shapiro, M.L. (2011). Memory-guided learning: CA1 and CA3 neuronal ensembles differentially encode the commonalities and differences between situations. The Journal of Neuroscience *31*, 12270–12281.

Bakker, A., Kirwan, C.B., Miller, M., and Stark, C.E.L. (2008). Pattern separation in the human hippocampal CA3 and dentate gyrus. Science *319*, 1640–1642.

Balderas, I., Rodriguez-Ortiz, C.J., Salgado-Tonda, P., Chavez-Hurtado, J., McGaugh, J.L., and Bermudez-Rattoni, F. (2008). The consolidation of object and context recognition memory involve different regions of the temporal lobe. Learning & Memory *15*, 618–624.

Barabasi, A., and Albert, R. (1999). Emergence of scaling in random networks. Science *286*, 509–512.

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

Bartsch, T., Alfke, K., Stingele, R., Rohr, A., Freitag-Wolf, S., Jansen, O., and Deuschl, G. (2006). Selective affection of hippocampal CA-1 neurons in patients with transient global amnesia without long-term sequelae. Brain *129*, 2874–2884.

Battaglia, F.P., Benchenane, K., Sirota, A., Pennartz, C.M.A., and Wiener, S.I. (2011). The hippocampus: hub of brain network communication for memory. Trends in Cognitive Sciences *15*, 310–318.

Baxter, M.G., and Murray, E.A. (2001). Opposite relationship of hippocampal and rhinal cortex damage to delayed nonmatching-to-sample deficits in monkeys. Hippocampus *11*, 61–71.

Bayley, P.J., Hopkins, R.O., and Squire, L.R. (2003). Successful recollection of remote autobiographical memories by amnesic patients with medial temporal lobe lesions. Neuron *38*, 135–144.

Bayley, P.J., Gold, J.J., Hopkins, R.O., and Squire, L.R. (2005). The neuroanatomy of remote memory. Neuron *46*, 799–810.

Bayley, P.J., Hopkins, R.O., and Squire, L.R. (2006). The fate of old memories after medial temporal lobe damage. The Journal of Neuroscience *26*, 13311–13317.

Beason-Held, L.L., Rosene, D.L., Killiany, R.J., and Moss, M.B. (1999). Hippocampal formation lesions produce memory impairment in the rhesus monkey. Hippocampus *9*, 562–574.

von Bechterew, W. (1900). Demonstration eines gehirns mit zerstorung der vorderen und inneren theile der hirnrinde beider schlafenlappen. Neurologisch. Zentralbl. *19*, 990–991.

Berman, D.E., Hazvi, S., Stehberg, J., Bahar, A., and Dudai, Y. (2003). Conflicting processes in the extinction of conditioned taste aversion: behavioral and molecular aspects of latency, apparent stagnation, and spontaneous recovery. Learning & Memory *10*, 16–25.

Biedenkapp, J.C., and Rudy, J.W. (2004). Context memories and reactivation: constraints on the reconsolidation hypothesis. Behavioral Neuroscience *118*, 956–964.

Biedenkapp, J.C., and Rudy, J.W. (2007). Context preexposure prevents forgetting of a contextual fear memory: implication for regional changes in brain activation patterns associated with recent and remote memory tests. Learning & Memory *14*, 200–203.

Bliss, T. V, and Collingridge, G.L. (1993). A synaptic model of memory: long-term potentiation in the hippocampus. Nature *361*, 31–39.

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

Bonnevie, T., Dunn, B., Fyhn, M., Hafting, T., Derdikman, D., Kubie, J.L., Roudi, Y., Moser, E.I., and Moser, M.-B. (2013). Grid cells require excitatory drive from the hippocampus. Nature Neuroscience *16*, 309–317.

Bonnici, H.M., Chadwick, M.J., Lutti, A., Hassabis, D., Weiskopf, N., and Maguire, E.A. (2012). Detecting representations of recent and remote autobiographical memories in vmPFC and hippocampus. The Journal of Neuroscience *32*, 16982–16991.

Bontempi, B., Laurent-Demir, C., Destrade, C., and Jaffard, R. (1999). Time-dependent reorganization of brain circuitry underlying long-term memory storage. Nature *400*, 671–675.

Bostock, E., Muller, R.U., and Kubie, J.L. (1991). Experience-dependent modifications of hippocampal place cell firing. Hippocampus *1*, 193–205.

Bouton, M.E. (1993). Context, time, and memory retrieval in the interference paradigms of Pavlovian learning. Psychological Bulletin *114*, 80–99.

Bozon, B., Davis, S., and Laroche, S. (2003). A requirement for the immediate early gene zif268 in reconsolidation of recognition memory after retrieval. Neuron *40*, 695–701.

Brandon, M.P., Bogaard, A.R., Libby, C.P., Connerney, M.A., Gupta, K., and Hasselmo, M.E. (2011). Reduction of theta rhythm dissociates grid cell spatial periodicity from directional tuning. Science *332*, 595–599.

Brewer, J.B., Zhao, Z., Desmond, J.E., Glover, G.H., and Gabrieli, J.D. (1998). Making memories: brain activity that predicts how well visual experience will be remembered. Science *281*, 1185–1187.

Bright, P., Buckman, J., Fradera, A., Yoshimasu, H., Colchester, A.C.F., and Kopelman, M.D. (2006). Retrograde amnesia in patients with hippocampal, medial temporal, temporal lobe, or frontal pathology. Learning & Memory *13*, 545–557.

Broadbent, N.J., and Clark, R.E. (2013). Remote context fear conditioning remains hippocampus-dependent irrespective of training protocol, training-surgery interval, lesion size, and lesion method. Neurobiology of Learning and Memory *106*, 300–308.

Brown, J.E., and Skaggs, W.E. (2002). Concordant and discordant coding of spatial location in populations of hippocampal CA1 pyramidal cells. Journal of Neurophysiology *88*, 1605–1613.

Brown, T.I., and Stern, C.E. (2013). Contributions of Medial Temporal Lobe and Striatal Memory Systems to Learning and Retrieving Overlapping Spatial Memories. Cerebral Cortex.

Brown, T.I., Ross, R.S., Keller, J.B., Hasselmo, M.E., and Stern, C.E. (2010). Which way was I going? Contextual retrieval supports the disambiguation of well learned overlapping navigational routes. The Journal of Neuroscience *30*, 7414–7422.

Brun, V.H., Otnass, M.K., Molden, S., Steffenach, H.-A., Witter, M.P., Moser, M.-B., and Moser, E.I. (2002). Place cells and place recognition maintained by direct entorhinal-hippocampal circuitry. Science *296*, 2243–2246.

Brun, V.H., Leutgeb, S., Wu, H.-Q., Schwarcz, R., Witter, M.P., Moser, E.I., and Moser, M.-B. (2008). Impaired spatial representation in CA1 after lesion of direct input from entorhinal cortex. Neuron *57*, 290–302.

Bunsey, M., and Eichenbaum, H. (1996). Conservation of hippocampal memory function in rats and humans. Nature *379*, 255–257.

Burnham, W.. (1903). Retroactive amnesia: Illustrative cases and a tentative explanation. American Journal of Psychology *14*, 382–396.

Buzsáki, G., Leung, L., and Vanderwolf, C. (1983). Cellular bases of hippocampal EEG in the behaving rat. Brain Research *287*, 139–171.

Buzsáki, G. (1986). Hippocampal sharp waves: their origin and significance. Brain Research *398*, 242–252.

Buzsáki, G. (2010). Neural syntax: cell assemblies, synapsembles, and readers. Neuron *68*, 362–385.

Calton, J.L., Stackman, R.W., Goodridge, J.P., Archey, W.B., Dudchenko, P.A., and Taube, J.S. (2003). Hippocampal place cell instability after lesions of the head direction cell network. The Journal of Neuroscience *23*, 9719–9731.

Cameron, K.A., Yashar, S., Wilson, C.L., and Fried, I. (2001). Human hippocampal neurons predict how well word pairs will be remembered. Neuron *30*, 289–298.

Carpenter, G.A. (1997). Distributed Learning, Recognition, and Prediction by ART and ARTMAP Neural Networks. Neural Networks *10*, 1473–1494.

Carr, M.F., Jadhav, S.P., and Frank, L.M. (2011). Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. Nature Neuroscience *14*, 147–153.

van Cauter, T., Poucet, B., and Save, E. (2008). Unstable CA1 place cell representation in rats with entorhinal cortex lesions. The European Journal of Neuroscience *27*, 1933–1946.

Chadwick, M.J., Hassabis, D., Weiskopf, N., and Maguire, E.A. (2010). Decoding individual episodic memory traces in the human hippocampus. Current Biology *20*, 544–547.

Chase, W.G., and Simon, H.A. (1973). Perception in chess. Cognitive Psychology *4*, 55–81.

Cheke, L.G., and Clayton, N.S. (2010). Mental time travel in animals. Wiley Interdisciplinary Reviews: Cognitive Science *1*, 915–930.

Chrobak, J.J., Lörincz, A., and Buzsáki, G. (2000). Physiological patterns in the hippocampo-entorhinal cortex system. Hippocampus *10*, 457–465.

Cipolotti, L., Shallice, T., Chan, D., Fox, N., Scahill, R., Harrison, G., Stevens, J., and Rudge, P. (2001). Long-term retrograde amnesia...the crucial role of the hippocampus. Neuropsychologia *39*, 151–172.

Clark, R.E., and Squire, L.R. (2010). An animal model of recognition memory and medial temporal lobe amnesia: history and current issues. Neuropsychologia *48*, 2234–2244.

Clark, R.E., and Squire, L.R. (2013). Similarity in form and function of the hippocampus in rodents, monkeys, and humans. Proceedings of the National Academy of Sciences of the United States of America *110*, 10365–10370.

Clark, R.E., Broadbent, N.J., and Squire, L.R. (2005). Impaired remote spatial memory after hippocampal lesions despite extensive training beginning early in life. Hippocampus *15*, 340–346.

Clayton, N.S., and Dickinson, A. (1998). Episodic-like memory during cache recovery by scrub jays. Nature *395*, 272–274.

Cohen, N.J., and Eichenbaum, H.B. (1993). Memory, amnesia, and the hippocampal system. (Cambridge, MA: The MIT Press).

Corcoran, K.A., and Maren, S. (2001). Hippocampal inactivation disrupts contextual retrieval of fear memory after extinction. The Journal of Neuroscience *21*, 1720–1726.

Corkin, S. (1965). Tactually-guided maze learning in man: Effects of unilateral cortical excisions and bilateral hippocampal lesions. Neuropsychologia *3*, 339–351.

Corkin, S. (1984). Lasting Consequences of Bilateral Medial Temporal Lobectomy: Clinical Course and Experimental Findings in H.M. Seminars in Neurology *4*, 249–259.

Corkin, S. (2002). What’s new with the amnesic patient H.M.? Nature Reviews. Neuroscience *3*, 153–160.

Coutureau, E., Killcross, A.S., Good, M., Marshall, V.J., Ward-Robinson, J., and Honey, R.C. (2002). Acquired equivalence and distinctiveness of cues: II. Neural manipulations and their implications. Journal of Experimental Psychology. Animal Behavior Processes *28*, 388–396.

Csicsvari, J., Hirase, H., Czurkó, A., Mamiya, A., and Buzsáki, G. (1999). Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving Rat. The Journal of Neuroscience *19*, 274–287.

Damasio, A.R. (1989). Time-locked multiregional retroactivation: a systems-level proposal for the neural substrates of recall and recognition. Cognition *33*, 25–62.

Davachi, L., and Wagner, A.D. (2002). Hippocampal contributions to episodic encoding: insights from relational and item-based learning. Journal of Neurophysiology *88*, 982–990.

Davachi, L., Mitchell, J.P., and Wagner, A.D. (2003). Multiple routes to memory: distinct medial temporal lobe processes build item and source memories. Proceedings of the National Academy of Sciences of the United States of America *100*, 2157–2162.

Davidson, T.J., Kloosterman, F., and Wilson, M.A. (2009). Hippocampal replay of extended experience. Neuron *63*, 497–507.

Day, M., Langston, R., and Morris, R.G.M. (2003). Glutamate-receptor-mediated encoding and retrieval of paired-associate learning. Nature *424*, 205–209.

Debiec, J., LeDoux, J.E., and Nader, K. (2002). Cellular and systems reconsolidation in the hippocampus. Neuron *36*, 527–538.

Debiec, J., Doyère, V., Nader, K., and Ledoux, J.E. (2006). Directly reactivated, but not indirectly reactivated, memories undergo reconsolidation in the amygdala. Proceedings of the National Academy of Sciences of the United States of America *103*, 3428–3433.

Dede, A.J.O., Wixted, J.T., Hopkins, R.O., and Squire, L.R. (2013). Hippocampal damage impairs recognition memory broadly, affecting both parameters in two prominent models of memory. Proceedings of the National Academy of Sciences of the United States of America *110*, 6577–6582.

Demb, J.B., Desmond, J.E., Wagner, A.D., Vaidya, C.J., Glover, G.H., and Gabrieli, J.D. (1995). Semantic encoding and retrieval in the left inferior prefrontal cortex: a functional MRI study of task difficulty and process specificity. The Journal of Neuroscience *15*, 5870–5878.

Dere, E., Kart-Teke, E., Huston, J.P., and De Souza Silva, M.A. (2006). The case for episodic memory in animals. Neuroscience and Biobehavioral Reviews *30*, 1206–1224.

Deshmukh, S.S., and Knierim, J.J. (2013). Influence of local objects on hippocampal representations: Landmark vectors and memory. Hippocampus *23*, 253–267.

DeVito, L.M., and Eichenbaum, H. (2010a). Distinct contributions of the hippocampus and medial prefrontal cortex to the “what-where-when” components of episodic-like memory in mice. Behavioural Brain Research *215*, 318–325.

DeVito, L.M., Kanter, B.R., and Eichenbaum, H. (2010b). The hippocampus contributes to memory expression during transitive inference in mice. Hippocampus *20*, 208–217.

Diana, R.A., Yonelinas, A.P., and Ranganath, C. (2007). Imaging recollection and familiarity in the medial temporal lobe: a three-component model. Trends in Cognitive Sciences *11*, 379–386.

Diana, R.A., Yonelinas, A.P., and Ranganath, C. (2008). The effects of unitization on familiarity-based source memory: testing a behavioral prediction derived from neuroimaging data. Journal of Experimental Psychology. Learning, Memory, and Cognition *34*, 730–740.

Diana, R.A., Yonelinas, A.P., and Ranganath, C. (2010). Medial temporal lobe activity during source retrieval reflects information type, not memory strength. Journal of Cognitive Neuroscience *22*, 1808–1818.

Douville, K., Woodard, J.L., Seidenberg, M., Miller, S.K., Leveroni, C.L., Nielson, K.A., Franczak, M., Antuono, P., and Rao, S.M. (2005). Medial temporal lobe activity for recognition of recent and remote famous names: an event-related fMRI study. Neuropsychologia *43*, 693–703.

Drachman, D.A., and Arbit, J. (1966). Memory and the hippocampal complex. II. Is memory a multiple process? Archives of Neurology *15*, 52–61.

Dragoi, G., and Buzsáki, G. (2006). Temporal encoding of place sequences by hippocampal cell assemblies. Neuron *50*, 145–157.

Dragoi, G., and Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. Nature *469*, 397–401.

Dragoi, G., and Tonegawa, S. (2013). Development of schemas revealed by prior experience and NMDA receptor knock-out. eLife *2*, e01326.

Dudai, Y., and Eisenberg, M. (2004). Rites of passage of the engram: reconsolidation and the lingering consolidation hypothesis. Neuron *44*, 93–100.

Dupret, D., O'Neill, J., Pleydell-Bouverie, B., and Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. Nature Neuroscience *13*, 995–1002.

Dupret, D., O'Neill, J., and Csicsvari, J. (2013). Dynamic reconfiguration of hippocampal interneuron circuits during spatial learning. Neuron *78*, 166–180.

Dusek, J.A., and Eichenbaum, H. (1997). The hippocampus and memory for orderly stimulus relations. Proceedings of the National Academy of Sciences of the United States of America *94*, 7109–7114.

Duvarci, S., and Nader, K. (2004). Characterization of fear memory reconsolidation. The Journal of Neuroscience *24*, 9269–9275.

Duvarci, S., Mamou, C. Ben, and Nader, K. (2006). Extinction is not a sufficient condition to prevent fear memories from undergoing reconsolidation in the basolateral amygdala. The European Journal of Neuroscience *24*, 249–260.

Eacott, M.J., and Norman, G. (2004). Integrated memory for object, place, and context in rats: a possible model of episodic-like memory? The Journal of Neuroscience *24*, 1948–1953.

Egorov, A. V, Hamam, B.N., Fransén, E., Hasselmo, M.E., and Alonso, A.A. (2002). Graded persistent activity in entorhinal cortex neurons. Nature *420*, 173–178.

Ego-Stengel, V., and Wilson, M.A. (2010). Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. Hippocampus *20*, 1–10.

Eich, E. (1985). Context, memory, and integrated item/context imagery. Journal of Experimental Psychology: Learning, Memory, and Cognition *4*, 764–770.

Eichenbaum, H. (1999). The hippocampus and mechanisms of declarative memory. Behavioural Brain Research *103*, 123–133.

Eichenbaum, H. (2004). Hippocampus: cognitive processes and neural representations that underlie declarative memory. Neuron *44*, 109–120.

Eichenbaum, H. (2006). The secret life of memories. Neuron *50*, 350–352.

Eichenbaum, H., Stewart, C., and Morris, R.G. (1990). Hippocampal representation in place learning. The Journal of Neuroscience *10*, 3531–3542.

Eichenbaum, H., Otto, T., and Cohen, N.J. (1994). Two functional components of the hippocampal memory system. Behavioral and Brain Sciences *17*, 449–517.

Eichenbaum, H., Yonelinas, A.P., and Ranganath, C. (2007). The medial temporal lobe and recognition memory. Annual Review of Neuroscience *30*, 123–152.

Eisenberg, M., and Dudai, Y. (2004). Reconsolidation of fresh, remote, and extinguished fear memory in Medaka: old fears don't die. The European Journal of Neuroscience *20*, 3397–3403.

Eisenberg, M., Kobilo, T., Berman, D.E., and Dudai, Y. (2003). Stability of retrieved memory: inverse correlation with trace dominance. Science *301*, 1102–1104.

Eldridge, L.L., Knowlton, B.J., Furmanski, C.S., Bookheimer, S.Y., and Engel, S.A. (2000). Remembering episodes: a selective role for the hippocampus during retrieval. Nature Neuroscience *3*, 1149–1152.

Eldridge, L.L., Engel, S.A., Zeineh, M.M., Bookheimer, S.Y., and Knowlton, B.J. (2005). A dissociation of encoding and retrieval processes in the human hippocampus. The Journal of Neuroscience *25*, 3280–3286.

Ellenbogen, J.M., Hu, P.T., Payne, J.D., Titone, D., and Walker, M.P. (2007). Human relational memory requires time and sleep. Proceedings of the National Academy of Sciences of the United States of America *104*, 7723–7728.

Ennaceur, A., Neave, N., and Aggleton, J.P. (1997). Spontaneous object recognition and object location memory in rats: the effects of lesions in the cingulate cortices, the medial prefrontal cortex, the cingulum bundle and the fornix. Experimental Brain Research *113*, 509–519.

Ergorul, C., and Eichenbaum, H. (2004). The hippocampus and memory for "what," "where," and "when". Learning & Memory *11*, 397–405.

Eschenko, O., and Mizumori, S.J.Y. (2007). Memory influences on hippocampal and striatal neural codes: effects of a shift between task rules. Neurobiology of Learning and Memory *87*, 495–509.

Etchamendy, N., Desmedt, A., Cortes-Torrea, C., Marighetto, A., and Jaffard, R. (2003). Hippocampal lesions and discrimination performance of mice in the radial maze: sparing or impairment depending on the representational demands of the task. Hippocampus *13*, 197–211.

Fagan, J.F. (1970). Memory in the infant. Journal of Experimental Child Psychology *9*, 217–226.

Falls, W.A., Miserendino, M.J., and Davis, M. (1992). Extinction of fear-potentiated startle: blockade by infusion of an NMDA antagonist into the amygdala. The Journal of Neuroscience *12*, 854–863.

Felleman, D.J., and Van Essen, D.C. (1991). Distributed hierarchical processing in the primate cerebral cortex. Cerebral Cortex *1*, 1–47.

Fenton, A., Lytton, W.W., Barry, J.M., Lenck-Santini, P.-P., Zinyuk, L.E., Kubík, S., Bures, J., Poucet, B., Muller, R.U., and Olypher, A. V (2010). Attention-like modulation of hippocampus place cell discharge. The Journal of Neuroscience *30*, 4613–4625.

Ferbinteanu, J., and Shapiro, M.L. (2003). Prospective and retrospective memory coding in the hippocampus. Neuron *40*, 1227–1239.

Fernández, G., Weyerts, H., Schrader-Bölsche, M., Tendolkar, I., Smid, H.G., Tempelmann, C., Hinrichs, H., Scheich, H., Elger, C.E., Mangun, G.R., et al. (1998). Successful verbal encoding into episodic memory engages the posterior hippocampus: a parametrically analyzed functional magnetic resonance imaging study. The Journal of Neuroscience *18*, 1841–1847.

Forcato, C., Argibay, P.F., Pedreira, M.E., and Maldonado, H. (2009). Human reconsolidation does not always occur when a memory is retrieved: the relevance of the reminder structure. Neurobiology of Learning and Memory *91*, 50–57.

Fortin, N.J., Agster, K.L., and Eichenbaum, H.B. (2002). Critical role of the hippocampus in memory for sequences of events. Nature Neuroscience *5*, 458–462.

Frank, L.M., Brown, E.N., and Wilson, M. (2000). Trajectory encoding in the hippocampus and entorhinal cortex. Neuron *27*, 169–178.

Frankland, P.W., Cestari, V., Filipkowski, R.K., McDonald, R.J., and Silva, A.J. (1998). The dorsal hippocampus is essential for context discrimination but not for contextual conditioning. Behavioral Neuroscience *112*, 863–874.

Frankland, P.W., Bontempi, B., Talton, L.E., Kaczmarek, L., and Silva, A.J. (2004). The involvement of the anterior cingulate cortex in remote contextual fear memory. Science *304*, 881–883.

Frankland, P.W., Ding, H.-K., Takahashi, E., Suzuki, A., Kida, S., and Silva, A.J. (2006). Stability of recent and remote contextual fear memory. Learning & Memory *13*, 451–457.

Fransén, E., Tahvildari, B., Egorov, A. V, Hasselmo, M.E., and Alonso, A.A. (2006). Mechanism of graded persistent cellular activity of entorhinal cortex layer v neurons. Neuron *49*, 735–746.

Frean, M., and Robins, A. (1999). Catastrophic forgetting in simple networks: an analysis of the pseudorehearsal solution. Network *10*, 227–236.

French, R.M. (1992). Semi-distributed representations and catastrophic forgetting in connectionist networks. Connection Science *4*, 365–377.

French, R.M. (1999). Catastrophic forgetting in connectionist networks. Trends in Cognitive Sciences *3*, 128–135.

Fried, I., MacDonald, K.A., and Wilson, C.L. (1997). Single neuron activity in human hippocampus and amygdala during recognition of faces and objects. Neuron *18*, 753–765.

Fyhn, M., Molden, S., Hollup, S., Moser, M.-B., and Moser, E. (2002). Hippocampal neurons responding to first-time dislocation of a target object. Neuron *35*, 555–566.

Fyhn, M., Molden, S., Witter, M.P., Moser, E.I., and Moser, M.-B. (2004). Spatial representation in the entorhinal cortex. Science *305*, 1258–1264.

Gabrieli, J.D., Milberg, W., Keane, M.M., and Corkin, S. (1990). Intact priming of patterns despite impaired memory. Neuropsychologia *28*, 417–427.

Gaffan, D. (1974). Recognition impaired and association intact in the memory of monkeys after transection of the fornix. Journal of Comparative and Physiological Psychology *86*, 1100–1109.

Gaffan, D. (1994). Scene-Specific Memory for Objects: A Model of Episodic Memory Impairment in Monkeys with Fornix Transection. Journal of Cognitive Neuroscience *6*, 305–320.

Geisler, C., Robbe, D., Zugaro, M., Sirota, A., and Buzsáki, G. (2007). Hippocampal place cell assemblies are speed-controlled oscillators. Proceedings of the National Academy of Sciences of the United States of America *104*, 8149–8154.

Gelbard-Sagiv, H., Mukamel, R., Harel, M., Malach, R., and Fried, I. (2008). Internally generated reactivation of single neurons in human hippocampus during free recall. Science *322*, 96–101.

Gilbert, P.E., Kesner, R.P., and DeCoteau, W.E. (1998). Memory for spatial location: role of the hippocampus in mediating spatial pattern separation. The Journal of Neuroscience *18*, 804–810.

Gilboa, A., Winocur, G., Grady, C.L., Hevenor, S.J., and Moscovitch, M. (2004). Remembering our past: functional neuroanatomy of recollection of recent and very remote personal events. Cerebral Cortex *14*, 1214–1225.

Gilboa, A., Winocur, G., Rosenbaum, R.S., Poreh, A., Gao, F., Black, S.E., Westmacott, R., and Moscovitch, M. (2006). Hippocampal contributions to recollection in retrograde and anterograde amnesia. Hippocampus *16*, 966–980.

Girardeau, G., Benchenane, K., Wiener, S.I., Buzsáki, G., and Zugaro, M.B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. Nature Neuroscience *12*, 1222–1223.

Glees, P., and Griffith, H.B. (1952). Bilateral destruction of the hippocampus (cornu ammonis) in a case of dementia. Monatsschrift Für Psychiatrie Und Neurologie *123*, 193–204.

Gluck, M.A., and Myers, C.E. (1993). Hippocampal mediation of stimulus representation: a computational theory. Hippocampus *3*, 491–516.

Gonsalves, B.D., Kahn, I., Curran, T., Norman, K.A., and Wagner, A.D. (2005). Memory strength and repetition suppression: multimodal imaging of medial temporal cortical contributions to recognition. Neuron *47*, 751–761.

Good, M., De Hoz, L., and Morris, R.G. (1998). Contingent versus incidental context processing during conditioning: dissociation after excitotoxic hippocampal plus dentate gyrus lesions. Hippocampus *8*, 147–159.

Goshen, I., Brodsky, M., Prakash, R., Wallace, J., Gradinaru, V., Ramakrishnan, C., and Deisseroth, K. (2011). Dynamics of retrieval strategies for remote memories. Cell *147*, 678–689.

Gothard, K.M., Skaggs, W.E., Moore, K.M., and McNaughton, B.L. (1996a). Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. The Journal of Neuroscience *16*, 823–835.

Gothard, K.M., Skaggs, W.E., and McNaughton, B.L. (1996b). Dynamics of mismatch correction in the hippocampal ensemble code for space: interaction between path integration and environmental cues. The Journal of Neuroscience *16*, 8027–8040.

Greene, A.J., Gross, W.L., Elsinger, C.L., and Rao, S.M. (2006). An FMRI analysis of the human hippocampus: inference, context, and task awareness. Journal of Cognitive Neuroscience *18*, 1156–1173.

Grossberg, S. (1976). Adaptive pattern classification and universal recoding: I. Parallel development and coding of neural feature detectors. Biological Cybernetics *23*, 121–134.

Grossberg, S., and Merrill, J.W. (1996). The hippocampus and cerebellum in adaptively timed learning, recognition, and movement. Journal of Cognitive Neuroscience *8*, 257–277.

Guillery-Girard, B., Desgranges, B., Urban, C., Piolino, P., De la Sayette, V., and Eustache, F. (2004). The dynamic time course of memory recovery in transient global amnesia. Journal of Neurology, Neurosurgery, and Psychiatry *75*, 1532–1540.

Gupta, A.S., Van der Meer, M.A.A., Touretzky, D.S., and Redish, A.D. (2010). Hippocampal replay is not a simple function of experience. Neuron *65*, 695–705.

Gupta, A.S., Van der Meer, M.A.A., Touretzky, D.S., and Redish, A.D. (2012). Segmentation of spatial experience by hippocampal θ sequences. Nature Neuroscience *15*, 1032–1039.

Guzowski, J.F., McNaughton, B.L., Barnes, C.A., and Worley, P.F. (1999). Environment-specific expression of the immediate-early gene Arc in hippocampal neuronal ensembles. Nature Neuroscience *2*, 1120–1124.

Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., and Moser, E.I. (2005). Microstructure of a spatial map in the entorhinal cortex. Nature *436*, 801–806.

Haist, F., Bowden Gore, J., and Mao, H. (2001). Consolidation of human memory over decades revealed by functional magnetic resonance imaging. Nature Neuroscience *4*, 1139–1145.

Halgren, E., Wilson, C.L., and Stapleton, J.M. (1985). Human medial temporal-lobe stimulation disrupts both formation and retrieval of recent memories. Brain and Cognition *4*, 287–295.

Hampson, R.E., Pons, T.P., Stanford, T.R., and Deadwyler, S.A. (2004). Categorization in the monkey hippocampus: a possible mechanism for encoding information into memory. Proceedings of the National Academy of Sciences of the United States of America *101*, 3184–3189.

Harris, K.D., Csicsvari, J., Hirase, H., and Dragoi, G. (2003). Organization of cell assemblies in the hippocampus. Nature *424*, 552–556.

Harvey, C.D., Collman, F., Dombeck, D.A., and Tank, D.W. (2009). Intracellular dynamics of hippocampal place cells during virtual navigation. Nature *461*, 941–946.

Hassabis, D., Kumaran, D., and Maguire, E.A. (2007). Using imagination to understand the neural basis of episodic memory. The Journal of Neuroscience *27*, 14365–14374.

Hasselmo, M.E. (1994). Runaway synaptic modification in models of cortex: Implications for Alzheimer's disease. Neural Networks *7*, 13–40.

Hasselmo, M.E. (2009). A model of episodic memory: mental time travel along encoded trajectories using grid cells. Neurobiology of Learning and Memory *92*, 559–573.

Hasselmo, M.E., and Eichenbaum, H. (2005). Hippocampal mechanisms for the context-dependent retrieval of episodes. Neural Networks *18*, 1172–1190.

Hasselmo, M.E., and Wyble, B.P. (1997). Free recall and recognition in a network model of the hippocampus: simulating effects of scopolamine on human memory function. Behavioural Brain Research *89*, 1–34.

Hasselmo, M.E., Schnell, E., and Barkai, E. (1995). Dynamics of learning and recall at excitatory recurrent synapses and cholinergic modulation in rat hippocampal region CA3. The Journal of Neuroscience *15*, 5249–5262.

Hasselmo, M.E., Wyble, B.P., and Wallenstein, G. V (1996). Encoding and retrieval of episodic memories: role of cholinergic and GABAergic modulation in the hippocampus. Hippocampus *6*, 693–708.

Hasselmo, M.E., Bodelón, C., and Wyble, B.P. (2002). A proposed function for hippocampal theta rhythm: separate phases of encoding and retrieval enhance reversal of prior learning. Neural Computation *14*, 793–817.

Hayman, R.M.A., Chakraborty, S., Anderson, M.I., and Jeffery, K.J. (2003). Context-specific acquisition of location discrimination by hippocampal place cells. The European Journal of Neuroscience *18*, 2825–2834.

Hebb, D. (1949). The Organization of Behavior (New York: John Wiley & Sons).

Heckers, S., Zalesak, M., Weiss, A.P., Ditman, T., and Titone, D. (2004). Hippocampal activation during transitive inference in humans. Hippocampus *14*, 153–162.

Hetherington, P.A., and Shapiro, M.L. (1997). Hippocampal place fields are altered by the removal of single visual cues in a distance-dependent manner. Behavioral Neuroscience *111*, 20–34.

Hill, A.J., and Best, P.J. (1981). Effects of deafness and blindness on the spatial correlates of hippocampal unit activity in the rat. Experimental Neurology *74*, 204–217.

Hinton, G.E. (1989). Connectionist learning procedures. Artificial Intelligence *40*, 185–234.

Hirano, M., and Noguchi, K. (1998). Dissociation between specific personal episodes and other aspects of remote memory in a patient with hippocampal amnesia. Perceptual and Motor Skills *87*, 99–107.

Hok, V., Lenck-Santini, P.-P., Save, E., Gaussier, P., Banquet, J.-P., and Poucet, B. (2007a). A test of the time estimation hypothesis of place cell goal-related activity. Journal of Integrative Neuroscience *6*, 367–378.

Hok, V., Lenck-Santini, P.-P., Roux, S., Save, E., Muller, R.U., and Poucet, B. (2007b). Goal-related activity in hippocampal place cells. The Journal of Neuroscience *27*, 472–482.

Hok, V., Chah, E., Save, E., and Poucet, B. (2013). Prefrontal cortex focally modulates hippocampal place cell firing patterns. The Journal of Neuroscience *33*, 3443–3451.

Holdstock, J.S., Mayes, A.R., Roberts, N., Cezayirli, E., Isaac, C.L., O’Reilly, R.C., and Norman, K.A. (2002). Under what conditions is recognition spared relative to recall after selective hippocampal damage in humans? Hippocampus *12*, 341–351.

Hollup, S.A., Molden, S., Donnett, J.G., Moser, M.B., and Moser, E.I. (2001). Place fields of rat hippocampal pyramidal cells and spatial learning in the watermaze. The European Journal of Neuroscience *13*, 1197–1208.

Honey, R.C., and Watt, A. (1998). Acquired relational equivalence: implications for the nature of associative structures. Journal of Experimental Psychology. Animal Behavior Processes *24*, 325–334.

Hopfield, J.J. (1982). Neural networks and physical systems with emergent collective computational abilities. Proceedings of the National Academy of Sciences of the United States of America *79*, 2554–2558.

Hopfield, J.J. (1984). Neurons with graded response have collective computational properties like those of two-state neurons. Proceedings of the National Academy of Sciences of the United States of America *81*, 3088–3092.

Howard, M.W., and Kahana, M.J. (2002). A Distributed Representation of Temporal Context. Journal of Mathematical Psychology *46*, 269–299.

Howard, M.W., Viskontas, I. V, Shankar, K.H., and Fried, I. (2012). Ensembles of human MTL neurons “jump back in time” in response to a repeated stimulus. Hippocampus *22*, 1833–1847.

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

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

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

Hyde, R.A., and Strowbridge, B.W. (2012). Mnemonic representations of transient stimuli and temporal sequences in the rodent hippocampus in vitro. Nature Neuroscience *15*, 1430–1438.

Iordanova, M.D., Burnett, D.J., Aggleton, J.P., Good, M., and Honey, R.C. (2009). The role of the hippocampus in mnemonic integration and retrieval: complementary evidence from lesion and inactivation studies. The European Journal of Neuroscience *30*, 2177–2189.

Itskov, V., Pastalkova, E., Mizuseki, K., Buzsaki, G., and Harris, K.D. (2008). Theta-mediated dynamics of spatial information in hippocampus. The Journal of Neuroscience *28*, 5959–5964.

Jackson, J., and Redish, A.D. (2007). Network dynamics of hippocampal cell-assemblies resemble multiple spatial maps within single tasks. Hippocampus *17*, 1209–1229.

Jeffery, K.J., Gilbert, A., Burton, S., and Strudwick, A. (2003). Preserved performance in a hippocampal-dependent spatial task despite complete place cell remapping. Hippocampus *13*, 175–189.

Jensen, O., and Lisman, J.E. (2005). Hippocampal sequence-encoding driven by a cortical multi-item working memory buffer. Trends in Neurosciences *28*, 67–72.

van der Jeugd, A., Goddyn, H., Laeremans, A., Arckens, L., D’Hooge, R., and Verguts, T. (2009). Hippocampal involvement in the acquisition of relational associations, but not in the expression of a transitive inference task in mice. Behavioral Neuroscience *123*, 109–114.

Jezek, K., Henriksen, E.J., Treves, A., Moser, E.I., and Moser, M.-B. (2011). Theta-paced flickering between place-cell maps in the hippocampus. Nature *478*, 246–249.

Ji, D., and Wilson, M.A. (2007). Coordinated memory replay in the visual cortex and hippocampus during sleep. Nature Neuroscience *10*, 100–107.

Johnson, A., and Redish, A.D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. The Journal of Neuroscience *27*, 12176–12189.

Kapur, N., and Brooks, D.J. (1999). Temporally-specific retrograde amnesia in two cases of discrete bilateral hippocampal pathology. Hippocampus *9*, 247–254.

Kelemen, E., and Fenton, A.A. (2010). Dynamic grouping of hippocampal neural activity during cognitive control of two spatial frames. PLoS Biology *8*, e1000403.

Kentros, C.G., Agnihotri, N.T., Streater, S., Hawkins, R.D., and Kandel, E.R. (2004). Increased attention to spatial context increases both place field stability and spatial memory. Neuron *42*, 283–295.

Kesner, R.P., Gilbert, P.E., and Barua, L.A. (2002). The role of the hippocampus in memory for the temporal order of a sequence of odors. Behavioral Neuroscience *116*, 286–290.

van Kesteren, M.T.R., Fernández, G., Norris, D.G., and Hermans, E.J. (2010). Persistent schema-dependent hippocampal-neocortical connectivity during memory encoding and postencoding rest in humans. Proceedings of the National Academy of Sciences of the United States of America *107*, 7550–7555.

van Kesteren, M.T.R., Beul, S.F., Takashima, A., Henson, R.N., Ruiter, D.J., and Fernández, G. (2013). Differential roles for medial prefrontal and medial temporal cortices in schema-dependent encoding: from congruent to incongruent. Neuropsychologia *51*, 2352–2359.

Kim, J.J., and Fanselow, M.S. (1992). Modality-specific retrograde amnesia of fear. Science *256*, 675–677.

Kirchhoff, B.A., Wagner, A.D., Maril, A., and Stern, C.E. (2000). Prefrontal-temporal circuitry for episodic encoding and subsequent memory. The Journal of Neuroscience *20*, 6173–6180.

Kirwan, C.B., Bayley, P.J., Galván, V. V, and Squire, L.R. (2008). Detailed recollection of remote autobiographical memory after damage to the medial temporal lobe. Proceedings of the National Academy of Sciences of the United States of America *105*, 2676–2680.

Kjonigsen, L.J., Leergaard, T.B., Witter, M.P., and Bjaalie, J.G. (2011). Digital atlas of anatomical subdivisions and boundaries of the rat hippocampal region. Frontiers in Neuroinformatics *5*, 2.

Knauer, B., Jochems, A., Valero-Aracama, M.J., and Yoshida, M. (2013). Long-lasting intrinsic persistent firing in rat CA1 pyramidal cells: a possible mechanism for active maintenance of memory. Hippocampus *23*, 820–831.

Knierim, J.J., McNaughton, B.L., and Poe, G.R. (2000). Three-dimensional spatial selectivity of hippocampal neurons during space flight. Nature Neuroscience *3*, 209–210.

Knowlton, B.J., and Squire, L.R. (1993). The learning of categories: parallel brain systems for item memory and category knowledge. Science *262*, 1747–1749.

Kobayashi, T., Nishijo, H., Fukuda, M., Bures, J., and Ono, T. (1997). Task-dependent representations in rat hippocampal place neurons. Journal of Neurophysiology *78*, 597–613.

Koene, R.A., and Hasselmo, M.E. (2008). Reversed and forward buffering of behavioral spike sequences enables retrospective and prospective retrieval in hippocampal regions CA3 and CA1. Neural Networks *21*, 276–288.

Koenig, J., Linder, A.N., Leutgeb, J.K., and Leutgeb, S. (2011). The spatial periodicity of grid cells is not sustained during reduced theta oscillations. Science *332*, 592–595.

Koh, M.T., and Bernstein, I.L. (2003). Inhibition of protein kinase A activity during conditioned taste aversion retrieval: interference with extinction or reconsolidation of a memory? Neuroreport *14*, 405–407.

Kohonen, T. (1972). Correlation matrix memories. IEEE Transactions on Computers *100*, 353–359.

Kohonen, T. (1984b). Self-Organization and Associative Memory Springer-Verlag Berlin (New York: Springer).

Kohonen, T. (1993). Physiological interpretation of the self-organizing map algorithm. Neural Networks *6*, 895–905.

Komorowski, R.W., Manns, J.R., and Eichenbaum, H. (2009). Robust conjunctive item-place coding by hippocampal neurons parallels learning what happens where. The Journal of Neuroscience *29*, 9918–9929.

Komorowski, R.W., Garcia, C.G., Wilson, A., Hattori, S., Howard, M.W., and Eichenbaum, H. (2013). Ventral hippocampal neurons are shaped by experience to represent behaviorally relevant contexts. The Journal of Neuroscience *33*, 8079–8087.

Kraus, B.J., Robinson, R.J., White, J.A., Eichenbaum, H., and Hasselmo, M.E. (2013). Hippocampal “time cells”: time versus path integration. Neuron *78*, 1090–1101.

Kreiman, G., Koch, C., and Fried, I. (2000). Category-specific visual responses of single neurons in the human medial temporal lobe. Nature Neuroscience *3*, 946–953.

Kumaran, D., and Maguire, E.A. (2006). An unexpected sequence of events: mismatch detection in the human hippocampus. PLoS Biology *4*, e424.

Kumaran, D., and McClelland, J.L. (2012). Generalization through the recurrent interaction of episodic memories: a model of the hippocampal system. Psychological Review *119*, 573–616.

Kumaran, D., Summerfield, J.J., Hassabis, D., and Maguire, E.A. (2009). Tracking the emergence of conceptual knowledge during human decision making. Neuron *63*, 889–901.

Kumaran, D., Melo, H.L., and Duzel, E. (2012). The emergence and representation of knowledge about social and nonsocial hierarchies. Neuron *76*, 653–666.

Kyd, R.J., and Bilkey, D.K. (2005). Hippocampal place cells show increased sensitivity to changes in the local environment following prefrontal cortex lesions. Cerebral Cortex *15*, 720–731.

Larry, S., Neil, C., and Lynn, N. (1984). Memory Consolidation. H. Weingartner, and E.S. Parker, eds. (Hillsdale, NJ: Psychobilogy of Cognition), pp. 185–210.

Lattal, K.M., and Abel, T. (2004). Behavioral impairments caused by injections of the protein synthesis inhibitor anisomycin after contextual retrieval reverse with time. Proceedings of the National Academy of Sciences of the United States of America *101*, 4667–4672.

Lavenex, P., and Banta Lavenex, P. (2013). Building hippocampal circuits to learn and remember: insights into the development of human memory. Behavioural Brain Research *254*, 8–21.

Lee, J.L.C. (2008). Memory reconsolidation mediates the strengthening of memories by additional learning. Nature Neuroscience *11*, 1264–1266.

Lee, J.L.C. (2009). Reconsolidation: maintaining memory relevance. Trends in Neurosciences *32*, 413–420.

Lee, J.L.C. (2010). Memory reconsolidation mediates the updating of hippocampal memory content. Frontiers in Behavioral Neuroscience *4*.

Lee, I., and Park, S.-B. (2013). Perirhinal cortical inactivation impairs object-in-place memory and disrupts task-dependent firing in hippocampal CA1, but not in CA3. Frontiers in Neural Circuits *7*.

Lee, H., Ghim, J.-W., Kim, H., Lee, D., and Jung, M. (2012). Hippocampal neural correlates for values of experienced events. The Journal of Neuroscience *32*, 15053–15065.

Lee, J.L.C., Di Ciano, P., Thomas, K.L., and Everitt, B.J. (2005). Disrupting reconsolidation of drug memories reduces cocaine-seeking behavior. Neuron *47*, 795–801.

Lehmann, H., Sparks, F.T., O'Brien, J., McDonald, R.J., and Sutherland, R.J. (2010). Retrograde amnesia for fear-potentiated startle in rats after complete, but not partial, hippocampal damage. Neuroscience *167*, 974–984.

Lesburguères, E., Gobbo, O.L., Alaux-Cantin, S., Hambucken, A., Trifilieff, P., and Bontempi, B. (2011). Early tagging of cortical networks is required for the formation of enduring associative memory. Science *331*, 924–928.

Leutgeb, S., and Mizumori, S.J. (1999). Excitotoxic septal lesions result in spatial memory deficits and altered flexibility of hippocampal single-unit representations. The Journal of Neuroscience *19*, 6661–6672.

Leutgeb, S., Leutgeb, J.K., Treves, A., Moser, M.-B., and Moser, E.I. (2004). Distinct ensemble codes in hippocampal areas CA3 and CA1. Science *305*, 1295–1298.

Leutgeb, S., Leutgeb, J.K., Barnes, C.A., Moser, E.I., McNaughton, B.L., and Moser, M.-B. (2005). Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. Science *309*, 619–623.

Lever, C., Wills, T., Cacucci, F., Burgess, N., and O'Keefe, J. (2002). Long-term plasticity in hippocampal place-cell representation of environmental geometry. Nature *416*, 90–94.

Levy, W.B. (1996). A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. Hippocampus *6*, 579–590.

Lewis, D.J. (1979). Psychobiology of active and inactive memory. Psychological Bulletin *86*, 1054–1083.

Lin, L., Osan, R., Shoham, S., Jin, W., Zuo, W., and Tsien, J.Z. (2005). Identification of network-level coding units for real-time representation of episodic experiences in the hippocampus. Proceedings of the National Academy of Sciences of the United States of America *102*, 6125–6130.

Lisman, J.E. (1999). Relating hippocampal circuitry to function: recall of memory sequences by reciprocal dentate-CA3 interactions. Neuron *22*, 233–242.

Liu, X., Ramirez, S., and Tonegawa, S. (2014). Inception of a false memory by optogenetic manipulation of a hippocampal memory engram. Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences *369*, 20130142.

Lopez, J., Herbeaux, K., Cosquer, B., Engeln, M., Muller, C., Lazarus, C., Kelche, C., Bontempi, B., Cassel, J.-C., and De Vasconcelos, A.P. (2012). Context-dependent modulation of hippocampal and cortical recruitment during remote spatial memory retrieval. Hippocampus *22*, 827–841.

Louie, K., and Wilson, M.A. (2001). Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. Neuron *29*, 145–156.

Lu, L., Leutgeb, J.K., Tsao, A., Henriksen, E.J., Leutgeb, S., Barnes, C.A., Witter, M.P., Moser, M.-B., and Moser, E.I. (2013). Impaired hippocampal rate coding after lesions of the lateral entorhinal cortex. Nature Neuroscience *16*, 1085–1093.

MacDonald, C.J., Lepage, K.Q., Eden, U.T., and Eichenbaum, H. (2011). Hippocampal "time cells" bridge the gap in memory for discontiguous events. Neuron *71*, 737–749.

Maguire, E.A. (2001). Neuroimaging studies of autobiographical event memory. Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences *356*, 1441–1451.

Maguire, E.A., Nannery, R., and Spiers, H.J. (2006). Navigation around London by a taxi driver with bilateral hippocampal lesions. Brain *129*, 2894–2907.

Mankin, E.A., Sparks, F.T., Slayyeh, B., Sutherland, R.J., Leutgeb, S., and Leutgeb, J.K. (2012). Neuronal code for extended time in the hippocampus. Proceedings of the National Academy of Sciences of the United States of America *109*, 19462–19467.

Manns, J.R., and Eichenbaum, H. (2006). Evolution of declarative memory. Hippocampus *16*, 795–808.

Manns, J.R., and Eichenbaum, H. (2009). A cognitive map for object memory in the hippocampus. Learning & Memory *16*, 616–624.

Manns, J.R., and Squire, L.R. (1999). Impaired recognition memory on the Doors and People Test after damage limited to the hippocampal region. Hippocampus *9*, 495–499.

Manns, J.R., Hopkins, R.O., and Squire, L.R. (2003). Semantic memory and the human hippocampus. Neuron *38*, 127–133.

Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. Science *275*, 213–215.

Markus, E., Qin, Y., Leonard, B., Skaggs, W., McNaughton, B., and Barnes, C. (1995). Interactions between location and task affect the spatial and directional firing of hippocampal neurons. The Journal of Neuroscience *15*, 7079–7094.

Marr, D. (1971). Simple memory: a theory for archicortex. Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences *262*, 23–81.

Martin, S.J., De Hoz, L., and Morris, R.G.M. (2005). Retrograde amnesia: neither partial nor complete hippocampal lesions in rats result in preferential sparing of remote spatial memory, even after reminding. Neuropsychologia *43*, 609–624.

Matus-Amat, P., Higgins, E.A., Barrientos, R.M., and Rudy, J.W. (2004). The role of the dorsal hippocampus in the acquisition and retrieval of context memory representations. The Journal of Neuroscience *24*, 2431–2439.

Maviel, T., Durkin, T.P., Menzaghi, F., and Bontempi, B. (2004). Sites of neocortical reorganization critical for remote spatial memory. Science *305*, 96–99.

Mayes, A.R., Holdstock, J.S., Isaac, C.L., Hunkin, N.M., and Roberts, N. (2002). Relative sparing of item recognition memory in a patient with adult-onset damage limited to the hippocampus. Hippocampus *12*, 325–340.

Mayes, A.R., Holdstock, J.S., Isaac, C.L., Montaldi, D., Grigor, J., Gummer, A., Cariga, P., Downes, J.J., Tsivilis, D., Gaffan, D., et al. (2004). Associative recognition in a patient with selective hippocampal lesions and relatively normal item recognition. Hippocampus *14*, 763–784.

McClelland, J.L. (2013). Incorporating rapid neocortical learning of new schema-consistent information into complementary learning systems theory. Journal of Experimental Psychology. General *142*, 1190–1210.

McClelland, J.L., McNaughton, B.L., and 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*, 419–457.

McCloskey, M., and Cohen, N.J. (1989). Catastrophic interference in connectionist networks : the sequential learning problem. The Psychology of Learning and Motivation *24*, 109–165.

McCulloch, W.S., and Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. The Bulletin of Mathematical Biophysics *5*, 115–133.

McDonald, R.J., Ko, C.H., and Hong, N.S. (2002). Attenuation of context-specific inhibition on reversal learning of a stimulus-response task in rats with neurotoxic hippocampal damage. Behavioural Brain Research *136*, 113–126.

McGaugh, J.L. (2000). Memory--a century of consolidation. Science *287*, 248–251.

McGraw, P.N., and Menzinger, M. (2003). Topology and computational performance of attractor neural networks. Physical Review. E, Statistical, Nonlinear, and Soft Matter Physics *68*, 047102.

McKenzie, S., and Eichenbaum, H. (2011). Consolidation and reconsolidation: two lives of memories? Neuron *71*, 224–233.

McKenzie, S., Robinson, N.T.M., Herrera, L., Churchill, J.C., and Eichenbaum, H. (2013). Learning Causes Reorganization of Neuronal Firing Patterns to Represent Related Experiences within a Hippocampal Schema. The Journal of Neuroscience *33*, 10243–10256.

McNaughton, B.L., and Morris, R.G. (1987). Hippocampal synaptic enhancement and information storage within a distributed memory system. Trends in Neurosciences *10*, 408–415.

McNaughton, B.L., and Nadel, L. (1990). Hebb-Marr networks and the neurobiological representation of action in space. In Neuroscience and Connectionist Theory, M. Gluck, and RumelhartDavid, eds. (Hillsdale, NJ: Lawrence Erlbaum Associates, Inc), pp. 1–63.

McNaughton, B.L., Barnes, C.A., and O’Keefe, J. (1983). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. Experimental Brain Research *52*, 41–49.

McNaughton, B.L., Barnes, C.A., Gerrard, J.L., Gothard, K., Jung, M.W., Knierim, J.J., Kudrimoti, H., Qin, Y., Skaggs, W.E., Suster, M., et al. (1996). Deciphering the hippocampal polyglot: the hippocampus as a path integration system. The Journal of Experimental Biology *199*, 173–185.

McNaughton, B.L., Battaglia, F.P., Jensen, O., Moser, E.I., and Moser, M.-B. (2006). Path integration and the neural basis of the “cognitive map”. Nature Reviews. Neuroscience *7*, 663–678.

Meunier, M., Bachevalier, J., Mishkin, M., and Murray, E.A. (1993). Effects on visual recognition of combined and separate ablations of the entorhinal and perirhinal cortex in rhesus monkeys. The Journal of Neuroscience *13*, 5418–5432.

Meunier, M., Hadfield, W., Bachevalier, J., and Murray, E.A. (1996). Effects of rhinal cortex lesions combined with hippocampectomy on visual recognition memory in rhesus monkeys. Journal of Neurophysiology *75*, 1190–1205.

Milekic, M.H., and Alberini, C.M. (2002). Temporally graded requirement for protein synthesis following memory reactivation. Neuron *36*, 521–525.

Miller, V.M., and Best, P.J. (1980). Spatial correlates of hippocampal unit activity are altered by lesions of the fornix and endorhinal cortex. Brain Research *194*, 311–323.

Miller, J.F., Neufang, M., Solway, A., Brandt, A., Trippel, M., Mader, I., Hefft, S., Merkow, M., Polyn, S.M., Jacobs, J., et al. (2013). Neural activity in human hippocampal formation reveals the spatial context of retrieved memories. Science *342*, 1111–1114.

Milner, B., Corkin, S., and Teuber, H.-L. (1968). Further analysis of the hippocampal amnesic syndrome: 14-year follow-up study of H.M. Neuropsychologia *6*, 215–234.

Milner, B., Squire, L.R., and Kandel, E.R. (1998). Cognitive neuroscience and the study of memory. Neuron *20*, 445–468.

Misanin, J.R., Miller, R.R., and Lewis, D.J. (1968). Retrograde amnesia produced by electroconvulsive shock after reactivation of a consolidated memory trace. Science *160*, 554–555.

Mishkin, M. (1978). Memory in monkeys severely impaired by combined but not by separate removal of amygdala and hippocampus. Nature *273*, 297–298.

Mishkin, M., Malamut, B., and Bachevalier, J. (1984). Memories and Habits: Two Neural Systems. In Neurobiology of Learning and Memory, G. Lynch, J. McGaugh, and N. Weinberger, eds. (New York: guilford press), pp. 65–77.

Moita, M.A.P., Rosis, S., Zhou, Y., LeDoux, J.E., and Blair, H.T. (2003). Hippocampal place cells acquire location-specific responses to the conditioned stimulus during auditory fear conditioning. Neuron *37*, 485–497.

Moita, M.A.P., Rosis, S., Zhou, Y., LeDoux, J.E., and Blair, H.T. (2004). Putting fear in its place: remapping of hippocampal place cells during fear conditioning. The Journal of Neuroscience *24*, 7015–7023.

Monaghan, D.T., and Cotman, C.W. (1985). Distribution of N-methyl-D-aspartate-sensitive L-[3H]glutamate-binding sites in rat brain. The Journal of Neuroscience *5*, 2909–2919.

Monfils, M.-H., Cowansage, K.K., Klann, E., and LeDoux, J.E. (2009). Extinction-reconsolidation boundaries: key to persistent attenuation of fear memories. Science *324*, 951–955.

Montaldi, D., Spencer, T.J., Roberts, N., and Mayes, A.R. (2006). The neural system that mediates familiarity memory. Hippocampus *16*, 504–520.

Morris, R.G., Garrud, P., Rawlins, J.N., and O'Keefe, J. (1982). Place navigation impaired in rats with hippocampal lesions. Nature *297*, 681–683.

Morris, R.G.M., Inglis, J., Ainge, J.A., Olverman, H.J., Tulloch, J., Dudai, Y., and Kelly, P.A.T. (2006). Memory reconsolidation: sensitivity of spatial memory to inhibition of protein synthesis in dorsal hippocampus during encoding and retrieval. Neuron *50*, 479–489.

Moser, E.I., Kropff, E., and Moser, M.-B. (2008). Place cells, grid cells, and the brain's spatial representation system. Annual Review of Neuroscience *31*, 69–89.

Müller G.E., P.A. (1900). Experimentelle Beiträge zur Lehre vom Gedächtnis. Z. Psychol. Ergänzungsband *1*, 1–300.

Muller, R.U., and Kubie, J.L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. The Journal of Neuroscience *7*, 1951–1968.

Muller, R.U., Kubie, J.L., and Ranck, J.B. (1987). Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. The Journal of Neuroscience *7*, 1935–1950.

Mumby, D.G., Gaskin, S., Glenn, M.J., Schramek, T.E., and Lehmann, H. (2002). Hippocampal damage and exploratory preferences in rats: memory for objects, places, and contexts. Learning & Memory *9*, 49–57.

Murray, E.A., and Mishkin, M. (1998). Object recognition and location memory in monkeys with excitotoxic lesions of the amygdala and hippocampus. The Journal of Neuroscience *18*, 6568–6582.

Myers, C.E., Shohamy, D., Gluck, M.A., Grossman, S., Kluger, A., Ferris, S., Golomb, J., Schnirman, G., and Schwartz, R. (2003). Dissociating hippocampal versus basal ganglia contributions to learning and transfer. Journal of Cognitive Neuroscience *15*, 185–193.

Nadel, L., and Moscovitch, M. (1997). Memory consolidation, retrograde amnesia and the hippocampal complex. Current Opinion in Neurobiology *7*, 217–227.

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

Nader, K., and Hardt, O. (2009). A single standard for memory: the case for reconsolidation. Nature Reviews. Neuroscience *10*, 224–234.

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

Nagel, T. (1974). What Is It Like to Be a Bat? The Philosophical Review *83*, 435–450.

Nagode, J.C., and Pardo, J. V (2002). Human hippocampal activation during transitive inference. Neuroreport *13*, 939–944.

Nakashiba, T., Young, J.Z., McHugh, T.J., Buhl, D.L., and Tonegawa, S. (2008). Transgenic inhibition of synaptic transmission reveals role of CA3 output in hippocampal learning. Science *319*, 1260–1264.

Nakashiba, T., Buhl, D.L., McHugh, T.J., and Tonegawa, S. (2009). Hippocampal CA3 output is crucial for ripple-associated reactivation and consolidation of memory. Neuron *62*, 781–787.

Navratilova, Z., Hoang, L.T., Schwindel, C.D., Tatsuno, M., and McNaughton, B.L. (2012). Experience-dependent firing rate remapping generates directional selectivity in hippocampal place cells. Frontiers in Neural Circuits *6*.

Nemanic, S., Alvarado, M.C., and Bachevalier, J. (2004). The hippocampal/parahippocampal regions and recognition memory: insights from visual paired comparison versus object-delayed nonmatching in monkeys. The Journal of Neuroscience *24*, 2013–2026.

Neunuebel, J.P., and Knierim, J.J. (2014). CA3 retrieves coherent representations from degraded input: direct evidence for CA3 pattern completion and dentate gyrus pattern separation. Neuron *81*, 416–427.

Neunuebel, J.P., Yoganarasimha, D., Rao, G., and Knierim, J.J. (2013). Conflicts between local and global spatial frameworks dissociate neural representations of the lateral and medial entorhinal cortex. The Journal of Neuroscience *33*, 9246–9258.

Newmark, R.E., Schon, K., Ross, R.S., and Stern, C.E. (2013). Contributions of the hippocampal subfields and entorhinal cortex to disambiguation during working memory. Hippocampus *23*, 467–475.

Norman, K.A. (2010). How hippocampus and cortex contribute to recognition memory: revisiting the complementary learning systems model. Hippocampus *20*, 1217–1227.

Norman, K.A., and O’Reilly, R.C. (2003). Modeling hippocampal and neocortical contributions to recognition memory: a complementary-learning-systems approach. Psychological Review *110*, 611–646.

Noulhiane, M., Piolino, P., Hasboun, D., Clemenceau, S., Baulac, M., and Samson, S. (2007). Autobiographical memory after temporal lobe resection: neuropsychological and MRI volumetric findings. Brain *130*, 3184–3199.

O’Keefe, J. (1999). Do hippocampal pyramidal cells signal non-spatial as well as spatial information? Hippocampus *9*, 352–364.

O’Keefe, J., and Burgess, N. (2005). Dual phase and rate coding in hippocampal place cells: theoretical significance and relationship to entorhinal grid cells. Hippocampus *15*, 853–866.

O’Keefe, J., and 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., and Nadel, L. (1978). The hippocampus as a cognitive map (New York: Clarendon Press).

O’Neill, J., Senior, T., and Csicsvari, J. (2006). Place-selective firing of CA1 pyramidal cells during sharp wave/ripple network patterns in exploratory behavior. Neuron *49*, 143–155.

O’Reilly, R.C., and McClelland, J.L. (1994). Hippocampal conjunctive encoding, storage, and recall: avoiding a trade-off. Hippocampus *4*, 661–682.

Ohara, S., Sato, S., Tsutsui, K.-I., Witter, M.P., and Iijima, T. (2013). Organization of multisynaptic inputs to the dorsal and ventral dentate gyrus: retrograde trans-synaptic tracing with rabies virus vector in the rat. PloS One *8*, e78928.

Oja, E. (1989). Neural networks, principal components, and subspaces. International Journal of Neural Systems *1*, 61–68.

Parkinson, J.K., Murray, E.A., and Mishkin, M. (1988). A selective mnemonic role for the hippocampus in monkeys: memory for the location of objects. The Journal of Neuroscience *8*, 4159–4167.

Pastalkova, E., Itskov, V., Amarasingham, A., and Buzsáki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. Science *321*, 1322–1327.

Paz-Villagrán, V., Save, E., and Poucet, B. (2004). Independent coding of connected environments by place cells. The European Journal of Neuroscience *20*, 1379–1390.

Pedreira, M.E., and Maldonado, H. (2003). Protein synthesis subserves reconsolidation or extinction depending on reminder duration. Neuron *38*, 863–869.

Pedreira, M.E., Pérez-Cuesta, L.M., and Maldonado, H. (2002). Reactivation and reconsolidation of long-term memory in the crab Chasmagnathus: protein synthesis requirement and mediation by NMDA-type glutamatergic receptors. The Journal of Neuroscience *22*, 8305–8311.

Pedreira, M.E., Pérez-Cuesta, L.M., and Maldonado, H. (2004). Mismatch between what is expected and what actually occurs triggers memory reconsolidation or extinction. Learning & Memory *11*, 579–585.

Pfeiffer, B.E., and Foster, D.J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. Nature *497*, 74–79.

Piaget, J. (1926). The language and thought of the child (New York: Harcourt, Brace.).

Piaget, J. (1928). Judgment and Reasoning in the Child (London: Kegan, Paul, Trench, and Trubner).

Piaget, J. (1929). The child’s conception of the world(J. Tomlinson & A. Thomlinson, Trans.) (New York: Harcourt, Brace & Company, Inc.).

Piolino, P., Giffard-Quillon, G., Desgranges, B., Chételat, G., Baron, J.-C., and Eustache, F. (2004). Re-experiencing old memories via hippocampus: a PET study of autobiographical memory. NeuroImage *22*, 1371–1383.

Pitkänen, A., Pikkarainen, M., Nurminen, N., and Ylinen, A. (2000). Reciprocal connections between the amygdala and the hippocampal formation, perirhinal cortex, and postrhinal cortex in rat. A review. Annals of the New York Academy of Sciences *911*, 369–391.

Power, A.E., Berlau, D.J., McGaugh, J.L., and Steward, O. (2006). Anisomycin infused into the hippocampus fails to block “reconsolidation” but impairs extinction: the role of re-exposure duration. Learning & Memory *13*, 27–34.

Preston, A.R., Shrager, Y., Dudukovic, N.M., and Gabrieli, J.D.E. (2004). Hippocampal contribution to the novel use of relational information in declarative memory. Hippocampus *14*, 148–152.

Quamme, J.R., Yonelinas, A.P., Widaman, K.F., Kroll, N.E.A., and Sauvé, M.J. (2004). Recall and recognition in mild hypoxia: using covariance structural modeling to test competing theories of explicit memory. Neuropsychologia *42*, 672–691.

Quirk, G.J., Muller, R.U., and Kubie, J.L. (1990). The firing of hippocampal place cells in the dark depends on the rat’s recent experience. The Journal of Neuroscience *10*, 2008–2017.

Quiroga, R.Q., Reddy, L., Kreiman, G., Koch, C., and Fried, I. (2005). Invariant visual representation by single neurons in the human brain. Nature *435*, 1102–1107.

Rajji, T., Chapman, D., Eichenbaum, H., and Greene, R. (2006). The role of CA3 hippocampal NMDA receptors in paired associate learning. The Journal of Neuroscience *26*, 908–915.

Ranganath, C. (2010). A unified framework for the functional organization of the medial temporal lobes and the phenomenology of episodic memory. Hippocampus *20*, 1263–1290.

Ranganath, C., and Rainer, G. (2003). Neural mechanisms for detecting and remembering novel events. Nature Reviews. Neuroscience *4*, 193–202.

Rao, V.A., and Howard, M.W. (2008). Retrieved context and the discovery of semantic structure. Advances in Neural Information Processing Systems *20*, 1193–1200.

Ravassard, P., Kees, A., Willers, B., Ho, D., Aharoni, D., Cushman, J., Aghajan, Z.M., and Mehta, M.R. (2013). Multisensory control of hippocampal spatiotemporal selectivity. Science *340*, 1342–1346.

Reed, J.M., and Squire, L.R. (1998). Retrograde amnesia for facts and events: findings from four new cases. The Journal of Neuroscience *18*, 3943–3954.

Rempel-Clower, N.L., Zola, S.M., Squire, L.R., and Amaral, D.G. (1996). Three cases of enduring memory impairment after bilateral damage limited to the hippocampal formation. The Journal of Neuroscience *16*, 5233–5255.

Ribot, T. (1882). Diseases of memory. In Essays in the Positive Psychology. The Iernational Scientific Series, (Kegan Paul, Trench & Co), p. 209.

Robinson, M.J.F., and Franklin, K.B.J. (2010). Reconsolidation of a morphine place preference: impact of the strength and age of memory on disruption by propranolol and midazolam. Behavioural Brain Research *213*, 201–207.

Robinson, M.J.F., Ross, E.C., and Franklin, K.B.J. (2011). The effect of propranolol dose and novelty of the reactivation procedure on the reconsolidation of a morphine place preference. Behavioural Brain Research *216*, 281–284.

Rodriguez-Ortiz, C.J., De la Cruz, V., Gutiérrez, R., and Bermudez-Rattoni, F. (2005). Protein synthesis underlies post-retrieval memory consolidation to a restricted degree only when updated information is obtained. Learning & Memory *12*, 533–537.

Rodriguez-Ortiz, C.J., Garcia-DeLaTorre, P., Benavidez, E., Ballesteros, M.A., and Bermudez-Rattoni, F. (2008). Intrahippocampal anisomycin infusions disrupt previously consolidated spatial memory only when memory is updated. Neurobiology of Learning and Memory *89*, 352–359.

Rolls, E.T. (1990). Theoretical and neurophysiological analysis of the functions of the primate hippocampus in memory. Cold Spring Harbor Symposia on Quantitative Biology *55*, 995–1006.

Rolls, E.T. (2013). A quantitative theory of the functions of the hippocampal CA3 network in memory. Frontiers in Cellular Neuroscience *7*.

Rolls, E.T., Miyashita, Y., Cahusac, P.M., Kesner, R.P., Niki, H., Feigenbaum, J.D., and Bach, L. (1989). Hippocampal neurons in the monkey with activity related to the place in which a stimulus is shown. The Journal of Neuroscience *9*, 1835–1845.

Rolls, E.T., Stringer, S.M., and Elliot, T. (2006). Entorhinal cortex grid cells can map to hippocampal place cells by competitive learning. Network *17*, 447–465.

Rose, J.K., and Rankin, C.H. (2006). Blocking memory reconsolidation reverses memory-associated changes in glutamate receptor expression. The Journal of Neuroscience *26*, 11582–11587.

Rosenbaum, R.S., Winocur, G., and Moscovitch, M. (2001). New views on old memories: re-evaluating the role of the hippocampal complex. Behavioural Brain Research *127*, 183–197.

Rosenbaum, R.S., McKinnon, M.C., Levine, B., and Moscovitch, M. (2004). Visual imagery deficits, impaired strategic retrieval, or memory loss: disentangling the nature of an amnesic person’s autobiographical memory deficit. Neuropsychologia *42*, 1619–1635.

Ross, R.S., and Eichenbaum, H. (2006). Dynamics of hippocampal and cortical activation during consolidation of a nonspatial memory. The Journal of Neuroscience *26*, 4852–4859.

Rudy, J.W., and Sutherland, R.J. (2008). Is it systems or cellular consolidation? Time will tell. An alternative interpretation of the Morris group’s recent science paper. Neurobiology of Learning and Memory *89*, 366–369.

Rumelhart, D.E., and Mcclelland, J.L. (1986). Parallel distributed processing: explorations in the microstructure of cognition. Volume 1. Foundations (Cambridge, MA: MIT Press).

Rutishauser, U., Mamelak, A.N., and Schuman, E.M. (2006). Single-trial learning of novel stimuli by individual neurons of the human hippocampus-amygdala complex. Neuron *49*, 805–813.

Rutishauser, U., Schuman, E.M., and Mamelak, A.N. (2008). Activity of human hippocampal and amygdala neurons during retrieval of declarative memories. Proceedings of the National Academy of Sciences of the United States of America *105*, 329–334.

Ryan, L., Nadel, L., Keil, K., Putnam, K., Schnyer, D., Trouard, T., and Moscovitch, M. (2001). Hippocampal complex and retrieval of recent and very remote autobiographical memories: evidence from functional magnetic resonance imaging in neurologically intact people. Hippocampus *11*, 707–714.

Sacco, T., and Sacchetti, B. (2010). Role of secondary sensory cortices in emotional memory storage and retrieval in rats. Science *329*, 649–656.

Sagar, H.J., Cohen, N.J., Corkin, S., and Growdon, J.H. (1985). Dissociations among processes in remote memory. Annals of the New York Academy of Sciences *444*, 533–535.

Sakai, K., and Miyashita, Y. (1991). Neural organization for the long-term memory of paired associates. Nature *354*, 152–155.

Samsonovich, A., and McNaughton, B.L. (1997). Path integration and cognitive mapping in a continuous attractor neural network model. The Journal of Neuroscience *17*, 5900–5920.

Sander, K., and Sander, D. (2005). New insights into transient global amnesia: recent imaging and clinical findings. Lancet Neurology *4*, 437–444.

Sara, S.J. (2010). Reactivation, retrieval, replay and reconsolidation in and out of sleep: connecting the dots. Frontiers in Behavioral Neuroscience *4*.

Save, E., Cressant, A., Thinus-Blanc, C., and Poucet, B. (1998). Spatial firing of hippocampal place cells in blind rats. The Journal of Neuroscience *18*, 1818–1826.

Schacter, D.L. (1999). The seven sins of memory. Insights from psychology and cognitive neuroscience. The American Psychologist *54*, 182–203.

Schendan, H.E., Searl, M.M., Melrose, R.J., and Stern, C.E. (2003). An FMRI study of the role of the medial temporal lobe in implicit and explicit sequence learning. Neuron *37*, 1013–1025.

Schiller, D., Monfils, M.-H., Raio, C.M., Johnson, D.C., Ledoux, J.E., and Phelps, E.A. (2010). Preventing the return of fear in humans using reconsolidation update mechanisms. Nature *463*, 49–53.

Schmitzer-Torbert, N., Jackson, J., Henze, D., Harris, K., and Redish, A.D. (2005). Quantitative measures of cluster quality for use in extracellular recordings. Neuroscience *131*, 1–11.

Schneider, A.M., and Sherman, W. (1968). Amnesia: a function of the temporal relation of footshock to electroconvulsive shock. Science *159*, 219–221.

Scoville, W.B., and Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. Journal of Experimental Psychology. Animal Behavior Processes *20*, 11–21.

Shapiro, M.L., Tanila, H., and Eichenbaum, H. (1997). Cues that hippocampal place cells encode: dynamic and hierarchical representation of local and distal stimuli. Hippocampus *7*, 624–642.

Sharp, P.E., and Koester, K. (2008). Lesions of the mammillary body region severely disrupt the cortical head direction, but not place cell signal. Hippocampus *18*, 766–784.

Shew, W.L., Yang, H., Yu, S., Roy, R., and Plenz, D. (2011). Information capacity and transmission are maximized in balanced cortical networks with neuronal avalanches. The Journal of Neuroscience *31*, 55–63.

Shohamy, D., and Wagner, A.D. (2008). Integrating memories in the human brain: hippocampal-midbrain encoding of overlapping events. Neuron *60*, 378–389.

Singer, A.C., Karlsson, M.P., Nathe, A.R., Carr, M.F., and Frank, L.M. (2010). Experience-dependent development of coordinated hippocampal spatial activity representing the similarity of related locations. The Journal of Neuroscience *30*, 11586–11604.

Skaggs, W.E., and McNaughton, B.L. (1998). Spatial firing properties of hippocampal CA1 populations in an environment containing two visually identical regions. The Journal of Neuroscience *18*, 8455–8466.

Skaggs, W.E., McNaughton, B.L., Wilson, M.A., and Barnes, C.A. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. Hippocampus *6*, 149–172.

Sloviter, R.S., and Brisman, J.L. (1995). Lateral inhibition and granule cell synchrony in the rat hippocampal dentate gyrus. The Journal of Neuroscience *15*, 811–820.

Smith, C., and Squire, L.R. (2005). Declarative memory, awareness, and transitive inference. The Journal of Neuroscience *25*, 10138–10146.

Smith, C.N., and Squire, L.R. (2009). Medial temporal lobe activity during retrieval of semantic memory is related to the age of the memory. The Journal of Neuroscience *29*, 930–938.

Smith, D.M., and Mizumori, S.J.Y. (2006). Learning-related development of context-specific neuronal responses to places and events: the hippocampal role in context processing. The Journal of Neuroscience *26*, 3154–3163.

Solstad, T., Moser, E.I., and Einevoll, G.T. (2006). From grid cells to place cells: a mathematical model. Hippocampus *16*, 1026–1031.

Song, E.Y., Kim, Y.B., Kim, Y.H., and Jung, M.W. (2005). Role of active movement in place-specific firing of hippocampal neurons. Hippocampus *15*, 8–17.

Song, Z., Jeneson, A., and Squire, L.R. (2011). Medial temporal lobe function and recognition memory: a novel approach to separating the contribution of recollection and familiarity. The Journal of Neuroscience *31*, 16026–16032.

Speakman, A., and O'Keefe, J. (1990). Hippocampal Complex Spike Cells do not Change Their Place Fields if the Goal is Moved Within a Cue Controlled Environment. The European Journal of Neuroscience *2*, 544–555.

Spear, N.E., Smith, G.J., Bryan, R.G., Gordon, W.C., Timmons, R., and Chiszar, D.A. (1980). Contextual influences on the interaction between conflicting memories in the rat. Animal Learning & Behavior *8*, 273–281.

Spiers, H.J., Hayman, R.M.A., Jovalekic, A., Marozzi, E., and Jeffery, K.J. (2013). Place Field Repetition and Purely Local Remapping in a Multicompartment Environment. Cerebral Cortex.

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

Squire, L.R., and Bayley, P.J. (2007). The neuroscience of remote memory. Current Opinion in Neurobiology *17*, 185–196.

Squire, L.R., and Zola, S.M. (1996). Structure and function of declarative and nondeclarative memory systems. Proceedings of the National Academy of Sciences of the United States of America *93*, 13515–13522.

Staresina, B.P., and Davachi, L. (2006). Differential encoding mechanisms for subsequent associative recognition and free recall. The Journal of Neuroscience *26*, 9162–9172.

Staresina, B.P., Gray, J.C., and Davachi, L. (2009). Event congruency enhances episodic memory encoding through semantic elaboration and relational binding. Cerebral Cortex *19*, 1198–1207.

Steinvorth, S., Levine, B., and Corkin, S. (2005). Medial temporal lobe structures are needed to re-experience remote autobiographical memories: evidence from H.M. and W.R. Neuropsychologia *43*, 479–496.

Stern, C.E., Corkin, S., González, R.G., Guimaraes, A.R., Baker, J.R., Jennings, P.J., Carr, C.A., Sugiura, R.M., Vedantham, V., and Rosen, B.R. (1996). The hippocampal formation participates in novel picture encoding: evidence from functional magnetic resonance imaging. Proceedings of the National Academy of Sciences of the United States of America *93*, 8660–8665.

Suh, J., Rivest, A.J., Nakashiba, T., Tominaga, T., and Tonegawa, S. (2011). Entorhinal cortex layer III input to the hippocampus is crucial for temporal association memory. Science *334*, 1415–1420.

Sutherland, R.J., and Lehmann, H. (2011). Alternative conceptions of memory consolidation and the role of the hippocampus at the systems level in rodents. Current Opinion in Neurobiology *21*, 446–451.

Suzuki, A., Josselyn, S.A., Frankland, P.W., Masushige, S., Silva, A.J., and Kida, S. (2004). Memory reconsolidation and extinction have distinct temporal and biochemical signatures. The Journal of Neuroscience *24*, 4787–4795.

Takashima, A., Nieuwenhuis, I.L.C., Jensen, O., Talamini, L.M., Rijpkema, M., and Fernández, G. (2009). Shift from hippocampal to neocortical centered retrieval network with consolidation. The Journal of Neuroscience *29*, 10087–10093.

Takehara-Nishiuchi, K., and McNaughton, B.L. (2008). Spontaneous changes of neocortical code for associative memory during consolidation. Science *322*, 960–963.

Tayler, K.K., Tanaka, K.Z., Reijmers, L.G., and Wiltgen, B.J. (2013). Reactivation of neural ensembles during the retrieval of recent and remote memory. Current Biology *23*, 99–106.

Teng, E., and Squire, L.R. (1999). Memory for places learned long ago is intact after hippocampal damage. Nature *400*, 675–677.

Teyler, T.J., and DiScenna, P. (1986). The hippocampal memory indexing theory. Behavioral Neuroscience *100*, 147–154.

Thompson, L.T., and Best, P.J. (1990). Long-term stability of the place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats. Brain Research *509*, 299–308.

Tort, A.B.L., Komorowski, R., Kopell, N., and Eichenbaum, H. (2011). A mechanism for the formation of hippocampal neuronal firing patterns that represent what happens where. Learning & Memory *18*, 718–727.

Treves, A., and Rolls, E.T. (1994). Computational analysis of the role of the hippocampus in memory. Hippocampus *4*, 374–391.

Tronel, S., Milekic, M.H., and Alberini, C.M. (2005). Linking new information to a reactivated memory requires consolidation and not reconsolidation mechanisms. PLoS Biology *3*, e293.

Trullier, O., Shibata, R., Mulder, A.B., and Wiener, S.I. (1999). Hippocampal neuronal position selectivity remains fixed to room cues only in rats alternating between place navigation and beacon approach tasks. The European Journal of Neuroscience *11*, 4381–4388.

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

Tse, D., Takeuchi, T., Kakeyama, M., Kajii, Y., Okuno, H., Tohyama, C., Bito, H., and Morris, R.G.M. (2011). Schema-dependent gene activation and memory encoding in neocortex. Science *333*, 891–895.

Tulving, E. (1972). Episodic and semantic memory. In Organization of Memory, (New York: Academic Press), pp. 381–402.

Tulving, E. (2005). Episodic Memory and Autonoesis: Uniquely Human? In The Missing Link in Cognition : Origins of Self-Reflective Consciousness, H.S. Terrace, and J. Metcalfe, eds. (New York: Oxford University Press), pp. 3–56.

Tulving, E., and Markowitsch, H.J. (1998). Episodic and declarative memory: role of the hippocampus. Hippocampus *8*, 198–204.

Valjent, E., Corbillé, A.-G., Bertran-Gonzalez, J., Hervé, D., and Girault, J.-A. (2006). Inhibition of ERK pathway or protein synthesis during reexposure to drugs of abuse erases previously learned place preference. Proceedings of the National Academy of Sciences of the United States of America *103*, 2932–2937.

Vargha-Khadem, F., Gadian, D.G., Watkins, K.E., Connelly, A., Van Paesschen, W., and Mishkin, M. (1997). Differential effects of early hippocampal pathology on episodic and semantic memory. Science *277*, 376–380.

Vazdarjanova, A., and Guzowski, J.F. (2004). Differences in hippocampal neuronal population responses to modifications of an environmental context: evidence for distinct, yet complementary, functions of CA3 and CA1 ensembles. The Journal of Neuroscience *24*, 6489–6496.

Vianna, M.R., Szapiro, G., McGaugh, J.L., Medina, J.H., and Izquierdo, I. (2001). Retrieval of memory for fear-motivated training initiates extinction requiring protein synthesis in the rat hippocampus. Proceedings of the National Academy of Sciences of the United States of America *98*, 12251–12254.

Viard, A., Piolino, P., Desgranges, B., Chételat, G., Lebreton, K., Landeau, B., Young, A., De La Sayette, V., and Eustache, F. (2007). Hippocampal activation for autobiographical memories over the entire lifetime in healthy aged subjects: an fMRI study. Cerebral Cortex *17*, 2453–2467.

Victor, M., and Agamanolis, D. (1990). Amnesia due to Lesions Confined to the Hippocampus: A Clinical-Pathologic Study. Journal of Cognitive Neuroscience *2*, 246–257.

Viskontas, I. V, Knowlton, B.J., Steinmetz, P.N., and Fried, I. (2006). Differences in mnemonic processing by neurons in the human hippocampus and parahippocampal regions. Journal of Cognitive Neuroscience *18*, 1654–1662.

Wais, P.E., Wixted, J.T., Hopkins, R.O., and Squire, L.R. (2006). The hippocampus supports both the recollection and the familiarity components of recognition memory. Neuron *49*, 459–466.

Wais, P.E., Squire, L.R., and Wixted, J.T. (2010). In search of recollection and familiarity signals in the hippocampus. Journal of Cognitive Neuroscience *22*, 109–123.

Walker, M.P., Brakefield, T., Hobson, J.A., and Stickgold, R. (2003). Dissociable stages of human memory consolidation and reconsolidation. Nature *425*, 616–620.

Wallenstein, G. V, and Hasselmo, M.E. (1997). GABAergic modulation of hippocampal population activity: sequence learning, place field development, and the phase precession effect. Journal of Neurophysiology *78*, 393–408.

Wang, S.-H., and Morris, R.G.M. (2010). Hippocampal-neocortical interactions in memory formation, consolidation, and reconsolidation. Annual Review of Psychology *61*, 49–79.

Wang, M.E., Wann, E.G., Yuan, R.K., Ramos Álvarez, M.M., Stead, S.M., and Muzzio, I.A. (2012). Long-term stabilization of place cell remapping produced by a fearful experience. The Journal of Neuroscience *32*, 15802–15814.

Wang, S.-H., De Oliveira Alvares, L., and Nader, K. (2009). Cellular and systems mechanisms of memory strength as a constraint on auditory fear reconsolidation. Nature Neuroscience *12*, 905–912.

Watts, D.J., and Strogatz, S.H. (1998). Collective dynamics of “small-world” networks. Nature *393*, 440–442.

Weinberger, N.M., Miasnikov, A.A., and Chen, J.C. (2009). Sensory memory consolidation observed: increased specificity of detail over days. Neurobiology of Learning and Memory *91*, 273–286.

Wiebe, S.P., and Stäubli, U. V (1999). Dynamic filtering of recognition memory codes in the hippocampus. The Journal of Neuroscience *19*, 10562–10574.

Wills, T.J., Lever, C., Cacucci, F., Burgess, N., and O’Keefe, J. (2005). Attractor dynamics in the hippocampal representation of the local environment. Science *308*, 873–876.

Wilson, M.A., and McNaughton, B.L. (1994). Reactivation of hippocampal ensemble memories during sleep. Science *265*, 676–679.

Wiltgen, B.J., and Silva, A.J. (2007). Memory for context becomes less specific with time. Learning & Memory *14*, 313–317.

Wiltgen, B.J., Zhou, M., Cai, Y., Balaji, J., Karlsson, M.G., Parivash, S.N., Li, W., and Silva, A.J. (2010). The hippocampus plays a selective role in the retrieval of detailed contextual memories. Current Biology *20*, 1336–1344.

Wimmer, G.E., and Shohamy, D. (2012). Preference by association: how memory mechanisms in the hippocampus bias decisions. Science *338*, 270–273.

Winocur, G., Moscovitch, M., Fogel, S., Rosenbaum, R.S., and Sekeres, M. (2005a). Preserved spatial memory after hippocampal lesions: effects of extensive experience in a complex environment. Nature Neuroscience *8*, 273–275.

Winocur, G., Moscovitch, M., Caruana, D.A., and Binns, M.A. (2005b). Retrograde amnesia in rats with lesions to the hippocampus on a test of spatial memory. Neuropsychologia *43*, 1580–1590.

Winocur, G., Moscovitch, M., and Sekeres, M. (2007). Memory consolidation or transformation: context manipulation and hippocampal representations of memory. Nature Neuroscience *10*, 555–557.

Winocur, G., Frankland, P.W., Sekeres, M., Fogel, S., and Moscovitch, M. (2009). Changes in context-specificity during memory reconsolidation: selective effects of hippocampal lesions. Learning & Memory *16*, 722–729.

Winocur, G., Moscovitch, M., and Bontempi, B. (2010). Memory formation and long-term retention in humans and animals: convergence towards a transformation account of hippocampal-neocortical interactions. Neuropsychologia *48*, 2339–2356.

Winters, B.D., Tucci, M.C., and DaCosta-Furtado, M. (2009). Older and stronger object memories are selectively destabilized by reactivation in the presence of new information. Learning & Memory *16*, 545–553.

Wirth, S., Yanike, M., Frank, L.M., Smith, A.C., Brown, E.N., and Suzuki, W.A. (2003). Single neurons in the monkey hippocampus and learning of new associations. Science *300*, 1578–1581.

Wiskott, L., Rasch, M.J., and Kempermann, G. (2006). A functional hypothesis for adult hippocampal neurogenesis: avoidance of catastrophic interference in the dentate gyrus. Hippocampus *16*, 329–343.

Witter, M.P., Wouterlood, F.G., Naber, P.A., and Van Haeften, T. (2000a). Anatomical organization of the parahippocampal-hippocampal network. Annals of the New York Academy of Sciences *911*, 1–24.

Witter, M.P., Naber, P.A., Van Haeften, T., Machielsen, W.C., Rombouts, S.A., Barkhof, F., Scheltens, P., and Lopes da Silva, F.H. (2000b). Cortico-hippocampal communication by way of parallel parahippocampal-subicular pathways. Hippocampus *10*, 398–410.

Wixted, J.T., and Squire, L.R. (2010). The role of the human hippocampus in familiarity-based and recollection-based recognition memory. Behavioural Brain Research *215*, 197–208.

Wood, E.R., Dudchenko, P.A., and Eichenbaum, H. (1999). The global record of memory in hippocampal neuronal activity. Nature *397*, 613–616.

Wood, E.R., Dudchenko, P.A., Robitsek, R.J., and Eichenbaum, H. (2000). Hippocampal neurons encode information about different types of memory episodes occurring in the same location. Neuron *27*, 623–633.

Wyss, J.M., Swanson, L.W., and Cowan, W.M. (1979). A study of subcortical afferents to the hippocampal formation in the rat. Neuroscience *4*, 463–476.

Xu, W., and Südhof, T.C. (2013). A neural circuit for memory specificity and generalization. Science *339*, 1290–1295.

Yamashita, K., Hirose, S., Kunimatsu, A., Aoki, S., Chikazoe, J., Jimura, K., Masutani, Y., Abe, O., Ohtomo, K., Miyashita, Y., et al. (2009). Formation of long-term memory representation in human temporal cortex related to pictorial paired associates. The Journal of Neuroscience *29*, 10335–10340.

Yassa, M.A., and Stark, C.E.L. (2011). Pattern separation in the hippocampus. Trends in Neurosciences *34*, 515–525.

Yonelinas, A.P., Kroll, N.E.A., Quamme, J.R., Lazzara, M.M., Sauvé, M.-J., Widaman, K.F., and Knight, R.T. (2002). Effects of extensive temporal lobe damage or mild hypoxia on recollection and familiarity. Nature Neuroscience *5*, 1236–1241.

Yonelinas, A.P., Otten, L.J., Shaw, K.N., and Rugg, M.D. (2005). Separating the brain regions involved in recollection and familiarity in recognition memory. The Journal of Neuroscience *25*, 3002–3008.

Zalesak, M., and Heckers, S. (2009). The role of the hippocampus in transitive inference. Psychiatry Research *172*, 24–30.

Zeithamova, D., and Preston, A.R. (2010). Flexible memories: differential roles for medial temporal lobe and prefrontal cortex in cross-episode binding. The Journal of Neuroscience *30*, 14676–14684.

Zeithamova, D., Maddox, W.T., and Schnyer, D.M. (2008). Dissociable prototype learning systems: evidence from brain imaging and behavior. The Journal of Neuroscience *28*, 13194–13201.

Zeithamova, D., Dominick, A.L., and Preston, A.R. (2012). Hippocampal and ventral medial prefrontal activation during retrieval-mediated learning supports novel inference. Neuron *75*, 168–179.

Zhang, K., Ginzburg, I., McNaughton, B.L., and Sejnowski, T.J. (1998). Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. Journal of Neurophysiology *79*, 1017–1044.

Ziv, Y., Burns, L.D., Cocker, E.D., Hamel, E.O., Ghosh, K.K., Kitch, L.J., El Gamal, A., and Schnitzer, M.J. (2013). Long-term dynamics of CA1 hippocampal place codes. Nature Neuroscience *16*, 264–266.

Zola, S.M., Squire, L.R., Teng, E., Stefanacci, L., Buffalo, E.A., and Clark, R.E. (2000). Impaired recognition memory in monkeys after damage limited to the hippocampal region. The Journal of Neuroscience *20*, 451–463.

Zola-Morgan, S., and Squire, L.R. (1985). Medial temporal lesions in monkeys impair memory on a variety of tasks sensitive to human amnesia. Behavioral Neuroscience *99*, 22–34.

Zola-Morgan, S., Squire, L.R., and Amaral, D.G. (1986). Human amnesia and the medial temporal region: enduring memory impairment following a bilateral lesion limited to field CA1 of the hippocampus. The Journal of Neuroscience *6*, 2950–2967.

CURRICULUM VITAE

[redacted]

[redacted]

[redacted]