THE UNIVERSITY OF CHICAGO

CONTEXT DEPENDENT ACTIVITY IN RAT HIPPOCAMPUS DURING PERFORMANCE OF A THREE ARM DELAYED SEQUENCE TASK

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE BIOLOGICAL SCIENCES
AND THE PRITZKER SCHOOL OF MEDICINE
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

BIOLOGY

BY
BRIAN R. LUSTIG

CHICAGO, ILLINOIS
DECEMBER 2018

Dedicated to Eli, Estelle, Norman, Hedy, Ron, Robbie, Tracey and Andy

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

This work would not have been possible without the training and support of Eva Pastalkova and Yingxue Wang. I am extremely grateful and proud to have been a member of the Pastalkova lab.

Thanks to Sandro Romani and Albert Lee for the continued support during exploration of SWRs.

I would also like to thank Jeff Magee, Kristin Branson, Nicho Hatsopoulos, Dan Margoliash, Stephanie Palmer, Anna Roe, Ford Ebner, and Robert Friedman.

ABSTRACT

The observed differential patterns of activity recorded in hippocampus during the delay period of spatial memory tasks has been proposed as representing a contextual memory being maintained for task completion. However, the majority of such evidence has been acquired using simple tasks that require only memory of experience from the previous trial across a delay in order to inform the choice on the next trial. Here we designed a novel three arm delayed spatial memory task that requires maintenance of context constructed from information beyond the immediate past choice. Surprisingly, we found that activity in the hippocampus across the delay period could not explain successful task performance because there was no differentiation during delays following the same immediate spatial history. In contrast, decoding of sharp wave ripple (SWR) events consistently predicted future behavior among multiple spatiotemporally remote goals. Finally, we show SWRs displayed future trajectories that included contextually specific versions of the same spatial location, providing the first evidence for the retention of contextual information within SWR recall. All together these findings suggest that despite evidence for contextually different plans across SWR activity as well as the preservation of contextual information in SWR content, contextual coding in the hippocampus across delay periods is primarily determined by immediate past spatial context alone and shows no evidence of differentiating the more global behavioral context.

CHAPTER 1

INTRODUCTION

1.1 Systems Neuroscience and Memory

Neuroscience is the study of the brain. In particular, the field is tasked with trying to understand how this large wet bundle of billions of neurons with trillions of synaptic connections all somehow use simple exchanges of ions to orchestrate complexities that emerge as behavior.

Perhaps one of the most fundamental attributes achieved by the brain is the ability to maintain information over time or more simply put: memory. From an evolutionary standpoint, it is arguable that memory is one of the more crucial skills the brain provides. Memory permits learning, which is fundamental to survival. The ability to remember past experiences permits organisms to make predictions about the future based on the past, which is a fundamental advantage for being successful and surviving. Ultimately memory is one of the key features of behavior, and thus one of the most prominent functions studied in systems neuroscience is understanding memory.

1.1.1 H.M. and episodic memory in the hippocampus

A major milestone in the study of memory was the work of Milner and Scoville in 1957, which revealed the hippocampus as an area important for memory as well as showed that there are different types of memory systems. During this time, it was common practice in patients with intractable seizures to remove the hippocampus from their brain because most often the hippocampus was the location where seizures originated. However, Henry Molaison was one of the first seizure patients to have a bilateral hippocampectomy. Unfortunately, after the procedure it was discovered that Henry was unable to form new memories. This tragic result was a major breakthrough in the study of memory because it localized a particular brain region, the hippocampus, as being required for the encoding of new memories. Furthermore, it helped refine our understanding of different memory systems (Fig 1.1). Henry Molaison's

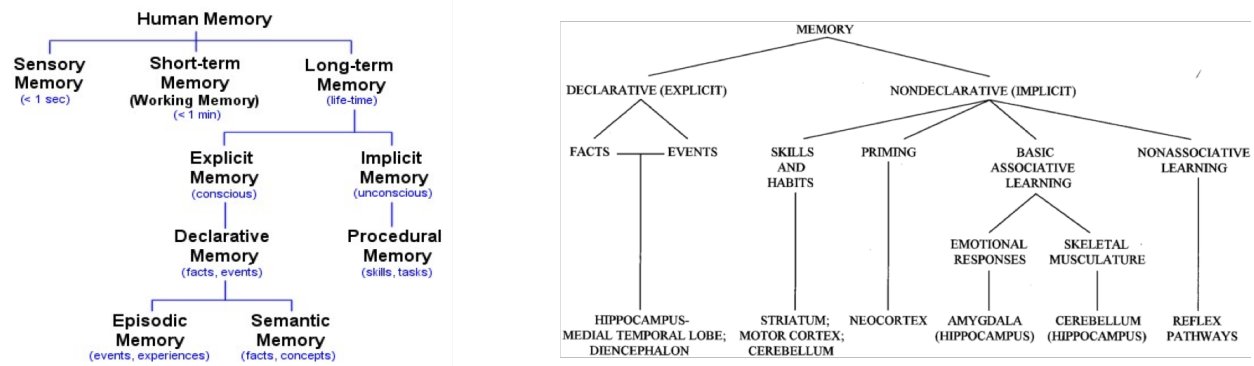

Figure 1.1: Types of memory
a) Different types of human memory
b) Speculative association of different types of memory with brain structures

memory impairment was specific to explicit memory, meaning he lacked the ability to form new declarative memories (particularly episodic memories; his ability to form new semantic memories remains controversial), while his ability to form new implicit memories such as unconscious motor skills / procedural tasks was intact (Scoville and Milner, 1957). This research along with more recent work in humans with hippocampal lesions showing that these patients were unable to imagine new experiences (Hassabis and Maguire, 2007), guided and shaped our understanding of what specific functional role the hippocampus has with a very specific memory system/type – that of episodic memories in particular.

1.1.2 Episodic memory

The term episodic memory was initially coined by Endel Tulving in 1972 and was used to make a distinction between remembering (episodic memory) and knowing (semantic memory) (Fig 1.1). There are plenty of discussions about the precise criteria that defines episodic memory as well as discussion about whether is exists in nonhuman animals (Clayton et al., 2007; Hassabis and Maguire, 2007; Templer and Hampton, 2013), but for the purposes of this thesis we will refer to a general statement in Tulving's original paper, "Episodic memory receives and stores information about temporally dated episodes or events, and temporal-spatial relations among these events " (Tulving et al., 1972).

1.1.3 Anatomy of hippocampus across species

Having localized the hippocampus as an area important for episodic memory, we can look in more detailed at its' structure and place in the brain. The hippocampus is a structure that is preserved across mammals . It is made up of three layers folded over in a seahorse shape. The classical trisynaptic pathway consists of the perforant path (entorhinal cortex (EC) to dentate gyrus(DG)), mossy fibers (DG to CA3) , and Schaeffer collaterals (CA1 to CA3). The performant pathway also projects from the EC to CA3 and CA1 (Fig 1.2) (Andersen et al., 2007).

More recent studies with improved tools are finding more pathways (Hitti and Siegelbaum, 2014; Kohara et al., 2014) and more subdivisions (Strange et al., 2014; Bienkowski et al., 2018). There is a strong gradient of the type of connectivity and information processed in the rodent hippocampus along the dorsal ventral axis (Strange et al., 2014). Finally, a comparative study found signatures of neural activity in hippocampus were conserved across species (Fig 1.3) (Mochizuki et al., 2016). All of these examples suggests a preserved functional role of hippocampus across mammals.

1.1.4 Animal model systems for neuroscience study

While one of the ultimate goals in systems neuroscience is to gain an understanding of the human brain, research performed in model systems can provide more detailed methods for studying the brain. The core assumption is that many of the fundamental principles of neural computation is preserved across species.

However, studying cognitive function is more difficult in animal models because unlike in human studies we cannot openly communicate with animals using language to gain insight into internal states and thoughts. Instead we must rely on well-designed behavioral tasks. Therefore, one of the first and most basic investigations into the hippocampus and memory was accomplished using lesion studies combined with behavioral tasks. The motivation is that with a well-designed behavioral task with a specific functional requirement (such as a

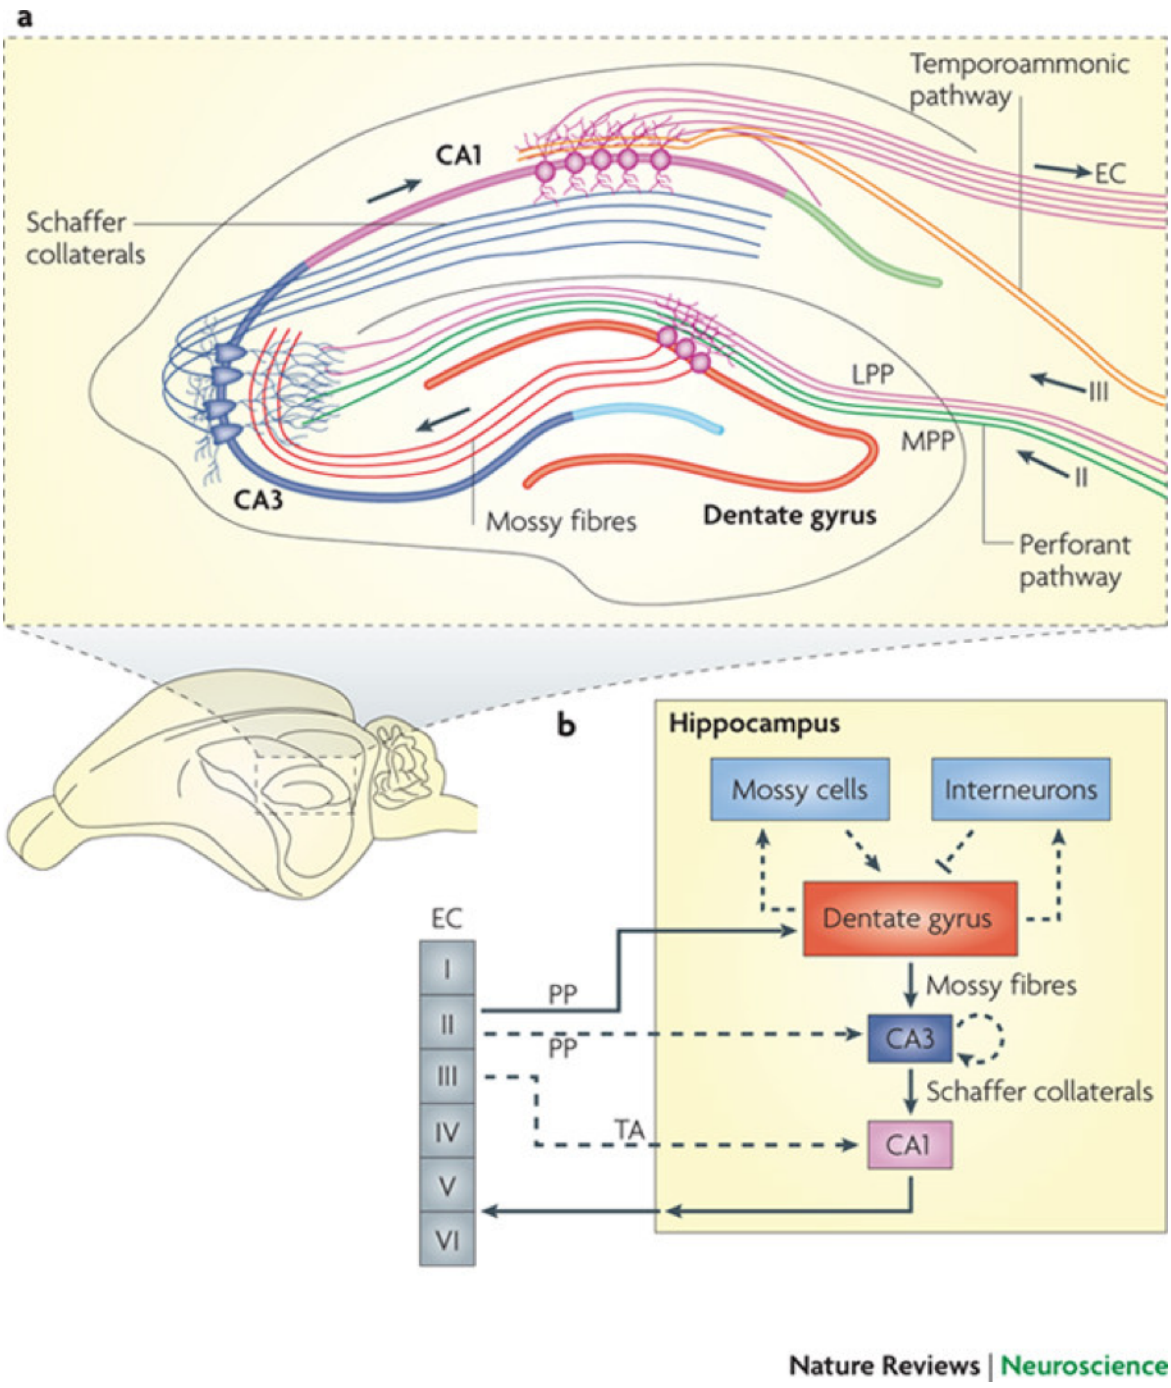

Figure 1.2: Anatomy of the rodent hippocampus

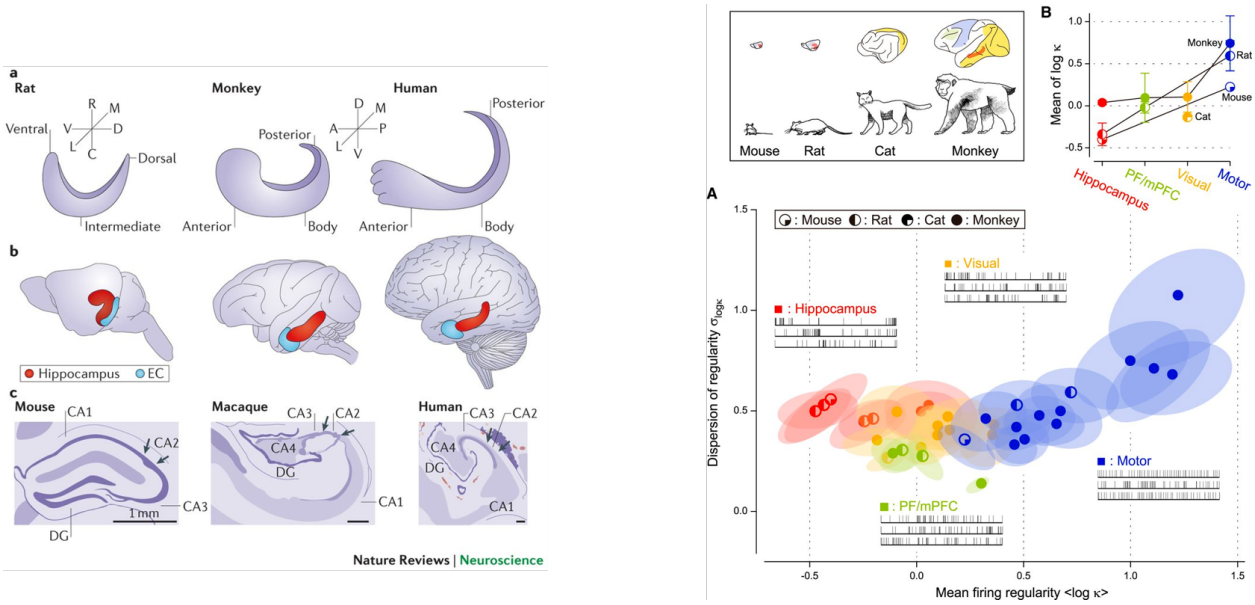

Figure 1.3: Hippocampus across species
Left from Strange et al 2014 shows the anatomical location and organization of hippocampus across Rat, Monkey, and Human brains.
Right from Mochizuki et al 2016 shows conservation of neural activity in the hippocampus across different species is preserved by using metrics that quantify the firing characteristics of neural activity in different areas within and across species.

requirement to remember a location or experience), experimenters can try to determine the role of an area in performing the function by lesioning it and observing how lesions affect performance in the task.

There have been several behavioral tasks used for studying memory in rodents (Dudchenko, 2004). A comprehensive review from Jarrard 1993 summarizes results from lesions of hippocampus in rodents:

The performance of rats with the hippocampus removed was compared with that of control animals in the acquisition and retention of spatial versus nonspatial information, forgetting of spatial and nonspatial information, contextual learning, recognition memory and concurrent discrimination learning, and complex representational learning (conditional discrimination and negative patterning learning). The general finding that rats without a hippocampus were impaired on those tasks that required the utilization of spatial and contextual information stands in contrast with the spared performance that was found in

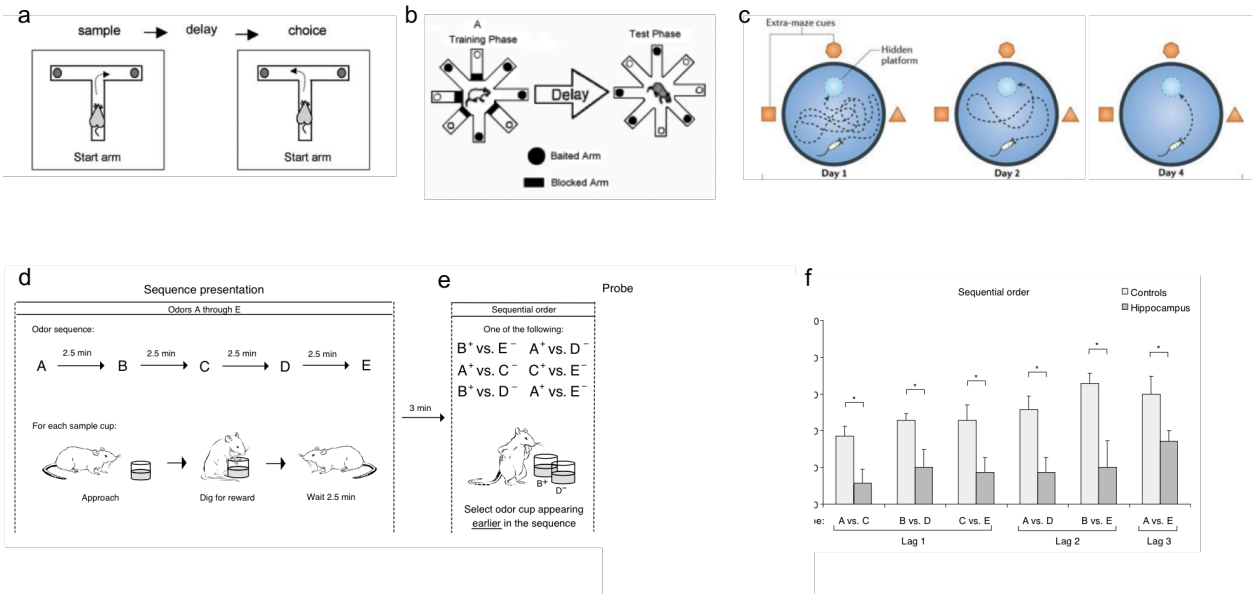

Figure 1.4: Examples of hippocampus dependent tasks
a) illustration of a delayed alternation task in the T-Maze
b) illustration of a radial arm maze spatial memory task
c) Illustration of the Morris watermaze task
d) Illustration of the ordered odor stimuli presented in the Fortin task
e) Illustration of the probe test for memory of sequential order
f) Fortin et al 2002 show that hippocampus is required for memory of sequential order in a non-spatial memory task

learning about and handling (even complex) nonspatial information. Rather than support for views that emphasize a role for the hippocampus in specific memory processes (working memory, declarative memory, temporary memory buffer, configural learning), the present results are more compatible with the idea that the hippocampus plays an especially important role in processing and remembering spatial and contextual information.

Most of the behavioral tasks involving tests for spatial memory are similar in concept such as the delayed alternation T maze task, the radial arm maze task, and the Morris water maze task which require remembering past spatial experience (Fig 1.4 a-c).

It is worth briefly covering an interesting non-spatial, hippocampus dependent memory task developed in Howard Eichenbaum's lab (1.4 d,f, (Fortin et al., 2002)). In the experiment rats were presented with a series of five random odors in cups of sand with reward buried inside. After odors were sampled, the rat was given a probe trial with two cups of sand with

odors from the previous sequence. The correct response was to dig into the cup with the odor that appeared earlier in the sequence. Fortin and colleagues found a clear impairment in memory of sequential order in rats with hippocampal lesions. This unique experiment provides an important example of how the hippocampus, while essential for spatial memory, is also important for memory of temporal context.

In summary, the conclusion presented by Jarrard in 1993 has remained true even after findings of function in non-spatial sequential memory tasks like Fortrin. The hippocampus in rodents is important for memory of spatial or temporal context.

1.2 Discoveries from recordings of neural activity in hippocampus

Having summarized the experiments aimed at attributing a function to the hippocampus, we now turn to observations made from electrophysiological recordings. Once again the major theme of neuroscience is to connect our understanding of behavior with the internal processes in the brain, therefore it is important to cover the knowledge gained from recordings of the internal activities in hippocampus. For the purpose of this introduction we will primarily focus on results reported from neural recordings in rodent hippocampus.

1.2.1 Regular Spiking activity and Large Irregular Activity states

The first and perhaps most prominent observation from recordings in hippocampus was the early description of two very distinct states of oscillations in local field potential (LFP) or EEG (electroencephalogram) activity. The two states are regular spiking activity (RSA) and large irregular activity (LIA) (Vanderwolf, 1969; Whishaw and Vanderwolf, 1973). Both states can be directly mapped onto two distinct behavioral states. RSA is seen when rats are moving, actively exploring, or even in a heightened attentional state without movement. In contrast the LIA state is observed in rats when they are stationary, grooming, and eating

or drinking. The two states of activity are mutually exclusive and each have strong neural correlates. The signature of RSA is the theta oscillation (6-10 Hz) which strongly modulates the neural activity during this state (Buzsáki, 2002). In contrast the LIA state is known for the absence of theta and irregularly disrupted by large amplitude fluctuations in LFP often accompanied or as a results of sharp wave ripple events (which will be covered in more detail below) (Buzsáki, 2015). In conclusion, RSA and LIA states are two prominent signatures of electrophysiology in rodent hippocampus. These two states may represent intrinsic default networks evidenced by the observation of mirrored versions of the two states observed at differing levels of isoflurane anesthesia (Lustig et al., 2016).

1.2.2 Place cells

Place cells where first recorded by O'Keefe and Dostrovsky in 1971 and have been the cornerstone of a significant amount of hippocampal research for the last few decades. First described as analogous to a higher order version of a Hubel and Wiesel receptive field, but one that could be described as being tuned to a multisensory higher order cognitive construct that appears to be based primarily on spatial position; place cells are neurons who have receptive fields for a particular location in space, meaning any time a rodent moves through the place field of a place cell that place cell will fire vigorously (Fig 1.5).

Place cells are considered to demonstrate a higher order cognitive function by displaying awareness of allocentric space – meaning despite the fact that a rodent may experience very different sensory cues depending on their orientation within a place field, the place cell can still be active – demonstrating an understanding of where the rodent is in relation to surrounding space in the environment. As methods, have allowed recording of more neurons simultaneously it has been shown that across the population of neurons in CA1 of dorsal hippocampus there is a place field representing all of space, i.e. that they tile any given environment (Wilson and McNaughton, 1993). Therefore, whenever a rat travels in space there is a sequential activation of the place cells whose receptive fields are along the

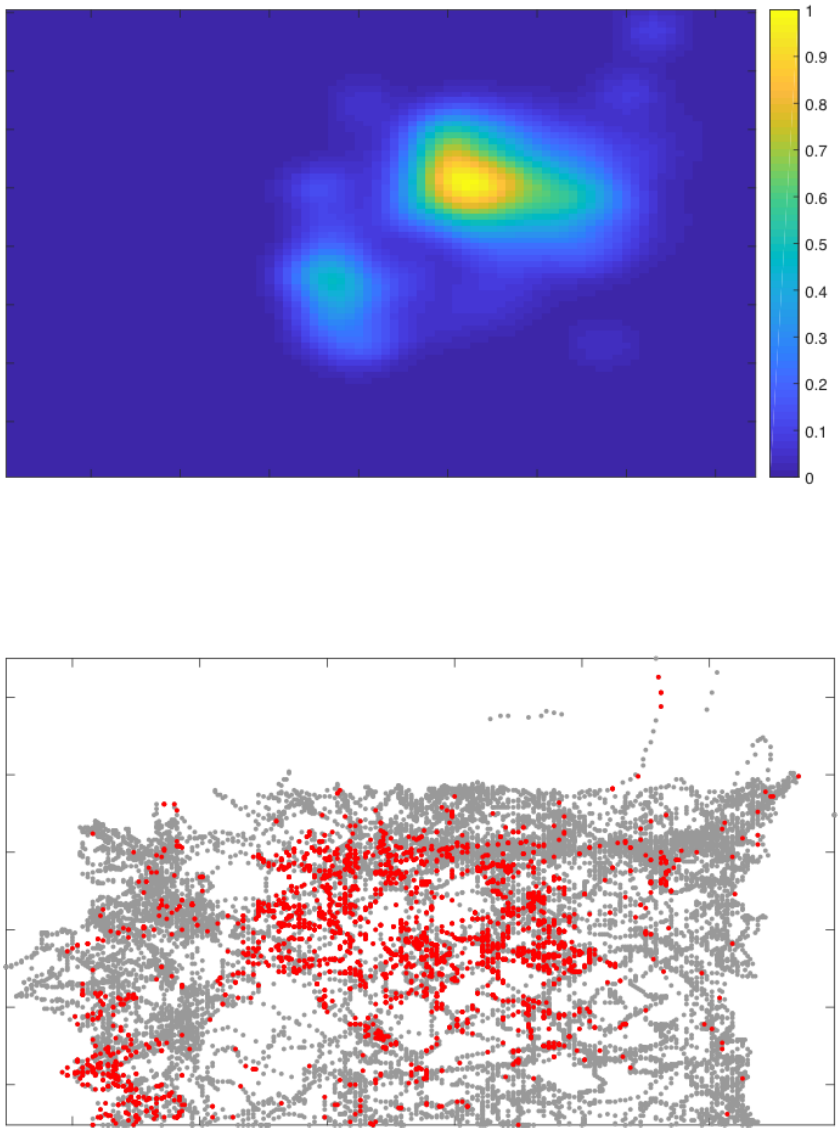

Figure 1.5: Place cells
An example of a place cell. On the bottom is tracking of the rat moving in an open field environment plotted in gray. In red are location where the neuron spikes. By binning the location of neural activity and dividing by binned occupancy a plot for the receptive field is shown above.

trajectory it moved through. Place fields themselves are thought to be fairly static for a given environment but importantly they are known to remap when there are changes in context of an environment (Leutgeb et al., 2005; Wills et al., 2005). There are two types of remapping described in place cells, rate remapping and global remapping (Leutgeb, 2005). Place fields are thought to undergo plasticity as they form in a novel environment and they become more stable (Cohen et al., 2017). There is evidence that the stability and formation of place cells is tied to sensory cues in the environment (Knierim, 2002; Wang et al., 2015). Recent research found that simply depolarizing a neuron in CA1 can uncover a place field (Lee et al., 2012) or using current injection to induce plateau potentials (Bittner et al., 2015). Finally, as even more novel methods are adapted, larger populations of neurons in hippocampus have been imaged over long timescales permitting new studies on the stability and long term dynamics of place cells (Hainmueller and Bartos, 2018; Ziv et al., 2013; Dombeck et al., 2010).

1.2.3 Episode cells and internally generated hippocampal activity

Place field firing in the hippocampus was flipped on its head with the pivotal discovery of Pastalkova et al (2008) where they found that the hippocampus was able to generate sequences of firing fields in the absence of changing external cues and motion through space. In this experiment rats were required to run in a wheel during the delay period of a spatial memory task. Pastalkova and colleagues found that looking at neural activity in hippocampus aligned to the start of wheel running showed the sequential activation of neurons in CA1 tiling the entire delay period, just as place cell do for movement through space (Fig. 1.6, Pastalkova et al 2008) However, in this situation the rat was running but clamped in space within the running wheel so there were no changing external cues. The neurons were named Episode fields because they appeared to code for episodes of time across the wheel run. Since this paper several reports have shown the hippocampus generating activity locked to behavior across delay periods ("time cells", (MacDonald et al., 2011)). These results shifted thinking about the possible role of activity in hippocampus from passive encoding of chang-

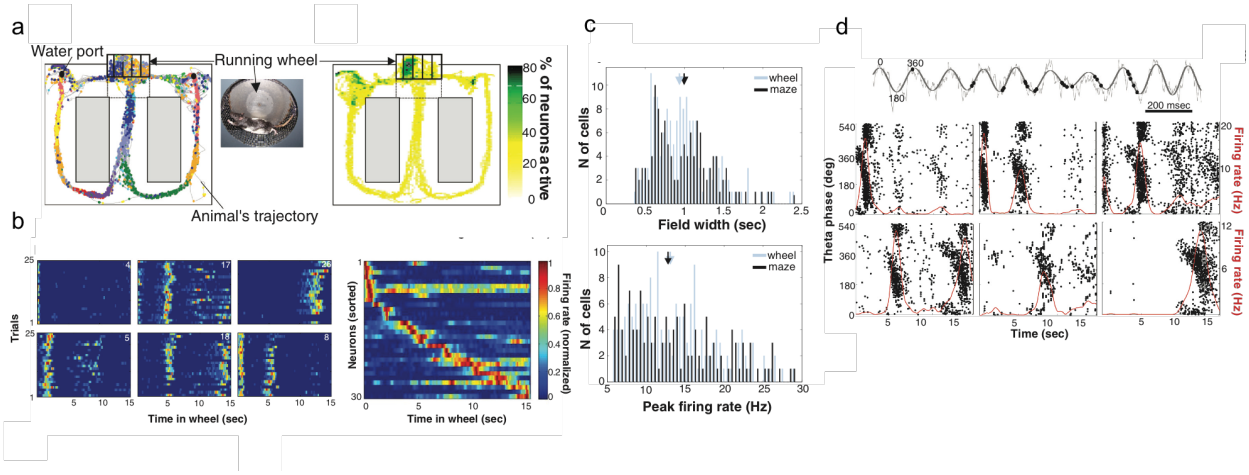

Figure 1.6: Episode Cells from Pastalkova et al 2008
a) The delayed alternation task with wheel running during the delay. Rats were required to run for 15 seconds during a delay period between alternating left and right arm choices for reward.
b) Examples of single episode cells recorded in CA1. Six cells are shown. For each cell firing rate over time during wheel running (x axis) is plotted for 25 trials (y axis). On the far left the mean activity in wheel running from 30 neurons are shown (y axis neurons, x axis time in wheel) spanning the entire delay period wheel run with the sequential activation of neurons.
c) Top : Comparison of the distribution of field width between place and episode cells recorded in the task. Bottom comparison of the distribution of peak firing rates for episode and place cells.
d) Demonstration of phase precession in a single episode cell at the top. Below: plots of spikes over time and relative phase of the local theta oscillation for six example episode cells show phase precession as seen in place cells.

ing sensory cues to one of generating sequential structure for sensory experience to be bound to for storage in memory (Buzsáki and Tingley, 2018; Pastalkova et al., 2008; Wang et al., 2015).

1.2.4 Sharp Wave Ripple events

The SWR event has been extensively studied (for a comprehensive overview see (Buzsáki, 2015)). Briefly, SWRs are mechanistically attributed to the synchronous discharge of a large population of CA3 neurons. The resulting signature sharp wave depolarization across CA1 is coupled with high frequency ripple oscillations (140-200Hz) embedded within the sharp wave seen in CA1. These are thought to be a consequence of the synchronous depolarization from the CA3 input (sharp wave) and the ripple component is thought to be a result of local

interneuron circuits (Stark et al., 2014).

SWRs have garnered strong interest in the field because it was found that during the relatively short burst of firing the neurons fire in a sequence that correlates with the sequences observed on the timescale of movement through place fields (Nádasdy et al., 1999; Lee and Wilson, 2002). Sharp wave ripple events are perhaps the largest most synchronous event in the brain and have been shown to have signatures across the entire brain network (Logothetis, 2015). They are thought to be important maybe for consolidation of memories because of their reach within the brain and their presence in sleep. In addition, selective perturbation of SWRs in sleep was shown to impaired learning (Girardeau et al., 2009; Ego-Stengel and Wilson, 2010). Theoretical work has suggested that SWRs might be important for solidifying connection strengths in an ensemble across the population by showing strong activation on the timescale optimal for short timescale synaptic plasticity mechanisms (Buzsáki, 2015). Interestingly the replay found during SWRs has been both of forward and reverse versions of place field sequences ((Diba and Buzsáki, 2007), Fig. 1.7). Also, it was shown that the temporal offset between the sequences correlates with the distance between place fields in space, so these are not just repeats of the ordered sequence, but actual replays of the timed sequence on a compressed timescale (Diba and Buzsáki, 2007),.

1.3 Overview of Thesis

The field of hippocampus research in systems neuroscience is considerably large. In this introduction, we briefly covered a subset of the relevant research for the purpose of providing sufficient background for the following chapters. Moving forward it will be important to link neural activity in hippocampus to function in behavior by recording hippocampus in rats performing memory tasks. The following thesis presents results from such an attempt. In particular, the research presented aimed at testing the prominent theory that the hippocampus encodes contextual information used for maintaining memories of past experiences across delay periods during delayed spatial memory tasks. By implementing a new three arm

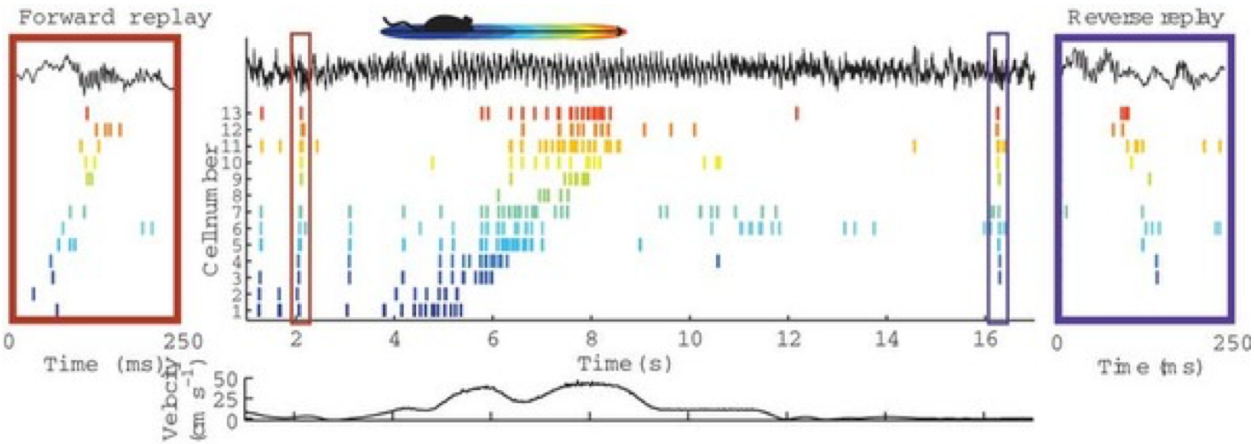

Figure 1.7: Forward and reverse replay during SWRs from Diba et al 2007
a) The figure shows a raster plot of 13 simultaneously recorded place cells while a rat runs down a linear track. During the short bursts of activity before and after running (expanded on the left and right sides of the figure) within Sharp Wave Ripples, the sequence of place cells it played in forward and reverse on a compressed timescale.

delayed sequence task (TADS) we were able to test if activity in the hippocampus across delay periods contains sufficient contextual information in a paradigm that requires contextual differentiation following the same immediate past spatial history. The first chapter of the thesis presents the results from this question. After finding a surprising lack of sufficient information to differentiate contexts following the same immediate past spatial history, chapter 2 shows results from further exploration of activity in hippocampus during behavior by decoding the content of sharp wave ripple events during reward consumption. These results provide some evidence for contextual differentiation of the center arm segment of the task in the form of spatially and temporally remote planning for the future trials. Finally chapter three will demonstrate that decoding of the sharp wave ripples included replay of contextually specific wheel run trajectories. The discovery of replay of contextually specific experiences provides an important demonstration of episodic memory in the neural activity of rodent hippocampus.

CHAPTER 2

ACTIVITY IN HIPPOCAMPUS DURING PERFORMANCE OF THE THREE ARM DELAYED SEQUENCE TASK

2.1 Introduction

The previous chapter provided an introduction to the study of memory in hippocampus. Moving forward the next step for the field is to gain an understanding of how neural activity in the rat hippocampus supports behavior in hippocampus dependent tasks. This chapter will present results that give important new insight into one of the major theoretical frameworks for how the hippocampus serves the function of maintaining and encoding context. We will first go over the key discoveries from recordings in hippocampus during behavior that have provided strong support for a neural implementation of maintenance and encoding of context. Then we will introduce a novel task that aims to confirm the hypothesis that contextual coding required for a spatial memory task is maintained in hippocampus. Finally, the results will be followed by a discussion about how the data changes a prominent theory of contextual coding in hippocampus.

In 2000 two papers reported observations of what they considered contextual coding in the firing of neurons in CA1 of the rodent hippocampus (Wood et al; Frank et al). In both tasks the rat needed to visit two different arms in an alternation, which meant needing to remember past experience to inform future choice. Importantly both tasks had a common segment between the alternation permitting the comparison of activity in the same location but under different behavioral contexts (Fig 2.1 a,c). When comparing activity in common arms between alternations they found examples of the firing rate of neurons in CA1 varying drastically depending on the surrounding context (Fig 2.1 b,d). For example, if the rat had just visited the left arm and was going to turn down the right arm next, then a neuron might have a place field with a 10Hz peak firing rate in the central arm during the behavior. However, on the next trial, under the opposite context where the rat just came from the right

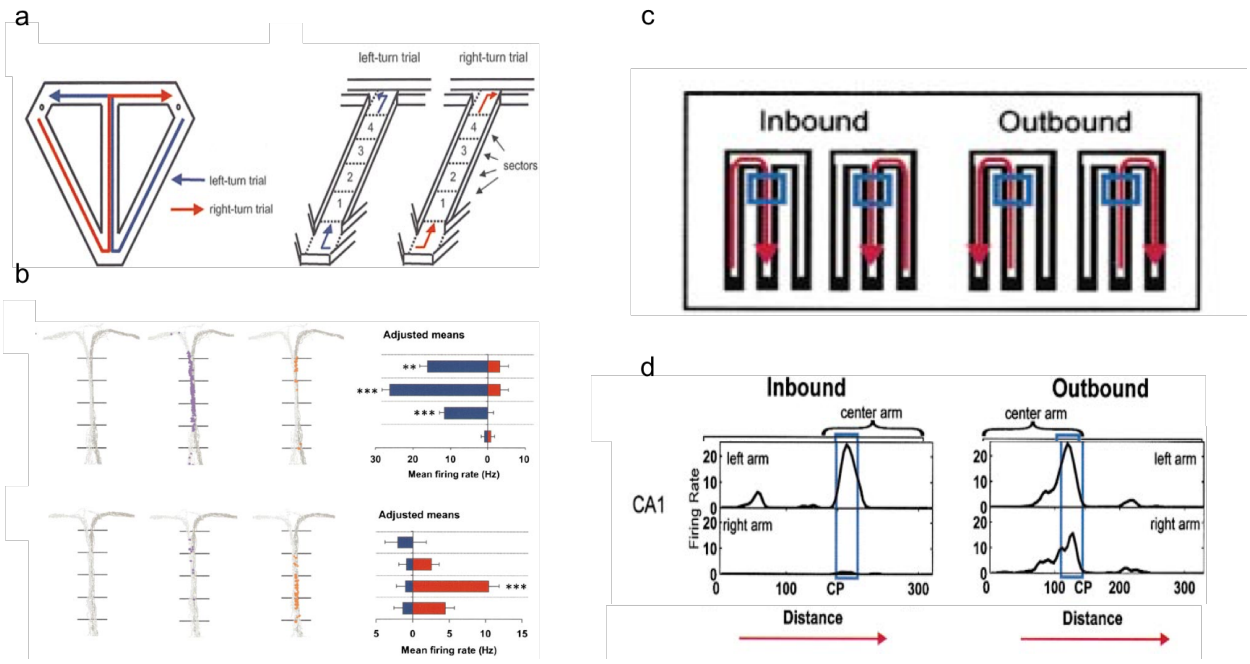

Figure 2.1: Evidence for contextual coding in hippocampus:
a) A diagram of the continuous t-maze used for the alternation task and segmentation of the center arm for analysis in Wood et al 2000.
b) Examples of neurons in CA1 with context dependent firing in the center arm. Trajectories from all trials on the far left column. Middle two column plots show location of spikes for examples neurons on left (purple) or right (orange) trials. Far right column shows the mean firing rate for segments of the center arm.
c) Diagram of the different trial segments as well as inbound and outbound comparisons from Frank et al 2000.
d) Examples of context dependent firing in CA1 from the inbound and outbound comparisons.

arm and plans on going to the left arm, the same neuron might show no activity at all in the center arm (Figure 2.1 b,d). Thus, based on the activity of neurons in hippocampus during traversal of a common location an observer could say something about the past experience or future behavior surrounding this moment in time. Therefore, these neurons in hippocampus appeared to encode contextual information.

The results of Wood and Frank et al were extended by Ainge et al in 2007. This paper confirmed the existence of contextual coding in a continuous alternation task (such as the modified continuous T-maze design used in Wood et al); however, removing the hippocampus had no effect on performance. Ainge and colleagues reported that the addition of a delay

as short as two seconds was required to make the alternation task hippocampus dependent (Fig 2.2 a, b). They also reported under a delayed version of the alternation task that the neurons that showed contextual coding (often referred to as "splitter cells" in the literature) were no longer as prominent in the center arm of the maze but instead shifted to the delay area (Ainge et al., 2007). This result fits with the theory that hippocampus is particularly important for the internal storage of contextual memory over time as many previous studies had reported a delay dependent effect in memory tasks that required hippocampus (Jarrard, 1993; Hampson et al., 1999; Ainge et al., 2007).

These papers supported the idea of contextual coding in hippocampus and while the evidence is attractive, it still faced two major limitations. The first was that the tasks where context dependent activity had been demonstrated were found after Ainge et to not be hippocampus dependent tasks. The second limitation was the fact that the idea of contextual coding rested on the important assumption that the difference in firing between the two conditions was not a result of any instantaneous behavioral difference. In several reports of these short mazes it is very difficult to discount behavioral differences. Note the examples in Frank 2000 are all near the choice point, and see the supplementary material from Fujisawa (Fujisawa et al., 2008) on head direction differences. If there are subtle behavioral differences that could explain the differential firing, then the report of context dependent activity would be easily explained. Therefore, it was especially relevant when in 2008 Pastalkova et al, besides reporting their discovery of internally generated activity during the wheel running of a delayed spatial memory task, also found context dependent firing during the delay. The details of Pastalkova et al 2008 were described in the introduction, but briefly in Figure 2.2 c one can see the design of the task behavior. Figure 2.2 d shows individual episode cells activity during the wheel running that showed differential activity depending on the context of the trial type. Figure 2.2 e shows that the neural population from the two trial types (ordered by the left trials latency peak to give an entire population sequence) displayed context dependence. The results in Pastalkova et al addressed the previous two problems

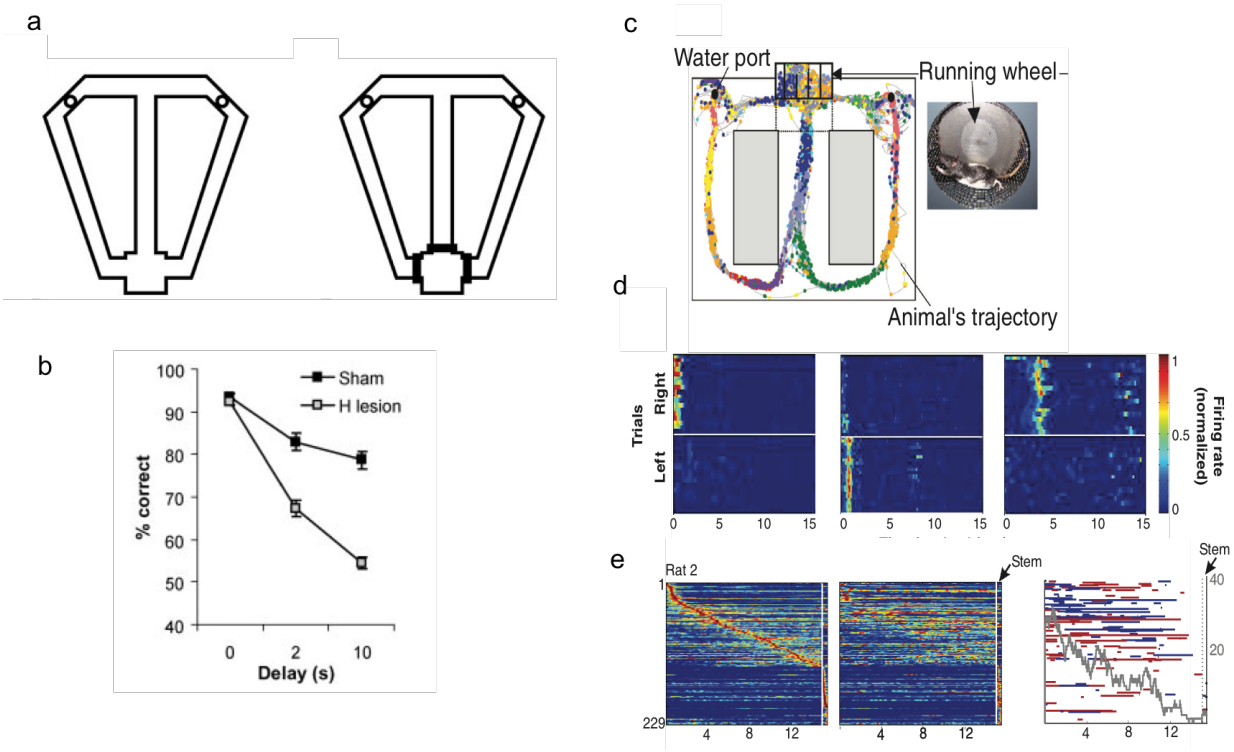

Figure 2.2: More evidence for contextual coding in hippocampus
a) Diagram of the continuous and delayed alternation tasks used in Ainge et al 2007.
b) Lesions of hippocampus in rats shows no impairment of performance on the continuous alternation task, while both the delayed alternation task with 2sec and 10 second delays show impairment.
c) Diagram of the delayed alternation task with wheel running from Pastalkova et al 2008.
d) Examples of context dependent episode cells.
e) Left, demonstration of episode cells across the population tiling the entire wheel run during the delay using activity averaged across left trials. Middle, population activity during right trial wheel running sorted by the order determined by using left trial activity. Right colored bars indicate neurons that showed significant differences in left or right trials. Grey line plots total number of significantly differentiating neurons in the population over time.

because they found context dependent coding in a hippocampus dependent task and during a controlled behavior (i.e. running in the wheel) that gave more confidence that the differential firing was not a result of differences in motor behavior or external cues. Pastalkova et al found that the context dependent activity was sufficient for predicting errors in the task, giving strengthened evidence for the correlation of this activity and behavior.

As mentioned in the first chapter, behavioral studies have pointed towards the hippocampus as being especially important for tasks that require maintaining a memory of context. The results introduced in this chapter found evidence for a possible neural implementation of contextual coding. Linking neural activity to function in behavior is one of the major goals in systems neuroscience, so confirmation of this possible example discovered in the hippocampus is worthy of rigorous study. The rest of this chapter will present the Three Arm Delayed Sequence task (TADS) as an attempt to confirm the role of internally generated contextual activity in rat hippocampus as being the substrate of contextual memory utilized in a delayed spatial memory task.

2.2 Methods

2.2.1 Three Arm Delayed Sequence Task

The Three Arm Delayed Sequence task (TADS) requires the rats to visit maze arms for reward in the following order: left, center, right, and center. Between each arm choice there is a forced delay where the rat must run in a wheel over a set speed threshold continuously for eight seconds. Once the eight second wheel run is completed all three doors will simultaneously open and the rat can make an arm choice for reward. Besides the required wheel run in the delay and the required sequence for visiting the arms, there are no other constraints on the rats' behavior. Figure 2.3 shows basic logic of the task: visiting the arms in the order left, center, right, and center for reward, with wheel runs during a delay period between each arm choice.

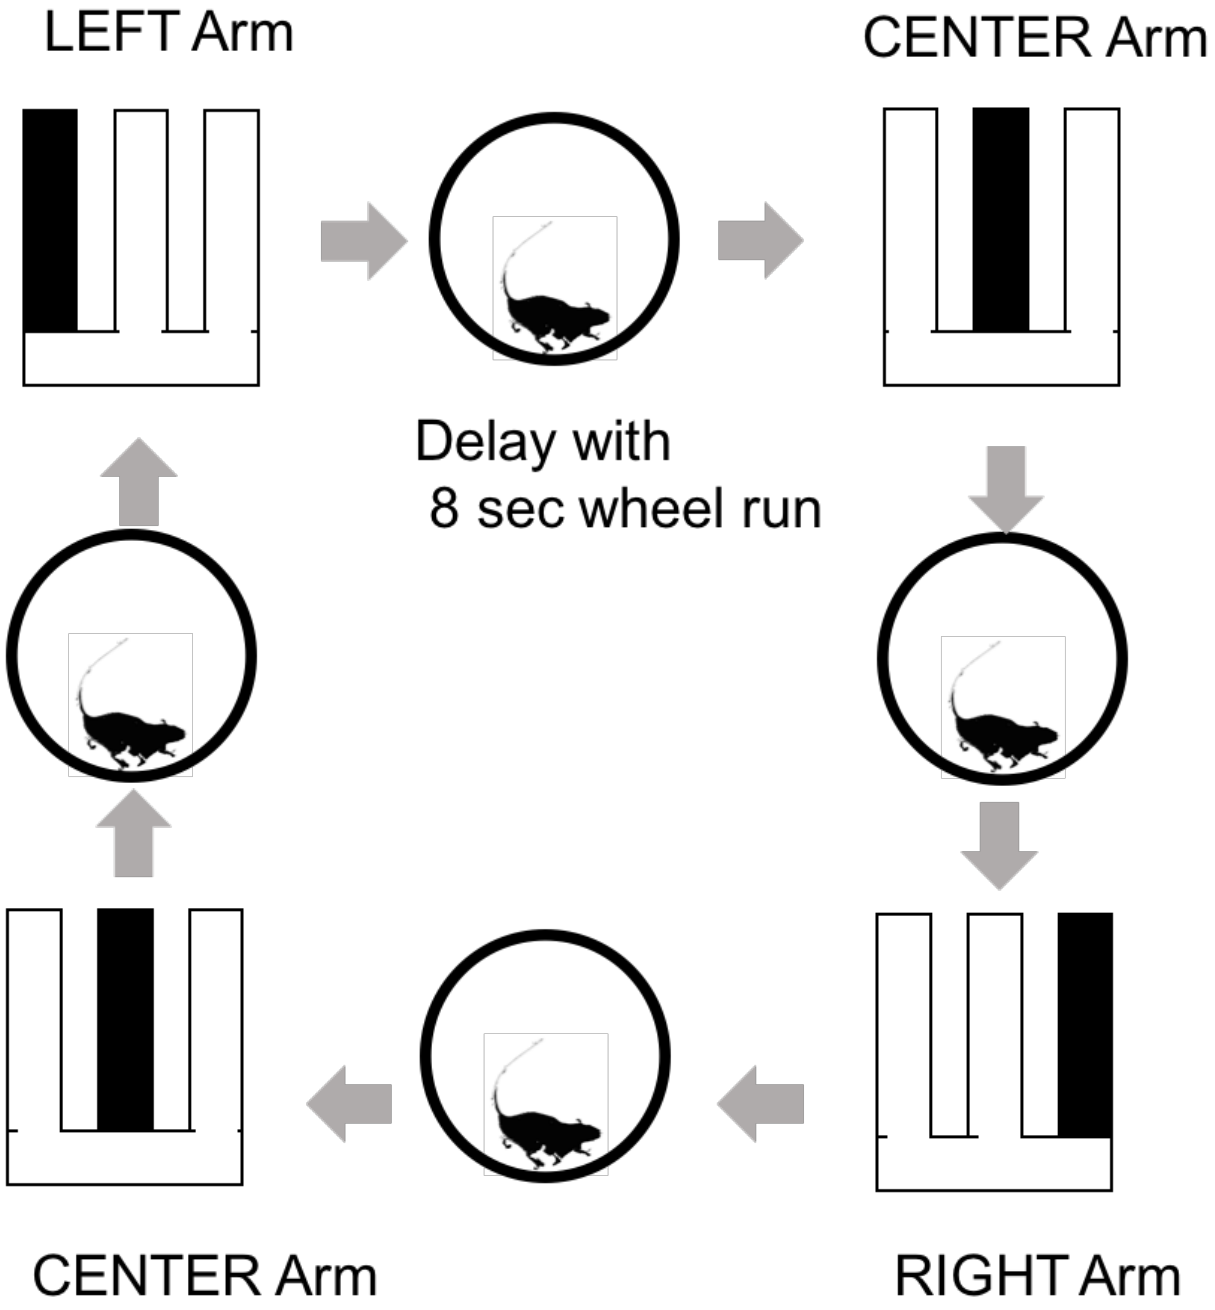

Figure 2.3: Three Arm Delayed Sequence Task logic
A cartoon diagram depicting the sequence of behavior in the TADS task. Rats must visit the left, center, right, and center arms in order with eight second wheel runs in-between.

The TADS task has some unique features compared to the typical delayed alternation tasks. First and most obvious is the addition of an extra arm, shifting the baseline performance at chance from 50% to 33%. Second is the amount of information that must be retained across the delay period. In the typical two choice delayed alternation task, maintenance of the memory of the past arm visited is sufficient to inform the next arm choice and obtain reward. Therefore, the task can be solved using a binary code that distinguishes between two states across the delay. In the TADS task, memory of where the rat last visited alone is insufficient to perform the task. Instead the task requires that the rat maintains four different contextual states. This is illustrated in Figure 2.3 where there are four distinct contexts around each of the delay wheel runs. If the rat only has two states to differentiate across the delay periods, then after visiting an outer arm (left or right) there is no way to know when the doors open if the next choice should be the center arm or the opposite outer arm. If the rat maintains three contextual states across the delay, for instance encoding the arm last visited, then the rat will be able to avoid repeating the last arm visited, however the rat will have no information indicating which arm to choose after a center arm choice. An alternative strategy that might seem sufficient would be to simply remember to go the center arm after visiting an outer arm. However, if the only information maintained after a delay was that you visited an outer arm and need to go to the center this would again leave you with no information for the next delay period after the center arm. Thus, the only strategy sufficient for completing the TADS task is to have four contextual states across the delay period. The rat must be able to hold a memory for four different contexts to distinguish which choice to make next. Thus, based on the current theory of contextual coding in hippocampus in the literature (Wood et al., 2000; Frank et al., 2000; Ainge et al., 2007; Pastalkova et al., 2008) the TADS is expected to reveal four distinct contextual states across the four distinct delay periods. A particularly novel and key comparison this task creates is the requirement for two different contextual versions of delay periods following a center arm visit, one in the context of having visited the left arm and heading to the right arm and the

exact opposite for the second version. We have covered the logic and functional requirement of the task, some of its unique features and the hypothesis for the expected results. Now we will describe the details of how the TADS task was implemented.

2.2.2 Maze implementation of TADS

The maze itself was designed by Eva Pastalkova and Tanya Tabachnik in the Instrument Design & Fabrication facility at Janelia Research Campus. Fig 2.4 shows a properly scaled diagram of the maze and its components. Briefly, there are three arms in the maze centered around a central delay area. The delay area contains a touchpad that is sensitive to the weight of the animal in order to detect the rats' presence in the delay. Also located in the delay area is a running wheel (10-cm depth, 29.5-cm diameter, Lafayette Instrument, 8086W) attached to an optical motion detector and break mechanism (HB6M-500-500-I-S-D, US Digital). The entrance to each arm from the delay area could be either opened or closed using automatic doors (hydraulic mechanism). Each maze arm is 15.25 cm wide at the base. Each maze arm was equipped with an inner and outer beam break (Omron E3F2-R4C4-P1) as well as a water port at the end equipped with an optical lick detector. Each water reward port consisted of a custom 3d printed housing for the lick detector and attached to Harvard Apparatus PHD-2000 for reward delivery. The entire maze system was controlled using an Arduino Mega 2560 based board. Behavioral events such as timestamped records of beam breaks, touchpad, wheel speed, opening and closing of doors, reward delivery, reward port licks, and tally of correct and error trials was saved into a .txt file.

Each trial contained the following structure (Figure 2.5): The rat begins in the delay area and must complete an eight second wheel run in order to open the doors. Then the rat is able to move about any and all arms; however, once the rat crosses one of the arms outer beam breaks the doors providing access to the two other arms he is not currently in are closed, locking in the rat's choice. If the rat chose the correct arm, at the same time the alternate arm doors close, the water reward pump is activated, starting the slow delivery of

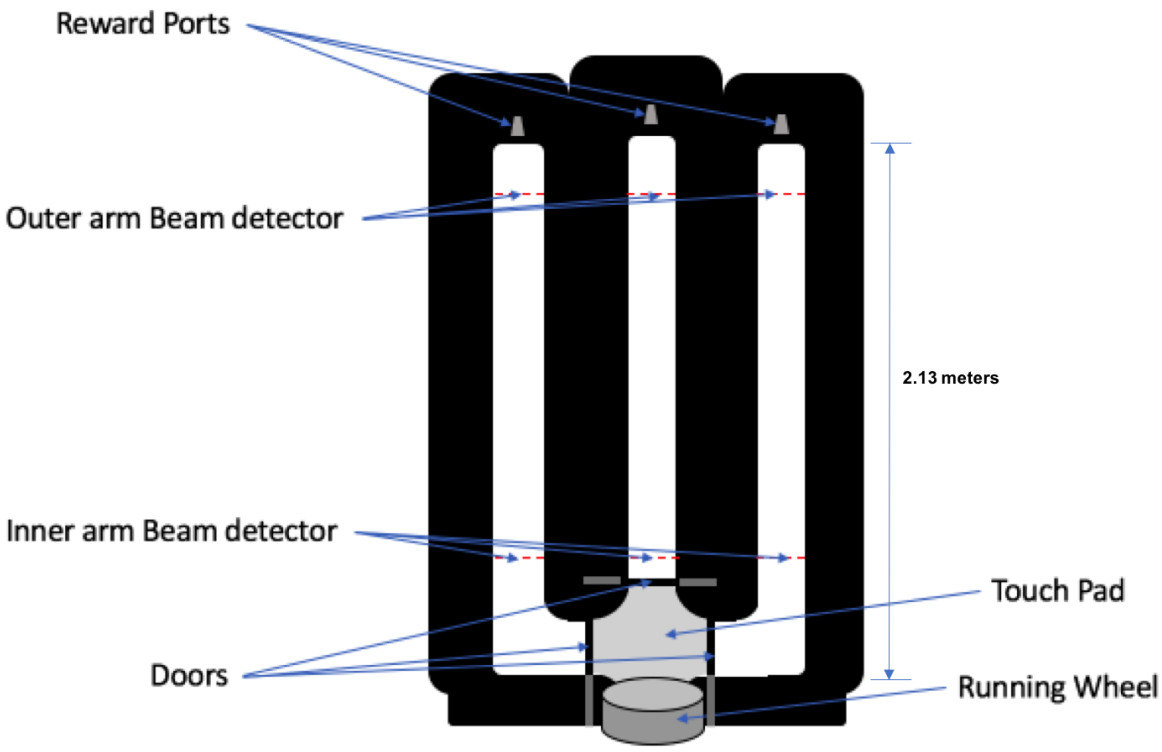

Figure 2.4: TADS Maze
A diagram of the maze used for TADS task.

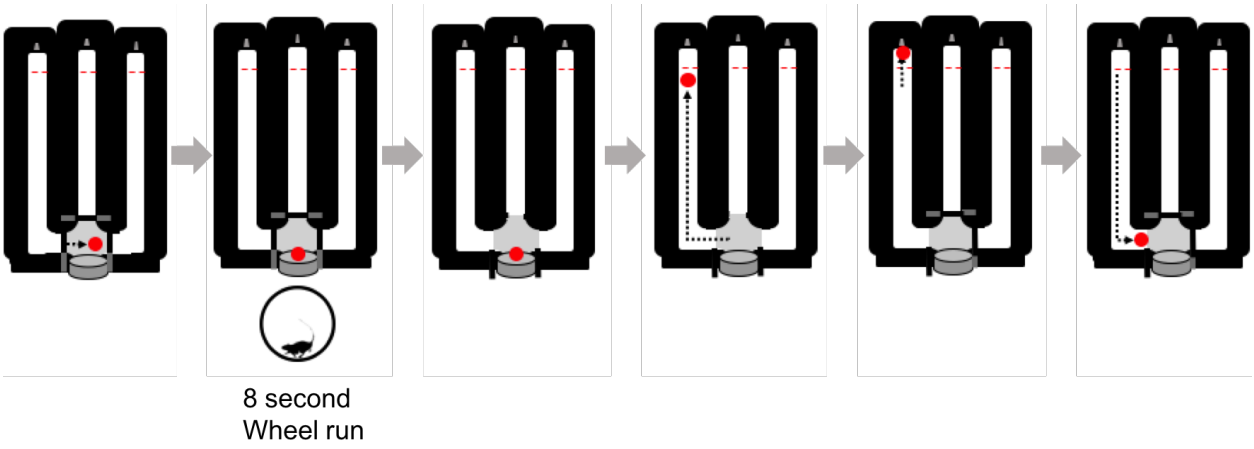

Figure 2.5: Single Trial of TADS task
Cartoon example of a single trial vising the left arm during performance of the TADS. Red dot represents the location of the rat.

100ul of water into the reward port. After consumption of water reward, or in the case of an error (in which the animal checks a port to find no water reward), the rat again has no time constraints, but once the rat crosses into the delay area and his position is detected by depression of the touchpad, the final door closes behind him, ending the trial and marking the start of the next delay period (Fig 2.5).

2.2.3 Recording neural activity

In order to link neural activity to function in behavior we performed chronic extracellular recording in freely behaving rats performing the spatial memory task. Chronic extracellular recording of neural activity was performed using silicon probes from Neuronexus (Buzsaki 64 and BuzsakiSP 64 probes, see figure 2.6 for probe dimensions and site arrangement).

Implant details – 3d printed ring, external faraday cage, nanodrives from Ronal Tools (http://www.ronal.com/). Ring designed by Eva Pastalkova and Yingxue Wang. Implant procedure by Eva Pastalkova. Faraday cage designed by Brian and Andy Lustig.

Neural signals and tracking of position was carried out using the Amplipex recording system (http://www.amplipex.com/). Importantly the Amplipex system is integrated with Intan multiplexing headstages (http://www.intantech.com/index.html) which permit record-

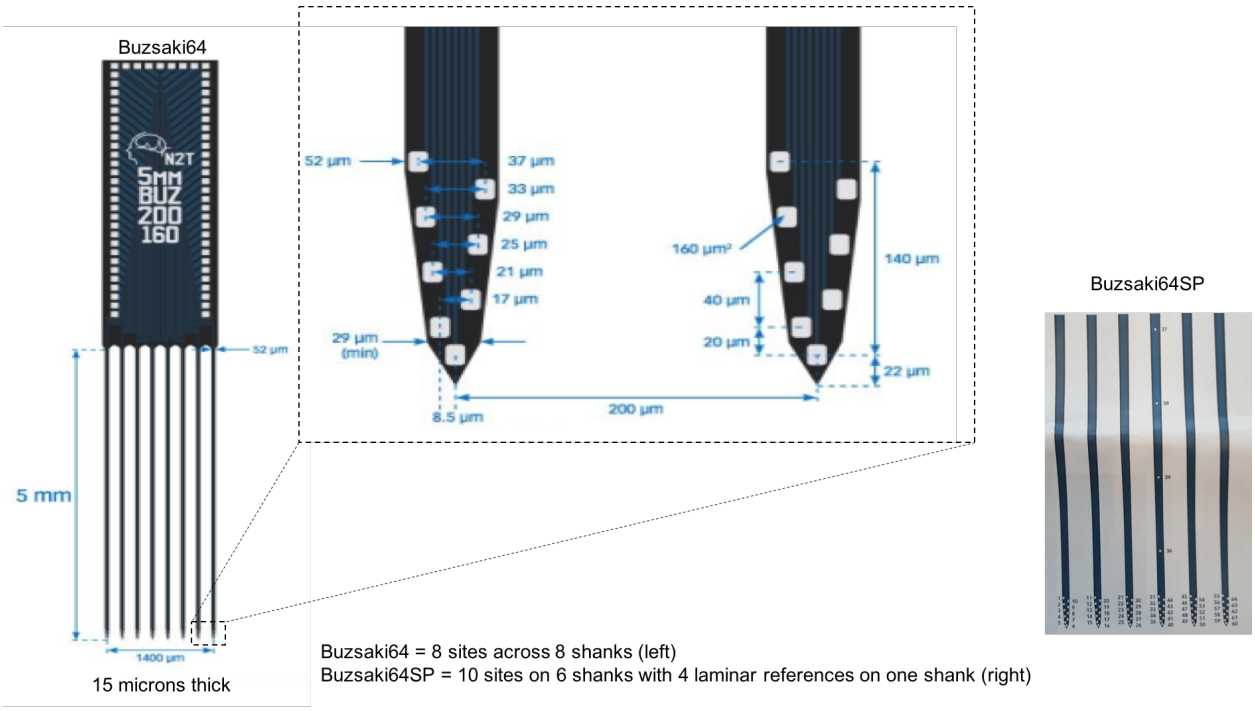

Figure 2.6: Diagram of silicon probe design and electrode layout for Buzsaki64 and Buzsaki64SP
Diagram from Neuronexus Product Catalogue

ing from large numbers of channels over a thin light cable to minimize interference with behavior in the freely moving rats. Commutators for recording were either custom from Dragonfly (www.dragonflyinc.com) or Doric Lenses (ERJ 12 HDMI-B2). Neural activity was recorded at 20 kHz.

Neural recordings were stored as .dat files. Behavioral location within the maze was recorded by tracking of LEDs on the rat's implant at 30Hz using a webcam integrated with the Amplipex system. A sync pulse from the maze behavioral control system was recorded on the Amplipex system in order to guarantee accurate syncing of behavioral information and neural activity recorded on Amplipex.

The NDManager and Klusters suite was used for analysis as well as MATLAB. Spike detection was performed using high pass filtered signals and threshold crossings above the mean + 1.5 SD. Then spike waveforms were stored and sorted using the KlustaKwik suite. After automatic sorting, clustering was further refined manually using Klusters.(Hazan et al.,

2006).

2.2.4 Data Processing Place field and Episode Field maps

Place field maps were computed using the following methods: All place field maps were computed using periods when the rat was moving over 5cm/s. Two-dimensional place fields were calculated for each neuron by taking the location of spikes in 2cm x 2cm bins (6x6 pixels) and also binning the time spent at each location. The binned spike count was divided by binned occupancy in order to obtain the binned firing rate. Then the firing rate was smoothed with a Gaussian kernel with standard deviation of 2.75 cm. Place cells in the maze showed strong directional selectivity, therefore for each trial the outbound and inbound segments were labeled and treated separately (outbound designated as the trajectory preceding reward collection, consisting of the rat moving away from the delay area and out towards the end of the maze arm). For instance, inbound place fields were computed using activity during the return period from the reward port back to the delay area. For calculating linearized place fields the X-Y tracking coordinates were projected onto the closest point on a manually drawn line for each arm. Then the linear location was determined as the index along the line. Linearized place fields were calculated using the linearized coordinates and they were calculated separately for the inbound and outbound segments of each trial. Thus, each neuron had a 2d place field map as well as linearized place field maps for the inbound and outbound trajectories of each of the three arms, for a total of six linearized maps (see figure 2.7 for an example). Depending on the analysis fields could be computed from the entire session or single trials.

Episode field maps were computed with the following methods: Episode field maps were computed by analyzing the spike times relative to the beginning of each wheel run. Activity was binned at 200ms for maps used for the classifier. Maps were made for single trials and averages across trials. As episode cells showed context specificity, maps were made for each contextual trial type separately (e.g. for the four different contexts, as described above).

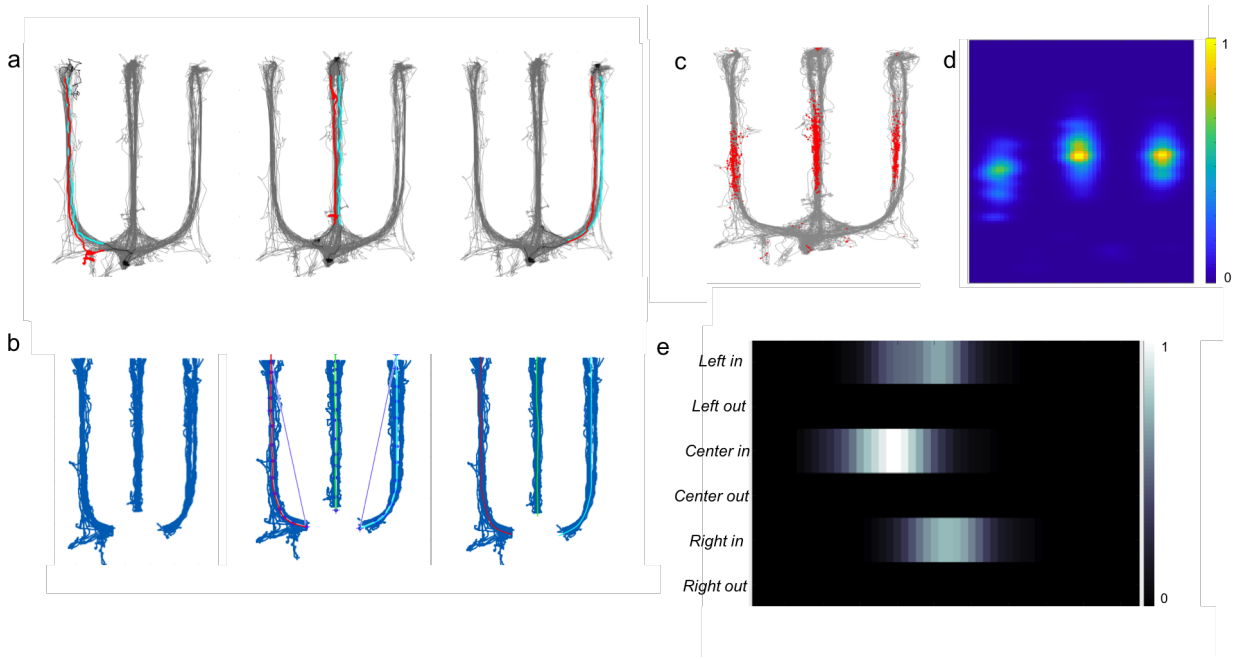

Figure 2.7: Place field analysis for Linearized and 2D maps
a) Trajectories from three examples trials are shown (left arm, center arm, right arm choices) with the inbound (red) and outbound (blue) segments indicated.
b) Lines used for projections to linearize coordinates for calculation of linearized inbound and outbound maps for each arm. c) Trajectories from a single recording session in gray with the location of all spikes from a neuron across the session plotted in red.
d) Place field map calculated for the neuron shown in c.
e) Six linearized place field maps for the neuron shown in c.

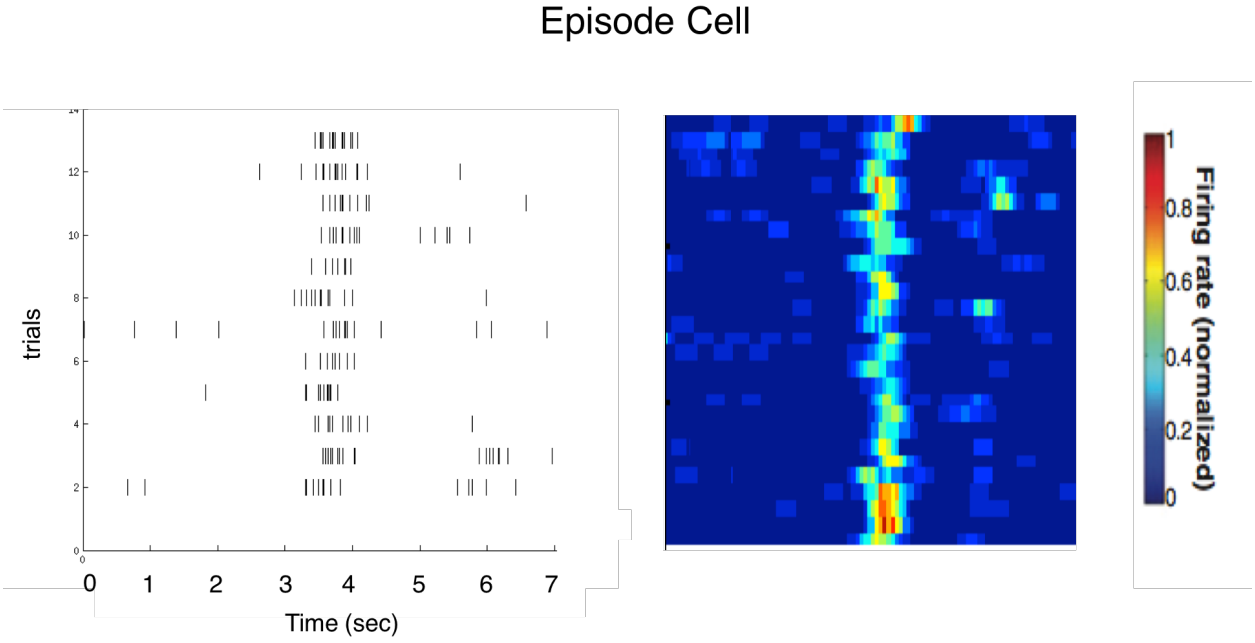

Figure 2.8: Episode cell calculation
An illustration of the activity from a episode cell with respect to the start of each trials wheel run on the left and the binned and smoothed activity on the right. (Right panel borrowed from Pastalkova et al 2008).

Figure 2.8 shows an example episode cell.

2.2.5 Linear Classifier

A matrix was constructed of the binned activity of all recorded neurons for each individual wheel run (N neurons for rows and 40 columns of activity binned at 200ms for the eight second wheel run). Then the population activity for each wheel run was transformed into a single vector for each trial. A binary linear classifier (fitclinear in MATLAB) was trained on the vectors for two trial types being compared (i.e. center to left and center to right trials) and the ability to successfully classify a trial was evaluated using leave one out cross validation. All results where replicated with other parameters, such as smaller bin sizes (20ms) and 10 fold cross validation. Classification was also tested with another software package (LABLINEAR) and showed the same results. For all trial comparisons, if there were unequal numbers of the two trial types being compared then all of the smaller group

of trial types was used and a matching number of trials was randomly chosen from the pool of the larger trial types. For each iteration of the classifier equal numbers of trial vectors were run in the classifier and the classification was then repeated using the same pools of trials but with the trial labels shuffled. This procedure was repeated 50 times. Evaluation of the performance of the linear classifier was achieved by comparing the distribution of cross-validated performance from all iterations to results using the exact same data and methods, but with the labels for each trial shuffled randomly. Then a two-sided t-test was performed on the two distributions of linear classifier performance to see if it would be appropriate to reject the hypothesis that the linear classifier performed no different than chance at distinguishing between the two trial types using information in the neural activity. Classifier performance greater than that from the shuffled iterations was taken as support that there was sufficient information in the neural activity to discriminate between the two contextual trial types.

2.3 Results

2.3.1 Behavior

Training

Because previous work has shown that there could be a relationship between the development of contextual coding and training procedure (Bower, 2005) it is important to have a standardized training procedure for each animal. The results presented are from three Long-Evans rats (300-400g). Rats were housed in cages (Techniplast) with custom made running wheels. Briefly all three rats were trained using the following method after having silicon probe implants inserted into the brain. After being given sufficient time to recover from the implant surgery the rat was started on acclimation to the maze and then training on the task. The two days of maze acclimation went as follows: on the first day, the rats were introduced to the maze by being placed inside the delay area with all doors open and allowed to freely explore and run in the wheel. The second day the rat was placed in the maze but the doors

were opened and closed with close supervision by the experimenter to acclimate the rat to the noise of the doors. After the two days of acclimation to the maze each rat was started on water regulation such that reward from the maze was the rats' only source of water. The next day after water regulation was started the rat was placed in the delay area with all arm doors closed. The door to the left arm was opened following a 0.5 second wheel run above the speed threshold. Once the door opened the rat was allowed to explore the left arm and the experimenter made sure the rat was aware that the reward port was dispensing water. Then the rat was allowed to return to the delay area where the door closed behind him. The rats quickly learned in 1-2 trials that turning the wheel caused a door to open. During this training period, only a single door opened each trial and the doors opened in what would be the correct order for the final task: left, center, right, center. The rat was given time in the maze everyday performing this "forced" version of the task that simply required running in the wheel to get a door to an arm to open, running down the arm to get reward, and returning to the delay area to start the next trial. Over time the required amount of time for continuous running in the wheel was gradually increased each day by 0.5 or 1 second. Typically, at the start of each day the rat ran the bare minimum amount in the wheel learned from the previous day (for example 2 seconds) and was surprised and frustrated when this caused nothing to happen. The rat would then try again until eventually running the extra amount of time to reach the time length needed for the day (say 2.5 seconds) causing the door to open. The rat would then collect water reward at the end of the arm and go on to the next trial. Typically, the next delay would also involve multiple wheel run attempts at first just short of the new time requirement; however, the rats quickly learned over a couple of trials the new required wheel run duration and typically ran the required time for the majority of the day's new trials. Once a rat was performing forced trials with a six second long delay he was moved on to the real task. This was simply done by starting the day with the maze set to the real task parameters, so after completing a wheel run all three door arms opened. Incorrect arm choices had no water reward delivered, thus over many days

the rats slowly learned the required correct behavioral procedure sequence of left, center, right, center by trial and error. Each rat had one behavioral session per day. They typically performed 105 trials over a total of 75 minutes and received 8.5 ml of water.

For each day's recording procedure, the rat was taken out of the cage and fed peanut butter treats in the experimenters' lap while the implant cover was removed and the electrophysiology recording cable and tracking LEDs were attached to the implant. Then the rat was placed into the delay area with one more peanut butter treat to eat while the experimenter returned to the recording system and started the recording of all behavioral parameters and electrophysiology data. Sessions were typically ended when the rats' behavior seemed to indicate lack of motivation to continue, with the exception of attempts to make sure the rat did enough trials in a single day to get sufficient water.

Quantifying performance

Fig 2.9 shows the behavioral performance of all three rats starting at the first day of the real task rules being implemented. Sub panels with the performance of each animal on the right have the recording sessions used for analysis indicated in red (Fig 2.9 a right panels).

Simulations were run in MATLAB to look at performance for a rat with no memory (mean of 33.8%), a rat that could remember the last arm choice (mean of 50.07%), a rat that remembered the rule to go to the center after visiting the outer arms as well as which arm was visited last (mean of 50.4%) , and a rat that performed using a fixed motor plan of always visiting each arm once in a fixed order of left, center, right (mean of 66%). The simulated performance was then compared to the real distribution of behavior from the last 20 behavior sessions of all three rats (mean of 82.34%, Figure 2.9 b).

Plotting the number of errors separated by type (classified as either repeat: returning to the last arm visited, inbound: error following a left or right arm choice and proceeding to the opposite outer arm; or outbound: error following a center arm choice and proceeding to the wrong outer arm) across all sessions moving in 50 trial steps shows the early increase

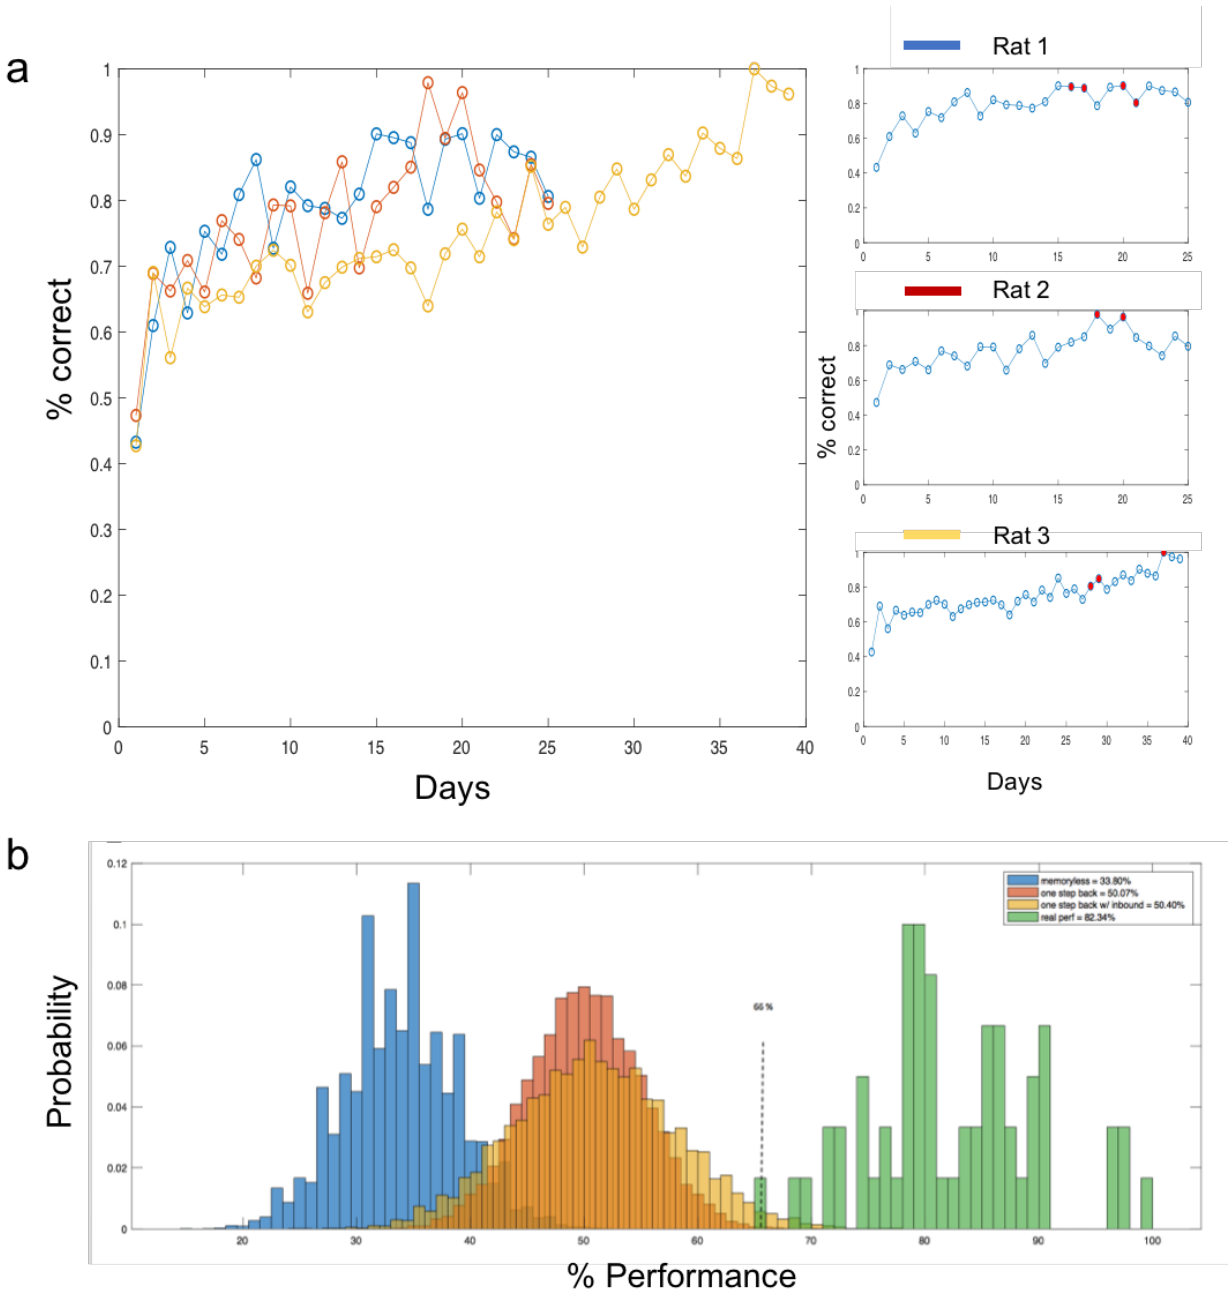

Figure 2.9: Behavioral performance in TADS
a) The performance for all three rats starting at the first day of the real task. Far right shows performance for each individual rat with sessions used for analysis marked in red.
b) Simulation of performance for a rat with no memory (blue) , a rat with memory of the last arm visited (orange), a rat with memory of the last arm visited and a rule to go to the center arm after a left or right arm visit (yellow) and a line indicating the performance for a rat with a fixed motor plan of visiting each arm once in the sequence LCR. Green shows the performance from all three rats using the last 20 sessions of behavior for each.

in performance is mostly due to a reduction in the number of repeat errors (Fig 2.10for one Rat T). All animals showed more outbound errors than inbound errors as expected due to differences in memory requirement (Kim and Frank, 2009; Jadhav et al., 2012).

Quantifying behavior

One of the important features of the task is the self-paced nature which avoids any interference by the experimenter. The maze doors, wheel, and reward delivery are all under computer control. In order to better quantify the behavior, we divided parts of each trial into three measures: the total delay time (TDEL), time to choice (T2C), and time to return (T2R).

Specifically, the TDEL is defined as the time from the start of the delay period (when all three doors are closed) until the point when the rat has completed the required eight second wheel run and all three doors are opened (Fig 2.11 top).

The T2C is defined as the moment all three doors open until the rat crosses an outer arm beam break triggering the closing of the doors for the two unchosen arms and activating the start of reward delivery for correct trials (Figure 2.11 middle).

The T2R is defined as starting at the point when the outer beam break is crossed/ alternative arm doors closed/ water reward delivery started until the rat triggers the touch pad in the delay area and closes the final door behind him to start the delay period for the next trial (Fig 2.11 bottom).

Figures 2.12, 2.13, 2.14, 2.15, gives four examples of quantification of behavior. For each trial example, in the top left the tracked motion of the animal for a single trial in the maze is plotted and colored by time on top of tracking from all other trial in the session in grey. A similar scheme below shows the entrance and exit into the wheel for wheel running during the delay. On the top right are behavioral measures of speed in the wheel (in clicks/sec) and in the maze (mm/sec). Underneath are the Y (middle) and X (bottom) pixel locations of the rat throughout the trial, also colored by time for reference to the 2d trajectory. Lastly, the

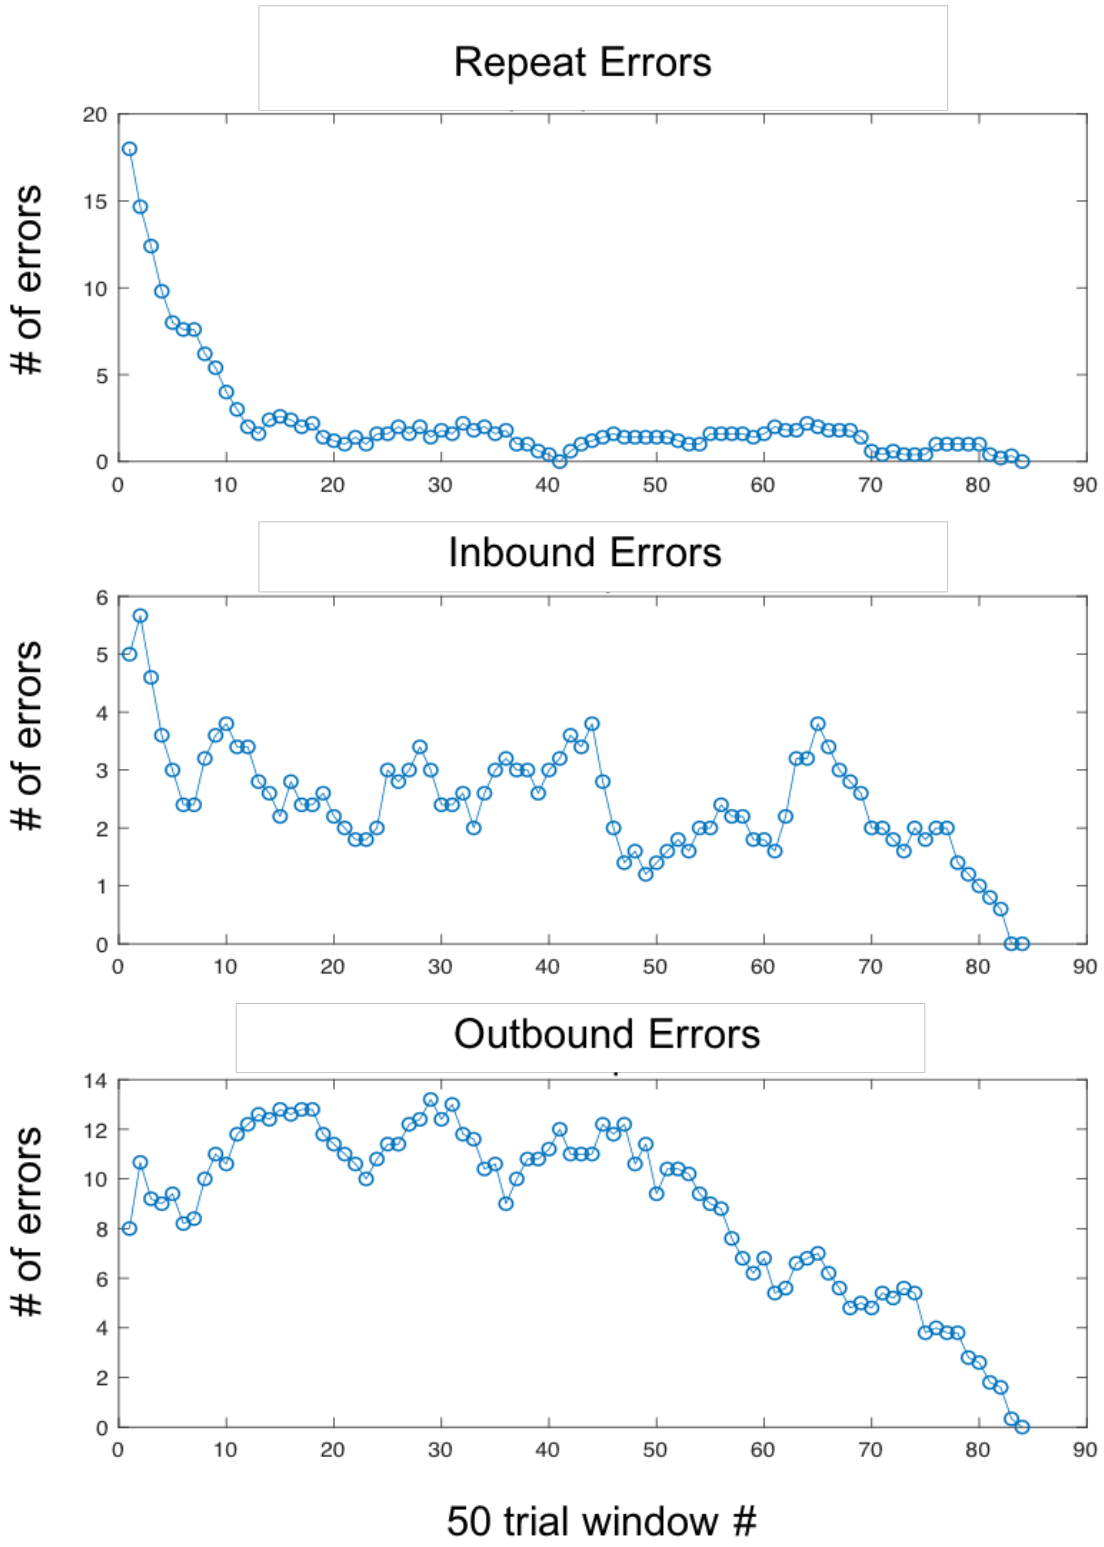

Figure 2.10: Error Analysis
Error analysis showing different error types using a moving window of 50 trials for Repeat, Inbound, and Outbound errors.

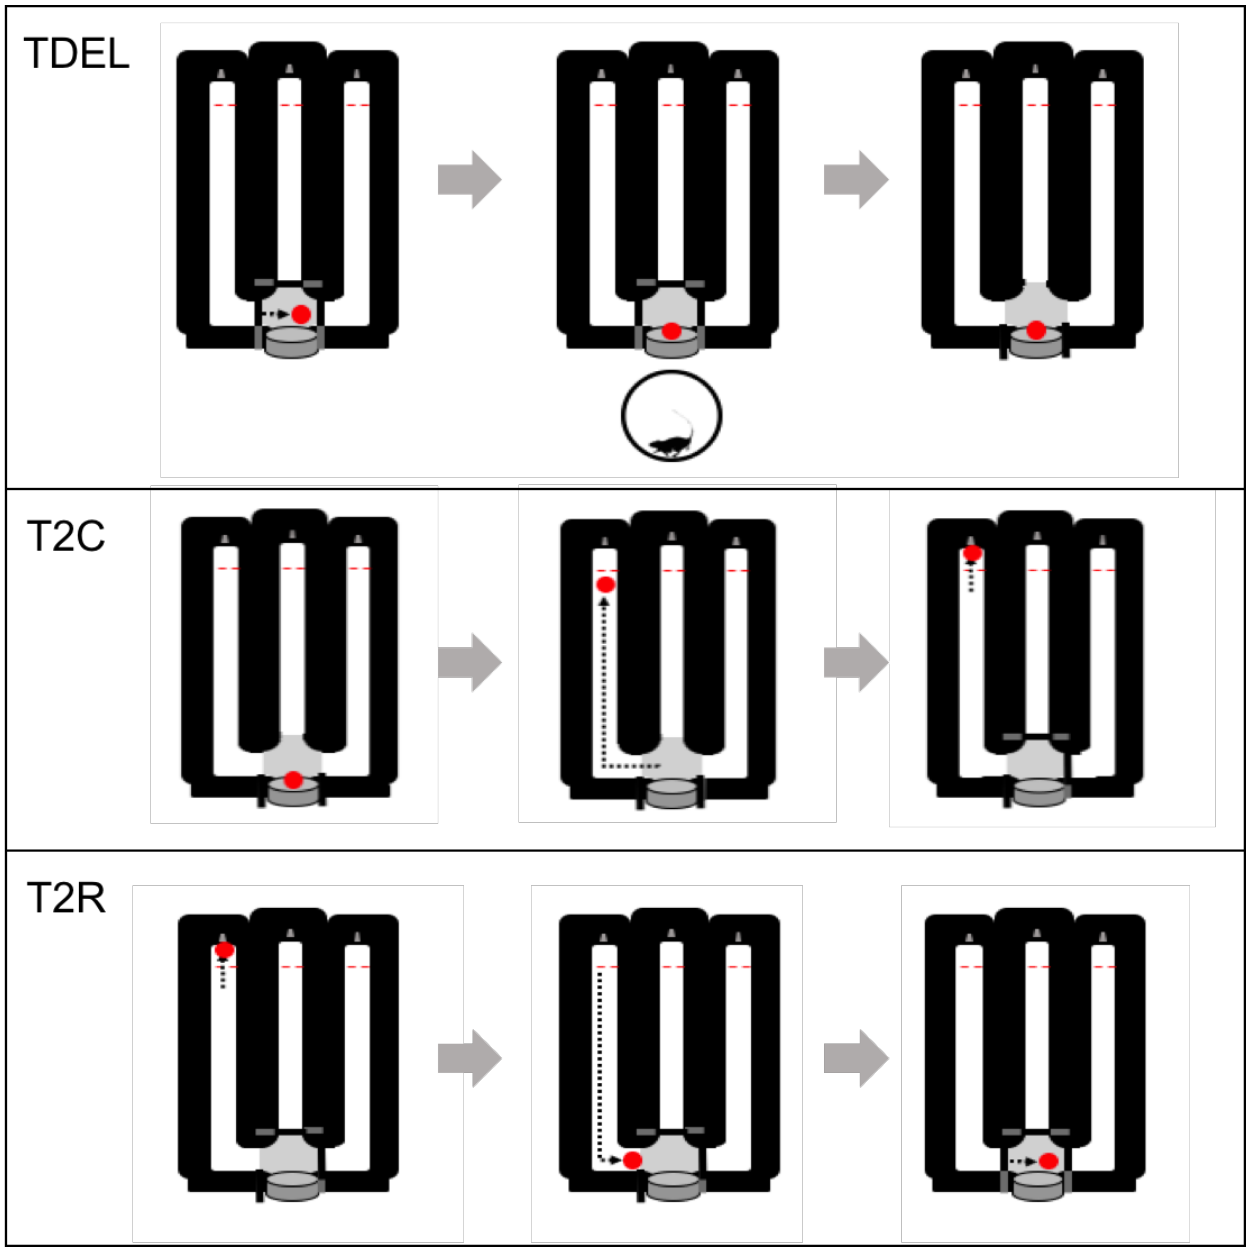

Figure 2.11: Quantification of behavior
Cartoons showing the behavior TDEL, T2C, and T2R.
Top: TDEL = total delay : the time from the start of the delay period (when all three doors are closed) until the point when the rat has completed the required eight second wheel run and all three doors are opened .
Middle: T2C = time to choose: the moment all three doors open until the rat crosses an outer arm beam break triggering the closing of the doors for the two unchosen arms and activating the start of reward delivery for correct trials.
Bottom: T2R = time to return: starting at the point when the outer beam break is crossed/ alternative arm doors closed/ water reward delivery started until the rat triggers the touch pad in the delay area and closes the final door behind him to start the delay period for the next trial.

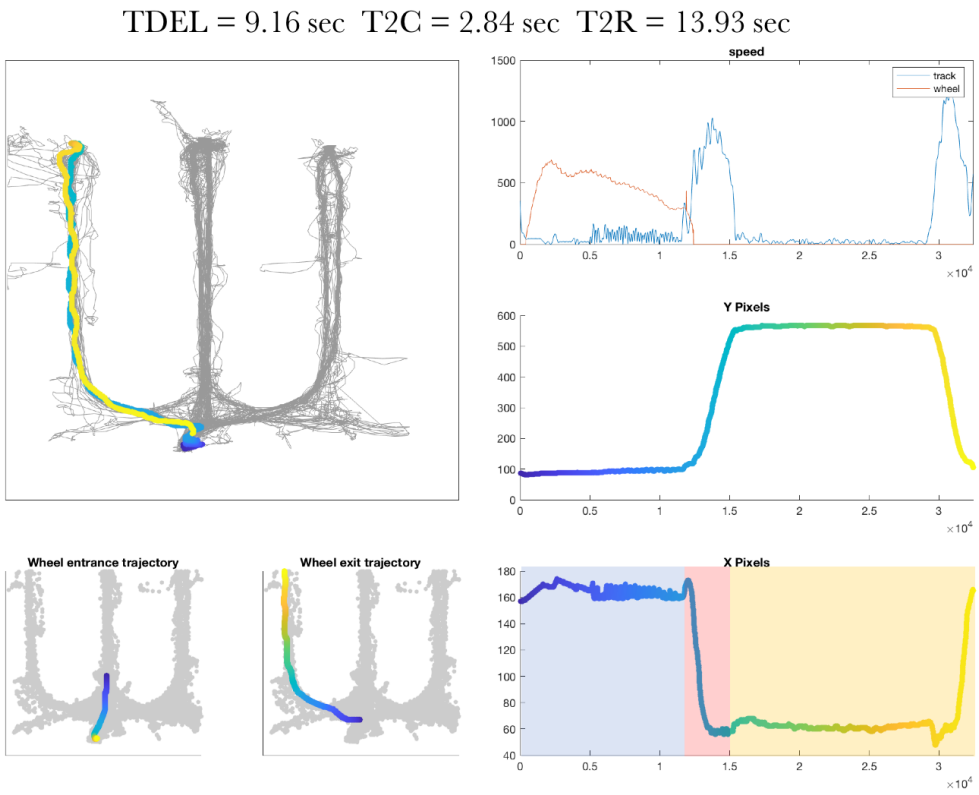

Figure 2.12: Single trial example in TADS with behavioral measures:typical
Example trial with typical behavior.

time periods designated for TDEL (blue), T2C (red) and T2R (orange) are overlaid on the X pixel location for the trial, and the times of each value are listed at the top of each trial panel. Typical behavior for most trials is reflected in the timings shown in Fig 2.12. Fig 2.13 shows an example of a trial with a long delay period (TDEL = 21.6 seconds) as a result of an initial wheel run attempt just under eight seconds followed by a successful eight second wheel run. Fig 2.14 shows an example of a trial with an extended time to return (T2R 40 seconds) as the rat reared on the wall of the left arm reward area after water consumption. Figure 2.15 shows an example of a trial with a longer time to choose (T2C = 6.04) indicated by the rat initially heading down the center arm before turning around to correctly choose the right arm.

All three behavior measures showed a significant difference between correct and error

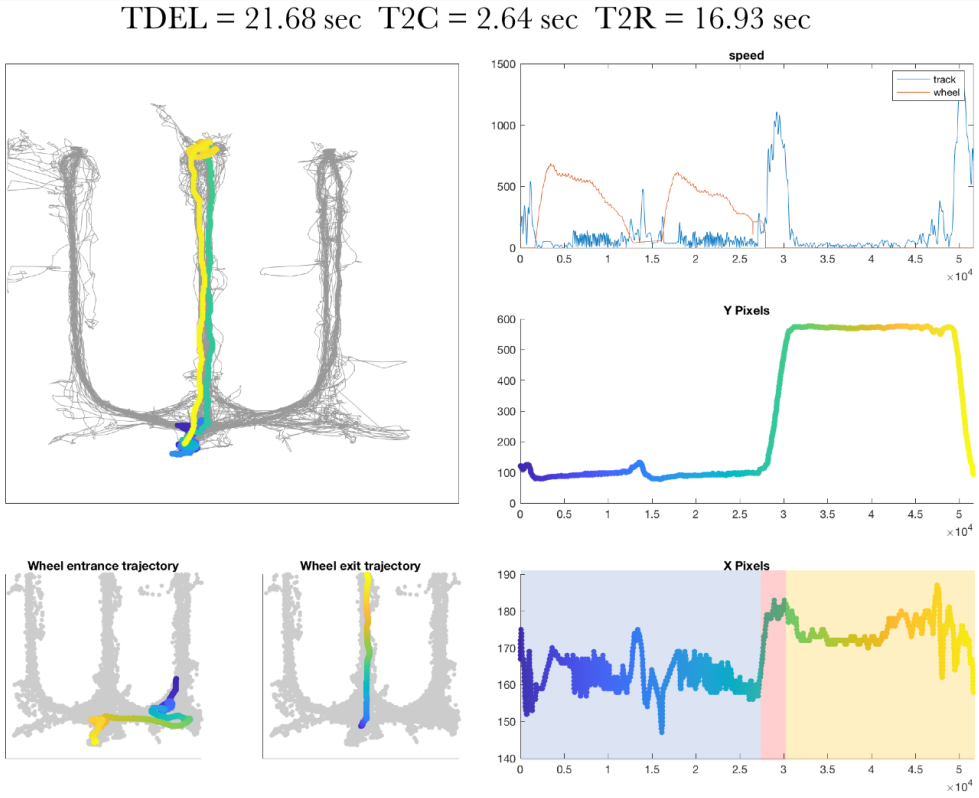

Figure 2.13: Single trial example TDEL
Example trial with an extended TDEL as a result of two wheel run attempts.

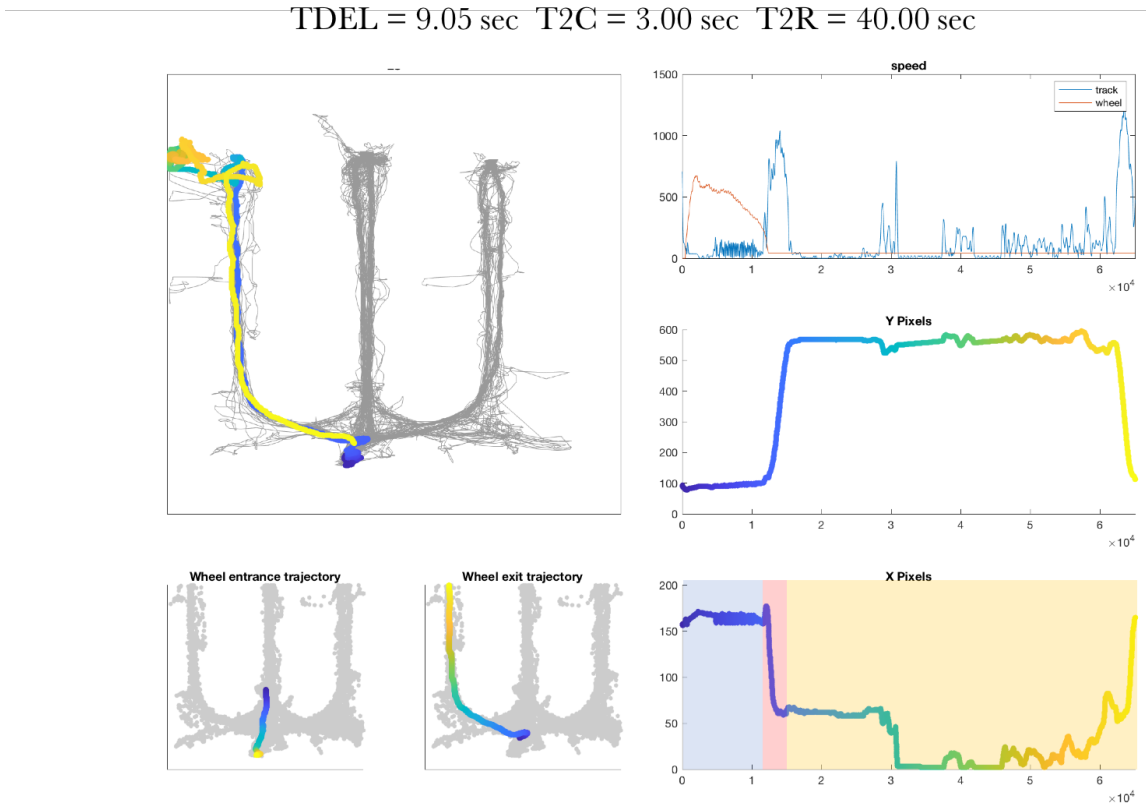

Figure 2.14: Single trial example T2R
Example trial with an extended T2R as the animal rears on the left arm wall after drinking reward and before returning.

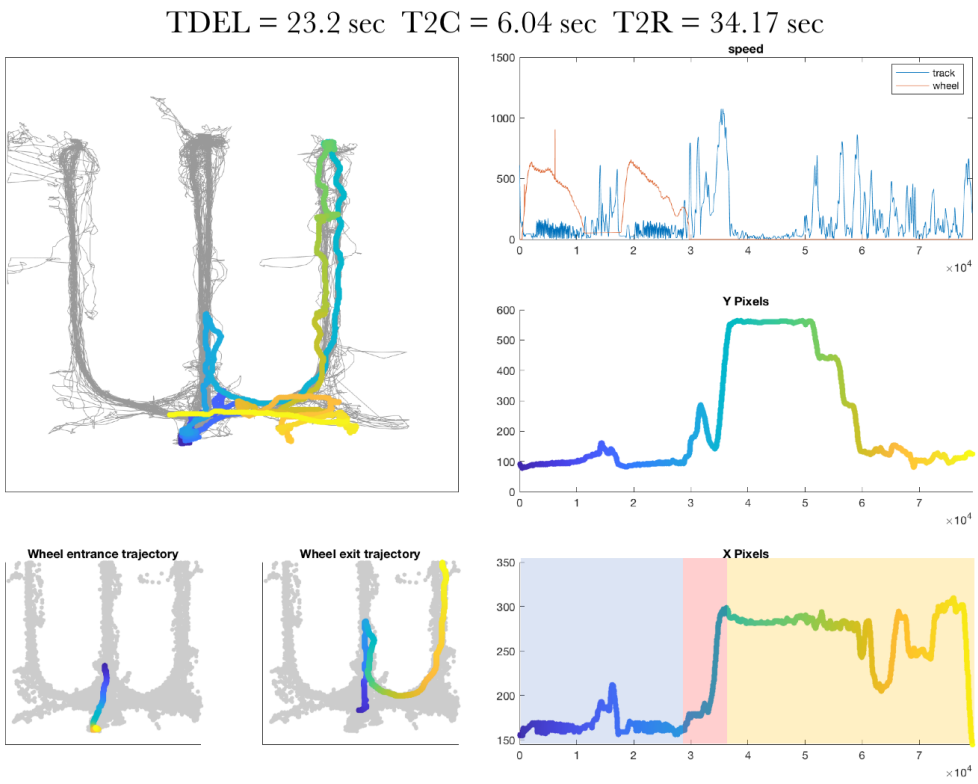

Figure 2.15: Single trial example T2C
Example trial with an extended TDEL and T2C as the rat exits the wheel and starts down the center arm before turning around and correctly going to the right arm.

trials. Error trials had significantly TDEL and T2C compared to correct trials; while T2R was significantly shorter for error trials compared to correct trials (which involved reward consumption).

2.3.2 Electrophysiology results

We now look at results from neural recordings in CA1 of the rats while performing the TADS.

RSA and LIA during behavior

As first described in the introduction there are two prominent states of activity in hippocampus (citation). RSA and LIA are clearly identified from just looking at the raw LFP on single trials during the TADS task. Using the format as Figure 2.12, Figure 2.16 shows the trajectory of a single trial. Figure 2.12b shows a spectrogram computed from LFP across a single trial for 1-24 Hz (multi taper method with window of 500ms) as well as the raw LFP trace from a single channel below. There is a clear increase in theta power during wheel and maze running indicative of RSA, while the time spent at the reward port without movement while the reward is consumed shows an absence of theta. Small segments of the LFP during running (red box) and reward consumption (blue box) are expanded on the right along with the signal bandpass filtered at 6-12 Hz (second order Chebychev Type II filter) overlaid in black (Fig 2.16 b,c).

Place cells and Episode cells in CA1 during TADS

Implanting multiple silicon probes bilateral in dorsal CA1 resulted in recordings of a large population of hippocampal neurons during TADS behavior. Fig 2.17 shows an example of the activity from 415 neurons recorded simultaneously and their activity across a single trial (sorted by the within trial firing rate). Note the clear distinction between activity during wheel and maze running (highly active theta modulated interneurons primarily at

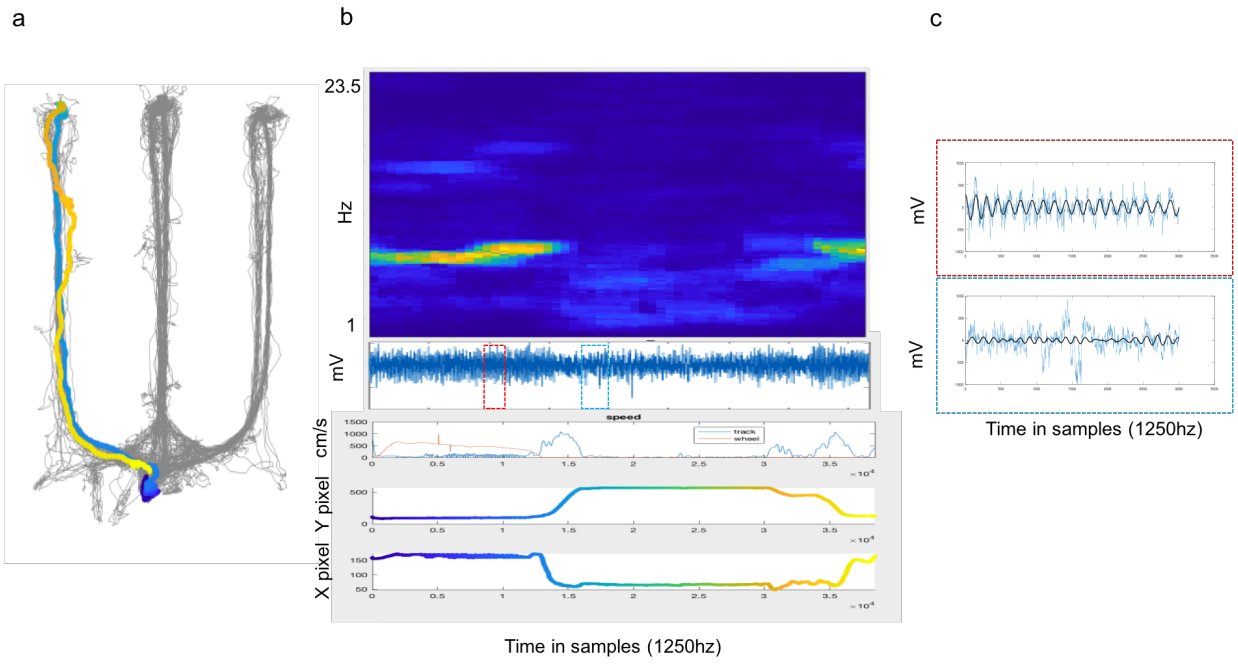

Figure 2.16: RSA and LIA in a single trial during TADS
a) Trajectory from a single trial visiting the left arm.
b) Top, Spectrogram of LFP during the single trial in a with raw LFP from a single channel, speed in the wheel and maze, and Y and X pixel location colored by time plotted below.
c) Expanded segments of the raw LFP from the trial in a with the signal bandpass filtered at 6-12 Hz overlaid in black.

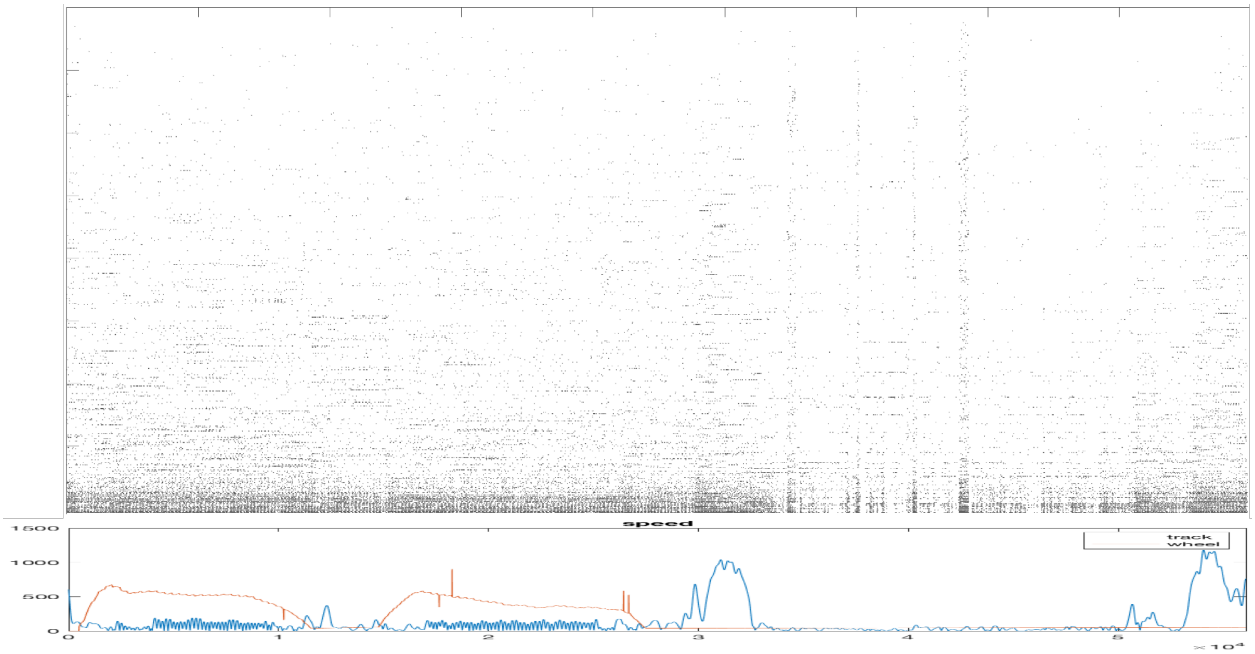

Figure 2.17: Simultaneous recording of 403 neurons for a single trial in TADS
Top shows a raster of all neurons with each row corresponding to a single neuron and spikes indicated with black tick marks. Neurons are sorted by mean FR for the trial. Bottom shows the aligned behavior for the trial with speed in the wheel (clicks/sec) and maze (mm/sec) plotted.

the bottom) and the activity during reward consumption (relatively silent with interruptions by large burst of activity – SWR events, Fig 2.17).

Figure 2.18 shows the place fields for all 415 neurons in the recording session, with each map normalized to the individual neuron's peak FR.

Figure 2.19 shows example neurons with their activity in the maze shown in 2d as well as activity sorted by trial type or arm visited for episode field wheel running and the three arms respectively.

Context specific activity during performance of TADS

Having recorded neural activity in rats successfully performing the TADS task and confirming the presence of known signatures of neural activity in CA1 of hippocampus, we will now characterize context dependent activity during the wheel running in the delay period.

In order to look at the population activity across the entire delay period we plotted the

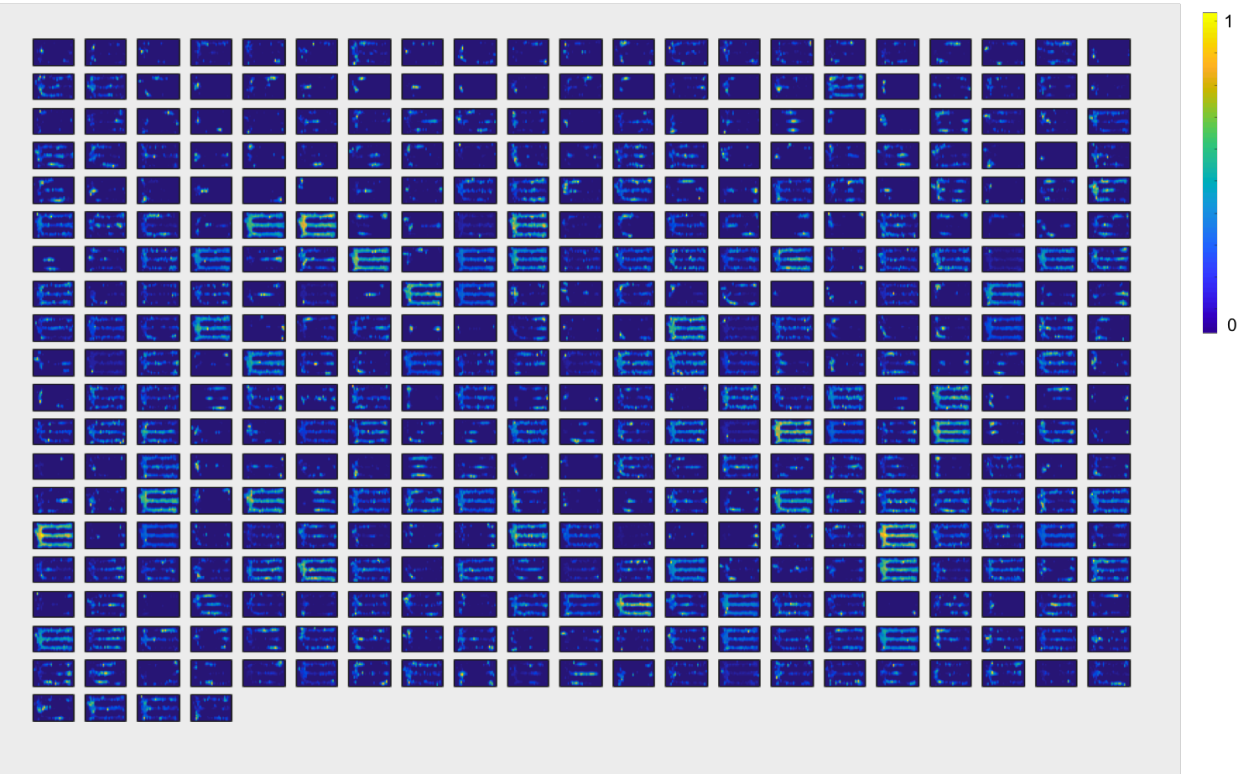

Figure 2.18: Place field maps from 403 simultaneously recorded neurons
Each map shows FR normalized to peak.

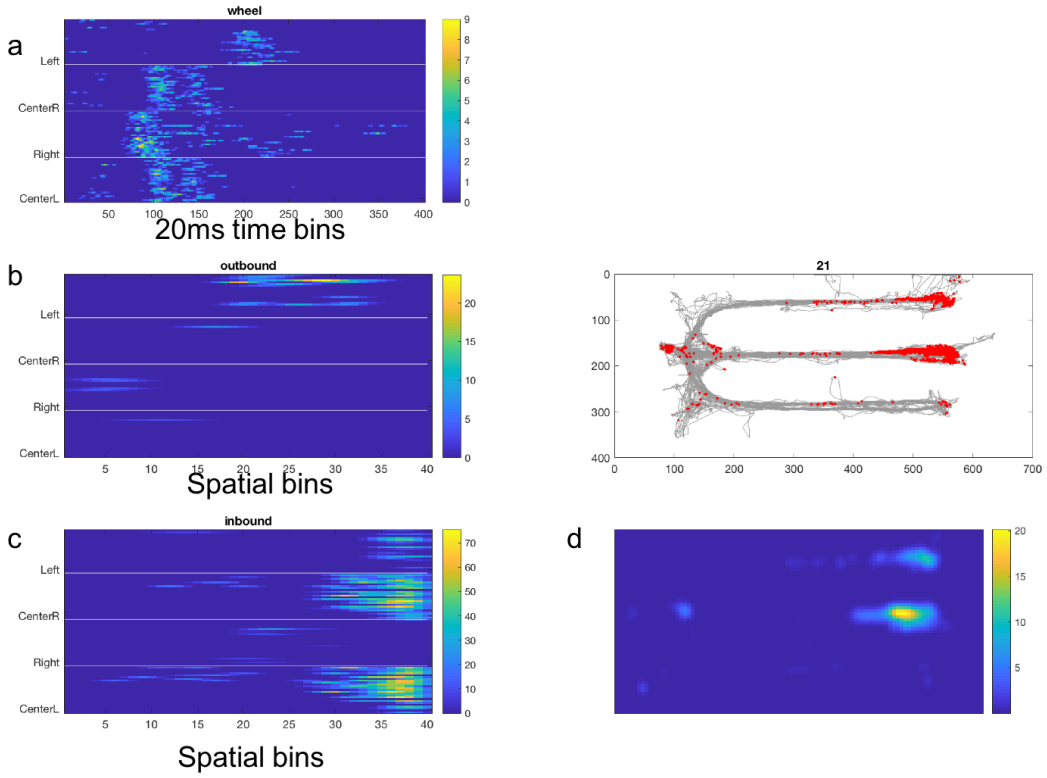

Figure 2.19: Example of Episode fields, linearized place fields, and 2d place field for a single neuron

a) Episode fields for the example neuron separated by trial type are shown for wheel run following Left, Center to Right, Right, and Center to Left arm choices. Each row is from a single trial.

b) Linearized outbound place field activity for all four trial types for example neuron.

c) Linearized inbound place field activity for all four trial types for example neuron.

d) 2d place field for example neuron.

mean firing rate over time for each neuron normalized to its peak from all the wheel runs of a single delay type (for instance all wheel runs after a left arm choice) and look at this for every neuron recorded with the order of the cells dictated by the latency to the peak firing of that cell. This shows a clear sequence of neural activity spanning the wheel run during the delay period (Figure 2.20 top left matrix). We can then repeat this method for each of the four trial types (shown across the diagonal, ordered population activity during the wheel runs for trials from left, center to right, right, and center to left arm trials). The snake like line of activity running across each plot on the diagonal is a result of the sorting process. For any matrix of activity, rows of the matrix (the firing of each different neuron) can be sorted according to latency to peak and show structure such structure. However, what can be informative is looking at activity for different trial types using common orderings. If the activity of across the population is preserved, then you would expect to see the maintenance of sequential structure across the matrix. However, if the population activity is different, then the structure will be degraded. The rest of the matrices in Figure 2.20 shows activity from different contextual trial types (columns) each sorted by the order dictated by sorting by latency to peak for a different trial type (rows).

While this visualization is informative, we would like to avoid averaging across trials and see if the sequence structure across the population holds on the single trial level. For example, we can look at 10 example trials in Figure 2.21 first with no sorting on the top row. Sorting the next row by applying the row ordering from one trial to the other trials reveals structure for similar trial types on the single trial level across the neural population during wheel runs.

In order to quantify the single trial context dependent activity during wheel running we can create a matrix looking at the Pearson's correlation value computed from all pairs of trials and their population activity matrices for the wheel runs (Figure 2.22).

Sorting the correlation matrix by trial type shows clear structure where trials with the same context are highly correlated in their population activity (Figure 2.23).

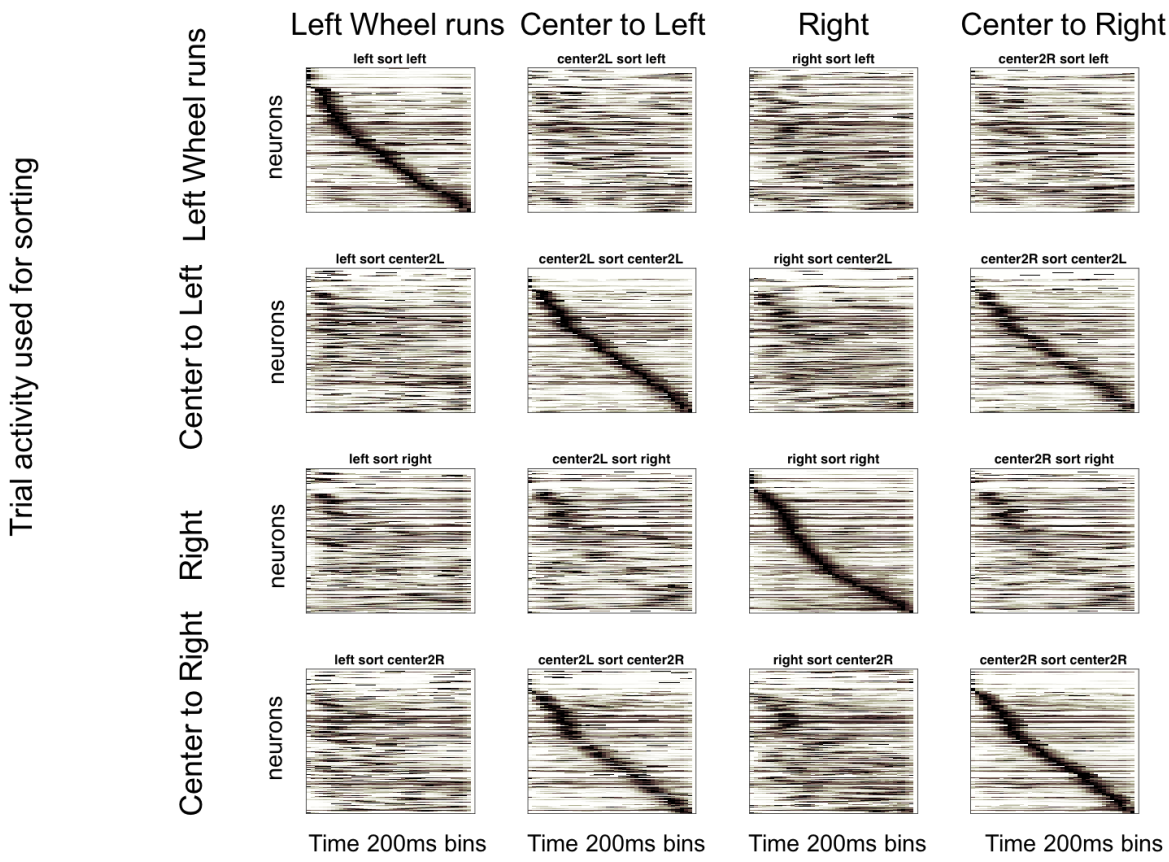

Figure 2.20: Ordered population sequences from mean activity of different wheel run contexts

Top row shows mean responses from 4 trial contexts sorted by the ordered latency to peak using left trials. Subsequent rows are the same as above but using order from Center to Left, Right, and Center to Right wheel runs.

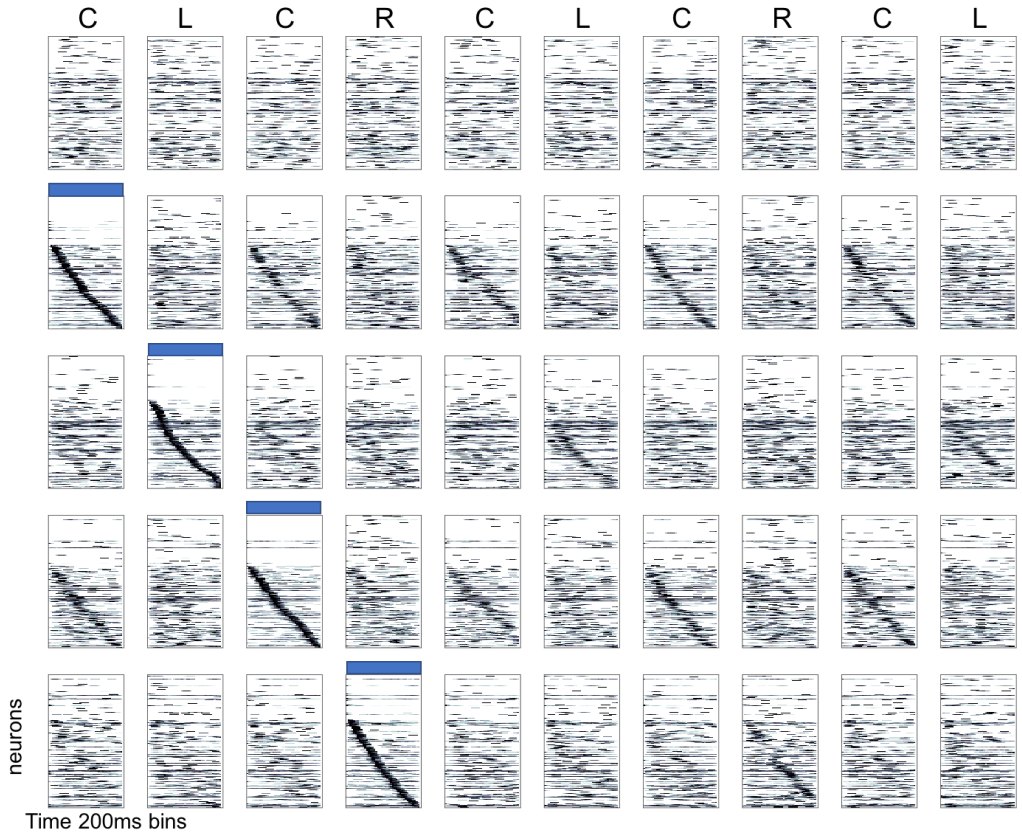

Figure 2.21: Population sequence over single trials
Same concept as figure 2.15 but each row shows the same 10 single trials and their population activity across the wheel run. Label at the top shows the arm that was visited before the wheel run. Top for shows the activity for the 10 wheel runs with no sorting of the neurons (rows) for each matrix. Each row below shows the same 10 trials but with the each trial using the sorting of neurons dictated by sorting the trial with a blue bar for neurons sorted by latency to peak. So 10 trials are shown (columns) with the same activity shown matrix shown for each column. Only change is for each row there is a different order used for sorting each individual matrix rows (neurons). Columns are the same trial with different orderings. Rows are the same ordering across different trials.

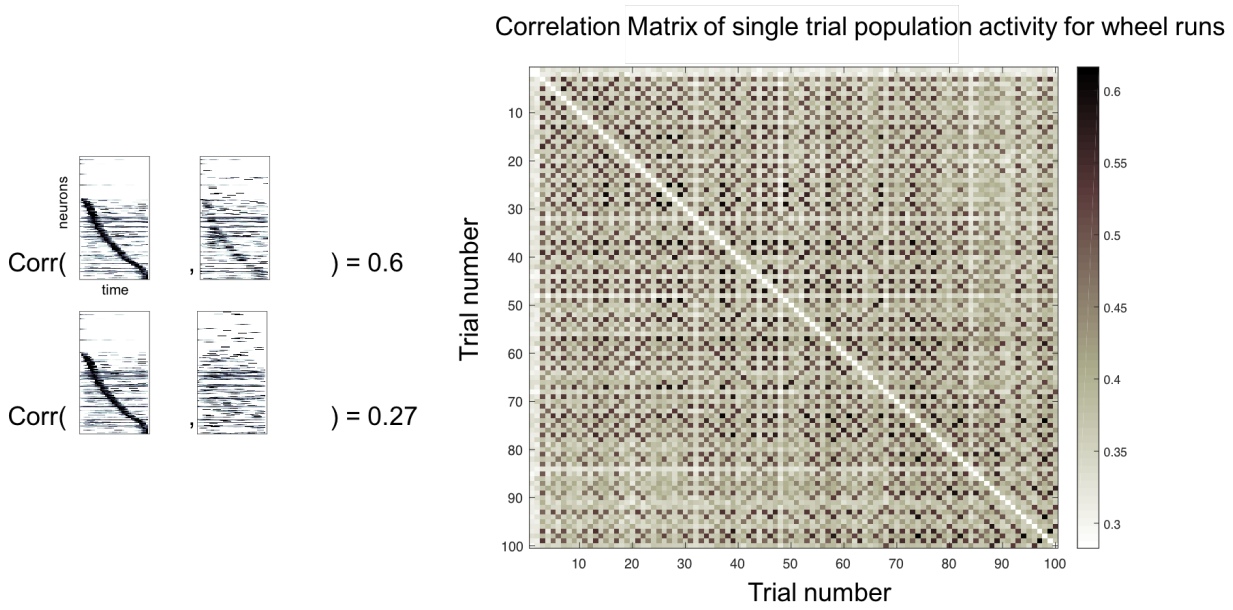

Figure 2.22: Correlation Matrix of single trials population activity during wheel runs. Example of the method used for calculating each value in the large matrix – the correlation between the population activity of each trials eight second wheel run is calculated. Each example uses the same ordering of neurons.

While the clusters along the diagonal represent similar trial types showing correlated activity there is a clear off diagonal correlation between the center to left and center to right trial types (Figure 2.23). Reordering the comparisons such that we plot all trials from the left, center, and right in order shows that both trial types from the center are equally correlated (Figure 2.24).

Classifier-based search for contextual coding

While these results do show context dependent activity during the delay period between the three spatial arms, as demonstrated in the introduction we know that the task could not be performed at the level we are observing without the rats having the ability to maintain two different contextual versions of the center arm delays. Therefore, we can attempt to check if there is any information in the neural activity that is sufficient to discriminate the multiple trial types (which is the definition of context dependent coding), in particular searching for signs of contextual coding for the two center arm delay periods. The classifier can look

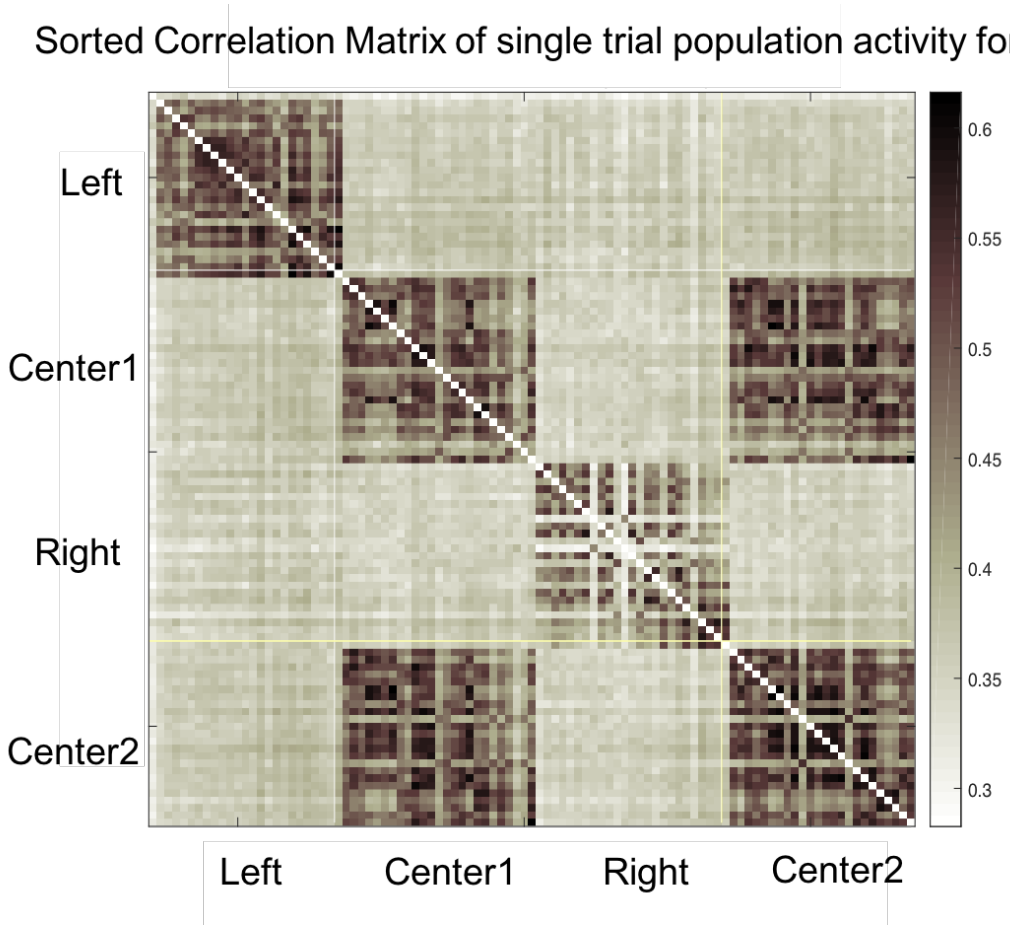

Figure 2.23: Sorted Correlation Matrix of single trial population activity for wheel runs
Correlation matrix from single trial population activity during wheel runs sorted by trial type (L, C1, R, C2).

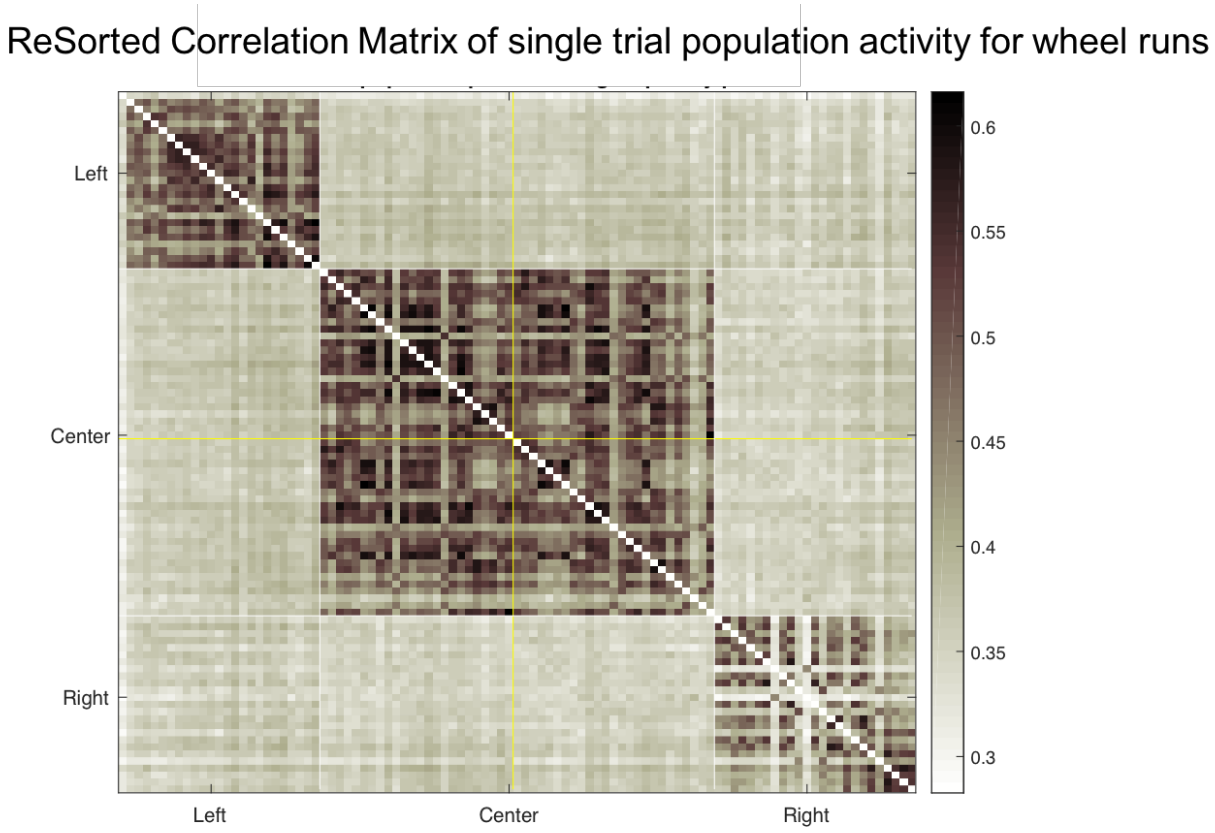

Figure 2.24: ReSorted Correlation Matrix of single trial population activity for wheel runs
Correlation matrix from single trial population activity during wheel runs Re-sorted by trial type. (L, C1,C2,R)

for possible changes in context that the correlation analysis may miss, such as systematic changes in firing rate. To do this we performed the following test using a linear classifier. We analyzed the performance of a linear classifier trained to distinguish between all four trial types, each against each other resulting in six comparisons (LvR; LvC1; LvC2; RvC1; RvC2; C1vC2). For confirmation that the existence of context dependent activity was not an artifact of differences in wheel running behavior, the same comparisons were made using the speed profiles from the wheel runs for each trial, which serves as a sensitive proxy for the motor behavior of the animal while running in the wheel. Finally, we created vectors of place field activity for the linearized inbound and outbound trajectories when running down each arm in order to look for any signs of splitter activity in the two types of center arm trials located in the maze arms. Figure 2.25 shows the results of classifier performance for the recording session from Rat T, while Figure 2.26 shows the collected results from the analyzed recording for all three rats.

The results from the linear classifier performance 1) confirmed that the presence of information in the neural activity on a trial by trial basis was sufficient to distinguish between delays following the three different maze arms; 2) confirmed neural activity during the delay was insufficient to distinguish between the delays following the two contextually different center arm trials; 3) confirmed that there was no information in the motor behavior sufficient to explain the contextual coding; and 4) found that the place field activity in the three arms were sufficiently distinct to distinguish each from the other; however there was no evidence for contextual coding in the place field activity located in the center arm for the two contextually different center arm trials.

2.4 Discussion

The TADS task required rats to visit three arms in the sequence of L,C,R,C with an eight second wheel run delay between visits. Consistent with previous reports, we did find evidence of internally generated activity during the wheel runs showing context dependent coding.

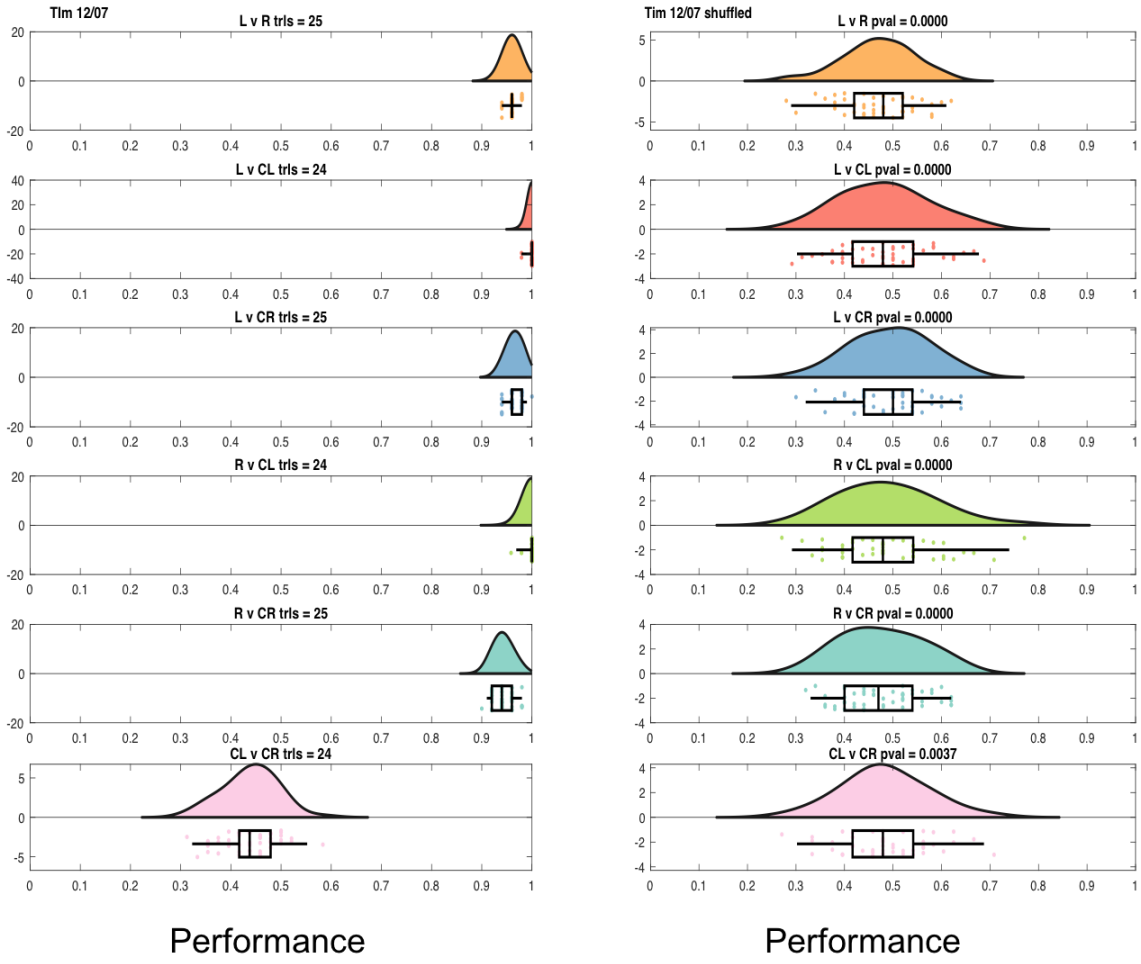

Figure 2.25: Classifier performance for six trial comparisons using 200ms binned neural activity across wheel runs

Left plots show classifier performance using single trial types. For example, the Top Left shows performance from the comparison Left versus Right trials with the number of each trial type available from the session shown (25 trials of each). On the top right the performance of a classifier with the same data, but with the trial labels shuffled. Pval shows the P-value from a two sided test on the actual and shuffled classifier performances.

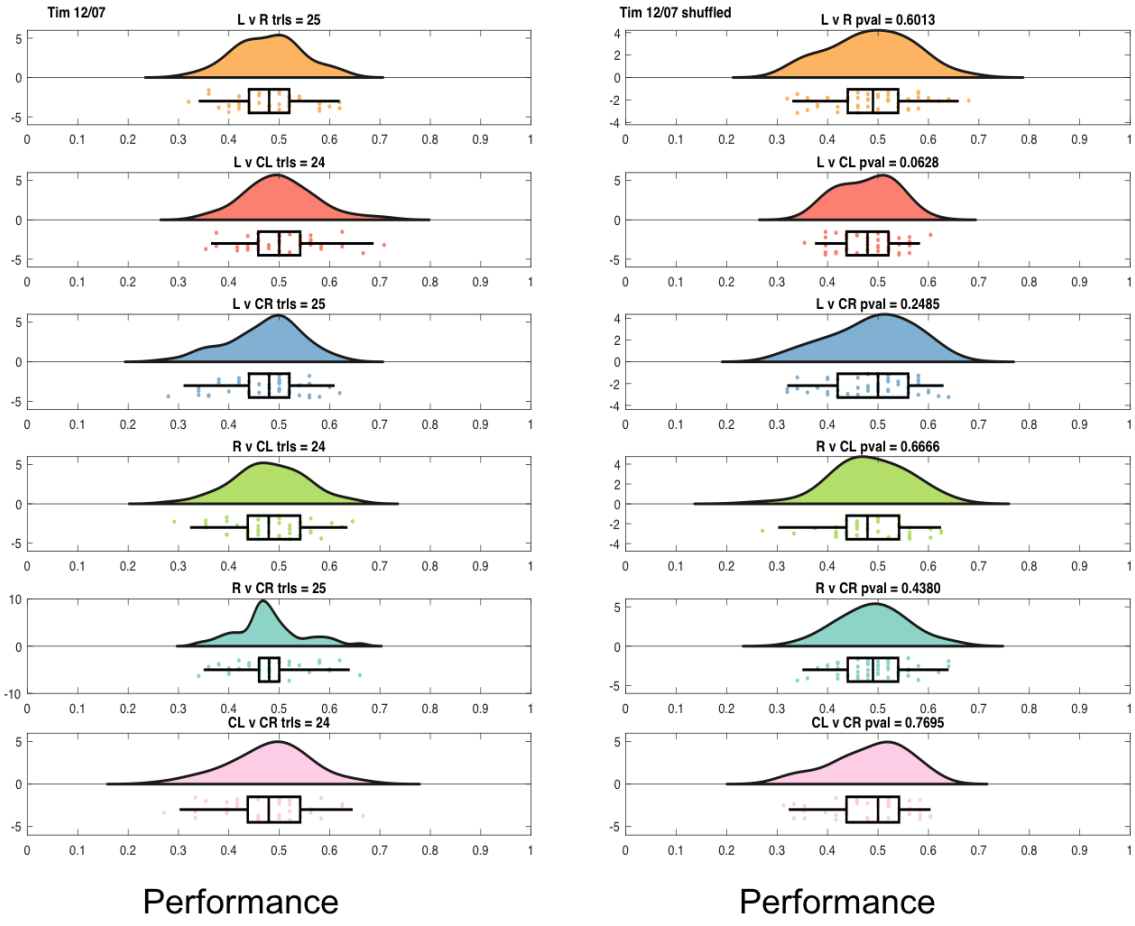

Figure 2.26: Classifier performance for six trial comparisons using 200ms binned wheel speed measurements
same format as 2.25

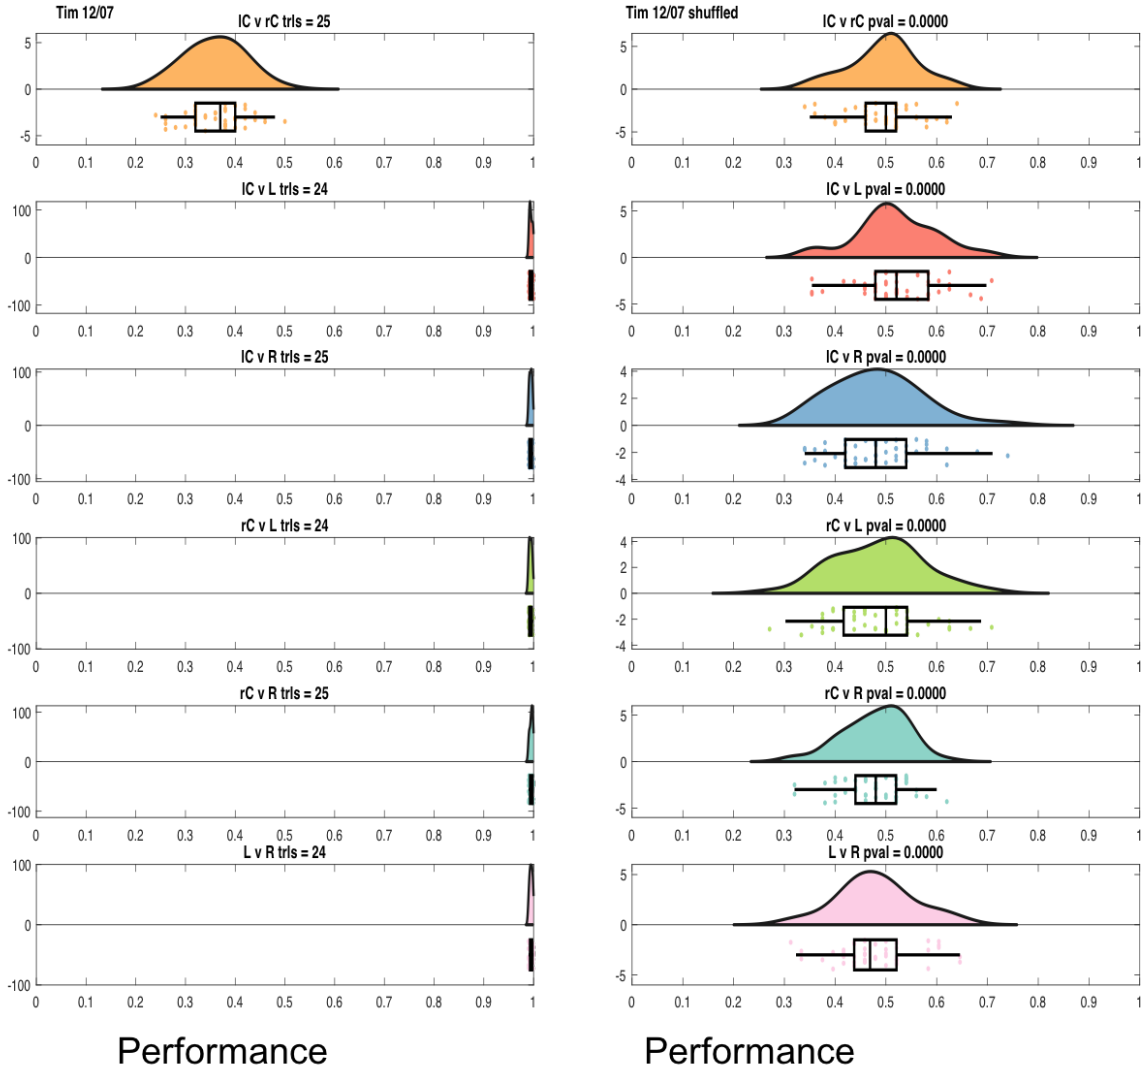

Figure 2.27: Classifier performance for six trial comparisons using binned linear place fields

same format as 2.25 Inbound and outbound place fields for each trial were concatenated to a single vector for each trial.

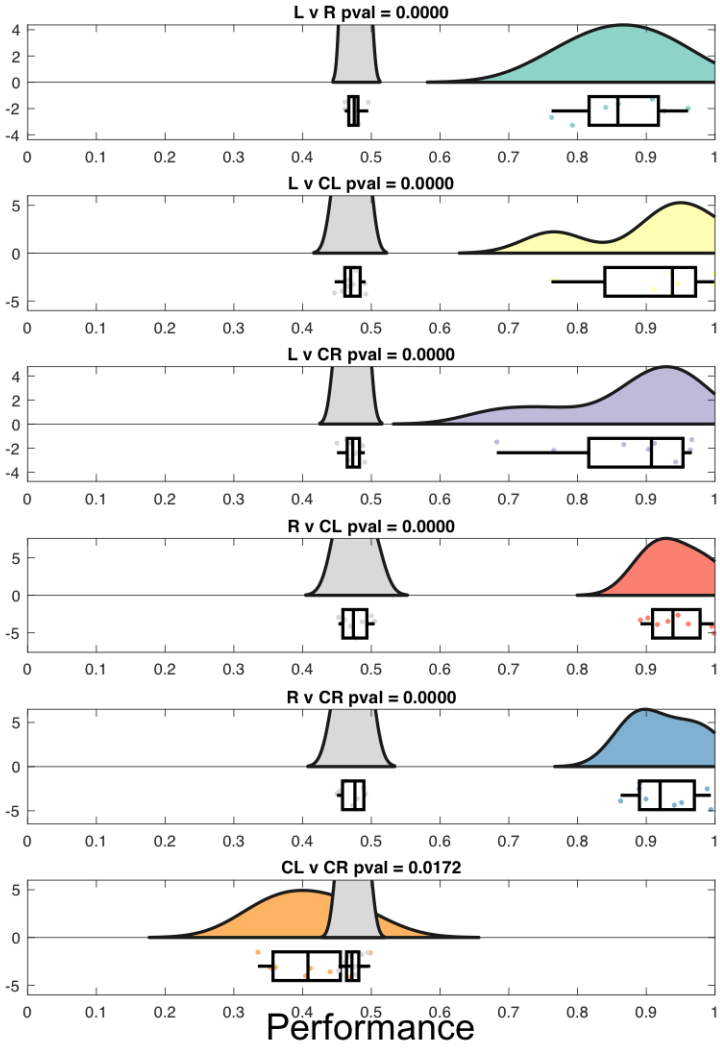

Figure 2.28: Classifier performance for six trial comparisons using 200ms binned neural activity across wheel runs for all sessions and rats
same format as 2.25 but with shuffled label version for each classification plotted in gray.

Surprisingly the context dependent coding was only able to distinguish between periods following the three spatial arm choices and there was no evidence for neural activity that could distinguish the two types of delay period following center arm choices. These results conflict with the known requirement that in order for the rats to achieve the levels of performance observed, they must have an internal distinction between the two types of delays following center arm choices, otherwise their performance on these outbound choices would be 50% (see performance simulations Fig 2.9).

One possible interpretation could be that the hippocampus is important primarily remembering the last arm visited, which although not the only requirement for this particular task, is still a basic component. For instance, avoiding repeats to the last arm chosen is one of the first lessons observed when the rats improve performance during learning, indicated by the increase in performance correlating with a drastic reduction in the number of repeat errors for all three rats (Fig 2.10). This interpretation would also be consistent with the results of Pastalkova et al 2008 in which they were able to predict errors. If the activity in the wheel represents a memory of the last arm visited, then in a two-choice task, mistakenly remembering the last arm visited would lead to making an error on the next trial ((Deadwyler et al., 1996) 'miscoding'). However, our results conflict with this interpretation because in this case we found consistent coding for the past arm choice before all error types including repeat errors. Thus, evidence for the context coding for the last arm visited being utilized by the rat in the task is sparse although it may be ill advised to rule out interpretation based on error trials alone.

In conclusion, these results suggest we rethink the dominant hypothesis in the field that context dependent coding seen in hippocampus is the substrate of contextual memory used in a delayed spatial memory task.

CHAPTER 3

BAYESIAN DECODING OF SWRS DURING TADS

3.1 Introduction

In the previous chapter, we explored neural activity in rat hippocampus during performance of the TADS memory task. In particular, we searched for context dependent coding within the internally generated activity found during delay period wheel running. Successful performance of the task required the creation of four contextually distinct delay periods. While we did observe contextually distinct internally generated activity, different sequences could only distinguish between the delays following each of the three arm choices. Any indication of contextual coding for the delays following the two types of center arm trials was not detectable. However, there is another window into internally generated activity in hippocampus besides that found during wheel running: the sequences generated during sharp wave ripple events (SWR). SWR events recorded during the LIA state are known to internally generate sequences that replay behavioral timescale sequential activity on a compressed timescale. In this chapter, we will start with a brief review of evidence that awake SWRs might be important for behavior as well as represent planning of future trajectories. Then we will use a Bayesian decoder to look at the content of SWRs detected during reward consumption in rats performing the TADS task. Finally, we will show that some of the decoded SWRs during TADS show spatially remote trajectories that correlate with future behavioral choices in a statistically significant manner. The presence of two different internally generated plans observed during reward consumption at the center arm reward port under different behavioral contexts provides support for the maintenance of context in the brain during center arm trials which was missing from the previous chapter.

As mentioned in Chapter 1, one of the more exciting discoveries from recordings in rat hippocampus was finding replay of behavioral sequences at a compressed timescale within SWR activity. We will now review two papers that provide important insight into the pos-

sible function of SWRs in the behaving rat. The first paper is Jadhav et al in 2012. While it had been previously shown using perturbation of SWRs that SWRs might play a role in consolidation of memories during sleep for the purpose of learning (Girardeau et al., 2009) , Jadhav and colleagues looked at the effect of perturbing the hippocampus and interrupting activity specifically during awake SWRs while rats performed a continuous three arm alternation task (W maze) used in the Frank 2000 paper (Fig 3.1 a). Although the continuous version of this task has been found to not require the hippocampus, selective perturbations may be more insightful into function compared to lesion studies which can suffer from compensatory mechanisms. Similar to design of the TADS task, in the Jadhav task the outbound trials from the center arm were considered the more difficult trials that required working memory. Subsequently, it was specifically when they looked at these outbound trials that they found an impairment in the learning and performance of rats which had their awake SWRs interrupted (Fig 3.1 b,c). This paper provided some of the first evidence that internally generated activity during awake SWRs might play an important role in a memory task.

The second paper looking at awake SWRs and behavior is Pfeiffer and Foster 2013. In this work, rats were trained on a task that involved exploring a large 2d environment with multiple reward wells (Fig 3.2a). On any given day, there would be a reward well that would be designated the "home" well and for each trial the rat could find a reward at this "home" well and then would need to explore the surrounding environment to search for reward at one of the randomly selected "away" wells. Then the rat knew after every "away" reward collected to return to the "home" well for another reward and then to once again explore the environment to find the next new "away" well with a reward. Pfeiffer and Foster were able to record from a large number of hippocampal CA1 neurons simultaneously during the behavior (Fig 3.2 b), which allowed them to decode the 2d content of SWRs that occur over a very short period of time (50-200ms). They looked at decoded trajectories from SWRs which occurred during pauses in the rats' movement (typically at the rewarded well

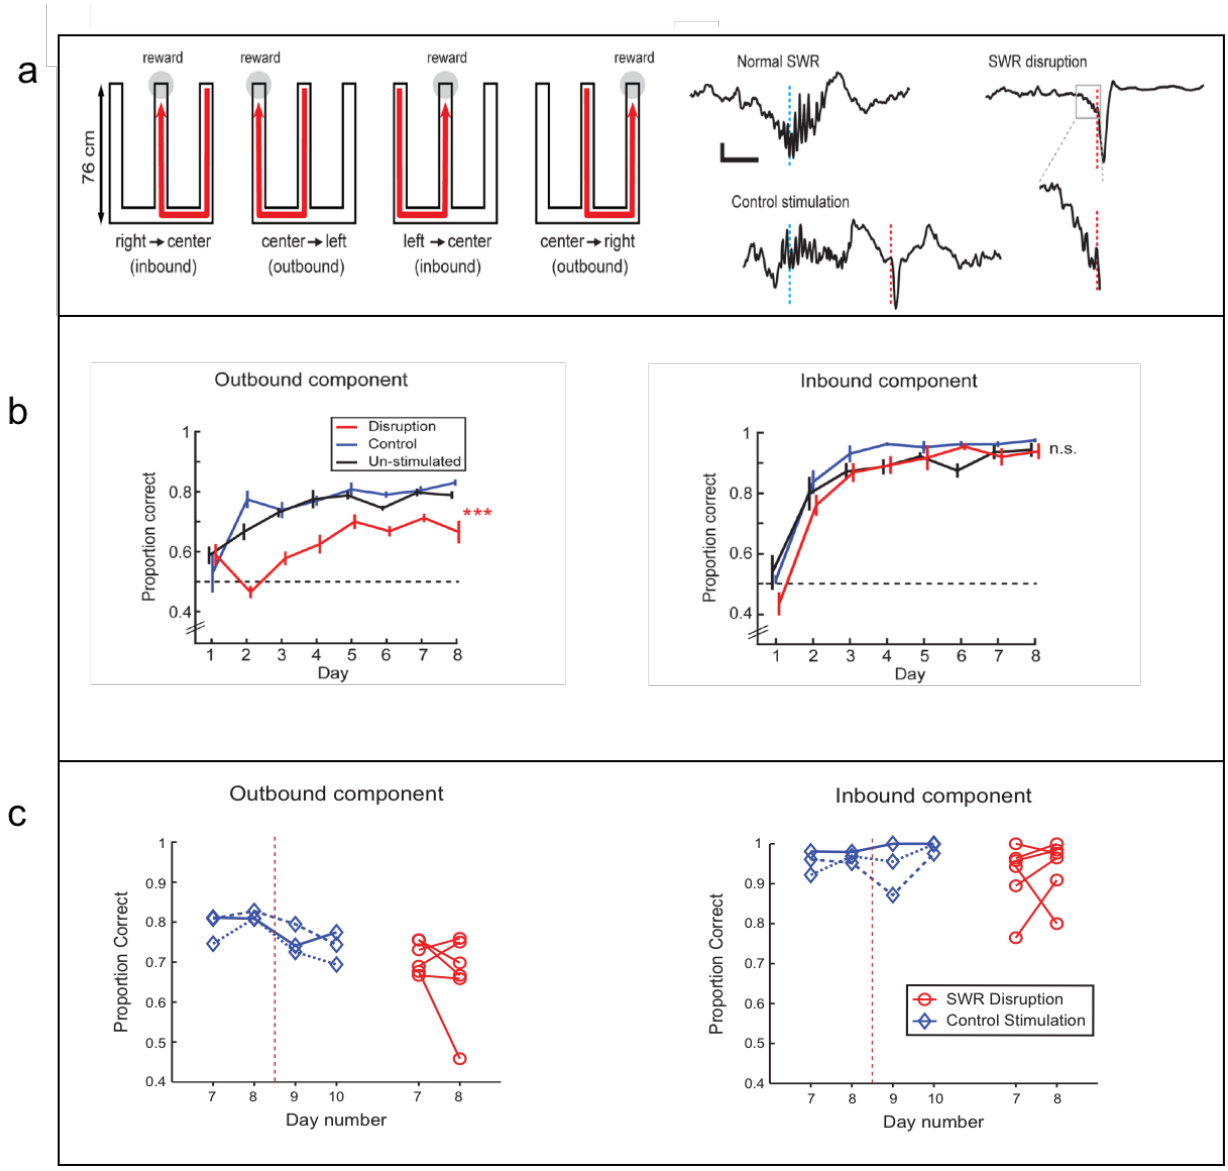

Figure 3.1: Perturbation of Awake SWRs impairs performance in a spatial memory task

a) Cartoon diagram showing the behavior task and examples of disrupting ripple from Jadhav et al 2012

b) Plots of behavior for rats with ripples disrupted (disruption), control (stimulation is applied at a random time after a ripple is detected), and un-stimulated. A significant impairment is seen in learning and performance on outbound trials for rats with ripples disrupted.

c) Performance on outbound trials shows an impairment even in Control animals who had learned the task and then were switched to real ripple disruption on days 9 and 10 (in blue last two days).

locations, during reward consumption) and asked if the trajectories matched the immediate past trajectory of the rat or the immediate future path (Fig 3.2 c-k). They found a significant difference in the absolute angular displacement of the decoded SWR trajectories compared with the past and future trajectories, with the difference between the future being smaller. These results suggested that the SWRs appear to correlate with the future behavior of the rat (Singer et al., 2013; Pfeiffer and Foster, 2013).

These papers both support the hypothesis that SWRs might serve a functional role in spatial memory tasks, and that they might be correlated with planned future behavior for the rat. With this in mind we decoded SWR events during performance of the TADS task to see if there was any correlation between the content of decoded SWRs and behavior.

3.2 Methods

Bayesian decoding of SWR content involves three important steps. First is the detection of putative SWR events typically referred to as population burst events (PBEs) during the LIA state. Second is the decoding of the SWR content using a model based on Bayes' rule and modeling of each neuron's activity as a Poisson process. Finally, there must be a principled way of determining which decoded trajectories are meaningful and which are merely noise or ones that we are unable to sufficiently decode. In the next few paragraphs we will cover the details used for each step.

3.2.1 PBE event detection

Detection of putative SWR events or PBEs during the task was based on the following criteria. First all neural activity was summed and binned (20 ms bins from a sliding window of 20ms moving in 10ms steps). This histogram was then smoothed using a moving average of five bins (100ms). Then for each trial peaks in the population activity histogram were detected using the findpeaks function in MATLAB with a minimum peak prominence deter-

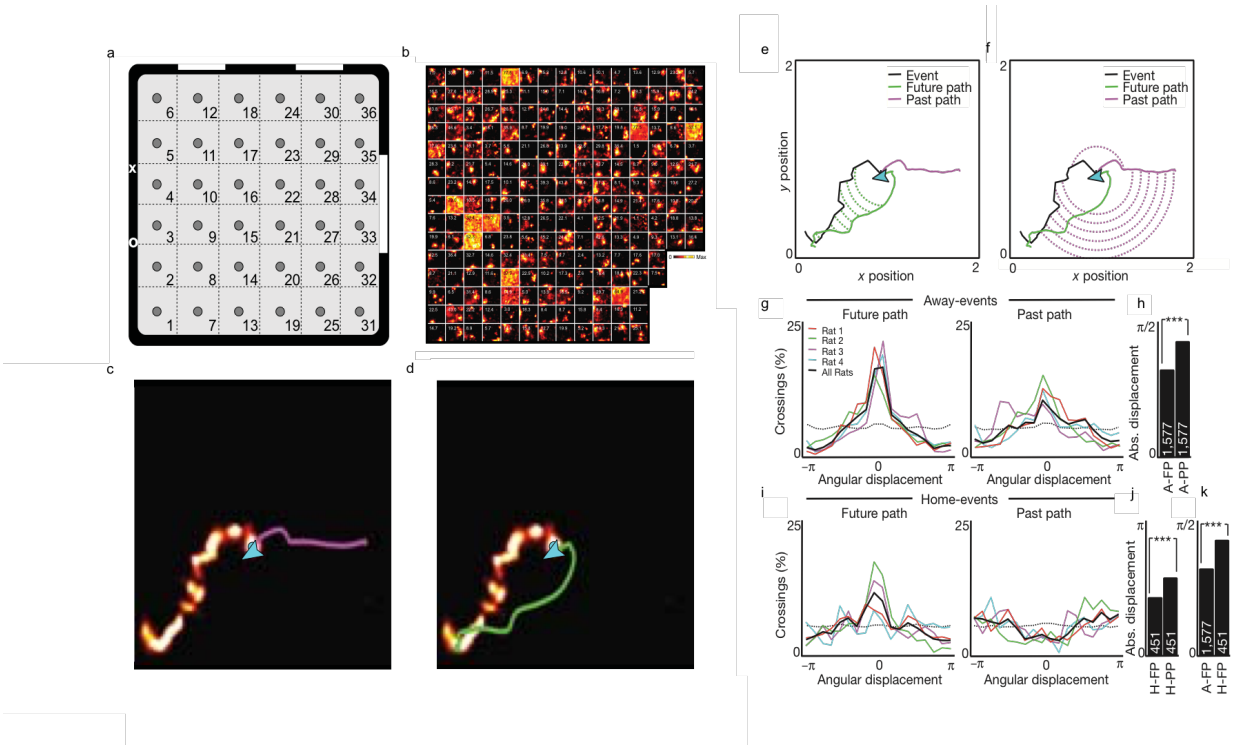

Figure 3.2: Awake SWRs correlate with future trajectories

a) Cartoon diagram of the 36 well 2d maze used in Pfeiffer and Foster 2013
b) Examples of 212 simultaneously recorded place cells
c) A decoded trajectory event collapsed over time to show the trajectory of the peak probability compared to the past behavior of the rat before the event
d) The same as C but with the future trajectory shown
e) Showing the method for quantifying the distance of angular displacement between the trajectory event and the future path f) Same as e but for past
g) Angular displacement for trajectories across all rats for events at away wells for future and past paths
displacement for future trajectories is significantly less than past trajectories
h) Same as g for events at the home wells
i) Trajectory events at the home well also show significantly less angular displacement with respect to future path compared to the past
j) Trajectory events at away well shows significantly less different from future paths than those are home wells and future paths

mined by the mean + standard deviation x 2 of the activity during the LIA state for each trial (see following). The mean activity of the LIA state was determined as the mean of activity when the rats' speed was below 5 cm/second and was located at one of the three arm reward ports (Fig 3.3, marked in brackets). Once peaks were detected, the windows defining a PBE was set by marking the start and end of the PBE as the point before and after the peak when the firing rate crossed below the mean of that trial's LIA period. If this did not happen before another peak was detected, then the PBEs were merged together. Finally, only PBEs that occurred with the following criteria were kept to ensure that they were indeed during consumption of reward, the rats speed had to be below the threshold of 3cm/second, indicating they were not moving, and the rat had to be located in the maze section defined as the reward zone – which is where the reward port was. These criteria yielded SWR events that occurred during reward consumption while the rats had their noses in the reward port, which further guaranteed a fixed head direction (Fig. 3.3). The average length of the PBEs detected was 250 ms.

3.2.2 Bayesian Decoding of PBE events

After detecting the PBEs the content of these putative SWRs was decoded using a Bayesian decoder with a uniform prior.

Bayes theorem states:

$$P(x|n) = \frac{P(n|x)P(x)}{P(n)}$$

Definitions:

- $P(n|x)$ is the likelihood of observing a vector of spike counts n when the spatial location encoded by hippocampus is in bin x. We are going to approximate this by assuming (1) independence (2) Poisson firing. - $P(x|n)$ is the posterior we want to estimate. We can write the posterior as $P(x|n) \propto P(n|x)P(x)$, where $P(x)$ is the prior (occupancy map of the animal). A normalization constant is chosen such that $\sum_x P(x|n) = 1$. This means that we

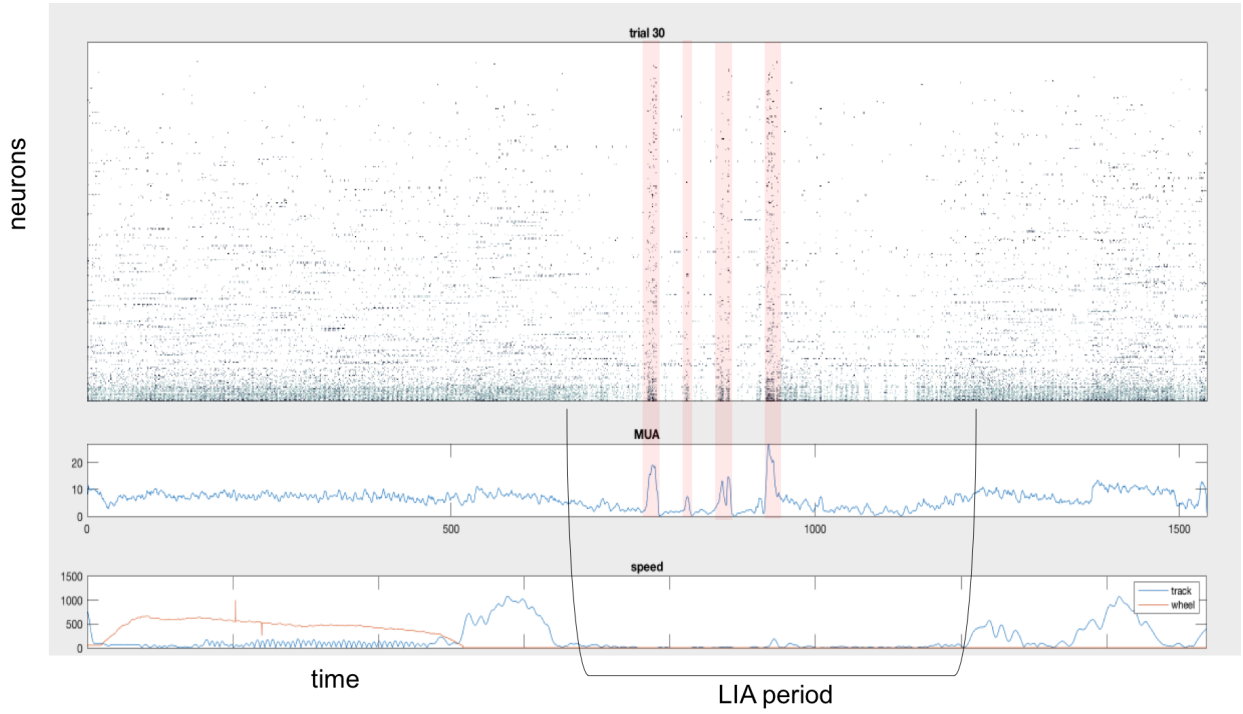

Figure 3.3: PBE detection
Activity of 403 neurons in 20ms bins for a single trial is shown on top. Below is the smoothed activity of the entire population. An example of the period defined as LIA period marked with brackets (When the rat is at the reward port and not moving. Mean of population activity across this period is used in the method for PBE detection. Bottom shows speed in the wheel and maze across the trial.

don't need to take care of factors that do not depend on x.

Likelihood approximation: (1) Independence: The likelihood factorizes in a product of likelihoods for the N neurons: $P(n|x) = \prod_i^N P(n_i|x)$

(2) Poisson firing: The likelihood for an individual neuron follows a Poisson distribution with rate $f_i(x)\tau$, where $f_i(x)$ is the standard place/episode field map estimated from the running data, and τ is the window used for decoding.

The likelihood is then written as $P(n_i|x) = \frac{e^{-f_i(x)\tau}(f_i(x)\tau)^{n_i}}{n_i!}$.

Posterior: Given the above, we can write the posterior as

$$P(x|N) \propto P(x) \prod_i \frac{e^{-f_i(x)\tau}(f_i(x)\tau)^{n_i}}{n_i!}$$

.

Each PBE was decoded in each frame using a 20ms window moving in 10ms steps. Only frames with more than 5 spikes were used for decoding. The PBE events were decoded using two collections of probability maps or encoding models. The first was built up using the 2d place fields for each neuron. The second was created by concatenating the outbound and inbound linearized place field maps for all three arms into a single vector, and thus the likelihood was determined for any of the three arms inbound or outbound trajectories for each arm. The linearized inbound and outbound maps for each arm gave 52 bins so all three arms meant 312 possible spatial locations. Therefore, the probability of the rat being located in any given location by chance is 1/312 or 0.0032 or 0.32%.

Following the decoding of each PBE onto the spatial vector using all three arms, the decoded maximum likelihood for each of the inbound and outbound trajectories of the three arms was extracted separately (Fig 3.4 a). Time is on the y-axis (20ms decoded frames moving in 10 ms steps). The map being decoded with is a vector constructed of the population place fields for the six different linearized arms all concatenated into the same probability map (binned space on the x axis). Thus for each row the probability for each location in binned space (Left arm outbound, Left arm inbound, center out, center in, right out, right

in) is decoded across a 20ms frame. Decoding of location in space proceeds in time from top to bottom. The dotted lines marking the map edges are for illustration purposes only. Color scale for entire decoded trajectory is normalized to the peak probability. Note directly above is the same decoder output but broken up by map and with the color scale normalized to the peak probability of each individual map separately. While you may see structure in the Right arm out decoded map, note it is much lower compared to the probability for the center arm map as shown by being undetectable in a when the color map commonly scaled. Then for each frame the maximum likelihood was identified within each map (Fig 3.4 b). Then maximum likelihoods with a probability higher than 0.1 were identified and all other peaks set to zero (Fig 3.4 c). Finally, for each map, if there were at least six frames with a peak probability over 0.1 left, then the Spearman correlation was calculated for the decoded trajectory within that map. Only decoded trajectories with a correlation value greater than 0.9 or less than -0.9 where kept as indicating a clean decoded ripple trajectory with a constant velocity (Davidson et al., 2009). Also included in the figure is at the top right shows the past, current, and future trial during the task when the detected and decoded PBE event occurs. Trajectories for the trials are plotted on top of the entire trial trajectories in gray. If the trial is correct the trajectory is blue, errors are in orange. Red dot shows the location of the rat when the event is happening. So in this example the rat just made an error to the left arm, and now is drinking water after a correct choice to the right arm. The next trial the rat will make a center arm choice. Below the trials is the peak of the maximum likelihood for all frames of a decoder using the 2d place cell maps projected onto a cartoon diagram of the maze. Dots are colored by time with cooler colors as earlier, and warmer as later in the decoded event.

3.3 Results

Decoding of SWRs during reward periods of the TADS task revealed several clean replay events. Fig 3.5, 3.6 ,3.7 shows examples. While there where ripples that originated at the

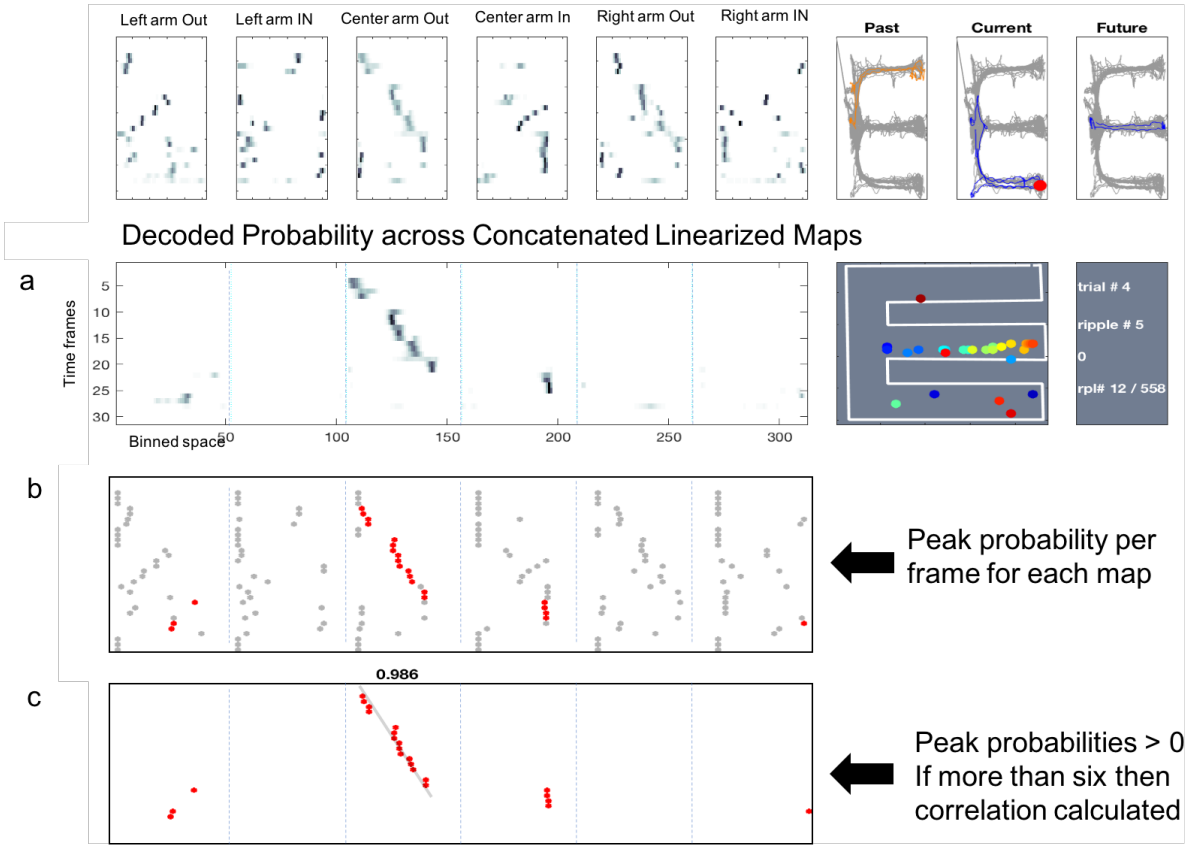

Figure 3.4: Method for detection of clean decoded SWR trajectories and layout format for subsequent figures

a) The output of a Bayesian decoder for a PBE event that lasts 310ms.

b) The same as above but the maximum likelihood is marked for each time frame in each of the six maps. Maximum likelihoods with a probability above 0.1 are colored red.

c) the same as above but with only the peaks above 0.1 probability kept. If the decoded output left in a given map has more than 5 frames above 0.1 probability then the correlation is computed for that event. The example shown has a correlation of 0.986.

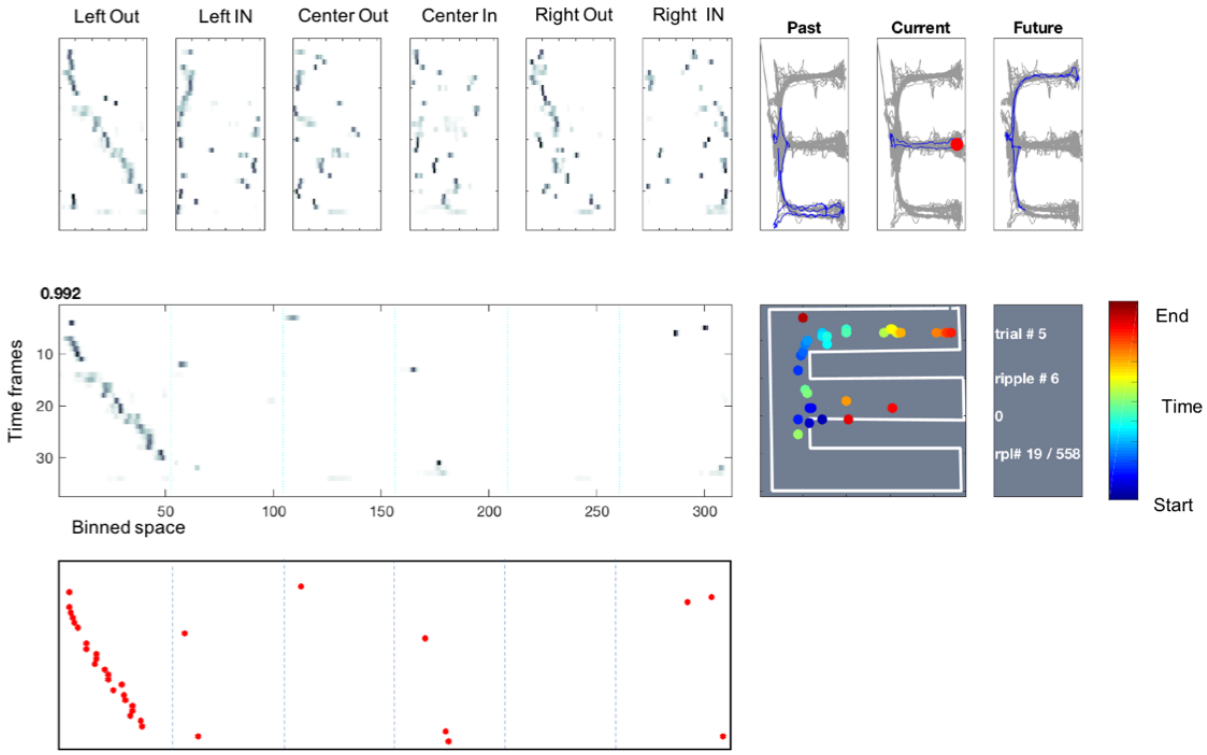

Figure 3.5: Remote trajectory event example 1

location of the rat in space, several of these trajectory events appeared to replay trajectories across arms other than the one the rat was located in (i.e. "remote trajectories) , which has only rarely been observed in other studies (Davidson et al., 2009; Karlsson and Frank, 2009).

Looking at the decoded trajectories within the context of the ongoing behavior during the task, it appeared that the remote trajectories matched the future behavior of the rat on the next trial. For example, in Figure 3.8 decoded ripples are shown for three consecutive trials during performance of the TADS task. The location of the rat when the ripple occurs is indicated on the maze in a large cyan circle. Decoded activity for the ripple is shown for each of the six linearized maps at the bottom of each trial. Also, the peak probabilities from each frame using the decoder with 2d probability maps are plotted onto the maze itself with decoded activity colored in time. For trial 8 the rat is located in the right arm consuming reward when decoding of a detected SWR shows a clear trajectory starting at the delay

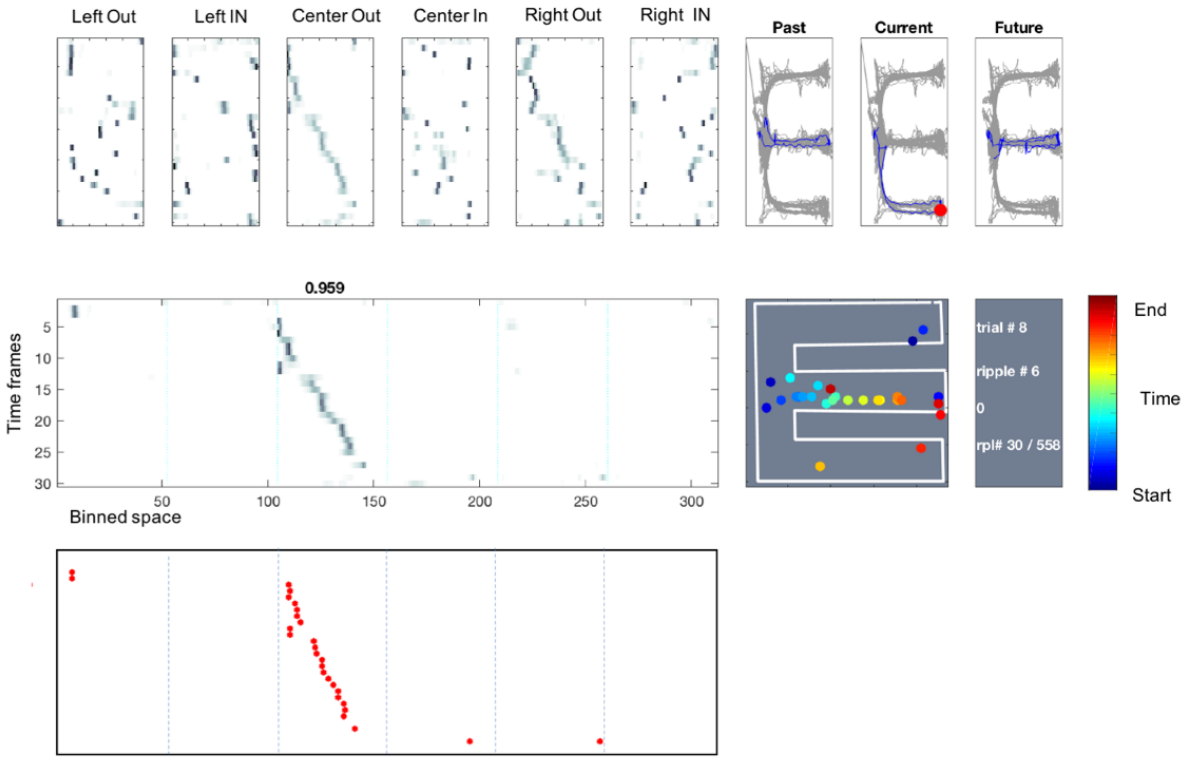

Figure 3.6: Remote trajectory event example 2

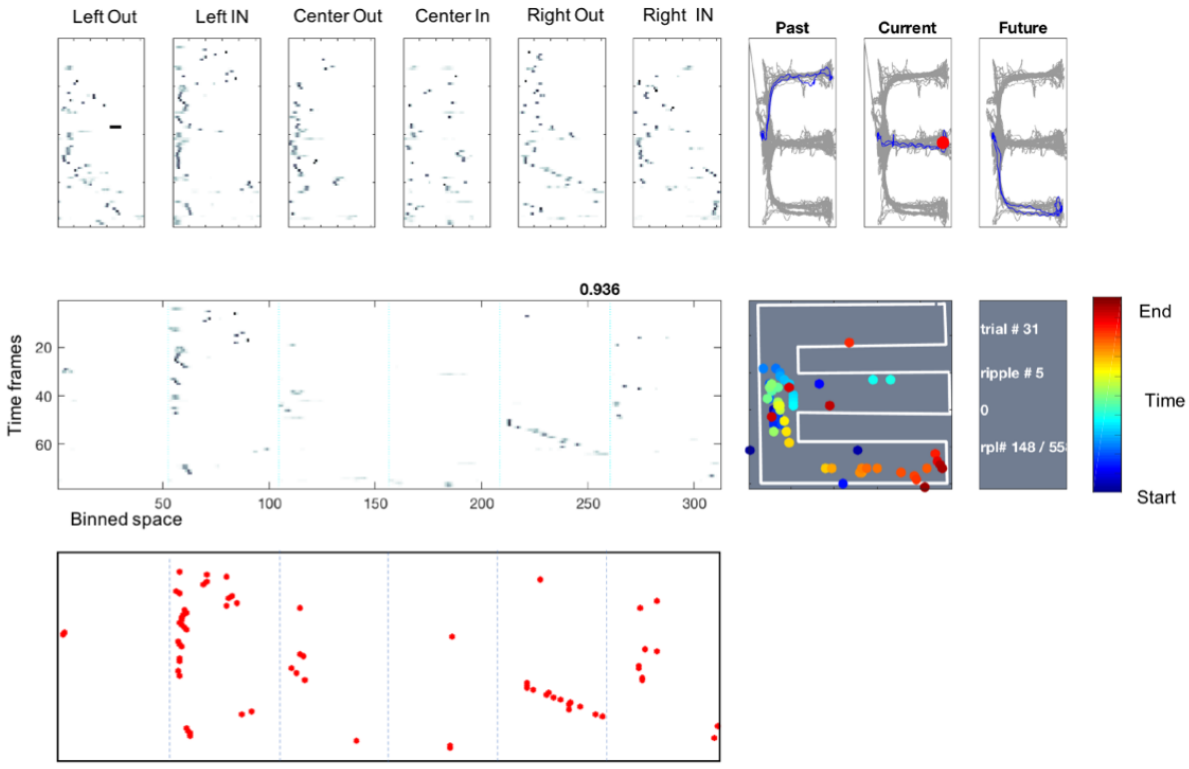

Figure 3.7: Remote trajectory event example 3

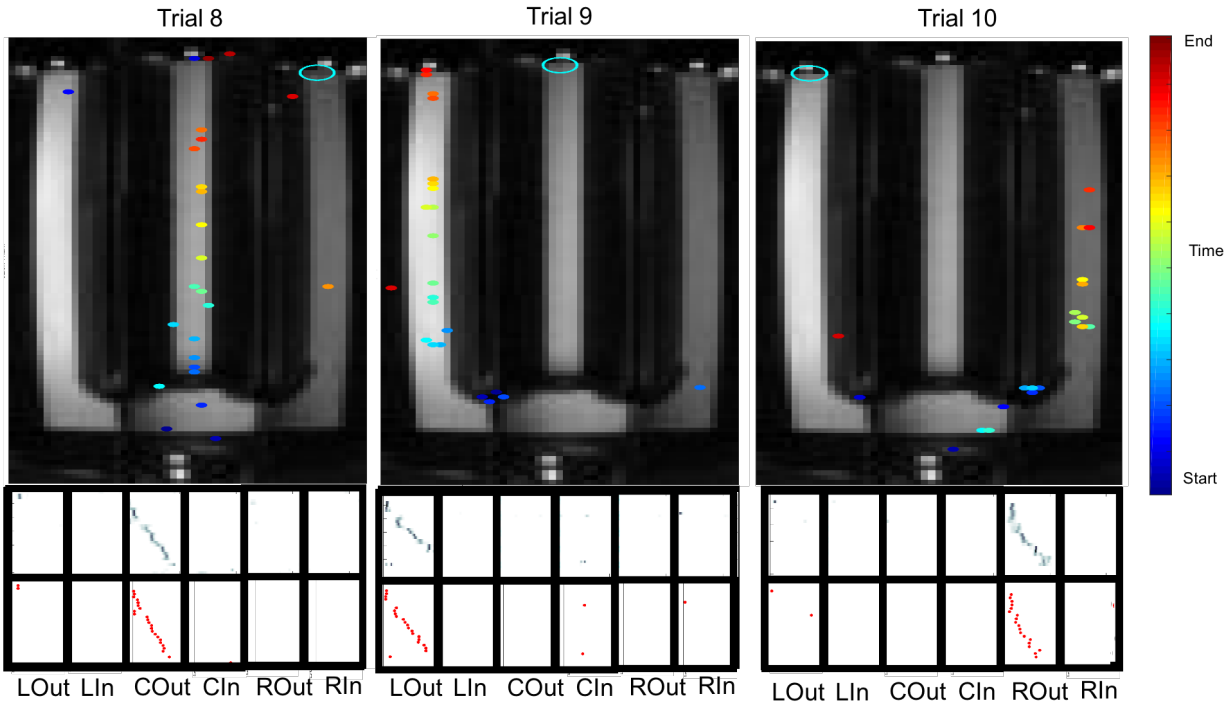

Figure 3.8: Decoded SWRs during behavior in the TADS task

area and proceeding outbound down the center arm (Fig 3.8 left) . After finishing reward consumption, the rat then runs back to the delay area, completes another eight second wheel run, and after the doors open makes a correct choice to go down the center arm for water reward on trial 9. While drinking during trial 9 another SWR is decoded showing activity originating at the delay area and this time proceeding outbound down the left arm (Fig 3.8 center). Subsequently the rat returns to the delay area, completes the wheel run and makes another correct choice proceeding down the left arm for trial 10. Finally, during trial 10 there is another SWR detected during reward consumption. The decoded SWR shows a trajectory starting at the delay area and proceeding outbound down the right arm (Fig 3.8 right). On the subsequent trial 11 after wheel running the rat proceeded down the right arm making an error and received no reward.

When the rat is located in the left or right arms it is difficult to interpret if remote trajectories are replaying the past or future; however, when the rat is in the center arm there is a clear separation between the past and future choices. With this is mind we asked: for

each decoded trajectory event while the rat was located in the center arm, how often did the trajectory match the future behavioral choice of the rat on the next trial? We found that 87% of the decoded trajectories while the rat was located at the center arm predicted the future choice of the rat on the next trial (All Rats = 87%, 49/56 trajectories, $p = 5 * 10^{-10}$; Rat 1 = 87%, 21/24, $p = 1.8 * 10^{-5}$; Rat 2 = 80%, 8/10, p = 0.0107; Rat 3 = 90%, 20/22, $p = 5.48 * 10^{-6}$). The prediction of future choice was significant for a wide range of ripple filtering parameters (Fig. 3.9). The data therefore suggests that the content of SWR events during reward consumption at the center arm reward port may represent future plans for the next trial.

3.4 Discussion

Using Bayesian decoding of SWR activity during the TADS task, we found several of the SWRs replayed spatiotemporally remote trajectories in the maze. These results give the first evidence of SWR content playing future planned trajectories in space when the rat has two different future goals to choose from. Unlike previous reports of replay during SWRs, these events started in a spatially remote but behaviorally relevant location and correlated with the rats' future behavior. During all decoded SWRs, the rat's head direction was well controlled as they were drinking water from the reward port. Unlike previous reports of correlation between decoded SWR trajectories and future behavior, the trajectory being played in these SWRs would match behavior carried out in the remote future after the rat returned to the delay period and completed another eight second wheel run. For this reason, every single remote replay played a prospective behavior that was at least twelve seconds in the future and sometimes much more.

In the previous chapter, it was shown that in order to complete the TADS the rat must internally maintain a contextual difference between the two types of center arm trials. However, when we looked at internally generated activity during the delay period wheel running after each center arm trial we found no difference. Here, by decoding the content of SWRs

Figure 3.9: Performance for predicting the future choice using decoded trajectory events when the rat is at the center arm reward port

Exploration of performance across different parameters for the clean decoded event criteria, using all recordings from all rats combined.
Y axis shows adjustments to the number of frames that must have a value greater than .01 (from greater than 2 (so 3 frames) to greater than 8 (so 9 frames)). X axis shows adjustments to the correlation value that all decoded events must be larger than (from 0.5 to 0.95).
Note as the criteria is reduced to accept what is most likely a large number of false positive events (top left), the performance approaches chance 50%. While at the other end when using only detected events of the highest quality to make predictions performance is highest (88 %)

during reward consumption at the center reward port, another form of internally generated activity, we found evidence for distinct contexts on different center arm trials; as the rat showed evidence of planning for the different future trajectories on the next trial. In conclusion, the decoded content of SWRs during reward consumption showed evidence of planned trajectories for the next trial, and in doing so provided evidence for the rats' having different internally generated activity between the two types of contextually different center arm trials.

CHAPTER 4

REPLAY OF CONTEXT IN THE CONTENT OF SWRS DURING TADS

4.1 Introduction

In the last two chapters, we looked at contextual coding in the hippocampus during the TADS task. In both cases, contextual coding meant looking at times when the rat was repeating the same behavior while activity generated in the hippocampus showed a difference depending on the surrounding context. First the same wheel run behavior showed different neural activity based on the context of which of the three arms was last visited. Second, during pauses for drinking water at the center arm reward port, decoding of SWR content revealed different remote trajectories depending on the context of where the rat planned to go for the next trial. Both cases support the idea that the hippocampus plays an important role in the maintenance of context. In this chapter, we will go over evidence that replay of internally generated sequences can retain episodic specificity by recalling contextually specific versions of a behavior. While there have been demonstrations of replay of spatial and non-spatial sequences, there has been no demonstration of replay of contextually specific non-spatial memory.

The first demonstration of replay of internally generated non-spatial sequences was in Wang et al 2016. In this paper replay was detected using the common template method. Briefly the mean responses of episode fields across all wheel runs where used to create an ordered sequence of firing by arranging them by latency to peak activity across neurons. Subsequently, in identified periods of SWRs, based on ripple power and the envelope of the reversal of the LFP signal across the pyramidal layer of the hippocampus, the same sequential ordering of neurons was searched for using the template constructed from the wheel run activity ((Wang et al., 2016) Fig 4.1 b, c) . If there was a sequence found with at least five neurons firing in the same order as the template sequence, then the probability

of such a sequence being seen by chance was estimated by a shuffling procedure. Based on this measure Wang et al found significantly higher than chance events that displayed replay of the internally generated sequence during wheel running. The second report of replay of internally generated activity was from Belchior et al 2018. As previously seen in the Eichenbaum lab paradigm using treadmill running between an alternation task (Fig 4.0 d), they were also able identify a clear sequence of internally generated activity and to create a template and search for replays ((Belchior et al., 2018) Fig 4.1 e, f). Both papers found replay of internally generated sequences using a template method. In our recordings, we have a considerable larger number of neurons recorded which permits searching for replay using a Bayesian decoder. Furthermore, we observed three contextually specific versions of internally generated activity during wheel runs. Thus, we looked at decoded SWR events using a Bayesian decoder and the probability maps for the three contextually specific wheel runs.

4.2 Methods

The methods are similar to those covered in chapter 3. However, in this case neural activity during SWR events was decoded using probability maps constructed from the three different wheel runs instead of the 6 inbound and outbound arm runs. The wheel run maps were created by binning the wheel run time period into 125ms bins. Then the activity was smoothed using a Gaussian with a 100ms SD.

As in chapter three we decoded the PBE activity across a vector that was the concatenated probability maps from all three wheel runs. After decoding such that the likelihood for each time frame gave a probability across all possible time bins from all three wheel runs that sums to 1 we then split up the maps. Then for each map frame the peak likelihood was identified and maintained if the probability was above 0.1. Then if the remaining decoded activity for each map had at least 5 frames of decoded activity above 0.1 probability then the correlation was calculated for those points to look for a smooth trajectory across time

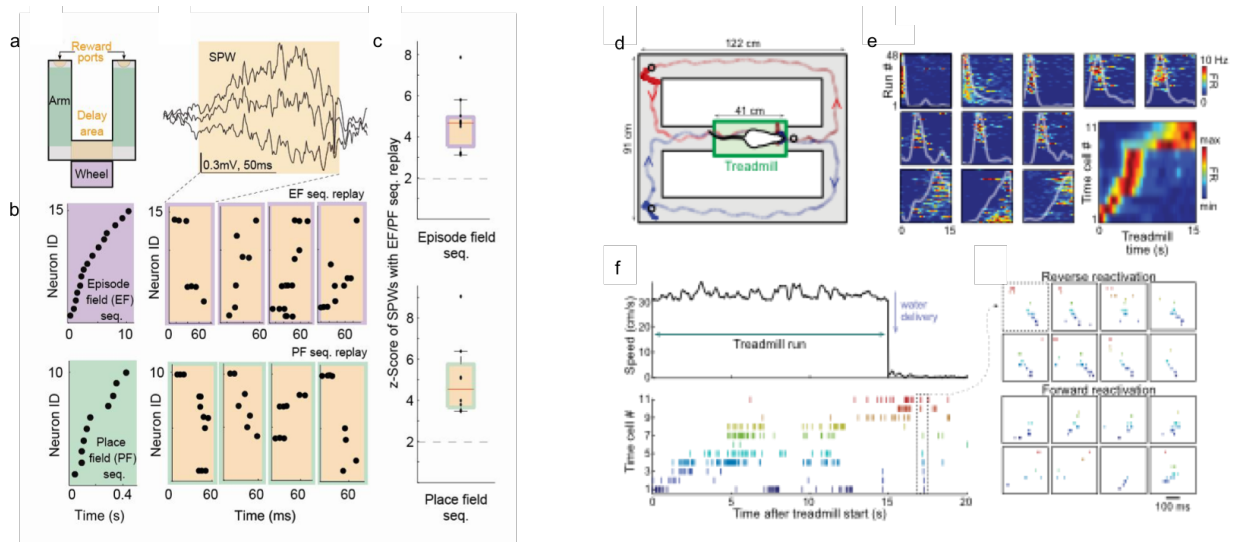

Figure 4.1: Detection of replay of internally generates sequences

a) From Wang et al 2016 a cartoon diagram of the arm used for delayed alternation with wheel running.

b) Example of a template constructed from location of the peak firing of 15 episode cells ordered by latency to peak for averages of all wheel runs (10 seconds long). To the right are examples of replay of neurons in the same order as during wheel running detected for SWR events (projection above shows a SWR event detected in the LFP)

c) Z-score of SPWs with episode field replay.

d) From Belchior et al 2018 shows a cartoon diagram of the alternation task with treadmill running in the common arm.

e) 11 time cells with their firing activity for 48 trials as well as the ordered sequence across treadmill running for all 11 in the bottom right.

f) An example of the firing of time cells during a 15 second treadmill run followed by their ordered replay in a SWR detected right after the treadmill run when the rat is given a water reward. Examples of forward and reverse replay were detected.

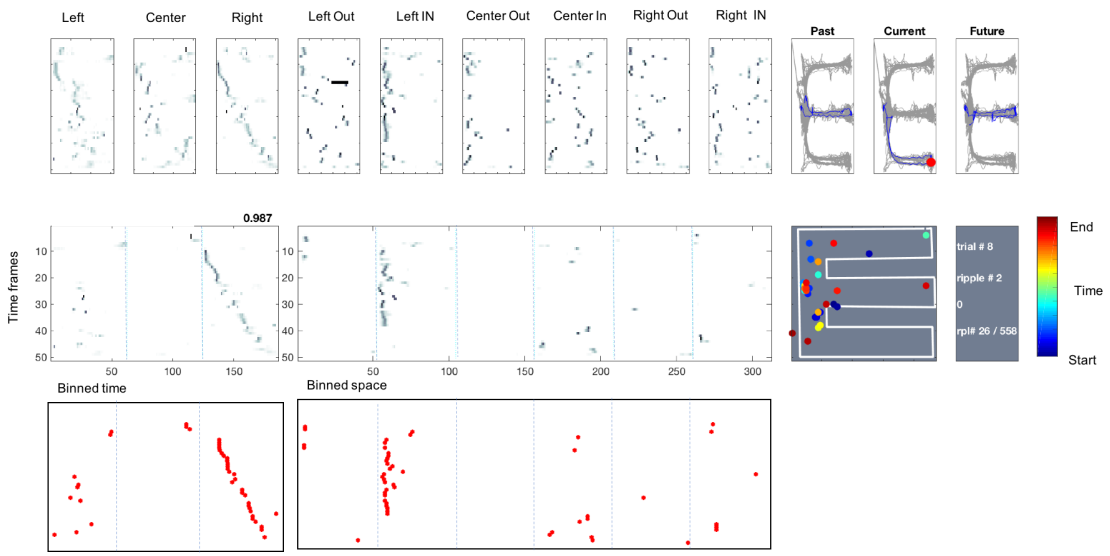

Figure 4.2: Decoded SWRs show replay of contextually specific wheel runs 1
The same format as 3.7 is used with the addition of three maps constructed from the three contextually distinct wheel runs. Note while the output of the decoded looks like noise on the maze maps (2d and linear maps) it shows beautiful structure replaying the entire sequence of running in the wheel after visiting the right arm.

and time bins. Replays with a correlation above 0.9 were considered wheel run trajectory events.

4.3 Results

Using a Bayesian decoder, we now looked at whether PBEs observed during the TADS task showed replay of the internally generated wheel running activity. However, in our case, unlike previous reports of replay of internally generated episode or time cell sequences we looked if there was replay of contextually specific internally generated population sequences. We found clear examples of replay of context specific sequences (Fig 4.2,4.3). This is the first evidence that the hippocampus can replay specific contextually generated behavioral sequences.

In some rare examples, we could see PBEs that when decoded showed complex trajectories that mapped onto both the maze and wheel in order (Fig 4.4 , 4.5 see arrows). These

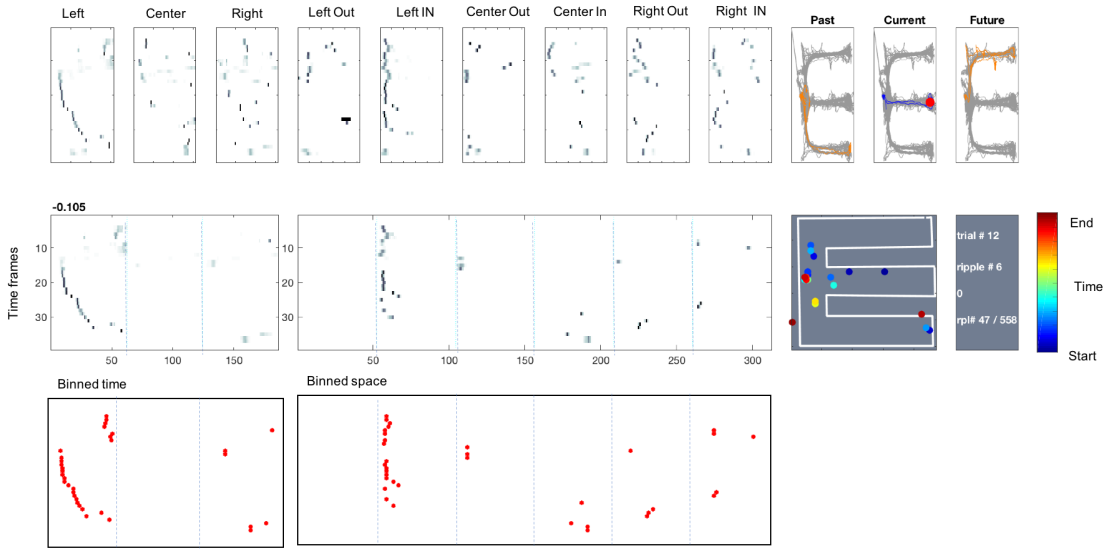

Figure 4.3: Decoded SWRs show replay of contextually specific wheel runs2
The same format as 3.7 is used with the addition of three maps constructed from the three contextually distinct wheel runs. Note while the output of the decoded looks like noise on the maze maps (2d and linear maps) it shows beautiful structure replaying the entire sequence of running in the wheel after visiting the right arm.

examples appeared to suggest that beyond just replaying contextually specific sequences, these replays as well might correlate with the future planned behavior of the rat.

It is very interesting to note that can be interpreted as the hippocampus encoding a pattern in the future that will specifically be modulated by the past experience (or current location during the ripple). To see if the contextually specific replay events were correlated with the future behavior of the rat (determined by the current location during the ripple) we asked for each event whether the activity matched the current arm location, thus agreeing with the activity expected in the future wheel run. The ripples show a 51% chance of playing the contextually specific wheel run sequence expected in the immediate future, which was significantly greater than expected by chance given the possibility of three different wheel run sequences (All rats = 51%, 83/164 trajectories, p = $1.1 * 10^{-6}$; Rat 1 = 51%, 41/80, p = $2.4 * 10^{-4}$; Rat 2 = 47%, 9/19, p = 0.061; Rat 3 = 50%, 33/65, p = 0.001). This result was also robust to a range of ripple event filter parameters (Fig 4.6).

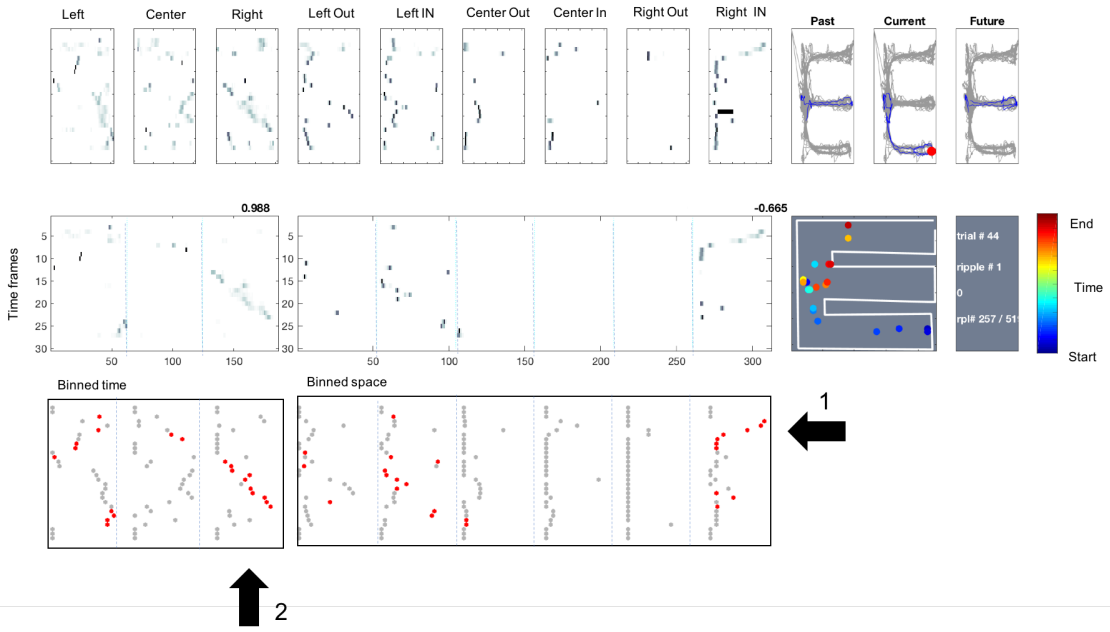

Figure 4.4: Decoded SWRs show replay of contextually specific wheel runs as well as maze trajectory 1

The same format as 4.3 In this particular event the decoded activity suggests a unified trajectory that starts at the location of the rat (at the right arm reward port) and moves inbound (playing out down the Right inbound map – see arrow 1) followed by a transition into playing the wheel sequence that is contextually specific for wheel runs after a visit to the right arm (arrow 2).

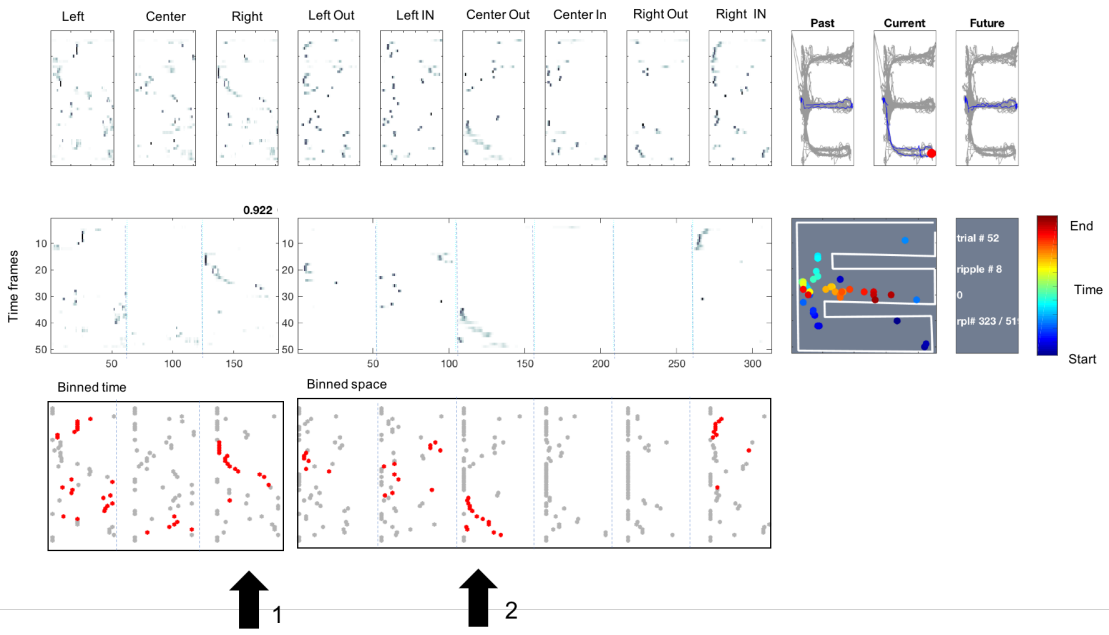

Figure 4.5: Decoded SWRs show replay of contextually specific wheel runs as well as maze trajectory 2

The same format as 4.3 In this particular event the decoded activity suggests a unified trajectory that starts with replay of the contextually specific wheel run seen after visiting the right arm (arrow 1) , followed by a trajectory that moves down the center outbound portion of the maze (arrow 2).

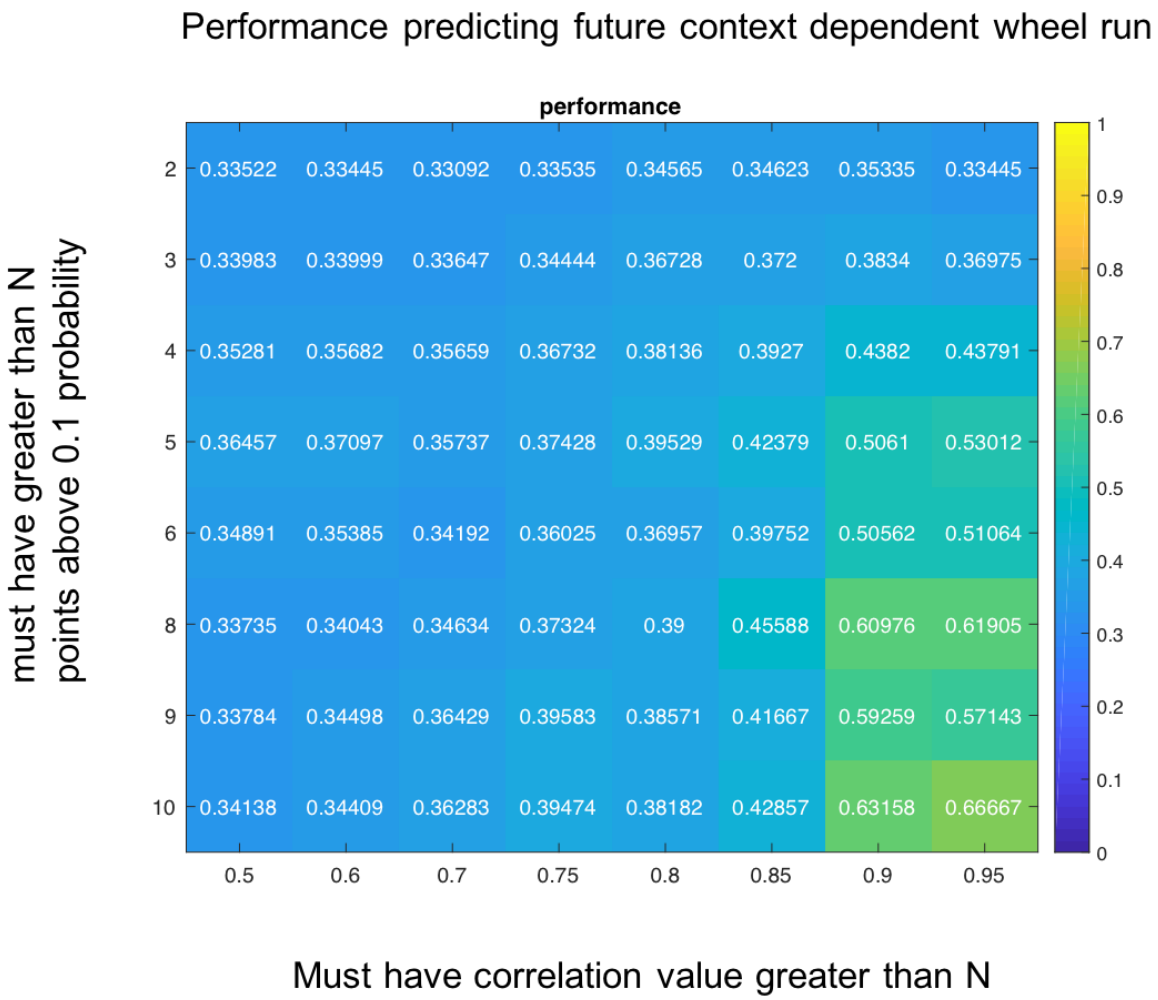

Figure 4.6: Performance predicting current location of the rat based on the context within decoded replays of contextually specific wheel runs.
As in 3.9 Note in this case chance is 33%. The top left approached chance while more selective criteria in the bottom right shows 67% prediction performance.

4.4 Discussion

In this chapter, we provide two novel findings. First we show replay of non-spatial internally generated sequences using a large population of neurons and a Bayesian decoding method. Second we show the first demonstration of replay of contextually specific internally generated sequences. Replay of sequences with specific contextual information provides new support for the maintenance of context within recall of experiences during replay. Furthermore, the time of contextually specific replay content with respect to the overall behavioral task suggest that these replays also are consistent with the data in chapter three showing that replay matches the future behavior and experience of the rat.

CHAPTER 5

CONCLUSION AND DISCUSSION

5.1 Context dependent activity during delay period wheel running

The work in this thesis started with the aim to test the current theory that context dependent activity observed in hippocampus across delay periods in a delayed spatial memory task is a contextual memory being maintained for the purpose of task performance. The TADS task is designed to create a situation where the rat needs to distinguish the different delay periods following center arm choices under different contexts, one proceeding to the left and another to the right. We found that although the hippocampus did display context specific internally generated sequences across different delay periods, there was only a contextual distinction between delay periods following the three spatial arm choices. Thus, there was no detectable information in hippocampal activity that could inform decisions following center arm choices.

In the discussion of chapter two we addressed possible interpretations with respect to the seemingly conflicting presence of contextual coding for the past arm visited, but not for different versions of visiting the center arm. One important lesson to take away from this result is the use of extremely simple behavioral paradigms in neuroscience, while necessary for several reasons, may have the unfortunate consequence of leading to overly simplified theories of function. Therefore it might not be surprising that as more complex behaviors are studied in systems neuroscience, there will be fewer examples of information in single regions of the brain being sufficient to understand cognition in the task. As a result, it will become especially important to try and understand the core functions that each area in the brain naturally provides to the animal in order to hopefully understand how they orchestrate more complex behaviors across regions.

Another important thought in light of the results in chapter two relates to the attempt to study memory in the brain. A large amount of literature has observed differences in persistent

activity during working memory tasks (Funahashi et al., 1989; Goldman-Rakic, 1995), but studying memory on longer timescales such as what is required in TADS, might force us to reconsider the mechanisms that might be at play for memory storage. For instance, if we assume that the information to distinguish the two center arm trials is held in synaptic weights, perhaps it is more beneficial to keep activity of the neurons holding the memory silent and only active when needed for readout. Under this hypothetical situation it would be extremely difficult to observe evidence for the maintenance of the contextual memory using extracellular recordings because there may not be a fixed point in time when the information is read out from the brain. Furthermore we have no idea the relevant timescale required for reading out such information. This highlights another major hurdle that will become important in systems neuroscience, which is designing tasks that constrain the time when mental computations must be made during complex behaviors (Mountcastle et al., 1990). Hopefully combining such well designed tasks with temporally specific perturbations using methods such as optogenetics, will help us tease apart the functional role of neural activity with respect to specific cognitive functions and behaviors.

5.2 Spatiotemporally remote planning in the content of decoded SWRs

Decoding the content of SWRs during reward consumption in the TADS task revealed strong evidence for a correlation between spatiotemporally remote trajectories and the future behavior of the rat. In the discussion of chapter three, we touched on how these results provided a glimpse into the internal representation of the two contextually different types of center arm trials, something we were unable to detect in activity across the delay.

The generation of differential sequences during SWRs but not during internally generated activity across the wheel run begs the question of what if any difference is there in the source of activation. It is interesting to note closed loop perturbations of ripple activity during a

non-delayed version of the three arm task showed a deficit specifically for outbound trials (Jadhav et al., 2012). It would be interesting to see if a similar perturbation in the TADS task would maintain outbound trial specificity or instead impair all trial types. It may be possible that hippocampus is important for delayed spatial memory tasks not for the function of holding a memory across the delay, but rather because of the need to plan spatial trajectories in the more remote future. For instance could the planned trajectories playing out in the hippocampus be the transfer of a plan to another area for storage across the subsequent delay (Yoon et al., 2008)?

5.3 What function does hippocampus provide in the TADS task?

It is well known that hippocampus becomes necessary in spatial memory tasks specifically when a delay is added. Thus it was expected that activity across the delay in hippocampus would be of particular importance. The results in this thesis do support the hypothesis that the hippocampus contains information about the most recent spatial arm choice. However, successful behavior in the TADS requires a memory for context within the larger behavioral sequence in order to differentiate different visits to the center arm. Moving forward it would be ideal to identify the area that is responsible for maintaining the contextual memory for the center outbound trials. Perturbation of this area while leaving hippocampus intact would be expected to create an increase specifically in outbound errors. Alternatively it would be predicted that impairing hippocampus would specifically lead to just as many repeat errors as other error types (Yoon et al., 2008).

REFERENCES

James A. Ainge, Matthijs A.A. van der Meer, Rosamund F. Langston, and Emma R. Wood. Exploring the role of context-dependent hippocampal activity in spatial alternation behavior. *Hippocampus*, 17(10):988–1002, 2007. ISSN 1098-1063. doi: 10.1002/hipo.20301. URL `http://dx.doi.org/10.1002/hipo.20301`.

Per Andersen, Richard Morris, David Amaral, John O'Keefe, and Tim Bliss. *The hippocampus book*. Oxford university press, 2007.

Hindiael A. Belchior, Rodrigo Pavao, Alan M.B. Furtunato, Howard Eichenbaum, and Adriano B.L. Tort. Reactivation of time cell sequences in the hippocampus. *bioRxiv*, 2018. doi: 10.1101/389874. URL `https://www.biorxiv.org/content/early/2018/08/10/389874`.

Michael S Bienkowski, Ian Bowman, Monica Y Song, Lin Gou, Tyler Ard, Kaelan Cotter, Muye Zhu, Nora L Benavidez, Seita Yamashita, Jaspar Abu-Jaber, et al. Integration of gene expression and brain-wide connectivity reveals the multiscale organization of mouse hippocampal networks. *Nature neuroscience*, page 1, 2018.

Katie C Bittner, Christine Grienberger, Sachin P Vaidya, Aaron D Milstein, John J Macklin, Junghyup Suh, Susumu Tonegawa, and Jeffrey C Magee. Conjunctive input processing drives feature selectivity in hippocampal ca1 neurons. *Nature neuroscience*, 18(8):1133, 2015.

M. R. Bower. Sequential-context-dependent hippocampal activity is not necessary to learn sequences with repeated elements. *Journal of Neuroscience*, 25(6):1313–1323, Feb 2005. ISSN 1529-2401. doi: 10.1523/jneurosci.2901-04.2005. URL `http://dx.doi.org/10.1523/JNEUROSCI.2901-04.2005`.

György Buzsáki. Theta oscillations in the hippocampus. *Neuron*, 33(3):325 – 340, 2002. ISSN 0896-6273. doi: https://doi.org/10.1016/S0896-6273(02)00586-X. URL `http://www.sciencedirect.com/science/article/pii/S089662730200586X`.

György Buzsáki. Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus*, 25(10):1073–1188, 2015.

György Buzsáki and David Tingley. Space and time: The hippocampus as a sequence generator. *Trends in Cognitive Sciences*, 22(10):853 – 869, 2018. ISSN 1364-6613. doi: https://doi.org/10.1016/j.tics.2018.07.006. URL `http://www.sciencedirect.com/science/article/pii/S1364661318301669`. Special Issue: Time in the Brain.

Nicola S Clayton, Lucie H Salwiczek, and Anthony Dickinson. Episodic memory. *Current Biology*, 17(6):R189–R191, 2007.

Jeremy D Cohen, Mark Bolstad, and Albert K Lee. Experience-dependent shaping of hippocampal ca1 intracellular activity in novel and familiar environments. *Elife*, 6:e23040, 2017.

Thomas J. Davidson, Fabian Kloosterman, and Matthew A. Wilson. Hippocampal replay of extended experience. *Neuron*, 63(4):497–507, Aug 2009. ISSN 0896-6273. doi: 10.1016/j.neuron.2009.07.027. URL http://dx.doi.org/10.1016/j.neuron.2009.07.027.

Sam A Deadwyler, Terence Bunn, and Robert E Hampson. Hippocampal ensemble activity during spatial delayed-nonmatch-to-sample performance in rats. *Journal of Neuroscience*, 16(1):354–372, 1996.

Kamran Diba and György Buzsáki. Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience*, 10(10):1241–1242, Sep 2007. ISSN 1546-1726. doi: 10.1038/nn1961. URL http://dx.doi.org/10.1038/nn1961.

Daniel A Dombeck, Christopher D Harvey, Lin Tian, Loren L Looger, and David W Tank. Functional imaging of hippocampal place cells at cellular resolution during virtual navigation. *Nature neuroscience*, 13(11):1433, 2010.

Paul A Dudchenko. An overview of the tasks used to test working memory in rodents. *Neuroscience & Biobehavioral Reviews*, 28(7):699–709, 2004.

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

Norbert J Fortin, Kara L Agster, and Howard B Eichenbaum. Critical role of the hippocampus in memory for sequences of events. *Nature neuroscience*, 5(5):458, 2002.

Loren M Frank, Emery N Brown, and Matthew Wilson. Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron*, 27(1):169–178, 2000.

Shigeyoshi Fujisawa, Asohan Amarasingham, Matthew T Harrison, and György Buzsáki. Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nature neuroscience*, 11(7):823, 2008.

S. Funahashi, C. J. Bruce, and P. S. Goldman-Rakic. Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *Journal of Neurophysiology*, 61(2):331–349, Feb 1989. ISSN 1522-1598. doi: 10.1152/jn.1989.61.2.331. URL http://dx.doi.org/10.1152/jn.1989.61.2.331.

Gabrielle Girardeau, Karim Benchenane, Sidney I Wiener, György Buzsáki, and Michaël B Zugaro. Selective suppression of hippocampal ripples impairs spatial memory. *Nature Neuroscience*, 12(10):1222–1223, Sep 2009. ISSN 1546-1726. doi: 10.1038/nn.2384. URL http://dx.doi.org/10.1038/nn.2384.

P.S Goldman-Rakic. Cellular basis of working memory. *Neuron*, 14(3):477 – 485, 1995. ISSN 0896-6273. doi: https://doi.org/10.1016/0896-6273(95)90304-6. URL http://www.sciencedirect.com/science/article/pii/0896627395903046.

Thomas Hainmueller and Marlene Bartos. Parallel emergence of stable and dynamic memory engrams in the hippocampus. *Nature*, page 1, 2018.

Robert E Hampson, John D Simeral, and Sam A Deadwyler. Distribution of spatial and nonspatial information in dorsal hippocampus. *Nature*, 402(6762):610, 1999.

Demis Hassabis and Eleanor A Maguire. Deconstructing episodic memory with construction. *Trends in cognitive sciences*, 11(7):299–306, 2007.

Lynn Hazan, Michaël Zugaro, and György Buzsáki. Klusters, neuroscope, ndmanager: A free software suite for neurophysiological data processing and visualization. *Journal of Neuroscience Methods*, 155(2):207–216, Sep 2006. ISSN 0165-0270. doi: 10.1016/j.jneumeth.2006.01.017. URL `http://dx.doi.org/10.1016/j.jneumeth.2006.01.017`.

Frederick L Hitti and Steven A Siegelbaum. The hippocampal ca2 region is essential for social memory. *Nature*, 508(7494):88, 2014.

S. P. Jadhav, C. Kemere, P. W. German, and L. M. Frank. Awake hippocampal sharp-wave ripples support spatial memory. *Science*, 336(6087):1454–1458, May 2012. ISSN 1095-9203. doi: 10.1126/science.1217230. URL `http://dx.doi.org/10.1126/science.1217230`.

Leonard E Jarrard. On the role of the hippocampus in learning and memory in the rat. *Behavioral and neural biology*, 60(1):9–26, 1993.

Mattias P Karlsson and Loren M Frank. Awake replay of remote experiences in the hippocampus. *Nature Neuroscience*, 12(7):913–918, Jun 2009. ISSN 1546-1726. doi: 10.1038/nn.2344. URL `http://dx.doi.org/10.1038/nn.2344`.

Steve M. Kim and Loren M. Frank. Hippocampal lesions impair rapid learning of a continuous spatial alternation task. *PLoS ONE*, 4(5):e5494, May 2009. ISSN 1932-6203. doi: 10.1371/journal.pone.0005494. URL `http://dx.doi.org/10.1371/journal.pone.0005494`.

James J Knierim. Dynamic interactions between local surface cues, distal landmarks, and intrinsic circuitry in hippocampal place cells. *Journal of Neuroscience*, 22(14):6254–6264, 2002.

Keigo Kohara, Michele Pignatelli, Alexander J Rivest, Hae-Yoon Jung, Takashi Kitamura, Junghyup Suh, Dominic Frank, Koichiro Kajikawa, Nathan Mise, Yuichi Obata, et al. Cell type–specific genetic and optogenetic tools reveal hippocampal ca2 circuits. *Nature neuroscience*, 17(2):269, 2014.

Albert K. Lee and Matthew A. Wilson. Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron*, 36(6):1183–1194, Dec 2002. ISSN 0896-6273. doi: 10.1016/s0896-6273(02)01096-6. URL `http://dx.doi.org/10.1016/S0896-6273(02)01096-6`.

Doyun Lee, Bei-Jung Lin, and Albert K Lee. Hippocampal place fields emerge upon single-cell manipulation of excitability during behavior. *Science*, 337(6096):849–853, 2012.

Jill K Leutgeb, Stefan Leutgeb, Alessandro Treves, Retsina Meyer, Carol A Barnes, Bruce L McNaughton, May-Britt Moser, and Edvard I Moser. Progressive transformation of hippocampal neuronal representations in "morphed" environments. *Neuron*, 48(2):345–358, 2005.

S. Leutgeb. Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science*, 309(5734):619–623, Jul 2005. ISSN 1095-9203. doi: 10.1126/science.1114037. URL http://dx.doi.org/10.1126/science.1114037.

Nikos K Logothetis. Neural-event-triggered fmri of large-scale neural networks. *Current opinion in neurobiology*, 31:214–222, 2015.

Brian Lustig, Yingxue Wang, and Eva Pastalkova. Oscillatory patterns in hippocampus under light and deep isoflurane anesthesia closely mirror prominent brain states in awake animals. *Hippocampus*, 26(1):102–109, 2016.

Christopher J. MacDonald, Kyle Q. Lepage, Uri T. Eden, and Howard Eichenbaum. Hippocampal "time cells" bridge the gap in memory for discontiguous events. *Neuron*, 71(4):737–749, Aug 2011. ISSN 0896-6273. doi: 10.1016/j.neuron.2011.07.012. URL http://dx.doi.org/10.1016/j.neuron.2011.07.012.

Yasuhiro Mochizuki, Tomokatsu Onaga, Hideaki Shimazaki, Takeaki Shimokawa, Yasuhiro Tsubo, Rie Kimura, Akiko Saiki, Yutaka Sakai, Yoshikazu Isomura, Shigeyoshi Fujisawa, et al. Similarity in neuronal firing regimes across mammalian species. *Journal of Neuroscience*, 36(21):5736–5747, 2016.

VB Mountcastle, MA Steinmetz, and R Romo. Frequency discrimination in the sense of flutter: psychophysical measurements correlated with postcentral events in behaving monkeys. *The Journal of Neuroscience*, 10(9):3032–3044, Sep 1990. ISSN 1529-2401. doi: 10.1523/jneurosci.10-09-03032.1990. URL http://dx.doi.org/10.1523/JNEUROSCI.10-09-03032.1990.

Zoltán Nádasdy, Hajime Hirase, András Czurkó, Jozsef Csicsvari, and György Buzsáki. Replay and time compression of recurring spike sequences in the hippocampus. *Journal of Neuroscience*, 19(21):9497–9507, 1999.

E. Pastalkova, V. Itskov, A. Amarasingham, and G. Buzsaki. Internally generated cell assembly sequences in the rat hippocampus. *Science*, 321(5894):1322–1327, Sep 2008. ISSN 1095-9203. doi: 10.1126/science.1159775. URL http://dx.doi.org/10.1126/science.1159775.

Brad E. Pfeiffer and David J. Foster. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, 497(7447):74–79, Apr 2013. ISSN 1476-4687. doi: 10.1038/nature12112. URL http://dx.doi.org/10.1038/nature12112.

W. B. Scoville and B. Milner. Loss of recent memory after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgery and Psychiatry*, 20(1):11–21, Feb 1957. ISSN 0022-3050. doi: 10.1136/jnnp.20.1.11. URL http://dx.doi.org/10.1136/jnnp.20.1.11.

Annabelle C. Singer, Margaret F. Carr, Mattias P. Karlsson, and Loren M. Frank. Hippocampal swr activity predicts correct decisions during the initial learning of an alternation task. *Neuron*, 77(6):1163 – 1173, 2013. ISSN 0896-6273. doi: https://doi.org/10.1016/j.neuron.2013.01.027. URL http://www.sciencedirect.com/science/article/pii/S0896627313000937.

Eran Stark, Lisa Roux, Ronny Eichler, Yuta Senzai, Sebastien Royer, and György Buzsáki. Pyramidal cell-interneuron interactions underlie hippocampal ripple oscillations. *Neuron*, 83(2):467–480, 2014.

Bryan A Strange, Menno P Witter, Ed S Lein, and Edvard I Moser. Functional organization of the hippocampal longitudinal axis. *Nature Reviews Neuroscience*, 15(10):655, 2014.

Victoria L Templer and Robert R Hampton. Episodic memory in nonhuman animals. *Current Biology*, 23(17):R801–R806, 2013.

Endel Tulving et al. Episodic and semantic memory. *Organization of memory*, 1:381–403, 1972.

C.H Vanderwolf. Hippocampal electrical activity and voluntary movement in the rat. *Electroencephalography and Clinical Neurophysiology*, 26(4):407–418, Apr 1969. ISSN 0013-4694. doi: 10.1016/0013-4694(69)90092-3. URL `http://dx.doi.org/10.1016/0013-4694(69)90092-3`.

Yingxue Wang, Sandro Romani, Brian Lustig, Anthony Leonardo, and Eva Pastalkova. Theta sequences are essential for internally generated hippocampal firing fields. *Nature neuroscience*, 18(2):282, 2015.

Yingxue Wang, Zachary Roth, and Eva Pastalkova. Synchronized excitability in a network enables generation of internal neuronal sequences. *Elife*, 5:e20697, 2016.

IQ Whishaw and C Hippocampal Vanderwolf. Hippocampal eeg and behavior: change in amplitude and frequency of rsa (theta rhythm) associated with spontaneous and learned movement patterns in rats and cats. *Behavioral biology*, 8(4):461–484, 1973.

Tom J Wills, Colin Lever, Francesca Cacucci, Neil Burgess, and John O'keefe. Attractor dynamics in the hippocampal representation of the local environment. *Science*, 308(5723):873–876, 2005.

M. Wilson and B. McNaughton. Dynamics of the hippocampal ensemble code for space. *Science*, 261(5124):1055–1058, Aug 1993. ISSN 1095-9203. doi: 10.1126/science.8351520. URL `http://dx.doi.org/10.1126/science.8351520`.

Emma R. Wood, Paul A. Dudchenko, R.Jonathan Robitsek, and Howard Eichenbaum. Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron*, 27(3):623–633, Sep 2000. ISSN 0896-6273. doi: 10.1016/s0896-6273(00)00071-4. URL `http://dx.doi.org/10.1016/S0896-6273(00)00071-4`.

T. Yoon, J. Okada, M. W. Jung, and J. J. Kim. Prefrontal cortex and hippocampus subserve different components of working memory in rats. *Learning and Memory*, 15(3):97–105, Feb 2008. ISSN 1072-0502. doi: 10.1101/lm.850808. URL `http://dx.doi.org/10.1101/lm.850808`.

Yaniv Ziv, Laurie D Burns, Eric D Cocker, Elizabeth O Hamel, Kunal K Ghosh, Lacey J Kitch, Abbas El Gamal, and Mark J Schnitzer. Long-term dynamics of ca1 hippocampal place codes. *Nature neuroscience*, 16(3):264, 2013.