## Title

A region-specific code for generalized representations across hippocampus and prefrontal cortex



## Author

Guidera, Jennifer Ann

## Publication Date

2023

Peer reviewed|Thesis/dissertation

A region-specific code for generalized representations across hippocampus and prefrontal cortex

by

Jennifer Guidera

DISSERTATION

Submitted in partial satisfaction of the requirements for degree of

DOCTOR OF PHILOSOPHY

in

Bioengineering

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, SAN FRANCISCO
AND
UNIVERSITY OF CALIFORNIA, BERKELEY

Approved:

DocuSigned by:
Loren Frank
9C247AE9D5474C3...

Loren Frank

Chair

DocuSigned by:
Karunesh Ganguly
[illegible]

Karunesh Ganguly

DocuSigned by:
[illegible]

Christoph Kirst

DocuSigned by:
Joni Wallis
0B8298CBBD20423...

Joni Wallis

Committee Members

## Dedication

To my parents, Ann Chien Guidera and Steven Allan Guidera

## Acknowledgements

The work presented in this thesis was supported by NIH award F30MH126483, a UCSF Discovery Fellowship, NIH grant T32GM007618, the Howard Hughes Medical Institute, and the Simons Collaboration for the Global Brain.

First and foremost, I am grateful to my PhD advisor, Loren Frank, for his mentorship and guidance throughout my time in the lab. Loren is an extraordinary mentor and human being. The first time I met with Loren, I remember being struck by how talking about scientific ideas felt, above all, fun. This feeling has continued throughout my time in the lab, and I've gotten a chance to see and understand more of Loren's values that influence how he thinks about science and life, and which so strongly enter into the experiences of those with whom he interacts. These values include a belief in sharing; in open-mindedness; in rigor and careful evaluation of ideas; and in reevaluation of one's mental models as new evidence and perspectives arise.

When it comes to how Loren tends to think about ideas, I have also been struck by the way in which Loren "thinks big". He wants to know what the data mean at a very high level, and this is fundamentally what he is deeply curious about. What conceptual models does a very specific finding support? Which models does it suggest against? As a trainee of Loren's, I have come to experience how thinking about what things mean at a very high level is one of the most exciting and fun experiences science can provide, in the potential for evolving and dismantling broad structures already in one's head.

Another one of Loren's special qualities is how he views his trainees as having inherent worth, independent of their scientific activities. Loren sees his trainees as

scientists, but at the end of the day also as people with unique desires and dreams. Loren often says that life is short: you should do exactly what it is that you want to do.

Loren is incredibly generous with his time and devoted to helping his trainees grow in all the ways that he can. I have been so fortunate to learn and grow from Loren's mentorship these last few years. As I go forward in science and life, Loren's generosity and values will remain forefront in my mind and inspire me daily to live out similar values.

I am grateful to Christoph Kirst for his teaching and inspiration in quantitative areas. Particularly in the middle portion of my training, Christoph introduced me to quantitative ideas that helped expand the ways I think about data. It has been a fun privilege to see through Christoph how a deep quantitative expertise paired with creativity can afford special insights.

I am grateful to Joni Wallis and Karunesh Ganguly for their valuable insights throughout the development of this project while serving on my qualifying exam and thesis committees. I feel fortunate to have been able to receive their feedback on this work, and to grow from getting a glimpse into their individual perspectives on science.

During my final two years in the lab, I was fortunate to work with a talented technician, Daniel Gramling. Daniel's insights, attention to detail, and steady and strong efforts have been critical to this project.

I am grateful to past members of the Frank lab for teaching and advice at early stages of my training, including Clay Smyth, Dani Astudillo Maya, Tom Davidson, Jason Chung, Hannah Joo, Anna Gillespie, and Mari Sosa. I am grateful to Viktor Kharazia,

Anya Kiseleva, and Michael Coulter for assistance with perfusions and histology. I am grateful to Abhilasha Joshi for taking me under her wing when I joined the lab and teaching me several aspects of electrophysiology; Alison Comrie for her mentorship on probe and hybrid drive implantation and her development of hybrid implant approaches; Eric Denovellis for support in using his decoding algorithm; and Abhijith Mankilli, Emily Monroe, and Rhino Nevers for post-surgical support. I thank Kyu Hyun Lee and Ryan Ly for their work on data preprocessing pipelines for the Frank lab. I am grateful to all members of the Frank lab for stimulating conversations, feedback, and support throughout this project.

Collaborators at the Lawrence Livermore National Laboratories designed, fabricated, and assembled the polymer probes that were used to collect the data in this dissertation, including Allison Yorita, Razi Haque, Jose Hernandez, Paige Thompson, and Jenny Zhou. These extraordinary devices were essential to this work.

I am grateful to the UCSF Medical Scientist Training Program (MSTP) leadership and staff for support over the years, including Catherine Lomen-Hoerth, Mark Anderson, Aimee Kao, Geri Ehle, Tiffani Quan, Andrés Zepeda, and Amanda Andonian.

I am also grateful to the UCSF Bioengineering Graduate Program leadership and staff, including Christoph Schreiner, Tejal Desai, Shuvo Roy, Todd McDevitt, Victoria Starrett, Rocio Sanchez, Kristin Olson, and SarahJane Taylor.

Lastly, I am grateful to my friends and family for their love and support. I am particularly grateful to my sisters, Katie Fisher and Laura Guidera, and my parents, Ann and Steven Guidera, for their love and support over the years.

## Contributions

Chapters were written by Jennifer A. Guidera (JAG) and edited by Loren M. Frank (LMF). JAG and LMF conceived the project and experiments. JAG and Daniel P. Gramling (DPG) carried out activities essential for data collection. JAG performed surgeries, with support from DPG, Alison E. Comrie, and Abhilasha Joshi. Jenny Zhou, Paige Thompson, Jose Hernandez, Allison Yorita, and Razi Haque designed, fabricated, and prepared 128-channel polymer probes. JAG performed analyses. LMF and Christoph Kirst provided input on analyses. Eric Denovellis provided software for decoding analyses.

**A region-specific code for generalized representations across hippocampus and prefrontal cortex**

Jennifer Ann Guidera

## Abstract

In novel situations, animals can leverage past experiences to learn rapidly. This ability is thought to depend on abstraction: the representation of the common structure across related experiences. In mammals, the hippocampus (HPc) and the prefrontal cortex (PFC), including its medial prefrontal (mPFC) and orbitofrontal (OFC) subregions, are thought to support abstraction by expressing neuronal firing patterns that represent generalized features of experiences. Whether these firing patterns reflect a single, distributed representation of generalized features of experience, or whether each area specializes to represent particular features at particular times, remains unknown. To address this, we continuously monitored large neural ensembles in the HPc, mPFC, and OFC of freely behaving rats performing a cognitive task. We found evidence for regional specialization in the coding of generalized task features. First, HPc firing patterns were consistent with a primarily route-based coding scheme, whereas mPFC and OFC firing patterns were organized around the act of traveling between goal locations and the specific actions required to reach goals. Second, task representations in mPFC and OFC were most reliable during distinct task phases, suggesting these areas specialize to express consistent task representations in distinct behavioral periods.

## Table of contents

## List of figures

# Chapter 1

## Cognitive maps in hippocampus and prefrontal cortex: a review of the literature

### 1.1 Overview

Over the last several decades, there has been a shift in thinking on how the brain is thought to give rise to behavior. In the early- to mid-20th century, a prevailing view was that the brain learns to associate particular sensory inputs with particular behavioral outputs to maximize reward and minimize punishment[1]. An alternative view to this "behaviorist" conception arose, which held that rather than learn to associate stimuli with responses, the brain constructs mental models of the underlying structure of experiences[2,3]. These mental models were thought to influence how an organism interprets incoming sensory information and plans responses and were a departure from the direct stimulus-response connections envisioned by behaviorists[3].

A conception of the brain as building mental models and deploying them to guide behavior prevails today. However, relatively little is known about how mental models are instantiated in the brain. In this introductory chapter, we review previous findings from the literature that address this unknown. We limit our scope to two brain regions that have been implicated in this function, the hippocampus (HPc) and the prefrontal cortex (PFC).

We organize our discussion as follows. First, we briefly review early conceptions of mental models from the field of psychology that influenced present-day views of the brain. We then proceed to consider the literature that speaks to how mental models are instantiated in the HPc and PFC, considering each region in turn. We next discuss

studies that provide a direct comparison of these areas in how they may implement mental models. We end by describing gaps in our understanding of how mental models are implemented in the brain, that we hope our study as well as future ones may be able to address.

### 1.2 Origins of the concept of a "mental model" in psychology

Present-day notions of mental models derive from ones advanced by psychologists in the early- to mid-20th century. Particularly influential were those of Bartlett, Piaget, and Tolman.

Among these individuals, Bartlett was first to advance a view of mental models[2]. Bartlett's view was meant to provide a theory of how we remember. According to Bartlett, remembering is a constructive process that draws upon organized groupings of past experiences and responses, which Bartlett termed "schemata" (plural of schema). One of Bartlett's research studies provides an example of this. In the study, subjects were told a ghost story that had typical and atypical elements and asked to recall it sometime later. In their recollection, subjects would often leave out atypical elements of the story, suggesting they had forgotten these. Conversely, subjects would falsely recall details that had not been in the original story but that one might expect to have been. Bartlett interpreted these patterns as suggesting that subjects had reconstructed their memory of the original story in terms of an organization of similar stories that they had heard before. This "schema" facilitated their remembrance of typical elements, their

forgetting of atypical elements, and their mistaken recollection of details that had not been in the original story but that had been in other, similar stories.

Piaget subsequently popularized the notion of schema in the context of a theory of cognitive development in children[4]. Piaget noted that the kinds of mistakes children would make in identifying objects changed over time, suggesting children organized information differently as they aged. For instance, younger children may tend to misidentify a zebra as a horse, on the basis of identifying all four-legged creatures with a horse. In keeping with Bartlett, Piaget termed these organizations of past experiences schemata. Piaget's account provided further support for the notion that humans organize past experiences, and these organizations influence how we process information and respond in everyday life.

Tolman formulated a notion of mental models through observing the behavior of rats[3]. Rats exhibited particular behaviors that suggested to Tolman they form a mental map of their environment, which Tolman termed a "cognitive map". Tolman envisioned similar maps may form of non-spatial entities as well. These cognitive maps differed from the stimulus-response connections envisioned by behaviorists in several ways. They appeared to form in the absence of reward or punishment, and they could be used to reason about the optimal response to entirely novel sets of stimuli. As evidence suggesting that rats form these kinds of cognitive maps, Tolman observed that rats learn to take a direct route through a maze from a starting location to a reward much faster if they had previously explored the maze in the absence of reward, than if they were encountering the maze for the first time[3]. Tolman interpreted this as suggesting

that rats that had explored the maze had constructed a mental model of it, which they were able to deploy to efficiently navigate to a reward that was later introduced.

Here, we take a simplified view and refer to notions of a schema as envisioned by Bartlett and Piaget and of a cognitive map as envisioned by Tolman together, as supporting a single notion of mental models.

### 1.3 Neural implementations of cognitive maps

Here, we review evidence in support of the theory that the hippocampus and prefrontal cortex serve a cognitive mapping function.

#### 1.3.1 The hippocampus as a cognitive map

Perhaps the earliest conceptualization of a neural instantiation of a cognitive map was of the hippocampus as a cognitive map of space. This view was inspired by the discovery of spatially selective cells in the hippocampus of rats in the early 1970's[5]. These so-called "place cells" discharged at particular locations within an environment. By providing information about an organism's physical location, place cells appeared capable of supporting a representation of physical space in the abstract[6]. The envisioned notion was of a three-dimensional Euclidean space that existed independent of an organism ("allocentric space")[6]. Thus, it was proposed that the hippocampus instantiates a cognitive map of the kind originally envisioned by Tolman, applied to the particular case of allocentric space[6].

In the years that followed, several observations arose that challenged this initial conception. As rats ran back and forth on linear alleyways, the majority of place cells were found to fire selectively as rats traveled in one direction or the other, inconsistent with the notion that place cells signal global location in space[7]. In a separate study, about a third of place cells were found to fire at distinct locations depending on whether animals took variable or more stereotyped trajectories through an environment[8]. This finding indicated that the place cell signal was sensitive to non-spatial factors. Further work identified a subset of place cells that fired differently as animals passed through the same location in the same direction depending upon where the animal had come from or where it was headed, suggesting that the place cell signal could reflect elements of the past or future[9,10].

An alternative view of the hippocampus thus arose that could account for the apparent modulation of hippocampal place cells by aspects of an animals' past, current, or future experience as identified by these and other studies. The view conceived of the hippocampus as implementing a "relational map" that links events experienced in succession[11–13]. Traveling through space provided sequential experiences and is thus mapped by the hippocampus, according to this view. The proposed relational map was envisioned as flexible, in that nodes on the map that had not been directly linked during experience could nevertheless be related by virtue of their placement in relation to others. Recent theoretical work has provided a rigorous account of this conception of the hippocampus as providing a relational map[14].

The kinds of deficits that result from inactivating the hippocampus are consistent with the view that the area provides a relational map. Inactivating the hippocampus impairs rats' ability to relate two stimuli never experienced together but which were both paired with a third stimulus[15,16]. This deficit in transitive inference is expected if the intact hippocampus supports a relational map of the stimuli. Inactivation of the dorsal hippocampus was recently shown to disrupt the ability of rats to plan optimal actions in a task that requires anticipating how a sequence of actions relates to a final outcome[17]. This impairment in planning is expected if the hippocampus supports a relational map of the action–action and action–outcome relationships in the task.

Electrophysiological studies are also consistent with the notion that the hippocampus serves as a relational map. Akin to how place cells map space by virtue of firing in select locations, hippocampal cells were found to map a sequence of auditory frequencies by virtue of firing at select frequencies in the sequence[18]. This demonstrates that hippocampal cells are capable of mapping a series of non-spatial experiences, consistent with the notion that the hippocampus can form relational maps of all kinds. A separate study found that as non-human primates navigate distinct mazes in virtual reality, a subset of hippocampal cells encode subjects' abstract position in the task invariant to spatial location, consistent with mapping of an abstract task coordinate[19]. Additional work has shown that hippocampal ensembles in primates have a capacity for representing non-spatial task variables in an abstract format[20]. These studies identify neurophysiological correlates in hippocampus that could support relational mapping of experiences, broadly speaking.

### 1.3.2 The prefrontal cortex as a cognitive map

A longstanding view of PFC function is that the PFC organizes behaviors directed towards goals ("goal-directed behaviors")[21–23]. Although accounts of this view have not made explicit reference to the notions of a "cognitive map" or "schema" from psychology, the view has been cast in terms that bear strong resemblance to these notions. For instance, an influential account proposed that the PFC's role in organizing goal-directed behavior depends upon its representation of goals and the actions required to obtain them[23]. This bears resemblance to a proposal that the PFC serves as a cognitive map or schema of goal-directed behaviors.

In more recent years, the PFC has been explicitly theorized to support a "cognitive map"[24] or "schema"[25]. Two subregions of the PFC, the medial prefrontal cortex (mPFC) and the orbitofrontal cortex (OFC), have been implicated through separate lines of evidence. These lines of evidence respectively draw upon established ideas from long-term memory formation in the case of the mPFC[25,26], and reinforcement learning in the case of the OFC[24,27]. These proposals go beyond a straightforward reinterpretation of the dominant view of PFC function as organizing goal-directed behaviors.

We will first review longstanding evidence suggesting a primary role for PFC in organizing goal-directed behaviors, and briefly discuss how this evidence is consistent with a conceptualization of the PFC as a cognitive map of goal-directed behaviors. We

then consider the lines of evidence which have led to explicit proposals that the mPFC and OFC serve as cognitive maps.

#### 1.3.2.1 A prevailing view of the PFC as organizing goal-directed behaviors, cast as a cognitive mapping function

A prevailing view of PFC function is that the area is dedicated to the organization of behaviors directed towards goals[22,23,21]. The PFC is thought to accomplish this overall function through supporting several subordinate abilities, including attending to relevant stimuli and conversely inhibiting automatic tendencies that would otherwise distract from this ("inhibition"), holding items in memory for short periods ("working memory"), and adapting to changes in circumstances ("cognitive flexibility")[21,28]. These subordinate functions are termed "executive functions" as they are viewed as exerting top-down control over other mental processes and an organisms's behavior[22,28]. This prevailing conceptualization of the PFC derives from findings from lesion and electrophysiological studies, which we review next.

Lesion studies support the view that the PFC is essential for organizing behavior directed towards goals. Patients with PFC lesions often maintain an ability to execute simple motor actions, but have difficulty carrying out sequences of actions towards goals[21,23,29]. The specific deficits underlying this impairment vary by the PFC subregion that sustained damage. In particular, three distinct "prefrontal syndromes" have been described that result from lesions to three distinct areas of the PFC in humans[29]. Lesions to the dorsolateral PFC (dlPFC) disrupt an ability to correctly sequence actions

and adapt to changes in rules[21,29,30]. By contrast, lesions to the medial prefrontal cortex (mPFC), and particularly its anterior cingulate cortex (ACC) region, result in an indifference to circumstances and a decreased will to act[31,32,29,21]. In extreme cases seen with bilateral lesions, patients will not talk, eat, or act on their own volition, despite an intact ability to do so[33]. By contrast, lesions to the orbitofrontal cortex (OFC) result in socially inappropriate and disinhibited behavior, vacillation, and suboptimal choices among options associated with distinct long-term values[34,35].

Electrophysiological studies also support a role for the PFC in organizing behavior towards goals. These studies typically provide subjects with a task: a set of conditions for attaining some goal, often a reward of some kind. To perform the task, subjects must direct their behavior towards attaining the goal. Seemingly all experimenter-defined aspects of a task have been found to be encoded by PFC neurons, including information about stimuli, responses, and outcomes[36–42]. Oftentimes PFC neurons exhibit complex responses to multiple aspects of a task, a feature that has been shown to increase the representational capacity of a network and may thereby be useful for the representation of complex entities like cognitive behaviors[42]. Also notable is that PFC neurons can encode information that is not accessible from current sensory inputs but is relevant for attaining a goal. As an example of this, early electrophysiological studies of PFC identified cells that appeared to encode the memory of a stimulus relevant for making a choice sometime later[43–45]. As another rather striking example, in a task that rewards monkeys for responding to stimuli according to an abstract rule, the most prevalent type of encoding observed was of the abstract rule[46].

In summary, lesion studies demonstrate an essential role for PFC in carrying out goal-directed behaviors. Electrophysiological studies demonstrate that PFC activity is structured around sequences of sensory input, actions, and hidden context that are relevant for this purpose. On this basis, the PFC has been conceived of as expressing representations of goals and the means required to attain them, which then support the structuring of actions to reach goals[23].

This conception bears resemblance to a proposal that the PFC supports a cognitive map or schema of goal-directed behaviors, of the kind described in psychology. The PFC representation is reminiscent of a model of how events will unfold over time during goal-directed behaviors. Aspects of the past and future important to attaining goals are represented. In this way, the past and future bear upon the present, as Bartlett had originally conceived of schema[2]. Meanwhile, the prominent representation of abstract rules governing which actions result in goal attainment demonstrates a capacity for invariance to sensory specifics in the PFC representation that may support inference to never-before-seen sets of stimuli, characteristic of the cognitive maps envisioned by Tolman[3,46].

More recently, the PFC has been explicitly proposed to represent a cognitive map or schema. Proposals have arisen separately for mPFC and OFC through distinct lines of evidence. We review these proposals next, beginning with the proposal that mPFC represents a cognitive map.

#### 1.3.2.2 The medial prefrontal cortex as a cognitive map

Recent work has cast the mPFC as encoding "schema" of the kind described by Bartlett and Piaget[25]. This conceptualization derives from an influential theory known as systems consolidation theory which provides an account of how new memories are transformed into long-lasting ones. The theory conceives of the PFC as storing long-term memories that are more semantic in nature[47,48]. We briefly review this theory to provide context for the proposal that the mPFC serves to represent schema.

Systems consolidation theory holds that memories are encoded parallelly in hippocampus and neocortex: initially strongly in the hippocampus and weakly in the neocortex[48,49]. Over time, the replay of hippocampal memories is thought to strengthen corresponding neocortical memory traces, eventually producing long-lasting neocortical memories which can be recalled independently of the hippocampus[48].

The basis for this theory was the observation that hippocampal damage impairs the formation of new memories while sparing remote ones[48]. Subsequent work in rodent models reproduced this deficit and identified the opposite pattern of deficits following damage to the prefrontal cortex[50,51]. For instance, in one such model, mPFC lesions made at successive timepoints following conditioning of an eye blink response impaired retention of the response more and more as time passed[50]. By contrast, HPc lesions resulted in the opposite pattern, initially impairing retention of the learned response but having less of an impact as time went on and eventually no effect by four weeks after the initial learning[50]. A follow-up electrophysiological study employing the same animal model identified putative neural correlates of the long-term memory in the mPFC: single

cell responses to the learned response formed on a gradual timescale that paralleled the increasing memory dependence identified by the original lesion study[52].

The gradual nature of the consolidation process is thought to promote the formation of representations that are more semantic in nature and reflect similarities across multiple related experiences in the PFC[47,48]. Single cell firing patterns observed in the mPFC of rats performing navigational tasks agree with this prediction[53,54]. Cells in the mPFC tend to fire at particular phases of a navigation task, with a diversity of preferred phases across the ensemble[53,54]. These firing patterns are in marked contrast to the spatially selective firing observed in the rodent HPc in spatial navigation settings[7]. These studies support the notion that relatively more semantic memories exist in the mPFC.

The initial study that proposed a specialized role for the mPFC in schema representation drew upon systems consolidation theory and the notion of a schema as introduced by Bartlett and Piaget[25,26]. The study employed a rat model of knowledge frameworks[26]. Rats learn to associate particular locations in an arena with particular flavors of foods: the flavor of food received in a starting box would predict where more of that same flavored food could be found out in the arena. Once rats were well-trained on an initial set of flavor-location associations, they learned a novel flavor-location association. This new learning was rapid, consistent with the idea that rats had formed a schema-like representation of the task during initial learning that could facilitate rapid assimilation of new information[26]. Expression patterns of two immediate early genes (IEGs) that serve as markers of synaptic plasticity, *Zif268* and *Arc*, were examined in

the brains of rats that learned a novel flavor-location association after having been well-trained on an initial set of flavor-location associations. It was thought that during the learning of this novel association, rats incorporated a new memory into their preexisting schema. IEGs were found to be upregulated in a set of cortical areas that included subregions of the mPFC. Additionally, disruption of neurotransmission in the prelimbic region of the mPFC, which had the greatest upregulation of IEGs among cortical areas examined, impaired learning of the novel association. On this basis, the authors proposed an essential role for the mPFC in supporting schema.

#### 1.3.2.3 The orbitofrontal cortex as a cognitive map

The OFC has historically been viewed as a region essential for choosing among options based on their costs and benefits ("value-guided decision-making")[55]. This view arose from the observation that humans with damage to the OFC exhibit deficits in value-guided decision-making while performing normally on classic tasks of executive function[35,56–60]. In parallel, non-human primates and rodents with OFC lesions are prone to make suboptimal choices following a reversal of the values of two different options, consistent with a role for the OFC in value-guided decision-making[61–65]. Complementing these behavioral findings, cells in the OFC encode subjective value[66] and other value-based quantities[39,41,67–69], constituting a potential neurophysiological basis for the OFC's role in value-guided decision-making[70].

More recently, however, the OFC has been theorized to have a much broader role as representing a cognitive map of tasks[24]. This theory arose as a parsimonious

explanation for the specific nature of the deficits in value-guided decision-making observed following OFC damage in animal models, and it draws upon ideas in reinforcement learning. In particular, non-human primates and rats with OFC damage are able to learn novel stimulus-reward associations at the same rate as control subjects, but are impaired in subsequently reversing those associations when a stimulus no longer predicts reward ("reversal learning")[62–65]. It was proposed that these deficits are well described by a failure of subjects to form a mental model of the task that incorporates "hidden" information about which stimulus-outcome relationship is operative at any given time[24]. The identity of the rewarding stimulus is hidden in that it is not explicitly signaled and must be inferred from experiencing how a stimulus leads to an outcome. It is thought that by representing this hidden information, subjects can maintain distinct sets of stimulus-outcome relationships for distinct identities of the rewarding stimulus. Applying this idea to the case of reversal learning, subjects faced with a reversal can maintain the originally learned stimulus-outcome relationships, and learn a separate, parallel set of stimulus-outcome associations. By contrast, subjects that do not incorporate this hidden information about the identity of the rewarding stimulus into their model of the task must first unlearn previous stimulus-outcome associations, a lengthier process. Reinforcement learning models implementing these different learning strategies could match the behavior of control and OFC-lesioned subjects from previous studies[24], providing support for the theory that OFC is important for representing hidden task states.

The idea that the OFC represents a cognitive map of tasks has since gained further support. Inactivation of the OFC impaired the ability of rats to predict whether a stimulus would result in reward when this required inference based on previously experienced stimulus-stimulus associations[71]. In a separate study, inactivation of the OFC disrupted planning based on previously experienced action-outcome associations[17]. These findings suggest an essential role for the OFC in inferring unseen relationships between task elements, as would be expected if the OFC supports a relational map of tasks that flexibly links stimuli, actions, and outcomes relevant to task performance.

Complementary findings from studies of brain activity also support the proposal that the OFC represents a cognitive map of tasks. In human subjects performing a task that requires representing hidden information inaccessible from current sensory input, hidden information could be decoded from the blood-level oxygen dependent (BOLD) signal in the OFC[72]. The ability to decode hidden information was positively correlated with task performance, suggesting the representation of hidden information in the OFC may support task performance[72]. In a brain-wide analysis of BOLD activity, a region within the OFC was the only area from which all hidden information could be decoded, suggesting a specialized role for the OFC in encoding hidden information[72]. These findings are consistent with the proposal that the OFC represents a "complete" cognitive map of tasks including hidden aspects that cannot be directly observed but are essential for task performance. Consistent with the proposal that OFC represents a cognitive map

of tasks, several studies have identified a variety of task-related correlates of single cell firing in the OFC[38,40,73,74].

### 1.3.3 Comparisons of HPc, mPFC, and OFC as cognitive maps

A limited number of neuropsychological and electrophysiological studies have directly compared the hippocampus, mPFC, and/or OFC. We next discuss studies that make these direct comparisons and provide insight into potential similarities and differences in how these regions may represent cognitive maps.

#### 1.3.3.1 HPc vs. mPFC

Neurophysiological studies suggest that single cell and population-level representations of tasks generalize across contexts to a greater extent in the mPFC relative to in the HPc. In a spatial navigation task, cells in the mPFC had a greater tendency to fire similarly across distinct paths in a maze in comparison to HPc cells[53]. In a separate study, as mice performed an alternation task across spatial and motor contexts, single cell and population-level representations of the task generalized across contexts to a greater extent in the mPFC than in the HPC[75]. These studies support a longstanding theory that more generalized memories exist in the PFC as compared to HPc[47,76] and suggest that cognitive maps in each area may respectively emphasize similarities and differences across experiences.

The mPFC and HPc also appear to differ in the time frame in which they support memory. The HPc supports more recent memories, whereas the mPFC supports more

distant ones[50,51,77]. Thus, the HPc cognitive map may serve as a temporary site for newly learned relations, whereas the mPFC cognitive map may serve as a long-term store for relations, in keeping with systems consolidation theory[48].

#### 1.3.3.2 HPc vs. OFC

A pair of neurophysiological studies examined firing properties of HPc and OFC ensembles as rats performed an associative memory task and identified a relative emphasis on spatial and reward related information in these respective areas[78,79]. In the task, rats were placed in rectangular environments and were allowed to retrieve a reward consistently paired with one of two objects in the environment. Rats had previously learned the associations between the environments and whether each object would or would not contain reward. The activity of HPc and OFC ensembles tended to emphasize distinct types of task information. HPc firing patterns most prominently distinguished environments, followed by the position of objects within an environment, followed by whether objects did or did not contain reward, and lastly by the identity of the object[78]. By contrast, OFC firing patterns most prominently distinguished whether objects did or did not contain reward, followed by the conjunction of object and environment, and lastly by the position of objects within an environment[79].

These findings reveal a distinct hierarchical preferencing of task information in HPc and OFC, with a relative emphasis on spatial context in HPc, and on reward outcomes in OFC, but representation of the other type of information in each area. Correspondingly, it has been proposed that the HPc and OFC contain cognitive maps

with distinct emphases: on biological relevance in OFC, and on spatially and temporally linked events in HPc[80].

#### 1.3.3.3 mPFC vs. OFC

Several lesion studies suggest that the mPFC and OFC make distinct contributions to decision-making. In non-human primates, lesions to the mPFC were found to disrupt decision making on the basis of action-outcome associations but not stimulus-outcome associations, whereas the opposite deficit was seen following OFC lesions[81]. These findings suggest that mPFC and OFC may play greater roles in decision making based on motor or sensory information, respectively. In rats, lesions to the ACC led to a reduction in the amount of effort subjects were willing to invest in rewards, whereas lesions to the OFC led to a reduction in how long subjects were willing to wait for rewards[82]. These findings suggest that these areas make distinct contributions to the processing of costs and benefits. Finally, in non-human primates performing a decision-making task, behavioral impairments following ACC sulcus or OFC lesions were consistent with selective roles for the ACC sulcus in action selection, and the OFC in rapidly updating associations following a change in reward contingency[83].

Some neurophysiological studies have identified distinct functional properties in mPFC and OFC cells during decision-making tasks. In rodents performing a spatial two-arm bandit task, signals related to upcoming choice were identified in the mPFC but not in the OFC[84]. By contrast, the value of chosen options and reward prediction error were represented more prominently in the OFC than in the mPFC[84]. These findings may be

consistent with specialized roles for mPFC in action selection and OFC in updating the values of options. Of note, however, a study in non-human primates found value signals for decision variables were more prevalent in ACC than in OFC, and reward prediction error signals only in ACC[85].

### 1.4 Remaining questions

Several questions remain about how cognitive maps are instantiated in the brain. One question relates to the nature of cognitive maps in mPFC and OFC. Lesions in these areas produce distinct behavioral impairments[82,83], yet multiple task variables are encoded in both regions[37,46]. It remains unclear whether these areas exhibit differences in how they represent experiences that could help explain the distinct behavioral deficits that result from their damage.

More generally, it is unclear whether there are organizational principles governing how generalized representations are distributed in the brain. One possibility is that there is a single, distributed representation of generalized features, and areas differ in the extent to which their representations are generalized. A distinct possibility is that different areas specialize to represent distinct generalized features at different times.

An additional question is how cognitive maps in distinct areas may interface with each other to support learning and behavior. A prevailing theory is that HPc memories are replayed, and this promotes the formation of long-lasting neocortical memories that are more semantic in nature[48]. Thus, one hypothesis is that the HPc cognitive map

primarily functions to construct long-lasting neocortical cognitive maps. Little is known about the precise neurophysiological mechanisms governing this hypothesized process. For example, hippocampal network events thought to reflect memory replay are known to modulate the activity of PFC cells[86], but the long-term impacts of this coordinated activity on PFC representations remains unclear. For instance, it remains to be determined whether these patterns of coordination result in an upregulation or downregulation of accompanying PFC memories, some complex combination of these, or neither. Multisite recording in HPc, mPFC, and OFC may help in addressing these unknowns about how cognitive maps get built up in the brain, and their precise implementation within and distribution across distinct brain areas.

# Chapter 2

## Heterogeneous task representations in the hippocampus and prefrontal cortex

### 2.1 Overview

In this chapter, we describe experimental procedures and analyses aimed at studying hippocampal and prefrontal representations of a cognitive task in rats. We then describe the results of analyses.

### 2.2 Methods

*Animals*. Data from five male Long-Evans rats (550–700g) (Charles River Laboratories) were included in this study. Prior to continuous recordings, rats were housed in a temperature- and humidity-controlled facility on a 12-hour light/dark cycle (6AM–6PM). Rats were initially housed with 1–2 cage mates and had unlimited access to chow. Rats were singly housed once food restriction began. During continuous recordings, rats were housed in a separate facility on a 12-hour light/dark cycle (7AM–7PM). All procedures were approved by the University of California San Francisco Institutional Animal Care and Use Committee.

*Pre-surgical linear track training.* Rats were food restricted to 85–90% of their free-feeding weight and trained to run on an elevated linear track (1.1m, 84cm elevation from floor) with reward wells affixed on either end[87]. This training served to familiarize subjects with running on an elevated surface and receiving reward at set locations in an

environment. The training also taught subjects to wait up to two seconds at locations to receive reward. In each training session, rats were placed on the linear track for 15 minutes and allowed to behave freely. They were rewarded with a drop of sweetened evaporated milk reward (Nestle sweetened condensed milk plus 25g of sugar; approximately 80–82µL) at reward wells. Milk reward was delivered automatically via syringe pump through plastic tubing (Tygon, 3.18mm inner diameter, 6.35 outer diameter). Reward became available at a well once rats visited the other well. Over successive training sessions, a temporal delay between when rats arrived to a well and when reward was delivered was gradually introduced. The duration of the delay was fixed within each training session, and increased across training sessions from 0s, to 0.5s, to 1s, to 2s in duration. Once rats achieved a criterion of receiving 30 or more rewards in a session, the delay increased in the next session. Upon reaching this criterion with the 2s delay, pre-surgical linear track training ended, and subjects returned to a diet of ad libitum chow.

*Tetrode microdrive*. Tetrodes were spun from 12.5µm diameter nichrome wire and annealed via heat gun using an automated tetrode spinner (SpikeGadgets). Twenty-four tetrodes were loaded into independently moveable shuttles within a 3D printed body[87].

*Polymer probes*. One hundred twenty-eight channel polymer probes were obtained from Lawrence Livermore National Laboratory. Each probe has four shanks spaced 250µm

apart. Each shank has 32 linearly arranged contacts (15µm diameter) spaced 26µm apart. Probes were sterilized with ethylene oxide prior to implantation.

*Polymer probe and tetrode microdrive implantation*. Anesthetized rats underwent stereotactic implantation of 128-channel polymer probes (Lawrence Livermore National Laboratory) and a tetrode microdrive under sterile conditions. Rats were deeply anesthetized with isoflurane, and a mixture of ketamine (50 mg/kg), xylazine (6 mg/kg), and atropine (0.14 mg/kg) was injected intraperitoneally. For the duration of the surgery, anesthesia was maintained using isoflurane and additional intraperitoneal injections of a mixture of ketamine (25 mg/kg) and atropine (0.07 mg/kg) as needed. Body temperature was maintained using a water-based heating pad. Hydration was maintained via subcutaneous delivery of lactated ringers. Lidocaine (0.2 mL) was injected locally into the scalp and an anterior-posterior incision was made to expose the skull. Connective tissue was carefully removed from the surface of the skull. A ground screw was placed over the cerebellum, and additional screws were placed in the rear and front of the skull to provide additional anchoring for the implant.

Polymer probes were targeted to one or both hemispheres of the dorsal medial prefrontal cortex (mPFC) (+3.2 mm AP, ±0.89–0.94 mm ML, -3.2– 3.6 mm DV, 0–10° tilt) and/or orbitofrontal cortex (OFC) (+3.94–4.5 mm AP, ±1.91–2.34 mm ML, -3.6–4.05 mm DV) using a previously published aproach[88]. In one animal, polymer probes were also targeted to the infralimbic region of the mPFC (+3.2 mm AP, ±0.41–0.44 mm ML, -4.28–4.33 mm DV). A silicon elastomer (Kwik-Sil) was used to seal craniotomies. A 24-

tetrode microdrive was placed over the dorsal hippocampus (-3.8mm AP, ±2.6mm ML). Hippocampal craniotomies were sealed with a silicon sheet. The microdrive was then anchored to the skull using dental cement. A custom-build hybrid headstage (SpikeGadgets) was attached to the electronic interface board on the tetrode microdrive and to the intan boards connected to probes. A 3D printed funnel was placed around the front of the implant and filled with silicon elastomer in order to stabilize the probe electronics. A 3D printed case was placed around the entire implant. Bupivacaine (0.2 mL) was injected locally in the scalp, and sutures were placed to approximate skin in front of and behind the implant. Buprenorphine (0.01–0.02 mg/kg) and meloxicam (2 mg/kg) were administered subcutaneously following surgery for analgesia, and enrofloxacin (5 mg/kg) was given as an antibiotic. Rats were closely monitored in the days following the surgery for signs of discomfort and additional doses of buprenorphine and/or meloxicam were administered as needed to achieve analgesia.

*Tetrode adjustment*. In the weeks following surgery, 1–2 tetrodes were targeted to corpus callosum to serve as a reference, and the remaining tetrodes were slowly advanced to dorsal hippocampus. The presence of sharp-wave ripples and the orientation of sharp waves were used to estimate tetrode depth relative to the hippocampal cell layer[89].

*Post-surgical linear track training.* Between 3 and 6 weeks after surgery, rats were food restricted to 85–90% of their free-feeding weight and reintroduced to running on the

linear track. Every 1–2 days for 1–1.5 weeks, rats performed 1–3 15-minute sessions per day on the linear track with a 2s delay between well arrival and reward delivery. During this period, rats were also habituated to the home cage that would serve as their overnight housing during 24-hour recordings.

*Alternation task behavior*

*Fork maze*. Behavior took place in fork-shaped mazes. Each maze was custom-built from acrylic (TAP Plastics). Each maze consists of four arms measuring 73.0cm attached a central segment measuring 82.2cm. Two of these arms, the "handle" and "center" arms, attach perpendicularly to the connecting segment at its midpoint and head in opposite directions. The other two arms, the "left" and "right" arms, attach to the connecting segment at its left and right endpoints and head parallel to the center arm. The passageways along the arms and connecting segment are 7.6cm wide and flanked by walls 3.7cm tall. At the ends of each arm is an expanded rectangular segment measuring 12.1cm by 11.1cm and containing a reward well[87]. Each reward well has an opening for milk reward delivery and an infrared diode to detect arrival to wells. We refer to reward wells as "goal locations" or when it is clear from context, simply "goals". Mazes were elevated approximately 81cm from the floor.

*Alternation task*. In the fork maze, rats learned and subsequently performed a memory-based spatial alteration task[9]. In the task, rats receive milk reward at goal locations according to an alteration rule (Figure 2.4A). The rule specifies that subjects should

make alternating visits to the left and right goals from a designated home goal. Thus, from the left and right goals, rats receive reward at the designated home goal, and from the designated home goal, rats receive reward at the least recently visited of the left and right goals. On any trial, the designated home goal is either the center or handle goal as dictated by the variant of the task described in *Context exposures*. On the first goal visit of each session, rats receive reward at any of the left goal, right goal, and designated home goal.

An automated behavioral program written in Statescript (Spike Gadgets) detects infrared beam breaks at reward wells and delivers reward via syringe pump through plastic tubing (Tygon, 3.18mm inner diameter, 6.35 outer diameter) following a 2s delay. The syringe pumps were estimated to deliver 81–85uL of reward, with the exception of the first animal, for whom pumps were estimated to deliver 64–82µL of reward due to a pump programming error.

Goal visits are deemed “correct”, “incorrect”, or “neither correct nor incorrect” according to the following convention: on trials in which rats could possibly receive reward according to the alternation rule, a goal visit is considered correct if rats receive reward and incorrect if they do not. On the other hand, if rats did not have the opportunity to receive reward per the alternation rule, as on the first goal visit of a session or when the animals are starting from the arm that is never rewarded in the operative sequence, a goal visit is considered neither correct nor incorrect.

*Fork maze environments*. Rats learn to perform the alternation task in three distinct environments. Each environment is an enclosed compartment of a room containing a dedicated fork maze surrounded by a distinct set of global cues. Two environments, “Haight left” (HL) and “Haight right” (HR), are located in a first room, and walls prevent the animal from seeing one environment while located within the other. The third environment, “San Anselmo” (SA), is located in a second room.

*Continuous neural recordings.* Throughout behavior on the alternation task, including the initial and all subsequent exposures to the task and the environments, continuous neural recordings were performed 7 days a week and 24 hours a day with the exception of time between recording sessions or in sporadic cases of recording hardware failure. Rats were weighed daily prior to the first sleep session and following the final sleep session and received chow after daily experiments to maintain 85–90% of their free-feeding weight. Occasionally, tetrodes were adjusted after the final sleep session. Rats were house overnight in a home cage (33cm x 65cm x 51cm) located in the first room.

*Daily training sessions on the alternation task*. Each day, rats performed up to eight 20-min behavioral sessions. In each session, the rat was placed on one of the three fork mazes and allowed to behave freely. Rats were rewarded according to a variant of the alternation task as described below. These behavioral sessions were flanked by rest sessions approximately 45-min in length in a small box (32cm x 33cm x 41cm) located in Haight.

*Context exposures*. Rats underwent a series of exposures to the alternation task in distinct contexts (Figure 2.1). These contexts varied along two dimensions: the environment in which the task was performed, and the identity of the designated home goal. We defined the performance in a session as "at criterion" if at least 80% of the total number of correct and incorrect goal visits in a session were correct goal visits. Rats were exposed to contexts in the following order.

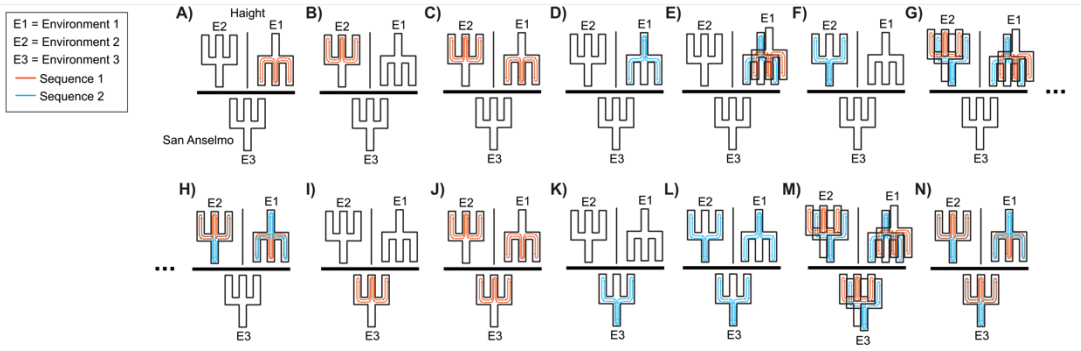


**Figure 2.1: Order of context exposures.** Context exposures are depicted for the case where the first experienced environment (E1) is "Haight right" and the first experienced sequence (S1) includes the center goal. Exposures are depicted in order from left to right, top to bottom.

*Initial learning of the alternation task*. Rats encounter the alternation task in a first environment (E1), either HL or HR, and with a first sequence of goal visits (S1), in which the designated home goal is either the handle or center goal (Figure 2.1A). We counterbalanced these assignments across rats. Once rats reached the performance criterion in each of two consecutive sessions, they performed an additional 2–3 sessions at criterion. These additional sessions were intended to ensure that rats had learned the task well before proceeding to novel contexts.

*Reacquisition of the alternation task in novel contexts*. Rats were then exposed to the alternation task in a variety of novel contexts, where they reacquired the task. Rats were said to have reacquired the task in a novel context once they performed at criterion in each of two consecutive sessions. Each exposure to a novel context occurred on a new day, with the exception of a single novel context exposure for the first rat to perform experiments. Novel context exposures were started on new days to match conditions during initial learning of the task to the extent possible, within the constraint that rats would first perform "reminder" sessions in familiar contexts as described below.

On the day of each novel context exposure, rats first performed reminder sessions in 1–2 familiar contexts. Rats typically performed these "reminder" sessions at criterion, but in the event they did not, they continued performing sessions in the familiar context until reaching criterion. These reminder sessions were meant to facilitate the study of memory generalization in two ways. First, we reasoned that for rats to generalize a memory of the task, they must have a memory of the task. We viewed rats performing at criterion as suggesting they maintained a memory of how to perform the task. Second, the reminder sessions facilitated the comparison of neural activity in familiar and novel contexts by increasing the chance of monitoring the same units across contexts due to close temporal proximity. Following reacquisition of the task in a novel context, in the next session rats entered a post-reacquisition phase in which they switched between familiar and novel contexts. Sessions from this phase allow

comparison of task representations across familiar and novel contexts following learning.

The first context in which rats reacquired the previously learned task consisted of the previously learned sequence (S1) in a novel environment (E2) (Figure 2.1B). For each rat, E2 was the environment among HL and HR that had not yet been encountered. Upon reaching criterion in the novel context, rats performed S1 in E1 and S1 in E2 on alternating sessions until reaching a performance criterion of 80% correct in two consecutive sessions, one in each environment (Figure 2.1C). Rats next learned S2 in E1 (Figure 2.1D). Upon reaching criterion, rats performed S1 in E1 and S2 in E1 on alternating sessions until reaching a performance criterion of 80% correct in two consecutive sessions, one on each sequence (Figure 2.1E). Rats were then exposed to S2 in E2 (Figure 2.1F). Upon reaching criterion, rats entered a multi-day period in which they switched among all four contexts according to several rotations (Figure 2.1G). Rotation 1 is defined as E2S1, E1S1, E1S2, E2S2. Rotation 2 is defined as E2S1, E2S2, E1S2, E1S1. Rotation 3 is defined as E2S1, E1S2, E2S2, E1S1. Rotation 4 is defined as E2S1, E1S1, E2S2, E1S2. Rats cycled twice through rotation 1, twice through rotation 2, once through rotation 3, and once through rotation 4.

Rats next learned to perform the task as the designated home goal reversed within a single session (Figure 2.1H). In this variant of the task, the home goal changed following eight consecutive correct trials. Across sessions, the environment and which sequence was operative first were varied. Rats completed this stage after encountering this reversal variant of the task in all sessions of a single day.

Rats next reacquired the first experienced sequence (S1) in a third novel environment (E3) located in a second room (Figure 2.1I). Upon reaching criterion, rats performed S1 in E1, E2, and E3 until reaching a performance criterion of 80% correct in three consecutive sessions, one in each environment (Figure 2.1J). Rats next reacquired S2 in E3 (Figure 2.1K). Upon reaching criterion, rats performed S2 in E1, E2, and E3 until reaching a performance criterion of 80% correct in three consecutive sessions, one in each environment (Figure 2.1L). Next, rats rotated among all six contexts in the order: E1S1, E2S2, E3S1, E1S2, E2S1, E3S2, repeating this three times (Figure 2.1M). Rats subsequently performed the variant of the task with the designated home well reversing within a session across environments for several days (Figure 2.1N).

*Histology*. Rats were deeply anesthetized with isoflurane, and electrolytic lesions were made at the ends of tetrodes. A day or more later, rats were deeply anesthetized with isoflurane, euthanized via injection of 1mL euthasol, and perfused with 4% paraformaldehyde (PFA). The head was placed in 4% PFA. After approximately one day in PFA, the brain was extracted and placed in 4% PFA for an additional day. The brain was then transferred to a sucrose solution for approximately five days for cryoprotection, then sliced in 60–80-micron coronal sections on a cryostat. Immunohistochemistry was performed to facilitate visualization of electrode tracts. Hippocampal sections were stained with 4',6-diamidino-2-phenylindole (DAPI) and cresyl violet or fluorescent Nissl (NeuroTrace 435/455 Blue, ThermoFisher). OFC and

mPFC sections were stained to visualize glial fibrillary acidic protein (GFAP) (primary antibody: mouse anti-GFAP antibody, Sigma-Aldrich; secondary antibody: donkey anti-mouse Alexa 594 antibody, ThermoFisher), and with fluorescent Nissl and DAPI. Stained sections were visualized using fluorescence or light microscopy (Nikon Ti-E Microscope) (Figure 2.2). We estimated the locations of electrodes by aligning sections containing tracts to the Paxinos and Watson rat brain atlas (6th edition)[90].

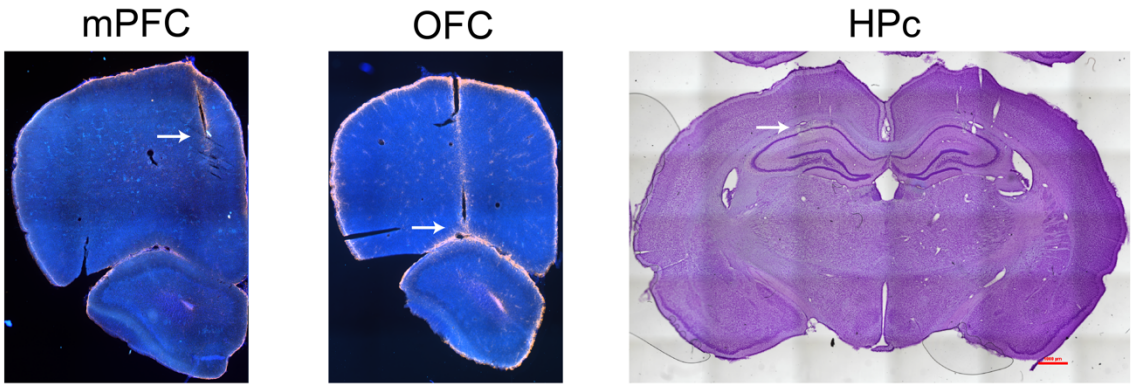


**Figure 2.2: Histological identification of electrode sites.** Fluorescence (mPFC and OFC) and brightfield (HPc) imaging of histological sections containing electrode tracts from one rat. Arrows point to tracts (mPFC and OFC) or a lesion site (HPc).

*Data acquisition.* Neural data (electric potential at polymer probe and tetrode contacts relative to a ground screw placed over the cerebellum) and behavioral events (infrared beam breaks at reward wells and syringe pump triggers) were continuously sampled at 30 KHz using a Spike Gadgets acquisition system. Video of behavior was obtained from an overhead camera at a rate between 41–50 frames per second. A set of red and green LEDs was mounted on the rat's head and used to track position.

*Data Processing*

*Position tracking.* The centroid of red and green LEDs was estimated using an offline tracking algorithm (SpikeGadgets). Head position measurements were obtained through smoothing and interpolation of these centroids over time.

*Spike sorting.* Spike sorting was performed to identify putative spikes of single neurons ("single-units"). In preparation for spike sorting, hippocampal tetrode recordings were referenced to a tetrode targeted to corpus callosum and bandpass filtered between 600 and 6000 KHz. Cortical recordings from 128 channel polymer probes were referenced to the median of recording traces on the same shank and bandpass filtered between 300 and 6000 KHz. Putative artifact times were identified as the two milliseconds around times when voltage exceeded 8 standard deviations or 500 mV on over one quarter of channels, and these times were excluded.

We performed spike sorting using MountainSort4, an automated spike sorting algorithm[91]. For hippocampal recordings, units were identified across tetrodes across concatenated run sessions within a day. For 128 channel polymer probes, units were identified in neighborhoods of channels within 115 microns of each other across concatenated run and sleep sessions within a day.

A first round of curation was performed on the units identified by MountainSort4 in an automated fashion using metrics calculated in SpikeInterface[92]. Units with metrics exceeding any of the following thresholds were rejected: a nearest neighbor noise

overlap of 0.03, an interspike interval violation ratio of 0.0025, and a time shift in the waveform peak of two samples[91,92].

A second round of curation was performed on the remaining units to further improve the identification of single-units. This round of curation involved visual inspection and had two objectives. First, different units that appeared to correspond to a single neuron whose spikes had been incorrectly split during spikesorting were merged. Second, individual units that appeared to correspond to more than one neuron were excluded. In the case of merging units that appeared to have been incorrectly split, we identified putative instances of two scenarios: the spikes of a single bursting neuron split into two clusters, and the spikes of a non-bursting neuron split into two clusters.

To identify unit pairs putatively corresponding to a single bursting neuron, we visually inspected unit pairs that met three criteria: 1) cosine similarity of average waveforms > 0.6; 2) interspike interval violation ratio less than 0.1; 3) cross correlogram constructed from spikes occurring within 100ms of each other has at least 100 spikes and meets an asymmetry criteria: the ratio of the sum of the density on the side of the correlogram with the greatest density and half the density at zero, to the total density, is greater than 0.6, and 4) the average waveform of the unit that tends to spike first versus second within a 200ms window has the larger peak amplitude. Of these candidate unit pairs, those whose spikes appeared to belong to a single neuron exhibiting successive bursts of spikes with decrementing amplitudes within tens of milliseconds were merged.

Second, to identify remaining cases of unit pairs that appeared to correspond to a single neuron, we visually inspected unit pairs with: 1) similarly shaped waveforms as

indicated by a cosine similarity of average waveforms of the two units greater than 0.7; 2) interspike interval violation less than 0.1; 3) cross correlogram constructed from spikes occurring within 100ms of each other has at least 100 spikes; 4) waveform amplitude densities have an overlap[93] greater than 0.7. Unit pairs that appeared to correspond to a single neuron were merged.

A very small number of units that did not pass these criteria, but appeared to correspond to a single neuron, were also merged. A third and final round of curation was performed on the resulting units using the same procedure employed during the first round of curation. Finally, units for which the maximal amplitude across waveforms occurred on a waveform measured on a channel outside of a target region (hippocampus, dorsomedial PFC including mPFC and ACC, and OFC) were excluded. The remaining units were used in downstream analyses.

*Data analysis*

*Software*. Analyses were performed in Python 3.7. Analysis pipelines were constructed using Datajoint[94].

*Data inclusion.* In this report, we focus on patterns of activity in dorsal mPFC, OFC, and HPc once rats have learned the task well. In one rat, the tetrode drive was found to have been placed too anteriorly, resulting in a large fraction of tetrodes mistargeted to areas anterior to the HPc. Due to concerns about the extent to which this data was representative of HPc, we excluded this data. For each rat, we included run sessions in

one day during which animals rotated among the four contexts in the first room (E1S1, E2S1, E1S2, E2S2) (6–8 run sessions per rat) (Figure 2.1G). By this time, rats were performing the task well above chance (Figure 2.3).

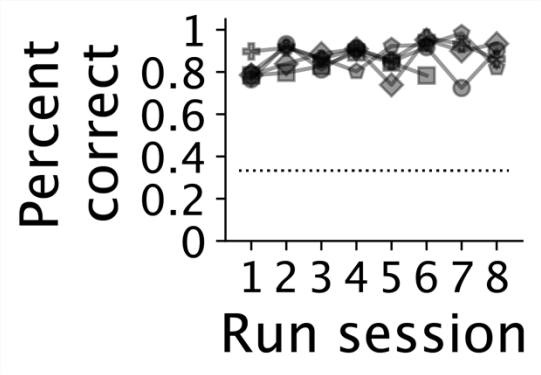


**Figure 2.3: Task performance.** Percent of trials performed correctly in the behavioral sessions analyzed in this work. Marker shapes correspond to individual rats and lines connect consecutive run sessions in a day.

*Single-unit inclusion*. All recorded units with an average firing rate of at least 0.1 Hz in the session being analyzed were included in analyses.

*Statistics*. We employed a hierarchical bootstrap procedure[95] to test for statistical significance and generate confidence intervals. The hierarchical bootstrap procedure provides a way to leverage power from repeated measurements at multiple levels of a dataset without increasing the rate of false positives[95]. The procedure consists of resampling with replacement at each level, then recomputing the test statistic of interest. This procedure is repeated several times to produce a "bootstrap distribution" of values that mimics what one may have observed from repeating an experiment several times. The $1 - \alpha$ confidence interval of the test statistic is formed from the $\alpha/2$ and $1 - \alpha/2$ percentiles of the bootstrap distribution.

To test for a significant difference between two group means, we performed the above hierarchical bootstrap procedure using the difference in group means as the test statistic and drawing one thousand bootstrap samples. A difference was deemed significant if the $1 - \alpha$ confidence interval of the resulting bootstrap distribution did not contain zero. To provide information about the strength of significance, we report the smallest significance level among $\alpha = 0.05, 0.01, 0.001, 0.0001$ at which results were significant. Confidence intervals shown in plots were constructed using a significance level of $\alpha = 0.05$.

*Definition of goal arrival and departure*. Arrival to a goal location was identified as the start of the first "down" state of the digital input detecting infrared beam breaks at the goal location, following a down state at another goal location. Departure from a goal location was identified as the end of the final down state at the location, prior to a down state at another goal location.

*Trials*. We define a trial as the time from the departure at one well to the departure at another well. We consider subsets of trials that correspond to distinct behavioral periods, as described in *Task periods*. When clear from context, we refer to these subsets as trials.

*Task periods*. For the purposes of analysis, we divided each trial into three task periods. The "path traversal" period begins upon a subject's departure from a goal location and

ends upon their arrival at a different goal location. The “delay” period begins upon arrival to a goal location and ends 2s later. At this time, a milk reward is delivered if the goal visit was correct. The “post-delay” period begins at the end of the delay period and ends when a subject departs from the goal location.

*Task progression*. During path traversals, we define task progression as the fraction of path rats have traversed. During the delay period, we define task progression as the time elapsed following goal arrival. During the post-delay period, we define task progression as the relative time between the end of the 2s delay period and departure from the goal location.

*Unit order in Figure 2.4E.* In Figure 2.4E, simultaneously recorded units from each brain region were sorted in ascending order by the location of average peak firing during a concatenation of the average firing rate in a session during traversals along the shown path, delay periods at the shown goal, and post-delay periods at the shown goal. For path traversals, we found the average firing rate in bins spanning 5% of total path length (approximately 3.65 cm) and convolved this with a Gaussian kernel of the same width. For the 2s-delay period, we found the average firing rate in 100ms bins spanning the delay period and convolved this with a 100ms Gaussian kernel. For the post-delay period, we found the firing rate in 100ms bins that spanned the period from 2s after goal arrival to departure and convolved this with a 100ms Gaussian kernel. Since this post-delay period can be of variable length depending on how long rats stay at the goal, we

converted time in each trial to the fraction of time elapsed between the end of the 2s delay period and departure from the goal. To form the firing rate map, we interpolated firing rate at twenty evenly spaced locations on this relative time axis.

*Firing rate vectors*. Population-level analyses, including principal component analysis (PCA), UMAP embeddings, and neural state trajectory analyses, were performed on vectors of single-unit firing rates estimated every 100ms. Single-unit firing rates were estimated by convolving spike times with a 100ms Gaussian kernel. To check robustness of results to z-scoring in Figure 2.12, single-unit firing rates were z-scored prior to forming firing rate vectors.

*PCA of population activity.* PCA was performed on firing rate vectors spanning individual behavioral sessions (Figure 2.5). For each brain region, the smallest number of principal components (PCs) needed to explain at least 80% of the variance in the data was identified (Figure 2.5A). The mean and associated 95% confidence intervals of this number across rats and sessions were found (Figure 2.5B). To assess how this number varies with ensemble size, we recomputed the number after randomly subsampling units without replacement to obtain ensemble sizes ranging from 5 to the number of units recorded in the brain region (Figure 2.5C).

*UMAP embeddings*. We used Uniform Manifold Approximation and Projection (UMAP) to visualize two-dimensional non-linear embeddings of the high-dimensional neural

activity space in which each axis corresponds to the firing rate of a single-unit[96]. UMAP embeds high-dimensional data in a lower dimensional space with the goal of preserving the local distance relationships between nearby data points. We embedded firing rate vectors (see *Firing rate vectors*) spanning single behavioral sessions.

We colored embedded neural states to visualize their organization with respect to task variables of interest, including task progression and path context (Figure 2.6, Figure 2.7). We visualized these variables separately in each of the three task periods. We defined task progression during each period as above (see *Task progression*). We defined path context during path traversals as the identity of the path the rat was on, and during the delay and post-delay periods as the identity of the path the rat had taken to the current goal. To obtain colors representing task progression, we discretized task progression values in one hundred equally spaced bins and mapped these to one hundred values from the jet color map. To obtain colors representing path context, we mapped path identity to a predefined color.

*Method for estimating levels of generality in single-unit representations of task progression.*

*Overview*. Progression through task periods occurs within distinct spatial contexts provided by the different paths and goal locations in a maze. In principle, single-unit representations of these progressions may be similar ("generalize") across any number of the distinct contexts. For instance, a single unit may exhibit characteristic firing rate modulations over the course of the delay period at a subset of goal locations. We use

the term “level of generality” to refer to the number of spatial contexts across which representations of task progression are invariant. Given $N$ contexts, there are $N$ levels of generality, one for each number from $1, \ldots, N$. At each level of generality, the number of ways to show generalized activity patterns across contexts is given by the binomial coefficient:

$$\frac{N!}{k!\,(N-k)!}$$

where $N$ is the number of available contexts and $k$ is the number of contexts generalized across, i.e., the level of generality.

We aimed to estimate the extent to which representations at different levels of generality contributed to single-unit representations of task progression. One approach to this could be to directly model single-unit firing as a function of covariates that represent task progression at distinct levels of generality. For instance, using a generalized linear model (GLM) framework, single-unit firing could be modeled as a function of a linear combination of covariates representing task progression at particular levels of generality. The fit coefficients for these covariates could be considered as reflecting contributions of representations of task progression at distinct levels of generality to the single unit firing. A potential pitfall of this approach, however, is that covariates at nearby levels of generality are likely to be correlated, and collinearity is a known challenge for coefficient estimation.

We therefore took an alternative approach to estimate the levels of generality in single-unit representations of task progression. Our approach relies on the idea that if the firing of a single unit represents task progression in the same manner across a

subset of contexts, then a model of the unit activity trained in any context in the subset should perform equally well in predicting activity in the other contexts in the subset, and perform poorly in predicting activity in contexts outside the subset. Consequently, distinct patterns of generalized representation across contexts should manifest as distinct patterns of cross-context model performance. From the observed pattern of cross-context model performance, we can estimate the contribution of representations at distinct levels of generality to single-unit representations.

To implement this approach, we first model the firing rates of individual units as a function of task progression in each context using a GLM and quantify the extent to which models trained in one context predict firing in others. We then model the pattern of cross-context generalization as a positive linear combination of patterns each reflecting generalization of firing across a particular subset of contexts. We consider the resulting coefficient for each pattern to be an estimate of its contribution to the observed single-unit representation of task progression. We describe this approach in detail next.

*GLM of single-unit firing as a function of task progression*. We modeled single-unit firing as a function of task progression using an established GLM framework[97–100]. In the framework, spikes are assumed to be derived from a Poisson process with a time-varying rate parameter that is function of a linear combination of covariates at each time. The function is chosen to be the exponential function with base $e$, which serves to guarantee that the rate parameter will be positive, as is required for the Poisson distribution. The rate parameter $\lambda$ is thus:

$$\lambda\left(c_1, \ldots, c_p, x_1(t), \ldots, x_p(t)\right) = e^{\sum_{j=1}^{p} c_j x_j(t)}$$

where $x_j(t)$ is the j[th] covariate at time t and $c_j$ is the coefficient for the j[th] covariate.

We built separate GLMs to predict firing during the path traversal and delay periods. We defined sets of covariates to capture task progression in each case using a raised cosine basis as described in Park et al.[99]. Compared to an indicator basis that represents the presence or absence of the rat in task progression bins with a binary variable, the raised cosine basis allows us to model task progression with relatively fewer covariates, which is expected to reduce overfitting. For GLMs of task progression during the path traversal period, we defined fourteen evenly spaced raised cosine bumps along each of the four potentially rewarded paths, for a total of 56 covariates. For GLMs of task progression during the delay period, we defined fourteen evenly spaced raised cosine bumps across the time in the delay at each of the three potentially rewarded goals, for a total of 42 covariates.

We trained and tested GLMs using the statsmodels python package[101]. During training, Broyden–Fletcher–Goldfarb–Shanno optimization[102] was used to find coefficients $c_1, \ldots, c_p$ that minimize the following cost function given a set of spike counts and covariates in 20ms time bins:

$$-\frac{l}{n} + \alpha \frac{\sum_{j=1}^{p} c_j^2}{2}$$

In the first term, $l$ is the model log likelihood and $n$ is the number of observations, which in our case is the number of time bins in the training set. The second term is an L2 penalty that we included to improve stability during training. In this term, $c_j$ is the coefficient for the j[th] covariate, and $\alpha$ is a free parameter that scales the impact of the

regularization term on the overall cost. We set $\alpha = 0.00005$. During testing, a trained model was used to predict spike counts based upon on covariate values in 20ms time bins.

*Training and testing GLMs across contexts*.

For path traversals and the delay periods at goal locations, we trained a model in each context and tested that model in the same and every other context. During path traversals, we defined context as path identity. During delay periods, we defined context as goal identity. For computational tractability, we considered only the subset of paths and goals that could yield reward according to the variant of the alternation task operative in a given session ("potentially rewarded").

These analyses were performed across data from many trials, where each trial corresponds to distinct traversals of a path in the case of the path traversal period, and distinct 2s-delays at a goal in the case of the delay period. When testing a model in the same context as that in which it was trained, we performed leave one out cross validation: we tested each trial on a model trained on all but that trial. We then combined results across trials. When testing a model on a different context from that in which it was trained, we tested on all trials from one context using a model trained on all trials from the other context.

For each pair of train and test context, we quantified model performance as the Pearson correlation coefficient between predicted and actual firing rates. We chose this metric as it in part reflects the informativeness of unit firing about progression through

the task period. As an extreme example of this, a cell that fires during a task period but exhibits no reliable modulations will have a low correlation value. By contrast, an area-based measure would yield a high value for such a cell. With regard to our dataset, this is particularly relevant for prefrontal cortex neurons as these tend to have nonzero baseline firing rates, resulting in high area-based values even when no significant modulations in firing rate are present.

To obtain firing rates, we convolved spike counts with a 100ms Gaussian kernel. We found the average and associated 95% confidence intervals of the correlation between predicted and actual firing rates for each brain region during the path traversal and delay periods across the following levels: rats, sessions, and contexts (Figure 2.9).

*Estimation of levels of generality in single-unit representations*.

For each single-unit, we approximated the vector of model performance values for pairs of train and test contexts with a positive linear combination of basis vectors that together capture all possible subsets of contexts across which a representation could generalize (Figure 2.10):

$$v_i \approx \sum_{b=1}^{B} a_b x_b$$

where $v_i$ is the observed vector of correlation coefficients between actual and predicted firing rates for each possible pair of train and test contexts for unit $i$, $x_b$ is the vector of correlation coefficients between actual and predicted firing rates for each possible pair of train and test contexts given generalization across the b$^{th}$ subset of contexts, $a_b$ are

positive coefficients, and $B$ is the total number of bases. The total number of bases is the sum of the number of ways to generalize across contexts at each level of generality:

$$B = \sum_{k=1}^{N} \frac{N!}{k!\,(N-k)!}$$

where $N$ is the total number of contexts.

Each basis contains model performance values that would be expected given a perfectly reliable representation of task progression that generalizes across a particular subset of contexts. Specifically, a basis that represents generalization of a task progression representation across exactly $j$ contexts would have model performance values of one for train/test pairs formed from the $j$ contexts, and model performance value of zero for any other train/test pair.

To approximate $v_i$, the observed vector of correlation coefficients between actual and predicted firing rates for pairs of train and test contexts for unit $i$, as a positive linear combination of basis vectors, we used non-negative least squares. To aid interpretation, we scaled the fit coefficients for each basis vector. Specifically, we multiplied each coefficient by the number of contexts to which it applies, and divided by the number of contexts in which there was at least one spike. As an illustration of the utility of this, consider two units that both represent task progression across all four possible path contexts, but in distinct ways. The first unit exhibits different firing rate patterns along different paths, whereas the second unit exhibits the same firing rate pattern along different paths. The first unit would have a coefficient of one for each of the four basis vectors corresponding to representation of a single context, resulting in a sum of coefficients of four. By contrast, the second unit would have a coefficient of one for the

single basis vector that corresponds to generalized representation across the four paths. After scaling coefficients as above, the two units would both have a coefficient sum of one, which accords with our intuition that both units reliably represent task progression across the maximum number of contexts.

Prior to performing this estimation, we set the prediction value to zero for pairs of train and test contexts for which no spikes occurred in the train context, to reflect the fact that no information about the testing context should be available from the training context. Similarly, for pairs of train and test contexts for which no spikes occurred in the test context, we set the prediction value to zero to reflect the fact that there are no events in the test context to predict.

We performed the above procedure for each unit. We show the resulting scaled coefficients for an example recording (Figure 2.8G). We also found the mean and associated 95% confidence intervals of scaled coefficients within each brain region and level of generality across sessions and rats (Figure 2.8H).

*Single-trial firing rate maps*. In Figure 2.8A–C, we computed single-trial firing rates as a function of path fraction in 20 bins spanning the path (approximately 3.65 cm bin width) and convolved the resulting values with a Gaussian kernel of the same width. In Figure 2.8D–F, we computed single-trial firing rates as a function of time in the delay period in 100ms bins and convolved the resulting values with a 100ms Gaussian kernel.

*Examination of the geometry and dynamics of population-level neural activity.*
We examined the geometry and dynamics of the population-level neural activity during task progression within the high-dimensional space in which each axis corresponds to the firing rate of a single-unit. We examined geometry in terms of the Euclidean distances between neural states and dynamics in terms of the directions in which neural states evolved.

First, we approximated the neural trajectory during task progression using a piece-wise linear curve[103]. The curve consists of straight line segments joined end to end. To define the curve, we specify the "knots": the locations where two line segments intersect. We define the knots in terms of the activity within task progression bins as follows. We define ten bins of equal width spanning each path, and ten bins of equal width spanning the relative time in the delay period at each goal location. We define knots as the average of consecutive neural states within each of these bins. The series of these knots over time and the line segments that connect them comprise the curve.

*Comparison of average dynamics across neural state trajectories along distinct paths*.
To test the hypothesis that the dynamics of neural state trajectories are more similar along paths with the same versus a different set of turn directions, we compared the direction in which neural states evolved from one task progression bin to the next on average (Figure 2.11A). Specifically, we found the average of the difference vectors along each path that connect each pair of contiguous task progression bins and computed the cosine similarity of average difference vectors at corresponding path

progressions along paths that did or did not share the same set of turn directions, with the exception of pairs of paths that branched from a common maze segment which we excluded in order to avoid confounding by spatial tuning. The cosine similarity between two vectors $a$ and $b$ is:

$$\frac{a \cdot b}{||a||\,||b||}$$

where $\cdot$ is the dot product, and $||\cdot||$ denotes vector length.

For each brain region and rat, we found the average and associated 95% confidence intervals of the cosine similarity of difference vectors across path pairs and sessions (Figure 2.11B, Figure 2.12). To summarize results across rats, we found the average and associated 95% confidence intervals of the cosine similarity of difference vectors in the first, middle, and last third of paths across path pairs, sessions, and rats (Figure 2.11C).

*Quantification of the consistency of neural state dynamics during task progression.* To quantify the consistency in the neural state dynamics as a function of task progression across trials, we computed the cosine similarity between pairs of difference vectors separated by more than ten seconds occurring at a given position in the task progression along each path or in the ensuing the delay period (Figure 2.14A). We found the average of this measure and associated 95% confidence intervals for each rat and brain region across contexts and sessions (Figure 2.14B).

*Quantification of the consistency of neural state geometry during task progression*. To quantify the consistency in the neural state geometry as a function of task progression across trials, we computed the Euclidean distances between firing rate states within the same and different task progression bins separated by at least 10s. We did this separately for path traversals and the ensuing the delay periods (Figure 2.15A). An example of the resulting values is shown for an individual session (Figure 2.15B). To compare the distances of neural states within the same task progression bin to those in different task progression bins, we defined a "similarity metric" as one minus the ratio of these distances. A greater value indicates that states in the same task progression bin tend to be closer in comparison to those in different task progression bins. For each rat and brain region, we found the average of this similarity metric and associated 95% confidence intervals across paths and sessions (Figure 2.15C).

*Population decoding*. We decoded task progression during path traversals or the delay period using a previously published Bayesian state space model[104]. Separately for the path traversal and delay period, we decode each trial using held out trials. In each case, we defined a set of contiguous 2ms time bins that spanned trials, denoted as $t_1, \ldots, t_T$ where $t_k$ is the k$^{th}$ time bin. We define latent variable $x_k$ that correspond to the neural representation of task progression at $t_k$. For the purposes of algorithm fitting, we scaled task progression to have similar range to linear position, which was the original use case for the algorithm. We perform the inverse scaling on the algorithm output to display results on the original scale of task progression.

We define $I_k$, a discrete latent variable that denotes whether the neural representation of task progression evolves according to continuous or fragmented dynamics. The model simultaneously estimates the posterior probability of task progression and the dynamics: $p(x_k, I_k \mid O_{1:T})$ where $O_{1:T}$ is the set of observed spike counts in each time bin from 1 through T. The estimation is carried out via a pair of equations. First, we apply a causal filter equation, beginning with initial conditions $p(x_0, I_0)$ and recursively iterating from $t_1$ to $t_T$:

$$p(x_k, I_k \mid O_{1:k})$$
$$\propto p(O_k \mid x_k, I_k) \sum_{I_{k-1}} \int p(x_k \mid x_{k-1}, I_k, I_{k-1}) Pr(I_k \mid I_{k-1})\, p(x_{k-1}, I_{k-1} \mid O_{1:k-1}) dx_{k-1}$$

Second, we apply an acausal smoother equation, recursively iterating backwards in time from $t_T$ to $t_1$:

$$p(x_k, I_k \mid O_{1:T})$$
$$\propto p(x_k, I_k \mid O_{1:k}) \sum_{I_{k+1}} \int \frac{p(x_{k+1} \mid x_k, I_{k+1}, I_k) Pr(I_{k+1} \mid I_k)}{p(x_{k+1}, I_{k+1} \mid O_{1:k})}\, p(x_{k+1}, I_{k+1} \mid O_{1:T}) dx_{k+1}$$

where:

$$p(x_{k+1}, I_{k+1} \mid O_{1:k}) = \sum_{I_k} \int p(x_{k+1} \mid x_k, I_{k+1}, I_k) Pr(I_{k+1} \mid I_k)\, p(x_k, I_k \mid O_{1:k}) dx_k$$

Four quantities must be defined or estimated in order to specify this model: the initial conditions $p(x_0, I_0)$, the dynamics movement model $Pr(x_k \mid x_{k-1}, I_k, I_{k-1})$, the dynamics transition matrix $Pr(I_k \mid I_{k-1})$, and the likelihood of the observations $p(O_k \mid x_k, I_k)$. To reflect our lack of prior knowledge about the initial latent positions and dynamics, we define the initial latent positions and initial dynamics as uniformly distributed:

$$p(x_0, I_0) = \frac{1}{number\ of\ available\ dynamics} U(\min x, \max x)$$

where $U$ denotes the uniform distribution.

The neural representations may evolve with distinct dynamics. The algorithm aims to model this explicitly. Two dynamics are specified that respectively capture continuous and non-continuous evolution of task progression. In the "continuous" dynamic, the next latent variable value is normally distributed around the current latent variable value. In the "fragmented" dynamic, the next latent variable value is uniformly distributed over all possible values of the latent variable. We assume that we have no information about the value of the latent variable when transitioning to or from the fragmented state, which we capture using the fragmented dynamic. These choices are reflected in the dynamics movement model:

$$\Pr(x_k \mid x_{k-1}, I_k, I_{k-1}) = \begin{bmatrix} N(x_{k-1}, 6.0) & U(min\ x, max\ x) \\ U(min\ x, max\ x) & U(min\ x, max\ x) \end{bmatrix}$$

with row indices corresponding to $I_{k-1} = [continuous, fragmented]$ and column indices corresponding to $I_k = [continuous, fragmented]$. We define the dynamics transition matrix to reflect the prior expectations that dynamics last 100ms on average, and that there is a small probability of changing to other dynamics:

$$\Pr(I_k \mid I_{k-1}) = \begin{bmatrix} 0.98 & 0.02 \\ 0.02 & 0.98 \end{bmatrix}$$

with row indices corresponding to $I_{k-1} = [continuous, fragmented]$ and column indices corresponding to $I_k = [continuous, fragmented]$.

We compute the likelihood of observations $p(O_k \mid x_k,\ I_k)$ using an encoding model:

$$p(O_k \mid x_k, I_k) \propto \prod_{i=1}^{C} (\lambda_i(t_k|x_k)\Delta_k)^{N_{t_k}^i} e^{-\lambda_i(t_k|x_k)\Delta_k}$$

where $N_{t_k}^i$ represents a spike in time bin $k$ from cell $i$, $\Delta_k$ is the size of the time bin, and $\lambda_i(t_k|x_k)$ is the instantaneous firing rate of cell $i$ given $x_k$. We estimate $\lambda_i(t_k|x_k)$ through kernel density estimation. We perform decoding using leave one out cross validation on trials during the path traversal or delay periods.

## 2.3 Results

We recorded from large ensembles of neurons in the mPFC, OFC, and HPc as rats performed a spatial alternation task (Figure 2.4). In the task, rats are rewarded for traveling between goal locations according to an alternation rule (Figure 2.4A). Subjects perform the task within different environments and with distinct sets of goal locations within the same environment (Figure 2.4C).

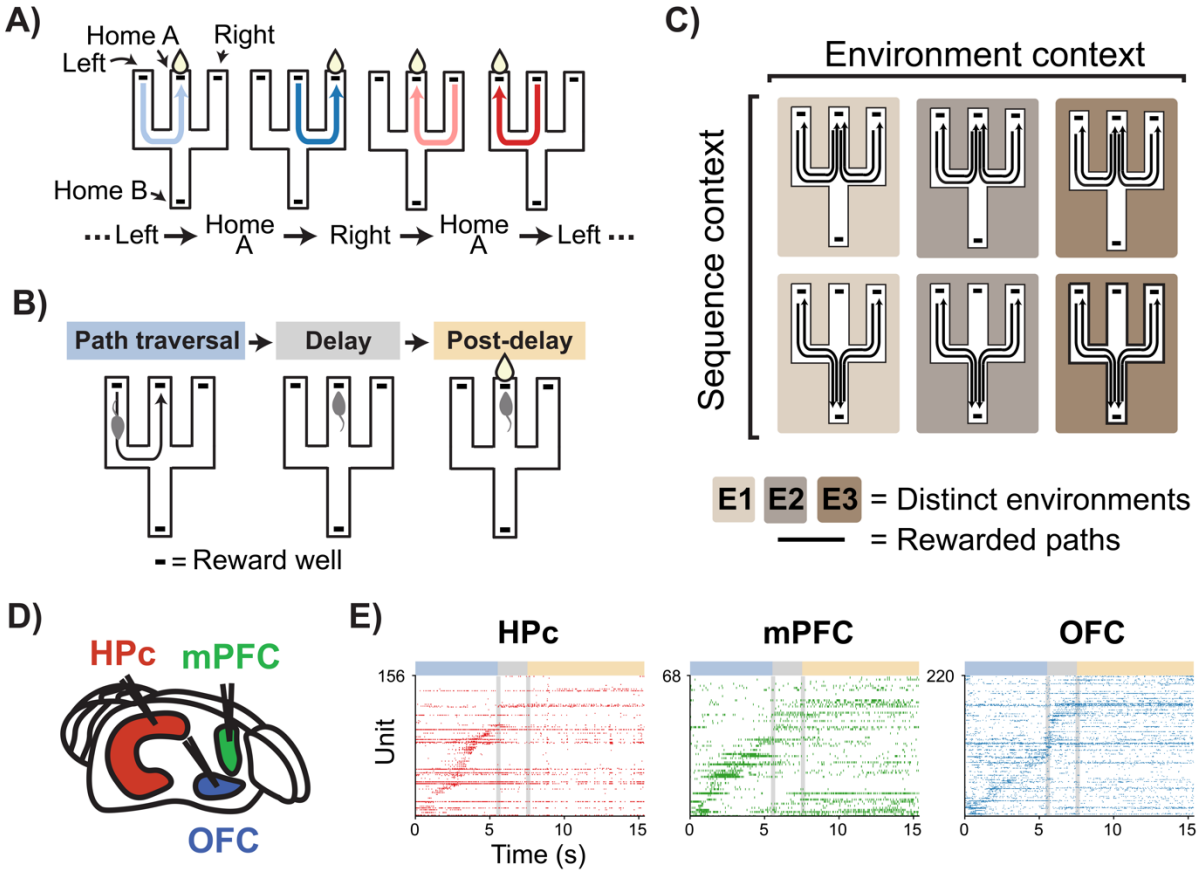


**Figure 2.4: Multisite recordings in rats performing a cognitive task. A:** Alternation rule. Rats receive a milk reward at a designated "home" goal from the "left" or "right" goals, and at the least recently visited of the left and right goals from the designated home goal. A sequence of correct well visits according to this rule is illustrated with home A as the designated home goal. **B:** Task periods defined for the purpose of analysis. The "path traversal" period begins upon departure from a goal location and ends upon arrival at a different goal location. The "delay" period begins upon arrival at a goal location and ends 2s later. At the end of this period, a milk reward is delivered if the correct goal location had been visited. The "post-delay" period begins

at the end of the delay period and ends upon departure from the goal location. **C:** After learning the alternation rule in a first environment and with an initial designated home goal, rats reacquire the rule in novel environments and with a different designated home goal. **D:** Recording sites. **E:** Spike times from simultaneously recorded HPc, mPFC, and OFC units during one trial of a 20-min behavioral session.

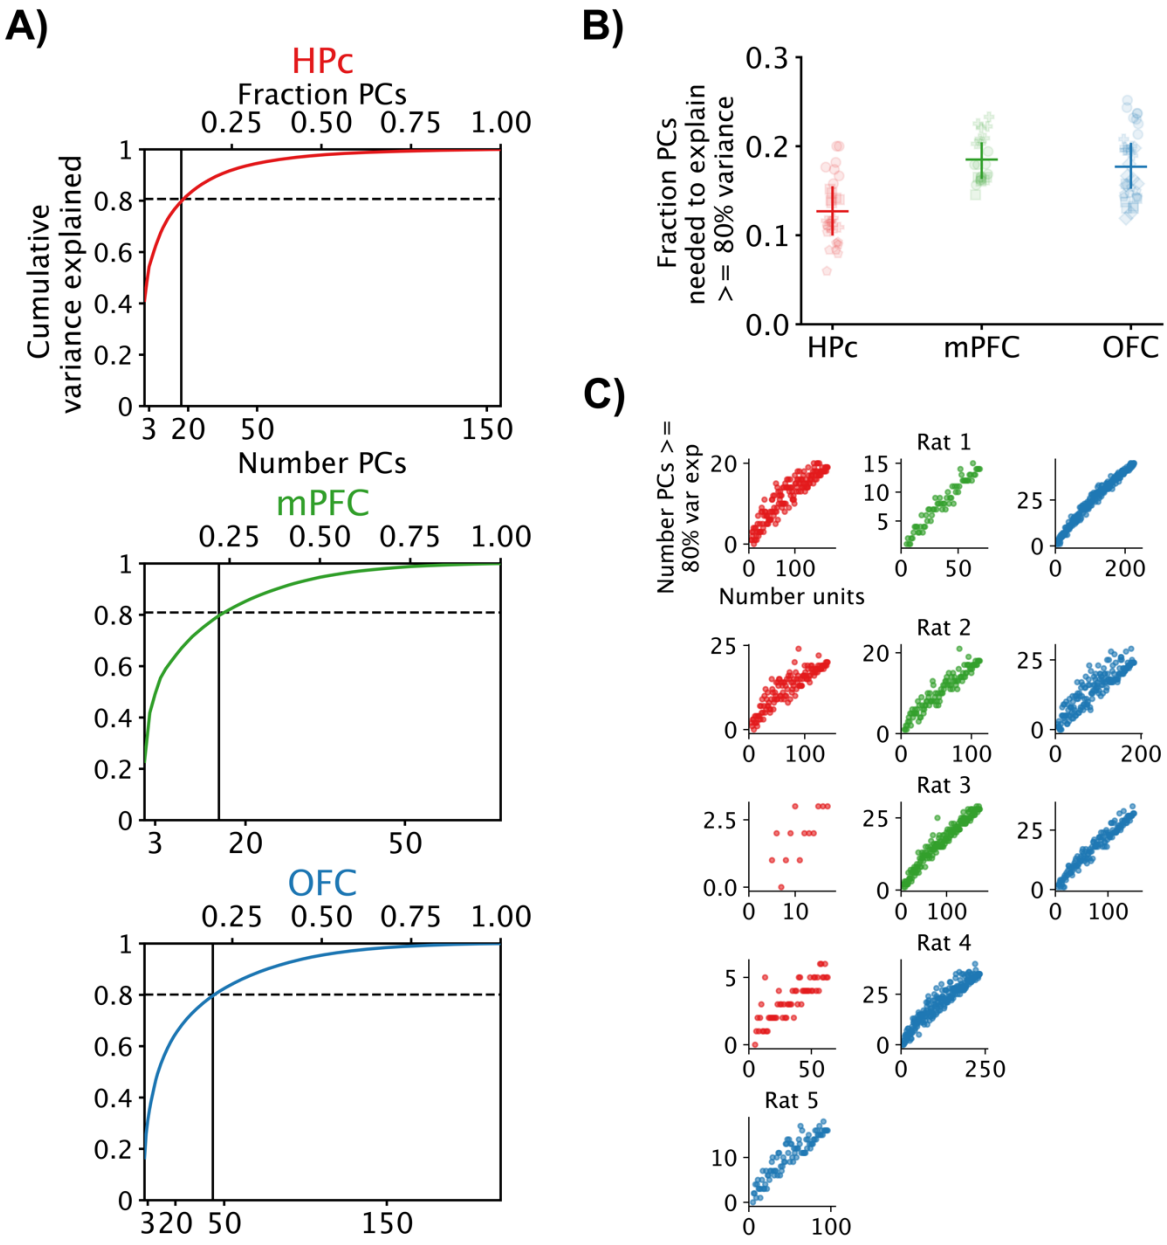


**Figure 2.5: Principal component analysis of population activity**. **A:** Cumulative variance explained as a function of the number (bottom axis) or fraction (top axis) of PCs, for HPc (top), mPFC (middle), and OFC (bottom) firing rate vectors during a 20-

min behavioral session. Vertical lines mark the smallest number of PCs needed to explain at least 80% of the variance. Horizontal lines mark the corresponding variance explained. **B:** Fraction of PCs needed to explain at least 80% of the variance in single sessions (markers) or on average across sessions and rats (horizontal lines) (associated 95% confidence intervals shown with vertical lines). **C:** Fraction of PCs needed to explain at least 80% of the variance as a function of subsampled ensemble size, shown for one representative session from each rat.

For this study we focused on activity recorded once animals were familiar with the environments and the task, allowing us to understand the relationship between neural activity and behavior under conditions of minimal behavioral variability (See *Data Inclusion* in methods).

We first set out to identify the key task-related variables represented in each region. In the context of recordings from large populations of neurons, the current standard approach is to use a linear dimensionality reduction technique such as PCA to identify a low-dimensional subspace of activity that can be visualized and that captures the majority of the variance in neuronal firing rates over time[105,106]. We therefore began by measuring neuronal firing rates in 100ms bins throughout each run session on the track and applied PCA to these data.

We found that the neural firing patterns in our task were not low dimensional (Figure 2.5). Instead, the number of dimensions required to explain 80% of the variance in the data grew approximately linearly with the number of neurons measured in a given region (up to ~250). Given the absence of a clear low-dimensional subspace that could be derived from linear methods, we turned to a method that can capture low-dimensional, non-linear structure in the data, UMAP (uniform manifold approximation and projection[96]), to gain insight into the structure of the neural activity patterns. UMAP produces lower dimensional spaces wherein the relative distances among points are

preserved to the greatest extent possible. As such, firing rate vectors that are close together in the full neural firing rate space are mapped to lower dimensional vectors that remain close together.

Applying a 2D UMAP to the same firing rate vectors revealed striking differences in the structure of embedded neural population activity in relation to the task across brain areas. Examining these visualizations led us to divide each trial into three non-overlapping periods (path traversal, defined as the time from when the animal left one reward well and arrived at another; delay, defined as the two second period beginning at arrival to a reward well and terminating, on rewarded trials, at the time of reward delivery; and post-delay, defined as a the time from the end of the delay until the animal left the reward well). Here we show 2D embeddings color coded to reflect the progression of firing rate vectors across each of these periods (Figure 2.6).

As expected from the well-established spatial and directional specificity of hippocampal "place cells"[7], embedded HPc activity was primarily organized around a subject's physical location and direction of travel. This organization persisted during the delay period and to some extent following the delay, with an additional progression consistent with the transition from movement- to immobility-related place activity[107,108].

By contrast, embedded mPFC and OFC activity was chiefly organized around the trial structure and the associated journey from one goal location to another (Figure 2.6A). Within this global organization, embedded neural states in both areas diverged along paths that required distinct turn sequences (Figure 2.6B). Thus, within the overall task-related progression, there appeared to be a nested representation of the actions

required to reach goals. These findings suggested a first organizational principle: that both mPFC and OFC show hierarchical structuring of activity around journeys to goals, with the journey between any two goal locations represented most prominently, followed by representations of specific kinds of journeys.

A close inspection of the embeddings further suggested an additional organizational principle. Namely, we observed a striking difference between mPFC and OFC subregions in when these activity patterns appeared most reliable. In mPFC, there was clear segregation between same turn paths along much of their extent. This segregation was less prominent in OFC. Conversely, in OFC, neural states appeared increasingly organized around task progression with proximity to the arrival at goal locations and into the delay period (Figure 2A). The same dynamic pattern was not evident in the mPFC activity (Figure 2A). This led us to hypothesize that mPFC and OFC might specialize to express reliable representations of task progression at distinct behavioral phases of the task.

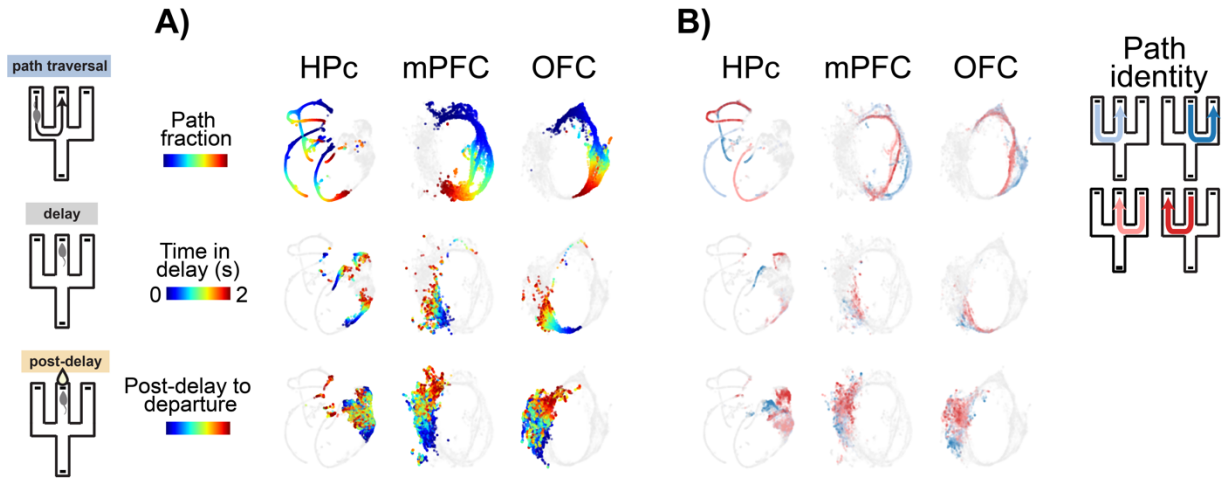


**Figure 2.6: Nonlinear embedding of ensemble activity.** UMAP embeddings of HPc, mPFC, and OFC firing rate vectors (left, middle, and right in each panel) during a single

behavioral session. Embedded neural states are colored according to task progression (A) or path identity (B) during path traversals (top), the delay period following goal arrival (middle) and the post-delay period (bottom). **A:** Embedded firing rate states are colored according to the fraction of any path traversed (top), time in the delay period (middle), and relative time in the post-delay period (bottom). **B:** Embedded firing rate states are colored according to the identity of the path rats were on (top) or had come from if at a goal location (middle and bottom).

We next set out to test these hypotheses at the single-unit and population levels. We included data from sessions in different environments and with different goal locations, because the patterns we had observed in population activity embeddings were evident across these conditions (Figure 2.7).

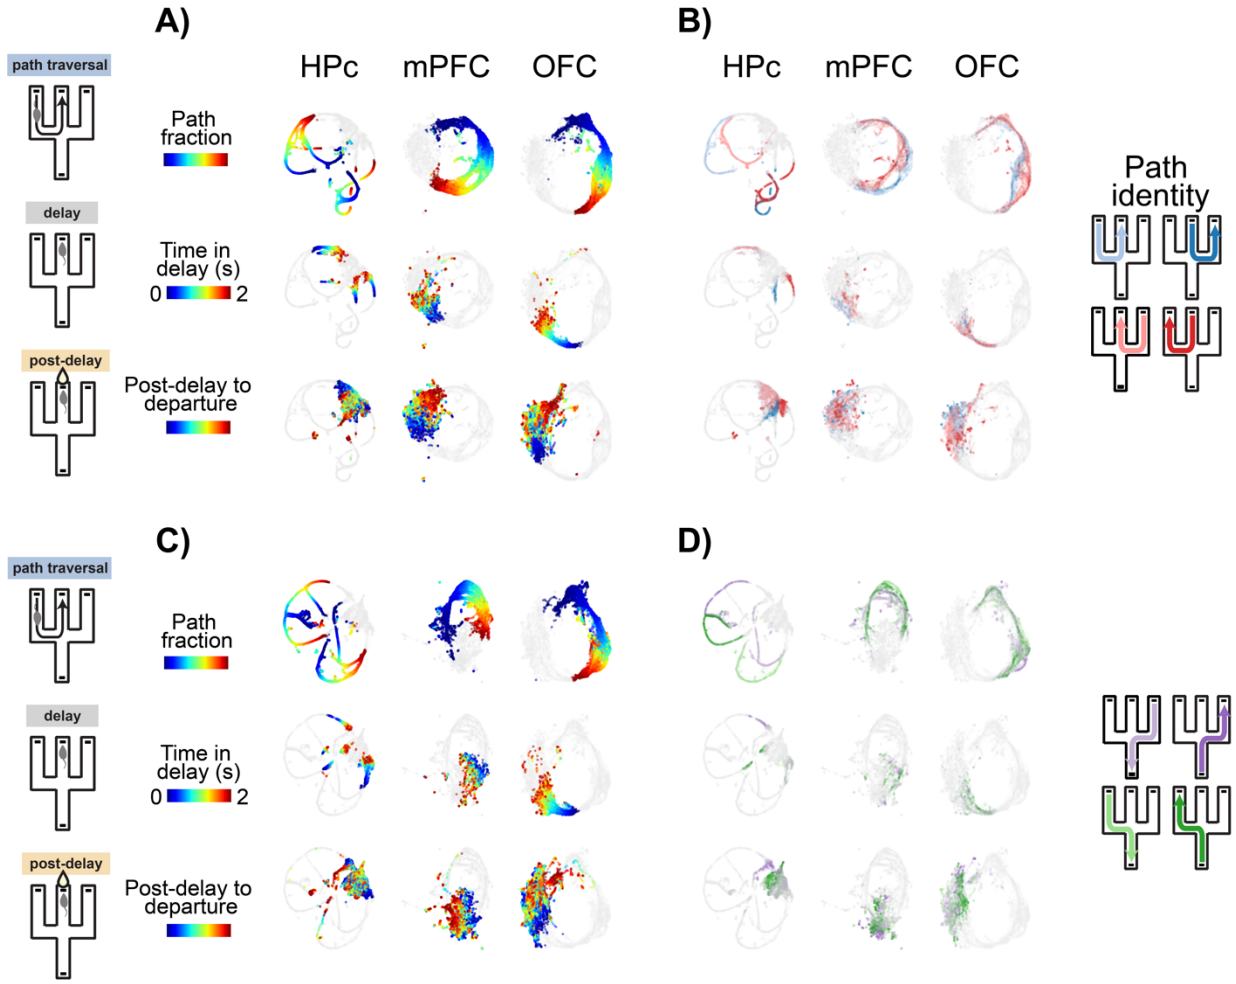


**Figure 2.7: Nonlinear embedding of ensemble activity in distinct environment and**

**sequence contexts.** As in Figure 2.6, UMAP embeddings of HPc, mPFC, and OFC firing rate vectors (left, middle, and right in each panel). Embeddings are of neural data from the same rat as in Figure 2.6, but from a session in a different environment (A–B) or with a different set of rewarded wells i.e., a different sequence (C–D). Embedded neural states are colored according to task progression (A, C) or path identity (B, D) during path traversals, the delay period and the post-delay period (top, middle, and bottom in each panel).

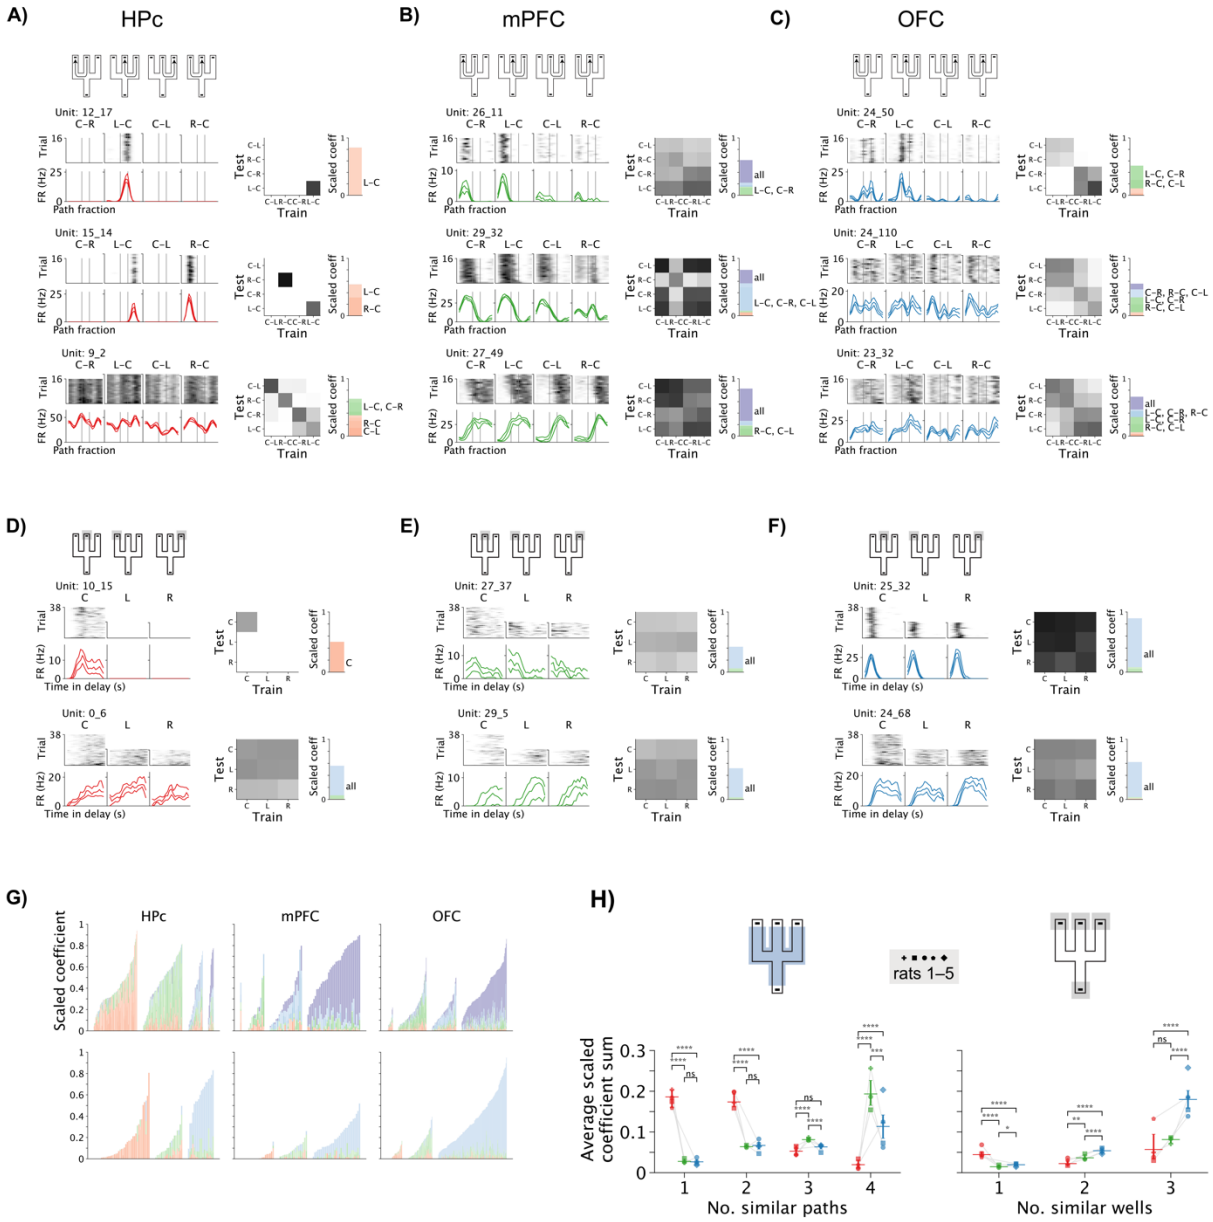


**Figure 2.8: Single-unit firing patterns during task performance**. **A–F:** single-unit firing rate patterns in HPc (A, D), mPFC (B, E), and OFC (C, F), during path traversals

(A–C) and the delay period (D–F). For each single-unit, the following is shown from left to right: first, single trial firing rates (top) and mean firing rates across trials with 25$^{th}$ and 75$^{th}$ percentiles of single trial firing rates (bottom) are shown for each path (A–C) or goal location (D–F) context. Second, the correlations between predicted and actual firing rates for each pair of train and test context are shown in matrix form. Third, scaled coefficients reflecting the estimated contribution of representations at distinct levels of generality are shown. For coefficients >= 0.1, corresponding contexts are abbreviated to the right. Mazes depict path and goal contexts. **G:** Stacked scaled coefficients for single-units from a single behavioral session, grouped by brain region. Each column corresponds to a single-unit. Scaled coefficients at lower levels of generality are lower in the stack. Within brain region, units are further grouped according to the level of generality estimated to contribute the most to the representation. Within this grouping, units are sorted by the magnitude of the scaled coefficient sum. **H:** Average scaled coefficient sum for each level of generality, for each brain region. Horizontal line and error bars show the mean across rats, sessions, and units and corresponding 95% confidence intervals. Markers show the mean across sessions and units for individual rats.

We began at the single-unit level, by inspecting single-unit firing patterns along the journey to goal locations and following arrival at these locations. HPc units tended to fire in a distinct manner along different paths and at different goal locations, consistent with the known location and directional selectivity of HPc neurons[7] and also with the organization around location and heading direction that we observed in the embedded HPc population activity (Figure 2.8A,D). By contrast, mPFC and OFC units tended to fire similarly across several paths or goal locations (Figure 2.8B–C, E–F).

To quantify these observations, we used a GLM-based approach to estimate the extent to which representations at distinct levels of generality contribute to the observed firing patterns during task progression (Figure 2.10). As expected, lower levels of generality contribute more to HPc unit firing patterns, whereas higher levels of generality contribute more to mPFC and OFC firing patterns (Figure 2.8G–H). These biases were present during path traversals as well as following arrival at goals (Figure 2.8G–H), suggesting their conservation across distinct behavioral phases of the task.

These results agree with previous reports demonstrating a relative bias towards context-specific representations in HPc and towards context-generalized representations in mPFC[53,75], and further suggest a relative bias towards context-generalized representations in the OFC. We next considered our results in relation to our first hypothesis that task progression is coded in a hierarchical manner across contexts in the mPFC and OFC. Consistent with our hypothesis, a representation of the journey between any two goal locations contributed the most to mPFC and OFC firing patterns, whereas more context-specific representations of particular kinds of journeys contributed less (Figure 2.8H).

We next considered our second hypothesis, that despite a similar overall organization of activity, mPFC and OFC specialize to express reliable representations of task progression in distinct behavioral periods. To test this hypothesis, we examined the correlations between predicted and actual firing rates during the path traversal and delay periods in these areas. Consistent with our second hypothesis, correlations were higher in mPFC than in OFC during path traversals, whereas they were higher in OFC than in mPFC during the delay period (Figure 2.9).

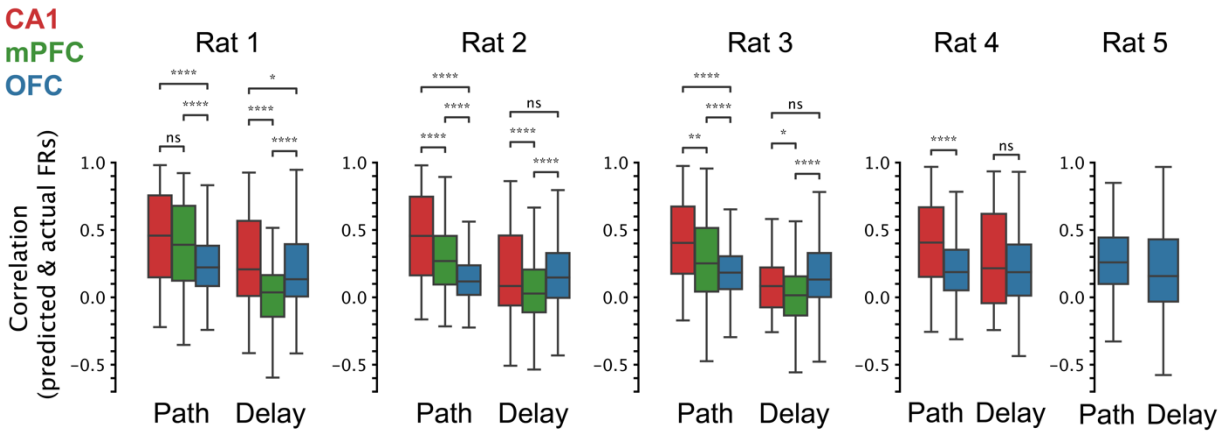


**Figure 2.9: Single-trial reliability of single-unit task progression representations.** Box plots showing the distribution of correlations between actual and GLM-predicted firing rates across trials during path traversals (left in each panel) and the delay period (right in each panel) for single-units in individual contexts (paths for path traversals and goals for the delay period).

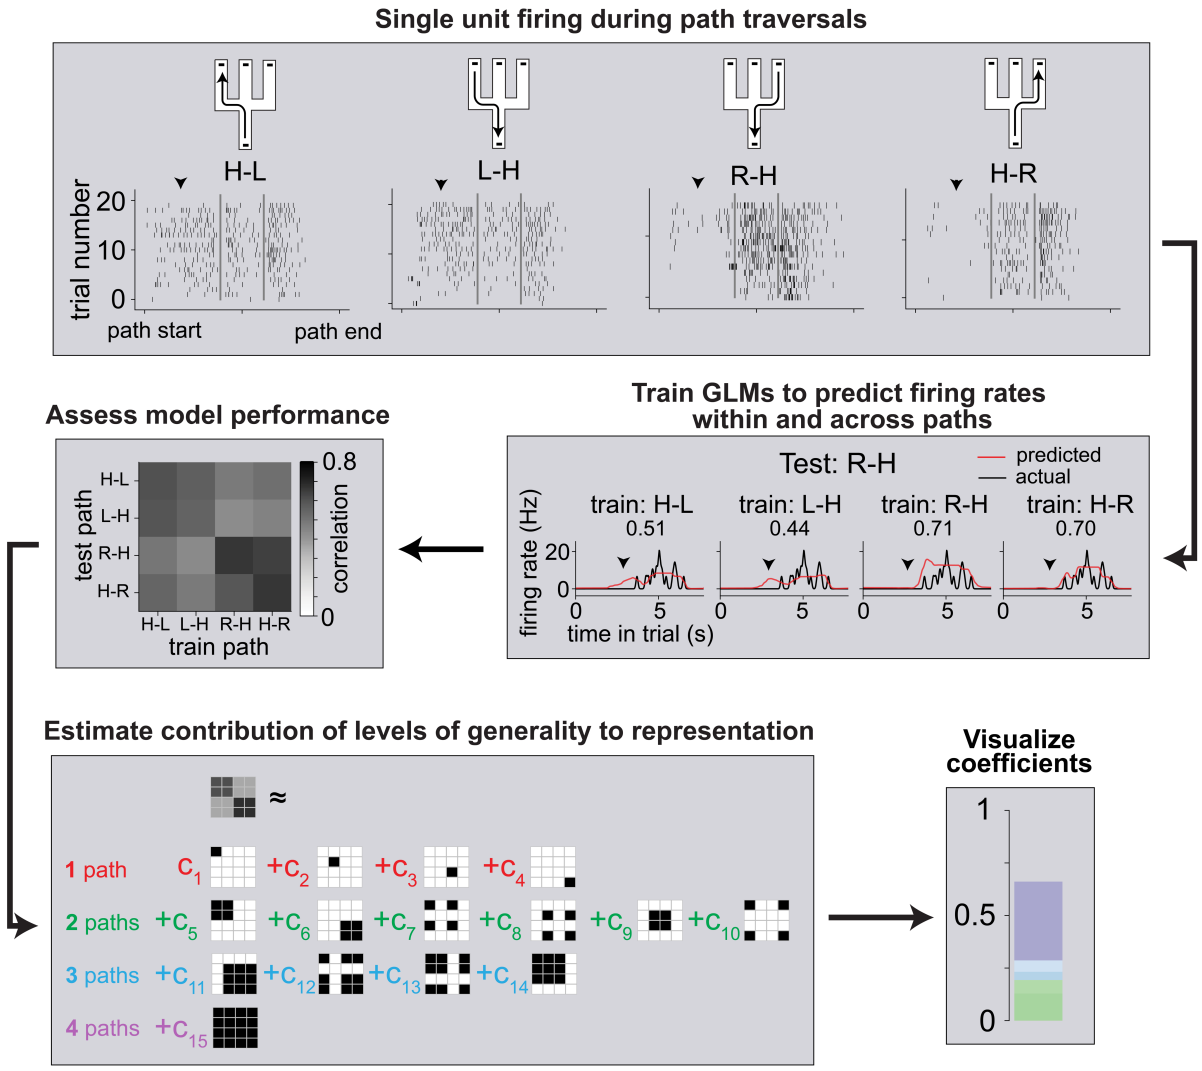


**Figure 2.10: Approach to quantifying levels of generality in single-unit representations of task progression**. Here, we illustrate the process for path traversals, although we performed an analogous process for the delay period. We begin with the spike times during individual trials along the four potentially rewarded paths in a behavioral session (top row). We find spike counts in 20ms bins. Using a Poisson GLM, we model spike counts as a function of fraction of path traversed. We quantify the extent to which models trained on one path can predict firing on the same and other paths in terms of the correlation between predicted and actual firing rates (middle row). We approximate the observed pattern of correlations as a positive linear combination of patterns reflecting representations at distinct levels of generality (bottom row, left). We consider the resulting coefficients an estimate of the contribution of distinct levels of generality to task progression representations (bottom row, right).

We next set out to examine our hypotheses at a population level. To do so, we returned to the original neural state space in which each axis corresponds to the firing rate of a single-unit. We began by examining our first hypothesis that representations of specific kinds of paths, such as those that share the same turn direction, are expressed by mPFC and OFC, in addition to a representation of the journey between any two goal locations. We examined this in terms of changes in neural state, or in other words, neural state dynamics. We asked whether neural state vectors evolve in more similar directions along paths that share the same sets of turns as compared to those that do not. We fit a piece-wise linear curve to the average neural state trajectory during the progression along each path. We quantified the similarity in average neural state dynamics across paths as the cosine similarity between corresponding average neural state difference vectors (Figure 2.11A). In both the mPFC and OFC, average neural state dynamics were more similar along paths that shared turn directions as compared to those that did not, and this difference was most prominent towards the middle of journeys where the turns are present (Figure 2.11B–C, Figure 2.12). However, even along paths that did not share turn directions, neural states proceed in similar directions on average towards the beginning and ends of traversals (Figure 2.11B–C, Figure 2.12). Together, these findings suggest that at the level of average population dynamics, the extent to which more generalized or specific representations of journeys are expressed by the mPFC and OFC changes over the course of journeys. We asked whether the greater differences in neural state dynamics that we observed across paths with different turn directions were likely explainable by differences in speed along these

paths and concluded that this seems unlikely, as the speed profiles were not consistently more dissimilar along these paths relative to those with the same turn directions (Figure 2.13). Our findings were also robust to the use of firing rates that were normalized within each unit (Figure 2.12).

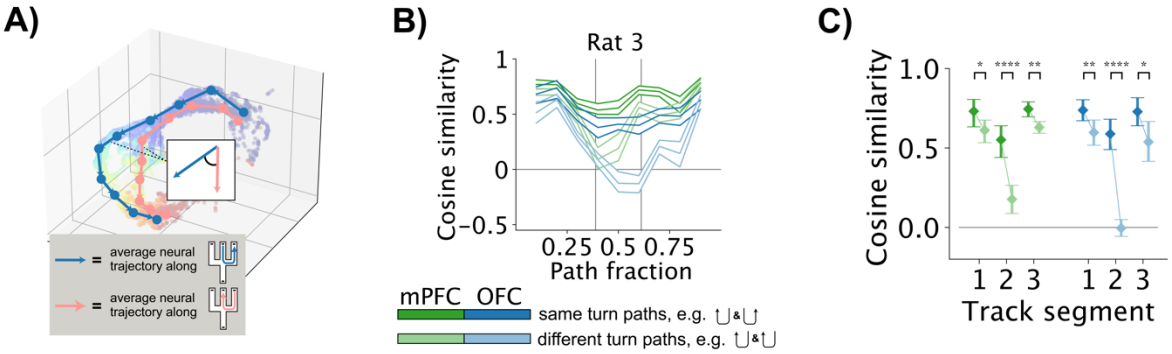


**Figure 2.11: Comparison of average neural state dynamics when similar or different actions are required to reach a goal. A:** Conceptual approach. Dots depict average neural states along two paths. Lines depict average neural state difference vectors. The similarity in neural state evolution across paths is quantified in terms of the cosine similarity of average neural state difference vectors at corresponding positions along paths. Note that although this approach is schematized with three dimensions, we performed the analysis in the full neural state space. **B:** Cosine similarity of average difference vectors along paths with the same (dark colors) or different (light colors) turn directions through different spaces. The mean across path pairs and sessions and associated 95% confidence intervals are shown for one subject. Vertical lines mark the locations of maze junctions. **C:** Cosine similarity of average difference vectors in the first, second, or final third of the track. The mean across path fraction bins in a track segment, path pairs, sessions, and rats and associated 95% confidence intervals are shown.

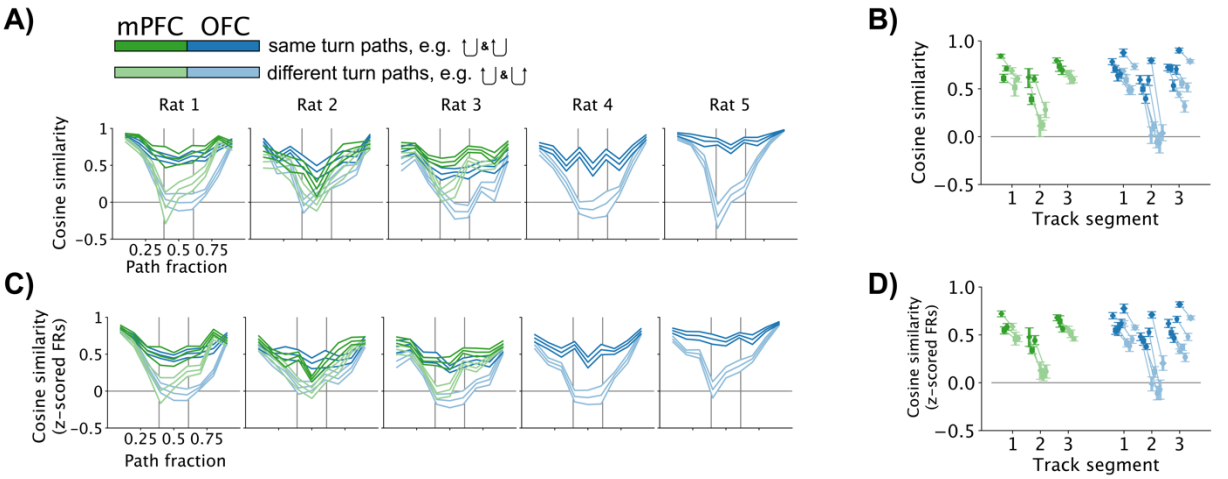


**Figure 2.12: Additional data for Figure 2.11. A, C:** Cosine similarity of average difference vectors formed from non-z-scored (A) or z-scored (C) firing rates along pairs of paths with the same (dark colors) or different (light colors) turn directions through different spaces. The mean across path pairs and sessions and associated 95% confidence intervals are shown for each rat. **B, D:** Cosine similarity of average difference vectors formed from non-z-scored (B) or z-scored (D) firing rates in the first, second, and final third of the track. The mean across path pairs and sessions and associated 95% confidence intervals are shown for each rat.

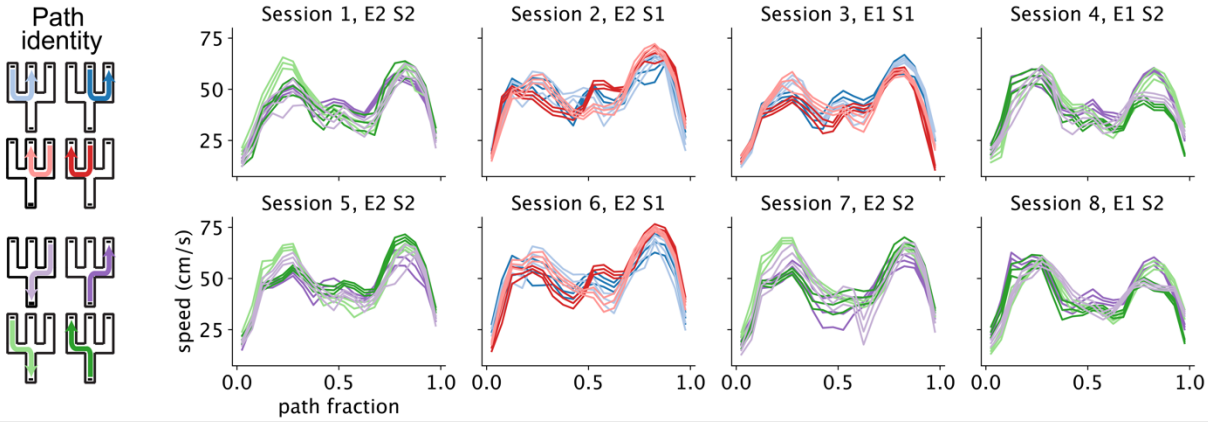


**Figure 2.13: Speed during traversals of different paths.** Mean head speed and associated 95% confidence intervals along potentially rewarded paths in each behavioral session of a day from the rat whose data is shown in Figure 2.11B.

We next examined our second hypothesis, that mPFC and OFC express reliable task progression representations at distinct behavioral phases of the task. To this end,

we quantified the consistency in neural state dynamics and geometry during task progression across single trials. Beginning with dynamics, we computed the cosine similarity between single trial neural state difference vectors spanning the progression along paths or in the delay period (Figure 2.14A). Consistent with our hypothesis, neural state dynamics were more consistent in the mPFC during path traversals, and in the OFC following arrival at goal locations (Figure 2.14B).

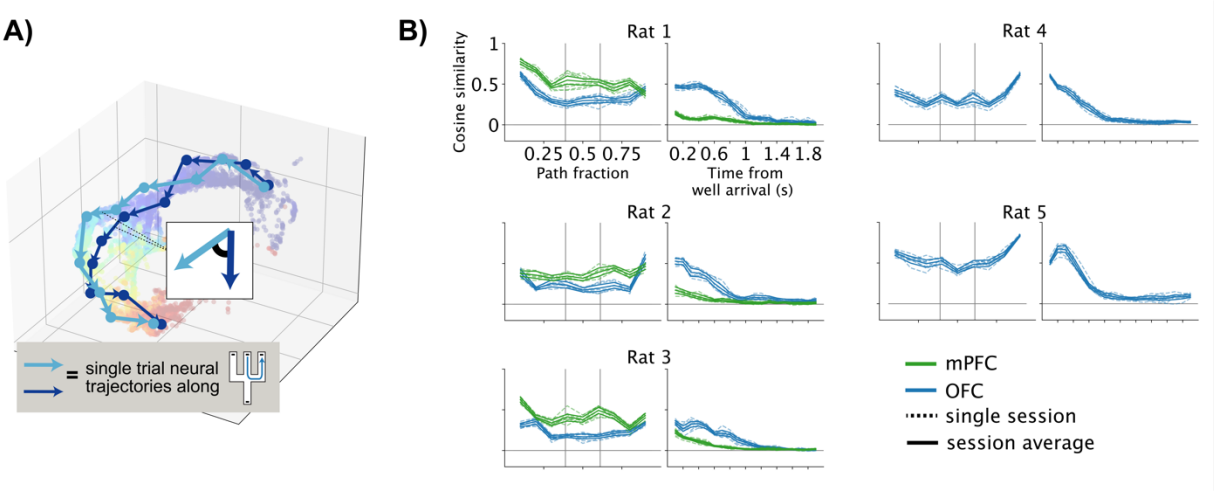


**Figure 2.14: Single trial reliability of neural state dynamics during task progression. A:** Conceptual approach illustrated in the case of path traversals. Blue dots depict neural states on different trials along the same path. Lines depict neural state difference vectors. Consistency of neural state dynamics across trials is quantified in terms of the cosine similarity of corresponding neural state difference vectors on pairs of trials. Note that although this approach is schematized with three dimensions, we performed the analysis in the full neural state space. **B:** Mean cosine similarity of single trial difference vectors during path traversals (left) or during the 2s delay period (right) for each subject. Vertical lines mark the locations of maze junctions.

We next assessed the single-trial consistency in the geometry of neural states during task progression. If neural activity patterns unfold in a reliable manner across single trials, neural states should be closer together at similar positions in the task relative to at different positions. To determine if this was the case, we compared the distances of neural states in the same task progression bin to the distances of neural

states in different task progression bins (Figure 2.15A). In both mPFC and OFC, neural states expressed at similar points in the task progression were closer together than those occurring at different locations (Figure 2.15B–C). However, this was more so the case in the mPFC relative to the OFC during path traversals, and in the OFC relative to the mPFC during the delay period (Figure 2.15B–C). These distinct patterns of reliability of neural state geometry in the mPFC and OFC are consistent with our second hypothesis that these areas express reliable representations of task progression at distinct behavioral phases. Overall, our population-level findings suggest that at both the level of neural state geometry and dynamics, there is a functional division of labor within the PFC in the reliable expression of representations of task progression during distinct behavioral periods.

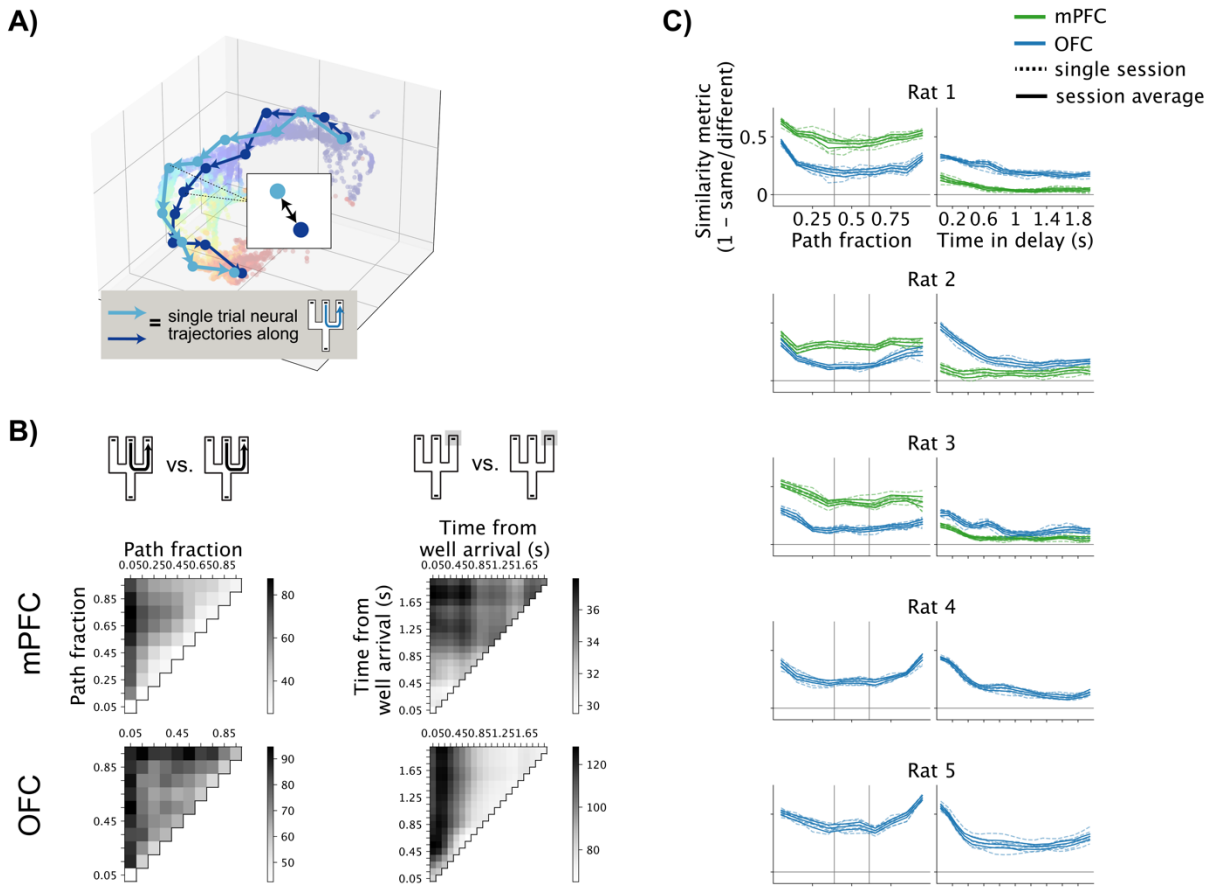


**Figure 2.15: Single trial reliability of neural state geometry during task progression**. **A:** Conceptual approach. As in 2.14A, except consistency of geometry is quantified. The quantification is in terms of the relative closeness of neural states within the same versus in different task progression bins. **B:** Average Euclidean distance between pairs of firing rate vectors for different pairs of task progression bins, during either traversals of the center to right path (left) or during the ensuing delay period at the right well (right) for mPFC (top) and OFC (bottom) neural states from a subject during one behavioral session. **C:** Similarity metric (one minus the ratio of firing rate vector distances within the same and different task progression bins) as a function of task progression during the path traversal (left) or delay period (right) for each subject. Vertical lines mark the locations of maze junctions.

We next sought to understand how the differential reliability of task progression representations in the mPFC and OFC during path traversals and the delay period might relate to how well task progression can be decoded from these areas during these distinct periods. To address this, we decoded task progression along paths and

during the delay period from the spiking activity in each area[104]. Progression along paths could be decoded more reliably from mPFC ensemble spikes, relative to spikes from an equally sized OFC ensemble, whereas the opposite was seen during the delay period (Figure 2.16). These results provide further support for our second hypothesis that the mPFC and OFC exhibit a dissociation in when during the task they express reliable information about task progression.

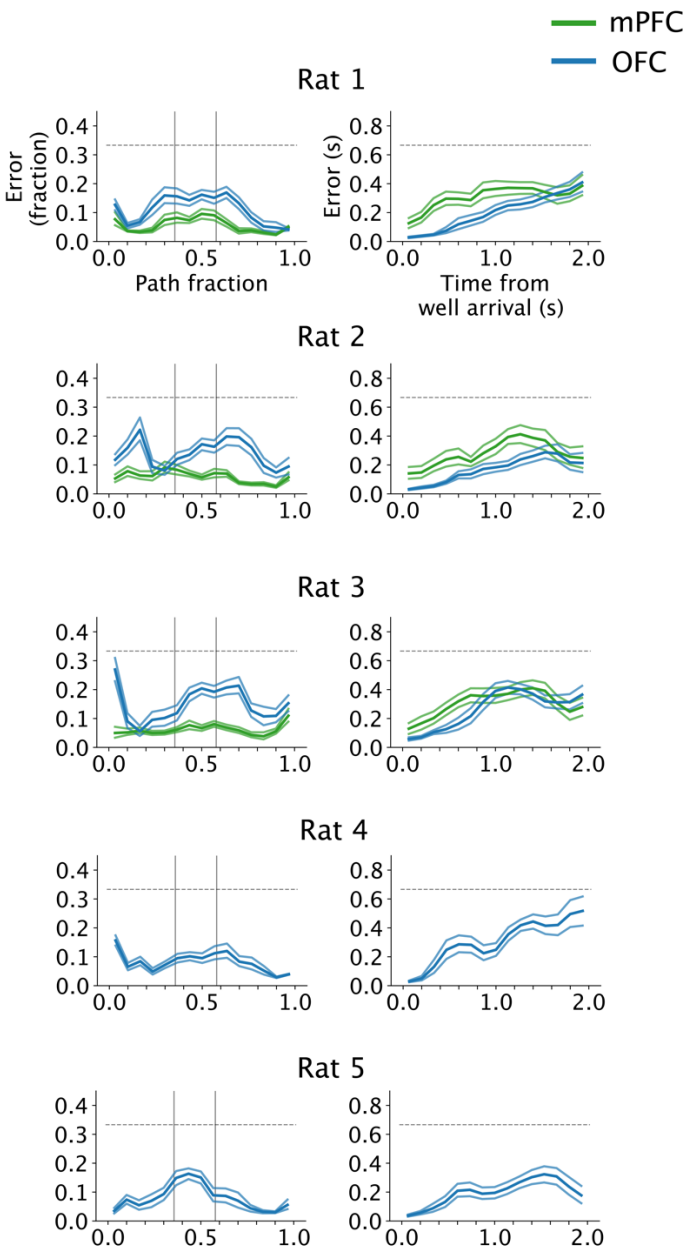


**Figure 2.16: Task progression decoding during distinct task phases.** Progression along paths (left) and in the delay period (right) was decoded from the spiking activity of a randomly selected 50 units in the mPFC and OFC each. **A:** Mean decoding error and associated 95% confidence intervals as a function of task progression during path traversals (left) or the delay period (right) during a representative behavioral session from each rat. Vertical lines mark the locations of maze junctions. Horizontal line marks an estimate of chance.

## Chapter 3

## Conclusions and Implications

The HPc, mPFC, and OFC are each thought to represent generalized features of experience, providing a "cognitive map" for guiding behavior. Generalized representations have been observed in each of these areas[19,75,109,110], with a relative bias towards more context-generalized representations in the PFC as compared to the HPc[53,75]. Here, we asked whether there are organizational principles that govern how generalized representations are distributed across these areas. To do so, we examined single-unit and population-level activity patterns across HPc, mPFC, and OFC as rats performed a spatial navigation task. As rats traversed paths, HPc activity signaled progression along particular spatial routes, consistent with decades of previous research on place cells[7]. By contrast, mPFC and OFC activity was chiefly organized around the journey between any two goal locations. Furthermore, activity in these areas was more similar across journeys requiring similar as opposed to different actions, suggestive of a concurrent representation of the set of actions required to reach goals. Our results are consistent with previous reports demonstrating that activity in the mPFC tends to generalize across contexts by way of representing distinct spatial routes as similar[53,54,111]. Our results further identify the OFC with this property. The turn-based representations that we observed in the OFC are consistent with a previous report describing prominent turn-related activity in the OFC[38]. Our results further identify the mPFC with turn-based representation of journeys.

Our findings also provide new insight into the extent to which mPFC and OFC represent progression through the task at distinct levels of generality. Using a GLM-

based approach, we estimated the contribution of representations of task progression at distinct levels of generality to single-unit firing during the journey to goals and in the delay period that followed. Our findings suggest that across the entire extent of journeys along paths, single cell firing patterns tend to reflect a generalized representation of the journey between any two goal locations, and, to a lesser extent, more specific representations of the journey between particular subsets of goals. We also examined the average evolution of population activity along journeys in mPFC and OFC, with the goal of understanding how population activity dynamics compared along journeys requiring the same or different turn directions. We found that dynamics were more similar across paths with different directions towards the beginning and end of journeys, and relatively more distinct towards the middle of journeys where turns are present, suggesting a dynamic modulation of journey representations at distinct levels of generality.

Although we found that mPFC and OFC exhibited similar kinds of generalized representation during the task, we further identified evidence that these areas specialize to express reliable representations of task progression at distinct behavioral phases. In comparison to OFC, mPFC activity more reliably represented progression along paths between goal locations. By contrast, in comparison to mPFC, OFC activity more reliably represented progression in the delay period following arrival at a goal location. We hypothesize that these distinct patterns of consistency in representation reflect the extent to which these areas are engaged in computations during these respective behavioral periods. If correct, this could help explain why lesions to these areas produce

distinct behavioral deficits. In rats, OFC lesions bias subjects towards small, quickly available rewards over much larger, delayed rewards but do not impair effort-based decision making, whereas the opposite pattern of deficits is seen with ACC lesions[82]. OFC lesions have also been found to reduce the extent to which rats modulate how long they will wait for rewards depending on task difficulty[112], consistent with a role for the OFC in representing decision confidence[113]. In our study, relative to the mPFC, the OFC expressed more consistent representations of task progression after rats arrived at goals and were faced with a choice of whether or not to wait for a potential upcoming reward. We hypothesize that this may relate to a specialized role for OFC in processing during this period and bears on rats' decision to stay at or leave a goal location.

More broadly, our findings raise the possibility that trial-by-trial consistency in activity patterns underlies regional specialization, a principle that could help reconcile why several task variables are coded in several brain areas, yet lesions to different areas can produce distinct behavioral deficits. It may be that the consistency of representations during a behavioral period, rather than their existence per se, is reflective of an area's involvement in computations essential to that period.

## References


1. Skinner, B. F. *Science and human behavior.* x, 461 (Macmillan, 1953).
2. Bartlett, F. C. Remembering: A study in experimental and social psychology. *Remembering: A study in experimental and social psychology.* xix, 317–xix, 317 (1932).
3. Tolman, E. C. Cognitive maps in rats and men. *Psychological Review* **55**, 189–208 (1948).
4. Piaget, J. *The language and thought of the child.* xxiii, 246 (Harcourt, Brace, 1926).
5. O'Keefe, J. & Dostrovsky, J. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research* **34**, 171–175 (1971).
6. O'Keefe, J. M., Nadel, L. & O'Keefe, J. *The hippocampus as a cognitive map*. (Clarendon Press, 1978).
7. McNaughton, B. L., Barnes, C. A. & O'Keefe, J. The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Exp Brain Res* **52**, (1983).
8. Markus, E. *et al.* Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *J. Neurosci.* **15**, 7079–7094 (1995).
9. Frank, L. M., Brown, E. N. & Wilson, M. Trajectory Encoding in the Hippocampus and Entorhinal Cortex. *Neuron* **27**, 169–178 (2000).

10. Wood, E. R., Dudchenko, P. A., Robitsek, R. J. & Eichenbaum, H. Hippocampal Neurons Encode Information about Different Types of Memory Episodes Occurring in the Same Location. *Neuron* **27**, 623–633 (2000).
11. Cohen, N. J. & Eichenbaum, H. *Memory, amnesia, and the hippocampal system.* xii, 330 (The MIT Press, 1993).
12. Eichenbaum, H. The hippocampus and mechanisms of declarative memory. *Behavioural Brain Research* **103**, 123–133 (1999).
13. Eichenbaum, H. & Cohen, N. J. Can We Reconcile the Declarative Memory and Spatial Navigation Views on Hippocampal Function? *Neuron* **83**, 764–770 (2014).
14. Whittington, J. C. R. *et al.* The Tolman-Eichenbaum Machine: Unifying Space and Relational Memory through Generalization in the Hippocampal Formation. *Cell* **183**, 1249-1263.e23 (2020).
15. Bunsey, M. & Eichenbaum, H. Conservation of hippocampal memory function in rats and humans. *Nature* **379**, 255–257 (1996).
16. Dusek, J. A. & Eichenbaum, H. The hippocampus and memory for orderly stimulus relations. *Proc. Natl. Acad. Sci. U.S.A.* **94**, 7109–7114 (1997).
17. Miller, K. J., Botvinick, M. M. & Brody, C. D. Dorsal hippocampus contributes to model-based planning. *Nat Neurosci* **20**, 1269–1276 (2017).
18. Aronov, D., Nevers, R. & Tank, D. W. Mapping of a non-spatial dimension by the hippocampal–entorhinal circuit. *Nature* **543**, 719–722 (2017).
19. Baraduc, P., Duhamel, J.-R. & Wirth, S. Schema cells in the macaque hippocampus. *Science* **363**, 635–639 (2019).

20. Bernardi, S. *et al.* The Geometry of Abstraction in the Hippocampus and Prefrontal Cortex. *Cell* **183**, 954-967.e21 (2020).
21. Fuster, J. M. The prefrontal cortex--an update: time is of the essence. *Neuron* **30**, 319–333 (2001).
22. Miller, E. K. The prefontral cortex and cognitive control. *Nat Rev Neurosci* **1**, 59–65 (2000).
23. Miller, E. K. & Cohen, J. D. An Integrative Theory of Prefrontal Cortex Function. *Annu. Rev. Neurosci.* **24**, 167–202 (2001).
24. Wilson, R. C., Takahashi, Y. K., Schoenbaum, G. & Niv, Y. Orbitofrontal cortex as a cognitive map of task space. *Neuron* **81**, 267–279 (2014).
25. Tse, D. *et al.* Schema-dependent gene activation and memory encoding in neocortex. *Science* **333**, 891–895 (2011).
26. Tse, D. *et al.* Schemas and memory consolidation. *Science* **316**, 76–82 (2007).
27. Takahashi, Y. K. *et al.* Expectancy-related changes in firing of dopamine neurons depend on orbitofrontal cortex. *Nat Neurosci* **14**, 1590–1597 (2011).
28. Diamond, A. Executive Functions. *Annu. Rev. Psychol.* **64**, 135–168 (2013).
29. Cummings, J. L. Frontal-Subcortical Circuits and Human Behavior. *Archives of Neurology* **50**, 873–880 (1993).
30. Milner, B. Effects of Different Brain Lesions on Card Sorting: The Role of the Frontal Lobes. *Arch Neurol* **9**, 90 (1963).

31. Nemeth, G., Hegedus, K. & Molnar, L. Akinetic mutism associated with bicingular lesions: Clinicopathological and functional anatomical correlates. *Eur Arch Psychiatr Neurol Sci* **237**, 218–222 (1988).

32. Darby, R. R., Joutsa, J., Burke, M. J. & Fox, M. D. Lesion network localization of free will. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 10792–10797 (2018).

33. Arnts, H. *et al.* On the pathophysiology and treatment of akinetic mutism. *Neuroscience & Biobehavioral Reviews* **112**, 270–278 (2020).

34. Harlow, J. M. & Miller, E. Recovery from the passage of an iron bar through the head. *Hist Psychiatry* **4**, 271–273 (1993).

35. Bechara, A., Damasio, H., Damasio, A. R. & Lee, G. P. Different Contributions of the Human Amygdala and Ventromedial Prefrontal Cortex to Decision-Making. *J. Neurosci.* **19**, 5473–5481 (1999).

36. Kennerley, S. W. & Wallis, J. D. Evaluating choices by single neurons in the frontal lobe: outcome value encoded across multiple decision variables. *Eur J Neurosci* **29**, 2061–2073 (2009).

37. Kennerley, S. W., Dahmubed, A. F., Lara, A. H. & Wallis, J. D. Neurons in the frontal lobe encode the value of multiple decision variables. *J Cogn Neurosci* **21**, 1162–1178 (2009).

38. Feierstein, C. E., Quirk, M. C., Uchida, N., Sosulski, D. L. & Mainen, Z. F. Representation of Spatial Goals in Rat Orbitofrontal Cortex. *Neuron* **51**, 495–507 (2006).

39. Schoenbaum, G. & Eichenbaum, H. Information coding in the rodent prefrontal cortex. I. Single-neuron activity in orbitofrontal cortex compared with that in pyriform cortex. *Journal of Neurophysiology* **74**, 733–750 (1995).

40. Zhou, J. *et al.* Complementary Task Structure Representations in Hippocampus and Orbitofrontal Cortex during an Odor Sequence Task. *Current Biology* **29**, 3402-3409.e3 (2019).

41. Schoenbaum, G., Chiba, A. A. & Gallagher, M. Orbitofrontal cortex and basolateral amygdala encode expected outcomes during learning. *Nat Neurosci* **1**, 155–159 (1998).

42. Rigotti, M. *et al.* The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**, 585–590 (2013).

43. Fuster, J. M. & Alexander, G. E. Neuron Activity Related to Short-Term Memory. *Science* **173**, 652–654 (1971).

44. Fuster, J. M. Unit activity in prefrontal cortex during delayed-response performance: neuronal correlates of transient memory. *Journal of Neurophysiology* **36**, 61–78 (1973).

45. Fuster, J. M., Bodner, M. & Kroger, J. K. Cross-modal and cross-temporal association in neurons of frontal cortex. *Nature* **405**, 347–351 (2000).

46. Wallis, J. D., Anderson, K. C. & Miller, E. K. Single neurons in prefrontal cortex encode abstract rules. *Nature* **411**, 953–956 (2001).

47. Marr, D. A Theory for Cerebral Neocortex. *Proceedings of the Royal Society of London. Series B, Biological Sciences* **176**, 161–234 (1970).

48. McClelland, J. L., McNaughton, B. L. & O'Reilly, R. C. 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 (1995).

49. Born, J. & Wilhelm, I. System consolidation of memory during sleep. *Psychological Research* **76**, 192–203 (2012).

50. Takehara, K., Kawahara, S. & Kirino, Y. Time-Dependent Reorganization of the Brain Components Underlying Memory Retention in Trace Eyeblink Conditioning. *J. Neurosci.* **23**, 9897–9905 (2003).

51. Maviel, T., Durkin, T. P., Menzaghi, F. & Bontempi, B. Sites of Neocortical Reorganization Critical for Remote Spatial Memory. *Science* **305**, 96–99 (2004).

52. Takehara-Nishiuchi, K. & McNaughton, B. L. Spontaneous Changes of Neocortical Code for Associative Memory During Consolidation. *Science* **322**, 960–963 (2008).

53. Yu, J. Y., Liu, D. F., Loback, A., Grossrubatscher, I. & Frank, L. M. Specific hippocampal representations are linked to generalized cortical representations in memory. *Nature Communications* **9**, 2209 (2018).

54. Jung, M. W., Qin, Y., McNaughton, B. L. & Barnes, C. A. Firing characteristics of deep layer neurons in prefrontal cortex in rats performing spatial working memory tasks. *Cerebral Cortex* **8**, 437–450 (1998).

55. Wallis, J. D. Cross-species studies of orbitofrontal cortex and value-based decision-making. *Nat Neurosci* **15**, 13–19 (2012).

56. Rogers, R. Dissociable Deficits in the Decision-Making Cognition of Chronic Amphetamine Abusers, Opiate Abusers, Patients with Focal Damage to Prefrontal Cortex, and Tryptophan-Depleted Normal Volunteers Evidence for Monoaminergic Mechanisms. *Neuropsychopharmacology* **20**, 322–339 (1999).

57. Bechara, A., Damasio, A. R., Damasio, H. & Anderson, S. W. Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition* **50**, 7–15 (1994).

58. Bechara, A., Damasio, H., Tranel, D. & Anderson, S. W. Dissociation Of Working Memory from Decision Making within the Human Prefrontal Cortex. *J. Neurosci.* **18**, 428–437 (1998).

59. Bechara, A., Tranel, D., Damasio, H. & Damasio, A. R. Failure to Respond Autonomically to Anticipated Future Outcomes Following Damage to Prefrontal Cortex. *Cereb Cortex* **6**, 215–225 (1996).

60. Bechara, A., Damasio, H., Tranel, D. & Damasio, A. R. Deciding Advantageously Before Knowing the Advantageous Strategy. *Science* **275**, 1293–1295 (1997).

61. Butter, C. M. Perseveration in extinction and in discrimination reversal tasks following selective frontal ablations in Macaca mulatta. *Physiology & Behavior* **4**, 163–171 (1969).

62. Schoenbaum, G., Setlow, B., Nugent, S. L., Saddoris, M. P. & Gallagher, M. Lesions of Orbitofrontal Cortex and Basolateral Amygdala Complex Disrupt Acquisition of Odor-Guided Discriminations and Reversals. *Learn. Mem.* **10**, 129–140 (2003).

63. Schoenbaum, G., Nugent, S. L., Saddoris, M. P. & Setlow, B. Orbitofrontal lesions in rats impair reversal but not acquisition of go, no-go odor discriminations: *Neuroreport* **13**, 885–890 (2002).

64. Izquierdo, A. Bilateral Orbital Prefrontal Cortex Lesions in Rhesus Monkeys Disrupt Choices Guided by Both Reward Value and Reward Contingency. *Journal of Neuroscience* **24**, 7540–7548 (2004).

65. Jones, B. & Mishkin, M. Limbic lesions and the problem of stimulus—Reinforcement associations. *Experimental Neurology* **36**, 362–377 (1972).

66. Padoa-Schioppa, C. & Assad, J. A. Neurons in the orbitofrontal cortex encode economic value. *Nature* **441**, 223–226 (2006).

67. Gallagher, M., McMahan, R. W. & Schoenbaum, G. Orbitofrontal Cortex and Representation of Incentive Value in Associative Learning. *J. Neurosci.* **19**, 6610–6614 (1999).

68. Tremblay, L. & Schultz, W. Relative reward preference in primate orbitofrontal cortex. *Nature* **398**, 704–708 (1999).

69. Thorpe, S. J., Rolls, E. T. & Maddison, S. The orbitofrontal cortex: Neuronal activity in the behaving monkey. *Exp Brain Res* **49**, (1983).

70. Knudsen, E. B. & Wallis, J. D. Taking stock of value in the orbitofrontal cortex. *Nat Rev Neurosci* **23**, 428–438 (2022).

71. Jones, J. L. *et al.* Orbitofrontal Cortex Supports Behavior and Learning Using Inferred But Not Cached Values. *Science* **338**, 953–956 (2012).

72. Schuck, N. W., Cai, M. B., Wilson, R. C. & Niv, Y. Human Orbitofrontal Cortex Represents a Cognitive Map of State Space. *Neuron* **91**, 1402–1412 (2016).

73. Zhou, J. *et al.* Evolving schema representations in orbitofrontal ensembles during learning. *Nature* **590**, 606–611 (2021).

74. Ramus, S. J. & Eichenbaum, H. Neural Correlates of Olfactory Recognition Memory in the Rat Orbitofrontal Cortex. *Journal of Neuroscience* **20**, 8199–8208 (2000).

75. Samborska, V., Butler, J. L., Walton, M. E., Behrens, T. E. J. & Akam, T. Complementary task representations in hippocampus and prefrontal cortex for generalizing the structure of problems. *Nat Neurosci* **25**, 1314–1326 (2022).

76. Marr, D. Simple memory: a theory for archicortex. *Phil. Trans. R. Soc. Lond. B* **262**, 23–81 (1971).

77. Bontempi, B., Laurent-Demir, C., Destrade, C. & Jaffard, R. Time-dependent reorganization of brain circuitry underlying long-term memory storage. *Nature* **400**, 671–675 (1999).

78. McKenzie, S. *et al.* Hippocampal Representation of Related and Opposing Memories Develop within Distinct, Hierarchically Organized Neural Schemas. *Neuron* **83**, 202–215 (2014).

79. Farovik, A. *et al.* Orbitofrontal Cortex Encodes Memories within Value-Based Schemas and Represents Contexts That Guide Memory Retrieval. *Journal of Neuroscience* **35**, 8333–8344 (2015).

80. Wikenheiser, A. M. & Schoenbaum, G. Over the river, through the woods: cognitive maps in the hippocampus and orbitofrontal cortex. *Nat Rev Neurosci* **17**, 513–523 (2016).

81. Rudebeck, P. H. *et al.* Frontal Cortex Subregions Play Distinct Roles in Choices between Actions and Stimuli. *Journal of Neuroscience* **28**, 13775–13785 (2008).

82. Rudebeck, P. H., Walton, M. E., Smyth, A. N., Bannerman, D. M. & Rushworth, M. F. S. Separate neural pathways process different decision costs. *Nat Neurosci* **9**, 1161–1168 (2006).

83. Buckley, M. J. *et al.* Dissociable Components of Rule-Guided Behavior Depend on Distinct Medial and Prefrontal Regions. *Science* **325**, 52–58 (2009).

84. Sul, J. H., Kim, H., Huh, N., Lee, D. & Jung, M. W. Distinct Roles of Rodent Orbitofrontal and Medial Prefrontal Cortex in Decision Making. *Neuron* **66**, 449–460 (2010).

85. Kennerley, S. W., Behrens, T. E. J. & Wallis, J. D. Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nat Neurosci* **14**, 1581–1589 (2011).

86. Jadhav, S. P., Rothschild, G., Roumis, D. K. & Frank, L. M. Coordinated Excitation and Inhibition of Prefrontal Ensembles during Awake Hippocampal Sharp-Wave Ripple Events. *Neuron* **90**, 113–127 (2016).

87. Roumis, D. https://figshare.com/articles/software/physiology/4747438/2.

88. Chung, J. E. *et al.* Chronic Implantation of Multiple Flexible Polymer Electrode Arrays. *Journal of visualized experiments : JoVE* **152**, (2019).

89. Sullivan, D. *et al.* Relationships between Hippocampal Sharp Waves, Ripples, and Fast Gamma Oscillation: Influence of Dentate and Entorhinal Cortical Activity. *Journal of Neuroscience* **31**, 8605–8616 (2011).

90. Paxinos, G. & Watson, C. *The Rat Brain in Stereotaxic Coordinates*. (Elsevier, 2007).

91. Chung, J. E. *et al.* A Fully Automated Approach to Spike Sorting. *Neuron* **95**, 1381-1394.e6 (2017).

92. Buccino, A. P. *et al.* SpikeInterface, a unified framework for spike sorting. *eLife* **9**, e61834 (2020).

93. Battaglia, F. P., Sutherland, G. R. & McNaughton, B. L. Local sensory cues and place cell directionality: additional evidence of prospective coding in the hippocampus. *J Neurosci* **24**, 4541–4550 (2004).

94. Yatsenko, D. *et al. DataJoint: managing big scientific data using MATLAB or Python*. http://biorxiv.org/lookup/doi/10.1101/031658 (2015) doi:10.1101/031658.

95. Saravanan, V., Berman, G. J. & Sober, S. J. Application of the hierarchical bootstrap to multi-level data in neuroscience. *Neuron Behav Data Anal Theory* **3**, (2020).

96. McInnes, L., Healy, J. & Melville, J. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. Preprint at http://arxiv.org/abs/1802.03426 (2020).

97. Truccolo, W., Eden, U. T., Fellows, M. R., Donoghue, J. P. & Brown, E. N. A Point Process Framework for Relating Neural Spiking Activity to Spiking History, Neural

Ensemble, and Extrinsic Covariate Effects. *Journal of Neurophysiology* **93**, 1074–1089 (2005).

98. Pillow, J. W. *et al.* Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* **454**, 995–999 (2008).
99. Park, I. M., Meister, M. L. R., Huk, A. C. & Pillow, J. W. Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nat Neurosci* **17**, 1395–1403 (2014).
100. McCullagh, P. & Nelder, J. A. *Generalized Linear Models*. (Routledge, 2019). doi:10.1201/9780203753736.
101. Seabold, S. & Perktold, J. Statsmodels: Econometric and Statistical Modeling with Python. *Proceedings of the 9th Python in Science Conference* **2010**, (2010).
102. 4. The BFGS Method. in *Iterative Methods for Optimization* 71–86 doi:10.1137/1.9781611970920.ch4.
103. Chaudhuri, R., Gerçek, B., Pandey, B., Peyrache, A. & Fiete, I. The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nat Neurosci* **22**, 1512–1520 (2019).
104. Denovellis, E. L. *et al.* Hippocampal replay of experience at real-world speeds. *eLife* **10**, e64505 (2021).
105. Kato, S. *et al.* Global Brain Dynamics Embed the Motor Command Sequence of Caenorhabditis elegans. *Cell* **163**, 656–669 (2015).
106. Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51–56 (2012).

107. Yu, J. Y. *et al.* Distinct hippocampal-cortical memory representations for experiences associated with movement versus immobility. *eLife* **6**, e27621 (2017).

108. Kay, K. *et al.* A hippocampal network for spatial coding during immobility and sleep. *Nature* **531**, 185–190 (2016).

109. Singer, A. C., Karlsson, M. P., Nathe, A. R., Carr, M. F. & Frank, L. M. Experience-Dependent Development of Coordinated Hippocampal Spatial Activity Representing the Similarity of Related Locations. *J. Neurosci.* **30**, 11586–11604 (2010).

110. Basu, R. *et al.* The orbitofrontal cortex maps future navigational goals. *Nature* **599**, 449–452 (2021).

111. Rubin, A. *et al.* Revealing neural correlates of behavior without behavioral measurements. *Nat Commun* **10**, 4745 (2019).

112. Lak, A. *et al.* Orbitofrontal Cortex Is Required for Optimal Waiting Based on Decision Confidence. *Neuron* **84**, 190–201 (2014).

113. Kepecs, A., Uchida, N., Zariwala, H. A. & Mainen, Z. F. Neural correlates, computation and behavioural impact of decision confidence. *Nature* **455**, 227–231 (2008).

**Publishing Agreement**

It is the policy of the University to encourage open access and broad distribution of all theses, dissertations, and manuscripts. The Graduate Division will facilitate the distribution of UCSF theses, dissertations, and manuscripts to the UCSF Library for open access and distribution. UCSF will make such theses, dissertations, and manuscripts accessible to the public and will take reasonable steps to preserve these works in perpetuity.

I hereby grant the non-exclusive, perpetual right to The Regents of the University of California to reproduce, publicly display, distribute, preserve, and publish copies of my thesis, dissertation, or manuscript in any form or media, now existing or later derived, including access online for teaching, research, and public service purposes.

DocuSigned by:
Jennifer Ann Guidera
B2AC621A7F56426...
Author Signature

3/21/2023
Date