

**Which way do I go? Strategic representations in
rat prefrontal cortex on spatial decision tasks**

**A THESIS
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY**

Nathaniel James Powell

**IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF SCIENCE**

A. David Redish, advisor

October, 2014

© Nathaniel James Powell 2014
ALL RIGHTS RESERVED

Acknowledgements

To Dave... I'll be forever in your debt. ¹

If it takes a Village to raise a child, it takes a Department to produce a PhD student, so I owe a great portion of my success to the Graduate Program in Neuroscience (GPN) at the University of Minnesota. The GPN has been a fabulous home over the years, and I'll miss all of you. In particular, as always, I have to thank John Paton and the late Sabina Deressa. Without you guys I never would have been anywhere I was supposed to be, or have been registered to be there. We always say the entire program would collapse without John, because its true.

To all my labmates in the Redish lab, past and present: it's been a priviledge, and you all taught me valuable lessons. Keep up the amazing work. In particular, Chris, Kelsey, Esther, and Ayaka, thanks for literally making my PhD work possible. To AdamS, Anoopum, Matt, and Jadin, thanks for showing me how all of this works. To AdamS, and Andrew, thanks for the endless converstations about bicycles, catapults, nerf guns, and, in case Dave is reading this, Science.

To my classmates, you have all been great from the first meeting before Itasca until this day. I'm sorry I haven't seen you all graduate yet, but I guess I had to go sometime. Of Itasca and my class, (borrowing from Abraham Lincoln):

Near seven years have passed away,
Since there I bid farewell,
To woods and lakes, and scenes of play,
and labmates loved so well.

¹ because I'll probably never finish cutting all my data.

Where many were, but few remain
Of Old familiar things,
But seeing them, to mind again
The lost and absent brings...

To the rest of the students in the program, Thanks for all the laughs and support, and hang on, if you aren't here yet, you'll get there.

To M-P, V, and V, I tried my best guys. It seems to be working out pretty well. And to Angela Hewitt and all the unofficial mentors, many thanks!

Specifically to Natalie, thanks for all the dinners and movies and reminding me to leave my office every so often, even if I just went back later...

To my other friends outside of Grad school both in town and far away, at least I know who to call when I can't take it anymore. I hope you all feel the same.

In particular, Colin, Ariel, and Is. Thanks for always being there.

To CAJIN (and the other letters), thanks for the occasional excuses to leave town for a little while... still hoping to bring you all out to Minnesota some day.

Thanks to Pat and Curt at CMRR for early lessons and encouragement with Matlab.

Paul Schrater, thanks for all the guidance, insights, and endless help with math methods. And the youtube videos. And for introducing me to Girl Talk! (also Thanks to Paul's lab for letting me hang out occasionally)

Matt Chafee, thanks for all the amazing questions that I hope to answer some day.

Jim Ashe, thanks for taking care of us for all those years. You had a challenging term as DGS, but you handled it very well. For that matter, thanks to Ginger as well and good luck to Dave.

Teresa Nick, wish you were still here, but I hope things are wonderful in California.

To Nantucket, thanks for being a lovely island. I doubt this means much to you, but I appreciate it.

This work was financially supported by:
NIH R01-MH080318
3M Science and Technology Fellowship
NIH 5T32 HD 007151 (Center for Cognitive Science Training Grant)
A huge note of thanks to the CCS for all the support and access to the office in
which this thesis was written.

Dedication

To my Parents... I hope I have the fortune and ability to end up half as successful and content as you have shown me how to be, and if I do I'll assume it was all because of you. Love and thanks.

And as always, "To our relatives, the Rodents, with apologies"
-Valentino Braitenberg (1998)

Abstract

The role of the Prefrontal Cortex (PFC) in animal behavior is both complex and subtle. For this dissertation, we reviewed the current state of knowledge about the role of the rat PFC in regard to behavior and decision-making (Chapter 1). We then described the spatial decision-making tasks and electrophysiological recording techniques we used to explore the role of PFC in rats (Chapter 2). We found overlapping populations of PFC neurons that simultaneously encoded multiple relevant task parameters, including some cases in which multiple parameters were encoded by single neurons (Chapter 3). We then described the spatial firing properties of PFC neurons on these tasks and concluded that although these cells did not seem to directly represent space *per se*, there were important differences in both single-cell and population representations that corresponded to the animal's location on spatial tasks (Chapter 4). Finally, using a population decoding approach that took advantage of the spatially coded information in the cells, we identified different representational states in the PFC between which the animals transition during these tasks. These states represented different cognitive strategies the animals were using to solve these tasks, and we demonstrated that the transitions between these states followed the time at which an animal received information that would cause him to change his strategy but preceded the actual change in the animal's behavior, and could not be accounted for solely on the basis of changes to either sensory information or motor output (Chapter 5).

Contents

Acknowledgements	i
Dedication	iv
Abstract	v
List of Figures	viii
1 Introduction	1
1.1 The Prefrontal Cortex	3
1.1.1 Functional Roles of the PFC	4
1.1.2 Anatomy of the PFC	7
1.2 Behavioral tasks in Rodent PFC	11
1.2.1 Spatial Firing Properties of rat PFC	11
1.2.2 Encoding of multiple task parameters in rat PFC	14
1.2.3 Strategic encoding in the rat PFC	18
1.2.4 The Euston-Cowen Hassle	24
2 Tasks and Recording Methods	28
2.1 Experimental Recording Techniques	29
2.2 The Multiple-T LRA task	35
2.3 The delay discounting task	38
2.4 The multi-track task	42

3	Simultaneous representation of multiple task-relevant parameters in PFC	45
3.1	Parameters of the MT-LRA task	46
3.1.1	Behavioral strategy encoding in PFC	46
3.1.2	Navigational decision encoding in PFC	50
3.1.3	Error encoding in PFC	52
3.1.4	Overlapping populations of task representation in PFC . .	52
3.2	Feeder Fire and Reward Responses in PFC	55
3.3	Consistency of cells across days in PFC	58
3.4	Mixed selectivity in rodent PFC	60
4	Spatial Representations in Rodent PFC	65
4.1	The significance of spatial firing in PFC	66
4.2	The importance of spatial aspects of firing in PFC	72
5	Strategic Representations in Rodent PFC	80
5.1	A method for detecting strategy transitions on the MT-LRA task	81
5.2	Strategy changes without behavioral change	87
5.3	Strategic Representations on the DD task	91
5.4	Wrestling with the Euston-Cowen Hassle	97
5.5	Discussion and Conclusion	101

List of Figures

2.1	Surgical Implantation Techniques. A) The first surgical technique used a tilted single bundle drive targeted at the PL at an angle to avoid the midline sagittal sinus. Histology shows tetrode tracks from this technique. B) The second surgical technique used a dual bundle drive targeted bilaterally around the mid-sagittal sinus. Histology shows final tetrode positions from this technique. Figures modified from (Paxinos and Watson, 1996)	32
2.2	The MT-LRA task. (A) The task consists of a central track containing three low-cost T-choices, a high-cost T-choice at the top of the central track, and two return-rails on which reward was delivered under either <i>leftward</i> , <i>rightward</i> , or <i>alternating</i> contingencies. The first reward site on the left side provided banana-flavored food pellets, while the first reward site on the right side provided fruit-flavored food pellets. The second reward site on both sides provided unflavored (white) food pellets. (B) Reward was only delivered if the animal made the correct choice under the active contingency. (C) For analysis, the maze was divided into six sections (start of maze [SoM], navigation sequence [NS], choice point [CP], top rails [Top], feeder sites [Fed], and bottom rails [Bot]). Reproduced with permission from Powell and Redish (2014)	37

2.3	(A) Average percent of correct laps over all switch sessions aligned to start of the session. Blue line indicates chance rate of responding. Because the first lap of the alternating contingency was always rewarded, chance is 66% on the first lap, but 50% after that. Rats started at chance and quickly learned to make correct choices. (B) Percent of correct laps aligned to the switch. On the switch lap, rat percent-correct dropped to the expected level given no knowledge of the oncoming switch, but then rose back up to above-chance levels. Blue line indicates chance level of behavior, green line indicates perseveration level. Reproduced with permission from Powell and Redish (2014)	39
2.4	The DD task. (A) The task consists of a central choice point and left and right reward locations, one of which provides a large reward after a delay while the other provides a small reward immediately (B) For analysis, the maze was divided into six sections (start of maze [SoM], central stem [Stem], choice point [CP], Delay Zone[DZ], feeder sites [FF], and return rails [RR]). C) Animals typically displayed three behavioral phases on this task, Exploitation, Titration, and Alternation, shown here for a typical session. . . .	41
2.5	The Multi-Track Task. A-D Examples of different task and track conditions presented to the animal on any given day. Four examples are shown, a full set of contingent examples includes HP track and small Track, DD task and LRA task, and delay left or delay right for DD or Left, Right, or Alternate strategies for LRA.	44

3.1	(A)	Examples of cells with strong changes in firing rate on laps before and after the switch. Red bars indicate laps before the switch, blue bars laps after the switch. Red and blue circles represent average firing rates over both types of laps. (B) z_{after} (before) and z_{before} (after) comparisons. Bars indicate overall distribution, with red bars indicating cells found to have a significant firing difference by KS test. Black line indicates expected distribution from ID shuffle. Reproduced with permission from Powell and Redish (2014)	49
3.2	(A)	Examples of cells with strong changes in firing rate on laps to the left vs. to the right of the track. Red bars indicate laps run to the right, blue bars laps run to the left. Red and blue circles represent average firing rates over both types of laps. (B) z_{right} (left) and z_{left} (right) comparisons. Bars indicate overall distribution, with red bars indicating cells found to have a significant firing difference by KS test. Black line indicates expected distribution from ID shuffle. Reproduced with permission from Powell and Redish (2014) . . .	51
3.3	(A)	Examples of cells with strong changes in firing rate on correct laps vs. laps on which the animal made an error. Red bars indicate error laps, blue bars correct laps. Red and blue circles represent average firing rates over both types of laps. (B) z_{error} (correct) comparison. Because there were too few error laps, we only calculated z_{Error} (Correct). Bars indicate overall distribution, with red bars indicating cells found to have a significant firing difference by KS test. Black line indicates expected distribution from ID shuffle. Reproduced with permission from Powell and Redish (2014) . . .	53

3.4	Distributions of multiple responses. (A) The actual distribution of cells responding to the three key binary factors and the overlap of cells that represent multiple factors. (B) Expected distribution overlap if the factors combined independently (i.e. $163/330 = 49\%$ before/after switch responding, $118/330 = 36\%$ correct/error responding, $98/330 = 30\%$ left/right responding). Reproduced with permission from Powell and Redish (2014)	55
3.5	Prelimbic cells show reward responses. Three example cells are shown in (A), (B), and (C). For each example, (top) Raster plot of cell firing by lap aligned to feeder fire time for all four feeder locations (Left Feeders 1 and 2, Right Feeders 1 and 2). (bottom) Peri-Event Time Histograms (PETH) aligned to feeder fire for all feeder locations.	57
3.6	Examples cells with matched waveforms across all six days. For each cell we show: (Top Left) comparison of waveforms across all 6 days; (Top Right) PETH plots of the average firing rate for each day at all four feeder locations; (Middle) Spatial tuning curves over the maze for all six days; (Bottom) 7x7 grid of cell firing. The 7x7 grids divide firing rate of the cell into bins dictated by maze location on the x-axis (using the seven maze locations identified in Figure 2.2) and strategic firing on the y-axis (overall firing rate (FR), firing rate on correct laps (C), firing rate on error laps (E), firing rate on laps to the left (L), firing rate on laps to the right (R), firing rate on laps before the switch (B), and firing rate on laps after the switch (A)). Partially reproduced with permission from Powell and Redish (2014)	59

3.7	Histogram of across-day correlations. (A) Histogram of the correlation coefficient of all pairs of the 7x7 grids of FR seen in Figure 3.6 for all pairs of cells not identified as being the same cell recorded across multiple days (Non-Matched). (B) The same histogram as in part (A) but for all pairs identified as being the same cell on consecutive days (Matched). Reproduced with permission from Powell and Redish (2014)	61
3.8	JAGS Mixed Selectivity Model Fit Examples. Examples of the fit parameter distributions for three neurons fit to the linear and nonlinear mixed selective models motified from Rigotti et al. (2013). Cell one is better described by the nonlinear model and displays only plausibly non-zero coefficients for the cross terms. Cell two displays only plausibly non-zero coefficients for the linear terms, and is better fit by the linear model, and cell 25 displays plausibly non-zero coefficients for both linear and nonlinear terms, and is better described by the nonlinear model.	64

4.1	<p>Prelimbic cells show non-uniform spatial firing. (A) Typical spatial tuning curves of prefrontal cells. Note the large firing response fields, typically covering a broad region of the maze and usually present on both left and right sides of the track. Color bars indicate firing rate in Hz. (B) Typical spatial tuning curves of hippocampal cells from the same task. Note the smaller firing response fields, more typical of hippocampal place cells. Color bars indicate firing rate in Hz. (C) Decoding confusion matrix generated from prefrontal cell firing, from 100 random assignments of laps to training and test sets averaged together, then averaged over all rats, all sessions, and left and right laps. Color bar indicates decoded probability at each location. (D) Decoding confusion matrix generated from hippocampal cell firing, from 100 random assignments of laps to training and test sets averaged together, then averaged over all rats, all sessions, and left and right laps. Color bar indicates decoded probability at each location. Reproduced with permission from Powell and Redish (2014)</p>	69
4.2	<p>Spatial firing patterns are conserved across days different sized tracks and tasks. Tuning curves across all 12 days of the multi-track task are shown for three different example cells, divided up by which maze and which task was being run on each day. Firing rate scales are the same across days for each cell, but differ between examples.</p>	71
4.3	<p>Parametric firing depends on maze location. Colors represent the number of cells with significantly different firing rates for each region of the MT track for each behavioral parameter (Before vs. After the switch, Correct vs. Error laps, and laps to the Left vs. Right side of the track. Maze and sections are depicted below for convenience.</p>	75

4.4	Parameters encoded on the DD task depend on maze location. Firing rate differences for each side of the maze are shown for four example cells recorded across multiple days on the DD task for each section of the maze. Z-scored firing rates for each day are represented by different markers. Marker internal color (red vs. black) represents delay side vs. non-delay side firing, marker external color (magenta vs. cyan) represents left vs. right side firing. All firing rates have been z-scored against all laps of that session. Arrows indicate sections of the maze with different consistent representations than the rest of the maze for that cell.	77
5.1	PFC Population Firing Representations. (A) Spatial Firing rate population representations for two sample laps from a single session. Note the differences between cell firing between laps. (B) A sample correlation plot on fake data demonstrating what gradual change in population representation over the course of a session would look like. (C) Another correlation plot on fake data demonstrating what a sudden change in population representation would look like. . .	83

5.2	An example of the generation of the transition probabilities from a single MT-LRA session. (A) The lap-lap correlation of the spatial population firing vectors for one session, with the switch lap marked (black line). (B) The average lap-lap correlation of the spatial population firing vectors over all sessions aligned to the switch lap for each session (C) The animal’s behavior on the same session. Red dots = error laps, black dots = correct laps before switch, blue dots = correct laps after switch. (D) An example k-means fit for the session with three clusters. Transitions between groups of clusters are indicated (black lines). (E) Example of the Change Point Algorithm on the k-means data. Change points are identified as the largest deviations of the slope of the cumulative sum from the average slope. (F) Sample Histogram of all detected k-means change points for 100 iterations with $k = 3$ clusters.	85
5.3	Transition scores for all MT sessions (A) A 21 lap window of averaged transition scores over all MT sessions aligned to the switch lap for each session (error bars SEM) (B) A 21 lap window of transition probabilities aligned to the lap of greatest behavior change for all MT sessions (error bars SEM).	86
5.4	Transition scores for all MT sessions from only single sides of the track (A) A 21 lap window of averaged transition scores aligned to the switch lap for each session (error bars SEM) (B) A 21 lap window of transition probabilities aligned to the lap of greatest behavior change (error bars SEM).	89

5.5	<p>Transition scores for all MT sessions from session start (A) Average transition scores of all sessions aligned to laps 5-30 (error bars SEM) (B) Lap by lap correlation matrix of population firing averaged over all sessions, aligned to start of session. (C) Lap by lap correlation matrix of population firing, averaged over all sessions, aligned to lap of highest transition score in the first 15 laps of each session. (D) Correlation plots for the spatial- population firing vector for all cells recorded consistently across several days for two rats (4 cells across six days for Rat 2 and 5 cells across five days for Rat 3). Magenta lines indicate the beginning and end of individual sessions, and black lines indicate the switch in each session. Arrows indicate areas of heightened self-correlation at the beginning of each session.</p>	92
5.6	<p>Finding Behavioral Changes on the DD task. (A) The track utilized in the DD task, for reference (B) Animals take four distinct paths on the DD task, which will help us answer the Euston-Cowen Hassle. Left to Right (LR), Right to Left (RL), Left to Left (LL), and Right to Right (RR). (C) Animal's behavior on a typical DD session plotted by delay value (for the delayed side) vs. lap number. A clear behavioral transition occurs between laps 41 and 45. (D) Calculating the behavioral change point (indicated by the red star) from the maximum deviation from the average slope on a plot of the cumulative sum of visits to the delayed side of the track. (E) The behavioral change point found in (D) matches the animal's behavioral transition seen in (C).</p>	94

5.7	An example of the generation of the transition probabilities from a single DD session. (A) The lap-lap correlation of the spatial population firing vectors for one session, with the detected behavioral change lap marked (dotted line). (B) The animal's behavior on the same session, delay value vs. time. Red dot indicates Behavioral change lap. (C) An example K-means fit for the session with three clusters. Transitions between groups of clusters are indicated (black lines). (D) Example of the change point algorithm on the K-means data. Change points are identified as the largest deviations of the slope of the cumulative sum from the average slope. F) Sample Histogram of all detected k-means change points for 100 iterations with $k = 3$ clusters.	95
5.8	Transition probabilities for all DD sessions. A 21 lap window of transition probabilities for all DD sessions aligned to the behavior change lap for each session (error bars: SEM)	96
5.9	Population Firing Correlations are not due to trajectory similarity. (A) Lap by lap population firing vector correlations for a sample session, as in Figure 5.7. Note here, black lines are drawn for illustration purposes to illustrate the different states present in the population firing, and reproduced on all plots. (B) The lap by lap trajectory correlation plot for the animal's running on the task. (C) The delay value of the delayed side by lap for this session. Note the black lines mark the same laps as in other panels, and were determined from cellular firing, NOT behavior. (D) The data plotted in (A) broken down by lap type to show correlations only among laps run to and from the same locations on the track. . . .	99

5.10	Transition probabilities for all DD sessions, for laps with the same trajectory type. A 21 lap window of transition probabilities for all DD sessions, calculated for a collection of only LR laps and only RL laps and averaged together, aligned to the behavior change lap for each session (error bars: SEM)	100
------	--	-----

Chapter 1

Introduction

On September 13th, 1848 near the town of Cavendish, Vermont, a railway worker named Phineas Gage suffered a horrific accident. While he was engaged in tamping down explosive for a blasting hole, his tamping rod accidentally sparked an explosion that sent the rod flying back out of the hole and through his skull. The rod, which was 1 1/4 inches in diameter and 3 feet 7 inches long, passed entirely through his skull from the top of his left jaw, behind his left eye, through the front of his cranium, and out through the top of his head, eventually landing 20 to 25 meters away (Macmillan, 2002). Such an accident would not be remarkable except that Gage survived the blast, not even losing consciousness according to his own report, and made his way with some assistance back to his house (reportedly claiming as he left the job site that he would be back at work the next day). He would not make good on that claim, but he did reach his apartment with minimal assistance and was able to climb the stairs up to his room, where he was treated by Dr. John Martyn Harlow (with assistance from Dr. Edward Higginson Williams). Despite extensive damage to his brain, he was awake and lucid throughout the entire ordeal.

Early newspaper reports of the Gage incident did not stress that anything astounding had occurred, because although it was remarkable that he was mentally coherent following severe brain trauma, he was not expected to survive the

expected infection and illness. However, Gage did survive and recover, which made his case truly extraordinary. Even to this day there are few incidents of such traumatic brain injury with positive outcomes (Macmillan, 2002). However, the full significance of the Gage case at the time (and to this day) for the fields of psychology and neuroscience involves the subsequent reports of changes to his personality and behavioral patterns. Reports vary, but many mention a man who had been described as a “smart worker” who was “a great favorite with his men” who following the incident had trouble holding a job for long periods of time. He was often also described as a man with “temperate habits” who “was quiet and respectful” but after the accident became “gross, profane, coarse, and vulgar.” Many of the descriptions of his behavior after the accident refer to him as being childlike, always following his impulsive desires, and having a greater fondness for drinking. In particular, he seemed to fail in social situations, behaving according to some friends “like an idiot”. A. R. Damasio in a 1994 study describes him as having lost “respect for social conventions”, as well his “neural basis for moral sense” (Damasio, 1994). All of these descriptions paint the picture of a man much changed as a result of his accident. However, equally fascinating to consider are the mental faculties he retained. He was not killed outright. He did not lose any of his powers of speech. He did not report significant difficulties with memory. While he seems to have lost some of the skills that made him a valuable worker (particularly his organizational skills), he seems to have suffered no gross motor deficits, nor is there any mention of his having lost particular motor skills he had before the accident. Nor is there any record of his senses being diminished in any significant way other than the loss of his left eye. In comparison to many similarly traumatic brain injuries, the particular blend of effects suffered by Phineas Gage is fascinating.

However, the Gage case is as fascinating as a tale of survival as it is frustrating as an actual scientific case study. So few details are known about the extent of the damage to Gage’s prefrontal cortex that it remains difficult to conclude much about the functions of that structure as a result. It is possible that all the changes

later attributed to Gage were unrelated to the destruction of a significant amount of brain tissue he suffered. It is also possible that subsequent infection and inflammation caused further damage in his brain beyond what could even be assumed as a result of the accident, and that these unknown changes were responsible for any personality changes. Furthermore, there is debate about the extent of the change in Gage's personality because not much was recorded about the man in detail either before or after his incident, and the veracity of the reports that have been made has been called into question (Macmillan, 2002). However, the Gage case contained two strong indications about the role of frontal cortex that would shape many future inquiries into this structure: 1. It is not strictly necessary to sustain life, or even relatively complex human behavior (no one subsequently described Gage as sub-human) 2. Nevertheless, the PFC does seem implicated in some of the more complex aspects of behavior. Prior to his injury Gage was described as a smart worker, capable of being a foreman, supervising other workers, and organizing complex projects. After his accident, these aspects of his personality seemed to change, leaving him "no longer Gage". These basic early conclusions are consistent with what is currently known about the function of the PFC.

1.1 The Prefrontal Cortex

In primates, the prefrontal cortex (PFC) sits on the anterior pole of the frontal cortical lobe, from which it derives the name pre-frontal. Though this name does not fit its location in some other species (of particular concern will be the rodent) the PFC itself is most closely associated with primates and in particular humans, in which it achieves its greatest size relative to the other cortical areas. All cortex located in front of the central sulcus is generally described as "action cortex" (Fuster, 2008) and PFC is no exception. Its function is generally described as managing higher order aspects of behavioral control, commonly called cognitive control. However, the PFC also receives extensive sensory information, which is critical to its role in identifying the current state of the world and selecting the

most appropriate behavioral response to that state. There is a parallel development of the complexity of behaviors displayed among mammalian species and the development of their prefrontal cortex, both of which reach their apex in the human. However, we believe aspects of the cognitive control inherent in this area are present, and may be effectively studied, in simpler species, specifically the rat.

1.1.1 Functional Roles of the PFC

The functional role of the prefrontal cortex is difficult to describe precisely because of the complexity and high level of its operation. Indeed, as with Phineas Gage many people and animals who have suffered prefrontal lesions appear to function normally on a large variety of every day tasks and only begin to demonstrate impairment in complex behaviors and situations. In addition, the functions of PFC tend to involve the interaction of a broad number of other cortical areas, making the attribution of fault throughout the system difficult. Nevertheless, many of the functions can be rather well characterized, and we will attempt to provide an overview of these functions below.

Joaquin Fuster described the PFC as necessary for “any series of purposive actions that deviated from rehearsed automatic routine or instinctual order” (Fuster, 2008, page 3). It is generally agreed that other cortical and subcortical areas are sufficient to allow instinctual actions and even well rehearsed automatic routines, but as action plans become more complicated and require greater planning the PFC becomes necessary. Planning may in fact be a critical concept, as the PFC tends to be extensively involved in situations that require the execution of many actions in order to achieve some goal of the animal, particularly when those actions need to be executed in a particular series (or specifically with a particular delay in time involved). In addition, Prefrontal cortex is commonly associated with the concept of executive attention, which itself consists of three specific aspects, working memory, preparatory set, and inhibitory interference control, all of which are to various degrees inter-connected.(Fuster, 2008)

Working memory refers to active mental representations necessary for some purpose in the immediate future. These representations could be of stimuli or patterns that are being held over a delay period (the classic case of working memory), specific anticipated stimuli or patterns that are expected and would release certain subsequent actions, or even motor action patterns themselves that are ready to be activated. This type of memory is very similar to the psychological concept of short term memory, itself often called working memory, which is thought of as a register of concepts currently under consideration and is classically thought to have a span of 7-9 items. Working memory in this sense, however, is distinct from the concept of short term memory in which items are stored before they are consolidated into long term memory, and may or may not be related to that process. In fact, it is quite likely that working memory representations are commonly extracted from long term memory representations elsewhere in cortex (Fuster and Alexander, 1973; Alexander and Fuster, 1973; Nishino et al., 1984). Working memory is often associated with PFC because various electrophysiological recording studies have found persistent signals in PFC that maintain firing over a delay period in various tasks requiring the preservation of a memory over that period of time (Fuster et al., 1971, 1982; Kojima and Goldman-Rakic, 1982; Funahashi et al., 1989). Additionally, inactivation studies demonstrate that animals with PFC lesions have difficulty with delay tasks despite being able to perform the same matching operations without a delay imposed (Jacobsen, 1935; Jacobsen and Nissen, 1937; Fuster and Bauer, 1974; Bauer and Fuster, 1976). In this way, working memory is related to the two more common aspects of PFC function, timing (memory traces are maintained over time) and goal directed behavior (memory traces are generally maintained in order to achieve some goal, or in some cases the goal itself is maintained over time).

The concept of preparatory set refers to a kind of anticipatory readiness in the motor system to execute some set of actions. An animal might have been trained to repeat some specific action on the appearance of a specific cue, for example making a saccade to the lower left side of the screen, running to the side

of the track above which a light appears, or throwing to a wide receiver running a stop and go route if the deep safety bites on the run fake. Note however that the preparatory set might not include just a single action and stimulus pair but rather a set of several actions to accomplish depending on the specific stimuli witnessed. Preparatory set is rightly thought of as a state of readiness to execute a specific action plan in response to a specific stimulus, and in this way it is similar to a working memory trace for a future action. Preparatory set is also similar to the concept of habitual behavior, and the precise difference between the two is difficult to determine as neither concept is always rigorously defined or used in precisely the same fashion (Gibson, 1941). However, in general a set could be established to prevent a habitual action in a certain case, for example suppressing the general tendency to turn left in response to a blue light in the case that a red light is also illuminated. In this way sets can be examples of more flexible behaviors than habits, but may also be thought of as habits that have not yet been sufficiently learned. The concept is also critical to the concept of set shifting, as in the case of classic tasks such as the Wisconsin card sorting task, in which a given set of rules or outcomes is suddenly changed and the subject must learn to change their behavior as a result. PFC has been shown to be involved in these sorts of set-shifts in many species (Dias et al., 1997; Rushworth et al., 2002; Vanderhasselt et al., 2006) including rats (Dalley et al., 2004).

Inhibitory interference control generally refers to the process of suppressing or inhibiting unnecessary actions or reactions. Inhibitory control is often necessary to suppress some actions that may be immediately tempting in order to be able to execute some longer term goal. For example, foregoing a smaller reward offered immediately in order to achieve a larger reward later on, not flinching from the path of an oncoming linebacker in order to complete a pass downfield to an open receiver (achieving a greater gain than simply “not getting hit), or suppressing the desire to go for a bike ride in favor of completing a chapter of one’s dissertation. Inhibitory interference control is generally thought of in terms of suppressing actions (in many cases actions that could otherwise be termed habitual), but it

could also be manifest as a decreased awareness of or attention to stimuli that are not the necessary cues in a given environment. Inhibitory control is commonly indicated in general lesion studies of PFC that often produce hyperactivity in animals (Jacobsen, 1931; Richter and Hines, 1938) , and also in some direct evidence that inactivation of PFC leads to decreased suppression of responses to nogo cues (Allen, 1940; Sakurai and Sugimoto, 1985a).

It is important to keep in mind that though these general concepts are all related to prefrontal function and are useful in characterizing the importance of the region, they remain rather unspecific. For example, the three described above can all be interrelated in ways that makes distinguishing between them difficult. A preparatory set, or example, may be thought of as a motor working memory. And a key component of a specific preparatory set may be simply to inhibit the performance of an action that would commonly be carried out. Also, successfully holding a particular pattern in working memory may be a case of suppressing (or inhibiting) further patterns from intruding, which underscores why all of these concepts are thought of as components of executive function. While the concepts are useful, in order to form a more mechanistic view of the role of PFC, we may need to turn to more specifics. We will begin with identifying the anatomical connections that define this area, then turn to specific experimental results that concern the topic of this dissertation.

1.1.2 Anatomy of the PFC

The simplest definition of prefrontal cortex is all areas that receive projections from the mediodorsal nucleus of the thalamus (Rose and Woolsey, 1948, 1949; von Monakow, 1895). This definition has the advantages of being precise and applicable to all mammalian brains (rodents included (Uylings and van Eden, 1990; Preuss, 1995)), but it also underscores the difficulty of identifying the primary locus and function of this area, as the projections from this nucleus are somewhat widespread. In primates (including humans), the overall prefrontal area can be

divided into three general regions, the medial, lateral, and orbital aspects of PFC, on both anatomical and functional grounds. Functionally speaking, the orbital region has been implicated in reward responses, motivational effects, and inhibition of internal or external factors that distract from goal oriented behaviors (Brutkowski and Dabrowska, 1963; Iversen and Mishkin, 1970; Klüver and Bucy, 1939). Lesions to this region in monkeys tend to produce decreased aggressiveness (Butter et al., 1970; Kling and Steklis, 1976; Peters and Ploog, 1976). The medial prefrontal cortex is linked to similar factors, perhaps with an increased role in emotional control, as lesions here are consistent with disinhibition of aggression and hunger (Brutkowski, 1965; Butter et al., 1970). Additionally, many lesions produce hyperactivity (Jacobsen, 1931) that may be more accurately be termed a hyper-reactivity as animals respond to all stimuli and show deficits in attention (Klüver, 1933; Konorski and Lawicka, 1964). The lateral prefrontal cortex is commonly associated with cognitive tasks and working memory. Lesions in this region cause animals to show great difficulty with delay tasks along with most behaviors that require execution of multiple actions related by temporally separable events (Jacobsen, 1935; Jacobsen and Nissen, 1937; Spaet and Harlow, 1943). Additionally, animals with dorso-lateral PFC lesions often display increased aggressiveness and muted emotional responses (Brody and Rosvold, 1952). Because of its association with cognitive aspects of behavior and temporal scheduling, we are most interested in the rodent analog of the lateral PFC. This region in the primate is also marked cytoarchitectonically by the presence of a granular layer IV, which is minimal or absent in the orbital and medial regions. The gross anatomical markers that separate these three regions (the presylvian fissure in most carnivores for example) are not present in the rodent, and thus of no use for identifying either the extent of the prefrontal cortex or the specific divisions between areas, so we must rely on primarily homologous connections to identify rodent PFC.

In the rodent, the mediodorsal thalamic nucleus projects primarily to two areas, a medial area in the upper edge and medial surface of the hemisphere (which

we will refer to as mPFC), and an inferolateral area in the lower lateral aspect of the hemisphere (which we will generally refer to as OFC)(Zilles, 1985). In rodents, contrary to the name, the mPFC generally seems more analogous to the primate lateral PFC, as it receives its input from the lateral (parvocellular) mediodorsal nucleus of the thalamus, while the OFC receives input from the medial (magnocellular) mediodorsal nucleus. The homology from mPFC to lateral prefrontal areas in the primate remains controversial for several reasons, primarily that all regions of the rodent PFC lack an identifiable layer IV, a feature shared by the primate medial and orbital aspects but not the lateral aspect. However, Uylings et al. (2003) point out that cytoarchitectonic differences across species are not uncommon, citing the believed homology of the motor cortex of the rat (which is granular) to the motor cortex of the primate (which is agranular in adults). For these reasons they believe that the correct homologies between primate PFC and rodent PFC need to be considered primarily on the basis of similar connections between various brain regions, similar functional properties, and similar distributions of neurotransmitters and receptors, and only secondarily on developmental and cytoarchitectonic cues. Thus we will now consider the connections and functions of the various rodent prefrontal regions.

The rodent mPFC is typically divided into three or four regions. The inclusion of the most dorsal of these regions as PFC proper remains controversial, as does the name of the region itself: it is alternately known as the frontal cortical area 2 (Fr2, (Zilles, 1985)), medial Precentral Area (PrCm, (Uylings et al., 2003)), medial agranular area (AGm, (Donoghue and Wise, 1982; Vertes, 2004)), or secondary motor area (abbreviated as both M2, (Paxinos and Watson, 1998), or MOs, (Swanson, 1992)). Previously this area was believed to be more similar to a pre-motor area due to assertions by Preuss (1995) and Condé et al. (1995) that it receives most of its thalamic input from the intralaminar, ventrolateral, and ventromedial nuclei rather than the mediodorsal nucleus. However more modern opinions (Uylings et al., 2003; Hoover and Vertes, 2007) view Fr2 as part of the PFC proper. We include it in this overview of the region, though our experiments

made no recordings from this area, so less attention is paid to it than the other regions. From dorsal to ventral the four regions of the rodent mPFC are frontal cortical area 2 (Fr2), the (dorsal) anterior cingulate cortex (ACC), the prelimbic cortex (PL), and the infralimbic cortex (IL). Due to the evidence given above, these regions are the most likely candidates to replicate the functional role of the primate lateral PFC, particularly the PL and ACC regions. As with prefrontal areas in all species, the regions of rat mPFC have diffuse projections to and from many cortical and sub-cortical regions. However, the afferent connections to the more dorsal areas (Fr2 and ACC) tend to include more input from the sensory cortices as well as sensory thalamic areas as opposed to the more ventral areas (PL and IL) which receive more input from limbic cortical and thalamic regions. The IL in particular receives most of its input from limbic and affective regions. Additionally, the IL seems to be the region of rodent mPFC least interconnected with the other prefrontal regions (Hoover and Vertes, 2007). These differential connections establish a gradient across the dorsal-ventral axis of mPFC from strong sensory connections in the Fr2 and dorsal ACC, through to more cognitive-like processing in the ventral ACC and PL, to eventually more affective and visceral control in the IL. Based on these connections, PL (along with the ventral aspects of ACC and perhaps the dorsal aspects of IL) is now thought to be the region best poised for cognitive processing in the rodent prefrontal cortex, and the most likely analog to lateral PFC in the primate (Vertes, 2006; Hoover and Vertes, 2007).

Functional studies of rodent PFC also support these homologies. PL lesions in rats have produced deficits in delayed response tasks (Brito and Brito, 1990; Seamans et al., 1995; Floresco et al., 1997) similar to those seen in lateral PFC of primates (Kolb, 1984; Goldman-Rakic, 1994). Additionally, PL has been implicated in many recent cognitive studies, particularly of the cognitive or strategic aspects of tasks (Rich and Shapiro, 2009; Jones and Wilson, 2005; Benchenane et al., 2010), as has the ACC (Durstewitz et al., 2010; Ma et al., 2014). Meanwhile, IL has been associated with fear processing (Quirk et al., 2000, 2006; Quirk

and Beer, 2006) and visceral control (Neafsey et al., 1986; Terreberry and Neafsey, 1983; Hardy and Holmes, 1988). Some studies also support the role of Fr2 as a sensory integration and motor planning region (Corwin et al., 1986). As a result of this evidence, we will concern ourselves primarily with the PL, with some recordings from the IL as well.

1.2 Behavioral tasks in Rodent PFC

From the earliest days, recording from rodent prefrontal cortex has been a rewarding but frustrating experience. Cells in rat prefrontal cortex vary their firing to a large number of circumstances, in ways that often seem suggestive of random variation (see Chapter 3). This complexity is not surprising given the general role of PFC as an area that coordinates behavioral outputs (motor responses), in response to changing stimuli (sensory inputs), using higher cognitive processing (neither directly sensory nor motor) to recognize latent variables and plan for future actions. Given the diverse sets of inputs, outputs, and functions of the region, we could expect to see cells responding to all of these different parameters, and that is what has been found. A common adage among PFC electrophysiologists is “if its present in the task, there’s a cell in PFC that codes for it.” Adding to the difficulty is the manner in which many variables are encoded. Despite early hopes that firing patterns as clear as those found in hippocampal place cells (which are themselves deceptively complicated (Redish, 1999; Eichenbaum, 2000; Eichenbaum and Cohen, 2014)) would be discovered in prefrontal cortex, it appears that single cell firing properties have remained perplexing. From these studies, however, several promising areas of research have emerged.

1.2.1 Spatial Firing Properties of rat PFC

The earliest electrophysiological studies in rat prefrontal cortex began appearing in the early 1980s (Sakurai and Hirano, 1983; Sakurai and Sugimoto, 1985b, 1986),

but major recording work in this region didn't take off until the late 1990s. This second wave of recording in PFC, informed by the successful work on place cells in the hippocampus (O'Keefe, 1976) and connections from CA1 hippocampus to the rodent prefrontal cortex (primarily prelimbic and infralimbic regions), concentrated heavily on identifying spatial firing correlates in rodent PFC. Poucet (1997) specifically designed his experiment to search for prefrontal spatial firing correlates. He recorded from prefrontal cortex on rats as they were performing a foraging task in a large (76 cm diameter) open field cylinder, hoping to find spatial correlates to cell firing on this task similar to those found in HPC. However, a large percentage of cells showed no consistent spatial firing patterns (very commonly firing all across the cylinder). Of the cells that did demonstrate some spatial firing, comparison of these cells across multiple sessions failed to reproduce similar firing patterns, prompting the author to check to see if spatial firing patterns in cells could be better explained by behavioral firing correlates that happened in specific locations on some sessions. Indeed, when behavioral firing properties were considered, these accounted for all spatial information. The author concluded that no spatial information was encoded by PFC neurons.

However, the very next year, work by Min Jung in Bruce McNaughton and Carol Barnes's lab further complicated the story (Jung et al., 1998). Jung et al. also set out to consider firing properties in PFC, with an emphasis on spatial firing properties. They recorded activity from PFC while rats navigated an 8-arm radial maze, a figure-8 task, and a square box (it is worth noting that while the Poucet's recordings were only from PL and IL (Poucet, 1997), Jung et al. recorded from virtually the whole rodent mPFC, though no comparisons were made between regions). On the tasks with limited tracks (the radial and figure-8 maze) they found cells with localized and specific spatial firing patterns. However, these patterns still tended to be rather larger than what would be expected from hippocampal place cells and often were repeated consistently across the track, for example firing at the same location on each arm of the radial arm maze. While these patterns were certainly more indicative of spatial firing patterns than

those seen by Poucet in the open field, the authors proposed that perhaps the firing represented not spatial information per se, but simply consistent cognitive processes, or consistent behavioral actions, that occurred at specific maze locations on these tasks.

In general most subsequent studies of the spatial firing properties of rat PFC have found similar results, with cells recorded on tasks that incorporate limited track environments with specific behavioral tasks typically displaying non-uniform firing properties (de Saint Blanquat et al., 2010) and tasks that use an open field foraging environment typically not finding spatial patterns (Gemmell et al., 2002). However, a study by Hok et al. (2005) further complicated the picture. Working in Bruno Poucet's lab, the authors employed the familiar large circular open field environment, but in this case they had animals perform a different task. Animals prompted the release of a food pellet by occupying a trigger zone location on the track. Pellets were consistently released into a landing zone, from whence they bounced around the open field and had to be tracked down and consumed. In general, animals spent much of their time at or headed to the trigger zone, until a pellet was released, at which point they reliably headed to the landing zone to start their search for it. This task design thus produced two specific locations, the trigger zone and the landing zone, that were spatially salient to the task as a whole, and also could be distinguished from consistent behaviors and the receipt of reward. On this task, the authors did find cells ($\sim 25\%$) in PL/IL that coded specifically for place more reliably than any other behavioral factors. The fields of these cells were concentrated at the trigger and landing zones, and they remained significantly larger than place fields detected from hippocampal place cells, but they did demonstrate spatial correlates that were not attributable to other factors.

The most common conclusion from these studies is that PFC is not optimized to encode for spatial locations, however it is capable of encoding spatial locations of high goal salience even in an open field environment. However, the precise role of spatial encoding in PFC remains controversial. Do these cells encode information about space independent of the specific structure of tasks? Or are the spatial

firing patterns in prefrontal cortex indicative of some other coding properties of significance, if not space per se? We will return to these questions in Chapter 4.

1.2.2 Encoding of multiple task parameters in rat PFC

Since it is apparent that spatial location is not the primary parameter encoded by neurons in rat prefrontal cortex, we are left with the question of what else might be encoded. Jung et al. (1998) did a nice job of exploring firing patterns in the single cells they recorded, and in many ways what they found is emblematic of the difficulty of characterizing single cell encoding patterns in PFC. After assessing the spatial firing patterns of their prefrontal cells, the authors searched for working memory cells on the two tasks. For their 8-arm radial maze task, rats had to visit each arm once per session in order to receive reward that had been placed there. This task tested the animal's memory of previously visited arms, because multiple visits to each arm would prove temporally costly. In addition, a random selection of four arms were initially blocked to prevent the animals from simply visiting all arms in sequential order. For this task, the authors surmised that some working memory cells would increase their firing rate as a specific arm was visited, then maintain that increased firing rate until the end of the session thereby marking that arm as visited. They further hypothesized that this behavior would lead to an increase in overall population firing rate in the PFC as the animal visited more arms until the end of the session. They were unable to find this population firing rate increase. Instead, they found that 18 of the cells increased their firing rate over the course of the session, and 22 cells decreased their firing rate, leading the population firing rate to remain much the same. While they concluded that the firing rate changes were not indicative of working memory cells as they had hypothesized them to exist, they did uncover a common tendency in PFC cell populations with the bimodal change in firing rate. Some subsequent studies (Ma et al., 2014; Powell and Redish, 2014) have discovered that firing rate changes among PFC cells seem to be balanced between cells that increase and decrease

their firing rates in response to each encoded parameter, causing overall population rates to remain rather constant. We will briefly consider this property in Chapter 3.

While Jung et al. (1998) were unable to find working memory cells in rat PFC, the area has been linked to working memory by inactivation studies in delayed alternation tasks, both spatial (Yoon et al., 2008) and relatively non-spatial (Horst and Laubach, 2009). Classical working memory cells (which maintain firing selective for a particular outcome) as reported in primate studies (Fuster et al., 1971, 1982; Kojima and Goldman-Rakic, 1982; Funahashi et al., 1989) seem to be more difficult to detect or simply more rare in rodents studies, though some evidence of their presence has been described (Horst and Laubach, 2012). It does appear that the rodent mPFC is involved in working memory tasks, though the specific firing patterns of the cells involved remains unclear. Additionally, some work has indicated the presence of cells in the rodent mPFC that are sensitive to interval timing (Kim et al., 2009; Narayanan and Laubach, 2009; Kim et al., 2013), which represents a slightly different type of sustained firing across delay periods. Kim et al. (2013) utilized a two choice decision maze with a retractable drawbridge that imposed a delay period on rats before they could reach the choice location and choose a reward arm. When animals approached the drawbridge, a tone sounded to indicate the beginning of a specific delay interval, at the end of which, the drawbridge would lower so animals could continue through the choice point. The length of the timing interval (from 3-5 seconds) determined whether the animal would receive reward on the left or right side of the track following the delay. Previous results (Kim et al., 2009) indicated that following temporary inactivation of the mPFC (using muscimol), an animal's ability to discriminate tone lengths on this task (and his ability to preform) dropped off significantly. In their 2013 study, the authors recorded a large population of neurons in the PFC with sustained firing rates throughout the delay, which also displayed some direct differentiation of long vs. short trials (necessary for the decision of which direction to go at the choice point). The authors were able to decode whether a trial was

long or short (and also which of six presented delay lengths was imposed on a given trial) from the populations of neurons they recorded. Classifications could be readily made from relatively small numbers of neurons (15-20) but greater accuracy was obtained from larger populations.

Another measure of delayed activation in rat prefrontal cells that has been documented is a persistent activation following error trials. The work of Narayanan and Laubach (2008) in particular has documented these responses. In this study, rats had to hold down a lever for a consistent time interval (1.0 sec) and then release it to receive reward. The authors discovered that after an error trial the behavior of the animals changed on the next trial. Specifically, they tended to take longer to release the lever following the end of the delay period after error trials. The authors further discovered that this post-error reaction time increase was abolished in animals whose medial PFC was inactivated by a muscimol infusion, strongly implying that the post error slowing was mediated by the PFC. They went on to demonstrate that 30% of their neurons had differential responses to error trials than correct trials. This effect could be due to the presence or absence of reward, but of these neurons 64% maintained a post-error firing rate difference throughout the inter-trial interval and on to the next trial. Coupled with the effect of PFC inactivation, these findings strongly indicate that PFC cells are modulated by correct or erroneous outcomes on trials, in other words by the outcome of goal directed behaviors.

Subsequent studies on the effect of error correlates in PFC have found some evidence for electrophysiological changes that might *predict* errors. As unlikely as this evidence sounds, in general it seems as though predictive error effects are more suggestive of failures in the system required for correct behavior. For example, Jones and Wilson (2005) discovered that theta frequency LFP coherence between CA1 hippocampus and PFC was increased as animals reached the decision point on a spatial navigation task, and that this coherence was higher on correct laps relative to error laps. Additionally, in the Kim et al. (2013) study described above, the authors were able to demonstrate that on error trials the time bin decoded

from the neuronal population was typically incorrect, indicating that likely these errors occurred because the PFC classified the length of the delay incorrectly. This point is specifically convincing when combined with the results of their previous study (Kim et al., 2009) which demonstrated the necessity of the PFC to solve this task. These studies provide evidence that prefrontal cortex has a role in coordinating goal-directed behavior, which perhaps implies that the characteristic firing patterns displayed on error laps are indicative of a signal granting salience to errors in a structure that is working to prevent them. We will briefly consider the role of error related firing in the PFC in Chapter 3.

A relationship between rodent prefrontal cortex and goal directed behaviors would make sense given the evidence that PFC is involved in goal directed behaviors in other species. In addition, we already saw that PFC could mediate spatial goal directed responses on an open field task (Hok et al., 2005), and further evidence of this fact was provided by Pratt and Mizumori (2001a). In this study, the authors employed the familiar 8-arm radial maze, and discovered cells with specific firing patterns indicating anticipation of rewards. They further showed that for some cells, these patterns could be separated from firing during consumptive behaviors (which were present in other cells), and that they were related to the anticipated amount of reward, even in cases when the actual amount of reward received was altered without the animal's knowledge. These cells therefore provided reward anticipation information independent of the actual receipt of reward and any specific stimuli that could be carrying information (because their firing tracked the anticipated larger reward, not the smaller reward actually available). We will also consider reward related firing on a PFC task in Chapter 3.

Taken together, it seems that neurons in PFC respond to a number of task relevant parameters, but in particular the region seems to play a part in organizing goal directed behavior. We now turn to a more specific consideration of how task relevant behavioral patterns are represented and executed in the PFC.

1.2.3 Strategic encoding in the rat PFC

The first careful examination of the role of rat PFC in encoding the strategic rules necessary for an animal to solve a decision making task came from the work of Rich and Shapiro (2009). These authors used a classic “plus maze” paradigm, in which animals followed either a place strategy (i.e. go east or go west, regardless of start position) or a egocentric response strategy (i.e. turn left or turn right, regardless of start position) to receive reward on any given trial. Animals were started at either the north or south end of the maze on any given trial and had to use the current response strategy to choose either the east or west arm to receive reward. After establishing a high level of performance on a given task criterion, animals were subjected to either a strategy switch or a criterion reversal without warning. In the strategy switch cases, the overall rule for correct laps was changed from a place to an egocentric strategy, or vice versa. In a reversal, the same general rule was kept in place, but the response criterion was reversed, i.e. go left became go right, or go east became go west. Of critical importance, for the strategy switches at least one physical path (i.e. from the north to west arm) would remain correct both before and after the switch.

The authors recorded neurons from PL and IL while animals performed these tasks, and discovered that the firing rate in a significant proportion of cells differed between paths taken before the strategy switch to after the strategy switch. The authors compared the proportion of cells that changed their firing rate on the same spatial paths before and after a strategy switch to the proportion of cells with different firing rates on the same path during stable performance, and found that a higher percentage of cells changed, even on the same path, in response to a strategy switch. Also, they compared the firing rate on paths that were not conserved before and after a strategy switch and found a larger percentage of cells with different firing rates on these two different paths than they found for similarly unmatched paths before and after a criterion reversal. In other words, cells changed their firing rates in response to a switch even when the animal was running the same path, and more cells changed their firing in response to

task rules than changed their firing rates in response to a reversal of the reward criterion within that rule. The authors further demonstrated that the change in the population firing in PL representing the change in strategy seemed to occur after the switch had occurred on the task, but before the animal's behavior changed to reflect the new strategy. Additionally, they showed that the change in firing patterns in PL seemed to precede the change in IL, which fits with our previous conclusion that the prelimbic cortex is the the most cognitive region of the rodent mPFC. This result was critical in establishing that the rodent mPFC is involved in the encoding of task rules that the animal uses to guide goal driven behavior.

Durstewitz et al. (2010) extended these results by examining some of the specific temporal dynamics of the transitions between different rule sets in PFC. They used an operant box setup for their experiment, with animals pressing levers for reward cues. The two different rule states they utilized involved switching between a stimulus guided decision process (press the lever with the illuminated light above it) and a place guided decision process (press the lever on the right side). The authors found again that both in single neurons and at the population level, discriminations were usually cleaner between the specific rule used (stimulus or place) than between the different cues within each rule (light right vs. light left, or right vs. left), although cells did distinguish between the different cue types as well at both the single cell and population levels. They used a Hidden Markov Model to model population transitions in a multi-dimensional firing rate space between one rule type and another and found that the transitions occurred rather rapidly from one state to another. In other words, the transition resembled a step change more than a gradual change over time. We should note, however, that these step changes still took on the order of 20 trials or so during transition periods, but relative to the 120 trials that were run in the course of a given session that change is relatively abrupt (though we will see in Chapter 3 that faster changes can occur between pre-trained strategies on a spatial task). The authors

also compared the neural and behavioral change points and found that the transitions in PFC tended to co-localize with changes in behavior, with perhaps a slight tendency for changes in PFC to precede those in behavior over all sessions they examined. Overall, Durstewitz et al. demonstrated that the PFC can also code for abstract rules on an operant box task as well as spatial tasks, and that the transitions encoded in the PFC population code occur rapidly, indicating a process that is recognizing a change in the environment, not a slow learning rate as had classically been seen in some psychology literature. Indeed, as the authors pointed out, it has been argued in more recent years that the rapid learning rates seen in classic literature may commonly be due to averaging behavior of individual animals that all show sharp behavioral transitions at different points in time together, which produces a gradual transition over the population as a whole (Gallistel et al., 2004).

Benchenane et al. (2010) also looked at the effect of PFC in managing rule acquisitions in association with hippocampus (through theta coherence) on a spatial task. They employed a spatial Y-maze track and a task that was essentially a spatial version of the rule shift task employed by Durstewitz et al. (2010). In this version, given the spatial nature of the task, the authors were able to replicate the results of Jones and Wilson (2005) that the theta coherence between mPFC and HPC in rats is heightened at the choice point of the task, but they also studied the temporal dynamics of this coherence as animals were forced to acquire a new rule governing task behavior by a strategy switch. They discovered that the coherence became even more heightened immediately following acquisition of a new rule (as the behavior under that new rule condition improved, implying learning). They also found a positive relationship between correct laps and theta coherence between HPC and PFC, implying that coherence at the choice point is important in order for animals to make the correct decision. Additionally, they discovered some evidence that this increased coherence is modulated by interneurons in PFC that become more active, suppressing the firing of some pyramidal cells during these periods, which may account for the heightened coherence. Their

finding confirms that spatially different zones can be treated differently by PFC on spatial tracks and suggests that these differences may be significant during rule acquisition. However, the timing of this coherence increase raises a fascinating question about the interaction between these two structures, as some evidence (Siapas et al., 2005; Jones and Wilson, 2005) indicates that the hippocampus leads the PFC in these coherent events, indicating a flow of information from HPC to PFC. However, the coherence increase they document comes AFTER the new rule has been acquired (i.e. after the animal's behavior has changed), and other studies have indicated that changes in PFC representations with respect to a new rule occur BEFORE the animal's behavior changes. The precise timing of these changes and the relationship between HPC and PFC on these rule changes remains to be determined.

A final paper to tackle the transition between abstract rule representations in the rodent was the work of Karlsson et al. (2012). This study employed a lever press task in an operant box setting, but unlike previous studies, the authors did not vary the rule that an animal had to learn to associate with changes in his environment but rather the animal's beliefs about likely rewards. Animals initiated a trial by pressing their noses into a nose poke port, at which point a tone would sound indicating whether the left or right hand lever would be rewarded. The animals could then press the lever on the appropriate side for reward. However, rewards were delivered stochastically, with a probability that varied on each side, such that one side was always rewarded more than the other (all probabilities between .15 and .8). Animals could also choose not to pull the lever associated with the tone, and instead return to the nose poke port to initiate another trial. After sufficient training, animals learned to press the lever primarily for the side that was rewarded at the higher probability, and tended to reject trials where the low rewarded lever was offered. Once behavior was stable in this state, the reward probabilities would be switched, leading animals to change their behavior. In general, they tended to enter a period of uncertainty in which they would sample from both sides more often than not, then they

usually transitioned to the opposite behavioral approach and began rejecting the side with lower reward probability (normally the side previously preferred). The authors examined the firing rates of cells from the mPFC of rodents on this task and discovered that there were generally abrupt transitions in firing rate following the laps where the animals behavior changed. These firing rate changes tended to occur rapidly, and usually were widespread across many of the neurons recorded simultaneously. The animals then entered a state for several laps that seemed suggestive of uncertainty, where the firing rate in PFC cells remained distinct and the animals accepted most of the offered trials, as though they were sampling the new reward probabilities. Finally, animals transitioned into a new state where they rejected the majority of trials on one side, and the firing pattern of the population of cells in their PFC reflected (in most cases) a new state of firing that seemed to not match the previous state. It is worth noting, however, that in several transition cases, after a period of the uncertainty related firing animals transitioned into a state that appeared very similar to the population firing state represented in PFC prior to the uncertainty period. The authors also found a reliable predictor of the onset of cellular transitions in an LFP increase in the high gamma band (65 to 140 Hz) that tended to occur in PFC shortly before a cellular transition in network state. The authors describe the transition that occurred in this study as reflecting a change in the beliefs of the animal, which is plausible from the presented data, but another interpretation would be a change in the behavioral state of the animal (from “generally accept the left side” to “generally accept the right side”). Experimentally, the distinction between these two interpretations is unclear, because there is not obviously a case in this study in which a belief change would occur without necessitating a concomitant change in the behavior, but it would be interesting to see cases where the probabilities of each side were changed, but the relative relationship was not (for example if the right side changed from 80% to 60% rewarded and the left side changed from 30% to 20% rewarded, would the resultant state in the PFC change, indicating a change in belief, or would it largely stay the same, indicating the behavioral representation

remained constant?) As we shall soon see, the distinction between task beliefs and rule representations and strategic representations may be important. However, the introduction of the intermediate state associated with the uncertainty behavior was a very interesting development (although some of the transitions described by Durstewitz et al. (2010) also were better explained by a transition into an intermediate state prior to settling into the final state).

All of these studies taken together have described the effects of PFC strategy transitions in great detail, but they all suffer from one critical limitation. In all cases, these studies forced animals to change their behavioral strategy by changing the parameters of the task. This is a very reasonable manipulation to make if searching for a change in a neural representation, but since we have already seen that PFC responds reliably to many different aspects of task structure, it raises the question of whether or not the representational changes being reported in these studies are truly changes in the strategy animals are using to solve these tasks, or merely the rules that govern their behavior on tasks. When the rules are changed to necessitate a strategy change it is very difficult to separate the two. However, on a task in which an animal chooses to change his own behavioral strategy at some point in order to maximize rewards, we would expect to see a transition in PFC even though the rules of the task have not explicitly been changed. This change would represent a true strategy change, and if we could detect such a change in cellular firing patterns, it should be generally correlated with behavior (although we would expect it to precede behavior or co-occur with it, not follow it). Similarly, if we would like to describe these strategic representations in PFC as truly cognitive, we should be able to separate them from behavioral changes as well as from sensory inputs. Therefore, we should be able to detect a change in the representations of PFC similar to a strategy change in situations where an animal's understanding of the underlying task changes without necessitating a related change in behavior. In Chapter 5 we will return to this issue and demonstrate examples of each of these types of strategy changes based on cellular population states in PFC. For now, however, we turn to one more critical aspect

of decoding cognitive representations from PFC that we have neglected to address thus far.

1.2.4 The Euston-Cowen Hassle

Amongst other correlates of firing properties attributed to mPFC by Jung et al. (1998) was a reference to various behavioral firing properties of cells. Some cells displayed firing properties related to specific patterns of running, and one cell distinctly appeared to be a right turn cell. These firing properties were little remarked on at the time, and PFC does have some outputs to motor cortex, so it was not surprising that it could have some motor pattern based firing. Indeed, some early studies on the rodent mPFC even considered the possibility that the ACC was more like a homolog of premotor region in primates (Neafsey et al., 1986; Passingham, 1986; Reep et al., 1987, 1990). However, a pair of significant papers (Euston and McNaughton, 2006; Cowen and McNaughton, 2007) cast doubt on many of the studies that had described cognitive roles for PFC and raised the possibility of another interpretation that now has to be accounted for in all studies.

The goal of the Euston and McNaughton (2006) paper was to examine whether working memory signals in PFC could serve as a marker to distinguish two identical spatial segments that were run as part of different overall sequences. Animals were trained to run paths that consisted of several journeys between different towers arranged around a circular track. The path sequences each consisted of several trips between towers, and always included two segments that were identical even though the steps before and after them in the sequence were different. The authors believed that these two segments would be run identically, and as such should be represented very similarly in spatial representations (in HPC), but that some brain signal must encode that they were traversed as part of two different overall sequences. They recorded from mPFC (at several different depths) and found several cells that fired differentially on the same path segments depending on which overall sequence they were a part of. However, they also found that

there were distinct behavioral differences on the two path segments, particularly on the first path segment, which also always had a higher percentage of cells firing differentially based on which overall sequence it belonged to. In fact, they discovered that the path differences on the segments between the different sequences accounted for the firing rate differences observed in the cells better than the identify of the overall sequence they belonged to. The authors concluded that in many cases, firing differences in PFC that are attributed to working memory or other cognitive factors may merely represent subtly different paths taken by the animals as they solve various behavioral tasks. Following this paper, the onus was placed on all future papers to demonstrate that task properties encoded by PFC could not simply be due to path differences taken by the animal.

Many subsequent studies attempted to counter this problem by employing operant box tasks, which are meant to constrain an animals movements so that they cannot account for firing rate differences. However, this approach has never really been a viable option because one year after the Euston and McNaughton (2006) paper another paper from the same lab (Cowen and McNaughton, 2007) showed a similar effect while utilizing an operant box setup. The apparatus in this experiment was carefully designed to prevent the animal from having to make any extraneous movements. Stimuli were provided by speakers arranged all around the animal as well as two vibration motors on the animals left and right sides that could operate independently to provide different sensory stimuli. The task design was a complex go-no go task, but responses only required the animal to reach a single nose-poke port, and either maintain his nose in the port or remove it to abort a trial that was unlikely to be rewarded. Animals had to recognize contingency relationships between different pairs of first stimuli (tones played on the speakers) and second stimuli (patterns of vibration) to determine whether it was worth waiting through the rest of the trial for reward or not. The setup was thus consistent with classic working memory protocols where the animal must maintain a memory trace of the initial stimulus over a delay period in order to understand the meaning of the second stimulus correctly. And again, the authors found cells

that had delay firing properties suggestive of a working memory code during the delay period. However, they also found subtle behavioral cues during delay periods that could also account for the differential firing of the neurons. Animals would position their heads differently based on the first cue during the delay period until the second cue occurred. The authors again found that these subtle positional changes better explained the firing differences seen among neurons in the PFC during the delay period, meaning that the neurons could have been not encoding a memory of the task, but merely sending a signal to maintain the behavioral shift over the delay.

The challenge presented by these two papers has come to be known as the Euston-Cowen Hassle, and simply put, it means that all papers that purport to show a cognitive firing rate difference in PFC must show that in fact neural firing pattern differences are not simply due to subtle differences in the positioning of the animal. There are a few mitigating factors that should be mentioned: First of all, in both studies the authors recorded at several different depths in mPFC starting at the most dorsal levels of ACC, and in both cases they found that the percentage of cells sensitive to the subtle positional changes decreased as the neurons were lowered deeper into mPFC. This means that recordings made from deeper structures, like the PL and IL, are less prone to difficulties with the Euston-Cowen Hassle. The second point is less technical and more philosophical: It is quite possible that the animals are simply employing an *embodied cognition* strategy to help them in solving a difficult behavioral task. The use of a physical strategy to help alleviate memory load by maintaining a physical representation through postural positioning is a reasonable tactic for animals to employ, especially when their nourishment often depends on these memories in rather complicated tasks. Even if the stakes were lower, any mechanism that can serve as a reminder is a useful adaptation to supplement imperfect biological processes. Furthermore, even if the role of PFC in these tasks is simply to engage and/or decode the meaning of the gesture used to code for a memory, that is still a significant role in the process of working memory. And while there is a strong correlation in

many cases between the postural deviations and prefrontal firing, that does not necessarily mean that PFC is incapable of supporting memory in the absence of postural deviations, if trials absent of postural deviations have not been recorded. Nevertheless, experiments that would demonstrate cognition in rodent mPFC will make a much more compelling case if they can demonstrate that firing properties of cells cannot be accounted for by postural differences in the animal, and so the Euston-Cowen Hassle is still a significant consideration. We therefore designed out experiments to allow us to account for the animals behavior as reasonably as possible, and will demonstrate whenever necessary that our firing rate differences are independent of the animals posture.

Chapter 2

Tasks and Recording Methods

All experiments described in this thesis were carried out on Male Fisher-Brown Norway Rats (three for the Multiple-T task, three for the Delay Discounting task, one for Multi-track task, seven total) aged 8-12 months at the start of behavior (because only male rats were used, the masculine article is used to refer to rats throughout this thesis). These animals were housed on a 12 hour light-dark cycle and all experiments for a given rat were run at the same time every day during the light phase. Animals were handled and trained to eat flavored food pellets (45 mg each, Research Diets, New Brunswick, NJ, USA) for one to two weeks prior to the start of task training. These same food pellets were used as a reward on all tasks. To motivate task running, animals were deprived of food in their home cages during their period of task running, but were always given free access to water in their home cages. Animals were weighed daily and were returned temporarily to free food access if their weights dropped below 80% of their initial free food weight or if they displayed any signs of illness. Additionally, any animals who did not receive enough food to maintain weight while running at criterion levels on the tasks were post-fed by hand after running to maintain their weight throughout the experiments. All training procedures were approved by the Institutional Animal Care and Use Committee at the University of Minnesota, and were in accordance with the National Institutes of Health guidelines.

2.1 Experimental Recording Techniques

We used tetrode electrophysiological recordings to address the role of rodent Prefrontal Cortex on spatial tasks and decision-making. Tetrode recordings allowed us to record a large ensemble of putative neurons (in the prefrontal cortex usually 20-30 cells simultaneously) from awake behaving animals. Tetrodes consisted of four single wires made from either nickel-chromium or platinum-iridium, wrapped together to create a single physically implantable unit with four different recording channels. These channels allowed us to measure extracellular voltage changes resulting from the activity of nearby neurons (up to approximately $50 \mu\text{m}$) in the brain (Buzsáki, 2004). We recorded two different forms of signals from the PFC, local field potentials (LFPs) and “spikes”, or putative action potentials from single neurons.

LFPs are overall changes in the external electrical potential throughout a local region of the brain that tend to oscillate at various frequencies. Many of these frequencies have been well characterized, tend to occur in certain regions, correspond to certain behavioral correlates, and are thought to be significant for communication between different brain regions over time. LFPs are created by the combined activity of many neurons in the local area rather than a single neuron, however much is still unknown about how LFPs are generated and exactly what role they play in neuronal computations.

In this thesis, we will concern ourselves primarily with spikes, and it is in the recording of spikes that the advantages of tetrodes over single wire recordings are most observable. Spikes are detected as sudden significant changes in the potential recorded on at least one channel of the tetrode. Spikes can be triggered by a single channel, although in general they will appear on multiple tetrode channels simultaneously, usually all four channels. These events are then recorded at high temporal resolution across all four channels. The advantage of having recorded multiple channels of the same event is that signal coming from the same neuron will be consistently modulated on each channel by the physical relationship between

the neuron that emitted the spike and the four different wires in space. For example, spikes will tend to have a higher amplitude on the nearest channel. Spikes can be categorized later off-line through a process known as spike-sorting or cluster cutting. This process involves organizing individual spike events into groups or clusters that all presumably originated from the same neuron. It is this process that allows scientists using tetrodes to record spike-trains from multiple different neurons simultaneously on a single tetrode, which allows for the recording of much larger simultaneous ensembles than can be achieved using single-wires.

For all of the studies described in this thesis, LFPs and spikes were recorded on a 64 channel analog Cheetah recording system (Neuralynx, Bozeman, MT, USA). Spikes were recorded online using the recording system's built in filters, then sorted offline by first separating them into putative clusters using KustaKwik (KD Harris) and then individually verifying and adjusting those clusters with the MClust 3.5 Software package (AD Redish). The Cheetah system registered all recorded activity to timestamps, which were co-registered to the animal's behavioral activity and the task based activity recorded through custom written Matlab (The Mathworks, Natick, MA) software which was used to control the tasks described in the next several sections. The 64 channel system allowed us to record spikes from 12 tetrodes ($12 \times 4 = 48$ channels), in addition to an LFP trace from each of the 12 tetrodes (12 channels), and two reference LFP channels (independent electrodes that could be used to assess the depth of the other tetrodes relative to measured LFP signatures that are typical of certain anatomical features in the brain). The 12 tetrodes and two reference electrode channels were housed in a hyperdrive (Kopf), an enclosure that contained the tetrodes, a screw driven mechanism for independently lowering each tetrode into the brain in small increments, and the circuitry for connecting all of the recording channels to the recording system.

A hyperdrive containing the 12 tetrodes was surgically implanted into each animal using similar surgical procedures, which had been well established in the lab. Animals were anesthetized with sodium Pentobarbital (Nembutal, 50 mg/kg,

delivered IP) then placed in a stereotax for surgery. While on the stereotax, anesthesia was maintained with isoflurane mixed at 0.5-2% with oxygen and delivered via a nosecone. The scalp and fascia were carefully resected to expose the skull and several anchor screws were placed in the skull to hold the drive in place. We then drilled a craniotomy to implant the electrodes over PFC, placed the hyperdrive carefully over the craniotomy with the electrodes lowered to the brain's surface, and secured the entire assembly in place by building up a structure with dental cement. One of the anchor screws was connected to the drive to serve as a ground reference for LFP signals.

Surgical implantation into medial PFC in the rat is a tricky prospect, because the target structure is quite narrow and located at the midline of the brain. In order to record successfully from the PL and IL, tetrodes must end up between the midline of the brain and the corpus callosum, which encompasses about 1.5 mm of tissue laterally in either hemisphere. Additionally, the sagittal sinus, a major blood vessel, runs roughly along the midline of the brain between the two hemispheres. Care must be taken during surgery to avoid damage to the sinus, or the animal may not survive the procedure. In general two techniques have been proposed to implant into the deep layers of PFC while avoiding the sinus. The first is to implant the tetrode bundle sufficiently lateral from the midline to avoid the sinus, but angled toward the midline such that the tetrodes eventually come to rest in the target regions(see Figure 2.1 panel A). The second technique, which we developed in our own lab independently but in parallel with others, is to use a split bundle drive that has two groups of tetrodes that are implanted on either side of the sinus(see Figure 2.1 panel B). The difficulty with this technique is that the position of the sinus must be carefully assessed before the drive can be placed, because anatomical variation between rats is too high to accurately drill craniotomies from the skull surface. In order to accomplish this, we had to grind away the skull layers slowly to first expose the sinus, then carefully open final craniotomies on either side. This process is described in more detail below.

The first rat was implanted with all the tetrodes and references grouped in a

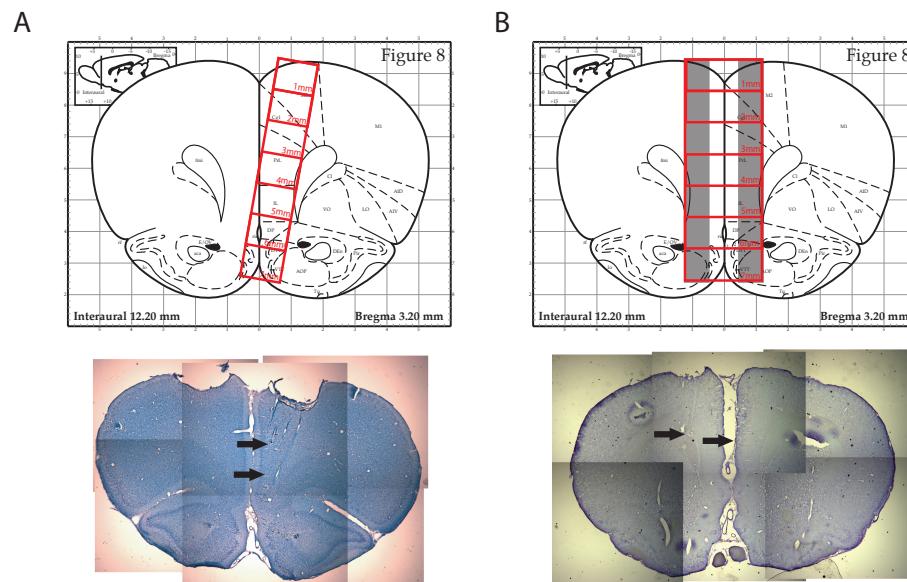


Figure 2.1: Surgical Implantation Techniques. A) The first surgical technique used a tilted single bundle drive targeted at the PL at an angle to avoid the midline sagittal sinus. Histology shows tetraode tracks from this technique. B) The second surgical technique used a dual bundle drive targeted bilaterally around the mid-sagittal sinus. Histology shows final tetraode positions from this technique. Figures modified from (Paxinos and Watson, 1996)

single circular bundle approximately 1.5 mm in diameter. The craniotomy for this surgery was located 3.0 mm anterior of Bregma, 1.3 mm lateral of the midline (on the right side), and angled at 9.5° from the vertical towards the midline (see Figure 2.1 panel A). Recordings from this animal were actually quite successful, but there were several downsides to this surgical technique. First of all, the angle of the drive had to be carefully assessed in surgery, which was challenging. However, the angle was critical because if it was too close to the vertical the tetrodes would end up in corpus callosum and miss the target region entirely. On the other hand, if the angle was too close to the horizontal the tetrodes would reach the midline before they reached the target region. In other words, they would reach the cortex at a dorsal location, for example ACC instead of PL. A related difficulty was that this technique could only target a single region (ACC, PL, or IL) since the trajectory of the tetrodes would only intersect one of those regions. Unless a wider bundle were constructed, we could not record from all regions with a single angled drive. Finally, although none of our studies found a difference between cells recorded from the right and left hemisphere, the angled implantation technique only targeted cells in one hemisphere for each rat.

All the subsequent rats used an alternate implantation technique with a split bundle design (see Figure 2.1 panel B). In these cases, the tetrodes were divided evenly (six tetrodes and one reference each) into two bundles with a 1.0 mm spacing between them. These bundles were then implanted such that one bundle was directly over the PL on either side of the sagittal sinus, 3.0 mm anterior to Bregma. In order to implant these bundles accurately, we had to expose the sinus by carefully removing layers of skull with a burr grinder in a 45,000 RPM drill (Freedom). Once we could visualize the sinus, we carefully removed the last layers of skull and dura and were able to position the two bundles to record from both hemispheres simultaneously. Another advantage of this implantation technique was that since tetrodes were lowered vertically to their targets in the PL, we could continue past the PL to record from the infralimbic cortex (IL) as well with the same tetrodes. In principle, we could also record from the ACC and even

FR2 on the way to PL, and we did observe some activity in these regions, but we did not optimize recordings for ACC, and none of those results are reported here. Due to the orientation of the rodent mPFC, tetrodes were implanted parallel to the laminar structure (especially in PL and IL). This orientation meant that a tetrode tended to remain in the same layer throughout the entire extent of both PL and IL. If the tetrode was in layer II/III or layer V, it would likely record strong signals from putative pyramidal cells at most depths. However, if it ended up in a relatively sparse layer, good recordings would be much less common and less dense. These tetrodes typically found one cell at a time and only at certain locations. The angled implantation technique has the potential to address this issue because the angled trajectory of tetrodes takes them across multiple cell layers. However, the difficulties with that technique outlined above still hold. An ideal system would allow us to target the dense cell layers with a vertical bilateral implant, but such precise targeting from the cortical surface during surgery is challenging.

After surgery, animals were given a three day course of antibiotics (Baytril) and returned to free food to recover for several days prior to being returned to task running. No animals were returned to running before their weights had stabilized and they appeared fully recovered from surgery. Tetrodes were turned in small increments every day following surgery until they reached the desired target coordinates. After reaching these areas, tetrodes were turned until they achieved good recordings. For the MT-LRA task, tetrodes with good cells were typically not turned throughout the switch sequence, which allowed us to record from the same cells across multiple days. On the DD task, tetrodes with good cells were usually left for multiple days of consistent recording, then turned down farther in an attempt to find more cells. As discussed above, tetrodes were implanted parallel to the laminar structure of the cortex, so tetrodes that ended up in a dense cell layer (such as layer V) would tend to have good ensembles of cells at multiple depths through PL and IL, while a tetrode that ended up in a sparse layer tended to have good cell recordings only at specific locations. Therefore we

could obtain the largest overall populations of independent cells by turning the tetrodes with good recordings relatively often.

2.2 The Multiple-T LRA task

The first task we used to study the role of rodent PFC was the Multiple-T Left Right Alternate (MT-LRA) task. This task allowed us to investigate many different behavioral parameters that had been individually connected to PFC simultaneously on the same task. Among these parameters were: differential encoding of correct and incorrect laps, shifts in behavioral strategies over time, navigational decisions, and reward related firing. Additionally, this task provided us with the ability to correct for subtle path and movement differences in behavior. The task has also been used many times in the Redish Lab to study a variety of brain structures (Schmitzer-Torbert and Redish, 2004; van der Meer et al., 2010; Blumenthal et al., 2011; Gupta et al., 2010; Steiner and Redish, 2012).

The MT-LRA task was run on an elevated sideways-8 shaped maze, depicted in Figure 2.2. On a single lap, rats started from the Start of Maze (SoM) region at the bottom of the track, and ran through three low cost choice points to a high cost choice point at the top of the track. The low cost choice points consisted of T-shaped regions that were pseudo-randomly arranged each day so that the animal would have to re-learn the sequence to turns. They were considered low cost because if the animal ran the wrong direction he was allowed to turn around and keep running; in addition, the animals could likely see which direction would allow him to keep running and which direction resulted in a dead end. The collection of low cost choice points is referred to as the Navigation Sequence (NS). At the end of the NS, the animal reached a high cost choice point where he had to choose whether to proceed to the left or to the right. Either side would take him past two feeder locations, where pellets would be dispensed automatically if he made the correct choice, and then back around to the SoM region again to begin another lap. We referred to the high cost choice point region as the Choice Point (CP),

the track between the CP and the feeder zones as the Top Rail (top), the area around and between the feeders as the Feeder Zone (Fed) and the area between the feeders and the SoM as the Bottom Rail (Bot). These divisions are depicted in Figure 2.2 panel C, and were used in our subsequent analysis.

At the beginning of a session, animals were placed on the track at the SoM location, then allowed to continuously run as many laps as possible for 40 minutes. Each successful lap resulted in the automated delivery of two reward pellets at each of the two feeder locations encountered, for a total of four pellets for a successful lap. The first feeder location on the left side of the track dispensed banana flavored food pellets, and the first feeder location on the right side of the track dispensed berry flavored food pellets. The second feeder location on either side dispensed unflavored, white-colored food pellets.

Whether or not the animal received a food reward at the feeders on a given lap was determined by whether or not he made the correct decision at the high cost choice point. There were three different reward criteria that specified the correct decisions on each lap, left reward (L), right reward (R) and alternating reward (A). In the left reward criterion, reward was provided only on the left side of the track. Similarly, in the right reward criterion reward was provided only on the right side of the track. In the alternating reward criterion, the animal received reward if he alternated laps to both sides of the track, so reward was only provided on a given lap if he went to the side opposite to the one he visited on the previous lap (under the alternate condition, the first lap of a session was always rewarded). During all training days, animals experienced only one pseudo-randomly chosen reward criterion per day. However, once animals were fully trained and behaving well, we began the 6-day switch sequence. During this period, animals began each session with a single reward criterion as with all previous sessions, but approximately halfway through the 40 minute session (randomly determined between 18-22 minutes) the reward criterion switched to one of the others. No external cue indicated this change in reward criterion, the animal had to learn it from failed reward incidents. The switch sequence consisted of six days so that all pairs

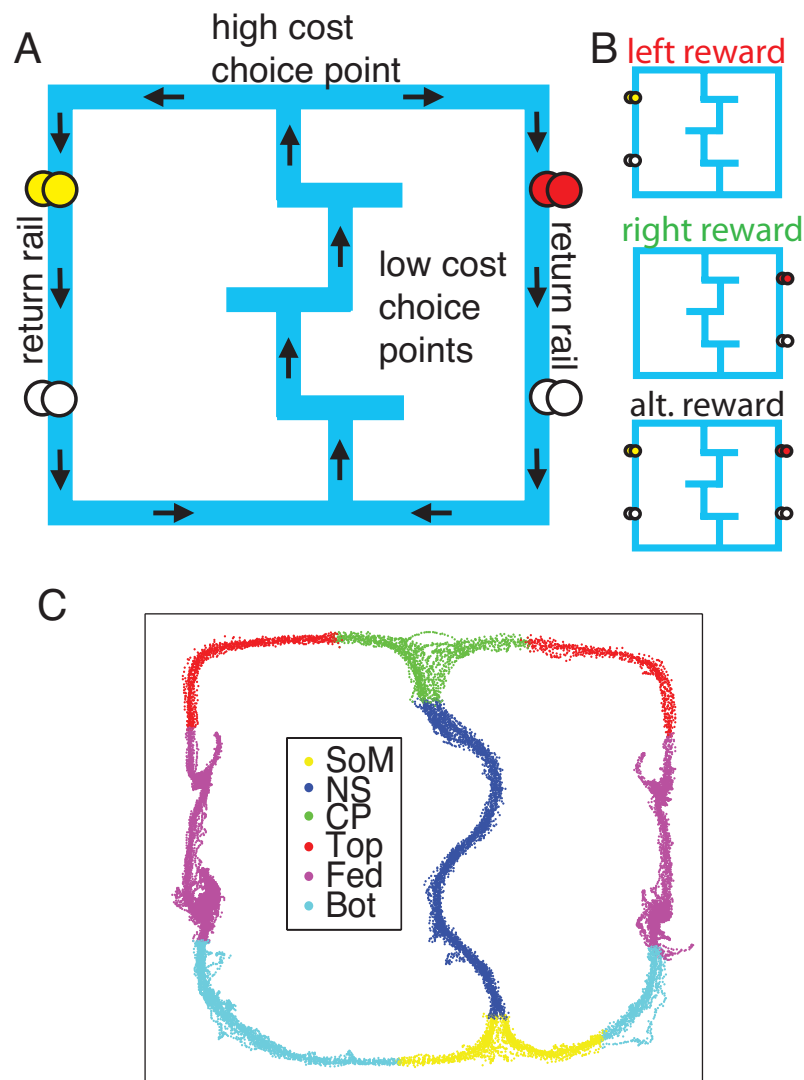


Figure 2.2: The MT-LRA task. (A) The task consists of a central track containing three low-cost T-choices, a high-cost T-choice at the top of the central track, and two return-rails on which reward was delivered under either *leftward*, *rightward*, or *alternating* contingencies. The first reward site on the left side provided banana-flavored food pellets, while the first reward site on the right side provided fruit-flavored food pellets. The second reward site on both sides provided unflavored (white) food pellets. (B) Reward was only delivered if the animal made the correct choice under the active contingency. (C) For analysis, the maze was divided into six sections (start of maze [SoM], navigation sequence [NS], choice point [CP], top rails [Top], feeder sites [Fed], and bottom rails [Bot]). Reproduced with permission from Powell and Redish (2014)

of initial and final criteria were presented.

The expected behavior on this task was for animals to rapidly identify the daily strategy over the course of the first several laps, then exploit that strategy reliably (i.e. perform laps at a high rate of success) up until the switch. At the switch lap, their behavior should drop down to perseveration levels (approximately 1/6 laps correct, the expected rate of continuing to use the now incorrect strategy) until they recognized their error, identified the new strategy, and returned to performing at a high level over the course of the next few laps. This is the exact behavior we saw for the rats we recorded from, as depicted in Figure 2.3, which displays the percentage of correct decisions averaged over all rats by laps aligned to the start of the session and the switch.

2.3 The delay discounting task

The Delay Discounting (DD) task is another task that has been utilized in the Redish lab recently (Papale et al., 2012) to examine animal behavior and the role of various brain regions in decision-making. DD is an inter-temporal choice task that takes advantage of the fact that delayed rewards are less valuable than immediate rewards to allow animals to reveal their own temporal discounting preferences by controlling the amount of time they are willing to wait to receive a larger delayed reward relative to a smaller immediately available reward. The task has been well described elsewhere (Papale et al., 2012) but we will briefly review the task structure and behavior.

The task itself took place on an elevated T-maze with two return rails that allowed animals to return to the start position after each lap so they could run continuously. This track is depicted in Figure 2.4 panel A. At the main T junction, animals ran to the left or right sides of the track, each of which had a feeder location that dispensed unflavored (white) food pellets. On one side of the track, the animals received one food pellet after a one second delay. On the other side of the track, animals received three food pellets after a variable delay. The

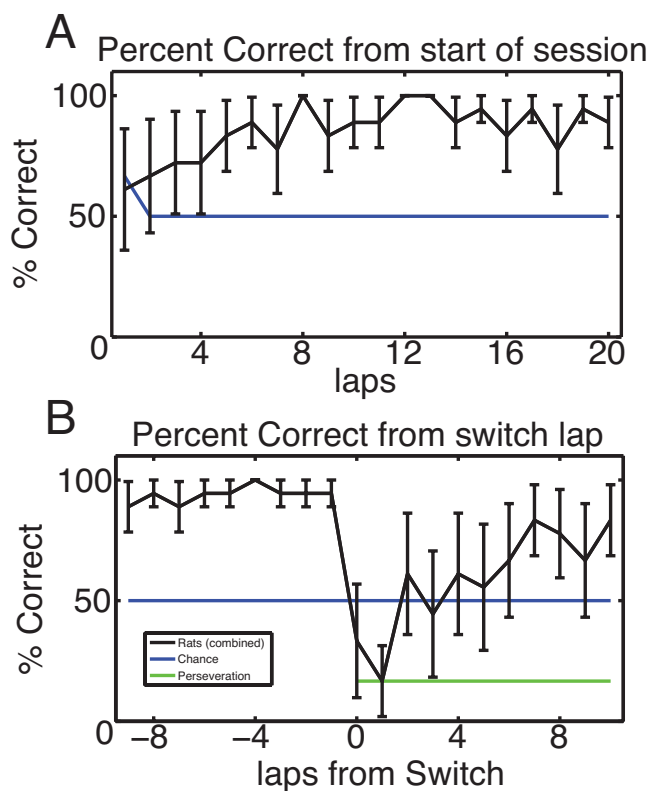


Figure 2.3: (A) Average percent of correct laps over all switch sessions aligned to start of the session. Blue line indicates chance rate of responding. Because the first lap of the alternating contingency was always rewarded, chance is 66% on the first lap, but 50% after that. Rats started at chance and quickly learned to make correct choices. (B) Percent of correct laps aligned to the switch. On the switch lap, rat percent-correct dropped to the expected level given no knowledge of the oncoming switch, but then rose back up to above-chance levels. Blue line indicates chance level of behavior, green line indicates perseveration level. Reproduced with permission from Powell and Redish (2014)

location of large and small rewards on the left or right side of the track was consistent within each session, but varied between sessions. The variable delay was initialized at the start of each session to a value pseudo-randomly chosen between one and 30 seconds (all values were sampled only once over the course of a 30-day training session, and were varied so as to not have too many high or low delays over several consecutive sessions). The variable delay also adjusted based on the animal's decisions, such that it increased by one second with every visit to the large reward side of the track, and decreased by one second with every visit to the low reward side of the track. This adjustment procedure allowed the animals to vary the delay by making multiple consecutive visits to either side of the track, or to maintain the current delay by alternating visits to each side of the track. All delays were indicated by a tone that counted down with decreasing frequency each second until reward was delivered by the automated feeder system (Med-Associates, St. Albans, VT, USA). These tones were standardized so that the same tone always indicated the same number of seconds of delay to wait. Longer delays were indicated by higher frequency tones, and all tones had a 175 Hz step size. Tones started counting down as soon as animals reached a delay zone region near the feeder on each side of the track, and continued as long as animals remained in this region until food was dispensed at the end of the delay. Animals were free to skip any delays by running past the feeder and on to the next lap (at which point the tones would stop and no food would be delivered). In practice, animals tended to wait out almost all delay periods.

The adjusting delay led to a typical pattern of behavior (depicted in Figure 2.4 panel C) seen over the course of a session which was typified by three distinct behavioral stages. First, animals would have a brief *exploratory* stage that lasted several laps during which they would determine which side had the large and small reward on a given session, and how long the initial delay was for that session. Then animals typically entered a *titration* phase, during which they made laps predominantly to one side of the track or the other to drive the delay value to a more preferable length. Finally, once the delay was adjusted so that the overall

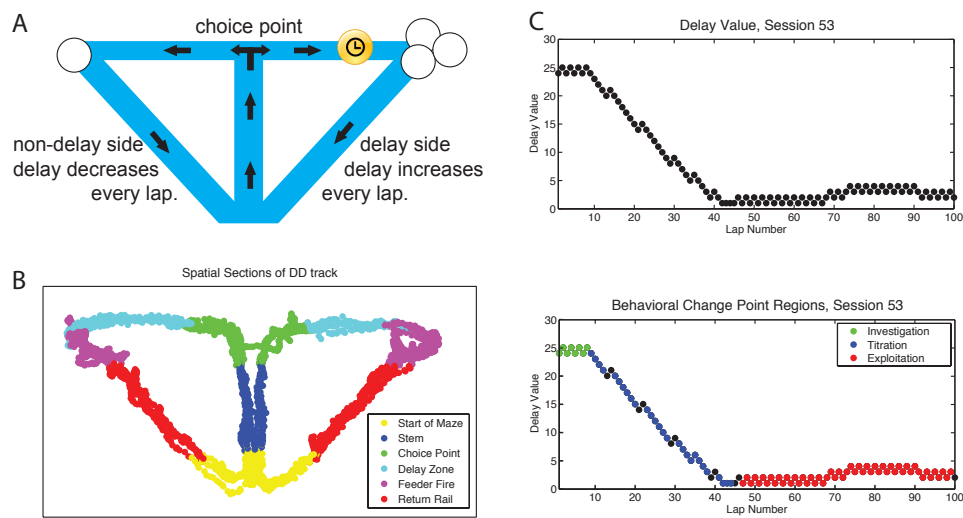


Figure 2.4: The DD task. (A) The task consists of a central choice point and left and right reward locations, one of which provides a large reward after a delay while the other provides a small reward immediately (B) For analysis, the maze was divided into six sections (start of maze [SoM], central stem [Stem], choice point [CP], Delay Zone[DZ], feeder sites [FF], and return rails [RR]). C) Animals typically displayed three behavioral phases on this task, Exploration, Titration, and Alternation, shown here for a typical session.

reward on each side of the track was equal, animals would alternate laps to either side of the track until they reached the maximum number of laps or time limit of the session. We called this last phase *exploitation*.

Animals ran daily sessions of either 100 laps or one hour, whichever elapsed first. The training phase of the task consisted of 30 sessions prior to surgical implantation so that every initial delay between one to 30 seconds could be sampled. Following recovery from surgery, animals ran another 30 sessions so that again every initial delay from 1 to 30 seconds could be sampled, while LFP and spiking data were recorded from PFC. In almost all cases, animals ran the full 100 laps before one hour had elapsed, and in general their behavior matched the typical session pattern (exploration, titration, exploitation) for most sessions. This task was particularly useful to us because the expected behavioral variation was suggestive of different strategies employed by the animal, but none of these changes were forced on the animal by changing their reward criteria or other external factors. Strategy changes on the Delay Discounting task reflect the animals internal decision process determining the optimal behavior over time.

2.4 The multi-track task

The Multi-Track task was designed to examine how non-uniform spatial firing in PFC neurons varied on the same task run on multiple different tracks, or with different cognitive tasks run on the same track. Thus far it has been run on only one animal (R236) who was already trained on the DD task on the standard DD track. This track was used as one of the two tracks in this task, the small track, in addition to a much larger track known as the HP maze, shown in Figure 2.5. The animal was also taught a variation of our MT-LRA task called LRA, where on each day his reward was contingent on following a single reward strategy (L,R, or A, described above) for the entire session. For the LRA task the animal ran for the same session criteria (100 laps or one hour, whichever came first) as DD, and received two pellets at the single feeder location on each side for a correct

lap. These task parameters were designed to make the two tasks as consistent as possible with respect to reward and session length.

The animal used for this task had previously run the DD task on the small maze, but had not run the task (or had his tetrodes turned) in approximately eight months. As a result, his tetrodes were still in their target locations, and the cells recorded were quite stable across sessions. These tetrodes were not turned any farther.

The rat began running the DD task on the small maze and after running well plugged in for a few days was introduced to the LRA task, one contingency at a time. After he rapidly learned these new contingencies, we introduced him to a series of interleaved sessions of both tasks on the small track, before eventually introducing the HP maze. Once the animal had learned to perform well on both tasks on the HP maze, we began reintroducing the small maze on some sessions. Once the animal was behaving well on both tasks and both tracks, we began a final 12 day test sequence, which interleaved sessions on the HP and small mazes, DD and LRA tasks, and all three reward contingencies for LRA and initial delays of 5, 15, and 25 seconds for the DD task psuedo-randomly. We used the procedure described in chapter 3 to identify cells consistently recorded across all 12 days and used these for our analysis of the multi-track task.




Figure 2.5: The Multi-Track Task. A-D Examples of different task and track conditions presented to the animal on any given day. Four examples are shown, a full set of contingent examples includes HP track and small Track, DD task and LRA task, and delay left or delay right for DD or Left, Right, or Alternate strategies for LRA.

Chapter 3

Simultaneous representation of multiple task-relevant parameters in PFC

As discussed in Chapter 1, there are many commonly ascribed roles for the Prefrontal Cortex (PFC) across species. Many of these roles (such as executive control, working memory, and temporal encoding) can be difficult to differentiate from each other for many tasks; what might seem to some observers to be a clear example of executive control might be interpreted by other researchers as working memory. In the specific case of rodent PFC, the picture is somewhat more clear as studies have tended to describe more precise encoding parameters rather than turning to broad terms such as executive control, but confusion still exists. PFC has been implicated in cognitive processing, working memory, interval timing, encoding of uncertainty, reward receipt, encoding of recently experienced errors, and behavioral strategies. Others have shown that neurons in rodent PFC display non-uniform firing patterns over space (which will be discussed more extensively in the following chapter), and of course the Euston-Cowen Hassle relates to the correlation of PFC firing to postural changes in the animals position. In general, these studies tended to examine a single aspect of PFC encoding at a time, and

thus not much is known about how these different parameters are encoded in the cortex. Are they all encoded simultaneously on a single task (if necessary), or does the cortex monitor only one of these task aspects at a time? If they are encoded simultaneously, are they encoded by separate populations of cells, or do these populations overlap? Can single cells in fact encode multiple aspects of a single task in the rodent PFC? We will attempt to answer these questions in this chapter by considering recordings obtained on the MT-LRA task.

The question of whether individual cells can simultaneously encode multiple task parameters has been generating increased interest recently, and has been addressed in the Primate PFC by a group of prominent researchers (Rigotti et al., 2013). These authors concluded that a significant population ($\sim 20\%$) of cells in primate PFC display firing that is modulated by multiple aspects of a task simultaneously. They described this phenomenon as “nonlinear mixed selectivity” and posited that it is a critical aspect of the encoding displayed in PFC. We will examine whether rat PFC cells display nonlinear mixed selectivity in the final section of this chapter.

3.1 Parameters of the MT-LRA task

3.1.1 Behavioral strategy encoding in PFC

Rodent PFC has been implicated in the encoding of strategic representations by several research groups. The basic established procedure for demonstrating that the PFC has encoded a strategic switch is to impose a change in the criterion for reward receipt on the animal performing the task and look for differential firing patterns between the epochs of the different reward criteria. If there is a difference in the firing of a cell, this is taken as evidence that it encodes the different strategies. We will explore the concept of strategy encoding by PFC in much greater detail in Chapter 5, but for now we will use this basic methodology to identify cells that encode strategy on the MT-LRA task (see Chapter 2 for a

thorough description of this task).

We recorded 330 putative units on different days from three different rats on MT-LRA. To assess whether any of these units encoded a strategic change on this task we compared the distribution of the firing rate for all laps prior to the strategy switch to the distribution of firing rates for all laps after the switch for each unit. We used a KS-test to assess whether these two distributions of firing rates were significantly different from each other and considered any cells with significantly different firing rates to encode for a strategy change. Using this method, we found 163 cells (49.3%) that significantly encoded for a strategic change based on their firing over the entire lap. Examples of the firing rates of some of these cells by lap can be seen in Figure 3.1 panel A.

We also examined the population of cells as a whole by calculating a parameter-based firing rate z-score for each cell. To do so we took the average of the firing rate of the cell on any given lap of type A (in this case laps after the switch), subtracted off the average firing rate of that cell for all laps of type B (in this case laps before the switch), and divided by the standard deviation of firing for all laps of type B:

$$z_B(A) = \frac{F_A - \mu(F_B)}{\sigma(F_B)} \quad (3.1)$$

Where F_A is the firing rate on lap type A, $\mu(F_B)$ is the mean firing rate from laps of type B, and $\sigma(F_B)$ is the standard deviation of the firing rate from laps of type B. We then took the average of this parameter over laps for each cell to get a final z-score for that cell. Because the standard deviations were not equal, $z_B(A)$ was not necessarily equal to $z_A(B)$, so these analyses were performed in each direction separately. Figure 3.1 panel B shows a histogram of the z_{after} (before) and z_{before} (after) values for the population of cells. The histograms are shown for both all cells (blue) and significantly encoding cells (as determined by KS test, red). It is clear that the cells that fire significantly differently on this task have higher z-scored values, but it also appears as though all cells tend to form a gaussian-like distribution of z-scored firing rates centered at zero. To compare

this distribution to what we would expect to record from randomly firing cells, we carried out the same z-score analysis on our data but with the parametric identity of each lap shuffled (in this case, randomly assigning laps to either before or after the switch). We compiled a distribution of z-scores based on 1000 such random shuffles, then fit a gaussian model to it, and plotted that distribution on top of the histograms in Figure 3.1 (black line). The actual distribution is clearly much broader than would be expected by chance, indicating that the population of cells does represent this before vs. after the switch differences on this task.

We mentioned above that we found 163 cells (49.3%) that significantly encoded for a strategic change on this task. In order to verify that this population of cells was above what could be expected from chance firing alone, we created a distribution of random sessions by shuffling the inter-spike-intervals (ISIs) of all cells recorded 500 times, and ran the same KS-test for significance on the firing of all cells in each of these random sessions. This procedure gave us an expected sampling distribution for the number of significant cells we would identify from a session with completely random firing. This sampling distribution had a mean and standard deviation of 18.6 ± 4.5 cells. We also calculated a sampling distribution for random firing before and after the switch by considering the actual firing pattern of all cells, but re-labeling each lap arbitrarily as occurring before or after the switch lap (with the same proportion as the original data for each session). We referred to this distribution as the identity control (ID), and it had a mean and standard deviation of 15.5 ± 4.8 cells. Thus the actual population of cells had significantly more cells that encoded for strategic changes than we would expect by chance. However, there is one more important control we need to consider to ensure that what we are seeing is indeed strategy encoding.

Some of our recorded cells show evidence of a continuous change in firing rate over the course of the session. These cells will likely be included in our count of strategy encoding cells because this pattern of firing can produce a significantly different firing rate before vs. after the switch. In order to control for this possibility, we ran a robust regression of firing rate against lap number for all of our

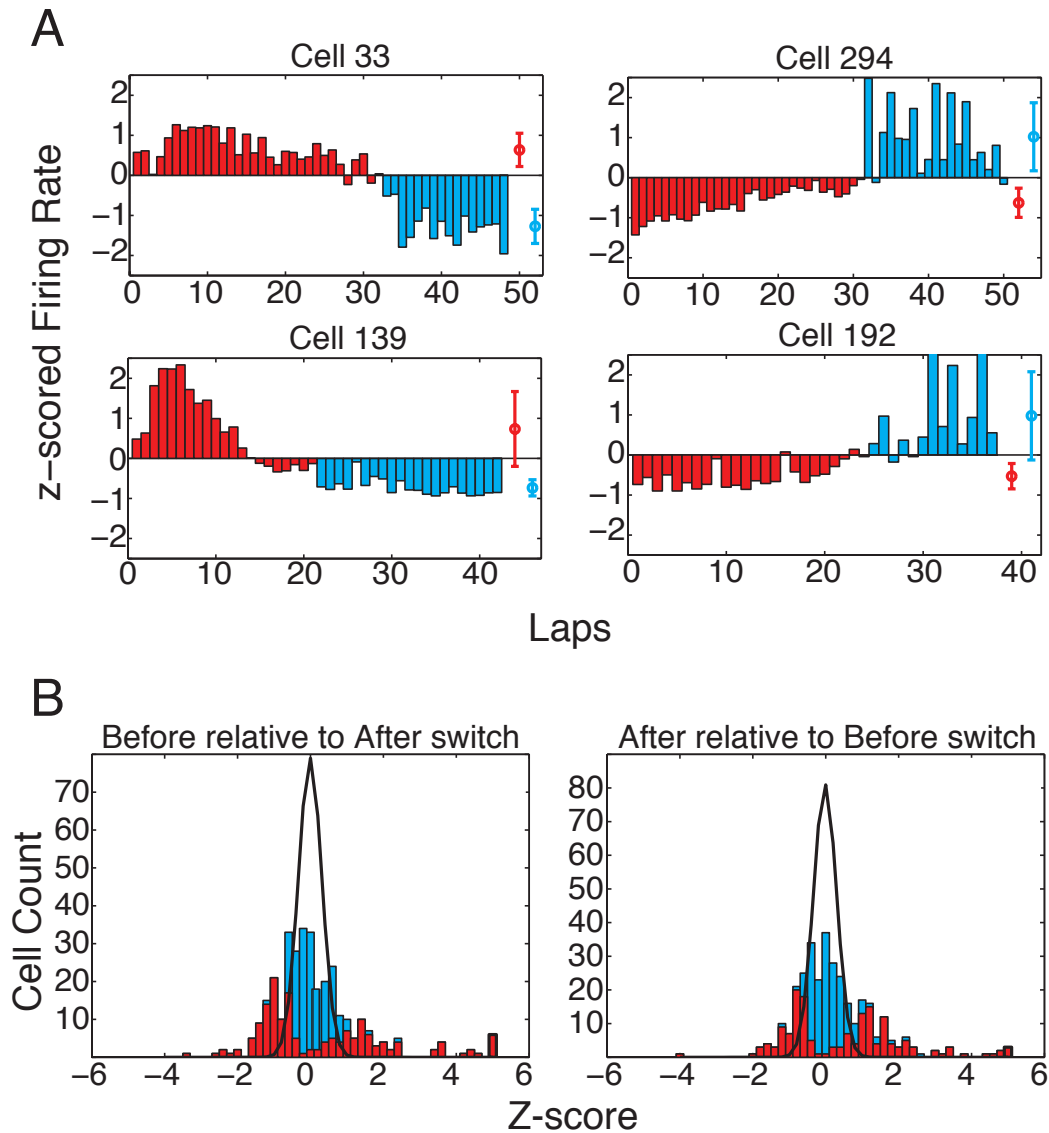


Figure 3.1: (A) Examples of cells with strong changes in firing rate on laps before and after the switch. Red bars indicate laps before the switch, blue bars laps after the switch. Red and blue circles represent average firing rates over both types of laps. (B) z_{after} (before) and z_{before} (after) comparisons. Bars indicate overall distribution, with red bars indicating cells found to have a significant firing difference by KS test. Black line indicates expected distribution from ID shuffle. Reproduced with permission from Powell and Redish (2014)

cells, and removed any from our population of strategy control cells that had a significant effect. This process left us with 109 cells encoding strategy changes, which is still significantly higher than the numbers of cells we found in our sampling distributions. Thus, we can conclude that this population of cells encodes the difference between laps before and after the switch on this task (our current proxy for strategy changes).

While it is important to establish that some cells that fire differently before vs. after the switch not merely as a result of continuously varying their firing rate over time, for the moment our purpose is simply to identify populations of cells with different parametric response patterns to compare to each other. Therefore we will consider the full population of 163 cells to differentially encode whether a lap occurs before or after the switch.

3.1.2 Navigational decision encoding in PFC

The most basic decision the animal has to make on any lap of the MT-LRA task is whether to go to the right or the left side of the track. If PFC represents parameters that may be of strategic value to solving the task, it is reasonable to believe that it might represent this distinction as well. So we compared the firing rates of cells on laps made to the left side of the track to laps made to the right, and again we found a significant population of cells (98 cells, 29.7%) that significantly changed their firing in accordance with this parameter. This was significantly more cells than we found in either of our shuffled sampling distributions (15.9 ± 3.9 for the ISI control and 15.5 ± 4.9 for the ID control). Example cells that separate firing on laps to the left side of the track vs. the right side of the track are shown in Figure 3.2, along with histograms of the $z_{\text{right}}(\text{left})$ and $z_{\text{left}}(\text{right})$ z-scores calculated as described above.

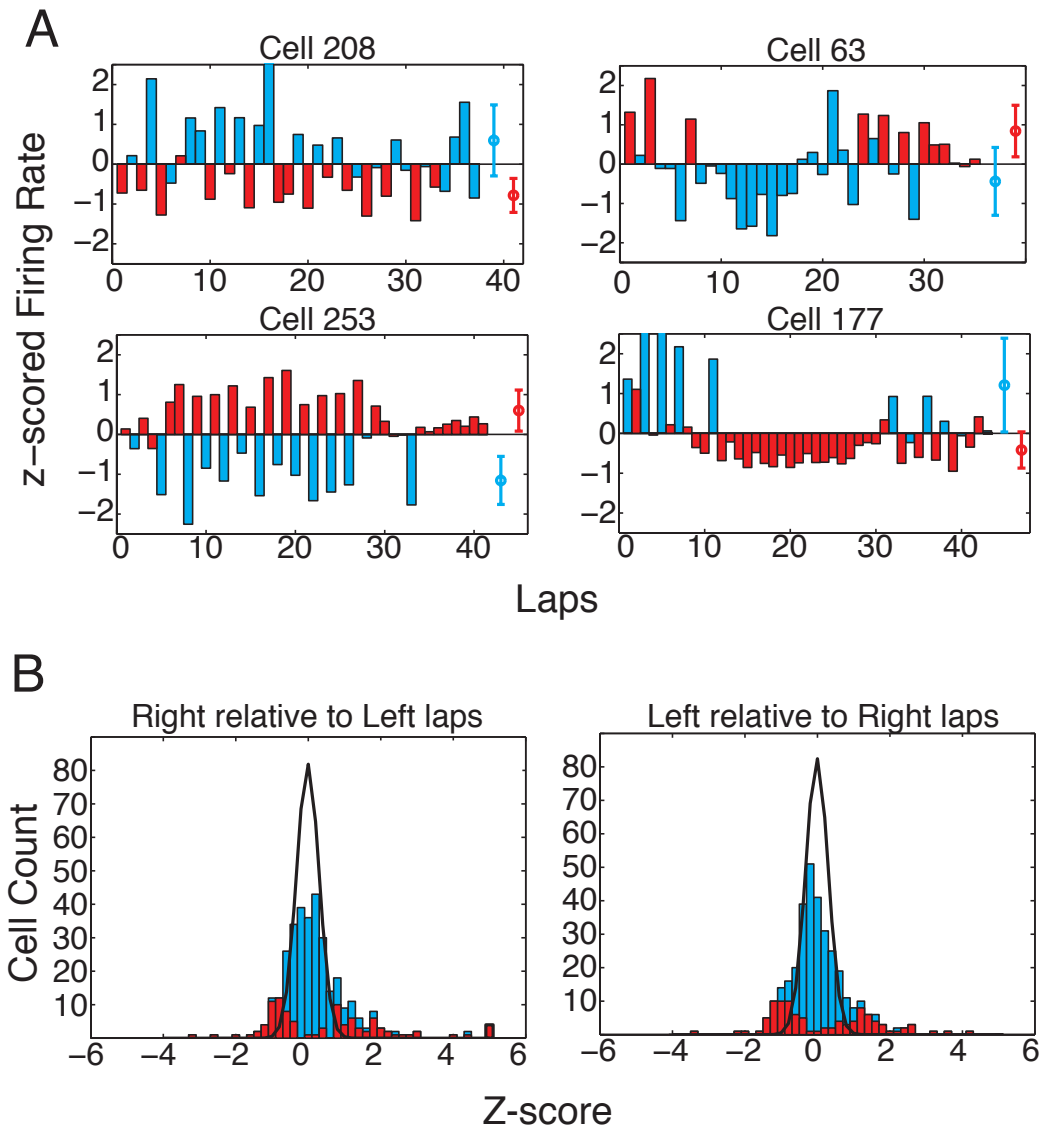


Figure 3.2: (A) Examples of cells with strong changes in firing rate on laps to the left vs. to the right of the track. Red bars indicate laps run to the right, blue bars laps run to the left. Red and blue circles represent average firing rates over both types of laps. (B) z_{right} (left) and z_{left} (right) comparisons. Bars indicate overall distribution, with red bars indicating cells found to have a significant firing difference by KS test. Black line indicates expected distribution from ID shuffle. Reproduced with permission from Powell and Redish (2014)

3.1.3 Error encoding in PFC

Several groups, most notably Laubach and colleagues (Narayanan and Laubach, 2008, 2009) have shown that the PFC is sensitive to errors that animals make running behavioral tasks. This distinction is not surprising given that cells in PFC have been shown to be responsive to reward receipt (Pratt and Mizumori, 2001b; Horst and Laubach, 2013; Hok et al., 2005), and if the region is involved with the encoding and processing of strategies, there ought to be some representation of whether these strategies were correct or erroneous. We found 118 cells (35.8%) with significantly different firing rates on correct laps compared to laps which resulted in errors, which again was significantly larger than the proportion of such cells found in the ISI shuffle control (19.1 ± 4.7) and the ID control (15.7 ± 4.8). Examples of cells with differential firing rates on correct vs. error laps are shown in Figure 3.3, again along with the $z_{\text{error}}(\text{correct})$ histogram. It is worth noting that while we found several cells with significant differences in firing rate on correct vs. error laps, our definition of a lap (for the MT-LRA task) included the feeder region and the return rail following the feeder as part of each lap. Thus cells that fire differentially in response to errors may be firing due to differences following rewards and not predicting rewards. Our current purpose is merely to identify parameters of the task that the cells respond to so that we can examine the representations of these different patterns on the same task, so for the moment it does not matter whether the cells respond to errors retrospectively or prospectively, but it should be noted that we make no claims that these cells are predicting errors on the MT task. We will consider this distinction in more detail in Chapter 4.

3.1.4 Overlapping populations of task representation in PFC

We have found that significant sub-populations of the same group of cells represent three different parameters on the same task. As a reminder, from a group of 330

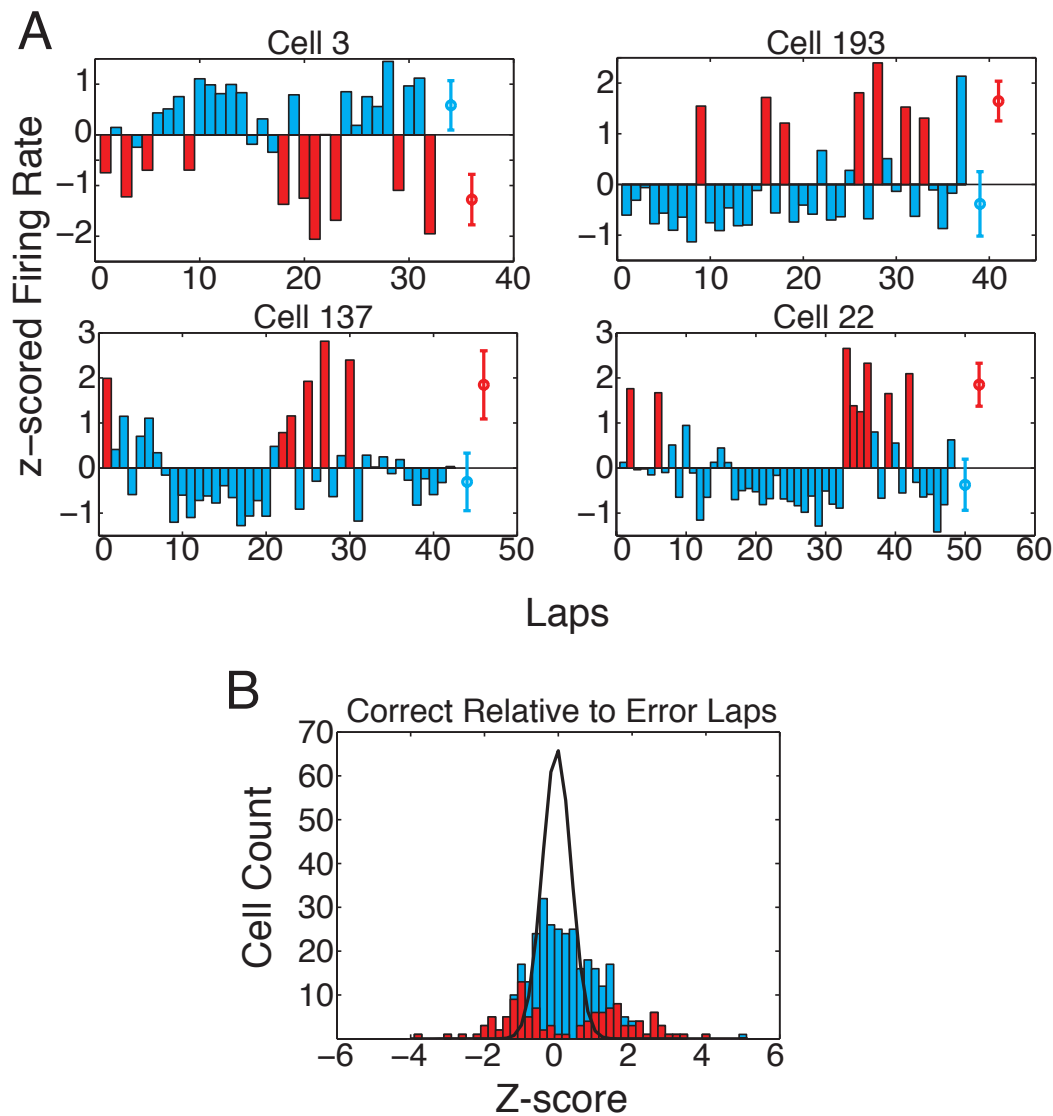


Figure 3.3: (A) Examples of cells with strong changes in firing rate on correct laps vs. laps on which the animal made an error. Red bars indicate error laps, blue bars correct laps. Red and blue circles represent average firing rates over both types of laps. (B) $z_{\text{error}}(\text{correct})$ comparison. Because there were too few error laps, we only calculated $z_{\text{Error}}(\text{Correct})$. Bars indicate overall distribution, with red bars indicating cells found to have a significant firing difference by KS test. Black line indicates expected distribution from ID shuffle. Reproduced with permission from Powell and Redish (2014)

cells we found 163 representing whether a lap occurred before or after the switch (B/A cells), 98 cells representing whether a lap was to the left or the right side of the track (L/R cells) and 118 cells representing whether a lap was correct or an error (C/E cells). From the number of cells encoded alone, it is clear that these populations of cells must intersect. By comparing the populations of cells labeled according to their individual firing rate tests, we found that there were 74 B/A only cells, 23 L/R only cells, 42 C/E only cells, 28 B/A and L/R cells, 15 L/R and C/E cells, 29 B/A and C/E cells, and 32 cells that responded differentially to all three parameters. These different populations are displayed as a Venn diagram in Figure 3.4, along with the population overlap we would expect from a set of three random independent processes sampling from the same population with hit rates determined by the overall percentage of cells we found for each parameter type, akin to flipping three different biased coins for each cell to determine whether or not it responds to each parameter. In many cases, the sets of population overlap numbers are quite similar, with the most obvious exception that there are more cells responsive to all three task parameters in the actual data than would be expected by random overlap and fewer cells responsive to any set of pairs of cells than would be expected.

In order to verify whether the overlapping population of cells we detected was credibly different from what we would expect from independent processes, we ran a Bayesian model comparison on our data. We fit overall rates of occurrence to the individual categorization of each cell as B/A, L/R, or C/E, then compared a model where the rates for each combination of cell types (including responsive to all and responsive to none) was a product of these base rates of occurrence to a model where each combination had its own independent rate of occurrence. Bayesian model comparison is an ideal technique for this comparison because the probability mass is normalized over the entire space, which automatically corrects the comparison for the larger number of parameters in model 2. Sampling over the resultant probability space with a Markov Chain Monte Carlo process using JAGS(Plummer, 2003; Kruschke, 2010) revealed strong evidence for the first

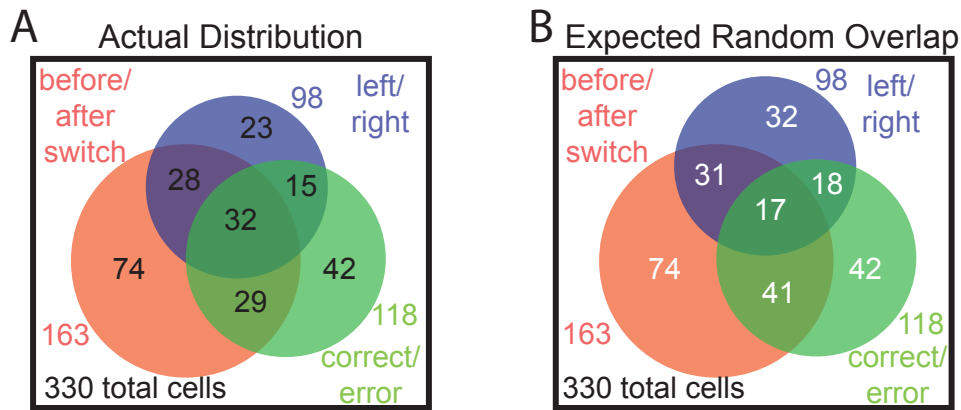


Figure 3.4: Distributions of multiple responses. (A) The actual distribution of cells responding to the three key binary factors and the overlap of cells that represent multiple factors. (B) Expected distribution overlap if the factors combined independently (i.e. $163/330 = 49\%$ before/after switch responding, $118/330 = 36\%$ correct/error responding, $98/330 = 30\%$ left/right responding). Reproduced with permission from Powell and Redish (2014)

model (4-20 times more likely than the second model over 3 independent comparisons), indicating that the rates at which cells are simultaneously responsive to multiple task parameters is consistent with the overlap expected from an independent combination of cell populations responsive to each individual parameter. In other words, the cells are distributed as though a certain percentage of cells will be responsive to each parameter, regardless of whether that makes individual cells responsive to one, two, or even three parameters. However, we will now show that cell responsiveness is not even limited to these three parameters.

3.2 Feeder Fire and Reward Responses in PFC

Previous studies (Pratt and Mizumori, 2001b; Hok et al., 2005; Horst and Laubach, 2013) have found evidence for reward-related responses in rat PFC, so we decided to examine whether our cells had differential firing rates at the feeder locations.

We measured the firing of cells from the time when the animal entered the “trigger zone” (the point at which the feeder would click if the lap was correct) until 2.5 seconds after this time at each feeder, and checked to see if firing during this period was significantly different from firing at all other locations on the maze for each of the 330 cells we recorded. We discovered that 298 cells (90.3%) responded differently to at least one feeder location, and 126 (38.2%) responded differently at all feeder locations. Raster plots of the firing responses of some example cells are shown in Figure 3.5, demonstrating some of the reward-response patterns observed. Pratt and Mizumori (2001b) found only about 31% of their recorded cells had a reward response, whereas almost all of our cells (90.3%) demonstrated a significant reward response. Additionally, there is some evidence (Horst and Laubach, 2013) that reward responses in PFC are driven by consumptive behaviors, specifically motor patterns (a corollary to the Euston-Cowen Hassle). Unfortunately, we do not have precise measurements of when our animals began consumptive behaviors, so it is difficult for us to definitively say that the firing we see is not driven by consumptive behaviors. However, considering the example cells in Figure 3.5, both A and B clearly show distinct firing changes that are aligned to the click of the feeder, which indicates that the animal is about to receive reward prior to his actual receipt of reward. Therefore it is unlikely that these cells are driven merely by consumptive behaviors.

We also examined the specific pattern of feeder responses among cells. As mentioned above, the largest set of cells responded to all feeder locations (126 cells), but beyond this there was no discernible pattern to the groups of feeders at which cells responded. We found groups of between 2 and 24 cells that responded to all other possible patterns of one, two, or three feeders, with the largest group (24 cells) responding to both right side feeders and left feeder 2, and the smallest group (2 cells) responding to left feeder 1 and right feeder 2. There were in general as many cells responding to seemingly non-sensical groupings of three feeders (14, 18, 20, and 24 cells) as there were responding to feeders on only the left (18) or right (19) sides of the track. There also did not seem to be notably more cells

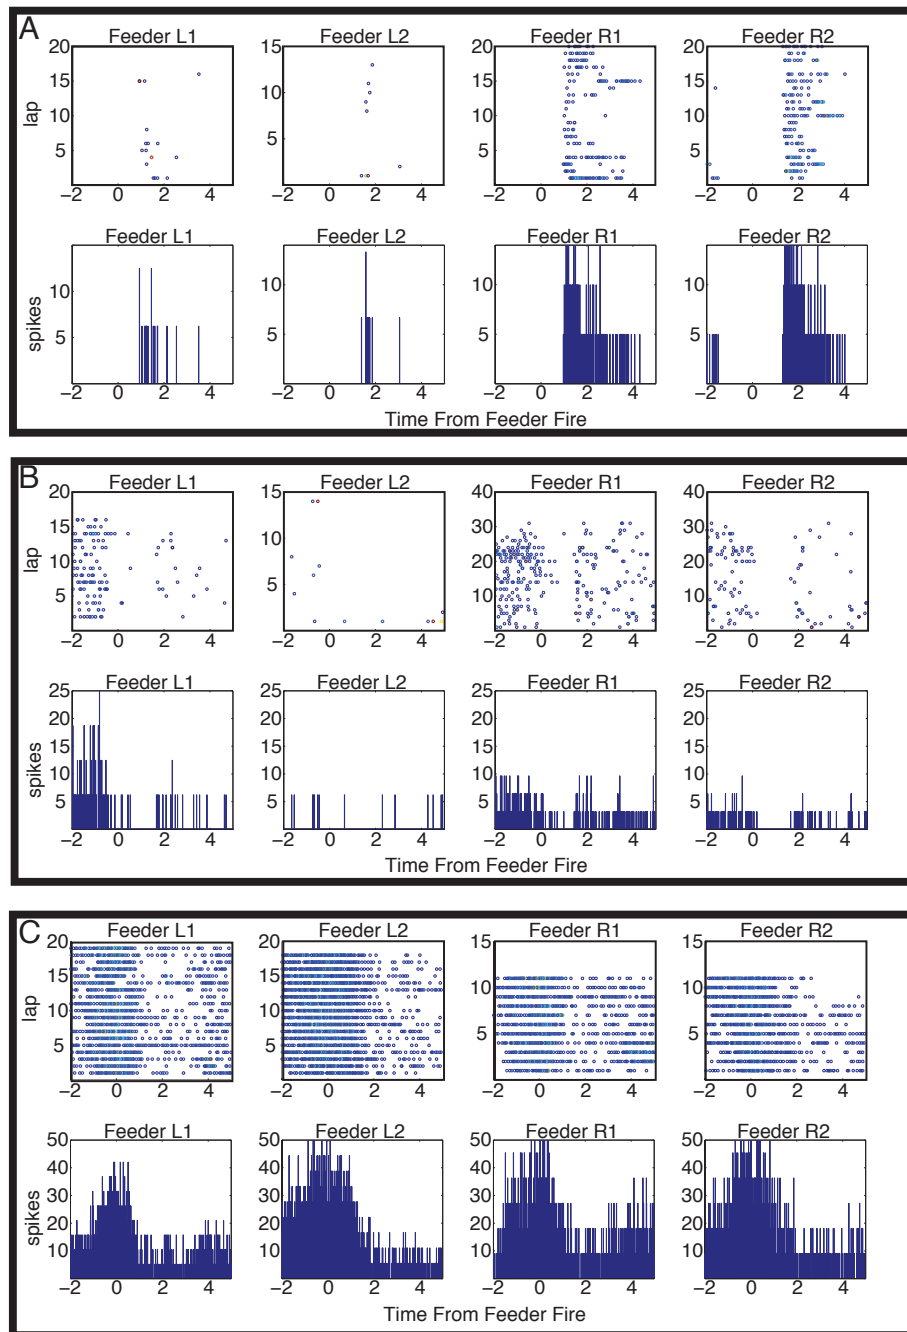


Figure 3.5: Prelimbic cells show reward responses. Three example cells are shown in (A), (B), and (C). For each example, (top) Raster plot of cell firing by lap aligned to feeder fire time for all four feeder locations (Left Feeders 1 and 2, Right Feeders 1 and 2). (bottom) Peri-Event Time Histograms (PETH) aligned to feeder fire for all feeder locations.

responding to the feeders with the same food flavors (left feeder 2 and right feeder 2, 6 cells) compared with other combinations (particularly left feeder 1 and right feeder 1, which were different flavors, 10 cells). This seemingly random set of feeder-fire associations paralleled many other findings suggestive of randomness in the firing properties of PFC neurons (including the tendency of population z-score distributions to resemble normally distributed cells with just the outliers of the distribution (both positive and negative) representing firing differences, and the fact that overlap of cells responsive to multiple parameters is consistent with the overlap of independent stochastic processes). The abundance of seemingly randomly ordered response properties led us to search for evidence that the firing properties of cells were not just randomly assigned on a day-by-day basis to fit some criteria for task performance.

3.3 Consistency of cells across days in PFC

In order to determine whether cells were randomly assigned or showed a measure of consistency in their firing patterns to various task parameters, we examined the firing properties of cells recorded across multiple days (and thus multiple sessions) while the animal was running the same task. We identified cells consistently recorded across days by combining waveform comparison methods used previously (Schmitzer-Torbert and Redish, 2004; Tolias et al., 2007) to classify units recorded on consecutive days as either matched (the same cell recorded on both days) or unmatched. Cells matched across multiple pairs of consecutive days were considered to be matched across all those days (for example if cell 5 on day one matched cell 17 on Day two and cell 17 on day two matched cell 35 on day three, then these three cells were considered to be the same cell). By this method we identified 145 cells recorded on just one day, and 60 cells recorded across more than one day, of which 32 were recorded across two days, 11 were recorded across three days, 5 were recorded across four days, 4 were recorded across five days, and 8 were recorded across all six switch days.

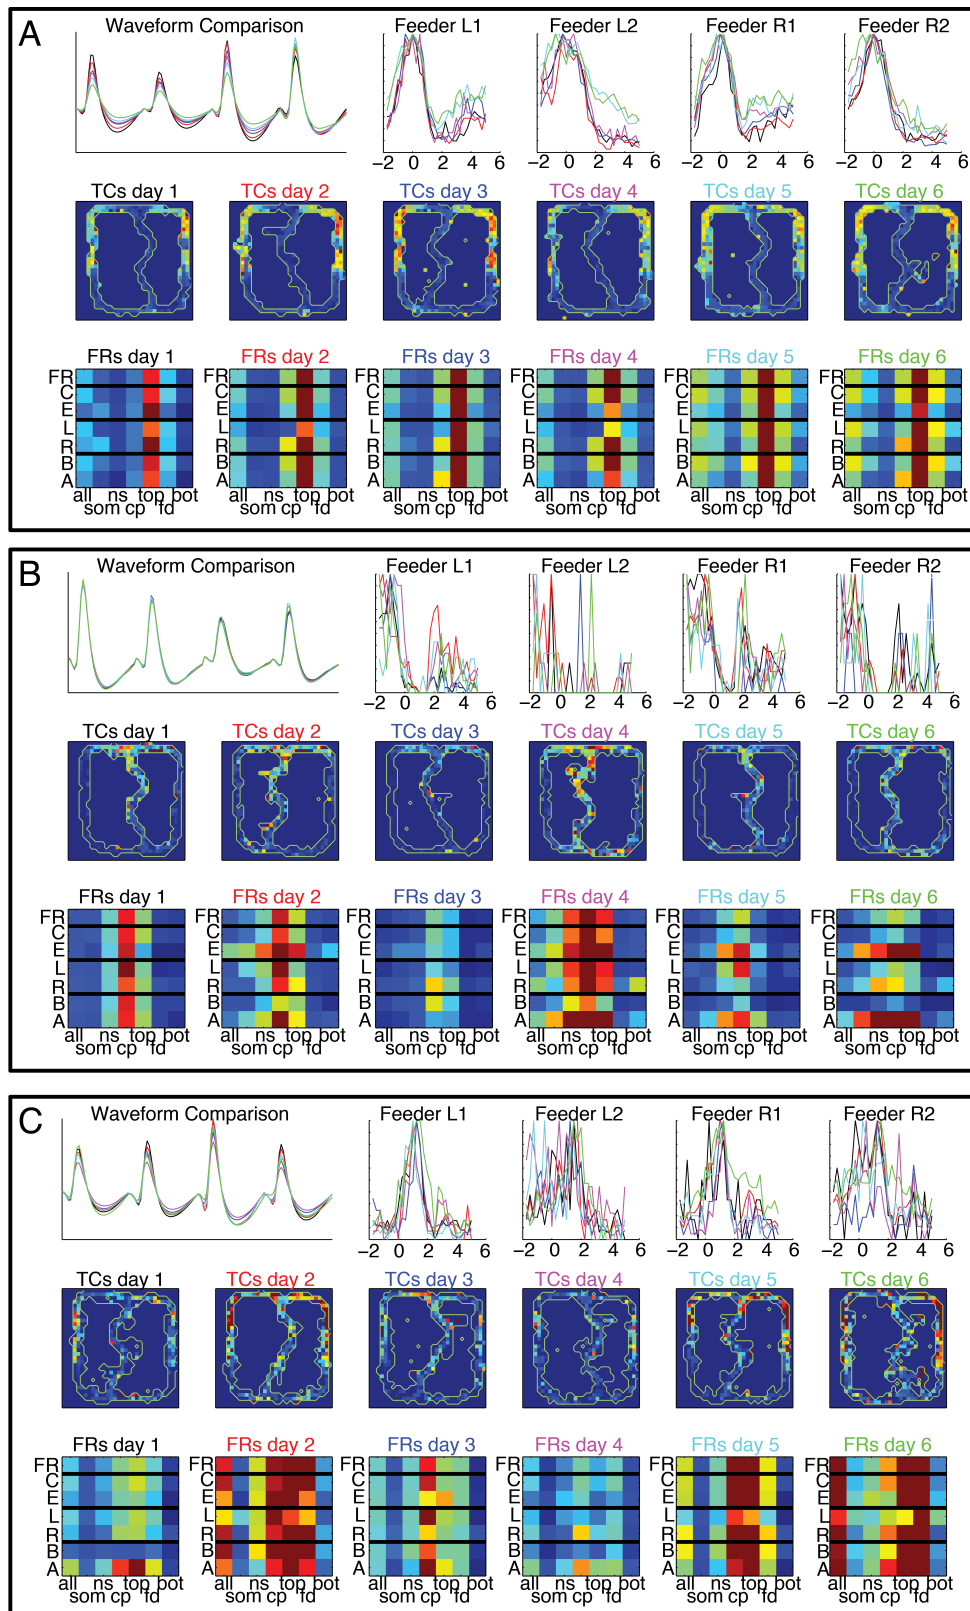


Figure 3.6: Examples cells with matched waveforms across all six days. For each cell we show: (Top Left) comparison of waveforms across all 6 days; (Top Right) PETH plots of the average firing rate for each day at all four feeder locations; (Middle) Spatial correlation maps for each day; (Bottom) 5.7.5.1

Cells recorded across multiple days in PFC display very stable firing characteristics on the MT-LRA task. In order to demonstrate this fact, we considered both the spatial and task-specific firing properties of these cells (for a more detailed description of the spatial firing of PFC cells, see Chapter 4). We divided the MT-LRA task into seven separate zones based on spatial regions of the track (including firing on the entire track as the 7th zone, see Figure 2.2). We then found the firing rate of each cell in each spatial zone for a specific type of laps (all laps to the left, all laps to the right, all laps before the switch, all laps after the switch, all correct laps, all error laps, and all laps overall) for each session. Figure 3.6 shows examples of three cells recorded across six consecutive days, with the cell's waveforms, a color coded grid representing the firing rate in each of these spatial and task-parameter-based bins, and spatial tuning curves from each day for comparison. The spatial firing patterns are highly conserved from one day to the next, and the spatial-parametric grids show a high degree of consistency as well. In order to classify this consistency across the population of cells, we correlated the spatial-parametric firing grids for all pairs of cells, then compared a histogram of the correlation coefficients for all matched pairs of cells to a histogram of the correlation coefficients for all unmatched pairs of cells. These two histograms are shown in Figure 3.7, and clearly the matched cells display a strong tendency towards positive correlations (median (Matched) = $.67 \pm 0.23(\text{SD})$) while the unmatched cells show no such tendency (median(Unmatched) = $0.01 \pm 0.36(\text{SD})$). So, even though the firing patterns of cells in PFC are quite varied and show a lot of randomness in their response patterns, they are quite consistent across days, implying that individual cells are quite consistent in their roles and not simply assigned response properties on a day by day basis.

3.4 Mixed selectivity in rodent PFC

The tendency of cells in PFC to represent multiple parameters of a task simultaneously has been described as “Mixed Selectivity”. In particular, Rigotti et al.

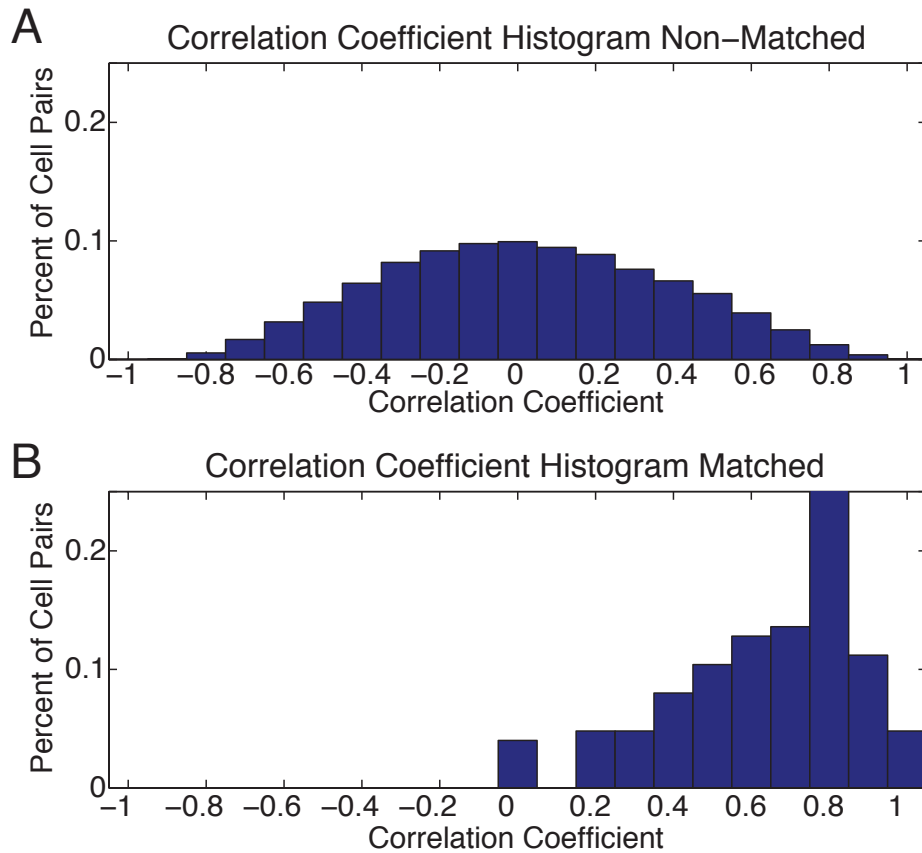


Figure 3.7: Histogram of across-day correlations. (A) Histogram of the correlation coefficient of all pairs of the 7x7 grids of FR seen in Figure 3.6 for all pairs of cells not identified as being the same cell recorded across multiple days (Non-Matched). (B) The same histogram as in part (A) but for all pairs identified as being the same cell on consecutive days (Matched). Reproduced with permission from Powell and Redish (2014)

(2013) described mixed selective cells as having “complex and diverse response properties that are not organized anatomically, and that simultaneously reflect different parameters” and observed that “the predominance of these mixed selectivity neurons seem to be a hallmark of PFC and other brain structures involved in cognition.” Multiple publications have demonstrated the presence of mixed selectivity neurons in the primate, but they have never been explicitly shown in the rodent PFC. However, if indeed the property is a hallmark of cognition and PFC function, we should expect to find evidence of non-linear mixed selectivity in our PFC cells on the MT-LRA task.

In order to check for mixed selectivity, we modified the approach described by Rigotti et al. (2013), who employed two different models of neuronal firing: one for simple linear selectivity to each of three different task parameters and one for mixed selectivity, in which case cells have significantly varying firing rates to specific combinations of the three task related parameters. We modified these models to fit our MT-LRA task using the task parameters we identified above (Before vs. After the switch, Left vs. Right side of the track, and Correct vs. Error laps). The resulting linear selectivity equation is:

$$FR = \beta_0 + \sum_{i=B,A} \beta_i [B/A = i] + \sum_{j=C,E} \beta_j [C/E = j] + \sum_{k=L,R} \beta_k [L/R = k] \quad (3.2)$$

To incorporate mixed selectivity, the authors added terms to the linear selectivity equation that represented firing in each specific session type, which in our case corresponds to individual lap types. Each lap type is defined as a specific combination of the other three parameters, so the resulting non-linear mixed selectivity equation is:

$$FR = \beta_0 + \sum_{i=B,A} \beta_i [B/A = i] + \sum_{j=C,E} \beta_j [C/E = j] + \sum_{k=L,R} \beta_k [L/R = k] \\ + \sum_{(i,j,k)=(B,C,L),(A,C,L)\dots(A,E,R)} \beta_{i,j,k} [B/A = i] \cdot [C/E = j] \cdot [L/R = k] \quad (3.3)$$

To identify non-linear mixed selective cells, we used the above equations to specify two different generalized linear models (GLMs) we could fit to each cell, and then

used a Bayesian model comparison to identify which model best described each neuron. In order to ensure that we had enough examples of the mixtures of parameters for each neuron, we used only cells that had been recorded across multiple days (since we verified in section 3.2 that these cells have consistent firing to task parameters from day to day). The model comparison was evaluated using MCMC methods in JAGS (Plummer, 2003; Kruschke, 2010), and neurons that were better described by the non-linear model (model 2) were classified as non-linear mixed selective. Examples of the fit parameters of typical linear and mixed selective neurons are shown in Figure 3.8. Of 60 neurons recorded across multiple days that we were able to classify, 13 (21.6%) were better described by the non-linear mixed selective model, which is in good agreement with the percentage found by Rigotti et. al. (47 of 237, 19.8%). This result demonstrates that mixed selective neurons are present in the rodent PL, in nearly the same proportions as they are found in Primate PFC, which bolsters the interpretation of the rodent PL as an analog of the primate PFC. Simultaneously it supports the interpretation that cognitive processes in the cortex employ similar hallmarks and mechanisms between the two species.



Figure 3.8: JAGS Mixed Selectivity Model Fit Examples. Examples of the fit parameter distributions for three neurons fit to the linear and non-linear mixed selective models modified from Rigotti et al. (2013). Cell one is better described by the nonlinear model and displays only plausibly non-zero coefficients for the cross terms. Cell two displays only plausibly non-zero coefficients for the linear terms, and is better fit by the linear model, and cell 25 displays plausibly non-zero coefficients for both linear and nonlinear terms, and is better described by the nonlinear model.

Chapter 4

Spatial Representations in Rodent PFC

The role of spatial firing patterns in the rodent PFC on spatial tasks remains unsettled. In other species, spatial representations in PFC are rarely discussed, though a significant reason for this absence may be due to the types of studies that are carried out in different species. Rodents tend to dominate many discussions of neural correlates of space, due to the prominence of place cells in hippocampus. In general, most primate and human studies are performed on subjects confined to head fixed or at least non-navigating task paradigms, making spatial variation minimal. However, in rodents it is very common to utilize spatial tasks when recording from PFC. Consistently these studies have found cells with non-uniform spatial firing patterns. However, the issue of whether these cells actually represent space *per se* or merely other task-based parameters that correlate with space is still under debate. Here we present evidence that non-uniform spatial firing patterns are present in the rodent PFC on spatial tasks and an accurate representation of space can be decoded from the firing of these cells. However, we believe that these firing patterns are still better described by other aspects of the task than by purely spatial representations. However, we will outline evidence that the spatial non-uniformity of PFC cell firing is an important and useful consideration in attempts

to decode task-based and cognitive representations from rat PFC.

4.1 The significance of spatial firing in PFC

Even the earliest reports on the nature of spatial representations in rodent PFC contrasted firing patterns with those of place cells recorded in hippocampus (HPC). Since PFC receives innervation from the HPC (directly in the PL/IL from ventral CA1 (Hoover and Vertes, 2007)), it is plausible that spatial locations could be represented in PFC as well. However, while hippocampal place cells tend to show very specific firing to localized areas of an environment and only fire at those locations, PFC spatial firing patterns tend to spread out over a larger area of the environment. Jung et al. (1998) assessed spatial firing patterns of rodent PFC cells on both an 8-arm radial maze and a simple figure-8 track (similar to the tracks used in our MT-LRA and DD tasks, see Chapter 2). They observed that on the radial maze most prefrontal cells tended to fire in relatively the same location on each arm of the maze (i.e. the firing was radially symmetric). While firing patterns were less consistently symmetric on the figure-8 track, it was also common to see firing patterns mirrored on the left and right sides of the track. Around the same time, Poucet (1997) reported on the spatial patterns recorded from PFC cells while rodents performed a task in an open environment with a circular wall. In this environment, no spatially selective PFC cells were reported after correcting for task behaviors that tended to occur in certain locations within the environment. Based on these and similar findings, the consensus of the field appears to be that in general firing in the PFC is not representative of space, but rather certain cognitive/behavioral aspects of task performance that tend to occur at certain locations in a spatial task. In the open field environments, the lack of a consistent alignment of cognitive/behavioral tasks with spatial locations accounts for the absence of clear spatial correlates of PFC cell firing. In further support of this interpretation, Hok et al. (2005) demonstrated that they found spatial firing patterns in PFC cells in an open field environment in which food

pellet rewards were consistently released when animals went to a specific but unmarked trigger location. Additionally, pellets released by the apparatus tended to land in a specific zone (though after landing they could roll to any location in the environment, making the landing zone somewhat independent of reward receipt). In this task, many PFC cells displayed place field-like firing patterns that overlapped the trigger zone and landing zone locations, likely because these zones acquired spatial significance for the task.

Our own MT-LRA task is similar to the figure-8 maze on which Jung et al. (1998) observed non-uniform firing, and we observed similar cellular responses. Examples are shown in Figure 4.1 panel A, along with examples of place cells recorded in HPC of a different set of rats who performed this same task in our lab (panel B)(van der Meer and Redish, 2011). There is a clear difference in both the extent of the spatial tuning curves and their tendency to be mirrored on both the left and right sides of the track for the prefrontal vs. the hippocampal cells. In addition, we employed a Bayesian decoding method (Zhang et al., 1998) to determine the most likely posterior position for the animal based on the firing of all cells recorded for our three rats on the MT-LRA task and a matched set of rats with HPC recordings on this task. For the decoding, we linearized the left and right paths through the track from SoM back to SoM, and normalized the length between landmarks (SoM, CP, Feeder 1, Feeder 2, and SoM again) for each session. The linearized paths were then divided into 25 spatial bins, and laps of each session were divided into training and test sets. We used a uniform prior distribution over space for all spatial decoding analyses. A decoder trained from a randomly assigned set of training laps was then used to decode the posterior probability of the animal being in any of the spatial bins for each of the 500 ms time bins in the test set, which were then compiled by spatial bin to produce a confusion matrix for decoded position by actual position. This process was repeated 100 times with different lap assignments to training and test sets and averaged to produce combined confusion matrixes. The left and right paths were very similar, so they were averaged to produce the plots seen in Figure 4.1 panel C for PFC and

panel D for HPC. From the figure we can clearly see that hippocampal decoding is more accurate, although decoding of space based on PFC cell firing also tended to produce a respectably accurate prediction of the rats position. However, one factor that stands out in particular from the PFC decoding is that the task tends to produce regions or chunks of space that decode similarly, for example the region prior to the Choice Point (CP on the figure), the region between CP and the first feeder (F1), and the region between the second feeder (F2) and the end of the lap. These regions display a heightened off-diagonal decoding, indicating that although there is still a heightened decoding at the actual position, the overall representation in cortex is more similar for all locations in this region than it is for locations outside of this region. This phenomenon suggests that the PFC may be essentially chunking the maze into specific processing zones. These zones certainly could represent the individual cognitive/behavioral regions hypothesized to account for the non-uniform spatial firing in prefrontal seen by many groups on spatially-based tasks. Additionally, both feeder regions (F1 and F2) are highly consistent with each other on the PFC confusion matrix, indicating that the PFC may represent these two regions very similarly. Given the common and specific firing patterns we discovered in PFC neurons on this task (see Figure 3.5) it is not surprising that these regions could be represented quite similarly based on the decoded firing pattern of the cells.

In Chapter 3 we examined the consistency of PFC representations across days and discovered that both parametric and spatial representations are consistent across days on the MT-LRA task. Examples of consistent spatial tuning curves for the same cell recorded across six days of performing the same task are shown in Figure 3.6. The fact that spatial firing patterns are constant from day to day does not necessarily support their being related to either space or cognitive functions correlated to space, but it does suggest a potential experiment to explore the question. We trained one rat in which we had stable recordings for several PFC cells to run two different behavioral tasks (our DD task and a variant of our MT-LRA task called LRA) each on two different spatial tracks, a large track and a small

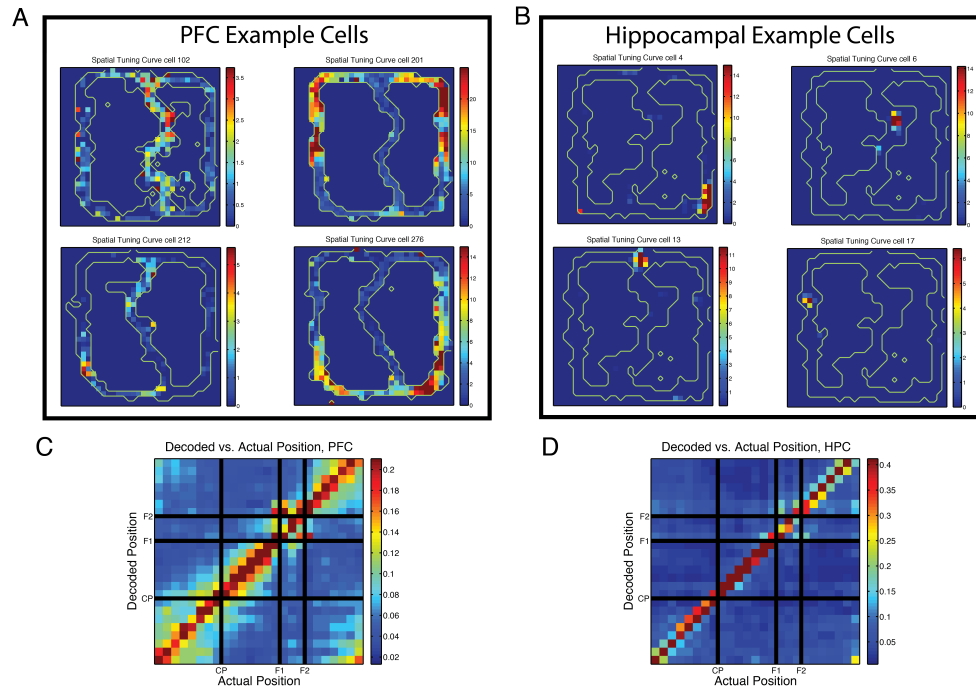


Figure 4.1: Prelimbic cells show non-uniform spatial firing. (A) Typical spatial tuning curves of prelimbic cells. Note the large firing response fields, typically covering a broad region of the maze and usually present on both left and right sides of the track. Color bars indicate firing rate in Hz. (B) Typical spatial tuning curves of hippocampal cells from the same task. Note the smaller firing response fields, more typical of hippocampal place cells. Color bars indicate firing rate in Hz. (C) Decoding confusion matrix generated from prelimbic cell firing, from 100 random assignments of laps to training and test sets averaged together, then averaged over all rats, all sessions, and left and right laps. Color bar indicates decoded probability at each location. (D) Decoding confusion matrix generated from hippocampal cell firing, from 100 random assignments of laps to training and test sets averaged together, then averaged over all rats, all sessions, and left and right laps. Color bar indicates decoded probability at each location. Reproduced with permission from Powell and Redish (2014)

track (the Multi-Track task, see Chapter 2). If the representations in PFC are coding for space, and not task related cognitive parameters, we would expect them to stay constant when the animal performs different behavioral tasks (presumably with different cognitive requirements) on the same spatial track. Conversely, if the cells represent simply cognitive parameters and not spatial information per se, we would expect them to show spatial firing patterns in roughly the same locations on both the large and the small track as long as they were running the same task. We recorded 12 consecutive days from the same animal and varied the track size, task type, and reward contingency for the task pseudo-randomly each day. We then examined tuning curves from each cell we could reliably record over all 12 days to determine whether they were consistent over spatial track types, or consistent over cognitive task types. Example tuning curves are shown in figure 4.2.

The results did not match either of our hypotheses. In all of the cells that had consistent spatial representations, the spatial tuning curves appeared to be consistent on both task types and, accounting for the difference in size, both track types. Interpreting this result is a little bit difficult. The fact that spatial regions activated by cells are consistent across the two different tracks for the same task implies that the spatial tuning curves are not due to spatial location. The consistent mapping of regions of the tracks across different physical tracks is inconsistent with place cell encoding, which tends to re-map significantly with even subtle changes in environment (Barnes et al., 1997; Fenton et al., 2008). Additionally, even if re-mapping did not occur, we would not expect to see such similar activations because the two tracks, while placed in the same room, were significantly different in size, so they occupied different locations in space. However, the fact that cells have similar spatial tuning curves on different task types implies that the representations are not strongly coupled to the specific task either. Ultimately, the most succinct explanation of the data seems to be that the two tasks are sufficiently similar in this case that they provoke comparable task-related spatial firing. In both tasks the feeders were at the same locations, the decisions about which direction the animal had to turn in order to receive

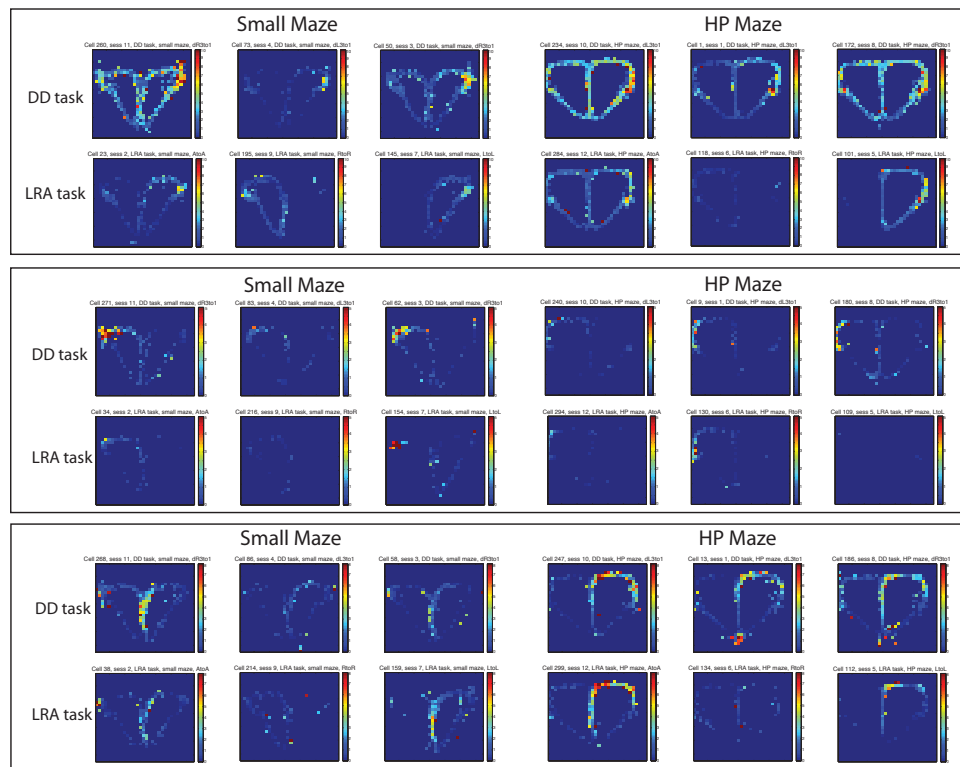


Figure 4.2: Spatial firing patterns are conserved across days different sized tracks and tasks. Tuning curves across all 12 days of the multi-track task are shown for three different example cells, divided up by which maze and which task was being run on each day. Firing rate scales are the same across days for each cell, but differ between examples.

reward occurred at the same location, and the dynamics of running were virtually the same. If the spatial tuning curves are indicative of spatially related sub-task goals, it is quite possible that many of these sub-tasks regions overlap on both the tracks and task types used here. A more illustrative manipulation may involve rotating the tracks into different orientations (or even recording in different rooms), moving the feeder locations to change the points of maximum salience on the task, or even finding a way to run a cognitively similar task on a completely different spatial layout. Regardless, the most reasonable explanation for all the current data is that cells in the PFC do not explicitly represent space *per se*, but rather occasionally demonstrate reliable spatial firing patterns because of reliable spatially related cognitive demands inherent to the structure of the task.

4.2 The importance of spatial aspects of firing in PFC

Given the conclusion of the previous section that cells in rodent PFC do not represent spatial locations, it may seem reasonable to conclude that spatial tasks in rodent PFC are of limited utility. Indeed, in light of the evidence provided by Euston and McNaughton (2006) and Cowen and McNaughton (2007), the so called Euston-Cowen Hassle (see Chapter 1), spatial tasks might seem like an undesirable complication to PFC research. However, we believe that spatial tasks in the PFC still provide a number of useful features for analysis of the representational properties of these cells, primarily because the non-uniform firing over the track is not driven only by spatial information.

We have shown that the firing pattern of individual PFC neurons differs over space, but we can also demonstrate that the representation of various task related parameters in PFC varies over space. On our MT-LRA task, we divided the track into six spatial regions (Start of Maze (SoM), Navigation Sequence (NS), Choice Point (CP), Top Rail (Top), Feeders (Fed), and Bottom Rail (Bot), as shown in

Figure 2.2). We then re-calculated the number of cells that were sensitive to each of the task based parameters we considered in Chapter 3 only on cell firing that occurred in these spatial regions, and for comparison purposes we also calculated the number of cells we would expect to find from a random firing distribution by the two shuffling methods. We found markedly different spatial patterns for the percentage of cells that differentially fired to each of the three task parameters.

For the behavioral strategy correlate (firing on laps before vs. firing on laps after the switch lap) we found relatively even numbers of cells significant for firing at each maze location (with the exception of the feeders): SoM :(ACTUAL=75, ISI-control= 10.4 ± 3.3 , ID-control= 9.3 ± 3.2); NS :(ACTUAL=88, ISI-control= 12.3 ± 3.4 , ID-control= 11.5 ± 3.7); CP :(ACTUAL=70, ISI-control= 9.3 ± 3.0 , ID-control= 8.6 ± 3.0); Top rail :(ACTUAL=62, ISI-control= 9.7 ± 3.0 , ID-control= 8.5 ± 3.0); Feeders :(ACTUAL=148, ISI-control= 18.3 ± 4.4 , ID-control= 15.2 ± 4.5); Bottom rail :(ACTUAL=75, ISI-control= 14.6 ± 4.0 , ID-control= 10.7 ± 3.4). There were notably more cells with firing rate differences before vs. after the switch on the entire track and at the feeders (where the animals spent the majority of their time on the track), but aside of these regions the number of cells responsive to this parameter was quite even across the track (between 62-88 cells in all regions). This fact seems to argue that spatial variation of cell firing patterns is not great, but it is also in contrast to the other two parameters we considered. For cells sensitive to differences in the animals navigational choices (differential firing on the left vs. right sides of the track) we found a quite different pattern: SoM :(ACTUAL=36, ISI-control= 9.7 ± 3.0 , ID-control= 9.5 ± 3.3); NS :(ACTUAL=34, ISI-control= 11.7 ± 3.3 , ID-control= 11.4 ± 3.5); CP :(ACTUAL=51, ISI-control= 9.0 ± 3.0 , ID-control= 8.5 ± 3.1); Top rail :(ACTUAL=81, ISI-control= 10.0 ± 3.2 , ID-control= 8.3 ± 3.1); Feeders :(ACTUAL=100, ISI-control= 17.0 ± 4.5 , ID-control= 15.2 ± 4.7); Bottom rail :(ACTUAL=73, ISI-control= 20.0 ± 5.6 , ID-control= 10.2 ± 3.3). Here there is a clear difference in the percentage of the cells that encode firing differences across space,

with more cells encoding differences on the areas of the track that are more spatially separated on left vs. right laps. This distinction makes sense, but it underscores the point that differential firing patterns in cells are not always uniform over the track. The difference is even more dramatic for the number of cells that have differential firing rates on Correct vs. Error laps: SoM :(ACTUAL=18, ISI-control= 9.4 ± 2.9 , ID-control= 9.6 ± 3.1); NS :(ACTUAL=17, ISI-control= 11.3 ± 3.4 , ID-control= 11.7 ± 3.5); CP :(ACTUAL=15, ISI-control= 8.8 ± 2.9 , ID-control= 8.9 ± 3.0); Top rail :(ACTUAL=21, ISI-control= 9.4 ± 3.0 , ID-control= 8.7 ± 3.1); Feeders :(ACTUAL=131, ISI-control= 24.3 ± 5.6 , ID-control= 15.6 ± 4.5); Bottom rail :(ACTUAL=50, ISI-control= 12.1 ± 3.5 , ID-control= 10.6 ± 3.5). For cells sensitive to correct vs. error laps, we only found significant numbers of cells with firing rate differences at the feeder and bottom rail regions. This was expected, as the animals should not know that they have made an error until they fail to receive reward at the feeder locations, although they seem to then retain this knowledge through the bottom rail region that brings them to the start of the next lap. It is clear though that the firing of the cells that carries information about certain behavioral parameters is limited to certain regions of the track (the percentages of cells carrying this information by region of the track is summarized in Figure 4.3).

The previous example from MT-LRA demonstrated that populations of neurons respond to specific task parameters differently over space, but can we show that individual PFC neurons encode parameters differently at different locations in space? In fact we can, but to do so we have to consider our Delay Discounting task. This task is described in great detail in Chapter 2, but briefly: on the DD task rodents again must decide on every lap whether to proceed to the feeder zone on the right side of the track or the left side of the track. On one side (varied between sessions but consistent within a session) the animal receives one food pellet of reward after a one second delay. On the other side, he receives three pellets of reward after an adjusting delay. This delay value is initialized pseudo-randomly for each session, and then varies with each lap based on the

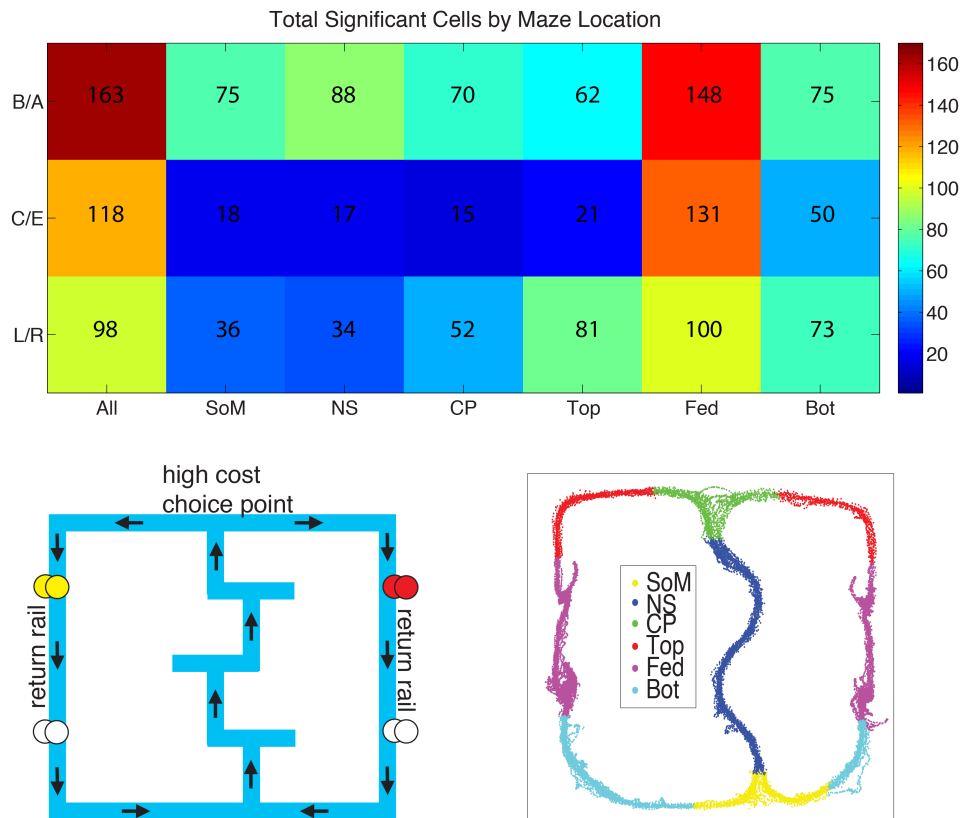


Figure 4.3: Parametric firing depends on maze location. Colors represent the number of cells with significantly different firing rates for each region of the MT track for each behavioral parameter (Before vs. After the switch, Correct vs. Error laps, and laps to the Left vs. Right side of the track). Maze and sections are depicted below for convenience.

rat's choices. If he chooses to go to the delayed side, this delay increases by one second on the subsequent lap; if he chooses to go to the non-delayed side, the delay decreases by one second on the subsequent lap. An important consequence of this task is that it is impossible to determine on an individual session basis whether a cell is encoding a preference for the left vs. right side of the track, or the delayed vs. non-delayed side of the track (because within an individual session, all laps to the left are also the delayed side, for example). However, since we have demonstrated that prefrontal neurons tend to encode the same task parameters on each day, we can examine the encoding of a particular cell recorded across several days, and see whether it more consistently fires to the spatial distinction (left vs. right) or to the delay distinction (Delayed vs. non-delayed). We examined this question for several cells recorded across days for all sub-regions of the track (divided similarly to the MT-LRA task, see Figure 2.4 for details). We found examples of cells that were consistently correlated to both possible side divisions (spatial or delay selective) but we also found several fascinating cells that were correlated to one parameter specifically at one location, and the other parameter at a different spatial location. Examples of each of these types of cells can be seen in Figure 4.4. The figure shows examples of two cells (90 and 116) that show a consistently elevated firing rate for one side of the maze (left) on most regions of the track prior to the choice point, but then show a consistently elevated firing rate for the delay side at the delay zone. A consistently different firing rate for the delay side of the track at the delay zone may not be surprising, and could be attributable to cells with delay-firing patterns, but there is also an example of a cell (100) that demonstrates elevated firing for laps to the delay side of the track at the stem region, which is far removed from the delay zone. Additionally, for much of the rest of the track, that cell showed elevated firing for the right side (a spatial distinction). Finally, another example cell (91) demonstrated a consistently elevated spatial firing for the left side of the track, until the last (feeder fire) zone, on which it showed elevated firing for the right side of the track. Cell 116 also showed a similar reversal of firing rate coding at the feeder region. The

presence of these types of cells demonstrates that at even at an individual cell level, firing selectivity can be related to the spatial region of the track on which the firing was recorded.

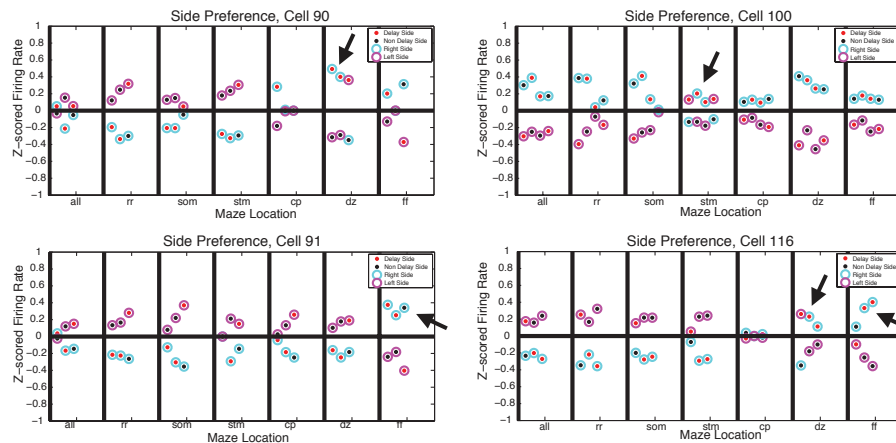


Figure 4.4: Parameters encoded on the DD task depend on maze location. Firing rate differences for each side of the maze are shown for four example cells recorded across multiple days on the DD task for each section of the maze. Z-scored firing rates for each day are represented by different markers. Marker internal color (red vs. black) represents delay side vs. non-delay side firing, marker external color (magenta vs. cyan) represents left vs. right side firing. All firing rates have been z-scored against all laps of that session. Arrows indicate sections of the maze with different consistent representations than the rest of the maze for that cell.

The finding that the encoding of information about behavioral parameters is not uniform over the track is consistent with our interpretation that the spatial firing patterns seen in cells represent the division of the task into separate cognitive chunks that are processed differentially by the cortex. However, the larger point of both findings should be that if the cortex is processing tasks in separate cognitive chunks, then cells should also be analyzed in reference to those chunks. In other words, if the firing of Cell A on the choice point represents a different cognitive sub-task than the firing of Cell A at the start of maze, then we can decode more information about what the cell firing represents by comparing firing in those two regions separately than we can by combining them. An analogous process is

commonly used by primate researchers in their behavioral recording tasks. These tasks are typically conducted with very regular timing intervals, and neural firing patterns are tracked over the timescale of the experiment. One of the common observations from these types of studies is that neurons tend to respond not only to specific parameters of the task, but they respond at specific times during the task, such as stimulus onset, movement initiation, or during a delay period. In most cases, firing rates are compared in time bins registered relative to the task structure. In fact, in some cases (Crowe et al., 2010; Rigotti et al., 2013) these studies employ decoding approaches to examine the predictive encoding of task parameters on the population level, and for most of these decoding approaches the task is broken into time bins and the encoded parameter is predicted for each time bin. This approach allows researchers to determine when the population most accurately represents a given parameter, but it also reflects the implicit assumption that there is a variation in firing properties of cells over time that makes it necessary to consider the context in which a particular spiking pattern occurred in order to determine what that pattern represents. For temporal based monkey tasks, this context is determined by time, but such a determination may not make sense for spatial rodent tasks.

On spatial tasks in rats alignment relative to time is not commonly possible across an entire session, because the animal is allowed to run at his own pace which makes the amount of time spent on each lap inconsistent (note that alignments to single events such as feeder fire or choice point entry are still often used on such tasks, for example through the construction of PETHs, as we did in Figure 3.5). However, for such a task alignment relative to space seems to be a reasonable alternative, especially given the evidence we have reviewed that the cells in cortex are spatially chunking this task into different cognitive sub-regions. In terms of these tasks then, dividing the track up into a series of spatial bins and using these as the context in which particular firing rate patterns are decoded across the population of cells makes a great deal of sense. We employ an approach informed by this general procedure in Chapter 5 to explore the role of strategic

representations in the rodent PFC.

Chapter 5

Strategic Representations in Rodent PFC

We have explored the response properties of individual rat PFC cells on various tasks, and determined that cells have non-uniform firing properties and representations over space on track-based tasks. We have even briefly considered strategic representations in individual PFC cells on the MT-LRA task. We are now ready to turn our attention to whether rodent PFC cells encode a truly cognitive representation of the strategy an animal is using to solve a particular task. In Chapter 1 we discussed the previous research indicating that cells in rodent PFC represent strategies that animals employ to solve behavioral tasks. However, in order to demonstrate that these representations really correspond to internal strategies, we have to show that they are not simply a result of the cells tracking the parameters of the task set up for the animal. Cognition (in the brain) has been defined as all spikes that are neither directly related to sensory input or motor output.¹ Therefore, in order to demonstrate that strategic representations in rat PFC are cognitive, we need to show that they can occur temporally removed from changes to both sensory input and motor output. That is our goal for this chapter, and

¹ I first heard this definition from Matt Chafee, and although he claims not to have originated it we haven't been able to find the originator, so he gets credit for it here

we will begin by establishing a method for detecting strategic representations in PFC population firing on the MT-LRA task.

5.1 A method for detecting strategy transitions on the MT-LRA task

On the Multiple-T Left, Right, Alternate (MT-LRA) task, animals received reward on every lap based on whether their choices matched the active reward criterion (see Figure 2.2). Animals were trained to expect a single reward criterion on each day, either left(L), right(R), or alternate(A), but for the switch sequence animals were forced to change their reward strategy halfway through the day's session by an unsignalled change in the reward criterion. We already demonstrated in Figure 2.3 that the behavior of the rats on this task indicates that they recognized the criterion change and altered their strategy accordingly in order to continue to receive reward on this task. If the PFC encodes the strategy animals are using to solve this task, we would expect to see a change in the population firing following the switch in reward contingency but before the animal's behavior changes on the task. As described previously (see Figure 2.2), we separated the firing of each cell into seven spatial bins mirrored on either side of the track. We then created a population code that consisted of an M -dimensional vector for each lap that represented the z-scored firing of each cell in each spatial bin ($M = N_{cells} \times 7$ (spatial bins)). Figure 5.1 panel A shows two examples of these population vectors represented in two dimensions. Each of these M -dimensional vectors represents the firing of each cell in the recorded population at each location of the track for a particular lap, or essentially the population representation in PFC for that lap. If the representation in PFC changes over time for this task, these vectors should change as well. So if the representation is different on lap 1 and lap 32 for example, the population vectors for these two laps should be quite different. On the other hand, if the representation in PFC is similar

between these two laps, the population vectors should also be similar. We can assess the similarity of a pair of population vectors by calculating the correlation coefficient between the two. In fact, if we plot the correlation coefficient between each lap's population vector with each other lap, we can get a snapshot of how the population representation across the PFC varies over the course of a session. An example of what these plots might look like is shown in Figure 5.1, which illustrates a subtle but important point. If the representation in PFC changes over time, early laps will be poorly correlated or even anti-correlated with late laps, as seen in both examples. However, if the representation changes gradually over the course of laps, we should see a diagonal stripe of highly correlated laps as shown in panel B. On the other hand, if the representation changes suddenly between two different representations, we should instead see a break point between two highly correlated regions which are poorly correlated (or even anti-correlated) with each other, as shown in panel C.

Figure 5.2 shows an example of the correlation of each of these vectors on a lap by lap basis for an actual session. There are several clear regions along the diagonal of the figure that show a high degree of self-correlation but are negatively correlated with other laps, similar to the example shown in Figure 5.1 panel C, implying that these regions represent different states of population firing in the PFC akin to strategy states. The break-points between these regions represent global transitions in the strategy represented in PFC. In Figure 5.2, it appears that the clearest break-point occurred around the same lap (32-33) as the contingency switch for this session.

In order to detect these changes between representations, we used a k-means clustering analysis on the population firing vectors for each session. This analysis classified each lap as belonging to one of a specified number (k) of clusters on the basis of their position in M -dimensional space. An example is shown in Figure 5.2 panel D. We then used change point analysis (Gallistel et al., 2004) on the sequence of cluster identities by lap to find the change points between clusters and marked these as putative transition points between states in PFC. An example of the

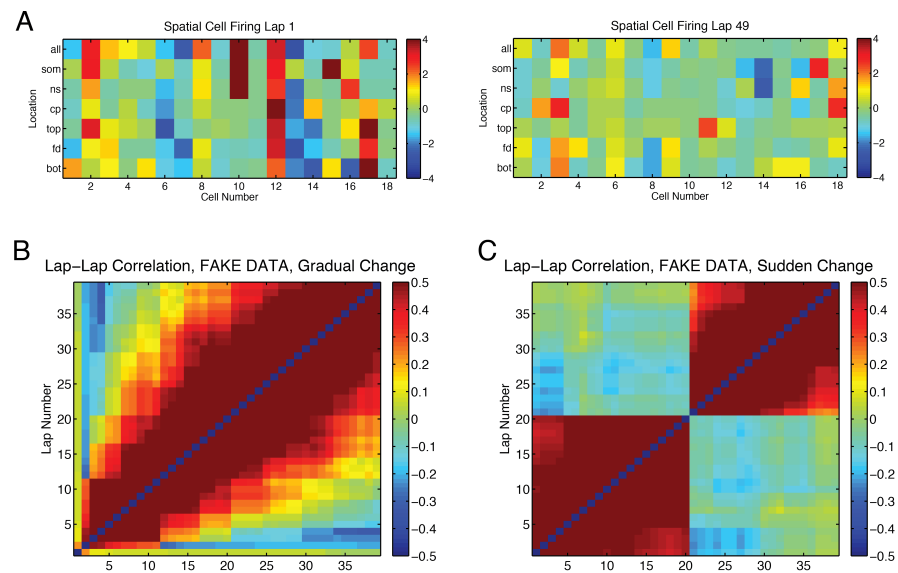


Figure 5.1: PFC Population Firing Representations. (A) Spatial Firing rate population representations for two sample laps from a single session. Note the differences between cell firing between laps. (B) A sample correlation plot on fake data demonstrating what gradual change in population representation over the course of a session would look like. (C) Another correlation plot on fake data demonstrating what a sudden change in population representation would look like.

process of finding change points from the cluster identities is shown in Figure 5.2 panel E. Because the k-means algorithm requires the number of clusters (k) to be specified and its output varies depending on its initial conditions we ran 100 trials for each value of k between two and five (the largest value we could reasonably expect to detect given the number of laps we recorded)² and combined the results into a histogram over laps for each session (as shown in Figure 5.2 panel F). The number of times a specific lap was detected in these histograms was compared to a distribution created from running the same analysis on the same session with the ISIs of all cells randomly shuffled 400 times. Finally we z-scored the number of times a lap was detected from the actual data against the distribution created for that lap from the randomly shuffled data, to create a transition score. The transition score represents the likelihood of a transition occurring on a given lap (the higher the transition score, the greater the probability of a transition on that lap). It is important to stress at this point that the transition score is calculated solely on the basis of the population firing from the recorded cells, without considering any information about the behavior of the animal or when the criterion switch actually occurred.

Figure 5.3 shows the average transition scores across all MT-LRA sessions for ± 10 laps aligned to the switch in reward criterion for a given session. The transition score was low for all laps prior to the switch, began to rise on the switch lap before peaking around one lap later, and then returned to normal levels following the switch. This is the pattern we would expect if the population firing rate encoded a strategic representation in PFC, because the switch provides the first evidence that a change in strategy is necessary. A change in strategy should also precede (or coincide with) a change in behavior on the task. In order to check this proposition, we determined the lap on which the animals behavior changed for each session by running a change point analysis (Gallistel et al., 2004)

² A rule of thumb is that the maximum k should be given by $k \approx \sqrt{n/2}$ (Mardia et al., 1979, page 365), and on average we recorded about 50 laps for MT giving us a max k of five. We used the same max for DD sessions even though we had more laps on those sessions to keep the analysis consistent.

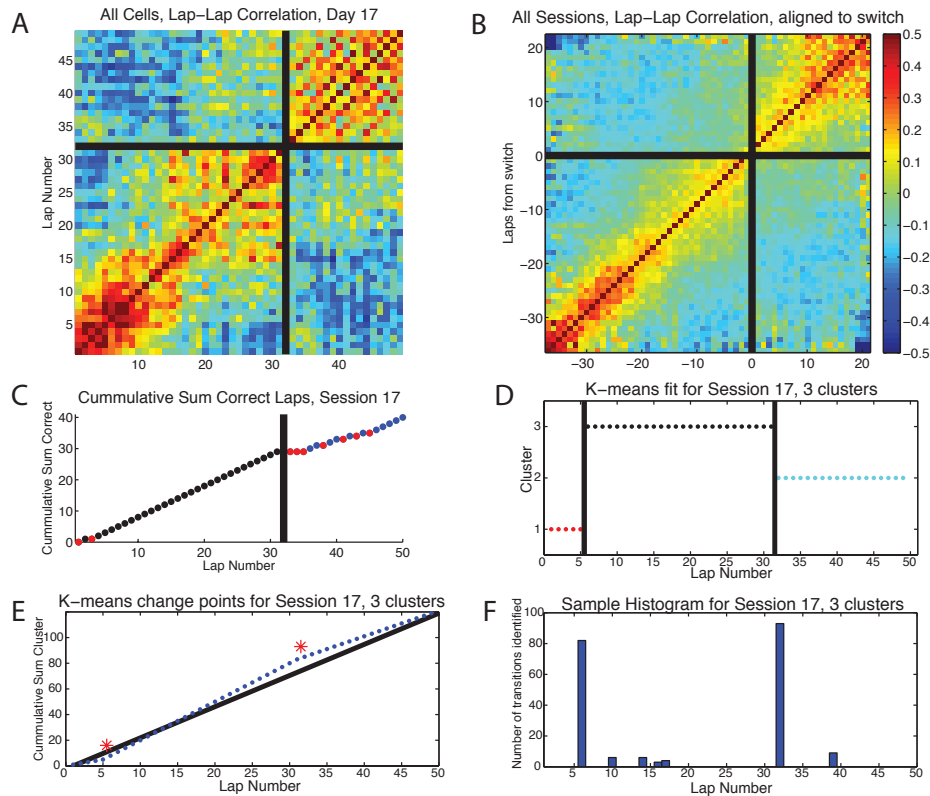


Figure 5.2: An example of the generation of the transition probabilities from a single MT-LRA session. (A) The lap-lap correlation of the spatial population firing vectors for one session, with the switch lap marked (black line). (B) The average lap-lap correlation of the spatial population firing vectors over all sessions aligned to the switch lap for each session (C) The animal's behavior on the same session. Red dots = error laps, black dots = correct laps before switch, blue dots = correct laps after switch. (D) An example k-means fit for the session with three clusters. Transitions between groups of clusters are indicated (black lines). (E) Example of the Change Point Algorithm on the k-means data. Change points are identified as the largest deviations of the slope of the cumulative sum from the average slope. (F) Sample Histogram of all detected k-means change points for 100 iterations with $k = 3$ clusters.

on the animals L/R side choices for each lap (for an example see Figure 5.2 panel C). Figure 5.3 also shows the average transition scores across all MT-LRA sessions aligned to the lap of greatest behavioral change for each session. Relative to the behavioral change, the transition score actually showed a peak above baseline four laps before the animal's behavior changed, and remained somewhat elevated before peaking again on the same lap that the animal's behavior changed, consistent with a new representation in the PFC driving the change in behavior. Indeed, most behavior change laps followed the switch lap by between zero and four laps, so the transition probability in PFC peaked in between the two.

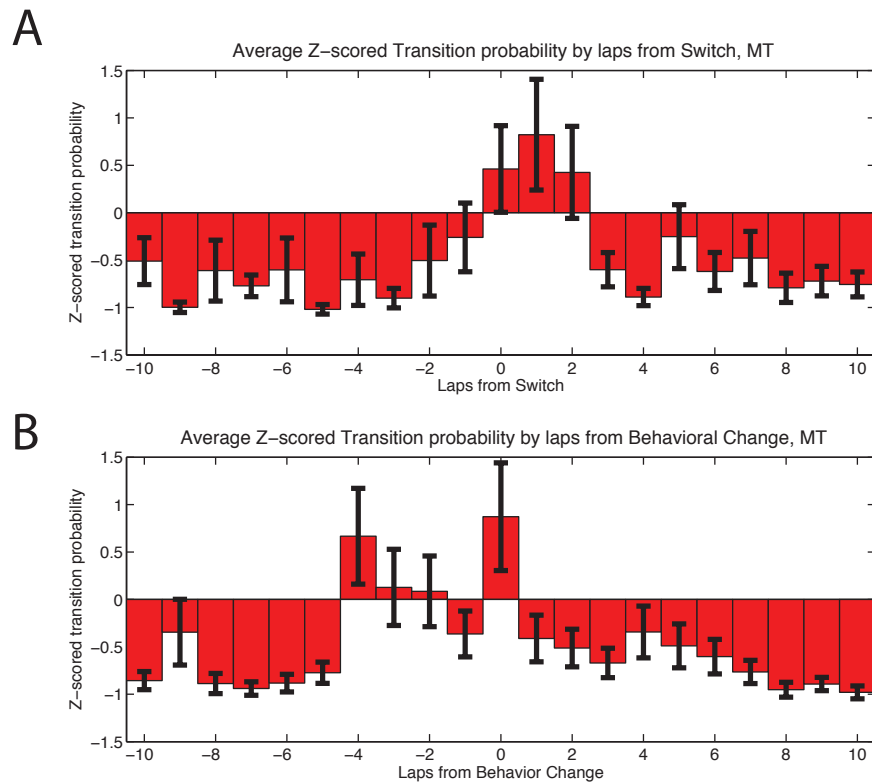


Figure 5.3: Transition scores for all MT sessions (A) A 21 lap window of averaged transition scores over all MT sessions aligned to the switch lap for each session (error bars SEM) (B) A 21 lap window of transition probabilities aligned to the lap of greatest behavioral change for all MT sessions (error bars SEM).

5.2 Strategy changes without behavioral change

As we have discussed (see Chapter 4) cells in rodent PFC have differential spatial firing patterns. Although for most spatial tasks these non-uniform firing patterns tend to be similar on both sides of the track, some cells do have different firing patterns on different sides of the track. Indeed, we have shown that some cells in PFC code for the side of the track that the animal visits on a given lap (see Chapter 3). Therefore, it is possible that the transition in the population firing that we see around the switch reflects the animal visiting a side of the track that they haven't visited before the switch (for example when switching from a Left strategy to an Alternate strategy the animal will begin making laps to the right only following the switch; if laps to the right produce different firing in some cells, that may account for the change in the population code following the switch). In order to control for this possibility, we considered all sessions in which the strategy before or after the switch was Alternate (which required laps to both sides) and re-ran the analysis considering only laps to the side of the track present both before and after the switch (so for a L→A session, we considered only left laps). In all of these cases, the transition scores identified on laps to only one side of the track matched those identified on both sides, indicating that the population firing on laps to the left side of the track differs between laps run during a Left strategy and laps run during an Alternate strategy. Figure 5.4 shows transition scores calculated from these restricted sets of laps to only one side of the track. It is important to point out two factors that make this control analysis less precise and accurate than the full version shown in Figure 5.3. First, since we used only sessions that include Alternate strategies, we only considered 12 sessions here instead of the 18 used in the full analysis. Second, because we only considered a subset of laps, the precision with which we could identify a particular change lap was weakened. For example, if the change occurred between lap 23, which is to the right, and lap 24, but we considered only laps to the left, we were only able to identify the change as occurring between laps 22 and 24, which means we had

to assign it a 50% chance of occurring between laps 22 and 23, and a 50% chance of occurring between laps 23 and 24. This problem was worse in cases where the change occurred in a series of several consecutive laps to the non-considered side. In general, the effect of these two factors was to broaden and weaken the effects seen in the aligned and averaged transition scores. Even still, there is a notable effect in Figure 5.4. For the criterion switch alignment, there is still a heightened transition score around and following the switch, but it is attenuated relative to the non-restricted version due to the factors mentioned above. For the behavioral change alignment, the effect several laps prior to the behavior change and around the behavior change are still evident, although in this figure the heightened peak around the behavior change now occurs a lap afterward. It is likely that the shift to a lap later is due to the lowered precision mentioned above. Regardless, this control demonstrates that there is a measureable strategy transition even among laps run to the same side of the track, so the changes we are seeing can't be accounted for solely by animals visiting spatial locations they have not before, or representing different side choices independent of an overall strategy.

The heightened transition score on the laps surrounding a forced strategy switch indicates that the transition score is a marker for strategic changes in the animal's PFC. However, previous results have indicated that a change in reward criterion should cause a change in the representation in PFC (Rich and Shapiro, 2009; Durstewitz et al., 2000; Benchenane et al., 2010; Karlsson et al., 2012). We still have not shown a change that is independent of either sensory input (the imposed switch in reward criterion that precedes cellular transition) or behavioral output (the change in the animal's choices that follows the cellular transition). What we have done is to outline a technique to identify strategic transitions solely based on cellular firing, which we can now use to look for changes that are independent of behavioral changes or sensory changes. We will consider a transition in the prefrontal representation independent of a behavioral change first.

On the first several laps of the MT-LRA task, animals had to identify the initial

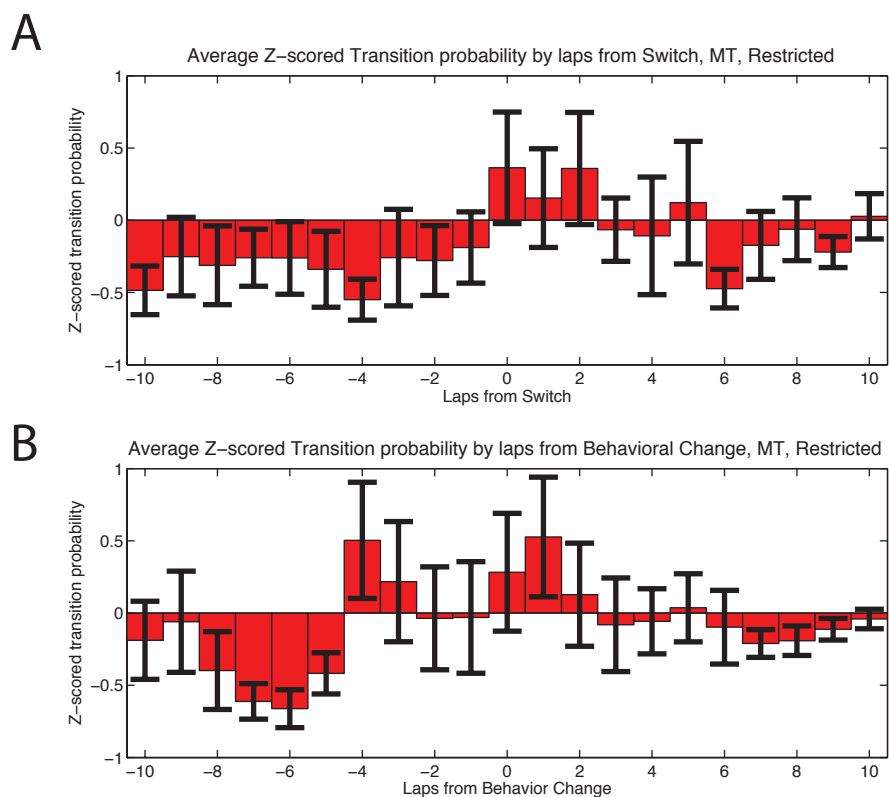


Figure 5.4: Transition scores for all MT sessions from only single sides of the track (A) A 21 lap window of averaged transition scores aligned to the switch lap for each session (error bars SEM) (B) A 21 lap window of transition probabilities aligned to the lap of greatest behavior change (error bars SEM).

reward criterion in effect on that day. In the single session example presented in figure 5.2, we can see that while the primary transition between areas of higher self-correlation occurs around the switch lap, a secondary transition occurs around laps 5-7. The timing of this transition is exactly when we would expect animals to transition from an initial search state to a state representing the pre-switch reward criterion. In Figure 5.5 Panel A we plot the average transition score for all MT-LRA sessions for laps 5-30. This score has peaks around laps 6, 8, and 11 before dipping well below zero for subsequent laps, indicating that while the precise lap on which animals make this latent transition from search state to hidden strategy state on each session is not consistent, it tends to occur within those first few laps (note that when averaged this way, the switch laps for each session do not line up, so on average the transition scores for later laps remain low). In addition, Figure 5.5 displays the average lap-lap population firing correlation over all sessions aligned to the first 40 laps of the session (panel B) and the average correlation for all session aligned to the lap with the greatest transition score in the first 15 laps of the session (panel C). In either case, it is clear that there is a heightened representational correlation in the beginning of the session that changes after the first several laps.

It is possible that this initial heightened representation is simply one of the three strategy states that the animal expects to find on a given session (L, R, or A) which happens to not match the state he found to be in effect. Because we have shown that cells on the MT-LRA task individually represent the same task parameters when recorded across several consecutive days (see Figure 3.6), we reasoned that the population of consistently recorded cells would represent the same states across days. Therefore, we examined the consistency of these initial state representation in the same population of cells recorded across several days of running on the task. In order to do so we simply collected the population of all cells recorded across a period of several days, calculated the spatial-population vectors for each lap across each day (as described earlier, Figure 5.2), and ran the same correlation analysis across this collection of lap-by-lap population vectors.

Examples of the resulting correlation plots from two rats are shown in Figure 5.5 panel D. These figures demonstrate that population firing codes for the initial state of each session are quite similar from one day to another but different from all the other states represented both before and after the switch for the rest of the session, indicating that this initial state is likely representative of a global search state distinct from the other strategic states represented.

Animals on the MT-LRA task consistently show a prefrontal representational change from the first few laps of running on a session until the time at which they appear to understand the reward criterion for that session. This initial state represented in the PFC seems to be independent of the actual criterion provided to the animal and is consistent on each day the animal runs the task. Furthermore, while the animal gains knowledge about the state of the task during these several laps, there has been no change in the actual experimental setup, only a change in the animal's knowledge of it. Thus the transition within the first several laps of the session represents a change in the animal's strategic representation that is not based on an overt change in the task parameters or the animal's behavior on the task. However, this change represents more of a change in the animal's understanding of the task than necessarily the strategy he is using to solve the task. Next we will demonstrate a transition in the animal's PFC representation driven solely by the animal's decision to employ a different strategy to solve various phases of the task. In order to do so, we will consider the Delay Discounting task.

5.3 Strategic Representations on the DD task

As we have discussed, behavior on a typical session of the Delay Discounting task tends to proceed in three different phases(see Figure 2.4). First animals explore their surroundings and learn the initial setup for this particular session (exploration), then they make a preponderance of their laps to one side of the track in order to adjust the delay value to a more palatable one (titration), then they tend to alternate laps to either side of the track, keeping the delay at the same

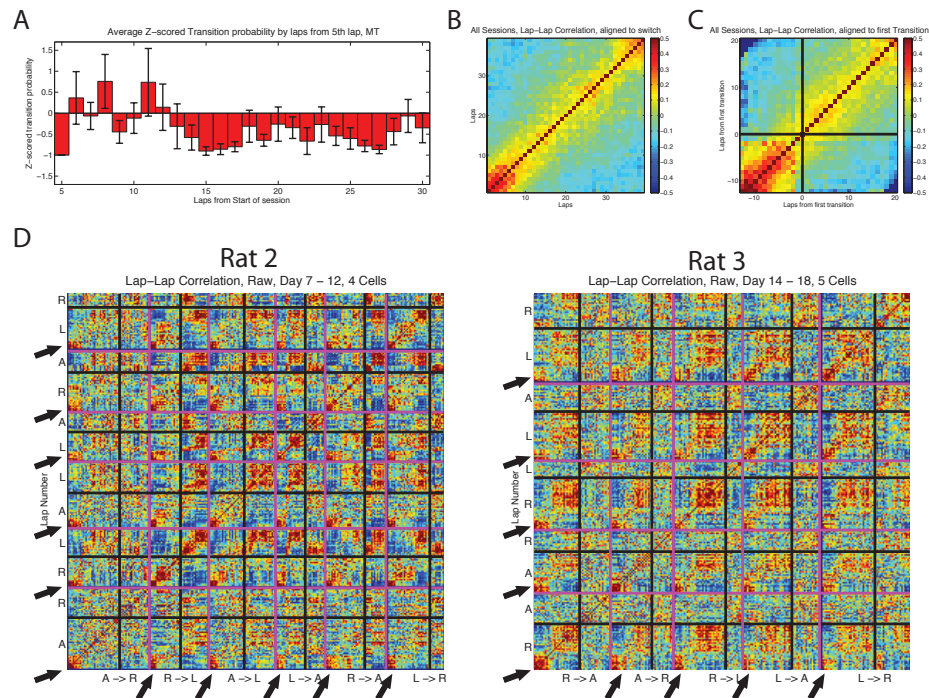


Figure 5.5: Transition scores for all MT sessions from session start (A) Average transition scores of all sessions aligned to laps 5-30 (error bars SEM) (B) Lap by lap correlation matrix of population firing averaged over all sessions, aligned to start of session. (C) Lap by lap correlation matrix of population firing, averaged over all sessions, aligned to lap of highest transition score in the first 15 laps of each session. (D) Correlation plots for the spatial- population firing vector for all cells recorded consistently across several days for two rats (4 cells across six days for Rat 2 and 5 cells across five days for Rat 3). Magenta lines indicate the beginning and end of individual sessions, and black lines indicate the switch in each session. Arrows indicate areas of heightened self-correlation at the beginning of each session.

value while receiving as many rewards as possible (exploitation or alternation). These behavioral patterns have been discussed and documented by previous experiments in the lab (Papale et al., 2012). However, it is reasonable to assume that these behaviors represent different strategies that the animal is employing to solve the task, in which case we would expect to find representations of them (and specifically the transitions between them) in the PFC. So we decided to test the hypothesis that these different behavioral modes are represented as different strategies in PFC by calculating the transition score for DD sessions as described above and checking to see whether there was a heightened transition score around the laps on which the animal's behavior changed.

In order to detect when the switch between behavioral phases occurred on individual sessions of the DD task, we again employed a change point analysis (Gallistel et al., 2004) on the animal's choice of large or small reward arm visits throughout the session to find the lap on which the greatest behavioral change occurred (see Figure 5.6). In general, this change corresponded to the transition from the titration to the alternation phase. Figure 5.7 shows the lap-lap correlation values for the population firing rate vectors (calculated as above for the MT-LRA task, using the maze sections identified in Figure 2.4) for an example session. In this case we have indicated the behavioral change point on the correlation plot to indicate the lap on which the animal's behavior changed, and again the population state in the PFC seems to represent two different states for the two different behavioral phases. We calculated the transition scores over all laps using the same methods described above for the MT-LRA task (for examples of the steps for DD, see Figure 5.7).

Figure 5.8 shows the average transition score over all DD sessions aligned to the behavioral change points we identified for each session. Similar to what we saw in Figure 5.3 for the MT-LRA sessions aligned to behavioral change, the transition score peaked several laps before the behavioral change (in this case eight laps before) then peaked again on the lap of the greatest behavioral change. From an analysis averaged over all sessions and all animals, it is possible that the change

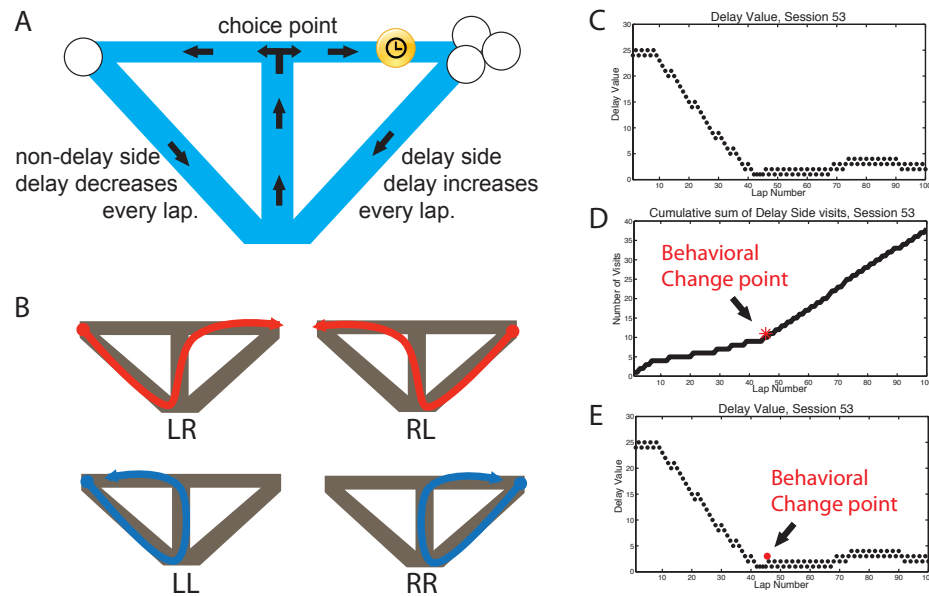


Figure 5.6: Finding Behavioral Changes on the DD task. (A) The track utilized in the DD task, for reference (B) Animals take four distinct paths on the DD task, which will help us answer the Euston-Cowen Hassle. Left to Right (LR), Right to Left (RL), Left to Left (LL), and Right to Right (RR). (C) Animal's behavior on a typical DD session plotted by delay value (for the delayed side) vs. lap number. A clear behavioral transition occurs between laps 41 and 45. (D) Calculating the behavioral change point (indicated by the red star) from the maximum deviation from the average slope on a plot of the cumulative sum of visits to the delayed side of the track. (E) The behavioral change point found in (D) matches the animal's behavioral transition seen in (C).

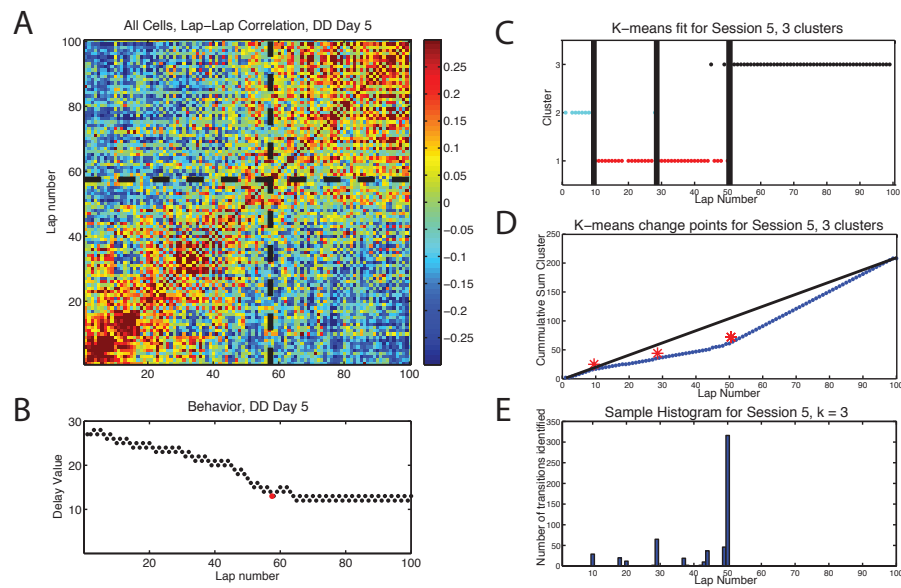


Figure 5.7: An example of the generation of the transition probabilities from a single DD session. (A) The lap-lap correlation of the spatial population firing vectors for one session, with the detected behavioral change lap marked (dotted line). (B) The animal's behavior on the same session, delay value vs. time. Red dot indicates Behavioral change lap. (C) An example K-means fit for the session with three clusters. Transitions between groups of clusters are indicated (black lines). (D) Example of the change point algorithm on the K-means data. Change points are identified as the largest deviations of the slope of the cumulative sum from the average slope. (E) Sample Histogram of all detected k-means change points for 100 iterations with $k = 3$ clusters.

precisely eight laps prior to the behavioral change is not uniform, but merely represents the most common change among all animals. However, the transition score clearly supports the interpretation that the representation in the animal's PFC changes to reflect a change in behavioral strategy, and this representational change precedes the change in the animal's behavior.




Figure 5.8: Transition probabilities for all DD sessions. A 21 lap window of transition probabilities for all DD sessions aligned to the behavior change lap for each session (error bars: SEM)

It is important to stress, again, that the behavioral change in the DD task, unlike all previous examples of behavioral changes demonstrated in rodent PFC, was not prompted by an experimenter changing the rules of the task the animal must solve, but was instead driven by internal processes responding to the demands of the task. In other words, the animal's own internal decision processes. In this way, the strategy change we have detected on the DD task represents a true strategic choice of the animal, independent of a direct sensory cue, and because it occurs prior to the behavioral change, also independent of a direct motor output. But before we can fully conclude that what we are seeing is a cognitive change, we have to make sure we have accounted for all other possible explanations.

5.4 Wrestling with the Euston-Cowen Hassle

As previously discussed, the Euston-Cowen Hassle describes the fact that cells in rodent PFC (particularly in dorsal regions) often display movement-related firing patterns, which can in many cases better account for their firing variation on tasks than the cognitive patterns that are often attributed to them. In addition to providing a structure in which animals are incentivized but not forced to make a strategic change in task execution, the Delay Discounting task provides us with an ideal mechanism for testing whether the population firing patterns we detected are unduly influenced by subtle differences in the path run by the animal on these tasks.

As depicted in Figure 5.6 panel B, animals on the DD task run four different types of laps, laps traversing the track from the left feeder to the right feeder, laps traversing the track from the right feeder to the left feeder, and laps back to the same side of the track, either from right to right or left to left. In general, the laps traversing from one side to the other are more common, although when the animal needs to adjust the delay value a large amount either up or down they will have to make many consecutive laps to the same side. Although the animal's trajectory varied grossly between these lap types, trajectories on the same type of laps are often highly consistent, especially when animals display some evidence of automation over sets of multiple laps with the same behavioral strategy. If the population firing variation which we are attributing to strategic representations is due to cognitive changes, and not postural changes, in the animal, we should see low correlations between laps of the same trajectory type across different strategic representations, but high correlations among laps of different trajectory types in the same strategic representation. However, if the population firing variation is better explained by the Euston-Cowen Hassle, we would expect to see consistently high correlations on laps of the same trajectory type regardless of strategic representation.

In order to test this question, we temporally normalized the animal's trajectories on each lap to specific start and end locations, then found normalized trajectory vectors for each lap. We then correlated these vectors with each other to produce lap-lap plots for trajectory correlations similar to our population firing rate trajectory plots. Figure 5.9 shows these plots for an example session. From the population firing correlation plot in panel A, we can see that this session has four different states of higher self-consistency and lower cross-consistency with other representations. These states are marked by black lines super-imposed on all plots (but note, in this case the states were determined by examination, and are provided simply for convenience). Panel B shows the overall spatial correlation across all laps. It is clear from examining panels A and B that the population firing rate correlations depicted in panel A don't match the trajectory correlations of the animal, depicted in panel B. Therefore it is very unlikely that the population firing rate is well explained by changes in the animal's trajectory. To further demonstrate this fact, panel D breaks down the cellular population correlations into individual trajectory types. It is quite clear that laps of the same trajectory type are not highly correlated across different strategic regions for any of the trajectory types. The difference is even quite clear in this example session for LL and RR type laps. Even though there are few of each of these types of laps, their population representations are quite different in the different strategic regions, but quite consistent within regions, indicating that the population firing is better explained by the strategic representations than by the animal's trajectory.

As an additional control, we reasoned that if population firing vectors were influenced by the trajectories animals ran, they would differ most for different trajectory types. Therefore, if we can still detect transitions around areas of behavioral change by considering only particular lap types, we would show that the population firing rate is different even on laps of the same type for the different strategic representations. This is very similar to the control analysis for the MT-LRA task in which we considered only laps to the left or right side of sessions. However, for the DD analysis we considered only LR laps and RL laps (because in

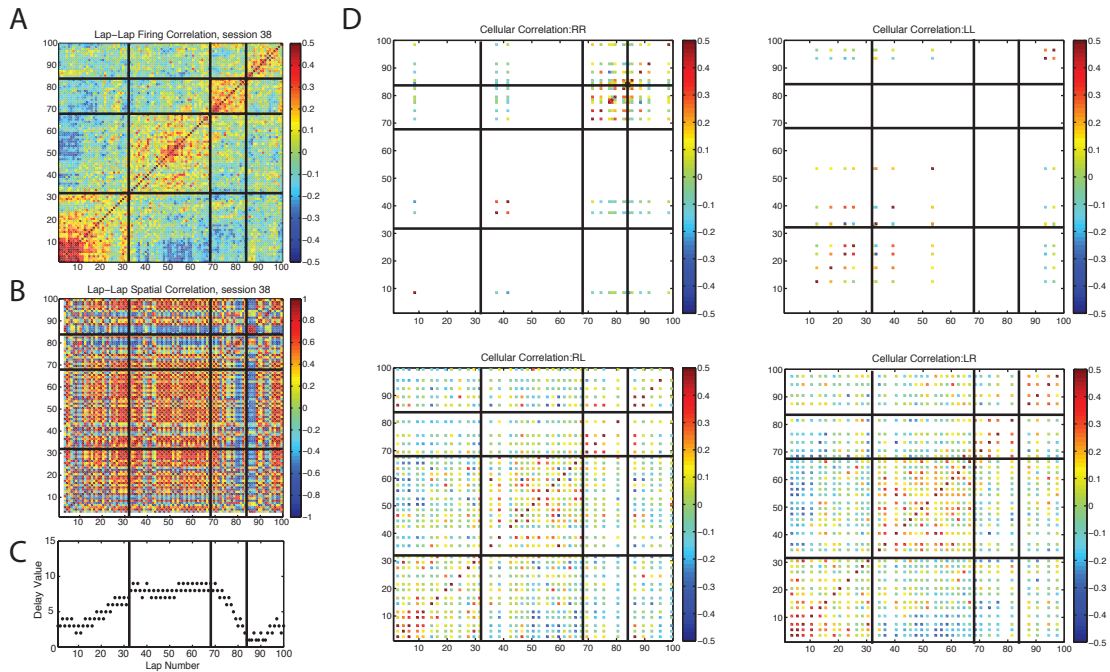


Figure 5.9: Population Firing Correlations are not due to trajectory similarity. (A) Lap by lap population firing vector correlations for a sample session, as in Figure 5.7. Note here, black lines are drawn for illustration purposes to illustrate the different states present in the population firing, and reproduced on all plots. (B) The lap by lap trajectory correlation plot for the animal's running on the task. (C) The delay value of the delayed side by lap for this session. Note the black lines mark the same laps as in other panels, and were determined from cellular firing, NOT behavior. (D) The data plotted in (A) broken down by lap type to show correlations only among laps run to and from the same locations on the track.

almost all sessions there were not enough RR or LL laps to reasonably determine any changes). As a result, we did not have to throw out any sessions in the DD case (because DD animals made a high number of LR and RL laps in all strategy cases), but we also have even more lap sparsity than for the MT-LRA control because there were some lap types that were excluded entirely. As before, we separated laps into groups of each type and then ran the transition score analysis on the separated groups, then combined them and aligned them to the behavioral change laps as we had previously. The result is shown in Figure 5.10. In this case, the heightened transition at the behavioral change is not as apparent, although there is a trend to increase the transition score at that point. However, the increase in transition score several laps prior to the behavioral change is still present, indicating that even considering only laps following the same overall trajectory we still see changes in the population representation on the DD task that preceded the change in the animal's behavior and are not forced on the animal by an experimental manipulation.

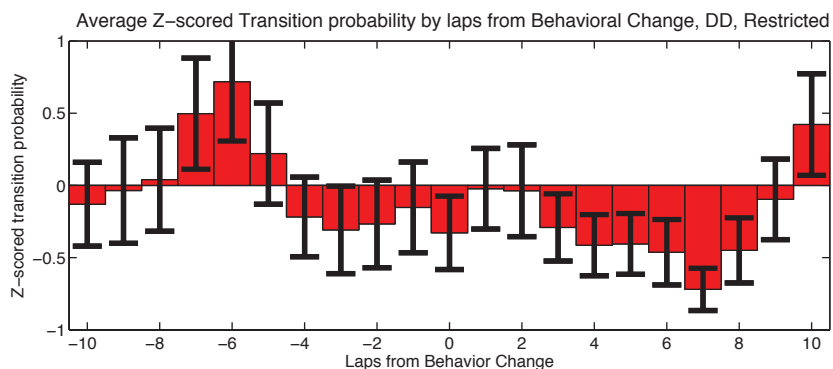


Figure 5.10: Transition probabilities for all DD sessions, for laps with the same trajectory type. A 21 lap window of transition probabilities for all DD sessions, calculated for a collection of only LR laps and only RL laps and averaged together, aligned to the behavior change lap for each session (error bars: SEM)

5.5 Discussion and Conclusion

We have shown that a representational change in the rat PFC corresponds to a strategic change in the animal's performance on various spatial decision-making tasks. In situations where the animals were prompted to adopt a new strategy by a change to the reward rule on the task, the strategy transition followed the receipt of information that the contingency of the task had changed, but preceded the change in the animals behavior on the task. Representational transitions with the same hallmarks also occurred on a task in which the animal was not forced to make a strategic change, but instead chose to in order to maximize his behavior. In that case, the representational transitions still preceded behavioral changes, implying that they were more likely guiding the change in behavior than reflecting it.

Some might question whether there is in fact a change in the stimuli presented on the DD task which could be the source of the population firing change in the PFC. Indeed, the delay value should be changing over the course of laps throughout the titration period and this steady change should cease when the animal begins his alternation phase. However, we do not believe that this transition accounts for the representational changes that we see in the PFC for several reasons. First of all, the transition scores we discovered in the DD task preceded the behavioral changes we detected by multiple laps, but the change between titration phase and alternation phase tended to coincide with the behavioral change, indicating that the representational change in PFC occurred *prior* to the change between these phases, not coincidentally with it. Further, the change in delay value is a gradual change, and as such it should produce a gradual change in the population representation in the PFC. A gradual change in population firing would produce a distinctive appearance on the population firing correlation plots we considered, as depicted in Figure 5.1 panel B. However, the lap-lap correlation plots we considered were more consistent with a sudden change in population representation, as depicted in Figure 5.1 panel C (for comparison to actual data, see

Figures 5.7 and 5.9). The important point is that the transitions we detected in PFC are more consistent with a sudden complete change in the state represented rather than a slow steady change from one state type to another.

Another interesting point to consider from our lap by lap correlation plots is that in many cases, we see not only a low correlation between different states, but actually an anti-correlation between different states. While this distinction might not appear immediately significant, if we merely scrambled the firing rates between cells we would expect the correlation to approach zero, implying that the two vectors are not related. A negative correlation actually implies that the two vectors are in fact opposites of each other in some ways. The implication of this fact is not abundantly obvious, but it may be related to another remarkable firing property of prefrontal neurons in general. It has been observed that cells in PFC have a tendency to represent various factors by both increasing and decreasing their firing rates (Ma et al., 2014). In fact, we saw this same property in our single-cell analysis of the MT-LRA task (see Chapter 3). On a population level, in most cases it seems that half of the cells will increase their firing in response to a particular parameter (for example the right side of the track) while the other half will decrease their firing to the same parameter. One result of this property is that in general the firing rate in the PFC as a whole tends to remain rather constant over behavior, with nearly equal numbers of cells increasing and decreasing their firing in response to nearly every variation. Another consequence could be the anti-correlated representations evident in some cases. Having a large population of cells increase their firing rate from one representation to another while a similarly-sized population of cells decreases its firing rate can give rise to anti-correlated representations that you would not see if cells simply increased or decreased their firing rates.

The fact that states are anti-correlated might lead us to believe that the population of cell firing tends to represent only two states, but this appears not to be the case, particularly for our MT-LRA task. We consistently found a transition between states around the contingency switch, but also we consistently found an

initial representation for the first several laps of the session that was distinct from the other two representations. This demonstrates that the PFC can in fact generate many independent representations, which is to be expected for such a high dimensional construct. However, it seemed as though the initial state representation was generally more likely to be anti-correlated with the two represented strategies on either side of the switch than they were with each other. If true, there is an implication that the initial state is in some way more distinct from the other two states. One possible explanation for this discrepancy would be that the initial state is involved in more deliberative processing while the others might represent a transition between two previously learned habitual states. Indeed, previous results have connected the PFC to establishing and controlling habitual behaviors (Killcross and Coutureau, 2003; Smith et al., 2012), and previous work has indicated that the first several laps of a session on the MT-LRA task are more likely representative of a deliberative decision-making system while later stages are more likely to be habitual (Johnson and Redish, 2007; van der Meer and Redish, 2010). However, the general hallmark of deliberative decision-making on such tasks is the presence of Vicarious trial and Error (VTE) behaviors at the choice point, and we were unable to find a clear relationship between representations in PFC and VTE on either task. Moreover, if the initial state representation in PFC were related to deliberative decision making, we would expect it to recur on the laps immediately following the contingency switch in MT-LRA sessions, and there does not seem to be much evidence that it does so (see Figure 5.5).

In conclusion, this thesis has demonstrated that through considering population firing of cells in PFC with particular attention paid to the spatial firing patterns of cells, and taking advantage of the consistency of cells across multiple recording days to maximize information, we can demonstrate that PFC displays cognitive representations of strategies animals are using to solve behavioral decision making tasks. Furthermore, the techniques and results developed herein provide a solid starting point for even more promising explorations of the role of the PFC in guiding and shaping rodent behavior, which may ultimately help us understand

the role of this complex structure across species.

References

- Alexander, G. E. and Fuster, J. M. (1973). Effects of cooling prefrontal cortex on cell firing in the nucleus medialis dorsalis. *Brain research*, 61:93–105.
- Allen, W. F. (1940). Effect of ablating the frontal lobes, hippocampi, and occipitoparieto-temporal (excepting pyriform areas) lobes on positive and negative olfactory conditioned reflexes. *Am. J. Physiol*, 128(7):54–777.
- Barnes, C. A., Suster, M. S., Shen, J., and McNaughton, B. L. (1997). Multistability of cognitive maps in the hippocampus of old rats. *Nature*, 388(6639):272–275.
- Bauer, R. H. and Fuster, J. M. (1976). Delayed-matching and delayed-response deficit from cooling dorsolateral prefrontal cortex in monkeys. *Journal of comparative and physiological psychology*, 90(3):293.
- Benchenane, K., Peyrache, A., Khamassi, M., Tierny, P. L., Gioanni, Y., Battaglia, F. P., and Wiener, S. (2010). Coherent theta oscillations and reorganization of spike timing in the hippocampal- prefrontal network upon learning. *Neuron*, 66(6):921–936.
- Blumenthal, A., Steiner, A., Seeland, K., and Redish, A. D. (2011). Effects of pharmacological manipulations of NMDA-receptors on deliberation in the Multiple-T task. *Neurobiology of Learning and Memory*, 95(3):376–384.
- Brito, G. N. and Brito, L. S. (1990). Septohippocampal system and the prelimbic sector of frontal cortex: a neuropsychological battery analysis in the rat. *Behavioural Brain Research*, 36(1-2):127–146.

- Brody, E. B. and Rosvold, H. E. (1952). Influence of prefrontal lobotomy on social interaction in a monkey group. *Psychosomatic Medicine*, 14(5):406–415.
- Brutkowski, S. (1965). Functions of prefrontal cortex in animals.
- Brutkowski, S. and Dabrowska, J. (1963). Disinhibition after prefrontal lesions as a function of duration of intertrial intervals. *Science*, 139(3554):505–506.
- Butter, C. M., Snyder, D. R., and McDonald, J. A. (1970). Effects of orbital frontal lesions on aversive and aggressive behaviors in rhesus monkeys. *Journal of comparative and physiological psychology*, 72(1):132.
- Buzsáki, G. (2004). Large-scale recording of neuronal ensembles. *Nature Neuroscience*, 7(5):446–451.
- Condé, F., Maire lepoivre, E., Audinat, E., and Crépel, F. (1995). Afferent connections of the medial frontal cortex of the rat. II. Cortical and subcortical afferents. *The Journal of Comparative Neurology*, 352(4):567–593.
- Corwin, J. V., Kanter, S., Watson, R. T., Heilman, K. M., Valenstein, E., and Hashimoto, A. (1986). Apomorphine has a therapeutic effect on neglect produced by unilateral dorsomedial prefrontal cortex lesions in rats. *Experimental neurology*, 94(3):683–698.
- Cowen, S. L. and McNaughton, B. L. (2007). Selective delay activity in the medial prefrontal cortex of the rat: The contribution of sensory-motor information and contingency. *Journal of Neurophysiology*.
- Crowe, D. A., Averbeck, B. B., and Chafee, M. V. (2010). Rapid Sequences of Population Activity Patterns Dynamically Encode Task-Critical Spatial Information in Parietal Cortex. *Journal of Neuroscience*, 30(35):11640–11653.
- Dalley, J. W., Cardinal, R. N., and Robbins, T. W. (2004). Prefrontal executive and cognitive functions in rodents: neural and neurochemical substrates. *Neuroscience & Biobehavioral Reviews*, 28(7):771–784.

- Damasio, A. (1994). *Descartes' error: Emotion, reason, and the human brain*. New York.
- de Saint Blanquat, P., Hok, V., Alvernhe, A., Save, E., and Poucet, B. (2010). Tagging items in spatial working memory: A unit-recording study in the rat medial prefrontal cortex. *Behavioural Brain Research*, 209(2):267–273.
- Dias, R., Robbins, T., and Roberts, A. C. (1997). Dissociable forms of inhibitory control within prefrontal cortex with an analog of the wisconsin card sort test: restriction to novel situations and independence from on-line processing. *The Journal of Neuroscience*, 17(23):9285–9297.
- Donoghue, J. P. and Wise, S. P. (1982). The motor cortex of the rat: cytoarchitecture and microstimulation mapping. *The Journal of Comparative Neurology*, 212(1):76–88.
- Durstewitz, D., Seamans, J. K., and Sejnowski, T. J. (2000). Dopamine-Mediated Stabilization of Delay-Period Activity in a Network Model of Prefrontal Cortex. *Journal of Neurophysiology*, 83(3):1733–1750.
- Durstewitz, D., Vittoz, N. M., Floresco, S. B., and Seamans, J. K. (2010). Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. *Neuron*, 66(3):438–448.
- Eichenbaum, H. (2000). A cortical-hippocampal system for declarative memory. *Nature Reviews Neuroscience*, 44(1):41–50.
- Eichenbaum, H. and Cohen, N. J. (2014). Can we reconcile the declarative memory and spatial navigation views on hippocampal function? *Neuron*, 83(4):764–770.
- Euston, D. R. and McNaughton, B. L. (2006). Apparent Encoding of Sequential Context in Rat Medial Prefrontal Cortex Is Accounted for by Behavioral Variability. *Journal of Neuroscience*, 26(51):13143–13155.

- Fenton, A. A., Kao, H.-Y., Neymotin, S. A., Olypher, A., Vayntrub, Y., Lytton, W. W., and Ludvig, N. (2008). Unmasking the CA1 Ensemble Place Code by Exposures to Small and Large Environments: More Place Cells and Multiple, Irregularly Arranged, and Expanded Place Fields in the Larger Space. *J. Neurosci.*, 28(44):11250–11262.
- Floresco, S. B., Seamans, J. K., and Phillips, A. G. (1997). Selective roles for hippocampal, prefrontal cortical, and ventral striatal circuits in radial-arm maze tasks with or without a delay. *Journal of Neuroscience*, 17(5):1880–1890.
- Funahashi, S., Bruce, C. J., Goldman-Rakic, P. S., et al. (1989). Mnemonic coding of visual space in the monkeys dorsolateral prefrontal cortex. *J Neurophysiol*, 61(2):331–349.
- Fuster, J. (2008). *The Prefrontal Cortex*. Elsevier, 4th edition edition.
- Fuster, J. M. and Alexander, G. E. (1973). Firing changes in cells of the nucleus medialis dorsalis associated with delayed response behavior. *Brain research*, 61:79–91.
- Fuster, J. M., Alexander, G. E., et al. (1971). Neuron activity related to short-term memory. *Science*, 173(3997):652–654.
- Fuster, J. M. and Bauer, R. H. (1974). Visual short-term memory deficit from hypothermia of frontal cortex. *Brain research*, 81(3):393–400.
- Fuster, J. M., Bauer, R. H., and Jervey, J. P. (1982). Cellular discharge in the dorsolateral prefrontal cortex of the monkey in cognitive tasks. *Experimental neurology*, 77(3):679–694.
- Gallistel, C. R., Fairhurst, S., and Balsam, P. (2004). Inaugural Article: The learning curve: Implications of a quantitative analysis. *Proceedings of the National Academy of Sciences, USA*, 101(36):13124–13131.

- Gemmell, C., Anderson, M., and O'Mara, S. M. (2002). Deep layer prefrontal cortex unit discharge in a cue-controlled open-field environment in the freely-moving rat. *Behavioural Brain Research*, 133(1):1–10.
- Gibson, J. J. (1941). A critical review of the concept of set in contemporary experimental psychology. *Psychological bulletin*, 38(9):781.
- Goldman-Rakic, P. S. (1994). Working memory dysfunction in schizophrenia. *The Journal of neuropsychiatry and clinical neurosciences*.
- Gupta, A. S., van der Meer, M. A. A., Touretzky, D. S., and Redish, A. D. (2010). Hippocampal Replay Is Not a Simple Function of Experience. *Neuron*, 65(5):695–705.
- Hardy, S. G. and Holmes, D. E. (1988). Prefrontal stimulus-produced hypotension in rat. *Experimental Brain Research*, 73(2):249–255.
- Hok, V., Save, E., Lenck-Santini, P. P., and Poucet, B. (2005). Coding for spatial goals in the prelimbic/infralimbic area of the rat frontal cortex. *Proceedings of the National Academy of Sciences, USA*, 102(12):4602–4607.
- Hoover, W. B. and Vertes, R. P. (2007). Anatomical analysis of afferent projections to the medial prefrontal cortex in the rat. *Brain Structure and Function*, 212(2):149–179.
- Horst, N. K. and Laubach, M. (2009). The role of rat dorsomedial prefrontal cortex in spatial working memory. *NSC*, 164(2):444–456.
- Horst, N. K. and Laubach, M. (2012). Working with memory: evidence for a role for the medial prefrontal cortex in performance monitoring during spatial delayed alternation. *Journal of Neurophysiology*, 108(12):3276–3288.
- Horst, N. K. and Laubach, M. (2013). Reward-related activity in the medial prefrontal cortex is driven by consumption. *frontiers in Neuroscience*, 7:56.

- Iversen, S. D. and Mishkin, M. (1970). Perseverative interference in monkeys following selective lesions of the inferior prefrontal convexity. *Experimental Brain Research*, 11(4):376–386.
- Jacobsen, C. F. (1931). A study of cerebral function in learning. the frontal lobes. *Journal of comparative Neurology*, 52(2):271–340.
- Jacobsen, C. F. (1935). Functions of frontal association area in primates. *Archives of Neurology & Psychiatry*, 33(3):558–569.
- Jacobsen, C. F. and Nissen, H. W. (1937). Studies of cerebral function in primates. iv. the effects of frontal lobe lesions on the delayed alternation habit in monkeys. *Journal of Comparative Psychology*, 23(1):101.
- Johnson, A. and Redish, A. D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *Journal of Neuroscience*, 27(45):12176–12189.
- Jones, M. W. and Wilson, M. A. (2005). Theta rhythms coordinate hippocampal-prefrontal interactions in a spatial memory task. *PLOS Biology*, 3(12):e402.
- Jung, M. W., Qin, Y., McNaughton, B. L., and Barnes, C. A. (1998). Firing characteristics of deep layer neurons in prefrontal cortex in rats performing spatial working memory tasks. *Cerebral Cortex*, 8:437–450.
- Karlsson, M. P., Tervo, D. G. R., and Karpova, A. Y. (2012). Network resets in medial prefrontal cortex mark the onset of behavioral uncertainty. *Science*, 338(6103):135–139.
- Killcross, S. and Coutureau, E. (2003). Coordination of actions and habits in the medial prefrontal cortex of rats. *Cerebral Cortex*, 13(8):400–408.
- Kim, J., Ghim, J. W., Lee, J. H., and Jung, M. W. (2013). Neural Correlates of Interval Timing in Rodent Prefrontal Cortex. *Journal of Neuroscience*, 33(34):13834–13847.

- Kim, J., Jung, A. H., Byun, J., Jo, S., and Jung, M. W. (2009). Inactivation of medial prefrontal cortex impairs time interval discrimination in rats. *Frontiers in Behavioral Neuroscience*, 3:38.
- Kling, A. and Steklis, H. (1976). A neural substrate for affiliative behavior in nonhuman primates. *Brain, Behavior and Evolution*, 13(2-3):216–238.
- Klüver, H. (1933). Behavior mechanisms in monkeys.
- Klüver, H. and Bucy, P. C. (1939). Preliminary analysis of functions of the temporal lobes in monkeys. *Archives of Neurology & Psychiatry*, 42(6):979–1000.
- Kojima, S. and Goldman-Rakic, P. S. (1982). Delay-related activity of prefrontal neurons in rhesus monkeys performing delayed response. *Brain research*, 248(1):43–50.
- Kolb, B. (1984). Functions of the frontal cortex of the rat: a comparative review. *Brain Research*, 320(1):65–98.
- Konorski, J. and Lawicka, W. (1964). Analysis of errors by prefrontal animals on the delayed-response test. *The frontal granular cortex and behavior*. McGraw-Hill Book Co., New York, pages 271–294.
- Kruschke, J. (2010). *Doing Bayesian data analysis: A tutorial introduction with R*. Academic Press.
- Ma, L., Hyman, J. M., Phillips, A. G., and Seamans, J. K. (2014). Tracking Progress toward a Goal in Corticostriatal Ensembles. *Journal of Neuroscience*, 34(6):2244–2253.
- Macmillan, M. (2002). *An odd kind of fame: Stories of Phineas Gage*. MIT Press.
- Mardia, K., Kent, J., and Bibby, J. (1979). *Multivariate analysis*. Probability and mathematical statistics. Academic Press.

- Narayanan, N. S. and Laubach, M. (2008). Neuronal correlates of post-error slowing in the rat dorsomedial prefrontal cortex. *Journal of Neurophysiology*, 100(1):520.
- Narayanan, N. S. and Laubach, M. (2009). Delay activity in rodent frontal cortex during a simple reaction time task. *Journal of Neurophysiology*, 101(6):2859.
- Neafsey, E. J., Bold, E. L., Haas, G., Hurley-Gius, K. M., Quirk, G., Sievert, C. F., and Terreberry, R. R. (1986). The organization of the rat motor cortex: a microstimulation mapping study. *Brain Research*, 396(1):77–96.
- Nishino, H., Ono, T., Sasaki, K., Fukuda, M., and Muramoto, K.-I. (1984). Caudate unit activity during operant feeding behavior in monkeys and modulation by cooling prefrontal cortex. *Behavioural brain research*, 11(1):21–33.
- O’Keefe, J. (1976). Place units in the hippocampus of the freely moving rat. *Experimental Neurology*, 51:78–109.
- Papale, A., Stott, J. J., Powell, N. J., Regier, P. S., and Redish, A. D. (2012). Interactions between deliberation and delay-discounting in rats. *Cognitive, Affective, and Behavioral Neuroscience*.
- Passingham, R. E. (1986). Cues for movement in monkeys (*Macaca mulatta*) with lesions in premotor cortex. *Behavioral neuroscience*, 100(5):695–703.
- Paxinos, G. and Watson, C. (1996). *The Rat Brain in Stereotaxic Coordinates*. Academic Press, New York, third edition.
- Paxinos, G. and Watson, C. (1998). *The Rat Brain in Stereotaxic Coordinates*. Academic Press, New York, fourth edition.
- Peters, M. and Ploog, D. (1976). Frontal lobe lesions and social behavior in the squirrel monkey (*saimiri*): a pilot study. *Acta biologica et medica Germanica*.

- Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. pages 20–22.
- Poucet, B. (1997). Searching for spatial unit firing in the prelimbic area of the rat medial prefrontal cortex. *Behavioural Brain Research*, 84:151–159.
- Powell, N. J. and Redish, A. D. (2014). Complex neural codes in rat prelimbic cortex are stable across days on a spatial decision task. *Frontiers in Behavioral Neuroscience*.
- Pratt, W. E. and Mizumori, S. J. Y. (2001a). Neurons in rat medial prefrontal cortex show anticipatory rate changes to predictable differential rewards in a spatial memory task. *Behavioural Brain Research*, 123(2):165–183.
- Pratt, W. E. and Mizumori, S. J. Y. (2001b). Neurons in rat medial prefrontal cortex show anticipatory rate changes to predictable differential rewards in a spatial memory task. *Behavioural Brain Research*, 123:165–183.
- Preuss, T. M. (1995). Do rats have prefrontal cortex? The Rose-Woolsey-Akert program reconsidered. *Journal of cognitive neuroscience*, 7(1):1–24.
- Quirk, G. J. and Beer, J. S. (2006). Prefrontal involvement in the regulation of emotion: convergence of rat and human studies. *Current Opinion in Neurobiology*, 16(6):723–727.
- Quirk, G. J., Garcia, R., and González-Lima, F. (2006). Prefrontal mechanisms in extinction of conditioned fear. *Biological Psychiatry*, 60(4):337–343.
- Quirk, G. J., Russo, G. K., Barron, J. L., and Lebron, K. (2000). The role of ventromedial prefrontal cortex in the recovery of extinguished fear. *Journal of Neuroscience*, 20(16):6225.
- Redish, A. D. (1999). *Beyond the Cognitive Map: From Place Cells to Episodic Memory*. MIT Press, Cambridge MA.

- Reep, R. L., Corwin, J. V., Hashimoto, A., and Watson, R. T. (1987). Efferent connections of the rostral portion of medial agranular cortex in rats. *Brain research bulletin*, 19(2):203–221.
- Reep, R. L., Goodwin, G. S., and Corwin, J. V. (1990). Topographic organization in the corticocortical connections of medial agranular cortex in rats. *The Journal of Comparative Neurology*, 294(2):262–280.
- Rich, E. L. and Shapiro, M. (2009). Rat prefrontal cortical neurons selectively code strategy switches. *J Neurosci*, 29(22):7208–7219.
- Richter, C. P. and Hines, M. (1938). Increased spontaneous activity produced in monkeys by brain lesions. *Brain*, 61(1):1–16.
- Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., and Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*, pages 1–6.
- Rose, J. E. and Woolsey, C. N. (1948). Structure and relations of limbic cortex and anterior thalamic nuclei in rabbit and cat. *The Journal of Comparative Neurology*, 89(3):279–347.
- Rose, J. E. and Woolsey, C. N. (1949). Organization of the mammalian thalamus and its relationships to the cerebral cortex. *Electroencephalography and Clinical Neurophysiology*, 1(1-4):391–404.
- Rushworth, M., Hadland, K., Paus, T., and Sipila, P. (2002). Role of the human medial frontal cortex in task switching: a combined fmri and tms study. *Journal of Neurophysiology*, 87(5):2577–2592.
- Sakurai, Y. and Hirano, T. (1983). Multiple unit response in reward areas during operant conditioning reinforced by lateral hypothalamic stimulation in the rat. *Behavioural Brain Research*, 8(1):33–48.

- Sakurai, Y. and Sugimoto, S. (1985a). Effects of lesions of prefrontal cortex and dorsomedial thalamus on delayed go/no-go alternation in rats. *Behavioural brain research*, 17(3):213–219.
- Sakurai, Y. and Sugimoto, S. (1985b). Effects of lesions of prefrontal cortex and dorsomedial thalamus on delayed go/no-go alternation in rats. *Behavioural Brain Research*, 17(3):213–219.
- Sakurai, Y. and Sugimoto, S. (1986). Multiple unit activity of prefrontal cortex and dorsomedial thalamus during delayed go/no-go alternation in the rat. *Behavioural Brain Research*, 20(3):295–301.
- Schmitzer-Torbert, N. C. and Redish, A. D. (2004). Task-dependent spatial encoding in the dorsal striatum. *Society for Neuroscience Abstracts*.
- Seamans, J. K., Floresco, S. B., and Phillips, A. G. (1995). Functional differences between the prelimbic and anterior cingulate regions of the rat prefrontal cortex. *Behavioral neuroscience*, 109(6):1063.
- Siapas, A. G., Lubenov, E. V., and Wilson, M. A. (2005). Prefrontal Phase Locking to Hippocampal Theta Oscillations. *Neuron*, 46(1):141–151.
- Smith, K. S., Virkud, A., Deisseroth, K., and Graybiel, A. M. (2012). Reversible online control of habitual behavior by optogenetic perturbation of medial prefrontal cortex. *Proceedings of the National Academy of Sciences*, 109(46):18932–18937.
- Spaet, T. and Harlow, H. (1943). Problem solution by monkeys following bilateral removal of the prefrontal areas. ii. delayed reaction problems involving use of the matching-from-sample method. *Journal of Experimental Psychology*, 32(5):424.
- Steiner, A. P. and Redish, A. D. (2012). The road not taken: neural correlates of decision making in orbitofrontal cortex. *frontiers in Neuroscience*, 6.
- Swanson, L. (1992). Brain maps: structure of the rat brain.

- Terreberry, R. R. and Neafsey, E. J. (1983). Rat medial frontal cortex: a visceral motor region with a direct projection to the solitary nucleus. *Brain Research*, 278(1-2):245–249.
- Tolias, A. S., Ecker, A. S., Siapas, A. G., Hoenselaar, A., Keliris, G. A., and Logothetis, N. K. (2007). Recording Chronically From the Same Neurons in Awake, Behaving Primates. *Journal of Neurophysiology*, 98(6):3780–3790.
- Uylings, H. B. and van Eden, C. G. (1990). Qualitative and quantitative comparison of the prefrontal cortex in rat and in primates, including humans. *Prog Brain Res*, 85:31–62.
- Uylings, H. B. M., Groenewegen, H. J., and Kolb, B. (2003). Do rats have a prefrontal cortex? *Behavioural Brain Research*, 146(1-2):3–17.
- van der Meer, M. A. A., Johnson, A., Schmitzer-Torbert, N. C., and Redish, A. D. (2010). Triple dissociation of information processing in dorsal striatum, ventral striatum, and hippocampus on a learned spatial decision task. *Neuron*, 67(1):25–32.
- van der Meer, M. A. A. and Redish, A. D. (2010). Expectancies in decision making, reinforcement learning, and ventral striatum. *Frontiers in Neuroscience*.
- van der Meer, M. A. A. and Redish, A. D. (2011). Theta phase precession in rat ventral striatum links place and reward information. *Journal of Neuroscience*, 31(8):2843–2854.
- Vanderhasselt, M.-A., De Raedt, R., Baeken, C., Leyman, L., and Dhaenen, H. (2006). The influence of rTMS over the right dorsolateral prefrontal cortex on intentional set switching. *Experimental brain research*, 172(4):561–565.
- Vertes, R. P. (2004). Differential projections of the infralimbic and prelimbic cortex in the rat. *Synapse*, 51(1):32–58.

- Vertes, R. P. (2006). Interactions among the medial prefrontal cortex, hippocampus and midline thalamus in emotional and cognitive processing in the rat. *Neuroscience*, 142(1):1–20.
- von Monakow, C. (1895). Experimentelle und pathologisch-anatomische Untersuchungen über die Haubenregion, den Sehhügel und die Regio subthalamica, nebst Beiträgen zur Kenntniss früh erworbener Gross-und Kleinhirn-defecte. *European Archives of Psychiatry and Clinical Neuroscience*, 27(1):1–128.
- Yoon, T., Okada, J., Jung, M. W., and Kim, J. J. (2008). Prefrontal cortex and hippocampus subserve different components of working memory in rats. *Learn. Mem.*, 15(3):97–105.
- Zhang, K., Ginzburg, I., McNaughton, B. L., and Sejnowski, T. J. (1998). Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *Journal of Neurophysiology*, 79(2):1017–1044.
- Zilles, K. J. (1985). *The cortex of the rat: A stereotaxis atlas*. Springer-Verlag, New York.