Title
Learning, Remembering, and Relating Sequences in the Hippocampus

Author
Singer, Annabelle C.

Publication Date
2009

Peer reviewed|Thesis/dissertation

Learning, Remembering, and Relating Sequences in the Hippocampus

by

Annabelle Singer

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Neuroscience

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, SAN FRANCISCO

[redacted]

Acknowledgements

I would like to thank my husband for his love and support throughout graduate school; my family and friends for reminding me how cool science is even when I was sick of it; and my lab for keeping me going. Mattias Karlsson taught be just about everything and the rest I learned from Steve Kim, Ana Nathe, Caleb Kemere, and Sen Cheng. Maya Chandru and Patricia Correia assisted with data collection and the members of the Frank Laboratory gave me scientific and editorial suggestions throughout my research. My thesis committee, Allison Doupe, Michael Brainard, Tricia Janak, and Jill Leutgeb have been an invaluable resource, asking all the right questions to push my research forward. Anne Smith developed the excellent adaptive algorithm we used for behavioral analysis. Howard Fields, Marianne Hafting-Fyhn, and Michael Stryker gave us fanstistic suggestions for our papers. The Keck Center, the Neuroscience Program, and the Physiology Department have been wonderful places to work. In particular, Louis Reichardt, Pat Veitch, and Carrie Huckaba have always been helpful and encouraging. And last, but not least, I thank Loren Frank without whom any of this would have been possible. Thank you for the endless mentorship, support, kindness, and enthusiasm. This work was supported by NIH grant MH080283, the John Merck Fund and the McKnight Foundation.

The text of Chapter 1 of this thesis is a reprint of the material as it originally appears in Neuron. The coauthor listed in this publication directed and supervised the research that forms the basis for the dissertation/thesis. Chapter 2 is primarily the work of Annabelle Singer with contributions from Mattias Karlsson, who collected and processed the 3 arm track data, Ana Nathe and Maggie Carr, who collected and analyzed the Multiple U track data, and Loren Frank who directed and supervised the research.

Learning, Remembering, and Relating Sequences in the Hippocampus

Annabelle C. Singer

As a crossroads between sensory inputs and long term memories, the hippocampus turns a plethora of information into concise episodes for us to remember. The hippocampus can employ different strategies to achieve this transformation. By selecting only notable experiences to transfer to long term memory storage, we can remember important experiences while forgetting the mundane. By encoding common principles among several experiences, we can remember appropriate general responses and predict future similar experiences. We considered ways the hippocampus might achieve these two possibilities by examining hippocampal activity while rats executed sequences for rewards.

Given that we must remember the experiences that lead to reward in order to exploit these rewards in the future, we asked if memory processes are enhanced by reward. In particular we examined hippocampal sharp wave-ripples (SWRs) because reactivation of previous experiences during SWRs is thought to be essential for event memory storage. We found that SWR activity increases when animals receive reward. This reward related SWR activity is further enhanced when animals have to learn new path-reward associations. Additionally, SWR activity reactivates neural patterns that occur as animals run to or from the reward. Because SWRs are implicated in memory consolidation, this enhanced SWR reactivation could be a mechanism to preferentially remember experiences associated with reward.

Furthermore, when navigating environments with many repeated elements, generalizing across elements can be advantageous to efficiently encode appropriate responses. Simultaneously, each element must also be differentiated from the others.

To study this, we then examined hippocampal activity as animals traversed environments with many repeated elements and had to distinguish between these elements to receive reward. We found that some hippocampal cells fire very similarly on multiple repeated elements, while other cells encode the elements differently. Cells that generalize across similar elements have correlated moment to moment activity, suggesting that they are part of functional ensembles. Furthermore, this generalizing / path equivalent activity increases as animals learn new relationships between repeated elements. This generalization across repeating elements could be a mechanism to extract general principles about related experiences.

Table of Contents

List of Tables

List of Figures and Illustrations

Introduction

We remember our lives in stories: episodes turn a jumble of information into sequences of places, events, actions, and outcomes. We use these stories to understand, to learn, and to predict our world and ourselves in it. The hippocampus is the brain's storyteller associating disparate information, like sounds, smells, expectations, actions, into episodes that can be later recalled (Squire, 1982). Sitting at a crossroads between incoming sensory information and long term memories, the hippocampus is essential for both memory formation, forming associations between incoming information, and memory consolidation, transferring those memories to the cortex for long term storage (Squire and Alvarez, 1995;Squire and Zola, 1996).

Not everything is remembered. We preferentially remember the "important" or "exciting" stories (Cahill and McGaugh, 1998;LaBar and Cabeza, 2006). While we can forget more mundane experiences, we remember what leads to positive or negative outcomes, like food or pain, presumably so that we can repeat or avoid these experiences in the future. Essential to the transition from memory formation to memory consolidation, the hippocampus seems an ideal place to modulate which episodes are stored for the long term and which are discarded. While much research has elucidated how we preferentially form conditioned responses to salient stimuli, like during fear conditioning, the mechanisms by which we preferentially remember some episodes over others is still unclear (Maren and Quirk, 2004). Episodic memory modulation could take place at either the memory formation stage or the memory consolidation stage and we will examine these possibilities.

An episode does not stand alone; it is formed and recalled in relation to other episodes. When similar experiences are formed into memories, we must separate them to remember each distinctly (Leutgeb et al., 2005c). But we can also generalize across similarities in these experiences, extracting and remembering general principles instead of storing similar information about each experience redundantly (Eichenbaum, 2000). These general principles can then be applied to new experiences. As a gateway through which all episodes must pass, the hippocampus is well situated to relate different episodes to each other. Indeed, the hippocampus is thought to merge information from related experiences to apply to new experiences (Eichenbaum, 2000). We will examine how the hippocampus balances the drive to separate and generalize when faced with experiences with many similar repeated elements.

Episodes and Sequences

An episode is, in many ways, a sequence: events, places, people, and our perspective on these things ordered in time and space. To grasp how we turn these distinct entities into episodes and episodic memories, we can examine how we learn and remember sequences. While we don't know if animals form episodic memories (Does Pongo remember last Christmas when he ate the roast ham off the table?), they do learn sequences, like sequences of places, actions, and stimuli. We can therefore investigate the neural mechanisms underlying sequence learning in animals to gain insight into the neural mechanisms of episodic memory formation and consolidation. In the research presented here, we consider neural activity in rats as the animals learn to execute sequences of places to receive a reward.

Hippocampal Activity and Sequence Learning

Several different types of hippocampal activity are thought to be involved in sequence learning and, perhaps, episodic memories. When animals run, neurons reliably fire in particular locations in space, called place fields, and the local field potential (LFP) oscillates at about 8 Hz, or theta (O'Keefe and Dostrovsky, 1971;McNaughton et al., 1983;O'Keefe, 1993). Place field activity during theta is theorized to underlie memory formation. Memory formation must occur rapidly, as we often form episodic memories from only a single experience, though additional experience can fortify memories. Similarly, place fields form rapidly in a new environment then undergo additional refinement with more experience (Frank et al., 2004;Karlsson and Frank, 2008)

Place field activity during theta can form organized sequences. As an animal runs from one location to another, a series of place cells fire such that individual spikes are ordered into sequences on short time scales. As an animal runs through a cell's place field, the neuron is more likely to fire progressively earlier and earlier in the phase of theta as the animal progresses through the place field (Skaggs et al., 1996;Harris et al., 2002;Dragoi and Buzsaki, 2006). Apply this principle to a population of neurons and, as the animal runs from one location to another, cells spike in precisely timed sequences (Figure 1).

This sequential activity is thought to be important because it organized neurons to fire in specific patterns on the short time scales required for plasticity. If cells fire in a particular order within tens of milliseconds of each other, the connections between neurons with neighboring spikes and receptive fields should strengthen according to

Hebbian mechanisms in which cells that fire together wire together (Hebb, 1949). If the connections are strengthened enough, then activity in the first cell could set off firing off the entire sequence. This, currently, is our best guess of how we encode sequences: neurons form receptive fields then form associations with each other creating long sequences. In this case, an animal traverses a series of places and a sequence of place fields is tied together such that one place field links to the next. We know that hippocampal neuronal can encode more general associations, not just places (Wiener et al., 1989;Young et al., 1994;Wirth et al., 2003). Therefore, a series of places, events, objects, and other associations could also be encoded by sequences of hippocampal neurons that "link" together.

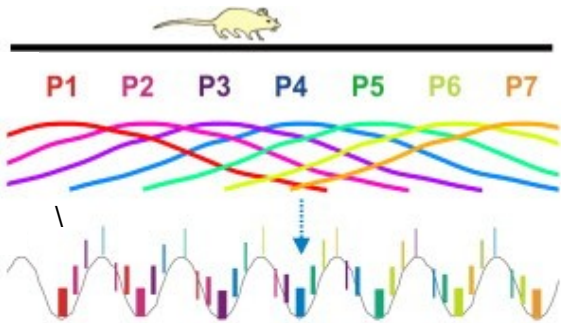

Adapted from Dragoi and Buzsaki 2006

Figure 1. Theta phase precession in the hippocampus. As a rat (top) runs, he crosses different cells' place fields (P1 through P7, middle). At the center of a cell's place field, the cell is most likely to fire at the trough of theta (P4). Theta is shown below and the probability of spiking during different phases of theta is represented by line thickness. As the animal progresses through the place field, the cell fires at earlier and earlier phases of theta. For cells with neighboring place fields, the spikes become ordered into precisely timed sequences.

These sequences of neural activity observed during theta are subsequently "replayed" during sharp-wave ripples (Fig. 2; (Buzsaki, 1986;Wilson and McNaughton, 1994;Skaggs and McNaughton, 1996;Sutherland and McNaughton, 2000;Foster and Wilson, 2006;Karlsson and Frank, 2009a) . Sharp wave ripples (SWR) are high frequency oscillations in the local field potential that correspond to bursts of population activity. SWRs are the corresponding bursts of activity are hypothesized to be involved

in memory consolidation. Memory consolidation occurs more slowly than memory formation, but it requires coordination between different brain regions because memories rapidly formed within the hippocampus must be transferred to the cortex for long term storage (Squire and Alvarez, 1995).

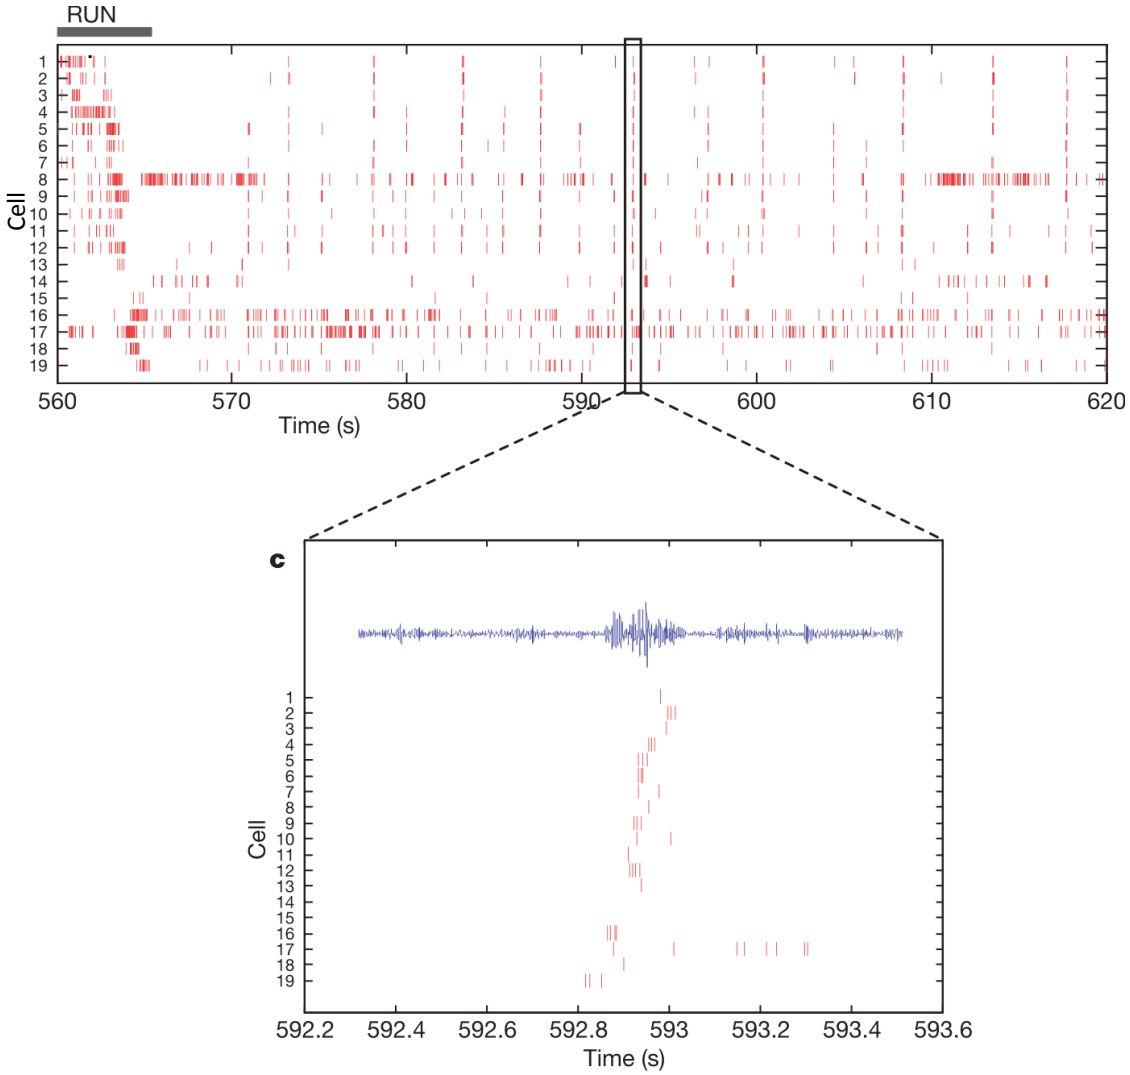

Adapted from Foster and Wilson 2006

Figure 2. SWR Replay. During behavior, sequences of cells fire spike trains (red lines) as the animals run through each cells' place field (top diagram on the left side). Cells are sorted in order of place field peak location, so that cells with earlier peaks are at the top. During pauses in behavior the same sequences are replayed during SWRs (bottom diagram). The blue line shows the LFP with a SWR (center), with simultaneously recorded spikes below. In this SWR, the cells fire in approximately reverse order to that seen during running.

Accordingly, SWR replay continues long after an experience is complete (Wilson and McNaughton, 1994;Karlsson and Frank, 2009a) and is coordinated across many brain regions (Sutherland and McNaughton, 2000;Pennartz et al., 2004;Ji and Wilson, 2007;Euston et al., 2007;Peyrache et al., 2009b). Indeed, recent research has shown that interrupting SWRs during sleep just after a memory task impairs later performance on that task, implicating SWRs in learning and memory (Girardeau et al., 2009).

SWRs occur not only during sleep, but also during pauses in behavior in waking and we will focus on the these SWRs during waking (Wilson and McNaughton, 1994;Foster and Wilson, 2006). During pauses in behavior, SWRs are enhanced in novel environments and by repeated passes through a familiar environment. These results reveal that SWR activity can be modulated by a variety of factors (Jackson et al., 2006;Cheng and Frank, 2008). The exact mechanisms of producing and modulating SWR activity is unknown but SWR activity may be modulated according to the need for memory consolidation of specific experiences.

Relating Different Sequences

Place cells and SWRs may allow us to learn and remember particular sequences, but we must also distinguish between similar sequences and associate related sequences. To distinguish between similar sequences, an entirely different set of place cells can encode each sequence. But in some cases, sequences should be related but distinguished, for instance when animals are in the same place but some things have changed. As animals run through the same location but come from or go to different destinations, place fields fire in the same place but at different rates (Figure 3, (Frank et al., 2000). This activity, termed prospective or retrospective coding, is thought

to encode not just the current location but the longer sequences of places the animal has or will traverse. In this sense, a place cell firing in a single location but at a different rate can signal two different but related sequences of places.

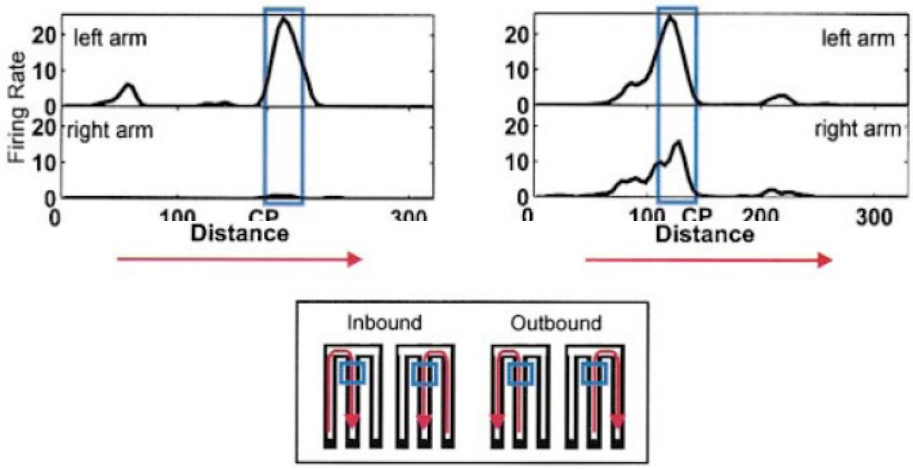

Adapted from Frank, Brown, and Wilson 2000

Figure 3. Trajectory Coding. A single CA1 cells fires much more when the animal is in the same location but coming from or going to the left arm (top row of top figure) than the right arm (bottom row of top figure). The top figure shows the linearized firing in hertz as a function of the animals position along the trajectory in centimeters. The track and the animal's trajectories through it are shown below. CP indicates the choice point, the point at the end of the middle arm where the animal had to select the next arm.

A similar pattern of activity has been observed if some cues in the environment change. When animals are in the same place but in different enclosures (e.g. the same cues in the distance but different local cues), place fields fire in the same location but at different peak rates (Figure 4, (Lever et al., 2002;Wills et al., 2005;Leutgeb et al., 2005a;Leutgeb et al., 2005b;Leutgeb et al., 2005c). This activity is called rate remapping and stands in contrast to global remapping, in which cells fire in different locations in different environments (Leutgeb et al., 2005b). In studies of rate remapping, animals foraged randomly through the environment so it is unknown how exact trajectories through space or trajectory coding affected firing rate. Nonetheless, both

trajectory coding and rate remapping have been proposed to encode different episodes or sequences in the same place.

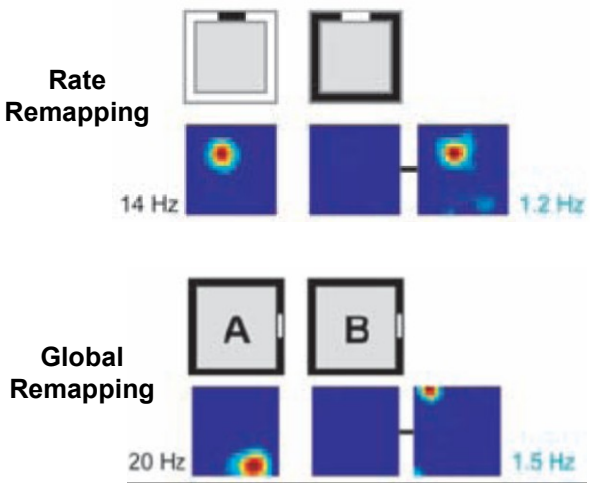

Adapted from Leutgeb et al 2005b

Figure 4. Rate Remapping and Global Remapping of Place Fields. In the same place, with a different enclosure, place cells in CA3 fire in the same place at different rates. (top figure). Above are overhead views of the different enclosures and below are rate maps for a single CA3 cell for the period when animals foraged in the enclosure. The left and center rate maps are scaled to a peak of 14 Hz, the peak rate in the left enclosure. The right rate map is scaled to 1.2 Hz, the peak rate in the right enclosure. In the same enclosure in different rooms (rooms A and B), place cells fire in different places (bottom figure).

Sequence Learning and CA3

Place fields, SWRs, and rate remapping can be found in several different hippocampal areas, but CA3 is postulated to play a central role in sequence learning. The hippocampus is often thought of as a single loop through a series of subregions, but it has a number of other connections (Figure 5). In particular CA3's recurrent collaterals, in which CA3 excitatory neurons synapse onto other CA3 excitatory neurons, allow direct connections between place cells in the same subregion. These recurrent connections may form the sequence activity that we think underlies memory formation. Furthermore, prior studies suggest that SWRs originate in CA3 (Csicsvari et al., 2000), perhaps also be due to these recurrent collaterals. CA3 receives input within the hippocampus from the dentate gyrus, CA3, and the entorhinal cortex and it sends output to CA1. CA3 also receives inputs from a number of regions outside the hippocampus:

the septum (shown to be involved in hippocampal theta oscillations), the VTA (implicated in reward encoding), the amygdala (engaged in processing emotions), the locus coeruleus (sends noradrenergic fibers to CA3), and the raphe nuclei (project serotonergic fibers to CA3; (Andersen et al., 2006). Furthermore, monoamine inputs are generally stronger in CA3 than CA1 (Andersen et al., 2006). Just upstream of CA3, the entorhinal cortex receives inputs from association cortices carrying a wide array of sensory information. Just downstream of CA3, CA1 sends projections to the prefrontal cortex and this connection is thought to be involved in long term memory storage. Given the central location of CA3 in the hippocampal network, its hypothesized role in sequence learning, and its strong modulatory inputs, we chose to examine CA3 activity as animal performed new and familiar sequences for reward.

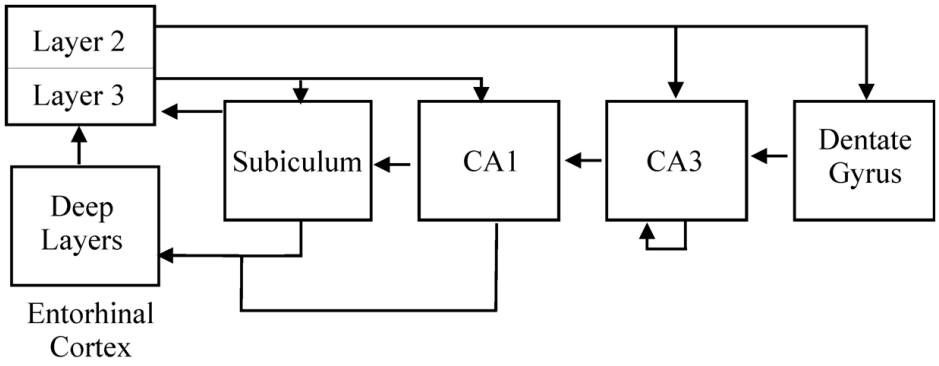

Figure 5. Circuit diagram of the hippocampus showing hippocampal subregions: the dentate gyrus, CA3, and CA1, and the subiculum and the entorhinal cortex. Arrows indicate excitatory connections of pyramidal cells. The entorhinal cortex receives inputs from and sends projections to parahippocampal and perirhinal cortices.

Many prior studies have examined hippocampal activity as animals searched for food randomly in an open field or as animals performed a well learned task but these approaches have a number of drawbacks for studying sequence learning. Because animal behavior is highly variable in open fields, isolating specific paths or sequences of

places is difficult. In studies using well learned tasks in linear environments, animals produce stereotyped behaviors, making sequences easier to isolate. But in these well learned tasks, sequence memories have already formed and undergone consolidation, so these memory processes cannot be observed in action. Furthermore, in most hippocampal studies, animals navigate through simple environments. While simplicity may be advantageous to answer some questions, in these simple environments there is limited opportunity to examine how similar or related sequences are encoded relative to each other.

These drawbacks motivated our study design. We examined hippocampal activity in CA3 as animals learned to perform new sequences in linear mazes that restricted the animals' path. Animals had to traverse and distinguish between several similar paths to correctly perform the task. Animals also had to learn to switch between sequences. We thereby were able to assess hippocampal activity as animal performed controlled sequential behaviors, continued to learn from the situation, and distinguished between similar paths.

Hypotheses

Experiences that lead to salient outcomes, like a reward, are more readily stored in memory. Although we might expect that rewarded outcomes would enhance memory storage, no neural mechanism has been found by which the outcome of an event enhances hippocampal memory formation for that event. Previous studies examining hippocampal responses to different outcomes have focused on place field activity. These reports found that the presence of reward or differences in motivational state can alter the firing rate or location of hippocampal place fields (Breese et al., 1989;Kobayashi

et al., 1997;Fyhn et al., 2002;Tabuchi et al., 2003;Holscher et al., 2003;Kennedy and Shapiro, 2009). Changes in visual cues have a similar effect on place cell firing (Hetherington and Shapiro, 1997), suggesting that reward can act like other sensory cues to alter the activity of place cells. Place field changes could signal the presence of something “interesting” when the animal is in the vicinity of the reward, but it is not clear how this activity would help the animal learn to navigate from distant locations to the reward.

We reasoned that, to learn and remember what experiences or paths lead to reward, the network must be able to encode associations between the path and the later outcome. While it is unclear how place fields might link a long path and an outcome, SWR reactivation is well suited to this task. Long paths that an animal traverses can be reactivated during SWRs (Davidson et al., 2009) and SWR reactivation is likely involved in memory consolidation (Buzsaki et al., 1992;Girardeau et al., 2009). We imagined that enhancing SWR reactivation in response to reward could be a mechanism to preferentially remember paths related to reward after the outcome of the path is known. Therefore, we hypothesized that reward would enhance SWR reactivation of paths related to that reward. We will discuss our findings in detail in Chapter 1.

While SWR reactivation may help us bind paths to their outcomes, understanding the general principles that predict reward could help animals navigate to new reward sites. For instance, if we understand how cities are organized into blocks, then we can better predict how to find a taxi or coffee shop. As described above, previous work has shown that in similar environments place cells rate remap (Lever et al., 2002;Wills et al., 2005;Leutgeb et al., 2005a;Leutgeb et al., 2005b). While some have argued that this rate remapping encodes different episodes in the same place, this sort of activity could reflect animal confusion. We theorized that rate remapping actually reflects generalizations across similar experiences.

We examined rate remapping between distinct but related paths as animals performed novel and familiar sequences in an environment with many repeating elements. Environments with many repeated elements carry both computational challenges, in distinguishing between similar elements, and advantages, in that the dimensions of one element predict the dimensions of another. We hypothesized that some place cells would show path equivalence, firing in the same locations on many different repeated elements of the environment, and that this path equivalence would increase as animal learned relationships between similar elements. We will discuss our findings in detail in Chapter 2.

Chapter 1

Rewarded Outcomes Enhance Reactivation of Experience in the Hippocampus

Abstract

Remembering experiences that lead to reward is essential for survival. The hippocampus is required for forming and storing memories of events and places, but the mechanisms that associate specific experiences with rewarding outcomes are not understood. Event memory storage is thought to depend on the reactivation of previous experiences during hippocampal sharp wave-ripples (SWRs). We used a novel sequence switching task that allowed us to examine the interaction between SWRs and reward. We compared SWR activity after animals traversed spatial trajectories and either received or did not receive a reward. Here we show that rat hippocampal CA3 principal cells are significantly more active during SWRs following receipt of reward. This SWR activity was further enhanced during learning and reactivated coherent elements of the paths associated with the reward location. This enhanced reactivation in response to reward could be a mechanism to bind rewarding outcomes to the experiences that precede them.

Introduction

How do we remember experiences that lead to reward? Although the hippocampus is required for storing memories of the places and events that make up these experiences (Squire, 1982), little is known about the mechanisms that associate

specific experiences with their outcomes. Studies in rodents examining hippocampal responses to different outcomes have generally focused on the presence or absence of a reward such as food or an escape platform in a watermaze. These reports analyzed place field activity, where hippocampal excitatory cells ("place cells") fire in particular locations in space during active exploration. These studies found that the presence of reward or differences in motivational state can alter the firing rate or location of hippocampal place fields (Breese et al., 1989;Kobayashi et al., 1997;Fyhn et al., 2002;Tabuchi et al., 2003;Holscher et al., 2003;Kennedy and Shapiro, 2009). As the presence or absence of visual cues has a similar effect on place cell firing (Hetherington and Shapiro, 1997), these studies suggest that reward can act like other sensory cues to alter the activity of place cells. Place field changes could signal the presence of something "interesting" when the animals is in the vicinity of the reward, but it is not clear how this activity would help the animal learn to navigate from distant locations to the reward.

Place cells are also active during high frequency network oscillations called sharp wave ripples (SWRs), in which sequences of cells activated during movement are reactivated (Skaggs and McNaughton, 1996;Lee and Wilson, 2002;Foster and Wilson, 2006;Ji and Wilson, 2007;Diba and Buzsaki, 2007). SWRs occur largely during sleep and awake immobility (Buzsaki et al., 1983;O'Neill et al., 2006;Cheng and Frank, 2008) and are thought to be important for spatial learning and memory formation (Redish and Touretzky, 1998;Redish, 1999;Samsonovich and Ascoli, 2005;Nakashiba et al., 2008). Hippocampal reactivation allows events that are experienced relatively briefly to be replayed over and over again on a short timescale compatible with synaptic plasticity (Buzsaki, 1986;Wilson and McNaughton, 1994;Sutherland and McNaughton, 2000). In particular, reactivation during pauses in waking behavior frequently results in the sequential activity of place cells active on paths to or from the animal's current location

(Foster and Wilson, 2006;Diba and Buzsaki, 2007;Davidson et al., 2009;Karlsson and Frank, 2009b). SWR reactivation can occur after traversing a path, potentially allowing the animal to learn the relationship between the path and its outcome (Johnson and Redish, 2005;Foster and Wilson, 2006;Diba and Buzsaki, 2007).

Given that rewarded events are often well remembered, we would predict that rewarded outcomes would modulate memory storage mechanisms for the associated events. In particular, we would expect that a rewarding outcome would facilitate reactivation of the experience that led to that outcome. While a recent report documented outcome related activity in the primate hippocampus (Wirth et al., 2009), the relationship between reward and reactivation has not been investigated. We therefore asked whether receipt of reward affects reactivation of place cells in the hippocampus.

Results

SWR events generally originate in hippocampal area CA3 (Csicsvari et al., 2000), so we focused our studies on this area. We recorded from principal neurons while animals learned to switch between two spatial sequences in response to changing reward contingencies (Fig. 1a,b). This sequence switching task allowed us to compare trials when the animals performed the same behavioral sequence and either did or did not receive reward (Fig. 1c). Animals first learned a spatial alternation sequence (S1) to criterion and then learned to switch between this sequence and a new sequence (S2). This task has four features that make it appropriate for examining the effect of reward on hippocampal memory processing. First, the rapid learning of the initial sequence requires the hippocampus (Kim and Frank, 2009) and the hippocampus is required for flexibly changing behavior in response to changing reward contingencies (Hsiao and

Isaacson, 1971;Hirsh et al., 1978;Ainge et al., 2007a). Second, because reward contingencies change during each run session, this task provides an adequate number of unrewarded trials to allow us to compare neural activity during rewarded and unrewarded trials. Third, the presence or absence of reward drives ongoing behavior. This is in contrast to tasks where reward is randomly omitted (Tabuchi et al., 2003) and animals must learn to behave continuously regardless of each trial's outcome. Fourth, in this task the animal learns to switch between sequences in a familiar environment, allowing us to control for the effects of spatial novelty on SWR activity (Cheng and Frank, 2008;Karlsson and Frank, 2008).

All six arms were open in both sequences and animals received a liquid chocolate reward at the end of an arm if that arm was the next correct arm in the sequence. No experimenter delivered cues indicated whether a trial was or was not rewarded other than the presence or absence of the reward itself. Here, as in other studies of reward (Tremblay et al., 1998;Fiorillo et al., 2008), receipt of reward consists of the entire reward or lack of reward experience including sensation, consumption and the affective states induced by reward presence or absence. Our goal was to determine if this reward experience changed memory processing in the hippocampus.

We focused our analysis on the sequence switching phase of the task. At the beginning of each session animals were placed in the home arm of the to-be-rewarded sequence. We distinguished between an accurate response, where the animal made a choice consistent with the rules of either S1 or S2, and a rewarded response where an animal made a choice consistent with the rules of the currently rewarded sequence. This allowed us to quantify the probability of an accurate response for both sequences on every trial (Fig. 1b). We found that when animals were first placed in the home arm of S2 they immediately performed the previously rewarded S1. Thus, animals used environmental cues and a track-based reference frame to perform the task, rather than

remembering a series of right or left turns based on their body reference frame. After executing S1 for several trials animals changed their behavior and eventually learned to perform S2. For the quantification of behavior used for our analyses we employed a dynamic state-space algorithm (Smith et al., 2004) to estimate the likelihood of an accurate response on each trial (Fig. S1). This algorithm allows us to compute confidence intervals which were essential for defining periods when one sequence was performed significantly more accurately than the other.

We recorded from three animals during the sequence switching phase of the task (n = 270 single neurons, Fig. S2 for histology and LFP in CA3). To determine whether hippocampal SWR activity varies with receipt of reward, we examined SWR activity when animals stopped at the well and were either rewarded or not rewarded. To identify SWRs, we filtered the local field potential from 150 - 250Hz, determined an envelope by Hilbert transform, and detected when the envelope amplitude exceeded 3 standard deviations from baseline for at least 15 msec (Cheng and Frank, 2008). We first restricted our analyses to putative excitatory neurons that were active in restricted spatial regions (place fields) on the track.

SWR activity on rewarded and unrewarded trials

We found that cells with place fields on the track (n = 107) were much more likely to be active during SWRs at the well (wSWRs) on rewarded trials than unrewarded trials (Fig. 2a, 2b, $p < 10^{-10}$, Fig. S3; all statistical tests were rank sum tests and n = 107 for activation probability per pass or per wSWR unless otherwise noted). This enhanced activity was associated with two differences in neural responses between rewarded and unrewarded trials. First, there were more wSWRs per unit time on rewarded trials (Fig. 2c, $p < 10^{-10}$, n = 3945 rewarded trials, n = 709 unrewarded for SWR rate unless otherwise noted). Second, we examined each wSWR individually to control for the

greater rate of wSWRs and found that place cells were more likely to be active in any given wSWR on rewarded trials (Fig. 2d, S4a, $p < 10^{-10}$). As expected given the greater activation probability per wSWR, the average number of spikes each neuron fired per wSWR, the mean firing rate in wSWRs and the proportion of cells active per wSWR were also larger in rewarded than unrewarded trials (Fig. 2e, 2f, and 2g, p's $< 10^{-4}$, n = 107 cells; prop. active n = 4427 SWRs following reward, n = 238 following no reward).

The increase in wSWR rate and activation probability within individual wSWRs accounted for a four-fold increase in total activation probability on rewarded trials. There was an additional two-fold increase that resulted from longer time spent at the well on rewarded trials ($p < 10^{-10}$, n = 3945 rewarded trials, n = 709 unrewarded). We controlled for this time difference by truncating the time on each rewarded trial to match the duration of immobility on a randomly selected unrewarded trial. Cells were still significantly more likely to be active on truncated rewarded trials than unrewarded trials (Fig. 2h, $p < 10^{-10}$). Similarly, the differences in wSWR rate and activation probability per wSWR remained significantly higher on truncated rewarded trials ($p < 10^{-10}$ and $p < 0.01$, respectively).

SWR activity during learning of new reward contingencies

If this enhanced reactivation is important for learning about experiences that lead to reward, we would expect stronger reactivation when the animal learns new path-reward associations. Consistent with this prediction, we found that wSWR rate, activation probability per wSWR, and the proportion of cells active per wSWR were higher when animals were first exposed to S2. We examined rewarded trials during periods when animals performed the rewarded sequence significantly more that the unrewarded sequence. On the first day of exposure to S2, wSWR rate was higher on rewarded trials during the first session of S2 than during the rewarded trials in the

previous S1 session (Fig. 3a, $p < 0.0005$, n = 471 S1 trials, n = 140 S2 trials). Similarly, both activation probability and proportion of cells active per wSWR were significantly greater in S2 than S1 on day 1. (Fig. 3b,c; activation prob. $p < 0.05$, Student's paired t-test, n = 14 cells; prop. active, $p < 10^{-4}$, S1 n = 428 SWRs, S2 n = 147). The increase in activation probability and proportion of cells active per wSWR in S2 was above and beyond the overall increases in rewarded trials compared to unrewarded.

Finally, as we would expect if these differences were related to learning a novel sequence, there were no significant differences by the third day of exposure to S2, when S2 was more familiar (Fig. 3b,c; p's > 0.1; SWR rate: S1 n = 353 trials, S2 n = 131; activation prob. n = 14 cells; prop. active S1 n = 455 SWRs, S2 n = 288). While these findings demonstrate clear differences between sessions and across days, the wSWR rate and the proportion of cells active per pass was relatively stable within individual sessions (Fig. S4b,c). Therefore this enhanced wSWR activity on S2 does not simply reflect sensitivity the changes in reward contingencies that occur at the beginning of each session. Furthermore, this increase in wSWR activity in S2 cannot be due to the presence of consumatory behaviors, as the animal consumed reward in both sequences. Taken together, these results demonstrate that reward related SWR activity is further enhanced when animal must learn new path-reward associations.

Reactivation during SWRs on rewarded trials

The increase in wSWR activity on rewarded trials was not simply encoding the presence of reward; instead wSWR activity reflected structured reactivation of neurons active on the paths associated with the rewarded location. Previous reports of reactivation during pauses in behavior have documented increased coordinated activity of pairs of CA1 place cells during SWRs (Kudrimoti et al., 1999;Cheng and Frank, 2008) as well as sequential replay of CA3 and CA1 place cells during SWRs (Foster and

Wilson, 2006;Diba and Buzsaki, 2007;Csicsvari et al., 2007;Davidson et al., 2009;Karlsson and Frank, 2009b). We therefore asked whether the wSWR activity we saw was specific to particular place cells or pairs of cells active on the track. We computed the probability that a neuron active during a wSWR at the well was also active during the run period leading up to or away from that well. Those probabilities were both significantly higher than the probability that the neuron was active during a randomly selected run period (see methods, Fig. 4a, $p < 10^{-10}$, n = 3945 trials).

If wSWR activity resulting from reward specifically reactivates meaningful patterns of place cell activity, we would also expect greater reactivation of cells with place fields on the track. We clustered cells during both run and rest sessions, allowing us to identify neurons that were active in the rest box but did not have place fields on the track. We found that cells with place fields on the track were much more likely to be activated than cells without (Fig. 4b, ranksum test, $p < 10^{-5}$). Further, the increase in activation probability per wSWR from unrewarded to rewarded trials was much larger for cells with place fields on the track, when measured as either the average across all cells ($p < 10^{-5}$) or as the increase within individual cells (Fig. 4c, $p < 10^{-10}$). We confirmed these effects with a 2-way ANOVA and found main effects of reward ($F(1,454) = 56.71$, $p < 10^{-5}$) and of the presence of a place field ($F(1,454) = 94.55$, $p < 10^{-5}$). There was also a highly significant interaction ($F(1,455) = 27.9$, $p < 10^{-5}$), due to a larger increase for cells with place fields. See the Supplementary Methods for further discussion of the measurements.

We then examined the place field locations of the cells active during wSWRs. One could imagine that SWRs preferentially reactivate cells that were most recently active on the run to the well. If so, we would expect cells that were active closest in space or time to the reward well would be more likely to fire during wSWRs. We found no such bias. Both the general population and cells that were active during wSWRs had

place field activity that congregated at the turns of the track. To visualize the spatial distribution of run period activity, we plotted firing rate maps constructed from all the spikes for every cell that was active during wSWRs at a particular well in a single session (Fig. 4d). The populations of cells that were active during wSWRs were active at multiple locations on the track and they tended to fire more during the turn leading to the reward well, consistent with previous reports that place fields congregate around relevant cues (Hetherington and Shapiro, 1997). We also noted that some cells had multiple place fields, perhaps accounting for the activity on more distant paths.

To quantify the spatial distribution of run period activity preceding wSWRs, we calculated the location of the peak of the occupancy normalized firing rate during the run period for all cells and only for cells that were active during wSWRs. The distributions for the entire population and for the cells that fired during wSWRs on rewarded trials were very similar (Fig. 4e and S5a,b; linear regression R^2's > 0.6, p's < 10^{-4}, n = 21 spatial bins). A complementary analysis examining the time between spiking and wSWR activity also failed to show a temporal bias for cells with place fields closer to the reward locations (Fig. S5c-f).

If reward enhances the reactivation of experiences associated with reward, we would predict that cells that fired together during the run would also fire together during wSWRs on rewarded trials. We computed the coactivity across cells pairs and found that pairs of place cells were more than twice as likely to fire together during wSWRs on rewarded than unrewarded trials (Fig. 5a, p < 10^{-10}, n = 498 pairs). This coactivation probability per wSWR was higher for cells with greater overlap between their place fields (R^2 = 0.1192, p < 10^{-10} for rewarded trials and R^2 = 0.0213, p < 0.005 for unrewarded trials, n = 498 pairs).

The coactivity of cells pairs on rewarded trials was greater than expected given the firing of the individual cells. We found a significant correlation between place field

overlap and the extent to which cells were more coactive per wSWR than expected by chance (Cheng and Frank, 2008);Fig. 5b, Fig. S6a,b, $R^2 = 0.0846$, $p < 10^{-5}$, n = 412 pairs on rewarded trials and $R^2 = 0.0335$, $p > 0.4$, n = 20 pairs on unrewarded trials; note that the measure is only defined if each cell is active at least once in a wSWR). Increased coactivation was present only in SWRs: there was no significant difference in coactivation during the run up to the food well for rewarded and unrewarded trials (see methods, Fig. 5c, $p > 0.34$, 95% confidence intervals for unrewarded: 0, 0.0191 and rewarded: 0, 0.0186, n = 498 pairs). We found similar results when we repeated these analyses by measuring the joint-surprise (Grun et al., 2002;Pazienti and Grun, 2006) of cell pairs' spiking (SFig. 6c, 6d) and by performing these analyses including only cells that were recorded from different tetrodes (SFig. 6e-h). Joint surprise also measures the extent to which two cells are more coactive than expected by chance based on the individual cells' firing, and has been used to help control for spike sorting errors. Taken together, these results confirm that SWR activity reactivated elements of an experience and that receipt of reward enhances this reactivation.

Finally, we found that reactivation in CA3 was consistent with ordered replay of place cell sequences observed in previous studies (Foster and Wilson, 2006;Diba and Buzsaki, 2007;Csicsvari et al., 2007;Karlsson and Frank, 2009b). Ordered replay implies that for two cells, the further apart their place fields, the longer the time between their spikes during SWRs. We therefore took each pair of cells and measured the extent to which the distance between their place field peaks predicted the inter-cell interspike intervals during wSWRs (Karlsson and Frank, 2009b); see methods). Considering only rewarded trials, we found a highly significant relationship, consistent with ordered replay (Fig. 5d, $R^2 = 0.0913$, $p < 10^{-10}$, n = 426 ISIs).

Previous reports of "reverse replay" argued that when cells were reactivated in an order opposite to that on the path to a reward, this pattern of activity could help the

animal learn the sequence of locations leading to the reward (Foster and Wilson, 2006;Diba and Buzsaki, 2007). We therefore asked whether the reactivation we saw was consistent with reverse replay. We found that in 337 of 495 cases (68.1%) where two cells were active in a wSWR, the order of activity was opposite that seen during the run (proportion > 0.5, $p < 10^{-10}$). At the same time, many cells were active at a given location in both directions of motion on the track, so these events are consistent with both replay of the path to the reward and "preplay" of paths from the reward. Nonetheless, these findings demonstrate that CA3 SWR activity following receipt of a reward reactivates coherent elements of the animal's path associated with the rewarded location.

SWR activity and place field activity across behavioral conditions

Finally, we carried out a series of controls to determine whether increased wSWR activity on rewarded trials could be attributed to activity during the run to the well, the relative timing of rewarded and unrewarded trials, behavioral variability, the behavioral sequence the animals executed, reward expectation, or wSWR properties. In no case were we able to identify a difference in these factors that could explain the differences between neural activity during wSWRs on rewarded and unrewarded trials. Here we discuss two of these controls; the rest can be found in the Supplementary Results, and SFig. 7 and 8.

We first asked whether the differences in rewarded trials could be explained by the specific spatial sequence the animal performed. We compared times when animals performed S1 when it either was or was not rewarded. For this analysis we took advantage of the fact that when reward contingencies changed to reward S2, animals continued to perform S1. We therefore compared periods when animals performed S1 accurately when reward was either delivered (e.g. Fig. 1c far left box) or omitted (e.g. Fig. 1c middle left box). Here the behavioral sequence was identical and only the receipt

of reward varied. Even when S1 was performed accurately, activation probability per wSWR was only elevated when animals received reward (Fig. 6a, Fig. S7i-l, $p < 10^{-5}$, n = 106 cells recorded when there were wSWRs following reward, n = 46 unrewarded). We found similar results if we truncated the time spent at the well on these rewarded trials to match the unrewarded trials ($p < 10^{-4}$). This shows that the enhanced wSWR activity was due to the receipt of reward, not the behavioral demands of the task.

These results also suggest that expectation of a reward did not lead to increased activation per wSWR on rewarded trials. Based on the history of rewards in S1, the animal would expect a reward for performing S1. If SWR activity was due to reward expectation, SWR activity would be higher when the animal performed S1 when it was no longer rewarded following a session where it was rewarded. However, SWR activity is higher on rewarded trials and lower on unrewarded trials regardless of the prior history of reward.

We similarly found that in cases where a reward was delivered at a previously unrewarded location, the presence of reward rather than the history of reward was the best predictor of wSWR activity. We examined trials on the first exposure to S2 when the animal received a reward in arm E, a previously unrewarded arm (animal 1:15 trials; animal 2: 4 trials; animal 3: 19 trials). Reward on these trials on arm E was initially unexpected given the animals' previous experience with S1. The number of unexpected reward trials included was similar to that used in other studies of reward expectation (Fiorillo et al., 2008;Schultz et al., 1992;Tremblay et al., 1998). Neurons were significantly more active on these unexpectedly rewarded trials than unrewarded trials (Fig. 6b, $p < 0.04$, n = 14 cells).

Discussion

We have shown that receipt of reward enhances SWR reactivation and this reactivation is further enhanced when new reward contingencies must be learned. The activation probability per SWR, the number of SWRs per time, the number of spikes per SWR, and the mean spike rate during SWRs were significantly higher at the end of rewarded trials than unrewarded trials. This increase in SWR activity only occurred after animals reached the reward well. Furthermore, on rewarded trials occurring when the animal was learning a new sequence (S2), the rate of SWRs and the probability that cells were active in each wSWR was even higher than during rewarded trials associated with the familiar S1. In this case the animal was receiving and consuming reward in both conditions, demonstrating that the increase in neural activity could not be explained solely as a result of the presence of consummatory behaviors. We also showed that spiking during wSWRs reactivates coherent elements of the experiences that are associated with the paths to and from the rewarded location. We performed an extensive series of control analyses examining differences in activity on the path to the reward, behavioral differences at the reward location and reward expectation. None of these potential confounds could explain the enhanced activation of CA3 neurons during wSWRs following receipt of reward, indicating that receipt of reward is a key determinant of wSWR rate and firing during wSWRs.

Our results are distinct from previous demonstrations of an interaction between reward and hippocampal activity. Previous studies found that place fields can change in response to reward (Breese et al., 1989;Kobayashi et al., 1997;Tabuchi et al., 2003;Holscher et al., 2003). Our findings, in contrast, indicate that reward plays a special role in modulating the reactivation of cells associated with recent experiences. Similarly, the increase in reactivation we saw is distinct from observations of outcome

selectivity in primate hippocampus (Wirth et al., 2009). That paper reported that hippocampal neuronal activity between trials was related to whether the animal made a correct or incorrect choice on the previous trial, irrespective of the specific stimuli the animal experienced during that trial. Our findings suggest that in the rodent hippocampus, activity following a reward specifically relates to the sequence of locations the animal traversed on the way to the reward. Finally, the enhanced reactivation we report is also distinct from reward "signals" in which single cells encode aspects of reward like reward expectation or reward prediction error as observed in several other brain regions (Schultz, 2000). Instead, enhanced SWR reactivation following reward is better understood as reactivating patterns of activity that reflect experiences associated with the reward location.

Mechanisms of SWR Reactivation

Previous findings have suggested a simple model whereby reactivation merely reflects recent activity within the hippocampal network. These studies have shown that reactivation during sleep reflects the structure of previous awake experience (Wilson and McNaughton, 1994), and the amount that two cells fire together during SWRs in sleep depends on the amount that the two cells fire together during prior experience(O'Neill et al., 2008). Similarly, repeated and regular traversals of the same path lead to increases in SWR rate and reactivation in familiar environments(Jackson et al., 2006).

Our results demonstrate that this simple model is not sufficient: we found an increase in SWR reactivation when animals are rewarded even though there is no increase in activity during the preceding run period. In addition, there was no apparent bias in which cells along the trajectory to the well were reactivated, indicating that on short timescales, the timing of place cell firing relative to the SWR did not determine the strength of reactivation. Finally, pairs of cells were more coactive during SWRs at the

well on rewarded trials, but there was no difference in coactivity during the run period of rewarded and unrewarded trials. Thus, reactivation is not simply a reflection of recent activity.

Instead, our results argue for a more complex model where an event's outcome modulates the strength of reactivation. We found that cells both with and without place fields were more likely to be activated during wSWRs on rewarded than unrewarded trials. These observations indicate that reward increases the likelihood of reactivation for all cells. At the same time, cells with place fields were much more likely to be reactivated than cells without place fields, and cells active on paths associated with the reward location were most likely to be reactivated. Therefore, the specific spatial sequence the animal traversed may strongly influence which cells will be active during SWRs, while outcomes like reward may modulate the amount and strength of reactivation.

Reward Enhanced Reactivation and Learning

We found that the rate of SWRs and the likelihood that cells would be active within SWRs increased when animals had to learn new path-reward associations, suggesting that SWR reactivation contributes to learning. More specifically, SWR reactivation is well suited to help the animal learn the paths that lead to reward. Unlike place field activity on the path up to the reward, activity in SWRs can activate specific patterns of place cells after the outcome of traversing the path is known. We found evidence for sequential activation in pairs of neurons, and other reports have established that these reactivation events frequently involve replay of entire paths along the track (Foster and Wilson, 2006;Diba and Buzsaki, 2007;Davidson et al., 2009;Karlsson and Frank, 2009b).

We propose two possible mechanisms by which this SWR reactivation could facilitate learning the paths that lead to reward. First, the enhanced reactivation of

rewarded paths could strengthen representations in neocortical areas of rewarded paths over unrewarded paths. Later, when animals are selecting between several possible paths, the rewarded paths could outcompete unrewarded paths in the decision process. Alternately, the enhanced SWR reactivation could facilitate an association between a path and the outcome of that path creating a set of path-outcome associations. Because SWR reactivation occurs after the path is complete and can be coherent with activity in neocortical and striatal areas (Chrobak and Buzsaki, 1996;Ji and Wilson, 2007;Wierzynski et al., 2009;Lansink et al., 2009;Peyrache et al., 2009a), reactivation could help link a specific path with reward information encoded in other brain regions. As place fields were commonly found around turns when animals have to make arm choices, reactivation of place cells could lead to associations between reward outcome and the activity related to choices the animal made to reach the reward.

In this latter case in which SWR activity facilitates the formation of path-outcome associations, we might wonder why reactivation occurs more on rewarded than unrewarded trials when a lack of reward is also informative. In our task, as in natural foraging, there are generally many more paths that do not lead to reward than paths that do lead to reward. Thus, it may be advantageous to preferentially encode the relatively small number of path-reward associations as compared to the large number of path-no reward associations. Indeed, when reward contingencies change and animals encountered no reward where they previously were rewarded, the lack of reward was not so significant that animals immediately changed their behavior. Instead, they persisted in performing the unrewarded paths for many trials, indicating that the presence of reward on that path in the past may continue to influence behavior despite the current lack of reward.

Overall, our results demonstrate a new link between reward and the reactivation of recent experience. This reward-related reactivation may be a mechanism to learn and remember experiences that lead to reward.

Methods

Data Collection

Male Long-Evans rats were handled and food deprived to 85-90% of baseline weight. Animals were initially trained to run back and forth on a linear track for liquid chocolate reward delivered in food wells at the ends of the track. Linear track pretraining took place in a different room from the recording room. One of the animals was pretrained on S1 (Fig. 1a) in the recording room, while two animals were not exposed to the behavioral task until recording began. Following pretraining animals were implanted with a microdrive array containing 16 independently movable tetrodes targeting CA3 (-3.6 mm AP; 3.4mm L) using previously described methods (Frank et al., 2004). Over the next 7 – 10 days tetrodes were lowered first to CA1 and then to CA3. Details on data collection can be found in Karlsson and Frank (2008). CA3 was identified by depth and the characteristic EEG waveforms on each recording tetrode. Electrode positions were additionally confirmed by histology. For one animal, electrode lesions were made at the end of each tetrode and later confirmed to be in the CA3 pyramidal cell layer (Fig. S2a). For two animals, the microdrive fell off before lesions could be made. In these animals we were able to confirm that the implant site was over dorsal CA3 and that the depths were consistent with CA3 recordings (Fig. S2b,c). Furthermore the EEG signatures characteristic of CA3 were similar in all animals (Fig. S2d), although it is possible that a small number of cells were recorded from the dentate gyrus. For all animals a reference tetrode was positioned in the corpus callosum. All neural signals were recorded relative to that reference to eliminate muscle artifacts from the recordings.

Behavior

During recordings animals were rewarded with liquid chocolate for performing the behavioral paradigm shown in Figure 1. The track included 4 sequence arms, B C D and E, and one extra arm on each end (A and F). Arms were separated by vertical walls (0.6 cm thick, 24 cm tall and 81 cm long). Distal cues were visible above these walls, at either end of each arm, and along the straight section connecting different arms. Circles indicate food wells where animals received reward in arms B through E. Colored arrows indicate trajectories included in S1 (blue) and S2 (red).

The task consists of two rules. First, a visit to the home arm (arm C is S1 and arm D in S2) was rewarded when the animal came from any other arm (inbound trajectories). Second, a visit to an arm adjacent to the home arm was rewarded when the animal came from the home arm after having previously visited the opposite adjacent arm (outbound trajectories). Consecutive visits to the same food well were never rewarded. Together, these rules defined a correct cyclical sequence of food-well visits (Fig. 1a): right, center, left, center, right, center, left, center, etc (Frank et al., 2000;Kim and Frank, 2009). The inbound trajectories require the animal to return to a single rewarded location, the home arm, from any other arm. The outbound trajectories require the animal to remember which arm he just came from. For the first outbound trajectory at the start of each session, the animal was rewarded for visiting either home-adjacent arm. If the animal visited an arm not included in the rewarded sequence (e.g. arm A, E or F for S1), the animal was rewarded for returning to the home arm. On the following outbound trajectory the animal was rewarded for visiting either home-adjacent arm. We have shown that during the initial learning of the task, animals learn the inbound component first and then learn to alternate on outbound trajectories (Kim and Frank, 2009). Therefore, once animals learn to perform the outbound trajectories with high

accuracy they are generally performing the entire sequence accurately. Note that we use the terms trajectory and path interchangeably throughout the manuscript.

During each run session the animal was placed in the home arm of the to-be-rewarded sequence (arm C for Sequence 1 and arm D for Sequence 2). Each run session was between 20 and 30 minutes long; one animal performed two sessions and two animals performed three sessions per day. Thirty to forty minute rest sessions in a high walled box in the same room preceded and followed each run session. Once the animal performed S1 with 80% accuracy, measured across a run session, or had 6 full days of training and was above 75% accurate, the sequence switching phase of the task commenced.

On the first day of sequence switching animals first performed one session where S1 was rewarded. Then in the second session, reward contingencies changed such that S2 was rewarded. All subsequent sessions alternated between rewarding S1 and S2 within each day (see Fig. S1 for details). Recording continued throughout the rest and run sessions.

Reward was delivered via an air pressure / solenoid system and was triggered via key press on a keyboard. The experimenter's back was to the animal such that the experimenter was between the animal and the keyboard during the experiment. The experimenter triggered reward release before the animal reached the reward well, and on correct outbound trials from the center arm the experimenter generally triggered both the reward in the outer arm and the subsequent center arm (inbound) reward. Thus the audible solenoid click occurred before the animal stopped at the well and there was no consistent temporal relationship between solenoid clicks and reward delivery. In rare cases reward was triggered just as the animal reached the well. Excluding these trials had no effect on the rewarded / unrewarded differences.

We distinguished between “accurate” responses that were consistent with the rules of S1 or S2, and rewarded responses. This allowed us to score behavior according to the rules of both sequences simultaneously. For Fig. 1 we used a 20 trial moving average applied to all trials to illustrate the behavior, but as this moving average does not provide confidence bounds, we also used a dynamic state-space smoothing algorithm (Smith et al., 2004;Smith et al., 2007) to estimate the animals’ probability of an accurate response for each sequence on each trial and to compute confidence intervals for the estimated probability. For the algorithm we focused on outbound passes because an outbound trajectory could be correct for S1 or S2 but not both. We scored outbound trajectories from the home arm of a sequence as accurate or inaccurate separately for S1 and S2. The estimated probability distribution produced by the algorithm was taken from the end of one run session and used as a starting value for the estimation of the next run session. This corresponds to the assumption that the animal began each session with some information from the previous session but also allows the learning state to “jump” if the animal behaves very differently at the beginning of the next session.

On the basis of these estimates we defined three stages of behavioral performance. First, animals performed the unrewarded sequence significantly more accurately than the rewarded sequence, as occurs when the reward contingences are first changed. In this period the mode of the estimated probability accurate response distribution for the newly unrewarded sequence is greater than the 95% confidence bounds of estimated probability accurate response of the rewarded sequence. Then, animals traversed a variety of spatial trajectories between food wells, and neither sequence was performed significantly more accurately than the other. In this the period neither mode is greater than the confidence bounds of the other sequence. Finally, animals performed the rewarded sequence significantly more accurately than the unrewarded sequence. In this period the mode for the rewarded sequence is greater

than the confidence bounds of the unrewarded sequence. We used these stages to compare accurate performance of S1 when S1 was rewarded to accurate performance of S1 when S2 was rewarded. We also calculated behavioral entropy as described previously(Jackson et al., 2006) so that we could identify any differences in behavior between rewarded and unrewarded trials (see Supplementary Methods and Results).

Data Processing

Only well isolated cells with tightly clustered spikes and clear refractory periods were included. Cells were clustered throughout run and rest periods, allowing us to identify cells that were active but did not have place fields on the track. As our results involved comparisons of spiking from the same clusters within a day, poor clustering cannot account for the effects we observed. We did not attempt to match cells across days, so in some cases the same cell may have been recorded across multiple days. All analyses were restricted to putative principal neurons (n = 100, 42, and 128 for animals 1,2 and 3 respectively;(Fox and Ranck, 1981;Frank et al., 2001). To identify cells with place fields we calculated the 'linearized' activity of each cell. Only times when animals were running forward at least 2 cm/sec were included. The behavioral data were separated into different spatial trajectories (e.g. A to B, B to A, B to C, …) and the animal's linear position was measured as the distance in cm along the track from the reward site on the start arm. We calculated occupancy normalized firing rate maps using 2 cm spatial bins smoothed with a 4 cm standard deviation Gaussian curve with a total extent of 20 cm, excluding bins with less than 0.2 sec/bin occupancy. Cells with a mean rate greater than 0.1 Hz and a peak spatial rate greater than 3 Hz were considered to have a place field on the track. Place field overlap was calculated according to a previously established method (Battaglia et al., 2004).

SWRs were identified as described previously (Cheng and Frank, 2008). Briefly, LFPs were recorded from one channel of each tetrode. On each day, the tetrode with the largest number of isolated neurons was used for SWR detection. The LFP signal was band pass filtered between 150-250 Hz and an envelope was determined by Hilbert transform. SWR events were detected if the envelope exceeded a threshold of mean + 3 stdev for at least 15 ms. Events included times around the triggering event during which the envelope exceeded the mean. SWR amplitude was measured in standard deviations from baseline.

We defined when animals were stopped at a food well as times when the animal was within 10 cm of the well with a linear speed (e.g. speed along the long axis of the track) equal to zero. These times include only periods when the animal had arrived at the well and could consume reward if it was present. SWRs that occur during these times are referred to as wSWRs. Linear speed was calculated as the change in linear distance per position samples divided by the time between samples (33 ms) and was not smoothed. We obtained essentially identical results when we used a two dimensional speed of zero, also not smoothed, to define periods of immobility.

We calculated a number of measures related to activity during SWRs. The activation probability per SWR was the number of SWRs in which a cell was active divided by the total number of SWRs. The mean rate during SWRs was the total number of spikes during SWRs divided by the total duration of SWRs. The proportion of cells active per SWR was the proportion of cells with place fields on the track that were active during the SWR. The proportion of cells active per SWR was calculated for each SWR; activation probability, mean rate, and number of spikes per SWR were per cell measures and SWR rate was measured per trial.

To control for differences in the timing of rewarded and unrewarded trials within the session we selected single pairs of adjacent rewarded and unrewarded trials. We

randomly selected the order of the pairs: rewarded followed by unrewarded or unrewarded followed by rewarded. We also controlled for differences in time spent at the well by truncating the time stopped at the well on rewarded trials to match that time spent on unrewarded trials.

We examined activity during the entire run period when animals were running at greater than 2 cm/sec and were more than 20 cm from the start or end well. For each cell we calculated peak firing rate and mean firing rate as described previously(Karlsson and Frank, 2008) for all rewarded or unrewarded trials. We then examined the coactivation probability during the run to the well. Binning the run periods into 100 msec bins, for each pair of cells we calculated the probability that the two cells were active together in a bin. We also computed two measures that quantified the extent to which the coordinated activity in SWRs was greater than that expected by chance, the coactivity z-score(Cheng and Frank, 2008), and joint surprise (Grun et al., 2002;Pazienti and Grun, 2006). Definitions of these measures are presented in the Supplementary Methods.

We determined whether SWR activity during periods when the animal was stopped at the well was consistent with previous reports of reactivation using two complementary analyses. For all of these analyses we identified run periods during each trial as times when the animal was moving forward at greater than 2 cm/s and more than 10 cm from the start or end well. We excluded cells with place fields within 10 cm of the wells at the beginning or end of the pass to focus on the cells with place fields active during running.

First, we calculated the probabilities that a cell active on the run towards the well or on the subsequent run away from the well was also active during wSWRs. We compared those probabilities to the probability that the cell was active on a randomly chosen run between wells. Second, we used a pair-wise measure to determine whether

the spiking during wSWRs was consistent with the ordered replay seen in downstream CA1. This pair-wise measure was necessary because CA3 tends to be sparsely active (Leutgeb et al., 2004) and, as our goal was to record for many days, we did not maximize the number of simultaneously recorded cells on a single day. The presence of coherent replay would predict that the time between spikes from different cells in a SWR should be related to the distance between the cells' place fields. For every pair of place cells active in a wSWR, we measured the absolute value of the time from each reference spike of one cell to all spikes from the other cell for each trial. We restricted the time axis to values between 0 and 500 ms and plotted each time between wSWR spikes from the pair of cells against the linear distance between the place field centers on the preceding run. We used the preceding run because some cells had multiple fields on the track, and we wished to focus on activity associated with the most recent run between food wells. We calculated the R^2 value of a linear fit to the points in the plot, which measured the degree to which the distance between the cells' place fields predicts the time between their wSWR spikes.

We characterized the run period activity of cells that fired during SWRs using three complementary analyses. We identified run periods during each trial as times when the animal was moving forward at greater than 2 cm/s and more than 10 cm from the start or end well. We excluded cells with place fields within 10 cm of the wells at the beginning or end of the pass to focus on the cells with place fields active during running. First, to visualize run period activity in the two dimensional track we collected all run period spikes from the cells that fired during wSWRs at a given well, identified the firing location of each spike in two dimensions, and created occupancy normalized rate maps as though these spikes had come from a single neuron.

Second, we computed the distribution of peak firing locations for all cells and cells that fired during wSWRs. For each trajectory for each cell, we computed the

location of the peak occupancy normalized firing rate and then created a distribution of peak locations from that set of cells. We computed these curves separately for rewarded and unrewarded trials.

Finally, we examined the location and timing of spikes during the run period leading up to a wSWRs. For each cell that fired during wSWRs on a pass, we determined the time between the cell's spikes during the run periods and the wSWRs in which the cell fired. This time was also decomposed into the time between the spikes during the run and when the animal reached the reward well and the time between when the animal reached the well and when wSWRs occurred.

Figures

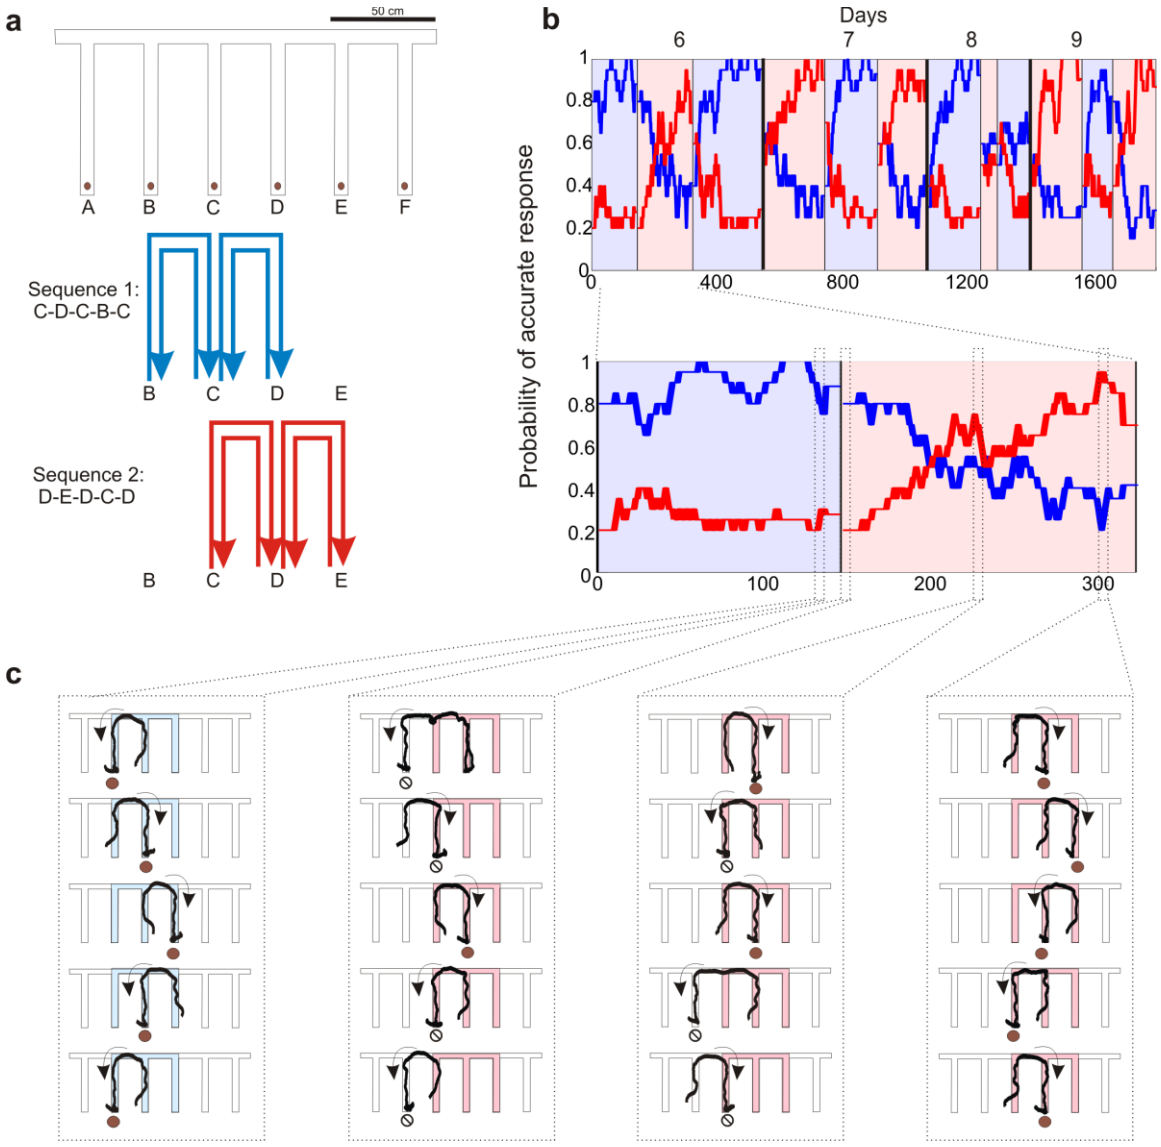

Figure 1. Task design and behavioral performance. a, Overhead view of the behavioral apparatus with reward sequences indicated by colored arrows (Sequence 1 (S1) in blue and Sequence 2 (S2) in red). Brown circles indicate the location of food wells; reward was delivered in arms B-E. Arms are 7.25 cm wide. Scale bar is 50 cm. **b,** Twenty trial moving average of correct responses for one animal when switching between performing S1 (blue) and S2 (red). Background color indicates which sequence

was rewarded. Black lines separate recording days. Top shows all sequence switching days, bottom shows first day of switching (day 6). Chance performance on this task is 0.2 as there are 5 arms the animal can choose from when leaving an arm. **c,** Examples of specific trials at the end of session one when S1 was rewarded (far left), the beginning of session 2 just after the reward contingencies change to reward S2 (middle left), the middle of session 2 (middle right), and the end of session 2 (far right). The animal's path through space is shown in black and the rewarded sequence arms are highlighted (S1 in blue and S2 in red). The arrow indicates direction. A brown filled circle indicates that the animal received chocolate at the end of the trajectory while a "Ø" indicates no reward was delivered. In the first two columns the animal performed the same sequence of trajectories but was rewarded differently. In session 1 all of the trajectories shown were rewarded as they are the correct performance of S1 (far left). Then at the beginning of session 2 (middle left), the animal was placed in the center arm of the newly rewarded sequence (S2) but performed the previously rewarded sequence (S1). In this case, only the third trajectory was rewarded (it is a correct return to the home arm of S2) while the other trajectories were unrewarded.

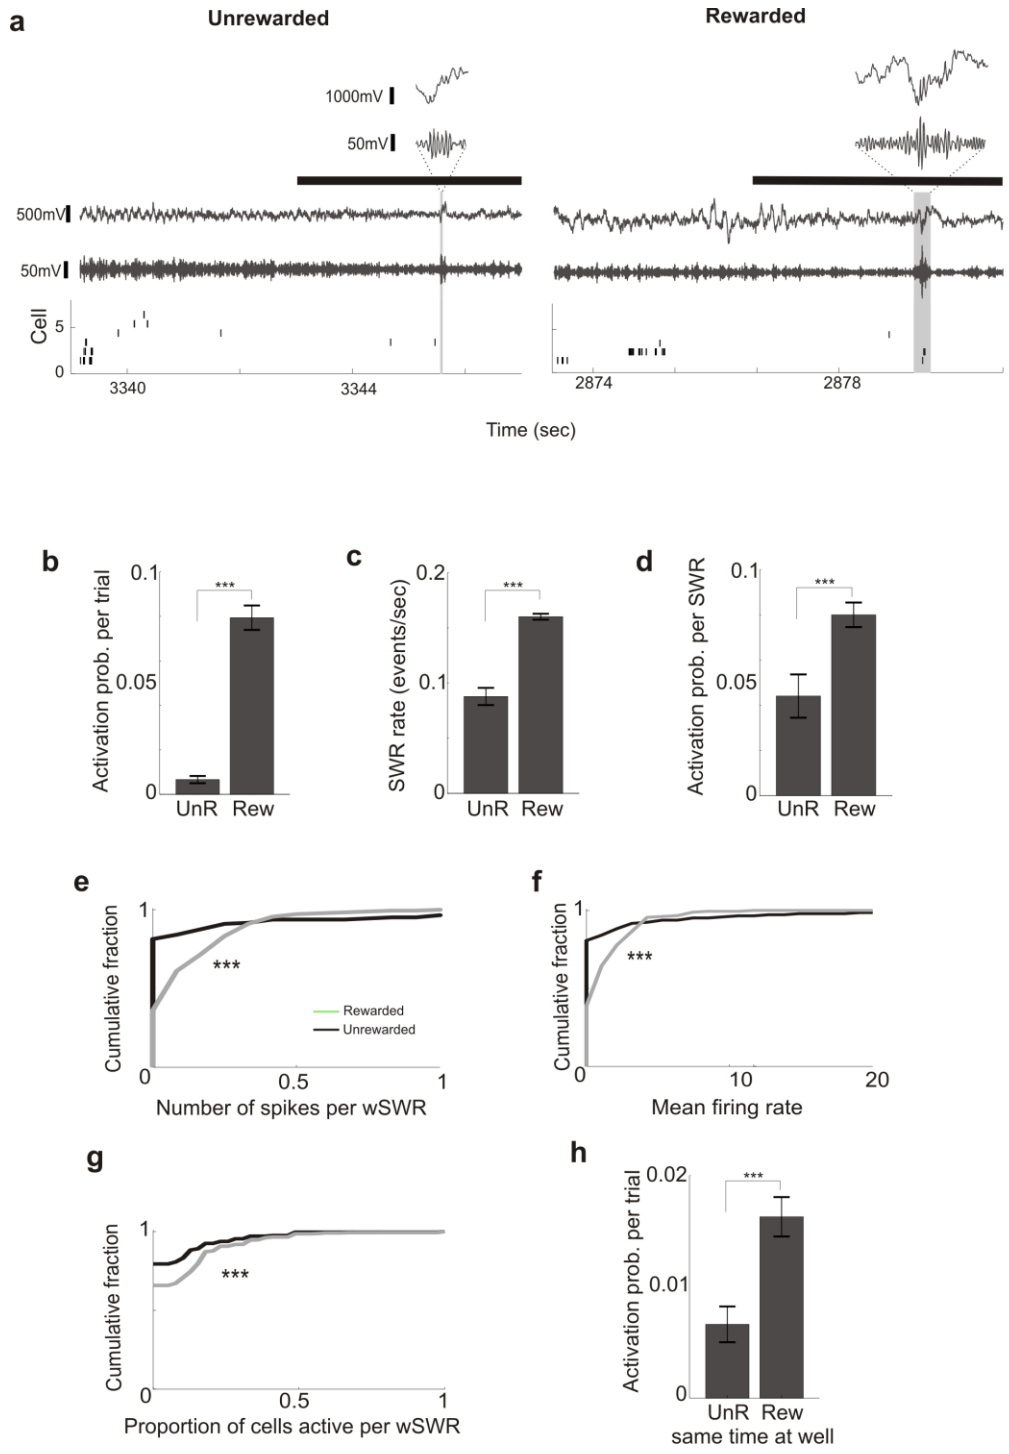

Figure 2. Enhanced SWR activity following reward. a, Examples of a single unrewarded (left) and rewarded (right) trial including times at the well (black lines) and the preceding run period. Trials are from the same epoch and contain the same cells.

Spikes that occurred during SWRs (grey bars on raster plot) when animals were stopped at the well are shown in red, all other spikes are shown in black. SWRs were detected using the simultaneously recorded EEG filtered at 150-250Hz (dark grey traces above raster plot, top: unfiltered, bottom: filtered 150-250Hz). The top traces show the SWR in unfiltered and filtered EEG at higher magnification. Activation probability per trial **(b)**, wSWR rate **(c)**, and activation probability per wSWR **(d)** for cells with place fields on the track when animals were stopped at the well in rewarded and unrewarded trials. Cumulative distribution of the mean number of spikes per wSWR **(e),** the mean firing rate **(f)**, and the proportion of cells active **(g)** during wSWRs on rewarded (green) and unrewarded trials (black). **h,** Activation probability per trial when rewarded trials were truncated to match the duration of immobility on unrewarded trials. Error bars represent standard errors. Bar graphs include only times when animals were stopped at the well. *** $p < 10^{-4}$.

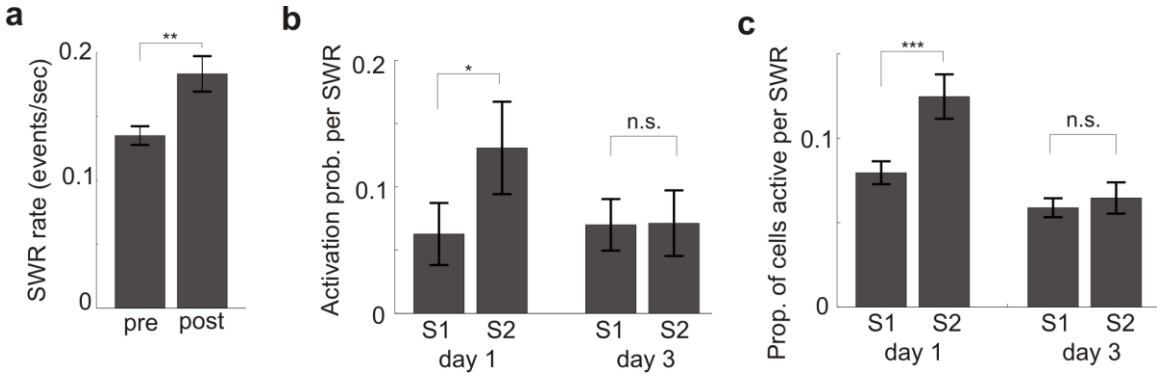

Figure 3. SWR activity following reward is further enhanced during learning. a, wSWR on rewarded trials before ('pre') and after ('post') reward contingencies changed for the first time (e.g. before and after the first exposure to S2). **b,** Activation probability per wSWR and **c,** proportion of cells active per wSWR in S1 (1^{st} and 3^{rd} bars) or S2 (2^{nd} and 4^{th} bars) on the first and third day of switching between sequences. All data include only rewarded trials when animals were stopped at the well. Error bars represent standard errors. *** indicates $p < 10^{-5}$, ** indicates $p < 0.005$, and * indicates $p < 0.05$.

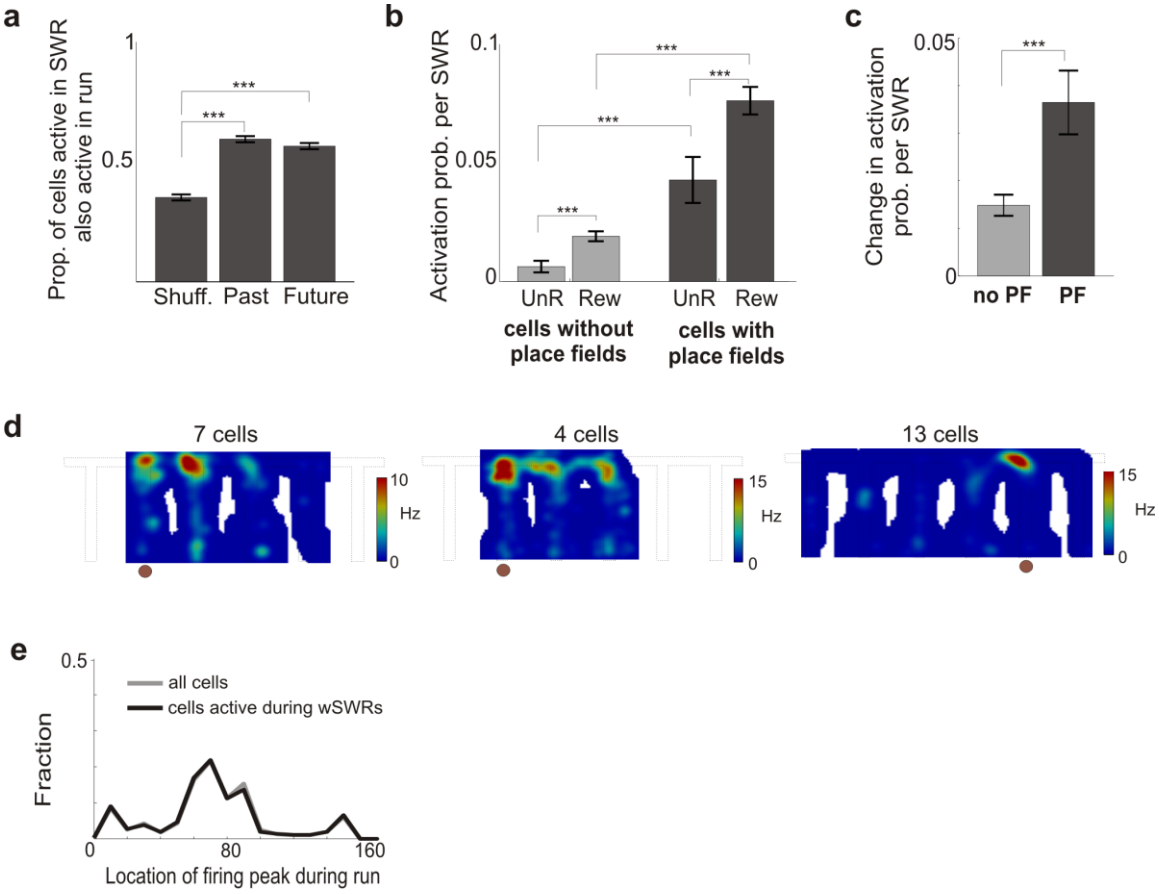

Figure 4. Reactivation of single cells during SWRs following reward. a, Proportion of cells active during wSWRs that are also active on the preceding run to the well (past) or the subsequent run from the well (future). **b**, Activation probability per wSWR for cells with (blue) and without (grey) place fields on the track. **c,** Within cell differences in activation probability per wSWR between rewarded and unrewarded trials for cells with (blue) and without (grey) place fields on the track. Activation probability for one cell on unrewarded trials was subtracted from that on rewarded trials. Positive differences indicate that rewarded trials had a greater activation probability than unrewarded trials. **d,** Firing rate maps of run period firing of all cells that fired during wSWRs at the well indicated by the brown circle. Spikes from all simultaneously recorded cells are shown together for each of 3 different epochs, one from each animal. **e**, Location of peak

occupancy normalized firing rate during run periods of all cells (grey) and cells that fired in ripples at the end of the trajectory (black) on rewarded trials. The distributions of peak firing locations on each trajectory were averaged together. Note that the two curves are almost entirely overlapping. Error bars represent standard errors. *** indicates $p < 10^{-5}$.

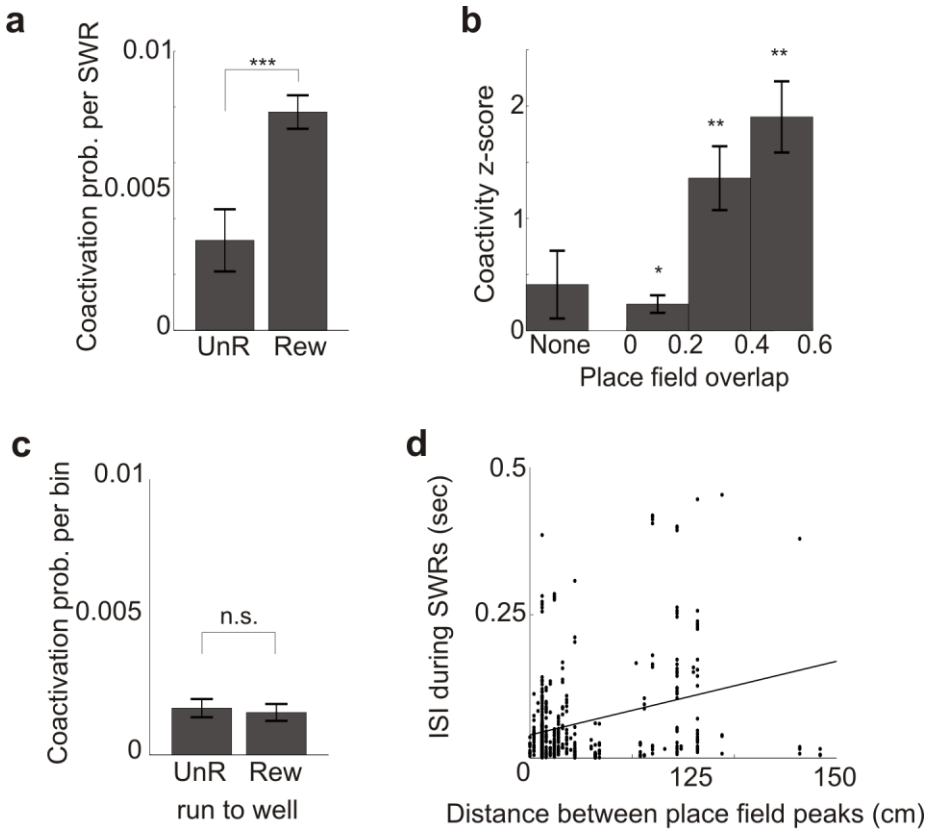

Figure 5. Structured reactivation of pairs of cells during SWRs following reward. a, Coactivation probability per wSWR on rewarded and unrewarded trials. **b,** Coactivity z-score per wSWR versus place field overlap on rewarded trials. Coactivity z-score was binned according to the cell pairs' place field overlap: no overlap, greater than 0 but less than 0.2, 0.2 to 0.4, and 0.4 to 0.6 overlap. Because many cells never fire in ripples on unrewarded trials, there were too few coactive pairs to perform this analysis on unrewarded trials. **c**, Coactivation among cell pairs per 100 msec bin during the run period of rewarded and unrewarded trials. **d,** Scatter plot of distance between place field peaks versus inter-cell inter-spike interval (ISI) during wSWRs. Each point is one inter-cell ISI during wSWR plotted against the distance between the cells' peaks on that pass. Only spikes that occurred during wSWRs when animals stopped at the well were included in measures of wSWR inter-cell ISI. A singer ISI could extend over multiple

SWRs, as previous work has shown coherent reactivation across SWRs (Davidson et al., 2009). Previous quantifications of replay have found replay in a subset of SWRs, and the R^2 values seen here would be expected to reflect the combination of events with replay and events where replay was absent. Error bars represent standard errors. * indicated $p < 0.01$, ** indicates $p < 0.0001$, *** indicates $p < 10^{-10}$.

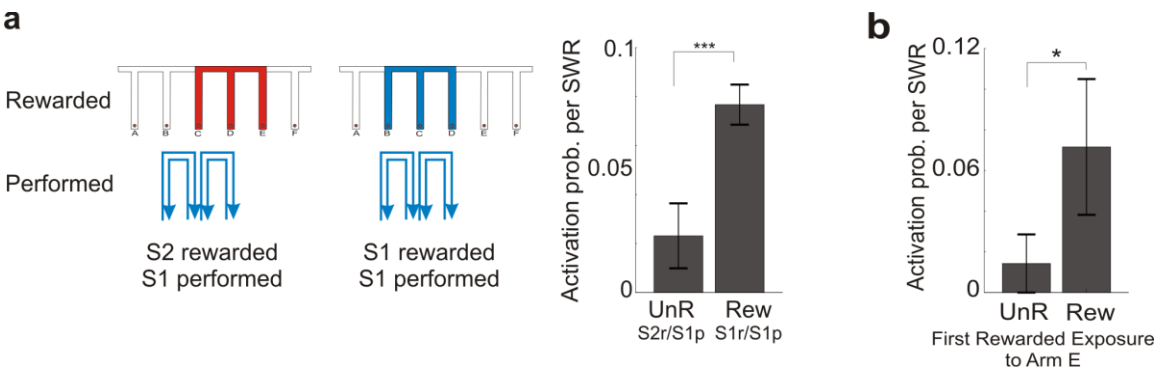

Figure 6. Comparison of SWR activity across behavioral conditions. a, Activation probability for the same behavioral sequence with different rewards. Left: schematic of sequence rewarded (highlighted on track, S1 in blue and S2 in red) and sequence performed (arrows below track) for data shown at right. Right: activation probability per wSWR when animals accurately performed S1 and reward was omitted (S2 rewarded, S2r/S1p) or delivered (S1 rewarded, S1r/S1p). **b,** Activation probability per wSWR during the first rewarded exposure to S2, when reward was unexpected based on the previous history of rewards, and unrewarded trials in all arms. Only times when animals were stopped at the well were included. *** indicates $p < 10^{-5}$.

Supplementary Figures

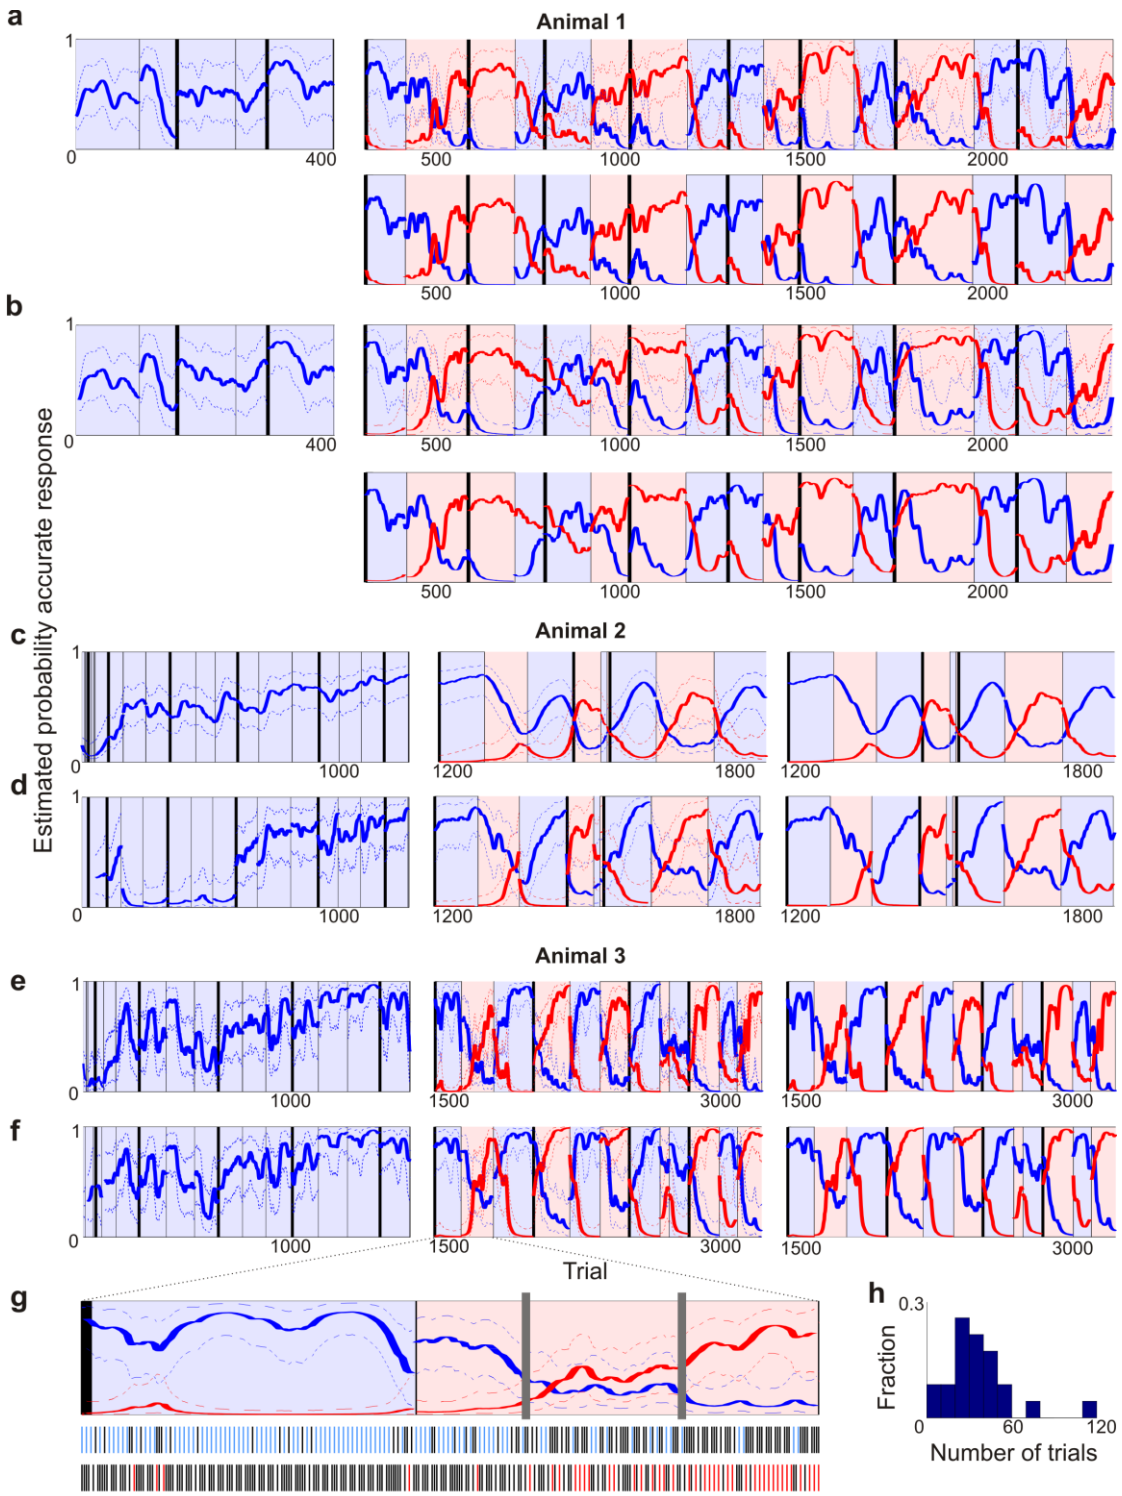

Supplementary Figure 1. Behavioral performance on outbound and inbound trajectories for each animal. **a,** Estimated probability of an accurate response across trials on outbound trajectories (trajectories starting in arm C in S1 or arm D in S2) for animal 1 during initial learning of S1 (left, solid line indicates mode, dashed lines indicate confidence bounds) and alternating between S1 and S2 (right, S1 in blue and S2 in red, background color indicates which sequence performance was rewarded). For clarity, the alternation phase is shown with confidence bounds (above) and without (below) **b,** Estimated probability of an accurate response across trials on inbound trajectories for animal 1. **c-f,** Estimated probability of an accurate response per trial on outbound (**c,e**) and inbound (**d,f**) trajectories for animal 2 and 3. For clarity, the alternation phase is shown without confidence bounds on the far right. All animals learned the initial S1 sequence and then learned to switch between S1 and S2 within each day. Black lines separate days. **g,** top: Modes and confidence bounds were used to separate the behavior into 3 periods (grey lines, see methods): when the animal performed the unrewarded sequence significantly more accurately than the rewarded sequence (to the left of the first grey line), when the animal performed neither sequence significantly more accurately than the other (between the grey lines), and when the animal performed the rewarded sequence significantly more accurately than the unrewarded sequence (to the right of the second grey line). Bottom: colored lines indicate the sequence of correct (blue for S1 and red for S2) or incorrect (black) trials that went into the estimates. **h,** The number of trials required to switch between sequences after the reward contingencies changed. The histogram shows the number of trials performed from the change in reward contingencies to when the animal performed the rewarded task significantly more than the unrewarded task. Animals took an average of 36 trials to switch. In 2 sessions it took more than 60 trials to switch and it never took more than 120 trials.

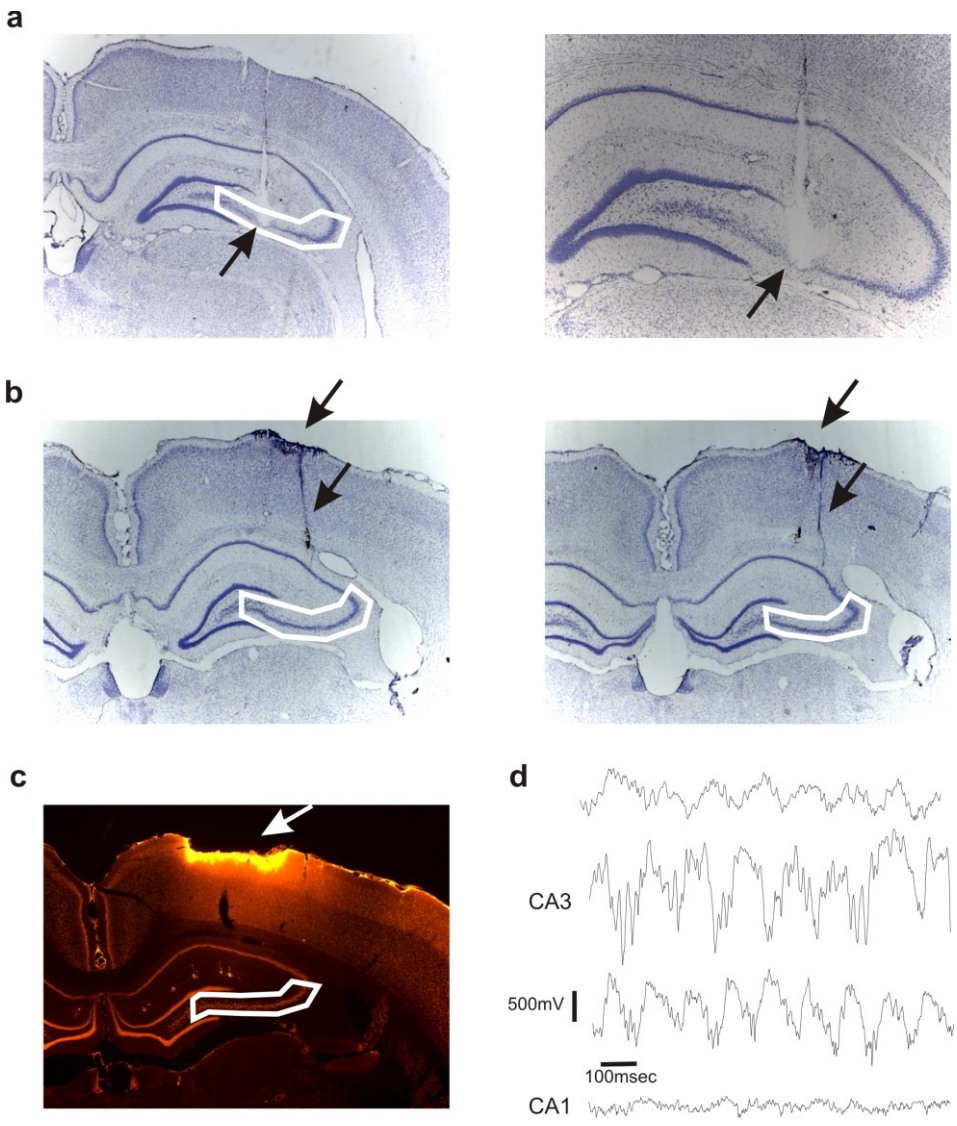

Supplementary Figure 2. Lesions and EEG signatures from CA3 recordings. a, Lesions (arrows) were made at the end of recording in animal 1 (left 20x, right 40x, Nissl stain). For animals without lesions (**b:** animal 3, 20x, Nissl stain**, c:** animal 2, 20x, NeuN stain), damage due to the drive implant and tetrodes were observed at the brain surface and tetrode tracks (arrows). Estimated areas of recording are outlined in white. **d,** All animals showed characteristic EEG signatures (top 3 traces) with large amplitude theta

modulation. For comparison, we included a recording from CA1(Karlsson and Frank, 2008) (bottom trace) which is much smaller in amplitude with less prominent theta modulation. Each trace is a total of one second, plotted on the same scale. As we lowered our tetrodes to CA3, we noted several "landmarks:" the quiet lack of cells in corpus callosum, the increase in ripple amplitude as we approached CA1, the densely packed cell layer in CA1, then the EEG amplitude increase on the other side of the CA1 cell layer, the relative quiet between CA1 and CA3, then again the increase in ripple amplitude again as we approached CA3 and finally the densely packed CA3 cell layer. Note that the larger size of the CA3 theta as compared to that seen in CA1 is likely to be due in whole or part to our use of a reference electrode that was in the corpus callosum above CA1.

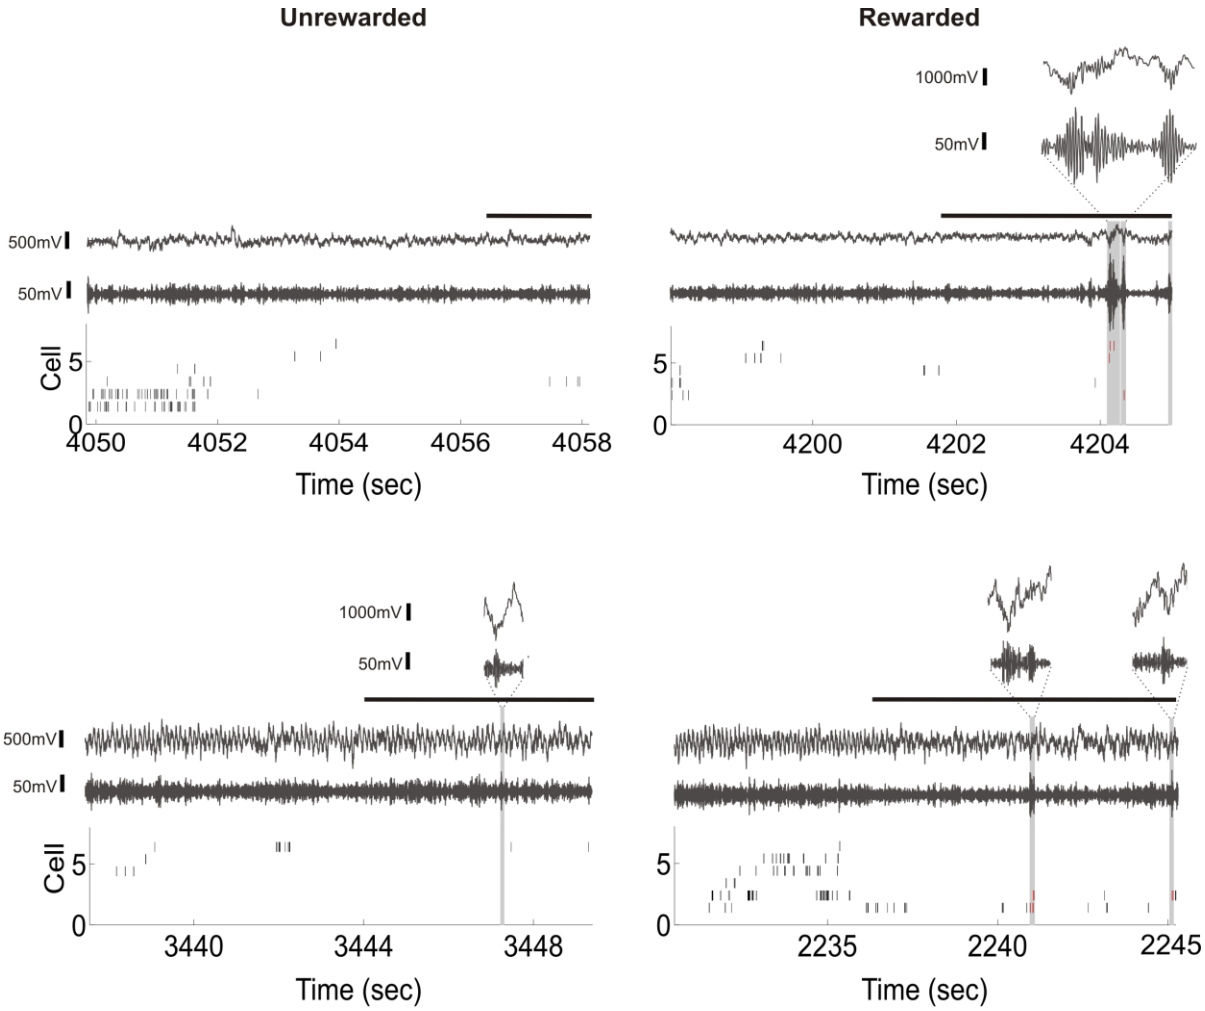

Supplementary Figure 3. Enhanced SWR activity following reward. Additional examples of single unrewarded (left) or rewarded (right) trials including times at the well (horizontal black lines) and the preceding run period. Trials in the same row are from the same run session and contain the same cells. This figure illustrates both the reduced spiking at the well for unrewarded trials and the reduced likelihood of seeing a SWR, (e.g. there was no SWR for the top left unrewarded trial). Spikes that occurred during SWRs when animals were stopped at the well (grey bars on raster plot) are shown in red, all other spikes are shown in black. SWRs were detected using the

simultaneously recorded EEG filtered at 150-250 Hz (grey, above raster plot, top: unfiltered, bottom: filtered for 150-250Hz). The top traces show the SWR in unfiltered and filtered EEG at higher magnification.

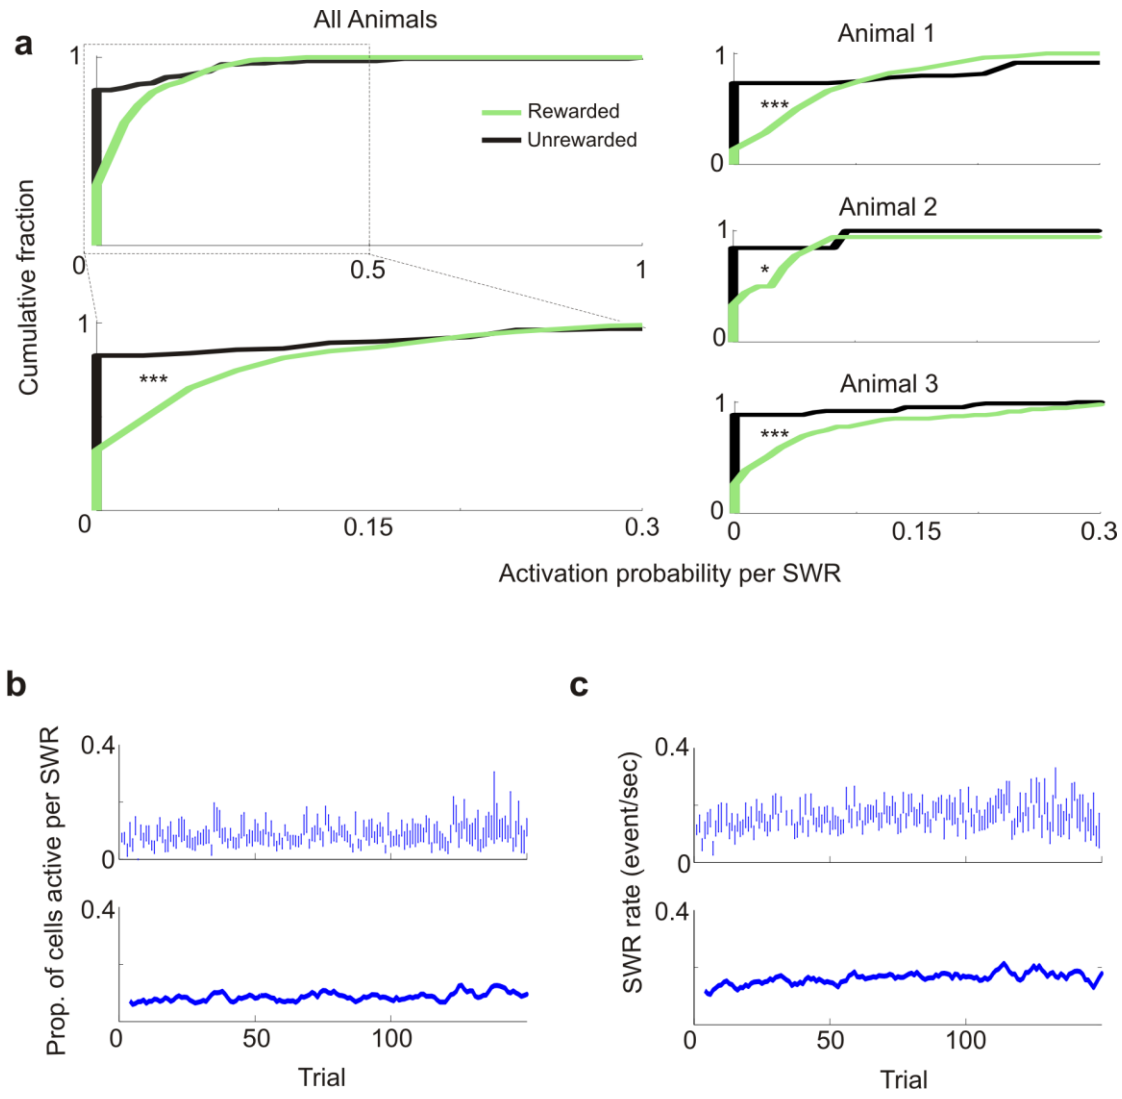

Supplementary Figure 4. SWR activity across animals and across trials for cells with place fields. **a,** Activation probability per wSWR is higher in rewarded trials in each animal. The graphs show cumulative distributions of activation probability per wSWR on rewarded (green) and unrewarded (black) trials for all animals (left, $p < 10^{-10}$) and each animal individually (right, Animal 1: $p < 10^{-5}$; Animal 2: $p < 0.05$; Animal 3: $p < 10^{-10}$). Overall, rewarded trials were associated with higher probabilities of activation and greater numbers of spikes than unrewarded trials. Proportion of cells active **(b)** and SWR rate **(c)** across trials for all sequence switching days combined. Mean and

standard error for each trial are shown above and 5 trial moving average is shown below. If SWR activity was enhanced by generally surprising events like the change in reward contingencies that occurred at the beginning of each session, we would expect higher activity on early trials. Instead, the overall proportion of cells active and SWR rates were stable or perhaps slightly increasing, indicating that any surprise related to changes in reward contingencies at the beginning of each session cannot explain our results. We should also note that these across trials values were highly variable within individual sessions, and we were therefore unable to quantify any within day trends for S2.

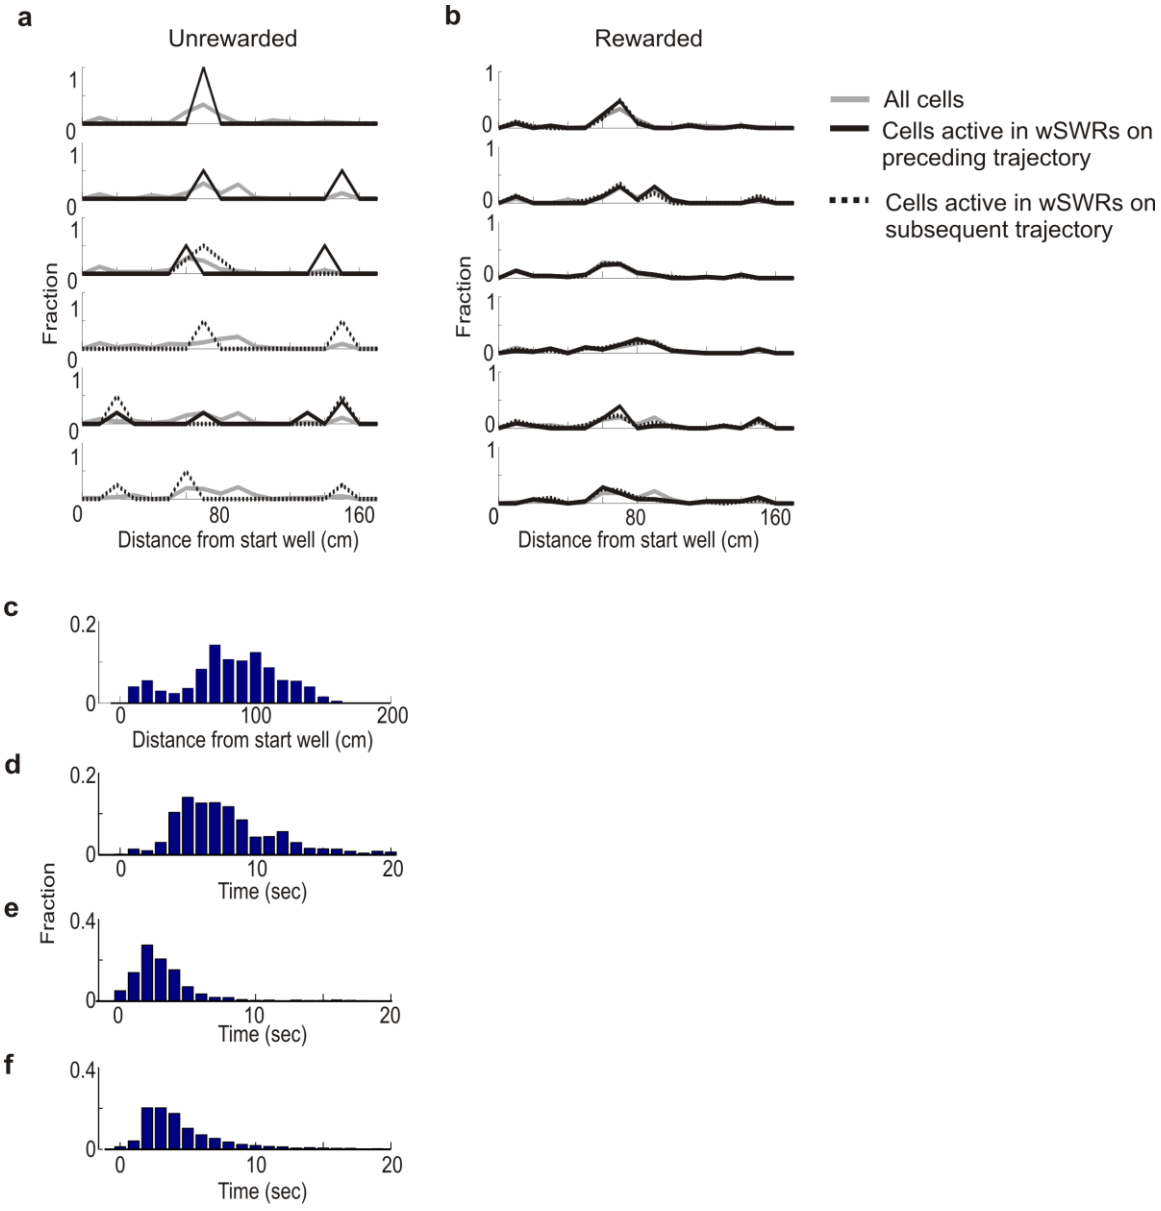

Supplementary Figure 5. Location and timing of spiking activity preceding wSWRs. Location of peak occupancy normalized firing rate during run periods for all cells (grey), for cells that fired in ripples at the end of that trajectory (black solid lines), and for cells that fired during wSWRs at the reward well and were active on the subsequent trajectory (black dashed lines) on unrewarded **(a)** and rewarded trials **(b)**. Each row is one trajectory. We also examined the location and time of each spike on the

trial leading up to the reward well for cells that fired during wSWRs. Again, we found that cells fired throughout the trajectory but more around the turns. These results suggest there is little or no spatial bias for cells close to the reward location to fire more than cells farther away. **c,** Location of spikes (distance from the start well of the trajectory) during the run period for cells that fired during wSWRs. **d,** Time between spikes during the run period and SWRs in which the cell fired. **e,** Time between when spikes fired during the run period and when the animal reached the reward well. **f,** Time between when the animal reached the reward well and when wSWRs occurred. The time between spikes during the run and wSWRs was a wide distribution with few cases when spiked occurred within 3 seconds or less and most cases greater than 5 seconds. This time between run and SWR spikes consisted of the time it took the animal to reach the well from the cells' place field area and the time from when the animal reached the well to when wSWRs occurred. A simple model whereby activity within SWRs was related to a smoothly decaying function of time from place field activity cannot explain these patterns, as cells active 5 or more seconds before the wSWR were more likely to be active in the wSWR than cells active less than five seconds before the wSWR.

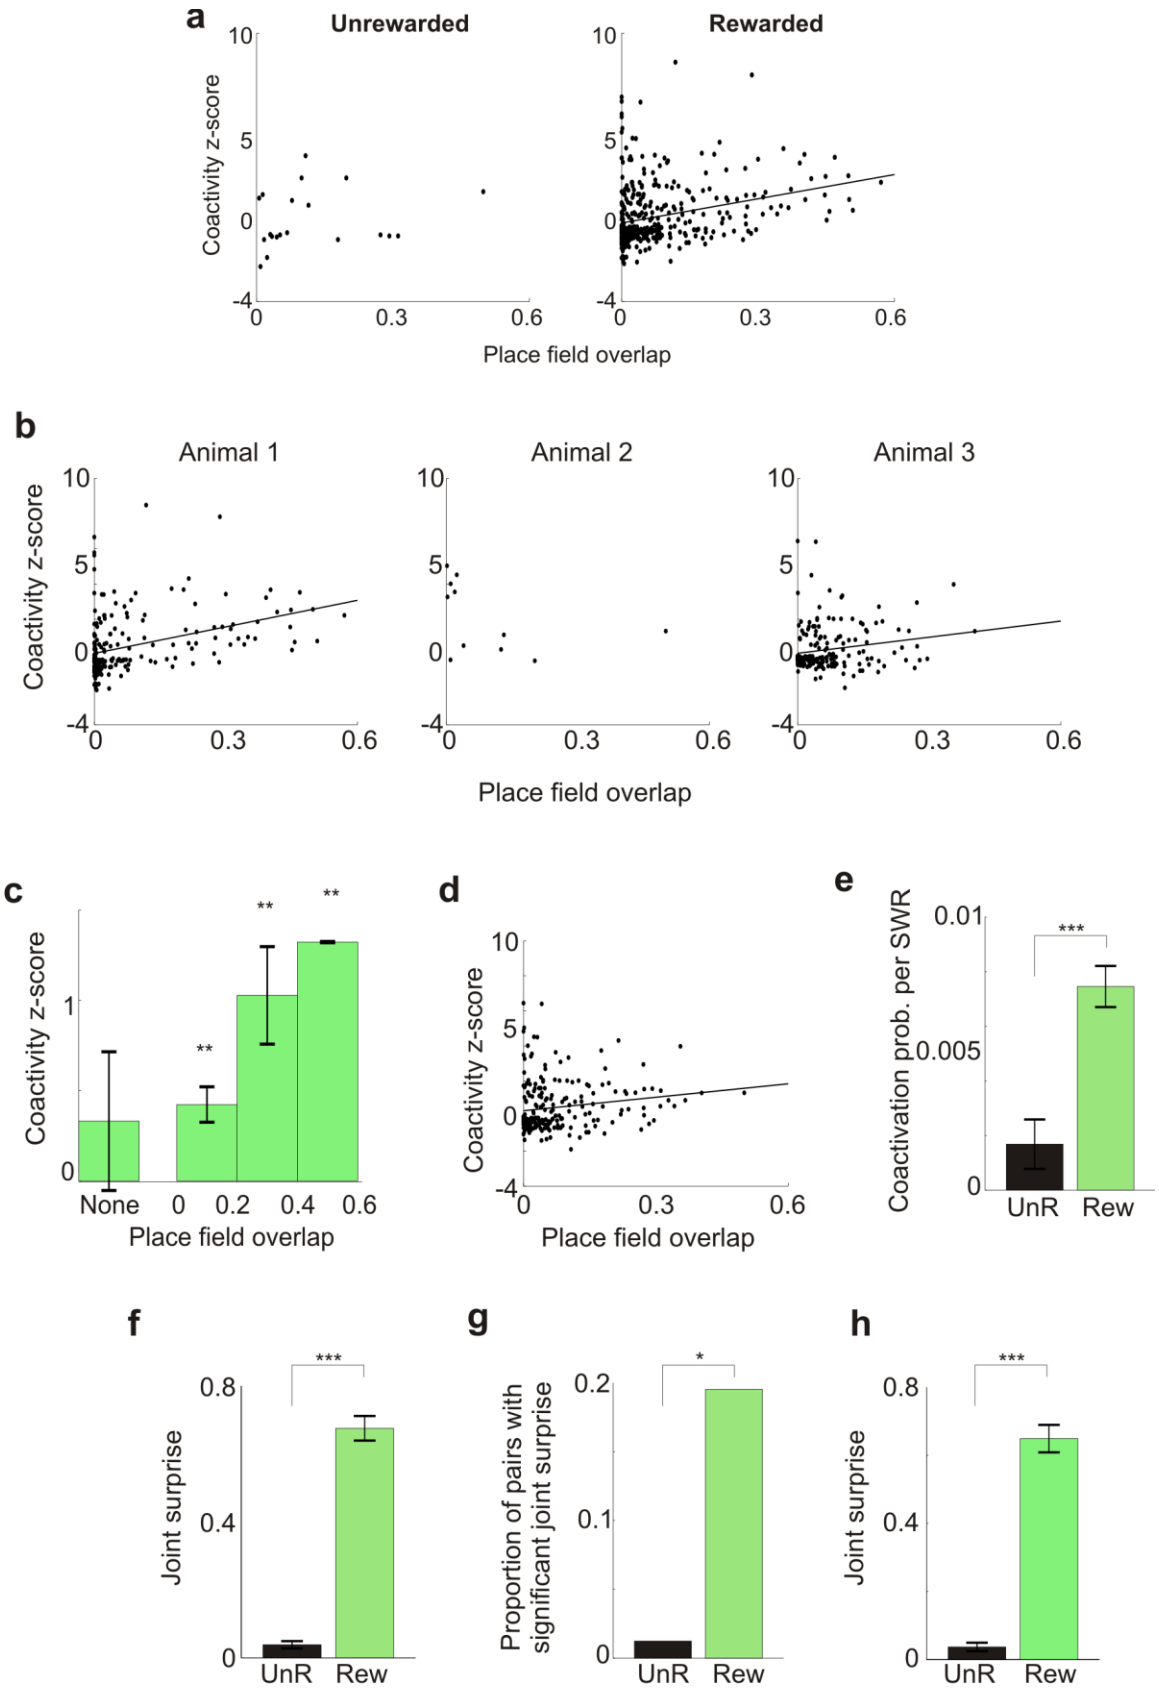
a
Unrewarded
Rewarded
Coactivity z-score
Place field overlap
b
Animal 1
Animal 2
Animal 3
Coactivity z-score
Place field overlap
c
Coactivity z-score
None
Place field overlap
d
Coactivity z-score
Place field overlap
e
Coactivation prob. per SWR
UnR
Rew
f
Joint surprise
UnR
Rew
g
Proportion of pairs with significant joint surprise
UnR
Rew
h
Joint surprise
UnR
Rew

Supplementary Figure 6. **Comparison of coactivity within wSWRs.** **a**, Coactivity z-score per wSWR versus place field overlap on unrewarded (left) and rewarded trials (right). Each point is one pair of cells. The regression was highly significant (R^2 = 0.0846, $p < 10^{-5}$). Coactivity z-score is undefined when a cell in the pair never fires in ripples or fires in all ripples. Because many cells never fire in ripples on unrewarded trials, there are fewer cell pairs in unrewarded trials. **b**, Coactivity z-score per wSWR versus place field overlap on rewarded trials for each animal. Each point is one pair of cells. The regression was highly significant for animals 1 and 3 (R^2 = 0.1382, $p < 10^{-5}$ and R^2 = 0.0346, $p < 0.01$, respectively). **c,** Coactivity z-score per wSWR versus place field overlap on rewarded trials only including cells recorded on different tetrodes. Coactivity z-score was binned according to the cell pairs' place field overlap: no overlap, greater than 0 but less than 0.2, 0.2 to 0.4, and 0.4 to 0.6 overlap. Because many cells never fire in ripples on unrewarded trials, there were too few coactive pairs to perform this analysis on unrewarded trials. Results were similar to those shown in Figure 4b. **d**, Coactivity z-score per wSWR versus place field overlap on rewarded trials only including cells recorded on different tetrodes(R^2 = 0.0272, $p < 0.01$). Each point is one pair of cells. **e,** Coactivation probability per wSWR on rewarded and unrewarded trials only including cells recorded on different tetrodes. Results were similar to those shown in Figure 4a. Joint surprise **(f)** and proportion of cells pairs with significant joint surprise, or joint surprise corresponding to $p < 0.05$ **(g),** for rewarded and unrewarded trials. The joint surprise between pairs of cells and the proportion of pairs with significant joint surprise was greater on rewarded than unrewarded trials. We further controlled for potential spike sorting errors by repeating these analyses including only cells that were recorded from different tetrodes. **h,** Joint surprise for rewarded and unrewarded trials only including cells recorded on different tetrodes ($p < 10^{-10}$). ** indicates $p < 0.005$. *** indicates $p < 10^{-10}$.

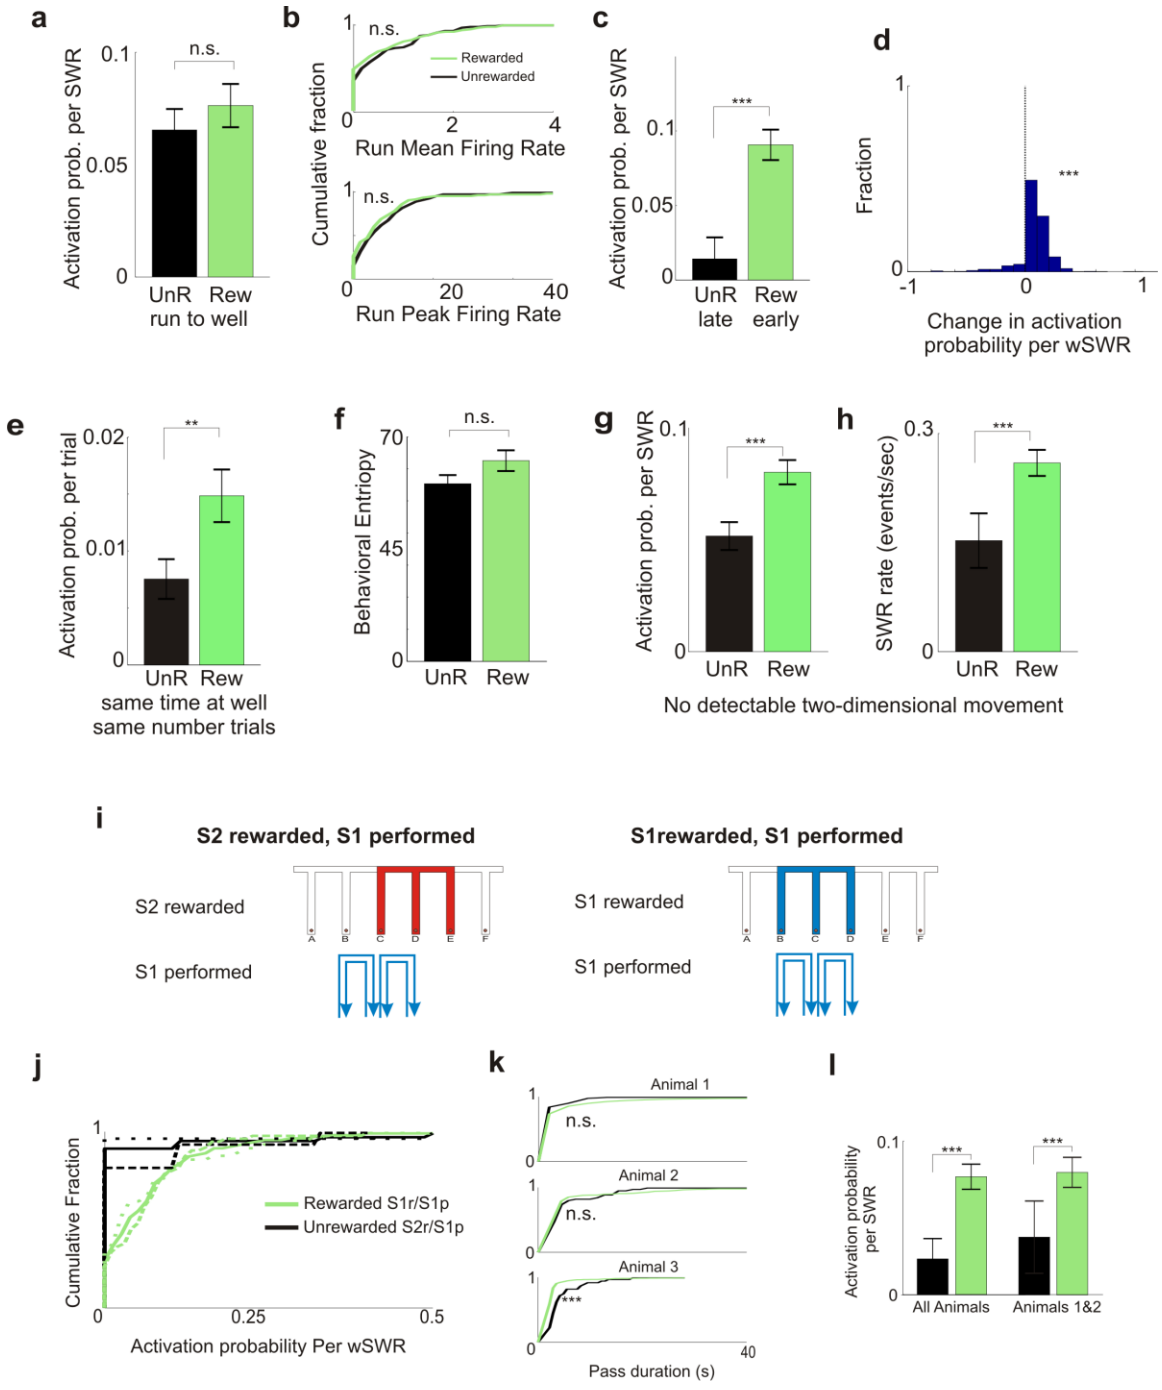

Supplementary Figure 7. Controls for behavioral or SWR differences across conditions. **a,** Activation probability per SWR when animals were more than 20 cm from the well on rewarded and unrewarded trials ($p > 0.12$). **b**, Cumulative distribution

of mean firing rate (top, $p > 0.11$) and peak firing rate (bottom, $p > 0.18$) during run periods. **c,** Activation probability per wSWR on rewarded trials from the first half of the session and unrewarded trials from the second half of the session. Only the first session of each day was included. **d,** Differences within cell in activation probability per wSWR between rewarded and unrewarded trials. Activation probability for one cell on unrewarded trials was subtracted from that on rewarded trials. Positive differences indicate that rewarded trials had a greater activation probability than unrewarded trials. The distribution was significantly different from zero (student's t-test, $p < 10^{-5}$). **e,** Activation probability per trial including only adjacent rewarded and unrewarded trials. Reward trials were also truncated to match the time since immobility onset and the duration of immobility on unrewarded trials ($p < 0.001$). **f,** Behavioral entropy on rewarded and unrewarded trials. **g,** Activation probability per wSWR ($p < 10^{-10}$) and **h,** wSWR rate ($p < 0.001$) on rewarded and unrewarded trials when the animal was stopped at the well and the animal's two-dimensional speed was zero. Therefore, detectable movements in two-dimensions do not explain differences in activation probability or SWR rate on rewarded and unrewarded trials. **i,** Activation probability for the same behavioral sequence with different rewards. Schematic of sequence rewarded (highlighted on track, S1 in blue and S2 in red) and sequence performed (arrows below track, S1 in blue) for data shown in **(j)**. **j,** Cumulative distribution of activation probability per wSWR in unrewarded trials when S1 was performed but S2 was rewarded (black) and rewarded trials when S1 was performed and S1 was rewarded (green) for each animal (animal 1: solid lines, animal 2: dashed lines, animal 3: dotted lines). Unrewarded trials were associated with lower activation probabilities during wSWRs even when the animal was accurately performing a learned sequence. **k,** To ensure that differences in wSWR activity were not associated with differences in behavior in this situation, we also measured the length of time required for the animal to complete each

pass from one food well to another. The distributions were not different for animals 1 and 2, but did differ for animal 3 (rank sum test, animal 1: $p > 0.15$, animal 2: $p > 0.05$, animal 3: $p < 10^{-5}$). We therefore recalculated the overall activation probability for the comparison including only animals 1 and 2 **(I)** and found that activation during wSWRs on rewarded trials was still significantly greater than activity on unrewarded trials ($p < 0.0001$). ** indicates $p < 0.001$; *** indicates $p < 10^{-5}$.

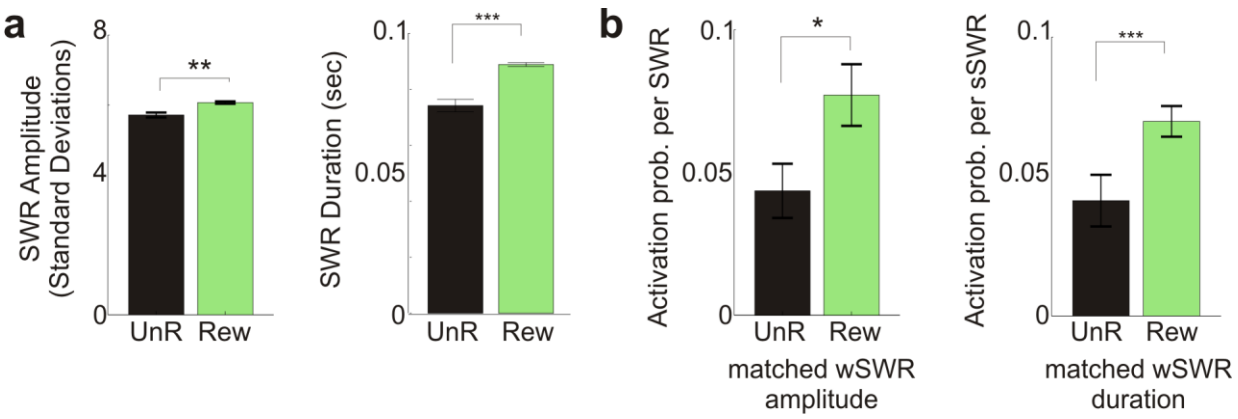

Supplementary Figure 8. Comparison of wSWR amplitudes and durations across reward conditions. **a,** wSWR amplitude (left) and duration (right) for rewarded trials and unrewarded trials. **b,** Activation probability per wSWR when rewarded wSWRs were selected to match the amplitude of wSWRs in unrewarded trials (left) and when rewarded wSWRs were truncated to match the duration of wSWRs in unrewarded trials (right). Error bars represent standard errors. * indicated $p < 0.01$, ** indicates $p < 0.0001$, *** indicates $p < 10^{-10}$. The relationship between SWRs and the associated neural activity is likely to be complex and non-linear, but these results indicate that differences in amplitude or duration cannot explain our results.

Supplementary Methods.

For comparing activity of cells with and without place fields we measured the within cell change in activation probability. This was calculated by subtracting the activation probability in unrewarded trials from that in rewarded trials for each cell. Therefore a positive change means the cell was more likely to be active in rewarded trials. The significance of this change was evaluated using a student's t-test. We measured the absolute increase as opposed to the increase multiplicatively for two reasons. First, many cells were entirely inactive during unrewarded SWRs. A single spike from one of these cells during an SWR on a rewarded trial would lead to an apparently infinite increase in activation, which means that the measure becomes very difficult to interpret. Second, the absolute measure captures the fact that, on rewarded trials, cells with place fields are about four times more likely to be activated than cells without place fields.

The coactivity Z-score was calculated as follows. Of a total of N events, if neurons A and B are independently active during n_A and n_B, respectively, the expected number of events during which both neurons were active, n_{AB}, follows a hypergeometric distribution (Sheskin, 2004) with mean

$$E\ n_{AB} \mid N, n_A, n_B\ = \qquad N$$

and variance

$$\sigma\ [n_{AB} \mid N, n_A, n_B] = \frac{n_A n_B (N - n_A)(N - n_B)}{N^2 (N - \ )}.$$

We can then compare expected and observed number of coincident events across neuron pairs with different activity levels with a z-score where we normalize the difference by the standard deviation:

$$z = \frac{n_{AB} - E[n_{AB} \mid N, n_A, n_B]}{\sigma[n_{AB} \mid N, n_A, n_B]}.$$

The joint surprise between cell pairs during rewarded and unrewarded trials was calculated as described previously (Grun et al., 2002;Pazienti and Grun, 2006). The joint surprise is defined as

$$js(\psi) = \frac{\log(1-\psi)}{\psi}$$

where the joint p-value ψ is computed from comparing the number of observed events, n_{emp}, and the number of predicted events, n_{pred}, in which both neurons were active:

$$\psi(n_{emp}, n_{pred}) = \sum_{n=n_{emp}}^{N} \frac{n^{-}}{n!}.$$

The number of predicted coactive events is based on the single cell activation probability per ripple, p_1 or p_2, during N observed ripples:

$$n_{pred} = N p_1 p_2.$$

If ψ was 1, as is the case when one or both neurons did not fire during SWRs, $js(\psi)$ was set to 0.

To ensure that there were no differences in overall behavior between rewarded and unrewarded trials we calculated the behavioral entropy as described previously(Jackson et al., 2006). Briefly, for each session, we binned the two-dimensional position data into 9.25 X 9.25 cm bins and positioned these bins to

encompass the 7 cm width of each arm, but the results were consistent across a range of bin sizes. The transition probability from each bin into every other bin was calculated for rewarded and unrewarded trials and all non-zero transitions were used to calculate the Shannon entropy:

$$H = \sum_{i=} \sum_{j=} \qquad _{ij}$$

where $p_{i,j}$ is the time-independent probability of transition from bin j to bin i. More variable behaviors will have a higher entropy.

We also examined the relationship between SWR properties and activation probability. To determine if differences in activation persisted when there was no significant difference in wSWR amplitude, we matched wSWRs in rewarded trials to the amplitude of wSWRs in unrewarded trials. For each wSWR in unrewarded trials, we selected the wSWR from rewarded trials in the same epoch with the closest amplitude. The amplitude of wSWRs from unrewarded trials and the selected wSWRs from rewarded trials did not differ significantly ($p > 0.12$). We then computed activation probability during wSWRs in unrewarded trials and during selected wSWRs in rewarded trials.

To control for differences in wSWR duration we truncated wSWRs in rewarded trials to match the duration of wSWRs in unrewarded trials. For each wSWR in rewarded trials, we randomly selected a wSWR from unrewarded trials in the same epoch. We truncated the wSWR in the rewarded trial to match that in the unrewarded trial. If the rewarded SWR was shorter than the unrewarded SWR, neither SWR was truncated. The duration of wSWRs from unrewarded trials and the selected wSWRs from rewarded trials did not differ significantly ($p > 0.43$). We then computed activation probability during wSWRs in unrewarded trials and during truncated wSWRs in rewarded trials.

Supplementary Results

SWR activity and place field activity across behavioral conditions

We performed a number of control analysis to determine whether factors other than reward could have lead to the differences we observed. First, we asked whether greater activation in rewarded wSWRs could have been caused by differences present before the animal reached the reward, as might be expected if animals' level of attention differed between rewarded and unrewarded trials. We isolated all SWRs that occurred when the animal was more than 20 cm from the well, thereby focusing on times where the presence or absence of reward was difficult or impossible for the animal to determine. We found no significant differences in activation probability per SWR on correct trials preceding reward or incorrect trials preceding a lack of reward (Fig. S7a, $p > 0.12$, 95% confidence intervals for unrewarded: 0, 0.4667 and rewarded: 0, 0.5, $n = 107$). We also examined individual place cells throughout the run period, restricting the analyses to cells that were active on runs in both rewarded and unrewarded trials. We found that neither their mean nor their peak firing rate during the run period differed significantly between rewarded and unrewarded trials (Fig. S7b, run mean rate $p > 0.19$, 95% confidence intervals for unrewarded: 0.008, 3.138 Hz and rewarded: 0.002, 3.026 Hz, run peak rate $p > 0.20$, 95% confidence intervals for unrewarded: 0.226, 38.61 and rewarded: 0.144, 49.75, $n = 107$).

We then controlled for the timing of rewarded and unrewarded trials to determine if this could explain the increased wSWR activity. Previous work has documented an increase in SWR activity over the course of the first ~20 trials in a familiar environment (Jackson et al., 2006). Animals in our study tended to make more errors early in each session, so unrewarded trials came, on average, earlier than rewarded trials. To ensure that increased wSWR activity was not due to this ordering, we compared rewarded trials

from the first half and unrewarded trials from the second half of the first session of each day. We found that cells had significantly higher activation probabilities per wSWRs on these rewarded trials early in the session than unrewarded trials late in the session (Fig. S7c, $p < 10^{-10}$) and the within cell increase in activation probability per wSWR from late unrewarded to early rewarded trials was significantly greater than zero (student's t-test $p < 10^{-5}$; Fig. S7d). We found similar results if we truncated the time spent at the well on these rewarded trials to match the unrewarded trials ($p < 10^{-4}$). We also controlled for the ordering of rewarded and unrewarded trials, the numbers of each trial type, and their timing within the session by selecting pairs of adjacent rewarded and unrewarded trials. Cells still had higher activation probabilities per trial on rewarded trials than unrewarded trials (see methods, Fig. S7e, $p < 0.001$). Thus, SWR activity increases due to repeated behavioral experience could not explain the increased activation probability during wSWRs on rewarded trials.

We then examined the animals' two-dimensional motion to determine if differences in the path to the well or behavior at the well could explain increased wSWR activity on rewarded trials. We examined the behavioral entropy of the transitions between two-dimensional positions in rewarded and unrewarded trials (Jackson et al., 2006). We found no significant difference in behavioral entropy between rewarded and unrewarded trials (Fig. S7f, $p > 0.08$, 95% confidence intervals for unrewarded: 19.0552, 93.1805 and rewarded: 21.3758, 100.9368, $n = 43$). To control for small movements when the animal was stopped at the well, we also selected times when the animal was stopped at the well with a two-dimensional speed equal to zero. We found that activation probability per wSWR and wSWR rate was still higher on rewarded than unrewarded trials when the animal was still (Fig. S7g,h, $p < 10^{-10}$ and $p < 0.001$, respectively). Thus, there was no evidence for behavioral differences that could explain our effects.

The initial learning of the task provides a period where the animal's behavior suggests imperfect ability to predict reward. Nonetheless, wSWR activity was enhanced on rewarded trials before the animal reached the learning criterion. We examined data from the initial period of S1 learning before the sequence switching phase of the task (days 1 through 5 or 6) using the same immobility criterion to identify times at the food wells. We once again found that rewarded trials were associated with significantly greater wSWR activation per SWR of CA3 place cells (n = 84) than unrewarded trials ($p < 10^{-10}$ for both activation probability per wSWR and per trial). Collectively, these results show that increased wSWR activity when animals receive reward is highly unlikely to be caused by differences in activity during the run to the well, the relative timing of rewarded and unrewarded trials, behavioral variability, the behavioral sequence the animals executed, or reward expectation.

Finally, we noted that confounds related to recording quality or clustering quality cannot explain our findings. Each cell was defined with a single set of cluster bounds during run and rest periods. Rewarded and unrewarded trials were intermixed during that period, so the same set of cluster bounds defined cells that were more active on rewarded trials and less active on unrewarded trials. Moreover, poor clustering would be expected to lead to differences throughout a recording period, not differences that were present only after the animal received or did not receive reward. The reward delivery system was also not a potential source of confounding signals, both because there were no electrical artifacts from this system and because the trigger for reward was issued before the animals stopped at the well.

Interaction between Reward and SWR Amplitude and Duration.

Lastly, we examined the relationship between wSWR amplitude and duration and the greater activity on rewarded trials. SWRs are recorded in the local field potential, which reflects the combined activity of single cells within several hundred microns of the recording tetrode. Because SWR activity reflects single cell activity, SWR amplitude and duration are likely related to the amount of spiking within the SWR. However the exact relation between SWR size and spiking within SWRs is unknown. We found that the amplitude and duration of wSWRs was significantly larger in rewarded than unrewarded trials (see methods; Fig. S8a, $p < 0.01$ and $p < 10^{-5}$ respectively, $n = 4427$). This is not surprising given that more cells were active during wSWRs on rewarded trials. Nonetheless, the increased activity on rewarded trials was still present after we subsampled wSWRs in rewarded trials to match the amplitude in unrewarded trials (see Supplementary methods; Fig. S8b, $p < 0.05$). Similarly, the increased activity on rewarded trials persisted when we truncated wSWRs on rewarded trials to match the duration of wSWRs on unrewarded trials (see Supplementary methods; Fig. S8b, $p < 10^{-10}$). While significant differences activation probability per wSWR persisted when we controlled for wSWR amplitude and duration, the differences between rewarded and unrewarded trials were smaller, suggesting that the wSWR amplitude and duration did have an effect.

Chapter 2

Coordinated hippocampal firing across related spatial locations develops with experience

Abstract

To learn we must identify and remember experiences uniquely but also generalize across experiences to extract common features. Hippocampal place cells can show similar firing patterns across locations, but the functional significance of this activity and the role of experience and learning in generating it are not understood. We therefore examined hippocampal place cell activity in the context of spatial tasks with multiple similar spatial trajectories. We found that path equivalent firing, where individual neurons are active in multiple similar locations, develops with experience and task learning. Furthermore, increased path equivalence was associated with increased moment-by-moment correlations between pairs of path equivalent neurons. Our results suggest that path equivalent activity represents both geometrical and learned, task-related similarities across distinct locations. Path equivalent ensembles could encode similarities among repeating elements, providing a framework for associating specific behaviors with multiple locations, while neurons without this repetitive structure maintain a distinct population code.

Introduction

The world is full of repeating elements, like city blocks, rolling hills, or trees evenly spread through a forest. However, we do not fully understand how neural

representations organize these spatial elements. On the one hand, neurons may encode the similarities among these elements to extract general principles about the environment and to facilitate the application of learned information to new experiences. On the other hand, neurons might encode each element very differently to easily distinguish between them and form unique associations with each element.

We know the hippocampus is required for spatial learning and separating between and generalizing across similar experiences (O'Keefe and Nadel, 1978;Morris et al., 1982;Nakazawa et al., 2002;McHugh et al., 2007). Previous studies of separation and generalization in the hippocampus have usually examined neural activity as animals navigated similar but distinct environments. In some cases, place cells have very different patterns of activity in similar environments (Muller and Kubie, 1987) and can even fire differently in visually identical places (Tanila, 1999). These distinct firing patterns have been termed “global remapping” (Leutgeb et al., 2005b). In other cases, place cells show “rate remapping” where they fire at comparable locations in each environment but at different peak rates (Lever et al., 2002;Leutgeb et al., 2004;Wills et al., 2005;Leutgeb et al., 2005b). Cells can also show a mixture of rate and global remapping (Skaggs and McNaughton, 1998).

Rate remapping appears to be very similar to a pattern of activity termed “path equivalence,” where neurons fire in similar locations within and across linear environments. Frank et. al. (2000) reported that entorhinal but not CA1 neurons showed “path equivalent’ activity in the context of a spatial alternation task. In contrast, recent results from recordings made while animals moved through a hairpin maze made up of multiple, connected U-shaped elements reported this sort of activity in both the entorhinal cortex (EC) and within the hippocampus (Derdikman et al., 2009). These authors further showed that this activity was associated with a resetting of the grid-cell map at the beginning of each segment of the maze.

While the presence of these patterns of similar coding is well established, important questions remain. First, it is not clear why these patterns appear in the hippocampus in some cases and not others. Neither the Derdikman et al. nor the Leutgeb et al. studies provided a behavioral readout demonstrating that the animals could distinguish among the similarly encoded places, leaving open the possibility that the presence of path equivalence or rate remapping in the hippocampus reflect confusion. Second, the functional significance of these patterns remains unknown. Frank et al. (2000) suggested that path equivalent activity represents similar behavioral demands across distinct locations. In contrast, Derdikman et al. (Derdikman et al., 2009) concluded that this type of activity could be entirely explained by the geometrical similarity of the locations and the associated grid cell inputs to the hippocampus, while Leutgeb et al. (2005b) suggested that similar coding in open field environments reflects the encoding of different episodes in the same place. We hypothesized that these similar coding patterns reflect learned generalizations across different places and episodes. To help distinguish among these possibilities we examined neural activity in CA3 and CA1 while animals performed tasks that required distinguishing among similar elements. Our results indicate that path equivalence in the hippocampus does not reflect confusion but instead represents the learned relationships among locations in the animal's environment.

Results

Behavior

Animals were trained to perform two alternation sequences in a 6 arm environment (Fig. 1). This task requires the animals to distinguish between the elements (arms). Therefore, when the animal correctly performs the task, similarities in neural coding between arms cannot be explained by a failure to distinguish among similar

locations. Furthermore, rapid learning of the alternation task requires the hippocampus, therefore hippocampal activity is likely involved in performing the task (Kim and Frank, 2009;Derdikman et al., 2009). Animals first learned a spatial alternation sequence (S1) to criterion and then learned to switch between this sequence and a second alternation sequence (S2, see methods). All six arms remained available for exploration. Animals received a liquid chocolate reward at the end of an arm if they selected the correct arm in the sequence. We initially focused our analysis on the sequence switching phase of the task where the entire environment was familiar.

When animals switched between sequences, we found that they were using a place-centered strategy as opposed to a body-centered strategy. At the beginning of each session animals were placed in the home arm of the to-be-rewarded sequence. We found that when animals were first placed in the home arm of S2 they immediately performed the previously rewarded S1 (n = 3 of 3 animals). Thus, animals used environmental cues and an allocentric reference frame to perform the task, rather than remembering a series of right or left turns based on their body reference frame. After executing S1 for several trials animals changed their behavior and eventually learned to perform S2 (Fig. 1d).

Path equivalent coding in the 6 arm maze.

We first asked whether hippocampal path equivalence was present in the context of the sequence switching task where animals must accurately represent their location to receive reward. We restricted our analysis to correct trials and found that many single cells fired in similar locations of multiple arms in the 6 arm environment. We observed repetitive activity for single cells in both 2 dimensional and linearized firing (Fig. 2a, b). Repetitive firing in different arms and trajectories became more obvious when firing rate maps were scaled so that all firing above 3 Hz is shown in the same color or linearized

firing was scaled to the peak firing on each trajectory. This scaling emphasizes firing location and disregards peak or total firing. 3 Hz was used for the firing rate maps because this criteria can be used to define the minimum peak firing rate for a place field (Karlsson and Frank, 2009b).

We found that the repetitive activity has periodic structure reflecting the structure of the environment: many cells fired at regular spatial intervals of about the length of a single trajectory. To quantify this activity for each cell (n = 107 cells with place fields on the track), we computed an autocorrelation of the cell's linearized firing rate. We normalized the linearized firing rate on each trajectory by its area and then concatenated the linearized firing for trajectories in the same turn direction. The resulting autocorrelations had peaks at zero and for many cells, at about the length of a single trajectory in the environment (~160 cm, Fig. 4c). Cells fired at the same spatial interval even though they fired at many different locations in the track. For instance, some cells fired at turns (Fig. 2a-c, left example), while other cells fired near reward wells (Fig. 2a-c, middle example) but in both cases these cells fired at ~160 cm intervals.

About half of the cells with place fields on the track had path equivalent activity. To determine the proportion of cells with significant peaks in the autocorrelation, we compared each cell's autocorrelation to 90% confidence bounds derived from the distribution of cross correlations between the linearized firing rates of different cells. Autocorrelations were computers separately for right turn and left turn trajectories. We identified the peak of regions where the autocorrelation exceeded the upper 5% confidence bound and plotted a histogram of peak locations (Fig. 2d). About half of the cells had significant peaks in either autocorrelogram at lags corresponding to the length of a single trajectory (140-180 cm) and we categorized them as "periodic" (Fig. 2e). The rest of the cells were either aperiodic, having a significant peak in either autocorrelogram at lags greater than 60 cm but not near the length of a single trajectory, or single peak,

having no significant peaks in either autocorellogram at lags greater than 60 cm. As a result we conclude that about half of the cells with place fields in the environment showed path equivalent coding. This mixture of path equivalent and non path equivalent coding was present within each session on the track, indicating that cells with different firing patterns were present simultaneously (Fig. 2f). Path equivalent activity persisted when we included all trials (not shown). To determine the number of trajectories that demonstrated path equivalence for each cell, we examined all task relevant trajectories, both correct and incorrect (see methods). Cells most often showed path equivalent coding in 2 or fewer trajectory pairs (Fig. 2g), indicating that individual cells were active in a subset of the possible path equivalent locations.

Path equivalence was not due to general trends in place field locations (e.g. place fields from many cells congregating around features like turns). We first noted that if path equivalent activity could be explained by general trends in place field locations, the cross correlation between different cells would also show peaks of activity at lags corresponding to the length of a trajectory which would have broadened the confidence bounds and therefore decreased the number of autocorrelation peaks near the length of a trajectory. We also examined this issue by computing the normalized overlap between linearized firing rates for the same cell in two different trajectories and comparing that to the normalized overlap of different cells in two different trajectories. Normalized overlap reflects the similarity of the firing location regardless of firing rate. High normalized overlap occurs when cells fire in very similar locations, while low normalized overlap occurs when cells fire in very different locations. Thus, normalized overlap measures the degree of overlap between place field locations and therefore the degree of path equivalence, as opposed to categorizing the cells as either path equivalent or not. We found that normalized overlaps of linearized firing from the same cell between different trajectories was much higher than between different cells on different trajectories (Fig.

2h, $p < 10\text{-}10$, n = 210 same, 998 different trajectory pairs, medians = 0.5574, 0.1661 from the same cell or different cells, respectively). Hence, the presence of path equivalent firing could not be explained by the clustering of place fields around particular features or more generally by many different cells firing in similar locations in different trajectories. Furthermore, these analyses also suggest that our results cannot be due to poor clustering. While poor clustering could result in cells that fired in multiple locations it would not result in cells that fire in similar locations on multiple arms of the track.

Path equivalence increases with experience

The presence of path equivalence in the context of a spatial alternation task in the 6 arm maze indicates that path equivalence does not imply confusion about location, but it is still unclear whether this pattern of firing is likely to serve a functional purpose. We therefore asked if the strength of path equivalence changed with experience, as would be expected if path equivalent activity reflects the encoding of common behavioral associations across related locations (Frank et al., 2000). We found that path equivalence increased with experiences relating the paths.

We first found that the prevalence of path equivalence increased with experience during the initial training on S1. Because animals did not know the correct sequence and performed poorly, we included all task relevant arms and both incorrect and correct trials, comparing trajectories in the same turn direction. On days 2 through 5, path equivalence for individual neurons, measured as the normalized overlaps of linearized firing across different trajectories increased as animals became more familiar with the environment and learned S1 (Fig. 3a, Kruskal-Wallis one way analysis of variance $\chi^2(4,162) = 27.6921$, $p < 0.0001$, medians = 0.2838, 0.3970, 0.5708, 0.5715 for days 2-5; day 2 was significantly smaller than days 4 and 5, $p < 0.05$ Tukey-Kramer post-hoc test). Consequently, path equivalence increased as animals learned the initial

sequence. Animals did not traverse the task arms enough on day 1 to compute linearized firing rates and normalized overlaps.

Similarly, the prevalence of path equivalence also increased when animals learned the new sequence (S2) in the now familiar environment. We examined the same set of trajectory pairs on the first, second, and third day that S2 was rewarded (the first, second, and third days of the sequence switching phase of the task). In S2, the entire environment and all of the task trajectories are familiar, as the animals fully explore the environment when initially learning S1. However some of the task trajectories are specific to S2 and are never rewarded during the initial learning of S1 (trajectories from arm D to E and E to D). As a result these trajectories do not initially have the same task relevance during the initial learning as trajectories that are included in S1. The distribution of normalized overlaps increased as the animal learned S2 and the new relationship between these trajectories and the task demands (Fig. 3b, Kruskal-Wallis one way analysis of variance $\chi^2(2,210) = 7.3580$, $p < 0.05$, medians = 0.2932, 0.2667, 0.4868 for days 1-3, day 2 was significantly smaller than day 3, $p < 0.05$ Tukey-Kramer post-hoc test). Thus, normalized overlap increases between related trajectories as animals learn each sequence.

Ensembles of cells are recruited together in the path equivalent code

The generalization evident in path equivalent activity led us to wonder whether CA3 place cells that develop path equivalent firing are recruited independently or whether these cells arise as part of functional ensembles that generalize across locations. We found that path equivalent cells that repeat together on multiple trajectories have correlated moment to moment activity, suggesting they are connected or receive similar inputs. To identify elements of cell ensembles we examined correlations between pairs of cells' moment to moment variability. High correlations

between cells' moment to moment variability suggest the cells receive similar inputs or are connected to each other either directly or indirectly (Lee et al., 1998;Shadlen and Newsome, 1998). We calculated the difference between each cell's expected number of spikes (based on linearized firing rates and the animal's position) and the actual number of spikes in each time bin (see methods, Fig. 4a). We then correlated these residuals between pairs of cells both in single trajectories and over the entire session when animals were running.

If cells are part of a functional ensemble and receive similar inputs or are connected to each other, we would expect that they have both similar place field locations and correlated fluctuations in moment to moment activity. Cells with only overlapping place field locations, however, may or may not be part of the same ensemble as they could have similar receptive fields by chance. Furthermore, if path equivalent cells are part of functional ensembles, we would expect the cells to have similar place field locations across multiple paths. Consistent with these expectations, we found that cell pairs with overlapping place fields and correlated moment to moment activity in one location are more likely to have overlapping place fields in another location. We identified pairs of cells with overlap in one trajectory greater than 0.3 and peak firing rates in both trajectories of at least 0.5 Hz and found strong evidence for this sort of relationship (see Fig. 4b for examples of cell pairs with low and high correlations of residuals). Across the population, we found that the residual correlation in the first trajectory was related to the overlap and normalized overlap in the other trajectory (Normalized overlap: Fig. 4c, rho = 0.3229, $p > 0.001$, robust fit slope = 0.24382, $p < 0.02$; overlap: rho = 0.27873, $p < 0.005$, robust fit slope = 0.24196, $p < 0.005$, n = 174 trajectory pairs). Therefore cells with overlapping place fields that vary together in one trajectory are more likely to have overlapping fields in another trajectory.

We also found that pairs of cells with highly correlated moment to moment variability throughout the session were more likely to have high overlap on multiple trajectories (see Fig. 6b for examples). We examined the distributions of the correlations of residuals between all cell pairs and pairs with low or high overlap (Fig. 4d-f). Cell pairs with overlapping fields (overlap > 0.3) in at least 3 trajectories (high overlap) had higher residual correlations than cell pairs with low overlap (overlap < 0.2) in at least 3 trajectories ($p < 0.0005$). Additionally, cell pairs with larger numbers of trajectories with overlapping fields tended to have higher correlations of residuals for each cell pair (Fig 4g, ANOVA: $F(4,30) = 2.9704$, $p < 0.05$, see methods). Only cells with a minimum total overlap across all trajectories of at least 0.1 were included for these analyses; cell pairs could overlap on a single trajectory or on multiple trajectories. While this excluded many cell pairs, it ensured that the cells that were included had some overlapping fields on at least one trajectory on which to compute residual correlation.

Furthermore, cell pairs with stronger moment to moment correlations were more likely to show path equivalent firing. If highly correlated pairs form coordinated, path equivalent ensembles, we would expect that cell pairs with positively correlated firing on a 500 ms time scale were likely to be made up of cells that both showed path equivalent firing. Similarly, negatively correlated pairs would be made of cells where one or both did not show path equivalent activity. Indeed, we found that approximately half of the cell pairs with positive residual correlations were pairs in which both cells were periodic, while only about one quarter of the cell pairs with negative residual correlations were such pairs (z test of proportions, $z = 2.2054$, $p < 0.05$, $n = 40$ cells with positive, 41 with negative residual correlations). Thus, cells whose rates co-vary on a moment by moment basis tend to show path equivalent activity across multiple trajectories.

The link between path equivalence and correlated activity was further strengthened when we examined changes in correlated firing with experience. Across

the population the correlation of residuals increased as animals gained more experience in S2. Including all cell pairs with an overlap of at least 0.1 across all trajectories, we computed the correlation of residuals for the entire session when animals ran faster than 3 cm/sec and performed S2 on the first two days of exposure to S2 or 3 or more days of exposure. We grouped days to include more data in each group. We also compared this to the last two days of exposure to S1 during the initial training on S1 before the sequence switching phase of the task began. We found that the correlations of residuals were higher when animals had 3 or more days of exposure to the sequence than on days 1 and 2 (Fig. 4h ANOVA $F(2,115)$: 5.9582, $p < 0.005$, medians = -0.1117, 0.0630, 0.0348 for S2 days 1-2, days 3+, and S1 last 2 days preswitch, S2 days 1-2 significantly less than S2 days 3+ and S1 last 2 days preswitch, $p < 0.05$ Tukey-Kramer post-hoc test). Because place fields form and fluctuate during the first few days of exposure to an environment (Wilson and McNaughton, 1993;Frank et al., 2004;Karlsson and Frank, 2008) we did not examine residual correlations during the first few days of exposure to S1.

We then performed three important control analyses. First, the above results could not be explained by a simple relationship between place field overlap and correlated moment to moment activity in the same location. We asked whether a high correlation of residuals was a general feature of overlapping place fields: covariance among cells with overlapping fields could better explain the relationship between correlations of residuals in one trajectory and overlap in a different trajectory. We found that this was not a likely explanation, as there was no significant relationship between the correlations of residuals and overlap in the same trajectory for cells with a peak firing rate of at least 0.5Hz in the trajectory (SFig. 1a, rho = 0.0456, $p > 0.7$, $n = 729$ trajectories). We also examined the relationship between overlap and correlation of residuals only in cases with overlap greater than 0.3, as we noticed that the distribution

of the residual correlation seemed qualitatively different at overlaps above and below 0.3. Again, we found no significant relationship between the correlation of residuals and overlap in the same trajectory (rho = 0.1273, $p > 0.18$, n = 174 trajectories). These results indicate that there is no relationship between place field overlap and the correlation of residuals within a single trajectory. Instead, this relationship is only evident when we examine path equivalent pairs across multiple trajectories.

Second, we determined that the relationship between overlap and correlation of residuals was not due to more time bins available for analyses in cells with high overlap. The larger the overlap between linearized firing curves, the greater the number of time bins in which we could compute a meaningful residual correlation. If cell pairs with little overlap had lower residual correlation simply because of fewer available time bins, we would expect a positive relationship between the number of time bins and the residual correlation value. Again we only examined cell pairs with an overlap of at least 0.1 across all trajectories, though we found similar results if we included all cell pairs. As we would expect to see this relationship with few time bins, we examined this relationship for cases with 20 (the minimum allowed number of time bins) to 100 time bins. We found no significant relationship between the number of time bins and the residuals correlation (SFig 1b, rho = 0.12102, $p > 0.37$, n = 97 cells). We conclude that correlation in a cell pairs' moment to moment variability was not due to measurement error or bias but is instead reflects the short time scale covariation of path equivalent cell pairs.

Finally, to control for possible contamination due to clustering errors (Quirk and Wilson, 1999), we repeated all of these residual correlation analyses for cells recorded on different tetrodes. We found similar results for all of the analyses, although one result was no longer significant. The relationship between cell pairs with larger numbers of trajectories with overlapping fields still tended to have higher correlations of residuals for each cell pair but the results were no longer significant (ANOVA: $F(4,13) = 0.66$, p =

0.6298). This is not surprising as restricting to cells on different tetrodes reduces the amount of data substantially. These correlations in moment to moment firing of path equivalent cells indicate that these cells may be part of functional ensembles.

Differences in firing rate and population coding associated with path equivalence

The above analyses suggest that path equivalence could encode similarities across distinct locations, providing a framework to form common behavioral associations with multiple places. However, animals would still need to distinguish among these places to perform spatial tasks. We found that the population as a whole could distinguish the arms of the six arm environment from one another despite the path equivalent coding of individual cells. We first noted that while many cells fired in very similar locations in different trajectories in a repetitive environment, they fired at different peak rates, similar to reports of rate remapping. We compared peak linearized firing rates between two trajectories when the firing on those trajectories was in similar locations (normalized overlap greater than 0.3, see methods). We took a ratio of the smaller peak divided by the larger peak firing rate. We found the peak rate ratio between different trajectories with overlapping firing locations was lower than if we compared the peak firing on the same trajectory in different run epochs on the same day (Fig. 5a, $p < 0.02$, medians = 0.3695, 0.6072, n = 150, 197 trajectory pairs for diff trajectories in the same epoch or the same trajectory in different epochs, respectively). We found similar results if we used a higher normalized overlap requirement to select cells with similar firing locations in different trajectories. Thus cells that fire in similar locations in different trajectories fire at different peak rates.

Additionally, cells with path equivalent coding but differences in peak firing rate and cells without path equivalent coding generate distinct population codes in each trajectory. To test this, we compared population activity in different trajectories by

computing a population vector of firing rates of multiple single cells in each 10 cm of bin of each trajectory. We then correlated vectors of corresponding bins from trajectories in the same turn direction and averaged the correlations for each trajectory. We found the resulting population vector correlations were significantly lower than if we compared the same trajectory in different epochs (Fig. 5b, $p < 0.02$, $n = 10$ averaged trajectory pairs), revealing significant differences between trajectories in the population code. The population vector correlations were also higher than if the cell identity was shuffled, revealing that the repetitive coding structure in some cells does result in higher correlation in the population activity (Fig 5b, $p < 0.0001$, $n = 10$ averaged trajectory pairs). This suggests that a population with both path equivalent (rate remapping) and non-path equivalent coding may simultaneously generalize and separate.

Path equivalence is less common in the 3 arm maze

If similar coding patterns reflect learned generalizations, then repetitive coding would be less common in a simpler, less repetitive environment. Indeed, a previous report has indicated that path equivalence was not prevalent in CA1 cells recorded while animals performed a spatial alternation task on a 3 arm or "W-track" maze (Frank et al., 2000). The environment in that study was highly familiar and no CA3 cells were sampled, however, and previous studies have reported differences in pattern completion and separation CA3 and CA1 (Leutgeb et al., 2004). Consequently we wondered if path equivalent activity is absent in a 3 arm maze or if it only occurs in CA3 or during learning in such a maze. We therefore examined 3 arm maze data from both CA3 and CA1 cells (Karlsson and Frank, 2008;Karlsson and Frank, 2009b). The alternation task in the three arm maze was the same as S1 or S2 in the six arm maze, however there were fewer arms and none of the arms were exactly geometrically identical (SFig. 2). In particular, in the six arm maze, all the arms end in a T junction, while in the 3 arm maze

the middle arm ends in a T, and the left and right arms end in right turn only or left turn only L shapes.

We found that most cells had single place fields in the 3 arm environment (Fig. 6a,b, left and right examples), though some had multiple fields in somewhat similar locations in different arms (Fig. 6a,b, middle example). Cells that had a single place field could fire in a single trajectory (Fig. 6a,b, left example) or multiple trajectories as the animal traversed the same location in multiple trajectories (Fig. 6a,b, middle and right examples). However, these cells generally did not have peaks in the autocorrelogram at the length of a single trajectory (average 180 cm), as we would expect if the cells' firing was periodic (Fig. 6c). Across the population, both CA3 and CA1 cells with multiple place fields had some tendency to be active at lags of ~180 cm (Fig. 7a,d), but the overall proportion of periodic cells (160-200cm) was lower than in the six arm maze (Fig. 7b, e; $p < 0.05$ χ^2 test for both CA3 and CA1 cells from the three arm maze compared to CA3 cells in the six arm maze).

Similarly, in the 3 arm maze, the spatial locations of firing from the same cell on different trajectories were more similar to firing from different cells on different trajectories. Unlike in the six arm maze, the distribution of normalized overlaps was not bimodal, instead generally decreasing from 0 to 1 (Fig. 7c,f). The same vs. different cell distributions did have slightly different medians (CA3: medians = 0.1697, 0.1211, n = 260, 1692 same cell and different cell distributions, respectively, $p < 0.05$; CA1: medians = 0.1557, 0.1129, n = 309, 2659 same cell and different cell distributions, respectively, $p < 0.005$), but the CA3 same cell distribution from the 3 arm maze had a much lower normalized overlap than that from the 6 arm maze (median = 0.5574 for 6 arm maze same cell, $p < 10^{-5}$). We conclude that path equivalent coding in CA3 and CA1 is present, but much less common in the 3 arm maze as compared to the six arm maze

even though the animals were executing the same alternation task. There were no differences in the prevalence of path equivalent coding between CA3 and CA1.

Furthermore in the 3 arm maze, normalized overlap did not change with experience. We examined normalized overlap of CA3 cells between trajectory pairs in track 2 (T2) in the 3 arm maze, from the animals first exposure to T2 to 5 days of exposure. We found no significant change in normalized overlap across days or between the first day of T2 exposure and exposures to the familiar track, track 1 (Fig. 7g, Kruskal-Wallis one way analysis of variance $\chi^2(4,108) = 7.8749$, $p > 0.09$, medians = 0.3454, 0.3146, 0.2682, 0.2333, 0.1636 for days 1-5 on T2, no significant differences between days 1-5 on T1, $p > 0.05$ Tukey-Kramer post-hoc test). In fact, the distributions of normalized overlap values for the highly familiar T2 was similar to those of the novel 6-arm maze (SFig. 3), indicating that both environments began with comparable overlaps across trajectories, but while the overlaps remained stable in the 3 arm maze, they evolved with experience in the 6 arm maze.

Path equivalence in CA1 in Multiple-U Environment

Finally, to better understand the role of geometrical versus local cue based similarity in path equivalent activity, we examined CA1 neurons in two animals traversing a multiple-U environment where each U had different local visual and olfactory cues (SFig. 4). We examined these parameters because previous studies have shown that subtle environmental manipulations can lead to distinct patterns of rate and global remapping (Skaggs and McNaughton, 1998;Lever et al., 2002;Leutgeb et al., 2004;Leutgeb et al., 2005a;Leutgeb et al., 2005b). In the multiple-U environment the local sensory cues differed substantially across each U (see methods) but despite these differences, path equivalent coding was common. We found many single cells fired in similar locations in the different U's that make up the track, while others only fired in a

single location (Fig. 8a,b, left and right example, respectively). The normalized overlap between linearized firing from the same cell on different trajectories was significantly higher than firing from different cells, though only 24 cells reached the criteria for inclusion in analysis (Fig 8c., $p < 0.001$). Because of the low number of simultaneously recorded cells, we did not perform the periodicity analysis as confidence bounds would be too noisy with so few cells. Nonetheless, these results establish that path equivalence also occurs in CA1, is not due to the specifics of the 6 arm paradigm, and, as each U had very different visual and olfactory cues, is not likely to be due solely to similarity of local sensory cues.

Discussion

We have shown that hippocampal neurons fire in multiple similar locations in repetitive environments even when animals must behaviorally distinguish among the arms. This path equivalent coding increases as animals learn the rewarded sequence and the relationships between paths and similar behaviors. Furthermore, path equivalent firing is not simply due to single cells acting independently. Rather pairs of cells that vary together from moment to moment repeat together in multiple segments, suggesting that ensembles of cells are repeatedly recruited together. These correlations also increased with experience, pointing to a role for learning in the development of path equivalent coding. While path equivalence is common in the two environments with repeating elements that we examined, cells that fire in multiple similar locations in an environment fire at different peak rates in different locations. Given a subpopulation of cells that continued to fire differently across different trajectories, our results argue that path equivalence could provide important information about the similarity across different locations while, at the population level, the system could still distinguish among those locations.

These patterns of activity cannot be explained by confusion, lack of distinguishing sensory cues or clustering errors. First, correct behavioral performance in the six arm task requires that the animal make different choices depending on current and past locations. While it is conceivable that the animal could make these choices without an accurate hippocampal representation of location, our behavioral data and previous results suggest that the hippocampus is involved in learning this task. When animals were first placed in the home arm of S2 they immediately moved to the adjacent arm and performed the previously rewarded S1. This indicates that the animals were using an allocentric strategy associated with hippocampal dependence (Packard and McGaugh, 1996). In addition, this task is based on a the simpler W-track alternation task where the hippocampus is required for rapid learning (Kim and Frank, 2009). Further the six arm task requires flexible changes in behavior in response to changing reward contingencies. Because this sort of flexibility generally requires the hippocampus (Hsiao and Isaacson, 1971;Hirsh et al., 1978;Ainge et al., 2007b), hippocampal activity is likely to be important for learning to switch between sequences. Thus there is good reason to believe that the hippocampus was engaged during the performance of the 6-arm task and that animals distinguished among the arms of the environment.

Second, the prevalence of path equivalent activity in the multiple-U track argues that this activity is not solely due to a lack of distinguishing sensory cues. Each U of that environment was bordered by walls with highly distinctive visual patterns and was scented uniquely. Given that animals can use local visual and odor cues to distinguish among locations (Anderson and Jeffery, 2003), it is unlikely that path equivalent activity in the hippocampus was due solely to sensory similarity.

Finally, there is no reason to suspect that errors in clustering could have led to our findings. Clustering errors would lead to errors such that a single clustered “unit”

would have multiple place fields corresponding to the place fields of the single cells that were erroneously combined. However fields from different single cells erroneously clustered as one cell would not tend to show path equivalence as there is little or no topographic organization of spatial receptive fields in the hippocampus (Redish et al., 2001).

While these cells fire in similar locations in multiple trajectories, they fire at different peak rates. This is consistent with the phenomenon of rate remapping (Leutgeb et al., 2005b;Leutgeb et al., 2005c;Leutgeb et al., 2006). We suggest that these two patterns of neural activity are manifestations of the same basic phenomenon where individual cells fire in related locations within and across environments. At the same time, we found a mixture of path equivalent, rate remapping activity as well as cells that fired in different location or not at all on different trajectories. It is not clear whether the same mixture of path equivalence and global remapping was seen in hippocampal neurons recorded in the context of the hairpin maze (Derdikman et al., 2009), as that study did not classify neurons individually.

Overall, our results provide new insights into both the generation and the functional significance of path equivalent activity. Derdikman (Derdikman et al., 2009) demonstrated that path equivalent activity in the hippocampus was associated with resets of the EC grid pattern across each geometrical element of the environment. Our observation of both path equivalence and global remapping across trajectories raises the interesting possibility that some elements of the entorhinal network may remain coherent across trajectories while other elements might shift. Our results also indicate that both the geometry of the environment and experience contribute to path equivalence. In the three arm maze, path equivalent activity in CA3 and CA1 was relatively uncommon, despite its prevalence in the EC in the same task (Frank et al., 2000). Thus, path equivalent activity does not always propagate from the EC to the hippocampus. Instead,

our results are consistent with the notion that the number of repeating elements determines whether this sort of similar activity across paths is seen in the hippocampus.

Further, our finding that path equivalence increases with experience, and in particular that it increased as animals learned the novel Sequence 2 in the familiar six arm maze, demonstrates that path equivalence in the hippocampus depends on experiences relating the paths. Given that the animals had full access to the geometrical structure of the environment for at least five days prior to their first exposure to S2, the change in path equivalence with experience is more closely aligned to changes in the behavioral significance of individual trajectories than to their stable geometrical character. Therefore path equivalence increased as animals learned relationships between paths and executed similar behaviors across paths, whether the environment was novel or familiar.

Our results also indicate that hippocampal cells with similar activity across locations are more correlated with one another, suggesting that experience drives the formation of path equivalent ensembles. We found that for a given pair of cells with overlapping place fields in one trajectory, higher correlated residual rates were associated with higher overlaps in another trajectory. Similarly, pairs of cells with overlapping, path equivalent fields across trajectories tended to have high overall short-time scale correlations of residuals. So while the similarity of cells' receptive fields in a single location was not predictive of moment to moment correlations, similarities over multiple locations were. To our knowledge, this is the first demonstration that 1) the firing properties of two place cells in one location can predict their patterns of spatial activity in another location and 2) that correlations in moment to moment activity is related to cells firing similarly over multiple locations but not a single location. High correlations of residuals restricted to a subset of cells are suggestive of common input or

direct synaptic connectivity, and thus our results raise the interesting possibility that hippocampal pairs with high residual correlations are part of specific neural ensembles.

Given the clear increases in path equivalence and correlation with experience, our data argue that the expression of path equivalence in CA3 is experience dependent. Thus, in environments with many repeating elements, a subset of CA3 neurons evolves with experience to represent general features of the environment. These findings suggest that this path equivalent activity could be a mechanism to generalize across related experiences in the hippocampus. General information about related experiences could then be used to encode appropriate behaviors across experiences or to predict new experiences.

Finally, as most path equivalent cells were active at different rates across a subset of the related locations and as about half of the cells we recorded fired in only one place or in unrelated locations, each location was still associated with a distinct population pattern of activity. In principle animals could therefore extract both a relative location within each trajectory in the environment as well as a global location from the ensemble of active place cells (Fenton et al., 2008). From this perspective, path equivalence or rate remapping would presumably reflect a sense of similarity across locations that could, in the context of specific behavioral demands, allow the animal to generalize across different elements.

Methods

We combined data from three different studies to examine remapping in linear environments. These studies varied the behavioral task, the geometry of the

environment and the recording location, allowing us to determine which factors influenced place cell activity.

Six-arm maze: Pretraining and data collection

Three male Long-Evans rats were handled and food deprived to 85-90% of baseline weight. Animals were initially trained to run back and forth on a linear track between food wells where liquid chocolate reward was delivered. Pretraining took place in a different room from the recording room. One of the animals was then pretrained on S1 in the recording room, while two animals were not exposed to the behavioral task until recording began (Fig. 1a, see Supp Table 1 for details). Following pretraining animals were implanted with a microdrive array containing 16 independently movable tetrodes targeting CA3 (-3.6 mm AP; 3.4mm L) using previously described methods (Karlsson and Frank, 2008). Over the next 7 – 10 days tetrodes were lowered first to CA1 and then to CA3. CA3 was identified by depth and the characteristic EEG waveforms on each recording tetrode. Tetrodes were then lowered to the expected depth of CA3 until the cell layer was reached. Ripples and theta were detected in CA3. For all animals a reference tetrode was positioned in the corpus callosum. All neural signals were recorded relative to that reference to eliminate muscle artifacts from the recordings.

Electrode positions were confirmed by histology (Fig. 1b). For one animal, electrode lesions were made at the end of each tetrode and later confirmed to be in the CA3 pyramidal cell layer. For two animals, the microdrive fell off before lesions could be made. In these animals we were able to confirm that the implant site was over dorsal CA3 and that the depths were consistent with CA3 recordings. Furthermore the EEG signatures characteristic during adjusting of passing through CA1 then traveling to CA3

were similar in all animals. In particular, as tetrodes traveled to the ventral side of the CA1 layer, the EEG amplitude greatly increased (Fig. 1c).

Six-arm maze: Behavior

These animals learned to perform a sequence switching task in a six arm maze (Fig. 1a). This task requires that the animal remember and select a correct arm out of a set of geometrically similar, although visually distinct arms. Therefore if the animal correctly performs the task, we can conclude that the animal can distinguish between these arms. Animals were rewarded for correct trajectories with liquid chocolate in reward wells at the end of the arm by an electrically triggered solenoid delivery system. The track included 4 sequence arms, B C D and E, and one extra arm on each end, A and F. Arms were separated by vertical walls (0.6 cm thick, 24 cm tall and 81 cm long). Distal cues were visible above these walls, at either end of each arm and along the straight section connecting different arms. Circles indicate food wells where animals received liquid chocolate reward in arms B through E. Colored arrows indicate trajectories included in Sequence 1 (purple) and Sequence 2 (orange).

The task consists of two rules. First, a visit to the home arm (arm C in S1 and arm D in S2) was rewarded when the animal came from any other arm (inbound trajectories). Second, a visit to an arm adjacent to the home arm was rewarded when the animal came from the center arm after having previously visited the opposite adjacent arm (outbound alternation). Consecutive repeat visits to the same food well were never rewarded. Together, these rules defined a correct cyclical sequence of food-well visits (Fig. 1a): right, center, left, center, right, center, left, center, etc (Frank et al., 2000;Kim and Frank, 2009). If the animal visited an arm not included in the rewarded sequence (e.g. arm A, E or F for S1), the animal was rewarded upon returning to the home arm. During the initial learning of the task, animals learn the inbound component first and then

learn to alternate on outbound trajectories. As a result, once animals learn to perform the outbound trajectories with high accuracy they are generally performing the entire sequence accurately. Rapid learning of the alternation task depends on an intact hippocampus (Kim and Frank, 2009), as does the ability to flexibly alternate among reward contingencies (Hsiao and Isaacson, 1971;Hirsh et al., 1978;Ainge et al., 2007). Thus, hippocampal activity is likely to be important for correct behavior in this task.

During each run session the animal was placed in the home arm of the to-be-rewarded sequence (arm C for Sequence 1 and arm D for Sequence 2) but no cues indicated which sequence was rewarded other than the presence or absence of reward at the food wells. Each run session was between 20 and 30 minutes long; one animal performed two sessions and two animals performed three sessions per day. Rest sessions in a high walled box preceded and followed each run session. Once the animal performed S1 with 80% accuracy, measured across a run session, or had 6 full days of training and was above 75% accurate, the sequence switching phase of the task commenced. On the first day of sequence switching animals first performed one session where S1 was rewarded. Then in the second session, reward contingencies changed such that S2 was rewarded. All subsequent sessions alternated between rewarding S1 and S2 within each day.

We distinguished between “accurate” responses that were consistent with the rules of S1 or S2, and correct (rewarded) responses. This allowed us to score behavior according to the rules of both sequences simultaneously. To illustrate the behavior, we plotted a 20 trial moving average applied to all (Fig. 1d).

3-arm maze: Pretraining and data collection

The data for the 3-arm maze was data recorded from animals 5, 6, and 7 of Karlsson and Frank (2008). Briefly, three male Long-Evans rats (500-600g) were food

deprived to 85 – 90% of their baseline weight and trained to run on a linear track with one reward well at each end of the track. Linear track pre-training was performed in a different room from the recording experiments. After the animals were accustomed to behaving for liquid reward (sweetened condensed milk), they were implanted with a microdrive array containing 30 independently movable tetrodes. After 5-6 days of recovery, animals were once again food deprived to 85% of their baseline weight. The tetrodes were arranged bilaterally in two 15 tetrode groups centered at AP -3.7 mm, ML ±3.7 mm. Each group was located inside an oval cannula whose major axis was oriented at a 45 degree angle to the midline with the more posterior tip of the oval closer to the midline. Tetrodes in the anterior and lateral portion of each group targeted lateral CA3 while more posterior and medial tetrodes targeted CA1.

3-arm maze: Behavior

Each recording day consisted of two or three 15-minute run sessions in W-shaped tracks, with rest sessions in a black box before and after each run. The two tracks were geometrically identical but visually distinct and were open to the room but separated from one another by a black barrier (Fig. 2). The tracks had one reward well at the end of each arm, and animals learned to perform a continuous alternation task where, starting from the center arm, they alternated visits to each outer arm for liquid reward (center-left-center-right, and so on (Frank et al., 2000;Frank et al., 2004;Kim and Frank, 2009)). The correct alternation sequence in the 3-arm maze was therefore identical to S1 or S2 in the 6-arm maze. Errors were not rewarded, and following an incorrect choice of an outer arm no rewards were given until the animal returned to the center arm. Animals were pretrained on track 1 (T1) for six days and then ran on both T1 and track 2 (T2) from days 7 onward. Recordings began on day 7. These animals'

familiarity with T1 was similar to the animals on the 6 arm maze's familiarity with that track and T2 was initially novel.

Multiple-U maze: Pretraining and data collection

Two male Long-Evans rats were handled and food deprived to 85-90% of baseline weight and trained to run on an exposed tabletop for liquid chocolate reward. Pretraining took place in the same room used for recording. After the animals were accustomed to behaving for liquid reward on the exposed tabletop, they were trained on the multiple-U (Fig. 3). Based on our experience, the multiple-U task is quite difficult for animals to learn quickly. We therefore introduced the animals to a very short multiple-U (50 cm long, not shown) and then expanded the environment in two stages. We trained the animals according to a 12-day protocol, and on each day the animal ran for either two or three 20-minute sessions on the track, and rested in a black box before and after each run session. Days 1-4 of training consisted of two 20-minute runs on the shortest version of the multiple-U (50 cm long). Days 5-8 consisted of one session on the shortest configuration of the multiple-U, a second session on the medium-length configuration (100 cm long), and a final session on the shortest configuration. On days 9-12 the animals ran the first session on the medium-length configuration, the second session on the longest configuration (150 cm long), and the final session on the medium length (Fig. 3).

Following pretraining animals were implanted with a microdrive array containing 16 independently movable tetrodes targeting CA1 (AV -3.6, L 2.2) using previously described methods (Karlsson and Frank, 2008). Over the next 7-10 days the tetrodes were lowered to CA1. 7 days after the electrodes were implanted, the animals were run on the open table for liquid chocolate reward until they would eat continuously for 2 5-minute sessions. We began the experiment after the animals reached this behavioral

criterion and the electrodes reached the cell layer. Tetrode positions were adjusted after daily recording sessions for all tetrodes that had poor unit recordings. On rare occasions some tetrodes were moved before recording sessions, but never within 4 h of recording.

Multiple-U track: Behavior

On each recording day the animals ran for 3 20-minute sessions in the medium – long – medium configuration (Fig. 3). The first session was the medium-length configuration, the second was the longest configuration, and the third was the medium length. All data included for analysis were from the long configuration. The animals were allowed to rest in a black box for 20 minutes before and after each session and we recorded continuously through the rest and behavioral sessions.

Data Collection and Processing

All data were collected using the NSpike data acquisition system (L. Frank, J. MacArthur). The animal's position was tracked with an infrared diode array attached to the animal's preamplifier and was reconstructed using a semi-automated analysis of digital video of the experiment with custom-written software. Spike data were sampled at 30 KHz, digitally filtered between 600 and 6 KHz (2 pole Bessel for high and low pass) and threshold crossing events were saved to disk. Continuous LFP data from all tetrodes was sampled at 1.5 KHz, digitally filtered between 0.5 and 400 Hz and saved to disk.

SWRs were identified as described previously (Karlsson and Frank, 2009b). Briefly, LFPs were recorded from one channel of each tetrode. The LFP signal was band pass filtered between 150-250 Hz and an envelope was determined by Hilbert transform. SWR events were detected if the envelope exceeded a threshold of mean + 3 stdev for at least 15 ms. Events included times around the triggering event during which the

envelope exceeded the mean. SWR amplitude was measured in standard deviations from baseline.

After neural data was collected, individual units were identified by clustering spikes using peak amplitude and spike width as variables. All spike sorting was done using custom software (MatClust, M. Karlsson). It was generally possible to use a single set of cluster bounds defined in amplitude and width space to isolate units across an entire recording session. In the minority of cases where there was a slight shift in amplitudes across time, units (putative single neurons) were clustered only when that shift was coherent across multiple clusters and where plots of amplitude vs. time showed a smooth shift. Only well isolated cells with tightly clustered spikes and clear refractory periods were included. We did not attempt to match cells across days, so in some cases the same cell may have been recorded across multiple days. All analyses were restricted to putative principal neurons identified using standard criteria (Fox and Ranck, 1981;Frank et al., 2001). The total number of cells recorded was as follows: for 6 arm maze during the switching phase of the task n = 100, 42, and 128 for animals 1, 2 and 3 respectively, during initial learning of sequence 1, n = 92 and 160 for animals 2 and 3 respectively; for 3 arm maze n = 407, 607, and 33 cells for animals 5, 6, and 7 respectively, for multiple-U track n = 22, 13 cells for animals MU1 and MU2.

To visualize neural activity across the environments, two-dimensional occupancy-normalized spatial rate maps were constructed with 1-cm square bins of spike count and occupancy, both smoothed with a two-dimensional Gaussian (1.5 cm s.d.). These maps include times when the animal traveled in both directions in each arm and were used for visualization only.

For analysis of place field activity, we calculated the 'linearized' activity of each cell. Only times when animals were running forward at least 3 cm/sec were included. The behavioral data were separated into different spatial trajectories (e.g. A to B, B to A,

B to C, …) and the animal's linear position was measured as the distance in cm along the track from the reward site on the start arm. All the trials when the animal was on that trajectory were included to calculate occupancy normalized firing rate maps. We used 2 cm spatial bins and smoothed with a 4 cm standard deviation Gaussian curve with a total extent of 20 cm. Bins with an occupancy less than 0.1 seconds were excluded. Cells with a peak spatial rate greater than 3 Hz were considered to have a place field on the track and included in further analyses (cells with place fields: n = 107 cells in the 6 arm maze during the switching phase of the task, n = 108 cells in the 6 arm maze during initial learning of sequence 1, n = 130 cells in T1 in the 3 arm maze, n = 131 cells in T2 in the 3 arm maze, n = 24 cells in the multiple-U maze).

We performed a number of analyses to compare single cells' firing in different trajectories. For the 3 and 6 arm mazes only correct trials were included. Results were similar if all trials were included. For all mazes only trajectories in the same turn direction were compared (e.g. right turn trajectories or left turn trajectories) as these are the most similar. This also eliminates the problem of comparing trajectories where the animal crosses the same place in the same direction for one of the arms in the trajectories, as is the case for trajectories that either both start or both terminate in the center arm of the sequence.

To determine if place fields occurred in similar locations in multiple trajectories, we computed the spatial autocorrelation of each cell's linearized firing. Linearized firing of each trajectory was normalized by its area so that firing on each trajectory had a total area of one. This allowed us to detect similarities in firing location regardless of firing rate. Trajectories in the same turn direction were concatenated, and the autocorrelation was calculated. Each correlation was scaled such that the autocorrelation at zero was one. For comparison, we computed the cross correlation between the concatenated normalized linearized firing rate of different cells recorded in the same session. If

trajectories had excluded bins in the linearized firing due to low occupancy, they were excluded from the analysis. We computed the mean and 90% confidence bounds for the cross correlations at each spatial lag. To determine if cells had peaks in the autocorrelation away from zero, we determined which spatial lags had autocorrelation values above the confidence bounds of the cross correlations. By definition, a place cell will have a peak in the autocorrelation at lags near 0, so we identified significant lags greater than 60 cm or less than -60 cm to detect peaks not due to autocorrelation near 0 lag for each cell. For each cell two autocorrelations were computed, one for right turn trajectories and one for left turn trajectories. Significant lags were identified on both autocorrelations to ensure that periodic cells that had directional firing were properly identified as periodic. Cells with peaks that occurred at lags within ± 20 cm of the length of a trajectory (140-180 cm in the 6 arm maze and 150-190 cm in the 3 arm maze) in either of the two autocorrelations (left turn or right turn trajectories) were considered "periodic" because these cells had firing peaks at distances similar to the length of trajectories in the environment. Cells with significant peaks greater than 60 cm but outside of the periodic range were deemed "aperiodic", while cells with no significant peak greater than 60 cm in both autocorrelations were categorized as "single peak" cells. When computing the proportion of periodic cells per session, only sessions with at least 5 isolated cells with place fields on the track were included and all session per day per included.

To determine the number of significantly path equivalent trajectory pairs per cell in the six arm maze, we included firing on all task relevant trajectories (trajectories from arm B to C, C to B, C to D, D to C, D to E, and E to D). We compared trajectories in the same turn direction, creating a total of 6 trajectory pairs. For each pair, the autocorrelations were computed as described above. Pairs with a significant peak in the periodic range were considered path equivalent.

We also computed the normalized overlap and peak rate ratio between trajectories in the same turn direction. For normalized overlap, each cell's linearized firing on each trajectory was normalized by its area and the overlap between trajectories in the same turn direction was computed. Overlap was calculated according to a previously established method (Battaglia et al., 2004). This revealed similarities in firing location in related trajectories regardless of firing rate. Normalized overlap values close to one indicated very similar firing locations while values close to zero indicated very different firing locations. To define cases when firing on two trajectories was in similar locations, we examined cases when the normalized overlap was greater than 0.3. We selected this criterion because the distributions of normalized overlaps of linearized firing rate from the same cell and different cells cross at this point (see Fig. 5d), but our results were similar for other cutoff values. To examine differences in firing rate, we computed the peak firing rate of the linearized firing of each trajectory. The smaller firing rate was divided by the larger firing rate to create a peak firing rate ratio. Ratios close to one had very similar peak firing rates while ratios close to zero had very different peak firing rates in the compared trajectories.

We also examined moment-by-moment variability in pairs of cells to determine whether there was evidence for sets of path equivalent cells firing together in organized ensembles. To determine if neurons' trial by trial variability was correlated, we adapted an approach from Schoppik et al. (Schoppik et al., 2008) where we examined the "noise" correlations defined as the correlations of firing rate residuals. The residuals were calculated as the difference, in each 500 ms window, between the predicted number of spikes from the linearized place field and the actual number of spikes recorded (see below). This approach has a number of advantages over standard cross-correlation techniques. In particular, by examining the correlations of the residuals, we avoid misclassifying increases in overall firing rate due to place field shape as correlated

firing. Instead, measuring the correlation of the residuals of firing asks whether the fluctuations of each neuron about its mean rate are related, as would be the case if the neurons are part of a functional ensemble that receives common inputs.

We divided each recording session into 500 ms bins and for each cell we calculated the expected number of spikes in each bin. The expected number of spikes was calculated by computing the expected firing rate in 33 ms bins based on the animal's location and the linearized firing curves and integrating that rate across each 500 ms bin. The 500 ms bin size was chosen to be large enough to effectively average over variability due to short time scale bursting and the modulation of the ~8 Hz theta rhythm.

We then calculated the residuals: the difference between the expected number of spikes and the actual number of spikes recorded in each time bin. Only bins where the expected firing rate exceeded 0.5 Hz were included. This cutoff was chosen to avoid floor effects associated with locations where few spikes occurred. We computed the correlation between residuals of cell pairs for each trajectory or for the entire session (all trajectories) when animals were running at greater than 3 cm/sec and no sharp wave ripples were detected on any tetrode in CA3. Correlations were only computed if there were at least 10 seconds of data to correlate, e.g. ≥ 20 bins in which both cells expected firing rate was greater than 0.5 Hz. Residual correlations were compared to the overlap between cells in each trajectory. We also determined if there was a relationship between the residual correlation for the entire session and the number of trajectories in which 2 cells overlapped. To compute the number of trajectories in which the cells overlapped, we counted that number of trajectories with overlap greater than 0.3. In all other trajectories that were not deemed overlapping (e.g. did not have overlap > 0.3), overlap had to be low: less than 0.1. Cells that had trajectories with overlaps in between these criteria (0.1-0.2) were excluded for this analysis to allow us to classify cell pairs as

having either high or low overlap on each trajectory. Finally, to control for possible contamination due to clustering errors, we repeated all of these residual correlation analyses for cells recorded on different tetrodes. We found similar results with these more restricted data.

The autocorrelation analysis provides a measure for assigned cells to path equivalent (periodic) or non-path equivalent (non-periodic) classes, but it does not provide an obvious way to measure changes in strength of path equivalent coding over time. We therefore used the normalized overlap measure to examine path equivalence during the initial training on sequence 1 (S1) and the first days of exposure to sequence 2 (S2) in the six arm maze. For two animals we recorded neural activity when the animals initially learned S1 from their first exposure to the track until they reached criterion before the sequence switching phase of the task. During this initial learning period animals did not know the correct sequence and performed poorly. We therefore included all task relevant arms and both incorrect and correct trials. We also examine the distributions of normalized overlap for the first days of exposure to T2 in the three arm maze. In both cases we compared trajectories in the same turn direction.

We determined the extent to which the hippocampal representation of space preserved information about the uniqueness of each trajectory using population vector correlations (Leutgeb et al., 2005b;Derdikman et al., 2009). Each trajectory was divided into 10 cm spatial bins and a population vector of the linearized firing rate for all cells was created. For each animal, recording days were grouped into pairs of consecutive days. Cells from each pair of days in that group were concatenated in the same vector to include more data to compute the correlation. Population vectors from corresponding spatial bins in different trajectories were correlated. Only trajectories in the same turn direction were compared (e.g. trajectory B-C was compared to C-D and C-B was compared to D-C for S1). The resulting correlations from all bins in a pair of trajectories

were averaged. We compared these correlations to cases when the cell identity in the population vector was shuffled or the same trajectory was compared across sessions. Only correct trials were included, though results were similar if all trials were included

Statistics

The measures discussed above tend to produce non-gaussian distributions, so rank sum tests were used for pairwise statistical tests unless otherwise noted. Similarly, the non-parametric Spearman's correlation was used for all correlation analyses unless otherwise noted. Standard analyses of variance was used when the distribution of the data was not different than a Gaussian (Lillie test, $p > 0.05$); otherwise a Kruskal-Wallis one way non parametric anova was used.

Figures

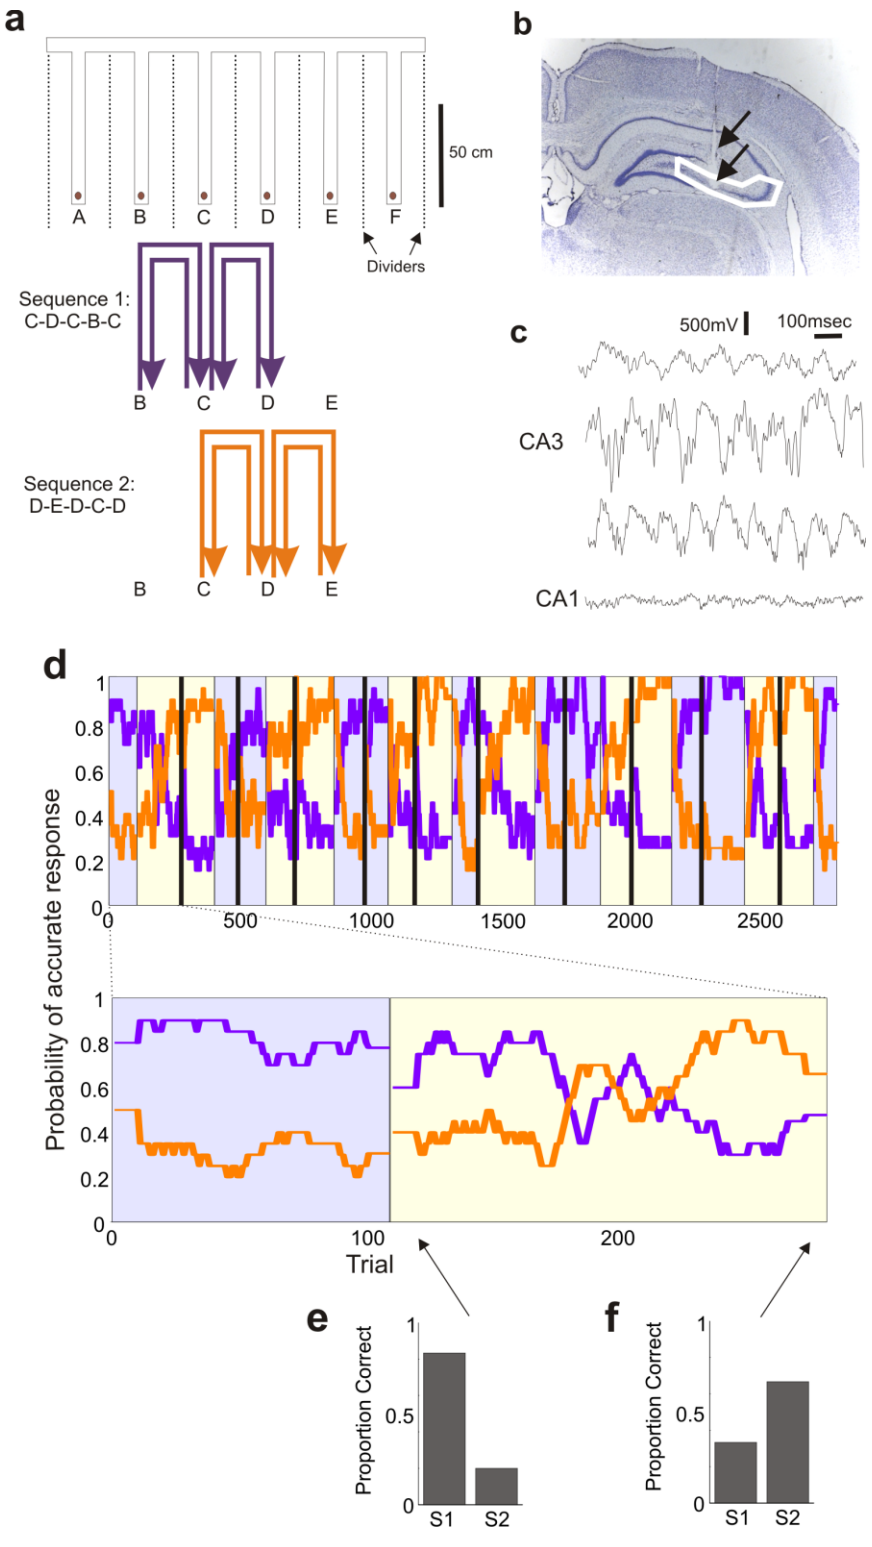
a
A
B
C
D
E
F
50 cm
Dividers
Sequence 1:
C-D-C-B-C
B
C
D
E
Sequence 2:
D-E-D-C-D
B
C
D
E
b
c
500mV
100msec
CA3
CA1
d
Probability of accurate response
0
0.2
0.4
0.6
0.8
1
0
500
1000
1500
2000
2500
0
100
200
Trial
e
Proportion Correct
0
0.5
1
S1
S2
f
Proportion Correct
0
0.5
1
S1
S2

Figure 1. Task design and behavioral performance. **a.** Overhead view of the behavioral apparatus with rewarded sequences indicated by colored arrows (Sequence 1 or S1 in purple and Sequence 2 or S2 in orange). Brown circles indicate the location of food wells; reward was delivered in arms B-E. Dashed lines indicate walls. **b.** Histology. Arrow indicates lesion at the end of a tetrode and the tetrode tracks. White box circles estimated recording area. **c.** All animals showed characteristic EEG signatures (top 3 traces) with large amplitude theta modulation. For comparison, we included a recording from CA1 from the 3 arm maze (bottom trace) which is much smaller in amplitude with less prominent theta modulation. Each trace is a total of one second plotted on the same vertical scale. **d.** Twenty trial moving average of correct responses for one animal when switching between performing S1 (purple) and S2 (orange). Background color indicates which sequence was rewarded. Dark black lines separate recording days. Top graph shows all sequence switching days, bottom graph shows the first day of switching. Chance performance on this task was assumed to be 0.2 because there are 5 arms the animal can choose from when leaving one arm.

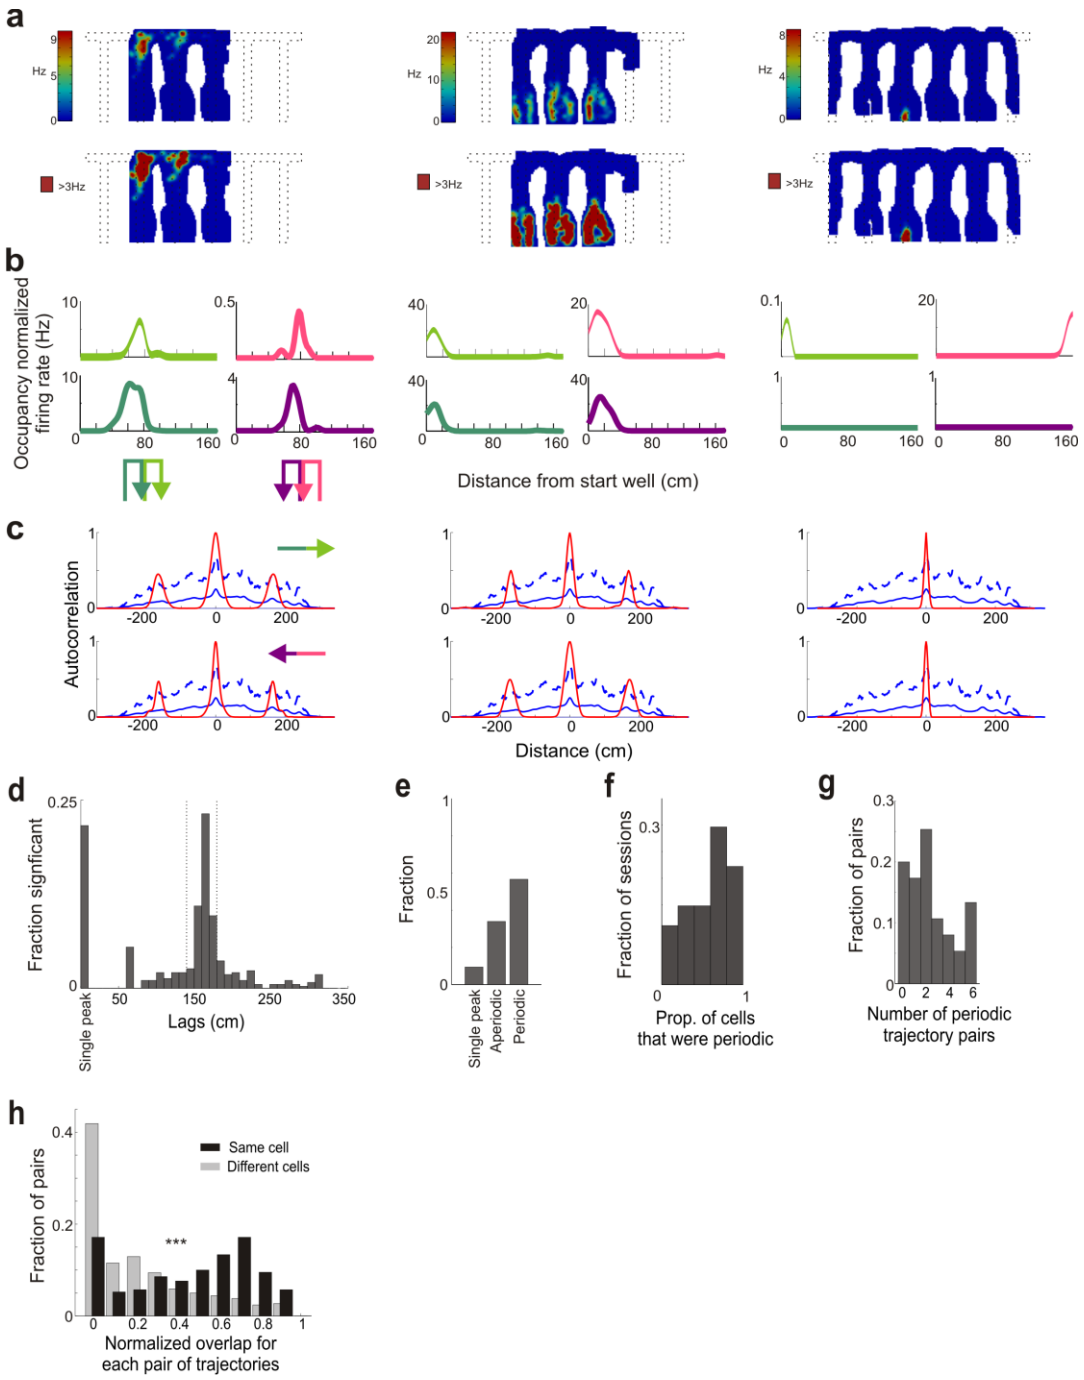

Figure 2. Path equivalent coding in the six arm maze. a. Spatial firing rate maps of three cells scaled to the cells' peak firing rate (above) or scaled so that all firing greater than 3Hz is red (below). The cells shown on the left and in the center have path equivalent firing patterns while the cell shown on the right does not. Only times when

the animal is running more than 3 cm/s are included. **b.** Occupancy normalized firing rates on each linearized trajectory for the same three cells shown in (**a**). All trials when the animal traveled on each trajectory are included together to create the linearized firing rates. **c.** Autocorrelation of the linearized firing for the same three single cells shown in **a** and **b** (red) and the mean cross correlation of the linearized firing between different cells and their confidence bounds (blue. mean, solid line. 90% confidence bounds, dashed line). Occupancy normalized firing rates on trajectories in the same turn direction were concatenated. Top shows right turns as the animal moves from the left to the right and bottom shows left turns as the animal moves from the right to the left. **d.** Histogram of significant peaks at lags greater than 60 cm for each cell's linearized firing autocorrelation. The region between the dotted lines corresponds to lags that are within 20 cm of the length of a single trajectory and were therefore classified as periodic. Autocorrelations and significant peaks were computed separately for trajectories in each turn direction. **e.** Proportion of cells that had periodic or aperiodic activity or had no significant peak at lags greater than 60 cm (single peak). Periodic cells had significant lags at 150-170 cm in either autocorrelogram. Single peak cells had no significant lags greater than 60 cm in either autocorrelogram. Aperiodic cells had significant lags greater than 60 cm and outside of 150-170 cm. **f.** Proportion of cells that had periodic activity in each session, including all sessions per day. **g.** Number of periodic trajectory pairs per cell for periodic cells. **h.** Distribution of normalized overlaps between the linearized firing of the same cell (black) on two different trajectories or different cells (light grey) on two different trajectories. The median for single cells was larger than for different cells ($p < 10^{-10}$, all tests of significance are rank sum tests unless otherwise noted). Only correct trials were included and only trajectories with the same turn direction were compared. Results were similar if all trials were included.

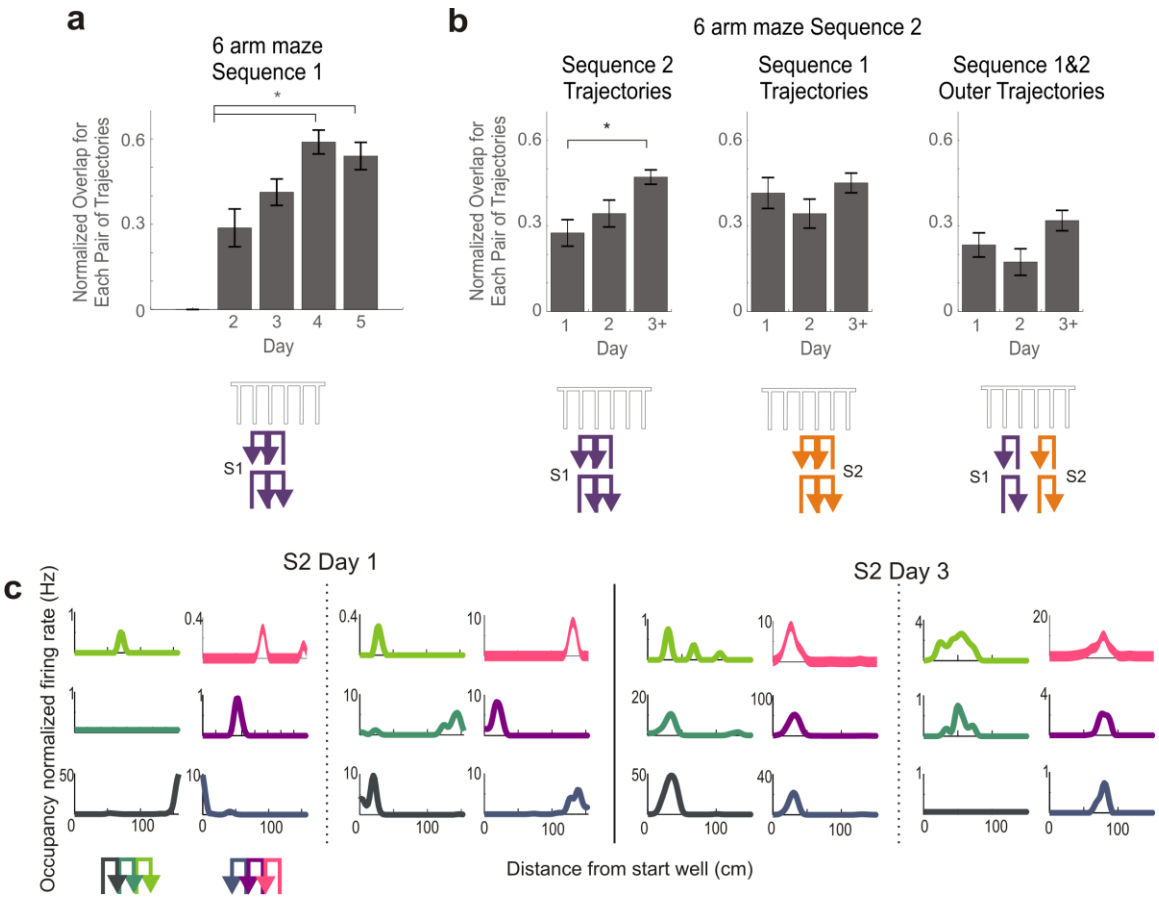

Figure 3. Path equivalence increases with experience. **a.** Normalized overlaps per trajectory pair during initial training on Sequence 1 (S1) in the 6 arm maze across days 2-5. The median normalized overlap increased for S1 in the six arm maze from day 2 to day 5 (Kruskal-Wallis one way analysis of variance $Chi^2(4,162) = 27.6921$, $p < 0.0001$, day 2 was significantly less than day 4 and 5, $p < 0.05$ Tukey-Kramer post-hoc test). There was insufficient data to compute the normalized overlap on day 1. **b.** Normalized overlaps during the first three days of training on Sequence 2 (S2) after initial training on S1 during the switching phase of the task in the 6 arm maze. The median normalized overlap increased for S2 in the six arm maze from day 1 to day 3 (Kruskal-Wallis one way analysis of variance $Chi^2(2,210) = 7.3580$, $p < 0.05$, day 2 was significantly smaller than day 3, $p < 0.05$ Tukey-Kramer post-hoc test) even though the animal had had many

days of experience on the track. Three days are shown because one animal had only three days of exposure to S2.

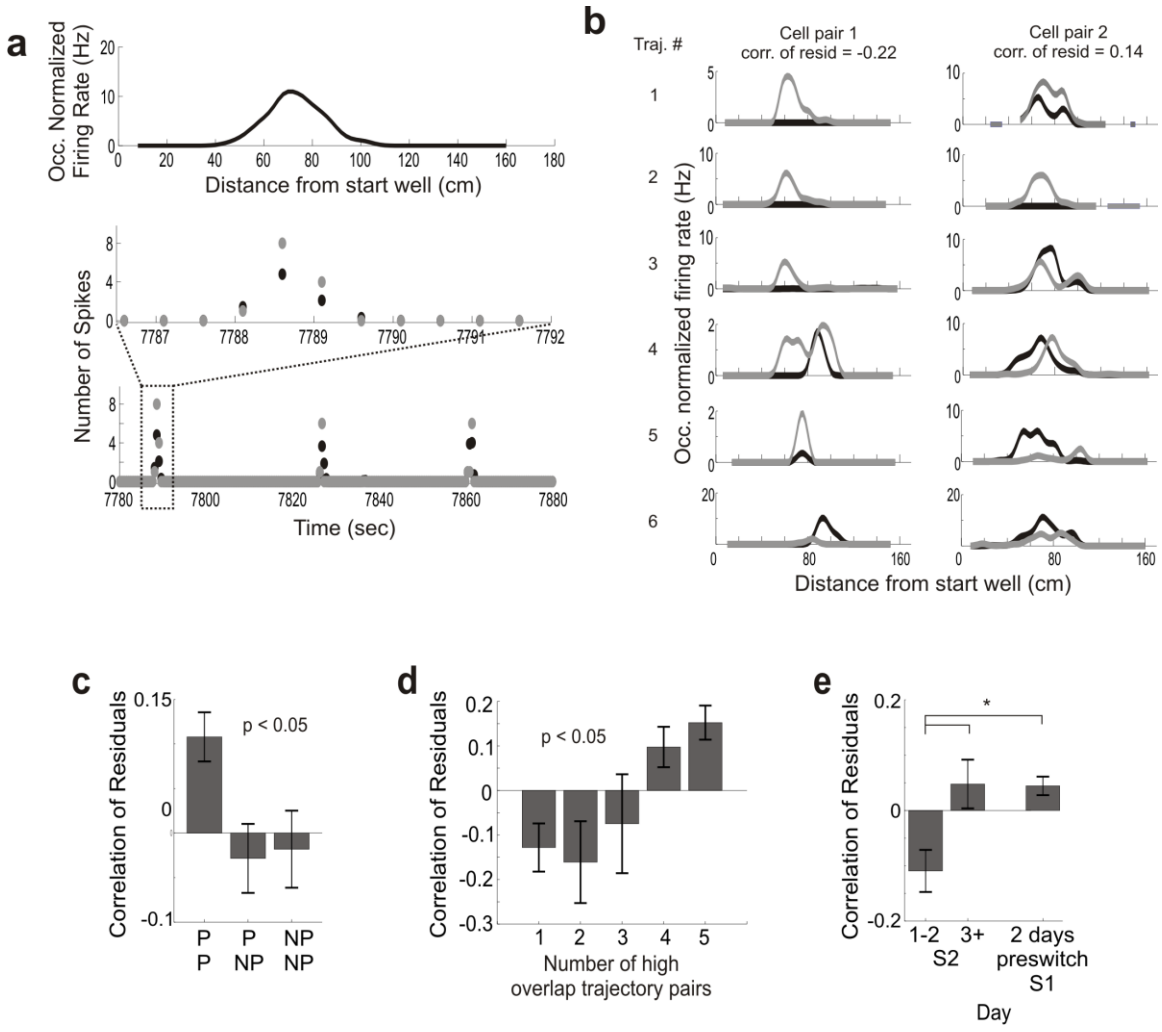

Figure 4. Ensembles of cells with overlapping place fields are correlated. a. The linearized firing rate on a trajectory (top) was used to compute the expected number of spikes (blue. middle and bottom graphs) in each time bin. The expected (black) and actual (light grey) number of spikes in each time bin for a single pass through the trajectory (middle graph) and through three passes (bottom graph). The residuals are the difference between the expected and actual number of spikes. **b.** Examples of linearized firing curves on all task trajectories for different pairs of cells. The cell pair on the left has a negative residual correlation of -0.21921 and has low overlap in most trajectories. The cell pair on the right has positive residual correlation of 0.13771 and

has high overlap in most trajectories. **c.** The cell pairs' residual correlation in one trajectory was related to the normalized overlap of linearized firing in a different trajectory (rho = 0.3229, $p < 0.001$, all Spearman's correlation unless otherwise noted). The line shows a robust fit (slope = 0.24382, $p < 0.02$). **d.** Histogram of residual correlation for the entire session for all cell pairs. **e.** Histogram of residuals' correlation coefficient for the entire session for cell pairs with low overlap (less than 0.2) in at least 3 trajectories. **f.** Histogram of residuals' correlation coefficient for the entire session for cell pairs with high overlap (greater than 0.3) in at least 3 trajectories. The median of the correlation of correlation for high overlap pairs is significantly higher than the medians of the other groups (from all pairs, $p < 0.05$, from pairs with low overlap, $p < 0.0005$). **g.** The number of trajectories with high overlap (overlap greater than 0.3) versus the correlation coefficient of the residuals for all cell pairs (ANOVA $p < 0.05$, $F = 2.9704$). Bars show mean and standard error. When comparing different trajectories, only trajectories in the same turn direction were compared. **h.** Correlation of residuals during the first two days of training on S2 (left bar), during three or more days of training on S2 (middle bar) after initial training on S1 during the switching phase of the task in the 6 arm maze and during the last two days of training on S1 before the switching phase commenced (right bar). Correlations increased from days 1-2 to days 3 or more (ANOVA $F(2,115)$: 5.9582, $p < 0.005$, S2 days 1-2 was significantly smaller than S3 days 3+ and S1 2 days preswitch, $p < 0.05$ Tukey-Kramer post-hoc test). Only cell pairs with an overlap of at least 0.1 across all trajectories were included. Bars show mean and standard error.

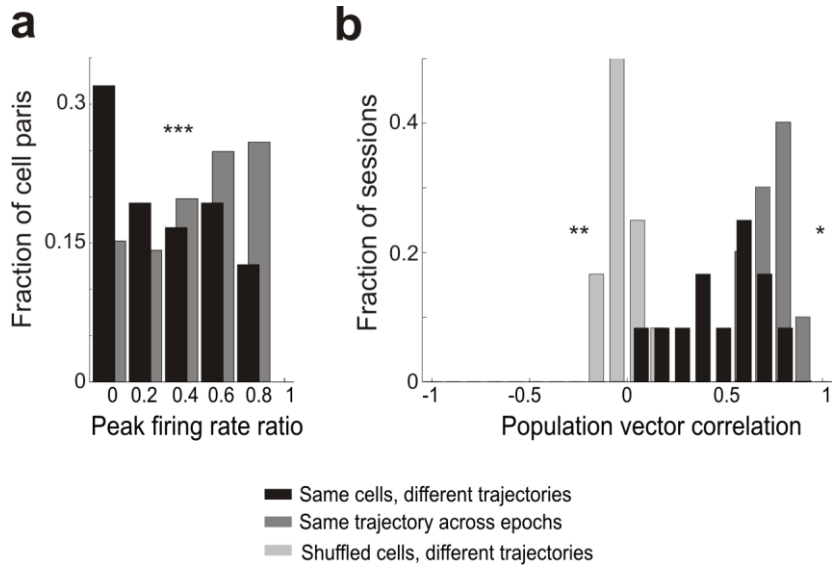

Figure 5. Firing rate differences create distinct population code. a. Histogram of peak rate ratio of cell pairs' linearized firing different trajectories with similar firing locations in the same run epoch (black; defined as normalized overlap greater than 0.3) and the same trajectory in different run epochs within a day (medium grey). The medians of the two distributions were significantly different ($p < 10^{-5}$). **b.** Histogram of population vector correlation of different trajectories in the same epoch (black), the same trajectory in different epochs (medium grey) and different trajectories in the same epoch with shuffled cell identity (light grey). The median population vector correlation across different trajectories in the same epoch was larger than that for the shuffled cell identities ($p < 0.0001$) but smaller than that for the same trajectory across epochs ($p < 0.02$). Only correct trials were included and when comparing different trajectories only trajectories in the same turn direction were compared.

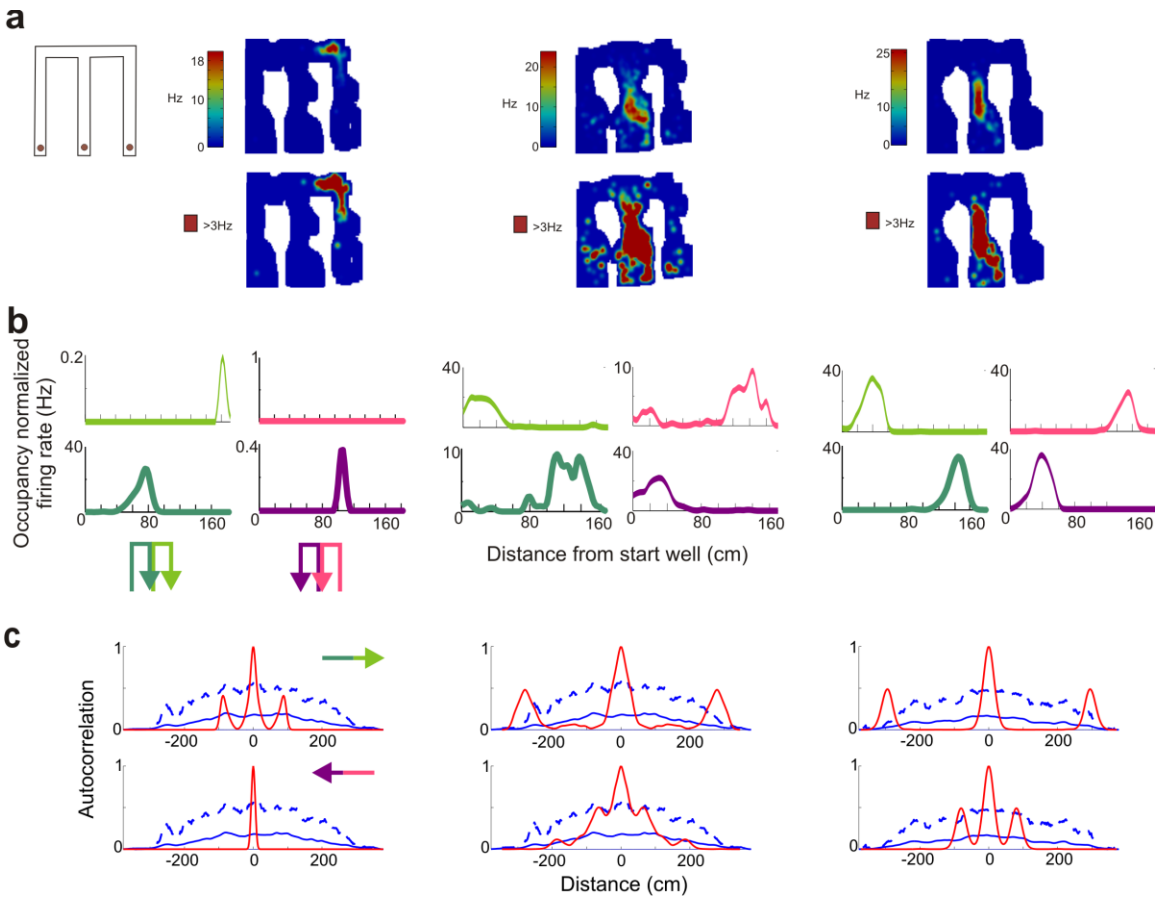

Figure 6. Path equivalent coding is less common in the three arm maze. a. Spatial firing rate maps of three cells scaled to the cells' peak firing rate (above) or scaled so that all firing greater than 3Hz is red (below). The cells shown on the left and center were recorded in CA3 and the cell shown on the right was recorded in CA1. The cell shown in the center fires in several arms in some similar locations, while the others have place fields limited to a single arm. Only times when the animal is running more than 3 cm/s are included. **b.** Occupancy normalized firing rates on each linearized trajectory for the same three cells shown in **a**. **c.** Autocorrelation of the linearized firing for the same three single cells shown in **a** and **b** (red) and the mean cross correlation of the linearized firing between different cells (blue. mean, solid line. 90% confidence bounds, dashed line). Occupancy normalized firing rates on trajectories in the same turn

direction were concatenated. Top shows right turns as the animal moves from the left to the right and bottom shows left turns as the animal moves from the right to the left.

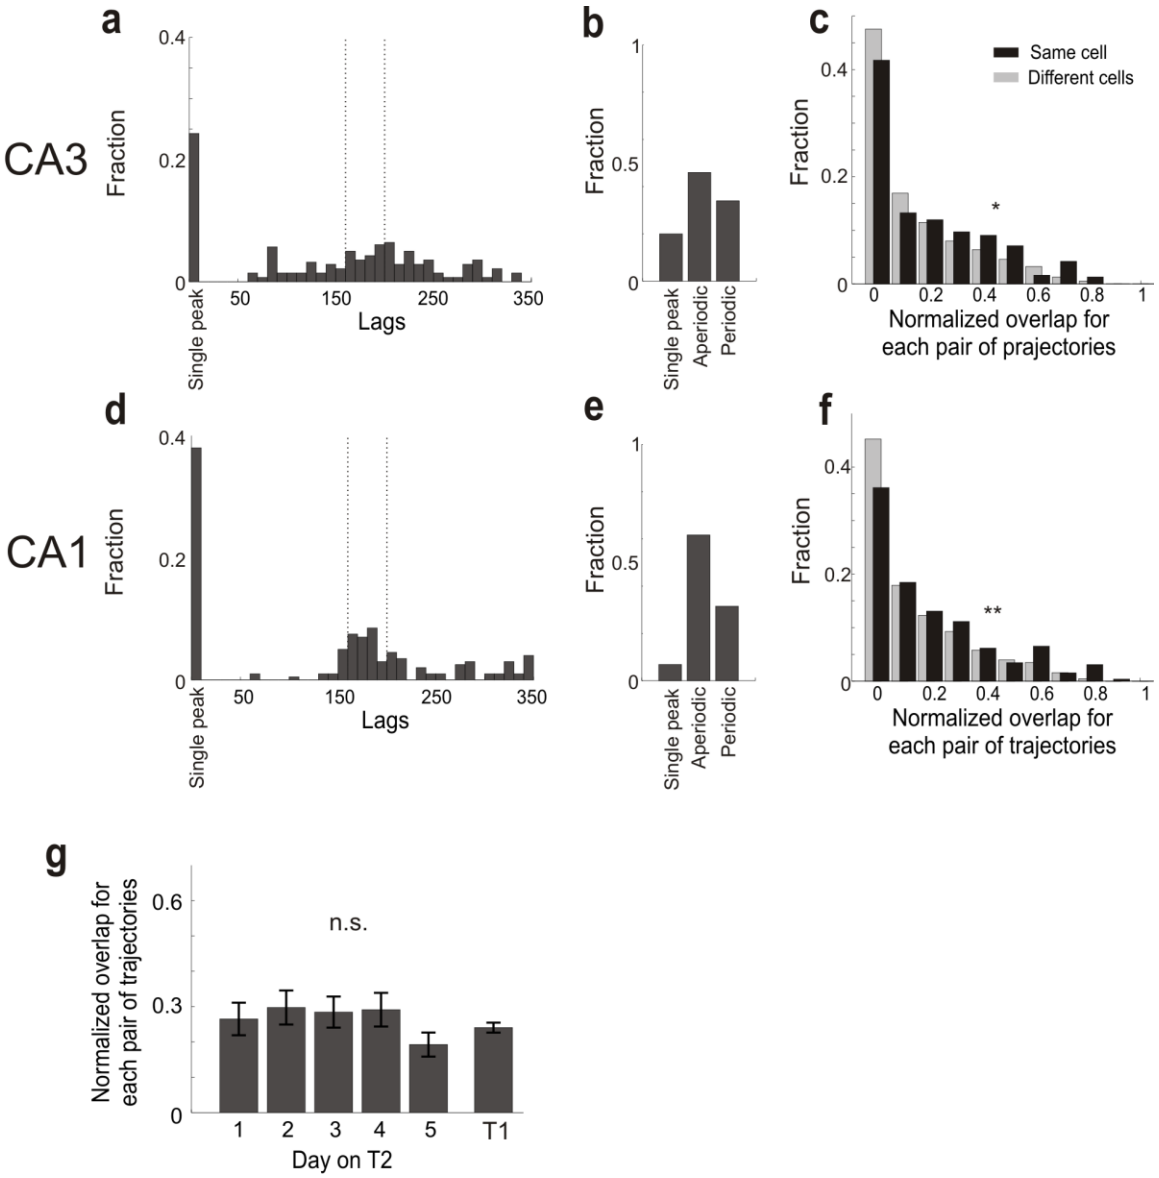

Figure 7. Quantification of path equivalent activity in the three arm maze. a. Histogram of significant peaks at lags greater than 60 cm for each cell's linearized firing autocorrelation (right) in CA3. **b.** Proportion of cells with periodic or aperiodic activity, or no significant peaks at lags greater than 60 cm in CA3 (single peak). **c.** Distribution of normalized overlaps between the linearized firing of the same cell (black) on two different trajectories or different cells (light grey) on two different trajectories in CA3. The median values were significantly different ($p < 0.05$) but were similar in magnitude

(medians = 0.1697 same cell, 0.1211 different cell). **d-f.** The same as **a, b,** and **c.** for CA1. Once again, the median normalized overlap values were significantly different ($p < 0.005$) but were similar is magnitude (medians = 0.1557 same cell, 0.1129 different cell). Only correct trials were included and only trajectories with same turn direction were compared. Results were similar if all trials were included. **g**. Normalized overlap in the initially novel 3 arm maze (Track 2) and from Track 1. There were no significant changes in the distributions across days (Kruskal-Wallis one way analysis of variance $Chi^2(4,108) = 7.8749$, $p > 0.09$, the days were not significantly different $p > 0.05$ Tukey-Kramer post-hoc test).

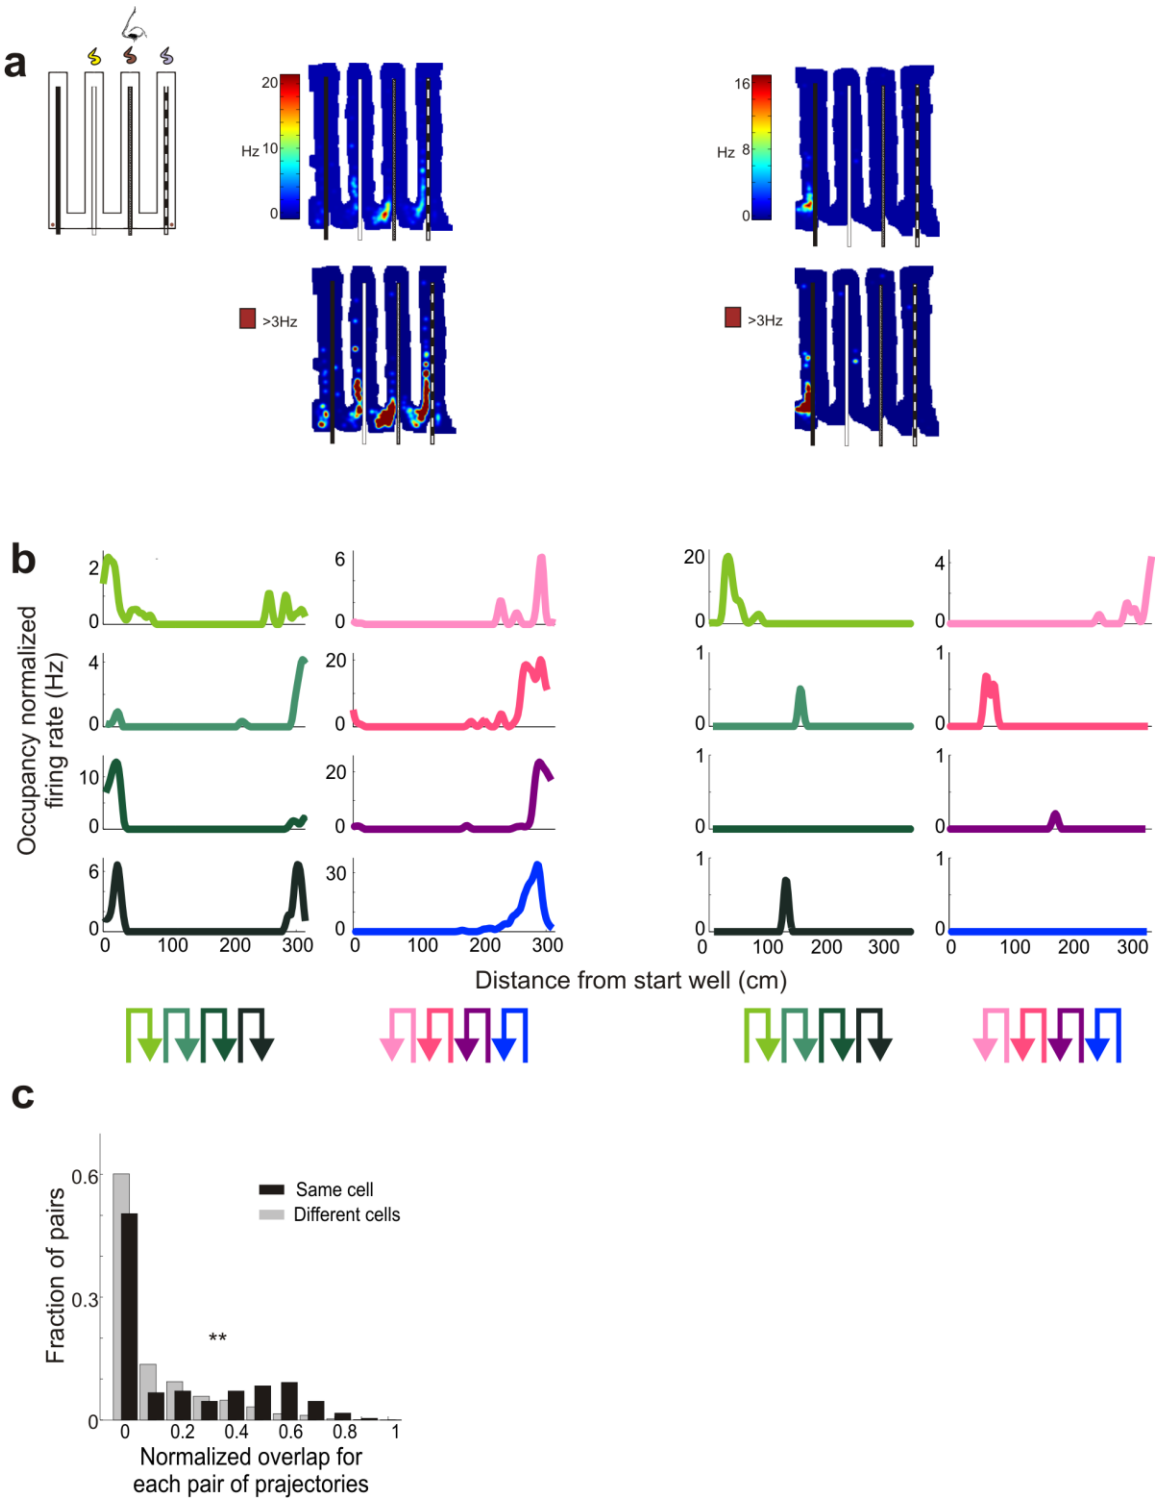

Figure 8. Path equivalent activity in CA1 neurons active on the multiple U track.

a. Spatial firing rate maps of two CA1 cells scaled to the cells' peak firing rate (above) or

scaled so that all firing greater than 3Hz is red (below). The cell on the left shows repetitive firing patterns while the cell on the right does not. Only times when the animal is running more than 3 cm/s are included. **b.** Occupancy normalized firing rates on each linearized trajectory for the same cells shown in **a**). **c.** Normalized overlap between the occupancy normalized firing rate of single cells (black) or different cells (light grey) on two different trajectories ($p < 0.001$). Only trajectories with the same turn direction were compared.

Supplementary Figures

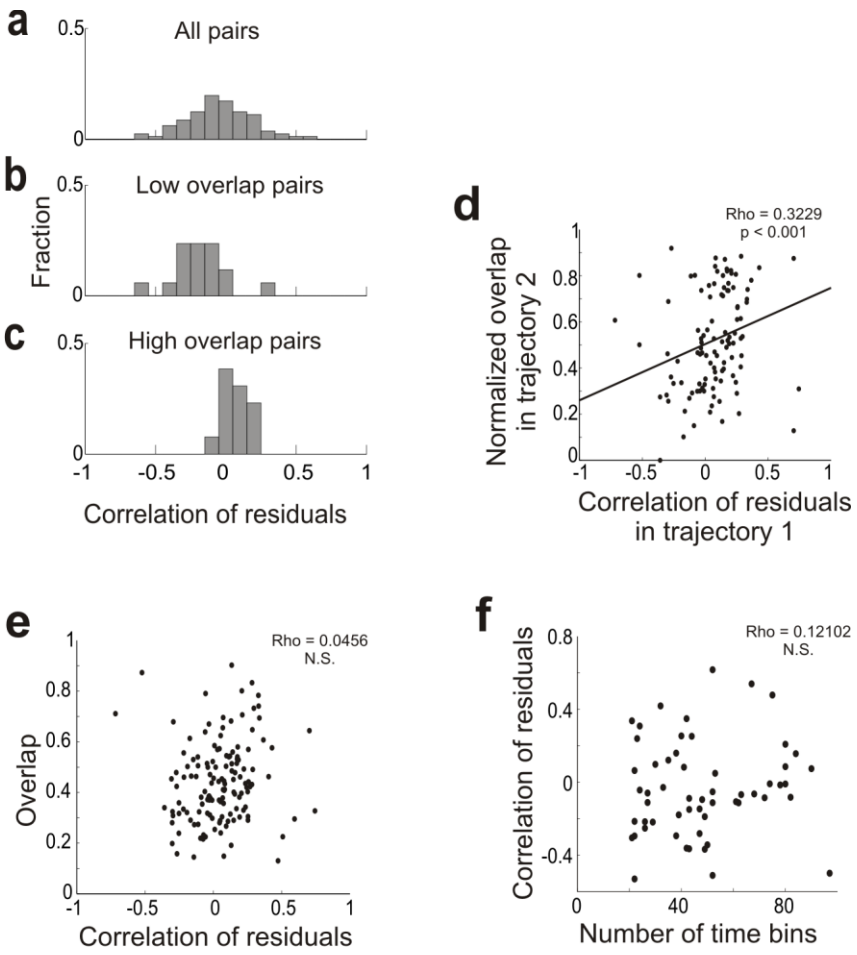

Supplementary Figure 1. Correlation of residuals not explained by overlap in the same trajectory time bins available for analysis. **a.** The cell pairs' correlation of residuals was not significantly related to overlap of linearized firing in the same trajectory (rho = 0.0456, p > 0.7). **b.** The number of time bins included to compute the residual correlation was not significantly related to the residual correlation for cases with less than 100 time bins (rho = 0.12102, p > 0.3s).

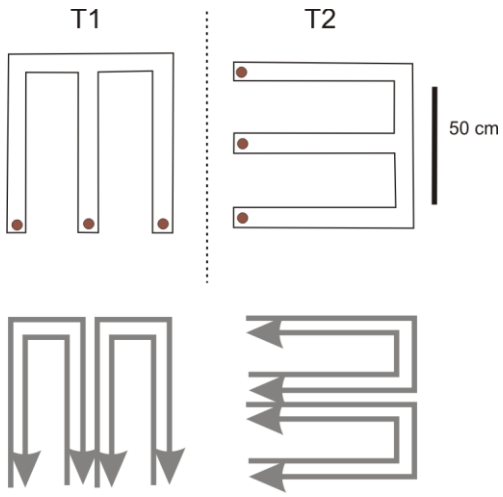

Supplementary Figure 2. Three arm maze. The overhead view of the two three arm W-tracks. Brown circles denote the location of food wells. Grey arrows indicate trajectories in the rewarded sequence. The animals were rewarded for performing the same continuous alternation task as in the 6 arm maze, which involved the following sequences of arms: center, left, center, right, center, and so on.

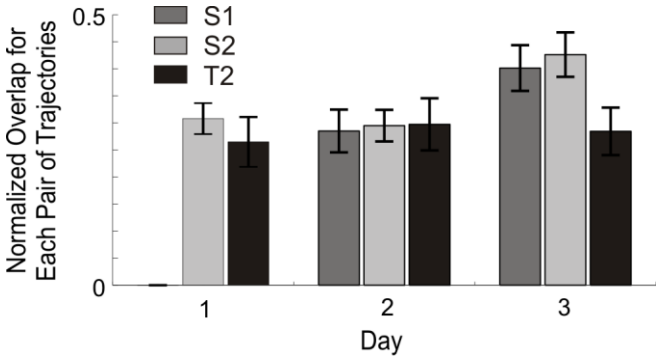

Supplementary Figure 3. Similar initial levels of overlap across trajectory pairs in the six and three arm mazes. Overlap per trajectory pair for days 1 to 3 in six arm maze sequence 1 (S1, dark grey, from Fig 8a), sequence 2 (S2, light grey, from Fig. 8b) in the 6 arm maze and T2 in the 3 arm maze (black, from Fig. 12a) for direct comparison. The similarities in the overlaps suggests that CA3 cells start out with more dissimilar firing across trajectories, but that as the animal learns the task in the six arm maze, the coding for these trajectories becomes more similar. Bars show mean and standard error.

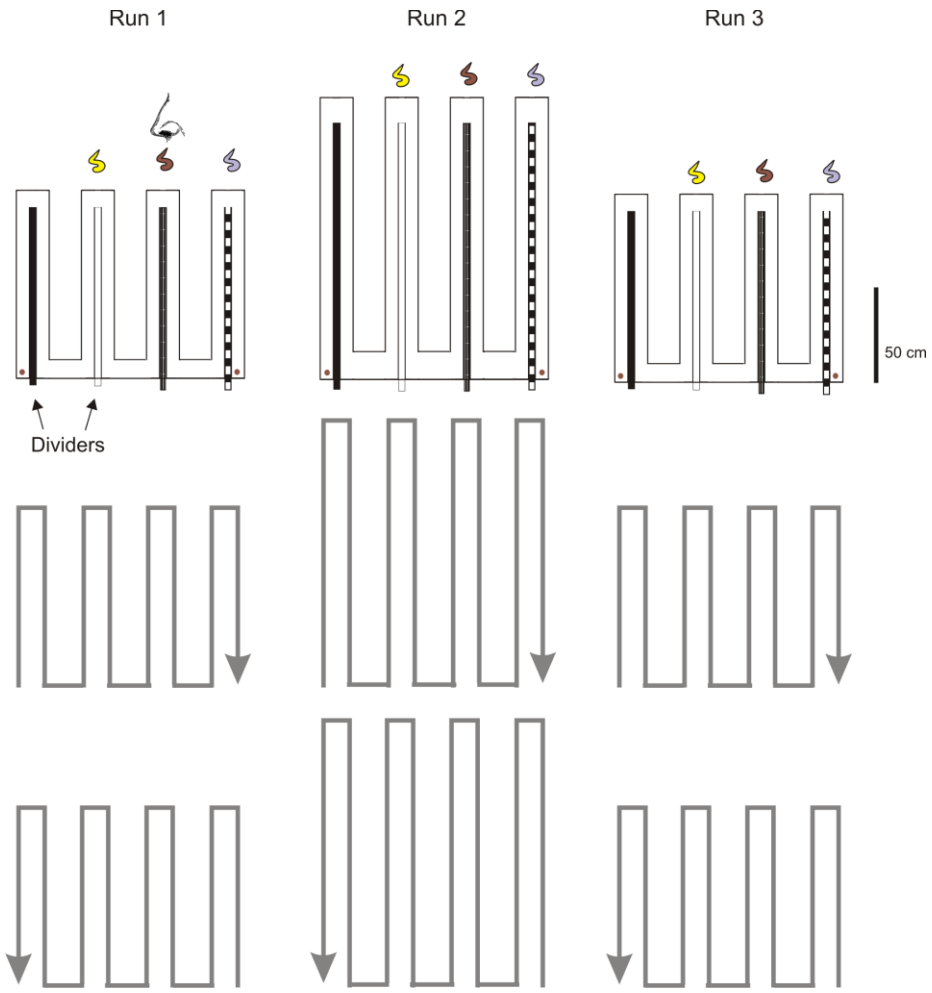

Supplementary Figure 4. Multiple U-track. The overhead view shows the medium and long configurations of the multiple-U track. Animals ran first on the medium configuration, then on the long configuration and then on the medium configuration with 20-30 minute rest sessions in between. The arms of each U were separated by a unique divider and each U was scented uniquely to ensure that the animals could distinguish between them. Brown circles denote the location of food wells. Grey arrows indicate trajectories in the rewarded sequence.

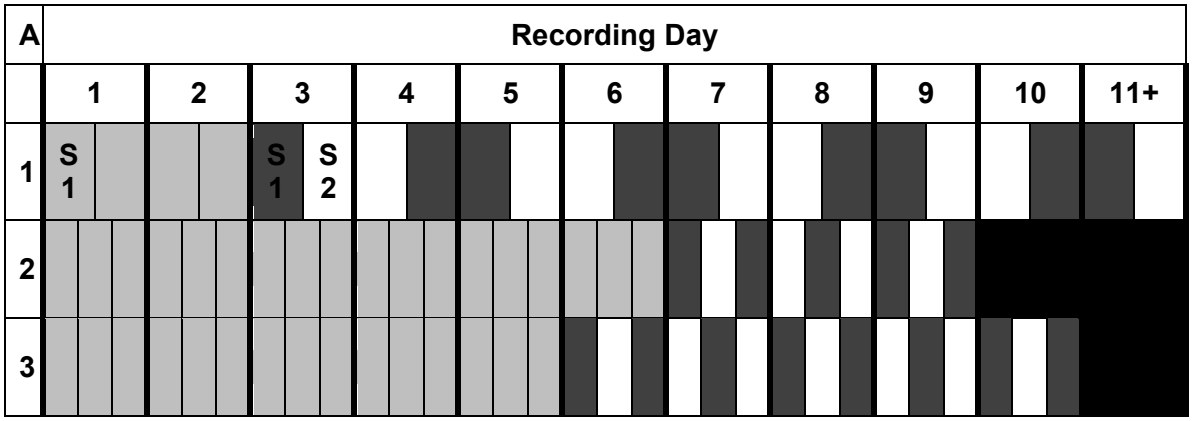

Supplementary Table 1. Recording schedule for each animal (A) organized by recording day. Each square is one recording session when the animal was rewarded for performing either Sequence 1 (S1, grey) or Sequence 2 (S2, white). Recording sessions were flanked by rest sessions (not shown). Light grey squares indicate initial training on S1. Dark grey squares indicate S1 and white squares indicate S2 during the sequence switching phase of the task when the animal was exposed to S2 and sessions alternated between rewarding S1 and S2. Black indicates no recording. Note: Animal 1 was pretrained on S1 before recording began and recording for animal 1 continued alternating between S1 and S2 up to recording day 18. Animal 2 and 3 had never been exposed to S1 or the behavioral track before recording began. All animals were pretrained on a linear track in a different room from the recording room.

Conclusion

As a crossroads between sensory inputs and long term memories, the hippocampus must turn a plethora of information into concise stories. The hippocampus can use several different strategies to achieve this transformation. By selecting only more “important” episodes to transfer to long term memory storage, we can remember important experiences while forgetting the mundane. By encoding general or common principles among several experiences, we can remember appropriate general responses or outcomes and predict responses and outcomes of future similar experiences. We considered ways the hippocampus might achieve both of these possibilities.

First, we examined hippocampal activity in response to reward. We must remember the experiences that lead to reward to exploit these rewards in the future. Because SWR reactivation is thought to be important for memory consolidation, we asked if reactivation is enhanced by reward. We found that SWR activity increases when animals receive reward. And this reward related SWR activity is further enhanced when animals have to learn new path-reward associations. Spiking during SWRs when animals are rewarded reactivates neural patterns that occur as animals run to or from the reward. Because SWRs are implicated in memory consolidation, we think this enhanced SWR reactivation could be a mechanism to preferentially remember experiences associated with reward.

Next, we examined hippocampal activity when animals navigate through an environment with many repeated elements. In environments with repeated elements, generalizing across elements can be advantageous to encode generally appropriate

responses. But each element must also be differentiated from the others, in particular in our task animals had to distinguish between similar trajectories to receive reward. We found that some hippocampal cells fire very similarly on multiple repeated elements, while other cells encode the elements differently. Cells that generalize across similar elements have correlated moment to moment activity, suggesting that they are part of functional ensemble. Furthermore, this generalizing / path equivalent activity increases as animals learn new relationships between repeated elements. We think this generalization across repeating elements could be a mechanism to extract general principles about related experiences, perhaps for more efficient memory encoding and later predictions.

Caveats

In interpreting these results, there are a number of caveats we should keep in mind. First, the evidence that SWRs are important for memory consolidation is strong but mostly circumstantial. SWRs have long been theorized to contribute to memory consolidation. During SWRs the same patterns of activity seen during behavior are reactivated on the short times scales required for plasticity (Buzsaki, 1986;Wilson and McNaughton, 1994;Sutherland and McNaughton, 2000). Furthermore, this activity is coherent across multiple brain regions (Sutherland and McNaughton, 2000;Pennartz et al., 2004;Ji and Wilson, 2007;Euston et al., 2007;Peyrache et al., 2009b), potentially allowing information to be transferred and coordinated between regions for memory consolidation. Recently, new evidence has suggested a causal link between SWRs and learning and memory. Girardeau et al. showed that disrupting SWRs during sleep after animals learn a memory task hinders subsequent task performance (Girardeau et al.,

2009). While this study presents a significant breakthrough, it is still unclear how specific the effects of SWR interruption were: was performance impaired on other tasks that did not involve memory or had been learned before SWR interruption? Did SWR interruption damage the hippocampus? Furthermore, no one has shown that SWR activity during waking contributes to learning or memory. Instead, our results agree with and provide further corroborating evidence for the theory that SWR activity during pauses in behavior is important for memory consolidation.

Second, correlations in moment to moment activity are thought to reveal common inputs or connections between neurons, however the evidence is indirect. Spiking activity is noisy and correlations in moment to moment activity are often called "noise correlations" as they are thought to represent covariance in activity not due to "signal" or the cells' receptive fields. Noise in spiking is thought to be transferred from one neuron to another, but the network is able to detect signal by averaging across many neurons with similar signal but different noise (Maynard et al., 1999;Carmena et al., 2005). Therefore by detecting covariation in activity not due cells' receptive fields, we could detect connections between neurons or common inputs as this noise travels from one neuron to the next but does not propagate across the entire population (Lee et al., 1998;Shadlen and Newsome, 1998). While this theory is likely to be correct, there has been no direct demonstration of it so far. Previous studies have shown that neurons in the same cortical column have correlated moment to moment activity (Lee et al., 1998). While neurons in the same cortical column tend to have similar inputs, previous studies have not shown the neurons with common input have correlated moment to moment activity while neurons without common inputs do not. Even with this caveat in mind, our results are still interesting. We have shown that cells that have similar receptive fields across multiple similar elements have more correlated moment to moment activity.

Whether these correlations are due to common inputs, common connections, or common sensitivity to some sensory stimulus, these results still suggest that these neurons have some common properties different from other neurons in the same region.

Third, an underlying assumption of many studies, including ours, is that averaging across many trials is equivalent to recording from many more neurons in only a single trial. Averaging both across many trials and across many neurons reduces noise. Since we cannot record from a significant portion of the cells in a region (even with large scale recordings), we compensate by averaging across trials. This issue is relevant to both our findings that (1) there is a relationship between SWR coactivity and place field overlap and (2) that even with path equivalent coding the population can distinguish between different elements of the environment. A superior analysis in both cases would be to examine these relationships trial by trial. For the SWR activity, we would ask if relationships between cells during SWRs reflect relationships between cells on the previous pass only, not the place fields' overlap which includes activity from the entire session. For the population activity, we would ask if the neural activity on each pass distinguishes between similar trajectories. We could perhaps answer these questions using a decoding analysis (Davidson et al., 2009;Karlsson and Frank, 2009a), however we do not have enough simultaneously recorded cells. Indeed, these particular analyses have not been successfully performed before. So instead, we have used place field activity averaged over the entire session and we must be careful in interpreting how these findings apply to a single SWR or a single pass.

Further Questions

Understanding how the hippocampus filters vast amounts of sensory information into succinct memories will require more study. And our findings have opened up a number of questions for further investigation.

We have shown that SWR reactivation is enhanced when animal are rewarded and further enhanced when animal must learn new path-reward contingencies. Previous work has shown that SWRs are also enhanced by novelty, when learning and memory consolidation may also be advantageous (Cheng and Frank, 2008). We then speculate that increased SWR reactivation could be a mechanism to preferentially remember some memories over others. This leads to the interesting possibility that SWR activity is enhanced by punishment or negative outcomes, as we can also learn what to avoid from negative experiences.

Further, we would like to understand the mechanisms that modulate increases and decreases in SWR reactivation. How SWRs occur is not known, but they are generally thought to occur when the inhibitory tone in the network decreases, permitting a burst of excitatory cell activity (Buzsaki, 1986). The burst of activity then shuts down because the excitatory cells activate inhibitory interneurons, which tamp down the network. Because SWR activity is enhanced by novelty and reward, we hypothesize that neuromodulators that have been implicated in reward and novelty signaling, like dopamine or norephinephrine, may modulate SWR activity (Harley, 2004;Schultz, 2007;Sara, 2009). Indeed the VTA and the locus coeruleus project to CA3, where SWRs are thought to originate (Andersen et al., 2006). Dopamine or norepinephrine, when released into the hippocampus, may change the network state to be more permissive to SWRs and bursts of population activity, perhaps by changing inhibitory aor

excitatory tone. These hypothesized mechanisms by which SWR activity is produce and modulated have yet to be proven.

We also wonder how SWR activity in the hippocampus influences and interacts with other brain regions. Activity during SWRs has been shown to be coherent across brain regions, including areas that are involved in decision making and reward processing (Pennartz et al., 2004;Lansink et al., 2008;Lansink et al., 2009;Peyrache et al., 2009b). As we discussed in Chapter 2, enhanced SWR reactivation during reward could enhance representations of rewarded paths over unrewarded or enhance associations between paths and outcome for later decision making. Exactly which of these possibilities is unclear as we do not understand how SWR reactivation is used in other brain regions.

We have shown that, when faced with experiences with repeating elements, some ensembles of hippocampal neurons encode these elements similarly while other cells differentiate between elements. Neurons with path equivalent activity on multiple elements had more correlated moment to moment activity, suggesting that they were part of functional ensembles. These results raise the interesting question: from where does this correlated moment to moment activity arise? These cells may have similar inputs from the entorhinal cortex or they may be directly connected to each other via CA3's recurrent collaterals. Alternately, they could be similarly receiving or sensitive to neuromodulators.

We also found that as animals learn relationships between elements, this generalizing, path equivalent activity and correlations in moment to moment activity increase. These findings have made us wonder, how do these ensembles form and create groups of cells that repeat together in multiple similar elements. We suggest two

possibilities. If correlated moment to moment activity between two cells indicates that they are connected or receive similar inputs then these common inputs may drive the cells to develop similar receptive fields across repeated elements. Or perhaps cells that have similar receptive fields in multiple elements fire together repeatedly, more than other cells that overlap in just one element, and therefore these cells are more likely to wire together (or wire more strongly to similar inputs) and develop correlated moment to moment activity.

Furthermore, if activity that generalized across elements is important for encoding response that are appropriate across situations as we suggest, we would expect this activity to be utilized in brain regions involved in behavioral responses and decision making. Indeed, the path equivalent activity that we observed is more similar to activity in regions like the striatum and prefrontal cortex (Jung et al., 1998;Mizumori et al., 2004). The striatum is involved in habit formation and encodes, among other things, behavioral responses like turning left or right (Mizumori et al., 2004;Graybiel, 2005). The prefrontal cortex is essential for behavioral flexibility and complex decision making and it encodes stimuli, responses, outcomes, and associations between these (Miller, 1999). Both of these regions receive projections from the hippocampus and have correlations in theta oscillations or theta phase locking with the hippocampus during task performance (Siapas et al., 2005;Jones and Wilson, 2005a;Jones and Wilson, 2005b;Andersen et al., 2006;DeCoteau et al., 2007). Accordingly, we would find it informative to know if path equivalent activity in the hippocampus is correlated with similar activity in these other regions.

Finally, we ask: what role might these processes, modulation of SWR reactivation and generalization across repetitive experiences, play in disease? If we remember every detail of every experience, we would soon be incapacitated by sensory

overload. Selectively remembering “important” episodes and generalizing across similar episodes, are essential to distilling a potential overload of information into a manageable amount of data to be stored for the long term. The failure of these processes could contribute to a variety of neurological and psychiatric diseases. While many studies have implicated the hippocampus is different psychiatric or neurological diseases, little is known about how hippocampal neural activity goes awry in these diseases (Burt et al., 1995;Sheline et al., 2002;Leppanen, 2006).

Recent work has shown that in a mouse model of Alzheimer’s diseases, mutant/diseased animals have subthreshold seizures (Palop et al., 2007). This breakthrough, if true in humans, suggests a potential explanation for fluctuations in cognitive function of those with Alzheimer’s: they may be having undetected seizures during periods when they are most cognitively impaired. This finding also suggests a new therapeutic approach: anti-epilectic drugs.

Unfortunately, such discoveries have been lacking in many other psychiatric and cognitive diseases because there are few good animal models of these diseases. Therefore, we are left to speculate how these hippocampal neural processes go awry in disease. For instance, post traumatic stress disorder (PTSD) could involve an overgeneralization across experiences such that a situation similar to a negative experience evokes a fear response even when that response is inappropriate. The hippocampus has also been implicated in depression, though its exact role is unknown (Burt et al., 1995;Sheline et al., 2002;Leppanen, 2006). Depression could also involve overgeneralization of negative experiences to neutral experiences. Or depression could be exacerbated by preferential storage of negative experiences in the hippocampus. Indeed, studies in humans have shown that depressed people tend to attend to and remember negative experiences more than controls (Leppanen, 2006). While the role of

the hippocampus in neurological and psychiatric disease is beyond the scope of this thesis, understanding different processes that contribute to hippocampal memory formation and consolidation may guide more pointed questions about what goes awry in these diseases.

In sum, understanding how we preferentially encode the significant experiences of our lives and relate these experiences to each other is essential for understanding how we form the stories that make up ourselves and our lives. To grapple with many unsolved neurological and psychiatric diseases, we will have to understand these different processes that make, relate, and store these stories.

Reference List

Ainge,J.A., Tamosiunaite,M., Woergoetter,F., and Dudchenko,P.A. (2007a). Hippocampal CA1 place cells encode intended destination on a maze with multiple choice points. J. Neurosci. 27, 9769-9779.

Ainge,J.A., van der Meer,M.A., Langston,R.F., and Wood,E.R. (2007b). Exploring the role of context-dependent hippocampal activity in spatial alternation behavior. Hippocampus. 17, 988-1002.

Andersen,P., Morris,R.G., Amaral,D.G., Bliss,T.V., and O'Keefe,J. (2006). The Hippocampus Book.

Anderson,M.I., and Jeffery,K.J. (2003). Heterogeneous modulation of place cell firing by changes in context. J Neurosci 23, 8827-8835.

Battaglia,F.P., Sutherland,G.R., and McNaughton,B.L. (2004). Local sensory cues and place cell directionality: additional evidence of prospective coding in the hippocampus. J. Neurosci. 24, 4541-4550.

Breese,C.R., Hampson,R.E., and Deadwyler,S.A. (1989). Hippocampal place cells: stereotypy and plasticity. J. Neurosci. 9, 1097-1111.

Burt,D.B., Zembar,M.J., and Niederehe,G. (1995). Depression and memory impairment: a meta-analysis of the association, its pattern, and specificity. Psychol. Bull. 117, 285-305.

Buzsaki,G. (1986). Hippocampal sharp waves: their origin and significance. Brain Res. 398, 242-252.

Buzsaki,G., Horvath,Z., Urioste,R., Hetke,J., and Wise,K. (1992). High-frequency network oscillation in the hippocampus. Science 256, 1025-1027.

Buzsaki,G., Leung,L.W., and Vanderwolf,C.H. (1983). Cellular bases of hippocampal EEG in the behaving rat. Brain Res. 287, 139-171.

Cahill,L., and McGaugh,J.L. (1998). Mechanisms of emotional arousal and lasting declarative memory. Trends Neurosci. 21, 294-299.

Carmena,J.M., Lebedev,M.A., Henriquez,C.S., and Nicolelis,M.A. (2005). Stable ensemble performance with single-neuron variability during reaching movements in primates. J. Neurosci. 25, 10712-10716.

Cheng,S., and Frank,L.M. (2008). New experiences enhance coordinated neural activity in the hippocampus. Neuron. 57, 303-313.

Chrobak,J.J., and Buzsaki,G. (1996). High-frequency oscillations in the output networks of the hippocampal- entorhinal axis of the freely behaving rat. J. Neurosci. 16, 3056-3066.

Csicsvari,J., Hirase,H., Mamiya,A., and Buzsaki,G. (2000). Ensemble patterns of hippocampal CA3-CA1 neurons during sharp wave-associated population events. Neuron 28, 585-594.

Csicsvari,J., O'Neill,J., Allen,K., and Senior,T. (2007). Place-selective firing contributes to the reverse-order reactivation of CA1 pyramidal cells during sharp waves in open-field exploration. Eur. J. Neurosci. 26, 704-716.

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

DeCoteau,W.E., Thorn,C., Gibson,D.J., Courtemanche,R., Mitra,P., Kubota,Y., and Graybiel,A.M. (2007). Learning-related coordination of striatal and hippocampal theta rhythms during acquisition of a procedural maze task. Proc Natl Acad Sci U S A 104, 5644-9.

Derdikman,D., Whitlock,J.R., Tsao,A., Fyhn,M., Hafting,T., Moser,M.B., and Moser,E.I. (2009). Fragmentation of grid cell maps in a multicompartment environment. Nat. Neurosci. 12, 1325-1332.

Diba,K., and Buzsaki,G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. Nat. Neurosci. 10, 1241-1242.

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

Eichenbaum,H. (2000). A cortical-hippocampal system for declarative memory. Nat Rev Neurosci 1, 41--50.

Euston,D.R., Tatsuno,M., and McNaughton,B.L. (2007). Fast-forward playback of recent memory sequences in prefrontal cortex during sleep. Science. 318, 1147-1150.

Fenton,A.A., Kao,H.Y., Neymotin,S.A., Olypher,A., Vayntrub,Y., Lytton,W.W., and Ludvig,N. (2008). Unmasking the CA1 ensemble place code by exposures to small and large environments: more place cells and multiple, irregularly arranged, and expanded place fields in the larger space. J. Neurosci. 28, 11250-11262.

Fiorillo,C.D., Newsome,W.T., and Schultz,W. (2008). The temporal precision of reward prediction in dopamine neurons. Nat. Neurosci.

Foster,D.J., and Wilson,M.A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. Nature. 440, 680-683.

Fox,S.E., and Ranck,J.B.J. (1981). Electrophysiological characteristics of hippocampal complex-spike cells and theta cells. Exp. Brain Res. 41, 399-410.

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

Frank,L.M., Brown,E.N., and Wilson,M.A. (2001). A Comparison of the Firing Properties of Putative Excitatory and Inhibitory Neurons From CA1 and the Entorhinal Cortex. J. Neurophys. 86, 2029-2040.

Frank,L.M., Stanley,G.B., and Brown,E.N. (2004). Hippocampal plasticity across multiple days of exposure to novel environments. J. Neurosci. 24, 7681-7689.

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.

Girardeau,G., Benchenane,K., Wiener,S.I., Buzsaki,G., and Zugaro,M.B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. Nat. Neurosci. 12, 1222-1223.

Graybiel,A.M. (2005). The basal ganglia: learning new tricks and loving it. Curr. Opin. Neurobiol. 15, 638-644.

Grun,S., Diesmann,M., and Aertsen,A. (2002). Unitary events in multiple single-neuron spiking activity: I. Detection and significance. Neural Comput 14, 43--80.

Harley,C.W. (2004). Norepinephrine and dopamine as learning signals. Neural Plast. 11, 191-204.

Harris,K.D., Henze,D.A., Hirase,H., Leinekugel,X., Dragoi,G., Czurko,A., and Buzsaki,G. (2002). Spike train dynamics predicts theta-related phase precession in hippocampal pyramidal cells. Nature 417, 738-741.

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

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. Behav. Neurosci. 111, 20-34.

Hirsh,R., Leber,B., and Gillman,K. (1978). Fornix fibers and motivational states as controllers of behavior: a study stimulated by the contextual retrieval theory. Behav. Biol. 22, 463-478.

Holscher,C., Jacob,W., and Mallot,H.A. (2003). Reward modulates neuronal activity in the hippocampus of the rat. Behav. Brain Res. 142, 181-191.

Hsiao,S., and Isaacson,R.L. (1971). Learning of food and water positions by hippocampus damaged rats. Physiol Behav. 6, 81-83.

Jackson,J.C., Johnson,A., and Redish,A.D. (2006). Hippocampal sharp waves and reactivation during awake states depend on repeated sequential experience. J Neurosci 26, 12415-12426.

Ji,D., and Wilson,M.A. (2007). Coordinated memory replay in the visual cortex and hippocampus during sleep. Nat Neurosci 10, 100-107.

Johnson,A., and Redish,A.D. (2005). Hippocampal replay contributes to within session learning in a temporal difference reinforcement learning model. Neural Netw 18, 1163-71.

Jones,M.W., and Wilson,M.A. (2005a). Phase precession of medial prefrontal cortical activity relative to the hippocampal theta rhythm. Hippocampus 15, 867-873.

Jones,M.W., and Wilson,M.A. (2005b). Theta rhythms coordinate hippocampal-prefrontal interactions in a spatial memory task. PLoS. Biol. 3, e402.

Jung,M.W., Qin,Y., McNaughton,B.L., and Barnes,C.A. (1998). Firing characteristics of deep layer neurons in prefrontal cortex in rats performing spatial working memory tasks. Cereb. Cortex 8, 437-450.

Karlsson,M.P., and Frank,L.M. (2008). Network Dynamics Underlying the Formation of Sparse, Informative Representations in the Hippocampus. J. Neurosci. 28, 14271-14281.

Karlsson,M.P., and Frank,L.M. (2009a). Awake replay of remote experiences in the hippocampus. Nat. Neurosci. in press.

Karlsson,M.P., and Frank,L.M. (2009b). Awake replay of remote experiences in the hippocampus. Nat. Neurosci. 12, 913-918.

Kennedy,P.J., and Shapiro,M.L. (2009). Motivational states activate distinct hippocampal representations to guide goal-directed behaviors. Proc. Natl. Acad. Sci. U. S. A. 106, 10805-10810.

Kim,S.M., and Frank,L.M. (2009). Hippocampal lesions impair rapid learning of a continuous spatial alternation task. PLoS. ONE. 4, e5494.

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

Kudrimoti,H.S., Barnes,C.A., and McNaughton,B.L. (1999). Reactivation of hippocampal cell assemblies: effects of behavioral state, experience, and EEG dynamics. J. Neurosci. 19, 4090-4101.

LaBar,K.S., and Cabeza,R. (2006). Cognitive neuroscience of emotional memory. Nat. Rev. Neurosci. 7, 54-64.

Lansink,C.S., Goltstein,P.M., Lankelma,J.V., Joosten,R.N., McNaughton,B.L., and Pennartz,C.M. (2008). Preferential reactivation of motivationally relevant information in the ventral striatum. J. Neurosci. 28, 6372-6382.

Lansink,C.S., Goltstein,P.M., Lankelma,J.V., McNaughton,B.L., and Pennartz,C.M. (2009). Hippocampus leads ventral striatum in replay of place-reward information. PLoS. Biol. 7, e1000173.

Lee,A.K., and Wilson,M.A. (2002). Memory of sequential experience in the hippocampus during slow wave sleep. Neuron 36, 1183-1194.

Lee,D., Port,N.L., Kruse,W., and Georgopoulos,A.P. (1998). Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. J. Neurosci. 18, 1161-1170.

Leppanen,J.M. (2006). Emotional information processing in mood disorders: a review of behavioral and neuroimaging findings. Curr. Opin. Psychiatry. 19, 34-39.

Leutgeb,J.K., Leutgeb,S., Treves,A., Meyer,R., Barnes,C.A., McNaughton,B.L., Moser,M.B., and Moser,E.I. (2005a). Progressive transformation of hippocampal neuronal representations in "morphed" environments. Neuron. %20;48, 345-358.

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

Leutgeb,S., Leutgeb,J.K., Moser,E.I., and Moser,M.B. (2006). Fast rate coding in hippocampal CA3 cell ensembles. Hippocampus. 16, 765-774.

Leutgeb,S., Leutgeb,J.K., Moser,M.B., and Moser,E.I. (2005c). Place cells, spatial maps and the population code for memory. Curr. Opin. Neurobiol. 15, 738-746.

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.

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.

Maren,S., and Quirk,G.J. (2004). Neuronal signalling of fear memory. Nat. Rev. Neurosci. 5, 844-852.

Maynard,E.M., Hatsopoulos,N.G., Ojakangas,C.L., Acuna,B.D., Sanes,J.N., Normann,R.A., and Donoghue,J.P. (1999). Neuronal interactions improve cortical population coding of movement direction. J. Neurosci. 19, 8083-8093.

McHugh,T.J., Jones,M.W., Quinn,J.J., Balthasar,N., Coppari,R., Elmquist,J.K., Lowell,B.B., Fanselow,M.S., Wilson,M.A., and Tonegawa,S. (2007). Dentate gyrus

NMDA receptors mediate rapid pattern separation in the hippocampal network. Science 317, 94-99.

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. Exp. Brain Res. 52, 41-49.

Miller,E.K. (1999). The prefrontal cortex: complex neural properties for complex behavior. Neuron 22, 15-17.

Mizumori,S.J., Yeshenko,O., Gill,K.M., and Davis,D.M. (2004). Parallel processing across neural systems: implications for a multiple memory system hypothesis. Neurobiol. Learn. Mem. 82, 278-298.

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.

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

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.

Nakazawa,K., Quirk,M.C., Chitwood,R.A., Watanabe,M., Yeckel,M.F., Sun,L.D., Kato,A., Carr,C.A., Johnston,D., Wilson,M.A., and Tonegawa,S. (2002). Requirement for hippocampal CA3 NMDA receptors in associative memory recall. Science. 297, 211-218.

O'Keefe,J. (1993). Hippocampus, theta, and spatial memory. Curr. Opin. Neurobiol. 3, 917-924.

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

O'Keefe,J., and Nadel,L. (1978). The hippocampus as a cognitive map (London: Oxford University 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'Neill,J., Senior,T.J., Allen,K., Huxter,J.R., and Csicsvari,J. (2008). Reactivation of experience-dependent cell assembly patterns in the hippocampus. Nat. Neurosci. 11, 209-215.

Packard,M.G., and McGaugh,J.L. (1996). Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. Neurobiol. Learn. Mem. 65, 65-72.

Palop,J.J., Chin,J., Roberson,E.D., Wang,J., Thwin,M.T., Bien-Ly,N., Yoo,J., Ho,K.O., Yu,G.Q., Kreitzer,A., Finkbeiner,S., Noebels,J.L., and Mucke,L. (2007). Aberrant excitatory neuronal activity and compensatory remodeling of inhibitory hippocampal circuits in mouse models of Alzheimer's disease. Neuron. 55, 697-711.

Pazienti,A., and Grun,S. (2006). Robustness of the significance of spike synchrony with respect to sorting errors. J. Comput. Neurosci. 21, 329-342.

Pennartz,C.M., Lee,E., Verheul,J., Lipa,P., Barnes,C.A., and McNaughton,B.L. (2004). The ventral striatum in off-line processing: ensemble reactivation during sleep and modulation by hippocampal ripples. J Neurosci 24, 6446--56.

Peyrache,A., Khamassi,M., Benchenane,K., Wiener,S.I., and Battaglia,F.P. (2009b). Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. Nat. Neurosci. 12, 919-926.

Peyrache,A., Khamassi,M., Benchenane,K., Wiener,S.I., and Battaglia,F.P. (2009a). Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. Nat. Neurosci. 12, 919-926.

Quirk,M.C., and Wilson,M.A. (1999). Interaction between spike waveform classification and temporal sequence detection. J. Neurosci. Methods 94, 41-52.

Redish,A.D. (1999). Beyond the cognitive map: From place cells to episodic memory (Cambridge, MA: MIT Press).

Redish,A.D., Battaglia,F.P., Chawla,M.K., Ekstrom,A.D., Gerrard,J.L., Lipa,P., Rosenzweig,E.S., Worley,P.F., Guzowski,J.F., McNaughton,B.L., and Barnes,C.A. (2001). Independence of firing correlates of anatomically proximate hippocampal pyramidal cells. J. Neurosci. 21, RC134.

Redish,A.D., and Touretzky,D.S. (1998). The role of the hippocampus in solving the Morris water maze. Neural Comput. 10, 73-111.

Samsonovich,A.V., and Ascoli,G.A. (2005). A simple neural network model of the hippocampus suggesting its pathfinding role in episodic memory retrieval. Learn. Mem. 12, 193-208.

Sara,S.J. (2009). The locus coeruleus and noradrenergic modulation of cognition. Nat. Rev. Neurosci. 10, 211-223.

Schoppik,D., Nagel,K.I., and Lisberger,S.G. (2008). Cortical mechanisms of smooth eye movements revealed by dynamic covariations of neural and behavioral responses. Neuron. 58, 248-260.

Schultz,W. (2000). Multiple reward signals in the brain. Nat. Rev. Neurosci. 1, 199-207.

Schultz,W. (2007). Behavioral dopamine signals. Trends Neurosci. 30, 203-210.

Shadlen,M.N., and Newsome,W.T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. J Neurosci 18, 3870--96.

Sheline,Y.I., Mittler,B.L., and Mintun,M.A. (2002). The hippocampus and depression. Eur. Psychiatry. 17 Suppl 3:300-5., 300-305.

Sheskin,D.J. (2004). Handbook of Parametric and Nonparametric Statistical Procedures (New York, NY: Chapman & Hall/ CRC).

Siapas,A.G., Lubenov,E.V., and Wilson,M.A. (2005). Prefrontal phase locking to hippocampal theta oscillations. Neuron 46, 141-51.

Skaggs,W.E., and McNaughton,B.L. (1996). Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. Science 271, 1870-1873.

Skaggs,W.E., and McNaughton,B.L. (1998). Spatial firing properties of hippocampal CA1 populations in an environment containing two visually identical regions. J. Neurosci. 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.

Smith,A.C., Frank,L.M., Wirth,S., Yanike,M., Hu,D., Kubota,Y., Graybiel,A.M., Suzuki,W.A., and Brown,E.N. (2004). Dynamic analysis of learning in behavioral experiments. J. Neurosci. 24, 447-461.

Smith,A.C., Wirth,S., Suzuki,W.A., and Brown,E.N. (2007). Bayesian analysis of interleaved learning and response bias in behavioral experiments. J. Neurophysiol. 97, 2516-2524.

Squire,L.R. (1982). The neuropsychology of human memory. Annu. Rev. Neurosci. 5:241-73, 241-273.

Squire,L.R., and Alvarez,P. (1995). Retrograde amnesia and memory consolidation: a neurobiological perspective. Curr. Opin. Neurobiol. 5, 169-177.

Squire,L.R., and Zola,S.M. (1996). Structure and function of declarative and nondeclarative memory systems. Proc. Natl. Acad. Sci. U. S. A. 93, 13515-13522.

Sutherland,G.R., and McNaughton,B. (2000). Memory trace reactivation in hippocampal and neocortical neuronal ensembles. Curr. Opin. Neurobiol. 10, 180-186.

Tabuchi,E., Mulder,A.B., and Wiener,S.I. (2003). Reward value invariant place responses and reward site associated activity in hippocampal neurons of behaving rats. Hippocampus 13, 117-132.

Tanila,H. (1999). Hippocampal place cells can develop distinct representations of two visually identical environments. Hippocampus. 9, 235-246.

Tremblay,L., Hollerman,J.R., and Schultz,W. (1998). Modifications of reward expectation-related neuronal activity during learning in primate striatum. J. Neurophysiol. 80, 964-977.

Wiener,S.I., Paul,C.A., and Eichenbaum,H. (1989). Spatial and behavioral correlates of hippocampal neuronal activity. J. Neurosci. 9, 2737-2763.

Wierzynski,C.M., Lubenov,E.V., Gu,M., and Siapas,A.G. (2009). State-dependent spike-timing relationships between hippocampal and prefrontal circuits during sleep. Neuron. 61, 587-596.

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. (1993). Dynamics of the hippocampal ensemble code for space. Science 261, 1055-1058.

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

Wirth,S., Avsar,E., Chiu,C.C., Sharma,V., Smith,A.C., Brown,E., and Suzuki,W.A. (2009). Trial outcome and associative learning signals in the monkey hippocampus. Neuron. 61, 930-940.

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.

Young,B.J., Fox,G.D., and Eichenbaum,H. (1994). Correlates of hippocampal complex-spike cell activity in rats performing a nonspatial radial maze task. J. Neurosci. 14, 6553-6563.

Publishing Agreement

It is the policy of the University to encourage the distribution of all theses, dissertations, and manuscripts. Copies of all UCSF theses, dissertations, and manuscripts will be routed to the library via the Graduate Division. The library will make all theses, dissertations, and manuscripts accessible to the public and will preserve these to the best of their abilities, in perpetuity.

Please sign the following statement:

I hereby grant permission to the Graduate Division of the University of California, San Francisco to release copies of my thesis, dissertation, or manuscript to the Campus Library to provide access and preservation, in whole or in part, in perpetuity.

Author Signature

12/21/09

Date