NEURAL CORRELATES OF POSITION, PROGRESS, AND ACTION DURING REWARD-GUIDED SEQUENCE NAVIGATION IN THE RAT MEDIAL PREFRONTAL CORTEX

KUSHAAN GUPTA
Bachelor of Engineering, Devi Ahilya Vishwavidyalaya, 2020

A thesis submitted
in partial fulfilment of the requirements for the degree of

MASTER OF SCIENCE

in

NEUROSCIENCE

Department of Neuroscience
University of Lethbridge
LETHBRIDGE, ALBERTA, CANADA

NEURAL CORRELATES OF POSITION, PROGRESS, AND ACTION DURING REWARD-GUIDED SEQUENCE NAVIGATION IN THE RAT MEDIAL PREFRONTAL CORTEX

KUSHAAN GUPTA

Date of Defence: July 29, 2024

Dr. David R. Euston Thesis Supervisor	Associate Professor	Ph.D.
Dr. Artur Luczak Thesis Examination Committee Member	Professor	Ph.D.
Dr. Masami Tatsuno Thesis Examination Committee Member	Professor	Ph.D.
Dr. Ian Q. Whishaw Chair, Thesis Examination Committee	Professor	Ph.D.

Dedication

To my family and friends, for your unwavering support, guidance, and encouragement.

Abstract

The complexity surrounding the seamless execution of skilled action sequences during sparse rewards necessitates an organization of behavior into modular, temporally extended sub-policies. The rodent Medial Prefrontal Cortex (mPFC) has emerged as a critical brain region implicated in orchestrating such action plans over extended time frames. Previous research has shed light on the mPFC's mechanistic contributions to the control of behavior through investigations into its numerous attributed roles, encompassing working memory, spatial representation, decision-making, and task-specific generalization. We investigated the neural correlates of position, task-specific features, and actions within the rat mPFC during a behaviorally elaborate, reward-guided sequence navigation paradigm. We identified a significant proportion of mPFC neurons exhibiting a robust task-specific code for tracking animal's progress between rewards, characterized most prominently by modulation departures from and approaches to reward zones. Furthermore, initial population-level analyses suggested a representation of prospective non-immediate actions, encoded as turn directions, extending beyond immediate temporal horizons. However, subsequent in-depth examination revealed that these initial findings could not hold up to predict future actions at single-trial level. Instead, the apparent prospective coding of turn directions seemed to be confounded by preparatory motor behaviors within the navigational environment. Consequently, these results call for a reassessment of the influence of various sensorimotor factors on mPFC activity during goal-directed decision-making, challenging the prevailing methodologies used to infer cognitive processes within this region.

Acknowledgments

I carry within myself the reverberation of experiences that both deeply humble and fuel my passion for understanding the brain. These are the moments with the people I have encountered along my journey, who shared their knowledge, experience, and wisdom with me. They made my problems their own, and for that, I am forever indebted to them. Whether in academia or life outside of it, I could rest easy, knowing that they held in their minds, my best interests—often more than I did: *Yagika Kaushik, Jiajie Liang, Kartik Iyer, Alisa Leshchenko, HaoRan Chang, Ritwik Das, Siyu X. Wang, Samsoon Inayat, Sara Citron, Garima Virmani, Arnab Nandi, Aatmika Barve, Soumen Das* and *Swananda V. Marathe*. To my advisor, *David R. Euston* for our discussions as we worked through the incessant hurdles of this scientific expedition and for the freedom to explore my ideas. To *Ian Q. Whishaw* for his teachings that deeply shaped my scientific thinking.

The experimental data presented in this thesis was designed and collected in the laboratory of *Bruce L. McNaughton* by *David R. Euston* and colleagues. I would like to thank *Artur Luczak* and *Masami Tatsuno* for their helpful feedback throughout the development of this work. I would also like to extend my gratitude to *Michael Kindley* and *Kevin Halili* for their contributions to discussions of the extant literature on the functional roles of mPFC, which significantly enriched the literature review component of this thesis.

Contents

List of Tables

List of Figures

List of Abbreviations

ACC	Anterior Cingulate Cortex
AIC	Akaike Information Criterion
BIC	Bayesian Information Criterion
CA1	Cornu Ammonis 1
CA3	Cornu Ammonis 3
CELU	Continuously Differentiable Exponential Linear Unit
Cg1	Cingulate Cortex Area 1
CS	Conditioned Stimulus
dlPFC	Dorsolateral Prefrontal Cortex
DMT	Dorsomedial Thalamus
DNN	Deep Neural Network
HPC	Hippocampus
IL	Infralimbic Cortex
IQR	Interquartile Range
LED	Light Emitting Diode
M2	Secondary Motor Cortex
MAE	Mean Absolute Error
MFB	Medial Forebrain Bundle
MLP	Multi-layer Perceptron
mPFC	Medial Prefrontal Cortex
MRI	Magnetic Resonance Imaging
MSE	Mean Squared Error
OFC	Orbitofrontal Cortex
PC	Principal Component
PCA	Principal Component Analysis
PFC	Prefrontal Cortex
PL	Prelimbic Cortex
PV	Population Vector
RMSE	Root Mean Squared Error
SGD	Stochastic Gradient Descent
SHAP	SHapley Additive exPlanations
SVM	Support Vector Machine
TTL	Transistor-Transistor Logic
vmPFC	Ventromedial Prefrontal Cortex

Chapter 1

Introduction

A tiger, dead in her tracks, awaits to pounce on a lone deer, quenching his thirst at a serene watering hole. She advances steadily, while the deer remains unaware. Suddenly, a misstep on a fallen branch betrays her presence. The deer startles, his head snapping up in alertness. With a burst of speed, the deer retreats into the safety of surrounding foliage. Acknowledging the lost opportunity, the tiger swishes her tail in frustration. The realization dawns that the chase, at this moment, will not be productive. The predator, undeterred and resilient, pauses to reassess her strategy. Aware that thirst will drive the deer to the watering hole again, contemplation gleams in her golden eyes as she plots a stealth attack from a new, strategic vantage point. Though it may not offer the same advantage as before, the element of surprise could prove worthwhile.

Let us examine the tiger's thought process step-by-step as the series of events unfolded while considering the cognitive functions required by these sub-processes. Observing her prey triggered a sudden spike in the tiger's *situational awareness*, upon which most parts of her brain were *reconfigured*, like train tracks switching to a new direction, towards the goal of hunting the deer.

Based on her action repertoire, the situation permitted her three affordances: (1) pounce directly towards the deer, (2) get surreptitiously closer to the deer, or (3) do nothing and walk away. Deciding among these affordances requires *generating predictive inferences about the consequences of each action*. In order to achieve it, the tiger must *represent the likely outcomes and potential risks associated with each action*, hence engaging in *value-*

based decision-making. This involves simulating familiar or novel scenarios in her mind based on *recalled memories of past experiences and learned knowledge*, processes that heavily rely on *working memory*. For instance, if she were to pounce directly, she must predict whether the deer is close enough for a successful catch or if it might flee before she reaches it, resulting in wasted energy and missed opportunity. Stalking the deer requires assessing her stealth abilities, the terrain, and the deer's alertness to ensure she can get close enough without being detected. Walking away or doing nothing might be considered if the conditions are not favorable for a successful hunt, such as if the deer is too alert or if the environment doesn't provide enough cover for stalking.

In addition to generating predictive inferences, the tiger must exercise *flexible and intentional inhibitory executive control* to suppress all but one option in response to the situational demands and internal goals, which, in this case, was the second option. Despite the temptation to pounce immediately, she opts to stalk the deer, predicting that a direct approach from her initial distance to the prey would likely fail.

Upon her misstep, the tiger's cognitive processes shifted from execution to reassessment as a rational agent. The tiger understood the causal relationship between the sound of the fallen branch breaking under her foot and the subsequent alerting of the deer. This realization necessitated her swift adaptation of strategy through *cognitive flexibility* and *contingency planning*. The auditory input immediately overrode her initial plan and necessitated a rapid recalibration of her approach. As the deer fled, the tiger realized that she must inhibit any impulsive reaction to chase immediately, as it would likely be futile given the deer's head start. Instead, she exercised top-down inhibitory executive control to suppress this impulse, indicative of advanced cognitive functions such as *self-regulation*.

The tiger then transitioned to a phase of *strategic planning*, involving a detailed assessment of her surroundings and the identification of a novel ambush site. She contemplated the likely return of the deer to the watering hole, driven by its need for water. In plotting her new approach, the tiger simulated potential outcomes in her mind. She envisioned dif-

ferent scenarios: the deer approaching from various directions, the availability of cover for a stealthy approach, and the timing required to execute a successful pounce. This mental simulation, supported by working memory, was rooted in her past hunting experiences and learned knowledge about her prey's behavior, integrated with knowledge gained from this recent experience.

To summarize, the tiger's thought process involved a complex interplay of sensory, cognitive, and motoric processing. Together, these processes enabled her to execute and adapt her behavior dynamically in response to environmental changes, ensuring her continued success as an apex predator. The Prefrontal Cortex (PFC) implicated in orchestrating all of these processes. It is hard to imagine that such a diverse body of cognitive functions can fit under a unifying framework of a general theory of PFC function. Nevertheless, several attempts have been made to consolidate the diverse functional roles of the PFC into a monosemic functional ascription.

The ambit of this thesis is to investigate the fundamental principles underlying the role of the rat Medial Prefrontal Cortex (mPFC) in orchestrating behavior in a goal-directed sequence navigation paradigm. We evaluate the existing theories of PFC function and the evidence supporting and challenging them in the context of rodent mPFC.

1.1 General theories of prefrontal cortical function

1.1.1 Complementary learning systems

Early observations, dating back to Ribot (1881), documented the increased susceptibility of recent memories to disease and injury compared to remote memories (Burnham, 1903). Subsequent research in healthy individuals revealed the phenomenon of interference, where learning is impaired due to conflict between new and previously memorized information, or when a familiar stimulus requires a response different from the previously learned one (Bergstrom, 1893; Greenberg & Underwood, 1950; Underwood, 1957). Collectively, these findings highlighted a contrast between a stable memory system and a fragile

one, and constraints in the ability to transform the fragile memory into a stable one.

Building on this, Müller and Pilzecker (1900) proposed that memory traces undergo a process of consolidation, rendering them resistant to disruption over time. This consolidation process was hypothesized to consist of two complementary systems: a labile storage, in which memories are vulnerable to interference, and a stable storage, during which memories are resistant to disruption. Over a significant period, a growing body of research in both laboratory and clinical settings has provided evidence supporting the distinction between labile and stable memories. Within the connectionist framework, McClelland et al. (1995) proposed the *standard model of memory consolidation* to describe this process. This theoretical neurocomputational model posits the localization, mechanisms, and computational advantages of the complementary learning systems underlying memory consolidation. Specifically, they posited Hippocampus (HPC) as a fast learning system that rapidly encodes new memories through fast weight changes in the synaptic connections, and the neocortex as a slow learning system that gradually integrates these memories through slow weight changes.

McClelland et al. (1995) proposed that complementary learning systems arose out of the need to overcome the problem of *catastrophic interference*. The solution led to the simultaneous maintenance of two distinct memory systems, of which one was more stable (the slow learning system with stable memories), but with a compromise in the fragility of the other system (the fast learning system with labile memories). They theorized that fast learning systems would reinstate the memory traces numerous times to the slow learning systems in order to consolidate the memory. During such reinstatements, the slow learning systems would change very gradually, as if they had a slow learning rate. They argued that such slow learning rates developed as a consequence of the objective to maintain a stable memory trace. Moreover, they proposed that this gradual integration happened in an interleaved manner, in which new information was reactivated in the slow learning systems while being interleaved with related domain information that was pre-existing in

the slow learning system (Hinton, 1989; Rumelhart, 1990). This theory suggests a data-driven, bottom-up generalization mechanism. This process involves extracting statistical regularities from existing memories, contrasting with a top-down approach that constructs conceptually distinct superordinate categories. The slow learning system would scale the memories based on the connotations of previously stored representations, thereby forming a basis for compositional knowledge (Kurth-Nelson et al., 2023).

While McClelland et al. (1995) used HPC and neocortex as general terms, they also suggested that regions beyond the HPC and neocortex could also act as fast and slow learning systems, such as the amygdala as an alternative fast learning system, and the basal ganglia and the cerebellum as slow learning systems.

1.1.2 Prefrontal cortex as a goal-based reconfigurator

To address the diverse functional roles of the PFC and consolidate them into a unified functional ascription, Miller and Cohen (2001) proposed that by providing the goal information, PFC reconfigures the stimulus-response mappings in other areas of the brain through an actively maintained control, aligning them with contextually appropriate representations of goals or task rules. For instance, when the tiger from our previous example noticed the deer, her PFC would have been responsible for generating the goal of hunting the deer, thereby reshaping the stimulus-response associations in the other regions of the brain to align with the goal of hunting the deer.

While primarily grounded in studies of the primate Dorsolateral Prefrontal Cortex (dlPFC), this theoretical framework has also informed numerous investigations of rodent mPFC function (Sakurai & Sugimoto, 1986; Orlov et al., 1988; Birrell & Brown, 2000; Duan et al., 2015; Ito et al., 2015; Guise & Shapiro, 2017). These line of investigations were supported by the evaluation of the anatomical evidence and the effects of mPFC lesions on various behavioral tasks, by Kolb and Whishaw (1990) suggesting potential parallels between the cognitive functions of the rodent mPFC and the primate dlPFC.

1.1.3 Somatic marker hypothesis

Bechara and Damasio (2005) theorized that the mPFC, specifically the Ventromedial Prefrontal Cortex (vmPFC), operates based on an internal frame of reference driven by somatic state signals and subjective values. According to their Somatic Marker Hypothesis, the mPFC guides decisions by anticipating emotional outcomes, driving changes in bodily states according to those predictions. This, in turn, generates automated responses to situations, thereby overriding extensive deliberation during decision-making.

It is important to note that Bechara and Damasio (2005) characterize value-based decision-making contingent on the vmPFC as a non-deliberative (Pavlovian) process, which they refer to as the primary decision-making process. According to their theory, post-decision rationalization merely gives the agent the impression that the decision was a result of intentional deliberation. Although these theories of PFC function were formed from evidence from humans and non-human primates, subsequent research has suggested that a similar framework may also apply to the study of rodent mPFC function (Euston et al., 2012).

1.1.4 Medial Prefrontal Cortex as a memory-guided decision system

A specific theory regarding the function of the mPFC, expanding on the Somatic Marker Hypothesis, was proposed by Euston et al. (2012). This theory integrates evidence from studies on both primates and rodents, suggesting that the mPFC receives contextual and event-related information from upstream regions and associates it with memory-supported action representations encoded within itself. Euston et al. (2012) suggest that the overarching goal of the mPFC is to make decisions, and that other cognitive phenomena associated with it, such as mnemonic capabilities, subserve its objective of decision-making. Within this theoretical framework, action information results from the dynamic interaction between multiple memory systems in the brain, which operate on different timescales. Initially, the rapidly acquired input-output mappings in the mPFC are supported by the HPC, but over time, they become independent.

Primarily, under this view, the mPFC may act in response to activations from upstream regions to produce decisions, as opposed to maintaining active control to direct stimulus-response associations in other areas, as suggested by Miller and Cohen (2001). These theories need not be mutually exclusive. Numerous studies in rodent reward-guided navigation fall under the umbrella of this theory. For instance, literature suggests that the HPC feeds spatial contextual information to activate representations of rules in the mPFC (Jung et al., 1998), motivational significance (Pratt & Mizumori, 2001), strategies (Rich & Shapiro, 2009), and trajectories (Ito et al., 2015).

1.1.5 Medial Prefrontal as a conflict resolver for multiple decision systems

Focusing specifically on the theory of rodent mPFC function, McLaughlin et al. (2021) propose that the mPFC functions as the locus of interaction for neurally distinct, multiple decision-making systems. They suggest that the mPFC arbitrates conflicts among these systems, thereby enabling the selection of appropriate single actions. This theory proposes multiple decision systems as independent yet hierarchically interacting entities. While a comprehensive evaluation of this framework is beyond the present study's purview, it aligns with Euston et al. (2012) in suggesting that the mPFC encodes action-related information potentially derived from interactions among multiple decision-making processes.

1.2 Investigations into specific roles of rodent mPFC

While extensive research had previously explored the effects of rodent mPFC lesions on various behavioral tasks, the correlation between rodent mPFC extracellular neurophysiology and behavior did not gain traction until 1986. Early studies in this area were inspired by existing literature on the effects of mPFC damage in humans, lesion and electrophysiology studies in non-human primates, lesion studies in rodents (Larsen & Divac, 1978), as well as electrophysiology studies in rabbits (Orona & Gabriel, 1983).

1.2.1 Role in working memory

The seminal studies on rodent mPFC electrophysiology often employed delayed response paradigms to investigate the relationship between neural activity, working memory, and behavior (Sakurai & Sugimoto, 1986; Orlov et al., 1988; Batuev et al., 1990; Orlov et al., 1991). These investigations sought to identify persistent neural activity analogous to that observed in the primate PFC (Fuster & Alexander, 1971), supporting the concept of active information maintenance as proposed by Miller and Cohen (2001) within the context of goal-directed behavior.

Following observations of impaired efficiency in response switching on a delayed response inhibition paradigm with a go/no-go alternation task in animals with lesions of the mPFC (Sakurai & Sugimoto, 1985), albeit less severe than those with Dorsomedial Thalamus (DMT) lesions, Sakurai and Sugimoto (1986) investigated the multiple-unit activity in the mPFC during the same task. They found increased firing rates in some PFC neurons during the delay period before the response, which differentiated the animal's choice to either go or not-go. This firing pattern was not unique to PFC neurons, as similar firing patterns were observed in some neurons of the DMT. Notably, these firing patterns in both PFC and DMT did not indicate the correctness of the choices, suggesting a prospective rather than a retrospective code. A confounding factor in their task design was the immediate transition from the end of one trial to the delay period of the next, potentially limiting the rats' time for reflection. Additionally, some DMT neurons exhibited selective firing during the delay period preceding correct choices, regardless of whether the choice was to go or not-go.

The series of studies using delayed matching-to-sample tasks (Orlov et al., 1988; Batuev et al., 1990; Orlov et al., 1991) where animals were locked in a transparent chamber before the presentation of a cue interpreted single-cell neuronal responses before the presentation of a cue as reflective of an expectation-like predictive signal. However, this interpretation was not directly tested against motoric confounds. To justify the claims, a complex analysis was required. This should have involved correlating neural firing prior to the cue (one

suspected of exhibiting the expectation predictions) with the animal's expectation for a particular direction, based on whether a reward had been experienced in that direction in the preceding trial. It is likely that due to the technological limitations at the time, such as the inability to perform multi-trial recordings and track the same neurons across all trials, achieving this level of analysis may not have been feasible. While these studies revealed diverse cellular specializations for distinct phases of the delayed-response task, with some neurons exhibiting activity unique to single task periods and others exhibiting conjunctive modulation across potentially non-adjacent periods, the lack of verification of multi-trial consistency of such neuronal firing patterns hinders confirmation of their role in maintaining delay-related working memory activity. Furthermore, their paradigm did not control for or account for potential confounds arising from spontaneous behavior-related neuronal modulation.

Later studies on rodent prefrontal correlates of working memory became primarily spatial, such as the delayed alternation task (Jung et al., 1998; Baeg et al., 2003; Euston & McNaughton, 2006). Of note, the delay period in these tasks usually involved traversal along a central arm in a maze after which a choice was supposed to be made. Activity during these periods has been associated with prospective tracking of probable trial outcomes (Hyman et al., 2012), encoding of reward-related feedback from preceding trials (Laubach et al., 2015), and spatial representations (Jung et al., 1998; Euston & McNaughton, 2006; Laubach et al., 2015). While Baeg et al. (2003) identified mPFC neurons exhibiting persistent activity during delay periods that predicted choices, they also reported the presence of non-continuously firing neurons. Notably, the authors demonstrated that highly trained rats could accurately predict previous and future choices without relying solely on the traditionally conceived working memory mechanism of persistent neural firing within the mPFC. Although these findings were claimed to support the notion of multiple dynamic representations of working memory within the mPFC, they can also be used to challenge the traditional view of working memory as exclusively reliant on persistent neural activity. The

fact that past and future choices can be predicted using any type of activity in the mPFC challenges the validity of methods conventionally employed to identify working memory-related neural correlates within the mPFC. It is possible that the persistent activity observed in the mPFC during delay periods may not be a direct correlate of working memory maintenance, but rather a reflection of some other confounding factor. Moreover, several studies (Euston & McNaughton, 2006; Cowen & McNaughton, 2007) further suggested that positional and sensorimotor factors confound the activity of mPFC neurons in spatial delayed alternation tasks. These studies collectively raised the fundamental question of the extent to which movement- or location-related factors underlie neural activity recorded during task performance.

The role of the rodent mPFC in working memory remains a subject of considerable debate, largely due to the inherent conflation of spatial navigation and working memory demands within experimental paradigms. Some studies suggest that the mPFC may not be directly involved in working memory maintenance, but rather in encoding and processing spatial information.

1.2.2 Memory consolidation in medial prefrontal cortex

The neural correlates of reinstatement of experience (or memory retrieval) have been extensively investigated, with a primary focus on the HPC. Key signatures of hippocampal reinstatement: reactivation (Wilson & McNaughton, 1994; Pavlides & Winson, 1989) and replay (Buzsáki, 2015) have been well-documented and are evidenced in supporting transfer of hippocampal information to the neocortex (Hoffman & McNaughton, 2002; Ji & Wilson, 2006; Khodagholy et al., 2017).

Consistent to these findings, the mPFC has shown evidence of temporally compressed sequential reactivation (Euston et al., 2007) and ripple oscillations (Shin & Jadhav, 2024). While the sleep replay is theorized to be crucial for memory consolidation (Wilson & McNaughton, 1994), mounting evidence suggests that awake replay plays a critical role in eval-

uating past and future choices (Pfeiffer & Foster, 2013; Singer et al., 2013; Xu et al., 2019). This hypothesis is further supported by studies demonstrating that inhibiting hippocampal sharp-wave ripples, periods known to be associated with heightened replay activity, during wakefulness leads to deficits in spatial task performance (Jadhav et al., 2012).

Recent work by Shin and Jadhav (2024) found that the rodent mPFC ripple oscillations can be detected both independently and simultaneously in both mPFC and HPC. The study revealed a pronounced PFC ripple-mediated suppression of Cornu Ammonis 1 (CA1) reactivation for neurons exhibiting higher reactivation. Contrary to complementary learning systems theory, these findings suggest that PFC ripples exert a top-down inhibitory influence on hippocampal reactivation rather than facilitating coordination. This top-down inhibition may serve to control the flow of information from the HPC to the mPFC, thereby regulating the consolidation of memories in the mPFC.

This bidirectional flow of information between the HPC and mPFC is further supported by the findings of Guise and Shapiro (2017). Their results suggested that during learning the mPFC's representations of task rules may help to minimize the interference in the HPC (or other fast learning systems). This hypothesized teacher-student dynamic between the mPFC and the HPC was demonstrated to be critical for hippocampal pattern separation and faster behavioral responses.

1.2.3 Position-correlated representations in medial prefrontal cortex

Despite the general columnar architecture of the neocortex implying functional homogeneity across its various areas (Edelman et al., 1982; Mountcastle, 1957), substantial evidence points towards regional specializations of functional heterogeneity encompassing the entire gamut of sensory, motor, and cognitive processing modalities (Kaas, 1987). A common thread tying these heterogeneous areas together appears to be the spatial coding in many neocortical areas. Notably, this phenomenon extends to large swaths of the dorsal cortex, encompassing both sensory and higher-order areas (Esteves et al., 2021). Spatial

tuning has been documented across diverse rodent neocortical areas, including the mPFC (Euston & McNaughton, 2006; Mashhoori et al., 2018; Kaefer et al., 2020; Sauer et al., 2022), secondary motor cortex (Chang et al., 2023), posterior parietal cortex (Harvey et al., 2012), retrosplenial cortex (Mao et al., 2017), visual cortex (Fiser et al., 2016; Saleem et al., 2018), somatosensory cortex (Long & Zhang, 2021), and Orbitofrontal Cortex (OFC) (Wikenheiser et al., 2021).

Aforementioned studies have predominantly investigated spatial tuning in the neocortex using reward-guided navigation tasks. These tasks necessitate the agent's ability to track its own spatial state or that of relevant objects for successful performance. An important caveat of most cortical "spatial tuning" studies in rodents is the use of highly constrained behavioral paradigms. These environments introduce potential confounds associated with movement patterns (e.g., turning behavior at path ends, changes in speed), reward consumption behaviors (e.g., chewing, licking), and other factors unrelated to spatial representation. These confounds may lead to the semblance of a spatial code or a generalized spatial code where there is none. This phenomenon stands in stark contrast to the well-established spatial tuning observed in hippocampal place cells, which are generally invariant to behaviors, aside from head direction, and are anchored to external spatial cues. However, Sauer et al. (2022) claims to have demonstrated that the mPFC in mice can exhibit environmental context-specific spatial encoding even in the absence of rule or reward-driven locomotion. This finding suggests a broader role for the mPFC in spatial processing beyond reward-guided navigation, potentially implicating place encoding as a more universal function. In contrast, many other studies have found no evidence of spatial representations at the single-unit level (Jung et al., 1998; Kaefer et al., 2020; Poucet, 1997), highlighting the need for further investigation to reconcile these seemingly discrepant observations. Sauer et al. (2022) identified two key factors that warrant further investigation to resolve the ambiguity surrounding the role of mPFC in spatial encoding. The first involves potential confounding effects arising from modulations in mPFC neuronal activity in response to non-spatial

behavioral variables, such as grooming or rearing behaviors. The second factor pertains to variations in the spatial complexity of the environments that animals spontaneously traverse, which in their case was a highly reduced setting of a one-dimensional circular track.

Further supporting the notion of spatial tuning as a default function of the mPFC is its anatomical and functional interconnectivity with the HPC. The hippocampal formation, with its spatially selective cells, is known to form the fundamental basis of cognitive maps. It was previously known that the ventral CA1 neuronal population in the HPC projects to the ventral PFC (Hoover & Vertes, 2007). In contrast, from their single-unit electrophysiological studies, Sauer et al. (2022) found a stronger spatial signal in the dorsal mPFC compared to the ventral region, which challenged this notion. This was a surprising result as the connectomics studies hypothesize a completely opposite spatial information encoding gradient in the mPFC. In an attempt to reconcile this, Draguhn (2022) suggested that mPFC spatial information may not be as hippocampal-dependent as previously expected, and that there may be alternative origins of spatial information in the mPFC, potentially from entorhinal or perirhinal inputs. An alternative explanation that accommodates the possibility of hippocampal-dependent spatial tuning of the mPFC emerges from recent work by Messanvi et al. (2023). Their findings in mice revealed a parallel pathway originating from the dorsal Cornu Ammonis 3 (CA3) region of the HPC, projecting to both the ventral and dorsal PFC. Notably, the previously documented projections from the ventral CA1 were found to anatomically relay these signals to both the ventral and dorsal PFC, with the dorsal lateral septum serving as an additional (though weaker) relay projecting to the ventral PFC. Both dorsal and ventral PFC have the capacity to influence activity in dorsal CA1 and ventral CA1, respectively, through relay pathways such as midline thalamic nuclei and the entorhinal cortex (Witter & Groenewegen, 1984; Prasad & Chudasama, 2013). Although the functional role of these connections is not yet established, this inter-regional connectivity may be prominent for the communication of spatial information to the mPFC.

The mPFC has also been implicated in the decision-making process, specifically for

flexibly adapting behavioral strategies following a rule change. Initial evidence for this role emerged from studies in humans and has since been corroborated in rats, where lesions to the mPFC result in behavioral inflexibility and deficits in rule-switching performance. The integration of memory and spatial processing functions in the mPFC is theorized to ultimately subserve the overarching function of decision-making (Euston et al., 2012).

Several studies proposed a specific mechanism for this integration, suggesting that the mPFC constructs a mental representation of the task space (Lapish et al., 2008; Durstewitz et al., 2010; Mashhoori et al., 2018). This "task space" is hypothesized to be a multi-dimensional construct that encompasses both spatial and choice-related information. The constituents of these task representations are hard to pin down as they seem to be used differently in different behavioral paradigms. Nonetheless, the task space is defined as a multiplexed representation of the choices, outcomes, and task contexts that may be embedded in a spatial framework. This mental model, they argue, can be navigated internally (Lapish et al., 2008; Mashhoori et al., 2018).

When Tolman (1948) proposed the concept of *cognitive maps*, he envisioned two utilities of such a construct: the map internally represents the spatial layout of the environment and is mentally navigable in order to optimize the behavioral policy of the agent (such as finding shortcuts in the environment). Evidence of both of these functions was claimed by Mashhoori et al. (2018) for rodent mPFC. This is not to diminish the importance of the brain regions in the HPC proper, which are theorized to form the basis of the cognitive maps. However, we speculate that the purported generalized spatial code in the mPFC found by Mashhoori et al. (2018) may form the highest level of hierarchical spatial abstraction. This cognitive construct could function as a computationally efficient data structure for mental search, analogous to the advantage of organizing files alphabetically over chronologically (Chalmers et al., 2016). Consequently, under this theoretical framework, the dynamics of spatial information decoded from the mPFC can be employed as a lens of high temporal and spatial resolution for investigating the cognitive processes underlying memory-guided

decision-making. Systematic deviations of decoded position can indicate recall of multimodal information regarding the value of distant locations within the task space. This information can be used to guide future decisions, potentially by biasing the animal's choice towards or away from specific locations based on the value of the information retrieved from those locations.

1.2.4 Decision-making

Relationship between spatial, working memory, and action representations in mPFC

Decisions do not emerge arbitrarily; rather, they are shaped against the backdrop of our memories. In natural environments, however, an animal seldom encounters environmental states that are an exact replica of a previous experience whose optimal action is precomputed in memory. While a conditioned reflex strategy may suffice for the survival of simple organisms such as the Venus flytrap, the complex evolutionary pressures faced by animals necessitate a more adaptable behavioral repertoire. Organisms exhibiting greater decision-making flexibility have demonstrated a selective advantage. Adaptive decision-making is characterized by the dynamic selection of suitable actions in response to continually evolving internal and environmental contingencies. To achieve this, animals must learn to meaningfully understand the causal links relating events to each other, encompassing not only factual information about their environment but also their own actions. Subsequently, they must distill such information into pliable neural representations of varying degrees of abstraction. The pliable representations, which may be referred to as cognitive maps or schematic representations (schemas), are structured bodies of prior knowledge that capture common patterns across related experiences in physical and social spheres and are used for particular functions. These mental models are essential for inferring the outcomes of actions, contemplating fictive scenarios, reasoning and planning in novel situations, and organizing knowledge based on experience. Numerous correlational and lesion studies conducted in primates provide evidence supporting the PFC as a central mediator of the

capacity to anticipate situations and predict how different actions may influence outcomes (Euston et al., 2012). Model-based decision systems, such as procedural and deliberative decision-making systems, rely on the evaluation of such constructed representations to predict the consequences of potential outcomes and select actions (McLaughlin et al., 2021).

In animals, decision-making or the devisement of strategies underlying behavior is mediated by the interplay of multiple decision systems that are both neurally and computationally distinct. Contingent on environmental pressures, these systems may come into conflict, complicating the convergence to singular actions. Evidence suggests that, in rodents, the mPFC may serve as a central arbitrator between different decision systems (McLaughlin et al., 2021). This suggests that regardless of the underlying decision-making system proposing the action, information about the finalized action must be present in the mPFC. By monitoring neural activity in the mPFC, we can discern information about the actions performed behaviorally driven by a particular action system. This information can then be utilized to deduce the actions carried out by an alternative decision system. The monitoring of mPFC activity thus serves as a unified method to infer the outputs of various decision-making systems via a common framework. The present study aims to test this theorized role of rodent mPFC by hypothesizing that during sequential decision-making tasks involving spatial navigation, the mPFC would represent fictive (imaginary) states when procedural (habitual) decision-making strategies are naturally disrupted, facilitating the cognitive exploration required to resolve conflicts between procedural and deliberative decision systems.

Inspired by the spatial discrimination deficits shown in rodents with mPFC lesions (Kolb, 1984; Kolb & Whishaw, 1990) and the discovery of ventral CA1 projections to ventral Prelimbic Cortex (PL) (Swanson, 1981; Ferino et al., 1987), Poucet (1997), using a random foraging paradigm similar to the one used to discover spatially tuned neurons in the hippocampal formation, did not find any spatially consistent or head-direction specific neuronal firing pattern in rat mPFC (specifically, PL) while performing a random pellet-chasing task.

Poucet (1997) argued that the apparent spatial selectivity observed in some firing rate maps arises from the co-occurrence of specific behaviors in specific locations, rather than from intrinsic spatial tuning independent of behavior. Later, Jung et al. (1998) investigated the neural correlates of working memory within the mPFC by employing a novel approach involving simultaneous recordings from a large population of neurons across three distinct spatial maze tasks. They found that mPFC neurons showed diverse (either single or multiple) behavioral correlates. They found only a handful of neurons that seemed to indicate any spatial specificity or delay-related sustained firing indicative of working memory. Jung et al. (1998) reported that nearly all mPFC units active on the eight-arm maze exhibited firing activity on all of the eight arms. This generalized neuronal activity might indicate a similarity in actions performed across these distinct paths, consistent with the hypothesis of Poucet (1997).

To explain the function of hippocampal ventral CA1 projections to the mPFC in the context of their findings and to account for the observations suggesting that HPC might model response components or behavioral contingencies (though these correlates were found on very constrained tasks, contrary to the one used by them), they theorized that HPC spatial information provides the framework in which action information is embedded to form something equivalent to a habit, like a sparse spatial representation of a behavioral policy. They assumed that the action information is also transmitted to the mPFC via some afferent connections called "detonator motor synapses," which would trigger an action when the spatial context is simultaneously provided by the HPC. If an action is highly repetitive at a certain location, then HPC presynaptic neurons would strengthen their functional coupling with mPFC neurons. With associative plasticity induced by adequate experience, the provision of hippocampal spatial information input alone becomes sufficient to initiate the action (Jung et al., 1998).

While the existence of "detonator motor synaptic input" to mPFC is unconfirmed and arguable, mounting evidence suggests that hippocampal-coordinated, if not hippocampal-

triggered, activation of mPFC neuronal representation of actions is a plausible mechanism. Miller and Cohen (2001) suggested that it may be suboptimal to actively maintain a persistent representation of a goal in some scenarios. Instead, they speculated that rapid learning systems, such as the HPC, could rapidly encode the association between the goal representations within the PFC and features of the circumstance under which the goal should be evoked, i.e., the context. Although Miller and Cohen (2001) posited this model based on evidence from investigations into primate cognition, this is in line with the theoretical model proposed by Jung et al. (1998), where rodent mPFC neurons encode the action information in the context of the spatial information provided by the HPC. Euston et al. (2012) expanded upon the concept of contextually activated action representations in the mPFC. The authors adopted a broad definition of context, encompassing spatial, temporal, and environmental or task-related information integrated within the mPFC to guide decisions. For instance, this framework could account for the ability to selectively retrieve the action of opening an umbrella from memory while walking out of the house, but only when it is a rainy day.

1.3 Thesis objectives

Different avenues of investigation into the role of the rodent mPFC seem to converge on similar conclusions. The studies reviewed here demonstrate a pivotal role of the positional and action representations in the mPFC in the context of reward-guided spatial navigation and decision-making. However, the precise mechanisms underlying the integration of these representations and their role in guiding behavior remain unclear.

In this work, we aim to investigate these roles in a complex spatial sequence navigation task that requires the integration of spatial, working memory, and action representations. Given previous evidence suggesting multiplexed task-relevant information within the mPFC (Mashhoori et al., 2018), we hypothesized that location information could be reliably decoded from mPFC population activity.

However, our positional decoding analysis failed to identify robust spatial representations within the mPFC population code. Consequently, we conducted a data-driven exploration of mPFC neural activity during task performance. Our results indicate that the precise motor control acquired through extensive task training may have obscured underlying spatial and predictive action representations within the mPFC. These findings align with those of Poucet (1997), Euston and McNaughton (2006), and Cowen and McNaughton (2007), challenging the hypothesis of a spatially organized code for actions and task-related information within the mPFC.

Chapter 2

Experimental Methods

Description of experimental methods were adapted from the original manuscript (Euston & McNaughton, 2006) and are described here in detail to provide context for the ensuing chapters. The subsequent two chapters will present specific analyses conducted on data collected using these methods.

2.1 Subjects

Four adult male Brown Norway/Fischer 344 hybrid rats, also abbreviated as BNF344F1 or BNF3F1 [for Brown Norway (BN) x Fischer 344 (F344) hybrid cross with BN (BN/Rij) mothers], were accommodated in Plexiglas home enclosures and maintained on a reversed light cycle (lights on: 10 p.m. to 10 a.m.). Limited handling of the rats began several months prior to surgery. Following the surgical implantation, the rats were handled and weighed on a daily basis. At the time of the surgical procedure, the rats were aged between 7 to 9 months and weighed between 350 g to 400 g. Training sessions and data recording were conducted during the nocturnal phase of this cycle (Euston & McNaughton, 2006). The individual rats were labelled as 7165, 8202, 8068, and 8482.

2.2 Surgery

Anesthetized rats underwent stereotactic implantation of tetrode micro-drive under sterile conditions. Rats were deeply anesthetized using isoflurane (1 % to 1.5 % by volume in oxygen at a flow rate of 1.5 L/ minute) as a general anesthetic, placed in a stereotaxic

holder, and injected with penicillin G (30,000 U per hindleg, i.m.). The use of a stereotaxic frame facilitated the accurate positioning of the recording devices according to stereotaxic coordinates (Euston & McNaughton, 2006).

2.2.1 Stimulation Electrode Implantation

The skull was cleared of skin and fascia, and craniotomies were conducted to place two stimulating electrodes targeting the Medial Forebrain Bundle (MFB) and a hyperdrive. Each stimulating electrode comprised two Teflon-coated stainless steel wires (coated diameter of 0.0045 inch) twisted together, with $\sim$0.5 mm of insulation removed from one tip. The stereotaxic coordinates for the electrode placements are detailed in Table 2.1. Although only one stimulating electrode was necessary, two were implanted to ensure that at least one would be optimally placed in order to produce adequate reinforcement (Euston & McNaughton, 2006).

Table 2.1: Stimulation Electrode Coordinates

Rat ID	Anteroposterior (AP)	Mediolateral (ML) bilaterally	Ventral from Dura
7165	-3.25 mm	1.65 mm	8.5 mm
8068	-3.8 mm	1.65 mm	8.4 mm
8202	-2.5 mm/-4.0 mm	1.8 mm/1.5 mm	8.5 mm/8.2 mm
8482	-4.0 mm/-2.5 mm	1.5 mm/1.8 mm	8.2 mm/8.5 mm

2.2.2 Hyperdrive implantation

Holes were drilled into the animal's skull in order to attach the implant anchoring screws, which serve to secure the implant firmly to the rat's head. A 2 mm circular craniotomy was performed over the left hemisphere of the Prefrontal Cortex (PFC), and the underlying dura was excised to accommodate the implant. The rats underwent chronic implantation of micro-drives containing 14 independently manipulable four conductor electrodes, known as tetrodes. These tetrodes consisted of four, polyimide-coated, fine nichrome wires (each measuring 14 µm in diameter) twisted together and fused by partial melting of in-

sulation. Each tetrode was encased in silica tubing (65 µm inner diameter, 125 µm outer diameter) which is further encased in a cannula extending to the end of the drive, making contact with the brain surface. The silica tubing was securely fastened to a drive cannula, connected to a precision screw that allowed the experimenter to adjust the depth of the tetrodes. Rotating the screw head by one full turn advanced the tetrode by 320 µm. Each wire from the tetrode was connected to a connector board positioned atop the hyperdrive. The hyperdrive was positioned over the left Medial Prefrontal Cortex (mPFC) and angled at towards the midline. The stereotaxic coordinates for the hyperdrive placements are detailed in Table 2.2. The bundle of cannulae holding the electrodes as they entered the brain, and consequently the electrodes themselves, extended approximately 1.4 mm to 2 mm in the anteroposterior dimension, sampling cells from a relatively wide region of the mPFC. The mediolateral extent varied between 0.5 mm to 1.4 mm, depending on the geometry of the cannula bundle (Euston & McNaughton, 2006).

Table 2.2: Hyperdrive Implant Coordinates

Rat ID	**Anteroposterior (AP)**	**Mediolateral (ML) bilaterally**	**Angle toward midline**
7165	3.0 mm	1.3 mm	9.5°
8068	3.2 mm	1.3 mm	9.5°
8202	2.9 mm	1.3 mm	9.5°
8482	3.0 mm	1.3 mm	9.5°

After surgery, twelve of the fourteen tetrodes were lowered into the mPFC. One of the tetrodes was lowered to an electrically silent area in the Infralimbic Cortex (IL) region in order to serve as a reference, while the other potential reference remained in the Secondary Motor Cortex (M2). Electrodes were advanced periodically during the experiment to ensure sampling from a large group of mPFC cells. mPFC cells were recorded from the Cingulate Cortex Area 1 (Cg1) and the Prelimbic Cortex (PL) subregions (Paxinos & Watson, 2007). This setup facilitated the concurrent recording of extracellular single unit activity in the mPFC unilaterally from the left hemisphere. Both Magnetic Resonance Imaging (MRI)

and histological imaging were used to validate electrode locations. Twisted-pair local field potential recording electrodes were also surgically implanted in the Hippocampus (HPC) (Euston & McNaughton, 2006).

Subsequently, the craniotomy site was filled with silastic, a protective polymer compound, and the hyperdrive was secured to support screws using dental acrylic. Following surgery, the rat underwent a postoperative recovery period lasting 3 to 4 days in the animal colony room with *ad libitum* feeding (Euston & McNaughton, 2006).

2.2.3 Spatial Sequence Behavioral Environment

Behavior took place in a circular arena (1.3 m in diameter) with computer-controlled Light Emitting Diodes (LED) edge-mounted at eight equidistant locations around the circumference to guide the animals. The LEDs were positioned 2 cm above the table surface and emitted flashes at a frequency of 2 Hz when activated. The onset of the light was concurrently cued by a 3 kHz, 200 ms tone (Euston & McNaughton, 2006).

Liquid feeders operated by gravity, comprising a 60 mL syringe connected through Tygon tubing (Formula B-44-4X; Saint-Gobain Performance Plastics, Akron, OH) to a solenoid valve (Series 3, two-way solenoid valve; Parker Hannifin Corporation, Pneutronics Division, Hollis, NH), and ultimately to a feeding port (⅛ inch inner diameter polypropylene male luer tubing connector; Small Parts, Miami Lakes, FL) situated within a small plastic weigh boat affixed to a 5 cm steel washer on the tabletop. The weigh boats were excluded for rats trained using electrical brain stimulation (Euston & McNaughton, 2006).

2.3 Training

For pre-training purposes, all rats were food deprived to 85 % of their *ad libitum* weight. The rats did not undergo a formal pre-surgical training regimen. Instead, they were acclimated to the maze over a three-week period by being placed on the table for 20 to 50 minutes, three times a week. Following the surgical procedure, the rats were trained over

several days to run towards flashing LEDs in a random order, with MFB stimulation used as reinforcement. The onset of the light was concurrently cued by a 4 kHz, 200 ms tone (Euston & McNaughton, 2006).

2.4 Spatial Sequence Behavioral Paradigm

All sessions consisted of two 50 minute behavior (maze) epochs, surrounded by three, 30 to 60 minute sleep epochs. During sleep epochs, the animals were placed in an elevated ceramic pot lined with towels and left undisturbed, allowing for the collection of baseline data on quiet wakefulness and sleep.

During maze epochs, the animals were trained to run towards a series of target locations around the circular arena, positively reinforced by MFB stimulation rewards awarded at each correct zone (Section 2.8.1). Sequences, consisting of either six or eight locations, were repeated throughout the course of the running maze session, alternating in blocks of three *cued* (light-guided) and three *non-cued* (memory-guided) sequences throughout the session. In non-cued condition the lights were invisible until 5 s timeout. This allowed an assessment of the degree to which rats had memorized the sequence.

Sequences were either 6-step circular sequences (simple), or complex 8-step circular sequences with overlapping segments requiring hierarchical coordinate maintenance to perform contextual disambiguation. Animals were trained on each sequence for three to five days. Each rat would be first trained to learn one sequence for a period of 1 or 2 recording sessions, with one session allocated for the simple sequence and two for the complex sequence. Subsequently, on the second or third session, the rat transitioned to a new sequence, typically around one fourth of the way through maze 1.

Afterwards the sequence was flipped and rotated to create a novel sequence. Rats typically ran the sequences with a high degree of accuracy, following highly stereotyped paths. Even without cue lights, rats behavior was highly accurate. Rats would typically attain between 800 and 1100 rewards during any given maze session (Euston & McNaughton,

2006).

2.5 Data Acquisition

Single units were recorded with respect to a reference electrode positioned deep in the mPFC (5000 µm from the brain surface). The tetrodes were moved gradually from 2000 µm to 4000 µm from the brain surface, keeping all electrodes at the same depth as nearly as possible. Typically, a full set of behavior (see below) was obtained at a given depth and then the electrodes were moved down 80 µm to 120 µm to ensure a fresh group of cells. After that, tetrodes were moved only as necessary to obtain good recordings until the next ensemble turning event. This procedure was repeated until electrodes reached a depth of 4000 µm (rat 7165) or 3200 µm (rat 8202). All tetrode positioning was done after a given recording session, to allow the tetrodes at least 18 h to stabilize. After all recordings were complete, the tips of the recording electrodes were marked by electrolytic lesions (5 A for 10 s, positive to electrode, negative to head-stage ground). Histological sections confirmed that the electrodes in rats 1 and 2 were in the Cg1 and PL regions, with most electrodes in the deep layers.

During the recording sessions, the hyperdrive was connected to a unity gain head stage (Neuralynx, Tucson, AZ) using connector pins, facilitating the transmission of neural data to the recording system with minimal noise interference. This head stage also incorporated a set of LEDs visible to an overhead camera, enabling real-time tracking of the rat's position on the maze at a frame rate of 60 frames/s. All data were acquired utilizing a Cheetah recording system (Neuralynx) in conjunction with a computer. Prior to digitization at 32 kHz, single-unit data from each tetrode underwent amplification (anywhere from 3000 to 9000 times) and bandpass filtering within the 0.6 kHz to 6 kHz range (Assembly Hunter amplifiers; Neuralynx). The video maintained a spatial resolution of approximately 3 pixels/cm (Euston & McNaughton, 2006).

Control of the experiment was executed by a BX-35 microcontroller card (Netmedia,

Tucson, AZ) along with a standard personal computer, which also handled data acquisition. A bespoke software monitored the rat's position and activated lights, tones, feeders, and electrical brain stimulation as needed. Solenoids were regulated through a bank of optical isolators (I/O module 70-ODC5; Grayhill, LaGrange, IL), while LEDs were powered directly from the microcontroller's Transistor-Transistor Logic (TTL) output. Electrical brain stimulation, administered via a stimulus isolator (model A365; World Precision Instruments, Sarasota, FL), was also controlled by the microcontroller's TTL output. Tones were emitted using the built-in audio controller of the personal computer (Euston & McNaughton, 2006).

For every control event initiated by the BX-35, a concurrent signal was dispatched to the TTL input port on the recording apparatus for time stamping, employing the same clock utilized for neural and video data (Euston & McNaughton, 2006).

2.6 Stimulation Electrode Calibration

Bipolar current pulses were applied to the MFB stimulation electrodes with varying durations, currents, and polarities to identify the optimal parameters. The effectiveness of the stimulation parameters was assessed by measuring the rate at which the rat self-administered MFB stimulation in an operant conditioning box via nose-pokes into a box equipped with an infrared trigger that activated the electrical stimulation. The optimal setting was determined to be a current of 70 to 100 µA at 150 Hz, with a stimulation train duration of 300 to 350 ms. Each biphasic current pulse was a bipolar square-wave with a width of 400 µs. In this setup, the rat received one train of stimulation for each correct goal attained (Euston & McNaughton, 2006).

2.7 Data analysis

All data analyses were performed using Python (Python Software Foundation. Python Language Reference, version 3.8). All code and the IPython Notebook containing the anal-

yses are available at https://github.com/kushaangupta/pfc-planning. Analyses specific to each investigation are mentioned in subsequent chapters.

2.8 Terminology

2.8.1 Reward zone

A *reward zone* is defined as a circular region, with its center located at a specific position on the circumference of the arena. The radius of each zone is 10 cm. The animal must enter this zone to receive a reward (Euston & McNaughton, 2006).

2.8.2 Sequence segment

Within a spatial sequence task, a *sequence segment* is defined as the consecutive pair of goal locations (source zone and goal zone) in the task sequence that the animal must navigate between to obtain a reward. A *sequence* is a predefined, ordered set of sequence segments within the task environment. Each maze session consists of an ordered sequence of n (where n = 6 or 8) such segments. These segments were arranged such that the goal zone of the final segment coincides with the source zone of the initial segment, thereby forming a closed loop. To achieve continuous reward, the animal must traverse these segments in the prescribed order. Notably, complex sequences with eight locations will contain two overlapping segments, where certain portions of the pathway were shared between the first and second half of the sequence.

2.8.3 Trial

We define a *trial* as the path traversal time period that begins upon the animal's departure from a goal zone and ends upon their arrival at the next correct goal zone in the sequence segment (Euston & McNaughton, 2006).

Chapter 3

Search for Position-Correlated Representations in the mPFC

3.1 Introduction

Drawing on previous studies that have demonstrated position-correlated neural representations in rodent Medial Prefrontal Cortex (mPFC) and their utility in elucidating the cognitive processes underlying task-solving strategies of the animals (Mashhoori et al., 2018), we sought to decode this spatial information in order to understand the cognitive mechanisms underlying the animal's performance in the spatial sequence task. Using a methodology akin to Mashhoori et al. (2018), we extended their approach to a more complex navigation paradigm. Given the significant working memory demands of the spatial sequence task, we hypothesized that neural representations in the rat mPFC would yield results tantamount to those reported by Mashhoori et al. (2018).

3.2 Analytical Methods

3.2.1 Preprocessing

The preprocessing technique employed mirrors that performed by Euston and McNaughton (2006) and Mashhoori et al. (2018), but it is elaborated upon here in greater detail. For stimulation rewarded rats (Rat IDs: 7165, 8202, 8068, 8482) in spatial sequence task data, artifacts resulting from Medial Forebrain Bundle (MFB) stimulation were eliminated by deleting all spikes recorded during the stimulation period. Putative single neurons, called single units, were isolated using MClust (Redish et al., https://redishlab.umn.edu/mclust),

which facilitates sorting based on the relative amplitudes of action potentials across different tetrode channels and other waveform parameters (McNaughton et al., 1983; Wilson & McNaughton, 1994). The final output is a collection of timestamps corresponding to each action potential from a specific unit (Euston & McNaughton, 2006).

The decoding analysis was conducted on a session-specific basis utilizing data from individual maze sessions. To minimize the influence of non-task behaviors on our neural data analysis, we employed a two-step filtering process. First, individual trials exceeding 1.5 times the duration of the median trial within each session were excluded. This approach aimed to remove neural activity potentially associated with behaviors unrelated to the task, such as grooming or rearing. Subsequently, neurons with an average firing rate below 0.5 Hz were eliminated. Initial assessments by Mashhoori et al. (2018) revealed that these low-firing neurons did not contribute meaningfully to decoding accuracy. Spike data were segregated into trials based on the timestamps marking the delivery of stimulation rewards as the end of each trial, and end of stimulation delivery indicating the start of the next trial. The rationale behind such division into trials beforehand ensures no data leakage occurs across trials due to further preprocessing steps. Within each isolated trial, the spike data was further segmented into non-overlapping bins of 1 ms using a box-car binning approach. A Gaussian kernel with a standard deviation of 150 ms was used for smoothing. The kernel size extended symmetrically up to three standard deviations. Finally, the area under the Gaussian kernel curve was normalized to integrate to 1. This normalization step guarantees that the smoothed data retains the same total firing rate as the raw data, preserving the overall activity level within each time bin. The resulting firing rates were then binned by summing them into non-overlapping and larger time windows of 50 ms each forming a Population Vector (PV) for each time bin. Subsequently, to mitigate the instability of spiking noise variance, the PVs were applied with an element-wise square-root transform, following previously proposed methods (B. M. Yu et al., 2009; Kihlberg et al., 1972).

Positional data of the behavioral experiment was acquired with a spatial resolution cor-

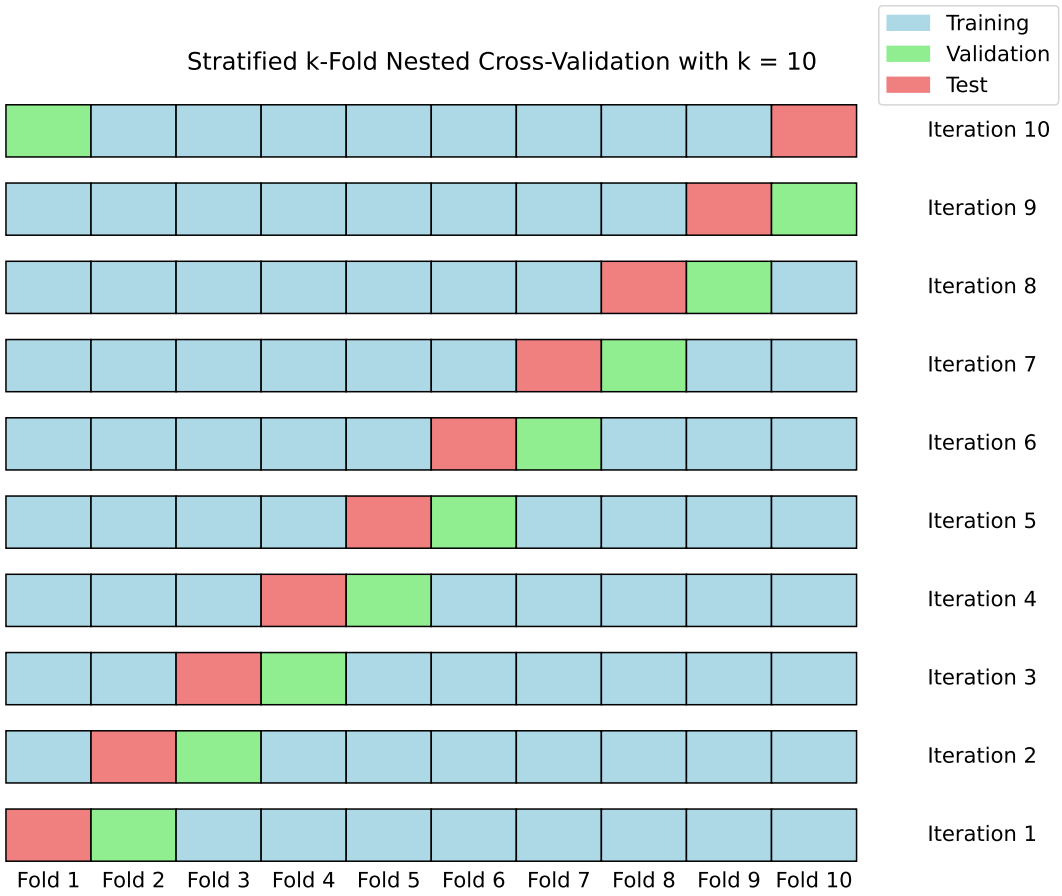

Figure 3.1: Schematic of stratified k-fold nested cross-validation process with 10 folds. Each row represents an iteration of the outer loop of cross-validation, where one fold is designated as the test set (red), another as the validation set (green), and the remaining folds as the training set (blue).

responding to the video tracker's pixel density (∼3 pixels/cm). A circle-fitting algorithm (Euston & McNaughton, 2006) was employed to determine the subject's position in each video frame captured at a temporal resolution of 15 ms. This data was subsequently linearly interpolated to match the 50 ms bin size used for the preprocessing of neural firing rates.

3.2.2 Cross-validation

To ensure robust validation and prevent overfitting, a stratified k-fold nested cross-validation strategy was employed throughout the training process. This method involves dividing the data into k folds (where $k = 10$) and iteratively using one fold for validation,

one for held-out testing, while training on the remaining folds (Figure 3.1). To mitigate potential confounding effects arising from the temporal proximity of PVs within smoothed time bins, we opted for a fold-selection strategy that did not rely on individual PVs. Instead, folds were constructed by randomly selecting trials corresponding to each sequence segment, employing stratification to ensure an equal number of trials from each sequence segment in each fold. This approach avoids biases introduced by the similarity inherent in temporally adjacent bins due to the preprocessing smoothing step. Instead of dividing the folds based on individual time-binned PVs—which could be confounded by the similarity introduced by temporally-adjacent bin smoothing during preprocessing—each fold consisted of an equal number of randomly selected trials (defined in Section 2.8.3) corresponding to each sequence segment (defined in Section 2.8.2) using stratification. This nested approach provides a more reliable estimate of the model's generalization performance. During each training iteration, the training, validation, and testing sets were independently applied with z-score normalization of the neuronal firing rates to enhance the efficiency of the gradient descent optimization process. This normalization technique stabilizes the error landscape, facilitating faster convergence during optimization (LeCun et al., 1998). Moreover, independent normalization of each dataset guarantees the absence of data leakage (Glaser et al., 2020).

3.2.3 Hyperparameter optimization and Model Architecture

Following the precedent of Mashhoori et al. (2018), the decoding analysis employed a Deep Neural Network (DNN) with a fully-connected Multi-layer Perceptron (MLP) architecture. In contrast to the arbitrary selection done in Mashhoori et al. (2018), the hyperparameters governing the model architecture were selected using a random subset of 4 sessions from the pool of all animals. All training procedures and decisions regarding hyperparameter selection were made using the validation partitions of k-fold nested cross-validation (for more details on the training procedure, see Section 3.2.4).

A range of hyperparameters to optimize over was manually deduced after manual training on a few of the subsampled sessions. The manual search also selected the number of training epochs; the activation function; the gradient clipping value; and initial learning rate, minima, and step size of the dynamic learning rate scheduler. The selected range of hyperparameters tested the number of hidden layers from 3 to 5 and dropout probability from 0 to 0.33. The number of units in each hidden layer could be varied from 64 to 512.

The optimized architectural hyperparameters were as follows: the input layer of the MLP received a N-dimensional vector representing the state of the neural population during a time bin in a session, where N corresponds to the total number of recorded single units. The hidden layer architecture consisted of two fully-connected layers, each containing 512 units. A Continuously Differentiable Exponential Linear Unit (CELU) activation function was applied after each hidden layer to introduce non-linearity. To mitigate overfitting, a dropout layer with a probability of 0.2 was incorporated following the second hidden layer. Finally, the output layer comprised 512 fully-connected linear units.

3.2.4 Model training procedure

The DNN was trained individually for each session to minimize the Mean Squared Error (MSE) between its estimated current location and the animal's actual current location.

$$\text{MSE} = \frac{1}{n}\sum_{i=1}^{n}\left((x_i - \hat{x}_i)^2 + (y_i - \hat{y}_i)^2\right) \tag{3.1}$$

where,

- n represents the total number of samples,
- x_i and y_i represents the actual coordinates of the i-th sample,
- $\hat{x}_i$ and $\hat{y}_i$ represents the predicted coordinates of the i-th sample based on the model.

The training partition of the dataset, obtained through stratified k-fold nested cross-validation, comprised 50 % of the trials in each of the four model training iterations con-

ducted on the different distribution of the session data. The model training procedure employed multifarious techniques in order to minimize inductive biases due in hyperparameter selection, accelerate the training process, and ensure effective generalization. The model was trained using mini-batch gradient descent where mini-batches repeatedly sampled randomly from the training dataset to produce independent and identically distributed subsets.

The AdamW optimizer—which decouples weight decay regularization from the gradient update—was chosen as the optimization algorithm for training the model. This weight decay penalizes large parameter values during training more effectively than L2-norm regularization by accounting for adaptive gradients, encouraging the model to learn generalized distributed representations (Loshchilov & Hutter, 2017). It offers advantages beyond the existing capabilities of Adam optimizer to dynamically adjust the learning rates for individual parameters throughout the training (Kingma & Ba, 2014).

A cyclical learning rate scheduler was implemented to dynamically adjust the learning rate during training. This approach has been shown to improve convergence and generalization compared to static learning rates (Smith, 2017).

Moreover, gradient clipping is a technique used to address issues of poor conditioning in the training of deep neural networks. Poor conditioning typically refers to situations where the gradient values vary significantly in magnitude, potentially causing numerical instability and making the optimization process difficult. Gradient clipping mitigates this by setting a threshold value for gradients, rescaling those that exceed this limit, thereby preventing excessively large updates to model parameters and avoiding the destabilization of the learning process (Pascanu et al., 2012). Additionally, it supports the effectiveness of adaptive optimizers by ensuring that gradient magnitudes remain controlled, allowing these optimizers to make consistent and stable updates (Kingma & Ba, 2014).

For the training loop, we employed mini-batch gradient descent which involves splitting the training dataset into smaller, randomly sampled batches (batch size = 4096 PVs). The order of the training data was shuffled at every epoch to introduce randomness and pre-

vent overfitting. Mini-batch gradient descent calculates the model error and updates model coefficients based on these batches. We specifically opted to utilize the sum of gradients across the mini-batch, which has the benefit of further reducing the variance of the gradient estimate.

The training loop ran for 10 epochs. This epoch selection was based on observations of asymptotic performance beyond this point, as determined through manual runs across various session datasets. To prevent overfitting and monitor model performance, validation loss was monitored after each epoch. The model parameters were checkpointed at each epoch. If the optimization process using the trained model parameters yielded a lower validation loss, the checkpoint was updated with these improved parameters.

3.2.5 Evaluation of decoder performance

Following the completion of training, the decoding performance of the decoders per session was assessed on their respective test sets utilizing the coefficient of determination and Root Mean Squared Error (RMSE) metrics. Coefficient of determination is a unit-agnostic metric of decoding performance that measures the fraction of variance explained by the decoder of the actual predicted variable. This value is only equivalent to the square of Pearson's correlation coefficient in the case of linear regression models. Coefficient of determination has a maximum value of 1, and value of 0 suggests that the given model has performance equivalent to a naive model predicting the mean of test data. Coefficient of determination does not have a lower limit, but it can be negative on the held-out test set due to the decoder performing worse than just predicting the mean of the predicted variable on the test set as a consequence of overfitting on the training set (Glaser et al., 2020). Figure 3.2 represents an illustrative example of a case where the coefficient of determination R^2 can be negative. RMSE error offers a unit-specific metric of decoding performance which can interpretably quantify the precision of the decoder that is comparable across datasets.

For the evaluation, the predictions on the test data for a session we concatenated across

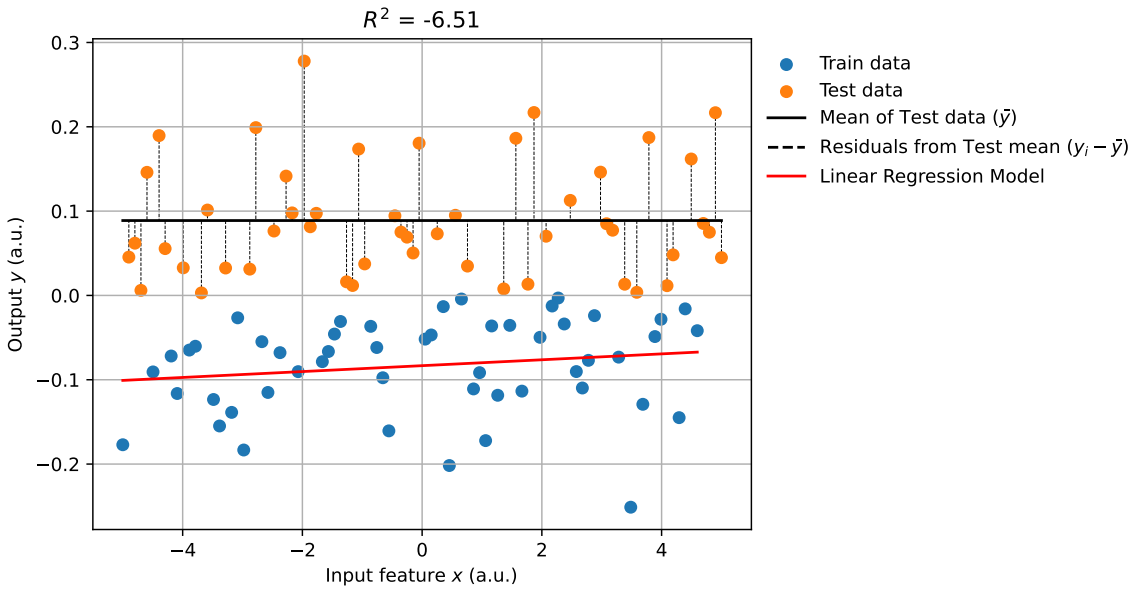

Figure 3.2: An illustrative example of a case where the coefficient of determination R^2 can be negative. The red line represents the actual data, while the blue line represents the predicted data. The dashed green line represents the mean of the actual data. In this case, the model performs worse than predicting the mean of the actual data, resulting in a negative coefficient of determination.

the individual training iterations, thus uniformly encompassing the whole session data. Subsequently, the mean coefficient of determination score for a session was calculated by taking a variance-weighted mean of the coefficient of determination scores for the predicted vs actual x and y coordinates separately. Similarly, the RMSE for a session was calculated by taking a uniformly-weighted mean of the RMSE errors for the predicted vs actual x and y coordinates separately. These aggregated metrics served as the final performance metric for each session. A higher value of coefficient of determination indicates better decoding performance, whereas a lower value of RMSE indicates better decoding performance. The stratified k-fold nested cross-validation technique ensured that testing could be done over the whole dataset with each testing set having a separate held-out training set.

We define the RMSE as:

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}((x_i - \hat{x}_i)^2 + (y_i - \hat{y}_i)^2)} \tag{3.2}$$

where,

- n represents the total number of samples,
- x_i and y_i represents the actual coordinates of the i-th sample,
- $\hat{x}_i$ and $\hat{y}_i$ represents the predicted coordinates of the i-th sample based on the model.

We define the coefficient of determination, denoted by R^2, as:

$$R^2 = 1 - \frac{\sum_{i=1}^{n} (y_i - \hat{y}_i)^2}{\sum_{i=1}^{n} (y_i - \bar{y})^2} \tag{3.3}$$

where,

- y_i represents the actual value of the i-th data point for the dependent variable,
- $\hat{y}_i$ represents the predicted value of the i-th data point for the dependent variable based on the model,
- $\bar{y}$ represents the mean value of the dependent variable across all test data points.

3.2.6 Shuffled null distribution

To establish null baseline distribution for decoding performance, we employed the same preprocessing steps detailed previously (see Section 3.2.1). However, after binning spikes into 1-millisecond windows, individual trial spikes were shuffled for each neuron. The preprocessing pipeline then continued as previously described. Subsequently, the shuffled data underwent the identical stratified k-fold nested cross-validation scheme outlined previously (see Section 3.2.2). This ensured that each model iteration was evaluated on a shuffled version of the held-out test set, allowing for fair performance comparisons between models and shuffled data. The resulting test set predictions were then concatenated and assessed using the evaluation metrics outlined previously (see Section 3.2.5). These metrics then served to

establish the session-specific null distribution of the decoder's performance. We repeated the procedure $n_{null} = 10$ times to obtain the null distribution for decoding performance.

3.3 Results

3.3.1 Spatio-temporal encoding of position by mPFC

Our initial investigation aimed to determine if the population activity within the mPFC accurately encoded the position of animals navigating a spatial sequence task. In this task, rats received MFB stimulation as a reward for traversing between designated goal locations in a predetermined sequence throughout a maze session (Figure 3.3C). The sequence alternated between cued and non-cued conditions at periodic intervals, with each cycle lasting three sequence traversals. Notably, the task environment remained constant, while the specific set of goal locations were periodically changed across sessions. During each behavioral session, rats completed between 500 to 1100 trials. Simultaneous spatial position of the animal and single-unit neural activity was recorded in the mPFC capturing the activity of 31 to 144 individual units per session (median = 75 units).

PVs were used as input to a DNN decoder, tasked with predicting the animals' current location every 50 ms throughout each session. Analysis of decoding performance across 120 sessions from four animals revealed substantial variability (Figure 3.4A–D). Analysis of the coefficient of determination[1] scores revealed limitations in decoding accuracy across the majority of sessions. Only $\sim$82 % (98 out of 120) of the sessions exhibited non-negative coefficient of determination scores, indicating that the decoding method failed to explain any variance in spatial location for the remaining $\sim$18 % (22 sessions) based on neural population activity. Among the sessions with non-negative coefficient of determi-

[1] The coefficient of determination reflects the proportion of variance in the test data explained by the model relative to a simple prediction based on the mean of the test data. For non-linear models (like artificial deep neural networks), this is not equivalent to the square of Pearson correlation coefficient. When evaluating a model on out-of-distribution data that the model was not designed to handle, the coefficient of determination can become negative. This outcome signifies that the model has overfit to the training data and performs worse than a naive model that simply predicts the average value of the target variable in the test set. See explanation in Section 3.2.5

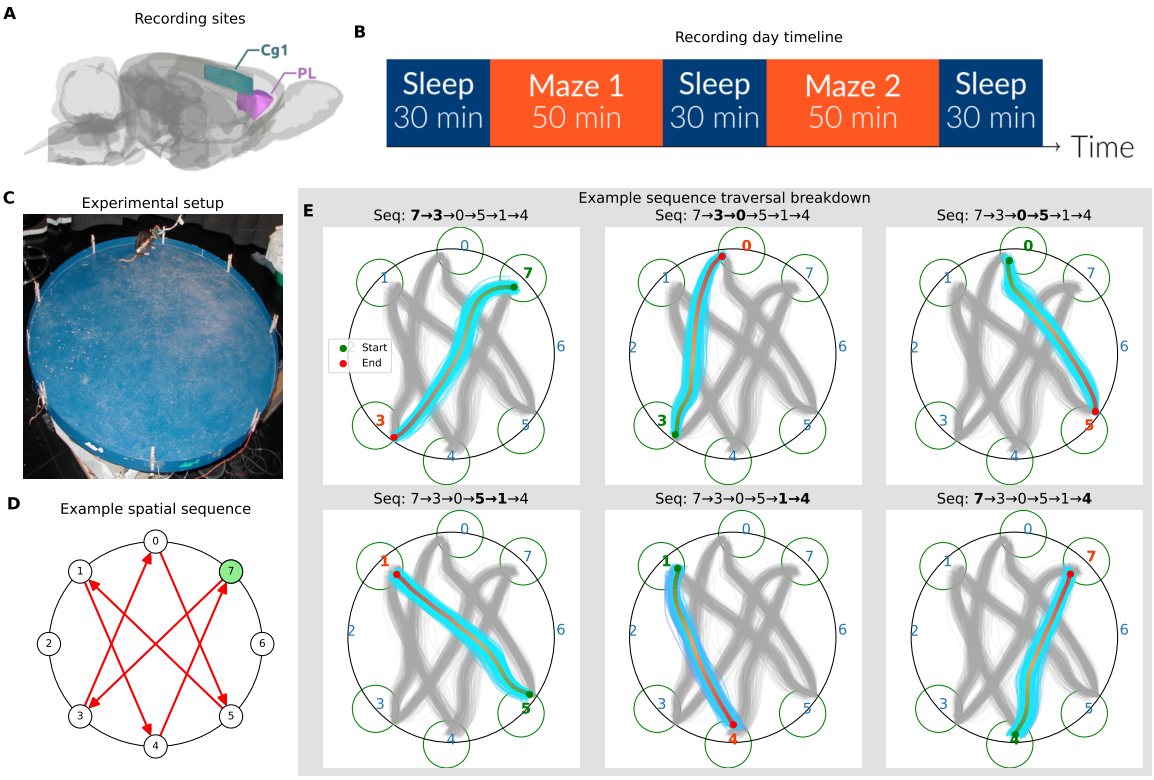

Figure 3.3: Experimental setup and task design. (**A**) Multi-unit recording sites (Cingulate Cortex Area 1 (Cg1) and Prelimbic Cortex (PL)) illustrated on the lateral view of a schematic of rat brain. (**B**) Timeline of events within a recording day. (**C**) Image of the experimental setup showing the implanted animal, maze and the location of the reward ports. (**D**) Illustration of an example spatial sequence across the maze. Green node represents the start location. (**E**) Breakdown of spatial sequence behavioral paradigm by sequence segments from a representative session. Each segment of the 8-element sequence is displayed separately, though the rat runs continuously through all. Light blue color highlights all preprocessed trials corresponding to the current sequence segment, while gray shows the trials for the entire sequence. Progress along a sequence segment is indicated by color gradient from green to red.

nation scores (n = 98), the mean score was 0.47 with a standard deviation of 0.22. This suggests moderate decoding performance on average, but with substantial variability. Furthermore, the average RMSE across all animals for these sessions was 24.33 cm, with a standard deviation of 5.26 cm. Notably, a subset of 38 sessions achieved high coefficient of determination scores (>0.6) coupled with low RMSE (18.85 cm; std = 1.43 cm), indicating successful decoding of spatial location in these instances. Conversely, the majority (52) of sessions exhibited low coefficient of determination scores (<0.4) alongside high RMSE (29.48 cm; std = 3.26 cm), suggesting limited explanatory power of the decoding method in these cases.

Although the obtained average RMSE appears to be approximately 20 % of the diameter of the tabletop, it is to be noted that the inter-reward distance that the animal traversed between two rewards was mostly smaller than the diameter of the tabletop. Moreover, the average RMSE value of shuffled data was found to be 45.68 cm with a standard deviation of 2.8 cm. All of the sessions exhibited RMSE below the minimum of a shuffled null distribution of RMSEs generated for each session (description in Section 3.2.6). However, the average RMSE of shuffle distribution is about 35 % of the diameter of the tabletop. This suggests that the decoding performance of the model is significantly better than chance, but the variability in the performance across sessions is substantial. Furthermore, the predominance of negative and weak coefficient of determination scores (< 0.4) underscores the overall limited decoding accuracy and its potential insufficiency in revealing the cognitive processes underlying spatial sequence task performance.

The inconsistent decoding performance obtained from our analysis was unexpected, particularly given the previous success of similar rat mPFC neural decoding methods in other spatial navigation tasks (Mashhoori et al., 2018). In particular, Mashhoori et al. (2018) reported decoding RMSEs of $\sim$11 cm in animals navigating a figure-eight maze engaged in an effort-based decision-making task. Using the same neural decoding methodology, Mashhoori et al. (2018) were able to reveal fictive trajectories toward alternative reward

sites in the figure-eight maze.

Our inability to achieve similar performance is noteworthy, considering that both our neural data and that of Mashhoori et al. (2018) underwent identical preprocessing and originated from functionally similar brain regions: Cg1 in Mashhoori et al. (2018) and Cg1 and PL in our study (Euston & McNaughton, 2006). Previous analyses on our dataset, as reported by Euston and McNaughton (2006), revealed highly similar response properties of the cells within these regions (Cg1 and PL). Given the precedent set by Mashhoori et al. (2018), we sought to compare our results with theirs by explicitly comparing the decoders and datasets used in both studies. Specifically, our analysis aimed to address two critical questions: (1) How does the performance of our decoder compare to that of Mashhoori et al. (2018)? (2) Are there any fundamental differences in the neural data that could explain the discrepancy in decoding performance between the two studies?

To test whether the low decoding performance in spatial sequence data was due to differences in DNN architecture or training procedure, we employed the identical DNN architecture and training loop reported by Mashhoori et al. (2018), leveraging their publicly available code. However, this replication did not yield any improvement in decoding errors across a total of 120 sessions (n = 4 animals). In fact, a one-tailed Wilcoxon signed-rank test with continuity correction and exact value reporting revealed a significant decrease in RMSE (hence, improvement in performance) using our decoder compared to the decoder from Mashhoori et al. (2018) ($W = 4065$, $p < 0.001$, alternative hypothesis: greater). The effect size, calculated using matched pairs rank-biserial correlation, was large ($r = 0.985$). This result suggests that the discrepancy in decoding performance between the two studies is not due to differences in the DNN architecture or training procedure (Figure 3.4E–F).

As a further confirmation, we evaluated our DNN decoder on the figure-eight maze task data from Mashhoori et al. (2018) and compared its performance to their decoder. Once again, our decoder outperformed the decoder reported in Mashhoori et al. (2018). A one-tailed Wilcoxon signed-rank test with continuity correction and exact value reporting

revealed a significant decrease in RMSE using our decoder compared to the decoder from Mashhoori et al. (2018) ($W = 21$, $p < 0.001$, alternative hypothesis: greater, $n = 7$ sessions). The effect size, calculated using matched pairs rank-biserial correlation, was large ($r = 1$). This result bolsters the conclusion that the discrepancy in decoding performance between the two studies is not due to differences in the DNN architecture or training procedure (Figure 3.4G–H).

Having established that the performance of our decoding method surpasses that of the decoder in Mashhoori et al. (2018), we conducted a Mann-Whitney U test to compare the decoding errors on both datasets using our decoding model. The results revealed that session RMSE in figure-eight maze were significantly lower than sequence tasks ($U = 540.0$, $p < 0.001$). The effect size, calculated using rank-biserial correlation, was large ($r = -1$). We observed significantly lower decoding errors in the Mashhoori et al. (2018) dataset compared to ours. This discrepancy in decoding performance between the two datasets suggests that the neural data from the spatial sequence task may not be as amenable to decoding as the figure-eight maze task data (Figure 3.4I–J). Of note, the architecture and training procedure of the decoder were identical in decoding both datasets.

Finally, to test if the discrepancy in decoding performance between the two studies is not specific to our decoder, we compared the performance of the decoder from Mashhoori et al. (2018) on both datasets. A Mann-Whitney U test revealed that the session RMSE in the figure-eight maze was significantly lower than in the sequence tasks ($U = 0.0$, $p < 0.001$). The effect size, calculated using rank-biserial correlation, was large ($r = -1$). This result suggests that the neural data from the spatial sequence task may not be as amenable to decoding as the figure-eight maze task data (Figure 3.4K–L).

3.3.2 Could the differences in number of cells or experience explain variations in decoding accuracy across sessions?

Having established that there is something fundamentally different about the data from the spatial sequence task that makes it less amenable to positional decoding, we sought to investigate potential explanations for this discrepancy. In order to statistically evaluate the potential factors contributing to the variability in decoding performance across sessions, we examined the relationship between the number of recorded neurons and decoding performance using linear regression analysis. One possible explanation is that the number of cells recorded in each session may influence decoding accuracy. Mashhoori et al. (2018) had reported that a minimum of 30 mPFC single-units were necessary for accurate reconstruction in their dataset. All sessions in our study included 31 or more single units, exceeding the threshold identified by Mashhoori et al. (2018). Moreover, the sessions from the spatial sequence task reported as having decoding performance worse than a mean predictor (coefficient of determination < 0) had sessions ranging from 31 to 120 single-units. This was one indication that the number of cells recorded in each session was not a limiting factor in decoding performance. Nevertheless, we extended our investigation into this factor by examining the relationship between the number of recorded neurons and decoding performance in our experiment by sorting the sessions for each rat on the basis of their neuron counts and visualizing decoder RMSE as a function of it. We observed ambiguous trends in the relationship between the number of recorded neurons and decoding performance across animals. Rats 7165 and 8202 exhibited a negative correlation between the number of recorded neurons and decoder RMSE ($F(1, 32) = 16.75$, $p < 0.001$, $R^2_{adj} = 0.34$; $F(1, 17) = 25.56$, $p < 0.001$, $R^2_{adj} = 0.61$, respectively), suggesting that the decoder performance improved with an increase in the number of recorded neurons. Conversely, rat 8482 exhibited an insignificant positive correlation between the number of recorded neurons and decoder RMSE ($F(1, 13) = 0.63$, $p = 0.45$, $R^2_{adj} = -0.03$). Rat 8068 exhibited an insignificant negative correlation between the number of recorded neurons and decoder RMSE ($F(1, 6) = 4.22$, p

= 0.11, R^2_{adj} = 0.39). These results suggest that the number of recorded neurons does not consistently explain the variability in decoding performance across sessions (Figure 3.5).

Another potential factor that could explain the variations in decoding accuracy across sessions is the experience of the animals. To investigate this, we examined the relationship between the number of sessions an animal had completed and the decoding performance. The decoding performance was averaged each day to ensure independence of data points and to account for the variability in the number of sessions completed each day. We again observed ambiguous trends in the relationship between the number of sessions completed and decoding performance across animals. Rat 7165 exhibited a significant positive correlation between the number of sessions completed and decoder RMSE (F(1, 32) = 7.42, p = 0.01, R^2_{adj} = 0.17), suggesting that the animal's performance worsened with experience. On the other hand, rat 8482 exhibited a significant negative correlation between the number of sessions completed and decoder RMSE (F(1, 13) = 7.3, p = 0.02, R^2_{adj} = 0.34), suggesting that the animal's performance improved with experience. Moreover, rats 8202 and 8068 exhibited insignificant correlations between the number of sessions completed and decoder RMSE (F(1, 17) = 1.12, p = 0.31, R^2_{adj} = 0.01; F(1, 6) = 0.24, p = 0.65, R^2_{adj} = -0.18, respectively). These results do not provide conclusive evidence that the experience of the animals explains the variability in decoding performance across sessions as the relationship between the number of sessions completed and decoding performance was completely contradictory across animals (Figure 3.6).

3.4 Conclusion

Despite meticulous efforts to decode the animals' instantaneous position in space using near identical recording sites and decoding methods, there were substantial discrepancies in decoding accuracy between animals traversing a figure-eight maze and those engaged in a spatial sequence task. These performance disparities could not be overcome by the use of advanced training strategies for the DNN decoding model, nor could they be explained

as a factor of limited sampling of neurons or learning in the spatial sequence task data. This suggests that the observed decoding performance differences are not a consequence of limitations in the decoding method, but rather reflect fundamental differences between the behavior in the two tasks and their underlying neural representations within the rat mPFC. Specifically, the neural correlates of position in the Mashhoori et al. (2018) dataset appear to be more distinct or more readily decoded. This discrepancy raises questions about the nature of the spatial representations in the mPFC and the cognitive processes underlying spatial sequence task performance.

It may be argued that subregion-specific differences in the mPFC neural code could account for the observed discrepancies in decoding performance. The recordings in Mashhoori et al. (2018) were conducted in the Cg1 subregion of the mPFC, whereas our recordings were conducted in both the Cg1 and PL subregions. The Cg1 and PL subregions have been shown to exhibit similar response properties (Euston & McNaughton, 2006), suggesting that the differences in decoding performance are not likely to be due to subregion-specific differences in the mPFC neural code. Moreover, our recording method only allowed a coarse estimate of the subregion-specific recording sites, making it difficult to ascertain the exact subregion-specific differences in the neural code. Nevertheless, it is possible that the differences in decoding performance between the two studies could be attributed to subregion-specific differences in the mPFC neural code.

Furthermore, given the extensive pre-training period, during which the majority of neural plasticity associated with learning is expected to occur, and the focus on behaviorally stable sessions in our analysis, it is unlikely that the observed variability in decoding performance can be attributed to ongoing learning processes. It is to be noted, however, that the interpretation of the effect of experience on decoding performance is potentially confounded by the fact that the animals were not given the same spatial sequence configuration with different goal locations in each session. This could have introduced variability in the neural data that is not accounted for in our analysis. Moreover, across the days of the ex-

periment, the depth of the electrodes was deliberately advanced to target deeper layers of the mPFC to sample a diverse population of neurons. This could have introduced variability in the neural data that is not accounted for in our analysis. The number of recorded single-units was also affected by the chronic nature of the recordings, which could have led to localized tissue damage and cell loss over time.

The ability to decode fictive trajectories toward alternative reward locations in the figure-eight maze, as reported by Mashhoori et al. (2018), alluded to the presence of an abstract cognitive map-like representation in the mPFC. The authors had argued that such representations could be leveraged to guide behavior toward alternative choices. Although Mashhoori et al. (2018) had inferred their success in precise decoding of spatial position as a potential consequence of the mPFC encoding task-specific abstractions across the spatial environment that gives the semblance of a spatial code, our results highlight the limitations of scaling up such decoding methods to more complex spatial tasks.

Beyond the task differences, the discrepancy in decoding performance between the two studies could also be attributed to the differences in experimenters that was necessary to carry out such a long-term project, specific environmental conditions, or other unforeseen variables. While attributing the spatial coding differences solely to uncontrolled variables might be an oversimplification, considering our previous arguments regarding the mPFC's potential for default spatial coding (Section 1.2.3), several questions remain unanswered. What could be the fundamental differences between the two tasks that lead to such vast differences in decoding accuracies? Could studying the differences between them reveal an insight into some functional underpinning of mPFC neural ensemble code? Thus, it is crucial to elucidate how task structure influences the neural code. This investigation forms the basis of the following chapter.

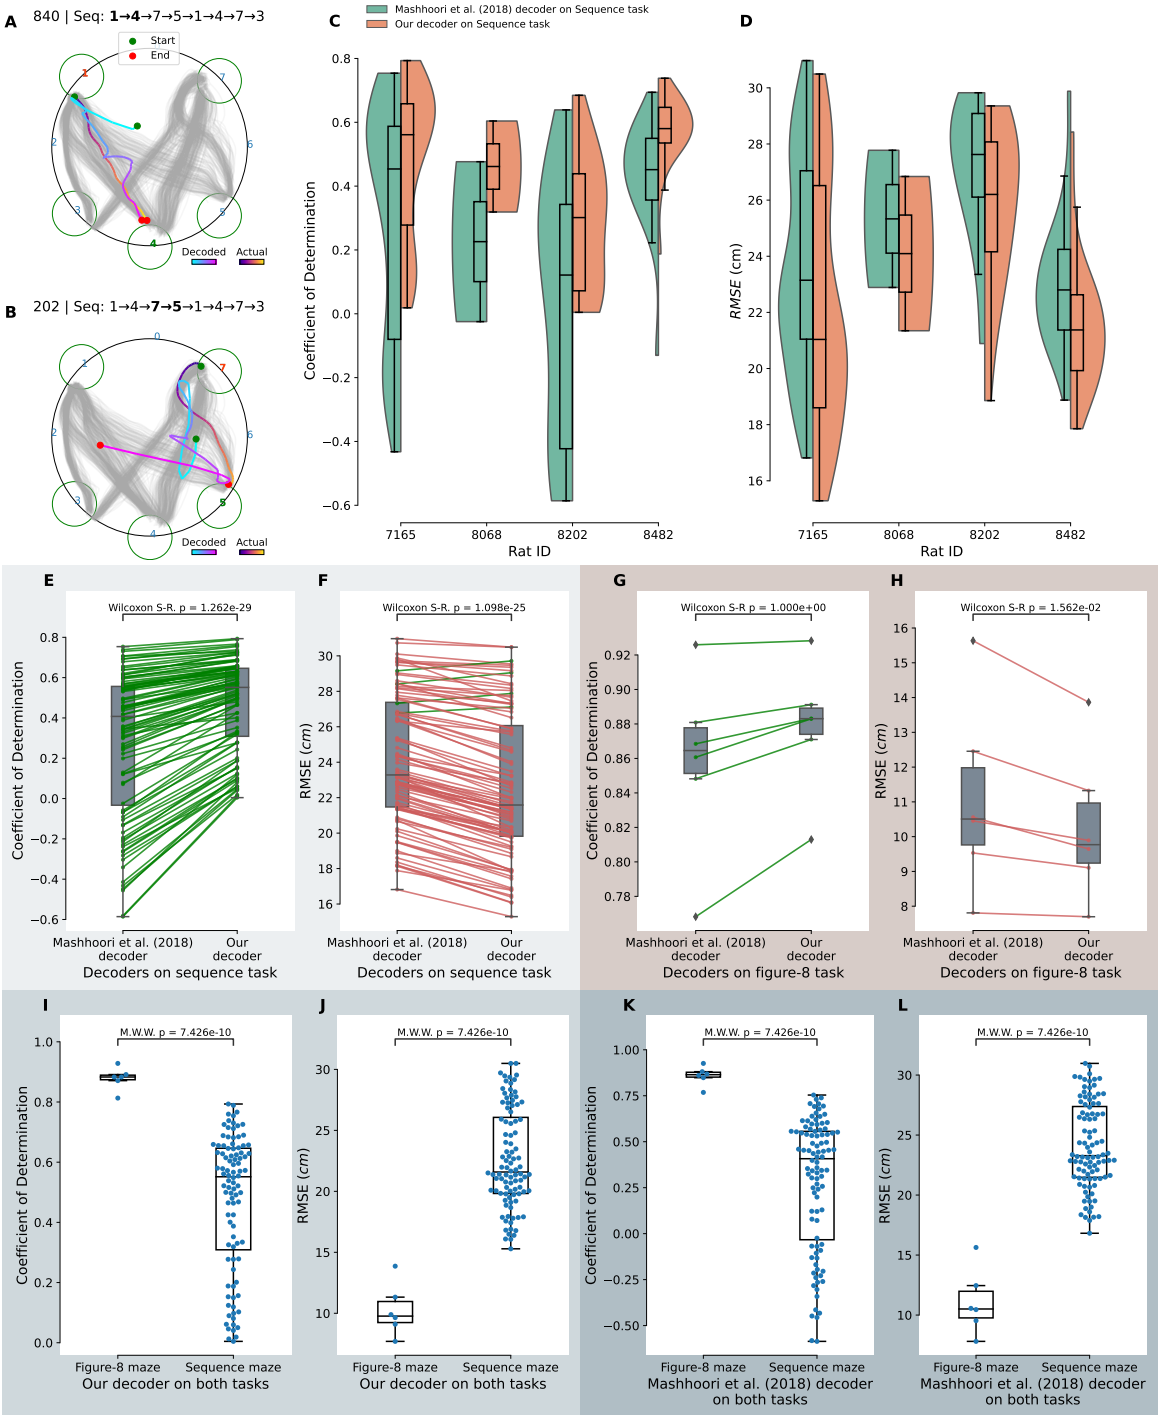

Figure 3.4: Decoding animal position using DNN. (**A**) A representative example with actual of a trial with best decoding accuracy, selected from a session with median overall decoding performance across sessions. (**B**) A representative example with actual and decoded trajectories of a trial with median decoding accuracy, selected from a session with median overall decoding performance across sessions. (**C** and **D**) The accuracy of the trained decoder for the animal's current location across all rats, assessed using held-out data from all maze sessions with a coefficient of determination score above zero. Measured by (C) coefficient of determination score and (D) RMSE error. Each point represents a single session. Whiskers

indicate the maximum or minimum values up to 1.5 times the Interquartile Range (IQR) from the box limits (**E**–**H**) Session-wise comparison of the decoders within each behavioral task. These box-and-whisker plots compare the performance metrics of two decoding models. The box represents the IQR, with the median depicted by the line within the box. The whiskers extend to the maximum and minimum values within 1.5 times the IQR from the edges of the box. Individual data points are represented by circles. Red circles and lines indicate sessions where our decoder produced a lower metric value compared to Mashhoori et al. (2018) decoder (n = 98 sessions). Conversely, green circles and lines indicate sessions where Mashhoori et al. (2018) decoder outperformed our decoder. (E) Comparison of the coefficient of determination scores of the two decoders on the spatial sequence maze task. (F) Comparison of the RMSE scores of the two decoders on the spatial sequence task. (G) Comparison of the coefficient of determination scores of the two decoders on the figure-eight maze task. (H) Comparison of the RMSE scores of the two decoders on the figure-eight maze task. (**I**–**L**) Session-wise comparison of each decoder's performance between both behavioral tasks. In the violin plot, each violin is drawn using a kernel density estimate of the underlying performance metrics' distribution (n = 7 sessions for figure-8 maze, n = 98 sessions for sequence maze). The inner box-and-whisker plot has the box representing the IQR, with the median depicted by the white point within the box. The whiskers indicate the the maximum and minimum values within 1.5 times the IQR from the edges of the box. (I) Comparison of the coefficient of determination scores of our decoder between the spatial sequence and figure-eight maze tasks. (J) Comparison of the RMSE scores of our decoder between the spatial sequence and figure-eight maze tasks. (K) Comparison of the coefficient of determination scores of Mashhoori et al. (2018) decoder between the spatial sequence and figure-eight maze tasks. (L) Comparison of the RMSE scores of Mashhoori et al. (2018) decoder between the spatial sequence and figure-eight maze tasks.

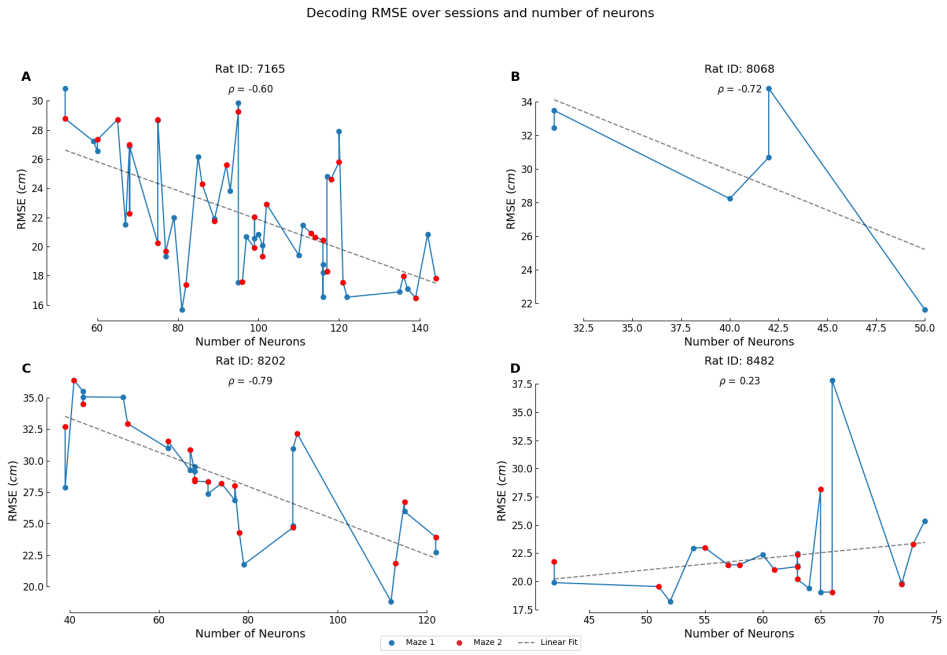

Figure 3.5: Relationship between the number of recorded neurons and decoding performance. (**A**–**D**) Each panel represents data from an individual rat, plotting the RMSE values against the number of neurons. Points depicted in blue represent the initial maze session conducted on the recording day, whereas red points correspond to the subsequent maze session on the same day. Linear fit lines are included to visualize potential trends. The Pearson's correlation coefficient (ρ) values displayed indicate the strength of the linear correlation between the number of neurons and decoding performance for each rat and maze. (A) Rat 7165. (B) Rat 8068. (C) Rat 8202. (D) Rat 8482.

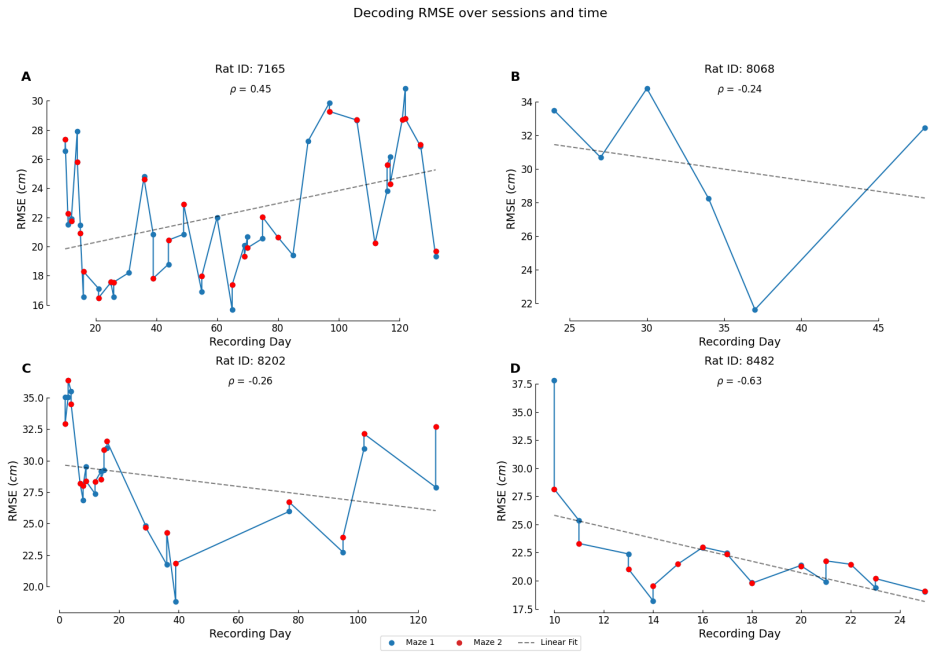

Figure 3.6: Relationship between the number of sessions completed and decoding performance. (**A**–**D**) Each panel illustrates the RMSE values plotted against the recording day for an individual rat. Points depicted in blue represent the initial maze session conducted on the recording day, whereas red points correspond to the subsequent maze session on the same day. The data points show fluctuations in RMSE across the recording days, without exhibiting a clear trend. Linear fit lines are included to visualize potential correlations, and the corresponding Pearson's correlation coefficient (ρ) values indicate the strength of the linear relationship between the RMSE and the recording day for each rat and maze. (A) Rat 7165. (B) Rat 8068. (C) Rat 8202. (D) Rat 8482.

Chapter 4

Behavioral Organization of Medial Prefrontal Computations

4.1 Introduction

The failure of our Deep Neural Network (DNN)-based decoding approach to consistently reveal location representations in the Medial Prefrontal Cortex (mPFC) neural ensemble during the execution of the spatial sequence task motivated us to adopt a data-driven analytical approach to understand the representational currency of mPFC neural ensembles at both the population and single-cell levels. Given the claims that higher cortical areas exhibit multiple selectivity and distributed coding properties (Rigotti et al., 2013), and in light of recent findings on single-cell level specialization of inter-goal progress encoding in the mPFC (Rubin et al., 2019; Kaefer et al., 2020; El-Gaby et al., 2023; Guidera et al., 2024), we began our data-driven analyses with a top-down approach, initially using unsupervised and supervised learning methods to understand neuron representations at the population level and subsequently at the single-cell level.

4.2 Analytical Methods

4.2.1 Preprocessing

Data preprocessing followed the procedures outlined in Section 3.2.1, with the exception of binning window sizes. For the analysis in the Figure 4.1, the Population Vectors (PV) were binned into 100 ms bins with Gaussian kernel smoothing using a kernel with standard deviation of 100 ms. For the remaining analyses in this chapter, we used a smaller

bin size of 50 ms while omitting the Gaussian kernel smoothing step.

4.2.2 Feature extraction

Path progress

To capture the progression of the animal along the spatial sequence, we calculated the path progress of the animal at each time bin. Within each trial, path progression at a specific bin is defined as the proportion of the distance traversed up to that point. Notably, this proportionate distance is not based on an idealized mean or median path, but rather on the total path length the animal covered in that specific trial. Therefore, the denominator is subject to variation due to any deviations from the stereotypical behavior in a sequence segment. We define the path progress, denoted by **P**, as:

$$\mathbf{P}_i = \frac{\sum_{j=1}^{i} d_j}{\sum_{k=1}^{T} d_k} \tag{4.1}$$

where,

– i represents the current time bin,

– d_j represents the distance traversed in between the j-th and $(j+1)$-th time bins,

– T represents the total number of time bins in the trial.

Path progression phases

To capture the progression of the animal along each segment of the spatial sequence, trials were segmented into three distinct phases based on path progress. The first phase corresponds to the *departure path progress phase*, the second phase to the *intermediate path progress phase*, and the third phase to the *approach path progress phase*. By analyzing each phase separately, we can identify potential differences in neural firing rates, population coding patterns, or behavioral strategies employed by the animal during each stage of the

journey. Hereafter, for brevity, we will omit references to path progress when discussing these phases, simply referring to them as departure phase, intermediate phase, and approach phase.

Turn direction

To identify the initial movement direction of the animal upon exiting reward zones and initiating subsequent sequence segments, the turn direction was computed. This direction is defined as the angle between the line fitted to the trajectory at the end of the current trial and the line fitted to the trajectory at the beginning of the next trial. For each trial, the behavioral data corresponding to the preceding and subsequent trials were identified. During the departure phase of the current trial, trajectory data were collected from the point where the animal leaves the source zone until a total distance of 50 pixels ($\sim$16.78 cm) is covered. Similarly, for the preceding trial, the inbound trajectory data covering the final 50 pixels before reaching the goal zone is collected (Figure 4.3B).

Linear regression was then performed individually on these respective trajectories to obtain line segments representing the animal's path. The angle between these two line segments was subsequently calculated to quantify the turn direction (clockwise or counter-clockwise) executed by the rat at the beginning of the current trial, referred to as the *start turn* of the current trial. In other words, the start turn refers to the angle between the line fitted to the trajectory at the end of the preceding trial and the line fitted to the trajectory at the beginning of the current trial (Figure 4.3C). Similarly, to compute the *next turn* of the current trial, the goal approach phase of the current trajectory and the outbound phase of the next trial were used. In other words, the next turn refers to the angle between the line fitted to the trajectory at the end of the current trial and the line fitted to the trajectory at the beginning of the next trial (Figure 4.3D). For simplicity, we refer to clockwise turns as right (R) and counter-clockwise turns as left (L).

4.2.3 Neural population state space

Neural population-level analyses, including Principal Component Analysis (PCA), Support Vector Machine (SVM), and neural population state geometry and dynamics analysis, were performed in a multi-dimensional space defined by simultaneously recorded single-unit time-binned firing counts as each dimension in a session. In this space, the concurrent activity of neurons at each time bin was represented as a point in the high-dimensional neural state space. This point is referred to as a PV (Section 4.2.1). As the activity of the neurons varied over time in a trial, the movement of the point forms a trajectory in the neural state space. This trajectory represents the evolution of the neural population activity over time.

4.2.4 Dimensionality reduction of population neural activity

PCA is a linear dimensionality reduction technique that identifies the dimensions of the data as orthogonal vectors (eigenvectors) that capture the maximum variance in the data. These orthogonal vectors are referred to as principal components. The principal components are ranked in order of the variance they explain, which is quantified by their corresponding eigenvalues. To prepare the neural population state vectors for dimensionality reduction, we performed z-score normalization on a session-wise basis. This normalization centered the mean of each neuron's activity to zero and addressed potential variations in firing rate variability across neurons. Following this preprocessing step, we employed PCA for unsupervised dimensionality reduction on a session-by-session basis. The individual principal components obtained were then analyzed with respect to individual task-relevant or physical behavior variables. Task-relevant variables included path progress and sequence segment. The PVs were either examined without further processing or were trial-averaged according to trial-specific classes defined by the behavioral variables, such as sequence segment index and turn direction class.

4.2.5 Phase-selective neuronal modulation analysis

The objective of this analysis was to identify neurons that showed selectively modulated activity (characterized by strong and robust neuronal variability) during specific path progression phases. We focused on two distinct phases: the departure phase and the approach phase. Approach cells are characterized by a pronounced increase/decrease in firing rate as the animal approaches the reward zone, while departure cells demonstrate decaying/increasing activity upon leaving the reward zone. The four-step algorithm enforces mutually exclusive categorization, assigning each neuron to either the approach or departure cell category based on the relative strength of modulation during the two path progression phases.

Neuron selection

To prepare the 50 ms binned PVs for classification, we employed session-wise rescaling. In contrast to the normalization method used for PCA analysis, in the current analyses, each neuron's binned data was scaled by dividing by its standard deviation, without subtracting the mean. This approach avoided introducing negative firing rates, thereby preserving the interpretability of the classifier's parameters. We focused our analysis on neural activity during specific path progression phases of the task, selecting only vectors associated with the departure and approach phases (initial and final thirds of path progress in a trial). For each trial, an equal number of bins were randomly sampled without replacement from both phase classes.

The SVM method selects the neurons that are have consistent firing rate differences in the two classes, making them most informative for classifying the path progress phase. A linear SVM classifier was employed to predict the path progress phase (Section 4.2.2) associated with each PV. SVMs are supervised learning algorithms that learn a hyperplane that maximizes the margin between categories by identifying key data points called support vectors (Boser et al., 1992).

We trained the SVM iteratively for 1000 epochs using the squared hinge loss func-

tion and Stochastic Gradient Descent (SGD) learning algorithm. Additionally, to promote model prediction and interpretation, we applied elastic net regularization (Zou & Hastie, 2005) to the classifier, a technique combining L1 and L2 weight penalties, with a ratio of L1 to L2 set at 0.15, and regularization strength of 0.0001. This technique also encourages the classifier to set the weights of irrelevant features (neurons) to zero and provide an estimate of neurons selective for each phase. Finally, we evaluated the classifier's performance using ten-fold cross-validation. The average test accuracy across all the cross-validation iterations was reported. The squared hinge loss function (L) was defined as follows,

$$L(y,\hat{y}) = \max(0, 1 - (y \cdot \hat{y}))^2 \tag{4.2}$$

where, y represents the true label, and $\hat{y}$ represents the predicted label.

Following cross-validation, we identified neurons informative for each phase classification based on the magnitude of their weights across folds in the linear SVM classifier. For each neuron, the weight value from each fold was first classified as non-zero or zero based on its magnitude. We then selected neurons that were consistently classified as non-zero across all ten folds. These consistently weighted neurons were considered the most informative for predicting the path progression phases.

After initial screening of neurons using SVM classification, *Boruta* feature selection was employed to refine the selection of the most informative neurons for phase classification. Boruta is a wrapper method that uses a random forest classifier to evaluate the importance of each feature in a dataset (Kursa & Rudnicki, 2010). The method compares the importance of each neuron (feature) to that of random shuffled version of those features. Features that are more important than the noise were retained, while those that are not were removed. This process was repeated iteratively until all features were either confirmed or rejected. The Boruta method was used to validate the neuron selection process and ensure that the most informative neurons were retained for subsequent analyses. In order for the feature importance results to be reproducible, we used SHapley Additive exPlanations

(SHAP) values (Lundberg & Lee, 2017) to evaluate the importance of each feature in the dataset. This method may select cells demonstrating significant modulation in at least one sequence segment as informative.

Preprocessing for analyzing representational fidelity across domains

After obtaining the time-binned neural activity following the preprocessing described in Section 3.2.1, a representative sequence segment trial having median number of bins was selected for each sequence segment. Subsequently, distances covered between successive time bins in the median trial were averaged for each sequence segment to compute the path length bin size for distance-binning of neuronal spikes. The neural activity was computed by distance-binning the 1 ms resolution time-binned PVs, using the previously determined bin size. To account for confounds associated with pauses or speed variations, the data was normalized for occupancy rate in each distance bin.

For each sequence segment, trial with minimum number of bins in both time and distance domain was identified. The minimum number of bins was used to clip (truncate) all trials in both domains in that sequence segment to the same length. Clipping was done separately on two different terminals of the trials, i.e., trials were clipped individually from the trial end and the trial start. This was done to ensure that the ramping activity of the cells locked to either of departure or approach phases was not completely lost.

This process resulted in a matrix of binned neural activity for both time and distance domains, structured as $\mathtt{T} \times \mathtt{B} \times \mathtt{N}$, where $\mathtt{T}$ is the number of trials in the sequence segment, $\mathtt{B}$ is the number of bins (same for both time and distance domains due to clipping), $\mathtt{N}$ is the number of neurons. It's important to note that the number of bins $(\mathtt{B})$ can vary across sequence segments due to segment-specific bin size calculation. All subsequent analyses based on this data need to be performed within the context of each individual sequence segment.

Behavioral Information rate

To quantify the modulation of neuronal activity with respect to both time and distance domains, we computed the behavioral information rate as defined by Skaggs et al. (1992). The behavioral information rate quantifies the amount of information a neuron conveys about the animal's current behavioral state. While this metric was originally introduced for quantifying spatial information, the authors acknowledged its applicability to any experimentally recorded variable (Skaggs et al., 1992). Therefore, we adapted it to assess the information a neuron's firing rate conveys about its temporal state (time domain) and its spatial state (distance domain).

The information rate is calculated in bits per second and is based on the neuron's mean firing rate across all trials in a sequence segment. The behavioral information rate is computed as the mutual information between the neuron's firing rate and the behavioral state variable, normalized by the overall mean firing rate of the neuron. The preprocessed time and distance binned PVs specific to each putative phase-selective neuron and sequence segment combination obtained from Section 4.2.5, of dimensions $\mathtt{T} \times \mathtt{B} \times \mathtt{N}$, were used for the behavioral information rate computation using the following formula,

$$I_n = \sum_{i=1}^{B} \lambda_n(i) \log_2 \left(\frac{\lambda_n(i)}{\sum_{j=1}^{B} \lambda_n(j) \mathbb{P}(j)} \right) \mathbb{P}(i) \tag{4.3}$$

where,

- I_n is the behavioral information rate of the neuron n in bits per second,
- i is the current behavioral state variable bin (here, time or distance domain),
- B is the total number of bins of the behavioral state variable,
- $\lambda_n(i)$ is the mean firing rate of the cell n when animal is in behavioral state i,
- λ_n denotes the tuning curve of neuron n with respect to the behavioral state variable,

- $\mathbb{P}(i)$ is the probability density for the animal being at behavioral state i,
- $\sum_{j=1}^{B} \lambda_n(j)\mathbb{P}(j)$ is the overall mean firing rate of the cell n.

We generated a baseline null distribution of behavioral information rate (I_{null}) by randomly shuffling the bin labels associated with each PV in a sequence segment. The behavioral information rate computation procedure was then repeated for each neuron after each shuffle, resulting in a null distribution value for each neuron and sequence segment combination. This shuffling process was repeated $n_{null} = 10$ times to obtain a distribution of null information rate values, yielding a mean null information rate ($\mu_{I_{null}}$) and standard deviation of the null distribution ($\sigma_{I_{null}}$). The p-value was derived by comparing the z-score of the original information rate data point to the null distribution.

Following the identification of significantly domain-tuned neurons across both temporal and spatial domains, the initial set of phase-selective cells was refined. Only those cells exhibiting significant tuning across both domains across all sequence segments were selected as the final phase-selective cells.

Modulation Consistency across domains

The phase-selective neurons were further classified into departure or approach phase cells based on the relative strength of their modulation during the two path progression phases. This classification was determined by comparing the standard deviation of the firing rate in each phase using path-progression binned trial-averaged neural activity. The phase with the higher standard deviation was considered the phase for which the neuron was selective.

Each of the $\mathtt{T} \times \mathtt{B} \times \mathtt{N}$ matrices obtained from procedure in Section 4.2.5 for each departure or approach phase cell and sequence segment combination, were sorted on the basis of trial duration. Subsequently, they were partitioned along the first dimension, resulting in three trial speed classes: fast, medium, and slow.

For each speed class within a sequence segment, the average neuronal activity across trials was calculated for each bin. This process yielded three local averages for each neuron within the segment, effectively estimating the empirical firing rates.

To quantify the consistency of modulation dynamics across the time and distance domains, we used the Mean Absolute Error (MAE). This involved calculating the absolute difference between all possible pairs of local averages within each sequence segment for each neuron. Subsequently, the average of these absolute differences (MAE) was computed. The MAE represents the overall variability between local averages within a segment for a specific neuron.

Finally, we compared the average variability of the time domain to that of the distance domain for each phase-selective neuron and sequence segment combination.

4.2.6 Turn direction selective neuron selection

We employed a two-step analysis to investigate the relationship between neural activity and turning behavior. For preprocessing, z-score normalization was applied across all PVs for the entire recording session.

Firstly, we trained a linear SVM classifier with preprocessed PVs as input from departure and approach phases to classify the two turn directions using cross-validation for both start and next turns, respectively (defined in Section 4.2.2). The turn directions were discretized to left and right turns using the sign of the turn angle. For the training we used the PVs of departure and approach phases for the start and next turn direction classification, respectively. The model was implemented with elastic net weight regularization (Regularization strength = 0.0001). Ten-fold cross-validation was employed during training using SGD learning for 1000 iterations, and performance was evaluated using the accuracy metric. This training procedure was conducted jointly for all sequence segments.

Finally, we used the Boruta feature selection method to further prune the neurons obtained from the SVM classifier for the informative neurons for each turn direction classifi-

cation (similar to Section 4.2.5).

4.2.7 Neural population activity geometry and dynamics

Similar to preprocessing in Section 4.2.5, we performed path progression binning of the neural spike data with the omission of clipping step. Within each sequence segment, all trials were binned to a fixed length corresponding to the median trial duration for that segment. This resulted in a PV matrix with dimensions $\mathtt{T} \times \mathtt{B} \times \mathtt{N}$, where $\mathtt{T}$ represents the number of trials within the sequence segment, $\mathtt{B}$ signifies the number of bins, and $\mathtt{N}$ denotes the number of neurons.

Subsequently, z-score normalization was applied to all of the PVs across the entire session's data. To enable the comparison of neural state trajectories across similar and dissimilar conditions (e.g., turn direction action sequences) within a representational domain (e.g., path progression), we averaged PVs across trials for each condition. This averaging was performed separately for even and odd-numbered trials within each condition, allowing for the internal comparison of neural state trajectories within the same condition.

Finally, we proceeded to compare the average neural state trajectories across various conditions of turn direction action sequences along the path progression domain. These conditions were defined based on the sequence of turn directions taken by the animal at the beginning of each trial within a segment and the subsequent trial in the following sequence segment. Specifically, the conditions comprised Left-Left (LL), Left-Right (LR), Right-Left (RL), and Right-Right (RR) turn sequences. Evaluating all pairwise combinations of these four conditions rendered sixteen comparisons for each metric. Adapted from (Guidera et al., 2024), this analysis examines two key aspects of how population activity evolves in a high-dimensional space during task performance: neural representational geometry and dynamics. For both of the following metrics, we compared the average PVs for odd-numbered trials in condition *"a"* to the average PVs for even-numbered trials in condition *"b"*.

Proximity

We leveraged a proximity metric (Guidera et al., 2024) to quantify the representational similarity of neural population state trajectories across two experimental conditions within a representational domain (e.g., path progression). This metric assesses the extent to which corresponding PVs in the two trajectories tend to be co-localized within similar bins across the representational domain (Guidera et al., 2024).

To compute the proximity of neural state trajectories across two conditions (e.g., turn direction action sequences), we employed a two-step distance metric for each path progression bin. First, we computed the within-bin distance, reflecting the variability in neural states across trials at a specific bin. This was achieved by calculating the distance between the average PVs for the two conditions within that bin.

Second, we assessed the between-bin distance, capturing the average distance between neural states at distinct points in the progression. This involved calculating the average distance between the current bin's average PV in one condition and a different bin's average PV in another condition.

Finally, a proximity metric was derived to evaluate the relative similarity within and between bins. This metric was defined as 1 minus the ratio of the average within-bin and between-bin distances. Consequently, higher proximity values indicate greater similarity between average neural activity across similar stages of path progression compared to those at distinct stages.

$$\pi_i(a,b) = 1 - \frac{\left\|\vec{v}_{i,a} - \vec{v}_{i,b}\right\|_N}{\frac{\sum_{j=1}^{B} \left\|\vec{v}_{i,a} - \vec{v}_{j,b}\right\|_N}{B}} \tag{4.4}$$

where,

- $\pi_i(a,b)$ represents the proximity of neural state trajectories in bin i between conditions a and b,

- N denotes the dimensionality of the neural state space (i.e., number of neurons),

- $\vec{v}_{i,a} \in \mathbb{R}^N$ is the average neural state trajectory in N-dimensional space at bin i for condition a (only odd trials),

- $\vec{v}_{i,b} \in \mathbb{R}^N$ is the average neural state trajectory in N-dimensional space at bin i for condition b (only even trials),

- B represents the total number of bins in both neural state trajectories, and

- $\|\cdot\|_N$ denotes the Euclidean norm in N-dimensional space.

Cosine similarity

To quantify the dynamics similarity of neural population state trajectories between two experimental conditions across a representational domain, we employed the cosine similarity metric. This metric assesses the extent to which neural state trajectories evolve in parallel or divergent directions within similar bins across the representational domain. For each pair of consecutive path progression bins, we define the cosine similarity of the average displacement vectors of the two conditions. The cosine similarity is calculated as the dot product of the two displacement vectors divided by the product of their magnitudes.

$$S_i(a,b) = \frac{\vec{d}_{i,a} \cdot \vec{d}_{i,b}}{\left\|\vec{d}_{i,a}\right\|_N \left\|\vec{d}_{i,b}\right\|_N} \tag{4.5}$$

where,

- $S_i(a,b)$ represents the cosine similarity of neural state trajectories in bin i between conditions a and b,

- N represents the dimensionality of the neural state space (i.e., number of neurons),

- $\vec{d}_{i,a} = \vec{v}_{i+1,a} - \vec{v}_{i,a}$ is the displacement vector in N-dimensional space between bins i and $i+1$ for condition a,

- $\vec{d}_{i,b} = \vec{v}_{i+1,b} - \vec{v}_{i,b}$ is the displacement vector in N-dimensional space between bins i and $i+1$ for condition b, and

- $\|\cdot\|_N$ denotes the Euclidean norm in N-dimensional space.

4.2.8 Progression phase neural population geometry and dynamics

Following the calculation of proximity (represented by π) and cosine similarity (represented by S) for each pairwise combination of turn direction action sequence conditions, these metrics were averaged across progression phases. Mirroring the approach in Section 4.2.2, π and S were further divided into three bins corresponding to path progression phases: departure, intermediate, and approach. This yielded phase-specific measures of proximity and cosine similarity for each pairwise comparison of turn direction action sequences.

4.3 Results

4.3.1 Dimensionality reduction reveals neural organization of behavior

In order to identify the core behavioral variables reflected in the activity patterns of the mPFC neural ensembles and to generate insights into the organization and evolution of neural activity in relation to these behaviors, we performed dimensionality reduction of the neural population state vectors. For this, we carried out PCA on the simultaneously recorded single-units from each maze session. PCA was applied to the z-scored spiking activity binned into 100 ms non-overlapping windows, referred to as PVs (Section 4.2.4).

To identify the behavioral variables most likely orchestrating the population-level activity patterns in the mPFC, we first employed multiple linear regression to decode sensorimotor and task-specific variables from the PVs (Figure 4.1A). Decoding performance revealed strong explanatory power for path progression (Section 4.2.2) and running speed within the sequence. Conversely, decoding of sequence progress and turn direction at reward zones (initial and upcoming turns) yielded weaker results. Interestingly, the decoder

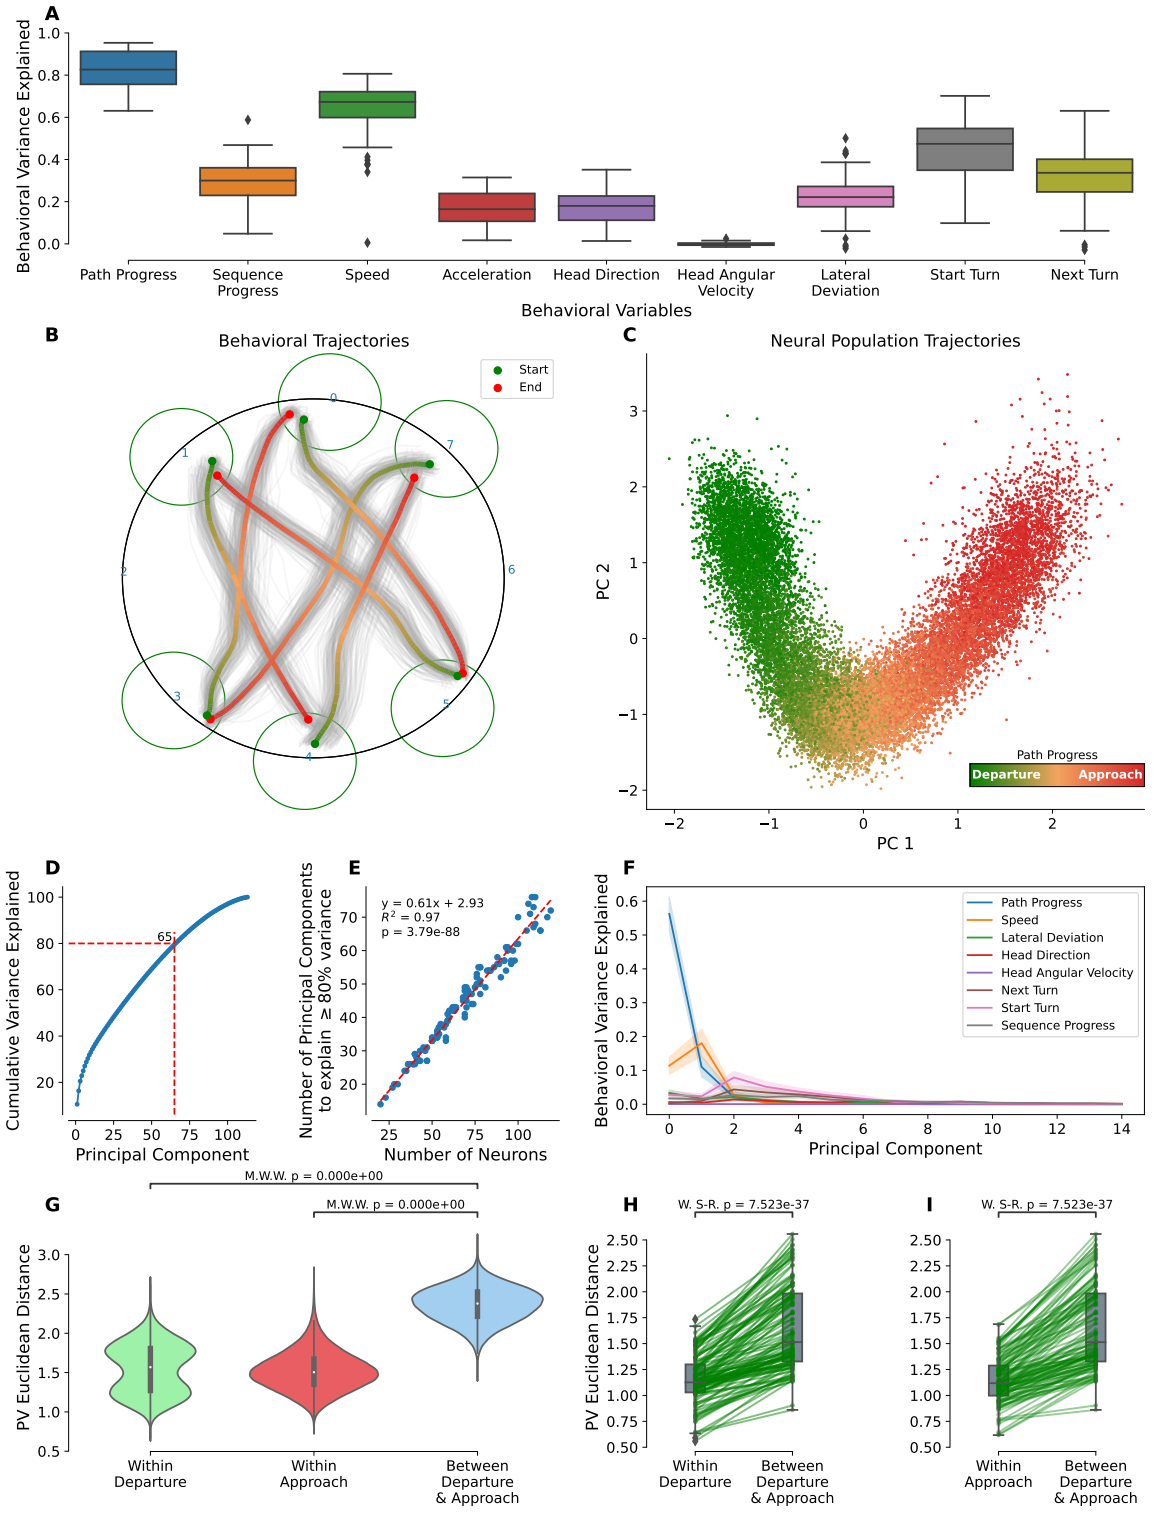

Figure 4.1: Neuronal population-level tracking of progress along intermediate reward zones. (**A**) Explained variance of linearly decoded behavioral variables computed via coefficient of determination scores from PVs. (**B**) Behavioral trajectories on a 6-step sequence session. The median trajectory is color-labeled by path progress and remaining single trial trajectories represented in gray. Green circles indicate the radius of the reward triggering

zone with numbers indicating the center of individual reward zones. (**C**) First two Principal Components (PC) of population activity from a representative session show segregation of departure phase activity from approach phase population activity. Phase division was made by dividing the proportionate distance to reward goal in 3 equal portions each corresponding to Departure, Intermediate, and Approach phases. (**D**) Cumulative variance explained by neurons in a representative session with first 65 PCs explaining 80 % of the variance. (**E**) Relationship of the minimum number of PCs needed to explain 80 % of the variance to the number of neurons across all sessions. The dashed line represents the linear regression fit (slope = 0.61, R^2 = 0.97). (**F**) Explained variance of linearly decoded behavioral variables from each of the first 15 PCs computed via coefficient of determination scores across all 120 sessions. The color coded line represents the mean and the associated shaded area represents the 95 % confidence interval. (**G**) Euclidean distance distributions within and between PVs of departure and approach phase classes for a representative session (two-tailed Mann-Whitney U test; $p < 0.001$). Violin plots represent the distribution of distances within and between phase classes. The white dot represents the median, and the thick black bar represents the Interquartile Range (IQR). The thin black lines extend to the most extreme data points not considered outliers (within 1.5 times the IQR). (**H** and **I**) Pairwise comparisons of average PV Euclidean distances within and between phase classes across all 120 sessions. Each line represents one session. Box plots represent the distribution of distances within and between phase classes. The box represents the IQR, the line within the box represents the median, and the whiskers extend to the most extreme data points not considered outliers (within 1.5 times the IQR). Outliers are represented by diamonds. (H) Within departure phase vs. between departure and approach phases (two-tailed Wilcoxon signed-rank test; $p < 0.001$). (I) Within approach phase vs. between departure and approach phases (two-tailed Wilcoxon signed-rank test; $p < 0.001$).

explained only 20 % of the variance in lateral deviation along the sequence path. This variable was previously reported in our data to confound putative decision or working memory signals in the mPFC that predicted upcoming turn direction (Euston & McNaughton, 2006), suggesting a motoric nature of the activity.

We examined the cumulative variance explained by the PCs. The observed pattern showed a steep monotonic increase in the explained variance with each additional PC (Figure 4.1D). To estimate the dimensionality of the mPFC population activity associated to our task, we calculated the minimum number of PCs required to explain 80 % of the variance across all sessions (Figure 4.1E). A strong positive correlation was observed between the two. This suggests that PCA fails to capture a low-dimensional subspace that explains majority of the variance in the neural data.

This finding shows that PCA is unable capture a low-dimensional subspace that explains the majority of the variance in the neural data. This implies a relatively low degree of linear correlation between mPFC neurons, which limits the effectiveness of PCA in identifying a subspace that best represents the underlying structure of the data. Alternatively, the relationships between neurons may be non-linear, forming a non-linear manifold in the neural state space. Recent studies have employed non-linear dimensionality reduction methods to address this limitation (Sun et al., 2023; Guidera et al., 2024), potentially offering a more accurate representation of the underlying neural dynamics. Notwithstanding, A qualitative inspection of the three most informative principal components revealed two consistent organizing principles across most sessions: (1) First, the neural population activity along the first two principal components was organized on the basis of progression between reward locations along individual segments of the larger sequence, which we refer to as *path progression* (Figure 4.1B-C). (2) Second, the neural population vectors were dissociated along the third principal component on the basis of the turn direction of the animal upon exiting reward zones to initiate the next sequence segment (Figure 4.3A).

We used linear regression to decode behavioral variables from each of the first 15 PCs (Figure 4.1F) to gain quantitative insights into how the organization of the neural population activity in the PC space relates to the behavioral variables. The results showed highly consistent decodability of behavioral variables from the first few PCs. Specifically, the first PC explained on average 56.08 % (standard deviation = 28.22 %) of variance of path progression. The third PC was consistently more predictive of turns taken at reward zones to continue for next path traversal in a both retrospective and prospective manner.

4.3.2 Progress within sequence segments

The organization of the PVs on the basis of path progression appears to evolve along a curved neural manifold in the first two principal components (Figure 4.1C,F). These results are consistent with the behavioral variables shown to govern the organization of neural ac-

tivity when visualized using non-linear dimensionality reduction methods (Guidera et al., 2024). Furthermore, our results of separability of departure and approach phase population activity was consistent for smaller bin size of 50 ms without any Gaussian smoothing. This sheds more light on similar progress encoding reported by Guidera et al. (2024) as due to the 100 ms standard deviation Gaussian smoothing window used in their analysis, it was unclear if the predictability of progress could be attributed to the Gaussian smoothing which could have introduced temporal correlations in their data, hence potentially portraying a position consistent activity as progress encoding. Our results suggest that the mPFC encodes progress in a manner that is consistent across sequence segments, and that this encoding is not an artifact of temporal correlations in the data.

We discretized the path progress into three phases: departure, intermediate, and approach (Section 4.2.2). One distinctive characteristic consistently observed in the first two PCs of the neural state space was the divergence of PVs associated with the departure and approach phases (Figure 4.1C). The separation between departure and approach phase PVs was quantified using Euclidean distances between PVs in the neural population state space. We observed significantly higher average distances between the departure and approach phases compared to within-phase distances (Figure 4.1G). This suggests that the neural population activity patterns associated with the departure and approach phases are distinct and well-separated in the neural state space. This separation was consistent across all sessions, as evidenced by the pairwise comparisons of average PV Euclidean distances within and between phase classes (Figure 4.1H,I).

4.3.3 Time vs Distance from/to goal: Neuronal modulation fidelity across trials

Encouraged by the success of the Euclidean distances in quantifying the phase-specific separation of PVs, we extended our analysis of progress encoding to the cellular level. One question that stemmed from the observed correlations with path progression phases concerned the underlying mechanism of the neural organization governing the progress

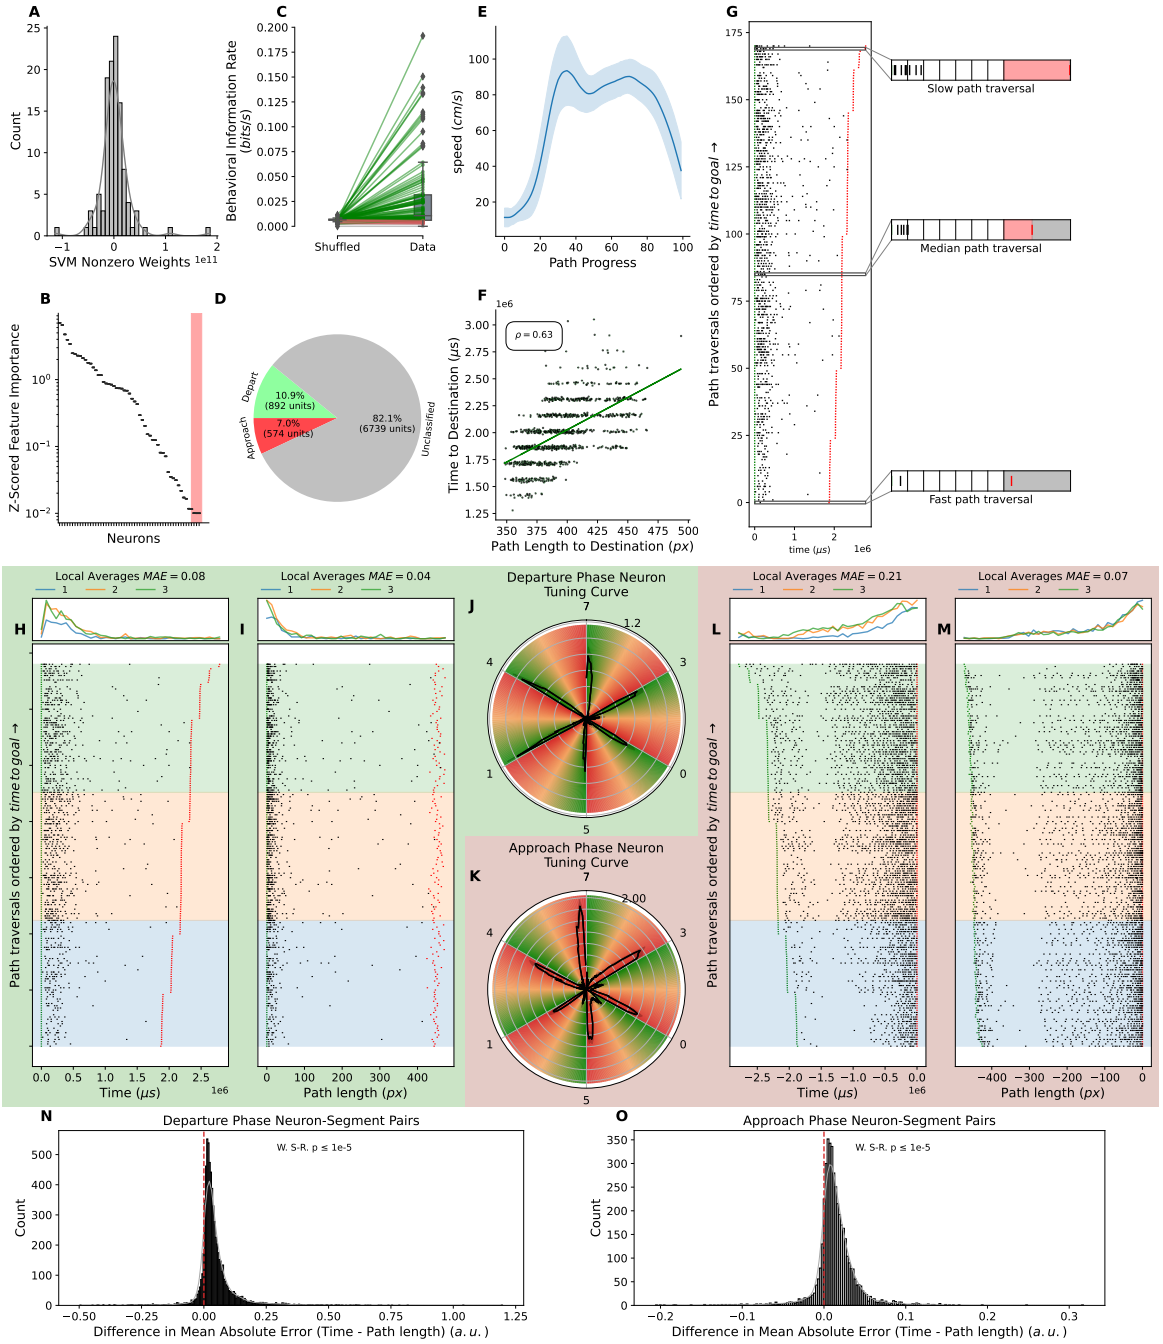

Figure 4.2: Comparison of representational consistency in domains of distance and time reveals more spatially consistent code. (**A**–**D**) Four step pipeline to select cells consistent for departure and approach phases and time and distance physical variables. (A) Representative feature weights obtained from SVM classification of departure and approach phase cells on a single session. (B) Feature significance obtained from Boruta feature selection method on a representative session to identify most informative neurons for each phase classification. Red region denotes rejected features. (C) Behavioral information rate criterion applied to the units binned according to both time and path length to identify neurons that were selective for both time and distance domains. (D) Obtained phase selective classification of units across all sessions (n = 120) after comparing the magnitude of standard deviation of

firing rates for each phase. (**E**) Dynamics of speed across path traversals from a representative session collapsed across all sequence segments. The line represents the mean and the shaded region represents the 95 % confidence interval. (**F**) Relationship between time and distance to cover to get to reward zone from a representative session across all sequence segments. The dashed line represents the linear regression fit (slope = 0.63). (**G**) Illustration of the binning method for the analysis. Raster plots of a representative departure phase selective neuron. Black lines denote the spikes and green and red lines denote the trial starts and end points, respectively. Insets: Red denotes clipped time window. (**H**–**J**) Representative departure phase cells (or reward departure cells) show differing levels of consistency in time and distance domains quantified by MAE between firing rates of local averages created from fast, intermediate, and slow trials. Corresponding trials in ramp plot and spike raster are similarly color-coded. (**K**–**M**) Representative approach phase cells (or reward approach cells) showing similar measures. (**N** and **O**) Paired differences of MAE between time and distance domains for phase-locked cells and sequence segment pairs across all sessions. Each data point represents one neuron across one sequence segment. (N) Departure phase cells (one-tailed Wilcoxon signed-rank test; $p < 0.001$; alternative hypothesis: MAE time < MAE path length). (O) Approach phase cells (one-tailed Wilcoxon signed-rank test; $p < 0.001$; alternative hypothesis: MAE time < MAE path length).

encoding. Specifically, what physical variables are likely to be tracked by the neurons to achieve this encoding? Distance and time emerged as prime candidates. In previous studies different assumptions about which of the physical variables of distance or time are being used to encode the progress between the rewards were made. For example, El-Gaby et al. (2023) used normalized time to define progress between rewards and Guidera et al. (2024) used uniform distance binning to define progress between rewards. This inconsistency in the definition of progress between rewards makes it difficult to precisely test the progress encoding hypothesis on navigational paradigms with variable distances between rewards. To address this, we need to pin down the physical variable that is being used by mPFC to encode progress between rewards.

On inspection of individual neurons we saw distinctive ramping of neurons upon departure from the reward zone and approach to the reward zone (Figure 4.2J,K). Although this reward-associated activity can be readily detected in a significant proportion of the neurons in mPFC, encompassing both reward departure and potentially, reward anticipation, a key question remains: whether this modulation in firing rate is occurs consistently in time

or distance? Elucidating this distinction is critical because these factors tend to co-vary in most behavioral paradigms and are often overlooked. In order to disentangle these influences, we leveraged the high trial count, well-defined reward epochs (Medial Forebrain Bundle (MFB) stimulation duration), and natural speed variations during path traversals across hundreds of trials in each session (Figure 4.2E,F). This paradigm provided a wide range of running speeds, enabling the effective separation of time and distance contributions to peri-reward neural activity.

To identify cells with ramping activity tied either to departure from or approach towards reward zones, we used a four step pipeline (Figure 4.2A–D). First, we selected the neurons (features) important for distinguishing between departure and approach phases in the neural state space using a pipeline of linear SVM classifier and *Boruta* feature selection trained to classify the two phases (Section 4.2.5). This served as an initial screening step to subsample only the neurons informative for encoding of path progression phases using a simple linear method, hence decreasing the computational cost of the subsequent complex steps.

For the selected informative neurons, we computed binned activity within both the temporal and spatial domains, ensuring an equal number of bins across trials for each domain (Figure 4.2G; Section 4.2.5)

To further refine our selected putative phase-selective cell population for possessing significantly consistent modulatory dynamics (ramp-like tuning) across both time and distance domains, we employed a behavioral information-rate metric (Skaggs et al., 1992) to compute the informational content of selected neurons for both time and path length for each sequence segment using the equivalent binned activity. Neuron selection was based on the criterion that its behavioral information rate significantly exceeded the null distribution obtained for each combination of neuron and sequence segment. The null distribution was determined using a shuffling procedure (n = 100 shuffles each). This test was done for each neuron in both time and path length domains and across each sequence segment. The obtained neurons were identified as putative neuronal subpopulations selective for each

phase (departure or approach) using the magnitude of standard deviation of their firing rates across individual neuron's trial-averaged activity in each of departure and approach phase class as a proxy measure for the modulation index. This assumed that the both phase selective classes of neurons were mutually exclusive which may not be the case. However, this assumption was made to simplify the analysis.

Building upon the identified departure- and approach-selective neuronal ensembles, we tried to unveil differences in modulation dynamics across temporal (time) and spatial (distance) domains using a custom modulation consistency domain analysis, detailed in Section 4.2.5. This method uniformly categorized trials in each sequence segment into fast, medium, and slow groups and computed the group-wise trial-averaged data. Pairwise comparison of the average MAE between group-wise averages obtained for each bin in the phase class revealed that the modulation dynamics of single cells selective for both departure and approach phases exhibited greater consistency across trials in the distance domain compared to the time domain (Figure 4.2H–O). Thus, upon testing the single unit-level modulatory dynamics responsible for such progression encoding in time and distance domain, we found that distance (rather than time) encoding provides a better explanation for the reward-associated cell modulation across the trials.

4.3.4 Neural population activity geometry and dynamics on the basis of turn direction action sequences

Another behavioral correlate of mPFC population activity appeared to be speed and turn direction. While locomoting between two reward zones, the speed of the animal varies quite stereotypically as shown in Figure 4.2E. The speed of the animal is lowest at the beginning of the trial, increases as the animal progresses towards the reward zone, and decreases as the animal approaches the reward zone. This stereotypical speed profile is consistent across all sequence segments and sessions. This, in turn, makes the speed closely entangled with the progression between rewards and disentangling the two is a challenging task.

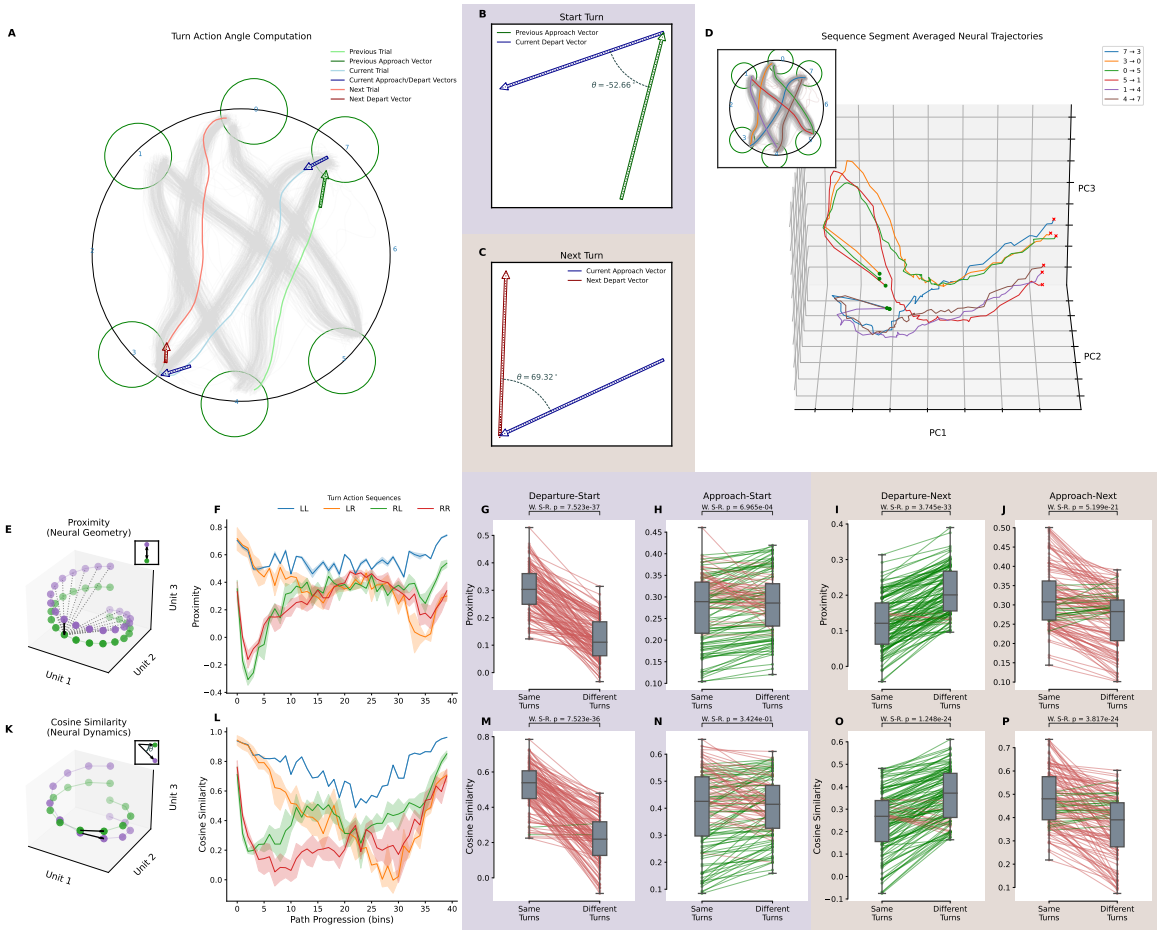

Figure 4.3: Turn action encoding in population geometry and dynamics. (**A**–**C**) Representative example of process of start and next turn direction computation. (**A**) Positional trajectories corresponding to three consecutive trials are shown in the maze environment. The trials are called Previous (colored: light green), Current (colored: light blue), and Next trials (colored: light red), referring to the chronological order. Each trajectory is overlaid with the departure and goal approach vectors used to compute the start and next turn, respectively. (**B**) Aligned previous trial goal approach and current trial departure vectors are shown. The angle between these vectors is computed to determine the start turn of the current trial. (**C**) Aligned current trial goal approach and next trial departure vectors are shown. The angle between these vectors is computed to determine the next turn of the current trial. (**D**) Trial averaged neural population trajectory in the first 3 PCs color coded by individual sequence segment for a representative 6-step sequence maze session. Green circles and red cross indicate the start and end of the trial averages, respectively. Inset: Spatial trajectories of the animal for the representative session in the maze environment color coded by individual sequence segment. Gray lines represent the animal's path in individual trials. Black circle denotes the circular tabletop's boundaries. Green circles represent the radius of the reward triggering zone with the blue numbers indicating the center of individual reward zones. (**E**–**P**) Proximity and cosine similarity metrics were used to quantify the representational organization and dynamics, respectively, of PVs with respect to different turn direction action sequences. (E and K) Illustration of the proximity metric (E) and cosine similarity metric (K) computation for a single bin between two neural trajectories.

(F and L) Values of the proximity (F) and cosine similarity (L) metrics across the path progression with PV trajectories groups according to the identity of the start and next turn action sequences (Turn sequences: LL (Start turn = Left; Next turn = Left), LR (Start turn = Left; Next turn = Right), RL (Start turn = Right; Next turn = Left), RR (Start turn = Right; Next turn = Right)). The metrics were compared using the average PV trajectory of group LL as baseline to average of all remaining turn action sequence grouped trajectories in a representative session. The proximity metric was computed for each bin and averaged across all bins. The color coded lines represents the mean proximity across all bins from the associated group's trajectory. The shaded region represents the 95 % confidence interval. (G–H and M–N) Pairwise comparisons of average proximity (G and H) and cosine similarity (M and N) metrics within same *start turn* and different start turn direction action across all sessions. Pairwise comparison of averaged metrics in departure phase bin (G, M). Pairwise comparison of averaged metrics in approach phase bin (H, N). Each line represents one session. The box represents the IQR, the line within the box represents the median, and the whiskers extend to the most extreme data points not considered outliers (within 1.5 times the IQR). Outliers are represented by diamonds. (I–J and O–P) Pairwise comparisons of average proximity (I and J) and cosine similarity (O and P) metrics within same *next turn* and different next turn direction action across all sessions. Pairwise comparison of averaged metrics in departure phase bin (I, O). Pairwise comparison of averaged metrics in approach phase bin (J, P). Each line represents one session. The box represents the IQR, the line within the box represents the median, and the whiskers extend to the most extreme data points not considered outliers (within 1.5 times the IQR). Outliers are represented by diamonds.

On the other hand, the turn direction of the animal upon exiting a reward zone to initiate the next sequence segment is a variable more amenable to interpret the neural activity. Hence, we focused on understanding the neural population structure and changes with respect to the turn direction. Literature suggests that the mPFC is involved in encoding an abstract representation of the animal's actions and these representations seem to depend both on prior and upcoming actions, suggesting both a prospective and retrospective code (Tang et al., 2023; Guidera et al., 2024). Thus, we hypothesized that the mPFC population code would be organized with respect to the directionality of the turns during both prior and upcoming turns.

We investigated how neural population activity organizes and evolves within the neural state space across different turn direction sequences. A turn action was defined as the direction of the animal's turn upon exiting a reward zone to initiate the next sequence segment.

The encoding for this turn was tested in two contexts: the departure and approach phases. These turns were referred to as start turn and next turn, respectively (Figure 4.3B–D). Our observations indicated that the animal executed motoric sequences associated with the turn during the departure phase, while potentially planning the turn during the approach phase. Building upon previous research suggesting an abstract (superordinate-level) action code in the mPFC (Tang et al., 2023; Guidera et al., 2024), we operationalized our previous hypothesis predicting that the mPFC population structure and dynamics would exhibit greater similarity with respect to turn directionality during both the start and next turns.

To quantify these properties, we employed proximity and cosine similarity metrics (Guidera et al., 2024) to assess the representational similarity and trajectories of neural states across various turn sequences (Figure 4.3E–T).

These bin-specific measures were then averaged in each of the departure and approach path progression phases (refer Section 4.2.2). Subsequent comparisons were conducted based on individual turn directions (start and next turn) within each turn action sequence at every path progression phase. The figure panel presents a detailed comparison of proximity and cosine metrics under different conditions ("Same Turns" and "Different Turns") across all combinations of progression phases and turn actions (Figure 4.3G–L,O–T). Each sub-figure represents a specific stage or comparison point, with individual data points connected by lines to illustrate changes between conditions. Each data point represents a single session, and the box plots show the distribution of the metric values across all sessions. The results are statistically analyzed using two-tailed Wilcoxon signed-rank tests, and p-values are provided for each comparison.

In the departure phase the PVs are consistently closer for the same start turn direction ("Departure-Start"), quantified via the proximity metric showing a significant difference between "Same Turns" and "Different Turns" conditions (Figure 4.3G; $p < 0.0001$), indicating lower proximity values in the "Different Turns" condition. Similarly, the cosine metric (Figure 4.3O) reveals a significant difference ($p < 0.0001$) with lower cosine val-

ues for "Different Turns." This suggests that the PVs move across the neural state space in a more parallel manner for the same start turn direction compared to different start turn directions.

In the approach phase, the proximity metric (Figure 4.3I) shows an opposite trend compared to the departure and intermediate phases, with a significant difference between "Same Turns" and "Different Turns" conditions ($p < 0.0001$), indicating lower proximity values for "Same Turns." Similarly, the cosine metric (Figure 4.3Q) also reveals a significant difference ($p < 0.0001$) with lower cosine values for "Same Turns." Hence, the synchrony in the neural state space for the same start turn direction in the departure phase is lost in the approach phase.

For the next turn direction, the departure phase shows a significant difference between "Same Turns" and "Different Turns" conditions for both the proximity metric (Figure 4.3J; $p < 0.0001$) and cosine metric (Figure 4.3R; $p < 0.0001$), with lower proximity and cosine values for "Same Turns." This suggests that organization and dynamics of PVs are not synchronized for the same next turn direction compared to different next turn directions in the departure phase.

In the approach phase, the proximity metric (Figure 4.3L) shows a significant difference between "Same Turns" and "Different Turns" conditions ($p < 0.0001$), with lower proximity values for "Different Turns." The cosine metric (Figure 4.3T) also reveals a significant difference ($p < 0.0001$) with lower cosine values for "Different Turns." This suggests that the PVs have been significantly synchronized on basis of the same next turn direction compared to different next turn directions in the approach phase.

To summarize, our results suggest that neural state trajectories in the mPFC are organized and evolve in a more synchronized manner for the same turn direction compared to different turn directions. This organization and evolution are more pronounced in the departure phase compared to the approach phase. These results indicate that mPFC neural activity encodes an abstract representation of the animal's actions, consistent with turn

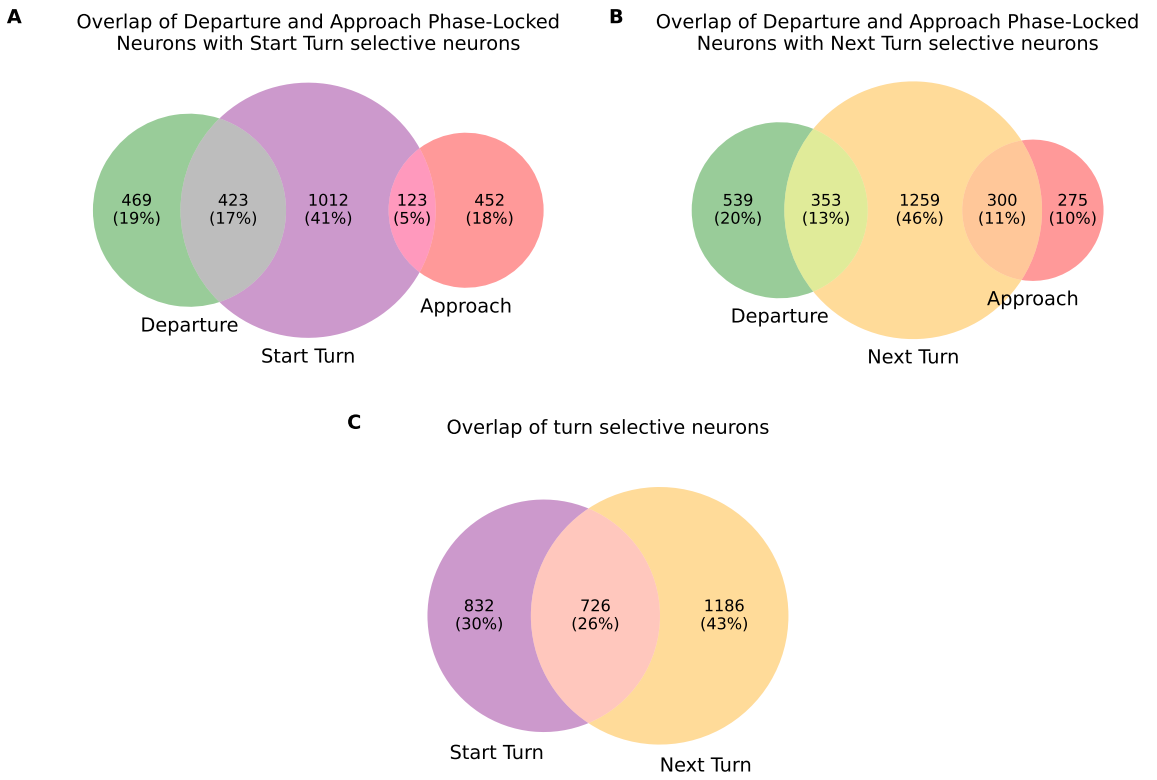

Figure 4.4: Venn diagram of units with turn and phase selectivity. (**A**) Overlap of putative neurons selective for start turn direction and phase across all sessions. (**B**) Overlap of neurons selective for next turn direction, departure phase, and approach phase across all sessions. (**C**) Overlap of neurons selective for start turn direction and next turn direction across all sessions. (A–C) The numbers and percentages in the Venn diagrams represent the number of putative neurons selective for each condition and the overlaps between them.

direction sequences. This activity organizes and evolves within the neural state space in a manner specific to the animal's current and upcoming turn actions, suggesting a prospective role for the mPFC in planning.

4.3.5 Comparison of neuronal turn selectivity with phase selectivity

In order to characterize the selectivity of individual neurons for path progression phase and turn direction, we classified the neurons most informative for the start and next turn directions and compared it with the neurons we previously found to be selective for the departure and approach phases.

We used the first two steps of the same four-step pipeline as in Section 4.3.3 to select

the neurons informative for predicting both start and next turn directions using as input PVs corresponding to departure and approach phases, respectively (Section 4.2.6).

We then compared the selected neurons with the neurons selective for the departure and approach phases. The results showed that a significant proportion of neurons selective for the start turn direction overlapped with the neurons selective for the departure phases than the approach phase (Figure 4.4A). However, the differences in the overlap between the neurons selective for the next turn direction and the departure and approach phases were less pronounced (Figure 4.4B). Moreover, in both comparisons there was only partial overlap between the turn selective neurons and the phase selective neuronal population (Figure 4.4A,B), as opposed to a complete overlap which would be expected if the phase selectivity was completely explained by the turn direction selectivity. The existence of a substantial population of neurons informative for turn direction but not for progression phase, and vice versa, suggests that the phase selectivity cannot be completely explained by the turn direction selectivity. Thus, it is possible that the two behavioral variables are encoded by partially overlapping yet distinct neural populations. It is possible that this lack of complete overlap may be due to limitations in feature selection methods as our criteria for selecting informative neurons for turn direction was more permissive than our criteria for selecting informative neurons for phase selectivity. It is possible that a more stringent feature selection or employing pruning of sparse firing neurons via threshold methods may reveal a complete overlap between the two populations.

A comparison of the neurons selective for the start and next turn directions revealed a significant overlap between the two populations (Figure 4.4C). This suggests that the neurons selective for the start and next turn directions are either not mutually exclusive, or that the turn action sequences made by the animal (owing to constraints in task structure) may lead to some predictability of the next turn direction from the start turn direction, or vice versa.

4.3.6 Motoric confounds to progress and abstract turn action encoding

While our task did not explicitly impose it, our animals showed natural variations in their trajectories in some sequence segments with respect to the start turn direction. On a qualitative inspection of neurons that were selective for the next turn direction, we observed the single neuronal activity on the sequence segment preceding such sequence segments with variations in turn at zones. We specifically analyzed neurons that were identified as important contributors in a classifier designed to predict the direction of the next turn (Section 4.3.5). Given the apparent geometrical and dynamical similarities at population level for the next turn direction encoding, we hypothesized that the single units predictive of next turn should be able to differentiate between these variations in turns made at the reward zone of next sequence segment in the special case where the animal made two different turns from the same reward zone.

Surprisingly, we observed that the single units predictive of the next turn direction were unable to differentiate between the variations in turn at the reward zone of the next sequence segment. This was evident from the lack of significant differences in the firing rates of these neurons across the different turn variations (Figure 4.5A).

To test this across the sessions, we quantified the probability of next turn direction (Left (L) or Right (R)) for each sequence segment and selected the sequence segments with sufficient variability in turn actions (Figure 4.5B). By sufficient variability, we mean that sufficient number of trials exist in which the animal makes either left or right next turns. We then computed the accuracy of the PVs to predict the next turn direction across all qualified sequence segments. The results showed that the distribution of predictive accuracy of next turn from PVs prominently peaked near the chance level (Figure 4.5B).

This suggests that the neural activity predictive of the next turn direction is minimally influenced by the future turn direction of the animal. Instead, it appears for some of the neurons that the activity may be influenced by the preparatory movements made by the animal before approaching the reward zone. Although there was a lot of heterogeneity,

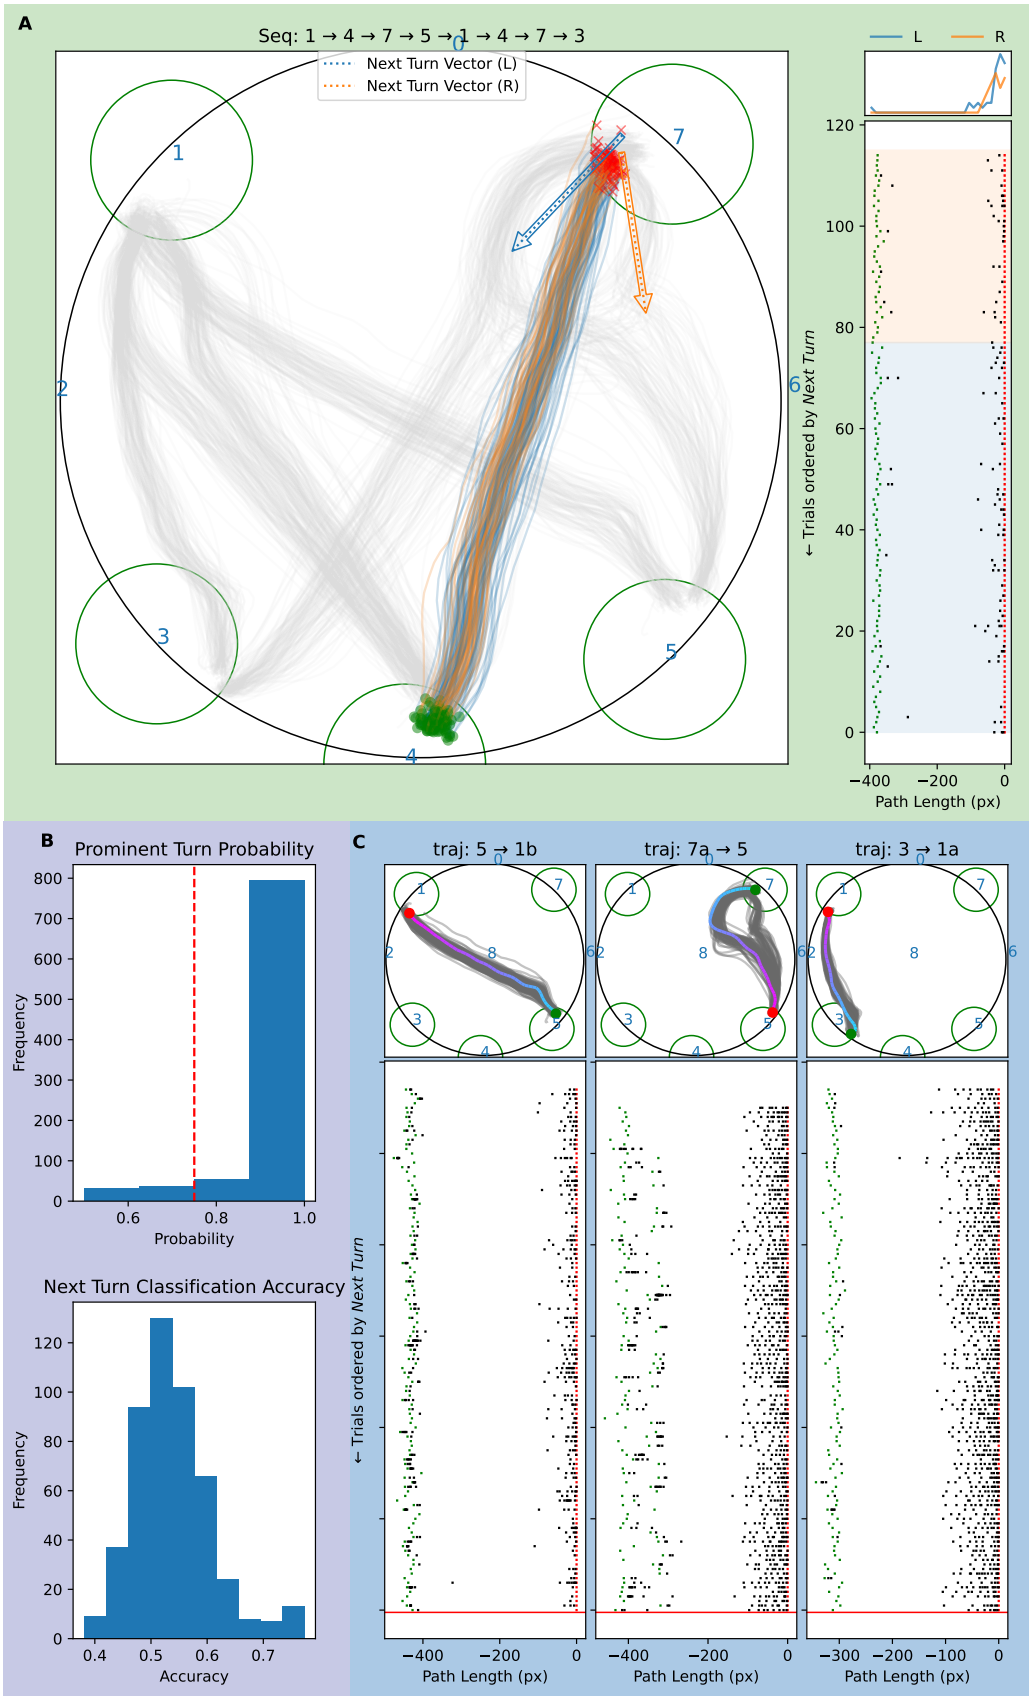

Figure 4.5: Next turn discriminability. (**A**) Single trials of a sequence segment preceding a

sequence segment with variations in turn at the reward zone from a representative session (rat: 7165; session: 44). Color code represents the side of turn in the next trial. *Left*: Positional trajectories taken by the animal. The color coded arrows represent the average direction of the next left and right turns from the reward zone. Green circles and red cross indicate the start and end of the trials, respectively. Gray lines represent the animal's path in remaining trials. Black circle denotes the circular tabletop's boundaries. Green circles represent the radius of the reward triggering zone with the blue numbers indicating the center of individual reward zones. *Right*: Raster plot of a representative next turn predictive neuron all trials of the corresponding sequence segment with trials sorted on the basis of the side of turn in the next trial. Black lines represent spike events and green and red lines represent the trial start and end points, respectively. The background color represents the groups corresponding to the side of turn in the next trial (Bottom). Average firing rates of the neuron across each group (Top). (**B**) *Top*: Histogram of the probabilities of making the most probable next turn action across sequence segments from all sessions. The dashed line represents the ceiling level of probability used to classify sequence segments as having sufficient variability in turn actions. *Bottom*: The distribution of accuracy of the PVs to predict the next turn direction across all qualified sequence segments that preceded sequence segments with significant variability in turn actions. The dashed line represents the chance level of accuracy. (**C**) Representative modulation of a next turn predictive neuron (same neuron as in A) classified as predictive of next turn across three sequence segments. Left to right shows sequence segments with variations in amount of preparatory turns made by the animal before approaching the reward zone. Top shows the positional trajectories and bottom shows the raster plot of the neuron.

qualitative inspections of few neurons seemed suggestive of this. For example, a neuron that was predictive of the next turn direction when evaluated for differences in firing patterns predictive of the next turn direction failed to show any differences for variations in next turn direction. However, the same neuron when evaluated across three sequence segments in the same session showed modulation that was consistent with the preparatory movements made by the animal before approaching the reward zone. This is apparent from the observation that the neuron increases its firing rate as the curvature of the trajectories approaching the reward zone increases (Figure 4.5C). Overall, this suggests that the neural activity predictive of the next turn direction may be confounded by the preparatory movements made by the animal before approaching the reward zone.

4.4 Conclusion

The present study aimed to investigate the neural representation of behavior within the mPFC during a spatial sequence task. We employed a combination of analyses previously used in the literature to uncover the neural organization of behavior within the mPFC with respect to progress between rewards and abstract turning action.

We found that the mPFC neural population activity is organized based on path progression, with distinct ensembles modulating during the departure and approach phases. These ensembles exhibit consistent modulation dynamics across distance domain, rather than time. Additionally, although our initial population-level analyses suggested that the mPFC encodes abstract turn direction action sequences in a manner that reflects the animal's current and upcoming actions, a deeper look suggests that, contrary to the population level analysis results, the single unit activity predictive of the next turn direction is minimally influenced by the future turn direction of the animal. Instead, it appears that all these interpretations are confounded by the preparatory movements made by the animal before approaching the reward zone. It remains to be tested if the preparatory movements can completely explain the neural activity predictive of the next turn direction.

Given the limitations in behavioral data, we were unable to further investigate the nature of these preparatory movements. It is possible that the precise motor control of the limbs of the animal was a major influence on the neural activity predictive of the next turn direction and due to highly stereotyped actions across the task such preparatory movements give a semblance of the neural activity being predictive of the next turn direction.

This is not to say that the neural activity was explained away by the motoric encoding, there was a significant diversity in the neural firing patterns. Some neurons appeared to be trajectory selective but it was hard to determine if this selectivity was due to the preparatory movements or representative of some task-related cognitive variable.

Chapter 5

Discussion

This thesis investigates the theories of abstraction of task-space and actions within the Medial Prefrontal Cortex (mPFC) during reward-guided spatial sequence navigation. Using a multi-pronged approach, we built upon the work of Mashhoori et al. (2018), who posited that apparent position-correlated activity in the mPFC reflects a generalized spatial code multiplexed with task structure. The present work aims to put their hypothesis and methodology to test on an elaborate navigation paradigm and re-evaluate their interpretations using descriptive analyses. Our results indicate that the mPFC does not encode a spatial map of the environment nor does it track the animal's progress throughout the entire task sequence. Instead, the position consistent firing of mPFC neurons appears to be more closely linked to its role in tracking inter-reward progress and actions. The action representations appear to be confounded by the motor variables which occur in a spatially-consistent manner during highly stereotyped trajectories, rather than an abstract-level action code. Our findings highlight the critical importance of considering consistent behavioral patterns when interpreting mPFC neural activity in well-trained animals. This emphasis aims to inform the design of future experiments that can disentangle the cognitive and motoric roles of the mPFC in spatial navigation tasks.

We started by attempting to extend the findings of Mashhoori et al. (2018) to a more complex spatial sequence navigation task. However, the methodology of decoding the animal's spatial position (as a proxy for abstract task-space position) from mPFC activity using Deep Neural Network (DNN) did not scale effectively to this more complicated task even

after significant improvements in DNN architecture and training procedure (Chapter 3). Following the inability of our decoding method to obtain performance consistent with previous reports, we adopted a data-driven approach to investigate the relationship between mPFC activity, behavior, and task-specific variables during the sequence navigation task. We employed both population and single-unit level analyses, revealing an apparent encoding of progress between subsequent rewards and abstract turning actions at reward zones. However, the observed sensitivity of mPFC neurons to the animal's motoric actions necessitates caution in interpreting these findings as evidence of an abstract task-space code or hierarchical action code in the mPFC (Chapter 4).

5.1 Rethinking the position-correlated representations in the medial prefrontal cortex

Contextualizing our results within the broader literature on the function of rodent mPFC, necessitate a reevaluation of existing interpretations of position-correlated representations. We see an immense potential for an alternative interpretation of the mPFC's position-correlated activity.

5.1.1 Spatial code in the medial prefrontal cortex

The inconsistent decoding of animal position from mPFC activity across sessions in our study, coupled with results inferior to previous reports, challenges the hypothesis that the rodent mPFC maintains a persistent spatial representation approximating the task environment. This finding contrasts with previous reports of successful position decoding (Mashhoori et al., 2018). A previously reported result from our data suggests that the mPFC may encode the systematic variations in the trajectory (Euston & McNaughton, 2006). However, we suspect that the motoric turn action encoding observed in our current results may be a better representative of the neural correlates of these systematic deviations. Specifically, our task's repetitive structure induces consistent motor actions, such as turns, at specific

spatial locations. Consequently, neural responses to these action-related variables become inherently linked to specific spatial coordinates, potentially mimicking spatial coding.

While Sauer et al. (2022) reported significant decoding accuracy of position information in the mPFC may seemingly undermine our position, we posit three key arguments against their findings directly contradicting ours:

1. Previous studies observed no spatial specificity in mPFC neurons during spontaneous behavior (Kaefer et al., 2020; Jung et al., 1998; Poucet, 1997), suggesting that the evidence of spatial information in their one-dimensional, circular environment requires further scrutiny.

2. The reported accuracy in Sauer et al. (2022), while statistically significant, is insufficient to definitively demonstrate a spatial code. Their reported accuracy of 0.1 translates to only 1 to 2 correctly predicted spatial bins out of 15 to 20 possible bins, given their bin size and arena size. Such a level of accuracy is inadequate to substantiate the presence of detailed spatial representations within the mPFC. Instead, these findings indicate a potentially coarse spatial code, if one exists at all.

3. It is likely that the mPFC neurons might encode the sensory perception of cues used for navigation. Several studies have reported sensorimotor information in mPFC (Cowen & McNaughton, 2007; Hyman et al., 2012) and cue-selective ensembles in other neocortical regions (Cowansage et al., 2014; Chang et al., 2023). Whether mPFC encodes cue perception (visual or tactile) or spatial location of cues remains to be elucidated. With the precise control of the trajectory of the mice in the virtual reality environment used by Sauer et al. (2022), the animals will always perceive the same cues at the same location. This presents a potential confound in the interpretation of the spatial selectivity of mPFC neurons. The sensory cue encoding hypothesis could also explain the restoration of spatial selectivity observed by Sauer et al. (2022) when the animal returned to the familiar environment. Nonetheless, further investigation is

needed to confirm this hypothesis.

In conclusion, our findings do not support the presence of an allocentric spatial code in the mPFC during reward-guided navigation tasks. Notwithstanding, the mPFC may still possess position-correlated activity that could have spatial origins, however, evidence of such a code should control for the sensorimotor information confounds arising in an over-simplified environment.

5.1.2 Task-space abstraction in the medial prefrontal cortex

Our findings, indicating the absence of consistent spatial information in the mPFC, also challenge claims of task-space abstraction within this region. Nevertheless, we acknowledge the necessity of further analysis to definitively challenge the notion of a task-space abstract code in the mPFC. Task-specific abstraction relevant to our paradigm has been proposed to be encoded in the mPFC through the representation of progress information between rewards (Ma et al., 2014) and sequence segment trajectory-specific encoding (Fujisawa et al., 2008; Ito et al., 2015). While we found evidence of progress encoding within intermediate rewards, the evidence was trajectory-identity encoding was seriously confounded by the motoric turn action encoding.

We demonstrated that distinct neuronal ensembles within the mPFC modulate during departure and approach phases of path progress between any two goals. This phase-specific coding reproduces the generalized spatial representations that have been reported numerous times in mPFC literature (J. Y. Yu et al., 2018; Shin et al., 2019; Rubin et al., 2019; Kaefer et al., 2020; El-Gaby et al., 2023; Guidera et al., 2024). Similar neuronal modulation was also observed in the Orbitofrontal Cortex (OFC) (Basu et al., 2021). In these studies different assumptions about which of the physical variables of distance or time are being used to encode the progress between the rewards were made. For example, El-Gaby et al. (2023) used normalized time to define progress between rewards and Guidera et al. (2024) used uniform spatial binning to define progress between rewards. Since these dimensions

typically covary in navigation of animals across the maze, it is difficult to decouple them. This ambiguity in the definition of progress between rewards makes it difficult to test the progress encoding hypothesis on navigational paradigms with more variable distances between rewards. However, our paradigm rendered high throughput of trials with variable speeds of progress between rewards, which allowed us to disentangle the two dimensions. Upon testing the single unit-level modulatory dynamics responsible for such progression encoding in time and distance domain, we found that cell modulation across the trials was consistent more in distance rather than time. Although it might be tempting to infer from this result that the mPFC potentially organizes its representations based on the anticipated location of the animal, rather than how much time it will take to get there (at least within the context of goal-directed spatial navigation), we refrain from making such a prediction. The presence of turn selectivity in these cells necessitates caution in drawing such an inference. A more conservative interpretation suggests that the mPFC may be reflecting spatially-consistent sensorimotor execution across stereotyped trajectories. Our analyses supports the latter explanation. It is likely that similar cells were potentially referred to as place encoding or position-correlated cells in the mPFC in previous literature (Kaefer et al., 2020; Hok et al., 2005).

To track the animal's progress in the task rather than just progress between intermediate reward zones, another piece of information on top of the path progress needed is the identity of the sequence segment. Evidence from several studies in the rodent mPFC literature suggests that the mPFC may encode the sequence segment identity (Fujisawa et al., 2008; Ito et al., 2015). Similar evidence exists for rodent OFC (Basu et al., 2021). Interestingly, we found that the mPFC neural variance in our task was explained both by the turn direction actions at the reward zones and the reward zone identity. The turn encoding is in line with the numerous findings of action encoding in mPFC (Cowen & McNaughton, 2007; Guidera et al., 2024). The interaction of the two variables and the constraints imposed by the insufficient samples to disentangle the two from analytical methods does not allow us

to confirm the actual variable being tracked by the mPFC among the two. It is possible that the mPFC may not encode the task-space in terms of the identity of the reward zone, but rather the motoric actions that the animal performs at these locations.

5.1.3 Action abstraction in the medial prefrontal cortex

The turn-selective cells we have identified exhibit multiple firing fields throughout the maze. We showed that this pattern is observed because the animal consistently makes the specific turn for which the cell is selective in these regions. Similar multi-peak firing characteristics for mPFC neurons, previously observed in constrained environments like the figure-eight maze (Ito et al., 2015), W-maze (Shin et al., 2019), plus maze (Kaefer et al., 2020) etc., have been attributed to the formation of generalized representations in the mPFC. Our findings suggest the necessity to address potential confounds arising from common turning directions in order to justify such an explanation for cellular responses with multiple firing fields within navigation paradigms. Further reevaluation is required to determine whether neuronal firing at these locations correlates with the animal's turning direction in subsequent trials. Accounting for this motor action encoding may necessitate a reevaluation of existing claims in the literature regarding mPFC function in goal-directed spatial navigation tasks.

Similar to our results, Guidera et al. (2024) reported that the mPFC ensembles represented the animal's turning actions more similarly regardless of the spatial location of the turn. They interpreted their findings as evidence of a hierarchical action code. However, our results suggest that the turn encoding in the mPFC may be more motoric in nature. We posit that the turn direction specific generalization and synchronization of neural dynamics reported by Guidera et al. (2024) and in our data may be a reflection of the animal's consistent turning behavior during departure from reward locations and minor preparatory movements for next turn during approach towards reward zones.

Our findings appears to align with results of Poucet (1997) which suggested that position-

correlated representations may reflect confounds arising from consistent performance of actions at specific locations in reward-guided navigation tasks. The start and next turn encoding we found in the mPFC may be a reflection of the animal's consistent turning behavior at reward locations. In support of this, we found that the apparent prospective turn encoding in the mPFC is confounded by the animal's preparatory turning movements to align itself for the upcoming turn action (next turn). In a similar vein, the start turn encoding we found in the mPFC may be a reflection of the animal's immediate turning behavior after receiving a reward, rather than an abstract action code. This is supported by the fact that the turn encoding was more prominent during the departure phase rather than the approach phase. Overall, this suggests that the mPFC may be encoding the animal's consistent turning behavior at reward locations instead of abstract actions.

If neurons are encoding a motoric action, we should see the same neuron specific to a particular turn direction during both approach phase and the departure phase of subsequent trial. However, our findings revealed only partial overlap between start-turn and next-turn encoding cells within the mPFC. We found a heterogeneous population of cells selective for the next turn but not the start turn, and vice versa. While limitations in our classification methodology might contribute to a portion of the non-overlapping population, the substantial size of this group suggests that the mPFC may perform additional computations on top of the motoric code.

Furthermore, the reward-departure and reward-approach specific cells we identified in the mPFC exhibited only partial overlap with turn-encoding cells. This finding suggests that not all position-correlated activity within the mPFC can be solely attributed to the animal's consistent turning behaviors at reward locations. Taken together, our data aligns with the view that position-correlated representations may arise from actions performed at specific spatial locations. However, we cannot definitively exclude the possibility of a more cognitive interpretation of the observed mPFC turn specificity, which would contrast the findings of Poucet (1997).

The position-correlated representations in the mPFC are theorized to be a part of a system that merges spatial context with task relevant information (Euston et al., 2012) like rules (Jung et al., 1998), motivational significance (Pratt & Mizumori, 2001), strategies (Rich & Shapiro, 2009), and routes (Ito et al., 2015). Specifically, our findings of the motoric turn encoding appear to be reminiscent of the trajectory-dependent selectivity in the mPFC reported by Ito et al. (2015). They found that the mPFC neurons were selective for the animal's trajectory in a continuous T-maze task. Their modified T-maze structure resembled the figure-eight maze used by Mashhoori et al. (2018), whereas the latter also incorporated elevated ramps at choice feeders. The presumed trajectory-dependent selectivity in mPFC may explain the successful decoding of animal position in the maze with similar geometry Mashhoori et al. (2018). We argue, however, that the limited maze geometry employed by Ito et al. (2015) might restrict the generalizability of their interpretations, as this theory falls short in explaining our inability to obtain the same results in a more complex maze. If this theory held true, we would expect clear differentiation between all sequence segments within the mPFC population activity. However, discriminability between segments was not apparent in mPFC code as shown by poor spatial and sequence progress decoding performance in our data. Although a limitation in sampling of trajectory-selective neurons could be argued, a more plausible explanation for this inconsistent segment discrimination may lie in the robust code of turning actions observed in the mPFC. We observed that when the animal was put in a position to make similar choice of left or right turn in different locations similar to the junction of the figure-eight maze, the mPFC neurons were selective for the animal's turn direction regardless of the identity of the reward zone location.

Reconciling our findings with those of Ito et al. (2015) presents a challenge. While their work reported trajectory-specific encoding in in mPFC cells, this specificity was apparently not reported beyond the first choice point, whether it was due to the authors' deliberate focus on the discriminative activity at the first choice point or due to the lack of trajectory-specific encoding beyond this point remains unclear. Their analysis did not explicitly control for

motoric turn actions—they focused on head direction and lateral position, which are not directly equivalent to the turn action encoding observed in our data. Consequently, there is a distinct possibility that a significant portion of the trajectory encoding reported by Ito et al. (2015) might be attributable to the animal's consistent turning behaviors at vertices of the maze. A neuron encoding a specific turn motor action would be expected to exhibit multiple firing fields across the task trajectory. For instance, a cell encoding left turns would ideally activate at every vertex within the trajectory where left turns are performed.

Assuming, for the sake of argument, that Ito et al. (2015) accurately captured exclusively trajectory encoding in mPFC neurons, it is possible that some non-turn-selective in their data might encode either action or trajectory information specific to the spatial context of the task. Disentangling these possibilities is difficult in their maze, as the animal's trajectory is inherently linked to its actions at each location. However, the lack of consistent trajectory decodability observed in our data suggests that the mPFC may not represent the animal's complete trajectory within the task. Alternatively, it is a plausible argument that our inconsistent decoding results might be due to insufficient sampling of neurons encoding individual trajectories in the mPFC, given the increased number of trajectories in our task compared to that of Ito et al. (2015).

It is also likely that the motoric turn action selective neurons we report are similar to the egocentric border cells reported by Long et al. (2024). However, egocentric border related activity seems an insufficient explanation for this cellular modulation as significant proportion of the cells selective for turn direction encoding have individualistic approach or departure related activity and our analysis of positions of spikes for these neurons suggest that they fire well away from the tabletop's border. However, one caution about its interpretation is that we could not capture neural activity close to the border as there were substantial artifacts of the brain stimulation occurring at reward zones on the border of the tabletop. Nevertheless, the turn direction specific encoding of these cells defined by us is defined on the basis of the activity that's independent of blackout period related activity. In

support of the argument against the egocentric border cell nature of these cells, population neural activity found similar among similar turns in Guidera et al. (2024) did not necessarily have a straight ahead border. Together, these suggest that the turn specific firing of the mPFC neurons found in our data are unlikely to be border cells.

Taken altogether, our findings underscore an often overlooked yet significant factor influencing mPFC neural representations. Our findings provide stronger support for a motoric interpretation of turn encoding in the mPFC. However, the possibility of a parallel associative code that integrates action and spatial context cannot be definitively excluded. Future investigations are necessary to disentangle the relative contributions of motoric and cognitive factors underlying turn encoding in mPFC.

5.1.4 Value associations in the medial prefrontal cortex

Our task design presents a challenge, as turning actions at reward zones coincide with reward delivery. Since mPFC is functionally implicated in assigning hedonic value to spatial states (Pratt & Mizumori, 2001), the reward approach neuronal modulation observed in our data might reflect reward expectancy in cells selective for the approach phase but not the next turn.

Supporting this possibility, Ma et al. (2014) investigated rat mPFC activity using a sequence task requiring generalization across actions (lever presses). Notably, their reward delivery port was positioned opposite the levers in the chamber. They found that mPFC representations exhibited a mutually orthogonal organization across first two principal components. The first component displayed a continuous ramp across all lever presses, followed by a slow decline towards baseline. The second component, however, showed a ramp during approach to the reward port. These findings suggest that mPFC neuronal populations might independently encode reward anticipation and behavioral policy. Our results, showing reward approach phase cells, could potentially reflect a mixture of both these signals, although we eschew calling it task-progress encoding cells as proposed by Ma et al. (2014)

since we found no evidence of encoding of progress in the entire spatial sequence.

Future experiments, potentially displacing reward location relative to turns, are necessary to definitively determine if encoding for both reward and action emerges in the same or distinct cells. While our analysis revealed an overlap between turn-encoding and reward-anticipatory cells, we also identified cells selective for the next turn but not reward anticipation. Displacing reward locations might provide a more conclusive test of this dissociation.

Interestingly, Ma et al. (2014) reported that mPFC cells were not selective for individual actions but rather the entire action sequence, whereas in our findings stronger support exists for individual motoric action coding. It is unclear how to reconcile both of these findings, but a key difference between both tasks seems to be the reward. In our sequence navigation task, the reward was delivered at the reward zones, whereas in the lever press task, the reward was delivered after the whole action sequence was executed correctly by the animal. Our findings, hence, suggest that the organization of behavior in the rat mPFC may be more short-sighted than previously hypothesized where progress between rewards, as opposed to progress between complete task, is tracked. It is also possible that a sub-population specific code, as proposed by El-Gaby et al. (2023) is being used to track the progress of animal across the entire task sequence, which we could not capture due to the limitations in our analyses. Future investigations are necessary to disentangle these possibilities.

While Ma et al. (2014)'s work provides valuable insights, it may not fully explain all our observations. We identified a substantial proportion of departure phase-selective neurons that exhibited no selectivity for the start turn. It could be argued that these cells might be a consequence of a brain-wide or local excitation induced by the reward delivery rather than tracking progress from the reward. However, the fact that we found a significant number of approach phase-selective cells which cannot be simply attributed to reward delivery induced excitation, suggests that the departure phase cells may be part of a larger population of cells tracking progress between rewards. Furthermore, if the departure cells were merely a consequence of excitation due to reward delivery, we would expect the dynamics of the

departure cells to be more consistent across trials in temporal domain. Our results of representational consistency analysis contradicted this hypothesis, thereby supporting the notion that these cells are tracking progress and utilizing a distance-based metric to achieve it.

Interestingly, the distance-based progress tracking we observed does not occur in previously hypothesized manner. The individual mPFC cells neither tesselate the inter-reward progress space, as posited by El-Gaby et al. (2023), nor do they have linearly increasing/decreasing firing across an entire trajectory between rewards. Instead, we find the modulation of cells seem to be localized to the reward locations. Thus, it suggests that reward-related modulation near goal departure and approach is due expected hedonic value encoding. However, a caution to this interpretation lies in our task design which has consistent reward delivery at each reward zone. This task structure makes it difficult to precisely decouple the contribution of expected value encoding from progress encoding. Future experiments producing precise variations in reward delivery locations and introduction of intermediate rewards may necessary to definitively determine the role of value associations in the mPFC during spatial navigation.

Moreover, while we found evidence of progress encoding in the mPFC, the fact that we saw a significant number of cells attuned to the motoric movements associated to turning actions also casts a doubt towards existence of a progress-specific code in the mPFC. The stereotypical nature of animals' speed variations while traveling between two intermediate rewards involves them starting from a slow speed at the departure phase and speeding up until approach to reward zone, where they tend to slow down. It is possible that the motoric control required by the animal to perform these speeding up and slowing down actions might be leading to such artifacts we misinterpret as a inter-reward progress tracking neural code. Although we tried to disentangle the obfuscation of turn action encoding and progress encoding, and we did find that not all progress-encoding cells were modulated by turn, the motoric code confounding may still be a valid hypothesis addressing which is beyond the scope of our analyses. In support of this line of reasoning, we point to a previous study in the

Prelimbic Cortex (PL) and Infralimbic Cortex (IL) cortices that reported goal-associated spatial tuning, potentially corresponding to spatially consistent turn encoding, even when rewards triggered after reaching a goal zone were dispensed remotely at random locations after reaching the goal (Hok et al., 2005).

Reconciling our findings with those of El-Gaby et al. (2023) presents a significant challenge. The task structure utilized by El-Gaby et al. (2023) includes numerous locations between intermediate rewards where the animal must perform turning actions. However, the authors did not account for potential motoric confounds in their task design and analyses. It is plausible that the tessellation of the inter-reward space observed in their data is attributable to the animal's consistent turning behaviors at these locations.

In conclusion, our findings do not rule out the possibility of spatially associated value associations in the mPFC. In parallel, we also acknowledge that the repetitive nature of the task creates consistent actions (turns and changes in speed) at specific locations. Neural activity corresponding to these action variables is hence consistently tied to the specific locations, which is misinterpreted as spatial or progress coding. Future research separating reward location from actions is needed to fully understand how mPFC integrates reward processing with movement coordination.

5.2 Methodological limitations and suggested best practices

Our investigation identified limitations in extending the methodology used by previous studies. One key concern lies in the application of the decoding method introduced by Mashhoori et al. (2018), which relies heavily on Gaussian smoothing with a standard deviation of 150 ms. This symmetric smoothing kernel applied across three standard deviations on either side results in a temporally smoothed firing rate window spanning up to 900 ms. While an empirical determination of better smoothing window size through metrics such as Akaike Information Criterion (AIC) or Bayesian Information Criterion (BIC) is feasible, our choice of a 150 ms was primarily driven by the need for methodological consistency

with the approach employed by Mashhoori et al. (2018).

Critically, we advise that this window size must be compared to the animal's actual traversal time within the trials. In our experiment, animals took an average of $\sim$2.1 s to traverse between rewards, with reward stimulation epoch excluded due to artifacts from the stimulation. Smoothing with such a large kernel risks excessive temporal spreading of firing rates and the creation of artificial firing rate ramps, potentially exploited by the decoding model to exaggerate the information of animal position in the neural data.

Conversely, Mashhoori et al. (2018) employed a smoothing window spanning the entire trial length (on average 13.6 s), which included reward consumption and post-reward pauses. While this smoothing kernel may have been suitable for their task, it is demonstrably unfit for our paradigm. Reducing the window size in our study resulted in a significant decrease in decoding accuracy, rendering the decoding approach ineffective. This highlights the inherent trade-off associated with smoothing: while necessary in some cases to achieve interpretable neural activity patterns, it can also introduce artifacts (in the form of ramping activity due to smoothing kernel) that compromise interpretation. Future analyses should hence consider the temporal scales of the task and the modulatory dynamics of neural activity under investigation to optimize smoothing parameters. As an alternative to smoothing pre-processing, the decoding model could utilize temporal covariates—historical and/or future neural activity bins (Glaser et al., 2020). This approach, however, may increase the dimensionality of the input data, necessitating appropriate scaling to avoid overfitting, and also require sophisticated techniques for model interpretation.

Furthermore, best practices for designing experiments that investigate task abstractions in the mPFC require addressing the entanglement of sensorimotor information with position-correlated representations. A synergistic combination of intelligent maze design and analytical methods can be employed to achieve this goal. For example, counterbalanced designs that ensure equal representation of potentially confounding actions at each trajectory can help to dissociate the influence of motoric confounds (Boyle et al., 2024).

Our data was not designed to address this issue, and thus sufficient sampling could not be ensured to decouple the motoric confounds from the abstract turn encoding. Another way to address this issue is to use highly reduced tasks that minimize the number of actions that the animal can perform at each location, such as linear track used by Basu et al. (2021) for sequence tasks. Recent advancements in deep reinforcement learning algorithms combined with precise behavioral tracking has surpassed conventionally computed behavioral variables in predicting neural activity within the dorsolateral striatum and motor cortex (Aldarondo et al., 2024). It might be necessary in future investigations to generate bio-mechanically realistic models of animal behavior to definitively isolate the cognitive contributions from motor influences on observed neural activity.

5.3 Conclusion

Overall, this study explored the rodent mPFC's role in spatial navigation using a reward-guided sequence task requiring generalization across different physical location configurations of the reward zones. We observed prominent encoding of progress between rewards and motor-related turn actions in mPFC neural code. However, limitations in our task structure prevent definitive conclusions about spatial codes for actions, task space, or reward value. Dissociating reward location from turns in future studies is crucial to understand how mPFC integrates reward processing with movement planning. Furthermore, our findings emphasize the need to consider motor confounds when interpreting position-related activity in mPFC. By addressing these considerations, future research can effectively disentangle the mPFC's cognitive and sensorimotor roles in spatial navigation tasks.

References

Aldarondo, D., Merel, J., Marshall, J. D., Hasenclever, L., Klibaite, U., Gellis, A., Tassa, Y., Wayne, G., Botvinick, M., & Ölveczky, B. P. (2024). A virtual rodent predicts the structure of neural activity across behaviors. *Nature 2024*, 1–3. https://doi.org/10.1038/s41586-024-07633-4

Baeg, E. H., Kim, Y. B., Huh, K., Mook-Jung, I., Kim, H. T., & Jung, M. W. (2003). Dynamics of population code for working memory in the prefrontal cortex. *Neuron*, *40*, 177–188. https://doi.org/10.1016/S0896-6273(03)00597-X

Basu, R., Gebauer, R., Herfurth, T., Kolb, S., Golipour, Z., Tchumatchenko, T., & Ito, H. T. (2021). The orbitofrontal cortex maps future navigational goals. *Nature*, *599*, 449–452. https://doi.org/10.1038/s41586-021-04042-9

Batuev, A. S., Kursina, N. P., & Shutov, A. P. (1990). Unit activity of the medial wall of the frontal cortex during delayed performance in rats. *Behavioural Brain Research*, *41*, 95–102. https://doi.org/10.1016/0166-4328(90)90145-5

Bechara, A., & Damasio, A. R. (2005). The somatic marker hypothesis: A neural theory of economic decision. *Games and Economic Behavior*, *52*, 336–372. https://doi.org/10.1016/J.GEB.2004.06.010

Bergstrom, J. A. (1893). Experiments upon physiological memory by means of the interference of associations. *The American Journal of Psychology*, *5*, 356. https://doi.org/10.2307/1410999

Birrell, J. M., & Brown, V. J. (2000). Medial frontal cortex mediates perceptual attentional set shifting in the rat. *Journal of Neuroscience*, *20*, 4320–4324. https://doi.org/10.1523/JNEUROSCI.20-11-04320.2000

Boser, B. E., Guyon, I. M., & Vapnik, V. N. (1992). Training algorithm for optimal margin classifiers. *Proceedings of the Fifth Annual ACM Workshop on Computational Learning Theory*, 144–152. https://doi.org/10.1145/130385.130401

Boyle, L. M., Posani, L., Irfan, S., Siegelbaum, S. A., & Fusi, S. (2024). Tuned geometries of hippocampal representations meet the computational demands of social memory. *Neuron*, *112*, 1358–1371.e9. https://doi.org/10.1016/J.NEURON.2024.01.021

Burnham, W. H. (1903). Retroactive amnesia: Illustrative cases and a tentative explanation. *The American Journal of Psychology*, *14*, 118. https://doi.org/10.2307/1412310

Buzsáki, G. (2015). Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus*, *25*, 1073–1188. https://doi.org/10.1002/HIPO.22488

Chalmers, E., Luczak, A., & Gruber, A. J. (2016). Computational properties of the hippocampus increase the efficiency of goal-directed foraging through hierarchical reinforcement learning. *Frontiers in Computational Neuroscience*, *10*. https://doi.org/10.3389/fncom.2016.00128

Chang, H. R., Esteves, I. M., Neumann, A. R., Mohajerani, M. H., & McNaughton, B. L. (2023). Cortical reactivation of spatial and non-spatial features coordinates with hippocampus to form a memory dialogue. *Nature Communications 2023 14:1*, *14*, 1–17. https://doi.org/10.1038/s41467-023-43254-7

Cowansage, K. K., Shuman, T., Dillingham, B. C., Chang, A., Golshani, P., & Mayford, M. (2014). Direct reactivation of a coherent neocortical memory of context. *Neuron*, *84*, 432–441. https://doi.org/10.1016/j.neuron.2014.09.022

Cowen, S. L., & McNaughton, B. L. (2007). Selective delay activity in the medial prefrontal cortex of the rat: Contribution of sensorimotor information and contingency [PMID: 17507507]. *Journal of Neurophysiology*, *98*, 303–316. https://doi.org/10.1152/jn.00150.2007

Draguhn, A. (2022). No single place for space: Neuronal representation of location beyond the hippocampus. *Pflügers Archiv - European Journal of Physiology*, *474*, 569–571. https://doi.org/10.1007/s00424-022-02699-3

Duan, C. A., Erlich, J. C., & Brody, C. D. (2015). Requirement of prefrontal and midbrain regions for rapid executive control of behavior in the rat. *Neuron*, *86*, 1491–1503. https://doi.org/10.1016/j.neuron.2015.05.042

Durstewitz, D., Vittoz, N. M., Floresco, S. B., & Seamans, J. K. (2010). Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. *Neuron*, *66*, 438–448. https://doi.org/https://doi.org/10.1016/j.neuron.2010.03.029

Edelman, G. M., Mountcastle, V. B., & Program., N. R. (1982). The mindful brain : Cortical organization and the group-selective theory of higher brain function, 100.

El-Gaby, M., Harris, A. L., Whittington, J. C. R., Dorrell, W., Bhomick, A., Walton, M. E., Akam, T., & Behrens, T. E. J. (2023). A cellular basis for mapping behavioural structure. *bioRxiv*, 2023.11.04.565609. https://doi.org/10.1101/2023.11.04.565609

Esteves, I. M., Chang, H., Neumann, A. R., Sun, J., Mohajerani, M. H., & McNaughton, B. L. (2021). Spatial information encoding across multiple neocortical regions depends on an intact hippocampus. *Journal of Neuroscience*, *41*, 307–319. https://doi.org/10.1523/JNEUROSCI.1788-20.2020

Euston, D. R., Gruber, A. J., & McNaughton, B. L. (2012). The role of medial prefrontal cortex in memory and decision making. *Neuron*, *76*, 1057–1070. https://doi.org/10.1016/J.NEURON.2012.12.002

Euston, D. R., & McNaughton, B. L. (2006). Apparent encoding of sequential context in rat medial prefrontal cortex is accounted for by behavioral variability. *Journal of Neuroscience*, *26*, 13143–13155. https://doi.org/10.1523/JNEUROSCI.3803-06.2006

Euston, D. R., Tatsuno, M., & McNaughton, B. L. (2007). Fast-forward playback of recent memory sequences in prefrontal cortex during sleep. *Science*, *318*, 1147–1150. https://doi.org/10.1126/science.1148979

Ferino, F., Thierry, A. M., & Glowinski, J. (1987). Anatomical and electrophysiological evidence for a direct projection from ammon's horn to the medial prefrontal cortex in the rat. *Experimental Brain Research*, *65*, 421–426. https://doi.org/10.1007/BF00236315

Fiser, A., Mahringer, D., Oyibo, H. K., Petersen, A. V., Leinweber, M., & Keller, G. B. (2016). Experience-dependent spatial expectations in mouse visual cortex. *Nature Neuroscience 2016 19:12*, *19*, 1658–1664. https://doi.org/10.1038/nn.4385

Fujisawa, S., Amarasingham, A., Harrison, M. T., & Buzsáki, G. (2008). Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nature Neuroscience 2008 11:7*, *11*, 823–833. https://doi.org/10.1038/nn.2134

Fuster, J. M., & Alexander, G. E. (1971). Neuron activity related to short-term memory. *Science*, *173*, 652–654. https://doi.org/10.1126/SCIENCE.173.3997.652

Glaser, J. I., Benjamin, A. S., Chowdhury, R. H., Perich, M. G., Miller, L. E., & Kording, K. P. (2020). Machine learning for neural decoding. *eNeuro*, *7*, 1–16. https://doi.org/10.1523/ENEURO.0506-19.2020

Greenberg, R., & Underwood, B. J. (1950). Retention as a function of stage of practice. *Journal of Experimental Psychology*, *40*, 452–457. https://doi.org/10.1037/H0062147

Guidera, J. A., Gramling, D. P., Comrie, A. E., Joshi, A., Denovellis, E. L., Lee, K. H., Zhou, J., Thompson, P., Hernandez, J., Yorita, A., Haque, R., Kirst, C., Frank, L. M., & Francisco, S. (2024). Regional specialization manifests in the reliability of neural population codes. *bioRxiv*, 2024.01.25.576941. https://doi.org/10.1101/2024.01.25.576941

Guise, K. G., & Shapiro, M. L. (2017). Medial prefrontal cortex reduces memory interference by modifying hippocampal encoding. *Neuron*, *94*, 183–192.e8. https://doi.org/10.1016/J.NEURON.2017.03.011

Harvey, C. D., Coen, P., & Tank, D. W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature 2012 484:7392*, *484*, 62–68. https://doi.org/10.1038/nature10918

Hinton, G. E. (1989). *Learning distributed representations of concepts.* Clarendon Press/Oxford University Press.

Hoffman, K. L., & McNaughton, B. L. (2002). Coordinated reactivation of distributed memory traces in primate neocortex. *Science*, *297*, 2070–2073. https://doi.org/10.1126/science.1073538

Hok, V., Save, E., Lenck-Santini, P. P., & Poucet, B. (2005). Coding for spatial goals in the prelimbic/infralimbic area of the rat frontal cortex. *Proceedings of the National Academy of Sciences*, *102*, 4602–4607. https://doi.org/10.1073/pnas.0407332102

Hoover, W. B., & Vertes, R. P. (2007). Anatomical analysis of afferent projections to the medial prefrontal cortex in the rat. *Brain Structure and Function*, *212*, 149–179. https://doi.org/10.1007/s00429-007-0150-4

Hyman, J. M., Ma, L., Balaguer-Ballester, E., Durstewitz, D., & Seamans, J. K. (2012). Contextual encoding by ensembles of medial prefrontal cortex neurons. *Proceedings of the National Academy of Sciences*, *109*, 5086–5091. https://doi.org/10.1073/pnas.1114415109

Ito, H. T., Zhang, S.-J., Witter, M. P., Moser, E. I., & Moser, M.-B. (2015). A prefrontal-thalamo-hippocampal circuit for goal-directed spatial navigation. *Nature*, *522*, 50–55. https://doi.org/10.1038/nature14396

Jadhav, S. P., Kemere, C., German, P. W., & Frank, L. M. (2012). Awake hippocampal sharp-wave ripples support spatial memory. *Science*, *336*, 1454–1458. https://doi.org/10.1126/science.1217230

Ji, D., & Wilson, M. A. (2006). Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nature Neuroscience 2007 10:1*, *10*, 100–107. https://doi.org/10.1038/nn1825

Jung, M. W., Qin, Y., McNaughton, B. L., & 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. https://doi.org/10.1093/CERCOR/8.5.437

Kaas, J. H. (1987). The organization of neocortex in mammals: Implications for theories of brain function. *Annual Review of Psychology*, *38*, 129–151. https://doi.org/10.1146/annurev.ps.38.020187.001021

Kaefer, K., Nardin, M., Blahna, K., & Csicsvari, J. (2020). Replay of behavioral sequences in the medial prefrontal cortex during rule switching. *Neuron*, *106*, 154–165.e6. https://doi.org/10.1016/J.NEURON.2020.01.015

Khodagholy, D., Gelinas, J. N., & Buzsáki, G. (2017). Learning-enhanced coupling between ripple oscillations in association cortices and hippocampus. *Science*, *358*, 369–372. https://doi.org/10.1126/science.aan6203

Kihlberg, J. K., Herson, J. H., & Schotz, W. E. (1972). Square root transformation revisited. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, *21*, 76–81. https://doi.org/10.2307/2346609

Kingma, D. P., & Ba, J. L. (2014). Adam: A method for stochastic optimization. *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*. https://arxiv.org/abs/1412.6980v9

Kolb, B. (1984). Functions of the frontal cortex of the rat: A comparative review. *Brain Research Reviews*, *8*, 65–98. https://doi.org/10.1016/0165-0173(84)90018-3

Kolb, B., & Whishaw, I. Q. (1990). *Fundamentals of human neuropsychology, 3rd ed.* W H Freeman/Times Books/ Henry Holt; Co.

Kursa, M. B., & Rudnicki, W. R. (2010). Feature selection with the boruta package. *Journal of Statistical Software*, *36*, 1–13. https://doi.org/10.18637/JSS.V036.I11

Kurth-Nelson, Z., Behrens, T., Wayne, G., Miller, K., Luettgau, L., Dolan, R., Liu, Y., & Schwartenbeck, P. (2023). Replay and compositional computation. *Neuron*, *111*, 454–469. https://doi.org/10.1016/j.neuron.2022.12.028

Lapish, C. C., Durstewitz, D., Chandler, L. J., & Seamans, J. K. (2008). Successful choice behavior is associated with distinct and coherent network states in anterior cingulate cortex. *Proceedings of the National Academy of Sciences*, *105*, 11963–11968. https://doi.org/10.1073/pnas.0804045105

Larsen, J. K., & Divac, I. (1978). Selective ablations within the prefrontal cortex of the rat and performance of delayed alternation. *Physiological Psychology*, *6*, 15–17. https://doi.org/10.3758/BF03326684

Laubach, M., Caetano, M. S., & Narayanan, N. S. (2015). Mistakes were made: Neural mechanisms for the adaptive control of action initiation by the medial prefrontal cortex. *Journal of Physiology-Paris*, *109*, 104–117. https://doi.org/10.1016/J.JPHYSPARIS.2014.12.001

LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, *86*, 2278–2323. https://doi.org/10.1109/5.726791

Long, X., Deng, B., Shen, R., Yang, L., Chen, L., Ran, Q., Du, X., & Zhang, S.-J. (2024). Border cells without theta rhythmicity in the medial prefrontal cortex. *Proceedings of the National Academy of Sciences*, *121*, e2321614121. https://doi.org/10.1073/pnas.2321614121

Long, X., & Zhang, S. J. (2021). A novel somatosensory spatial navigation system outside the hippocampal formation. *Cell Research 2021 31:6*, *31*, 649–663. https://doi.org/10.1038/s41422-020-00448-8

Loshchilov, I., & Hutter, F. (2017). Decoupled weight decay regularization. *7th International Conference on Learning Representations, ICLR 2019*. https://arxiv.org/abs/1711.05101v3

Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, *2017-December*, 4766–4775. https://arxiv.org/abs/1705.07874v2

Ma, L., Hyman, J. M., Lindsay, A. J., Phillips, A. G., & Seamans, J. K. (2014). Differences in the emergent coding properties of cortical and striatal ensembles. *Nature Neuroscience 2014 17:8*, *17*, 1100–1106. https://doi.org/10.1038/nn.3753

Mao, D., Kandler, S., McNaughton, B. L., & Bonin, V. (2017). Sparse orthogonal population representation of spatial context in the retrosplenial cortex. *Nature Communications 2017 8:1*, *8*, 1–9. https://doi.org/10.1038/s41467-017-00180-9

Mashhoori, A., Hashemnia, S., McNaughton, B. L., Euston, D. R., & Gruber, A. J. (2018). Rat anterior cingulate cortex recalls features of remote reward locations after disfavoured reinforcements. *eLife*, *7*. https://doi.org/10.7554/ELIFE.29793

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*, 419–457. https://doi.org/10.1037/0033-295X.102.3.419

McLaughlin, A. E., Diehl, G. W., & Redish, A. D. (2021). Potential roles of the rodent medial prefrontal cortex in conflict resolution between multiple decision-making systems. *International Review of Neurobiology*, *158*, 249–281. https://doi.org/10.1016/BS.IRN.2020.11.009

McNaughton, B. L., O'Keefe, J., & Barnes, C. A. (1983). The stereotrode: A new technique for simultaneous isolation of several single units in the central nervous system from multiple unit records. *Journal of Neuroscience Methods*, *8*, 391–397. https://doi.org/10.1016/0165-0270(83)90097-3

Messanvi, K. F., Berkun, K., Perkins, A., & Chudasama, Y. (2023). Parallel pathways provide hippocampal spatial information to prefrontal cortex. *Journal of Neuroscience*, *43*, 68–81. https://doi.org/10.1523/JNEUROSCI.0846-22.2022

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, *24*, 167–202. https://doi.org/10.1146/annurev.neuro.24.1.167

Mountcastle, V. B. (1957). Modality and topographic properties of single neurons of cat's somatic sensory cortex. *https://doi.org/10.1152/jn.1957.20.4.408*, *20*, 408–434. https://doi.org/10.1152/JN.1957.20.4.408

Müller, G. E., & Pilzecker, A. (1900). *Experimentelle beiträge zur lehre vom gedächtniss* (Vol. 1). Z Psychol Ergänzungsband.

Orlov, A. A., Kurzina, N. P., & Shutov, A. P. (1988). Activity of medial wall neurons in frontal cortex of rat brain during delayed response reactions. *Neuroscience and Behavioral Physiology*, *18*, 31–37. https://doi.org/10.1007/BF01186902

Orlov, A. A., Shutov, A. P., Kurzina, N. P., & Selezneva, E. V. (1991). Reflection of the prognostic activity of the brain in the activity of the neurons of the frontal cortex of rats. *Neuroscience and Behavioral Physiology*, *21*, 39–41. https://doi.org/10.1007/BF01184237

Orona, E., & Gabriel, M. (1983). Multiple-unit activity of the prefrontal cortex and mediodorsal thalamic nucleus during reversal learning of discriminative avoidance behavior in rabbits. *Brain Research*, *263*, 313–329. https://doi.org/10.1016/0006-8993(83)90323-2

Pascanu, R., Mikolov, T., & Bengio, Y. (2012). On the difficulty of training recurrent neural networks. *30th International Conference on Machine Learning, ICML 2013*, 2347–2355. https://arxiv.org/abs/1211.5063v2

Pavlides, C., & Winson, J. (1989). Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. *Journal of Neuroscience*, *9*, 2907–2918. https://doi.org/10.1523/JNEUROSCI.09-08-02907.1989

Paxinos, G., & Watson, C. (2007). *The rat brain in stereotaxic coordinates* (6th edn.). Elsevier Press.

Pfeiffer, B. E., & Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature 2013 497:7447*, *497*, 74–79. https://doi.org/10.1038/nature12112

Poucet, B. (1997). Searching for spatial unit firing in the prelimbic area of the rat medial prefrontal cortex. *Behavioural Brain Research*, *84*, 151–159. https://doi.org/10.1016/S0166-4328(96)00144-1

Prasad, J. A., & Chudasama, Y. (2013). Viral tracing identifies parallel disynaptic pathways to the hippocampus. *Journal of Neuroscience*, *33*, 8494–8503. https://doi.org/10.1523/JNEUROSCI.5072-12.2013

Pratt, W. E., & Mizumori, S. J. (2001). 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. https://doi.org/10.1016/S0166-4328(01)00204-2

Ribot, T. (1881). Les maladies de la mémoire, 1–169.

Rich, E. L., & Shapiro, M. (2009). Rat prefrontal cortical neurons selectively code strategy switches. *Journal of Neuroscience*, *29*, 7208–7219. https://doi.org/10.1523/JNEUROSCI.6068-08.2009

Rigotti, M., Barak, O., Warden, M. R., Wang, X. J., Daw, N. D., Miller, E. K., & Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature 2013 497:7451*, *497*, 585–590. https://doi.org/10.1038/nature12160

Rubin, A., Sheintuch, L., Brande-Eilat, N., Pinchasof, O., Rechavi, Y., Geva, N., & Ziv, Y. (2019). Revealing neural correlates of behavior without behavioral measurements. *Nature Communications 2019 10:1*, *10*, 1–14. https://doi.org/10.1038/s41467-019-12724-2

Rumelhart, D. E. (1990). Brain style computation: Learning and generalization. Academic Press Professional, Inc.

Sakurai, Y., & Sugimoto, S. (1985). Effects of lesions of prefrontal cortex and dorsomedial thalamus on delayed go/no-go alternation in rats. *Behavioural Brain Research*, *17*, 213–219. https://doi.org/10.1016/0166-4328(85)90045-2

Sakurai, Y., & 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*, 295–301. https://doi.org/10.1016/0166-4328(86)90229-9

Saleem, A. B., Diamanti, E. M., Fournier, J., Harris, K. D., & Carandini, M. (2018). Coherent encoding of subjective spatial position in visual cortex and hippocampus. *Nature 2018 562:7725*, *562*, 124–127. https://doi.org/10.1038/s41586-018-0516-1

Sauer, J.-F., Folschweiller, S., & Bartos, M. (2022). Topographically organized representation of space and context in the medial prefrontal cortex. *Proceedings of the National Academy of Sciences*, *119*, e2117300119. https://doi.org/10.1073/pnas.2117300119

Shin, J. D., & Jadhav, S. P. (2024). Prefrontal cortical ripples mediate top-down suppression of hippocampal reactivation during sleep memory consolidation. *Current Biology*, *34*, 2801–2811.e9. https://doi.org/10.1016/j.cub.2024.05.018

Shin, J. D., Tang, W., & Jadhav, S. P. (2019). Dynamics of awake hippocampal-prefrontal replay for spatial learning and memory-guided decision making. *Neuron*, *104*, 1110–1125.e7. https://doi.org/10.1016/J.NEURON.2019.09.012

Singer, A. C., Carr, M. F., Karlsson, M. P., & Frank, L. M. (2013). Hippocampal swr activity predicts correct decisions during the initial learning of an alternation task. *Neuron*, *77*, 1163–1173. https://doi.org/10.1016/J.NEURON.2013.01.027

Skaggs, W. E., Mcnaughton, B. L., Gothard, K. M., & Markus, E. J. (1992). An information-theoretic approach to deciphering the hippocampal code. *Advances in Neural Information Processing Systems*, *5*.

Smith, L. N. (2017). Cyclical learning rates for training neural networks. *Proceedings - 2017 IEEE Winter Conference on Applications of Computer Vision, WACV 2017*, 464–472. https://doi.org/10.1109/WACV.2017.58

Sun, W., Winnubst, J., Natrajan, M., Lai, C., Kajikawa, K., Michaelos, M., Gattoni, R., Stringer, C., Flickinger, D., Fitzgerald, J. E., & Spruston, N. (2023). Learning produces a hippocampal cognitive map in the form of an orthogonalized state machine. *bioRxiv*, 2023.08.03.551900. https://doi.org/10.1101/2023.08.03.551900

Swanson, L. W. (1981). A direct projection from ammon's horn to prefrontal cortex in the rat. *Brain Research*, *217*, 150–154. https://doi.org/10.1016/0006-8993(81)90192-X

Tang, W., Shin, J. D., & Jadhav, S. P. (2023). Geometric transformation of cognitive maps for generalization across hippocampal-prefrontal circuits. *Cell Reports*, *42*, 112246. https://doi.org/https://doi.org/10.1016/j.celrep.2023.112246

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, *55*, 189–208. https://doi.org/10.1037/H0061626

Underwood, B. J. (1957). Interference and forgetting. *Psychological Review*, *64*, 49–60. https://doi.org/10.1037/H0044616

Wikenheiser, A. M., Gardner, M. P., Mueller, L. E., & Schoenbaum, G. (2021). Spatial representations in rat orbitofrontal cortex. *Journal of Neuroscience*, *41*, 6933–6945. https://doi.org/10.1523/JNEUROSCI.0830-21.2021

Wilson, M. A., & McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, *265*, 676–679. https://doi.org/10.1126/SCIENCE.8036517

Witter, M. P., & Groenewegen, H. J. (1984). Laminar origin and septotemporal distribution of entorhinal and perirhinal projections to the hippocampus in the cat. *Journal of Comparative Neurology*, *224*, 371–385. https://doi.org/10.1002/CNE.902240305

Xu, H., Baracskay, P., O'Neill, J., & Csicsvari, J. (2019). Assembly responses of hippocampal ca1 place cells predict learned behavior in goal-directed spatial tasks on the radial eight-arm maze. *Neuron*, *101*, 119–132.e4. https://doi.org/10.1016/J.NEURON.2018.11.015

Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. V., & Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Journal of Neurophysiology*, *102*, 614–635. https://doi.org/10.1152/JN.90941.2008

Yu, J. Y., Liu, D. F., Loback, A., Grossrubatscher, I., & Frank, L. M. (2018). Specific hippocampal representations are linked to generalized cortical representations in memory. *Nature Communications 2018 9:1*, *9*, 1–11. https://doi.org/10.1038/s41467-018-04498-w

Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, *67*, 301–320. https://doi.org/10.1111/J.1467-9868.2005.00503.X