Monkey Hippocampal Neurons Track Ongoing Experiences

Yoni (Jonathan) Ivens Browning

A dissertation

submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

University of Washington

2021

Reading Committee:

Elizabeth Buffalo, Chair

Adrienne Fairhall, Co-Chair

Sheri Mizumori

Eric Shea Brown

Program Authorized to Offer Degree:

Neuroscience

University of Washington

Abstract

Monkey Hippocampal Neurons Track Ongoing Experiences

Yoni (Jonathan) Ivens Browning

Chairs of Supervisory Committee:

Elizabeth Buffalo

Adrienne Fairhall

Physiology and Biophysics

The hippocampus and surrounding structures of the medial temporal lobe are necessary for the formation of new episodic memories, that is, memories of experiences. It remains unclear, however, how neural activity within the hippocampus supports this mnemonic function. Decades of recordings from the rodent hippocampus have documented neural activity that tracks the animal's location in space, though it has yet to be determined if or how such navigational correlates might contribute to memory. An increasing body of work suggests that the precise navigational readout seen in rodent work may not be conserved across species. Navigational correlates in the monkey, for example, can be seen under the right conditions but often either offer poor spatial resolution or are better described by other features of a behavioral task. Together with results showing coding outside the spatial domain in rodents and other species, these data suggest hippocampal neurons reliably respond as an animal progresses through an orderly task, suggesting that salient task events are sufficient to structure hippocampal activity. This raises the compelling possibility that the ongoing activity in the hippocampus functions to link salient events into a continuous, unitary episode.

To characterize task-relevant activity in the monkey hippocampus, I trained monkeys to perform a variety of tasks in virtual reality while using chronic recording techniques to record from their

hippocampus. In my primary experiment, monkeys performed a spatial delayed alternation task in virtual reality. I used a chronically implanted hyperdrive targeting the medial temporal lobe with 124 independently movable single wire electrodes to record neural activity from thousands of cells over hundreds of behavioral sessions. I find neurons with significantly reliable and informative about the animal's position in virtual space. However, the majority of these spatially correlated neurons cluster around the key events in the maze, including the start zone, the choice point and the goals. In fact, the population activity in this task appears to be better described in terms of maze segments or task phase than it is in terms of spatial location. Some of these neurons maintain their task-specific responses irrespective of the visual cues defining the virtual environment, while other neurons "remap" when the monkey performs the same task in a visually distinct context. The stability of activity across mazes suggest that the hippocampus may abstract task structure across individual exemplars. In preliminary data from other virtual reality tasks, we find additional examples of task relevant information encoded even in the absence of explicit spatial coding. Taken together, these findings emphasize the important role that task structure plays in shaping hippocampal activity, and are consistent with the theory that activity in the hippocampus continuously tracks task-relevant events, joining them into a cohesive episode.

Table of Contents

Table Of Figures

Acknowledgements

Clare: At the risk of soundly cliché- I am not sure how I would have made it this far without you. From taking care of me when I broke my hands to making sure I had my facts straight about interneurons, from late night writing to trips on our tandem. Grad school has been just one part of the life we have started to build together, but you have made both the good and the bad parts of it so much better. Thank you.

My Family: I am so lucky to be have and be near to such a supportive family, particularly my parents, John and Elizabeth, and my sister Aviva. Thanks to all of you for putting up with my endless talk about the hippocampus, for proof reading my grants late at night, and for always be ready to climb a mountain when I needed a break.

My Friends: There are far too many people to name here, but I could not have finished this without your support. Thanks especially to those of you in my graduate school cohort. Just in case you are bored enough to read my thesis; Thank you.

Beth: I thought I was going to be a theorist when I entered grad school, until I got it in my head that maybe lab work might be something I wanted to learn you took the time and the effort take me into her lab and train me on just about every aspect of being a scientist. Thank you.

Adrienne: Somehow my PhD when from retinal adaptation to monkey navigation, and you always made sure I know how to model it. Even when I concluded we really needed a detour into mouse visual cortex, you took the time to help me think it through. Thank you

Jon: This thesis is really a manifestation of teamwork; most of the work here is a close collaboration between us. To do this, though, you not only had to pull your weight on this project but also teach me how to hold up my end. Thank you.

Buffalo Lab: After I broke my hands and my kneecap, food deliveries from “Giuseppe” started showing up at our house. Just about everything I we pretty difficult for me right then, and the food made it that little bit easier. Giuseppe may have been good at video games, but I have a feeling it was you who ordered the food. Thank-you especially AJ and Brian for running so many of the day to day aspects of this project, and for making sure things stayed on track. Thank you to Ian, for figuring out how to store

all the data we generated. Thank you to Kelly for training me to train monkeys. Thank you to Sierra for keeping our virtual environments well decorated. And thank you to Megan, not only for keeping our lab running smoothly but for driving me to campus every day for several months after I broke my hands.

Fairhall Lab: Everyone in the Fairhall lab does slightly different work, and I have learned so much about it. It's amazing how often some knowledge picked up in a meeting about hydra or song birds ends up being useful in understanding the hippocampus. Thanks especially to Kenneth for repeatedly explaining GLMs to me until I finally understood.

My Committee: I think some people are scared of their committees, but somehow, I only ever got helpful and encouraging feedback from you. Thank you. Thanks also to my reading committee, for making sure this thesis isn't too terrible and for putting up with my horrendous spelling. And thanks especially to Eric, who was my undergraduate advisor, encouraged me to go to graduate school, and continues to give positive and insightful input on my work.

Thank you to also to the NIH, the NSF, the UW CNC, the WANPRC, the Simons foundation, and the McKnight foundation for funding this work and my education. Thanks for to the UW graduate program and community for letting me have this grad school opportunity, and for making sure I got through to the end.

1 Introduction

The hippocampus and surrounding cortex are necessary for the formation of new episodic memories (Buzsáki and Moser, 2013). It is unclear, however, how neural activity in the hippocampus supports this mnemonic function. Despite the profound memory deficits suffered by human patients with damage to the hippocampus (Corkin, 2013; Scoville and Milner, 1957), studies in monkeys have had remarkable difficulty in demonstrating the behavioral impact of hippocampal lesions in the monkey (Rueckemann and Buffalo, 2017). Perhaps as a consequence of this behavioral ambiguity, the field also lacks a clear understanding of behavioral features that can drive hippocampal activity in the monkey.

Research using rodents, on the other hand, has established a number of factors that can well predict the firing of individual cells in rats and mice. Chief amongst these is an animal's physical location in space; the firing of so-called "place cells," as well as other spatially correlated firing in the rodent medial temporal lobe, has led to the hypothesis that the rodent hippocampus may encode a "cognitive map" that can be flexibly used for navigation (O'Keefe and Nadel, 1978). It remains unknown how such a map might facilitate memory (Bellmund et al., 2018; Whittington et al., 2020). Further, neurophysiological studies in rodents have shown that the best known predictor of an hippocampal cell's firing is actually the firing of other hippocampal cells (Meshulam et al., 2017), indicating that understanding population dynamics within the hippocampus is crucial for understanding how individual cells support the structure's role in memory.

It is not clear to what extent these findings in rodents hold true in the primate. Under the right circumstances, monkey cells can show spatial firing (Courellis et al., 2019), but this activity may be convolved with or better described by the animal's view of its world (Feigenbaum and Rolls, 1991; Rolls and O'Mara, 1995). Further, many of the population-level predictors present in rodents (most notably, the presence of a consistent theta-band oscillation) have a very different phenotype in the monkey. This disconnect presents a challenge for the hippocampus field; although information gleaned from research with rodents is invaluable, current monkey studies present a counterexample to the assertion that all of the activity observed in rodents can be construed as general properties of the mammalian hippocampus. This challenge is compounded by an inadequate understanding of the monkey hippocampus; although it

is clear that the monkey hippocampus is very different from the rodent, it is not yet clear what drives or organizes monkey hippocampal activity. For example, where rodent work has progressed to a carful dissection of well-studied phenomena routinely observed across multiple labs, the relative scarcity of monkey studies means that this work often focuses on showing a phenomena exists before it can be intricately understood. Perhaps most notably, where it is well established that rodent hippocampal cells will reliably fire for an animal's location, the monkey field has not yet found a similarly reliable predictor of hippocampal activity that is routinely replicated across labs. Without such an established "ground-truth" in monkeys, the field is limited in its ability to perform constrained studies or manipulate population dynamics.

In this introduction, I will discuss the limited body of existing work conducted in the monkey hippocampus using electrophysiological techniques. I will discuss evidence that, as a whole, these data suggest that the hippocampus tracks task events as the animal experiences them. This can include tracking of continuous or spatial task variables, but extend more generally to a demarcation of salient task events like trial starts or reward. I will further suggest that eye movements provide insight into the temporal organization of monkey hippocampal activity. In the absence of a consistent theta-band oscillation in the monkey, this organization is potentially analogous to theta-modulation in rodents.

1.1 What is the function of the hippocampus?

Despite strong evidence linking the hippocampal formation to amnesia (Clark and Squire, 2010), the proximity and anatomical connectivity of the hippocampus to the limbic system (Gergen and MacLean, 1964; MacLean et al., 1968) and to olfactory areas (Green and Arduini, 1953; Yokota et al., 1967) led early physiology studies to focus on potential ties between the hippocampus and emotion and smell (reviewed in (Colgin, 2020)). This early work, however, found that it was difficult to drive hippocampal activity though simple sensory stimulation, particularly in the monkey. For example, "rhythmically slow activity" (what would now be termed a theta oscillation (Colgin, 2020)) could be elicited in the hippocampus of rabbits or cats through presentation of an olfactory, visual, auditory or tactile stimulus (Green and Arduini, 1953). In the monkey, however, simple sensory stimulation was insufficient to induce the same kind of activity; to elicit a similar oscillation, the monkey had to be presented with an unfamiliar stimulus. Later olfactory-focused work faced similar challenges. Though stimulation of the olfactory bulb

was sufficient to drive sub-threshold excitatory postsynaptic potentials in hippocampal neurons (Yokota et al., 1967), researchers had to undertake the heroic effort of intracellular recording from the hippocampus of awake squirrel monkeys in order to see this effect. In contrast, stimulation of the septum (Gergen and MacLean, 1964) proved to be a considerably more effective method of driving hippocampal activity.

These early hypotheses about hippocampal function were ultimately eclipsed by the discovery that bilateral lesions of hippocampus and surrounding medial temporal lobe structures lead to the inability to form certain kinds of new memories (Scoville and Milner, 1957). The most famous of these patients, H.M., underwent these lesions as part of a surgical attempt to stop otherwise intractable epilepsy. Other patients had lesions performed as an attempt by 1950's physicians to mitigate various "psychoses," an idea born out of the hypothesized role of the hippocampal formation in emotion (MacLean, 1952). H.M.'s surgery was successful in preventing further seizures, though temporal lobe resection was deemed (except in a few patients) to be ineffective in treating "psychoses." However, all patients tested after medial temporal lobe lesion suffered severe memory impairments. For example, following his surgery, H.M "could no longer recognize the hospital staff nor find his way to the bathroom, and he seemed to recall nothing of the day-to-day events of … life (Scoville and Milner, 1957)." Although H.M. could not remember certain elements of his past immediately before the surgery, by far the most profound deficit he and the other documented patients suffered was an inability to form memories of events after their operations. Common amongst the damage in all amnesic patients was bilateral damage to the hippocampal formation, leading to the conclusion that this structure was necessary for the formation of memories. Further work on H.M. and amnesic patients with other sources of hippocampal damage would go on to pinpoint the necessity of the hippocampal formation in episodic memory, that is, specific memory of events (Corkin, 2013).

In principle, the monkey provides a promising model for further studies of the neural underpinnings of memory. Though extremely valuable behaviorally, human work is limited in assessing the function of any specific structure for the simple reason that lesions are not strictly limited to any particular structure; even in Scoville and Millner's initial study, lesions in amnesiac patients (including HM) extend beyond the hippocampal formation into large sections of the medial temporal lobe (Annese et al., 2014). Monkey studies should provide the ability to target specific brain structures in a behavioral tractable animal model.

In practice, however, there has been confusion in the monkey lesion literature over the extent and type of hippocampal lesion used, and thus interpretability of any behavioral deficits (Squire et al., 2007). Different labs have adopted different approaches in making lesions that are difficult to compare. One lesion approach, for example, was to do a combined lesion of the fornix and amygdalofugal pathways together, destroying the main white matter innervating the hippocampus and amygdala. Separately severing either of these pathways, however, was found to be insufficient to produce behavioral deficits, challenging any study that attempted to pinpoint functional necessity on any one structure (Bachevalier et al., 1985). Similarly, though a number of behavioral deficits were achieved through relatively non-selective lesions that aspirated the hippocampus and some of the surrounding cortex, many of these same deficits could not be replicated in monkeys with more selective ibotenic acid lesions to the hippocampus and amygdala (Murray and Mishkin, 1998). In fact, damage to the perirhinal cortex – a structure often co-lesioned with the hippocampus – has very similar effects (in humans) to damage to the hippocampal formation (Buffalo et al., 1998) and even non-selective lesions that spared portions of the perirhinal cortex may have damaged the white matter innervating it (Murray and Mishkin, 1998). These results make it difficult to draw specific conclusions from literature aimed at attributing specific deficits to the hippocampus.

This uncertainty over lesion extent extends into an uncertainty about the hippocampal dependence of specific behaviors. This is confounded by difficulty in creating a task that requires “memory” in the monkey. Many studies have adopted a recognition memory task known as “Delayed Matching to Sample (DMS),” in which the monkeys see a cue (it may be an object, image, light at a particular location, etc.) and, after a delay, are required to report if a second cue matches the first (Squire et al., 2007). A complementary “Delayed Non-Matching to Sample (DNMS)” task requires the monkey to report when two cues do not match. Though lesion work initially pointed to deficits in monkeys with hippocampal lesions as this delay became long (on the order of minutes), monkeys with the more selective lesions did not show such deficits ((Murray and Mishkin, 1998) but see also (Zola et al., 2000)). This finding is further complicated by the observation that the monkeys trained pre-lesion were unimpaired in the DNMS task while those trained post-lesions suffered a deficit, suggesting task acquisition, rather than some specific feature of the delayed matching task, was interrupted by hippocampal damage. These negative results (though not without controversy, see (Clark and Squire, 2010; Squire et al., 2007)) are still important to

mention because many early studies of hippocampal electrophysiology were designed around the assertion that delayed matching tasks were specifically dependent on the hippocampus (Cahusac et al., 1989; Miyashita et al., 1989; Opris et al., 2015; Riches et al., 1991; Rolls et al., 1989; Vidyasagar et al., 1991; Watanabe and Niki, 1985; Wilson et al., 1990). Though this assertion appears untrue, these studies did find task-related hippocampal activity, and I will describe in the following section what these data might tell us about understanding hippocampal function in light of the fact that these signals are not necessarily essential for solving this task.

Despite this confusion, it must be noted that studies have identified deficits in monkeys with selective hippocampal damage (Rueckemann and Buffalo, 2017). First, deficits in task acquisition noted in the delayed non-matching task (Murray and Mishkin, 1998) were corroborated in other DNMS paradigms (Baxter and Murray, 2001; Beason-Held et al., 1999) as well as in tasks involving object-place associations (Belcher et al., 2006) and learned discrimination (Teng et al., 2000). These results point to the observation that it is task acquisition itself, rather than some feature of the task, that may be impaired by hippocampal damage (Rueckemann and Buffalo, 2017). Further, transient inactivation of the hippocampus (which eliminates concerns about hippocampal lesions disrupting task acquisition) have shown that monkeys are unable to remember recently visited reward boxes in a cue-poor, non-navigational foraging task (Forcelli et al., 2014). This result is dependent on the delay between trials, pairing with other data from lesion studies in infant monkeys to suggest that the hippocampus is necessary for “holding” task information across a delay (Rueckemann and Buffalo, 2017). However, performance on the box-memory task can be rescued during hippocampal inactivation by coloring the reward boxes, implying that the role of the hippocampus in buffering task variables and relationships could be mitigated by increasing the information available to solve the task.

What, then, is the function of the hippocampus? Despite the profound losses seen in human amnesic patients, the difficulty seen in monkey studies with respect to pinpointing specific hippocampal deficits suggests an as yet incomplete understanding of what specific features of behavior are impacted by hippocampal damage. In the subsequent section I will move away from causal lesion studies to discuss the behavioral correlates of hippocampal neural activity in the monkey, and what these correlates might tell us about hippocampal function.

1.2 Behavioral correlates in the monkey hippocampus

Like much of the literature surrounding the hippocampus, electrophysiological studies of the monkey hippocampus can be broadly divided into two categories. The first is work concerning memory, stemming largely from the human and monkey work discussed in the previous section. This includes both working memory tasks (DMS and other tasks built around a delay) and studies of recognition memory that characterize differences in hippocampal representations of first and subsequent presentations of a stimulus. The second is work concerning spatial representations. This branch of the field stems from the Nobel-prize-winning discovery that hippocampal cells encode representations of an animal's physical location in space (O'Keefe, 1976; O'Keefe and Dostrovsky, 1971). Research in rodents has led to the hypothesis that the neural basis for a "cognitive map" (Tolman, 1948) that can be flexibly used for navigation (Buzsáki and Moser, 2013; O'Keefe and Nadel, 1978). Even as theoretical work has sought to reconcile these two models of the hippocampus (Bellmund et al., 2018; Whittington et al., 2020), studies in monkeys have had a difficult time replicating many of the most basic findings from the rodent navigational literature. Instead, studies of the monkey hippocampus have identified "view cells" that encode the location of a monkey's gaze. In this section, I will review the literature surrounding both spatial and mnemonic responses in the monkey hippocampus. I will suggest, however, that many of these responses are better understood in terms of the task the monkey is performing than in terms of a specific spatial or mnemonic response. Data from these studies, then, point to the idea that the hippocampus tracks the immediate experience of the animal (Rueckemann and Buffalo, 2017) and that many of the identified differences between rodent and monkey hippocampal responses are a reflection of different tasks and ethological considerations between animals rather than a fundamental difference in hippocampal function.

1.2.1 Mnemonic responses

As has been mentioned previously, it is now doubtful that many of the working memory tasks used in early monkey hippocampal studies were explicitly hippocampal-dependent. These tasks, however, were still able to elicit hippocampal neural firing. For example, Watanabe and Niki (1985) used a spatial DMS task in which monkeys had to respond with differential lever presses if a cue stimulus matched a second stimulus presented 3 seconds after the first. Hippocampal neurons responded during all epochs of this

task, i.e., there were neurons that fired during the cue period, neurons that fired during the delays between stimuli, neurons that responded to the second cue, and neurons that responded to the choice. Rather than spanning the entire delay period, delay neurons tended to be selective for a particular moment during the delay. What is more, a subset of delay neurons appeared to be trial-type specific, with selective firing for the cue used. Similar responses were later found in a similar task where a monkey had to respond to a match between two visual stimuli (Wilson et al., 1990), between auditory-visual pairings (Colombo and Gross, 1994), or between object-place pairings (Cahusac et al., 1993). Neurons recorded during object-place tasks showed analogous trial-type specificity to either specific stimuli or stimulus-response combinations (Miyashita et al., 1989). When monkeys trained to do delayed matching within categories of images (e.g. humans vs not humans, colorful vs not colorful) hippocampal cells showed selectivity for the entire trained category (Hampson et al., 2004), indicating that the stimulus specificity noted in previous studies can generalize to more abstract stimulus features. All of these studies used consistent delays. Notably, when the delay period was varied during training and recording, delay period activity was no longer seen in the hippocampus although activity correlated with cue, choice, and other task phases remained intact (Vidyasagar et al., 1991). Together, these results suggest that responses to particular epochs can be specific to particular trial types, but that the presence of these responses are largely independent of the stimuli used. Hippocampal delay period activity, however, is dependent on the monkey's familiarity with a consistent task structure.

A particular challenge with electrophysiological studies of memory is to create a controlled, repeatable framework in which to look for consistent neural responses. While DMS and similar tasks do this for working memory, such studies do not capture the "one shot" leaning that is indicative of episodic memory, that is, the ability of animals to recall an event after it only happened once. Laboratory frameworks for this phenomenon have revolved around "recognition memory" by testing for differences in neural responses or behavior during and after a single presentation of a stimulus. Early work explicitly queried a monkey's memory by asking the animal to respond to whether it was the first time a stimulus was shown (i.e., it is novel) or if the monkey had seen the stimulus before (the stimulus was familiar or repeated). Under this paradigm, a small (as little as 2.3%) but notable fraction of hippocampal neurons showed selectivity for novel stimuli (Rolls et al., 1989, 1993). The challenge in interpreting this result is

that the task convolved stimulus novelty with task response (Thome et al., 2012). Subsequent work aimed to dissociate recognition memory signals from task correlates by rewarding a monkey simply for viewing natural images (Jutras and Buffalo, 2010; Jutras et al., 2009; Sakon and Suzuki, 2019; Thome et al., 2012) or movies (Sakon and Suzuki, 2021) and looking for differences in hippocampal responses according to whether a stimulus was novel or familiar. A subset of monkey hippocampal neurons tend to have higher firing rates for novel stimuli (Jutras and Buffalo, 2010; Sakon and Suzuki, 2019, 2021; Thome et al., 2012; Xiang and Brown, 1998). The same is true in the monkey entorhinal cortex (Killian et al., 2012; Xiang and Brown, 1998). These changes in firing rate are not necessarily indicative of an explicit or independent novelty signal; neurons with novelty responses tend to be selective for both novel and repeated stimuli, but with higher firing rates for the novel stimuli rather than an exclusive novelty response (Thome et al., 2012). Yet the novelty preference of these neurons seems to be behaviorally relevant. In some recognition memory tasks, monkeys were allowed to look away from images at will and did so more quickly for familiar than novel stimuli, providing a behavioral indication of recognition without an explicit response from the animal (Jutras et al., 2009). In this case, increased firing for novelty was predictive of the animals' later memory (Jutras and Buffalo, 2010). Similarly, novelty responses differentiated truly novel images from "lure" images designed to resemble them (Sakon and Suzuki, 2019).

One challenge in the recognition memory studies mentioned thus far is that they rely on recognition of novel and repeated stimuli within a given recording session. In this case, novelty must be treated as a broad category while ignoring stimulus specificity. In humans, responses to well known or familiar people are actually greater than responses to unfamiliar people (Viskontas et al., 2006, 2009). These responses are also more selective for particular people and would likely be missed by the recognition memory assays used in monkeys (Sliwa et al., 2016). Although not directly analogous, monkey studies have found similar response patterns when monkeys were tasked with learning new associations. As monkeys learn the relationship between a rewarded fixation location and a particular image stimulus, a subset of cells broadly changed their responses over the course of a single learning session (Wirth et al., 2003). In this framework, each recording session included several new stimulus-response associations, which were learned within that session and compared to several previously learned "reference" associations that were held constant across days. Although some cells showed a broad preference for new stimuli, a

separate subset showed strong selectivity for one particular well-learned stimulus-response pairing (Yanike et al., 2004). Together, these results suggest that hippocampal responses to novel stimuli are broad and non-specific, while responses to well learned stimuli are sparse and selective.

Such a coding scheme may not be unique to the hippocampus. Work in structures as disparate as mouse primary visual cortex, for example, has suggested that neural responses to images within a task structure become sparser and more specific as the task is learned (Garrett et al., 2020; Poort et al., 2015)[1]. Though further work is required to identify what might separate hippocampal responses from those found elsewhere in the brain, these differences likely include the complex nature of hippocampal responses (e.g. to classes of stimuli or stimulus-response associations rather than the stimuli themselves (Hampson et al., 2004; Wirth et al., 2003)) and the relative rapidity with which within-session changes in neural responses take place (Wirth et al., 2003).

1.2.2 Spatial Responses

Inspired by the discovery of place cells in rodents (O'Keefe and Dostrovsky, 1971), a subset of early memory experiments incorporated the spatial location of stimuli as a factor in their delayed matching tasks. Though these studies did find spatially-selective firing in hippocampal neurons (Cahusac et al., 1989; Miyashita et al., 1989; Watanabe and Niki, 1985), there is little evidence showing that space is distinct from other factors (e.g. object identity) that the monkey might use to solve the task. To address this, (Rolls et al., 1989) used an object-place task that employed many spatial locations tiling a screen. Instead of finding responses specific to individual stimulus locations, they found that a subset of hippocampal neurons responded to stimuli in contiguous regions of the screen, providing early evidence for specific spatial coding in the monkey.

Whereas studies with rodents typically use a chronic implant that allows the animal to move freely, research using monkeys has focused on acute recordings that require a head-fixed animal. Until the relatively recent proliferation of chronic recording techniques for deep structures in the monkey (see below), this constraint made it impossible to have monkeys perform analogous tasks to those used in

[1] Further evidence comes from my own unpublished work using the data from (Garrett et al., 2020), where we show differences in the progression of adaptation over repeated stimulus presentations for novel and familiar images in primary visual cortex at both the single cell and population level.

rodents. Instead, researchers placed head-restrained monkeys in movable chairs and maneuvered them around a room. These studies identified subpopulations of hippocampal cells that responded to the monkey's motion (O'Mara et al., 1994) or orientation (Tamura et al., 1990). The only spatial responses that were identified, however, consisted of spatially-dependent modulation of hippocampal responses to elements of the task (e.g. reward) (Ono et al., 1993a). In other words, these studies failed to find consistent responses that were not better explained by some other aspect of the monkey's behavior. Instead, Rolls and O'Mara (Rolls and O'Mara, 1995) identified neurons that responded to direction of view. In a follow-up study, Rolls et al. (Rolls et al., 1997) approximated free motion by having monkeys pull their recording apparatus behind them on the cart while searching for rewards. By tracking their eye movements, this study found that just over 10 percent of neurons encoded information about where in the room the monkey was looking but not about the monkey's physical location in the room (Rolls et al., 1998). The activity of these "view" cells could not be wholly explained by the animals' head direction, and cells maintained roughly the same firing fields for a short period of time even when the lights in the room were turned off (Rolls et al., 1997). Subsequent studies tracked the eye movements of monkeys looking at a screen to confirm that monkey hippocampal cells are selective for particular parts of visual space either in response to guided saccades (Nowicka and Ringo, 2000b) or natural eye movements as the monkey freely explores natural images on the screen (Konig and Buffalo, 2021). Similar work in the entorhinal cortex has identified cells that have both grid (Killian et al., 2012) and non-grid (Meister and Buffalo, 2018) representation of gaze location on a screen. Taken together, these studies indicate (1) that the primate hippocampal formation does encode spatial information but that (2) in certain situations, the monkey hippocampus represents not where a monkey is in a physical space, but rather where in that space the monkey is directing his gaze.

One important comment is that, in most tasks, the monkey's gaze can reasonably be assumed to focused on where the animal is attending. This makes it ambiguous whether spatial responses are simply a function of the monkey's eye position in the environment, or if they indicate the location of the monkey's attention. Two studies (Feigenbaum and Rolls, 1991; Wilming et al., 2018) deconvolved gaze from attention by having monkeys fixate a particular location on a screen while responding to a stimulus elsewhere on the screen. Spatially responsive cells in both the hippocampus (Feigenbaum and Rolls,

1991) and entorhinal cortex (Wilming et al., 2018) responded better to the location of the task-relevant stimulus than to the explicit location of the fixation, suggesting that hippocampus is better described as representing visual attention than as an explicit readout of gaze location.

In light of the difficulty in studying freely moving monkeys, an alternative method for studying spatial representation is to train monkeys to navigate in virtual reality. A typical approach is to train monkeys to use a joystick to navigate a first-person avatar around a virtual environment populated with relatively distal landmarks and cues. Hippocampal cells recorded while monkeys navigate in virtual reality show correlates to the position of the monkey's avatar (Baraduc et al., 2019; Furuya et al., 2014; Gulli et al., 2020; Hori et al., 2005a; Wirth et al., 2017), suggesting that the hippocampus encodes information about the avatar's position. Studies of humans navigating in VR has shown similar results (Ekstrom et al., 2003a). The responses found in these experiments, however, bear important qualitative differences that distinguish them from the spatial responses so often found in rodents. Responses in the monkey hippocampus tend show multiple firing fields, and these fields tend to be less precise than the punctate "place" fields so often described in rodent studies (Hori et al., 2005a) or require additional non-spatial information to provide statistically significant prediction (Baraduc et al., 2019; Wirth et al., 2017). What is more, if the monkey performs a similar task at multiple locations in a virtual environment, selective cells will tend to have their clearest responses around the salient task locations and to have similar or symmetrical responses for each of these locations (Gulli et al., 2020; Hori et al., 2005a). When the monkey performs a different task in the same space, spatial responses change (Gulli et al., 2020; Hori et al., 2005a). Conversely, when task structure is maintained but spatial cues are altered, a majority of cells preserve their spatial firing (Baraduc et al., 2019). These data suggest that, though there exist neural correlates to avatar position in virtual reality, these responses are more indicative of whatever task the monkey is performing than an explicit readout of allocentric space.

It remains difficult to interpret these spatial responses in monkey hippocampus in the context of the broader rodent hippocampal literature; primate studies have focused on restrained monkeys, and it is unclear whether the differences in responses between rodent and monkey studies are a result of the constrained nature of this monkey work, or true difference between species. A limited set of studies has recorded from freely moving monkeys in navigation tasks. Squirrel monkeys searching for food in a 3D

environment show consistent, location-specific hippocampal activity (Ludvig et al., 2004a)[2]. More recent work with chronic implants has shown that both marmoset (Courellis et al., 2019) and macaque monkeys (Hazama and Tamura, 2019) have discrete and reliable spatial activity when shuttling down a linear track. These studies suggest that when monkeys are presented with rodent-like tasks, they can elicit rodent-like hippocampal responses. However, interpretation is complicated by the fact that, like a majority of rodent studies, experiments did not track other features known to influence responses in the primate hippocampus such as gaze. Modern work considering these co-factors suggests that including other behavioral features in addition to space improves predictions of hippocampal activity even in unrestrained monkeys (Mao et al., 2021) making it unclear to what extent the ostensibly spatial responses seen in freely moving monkeys can be better explained by other behavioral features.

An alternative approach to understand the difference between rodent and primate studies is to consider situations in which rodents are placed in constrained or virtual reality tasks that resemble those used in monkeys. Notably, studies with head-restrained rodents in VR paradigms have had a similarly difficult time recreating the reliable “place cells” seen in freely moving animals. Some have found, for example, that head fixing rats greatly reduces spatial activity in a virtual environment (Acharya et al., 2016; Aghajan et al., 2014; Aronov and Tank, 2014; Dombeck et al., 2010). Restoring a limited degree of motion, for example allowing head movements, rescues some of this spatial activity (Aronov and Tank, 2014). More reliable spatial activity is seen when the rat is given an explicit navigational task or goal (Aghajan et al., 2014), suggesting that the addition of task structure is sufficient to elicit spatial activity. Strikingly, spatial correlates in these studies bear a resemblance to those seen in monkey VR in that they tend to have a repeating structure in disparate locations where the animal performs a similar task. It is also noteworthy that tracking additional navigational variables in the entorhinal cortex of mice in virtual reality reveals correlates to eye and head movements (Mallory et al., 2021), leaving open the possibility that these behavioral features somehow influence how information is represented in the rodent hippocampal formation.

[2] As yet unpublished work from Daniel Huber present at the Ascona Circuits Meeting in 2019 show similar responses in freely moving mouse lemurs. See (Ho et al., 2021) for published work using similar techniques to record from mouse lemur V1.

1.2.3 Task correlates as an alternative model

One of the ironies of studying the monkey hippocampus is that many studies had trouble driving hippocampal activity, much less differentiating which behavioral features best correlated with it. Early work on spatial responses, for example, found that less than 10% of neurons showed spatial (Cahusac et al., 1989) or visual (Rolls and O'Mara, 1995) responses. This number is considerably smaller than the activity observed during tasks with inherent structure. Early DMS studies reported as many as 40% of hippocampal neurons selective for particular task epochs (Vidyasagar et al., 1991; Wilson et al., 1990). Even here, these responses were lackluster; Vidyasagar et al. (Vidyasagar et al., 1991) found that simply by showing their monkeys a raisin (a high value reward), they were able to elicit a greater degree of firing than in their DMS memory paradigm. In fact, this study went on to observe hippocampal neurons during stereotyped presentation of a raisin to the monkey, and found the neurons responded reliable responses to different presentations epochs of this "raisin show." This study, then, provides an early example of monkey hippocampal neurons tracking the salient experiences of the monkey, and suggests that this activity is most easily studied by imposing structure on the monkey's behavior. Further evidence is seen in the reward and timing representations consistently observed in the hippocampus across tasks, and in the task-dependent nature of monkey hippocampal spatial representations.

When a task has a distinct reward period, as many as 50% of monkey hippocampal neurons will show selective activity for this part of the task (Fuster and Uyeda, 1971; Wirth et al., 2009). Although one study noted that these task-epoch-specific responses were often the same regardless of task condition (in their case, responses were the same for reward or a mild electric shock) (Fuster and Uyeda, 1971), others have suggested that activity is specific to rewarded responses or epochs. Hippocampal neurons fire selectively for appetitive stimuli (Tamura et al., 1991). Separate subpopulations respond to correct or incorrect responses (Yanike et al., 2009). What is more, response to rewarded epochs appears to be a learned feature of a task response, rather than simply a feature of the reward itself. Responses to rewarded epochs become stronger as monkey learns an object-place association task (Yanike et al., 2009), and untrained monkeys show notably fewer reward-epoch responses (Fuster and Uyeda, 1971). Further, hippocampal responses overrepresent rewarded spatial locations both on a screen and in virtual reality (Bretas et al., 2019; Rolls and Xiang, 2005) and changing the rewarded location shifts spatial

response preferences (Rolls and Xiang, 2005). Taken together, these data demonstrate robust responses for the reward phase of many hippocampal tasks. These responses are not reward-specific, but rather reflect a learned relationship between tasks and rewarded outcomes.

Further indication that task structure drives hippocampal activity in the monkey can be seen in observed reliable responses to task timing. On a broad scale, hippocampal neurons carry timing information about epochs of a structured task, with neurons firing broadly across multiple phases of the task (called “incremental” timing by (Naya and Suzuki, 2011; Naya et al., 2017; Sakon et al., 2014)). Finer-scale timing is also observed. When a task has a consistently spaced delay period, a sub-population of hippocampal neurons fire at particular moments during that delay (Watanabe and Niki, 1985). This latter form of time coding bears striking resemblance to the “episode” or “time” cells seen in rodent studies (Kraus et al., 2013; Pastalkova et al., 2008), particularly as studies in rodents also find timing signals to be task-dependent. Studies that looked for reliable delay activity in the monkey without long-term training at a consistent delay failed to find such punctate delay activity (Vidyasagar et al., 1991), suggesting a similar task dependence in the monkey. Further task dependence of timing signals is seen at all levels in that both incremental signals (Sakon et al., 2014) and fine-grained delay signals (Hampson et al., 2004; Miyashita et al., 1989; Watanabe and Niki, 1985) show context dependence, with differential responses for different task conditions or trial types.

Considering an animal’s task may also help disambiguate apparent mismatches in spatial representations between different monkey studies. For instance, when a monkey was tasked with doing a particular task on a screen, neurons maintained their spatial relationship to the screen even as the screen was moved around a room (Feigenbaum and Rolls, 1991). This “locally allocentric” behavior is contrast to the later discovery of “globally allocentric” activity of gaze-responsive cells that fired wherever the animal looked at a specific location within a room during an exploration task (Rolls and O’Mara, 1995). This apparent mismatch in reference can be reconciled by considering the behavioral task the monkey was performing; neurons coded information in terms of the screen when screen information was necessary to complete the task, and in terms of allocentric visual space when exploration of the whole room was necessary to complete the task. A similar explanation can be applied to the apparent mismatch between neurons that fire based on where a monkey is looking on a screen (Cahusac et al., 1993; Feigenbaum

and Rolls, 1991; Killian et al., 2012; Konig and Buffalo, 2021; Meister and Buffalo, 2018; Nowicka and Ringo, 2000b) and responses noted in virtual reality projected on a screen (Baraduc et al., 2019; Gulli et al., 2020; Wirth et al., 2017) or between gaze-related response in an exploration task (Rolls and O'Mara, 1995) and apparent spatial responses on a linear track (Courellis et al., 2019). In these situations, neuronal responses span the space that the monkey needs to consider in performing the task at hand, while hippocampal correlates do not appear to track task-irrelevant spatial reference frames. Further, in the few cases where a task does not explicitly require a particular reference frame (Meister and Buffalo, 2018; Tamura et al., 1990) studies of the hippocampal formation have noted presence of both egocentric and allocentric neural correlates suggesting that one default reference frame may not exist in the absence of a task.

Changing task structure can change hippocampal responses. This result holds true across spatial reference frames. When monkeys switch from image viewing to a guided saccade task, the spatial responses of hippocampal neurons change (Konig and Buffalo, 2021). Spatial responses are also eliminated when visual stimuli were removed from a guided saccade task (Nowicka and Ringo, 2000b). Although such data do not yet exist for freely moving monkeys, one group (Ono et al., 1991, 1993b, 1993a) developed a framework in which a monkey was required to drive himself around a track on a cart. The monkey was cued to manipulate a lever to move the cart, then rewarded for correct manipulation. Neurons in this task responded to particular task epochs, such as lever presses and reward. What spatial responses were seen consisted of spatial modulation of e.g., reward signals, rather than responses to navigation or motion. This an effect that was particularly notable when the monkey's response depended on the spatial location of the cart (Eifuku et al., 1995). When the control of the monkey's cart system was removed and the animal was instead passively driven around the environment, whatever spatial preference these neurons had was eliminated (Nishijo et al., 1997a), demonstrating that any spatial modulation was task-specific. Studies in virtual reality have extended this result by showing that changing, rather than removing, the task a monkey is doing while holding the environment constant is similarly sufficient to eliminate any spatial response (Gulli et al., 2020; Hori et al., 2005b).

Rodent studies have well characterized the phenomenon of remapping, in which place cells either change their responses in new environments or in new tasks (Leutgeb and Leutgeb, 2014; Pastalkova et

al., 2008). These studies are consistent with the observation of remapping in response to a task change. In the monkey, however, changing the spatial appearance of an environment is not sufficient to drive "remapping" in all spatially-correlated cells (Baraduc et al., 2019; Hori et al., 2011), indicating that it may be the task structure rather than the relationship of the task to specific spatial co-variates that drive hippocampal activity. This latter observation marks one of the starkest differences between rodent and primate hippocampal responses, and suggests that remapping needs to be considered in terms of behavioral context rather than explicit metric space; whereas changing cues in an environment may be sufficient for a rodent to consider itself in a new context, monkeys (or at least the monkey hippocampus) appears to encode context in terms of the task at hand rather than as a function of the details of its environment.

1.3 Physiological hallmarks of the monkey hippocampus

Thus far I have focused on the work looking at the behavioral correlates of spiking activity in the monkey hippocampus. This work has concerned extrinsic factors that relate to hippocampal activity, that is, what features of given behavioral task relate to activity of hippocampal neural activity. Rodent work, however, has pinpointed a number of oscillatory dynamics visible in hippocampal local field potentials (LFPs) that structure hippocampal activity. Although the focus of this thesis is on the spiking behavior of monkey hippocampal neurons, in this section I will briefly address two of these oscillatory phenomena: (1) sharp wave ripples, which are often seen as a correlate of memory (Buzsáki, 2015) and appear to be broadly conserved across species (Colgin, 2016) and (2) theta oscillations, which structure rodent place cell activity but have notably different behavior in the monkey. Further, I will consider the hypothesis that eye movements provide structure to activity in the monkey hippocampus that may be analogous to theta oscillations in the rodent.

1.3.1 Sharp wave ripples

Sharp wave ripples are brief, synchronous bouts of high frequency hippocampal activity thought to be correlated with memory consolidation and often used as a biomarker for memory and planning (Buzsáki, 2015). They typically occur when an animal is relatively inactive and hippocampal activity becomes otherwise desynchronized (Green and Arduini, 1953). In rodents, sharp wave ripples are known to co-occur with hippocampal replay events, where the same sequence of cell activity seen during a task

occurs in the same order as seen during the task, but in a temporally compressed pattern. Similar replay activity can be observed synchronously elsewhere in the brain, suggesting that sharp wave ripples provide a physical basis for coordinating cross-area replay events. Sharp wave ripple patterns are preserved as a physiological hallmark of hippocampal activity across species (Buzsáki et al., 2013). As in rodents, monkey sharp wave ripples are most common when the monkey is in an inattentive behavioral state, drowsy, or even asleep, and are accompanied by bursts of otherwise sparsely firing neurons (Skaggs et al., 2007a). One study (Leonard et al., 2015) has suggested that sharp wave ripples may be a feature of active exploration in the monkey, though these results are in conflict with the remainder of the literature[3]. To my knowledge, no study with monkeys has yet investigated the explicit content of sharp wave ripple associated spiking, making it difficult to understand if there are two modes of sharp wave ripple activity or if this one study was misguided in its ripple characterization. Further, monkey sharp wave ripples are associated with the same synchronous activity in other brain areas that is seen in rodents (Buzsáki et al., 2013; Logothetis et al., 2012), providing additional evidence that monkey sharp wave ripples fulfill a similar physiological role to those in other species.

1.3.2 Intermittent theta oscillations

Theta oscillations represent one of the dominant modes of LFP in the rodent hippocampal formation (Colgin, 2016). In contrast to the desynchronized hippocampal LFP when an animal is inactive (with the exception of sharp wave ripples (Buzsáki, 2015)), theta oscillations consist of an ~8Hz (though ranging from 6-12Hz) rhythm that occurs when an animal is active, e.g., running, exploring, or otherwise alert (Colgin, 2016). As with many LFP phenomena, the explicit “purpose” of hippocampal theta remains unclear. With that said, theta does provide a mechanism for organizing sequences of spiking in the rodent hippocampus. As a rat moves down a linear track, for example, place cells representing current and upcoming locations will fire at successive moments relative to the phase of the theta oscillation (Skaggs et al., 1996). Hippocampal theta rhythms and their timing are entrained to theta in the medial septum (Petersen and Buzsáki, 2020; Vandecasteele et al., 2014), and disrupting septal theta subsequently disrupts these sequences in the hippocampus (Wang et al., 2015). This same manipulation impairs

[3] Including our own unpublished work from the same dataset described in Chapters 2,3 and 4 of this thesis.

behavior on a delayed alternation task, suggesting that rodent theta – or at least the cellular organization it provides – contributes substantively to behavior.

Given the ubiquity and apparent importance of hippocampal theta in rodents, it is striking that theta oscillations appear to have a very different structure in other animals, including monkeys. In contrast to the consistent oscillation seen in alert and/or moving rats, work in monkeys (Courellis et al., 2019; Green and Arduini, 1953; Jutras et al., 2013), bats (Eliav et al., 2018; Yartsev et al., 2011a), and humans (Goyal et al., 2020) has failed to find a consistent theta oscillation. Instead, theta in these animals occurs in intermittent bouts lasting as little as a few theta cycles (Killian et al., 2012; Ulanovsky and Moss, 2007). These bouts have some behavioral correlation; novel smells, for example, can drive theta bouts in monkeys (Green and Arduini, 1953). Further, theta within a bout seems to have similar behavior to rodent theta. In monkeys, for example, grid cells in the entorhinal cortex show within-bout entrainment to theta (Killian et al., 2012). There is some evidence to suggest that theta behaves as a traveling wave in humans (Zhang and Jacobs, 2015), suggesting a similar wave pattern to that seen in rodents (Colgin, 2016; Lubenov and Siapas, 2009). Despite these similarities, there remain striking differences between these animals with theta bouts and those with consistent theta. For example, the monkeys discussed later in this thesis are able to perform the delayed matching task, although (Wang et al., 2015) note that rodents with disrupted theta are impaired on a similar task. Moreover, although grid and place cell firing is disrupted when theta is altered through septal manipulation in rodents, spatial firing in hippocampus and entorhinal cortex of bats and monkeys remains intact outside of theta bouts (Killian et al., 2012; Yartsev et al., 2011a). Together, these data challenge existing models for the explicit function of theta shaping representations in the hippocampal formation (see (Burgess and Burgess, 2014; Burgess et al., 2007)), but leave open the possibility that some other process provides analogous organization of hippocampal activity in monkeys and other animals with non-continuous bouts of theta.

1.3.3 Eye movement correlates structure monkey hippocampal activity.

It remains unclear what might coordinate hippocampal activity in monkeys and other animals that lack consistent theta. Although the large-scale recordings needed to identify cell sequences in the rodents have not yet been achieved in the hippocampus of monkeys or other animals, there is evidence to suggest similar mechanisms may organize cell firing. Work in bats has shown phase precession relative

to non-periodic features of the LFP (Eliav et al., 2018), suggesting a similar organization mechanism despite the absence of an identifiable oscillation within a consistent frequency band. To the best of my knowledge, analogous research in monkeys has not yet been performed. Instead, studies in the monkey have noted correlates in the hippocampal LFP that correspond to the onset of saccadic eye movements. Theta oscillations that are present in the monkey are strongest around the timing of saccades (Hoffman et al., 2013). More specifically, theta oscillations undergo a saccade triggered phase reset (Jutras et al., 2013), suggesting that saccades may provide temporal structure to the activity within the hippocampal formation. Although these results are complicated by concerns about recording contamination by ocular-muscle artifacts that cannot be wholly controlled for with the recording techniques used (Buzsáki et al., 2012; Kovach et al., 2011), the idea that eye movements might structure hippocampal activity is strongly supported by the observation that eye movements also correlate reliably with the timing of individual neurons in the hippocampal formation.

A subset of cells in the monkey hippocampus lock to eye movements. Hippocampal activity locked to the onset of saccades was observed when the monkey was not doing a task both in darkness and in light (Ringo et al., 1994), suggesting that hippocampal activity related to eye movements is not a response to some specific visual stimulus. Subsequent work noted similar responses in tasks involving guided saccades (Nowicka and Ringo, 2000b; Sobotka et al., 1997), when a monkey was freely viewing pictures (Konig and Buffalo, 2021), or freely moving around an arena (Mao et al., 2021). These data are particularly notable because the primate hippocampus is typically considered a “sensory” or “visual” area, although olfactory (Green and Arduini, 1953; Yokota et al., 1967) and auditory (Tamura et al., 1990) stimuli can also elicit hippocampal activity. Further, humans (Scoville and Milner, 1957) and monkeys (Rueckemann and Buffalo, 2017) with hippocampal damage are not described as having visual impairments. Yet the number or fraction of neurons showing eye movement modulation (20% or greater (Nowicka and Ringo, 2000b; Sobotka et al., 1997)) dwarfs the fraction of neurons with specific responses in memory and spatial tasks more often attributed to the hippocampus (Rolls and O’Mara, 1995; Rolls et al., 1993, 1997). Moreover, eye movement correlates are not specific to the hippocampus but rather distributed over the medial temporal lobe (Minxha et al., 2017; Nowicka and Ringo, 2000b; Ringo et al.,

1994; Sobotka et al., 1997), suggesting that eye movement related responses may underlie some common process not unique to the hippocampus.

Although some cells show eye movement correlates in the absence of other clear behavioral correlates, cells that carry spatial or task information can also do so in a manner that is convolved with eye movement correlates. Hippocampal neurons recorded during a guided saccade task show selectivity for the direction of a saccade (Sobotka et al., 1997), with a higher amplitude eye-movement triggered response in a cell's preferred direction. Similar saccade-direction selectivity is seen in cell in the entorhinal cortex during image viewing (Killian et al., 2015). In a separate task, eye movement correlates were more pronounced for a cell's preferred spatial location on a screen (Nowicka and Ringo, 2000b). This latter result has been corroborated by work in our own lab identifying "view" cells as monkeys freely viewed images on a screen. In this case, some neurons showed eye-movement locked activity only when the monkey fixated a particular region on the screen (Konig and Buffalo, 2021). Further, like other task related responses, responses with eye movement correlates can be context dependent, losing both their firing fields and their eye-movement-triggered firing when the task is changed (Konig and Buffalo, 2021; Nowicka and Ringo, 2000b; Sobotka et al., 1997). Thus, eye-movement related firing does not necessarily reflect a separate population from the spatial, mnemonic, and task responses discussed in the previous section. Instead, it appears that many of these responses are structured by eye movement correlates (Minxha et al., 2017). What is more, not all of these eye movement correlates have common timing relative to eye movement onset. Peaks or troughs in firing rate can begin as early as the onset of the saccade or late enough as to span into a subsequent saccade (Konig and Buffalo, 2021). Such delays are fixed relative to the timing of the eye movement, and do not rescale with fixations of different durations. Similar delays in response to visual stimuli spanning the range of a fixation, some 50-450 ms after the stimulus presentation, have been observed in studies that did not explicitly check for eye movement correlates (Cahusac et al., 1993; Miyashita et al., 1989; Riches et al., 1991; Rolls et al., 1989, 1993). In fact, neurons with object- or task-specific coding may tend towards longer delays (Eifuku et al., 1995), suggesting that delay time could be indicative of time needed for information processing. Although no study has as yet examined eye movement correlates in population recordings, these results leave

open the possibility that eye movements coordinate a sequence of neural activity within the hippocampus that produces ordered firing, a mechanism reminiscent of the theta sequences seen in rodents.

It is worth noting that eye movement signals are not necessary for specific response phenotypes in the hippocampal formation. When monkeys are trained to fixate a particular location while making responses about events elsewhere on the screen, cells in the hippocampus (Feigenbaum and Rolls, 1991) and entorhinal cortex (Wilming et al., 2014) show responses specific to where events on the screen occur despite the animals holding their eyes stationary. This is not inconsistent with observations from rodents; medial septal manipulations that alter rodent theta do not destroy place field responses, but rather affect the temporal order in which those place cells fire (Petersen and Buzsáki, 2020) and/or change the information content of those responses (Wang et al., 2015).

The non-vergence nature of rodent eye movements and the wide rodent field of view (Wallace et al., 2013), combined with the fact that rodent gaze is rarely tracked experimentally, make it difficult to explicitly compare the timing of eye movement related hippocampal activity in monkeys to the more abundant rodent hippocampal literature. One study (Mallory et al., 2021) has noted eye movement correlates in the rodent entorhinal cortex, but this study does not address whether the striking timing relationship between eye movements and hippocampal activity is seen in the rodent. Instead of eye movements, rodent work has demonstrated a relationship between rodent sniffing and whisking behavior and hippocampal theta (Colgin, 2016). Early work recording rodent facial movements identified coherence between sniffing and whisking behavior (which are themselves correlated) and hippocampal theta (Komisaruk, 1970; Macrides et al., 1982). This relationship does not always hold; the phase relationship between theta and sniffing breaks down when animals are not performing a task or are exploring in an unstructured manor (Berg et al., 2006). It is specifically when the sniffing (Macrides et al., 1982) or whisking (Grion et al., 2016) is relevant to the animals' behavioral task that sensory-theta coherence is present, both in the hippocampal LFP and in the activity of individual neurons. Importantly, differences in sensory behavior alone cannot explain the differences in monkey and rodent theta. Rodent theta is ubiquitous during periods of exploration even when it is not apparently phase-locked to any sensory behavior (Berg et al., 2006). Thus, rodent theta can be entrained to sensory behavior, but theta itself is not dependent on any sensory relationship. While it is not yet clear what mediates the theta bouts seen in

monkeys, it is unlikely that behavioral differences between species will be sufficient to explain the differential behavior of LFP activity in the hippocampus.

Although the link between sensory activity and hippocampal theta is not explicitly causal, the striking relationship between these the two in monkeys, rodents, and other animals (e.g. weakly electric fish (Fotowat et al., 2019)) suggests that active sensory behavior can organize neural activity in the hippocampus (Meister and Buffalo, 2016). Although differences in theta between species cannot be explained by differences in sensory behavior, eye movements may prove an important tool for understanding the organization of hippocampal activity in the monkey. As increasingly large-scale recordings from the monkey become possible, such an understanding may prove useful as a starting point for determining whether the same sequential activity that organizes activity in the rodent hippocampus has an analogue in other species.

1.4 Are monkeys just giant rats?

It would be willfully naive to consider electrophysiological studies in monkeys in isolation; the fact remains that the vast majority of hippocampal electrophysiology has been done in rodents. One likely experimental difference between studies in rodents and monkeys is the type of electrodes used in recordings from each animal. Recordings in rodents often use tetrodes (Gray et al., 1995) or other multi-contact devices (e.g. Neuropixels, (Steinmetz et al., 2021) that allow for picking up a single neuron on multiple contacts, while monkey researchers almost always employ single wire electrodes in their recordings. The advantages of multi-contact probes are that they pick up neural signals simultaneously on many electrodes, allowing timing information to be used for cell isolation, and the small differences in electrode locations allow for use of spatial information in differentiating individual neurons. This point is far from arcane; the few studies that use tetrode recordings see a higher fraction of sparse firing, complex spiking cells in the monkey hippocampus (Leonard et al., 2015; Skaggs et al., 2007a) suggesting the ‘units’ recorded in single wire studies may often consist of multiple cells rather than the clean, isolated neurons characterized in rat studies. In fact, data from these same rat studies show that isolation on a single tetrode wire is greatly reduced from that of the full tetrode and may include some degree of bias towards high firing rate neurons (Chung et al., 2017a; Gray et al., 1995). Further, rodent studies are far more likely to employ chronic implants, meaning they are less likely to suffer from sampling biases that

may drive researchers towards high firing rate cells that are more likely to be interneurons (Sliwa et al., 2016). While this technical note is in no way meant to diminish work in monkeys, It is important to underscore that single unit detection in monkeys is in a more primitive state, and so care should be taken when comparing the temporal properties and response specificities of monkey hippocampal experiments to those in rodents.

Just as neural recording techniques very between species, so do the behavioral variables that are recorded. The most consequential of these differences is that monkey studies often track the monkey's eye movements, making it possible for identification of view cells and eye movement-related activity. With rare exceptions ((Mallory et al., 2021; Wallace et al., 2013) are examples, though by no means an exhaustive list) studies in freely moving rodents almost never track rodent eye movements. Visual manipulations in the rodent tend to involve large manipulations of visual cues (e.g. (Aghajan et al., 2014)). These manipulation studies do show changes in hippocampal activity, leaving open the possibility that view-cell-like behavior may explain some rodent hippocampal activity. Similarly, studies have identified spatial responses as a rat or bat watches a conspecific perform a task, suggesting that gaze can inform hippocampal firing in these animals (Danjo et al., 2018; Omer et al., 2018; Rao et al., 2019). However, until more specific eye tracking is used in rodents, direct comparisons will remain difficult. A similar critique could be offered for the lack of monitoring of sniffing or breathing in modern monkey studies, making it unclear if non-visual sensory inputs have any organizational role in hippocampal activity.

I have alluded to a number of key differences between rodent and monkey behavioral studies. Although some have suggested that the hippocampus has fundamentally distinct roles in different species (Quiroga, 2019, 2020), it is likely that the observed differences in response phenotype are simply a result of different behavioral tasks being used. As I have discussed in previous sections, when similar experimental methods are used across species, remarkably similar hippocampal representations are observed. When cross-species differences are seen, they can often be explained by differences in

ethological niches across species (Payne et al., 2021)[4]. As a result, studying different animals can give us different behavioral lenses for examining the central tenets of hippocampal function.

This is not meant to imply that the primate and rodent hippocampus are identical. I have already discussed the lack of sustained theta oscillations in monkeys and humans. There is additional histological evidence to suggest organizational differences in the structure of rodent and primate hippocampus. Notable among these is the observation that the rodent hippocampus exhibits anatomical distinction between its dorsal and ventral sections, such that these areas are more interconnected within themselves than between each other (Strange et al., 2014). No such divide exists in the monkey, although it is seen in other species including food caching birds (Payne et al., 2021). Importantly, spatial representations in the dorsal component of the rodent hippocampus are finer grained than those in the ventral section, and this anatomical distinction along this longitudinal axis of the hippocampus offers a potential mechanism for the different behavior of these cells. In the monkey, the anterior-posterior axis of the hippocampus is considered to be analogous to the dorsal-ventral axis in the rodent. It is not clear to what extent monkey responses exhibit a similar gradient. One dataset from the entorhinal cortex suggests that there are correlations between grid spacing and saccade modulation and anatomical location (Killian et al., 2012, 2015), with tighter grids and greater saccadic modulation in the posterior cortex. Data from the hippocampus are sparser. Most studies report no significant anatomical distribution of responses in task responses. One notable exception is (Colombo et al., 1998), which reports a concentration of spatial responses to the posterior hippocampus, though this study and its predecessor (Colombo and Gross, 1994) did not show explicit examples of this gradient, making full evaluation of this distinction difficult. Moreover, most of these studies examined the presence or absence of hippocampal responses rather than differences in the specific responses themselves. There is preliminary evidence from the marmoset that anterior and posterior spatial representations may differ in specificity (Courellis et al., 2019), but this

[4] (Payne et al., 2021) is actually a study of activity in the hippocampal analogue of birds, rather than rodents or primates. They find that a food-caching bird species has exquisite place cell representations in a food caching task, while a not food-caching species does not. In addition to providing further evidence that spatial representations only occur the hippocampus when those spatial representations might be in some way relevant to the animal, these data suggest an ethological dependence of hippocampal representations.

study examined responses at different points along the longitudinal axis in a small selection of monkeys, making it impossible to deconvolve individual difference and anatomical specificity. Further large scale, anatomically specified studies are needed to determine if there is an anatomically organized distribution to primate hippocampal responses. I will present one such survey in the next chapter of this thesis.

1.5 Conclusions

Monkeys are not simply giant rats. Though it is unclear if the difference between the rodent and monkey hippocampus are a function of these species' behavioral difference's, anatomical differences, or both, it is clear that findings in rodents do not completely account for hippocampal activity in the monkey. A complete understanding of the hippocampus – particularly one with a bent towards clinical relevance – will require establishing reliable predictors of hippocampal activity in the monkey that allow for targeted experimentation that to date remains possible only in rodents. In this thesis, I will present results from one of the largest surveys of monkey hippocampal activity performed to date. Using results from a number of virtual reality tasks, I will argue that task-related firing may provide one such reliable predictor. Additionally, I will show preliminary evidence suggesting that eye movements provide some level of temporal organization in the monkey hippocampal population even in relatively unconstrained viewing behaviors.

2 Hippocampal Neurons Track Task Progress in Virtual Reality

2.1 Introduction

We do not yet know what structures neural activity in the hippocampus. A half century of work has highlighted the exquisite spatial correlates of individual neurons in the hippocampus and surrounding cortex, leading to the hypothesis that the hippocampal formation provides a neural substrate for a cognitive map that can be queried by an animal for navigation (McNaughton et al., 2006; O'Keefe and Nadel, 1978). However, recent evidence suggests that activity in the hippocampus may be better described in terms of the task an animal is performing than the physical space in which that task takes place (Aronov et al., 2017; Pastalkova et al., 2008).

Here, we present data from an investigation of the activity of individual neurons in the hippocampus of monkeys trained to navigate an ordered task in virtual reality. We show that, although there are spatial correlates in the hippocampus of monkeys performing this task, these data favor the hypothesis that task structure is driving hippocampal responses.

2.2 Results

We trained three monkeys to use a joystick to traverse a virtual delayed spatial alternation task (a Y-maze, Figure 2-1 (a, b)). Monkeys were rewarded for finding a hidden target (a floating banana) at the end of either the right or left arm of the maze, before being teleported back to the start point after a delay. The correct choice was always the direction not chosen on the previous trial; incorrect trials were unrewarded, and the target location did not change. Once trained, monkeys have high accuracy in this task (mean/median 95.5/97% correct, Figure 2-2) precluding separate analysis of error trials. Importantly, the monkey's avatar's movements are constrained to a track system in this task, allowing us to look for consistent cell activity while the monkey performed an ordered set of actions. We used a monocular infrared eye tracker to estimate the monkey's eye position on the VR display screen and extrapolate this into an estimate of the monkey's virtual gaze (see Section 2.4, Methods).

We used chronically implanted hyperdrives, each loaded with 124 high impedance single-wire electrodes, to record activity from the right hippocampus of each of the three monkeys as they performed this task. The large footprint of these hyperdrives allowed us to sample the full anterior-posterior and

medial-lateral extent of each monkey's right hippocampus, including all major subfields. Figure 2-3 shows an illustration of our drive and target recording locations.

This chapter contains data from 555 recording days, collected from 3 monkeys. Because of the sheer volume of these data and the time required to process it, many of the analyses presented here include different subsets of the population data. All analysis will be confirmed using the complete dataset before these data are submitted for peer review.

2.2.1 Spatial Reponses

We first asked whether hippocampal units showed activity related to the monkey's avatar's position in this virtual task. To avoid confounds with reward-related activity, we confined this analysis to the periods of the task when the monkey was able to drive or turn, that is, excluding times when he was frozen to receive reward. Intuitively, any unit carrying an interpretable spatial code must, first, carry some information about the monkey's avatar's position in the maze and, second, be sufficiently consistent in its firing that this information might be reliably decoded. We explicitly tested for these factors by quantifying (a) the spatial information carried by each spike train in terms of information per spike (Skaggs information), (b) the spatial information per unit time (Shannon's mutual information), and (c) the consistency of each cell's rate map (the Spearman's correlation of the spatial responses between halves of each behavioral session). We then compared each of these quantities to distributions generated from rotated spike trains to establish a distribution of "chance" values for each of these metrics. This procedure is discussed in detail in Chapter 4; briefly, this rotation consisted of circularly shifting spike times within each trial to preserve spike train statistics but decouple spiking from any behavioral feature. Distributions for these metrics are seen in Figure 2-1. Of 1344 recorded hippocampal units, 455 (34%, 2 monkeys, including cells within 2 mm of hippocampal STL) carried more information and were more consistent than expected by chance (threshold $p<.01$)[5]. Example rate maps are shown in Figure 2-5, with additional examples in Figure 2-21.

[5] Note that this fraction of responsive cells depends on what fraction of space is included in any given analysis. The number reported here considers all space that the monkey can traverse. Including only the 3 principal arms of the maze (i.e. excluding the nubs the monkey drives down to receive bananas – a relevant distinction because the monkey only learns that whether his choice was correct on a given trial

We next sought to examine the distribution of spatial responses within our population. Hypothetically, an efficient spatial code would have uniform coverage of all maze locations, so particular locations could be discriminated from the population response. Figure 2-7 shows the normalized rate-maps for all recorded hippocampal units. Importantly, peak response does not evenly tile the full maze. Instead, responses cluster around the beginning of the center arm, the end of arms, and the approach to the reward. To quantify the degree to which the population responses distinguish between locations in the maze, we computed the correlation between population response vectors at each point in the maze space. Effectively, this gives a degree to which location might be "decoded" from the neural population, though instantaneous decoding is not possible because our population was not recorded simultaneously. Figure 2-6 shows possible outcomes for this correlation analysis, while Figure 2-7 shows population correlation matrices for our spatial data. Notably, only the beginnings and end of arms, as well as the small segments at the end of the arm proceeding the reward (the "nubs"), can be reliably distinguished from other locations. Between these salient task locations, it appears that the population divides the maze into "epochs," such that each epoch can be distinguished from one another, but the monkey's specific location within that epoch cannot.

One challenge with our chronic drives is that the fact that recordings are made on single wires, combined with the inability to move electrode locations in real time, means that cell isolation can be relatively mediocre. To confirm that this result was not the explicit by-product of e.g., poor or selective sorting, we used the automated "Mountain-Sort" algorithm to automatically identify cells. Figure 2-9 shows a comparison between mountain-sorted and manually sorted cells. The lack of clear distinction between the population responses identified using these sort types suggests that this result is not a byproduct of selective spike sorting.

Our drives span the longitudinal axis of the monkey hippocampus, as well as all of the major subfields of the hippocampus (Saleem and Logothetis, 2012). We see no relationship between identified hippocampal subfields and the likelihood of a cell being spatially responsive (see Figure 2-8). In the rodent, a gradient of spatial firing along the longitudinal axis of the hippocampus has been interpreted as

---

once he is able to turn into the nub) reduces the fraction of responsive cells to ~27%. This is the method used for the analysis in Figure *2-8*.

anatomical evidence for a role of this structure in spatial processing (Strange et al., 2014). However, Figure 2-8 shows that we see no relationship between location along the longitudinal axis of the hippocampus and the amount of spatial information carried by units, nor any quantitative change in field size or number, suggesting that this gradient does not occur in monkeys[6].

[6] Figure 2-8, as well as this result, are built from preliminary data from monkey SP. I will redress this section prior to journal submission.

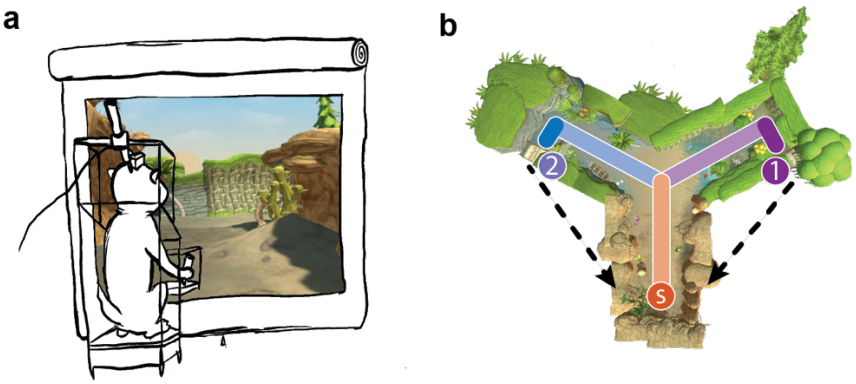

Figure 2-1 Virtual Reality Setup and Y-maze Task

(a) Monkeys use a joystick to navigate a virtual environment.

(b) Spatial Delayed Alternation Task. The correct response is always the opposite direction of the previous trial.

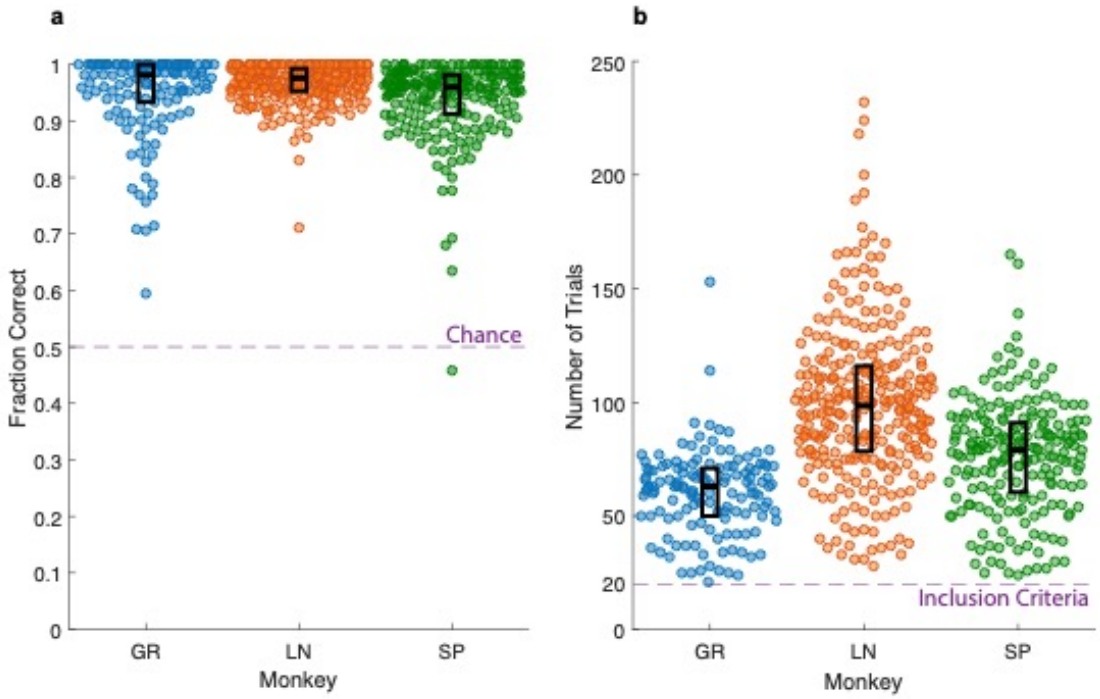

Figure 2-2 Y-maze Behavioral Performance

(a) Performance in the Y-maze task for included sessions. Note that each recording day can include 1-4 sessions.

(b) Number of trials completed per session. Sessions with fewer than 20 trials were excluded from analysis.

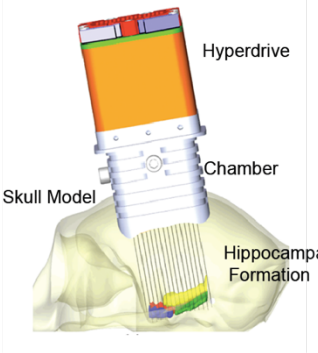

Figure 2-3 Illustration of Gray Drive Recording System

Recording targets illustrated for monkey GR. Skull model and structures are reconstructed from pre-surgical MRI. Yellow shows the hippocampus proper, with the adjacent medial temporal lobe cortex also shown.

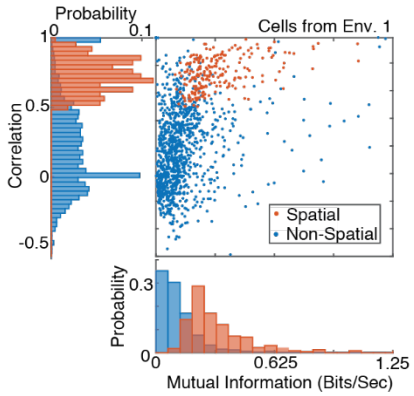

Figure *2-4* Distributions of Correlation and Information for Spatially Responsive Cells.

Correlation and mutual information distributions for manually sorted hippocampal cells. Orange denotes cells with significant spatial activity.

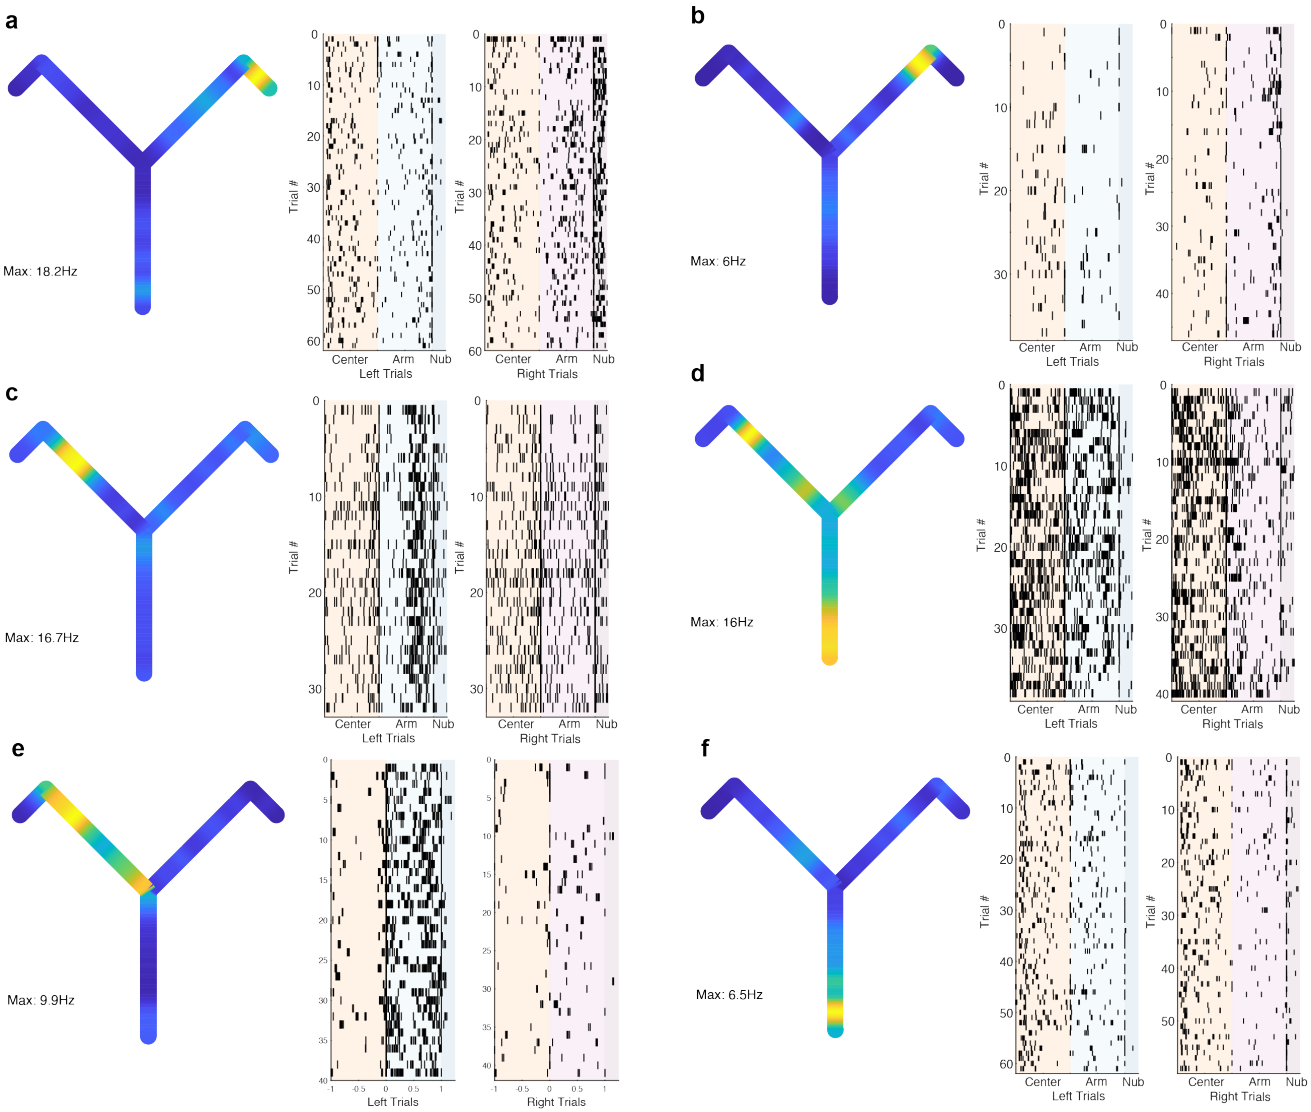

Figure 2-5 Example Spatially Selective Cells

(a-f) Examples of spatially responsive neurons. (Left) Rate-map (Right) raster plot separated by trajectory.

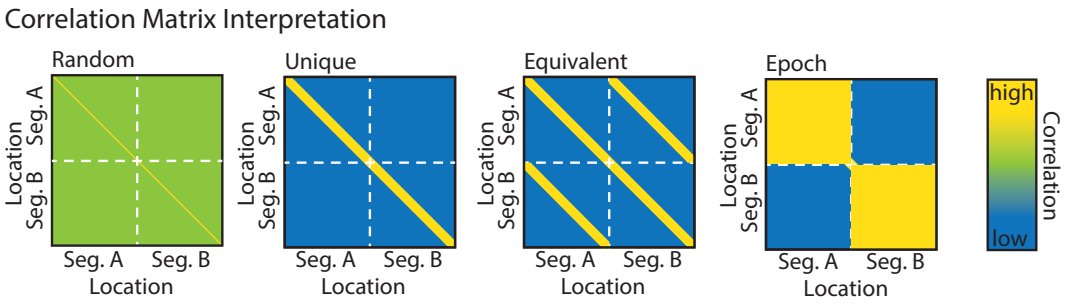

Figure 2-6 Possible Correlation Matrices

Idealized versions of potential population correlations, shown here in terms of an animal's location.

(Random) If no spatial relationship exists, population correlation will be relatively low over the space. If the matrix in question is an autocorrelation, then by definition the diagonal must equal one.

(Unique) If each location has a unique spatial code, correlations will be low everywhere but high along the diagonal. The width of the diagonal band indicates how specific the code is, that is, the degree of uncertainty in the spatial location given the population activity.

(Equivalent) If multiple parts of the maze have equivalent representations, off diagonal bands of high correlation will emerge.

(Epoch) If a particular part of the maze has a constant representation, this area can appear as a "box" of high correlation. This is not wholly different from the "unique" diagonal band, except that the highly correlated/uncertain region extends over the entire segment of the task.

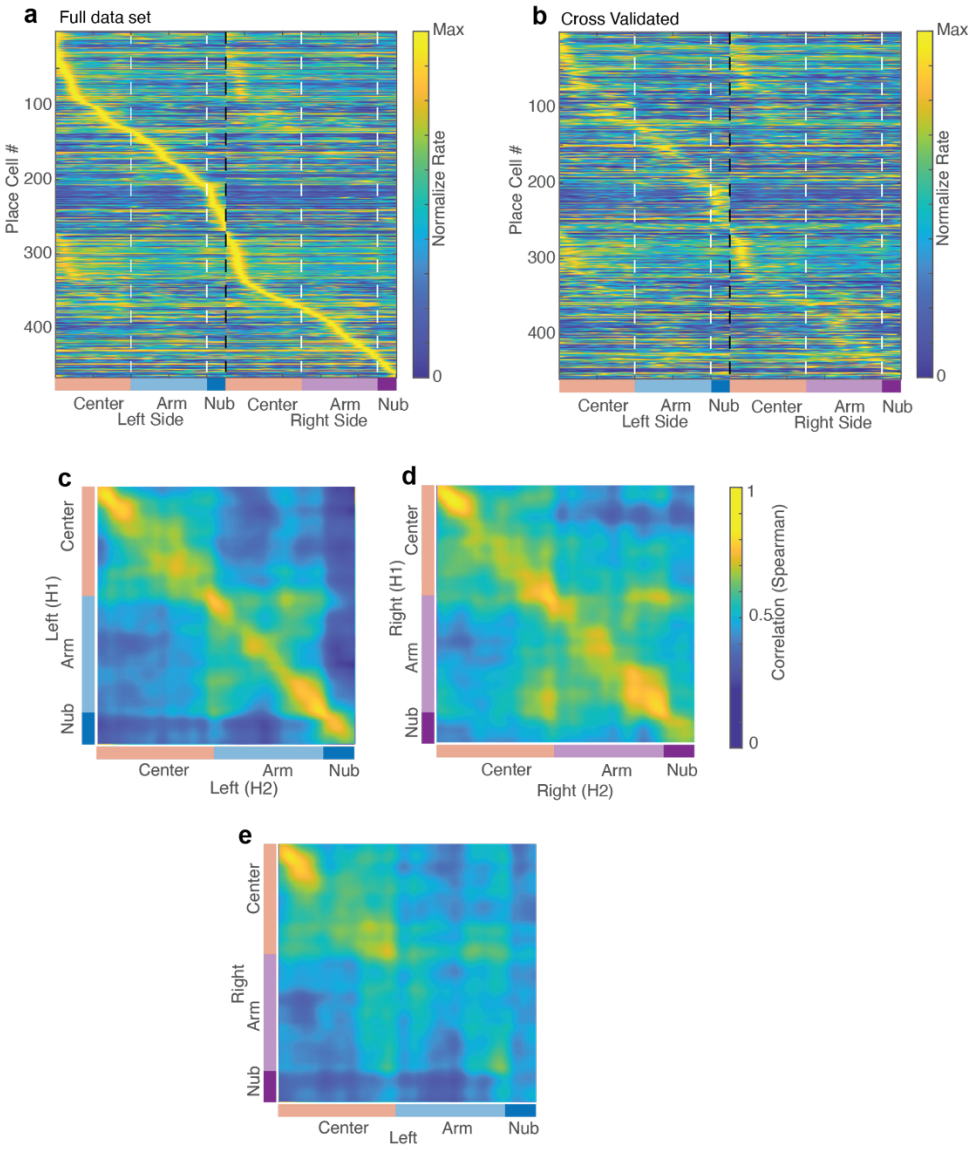

Figure 2-7 Population of Spatial Responses

(a) Population of spatially responsive, manually sorted cells. Plot is sorted by location of max response in the maze. Each row is normalized by the cell's maximum firing rate.

(b) As in (a), but cross validated (sorted on ½ session, with the other ½ session shown)

(c,d) Population correlation between the 1st and 2nd halves of population data for left (c) and right (d) going trajectories.

(e) Population Correlation between left and right trajectories.

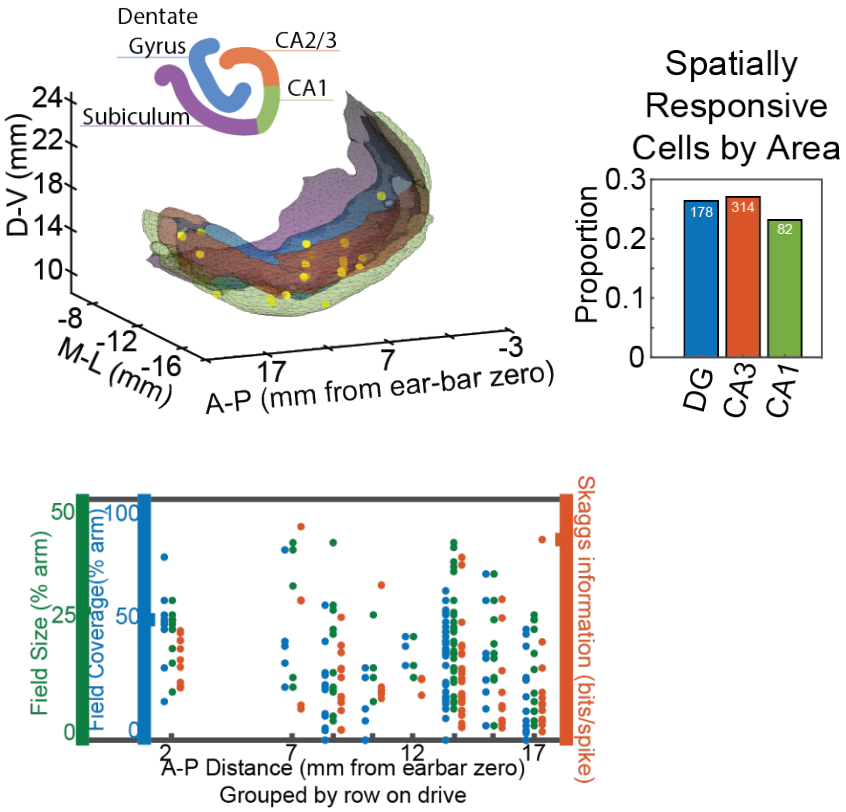

Figure 2-8 No relationship between Longitudinal Axis and Spatial Properties

(Top) Estimates of recording locations of spatially responsive cells from monkey SP, shown over the hippocampus proper reconstructed from an MR image. Dividing by subfield causes no substantive change in the fraction of spatially responsive cells, nor does distance along the longitudinal axis.

(Bottom) A variety of metrics on the spatial properties of hippocampal neurons show no difference along the Longitudinal (anterior=posterior) axis of the hippocampus.

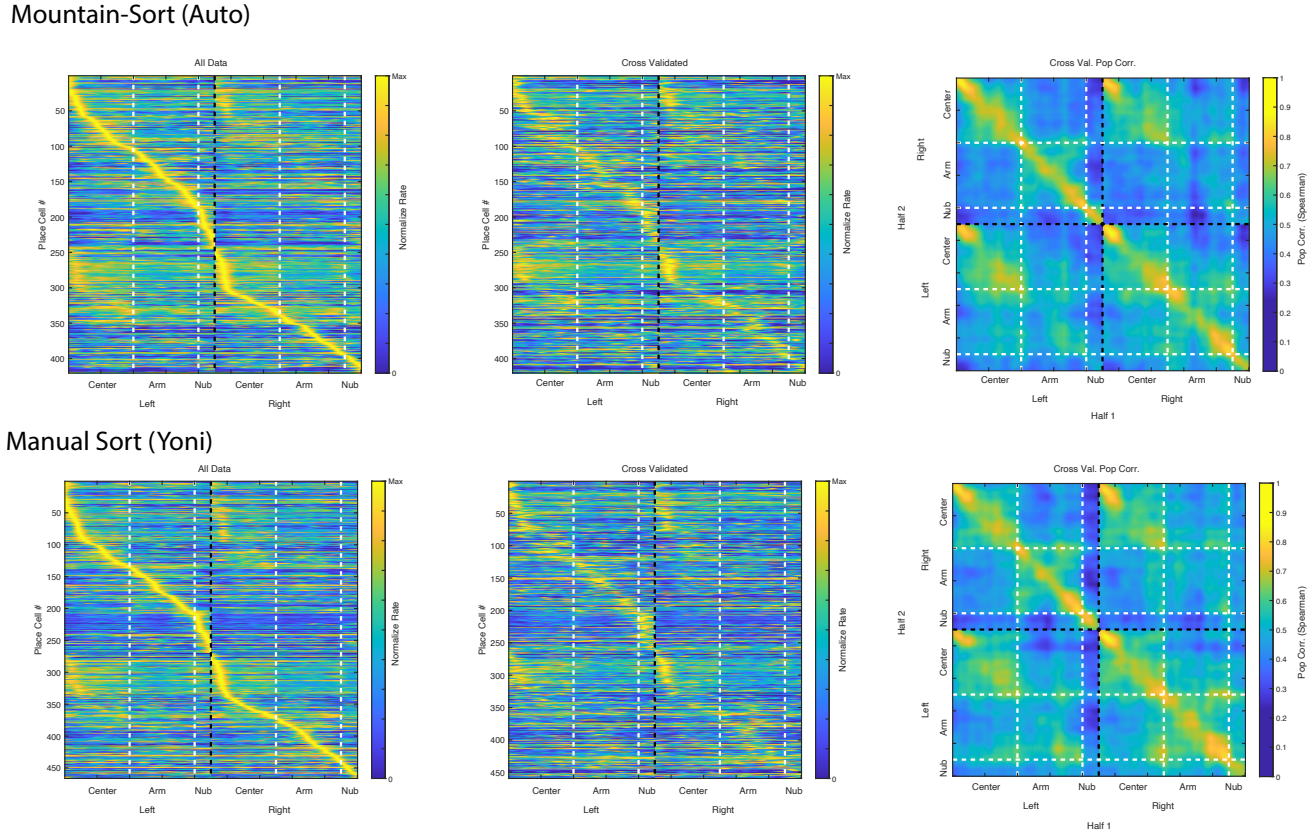

Figure 2-9 Spike sorting method does not change population response

(Top row) Automatically Sorted Cells

(Bottom Row) Manually sorted cells.

(First Column) Normalized rate-maps for spatially responsive neurons, sorted by location of peak rate)

(Second Column) As in first column, but cross validated (sorted on ½ data, other half shown)

(Third Column) Population Correlation of second column.

2.2.2 Responses to gaze location

Under certain circumstances, monkey hippocampal activity can correlate with the location of an animal's "view" or "gaze" either on a screen (Feigenbaum and Rolls, 1991; Killian et al., 2012; Konig and Buffalo, 2021) or in a room (Rolls and O'Mara, 1995). It is reasonable to consider the possibility that the activity seen thus far could be explained as a function of the animal's gaze; the constrained nature of our task means that an animal's location in virtual space is heavily correlated with virtual gaze. We used an infrared eye tracker to estimate the location of the monkey's eye position on the screen (see Methods; Eye tracking). We then projected a ray from this screen location into the 3-dimensional model of the virtual world to estimate the location of the monkey's virtual gaze. This process is illustrated in Figure 2-10.

Of 1268 units tested, spiking for 270 (20.5%, taken from 2 monkeys within 1mm of hippocampus) of these units carried significant and reliable information about the instantaneous location of the animal's gaze in the virtual world. It is important to specify instantaneous information here, because it is not clear that hippocampal activity should be a direct reflection of the exact location of the animal's gaze. The hippocampus is not classically thought of as a visual area, but rather a recipient of processed visual information (Felleman and Van Essen, 1991). If information takes time to reach the hippocampus, we would expect to see a delay between when the monkey looks at a location or object and when a cell fires in response. Conversely, if the hippocampus participates in prediction or planning (Stachenfeld et al., 2017), we might expect to see hippocampal firing in advance of visual behavior. To test these possibilities, we systematically shifted each spike train relative to the animal's visual behavior (see Figure 2-11) and reevaluated the correlation. Figure 2-12 shows examples of two gaze selective neurons, each of which is best correlated with behavior about 200ms prior to the actual spike time. Overall, we find that gaze-related activity in the hippocampus best correlates with the animal's eye position in the near past. As is illustrated in Figure 2-13, this delay in response varies between cells but falls largely within 0-250 milliseconds. Since the average timescale of a foveation in our virtual reality task is between 200 and 250 milliseconds for our monkeys, this suggests that the processing of visual information of gaze information is typically constrained within a foveation. Figure 2-14 shows examples of additional visually responsive neurons. Here, firing fields for each neuron are identified by noting which regions in visual space the cell

fires with rate greater than expected from shuffled data (95th percentile, 1000 shuffles). Foveations within each field are indicated for each cell. See Chapter 4 for further analysis of activity related to the onset of foveations.

Because gaze and spatial position are so highly correlated, we next sought to determine which behavioral feature is a better predictor of hippocampal neural activity. To do this, we turned to a Generalized Linear Modeling approach (Gareth James et al., 2014; McCullagh and Nelder, 1989), giving us a consistent set of metrics with which to compare the ability of each model to predict neural activity (See Chapter 4). Specifically, we used bootstrapped cross validation (see Section 5.4.3). We used two visual models; (1) parameterizing the monkey's view using a basis built along walls of the Y-maze and (2) using a basis consisting of the set of objects (3d-models) used to construct the Unity scene. Space was parameterized using a set of gaussian b-splines fit over a 1-dimensional linearization maze. Figure 2-16 shows the likelihood of a cell being well fit by one or more of these models with zero time offset. Note that the precise fraction of cells in a given category may differ slightly from the information theoretic analysis above due to use of a model-based test. Although these analyses identified about 21% of hippocampal cells with firing predicted by a view model, only a small fraction of these cells were not also significantly predicted by a spatial model. The object model was less successful, and all cells that were predicted by the object model were also predicted by the view model.

We next compared the predictive ability of a parameterized view and space. We did this both with zero shift in time and allowing for the shift that demonstrated by maximal visual information for each cell. It is worth noting that the validity of our shifting analysis relies on a degree of randomness in the monkey's visual behavior; because our monkey is locked into a track system, an equivalent time-shift analysis in space would simply reposition any spatial activity. We therefore shifted spike trains in both our spatial and visual models to the "best" time for visual information. As seen in Figure 2-17, almost all cells are better described by the spatial model even in cases where view provides an effective prediction. This is true even when time-shifting is allowed to optimize visual information, although a small population for which view is the only significant predictor exist.

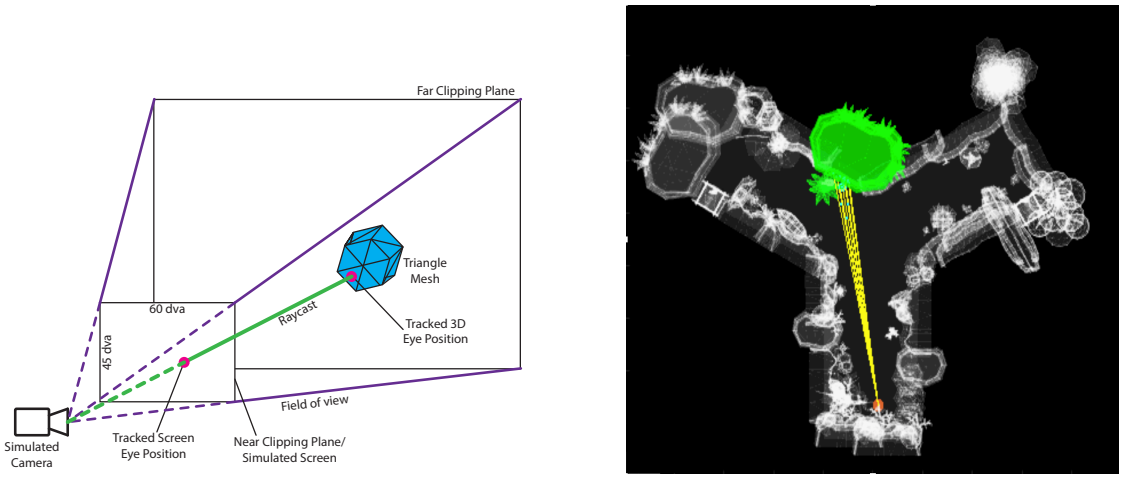

Figure 2-10 Schematic of ray-casting procedure

(Left) Illustration of post-hoc ray-casting. The monkey's tracked eye location is used to project a ray from the camera's near clipping plane into the 3D virtual space.

(Right) Top-down view of raycasting in 3D virtual space. Orange dot represents the monkey's avatar, while the yellow lines represent his field of view. Shown is an approximation of a fovea, constructed by casting multiple rays with a 2 degree-visual-angle radius. Model objects shown in green are those that the monkey is foveating.

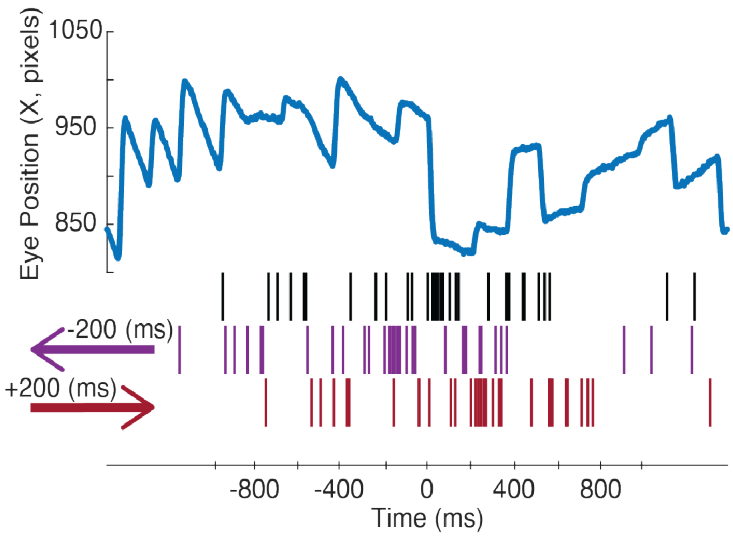

Figure 2-11 Illustration of shift analysis

Spike trains are shifted forwards and backward in time relative to the eye-movement trace.

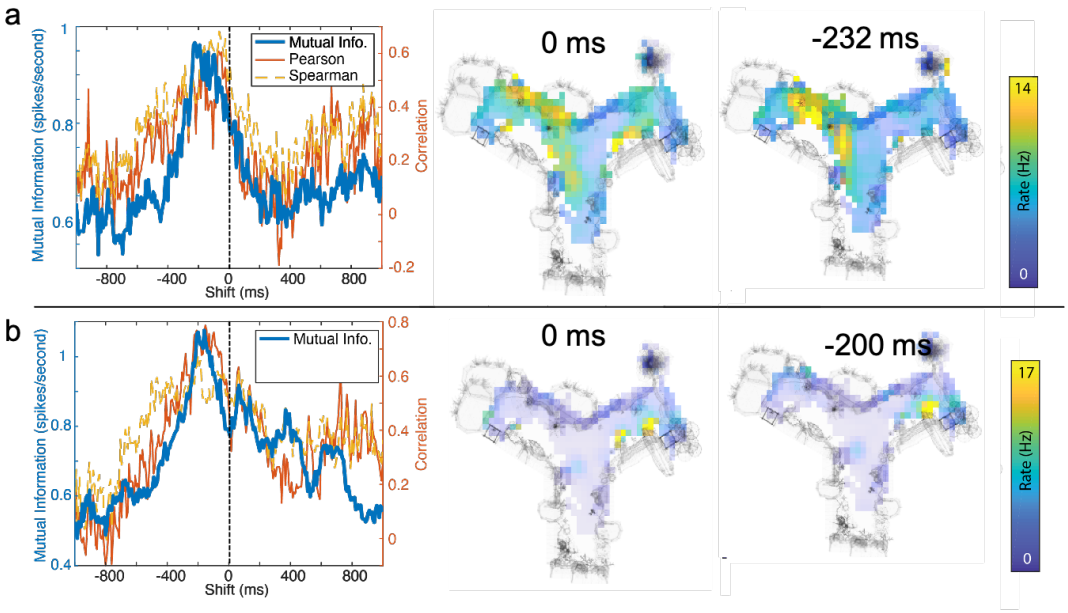

Figure 2-12 Examples of gaze correlated neurons

(a, b) Each row represents a single cell.

(Left) Mutual information for incremental shifts of the spike train relative to time (blue) peaks when the spike train is compared to the monkey's behavior in the past. Correlation between 1st and 2nd half of the spike train peaks at a similar offset.

(Center) Rate-map of firing with respect to the monkey's gaze at 0 offset.

(Right) Rate-map of firing at offset with maximal information.

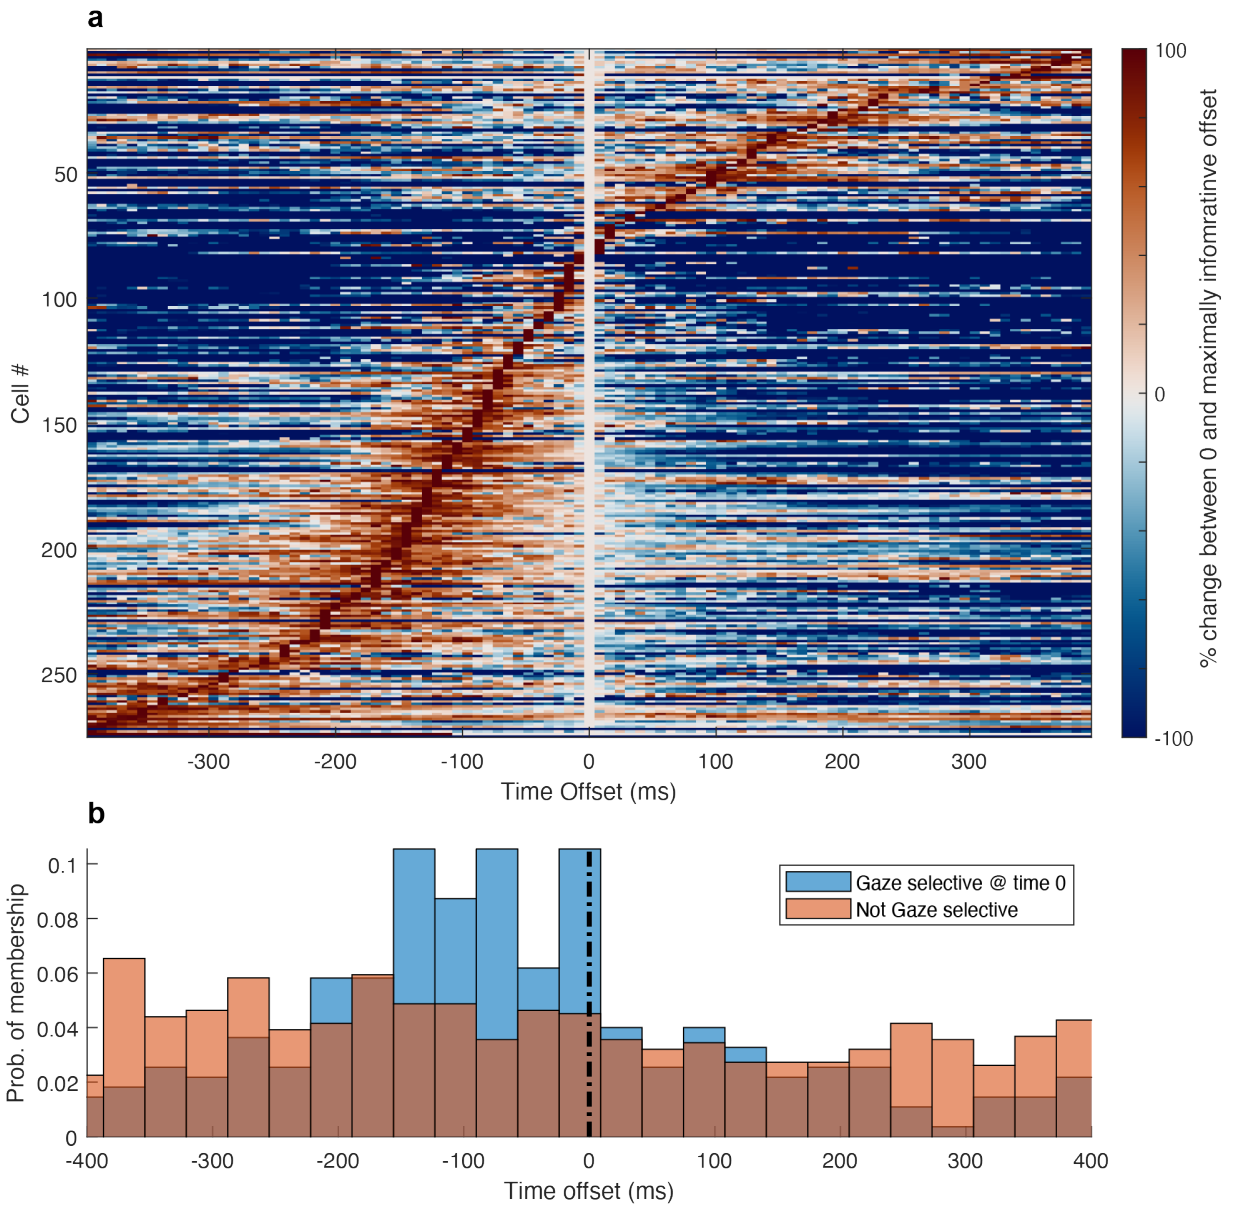

Figure 2-13 Visual response delay in the hippocampal population

(a) Population quantification of time-shifted information. Each row represents a single cell that caries significant visual information at time 0 (no shift), normalized by the difference in information between time 0 and the maximum for that cell. Rows are sorted by time with maximal information.

(b) Histogram of peaks in (a) is shown in blue. For comparison, the histogram of peaks for non-visually informative cells is shown in orange.

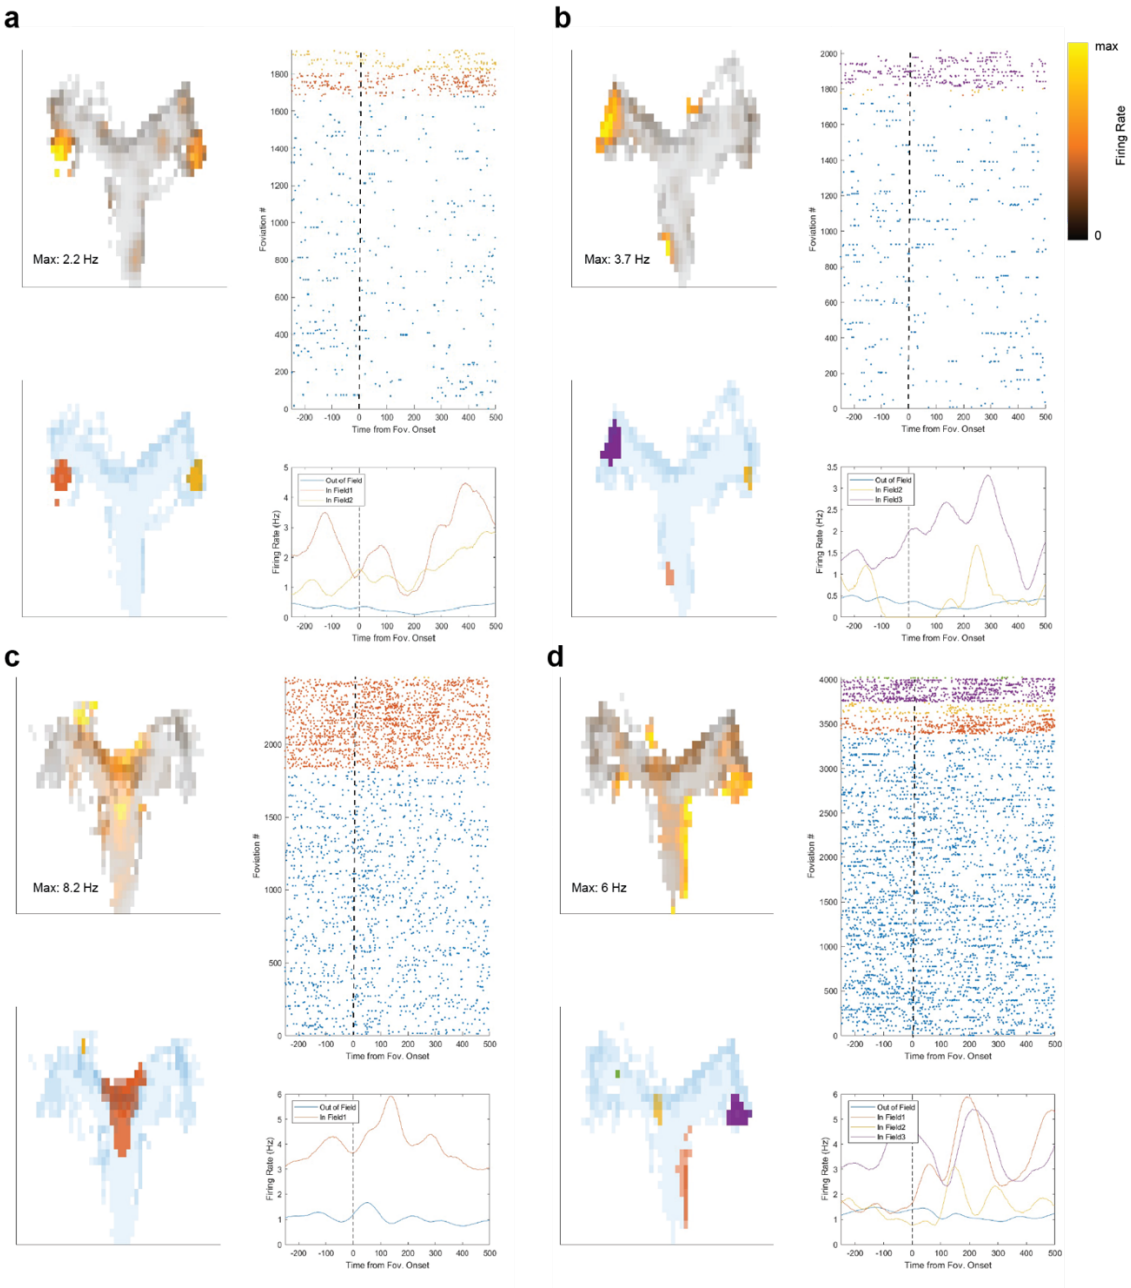

Figure 2-14 View-selective neurons with foveation triggered raster plots

(a-d) View responses of individual hippocampal neurons

(top left) Visual rate map for each neuron

(bottom left) Continuous regions where the cell fired greater than expected by chance (i.e., the cells firing field; 95th percentile, 1000 shuffles).

(right) Raster plot and average rate for each neuron aligned to foveation onset. Foveations within a given field and color matched to the field. Extra-field firing is shown in blue.

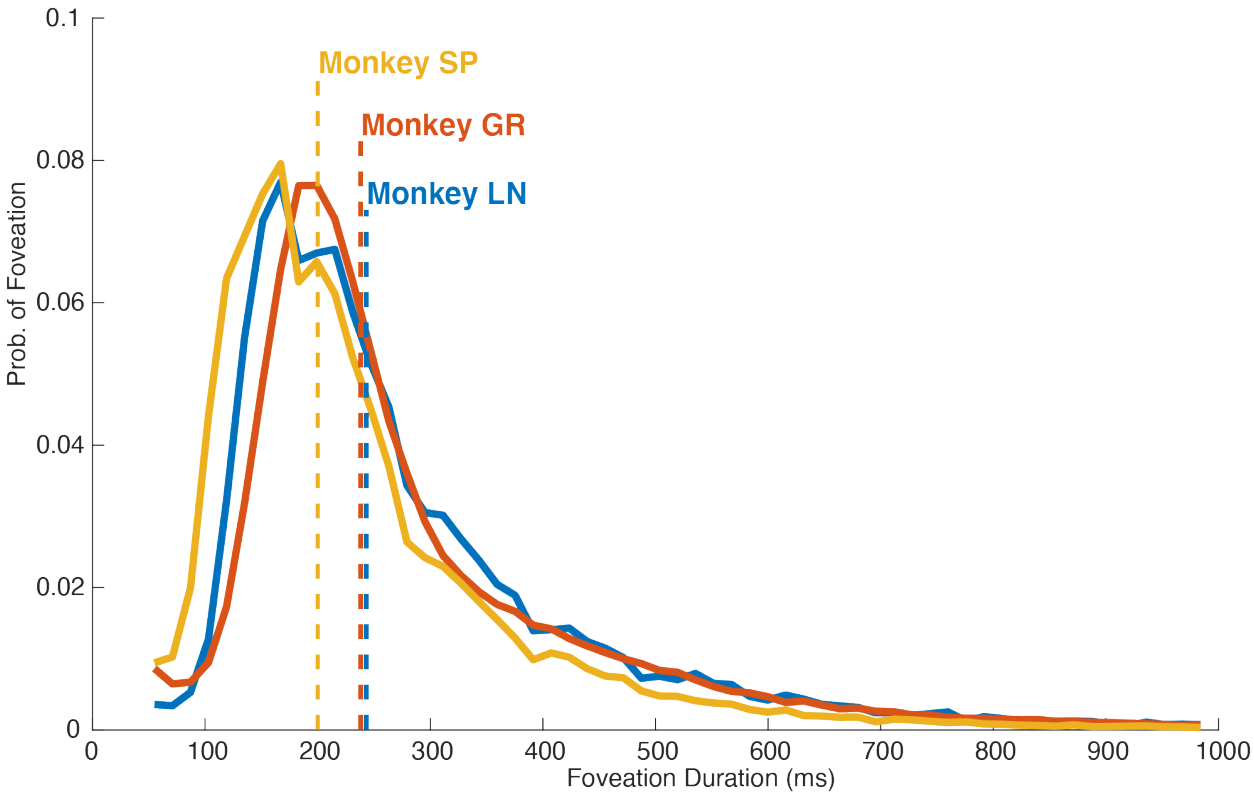

Figure 2-15 Distribution of Foveation Durations

Distribution of foveation durations during the Y-maze task using one month's data from each of our 3 monkeys. Solid lines show the probability density for each duration; dotted lines give the median for each distribution. All fixations less than 55 ms were discarded as likely tracking errors.

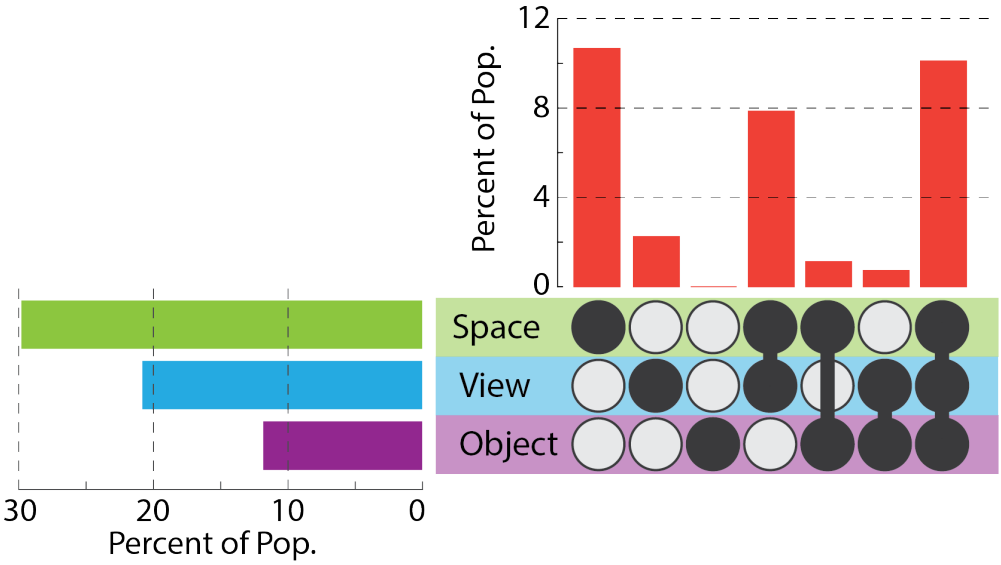

Figure 2-16 Most view cells are also spatially responsive

Fractions of hippocampal cells with firing predicted by each model type. Note that “view” and “object” are both parcellations of visual space.

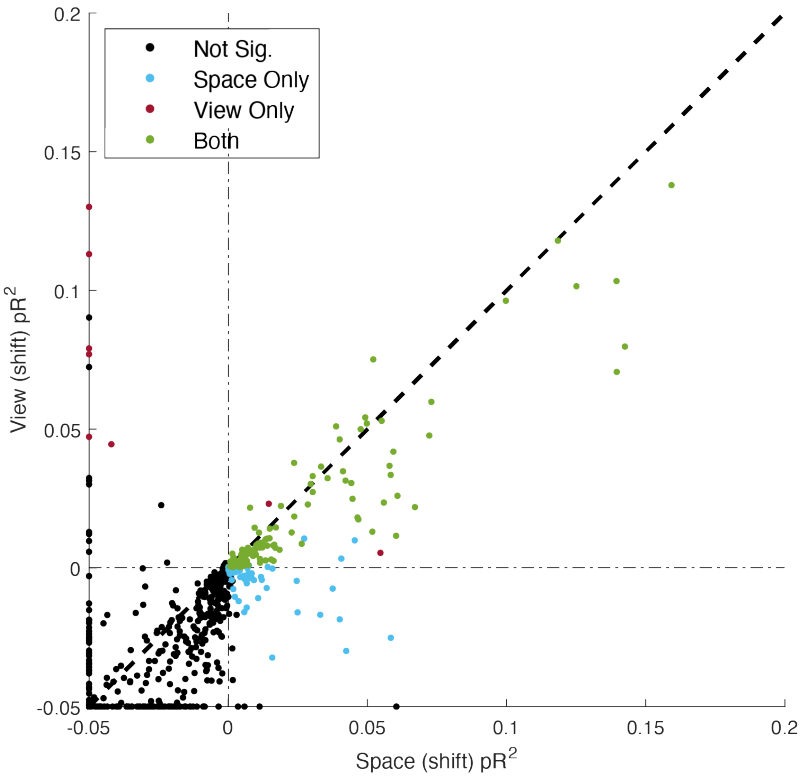

Figure 2-17 Comparison of View and Space Models

Each dot represents the response properties of a hippocampal cell, while axes compare the median predictive ability of our space and view models. All cells are shifted to their peak view-response delay. In almost all cases, space provides a better description of a cell's activity, but this is more-often-than-not correlated with visual information.

2.2.3 Responses to task phase

Because spatial responses seem to be largely driven by salient task events, we next asked whether non-spatial task features were similarly correlated with hippocampal activity. To do this, we looked explicitly at a neural activity during reward delivery and during the inter-trial-interval. A subset of recorded neurons show activity locked to the collision of the monkey with the target banana and subsequent onset of food reward. Figure 2-18 (first row) shows examples of two such neurons, while Figure 2-19 shows a summary of these population responses for the nub-reward period for a subset of our data (manually sorted cells). Similarly, Figure 2-18 (second row) shows examples of two neurons that were active during the four second inter-trial interval.

Rather than separately identify cells with responses to each of these epochs, we chose to analyze the Y-maze as a continuous “trial” space. This was achieved by rescaling each identifiable task epoch (driving down an arm, turning, receiving reward, delay, etc.) to occupy this space in proportion to the average time the monkey spends in each segment. Normalized rate-maps plotted over the task are shown in Figure 2-19 (a), while Figure 2-19 (b-c) shows the correlation matrices for these data within and between left and right trajectories. The “driving,” “reward,” and “delay” sections of the maze are all eminently distinguishable from one another, implying that the population activity broadly distinguishes between these task segments (again, see Figure 2-6 for interpretation). Notably, the end of the “drive” segments (when the monkey is turning towards the banana and driving down the nub) show relatively low correlation between the left and right sides of the maze, whereas the “reward” and “delay” portions of the task have similarly high within- and between-side correlations. This suggests that, although units active during the navigation portions of the task may show trajectory specific (i.e., spatial) firing, units active during reward and delay share common activity between trajectories.

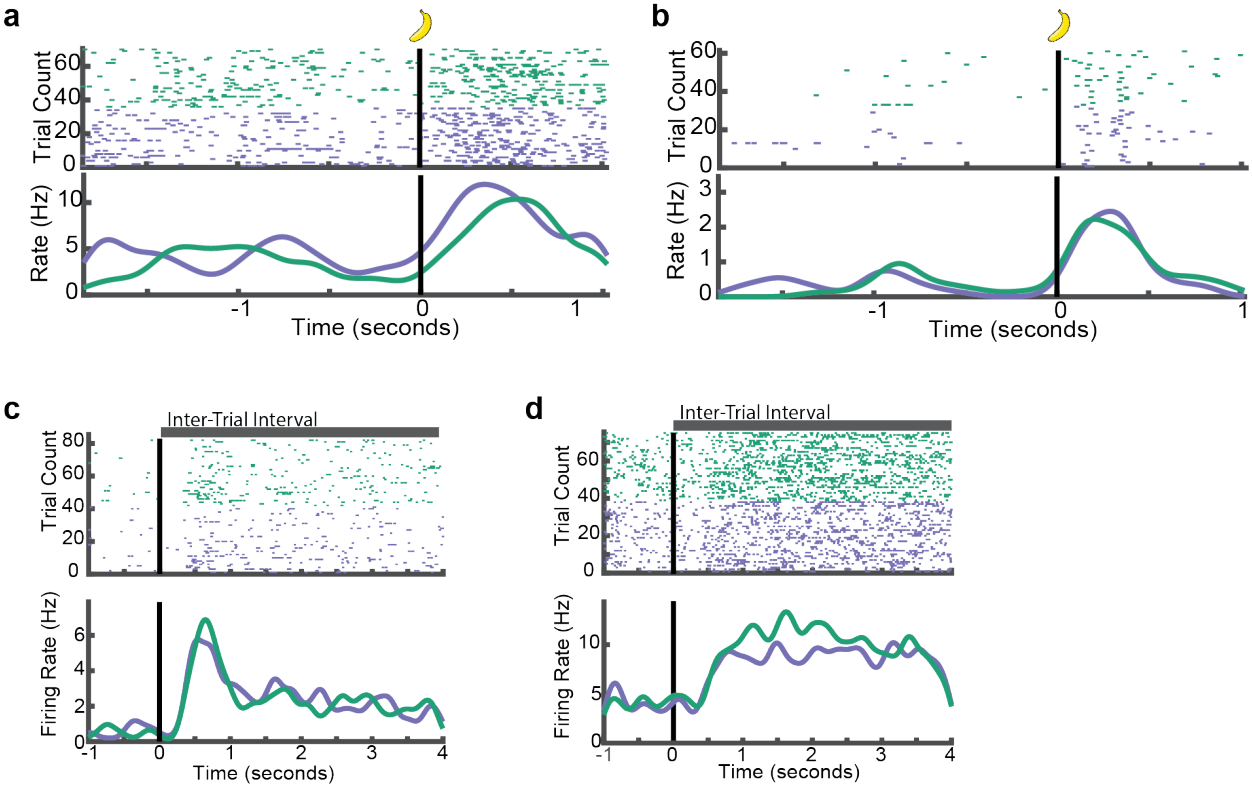

Figure 2-18 Example Raster Plots from Task Responsive Neurons

(a-b) Example reward selective neurons. Colors indicate right-(green) or left-(purple) ward going trials.

(c-d) Examples of neurons active during the inter-trial interval.

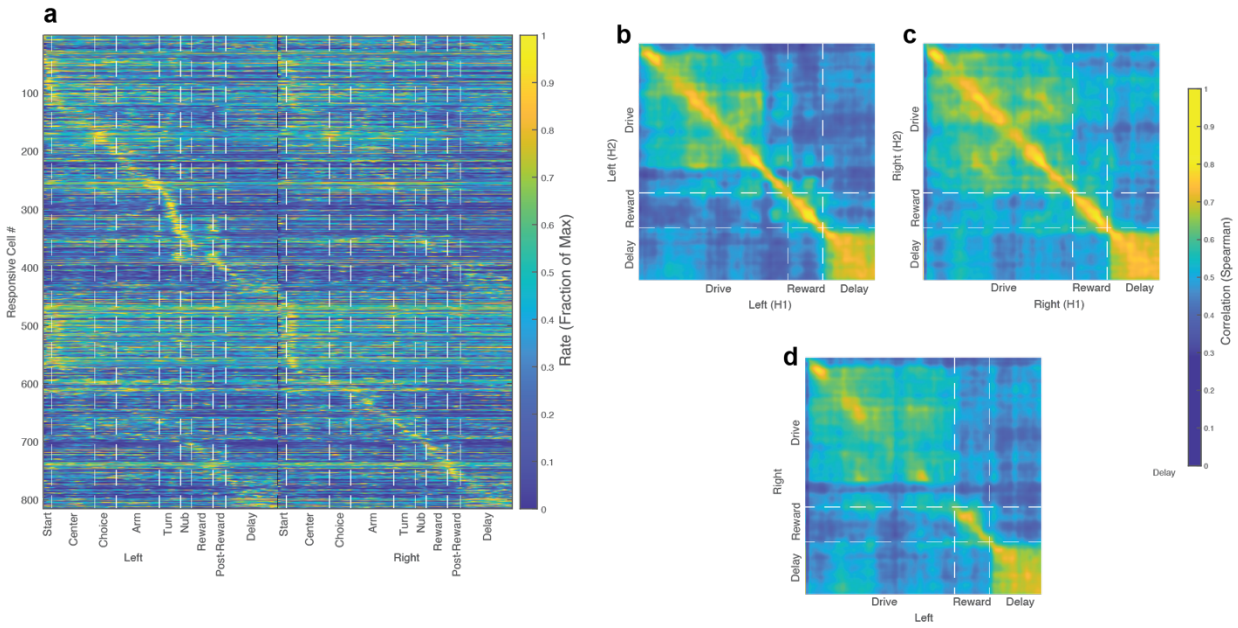

Figure 2-19 Population activity segments task phases

(a) Cross validated population plot for all neurons selective for some part of the task.

(b, c) Correlations between 1st and 2nd halves of population data on the right (g) and left (f) sides.

(d) As in (b) and (c), but for the cross correlation between the right and left sides.

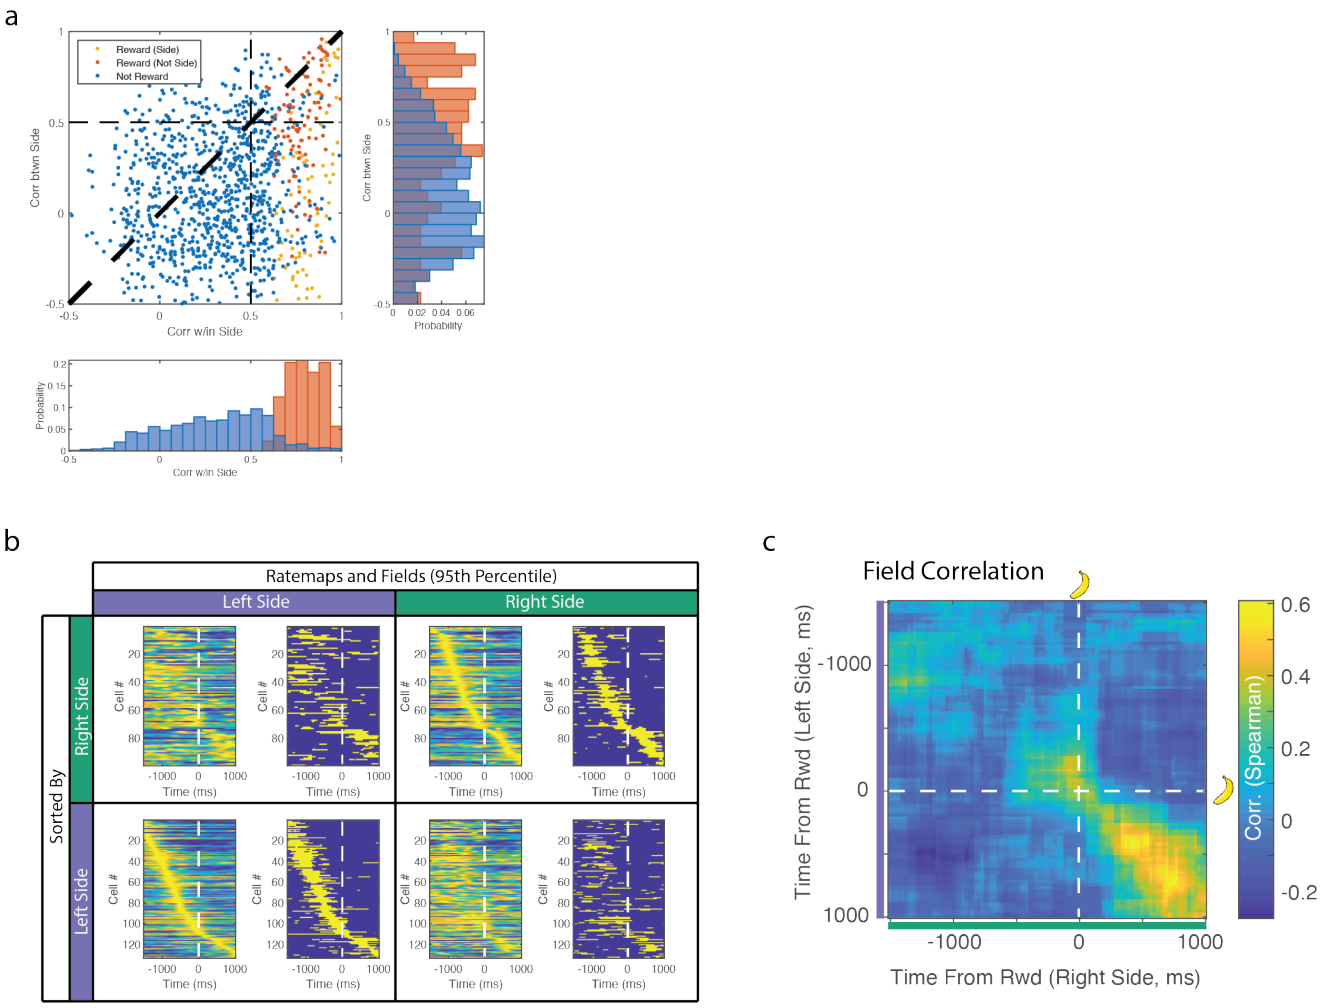

Figure 2-20 Side specificity of reward

(a) Comparison of within- and between- side correlation for cells active during the nub/reward period. Side panels show correlations for cells responsive during this epoch (orange) and all other cells (blue). In the scatter plot, responsive cells are divided between side selective (yellow) and non-side selective (orange).

(b) Top: Cells with significant responses on the right side, sorted by max firing rate. Fields (where rate is greater than expected by the 95th percentile of shuffles) are also shown. Bottom: Same, but for cells with left side responses. Note that many cells are significant on both sides (see a) and are shown twice.

(c) Correlation of isolated field for responsive cells between sides. Activity >1/2 a second before the banana contact have low correlation indicating a high degree of side specificity, with population responses most similar immediately before and during reward consumption.

2.2.4 Responses can be stable across visual contexts

To better differentiate visual- and task-driven activity, we changed the visual appearance of our virtual Y-maze while keeping all other aspects of that task the same. Where possible, we used an A-B-A presentation format for these visual contexts (Figure 2-21), so that we could (1) differentiate responses between environments and (2) confirm that a given cell held its relationship to a particular environment throughout a recording day. We used a total of 7 distinct virtual environments, meaning that A and B were not necessarily the same across days. A subpopulation of our recorded hippocampal neurons (37/214 neurons with the opportunity to remap) showed significantly different activity between different visual contexts under a gain-in-information-with-partition test[7] (see Section 5.2 for a discussion of this test, and Section 5.5 for a comparison to other remapping methods). Examples of this activity are shown in Figure 2-21 (b, c). The remaining 83% of cells had consistent responses between visually distinct environments, further indicating that the majority of cells can be better defined as task-responsive than in terms of specific visual responses. Examples of this activity are shown in Figure 2-21 (d, e).

If is not entirely clear what differentiates a remapping cell from a non-remapping cell. Spatial responses for remapping cells tend to carry more information about the animal's position (see Figure 2-22). Similarly, they are more likely to be correlated with the animal's gaze location in a particular environment, as seen in Figure 2-23. This presents the possibility that there exists a subset of highly selective, visually-driven cells that is distinct from the more broadly tuned population of task-responsive cells.

[7] Note that this analysis included only the 3 primary arms of the maze, not the full navigable space (i.e. the arm nubs were not included, nor were the other task events). Further analysis on these other maze sections will take place prior to submission for peer review.

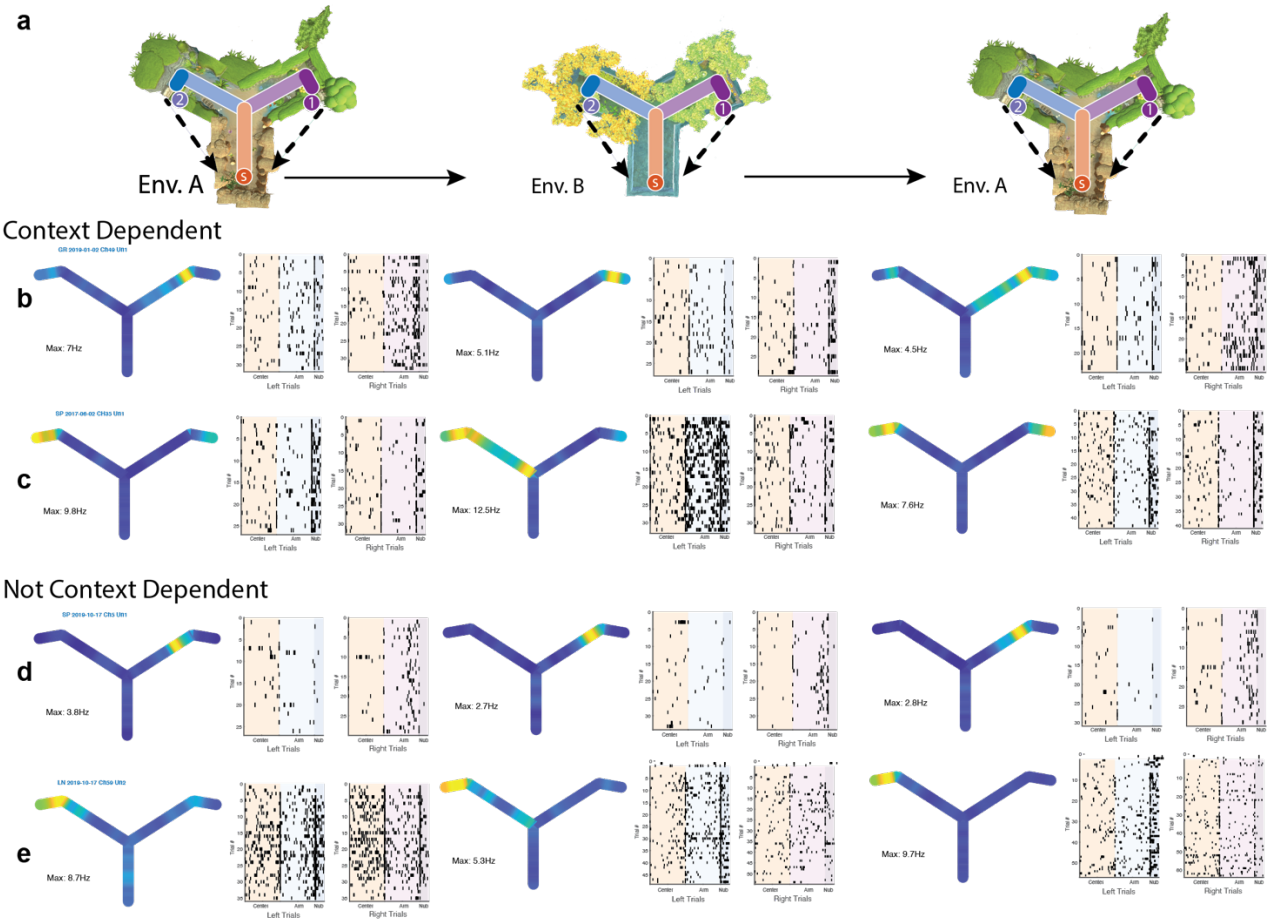

Figure 2-21 Responses can be dependent on visual context

(a) A-B-A Behavioral paradigm. Monkey performed the same task in two visually distinct environments.

(b-c) Context-dependent cells change their firing responses between sessions in different visual contexts, but show consistent responses between different sessions in the same visual context. Two example cells are shown, one in each column.

(d-e) A majority of neurons show consistent responses regardless of the visual environment. Again, two example cells are shown.

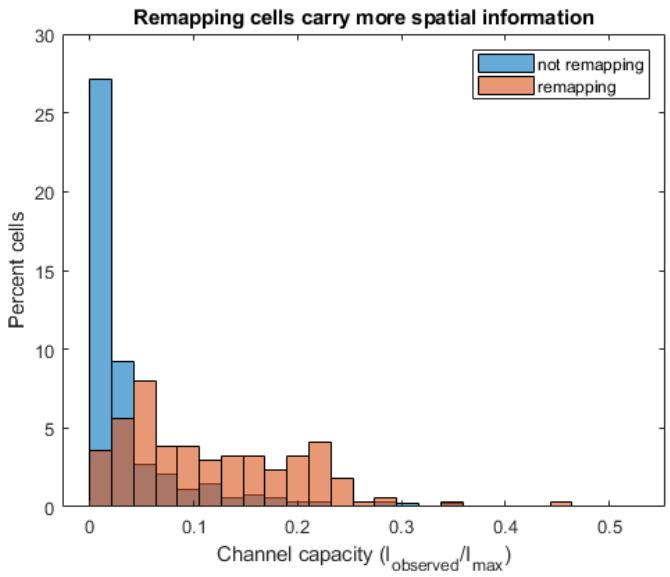

Figure 2-22 Remapping cells carry more spatial information

Histograms are shown in units of observed Skaggs Information divided by the maximum channel capacity.

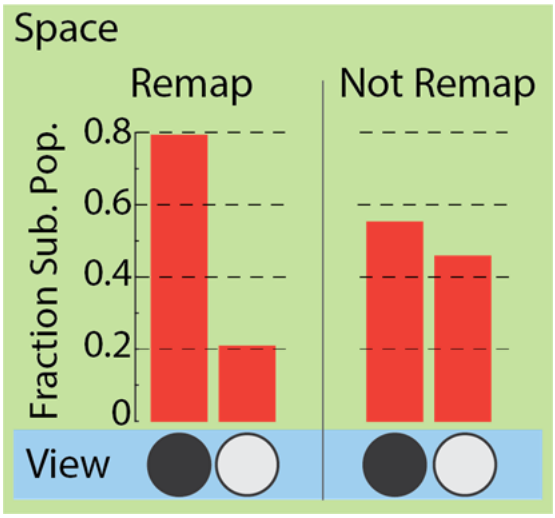

Figure 2-23 Remapping cells are more likely to be predicted by the animal's gaze.

Fractions of remapping and non-remapping spatial cell (bootstrapped GLM) that do (black) and do not (white) carry significant visual information.

2.3 Discussion

We are not the first to demonstrate spatial activity in the monkey hippocampus. Although seemingly more difficult to elicit than in rodent studies (Rolls and Wirth, 2018), spatial specificity has been observed in the hippocampus of freely moving monkeys (Courellis et al., 2019; Hazama and Tamura, 2019; Ludvig et al., 2004b), monkeys constrained on track systems (Nishijo et al., 1997b), and monkeys navigating in virtual reality (Baraduc et al., 2019; Gulli et al., 2020; Hori et al., 2003; Wirth et al., 2017). Strikingly, while the cells in these studies do carry spatial information, they often lack the punctate responses that are the hallmark of rodent spatial representations (McNaughton et al., 2006). Although some of this difference may be induced by the differential recording techniques used in monkeys and rodents (Chung et al., 2017b; Skaggs et al., 2007a), similar results from other species (Ekstrom et al., 2003a; Payne et al., 2021) imply that this phenotype of spatial activity is not unique to monkeys nor is it solely a result disparate recording techniques. While the cell isolation and yield of high-density probes are still needed to fully understand population activity in the monkey hippocamps, our data provide a survey of hippocampal activity and allow us to query the structure of monkey hippocampal responses in a manner not previously possible.

A distinction must be made, however, between spatial activity related to an animal's gaze and activity related to that animal's physical location. When a monkey is exploring a room (Rolls and O'Mara, 1995) or a screen (Feigenbaum and Rolls, 1991; Georges-François et al., 1999; Killian et al., 2012; Konig and Buffalo, 2021; Meister and Buffalo, 2018; Rolls and O'Mara, 1995; Wilming et al., 2018) neurons in the hippocampal formation encode the animal's gaze location. This is in striking contrast to our result that neurons encode the location of the monkey's virtual avatar, irrespective of where he is looking on the screen. Even if we consider where the animal is looking in the virtual environment, the vast majority of neurons are better described by the monkey's avatar's location and many do not change their spatial specificity even if the visual scene is changed. How, then, do we reconcile this seeming disconnect? Perhaps the most parsimonious explanation is that hippocampal responses are bounded by the reference frame of the task at hand. Others (Feigenbaum and Rolls, 1991; Meister et al., 2013) have shown that neurons in the hippocampal formation can maintain their spatial relationships to a screen if that screen is moved around the room, while changing their responses if the same screen is used with a different task

(Gulli et al., 2020; Konig and Buffalo, 2021) implying that ostensibly spatial responses are really a reflection of task relevant information in each of these situations. Our result can be seen as an extension of these past findings; here, our virtual navigation task itself drives hippocampal activity. This can include visual information about the task, but it is not limited to gaze-specific information. Thus, spatial correlates in this and other virtual reality paradigms can be thought of as reflections of the task at hand, rather than a necessary byproduct of spatial mapping being a core hippocampal function.

Moreover, our data suggest how such task-specific information is organized. Spatial responses cluster around key task events, including the beginnings and ends of the maze. Further, the population of responses seems to divide relevant segments of the task, even if specificity within any particular segment is lower than might be expected from rodent work. This is true both in the spatial domain and at the level of task phases. Because these responses were not simultaneously recorded, we are limited in our ability to explicitly probe the “decodability” of spatial location or task progress in our hippocampal population. However, these data suggest that task phase can be broadly deciphered by hippocampal activity, even if explicit location or timing cannot. Again, while this finding lacks the high spatial and temporal specificity reported in some rodent work (Eichenbaum, 2014; Kraus et al., 2013; Pastalkova et al., 2008), it is in keeping with task specific firing in other monkey studies (Vidyasagar et al., 1991; Watanabe and Niki, 1985) as well as in rodent studies with discernable transitions between task phases (Aronov et al., 2017; Terada et al., 2017).

Task information is not the sole variable encoded in the monkey hippocampus. Just as others have reported gaze correlates in the monkey hippocampus, we see cells whose firing correlates with where the monkey is looking in our virtual space. The high correlation between view and space, however, convolutes this result. On one hand, the striking delay in visually-correlated spiking suggests what the monkey is looking at plays a meaningful role in shaping hippocampal activity, possibly by delivering task information from the most recent eye movement, for at least a subset of cells. On the other hand, view is rarely the best predictor of neural activity, suggesting that apparently visual correlates might be better discussed in terms of other behavioral features. Since our task also plays out entirely in a virtual space that must be visually explored, it is further possible that at least some of visual responses are a reflection

of task information reaching the hippocampus, rather than an explicit spatial “mapping” of the monkey’s gaze location in the virtual world.

Moreover, many of our hippocampal cells do not change their responses when the visual environment is changed. This suggests that for these cells, the visual appearance of the virtual task is not what is driving hippocampal activity. Since task structure is held consistent between environments, this result provides further evidence that task structure is shaping this firing. Not all cells are consistent between environments; a notable minority “remap” in the new visual context. Since the vast majority of these cells are visually responsive in at least one environment, this result suggest that task is represented by a subpopulation of neurons distinct from those that explicitly represent gaze location.

While we have described what might drive monkey hippocampal activity, it remains unclear how these representations might contribute to hippocampal function. The hippocampus has a well-established role in the formation of new episodic memories (Scoville and Milner, 1957) and recent work has focused on how spatial representations might contribute to this mnemonic function (Bellmund et al., 2018; Whittington et al., 2020). By showing that the hippocampus parcellates experience into task-relevant segments, our work supports the possibility that the hippocampus contributes to memory by tracking ongoing features of experience.

2.4 Methods

2.4.1 Y-Maze task

Monkeys were trained to use a joystick (a modified and ruggedized Logitech Extreme 3d Pro) to navigate a virtual spatial-delayed-alternation-task (Y-maze) custom built in the Unity game engine (2017.f1a, .Net framework 2.0 complete for serial compatibility) and running on purpose build computers (Windows 7/10). Monkeys viewed the maze from the perspective of a first-person avatar (45 dva field of view), projected onto a large (127x97cm) screen at 120 Hz. Their movements were locked onto a track system in this maze with turning only allowed at specific locations. The only deviations from this track were to smooth the monkey’s perspective if his turns did not perfectly align with the track system.

Monkeys were given food reward for driving into bananas in the virtual maze. On each trial, a single banana was hidden in a hut at the end of the side arm. Each trial started in the center stem of the maze, and on each trial the monkey had to drive down this stem, choose a side arm, drive down that arm, then

turn into the hut to get reward. The correct choice on any given trial was always the direction not chosen on the previous trial (with the exception of the first trial in a session, for which the correct answer was always to the right) and the monkey could not see the banana until he began his turn into the hut. After the trial, monkeys were teleported back to the start of the center stem with a 4 second delay (dark screen). On incorrect trials, this teleportation happened without turning into the hut at the end of the arm. The monkey's day was divided into sessions, each consisting of about 20 minutes of game play with a break of approximately 10 minutes in between. For the context-dependence experiments, the second session used a virtual environment with a different appearance from the first and third blocks, with all other task features held constant. Each environment was made by arranging assets from different packages purchased from the Unity Assets store.

2.4.2 Training

Monkeys were briefly rewarded for interacting with the joystick in their home cage. The joystick was then affixed to that monkey's behavioral chair, and the monkey was brought into a behavioral booth. Here, he was trained first to interact with the joystick, then to center a target banana on the screen (to turn towards and stop when facing a banana) and to drive towards a centered banana. Once proficient, the monkey was placed in a "free foraging" task where he had free motion over a field populated with randomly distributed bananas. His task was to drive around the field, collecting each banana. The field would repopulate once all bananas were collected. Each monkey was monitored to confirm he was routinely turning both left and right, and side biases were addressed by a return to the "target to center" task. Monkeys were deemed proficient once they could consistently clear a field of 10 bananas for several >20 min behavioral sessions in a day.

Once each monkey was proficient in the foraging task, he was placed in the "classic" Y-maze environment. Initially, a "Lure" banana was situated at the end of the center stem, with "Side" bananas placed such that they were visible from the central choice point. After each correct trial, the "Side" banana would change sides. Once the monkey became proficient in collecting these bananas, the "Side" bananas were systematically moved down the arms of the maze and into the hut, and the "Lure" banana was removed. All monkeys were proficient (>90% accuracy, Figure 2-2) in the final Y-maze task before their drives were implanted.

2.4.3 Eye tracking

We tracked each monkey's right eye using ISCAN infrared eye tracking (244 Hz-1000Hz, depending on session) with a designated computer (Windows XP). Pupil position (Horizontal and vertical) was sent via serial port (115200 Baud) to the Unity computer (https://www.alanzucconi.com/2015/10/07/how-to-integrate-arduino-with-unity/). A simultaneous copy of this signal, as well as an estimate of pupil diameter, was converted to +/- 5 V analogue signal and sent to the Neuralynx Cheetah system. These replicated signals were used for system alignment (see below). All behavioral analysis used the Unity (digital) version of the eye signal.

2.4.3.1 Calibration

During each recording session (typically at the beginning), monkeys performed an eye calibration task which was later used to relate eye-signal values to two-dimensional locations on the screen. Each monkey was trained to fixate a gray square, and to hold fixation until that square changed to a yellow color. The color change took place with a random timing uniformly distributed over 0.5-1.5 seconds. Squares appeared on a 5x5 grid evenly spaced over the screen. Monkeys also performed an image viewing task interspersed with these calibration trials, although these data are not discussed here.
We related eye position to screen position using a two-dimensional, second order polynomial interpolation (A 2nd order polynomial was used to approximate the apparently spherical errors in tracking apparent at the edge of our large screen). We achieved this correction by assuming the "true" location the monkey was looking at was the target on the screen, t, and relating it to the value recorded by the eye tracker, s, by solving for a fixed set of parameters, A and B, as follows:

$$x_t = A_1 + x_s A_2 + y_s A_3 + x_s y_s A_4 + x_s^2 A_5 + y_s^2 A_6$$

$$y_t = B_1 + x_s B_2 + y_s B_3 + x_s y_s B_4 + x_s^2 A_5 + y_s^2 A_6$$

These equations were solved in real time using the Math.Net package into Unity, and post hoc using the "transform" and "inverseTransform" commands in MATLAB.

2.4.3.2 Three-Dimensional Ray-casts

We used each day's calibration transform to approximate a mapping from eye-tracker coordinates to screen-viewing location coordinates on any given day. We then used the post-hoc ray-casting in MATLAB to identify where in the virtual world the monkey's gaze was centered. This procedure is shown in Figure

2-10. First, we identified the field of view of the monkey (in the virtual space) from his present virtual position. We then projected a ray into the virtual game-space at the angle described by the monkey's present on-screen eye position, the coordinates where this ray first makes contact with a game object, as well as the identity of that object. Ray-cast location and object identity were recorded once per eye-tracker sample, meaning that they can update more frequently than the framerate of the game. We checked our MATLAB ray-casting against locations and object identities found from ray-casting by replaying data through the Unity Engine and found them to be the same (up to the point of numerical error). We ultimately used the MATLAB instantiation of this algorithm due to the profound speed increase available from GPU-acceleration in MATLAB.

2.4.4 Data Logging

All Unity data were logged by asynchronously writing these data to text files. Task information was logged using Unity's Update timing loop, meaning that logging took place once per frame and was tagged with the timestamp at which it took place. Eye tracker information was logged with the timestamp at which it was received, although ray-cast transform and object identification only took place once per frame.

2.4.5 Electrophysiological Recordings

Each monkey was implanted with a 124 channel hyperdrive (Grey Matter Research) loaded with 100micron tungsten wires (FHC). Wires were arranged in a square grid with 1.25 mm spacing. Wide band (0.1-8000Hz) neural signals were sampled from all wires at 32kHz using a Neuralynx Cheetah system (software version 6.0, windows 7/10).

2.4.5.1 Spike Sorting

Wide band neural recordings were high-pass filtered at 450Hz. Candidate spike waveforms - those with amplitudes lower than 3.5 standard deviations below the signal mean – were extracted for further analysis. To maximize uniformity during sorting, waveforms were up-sampled by a factor of 4 using optimal sinc function interpolation (Oppenheim and Shafer, 2009; Shannon, 1949), aligned to their trough, then down-sampled back to the original 32kHz. Some of these data were then brought into Plexon's Offline Sorter software for manual cluster cutting using primarily the PCA and "slice" (template) views. Remaining data were too numerous to manually sort, and cells were identified by running our extracted

waveforms through the Isosplit automated clustering algorithm that is typically used in the context of the Mountain Sort data processing pipeline (Chung et al., 2017b).

2.4.5.2 Time synchronization

To synchronize our recordings with the behavior, we made use of the common eye signal split between our Unity and Neuralynx computers. At the start of each behavioral session the VR computer sent a text command via serial port to the Neuralynx computer. This signal was not consistent enough in timing to be used for exact synchronization. Instead, the first blink – a large, sudden discontinuity in the eye signal – after this text command was identified both in the Neuralynx and Unity data streams, and the start of this blink was used to synchronize both systems post hoc. Each behavioral session was synchronized separately, and time drift over the course of any given session was small enough to be unmeasurable.

2.4.6 Anatomical localization

Recording locations reported here are all pre-histology estimates. Hyperdrives were implanted in several steps, with separate surgeries for chamber and drive implantation. All monkeys underwent a presurgical MRI, as well as an MRI after the chamber implant surgery. This later MRI included a set of fiducial markers to provide an initial estimate of drive position and angle. The titanium chamber caused distortion such that this co-registration was insufficient on its own. As a result, drive alignment was confirmed using a combination of (a) post-surgical measurement (b) X-ray images taken in stereotaxic coordinates and (c) 3D scans of the monkey's face and drive.

Drive alignment was used to estimate starting positions for each electrode. Each time any electrode was turned, it was manually entered into a log and these values were integrated to estimate electrode position assuming straight trajectories. X-ray images were used to confirm that electrodes maintained generally straight trajectories, and we stopped turning any electrodes that strayed from a straight trajectory. Target structures were identified as volumetric STL files using 3D-slicer, an MRI analysis software package. Estimates of location relative to these structures was determined in MATLAB. To account for errors in alignment, many of the analyses presented in this chapter allowed 1-2 mm of error around the hippocampal STL. Changing this bound has only small impact on the number of cells included, and in no case did it induce a qualitative change in population activity.

2.4.7 Analysis

The analyses used to characterize the predictor relationships of single-cell spike trains are described in detail in Chapter 4. In the present draft, at least 1000 Monte-Carlo shuffles were used in computing any reported significance.

3 Navigation in Abstract Tasks

3.1 Free Foraging in Virtual Reality

From the beginning, one of the principal goals of my dissertation project has been to assess to what extent spatial responses exist in the monkey hippocampus. Our initial approach was to adopt canonical rodent tasks into monkey-friendly virtual reality, and then to record from the hippocampus of monkeys playing these games using a novel type of chronically implanted flexible electrodes. From the perspective of finding spatial responses in the monkey hippocampus, this approach failed completely. Piloting these flexible electrodes was more fraught than anticipated, and we were unable to find any spatial correlates whatsoever in these data. However, this pilot study was informative in designing our later studies and in shaping our thinking about how the hippocampus processes task information, and so I have included it here.

The quintessential method for recording rodent place cells is to do a “free foraging” task, in which food (chocolate sprinkles or the equivalent) is randomly distributed around a box or room, and animals run around collecting this food (Cells et al., 2015; Fyhn et al., 2004). This apparently simple experiment has yet to be replicated in monkeys. Work looking at freely moving monkeys has instead tended to have the monkeys explore relatively high-reward target locations (Mao et al., 2020; Rolls et al., 1997), and these studies have tended to find correlates of the monkey’s visual search in looking for these targets rather than his absolute physical location during this search.

We sought to create a more direct comparison between monkey and rodent work by directly training a monkey on an analog of this free foraging task. We trained one monkey (GS) to use a joystick to navigate around a virtual field, bounded by walls on all sides. Floating bananas were randomly scattered around this field, and the monkey was rewarded for “collecting” each banana by running into it.

In this section, I will describe our results from a small set of recordings in this virtual foraging task. I will show that our results differ strikingly from what is expected in rodent data. Our results are, however, consistent with a task-focused view of hippocampal function, and are reflective of the monkey’s behavior in this task.

3.1.1 Methods

Whereas all other sections of this thesis use our home-grown VR platform built in the Unity Game engine, the foraging task used in this experiment was created in collaboration with the Kahana and Jacobs labs using their PandaEPL game engine, an extension on the open source Panda3D engine that allowed for continuous data logging (Jacobs et al., 2013; Solway et al., 2013). The logging feature in this engine is not asynchronous, meaning data is written to disk every frame. When used on the high framerate (120Hz) monitors typically used in monkey electrophysiology, this led to PandaEPL games tending to skip frames in a manner upsetting to the monkey, particularly when complex scenes are used. As a result, PandaEPL was stable in this minimal foraging task but unsuitable for the more complex tasks found elsewhere in this thesis.

The behavioral setup was otherwise the same as that used in Chapter 2. The monkey saw the VR task on a 45x30 dva projection screen, and used a customized joystick to navigate his virtual environment while we tracked his eye movements using an ISCAN infrared tracker.

These recordings used a “flex-shaft” probe, developed by Pascal Fries and Chris Lewis, illustrated in Figure 3-1[8]. Each probe was constructed from layers of polyimide with 16 iridium oxide contacts spaced linearly arranged with 130 micron spacing along the shank. The defining feature of these electrodes is their flexibility: their 10 micron polyimide construction makes them extremely flexible, giving them the possibility of moving with the monkey’s brain and theoretically (a) providing long term recordings and (b) doing minimal damage to the surrounding tissue. We recorded from these probes using a Blackrock Microsystems Cerebus system, sampling at 30kHz.

Monkey GS had two sets of flex-shaft implants. The first implant of 3 probes was apparently successful, yielding some 21 putatively hippocampal single units many of which were held across multiple recording days. Figure 3-1 (a) shows the target locations for these recordings, although we do not yet have histology to corroborate this targeting. The second implant was less successful, yielding no single unit data. This low success rate was a major contributing factor to our ultimate switch to the Grey Matter Hyperdrives for further recordings.

[8] These probes are now available under the name “MicroFlex Array” from Blackrock Microsystems

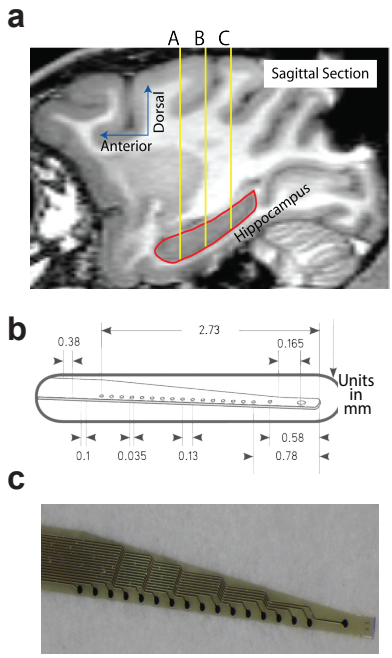

Figure 3-1: Flexshaft Recording

(a) Presurgical targeted locations for hippocampal recordings, show plotted over the MRI from monkey GS.

(b) CAD drawing for “flexshaft” polyamide probe, taken from marketing material.

(c) Image of “flexshaft” probe used in recordings.

3.1.2 Results

We did not see distinct spatial responses during virtual foraging. Figure 3-2 shows activity of four putative hippocampal neurons (b-e), plotted as a function of the location of the monkey's avatar in our virtual environment. Despite having good coverage of our virtual environment (right), we found no cells that had spatial responses that resembled the punctate field seen in rodents or could not be better explained by some other behavioral feature.

Given the research into the influence of gaze location on hippocampal firing in the monkey, it might also be expected that we would see spatial firing with respect to where on the virtual screen the monkey was looking, rather than his avatar's location in the virtual world. Our virtual reality task, however, did not drive the monkey to fully explore the screen with his eyes. Instead, Figure 3-2 (f) shows that the monkey's gaze was primarily located along the horizontal plane of the screen where the majority of task-relevant objects (bananas, walls) appeared during the VR task. Even with this caveat, we did not see cells with reliably specific view-space responses to any of the screen areas well explored by the monkey. Examples of these non-specific responses are shown in Figure 3-2 (g) and (h).

This is not to say that we do not see reliable, task-related responses in these data. To determine whether other navigational features might drive responses, we looked for correlations between the monkey's forward and turning velocity as he performed this task. Of 21 putatively hippocampal neurons included in this analysis, 9 (40%) showed tuning to whether the animal was turning (Figure 3-3 (a-e)) and/or his forward moving speed (Figure 3-3 (f-h)). Although the controls in this task gave the monkey the ability to modulate his speed by changing the pressure he applied on the joystick, for the most part this monkey (and all other monkeys trained in this task) drove with an "all in" strategy, maximizing speed and joystick pressure. With this in mind, turning responses appear to be binary, with cells selective for one turning direction or another. Although speed responses appear more graded, there appears to be a similar jump between high and low velocities for these cells as well.

To make sense of these responses, we looked further at the monkey's behavior. Although this foraging task takes place on a virtual-spatial stage, it is not in and of itself a spatial task. Bananas originate randomly in groups of 10, with 10 more appearing after the monkey has collected the existing set. There is, however, no benefit to collecting bananas in a particular order nor is there a necessary

memory component. Thus, there is no need to attend to the explicit spatial location of any one banana. In this sense, the task we have asked the monkey to perform is “align the avatar with a banana, then drive into it to receive reward.” This latter description better suits observed data; neural activity correlated with the maneuvers the monkey is making in retrieving the banana, rather than the spatial location of the avatar.

To ask if this intuition matches the monkey’s behavior, we looked at the maneuvers made by the monkey in collecting bananas. Specifically, we asked what his strategy was in collecting the first banana in each “trial” or set of 10. The advantage of using this first banana is that banana positions are random, and therefore for the first banana the monkey has no prior knowledge about the distribution of target locations. Figure 3-4 shows the probability distribution for monkey GS when different numbers of bananas appear in his 45-degree (pi/4 radian) field of view. This monkey had a right turn bias; when no banana was visible, he almost always turned to the right and collected the first banana he saw. If only one banana was visible, he collected this banana before continuing on to those outside his field of view. Interestingly, if more than one banana was visible, he would drive to the closest and leftmost of his options. Though this is apparently in contrast with this monkey’s right-side bias, this initial left turn actually sets the stage to optimize the number of right turns in collecting subsequent bananas. These data show that the monkey’s strategy revolves around the maneuver he is making, rather than the explicit spatial location of the monkey. The maneuver-specific responses seen in this task are therefore consistent with the idea that the monkey hippocampus is tracking the animal’s task progress.

In addition to looking at the specific responses of individual neurons, one of the promised advantages of the flexshaft recordings is the ability to achieve stable recordings from an individual neuron across long periods of time. In a few cases, we were able to identify that a particular cell held its navigational response across (1) different virtual environments and (2) across recording days. In rodents, global remapping (Leutgeb and Leutgeb, 2014) would lead to different spatial responses in different environments. The stability of maneuver responses across environments further suggests that it is the maneuver itself, rather than something about the environment, that is driving hippocampal responses. The stability across days further demonstrates the stability of these task responses.

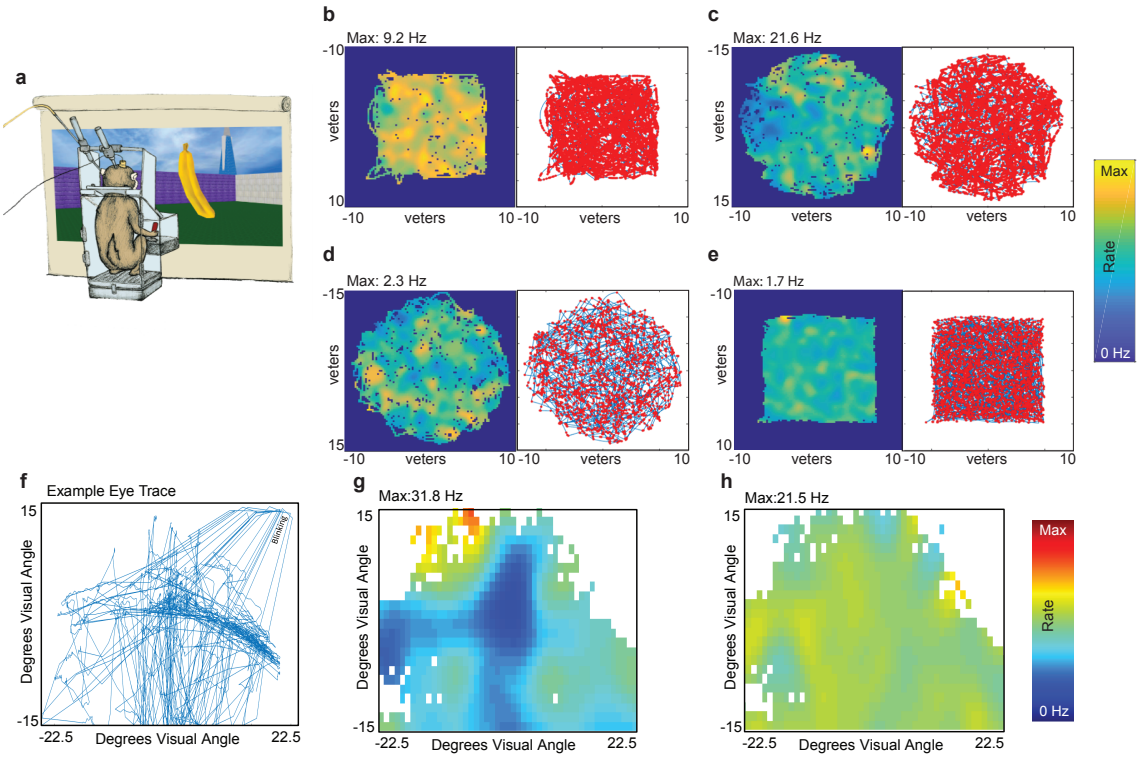

Figure 3-2: No Responses to Spatial Position or View

(a) Virtual Foraging environment. Monkeys navigate a visual sparse environment (distal landmarks and wall colors are only cues). Distances are in “virtual meters,” which I contract to “veters.”

(b-e) Rate maps (left) and spike locations plotted over trajectories (right) from 4 putative hippocampal neurons recorded during foraging.

(f) Example eye trajectories. The monkey explores primarily the horizontal plane of the screen during foraging. Blinks are excluded from further analysis.

(g-h) Rate maps for firing of two example neurons based on eye position on the screen

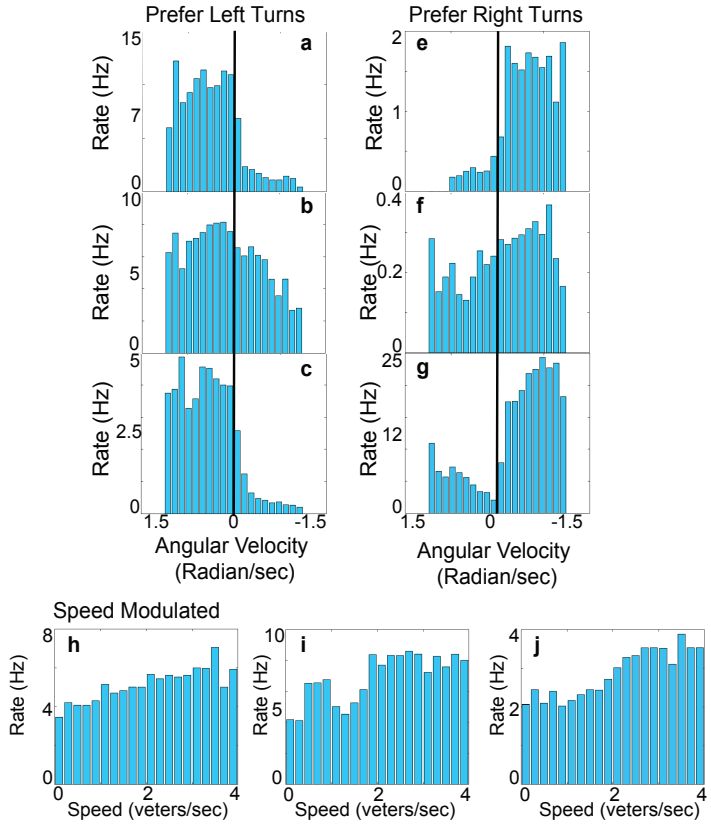

Figure 3-3: Responses to Navigation

(a-g) 6 distinct neurons with turning responses, with spike rate shown as a function of angular velocity.
(h-j) 3 distinct neurons with speed responses, with spike rate shown as a function of velocity.

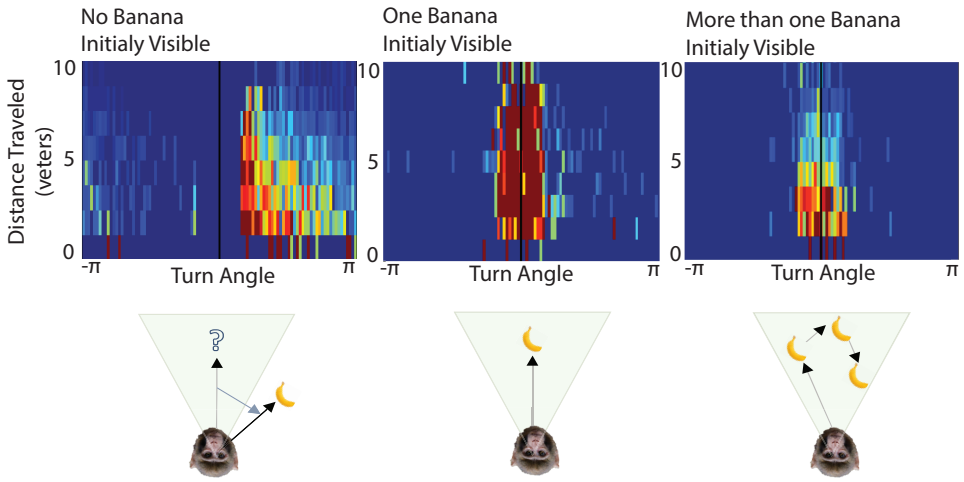

Figure *3-4*: Behavior in Retrieving First Banana in Trial

(Top) Probability distribution for the location of the first banana collected by the monkey GS within a trial. Locations are expressed in monkey-centric coordinates, with the X axis denoting angle turned and the Y axis denoting distance to the banana. Trials are sorted by the number of bananas visible.
(Bottom) Cartoon summarizing the monkey's strategy in each situation.

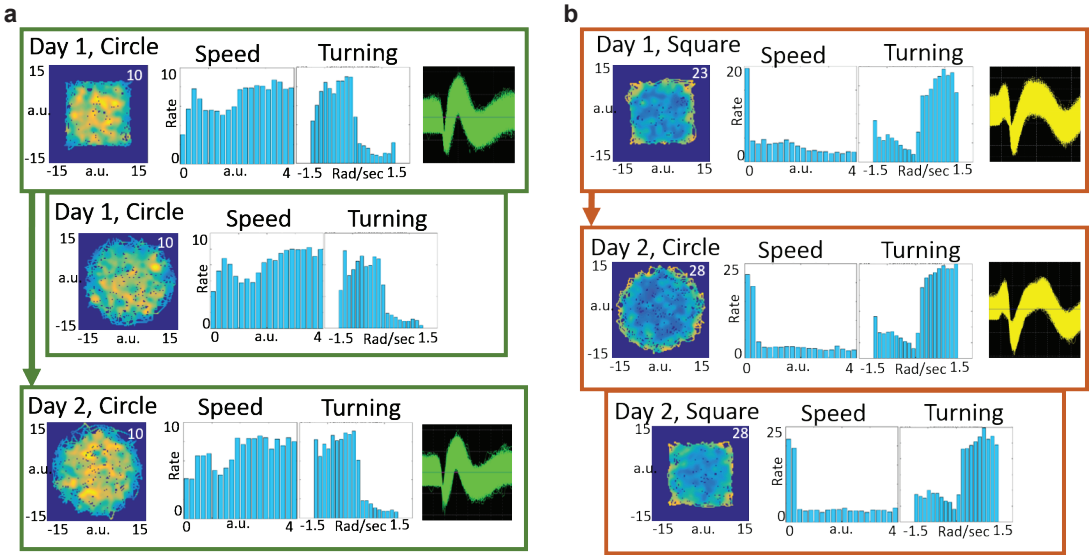

Figure *3-5*: Reponses can be Stable Across Sessions, Environments, and Days.

Responses of two neurons held across days. Plots show the rate map, speed and tuning curve for each neuron within a day in different environments and across days with the same environment. Inset shows the waveform shape which, together with having a consistent response on the same channel, was used to estimate if a neuron was the same across days. Note both cells shown here are included in Figure 3-3.

3.1.3 Conclusions

We did not identify allocentric spatial responses during a foraging task in the monkey hippocampus. On one hand, this result is surprising in and of itself; given the well-documented prevalence of place cells in rats, we would have expected to find a few cells with spatial responses if similar processing takes place in the monkey hippocampus. On the other hand, these data fit nicely with the story told by more complete datasets elsewhere in this thesis. Rather than seeing exclusively spatial responses, the responses that we see in these recordings track the events relevant to the monkey as he performs his task. These results are consistent with other recordings in monkeys (Mao et al., 2020; Rolls et al., 1997) that struggled to find punctate place cells in free navigation situations, as well as work in rodents (Aghajan et al., 2015) that failed to find spatial responses in unguided virtual reality tasks. Instead, when monkeys (Courellis et al., 2019) or rodents (Aghajan et al., 2015) are given goal directed tasks, spatial responses emerge that track the steps required to attain that goal. With this in mind, later task iterations were designed with more thought to targets and repeatable trial structures.

Unfortunately, the extremely small and un-replicated nature of the experiment in this section makes it difficult to assess to what extent the data shown here reflect the true underpinnings of the hippocampus, or to what extent they are a symptom of an under-sampled or mistargeted dataset. Others (see (Nishijo et al., 1997b; Rolls et al., 1997; Watanabe and Niki, 1985) et al.) have reported remarkably low numbers of hippocampal neurons engaged in task related activity, with often less than 5 or 10 percent of cells having identifiable responses in a given task. If this is the case, 21 neurons are hardly enough to reliably identify any one response phenotype. What is more, the relatively novel nature of the flexshaft recording technology means that we do not know (1) what was the likelihood that we were recording multi-units and thus losing whatever spatial specificity might have been present and (2) what the risk might be that our electrodes partly or entirely missed the hippocampus.

A final concern is that the virtual environment used in this study was, quite frankly, boring. The environment is cue-poor to the extreme that the monkey may have ignored the few spatial cues (all distal landmarks or wall colors) available to him. When, for example, we increased the probability of reward in a particular part of the virtual maze, we were unable to find behavioral evidence that this affected his strategy in any way. While this may just have been an idiosyncrasy of our task design (he still had to

collect all 10 bananas eventually, so there was no incentive for “greedy” searching), it could also reflect an under-developed visual landmark that led the monkey to simply believe that we had randomized reward. This concern was addressed in later experiments through the inclusion of much more obvious landmarks and cues in the virtual environment.

3.2 Navigation in a color world

Despite the decades of work characterizing spatial responses in rodents, it remains unclear whether or how these spatial responses might contribute to the structure’s role in episodic memory. One hypothesis proposed to bridge this disconnect is that physical space is not the only space represented by the rodent hippocampus. An alternative model posits that it is task structure that drives hippocampal activity, and that this tracking of ongoing experience allows for memories of these experiences to be formed. In this view, spatial correlates are then a natural result of any spatial correlates that are inherent to the task. Confirming or differentiating either of these hypotheses requires the development of tasks that circumvent the inherent physical features of a task involving freely movement. In this section, I will discuss our work developing tasks that require navigation through an abstract space, in our case a sequence of arbitrarily connected colors, and show some preliminary data from recordings in these tasks.

Just as place cells might form a “cognitive map” (Tolman, 1948) of an animal’s physical environment that animal can use to navigate (O’Keefe and Nadel, 1978), hippocampal neurons might provide similar coverage for more abstract or conceptual variables (Bellmund et al., 2018; Constantinescu et al., 2016; Knudsen and Wallis, 2021; Schafer and Schiller, 2018; Schiller et al., 2015). Under this model, such conceptual maps can be used by the animal to explore learned relationships in abstract variables, in turn facilitating memory. In an alternative formulation of this idea, others (Whittington et al., 2020) have suggested that the hippocampus might learn relationships between objects and other features of both abstract and physical space, allowing for a similar mnemonic mechanism without explicit construction of a spatial coordinate system.

A particular challenge with this abstract “map” hypothesis is that it has been difficult to test empirically. There are, however, a few notable examples of rodent hippocampal activity spanning a non-spatial variable. When a task has a built in time delay (often in the form of a rat running on a treadmill), neural activity in the hippocampus will tile that delay (Pastalkova et al., 2008) independently of the

animal's spatial location (Kraus et al., 2013), leading to the suggestion that time is one such abstract variable that can be tracked by the hippocampus. Additionally, when a rat is asked to respond at a particular moment in a sequence of tones, neurons in the hippocampus will respond to particular tones in the sequence (Aronov and Tank, 2014).

The caveat with these results is that they are wholly task dependent. Neurons in the hippocampus do not, for example, respond whenever a rat runs on a treadmill, but rather only when the treadmill run is associated with a specific task (Pastalkova et al., 2008). Responses to tones similarly require the animal to be engaged in the tone response task (Aronov and Tank, 2014). These subtleties promote the hypothesis that it is not necessarily some internal map of these task variables that drives neural activity, but the specific structure of the task itself. In many tasks, neural firing does not strictly tile the task space but rather clusters around key moments in the task like trial start or reward (Aronov and Tank, 2014; Butler et al., 2019; Gauthier and Tank, 2018). An example of this uneven tiling is also seen in my own work (see Chapter 2), in our case exemplified by inconsistent coverage of our virtual maze. Further, tasks with discrete rather than continuous events show analogous firing to task events (Sun et al., 2020; Terada et al., 2017; Watanabe and Niki, 1985), suggesting that the continuous "map" is not needed to drive activity.

To study which features drive hippocampal activity in a non-physical space, we trained monkeys to navigate through a sequence of colors. Here, the monkey drives through a series of colors to arrive at cued "goal" colors. The monkey sees only the color of his current "location," although a set of continuous metric relationships govern both behavior of this color "space" and the monkey's "movement" through it. Although virtual reality, technically, is a non-physical representation of a space, this task allows us to study hippocampal activity in a situation where not even the illusion of physical space exists, even if the metric relationships inherent in physical motion still do.

The map and task models for what might drive activity in the hippocampus each make specific predictions of what these data will ultimately show. If monkeys construct a "cognitive map" of our color space, we would expect to see even tiling of this space by a sizable constituency of hippocampal neurons. This titling would likely change, or "remap", if the metric relationships underpinning these colors were to change. However, if task is the primary driver, then we would expect to see neural activity cluster

around key task events irrespective of the specific color sequence used. At the time of my writing this thesis, recordings in this task are ongoing and data analysis is still underway. Present preliminary data, as with our results in Y-maze environments approximating physical space, point to the idea that the task the monkey is performing is the better model for describing hippocampal firing.

3.2.1 Behavior

We trained two versions of this color navigation task, both of which are highly similar in structure. In both cases, the monkey learns the relationships of a sequence of colors that are mapped onto an underlying metric space. On a given trial, the monkey is cued with a color, then must use a joystick to "navigate" through the color "map" until he thinks he has found a match. The monkey is trained to acknowledge this match by releasing the joystick, at which point he is either rewarded for a correct response or reset to the start of the task. While he is navigating, the monkey sees only the color representing his current location on the map. Because colors typically fade continuously together, a range or "bubble" of space around the target color is accepted as a correct response. Figure 3-6 gives an outline of this task. In both cases, the color task was programmed and run on the same Unity VR system used elsewhere in this thesis.

3.2.1.1 Linear track analog

Two of our monkeys were trained on a version of this color navigation task designed to be an analog of a "linear track" in a physical space. Here, there is a colormap consisting of 6-8 color segments with black fades in between, and the monkey moves through this map in a fixed direction. At the beginning of each session, the monkey is shown a cue indicating the target color on a fixed set of "training trials," after guided fixation on a target. He is trained to release when he thinks he has arrived at the target. During the initial training trials an additional border indicating the target color remains visible to the monkey during navigation, but this aid is removed once the training trials are complete. Speed is varied within trials so that the match cannot be made solely on drive time, effectively requiring the monkey to attend to the colormap. A schematic of this task layout is shown in Figure 3-6 (c).

One advantage of this framework is that monkeys can rapidly learn new color maps. Within a 20-30 min session, a well-trained monkey can be shown ~15 training trials, learn the stop point on the map, and successfully perform task for the remainder of the session without being re-cued as to the color.

This task is extremely similar to the one used by (Aronov and Tank, 2014), with the exceptions that: (1) we use colors to define our metric space, rather than sound (2) we routinely change the color maps, allowing us to ask questions about learning, context dependence, and remapping and (3) in many of our maps we repeat non-target colors, allowing us (at least in principle) to differentiate color-specific responses from task-location responses.

3.2.1.2 Two-directional Task

One monkey (SP) was initially trained on a more complex color navigation task. Here, there were four available target colors evenly spaced throughout the map. The monkey starts in the center of the map, and is cued towards one of these targets. He must then drive to whichever target is requested on a given trial. Importantly, these different target colors can share overlapping elements of their trajectories, allowing us to compare location and target/trajectory specific responses. Further, once a direction is chosen the monkey cannot move “backwards,” so that the initial movement on the trial provides a “two-alternative, forced choice” metric for the monkey’s understanding of that map. A schematic of this task design is shown in Figure 3-6 (d).

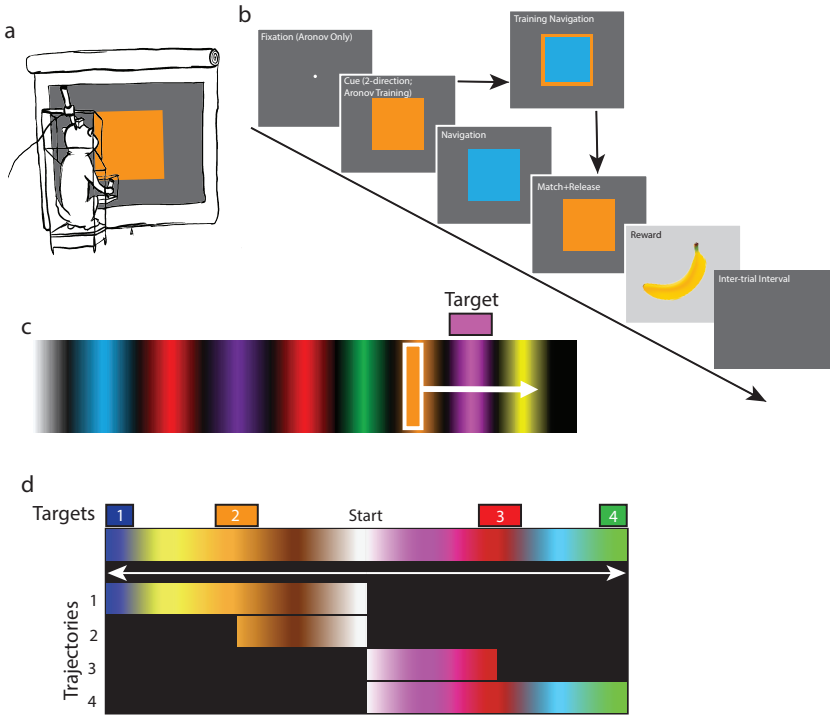

Figure 3-6 One and Two Direction Versions of the Color Navigation Task

(a) The monkey plays the color game in the same VR rig used for other tasks

(b) Schematic of task flow.

(c) The “linear track” or “Aronov” task. Colors are separated by transitions to black. The monkey only sees the color he is currently on.

(d) The “2-directional” task. Here, the map has 4 potential targets and the monkey must navigate to the correct target based on his initial cue.

3.2.2 Results

3.2.2.1 Task correlates in the color game.

We see putatively hippocampal neurons with responses to the monkey's progress through the linear color map. Figure 3-7 shows four examples of neurons with map-correlated responses. These neurons tend to be tuned to either the beginning or end of the monkey's trajectories (Figure 3-7 (b) and (c), also see Figure 3-9). When neurons do show mid-maze tuning, it is sometimes better explained in terms of transitions between the semi-discrete colors (as in Figure 3-7 (a)) than in terms of specific tuning to particular points along the maze.

In rare cases neurons show tuning for a location or color in the map. One example of this is in Figure 3-7 (d), with another example in (c). Although these examples are too few and far between to draw explicit conclusions, it is worth noting that in both cases the response overlaps with color that is repeated in the map topology. In both cases the neurons respond differently at each iteration of the color. These data are a preliminary indication that it is the map organization, rather than particular colors, that underlies this specific activity.

The relative overrepresentation of response at the starts and ends of the monkey's trajectories led us to look more carefully at these periods. Figure 3-8 shows examples of responses to these salient task events in both the linear track (black raster) and 2-direction (colored raster) versions of the task. The cells in Figure 3-8 (a-c) are correlated when the monkey releases his joystick and receives reward, while those in Figure 3-8 (d,e) respond just as the monkey is about to initialize movement. Although the linear track version of the task does not have a cue-period on every trial, the 2-direction has a 3-second cue period at the start of each trial. A further subset of cells seems to be selectively active during particular parts of this cue period, as seen in Figure 3-8 (f-h).

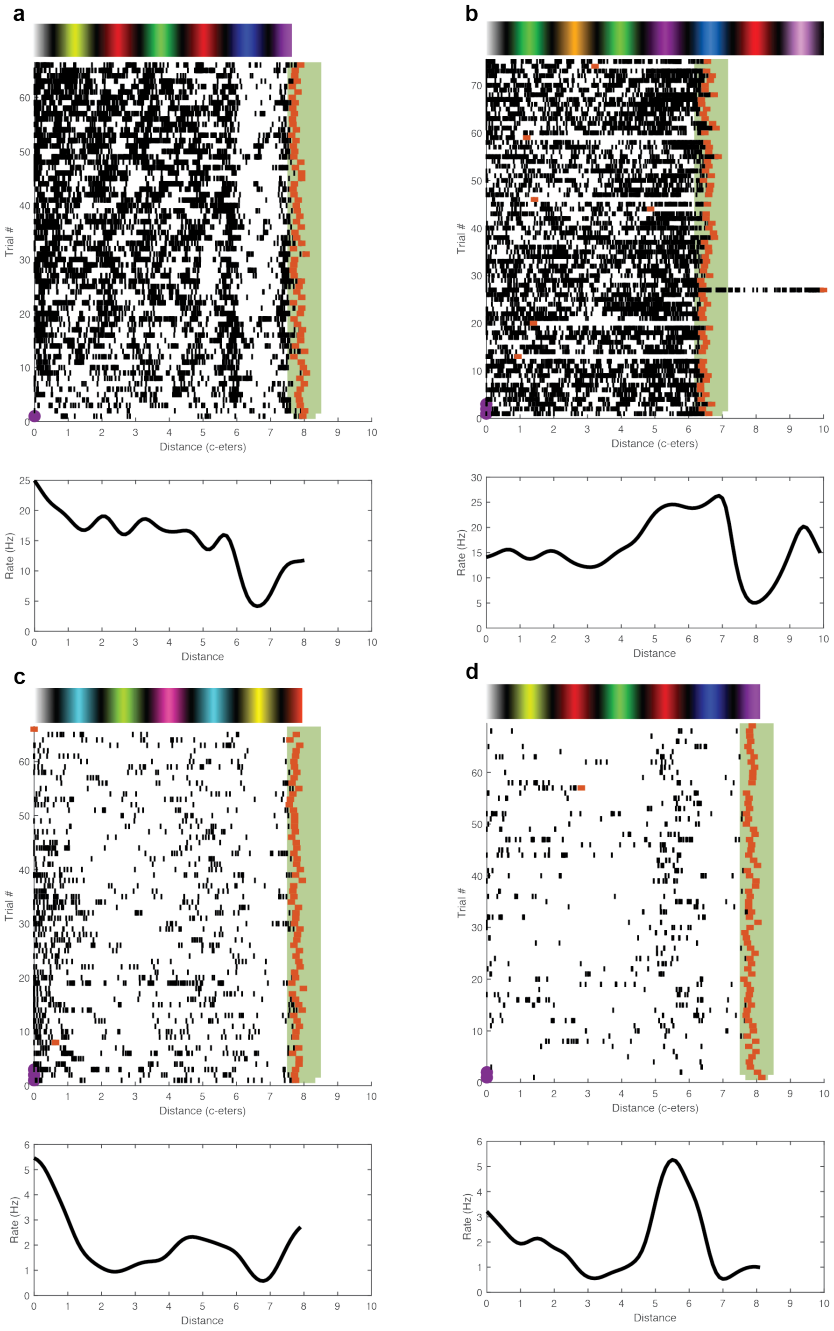

Figure 3-7 Responses in the Linear Color Game

(a-d) Examples of single neuron responses in the linear color game. Four cells are shown for each:
(Top) The color sequence the monkey saw while driving. In cases where the sequence does not span the full axis, this is because the monkey did not traverse the full map
(Middle) Per-trial raster of each neuron, with each spike shown as a black tick. The green field denotes each trial's target, orange ticks represent the monkey's stopping point on each trial. Purple dots indicate trials with a cue.
(Bottom) The average rate for the neuron in question.

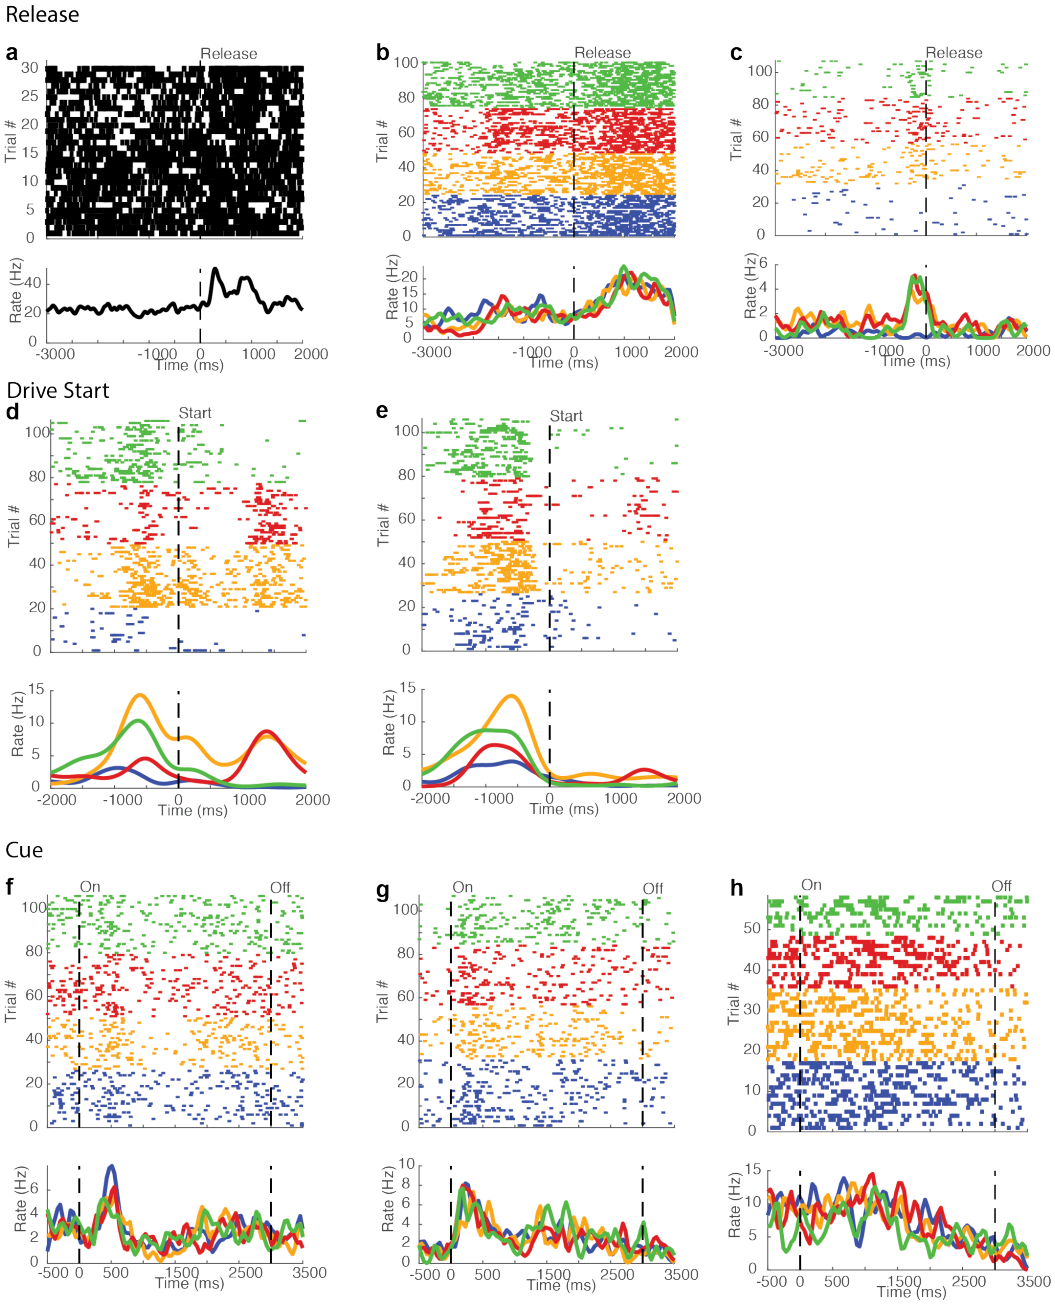

Figure 3-8 Event responses in the color game

(a-c) Activity of three neurons, triggered on the monkey's joystick release at the end of the trial. (a) is from the linear maze, (b-c) are from the 2-direction maze and sorted/colored by target color.

(d-e) Activity of two neurons, triggered on the monkey's initial joystick press. Both are from the 2-direction maze and sorted/colored by target color.

(f-h) Activity of 3 neurons during the cue period in the 2-direction maze. "On" at time zero indicates the start of the cue, and "off" at 3 seconds indicates the shortest possible cue period (cue was extended if monkey pushed the joystick too soon).

3.2.2.2 Responses can be trajectory specific.

In certain cases, the responses in the color game can be specific to the particular trajectory the monkey is taking through the color space. Our two different versions of the color task offer two opportunities to observe this phenotype of response. In the linear task, the monkey drives a block of trials using a single set of colors to reach a single target. In much the same way as in Section 2.2.4, we can have multiple blocks of trial using different maps. Figure 3-9 (a) shows the response of a single neuron with maps presented in A-B-A pattern. In the first color map (A), this cell fires primarily as the monkey starts driving, and this pattern is qualitatively preserved when the monkey is returned to this map later in the day. In the second map (B), however, the cell fires very differently, effectively ramping toward the end of the drive.

In the two directional color map trials to different targets do not occur in blocks but are instead interspersed with each other. Despite this difference, location specific coding still depends on the monkey's target. Target-specific firing of three cells is shown in Figure 3-9 (b-d). Importantly, this 2-direction task requires traversal of the same color segments for multiple targets, e.g., the monkey must pass through brown on its way to both the orange and the blue targets. Despite the overlap in the color sequence, the firing of our putatively hippocampal neurons appears to depend on the monkey's target. In Figure 3-9 (b,d), for example, the cell fires selectively on orange-going trails but not blue-going trials that cover the same locations in the map topology. Figure 3-9 (c) shows similar behavior with a preference for blue-going trials. At this point in our analysis, the vast majority of cells with color-specific firing do so in a trajectory specific manner.

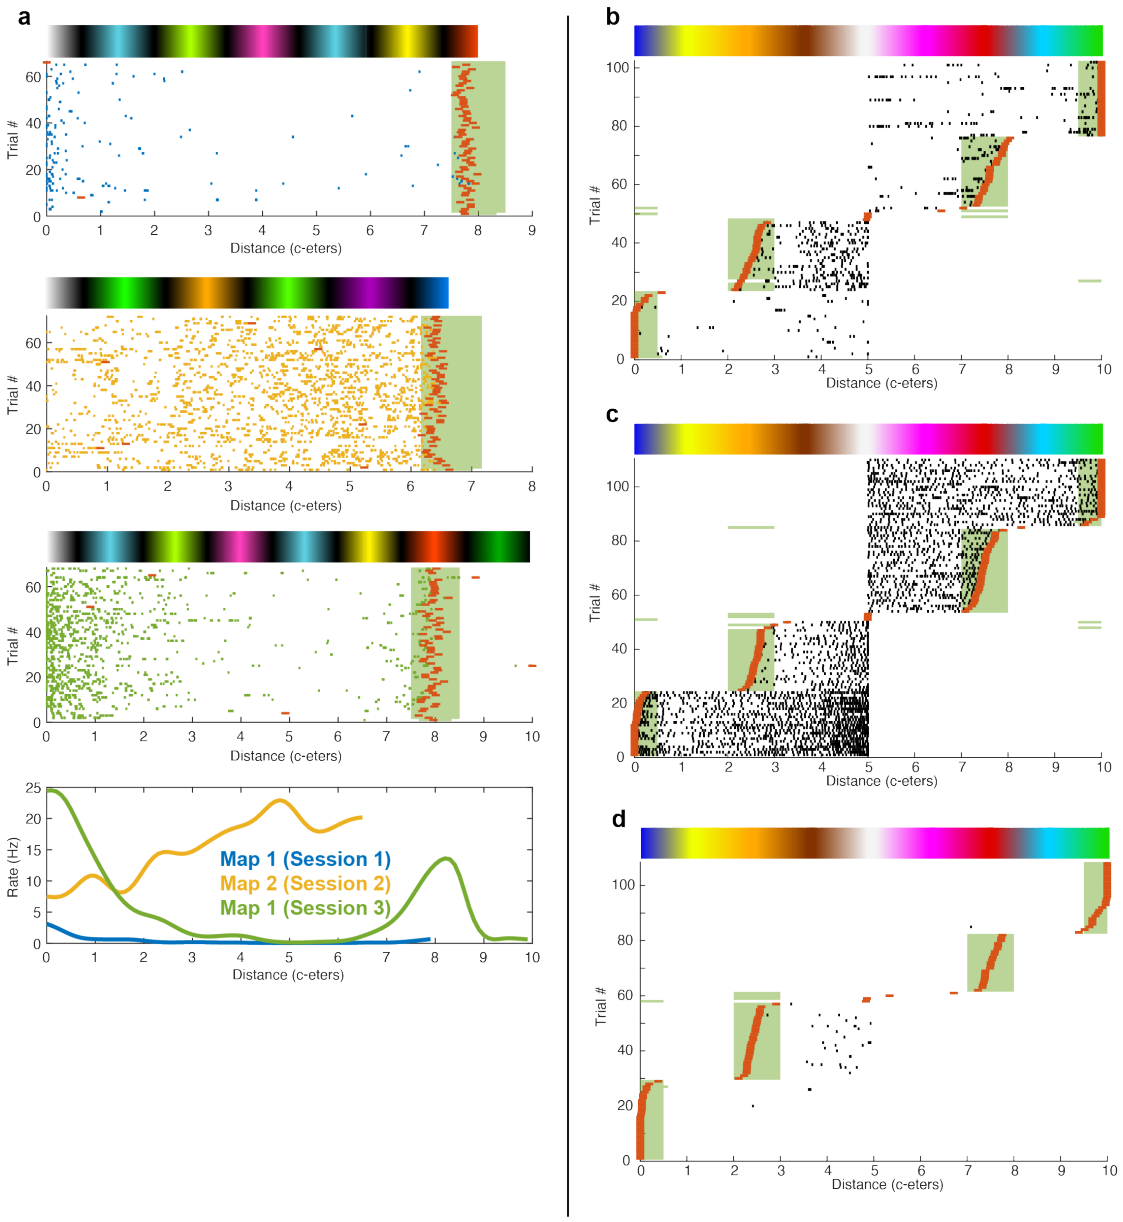

Figure 3-9 Color Task Responses can be Trajectory Dependent

(a) Activity of a single neuron over sessions of the linear color task, including two color maps presented in an A-B-A format. Rasters are as in Figure 3-7, with aligned rates for all three maps shown in the bottom panel. The response of the neuron is broadly map-dependent, albeit with a gain change between sessions 1 and 3.

(b-d) Activity from three neurons in the 2-direction color task. Notation is similar to Figure 3-7, except the monkey starts at the center of the maze (5) and drives outward towards the indicated target. Here, neurons show trajectory-dependent firing even while traversing the same set of colors between targets.

3.2.3 Conclusions

First, it must again be stressed that the data presented in this chapter are preliminary; as with the recordings discussed in Section **Error! Reference source not found.**, our color game data come from one monkey with a small population of sorted cells. Until more spikes are sorted and data are collected from a second monkey (both efforts that are underway at the time of writing), it is impossible to estimate if the responses shown here are ubiquitous in the hippocampal population or if they represent a relatively uncommon phenomenon.

Even with these caveats, it is worth noting the striking similarity between the task-selective responses we see in the color game and those we see in other VR experiments, such as those seen in the Y-maze in Chapter 2 or in Foraging in Section 3.1. Our color task, we rarely see the punctate, "location"-specific responses that one might expect from rodent studies in similarly abstract tasks (Aronov et al., 2017; Pastalkova et al., 2008). Instead, while we do see task-specific responses, they may be better described in terms of task events -the start of the trial, the end of the trial, or the presentation of a cue – than as a function of the color-topology learned by the monkey.

In our 2-directional task, location-specific firing is further convolved with the trajectory the monkey is navigating along. Were cells to represent particular locations along a maze, then we would expect common representations regardless of the monkey's target. Instead, this trajectory-specific firing mimics the phenomenon known as "splitting" in navigating rodents, where neurons will fire differently down a particular path depending on their upcoming trajectory (Kinsky et al., 2020). Having said this, it is notable that we rarely see splitting in the Y-maze discussed in Chapter 2 (see specifically Figure 2-7, which shows common activity during the center stem on left-ward and right-ward going trials).

Although our color task has obvious similarities to the navigation literature because of the requirement that the monkey traverses a metric map of some abstract space, I would be remiss if I did not also mention the similarly of this task to many of the delayed-match-to-sample (DMS) tasks discussed in Chapter 1. Like in those studies, we ask the monkey to match his position to a particular target. Interestingly, a number of these DMS studies showed activity that was both specific to task events and that could be selective to trial type or category (Hampson et al., 2004; Watanabe and Niki, 1985).

Although our color task builds on these findings by incorporating a traversal of a continuous task dimension, our findings of task-phase and trajectory-specific firing are reminiscent of these DMS results.

4 Eye movement modulation of single-cell hippocampal activity

4.1 Introduction

Eye movements modulate neural activity in the hippocampus. After both guided saccades (Nowicka and Ringo, 2000b; Ringo et al., 1994; Sobotka et al., 1997) and natural saccades during free image viewing (Konig and Buffalo, 2021), neurons in the hippocampus show reliable changes in firing rate on the timescale (<250ms) of natural eye movement behaviors. Although an explicit connection to the activity of single neurons remains unclear, similar timescales of eye movement correlates are also visible in hippocampal local field potentials, where eye movements appear to reset underlying theta-band oscillations (Hoffman et al., 2013; Jutras et al., 2013).

The precise function of this eye-movement modulation remains unclear. One intriguing possibility is the suggestion that eye movements “gate” information flow to the hippocampus and other medial-temporal-lobe structures (Minxha et al., 2017). In this model, neurons may have specific sensory and behavioral responses (e.g. to faces in (Minxha et al., 2017) or to spatial location in (Konig and Buffalo, 2021)), but selective firing of these responses is timed relative to the eye movement and subsequent arrival of sensory information. A second (and certainly not mutually exclusive) model suggests that eye movements provide a similar organizational mechanism to that of rodent theta (see Chapter 1, Section 1.3.3).

To date, studies of eye-movement triggered activity have been limited to simple tasks in dark rooms (Ringo et al., 1994) or, at their most ethological, to free image viewing (Konig and Buffalo, 2021). While informative, the constrained nature of these studies limits understanding of the extent to which hippocampal activity might show eye movement relationships in more naturalistic settings. Virtual reality presents itself as a convenient tool for studying this. In VR, for example, monkeys make a wider array of foveations beyond a simple fixation, i.e., smooth pursuit during motion and optokinetic reflex during turning. Because these different foveation types take place at different locations in the virtual maze, I will not attempt to discriminate between them. However, their presence indicates the rich gaze behavior we are able to evoke within the controlled setting of virtual reality. Further, although we will not claim to describe the internal strategy monkeys use to solve our VR task, VR also allows us to distinguish between onscreen, task related foveations and foveations that are directed offscreen or take place during

the intertrial interval. Finally, as with spatial responses, we can compare task foveations in different visual environments.

Here, I will describe eye movement modulation in our virtual reality task. I will show that we see eye movement modulation reminiscent of what is observed in more constrained paradigms. I will show that responses can change or “remap” (1) between tasks, (2) between different conditions within a task, and (3) between different contexts within a task framework. Finally, I will show that in a task where eye movements cannot provide the monkey with task-related information, we have yet to see eye-movement modulation in the hippocampus.

4.2 Methods

First, I want to provide a few notes about the neural data shown in this chapter. While automated spike sorting has provided (elsewhere in this thesis) access to larger longitudinal studies of neural recordings, for the most part the sample sizes in this section will be smaller and involve only the neural subpopulation that I have isolated by hand. Further, I will constrain my analysis of the hippocampus to those cells that are identified via MRI registration as being within 1-2 mm of the outlined border of the hippocampus. Given the inherent uncertainty in MR guided position identification (voxel size is >=.35 mm^3, titanium implants distort images, monkeys grow, and the assumption of perfectly straight electrodes or perfect depth estimates would likely reflect an overly optimistic picture of mechanical and human error), these data should be treated as preliminary pending histology to better identify recording sites. Finally, I have made little effort to determine whether reported cells are the same cell held across days; while there is the potential for interesting analysis on the stability of eye movement responses, my goal here is to provide a survey of the nature of this activity in the hippocampal population. Such further analysis remains a clear starting place for future work.

4.2.1 Behavioral Tasks

This chapter contains secondary analysis of data from 3 tasks.

(1) A calibration and image viewing task, in which the monkeys freely explore natural scenes. Unrewarded image trials are preceded by three calibration dots, and the monkey is rewarded for fixating each dot. Ultimately, this guided fixation is used to establish eye tracking for other tasks.

(2) The virtual Y-maze, a delayed spatial alternation task (see Chapter 2). This task requires navigation through a complex visual scene, and we track eye movements during this exploration. Each trial is separated by a 4 second ITI in which monkeys continue to make eye movements, but no stimulus is presented on the screen.

(3) The color-game task, designed to require monkeys navigate through a non-physical space (In our case a sequence of colors, see Chapter 3, Section 3.2).

4.2.2 Eye Tracking

In all three tasks, we used the same eye tracking algorithm. At its core, this is a simple velocity thresholding algorithm in which saccades are considered moments of high eye velocity and foveations are epochs of low eye velocity. There is, however, noise in our velocity estimate. To account for this, we use the following algorithms. This is illustrated with a Y-maze eye trace in Figure 4-1:

(1) All data are up-sampled to 1000Hz (using hold over interpolation). Blinks and other outside-of eye-tracker timestamps are marked as missing data. Although the algorithm is run by interpolating data across these epochs, any classification made will be discarded as a final step. The unsampled data are then low-pass filtered at 60Hz to limit noise.

(2) A velocity threshold is taken as 100 degrees of visual angle/second (dva/s). Points with a velocity above this threshold are tagged as either “saccades,” with those below denoted as “foveations.”

(3) For each block of saccade points, the peak is identified and gradient descent is used to identify all points contributing to that peak, even if they are below the initial velocity threshold. A minimum velocity of 17 dva/s is still required to be considered part of a saccade.

(4) For each remaining foveation, the standard deviation of velocities within that foveation is found. Any points with a velocity >3 standard deviations above 0 are flagged as saccades, and the peak detection is repeated.

(5) For VR data, an additional step is needed. Whereas in, for example, image viewing, the terms foveation and fixation can reasonably be used interchangeably, this is not the case in VR where we classify foveations as smooth pursuit, optokinetic reflex, or fixation. Here, anything flagged as a foveation undergoes an additional processing step. If the monkey is moving forward and the eye velocity exceeds 17 dva/s, that point is flagged as a smooth pursuit. If the monkey is turning

at >1 degree/second and eye velocity exceeds 17 dva/s, that point is flagged as optokinetic reflex. All other points are classed as fixations. Finally, all points within a foveation are given a “vote” on the classification of that foveation in a winner take all decision.

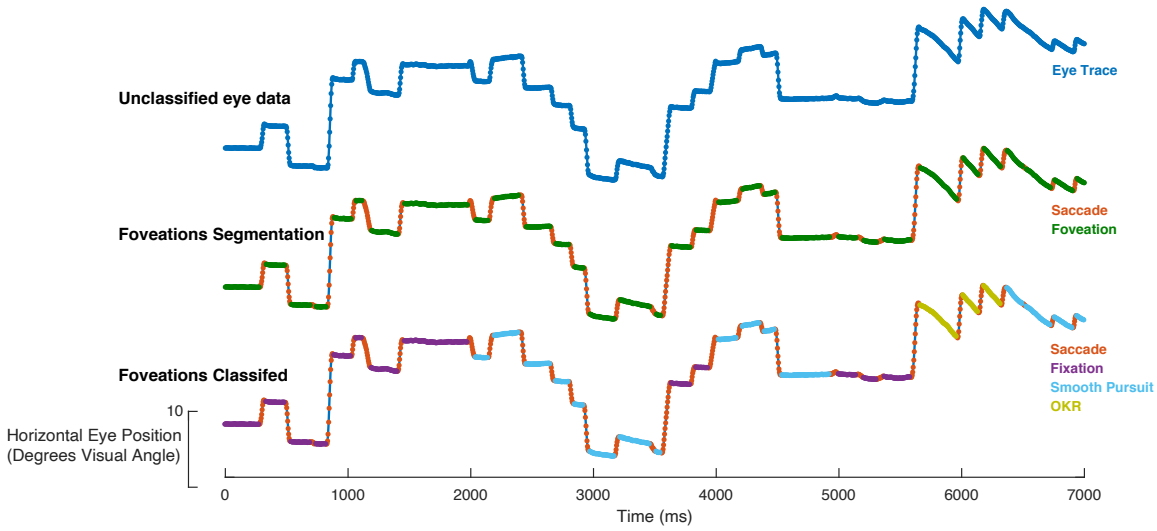

Figure 4-1: Classification of eye movement traces.

Each row shows the same eye movement trace from monkey SP as he navigates the center stem of our virtual Y maze. Dots indicated samples recorded by the eye tracker, and the line connecting them is a linear interpolation to guide the reader.

(Top) Unclassified eye data

(Middle) Data divided into saccades and fixations (up to step 4 in the text)

(Bottom) Data further divided by eye movement class (Step 5 in the text). OKR: optokinetic reflex

4.3 Task specific eye movement modulation in virtual reality

We see eye-movement modulation during navigation in virtual reality. Figure 4-2 shows example raster plots from foveation triggered spiking (Konig and Buffalo, 2021) for 3 cells, sorted by where (on- or off-screen or inter-trial interval) the monkey was looking when the triggering foveation began. The cells in Figure 4-2 (a) and (c) have notable firing modulation at precise timings after the start of onscreen foveations, with either different or undetectable modulation when the foveations began elsewhere. The cell in (b) has modulation specifically in the inter-trial interval of the task, though this latter mode of response is less common and will not be analyzed further.

Twenty percent of included hippocampal neurons (254 of 1277 neurons meeting inclusion criteria of having a minimum eye triggered rate of 0.5 Hz) showed significant and reliable eye movement modulation in a given environment when the monkey was looking at the screen (Monte Carlo simulations with 10000 shuffles testing for information and stability greater than expected by chance, see Chapter 4, tested on the first virtual environment of the day). Further examples of this sort of activity (all using different cells), focusing specifically on the “on screen’ foveations, can be seen in Figure 4-5 and Figure 4-6.

Even though eye movement modulation is quite prevalent in our hippocampal data, the actual amplitude of the modulation is relatively small. Unlike, say, the punctate response of a rodent place cell, most of the eye movement modulation we see is simply a small but significant shift away from some baseline eye-triggered firing rate. This is demonstrated in Figure 4-3, which shows the normalized firing rates of all our eye movement modulated neurons. The large amount of homogenous yellow suggests that most cells’ average firing rate is near their peak. However, Figure 4-3 (b) shows the population cross correlation for these data. This plot notes how similar population activity is between any two instances in time, effectively giving a metric for how well one might “decode” time after fixation. Here, despite the relatively small changes in rate post eye movement, we see pronounced diagonal stripes suggesting that activity at different offsets from the fixation start is fairly distinguishable. Although the majority of neurons have peak rates in the 50-250ms window, a notable subpopulation with rate troughs immediately after the fixation start make this the most distinguishable epoch in the fixation. Further, we see an off-diagonal component with an offset of about 200ms. This timing corresponds to the approximate duration of a fixation, and suggests that activity is consistent across neighboring fixations. To better illustrate this

phenomenon, Figure 4-3 (c,d) shows the same data with the mean rate subtracted prior to normalization, while Figure 4-3 (e,f) show locations where firing is greater than expected by chance (i.e., greater than the 99th percentile rate from shuffled data).

It is worth noting that the hippocampus is not alone in showing eye movement modulation in its firing. For example, more objectively visual areas, such as LGN and V1, show eye movement modulation in their spiking or visually-evoked LFPs, respectively (Gavornik and Bear, 2014; Reppas et al., 2002). Other temporal lobe areas also show eye movement modulation similar to that seen in the hippocampus (Ringo et al., 1994; Sobotka et al., 1997). Figure 4-4 is a demonstration of a fixation selective cell recorded from an (as yet unidentified) region of frontal cortex during image viewing.

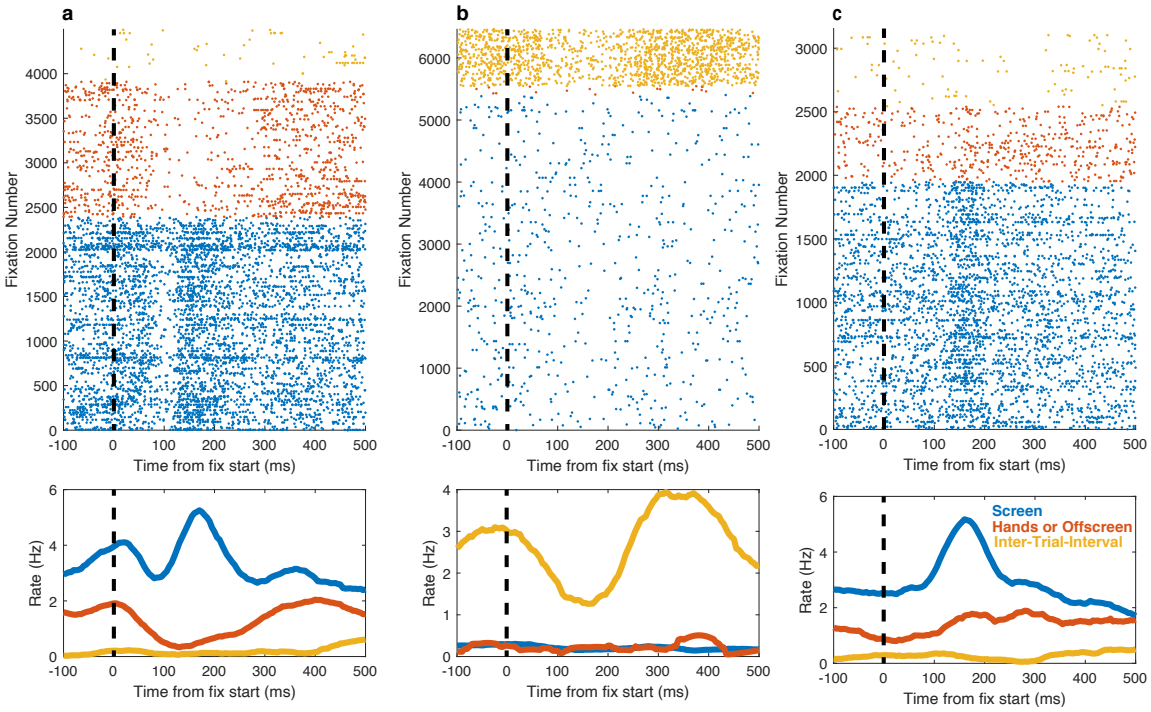

Figure *4-2*: Example foveation triggered raster plots under different task conditions

(a-c) Spikes aligned to the onset of foveations. Spike are sorted by whether the triggering foveation was on to the VR screen during navigation (blue) off-screen and typically onto the monkey's hands (orange) or took place during the intertrial interval (yellow).

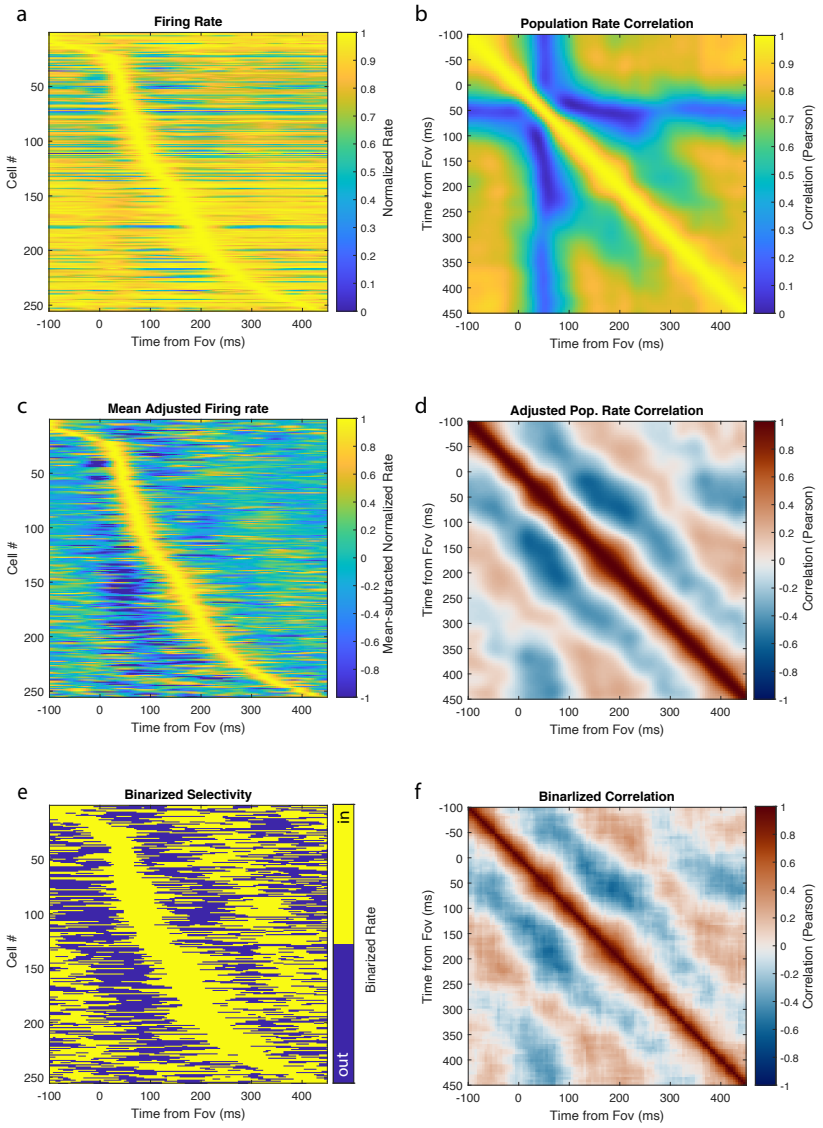

Figure *4-3* Population responses triggered on eye movements.

(a) Normalized rate-maps for 254 neurons with significant eye movement modulation, arranged by peak firing rate.

(b) The population correlation matrix for a.

(c, d) As in (a, b), but with rate-maps normalized after mean subtraction

(e, f) As in (a,b), but with rate-maps binarized, with “in field” regions passing the 99th percentile of firing rates in shuffled data.

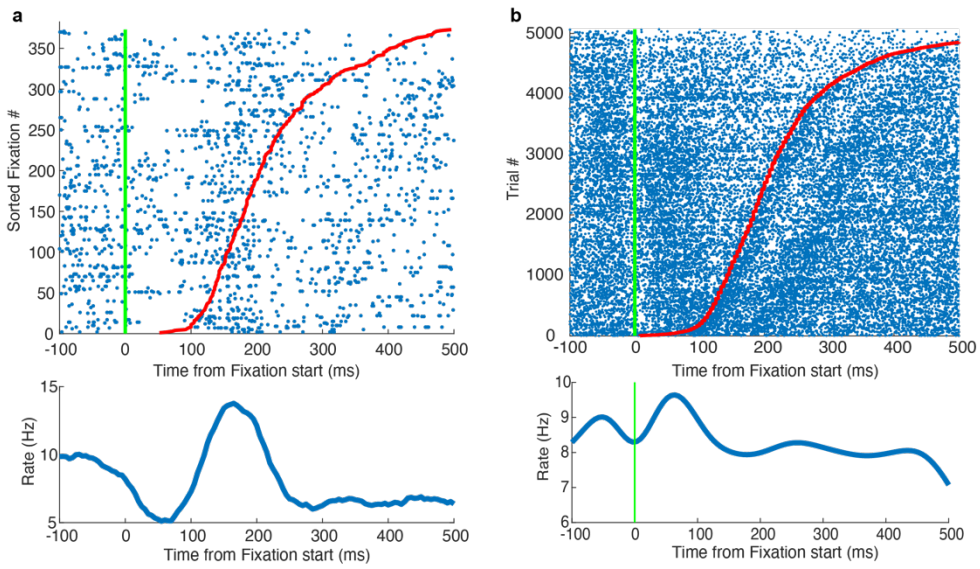

Figure 4-4: Example of cortical cells

Two examples of non-hippocampal, putative cortical cells showing eye movement related activity. Both electrodes are too anterior and too shallow to be in the temporal lobe; the precise location of this cell has not yet been determined.

4.4 Eye movement modulation can be task and context dependent

So far, we have seen that eye movement modulation can be specific to the content of the foveation, even within a given parent task. Others have shown that neurons with both general (Sobotka et al., 1997) and stimulus specific (Konig and Buffalo, 2021) foveation modulation can also be specific to the task a monkey is performing. To determine whether neurons have different responses in different tasks in our data, we compared foveation triggered responses in our Y-maze task to those seen for that same neuron during image viewing. Similar to the previous work, we see neurons that have patterns of eye movement modulation specific to specific tasks. Two example neurons (simultaneously recorded) are shown in Figure 4-5.

We also asked if, in addition to task dependence, hippocampal neurons might show changes in eye movement modulation between visual contexts within a task. To do this, we compared eye movement modulation of individual neurons when the monkey was looking at the screen in visually distinct Y-maze environments, as seen in Chapter 2. A notable population of neurons show context dependent changes in foveation-triggered firing, as seen in Figure 4-6. There responses were stable within a session and across sessions in the same virtual environment. Notably, the cells in Figure 4-6 (a) and (b) were simultaneously recorded but have seemingly uncorrelated responses between environments, suggesting that cells do not hold their relationships between environments. This result suggests that these changes in firing may be analogous to the “global remapping” seen in place cells when rodents are placed different environments.

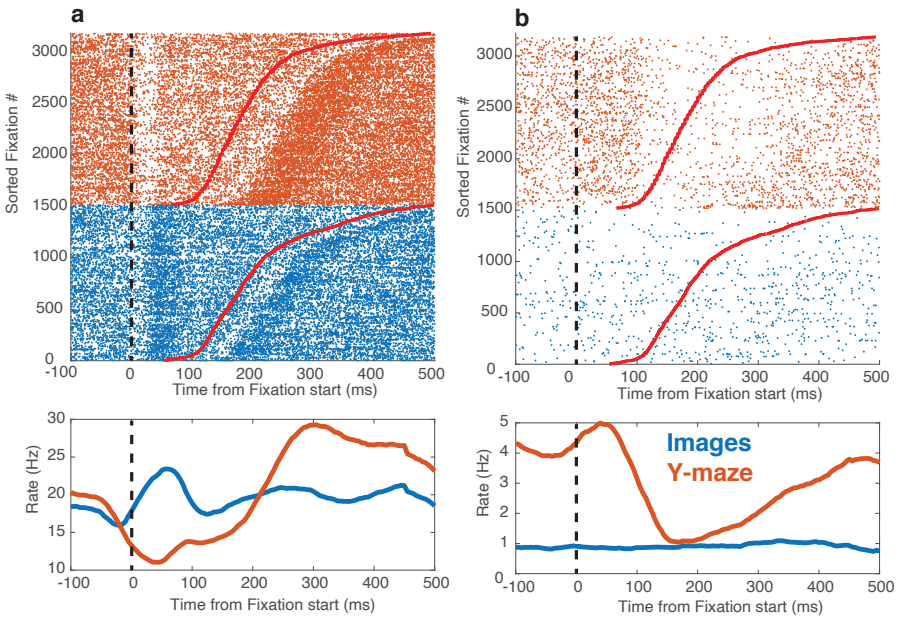

Figure 4-5: Task Specific Eye Movement Modulation

Foveation triggered raster plots and rate maps for two cells during image viewing (blue) and for the same cell during onscreen Y-maze foveations (orange). The red line in the upper plot indicates the end of the foveation in question.

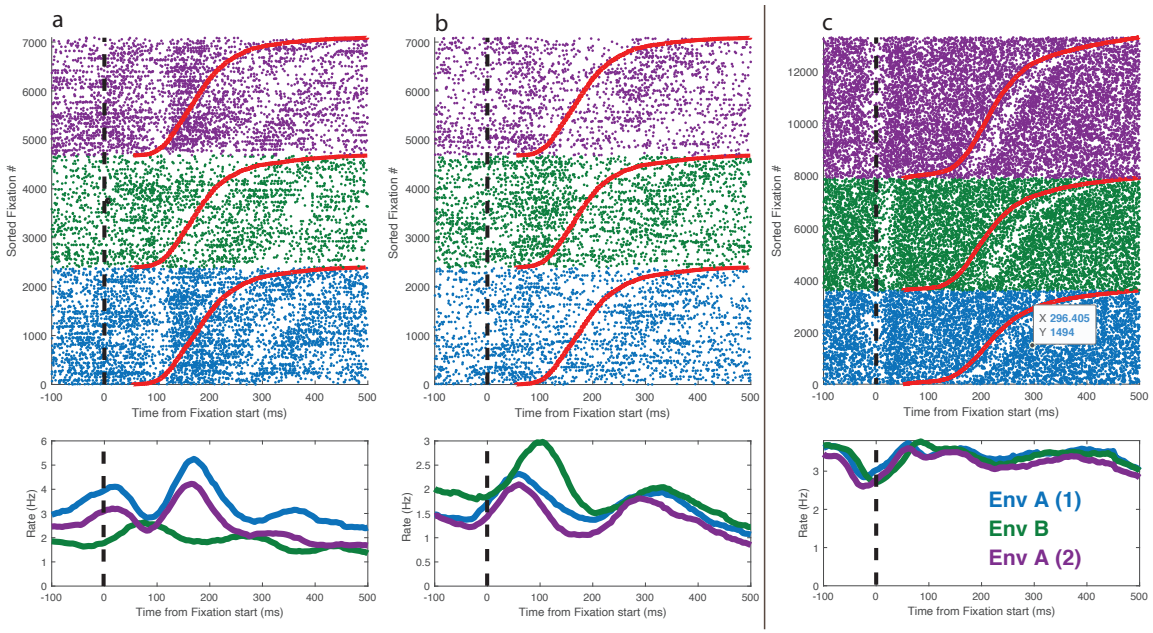

Figure *4-6*: "Remapping" of eye movement responses in different visual contexts

(a-c) Fixation-triggered responses of each cell during on screen fixations in the Y-maze. Each color denotes a distinct session, with blue and purple sessions taking place in the same environment (Env A) and the green session in a different environment (Env B). Fixations are sorted by duration, indicated in red. Note that (a, b) were simultaneously recorded.

4.5 Conditions without eye movement modulation

In a survey of automatic sorting and limited manual sorts from one monkey well trained in our color navigation task, we have yet to see any cells that show eye movement modulation during the navigation portion of the color task. This is in spite of the fact that we do see cells with specific responses in this task, discussed in Chapter 3. Importantly, the nature of this color-navigation task is such that the monkey sees a screen with a solid color that changes as he "drives" through the space. Thus, while visual input is necessary to complete the task, eye movements themselves do not add task information. This negative result, therefore, suggests that eye movements may gate activity only when the eye movement itself is relevant to the task at hand.

4.6 Conclusions

These data corroborate previous findings of eye movement modulation in the hippocampus. The observation that modulation can vary across tasks and/or even contexts within the same task supports the hypothesis that this eye movement modulation is not a direct motor efference copy of eye movement motor commands.

Instead, changes in responses across tasks and conditions suggest differential activity of a neuron in different situations, that is, that the neurons are engaged in different network level activity. In many ways, this change (or "remapping") in firing relative to eye movements is reminiscent of remapping of place cells in the rodent hippocampus, which can take place across tasks in the same environment or across environments even when the task is held constant (Buzsáki and Moser, 2013). Here, I have shown that eye movement modulation can follow a similar phenotypical pattern. Further, in Figure 4-5 and Figure 4-6, I have shown examples of cells that were simultaneously recorded changing their responses. In both cases, the changes by each cell are different. This suggests that, similar to the phenomenon of "global remapping" in the rodent, cells with a relationship in one network state do not necessarily maintain this relationship in a different network state.

What is more, the existence of eye movement modulation outside the hippocampus suggests that this mode of network activity may not be hippocampus specific. On the one hand, this result is not surprising; eye movement modulation of individual neurons has been noted as early in the visual stream as the LGN (Reppas et al., 2002), and there is no reason to suspect that such activity would remain isolated to the

hippocampus. However, the possibility of eye movement correlates in non-temporal cortical areas (e.g in Figure 4-4), for example, leave open the possibility of using eye movement triggering as a tool for studying interactions between brain areas. Further work will be needed to quantify this non-hippocampal activity, and determine to what extent it is phenotypically similar to activity seen in the hippocampus.

It should be said that with the exception of Figure 4-3, this chapter has focused on specific examples of eye movement correlates and the impact of task and task content. In Chapter 2, we saw evidence that while a notable subpopulation of neurons remap between Y-maze task contexts, a majority of neurons do not. Further analysis is needed to see (a) if these neurons that change their eye movement modulation come from this pool of remapping cells (b) if these neurons preferentially carry visual information or if their eye movement modulation is network rather than stimulus dependent (c) if similar populations of task or even eye-movement stable cells exist in the hippocampus and (d) if neurons showing eye movement modulation primarily draw from the same subpopulation across all tasks and contexts, or if a recruitment can be at random from the hippocampal population.

In a similar vein, I have focused here on neurons with eye movement modulation that can be seen across a predefined phase of our Y-maze task. If anything, this mode of investigating is under-sampling the role of eye movement modulation in the monkey hippocampus. (Nowicka and Ringo, 2000a) and (Konig and Buffalo, 2021), for example, have shown situations in which eye-movement triggered activity is specific to when a monkey is attending a particular part of a screen. In our case, this leaves open the possibility that eye movement relationships only hold when the animal is attending to the stimulus that might drive any particular cell. Evidence for a similar phenomenon has also been seen in the amygdala (Minxha et al., 2017). In fact, the delay in correlates with visual information that we see in our Y-maze task could be the result of stimulus-specific activity occurring with a reliable offset from the stimulus that drives it. Although far from conclusive, support for this idea comes from the observation that the delay in visual information and delay in eye movement modulation take place over similar timescales.

Though extremely preliminary, our present inability to find eye movement modulation in a task where eye movements cannot contribute usable task information (i.e. our color game) provides additional support for the idea that a core function of the hippocampus is tracking task progress. When this progress observed through eye movements then eye movements shape hippocampal activity, and when it arrives

through other perceptual means then eye movements and hippocampal activity de-couple. These results should be interpreted with caution: automatic sorting does not provide as complete a picture of our recorded cells, and our limited population size make it difficult to rule out the possibility that we have simply missed the relatively small subpopulation of cells with eye-movement relationships. Nevertheless, a low incidence of eye-movement related activity during our color task is consistent with rodent results showing that theta oscillations and spike timing are phase coherent with sniffing only when that sniffing is relevant to the rodent's task (Berg et al., 2006; Grion et al., 2016).

It is worth noting a possible disconnect between our findings of the lack of eye-movement correlation in certain conditions and the observation by Ringo and others (Ringo et al., 1994) of eye-movement related activity in the hippocampus under all conditions, even those in which no explicit stimulus was present. One potential explanation for this difference is the possibility that monkeys in these studies, although not explicitly doing a task or fixating a target, are still engaged and actively exploring their visual environment. We find that monkeys sitting in the dark, for example, frequently attend to noises, look at the door to their room to see if an experimenter might pay a visit, or otherwise engage in a "covert" task not specifically mandated by the experimenter. In our color task we have the ability to, at least in part, command the monkey's attention in an eye-movement independent way. With that said, further data from hand sorting and a second monkey are needed to validate our negative result.

Taken together, the data in this chapter add support to the notion that eye movements can organize information in the hippocampus. Whether this organization is simply a result of mediating the arrival of information to the hippocampus, or whether it is the result of some broader network process, merits further experimentation. Monkeys lack the consistent theta oscillation seen in rodents, and either case may point to the possibility that eye movement modulation serves a similar organizational mechanism for the hippocampal network activity, although further more direct causal manipulation is needed to empirically test the role of this modulation in hippocampal processing.

5 Methods for the Analysis of Spike Trains

Even though behavioral electrophysiologists often discuss behavioral "correlates" of neural activity, there is strikingly limited agreement on how to statistically characterize these correlates. This problem is

particularly acute in studies of the monkey hippocampus, where it is not immediately clear what, if any, behavioral features a neuron's activity might encode. While neurons in rats are famous for their punctate and reliable spatial correlates (i.e., "place fields"), work in the monkey has not yet identified such a reliable predictor of hippocampal activity. In fact, even studies that have identified spatial activity in the monkey hippocampus have failed to identify the discrete responses so well-characterized in rodents. From a practical standpoint, this means that many of the statistical tools that are often used in rodent studies lack the nuance or noise-resilience to be useful in characterizing monkey recordings.

This chapter serves two simultaneous purposes. On one hand, I have written it as a broad overview of certain statistical techniques. Bootstrapping and Monte Carlo simulations of information theoretic and correlation metrics are broadly used to characterize spike-train-predictor relationships (Baraduc et al., 2019; Gulli et al., 2020; Konig and Buffalo, 2021; Wirth et al., 2017). Similarly, cross validated Generalized Linear Models (here I focus on the Poisson case) are increasingly becoming commonplace in identifying single-predictor relationships, as well as relationships to co-occurring predictors (Hardcastle et al., 2017; Mallory et al., 2021; Tsao et al., 2018). However, in the course of the work presented elsewhere in this thesis, I have expanded on a number of these commonly used statistical techniques. This work is discussed concurrently in this chapter and includes (1) using information gain across conditions as a way to characterize changes in a given cell's response (i.e., remapping) (2) implementing a Monte Carlo simulation to determine p-values for additional contributions of co-occurring predictors and (3) adapting the time rescaling theorem from a goodness-of-fit metric into an alternative, model-based significance test for remapping. Elsewhere in this thesis I use many of these statistical techniques on my own monkey data. Here, however, I present only illustrative set of examples.

5.1 Information theory and descriptive statistics

We wish to know if a neuron's spiking co-varies in any way with a given predictor. The most common descriptive correlation metrics assume a monotonic relationship between the two variables in question[9]. In the case of spike train-predictor relationships, there is no reason to assume that a neuron's spike rate

[9] Pearson's *r*, for example, assumes a linear relationship while Spearman's *rho* and Kendall's *tau* rely on rank-ordering to relax the assumption of linearity but preserve the assumption of a monotonic relationship

will maintain either a linearly or non-linearly monotonic relationship with a predictor. On the contrary, a well-characterized hallmark of many kinds of neural activity is that neurons will show punctate relationships with a particular value of a predictor. Neurons in V1, for example, will fire for a particular orientation or a horizonal bar but be silent elsewhere (Hubel and Wiesel, 1959). Similarly, the canonical rodent place cell will have a punctate firing field at a particular part of a track, but be unresponsive elsewhere (O'Keefe, 1976). Shannon (1948) provides a framework known as information theory for quantifying these non-linear, non-monotonic relationships. Skaggs et al. (1996, 2007) extend this idea to consider only times when a neuron is firing. Similarly, Treves and Rolls (1994) quantify the limited span of a spike train-predictor relationship with a metric known a "lifetime sparseness." In this section I will provide an informal background on these information theoretic approaches, their underlying assumptions, and their use in describing neural data. I will further demonstrate the use of Monte-Carlo simulations to provide a meaningful significance test for the relationship between a spike train and predictor. This section is largely review - these methods are well established and widely used in modern neuroscience - but it provides the theoretical groundwork for discussion of novel approaches in the subsequent sections.

5.1.1 Shannon's mutual information

For a moment, let us move away from neurons and think about two random variables, X and Y. Shannon's Information theory was developed to answer the question: how much does knowing about Y tell me about X? This concept is known as "mutual information" (Shannon, 1948).

To define mutual information, Shannon posed the specific question: how much does knowledge of Y decrease uncertainty in X. In order to measure uncertainty, Shannon borrowed the concept of entropy – that is, disorder – from physics to give the average "surprisingness" of a particular observation being drawn from X. In the discrete case (where X is made up of possible observations, $\mathrm{x_i} = \{\mathrm{x_1}, \mathrm{x_2}, \ldots \mathrm{x_n}\}$ each occurring with probability $\mathrm{p(x_i)}$, the entropy (H) of X is defined as:

$$\mathrm{H(X)} = -\sum_{\mathrm{i}}^{\mathrm{N}} \mathrm{p(x_i)} \log\left(\mathrm{p(x_i)}\right)$$

(Dyan and Abbott, 2005) give a summary of the derivation of entropy, and provide an intuitive explanation for why the log function is included. A log with base-2 is typically employed because it provides direct interpretability; a fair coin toss (with 50-50 odds) yields an entropy of 1. While entropy is itself

dimensionless, a unit of “bits” is traditionally employed with one “bit” being the information gained from knowing the result of such a coin toss. It should also be noted that the negation of this expression is simply a convention to ensure that $\mathrm{H}(\mathrm{X}) > 0$, because all $\mathrm{p}(\mathrm{x_i})$s are probabilities and must occur between 0 and 1.

If entropy is the uncertainty inherent in X, then the mutual information between X and Y can be stated as the reduction in the entropy of X that occurs when we have knowledge of Y. This can be expressed mathematically by considering the distribution of X conditioned on Y ($\mathrm{P}(\mathrm{X}|\mathrm{Y})$), and subtracting its entropy from the entropy of X:

$$\mathrm{I}(\mathrm{X};\mathrm{Y}) = \mathrm{H}(\mathrm{X}) - \mathrm{H}(\mathrm{X}|\mathrm{Y})$$

For historical reasons, examples of mutual information often have to do with lines of communication. I will borrow from an example to illustrate the limits of mutual information (Dyan and Abbott, 2005; Stephenson, 2012). Consider a telephone, where a sender speaks a set of words that have distribution X and a receiver hears the resulting set of words Y. In an ideal case, there would be a one-to-one mapping between what the receiver hears and the sender sends, meaning there would be no uncertainty in X once Y is known and $\mathrm{H}(\mathrm{X}|\mathrm{Y})$ would equal 0. In this ideal case, the information between X and Y simply be equal to the entropy of X itself. There are times, however, when this ideal communication line would pose a problem; it would be easy for a 3rd party to listen to the conversation between our sender and receiver. By recording Y, this eavesdropper would have access to maximal information about X. Imagine, however, that our original sender and receiver would rather not be overheard. In this case, they would like a channel whereby an eavesdropper gets no information about X. Such a channel would be achieved by encrypting Y such that $\mathrm{H}(\mathrm{X}) = \mathrm{H}\left(\mathrm{X}\middle|\mathrm{Y}_{\mathrm{encrypted}}\right)$, that is, $\mathrm{I}\left(\mathrm{X};\mathrm{Y}_{\mathrm{encrypted}}\right) = 0$. This plan is only worthwhile if the intended receiver knows the transform from $\mathrm{Y}_{\mathrm{encrypted}}$ back to Y, that is, if the intended receiver can decode the original message.

For our purposes, we are not especially concerned with telephones. Instead, our goal is to understand information as it is encoded by neurons. Let’s say we have some stimulus or behavioral feature, and that it occurs with distribution X. Similarly, we record a spike train from a given neuron, which we will say has distribution Y. Mutual information quantifies how much information a cell’s firing contains about a particular feature predictor.

5.1.2 Information in Spike train

Imagine that we have trained a monkey to drive down a linear track and are recording the spiking of an idealized “place cell.” The word “imagine” should be taken literally here; Figure 5-1 shows a raster plot for a simulated version of such a neuron, as well as a simulation of a neuron with only a weak spatial correlate.

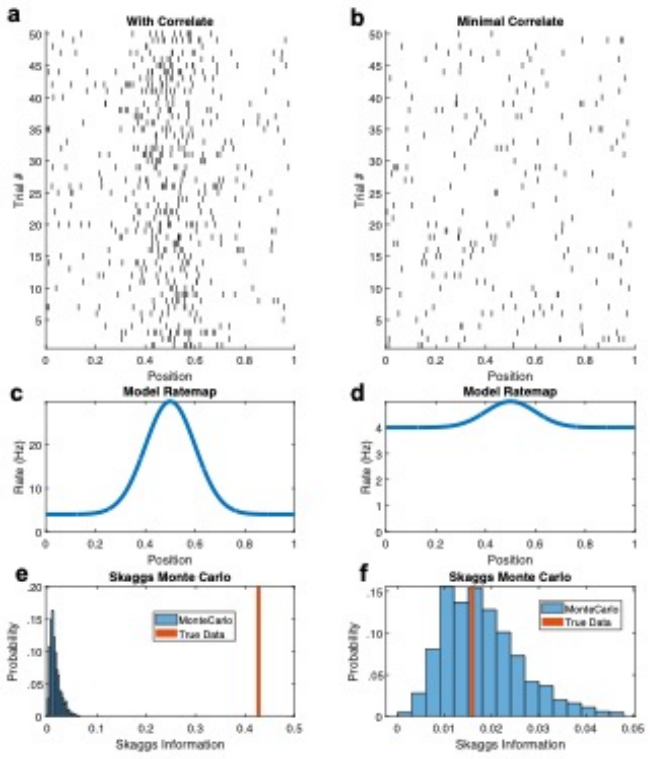

Figure 5-1 Simulated spatial firing

(a) An example raster plot for a strong place cell, simulated using a Poisson GLM from rate map shown in (c).

(b) An example raster plot from a neuron in which the place-specificity is too weak to be visible, with (d) giving the underlying rate map.

(e & f) Histograms of Monte Carlo Simulations (1000 shuffles) of Skaggs information with shuffled spike trains, with the Skaggs information of the “true” data indicated by the orange line. The cell shown in (a & e) has significant Skaggs information significantly greater than shuffled data ($p<.001$) while the cell in (b & f) does not ($p>.1$).

It is worth pointing out a few details about this conversion from continuous stimulus and spike train to probability distributions. First, the definitions of entropy and information used here are discrete. This means that even for a continuous predictor, it is necessary to bin X and calculate the probability of the predictor occupying each bin. Figure 5-2 (a) illustrates the probability that our simulated monkey will occupy any particular bin at any particular moment in time.

Second, whereas our telephone example considered the distribution of "words" sent over the telephone, it is not immediately clear how to define a word in most experiments' spike trains or predictor. The typical approach to this is to further bin time into small, discrete chunks. A word can then be made of one or several of these discrete time bins (for our purposes here, we will consider a single bin at a time). It should be noted that the unit "bits" is applied to the information gleaned (on average) from a single word. The unit for information in a spike train now becomes "bits/word," which is directly scalable to "bits/second."

Third, because spikes are discrete events, our spike train, Y, is naturally discretized. If time bins are taken that are sufficiently small, then within any given bin the cell will either be firing ($\mathrm{y_t} = 1$) or silent ($\mathrm{y_t} = 0$). Larger time bins would permit multiple spikes (e.g., $\mathrm{y_t} = 2,3,\ldots$) but would not change this idea. The conditional entropy, $\mathrm{H(X|Y)}$, can be written as the sum of the conditional entropies for each value of Y, scaled by the probability of that value of Y occurring:

$$\mathrm{H(X|Y)} = \sum_{\mathrm{j}}^{\mathrm{M}} \mathrm{H(X|y = j)} \cdot \mathrm{p(y = j)}$$

Figure 5-2 (b, c) gives the distributions p(y=0) and p(y=1) for our example cell.

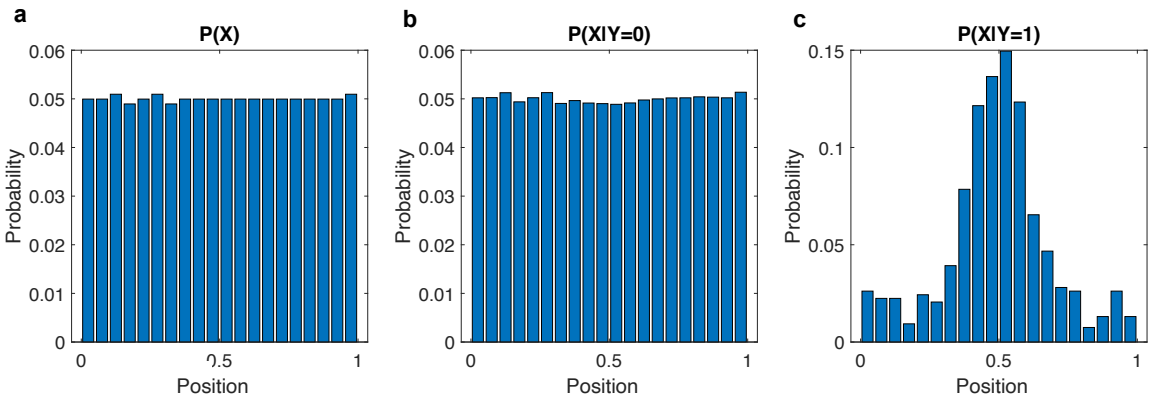

Figure 5-2 PDFs needed to compute information

Shown for the example cell shown in Figure 5-1(a).

(a) Binned Probability distribution for X.

(b) Binned probability distribution for P(X|Y=0), the times when the cell is not firing.

(c) Binned probability distribution for P(X|Y=1), the occupancy of the space when there is a spike.

The actual values of entropy and mutual information are heavily dependent on the number of bins chosen when discretizing data. As a result, information values computed with different binnings cannot be directly compared with one another. Figure 5-3 shows the near-logarithmic increase in computed information as the number of bins increases for our example cell. Section 5.1.5 provides methods both for scaling information scores to make them interpretable, and also demonstrates a procedure for a simulation-based significance test.

Rescaling of information values similarly means information values cannot be compared between different predictors. Section **Error! Reference source not found.** lays out a model-based procedure that allows for comparisons between variables that cannot share a common set of bins.

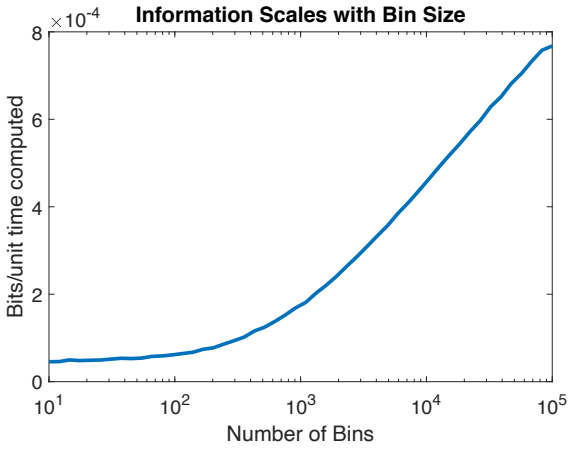

Figure 5-3 Information scales with number of bins.

Information values denote mutual information and are shown for varying bin counts tiling “Position” for the simulated cell in Figure 5-1(a).

5.1.3 Skaggs Information

Shannon's mutual information considers a channel as a whole: each "word" or time bin is considered to be equivalent, and classical mutual information gives the average information carried by a single "word." While this is perfectly reasonable in some cases, it is not clear that this is a good assumption when characterizing a neuron. In most cases, neural firing is extremely sparse; a typical pyramidal cell fires at only a few Hertz. Moreover, neuroscientists typically care about events on a milli-second timescale, meaning that only one in every few hundred bins might actually contain a spike, while the rest have a spike-count of 0.

Instead of considering all time, Skaggs et al. (Skaggs and McNaughton, 2007; Skaggs et al., 1996) provide a formulation of information that considers only the times when the neuron is firing. It is important to highlight this difference in assumption between Shannon and Skaggs information: by considering only times when a neuron is spiking, Skaggs assumes that spiking is all that matters to any reader of that neuron's activity. Shannon's formulation assumes that all times matter equally.

The examples in the previous section provide an illustration of this difference in a neuron: our simulated cell has a relatively low firing rate, meaning that it is silent the vast majority of the time. This means that P(X|y=0) looks almost the same as P(X) and has very similar entropy (in this case, our example cell has H(X) = 4.32 bits and H(X|Y=0) is very nearly equal). Similarly, when we scale our conditional entropies, p(y=0) will be very large (.99 in this example) while p(y=1) will be very small (here 0.01). The result is that the Shannon mutual information between most spike trains and most stimuli will be quite low (I(X,Y)=0.01).

On the other hand, Skaggs information excludes times when the neuron is not spiking from mutual information computed above, so that Skaggs information is given by:

$$\mathrm{I}(\mathrm{X}, \mathrm{Y} = 1) = \mathrm{H}(\mathrm{X}) - \mathrm{H}(\mathrm{X}|\mathrm{y} = 1)$$

The effective unit here is still in "bits/time bin," but might be better stated "bits/time bin with a spike in it." For increased interpretability, Skaggs information is often given in its continuous form:

$$\mathrm{I}(\mathrm{X}, \mathrm{Y} = 1) = \int_{\mathrm{x}} \lambda(\mathrm{x}) \log_2 \frac{\lambda(\mathrm{x})}{\bar{\lambda}} \mathrm{p}(\mathrm{x}) \mathrm{dx},$$

so that this unit becomes “bits/spike.” An incomplete derivation of this equation is given in appendix 5.6.1. For our example cell, H(X|y = 1) =3.17 bits, and so observing a time bin with a spike in it provides 1.15 bits of information about X.

Formulating Skaggs information in terms of firing rates allows for standard rate map assumptions - namely the use of smoothing kernels - to be easily incorporated. However, using kernels will change the computed value of Skaggs information. Similarly, Skaggs information is susceptible to the same interpretability problems inherent in Shannon information, and requires similar considerations in its interpretation (again, see section 5.1.5).

5.1.4 Lifetime Sparseness

Skaggs information is similar to another frequently used descriptive statistic, “lifetime sparseness,” which was introduced by (Treves and Rolls, 1994). Roughly, sparseness measures the dispersion of a rate map. It is given by taking the variance of the rate map, $\sigma^2 = \langle\lambda(x)\rangle^2$, and dividing it by the expected value of the squared firing rate:

$$a = \frac{\langle\lambda(x)\rangle^2}{\langle\lambda(x)^2\rangle}$$

In the discrete case with n bins over the stimulus, this takes the form:

$$a = \frac{\left(\sum_{i=1,n}\left(\frac{\lambda(x_i)}{n}\right)\right)^2}{\sum_{i=1,n}\left(\frac{\lambda(x_i)^2}{n}\right)}$$

Baked into lifetime sparseness is an implicit assumption that all stimulus bins are occupied equally, something not assumed by the information metrics in the previous sections. However, the ratio sparseness has the distinct advantage of being qualitatively interpretable. A neuron with a very narrow firing field will have a low sparseness because σ^2 will be small. A neuron with uniformly distributed firing will have $\sigma^2 = \langle\lambda(x)^2\rangle$, and thus have a sparseness 1.

Because this definition of sparseness scales oppositely to metrics like Skaggs information (where high values imply a sparse firing field) and because sparseness is bounded between 0 and 1, some authors choose to rescale sparseness to avoid confusion:

$$\text{sparseness} = 1 - a$$

Further, lifetime sparseness should not be confused with a similar metric, population sparseness, which considers the relative responses of all neurons in a population to a single stimulus.

5.1.5 Interpretability and Monte-Carlo Simulations

The challenge with all these metrics is that there is not a known distribution or expected value for them. This problem is confounded by the fact that changing the number of bins can change the computed value for each metric, meaning that the way in which data are processed can change the outcome of these values. The most common solution to this problem is to empirically simulate the distribution of values for what each metric would be a in a “random” case and ask if the observed value might have come from that distribution.

What “random” means depends on the question at hand. For example, we often use information to ask if a cell has a relationship with some predictor. A given cell’s spike train, however, is going to have internal features and statistics that could impact the observed information. A bursting cell, for example, will have many spikes together and therefore appear to carry information even in a spurious burst. It is therefore important that the random distribution preserve the firing statistics of the original data, with only the spike-predictor relationship randomized. This is most often accomplished by shifting the cell’s spike-train relative to the predictor in question (or visa-versa), while preserving as much as possible the spike timing relationships within the spike train. This process is illustrated in Figure 5-4. For data with a rigid trial structure, it is often more useful to rotate within each trial, so consistent activity at one part of the trial does not simply become shifted to a different part of the trial. Similarly, trial identities might also be shuffled when the trial identity also carries information about the predictor (for example, that the right side of a maze might only be traversed on some trials).

This process of rotating data is known as Monte-Carlo simulation. Some literature will refer to it as the bootstrap. Here we will differentiate these methods by using “bootstrapping” to specifically refer to the process of estimating the variance of a statistic using resampling, in contrast to “Monte-Carlo simulations”, which are used to empirically simulate the null distribution for a statistic.

The result of many Monte-Carlo simulations is a null distribution for what the statistic in question would have been if it had no relationship with the randomized predictor. This value can be used to ask the probability of observing the true data given this null distribution, a quantity commonly known as a p-value.

This p-value can be compared to a threshold as a significance criterion for a statistic (typically .01 or .05), allowing us to say that the value of the statistic is significantly different from what was expected by chance. This test is demonstrated graphically for our example neurons in Figure 5-1.

It is worth noting that, when data are binned in the same way, information theoretic statistics can be directly compared to one another; we can, e.g., say that one cell carries more information about a predictor than another. Having said this, it is also true that computed information value still depends on the statistics of the spike train in question. Thus, by comparing information one can learn something about the behavior of the cell, but not whether a given cell is behaving differently than expected by chance.

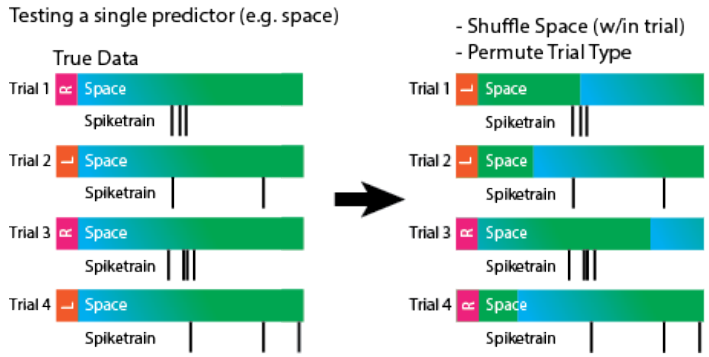

Figure 5-4 Shuffling procedure

Illustration of a shuffling procedure for the Y-maze experiment. In this case, spike trains are held constant while trial type and location data are permuted randomly.

5.2 Information Gained with a Partition

Another way of looking at Shannon information is to view it in terms of partitioning data. Each conditional entropy can be viewed as the entropy for a subset of the data. In the case of spikes, this means that we are considering entropy separately where there are no spikes, one spike, etc. and computing information by weighting each of these conditions. This logic can also be applied to information theoretic metrics for data that are collected under different conditions. Specifically, we can ask if considering multiple conditions separately provides more information than ignoring the conditional distinction, that is, the gain in information by considering the partition.

Mathematically, this gain in information is expressed:

$$I_{gain} = \sum_{a=1}^{A} \frac{n_a}{N} I_a - I_{All},$$

where each A is the set of partitions in question of the data, such that n_a is the number of datapoints in partition a and I_a is the information within said partition. I_{All} is the total information if all data were considered as a single condition. Just as Shannon information must always be greater than zero, the information gained must also be greater than or equal to zero, meaning that information cannot be lost from the partition.

There are, of course, many possible applications for considering information between partitions. The most intuitive example is remapping experiments. Here, we have two visual contexts, and we want to know if changing the context impacts a given cell's response. Another way of phrasing this is that we want to know if we gain information about the cell's response by considering these two conditions separately versus considering them together.

As with other information theoretic metrics, there is no underlying model that provides an expected distribution for information gain. Because the gain must always be greater than zero and must be proportional to the information carried by the cell in the first place, it again becomes necessary to perform a Monte Carlo simulation to produce a null distribution to which the observed value can be compared. In this case, our question is no longer about spike-predictor relationships, but instead about spike-condition relationships. As a result, we no longer rotate spike trains within trials but instead shuffle the conditional identities between trials to produce a null distribution for I_{gain} when the condition is randomized. This can be compared to the observed value of I_{gain} by way of a p value. Figure 5-5 shows example raster plots

for simulated remapping and non-remapping cells, as well as a graphical demonstration of this permutation-based statistical test.

Information gain is a longstanding metric in physics and other applications of information theory. However, this is (to the best of our knowledge) its first application in a remapping study.

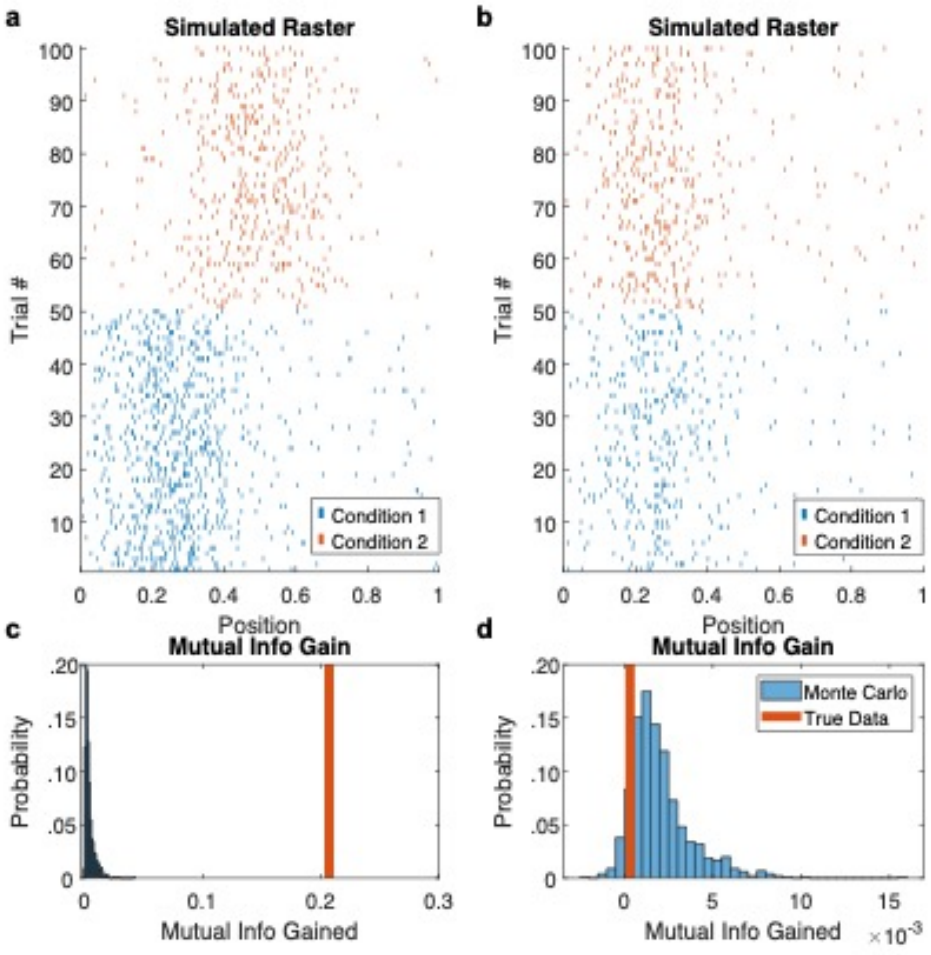

Figure 5-5 Simulated remapping Cells

Simulated Remapping (a) and non-remapping (b) cells generated using a Poisson Generalized Linear Model.

(c) Distribution of gains in mutual information given condition for the cell shown in (a), with the true gain denoted by the orange line. There is a significant gain for this cell with $p<0.001$.

(c) Same for (b). Here the gain is not significant ($p>0.1$). Although Mutual information must be greater than 0, the very small values of <0 here are a result of numerical error given the bin size used.

5.3 Generative Models

At their core, the methods described thus far are descriptive; metrics of information or sparseness describe existing data, but they cannot tell anything about data that have not yet been seen by the metric or model. Often, however, we wish to have a generative model, that is, a model that predicts new data from the data that it has been trained on. To do this, it is common to employ generative models, that is, models that provide explicit predictions for data not used in their construction.

A further drawback to the information theoretic approach discussed above is that it makes comparisons between different predictive variables difficult. I have previously demonstrated that information is not comparable between different parcellations or binning of data, even of the same predictor. This criticism extends to comparisons between predictors where using different binning becomes a necessary feature of different predictors existing in different spaces. In what follows, I will demonstrate a predictor-agnostic method for model evaluation that allows for direct comparison between models trained using different features.

5.3.1 General linear models

Let's say we have a neuron that fires with spike rate λ. In a linear model, we would model the spike rate of this neuron as a linear combination of some set of N features or predictors, $\mathrm{X} = \{\mathrm{x}_1, \mathrm{x}_2, \ldots, \mathrm{x}_\mathrm{N}\}$, weighted by a set of parameters $\beta = \{\beta_1, \beta_2, \ldots, \beta_\mathrm{n}\}$, and with optional constant β_0. Mathematically, we would express this function:

$$\lambda = \mathrm{X}\beta = \mathrm{x}_1\beta_1 + \mathrm{x}_2\beta_2 + \cdots + \mathrm{x}_\mathrm{N}\beta_\mathrm{N} + \beta_0$$

Note that, although there is no explicit time component in this formulation, both λ and X are typically functions of time (e.g., $\lambda = \lambda(\mathrm{t})$). Practically, time can be discretized with the time unit representing the feature of interest. We might, for example, model spiking of some pre-determined period of interest in a trial. Here, we will focus on the case where the predictor, X, represents a task variable that can change continuously (such as an animal's location in space). In this case, it is useful to define arbitrarily small time bins and model the firing rate in each bin as a separate observation of the model.

In the linear case, the estimate for each β value can be solved by arranging all observations of λ and X as follows:

$$\begin{bmatrix}\lambda_1\\ \lambda_2\\ \vdots\\ \lambda_T\end{bmatrix} = \begin{bmatrix}1 & x_{1,1} & x_{1,2} & \cdots & x_{1,N}\\ 1 & x_{2,1} & x_{2,2} & & x_{2,N}\\ \vdots & & & \ddots & \\ 1 & x_{T,1} & x_{T,2} & \cdots & x_{T,N}\end{bmatrix}\begin{bmatrix}\beta_0\\ \beta_1\\ \beta_2\\ \vdots\\ \beta_N\end{bmatrix}$$

In this formulation, we can find the cumulative error between our model and all observations, which is given as the sum of squared residuals (i.e., residual sum of squares, RSS) between the model and spike rate:

$$\mathrm{RSS} = \sum_{t=1}^{T}(\lambda_t - x_t^T\beta)^2 = (\lambda - X\beta)^T(\lambda - X\beta)$$

Using this, we can find the estimate of β, $\hat{\beta}$, that minimizes RSS for any given sets of λ and X:

$$\hat{\beta} = (X^TX)^{-1}X^T\lambda$$

Fundamentally, this model can now do something that the information theoretic methods above could not do. For any new observation of X, we can use $\hat{\beta}$ to generate a prediction of λ, $\hat{\lambda}$. This prediction, however, is predicated on the assumption that λ is a linear function of X and of β. More specifically, we are assuming the spiking variance of our model is normally distributed around a mean value of λ, $\overline{\lambda}$ (if a constant is used, $\beta_0 = \hat{\lambda}$). This assumption lends itself nicely to error metrics to assess how well our model works; by taking the ratio between variance (under the normal model) explained by our model and the variance actually present in our data, we can determine the fraction of overall variance that we have captured with our linear fit. This quantity is known as the coefficient of determination, R^2:

$R^2 = \frac{\Sigma(\hat{\lambda_t}-\bar{\lambda})^2}{\Sigma(\lambda_t-\bar{\lambda})^2}$.

5.3.2 Generalized linear models

In the vast majority of cases, normality is an inappropriate assumption for neural spike trains. Any truly normally distributed variable would allow for negative values, and yet neural firing rates can never be negative. Further, action potentials are discrete, so that only whole numbers (i.e., 0,1,2,3,…) of action potential counts can be observed over a given time interval. Given these non-linearities, the general linear model discussed above can rarely provide a good estimate of a single neuron's firing over a small time interval.

Rather than a normal model, it is generally considered more appropriate to approximate the firing of a single neuron with a Poisson model. Under this model, the probability of observing k spikes over a given time interval is expressed:

$$\mathrm{P}(\mathrm{count} = \mathrm{k}) = \frac{\lambda^{\mathrm{k}} \exp(-\lambda)}{\mathrm{k!}}$$

where λ is again the instantaneous firing rate of the neuron and ! denotes the factorial function. The Poisson model assumes that the number of spikes is a whole number (i.e., that it is discrete and non-negative), and that spike events cannot occur simultaneously. These are sufficiently reasonable as assumptions for spike-train models for ourselves and others (Hardcastle et al., 2017; Mallory et al., 2021; Park et al., 2014; Pillow et al., 2008) to use Poisson models in characterizing single neuron firing. However, these models also assume that spiking events are independent such that one spike does not affect another, nor does it affect the overall firing rate. This assumption is known to be inaccurate; bursting neurons provide an obvious example for a case where this assumption fails (Pillow et al., 2005, 2008; Weber and Pillow, 2017). Although methods exist for successfully capturing history dependent non-Poisson features of spike trains using autoregressive filters (see (Weber and Pillow, 2017)), these methods are beyond the scope of this thesis.

A generalized linear model modifies the assumption of the normal “general linear model” by imposing a non-linearity, sometimes called a “link function”, in the equation for linear regression above. This is done by applying a non-linear function, $\mathrm{g}(\mathrm{x})$, to the left-hand side.

$$\mathrm{g}(\lambda) = \mathrm{X}\beta = \mathrm{x}_1\beta_1 + \mathrm{x}_2\beta_2 + \cdots + \mathrm{x}_\mathrm{N}\beta_\mathrm{N} + \alpha$$

Confusingly, this non-linearity is sometimes referred to as a “generalization” of this modeling technique, leading to such a model being dubbed a “Generalized Linear Model.” This terminology is so confusing that its creators are on record as apologizing for its broad adoption in the field (McCullagh and Nelder, 1989). To avoid confusion some (see (Aljadeff et al., 2013; Shapley and Victor, 1980), or (Hardcastle et al., 2017) for an MTL example) use an equivalent terminology in which they specify this generalization as a linear-nonlinear model, denoting the specific model components. In this document, we will discuss Poisson generalized linear models as Poisson GLMs or simply GLMs, which would be considered by other authors as linear-nonlinear-Poisson models. Notation is further complicated by fact that models are sometimes more useful when expressed in terms of the inverse-link functions, $\mathrm{g}^{-1}(x)$, i.e.:

$$\lambda = g^{-1}(X\beta = x_1\beta_1 + x_2\beta_2 + \cdots + x_N\beta_N + \alpha)$$

A Poisson GLM imposes a log link function (McCullagh and Nelder, 1989). The resulting model is:

$$\ln(\lambda) = X\beta$$

Whereas the β values for the truly linear model can be fit using ordinary least squares, βs for the Poisson GLM need to be determined using maximum likelihood optimization. This approach is discussed in detail in Appendix 5.6.2. Briefly, the likelihood of the model is product of the probability (under the Poisson assumption) of observing all given data points. A model is trained by selecting βs that maximize this likelihood function. In practice, optimization takes place on the log likelihood function, $\mathcal{L}$, though this approach will converge on the same β values.

Much like the residuals of linear model, the log likelihood function lends itself to model evaluation. In this case, the mean rate model would be a "null" model fit using one constant parameter. This would be the same as estimating the firing rate of a neuron as that neuron's mean firing rate, regardless of the influence of any predictors (that is, as a homogenous Poisson process):

$\ln(\lambda_{\text{null}}) = \alpha$.

Effectively, $\mathcal{L}_{\text{null}}$ gives the worst log likelihood we could achieve for this data. What we want to know, though, is what fraction of the available likelihood we capture using our model. To do this, we also need to define $\mathcal{L}_{\text{Saturated}}$, the best log likelihood we could hope to achieve. This would be the equivalent of having X equal the full spike train, that is, the equivalent of having one unique parameter for every data sample. Combining these two measurements with the overall log likelihood for our model, $\mathcal{L}_{\text{model}}$, we can define a "fraction of likelihood explained" metric that is similar in spirit to R^2 for a linear model. Because of this model, this metric is often known as pseudo R^2 (Cameron and Trivedi, 2013; Kraus et al., 2013):

$$\text{psuedo}R^2 = \frac{\mathcal{L}_{\text{model}} - \mathcal{L}_{\text{null}}}{\mathcal{L}_{\text{Saturated}} - \mathcal{L}_{\text{null}}}$$

Note that the highest pseudo R^2 possible is 1, which would imply a perfect model fit. There is, however, no lower bound to this metric because poorly fit models can do infinite damage to estimating the data.

The key advantage of this pseudo R^2 metric is that it is dataset agnostic: pseudo R^2 values can be compared between cells, datasets, and even different size subsets of the same data. The metric is, however, agnostic to the number of parameters used in fitting a model. While it assumes at least one

parameter (the mean rate), it cannot differentiate between models with different numbers of parameters. This means that it is a very useful metric for e.g., cross validation, but that relying on pseudo R^2 as a descriptive statistic risks overfitting or inadvertently penalizing under-parameterized models.

It is worth mentioning that there are a number of other commonly used, likelihood-based metrics that take into account the number of model parameters used. Chief among these is Akaike's information Criteria (AIC), which penalizes models linearly for their number of parameters, m:

$$AIC = 2m - 2\mathcal{L}_{model}$$

Here, convention dictates that the model with the lowest AIC is considered the best. Similarly, the Bayesian Information Criteria (BIC) penalizes as a function of the log of the number of data points included, T:

$$BIC = m \ln(T) - 2\mathcal{L}_{model}$$

Both of these metrics, however, include a non-normalized $\mathcal{L}_{model}$ term. Because the value of this term is dependent on the data in question, these metrics are of limited utility for comparing data from different time epochs or different neurons. They are, however, applicable in selecting from amongst multiple candidate models for a single dataset.

5.3.3 Model parameterization

Before moving into a discussion of GLM evaluation, it is worth mentioning the construction of predictor matrix X, particularly with respect to continuous variables. Similar to the information metrics above, X must be broken down into a set of distinct components, each of which can be fit with its own β value, so that each row of X is a discrete sample and each column represents one of its constituent components. One option for doing this is simply divide a predictor into discrete bins, and then to assign a unique β to each bin. This would, effectively, transform each column of X into a categorical membership variable:

$$x_{n,t} = \begin{cases} 1 & \text{in bin } n \\ 0 & \text{otherwise} \end{cases}$$

The trouble with this method is that is it can lead to an explosion in parameters; to sufficiently tile any particular space, one could need a large number (tens to hundreds) of bins to capture it (in our Y maze, we use between 64 and 256 bins to tile avatar position space and hundreds to tile avatar-position space).

Because the maximum likelihood algorithm optimizes over all parameters, this large number of bins lead to long model runtimes, potentially rendering fits on large numbers of cells or Monte Carlo simulations too computationally expensive to be useful.

An alternative approach (Pillow et al., 2008) is to define a set of basis functions that span the predictor in question. Bins are, in effect, square basis functions. However, it is more useful to define a set of basis functions that overlap, so that the interaction between these functions allows the space to be represented with fewer overall basis functions. One easy candidate for such a function is to tile the space with gaussian bumps representing each component of x as a function of the predictor, w, the center of the bump, μ, and a width parameter, σ^2:

$$x_n(t) = ae^{\left(-\frac{1(w(t)-\mu_n)^2}{2\ \sigma^2}\right)}$$

Centers and width can be chosen so as to span the predictor in question. The constant a is unimportant, because it will simply rescale β values during model fitting.

Though useful, the long tails on gaussian bumps can lead to unexpected behavior during model fitting, with under-sampled or low spike rate neurons producing λ values that approach infinity near the edges of sampled predictor w. One approach to fix this problem is to employ a “raised cosine” bump (Pillow et al., 2008). This function has similar behavior to a gaussian, but without infinite tails. Here σ is a width parameter and r specifics the rate of roll-off from around the bump’s edges.

$$x_n(t) = \begin{cases} 1 & |w(t)| \leq \frac{1-r}{2\sigma} \\ \frac{1}{2}\left(1 + \cos\left(\frac{\pi\sigma}{r}\left(|w(t)| - \frac{1-r}{2\sigma}\right)\right)\right) & \frac{1-r}{2\sigma} < |w(t)| \leq \frac{1+r}{2\sigma} \\ 0 & \text{otherwise} \end{cases}$$

An alternative approach, and the one used elsewhere in this thesis, is to re-normalize all gaussian bumps so that there is even coverage of the space. This has the effect of fixing edge artifacts created by long tails. This re-normalization is achieved by dividing each basis function by the sum of all other basis functions for each value of the predictor:

$$x_n(t) = \frac{e^{\left(-\frac{1(w(t)-\mu_n)^2}{2\ \sigma^2}\right)}}{\sum_{m=1}^{N} e^{\left(-\frac{1(w(t)-\mu_m)^2}{2\ \sigma^2}\right)}}$$

with the result that, for any value of $w(t)$, $\sum_{n=1}^{N} x_n(t) = 1$. This is a version of spline interpolation, and is sometimes referred to as “basis-spline” or “b-spline” fitting.

Figure 5-6 shows examples of some of the basis types employed elsewhere in this thesis. Figure 5-7 shows the result of using these basis functions to estimate the rate-map from the cell in Figure 5-1(a).

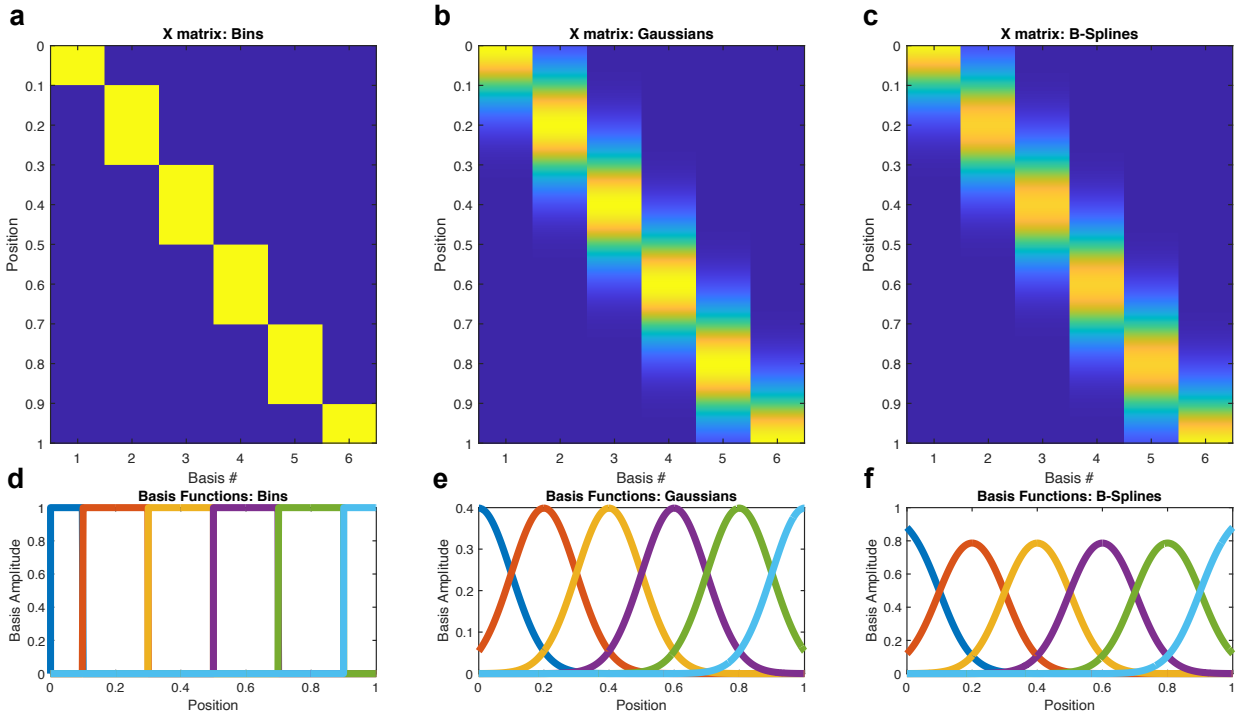

Figure 5-6 Basis Function Examples

Examples of X matrices (a-c) and their underlying basis functions (d-f). In each case, 6 basis functions are used to tile the space.

(a, d) Discrete Bins

(b, e) Gaussian Bumps

(c, f) Basis splines made from re-normalizing gaussian bumps.

Raised Cosines are not used in this thesis and are not shown. See Agarwal et al., 2014 or Pillow et al., 2008 for an examples.

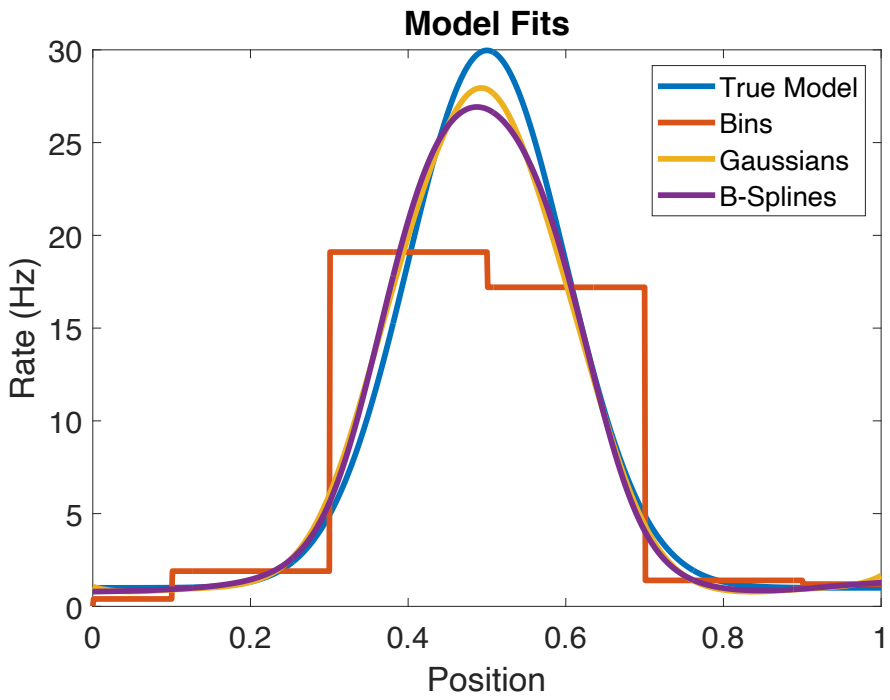

Figure *5-7* Examples of model fits with different basis functions

Examples of rate-maps fit using basis functions in Figure 5-6 to represent the simulated neuron in Figure 5-1(a). The true model is the rate-map underlying the simulation (i.e., in Figure 5-1 (c)).

5.4 Using generalized linear models

5.4.1 Monte Carlo Simulations of one predictor

We can use pseudo R^2 like any other descriptive statistic by training a model on a given dataset and then asking how well our model describes this data. This approach suffers from the same drawback as many of the information theoretic metrics described above; when training and testing on the same dataset, we get no information about how well a model predicts new data. Instead, what we get is an upper bound (determined by maximizing likelihood) on model performance. It would be arrogant, however, to directly interpret this upper bound. If we have too little data or too many parameters, then we risk overfitting our data with our model and giving the false impression that a predictor can explain variance in our data.

The solution to this problem is the same as for the information theoretic metrics. If we repeatedly shuffle the spike train relative to the predictor matrix, , and fit new models for each shuffle, we can build a "null" distribution for pseudo R^2. Because the data used in fitting these models have the same number of parameters and data points, we know that any measured relationship (positive pseudo R^2) between our shuffled data and the predictor in question is the result of overfitting our data. Even though specific values of pseudo R^2 are themselves uninterpretable, this distribution allows us to determine a p-value for our observed pseudo R^2 and assess whether our model documents a "significant" data-predictor relationship using a predefined threshold.

Although I use pseudo R^2 as an example here, it is worth noting any metric with a linear log-likelihood term (AIC and BIC, for example) will yield the same p-value in these simulations because the size of the dataset and number of parameters are kept constant, as are $\mathcal{L}_{\text{Saturated}}$ and $\mathcal{L}_{\text{null}}$.

5.4.2 Comparing Predictors

Using Monte Carlo simulations to determine how well a single predictor describes a dataset is, frankly, a little silly. Compared to the information-theoretic metrics described in section [1.1], the iterative maximum likelihood approach needed to fit a model with a Poisson GLM is extremely slow. Techniques like GPU acceleration can offer non-trivial improvements in model fitting time, but they cannot approach the computational speed of the relatively simple information-theoretic metrics (Latimer et al., 2015, 2019).

Though these approaches rest on slightly different assumptions and can therefore yield different answers, in practice these differences rarely merit the extra computational cost of a GLM Monte-Carlo simulation.

Why, then, have just I devoted so much of my thesis to describing this method? Even though the single-predictor use case for GLM Monte-Carlo simulations is of limited utility, the same method can be used to compare multiple correlated predictors in a way that information-theoretic metrics cannot. Consider a situation in which we have predictors. For want of a better notation, I will call these א and ב, so that $X = [X_א, X_ב]$ and $\beta = [\beta_א, \beta_ב]$ (Typically for our data, א and ב and space and view information, respectively). If א and ב are independent, we can simply assess the pseudo R^2 for each variable separately to determine whether one, both, or neither of them predicts spiking. Things are trickier if א and ב are not known to be independent. In this case, determining that one feature is predictive of spiking leaves open the possibility that any likelihood or variance we explain is really a byproduct of this correlation and cannot be wholly attributed to that particular feature. Here, it is more useful to ask "does one predictor describe the data beyond what would be expected given all other predictors?"

Consider the model:

$$\ln(\lambda) = X_א\beta_א + X_ב\beta_ב + a$$

Now, say we want to know if ב provides a better description of spiking than we would expect from א alone. In this case, our null likelihood is now the likelihood of a model without any ב terms, i.e., $\ln(\lambda) = X_א\beta_א + a$. This allows us to define:

$$\text{pseudo}R^2_{\text{gain}}(ב, א) = \frac{\mathcal{L}_{\text{full model}} - \mathcal{L}_א}{\mathcal{L}_{\text{Saturated}} - \mathcal{L}_א}$$

Which can be interpreted as the gain in explained likelihood of a model that includes ב over a model that does not.

Again, pseudo R^2_{gain} is not directly interpretable on its own when training and testing data are the same. However, we can compare these predictors using Monte-Carlo simulations. In this case, we want to determine a null distribution where ב has no predictive power. To do this, we build a distribution of models in which we shuffle ב relative to all other model components (א and the spike-train). This distribution can then be used to establish p-value for the observed data, and assess the significance of the specific contribution of predictor ב. This process is illustrated in Figure 5-8.

One advantage of this method is that it protects against poor parcellation of predictor variables. Although useful, a head-to-head assessment of $\mathrm{pseudoR}^2_{א}$ and $\mathrm{psuedoR}^2_{ב}$ risks a situation in which one predictor is better described by the chosen basis set, is given an advantage in this comparison, and is thus more reflective of experimenter bias than actual data explanation. Because the same basis sets are used in each iteration of this Monte-Carlo simulation, this method still provides a valid p-value even if the computed $\mathrm{pseudoR}^2_{\mathrm{gain}}$ values do not reflect the full potential of a given predictor. This final comment ceases to be true if e.g., $\mathrm{X}_{א}$includes information about ב that is other not included in $\mathrm{X}_{ב}$.

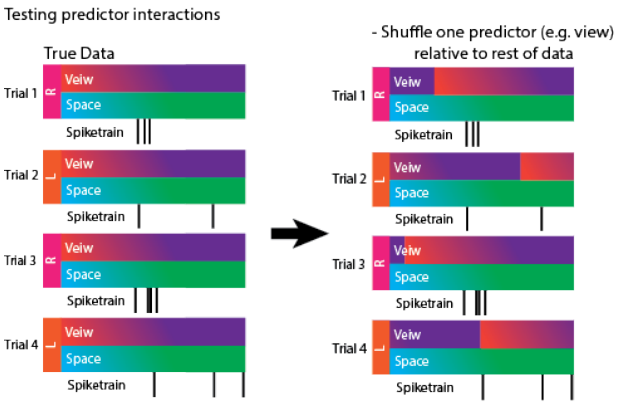

Figure 5-8 Shuffling Predictors

Shuffling procedure with multiple predictors. Here, view is shuffled relative to all other variables to determine what information is lost when visual information is no longer meaningful.

5.4.3 Cross-Validation

A key feature of the generative GLM approach is that we get a method for assessing how well a model predicts new data, in addition to describing goodness-of-fit for existing data. Cross-validation is a widely used model evaluation technique that explicitly characterizes the predictive ability of a given model. In cross-validating a model, we divide our model into training and testing datasets. We then fit a model to the training dataset and evaluate the performance of this model on the testing dataset. Here, again, any of the metrics described above can be used to assess model performance, but I will favor $\mathrm{psuedoR}^2$ because (1) we are now making predictions of new data, which means that we are less concerned about overfitting and (2) it allows values to be directly compared across predictors, across cross validation folds, and across cells.

In order to fully use all available data, model cross-validation rarely involves a single division between training and testing datasets. Instead, it is common to take multiple partitions, or "folds", that split data multiple times and allow multiple samples of model performance within a given dataset. There are many folding methods, and the "correct" method to use can depend greatly on the modeling context. Example folding methods include:

1. K fold: Here, the data are divided into k segments. On each of k iterations, $\frac{1}{\mathrm{k}}^{\mathrm{th}}$ of the data held out for testing while the model is trained on the remaining $\left(\frac{\mathrm{k}-1}{\mathrm{k}}\right)^{\mathrm{th}}$ of the data. Each data segment is only included the testing dataset only one time. Five is a commonly used value of k (20% held out on each iteration).
2. N-1 fold: Similar to k fold, except that on each iteration a single trial or datapoint is held out as the testing data while the remainder is used for training.
3. Bootstrapped: again, as with k fold, a specified fraction of the data is held out on each fold. Unlike k folds, this training data is sampled randomly and with replacement, meaning that greater than k iterations are needed to fully encapsulate all available data. The advantage of this method is that is allows for move comprehensive sampling of the variance of goodness-of-fit measures, while still having enough data to have substantive testing and training sets.

All of these folding methods produce a distribution of psuedoR^2 values, rather than a single value. Using cross-validation to evaluate model fit over a whole dataset, then, requires assessing this distribution. We can use a non-parametric Wilcoxon signed-rank test to evaluate the median performance of a given predictor, or to compare between predictors (Hardcastle et al., 2017). Note that because psuedoR^2 values can be arbitrarily low but have an upper bound at positive 1, the median is a more applicable statistic than the mean when assess these values. Consider the following use cases:

- Say we want to know if a single feature has any predictive power. In this case, we simply care that psuedoR^2 is greater than 0, that is, that we can explain a non-trivial fraction of the variance in our data use the predictor in question. Here, we would use a one sided signed-rank test to assess the probability that the true median is greater than 0.
- Now, say we want to compare predictors א and ב. Specifically, we want to know if ב provides a better prediction than א. In this case, we would build a distribution of psuedoR^2 values for each predictor, then use a one sided signed-rank test to see if ב's distribution has a larger median than that of א. Alternatively, we could assess whether the two predictions were substantively different using a two-sided test.
- Finally, say we want to know if adding a new predictor, ב, into a model improves predictive power over a model with only one predictor, א. In this case, we would again compare the distributions for two models, this time testing to see if the median for models including both א and ב is greater than the median for models fit using א alone.

5.5 Using generative models for remapping

Just as we can use the generative GLM to evaluate how well a model predicts held-out data, we can evaluate how well a GLM trained on one condition predicts another. Because of this, generative models lend themselves very naturally to assessing if the behavior of a cell changes under different conditions, for example, if a cell "remaps" in a new environment.

5.5.1 Via Cross-Validation

The same cross-validation approach discussed in the previous section lends itself nicely to comparing between conditions. Each fold in the data produces a model fit to that fold's training data. In addition to cross-validating within condition, we measure the goodness-of-fit for this model on data collected in the

alternate condition. We can then use a sign-rank test to assess whether these paired observations are drawn from a distribution with the same median, that is, whether or not the distribution of models has similar performance in both conditions. Finding that these two distributions are significantly different implies that the cell had a different response between conditions.

5.5.2 Via the time rescaling theorem

So far, we have primarily used $\mathrm{psuedoR}^2$ to assess model goodness-of-fit. (Brown et al., 2002) offer an alternative method for assessing the goodness of fit of a Poisson model using the time rescaling theorem. A key advantage of this method is that it lends itself to a convenient statistical test. Here, I will adapt this technique to assess whether a cell's response changes across conditions.

Consider a uniform Poisson process with unit rate, that is, one with a constant rate $\lambda(t) = 1$. In this case, the amount of time between any two events can be shown (see (Brown et al., 2002)) to be exponentially distributed. The time rescaling theorem takes any Poisson model conditioned on a set of predictors, $\lambda(t) = \lambda(t|X, \beta)$, and transforms it so that the time between events is an exponential distribution. Intuitively, this can be thought of as rescaling the timing of this processes so that every rescaled time-bin now has equal probability of containing a spike (see figure X). Let's say $0 < u_1 < u_2 < \cdots < u_K < T$ describe the events observed from an inhomogeneous process. Brown et al. transform this inhomogenous Poisson process $\Lambda(u_k)$ into a homogenous Poisson process with unit rate:

$$\Lambda(u_k) = \int_0^{u_k} \lambda(u|X_u, \beta)du$$

The time interval between points from this process, $\tau_k = \Lambda(u_k) - \Lambda(u_{k-1})$, will now be drawn from an exponential distribution. For convenience, this can be transformed into a uniform distribution:

$$z_k = 1 - \exp(\tau_k)$$

The advantage of this uniform distribution being that its cumulative density function is a unity line, lending it to an intuitive graphical interpretation (see Figure 5-9).

What the time rescaling theorem gives us, then, is the expected distribution of intervals between events in a Poisson process. For a neuron, this translates into an expected distribution of inter-spike-intervals were that neuron to be perfectly described by our Poisson model. Model goodness-of-fit can therefore be tested by asking how closely the distribution of observed spike times match the distribution

of a proposed model. Brown et al. recommend using a non-parametric statistical test, such as a KS test, to determine a p-value for the hypothesis that the inter-spike-interval distribution for a given spike train matches what is expected by a model.

We extend the time rescaling theorem to remapping by asking a slightly different question. Just as we can ask whether the model-transformed spikes match a uniform distribution, we use a KS-test to ask whether model-transformed spikes from two different experimental conditions have the same distribution. We are no longer asking whether a model produces a good prediction of spiking, but rather whether the model produces equivalent predictions of spiking across environments. A general procedure for this test goes as follows: (1) fit a model on training data from condition 1 (typically ½ of data), (2) use this model to develop distributions ()(z_k) from the remaining data in condition 1 and data from condition 2 and (3) use a KS test to see if these distributions are the same. Figure 5-10 shows examples of theses distributions and the resulting KS test using the simulated cells in Figure 5-5.

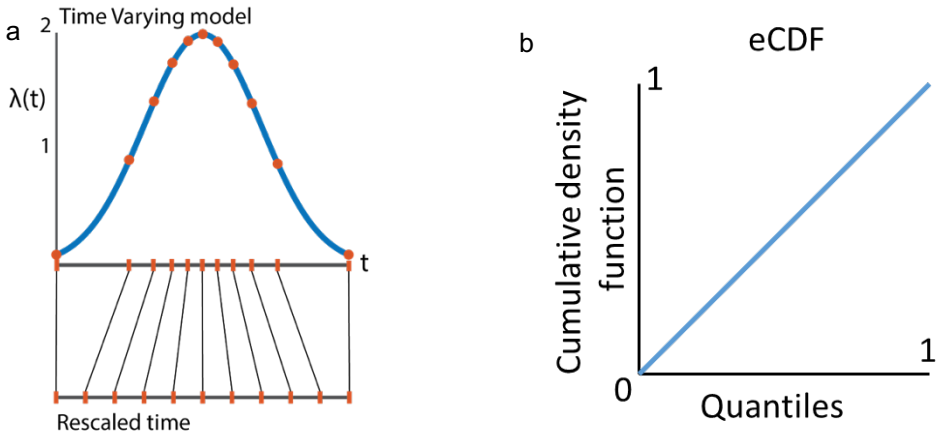

Figure 5-9 Illustration of the time rescaling theorem.

(a) Schematic showing rescaled time as creating equivalent chunks of $\int \lambda(t)$

(b) Expected cumulative density function for transformed inter-spike-intervals for an ideal Poisson

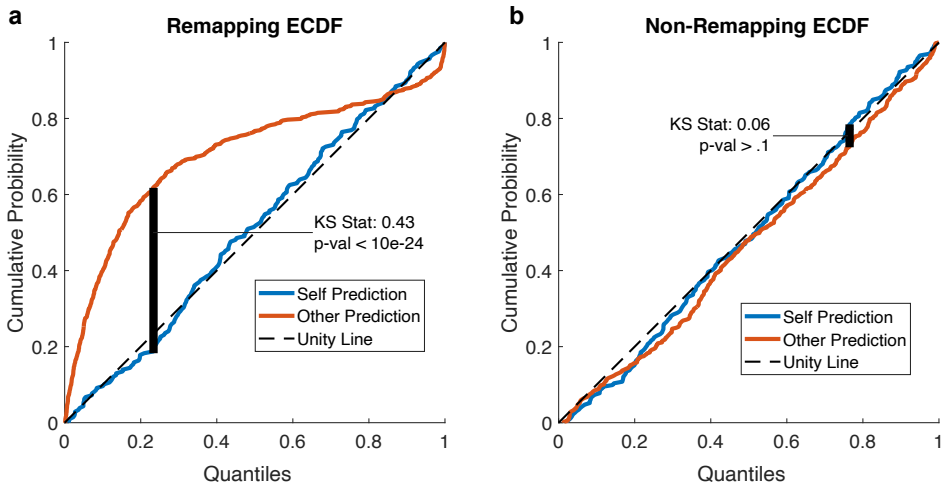

Figure 5-10 Time Rescaling examples

Example Cumulative Density functions of transformed inter-spike-intervals for the data in Figure 1-4 showing application of the time rescaling theorem. Blue represents the prediction within condition (Half of data from condition 1 predicting the other half data of from condition 1), and orange represents prediction across condition (half data from condition 1 predicting condition 2). The KS-statistic (maximum distance between density functions) and p-value are indicated.

(a) Remapping cell, taken from example 1-4a.

(b) Non-remapping cell, taken from example 1-4b.

5.6 Appendix

5.6.1 A loose derivation of Skaggs information

Skaggs notes that mutual information is symmetrical, thus $I(X, Y = 1) = I(Y = 1, X)$. This can be shown to have the following alternate form:

$$I(X, Y = 1) = \sum_{i,j} p(y_j = 1|x_i) \log_2 \left(\frac{p(y_j = 1|x_i)}{p(y_j = 1)} \right) p(x_i)$$

Expressing the equation for information this way is convenient because $p(y_j = 1|x_i)$ is the probability of the neuron firing at any given location in space. This can easily be converted to the firing rate at that particular location in space, λ_i, by scaling this probability relative to the considered time interval, Δt:

$$\lambda_i = \frac{p(y_j = 1|x_i)}{\Delta t}$$

Similarly, the average firing rate for the neuron can be found:

$$\bar{\lambda} = \frac{p(y_j = 1)}{\Delta t} = \sum_i \lambda_i p(x_i)$$

Substituting these into the discrete formulation of $I(X, Y = 1)$ and taking the limit of very small timestamps yields the equation for Skaggs information (the discrete approximation is often more practical for real data):

$$I(X, Y = 1) = \int_x \lambda(x) \log_2 \left(\frac{\lambda(x)}{\bar{\lambda}} \right) p(x) dx$$

5.6.2 Poisson Regression via maximum likelihood

A Poisson generalized linear model takes the form:

$$\ln(\lambda) = X\beta$$

Here, $\lambda = \exp(X\beta)$ is the expected number of events, that is the mean number of spikes, for a distribution taken over specific values of X. The probability of observing y events during any given sample from this distribution is:

$$p(y|X, \beta) = \frac{\lambda^y e^{-\lambda}}{y!} = \frac{e^{y(X\beta)} e^{-e^{X\beta}}}{y!}$$

In a typical experiment, we know a set of observations of spike counts at multiple instances in time $\mathrm{Y} = \{\mathrm{y}_1, \mathrm{y}_2, \dots, \mathrm{y}_\mathrm{T}\}$, each of which is paired with a specific observation of $\mathrm{X} = \{\mathrm{x}_1, \mathrm{x}_2, \dots, \mathrm{x}_\mathrm{T}\}$. By multiplying the probability of each of these observations together under the Poisson model, we can determine the overall probability of these data being observed under this model:

$$\mathrm{p}(\mathrm{y}_1, \mathrm{y}_2, \dots, \mathrm{y}_\mathrm{T} | \mathrm{x}_1, \mathrm{x}_2, \dots, \mathrm{x}_\mathrm{T}, \beta) = \prod_{\mathrm{t}=1}^{\mathrm{T}} \frac{\mathrm{e}^{\mathrm{y_t}(\mathrm{x_t}\beta)} \mathrm{e}^{-\mathrm{e}^{\mathrm{x_t}\beta}}}{\mathrm{y_t}!}$$

This overall probability is known as the likelihood, $\mathrm{L} = \mathrm{L}(\mathrm{Y}|\mathrm{X}, \beta) = \mathrm{L}(\beta|\mathrm{Y}, \mathrm{X})$, of the model for the given data.

Because we know Y and X from our experiment, our remaining task is to choose values of β that best describe these data. Mathematically, this equates to selecting the value of β that maximizes the likelihood of our data being observed. It is generally considered easier to optimize the log likelihood function rather than the likelihood function itself:

$$\ln(\mathrm{L}) = \mathcal{L} = \ln\left(\prod_{\mathrm{t}=1}^{\mathrm{T}} \frac{\mathrm{e}^{\mathrm{y_t}(\mathrm{x_t}\beta)} \mathrm{e}^{-\mathrm{e}^{\mathrm{x_t}\beta}}}{\mathrm{y_t}!}\right) = \sum_{\mathrm{t}=1}^{\mathrm{T}} \left(\mathrm{y_t}(\mathrm{x_t}\beta) - \mathrm{e}^{-\mathrm{e}^{\mathrm{x_t}\beta}} + \ln(\mathrm{y_t}!)\right)$$

Because we are only concerned with optimizing over beta, this formulation has the distinct advantage that we can now ignore the intimidating factorial term in our optimization problem. We now find ourselves maximizing:

$$\mathcal{L} = \sum_{\mathrm{t}=1}^{\mathrm{T}} \left(\mathrm{y_t}(\mathrm{x_t}\beta) - \mathrm{e}^{-\mathrm{e}^{\mathrm{x_t}\beta}}\right)$$

For Poisson and other exponential family distributions, this optimization problem is guaranteed to be convex. This includes the normal distribution, which yields an alternate derivation of the ordinary least squares method. Tradition dictates that one seeks the minimum of the negative log likelihood function, $-\mathcal{L}$, which can be done using any of a myriad of existing gradient descent code. For the entirety of this thesis I borrow the MATLAB optimization method used by (Park et al., 2014).

For a more complete discussion of likelihood approaches to modeling Poisson data, please refer to (Cameron and Trivedi, 2013).

6 A Semi-Supervised approach to Spike Sorting

Despite almost a century of extracellular neural recordings, distinguishing signals from individual neurons remains one of the largest technical problems in this microcosm of neuroscience. Though there have been numerous attempts to automate this process through application of modern clustering methods, these methods often suffer idiosyncratic failures. As a result, the “gold standard” for spike sorting remains manual curation of data by an expert user. Such manual curation, however, is extremely time consuming. To make matters worse, the time required scales linearly with the number of recording channels involved so that - for large datasets like the one discussed in this thesis - it becomes virtually impossible for a single person to sort all of the recorded data. In this chapter, I present a novel “hybrid” approach to spike sorting. Rather than relying on the assumptions of a fully automated method, I adopt a “semi-supervised” approach whereby a user seeds a clustering algorithm, but the time-consuming parts of spike sorting (e.g., forming careful divisions between clusters) are filled in by the algorithm based on the initial seeds. This method has two large advantages. First, it can be considerably faster than fully manual curation; a necessity as high channel count recording systems become increasingly ubiquitous. Second, it increases reproducibility by ceding the cluster-boundary decisions to a density-based algorithm that is robust to small changes in input. In this vein, I will conclude by suggesting the use of “fuzzy logic” in cluster identity as a method to control uncertainty in cluster identity during analysis.

6.1 Why focus on spike sorting?

There are two commonly used approaches towards addressing the problem of neuron isolation. The first is to modify the design of probes to provide better isolation of the neurons being recorded. The classic approach to this is to use a four contact “tetrode,” where each probe consists of 4 closely spaced contacts. Here, the waveform from a given neuron is simultaneously observed on multiple contacts, allowing differences in timing, shape, and spatial location to be used in distinguishing neurons from one another (Gray et al., 1995). More modern probes, most notably the Neuropixels probe, expand on this idea by having hundreds of tightly spaced contacts such that most neurons are picked up on several contacts at once (Steinmetz et al., 2021). Such techniques improve spike sorting by increasing the

available information about the neurons being identified, and this increased information lends itself towards effective waveform discrimination using automated methods (Pachitariu et al., 2016).

Multi-contact probes are not (yet) the norm for collecting data from monkeys or humans. The rationale for this seems, in part, to be due to concerns about the larger spacing and size of monkey neurons relative to those of rodents (Anderson et al., 2006), and fear that this change may limit the utility of multi-contact probes. Difficulties with manufacture and insertion of longer multi-contact probes needed for primate recordings likely also play a role in their limited use. In the rare case where a study has done systematic tetrode recordings from the monkey hippocampus, cell isolation is noticeably improved (Skaggs et al., 2007b). What is more, rodent studies comparing single wire to multi-contact probes find the single wire probes noticeably insufficient for effective cell isolation (Chung et al., 2017b; Magland and Barnett, 2015). Such difference in recordings may explain qualitative differences in firing rate and putative cell types observed in monkey and rodent studies.

The second approach towards neuron isolation is through post-hoc processing. This is particularly important as the scale of neural recordings increases. Historically, spike sorting could take place online using manual curation methods, but this ceases to be a feasible solution when a recording consists of more than a few electrodes. Instead, spike sorting takes place offline using waveforms extracted from the broadband signal. These waveforms are projected into a sorting space (e.g., principal component space), with the idea that neurons from any given unit will form a cluster within this space. Thus, the challenge becomes clustering these data.

For well isolated neurons, this clustering step is relatively straightforward. As the separation between clusters increases, it becomes easy for relatively simple clustering algorithms to identify groups in the data. In fact, the key advantage of multi-contact probes is that they expedite this clustering step by providing multi-faceted information about any given neuron. For poorly isolated neurons, however, there is often insufficient distance between units in sorting space for an algorithm to differentiate between clusters. This is a particular challenge with our chronic drives, where precise electrode positioning is not always possible.

6.2 Data processing and alignment

Before we begin our discussion, it is worth making a few brief comments about spike extraction. Post-hoc spike sorting methods typically consider an extracted short “clip” of time that may have come from a neuron, then use either manual or automated curation. Many clustering methods depend on dimensionality reduction techniques (PCA, etc.) to differentiate shapes of waveforms unique to each cluster, and it is therefore extremely important that waveforms share a common alignment. In our case, we aligned to the trough of each waveform using the following procedure:

- We high pass filtered our data at 450 Hz to isolate activity that might have originated from a single neuron.
- We then thresholded this to identify where the filtered signal fell below 3 standard deviations in voltage from 0 (standard deviation provides a threshold that scales for each channel). From each threshold crossing, we identified the next local trough in voltage and extracted the snippets of filtered signal around this trough as our waveform.
- To assure that all waveforms were properly aligned to their trough, we next up-sampled each waveform by a factor of 4 using optimal sinc interpolation, such that no information was lost in this up-sampling step (Oppenheim and Shafer, 2009). We re-aligned all up-sampled signals to their minimum, then down-sampled back to the original sampling rate.

These aligned waveforms were ultimately used for all sorting.

6.3 Label Propagation and Spreading

For our semi-supervised clustering I used the Label Propagation and Label Spreading algorithms. Semi-supervised learning represents a broad class of algorithms many of which may be useful here, but we have chosen Label Propagation and Spreading because they are already implemented in the SciKit-Learn python package (Pedregosa et al., 2011). As this is not a machine learning thesis, I will limit my description of these algorithms to their practical application for spike sorting. For a more technical description of these methods, please refer to (Zhou et al., 2004).

Before diving into label spreading, it is helpful to consider the failure modes of common clustering algorithms. Two such examples are shown in Figure 6-1, where these algorithms (columns) are applied to standard SK-Learn example datasets. First, consider “Agglomerative” clustering or “Ward” clustering. This

is a hierarchical clustering method that considers the distance between points or clusters in Euclidean space. This algorithm does extremely well for ball shaped clusters, but fails when part of one cluster is closer to a different “true” cluster than to part of itself. Next, consider “gaussian mixture model” clustering, which attempts to fit N gaussian models over the data. Again, this algorithm does very well on ball shaped clusters. However, this is because the assumption of this model is that clusters themselves are gaussian in shape and density. The algorithm fails on non-gaussian shaped clusters. These failures make these algorithms relatively unpredictable when clustering spiking data, where gaussian shaped clusters are not guaranteed (in fact, bursting neurons are known to have non-gaussian clusters). In spite of this fact, Klustakwik, a popular automated method, relies heavily on fitting Gaussian Mixture Models (Rossant et al., 2016).

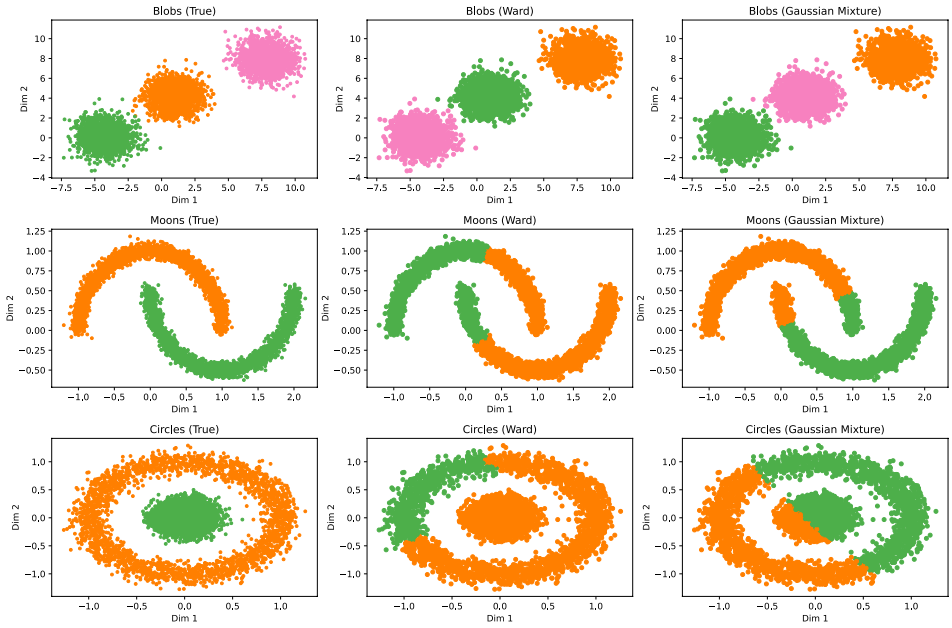

Figure *6-1* Typical clustering algorithms fail on “oddly” shaped clusters

(First Column) Three examples of cluster shapes (defaults in the SK-Learn package).

(2nd Column) “Ward” Clustering, a type of agglomerative hierarchical clustering, struggles with unconventional cluster shapes.

(3rd Column) Gaussian mixture models cannot capture non-round (gaussian) clusters.

Label propagation works very differently from either of these algorithms. Rather than being wholly unsupervised, this algorithm requires that at least some of the data be tagged with an initial cluster identity. The algorithm then follows the nearest-neighbor density or point density function of the dataset to identify which points best match the seed categories. By iteratively including neighboring points into the cluster, this algorithm is able to follow clusters with highly irregular shapes. These labels are used to construct a labeling model, and new or existing data can be labeled by querying nearest neighbor points. See Figure 6-2: the right column shows initial seed points, the left column shows the results of the label propagation algorithm. Label spreading works similarly to Label propagation, but includes a regularization term that limits the impact of noise (Figure 6-2 uses Label Propagation because these synthetic datasets have low noise. The remainder of this chapter will use label spreading).

In addition to irregular clusters, the other common failure mode of many automated algorithms is in failing to separate clusters with some overlap. An example of this is shown in Figure 6-3. In cases where algorithms estimate the number of clusters, this can often mean completely missing the divide between clusters (see (Chung et al., 2017b)). In the case where algorithms are given the number of clusters, they can fail to identify the correct split. In Figure 6-3, Ward clustering incorrectly draws a boundary between the overlapping clusters, while the Gaussian Mixture model divides the clusters in approximately the correct location.

This same toy example gives us an opportunity to examine how label spreading performs on overlapping data. In Figure 6-3, I have selectively provided seed labeling for an orange and a green cluster. Here, the label spreading algorithm fills out the divide between clusters with a good approximation. In this case, the semi-supervised algorithm arguably outperforms the fully automated algorithms. Further, this approach allows the user to explicitly specify which points to include in given clusters, while automating the decision process for ambiguous points.

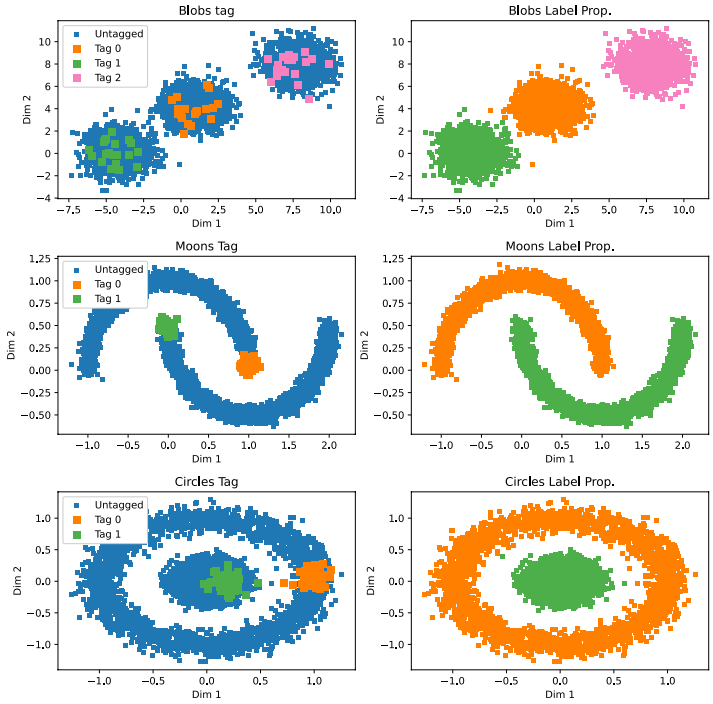

Figure *6-2* Label Spreading on unconventional cluster shapes

(First column) Same data as in Figure 6-1, with a subset of data marked.

(Second column) Result of label spreading algorithm starting with the tags in the first column.

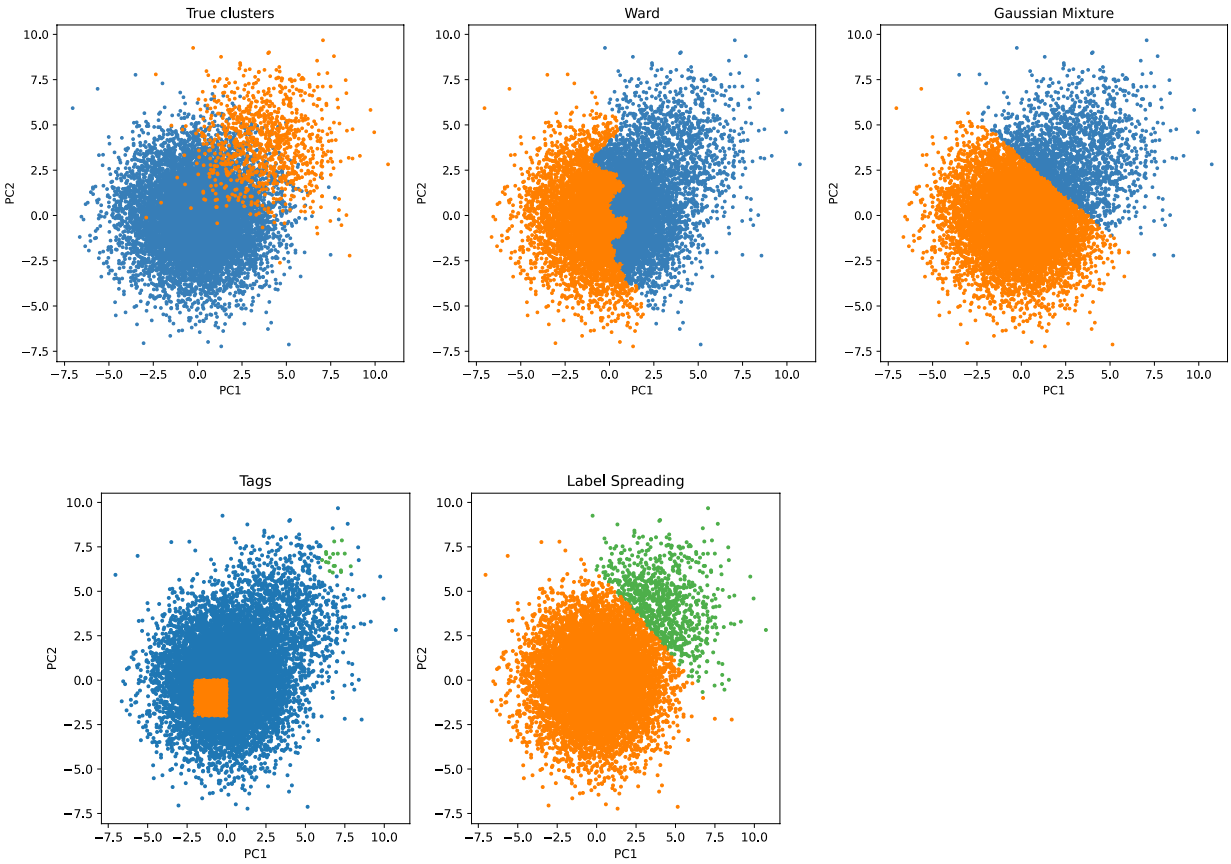

Figure 6-3 Algorithm performance when clusters overlap

(Top, Left) Example data drawn from two overlapping gaussians. This example is not atypical of the overlap between “noise” (blue) and a “cell” (orange).

(Center) Hierarchical clustering fails to properly identify the “cell.”

(Right) Gaussian mixture model can identify the “cell”. Note that this is to be expected, because these data were generated using a gaussian mixture.

(Bottom, Left) Tags for two clusters (orange, green) for label spreading. Blue dots are untagged.

(Right) Result of Label Spreading algorithm using tags.

6.4 Semi-Supervised spike sorting

We applied label-spreading to partially labeled data from our Gray Drive recording systems. A manual user used Plexon's Offline Sorter application to tag? (a) the noise cloud for a given channel and (b) the center of any units. Although the user typically sorted in either principal component or a template view, label spreading was run using the entire 33 sample waveform. Figure 6-4 shows an example of the label spreading algorithm applied to one channel of data. For comparison, the result of manual sorting using only principal components space is shown. Notably, the label spreading sort took <1 minute of user time, whereas fully partitioning these data can take an experienced user tens of minutes.

Unfortunately, building a density map and running label spreading on the full dataset can be computationally expensive for channels with large numbers of waveforms. The label spreading algorithm, however, sorts based on the kernel density map for the space at hand. This means densely packed points are likely to be all from the same cluster, where decisions need to be made about less densely packed points. To illustrate this, Figure 6-5 shows the class boundaries generated by the label spreading model. Because of this, we increase computational efficiency by randomly discarding a subset of the most densely packed points in a given sorting dataset. So long as the density in the area of the discarded points remains higher than the density around the boundary decisions, excluding these data while building the label spreading model will not impact the model itself. For our data, we find that limiting the model fit to 50,000 waveforms provides a considerable increase in computational speed without impacting the model output.

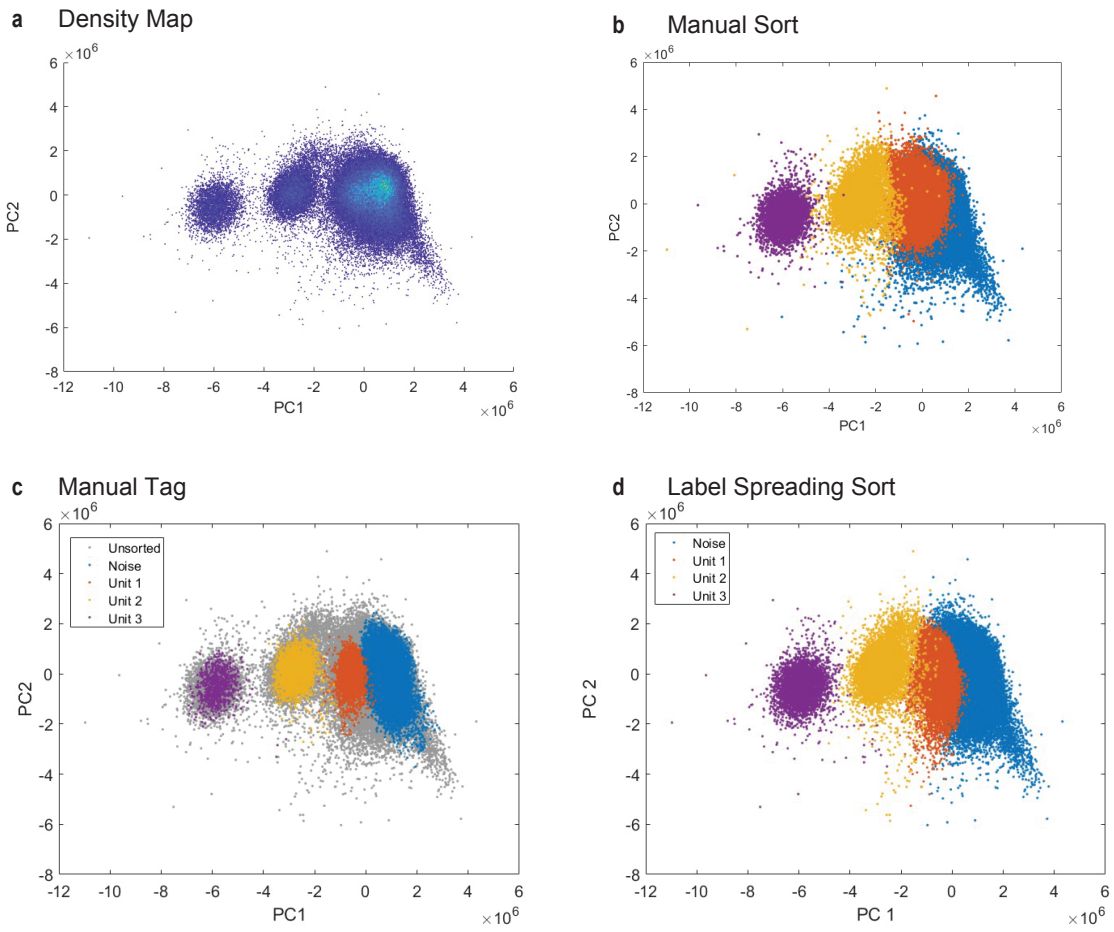

Figure 6-4 Example of Label Spreading for Spike Sorting

(a) Density map for example cell.

(b) Result of manual sorting.

(c) Manual tagging for label spreading, with no specific choices around cluster edges.

(d) Output of semi-supervised sorting.

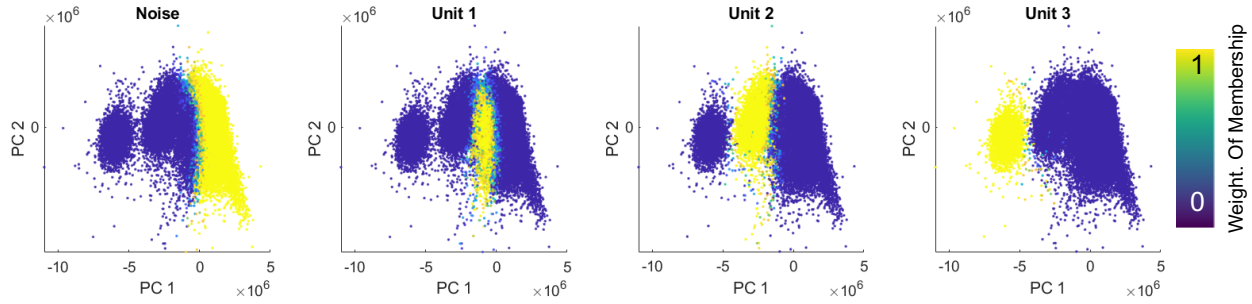

Figure 6-5 Label spreading decision values

Data are the same as in Figure 6-4. For each class, the probability of membership under the label spreading model is shown. Yellow represents points with high certainty of membership, blue represents points below. Points along the decision boundary (green) are shaded according to their membership likelihood.

6.5 Fuzzifiers

Thus far, I have discussed clustering methods, that is, methods that explicitly divide data into discrete categories each of which we can analyze as a single cell. This is the ideal if cells are very well isolated. Sometimes, however, there is no clear divide between any two clusters. In this latter case, we would like a method that allows a certain degree of uncertainty in the division of clusters. “Fuzzy” or “soft” clustering describes a class of methods that allows data points to belong to multiple clusters, with each point assigned a degree of certainty of its membership in any possible cluster.

One common method of fuzzy clustering is C-means (Bezdek, 1981). C-means is an extension of K-means that relaxes the requirement for “hard” cluster identities. Like K-means, the algorithm optimizes the center-of-mass locations for the clusters. In the “fuzzy” case, each point is given a weighted membership to the cluster. This weighting, bounded between 0 and 1, effectively gives a degree of membership to any given cluster. The membership, w, of the i^{th} point to the j^{th} cluster centroid, c, is given:

$$w_{i,j} = \frac{1}{\sum_{k=1}^{c} \left(\frac{\left\| x_i - c_j \right\|}{\left\| x_i - c_k \right\|} \right)^{\frac{2}{m-1}}}$$

This weighting function is referred to as a “Fuzzifier,” because it introduces the fuzziness between cluster identifies. The value of the free parameter, m, determines the degree of overlap of cluster weightings.

Even when manually seeded with the correct number of clusters, fuzzy clustering - specifically the C-means algorithm implementation in MATLAB - never converged to a “good” spike sort in our hands. However, we find the fuzzifier to provide a useful tool for addressing clustering uncertainty. The fuzzifier gives a quantitative way to handle waveforms on the edge between two or more clusters. Rather than needing to make an explicit decision on each data point, weighted membership allows for a given waveform to be systematically included (or excluded) from analysis depending on the question at hand. Concerns about poor sorting, for example, might be addressed by excluding all but the 90% of waveforms most included in a given cluster. Similarly, an LFP-relationship or multi-unit analysis might include all waveforms that have at least a 50% weighting on a given cluster.

Figure 6-6 shows an example of this fuzzifier applied to the data from Figure 6-6, with m=2. Points very near the center of identified clusters are strongly weighted to that cluster, points near cluster boundaries have more diversified weights.

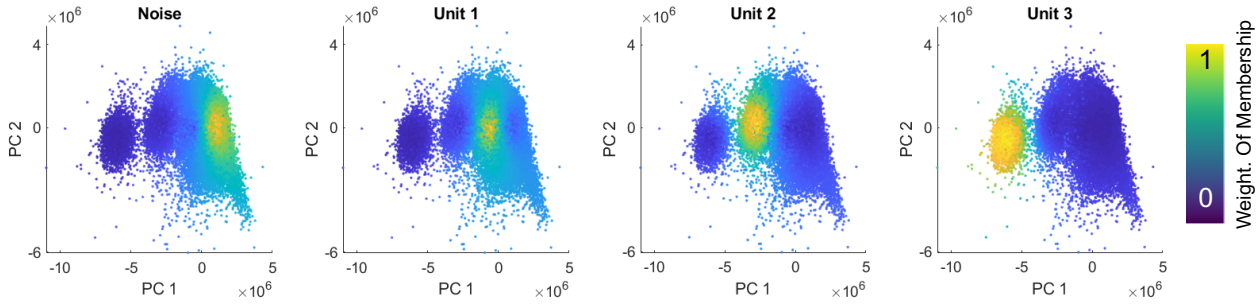

Figure 6-6 Fuzzy classification.

Data are the same as in Figure 6-4. Here, a fuzzifier is used to weight membership in place of model - boundaries, with each point colored by the weight of its membership to the cluster in question

6.6 Code availability

Code for semi-supervised spike sorting, as described here, is available from https://bitbucket.org/yonibrowning/semisupervisedspikesorting/.

6.7 Conclusions

Elsewhere in this thesis, I use both fully manual and fully automated spike sorting to characterize activity recorded from our hippocampal hyperdrives. Both of these methods have pitfalls. The former is onerously time consuming, to the point that recording and sorting data from single monkey must occupy the vast majority of any one researcher's time. The latter, although requiring far fewer person-hours, is extremely error prone. Both of the automated algorithms demonstrated in this chapter were ruled out because they tend to mis-sort data (as shown). The Mountain-Sort algorithm ultimately used in Chapter 2 has the opposite problem; it tends to sort well isolated neurons very well, but it will fail to identify more poorly isolated neurons. In the absence of better recording techniques, neither of these methods provides a complete spike sorting solution.

Here, I demonstrate the viability of a hybrid approach. The semi-supervised method requires manual input for each cell that is sorted, but it is precisely for this reason that it avoids many of the problems associated with typical automated approaches. Importantly, the semi-supervised approach automates decisions about where to draw the boundaries between clusters, effectively eliminating one of the more arbitrary and time-consuming steps of manual sorting. The idea of a "fuzzifier" can be further applied capture the ambiguity of points whose true identity might fall in multiple clusters.

It is worth commenting on the similarity between the label propagation algorithm used here and the class of broadly accepted manifold discovery algorithms (Isomap, T-SNE, UMAP) that are often used in non-linear dimensionality reduction. These algorithms attempt to learn underlying structure in data based on the relationships of data points, much like label propagation uses the K-nearest neighbors density map to determine how user-identified clusters should relate to one another. This means that the label spreading algorithm, effectively, has access to non-linear information that might be lost to researchers "cluster-cutting" only in, for example, a linear principal component space. An alternative way of describing this algorithmic approach to spike sorting is that label spreading operates as a sort of pattern completion, which uses non-linear relationships between points to mimic the pattern seeded by the user.

While my data were collected entirely from single wire electrodes, another important point is that this semi-supervised method is not specific to any one particular data type. Label-spreading can and will work on any feature space, and could easily incorporate data recorded with tetrodes or high-density probes. Adding information with better cell isolation, as with fully automated methods, should only improve the spike sorting.

Finally, an astute reader will have noticed that I have not made use of this novel spike sorting technique in the results section(s) of this thesis. The simple reason for this is that we chose not to impede publication of our results waiting for peer-review and publication of this novel spike sorting approach.

7 Conclusions

In this thesis, I have demonstrated the presence of spatially reliable neural activity in the monkey hippocampus as a monkey performs a structured task in virtual reality. However, I have shown that this ostensibly spatial activity is convolved with other behavioral predictors, providing at best a poor description of what drives hippocampal activity. Specifically, I have shown that hippocampal activity may better correlate to the monkey's progress through its task than to the animal's explicit spatial coordinates. Although the majority of these data are from our virtual Y-maze, I have further demonstrated preliminary evidence of task-related firing in multiple tasks, including presenting preliminary data from a task that explicitly probes hippocampal relationships to abstract and non-physical task spaces.

7.1 How is this study different from all other studies?

At the most basic level, what sets this study apart from this past work is its scale. As I have noted previously, we are not the first to identify spatial firing in monkeys. Yet our chronic recordings have allowed us to record from more cells over more recording sessions using more electrodes and covering a greater anatomic extent of the hippocampus than any recent study of the rhesus monkey hippocampus. Specifically, my estimate (Figure 7-1) says that we have recorded more hippocampal cells than any macaque monkey study in more than 30 years. This technological progress is more than simply a matter of increasing N over past monkey hippocampal work; the scale of our recordings allows us to explicitly query features of hippocampal responses that been historically inaccessible in the monkey. For example, where others have been able to identify spatial responses, we are able to characterize response phenotypes and locations at the level of a pseudo-population. Further, our data give us access to questions about anatomical specificity and longitudinal stability of hippocampal recordings to a degree not seen in past studies. While these differences are important to this thesis, the benefits from our Gray Drive recordings will be even more apparent in surveys of monkey hippocampal oscillations that are currently underway by my collaborators in the Buffalo Lab.

An additional benefit of our chronic implants is that they eliminate some biases that have classically been baked into deep-structure recordings in the monkey. Classical methods favor acute recordings, whereby an electrode is lowered into the monkey's brain for a few hours at a time. Although acute

recordings are extremely effective, they rely on real time unit isolation by the researcher. This means that acute recordings are more likely to find high firing rate cells with large waveforms that can be quickly identified and isolated. In our chronic approach, we would allow electrodes at least a day to settle before recordings, which (we think) is more likely to yield sparsely firing cells. This method is not without its drawbacks. For one thing, it takes considerably longer to sample cells. For another, the inability to adjust electrode position in real time can lead to poorer cell isolation. Despite these caveats, we believe our study to be an important step forward in understanding the activity of hippocampal cells in the monkey.

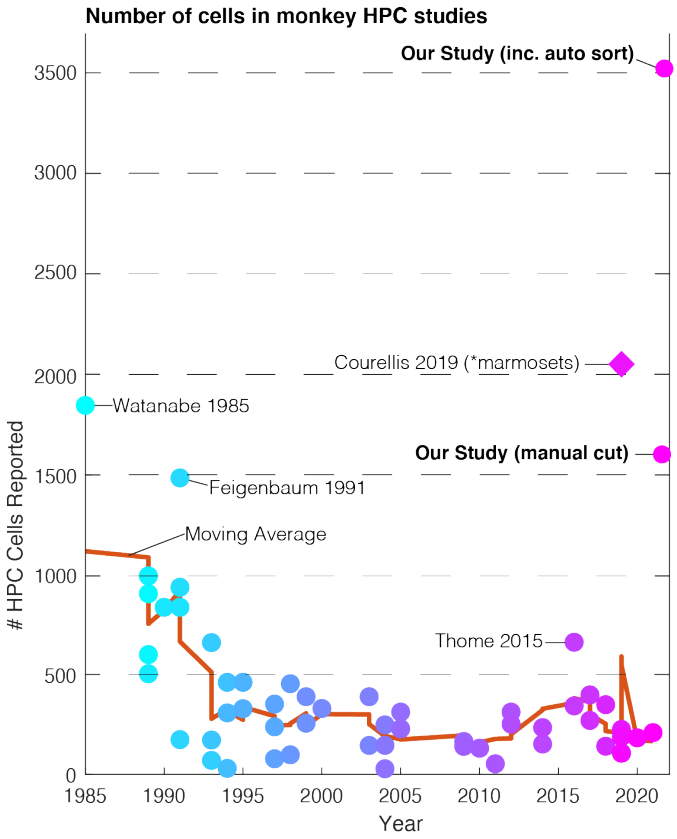

Figure *7-1* Number of hippocampal cells reported in monkey studies.

The number of cells involved in monkey hippocampal studies has been gradually decreasing since 1985. New, chronic recording techniques (including those used in this study and (Courellis et al., 2019)) have greatly increased the number of recorded cells. Note that (Courellis et al., 2019) used marmosets and is therefore not directly comparable to the macaque studies that make up the rest of this graph.

The final major difference between this thesis and other monkey work is its behavioral simplicity. A criticism of past monkey work might be that the possibility of using complex monkey behaviors drives monkey researchers towards complex or multifaceted tasks. These tasks can, in turn, complicate data interpretation and limit the reproducibility of particular behavioral conditions. Despite the technological advances inherent in our virtual reality task and recording technique, our virtual tasks are specifically designed to be as simple as possible. Our Y-maze involves repeated trajectories down a constrained track, and our foraging task involves a limited repertoire of maneuvers the monkey can make. Even our color task, which is specifically designed to allow the monkey to explore an abstract space, does so within a constrained analog of a linear track. Consolidating these behaviors allows us to specifically examine the specificity and reliability of hippocampal responses with a more reproducible behavior than past studies.

These technical and behavioral advances matter because the field is in many ways lacking a set of ground truths underlying neural activity in the monkey hippocampus. A set of accepted observations underpins what we know about the rodent hippocampus; cells have discrete place fields, firing is organized by theta rhythms, etc. Even as there is debate amongst neuroscientists about what drives these phenomena (such as whether place cells are truly a reflection of a "cognitive map (Whittington et al., 2020)"), these observations provide a framework in which to manipulate and study hippocampal neural activity. We lack such basic knowledge about reliable predictors for hippocampal activity in the monkey; in fact, there is increasing evidence that many of the canonical rodent results do not directly translate to the primate. Large survey studies such as this one provide a picture – though by no means a complete one – of what shapes neural activity in the monkey hippocampus.

7.2 What did we learn about the hippocampus?

Hippocampal research is dominated by the hypothesis that the hippocampus stores a cognitive map of an animal's environment, and that this map can be used by the animal to flexibly navigate its environment. This idea stems from the observation that "place cells" in the hippocampus encode an animal's spatial location with discrete, punctate place fields. The presence in the hippocampal-entorhinal system of "grid cells," "speed cells," "head direction cells," "border cells," and other cells with firing that correlates with their titular navigational predictor has been interpreted as further evidence of a hippocampal role in navigation. The prevalence of this navigational hypothesis is driven home by the fact

that John O'Keefe, who initially discovered place cells, and May-Britt and Edvard Moser, who together discovered grid cells, were awarded the 2014 Nobel Prize in Physiology or Medicine for "for their discoveries of cells that constitute a positioning system in the brain."

For place cell responses to be an integral component of hippocampal function, we would expect them to be ubiquitous across species and behavioral conditions. Most studies documenting spatial responses, however, are in rodents. Bats also show punctate spatial firing in crawling bats (Yartsev et al., 2011b) as well as 3 dimensional spatial responses in flying animals (Geva-sagiv et al., 2015; Yartsev and Ulanovsky, 2013). Work in other species has been more mixed. Recording from the hippocampal analogue of birds shows that the degree of spatial responsiveness in a particular task varies by species and likely depends on the ethological relevance of the task at hand (Payne et al., 2021). Further, although spatial responses have been documented in humans (Ekstrom et al., 2003b) and monkeys (as seen here, but also notably in (Baraduc et al., 2019; Courellis et al., 2019; Gulli et al., 2020; Killian et al., 2012; Konig and Buffalo, 2021; Ludvig et al., 2004a; Rolls et al., 1997; Wirth et al., 2017), amongst others, these primate responses are qualitatively different from the Nobel-Prize-winning cells described in rodents. Yes, primates have hippocampal activity that is spatially selective, but these responses lack the discrete nature that is inherent in a rodent "place" cell.

Importantly, the nature of monkey spatial responses reported in this thesis supports the possibility that spatial responses are not a core function of the hippocampus. I do not mean to imply that spatial variables cannot be decoded from the hippocampus, but rather that forming a "cognitive map" does not explain hippocampal activity in the same way that, say, visual activity might describe activity in visual cortex. The apparent difficulty in recording spatial responses has confused the authors of older studies, who struggled to identify spatial responses even as they found that hippocampal neurons could be easily driven by seemingly innocuous events. Vidyasagar et al. (1991), for example, their strongest responses outside their spatial task occurred when the researcher showed their animal a raisin. Even as we report some of the clearest spatial responses seen to date in monkeys, our statistical analysis and behavioral manipulations suggest that this activity is likely better explained by the animal's progress through the virtual task. Further, our data fail to corroborate certain data often taken as evidence that space is an important hippocampal function, such as the anatomical gradient of "place cells" seen in the rodent.

What is more, using task rather than space as a model for hippocampal activity may ultimately prove more parsimonious for interpreting results from rodents and other species. Cells representing “time” or “space” can similarly remap within a given environment if the task the animal is performing changes (Gulli et al., 2020; Pastalkova et al., 2008; Ulanovsky and Moss, 2011). This is particularly pronounced when rodents are performing virtual navigation tasks. Virtual reality studies involving free exploration or foraging have a particularly difficult time eliciting a spatial response. If, however, the rat or mouse performs a constrained task with reliable paths to reward, consistent neural activity emerges (Acharya et al., 2016; Aghajan et al., 2014). Together with our data, it seems likely that it is the task that drives this activity rather than some perception of space.

Why, then, is task so important to the hippocampus? The behavioral tasks used here are very far from natural behaviors an animal might experience outside a laboratory setting; as a result, neurons representing task might well be as artificial as the task itself. Representations of task, then, are really a reflection of the ongoing experience of the monkey. Given that the hippocampus is necessary for the formation of new memories of events and experiences, it is possible that the apparently “task” related activity we see is a reflection of tracking these experiences so that they can be remembered later. While this thesis has shed light on the structure of these representations, further work is needed to see if or how task-responses might contribute to memory.

7.3 Future work

Despite using state of the art methods, there are a number of notable drawbacks to our chronic recording techniques. The drawbacks broadly come down to the fact that our Gray drives offer a particularly slow way to record. The time demands of turning electrodes and isolating units over course of days means that on a good day, 1 in 10 electrodes might have a distinguishable unit. This problem is compounded by the fact that our single wire electrodes have particularly poor unit isolation, meaning that it is often difficult to isolate what neural activity we can detect. The consequence of these drawbacks is that it can take months to build even a basic survey of single unit hippocampal responses in a given task. What is more, this sparse sampling combined with the large footprint of our chronic drives means that it is relatively unlikely that any two simultaneously recorded cells will have firing relationships or even similar

responses, making it difficult to draw inferences about cell interactions or population dynamics from these data.

Fortunately, the solution to this problem may be near fruition. The use of multi-contact probes in the rodent has been increasing as recording devices like the Neuropixels probe become more readily available (Steinmetz et al., 2021), and multiple instantiations of high yield, multi-contact probes capable of reaching deep structures in the monkey are under development. These probes will bring two major benefits. Recording large number of neurons simultaneously will make it possible to:

1. Rapidly iterate through behavioral manipulations, whilst still ensuring sufficient cell counts in each condition and
2. Study the interactions between multiple neurons at scale.

High yield recordings present an exciting opportunity to build on the work in this thesis. Our study is largely correlative, keeping a consistent task structure and observing how this structure shapes hippocampal activity. Future studies will be able to use virtual reality as a framework to train relatively simple tasks, then to iterate through many variations in these tasks to probe which task features play a prominent role in hippocampal representation. Where our sole manipulation was the visual appearance of our virtual environment, higher daily cell yields will allow future studies to manipulate speed, reward amounts, direction, and numerous other variables to identify the specific aspects of a given task that drive hippocampal activity. Such behavioral-perturbation studies will be essential in isolating the specific predictors of hippocampal activity in the monkey and, by extension, establishing the set of “ground-truths” on which to build more complex experiments.

What is more, high density recordings will make it possible to study how populations of cells interact. This is particularly exciting because there is reason to suspect that the hippocampal population behaves differently in rodents and monkeys. In rodents, for example, the hippocampal population of animals placed in a new environment undergoes “global remapping,” in which each cell will change its spatial response in seemingly random manor. Our work and that of others (Baraduc et al., 2019) show that this is not the case in monkeys; instead, some cells maintain their relationship to the task at hand, while others “remap” in response to environmental changes. It remains an open question how this subset of

remapping cells might relate to the stable, task-dependent cells, as well as if there is coordination in the way in which these cells change their firing.

Another exciting target of large-scale recordings is the possibility of understanding what organizes population activity in the monkey hippocampus. Because monkeys lack an obvious theta rhythm, it remains unclear if there exists a theta-like mechanism that coordinates the firing of hippocampal cells. This thesis adds to past work from the Buffalo lab in suggesting that eye-movements provide a modicum of structure to hippocampal activity. Recording many eye-movement modulated cells simultaneously could potentially uncover if and how this single cell phenomena plays out as coordinated activity at the population level. What is more, building an understanding of hippocampal processing on-line will be essential to searching for offline phenomena, such as hippocampal replay, that might show how the hippocampus engages in memory formation.

8 Works Cited

Acharya, L., Aghajan, Z.M., Vuong, C., Moore, J.J., and Mehta, M.R. (2016). Causal Influence of Visual Cues on Hippocampal Directional Selectivity. Cell *164*, 197–207.

Agarwal, G., Stevenson, I.H., Berényi, A., Mizuseki, K., Buzsáki, G., Sommer, F.T., G. Agarwal, I. H. Stevenson, A. Berenyi, K. Mizuseki, et al. (2014). Spatially Distributed Local Fields in the Hippocampus Encode Rat Position. Science (80-.). *344*, 626–630.

Aghajan, Z.M., Acharya, L., Moore, J.J., Cushman, J.D., Vuong, C., and Mehta, M.R. (2014). Impaired spatial selectivity and intact phase precession in two-dimensional virtual reality. Nat. Neurosci. *18*, 121–128.

Aghajan, Z.M., Acharya, L., Moore, J.J., Cushman, J.D., Vuong, C., and Mehta, M.R. (2015). Impaired spatial selectivity and intact phase precession in two-dimensional virtual reality. Nat. Neurosci. *18*, 121–128.

Aljadeff, J., Segev, R., Berry, M.J., and Sharpee, T.O. (2013). Spike triggered covariance in strongly correlated gaussian stimuli. PLoS Comput. Biol. *9*, e1003206.

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

Annese, J., Schenker-Ahmed, N.M., Bartsch, H., Maechler, P., Sheh, C., Thomas, N., Kayano, J., Ghatan, A., Bresler, N., Frosch, M.P., et al. (2014). Postmortem examination of patient H.M.'s brain based on histological sectioning and digital 3D reconstruction. Nat. Commun. *5*, 1–9.

Aronov, D., and Tank, D.W. (2014). Engagement of neural circuits underlying 2D spatial navigation in a rodent virtual reality system. Neuron *84*, 442–456.

Aronov, D., Nevers, R., and Tank, D.W. (2017). Mapping of a non-spatial dimension by the hippocampal-entorhinal circuit. Nature *543*, 719–722.

Bachevalier, J., Parkinson, J.., and Mishkin, M. (1985). Visual recognition in monkeys: effects of separate vs. combined transection of fornix and amygdalofugal pathways. Exp. Brain Res. 532–543.

Baraduc, P., Duhamel, J.R., and Wirth, S. (2019). Schema cells in the macaque hippocampus. Science (80-.). *363*, 635–639.

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

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

Belcher, A.M., Harrington, R.A., Malkova, L., and Mishkin, M. (2006). Effects of hippocampal lesions on the monkey's ability to learn large sets of object-place associations. Hippocampus *16*, 361–367.

Bellmund, J.L.S., Gärdenfors, P., Moser, E.I., and Doeller, C.F. (2018). Navigating cognition: Spatial codes for human thinking. Science (80-.). *362*.

Berg, R.W., Whitmer, D., and Kleinfeld, D. (2006). Exploratory whisking by rat is not phase locked to the hippocampal theta rhythm. J. Neurosci. *26*, 6518–6522.

Bezdek, J.C. (1981). PatternRecognition with Fuzz Objective Function Algorithms.

Bretas, R.V., Matsumoto, J., Nishimaru, H., Takamura, Y., Hori, E., Ono, T., and Nishijo, H. (2019). Neural Representation of Overlapping Path Segments and Reward Acquisitions in the Monkey Hippocampus. Front. Syst. Neurosci. *13*, 1–16.

Brown, E.N., Barbieri, R., Ventura, V., Kass, R.E., and Frank, L.M. (2002). The time-rescaling theorem and its application to neural spike train data analysis. Neural Comput. *14*, 325–346.

Buffalo, E. a., Reber, P.J., and Squire, L.R. (1998). The human perirhinal cortex and recognition memory. Eur. J. Neurosci. *Hippocampu*, 330–339.

Burgess, C.P., and Burgess, N. (2014). Controlling phase noise in oscillatory interference models of grid cell firing. J Neurosci *34*, 6224–6232.

Burgess, N., Barry, C., and O'Keefe, J. (2007). An Ocsilatory Interference Model of Grid Cell Firing. Hippocampus *17*, 801–812.

Butler, W.N., Hardcastle, K., and Giocomo, L.M. (2019). Remembered reward locations restructure entorhinal spatial maps. Science (80-.). *363*, 1447–1452.

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

Buzsáki, G., and Moser, E.I. (2013). Memory, navigation and theta rhythm in the hippocampal-entorhinal system. Nat. Neurosci. *16*, 130–138.

Buzsáki, G., Anastassiou, C.A., and Koch, C. (2012). The origin of extracellular fields and currents-EEG, ECoG, LFP and spikes. Nat. Rev. Neurosci. *13*, 407–420.

Buzsáki, G., Logothetis, N., and Singer, W. (2013). Scaling brain size, keeping timing: Evolutionary preservation of brain rhythms. Neuron *80*, 751–764.

Cahusac, P.M.B., Miyashita, Y., and Rolls, E.T. (1989). Responses of hippocampal formation neurons in the monkey related to delayed spatial response and object-place memory tasks. Behav. Brain Res. *33*, 229–240.

Cahusac, P.M.B., Rolls, E.T., Miyashita, Y., and Niki, H. (1993). Modification of the responses of hippocampal neurons in the monkey during the learning of a conditional spatial response task. Hippocampus *3*, 29–42.

Cameron, A.C., and Trivedi, P.K. (2013). Regression Analysis of Count data (Cambridge).

Cells, P., Cells, G., Moser, M., Rowland, D.C., and Moser, E.I. (2015). Place Cells, Grid Cells, and Memory. 1–16.

Chung, J.E., Magland, J.F., Barnett, A.H., Tolosa, V.M., Tooker, A.C., Lee, K.Y., Shah, K.G., Felix, S.H., Frank, L.M., and Greengard, L.F. (2017a). A Fully Automated Approach to Spike Sorting. Neuron *95*, 1381-1394.e6.

Chung, J.E., Magland, J.F., Barnett, A.H., Tolosa, V.M., Tooker, A.C., Lee, K.Y., Shah, K.G., Felix, S.H., Frank, L.M., and Greengard, L.F. (2017b). A Fully Automated Approach to Spike Sorting. Neuron *95*, 1381-1394.e6.

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

Colgin, L.L. (2016). Rhythms of the hippocampal network. Nat. Rev. Neurosci. *17*, 239–249.

Colgin, L.L. (2020). Five decades of hippocampal place cells and EEG rhythms in behaving rats. J. Neurosci. *40*, 54–60.

Colombo, M., and Gross, C.G. (1994). Responses of inferior temporal cortex and hippocampal neurons during delayed matching to sample in monkeys (Macaca fascicularis). Behav. Neurosci. *108*, 443–455.

Colombo, M., Fernandez, T., Nakamura, K., and Gross, C.G. (1998). Functional differentiation along the anterior-posterior axis of the hippocampus in monkeys. J. Neurophysiol. *80*, 1002–1005.

Constantinescu, A.O., O'Reilly, J.X., and Behrens, T.E.J. (2016). Organizing conceptual knowledge in humans with a gridlike code. Science (80-.). *352*, 1464–1468.

Corkin, S. (2013). Permanent Pressent Tense (Basic Books).

Courellis, H.S., Nummela, S.U., Metke, M., Diehl, G.W., Bussell, R., Cauwenberghs, G., and Miller, C.T. (2019). Spatial encoding in primate hippocampus during free navigation.

Danjo, T., Toyoizumi, T., and Fujisawa, S. (2018). Spatial representations of self and other in the hippocampus. Science (80-.). *359*, 213–218.

Dombeck, D. a, Harvey, C.D., Tian, L., Looger, L.L., and Tank, D.W. (2010). Functional imaging of hippocampal place cells at cellular resolution during virtual navigation. Nat. Neurosci. *13*, 1433–1440.

Dyan, P., and Abbott, L. (2005). Theoretical Neuroscience (MIT Press).

Eichenbaum, H. (2014). Time cells in the hippocampus: a new dimension for mapping memories. Nat. Rev. Neurosci.

Eifuku, S., Nishijo, H., Kita, T., and Ono, T. (1995). Neuronal activity in the primate hippocampal formation during a conditional association task based on the subject's location. J. Neurosci. *15*, 4952–4969.

Ekstrom, a D., Kahana, M.J., Caplan, J.B., Fields, T. a, Isham, E. a, Newman, E.L., and Fried, I. (2003a). Cellular networks underlying human spatial navigation. Nature *425*, 184–188.

Ekstrom, A.D., Kahana, M.J., Caplan, J.B., Fields, T. a, Isham, E. a, Newman, E.L., and Fried, I. (2003b). Cellular networks underlying human spatial navigation. Nature *425*, 184–188.

Eliav, T., Geva-Sagiv, M., Yartsev, M.M., Finkelstein, A., Rubin, A., Las, L., and Ulanovsky, N. (2018). Nonoscillatory Phase Coding and Synchronization in the Bat Hippocampal Formation. Cell *175*, 1119-1130.e15.

Feigenbaum, J.D., and Rolls, E.T. (1991). Allocentric and egocentric spatial information processing in the hippocampal formation of the behaving primate. Psychobiology *19*, 21–40.

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

Forcelli, P.A., Palchik, G., Leath, T., DesJardin, J.T., Gale, K., and Malkova, L. (2014). Memory loss in a nonnavigational spatial task after hippocampal inactivation in monkeys. Proc. Natl. Acad. Sci. U. S. A. *111*, 4315–4320.

Fotowat, H., Lee, C., Jun, J.J., and Maler, L. (2019). Neural activity in a hippocampus-like region of the teleost pallium is associated with active sensing and navigation. Elife *8*, 1–25.

Furuya, Y., Matsumoto, J., Hori, E., Boas, C.V., Tran, A.H., Shimada, Y., Ono, T., and Nishijo, H. (2014). Place-related neuronal activity in the monkey parahippocampal gyrus and hippocampal formation during virtual navigation. Hippocampus *24*, 113–130.

Fuster, J.M., and Uyeda, A.A. (1971). Reactivity of limbic neurons of the monkey to appetitive and aversive signals. Electroencephalogr. Clin. Neurophysiol. *30*, 281–293.

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

Gareth James, Witten Daniela, Hastie, T., and Tibshirani, R. (2014). An Introduction to Statistcal Learning.

Garrett, M., Manavi, S., Roll, K., Ollerenshaw, D.R., Groblewski, P.A., Ponvert, N.D., Kiggins, J.T., Casal, L., Mace, K., Williford, A., et al. (2020). Experience shapes activity dynamics and stimulus coding of VIP inhibitory cells. Elife *9*, 1–25.

Gauthier, J.L., and Tank, D.W. (2018). A Dedicated Population for Reward Coding in the Hippocampus. Neuron *99*, 179-193.e7.

Gavornik, J.P., and Bear, M.F. (2014). Learned spatiotemporal sequence recognition and prediction in primary visual cortex. Nat. Neurosci. *17*, 732–737.

Georges-François, P., Rolls, E.T., and Robertson, R.G. (1999). Spatial view cells in the primate hippocampus: Allocentric view not head direction or eye position or place. Cereb. Cortex *9*, 197–212.

Gergen, J.A., and MacLean, P.D. (1964). the Limbic System: Photic Activation of Limbic Cortical Areas in the Squirrel Monkey. Ann. N. Y. Acad. Sci. *117*, 69–87.

Geva-sagiv, M., Las, L., Yovel, Y., and Ulanovsky, N. (2015). Spatial cognition in bats and rats : from sensory acquisition to multiscale maps and navigation. Nat. Publ. Gr. *16*.

Goyal, A., Miller, J., Qasim, S.E., Watrous, A.J., Zhang, H., Stein, J.M., Inman, C.S., Gross, R.E., Willie, J.T., Lega, B., et al. (2020). Functionally distinct high and low theta oscillations in the human hippocampus. Nat. Commun. *11*, 1–10.

Gray, C.M., Maldonado, P.E., Wilson, M., and McNaughton, B. (1995). Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex. J. Neurosci. Methods *63*, 43–54.

Green, J. d., and Arduini, A. a. (1953). Hippocampal electrical activity in arousal. J. Physiol.

Grion, N., Akrami, A., Zuo, Y., Stella, F., and Diamond, M.E. (2016). Coherence between Rat Sensorimotor System and Hippocampus Is Enhanced during Tactile Discrimination. PLoS Biol. *14*, 1–26.

Gulli, R.A., Duong, L.R., Corrigan, B.W., Doucet, G., Williams, S., Fusi, S., and Martinez-Trujillo, J.C. (2020). Context-dependent representations of objects and space in the primate hippocampus during virtual navigation. Nat. Neurosci. *23*, 103–112.

Hampson, R.E., Pons, T.P., Stanford, T.R., and Deadwyler, S.A. (2004). Categorization in the monkey hippocampus: A possible mechanism for encoding information into memory. Proc. Natl. Acad. Sci. U. S. A. *101*, 3184–3189.

Hardcastle, K., Maheswaranathan, N., Ganguli, S., and Giocomo, L.M. (2017). A Multiplexed, Heterogeneous, and Adaptive Code for Navigation in Medial Entorhinal Cortex. Neuron *94*, 375-387.e7.

Hazama, Y., and Tamura, R. (2019). Effects of self-locomotion on the activity of place cells in the hippocampus of a freely behaving monkey. Neurosci. Lett. *701*, 32–37.

Ho, C.L.A., Zimmermann, R., Flórez Weidinger, J.D., Prsa, M., Schottdorf, M., Merlin, S., Okamoto, T., Ikezoe, K., Pifferi, F., Aujard, F., et al. (2021). Orientation Preference Maps in Microcebus murinus Reveal Size-Invariant Design Principles in Primate Visual Cortex. Curr. Biol. *31*, 733-741.e7.

Hoffman, K.L., Dragan, M.C., Leonard, T.K., Micheli, C., Montefusco-Siegmund, R., and Valiante, T.A. (2013). Saccades during visual exploration align hippocampal 3–8 Hz rhythms in human and non-human primates. Front. Syst. Neurosci. *7*, 1–10.

Hori, E., Tabuchi, E., Matsumura, N., Tamura, R., Eifuku, S., Endo, S., Nishijo, H., and Ono, T. (2003). Representation of place by monkey hippocampal neurons in real and virtual translocation. Hippocampus *13*, 190–196.

Hori, E., Nishio, Y., Kazui, K., Umeno, K., Tabuchi, E., Sasaki, K., Endo, S., Ono, T., and Nishijo, H. (2005a). Place-related neural responses in the monkey hippocampal formation in a virtual space. Hippocampus *15*, 991–996.

Hori, E., Nishio, Y., Kazui, K., Umeno, K., Tabuchi, E., Sasaki, K., Endo, S., Ono, T., and Nishijo, H. (2005b). Place-related neural responses in the monkey hippocampal formation in a virtual space. Hippocampus *15*, 991–996.

Hori, E., Tabuchi, E., Matsumura, N., Ono, T., and Nishijo, H. (2011). Task-dependent and independent synchronous activity of monkey hippocampal neurons in real and virtual translocation. Front. Behav. Neurosci. *5*, 1–19.

Hubel, D.H., and Wiesel, T.N. (1959). Receptive Fields of Single Neurons In Cat's Striate Cortex. J. Physiol. *148*, 574–591.

Jacobs, J., Weidemann, C.T., Miller, J.F., Solway, A., Burke, J.F., Wei, X.-X., Suthana, N., Sperling, M.R., Sharan, A.D., Fried, I., et al. (2013). Direct recordings of grid-like neuronal activity in human spatial navigation. Nat. Neurosci. *16*, 1188–1190.

Jutras, M.J., and Buffalo, E.A. (2010). Recognition memory signals in the macaque hippocampus. Proc. Natl. Acad. Sci. U. S. A. *107*, 401–406.

Jutras, M.J., Fries, P., and Buffalo, E. a. (2009). Gamma-Band Synchronization in the Macaque Hippocampus and Memory Formation. J. Neurosci. *29*, 12521–12531.

Jutras, M.J., Fries, P., and Buffalo, E. a. (2013). Oscillatory activity in the monkey hippocampus during visual exploration and memory formation. Pnas 1–6.

Killian, N.J., Jutras, M.J., and Buffalo, E. a. (2012). A map of visual space in the primate entorhinal cortex. Nature *491*, 761–764.

Killian, N.J., Potter, S.M., and Buffalo, E.A. (2015). Saccade direction encoding in the primate entorhinal cortex during visual exploration. Proc. Natl. Acad. Sci. U. S. A. *112*, 15743–15748.

Kinsky, N.R., Mau, W., Sullivan, D.W., Levy, S.J., Ruesch, E.A., and Hasselmo, M.E. (2020). Trajectory-modulated hippocampal neurons persist throughout memory-guided navigation. Nat. Commun. *11*, 1–14.

Knudsen, E.B., and Wallis, J.D. (2021). Hippocampal neurons construct a map of an abstract value space. Cell *184*, 4640-4650.e10.

Komisaruk, B.R. (1970). Synchrony between limbic system theta activity and rhythmical behavior in rats. J. Comp. Physiol. Psychol. *70*, 482–492.

Konig, S., and Buffalo, E.A. (2021). Eye Movements Temporally Organize Spatial Representations in the Primate Hippocampus.

Kovach, C.K., Tsuchiya, N., Kawasaki, H., Oya, H., Howard, M.A., and Adolphs, R. (2011). Manifestation of ocular-muscle EMG contamination in human intracranial recordings. Neuroimage *54*, 213–233.

Kraus, B.J., Robinson II, R.J., White, J. a, Eichenbaum, H., and Hasselmo, M.E. (2013). Hippocampal ``Time Cells'': Time versus Path Integration. Neuron *78*, 1–12.

Latimer, K.W., Yates, J.L., Meister, M.L.R., Huk, A.C., and Pillow, J.W. (2015). Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. Science (80-.). *349*, 184–187.

Latimer, K.W., Rieke, F., and Pillow, J.W. (2019). Inferring synaptic inputs from spikes with a conductance-based neural encoding model. Elife *8*, 1–28.

Leonard, T.K., Mikkila, J.M., Eskandar, E.N., Gerrard, J.L., Kaping, D., Patel, S.R., Womelsdorf, T., and Hoffman, K.L. (2015). Sharp wave ripples during visual exploration in the primate hippocampus. J. Neurosci. *35*, 14771–14782.

Leutgeb, S., and Leutgeb, J.K. (2014). Remapping to Discriminate Contexts with Hippocampal Population Codes. In Space, Time and Memory in the Hippocampal Formation, D. Derdikman, and J.J. Knierim, eds. (Springer), pp. 227–251.

Logothetis, N.K., Eschenko, O., Murayama, Y., Augath, M., Steudel, T., Evrard, H.C., Besserve, M., and Oeltermann, A. (2012). Hippocampal-cortical interaction during periods of subcortical silence. Nature *491*, 547–553.

Lubenov, E. V., and Siapas, A.G. (2009). Hippocampal theta oscillations are travelling waves. Nature *459*, 534–539.

Ludvig, N., Tang, H.M., Gohil, B.C., and Botero, J.M. (2004a). Detecting location-specific neuronal firing rate increases in the hippocampus of freely-moving monkeys. Brain Res. *1014*, 97–109.

Ludvig, N., Tang, H.M., Gohil, B.C., and Botero, J.M. (2004b). Detecting location-specific neuronal firing rate increases in the hippocampus of freely-moving monkeys. Brain Res. *1014*, 97–109.

MacLean, P.D. (1952). Some psychiatric implications of physiological studies on frontotemporal portion of limbic system (Visceral brain). Electroencephalogr. Clin. Neurophysiol. *4*, 407–418.

MacLean, P.D., Yokota, T., and A., K.M. (1968). Photically Sustained On-Responses of Units in Posterior Hippocampal Gyrus of Awake Monkey. J. Physiol.

Macrides, F., Eichenbaum, H.B., and Forbes, W.B. (1982). Temporal relationship between sniffing and the limbic θ rhythm during odor discriminatin reversal learning. J. Neurosci. *2*, 1705–1717.

Magland, J.F., and Barnett, A.H. (2015). Unimodal clustering using isotonic regression: ISO-SPLIT. 1–29.

Mallory, C.S., Hardcastle, K., Campbell, M.G., Attinger, A., Low, I.I.C., Raymond, J.L., and Giocomo, L.M. (2021). Mouse entorhinal cortex encodes a diverse repertoire of self-motion signals. Nat. Commun. *12*.

Mao, D., Avila, E., Caziot, B., Laurens, J., Dickman, J.D., and Angelaki, D.E. (2020). Spatial Modulation of hippocampal activity in freely moving macaques. BioRxiv.

Mao, D., Avila, E., Caziot, B., Laurens, J., Dickman, J.D., and Angelaki, D.E. (2021). Spatial modulation of hippocampal activity in freely moving macaques. Neuron 1–14.

McCullagh, P., and Nelder, J.A. (1989). Generalized Linear Models (Chapman & Hall\CRC).

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

Meister, M.L.R., and Buffalo, E.A. (2016). Getting directions from the hippocampus: The neural connection between looking and memory. Neurobiol. Learn. Mem. *134*, 135–144.

Meister, M.L.R., and Buffalo, E.A. (2018). Neurons in primate entorhinal cortex represent gaze position in multiple spatial reference frames. J. Neurosci. *38*, 2430–2441.

Meister, M.L.R., Killian, N.J., and Buffalo, E.A. (2013). Allocentric representation in primate entorhinal neurons. In SFN Abstract, p.

Meshulam, L., Gauthier, J.L., Brody, C.D., Tank, D.W., and Bialek, W. (2017). Collective Behavior of Place and Non-place Neurons in the Hippocampal Network. Neuron *96*, 1178-1191.e4.

Minxha, J., Mosher, C., Morrow, J.K., Mamelak, A.N., Adolphs, R., Gothard, K.M., and Rutishauser, U. (2017). Fixations Gate Species-Specific Responses to Free Viewing of Faces in the Human and Macaque Amygdala. Cell Rep. *18*, 878–891.

Miyashita, Y., Rolls, E.T., Cahusac, P.M.B., Niki, H., and Feigenbaum, J.D. (1989). Activity of hippocampal formation neurons in the monkey related to a conditional spatial response task. J. Neurophysiol. *61*, 669–678.

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

Naya, Y., and Suzuki, W.A. (2011). Integrating what and when across the primate medial temporal lobe. Science (80-.). *333*, 773–776.

Naya, Y., Chen, H., Yang, C., Suzuki, W.A., and Squire, L.R. (2017). Contributions of primate prefrontal cortex and medial temporal lobe to temporal-order memory. Proc. Natl. Acad. Sci. U. S. A. *114*, 13555–13560.

Nishijo, H., Ono, T., Eifuku, S., and Tamura, R. (1997a). The relationship between monkey hippocampus place-related neural activity and action in space. Neurosci. Lett. *226*, 57–60.

Nishijo, H., Ono, T., Eifuku, S., and Tamura, R. (1997b). The relationship between monkey hippocampus place-related neural activity and action in space. Neurosci. Lett. *226*, 57–60.

Nowicka, a, and Ringo, J.L. (2000a). Eye Position-Sensitive Units In Hippocampal-Formation And In Inferotemporal Cortex Of The Macaque Monkey. Eur. J. Neurosci. *12*, 751–759.

Nowicka, A., and Ringo, J.L. (2000b). Eye position-sensitive units in hippocampal formation and in inferotemporal cortex of the Macaque monkey. Eur. J. Neurosci. *12*, 751–759.

O’Keefe, J. (1976). Units in the Hippocampus Moving Rat. *109*.

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

O'Keefe, J., and Nadel, L. (1978). The Hippocampus as a Cognitive Map (Oxford University Press).

O'Mara, S.M., Rolls, E.T., Berthoz, A., and Kesner, R.P. (1994). Neurons responding to whole-body motion in the primate hippocampus. J. Neurosci. *14*, 6511–6523.

Omer, D.B., Maimon, S.R., Las, L., and Ulanovsky, N. (2018). Social place-cells in the bat hippocampus. Science (80-.). *359*, 218–224.

Ono, T., Tamura, R., and Nakamura, K. (1991). The hippocampus and space: are there "place neurons" in the monkey hippocampus? Hippocampus *1*, 253–257.

Ono, T., Nakamura, K., Nishijo, H., and Eifuku, S. (1993a). Monkey hippocampal neurons related to spatial and nonspatial functions. J. Neurophysiol. *70*, 1516–1529.

Ono, T., Eifuku, S., Nakamura, K., and Nishijo, H. (1993b). Monkey hippocampal neuron responses related to spatial and non-spatial influence. Neurosci. Lett. *159*, 75–78.

Oppenheim, A. V., and Shafer, R.W. (2009). Discrete-Time Signal Processing.

Opris, I., Santos, L.M., Gerhardt, G.A., Song, D., Berger, T.W., Hampson, R.E., and Deadwyler, S.A. (2015). Distributed encoding of spatial and object categories in primate hippocampal microcircuits. Front. Neurosci. *9*, 1–11.

Pachitariu, M., Steinmetz, N., Kadir, S., Carandini, M., and Harris, K. (2016). Fast and accurate spike sorting of high-channel count probes with KiloSort. Adv. Neural Inf. Process. Syst. 4455–4463.

Park, I.M., Meister, M.L.R., Huk, A.C., and Pillow, J.W. (2014). Encoding and decoding in parietal cortex during sensorimotor decision-making. Nat. Neurosci. *17*, 1395–1403.

Pastalkova, E., Itskov, V., Amarasingham, A., and Buzsáki, G. (2008). Internally Generated Cell Assembly Sequences in the Rat Hippocampus. Science *321*, 1322–1328.

Payne, H.L., Lynch, G.F., and Aronov, D. (2021). Neural representations of space in the hippocampus of a food-caching bird. Science (80-.). *373*, 343–348.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine Learning in Python. J. Mach. Learn. Res. *12*.

Petersen, P.C., and Buzsáki, G. (2020). Cooling of Medial Septum Reveals Theta Phase Lag Coordination of Hippocampal Cell Assemblies. Neuron *107*, 731-744.e3.

Pillow, J.W., Paninski, L., Uzzell, V.J., Simoncelli, E.P., and Chichilnisky, E.J. (2005). Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model. J. Neurosci. *25*, 11003–11013.

Pillow, J.W., Shlens, J., Paninski, L., Sher, A., Litke, A.M., Chichilnisky, E.J., and Simoncelli, E.P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. Nature *454*, 995–999.

Poort, J., Khan, A.G., Pachitariu, M., Nemri, A., Orsolic, I., Krupic, J., Bauza, M., Sahani, M., Keller, G.B., Mrsic-Flogel, T.D., et al. (2015). Learning Enhances Sensory and Multiple Non-sensory Representations in Primary Visual Cortex. Neuron *86*, 1478–1490.

Quiroga, R.Q. (2019). Neural representations across species. Science (80-.). *363*, 1388–1389.

Quiroga, R.Q. (2020). Searching for the neural correlates of human intelligence. Curr. Biol. *30*, R335–R338.

Rao, R.P., von Heimendahl, M., Bahr, V., and Brecht, M. (2019). Neuronal Responses to Conspecifics in the Ventral CA1. Cell Rep. *27*, 3460-3472.e3.

Reppas, J.B., Usrey, W.M., and Reid, R.C. (2002). Saccadic eye movements modulate visual responses in the lateral geniculate nucleus. Neuron *35*, 961–974.

Riches, I.P., Wilson, F.A.W., and Brown, M.W. (1991). The effects of visual stimulation and memory on neurons of the hippocampal formation and the neighboring parahippocampal gyrus and inferior temporal cortex of the primate. J. Neurosci. *11*, 1763–1779.

Ringo, J.L., Sobotka, S., Diltz, M.D., and Bunce, C.M. (1994). Eye movements modulate activity in hippocampal, parahippocampal, and inferotemporal neurons. J. Neurophysiol. *71*, 1285–1288.

Rolls, E.T., and O'Mara, S.M. (1995). View-responsive neurons in the primate hippocampus. Hippocampus *5*, 409–424.

Rolls, E.T., and Wirth, S. (2018). Spatial representations in the primate hippocampus, and their functions in memory and navigation. Prog. Neurobiol. *171*, 90–113.

Rolls, E.T., and Xiang, J.Z. (2005). Reward-spatial view representations and learning in the primate hippocampus. J. Neurosci. *25*, 6167–6174.

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

Rolls, E.T., Cahusac, P.M.B., Feigenbaum, J.D., and Miyashita, Y. (1993). Responses of single neurons in the hippocampus of the macaque related to recognition memory. Exp. Brain Res. *93*, 299–306.

Rolls, E.T., Robertson, R.G., and Georges-François, P. (1997). Spatial view cells in the primate hippocampus. Eur. J. Neurosci. *9*, 1789–1794.

Rolls, E.T., Treves, A., Robertson, R.G., Georges-François, P., and Panzeri, S. (1998). Information about spatial view in an ensemble of primate hippocampal cells. J. Neurophysiol. *79*, 1797–1813.

Rossant, C., Kadir, S.N., Goodman, D.F.M., Schulman, J., Hunter, M.L.D., Saleem, A.B., Grosmark, A., Belluscio, M., Denfield, G.H., Ecker, A.S., et al. (2016). Spike sorting for large, dense electrode arrays. Nat. Neurosci. *19*, 634–641.

Rueckemann, J.W., and Buffalo, E.A. (2017). Spatial responses, immediate experience, and memory in the monkey hippocampus. Curr. Opin. Behav. Sci. *17*, 155–160.

Sakon, J.J., and Suzuki, W.A. (2019). A neural signature of pattern separation in the monkey hippocampus. Proc. Natl. Acad. Sci. U. S. A. *116*, 9634–9643.

Sakon, J.J., and Suzuki, W.A. (2021). Neural evidence of recognition-but not recall-for videos in monkey hippocampus. Hippocampus 1–17.

Sakon, J.J., Naya, Y., Wirth, S., and Suzuki, W.A. (2014). Context-dependent incremental timing cells in the primate hippocampus. Proc. Natl. Acad. Sci. U. S. A. *111*, 18351–18356.

Saleem, K., and Logothetis, N. (2012). Atlas of the Rhesus Monkey Brain n Sterotaxic Coordinates.

Schafer, M., and Schiller, D. (2018). Navigating Social Space. Neuron *100*, 476–489.

Schiller, D., Eichenbaum, H., Buffalo, E. a., Davachi, L., Foster, D.J., Leutgeb, S., and Ranganath, C. (2015). Memory and Space: Towards an Understanding of the Cognitive Map. J. Neurosci. *35*, 13904–13911.

Scoville, W.B., and Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. J. Neuropsychiatry Clin. Neurosci. *12*, 103–113.

Shannon, C.E. (1948). A Mathematical Theory of Communication. Bell Syst. Tech. J. *27*, 623–656.

Shannon, C.E. (1949). Communication in the presence of noise. Proc. IEEE.

Shapley, R.M., and Victor, J.D. (1980). The effect of contrast on the non-linear response of the Y cell. J. Physiol. *302*, 535–547.

Skaggs, W.E., and McNaughton, B. (2007). information theoretic approach to deciphering the Hippcampal code. 1030–1038.

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

Skaggs, W.E., McNaughton, B.L., Permenter, M., Archibeque, M., Vogt, J., Amaral, D.G., and Barnes, C.A. (2007a). EEG sharp waves and sparse ensemble unit activity in the macaque hippocampus. J. Neurophysiol. *98*, 898–910.

Skaggs, W.E., McNaughton, B.L., Permenter, M., Archibeque, M., Vogt, J., Amaral, D.G., and Barnes, C.A. (2007b). EEG sharp waves and sparse ensemble unit activity in the macaque hippocampus. J. Neurophysiol. *98*, 898–910.

Sliwa, J., Plante, A., Duhamel, J.R., and Wirth, S. (2016). Independent Neuronal Representation of Facial and Vocal Identity in the Monkey Hippocampus and Inferotemporal Cortex. Cereb. Cortex *26*, 950–966.

Sobotka, S., Nowicka, A., and Ringo, J.L. (1997). Activity linked to externally cued saccades in single units recorded from hippocampal, parahippocampal, and inferotemporal areas of macaques. J. Neurophysiol. *78*, 2156–2163.

Solway, A., Miller, J.F., and Kahana, M.J. (2013). PandaEPL: A library for programming spatial navigation experiments. Behav. Res. Methods *45*, 1293–1312.

Squire, L.R., Wixted, J.T., and Clark, R.E. (2007). Recognition memory and the medial temporal lobe: A new perspective. Nat. Rev. Neurosci. *8*, 872–883.

Stachenfeld, K.L., Botvinick, M.M., and Gershman, S.J. (2017). The hippocampus as a predictive map. Nat. Neurosci. *20*, 1643–1653.

Steinmetz, N.A., Aydin, C., Lebedeva, A., Okun, M., Pachitariu, M., Bauza, M., Beau, M., Bhagat, J., Böhm, C., Broux, M., et al. (2021). Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. Science (80-.). *372*.

Stephenson, N. (2012). Cryptonomicon.

Strange, B.A., Witter, M.P., Lein, E.S., and Moser, E.I. (2014). Functional organization of the hippocampal longitudinal axis. Nat. Rev. Neurosci. *15*, 655–669.

Sun, C., Yang, W., Martin, J., and Tonegawa, S. (2020). Hippocampal neurons represent events as transferable units of experience. Nat. Neurosci. *23*, 651–663.

Tamura, R., Ono, T., Fukuda, M., and Nakamura, K. (1990). Recognition of egocentric and allocentric visual and auditory space by neurons in the hippocampus of monkeys. Neurosci. Lett. *109*, 293–298.

Tamura, R., Ono, T., Fukuda, M., and Nishijo, H. (1991). Role of monkey hippocampus in recognition of food and nonfood. Brain Res. Bull. *27*, 457–461.

Teng, E., Stefanacci, L., Squire, L.R., and Zola, S.M. (2000). Contrasting effects on discrimination learning after hippocampal lesions and conjoint hippocampal-caudate lesions in monkeys. J. Neurosci. *20*, 3853–3863.

Terada, S., Sakurai, Y., Nakahara, H., and Fujisawa, S. (2017). Temporal and Rate Coding for Discrete Event Sequences in the Hippocampus. Neuron *94*, 1248-1262.e4.

Thome, A., Erickson, C.A., Lipa, P., and Barnes, C.A. (2012). Differential effects of experience on tuning properties of macaque MTL neurons in a passive viewing task. Hippocampus *22*, 2000–2011.

Tolman, E.C. (1948). Cognitive maps in rats and men. Psychol. Rev. *55*, 189–208.

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

Tsao, A., Sugar, J., Lu, L., Wang, C., Knierim, J.J., Moser, M.B., and Moser, E.I. (2018). Integrating time from experience in the lateral entorhinal cortex. Nature *561*, 57–62.

Ulanovsky, N., and Moss, C.F. (2007). Hippocampal cellular and network activity in freely moving echolocating bats. Nat. Neurosci. *10*, 224–233.

Ulanovsky, N., and Moss, C.F. (2011). Dynamics of hippocampal spatial representation in echolocating bats. Hippocampus *21*, 150–161.

Vandecasteele, M., Varga, V., Berényi, A., Papp, E., Barthó, P., Venance, L., Freund, T.F., and Buzsáki, G. (2014). Optogenetic activation of septal cholinergic neurons suppresses sharp wave ripples and enhances theta oscillations in the hippocampus. Proc. Natl. Acad. Sci. U. S. A. *111*, 13535–13540.

Vidyasagar, T.R., Salzmann, E., and Creutzfeldt, O.D. (1991). Unit activity in the hippocampus and the parahippocampal temporobasal association cortex related to memory and complex behaviour in the awake monkey. Brain Res. *544*, 269–278.

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

Viskontas, I. V., Quiroga, R.Q., and Fried, I. (2009). Human medial temporal lobe neurons respond preferentially to personally relevant images. Proc. Natl. Acad. Sci. U. S. A. *106*, 21329–21334.

Wallace, D.J., Greenberg, D.S., Sawinski, J., Rulla, S., Notaro, G., and Kerr, J.N.D. (2013). Rats maintain an overhead binocular field at the expense of constant fusion. Nature *498*, 65–69.

Wang, Y., Romani, S., Lustig, B., Leonardo, A., and Pastalkova, E. (2015). Theta sequences are essential for internally generated hippocampal firing fields. Nat. Neurosci. *18*, 282–288.

Watanabe, T., and Niki, H. (1985). Hippocampal unit activity and delayed response in the monkey. Brain

Res. *325*, 241–254.

Weber, A.I., and Pillow, J.W. (2017). Capturing the dynamical repertoire of single neurons with generalized linear models. Neural Comput. *29*, 3260–3289.

Whittington, J.C.R., Muller, T.H., Mark, S., Chen, G., Barry, C., Burgess, N., and Behrens, T.E.J. (2020). The Tolman-Eichenbaum Machine: Unifying Space and Relational Memory through Generalization in the Hippocampal Formation. Cell *183*, 1249-1263.e23.

Wilming, N., Konig, P., and Buffalo, E. a. (2014). Grid Cells reflect the locus of attention, even in the absence of movement. In SFN Abstract, p.

Wilming, N., König, P., König, S., and Buffalo, E.A. (2018). Entorhinal cortex receptive fields are modulated by spatial attention, even without movement. Elife 183327.

Wilson, F.A.W., Riches, I.P., and Brown, M.W. (1990). Hippocampus and medial temporal cortex: Neuronal activity related to behavioural responses during the performance of memory tasks by primates. Behav. Brain Res. *40*, 7–28.

Wirth, S., Yanike, M., Frank, L.M., Smith, A.C., Brown, E.N., and Suzuki, W.A. (2003). Single neurons in the monkey hippocampus and learning of new associations. Science (80-.). *300*, 1578–1581.

Wirth, S., Avsar, E., Chiu, C.C., Sharma, V., Smith, A.C., Brown, E., and Suzuki, W.A. (2009). Trial Outcome and Associative Learning Signals in the Monkey Hippocampus. Neuron *61*, 930–940.

Wirth, S., Baraduc, P., Planté, A., Pinède, S., and Duhamel, J.R. (2017). Gaze-informed, task-situated representation of space in primate hippocampus during virtual navigation. PLoS Biol. *15*, 1–28.

Xiang, J.Z., and Brown, M.W. (1998). Differential neuronal encoding of novelty, familiarity and recency in regions of the anterior temporal lobe. Neuropharmacology *37*, 657–676.

Yanike, M., Wirth, S., and Suzuki, W.A. (2004). Representation of well-learned information in the monkey hippocampus. Neuron *42*, 477–487.

Yanike, M., Wirth, S., Smith, A.C., Brown, E.N., and Suzuki, W.A. (2009). Comparison of associative learning-related signals in the macaque perirhinal cortex and hippocampus. Cereb. Cortex *19*, 1064–1078.

Yartsev, M.M., and Ulanovsky, N. (2013). Representation of three-dimensional space in the hippocampus of flying bats. Science *340*, 367–372.

Yartsev, M.M., Witter, M.P., and Ulanovsky, N. (2011a). Grid Cells without theta oscillations in the entorhinal cortex of bats. Nature *479*.

Yartsev, M.M., Witter, M.P., and Ulanovsky, N. (2011b). Grid cells without theta oscillations in the entorhinal cortex of bats. Nature *479*, 103–107.

Yokota, T., Reeves, A.G., and MacLean, P.D. (1967). Intracellular olfactory response of hippocampal neurons in awake, sitting squirrel monkeys. Science (80-.). *157*, 1072–1074.

Zhang, H., and Jacobs, J. (2015). Traveling Theta Waves in the Human Hippocampus. J. Neurosci. *35*, 12477–12487.

Zhou, D., Bousquet, O., Lai, T.N., Weston, J., and Scholkopf, B. (2004). Learning with local and global consistency. Adv. Neural Inf. Process. Syst.

Zola, S.M., Squire, L.R., Teng, E., Stefanacci, L., Buffalo, E.A., and Clark, R.E. (2000). Impaired recognition memory in monkeys after damage limited to the hippocampal region. J. Neurosci. *20*, 451–463.