UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

The Contributions of CA1, CA2, and CA3 to Hippocampal Coding for Temporal and Spatial Context

Permalink

https://escholarship.org/uc/item/8sd27587

Author

Mankin, Emily A.

Publication Date

2014

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

The Contributions of CA1, CA2, and CA3 to Hippocampal Coding for
Temporal and Spatial Context

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Neurosciences,
with a specialization in Computational Neuroscience

by

Emily A. Mankin

Committee in charge:

        Professor Stefan Leutgeb, Chair
        Professor Timothy Gentner
        Professor Jill Leutgeb
        Professor Terrence Sejnowski
        Professor Lawrence Squire

2014

The Dissertation of Emily A. Mankin is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

_____

_____

_____

_____

_____

Chair

University of California, San Diego

2014

iii

DEDICATION

This dissertation is dedicated to my parents,

who instilled in me a sense of curiosity about the world and taught me from the youngest age

that for every fascinating question there is an experiment waiting to be done,

to my brother,

whose love of all things math and science inspired me to love them too,

and to my husband,

who believes in me unfailingly and makes every day a joy.

EPIGRAPH

Time is what keeps everything from happening all at once.

—Ray Cummings

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

I would like to acknowledge Professor Stefan Leutgeb for his support as chair of my committee. Under his guidance I have learned to think as a neuroscientist and to persevere when times are hard or data are confusing. I would also like to acknowledge Professor Jill Leutgeb for allowing me to join "team time", and for teaching this mathematician to perform neurophysiology experiments.

I would like to acknowledge my fellow lab-mates who have worked alongside me, sharing the joys and frustrations of this research. They have kept me passionate about science and have helped me immeasurably on this journey.

I would also like to thank my committee for sharing their time with me, and for the many ways in which they helped me get here. Their ideas and advice have been critical to the success of this project.

VITA

Bachelor of Arts, Rice University, Houston, TX

Assistant Language Teacher, Tagajo High School, Miyagi, Japan

Teaching Assistant, Department of Mathematics, University of Colorado, Boulder

Master of Arts, University of Colorado, Boulder

Research Assistant, University of California, San Diego

Doctor of Philosophy, University of California, San Diego

PUBLICATIONS

Mankin, Emily A., Diehl, Geoffrey W., Sparks, Fraser T., Leutgeb, Stefan, and Leutgeb, Jill K. "Hippocampal CA2 activity patterns change over time to a larger extent than between spatial contexts." *Neuron*. In Press.

Leutgeb, Jill K., Mankin, Emily A., and Leutgeb, Stefan. "Population Coding by Place Cells and Grid Cells." In R.Q. Quiroga and S. Panzero (Eds), *Principles of Neural Coding* (pp. 285-302). CRC Press. 2013

Mankin, Emily A., Sparks, Fraser T., Slayyeh, Begum, Sutherland, Robert J., Leutgeb, Stefan, and Leutgeb, Jill K., "Neuronal Code for Extended Time in the Hippocampus." *Proceedings of the National Academy of Sciences*. Vol. 109:47, p 19462-19467. 2012.

FIELDS OF STUDY

Major Field: Biology (Neuroscience and Computational Neuroscience)

Studies in Neuroscience
Professors Stefan Leutgeb and Jill Leutgeb

Studies in Mathematics
Professor Kent Goodrich

Studies in Cognitive Science
Professor Randi Martin

ABSTRACT OF THE DISSERTATION

The Contributions of CA1, CA2, and CA3 to Hippocampal Coding for
Temporal and Spatial Context

by

Emily A. Mankin

Doctor of Philosophy in Neurosciences, with a Specialization in Computational Neuroscience

University of California, San Diego, 2014

Professor Stefan Leutgeb, Chair

Memories for autobiographical events contain information not only about what happened, but also about when and where. The hippocampus is known to be critical for organizing autobiographical memories, yet little is known about how the temporal context of a memory is represented there, especially on the timescale of hours to days. We sought to elucidate the coding mechanisms used by the hippocampus to represent the spatio-temporal context of events. In our

experimental paradigm, rats performed random foraging tasks twice a day for several days, as we conducted electrophysiological recordings in three subregions of the hippocampus, CA1, CA2, and CA3. By recording from the same cells for 30 or more hours, we were able to evaluate the impact of passing time on the representation of identical events in each of these subregions. We performed these experiments within two distinct arenas, allowing us to explore coding for spatial context in addition to temporal context. We found that the population of cells in CA3 showed robust coding for spatial context that was very stable for a period of at least 30 hours. Conversely, neuronal ensembles in CA2 were highly variable on the time scale of hours to days, with the change in time overshadowing change with spatial context. Finally, CA1, which receives input from each of these regions, showed a gradual change with time while also coding strongly for spatial context. These findings correspond well with theoretical predictions that time-stamped memories may be generated by combining a stable *content* signal with a time-varying *temporal context* signal. We thus propose that CA2 provides temporal context information to CA1, while CA3 provides stable information about space and spatial context, allowing CA1 to represent both types of information simultaneously.

INTRODUCTION

What-Where-and-When: Episodic-like Memory and Time

>Space is one fundamental behavioral matrix; time is the other. Animals ranging from insects to humans record the time of occurrence of an event or observation. They also extract elapsed intervals. This capacity is the foundation of the capacity to represent the temporal structure of experience. A representation of the temporal structure of experience makes it possible for an animal to predicate its behavior on what may be predicted to occur, that is, to adapt its behavior to the temporal structure of the events in its behavioral space (Gallistel, 1990).

In 1972, Tulving coined the term "episodic memory" to distinguish between memories that have specific reference to autobiographical events and "semantic memories," those that have to do with general knowledge about the world (Tulving, 1972). Because our lives unfold over a one-dimensional timeline, a critical and well-defined aspect to any autobiographical memory is its temporal relationship to other events and memories. Specifically, Tulving suggested that an episodic memory "can be reasonably completely described in terms of (a) its perceptible properties and (b) its temporal-spatial relation to other experienced events." Or more simply, episodic memories contain information about not only *what* happened but also *when* and *where*. Though his term, "episodic memory" has stuck, and its definition has subsequently come to include "mental time travel" and other facets that are difficult to assess in non-verbal animals (Suddendorf & Busby, 2003; Tulving & Markowitsch, 1998), Tulving was neither the first nor the last to recognize that memories containing information about *what*, *where*, and *when* something happened hold a special place for humans, as they allow reminiscing about previous experience, and are likely critically important, even for animals, for making appropriate behavioral choices based on the learned spatio-temporal structure of the world (Clayton & Dickinson, 1998; Gallistel, 1990; Reiff & Scheerer, 1959, cited in Tulving, 1972).

Gallistel distinguished between two types of time sense that an animal may employ: "phase sense" and "interval sense." Phase sense is the ability to read out the phase of an

endogenous oscillator, for example to estimate the time of day from the circadian oscillator. Interval sense is the ability to measure elapsed time or the duration of a stimulus (Gallistel, 1990). Rats are able to use each of these timing mechanisms.

Although rats are known to have endogenous ultradian oscillators (Gallistel, 1990), most behavioral studies of phase sense focus on circadian oscillators. In conditions of food deprivation, when rats were fed at a specific time of day and had access to a running wheel, they developed anticipatory running behavior, such that their activity score ramped up steeply in the hours before feeding. They did not adapt this behavior to 19- or 29-hour feeding cycles, even when born and reared in congruent environmental conditions, with lights turned on and off on a 19- or 29-hour cycle (Bolles & De Lorge, 1962; Bolles & Stokes, 1965), suggesting that anticipatory running is likely dependent on rats' endogenous circadian rhythm. Rats could keep track of two daily feeding times, indicated by anticipatory running preceding each feeding. This behavior developed over a week of training, was abolished when rats were returned to ad lib feeding, and was reinstated after one day of return to food deprivation with two daily feedings (Bolles & Moot, 1973). This could be indicative of long-term memory for time of day (Gallistel, 1990), but because the daily feedings were given at their usual time on the first day of return to food deprivation, it may indicate only long-term memory for the "task," and not for the specific time of day previously learned. Rats also show a cyclical response to single-trial learning tasks; in both active and passive avoidance tasks, regardless of the time of day of training, their probability of expressing avoidance learning after a single retention interval is highest immediately following the learning or 24 hours later, lowest after 6 or 18 hours, and intermediate after 12 hours (Holloway & Wansley, 1973a, 1973b; Wansley & Holloway, 1976). Though the authors interpret this as an oscillatory deficit in retention, it could also reflect that rats associate the time of day of with the shock experience and estimate that the probability of being re-shocked is the greatest at the same time of day.

In addition to circadian timing, rats are capable of performing a variety of interval timing tasks, from the simple to the complex: wait for a given interval after some signal then press a lever to receive food reward (Meck & Church, 1984; S. Roberts, 1981), press a lever if the duration of a signal matches a learned reference duration (Church & Gibbon, 1982), or press one lever continuously for a given duration and then switch to a different lever to minimize the waiting time until reward (Gibbon & Church, 1981). Animals have been successful on these tasks on time scales of seconds to nearly an hour (reviewed in Buhusi & Meck, 2005).

The studies discussed above (and others like them) indicate that rats (as well as birds and bees) have access to temporal information to guide their behavior. Yet none of those tasks required the integration of *when* information with *what*-and-*where* information. Clayton and Dickinson (1998) established an experimental paradigm for animals for which success required the integration of *what-where-and-when* into memory traces. Scrub jays, birds that naturally cache food for later retrieval, were given the opportunity to cache highly desirable wax worms, as well as less desirable peanuts. A series of training experiences demonstrated that the worms would be fresh after 4 hours but would degrade after 124 hours, while the peanuts never degraded. At test, the birds were given the opportunity to cache one food, wait 120 hours, and then cache the other in a visuo-spatially distinct tray. After another 4 hours, they had the chance to retrieve their caches. Birds who cached worms first, such that it had been 124 hours since they had been cached and experience would indicate that the worms should have degraded, preferentially searched for their caches in the location where they had previously stored peanuts. Conversely, birds who cached worms second, such it had been only four hours since caching, preferentially searched for their caches of worms. To successfully complete this task, as the scrub jays were able to, required that the birds remember which of the two foods they had stored in each location and know how long it had been since caching the wax worms. Thus it appeared that

the scrub jays were able to incorporate information about *what-where-and-when* into memories in order to make good behavioral choices.

The question of how and whether rats incorporate timing cues with *what-and-where* information to modulate their behavior has received extensive attention. Of particular note is a series of experiments by Babb and Crystal, who adapted the experimental paradigm of Clayton and Dickinson for rats (Babb & Crystal, 2005, 2006a, 2006b). In these experiments, rats searched for food reward on an 8-arm radial arm maze. During the sample phase, 4 arms were baited with rat chow and one was baited with chocolate, which rats prefer. During the test phase, the previously un-baited arms were always baited with rat chow. The arm previously baited with chocolate was either empty or re-baited with chocolate, depending on whether the interval between the sample and test phase was 30 minutes (empty) or 4 hours (chocolate). After training, rats showed a higher probability of re-visiting the chocolate arm after a 4-hour interval than a 30 minute interval (Babb & Crystal, 2005), indicating that rats could use the retention interval as a cue to find food reward. In a similar task, two arms were baited with treats (chocolate or banana) during the sample phase and were re-baited for the test phase only if 4 hours had elapsed. Again, rats showed stronger preference for the treat arms on the long retention interval trials. Next, chocolate was paired with lithium chloride to induce taste aversion during a long retention interval; at test rats revisited the arm where banana had previously been present but now avoided the chocolate arm (Babb & Crystal, 2006b). Success on this task required that the rats remember what food they had encountered on each arm (*what* and *where*) and integrate that information with how long ago they had done the task (*when*). Because the pairing of chocolate with LiCl occurred between the study and test phase, it was clear that rats were not using a strategy of encoding only whether an arm was distinctive or not, but rather that they remembered specifically which arm had contained which food.

What kind of timing cue were rats using on this radial arm maze task? Because the sample phase occurred at the same time each day, the test phases were distinctive not only in the duration of the retention interval but also in the time of day at which they occurred, making either a "phase sense" or an "interval sense" explanation plausible. Babb and Crystal showed that rats could perform the radial arm maze task when the retention intervals were 1 and 25 hours, such that the test phase occurred at the same time of day (Babb & Crystal, 2006a), indicating that distinct times of day were not necessary for performance of the task. Roberts et al followed up with a similar task, manipulating the conditions such that for one group of rats (the "when group") the time-of-day of the sample phase was the relevant cue, while for the other group (the "how-long-ago group") the relevant cue was the time elapsed since the sample phase. Rats in the when group could not learn the task, while rats in the how-long-ago group succeeded (Roberts et al., 2008). These results were consistent with earlier studies by Wilkie and colleagues, who had shown that on a wide variety of tasks, rats have difficulty with time-place-event associations when time-of-day is the relevant cue (Carr & Wilkie, 1997; 1999; Thorpe, Bates, & Wilkie, 2003). Thus, there is converging evidence that when a behavior requires integration of *what-where*-and-*when*, rats are more likely sensitive to elapsed time than to time-of-day. It is worth noting, however, that when the stakes are high, such as in contextual fear conditioning or increased response-cost, or when initial training paradigms are shifted toward making time-of-day the only relevant cue, rats are able to use time-of-day as a cue in *where-when* or *what-where-when* memory tasks (Iordanova, Good, & Honey, 2008; O'Brien & Sutherland, 2007; Widman, Gordon, & Timberlake, 2000; Widman, Sermania, & Genismore, 2004; Zhou & Crystal, 2009).

Models of Time and Memory

One of the first efforts to study memory systematically was done by Ebbinghaus (1885, translated to English in 1913), who, testing his own ability to remember lists of nonsense syllabi,

made the not-so-surprising discovery that the longer he waited between learning and testing, the more he had forgotten. Though this result seems intuitively obvious, since it so closely reflects our own experiences, the mechanism of this forgetting was not as immediately clear, especially when juxtaposed with another of his results, that a previously forgotten list of syllables may be relearned more quickly than it was initially learned, which indicates that apparently forgotten memory traces have not simply vanished with time. Ebbinghaus also documented the distributed learning effect, in which studying the same material at more spread-out time intervals resulted in better memory performance than the same amount of studying done in a condensed time period. Though these ideas were not new at the time, Ebbinghaus's approach to studying memory experimentally was. Thus began a rich tradition of studying (what later became known as) episodic memory by presenting subjects with lists of words and measuring how various manipulations affect their recall.

Many theories addressing time and memory begin by trying to account for some the peculiarity of how memory changes over time. Why do we forget? Why do we spontaneously recover a memory that has been subject to extinction? Why do we remember the first and last elements of a list better than those in the middle? If we remember a word from the middle of the list, why are we much likelier to remember the word presented immediately afterward? And why does this effect depend on the time between item presentations?

Estes (1955) sought to explain the phenomena of spontaneous regression and recovery from extinction with a rigorous mathematical model that would make predictions that could be compared to empirical data. He proposed a simple model whereby "elements" (which he left undefined but could probably be considered neurons or ensembles of neurons) could be part of an available set or an unavailable set, and they would move spontaneously between these sets during rest intervals with some fixed probability. The available elements could be modified by training during learning and retrieved during recall, while the unavailable elements could not.

Spontaneous regression–that is, forgetting something previously learned–could be accounted for by an element that was trained while in the available set moving into the unavailable set, and conversely, recovery from extinction could be accounted for by an element that was once trained moving into the unavailable set during extinction, and then returning to the available set. This rather simple theory accounted not only for the existence of spontaneous regression and recovery from extinction, but also for the shape of their curves and the ways in which the asymptotes depended on training procedures.

Raaijmakers and Shiffrin (1981) and Murdock (1982) independently put forward the Search of Associative Memory (SAM) and the Theory of Distributed Associative Memory (TODAM) models, respectively, to account for experimental data regarding humans' ability to learn lists of associated word-pairs. Though somewhat different in their implementation, each model required not only item information (a vector of features specific to each word studied) but also context information (a vector of other features, which may include such things as physical location, motivational state, and which list is being studied). Upon considering certain experimental effects the models failed to describe, both models were later modified to include a contextual component that drifted with time. To account for differences in rates of forgetting for items and associations, Murdock introduced TODAM2, in which the values in the context vector changed by a certain (small) percent at each time step (1997). Mensink and Raaijmakers introduced a theory of contextual fluctuation (1989), quite similar to the stimulus fluctuation of Estes (1955), in which contextual components moved between active and inactive sets. They combined contextual fluctuation with the SAM model to account for the time course of proactive and retroactive interference in associative memory tasks (G.-J. Mensink & Raaijmakers, 1988).

Another time-based peculiarity of episodic memory is seen in free-recall paradigms. If a subject studies a list of words and then is given an opportunity to spontaneously recall as many as possible, words studied at the end of the list are usually recalled first, and once one word is

recalled, the words that were studied at similar times, and in particular the word that was studied next, are preferentially recalled. Howard and Kahana (1999) documented these effects, naming them the recency and lag-recency effects. To account for them, they added a criterion to the SAM+contextual fluctuation model, requiring that retrieval of an item also retrieve its associated context, thereby forcing the network to "jump back in time." In a subsequent update of the model, they replaced the "random fluctuations" of context with a contextual update based on retrieved context. In this "Temporal Context Model," context still changes gradually over time (Howard & Kahana, 2002).

Though an evaluation of the precise implications of the differing implementations of these models is outside the scope of this chapter, it is worth noting themes common to all of them. First, there is a representation of each "item," the *what*-component of each memory. There is also a contextual component that is somehow combined with the item representation. Finally, to account for many types of experimental data, the contextual component must vary with time in such a way that the similarity of two contextual representations decreases as a function of time between them. These theories were all generated from data related to human memory for lists of words. Nonetheless, Bouton points out that many effects of proactive or retroactive interference observed in classical conditioning paradigms with animals can also be parsimoniously explained with a few assumptions, chief among them that contextual stimuli guide memory retrieval and that time itself serves as a context (Bouton, 1993).

Potential Mechanisms that could be Used for Timing

We have seen that models of episodic memory seem to require a context signal that gradually varies with time. This does not necessarily indicate how estimates of time are performed. To query how one estimates time, it is again necessary to think about to what kind of time and to what timescale we are referring. Time-of-day, recency, duration of an interval,

relative order or events, absolute dates or time-stamps, and even whether two events happened as part of the same or different occasion may be considered part of the *when* component of episodic memory. Yet they may have distinct neural implementations. And even the same type of timing may differ in its implementation when considered on differing time scales.

Gallistel distinguished between two putative mechanisms by which an animal may estimate time (Gallistel, 1990). With an oscillatory timer, an animal marks the phase of one or more endogenous oscillators, such as a circadian oscillator for marking the time of day. Higher frequency oscillators provide higher precision, and a combination of oscillators with several frequencies could allow for precise measurement across broad time scales, in much the same way that the phases of the second hand, minute hand, and hour hand of an analog clock provide sub-minute temporal resolution across a period of 12 hours. Along these same lines, yet another model for associative memory, which tries to account for many experimental observations related to serial order effects, is OSCAR (OSCillator-based Associative Recall) (Brown, Preece, & Hulme, 2000). Like the models already discussed, this model combines a "learning-context vector" with an "item vector." In OSCAR, the values in the temporal context vector are products of a subset of endogenous oscillators, each oscillating at a different frequency.

As with a clock, time points at the same phase of the longest oscillator would be indistinguishable. This could be advantageous for animals for whom the memory of yesterday need not be distinguished from the memory for the day before yesterday. For example, bees need only remember the time of day at which a flower produces pollen to make appropriate decisions about when and where to forage (Gallistel, 1990; Moore, 2001) and have not been shown to use timing on scales longer than their circadian oscillator (Gallistel, 1990). Timing with oscillators could also be sufficient to describe effects of word-list learning, whose duration is less than a circadian period. But for longer-term episodic-like memories, oscillator phase is unlikely to be the only solution employed for the representation of time. Decay or accumulation processes are

Gallistel's second proposed timing mechanism. Again, a combination of these processes on multiple time scales would be required to maintain a precise representation of time over a long period. To achieve "interval sense," an accumulator could count the number of times and oscillator completes a cycle (Gibbon, Church, & Meck, 1984) or by comparing start and stop times that were recorded by either an oscillator or an accumulator or decay process (Gallistel, 1990).

While Gallistel, Bolles, Gibbon, Church, and others were interested in how animals compute time of day and measure temporal intervals, others have been particularly interested in understanding retrospective judgments of elapsed time. Friedman (1993) reviewed many experiments testing these retrospective judgments in humans and grouped the theories that had been put forth to account for temporal judgments into three broad categories: those based on distance, those based on location, and those based on relative times. In distance-based theories, "memory for time depends on information about the age of a memory, and this information is derived from a basic process associated with the passage of time (or the occurrence of subsequent events)" (Friedman, 1993). Such theories include the strength theories, which posit that more recent events have stronger memory traces, so that the recency of an event may be read out by the strength of the related trace; chronological organization theory, which supposes that as new events are encoded, memories for older events get pushed backward on a timeline like bags on a conveyor belt; and the contextual overlap theory, which suggests that context changes over time and the amount of contextual overlap between the remembered event and the time of retrieval can be used as a cue for recency. Friedman describes the changing context as primarily external, such as the difference between winter and spring. Though not addressed by Friedman directly, Block's contextual change theory (Block, 1985; Block & Reed, 1978) also fits within a contextual overlap framework, but Block thinks of the changing context as internal changes in context, which may be accumulated by doing multiple types of tasks or even by increasing boredom. Gallistel's decay

processes could probably be grouped among Friedman's distance-based theories, as well, though Gallistel seems to imagine a separate timer that can be started and stopped for each event, while the theories Friedman describes seem instead to suppose that the decay process resides within the memory trace itself.

Friedman's location based theories are based on the idea that instead of measuring recency by looking at an associated process that changes with time, information about time is directly encoded as part of a memory trace in the same way that a fixed spatial location could be encoded. Time-tagging theories suppose that a time-stamp is directly encoded into a memory trace, which is potentially consistent with Gallistel's circadian oscillator time stamp proposal (Gallistel, 1990). Reconstruction theories suppose instead that the nature and context of a stimulus, along with semantic knowledge about how time progresses, can be used to reconstruct the time of a memory. Finally, two classes of theory are related to relative times of occurrence. Associative chaining theories suppose that temporal information is stored by linking subsequent events together in memory, while order code theories suppose that whenever one event causes the retrieval of an older memory, the relative order of the two events is encoded.

These theories should not be seen as mutually exclusive. In fact, if multiple mechanisms existed side by side, they would probably complement each other and improve temporal judgments. Distinct mechanisms may also be used for distinct time scales. For example, associative chaining could connect events on a short time scale to keep track of the order of events within an episode, but judging the recency of distinct episodes may require a different kind of timing all together. Experiments on retrospective estimate of time are conducted with humans, generally by asking them to recall an event and indicate how long ago the event occurred. These experiments are typically done either by asking about historical or life events that happened up to many years ago, or by presenting a list of words and then asking how long ago a particular word occurred (Friedman, 1993). Very little research was cited for episodic memory for time scales

between minutes and years. Friedman presents mixed evidence for each theory, with some of the negative evidence pointing out that various theories do not work at particular time scales. This does not seem to be a reason to reject a theory entirely, however. Rather this is an indication that questions about the mechanisms involved in temporal memory need to be studied on multiple time scales.

Computational models for episodic memory generally contain a stable item signal combined with a context signal that drifts with time. This may take the form of elements that turn on and off (Estes, 1955; G. Mensink & Raaijmakers, 1989; G.-J. Mensink & Raaijmakers, 1988), combine endogenous oscillators (Brown et al., 2000; Gallistel, 1990), or otherwise vary with time (Howard & Kahana, 2002; Murdock, 1997). Where in the brain might such a time-varying contextual signal be found? We turn now to the hippocampus, a brain region thought to be critically involved in binding context to memories.

The Hippocampus: its role in memory and the cognitive map

Ever since patient H.M.'s bilateral resection of the medial temporal lobe (MTL) resulted in his complete inability to form new memories, along with severe retrograde amnesia (Scoville, 1954; Scoville & Milner, 1957), it has been known that the MTL is critically important for certain forms of memory. Specifically, damage to the MTL causes severe impairment in declarative memory, including episodic memory, while non-declarative memory (such as skill-learning) remains unaffected (Squire, 1992). The MTL includes the hippocampal formation (the dentate gyrus, the CA fields, and subiculum), the entorhinal cortex, the parahippocampal cortex, amd the perirhinal cortex. Teasing out the role of the hippocampus vs. the other areas of the MTL has involved further study with human patients (many of whom experienced ischemic damage to a sub-portion of the MTL), along with lesion studies in rodents and non-human primates. Converging evidence from these lines of study has indicated that the hippocampus-proper plays a

critical role in binding together the content and context of an episodic memory (for a comprehensive review, see Squire, 1992; Zola-Morgan & Squire, 1993). Rats with hippocampal lesions are impaired on a wide array of memory tasks, including spatial tasks such as the Morris water maze, radial arm maze tasks, and spatial alternation tasks, and non-spatial tasks such as odor discrimination, cue relationships, and novel object preference (Squire, 1992).

The role of the hippocampus in spatial navigation and other spatial memory tasks has received abundant attention. This is largely because of the discovery that most principal cells in rat hippocampus function as "place cells" (J O'Keefe, 1976; J O'Keefe & Dostrovsky, 1971). That is, as a rat explores an arena, the firing of any single principal cell tends to cluster within one (or a few) specific regions within the arena. Based on this observation, it is not a far stretch to suppose that hippocampal firing in the rat indicates the rat's precise location in space, and O'Keefe and Nadel soon put forward the theory that the hippocampus functions as the basis for the "cognitive map" (J O'Keefe & Nadel, 1978, 1979). Only a few years later, the "water maze" task was introduced by Morris (Morris, Garrud, Rawlins, & O'Keefe, 1982), who demonstrated that rats with hippocampal lesions are severely impaired at learning/navigating to the fixed location of a hidden escape platform in a pool of opaque water. This has become one of the most widely recognized hippocampus-dependent memory tasks. Furthermore, a number of elegant models have demonstrated how the physiology of the entorhinal cortex and hippocampus allow the hippocampus to update its representations based on self-motion cues, suggesting that the hippocampus performs primarily as a path-integrator (B. McNaughton et al., 1996; Samsonovich & Mcnaughton, 1997), such that the ensemble response of hippocampal cells may give high-precision information about the current position of the rat (M. A. Wilson & McNaughton, 1993).

Space is not the only domain of the hippocampus, however. Even before the cognitive map theory was proposed, Hirsh (1974) put forward the idea that an important role of the hippocampus is to provide contextual information that is separate from a direct stimulus-response

behavior. That is, the hippocampus could help an animal select the relevant pair from a variety of possible stimulus-response relationships. For example, he suggested that without a hippocampus a rat could easily learn that location X is rewarded, but to learn that location X is rewarded with water, location Y is rewarded with food, and the appropriate behavior when hungry is to seek location Y would require the hippocampus. Indeed, rats with hippocampal lesions are unable to pair their motivational context (hungry or thirsty) with external cues for where to find food or water (Kennedy & Shapiro, 2004). The hippocampus is required for the expression of freezing in contextual fear conditioning (Kim & Fanselow, 1992; Penick & Solomom, 1991; Phillips & LeDoux, 1992), which is generally regarded as evidence that the hippocampus is required for memory of a context or for binding the context and shock into a single representation (but see Gewirtz, McNish, & Davis, 2000 for an argument that the context is still remembered; Anagnostaras, Gale, & Fanselow, 2001 for an alternative view; and Frankland, Cestari, Filipkowski, McDonald, & Silva, 1998 for a proposal that the hippocampus is necessary for contextual discrimination even if not for fear conditioning).

Additionally, the hippocampus appears to be required for some temporal context tasks. Rats with hippocampal lesions can perform timing tasks such as the peak procedure or the peak procedure with gaps (Dietrich & Allen, 1998; though fornix lesions impair performance: Olton, Meck, & Church, 1987). However, when timing information must be combined with other information, the hippocampus is necessary. For example, rats with hippocampal lesions can recognize previously experienced odors but cannot maintain a representation of the order of their presentation (Fortin, Agster, & Eichenbaum, 2002). Similarly, mice with hippocampal lesions are impaired on a *what-where-when* task designed to assess episodic-like memory with minimal training. This task takes advantage of a natural preference in mice to explore novel items. After two study sessions spaced 50 minutes apart, intact mice explore "old" objects more than "recent" objects and "displaced" objects more than "stationary" objects at a test session 50 minutes later

(Dere, Huston, & De Souza Silva, 2005), while those with hippocampal impairments do not show these trends (DeVito et al., 2009; Good, Barnes, Staal, McGregor, & Honey, 2007). These observations support the idea that the hippocampus is critically involved in binding together information from a variety of modalities (time, space, odors, visual cues, internal motivational state, etc.) in order to make behavioral choices that are appropriate for a given context (based on episodic-like memory for similar events).

It is clear that the hippocampus is necessary for episodic-like memory and also for spatial navigation, yet how those functions interact has remained unresolved. For many years these functions have been considered separately, and have even been supposed to be mutually exclusive (Morris, 1990, cited in Eichenbaum, 2013; O'Keefe, et al, 1998; O'Keefe, 1990), though this was not originally the authors' intended interpretation of the cognitive map theory (O'Keefe & Nadel, 1978; cited in Derdikman & Knierim, 2014). Recently, however, views that reconcile the two have become more popular (e.g. Allen and Fortin, 2013; Buzsáki and Moser, 2013; Eichenbaum, 2013; Smith and Mizumori, 2006a). Generally, these suggest that spatial navigation mechanisms that are evolutionarily old have been co-opted and adapted to serve episodic memory demands (though the argument may also be made that the spatial navigation system instead co-opted earlier episodic memory mechanisms (H Eichenbaum, MacDonald, & Kraus, 2014)). Location-specific firing still undergirds the structure of hippocampal representations, but changes in spiking intensity within a place field may reflect both spatial and non-spatial aspects of context (reviewed in Leutgeb and Leutgeb, 2014 and below). Temporal mechanisms such as phase precession (described below) that may have initially served in support of maintaining spatial sequences for navigation could also be used to impart temporal structure to episodic memories so that the order of events may be recalled (Buzsáki & Moser, 2013; Howard Eichenbaum, 2013).

Since O'Keefe's first description of place cells (J O'Keefe, 1976), intense effort has been devoted to understanding what cues govern the locations and firing rates of place fields. Early

work was centered on questions relating to whether firing was determined by the animal's trajectory or by nearby visual cues, and what would happen to the fields if cues came into conflict. Though fields appeared to be under the control of visual cues, rotating their firing if the visual cues were rotated, they typically continued firing in the same location when a subset of cues was removed (J O'Keefe & Conway, 1978), and even in complete darkness when no visual cues were available (Markus, Barnes, McNaughton, Gladden, & Skaggs, 1994; B. L. McNaughton, Leonard, & Chen, 1989; J O'Keefe, 1976; Quirk, Muller, & Kubie, 1990). Muller and Kubie introduced the "random foraging" task, in which rats chased food pellets delivered randomly within an open arena, in order to query whether the exact path of the rat was a correlate of place field firing – they found that it was not, and place field location was stable across multiple exposures to the environment (Muller, Kubie, & Ranck Jr, 1987). They further confirmed that fields followed a cue card when it was rotated, and coherently rotated to a random degree if the cue card was removed altogether (Muller & Kubie, 1987).

These findings seem to support a theory that place field firing may be locked to visual environmental stimuli originally but can be maintained by integrating self-motion cues (known as path integration) (B. McNaughton et al., 1996; Samsonovich & Mcnaughton, 1997). A clever demonstration of this involved rats running between a "start box" and the end of a linear track. While the animal was running, the start box could be moved, so that there was misalignment between path-integration coordinates and real-world coordinates. If the mismatch was large enough, the hippocampal representation underwent an abrupt shift from the representation expected from run-distance to the representation expected from physical location (Gothard, Skaggs, & McNaughton, 1996). Taken together, these findings seemed to confirm that spatial location is the dominant correlate of firing in hippocampal complex spike cells and justify the view that, at least in rats, the hippocampus is purely spatial (O'Keefe et al., 1998; O'Keefe, 1990).

A more nuanced view has emerged, however, as understanding has grown and experimental paradigms have expanded. Small changes to training protocols, task demands, or environmental context can lead to surprisingly distinct hippocampal representations if one imagines that spatial location is the only determinant of place field firing. For example, place fields may or may not follow the rotation of a cue card, depending on whether a rat been trained that the cue card's location in space is reliable (Jeffery, 1998; J. Knierim, Kudrimoti, & McNaughton, 1995). Minor visual changes to the spatial context such as changing the color of the cue card or box walls, or more major changes such as switching the shape of the box between square and circle, lead to divergent representations despite the arena remaining in the same spatial location (Bostock, Muller, & Kubie, 1991; S. Leutgeb et al., 2005; Lever, Wills, Cacucci, Burgess, & Keefe, 2002). Conversely, visual cues are not the sole determinant of place field representations: when rats explored visually identical boxes in distinct locations, the hippocampal representations of the two boxes were independent (S. Leutgeb et al., 2005; W. E. Skaggs & Mcnaughton, 1998; Tanila, 1999).

Although the hippocampal representations in distinct environments differ in each of these cases, an important distinction must be made about the way in which they differ (S. Leutgeb & Leutgeb, 2007). In the second case, the boxes are in physically distinct locations, and the cells that are involved in the representations overlap no more than would be expected by chance. If one imagines place fields creating a "map" of an environment (J O'Keefe & Nadel, 1978), it is reasonable to expect that distinct elements create the maps for distinct locations. This phenomenon has been termed "global remapping". In the first case described above, the boxes occupy the same physical location. Although it often appears that a cell has lost a field or the field location has moved, careful analysis may show that the original field is still present but with very low firing rate, and the "new" field may have existed previously, but with a firing rate too low to detect (S. Leutgeb et al., 2005), a phenomenon known as "rate remapping". This may sound like

semantics, as it can be difficult to distinguish a set of cells with independent representations from a set of cells with fields whose firing rates have diminished so much as to be undetectable. However, there are important implications of having both systems of remapping. Furthermore, the distinction can be shown truly to exist by varying the contexts in incremental steps and observing incremental increases or decreases in firing rate in the rate remapping paradigms (Colgin et al., 2010; J. K. Leutgeb et al., 2005). One important implication of rate remapping is its ability to bridge the gap between the spatial and the episodic/contextual theories of hippocampal function. We see that the fixed locations of firing fields in a given physical location is the structural backbone of hippocampal representations, but that firing rate modulation on top of this structure allows for contextual information to be integrated into these representations (S. Leutgeb & Leutgeb, 2007).

The appearance or spatial geometry of a random foraging box is not the only kind of contextual information that can be encoded by the hippocampus. In addition to varying box color, Anderson and Jeffery (2003) varied the dominant odor in a random foraging box, creating four odor-color combinations. While some cells varied along only one dimension (distinguishing, for example between the square and circle boxes, but having the same response regardless of whether the odor was lemon or vanilla), many cells exhibited conjunctive responses, treating the four boxes as separate contexts. The possibility of conjunctive representations within single cells seems to be a common feature of the hippocampus. In a task where rats searched for a fixed reward location after being released from a start box whose location varied within an open field, some hippocampal cells functioned as purely spatial, while others followed the start box or represented the rewarded location. Still other cells had multiple fields, and these fields could be controlled independently by any of these cues (Gothard, Skaggs, Moore, & McNaughton, 1996). Several tasks have been developed in which a rat must cover the same portion of a maze but some non-spatial contextual element determines appropriate behavior (K. Allen, Rawlins, Bannerman,

& Csicsvari, 2012; Ferbinteanu & Shapiro, 2003; L. Frank, Brown, & Wilson, 2000; Smith & Mizumori, 2006b; E. R. Wood, Dudchenko, Robitsek, & Eichenbaum, 2000). In all cases, the hippocampal response on the overlapping portion of the maze depended on the context that signaled the next correct decision. In many cases, the hippocampal activity on the overlapping portion could be seen to reflect either prospective or retrospective coding, suggesting that the rat was either remembering where it had just been or considering where it was about to go. In at least one case, this was correlated with correct behavioral choices (Ferbinteanu & Shapiro, 2003). In another, rats were presented with an odor cue to indicate which arm of a T maze would be rewarded on the next trial. Place fields on the stem of the maze were spatially stable, but responded with differing rates based on the odor just presented, the most recently visited location, or in a conjunctive manner, with four possible firing rate profiles (K. Allen et al., 2012). Differing hippocampal responses on the overlapping portion of a T-maze have even been observed when that portion was replaced with a treadmill, such that the spatial location of the rat was held constant, while one of two sequences of hippocampal cells fired, depending on which direction the rat was heading next (Pastalkova, Itskov, Amarasingham, & Buzsáki, 2008).

Finally, it is important to note that hippocampal cells can represent completely non-spatial items in non-spatial tasks. For example, on a delayed non-match to sample task in which rats were presented with one odor, waited for a delay period, and were presented with another odor, rats had to determine whether the new odor was the same as the previous odor. In this task, many cells showed non-spatial responses, including cells that were odor-selective and cells that were selective for match or non-match (Wiebe & Stäubli, 1999; E. Wood, Dudchenko, & Eichenbaum, 1999). Finally, in another task, rats were presented with two "items" (odor cups), X and Y, in each of two spatial contexts, A and B. In context A, item X was rewarded, while in context B, item Y was rewarded, regardless of the position within the context where the items were presented. At the beginning of learning, most cells functioned as place cells, encoding one

position within one context. As the rat learned the associations, however, many of these cells became "item-position" conjunctive cells. Interestingly, even before the correlates of hippocampal firing had been worked out, Hirsh first proposed that the primary role of the hippocampus would be to provide the context that would allow an animal to learn precisely this type of task (Hirsh, 1974).

Physiology and Function in the Different Hippocampal Subregions

So far, we have discussed "the hippocampus" as if it were a single entity. It is, however, made up of several different subregions, each of which has special properties that distinguish it from the others. In the majority of the studies described above, cells were recorded only from CA1, or cells were recorded from both CA1 and CA3 but differences were not reported. The results described above, then, might better be seen as describing the function of (dorsal) CA1 than of the hippocampus as a whole. We will therefore briefly review some of the properties that make each subregion unique.

Though the anatomy of the hippocampus is more complex than a simple loop of feed-forward information processing, the classical model of information flow through the hippocampus is the "tri-synaptic loop." In this circuit diagram, information from the entorhinal cortex (EC) is passed to the dentate gyrus (DG) via the perforant path. It flows from the dentate gyrus to CA3 via mossy fibers, and finally, from CA3 to CA1 via the Schaffer collaterals. In reality, the EC has direct connections not only with DG, but also with CA3 (via the performant path), CA1 (via the temporo-ammonic path), and CA2 (whether this connection is part of the temporo-ammonic path or is separate has not yet been defined).

In the tri-synaptic loop, the dentate gyrus is the first hippocampal subregion to receive input from the EC. The DG has three properties that are exceptional relative to most other brain regions (reviewed in Aimone et al., 2014; Piatti et al., 2013). First, the dentate has many more

cells than either its primary input region (EC) or its primary output region, yet it is estimated that only 1-2% of DG granule cells are active in any given environment due to a continual, strong feedback inhibition mechanism. Thus the dentate is said to have an extremely sparse code. Second, because of the extremely sparse code, the DG cannot rely on convergent input to affect its downstream targets, so a single bursting granule cell can bring a downstream CA3 cell to spiking threshold. Finally, the DG is one of only two brain regions that undergoes neurogenesis in adulthood. The function of the newborn neurons on the DG network remains poorly understood, but behavioral studies support the idea that it is important for learning and memory. These unique features of the dentate gyrus make it well situated to perform a role in "pattern separation." That is, the representations of two similar events are typically less similar to each other than the similarity found in the events themselves.

The hallmark feature of CA3, the DG's downstream target, is its richly recurrent connections. In fact, the largest number of synapses onto CA3 cells comes from other CA3 cells (Treves & Rolls, 1994). This allows CA3 to function as an auto-associative network, which has been proposed to allow CA3 to perform the complementary function to pattern separation, pattern completion (Treves & Rolls, 1994). Pattern completion would mean that a partial cue set or degraded input could nonetheless reinstate the complete memory trace. If a portion of a memory trace is activated, the recurrent connections between those cells and other cells that belong in the representation could excite those cells and bring them into the active set of cells. In addition to pattern completion, the other major role for CA3 proposed by Treves and Rolls based on its anatomy is to bind together "snapshots" of memory (Treves & Rolls, 1994), making it important for single-trial learning. The powerful mossy-fiber to CA3 connection could force a new firing pattern on CA3 in novel circumstances, and the recurrent connections could form that pattern into an attractor state. Thus, alongside CA3's pattern completion function, CA3 could also perform pattern separation, creating orthogonal representations of new experiences. In a number of

experimental studies, both involving knock-down of hippocampal function and recording from CA3 cell populations, these predictions have borne out. In mismatch trials, where subsets of cues are varied from a learned configuration, the CA3 population coherently follows one set of cues, while CA1 is likelier to respond heterogeneously or by remapping entirely (J. J. Knierim, Lee, & Hargreaves, 2006; Nakazawa et al., 2002; Vazdarjanova & Guzowski, 2004). Previously we have seen that if a subset of cues is removed (as opposed to put in conflict with other cues), CA1 is still able to follow the remaining cues (J O'Keefe & Conway, 1978). If NMDA receptors in CA3 are deleted, however, CA1 functions normally in well-trained environments, but fails to follow a subset of cues, indicating that this ability was dependent on CA3's role in pattern completion (Nakazawa et al., 2002). Similarly, deletion of NMDA receptors in CA3 impairs behavioral performance on pattern-completion dependent memory tasks (Nakashiba, Young, McHugh, Buhl, & Tonegawa, 2008; Nakazawa et al., 2002). If an animal has been trained in one environment and is then exposed to a novel environment, the cells that are involved in the new representation have low overlap with the cells that were part of the familiar representation (S. Leutgeb, Leutgeb, Treves, Moser, & Moser, 2004; Vazdarjanova & Guzowski, 2004; I. a Wilson, Ikonen, Gallagher, Eichenbaum, & Tanila, 2005). Supporting the idea that CA3 is important for single trial learning, it has been shown that CA3 stabilizes its representations in a novel environment to a greater extent on day 1 than CA1 does (J. J. Knierim et al., 2006; I. Lee, Rao, & Knierim, 2004), and impairing CA3 function causes spatial tuning in CA1 to develop more slowly and impairs single-trial learning (Nakashiba et al., 2008; Nakazawa et al., 2003).

CA1 firing properties have largely been described in the previous section, and CA1 may be defined largely by the ways in which it is not like DG or CA3. CA1 has neither neurogenesis nor (many) recurrent connections. Representations in CA1 are less sparse than in DG or CA3. Rather than performing pattern completion and pattern separation, CA1 tends to have a more linear relationship to external stimuli: small changes to the environment typically incur small

changes to the CA1 response, while larger changes to the environment incur larger changes to the CA1 response (Guzowski, Knierim, & Moser, 2004). As such, CA1 likely serves as a comparator between stored representations in CA3 and the current situation, enabling CA1 to detect novelty and update its representations of the world at a more gradual pace (I. Lee, Yoganarasimha, Rao, & Knierim, 2004).

Nestled between CA3 and CA1 there is CA2, a small subregion of the hippocampus that is not part of the tri-synaptic loop and has been largely overlooked in studies of hippocampal function. Originally defined as a subregion of the hippocampus by Lorente de Nó (1934), CA2 has often been considered as a transition zone between CA1 and CA3, likely because the region is small and the borders difficult to define by eye. Yet, it is apparent that CA2 differs from other regions its anatomical connections (R Bartesaghi, Migliore, & Gessi, 2006; Renata Bartesaghi & Gessi, 2004; Chevaleyre & Siegelbaum, 2010; Maglóczky, Acsády, & Freund, 1994; Mercer, Trigg, & Thomson, 2007; Pan & McNaughton, 2004), gene expression profiles (Lein, Callaway, Albright, & Gage, 2005; X. Zhao et al., 2001), basic physiological properties (Simons, Escobedo, Yasuda, & Dudek, 2009; M. Zhao, Choi, Obrietan, & Dudek, 2007), synaptic plasticity (Caruana, Alexander, & Dudek, 2012; S. Lee et al., 2010; Simons, Caruana, Zhao, & Dudek, 2012; M. Zhao et al., 2007), and susceptibility to disease state (Benes, Kwok, Vincent, & Todtenkopf, 1998; Sloviter, 1983).

As improved genetic-based tools for tracing and manipulating circuits have been developed, interest in CA2 has grown. New studies have worked to improve our understanding of the afferent and efferent connections with CA2, though consensus has not been reached; connections observed in one study or another were found to be lacking in other studies (Cui, Gerfen, & Young, 2013; Kohara et al., 2014; Rowland et al., 2013). There is agreement, however, that CA2 receives input from entorhinal cortex, layer II and from CA3 (though this may be net inhibitory in nature), and CA2 projects to CA1.

What do the differences in CA2 mean for behavior? Does CA2 express place fields, and how do they compare to DG, CA3, and CA1? These questions are still being worked out. Vasopressin 1b receptors (Avpr1b) are selectively enriched in the CA2 region of the brain, so knocking them out has been used as a first pass approximation to a CA2 knockout. Avpr1b knockout mice have reduced aggressive behavior, are deficient in social recognition and novelty tasks, and are impaired on temporal tasks, including remembering odor identity across a 10-second delay period and the *what-when* portion of the *what-where-when* novel object preference task described above (DeVito et al., 2009; Wersinger, Ginns, O'Carroll, Lolait, & Young, 2002). Though Avpr1b knockout is not equivalent to silencing CA2, some of these results have been replicated to indicate that CA2 is directly related. Specifically, by replacing Avpr1b to CA2 locally in a total Avpr1b knockout, aggressive responses were restored (Pagani et al., 2014), and a mouse in which CA2 itself was silenced genetically showed social memory impairment but no impairment in spatial or contextual memory (Hitti, 2014). As for place fields, CA2 does have them (Martig & Mizumori, 2011), but they have not yet been characterized in depth. Experiments that describe the firing patterns of CA2 cells are presented in chapter 2 of this dissertation.

Time in the Hippocampus, on multiple scales

Although the primary coding modality of the hippocampus is spatial, many additional attributes of episodic memory have now been attributed to coding within the hippocampus, including long-term stability of learned maps (Thompson & Best, 1990), representations of "*what*" (including odor-selective cells, and cells that respond to reward locations), representations of "*where*" (including precise spatial location and also spatial context), and conjunctive representations of learned associations. Does the hippocampus also somehow represent "*when*"?

As described in previous sections, time passes on multiple scales, and "when" can be defined in many ways. These may include being able to remember or reproduce a sample interval,

keeping sequences in their correct temporal order, associating an event to a particular time of day, knowing how long ago something happened or whether two events happened at similar times, disambiguating between similar events that happened at different times, or even "mental time travel." For this section, I will present evidence that the hippocampus is involved in representing temporal information at several time scales, irrespective of which type of timing information is involved.

The theta rhythm is an approximately 8 Hz oscillation in the local field potential of the hippocampus (and many surrounding areas) and is prominent during exploratory behaviors in rats. Interneurons in the hippocampus are generally phase-locked to the theta rhythm (Ranck Jr, 1973), but principal cells display a much less pronounced phase preference. This is due to a phenomenon that has been called theta phase precession. Principal cells typically fire in bursts as an animal passes through a place field, but the frequency of the bursts slightly outpaces the frequency of the theta rhythm. Consequently, each burst occurs at a somewhat earlier phase of the theta rhythm than the previous burst (J O'Keefe & Recce, 1993). Strikingly, the initial phase of theta is relatively constant across passes through the field and across cells; spikes may precess almost an entire theta cycle but they virtually never exceed a cycle (Maurer, Vanrhoads, Sutherland, Lipa, & McNaughton, 2005). Phase precession has an important consequence for rats running on a linear track. Suppose two place fields overlap such that field A occurs somewhat before field B. When the rat enters field A, cell A will begin bursting. By the time the rat enters field B, the bursts of cell A will have precessed relative to their starting point, and hence relative to the bursts of field B. Thus, even though the fields overlap, the order of the two fields is maintained in their firing representations within a single theta cycle, and a second or so of real-time activity (passing through multiple place fields) has been compressed and represented on the millisecond time scale (W. Skaggs & McNaughton, 1996). These compressed sequences have been seen to be replayed during sleep, potentially as a manner of consolidating the memory (M.

Wilson & McNaughton, 1994), and during waking, potentially as a way of playing out and deciding between different behavioral choices (Singer, Carr, Karlsson, & Frank, 2013). Though studies of this sort have focused on place fields on tracks or mazes with fixed trajectories, this has been proposed as a manner by which the hippocampus could chain sequences of events together and maintain the temporal order of events in its representations (Buzsáki & Moser, 2013).

Another possible substrate of maintaining temporal organization on the order of seconds is a sequence of "time cells." Pastalkova found that if rats ran on a treadmill while on the center stem of a T-maze, hippocampal CA1 cells would fire in a reliable sequence while the rats were running, despite the fact that the animal's spatial location was not changing, and the particular sequence that was activated depended on what trajectory the rat was about to take (Pastalkova et al., 2008). These sequences only developed when there was a memory demand (which direction to turn on the maze), suggesting that they were important for maintaining the rat's plan during the delay period. Eichenbaum and colleagues performed several follow-up studies, noting that the firing of each cell from a sequence occurred at a specific time point within the delay. They thus named these cells "time cells," in analogy with "place cells" that fire in specific places within an environment (MacDonald, Lepage, Eden, & Eichenbaum, 2011).

The exact relationship between time cells and behavioral expression of timing is yet unexplored. If a delay period is changed (from one reliable length to another), time cells can "re-time" like place cells remap (MacDonald et al., 2011). If the relationship between time and distance traveled on a treadmill is dissociated, some cells follow time while others follow distance, and many have a conjunctive representation (H Eichenbaum et al., 2014). There is no need for the animal to be running on a treadmill for time cells to emerge, however. In fact the original "time cell" paper examined firing during a delay period in a go-no go task (MacDonald et al., 2011), and time cells may even be observed in head fixed animals performing a delayed match to sample odor task (MacDonald, Carrow, Place, & Eichenbaum, 2013). Are time cells

used to "bridge the gap" over delay periods, or are they actually performing timing? For example, could these cells play a critical role in interval timing procedures? Or might they be important for remembering the order of information? Further study is needed, but time cells could be a plausible mechanism for measuring time on the scale of seconds.

On longer time scales, having increasingly many time cells that fire at a precise moment after delays of several minutes seems less plausible, although it is known that the spatial resolution of place cells increases along the dorsal-ventral axis of the hippocampus, so perhaps there are time cells that fire over much longer periods in the ventral region of the hippocampus as well. Regardless, time cells always fire with respect to a clear start time, and so far have only been shown for well-trained intervals.

For longer time scales and without a clearly defined period, the idea of a gradually drifting temporal context may be more suitable for providing temporal information. Indeed, work from Eichenbaum's lab indicates just that. Rats were presented sequences of 5 odors with 10-second delays between odor presentations. Fifteen seconds after the end of the sequence, they were given access to two odor cups from the sequence and were rewarded for digging in the cup that had been presented earlier. During several trials of this task, cells were recorded from CA1, and the similarity of the population response for pairs of odors was plotted as a function of the number of trials between them. In line with models of a time-varying context, the similarity decreased as a function of time. Of particular note, when examining error trials, the population similarity did not differ as a function of time. This suggests that the changes invoked by temporal context were important for reading out the temporal order of the two events (Manns, Howard, & Eichenbaum, 2007). Additional evidence that the temporal context model may function at the time scale of minutes comes from recordings in human patients who have had electrodes placed in their brains temporarily for clinical purposes. When performing free recall of lists of words, there is evidence in both single unit recording and also LFP that the previous context is reinstated

when a word is recalled (Howard, Viskontas, Shankar, & Fried, 2012; Manning, Polyn, Baltuch, Litt, & Kahana, 2011).

Neurogenesis in the dentate gyrus has been proposed to be important for creating temporal context in memories on much longer time scales (Aimone, Wiles, & Gage, 2006). Newborn neurons mature over the course of several weeks and are much more excitable than their mature counterparts during this period. The theory, then, is that they would be likely to respond to many episodes during their maturation period, while older neurons would respond only when an association to a stimulus already existed. Two episodes that occurred in close temporal proximity would share substantial overlap in their representations, even if the episodes were quite different, because they would share the population of granule cells that were young and excitable at the time the episodes took place. Conversely, if many weeks elapsed between two events, their DG representations would show minimal overlap because the population of young granule cells would have turned over in the time between the events.

For even longer time scales, it may be instructive to return to patient H.M., and other amnesics, who show a temporally graded retrograde amnesia for events prior to their brain injuries. That is, events for up to a few years immediately prior to the injury were completely forgotten, while events from long before were still remembered vividly. These findings led Squire to propose that after some time, episodic memories become consolidated to cortical representations and become independent of the hippocampus (Alvarez & Squire, 1994; Squire, 1992), and consolidating has been suggested to be a gradual process (McClelland, McNaughton, & O'Reilly, 1995; O'Reilly & Norman, 2002). A speculative idea, then, is that the degree to which a memory has been consolidated could serve as an assay of its recency on scales up to many years.

It is interesting to note that the various proposals for temporal coding come from different experimental paradigms, different time scales, and feature different brain regions as the primary

source of the temporal organization. Nonetheless, they share one feature in common. Specifically, in each case, the temporal signal is based on the cell set that is active in encoding an event, and the overlap of these cell sets decays over some time scale, which determines the time scale over which the coding mechanism may be useful for estimating elapsed time. The variation in the active cell sets also ensure that even identical events that occur at dissimilar times have non-identical representations. These themes are in-line with the stimulus fluctuation model and other models of temporal context described previously and will be seen again throughout the remaining chapters of the dissertation.

Preview of the Dissertation

We have seen that animals are capable of performing a wide variety of timing tasks, including tasks that require associating temporal information with other types of information such as object identity or spatial location, and many of these tasks are performed on time scales between hours and days. Many models suppose that time can be estimated via a time-varying context signal, and the hippocampus is required for forming contextual associations. Yet there has been very little research on hippocampal representations over hours and days. In this dissertation, we seek to fill that gap.

We followed the long tradition of using random foraging trials to assay hippocampal representations of place, and considered whether any type of remapping would occur as time passed, thus imparting temporal context. We had rats run in an arena with flexible wall four times each morning and four times each afternoon. At each time point, the box was shaped as a square for two sessions and as a circle for the other two. This allowed us to consider the interplay of spatial and temporal context. During these sessions we recorded extracellular spiking of principal cells from CA1, CA2, and/or CA3, and we followed the same cell sets across multiple days, such

that we could observe hippocampal representations for events that occurred anywhere between 10 minutes and 60 hours apart in time.

In chapter 1, we will look at the population dynamics of CA1 and CA3. Consistent with many previous studies, we find that the cells in each of these regions undergo strong rate remapping between box shapes. The two regions differ in their response to time, however. Consistent with its role as a pattern-completer and site of long-term stable memory, CA3 showed virtually no remapping over time. In CA1, however, we found that place fields sometimes turned on and off, indicating a temporary rate remapping effect, quite reminiscent of the field fluctuating between available and unavailable sets (Estes, 1955). The result is that the population vector correlation between representations of two identical boxes is a monotonically decreasing function of the time between them. Though we expected to find some effect of the circadian oscillator on CA1 representations, we found no rebound in similarity at 24 hours in the population code, nor did we observe any cells whose firing rates were reliably predicted by the time of day. This is consistent with empirical data described above that, when the stakes are low, rats use how-long-ago instead of time-of-day as their primary temporal cue.

Chapter 2 expands the experimental paradigm of the previous chapter to recordings in CA2. We first characterize many basic firing properties of this region, as that had not been systematically done before. One critical observation from this endeavor is that cells in CA2 have a much higher propensity to display multiple place fields in a single environment than cells in CA1 or CA3. Looking over temporal intervals of hours to days, we find that place fields in CA2 undergo dramatic rate remapping, turning on and off with greater frequency than the cells in CA1. Because CA2 cells tend to express multiple fields, any one of which may be on or off at any time point, this gives the impression that CA2 maps are highly unstable. This results in a population vector correlation that decays much more rapidly than in CA1, reaching a floor value after 18-24 hours. Although maps in CA2 appear unstable over time, they are surprisingly stable

between the square and circle boxes. We suggest that on time scales of hours to days, CA2 fills the role of the temporal context signal proposed in the computational models, while CA3 provides reliable content information in the form of a stable spatial context signal. Finally, CA1 codes conjunctively and simultaneously represents spatial and temporal information.

In chapter 3 we provide preliminary results from computational modeling studies. The goals for this study were two-fold: first, we desired to describe succinctly the session-by-session changes in place field firing that can account for the patterns of population vector correlations observed in CA2 and CA3. Next, we wished to assess to what extent CA1 may passively inherit spatial and temporal context discrimination from these regions.

Finally, chapter 4 will summarize these findings provide some thoughts on future directions this research may take.

CHAPTER 1: NEURONAL CODE FOR EXTENDED TIME IN THE HIPPOCAMPUS

Abstract

The time when an event occurs can become part of autobiographical memories. In brain structures that support such memories, a neural code should exist that represents when or how long ago events occurred. Here we describe a neuronal coding mechanism in hippocampus that can be used to represent the recency of an experience over intervals of hours to days. When the same event is repeated after such time periods, the activity patterns of hippocampal CA1 cell populations progressively differ with increasing temporal distances. Coding for space and context is nonetheless preserved. Compared to CA1, the firing patterns of hippocampal CA3 cell populations are highly reproducible, irrespective of the time interval, and thus provide a stable memory code over time. Therefore, the neuronal activity patterns in CA1 but not CA3 include a code that can be used to distinguish between time intervals on an extended scale, consistent with behavioral studies showing that the CA1 area is selectively required for temporal coding over such periods.

Introduction

Autobiographical memories include a detailed record about the place where events occurred and about when the events took place(Friedman, 1993; O'Keefe, J, Burgess, N. Donnett, JG, Jeffery, KJ, Maguire, 1998; Tulving, 1972, 1983). The temporal aspect of these memories consists of detailed sequences of events as they occurred in real-time, but many memories are also part of an extended timeline that relates the different episodes to each other, such as in a storyboard for film making. The hippocampal system is required for remembering events over time intervals that exceed the capacity of short-term or immediate memories (Zola-Morgan & Squire, 1993). Because these memories include temporal information over extended time scales, hippocampal neuronal activity should generate a code that can represent when or how long ago an

event occurred within the remembered time range. However, hippocampal coding for time has only been described on a scale up to minutes. Temporal coding over these shorter intervals is thought to occur by the sequential or cumulative activation of cell populations or by the activation of delay-dependent cells(Gill, Mizumori, & Smith, 2011; Howard et al., 2012; MacDonald et al., 2011; Manns et al., 2007; Naya & Suzuki, 2011; Pastalkova et al., 2008). These mechanisms depend on ongoing activity patterns, which can be interrupted by unrelated neuronal firing patterns or by a change in brain state, and thus are not amenable to coding over longer time scales. It is known that humans and other animal species can remember temporal relations on the time scale of hours and days(Babb & Crystal, 2005, 2006a, 2006b; Clayton & Dickinson, 1998; Clayton, Yu, & Dickinson, 2001; Iordanova et al., 2008; Kesner & Hunsaker, 2010; O'Brien & Sutherland, 2007; W. a Roberts et al., 2008), and theoretical considerations have suggested that the coding of such extended temporal relations may emerge from fluctuations within the memory network (Estes, 1955; Gallistel, 1990; Howard & Kahana, 2002; G.-J. Mensink & Raaijmakers, 1988; Murdock, 1997). Nonetheless, a neural code for remembering when or how recently an episode occurred within the last hours or days has not yet been identified experimentally.

Beyond the time scale during which sustained firing patterns of hippocampal neurons can maintain a continuous representation of events, each new memory is thought to be stored as a change in synaptic strengths resulting from neuronal firing patterns at the time of encoding (Hebb, 1949; B. McNaughton & Morris, 1987). Later retrieval would be achieved by reinstating the corresponding neuronal activity from the stored pattern of synaptic strengths (Liu et al., 2012; Reijmers, Perkins, Matsuo, & Mayford, 2007). Under such a coding scheme, the activation of neuronal firing patterns for familiar, repeated events would be expected to be highly similar. If the reinstated firing patterns were exactly the same during each repetition, the event would be represented with high fidelity, but the neuronal code would not contain any information about when the events occurred or about how much time had passed between the events. Alternatively,

random fluctuations or a continuous change in neuronal firing patterns could result in a somewhat distinct neuronal firing pattern for each occurrence of an otherwise highly similar event and thereby generate information about the temporal distances or the temporal context of events (Estes, 1955; Gallistel, 1990; Howard & Kahana, 2002; G.-J. Mensink & Raaijmakers, 1988; Murdock, 1997). We therefore asked whether these proposed neural coding schemes could be observed in hippocampal neural networks for events that were identical except for the time when they occurred.

If differences in neuronal activation emerge over extended time periods, this might result in a loss of detailed information about past events. We therefore also asked how encoding of time might be complementary to retaining precise representations for other types of contextual information. To address these questions, we used a behavioral paradigm in which the firing patterns of the same hippocampal cell populations were repeatedly measured in the same behavioral apparatus over extended time intervals. To examine if coding differences over time might result in a loss of precise information about other aspects of the context, we tested, at each time point, to what extent the firing patterns were able to code for differences between enclosures that only differed in their shape.

Results

We trained six rats to randomly forage in a square and in a circular enclosure in the morning (AM) and, after an interval of 6 hours, in the afternoon (PM) of each day. The training at each time of the day consisted of four 10-minute sessions, two in each enclosure shape, presented in random order (Figure 1.1A). The training schedule remained consistent over many days. Rats were therefore highly familiar with the behavioral paradigm when recordings began such that any observed differences in firing patterns could not be attributed to novelty-related effects.

Because it was critical for all aspects of the data analysis that each recorded cell could be identified accurately throughout the entire sequence of sessions, we used a technique for electrode placement that was optimized for recording stability across time, and we applied stringent standards to the acceptance of cells for further analysis (see Materials and Methods and Figure 1S.1). In particular, we flanked the recordings in behavior with recording periods during rest. Most hippocampal cells fire bursts of action potentials during rest periods (Olmstead, Best, & Mays, 1973; Ranck Jr, 1973), and these recordings can thus be used to confirm that the signals of each cell could be reliably recorded before, after, and throughout the entire experimental paradigm, even if the cell was not active during all periods of random foraging.

In the dataset with confirmed cell identity within an entire recording day (n = 99 CA1 cells and n = 126 CA3 cells), we found in both the hippocampal CA1 and CA3 subregions that repetitions of the same enclosure shape resulted in highly reproducible firing of hippocampal cells at well-defined spatial locations, as reported consistently since hippocampal place fields were first described (Muller et al., 1987; J O'Keefe, 1976; J O'Keefe & Dostrovsky, 1971; Thompson & Best, 1990). However, when considering the network activity pattern of the CA1 cell population, we found that the similarity in neuronal activity decreased for the 6-hour interval compared to intervals of less than 1 hour (t = -7.71, $P < 0.001$), indicating that events that are farther separated in time have more distinct neuronal representations (Figure 1.1B-C). In contrast, the similarity of firing patterns in the CA3 cell population decreased to a much smaller extent between the < 1-hour interval and the 6-hour interval (t = -2.95, $P < 0.05$ for the comparison between time intervals; t = -9.92, $P < 0.001$ for the comparison between CA1 and CA3).

Differences between hippocampal subregions were also observed for the firing rates of single cells such that subsets of CA1 but not CA3 cells showed a high variability between repetitions of the same enclosure shape in the morning and afternoon. Thus we find that CA3 maintains highly similar representations for repeated events, irrespective of the elapsed time
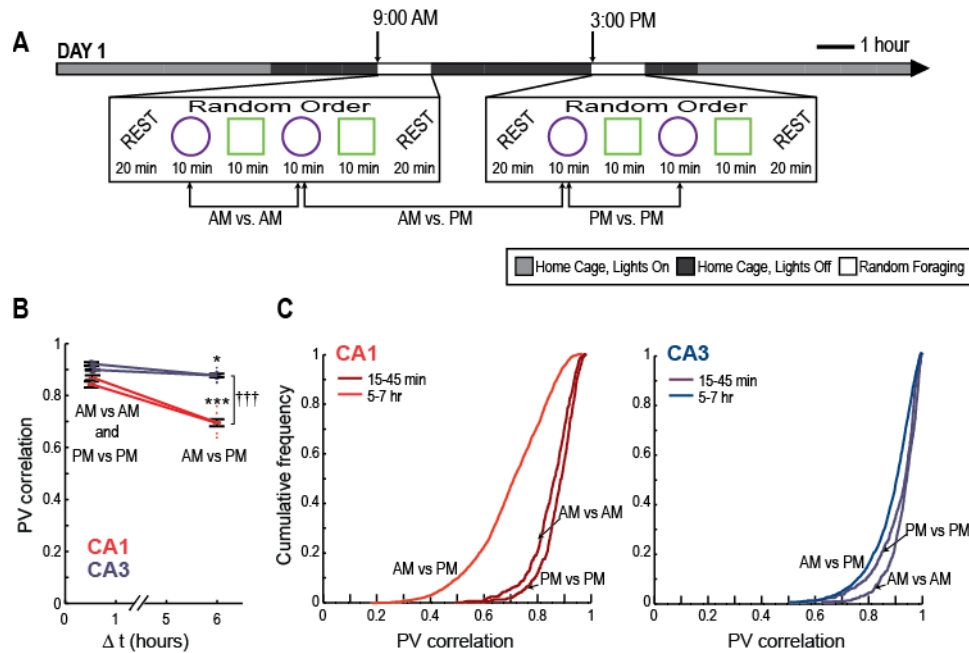
Figure 1.1. Decorrelated neuronal firing patterns between morning and afternoon sessions in hippocampal cell populations. (A) Behavioral design and experimental timeline. A series of four random-foraging sessions was conducted in the morning (AM), and a second series of four sessions was conducted in the afternoon (PM). Each series consisted of a random sequence of two sessions in a square and of two sessions in a circular enclosure. The enclosure was placed in the same location within the room, and extra-maze cues remained constant. Recordings began after 9-26 days of pretraining in this paradigm. (B) The firing patterns of simultaneously recorded CA1 and CA3 neuronal populations were compared across repetitions of the same enclosure shape [as illustrated for circle comparisons in (A)]. A population vector (PV) represents the firing rates of all active cells within a 5 cm by 5 cm pixel in the spatial map. The activity patterns of an active cell population were compared between sessions by calculating the correlation coefficients of PVs from corresponding pixels between two sessions and by taking the average over all pixels. The values for pair-wise comparisons at short time intervals (within AM or within PM) and at long time intervals (AM vs. PM) were grouped. Each of the comparisons is shown as a red or blue dot, and black error bars report the mean +/- SEM of between-session comparisons at each time interval. CA1 cell populations (red) showed a decreased correlation between sessions that occurred at 6-hour intervals compared to sessions within < 1 hour. PVs of CA3 cells (blue) also showed a small degree of decorrelation between time intervals, but the degree of decorrelation was substantially lower in CA3 compared to CA1. (C) Cumulative distribution functions for PV correlations between pairs of recordings in the same enclosure shape at different time intervals. * $p < 0.05$; *** $p < 0.001$ for comparisons between time intervals; ††† $p < 0.001$ for comparisons between CA1 and CA3; see text for statistics.

between them, whereas CA1 representations vary to a larger extent between repeated events that are separated over intervals of several hours [F(103,49) = 1.98, $P < 0.01$] (Figure 1.2).

Next, we asked whether differences in hippocampal firing patterns were circadian. Pronounced fluctuations in hippocampal gene expression patterns, neurotransmitter release, and in hippocampal synaptic excitability occur with the circadian clock (Barnes, McNaughton, Goddard, Douglas, & Adamec, 1977; Hut & Van der Zee, 2010; Tononi & Cirelli, 2001, 2006; Valnegri et al., 2011). Any differences in the population code for an identical environment after a time interval of six hours might therefore reflect a circadian or ultradian modulation of activity patterns. If hippocampal neural network activity were controlled by internal states associated with circadian or other endogenous clocks, one would expect to find a cyclical pattern in neuronal firing such that the similarity at matching times of the day is higher than at non-matching times within a day. To examine this possibility, we identified the hippocampal place cells that could be reliably tracked through the morning and afternoon of two consecutive recording days (n = 50 CA1 cells and n = 71 CA3 cells). We first looked for circadian fluctuations in normalized firing rates. Consistent with a previous study (Munn & Bilkey, 2012), we found no differences across time in the average normalized firing rate [F(15) = 0.90 for CA1 and 0.32 for CA3; n.s.] (Figure 1.3A-B). We then asked whether the differences in neuronal activity patterns in the hippocampal CA1 area, when comparing sessions recorded 24 hours apart and thus matched for time of day, showed any increase in similarity compared to those at 6-hour time intervals, as would be expected for a circadian cycle. Instead, the recordings after 24 hours exhibited a further decrease in similarity compared to the 6-hour time point, such that the similarity in hippocampal firing in CA1 monotonically decreased for time intervals between 30 minutes up to at least 30 hours [F(4) = 46.3, $P < 0.001$] (Figure 1.3C-D; Figure 1S.2). The network activity in CA3 also showed a gradual decrease in similarity with time [F(4) = 50.5, $P < 0.001$], although to a much smaller extent than the decrease in CA1. The correlation at 30 hours is 0.80 +/- 0.003 in CA3 compared
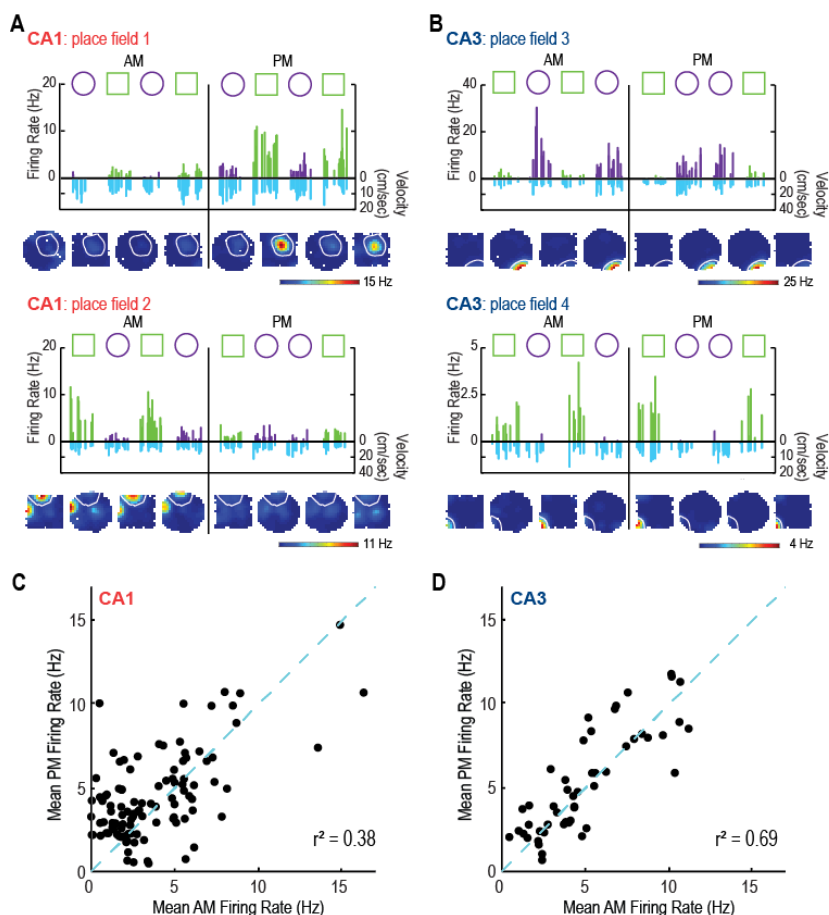
Figure 1.2. CA1 place fields show variability in firing rate between morning and afternoon sessions. (A) Firing rates for two representative CA1 place fields and (B) for two CA3 place fields. Fields 2, 3, and 4 were recorded simultaneously. Each 10-minute recording session throughout the morning (AM) and afternoon (PM) is shown. Symbols above each graph indicate the order of enclosure shapes. Each bar (green for the square shape and purple for the circular shape) represents the firing rate of the cell during a pass of the animal through the place field (see Figure 1S.6 for methods and Figure 1S.7 for additional examples). For each pass, the corresponding running speed of the animal is plotted downwards below the x-axis in cm/s (in blue). The variability in firing rates can neither be explained by movement velocity nor by the proximity of the path to the field center (Figure 1S.8). For each firing field the corresponding color-coded rate maps (averaged across each 10-minute recording session) are shown below the line graph. The color scale for rate maps is from 0 Hz (blue) to the peak rate of the day (red). Coding differences emerged, in part, from CA1 place fields that changed or became silent at a subset of time points. In all cases in which cells became silent, it was verified in preceding or subsequent rest sessions that spikes from these cells could be detected (Figure 1S.1). (C) For each place field, mean rates within each enclosure shape were calculated in the morning and in the afternoon. These rates were compared within each hippocampal subregion. Place fields that were not active (mean rate < 2 Hz) at either time point were excluded. For active cells, the firing rates between morning and afternoon were more variable in CA1 compared to CA3 (see text for statistics and Supplementary Methods).
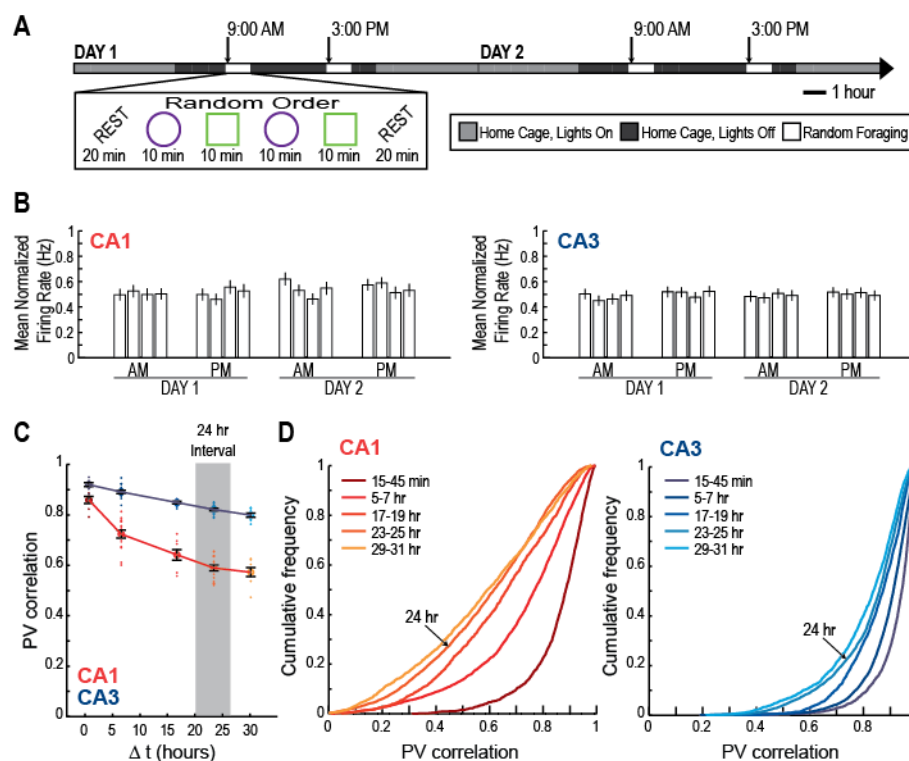
Figure 1.3. The decorrelation in CA1 network activity over extended time periods does not repeat cyclically across days. (A) To determine whether differences in CA1 network activity patterns can be explained as a circadian effect, we extended the hippocampal recordings across two days. (B) The mean +/- SEM normalized firing rate for all active CA1 (left) and CA3 (right) principal neurons is shown for each ten-minute recording session across the two day experiment. For each cell, the normalized firing rate was calculated by dividing the average firing rate for each session by that cell's maximum average firing rate in any of the sessions. A circadian variation in firing rate was not observed (see text for statistics). (C) Population vector (PV) correlations between pairs of recordings in the same enclosure shape are shown as dots. Every pair-wise comparison is aligned to its time interval, so that, for example, the AM-PM comparisons on day 1 and the AM-PM comparisons on day 2 are all aligned to the six-hour interval (see Figure 1S.9 for the complete pair-wise correlation matrix). The black error bars report the mean +/- SEM for pair-wise comparisons at each time interval. The correlation coefficients for the CA1 population activity (red) decreased monotonically as a function of elapsed time between recording sessions up to at least 30 hours (see text for statistics and Figure 1S.2 for comparisons of up to 60 hours). Repeated CA1 recordings at matching times of day on two consecutive days (24-hour interval) show a smaller correlation than recordings at shorter intervals but at different times of the day ($P < 0.001$ for the *posthoc* comparison). Thus, the effect we observed is not due to circadian fluctuations. (D) Cumulative distribution functions for PV correlations between pairs of recordings in the same enclosure shape at different time intervals.

to 0.57 +/- 0.006 in CA1 (t = -11.5, $P < 0.001$). Therefore, we do not find a distinct coding scheme in the CA1 or CA3 network for representing the exact time of day at which an event occurred, but rather find that the degree of similarity between two representations differs depending on the temporal distance between them, in particular within the CA1 cell population

If the decrease in similarity over time reflected merely an unconstrained random drift in firing patterns, such continuous change could eventually result in a decrease in the accuracy with which hippocampal neuronal firing patterns code for a repetition of the same event after a long time interval, and consequently, in a change of the ability to represent the relative similarity between contexts (J. K. Leutgeb, Leutgeb, Moser, & Moser, 2007; Muller & Kubie, 1987). We observed that, at any time point, there were many CA1 cells that discriminated between the two different enclosure shapes, but that individual CA1 cells showed inconsistency across extended time periods in their strength of preference for firing in one of the enclosure shapes (Figure 1.4A and B; see figure legend for statistics). In contrast, individual CA3 cells showed a striking consistency in enclosure shape coding by firing rate over extended time periods. Although individual CA1 cells showed high variability in preferentially firing in one of the enclosure shapes, the average degree of shape preference within the CA1 cell population did not change over time periods of 6 hours nor between days [Absolute shape preference: CA1, $F(3) = 0.28$, n.s., CA3, $F(3) = 0.33$, n.s.; Spatial correlation: CA1, $F(3) = 0.30$, n.s., CA3, $F(3) = 2.59$, n.s.] (Figure 1.4C and D). This pattern of coding in CA1 can result in consistent representations of contextual differences, even though the population vector correlations for repetitions of identical shapes were as low as 0.57 between days (see above and Figure 1.3).

Did the dissimilarity in the CA1 hippocampal population code emerge as a function of elapsed time, or could differences in neuronal firing patterns be the outcome of cumulative rate changes during the behavioral sessions such that larger differences were reached with increased random foraging experience? If experience in the recording enclosure were primarily responsible

for the reorganization of the network, then the changes in network activity would be expected to occur predominantly during running in the enclosure, the period while place cells are active and when activity-dependent network plasticity occurs (Mehta, Barnes, & McNaughton, 1997). To examine this possibility, we established several control conditions. First, we tested whether the amount of random foraging experience between two time points had an effect on the similarity of CA1 population coding. At the 24-hour and at the 30-hour interval, we found the same degree of decorrelation irrespective of the number of intervening four-session blocks (t = -0.64, n.s. for 24-hour intervals; t = -1.82, n.s. for 30-hour intervals) (Figure 1S.3A). Second, we observed decorrelation with time even when limiting our analysis to time intervals without any intervening AM or PM blocks [F(3) = 41.6, p values for all posthoc comparisons < 0.001 except that the comparisons between the 18 and 24 hours interval is n.s.] (Figure 1S.3B). Third, we asked whether the continued switching between two different enclosure shapes might explain why CA1 firing patterns accumulated firing differences across a series of recording sessions. To examine this possibility, we performed a two-day series of AM and PM recordings in only one enclosure shape (Figure 1.5A). We found that this paradigm did not change the low degree of decorrelation that we had observed over long time periods in CA3 (all mean PV correlations > 0.82; F(5) = 10.1, $P < 0.001$; all *posthoc* comparisons between time intervals ≥ 0.5 hrs are n.s. except $P <$ 0.001 for 6 hours compared to 24 hours) (Figure 1.5B and C). At the same time, we continued to observe a monotonic decrease in population similarity with extended time in the hippocampal CA1 cell population [F(5) = 54.0, $P < 0.001$; $P < 0.001$ for all *posthoc* comparisons, except $P <$ 0.01 for < 1 hour compared to 6 hours, and n.s. for 6 hours compared to 18 hours, 18 hours compared to 24 hours, and 24 hours compared to 30 hours], and the decrease over time remained larger in CA1 than in CA3 (at 30 hours, t = -9.0, $P < 0.001$). Finally, we examined whether a repetitive increase or a repetitive decrease in the firing rate of place fields within each 10-min
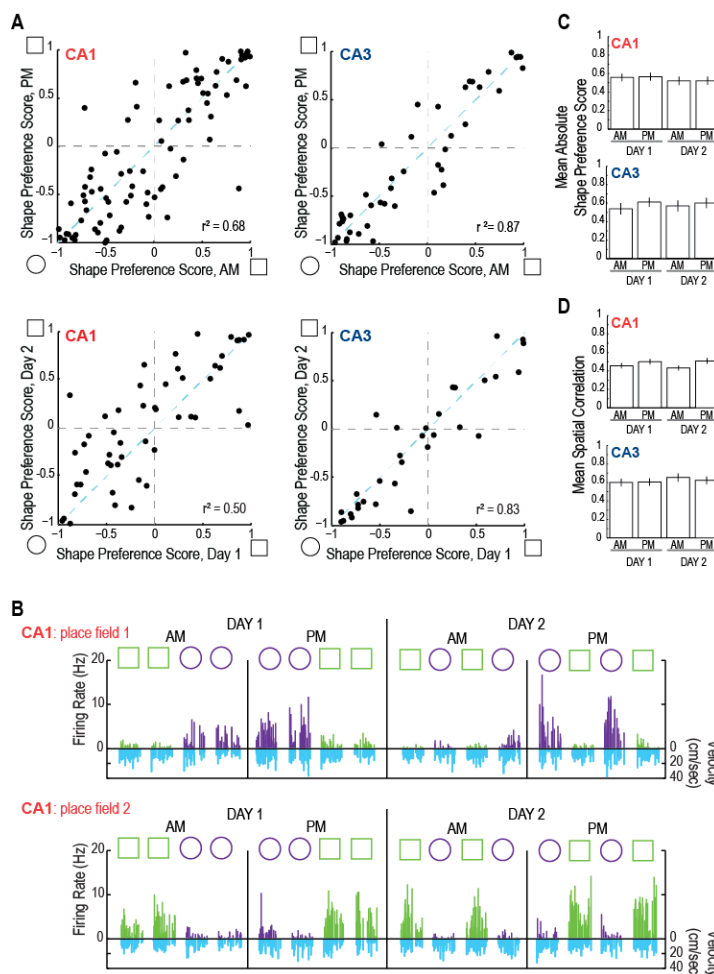
Figure 1.4. Differences in firing patterns over extended time periods did not preclude the encoding of spatial information or of contextual differences. (A) For each place field a shape preference score was calculated as a measure of the difference in firing rates between the circular and the square enclosure (scores of -1 or +1 indicate that the cell fired only in the circle or only in the square, respectively). This score was compared between recording blocks (left, CA1; right, CA3) in the morning and afternoon (top panel) and between days (bottom panel). For calculating each day's score, all AM and PM recording sessions within a day are used. Individual CA1 fields show more variable shape coding over 6-hour intervals and over one-day intervals than individual CA3 fields [$F_{(82,43)} = 2.51$ and $F_{(50,31)} = 2.55$, $P < 0.01$ for comparisons of shape preference scores at both time intervals]. (B) Firing rates within two representative CA1 place fields from the same cell over two days (data presented as described in Figure 1.2) (see Figure 1S.10 for additional examples). The CA1 place fields showed a change in the degree of discrimination between the square and circular enclosure between time points (see Figure 1.2A). Field 1 normally fired at its highest rate in the circular shape, but became silent in the AM session on day 2. The field resumed its firing in the circular enclosure during the PM session on day 2. It was therefore observed that shape preferences could be lost and regained between blocks of sessions at different time points. (C and D) Even though individual CA1 cells show coding differences between time points, the average degree of context and place coding is consistent within CA1 and within CA3 cell populations (see text for statistics).
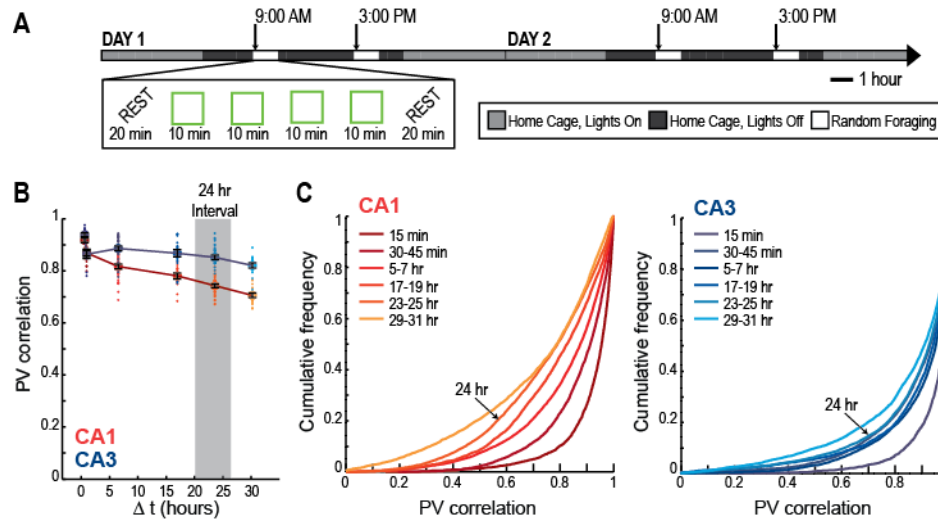
Figure 1.5. When testing with a single enclosure shape, firing patterns of the CA3 network remained highly consistent for repetitions of the same environment over extended time intervals while activity patterns in the CA1 network changed. (A) An experimental design with a single enclosure shape was used to test whether the decorrelation of hippocampal activity patterns could have been an effect of intervening experiences in a different context (see Figure 1.3). The mean PV correlation between pairs of recordings in the same enclosure shape (B) and the corresponding cumulative distribution function for the PVs (C) are shown as described in Figure 1.3. Highly consistent firing patterns in the CA3 population were observed over time intervals of 30 minutes to 30 hours. In contrast, the CA1 network continued to show a more pronounced monotonic decrease in firing similarity with time (see text for statistics).

session might contribute to progressive changes over long time periods. We observed that a large fraction of CA3 place fields (57.9 %) exhibited either a consistent decrease or a consistent increase in firing rate within each ten-minute behavioral session. However, consistent rate changes were uncommon in CA1 cells (11.1 %) (Figure 1S.4). Systematic changes in firing patterns within recording sessions – as expected for activity-dependent plasticity mechanisms – were thus primarily observed in the CA3 hippocampal subregion, which showed only minor differences in activity patterns over long time periods, and not in the CA1 subregion, where marked changes occurred across longer time scales.

Discussion

Theories of long-term memory coding require that stored firing patterns are accurately reinstated during later retrieval (Hebb, 1949; Liu et al., 2012; B. McNaughton & Morris, 1987; Reijmers et al., 2007). In contrast, theoretical considerations for representing temporal aspects in long-term memory require fluctuations of activity patterns in neuronal networks such that relative temporal distances or temporal order can be represented even when the events are otherwise identical(Estes, 1955; Gallistel, 1990; Howard & Kahana, 2002; G.-J. Mensink & Raaijmakers, 1988; Murdock, 1997). We provide evidence for both of these neuronal coding schemes within distinct subregions of hippocampus on a time scale of hours and days. In the CA1 cell population, the degree of similarity of neuronal responses for identical locations in the same context decreases monotonically as a function of the time between experiences for at least 30 hours. These changes in neuronal network activity were measured in highly familiar environments and thus appear unrelated to changes that have been reported during new learning (L. M. Frank, Stanley, & Brown, 2004; S. Leutgeb & Leutgeb, 2007; Lever et al., 2002; Munn & Bilkey, 2012). The fluctuations also seem not to correspond merely to noise, because it would be unlikely for a cell to stop firing for a time and then resume its prior firing field, firing rate, and

shape preference by chance. Finally, we could not detect a circadian component in the coding difference within the hippocampal CA1 cell populations. The increasing decorrelation with longer time intervals in CA1, but not in CA3, is consistent with behavioral studies showing that the hippocampal CA1 area is selectively required for temporal coding over extended time periods (Kesner & Hunsaker, 2010) and that rats are able to use "how long ago" but not "when in a day" as a cue for locating a food reward (Babb & Crystal, 2006a; W. a Roberts et al., 2008).

In contrast to CA1, CA3 cell populations showed highly reproducible firing patterns over extended time periods between repeated recordings. We therefore find a complementary neural code in CA3, which provides a highly stable representation of space and context without the possibility of contributing information about extended time. The stability of firing patterns in CA3 compared to the pronounced decorrelation in CA1 suggests that differences in neuronal activity over time are not an inevitable consequence within highly plastic hippocampal networks. To the contrary, CA3 generates highly similar firing patterns for repeated events over time intervals when changes in synaptic strengths are expected (Barnes et al., 1977; Hut & Van der Zee, 2010; Tononi & Cirelli, 2001, 2006). This function is consistent with the proposed role of the recurrent network architecture in CA3 for pattern completion (Marr, 1971; B. McNaughton & Morris, 1987; Rolls, 1989; Treves & Rolls, 1992). Although the typical definition of pattern completion proposes that a partial sensory input pattern is expanded to a pattern that is stored in memory, our finding of network stability for exact repetitions across long time intervals suggests that accurate neuronal firing patterns can be generated not only from degraded sensory inputs, but also when fluctuation or degradation in synaptic strength may emerge within neural circuits over extended time intervals. Ongoing changes in synaptic strength could emerge not only from random variability, but also by network reorganization over long time periods, as predicted by theories of consolidation and reconsolidation (S.-H. Wang & Morris, 2010). Our results indicate that hippocampal memory circuitry includes network mechanisms within CA3 that provide

consistency of neuronal representations despite fluctuations or circuit reorganization over extended time intervals within CA1 and, possibly, within a wider cortical network.

Even though neuronal firing patterns in CA3 remained consistent, the input from CA3 to CA1 did not result in equally consistent firing patterns in CA1. This is perhaps not unexpected because CA1 also receives major direct input from entorhinal cortex (Witter & Moser, 2006) and is thus the site for convergence of firing patterns from layer III of entorhinal cortex – which may include temporal information (Howard & Natu, 2005; Suh, Rivest, Nakashiba, Tominaga, & Tonegawa, 2011) and be critical for memory consolidation (Remondes & Schuman, 2004)– with the highly consistent firing patterns it receives from CA3. Even though the CA3 firing patterns are not completely transferred to CA1, they may nonetheless provide substantial firing stability to the CA1 cell population. This notion is supported by the finding that although each individual subregion became more decorrelated in the two-shape compared to the single shape condition, a difference in decorrelation over time occurred between CA3 and CA1 irrespective of the experimental condition (Figures 1.3, 1.5 and 1S.5). This further supports the idea that a neuronal code for temporal distances might emerge in CA1 by adding fluctuations to a more stable representation of other aspects of context it receives from CA3.

While temporal coding over periods of seconds and minutes can use delay dependent and sequential coding (Gill et al., 2011; MacDonald et al., 2011; Naya & Suzuki, 2011; Pastalkova et al., 2008), and while time stamps over weeks might be explained by long-term structural reorganization (Aimone et al., 2006), the mechanism that we describe can encompass temporal coding in the behaviorally relevant domain of hours and days. The observed pattern of neural activity rules out the possibility that the differences across time are due to a small subpopulation of hippocampal cells that is dedicated to temporal coding. Rather, differences in neuronal coding with extended time, along with sustained precision for other aspects of the context unrelated to time, are achieved by a high but balanced variability in the firing patterns of individual CA1 cells

such that the average context discrimination remains consistent. Even though such a network pattern may, at least in part, be generated by random variability and decay in synaptic strength, by intervening modifications of synaptic strength, or by random fluctionations in network patterns this would nonetheless correspond to theoretical considerations (Estes, 1955; Gallistel, 1990; Howard & Kahana, 2002; G.-J. Mensink & Raaijmakers, 1988; Murdock, 1997) that predict that temporal information can be retrieved from the resulting variability in the population code over time. Our findings are the first experimental evidence of a neuronal code that can be used for encoding temporal distances on a time scale of hours and days and that can co-occur with a precise neuronal code for other aspects of memory such that the when, what, and where aspects of memories can be simultaneously represented within the hippocampal CA1 cell population.

Materials and Methods

      Subjects. Six male Long Evans rats with a preoperative weight of 400-485 g were housed individually and maintained on a 12-h light/12-h dark schedule with lights off at 6:00 am. All behavioral testing occurred in the dark phase with one exception, as described in the behavioral procedures section. All experimental procedures were performed as approved by the Institutional Animal Care and Use Committee at the University of California, San Diego and according to National Institutes of Health and institutional guidelines.

      Surgical procedures. At the time of surgery, the rats were anesthetized with isoflurane (2-2.5% in O2) and an electrode assembly that consisted of fourteen independently movable tetrodes was implanted above the right hippocampus (AP 3.8-4.0 mm posterior to bregma, ML 3.0 mm). Tetrodes were prepared as described previously (J. K. Leutgeb et al., 2007), and we used a technique for electrode placement in the hippocampal cell layers that was optimized for recording stability across time.

Behavioral Procedures. After one week of recovery from surgery, rats were partially food-deprived and trained to forage for randomly scattered cereal crumbs in an enclosure with walls that could be shaped either as a square (80 cm by 80 cm) or as a 16-sided polygon (50 cm radius; referred to as a 'circle')(J. K. Leutgeb et al., 2007). A polarizing white cue card (20 cm wide) was placed on an inside wall of the enclosure. The center of the enclosure was always located at the same place in the room, and the angle of the cue card compared to external room cues was kept constant. Training was performed in two daily blocks. For all rats, the first block started between 8:30 and 9:30 am and the second block between 2:30 and 3:30 pm. For each individual rat, the daily start time of each block varied by less than 30 minutes. Rats were returned to the animal housing room between the morning and afternoon training blocks.

Rats were trained to run for four 10-minute sessions during each block, with two sessions in the square enclosure and two sessions in the circle enclosure. The order of the shapes varied randomly within each training block. The rats were allowed to rest for five minutes between sessions, and training blocks were flanked by sleep sessions (10-20 minutes before and after each block). The floor of the enclosure was cleaned with water between each session. Following the sleep session at the end of the afternoon training block, rats were screened for single unit activity. Electrophysiological recordings throughout the morning and afternoon sessions were initiated when multiple well-isolated cells (>300 µV) were observed on most tetrodes. The recording phase of the experiment began after 14 to 26 days of behavioral training, except in one rat in which the recordings commenced after 9 days.

Recordings were first conducted for two days in the training paradigm (referred to as 'Day 1' and 'Day 2' in the text). In three animals, recordings were performed on a third day. The third day was identical to the first two recording days, except that the start times of the blocks were shifted by 6 hours, so that the first block occurred at 3:00 PM and the second at 9:00 PM. The second block was thus during the light phase of the light-dark cycle. Two animals returned

for one day to the standard training conditions before two additional days of recordings were performed in which all four random foraging sessions in the morning and afternoon block were conducted in one enclosure shape ('Single Shape, Day 1' and 'Single Shape, Day 2'). This paradigm was identical to the standard paradigm, except that only one of the two enclosure shapes was used in all behavioral sessions throughout both days. One rat was tested in the square and the other in the circle enclosure. For each animal, we selected the shape in which we identified the larger number of active cells during the recording on the preceding day.

Cell-tracking. Because our study depended on being able to follow the same set of principal cells over an extended time period, we developed a customized version of MClust (Redish, n.d.) with added functions that allowed for the comparison of the cluster boundaries of each cell throughout a series of 10-minute recording sessions. Clusters that persisted in the same region of parameter space throughout a day (or multiple days) were accepted for further analysis. Care was taken to accept only cells that could be precisely followed from the beginning to the end of the data analysis (Figure 1S.1), and for which all spikes were included in the cluster boundary such that observed rate changes could not be attributed to the definition of the cluster boundaries.

Data analysis. For tracked cells, we calculated the spatial map of each 10-minute session and the spatial correlations between maps. For a more detailed analysis of firing rates within the place field, we also analyzed the firing during individual passes through the place field (Figure 1S.6). From firing rate measurements within the place field, we derived scores for differences between square and circular enclosures for different times throughout the experiment. For the entire population of CA1 cells and for the population of CA3 cells, we calculated population vector correlations between 10-minute sessions that were recorded at different times, but within the same enclosure shape.

Histology. Rats received an overdose of sodium pentobarbital and were perfused intracardially with saline and 4 % formaldehyde. The brains were extracted and stored in

formaldehyde. Frozen coronal sections (40 µm) were cut and stained with cresyl violet. Each section through the segments of the hippocampus with electrode tracts was collected for analysis. Final tetrode positions were determined by subsequent three-dimensional reconstruction of the tetrode array in serial sections. Recordings from a tetrode were included in the data analysis if the tetrode's deepest position was in the CA1 or CA3 pyramidal cell layer. The electrodes had not been moved after the recordings. Tetrodes with tips in or near the border of the CA2 region were excluded from the analysis.

Appendix 1.1: Supplemental figures

Figure 1S.1. Clusters of spikes from single cells were tracked across multiple days. (A) The behavioral design and the experimental timeline are shown for the two-day, single-shape experiment including rest and behavioral sessions. The sessions that are shown in (B) and (D) are highlighted in gray. (B) Each panel shows the projection of the peak amplitude recorded on one channel of a tetrode located in the CA1 pyramidal layer compared to the peak amplitude recorded on another channel of the same tetrode. The same projection is shown for each cluster diagram, and different diagrams correspond to different time points in the experiment. Each dot represents one sampled spike. Colored dots of a single color represent spikes that were assigned to a single cell. After defining clusters on a single day, the cluster boundaries were applied to data recorded from the same tetrode on subsequent days. If clusters persisted in the same region of parameter space where they had been identified on the previous day, boundaries were adjusted to assure that all spikes that were part of the cluster were included within the boundaries. If the spikes for a set of clusters on one tetrode could be included with minor adjustments of the boundaries from the day before, the clusters were considered to be the same set of cells as on the day before, and the cells were included in our analysis as trackable across days. Because we occasionally found cells that appeared to 'turn on' or 'turn off' during the day, it was imperative to confirm that these effects were not simply a consequence of a tetrode having moved into or out of the range where the cell could be recorded. We thus only included cells in the analysis that were active at the beginning of the day (in the first sleep session or the first behavioral session in either shape) and at the end of the day (during the last sleep session or the last behavioral session in either shape). The highlighted (blue) cluster depicts a cell that fired robustly during all behavioral sessions on the first day and in the morning of Day 2. It then fired very few spikes during the second afternoon (PM3-4, Day 2) before regaining robust activity during the PM Rest2 session, indicating that the small number of spikes recorded during the afternoon was not due to the cell having moved out of the assigned parameter range, but was rather due to a genuine decrease in firing rate. Note that the pattern of clustered spikes remained stable throughout the entire recording sequence, conforming to the criterion for assessing stability of the recording electrode as described above. (C) Average waveforms from four clusters shown in (B); the color of the waveform corresponds to the color of the corresponding spike cluster. The average waveform +/- standard deviation is shown for the first behavioral session (left) and the last behavioral session on the second day (right). The average waveforms and the average amplitudes remained consistent throughout the experiment as demonstrated by the comparison between AM1, Day1 and PM4, Day2. (D) Spatial maps from the same four cells in (C) during the behavioral sessions in (B). The average firing rate in each 5 by 5 cm pixel is color coded with a color scale from 0 Hz (blue) to the maximum peak rate within the day (red). The peak firing rate for each cell on each day is indicated to the right of the PM4 map. Although spatial maps from tracked cells were not used to determine whether clusters were the same from one day to the next, we visually inspected spatial maps from tracked cells. The location of place fields was identified to be the same from one day to the next, confirming that our method for tracking cells was reliable. (E) For each cluster that had a place field, isolation distance was used as a metric of cluster quality. Higher isolation distance indicates better cluster quality. (Left) There was no difference in the distribution of cluster quality between CA1 and CA3 on either the single day (top, KS test statistic = 0.21, n.s.) or the clusters that were tracked across two days (bottom, KS test statistic = 0.29, n.s.). (Middle) The variability in shape preference observed in CA1 was not related to cluster quality. The same data as presented in Figure 1.4 are shown with dots color-coded by cluster-quality quartile. The clusters with highest quality are shown as blue circles, the clusters with lowest quality are shown as yellow squares, the middle 50% of clusters are in gray (clusters

Figure 1S.1. Clusters of spikes from single cells were tracked across multiple days, continued.

whose quality metric was undefined are omitted from the plot). The degree to which clusters vary with reference to the identity line is not different between the best and worst quartiles for either single day [top, $F(20,20) = 0.97$, n.s.] or two-day [bottom, $F(11,13) = 1.37$, n.s.] data. (Right) The same measures are shown for CA3 data. Variability was not different between the best and worst quartiles for single day [top, $F(9,9) = 2.79$, n.s.] or two-day data [bottom, $F(6,6) = 2.43$, n.s.]. (F) Cluster quality was assessed using a second metric, L-Ratio. For this metric, lower scores indicate better clusters. Panels are laid out as described in (E). There is no difference in the distribution of cluster quality between CA1 and CA3 for single day data (KS test statistic = 0.24, n.s.). For clusters tracked across two days there is a small but statistically significant difference in distributions with CA1 clusters having lower scores than CA3 ($P < 0.05$, KS test statistic = 0.38), indicating that cluster quality in CA1 is slightly better than in CA3. There is no difference in variability between the clusters in the best quartile and worst quartile in either brain region [CA1, single day: $F(20,20) = 1.21$, n.s.; CA1, two day: $F(12,12) = 1.21$, n.s.; CA3 single day: $F(9,9) = 0.63$, n.s.; CA3 two day: $F(6,6) = 0.38$, n.s.]
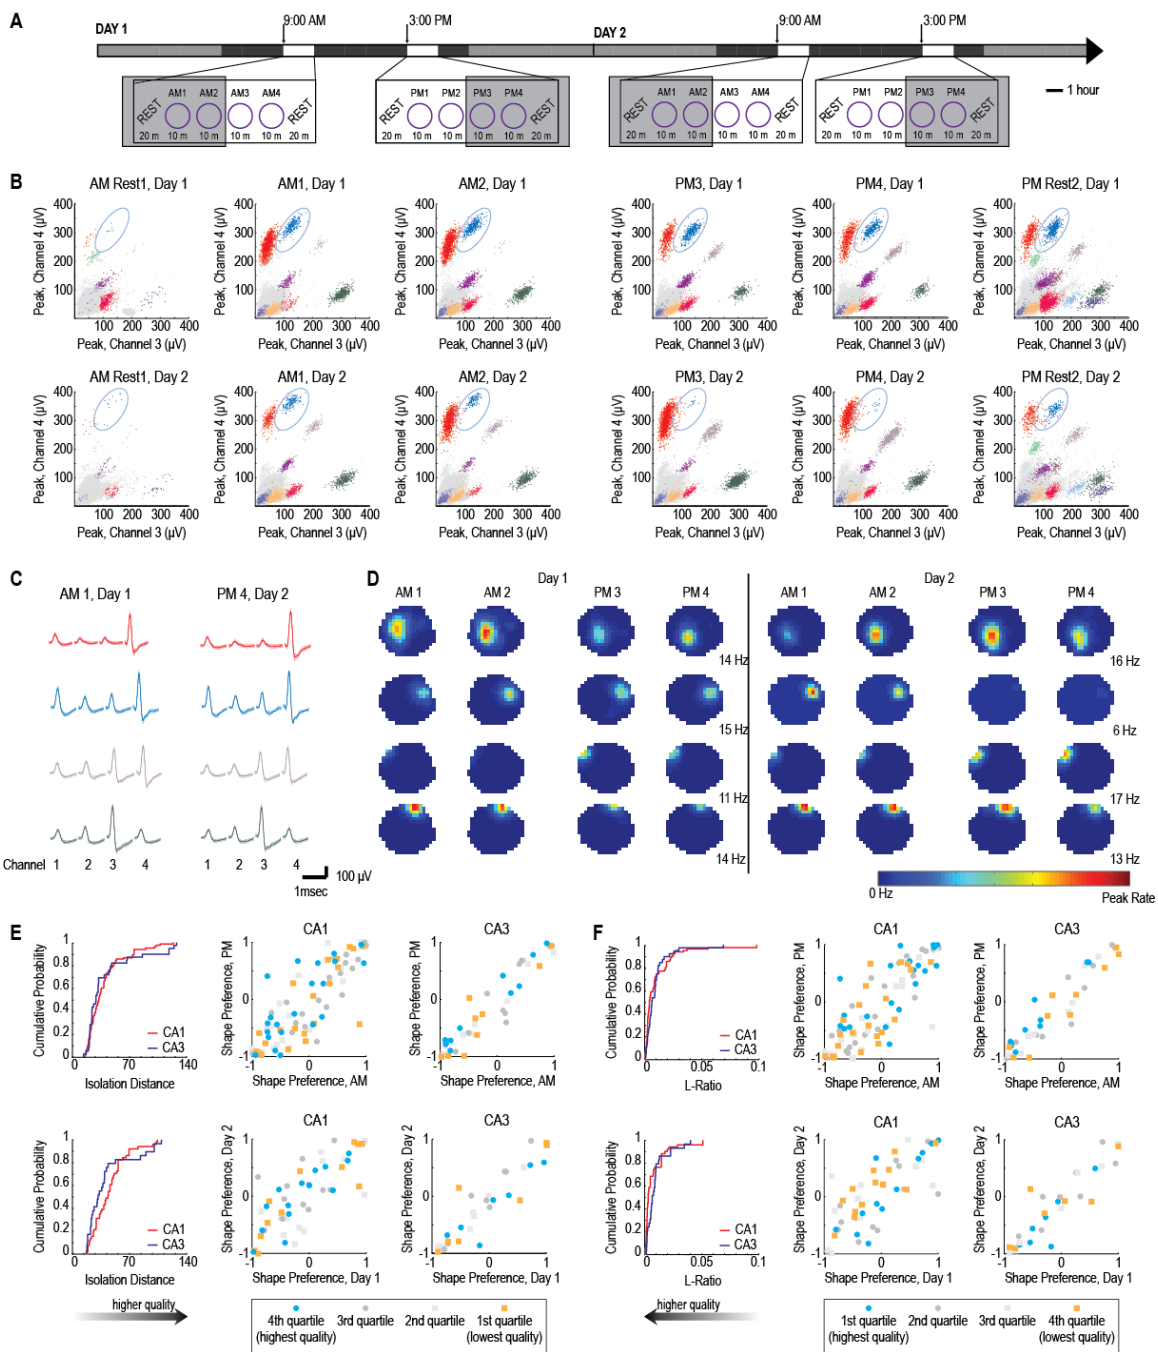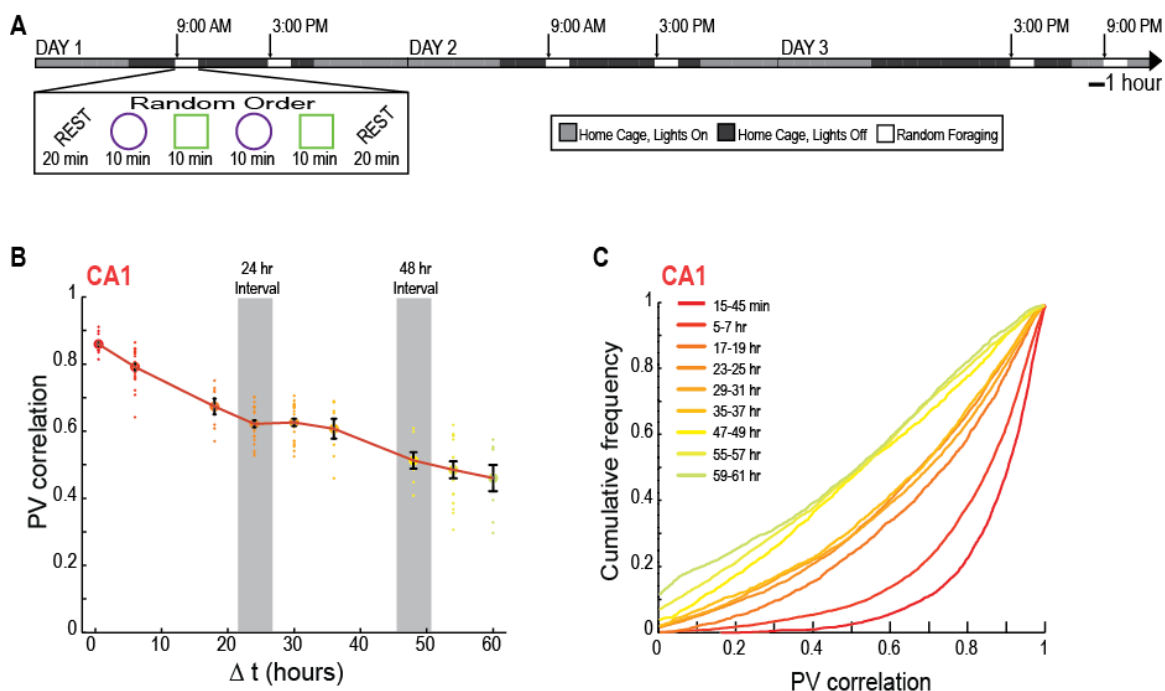
Figure 1S.2. The similarity in the activity patterns of the CA1 neural network continued to decrease on a third recording day. (A) To assess differences in CA1 network activity patterns across time intervals of up to 60 hours, we extended the hippocampal recordings across three days in three animals. On the third day, behavioral sessions were shifted by six hours, providing additional comparisons for the population of cells that was tracked across all three days (n = 30 CA1 cells). (B) Population vector correlations between pairs of recordings in the same enclosure shape are shown as dots. The black error bars correspond to the mean +/- SEM for recordings at a particular time interval. The correlation coefficients for the CA1 population activity decreased as a function of elapsed time between recording sessions [F(8) = 78.8, P < 0.001]. Twenty-four and forty-eight hour intervals, which represent recordings at matching times of day, are indicated in gray. Repeated CA1 recordings at the same time of day on two consecutive days show a smaller correlation than repeated recordings at shorter intervals, but at a different time of day (P < 0.001 for the posthoc comparison). The decrease in the population vector correlation therefore does not show a circadian pattern. (C) Cumulative distribution functions for population vector correlations between pairs of recordings in the same enclosure shape.
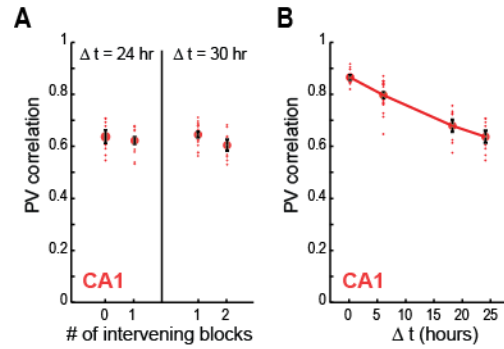
Figure 1S.3. The activation of hippocampal neurons during intervening behavioral testing is unrelated to the pattern of decorrelation with extended time. (A) Population vector (PV) correlations for comparisons between recordings in the same enclosure shape at intervals of 24 hours (left) or 30 hours (right). Each comparison between two sessions is shown as a dot, and the black bars are the mean +/- SEM for a time interval. Differences in the number of intervening blocks did not result in differences in the degree of decorrelation (comparisons between 24-hour intervals: t = -0.64, n.s.; comparisons between 30-hour intervals: t = -1.82, n.s.). (B) PV correlations between pairs of recordings in the same enclosure shape without intervening blocks. The twenty-four hour comparisons are from recordings without an intervening AM block (see Figure 1S.2). Each comparison between two sessions is shown as a dot, and the black bars are the mean +/- SEM for a time interval. A decorrelation of the CA1 population activity with time was observed [F(3) = 41.6, p values for all posthoc comparisons < 0.001 except that the comparison between the 18 and 24 hour interval is n.s.].
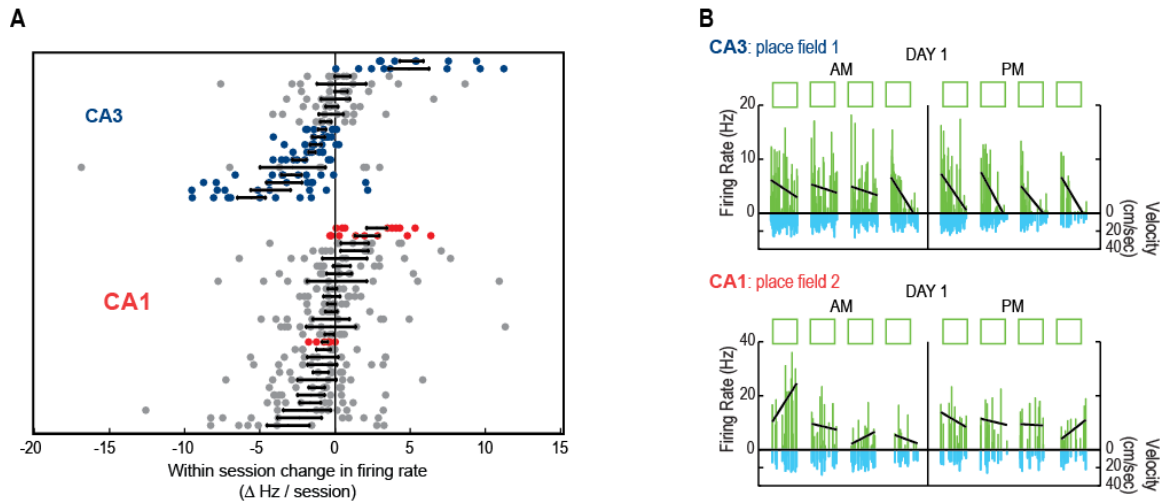
Figure 1S.4. Changes in network activity patterns at shorter time scales (A) The change in firing rate within a session is shown for individual fields in CA3 and CA1, recorded across eight sessions in the same shape (four during the AM and four during the PM block).To obtain estimates for the firing rate at the beginning and at the end of each session, a regression line was fitted to the firing rates during single passes through the place field. For each field, the difference in firing rate between the beginning and the end of each of the eight 10-minute sessions is shown as a dot. Colored dots depict rate differences when the mean rate change was significantly different (at the 0.05 level) from zero, which is the expected value in the case of random fluctuation in firing rate. The black bars show the mean +/- SEM for each field. The fields of each hippocampal subregion are sorted in the order of their mean firing rate difference. Many CA3 cells consistently changed their firing rates within a series of 10-minute recording sessions, often resetting towards the previous start-value. Together with the finding that CA3 changed to a small degree over extended time periods, these results show that changes in firing within a session do not result in corresponding changes in population activity over long time intervals. (B) Representative CA3 and CA1 place fields for the effects described in (A). Place cells were repeatedly recorded within the same shape. The regression line for firing rate versus time in each session is shown in black.
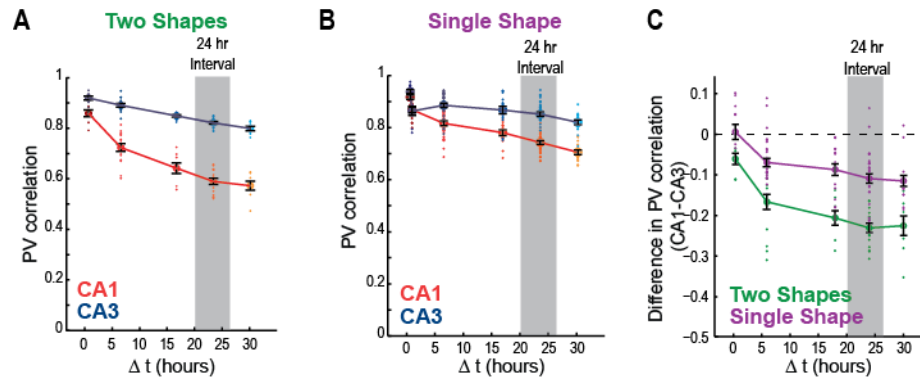
Figure 1S.5. The difference between population vector correlations in CA1 and CA3 is consistent across experimental paradigms. (A-B) For comparison, data from Figure 3C (A) and Figure 5B (B) are redrawn. In (A) the PV correlation is plotted for the two shape experiment, in which animals were run in two enclosure shapes during each recording block. In (B) the same analysis is shown for the single shape experiment, in which animals were exposed to only one shape for the entire 30-hour experiment. (C) The difference between CA1 and CA3 was larger in the two-shape condition compared to the single shape condition, but increased in parallel over time for the two conditions [Main effects for time and paradigm: $F(4) = 22.3$, $P < 0.001$ and $F(1) = 97.4$, $P < 0.001$; no significant interactions]. This suggests that the decorrelation due to time may be added in CA1 to a small decorrelation that occurs in both subregions in response to other aspects of the experience, such as the ongoing changes in the shape of the recording enclosure.
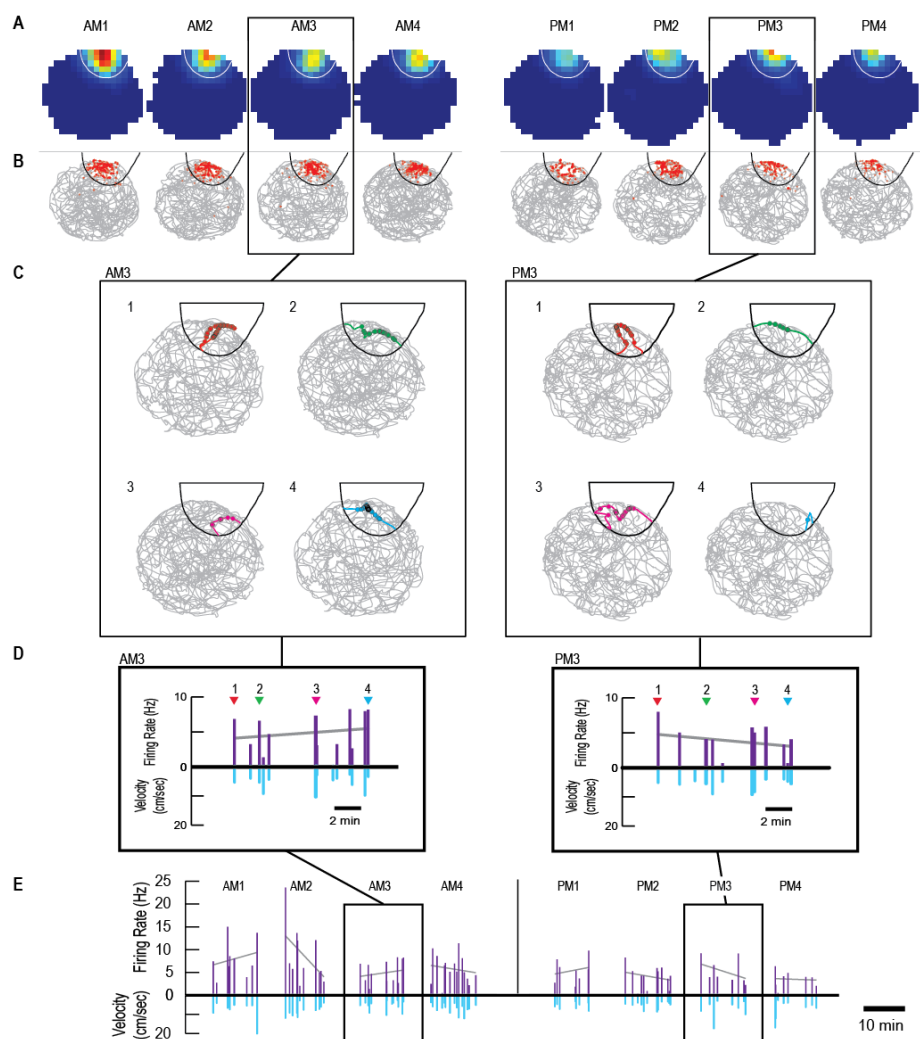
Figure 1S.6. Method for calculating firing rates during individual passes through a place field. (A) The spatial maps from a representative cell are shown for the four behavioral sessions in the morning (AM sessions 1-4) and for the four sessions in the afternoon (PM sessions 1-4). The boundary of the place field is outlined in white (see Figure 1S.11 for the method that is used for defining boundaries) (B) For each 10-minute behavioral session, the path of the rat is shown in gray. The spikes fired by the cell are superimposed on the path in red. As expected, most spikes occur within the place field boundary, shown in black. (C) Each of the four panels shows spikes of an individual pass through the place field during a 10-minute session (AM3 on the left, PM3 on the right). The path of the rat from the entire 10-minute session is shown, and single passes (from crossing the boundary into the field until exiting the field) are highlighted in color. Spikes that occurred during the pass through the field are superimposed on the path as dots. (D) For each pass through the field, we calculated the average firing rate as the number of spikes that occurred during the pass divided by the duration of the pass. We also calculated the average velocity at which the rat traveled. The average firing rate for each pass is shown as a colored bar (purple for foraging sessions in the circle), with firing rate on the y-axis and the time within the behavioral session on the x-axis. The average velocity of each pass is represented as a bar plotted downwards from the time axis (in blue). The representative passes shown in (C) are marked with arrows. (E) Firing rates for all passes though the place field throughout the entire recording day.

Figure 1S.7. Firing rates of complete cell samples recorded simultaneously on the same tetrode as the representative examples shown in Figure 1.2A and B. The examples that are also shown in Figure 1.2 are highlighted in gray. (A) CA1 place fields recorded simultaneously on the same tetrode as field 1 in Figure 1.2A. (B) CA1 place fields recorded simultaneously on the same tetrode as field 2 in Figure 1.2A. (C) CA3 place fields recorded simultaneously on the same tetrode as field 3 in Figure 1.2B. (D) CA3 place fields recorded simultaneously on the same tetrode as field 4 in Figure 1.2B. For each place field in (A-D), all 10-minute recording sessions throughout the morning (AM) and afternoon (PM) are shown. The order of the enclosure shapes is indicated above the plot on the top left. Each bar represents the firing rate of the cell during a pass of the animal through the place field (see Figure 1S.6). Time within the recording session is shown on the x-axis. The bars are colored green for square enclosures and purple for circular enclosures. For each pass, the corresponding running speed of the animal is plotted downwards below the x-axis (blue bars). A pass during which the cell does not fire is indicated by the presence of a velocity bar for that pass without a corresponding rate bar. A regression line for firing rate versus time is fitted to each session (shown in black). (E) Cresyl violet stained section showing recording sites in the hippocampus. Final recording sites in CA3 are outlined in blue, for CA1 in red. Scale bar, 500 μm.
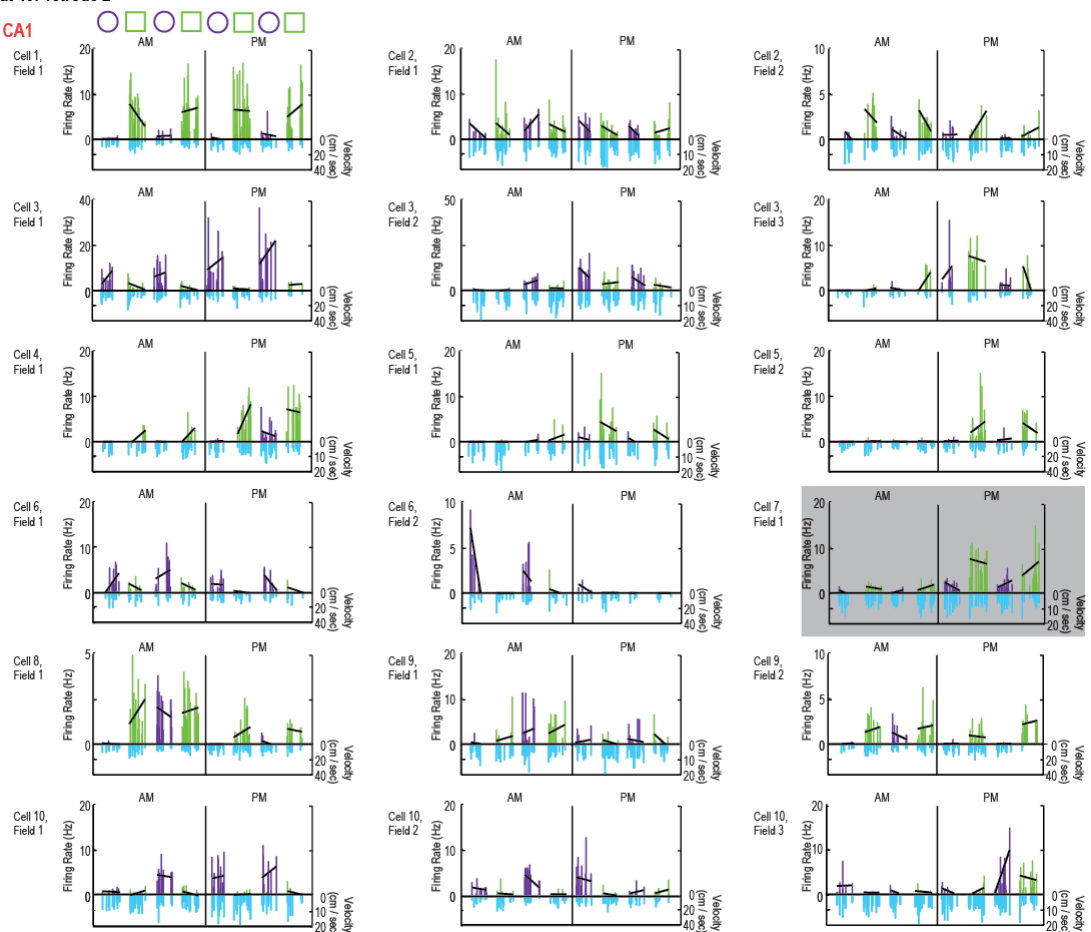
Figure 1S.7 continues on the next page

Figure 1S.7 continued
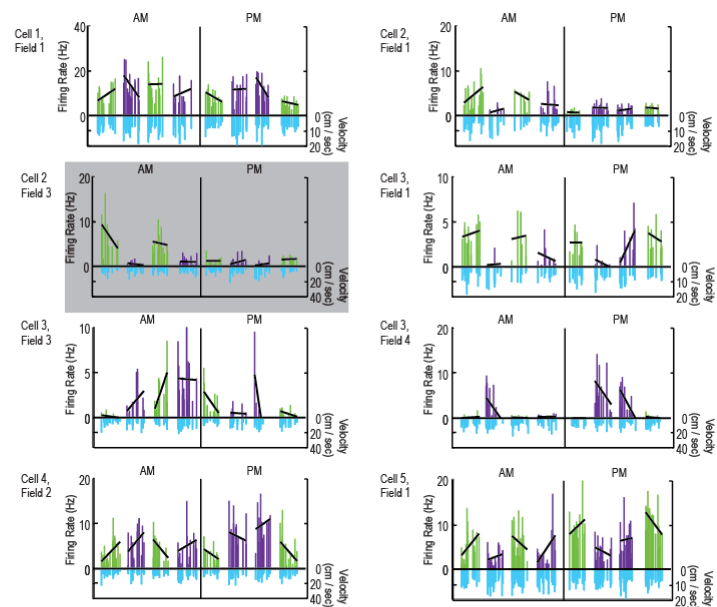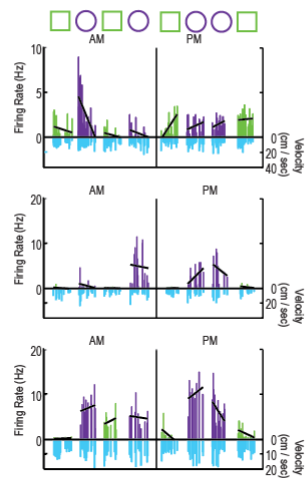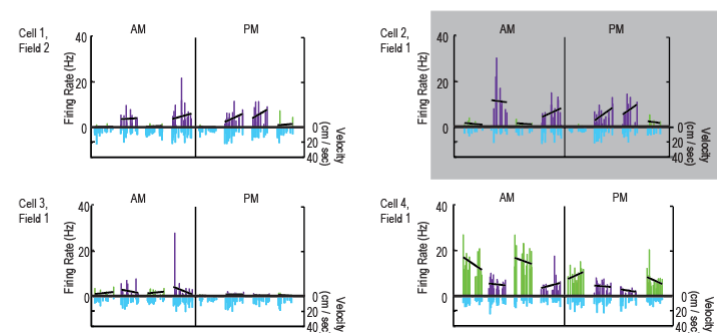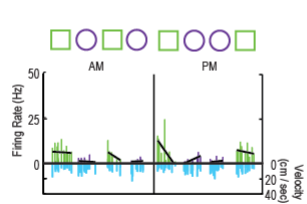
Figure 1S.8. Firing rates during individual passes through place fields were weakly correlated with the animal's velocity and with the minimum distance to the center of the place field. (A) For each field, we calculated the correlation (Spearman's rank correlation coefficient) between the animal's average velocity and the cell's firing rate on individual passes through the field. All passes in all behavioral sessions from a single enclosure shape (square or circle) were combined. Fields that were silent in a particular configuration were excluded from analysis, as were fields through which the rat passed fewer than 5 times in any 10-minute session. The distribution of correlation values is shown for CA1 (top panel, median = 0.16) and CA3 (lower panel, median = 0.23). Examples of representative fields in (B) CA1 or (C) CA3. Each panel shows a scatter plot of rate vs. velocity for all passes through the place field. Correlation coefficients are noted for each field, and coefficients that are significantly larger than zero are marked in blue. (D-F) Panels are laid out as in (A-C), but here the correlation between firing rate and the minimum distance of the path to the center of the place field is reported [CA1, median = -0.18; CA3, median = -0.23; in (E) and (F), fields with a correlation coefficient significantly less than zero (at the 0.05 level) are marked in pink]. We expect a negative correlation because the center of the place field is defined as the place at which the cell's peak firing occurred. Traveling closer to the center should thus lead to increased spiking. While there is a moderate correlation in firing rate with both velocity and distance to the center of the place field, this correlation is not sufficiently high to explain the variability in our data. Of particular relevance to our data, rats did not show a consistent trend over extended time in either their velocity profiles or in the average distance to the center of the place field. The difference in the mean velocity between morning and afternoon sessions never exceeded 2 cm/sec and, on average, was less than 0.5 cm/sec.
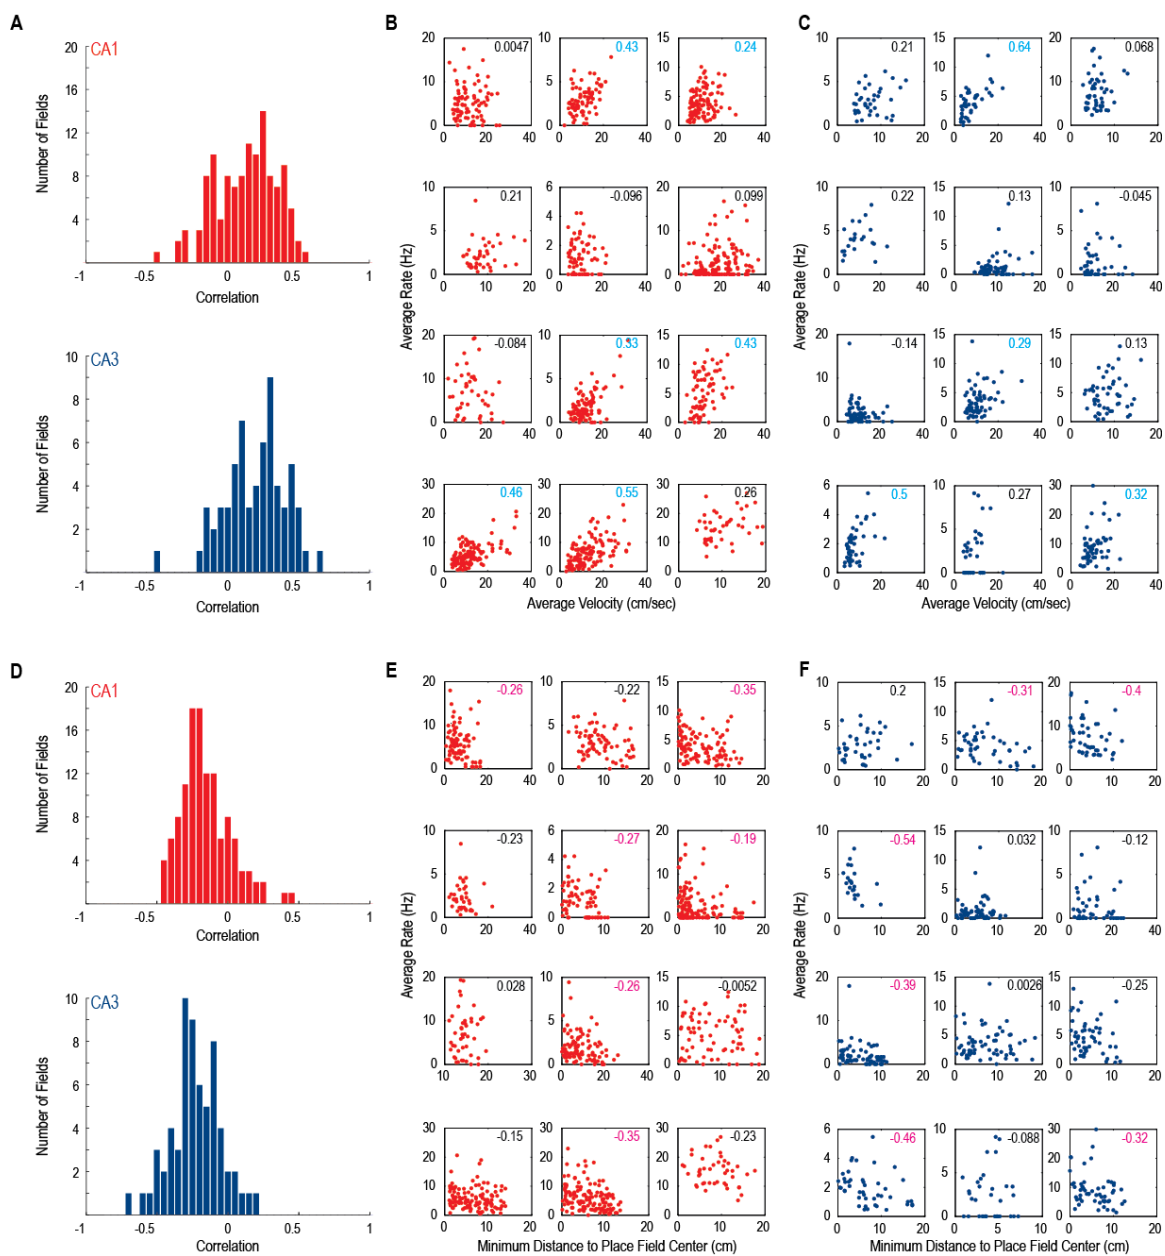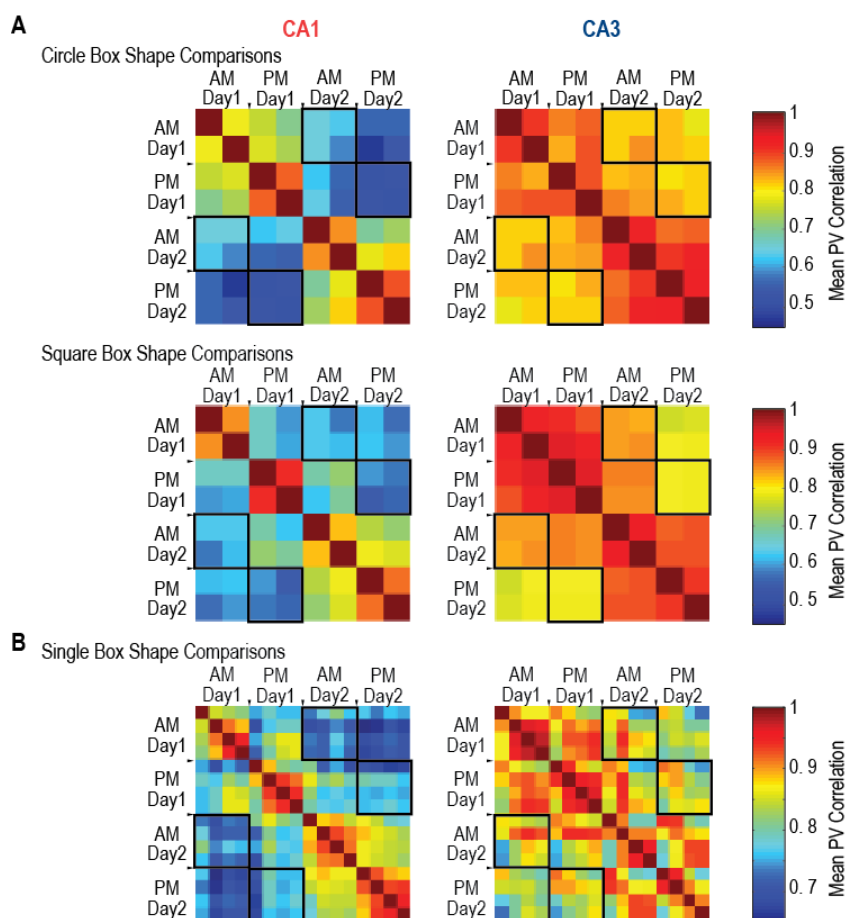
Figure 1S.9. Correlation matrices for repeated recordings in the same enclosure shape. (A) Each AM and PM recording block during two recording days (Day 1, Day 2) included two sessions in the square enclosure and two sessions in the circle enclosure (see Figure 1.3). Over two days, eight recording sessions were therefore performed in enclosures of the same shape. The correlation matrices depict all possible comparisons between repeated recordings in the same shape at different time points. The mean population vector correlation (i.e., the average across the spatial bins of the recording enclosure) of each pairwise comparison between identical enclosures is shown. The scale bar to the right of the matrices shows the color code for the correlation coefficients. All four matrices use the same color scale. Comparisons between the same sessions are shown along the diagonal, and their correlation coefficient is, by definition, 1. Comparisons at 24-hour intervals are highlighted by black boxes. The matrices are symmetrical along the diagonal and the same values are thus shown in the top right and bottom left. In CA1, the highest correlation coefficients were measured at the time intervals closest to the diagonal and decreased correlation coefficients at longer time intervals. A small decorrelation was also observed in CA3, but the effects are minor compared to CA1 (see Figure 1.3 for statistics). (B) For the single-shape control days, each AM and PM block during two consecutive recording days included four recording sessions in the same shape. The correlation matrices therefore include comparisons between each of the sixteen sessions. As observed in (A), population vectors in CA1 (left) became substantially decorrelated with increasing time intervals. In contrast, population vectors in CA3 remained highly correlated for intervals of up to 30 hours (see Figure 1.5 for statistics).

Figure 1S.10. CA1 place fields that were recorded simultaneously on the same tetrode as those in Figure 1.4B. The firing fields that are identical to those in Figure 1.4B are highlighted in gray. For each place field, firing correlates are shown for all 10-minute recording sessions within the morning (AM) and within the afternoon (PM) for each of the two recording days, with time within each recording session on the x-axis. The order of the enclosure shapes throughout the sessions is shown above the plot on the top. Each bar represents the firing rate of the cell during a pass of the animal through the place field as described in Figure 1S.6.

Rat 102: Tetrode 5

CA1
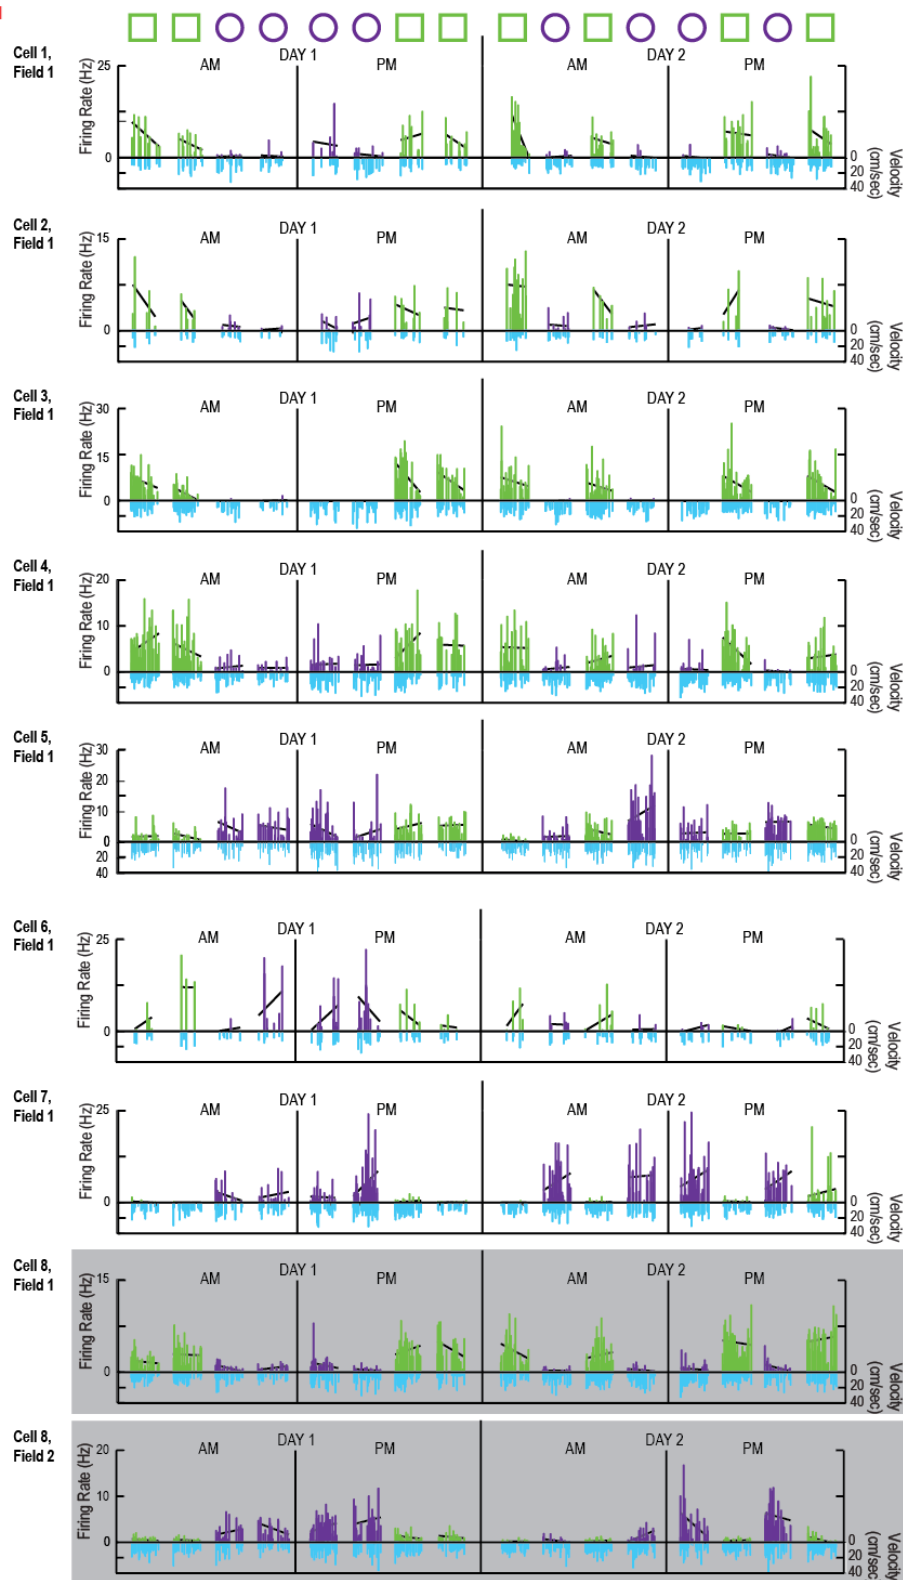
Figure 1S.11. Illustration of the method for calculating boundaries of place fields. To illustrate that the analysis procedure separates adjacent fields, a cell with multiple fields is shown as an example. (A) Spatial maps are first calculated for each of the eight behavioral sessions in a single day. The firing rates within each map are color-coded according to the scale bar on the right, and the shape of the box is superimposed on each map. (B) A reference map was calculated as the average of the eight maps in (A). Note that the place fields in the reference maps are overlapping, but that the maps of individual sessions suggest that the firing rates in different areas appear to be modulated independently. (C) Contours were calculated at twenty levels between zero and the peak rate. The contours are shown superimposed on the reference map. (D) Peak contours (shown in red, green, and blue in the three panels) were found. Surrounding contours are in gray. The three fields share many of the contours at lower levels. (E) Each shared contour is divided into segments at inflection points of the contour, and each segment is assigned to the nearest field. Each segment is illustrated in the color of the field to which it is assigned. (F) The end points of the segments and the outside contours of different colors in (e) are combined into outer boundaries of the fields. (G) The maps in (A) are redrawn with the boundaries of the three identified fields superimposed. (H) Two additional examples of cells with multiple, adjacent fields. The large panel on the left depicts the reference map, the smaller panels show the eight individual sessions with the field boundaries superimposed [arranged as in (A)]. (I) The method also identifies the boundaries of single fields. Three examples of cells with single fields are shown. (J) Histograms of the number of fields of CA1 cells (top) and CA3 cells (bottom) in the single-shape condition. (K) Histograms of the number of fields of each cell in the two-shape condition. The shift towards a higher number of fields per cell is a result of recording in two box shapes, because fields occasionally appear in one box shape but not in the other.
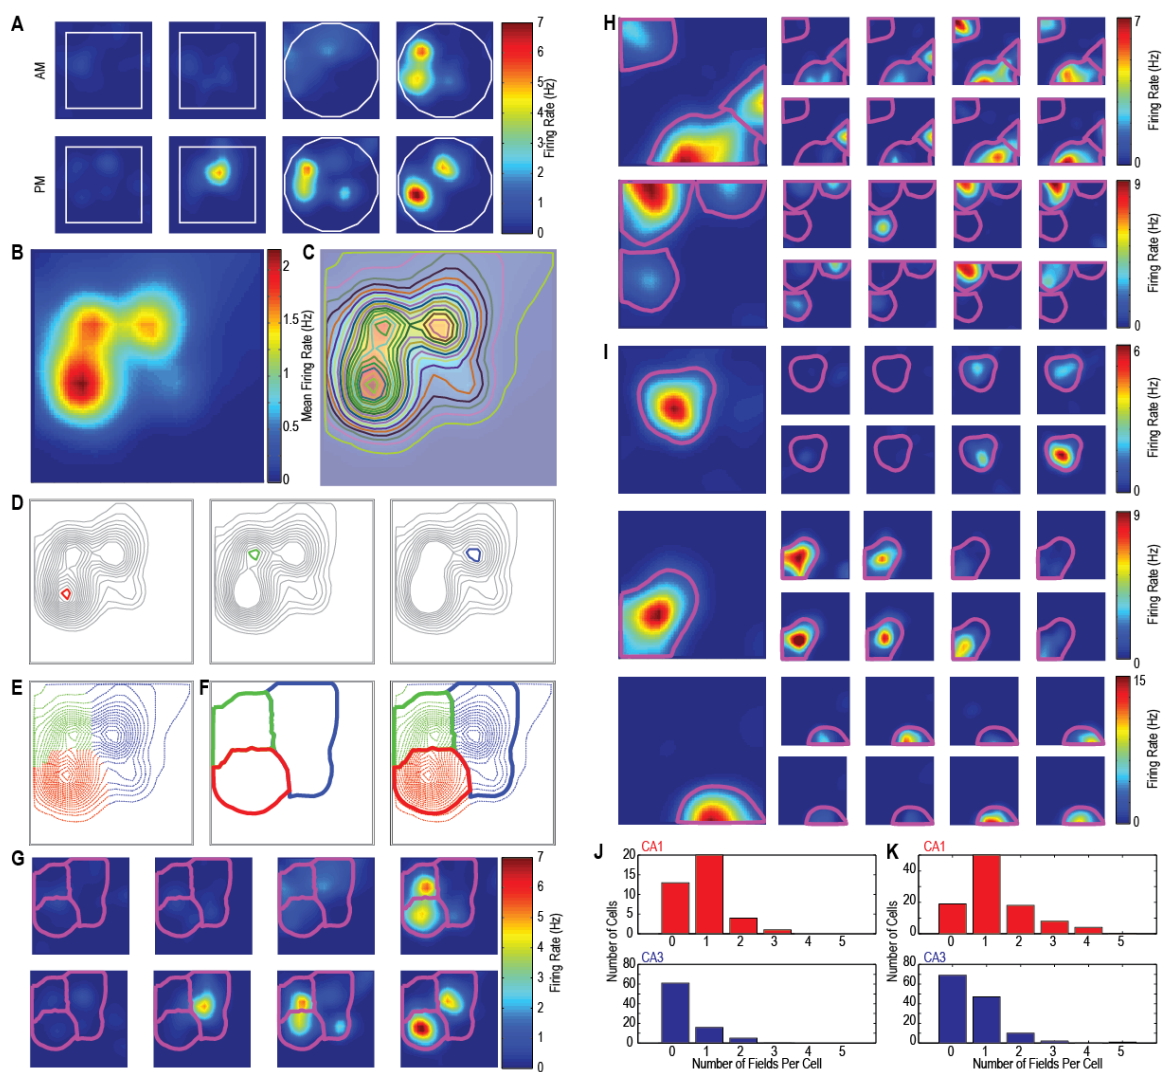
Appendix 1.2: Supplemental methods

Electrode turning. One tetrode remained in the cortex and was used as a reference for all recordings. Another tetrode was lowered to the stratum lacunosum-moleculare to record hippocampal local field potentials. Of the remaining twelve tetrodes, approximately half were advanced towards the CA1 cell layer and the other half towards the CA3 cell layer over a period of at least two weeks after surgery. The depth profile of the local field potential on different tetrodes (Csicsvari, Hirase, Czurkó, Mamiya, & Buzsáki, 1999) was used as a guide for electrode movement. When the tetrodes approached the cell layers, further movement of tetrodes was done in small increments over several days. Once the tetrodes were adjacent to the cell layer, as indicated by the presence of low amplitude, multi-unit activity, tetrodes were not turned again and were allowed to settle into the cell layer for stable recordings over a period of many days.

Recording procedures. For recording spikes and local field potentials, the electrode assembly was connected to a multichannel head-mounted preamplifier. The x-y position of light emitting diodes on the preamplifier was tracked at 30 Hz by processing video images. Unit activity was amplified and band-pass filtered at 600 Hz to 6 kHz. Spike waveforms above a trigger threshold (40-60 μV) were time-stamped and recorded at 32 kHz for 1 ms.

Spike sorting. Spike sorting was performed offline using the graphical cluster-cutting software MClust (Redish, n.d.). Clustering was performed manually in two-dimensional projections of the multidimensional parameter space (consisting of waveform amplitudes, the difference between peak and trough of the waveform, and waveform energies). Autocorrelation and cross-correlation functions were used as additional separation criteria. Putative excitatory cells were distinguished from putative interneurons by spike width and average rate. Only putative excitatory cells were included in analysis. Recording stability within a day and across days was confirmed as described in Figure 1S.1. In particular, sleep sessions before and after the behavioral sessions were used in addition to the recordings in behavior to assess recording

stability, and a modified version of MClust was used that allows for the confirmation of the cluster boundaries of each cell in each 10-minute recording session.

Cluster Quality Metrics. For each cluster, we calculated the L-ratio and Isolation Distance as described (Schmitzer-Torbert, Jackson, Henze, Harris, & Redish, 2005).When calculating the Mahalanobis Distance, we included the features that we used for clustering (see above). We calculated the cluster quality for each experimental block by including all putative spikes that were detected during all behavioral sessions of the block. Full day cluster quality was reported as the mean cluster quality between the two blocks of a given day. Cluster quality was compared across hippocampal subregions using a two-sample Kolmogorov-Smirnov (KS) test.

Spatial Maps. Spatial firing rate distributions for each well-isolated neuron were constructed in the standard manner, by summing the total number of spikes that occurred in a given location bin (5 cm by 5 cm), dividing by the amount of time that the animal spent in that location, and smoothing with a Gaussian centered on each bin (J. K. Leutgeb et al., 2007).

Spatial correlations. Spatial firing patterns of individual cells were compared between sessions with a spatial correlation procedure (Fenton & Muller, 1998; Hetherington & Shapiro, 1997; J. K. Leutgeb et al., 2007; Save, Nerad, & Poucet, 2000). The Pearson correlation coefficient for firing rates in corresponding pixels was calculated for pairs of spatial maps of each cell. Pixels visited less than 150 msec in either session were excluded to avoid artifacts in the correlation measure.

Place field boundaries. Standard practice has been to define boundaries by first finding the pixel with the maximum firing rate and then iteratively extending the field to any adjacent pixels that exceed a predefined threshold (Fenton & Muller, 1998; Hetherington & Shapiro, 1997; J. K. Leutgeb et al., 2007; Save et al., 2000). In our data, we found that a subset of fields had clearly separated peak values that were nonetheless joined by adjacent pixels and could thus not

be separated by this method. We therefore developed a contour-based analytical method to separate fields (Figure 1S.11).

Individual passes through place fields. For each field and each experimental session, the path segments inside of field boundaries were identified. Time periods when the path segments were continuously inside the field boundary for at least 200 ms were considered a pass through the field if the path crossed inside of the contour at 60% of the field's peak value. The average firing rate for a pass was calculated as the number of spikes during the pass divided by the duration of the pass. In addition, we calculated the average velocity during the pass and the minimum distance to the center of the field (Figure 1S.6). For any analyses reported by fields, we discarded fields through which the rat passed fewer than five times in any individual session.

Rate changes within sessions. For each place field recorded on Single Shape, Day 1, the rate change per session was calculated by fitting a regression line to the rates during individual passes through the field and by taking the difference of the *y*-values of the line between the end and start of the session. If rates fluctuated randomly in time, the average difference over repeated sessions should be close to zero. To measure whether a place field showed consistent rate changes, a Student's t-test was performed on the set of eight rate changes (one from each of the eight sessions within the day). If the mean rate change was found to be different from zero at the 0.05 level, that field was considered to change its rate consistently.

Shape preference score. To measure whether neuronal firing within a place field occurred preferentially in one enclosure shape, we calculated a shape preference score. For each field, the rates from the individual passes were combined to yield a distribution of firing rates for the square enclosure and a distribution of firing rates for the circular enclosure. The degree to which these distributions overlapped was calculated as the area under the ROC curve for the distribution in the square compared to the distribution in the circle. If all firing rates in the square enclosure were higher than in the circle enclosure, the area under the curve was 1. If all firing

rates in the square enclosure were smaller than in the circle enclosure, the area was 0. If the distributions were identical the area was 0.5. We then did a linear transformation on this measure (subtracting 0.5 and multiplying by 2), so that the shape-preference score would vary between -1 and 1, with 0 indicating no shape preference. The signed value of this score provides a nonparametric estimate of preferred firing in the square enclosure shape compared to the circle enclosure shape. The absolute value of the score gives an indication of the extent of firing preference for either shape.

Rate and shape-preference variability. To compare the variability in firing rate and in shape-preference between CA1 and CA3 across time intervals, we calculated the average firing rate of each cell and the shape preference score of each place field at different time points. Morning-firing rates were compared to afternoon firing rates, morning preference scores were compared to afternoon preference scores, and day 1 preference scores were compared to day 2 preference scores. The values for the two different time points were plotted against each other. If cells showed the same firing rate or the same degree of shape coding at the two time points, the data points would fall on the identity line. We used an F-test to determine whether the variance around the identity line differed between the CA1 and CA3 subregions. Because the average shape preference could systematically change over time, resulting in a best fit line that is different from the identity line, we repeated our analysis with a standard linear regression line. The results with the regression line replicated the findings with the identity line. To evaluate whether the variability in shape preference could be related to cluster quality, we repeated F-tests within each brain region for place fields with clusters in the highest quartile of cluster quality compared to place fields with clusters in the lowest quartile of cluster quality.

Population vector correlations. For each behavioral session, rate vectors were constructed by arranging the spatial maps of all cells recorded from all animals in an *x-y-z* stack, where *x* and *y* represent the two spatial dimensions in 5 cm by 5 cm bins and *z* represents the cell-

identity index (J. K. Leutgeb et al., 2005; S. Leutgeb et al., 2005). CA1 and CA3 cells were

analyzed in separate stacks. For multiple-day analysis, only cells that were tracked from the

beginning of the first day through the end of the last day were included in the population vectors.

The distribution of mean rates along the z-axis for a given x-y location represents the composite

population vector for that spatial bin and hippocampal subregion. For each pair of sessions, the

Pearson correlation coefficient was calculated for spatial bins at corresponding locations. Cells

with firing below 1 Hz in all bins of the two sessions were excluded from the analysis. To allow

for comparisons between the square and circle enclosure shape, the analysis was restricted to the

16 by 16 bins that were common to both shapes, yielding 256 correlation estimates for a pair of

sessions. The correlation coefficients of all spatial bins were averaged to estimate the average

population vector correlation for a pair of sessions. All average correlations for a particular time

interval (e.g., < 1 hr, 6 hrs) were compared using Student's t-tests or, for experiments with more

than two time points, using ANOVA with time interval and enclosure shape as factors or, in Fig

S5, with time interval and number of enclosure shapes as factors. Bonferroni corrections were

applied to correct t-tests and *posthoc* analysis for multiple comparisons. To create the cumulative

density plots, the population vectors of all spatial bins from pairs of maps with the same time lag

were combined, and the functions were plotted for each time lag. The complete matrix of

population vector correlations between all 10-minute sessions in which cells were recorded over

two days are shown in Figure 1S.9.

Three-dimensional reconstruction of the tetrode array in serial sections. All tetrodes

of the 14-tetrode bundle were identified by finding their electrode tracks across sections. A small

angular deviation between the plane of sectioning and the electrode tracks resulted in an apparent

downward shift of the tissue damage. The electrode tip was considered to be located in the

section where the tissue damage became negligible (see Figure 1S.7E).

Acknowledgements

Chapter 1, in full, is material as it appears in the *Proceedings of the National Academy of Sciences*, 2012. Mankin, Emily A., Sparks, Fraser T., Slayyeh, B., Sutherland, Robert J., Leutgeb, Stefan, and Leutgeb Jill K. The dissertation author was the primary investigator and author of this paper.

CHAPTER 2: HIPPOCAMPAL CA2 ACTIVITY PATTERNS CHANGE OVER TIME TO A LARGER EXTENT THAN BETWEEN SPATIAL CONTEXTS

Abstract

The hippocampal CA2 subregion has a different anatomical connectivity pattern within the entorhino-hippocampal circuit than either the CA1 or CA3 subregion. Yet major differences in the neuronal activity patterns of CA2 compared to the other CA subregions have not been reported. We show that standard spatial and temporal firing patterns of individual hippocampal principal neurons in behaving rats, such as place fields, theta modulation, and phase precession, are also present in CA2, but that the CA2 subregion differs substantially from the other CA subregions in its population coding. CA2 ensembles do not show a persistent code for space or for differences in context. Rather, CA2 activity patterns become progressively dissimilar over time periods of hours to days. The weak coding for a particular context is consistent with recent behavioral evidence that CA2 circuits preferentially support social, emotional, and temporal rather than spatial aspects of memory.

Introduction

The hippocampal CA fields are subdivided into the CA3, CA2, and CA1 subregions based on unique cytoarchitecture, connectivity, physiology, and gene expression patterns (Kjonigsen, Leergaard, Witter, & Bjaalie, 2011; Lein et al., 2005; Lorente de No, 1934; Woodhams, Celio, Ulfig, & Witter, 1993; X. Zhao et al., 2001). Standard circuit diagrams of the hippocampal formation include a trisynaptic loop from the entorhinal cortex to the dentate gyrus, from the dentate gyrus to CA3, and from CA3 to CA1, as well as additional direct connections from entorhinal cortex to the dentate gyrus, the CA3 subregion, and the CA1 subregion. Although it has long been recognized that the hippocampal CA2 subregion is distinct from the other CA subregions in that it receives inputs from the supramammillary nucleus (Cui et al., 2013; Jones &

75

McHugh, 2011; Maglóczky et al., 1994; Pan & McNaughton, 2004; Woodhams et al., 1993), it has primarily been considered as a transition zone between CA1 and CA3. However, major differences from CA3 and CA1 in CA2 connectivity within the hippocampal circuit and with entorhinal cortex have recently been described (Cui et al., 2013; Hitti & Siegelbaum, 2014; Kohara et al., 2014; Rowland et al., 2013). Notably, CA2 neurons are strongly excited by distal dendritic inputs from the entorhinal cortex and only weakly activated by CA3 inputs (R Bartesaghi et al., 2006; Renata Bartesaghi & Gessi, 2004; Chevaleyre & Siegelbaum, 2010; Kohara et al., 2014; M. Zhao et al., 2007). Thus, entorhinal information arrives in CA1 via the CA2 pathway in parallel to the direct pathway to CA1 and the indirect pathway through the dentate/CA3 subregions (**Figure 2.1A**).

In addition to these major differences in connectivity, CA2 is unique among hippocampal subregions in its mechanisms for long-term plasticity and in the baseline membrane properties of its principal cells (Caruana et al., 2012; Chevaleyre & Siegelbaum, 2010; Jones & McHugh, 2011; Pagani et al., 2014; M. Zhao et al., 2007). Furthermore, behavioral studies support a potentially unique functional role for CA2 in memory by demonstrating that the vasopressin 1b receptor, which is selectively enriched in CA2 neurons (Young, Li, Wersinger, & Palkovits, 2006), is necessary for social recognition and for discriminating the recency of an event (DeVito et al., 2009; Wersinger et al., 2002). In addition, CA2 has been directly found to be necessary for aggression towards intruders and for social memory(Hitti & Siegelbaum, 2014; Pagani et al., 2014). Neither vasopression 1b receptor knockout nor genetic silencing of CA2, however, affects spatial or contextual memory (Wersinger et al., 2002; Hitti and Siegelbaum, 2014; DeVito et al., 2009).

Major differences in anatomical and functional characteristics between hippocampal subregions do not *a priori* enable predictions of whether or how neural network firing patterns

will differ in behaving animals. For example, standard spatial and temporal firing patterns of hippocampal principal cells, such as place fields, theta modulation, and phase precession, are remarkably similar between CA1 and CA3, despite the substantial differences in connectivity and function between these subregions. Differences in neuronal activity patterns between these subregions only become apparent when considering how activity across the entire population of neurons responds to different behavioral situations. For example, when conflicting cues are presented, CA1 cells show a heterogeneous response, with different subpopulations responding to each aspect of an environment or memory task, while the cell population in the CA3 subregion more coherently follows one set of cues(I. Lee, Yoganarasimha, et al., 2004; J. K. Leutgeb et al., 2007; S. Leutgeb et al., 2004; Vazdarjanova & Guzowski, 2004). Additionally, firing patterns change over time in the CA1 population (Ludvig, 1999; Mankin et al., 2012; Manns et al., 2007; Ziv et al., 2013) while they remain more consistent within the CA3 network (Mankin et al., 2012). These differences in population responses indicate that each hippocampal subregion performs specialized computations that, in concert, can support the acquisition and retrieval of the different aspects of episodic memories (Marr, 1971; B. McNaughton & Morris, 1987; Rolls, 1989; Treves & Rolls, 1994). We thus asked whether the CA2 network might show neuronal coding at the population level that is distinct from CA1 and CA3 and, consistent with behavioral studies (DeVito et al., 2009; Hitti & Siegelbaum, 2014; Wersinger et al., 2002), may show less specialized network coding for spatial compared to temporal aspects of memories.

Results

   To examine how time and contextual change effect firing patterns in CA2, we obtained single unit and local field potential recordings in an experimental design in which rats randomly foraged in highly familiar environments in the morning and again, after an interval of 6 hours, in

the afternoon (**Figure 2.1**). Each morning and afternoon block consisted of four 10-min sessions, two in a square and two in a circular enclosure, and the enclosure shapes were presented in random order within each testing block. The identity of hippocampal CA2 cells (n = 62 cells in 5 rats) was tracked for a total of sixteen sessions (**Figure 2.S1**) from the morning block of one recording day through the afternoon block of the next recording day (i.e., 4 recording blocks with 4 sessions each). The recordings from the CA2 region were simultaneous with recordings from tracked CA1 cells (n = 43 in 4 rats) and/or tracked CA3 cells (n = 42 in 3 rats) (see **Table 2S.1** for the number of cells per rat). For the comparisons with CA1 and CA3, we also included additional simultaneous recordings from these two subregions in the same experimental design (n = 46 CA1 cells and 29 CA3 cells in 3 rats; Mankin et al., 2012). Recording locations were confirmed using anatomical criteria and immunohistochemical markers specific for CA2 neurons (**Figures 2.1B-C; Figure 2.S2**).

We first analyzed the spatial firing patterns of hippocampal cells within each of the sixteen 10-min sessions and, for each cell, averaged across the sixteen sessions. CA2 cells had a higher mean firing rate than either CA1 or CA3 cells (Mann-Whitney U: CA1 vs. CA2, $z = -2.96$, $P = 0.0062$; CA1 vs. CA3, $z = 1.51$, $P = 0.13$; CA2 vs. CA3, $z = 3.43$, $P = 0.0018$; see **Table 2S.2** and **2S.3** for detailed statistics for all Mann-Whitney U tests) (**Figure 2.2A** and **2.2B**). However, when considering the peak firing among all spatial locations in the enclosure, there were no differences between subregions (Mann-Whitney U: CA1 vs. CA2, $z = -0.30$, $P = 0.77$; CA1 vs. CA3, $z = 2.09$, $P = 0.087$; CA2 vs. CA3, $z = 2.19$, $P = 0.087$) (**Figure 2.2C**). A higher mean rate without a difference in peak rate could emerge from broader spatial firing in CA2 compared to the other hippocampal subregions. Consistent with this notion, the amount of spatial information per cell was lower in CA2 compared to CA1 and CA3 (Mann-Whitney U: CA1 vs. CA2, $z = 5.85$, $P < 0.001$; CA1 vs. CA3, $z = 0.93$, $P = 0.35$; CA2 vs. CA3, $z = -4.95$, $P < 0.001$) (**Figure 2.2D**).

Figure 2.1. Behavioral paradigm and the identification of recording sites in CA1, CA2, and CA3. (**A**) Schematic of the entorhino-hippocampal circuitry. Dotted lines denote CA2 connections that have recently been described but have not been confirmed in additional anatomical studies (Cui et al., 2013; Hitti and Siegelbaum, 2014; Kohara et al., 2014; Rowland et al., 2013). EC, entorhinal cortex; DG, dentate gyrus; S, subiculum; pp, perforant path; ta, temporoammonic path. (**B**) The hippocampal CA2 area (demarcated by yellow lines) is defined by positive α-actinin-2 immunoreactivity (brown), and cell bodies that are larger and less densely packed than in CA1, as indicated with a cresyl violet counterstain (purple). The locations of all recording tetrode positions along the A-P axis are projected onto a representative section according to their proximal to distal position within each subregion, but note that tetrode placement spans up to 1 mm along the A-P axis. Tetrodes that were more anterolateral were targeted to either CA2 or CA3 while tetrodes that were more postermedial were targeted to CA1 or CA2. Because the orientation of the dorsal hippocampus is from anteromedial to posterolateral, this strategy resulted in electrode positions in CA1/CA2 and CA3 that were approximately matched for the longitudinal position within the hippocampus. Along the transverse axis, most recordings were in proximal CA1 while few recordings were in distal CA3 (i.e., close to CA2). Although this increased our confidence that CA3 recordings could not have been misassigned to CA2, this resulted in recording sites that were not precisely matched for connectivity between CA3 and CA1, which is strongest from distal CA3 to proximal CA1 (Witter, 2007). (**C**) Tetrode tracks in a section with α-actinin-2 and cresyl violet staining. Overview (left, scale bar = 500 μm) shows tetrode tracks (red oval, CA1; teal oval, CA2) with areas shown at high magnification to the right (red and teal boxes). In CA1 (middle), cell bodies are small, the cell layer is compact, and there is minor co-staining for α-actinin-2 (scale bar = 50 μm). In CA2 (right), cell bodies are larger and less densely packed, and there is strong α-actinin-2 staining in cell bodies and proximal dendrites (white arrowheads; scale bar = 50 μm). See **Figure 2S.2** for further illustration. (**D**) Behavioral design. A series of four 10-min random foraging sessions were performed in the morning and again in the afternoon, over multiple days. Each time block consisted of a random sequence of 2 sessions with the recording enclosure in a square configuration and 2 sessions in a circle configuration. Twenty-min rest sessions flanked the behavioral sequence. Single unit recordings commenced after 9-20 days of pretraining ('Day 1' indicates the first day of electrophysiological recording).

A

EC
III
V VI ta
II
S
pp
DG
CA1
CA2
CA3

B

CA1
DG
CA2
CA3

C

D

DAY 1 9:00 AM 3:00 PM DAY 2 9:00 AM 3:00 PM — 1 hour

Random Order

REST □ □ ○ ○ REST
20' 10' 10' 10' 10' 20'

□ Random Foraging
□ Home Cage, Lights On
■ Home Cage, Lights Off
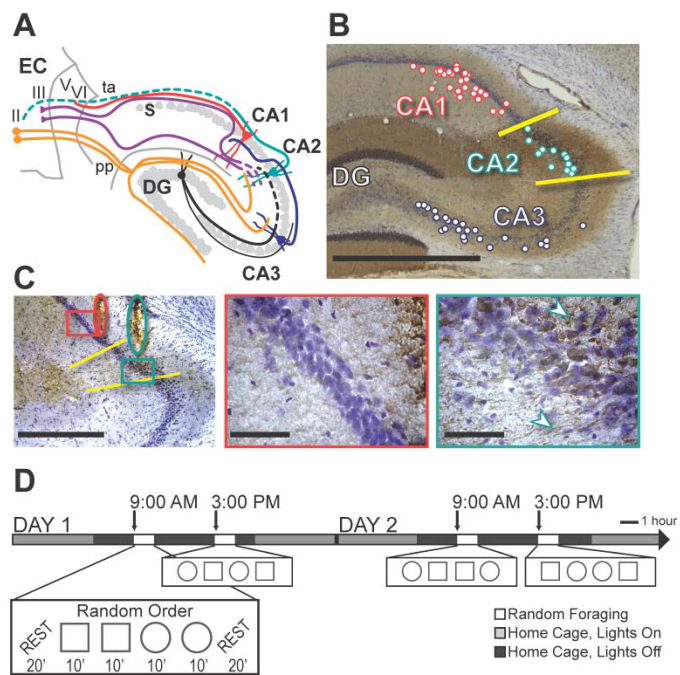
We then asked whether the lower spatial information in the CA2 cell population indicated that many of the CA2 cells were somewhat less spatially tuned or whether the lower average may have emerged from a heterogeneous population of CA2 cells in which some cellsremained spatially tuned while others had extremely weak spatial tuning. Of the CA2 principal neurons that were active during any of the recorded 10-min sessions (n = 54 of 62), we found that all had a spatial information score that was higher than 0.75 and place fields smaller than 25 % of the recording enclosure in at least one session, which indicates that each cell showed at least moderate spatial tuning.

To further characterize the spatial firing of CA2 cells, we calculated the number of place fields per cell and the size of each place field. The number of place fields per cell was higher in CA2 compared to the other subregions (Mann-Whitney U: CA1 vs. CA2, $z = -2.37$, $P = 0.035$; CA1 vs. CA3, $z = 2.30$, $P = 0.035$; CA2 vs. CA3, $z = 3.73$, $P < 0.001$) (**Figure 2.2E**). Because many cells in CA1 and CA3 were either silent during behavior and hence did not have a place field or had only one place field during behavior, this resulted in an average of less than one field per cell in CA1 and CA3. To examine whether a higher proportion of active cells per session in CA2 (CA1, 59.2 %; CA2, 68.2 %; CA3, 48.1 %) may have resulted in the larger number of fields, we restricted the analysis to cells that had at least one place field. Even when considering only cells with at least one field, CA2 had more fields per cell than the other hippocampal subregions (**Table 2S.2**) (Mann-Whitney U: CA1 vs CA2, $z = -2.85$, $P = 0.0088$; CA1 vs. CA3, $z = 1.49$, $P = 0.14$; CA2 vs. CA3, $z = 3.28$, $P = 0.0032$). When measuring field size, we found that the fields of CA2 cells were 24.5 % larger than those of CA1 cells (Mann-Whitney U: CA1 vs. CA2, $z = -2.54$, $P = 0.034$) (**Figure 2.2F**). The difference between CA2 and CA3 did not reach statistical significance (Mann-Whitney U: CA2 vs. CA3, $z = 2.2$, $P = 0.053$). The less pronounced difference in field size compared to spatial information can be explained by the fact that the
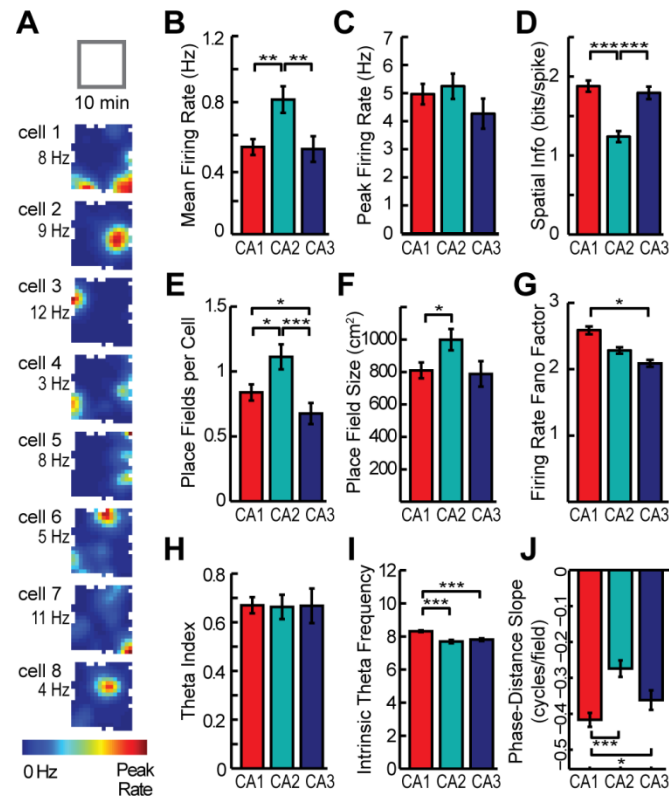
Figure 2.2. The spatial and temporal firing patterns of individual hippocampal CA2 principal neurons in 10-min sessions are largely consistent with those of CA1 and CA3, but with quantitative differences. (**A**) The firing rate maps of eight CA2 cells that were recorded simultaneously during a 10-min random foraging session in a square-shaped box. Average firing rate in each spatial location is represented from 0 Hz (dark blue) to the peak rate for the cell (red, noted to the left of each map). (**B-F**) Rates were higher and spatial tuning in CA2 was broader than in CA1 and CA3. The broader spatial tuning resulted from an increase in place field size and place field number per cell. (**G**) The variability in firing rate during individual passes through each place field did not differ between CA2 and the other CA subregions. (**H, I**) CA2 cells are modulated by the hippocampal theta rhythm to a similar extent as CA1 and CA2 cells and show intrinsic theta frequency comparable to CA3. (**J**) Place fields in CA2 showed phase precession, but to a lesser degree than in CA1. * $P < 0.05$, ** $P < 0.01$, *** $P < 0.001$ (Mann-Whitney U test). Bars are the mean ± SEM. See **Table 2S.2** and text for detailed statistics and **Figure 2S.3** for examples of phase precession.

reduction of spatial information in CA2 is caused by the combination of a larger number of fields per neuron and an increase in field size.

Because field size in CA2 was moderately larger than in CA1, we considered whether the increased field size in CA2 might be a result of slow spatial drift throughout the 10-min recording session. To examine this possibility, we first calculated field size using spatial maps that were obtained from either the first or the second 5-min half of the session. CA2 had larger fields than either CA1 or CA3 even over 5-min periods (Mann-Whitney U: CA1 vs. CA2, $z = -2.31$, $P = 0.042$; CA1 vs. CA3, $z = 1.94$, $P = 0.052$; CA2 vs. CA3, $z = 3.62$, $P < 0.001$). To then directly examine whether fields became larger by drift, we constructed spatial maps from 5 minutes of recording data that were sampled by including either only the odd or only the even minutes of the 10-min recording session. We then compared the field sizes from the continuous 5-min periods with the field size from the interleaved samples over 10 minutes and found no difference (Mann Whitney U: CA1, $z = 0.38$, $P = 0.71$; CA2, $z = 0.78$, $P = 0.44$; CA3, $z = -0.64$, $P = 0.52$). This is evidence that spatial drift on a time scale of minutes does not account for the larger fields in CA2.

After confirming that CA2 fields showed no evidence of greater spatial variability within a 10-min recording session than those in CA1 or CA3, we also examined the variability of the firing rates throughout the session. The variability in the firing rates between passes through the field did not differ between CA2 and the other subregions (Mann-Whitney U: CA1 vs. CA2, $z = 1.90$, $P = 0.11$; CA1 vs. CA3, $z = 2.85$, $P = 0.013$; CA2 vs. CA3, $z = 1.20$, $P = 0.23$) (**Figure 2.2G**). The standard measurement of variability is not sensitive to a systematic drift in firing rate throughout the 10-min recording session. We therefore estimated by how much the firing rate within each place field changed between the beginning and the end of a 10-min session, and found that the change was smallest in CA2 (Mann-Whitney U: CA1 vs. CA2, $z = 3.5$, $P = 0.0016$; CA1 vs. CA3, $z = 0.57$, $P = 0.57$; CA2 vs. CA3, $z = -2.4$, $P = 0.033$). Taken together, we found

no evidence that place field location or firing rate in CA2 is less stable than in CA1 or CA3 cells during a single 10-minute random foraging session.

We also examined the relationship between firing in CA2 cells and the hippocampal theta rhythm and found that the depth of theta modulation of CA2 cells was not different from CA1 and CA3 (Mann-Whitney U: CA1 vs. CA2, $z = 0.79$, $P = 0.85$; CA1 vs. CA3, $z = 1.28$, $P = 0.60$; CA2 vs. CA3, $z = 0.65$, $P = 0.85$) and that the intrinsic theta frequency of CA2 cells was not different from CA3 (Mann-Whitney U: CA2 vs. CA3, $z = -0.36$, $P = 0.72$) though slower than in CA1 (Mann-Whitney U: CA1 vs. CA2, $z = 4.75$, $P < 0.001$). To determine to what extent the phase at which cells fire within the theta cycle precesses during running through the place field, we calculated the slope of the phase-distance relationship for each place field (O'Keefe and Recce, 1993). The slopes of CA2 fields were significantly less than 0 (sign test: n = 53, sign = 4, $P < 0.001$), indicating that the majority of CA2 place cells phase precessed, although the magnitude of the precession was less than in CA1 (Mann-Whitney U: CA1 vs. CA2, $z = -4.97$, $P < 0.001$). This difference is consistent with the well-established relation between a larger field size and a less pronounced phase precession (Shen, Barnes, Mcnaughton, Skaggs, & Weaver, 1997). The difference in phase precession between CA2 and CA3 did not reach statistical significance (Mann-Whitney U: CA2 vs. CA3, $z = 1.84$, $P = 0.066$) (**Figure 2.2H-J** and **Table 2S.2**; see **Figure 2S.3** for examples of phase precession plots from individual CA2 fields). After finding that CA2 cells had the basic firing characteristics of hippocampal place cells with only minor quantitative differences during single 10-min random foraging sessions, we asked whether CA2 ensembles exhibited additional population coding features that are typical of neural networks in CA1 and CA3 and next analyzed neuronal activity patterns during the four recording sessions within a block, two in a square enclosure and two in a circular enclosure (**Figure 2.3A**). As expected (S. Leutgeb et al., 2005; Lever et al., 2002; Muller & Kubie, 1987), the activity patterns of CA1 and CA3 cells were distinct between these contexts. However, the shape

preference of CA2 cells was considerably lower than in the other CA regions (Mann-Whitney U: CA1 vs. CA2, $z = 7.59$, $P < 0.001$; CA1 vs. CA3, $z = 0.62$, $P = 0.54$; CA2 vs. CA3, $z = -4.53$, $P < 0.001$) (**Figure 2.3B-C;** see also **Figure 2S.4** for a description of the shape preference score accompanied by individual CA2 examples). Weak discrimination between spatial contexts by hippocampal network activity would typically be a result of unchanged network representations for different box shapes, but could also emerge when there is low baseline reproducibility of activity patterns for repetitions of the same shape. To distinguish between these alternatives, we computed population vector correlations between pairs of sessions (**Figure 2.3D**). We first tested whether spatial firing patterns were consistent between repeated visits to the same box shape. When selecting two consecutive sessions in the same box shape, CA2 showed activity patterns that were as consistent as in CA1 or CA3 [Mann-Whitney U: CA1 vs. CA2, $U(n1 = 7, n2 = 5) = 4$, $P = 0.091$; CA1 vs. CA3, $U(n1 = 7, n2 = 8) = 26$, $P = 0.87$; CA2 vs. CA3, $U(n1 = 5, n2 = 8) = 13$, $P = 0.71$] (**Figure 2.3D-E**). However, for repeated sessions in the same box shape that were separated by an intervening session of the other shape, CA2 was less consistent than CA1 or CA3 [Mann-Whitney U, CA1 vs. CA2, $U(n1 = 6, n2 = 4) = 1$, $P = 0.038$; CA1 vs. CA3, $U(n1 = 6, n2 = 4) = 0$, $P = 0.029$; CA2 vs. CA3, $U(n1 = 4, n2 = 4) = 0$, $P = 0.038$]. In fact, at this lag, the coding differences between repetitions of the same context were as pronounced as the coding differences between sessions in different contexts in CA2 [Mann-Whitney U: same shape vs different shape at lag 2, $U(n1 = 4, n2 = 8) = 10$, $P = 0.37$]. This suggests that any contextual coding that may be present in CA2 would be masked by temporal changes in network activity, even for intervals as short as 20 minutes.

To examine how time and contextual changes affect firing patterns in CA2 over longer time intervals, we analyzed the full experimental design in which rats randomly foraged in highly familiar environments in the morning and afternoon over 2 days (**Figure 2.4A**). The similarity between the CA2 population representations in identical enclosure shapes decreased
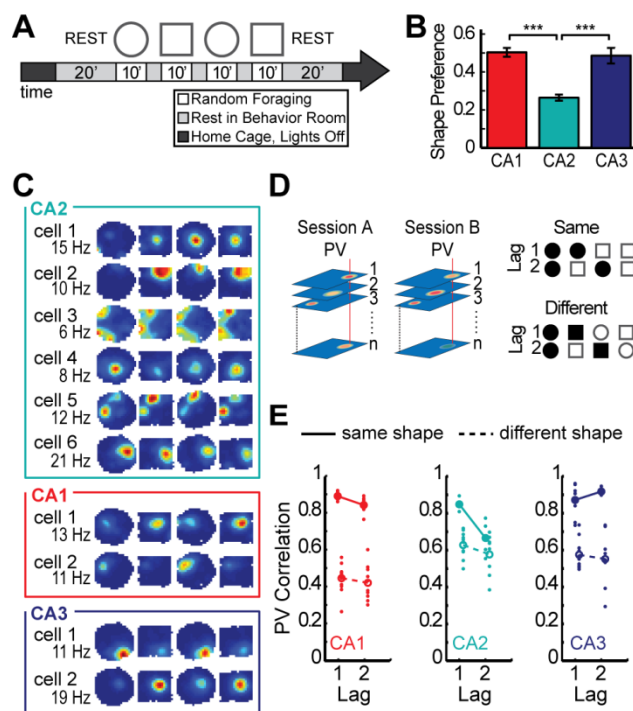
Figure 2.3. Place fields in CA2 are weakly modulated by spatial context. (**A**) Experimental timeline. Four 10-min sessions of random foraging in a square and a circle-shaped box. (**B**) Place fields in CA2 had lower shape preference scores than fields in CA1 or CA3. See **Figure S4** for a description of the shape preference score and individual examples. (**C**) Spatial firing rate maps for six representative CA2 cells and, for comparison, two CA1 cells and two CA3 cells. Maps are color coded as described in **Figure 2.2A**. CA2 cells 1-4 were recorded simultaneously with CA1 cells 1-2, and CA2 cells 5-6 were recorded simultaneously with CA3 cells 1-2. The changes in the spatial firing patterns of CA2 cells were not correlated with the switching between box shapes. (**D**) The schematic on the left shows how population vectors (PVs) were calculated. The spatial maps of all cells in corresponding sessions were arranged into *x-y-z* stacks, where *x* and *y* represent the two spatial dimensions and *z* represents the cell identity. In each stack, the distribution of firing rates along the *z* axis for a given *x-y* location represents the population vector for that spatial bin (examples are denoted by the red vertical lines in each stack). To compare two recording sessions, the Pearson correlation coefficient was calculated between each pair of population vectors at corresponding locations, and the correlation coefficients of all spatial bins were averaged. A PV correlation of 1 indicates identical activity patterns and 0 indicates independent patterns. The schematic to the right gives examples of comparisons between pairs of sessions (filled shape symbols) for each time lag in either the same-shape or different-shape category. (**E**) Each pairwise population vector correlation is shown as a dot, and the mean correlation for each lag is shown as a circle (filled, same-shape; open, different-shape). These measures revealed that same-shape comparisons in CA2 were as stable as in CA1 or CA3 only for adjacent sessions (lag 1). *** $P < 0.001$. Symbols and error bars are the mean ± SEM.

monotonically as a function of the temporal distance between exposures for time intervals up to 18 h and then reached an asymptote of approximately 0.35 [ANOVA: F(4) = 103.8, $P < 0.001$; Tukey's HSD: all P-values < 0.001, except 18, 24, and 30 hour time points were not significantly different from each other]. The asymptotic value is larger than the correlation when cell identity was shuffled ($> 99.9\%$ of shuffled values for each time point were smaller than the mean of the actual values) (**Figure 2.4B-C**; see **Figure 2S.5** for timescales of up to 60 hours and **Figure 2S.6** for example cells from each rat). The time-dependent effect in CA2 was sufficiently pronounced that the amount of change due to time after six hours already exceeded the amount of change produced by distinct spatial contexts without a time lag (see **Figure 2.4C**). There were no circadian fluctuations in CA2 population similarity (**Figure 2.4C**) or in normalized firing rates (**Figure 2S.5**), although there was a significant increase of firing rates within each recording block [two-way ANOVA: between blocks, F(3) = 1.53, $P = 0.21$; session number within blocks, F(3) = 5.23, $P = 0.0014$; Tukey's HSD between session 1 and session 4, $P < 0.01$; all other comparisons, n.s.].

Next, we compared the pattern of population similarity in CA2 with that of CA1 and CA3 in the same behavioral paradigm. The change in neuronal activity as a function of time was more pronounced in CA2 than in either the CA1 or CA3 cell populations [two-way ANOVA: brain region, F(2) = 1061.1, $P < 0.001$; time difference, F(4) = 184.8, $P < 0.001$; interaction, F(8) = 24.6, $P < 0.001$; Tukey's HSD between brain regions, $P < 0.001$]. We confirmed that the larger difference in firing patterns with time in CA2 compared to the other hippocampal subregions could not be attributed to the quality of the isolation of single units (**Figure 2S.1**) and that it was found in all but one single rat (**Figures 2S.5** and **2S.6**). We also confirmed that the decrease in correlation over time was found irrespective of the number of intervening recording blocks (**Figure 2S.5**). The passage of time rather than the amount of exposure to the recording environment therefore best explained the difference in the CA2 firing patterns. Finally, we
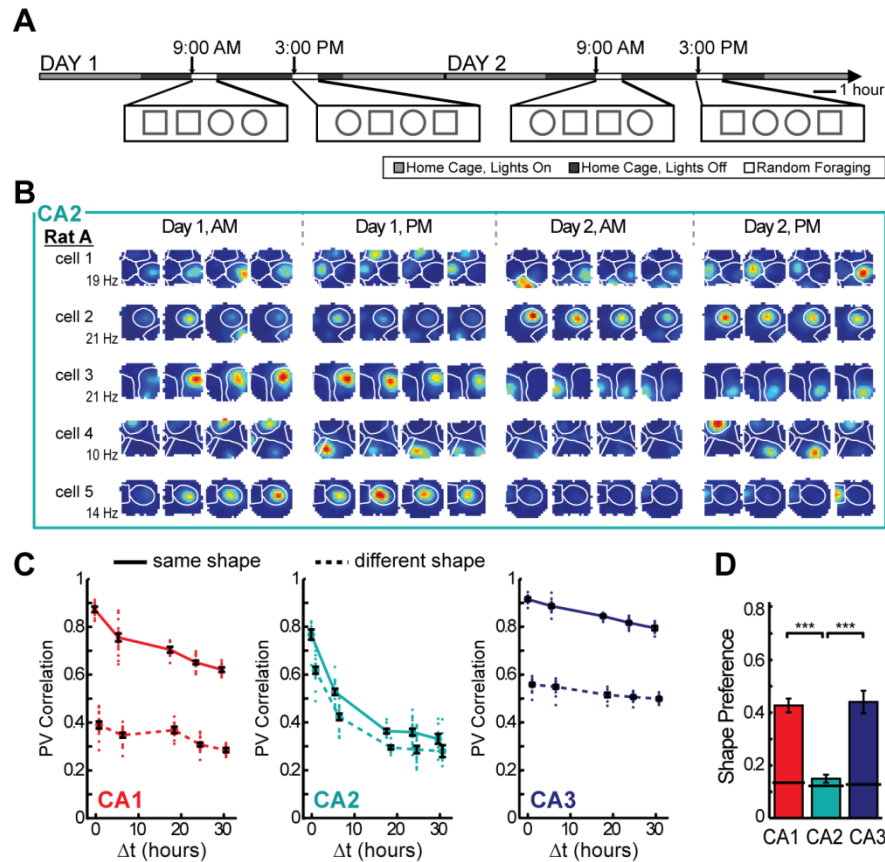
Figure 2.4. Of the three hippocampal CA areas, CA2 is the only one that shows more pronounced change over time than between spatial contexts. (**A**) To examine the effect of temporal distance on spatial firing patterns in CA2, we recorded CA2 ensembles across two days during four blocks of four 10-min sessions. (**B**) Spatial firing rate maps for five simultaneously recorded CA2 cells. Place field boundaries, calculated from the average of the 16 spatial maps for each cell, are superimposed in white. Note that individual place fields can be off for several sessions before reappearing at the same location and that the firing rates of individual place fields from single cells are modulated independently (see **Figure 2S.8**). (**C**) The population vector correlation was calculated between pairs of sessions of either the same or different shape, and the comparisons were grouped by the time interval between sessions (Δt). Each dot represents a pairwise comparison, and symbols and error bars represent the mean +/- SEM for each time lag. The mean correlations for same shape comparisons are connected by a solid line, while the mean correlations for comparisons between square and circle are connected by a dotted line. (**D**) Place fields in CA2 had lower shape preference scores than fields in CA1 or CA3, and their scores were not different from a shuffled distribution across all 16 sessions (solid black line). See **Table 2S.3** and text for detailed statistics, **Figure 2S.1** for cluster stability, **Figure 2S.5** for additional analysis, and **Figure 2S.6** for examples from each rat. *** $P < 0.001$.

calculated shape preference for each field across all 16 recording sessions, of which eight were in the square and eight were in the circular enclosure. CA2 showed much lower shape preferences than either CA1 or CA3 (Mann-Whitney *U*: CA1 vs. CA2, $z = 8.16$, $P < 0.001$; CA1 vs. CA3, $z = -0.23$, $P = 0.82$; CA2 vs. CA3, $z = -6.09$, $P < 0.001$), and the shape preference scores in CA2 were not significantly different than scores after randomly shuffling shape identity (only 67.9% of shuffled scores were lower than the actual mean score) (**Figure 2.4D**; see **Figure 2S.4** for individual examples). CA2 is therefore the only hippocampal subregion in which the population code more prominently differs between highly similar experiences at different time points than between different spatial contexts in close temporal proximity.

To determine whether the emergence of inconsistency in coding for the same box shape in CA2 required the intervening experience in a different box shape, we also performed recordings in a paradigm in which all 10-minute random foraging sessions were performed in the same box shape (n = 62 CA1 cells in 4 rats, 34 CA2 cells in 2 rats, and 70 CA3 cells in 2 rats; **Figure 2.5A-B**, **Table 2S.1**). In this paradigm, we found that the population vector correlation between sessions within a block was generally lower in CA2 than in the other subregions, but that the correlation decreased in all three hippocampal subregions with an increasing lag [Two-way ANOVA: brain region, $F(2) = 19.8$, $P < 0.001$; lag, $F(2) = 40.1$, $P < 0.001$; interaction, $F(4) = 1.73$; $P = 0.15$; Tukey's HSD between CA2 and CA1 and between CA2 and CA3, $P < 0.001$] (**Figure 2.5C**). The similar trend for CA1 and CA3 as for CA2 within a block of four recording sessions raised the question whether the more pronounced decline in CA2 compared to the other CA subregions, which we had observed in the two-shape paradigm, would at longer time intervals also emerge in the single-shape paradigm (**Figure 2.6A**). When comparing CA2 population vectors between blocks of recordings in a single shape, the similarity decreased as a function of the temporal distance between recording sessions for time intervals up to 18 h [ANOVA: $F(4) = $
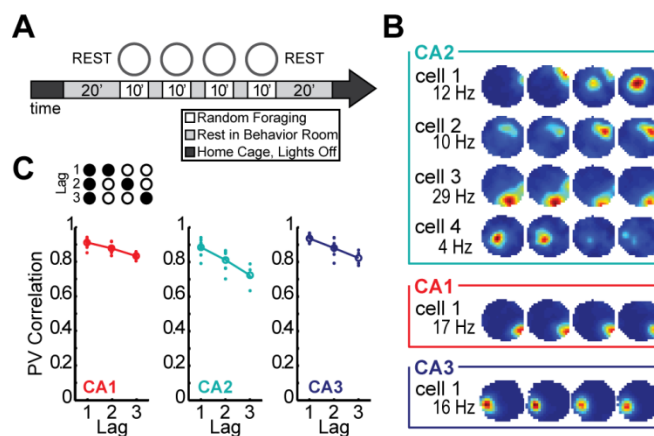
Figure 2.5. In the single-shape paradigm all hippocampal subregions are characterized by a short-term decrease in the correlation of population activity. (**A**) To determine whether the inconsistency in coding for the same box shape in CA2 required the intervening experience in a different box shape, we performed recordings in a paradigm in which all 10-min random foraging sessions were in the same box shape. (**B**) Spatial firing rate maps for four representative CA2 cells are shown and, for comparison, one representative CA1 and CA3 cell. Maps are color coded as described in **Figure 2.2A**. Variability in spatial firing patters of CA2 cells occurred despite the consistent repetition of the same box shape over time. (**C**) An increase in the temporal distance between sessions within a block was accompanied by a decrease in the PV correlation in all three subregions, but the overall PV correlation was lowest in CA2. Symbols and error bars represent the mean +/- SEM for each time lag. See text for statistics and **Figure 2S.7** for additional analysis.

202.6, $P < 0.001$; Tukey's HSD, P-values < 0.001 for all comparisons except comparisons between the 18, 24, and 30 hour time points were n.s.] (**Figures 2.6B**-**D**), and the decrease over time was more pronounced in CA2 than in either CA1 or CA3 [Two-way ANOVA: region, F(2) = 1204.2, $P < 0.001$; time difference, F(4) = 263.5, $P < 0.001$; interaction, F(8) = 79.5, $P < 0.001$; Tukey's HSD between brain regions, $P < 0.001$]. By charting the PV correlations between the first session of each block and all the other sessions within the two-day recording sequence, we could directly compare the population vector correlation within a block with the correlation between blocks (**Figure 2S.7**). All CA subregions showed a short-term decrease in their correlation within a block. Between blocks, the correlation reset to a higher value in CA3 while it typically continued to decrease in CA2. CA1 was intermediate between CA2 and CA3. The recordings with only a single shape therefore confirmed that CA2 is the hippocampal subregion in which the change of the population code over time periods of hours is most pronounced.

Differences in the CA2 population code over time may result from various sources of variability in the firing patterns, such as from a loss or gain of firing fields, from a drift in place field location, or from rate changes within single firing fields. These possibilities can be distinguished by measuring the number and location of place fields over different time periods (**Figure 2.7A**). We compared the number of active firing fields of CA2 cells with those of CA1 and CA3 cells when averaging over an increasing number of sessions (i.e. one session, a block of four sessions, the eight sessions in a single day, and sixteen sessions over two days). CA1 and CA3 showed no difference in the number of fields over different time periods, whereas in CA2, the mean number of place fields per cell increased for increasingly longer analysis periods [Two-way ANOVA: brain region, F(2) = 214.1, $P < 0.001$; time scale, F(3) = 13.7, $P < 0.001$; interaction: F(6) = 4.24, $P < 0.001$; Tukey's HSD between brain regions, P < 0.001] (**Figure 2.7B**). Furthermore, we found that place fields from a single cell modulated their firing rates independently (**Figure 2S.8**). Thus, the transient presence and independent modulation of
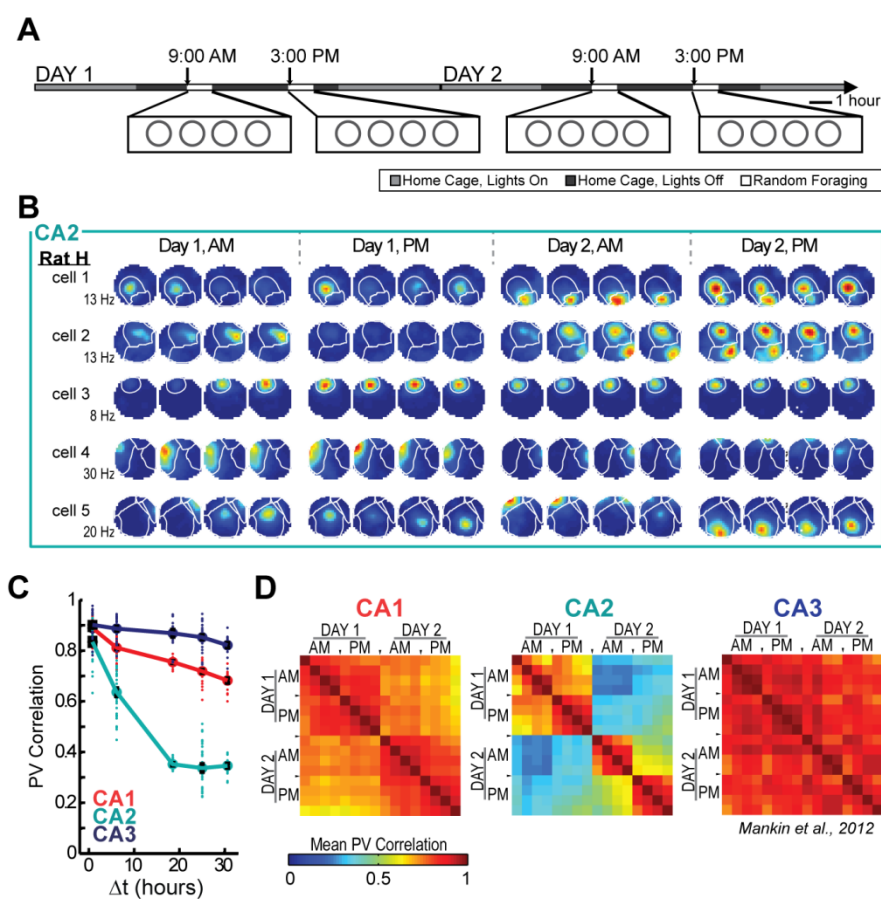
Figure 2.6. The change in population activity over extended time periods was most pronounced in CA2 even when box shape was held constant. (**A**) To test whether the change in CA2 representations required two spatial contexts or would also be observed during testing in a single context, we recorded cells in only a single environment shape over two days. (**B**) Spatial firing rate maps for five simultaneously recorded cells in CA2 during the single-shape behavioral paradigm, with place field boundaries superimposed (white lines). As in the two-shape paradigm, place fields in CA2 cells appeared, vanished, and could reappear. (**C**) In CA2, the decrease in PV correlation over time reached the same asymptotic level in the single-shape paradigm as in the two-shape paradigm, indicating that the change was predominantly a function of temporal distance and did not require switching between box shapes. In contrast, representations in CA3 have previously been shown to remain highly correlated over longer time intervals. In this paradigm, the CA2 and CA3 recordings are from different animals, and the CA3 recordings correspond to those reported in Mankin et al. (2012). Each dot is a pairwise comparison, and symbols and error bars represent the mean +/**-** SEM for each time interval. (**D**) Pairwise PV correlation matrices for repeated recordings in the same enclosure shape. Correlation matrices depict all possible comparisons between each of the sixteen recording sessions. Comparisons between the same sessions are shown along the diagonal, and their correlation coefficient is, by definition, 1. The lowest correlation coefficients were observed in the CA2 population for comparisons at intervals of at least 18 hours (see **Figure 2S.7** for additional plots).

each of the multiple firing fields of CA2 neurons is a source of the decorrelation within the CA2 network over time. In addition, changes in field locations could also cause decorrelation. To test this directly, we estimated the center of each place field in each session and traced the trajectory of the centers across the sixteen 10-min recording sessions in the single-shape experiment. CA2 place field centers drifted considerably more than those in CA1 and in CA3 (Mann-Whitney U: CA1 vs. CA2, $z = -4.61$, $P < 0.001$; CA1 vs. CA3, $z = 1.74$, $P = 0.073$; CA2 vs. CA3, $z = 4.52$, $P < 0.001$) (**Figure 2.7C**). We therefore detected that both spatial drift and firing rate variability were much higher in CA2 than in the other subregions over long time intervals but not over short time intervals. These findings indicate that a combination of loss or gain of firing fields, changes in firing rate, and moderate drift in precise firing location of CA2 place cells resulted in the substantial change in neuronal activity patterns in the CA2 network over time.

Discussion

The distinct connectivity, gene expression profiles, and cellular plasticity of CA2 (Caruana et al., 2012; Chevaleyre & Siegelbaum, 2010; Cui et al., 2013; Hitti & Siegelbaum, 2014; Kohara et al., 2014; Lein et al., 2005; Pagani et al., 2014; Rowland et al., 2013; Woodhams et al., 1993; M. Zhao et al., 2007) suggest that its neuronal computations differ from the other hippocampal subregions. To test for specialized neural network activity, we recorded activity patterns from CA2 cells in behaving animals. We found that standard spatial and temporal firing patterns in CA2 at the level of single neurons, such as place fields, theta modulation, and phase precession, are comparable to the other CA fields, with only minor quantitative differences. This confirms a previous report in which differences in location-selective firing between CA2 and CA1 cells were not detected (Martig & Mizumori, 2011). However, when comparing activity patterns during repeated visits to the same environment over extended time periods, major differences in the CA2 firing patterns emerged. Rather than faithfully coding for features of an
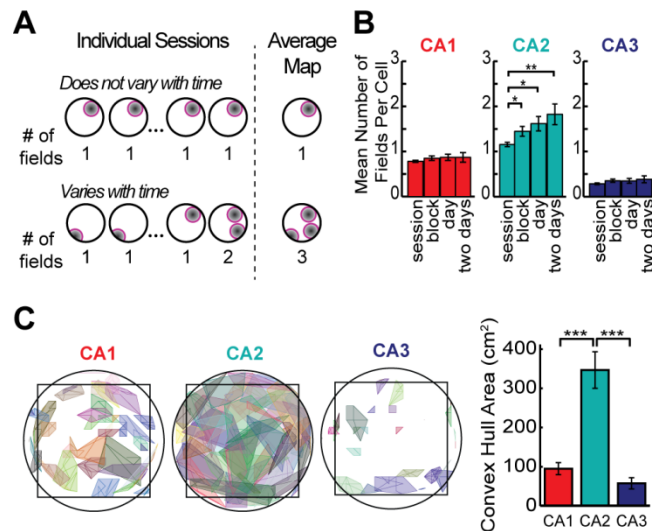
Figure 2.7. Dissimilarity in spatial firing patterns in CA2 emerges from transiently silent firing fields accompanied by a drift in the center of each place field location. (**A**) The schematic shows that the number of fields in the average firing rate map remains constant over a series of sessions when cells have a consistent place field, but that the number increases when the firing rate switches on and off at multiple place field locations. See **Figure2S.8** for additional analysis of firing rates. (**B**) Number of fields per cell in the single-shape paradigm after averaging over different time periods. In CA2, the number of fields per cell increased when including longer time periods, consistent with the idea that each cell can be transiently active at multiple firing locations (see **Figures 2.4, 2. 6,** and **2S.6** for examples). (**C**) To evaluate the degree to which firing within a place field was retained at a consistent location, place field boundaries were calculated from the average map (over 16 sessions). For each place field, the trajectory of the field center was then tracked across sessions, and the convex hull of the trajectory is shown. The bar graph to the right shows the average area of the fields' convex hulls, which was largest in CA2, indicating that the exact firing distribution within the field varied from session to session. See text for detailed statistics. * $P < 0.05$, ** $P < 0.01$, *** $P < 0.001$. Bars represent the mean +/- SEM.

environment and for differences between environments, as is characteristic of CA1 and CA3, our data show pronounced variability in the spatial firing patterns of CA2 cells over hours and days. The major time dependent differences were a consequence of the fact that CA2 cells could exhibit place fields at multiple locations, of which only a subset was selectively active at any point in time. The firing rates within each of the fields of a CA2 cell varied independently, and each field showed drift around its central firing position. Through these combined changes in the firing patterns of each cell, the population coding of CA2 changed over time, and the amount of change after 6 hours already exceeded the amount of change as a consequence of presenting different environments. Together with behavioral evidence that neither silencing of CA2 nor ablating vasopressin 1b receptors, which are selectively enriched in CA2, impacts performance on spatial and contextual tasks (DeVito et al., 2009; Hitti & Siegelbaum, 2014), this suggests that CA2 is less specialized for representing space and for distinguishing between spatial contexts than the other hippocampal subfields.

The generation of distinct neuronal codes for different contexts is a prominent feature of hippocampal neuronal activity patterns in CA1 and CA3 (Anderson & Jeffery, 2003; S. Leutgeb et al., 2004; Lever et al., 2002; Muller & Kubie, 1987; Vazdarjanova & Guzowski, 2004). A study using immediate early gene labeling as a marker for neuronal activity in the mouse recently reported that the CA2 region is similar to CA1 and CA3 in that it generates distinct neural codes for two different environments that are presented with an interval of 20 minutes (Wintzer, Boehringer, Polygalov, & McHugh, 2014). At temporal distances on the order of minutes and without any intervening experiences, we also found that two different contexts resulted in a more distinct CA2 firing pattern than a repetition of the same context (see **Figure 2.3E**). However, when we extended our recordings to longer time intervals, the changes in firing patterns with time were much more pronounced than the component of the decorrelation that was context dependent. These major changes over time rather than in response to distinct contexts are contrary to what

we observed in the same paradigm in the CA1 and CA3 networks, where network similarity for repeated presentations of the same environment, even over intervals of 30 hours, is higher than for distinct contexts at close temporal proximity (Mankin et al., 2012; see also **Figure 2.4C**).

The lower stability of CA2 firing patterns could originate from the unique connectivity and physiology of this hippocampal subregion. For example, long-term potentiation (LTP) at the synapses between CA3 and CA2 is not inducible by standard protocols in hippocampal slices while these synapses can be potentiated by neuropeptides (Caruana et al., 2012; Chevaleyre and Siegelbaum, 2010; Pagani et al., 2014; Zhao et al., 2007). In CA1, pharmacological blockade of LTP reduces place field stability while conditions that enhance LTP result in more stable CA1 place fields (C. G. Kentros, Agnihotri, Streater, Hawkins, & Kandel, 2004; C. Kentros, 1998). Thus, one source of place field instability in CA2 could potentially be the more limited LTP of inputs from CA3, and stability may increase by peptide release during behaviors that depend on vasopressin 1b receptor activation (DeVito et al., 2009; Pagani et al., 2014; Wersinger et al., 2002). However, it is currently unknown whether plasticity in CA2 can be modulated during behavior and, because CA3 inputs to CA2 are at baseline already weaker than entorhinal inputs (Chevaleyre & Siegelbaum, 2010), it is uncertain whether modulating plasticity at the CA3 inputs to CA2 would have major effects on CA2 firing patterns. Rather, from the findings that entorhinal inputs to CA2 are strong and that the resting membrane potential of CA2 cells is lower than in other hippocampal subregions (Chevaleyre & Siegelbaum, 2010; M. Zhao et al., 2007), it appears that CA2 activity may be more directly dependent on the convergence of inputs from entorhinal subdivisions. Medial entorhinal inputs to the hippocampus consist of grid cells, head direction cells, border cells, and nonspatial cells (Hafting, Fyhn, Molden, Moser, & Moser, 2005; Sargolini et al., 2006; Zhang et al., 2013) while lateral entorhinal inputs are generally less modulated by spatial features than those from the medial entorhinal cortex (Deshmukh & Knierim, 2011; Hargreaves, Rao, Lee, & Knierim, 2005; Tsao, Moser, & Moser, 2013). Furthermore, grid cells

were found to not be context selective (Fyhn, Hafting, Treves, Moser, & Moser, 2007). Taken together, this raises the possibility that the reduced context selectivity and high variability of CA2 firing patterns results from the convergence of spatial and nonspatial entorhinal inputs, which have not been processed by the dentate gyrus and/or CA3. We also observed that CA2 cells can become silent within a particular firing field to only later reemerge at the same location. This observation suggests a stable spatial input over time from either the entorhinal cortex or, alternatively, from CA3, which has weaker input to CA2 (Chevaleyre and Siegelbaum, 2010) but has previously been found to retain consistent spatial representations in the same experimental paradigm (Mankin et al., 2012).

The observation that there is a strong time-varying signal in the CA2 network compared to other hippocampal subregions raises questions about the function of neuronal firing patterns that vary over time within a brain structure that is required for long-term memory. It has been found that noise or variability over time can be used as a neural coding mechanism. For example, a time-varying signal in memory circuitry has been shown to be necessary in brain circuits for motor learning (Stepanek & Doupe, 2010; Wu, Miyamoto, Gonzalez Castro, Ölveczky, & Smith, 2014). Furthermore, findings in rats and human subjects demonstrate that a time-varying code in the hippocampus and medial temporal lobe can predict subjective estimates of elapsed time, as well as performance on temporal order and sequence memory tasks (Ezzyat & Davachi, 2014; Hsieh, Gruber, Jenkins, & Ranganath, 2014; Manning et al., 2011; Manns et al., 2007). These experiments demonstrate that neural drift on a time scale of up to minutes is informative and that gradually changing activity patterns in the hippocampus can be integrated into a neural code that contains memory for temporal context. A particularly clear manifestation of a temporal code are the recently discovered sequence and time cells in the hippocampus, which fire in a stereotyped order while animals are stationary over periods of up to tens of seconds during each delay period(MacDonald et al., 2011; Pastalkova et al., 2008). Here we find a pronounced gradual

change in CA2 ensemble activity over intervals of hours, but it remains to be determined whether neuronal firing patterns that fluctuate over this time scale could become repeated. Although there is no theoretical reason why temporal coding with repeated sequences would be limited to a particular time scale, it is likely that sequential neuronal activity on a much longer time scale would require different underlying cellular and circuit mechanisms than the sequential activation of CA1 cells over much shorter intervals. In contrast to a mechanism that relies on fixed sequences to be informative about elapsed time, it is also feasible that the time-varying neuronal firing patterns do not become informative by direct repetition during memory recall, but that it is rather a transition from changing to fixed neuronal firing patterns that supports memory, as has been suggested for neuronal activity in the mouse CA1 subregion (C. G. Kentros et al., 2004; M. E. Wang et al., 2012).

Alternatively, the CA2 cell population may contribute to memory coding neither by showing a sequence code nor by becoming stable, but by continuing to fluctuate and by thus providing a unique input pattern to CA1 at different time points. In this coding scheme, the variability over time in CA2 would be a prerequisite for providing temporal context, but it would not by itself constitute the temporal code. Rather, the unique inputs from CA2 would be associated with other stable inputs to CA1, such that CA1 activity patterns at one time differ somewhat from the activity patterns at a later time point. Such convergence of time-varying and stable inputs would provide a time-stamped neural code that differs between similar events at different times while it has higher overlap for events that occurred in close temporal proximity (Estes, 1955; Howard & Kahana, 2002; G.-J. Mensink & Raaijmakers, 1988). In support of such combinatorial coding, the neural population code in CA1 has previously been identified to gradually vary over intervals of hours to weeks (Mankin et al., 2012; Manns et al., 2007; Ziv et al., 2013) while also faithfully continuing to discriminate between spatial contexts (Mankin et al., 2012). Yet it has not been apparent how reliable representations of different environments could be retained in CA1 while

also allowing the network activity in the same cell population to drift over time. We now show that the dissimilarity in CA1 population activity over time is intermediate between CA2 and CA3, and it is known that CA1 receives strong excitatory inputs from both the CA3 and the CA2 subregion (Bartesaghi and Gessi, 2004; Bartesaghi et al., 2006; Chevaleyre and Siegelbaum, 2010; Kohara et al., 2014). This suggests that the CA1 network can integrate and/or compare the consistently precise information about spatial context it receives from CA3 with the slowly changing firing patterns we characterized in CA2 (**Figure 2.8**). The intermediate response of CA1 could thus indicate that the final processing stage of the hippocampus integrates information from not only CA3 and entorhinal cortex, but also from CA2 such that the stability of the CA1 firing patterns is dynamically regulated to determine the persistence and temporal context of hippocampal memory signals.

Materials and Methods

      Subjects and surgical procedures. Eight male Long-Evans rats (400-510 g) were implanted with a multitetrode drive assembly ('hyperdrive') aimed at the right hippocampus (AP 3.9-4.0 mm posterior to bregma, ML 3.0-3.5 mm). Tetrodes were prepared as described previously (Leutgeb et al., 2007) and were placed in the hippocampal cell layer using techniques optimized for recording stability across days (Mankin et al., 2012).

      Behavioral procedures. After one week of recovery from surgery, rats were partially food-deprived and trained to forage for randomly scattered cereal crumbs in an enclosure with walls that could be shaped either as a square (80 cm by 80 cm) or as a 16-sided polygon (50 cm radius; referred to as a 'circular enclosure'). Training was performed in two daily blocks. The first block started at approximately 9:00 am and the second block at approximately 3:00 pm. Rats were trained to run for four 10-minute sessions during each block, with two sessions in the square enclosure and two sessions in the circular enclosure, presented in random order. The recording
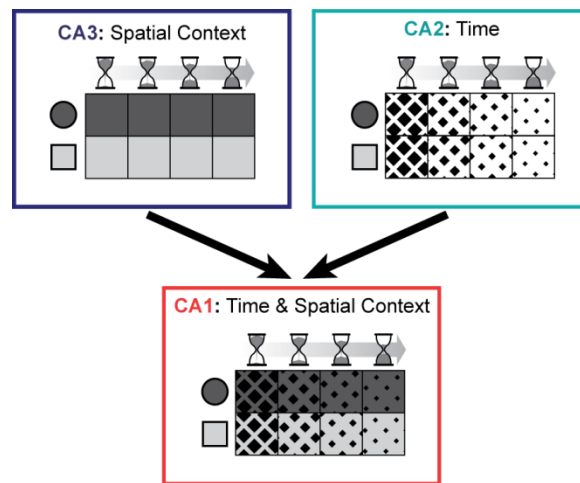
Figure 2.8. A schematic of the coding in CA1, CA2, and CA3 for context and space at different times and how inputs from CA2 and CA3 could be combined to jointly reflect this information in CA1. The two-by-four grid for each hippocampal subregion depicts a population representation for events at different times (left to right) and in different spatial contexts (top and bottom). Firing patterns in CA3 differ depending on context (shades of gray) and firing patterns in CA2 differ depending on elapsed time (diamond size). CA1 shows coding for both aspects, possibly by integrating or comparing inputs from the other hippocampal subregions.

phase of the experiment began after 9 to 20 days of behavioral training. Recordings were first conducted for 2 days in the standard training paradigm (referred to as two-shape, day 1 and day 2). Additionally, a subset of animals was tested in a paradigm in which all random foraging sessions were conducted in a single enclosure shape (single-shape, day 1 and day 2).

Cell-tracking. Because our study depended on tracking the same set of principal neurons over an extended time period, we developed a customized version of MClust (Redish, A.D., http://redishlab.neuroscience.umn.edu/MClust/MClust.html) with added functions that allowed for the comparison of the cluster boundaries of each cell throughout a series of recording sessions. Clusters that persisted in the same region of parameter space throughout two days were accepted as single cells for further analysis. Care was taken to accept only cells that could be precisely followed from the beginning to the end of the data analysis (**Figure 2S.1**).

Data Analysis. For tracked cells, we calculated spatial maps and identified place fields. For each place cell and field, we determined standard characteristics (e.g., mean rate, peak rate, spatial information, phase precession), and we analyzed the firing during individual passes through the place field. From the firing rate distribution within the place field in each enclosure shape, we derived a shape preference score (see **Figure 2S.4**). For the entire population of cells recorded in each subregion, we calculated all pairwise population vector correlations between 10-minute sessions, and grouped them by the elapsed time between sessions and by comparisons between either different shapes or the same shape.

Statistical Analysis. Comparisons between the firing characteristics of hippocampal subregions were performed using the Mann-Whitney *U* test. Holm-Bonferroni corrections for multiple comparisons were applied to the *P*-values. Comparisons between population vector correlations over different time intervals were performed using the Mann Whitney *U* test when there were two conditions and using ANOVA when there were three or more time intervals. If

comparisons were between multiple time intervals as well as between brain regions, two-way ANOVA was used. Tukey's HSD method was used for all post-hoc comparisons.

Histology. Tetrode locations were confirmed postmortem in histological material. Immunostaining for α-actinin-2 (i.e., a CA2 marker) (Ratzliff & Soltesz, 2001; Wyszynski et al., 1998) and cresyl violet were used to determine whether the final recording site for each tetrode was in or near the principal cell layers of the CA3, CA2, or CA1 subregion (see **Figure 2S.2**).

Detailed descriptions on cell tracking, data analysis, and histology can be found in **Appendix 2.3.**

Approvals. All experimental procedures were performed as approved by the Institutional Animal Care and Use Committee at the University of California, San Diego and according to National Institutes of Health and institutional guidelines.

Appendix 2.1: Supplemental Tables

Table 2S.1. Number of trackable cells per animal for each experimental condition. CA1 and CA3 data from rats A-F have been reported previously (Mankin et al., 2012). * Rat H performed the experimental sequence twice, separate sets of unique cells were recorded between the first and second set of experiments. The contribution of cells from each set of experiments is indicated in parentheses (first set/second set). ** Due to the small number of active CA3 place cells across the Third Day of experiments this data was not reported in the analysis.

	Two Shape			Single Shape			Third Day/Shift Day		
	CA1	CA2	CA3	CA1	CA2	CA3	CA1	CA2	CA3
Rat A	3	5	9	0	0	0	0	0	0
Rat B	0	7	16	0	0	0	0	0	0
Rat C	1	8	17	0	0	0	0	0	0
Rat D	30	0	10	0	0	0	17	0	8
Rat E	3	0	11	9	0	23	2	0	5
Rat F	13	0	8	16	0	47	11	0	5
Rat G	15	12	0	26	6	0	10	7	0
Rat H*	24 (16/8)	30 (9/21)	0	11 (5/6)	28 (6/22)	0	17 (9/8)	25(7/18)	0
Total	89	62	71	62	34	70	57	32	18**

Table 2S.2. Descriptive statistics for the spiking properties of CA1, CA2, and CA3 neurons during 10-minute random foraging sessions. All reported statistics were first calculated for each of the sixteen 10-minute sessions (see Figure 1D) and were then averaged across sessions to yield a single value per cell or per field. For measures reported in the larger table, comparisons between regions were performed using a Mann-Whitney U Test. The Holm-Bonferroni correction for multiple comparisons was applied to obtain adjusted P values. The measure reported in the smaller table compares two methods of calculating field size, and comparisons were performed within each region using a Mann-Whitney U Test.

Measure	N calculated on		Summary			Mann-Whitney Test					
			N	Mean	SEM		z-value	U	P	Adjusted P	Significance level
Mean firing rate (Hz)	Number of cells	CA1	89	0.53	0.048	CA1 vs CA2	-3.0	1977	0.0031	0.0062	P<0.01
		CA2	62	0.82	0.080	CA1 vs CA3	1.5	2720	0.13	0.13	n.s.
		CA3	71	0.52	0.078	CA2 vs CA3	3.4	1441	6.1 e-4	0.0018	P<0.01
Peak firing rate (Hz)	Number of cells	CA1	89	4.9	0.36	CA1 vs CA2	-0.30	2680	0.77	0.77	n.s.
		CA2	62	5.2	0.45	CA1 vs CA3	2.1	2551	0.037	0.087	n.s.
		CA3	71	4.2	0.53	CA2 vs CA3	2.2	1716	0.029	0.087	n.s.
Thresholded peak firing rate (Hz)	Number of cells with at least one session with peak rate > 2 Hz	CA1	76	7.5	0.34	CA1 vs CA2	0.85	1871	0.39	1	n.s.
		CA2	54	7.0	0.39	CA1 vs CA3	0.93	1712	0.35	1	n.s.
		CA3	50	7.4	0.58	CA2 vs CA3	0.22	1316	0.83	1	n.s.
Spatial information (bits/spike)	Number of cells	CA1	89	1.9	0.072	CA1 vs CA2	5.9	1211	4.8 e-9	1.4 e-8	P<0.001
		CA2	62	1.3	0.071	CA1 vs CA3	0.93	2888	0.35	0.35	n.s.
		CA3	71	1.8	0.080	CA2 vs CA3	-5.0	1102	7.2 e-7	1.4 e-6	P<0.001
Number of fields per cell	Number of cells	CA1	89	0.84	0.062	CA1 vs CA2	-2.4	2132	0.018	0.035	P<0.05
		CA2	62	1.1	0.096	CA1 vs CA3	2.3	2494	0.022	0.035	P<0.05
		CA3	71	0.68	0.082	CA2 vs CA3	3.7	1378	1.9 e-4	5.7 e-4	P<0.001
Number of fields per active cell	Number of cells with at least one field in at least one session	CA1	76	1.3	0.046	CA1 vs CA2	-2.8	1456	0.0044	0.0088	P<0.01
		CA2	54	1.5	0.071	CA1 vs CA3	1.5	1608	0.14	0.14	n.s.
		CA3	50	1.3	0.067	CA2 vs CA3	3.3	853	0.0011	0.0032	P<0.01
Field size (cm²)	Number of cells with at least one field in at least one session	CA1	76	803	49	CA1 vs CA2	-2.5	1515	0.011	0.034	P<0.05
		CA2	54	1000	64	CA1 vs CA3	0.38	1823	0.70	0.70	n.s.
		CA3	50	788	78	CA2 vs CA3	2.2	1008	0.026	0.053	n.s.
Field size (cm²), continuous half sessions	Number of cells with at least one field in at least one half-session	CA1	74	677	44	CA1 vs CA2	-2.3	1580	0.021	0.042	P<0.05
		CA2	56	839	56	CA1 vs CA3	1.9	1437	0.052	0.052	n.s.
		CA3	49	564	61	CA2 vs CA3	3.6	808	3.0 e-4	8.9 e-4	P<0.001
Fano factor of within-session firing rates per pass through field (Hz)	Number of place fields; analyzed in sessions with mean rate > 0.25 Hz	CA1	132	2.6	0.11	CA1 vs CA2	1.9	7590	0.057	0.11	n.s.
		CA2	133	2.3	0.093	CA1 vs CA3	2.8	3769	0.0044	0.013	P<0.05
		CA3	75	2.1	0.11	CA2 vs CA3	1.2	4488	0.23	0.23	n.s.
Within-session absolute change in firing rate	Number of place fields; analyzed in sessions with mean rate > 0.25 Hz	CA1	132	2.8	0.16	CA1 vs CA2	4.4	6033	1.1 e-5	0.3.3 e-5	P<0.001
		CA2	133	1.9	0.10	CA1 vs CA3	-0.76	4633	0.44	0.44	n.s.
		CA3	75	2.5	0.18	CA2 vs CA3	2.8	3803	0.0047	0.0093	P<0.01
Theta index	Number of cells with at least one session with more than 100 spikes	CA1	75	0.67	0.033	CA1 vs CA2	0.79	1858	0.43	0.85	n.s.
		CA2	54	0.66	0.050	CA1 vs CA3	1.3	1586	0.20	0.60	n.s.
		CA3	49	0.67	0.071	CA2 vs CA3	0.65	1224	0.52	0.85	n.s.
Intrinsic theta Frequency (Hz)	Number of cells with at least one session with more than 300 spikes	CA1	69	8.3	0.053	CA1 vs CA2	4.8	886	2.0 e-6	6.0 e-6	P<0.001
		CA2	52	7.7	0.10	CA1 vs CA3	4.6	567	5.3 e-6	1.1 e-5	P<0.001
		CA3	36	7.8	0.086	CA2 vs CA3	-0.36	892	0.72	0.72	n.s.
Phase-distance Slope (cycles/field)	Number of cells with at least one field with more than 100 spikes and a well-fit phase-distance line	CA1	68	-0.42	0.020	CA1 vs CA2	-5.0	851	6.8 e-7	2.1 e-6	P<0.001
		CA2	53	-0.27	0.023	CA1 vs CA3	-2.5	830	0.012	0.025	P<0.05
		CA3	35	-0.36	0.027	CA2 vs CA3	1.8	711	0.066	0.066	n.s.

Measure	N calculated on		Continuous Halves			Interleaved Halves			Continuous vs Interleaved			
			N	Mean	SEM	N	Mean	SEM	z-value	U	P	Significance level
Field size (cm²) for half sessions	Number of cells with at least one field in at least one half-session	CA1	74	677	44	77	624	36	0.38	2747	0.71	n.s.
		CA2	56	839	56	55	762	45	0.78	1408	0.44	n.s.
		CA3	49	564	61	52	602	54	-0.64	1179	0.52	n.s.

Table 2S.3. Full statistics for all comparisons using a Mann-Whitney U Test in analysis over time scales longer than 10 minutes. The first segment of the table reports comparisons between regions. The Holm-Bonferroni correction for multiple comparisons was applied to obtain adjusted P values. The second portion of the table compares PV correlations at Lag 2 between sessions recorded in the same box shape and sessions recorded in different box shapes; comparisons were performed within each region. Gray boxes indicate tests for which the sample size was too small to calculate a z-value.

Measure	N calculated on		Summary			Mann-Whitney U Test					
			N	Mean	SEM		z-value	U	P	Adjusted P	Significance level
Shape preference (within single block)	Place fields with at least 5 passes through the field in each of the 16 sessions	CA1	105	0.50	0.023	CA1 vs CA2	7.6	2372	3.2E-14	9.7E-14	P < 0.001
		CA2	112	0.27	0.016	CA1 vs CA3	0.62	2312	0.54	0.54	n.s.
		CA3	47	0.49	0.041	CA2 vs CA3	-4.5	1432	6.0E-06	1.2E-05	P < 0.001
PV corr at lag 1 (same box shape)	Comparisons between population vectors with at least 25 cells	CA1	7	0.90	0.0087	CA1 vs CA2		4	0.030	0.091	n.s.
		CA2	5	0.85	0.014	CA1 vs CA3		26	0.87	0.87	n.s.
		CA3	8	0.87	0.031	CA2 vs CA3		13	0.35	0.71	n.s.
PV corr at lag 2 (same box shape)	Comparisons between population vectors with at least 25 cells	CA1	6	0.85	0.019	CA1 vs CA2		1	0.019	0.038	P < 0.05
		CA2	4	0.67	0.052	CA1 vs CA3		0	0.010	0.029	P < 0.05
		CA3	4	0.92	0.011	CA2 vs CA3		0	0.029	0.038	P < 0.05
Shape preference (across two days)	Place fields with at least 5 passes through the field in each of the 16 sessions	CA1	105	0.43	0.026	CA1 vs CA2	8.2	2106	3.2E-16	9.3E-16	P < 0.001
		CA2	112	0.15	0.015	CA1 vs CA3	-0.23	2409	0.82	0.82	n.s.
		CA3	47	0.44	0.043	CA2 vs CA3	-6.1	1017	1.1E-09	2.2E-09	P < 0.001
Area covered by spatial drift of place field center	Place fields in the single-shape paradigm that were active in at least 3 sessions	CA1	51	95	16	CA1 vs CA2	-4.7	743	3.3E-06	9.8E-06	P < 0.001
		CA2	61	350	47	CA1 vs CA3	1.8	475	0.073	0.073	n.s.
		CA3	25	57	15	CA2 vs CA3	4.5	278	5.4E-06	1.1E-05	P < 0.001

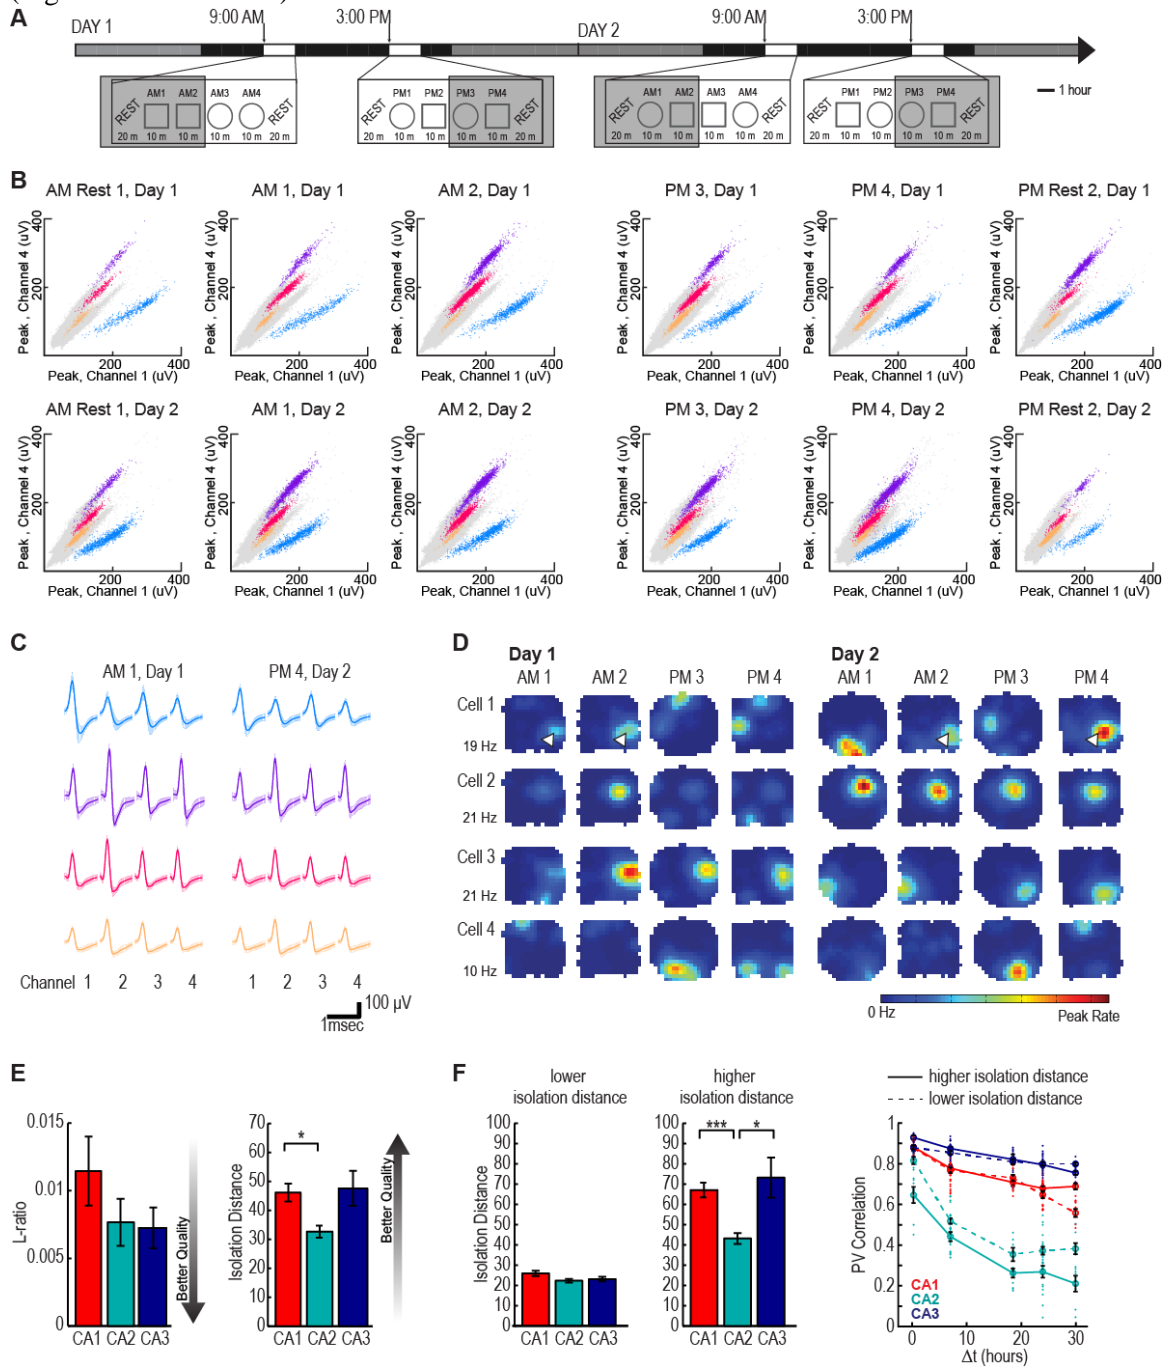| Measure	N calculated on		Same Box Shape			Different Box Shape			Same vs Different			
			N	Mean	SEM	N	Mean	SEM	z-value	U	P	Significance level
PV correlation at lag 2	Comparisons between population vectors with at least 25 cells	CA1	6	0.85	0.019	10	0.43	0.030		0	2.5 e-4	P < 0.001
		CA2	4	0.67	0.052	8	0.58	0.039		10	0.37	n.s.
		CA3	4	0.92	0.011	8	0.55	0.054		0	0.0040	P < 0.001

Appendix 2.2: Supplemental Figures

Figure 2S.1. Changes in CA2 activity are not a result of tetrode recording instability over time. (A) Complete experimental timeline. The gray boxes mark sessions that are highlighted in (B-D). (B) In each panel, the peak spike amplitude recorded on one channel of a tetrode is plotted against the peak spike amplitude recorded on another channel of the same tetrode. The same two tetrode channels are used for each plot, and different plots correspond to different time points in the experimental timeline as indicated by the gray boxes in (A). Each dot represents one sampled spike, and each dot color represents spikes that were assigned to a cluster. In amplitude plots, clusters are known to be generated by the regularity of the amplitude distribution of extracellular spikes from single cells. After defining clusters for all the recording sessions from one day, the cluster boundaries were applied to data recorded from the same tetrode on subsequent days. If clusters persisted in the same region of parameter space where they had been identified on the previous day, boundaries were adjusted to assure that all spikes that were determined to correspond to a cell were included within the boundaries. If the spikes for the set of clusters of a tetrode could be included with minor adjustments of the boundaries from the day before, the clusters were considered to correspond to the same set of cells as on the previous day, and the cells were included in analysis. Because we occasionally found that cells appeared to turn their spiking activity on or off during a subset of the recorded sessions, it was imperative to confirm that these effects were not simply a consequence of instability in the tetrode location. We thus only included cells in the analysis for which clusters could be identified within the same amplitude space at the beginning of the day (in the first sleep session or the first behavioral session in either shape) and at the end of the day (during the last sleep session or the last behavioral session in either shape). The cluster plots that are depicted correspond to cells 1-4 in Figure 2.4B. (C) Average waveforms from the clusters in (B). The color of the waveform matches the color of the corresponding spike cluster. The mean waveform (± standard deviation) is shown for the first behavioral session on the first day (left) and the last behavioral session on the second day (right). (D) Spatial maps from the same cells as in (B). The average firing rate in each 5×5 cm pixel is color-coded with a color scale from 0 Hz (blue) to the peak rate within the day (red). The peak firing rate for each cell is indicated to the left of the maps. Although spatial maps from tracked cells were not used to determine whether clusters were the same from one day to the next, we visually inspected spatial maps from tracked cells. In CA2, place fields frequently changed over the course of two days, but also reappeared at former locations after extended time intervals. For example, white arrowheads on the maps of Cell 1 indicate one location where a place field was often but not consistently present, including during the first and last behavioral session. This can be taken as further confirmation that the same cell was recorded across both days, even though a different spatial firing pattern appeared during intervening periods. (E) For each cell with more than 13 spikes per block, we measured cluster quality with two metrics: L-Ratio and isolation distance. Lower L-Ratios indicate better cluster quality, while higher isolation distances indicate better cluster quality. There was no statistical difference between regions for L-Ratio [Mann-Whitney U: CA1 (n=71 cells) vs. CA2 (n = 50 cells): z = 0.67, P = 1; CA1 vs. CA3 (49 cells): z = 0.74, P = 1; CA2 vs. CA3: z = 0.13, P = 1]. CA2 had moderately lower mean isolation distance compared to CA1 [Mann-Whitney U: CA1 (n=71 cells) vs. CA2 (n = 50 cells): z = 2.54, P = 0.033; CA1 vs. CA3 (n = 49 cells): z = 1.09, P = 0.52; CA2 vs. CA3: z = -1.12, P = 0.52]. (F) (Left) To confirm that differences in population coding between hippocampal subregions could not be attributed to differences in the isolation distance between spike clusters, we separated the cell sample into two groups based on whether each cell's isolation distance was greater than or less than the median isolation distance in that subregion (CA1, 40.6; CA2, 28.9; CA3, 34.5).

Figure 2S.1: Changes in CA2 activity are not a result of tetrode recording instability over time, continued

There was no difference between regions when comparing the mean isolation distance for cell clusters in the lower isolation distance group [Mann-Whitney U: CA1 (n=36 cells) vs. CA2 (n=25 cells): z = 1.49, P = 0.41; CA1 vs. CA3 (n = 25 cells), z = 1.40, P = 0.41; CA2 vs. CA3: z = -0.33, P = 0.74], while CA2 had lower mean isolation distances than the other regions when comparing cells included in the higher isolation distance group [Mann-Whitney U: CA1 (n = 35 cells) vs. CA2 (n = 25 cells): z = 4.53, P < 0.001; CA1 vs. CA3 (n = 24 cells) z = 1.10, P = 0.27; CA2 vs. CA3: z = -2.55, P = 0.022]. These results indicate that CA2 clusters were not considerably less isolated than those in the other subregions, but rather that CA2 had fewer exceptionally well-separated clusters. (Right) Population vector correlations for the same box shape across time (as in Figure 2.4C) are plotted for cells from each group of cluster quality (high quality, connected by solid lines; lower quality, connected by dashed lines) for each region. The PV correlation declined to the greatest extent across time in CA2 regardless of whether we included only the cells with the lowest isolation distances or only the cells with the highest isolation distances [two-way ANOVA for lower isolation distance: brain region: $F_{(2)} = 473.2$, P < 0.001; elapsed time: $F_{(4)} = 107.2$, P < 0.001; interaction: $F_{(8)} = 21.1$, P < 0.001; Tukey's HSD for all comparisons between regions, P < 0.001; two-way ANOVA for higher isolation distance: brain region: $F_{(2)} = 631.4$, P < 0.001; elapsed time: $F_{(4)} = 67.1$, P < 0.001; interaction: $F_{(8)} = 6.9$, P < 0.001; Tukey's HSD for all comparisons between regions, P < 0.001]. Although the two curves for CA2 differ [two-way ANOVA: cluster quality group, $F_{(1)} = 48.4$, P < 0.001; elapsed time, $F_{(4)} = 78.1$, P < 0.001; interaction, $F_{(4)} = 1.3$, P = 0.28], PV correlation values were lower in the higher isolation distance group than in the lower isolation distance group. Taken together, these findings indicate that the larger change in CA2 population activity over time could not be attributed to poorer cluster quality.

(Figure 2S.1 continued)

A DAY 1 · 9:00 AM · 3:00 PM · DAY 2 · 9:00 AM · 3:00 PM — 1 hour

B AM Rest 1, Day 1 · AM 1, Day 1 · AM 2, Day 1 · PM 3, Day 1 · PM 4, Day 1 · PM Rest 2, Day 1

AM Rest 1, Day 2 · AM 1, Day 2 · AM 2, Day 2 · PM 3, Day 2 · PM 4, Day 2 · PM Rest 2, Day 2

C AM 1, Day 1 · PM 4, Day 2

D Day 1 · Day 2

E L-ratio · Isolation Distance

F lower isolation distance · higher isolation distance · PV Correlation

Figure S2. Confirmation of recording sites in each of the hippocampal subregions by using cytoarchitectural and immunohistochemical criteria. (A) Sections were immunostained for α-actinin-2 and counterstained with cresyl violet to compare the boundaries that were defined by each of the methods. Example images were taken along the hippocampal cell layer (20x, scale bar = 200 μm) and were stitched together throughout the region of interest to determine the oundaries between subregions (top: CA1 to CA2; bottom: CA2 to CA3). Through the region of interest elineated by dashed lines, additional images of the CA1 to CA2 cell border are shown at higher magnification (63x, scale bar = 50 μm). The CA2 area (demarcated by teal lines) was defined as the region with the most dense α-actinin-2 immunostaining (brown; see Wyszynski et al., 1998) of cell bodies and proximal dendrites (teal arrows). CA1 and CA3 (demarcated by white lines) were defined as the adjacent areas with weak immunolabeling. In addition, CA1 is characterized by smaller, densely packed cell bodies while CA3 shows larger cell bodies and a large fraction of cells outside of a densely packed inner sublayer. A transition zone between CA1 and CA2 (between the white and teal line) was defined as the area in which the cells were almost as densely packed as in CA1 and in which there was only sparse immunolabeling of cell bodies and dendrites. A transition zone between CA3 and CA2 (between the teal and white line) was defined as the area with larger cell bodies, but few cell bodies outside of the densely packed inner sublayer. (B) Cytoarchitectural criteria that matched the criteria from the α-actinin-2/cresyl violet co-staining were used in cresyl violet stained tissue to identify the hippocampal subregions and the transition zones. Only recordings from tetrode tracks (circled in color) that could be confidently assigned to a hippocampal subregion outside of transition zones were used in the analysis, and those in transition zones were excluded (see Figure 1B for the final recording locations of cells included in analysis).

A

α-actinin-2 Immunohistochemistry + Cresyl Violet
RAT G:

B

Cresyl Violet
RAT C:
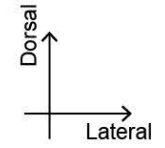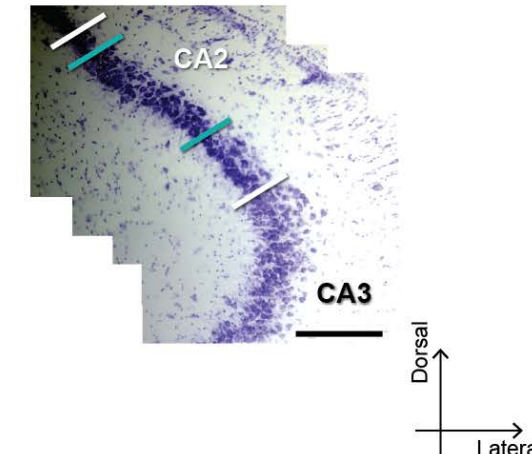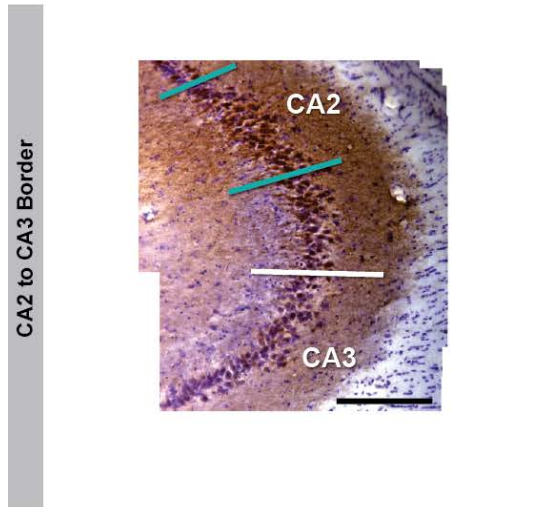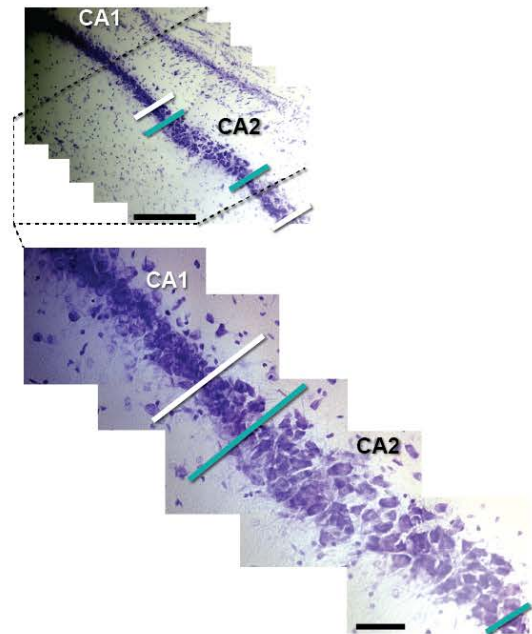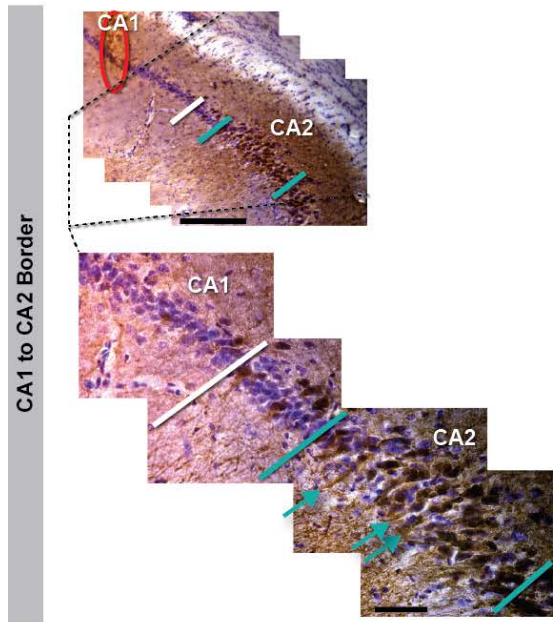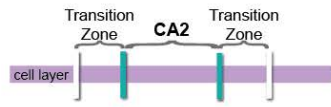
Figure 2S.3. Examples of phase precession in CA2 place fields. (A) For each rat, the activity pattern is shown for 2-3 CA2 cells. (Upper left of each panel) Color-coded firing rate map (0 Hz, blue; peak rate, red) with the field boundary superimposed in white. The cell number and session number correspond to cell numbers and sessions numbers in Figure 2S.6. (Lower left) Spikes within the place field (red dots) superimposed on the rat's trajectories through the field (gray lines). The field boundary is superimposed in black. (Right panel) Scatter plot of the normalized run distance against the theta phase at which each spike occurred. For better visualization, each spike is replotted in a second cycle with the beginning and end of each plotted cycle (0, 1, and 2) representing the trough of the theta cycle recorded in the hippocampal fissure. The regression line is superimposed in blue when the slope is significantly different from 0 (P < 0.05) and in gray when n.s., and its slope (in theta cycles/normalized run distance) is reported above each scatterplot. As common for the firing patterns of CA2 cells (Figure 2.2), Cell 2 from Rat H fired action potentials in multiple spatial locations within the environment. Phase precession was found in each of the three fields. Each regression line has a y-intercept between 1.0 and 1.1, which indicates that the phase precession cycle did not continue between fields but rather began at the boundary of each field. (B) For comparison, one example cell per rat is shown from either CA3 (rats A-C) or CA1 (rats G-H). Panels are arranged as described in A. Cell number is not indicated, as these cells do not correspond to those in other supplemental figures.
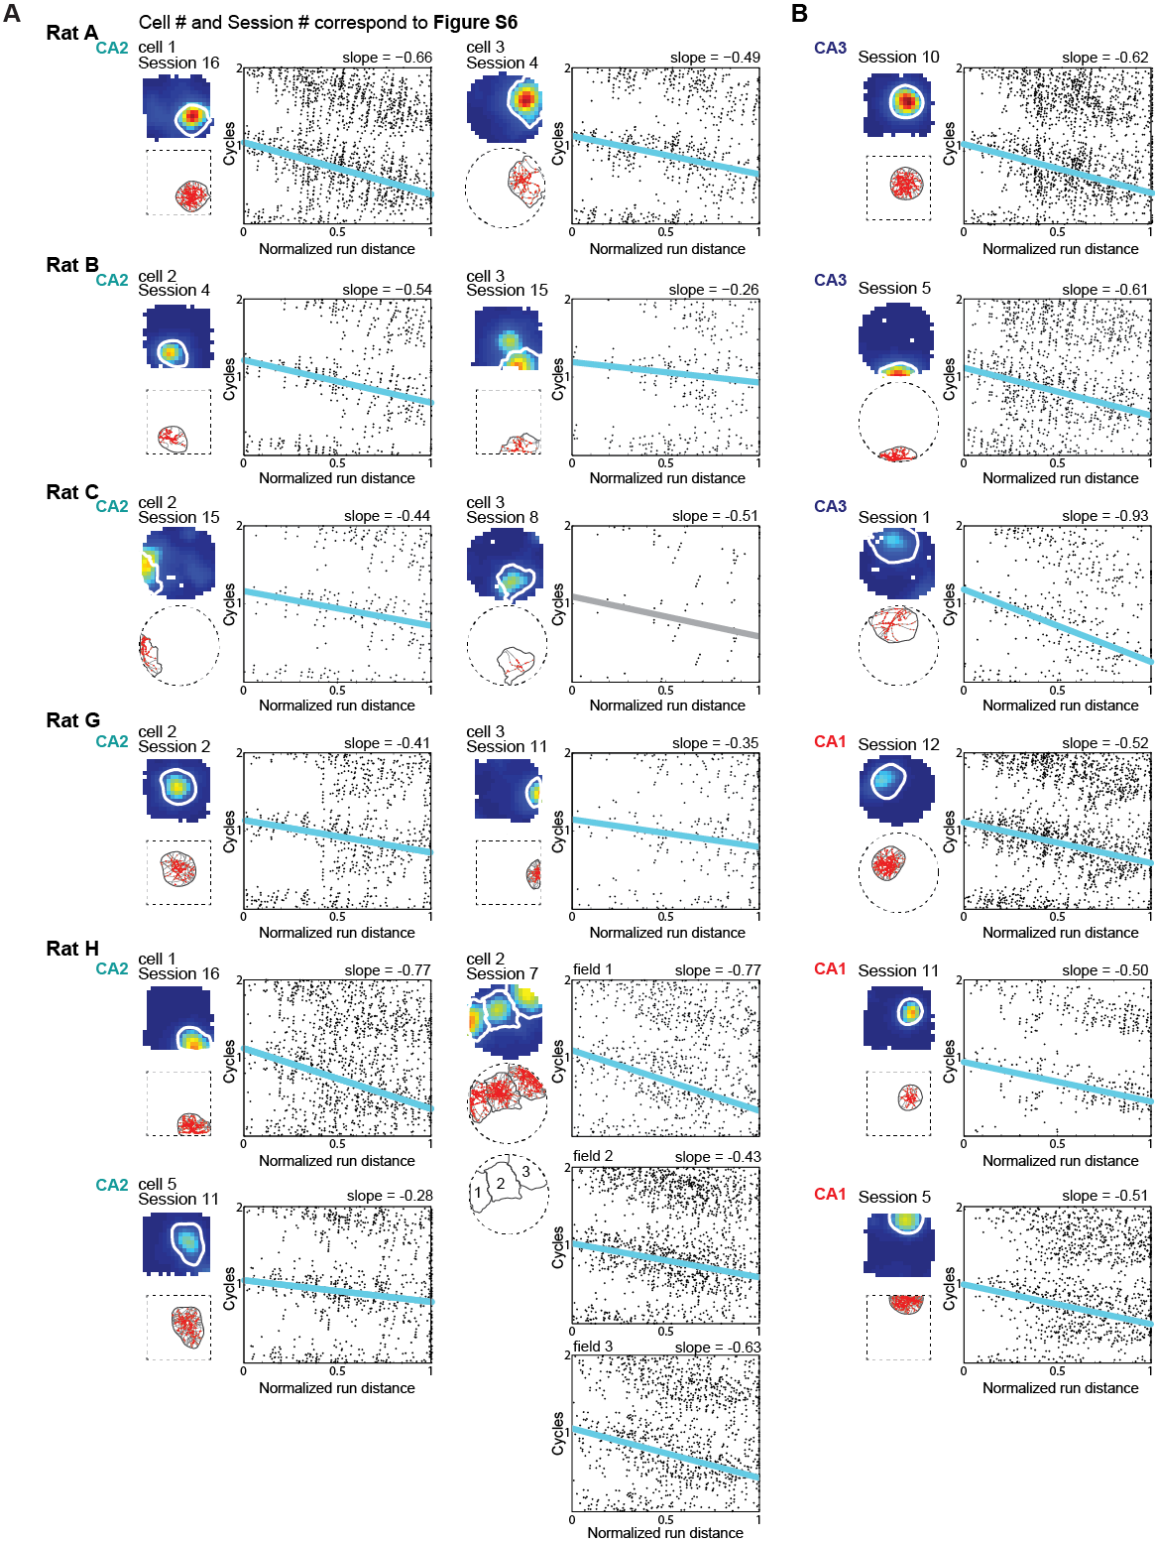
Figure 2S.4. CA2 cells with low shape preference scores are found in all rats. (A-D) Calculation of shape-preference scores. (A) The activity pattern of two example cells (top, CA1; bottom, CA3) are shown for the 16 sessions of the two-day different-shape experiment. For each cell, color-coded firing rate maps (0 Hz, dark blue; peak rate, red) are shown for each 10-minute random foraging session, and the place field boundaries that were calculated from the average of the 16 spatial maps for each cell are superimposed in white. Below each rate map, the firing rate for each pass of the rat through the place field is shown in purple when the enclosure was in the circle configuration and in green when the enclosure was in the square configuration. The running velocity during each pass is plotted downwards in light blue. The presence of a velocity bar without a corresponding firing rate bar indicates passes without spikes. The CA1 cell was silent during the morning of the first day and became active in subsequent sessions, but only when the enclosure was shaped as a square. The CA3 cell exhibited higher firing rates in the circular shape than in the square shape throughout the entire recording sequence. Each cell demonstrated a clear firing rate preference for one of the box shapes, which resulted in a high shape preference score as described in B-D. (B) The firing rate distribution for passes through the place field when the box was in a square configuration (left, green) and for passes when the box was in a circle configuration (right, purple). The horizontal axes begin at -2 Hz in order to clearly visualize values at 0. (C) Cumulative distribution plots for the firing rates in B. Lines that are farther to the right indicate distributions in which higher rates were more frequent. (D) For each pair of firing rate distributions in C, the Receiver Operant Characteristic (ROC) curve was calculated (thick gray line) and the difference between the area under the ROC curve and area under the identity line was calculated (shaded in gray). This value was multiplied by 2, which results in a shape preference score that equals -1 for firing in only the circle, 1 for firing in only the square, and 0 for equal firing in the square and circle. The absolute value of the score measures the magnitude of the preference for either shape. The absolute shape preference score for the example cells is indicated in bold script. The same procedure was followed to calculate within-block absolute shape preference scores, except that only firing rates from sessions within a single block were used and that the scores for each of the four blocks were averaged. (E) Firing rate histograms, firing rate distributions, and shape preference scores for three CA2 place fields from each rat (A-G) and for six place fields from rat H. The firing rates correspond to those in the fields in Figure S6 where the rate maps for all 16 sessions are depicted. The color scheme and layout of all plots is as described in A-D. Note that CA2 fields frequently have similar CDFs for each of the shapes, yielding shape preference scores close to 0.

A

CA1
cell 1

AM 1   PM 1   AM 2   PM 2

Firing Rate (Hz)

velocity (cm/sec)

CA3
cell 1

Firing Rate (Hz)

velocity (cm/sec)

B

Square   Circle

Number of Passes

Firing Rate per Pass

Square   Circle

Number of Passes

Firing Rate per Pass

C

CDFs

Firing Rate

CDFs

Firing Rate

D

ROC curve

Shape Pref = 0.64

ROC curve

Shape Pref = 0.86

E

CA2

Rat A

cell 1   AM 1 PM 1 AM 2 PM 2

FR   vel

Shape Pref = 0.16

cell 2

FR   vel

Shape Pref = 0.12

cell 3

FR   vel

Shape Pref = 0.04

Rat B

cell 1

FR   vel

Shape Pref = 0.03

cell 2

FR   vel

Shape Pref = 0.04

cell 3

FR   vel

Shape Pref = 0.04

Rat C

cell 1

FR   vel

Shape Pref = 0.07

cell 2

FR   vel

Shape Pref = 0.11

cell 3

FR   vel

Shape Pref = 0.18

Rat G

cell 1

FR   vel

Shape Pref = 0.11

cell 2

FR   vel

Shape Pref = 0.24

cell 3

FR   vel

**Shape Pref = 0.32**

(Figure 2S.4 continued)

Figure 2S.5. Additional analysis of the similarity in hippocampal firing patterns over time. (A) To confirm that an asymptote had been reached in CA2 after approximately one day, we extended hippocampal recordings across three days (see Supplemental Experimental Procedures). On the third day, all behavioral sessions were shifted by six hours, which makes comparisons for up to 60 hours possible for the population of cells that were tracked across all three days (n = 57 CA1 cells from 5 rats; n = 32 CA2 cells from 2 of 5 rats). (B) Population vector correlations between pairs of recordings in the same enclosure shape are shown as dots. Each dot corresponds to a pairwise comparison at the interval that is indicated on the horizontal axis. Symbols and error bars correspond to the mean +/- SEM for all comparisons at a particular time interval. Solid black lines represent best exponential fits to the data, and dotted black lines represent best linear fits. The explained variance was higher for the exponential compared to the linear fits in CA1 (R^2 exponential = 0.83, and R^2 linear = 0.81) and CA2 (R^2 exponential = 0.84, R^2 linear = 0.73). For the exponential fit, the time constant for the decay of the PV correlation is 76.3 hours in CA1 and 20.96 hours in CA2. (C) A two-way ANOVA on brain region and elapsed time revealed main effects of brain region [F(1) = 662.1, P < 0.001] and elapsed time [F(8) = 172.7, P < 0.001], as well as a significant interaction [F(8) = 9.97; P < 0.001] indicating that the way in which correlation changed with time varied between regions. For each region, a grid of Tukey's HSD comparisons is shown, and darker colors indicate more highly significant differences. Correlation levels in CA2 reached an asymptote by 30 hours. The slower change in CA1 compared to CA2 could be explained by proposing that CA1 integrates the continually changing signal from CA2 with a stable signal from CA3, which would presumably result in an intermediate level of place field stability in CA1 (see Figure 2.8). (D-E) Hippocampal CA2 and CA1 population patterns become decorrelated over time irrespective of intervening experience. To determine the possible role of intervening blocks on the population vector decorrelation, we analyzed hippocampal recordings on the third recording day. The 6-hour shift on the third day resulted in blocks between the second and third day that are separated by 24 hours without intervening behavioral testing and by 30 hours with one intervening block. (D) Population vector (PV) correlations for comparisons between recordings in the same enclosure shape at intervals of 24 hours or 30 hours (left, CA1; right, CA2). Each comparison between two sessions is shown as a dot, and the symbols and error bars are the mean +/- SEM for a time interval. Different numbers of intervening blocks did not result in differences in the degree of decorrelation [comparisons between 0 and 1 intervening blocks at the 24-hour interval: CA1, t(22) = -2.11, P = 0.094; CA2, t(22) = -1.07, P = 0.60; comparisons between 1 and 2 intervening block at the 30-hour intervals: CA1, t(22) = -0.49, P = 0.63; CA2, t(22) = -0.052, P = 0.96]. (E) PV correlations between pairs of recordings in the same enclosure shape. Only comparisons without intervening recording blocks are included. Each comparison between two sessions is shown as a dot, and the symbols and error bars are the mean +/- SEM for a time interval. A stronger decorrelation of CA2 population activity compared to CA1 population activity with time was also observed when there were no intervening recording blocks [Two-way ANOVA: brain region, F(1) = 203.0, P < 0.001; time difference, F(3) = 125.2, P < 0.001; interaction, F(3) = 10.9, P < 0.001]. (F) The population of neurons that were recorded in CA2 over three days reached asymptote after 30 hours (see B and C) compared to 18 hours in the larger dataset in Figure 4C. To test whether this was a result of including a behavioral session that was shifted by 6 hours on Day 3, we analyzed the same population of cells that was used for the analysis in (B) and (C) but only included comparisons for sessions that were recorded on Days 1 and 2. When comparing both data sets, the same level of decorrelation was reached at 30 hours. This result indicates that the variability in the time interval when asymptote was reached in CA2 was due to selecting a different set of animals for the analysis rather than due to selecting a different set of recording days for the analysis. (G) To further test the variability between individual animals in the two-day, two-shape condition

(Figure 2S.5: Additional analysis of the similarity in hippocampal firing patterns over time, continued)

(Figure 2.4), we performed a jackknife procedure in which PV correlations were calculated for sets of cells while leaving out one animal at a time. The mean correlation values for these reduced datasets are shown. In each case, there was a much more substantial decrease in coding similarity with time in CA2 compared to CA1 and CA3, although the decorrelation levels in CA2 after one day (18 to 30 hour intervals) varied from approximately 0.5 to 0.3. Despite the variability in the exact time course and decorrelation level, the full effect size in CA2 was reached within approximately one day (18-30 hrs) for all data sets. (H) Because one rat performed the experimental sequence twice with two different sets of independent cells, we also checked whether the observed results would differ substantially if the rat had been excluded (solid lines), or if the set of cells from one or the other repetition had been removed (dashed lines). In each case the decorrelation progressed much more rapidly in CA2 than in CA1, and reached a level between 0.3 and 0.4 in CA2 by 30 hours. (I) Finally, we only included CA1 and CA3 ensembles that were recorded simultaneously with CA2 ensembles in the analysis and found a pronounced difference in the rate of decorrelation between CA2 and each of the two other hippocampal subregions [Two-way ANOVA: brain region, $F_{(2)} = 1251.0$, $P < 0.001$; time difference, $F_{(4)} = 167.3$, $P < 0.001$, interaction, $F_{(8)} = 29.6$, $P < 0.001$; Tukey's HSD between brain regions, $P < 0.001$ for each comparison]. (J) Firing rates in CA2 did not show circadian fluctuations but varied with recording blocks. The mean +/- SEM of the normalized firing rate for all active CA1, CA2, and CA3 principal neurons is shown for each ten-minute recording session across the two-day, two-shape experiment. For each cell, the normalized firing rate was calculated by dividing the average firing rate for each session by that cell's average maximum firing rate throughout all 16 sessions. An ANOVA with day and session as factors did not reveal differences between Day 1 and Day 2 for any of the subregions, which is expected because the recording environments were already highly familiar on the first recording day [CA1: $F_{(1)} = 1.56$, $P = 0.21$; CA2: $F_{(1)} = 0.061$, $P = 0.80$; CA3: $F_{(1)} = 0.035$, $P = 0.85$]. An effect of session identity was not found in CA3 [CA3: $F_{(7)} = 0.36$, $P = 0.92$], was close to significance in CA1 [CA1: $F_{(7)} = 1.97$, $P = 0.055$], and was found in CA2 [$F_{(7)} = 2.41$, $P = 0.019$, Tukey's HSD revealed that the significant differences ($P < 0.05$) were between session 5 and sessions 4 and 8]. This result suggests that the effect was not circadian but that there were firing rate differences between the beginning and the end of recording blocks. An ANOVA with block and session within a block as factors confirmed that normalized firing rates did not vary between blocks in any region [CA1: $F_{(3)} = 1.04$, $P = 0.37$; CA2: $F_{(3)} = 1.53$, $P = 0.21$; CA3: $F_{(3)} = 0.53$, $P = 0.65$] but varied within blocks in CA1 and CA2, but not CA3 [CA1: $F_{(3)} = 3.82$, $P = 0.0097$; CA2: $F_{(3)} = 5.23$, $P = 0.0014$; CA3: $F_{(3)} = 0.21$; $P = 0.89$]. In CA1, sessions 1 and 2 had lower firing rates than session 4 (Tukey's HSD, $P < 0.05$), and in CA2 session 1 had lower rates than session 4 (Tukey's HSD, $P < 0.01$).
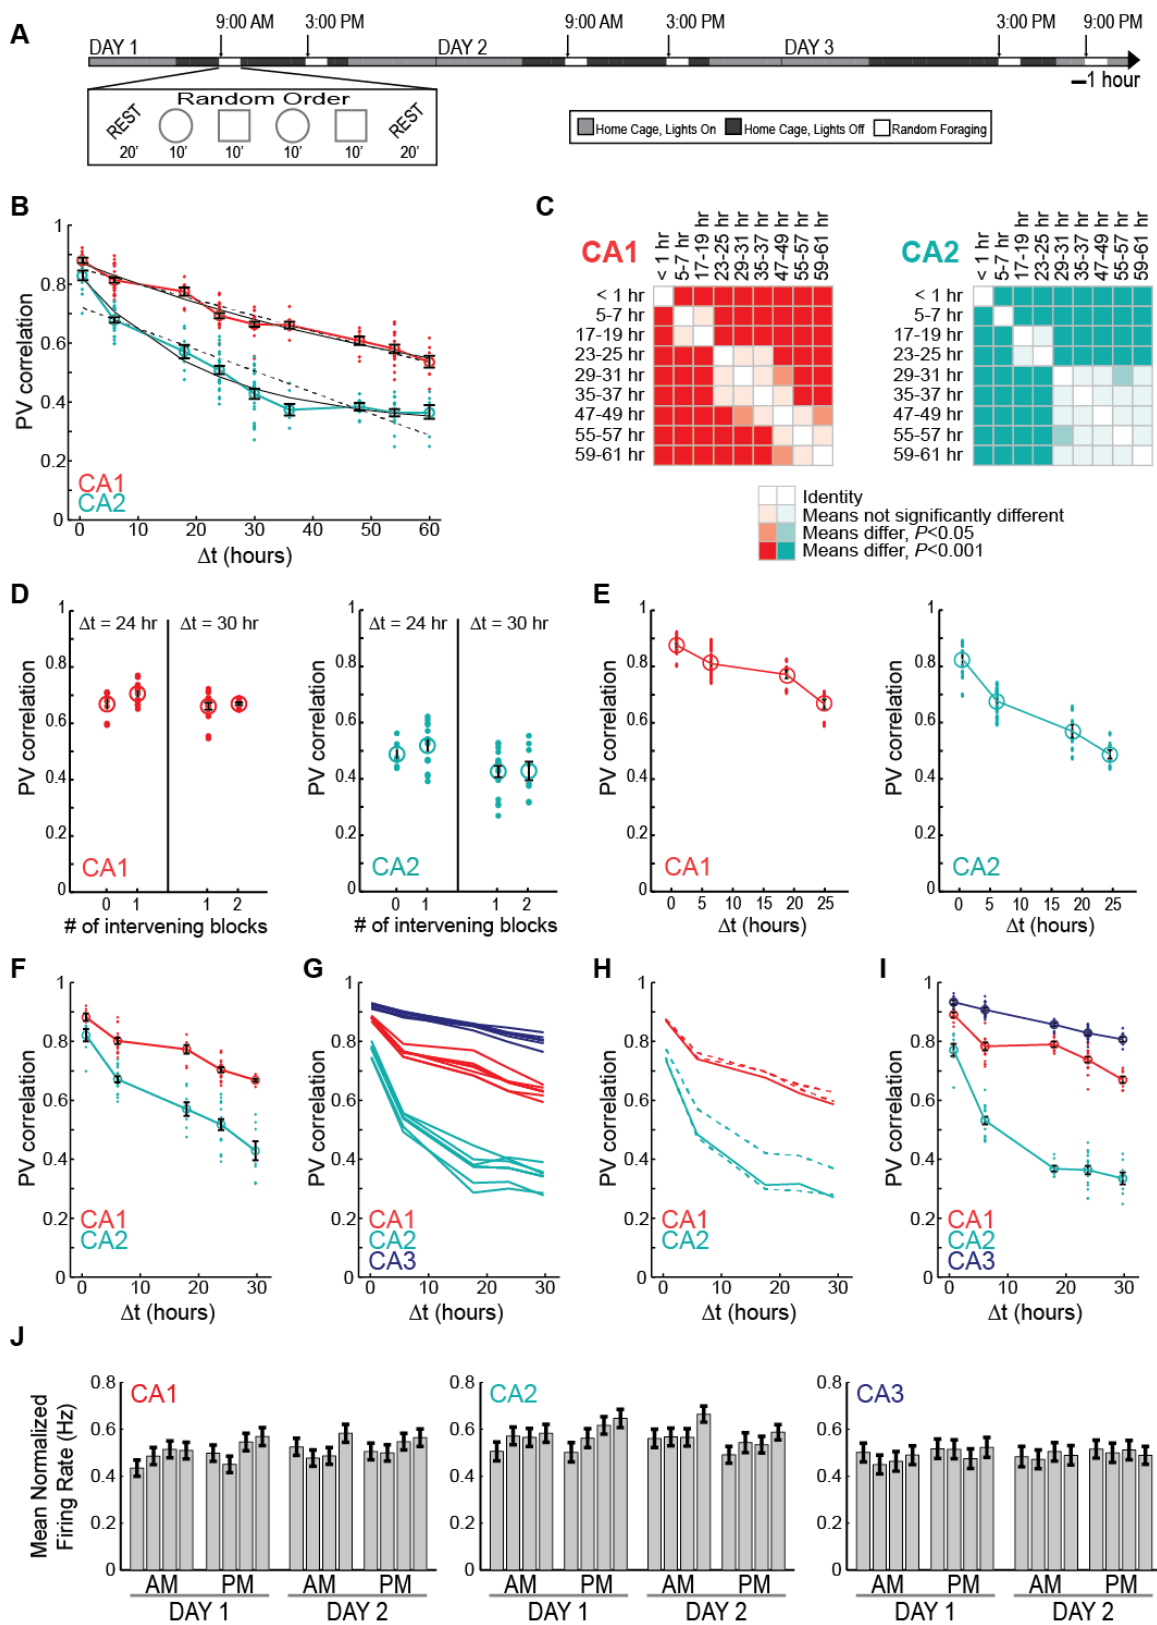
Figure 2S.6. CA2 cells with profound changes in spatial firing patterns over time were found in each rat. (A) Three example CA2 cells are shown from each animal. (Each row) Color-coded firing rate maps of one cell that was recorded for 16 behavioral sessions throughout the course of the two day experiment. The color scale is from 0 Hz (blue) to peak rate within the entire testing sequence (red, indicated to the left) with pixels not visited in white. For each cell, the place field boundaries were computed on the average map of all 16 sessions in square and circular nclosures. The field boundary for one place field from each cell is highlighted in white. Rat H performed the experimental sequence twice with a unique set of cells during each repetition; cells 1-3 are from the first repetition and cells 4-6 are from the second repetition. (B) Mean population vector orrelations (as described in Figure 4C) for simultaneously recorded cell ensembles from each rat. Pairs of recordings in the same enclosure shape at different time points were compared. The number of cells in the analysis is indicated at the lower left corner of each plot (in parentheses). Red, CA1; teal green, CA2; dark blue, CA3. Brain regions with fewer than 5 cells were not analyzed. The smallest difference between CA1 and CA2 in the population measure was observed in Rat G and cell set 1 of Rat H, nonetheless these recordings had clear examples of multiple cells that changed over time.
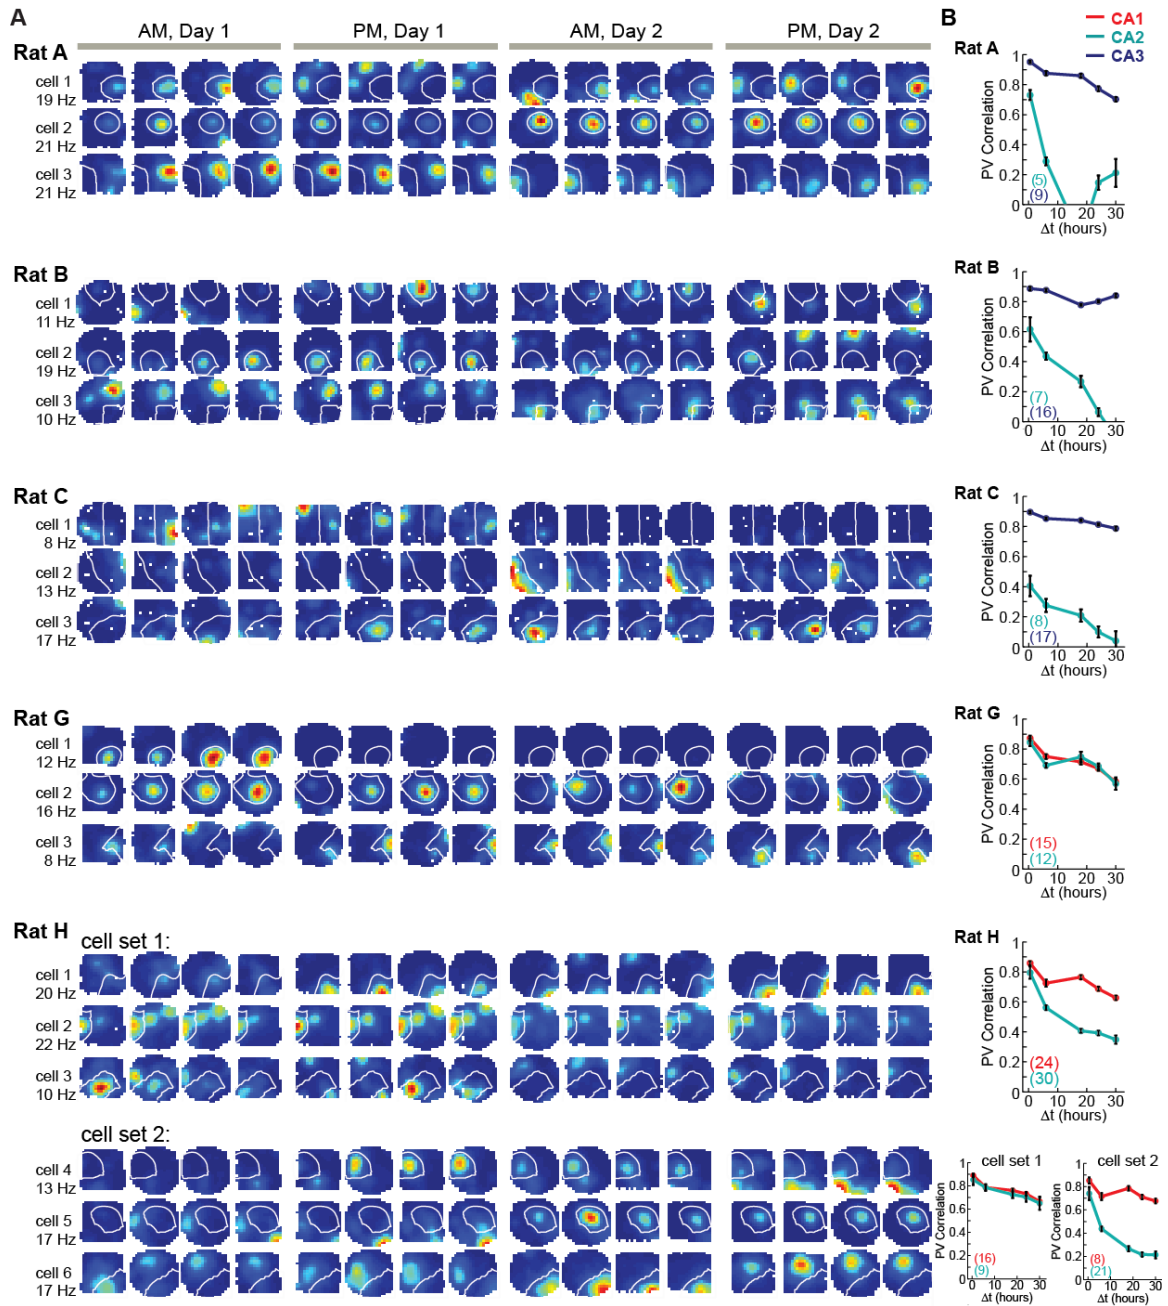
Figure 2S.7. In the two-day, single-shape paradigm, all three hippocampal subregions show a moderate decorrelation within a recording block, but major differences between subregions emerge over longer time intervals. (A) In the timeline for the two-day single-shape experiment, the first 10-minute session of each AM or PM block is denoted with a lower case letter. (B) The complete correlation matrix for each hippocampal subregion is redrawn from Figure 6D. Correlation coefficients between sessions are represented on a color scale (0 to 1, dark blue to red). The pixel in the ith row and jth column in the correlation matrix represents the mean Population Vector (PV) correlation between the ith and jth sessions. Consequently, the matrix is symmetric with comparisons between the same sessions along the diagonal (their correlation coefficient is, by definition, 1). The four rows that consist of comparisons with the first session of each block (a-d, see A) are highlighted on the CA3 matrix (black box). (C) PV correlations between the first session of each block (a-d) and all other sessions. Each hippocampal subregion is shown in a column (CA3, left; CA2, middle; CA1, right). The time axis represents elapsed time from the beginning of the experiment. In CA2, correlation falls off sharply in either direction from the reference point. In CA3, the PV correlations decrease within a block but reset at the beginning of each block. CA1 is intermediate between CA2 and CA3.

Figure 2S.8. Peak firing rates in place fields of multi-peaked CA2 cells are modulated independently. (A) Five representative CA2 cells with multiple place fields are shown. For each cell, the place field boundaries were computed on the average map of all 16 sessions in square and circular enclosures. In the overview image on the left, the place field with the highest overall peak rate is shown in blue, the place field with the next highest rate is shown in green, and if present, a third place field is shown in red. To the right of the overview image, the firing rate maps of all 16 sessions over 2 days are depicted. Maps are ordered by time, but squares and circles are segregated into separate panels. Thus, sessions that are labeled 1-2 in the line graph were recorded on the morning of day 1, 3-4 were recorded on the afternoon of day 1, 5-6 were recorded on the morning of day 2, and 7-8 were recorded on the afternoon of day 2. For each cell, the peak rate of each place field in each behavioral session is plotted below the corresponding map. Line colors correspond to the place field depicted in the overview. (B) The set of 8 firing rates from a single field in each enclosure shape, ordered by time, constitutes a firing rate trajectory. For pairs of fields from the same cell, the Spearman correlation coefficient between firing rate trajectories was calculated. The same calculation was performed for pairs of fields from separate cells within the same subregion. To test whether firing rates in fields of single cells were modulated independently, we compared the distribution of correlation coefficients from within cell comparisons with the distribution of correlation coefficients from distinct cells. In all three CA regions, the distributions did not differ (Kolmogorov-Smirnov test: CA1, $K = 0.05$, $P = 0.76$; CA2, $K = 0.05$, $P = 0.50$; CA3, $K = 0.08$, $P = 0.54$). This indicates that the firing rates of individual place fields of the same cell are modulated as independently as firing fields from different cells.

Appendix 2.3: Supplemental Methods

     Subjects and surgeries. CA2 data were collected from five male, experimentally naïve, Long Evans rats with a preoperative weight of 400-510 g. Animals were housed individually and maintained on a 12-h light/12-h dark schedule with lights off at 6:00 am. All behavioral testing occurred in the dark, except when noted. The data from three of the five animals were recorded simultaneously with previously reported CA1 and CA3 recordings (Mankin et al., 2012). For comparisons of the CA2 data with all available CA1 and CA3 recordings from the same experimental design, previously published data from three additional animals without CA2 recordings were also included (Mankin et al., 2012).

     All experimental procedures were performed as approved by the Institutional Animal Care and Use Committee at the University of California, San Diego and according to National Institutes of Health and institutional guidelines. At the time of surgery, rats were anesthetized with isoflurane (2-2.5 % in O_2) and buprenorphine was administered as an analgesic. An electrode assembly consisting of fourteen independently movable tetrodes was implanted above the right hippocampus (AP 3.9-4.0 mm posterior to bregma, ML 3.0-3.5 mm). Tetrodes were prepared as described previously (Leutgeb et al., 2007) and were advanced to the hippocampal cell layer using techniques optimized for recording stability across time (Mankin et al., 2012).

     Behavioral Procedures. After one week of recovery from surgery, rats were partially food-deprived and trained to forage for randomly scattered cereal crumbs in an enclosure with walls that could be shaped either as a square (80 cm by 80 cm) or as a 16-sided polygon (50 cm radius; referred to as 'circular enclosure'). A polarizing white cue card (20 cm wide) was placed on an inside wall of the enclosure. The center of the enclosure was always located at the same place in the room, and the angle of the cue card compared to external room cues was kept constant. Training was performed in two daily blocks. For all rats, the first block started between

8:30 and 10:00 am and the second block between 2:30 and 4:00 pm. For each individual rat, the daily start time of each block varied by less than 30 minutes. Rats were returned to the animal housing room between the morning and afternoon training blocks.

Rats were trained to run for four 10-minute sessions during each block, with two sessions in the square enclosure and two sessions in the circular enclosure. The order of the shapes was varied randomly within each training block. The rats were allowed to rest for five minutes between sessions, and training blocks were flanked by 20 minute sleep sessions. The floor of the enclosure was cleaned with water between each session. Following the sleep session at the end of the afternoon training block, rats were screened for single-unit activity. One tetrode remained in the cortex and was used as a reference for all recordings. Another tetrode was lowered to the hippocampal fissure to record hippocampal local field potentials. The remaining twelve tetrodes were slowly advanced towards the hippocampal cell layer. Prior to the implantation of the recording array, the relative spatial arrangement of all tetrodes with respect to each other was noted. This allowed for the targeting of more anterolateral tetrodes towards CA2/CA3 and the more posteromedial tetrodes towards CA1/CA2. The placement of all tetrode recording sites in each of the hippocampal subregions was determined from histological material that was prepared as described below.

Electrophysiological recordings throughout the morning and afternoon sessions were initiated when multiple well-isolated cells ($> 300 \, \mu V$) were observed on most tetrodes. The recording phase of the experiment began after 9 to 20 days of behavioral training. Recordings were first conducted for 2 days in the standard training paradigm that included recordings in the square and circular enclosure (referred to as two-shape, day 1; two-shape, day 2). In five animals (two with CA2 recordings), recordings were also performed on a third day that was identical to the first two recording days, except that the start times of the blocks were shifted by 6 hours, so that the first block occurred at 3:00 PM and the second at 9:00 PM. The second block was thus

during the light phase of the light/dark cycle (**Figure 2S.5**). Additionally, four animals (two with

CA2 recordings) performed two days of behavior in which all sixteen random foraging sessions

were conducted in a single enclosure shape (single-shape, day 1; single-shape, day 2). For each

animal, we selected the shape in which we identified the larger number of active cells during the

recording on the preceding day. This paradigm was otherwise identical to the two-shape

paradigm. In one animal, the five-day recording sequence that included the three days in the two-

shape paradigm and the two days in the single-shape paradigm was repeated once. In this animal,

the first set of recordings began after 16 days of training and the second set began after 22 days of

training and the set of recorded cells during each sequence were from different tetrodes or, when

from the same tetrode, determined to be unique by inspecting the clusters. Therefore, the cells

that were recorded during each repetition were included as separate cells in the analysis.

Recording procedures. For recording spikes and local field potentials, the electrode

assembly was connected to a multichannel head-mounted preamplifier. The *x-y* position of light

emitting diodes on the preamplifier was tracked at 30 Hz by processing video images. Unit

activity was amplified and band-pass filtered at 600 Hz to 6 kHz. Spike waveforms above a

trigger threshold (40-60 μV) were time-stamped and recorded at 32 kHz for 1 ms.

Cell-tracking and single unit isolation. Because our study depended on being able to

follow the same set of principal cells over an extended time period, we developed a customized

version of MClust (Redish, MClust. http://redishlab.neuroscience.umn.edu/MClust/MClust.html)

with added functions that allowed for the comparison of the cluster boundaries of each cell

throughout a series of rest and 10-minute recording sessions. Clusters that persisted in the same

region of parameter space throughout two days were accepted as single cells for further analysis.

Rest sessions at the beginning and end of each behavioral session were analyzed to confirm that

changes in cell activity patterns in behavior were not a result of tetrode instability. Care was taken

to accept only cells that could be precisely followed from the beginning to the end of the data

analysis, and all clusters were inspected visually to ensure that rate changes that were observed

could not be attributed to clusters drifting outside of defined cluster boundaries (**Figure 2S.1**).

Postmortem confirmation of recording location. Rats received an overdose of sodium

pentobarbital and were perfused transcardially with saline and either 4 % formaldehyde or, when

immunohistochemistry was performed, 4 % paraformaldehyde. The brains were extracted and

stored in the fixative with 30 % sucrose. Frozen coronal sections (40 µm) were cut, and each

section through the part of the hippocampus with electrode tracks was collected and stained with

cresyl violet. To confirm our delineation of CA field boundaries, immunohistochemistry using

antibodies against α-actinin-2 was performed as previously described (Ratzliff and Soltesz, 2001)

on tissue from two recording animals and two additional animals in which recording was not

performed. Immunoreactivity for α-actinin-2 has been shown to be high in CA2, in the molecular

layer of the dentate gyrus, and in the cell bodies of interneurons scattered throughout the

hippocampus (Wyszynski et al., 1998). Free-floating sections were incubated with mouse anti α-

actinin-2 antibody (1:50,000; clone EA-53, Sigma) overnight at 4°C. On the following day, the

sections were incubated for 2 hours with HRP-conjugated horse anti-mouse secondary antibody

(1:1,000; Sigma). The secondary antibody was visualized by reaction with diaminobenzidine, and

sections were counterstained with cresyl violet.

The final tetrode positions were determined by three-dimensional reconstruction of the

tetrode array in serial sections. Because the brain was cut at a small angle compared to the

orientation of the electrode tracks, the path of each electrode through the brain could be

visualized as elongated segments of tissue damage in each section. The tip of each electrode was

found by tracking the damage throughout the series of sections that was arranged from anterior to

posterior (see Figure S1 in Leutgeb et al., 2007 for the appearance of the tracks in serial sections).

Tetrodes were not moved after electrophysiological recordings, and recordings from a tetrode

were included in the data analysis if the tetrode's deepest position was in or just below the

pyramidal cell layer of CA1, CA2, or CA3. In tissue for which immunohistochemistry was performed, CA2 was defined as the region with dense α-actinin-2 staining in cell bodies and proximal dendrites. We confirmed that this definition coincided with morphological criteria, so that the boundaries of CA2 could also be identified in tissue in which we did not stain for α-actinin-2. The boundary between CA1 and CA2 was drawn at the transition where the cell layer thickens and where cell bodies increase in size and become less densely packed (Lorente de No, 1934; Woodhams et al., 1993). The boundary between CA2 and CA3 was drawn at the transition where the cell layer becomes less compact and where large cell bodies are scattered outside the continuous cell layer (**Figure 2S.2**).

To further increase confidence that tetrodes were correctly assigned, we excluded tetrodes that were at transitions between CA regions. The transition zone between CA1 and CA2 was defined as the area in which there was only a minor decrease in packing density compared to CA1. This zone corresponded to an area in which the immunolabeling of cell bodies and dendrites was sparse in the α-actinin-2 stained section. The transition zone between CA3 and CA2 was defined as the area with larger cell bodies, but with few cell bodies outside of the densely packed inner sublayer. Because the CA2/CA3 border shows a more gradual transition than the CA1/CA2 border, we obtained further confirmation of tetrode assignment to CA2 by measuring the distance of all CA2 tetrodes from the CA1/CA2 border. At the anteroposterior level where the majority of recording tetrodes was located, the proximal to distal extent of the principal cell layer within the CA2 region, as defined by α-actinin-2 staining in proximal dendrites, ranged from approximately 350 μm to 425 μm. All tetrodes assigned to the CA2 region were less than 350 μm from the CA1/CA2 border. We also confirmed that tetrodes that were assigned to CA3 were well beyond the CA2/CA3 border, and the nearest tetrode was at a distance of more than 700 μm.

Spatial tuning and correlation. For each well-isolated neuron, a spatial firing rate distribution was constructed in the standard manner, by summing the total number of spikes that occurred in a given location bin (5 cm by 5 cm), dividing by the amount of time that the animal spent in that location, and smoothing with a Gaussian centered on each bin (Leutgeb et al., 2007). Spatial information per spike was calculated for each spatial firing rate map as

$$I = \sum_i P_i \frac{R_i}{R} log_2 \frac{R_i}{R}$$

, where *i* indexes the spatial bins, P_i is the probability of occupancy in each bin, R_i is the mean firing rate in each bin, and *R* is the mean firing rate across the spatial map (W. E. Skaggs, McNaughton, Gothard, & Markus, 1993).

Place field boundaries, size, number, and stability. Place field boundaries for each cell were calculated as described previously (Mankin et al., 2012) by using maps from either single sessions or the average of 4, 8, or 16 recording sessions. Field boundaries from single sessions were used for the initial analysis to avoid bias that could result from drifting fields (**Figure 2.2** and **Table 2S.2**), and field boundaries from averages across sessions were subsequently used to allow for comparisons over different time intervals. To consider whether short-term spatial drift accounted for differences in field sizes between regions, boundaries were also defined on maps created from 5 minutes of data that were either sampled continuously (the first or second half of each session) or in an interleaved manner (using either the even or the odd minutes of each session). To ensure that broad firing fields were not excluded from the analysis, we used a low threshold and included all fields with peak firing rates of greater than 2 Hz in a single session. Because this criterion was not stringent, we obtained higher proportions of active cells (59.2 % in CA1, 68.2 % in CA2, and 48.1 % in CA3) than typically reported for CA1 and CA3 in similar recording environments (Leutgeb et al., 2004). To confirm that there was also a higher proportion of active cells in CA2 with a more stringent criterion, we determined the proportion of active cells with peak rates greater than 5 Hz, which resulted in 41.6 % active cells in CA1, 45.8 %

active cells in CA2, and 29.5 % active cells in CA3. These values are comparable to those that are typical and confirm that CA2 has more active cells than the other CA subregions irrespective of a particular threshold. In all analysis that averaged maps from recordings in square and in circular shapes, fields with peak firing rates that exceeded 1 Hz were included. The lower threshold for analysis that included the two shapes was used to correct for the fact that fields that were active in only one of the two box shapes would yield a reduced average firing rate in average maps that also include the box shape in which the cell was silent.

After defining field boundaries, the area within the field boundary was taken as the size of the place field, and the bin with the peak firing rate was taken as the field center. For each field that was active in at least three sessions, the area of the convex hull traced out by the field center was used as an estimate for the spatial stability between recording sessions. For each place field, a 'pass' through the field was defined as a trajectory through the place field that included a section through the area of the field in which the firing rate was > 50 % of the maximum firing rate. The mean rate per pass was taken as the number of spikes during the pass divided by the duration of the pass. To estimate the variability of the firing rates, the Fano factor per field was calculated as the variance in firing rate per pass divided by the mean firing rate. Additionally, the change in firing rate throughout a recording session was estimated by calculating the regression for predicting the firing rate from the elapsed time within the session. The absolute value of the slope of the regression line was taken as the change in firing rate over time, and a slope of 0 indicated that there was no increase or decrease in rate. To avoid that the results reflect variability of cells with very low firing rates or very few passes, sessions during which the mean rate per pass in a field was < 0.25 Hz or during which there were fewer than 5 passes were excluded from the analysis of the Fano factor and the within-session change in firing rate.

Theta index and intrinsic frequency. For each cell and each 10-minute session, we calculated the spike autocorrelogram using a time bin size of 2 ms. For the range between -500 ms and 500 ms, each autocorrelogram was fit with the following equation

$$y(t) = [a(\cos(\omega t) + 1) + b]e^{-|t|/\tau_1} + ce^{-t^2/\tau_2^2},$$

where t is the time lag, y is the correlation coefficient at that time lag, and a, b, c, τ_1, τ_2, and ω are the fit parameters. The theta index was defined as the ratio a/b (Royer, Sirota, Patel, & Buzsáki, 2010). We also calculated the spectrum of the autocorrelogram using the multi-taper method (http://chronux.org) (Mitra & Bokil, 2008) and defined the intrinsic theta frequency for each cell as the frequency within the 4-11 Hz range with the largest power.

LFP, phase of firing, and phase precession. A local field potential (LFP) was recorded from the hippocampal fissure of each animal and then band-pass filtered in the theta range (4-12 Hz). A Hilbert transform was performed on the theta signal to decompose it into theta phase and amplitude. Periods when theta amplitude fell below two standard deviations of the mean theta amplitude were designated as low-theta-power periods and were not included in the analysis. For each identified firing field in a 10-minute session, the theta phase for each spike within the field was calculated. To analyze phase precession, all spikes from all passes through a field were pooled. For fields with at least 100 spikes, a best-fit line between phase and distance traveled within the place field was found using the circular-linear fit method (Kempter, Leibold, Buzsáki, Diba, & Schmidt, 2012). The method is iterative, and fields for which the fit did not converge (CA1: 2.75%, CA2: 3.22%, CA3: 1.18%) were excluded from the analysis. The median slope of all fitted fields was compared to zero using the sign test for medians.

Single-Session Firing Properties. To characterize single-session firing properties and compare them between CA regions, we calculated the average firing rate, the spatial information, the number of fields, the theta index, and the intrinsic theta frequency for each cell in each of the 16 sessions recorded across two days. In addition, each cell's firing fields were defined in each

session, and the peak rate and field size of each field were measured. If multiple fields were identified, the peak firing rate and field size of all fields within a session were averaged. The statistics for each CA region were calculated after averaging the values for each cell over the 16 recording sessions, and one value per cell was thus used for statistical analysis. To quantify the variability of firing within a field, we calculated the Fano factor for the firing rates during individual passes through the field within each session, and we calculated the slope of the regression line between the firing rate per pass and time when the pass occurred within each recording session. For these calculations, place field boundaries were defined using the average firing map of all 16 behavioral sessions, and the statistics for each CA region were calculated after averaging the values for each field over the 16 recording sessions. Hippocampal CA regions were compared using the Mann-Whitney U test, and the conversion of the U value to the z statistic is reported if the number of observations in each group was greater than 10. If the number of observations was smaller, the conversion to the z statistic is not accurate and the U value is reported in the text. Significance levels were corrected for multiple comparisons by applying the Holm-Bonferroni correction. To find the proportion of active cells, a cell was considered active in a session if its maximum firing rate in any spatial bin exceeded a threshold (either 2 Hz or 5 Hz). The proportion of active cells in each region was calculated for each session and then averaged across the 16 sessions.

Shape Preference. A shape preference score was calculated for each firing field by measuring the degree to which the distribution of firing rates in one shape was separated from the distribution of firing rates in the other, as previously described (Mankin et al., 2012). Briefly, for every pass through a place field, the average firing rate was computed, and the rates from all passes in each box shape were pooled across either a single block across 16 sessions to yield the firing rate distributions in the square and in the circular enclosure for that place field. The Receiver Operant Characteristic (ROC) curve was computed for these distributions. The area

under the ROC curve may take values between 0 and 1 and can be used as a measure of the overlap of two distributions. We scaled this score by subtracting 0.5 to center it at zero and by multiplying it by 2 so that its range was between -1 and 1. We then took the absolute value so that equal firing rate distributions across the two sets of shapes would have a score of zero and that preference for either the square or the circular enclosure would both have a positive score (see **Figure 2S.4**). For shape preference scores across 16 sessions, chance values were obtained for each CA region by performing 100 random shuffles of shape identity across the 16 recording sessions. To retain the actual number of observations for statistical comparisons between the random and actual scores, we compared the real data with a matching number of subsamples from the shuffled data. To avoid that a particular selection of subsamples yielded unexpected results, we repeated the comparison between randomly selected subsamples from the shuffled distribution (each consisting of the same number of observations as the real data) and the real data 500 times, and reported the median significance of these comparisons.

Population Vector Correlations. For each behavioral session and each hippocampal CA region, rate vectors were constructed by arranging the spatial maps of all cells recorded from that region from all animals in an *x-y-z* stack, where *x* and *y* represent the two spatial dimensions and *z* represents the cell-identity of tracked cells (Leutgeb et al., 2005a; Leutgeb et al., 2005b). To allow for comparisons between the square and circular enclosure shapes, the analysis was restricted to the 16 by 16 bins that were common to both shapes, yielding 256 *x-y* locations. Population vector correlations were obtained by calculating, for each *x-y* location, the Pearson correlation coefficient for firing rates along the *z*-dimension between pairs of sessions. Cells with firing below 1 Hz in all *x-y* bins of the two sessions that were compared were excluded from the population vectors before calculating the correlation coefficients. The correlation coefficients of all spatial bins were averaged to estimate the population vector correlation for a pairwise comparison between sessions. For comparisons by time lag within a block, pairs of recording

sessions were analyzed based on the number of recording sessions between them (e.g., no intervening session corresponds to lag 1), the experimental paradigm used (two-shape, single-shape), and whether sessions occurred in the same or different shape. Because the sequence of shapes was chosen randomly for each block in the two-shape condition, possible comparisons in a comparison group (e.g., lag 1, same-shape) could include instances in which the shapes were presented at different positions within the recording sequence. For example, same shape comparisons at lag 1 could include instances when the first and second recording session were squares (square 1-square 2), when the second and third recording session were squares (square 2-square 3), or when the third and fourth recording session were squares (square 3-square 4) as well as instances when both of the recordings in a pair were in circles (circle 1-circle 2, circle 2-circle 3, circle 3-circle 4). From each of these possible pairwise comparisons, population vectors were created from cells that were recorded in corresponding shapes over a corresponding lag, and their correlations were included as data points when the vectors consisted of more than 25 cells. Because we only compared lag 1 and lag 2 in the two-shape design, we used Mann-Whitney tests for comparisons that differed by lag. The Holm-Bonferroni method was applied to correct for multiple comparisons. In the single-shape comparisons, three different lags (1, 2, and 3) within a block were compared, and we therefore used a two-way ANOVA by brain region and lag. For comparisons on longer time scales, the population vectors were generated by stacking the cells that were recorded in all animals in each hippocampal subregion in each of the 16 recording sessions. The sessions were sorted by their temporal order after separating squares and circles. Pairwise comparisons between the 8 stacks in the square and the 8 stacks in the circular enclosure were grouped by time interval between blocks (e.g., 6 hours) and by same/different shape comparison. Population vector correlations were then compared using an ANOVA, with Tukey's Honestly Significant Difference (HSD) criterion for *posthoc* analysis.

Acknowledgements

Chapter 2, in full, has been accepted for publication of the material as it will appear in *Neuron*, 2014, Mankin, Emily A., Diehl, Geoffrey W., Sparks, Fraser T., Leutgeb, Stefan, and Leutgeb, Jill K. The dissertation author was the primary investigator and author of this paper.

CHAPTER 3: A MODEL OF TEMPORAL AND SPATIAL CONTEXT INHERITANCE IN CA1

Abstract

Place fields in CA2 show markedly greater variability over time in both their location and firing rate patterns than those in CA3. The correlation between representations of two identical events (PV corr) in CA2 thus decays as a function of time to a much greater extent than in CA3. On the other hand, on short time scales, representations of distinct spatial contexts are more similar in CA2 than in CA3. We have proposed input from CA2 and CA3 may be combined in CA1 in order to generate population activity changes over time but more slowly than in CA2 and is simultaneously able to discriminate between spatial contexts.

To test the hypothesis that spatial and rate variability were critical to the amount of change in PVcorr observed in each hippocampal subregion, we generated pseudodata for each region in which spatial and rate variability could be modified. We found that limiting spatial variability shifted the PV corr curves up with little impact on the slope, while limiting rate variability decreased the slope of the curves. We then identified a model for spatial variability in place fields that, combined with true rate profiles, generated PV corr curves that reflected the original data. We found, however, that rate profiles were not amenable to a similar effort.

To test the hypothesis that CA1 may passively inherit its temporal and spatial context codes from CA2 and CA3, we created a feed-forward network of CA2 and CA3 cells and generated a weight matrix to create CA1 cells with desired place fields in one session, then allowed the representations to evolve over time to examine whether the generated network changed over time and across spatial contexts to the same degree as the true CA1 network. We found that a network in which each cell received connections from three CA3 cells and one CA2

cell showed same-shape PV corrs similar to real CA1 data, and inherited some spatial context

discrimination, though less than observed in the actual CA1 network.

Data in this chapter should be considered preliminary.

Introduction

Many theoretical models of memory suppose memory traces contain information not only

about items learned but also about the context in which they were learned, such that items that are

learned in similar contexts have more similar representations than items learned in distinct

contexts (Murdock, 1982; Raaijmakers & Shiffrin, 1981). In order to account for temporal effects

of memory, such models must be expanded to include a component of the contextual signal that is

time-dependent (G. Mensink & Raaijmakers, 1989; G.-J. Mensink & Raaijmakers, 1988;

Murdock, 1997). Thus a stable content signal is combined with a gradually varying context signal

to form a memory trace. Though different models implement this requirement differently, one

consistent feature of the models is that the similarity of the context signal at two different time

points decays exponentially with the time between them (Estes, 1955; G. Mensink &

Raaijmakers, 1989; Murdock, 1997).

We recently showed that the population similarity between representations of a given

spatial context in hippocampal subregion CA2 decays exponentially as a function of the time

between exposures, such that it reaches its floor level of correlation (~0.35) between 18 and 30

hours, whereas population similarity in CA3 remains high for at least 30 hours (Mankin et al,

2014; see chapter 2 of this dissertation). Conversely, population similarity between distinct

environments (an arena shaped as either a square or a circle) was high in CA2 at short time

scales, but low in CA3, indicating that CA3 but not CA2 codes reliably for precise location

spatial context. CA1, the primary output region of the hippocampus, receives information from

CA2 and CA3 (R Bartesaghi et al., 2006; Renata Bartesaghi & Gessi, 2004; Chevaleyre &

Siegelbaum, 2010; Kohara et al., 2014; M. Zhao et al., 2007), The degree of change over time in CA1 population representations was intermediate between CA2 and CA3, while the degree of contextual coding was moderately greater than in CA3. We proposed that CA3 and CA2 may be playing the predicted roles of the stable content signal and the time varying temporal context signal, respectively, and the CA1 could then be combining these two inputs in order to simultaneously encode space, spatial context, and time.

The goal of this chapter is to ask whether, with a very basic network architecture, CA1 could passively inherit the essential features of temporal and spatial context coding from CA2 and CA3. This question was pursued in two steps. First, we explored the ways in which the variability in the size, shape, location, and amplitude of individual place cells over time could account for the patterns of population vector correlation observed in data. CA2 and CA3 each receive input from the entorhinal cortex and other upstream brain regions which are certainly critical in shaping the output of these regions, but for the purposes of this chapter, we assumed that their general behavior could be described as following rules in space and time. For each network we were able to discover a reasonably simple set of spatial parameters that could be used to generate spatial maps representing the activity of place cells from CA2 and CA3 at different time points and in different spatial contexts that varied realistically, both on the individual cell level and on the population level. Unfortunately, similar efforts to find a simple description of how firing rates vary with time has not been equally successful in generating models that reflect the original data.

Next we set up a network architecture with cells from CA2 and CA3 connected with feed-forward weights to the set of cells, N, which we hoped would reflect CA1. Instead of beginning with random weights and training the network, we made use of the observation of Cheng and Frank (2011) that virtually any training paradigm intended to create a place field in a particular location will generate weights that vary inversely with the distance of the maximum firing rate of the input cell to that location. We thus started by defining target place fields for each

cell of N. For half of the cells, we specified weights based on the place fields exhibited by CA2 and CA3 cells in a square session, and for the other half we specified weights based on a session in a circle. Having defined a static weight matrix from CA2/CA3 to N based on input from only a single session (per cell in N), we then asked if the N network would show coding for temporal and spatial context in a manner similar to CA1, such that we could consider the feed forward model from CA2 and CA3 a realistic depiction of activity in CA1.

Results

It has been observed that in 2-dimensional environments the firing rates of most place fields in the hippocampus may be approximated by a 2-dimensional Gaussian surface (Muller et al., 1987; John O'Keefe & Burgess, 1996; Solstad, Moser, & Einevoll, 2006). After identifying individual place fields from each cell (see method in Mankin, et al, 2012, or this dissertation Figure 1S.11), each field was fit individually in each session according to the equation

$$F(x,y) = a + b \exp\left[-\left(\frac{(x - c_1)\cos(\theta) + (y - c_2)\sin(\theta)}{w_1} \right)^2 - \left(\frac{-(x - c_1)\sin(\theta) + (y - c_2)\cos(\theta)}{w_2} \right)^2 \right],$$

where *F(x,y)* gives the firing rate at each location, with 7 fitted parameters: *a,* the overall offset from 0; *b,* the peak rate; (c_1, c_2), the *(x,y)* coordinate of the center; w_1 and w_2, the spread of the Gaussian in the *x* and *y* direction; and *θ,* the angle of rotation (**Figure 3.1A**). Creating spatial maps by adding together the fitted Gaussians from each field of a place cell resulted in excellent approximations of the original spatial maps (median spatial correlation between original map and reconstructed map; CA1: 0.958; CA2: 0.942; CA3: 0.947). Of the maps that were not well-fit, the majority had peak rates less than 2 Hz (CA1: 81.7%, CA2: 71.3%, CA3: 92.2%) (**Figure 3.1 B-C**).
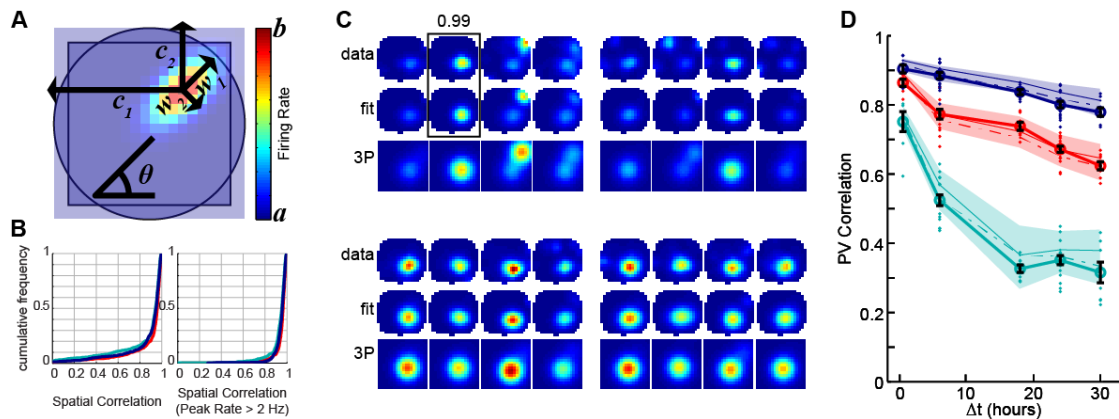
Figure 3.1. Spatial firing rate maps were well-fit by sums of 2-dimensional Gaussian surfaces. **(A)** Each firing field was fit with a 2-dimensional Gaussian surface with 7 free parameters: a and b determined the maximum and minimum values of the curve, c_1 and c_2 determined the center, w_1 and w_2 determined the width of the field in each direction, and θ determined the angle of rotation of the field (see text for full equation). In the example field shown, $a = 0$, $b = 8$, $c_1 = 14$, $c_2 = 7$, $w_1 = 4$, $w_2 = 2$, and $\theta = \pi/4$. Note that all lengths are reported in units of pixels of the spatial map. 1 pixel = 5 cm. **(B)** Maps were well fit by this method. (Left) Cumulative distribution plots for each subregion of the spatial correlation between the original spatial map and the map generated by summing Gaussians. In all three regions, the median correlation was greater than 0.94. CA1: red, CA2: teal, CA3: dark blue. (Right) Cumulative distribution plots for each subregion of the spatial correlation between original and best-fit maps, including cells only when their peak firing rate was greater than 2 Hz. Much of the left-skew of the distribution is eliminated by removing maps with low firing rates. **(C)** Top: Example firing rate maps from one CA2 cell across one day of recording. Each pixel represents the mean firing rate for that cell when the animal was in that location. The color scale goes from 0 Hz (dark blue) to 19 Hz (red; the cell's peak rate across all sessions). The top row shows spatial firing rate maps as recorded. The eight maps represent 8 behavioral sessions, 4 in the morning and 4 in the afternoon, sorted so that the sessions in the square are shown first within each block of 4. The second row, labeled 'fit' shows the firing rate maps generated from each of the data maps by fitting with a Gaussian as described in A. The spatial correlation was generally quite high. The pair with the best spatial correlation is highlighted by a black square, and the correlation was 0.99. In the third row, labeled 3P, maps have been generated with the same values for c_1, c_2, and b that were used in the "fit" row, however all other parameters have been set to a fixed value; $a = 0$, $w_1 = w_2 = 4.5$, $\theta = 0$; Bottom, maps from the same cell on the second day of recording, arranged as described. **(D)** Replacing maps by their best fit versions or simplified 3-parameter fit versions does not dramatically change the PV correlation values observed in the data. PV correlation values from original data are shown for each region with light dashed lines. PV correlation values from the best fit data are shown in light solid lines. The PV correlations for the 3P model shown in the third rows of C are displayed in heavy lines. This model received a score of 3 in each region (see text for definition of score). The shaded portion in each color shows the 95% confidence intervals generated from 100 subsamples of the data. Requiring all mean PV correlation values from a model to lie within the confidence intervals resulted in an even stricter standard than the requirements for a 3 point model.

Parameterizing Spatial and Rate Variability

The primary result we wanted to approximate with modeling in the first half of the project was the trend in population vector correlation (PVcorr) as a function of time in each subregion. We thus calculated the PVcorr between the reconstructed spatial maps and compared them to the PVcorr in the original data. In each region, we found that the PVcorr in the reconstructed data was very slightly greater than in the original data (ANOVA on data type and ΔT, main effect of data type, CA1: $F(1) = 4.62$, $P = 0.034$, CA2: $F(1) = 4.14$, $P = 0.045$, CA3: $F(1) = 14.4$, $P < 0.001$; maximum difference of mean PVcorr values, CA1: 0.024, CA2: 0.046, CA3:0.020), but there was no interaction, indicating that the overall trends were the same (interaction term, CA1: $F(4) = 0.027$, $P = 0.999$, CA2: $F(4) = 0.65$, $P = 0.626$, CA3: $F(4) = 0.057$, $P = 0.999$) (Figure 3.1D). This is consistent with the idea that the reconstructed maps replicate the original data well but smooth out a small amount of extra-field noise, leading to moderately higher correlation values.

A model should closely recreate the data that it is designed to represent. But as simplifying assumptions are introduced, some amount of change is bound to occur. We have already observed that replacing spatial maps by their best reconstruction as a sum of Gaussians introduced a small degree of error in the mean PVcorrs, but this degree of error did not seem sufficient to rule this out as a model. This is because the differences in mean PVcorr were small at each time point. Additionally, the main feature of the data that we set out to capture with the model was the degree of change in PVcorr over time in each subregion. Because the interaction term in the ANOVA was not significant for any of the regions, we were able to say that this model performed well in that regard. Based on these ideas, we developed a rating system for each model tested. A model received "gold" if the ANOVA between the generated PVcorrs and either the PVcorrs from the actual data or from the best-fit reconstructed data showed no significant effect of data type (modeled vs actual data or modeled vs best fit data) and there was no

significant interaction between data type and ΔT. A "silver" model was also required to have no

significant interaction between data type and ΔT but the requirement that there be no main effect

of data type was relaxed to a requirement that the maximum difference between mean PV corrs at

any time point was less than 0.05. Finally, a "bronze" model was one that did not meet the

standards of gold or silver, but nonetheless the maximum difference between mean PVcorrs at

any time point was less than 0.1. Any model that did not meet at least one set of these criteria

was deemed poorly fit. To check whether these requirements were too stringent, requiring models

to be overfit, we performed 100 subsamples of the data, taking 80% of the cells in each

subsample, and recalculated the PVcorrs for each subsample. For each ΔT point, we generated the

96% confidence interval for PV corr by selecting the 3[rd] and 98[th] rank-ordered value from the 100

sampled values (Figure 3.1D). It was our intention that requiring the mean PV corrs of a model to

fall within these confidence intervals would be an easier standard to meet than the silver standard

described above and would thus prevent over fitting our models. Instead, many models that were

ranked as gold models failed to meet this criterion. We thus focused only on the gold-silver-

bronze scoring system to evaluate each model.

Many model specifications were not deterministic. If there was some element of

randomness in a model, different instantiations of the same model could have returned different

results. To test the soundness of such a model, we simulated each model 5 times, gave each result

a score (gold = 3, silver = 2, bronze = 1, poorly fit = -1), and took the mean score as an estimate

of how well the model performed. Models with mean scores greater than 2 were considered to

have performed well.

The hypothesis put forth in the previous chapter was that the variability in rate and in the

centers of the place fields, combined with the presence of multiple independently-modulated

place fields from a single cell, were the primary drivers of the decline in PVcorr over time

observed in CA2. To test this, we asked whether all other parameters for each place field could be

set to a fixed value without dramatically changing the results. We set $a = 0$, $\theta = 0$, and required w_1 = w_2, such that all place fields in each region would now be circular and of the same size. To choose the correct width, we created maps using a fixed width between 2 and 8 pixels and scored each model (each pixel represents a 5 cm2 bin, so this translates to field radii of between 10 and 40 cm). The goodness of fit was robust across many width values in each region In CA1, models received gold ratings for widths between 2 and 5.5 pixels. In CA2, models with field widths greater than 4 pixels all received gold ratings, and in CA3, taking the width between 3 and 7 resulted in gold fits. Because many width values captured the PVcorrs that we were trying to fit, we chose values that reflected the approximate field sizes observed in each region. So for subsequent modeling efforts, we used $w=3$ in CA1 and CA3 and $w=4.5$ in CA2 (**Figure 3.1 C-D**). Based on these results, we simplified the field equation above to have only three free parameters: the 2-dimensional coordinates of the field center (c_1, c_2) and the peak rate, b, each of which is a function of session number, s:

$$F(x, y, s) = b(s)\exp\left[-\left(\frac{x - c_1(s)}{w}\right)^2 - \left(\frac{y - c_2(s)}{w}\right)^2\right].$$

We asked whether instability in the location of the center of the place field over time was important for the decrease in PVcorrs over time. For each field, we fixed its center in each session at its mean location across the sixteen sessions and then calculated PVcorrs. The slope of each PV curve was approximately the same as the original in this case (only CA2 showed a significant interaction between data type and ΔT. Interaction in CA1:$F(4) = 0.82$, $P = 0.52$, CA2: $F(4) = 3.59$, $P = 0.0089$; CA3: $F(4) = 0.18$, $P = 0.95$), but PVcorrs were markedly shifted up (main effect of data type, CA1: $F(1) = 69.25$, $P < 0.001$; CA2: $F(1) = 232.2$, $P < 0.001$; CA3: $F(1) = 17.2$, $P < 0.01$) (**Figure 3.2A**). It seems, then, that the slope of the decorrelation is likely driven by variability in firing rate, but spatial variability adjusts the overall starting level. To further

Figure 3.2. Spatial variability can be well-modeled with a 2 or 3 stage random selection process. **(A)** For each place field, location was computed from the Gaussian fits in each session. The field's location was defined as the mean of these locations. PV correlations were computed when the place field center was held constant across time, and are shown in thick lines. For comparison, mean PV corrs for actual data are shown in light dashed lines and for data generated with the best Gaussian fit are shown in light, solid lines. PV correlations in CA3 were well-fit with stationary place fields. In CA1, PV correlations at $\Delta T = 0$ and $\Delta T = 6$ hours are close, but PV correlations do not decrease as much as expected on longer time scales. CA2 is well-fit only for $\Delta T = 0$ hours. **(B)** Place field centers were selected in each session from a normal distribution centered on the mean location and with a standard deviation (std) that was varied systematically. PV correlation curves are shown for three different standard deviation values. The CA1 curve could be approximated by selecting for std = 0.625, but the correlation was too low at $\Delta T = 0$ and too high for $\Delta T > 18$ hours. When std = 1, CA2 was well fit at short time scales but too correlated at long time scales. For parameters when CA2 reached sufficient decorrelation at long time scales, such as when std = 1.75, it was not well fit at short time scales. **(C)** Top left: an example place field at the center of the box and with a radius of 3 pixels, the size used to model CA1 and CA3 cells. Black circles in the middle show 100 possible location choices with std = 0.35, the standard deviation that best fit CA3. Approximately two-thirds of new centers would be selected from within the white circle. Top right, magnification of the center of the box. The large black circle shows the region from which field centers could be selected as described in B with std = 0.625. With that process, two sessions that occurred adjacent to each other in time could have field centers all the way across the circle from each other, which may explain why the PV correlation in CA1 was too low for this std at $\Delta T = 0$, even though the rest of the curve was well fit. To account for this, a two-stage process may be used instead. A smaller black circle shows the region from which block centers may be selected. 4 red dots give example block centers. For each block center, a new region is defined by a red circle, indicating the region from which place field centers may be selected for the sessions in that block. Using std = 0.35 for the small black circle and 0.5 for the red circles resulted in spread in the place field centers that would cover the same region as using std = 0.625 in the singe stage process, but it maintained some temporal structure that was lost in the single stage process, and allowed a better fit for CA1. Lower left: an example CA2 place field with radius = 4.5 pixels and 100 place field centers chosen with std = 1. Two thirds of the place field centers should lie within the white circle. Lower right: A three-stage selection process was required to maintain sufficient temporal structure to fit CA2 data. The black cross shows the place field center and the black circle outlines the region from which centers for each day may be selected. Two red dots indicate possible centers. Red circles outline the region from which block centers may be selected. From each red circle two block centers are selected and individual session centers may be selected from within each blue circle. Blue circles connected with a line indicate blocks from the same day. In this way, field centers from a single block are constrained to be close together, and field centers from a single day are also more constrained than fields from across days. For CA2, the best parameters were 0.975 for the standard deviation for each day (black circle), 1.35 for each block (red circles), and 0.75 for individual sessions (blue circles). **(D)** Mean scores for two-stage models are shown for various R-S combinations. Based on the best fit single-stage parameter for each region, nearby R-S space was tested. Five iterations for each combination were generated and scored from -1 to 3, as described in the text. The mean score is shown, with lighter colors indicating scores closer to 3. CA1 and CA3 each had broad regions where the mean score for the model was 3, indicating that the two-stage process reliably produced good fits anywhere in this region. For CA2, the best R-S combination received a score of 2.6. Adjacent R-S combinations, however, scored poorly, suggesting that there was not a region of R-S space where the two-stage model was reliably well-

fit, which was why a 3-stage model was employed for CA2. **(E)** Top: example cell from CA1 recorded across 2 days. Laid out as in Figure 3.1. The top row shows the maps with actual field centers (3P is the three-parameter model described in Figure 3.1), the row below shows maps with field centers selected with the two-stage process described in C (2S indicates the two-stage selection process). Bottom: The same CA2 cell shown in Figure 3.1. 3P maps are repeated, beneath them examples of fields selected with the three-stage process described in C. **(E)** PV correlation curves for five repetitions of selecting place field centers according to the models described. For comparison, the 96% confidence interval is shown. The three-stage model reliably produced good fits for CA2 data.
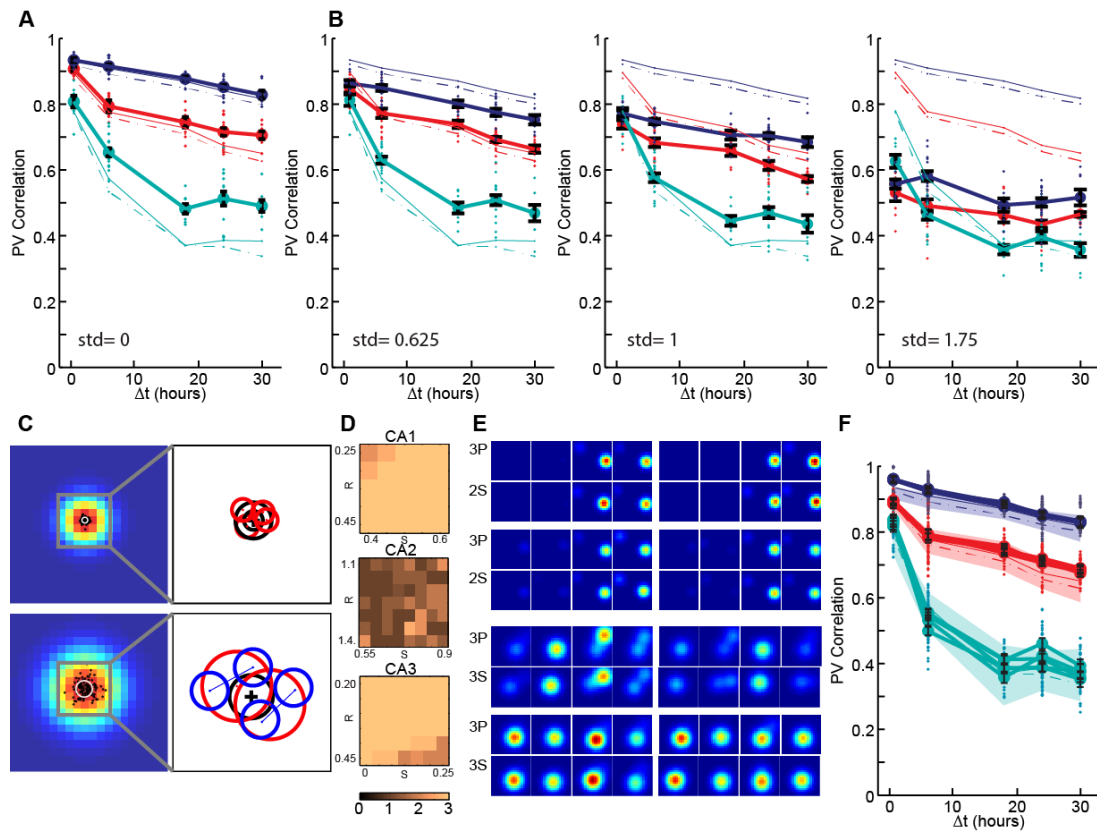
explore the spatial variability, we allowed each field's center to fluctuate randomly around its mean location. That is, for each time point, we drew a new x and y location for the center from a normal distribution centered around the field's mean location. We repeated this process for several different variances to find the value that most closely matched the data.

We found that introducing a small amount of variance in place field centers was sufficient to capture the data of CA1 and CA3 (CA1 scored 3 points for standard deviation between 0.5 and 0.625, and CA3 scored 3 points for standard deviations less than 0.375) . CA2, however, was never well fit by this model. With small levels of variance, the PVcorrs for short ΔT (1-6 hours) could be well fit, but the PVcorrs did not reach the appropriate decorrelation level at longer ΔTs. Conversely, larger levels of variance brought the correlation levels down, but to the point that short ΔT values were not well fit. This result is not entirely surprising, as it fails to account for the temporal structure of spatial variability. That is, under the assumptions of this model, sessions that were close together in time were just as likely to have distant place fields as sessions that were far apart in time. To account for this, we developed a two-stage random sample model. A center point for each 4-session block was selected from a normal distribution around the field's center, with standard deviation R, and then the location in each session was selected from a normal distribution around its block's center point with standard deviation S, where R and S were parameters that could be varied to find the best model fit (**Figure 3.2B**).

The two-stage random sample model improved the fits in all three subregions. In CA1 and CA3 there were several (R,S) pairs that produced scores of 3 (indicating gold standard fits across all 5 tests), and these appeared to cover a continuous region of R-S space. The fits for CA2 improved substantially, and there was one pair of (R,S) values that scored > 2.5, but it did not appear to be inside a stable within the R-S space, as several values immediately adjacent scored < 2 (**Figure 3.2D**). The two-stage model, while accounting for some temporal structure in the data, may not have accounted for it all. The four blocks were not equally spaced in time. We thus

instated a three-stage random sample model, where the center point for each day was initially chosen from a normal distribution around the field's center point with standard deviation Q, and then block and session center points were selected around those points as described above (**Figure 3.2C**). Because CA1 and CA3 were well fit with the two-stage process, we let Q = 0 for these regions and only explored the parameter space for Q for CA2 with the (R,S) pair that had already been selected as the best fit for the two stage model. The best (Q,R,S) parameters for each region were CA1: (0 0.35 0.5), CA2: (0.975, 1.35, 0.75) and CA3: (0, 0.3, 0).

Next we asked whether the greater number of place fields in CA2 compared to CA3 was critical for the difference in PV corr decay. For each field, we took only the *n* place fields with the highest mean firing rate. Surprisingly, when *n*=2, CA2 received a score of 3, indicating that it met the gold standard in each of 5 repetitions. This indicates that using two place fields per cell is sufficient to induce the degree of change observed in the CA2 population over time. CA1 also required two or more fields, while CA3 was well fit with one or more fields. These results were consistent whether we used actual data values for the parameters of the Gaussians or we used the 3-parameter model with spatial variability induced as described above.

Finally, we wanted to understand the structure of the rate variability we observed over time. We wondered whether rates were random, and whether we could find subsets of cells that may have been functioning as cell assemblies such that their rates varied together. To start, we checked that rate variability was important by fixing the rate for each place field as its respective mean rate across the 16 sessions. As predicted from the results with fixed place field centers, removing rate variability dramatically decreased the slope of the PV corr curves in all three regions (**Figure 3.3A**). Previous work, not shown, had suggested that none of the following produced well-fit models: selecting weights randomly around a mean rate, inducing a random walk in the rate profile in which each successive rate is sampled from a distribution around the previous weight, using a Markov model in which a cell transitions with some probability between
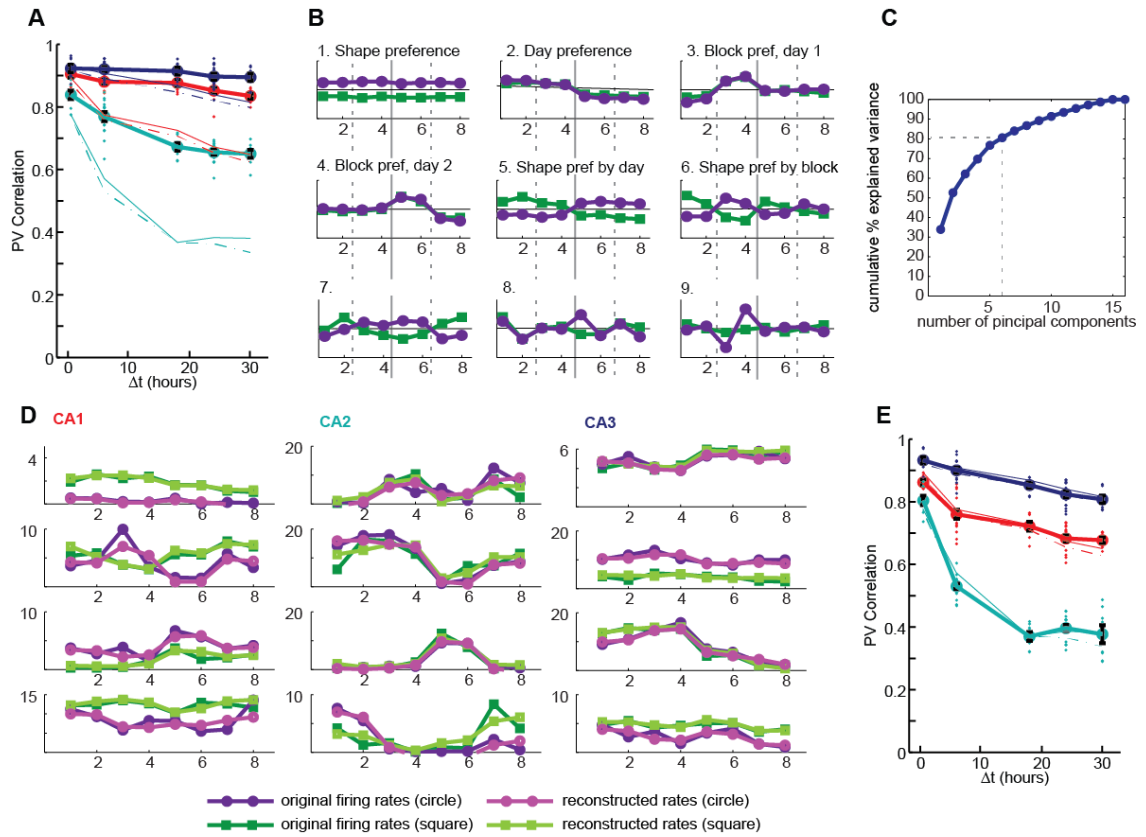
an on and off state and each successive rate in the on state is modeled as a random walk. Instead, we sought to describe more accurately the rates that were actually present in the data.

For each field, we collected its rate trajectory, the 16 values of b in the Gaussian fits across sessions, and then performed principal components analysis (PCA) on the collected rate trajectories to find a set of basis functions for the rate trajectories (shown in **Figure 3.3B**). We found that the first 4 principal components were remarkably in line with what we had already observed. The first component had almost no dependence on time, but the rates for one shape were positive while the rates for the other were negative. Since these are basis functions, signs are arbitrary, so this component would elevate the rates of one shape while suppressing the weights of the other, indicating that rate differences between shapes accounts for a lot of the data. The second, third, and fourth components, conversely, had no dependence on shape, but instead varied with time. In the second component, rates on one day were elevated while rates on another were suppressed. In the third component, rates in one block of day 1 were elevated while rates in the other block of day 1 were suppressed, and rates on day 2 were unaffected. The fourth component was essentially the mirror image of that, with rates on day 1 unaffected and rates on day 2 modulated by block. Components 5 and 6 show some interaction between shape preference and day or between shape preference and block. Together these 6 principal components form a basis for the rate trajectories that accounts for approximately 80% of the variance in firing rates across all three regions (**Figure 3.3C**). Using only the scores for those 6 components, we reconstructed the firing rate trajectories for each cell, and found that the mean square error was less than 2 Hz in each region (**Figure 3.3D**). Finally, we modeled the cells of each subregion according to the model built so far, but with the rates constructed from only 6 principal compnents and found that the model was well-fit (gold) in each region (**Figure 3.3E**).

Having established that firing rate trajectories could be well described by combining the first six basis functions, we asked whether firing rate trajectories clustered in PCA space based on

Figure 3.3. The first principal components of firing rate trajectories reveal basic patterns of spatial and temporal context discrimination. **(A)** PV correlation curves are shown in thick lines for data in which the firing rate for each field is fixed at its mean rate across time. For comparison, thin lines show mean PV corrs for actual data (dashed line) and best Gaussian fit (solid line). The PV corr curves are nearly flat for CA1 and CA3, and the slope of the CA2 curve is much less steep. This indicates that rate variability is critical for generating PV corr curves that become decorrelated with the passage of time. **(B)** For each place field, a list of the peak firing rate in each session composes a firing rate trajectory. For ease of visualization, trajectories are split by shape of the box, so that two 8-session curves are plotted. Purple indicates rates in circular boxes, and green indicate rates in square boxes. These firing rate trajectories were collected for each place field across all three hippocampal subregions, and principal components analysis was applied in order to find the basic "building blocks" for firing rate trajectories observed in the data. They are shown in decreasing order of explained variance. Because components function as a basis set for building firing rate trajectories, units and sign are arbitrary. The first component infuses shape preference: one shape is elevated while the other is suppressed to an approximately equal degree, and there is virtually no temporal dependence. The second component elevates rates on one day while suppressing rates on the other, with virtually no shape difference. The third component elevates rates in one block on day 1 while suppressing rates in the other block of that day and leaving day 2 unaffected. The fourth component does the equivalent for day 2. Interestingly, the 5th component has higher rates for one shape on day 1 and the other shape on day 2. And the 6th component has shape preference switching each block. It is important to note, however, that these components are likely to be combined with one of the earlier components such that the overall shape or time preference overrides the effect of these; it is rare for shape preference to switch cleanly as these components might suggest. Subsequent components are less clearly summarized with words, but an additional three are shown. **(C)** The cumulative amount of explained variance from principal components is shown. The first six components account for just over 80% of the total variance. **(C)** Rate trajectories can be reconstructed closely with only 6 principal components. The firing rate trajectories from 4 CA1, 8 CA2, and 4 CA3 cells are shown. The original trajectories are shown in dark green and purple. Trajectories reconstructed from the scores of the first six principal components are shown in light green and purple. Though there are some points that are not in perfect agreement, overall the reconstructions appear very similar. The mean square error between the original rate trajectories and the reconstructed trajectories was less than 2 Hz in each region (CA1: 1.61 Hz, CA2: 1.86 Hz, CA3: 1.75 Hz). **(E)** The PV corr curves generated from rate trajectories that were reconstructed with 6 principal components were very similar to original data.

A

B
1. Shape preference
2. Day preference
3. Block pref, day 1
4. Block pref, day 2
5. Shape pref by day
6. Shape pref by block
7.
8.
9.

C

D CA1   CA2   CA3

E

original firing rates (circle)   reconstructed rates (circle)
original firing rates (square)   reconstructed rates (square)

the set of coefficients that described each trajectory. We used K-means clustering and applied the Calinski-Harabasz and silhouette measures to search for the optimal number of clusters. Unfortunately, neither measure resulted in a global extremum that was not on one end of the possible number of clusters or the other. That is, Calinski-Harabasz indicated that the optimal number of clusters was 2, while the silhouette measure indicated that more clusters was always better, up to the number of fields we were attempting to cluster. Together these results suggest that the data may not be naturally clustered. We therefore do not find evidence that subsets of cells are modulated together.

Although we did not find a clearly correct number of clusters to use, we still wanted to find out if clustering the data could give us any insight. We thus performed k-means clustering on the PCA coefficients of each field with 10, 20, 30, 40, 50, 60, 70, and 80 clusters, and computed PV corrs for each region. We also tried clustering based on 6, 8, 10, or 12 PCA components included in the analysis for each of these values. **Figures 3.4A-D** illustrate results from using 30 clusters with 6 PCA components, but the results were remarkably similar across all conditions.

For each cluster, we constructed the rate profile of the cluster's center, which we call the cluster exemplar. These are shown in **Figure 3.3A**, sorted by the degree to which the rates in shapes at corresponding times differed. For each region, we computed the fraction of cells that ended up in each cluster. The majority of place fields in CA2 were in clusters whose exemplars had low shape difference, while the majority of place fields in CA3 were in clusters whose exemplars had high degrees of shape difference, and CA1 cells were relatively evenly distributed across the clusters (**Figure 3.3B**). For each place field, we replaced its firing rate trajectory with the trajectory of its cluster's exemplar, scaled such that the maximum rate across all 16 sessions was maintained (**Figure 3.3C**), and then we calculated PV corrs. PV correlations were too high in all three regions, especially in CA2 (**Figure 3.3D**). Because we were replacing rate trajectories with exemplar trajectories, and these were built without the smallest principal components which

Figure 3.4. Replacing rate profiles with exemplars from a corresponding cluster does not succeed in recreating PV correlation curves. **(A)** Using k-means clustering, rate profiles were grouped by similarity in their principal component scores. The central point from each cluster was reconstructed to yield an exemplar rate trajectory. In this example, 30 clusters were created from the PCA scores of 6 components. The clusters are sorted in increasing order of separation by box shape. The dark purple line corresponds to the firing rates of the eight sessions recorded in a circle, while the dark green line corresponds to the firing rates of the eight sessions recorded in a square. Each box is outlined in a color that corresponds to how much the firing rates were separated by shape. The blue end of the color spectrum represents clusters with low separation by shape and the red end represents clusters with high separation by shape. **(B)** Pie charts show how many place fields from each hippocampal subregion ended up in each cluster. Accordingly, the pie chart for CA2 is largely blue, the pie chart for CA3 is largely red, and the chart for CA1 is more uniformly distributed. **(C)** Rate trajectories were created by multiplying each field's cluster's exemplar trajectory by the desired maximum rate. Dark green and purple represent original rates in square and circle, respectively; light green and purple are the reconstructed rates. **(D)** PV correlation curves from five replicates of 30 clusters from 6 PCA components. Correlation values in all three regions were higher than appropriate. Increasing the number of clusters, increasing the number of components, or artificially adding noise to the rates did not significantly impact these results. **(E)** Rate trajectories were randomly sampled from actual rate trajectories of place fields within the same region, such that the spatial properties of the field were decoupled from the rate profile, and the exact distribution of rate profiles changed. In this case, PV correlations were substantially lower than expected in CA1 and CA3.
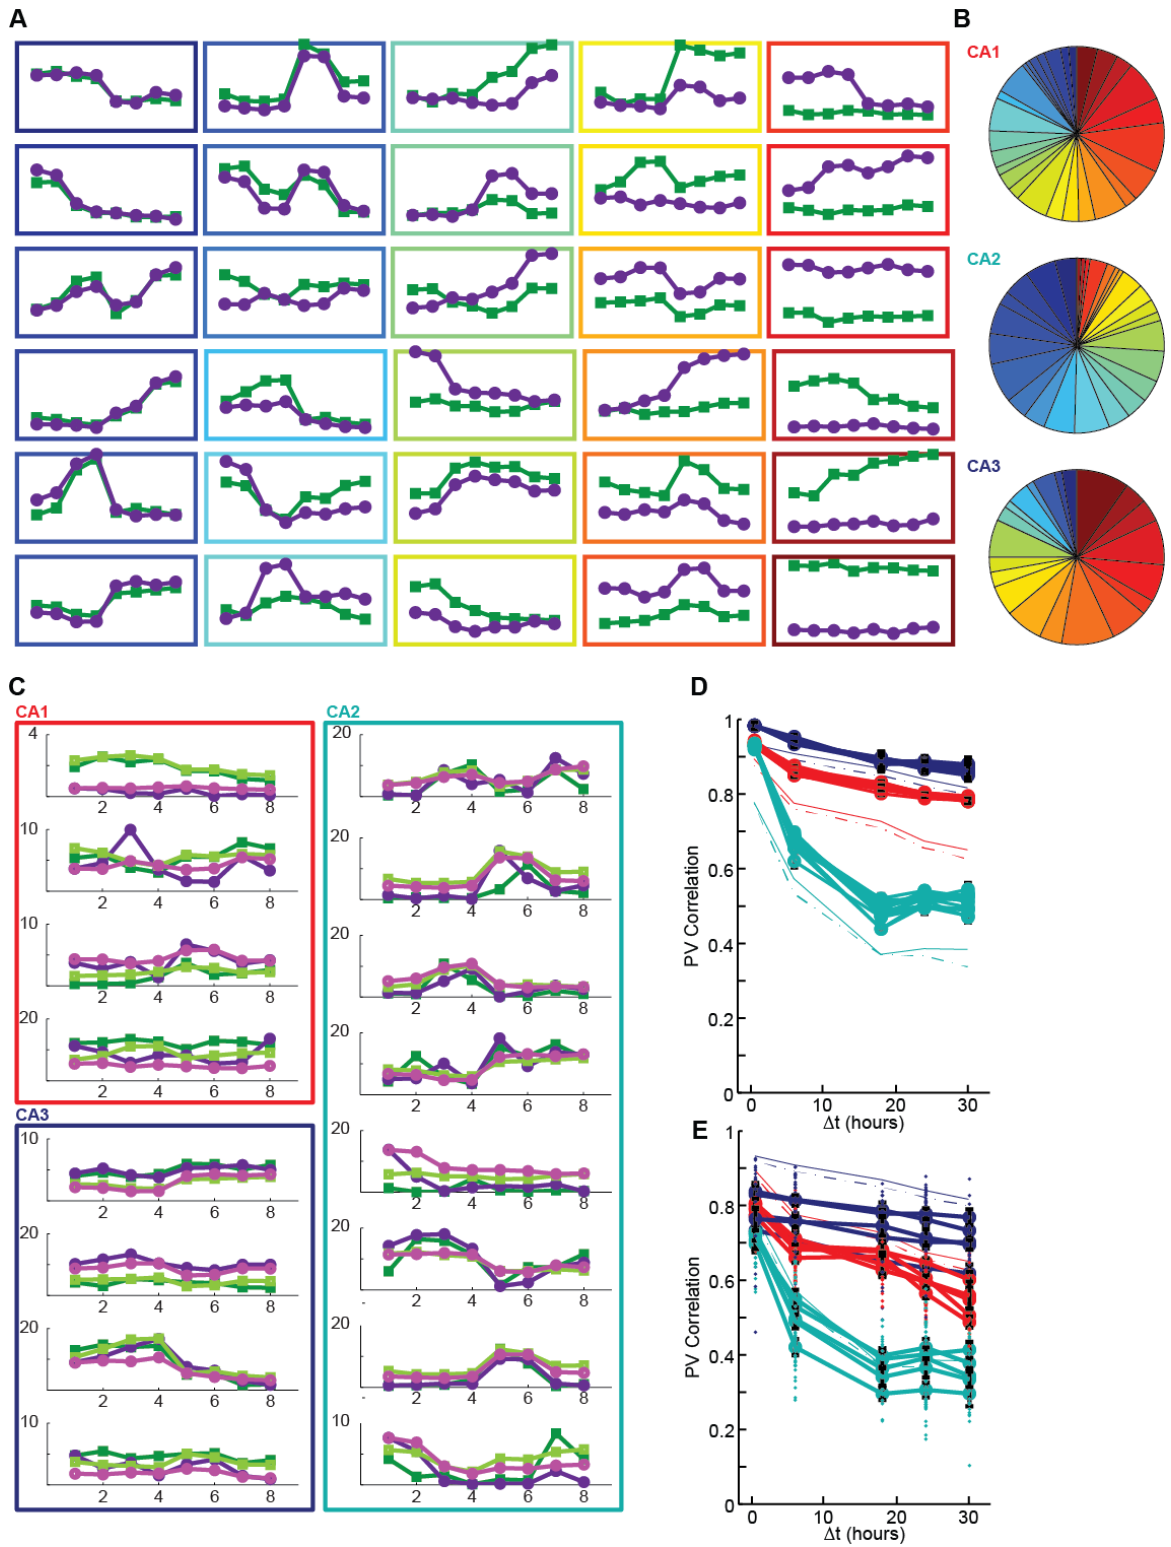
are thought of as the noise, we thought that the too-high correlations might have been because we introduced correlation amongst cells and smoothed out the noise. To test this, we added jitter, uniformly distributed between 0 and 2 Hz, to each firing rate. This did not, however, qualitatively change the results.

Taken together, these results indicate the firing rate trajectories of place fields do not naturally fall into clusters but instead are independent of each other, and this independence may be critical for the level of decorrelation observed in the data.

Finally, based on the idea that each place field is its own cluster, we wondered whether the particular firing rate trajectory mattered for a particular cell, or whether it just mattered that a realistic trajectory was chosen. We thus created a model in which, instead of choosing an exemplar trajectory from a cell's own cluster, we instead replaced each cell's firing rate trajectory with the firing rate trajectory of another cell randomly sampled from the same hippocampal subregion. We hypothesized that this should have minimal impact. We have already seen that subsampling to any 80% of the data generates narrow confidence intervals, and that the precise spatial locations of place field centers may be replaced with centers sampled with a 2 or 3 stage process. It seemed unlikely, then, that disrupting the precise combination of firing rates and place field centers evident in the data would have any impact. If anything, because firing rate trajectories were sampled with repetition, we thought this could introduce spurious correlation in the same way that repeating firing rates from a single cluster seemed to. However, the correlation values were brought down significantly with this model. CA2 actually became well-fit, with the correlation values at longer ΔTs being brought into the appropriate range. CA3, however, lost a lot of correlation, especially at short ΔTs (**Figure 3.4E**). This result indicates that the relationship between firing rate profiles and other spatial properties of place fields may not be easily decoupled, at least in CA1 and CA3. Further exploration is needed to understand these relationships.

An inheritance model for CA1

The second goal of this project was to ask whether CA1 could passively inherit the properties of spatial and temporal context discrimination from CA2 and CA3. To do this, we set up a simple network in which cells from CA2 and CA3 projected onto cells from CA1, and spatial maps in CA1 were generated as a simple sum of the spatial maps from CA2 and CA3 cells multiplied by a weight from a weight matrix. We present here preliminary results, acknowledging that the assumptions we made were over-simplified and many additional parameters might improve network performance.

In particular, we chose not to set up the model as a learning network that would develop spatial selectivity through many iterations of refining weights, but rather imposed a weight matrix that would encourage spatial selectivity from the outset, at least in one session. Thus, there was no competition for weights or anything else amongst CA1 cells. In fact, the cells of CA1 did not interact with each other in terms of defining the weight matrix. Instead, we created 100 desired maps by sampling centers uniformly across the box and by choosing goal rates from a normal distribution with mean 6 Hz and standard deviation 3 Hz. Goal weights that were sampled at less than 2 Hz were set to 2 Hz. To define the input weights to each field, we took advantage of an observation made by Cheng and Frank when they were setting up a grid cell-to-place cell model. Specifically, given a desired location for a place field, any training algorithm will converge on a set of weights in which the weight from an input cell is inversely related to the distance from the input cell's peak location to the desired location (Cheng & Frank, 2011). Thus we created a matrix in which the weight from input cell i to output cell j was given by one over the squared distance between the two place field centers. We then imposed competition among input elements by limiting the number of input cells that could have non-zero weights from each region to any given cell and replacing all but the n highest weights with 0. For defining the weights, we needed to specify from which session we would calculate a cell's peak location, as we have seen that that

is variable over time and across spatial contexts. For half of the cells, we used data from the first

square session to "train" our weight matrix. For the other half, we used data from the first circle

session (**Figure 3.5A**). By defining the weights to each cell based on only one session of input,

we could examine to what extent the variability in the input regions was transfered to the output

region. Using the static weight matrix, we calculated output maps in all 16 regions. We then

scaled the output maps of each cell so that the relative rates were maintained between sessions but

the peak rate was aligned to the goal rate selected for the cell.

First, we asked whether the output network resembled its input network if only one

region contributed positive weights (**Figure 3.5B**). When only input from CA3 was given (we

allowed 3 non-zero inputs per cell), the PV correlations of output network resembled those of

CA3. Actually, PV corrs were even more stable than expected for CA3. Perhaps any small degree

of variability in one input cell got smoothed out by the stability of the other cells that had input to

the output cell. We also looked at the PV corrs between sessions in different box shapes (compare

to Figure 2.4 of this dissertation). We expect that in CA3, the cross-shape PV corrs should be

around 0.5 and in CA1, they should be around 0.4. We did not observe those values in this data.

Instead, we found that cross shape PV corrs in the output model were stayed between

approximately 0.7 and 0.8. Though the magnitude of the cross-shape decorrelation was much less

than in the real data, it is clear that the output network did inherit some degree of shape

discrimination. Conversely, when we allowed only input from CA2, the output network showed a

decay in correlation over time that was between what would have been expected for CA2 and for

CA1, and the output network did not show reliable shape discrimination (**Figure 3.5B**).

Finally, we asked whether the output network could simultaneously inherit temporal and

spatial contextual coding. We generated networks with different ratios of input from CA2 and

CA3. Because we had seen that restricting the number of non-zero weights was useful for

network performance (data not shown), we created 1-3, 3-3, and 3-1 network architectures, where

Figure 3.5. A network with feed-forward input from CA2 and CA3 replicated some properties of CA1 over time. **(A)** The network architecture. One hundred "CA1 cells" were generated with input from CA2 and CA3. The small black squares represent individual CA1 cells. The network is shown as divided, because weights to half of the cells were determined by network activity in CA2/CA3 during the first session in a square box and the other half were determined by network activity during the first session in a circular box. The network stack below the CA1 network represent the activity of CA2 and CA3 networks that changed over time and across different shapes. The first and third network in each stack represent the first square and circle sessions, respectively. These are shown in darker lines than the other sessions, as they were the only sessions used to define the weight matrix from the input regions to CA1. The arrows from CA2 to CA1 and from CA3 to CA1 are shown in different stroke weights, indicating that the relative influence from each region could be parameterized. See text for further description of how weights were determined. After the weight matrix had been established, CA1 firing rate maps were generated for all 16 sessions, and their PV correlations were computed. **(B)** Left: PV correlation curves for 25 replicates of a network with input from only CA3 cells. Thick solid lines show PV correlation between sessions of the same shape, thick dotted lines show PV correlation between sessions in different shapes. Thin lines show best fit and actual data PV correlations from each subregion, as in previous figures. The network that received input only from CA3 shows high PV correlations for each ΔT, consistent with the CA3, and also shows some distinction between same-shape and different-shape correlations, though to a much lesser degree than the actual CA3 network, in which cross-shape PV correlation is approximately 0.5. Right: PV correlation curves for 25 replicates of a network with input only from CA2 cells. All lines are as described for the CA3-only network. In this case, PV correlation drops off with time, though to a lesser degree than in the CA2 network itself. There is also no indication of differentiation between same-shape and different-shape PV correlations, consistent with the CA2 network. **(C)** Network performance for three ratios of CA2:CA3 input. Left: each CA1 cell received positive weight from 3 CA2 cells and 1 CA3 cell. Middle: each CA1 cell received positive weight from 3 CA2 and 3 CA3 cells. Right: each CA1 cell received positive weight from 1 CA2 cell and 3 CA3 cells. The network on the right most closely replicates CA1's actual PVcorr curve while still imparting at least some degree of differentiation between same-shape and different-shape correlations. **(D)** Spatial maps across 16 sessions of 4 example CA1 cells generated by a network in which each cell received positive weight from 1 CA2 cell and 3 CA3 cells. Each row of 8 shows maps from a single day, each set of 4 maps represents a block, and maps are ordered within the block with the two square sessions first in order to facilitate observation of rate coding for box shape. The color scale is from dark blue (0 Hz) to red (peak rate for the cell across all 16 sessions). Peak rates are indicated to the left of each cell. **(E)** For each place field, shape preference was calculated as abs((A-B)/(A+B)), where A was the mean across rates in square boxes and B was the mean across rates in circles. Preference for day or block was calculated analogously, except that for day preference A and B were the means of rates on day 1 and day 2 respectively, and for each day, block preference was calculated with A and B as the means across the first four and last four sessions of the day. Time preference for each place field was taken as the maximum across day preference, block preference on day 1, and block preference on day 2. For each field in actual CA1 data, time preference is plotted against shape preference as a white dot. The colors indicate the density of fields in each subregion of this space. **(F)** The same calculation done in (E) is shown for each of 24 instantiations of the model exemplified in (D). In general, shape and time preference were somewhat diminished compared to actual data. There were very few cells that showed high scores for both shape and time preference. Compare upper right region of each plot where the density is close to 0 in the modeled data but close to maximal density in the real data.
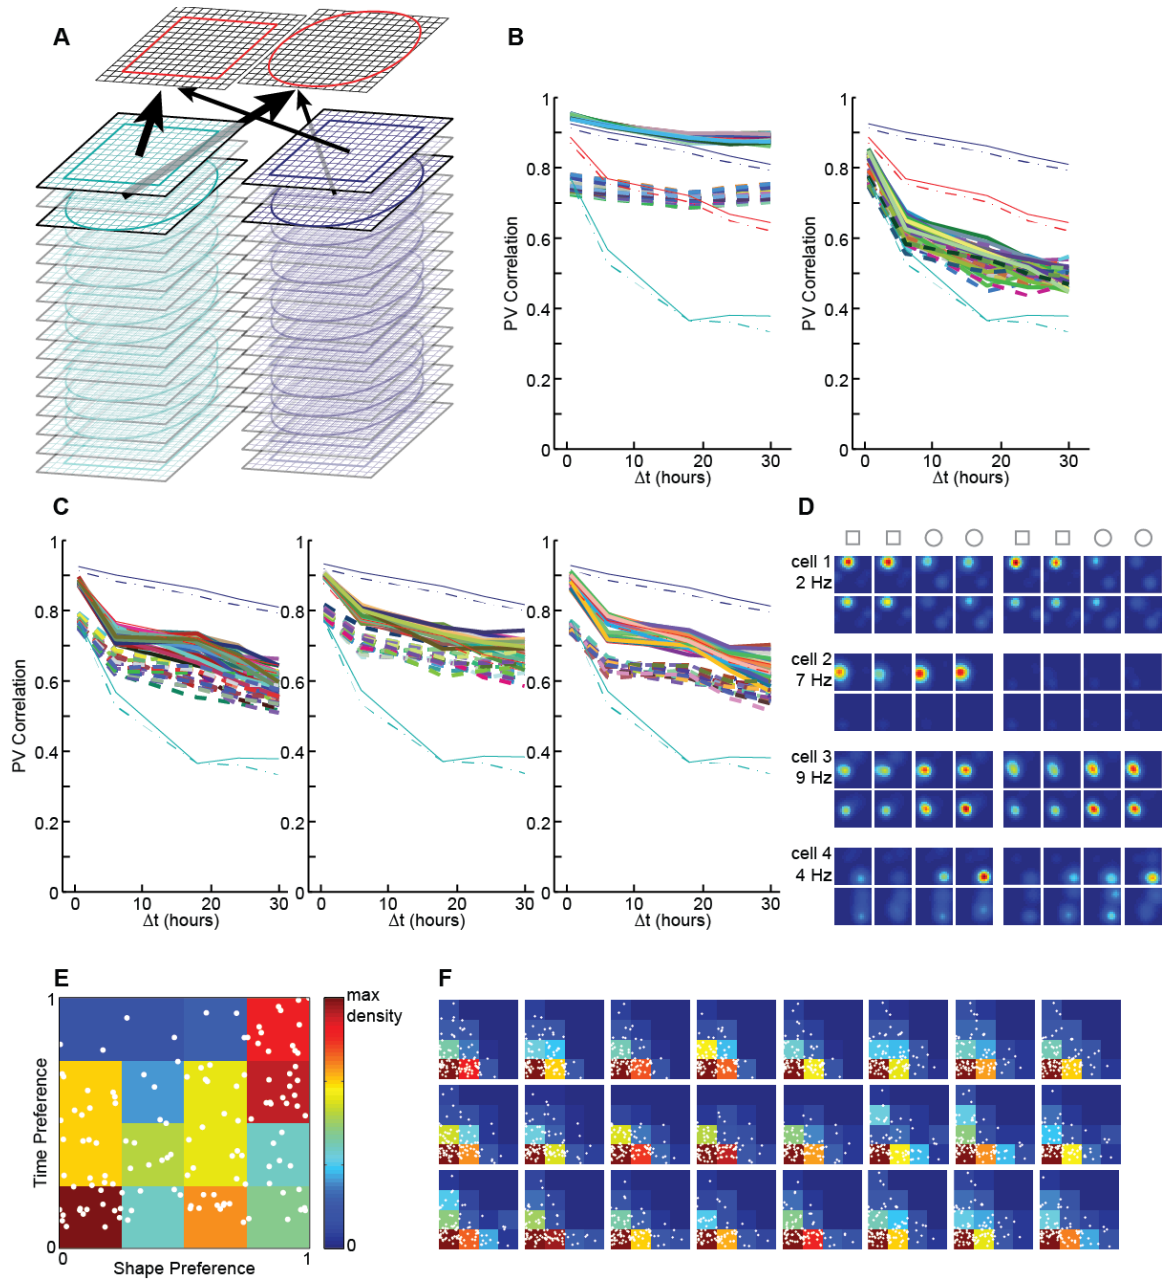
the first number indicates the maximum number of nonzero weights from CA2 onto each element of the output network, and the second number is analogous for CA3. In all three architectures, we found that the PV correlations for same-shape boxes in the output network resembled the expected output of CA1. There was also evidence of spatial context discrimination, but as observed in the CA3-only network, the separation between same-shape and cross-shape PV correlation was not as strong as in the actual CA1 network (**Figure 3.5C**). Nonetheless, in all three models with input from each subregion, the output region simultaneously inherited some degree of spatial and contextual coding.

We also looked at the spatial maps of the output cells. Sixteen spatial maps are shown for each of the first four output cells from a 3-1 network (**Figure 3.5D**). They are ordered with time increasing from left to right on each day (top row, day 1; bottom row, day 2), and with both square sessions shown before the circle sessions in each block. These cells resemble different aspects of true CA1 cells. Cell 1 has two fields; the high rate field fires preferentially in squares, and its peak rate diminishes slowly across time, while the low-rate field fires preferentially in circles. Cell 2 shows minimal preference for one box shape or the other, but is primarily active only in the first block. Cell 3 shows a moderate degree of shape preference. And cell 4 fires more in circles and on day 1. Thus, individual cells in the network seem to inherit properties common to CA2, such as temporal coding or multiple place fields, or CA3, such as strong shape preference. Individual cells may even express these properties simultaneously, such as cell 4, which has both shape and time preference, which is in line with observed CA1 cells (see, for example, Figure. 1.2A). Because most individual cells take on shape or temporal coding patterns, the network shows population coding for both spatial and temporal context.

Discussion

We proposed in the previous chapter that spatial and rate variability were likely responsible for the change in PV correlation over time observed in CA2. It has been shown previously that rate changes in CA3 are critical for coding differences between spatial contexts. Thus, we sought to characterize the ways in which place fields changed their central locations and firing rates over time in each hippocampal subregion.

We found that place fields in CA3 underwent very small degrees of spatial variability, with all central locations across 16 sessions staying within a couple centimeters of each other. Place fields in CA2, on the other hand, moved enough as time passed to require a 3-stage model of place field locations to account for sessions that are nearby in time having place fields that are nearby in space.

We found that the basic building blocks of firing rate trajectories included shape preference and temporal preference on time scales of days and blocks. However, it does not appear that firing rate trajectories form natural clusters, suggesting that place field modulation across long time scales is not coordinated across cell assemblies that fluctuate together. Further exploration is necessary to elucidate the relationship between the components identified by PCA and how these are translated into firing rate trajectories. Would a random selection of coefficients from some distribution succeed in recreating the PV corr curves of the original data? How would the presence of cell assemblies affect the PV corr values? That is, is it necessary that all firing rate curves be independent in order to reach PV corr values around 0.35? Why did correlation values drop in CA1 and CA3 when firing rate trajectories were decoupled from the spatial aspects of place fields? Although this result makes generating new cell sets that approximate true data more difficult, it may also indicate that the change in PV corr over time is not merely an emergent property of a noisy network, but rather that there is an underlying structure to the way the place

fields vary that is optimized for producing representations whose similarities change gradually as a function of time.

Finally, we provided a proof of principle that a network of cells receiving input from CA2 and CA3 cells may inherit temporal and spatial context coding simultaneously. A more sophisticated network architecture, competition process, or learning rule may provide results that more closely fit observed data, but these preliminary results provide evidence that the model proposed at the end of chapter 2 of this dissertation is plausible. That is, that CA2 provides a gradually changing signal alongside a stable signal from CA3 that may be combined in CA1, such that CA1 may simultaneously represent information about spatial and temporal context.

Acknowledgements

CHAPTER 4: CONCLUDING REMARKS AND FUTURE DIRECTIONS

Memories that contain information about the spatial and temporal context of an event are important for guiding behavior. It is clear that humans and animals are able to synthesize *what-where-and-when* information as part of their memories, and can use this information to flexibly modify future decisions. Many experimental paradigms have required animals to remember the timing of an even on the time scale of hours to days (Babb & Crystal, 2005, 2006b; Clayton & Dickinson, 1998; Dere et al., 2005; DeVito et al., 2009; W. a Roberts et al., 2008), yet little previous work has focused on understanding the coding mechanisms in the brain that might enable them to do that.

In this dissertation, we have attempted to fill the gap between experimental paradigms demonstrating contextual memory on time scales of hours to days, and both the physiology experiments and theory that have focused either on much shorter or much longer time scales (Aimone et al., 2006; Block & Reed, 1978; Friedman, 1993; MacDonald et al., 2011; Naya & Suzuki, 2011; Pastalkova et al., 2008). The hippocampus is a brain region thought to be critical for associating context to memories (Squire, 1992). By recording the activity of principal cells in the hippocampus as rats explored two spatial contexts at different times across two to three days, we were able to look at the impact of changing temporal and spatial context on hippocampal representations of space. We distinguished between cells recorded in three subregions of the hippocampus, CA1, CA2, and CA3, which allowed us to probe whether the distinct subregions performed distinct computations in support of temporal and spatial context coding.

We found that population representations in CA3 were highly stable, even for long periods of time, but they showed strong discrimination between two spatial contexts. Place fields that responded strongly in one box shape often responded much less to a second box shape, even to the point of having appeared to turn off in the second box. This confirmed previous reports that CA3 uses rate remapping to encode differences between spatial contexts (S. Leutgeb et al., 2005).

163

In a strikingly different pattern, we found that population representations in CA2 discriminated minimally between spatial contexts but changed dramatically over periods of a few hours. CA2 appeared to use rate-remapping to code for temporal context, as different place fields in CA2 cells seemed to turn off for some amount of time but then reappear.

CA1, the final output region of the hippocampus, often shows more graded responses than the dentate gyrus or CA3, responding linearly to changes in its inputs rather than performing pattern completion or pattern separation (Guzowski et al., 2004; S. Leutgeb et al., 2004; Vazdarjanova & Guzowski, 2004). For example, when subsets of well-known cues are placed in conflict, instead of an all-or-none response as seen in CA3, some place cells follow one set of cues while others follow the other set of cues (Vazdarjanova & Guzowski, 2004). In a similar way, we found that some cells in CA1 showed temporal remapping, turning on and off for blocks of time, as if they were following CA2, and some cells showed spatial remapping, turning on or off in one box shape, as if they were following CA3. A single cell in CA1 could even follow both; a cell that showed rate remapping between spatial contexts in most blocks might be off in both contexts for one block of time. On the population level, CA1 showed strong discrimination for spatial context, much like CA3. Like CA2, however, the population vector correlation for identical events decreased as a function of the time between the events in CA1, although this decrease was slower than the decrease in CA2. Thus CA1 seems to integrate inputs from both CA2 and CA3.

Theoretical considerations of how time might be incorporated into memory have suggested that "contextual change" could be a critical component of a temporal code. (Block & Reed, 1978; Friedman, 1993). Furthermore, many computational models of associative memory or free recall incorporate at time-dependent context signal in order to account for data on the temporal structure of how words are recalled (Estes, 1955; Howard & Kahana, 2002; G. Mensink & Raaijmakers, 1989; G.-J. Mensink & Raaijmakers, 1988; Murdock, 1997). Additionally, it is

assumed that the similarity of the context signal at two points in time decays as a function of the elapsed time between them. The time scale of this decay is probably indicative of the time scale over which a context signal may provide information about time. The hippocampus seems to house temporal context codes at several time scales. For example, "time cells" fire at specific time points during a short delay period for up to a few seconds (MacDonald et al., 2011). The similarity in the population activity of time cells thus falls off over a few seconds, indicating that time cells could be useful for providing timing information on this scale but are unlikely to provide timing information for many minutes or longer. The time scale of neurogenesis in the dentate gyrus, on the other hand, means that the similarity of the set of newborn neurons changes gradually over weeks, and could provide temporal context information on long time scales (Aimone et al., 2006).

We found that the population similarity in CA2 decayed over the course of 18-24 hours. This suggests that CA2 could be providing a temporal context signal that would support timing decisions on the scale of hours to days. CA1 also shows changes on this time scale, but its representations change more slowly, suggesting that CA1 may support timing decisions across intervals of up to several days. We proposed that the slower change observed in the CA1 population might come about from integrating temporally-changing input from CA2 with much more stable input from CA3, and preliminary modeling work provides support for that idea. Thus, we find that the CA3-CA2-CA1 network architecture maps nicely onto theoretical models in which a stable signal (from CA3) is combined with a temporal context signal (from CA2) in order to represent content and context simultaneously (in CA1). In Estes's original stimulus fluctuation model (1955), he proposed that temporal context might come about as elements randomly move in and out of a set of "available elements." It is interesting that the temporal context observed in CA2 comes largely from rate variability, in which place fields seem to turn on and off; the "off" place fields could be analogous to the elements that are unavailable in Estes's model.

We are proposing that CA2 plays an important role in providing temporal context to CA1 on the scale of hours to days. It is worth noting, however, that even though the data are consistent with theoretical models, there is much work left to do to prove that the fluctuations in CA2 response to a given event is truly important for an animal's ability to do timing tasks on this time scale. In particular, two lines of research could be pursued to address this question. First, a causal role for CA2 in timing should be established. A transgenic mouse model in which CA2 can be silenced was recently developed, which allows for optogenetic "lesioning" of CA2 (Hitti & Siegelbaum, 2014). Does silencing CA2 impair an animal's ability to perform *what-where-when* tasks on the time scale of hours to days? Knocking out vasopressin1b receptors, which are selectively enriched in CA2, seems to impair the *when*-component of a memory task that relies on novel object recognition (DeVito et al., 2009), providing a proof of principle. However, further work in which CA2 output is not conflated with a particular receptor that is expressed there will be necessary. A task in which animals receive reward for successfully encoding time, such as the task of Babb and Crystal (2005), would also provide more solid evidence. If the optogenetic silencing of CA2 is reversible, it would be especially interesting to test mice that have been well trained on the Babb and Crystal task to see whether silencing CA2 impairs their temporal perception. One especially convincing result would be if silencing CA2 during the delay period could cause a well-trained mouse to treat a 4-hour period as a 30-minute period. Potentially this could be done either by silencing CA2 for 3.5 hours so that it has only 30 minutes during which it can evolve, or perhaps silencing CA2 very briefly 30 minutes before the end of the delay period could reset the network somehow, again leaving only 30 minute for the CA2 and CA1 network representations to evolve..

It would also be informative to record from cells in CA1 while CA2 is silenced. Do the PV correlations in CA1 stay high, reflecting the stable input from CA3 and missing decorrelating influence of CA2?

Another line of research would involve further recordings of cells in CA2 while animals are trained to attend to temporal, spatial, or social cues. Mice that have no incentive to pay attention to spatial cues have less stable place fields than those who must learn to navigate to a remembered location (C. G. Kentros et al., 2004). It is possible, then, that we observe a lack of spatial stability in CA2 cells because the random foraging task does not require rats to encode space. It would be informative to see whether the PV correlation curves change in CA2, or any hippocampal subregion, if a spatial memory demand is imposed. This would be somewhat surprising, as CA2 has not been implicated in playing a special role in spatial context or navigation. In fact, silencing CA2 has no impact on contextual fear conditioning or learning of the Morris water maze (Hitti & Siegelbaum, 2014), so we hypothesize that a spatial memory demand would not impact the representations of CA2. On the other hand, CA2 has been implicated in social recognition memory (Hitti & Siegelbaum, 2014; Pagani et al., 2014). Thus, it would be informative record from CA2 during a task in which social memory is required. Would CA2 representations become stable after an experience with another rat? Would CA2 representations show rate remapping between experiences in the presence of different animals? If CA2 representations either became stable or changed at an even faster pace in the presence of other animals or a task requiring attention, it would be interesting to test whether these tasks also affect temporal perception.

The idea that other experiences might change coding in CA2 and consequently change perception of elapsed time is actually consistent with the subjective experience that "time flies when you're having fun", or more specifically, when a lot is happening. Block's contextual change hypothesis was predicated on the notion that subjective experience of elapsed time is based on the degree to which internal context has changed since the event being recalled (Block & Reed, 1978). He found that asking subjects to do two types of tasks in between learning a list of words and judging the recency of that learning increased subjects' estimates of elapsed time

more than doing one task twice as much. He proposed that doing the two types of task increased internal context change, though he was unable to measure that directly. We found that when rats performed random foraging in two box shapes, representations in CA1 and CA2 showed stronger changes than when they performed random foraging in only one box shape (compare Figures 2.4C and 2.6C), consistent with Block's hypothesis, though we did not measure the subjective judgment of elapsed time directly.

Autobiographical memories contain information not only about what happened but also where and when. Many autobiographical memories have largely overlapping sensory content, since much of people's lives tend to be played out within only a few locations, such as home and work. Thus, the sensory content of an experience is usually not sufficient to provide information about the time at which the event occurred. Nonetheless, humans and animals are able to use relative recency of events on time scales of hours to days to make appropriate behavioral choices, indicating that the brain must have a method for processing temporal information that is distinct from its processing of sensory content. In this dissertation we provide evidence from direct physiological recordings in behaving rats that representations of identical events change over time in hippocampal subregions CA1 and CA2. Thus, the degree of similarity between representations of two events could serve as a measure of the elapsed time between them for periods of up to approximately 24 hours.

It has been said that "time is what keeps everything from happening all at once" (Ray Cummings). Perhaps, then, CA2 is what keeps our memories of everything from seeming to have happened all at once.

REFERENCES

Aimone, J. B., Deng, W., & Gage, F. H. (2014). Adult Neurogenesis in the Dentate Gyrus. In D. Derdikman & J. J. Knierim (Eds.), *Space, time and memory in the hippocampal formation2*. Heidelberg: Springer.

Aimone, J. B., Wiles, J., & Gage, F. H. (2006). Potential role for adult neurogenesis in the encoding of time in new memories. *Nature Neuroscience*, *9*(6), 723–7.

Allen, K., Rawlins, J. N. P., Bannerman, D. M., & Csicsvari, J. (2012). Hippocampal place cells can encode multiple trial-dependent features through rate remapping. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *32*(42), 14752–66.

Allen, T. a & Fortin, N. J. (2013). The evolution of episodic memory. *Proceedings of the National Academy of Sciences of the United States of America*, *110 Suppl* , 10379–86.

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

Anagnostaras, S. G., Gale, G. D., & Fanselow, M. S. (2001). Hippocampus and contextual fear conditioning: recent controversies and advances. *Hippocampus*, *11*(1), 8–17.

Anderson, M. & Jeffery, K. (2003). Heterogeneous modulation of place cell firing by changes in context. *The Journal of Neuroscience*, *23*(26), 8827–8835.

Babb, S. J. & Crystal, J. D. (2005). Discrimination of what, when, and where: Implications for episodic-like memory in rats. *Learning and Motivation*, *36*(2), 177–189.

Babb, S. J. & Crystal, J. D. (2006a). Discrimination of what, when, and where is not based on time of day. *Learning & Behavior*, *34*(2), 124–130.

Babb, S. J. & Crystal, J. D. (2006b). Episodic-like memory in the rat. *Current Biology : CB*, *16*(13), 1317–21.

Barnes, C., McNaughton, B., Goddard, G. V., Douglas, R. M., & Adamec, R. (1977). Circadian rhythm of synaptic excitability in rat and monkey central nervous system. *Science*, *197*(4298), 91–92.

Bartesaghi, R. & Gessi, T. (2004). Parallel activation of field CA2 and dentate gyrus by synaptically elicited perforant path volleys. *Hippocampus*, *14*(8), 948–63.

Bartesaghi, R., Migliore, M., & Gessi, T. (2006). Input-output relations in the entorhinal cortex-dentate-hippocampal system: evidence for a non-linear transfer of signals. *Neuroscience*, *142*(1), 247–65.

Benes, F. M., Kwok, E. W., Vincent, S. L., & Todtenkopf, M. S. (1998). A reduction of nonpyramidal cells in sector CA2 of schizophrenics and manic depressives. *Biological Psychiatry*, *44*(2), 88–97.

Block, R. (1985). Contextual coding in memory: Studies of remembered duration. *Time, Mind, and Behavior*.

Block, R. & Reed, M. (1978). Remembered duration: Evidence for a contextual-change hypothesis. *Journal of Experimental Psychology: Human Learning & Memory*, *4*(6), 656–665.

Bolles, R. C. & De Lorge, J. (1962). The rat's adjustment to a-diurnal feeding cycles. *Journal of Comparative and Physiological Psychology*,*55*(5), 760–762.

Bolles, R. C. & Moot, S. a. (1973). The rat's anticipation of two meals a day. *Journal of Comparative and Physiological Psychology*, *83*(3), 510–514.

Bolles, R. C. & Stokes, L. W. (1965). Rat's anticipation of diurnal and a-diurnal feeding. *Journal of Comparative and Physiological Psychology*, *60*(2), 290–294.

Bostock, E., Muller, R. U., & Kubie, J. L. (1991). Experience-dependent modifications of hippocampal place cell firing. *Hippocampus*, *1*(2), 193–205.

Bouton, M. (1993). Context, time, and memory retrieval in the interference paradigms of Pavlovian learning. *Psychological Bulletin*, *114*(1), 80–99.

Brown, G. D. a., Preece, T., & Hulme, C. (2000). Oscillator-based memory for serial order. *Psychological Review*, *107*(1), 127–181.

Buhusi, C. V & Meck, W. H. (2005). What makes us tick? Functional and neural mechanisms of interval timing. *Nature Reviews. Neuroscience*, *6*(10), 755–65.

Buzsáki, G. & Moser, E. I. (2013). Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nature Neuroscience*, *16*(2), 130–8.

Carr, J. a. . & Wilkie, D. M. (1999). Rats are reluctant to use circadian timing in a daily time–place task. *Behavioural Processes*, *44*(3), 287–299.

Carr, J. a. R. & Wilkie, D. M. (1997). Rats use an ordinal timer in a daily time-place learning task. *Journal of Experimental Psychology: Animal Behavior Processes*, *23*(2), 232–247.

Caruana, D. a, Alexander, G. M., & Dudek, S. M. (2012). New insights into the regulation of synaptic plasticity from an unexpected place: hippocampal area CA2. *Learning & Memory (Cold Spring Harbor, N.Y.)*, *19*(9), 391–400.

Cheng, S. & Frank, L. M. (2011). The structure of networks that produce the transformation from grid cells to place cells. *Neuroscience*, *197*, 293–306.

Chevaleyre, V. & Siegelbaum, S. a. (2010). Strong CA2 pyramidal neuron synapses define a powerful disynaptic cortico-hippocampal loop. *Neuron*, *66*(4), 560–72.

Church, R. & Gibbon, J. (1982). Temporal generalization. *Journal of Experimental Psychology: Animal Behavior Processes.*, *8*(2), 165–186.

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

Clayton, N. S., Yu, K. S., & Dickinson, A. (2001). Scrub jays (Aphelocoma coerulescens) form integrated memories of the multiple features of caching episodes. *Journal of Experimental Psychology: Animal Behavior Processes*, *27*(1), 17–29.

Colgin, L. L., Leutgeb, S., Jezek, K., Leutgeb, J. K., Moser, E. I., McNaughton, B. L., & Moser, M.-B. (2010). Attractor-map versus autoassociation based attractor dynamics in the hippocampal network. *Journal of Neurophysiology*, *104*(1), 35–50.

Csicsvari, J., Hirase, H., Czurkó, A., Mamiya, A., & Buzsáki, G. (1999). Fast network oscillations in the hippocampal CA1 region of the behaving rat. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *19*, RC20.

Cui, Z., Gerfen, C. R., & Young, W. S. (2013). Hypothalamic and other connections with dorsal CA2 area of the mouse hippocampus. *The Journal of Comparative Neurology*, *521*(8), 1844–66.

Derdikman, D. & Knierim, J. J. (2014). *Space,Time and Memory in the Hippocampal Formation*. (D. Derdikman & J. J. Knierim, Eds.). Vienna: Springer Vienna.

Dere, E., Huston, J. P., & De Souza Silva, M. a. (2005). Integrated memory for objects, places, and temporal order: evidence for episodic-like memory in mice. *Neurobiology of Learning and Memory*, *84*(3), 214–21.

Deshmukh, S. S. & Knierim, J. J. (2011). Representation of non-spatial and spatial information in the lateral entorhinal cortex. *Frontiers in Behavioral Neuroscience*, *5*(October), 69.

DeVito, L. M., Konigsberg, R., Lykken, C., Sauvage, M., Young, W. S., & Eichenbaum, H. (2009). Vasopressin 1b receptor knock-out impairs memory for temporal order. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *29*(9), 2676–83.

Dietrich, A. & Allen, J. D. (1998). Functional dissociation of the prefrontal cortex and the hippocampus in timing behavior. *Behavioral Neuroscience*, *112*(5), 1043–1047.

Ebbinghaus, H. (1913). *Uber das Gedachtnis: Untersuchungen zur experimentellen Psychologie. (Memory: A Contribution to Experimental Psychology)*. (H. A. Ruger & C. E. Byessennine, Eds.) (Translated.). New York: Dover.

Eichenbaum, H. (2013). Hippocampus: remembering the choices. *Neuron*, *77*(6), 999–1001.

Eichenbaum, H., MacDonald, C. J., & Kraus, B. J. (2014). Time and the Hippocampus. In D. Derdikman & J. J. Knierim (Eds.), *Space, time and memory in the hippocampal formation2*. Heidelberg: Springer.

Estes, W. (1955). Statistical theory of spontaneous recovery and regression. *Psychological Review*, *62*(3), 145–154.

Ezzyat, Y. & Davachi, L. (2014). Similarity breeds proximity: pattern similarity within and across contexts is related to later mnemonic judgments of temporal proximity. *Neuron*, *81*(5), 1179–89.

Fenton, A. A. & Muller, R. U. (1998). Place cell discharge is extremely variable during individual passes of the rat through the firing field. *Proceedings of the National Academy of Sciences of the United States of America*, *95*, 3182–3187.

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

Fortin, N. J., Agster, K. L., & Eichenbaum, H. B. (2002). Critical role of the hippocampus in memory for sequences of events. *Nature Neuroscience*, *5*(5), 458–62.

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

Frank, L. M., Stanley, G. B., & Brown, E. N. (2004). Hippocampal plasticity across multiple days of exposure to novel environments. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *24*(35), 7681–9.

Frankland, P. W., Cestari, V., Filipkowski, R. K., McDonald, R. J., & Silva, A. J. (1998). The dorsal hippocampus is essential for context discrimination but not for contextual conditioning. *Behavioral Neuroscience*, *112*(4), 863–874.

Friedman, W. J. (1993). Memory for the time of past events. *Psychological Bulletin*, *113*(1), 44–66.

Fyhn, M., Hafting, T., Treves, A., Moser, M.-B., & Moser, E. I. (2007). Hippocampal remapping and grid realignment in entorhinal cortex. *Nature*, *446*(7132), 190–4.

Gallistel, C. R. (1990). *The Organization of Learning*. The MIT Press.

Gewirtz, J. C., McNish, K. a., & Davis, M. (2000). Is the hippocampus necessary for contextual fear conditioning? *Behavioural Brain Research*, *110*(1-2), 83–95.

Gibbon, J. & Church, R. (1981). Time left: linear versus logarithmic subjective time. *Journal of Experimental Psychology: Animal Behavior Processes*, *7*(2), 87–107.

Gibbon, J., Church, R. M., & Meck, W. H. (1984). Scalar Timing in Memory. *Annals of the New York Academy of Sciences*, *423*(1 Timing and Ti), 52–77.

Gill, P. R., Mizumori, S. J. Y., & Smith, D. M. (2011). Hippocampal episode fields develop with learning. *Hippocampus*, *21*(11), 1240–9.

Good, M. a, Barnes, P., Staal, V., McGregor, a, & Honey, R. C. (2007). Context- but not familiarity-dependent forms of object recognition are impaired following excitotoxic hippocampal lesions in rats. *Behavioral Neuroscience*, *121*(1), 218–23.

Gothard, K., Skaggs, W. E., & McNaughton, B. (1996). Dynamics of mismatch correction in the hippocampal ensemble code for space: interaction between path integration and environmental cues. *The Journal of Neuroscience*, *16*(24), 8027–8040.

Gothard, K., Skaggs, W., Moore, K., & McNaughton, B. (1996). Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. *The Journal of Neuroscience*, *16*(2), 823–835.

Guzowski, J., Knierim, J., & Moser, E. (2004). Ensemble dynamics of hippocampal regions CA3 and CA1. *Neuron*, *44*(4), 581–584.

Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, *436*(7052), 801–6.

Hargreaves, E. L., Rao, G., Lee, I., & Knierim, J. J. (2005). Major Dissociation Between Medial and Lateral Entorhinal Input to Dorsal Hippocampus, (June), 1792–1795.

Hebb, D. O. (1949). *The Organization of Behavior. The Organization of Behavior* (Vol. 911, p. 335).

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

Hirsh, R. (1974). The hippocampus and contextual retrieval of information from memory: A theory. *Behavioral Biology*, *12*, 421–444.

Hitti, F. L. (2014). Genetically targeted anatomical and behavioral characterization of the cornu ammonis 2 (CA2) subfield of the mouse hippocampus, *2*.

Hitti, F. L. & Siegelbaum, S. A. (2014). The hippocampal CA2 region is essential for social memory. *Nature*.

Holloway, F. & Wansley, R. (1973a). Multiphasic retention deficits at periodic intervals after passive-avoidance learning. *Science*, *180*(4082), 208–210.

Holloway, F. & Wansley, R. (1973b). Multiple retention deficits at periodic intervals after active and passive avoidance learning. *Behavioral Biology*, *9*(275), 1–14.

Howard, M. W. & Kahana, M. J. (1999). Contextual variability and serial position effects in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*(4), 923–941.

Howard, M. W. & Kahana, M. J. (2002). A Distributed Representation of Temporal Context. *Journal of Mathematical Psychology*, *46*(3), 269–299.

Howard, M. W. & Natu, V. S. (2005). Place from time: Reconstructing position from a distributed representation of temporal context. *Neural Networks : The Official Journal of the International Neural Network Society*, *18*(9), 1150–62.

Howard, M. W., Viskontas, I. V, Shankar, K. H., & Fried, I. (2012). Ensembles of human MTL neurons "jump back in time" in response to a repeated stimulus. *Hippocampus*, *22*(9), 1833–47.

Hsieh, L.-T., Gruber, M. J., Jenkins, L. J., & Ranganath, C. (2014). Hippocampal activity patterns carry information about objects in temporal context. *Neuron*, *81*(5), 1165–78.

Hut, R. A. & Van der Zee, E. A. (2010). The cholinergic system, circadian rhythmicity, and time memory. *Behavioural Brain Research*, *221*(2), 466–80.

Iordanova, M. D., Good, M. a, & Honey, R. C. (2008). Configural learning without reinforcement: integrated memories for correlates of what, where, and when. *Quarterly Journal of Experimental Psychology (2006)*, *61*(12), 1785–92.

Jeffery, K. J. (1998). Learning of landmark stability and instability by hippocampal place cells. *Neuropharmacology*, *37*(4-5), 677–687.

Jones, M. W. & McHugh, T. J. (2011). Updating hippocampal representations: CA2 joins the circuit. *Trends in Neurosciences*, *34*(10), 526–35.

Kempter, R., Leibold, C., Buzsáki, G., Diba, K., & Schmidt, R. (2012). Quantifying circular-linear associations: Hippocampal phase precession. *Journal of Neuroscience Methods*, *207*, 113–124.

Kennedy, P. J. & Shapiro, M. L. (2004). Retrieving memories via internal context requires the hippocampus. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *24*(31), 6979–85.

Kentros, C. (1998). Abolition of Long-Term Stability of New Hippocampal Place Cell Maps by NMDA Receptor Blockade. *Science*, *280*(5372), 2121–2126.

Kentros, C. G., Agnihotri, N. T., Streater, S., Hawkins, R. D., & Kandel, E. R. (2004). Increased Attention to Spatial Context Increases Both Place Field Stability and Spatial Memory. *Neuron*, *42*(2), 283–295.

Kesner, R. P. & Hunsaker, M. R. (2010). The temporal attributes of episodic memory. *Behavioural Brain Research*, *215*(2), 299–309.

Kim, J. & Fanselow, M. (1992). Modality-specific retrograde amnesia of fear. *Science*, *256*(5057), 675–677.

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

Knierim, J. J., Lee, I., & Hargreaves, E. L. (2006). Hippocampal place cells: parallel input streams, subregional processing, and implications for episodic memory. *Hippocampus*, *16*(9), 755–64.

Knierim, J., Kudrimoti, H., & McNaughton, B. (1995). Place cells, head direction cells, and the learning of landmark stability. *The Journal of Neuroscience*, *15*(3), 1648–1659.

Kohara, K., Pignatelli, M., Rivest, A. J., Jung, H.-Y., Kitamura, T., Suh, J., Frank, D., Kajikawa, K., Mise, N., Obata, Y., Wickersham, I. R., & Tonegawa, S. (2014). Cell type-specific genetic and optogenetic tools reveal hippocampal CA2 circuits. *Nature Neuroscience*, *17*(2), 269–79.

Lee, I., Rao, G., & Knierim, J. J. (2004). A double dissociation between hippocampal subfields: differential time course of CA3 and CA1 place cells for processing changed environments. *Neuron*, *42*(5), 803–15.

Lee, I., Yoganarasimha, D., Rao, G., & Knierim, J. (2004). Comparison of population coherence of place cells in hippocampal subfields CA1 and CA3. *Nature*, *430*(July), 456–459.

Lee, S., Simons, S., Heldt, S., Zhao, M., Schroeder, J., Vellano, C., Cowan, D., Ramineni, S., Yates, C., Fent, Y., Smith, Y., Sweatt, J., Weinshenker, D., Ressler, K., Dudek, S., & Hepler, J. (2010). RGS14 is a natural suppressor of both synaptic plasticity in CA2 neurons and hippocampal-based learning and memory. *Proceedings of the National Academy of Sciences of the United States of America*, *107*(39), 16994–16998.

Lein, E. S., Callaway, E. M., Albright, T. D., & Gage, F. H. (2005). Redefining the boundaries of the hippocampal CA2 subfield in the mouse using gene expression and 3-dimensional reconstruction. *The Journal of Comparative Neurology*, *485*(1), 1–10.

Leutgeb, J. K., Leutgeb, S., Moser, M.-B., & Moser, E. I. (2007). Pattern separation in the dentate gyrus and CA3 of the hippocampus. *Science (New York, N.Y.)*, *315*(5814), 961–6.

Leutgeb, J. K., Leutgeb, S., Treves, A., Meyer, R., Barnes, C. a, McNaughton, B. L., Moser, M.-B., & Moser, E. I. (2005). Progressive transformation of hippocampal neuronal representations in "morphed" environments. *Neuron*, *48*(2), 345–58.

Leutgeb, S. & Leutgeb, J. K. (2007). Pattern separation, pattern completion, and new neuronal codes within a continuous CA3 map. *Learning & Memory (Cold Spring Harbor, N.Y.)*, *14*, 745–757.

Leutgeb, S. & Leutgeb, J. K. (2014). Remapping to discriminate contexts with hippocampal population codes. In D. Derdikman & J. J. Knierim (Eds.), *Space, time and memory in the hippocampal formation*. Heidelberg: Springer.

Leutgeb, S., Leutgeb, J. K., Barnes, C. a, Moser, E. I., McNaughton, B. L., & Moser, M.-B. (2005). Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science (New York, N.Y.)*, *309*(5734), 619–23.

Leutgeb, S., Leutgeb, J. K., Treves, A., Moser, M.-B., & Moser, E. I. (2004). Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science (New York, N.Y.)*, *305*(5688), 1295–8.

Lever, C., Wills, T., Cacucci, F., Burgess, N., & Keefe, J. O. (2002). Long-term plasticity in hippocampal place-cell representation of environmental geometry, *416*(March), 236–238.

Liu, X., Ramirez, S., Pang, P. T., Puryear, C. B., Govindarajan, A., Deisseroth, K., & Tonegawa, S. (2012). Optogenetic stimulation of a hippocampal engram activates fear memory recall. *Nature*, *484*(7394), 381–5.

Lorente de No, R. (1934). Studies on the Structure of the Cerebral Cortex. II. Continuation of the Study of the Ammonic System. *Journal Fur Psychologie Und Neurologie*, 113–177.

Ludvig, N. (1999). Place Cells Can Flexibly Terminate and Develop Their Spatial Firing. A New Theory for Their Function. *Physiology & Behavior*, *67*(1), 57–67.

MacDonald, C. J., Carrow, S., Place, R., & Eichenbaum, H. (2013). Distinct hippocampal time cell sequences represent odor memories in immobilized rats. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *33*(36), 14607–16.

MacDonald, C. J., Lepage, K. Q., Eden, U. T., & Eichenbaum, H. (2011). Hippocampal "time cells" bridge the gap in memory for discontiguous events. *Neuron*, *71*(4), 737–49.

Maglóczky, Z., Acsády, L., & Freund, T. (1994). Principal cells are the postsynaptic targets of supramammillary afferents in the hippocampus of the rat. *Hippocampus*, *4*(3), 322–334.

Mankin, E. a, Sparks, F. T., Slayyeh, B., Sutherland, R. J., Leutgeb, S., & Leutgeb, J. K. (2012). Neuronal code for extended time in the hippocampus. *Proceedings of the National Academy of Sciences of the United States of America*, *109*, 19462–7.

Manning, J. R., Polyn, S. M., Baltuch, G. H., Litt, B., & Kahana, M. J. (2011). Oscillatory patterns in temporal lobe reveal context reinstatement during memory search. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(31), 12893–7.

Manns, J. R., Howard, M. W., & Eichenbaum, H. (2007). Gradual changes in hippocampal activity support remembering the order of events. *Neuron*, *56*(3), 530–40.

Markus, E. J., Barnes, C. A., McNaughton, B. L., Gladden, V. L., & Skaggs, W. E. (1994). Spatial information content and reliability of hippocampal CA1 neurons: effects of visual input. *Hippocampus*, *4*(4), 410–21.

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

Martig, A. K. & Mizumori, S. J. Y. (2011). Ventral tegmental area disruption selectively affects CA1/CA2 but not CA3 place fields during a differential reward working memory task. *Hippocampus*, *21*(2), 172–84.

Maurer, A. P., Vanrhoads, S. R., Sutherland, G. R., Lipa, P., & McNaughton, B. L. (2005). Self-motion and the origin of differential spatial scaling along the septo-temporal axis of the hippocampus. *Hippocampus*, *15*(7), 841–52.

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

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

McNaughton, B. L., Leonard, B., & Chen, L. (1989). Cortical-hippocampal interactions and cognitive mapping: A hypothesis based on reintegration of the parietal and inferotemporal pathways for visual processing. *Psychobiology*, *17*(3), 230–235.

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

Meck, W. & Church, R. (1984). Simultaneous temporal processing. *Journal of Experimental Psychology: Animal Behavior Processes*, *10*(1), 1–29.

Mehta, M., Barnes, C. A., & McNaughton, B. L. (1997). Experience-dependent, asymmetric expansion of hippocampal place fields. *Proceedings of the National Academy of Sciences of the United States of America*, *94*(August), 8918–8921.

Mensink, G. & Raaijmakers, J. (1989). A model for contextual fluctuation. *Journal of Mathematical Psychology*, *33*(2), 172–186.

Mensink, G.-J. & Raaijmakers, J. G. (1988). A model for interference and forgetting. *Psychological Review*, *95*(4), 434–455.

Mercer, A., Trigg, H. L., & Thomson, A. M. (2007). Characterization of neurons in the CA2 subfield of the adult rat hippocampus. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *27*(27), 7329–38.

Mitra, P. & Bokil, H. (2008). *Observed Brain Dynamics*. *Convergence* (p. xxii, 381 p.).

Moore, D. (2001). Honey bee circadian clocks: behavioral control from individual workers to whole-colony rhythms. *Journal of Insect Physiology*, *47*(8), 843–857.

Morris, R. (1990). Does the hippocampus play a disproportionate role in spatial memory. *Discussions in Neuroscience*, *6*, 39–45.

Morris, R., Garrud, P., Rawlins, J., & O'Keefe, J. (1982). Place navigation impaired in rats with hippocampal lesions. *Nature*, *297*(24), 681–683.

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

Muller, R., Kubie, J., & Ranck Jr, J. (1987). Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *J Neurosci*, *7*(July), 1935–1950.

Munn, R. G. K. & Bilkey, D. K. (2012). The firing rate of hippocampal CA1 place cells is modulated with a circadian period. *Hippocampus*, *22*(6), 1325–37.

Murdock, B. B. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review*, *89*(6), 609–626.

Murdock, B. B. (1997). Context and mediators in a theory of distributed associative memory (TODAM2). *Psychological Review*, *104*(4), 839–862.

Nakashiba, T., Young, J. Z., McHugh, T. J., Buhl, D. L., & Tonegawa, S. (2008). Transgenic inhibition of synaptic transmission reveals role of CA3 output in hippocampal learning. *Science (New York, N.Y.)*, *319*(5867), 1260–4.

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, & Tonegawa, S. (2002). Requirement for hippocampal CA3 NMDA receptors in associative memory recall. *Science (New York, N.Y.)*, *297*(5579), 211–8.

Nakazawa, K., Sun, L. D., Quirk, M. C., Rondi-Reig, L., Wilson, M. A., & Tonegawa, S. (2003). Hippocampal CA3 NMDA Receptors Are Crucial for Memory Acquisition of One-Time Experience. *Neuron*, *38*(2), 305–315.

Naya, Y. & Suzuki, W. a. (2011). Integrating what and when across the primate medial temporal lobe. *Science (New York, N.Y.)*, *333*(6043), 773–6.

O'Brien, J. & Sutherland, R. J. (2007). Evidence for episodic memory in a pavlovian conditioning procedure in rats. *Hippocampus*, *17*(12), 1149–52.

O'Keefe, J. (1976). Place units in the hippocampus of the freely moving rat. *Experimental Neurology*, *51*, 78–109.

O'Keefe, J. (1990). A computational theory of the hippocampal cognitive map. *Progress in Brain Research*, *83*, 301–312.

O'Keefe, J. & Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature*, *381*(5), 425–428.

O'Keefe, J. & Conway, D. (1978). Hippocampal place units in the freely moving rat: why they fire where they fire. *Experimental Brain Research*, *30*, 573–590.

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

O'Keefe, J. & Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. Clarendon Press Oxford.

O'Keefe, J. & Nadel, L. (1979). Precis of O'Keefe & Nadel's The hippocampus as a cognitive map. *Behavioral and Brain Sciences*, 487–533.

O'Keefe, J. & Recce, M. L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, *3*(3), 317–30.

O'Keefe, J, Burgess, N. Donnett, JG, Jeffery, KJ, Maguire, E. (1998). Place cells, navigational accuracy, and the human hippocampus. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 1333–1340.

O'Reilly, R. C. & Norman, K. A. (2002). Hippocampal and neocortical contributions to memory: advances in the complementary learning systems framework. *Trends in Cognitive Sciences*, *6*(12), 505–510.

Olmstead, C., Best, P., & Mays, L. (1973). Neural activity in the dorsal hippocampus during paradoxical sleep, slow wave sleep and waking. *Brain Research*, *60*, 381–391.

Olton, D., Meck, W., & Church, R. (1987). Separation of hippocampal and amygdaloid involvement in temporal memory dysfunctions. *Brain Research*, *404*, 180–188.

Pagani, J. H., Zhao, M., Cui, Z., Williams Avram, S. K., Caruana, D. a, Dudek, S. M., & Young, W. S. (2014). Role of the vasopressin 1b receptor in rodent aggressive behavior and synaptic plasticity in hippocampal area CA2. *Molecular Psychiatry*, (November 2012), 1–10.

Pan, W.-X. & McNaughton, N. (2004). The supramammillary area: its organization, functions and relationship to the hippocampus. *Progress in Neurobiology*, *74*(3), 127–66.

Pastalkova, E., Itskov, V., Amarasingham, A., & Buzsáki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. *Science (New York, N.Y.)*, *321*(5894), 1322–7.

Penick, S. & Solomom, P. (1991). Hippocampus, context, and conditioning. *Behavioral Neuroscience*, *105*(5), 611–617.

Phillips, R. G. & LeDoux, J. E. (1992). Differential contribution of amygdala and hippocampus to cued and contextual fear conditioning. *Behavioral Neuroscience*, *106*(2), 274–285.

Piatti, V. C., Ewell, L. a, & Leutgeb, J. K. (2013). Neurogenesis in the dentate gyrus: carrying the message or dictating the tone. *Frontiers in Neuroscience*, *7*(April), 50.

Quirk, G., Muller, R., & Kubie, J. (1990). The firing of hippocampal place cells in the dark depends on the rat's recent experience. *The Journal of Neuroscience*, *10*(6), 2008–2017.

Raaijmakers, J. & Shiffrin, R. (1981). Search of associative memory. *Psychological Review*, *88*(March), 93–134.

Ranck Jr, J. (1973). Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats. Part I. Behavioral Correlates and Firing Reprtoires. *Experimental Neurology*, *41*(2), 462–531.

Ratzliff, A. D. H. & Soltesz, I. (2001). Differential immunoreactivity for alpha-actinin-2, an N-methyl-D-aspartate-receptor/actin binding protein, in hippocampal interneurons. *Neuroscience*, *103*, 337–349.

Redish, A. (n.d.). MClust.

Reiff, R. & Scheerer, M. (1959). *Memory and Hypnotic Age Regression*. Internat. University Press.

Reijmers, L. G., Perkins, B. L., Matsuo, N., & Mayford, M. (2007). Localization of a stable neural correlate of associative memory. *Science (New York, N.Y.)*, *317*(5842), 1230–3.

Remondes, M. & Schuman, E. M. (2004). Role for a cortical input to hippocampal area CA1 in the consolidation of a long-term memory. *Nature*, *431*, 699–703.

Roberts, S. (1981). Isolation of an internal clock. *Journal of Experimental Psychology: Animal Behavior Processes.*, *7*(3), 242–268.

Roberts, W. a, Feeney, M. C., Macpherson, K., Petter, M., McMillan, N., & Musolino, E. (2008). Episodic-like memory in rats: is it based on when or how long ago? *Science (New York, N.Y.)*, *320*(5872), 113–5.

Rolls, E. T. (1989). Functions of neruonal networks in the hippocampus and cerebral cortex in memory. In R. Cotterill (Ed.), *Models of Brain Function* (pp. 15–33). Cambridge: Cambridge University Press.

Rowland, D. C., Weible, A. P., Wickersham, I. R., Wu, H., Mayford, M., Witter, M. P., & Kentros, C. G. (2013). Transgenically targeted rabies virus demonstrates a major monosynaptic projection from hippocampal area CA2 to medial entorhinal layer II neurons. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *33*(37), 14889–98.

Royer, S., Sirota, A., Patel, J., & Buzsáki, G. (2010). Distinct representations and theta dynamics in dorsal and ventral hippocampus. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *30*(5), 1777–87.

Samsonovich, A. & Mcnaughton, B. L. (1997). Path Integration and Cognitive Mapping in a Continuous Attractor Neural Network Model, *17*(15), 5900–5920.

Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B., Witter, M., Moser, E., & Moser, M. (2006). Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science*, *312*(5), 758–762.

Save, E., Nerad, L., & Poucet, B. (2000). Contribution of multiple sensory information to place field stability in hippocampal place cells. *Hippocampus*, *10*, 64–76.

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

Scoville, W. (1954). The Limbic Lobe in Man*. *Journal of Neurosurgery*, *11*(1), 64–66.

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

Shen, J., Barnes, C. A., Mcnaughton, B. L., Skaggs, W. E., & Weaver, K. L. (1997). The Effect of Aging on Experience-Dependent Plasticity of Hippocampal Place Cells. *The Journal of Neuroscience*, *17*(17), 6769–6782.

Simons, S. B., Caruana, D. a, Zhao, M., & Dudek, S. M. (2012). Caffeine-induced synaptic potentiation in hippocampal CA2 neurons. *Nature Neuroscience*, *15*(1), 23–5.

Simons, S. B., Escobedo, Y., Yasuda, R., & Dudek, S. M. (2009). Regional differences in hippocampal calcium handling provide a cellular mechanism for limiting plasticity. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(33), 14080–4.

Singer, A. C., Carr, M. F., Karlsson, M. P., & Frank, L. M. (2013). Hippocampal SWR activity predicts correct decisions during the initial learning of an alternation task. *Neuron*, *77*(6), 1163–73.

Skaggs, W. E. & Mcnaughton, B. L. (1998). Spatial Firing Properties of Hippocampal CA1 Populations in an Environment Containing Two Visually Identical Regions. *The Journal of Neuroscience*, *18*(20), 8455–8466.

Skaggs, W. E., McNaughton, B. L., Gothard, K. M., & Markus, E. J. (1993). An Information-Theoretic Approach to Deciphering the Hippocampal Code. *Advances in Neural Information Processing Systems*, *5*, 1030–1037.

Skaggs, W. & McNaughton, B. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*, (6), 149–172.

Sloviter, R. S. (1983). "Epileptic" brain damage in rats induced by sustained electrical stimulation of the perforant path. I. Acute electrophysiological and light microscopic studies. *Brain Research Bulletin*, *10*(5), 675–697.

Smith, D. M. & Mizumori, S. J. Y. (2006a). Hippocampal place cells, context, and episodic memory. *Hippocampus*, *16*(9), 716–29.

Smith, D. M. & Mizumori, S. J. Y. (2006b). Learning-related development of context-specific neuronal responses to places and events: the hippocampal role in context processing. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *26*(12), 3154–63.

Solstad, T., Moser, E. I., & Einevoll, G. T. (2006). From grid cells to place cells: a mathematical model. *Hippocampus*, *16*(12), 1026–31.

Squire, L. R. (1992). Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychological Review*, *99*(2), 195–231.

Stepanek, L. & Doupe, A. J. (2010). Activity in a cortical-basal ganglia circuit for song is required for social context-dependent vocal variability. *Journal of Neurophysiology*, *104*(5), 2474–86.

Suddendorf, T. & Busby, J. (2003). Mental time travel in animals? *Trends in Cognitive Sciences*, *7*(9), 391–396.

Suh, J., Rivest, A. J., Nakashiba, T., Tominaga, T., & Tonegawa, S. (2011). Entorhinal cortex layer III input to the hippocampus is crucial for temporal association memory. *Science (New York, N.Y.)*, *334*(6061), 1415–20.

Tanila, H. (1999). Hippocampal place cells can develop distinct representations of two visually identical environments. *Hippocampus*, *9*(3), 235–46.

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

Thorpe, C. M., Bates, M. E., & Wilkie, D. M. (2003). Rats have trouble associating all three parts of the time–place–event memory code. *Behavioural Processes*, *63*(2), 95–110.

Tononi, G. & Cirelli, C. (2001). Modulation of brain gene expression during sleep and wakefulness: a review of recent findings. *Neuropsychopharmacology*, *25*(S5), S28–S35.

Tononi, G. & Cirelli, C. (2006). Sleep function and synaptic homeostasis. *Sleep Medicine Reviews*, *10*(1), 49–62.

Treves, a & Rolls, E. T. (1992). Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus*, *2*(2), 189–99.

Treves, a & Rolls, E. T. (1994). Computational analysis of the role of the hippocampus in memory. *Hippocampus*, *4*(3), 374–91.

Tsao, A., Moser, M.-B., & Moser, E. I. (2013). Traces of experience in the lateral entorhinal cortex. *Current Biology : CB*, *23*(5), 399–405.

Tulving, E. (1972). Episodic and semantic memory. In *Organization of memory* (Vol. 1, pp. 381–403).

Tulving, E. (1983). *Elements of Episodic Memory*. Oxford: Clarendon Press Oxford.

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

Valnegri, P., Khelfaoui, M., Dorseuil, O., Bassani, S., Lagneaux, C., Gianfelice, A., Benfante, R., Chelly, J., Billuart, P., Sala, C., & Passafaro, M. (2011). A circadian clock in hippocampus is regulated by interaction between oligophrenin-1 and Rev-erbα. *Nature Neuroscience*, *14*(10), 1293–301.

Vazdarjanova, A. & Guzowski, J. F. (2004). Differences in hippocampal neuronal population responses to modifications of an environmental context: evidence for distinct, yet complementary, functions of CA3 and CA1 ensembles. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *24*(29), 6489–96.

Wang, M. E., Wann, E. G., Yuan, R. K., Ramos Álvarez, M. M., Stead, S. M., & Muzzio, I. a. (2012). Long-term stabilization of place cell remapping produced by a fearful experience. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *32*, 15802–14.

Wang, S.-H. & Morris, R. G. M. (2010). Hippocampal-neocortical interactions in memory formation, consolidation, and reconsolidation. *Annual Review of Psychology*, *61*, 49–79, C1–C4.

Wansley, R. & Holloway, F. (1976). Oscillations in retention performance after passive avoidance training. *Learning and Motivation*, *7*, 296–302.

Wersinger, S. R., Ginns, E. I., O'Carroll, a-M., Lolait, S. J., & Young, W. S. (2002). Vasopressin V1b receptor knockout reduces aggressive behavior in male mice. *Molecular Psychiatry*, *7*(9), 975–84.

Widman, D. R., Gordon, D., & Timberlake, W. (2000). Response cost and time-place discrimination by rats in maze tasks. *Animal Learning & Behavior*, *28*(3), 298–309.

Widman, D. R., Sermania, C. M., & Genismore, K. E. (2004). Evidence for time-place learning in the Morris water maze without food restriction but with increased response cost. *Behavioural Processes*, *67*(2), 183–93.

Wiebe, S. & Stäubli, U. (1999). Dynamic filtering of recognition memory codes in the hippocampus. *The Journal of Neuroscience*, *19*(23), 10562–10574.

Wilson, I. a, Ikonen, S., Gallagher, M., Eichenbaum, H., & Tanila, H. (2005). Age-associated alterations of hippocampal place cells are subregion specific. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *25*(29), 6877–86.

Wilson, M. A. & McNaughton, B. L. (1993). Dynamics of the Hippocampal Ensemble Code for Space. *Science*, *261*(5124), 1055–1058.

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

Wintzer, M. E., Boehringer, R., Polygalov, D., & McHugh, T. J. (2014). The hippocampal CA2 ensemble is sensitive to contextual change. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *34*(8), 3056–66.

Witter, M. P. & Moser, E. I. (2006). Spatial representation and the architecture of the entorhinal cortex. *Trends in Neurosciences*, *29*(12), 671–8.

Wood, E., Dudchenko, P., & Eichenbaum, H. (1999). The global record of memory in hippocampal neuronal activity. *Nature*, *397*(18 February), 613–616.

Wood, E. R., Dudchenko, P. a., Robitsek, R. J., & Eichenbaum, H. (2000). Hippocampal Neurons Encode Information about Different Types of Memory Episodes Occurring in the Same Location. *Neuron*, *27*(3), 623–633.

Woodhams, P., Celio, M., Ulfig, N., & Witter, M. (1993). Morphological and functional correlates of borders in the entorhinal cortex and hippocampus. *Hippocampus*, *3*, 303–312.

Wu, H. G., Miyamoto, Y. R., Gonzalez Castro, L. N., Ölveczky, B. P., & Smith, M. a. (2014). Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nature Neuroscience*, *17*(2), 312–21.

Wyszynski, M., Kharazia, V., Shanghvi, R., Rao, A., Beggs, A. H., Craig, A. M., Weinberg, R., & Sheng, M. (1998). Differential regional expression and ultrastructural localization of alpha-actinin-2, a putative NMDA receptor-anchoring protein, in rat brain. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *18*, 1383–1392.

Young, W. S., Li, J., Wersinger, S. R., & Palkovits, M. (2006). The vasopressin 1b receptor is prominent in the hippocampal area CA2 where it is unaffected by restraint stress or adrenalectomy. *Neuroscience*, *143*(4), 1031–9.

Zhang, S.-J., Ye, J., Miao, C., Tsao, A., Cerniauskas, I., Ledergerber, D., Moser, M.-B., & Moser, E. I. (2013). Optogenetic dissection of entorhinal-hippocampal functional connectivity. *Science (New York, N.Y.)*, *340*(6128), 1232627.

Zhao, M., Choi, Y.-S., Obrietan, K., & Dudek, S. M. (2007). Synaptic plasticity (and the lack thereof) in hippocampal CA2 neurons. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *27*(44), 12025–32.

Zhao, X., Lein, E. S., He, a, Smith, S. C., Aston, C., & Gage, F. H. (2001). Transcriptional profiling reveals strict boundaries between hippocampal subregions. *The Journal of Comparative Neurology*, *441*(3), 187–96.

Zhou, W. & Crystal, J. D. (2009). Evidence for remembering when events occurred in a rodent model of episodic memory. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(23), 9525–9.

Ziv, Y., Burns, L. D., Cocker, E. D., Hamel, E. O., Ghosh, K. K., Kitch, L. J., El Gamal, A., & Schnitzer, M. J. (2013). Long-term dynamics of CA1 hippocampal place codes. *Nature Neuroscience*, *16*(3), 264–6.

Zola-Morgan, S. & Squire, L. R. (1993). Neuroanatomy of memory. *Annual Review of Neuroscience*, *16*, 547–63.