

Neural maps of world and behaviour in frontal-entorhinal circuits



Adam Loyd Harris

Department of Experimental Psychology

Merton College


University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2025

This work is licensed under a [Creative Commons “Attribution-NonCommercial 4.0 International”](https://creativecommons.org/licenses/by-nc/4.0/) license.



Preferred citation: Harris, AL. *Neural maps of world and behaviour in frontal-entorhinal circuits*. (2025) DPhil Thesis. University of Oxford.

Abstract

Intelligent, flexible behaviour requires an understanding of both the structure of the world around us and the structure of the problem we are trying to solve. Decades of research have made great progress in understanding how the brain maps the structure of physical space, but the neural mechanisms for mapping the structure of behaviour remain relatively enigmatic. In this thesis, we uncover a neural implementation of an algorithm for mapping abstract behavioural structure and transferring it to new scenarios.

We develop a behavioural paradigm in which mice learn many tasks that shared a common structure (organising a sequence of goals) but differed in the specific goal locations. The spatial locations of the rewards change across tasks, cleanly decoupling physical and task space distances. Mice discover the underlying task structure: they learn new task instances more efficiently over time and, critically, infer unseen transitions on the very first trial of a new problem.

To understand the neural mechanisms that support this generalisation behaviour, we perform silicon probe recordings in medial frontal cortex (mFC). We find that the majority of neurons in mFC show progress-to-goal tuning, stretching and compressing their tiling to fit different goal distances. Crucially, we do not find an absolute representation of task state in mFC. Instead, neurons remap quasi-coherently in modules that preserve their relative tuning within a module, but not across modules or relative to the overall task sequence.

We explain these neural findings by introducing the (task)-Structured Memory Buffers model (SMBs): a set of ring-like modules anchored to specific behavioural steps. Neurons in an SMB fire with a fixed task-space lag from their behavioural anchor, effectively storing a memory of a previously visited location at a particular progress increment. The combined activity across SMBs therefore encodes the entire history of behavioural steps at every time point. SMBs' circular activity structure persists offline during sleep and rest, arguing for an intrinsic network organisation. Further, SMB activity is predictive of distal upcoming choices, meaning that on top of storing memories of past actions, SMB dynamics compute *future* actions.

Looking deeper at mFC's spatial code, we find simultaneous concrete and generalised spatial representations. A set of complementary population analyses recover maze topology while

simultaneously grouping locations by abstract features (e.g. “corners”, “centre”). Neural subspaces for distal goal-progress bins are near-orthogonal, consistent with mFC having separate spatial maps for each goal-progress increment.

Motivated by recent theoretical work, we next record from both medial and lateral entorhinal cortices to test for an explicit, factorised code for position in the behavioural sequence. Early analyses do not provide evidence that either region encodes absolute task state. Instead, medial entorhinal cortex has a stable spatial map across tasks and lateral entorhinal cortex shows task-general goal-progress tuning, building a picture of a multi-region functional hierarchy.

Our findings suggest that schemata of complex behavioural structures can be generated by sculpting progress-to-goal tuning into task-structured buffers of individual behavioural steps. This architecture allows new example sequences to be programmed into neural dynamics – rather than synaptic weights – and these dynamics automatically compute appropriate future actions. Ultimately, this work contributes to our knowledge of the neural basis of flexible intelligent behaviour.

Acknowledgements

They say it takes a village to raise a child. Apparently, it takes one to produce a thesis, too. So many people helped me in one way or another to get to the point of finishing this thing that it's impracticable to list you all by name. If you're reading this and feel like you ought to have been mentioned, just know that it's not because you weren't noticed or appreciated.

I couldn't have hoped for a finer gang of supervisors: Mohamady El-Gaby, Mark Walton, and Tim Behrens. It's probably used up a lifetime's worth of luck that you took a punt on a slightly lost grad student and let me work on some of the most interesting scientific questions it's possible to spend your time thinking about. To do so with such generous, knowledgeable, and creative mentors made it all the better. I've learned so much from each of you outside of the science. From Mohamady, your ability to see the best in a situation when things go wrong; from Mark, your skill at saying just the right thing at just the right time; and from Tim, your world-leading knack for vaping in the workplace so discreetly that it doesn't even seem inappropriate.

It's been a particular privilege to be the first DPhil to come through the El-Gaby lab. Systems-neuro projects are journeys of contrasts: from days at a time* down in the trenches dealing with rig noise, dead channels, chewed wires, and broken probes – to hours-long chats in Lisbon, dreaming of cognition. Every bit's been a laugh. Cheers, Mohamady.

Over the last few years I've been amazed and inspired by all members of the Behrens, Walton, and El-Gaby labs, but special mentions to Peter, Mingyu, Svenja, Alon, Lauren, Rebecca, Lillian, Noor, and Charles. Your impossible mixture of equanimity, immaculate vibes, and sheer competence is the stuff of envy. And of course, credit has to go to Takam. Jack of all trades and master of all of them. None of this would have happened without the things you built. You take open-source to its logical, personal extreme.

People's paths in science are varied; it basically all hinges on who you get to work with. I'm in the lucky position to have never had a supervisor that can be considered less than excellent. So, I give thanks to all of my past mentors: Jo Jackson, Crawford Winlove, Adam Packer, and Chris

*Weeks[†]

[†]Months

Summerfield. You gave me a deep appreciation for the beauty of the brain and mind.

Many people in Oxford have had my back in a pinch or two, but few more so than Debbie Clarke and Andy King. A big thanks also to Katherine and Jonathan, whose working knowledge of the human mind has helped me greatly in figuring out these little mouse ones.

To all my friends: the Faithfuls, the Tour Baes, the Bass Defence League, the Math Rock Devotees. You are, have been, and will be, a source of immeasurable joy. A big thanks to Becky, who always seems to know when I need to talk about wacky science or sad South London post-punk. And Justine, your well-timed Banh Mi in the depths of data collection reminded me, right when I needed it, that life is quite a lot better if not all of my calories come from HuelTM and MSG on toast.

To my dad, Loyd. Thank you for putting up with 40-minute long monologues on structure learning when you only asked for two, for talking me through changing paths to neuroscience over dinner at an open day (what an impact that turned out to have), and, well, just for being my Dad.

Ann, Rob, Dave, Claire, Tasha, Andrew, John, Yvonne, Jim, and Susan. You all know what you mean to me. Our spirited debates on such deep scientific questions as 'how far away is that boat?' and 'what is the volume of that phone box in units of Labradors?' sharpened my analytic thinking early on. Yvonne, my mate, those documentaries we watched together when I was a young lad – about mysterious limbs, flesh-eating bacteria, and jungle parasites – were probably what first got me into biology. And my brother, Ed. Dr Mr \$wag himself. You don't know quite how useful it is to have a younger, smarter version of myself around. Those long nights in the lads' flat playing shots-chess made lockdown a blast.

Finally, Georgie. No thanks could possibly be adequate, but I'll give it a go. Thank you for your patience, jokes, eye for detail, being a fixed point, and making us a home. This is as much yours as mine. Now, consider this a formal acceptance that my older code wasn't 'production standard' and let's put the issue to bed.

Author contributions

I declare that this thesis is my own, and where work has been contributed by other sources I have properly attributed credit. Where the pronoun *we* is used in this text, it refers to the team below.

Adam Loyd Harris: Conceptualised study; developed the ABCD task; collected behavioural and neural data (cohort 3 onwards); performed $\sim 25\%$ of medial frontal cortex surgeries; performed all medial entorhinal cortex surgeries; developed protocol and performed all lateral entorhinal cortex surgeries; wrote preprocessing pipelines; analysed and interpreted data for all chapters; contributed to the development of the Structured Memory Buffers model.

Mohamady El-Gaby: Conceptualised study; developed the ABCD task; performed $\sim 75\%$ of medial frontal cortex surgeries; collected behavioural and neural data (cohorts 1 and 2); analysed and interpreted data for Chapter 3, Chapter 4, and Chapter 5; formulated the Structured Memory Buffers model.

Tim Behrens: Conceptualised study; developed the ABCD task; interpreted data analysis; contributed to the development of the Structured Memory Buffers model.

Mark Walton: Conceptualised study; developed ABCD task; supported with data collection and analysis interpretation.

Thomas Akam: Contributed to study conceptualisation, designed the maze apparatus; wrote the PyControl software; provided technical support during data collection.

Figure 5.1 is an edited version of a schematic provided by **Gil Costa**. Refinements to the preprocessing pipeline in the lateral entorhinal cortex experiment were provided by **Peter Doohan**, **Charles Burns** and **Mingyu Zhu**. Finally, I was supported by research assistants **Yuxing Yang** during data collection for the lateral entorhinal cortex experiment, **Ben Pendry** for behaviour pretraining in the medial entorhinal cortex experiment, and **Arya Bhomick** for the final cohort of the medial frontal cortex experiment (and working through teething problems during the transition to Neuropixels).

The first three results chapters (Chapter 3, Chapter 4, and Chapter 5) are a subset of the work presented in a paper published in *Nature*:

*El-Gaby, M., **Harris, A. L.**, Whittington, J. C. R., Dorrell, W., Bhomick, A., Walton, M. E., Akam, T., and Behrens, T. E. J. (2024). A cellular basis for mapping behavioural structure. Nature, 636(8043):671–680.*

For these chapters, credit is to be attributed 50:50 between myself and **Dr Mohamady El-Gaby**.

Contents

List of Figures	xi
List of Selected Abbreviations	xiii
1 Introduction	1
1.1 Cognitive maps, learning sets, schemata: three sides of the same coin	3
1.2 Structured neural representations: the example of space	12
1.3 Prefrontal cortex and behavioural flexibility	17
1.4 Storing memories in weights or activations	27
1.5 Entorhinal cortex and factorised representations	30
1.6 Summary	33
2 General methods	35
2.1 Mice	35
2.2 Behavioural training	36
2.3 Surgery	41
2.4 Recording and preprocessing	43
2.5 Behavioural scoring	45
2.6 Tuning to task variables	47
2.7 Task space generalisation and modularity	49
2.8 Mnemonic fields in task space	51
2.9 Predicting distal choices	54
2.10 Activity structure in sleep	55
2.11 Generalised spatial tuning	57
3 The ABCD task: a behaviour paradigm for task structure abstraction	63
3.1 The ABCD Task: a behavioural paradigm for probing task structure abstraction .	64
3.2 Mice learn efficient routes between rewards, and acquire new examples more readily across learning	69
3.3 Inference of unseen task-space transitions: a strong test of task structure knowledge	73
3.4 Summary	75

4	Characterising task structure representations in medial frontal cortex	77
4.1	High-yield chronic <i>in vivo</i> electrophysiology with silicon probes	78
4.2	Space and goal-progress tuning are dominant features of mFC representations	80
4.3	mFC does not have generalising state tuning in the ABCD task	89
4.4	Summary	95
5	mFC neurons are organised into task-Structured Memory Buffers	97
5.1	An introduction to the task-Structured Memory Buffers (tSMB) model	99
5.2	Neurons in a tSMB have mnemonic fields in task space	104
5.3	Summary	113
6	Structured memory buffers are internally organised and predict distal choices	115
6.1	Neurons in a tSMB retain their ring-like dynamics offline	116
6.2	Neurons in a tSMB predict distal behavioural choices	122
6.3	Summary	132
7	Signatures of generalised spatial representations in mFC	135
7.1	Euclidean distance and node degree are dominant features of neural subspaces	137
7.2	Representational similarity analysis supports simultaneous concrete and generalised spatial coding	141
7.3	Decoder generalisation supports simultaneous concrete and generalised spatial coding	144
7.4	Summary	147
8	Entorhinal cortex in the ABCD task: preliminary data	149
8.1	Neuropixels implants in medial and lateral entorhinal cortices	151
8.2	Example MEC and LEC neurons in the ABCD task	153
8.3	Goal-progress tuning is stronger in LEC-targeted neurons than in MEC	157
8.4	Neither LEC nor MEC neurons consistently align to task state	159
8.5	Neither LEC nor MEC neurons remap coherently across tasks	161
8.6	Summary	165
9	General discussion	167
9.1	Limitations and links to wider literature	169
9.2	Future directions	177
10	Outro	181
	Bibliography	183
11	Appendix: Mapping variables with attractors	199

List of Figures

1.1	Tolman’s sunburst maze	6
1.2	Harlow’s learning sets	8
1.3	A “zoo” of spatially-tuned functional cell types in the hippocampal formation . .	15
1.4	PFC acts as a ‘switchboard’ to route signals throughout the brain in service of a goal	18
3.1	The ABCD task - a behaviour paradigm to probe task structure abstraction	66
3.2	Decoupling the structure of the environment from the structure of the task	68
3.3	Mice learn efficient routes between goals and task acquisition improves across examples	72
3.4	Mice use the abstract task structure to infer the d→a transition on the first trial .	74
4.1	Chronic high-yield electrophysiology with silicon probes in medial frontal cortex	80
4.2	Example mFC neurons tuned to place, state, and goal-progress	83
4.3	A generalised linear model analysis of mFC representations	85
4.4	Progress-to-goal is a fundamental feature of mFC representations	88
4.5	mFC neurons do not code explicitly for positions in the ABCD task structure . .	91
4.6	Modular organisation of coherently remapping neurons in mFC	94
5.1	The (task)-Structured Memory Buffers model	103
5.2	Mnemonic fields in task space 1 - aligning neural data to the best-fitting anchor .	107
5.3	Mnemonic fields in task space 2 - neurons form place cell-like maps of previously visited locations	110
5.4	Mnemonic fields in task space 3 - an elastic-net approach to predicting task space maps with multiple competing anchor×lag conjunctions	113
6.1	Correlation structure in sleep sessions suggests ring-like dynamics persist offline within coherently remapping modules	119
6.2	GLM analysis suggests ring-like dynamics persist offline during sleep/rest within coherently remapping modules	121
6.3	Schematic - distal prediction of behavioural choices from neurons in a tSMB . . .	124
6.4	Activity of mFC neurons is increased at bump time before distal anchor visits . .	127
6.5	Activity of mFC neurons predict distal future choices at the time specified by their position in a tSMB	131

7.1	Example neurons that generalise over locations of the same abstract type	137
7.2	Concrete locations and node-degree are dominant features of neural subspace, and subspaces of distal goal-progress bins are near-orthogonal	139
7.3	A confound in detecting node-degree tuning with neural subspace analysis . . .	141
7.4	Both Euclidean distance and node-degree explain neural dissimilarity	143
7.5	Decoding analysis supports simultaneous abstracted and concrete spatial tuning	146
8.1	MEC histology and theta rhythm	151
8.2	Histology - Neuropixels 2.0 targeted to lateral entorhinal cortex	152
8.3	High yields of neurons stable across double days in the LEC-targeted cohort . . .	153
8.4	Example MEC task-space rate maps across tasks	154
8.5	MEC spatial tuning is stable across tasks	156
8.6	Example LEC task-space rate maps across tasks	157
8.7	Goal-progress tuning is stronger in LEC-targeted units than in MEC	159
8.8	No evidence of task-aligned generalising state cells in MEC or LEC	161
8.9	Above chance coherent remapping for task-space proximal, but not distal, MEC and LEC neurons	163
8.10	MEC and LEC task-space remapping is consistent with pure spatial tuning . . .	165
11.1	Line attractor dynamics in a biological circuit	203
11.2	A ring attractor network with left and right velocity inputs	205
11.3	Toroidal topology of the grid cell network	206

List of Selected Abbreviations

1D, 2D, ...	1-Dimensional, 2-Dimensional, ...
ACC	Anterior Cingulate Cortex
AI	Artificial Intelligence
BOLD	Blood Oxygen Level Dependent
CAN	Continuous Attractor Network
dlPFC	Dorsolateral Prefrontal Cortex
dmPFC	Dorsomedial Prefrontal Cortex
EM	Episodic Memory
fMRI	Functional Magnetic Resonance Imaging
GLM	General Linear Model
GP	Goal Progress
HPC	Hippocampus
LEC	Lateral Entorhinal Cortex
M1, M2	Primary/Secondary Motor Cortices
MEC	Medial Entorhinal Cortex
MD	Mediodorsal Thalamic Nucleus
mFC	Medial Frontal Cortex
ML	Machine Learning
MTL	Medial Temporal Lobe
OFC	Orbitofrontal Cortex
PCA	Principal Component Analysis
PFC	Prefrontal Cortex
PrL	Prelimbic Cortex
RDM	Representational Dissimilarity Matrix

RL	Reinforcement Learning
RNN	Recurrent Neural Network
RPD	Relative Path Distance
RSA	Representational Similarity Analysis
(t)SMB	(task) Structured Memory Buffer
SR	Successor Representation
SVM	Support Vector Machine
SWM	Sequence Working Memory
TEM	Tolman-Eichenbaum Machine
UMAP	Uniform Manifold Approximation and Projection
vmPFC	Ventromedial Prefrontal Cortex
WM	Working Memory

For Maggie.

"Funes remembered not only every leaf of every tree of every wood, but also every one of the times he had perceived or imagined it... I suspect, however, that he was not very capable of thought. **To think is to forget differences, to generalise, to abstract.** In the teeming world of Funes, there were only details, almost immediate in their presence."

— Borges, *Funes the Memorious* (1942)

"All jokes are the same, it's just the words that are different."

— Stewart Lee

1

Introduction

Contents

1.1	Cognitive maps, learning sets, schemata: three sides of the same coin	3
1.2	Structured neural representations: the example of space	12
1.3	Prefrontal cortex and behavioural flexibility	17
1.4	Storing memories in weights or activations	27
1.5	Entorhinal cortex and factorised representations	30
1.6	Summary	33

In a short story by Jorge Luis Borges, Ireneo Funes gains the ability to remember everything he experiences following a hard fall from his horse (Borges, 1999). Although this is not the worst recorded outcome of such an injury (Krüger et al., 2018), he is left, we are told, without the ability to abstract across these memories. To Funes, it is not only inconceivable that the category of *dogs* encompasses so many disparate creatures of such diverse sizes and forms, but also that a dog as viewed from the front is not a different entity to the same dog as viewed from the side. This inability to handle abstractions – to extract common structures in the world by forgetting differences and emphasising similarities – precludes the use of the higher cognitive functions characteristic of human intelligence (James, 1890). If each experience is totally new, and every detail is of equal importance, then there are no means by which you can draw from your past to steer your future.

Our experiences of the world are rich with structure. This is true of the world itself, with simple laws of geometry and topology constraining possible connections between locations no matter where you are, but also of our behaviour within that world. From cooking a meal to solving a maths problem, we compose elaborate sequences of actions to achieve our goals. When elements of this structure are common across experiences, we can build schemata: generalised representations of task states that enable us to infer new behavioural sequences (Piaget, 1952; Bartlett, 1933).

Consider the example of attending a wedding. In the UK at least, a wedding schema would look something like: dress smartly, arrive at the venue, mingle politely, sit for the ceremony, eat dinner, mingle less politely, dance, then leave. When we go to a wedding at an unseen venue with an unfamiliar group of people, we can immediately infer that after dinner we are expected to mingle and then dance without having to experience it, because we *know how weddings work*. Such generalisation of behavioural schemata also applies in domains totally separate to space. When shopping online for the wedding gift, we know we have to go to our basket and click checkout even if we have never shopped on this particular site before. In cases where the schema we have learnt through experience differs radically from what actually unfolds in the world – such as a different culture’s wedding, or online shopping through a university procurement system – we consider ourselves to be lost. But we are not lost in the typical, navigational sense. We are lost in the task.

These examples are tightly contained in time, from minutes to a single day, but schemata can be extended and abstracted almost arbitrarily. Consider courtship, negotiation, getting a degree. A mental schema for these dynamic and extended sequences of behaviours provides a high-level scaffold that organises our expectations, guides decision making, and allows us to interpret events as part of a coherent process, even when individual steps are unfamiliar or unfold over months or years.

This thesis concerns the neural circuits (and the algorithms they implement) that underpin the ability of intelligent animals to generalise the structure of problems and transfer them to new scenarios. In the absence of experimental paradigms that allow single-neuron recording in humans across the rich, temporally extended behaviours detailed above, we will be taking the underlying logic of behavioural structure generalisation and implementing it in a way that can be performed by mice, the model-organism *du jour* of cognitive systems neuroscience. We will argue (hopefully it will be considered not just an argument, but a demonstration) that mice can learn a task schema to such a degree that they can perform inferences upon first exposure to a new problem, show neural representations from multiple brain regions while the mice perform this complex cognitive behaviour (frontal and entorhinal cortices), and develop a model that explains these behavioural and neural observations.

Before detailing our experimental findings, we will discuss the history and state-of-the-art of research into learning and representing the structure of the world and the structure of tasks. We will focus the neural discussion on frontal cortex and the medial temporal lobe: the regions from which we record. This, I hope, will provide sufficient context to understand the inspirations for the current study.

1.1 Cognitive maps, learning sets, schemata: three sides of the same coin

We owe so much of our understanding of the mind to the fact that monkeys like fruit, people like stories, children like play, and rats like anything so long as they can eat it. In this section, we will cover three classic strands of research that, on the surface, tackled very different problems – how rats navigate mazes, how monkeys *learn how to learn*, and how humans recall stories – but which all converged on the idea that animals construct internal representations of structure that go beyond specific stimuli or responses.

1.1.1 Tolman and the Cognitive Map

In the twentieth century, there was a tension in psychology between behaviourism and cognitivism. The hard-line behaviourists held that even very complex behaviour could be explained as long chains of reinforced stimulus–response associations (S–R chains), learned through trial and error (Thorndike, 1911; Watson, 1913; Hull, 1943; Skinner, 1938). The cognitivists, by contrast, sought to re-centre the role of the internal: mental models, thoughts, and expectations (Broadbent, 1958; Miller, 1956; Neisser, 1967; Chomsky, 1959). It is easy to caricature strict behaviourism as claiming the internal is unimportant, or even that it does not exist. The more reasonable interpretation is that internal states were deemed unobservable and thus off-limits for scientific explanation (Watson, 1913; Skinner, 1953). The inputs to the animal (environmental stimuli, rewards/reinforcements) and the corresponding behavioural outputs were considered the primary rigorous basis for psychology (Skinner, 1953).

In the 1930s, however, Edward C. Tolman and others demonstrated forms of learning in rats that could not be explained by simple S-R chains or by reinforcement alone. Tolman's work provided clean behavioural evidence for what he called a cognitive map: an internal representation of the layout of the world that the animal could exploit in a goal-directed, flexible way. Crucially, these internal representations could be acquired without explicit rewards, and were relational in the sense that they encoded relationships among stimuli or locations, rather than specific motor commands. Here, I will summarise the key experiments that led to this conceptual development.

Latent learning

Tolman and Honzik (1930) showed that rats can learn about the environment even without reinforcement. In their experiment, one group of rats was allowed to freely explore a maze for several days with no rewards present. A second group received a food reward at the end of the maze on every training trial. As expected, the rewarded group gradually improved (fewer errors, faster completion) over repeated trials, whereas the unrewarded group showed little improvement. In the next phase, a reward was placed in the maze for the previously unrewarded rats. On the very next trial, their performance jumped dramatically, matching or

exceeding that of the always-rewarded group (Tolman and Honzik, 1930). This is to say that the rats that had explored with no reward had nevertheless learned the maze layout, forming an internal model of which corridors connected to where. As soon as there was motivation to run to the goal, they could exploit this latent knowledge. This result was a challenge to pure behaviourist doctrine because it implied that learning of the maze's spatial structure had occurred without reinforcement. The reward affected the rats' performance (motivating them to use what they had learned), but not the learning itself.

Place vs. response learning

In another study, Tolman et al. (1946) investigated specifically whether the rats had learned *actions* to take or *locations* in space. For this, they trained rats to run from a start arm in a maze to a goal arm for food. They then changed the starting position of the rat relative to the room. If the rats had simply learned a reinforced motor response (say, turn left then right), they should repeat that same sequence of turns from the new start position, which would lead them to the wrong place. If, however, they had learned the place where the food was located in the room's coordinates, they should navigate to that same place even from a novel start position. Tolman found that, under many conditions with rich environmental cues, rats favoured the place strategy over the simple response strategy: they had encoded the goal in an allocentric coordinate frame (linked to external room cues) rather than memorising a series of egocentric turns. This suggested that the internal representation was spatial and relational (a map of the environment) rather than a chain of movements. Later research showed that with certain overtraining or cue-poor environments, animals can revert to response strategies (Restle, 1957), but Tolman's core point was that an allocentric place code allows much greater flexibility because it enables immediate transfer of knowledge to new situations, like starting from a different point in the maze.

Detours, shortcuts, and novel routes

Perhaps the most vivid evidence for a cognitive map came from experiments where rats had to take novel routes in a familiar space. If a well-learned path to a goal was suddenly blocked, rats often took another path that led to the goal, even if that route was never directly reinforced before. In Tolman's sunburst maze experiment (Figure 1.1), after rats learned to run down a corridor to

a goal, the maze was rearranged into a sunburst of alternative paths, with the corridor they took during the training period now blocked. Many rats chose a path that led in the direction of the goal's previous location, rather than simply backtracking or picking paths at random (Tolman, 1948). Such behaviour is hard to explain with strictly habit-based accounts, which would predict perseveration when the usual path is blocked. It is, however, naturally explained if the rat has an internal map of the environment and knows about spatial relationships, allowing it to compute a new route on-the-fly.




Figure 1.1: Tolman's sunburst maze

Adapted from Tolman (1948). Left: During training, rats start in a circular arena and follow a corridor to a reward location, marked with an arrow. Right: During test, the original corridor is blocked, and the rat is instead presented with 18 radial arm corridors, with the corridor from the training period blocked. Rats immediately infer which of the new corridors (5) take them to the reward location.

Taken together, these findings led Tolman to propose that rats (and by extension, other intelligent animals) form cognitive maps of their environments. A cognitive map, in Tolman's view, is an internal model encoding the topology and geometry of the environment, allowing the animal to plan novel paths and infer outcomes beyond its direct experience. The cognitive map model was influential in part because it challenged S-R theories on their own empirical ground, behaviour. Classic S-R reinforcement theories had great success in explaining phenomena like gradual learning curves, generalisation gradients, and simple discrimination learning (Thorndike, 1911;

Skinner, 1938; Ferster and Skinner, 1957; Spence, 1936; Guttman and Kalish, 1956). But latent learning and novel-route inference are difficult to reconcile with S-R chains. It is, in principle, possible to add complexity to S-R models (for instance, assume the rat learned a very long series of conditional rules covering each possible contingency), but such accounts are clumsy and unparsimonious.

While navigation in physical space was the domain in which Tolman introduced the cognitive map, he took it to be a far more domain-general organising principle for encoding relational structure (Tolman, 1948). Later neuroscience grounded the spatial case in hippocampal *place cells* and further proposed that the hippocampus supports general relational memory that extends beyond the geometry of the environment to abstract state relations (O’Keefe and Dostrovsky, 1971; O’Keefe and Nadel, 1978; Eichenbaum et al., 1999; Eichenbaum, 2017a).

With this in mind, we will now turn to early influential work on how the generalisable relational structure of *problems* is acquired and used in new situations.

1.1.2 Harlow and the Learning Set

In parallel with Tolman’s challenge to S-R chaining in spatial navigation, Harry Harlow showed that rhesus monkeys rapidly acquire a capacity to *learn how to learn* across series of object-discrimination problems (Harlow, 1949). In Harlow’s paradigm, animals solve many two-choice discriminations in succession (e.g., A⁺ vs. B, then C⁺ vs. D, then E⁺ vs. F, with novel pairs each time). Within a problem, the monkeys get immediate feedback through fruit reward for correct responses. Critically, monkeys become dramatically faster at reaching high performance on later problems despite the specific objects changing each time (Figure 1.2). This improvement reflects the learning of an abstract problem schema: the monkeys discover a general strategy (e.g., sample-test, then adopt a *win-stay/lose-shift* rule) and reuse it on new instances (Harlow, 1949).

A closely related phenomenon appears in *reversal learning sets*, where after mastering a discrimination, reward contingencies are flipped (A^- vs. B^+). With experience across many reversals, monkeys approach near-one-trial reversals wherein a single piece of negative feedback is enough to trigger switching, indicating that the animal has internalised the structure of the reversal task rather than memorising specific stimulus-reward pairings (Harlow, 1949).




Figure 1.2: Harlow's learning sets

Adapted from Harlow (1949). Monkeys perform a long series of novel two-choice object discriminations. Although each problem uses new objects, the rate of learning (as measured by the fraction of correct responses on a specific trial averaged over a consecutive block of problems) rises sharply across problems as animals acquire a *learning set*: a reusable strategy such as win-stay/lose-shift.

Strikingly, even extreme cortical lesions do not abolish the capacity to acquire abstractions about task structure. In hemidecorticated (single-hemisphere) monkeys, performance on discrimination and delayed-response problems remains sufficiently intact to demonstrate new learning (Thompson and Bucy, 1969).

Harlow's learning set experiments provide behavioural evidence that animals form abstract task representations that generalise across changing sensory/motor features. This is a non-spatial, 'task-space' analogue to Tolman's cognitive maps. Where Tolman emphasises a relational map of places, Harlow demonstrated a relational map of problems: rules and contingencies that transfer to new task instances (Harlow, 1949).

An aside: Bateson's Deutero-learning

In Harlow's work, monkeys demonstrate two forms of learning. At one level, the monkeys learn which of two stimuli lead to them getting a reward. On a higher level, they learn a strategy that generalises across problems of the same class, be it simple discrimination or reversal learning. Separately to Harlow, the concept of hierarchies of learning was explored by the anthropologist Gregory Bateson (Bateson, 1972). He distinguished:

- Learning 0: No change from experience (pure reflex/repetition).
- Learning I: (Protolearning): Trial-and-error changes in specific responses (e.g. forming or extinguishing associations).
- Learning II: (Deuterolearning): Changes to the procedures of Learning I.
- Learning III: Rare, transformative shifts in one's repertoire of Learning II patterns.

In this sense, Harlow's learning sets instantiate Bateson's Learning II, in that animals track their learning of individual stimulus pairs to abstract task structure and reuse it across superficially novel problems (Harlow, 1949; Bateson, 1972). These ideas link to modern *meta-learning*, where a slow across-task process shapes priors so that within-task acquisition is rapid: an inner loop for learning a specific problem, and an outer loop for learning across problems (Finn et al., 2017; Wang et al., 2016, 2018). We will touch on meta-learning in modern machine learning (ML) and artificial intelligence research in later sections.

1.1.3 Piaget, Bartlett, and the Schema

Tolman's cognitive map and Harlow's learning sets point to structured, relational representations of the physical world, and the structure of problems. Some early work about such schematic structure in human cognition came from Jean Piaget and Frederic Bartlett (Piaget, 1952; Bartlett, 1933). Although developed in different traditions (developmental psychology/epistemology for Piaget, experimental memory psychology for Bartlett) their work converges on the view that cognition relies on organised frameworks that guide interpretation, learning, and recall. In this section, we will briefly sketch these complementary strands of thinking.

Bartlett's *War of the Ghosts*

Bartlett approached schemas through the lens of reconstructive memory. In his *War of the Ghosts* experiments, student participants recalled and retold an unfamiliar folk story across multiple delays (serial reproductions). Upon these retelling events, systematic distortions emerged. Participants omitted atypical information, changed details toward more culturally familiar forms, and rationalised parts of the story to increase its coherence (Bartlett, 1933). Bartlett inferred that remembering is not playback of stored traces but an active reconstruction guided by organised knowledge structures that encode typical relations among events, objects, and causal roles. Schemas both enable efficient inference from sparse evidence and bias recollection toward prior expectations.

According to this perspective, learning integrates new episodes into existing schematic organisation and the act of recollection generates a plausible narrative by fitting cues to schematic relations. The same process that supports efficient prediction and interpretation can also explain systematic error. When incoming information conflicts with prior organisation, reconstruction can resolve ambiguity in favour of the schema.

Piaget and play

Piaget introduced schemas as organised patterns of action or thought that structure how an organism engages with the world. Two core adaptive processes govern schema dynamics: *assimilation* (integrating new experience into extant schemas) and *accommodation* (modifying schemas to fit incongruent input). Cognitive development reflects progressive reorganisation of schemas across stages, starting from sensorimotor action patterns up to operational structures. The result is a set of increasingly abstract combinations and transformations (Piaget, 1952).

For Piaget, play is central to this process. In play, children actively assimilate reality to existing schemas, relaxing external constraints to explore combinations of actions and roles. This "safe sandbox" supports hypothesis testing and schema refinement in a low-risk setting. Crucially, Piaget argues that the same mechanisms that scaffold the child's manipulation of objects and rules also support later symbolic and logical operations. These schemas are more than memory templates, like those that bias the recollection of stories in Bartlett's experiment. They generate predictions, constrain interpretation, and enable transfer across superficially different contexts, ultimately in service of flexible intelligent behaviour.

In our terms, a schema is a reusable representation of task structure. Tolman's map is a spatial schema, Harlow's learning set is a problem schema, and Piaget and Bartlett articulate the same principle at the levels of conceptual development and memory. All converge on a single deep functional claim. Cognition depends on internal models that encode relations (among places, objects, events, and rules) and can be applied to novel instances with minimal data.

1.1.4 In this thesis – decoupling spatial structure from task structure

Flexible behaviour rests on at least two forms of knowledge: the structure of the world (like Tolman's cognitive map) and the structure of the task (like Harlow's learning sets). To understand their neural representations and computations, the two must be dissociated. Our approach is to design a paradigm that decorrelates abstract structure from environmental geometry, allowing

us to examine each code on its own terms and to test how they combine to guide behaviour.

With this separation in mind, in the coming sections we will explore the neural representations of spatial and behavioural structure, and what we have learned about how these afford flexible intelligent behaviour.

1.2 Structured neural representations: the example of space

A central theme in contemporary systems neuroscience, and in this thesis, is that the brain builds structured internal representations that capture the relational organisation of the world and our behaviour. These representations allow agents to generalise, to infer unexperienced relationships, and to flexibly adapt behaviour when goals or contexts change (Tolman, 1948; Behrens et al., 2018; Whittington et al., 2022). Although cognitive maps may span many domains (social, conceptual, task-structural), space provides the most concrete example. Here, decades of work have identified specific cell types that together instantiate a richly structured code for locations, directions, and distances. The resulting picture has become a model system for how the brain might represent structure more broadly.

1.2.1 Space in the medial temporal lobe

Functions of the medial temporal lobe: insights from lesion studies

A rich corpus of literature has shown that the medial temporal lobe, particularly the hippocampus and entorhinal cortices, are critical for flexible behaviour, particularly in – but not limited to – spatial cognition. This functional role is neatly demonstrated by lesion studies. For instance, in an experiment in which rats had to take a familiar route from a start to a goal in a water maze, lesions to the fornix (a major hippocampal input and output nerve bundle) do not prevent rats from following this familiar path. However, when the start location or global cues changed, fornix lesions degraded the rats' ability to navigate to that goal. That is, they were inflexible relative to control rats and struggled to form new place associations (Eichenbaum et al., 1990). Complementary findings come from more complex spatial environments. In a 'honeycomb' maze, rats with hippocampal lesions have significantly diminished ability to

navigate, particularly in cases where low-level strategies, like using a simple heading vector to the goal, are not useful (Wood et al., 2018). Further, hippocampal lesions have specific relational deficits in non-spatial problems. A particularly salient case is transitive inference, with rats being able to learn simple pairwise associations (like $A > B$, and $B > C$) but are unable to perform inferences based on this knowledge ($A > C$; Dusek and Eichenbaum (1997); DeVito et al. (2010)). Outside of the hippocampus, lesions to lateral entorhinal cortex (LEC) impair the formation of object-place associations (Kuruvilla et al., 2020), and medial entorhinal cortex (MEC) lesioned mice show impaired path integration (Gil et al., 2018).

Together, these lesion results argue that medial temporal lobe structures support flexible behaviour by providing hippocampal relational maps that permit new associations and inferences, MEC computations for integrating self-motion, and LEC mechanisms for context/item binding. We will now turn to the representations and computations in the medial temporal lobe that support these functions.

Spatial representations in the medial temporal lobe

The study of the brain's spatial *representations* begins with the discovery of place cells – hippocampal neurons fire selectively at particular locations in an environment (O'Keefe and Dostrovsky, 1971; O'Keefe and Nadel, 1978). This finding established the hippocampus as a candidate substrate for a map-like representation of space with enough precision to support navigation and memory. Looking just one synapse upstream in medial entorhinal cortex, *grid cells* exhibit strikingly regular, triangular-lattice firing fields that provide a multi-scale metric for space (Hafting et al., 2005). These neurons form modules, with neurons within a module having the same spacing between firing fields, and orientation of the grid relative to the environment, but are phase-offset from one another in space. Across modules, grid spacing (and to a lesser extent, orientation) varies. The scaling ratio is discrete and follows steps of roughly $\sqrt{2}$. Further, the modules are arranged anatomically along the dorsoventral axis (Stensola, 2012).

While grid cells and place cells are the most famous, there is a large “zoo” of interpretable spatially-selective neurons in the hippocampal formation (Figure 1.3). Head-direction cells in postsubiculum and related regions fire as a function of the animal’s orientation relative to allocentric cues, providing a circular code for orientation (Taube et al., 1990). Boundary vector and object vector cells fire at a fixed distance and orientation from objects and boundaries – behaviourally-salient features of the world (Lever et al., 2009; Høydal et al., 2019). Further, we observe social place fields for other animals’ locations (Danjo et al., 2018; Omer et al., 2018).

The organisation is pretty elegant. Periodic grid codes can, in principle, support Euclidean operations and path integration (Burak and Fiete, 2009; McNaughton et al., 2006), head-direction codes act as a ring attractor for heading (see Figure 11.2), and place codes bind this metric structure to landmarks and episodic content (Whittington et al., 2020; George et al., 2021). Together these neuron types provide a hyper-flexible mechanism for mapping out a state spaces relevant for behaviour (Bakermans et al., 2025).

Two computational/theoretical developments have furthered the link between spatial map structure and its use. First, the *predictive map* or successor-representation (SR) view reframes hippocampal codes as learning, for each location, the expected future occupancy of other locations under a policy (Stachenfeld et al., 2017). This emphasises transition structure rather than pure geometry, explaining things like reward sensitivity and policy dependence of hippocampal activity. In SR terms, the hippocampal place codes are a map of how states lead to other states. Value and planning in RL problems become readouts from this representation. Second, work in artificial agents has shown that grid-like units emerge as an efficient basis for vector navigation, enabling shortcutting and direct goal vectors (Banino et al., 2018). This provides computational support for the hypothesised role of entorhinal grid cells as a metric code for space.




Figure 1.3: A “zoo” of spatially-tuned functional cell types in the hippocampal formation

Adapted from Behrens et al. (2018). Neurons in the hippocampal formation show diverse and interpretable representations of spatial variables.

1.2.2 Spatial representations in prefrontal cortex

Outside of the hippocampal formation, spatial responses have been discovered in frontal regions across species and of different formats. I will introduce it briefly here, but we will revisit spatial tuning in frontal cortex in much more detail in Chapter 7.

In humans, hexagonally symmetric grid-like codes have been found in frontal cortex of people navigating a virtual arena (Doeller et al., 2010). In rodents, mFC and anterior cingulate cortex (ACC) contains place cells (Sauer et al., 2022; Bota et al., 2021). Though, interestingly, in ACC the *formation* but not *maintenance* of these place cells is dependent on hippocampal input. If the hippocampus is lesioned before an environment is explored, the place cells never form. But once the spatial map has been established in ACC, hippocampal lesion does not degrade these place cells, suggesting the ACC’s spatial map is learned from hippocampal input (Bota et al., 2021).

Intriguingly, the ability of frontal regions to form full spatial maps is not found in all cases. A parallel literature finds that neurons in frontal cortex instead generalise across locations with similar sets of features. For instance, in a structured environment consisting of four reward ‘wells’ connected by ‘corridors’, frontal neurons reliably code for wells vs corridors rather than specific locations (Yu et al., 2018). Similarly, in a cross-shaped maze, frontal neurons do not form place-cell type responses for concrete locations but instead code for the start, middle, or end of a run, collapsing over the specific locations of that run (Kaefer et al., 2020; Nardin et al., 2023, 2021; Berners-Lee et al., 2021; den Bakker and Kloosterman, 2025).

These conflicting findings raise questions about the situations in which frontal cortex expresses a spatial map and whether it is driven by task demands.

1.2.3 Spatial maps for non-space

When Tolman formulated the cognitive map idea he used spatial maps as his working example, but intended the idea to be a much more general principle for relational coding (Tolman, 1948). Mirroring this intention, recent neuroscience has put forward that the neural codes and algorithms that underpin spatial cognitive maps can be used to map arbitrary non-spatial state spaces (Eichenbaum et al., 1999; Dusek and Eichenbaum, 1997; Eichenbaum, 2017b; Whittington et al., 2020; George et al., 2021). In a recent review, Behrens et al. (2018) argue that the key components (states, relations among states, and operations that move an internal pointer through that structure) are not unique to space. Spatial cell responses are instead flexible basis sets that can act as a substrate for rapid and efficient generalisation across domains (Dordek et al., 2016).

Owing to the discovery of grid $\times$ direction conjunctive cells in the deep layers of medial entorhinal cortex, hexagonally symmetric grid-like signals are detectable in coarse-grained fMRI BOLD signals (Doeller et al., 2010). This led to the discovery of grid-like signals, map-like generalisation, and relational coding in tasks that vary non-spatial features (conceptual dimensions, latent task states, social hierarchies). This suggests a common representational format for conceptual spaces that matches that of physical space (Aronov et al., 2017; Constantinescu et al., 2016; Bao et al.,

2019; Chen et al., 2024a; Liang et al., 2024; Park et al., 2021; Bongioanni et al., 2021). Notably, these signals are located in *medial prefrontal cortex* (and for more abstract codes, its ventral region), as well as entorhinal cortex. Even in the spatial domain, we observe single neuron responses for task-related variables like vectors to goals (Sarel et al., 2017), reward location (Gauthier and Tank, 2018), reward probability contingencies (Knudsen and Wallis, 2021), and accumulated evidence (Nieh et al., 2021). Further, the highly consistent spatial metric provided by grid cells is warped by task variables like goal locations, both in rate remapping (Butler et al., 2019) and distortion of spatial fields (Boccaro et al., 2010).

It seems then that mapping space is just one specific case of a more general set of computations: identify latent variables and possible transitions; learn a geometry (or graph) over them; and bind that abstract scaffold to observations and goals so that new inferences become possible (Whittington et al., 2022, 2020; George et al., 2021; Behrens et al., 2018).

In sum, the spatial navigation literature provides a uniquely transparent instance of structured neural representation that allows us to understand the algorithms neural circuits may be implementing, and how these may be used in other non-spatial problems.

1.3 Prefrontal cortex and behavioural flexibility

The aim of this thesis is to identify the neural mechanisms supporting the storage and re-use of schematic knowledge to guide behaviour in new situations. A century of research has pointed to the prefrontal cortex (PFC) being crucial for flexible behaviour, particularly in cases where rules or contingencies change frequently (Walton et al., 2010; Birrell and Brown, 2000; Owen et al., 1991). An influential theory of PFC's function posits that its job is to generate and maintain representations of goals, and to flexibly route information in other brain areas to achieve that goal (Miller and Cohen, 2001) – see Figure 1.4. In keeping with this broad functional remit, damage to various PFC subregions has been linked to wide-ranging cognitive effects. Furthermore, its representations are mixed, often difficult to interpret at the single-neuron level, and vary depending on task demands. The majority of the data in the current

work are recordings from mouse medial frontal cortex (mFC), a rough homologue of PFC in rodents (I will discuss anatomical homology at the end of this section). We will now cover the behavioural effects that arise from PFC damage – and the sorts of variables PFC subregions represent – that justify our focus on this area.




Figure 1.4: PFC acts as a 'switchboard' to route signals throughout the brain in service of a goal

Adapted from [Miller and Cohen \(2001\)](#); [Cohen et al. \(1990\)](#). Example setting in which PFC acts as a 'switchboard' to route signals to generate goal-directed behaviour. In a Stroop task, participants must report either the word or the colour of the text. Here, PFC represents the current rule (colour or word). Lines denote synaptic connections with thicker lines indicating stronger weights. Loops ending in small black circles mark mutual inhibition within a layer (e.g., between the 'red' and 'green' output units). **(a)** In the absence of PFC input, when a conflict stimulus is presented – the word "green" coloured in red – the strong connections in the reading pathway between the word stimuli and the verbal output result in a "green" response. **(b)** With PFC representing the *colour* rule, the colour units are activated, biasing the response in favour of "red", despite the word-reading pathway having stronger synaptic weights.

1.3.1 What happens when you damage PFC?

Broadly, perturbations of prefrontal circuitry across species do not abolish simple, overlearned habits. Instead, they selectively disrupt flexible, goal-directed control – particularly when rules change, when attention must be reallocated to new task-relevant features, when credit must

be assigned correctly after ambiguous outcomes, or when choices must be inferred rather than retrieved from direct experience.

Early causal evidence came from primate and human work. Lesions of frontal cortex in monkeys abolished performance on delayed-response problems (a test of maintaining task rules and goals over short delays) whilst sparing simple habits (Jacobsen, 1935). In parallel, clinical frontal resections in humans produced disproportionate impairments in planning, initiative, and social/behavioural control despite relatively preserved basic intellect, consistent with a deficit in organising behaviour around internal goals rather than in basic perceptual or mnemonic capacity (Penfield and Evans, 1935).

Later rodent studies build on this with sub-region specificity. Damage or inactivation of medial PFC (prelimbic cortex – the region in which most of the mFC neurons recorded in this thesis are located) causes strong deficits when tasks demand sustained attention or a shift away from a prepotent strategy. Notably, prelimbic perturbations selectively impair extra-dimensional (feature-set) shifts – rodent analogues of the human Wisconsin Card Sorting Test – while leaving simpler discriminations relatively intact (Granon and Poucet, 2000; Birrell and Brown, 2000). During concept learning, the prelimbic cortex appears to maintain attention to category-relevant features and to update category representations in a recency-weighted manner. Lesions or inactivation blunt this trial-by-trial updating, implicating prelimbic circuits in the on-line control of what is learned and from which features (Broschard et al., 2021).

Orbitofrontal (OFC) and ventromedial (vmPFC) regions contribute complementary computations for flexible choice. In a task in which macaques track varying reward contingencies in a multi-armed bandit, OFC lesions spare basic acquisition but disrupt the correct credit assignment to the choice that actually produced reward when outcomes are temporally confusable, revealing a role in contingent learning over and above just tracking reward history (Walton et al., 2010). Similarly, timed causal perturbation experiments support OFC's role in contingency learning. Closed-loop, theta-locked microstimulation of macaque OFC disrupts reward-guided learning,

indicating that (hippocampus-driven) OFC rhythms support credit assignment (Knudsen and Wallis, 2020).

In humans, ventromedial PFC (vmPFC) lesions show a related but distinct vulnerability. Local pairwise learning can be preserved while inferential choice is impaired (e.g., in transitive inference between items never directly compared) implicating vmPFC in building and querying *relational value structures* needed for inference (Koscik and Tranel, 2012). Transient ultrasound disruption of medial frontal cortex in primates selectively degrades novel, inferential choices and, intriguingly, reveals grid-like population codes for comparing multi-attribute options, suggesting that prefrontal circuits construct structured, reusable representations to support generalisation (Bongioanni et al., 2021).

Taken together, these lesion and perturbation studies build a picture of PFC sustaining abstract goal/context representations and exerting top-down control over attention, learning, credit assignment, and inference. This is especially true when environments change or when generalisation to unseen conditions is required. In rodents, prelimbic cortex biases which features are attended to and updated; in primates, OFC ensures that outcomes are assigned to their causal choices; vmPFC supports relational inference over latent structures. These functions align with the theoretical perspective that PFC flexibly routes information in the service of current goals (Miller and Cohen, 2001), and it motivates our focus on mouse medial frontal cortex as a substrate for storing and re-using schematic knowledge to guide behaviour in new situations.

1.3.2 PFC representations for flexible behaviour

Clearly, perturbing the PFC degrades a range of functions needed for flexible behaviour. But what does PFC represent that affords these capacities in the first place? In contrast to the striking (broadly) single-variable selectivity observed in hippocampal–entorhinal spatial codes, PFC neurons typically exhibit high mixed selectivity, with individual units combining diverse information about stimuli, rules, context, actions, and outcomes in ways that can resist interpretation at the single-cell level (Tye et al., 2024; Rigotti et al., 2013; Dang et al., 2021).

PFC represents rules

A central ingredient of flexibility is the ability to represent abstract task rules: mappings that specify which stimulus features, responses, or outcomes are currently relevant. Classic single-unit work in primates showed that individual PFC neurons encode such rules at an abstract level. The same physical stimuli elicit different neural responses depending on the currently relevant rule, indicating that PFC activity implements a context-dependent gating of stimulus–response mappings (Wallis et al., 2001). In these tasks, neurons discriminated whether to apply, for example, a “match” versus “non-match” rule, generalising across specific stimuli and motor outputs. This established that rule information can be explicitly readable from single PFC units, despite the broader backdrop of mixed selectivity.

Beyond instantaneous rules for isolated trials, PFC also abstracts rules at the level of *action structure*. In monkeys planning four-step action sequences, lateral PFC neurons categorised whole behavioural sequences according to their temporal structure (e.g., push–push–pull–pull versus pull–pull–push–push), with selectivity emerging during planning, not just execution. This suggests that PFC can represent macro-structured action knowledge that organises forthcoming behaviour under a compact rule-like code (Shima et al., 2007).

In mice, longitudinal two-photon imaging demonstrates how rule codes are acquired and stabilised with learning. During rule-based visual categorisation, medial PFC neurons progressively develop category selectivity that supports generalisation to novel stimuli; subsets of cells are differentially recruited when the relevant rule later switches, consistent with a dynamic, reusable rule representation that can be reconfigured as task demands change (Reinert et al., 2021). In a similar task requiring context-dependent mappings, frontoparietal networks in humans (including PFC) adopt *orthogonalised* or factorised representational geometries, separating context/rule from stimulus dimensions so the correct decision boundary can be applied across conditions (Flesch et al., 2022).

Together, these data suggest that PFC not only encodes which rule is currently in play, but also shapes representational geometry in service of applying that rule and generalising it to new instances.

PFC represents problem structure

A body of work across species has suggested that explicitly representing the latent structure of tasks is crucial for flexible behaviour (Tse et al., 2007; Gershman and Niv, 2010; Swaminathan et al., 2023; Bernardi et al., 2020). This is a powerful representational format for generalisation to new scenarios that permits flexible behaviour, as new examples can be solved by merely binding state-tuned neurons with those representing the specifics of a behavioural sequence (Barone and Joseph, 1989; Whittington et al., 2022; Swaminathan et al., 2023). We will cover here examples of this phenomenon in frontal regions.

vmPFC and entorhinal cortex in humans build abstract representations of *problem structure* across superficially different RL problems, such that representational dissimilarity depends on the latent relational structure rather than concrete sensory particulars (Baram et al., 2021). A rich series of work in rats has pointed to OFC representing a problem schema (i.e., an internal model of the latent states and transitions that define a task’s structure) rather than encoding instantaneous stimulus–reward associations (Wilson et al., 2014). Crucially, these OFC state codes are not static. Across learning they evolve to better capture the schema of the task, aligning population geometry with the hidden structure that predicts outcomes, thereby supporting generalisation to new instances (Zhou et al., 2021). Paradoxically, recent work shows that once a schema is learned, the persistence of its representation in OFC is *positively* correlated with the acquisition of a new conflicting schema, rather than being a source of interference (Maor et al., 2025).

Work on task structure representation in frontal cortex in rodents extends across tasks to probabilistic reversal learning. Joint recordings in mFC and hippocampus show a complementary code, with hippocampal neurons coding for specific ports on a nose poke array while mFC neurons

generalise across ports from the same *task stage* (e.g., initiation vs choice port). Contrasting with this, in the odour sequence task both HPC and OFC neurons failed to distinguish states in the task when it was not necessary for behaviour (Zhou et al., 2019) – see also Morrissey et al. (2017). Both regions show signatures of the task schema, not a full hippocampal-style mapping of the full task state space. Instead of extracting the generalisable features from hippocampal representations through consolidation – see Sun et al. (2023) – recent evidence using optogenetic suppression of HPC output via ventral subiculum finds that OFC learns its schematic representations in parallel to HPC. In fact, HPC suppression actually accelerates OFC schema cell formation, which could indicate some compensatory mechanism or conflict between schema formation and detailed mapping of episodes (Zong et al., 2025).

Together, these findings support that under certain conditions, prefrontal regions can encode latent task states and generalise them to new scenarios. In a schema learning task, a prediction of this view is that we will see an explicit representation of sequence position that generalises across task examples. We will discuss this in detail in Chapter 4.

Progress towards goals

Maps of task structure should comprise neuronal dynamics that evolve as a function of *progress in a task*, independent of related variables such as elapsed time or the number of actions taken. A progress coordinate would naturally support generalisation across goal-directed sequences that vary in length and duration, because trajectories through the neural state space could be stretched or compressed without losing their alignment to the underlying behavioural objective.

Several lines of evidence from frontal circuits are consistent with this view. First, frontal ensembles carry sufficient information to recover latent behavioural variables like progress through a running trajectory without explicit measurements, suggesting that internal progress-like states are encoded in ongoing population activity (Rubin et al., 2019). Second, medial prefrontal neurons show anticipatory rate changes to predictable outcomes in spatial memory tasks (Pratt and Mizumori, 2001), and orbitofrontal populations map *future* navigational goals rather than

current position, indicating coding relative to upcoming subgoals independent of absolute distance or travel time (Basu et al., 2021). Third, prefrontal activity is organised into behaviourally meaningful sequences that can be replayed both during sleep and when task rules change. (Peyrache et al., 2009; Kaefer et al., 2020). Fourth, the mouse mFC ‘task stage’ code we discussed previously in a probabilistic reversal learning task can be recast as a progress code relative to a single goal, instead of discrete task states (Samborska et al., 2022).

Real-world complex behaviours, however, are rarely single-goal. They are hierarchical and composed of nested sub-goals and stages (Tenenbaum et al., 2011; Tomov et al., 2020). Extending the logic above implies that frontal population dynamics should track progress *within* and *between* subgoals, giving us a multi-scale progress representation that preserves invariances across differing durations and action counts at each level of the hierarchy. The above findings are consistent with such a progress code relative to individual goals but it remains an open question how neurons in frontal circuits integrate these codes to represent progress through complex, multi-goal tasks.

PFC encodes sequences

Many goal-directed behaviours are multi-goal and therefore inherently *sequential*. Understanding how prefrontal circuits represent, maintain, and manipulate sequences in general will therefore be crucial for parsing the specific case of representing sequential goals. Causally, medial PFC contributes specifically to order memory. Rats with mPFC lesions cannot remember the temporal order of odours in a list despite intact recognition of which odours were encountered, dissociating sequence coding from item familiarity (DeVito and Eichenbaum, 2011). This selective deficit points to prefrontal mechanisms that bind items to their positions in a sequence, rather than supporting item memory itself.

Recordings in nonhuman primates reveal how such bindings can be implemented in dorsolateral PFC (dlPFC). In a sequence working-memory (SWM) task requiring monkeys to remember and reproduce a list of spatial locations via saccades, dlPFC population activity during the delay

period contains a representation of the *entire* sequence, organised into near-orthogonal subspaces indexed by sequence position. Within each position-specific subspace, the code recovers the positions of the full spatial array (Xie et al., 2022). This representational geometry amounts to a systematic conjunction of *position* $\times$ *item*. The network carries, in parallel, a map for “first item,” a map for “second item,” and so on, with minimal interference between them.

Crucially, these subspaces are not static and are controllable by task context (Chen et al., 2024b). When task demands require it, dlPFC can recruit additional subspaces to support flexible reordering, maintaining a memory for the presented list while assembling an alternative output order (Tian et al., 2024).

Together, these findings outline a general set of principles for sequence memory in PFC. There can be high-dimensional mixed selectivity at the single-unit level, structured at the population level into position-labelled subspaces that each host a full map of items, with executive control signals selecting or recombining those subspaces as behaviour demands.

Rather than explicit inference over latent task states, this sequence-memory architecture offers an alternative method for generalisation across similarly-structured tasks with sequential goals. Concretely, a prefrontal network can *directly* store conjunctive position-by-item codes in its dynamics, without first inferring a latent state and then binding it to sensory specifics. Because the position axes are near-orthogonal, the same items (locations, actions, outcomes, etc.) can be represented in different positions (first, second, third, ...) with little interference. Further, because each position subspace tiles the entire item set, the scheme naturally accommodates changes in item identity without changes in synaptic weights.

This mechanism has importantly different neural predictions to those offered by explicitly representing sequence position. SWM-like algorithms would predict that mFC populations carry separable subspaces for locations in task-space coordinates. When a new task instance permutes where A–D are in space, the network need not relearn a task state $\times$ goal mapping.

It can reuse the same subspaces and simply populate them with the new item identities. In sum, PFC SWM codes provide a dynamical mechanism by which prefrontal circuits can store and recombine structured behaviours, supporting schema-like generalisation without requiring explicit new bindings for each task variant.

Summary

Frontal cortex, across species, carries powerful neural representations necessary for flexible behaviour: transition structures, goals and progress towards them, rules, and sequences. A combination or modification of these representational formats ought to be able to support generalisation of behavioural structure.

1.3.3 A note on "PFC" anatomy in rodents

In the current work, we are recording from frontal cortex of mice performing a schema learning task. There is little consensus about whether rodents have a 'prefrontal' cortex, and if they do, what structures that label refers to (Laubach et al., 2018). Classically, PFC was defined based on cytoarchitecture, following Brodmann's observation of a distinct, granular layer IV in anterior frontal regions of the primate brain (Brodmann, 1909). Under this strict criterion, many non-primate mammals, whose frontal cortices are largely agranular, were argued to lack a true PFC. This view shaped early comparative work and fostered the claim that rodents, in particular, do not possess prefrontal cortex (Preuss, 1995; Passingham, 2002). A parallel line of work instead defined PFC by its thalamic connectivity. Rose and Woolsey (1948) proposed that PFC is the principal cortical target of the mediodorsal (MD) thalamic nucleus, a definition later elaborated in anatomical studies (Leonard, 1969; Guldin et al., 1981). Under this MD-projection criterion, rodents do have candidate prefrontal fields, since medial frontal areas (e.g. prelimbic, infralimbic, anterior cingulate) receive dense MD input. This broadened the taxonomic scope of PFC beyond primates to other mammalian species. However, neither definition straightforwardly establishes homology. A granular layer IV is not a necessary or sufficient marker of functional equivalence across species, and MD connectivity is likewise

insufficient to guarantee one-to-one area correspondences.

Our frontal cortex recordings are primarily from prelimbic and infralimbic regions, with some recording sites located in anterior cingulate and M2. To avoid making too strong a claim about the existence of PFC in rodents, we will be referring to these regions collectively as medial *frontal* cortex rather than medial *prefrontal* cortex.

1.3.4 An aside: Mapping variables with attractors

A central aim of this thesis is to identify circuit mechanisms that map the structure of behaviour into neural activity. *Attractor* models offer a useful language for discussing such structure in recurrent networks (Khona and Fiete, 2022). An appreciation for various classes of attractor may help build an intuition for some of the results in this thesis. A detailed discussion of these models falls outside the remit of this introductory chapter. Those interested in building this intuition can turn to an appendix (Chapter 11) for an overview of attractor networks, examples of things in the brain that look like them, and how they could plausibly be used for mapping behavioural structure. I should be careful to note that I am not asserting that the neural results in this thesis arise due to literal mathematical attractors. However, I believe the framework is valuable as an intuition pump for some general principles of having an approximately-stable structure to neural representations that mirrors the structure of a real-world variable, which attractor networks can provide.

1.4 Two algorithms for structure learning: encoding memories in weights or activations

Intelligent behaviour hinges on the ability to predict the consequences of actions and to retrieve the right memory at the right time. A useful way to frame this problem computationally is through cognitive maps learned from sequential experience, which bind latent positions to observations so that behaviour can be flexibly guided in novel situations (e.g. Whittington et al. (2020)). In recent work, Whittington et al. (2025) outline two complementary but quite distinct algorithmic solutions to this problem, one aligned with factorised hippocampal-entorhinal codes,

the other with prefrontal working memory (WM). The two have a mathematical equivalence, and suggest that the same computation can be implemented either by storing memories in *synapses* or in *activity*. But, there are distinct representational signatures and control mechanisms in each case that we should be able to parse in a carefully designed schema learning paradigm.

1.4.1 The factorised solution: storing memories in synapses with an abstract position index

In the factorised solution, sequential experience is converted into a representation of latent position (ordinal, spatial, or abstract) maintained by a recurrent network (putatively *entorhinal*) that integrates velocity or action signals. This position code serves as an index into a set of synaptic memories (for instance, in a Hopfield network) that bind observations to positions via learned weights. While this mechanism is primarily associated with the hippocampal formation, there is evidence for frontal regions performing the kind of latent state tracking performed by the entorhinal RNN in this formulation (Samborska et al., 2022; Zhou et al., 2021; Baram et al., 2021).

A direct consequence of this architecture would be a robust, abstract position code that generalises across problems. Because the index is independent of the specific observation identities, pairs of "position cells" would preserve their phase relationships when a task is re-instantiated with different item $\times$ state bindings (an analogy to grid-cell phase stability). This stability offers a practical representational prediction of factorised solutions in a schema learning paradigm.

1.4.2 Working-memory solution: storing memories in activity via structured, controllable slots

Rearranging the same computation gives us a WM solution in which memories are kept in neural *activity* rather than synapses. Here, the recurrent state decomposes into structured, largely non-overlapping subspaces ("activity slots"), each capable of holding an observation. Fixed readout weights attend to a designated output slot. The role of recurrence is to *copy and shift* slot contents so that the correct item arrives in the readout slot at the right time. Velocity-dependent recurrent matrices (effectively expanded versions of the *item* and *position* representations in the factorised

solution) implement this controlled routing of content across slots. Once the mechanism is learned, the system can solve novel problem realisations without further synaptic plasticity. The slots are flexible, while readout is fixed.

Unlike the factorised scheme – where binding new items to positions typically requires new synaptic assignments – the WM solution naturally conjuncts item and position in activity, so no weight changes are needed to place different items into slots. Echoing this idea, recent computational work models PFC as a recurrent network that *meta-learns* across families of related tasks. After training, the RNN can acquire novel problem instances even with weight updates turned off, indicating that the learning algorithm itself is embedded in recurrent dynamics (Wang et al., 2018). This is exactly in the spirit of the activity-slot WM mechanism.

This activity-slot mechanism leads to a distinct code. Rather than representing a single abstract position, the WM network maintains separate subspaces (possibly, but not necessarily, localised to distinct populations) for each abstract position, with each subspace able to store a specific item in memory. The joint activity of the subspaces can therefore encode item $\times$ position pairs simultaneously. For an example in a sequence working memory task, monkeys are asked to remember the order of three cued locations on a screen before recalling them with saccades. During a delay period, subspaces for sequence position 1, 2, and 3 store the concrete locations associated with those sequence positions all at once (Xie et al., 2022; Tian et al., 2024; Chen et al., 2024b). Because slot identity carries the relative position label, single-unit relationships need not be preserved across task instances that reshuffle observation $\times$ state pairings. Accordingly, WM models tend to show strong decodability of slot sequences and weak abstract position coding.

1.4.3 Link to the current work

The factorised and WM algorithms offer two possible solutions to generalising a task schema. In a factorised solution, frontal or entorhinal cortices would contain an absolute, abstract, and generalisable representation of sequence position that indexes a memory of concrete goal locations elsewhere in the brain (likely the hippocampus). By contrast, a WM-like solution

would contain no such abstract representation. Instead, neurons would carry a conjunctive representation of location $\times$ sequence position, much like the mechanism for immediate serial recall proposed by Botvinick and Plaut (2006). Importantly, the 4-state sequence structure would still be present, but embedded in the network’s recurrent connectivity rather than expressed by individual neurons.

Here, we will test for neural signatures of both the WM and factorised algorithms in frontal and entorhinal cortices of mice performing a schema learning task.

1.5 Entorhinal cortex and factorised representations

In the previous section we introduced two complementary algorithms for generalising the structure of problems: a factorised code, and dynamics-based sequence working memory code.

Entorhinal cortices are well positioned to support a factorised representation. Computational work has posited that in the medial temporal lobe episodic memory system, an entorhinal RNN tracks position in a structured environment and controls an external memory system in the hippocampus, essentially telling the hippocampus which memory to recall, or store, and when (George et al., 2021; Whittington et al., 2020). Under this regime, medial entorhinal cortex provides ‘where’ information (not strictly a spatial ‘where’, but a position on an arbitrary graph, including in principle a task structure), and lateral entorhinal provides ‘what’ information. Implementing this in artificial neural networks recovers many famous single-cell representations found in the medial temporal lobe including place cells, grid cells, landmark cells, splitter cells, among others. Crucially, these representations fall out when the transition structure of the environment is a spatial one. The entorhinal RNN component can track position on any graph structure, lending the regime nicely to a multi-goal task structure: the entorhinal RNN would track position in a task sequence.

We will now look briefly at the cases that can be made for LEC and MEC being candidate brain regions supporting a factorised representation of task structure.

1.5.1 Candidate 1: Medial Entorhinal Cortex

Let's first consider medial entorhinal cortex. There are many reasons to suspect that MEC might hold an explicit representation of task structure. First, it is densely and bi-directionally connected to various subregions of the hippocampus and its surroundings (van Groen et al., 2003; Witter et al., 2017; Canto et al., 2008; van Strien et al., 2009): layer 2/3 projecting to dentate gyrus CA1, CA3 and subiculum; and layers 5/6 receive hippocampal outputs and project to broad neocortex. Second, population activity in grid-cell modules forms a low-dimensional manifold with 2D periodic boundary conditions (a torus), providing a biologically grounded continuous-attractor-(like) substrate to support a periodic task structure (Gardner et al., 2022). Although, whether such a system could simultaneously map task structure *and* space simultaneously is questionable. While it is famous for its spatial representations, the hippocampal-entorhinal system can map non-spatial task dimensions when those dimensions provide a stable metric. In rodents, MEC and hippocampus represent an auditory frequency axis as a navigable space (Aronov et al., 2017). In humans, entorhinal cortex carries grid-like codes for abstract conceptual relations (Constantinescu et al., 2016; Bao et al., 2019) and learned graph structure (Garvert et al., 2017). Moreover, during reinforcement learning, human entorhinal and vmPFC appear to encode and generalise latent task structure across problems with different sets of cues, consistent with a reusable schematic map (Baram et al., 2021).

1.5.2 Candidate 2: Lateral Entorhinal Cortex

There has been a rich and fertile research stream surrounding MEC, due in large part to its striking and interpretable spatial representations. By comparison, lateral entorhinal cortex has been fairly negLECTed*. We will briefly summarise the features of LEC that make it a strong candidate for representing a factorised task schema.

First, anatomically, LEC sits at an ideal place to bind abstract task structure to spatial specifics. Superficial LEC projects to hippocampal dentate gyrus/CA3 and to CA1, while deep layers

*Stylisation enforced by Dr Mohamady El-Gaby

of LEC receive hippocampal output – CA1/subiculum (Witter et al., 2017; Canto et al., 2008; Amaral et al., 2014). Further, LEC receives inputs from perirhinal and olfactory/piriform cortices – for object and associative content (Nilssen et al., 2019) – and has monosynaptic prefrontal afferents, including from the prelimbic/mFC areas (Hisey et al., 2023), although most prefrontal influence likely reaches LEC via perirhinal/postrhinal and thalamic routes (Hwang et al., 2017). This positioning between mFC and hippocampus makes LEC a plausible interface where a task schema could be held and bound to the concrete goals.

Regarding its representations, LEC carries a rich, conjunctive code for *what happened where and when*. First, it exhibits spatially modulated firing that is generally weaker and more diffuse than hippocampal place coding (Wang et al., 2024; Huang et al., 2023; Deshmukh and Knierim, 2011; Keene et al., 2016), yet remaps robustly across contexts within the same arena (Huang et al., 2023), consistent with a context-sensitive spatial representation. Second, LEC neurons are strongly object-related: cells develop fields anchored to objects, shift when objects are moved, and some retain trace activity at previous object locations (Tsao et al., 2013), linking present experience to recent spatial history. Third, LEC populations carry a pronounced temporal signal in both individual neurons and broad population activity (Wang et al., 2024; Montchal et al., 2019; Tsao et al., 2018, 2022; Kanter et al., 2025; Issa et al., 2024). Individual neurons track experiential and elapsed time within and across trial epochs (e.g. ramping before or after event boundaries such as reward), effectively segmenting experience into pre-event, event, and post-event phases (Issa et al., 2024). In humans, LEC has been shown to encode temporal *event* structure (Bellmund et al., 2019). Furthermore, in the absence of behavioural demands, LEC population activity drifts slowly along a single latent dimension over time, with discrete shifts at event boundaries (Kanter et al., 2025).

Coordinated activity between LEC and various prefrontal subregions has been shown to be crucial for various forms of associative memory (Kinnavane et al., 2025), with LEC ablation disrupting the stability of mPFC stimulus sequence (Pilkiw et al., 2022; Yu et al., 2021), and deep LEC (layer 5/6) inhibition prohibiting mPFC from splitting stimuli by their behavioural salience

(Jun et al., 2024).

The most promising recent evidence for LEC holding an explicit factorised task representation comes from an experiment in which a rat is exposed to multiple sequences of odours throughout a day (e.g. odour A, B, C, D; odour S, T, U, V; odour W, X, Y, Z). Population activity in LEC partitions the representations of this odour sequence into three parallel subspaces, with one dimension separating the sequences by their sensory specifics, with another respecting sequence position (Kanter et al., 2025). Such a coding scheme could easily map onto multiple instances of tasks with a similar abstract multi-goal structure.

In sum, two competing hypotheses motivate our interest in entorhinal cortices:

1) In Tolman-Eichenbaum-Machine style accounts, MEC provides a reusable map of latent structure, LEC inputs processed sensory information about objects, and hippocampus binds these two (Whittington et al., 2020). The prediction from this body of work is therefore that we should find a factorised task schema in MEC.

2) Recent empirical work emphasises LEC’s robust structured temporal coding (Tsao et al., 2018), shifts at event boundaries (Issa et al., 2024; Kanter et al., 2025), and sequence position (Bellmund et al., 2019), in a manner that generalises across sensory realisations (Kanter et al., 2025). A prediction of these findings, if they extend to multiple elements in a structured sequence, is that LEC should hold a factorised task schema.

1.6 Summary

The above discussion highlights that PFC is crucial for flexible behaviour. On top of that, we know that PFC, across species, has representations for abstractions, schemas, rules and sequences. What remains unclear are the *algorithms that use these representations* to rapidly

generalise behavioural structure to new, but structurally related, situations.

For spatial computations in the hippocampal-entorhinal network, we have a comparatively mature algorithmic understanding. We have a firmer grasp on the algorithms that use its various spatial representations for functions like path integration and generalisation of spatial structure across environments. What is less clear is whether the same algorithms that map and compute over space can be extended to *task structures*.

The lofty aims of this thesis are to (i) bring the algorithmic understanding of task structure abstraction in both PFC and MTL up to that achieved for spatial navigation in MTL, and (ii) test whether the computations for mapping space in entorhinal cortex can extend to the structure of tasks. To achieve this, our overarching approach is to explore both spatial and task representations in PFC and entorhinal cortices using *the same behavioural paradigm*. We want to know both “what is represented?” and “how are those representations used computationally?”

We approach this problem in the following way:

1. In Chapter 3 we aim to develop a behavioural paradigm that orthogonalises the structure of the task and the structure of the world, thus enabling us to study the two components separately.
2. Next – in Chapter 4, Chapter 5, and Chapter 6 – we aim to test for signatures of a factorisation-based algorithm (Whittington et al., 2020, 2022) or a dynamics-based sequence working memory algorithm (Wang et al., 2018; Xie et al., 2022).
3. In Chapter 7, we aim to find out whether PFC has a spatial model, and if so, in what form.
4. Finally, in Chapter 8, we aim to look for signatures of the factorisation-based algorithm in both lateral and medial entorhinal cortices.

The results in this thesis begin to clarify what is represented, where it is represented, and what algorithms are being implemented that support flexible intelligent behaviour.

2

General methods

Contents

2.1 Mice	35
2.2 Behavioural training	36
2.3 Surgery	41
2.4 Recording and preprocessing	43
2.5 Behavioural scoring	45
2.6 Tuning to task variables	47
2.7 Task space generalisation and modularity	49
2.8 Mnemonic fields in task space	51
2.9 Predicting distal choices	54
2.10 Activity structure in sleep	55
2.11 Generalised spatial tuning	57

2.1 Mice

Adult male C57BL/6J mice (Charles River Laboratories) were used. Animals were maintained in a dedicated housing facility on a 12:12 h light/dark cycle (lights on 07:00–19:00) at 20–24 °C and 55 % $\pm$ 10 % relative humidity.

For the mFC experiments, mice were organised into four cohorts ($n = 13$; Cohort sizes: 4, 4, 3, and 2), with inclusion determined by the criteria described under *Behavioural training*. For the

entorhinal recordings, we studied one cohort targeting medial entorhinal cortex (MEC; $n = 4$ from which we show neural data for $n = 1$) and one cohort targeting lateral entorhinal cortex (LEC; $n = 5$). No a priori statistical method was used to determine sample sizes.

For mFC cohorts, mice were group-housed with littermates until the start of the experiment with free access to water. In the entorhinal study, mice remained group-housed with littermates throughout behavioural training (typically 3–4 weeks) and were implanted and singly housed thereafter. Food was available *ad libitum* throughout; water was available *ad libitum* prior to experimental procedures. Mice were 2–5 months old at the time of testing. All procedures conformed to University of Oxford animal-use guidelines and were performed under UK Home Office Project Licence P6F11BC25.

2.2 Behavioural training

Overview of the ABCD task

Animals performed a family of tasks in a 3×3 grid maze in which four rewarded locations (labelled a, b, c, d) had to be visited in a fixed order that repeated cyclically. A reward became available at location b only after the animal collected the reward at a , and likewise for c and d . Completion of d reinstated availability at a , forming a repeating loop. On every trial, delivery of reward at a was accompanied by a brief tone, marking the start of state A and aligning the beginning of the loop across different task instances.

Spatial layout (Location)

The maze comprised nine circular platforms (“nodes”; numbered 1–9; see Figure) arranged in a 3×3 grid, connected by narrow bridges (“edges”). Connectivity was restricted to the cardinal directions; diagonal links were absent. Thus, for example, node 1 connected only to nodes 2 and 4 (see photograph and schematic in Figure 3.1, a).

Task instance (Task)

A specific task instance is defined by the ordered sequence of rewarded nodes for the ABCD loop (e.g., 1–9–5–7; Figure 3.1, b), such that rewards a , b , c , d are at nodes 1, 9, 5 and 7, respectively.

Session

A session was a continuous 20 min block of the same task. Mice could complete as many trials as possible within this period. Multiple sessions of the same task could be run on the same day. Between sessions, animals were removed from the maze and placed either in the home-cage or an enclosed rest/sleep box.

Trial

A trial comprised one complete pass through the ABCD loop for a given task, beginning with reward at location a and ending at the subsequent collection of reward at a . For instance, in a 1–9–5–7 task, trial 12 spans from the 12th reward at node 1 to the 13th reward at node 1.

State

A state was the interval between consecutive rewards in the sequence. State A ran from reward a to reward b ; state B from b to c ; state C from c to d ; and state D from d back to a .

Transition

We use “transition” as a general term for the change from one rewarded location to the next (e.g., $a \rightarrow b$ or $c \rightarrow d$).

Goal progress

Goal progress quantified how far through a state the animal had advanced, expressed as the fraction of elapsed time between successive rewards. Unless otherwise specified, we discretised progress into early, intermediate, and late bins, each spanning one third of the state duration. For example, if 9 s elapsed between two rewards, each bin covered 3 s; if 15 s elapsed, each bin covered 5 s. Because the ABCD loop comprises four rewards, these progress bins recur four

times per trial (one set per state; Figure 4.2).

Choice

At each node, the animal could choose among the immediately adjacent nodes (two or more options, depending on location). For example, from node 1 the animal could move to node 2 or node 4 (Figure 3.1).

Apparatus

The inner maze dimensions were $45 \times 45 \times 30$ cm (length $\times$ width $\times$ height). Outer dimensions were $66 \times 66 \times 125$ cm. Instead of solid outer walls, the perimeter was enclosed with electromagnetic shielding curtains (Electro Smog Shielding, product 4260103664431) custom-cut to cover the maze exterior. Nodes were 11×11 cm; bridge length was 7 cm. Water reservoirs were positioned 80–85 cm above the floor (measured from the bottom of the syringe), filled to 30 ml. Full design files and materials are available at <https://github.com/pyControl/hardware/tree/master/GridMaze>.

Animal selection (Pre-selection)

We used 21 mice across six cohorts (13 mFC, 3 mEC, 5 LEC). For each cohort, 3–6 mice were selected from an initial pool of 10–20 based on:

1. Body weight > 22 g.
2. No overt stress during a 20 min unrewarded exposure to the maze (stress indicated by thigmotaxis or defecation).
3. Habituation sessions in which mice learned that nose-poking at wells produced a water drop and that, after collecting a reward, they had to move to a different node. The number of accessible nodes was fixed within a session, and animals achieving ≥ 40 rewards in 20 min advanced to pretraining (more details below).

4. If more than four animals met these criteria in a cohort, we selected those with the highest exploration entropy (see “Behavioural scoring”).

Habituation and pre-training

Following at least one week of post-surgical recovery (see “Surgeries”), animals were placed on water restriction and maintained at 88–92% of their post-operative baseline body weight to maintain motivation without excessive drive which can impair model-based behaviour (Tolman, 1948). Mice were habituated to tethering for electrophysiological recording both on the maze and in sleep boxes for at least three days before experiments commenced.

During this period, mice performed sessions on a partially connected maze (5–7 accessible nodes; not all connected), while tethered. Available nodes were pseudo-randomly connected to guarantee access to all accessible nodes while preserving dead-ends; identity, number, and connectivity of accessible nodes changed each session to minimise spatial biases. A water drop was delivered upon nose-poke at any node, but after receiving a drop at one node, a subsequent reward required visiting a different node (matching step 3 in pre-selection, now under implant+tether conditions). Animals transitioned to the main task after earning > 40 rewards in a 20 min session. Three mice in cohort 1 were not implanted. For these animals, rewarded sessions during pre-selection and pre-training were identical. Pre-training reward volumes were typically 10–15 μ l (larger than during training).

Across pre-selection and pre-training combined, animals completed a minimum of two and a maximum of eight 20 min sessions, with a distinct connectivity configuration in each session.

Training procedure

Behaviour was controlled with pyControl (Akam et al., 2022). In each 20 min session, water rewards (3–4 μ l) were available sequentially at four nodes. A reward was triggered when the animal broke an infrared beam inside the well with a nose-poke. After collecting *a*, the next available reward was at *b*, then *c*, then *d*, before looping back to *a*. No time-outs were imposed

for incorrect choices, and animals were allowed to collect as many rewards as possible.

Mice were given at least a 20 min water-free interval between sessions. To avoid side preferences and reduce reliance on fixed orienting cues, the entry side for each session was pseudo-randomised via the shielding curtains (Electro Smog Shielding, product 4260103664431). A single cue card was placed high on one corner of the surrounding structure (≥ 50 cm above the maze); no low-level cues were visible.

Although all rewarded locations were identical in water delivery, a pure tone (5 kHz, 2 s) was played upon consumption at *a* to mark the onset of state A uniformly across tasks. Continuous white noise was presented to mask external sounds.

Task-set design and counterbalancing

For each mouse, task sequences (ABCD location orders) were selected pseudo-randomly under the following constraints:

1. Physical distance between consecutive rewards (number of steps on the maze) and “task distance” (number of states between two rewards) were orthogonalised (Figure 3.2).
2. The task could not be solved (i.e., $\geq 75\%$ optimal performance) by simply circling clockwise or anticlockwise around the maze.
3. For the first two tasks only, the central node (node 5) was included as a rewarded location to prevent trivial outer-ring strategies. (This constraint did not apply to later tasks or to the MEC/LEC cohorts.)
4. Consecutive tasks did not share a transition (i.e., no pair of adjacent rewarded locations was repeated across successive tasks). (Not enforced for MEC/LEC cohorts.)
5. Chance performance was matched across $a \rightarrow b$, $b \rightarrow c$, $c \rightarrow d$, $d \rightarrow a$ transitions and control transitions ($c \rightarrow a$, $b \rightarrow a$), computed either analytically (diffusive movement model) or empirically from each animal’s unrewarded exploration session (see “Behavioural scoring”).

Across the first 10 tasks, mice advanced to a new task when they achieved $\geq 70\%$ optimal transitions (i.e., selecting one of the shortest paths between rewards) over ≥ 10 consecutive trials, or if performance plateaued for ≥ 200 trials. During this phase, a maximum of four sessions per day were run, either all on the same task or split between the current and next task once criteria were met (two sessions each). From task 11 onward, mice learned three new tasks per day with the first task repeated at day's end (daily pattern: X–Y–Z–X'; four sessions total).

To test for unintended structure, we computed cross-task correlations within each mouse, comparing each sequence against all cyclic shifts of the others to exhaustively capture sequence similarity; the average correlation across shifts was reported in Figure 3.2. No two tasks presented to a given animal were rotations of one another in physical space. Task orders were randomised and counterbalanced across the experiment.

2.3 Surgery

General procedures.

Animals were removed from water restriction 48 h prior to surgery. Anaesthesia was induced with isoflurane (3 %) and maintained at 0.5–1 % throughout. Pre-emptive analgesia was administered (buprenorphine, 0.1 mg kg^{-1} ; meloxicam, 5 mg kg^{-1}). Mice were then secured in a stereotaxic frame for implantation.

Frontal cortex implantations.

A silicon probe mounted on a microdrive (Ronal Tool) and enclosed within a custom recoverable housing (ProtoLabs) was targeted to medial frontal cortex at anterior–posterior (AP) +2.00 mm, medial–lateral (ML) –0.40 mm relative to bregma, and dorsal–ventral (DV) –1.0 mm relative to the cortical surface. A ground screw was placed over the cerebellum. Cohorts 1–3 were implanted with Cambridge Neurotech F-series probes (64 channels, 6 shanks) oriented along the AP axis. Cohort 4 received single-shank Neuropixels 1.0 probes at the same AP/ML coordinates, lowered to DV –3.0 mm to utilise the depth span of the lower channel bank.

Entorhinal cortex implantations.

MEC: Single-shank Neuropixels 1.0 probes were placed at AP -4.6 mm, ML -3.3 mm (relative to bregma). DV depth was set by lowering the shank until slight deformation indicated contact with the ventral skull surface, then retracting by $50\text{ }\mu\text{m}$. Owing to the enclosure geometry for these implants, the ground screw was positioned over barrel cortex rather than cerebellum. *LEC*: Four-shank Neuropixels 2.0 probes were targeted at AP -3.5 mm, ML -3.5 mm and advanced to DV -4.0 mm from the brain surface at a 10° lateral approach. The probe was aligned so that the shanks spanned the AP axis. Owing to the enclosure geometry for these implants, the ground screw was positioned over barrel cortex rather than cerebellum.

Post-operative care and return to training.

Meloxicam was provided once daily for three days after surgery. Animals were monitored closely for seven days and then reinstated on water restriction 24 h before pretraining commenced.

Histology and probe localisation.

At the end of experiments, animals were perfused and brains were fixed, sectioned, and imaged to identify probe trajectories. Probe locations were assigned to anatomical regions using HERBS (version 0.2.8) based on the histological images. For Neuropixels recordings, the full probe track was traced directly. For Cambridge Neurotech 6-shank probes, the pial entry point of each shank was marked and the shank tip estimated from the final DV position recorded post-surgery together with subsequent nanodrive lowering. For each mouse, these marks were registered to a standard template to derive anatomical labels.

For the mFC cohort, using this approach, 90.7 % of recorded units were localised within peri-mFC regions: 68.3 % prelimbic cortex, 11.3 % anterior cingulate cortex, 6.1 % infralimbic cortex, and 5.0 % secondary motor cortex (M2). Of the remaining 9.3 %, 4.8 % could not be assigned to a specific peri-mFC subdivision because they mapped to peri-mFC white matter in the template (likely reflecting differences between true boundaries and those in the standard atlas used by HERBS), 2.2 % were in dorsal peduncular nucleus, 1.1 % in striatum, 0.6 % in medial orbital

cortex, 0.3 % in lateral septal nucleus, and 0.3 % in olfactory cortex. All recorded neurons were included in analyses except where otherwise stated explicitly.

At the time of writing, the precise alignment of MEC and LEC probe tracks to a reference atlas has not been completed. Approximate anatomical locations are discussed in the relevant results chapter (Chapter 8).

2.4 Recording and preprocessing

Acquisition and synchronisation.

Neural signals were recorded at 30 kHz using: for Cambridge Neurotech mice, a 32-channel Intan RHD2132 amplifier (Intan Technologies) connected to an Open Ephys acquisition board; for Neuropixels mice, Neuropixels (or Neuropixels 2.0) head-stage connected to the Neuropixels PXIe acquisition module. Behavioural events, video, and electrophysiology were synchronised via TTL sync pulses generated by the pyControl system.

Spike sorting and unit inclusion.

Recordings were spike-sorted with Kilosort - versions 2.5 and 3 for mFC, version 4.0 for MEC and LEC (Pachitariu et al., 2016, 2024) - and manually curated in phy (<https://github.com/kwikteam/phy>). Clusters were classified as single units and retained for analysis when they:

1. displayed a characteristic extracellular waveform,
2. exhibited a clear refractory period in the autocorrelogram, and
3. were stable across time within the recording epoch(s).

Cross-day concatenation and curation criteria.

To increase the number of tasks per neuron, pairs of consecutive recording days were concatenated to form “double-day” datasets (yielding six tasks). For each concatenated pair, all raw binary files across the two days were concatenated, processed as a single dataset through

the standard Kilosort pipeline, and then manually curated. Units from these concatenated files were screened using the following criteria:

1. **Refractory contamination.** Violations within a 2 ms refractory period were required to be $\leq 10\%$ of the baseline bin, where the baseline is defined as the count in the 25 ms bin (the maximum of the displayed autocorrelogram in Kilosort/phy).
2. **Rate stability across sessions.** Units were excluded if the firing rate in ≥ 3 sessions fell below 20% of that unit's peak session firing rate.

Curation summary metrics.

- Mean fraction of spikes within 1 ms and 2 ms refractory bins, respectively: $0.033 \pm 0.003\%$ and $0.076 \pm 0.004\%$.
- Proportion of neurons with a maximal drop-off in the autocorrelogram occurring at ≥ 2 ms: 99.2%.
- Dropout rate (Kilosort "good" pre-curation vs. post-curation yield): 51.4% for concatenated double-days versus 29.7% for single-day datasets.

Assessing cross-day unit identity via goal-progress tuning.

Because goal-progress tuning is highly conserved in mFC neurons, we leveraged within-day versus across-day correlations of goal-progress tuning to assess whether concatenation captured the same units across days. For each neuron, we computed:

1. A strong relationship between within-day and across-day goal-progress correlations (Pearson: $n = 1,540$ neurons, $r = 0.88$, $P = 0.0$).
2. Indistinguishable correlation magnitudes within versus across days (within-day: 0.63 ± 0.01 ; across-day: 0.62 ± 0.01 ; Wilcoxon test: $n = 1,549$, $W = 567,550$, $P = 0.14$, d.f. = 1,540).
3. Retention of significant goal-progress tuning across days for the vast majority of tuned units (95.4%; two-proportions test: $n = 1,249$ neurons, $z = 45.2$, $P = 0.0$). Significance was determined by comparing each neuron's empirical goal-progress correlation to the 95th percentile of a circularly shifted null distribution.

Behavioural tracking and region-of-interest assignment.

For mFC cohorts, pose estimation was performed with DeepLabCut (v2.0; [Mathis et al. \(2018\)](#)). For MEC and LEC cohorts, SLEAP was used ([Pereira et al., 2022](#)). The (x, y) pixel position of the back of the head was extracted per frame and used to compute kinematic variables (speed, acceleration). Frame-wise positions were then mapped to maze regions using binary masks (nodes and edges) created in ImageJ, yielding an ROI label for each frame.

Data analysis

Software. All analyses were performed in Python 3 using custom code alongside the following libraries and versions: numpy (1.22.0), scipy (1.10.1), pandas (2.0.3), scikit-learn (1.3.2), matplotlib (3.7.3), and seaborn (0.13.2).

2.5 Behavioural scoring

Performance metrics.

Behavioural performance was quantified using:

1. **Shortest-path transitions:** the percentage of state transitions (e.g., $a \rightarrow b$) for which the animal took one of the shortest available routes.
2. **Shortest-path trials:** the percentage of complete trials (ABCD loops) where the concatenated route across all four transitions was globally shortest.
3. **Relative path distance (RPD):** for each trial, the total path length taken between successive rewards divided by the corresponding shortest possible length.

$$\text{RPD} = \frac{L_{\text{taken}}}{L_{\text{min}}},$$

where L_{taken} is the number of steps actually traversed and L_{min} is the minimal number of steps required for that trial's reward sequence. It should be noted that because the maze has no diagonal connections, a single error in the path will increase the number of steps by 2. For example, a single mistake on a trajectory of one step from node 1 to 2 will take two extra steps to correct, giving a relative path distance of 3.

Chance levels.

Chance benchmarks for shortest-path measures were computed either analytically or empirically:

1. **Analytical chance (random-diffusion model):** assuming a randomly diffusing agent, we computed the probability of moving from node X to node Y in N steps and thereby the probability of selecting a shortest path. This was used to ensure fair comparison of the DA transition with BA/CA in zero-shot analyses.
2. **Empirical chance from exploration:** for each animal, we constructed a location-to-location transition matrix from a pre-task exploration session (no rewards) and used these empirical transition statistics to estimate the probability of shortest-path choices. This was applied to fairly compare DA with BA/CA in zero-shot quantification.

Chance level for the RPD measure. For relative path distance, chance was defined as the mean first-trial RPD averaged across the first five tasks and across all animals.

Correct-transition entropy (stereotypy of optimal choices). For transitions with multiple equally short routes, we quantified the distribution over the set of shortest paths and computed entropy to index stereotypy. $\mathcal{P} = \{1, \dots, M\}$ denotes the set of M shortest routes between two successive rewards, and p_k denotes the observed probability of selecting route k . We report the base- M normalised entropy:

$$H = - \sum_{k=1}^M p_k \log_M p_k.$$

An entropy of $H = 1$ indicates no bias among shortest routes (maximal diversity), whereas $H = 0$ indicates exclusive use of a single shortest route (maximal stereotypy).

Pre-configured exploration biases and task performance. To assess how baseline exploration tendencies influence task performance, we computed, per mouse, the correlation between:

1. **Relative path distance** on each ABCD trial (ratio of taken to optimal path length), and

2. **Mean baseline probability** of all steps actually taken by the animal on that same trial (measured from an exploration session before exposure to any ABCD task)—that is, how likely was the animal to take this path on a given trial before task exposure?

This correlation was positive overall, indicating that when animals deviated from the optimal path (higher RPD), they tended to do so via steps that were already high-probability during unrewarded exploration—that is, errors reflect persistence of pre-existing behavioural biases.

2.6 Tuning to task variables

Activity normalisation

We defined a common task space for projecting neural activity. For each trial, we aligned activity to state boundaries: consumption of reward *a* marks the start of a row; the next *a* marks the start of the next row. We time-normalised each state to 90 bins, yielding 360 bins per trial (*a*→*b*, *b*→*c*, *c*→*d*, *d*→*a*). For each neuron we averaged firing within bins, producing an $n \times 360$ matrix (trials $\times$ bins), then averaged across trials. We smoothed the circular mean trace with a Gaussian kernel ($\sigma = 10^\circ$) by concatenating the mean trace to itself 3 $\times$, smoothing the whole concatenated array, and extracting the middle third. This was to avoid edge artefacts introduced by smoothing. To reflect the circularity of the ABCD task, we plotted the mean and SEM on polar axes.

Generalised linear model

To quantify tuning to task variables, we fit a linear model to each neuron and assessed significance with permutation tests. Our goal was to measure how consistently goal-progress and place tuning generalise across tasks and states. We used leave-one-out cross-validation across all task $\times$ state combinations: hold out one combination (e.g., task 2, state C) for testing and train on the rest. From the training data, we computed two rate maps per neuron: mean firing in five goal-progress bins, and mean firing at each of nine maze nodes (edges excluded because they are never visited in the earliest progress bin). In the held-out data we predicted firing rates using the corresponding rate map value for the current bin. This means that location and goal-progress have just one regressor each which tests for consistency, rather than having a

regressor per location. Further, we included regressors for low-level kinematic variables (speed and acceleration), as well as time-from-reward and distance-from-reward. We repeated this for all folds and obtained a regression coefficient per variable and fold.

A neuron was only deemed to be significantly tuned to a variable if it passed both of the following criteria:

1. Regression coefficient exceeding the 95th percentile of a null distribution generated by running the regression with 100 circular-shift permutations of the neuron's activity.
2. Cross-task correlations of activity maps (place: 9-vector; progress: 5-vector) significantly greater than 0 (two-sided one-sample t-test). To ensure power, we analysed double-day data (up to six unique tasks).

We used two-proportion z-tests to ask whether the population fraction of significant neurons exceeded the 5% chance level.

State tuning

We assessed within-task state selectivity independently of the GLM (which tests cross-task consistency). For each trial, we took the peak firing rate in each state (A–D), giving us an $n \times 4$ max-activity matrix (trials $\times$ states). We then z-scored each row (transforming the data so that mean = 0, s.d. = 1) to remove between-trial fluctuations. For each neuron, we identified the preferred state in that task and took its z-scores across trials. We defined state tuning as a significant deviation of this distribution of preferred-state z-score from zero (as determined by a two-sided t-test). Unless stated otherwise, the threshold was $P < 0.05$; for stricter subsets we used $P < 0.01$. These criteria were used when quantifying cross-task generalisation and pairwise coherence, and for testing anchoring invariance in lagged task-space analyses.

UMAP manifold analysis

We visualised within-task population structure with Uniform Manifold Approximation and Projection (UMAP), a technique for non-linear dimensionality reduction. For this we included only concatenated double-days with (i) at least six tasks and (ii) at least ten simultaneously recorded neurons. We used as inputs each neuron’s z-scored, averaged and time-distorted activity (see the previous section on Activity Normalisation). We get from this an $n \times 360$ matrix per task ($4 \text{ states} \times 90 \text{ bins}$). We embedded this matrix into 3 dimensions using the cosine distance metric, $n_neighbours = 50$, and $min_dist = 0.6$.

Colouring the embedding by goal-progress or by state revealed a hierarchical geometry: progress sequences form loops that cluster by state, producing a four-petal floral pattern. To quantify this per double-day, we measured distances in the low-dimensional space between: (1) across-state, opposite progress bins (“across-progress”); (2) across-state, same progress bins (“within-progress”); and (3) a shuffled control with state labels randomised across trials. These embeddings are intentionally within-task and do not show generalisation of the manifold structure across task instances.

2.7 Task space generalisation and modularity

Inclusion criteria.

All analyses were restricted to state-tuned neurons. To ensure robustness, we further required neurons to be state-tuned in more than one-third of recorded tasks; this subset is used throughout where state-tuned cells are investigated.

Single-neuron remapping angle (best-rotation).

For each neuron and task pair (e.g. X vs. Y), we rotated the neuron’s firing-rate vector in task Y in 10° steps and computed the Pearson correlation with the mean firing-rate vector in task X. The rotation yielding the maximum correlation defined the remapping angle. Cross-task histograms pooled angles from X vs. Y and X vs. Z; within-task histograms compared X vs. X' .

To quantify generalisation, we took the *maximum* absolute rotation across the two cross-task comparisons (X vs. Y and X vs. Z) and classified a neuron as generalising if this value lay within $\pm 45^\circ$ of 0. Because $\pm 45^\circ$ spans one quarter of the circle, the chance level is $(\frac{1}{4})^{n_c}$, where n_c is the number of independent comparisons. Thus, chance is 25% for a single comparison and 1/16 (6.25%) for two comparisons. This definition is insensitive to goal-progress tuning strength, which can compress/stretch firing fields but does not alter the circular fraction.

Between-neuron coherence (pairwise angle).

We asked whether pairs of neurons preserve their *relative* spacing in task space across tasks. For each pair (j, k) and for each task pair, we rotated neuron j in 10° steps and computed the Pearson correlation with the mean firing vector of neuron k ; the rotation maximising the correlation defined the between-neuron angle. We then compared this angle from task X to the same pairwise angle in tasks Y and Z. As above, we used the maximum absolute rotation across the two cross-task comparisons and applied the same $\pm 45^\circ$ coherence threshold. Chance levels are again $(\frac{1}{4})^{n_c}$ (25% for one comparison; 6.25% for two).

Clustering into coherently rotating modules.

To assess modular organisation, we clustered neurons by their cross-task coherence relationships: *Step 1:* For each recording day, compute the maximum difference in pairwise between-neuron angles across all comparisons and convert to a circular distance, $d = 1 - \cos(\Delta\theta)$, yielding an incoherence matrix.

Concretely:

For neurons i and j and each comparison c (task pair), we define

$$\Delta^{(i,j,c)} = \text{wrap}(\phi^{(i,j)} - \psi^{(i,j)}).$$

Where $\phi^{(i,j)}$ is the angle between neurons i and j in task X, and $\psi^{(i,j)}$ is the angle between neurons i and j in task Y/Z. We convert this relative angle change into a circular distance:

$$d^{(i,j,c)} = 1 - \cos(\Delta^{(i,j,c)}).$$

Finally, we take the maximum across comparisons to obtain the incoherence matrix entry:

$$D^{(i,j)} = \max_c d^{(i,j,c)}.$$

Step 2: Embed this matrix using t-SNE (scikit-learn, TSNE, perplexity = 5).

Step 3: Cluster the embedded points with hierarchical agglomerative clustering (scikit-learn, AgglomerativeClustering, distance_threshold = 300).

Note that across three tasks there are two cross-task comparisons (X vs. Y, X vs. Z). In the trivial case with perfect goal-progress quantisation (exact 90° steps) and random remapping, one expects $4^2 = 16$ clusters; with imperfect quantisation, more clusters are expected. To avoid assumptions, we generated a null by preserving state tuning in task X and goal-progress tuning in all tasks while randomly permuting each neuron’s state preference across tasks.

For each day we computed the silhouette score,

$$s = \frac{b - a}{\max(a, b)},$$

where a is the mean intra-cluster distance and b is the mean distance to the nearest other cluster. We computed the same metric for null data to compare real vs. permuted clustering.

For illustration, we show example neurons from a single day with state tuning on the in-plane axes (e.g. tasks X and Y) and cluster membership encoded by the third dimension/order. Ordering within clusters is arbitrary.

2.8 Mnemonic fields in task space

The tSMB model predicts neurons that maintain an invariant task-space *lag* from a specific behavioural step (anchor; place $\times$ goal-progress), irrespective of task sequence. We tested this with three complementary analyses. Unless noted, all analyses used *state-tuned* neurons (see *State tuning*) and concatenated two recording days spike-sorted jointly, yielding six tasks per

animal (two animals had four and five tasks). We used concatenated days to increase power for cross-task generalisation: with six tasks, each of the nine locations is typically encountered in ~ 2.67 distinct task contexts (versus ~ 1.33 with three tasks), allowing cross-validated tests of anchor-lag invariance.

Method 1: Single-anchor alignment.

This method assumes each neuron has a single place $\times$ goal-progress anchor and asks whether its task-space lag relative to that anchor is conserved across tasks.

1. For each candidate anchor (3 goal-progress bins $\times$ 9 locations = 27), identify all visit times and extract 360 bins (one full trial) starting at each visit. For each task, align to anchor visits, average, and smooth (see *Activity normalisation*).
2. For a given train-test split (typically 5 training tasks, 1 test), compute pairwise angles θ between anchor-aligned mean firing vectors for all training-task pairs by rotating one vector in 10steps and taking the rotation maximising Pearson correlation.
3. Convert angles to distances $d = 1 - \cos \theta$ to form a distance tensor of size $N_{\text{train}} \times N_{\text{train}} \times 27$ (typically $5 \times 5 \times 27$). Average across pairs to obtain a mean-distance matrix M_{mean} of size $N_{\text{train}} - 1 \times 27$.
4. Select the (training-task, anchor) entry with the minimum distance as the best anchor and reference training task for that neuron. Align the left-out test task to this best anchor and compute the remapping angle (test vs. reference). Angles within $\pm 45^\circ$ of 0 indicate conserved anchoring.
5. For visualisation, histogram angles per split and average across splits. For effect size, compute the Pearson correlation between anchor-aligned activity in the test and reference tasks, restricting activity to each neuron's preferred goal-progress bin.
6. Define the neuron's *lag from anchor* as the angle at which the anchor-aligned activity is maximal in training; this lag is used below for behaviour prediction.

Method 2: Lagged spatial similarity.

If a neuron encodes a memory of a place $\times$ progress anchor, its cross-task *spatial* maps should correlate most strongly at a non-zero past lag. We are asking here if mFC neurons form spatial maps of *previously visited locations*. This method assumes a single lag, but not a single anchor.

1. For each neuron, and within its preferred goal-progress bin, compute spatial rate maps using positions from 12 past-lag bins spanning one full trial (4 states $\times$ 3 progress bins).
2. Using N_{train} tasks, find the lag that maximises cross-task spatial correlation. Apply this lag to the left-out task and measure the Pearson correlation between its spatial map and the training maps at that lag.
3. Note: Firing rates are always computed within the neuron's preferred goal-progress bin; positions are taken either at the present (0 or successively further back in time for the 12 lag bins).

Method 3: Model fitting (elastic net).

Here, we drop the assumptions about neurons having a single anchor or a single lag by modelling their tuning as a function of all anchors $\times$ lag conjunctions. We do this using regularised elastic net regression.

1. Regressors: 324 anchor $\times$ lag terms (9 locations $\times$ 3 goal-progress bins $\times$ 12 lags; 4 states $\times$ 3 progress bins).
2. Model: Elastic net (scikit-learn `ElasticNet`) with L1:L2 = 1:1 (`l1_ratio` = 0.5), α = 0.01. Train on five tasks; predict the left-out task. Training and validation use only the neuron's preferred goal-progress bin to avoid goal-progress confounds.
3. Non-zero-lag selections (when specified): include only neurons whose three largest coefficients occur at non-zero lags (thresholds: $\geq 30^\circ$ for 'all non-zero'; $\geq 90^\circ$ for 'distal non-zero'). For prediction, use coefficients at least 30° (or 90°) away from the anchor to avoid contamination by direct anchor tuning.

Excluding trajectory-tuned neurons.

To exclude neurons with trajectory tuning from the lagged-tuning analyses, we fit a separate GLM in which “place” was replaced by conjunctions of current place and next step (24 place→next-place categories). For each neuron, regressors were the mean firing rates for each category computed from the training task–state folds; these were used to predict binned firing in the held-out fold while co-regressing goal progress, speed, acceleration, time from reward, and distance from reward. We repeated this for all task–state folds and computed a regression coefficient per fold; the mean coefficient was then compared to a circular-shift null (spike trains shifted relative to behaviour) to obtain a 95th-percentile threshold. This single, intentionally permissive criterion was used to flag and exclude any neuron with even weak trajectory tuning.

2.9 Predicting distal choices

The SMB model predicts that “bumps” of activity along a buffer bias future choices at a precise lag from the buffer’s anchor. We therefore asked whether distal neural activity predicts upcoming choices with logistic regression.

Inclusion and anchoring.

We analysed only *consistently anchored* neurons. These are cells whose best anchor (place×goal-progress) and lag from that anchor were identical in at least half of the tasks (estimated with the single-anchor alignment; see *Lagged task-space tuning*, Method 1). To avoid contamination by direct spatial tuning, we excluded near-zero-lag activity. That is, we required at least $\pm 30^\circ$ separation from the anchor (one third of a state), and we repeated the analysis with a stricter $\pm 90^\circ$ criterion (one full state).

Decision opportunities and outcome.

A decision opportunity was defined whenever the animal was one state transition away from a neuron’s anchor (e.g. if the anchor is early progress at place 2, the preceding opportunities are late progress at places 1, 3, or 5). For each opportunity t , the binary outcome $y_t \in \{0, 1\}$ indicated whether the anchor was revisited on the next step ($y_t = 1$) or not ($y_t = 0$).

Neural predictors (bump time).

For each neuron, we measured *bump-time* activity at its anchor-relative lag X as the mean firing rate in a window spanning $[X, X+30^\circ]$ for ABCD (i.e. one third of a state). This value was entered once per decision opportunity.

Logistic regression model and history controls.

We first fit, per task, a logistic model with separate regressors for the previous 10 choices to control for behavioural autocorrelation:

$$\text{logit } P(y_t=1) = \beta_0 + \beta_{\text{bump}} r_t^{\text{bump}} + \sum_{k=1}^{10} \gamma_k c_{t-k},$$

where r_t^{bump} is bump-time firing and $c_{t-k} \in \{0, 1\}$ indicates whether the anchor was chosen k trials back. We then fit an exponential decay to the $\{\gamma_k\}$ and, in all subsequent regressions, replaced the 10 separate history terms with a single kernel-weighted history regressor $\tilde{c}_t = \sum_{k=1}^{10} w_k c_{t-k}$, keeping the same bump-time predictor.

Control times.

To test temporal specificity, we repeated the regression after replacing bump-time activity with: (i) *decision time* activity (30° before the potential anchor visit), (ii) *random time* activity (matched in number of samples), and (iii) activity at one-state offsets from bump time ($X + \{90^\circ, 180^\circ, 270^\circ\}$), which preserve goal-progress preference but disrupt the anchor-relative lag.

We repeated the procedure (i) restricting to distal neurons ($\geq 90^\circ$ from the anchor).

2.10 Activity structure in sleep

To probe whether task-shaped structure in mFC is internally organised (rather than purely evoked by sensory/behavioural correlates during task execution), we analysed neural activity during sleep/rest in an enclosure with home-cage bedding and no task cues. We asked two questions: does sequential structure persist offline, and does the offline state space behave like a *closed ring* or an *open delay line*?

Recording and session structure.

Animals were pre-habituated to “sleep boxes.” Each recording day comprised a 1 h pre-task sleep/rest session, 20 min inter-session sleep/rest sessions after each task, and a 30–45 min post-task session. All sessions except the first were designated “post-task.”

Binning and inclusion criteria.

Spike trains were binned at 250 ms. Analyses were restricted to *consistently anchored* neurons (cells whose best place $\times$ progress anchor and anchor-relative lag matched in at least half of tasks).

Awake-derived distances (state-space geometry).

For each pair of co-anchored neurons we computed two awake-task distances: *circular distance* on the putative ring (anchor-relative angular separation), and *forward distance* along the ABCD sequence (sequence-lag). Angles for pre-task sleep were taken from the first task of the day; for post-task sleep, from the immediately preceding task.

GLM: linking distances to offline coactivity.

Per sleep/rest session, we computed pairwise cross-correlations and regressed them onto circular and forward distances *in the same model*. To control for tuning similarity not specific to our putative state-space structures, we included spatial map correlation and circular goal-progress distance as co-regressors. We ran these regressions (a) within co-anchored pairs (putative *within-module*) and (b) across pairs with different anchors (*between-modules*).

Anchor specificity (modular signature).

To test whether offline organisation is modular, we compared the regression coefficients for *circular distance* between co-anchored and cross-anchor pairs. A stronger negative dependence within co-anchored pairs indicates ring-like structure preserved *within* modules.

Topology test via forward distance.

As a complementary, assumption-light test, we examined sleep cross-correlations as a function of forward distance for co-anchored pairs. A ring predicts a V-shape: correlations decrease with forward distance over 0° – 180° and increase over 180° – 360° (wrapping back). A delay line predicts a monotonic decrease. We therefore computed, per session, the correlation between forward distance and cross-correlation separately for the two ranges and compared them.

Together, the GLM (with controls) and the V-shape test assess whether the task-shaped, anchor-relative geometry inferred during behaviour is *intrinsically* maintained offline and whether its topology is circular rather than linear.

2.11 Generalised spatial tuning

To query the structure of spatial representations in mFC while controlling for lagged spatial influences, we used a neural-subspace approach that isolates the *present-location* code and compares its geometry across goal progress.

Data assembly and regressors.

For each recording *double-day* we pooled all awake (task) sessions and segmented activity into 250 ms bins. In each bin we extracted the animal's *current* location (one-hot over 9 nodes), *prior* locations at 20 task-space lags (5 lags per state across the 4-state loop; each lag encoded as a separate 9-dimensional one-hot vector), and speed and acceleration kinematic variables (both z-scored). Analyses were performed separately in **five** goal-progress bins.

Lasso regressions and split-half stability.

Within each goal-progress bin, we fit $L1$ -penalised *Lasso* regressions ($\alpha = 0.005$) predicting binned firing from the full design matrix (current location, $20 \times$ lagged locations, speed, acceleration). For stability, we ran five iterations. In each iteration, random 50/50 splits of time bins yielded two independent fits, producing β_1 and β_2 . We then averaged across iterations to obtain a mean β_1 matrix and a mean β_2 matrix *for each goal-progress bin and each recording day*.

Subspace construction.

Our subspace analyses focus on the *instantaneous spatial code*. Therefore, subspaces were constructed separately for each recording day using only the *current-location* coefficients from the corresponding β matrices (lagged-location coefficients remain in the model purely as controls). For each day and goal-progress bin, we performed PCA on the $9 \times (\text{neurons})$ block of current-location coefficients and retained the top three principal vectors to define a 3D “location subspace.”

For visualisation, we concatenated the day-wise current-location coefficient matrices across recording days and computed a global PCA, projecting onto the first three PCs to obtain nine points (one per node) in 3D.

Principal-angle comparisons.

We quantified subspace relationships using principal angles in two settings:

1. **Across goal progress (test):** compare the 3D subspaces derived from β_1 across different goal-progress bins (each fit in a separate regression), within the same recording day.
2. **Within goal progress (control):** compare β_1 vs. β_2 subspaces from the *same* goal-progress bin (split-half consistency), within the same recording day.

For any pair of conditions A and B , we say $U_A, U_B \in \mathbb{R}^{9 \times 3}$ contain the three principal vectors (PC1–PC3) spanning the current-location subspaces. We formed the inner-product matrix $M = U_A^\top U_B$, computed its SVD $M = W\Sigma V^\top$, and defined the principal angle as

$$\theta_1 = \arccos(\sigma_1),$$

where σ_1 is the first singular value of Σ . Repeating this for all relevant pairs yields day-wise distributions of principal angles for across-progress (test) and within-progress (control) comparisons.

Rationale.

This framework controls for lagged spatial tuning by explicitly modelling 20 lagged one-hot location sets, isolates a compact, day-specific low-dimensional spatial subspace for each goal-progress bin, and uses a principled angle-based metric to test whether present-location codes are stable or rotate across goal progress relative to a split-half baseline.

Representational similarity analysis (RSA) of the population place code

To probe how mFC population activity represents space, we performed an RSA across maze locations. We first applied a speed threshold to exclude stationary bins and tested for potential speed confounds across locations: an ANOVA on mean speed by node showed no significant differences, indicating that any representational structure is unlikely to be driven by systematic speed variation across places.

Data RDMs (per recording day).

For each recording day, we computed a mean population firing-rate vector for each of the nine maze locations (averaging across speed-filtered bins). We then constructed, for each goal-progress bin, a 9×9 representational dissimilarity matrix (RDM) using Euclidean distances between these location-specific population vectors:

$$D_{\text{data}}(i, j) = \|\mathbf{r}_i - \mathbf{r}_j\|_2 = \sqrt{\sum_{k=1}^p (r_{ik} - r_{jk})^2}.$$

Here, i and j index maze locations (nodes), $\mathbf{r}_i \in \mathbb{R}^p$ is the (speed-filtered) mean population firing-rate vector at location i across the p simultaneously recorded neurons, and r_{ik} denotes the firing rate of neuron k in that vector. The norm $\|\cdot\|_2$ is the Euclidean norm.

Model RDMs.

We compared the empirical RDMs to three hypothesis-motivated model RDMs:

1. **Geometric distance:** Euclidean distance between locations in maze coordinates.

2. **Node-degree difference:** absolute difference in graph degree (corner= 2, edge= 3, centre= 4) between location pairs.
3. **Place-type category:** a nuisance categorical RDM with entries 0 when two locations share the same place type (corner/cardinal/centre) and 1 otherwise.

Fitting procedure and nulls.

For each recording day, we vectorised the upper triangle of the data RDM and regressed it simultaneously onto the vectorised model RDMs (linear regression). We performed two complementary analyses: (i) *Aggregate* - we averaged data RDMs across days to obtain a mean RDM and fit the model, assessing each β against a row/column-wise permutation null (100 label shuffles of the data RDM, preserving symmetry). (ii) *Per-day* - we fit the model separately per day and tested the distribution of β weights across days against zero (two-sided t -tests).

Low-dimensional embedding of data RDM.

To visualise the representational geometry, we applied (non-metric) multidimensional scaling to the aggregate data RDM and projected locations into 2D. The resulting configuration revealed a clear node-degree structure, consistent with the regression results (e.g., graded separation aligned with centre $\rightarrow$ cardinal $\rightarrow$ corner).

Decoding analyses: concrete and generalised spatial code

To complement the subspace and RSA results, we asked whether mFC population activity supports *generalised* readouts of spatial structure using support vector machines (SVMs; radial-basis-function kernel, balanced class weightings, with regularisation parameter $C = 1$). All analyses used speed-filtered time bins. Balanced accuracy (BA) was used to aggregate across folds,

$$\text{BA} = \frac{1}{2} \left(\frac{\text{TP}}{\text{P}} + \frac{\text{TN}}{\text{N}} \right),$$

computed by pooling predictions and labels across folds within each recording double-day (chance = 0.5).

Held-out place type (corner vs. cardinal).

We trained an SVM to classify the *place type* of a location (corner vs. cardinal edge node; the centre was excluded). In each fold, a single node was held out of training entirely, and the trained model was asked to label that held-out node (8-fold CV over the eight peripheral nodes). The held-out test set contained only a single class per fold.

Held-out topology class (node vs. edge).

We repeated the analysis adding the maze edges, as well as nodes (9 nodes + 12 edges). Using leave-one-location-out CV, the SVM was trained to discriminate node vs. edge.

Cross-task location identity and structured errors.

Finally, we tested generalisation across tasks by training the SVM to predict node identity (absolute location) using data from all but one task and evaluating on the held-out task (leave-one-task-out). A one-vs-rest decision function with balanced class weighting was used in training the SVM. Confusion matrices showed a pronounced diagonal, indicating a stable absolute place code that transfers to a new task instance. To identify structure in the decoding errors, we removed the diagonal and node 5 (centre), renormalised each matrix, and compared misclassification rates to nodes of the same feature set (e.g. corner to corner misclassification) versus different feature set (e.g. corner to cardinal) within each recording double-day. Misclassifications were *significantly more likely* to fall on locations of the same degree than across degrees (statistical test across days), indicating that when the decoder errs, errors respect the maze's abstract geometry.

3

The ABCD task: a behaviour paradigm for task structure abstraction

Contents

3.1 The ABCD Task: a behavioural paradigm for probing task structure abstraction	64
3.2 Mice learn efficient routes between rewards, and acquire new examples more readily across learning	69
3.3 Inference of unseen task-space transitions: a strong test of task structure knowledge	73
3.4 Summary	75

In order to investigate the formation and use of abstract task structure representations (schemas), we need a behavioural paradigm that decouples the structure of the task from the structure of the environment in which that task takes place. Prior studies have used tasks that get to this in varying degrees. For instance, in alternating T-maze tasks in which animals have to either turn left or right upon reaching a junction, the task state of the environment in the corridor leading up to the junction is different in left trials vs right trials in terms of behavioural consequences, despite the spatial structure being identical. This state aliasing is reflected in hippocampal representations, where a subset of place cells, termed splitter cells, display consistent spatial tuning but only for trials for a specific future direction (Frank et al., 2000; Wood et al., 2000). A more extreme version of this can be found in a task where animals must complete multiple laps

of a square arena before receiving a reward. Again, hippocampal place cells display consistent spatial tuning, but with a subset of cells possessing receptive fields for a particular lap of the trajectory (Sun et al., 2020).

For our purposes, we need more flexibility in how task examples can be instantiated, ideally with: (i) uncorrelated distances between task and world states, allowing us to properly disentangle representations of the world from behaviour (ii) different and rich sets of behavioural policies needed to solve each task instance, confirming that representations we see are not due to low-level action sequences that correlate between examples, and (iii) multiple hierarchically-structured task states more closely resembling real-world complex behaviour than single-goal tasks.

3.1 The ABCD Task: a behavioural paradigm for probing task structure abstraction

Here we present the ABCD task: a behavioural paradigm in which multiple goals are hierarchically organised by an abstract task structure (Figure 3.1). The ABCD task involves knowledge of two structures. The first - the world structure - is a three x three grid maze arranged on a square lattice (Figure 3.1 a, right). The maze comprises nine towers (nodes) connected by bridges (edges), each with a water reward port in the centre that can be activated by the breaking of an infrared beam. The natural discretisation of the maze structure is reminiscent of gridworld environments often used in reinforcement learning, and allows the maze to be treated as a graph. The second structure - the task structure - is a repeating sequence of four uncued reward locations on that maze, a , b , c and d (Figure 3.1 a, left). Mice must learn to navigate efficiently between these reward locations, and encounter multiple task instances where the specific reward locations change. We define *task states* as the time period between collection of a pair of rewards. That is, task state A begins when the mouse collects reward at location a and ends when it collects reward at location b , which in turn marks the beginning of task state B , and so on. When the mouse collects reward at location d , it has to return to location a to begin the next trial. The onset of a new trial is cued by a brief two-second tone that sounds immediately upon collection of reward at location a , allowing the mice to learn an equivalence across trials in a way that does

not require a memory of the first reward location in the whole session. Sessions were 20-minutes long, with the mice being allowed to retrieve as many rewards as possible within that time limit. There were no penalties or time-outs for attempts at reward retrieval at incorrect locations.

Across task instantiations (referred to hereafter as simply 'tasks') the specific locations corresponding to those task states change pseudorandomly (Figure 3.1 b). For example, in task 1 reward *a* is in location 1, reward *b* is in location 9, reward *c* is in location 5, and reward *d* is in location 7. In task 2, the rewards for *a*, *b*, *c*, and *d* are in locations 5, 7, 3, and 9. There were some constraints on possible consecutive pairs of task configurations that are detailed in full in General Methods, with the most important being that a task cannot be solved by looping in one direction around the outer nodes due to mice's well-documented natural wall-hugging behaviour. Concretely, a task is considered to be solved by this degenerate solution, and hence not included, if an outside loop comprises >75% optimal (shortest) paths between reward location pairs.

To encourage mice to disentangle the task structure from the maze structure, we ensured that the two were orthogonal throughout the set of tasks (Figure 3.2). An example of where this would not be the case, and would therefore discourage the successful separation of the two representations neurally, would be if positions in the ABCD sequence unfolded only via adjacent location in the maze. In this case there would be a tight correlation between the two structures. Three example tasks and their optimal solutions can be seen in (Figure 3.2, a). Notably, task states differ by the number of steps in their optimal paths both within tasks (e.g. four steps for state *A* in task 1, and two steps for state *B* in task 1) as well as across tasks (e.g. four steps for state *A* in task 1 and just one step for state *A* in task 2). We quantified this in two ways. The first involved correlating task-space distances (how many task states are between the rewards) with physical distances between reward locations (number of steps taken to reach those rewards (Figure 3.2, b). Across the set of tasks, for all mice, there was no significant correlation between task-space and physical-space distances (Pearson correlation: $r=6.04 \times 10^{-18}$, $P=1.0$; one-way ANOVA statistic=2.02, $P=0.147$, $df=12$). The second measure involved correlating the reward




Figure 3.1: The ABCD task - a behaviour paradigm to probe task structure abstraction

(a) The ABCD task involves learning both an abstract task structure (a repeating sequence of four identical uncued rewards, ABCD), and a spatial maze structure (a three x three grid maze). A brief two-second tone marks the start of a new trial. **(b)** Across tasks the physical locations of the four rewards changes, while the abstract task structure, a four-state repeating sequence, stays the same.

sequences themselves (Figure 3.2, c). Here we show correlations in task sequences for three different partitions of the data: all days for all mice, all days for the mice from whom we recorded neural data, and all the 'late 3-task' days (tasks 20-40). The late 3-task days are the only days from which we will show neural data in the following chapters. Performing a T-test (two-sided) against 0 correlation for all tasks ($N=13$ mice, $r=0.002$, statistic=2.91, $P=0.013$, $df=12$) showed minimal cross-task correlation. The neural mice ($N=7$ mice, $r=0.002$, statistic=1.81, $P=0.120$, $df=6$) and late 3-task days from neural mice ($N=7$ mice, $r=4.5 \times 10^{-4}$, statistic=0.22, $P=0.634$, $df=6$) both showed no significant correlation in reward sequences across tasks. The small, albeit significant, correlation in the 'all tasks' condition is due to a criterion that forces the centre (node 5) to be rewarded on the first two tasks to combat thigmotaxis. These data are not included in subsequent analysis.

Taken together, these measures show that using the ABCD paradigm we have successfully disentangled task structure from world structure.




Figure 3.2: Decoupling the structure of the environment from the structure of the task

(a) Schematic showing three tasks with optimal policies for solving them demonstrating that optimal path lengths differed both within and across tasks. **(b)** Top - schematic of the notions of distance in task-space and physical space. The red arrows show that a distance of 3 in task space corresponds to a physical distance of 1. Bottom - there is no significant difference between distances in task space and physical distances across the set of tasks for all mice. **(c)** Cross-task correlations of reward sequences shows minimal correlation for all task days, and no significant correlation for neural mice and neural mice late 3-task days (from which the electrophysiology data shown in later chapters was collected).

3.2 Mice learn efficient routes between rewards, and acquire new examples more readily across learning

We sought to test the degree to which mice could learn to efficiently navigate between reward locations in the ABCD task. We first pre-selected mice using some minimal criteria. Mice retained after applying these criteria moved onto the training portion of the experiment. During this phase, mice were repeatedly given sessions of the same task (four sessions per day) until they either reached 70% shortest path transitions, or a 200 trial plateau. Once one of these conditions was met, they were moved onto the next task. The training period continued for 10 tasks, after which the mice were moved onto the *high-task regime*, wherein they were given three new tasks per day, followed by a repetition of the first task (task X, task Y, task Z, task X').

By visualising trajectories on single trials for multiple tasks, we can see that mice reliably learn efficient routes between reward locations (Figure 3.3, a). To quantify this we introduce a metric of task performance that is better at capturing near-optimal behaviour than the fraction of transitions that take the shortest paths metric we used to determine when to progress mice to the next task. This measure, which we refer to as *relative path distance*, is the number of steps taken between two goals divided by the theoretical optimal. This means that optimal performance would have a relative path distance of 1. Taking 4 steps when the optimal path only required 2 would have a relative path distance of 2. When we compute the mean relative path distance in the last 20 trials of tasks 1-10 across all mice, we see that they robustly use near-optimal routes between goals (Two-sided t-test against chance (6.44): $n=13$ mice, $t\text{-statistic}=-43.2$, $P=1.54\times 10^{-14}$, $d.f.=12$) (Figure 3.3, b). We calculate chance here as the mean relative path distance for first-trial transitions averaged across tasks 1-5 for all animals. Thus, mice learn efficient transitions in ABCD task examples.

The example trajectories in Figure 3.3 (a) show that mice have some preference for certain stereotyped routes. To quantify this, we computed the entropy of their chosen routes on correct transitions. We find that this entropy is significantly lower than what is expected if they sample correct routes uniformly (two-sided T-test against 1: $N=13$ animals, $\text{statistic}=-23.6$, $P=1.96\times 10^{-11}$,

df=12; Figure 3.3, c). This means that the mouse behaviour is not maximally entropic, but nor is it totally stereotyped.

Extending this, we tested the degree to which suboptimal performance (transitions with a relative path distance > 1) could be explained by animal-specific behavioural biases (Figure 3.3, f). Using data from exploration sessions on the maze before any training and in the absence of rewards, we are able to calculate the probability that the steps in a particular suboptimal trajectory would have been taken by chance. When we correlate mean relative path distance between goals with these transition probabilities, we see a significant net positive correlation, indicating that when mice take longer-than-optimal routes, they likely preferred these routes before any learning took place (T-test (two-sided) against 0 - N=13 animals, statistic=2.70, P=0.019, df=12). This shows that errors can be partially explained by 'stickiness' of their pre-configured behavioural biases.

Having ascertained that mice can learn individual examples of the ABCD task, we sought to determine whether mice learned new example tasks faster through training. This effect – sometimes called *transfer learning* in the ML literature (Caruana, 1997; Zhuang et al., 2020) – is indicative of having learned generalisable features of a problem set that enables data-efficient learning downstream. For instance, an image recognition model trained to distinguish cats from dogs will require less training data or fewer iterations to go on to distinguish cows from sheep than a network trained from scratch, because it has already learned generalisable features of natural images.

We go on to quantify this effect in our data in a few ways. Firstly, we observe that across the first 20 trials within a task the relative path distance quickly falls to near optimal (Figure 3.3, d). This effect was more rapid for the last five tasks than the first five tasks, with the animals solving the task more efficiently from the very first trial. Further, they improve to a lower relative path distance by the end of the 20-trial window than in the early tasks (two-way repeated-measures ANOVA (n=13 mice) showed a main effect of trial: $F=11.7$, $P=1.5 \times 10^{-5}$, d.f.1=19, d.f.2=228; task: $F=35.0$, $P=7.1 \times 10^{-5}$, d.f.1=1, d.f.2=12; and trial $\times$ task interaction: $F=2.99$, $P=0.030$, d.f.1=19,

d.f.2=228). When limiting our relative path analysis to just the first trial, we see that for the 9 mice that completed all 40 tasks, the relative path distance improves markedly across tasks (one-way repeated measures ANOVA shows a main effect of task: $F=2.73$, $P=0.016$, d.f.1=7, d.f.2=42; Figure 3.3, e). For the subsequent analyses, we aggregate the results of the first 20 trials of a new task for early vs late tasks. In doing so, we observe that the mean relative path distance is significantly higher in early tasks than in late tasks (two-sided Wilcoxon test $N=13$ animals, Statistic=0.0, $P=2.44 \times 10^{-4}$; Figure 3.3, g). When we look at all transitions in these first 20 trials, there is a significantly higher proportion of shortest-paths taken in late trials than in early trials (two-sided Wilcoxon test $N=13$ animals, Statistic=10.0, $P=0.010$, Figure 3.3, h)). Finally, when we compare the proportion of perfect trials (trials where all transitions - $a \rightarrow b$, $b \rightarrow c$, $c \rightarrow d$, and $d \rightarrow a$ - took the shortest path), we see a significantly higher proportion in late tasks than in early tasks (Two-sided Wilcoxon test for $N=13$ animals, Statistic=11.0, $P=0.028$, Figure 3.3, i left)). However, for both early and late tasks, the proportion of perfect trials is significantly above chance (two-tailed T-test against chance (0.007): Early tasks statistic=2.55, $P=0.025$; Late tasks - statistic=4.06, $P=0.002$, Figure 3.3, i right)).




Figure 3.3: Mice learn efficient routes between goals and task acquisition improves across examples

Figure caption on next page.

Figure 3.3: Mice learn efficient routes between goals and task acquisition improves across examples.

(a) Example single-trial trajectories for three ABCD task configurations. Mice show a preference for stereotyped solutions across trials. (b) In the last 20 trials of a task, mice routinely take routes whose relative path distance is close to optimal and is significantly lower than expected by chance. (c) Quantification of mouse stereotypy as the entropy of their correct transitions against a uniform sampling of all possible correct transitions. All mice show significantly lower entropy than the maximum, indicating a preference for a specific subset of possible optimal paths. (d) Relative path distance across the first 20 trials of the first 5 tasks (blue) and last 5 tasks (red). Lighter coloured lines represent data from individual mice. (e) Relative path distance on the first trial of tasks 1-40. (n=9; 4 mice did not complete all 40 tasks). (f) Suboptimal transitions are correlated with exploratory biases in the maze prior to any task exposure. (g) The mean relative path distance for all transitions in the first 20 trials is significantly higher in early tasks than in late tasks. (h) The proportion of transitions taking the shortest path in the first 20 trials is significantly higher in late tasks compared to early tasks. (i) The proportion of perfect whole trials is significantly greater in late tasks than in early tasks. All data mean $\pm$ SEM.

3.3 Inference of unseen task-space transitions: a strong test of task structure knowledge

The transfer learning effect discussed in the previous section is just one measure of assessing whether animals have learned an abstract task structure applied to new tasks that share that structure. However, these improvements across tasks could be due to factors other than task structure generalisation (e.g., reduced anxiety, reduced thigmotaxis, stronger spatial map). The cyclical nature of the ABCD task allows us to look for a much stronger test of task structure knowledge: inference of an unobserved transition on the first trial of a new task (Figure 3.4, a). If the mice understand that the task is a four-state loop, then upon collecting reward d they should, on the very first trial, return directly to a via one of the available shortest paths. This cannot be attributed to mere memory retrieval, as this $d \rightarrow a$ transition is yet to be experienced. This inference is only possible with task structure knowledge.

To test this, we compared the proportion of tasks in which the mice completed a successful $d \rightarrow a$ transition against premature returns to a from rewards b or c (Figure 3.4, b). Strikingly, we find that in late tasks, but not early tasks, mice performed significantly more shortest path transitions from d to a on the very first trial compared to the premature-return control conditions (Two-sided Wilcoxon test; early tasks: n=13 mice, W-statistic=17.5, P=0.168; late tasks: n=13 mice,




Figure 3.4: Mice use the abstract task structure to infer the $d \rightarrow a$ transition on the first trial

(a) Schematic. On the first trial of a new task, knowledge of the ABCD loop will enable mice to infer that after reward D they have to return at A without having to experience it. **(b)** Mice are more likely to take the shortest path from reward $d \rightarrow a$ than premature returns to a from rewards b or c in late tasks, but not early tasks. This effect is not explained by potential confounding factors such as: **(c)** analytic chance levels, **(d)** animal-specific behavioural biases, or **(e)** spurious differences in goal distances on the maze.

W-statistic=3.0, $P=0.004$). This strongly supports mice having learned the abstract task structure.

For further confirmation, we quantified the degree to which this effect could be explained by various potential confounds. When comparing analytic chance levels (the probability of a transition from node X to node Y by the shortest path in a randomly diffusing animal), we see no significant difference between $d \rightarrow a$ and $c \rightarrow a/b \rightarrow a$ transitions in either early or late tasks (Wilcoxon test (two-sided); early tasks: $N=13$ animals, statistic=20.0, $P=0.080$; late tasks: $N=13$ animals, statistic=32.0, $P=0.376$; Figure 3.4, c). Using data from mice exploring the maze environment before any training in the absence of any task, we can also define chance levels using animal-specific behavioural biases (Figure 3.4, d). Using this method, we observe no differences in the empirical chance levels between $d \rightarrow a$ and $c \rightarrow a/b \rightarrow a$ transitions in either

early or late tasks (Wilcoxon test (two-sided); early tasks: $N=13$ animals, statistic=42.0, $P=0.839$; Late tasks: $N=13$ animals, statistic=43.0, $P=0.893$). Finally, we compared physical distances on the maze between $d \rightarrow a$ nodes against $b \rightarrow a / c \rightarrow a$ (Figure 3.4, e), using the rationale that if $d \rightarrow a$ were spuriously closer than the controls it would inflate the proportion of $d \rightarrow a$ transitions. Doing this, we see no differences in the shortest physical distances between $d \rightarrow a$ and $c \rightarrow a / b \rightarrow a$ transitions in either early or late tasks (Wilcoxon test (two-sided); early tasks: $N=13$ animals, statistic=16.0, $P=0.071$; Late Tasks: $N=13$ animals, statistic=25.0, $P=0.477$).

3.4 Summary

In this chapter we have introduced the ABCD task, a behavioural paradigm for task structure abstraction (Figure 3.1). We have demonstrated that by scrambling reward locations across task instances we can decouple the structure of the task from the structure of the spatial environment in which that task takes place (Figure 3.2) which will allow us to study the neural representations of these structures separately in later chapters. We have shown that mice can learn to efficiently navigate between goal locations after being trained on ABCD task examples, and that they show a transfer-learning-like improvement across tasks indicative of having learned generalisable features (Figure 3.3). Finally, we showed that mice can infer an unseen transition on the very first trial of a new task (Figure 3.4), which is a strong test of task structure knowledge and cannot be explained by mere memorisation.

Together, these results define a behavioural phenotype of task structure abstraction: efficient paths, transfer across examples, and first-trial inference in a new configuration. In the next chapters we leverage the space–task dissociation built into the ABCD paradigm to test whether, and how, neural populations encode this task structure and generalise it across task instances.

4

Characterising task structure representations in medial frontal cortex

Contents

4.1	High-yield chronic <i>in vivo</i> electrophysiology with silicon probes	78
4.2	Space and goal-progress tuning are dominant features of mFC representations	80
4.3	mFC does not have generalising state tuning in the ABCD task	89
4.4	Summary	95

To look for neural representations of task structure that permit the complex goal-directed behaviours we described in Chapter 3, we focus first on medial frontal cortex. Causal perturbations and recordings across species converge on its role in representing behavioural structure (rules, stages, and sequence order) and using these representations to guide flexible behaviour. Because mice in our paradigm generalise to different ABCD configurations from the first trial (based on their inference behaviour), we are specifically looking for a neural code that permits such generalisation behaviour.

I will here detail our primary hypothesis about the representations we expect to observe in mFC of mice performing the ABCD task. Previous studies have found that frontal regions represent a full (factorised) abstraction of the task structure (Baram et al., 2021; Zhou et al., 2021).

For an example in humans, medial-frontal and entorhinal regions represent the structure of a reinforcement learning task in a way that is invariant to sensorimotor specifics. In mice, mFC and hippocampus form complementary representations of a reversal learning task, wherein mFC neurons code for task stages (e.g. in the initiation port regardless of its location; or the choice port regardless of its location), while the specific port locations are represented by CA1 neurons (Samborska et al., 2022).

Based on these studies, we hypothesise that in the ABCD task mFC neurons will represent absolute position in the ABCD sequence, and that these absolute sequence position representations will flexibly recouple with a representation of goal locations elsewhere (presumably hippocampus, though we will not be showing any data from the hippocampus). This type of solution is reminiscent of computation models of the hippocampal entorhinal circuit which involve representing the structure of arbitrary graphs in MEC in a way that permits flexible association with sensory details (Whittington et al., 2022, 2020).

To test this hypothesis, we will record from mFC in mice performing many ABCD tasks and employ analytical methods to track single neurons across task days.

4.1 High-yield chronic *in vivo* electrophysiology with silicon probes

To test how mFC encodes abstract behavioural structure across distinct instances of the ABCD task, we implanted high-throughput silicon probes in mice (N=7 total; Cambridge Neurotech F-series 64-channel 6-shank probes - N=5; Neuropixels 1.0 - N=2) performing the ABCD task (Figure 4.1). For mice implanted with the 6-shank Cambridge Neurotech probes, the shanks were arranged on the anterior-posterior axis (Figure 4.1, a). We were able to recover recording site locations by registering histology images to the Allen mouse brain atlas using the python-based histology software HERBs (Figure 4.1, a - c). Looking at the laminar profile of channel positions we are able to infer that 90.7% of neurons we recorded from were located in mFC. Most of these mFC sites were localised to the prelimbic cortex (68.3%), with additional placements in

anterior cingulate (11.3%), infralimbic (6.1%), and M2 (5.0%). The remaining 9.3% were located in peri-mFC white-matter (4.8%, likely reflecting atlas–brain boundary mismatches), dorsal peduncular nucleus (2.2%), striatum (1.1%), medial orbital cortex (0.6%), lateral septal nucleus (0.3%), and olfactory cortex (0.3%).

Since we need to be able to track individual neurons across many tasks, we aggregate data from two recording days to form ‘recording double-days’ (Figure 4.1, d). Each recording day consisted of four ABCD sessions in the structure task X, task Y, task Z, task X’ - where task X’ is a repeat of the first task of the day but separated by many hours. The mouse therefore sees 3 new tasks per day; 6 in a double day. We record sleep/rest sessions between each ABCD task, as well as at the start of the day. We combined neural data across double-days by concatenating the raw data recordings (see the ‘Electrophysiological recording, spike sorting, and behavioural tracking’ subsection of General Methods for more detailed information about the spike sorting procedure and criteria). Looking at individual drift traces and example spike waveforms we can see that there is broad stability of recordings both across sessions within a day and across days (Figure 4.1, e).

Together, accurate anatomical localisation, double-day concatenated spike sorting with conservative curation, and broad drift stability yielded single-unit populations suitable for longitudinal analyses across multiple ABCD tasks. Despite expected attrition (lower for single days, higher for double-days), the resulting yields were sufficient to perform robust single-neuron level analysis of task structure representations.




Figure 4.1: Chronic high-yield electrophysiology with silicon probes in medial frontal cortex

(a) 3D reconstruction of probe locations (channel sites only) rendered and tracked using the HERBs python package. Inset shows medial frontal cortex regions. (b) Anatomical regions that correspond to recording sites on all shanks of all probes in this dataset. The majority of recording sites are located in mFC regions, with the majority of those falling within prelimbic cortex. (c) Coronal histology slice from mouse implanted with Cambridge Neurotech F-series probe terminating in prelimbic cortex. (d) Recording from individual sessions are concatenated across two days into 'recording double-days'. Each day consists of four ABCD sessions separated by sleep sessions in a separate environment to both the maze and the home cage. A single day involves the mice completing three different tasks (task X, task Y, task Z) followed by a repetition of the first task (task X'). A recording double day therefore includes 6 different task configurations. (e) Top: Estimated drift trace for a concatenated recording double day. The dashed line indicates the split between the day 1 recordings and the day 2 recordings. The drift trace shows that the recording is largely stable across sessions and days. Bottom: Three example mean spike waveforms from three different mice. The black traces show the mean waveforms from the first 100 spikes extracted from day 1 and the red traces are the mean of the last 100 spikes on day 2. Scale bars: Vertical: 200 μ V, Horizontal: 0.5 ms.

4.2 Space and goal-progress tuning are dominant features of mFC representations

As we demonstrated in Chapter 3, mice are able to track their position in the ABCD task sequence. We have multiple behavioural measures of this: near-optimal paths between reward locations

(Figure 3.3, a, b); transfer-learning like improvements across task instances (Figure 3.3, d, e, g, h, i); and the strongest measure, inference of an unexperienced transition on the very first trial of a new task (Figure 3.4). To understand the neural representations that underpin these signatures of flexible model-based behaviour, we performed chronic electrophysiological recordings from mFC (Figure 4.1, a, b) of expert mice during the late task stage (tasks 21-40) in a regime where they perform three new tasks per day (Figure 4.1, d).

It is important to define here some task structure variables of interest (Figure 4.2, a). The first of these variables is *goal-progress*, which is how much progress the mouse has made towards its next goal. Here we define goal-progress as the percentage elapsed time between a pair of goals (Figure 4.2, a-Top), rather than absolute time or distance travelled. This goal-progress measure stretches and compresses depending on the time taken to reach the next goal (and the correlate of this, how far apart the goals are). For example, a goal-progress of 50% for a pair of distal goals that takes the mouse 40 seconds to reach will be at the 20 second mark, whereas a pair of proximal goals that takes only 10 seconds to reach will be at the 5 second mark. Goal-progress is not a measure of how much progress the mouse has made through the entire sequence - just a pair of consecutive goals. The next variable of interest is what we refer to as *state* (Figure 4.2, a - bottom). The state variable is which goal in the ABCD sequence the mouse is currently working towards. For instance, task state A begins when the mouse receives reward at location *a* and ends when the mouse gets reward at location *b*, and so on. The final variable is *space*, which is simply the location of the mouse on the 3×3 grid maze.

We see example neurons that are strongly tuned to goal-progress, and do not distinguish between which reward the mice are currently working towards (Figure 4.2, b). The raster plot shown in Figure 4.2 (b - top) is from a neuron that consistently increases its firing rate as the mouse approaches the reward. It is a late goal-progress cell. Note that the neuron fires at late goal-progress for trials of different lengths in terms of absolute elapsed time. We can see that neurons fire not only for late goal-progress, but also for early and middle goal-progress (Figure 4.2, a - bottom). Due to the periodic structure of the ABCD task, we are able to represent a trial using

polar plots. The angle around the plot represents the progress through the ABCD task. For these plots, individual states have been separately stretched/compressed into $90 \times 1^\circ$ goal-progress bins. The radial distance represents the firing rate of the neuron at that goal-progress bin.

A subset of these goal-progress tuned neurons are selective for particular states (Figure 4.2, c - bottom). These state-tuned cells do not necessarily code for just one state; some have peaks in two or three states. It should be noted that these example state tuned neurons still respect their particular goal-progress tuning. These state selective example neurons can be both spatially tuned (Figure 4.2, c - top left) or non-spatially tuned (Figure 4.2, c - top right). The example spatially tuned example neuron has firing peaks when the mouse receives reward at the centre (node 5) when it is both reward *c* and reward *a*, both of which at early goal-progress (around 5°). The example non-spatially tuned neuron has two firing peaks in both tasks but no generalising spatial selectivity.




Figure 4.2: Example mFC neurons tuned to place, state, and goal-progress

(a) Schematic illustrating what we mean by ‘goal-progress’ (% elapsed time between a pair of goals) and ‘state’ (sequence position in the ABCD loop). (b) Bottom - Example neurons that are tuned to relative progress to goal. Angle on the polar plots corresponds to progress through the trial, normalised separately for each state. Maximum firing rate indicated on upper right quadrants. Top: raster plots for single neuron for all trials of state C. The neuron becomes ‘late’ goal-progress tuned, firing as the mouse approaches the reward at *d*. (c) Bottom: example neurons that are selective to a particular state while retaining goal-progress tuning. Top left: example spatially tuned neuron. This neuron is selectively active when the mouse gets reward at node 5 in two tasks when it is both reward *c* and reward *a*. Top right: Example neuron with state selectivity that is not explained by spatial selectivity.

We next sought a principled method to categorise neurons by their tuning properties. For this we employed a GLM approach wherein task variables, as well as lower-level kinematic variables, compete to explain variance in the firing rates of individual neurons (Figure 4.3, a). The specific regressors we included were: goal-progress, place, speed, acceleration, absolute elapsed time from reward, and absolute distance travelled from reward. Full details of the GLM setup can be found in General Methods, but I will summarise them briefly here. We quantified how consistently goal-progress and spatial tuning generalise across tasks and states using leave-one-out cross-validation over all task $\times$ state combinations. For each fold, we held out one task–state (e.g., task 2, state A) and, from the remaining data, built per-neuron rate maps: mean firing in five goal-progress bins and mean firing at each of the nine maze nodes (we exclude edge bins as these are almost never occupied in the earliest goal-progress bin). In the held-out data, we predicted binned firing rates using the corresponding rate map value for the current bin, while

controlling for speed, acceleration, time-from-reward, and distance-from-reward. The resulting regression coefficients provide an index of spatial and goal-progress tuning that generalises across tasks $\times$ states. We determined that a given neuron was place or state selective if: it has a mean β greater than 95% of its own null distribution (determined by a 100-iteration permutation test), and cross-task correlation coefficient significantly greater than 0 (see General Methods for specifics). Importantly, we do not use this GLM to assess state tuning. Instead, we use a z-scoring approach detailed in General Methods.

Surprisingly, given the broad heterogeneity and difficult-to-interpret mixed selectivity of frontal neurons reported in previous work, we find that the large majority of neurons in our dataset are significantly tuned to goal-progress (74% of 1252 neurons; Figure 4.3, b). However, a goal-progress signal does not tell you how far through the *sequence* you are, just how close you are to your next goal. For this you need a kind of state representation. Of the 931 goal-progress neurons, 64% were significantly tuned to state, and of these 597 goal-progress/state neurons, 63% were significantly tuned to place. It is notable that as the remaining 37% of neurons do not have spatial tuning, it is unclear at present what the determinant of their state tuning is (but we will return to this mystery in Chapter 5). When we look at the total population of 738 state-selective cells, 81% are also goal-progress cells; by-and-large state cells are a subset of goal-progress cells. This could imply that goal-progress is a scaffold onto which state tuning is added.

To determine the degree to which the population as a whole is sensitive to the task related variables place and goal-progress, we plot histograms of mean β values across task $\times$ state combinations for each neuron (Figure 4.3, c). We find that the population as a whole is highly significantly sensitive to goal-progress (one-sample two-sided T-test against 0; $N=1252$ neurons; statistic=21.7; $P=8.93 \times 10^{-89}$, $df=1251$; Figure 4.3, c - left). Likewise, the population as a whole is highly significantly sensitive to place (two-sided one-sample T-test against 0, $N=1252$ neurons, statistic=24.9; $P=3.31 \times 10^{-111}$, $df=1251$; Figure 4.3, c - right). We further quantified the population tuning to the kinematic co-regressors speed and acceleration (Figure 4.3, d). For speed, we

observe a slight but significant rightward shift in the histogram of mean β values (two-sided one-sample T-test against 0; $N=1252$ neurons, statistic=3.36, $P=8.01 \times 10^{-4}$, $df=1251$). For acceleration, we do not observe any significant shift from 0 in the population ($N=1252$ neurons, statistic=-0.78, $P=0.438$, $df=1251$.)

To summarise, cross-validated GLMs showed that goal-progress is the dominant signal in mFC (74% of 1252 neurons), with robust but secondary spatial tuning. State selectivity typically sits atop these axes (81% of state-selective neurons are also goal-progress cells), while kinematic covariates contribute only weakly at the population level.




Figure 4.3: A GLM analysis of mFC representations

(a) Schematic of generalised linear model used to characterise tuning properties of mFC neurons in the ABCD task. Goal-progress, place, speed, acceleration, time from reward and distance from reward compete to explain variance in firing rates of individual neurons in a recording double-day. **(b)** Proportions of neurons significantly tuned to goal-progress, state and place variables. Neurons are categorised based on GLM results alongside additional tests detailed in the main text and in General Methods. **(c)** Distribution of regression coefficients for the task-related variables goal-progress and place across for all neurons. **(d)** Distribution of regression coefficients for kinematic variables speed and acceleration for all neurons.

We next dig deeper into goal-progress tuning, given the surprising dominance of this signal in the mFC population (Figure 4.4). Firstly, we generate a mean rate map for goal-progress bins in four tasks in a single day, (Figure 4.4, a). These rate maps aggregate goal-progress bins from all four states in a single task. By rank ordering the neurons by the goal-progress bin with the maximum firing rate in task X and then applying that ordering to the three subsequent sessions, we observe that 1) there is a remarkable degree of consistency in individual neurons' goal-progress tuning and 2) there are neurons with goal-progress peaks along the entire span of goal-progress values. This is despite the differing goal distances and locations across states within a task, as well as across tasks.

We go on to quantify these two phenomena. To test the consistency of goal-progress tuning, we calculate the mean correlation coefficients between goal-progress rate maps across tasks for each neuron (Figure 4.4, b). This reveals a highly significant degree of consistent goal-progress tuning in the population (two-sided one-sample T-test against 0: $n=2,461$ neurons; $t\text{-statistic}=104.3$; $P=0.0$, $d.f.=2,460$). Next, for those neurons with at least one significant goal-progress peak ($N=873$), we plot a histogram of the peak goal-progress bins (Figure 4.4, c). There are neurons with peaks at every goal progress bin, though they are not uniformly distributed, with there being a bump in the histogram at around the 60° bin (mid-late goal-progress). Further, we can see that goal-progress is invariant to changes in location and goal distances in the spiking activity of individual neurons on single trajectories (Figure 4.4, d).

The GLM analyses suggested that state tuning was layered on top of goal-progress coding. For more direct confirmation, we sought a more direct, geometry-level test. For this, we first generate a non-linear embedding using Uniform Manifold Approximation and Projection (UMAP) on the aggregate mean task rate maps across mice and recording days (360 bin 4-state goal progress map $\times$ neurons). We show the same manifold two ways (Figure 4.4, e) - one coloured by goal-progress (left) and one coloured by state (right). A clear pattern emerges using this procedure: the task manifold is made up of four distinct goal-progress sub-loops – one for each state – into

a four-petal floral structure. We repeated this procedure for recordings from individual double-days ($n=20$) and quantify this emergent structure using pairwise distances on the manifold (Figure 4.4, f). Specifically, we compared distances between points across different states for the same goal-progress (e.g., state B, goal-progress bin 45° vs state D, goal-progress bin 45°), across different states for the opposing goal-progress bin (e.g., state D, goal-progress bin 0° vs state B, goal-progress bin 45°), and, as a control, across different states and with the goal-progress bin for but for shuffled data. In doing so, we observe that the across goal-progress distances are robustly further than within goal-progress distances between states within a task. In turn, both of these conditions had greater pairwise distances than the shuffled data control (two-sided Bonferroni-corrected T-tests: across goal-progress versus within goal-progress: $t\text{-statistic}=6.09$, $P=2.25\times 10^{-5}$, $d.f.=19$; across-goal progress versus permuted control: $t\text{-statistic}=26.0$, $P=7.85\times 10^{-16}$, $d.f.=19$; within-goal progress versus permuted control: $t\text{-statistic}=8.63$, $P=1.60\times 10^{-7}$, $d.f.=19$; mean $\pm$ SEM, arbitrary units). This is a direct corroboration of what was suggested by the GLM results: state tuning is expressed *on top of* goal-progress. Goal-progress is the scaffold.

In this section, we recorded from mFC from expert mice performing many ABCD task examples and defined three key variables: spatial position on the 3×3 maze, the current sequence state (A-D), and goal-progress (percent of elapsed time within a goal-to-goal segment). Many single neurons showed robust goal-progress tuning that was invariant to absolute time and distance, with state selectivity often modulating that progress code and spatial selectivity present in a substantial subset. A leave-one-out, cross-validated GLM – controlling for speed, acceleration, time-from-reward and distance-from-reward – showed that goal-progress was the dominant signal (74% of neurons), with strong but secondary place effects and weaker contributions from kinematics. Across tasks, neurons' goal-progress rate maps were highly consistent, and a population embedding revealed four progress “loops” (one per state) forming a flower-like manifold. Overall, mFC representations are organised primarily by goal-progress, with state and space shaping that scaffold rather than replacing it.




Figure 4.4: Progress-to-goal is a fundamental feature of mFC representations

(a) Goal-progress tuning is stable across tasks despite changes in reward locations. For task X, neurons' average firing rates relative to each goal (pooled across states) are arranged by their peak goal-progress bin. The same alignment is largely maintained in tasks Y and Z and in a subsequent session of task X (X'). **(b)** Histogram of correlation coefficients, where each value is a neuron's mean goal-progress vector correlation across tasks. Goal progress is a highly correlated variable across tasks for the vast majority of neurons. **(c)** Distribution of significant goal-progress peak bins across neurons, tasks, and states. Peaks populate the entire span of goal-progress values. **(d)** Spikes from three goal-progress tuned neurons on one trajectory from reward *a* to reward *b*, and another trajectory from reward *c* to reward *d*. The three neurons keep their goal-progress tuning along these runs of different distances between different locations. **(e)** Mean task manifold obtained using Uniform Manifold Approximation and Projection (UMAP). Left: manifold coloured by goal-progress. Right: manifold coloured by state. The manifold is made up of four goal-progress subloops. **(f)** Manifold distances are measured for (i) across-state, opposite goal-progress bins; (ii) across-state, same goal-progress bin; and (iii) a shuffled control. $n = 20$ double days. Data are mean $\pm$ SEM (a.u.)

4.3 mFC does not have generalising state tuning in the ABCD task

At the start of this chapter, I outlined our primary hypothesis about the role that mFC neurons will serve in the ABCD task. Specifically, we hypothesised that mFC neurons will code explicitly for absolute position in the ABCD sequence. That is, a neuron that fires selectively for state B in task X will continue to fire selectively for state B in tasks Y and Z. We observed in the previous section that a subset of goal-progress tuned neurons was further tuned to task state in a single task (Figure 4.2, 4.3). However, we have not yet tested the consistency of this state tuning across tasks.

Looking at task rate maps for a selection of five neurons across tasks X, Y, Z and X', we do not see example neurons that retain their task space tuning throughout all four examples (Figure 4.5, a). Instead, these neurons remap across tasks, changing the task space position of their firing rate peak (or peaks) as the task changes. We go on to quantify this remapping across tasks. To find the 'remapping angle' from task X to task Y/Z/X', we find the circular shift of the task X rate map that maximises the correlation between rotated map X and the rate map from task Y/Z/X'. We use this as a measure of remapping angle rather than simply taking the task space bin with the highest firing rate for the two tasks and calculating the difference. This is because such a method would be vulnerable to multi-peaked cells with small changes in peak heights affecting the remapping angle where it ought not – although we show that the result holds for the peak-based remapping analysis in a published version of this work (El-Gaby et al., 2024). To see if there was an over-representation of neurons with stable state tuning (with a remapping angle of approximately 0), we calculated the remapping angle between task X and task Y/Z for all neurons and plotted the results as a polar histogram (Figure 4.5, b - bottom left). We find no peak at 0 compared to other remapping angles, which is evidence against the existence of generalising state cells (two proportions test against 25%; $n=1,594$ neurons; mean proportion of generalising neurons across one comparison (mean of X versus Y and X versus Z)=24%, $z=0.91$, $P=0.363$). Note that the distribution of remapping angles showing no generalisation of state cells is not uniform across all remapping bins but instead has four large peaks at the cardinal

directions (0° , 90° , 180° , and 270°). This is due to the strong goal-progress tuning we described in the previous section. There are effectively only four possible remapping angles that a neuron can have that preserves goal-progress tuning. We can next test whether remapping is random across sessions, or is due to a constant drift in the representation through time, using the comparison between task X and task X' (Figure 4.5, b - bottom right). These sessions have the same reward sequence but are separated by many hours in the day. If remapping is random across sessions or due to drift over time, then we would expect to see a similar 4-peaked remapping angle distribution to the X vs Y/Z comparison. Remarkably, we observe that the population of neurons has the same task tuning in task X to X', with a single large peak at 0° (two proportions test against a chance level of 25%; $n=1,160$ neurons; proportion generalising=78%, $z=25.7$, $P=0.0$). This can be interpreted as evidence against remapping being due to drift.

Although we did not find single neurons that preserved an absolute state preference across tasks, the population might still stably encode state if relative tuning relationships are maintained, analogous to grid cell modules remapping across environments. Accordingly, we next asked whether the population remapped but did so coherently, meaning that all neurons remapped by the same amount across sessions. This would be conceptually the same as having the absolute ABCD sequence representation but in a way that is not aligned to state A, or the tone that signals the start of a new trial. For this, we use a similar correlation-maximising task space rate map rotation procedure to what we used for the single-neuron remapping analysis, only this time we use it to compute the relative angle between a pair of neurons within the same task (Figure 4.5, c - top). We denote the between-neuron angle in the two tasks we are comparing as ϕ and ψ . If we do this separately for task X and task Y/Z, we can compare the between-neuron angle across sessions. If the neurons preserve their relationship, this difference-of-differences will be approximately 0, and if that is the case, we say that this is a coherently-remapping neuron pair. We compute this remapping coherence measure for each pair of neurons and plot the results in a polar histogram (Figure 4.5, c - bottom). If the entire population remapped coherently, there should be a single peak around 0° . Importantly, we do not find that the whole population, or close to that, remaps coherently. However, it is not totally incoherent, with a

slight but significant over-representation of neuron pairs at the 0° coherently remapping bin (two-proportions test against a chance level of 25%; $n=35,164$ pairs; mean of X vs Y coherence and X versus Z coherence=29.3%, $z=13.0$, $P=0.0$). That is to say that there are more neuron pairs that remap coherently than expected by chance, but it is not the entire population.




Figure 4.5: mFC neurons do not code explicitly for positions in the ABCD task structure

(a) Example task polar rate maps for five simultaneously recorded neurons. State tuning remaps but goal-progress preference remains stable. Each polar plot is annotated on the right with that cell's rotation angle ($^\circ$) relative to task X. **(b)** Top - Schematic illustrating how tuning angles are quantified across sessions. Generalising neurons retain their task tuning angle from task X in task Y ($\alpha - \beta \approx 0$), while remapping neurons change their state tuning ($\alpha - \beta > 0$). Bottom left: Distribution of remapping angles ($\alpha - \beta$) from task X to tasks Y/Z. There is no peak at 0, indicating that there is not an over-representation of neurons that retain their state tuning across different tasks. Bottom right: Distribution of remapping angles from task X to task X'. Neurons have the same state tuning when given a repeated session of the same task. **(c)** Top: Schematic illustrating how pair-wise relationships in task tuning are calculated for pairs of neurons across tasks. In task X, neuron 1 and neuron 2 have a tuning difference ϕ of around 170° . If the neurons retain their pair-wise relationship - that is, they remap coherently - their task Y tuning difference ψ will also be 170° ($\phi - \psi \approx 0$). If they do not remap coherently, their tuning difference will not be preserved ($\phi - \psi > 0$). Bottom left: Polar histograms show that the fraction of coherent state-tuned pairs (the peak at 0°) exceeds chance yet is far from 100%. Bottom right: most state-tuned pairs preserve their relative angles across sessions of the same task, as predicted by the single-neuron remapping analysis.

We have shown that there are more neuron pairs that remap their task-space tuning coherently across tasks that would be expected by chance (Figure 4.5, c). This finding could indicate the existence of a modular structure to the population, reminiscent of grid cell modules in MEC (Stensola, 2012). That is, there could be discrete subpopulations of neurons, within which neurons retain their task tuning relative to each other, but across which the task tuning is incoherent and changeable. Here, we test the hypothesis that state-tuned mFC neurons are organised into task-space modules that remap coherently across tasks. To do this, we apply a dimensionality reduction and clustering approach (Figure 4.6, a). For each recording day, we construct a neuron $\times$ neuron distance matrix D , where D_{ij} quantifies pairwise remapping coherence between neurons i and j across tasks; smaller values indicate neurons that preserve their relative task-space tuning, larger values indicate an inconsistent relationship. That is, we quantified pairwise remapping *incoherence* by converting between-neuron angle changes into a circular distance and taking the worst case across comparisons (highest incoherence).

We then obtained a low-dimensional representation (2D for visualisation) by applying t-SNE and identified putative modules via agglomerative hierarchical clustering on the embedded coordinates. An example clustering that appears to show seven distinct modules in a single recording day can be seen in Figure 4.6, a - left. To assess this clustering statistically, we compute the silhouette score for each recording day (a metric of clustering quality). We repeated the analysis on permuted data that kept each neuron's state tuning in task X and its goal-progress tuning in all tasks intact, but randomly remapped state preference across tasks. Comparing silhouette scores from the real and permuted datasets provides a baseline for chance clustering (Figure 4.6, a - right). Clustering of real neural data gives silhouette scores significantly greater than permutations, indicative of genuine modularity in the population (two-sided Wilcoxon test: $n=38$ recording days; W -statistic=126.0, $P=2.13 \times 10^{-4}$). We next visualise the remapping of the seven putative modules identified from the clustering analysis in the example recording day (Figure 4.6, b). The neurons in these modules broadly retain their internal relationships, remap together across tasks Y and Z, and return to their original orientation in task X'.

Further, we aimed to quantify whether it was only similarly tuned neurons that remapped coherently. For the modular account to hold, we need coherent pairs to span the whole task space, whereas only similarly tuned neurons being coherent can be explained by mere spatial tuning. For this, we calculated the proportion of coherently remapping neurons across a recording day split out by their tuning difference (0-45°, 45-90°, 90-135°, 135-180°; Figure 4.6, c). Comparing these proportions to a chance level of 1/16 (the probability that randomly remapping neurons with goal-progress tuning will remap together by chance across two tasks: $1/4 \times 1/4$) we find significantly above chance proportions for all conditions (two-sided T-test with Bonferroni correction 1/16: N=38 recording days: 0-45 degrees statistic=8.17, $P=3.33 \times 10^{-9}$, df=37; 45-90 degrees, statistic=2.84, $P=0.013$, df=37; 90-135 degrees statistic=3.88, $P=0.001$, df=37; 135-180 degrees statistic=2.89, $P=0.013$, df=37).

Finally, inspired by the anatomical proximity of neurons within a grid cell module compared to across modules, we used the probe geometry and channel IDs of individual neurons' spike waveforms to approximate the anatomical distance between neurons that remap coherently vs those that are incoherent (Figure 4.6, c). To reduce noise introduced by including neurons that remap together merely by chance, for this analysis we only consider neurons that remap together across double-days (6 tasks) to be coherent. Remarkably, we find that coherently remapping neurons are slightly, but significantly, closer together anatomically than incoherently remapping neurons (Two-sided Mann-Whitney U test: N=3567 pairs – 53 coherent, 3514 incoherent – statistic=72872, $P=0.006$).




Figure 4.6: Modular organisation of coherently remapping neurons in mFC

(a) Left - Clusters of coherently remapping neurons in a single session revealed by *t*-distributed stochastic neighbour embedding (tSNE) and hierarchical clustering on a distance matrix summarising cross-task remapping similarity. Each dot is a neuron, and dots are coloured by cluster. Inset is the silhouette score of this session's clustering. Right - Plot of actual silhouette scores for individual recording days against permutations that preserve goal-progress tuning but randomly remap state. **(b)** Seven modules of coherently remapping neurons in tasks X-Y-Z-X' in a single recording day. Each ring corresponds to a single module, each dot is a single neuron, and neurons are coloured by their state tuning in task X. Within-module, but not between-module, tuning relationships are preserved across tasks. **(c)** Proportions of coherently remapping cell pairs by tuning difference. There is significantly above chance cell pair proportions for all tuning difference bins. Chance is 1/16, which is the probability that two goal-progress tuned neurons will appear to remap coherently by chance across tasks in a recording day ($1/4 \times 1/4$). Dots corresponding to single recording days. **(d)** Coherently remapping cell pairs are slightly but significantly closer anatomically than incoherently remapping cell pairs. Unlike the previous analysis, this only uses neuron pairs that remap coherently over double-days (6 tasks).

4.4 Summary

In this chapter we used high-yield, chronic silicon probe implants to record mFC activity in mice performing multiple examples of the ABCD task (Figure 4.1). We introduced three variables that these mFC neurons represent — space (location on the 3×3 maze), state (sequence position in the ABCD loop), and goal-progress (percent elapsed within a goal-to-goal segment; (Figure 4.2). We showed that goal-progress tuning is robust, often modulated by state, and sometimes accompanied by spatial selectivity. This goal-progress tuning stretches and compresses to accommodate different goal distances (Figure 4.4).

A cross-validated GLM in which task and kinematic variables competed to explain firing revealed that goal-progress is the dominant, generalisable signal in mFC activity (74% of 1,252 neurons), with strong but secondary spatial effects and only weak contributions from speed and acceleration (Figure 4.3). At the single-cell level, goal-progress rate maps were highly consistent across tasks, and tuning peaks tiled the entire progress axis. At the population level, a UMAP manifold analysis uncovered four progress sub-loops (one per state) forming a four-petal structure. Quantifying distances on this manifold indicate that separation is primarily organised by goal-progress rather than state identity (Figure 4.4).

Regarding state tuning, neurons did not preserve absolute state preferences across different tasks (no excess at zero remapping), disconfirming our hypothesis that mFC would form a generalised code for the absolute ABCD loop (Figure 4.5). However, mFC neurons reinstated state tuning when the identical task was repeated later the same day (X vs X' peak at zero), which appears to rule out mere drift. Pairwise analyses showed an over-representation of neuron pairs that remapped coherently (relative angles preserved) without the entire population rotating as one. This quasi-coherence inspired a modular account: subpopulations of neurons (modules) maintain state tuning relative to each other across tasks, but this relationship is not conserved across modules or relative to the absolute task structure (Figure 4.6).

We tested this by constructing neuron–neuron incoherence matrices from worst case between-neuron angle changes across tasks, embedding them using tSNE, and clustering with agglomerative hierarchical methods. Real data exhibited significantly higher silhouette scores than permuted controls, indicating a genuine modular structure. Within identified modules, relative tuning was preserved across tasks. This coherence held for neuron pairs spanning a wide range of tuning offsets, not just for similarly tuned neurons. Finally, we showed that coherent neuron pairs were slightly but significantly closer anatomically than incoherent neuron pairs, reminiscent of grid cell modules in mEC (Figure 4.6).

Overall, mFC activity appears to be scaffolded by a strong goal-progress code that stretches and compresses with behaviour, onto which state and space are added. Rather than a single absolute state map, the population is made up of multiple task-space modules that remap quasi-coherently across tasks, but which return to their original configuration when the same task recurs. In the following chapter, we will introduce a model that can explain our neural and behavioural observations.

5

mFC neurons are organised into task-Structured Memory Buffers

Contents

5.1	An introduction to the task-Structured Memory Buffers (tSMB) model	99
5.2	Neurons in a tSMB have mnemonic fields in task space	104
5.3	Summary	113

If modules are the answer, then what is the question?

We initially hypothesised that mFC neurons would form an explicit, fully abstracted factorised representation of the position in the ABCD sequence, and that this representation would be transferred to new tasks by some associative binding. That is, we expected that a neuron representing state A in task 1 would persist in representing state A through subsequent tasks, regardless of where goal A was in space. However, this hypothesis was not supported by our remapping analyses detailed in Chapter 4. This hypothesis had its roots in factorised models of structure learning (Whittington et al., 2020; George et al., 2021). In these models, structures (be they spatial, ordinal, arbitrary graphs) are represented explicitly, and flexibly rebound to the items that populate those structures via transient synaptic connections.

There are several features of our data that point instead to a conjunctive, dynamics-based mechanism by which mFC may solve the generalisation problem. This account of the neural data can explain our key behavioural results: the transfer-learning-like improvements across tasks, and the zero-shot inference at the $d \rightarrow a$ transition on the first trial of a new problem (Chapter 3).

To summarise our key neural findings:

1. The majority of neurons tile progress-to-goal (defined here as the percentage of elapsed time between two rewards) with striking consistency across states within a task - as well as across tasks - compressing and stretching their firing fields to fit different goal distances. This was not explained by low-level kinematic features like speed, acceleration, true elapsed time, or distance travelled.
2. There is spatial tuning in a subset of neurons, often in conjunction with progress-to-goal tuning.
3. Neurons do not consistently code for the same position in the sequence across task examples, but instead ‘remap’ in task space in a way that respects their goal-progress tuning.
4. When we record the same task configuration twice, separated by two different task examples (X, Y, Z, X'), the neurons return to their initial task state representations. That is, the remapping angle of a neuron in task X and task X' is approximately 0.
5. There is a slight over-representation of neuron pairs that consistently remap by the same amount, meaning their between-neuron relationship is preserved. These neurons remap coherently. However, there is no evidence that the whole population remaps coherently, which would be conceptually the same as there being a fully abstract representation, but in a way that does not align itself to the tone that signals the start of a new trial.
6. **We find groups of neurons that remap together across tasks. We have termed these groups of coherently remapping neurons modules.**

Using these pieces of information, we developed the (task)-Structured Memory Buffers model. For the sake of clarity, I will give a broad overview of the model and how it was derived before discussing its implications and showing data in support of some of its predictions.

5.1 An introduction to the task-Structured Memory Buffers (tSMB) model

A change in reference frame

We start by observing the presence of modules of coherently remapping neurons that return to their original tuning when a previously experienced task configuration is presented for a second time. This indicates that remapping is not randomly drifting, and that modules are *anchored to a particular reference frame*. The anchoring may be to some variable in the environment, some variable in the behaviour, or a combination thereof. The modules must have separate reference frames, as if they did not, it would appear that the whole population remapped coherently and we would not see a modular structure at all. The game then becomes figuring out what these reference frames are.

Anchoring modules to behavioural steps

We proceed by acknowledging that space and goal-progress are dominant features of the representations of our recorded mFC neurons. We reasoned that the conjunction of these variables may constitute a *behavioural step*, and that these behavioural steps could act as the reference frames for coherently remapping modules. For example, one module could be anchored to when the mouse has just collected a reward (early goal-progress) at node 9. Another module could be anchored to the mouse being in the middle of a run (mid goal-progress) while crossing over node 5. Under these assumptions, the spatially-tuned state neurons we observed in Chapter 4 would serve as the inputs to the module (from some upstream region like hippocampus CA1). The non-spatial state selective (*non-anchor*) neurons could therefore carry information from this input neuron. Each module then acts as a *memory buffer* for a previous occurrence of a

behavioural step (Figure 5.1, a)

According to this model, if a neuron's activity is aligned to its module's anchor, then its activity will appear consistent across tasks (Figure 5.1, a), even when reward locations change (Figure 5.1, b). Conversely, the changing reward locations will cause apparent remapping if activity is aligned to task state (Figure 5.1, c), due to the anchor being visited (and the module's input triggered) at a different point in the task. The prediction that neurons in an SMB carry a memory of a past anchor visit can be tested without requiring many neurons from the same module to be recorded simultaneously. We can do this on a single neuron level by finding both the anchor *and* task-space lag that explain neuron firing rates that generalise across tasks. This will be the focus of Section 5.2

Tracking progress from reward locations

As the non-anchor neurons in a module are still goal-progress tuned, activity will move around the whole memory buffer as a function of task progress, returning to the anchor after four rewards. Modules therefore track elapsed task progress from their particular behavioural step. Imagine, for simplicity, four structured memory buffers anchored to the reward locations for a particular task configuration (Figure 5.1, d). When the mouse gets reward *a* at node 1, an activity bump is initialised at the anchor neuron for memory buffer 1 (Figure 5.1, b left). By the time the mouse has reached reward *b* at node 5, the activity on the ring for memory buffer 1 has progressed by 90° (one state), and the module for memory buffer 5 has an activity bump initialised at its anchor neuron. This process repeats – bumps initialised at anchors and progressing around the SMB rings as the mouse completes the task – for a whole trial up to reward *d*. The proposed circular structure of the SMB rings provides a clear substrate for biasing behaviour. Take the $t = 3$ panel from Figure 5.1, b. The activity on buffer 1, tracking task progress from the behavioural step of obtaining reward at node 1, has traversed 270° around the ring. A downstream area receiving input from the neurons at these high lags could read information from this ring to bias the animal towards returning to node 1 after the 3 intervening rewards have

been collected. This provides a mechanistic explanation for the zero-shot inference behaviour demonstrated in Chapter 3; a memory buffer structured by the ABCD task can store a memory of the first reward location and tell the mouse to return to that location a particular number of states in the future. This idea, that SMB activity can be used to bias behaviour, is the focus of Section 6.2.

Tracking progress from intermediate steps

The situation outlined above, with four active buffers in a task tracking progress from just the reward locations, is a simplification. After all, we defined the behavioural step that anchor each module as a conjunction of place $\times$ goal-progress, as these were dominant features of mFC representations we showed in Chapter 4. This implies that SMBs can track progress from steps *in between goals* (Figure 5.1, e, f). In Figure 5.1 (e), we outline 5 different SMBs anchored to node 5, each for a different goal-progress increment. In the context of wider behaviour, in Figure 5.1 (f), a mid-goal progress anchored buffer for node 5 tracks goal progress from a behavioural step between reward *c* at node 8 and reward *d* at node 6. The implication of this extension is striking: across the set of SMBs, the entire history of behavioural steps is stored in memory *at every point in time*. Combining this with the hypothesis that SMB activity can be used to bias behaviour, tSMBs could therefore compute the appropriate future action at every step.

Memory buffers are internally structured

Finally, a broader point is that there is a task-structured internal organisation to SMBs, rather than the results we observe arising due to temporal correlation of externally-imposed task structure motor or sensory drives. We can get at this by analysing the activity structure of mFC neurons offline during sleep, in total absence of task inputs. This will be the focus of Section 6.1.

Having introduced the tSMB model, across the following two chapters we will test its various predictions. Specifically, we will ask: (i) Can the activity of individual neurons be explained by tuning to an anchor $\times$ task-lag conjunction? (ii) Are circular SMB dynamics internally structured,

such that they persist offline during sleep and rest? (iii) Can we use the firing rates of SMB neurons to predict distal upcoming choices at specific future time steps?




Figure 5.1: The (task)-Structured Memory Buffers model

(a) A hypothetical ring-structured module of neurons (*task-structured memory buffer - tSMB*) anchored to behavioural step x : reward in the top left node. An activity bump moves around the tSMB as a function of goal progress. The neurons that are not directly anchor-tuned directly compute progress in the overall loop relative to the anchor. Each dot is a neuron, the blue neurons represent lags of one entire goal-goal progress cycle: 0° , 90° , 180° and 270° , whereas white neurons represent lags not a multiple of 90° . After four rewards, the activity bump returns to the start of the tSMB (b) Two ABCD task configurations with anchor x as the top left node. In task 1, anchor is reward a , whereas in task 2 the anchor is reward c . (c) If we align the activity of the tSMB to reward a instead of anchor x , the module appears to rotate by 180° across tasks, just like the real data in Figure 4.2. (d) A set of 9 memory buffers anchored to rewards at each of the 9 locations, of which 4 are highlighted (the rewards for task 1). We show the activity bumps along this tSMB for four time points at 90° increments through a trial. An activity bump starts at the anchor neuron (top of the ring) and moves along the tSMB as a function of goal-progress. Each ring tracks position from a particular anchor, so the population tracks position from all anchors simultaneously. We put forward that when the activity bump returns to the anchor neuron, the behaviour of the animal is biased towards revisiting that anchor. (e) Behavioural steps that anchor tSMBs need not just be reward location, but intermediate steps as well (due to goal-progress:place conjunctions for middle + late goal-progress tuning). Left: SMB anchored to reward at node 5. Right: tSMBs anchored to node 5 but for non-zero goal-progress values (at various points on a run). (f) Rings anchored at non-zero goal-progress (e.g., the purple outline) track what happens between goals, not just at reward. This means that the combination of all tSMBs provide, at any given moment, a record of the entire sequence of steps the animal has taken, not just the order of reward locations. (Graphics adapted with permission from Gil Costa)

5.2 Neurons in a tSMB have mnemonic fields in task space

According to the tSMB model, in the ABCD task, modules form ring-like memory buffers wherein neurons track position in task space from a particular anchor (e.g. reward at node 5; middle of a run over node 9, and so on). Other than the neurons that have a bump of activity directly at the anchor, the other neurons encode a memory of having visited that anchor at a certain task-space lag (e.g. reward at node 5 two rewards ago - 180° ; middle of a run over node 9 one and a half rewards ago 135°). We can say then that the neurons with non-zero lag (non-anchor neurons) have a *mnemonic field in task space*. There are a number of ways that we can get at this claim by analysing the activity of single neurons, rather than requiring activity to be recorded from entire tSMB rings simultaneously. In this section, we will provide evidence from three separate but related analyses that mFC neurons have mnemonic fields in task space. For all of the following analyses, to account for the strong goal-progress signal present in most mFC neurons, we only include data from each neuron's preferred goal-progress bin.

5.2.1 Aligning to a single anchor

For the first analysis, we will attempt to find, for each state-selective neuron, the single best-fitting anchor (place $\times$ goal-progress conjunction). This analysis assumes that neurons are tuned to a single anchor *and* a single lag. Specifically, we find the anchor which maximises the cross-task correlation when the neural activity is aligned to that anchor rather than the abstract task state (Figure 5.2). For this we used a 6-fold cross-validation approach where we use 5 tasks to fit the anchor which is then tested on the held-out task. The details of this analysis can be found in General Methods but I will provide a brief summary here. For each anchor ($N=27$; nodes $1-9 \times$ goal-progress $1-3$), we find the bins when the mice visit that anchor and sample 360 task bins from that time point. This equates to one full trial after the anchor visit. We then make polar rate maps of anchor-aligned activity for each neuron for all 5 training tasks. Using the correlation-maximising rotations procedure we described previously (Figure 4.5), we can find the angle between these polar rate maps for all task-task comparisons. We convert these angles to a distance $1 - \cos(\theta)$ and use these distances to find the best anchor alongside the training

task against which the held-out test task will be compared.

We can see from example neurons that finding the best fitting anchor, and aligning neural activity to visits to that anchor, can produce stable anchor-aligned rate maps across tasks, even when the task-aligned rate maps re-map significantly (Figure 5.2, a). We can test this alignment consistency across the population by computing the rotation of the candidate anchor-aligned training task for each neuron that is required to maximise correlation with the test task (the angular mis-match). Here, 0° indicates perfect alignment between the training and test anchor-aligned rate maps. We display the distribution of mis-match angles across mFC state neurons using polar histograms (Figure 5.2, b). For the distribution across all state neurons, as well as the distribution of only non-zero lag state cells (neurons that are state tuned but do not have a firing field at their anchor), there is a marked peak at 0° , indicative of cross-task preservation of anchor preference. Quantitatively, the fraction of neurons whose training vs test anchor-aligned rate maps deviated by up to $\pm 45^\circ$ (generalising anchors; chance = 25%) was 39% for all state-tuned neurons ($n=738$; two-proportions test vs 25%: $z=5.69$, $P=1.24 \times 10^{-8}$). When we restricted the analysis to non-zero-lag state-tuned neurons, this fraction was 36% ($n=545$; $z=3.87$, $P=1.08 \times 10^{-4}$). We go on to quantify the cross-validated correlation between the anchor-aligned test and training tasks (Figure 5.2, c). There is a rightward shift in the distribution of correlation coefficients for all state neurons (left, two-sided t-test against 0, $n=737$ neurons, $t\text{-statistic}=9.86$, $P=1.32 \times 10^{-21}$, $d.f.=736$). This effect remained when only considering non-zero lag state neurons (right, two-sided t-test against 0, $n=544$ neurons, $t\text{-statistic}=7.55$, $P=1.86 \times 10^{-13}$, $d.f.=543$).

We use the results of this single best-fitting anchor analysis to extract some basic properties of the mFC state neurons. Specifically, we examine (i) the distribution of lags from anchor, and (ii) the distribution of anchor locations. Since the above analysis was restricted to data within each neuron's preferred goal-progress bin, the goal-progress component of the anchor was assumed before the alignment procedure. Across the population of *consistently-anchored* state neurons, there are neurons with lags from anchor that span the entire span of possible values (0° - 360° ; Figure 5.2, d). This is true when considering the lag to be both the mode of lags across

folds of the cross-validation (left) as well as the circular mean across folds (right). For both of these measures, there is a peak around 0° , indicative of a slight over-representation of anchor-tuned neurons. We then mapped the spatial distribution of anchors for consistently-anchored state neurons (Figure 5.2, e). Looking at the 2D spatial histogram of anchor locations for early, mid, and late goal-progress cells, we observe that consistently anchored cells populate all goal-progress/place combinations. Note that these are not distribution of spatial firing peaks, but the spatial distribution of anchors. The spatial firing patterns of non-zero lag neurons are likely to be low / non-existent depending on the animal's behavioural stereotypy, despite having a strong spatial anchor. For both mid and late goal-progress neurons, there appears to be a slight over-representation of neurons anchored to the centre node. The anchor and lag parameters obtained using this set of analyses will be used in further analyses downstream, specifically those in which we predict future behavioural choices, and characterise neuron correlation structure during sleep.

To summarise this analysis, aligning activity to a single best-fitting place $\times$ goal-progress anchor reveals generalising cross-task structure that is obscured in task state-aligned maps. Anchor-relative lags tile the full 0° – 360° range, and anchors themselves cover all place $\times$ goal-progress combinations (with a mild centre bias at later goal-progress bins).




Figure 5.2: Mnemonic fields in task space 1 - aligning neural data to the best-fitting anchor

(a) Two example neurons with activity aligned to the abstract states in the ABCD task (top, blue) or the best fitting anchor (bottom, green) across 6 tasks in a recording double-day. For the anchor-aligned rate maps, we put the anchor at 0° (North) as indicated by the dashed vertical line. The task-aligned rate maps remap across tasks, whereas there is a high degree of consistency across the anchor aligned rate maps, with consistent peaks at $\approx 135^\circ$ for neuron 1 and $\approx 180^\circ$ for neuron 2. (b) Quantifications of the angle between anchor-aligned rate maps in held-out test tasks against training tasks for all state neurons (left) and non-zero lag neurons (right). For both conditions there is a peak at 0° , indicating a high degree of alignment in anchor-aligned rate maps across tasks. (c) Histograms of the cross-validated task map correlations between training and test tasks for all state neurons (left) and non-zero lag neurons (right). There is a significant rightwards shift in both of these distributions, indicating stable anchor alignment across tasks. (d) Distribution of lags from anchor (angle of the firing rate peak in the anchor-aligned rate maps) from all consistently tuned neurons identified in the previous analyses. For both the mode anchor across tasks (left) and mean anchor across tasks (right), there are neurons with anchors that span all 360° bins but with both having a peak at 0° : anchor neurons. (e) 2D spatial histograms of anchor locations for early, mid, and late goal progress. The maximum count is shown above each plot. For all goal progress bins there are neurons significantly anchored to all locations. There is an over-representation of node 5 for mid and late goal progress anchors.

5.2.2 Spatial maps for past locations

In the previous analysis, we demonstrated that a large fraction of mFC neurons have firing fields at a fixed task-space lag from a place $\times$ goal-progress anchor. We reasoned that if this were true, then these neurons are in effect place cells but for previously visited locations. Let's take the example of a neuron anchored to node 1 (corner), with a task lag of 180° . If we were to build a spatial rate map for this neuron using the firing rates and locations the mouse is currently in (zero-lag), spatial tuning will likely look diffuse. However, if we instead build a rate map using locations where the animal was *two states ago* (180°), the resulting rate map should sharpen into a place cell for node 1. This is the motivation for the following analysis.

To test the idea that neurons in a tSMB form place cells of the past, we constructed lagged spatial rate maps for neurons using locations from up to a full trial in the past, broken into 12 goal-progress bins (Figure 5.3). If neurons are coding for a memory of a previous visit to a specific place in a way that generalises across tasks, then there should be a specific lag at which the spatial rate maps are highly correlated across tasks. The peak of the spatial distribution at this correlation-maximising lag would be the neuron's spatial anchor. For this analysis we only took neurons from double-days meaning we have 6 task instances to provide robust cross-task comparisons. As with single anchor alignment, this analysis assumes that neurons have a single lag from its anchor. But here, we are *not* assuming single anchor.

Following this procedure, we show lagged spatial rate maps across 6 tasks for 3 example neurons, with the accompanying cross-task correlations (Figure 5.3, a). The three neurons shown have correlation peaks at three different lags. Neuron 1 has a cross-task correlation peak at -210° , with the inset showing a spatial field at node 7 (bottom left corner). Neuron 2 has a correlation peak at 0° : it is an anchor cell tuned to node 9. Neuron 3 has a correlation peak at 180° with a spatial field at node 6. These examples illustrate the general principle we showed with the anchor alignment analysis: neurons in tSMBs are tuned to anchor $\times$ lag conjunctions.

To further highlight this at the single-trajectory level, we can look at the spikes of a pair of neurons on single trials of two tasks (Figure 5.3, b). The neuron in red has a task lag of 0° anchored to node 6. Accordingly, it has a cluster of spikes on node 6 across the two tasks when node 6 is reward *b* and reward *a*. The neuron in green has a task lag of 270° and is also anchored to node 6, meaning it is in the same tSMB (module) as the red neuron. In Task 1, when the red neuron fires at reward *b*, the green neuron has a cluster of spikes at reward *a* on node 5. In task 2 however, when the red neuron's spike cluster is at reward *a*, the green neuron has a cluster of spikes at reward *a* on node 2. The green neuron has remapped in spatial tuning for *now*, but has retained its lagged spatial tuning relative to its anchor.

Finally, we assessed the consistency of lagged spatial tuning across the population of state sensitive neurons in mFC. For this we used a leave-one-out cross-validation approach wherein 5 training tasks were used to estimate the correlation-maximising lag. We then applied this lag to the test task and took the correlation between training and test lagged spatial maps. Plotting a histogram the mean correlation across folds of the cross-validation for all state tuned neurons (Figure 5.3, c), we observe a strong significant rightward shift from 0, indicating conserved lagged spatial tuning across tasks ($n = 738$; two-sided t-test: $t = 22.6$, $P = 1.45 \times 10^{-86}$, d.f. = 737). We observe a similar rightward shift even when restricting the analysis to non zero-lag neurons ($n = 285$; $t = 7.48$, $P = 9.07 \times 10^{-13}$, d.f. = 284; Figure 5.3, c).

To summarise, many mFC neurons behave like place cells of the past: their fields sharpen when locations are taken at a fixed task-space lag relative to a preferred place×progress anchor, and those lagged spatial maps reappear in left-out tasks. Single-trial examples show that cells within the same tSMB preserve anchor-relative lag even when their zero-lag spatial fields remap across tasks. Population-wise, cross-validated correlations at each neuron's preferred lag are strongly positive (and remain so for non-zero-lag cells), confirming that lagged spatial tuning generalises across task instances.




Figure 5.3: Mnemonic fields in task space 2 - neurons form place cell-like maps of previously visited locations

(a) Bottom: Three example neurons with spatial rate maps using locations the animal was in at various task lags (columns) for 6 tasks (rows). The column on the furthest right, at last lag 0, is just a place cell map for where the animal currently is. The maps use locations progressively further into the past for 12 task lag bins (one full trial) the further to the left they are. Top: The correlation between spatial rate maps across tasks for the three example neurons. The task lags with maximal cross-task spatial correlation are highlighted in black squares. For neuron 1, there is a peak correlation across tasks at task lag 210°; for neuron 2, the peak is at lag 0° (present); for neuron 3, the peak is at 180°. (b) Example neurons with spikes that occur on a single trial for two tasks. The behavioural trajectories are marked with the black lines. The spikes from neuron 1 are in red, the spikes from neuron 2 are in green. Neuron 1 is anchored to reward at node 6 with lag 0°. It has a cluster of spikes at that location when node 6 is reward *b* and reward *a*. Neuron 2 is also anchored to reward at node 6 but has a task lag of 270°. It has clusters of spikes at reward *a* in task 1 when neuron 1 fires at reward *b*, and at reward *d* when neuron 1 fires at reward *a*. (c) Histogram of mean cross-validated correlations between spatial map at preferred lag from training task and the data from the held out task aligned to that lag. When all state neurons are included there is a significant rightward shift against 0. (d) Histogram of lagged spatial map correlations but limiting the data for non-zero lag neurons. The rightward-shift against 0 persists despite the exclusion of the 0-lag neurons.

5.2.3 Allowing multiple anchor $\times$ lag conjunctions with elastic-net regression

For the final analysis testing whether mFC neurons represent task-space lags from a place $\times$ goal-progress behavioural step, we sought a method that would allow for the potential case where neurons are not cleanly tuned to a single anchor $\times$ lag conjunction. Single anchor alignment assumes a single lag *and* single anchor; lagged spatial tuning assumes a single lag *but not* single anchor; this analysis assumes neither. Instead, they could be some mixture of fields that contribute to a neuron's firing rate. For this, we used a linear regression approach to ask what combination of all the possible anchor $\times$ lag conjunctions explain the firing rates of mFC state neurons.

Specifically, we used elastic-net regression, which fits a linear model with an added penalty term for regularisation. This penalty term is a combination (1:1) of the $L1$ and $L2$ norms of the vector of model parameters. This model had a total of 324 regressors (9 locations $\times$ 3 goal-progress bins $\times$ 12 task-space lags, one full trial). As in the previous analyses, for each neuron we only considered data from each neuron's preferred goal-progress bin to ensure, to the highest extent practicable, that the state preferences we extract are not driven by strong goal-progress tuning.

We use these regression coefficients that explain firing rates, along with the mouse behaviour in a particular task, to predict task-space rate maps (Figure 5.4). We once again use leave-one-out cross-validation, wherein 5 tasks are used to fit the elastic-net model which is then used to predict the task-space rate map in a held out task. We show two example neurons' matrices of regression coefficients, alongside their actual vs predicted task-space rate maps (Figure 5.4, a). It is notable that both of these neurons have not just one, but multiple peaks in their regression coefficients. This means that variance in the neural data is not just explained by a single anchor or anchor $\times$ lag conjunction, as assumed by the previous analyses. Moreover, we observe a remarkable degree of correspondence between the actual rate maps and the rate maps predicted by the state tuning regression coefficients and the actual behaviour in a held out task.

Finally, we quantify the correlation between actual and predicted task space rate maps in mFC state neurons (Figure 5.4, b). We quantify this in four conditions and display the distribution of actual vs predicted task map correlations in histograms. Across all state-tuned neurons with non-zero regression weights, correlations were significantly greater than zero ($n = 489$; two-sided t-test: $t = 9.3$, $P = 5.3 \times 10^{-19}$, d.f. = 488). The effect remained when restricting to non-zero lag neurons ($n = 329$; $t = 3.9$, $P = 1.08 \times 10^{-4}$, d.f. = 328), as well as only distal non-zero lag neurons (lag > 1 state away from the anchor, $\pm 90^\circ$; $N=224$ neurons, $t=2.53$, $P=0.012$, d.f.= 223).

For the final condition, we aimed to rule out residual *trajectory coding*: neurons sensitive to the specific path the mouse is taking, rather than just location (Figure 5.4, b - right). For this, we ran a separate GLM in which place was replaced by conjunctions of current node $\times$ next node (24 place $\rightarrow$ next-place combinations). For each neuron, we built regressors from the mean firing rates of these conjunctions in the training task-state periods and predicted binned firing in the held-out period, while co-regressing goal-progress, speed, acceleration, time-from-reward, and distance-from-reward. We repeated this for all task-state folds and compared each neuron's mean trajectory coefficient to a circular-shift null. Any neuron with a coefficient above the 95th percentile of the null was determined to be trajectory-tuned and therefore excluded. This intentionally permissive threshold removes even weakly trajectory tuned neurons. In the remaining trajectory-free state neurons, the elastic-net regression still predicted task-space rate maps in held-out tasks ($n = 112$; two-sided t-test: $t = 4.27$, $P = 4.13 \times 10^{-5}$, d.f. = 111).




Figure 5.4: Mnemonic fields in task space 3 - an elastic-net approach to predicting task space maps with multiple competing anchor $\times$ lag conjunctions

(a) Linear regression with elastic-net regularisation reveals anchor:lag conjunctions predictive of neural activity for two example neurons. This method allows neurons to be sensitive to multiple anchors and lags. Left: matrix of regression coefficients for two neurons. Both neurons have multiple peaks, meaning multiple anchor:lag conjunctions explained variance in the activity of these neurons. Right: actual state tuning for neurons 1 and 2 (blue) and predicted state tuning based on the anchor:lag regression coefficients and the behaviour in that task. There is a high degree of similarity between the actual and predicted task rate maps using this regression method. **(b)** Correlations between actual and predicted task maps for all state tuned neurons, non-zero lag neurons only, distal non-zero lag neurons only, and non-trajectory tuned neurons only. For all conditions there is a significant rightward shift against 0.

5.3 Summary

In this chapter, we showed that mFC neurons have receptive fields at specific task-space lags from particular *behavioural steps* (place $\times$ goal-progress conjunctions). Firstly, by aligning activity

to a single, best-fitting place $\times$ goal-progress anchor, we observed cross-task structure that is obscured in absolute task state-aligned maps. Anchor preferences reinstated in left-out tasks, anchor-relative lags tiled the full 0° – 360° range, and anchors covered all place $\times$ progress combinations (with a mild centre bias at later progress bins). Secondly, task-space lagged spatial maps demonstrated that mFC neurons form *place cells of the past*. Neurons' spatial tuning sharpened when positions were taken at a fixed task-space lag, and these lagged maps generalised across tasks, including for non-zero lag cells. Thirdly, elastic-net regression showed that firing is explained by weighted combinations of anchor $\times$ lag terms. These parameters, when combined with mouse behaviour, predict task-space rate maps on held-out tasks. This effect holds when restricting to distal lags and after conservatively removing any neuron with even weak trajectory tuning.

Taken together, these results support the tSMB model (Figure 5.1). mFC implements memory buffers that are anchored to specific place $\times$ goal-progress steps, with neurons encoding stable, task-general lags around those anchors. If these buffers are truly circular (with intrinsic dynamics that wrap around rather than these responses being driven by task-structured sensorimotor inputs), then they provide a potential substrate for a predictive readout. That is, that they can be used to drive behaviour. These are the questions we will be answering in the following chapter.

6

Structured memory buffers are internally organised and predict distal choices

Contents

6.1	Neurons in a tSMB retain their ring-like dynamics offline	116
6.2	Neurons in a tSMB predict distal behavioural choices	122
6.3	Summary	132

In the previous chapter we introduced the tSMB model. According to this framework, mFC neurons are organised into memory buffers whose structure is shaped by behaviour. In our case, this structure is a four-state looping sequence. Rather than tracking progress through the absolute ABCD loop, multiple SMBs run in parallel, each tracking progress from a particular behavioural step (place $\times$ goal-progress). Because SMBs are anchored not only to rewarded locations but also to intermediate steps (mid/late goal-progress), the joint population activity across SMBs encodes the entire history of prior steps at every point in time.

We posited that tSMB activity can drive behaviour. If this is true, then we should be able to read tSMB activity at specified time to *predict* upcoming behaviour. This prediction will be the focus of this chapter, and we tackle the problem in two parts.

First, for prediction to be plausible, the SMBs must exhibit internally organised, task-structured circular dynamics independent of immediate inputs. To demonstrate why, consider the example of a neuron anchored to reward at node 5 with a task lag at 270° . This neuron's activity bump will be most active 3 states after a reward at node 5. But if the tSMB dynamics are internally organised and circular, we would be able to read from this neuron *if the anchor will be revisited in 90° time*. We therefore establish that this structure exists by showing that the SMB ring-like geometry is preserved offline during sleep and rest. This provides the necessary substrate for a prospective computation.

Second, having shown that the buffers are intrinsically structured, we then turn to a direct test — whether activity at tSMB neurons' activity bump time is predictive of upcoming choices.

6.1 Neurons in a tSMB retain their ring-like dynamics offline

A central claim of the tSMB model is that mFC is organised into ring-like buffers. That is, the internal structure of the network mirrors the structure of the behavioural task on which the mice have been trained: a 4-state loop. The results shown so far are not explained only by circularity. They could also be explained by line-attractor type dynamics and working memory. A strong test of intrinsic ring structure is to look offline, when sensory and motor drive are absent. If the buffers are a real, intrinsic architecture (rather than a convenient description of motor/sensory evoked patterns), their internal geometry should be detectable during rest and sleep, much like other cortical–hippocampal attractor systems whose relative phases and module structure are preserved offline and support replay/compression of recent experience (Kaefer et al., 2020; Peyrache et al., 2009; Trettel et al., 2019; Gardner et al., 2022; Peyrache et al., 2015; Dragoi and Tonegawa, 2011; Grosmark and Buzsáki, 2016). The 2D toroidal manifold of a grid cell module is detectable during sleep, as are the ring-attractor like dynamics of the head-direction circuitry. In this section, we will investigate whether the ring-like dynamics predicted by the SMB model are retained during rest and sleep sessions without any sensorimotor task-structured inputs.

We ask a deliberately simple question: do pairs of neurons that belong to the same putative buffer retain their ring-like relationship when the animal is not performing the task? Concretely, if two cells are close together on the ring (that is, they have small circular distance in task space), they should be more coactive during sleep than cells far apart on the ring.

We have already demonstrated that neurons track task progress from a particular behavioural step, but the analyses so far do not directly speak to the modules being structured as a ring. Concretely, these results so far could be either explained by a ring-like neural state space or a *delay line* neural state space (Figure 6.2, a). According to the delay line model, neurons can still track task progress from an anchor and permit readout to compute actions at a set task-space lag from that anchor, but the activity bump will not return to the anchor neuron after a complete ABCD loop. If the anchor is revisited, an activity bump will be re-initiated at the start of the delay line, but the bump will not loop around otherwise.

These two models make importantly different predictions about the sort of activity structure that should be preserved offline. Specifically, they make different predictions about the *neural distances* in pairwise neuron-neuron relationships. In Figure 6.2 (a), we demonstrate this with a pair of neurons, and their different relative positions on a ring and delay line. On the ring-structured neural state space, neuron 1 has a firing field around 30° after the anchor, while neuron 2 has a firing field around 30° before the anchor. The neural distance between these neurons is therefore 60° of *circular distance*. However, on a delay line there are no connections that bridge the high-lag neurons back to the anchor, meaning the neural distance between these neurons is 300° of *forward distance*. These two types of pair-wise distances will be critical for distinguishing between the ring and delay-line models of internal organisation.

The first analysis is conceptually very simple; we will be looking at pairwise sleep cross-correlations directly (Figure 6.1). We examined offline coactivity for neuron pairs sharing the same anchor by plotting their sleep cross-correlations against circular distance (Figure 6.1, a, top x-axis) and forward distance (bottom x-axis). The example schematics in Figure 6.1 (a)

illustrate how each distance is defined for representative pairs of neurons on a ring and on a delay line. The pairwise cross-correlation showed a clear V-shaped dependence on forward distance, with the highest correlation near $0^\circ/360^\circ$ and lowest near 180° . This is consistent with a ring-like neural state space. Additionally, we show kernel density estimates (Figure 6.1, a inset) to summarise the distribution of cross-correlations within each distance bin.

To quantify this effect, we computed, per sleep session, the correlation between pairwise forward distance and pairwise sleep cross-correlation, splitting neuron pairs by forward distance range. Correlation coefficients were negative for $0-180^\circ$ and positive for $180-360^\circ$ forward distances. Two-sided t-tests vs 0: $0-180^\circ$, $n=59$ sessions, $t=-5.16$, $P=1.07\times 10^{-4}$, d.f.=58; $180-360^\circ$, $n=59$ sessions, $t=4.16$, $P=1.08\times 10^{-4}$, d.f.=58. The difference between these two ranges was significant (paired t-test: $n=59$, $t=-5.55$, $P=7.35\times 10^{-7}$, d.f.=58). Bars show mean $\pm$ SEM.; individual session points are shown in the scatter.




Figure 6.1: Correlation structure in sleep sessions suggests ring-like dynamics persist offline within coherently remapping modules

(a) Sleep cross-correlations for neuron pairs that share the same anchor, plotted against forward distance (bottom x-axis) and circular distance (top x-axis). The schematics illustrate example neuron pairs and how they map to each distance measure. Cross-correlations follow a V-shaped profile with forward distance. That is, correlations are high near zero (and a full turn) and lowest around half a turn. This is consistent with a circular state space. Inset: kernel density estimates of cross-correlation values per bin. **(b)** Top - correlation coefficients of pairwise forward distance onto pairwise sleep cross-correlation for co-anchored neuron pairs. Neuron pairs with forward distances between 0° - 180° have significantly negative cross-correlation, whereas forward distances between 180° - 360° have significantly positive cross-correlation. Bottom - scatter plot of data from top. $N=59$ sleep sessions.

Our next analysis to test whether mFC contains ring-like subnetwork structures that persist offline, we employed a GLM-based method to account for other, non task-structure variables (Figure 6.2, b). Specifically, for a given pair of neurons, we regress their pairwise circular distance and forward distance (as determined by their task-space lags calculated in Figure 5.2) onto their

cross-correlation during sleep/rest. To account for other potential drivers of coactivity during sleep, we add as co-regressors the correlation between their spatial rate-maps as well as their circular goal-progress tuning distance. When we compare the regression coefficients for pairwise circular distance for neurons within a putative buffer against neurons between putative buffers, we find an inverse relationship between circular distance and sleep co-activity that was stronger within modules (Figure 6.2, c). That is, regression coefficients were significantly more negative for co-anchored pairs than for cross-anchor pairs (Welch's one-sided t-test; within: $n=430$ pairs; between: $n=13,932$ pairs; $t=-1.80$, $P=0.036$, $d.f.=14,360$). The negative coefficients here are easy to interpret. The further apart on a ring the neurons are, the less co-active they are during sleep.

Looking specifically at only co-anchored pairs (Figure 6.2, d), regression coefficients relating circular distance to sleep cross-correlation were significantly negative (two-sided one-sample t-test vs 0: $N=430$ pairs, $t=-2.66$, $P=0.008$, $d.f.=429$), whereas coefficients for forward distance did not differ from zero ($N=430$, $t=1.61$, $P=0.108$, $d.f.=429$). This indicates that during sleep, co-anchored neurons that are closer in the ring (small circular distance) are more co-active, whilst forward distance is not predictive.

These GLM results support the ring-structured state space account rather than the delay line account, and more generally, support the existence of internal structure to the representations rather than the results so far being purely driven by external sensory or motor drives during task performance.




Figure 6.2: GLM analysis suggests ring-like dynamics persist offline during sleep/rest within coherently remapping modules

(a) Two candidate circuit organisations for the ABCD task. Top: a ring-attractor arrangement in which neuron-neuron relationships are captured by *circular distance*. Bottom: a delay-line arrangement in which relationships are captured by *forward (sequence-lag) distance*. (b) Schematic of the generalised linear model used to dissociate the contributions of circular versus forward distance to pairwise co-activity. Goal-progress bin distance and spatial map correlation are included as co-regressors to absorb variance due to tuning similarity that is not specific to the neural state-space topology. (c) Regression coefficients for circular distance on sleep cross-correlation between neuron pairs that share an anchor (within-module) and pairs that do not (between-modules). $N=430$ pairs within-module, 13,932 pairs between-modules. Regression coefficients were significantly more negative for within-module pairs than between-module pairs. Negative coefficients here mean that neuron pairs with higher circular distance have lower sleep cross-correlation. (d) Regression coefficients for circular distance (left) and forward distance (right) for all co-anchored neuron pairs (within-module; $N=430$ pairs). Coefficients for circular distance are significantly negative; forward distance is non-significant against 0.

Interim summary: Neurons in a tSMB retain their ring-like dynamics offline

In sum, our offline sleep/rest analyses support the view that mFC contains internally-organised, ring-like subnetworks whose task-structured activity structure persist in the absence of task inputs. A GLM relating sleep/rest coactivity to pairwise distances on ring and delay-line state spaces showed that circular distance (on a ring) negatively predicted sleep cross-correlation, and this effect was significantly stronger for co-anchored pairs than for pairs spanning different anchors. Within co-anchored pairs, circular distance coefficients were significantly negative, whereas forward (delay-line) distance was not predictive, even after controlling for spatial map similarity and goal-progress proximity. A complementary, assumption-light analysis of cross-correlations binned by forward distance revealed the expected V-shape: negative session-wise correlations for $0\text{--}180^\circ$ and positive for $180\text{--}360^\circ$, consistent with a circular rather than linear structure. Together, these converging results argue against a simple delay-line account and indicate that the modular, ring-like geometry inferred during behaviour is preserved offline. This preservation mirrors other cortical–hippocampal attractor systems and suggests that mFC’s task-shaped buffers reflect intrinsic network structure rather than merely stimulus- or motor-evoked patterns. In the next section, we turn directly to predicting future actions.

6.2 Neurons in a tSMB predict distal behavioural choices

As these buffers are cyclical, activity in SMBs that have nearly completed a full rotation from their anchor can be read out to bias behaviour at the next step. Since different SMBs anchor different prior steps, their combined activity provides a computation for choosing an appropriate action at *every* step. The tSMB model makes a precise, testable behavioural prediction: if neurons form ring-like buffers anchored to specific behavioural steps, then the instantaneous population activity at time t contains information about which step will be revisited at a specific future lag. Concretely, for a neuron whose activity peaks at a lag X relative to its anchor, that neuron should predict a return to the same anchor at lag $(360^\circ - X)$ *in the future*. We refer to X as the neuron’s bump time.

We run into a problem if we view this model in purely deterministic (almost mechanical) terms. That is, as the SMBs are set off in a particular order as determined by behaviour (the sequence of anchor visits), a pure deterministic influence on future behaviour using these SMBs (a hard read-out) would result in exact repetitions of the first trial behaviour. We know from the actual observed behaviour that task performance improves over trials, with more mice taking more optimal routes as the sessions progress (Chapter 3). Conceptually, this can be resolved to some degree by simply acknowledging that actual neural implementations of this mechanism are likely to be imperfect and noisy, preventing such hard-readout-induced stereotypy. Further, top down controls on the population activity can also resolve this issue (such as gating or attentional controls, neuromodulatory inputs, or inputs to SMBs from sampled/imagined anchor visits). From an analytical standpoint, careful decoupling of the history of past choices from the influence of SMB activity on future behaviour will be crucial when attempting to test this claim.

Mice do not always take the same route between goals across trials. This is useful for analyses, as we can use this variability to disentangle representations of the past from the representation of the future. I will now detail the components of the tSMB model and how they will be exploited to analyse neural influences on future choices in the regime where there will be multiple SMBs competing to bias behaviour. In the schematic in Figure 6.3 we show two possible situations of what can happen when a mouse is at a 'decision point' on node 1. The mouse has two options, move east to node 2 (brown), or south to node 4 (red). In this situation, the SMB anchored to node 2 and node 4 for this particular goal progress bin will be competing to bias behaviour to their respective anchors. We put forward here that in this situation, the SMB with the highest activity just before the anchor neuron ($t = 3$ in the schematic) at the decision point will win out. In Figure 6.3 a, at $t = 1$, the mouse visits the anchor at node 2, initiating an activity bump in the brown SMB. The activity bump propagates around the ring, meaning that at $t = 3$ the activity bump in the node-2-anchored SMB is higher than the node-4-anchored SMB. Accordingly, node 2 is more likely to be chosen at the decision point. The opposite is true in Figure 6.3, b.

Importantly, according to this model, for both the SMB with the large and the small activity bumps, the activity around the ring is stable. This allows, in principle, the future behavioural bias to be detected by looking at any of the neurons on the ring, including those far from both the bump initiation and the decision point. For instance, by looking at the SMB activity from the right neurons at $t = 2$ in Figure 6.3 a, one would be able to determine that the activity bump is larger for the node-2-anchored SMB and surmise that the choice will be biased to that anchor revisit.

This is the motivating principle for the following analyses: that the activity of neurons at their bump time is predictive of revisits to their anchors. For what follows, we will use the task-space lags (synonymous with bump times), and their spatial anchors, determined by the single-anchor alignment analysis in Figure 5.2.




Figure 6.3: Schematic - distal prediction of behavioural choices from neurons in a tSMB

Two example tSMBs anchored at node 2 (brown) and node 4 (red), each with the same goal-progress tuning and task lag. At the decision point on node 1 (grey), each buffer's activity bump has nearly completed a full loop back to its anchor neuron. The two buffers' activity compete to bias the mouse to revisit their respective anchors (their goal-progress:location conjunction). With two competing tSMBs, the larger bump wins: in (a), the activity in the brown tSMB is greater than the red tSMB, meaning the mouse is biased towards choosing node 2. In (b) the activity bump in the red tSMB is greater than the brown tSMB, meaning the mouse is biased towards choosing node 4. Because bump amplitude is persistently larger around the ring within a trial (e.g., $t=1$, lag 0° ; $t=2$, lag 135° ; $t=3$, lag 340°), the upcoming choice can, in principle, be predicted from anywhere on the ring so long as activity is sampled at the neuron's bump time.

Bump time firing rates preceding anchor visits

The first analysis to get at this question is simple. Do neurons have higher bump-time firing rates in trials where they do revisit their anchor than when they do not? (Figure 6.4)

To control for choice history within a task, we split out consecutive trial pairs ($N - 1 \rightarrow N$) by anchor visits: 0:0 (no visit $\rightarrow$ no visit), 0:1 (no visit $\rightarrow$ visit), 1:0 (visit $\rightarrow$ no visit), and 1:1 (visit $\rightarrow$ visit) (Figure 6.4, a). By comparing bump-time activity in trial N across these four cases, we decouple SMB readout from simple repetition of the previous choice and ask whether neurons show elevated bump-time firing specifically on trials that end in a return to their anchor.

We can do this for single neurons using the anchors and task-space lags for significantly anchored neurons we computed in Figure 5.2. When we compare the normalised bump-time activity of all non-zero lag neurons (Figure 6.4, b), we find that trials ending in anchor visits have higher activity both when the anchor was *not* visited in the current trial ($0:1 > 0:0$; $n = 131$ tasks; Wilcoxon two-sided $W = 2749$, $P = 3.0 \times 10^{-4}$) and when it *was* visited ($1:1 > 1:0$; $n = 123$ tasks; $W = 2754$, $P = 0.008$). Further, a two-way ANOVA (factors: Past visit in N , Future visit in $N+1$; $N = 123$ tasks) showed a main effect of *Future* ($F(1, 122) = 6.06$, $P = 0.015$), a trend for *Past* ($F(1, 122) = 3.51$, $P = 0.063$), and no *Past* $\times$ *Future* interaction ($F(1, 122) = 0.80$, $P = 0.373$).

Next, we restricted the analysis to only include neurons with distal lags (at least one state, $\geq 90^\circ$ from the anchor; Figure 6.4, c). Like in the all non-zero lag neuron case, normalised bump-time firing was higher on trials leading to an anchor visit both when the anchor was *not* visited in trial N ($0:1 > 0:0$; $n = 128$ tasks; Wilcoxon two-sided $W = 2917$, $P = 0.004$) and when it *was* visited ($1:1 > 1:0$; $n = 115$ tasks; $W = 2534$, $P = 0.025$). A two-way ANOVA ($N = 115$ tasks) showed no main effect of *Past* ($F(1, 114) = 2.43$, $P = 0.122$), a main effect of *Future* ($F(1, 114) = 5.74$, $P = 0.018$), and no *Past* $\times$ *Future* interaction ($F(1, 114) = 0.33$, $P = 0.566$).

Together, these results show that non-anchor mFC neurons fire more strongly at their bump time on trials that culminate in a return to their anchor, irrespective of whether the anchor was

visited on the current trial, with the effect remaining when only considering neurons with distal lags. This is consistent with a prospective SMB-based bias to the behaviour, rather than mere perseveration. Next, we analyse this at the single-trial level using a logistic regression that predicts return-to-anchor on $N+1$ from each neuron's bump-time activity while controlling for choice history and contrasting against non-bump-time baselines.




Figure 6.4: Activity of mFC neurons is increased at bump time before distal anchor visits

(a) Distal choice prediction from SMBs while controlling for choice history. We exploit variability in how strongly an anchor visit drives the SMB bump to test whether bump-time activity predicts the next choice beyond trial history. Example SMB: anchor at location 2, intermediate goal progress. We show four scenarios across consecutive trials $N-1 \rightarrow N$:

- 0:0 no visit in $N-1$, low bump $\Rightarrow$ no visit in N ;
- 0:1 no visit in $N-1$, high bump (due to noise/top-down control) $\Rightarrow$ visit in N ;
- 1:0 visit in $N-1$, low bump $\Rightarrow$ no visit in N ;
- 1:1 visit in $N-1$, high bump $\Rightarrow$ visit in N .

This variability decouples SMB activity from prior choices, allowing us to ask whether bump-time activity predicts a return to the anchor. Time stamps t_1, t_2, t_3 are within trial N . **(b)** Left - The bump-time firing rate of non-zero lag neurons is higher when the mouse will visit that neuron's anchor in the upcoming trial than when it does not (0:0 vs 0:1). Both of these conditions are taken from trials when the mouse *had not* visited that anchor in the previous trial. Right - Normalised bump-time activity is higher when the mouse is about to visit the anchor on trials where the anchor *was* previously visited (1:0 vs 1:1). **(c)** Same analysis but only using neurons with distal non-zero lag from the anchor (>1 state - 90°). Left - Normalised bump-time activity of distal non-zero lag neurons is higher when the mouse is about to visit the anchor when the anchor *was not* visited on the previous trial (0:0 vs 0:1). Left - Normalised bump-time activity of distal non-zero lag neurons is higher when the mouse is about to visit the anchor when the anchor *was* visited on the previous trial (1:0 vs 1:1). $N=128$ tasks for distal non-zero lag neurons analysis; $N=131$ Tasks for all non-zero lag neurons analysis.

Logistic regression analysis

The previous analysis allowed us to assess bump time activity of neurons on trials with an anchor visit while controlling for one trial's behaviour (was an anchor just visited, or not?). To expand this and let us control for even more of the behavioural history, we next used a logistic regression approach to quantifying the predictive power of bump-time neural activity of tSMB neurons on future behavioural choices.

We first illustrate two example tasks where we are monitoring the activity of a single neuron in a tSMB anchored to a specific location at a fixed lag across two tasks (Figure 6.5, a). In task X, the firing rate of the neuron (green) with a task lag of 135° from its anchor at node 2 is predictive of a revisit to node 2 at a specific time point 225° in the future: when the bump-time firing rate is high, the anchor is revisited; when it is low, it is not. This remains true in task Y, despite the lag 135° bump time occurring when the mouse is in a different location on the maze. For the trials in these schematics, and in the analysis proper, the decision point occurs when the mouse is one location away from the anchor (a one-step transition). Trials in which the mouse is at a more distal location from the anchor at the decision point are not included.

We fit logistic regression models of bump-time neural activity onto future behavioural choices (Figure 6.5, b). As in the previous analysis, we only included neurons with consistent anchor $\times$ lag tuning across tasks (same anchor $\times$ lag in more than half of tasks; Figure 5.2). We labelled trials in which the anchor could have been revisited but was not as 0 (that is, the mouse was a one-step transition away from the anchor at the decision point, but moved on to another location instead). Trials in which the anchor was revisited we labelled as 1. As discussed previously, it is important to disentangle the impact of tSMB activity on upcoming behavioural choices from the history of previous behaviour, which can be stereotyped and autocorrelated in time. To account for this, we included the choice history from the previous 10 trials as co-regressors, allowing us to detect the variance explained by neurons on upcoming choices over and above these behavioural correlations.

When we fit logistic regression models on all non-zero lag consistently anchored neurons (Figure 6.5, c), we see significantly positive regression coefficients for bump-time neural activity on upcoming choices (two-sided t-test against 0, $N=131$ tasks, $t=2.74$, $P=0.007$, d.f.=130). Expectedly, we also see positive coefficients for the regressors of choices, but these decrease across trials in to the past (**n-1**: $N=131$, $t=7.39$, $P=1.60 \times 10^{-13}$, d.f.=130; **n-2**: $N=131$, $t=8.03$, $P=5.04 \times 10^{-13}$, d.f.=130; **n-3**: $N=131$, $t=4.77$, $P=4.80 \times 10^{-6}$, d.f.=130; **n-4**: $N=131$, $t=3.38$, $P=9.45 \times 10^{-4}$, d.f.=130; **n-5**: $N=131$, $t=4.36$, $P=2.58 \times 10^{-5}$, d.f.=130; **n-6**: $N=131$, $t=1.76$, $P=0.080$, d.f.=130; **n-7**: $N=131$, $t=2.83$, $P=0.005$, d.f.=130; **n-8**: $N=131$, $t=2.19$, $P=0.030$, d.f.=130; **n-9**: $N=131$, $t=2.42$, $P=0.017$, d.f.=130; **n-10**: $N=131$, $t=0.40$, $P=0.691$, d.f.=130).

According to the SMB model, the precise timing of the activity bump is crucial for predicting when the upcoming choice will be made. That is, if a neuron has a bump-time of lag 90° from its anchor, an increase in activity of that neuron should signal a return to its anchor in 270° time, and *not* at any other time. To test this component of the model, we fit additional regression models that replace bump-time firing rates with control non bump-time firing rates: specifically, from (i) the decision point, (ii) random times (iii, iv, v) one-state increments from the bump-time, 90° , 180° , 270° (Figure 6.5, d, e). Remarkably, across the population of all non-zero lag neurons, the only condition in which neural activity is positively predictive of upcoming choices is bump-time, and not any of the control times (two-sided t-tests against 0 for **bump-time**: $n=131$ tasks, $t=2.75$, $P=0.007$, d.f.=130; **decision point**: $n=131$, $t=-0.77$, $P=0.446$, d.f.=130; **random time**: $n=131$, $t=-0.79$, $P=0.433$, d.f.=130; **90° shifted time**: $n=131$, $t=-2.74$, $P=0.007$, d.f.=130; **180° shifted time**: $n=131$, $t=-1.47$, $P=0.143$, d.f.=130; **270° shifted time**: $n=131$, $t=-2.54$, $P=0.012$, d.f.=130.)

Furthermore, this effect remains even when we restrict the analysis to neurons with a lag of greater than $\pm 90^\circ$ from their anchor (Figure 6.5, e; two-sided t-tests against 0 for: **bump time**: $N=128$ tasks, $t=3.68$, $P=0.0$, d.f.=127; **decision time**: $N=128$ tasks, $t=-0.56$, $P=0.577$, d.f.=127; **random time**: $N=128$ tasks, $t=-0.29$, $P=0.769$, d.f.=127; **90° shifted time**: $N=128$ tasks, $t=-1.61$, $P=0.111$, d.f.=127; **180° shifted time**: $N=128$ tasks, $t=-0.61$, $P=0.543$, d.f.=127; **270° shifted time**:

N=128 tasks, $t=-2.75$, $P=0.007$, d.f.=127.)

In both the all non-zero lag and distal non-zero lag conditions, there are time bins wherein there are significantly *negative* coefficients (90° and 270° for all non-zero lag; 270° for distal non-zero lag). These findings are also consistent with the temporally-precise behavioural biasing predicted by the SMB model. If a neuron has high activity 90° (or 270°) out of step with what it ought to based on its task space lag, it implies that at choice time, the activity bump will not be in the correct position on the SMB to be read out to bias behaviour. A mis-positioned bump should reduce the influence of that SMB on choice, allowing other SMBs to win the competition and bias behaviour in other directions.

Interim summary: Neurons in a tSMB predict distal behavioural choices

In this section, we asked whether mFC activity carries a prospective, time-specific behavioural bias consistent with ring-like SMBs. Using neuron-specific anchors and lags estimated in previous analysis, we first showed that non-anchor neurons fire more at their bump time on trials that will return to their anchor, regardless of whether the anchor was just visited. This held even for distal lags ($>\pm 1$ state).

We then modelled single-trial choices with logistic regression, predicting revisits to anchors on the upcoming trial from bump-time firing while controlling for 10 trials of choice history. Bump-time weights were significantly positive, whereas the same neurons' activity at decision time, a random time, or at 90° , 180° , 270° relative to bump time did not predict anchor returns. The one-state increment offsets sometimes carried negative weights, which is what the SMB readout predicts if a misaligned bump fails to bias choice at the right time.

Together, these results show that mFC neurons provide an anchor-relative signal that predicts actions at the moment specified by each neuron's task-space lag. In other words, SMBs form a basis for selecting future actions, rather than merely storing a memory of previous ones.




Figure 6.5: Activity of mFC neurons predict distal future choices at the time specified by their position in a tSMB

(a) Schematic demonstrating the principle of predicting future choices from the activity of neurons in a tSMB. We show two tasks (X and Y) and two trials within those tasks. This tSMB is anchored to node 2 (brown) at middle goal-progress. The non-anchor neuron (green) has a task lag of 2.5 states (135° , t2) from the anchor, meaning the activity of that neurons should be predictive of whether the mouse revisits that anchor but specifically at that 135° bump-time. In example trial 1 (top), the activity bump for this tSMB is large, meaning the probability of a revisit to the anchor is high, whereas in example trial 2 (bottom) the activity bump is smaller and the anchor is not revisited. The tSMB structure generalises across tasks, meaning the non-anchor neuron's activity is predictive even when the location in which that neuron is active changes (node 8 in task 1, node 4 in task 2) **(b)** Schematic of logistic regression analysis. We regress neural activity onto the upcoming choice (0: anchor not visited; 1: anchor visited). We add choices from previous 10 trials as co-regressors to account for the effect of behavioural stereotypy. **(c)** Bump time neural activity regression coefficients are significantly positive. Co-regressors for previous choices are also significantly positive, but the effect decreases across trials into the past. **(d)** Using all non-zero lag neurons, regression coefficients are only significantly positive for neural activity taken from bump times, not any other control times. **(e)** Using only distal non-zero lag neurons, regression coefficients are only significantly positive for neural activity taken from bump times, not any other control times. N=128 tasks for distal non-zero lag neurons analysis; N=131 Tasks for all non-zero lag neurons analysis.

6.3 Summary

We set out expecting an abstract, task-state code in mFC that would generalise across tasks (fixed A/B/C/D tuning). Instead, remapping analyses showed that single neurons do not preserve absolute state identity. This led us to a different account wherein mFC contains modules of coherently remapping neurons anchored to concrete behavioural steps defined by a conjunction of place $\times$ goal-progress. We formalised this as task-Structured Memory Buffers (tSMBs): ring-like subnetworks in which activity bumps track task progress relative to an anchor and return after a full cycle.

Three lines of evidence show that neurons have mnemonic fields in task space (carry a memory of a previous behavioural step a specific number of task states ago):

- 1) Single-anchor alignment: aligning each neuron's activity to its best-fitting place $\times$ progress anchor revealed stable, cross-task structure that is obscured in absolute state coordinates. Anchor preferences reinstated in left-out tasks; anchor-relative lags tiled 0–360°, and anchors covered all place $\times$ progress combinations (mild centre bias late in the run).
- 2) Lagged spatial maps (“place cells of the past”): spatial maps computed at fixed task-space lags sharpened and generalised across tasks, including for non-zero-lag cells, demonstrating that many neurons encode locations visited a specific number of states ago.
- 3) GLM/elastic-net: firing was explained by weighted combinations of anchor $\times$ lag terms, and these weights predicted held-out task-space rate maps. Effects persisted for distal lags and after conservatively excluding any neuron with even weak trajectory tuning.

Offline sleep/rest analyses point to an intrinsic, ring-like organisation. In a GLM, circular distance negatively predicted sleep co-activity (stronger for co-anchored pairs) and this held after controlling for forward distance, place-map similarity, and goal-progress proximity. A complementary binning analysis showed the expected V-shape with forward distance (lowest

near 180° , higher near $0^\circ/360^\circ$), arguing against a delay-line and indicating that the modular, task-shaped geometry persists offline.

Finally, we found that tSMBs carry prospective information about behaviour. Using neuron-specific anchors and lags, non-anchor cells fired more at their bump time on trials that involve a return to their anchor, regardless of whether their anchor was visited on the current trial. Logistic regression at the single-trial level confirmed that bump-time activity predicted upcoming choices over and above 10 trials of choice history. Crucially, the same neurons' activity at decision time, random times, or 90° , 180° , 270° rotations from bump time, did not carry positive predictive power, and some offsets were negatively predictive. These results are consistent with a temporally-precise readout to bias behaviour.

Together, these results support an internally-organised, task-shaped architecture in mFC. Multiple ring-like buffers are anchored to distinct behavioural steps, whose constituent neurons encode stable lags around their anchors. This scaffold explains cross-task generalisation, enables temporally precise, distal prediction of choices, and persists offline, pointing to a genuine network organisation rather than mere sensory-motor drive.

7

Signatures of generalised spatial representations in mFC

Contents

7.1	Euclidean distance and node degree are dominant features of neural subspaces	137
7.2	Representational similarity analysis supports simultaneous concrete and generalised spatial coding	141
7.3	Decoder generalisation supports simultaneous concrete and generalised spatial coding	144
7.4	Summary	147

In the previous chapters, we have described mechanism by which mFC neurons solve new examples of the ABCD task using a conjunction of a spatial and task variables. For example, a neuron selective for an animal having collected reward at node 5 (a spatial variable), two and a half states ago (a task-structure variable). The inclusion of a spatial component in this code sits within a mixed literature on prefrontal spatial tuning. Some studies report clear spatial responses in human (Doeller et al., 2010), primate (Xie et al., 2022), and rodent (Bota et al., 2021) PFC, with units or population states for specific locations within an environment. By contrast, others argue that prefrontal activity “collapses” across environmental features and instead tracks abstract task variables that generalise across contexts. For the start, middle and end of a run in a cross-maze, but not the absolute ID of the North, East, South, and West arms of the maze

(Kaefer et al., 2020; Nardin et al., 2021, 2023), or corridors vs reward wells in a structured maze (Yu et al., 2018). However, these studies have various confounds that mean generalised spatial tuning-like responses can be explained by behavioural variables like progress-to-goal. The strong decoupling of task from world structure in the ABCD task addresses these confounds. Here, we will use various population-based methods to query the precise nature of our mFC spatial code.

Firstly, we can look at the spatial firing fields of individual neurons (Figure 7.1). We can see from these example neurons that they variously generalise over different environmental features. For instance, the *corner cell* has firing fields at nodes 1, 3, 7, and 9 but no other locations, and this preference is the same across two tasks. A *cardinal axis cell* has firing fields at nodes 2, 4, 6 and 8: the nodes on the perimeter of the maze that are not corners. The *centre cell* has a firing field on node 5 (though this is difficult to attribute to generalisation, because node 5 is the only node of its feature set so this is indistinguishable from a place cell). Furthermore, we see neurons sensitive to all nodes (*node cell*) and all edges (the bridges connecting the nodes; *edge cell*). For the three node-type sensitive neurons there are multiple different sets of features that could be driving this generalisation that are difficult to disentangle without further experiments. For example, they could be sensitive to *node degree* (the number of connections to other locations that location has) or they could be tuned to *centrality*. For now, we will not be making claims about which specific features are driving these effects.

These example cells hint at the existence of a generalised spatial code in mFC. But single units alone do not tell us which structure dominates. We will now begin using various population-based methods to determine the dominant spatial representations in mFC. Specifically, we will (i) investigate the structure of low-dimensional neural subspaces when controlling for low level kinematics, (ii) determine the spatial features that explain neural dissimilarity using model-based representational similarity analysis (RSA), and (iii) train support-vector-machine (SVM) decoders to read out both absolute location and feature classes in held-out data.




Figure 7.1: Example neurons that generalise over locations of the same abstract type

Five example neurons tune for multiple locations of the same type across two tasks. The *cardinal axis cell* fires at nodes 2, 4, 6, and 8 across two tasks. These nodes all have a degree of 3. The *corner cell* has peaks at nodes 1, 3, 7, and 9 across two tasks. These nodes all have a degree of 2. The *centre cell* has a peak selective for node 5 across two tasks - a location with degree of 4. Note that as this is the only location of degree 4, centre cells are indistinguishable from place cells. The *edge cell* has firing peaks at all of the bridges that connect the nodes, but not at the nodes themselves. The *node cell* is the opposite, with firing peaks at every node, but not the edges.

7.1 Euclidean distance and node degree are dominant features of neural subspaces

We next asked what form spatial information takes in mFC once known confounds are removed. To do this, we fit sparse linear models that predicted each neuron's firing from (i) the animal's current node, while (ii) explicitly controlling for one full trial of recent location history and (iii) instantaneous speed and acceleration (Figure 7.2a, left). Models were fit separately within goal-progress bins so that any structure we report here is not a by-product of progress tuning.

From each fit we extracted the coefficients on the nine current-location indicators, which can be read as history/kinematics-controlled spatial maps for each neuron (Figure 7.2a, middle). Embedding these maps with PCA (Figure 7.2, a right) revealed a compact, low-dimensional geometry in which locations arranged into clear, feature-based groupings (Figure 7.2, b). In particular, corner nodes (1, 3, 7, 9) clustered together, cardinal nodes (2, 4, 6, 8) formed a separate

cluster, and the central node (5) occupied a distinct position in the embedding. Indeed, for all 5 goal-progress bins, this feature-based separation dominated the first principal component of the embedding: node-degree/centrality explains the highest variance of the spatial subspace.

This feature-based separation is not the only spatial information displayed in these low-dimensional embeddings. When we include the second and third principal components, we are able to recover the true maze topology for all goal-progress bins (though this is admittedly weaker for the earliest goal-progress bin at the very start of a new trial). This suggests a simultaneous generalised and concrete representation of space in the mFC population.

Two immediate points are raised. First, controlling for one-trial history and kinematics leaves a robust spatial signal, arguing against a purely trajectory- or speed-based explanation. Second, the emergent subspace matches the single-unit patterns (corner-, cardinal-, and centre-preferring examples), co-existing simultaneously with a place-cell-like representation of absolute position.

Finally, in previous chapters we showed that many spatially-tuned neurons are conjunctive with a goal-progress code. We can test this explicitly here. A prediction of this is that spatial maps for different goal-progress bins should be orthogonal (or near-orthogonal) to each other. Here, we calculate the angles between the 3D spatial subspaces across goal-progress bins (Figure 7.2, d). In doing so, we find that angles between subspaces are near-orthogonal for distal pairs of goal progress bins (e.g. GP1 vs GP5), and this angle decreases as the GP bins come together. This is consistent with mFC having separate spatial maps for each goal-progress increment.




Figure 7.2: Concrete locations and node-degree are dominant features of neural subspace, and subspaces of distal goal-progress bins are near-orthogonal

Figure caption on next page.

Figure 7.2: Concrete locations and node-degree are dominant features of neural subspace, and subspaces of distal goal-progress bins are near-orthogonal

(a) Left - Schematic of Lasso ($L1$ -penalised) GLM analysis used to determine the population structure of mFC spatial representations. We regress current location (node 1-9) as well as the locations at 19 previous task lags (one full trial) onto neural activity. Location at each task lag (0-19) is encoded as a length-9 one-hot vector. We include speed and acceleration kinematic co-regressors to control for systematic differences in these variables across locations. The total number of regressors is therefore $2_{kinematic} + 9_{nodes} \times 20_{lags} = 182$. For each neuron, we concatenate activity across all tasks in a double day, and split the data into 5 goal-progress bins. Each regression uses data from just one goal-progress bin. Middle - We get for each neuron-double-day a $neuron \times regressors$ matrix of coefficients. Right - To find the axes in this high-dimensional space that explain the highest variance in spatial tuning, we perform PCA on the regressors for task lag 0 (the place code for *now*). Key: maze locations coloured by their node-degree/centrality. (b) Matrix of regression coefficients from aggregated data across all recording double days projected onto the first 2 principal axes. For all goal-progress bins, the first principal component separates locations based on their node-degree/centrality. (c) Projection of the regression coefficients onto the first 3 principal axes recovers the true maze topology. (d) Angles between 3D neural subspaces for pairs of goal-progress bins. Points in red are comparisons across goal-progress bins (e.g. goal-progress 1 subspace vs goal-progress 3 subspace). Points in grey are comparisons between subspaces built from regression coefficients of 50:50 splits of data from the same goal-progress bin. Subspaces for distal pairs of goal-progress bins are near-orthogonal. Data are mean $\pm$ SEM. $N=25$ recording double-days.

It is notable that there is a potential confound that can limit the interpretation of generalised spatial coding in a 3×3 grid maze layout with this subspace analysis. Specifically, a place cell-like code in which neurons have wide firing fields (but no generalised spatial coding) can give rise to the feature-based separation in the lower principal components of the spatial subspace. The rationale for this is simple: for a place-cell like population with wide firing fields, but not necessarily narrow firing fields, the more central locations will fall within the outer edges of the firing rate distributions even for neurons whose peak is at one of the outer nodes. This will make the overall population activity higher at the central nodes than the more peripheral nodes. We test this using simulated place cells. In the place cell population with wide firing fields, a feature-based separation comes out cleanly in PC3, while the maze topology is recovered in PC1 and PC2 (Figure 7.3). While this does not guarantee that the separation we observe in Figure 7.2 is artefactual – it is, after all, PC1 in the real data and only PC3 in the simulated data – it does motivate further population-based analyses that are not sensitive to this confound. This will be the focus of the following two sections.




Figure 7.3: A confound in detecting node-degree tuning with neural subspace analysis

Node-degree/centrality appears in neural subspaces with no generalisation of location types for place cell populations with wide firing fields. Left: Set of simulated place cells with narrow firing fields. PC1 vs PC2 roughly recovers the maze topology, where there is little discernable structure to PC2 vs PC3. Right: Set of simulated place cells with wide firing fields. PC1 vs PC2 captures the maze topology, while PC3 cleanly separates locations by node-degree/centrality despite this signal not being directly present in the simulated cell representations.

7.2 Representational similarity analysis supports simultaneous concrete and generalised spatial coding

Next, we will be using a representational similarity approach to confirm the suggestions of the subspace analysis: that there exists simultaneous concrete and generalised spatial coding in mFC. This method is not sensitive to the potential wide place field artefact (Figure 7.3). RSA asks whether the geometry of neural population responses matches the geometry predicted by a hypothesis (Kriegeskorte et al., 2008). Instead of decoding single conditions, RSA compares all pairs of conditions to build a representational dissimilarity matrix (RDM). Given activity patterns $\mathbf{r}_i$ and $\mathbf{r}_j$ for conditions i, j , a neural RDM stores $D_{ij} = \delta(\mathbf{r}_i, \mathbf{r}_j)$, where $\delta(\cdot)$ is a distance measure - in this case, we used the Euclidean distance between pairs of population vectors. Competing hypotheses are expressed as model RDMs (e.g., same place vs. different place; corner/cardinal/centre; node degree difference). Model fit is quantified by correlating or regressing the neural RDM with each model RDM. Because RSA operates on relative pattern

distances, it is agnostic to overall firing rates (or number of neurons, or even measure of neural activity - e.g. neuron firing rates vs BOLD signal in fMRI voxels), making it well-suited for testing abstract, distributed codes (including across species, modalities, and even against computational models). In short, RSA turns “which condition activates which neurons?” into “do the relationships among conditions in neural space match those implied by our hypothesis?”.

We constructed neural RDMs using the mean population vector at each node and calculating the Euclidean distance between pairs of mean population vectors across locations (Figure 7.4, a). To account for potential confounds, we (i) computed a separate RDM for each goal-progress bin and took the mean, and (ii) set a minimum speed threshold that removes stationary time bins and equalises the mean speed at each node type (one-way ANOVA results: F-statistic = 0.76, p-value = 0.46, d.f. (between-groups) = 2; Figure 7.5, a). Firstly, to visualise inter-location relationships, we perform multi-dimensional scaling on the aggregate data RDM. This RDM is made up of the mean across all mice, goal-progress bins, and recording double days. (Figure 7.4, c). Reassuringly, the feature-based separation of maze locations we observed in the regression subspace analysis (Figure 7.2, b) is replicated here. This is important as, in contrast to PCA (which could potentially be dominated by a global rate gradient due to wide place fields), this RSA setup focussing on patterns of neural dissimilarity is not sensitive to this artefact.

Next, we constructed model RDMs that capture the types of dissimilarity between locations predicted by concrete and generalised spatial codes (Figure 7.4, a right). Specifically, we produce RDMs of Euclidean distances between maze locations (absolute locations); difference in node degree (feature-based similarity); and a categorical place-type model (which is a more basic version of the feature-based code but with less structure). To quantify the degree to which these similarity structures contribute to the similarity of neural representations, we used a multivariate linear regression approach. Here, the model RDMs compete to explain variance in the neural data. Completing this regression on the combined data RDM, we extract regression coefficients for the three model RDMs and compare each to a null distribution as determined

by 100 row-wise permutations. We find that the coefficients for Euclidean distance and node-degree are higher than the 95th percentile of the null (Euclidean, $p < 0.001$; node-degree, $p = 0.004$), while the coefficient for the categorical place-type falls below this threshold ($p = 0.429$). We further quantify this effect across recording double days by fitting a separate regression for each (Figure 7.4, b). We find that the distribution of regression coefficients are significantly above 0 for Euclidean distance and node-degree (one sample, two-sided t-test against 0; Euclidean model, t -statistic=5.45, $p < 0.001$, d.f.=24; node-degree model, t -statistic = 3.51, $p = 0.002$, d.f.=24).

These RSA results suggest, like the previous subspace geometry analysis, the existence of simultaneous concrete and generalised spatial representations in mFC.




Figure 7.4: Both Euclidean distance and node-degree explain neural dissimilarity

(a) Representational similarity analysis setup. Left - data representational dissimilarity matrix computed by taking the Euclidean distance between pairs of mean population vectors for each location. The combined data RDM is the mean across recording double-days ($N=25$). Three model RDMs compete to explain variance in the combined data RDM in a linear regression (Euclidean Model: Euclidean distances between maze locations; Node-degree model: differences in node degrees between maze locations; Place-categorical model: binary matrix of 1 when locations are the same type - centre, cardinal, corner - and 0 elsewhere). Regression coefficients for Euclidean distance and node-degree are significantly above the null as determined by 100 row permutations of the data (seen beneath model RDMs). (b) Regression of model RDMs onto data RDMs are completed for each recording day separately ($N=25$). Across recording days, regression coefficients for Euclidean distance and node degree are significantly above 0. (c) Using multi-dimensional scaling for non-linear dimensionality redirection of the aggregate data RDM, maze locations are split out by node degree. This method is not sensitive to the place field artefact shown in Figure 7.3.

7.3 Decoder generalisation supports simultaneous concrete and generalised spatial coding

The subspace and RSA analyses suggest that mFC simultaneously carries a concrete place code (Euclidean layout) and a generalised code (node-degree). A complementary test is *decoding*: can these variables be read out from instantaneous population activity in held-out data, and how well do they generalise? The logic of this section is simple. First, if mFC has a spatial code that generalises across locations with similar feature sets (e.g. node-degree), then we should be able to decode the generalised features (e.g. corner vs cardinal; node vs edge) of a location that has been completely left out of the training data. Second, if mFC has a stable spatial code, then we should be able to decode the absolute location of a mouse in a held-out task. These two analyses are strong tests of the two spatial representation formats.

For all the analyses in this section, we used Support Vector Machine (SVM) decoders with radial basis function kernels and balanced class weightings. We schematise the three decoding analysis we performed in Figure 7.5, b. Furthermore, we reused the speed threshold we employed in the RSA analysis to remove the effect of systematic differences in running speed across locations (Figure 7.5, a).

First, we aimed to predict whether a held-out node is a corner (degree-2), or cardinal (degree-3) Figure 7.5 (b-top, c). The centre node is excluded from this analysis as it is the only location of its type (degree-4) and therefore cannot be decoded in a cross-validated manner. For each recording double day, we totally exclude one location from the training data (e.g. node 1) in a leave-one-out CV setup, and train on everything else across all 6 tasks. Each CV fold therefore only has test data with one label. To get around this, we report balanced accuracy of the predictions pooled across all CV folds. Remarkably, we see significantly above-chance balanced decoding accuracy for corners vs cardinals across mouse recording double-days (one-sample two-sided t-test against 0.5: statistic = 3.36, $p=0.002$, d.f. = 24). This is strong support for generalised spatial tuning.

Inspired by the example single units that appear to show tuning to nodes vs edges (Figure 7.1), we repeat this procedure but instead ask if we can decode whether a left-out location is a node or an edge (Figure 7.5, d). Using the same leave-one-location-out scheme (now withholding a single node or edge from training), we trained an SVM to classify withheld samples as node vs. edge and evaluated balanced accuracy pooled across folds (Figure 7.5 b, middle). As with corner–cardinal decoding, performance exceeded chance across recording double-days (one-sample two-sided t-test against 0.5; t-statistic = 4.63, $p = 0.0001$, d.f. = 24; Figure 7.5, d).

Finally, to test for a stable, concrete place code, we trained 9-way one-vs-rest SVM ensembles to decode absolute location in a leave-one-task-out regime (train on five tasks, test on the held-out sixth; Figure 7.5b, bottom). The decoding accuracy in the held out task was robustly above the chance level $1/9$, with a strong diagonal in the prediction confusion matrix (Figure 7.5, e-left), supporting the existence of a stable concrete place code. Further, we can also use the erroneous predictions of the place decoder to see whether mis-localisations are biased towards places of the same feature set (e.g. mistaking node 1, a corner, for node 3, also a corner). By removing the diagonal and node 5 and renormalising the confusion matrices (Figure 7.5, e-right), we can retrieve *error confusion matrices* that tell us about the structure of the decoder’s mistakes. As predicted by the previous analyses, SVM mis-localisations are significantly more likely to occur within place-types than between place-types across recording double-days (one-sample two-sided t-test against $chance = 3/7$: t-statistic = 4.32, $p=0.000$, d.f. = 24).

In sum, the decoding results concur with the single units, subspace geometry, and RSA analyses: mFC appears to hold a generalised and concrete spatial code simultaneously.




Figure 7.5: Decoding analysis supports simultaneous abstracted and concrete spatial tuning

(a) To account for the potential confound of systematic differences in running speed across locations, we set a minimum speed threshold. In doing so, we observe no significant differences in mean running speed across the corner, cardinal, or centre locations. (b) Schematics outlining the decoding analyses to probe generalisation across locations. Top - train a support vector machine (SVM) to decode whether a location is a corner or cardinal with one location totally held out of the training data. Test the model on this held-out location (one fold per held-out location). The centre node (node 5) was excluded from this analysis as it is the only location of its degree. Middle - train SVM to decode whether a location is a node or an edge. Test on a held-out location (one fold per held-out location). Bottom: Train SVM to decode location on the maze from all data except for one session. Test on a held-out session (one fold per held-out session). (c) Balanced accuracy significantly above chance when predicting corner vs cardinal ($N=25$ recording days). (d) Balanced accuracy significantly above chance when predicting node vs edge ($N=25$ recording days). (e) Left - mean confusion matrix across recording days shows maze location can be reliably decoded in held-out session. Right: mean 'error confusion matrix' - renormalised confusion matrix excluding the middle node and the diagonal. (f) Significantly above-chance proportion of decoding errors to locations of the same degree/type (i.e node 1 misclassified as node 3; both corners; $N=25$ recording days).

7.4 Summary

Here we used multiple complementary analyses to ask in what format mFC represents current spatial location. In doing so, we demonstrated that mFC holds co-existing generalised and concrete spatial codes.

First, after regressing out one-trial location history and kinematics, the coefficients of current location formed compact neural subspaces (Figure 7.2). Across goal-progress bins, the first principal axis reliably separated nodes by *node-degree/centrality* (corner/cardinal/centre), while the second–third axes recovered the maze geometry. Subspace maps derived from distal goal-progress bins were near-orthogonal, consistent with the conjunctive place×progress coding observed in single cells in Chapter 4.

Second, RSA confirmed this geometry while avoiding the wide–place-field confound. Neural dissimilarities were jointly explained by Euclidean distance and node-degree, whereas a coarse categorical place-type model alone did not account for the data (Figure 7.4). Multidimensional scaling of the neural RDM recapitulated the feature-based separation seen in the subspace analysis.

Third, SVMs generalised to withheld locations for corner vs. cardinal and node vs. edge classification, and also decoded absolute location across held-out tasks (Figure 7.5). Mislocalisations in the place decoder clustered within place-type, revealing structured, feature-consistent errors.

Taken together, these results indicate that mFC carries a mixed spatial code: a stable concrete place code suitable for precise localisation, and a *generalised* code aligned to structural properties of the environment (for example, node-degree/centrality).

A possible interpretation could be that mFC is representing the transition structure of the environment at multiple levels of abstraction. Planning a route requires knowing the connections

between locations. In our particular 3×3 grid maze, transitions can only be made between locations that are neighbours on the apparently ordinal node-degree axis. For example, from a route through the maze across node $1 \rightarrow \text{node } 2 \rightarrow \text{node } 5 \rightarrow \text{node } 8 \rightarrow \text{node } 9$ has a node degree sequence of $2 \rightarrow 3 \rightarrow 4 \rightarrow 3 \rightarrow 2$. There are no connections on the maze that violate this rule. A plan using such an abstracted transition structure would go from being a precise sequence of locations (e.g. $6 \rightarrow 9 \rightarrow 8 \rightarrow 5$) to more general steps (go to the corner, follow it round, head to the middle). A future experiment could test the idea that mFC's generalised spatial code is due to an abstracted transition structure by altering the connectivity of the maze so that, say, a degree 2 location is directly connected to a degree 4 location. Such a change should cause a reshuffling of the locations in the neural subspace, reconfiguration of the neural dissimilarity across locations, and lead to structured changes in decoder predictions. An alternative, speculative interpretation is that the generalised axis captures affordances or control-relevant similarities (e.g. branching options).

8

Entorhinal cortex in the ABCD task: preliminary data

Contents

8.1	Neuropixels implants in medial and lateral entorhinal cortices	151
8.2	Example MEC and LEC neurons in the ABCD task	153
8.3	Goal-progress tuning is stronger in LEC-targeted neurons than in MEC	157
8.4	Neither LEC nor MEC neurons consistently align to task state	159
8.5	Neither LEC nor MEC neurons remap coherently across tasks	161
8.6	Summary	165

In the early chapters of this thesis, we outlined our initial hypothesis that mFC neurons would track absolute position in the ABCD sequence and that this abstract representation would be bound to reward locations in other brain regions in each new task (Chapter 4). Theoretical work has suggested that there is reason to suspect that this representation does exist elsewhere in the brain (Whittington et al., 2025, 2020). By comparison to the factorised, associative solution, the SMB solution is radically more neuron-hungry. Let’s take a cartoon example wherein the four ABCD task states and the nine locations are represented discretely (i.e., no edges, goal-progress coding, or intermediate steps). The factorised solution requires four neurons for task state (one neuron for A, one neuron for B, one neuron for C, and one neuron for D), and nine neurons for locations. The solution to a new task is encoded in new *weights* between these sets

of neurons. This task solution uses just $N_{taskstates} + N_{locations} = 13$ neurons. By comparison, the SMB solution uses $N_{taskstates} \times N_{locations} = 36$ neurons. The SMB solution is effectively the product of the neurons used in the factorised solution, with one neuron per task state $\times$ location conjunction. Instead of the task being solved by changing synaptic connections, the solution is formed in the activations of the neurons in the SMB. As the resolution of the spatial and task state representation increases (e.g., by introducing goal progress), the number of neurons required for the SMB solution – even in the noiseless cartoon case with no regard for robustness – gets very large very quickly.

In the General Introduction (Chapter 1), we outlined two – potentially conflicting – bodies of evidence that implicate medial and/or lateral entorhinal cortices in representing absolute position in a task sequence. To briefly reiterate:

1. In Tolman-Eichenbaum-Machine style accounts, MEC provides a reusable map of latent structure, LEC inputs processed sensory information about objects, and hippocampus binds these two (Whittington et al., 2020). The prediction from this body of work is therefore that we should find the ABCD task schema in MEC.
2. Recent empirical work emphasises LEC’s robust structured temporal coding (Tsao et al., 2018), shifts at event boundaries (Issa et al., 2024; Kanter et al., 2025), and sequence position (Bellmund et al., 2019), in a manner that generalises across sensory realisations (Kanter et al., 2025). A prediction of these findings, if they extend to multiple goals in a structured sequence, is that LEC may hold the ABCD task schema.

In the remainder of this chapter we adjudicate between these possibilities. We show preliminary recordings from medial and lateral entorhinal cortices in mice performing the ABCD task to test explicitly whether either of these populations carry the factorised ABCD task schema representation.

8.1 Neuropixels implants in medial and lateral entorhinal cortices

We performed Neuropixels implants targeted to EC in two cohorts. In the first cohort (four mice attempted), one mouse yielded stable multi-session recordings with a Neuropixels 1.0 probe (Jun et al., 2017) targeted to MEC during many ABCD task instances. Post-hoc histology allowed us to identify the cortical entry point, but, owing to suboptimal sectioning, not the most distal tip of the track (Figure 8.1, a). While we are unable to track the full probe extent, the medial-lateral insertion coordinate and the maximum insertion depth are consistent with MEC localisation. Furthermore, we observe a strong 6-10 Hz theta rhythm in our recordings (Figure 8.1, b), consistent with – but not totally diagnostic of – MEC probe localisation. Given the uncertainty in the distal trajectory and the proximity of parasubiculum and adjacent cortices – which also exhibit strong theta (Sammons et al., 2019; Cacucci et al., 2004) and have similar spatial representations (Boccaro et al., 2010) – partial sampling of neighbouring regions cannot be ruled out. It is therefore appropriate to describe these units as *putative MEC*.




Figure 8.1: MEC histology and theta rhythm

(a) Histology slice for single mouse implanted with single-shank Neuropixels 1.0 probe targeting MEC. The probe insertion point is marked with the white arrow. Due to histology quality we are not able to locate the position of the probe tip in the brain. (b) Example 2.5 second window of firing rates for all neurons recorded in a double day. We see a strong theta frequency rhythm across neurons, consistent with - but not totally diagnostic of - the probe being located in MEC.

In a separate cohort, we targeted LEC with Neuropixels 2.0 probes (Steinmetz et al., 2021) in five

mice while they performed many ABCD task instances. For these implants, fluorescent DiI signal and track reconstruction provided clearer confirmation that the most ventral probe channels were located in LEC for all 5 mice (Figure 8.2). In the highlighted mouse, the shanks extended into the superficial layers of LEC (putative layer 2/3), but the probe terminus for the remaining 4 mice appears to lie in the deep layers (putative layer 5/6) based on the depth of the labelled track relative to the rhinal sulcus. While we will not be making laminar distinctions in the downstream analyses, the differences in anatomical projections to and from hippocampal and prefrontal subregions (Jun et al., 2024; Hisey et al., 2023; Ohara et al., 2021) are likely important for interpreting our LEC representations.




Figure 8.2: Histology - Neuropixels 2.0 targeted to lateral entorhinal cortex

Coronal sections for 5 mice implanted in lateral entorhinal cortex with 4-shank Neuropixels 2 probes. White trace is signal from DiI painted onto the shanks before implantation. Shanks are aligned along the A-P axis. Inset: recovered probe trajectories for highlighted mouse. LEC is labelled in blue on schematic.

While we were only able to obtain data from one mouse in the MEC cohort, we achieved high stable yields for all 5 mice in the LEC-targeted cohort (Figure 8.3), with a total of 2742 quality-controlled single units stable across double days (with an average of 18 concatenated recordings per double day). Note that this yield refers to all units detected from these recordings,

including units from more dorsal channels likely to fall outside of LEC. Later remapping and coherence analyses use a spatial filter to exclude these neurons.




Figure 8.3: High yields of neurons stable across double days in LEC-targeted cohort.

Total yields from LEC-targeted cohort, including units from more dorsal recording sites that likely fall outside of LEC itself.

8.2 Example MEC and LEC neurons in the ABCD task

Figure 8.4 illustrates task-space rate maps for example putative MEC neurons recorded across six ABCD tasks. Within each task, units showed clear modulation across the four-state cycle, often with multiple peaks per trial, indicating reliable within-task structure. However, these rate maps reorganised idiosyncratically across tasks. In these examples, we do not see consistent state tuning (e.g. no neurons that fire selectively for state *A* across all tasks), or goal-progress tuning. For example, Neuron 2 has a peak in state *D* late goal-progress in task 2, but state *B* early goal-progress in task 4.




Figure 8.4: Example MEC task-space rate maps across tasks

Example task-space rate maps across 6 tasks for neurons recorded from MEC. Rate maps are multi-peaked, remap across tasks, and do not appear to show conserved goal-progress tuning even within a task. Inset: anatomical location of MEC in mouse brain rendering.

Given MEC's well-documented spatial role, we next asked whether spatial tuning remained stable across ABCD task instances. Example units exhibited crisp single- and multi-field rate maps that persisted across tasks (Figure 8.5, a). To quantify stability, for each neuron we computed the mean Pearson correlation between its discretised spatial firing-rate maps across all task pairs. The distribution was markedly right-shifted from zero, indicative of robust cross-task stability (mean across neurons = 0.65, $n = 155$ MEC neurons; one-sample two-sided t -test against 0: t -statistic=44.34, $p < 0.0001$; Figure 8.5, b).

To contextualise these task recordings, a subset of recording days included sessions in a large open field arena ($1.2m \times 1.2m$). In these data, we observed neurons with multi-peak, approximately periodic spatial responses and others with boundary-concentrated firing consistent with border-like tuning (Figure 8.5, c). Although we do not yet quantify grid scores or border indices here, in future analyses these recordings will enable comparisons of how mEC spatial codes are shaped by our environment and task structure, and whether grid/border signatures are sharpened, distorted, or otherwise reorganised in the ABCD setting.




Figure 8.5: MEC spatial tuning is stable across tasks

(a) Example spatial rate maps for four MEC neurons across four tasks. Spatial firing peaks in these examples show a high degree of consistent spatial selectivity across time. (b) Distribution of mean cross-task spatial rate map correlations for all MEC neurons. There is a marked rightward shift from 0, indicating robust stable spatial tuning. (c) Four firing rate maps for example neurons recorded while a mouse forages for food reward in a large $1.2m \times 1.2m$ open-field arena. We see examples of border cells and multi-peak spatial cells.

Next, we examined task-space rate maps from the LEC-targeted cohort (Figure 8.6). Across animals, example neurons showed multi-peaked profiles within a task and remapping of peak locations across tasks, indicating that absolute ABCD state preference was not preserved. Despite this task-state remapping, the neurons appear to exhibit a stable goal-progress preference, similar to the majority of neurons in mFC.

In the following sections, we will look deeper into goal progress tuning and remapping across both brain regions.




Figure 8.6: Example LEC task-space rate maps across tasks

Task-space rate maps for four example neurons in LEC. Neurons are multi-peaked, remap across tasks, but do appear to have consistent goal-progress tuning. Inset: anatomical location of LEC in rendered mouse brain.

8.3 Goal-progress tuning is stronger in LEC-targeted neurons than in MEC

In our mFC recordings, most neurons tiled progress-to-goal; their firing fields spanned the 0–100% progress continuum with peaks distributed across bins and preserved across tasks (Figure 4.4). We next asked whether entorhinal cortex carries a comparable progress code and, if so, whether it differs between MEC and LEC (Figure 8.7).

In MEC (Figure 8.7, a), population heat maps showed weak cross-task preservation of progress phase: activity was generally low during stationary reward consumption and higher while the animal was running, yielding profiles consistent with kinematics rather than an abstract progress signal. While a more considered GLM analysis to confirm this hypothesis is yet to be completed, the dependence of MEC activity on running kinematics is well documented (Hinman et al., 2016; Dannenberg et al., 2019). Accordingly, peaks were sparse at very early and very late progress

bins (Figure 8.7, c).

By contrast, neurons from the LEC-targeted cohort exhibited robust, task-general progress tuning (Figure 8.7, b): peak bins collectively tiled the entire progress range with an over-representation at early and late phases (reward departure/approach; see Figure 8.7, c). Moreover, mean cross-task goal-progress correlations were markedly higher in LEC-targeted neurons than in MEC (Mann-Whitney U test: $U = 251324$, $p < 0.0001$; Figure 8.7, d). The over-representation of neurons with peaks at reward approach and reward departure is consistent with recent findings that LEC subpopulations segment experiences into distinct epochs surrounding rewards (Issa et al., 2024).

Together, these results suggest, but do not yet confirm, that broad MEC activity during ABCD is dominated by locomotor state (running vs stationary), whereas LEC carries a stable goal-progress code more akin to that seen in mFC (Figure 4.4).




Figure 8.7: Goal-progress tuning is stronger in LEC-targeted units than in MEC

(a) Normalised goal-progress rate maps for all neurons recorded in MEC ($n=155$). Maps are shown for four tasks and neurons are sorted by the bin of their peak firing rate in task 1. Goal progress tuning does not appear to be conserved. Instead, the firing rates are consistent with simple kinematics: low during stationary reward-consumption period and high during the period when the mouse is running. **(b)** Goal-progress heat maps for all neurons from the LEC-targeted cohort ($n=2742$ in total, including putative non-LEC neurons). Goal progress tuning is preserved, with some over-representation of early goal-progress and late goal progress (reward departure and reward approach). **(c)** Distribution of goal-progress peaks (averaged heat maps over all sessions). For LEC-targeted neurons, peak bins span the entire goal-progress spectrum, with over-representation around 20° and 70° . For MEC, there is a scarcity of neurons with peaks at early goal-progress ($<30^\circ$) and the latest goal-progress bins. **(d)** Distribution of mean cross-task Pearson correlation coefficients for MEC and LEC-targeted units. The LEC-targeted population has significantly higher goal-progress tuning than the MEC population.

8.4 Neither LEC nor MEC neurons consistently align to task state

As our first test of the hypothesis that LEC or MEC could hold the factorised ABCD sequence representation, we performed the cross-task remapping analysis we completed in Figure 4.5. Here, we use two filters to subset the neurons used. First, using the mean z-scored firing rate across task states for each neuron, we set a minimum entropy to remove neurons with no sensitivity to task state in over 50% of sessions. Second, for mice in the LEC-targeted cohort, we set a stringent spatial threshold to include only neurons in cortex rather than more dorsal regions

like ventral hippocampus. Specifically, we only included neurons with maximum signal on the bottom $150\mu\text{m}$ of the probe which, based on the probe placement identified in the histology, must be in LEC. These filters yielded 732 LEC neurons and 113 MEC neurons.

For both MEC and LEC, we calculated the distribution of cross-task remapping angles (Figure 8.8a, c). To reiterate, this aims to find the circular shift that maximises the correlation between the task X tuning and task Y tuning (Figure 8.8 inset). Neither the MEC nor LEC distributions have a peak at 0° , indicating that there is not an over-representation of generalising state cells. Interestingly, both the MEC and LEC distributions have just four strong peaks at the cardinal directions, rather than a more uniform circular distribution, despite the lack of goal-progress tuning in MEC. This is well explained by the speed modulation: rotating the map by an increment that is not a multiple of 90° is likely to place a peak into a stationary period, lowering the correlation.

In a complementary analysis, we quantified, for each neuron, the mean pairwise correlation of its 4D state-tuning vectors across tasks. Because the null distribution of the mean correlation depends on how many task vectors (K) are averaged over, we generated an empirical, K -matched null for each double-day by drawing K random 4-vectors (entries i.i.d. L2-normalised), computing their mean pairwise correlation, and repeating this 10,000 times. This yields, for each recording, the expected distribution of mean correlations if state vectors were unrelated across tasks. Comparing the observed neuron-wise means to the corresponding empirical nulls revealed no enrichment of high cross-task state consistency in either region: MEC (Mann–Whitney $U = 2,123,615$, $p = 0.912$; Figure 8.8, b) and LEC ($U = 57,224,045$, $p = 0.616$; Figure 8.8, d). Thus, we find no evidence for an over-representation of neurons that preserve an absolute ABCD state identity across tasks in entorhinal cortex.




Figure 8.8: No evidence of task-aligned generalising state cells in MEC or LEC

Inset: Schematic of cross-task remapping angle calculation. We find the circular rotation for a neuron's rate map in task X that maximises its correlation to the rate map in task Y. This is the neuron's remapping angle for that task pair. **(a)** Distribution of cross-task remapping angles for state-selective MEC neurons ($n = 113$). There is no peak around 0° indicating there is no over-representation of consistently state tuned neurons. **(b)** Distribution of mean cross-task state vector correlations for all state-selective MEC neurons. There is no significant difference between the data distribution (blue) and the null distribution (grey). **(c)** Distribution of cross-task remapping angles for state-selective spatially-filtered LEC neurons ($n = 732$). There is no peak around 0° indicating there is no over-representation of consistently state tuned neurons. **(d)** Distribution of mean cross-task state vector correlations for all spatially-filtered state-selective LEC neurons. There is no significant difference between the data distribution (blue) and the null distribution (grey).

8.5 Neither LEC nor MEC neurons remap coherently across tasks

Given the absence of single-neuron generalisation to absolute ABCD states demonstrated above, we asked whether pairs of neurons preserve their relative task-space relationship across tasks. That is, whether the population might rotate coherently as a whole. Such a coding scheme is conceptually the same as having an abstracted factorised task schema representation but does not lock to absolute task states, much like how grid cell modules remap relative to the environment but maintain their internal relationships. To do this, for each neuron pair we

estimated the between-neuron angle in task X via a correlation-maximising circular shift and compared it to the angle in task Y and Z .

We first generate distributions of remapping coherences for MEC (Figure 8.9, b) and LEC (Figure 8.9, f). The cross-task distributions of pairwise angle changes in both MEC and LEC showed no global alignment. Although the histograms were peaked at 0° , they also contained substantial mass elsewhere. Thus, there is no evidence for a single manifold that rotates rigidly across tasks.

However, for both of these distributions, we see a peak at 0° (reminiscent of the mFC data in Chapter 4). If this is truly partial coherence of task space remapping, we would expect to see coherent remapping for neurons for all circular distances (e.g. neuron pairs 180° apart in task space, as well as those 10° apart should rotate coherently). Indeed, we observe this phenomenon in mFC, which is one of the results that led to the modular account of mFC structure (Figure 4.6). If instead only closely tuned neurons remap together, this can be driven by purely spatial tuning (see Figure 8.10).

To test this, we labelled neuron pairs to be coherent if the change in angle was within $\pm 45^\circ$. We then tested, per recording day, whether the fraction of pairs that rotated coherently across 2 tasks exceeded the chance level of 6.25% ($25\% \times 25\%$) separately for pairs that were task-proximal in the reference task (“close”, initial separation $< 90^\circ$) versus task-distal (“far”, $\geq 90^\circ$).

In MEC (Figure 8.9, c), close pairs showed a strong above-chance coherence (two-sided one-sample t -test $t = 7.551$, $p < 0.0001$; $n = 6$ single recording days), whereas far pairs did not ($t = 0.396$, $p = 0.097$). Accordingly, the corresponding coherence distributions exhibited a clear peak at 0° for close pairs but not for far pairs (Figure 8.9, d).

In LEC (Figure 8.9, g), close pairs were likewise above chance ($t = 4.614$, $p < 0.001$; $n = 43$ single recording days). Far pairs showed only a small effect ($t = 2.034$, $p = 0.048$) that was weak

relative to close pairs and did not survive Bonferroni multiple-comparisons correction. Like in MEC, the distribution for far pairs lacked a pronounced 0° mode (Figure 8.9, h).

Together, these results argue against a global rotation of the entire EC population across tasks.




Figure 8.9: Above chance coherent remapping for task-space proximal, but not distal, MEC and LEC neurons

Inset: Schematic illustrating coherent vs incoherent remapping of neuron pairs. We use the correlating-maximising circular shift procedure to test whether between-neuron angles are preserved across tasks. **(a)** 3D rendering of MEC location in mouse brain. **(b)** Distribution of cross task between-neuron angle differences in MEC. There is a peak at 0° , which may imply a ring of task-aligned neurons rotating together. **(c)** Proportions of coherently remapping MEC neuron pairs separated by the initial between-neuron angle ($n = 6$ single recording days). Only 'close' neurons ($<90^\circ$ in task 1) have significantly above chance proportion of coherently remapping pairs. **(d)** Remapping coherence distributions for close and far neurons. There is only a peak at 0° in the close neurons distribution. **(e)** 3D rendering of LEC location in mouse brain. **(f, g, h)** LEC coherence remapping results as for MEC. Only close neurons have significantly above-chance coherently remapping pairs. ($n = 43$ LEC single days). Data are mean $\pm$ SEM.

Why do EC neuron pairs that are task-proximal (small circular distance) show strong remapping coherence, whereas task-distal pairs do not? We asked whether this pattern could arise from purely spatial tuning plus speed modulation, without any representation of the ABCD structure (Figure 8.10). We simulated a population of $N = 200$ neurons with stable spatial fields (each unit had 1–4 place fields randomly positioned on the maze, fixed across tasks), but no task-state tuning (Figure 8.10, a). For each of six task layouts, we generated optimal policy trajectories through the four rewarded locations and sampled spike trains from the spatial fields. To capture

well-known EC speed effects, firing was multiplicatively modulated by a simple speed profile (low during reward approach/consumption; high during post-reward acceleration). We then computed task-space rate maps and applied the same correlation-maximising circular-shift analyses as for the real data to obtain a cross-task remapping angle histogram (Figure 8.10, b), cross task neuron-pair coherence histogram (Figure 8.10, c), and the fraction of pairs with angle-change within $\pm 45^\circ$, split by initial between-neuron separation (“close” $< 90^\circ$ vs. “far” $\geq 90^\circ$; Figure 8.10, d).

Despite lacking any task-variable code, the simulation reproduced key features of the EC data. First, we observe a multi-modal remapping-angle distribution with modes at quarter turns but no excess alignment at 0° across tasks. Second, we see a sharp 0° peak in the coherence distribution. Third, we see above-chance coherence for task-proximal pairs, but at-chance coherence for task-distal pairs. Intuitively, neurons with overlapping spatial fields tend to be co-active along the same path segments and therefore show similar task-space firing patterns. This means that when the same paths are executed in different tasks (with similar approach/departure kinematics), their task-aligned rate maps rotate together, yielding local, but not global, coherence.

These results show that the pronounced close-pair coherence observed in EC can emerge from spatial tuning plus speed modulation alone, cautioning against interpreting EC coherence as evidence for a task-structured code.




Figure 8.10: MEC and LEC task-space remapping is consistent with purely spatial tuning

(a) A population ($n = 200$) of simulated neurons with speed modulation and random spatial peaks that are consistent across tasks. We generated task-space rate maps by assuming optimal policies for 6 tasks. (b) Distribution of cross-task remapping angles for purely spatial simulated neurons. (c) Distribution of pairwise remapping coherences. There is a peak at 0° despite there being no task structure representations built into the simulation. (d) Proportion of coherently remapping pairs across sessions for 'close' vs 'far' neurons. As in the MEC and LEC data, only close neurons have an above-chance proportion of coherently remapping pairs.

Taken together, these remapping results argue against the existence of a factorised ABCD schema in either medial or lateral entorhinal cortex, in either the task-state aligned format, or task-state unaligned but internally consistent format.

8.6 Summary

In this chapter, we collected Neuropixels data from both lateral and medial entorhinal cortex. We have not yet found evidence that either MEC or LEC expresses a task-invariant representation of the ABCD sequence that can be interpreted at the single-neuron level, although population manifold analyses, or accounting for temporal drift in LEC, may still recover the task schema.

Putative MEC units showed robust, cross-task *spatial* stability (Figure 8.5) but weak and kinematics-consistent goal-progress profiles (Figure 8.7 a, c). The LEC-targeted cohort yielded large, stable datasets and clear *goal-progress* tuning that generalised across tasks (Figure 8.7 b-d). However, neither region contained an over-representation of neurons that preserved absolute ABCD state identity across tasks (Figure 8.8). At the population level, MEC and LEC both exhibited *local* remapping coherence: task-proximal neuron pairs tended to rotate together, whereas *global* coherence was absent (Figure 8.9). A simple simulation with stable spatial fields

plus speed modulation reproduced these EC coherence signatures (Figure 8.10), indicating that local coherence can arise without an explicit task-structured code.

Important caveats remain, most notably the single-animal MEC dataset (though refinements to data preprocessing may yet be able to rescue data from the remaining mice from that cohort), laminar heterogeneity in LEC, and the absence of controls for temporal drift, or a task schema *manifold* rather than neuron-aligned representation.

A division of labour across MEC, LEC, and mFC

Our preliminary EC-mFC comparisons suggest a hierarchy for mapping behavioural structure. First, MEC carries a robust *where-you-are-now* signal: spatial rate maps are stable across ABCD tasks, with clear place-like responses and occasional open-field signatures (multi-peak/grid-like and border-like fields). Notably, our preliminary analyses find no evidence in MEC for a goal-progress code. MEC appears to localise the animal in allocentric space but not its progress to a goal or progress through a multi-goal sequence.

LEC shows goal-progress tuning, consistent with LEC's role in coding *what/when* features of episodes. However, this signal does not extend into a stable representation of multi-goal sequence position: across tasks we do not see coherent remapping or task-state alignment that would encode absolute ABCD state. In other words, LEC knows about goals, but not the *order* of multiple goals in a reusable, task-abstract format.

Medial frontal cortex holds both of these representations, with strong spatial tuning and goal-progress tuning. However, rather than just integrating these two representations, mFC has a component that neither entorhinal region shows: dynamics *sculpted by the 4-state loop task structure*. How this 4-reward structure is acquired by mFC remains an open question.

9

General discussion

Contents

9.1 Limitations and links to wider literature	169
9.2 Future directions	177

In this thesis, we have argued that flexible behaviour depends on neural representations – and algorithms that compute over them – that map both the structure of the world and the structure of behaviour. We recorded from multiple brain regions from mice performing a complex cognitive behaviour, and discovered a neural implementation of an algorithm for mapping behavioural structure that allows new sequences to be programmed into structured neural dynamics.

In Chapter 3 we introduced the ABCD task, a behaviour paradigm in which mice navigate between a repeating sequence of four reward locations in a grid maze. We showed that mice acquired a reusable schema for this four-step goal sequence: (i) they showed transfer learning-like improvement across problems, and (ii) inferred an unseen transition on the very first trial, demonstrating genuine structure transfer rather than rote memorisation (Chapter 3).

In Chapter 4 we started looking for neural representations that can support this generalisation behaviour. Large-scale recordings in medial frontal cortex revealed that most neurons tile

progress-to-goal, stretching and compressing their tuning to accommodate different inter-goal distances. Secondly, there were neurons with tuning to specific states within a task (but not across tasks). For a fraction of these state-selective neurons, state preference was explained by spatial tuning. Notably, we did not see neurons whose tuning is locked to specific task states across tasks. Instead, task tuning in mFC neurons remaps quasi-coherently in modules, wherein neurons preserve their relative arrangement within a module without locking to absolute ABCD states across tasks.

Motivated by these observations, in Chapter 5 we introduced *task-Structured Memory Buffers* (tSMBs): ring-like modules anchored to specific behavioural steps. Neurons in these modules fired at a fixed task-space lag from their anchor, rather than an explicit task state: a change in reference frame. This organisation compactly stores a rolling history of recent steps. Using single-anchor alignment, lagged spatial maps, and an elastic-net regression, we showed that neural activity in held out data can be predicted by anchor $\times$ lag conjunctive tuning, providing evidence for this ring-like modular structure.

In Chapter 6 we tested predictions of the tSMB model, specifically related to SMBs influencing behaviour. We find that the ring-like structure of SMBs persisted offline during sleep, consistent with intrinsic circuitry rather than task-structured sensorimotor drive. Critically, SMB activity predicts distal choices, with this predictive power being over and above the history of past choices. This means that instead of just storing a memory, the SMBs are *computing future actions*. SMBs therefore offer a programmable solution to mapping new ABCD task examples in neural dynamics without requiring any synaptic weight changes.

In Chapter 7 we dug deeper into mFC's spatial tuning. Using a set of complementary population-based analyses we find that mFC carries a simultaneous concrete spatial code that cares about absolute location, and a generalised spatial code that collapses over locations that share feature sets (e.g., "corners" vs "centre"). The computational relevance of this generalised coding is unclear.

Finally, in Chapter 8 we recorded from medial and lateral entorhinal cortex to look for signatures of the fully abstracted ABCD task schema. Preliminary analyses have not yet found this representation in either region. Instead, MEC showed a stable spatial code and LEC showed goal-progress tuning (like in mFC), suggestive of a cross-region functional hierarchy.

9.1 Limitations and links to wider literature

Immediate serial recall

tSMB-like algorithms have been proposed for other problems. For an example, it is worth at this point focusing in on a specific case of structure learning: immediate serial recall (ISR). ISR is the process of remembering a sequence of items and reproducing them in the order in which they were presented, like recalling a phone number or steps in recipe. Initial hypotheses about the computation being performed during this process suggested a form of chaining: each item cues the next through learned item→item links (Lewandowsky and Murdock, 1989). However, key behavioural results cast doubt on this solution. Specifically, the chaining solution implies that recall errors should cascade through the chain. If you get item 3 wrong, you should get the subsequent items wrong too. Instead, errors are typically local and recall can recover downstream.

This led to the introduction of factorised item $\oplus$ context models. Here, an ordinal ‘context’ signal exists independently of the representation of the items, rather than the items being directly linked to one another (Burgess and Hitch, 1992; Brown et al., 2000). That is, the ‘context’ represents start→middle→end (or A→B→C→D), and individual items bind to that context representation. This model does not suffer from the same cascading errors problem, but it failed to capture other features of recall error profiles, like errors being more likely for phonetically similar items than dissimilar ones. This setup is conceptually equivalent to our initial hypothesis (that we have found no support for) that mFC would house the full factorised ABCD abstraction.

In response, Botvinick and Plaut (2006) introduced a model that, instead of solving each new recall problem with associative binding, stored and recalled sequences using sustained patterns

of activation in a recurrent network without requiring new plasticity. This model learns mixed, conjunctive representations of item and position (e.g. a '4' at position 3 is represented differently to a '4' at position 2). The Botvinick and Plaut model is effectively an SMB solution to ISR. Therefore, SMBs can be considered to be providing a circuit basis for a kind of *structured sequence working memory*, with explicit representations of items at all future time steps or sequence positions.

Some direct evidence for this kind of coding on the single-neuron level exists in primate PFC. [Chiang and Wallis \(2018\)](#) recorded from macaque dlPFC during a serial visual search task and found that neurons had "spatiotemporal receptive fields". That is, they encoded both the spatial location of a target and its ordinal position within the search sequence. This has a conceptual equivalence to the anchor $\times$ lag conjunctive tuning we observe in mFC SMBs. The convergence across species (macaque dlPFC, mouse mFC) and task domains (visual search, spatial navigation) strengthens the case that item $\times$ future position conjunctive coding could be a general prefrontal mechanism for structured sequence representation, rather than a task-specific phenomenon.

Spacetime attractors

We initially framed our modular structure of mFC neurons as a structured mechanism for storing memories: representations of previously visited locations at a fixed task-space lag. However, the predictive power of neurons in these buffers on choices over and above the history of previous actions reframes this as being an explicit representation of the future (Figure 6.5). Additionally, the bump-time firing rates of SMB neurons is higher on trials which will include an anchor revisit regardless of whether the anchor was visited previously, effectively decoupling the past from the future (Figure 6.4). Recent theoretical work has demonstrated that having an explicit representation of the world at future steps, with each time-step in a separate activation slot, provides a natural mechanism for planning using simple attractor dynamics ([Jensen et al., 2025](#)). In the appendix (Chapter 11), we outlined two possible use cases for discrete attractor dynamics in the ABCD task. This is a third. All that is required for this to work is building in an extra layer of structure within and between activation slots. By wiring the world transition structure

in the connections between slots (e.g., node 1 at time-step 1 has excitatory connections to nodes 2 and 4 at time-step 2), then a simple representation for current location, and inputs that excite goal locations at future time-steps, will enable attractor dynamics to recover the shortest path between the current location and goal location. Here, the network's fixed points are entire paths through space and time; plans are represented as static assemblies (Buzsáki, 2010). In future work we will be looking for signatures of these attractor-like planning dynamics in SMBs, as well as for the specific connectivity profile such a mechanism requires.

Predicting future choices

In Chapter 5 we showed, using logistic regression, that the *bump-time* activity of consistently anchored SMB neurons carries information about upcoming choices (Figure 6.5). By embedding concrete behavioural steps within abstract task coordinates, structured memory buffers (SMBs) allow the medial frontal network to “programme” novel action sequences that share a common schema. In this view, the network's dynamics implement a mapping from task coordinates to forthcoming states, enabling rapid composition of new sequences without rebuilding the code from scratch. A critical point is that the SMB code that forecasts future steps is separable from its mnemonic (retrospective) tuning. We show this by adding 10 trials of past choices as co-regressors into the model. Bump time neural activity is predictive *over and above* these past choices. Although SMBs can be ‘loaded’ through experience in periodic tasks (such as ABCD-style paradigms), they could just as readily be loaded by internal simulation.

In the current work, we have not answered a complementary question: can we use instantaneous SMB activity to explicitly predict the entire future path of an animal in held-out data? A straightforward test would train multi-class decoders (e.g., one-vs-rest SVM) for each task-space future lag (e.g., one model for 10 degrees in the future, another for 70 degrees in the future, and so on). Performance would then be evaluated on a held-out session, with care taken to account for behavioural stereotypy and reward location constraints when computing chance levels. Practically, our current recordings may undersample buffers too much for this analysis to be fruitful, if we go by typical standards for (much simpler) decoding of current location from

hippocampal place cells (e.g., [Duvelle et al. \(2021\)](#)).

Unlike hippocampal splitter cells ([Frank et al., 2000](#); [Wood et al., 2000](#)) and other latent codes that disambiguate the same place across stereotyped sequences ([Dupret et al., 2010](#)), and unlike sequential encodings such as theta sweeps [Johnson and Redish \(2007\)](#) or awake replay in which a neuron often represents the same location now and moments ahead ([Pfeiffer and Foster, 2013](#)), our mFC code is explicitly prospective. At any instant, distinct mFC neurons are concurrently active for different future locations at different future lags, so the entire planned trajectory is present in the instantaneous population pattern. This has powerful consequences for downstream readouts. Using this mechanism, a downstream circuit can compute over the whole plan at once. Further, when a new problem shares the same abstract structure, the same SMB scaffold can be rapidly reconfigured by merely slotting new items into existing structured activity slots with no new plasticity. Such fast, schema-preserving reconfiguration provides a natural account of cross-task generalisation which may help explain why mFC is so critical for rapid task switching ([Birrell and Brown, 2000](#); [Owen et al., 1991](#)).

Internally-organised modular structure

In Chapter 6 we provided evidence that individual memory buffer modules retain their circular geometry offline during sleep. We did this with two separate but related analyses. (i) By directly computing cross-correlation from neurons with a shared anchor, we see a V-shaped curve with forward distance between neuron pairs. That is, the neurons with the *lowest* cross-correlation are those with 180° anchor-relative tuning difference, and cross-correlation increases as tuning forward difference decreases to 0° or increases to 360° . (ii) using a GLM, we regress forward and circular distance between neuron pairs onto sleep cross-correlation while controlling for spatial and goal-progress tuning. We see negative coefficients for circular distance, indicative of a circular neural state space.

Some of the most striking results for offline geometry preservation involve the ability to recover the whole manifold – for example, the torus of a grid cell module (Gardner et al., 2022). Such analysis in our mFC data is impracticable. Extracting the manifold structure of a grid module requires a very large number of neurons within that module to be recorded simultaneously (>100). The number of presumed SMBs in the ABCD task (even assuming no overlap) is vastly greater than the number of grid cell modules. Thus, the number of concurrently recorded neurons to extract the total SMB structure is orders of magnitude larger than our best cell yields.

Generalised spatial tuning

In Chapter 7 we detailed three separate population-based analyses that suggest mFC carries, simultaneously, a concrete spatial code that respects the true geometry of the maze, and a generalised spatial code that collapses over locations with similar features.

In both the subspace analysis (Figure 7.2) and multidimensional scaling of the aggregate data RDM (Figure 7.4) we observe an ordered structure to the feature-based separation that follows centrality/node degree. That is, an arrangement that follows corners→cardinal→centre. There is a risk of over-interpreting the ordered arrangement, over and above the mere separation. This is because, with the current set of analyses, we cannot fully include the potential for some of this order being driven by autocorrelation through time. A way of getting at this would be, as with the SMB representations, looking at coactivity structure in sleep in the absence of behavioural autocorrelation.

Our population analyses address, but do not fully eliminate, potential kinematic confounds. In the subspace analysis (Figure 7.2), movement variables were included as co-regressors, which reduces the chance that the low-dimensional structure we observe merely reflects speed or acceleration. However, in our RSA and decoding analyses, our kinematic control was to remove stationary bins with a speed threshold. This resulted in there being no significant differences in mean running speed across location types (Figure 7.5, a), but such filtering therefore attenuates only the most obvious discrepancies and may leave subtler dependencies intact. A more

considered approach would be to residualise each neuron's activity with a linear model as a preprocessing step. That is, we could use a linear model with kinematic regressors to subtract the component of the neural firing rates explained by those variables. Such an approach would bolster our claim that we are able to decode, say, 'corners' in held out data, rather than inadvertently just decoding running speed.

Dependence on external memory

Throughout this thesis we have drawn a distinction between (i) a factorisation-based algorithm, wherein an explicit representation of task state is flexibly re-bound to specifics in each new task example, and (ii) the SMB/activity slots algorithm, where neurons – or population subspaces – carry a conjunctive item $\times$ position conjunction that stores new examples in activations rather than weights. Recent theoretical work has put forward that which of the two algorithms a network learns depends, broadly, on access to an external memory system (Whittington et al., 2025). Specifically, if a recurrent network has access to a 'hippocampal' Hopfield network, then the recurrent network learns an abstract representation to guide memory retrieval. If, instead, it has no such access, the recurrent network learns the activation slots algorithm. Under this view, the relative anatomy of EC and PFC would predict that we would find the abstraction in EC, and SMBs in PFC, due to the presence or absence of dense connectivity with hippocampal CA3.

There is evidence supporting the view that entorhinal cortex represents the factorised ABCD code in currently unpublished data from humans performing the ABCD task. In fMRI representational dissimilarity structure we find signatures of the factorised ABCD representation in EC and OFC (Küchenhoff et al., 2025). These results are corroborated by single-unit recordings, where generalising state cells have been found in both the frontal and medial-temporal lobes (Küchenhoff et al., personal communication, 2025). Furthermore, in exciting but preliminary analyses, fMRI signatures of future position representation have been found in cingulate cortex (à la SMBs).

In our experiments in mice, we have not found factorised ABCD representations in LEC or MEC. This is despite both regions having dense connectivity with hippocampal CA3. These preliminary results point to an interesting cross-species distinction.

Why don't we see an ABCD schema in MEC and LEC?

In Chapter 8 we show that neither LEC nor MEC hold a representation of the factorised ABCD task schema. This was the case both for generalising state neurons that explicitly code for the same state across tasks, as well as a single coherently rotating population.

Several non-exclusive explanations could account for the absence of a clear, task-invariant ABCD representation in our MEC and LEC data. First, entorhinal cortex may participate primarily during learning, contributing to the formation of task-structured buffers in mFC. Once those buffers are established, there may be little computational pressure to retain an explicit, stable ABCD code in LEC and its contribution could revert to segmenting episodes and binding item- and event-specific details. The data presented in this thesis are from expert mice performing tasks 20-40. While we do not have neural data from the very first task, we do for tasks 10-20 (the start of the high-task regime) which could speak to this possibility.

Second, an abstract ABCD signal could exist at the population/manifold level but not align with single-cell axes. In that case, cell-wise remapping analyses could miss it. Recovering such a latent code would require across-task manifold alignment (e.g., Procrustes/CCA) and testing for task-state correspondences after alignment.

Third, pronounced temporal drift, a hallmark of LEC thought to reduce interference between episodes (Pilkiw et al., 2022), could decorrelate activity patterns across tasks despite a stable underlying computation. If LEC intentionally changes representations over time, cross-task state similarity would appear weak even if the same operations are performed. Analyses that factor

out slow drift could clarify this possibility.

Fourth, for MEC especially, a sampling limitation cannot be fully excluded. Our MEC dataset comprises a single high-yield mouse with uncertain distal track localisation, and although LEC targeting was stronger, laminar sampling varied across animals. It remains possible that a robust ABCD code resides in regions we under-sampled. Nevertheless, the convergent null results across multiple LEC mice make a strong schema encoded at the single-neuron level less likely.

Of course, the possibility that the representation does not exist at all in these regions cannot be ignored. This could be explained by evolutionary constraints on the sorts of inputs available to EC in mice, and these constraints preventing the sorts of arbitrary structure mapping suggested by theoretical work ([Whittington et al., 2020](#)).

Relating to LEC's goal-progress tuning, we have not yet completed the GLM-based analysis that accounts for lower-level potential confounds like total elapsed time or distance. Both LEC and HPC (from which we may have some units in the LEC-targeted dataset despite our filter to retain only units on the bottom-most $150\mu\text{m}$ of the probe) have temporal coding at the single unit level ([MacDonald et al., 2011](#); [Eichenbaum, 2014](#)). Thus, we need a more considered approach to fully confirm a progress code, rather than a temporal code.

It should also be noted that, as these results are only preliminary, we have not run the full suite of SMB-style analyses on the EC data that we did for the mFC data. Of course, the lack of coherent remapping in MEC and LEC makes the existence of SMBs in entorhinal cortex unlikely, but until these analyses are performed, it cannot be disregarded.

In sum, the current most parsimonious interpretation is that LEC emphasises goal progress signals rather than a stable, cell-aligned ABCD state code, and MEC codes strongly for current position. Future work should combine manifold alignment, drift-aware analyses, and

layer-resolved recordings to adjudicate between a latent, population-level schema and/or a learning phase contribution.

9.2 Future directions

Our results point to a set of tractable next steps that would deepen our understanding of the formation, computational power, and circuit mechanisms of structured memory buffers. We will summarise just a few of these directions below.

Hierarchical structure: nested buffers and composition

Do SMBs compose hierarchically (e.g., rings within rings)? An experimental extension training mice on *nested sequences*, like ABABCD_{CD}, would get at to what degree the SMB algorithm can be extended across hierarchical behavioural structures. If SMBs are compositional, we may predict concurrent bumps on multiple rings with coupling across rings.

Innate versus task-sculpted dynamics

Are some ring-like subnetworks preconfigured and later hijacked by learning, or are they entirely sculpted by experience? Two approaches can dissociate these possibilities. First, developmental/naïve recordings in juvenile mice during spontaneous locomotion and sleep can test for latent ring structure or proto-goal-progress signals before task training. Second, we could try structure-mismatch pretraining (train on non-circular tasks, then switch to ABCD) and ask whether pre-trained structures reconfigure into rings, or whether new modules are recruited. A strong prior for rings (due to, say, an ethological bias for periodic behaviours) would manifest as rapid emergence of circular geometry even under sparse task exposure.

The origin and computation of goal progress

Where does goal progress come from, and how is it computed? Candidate sources include medial temporal signals about distance-to-reward (time, steps, path distance). The stretching/compressing of our goal-progress signal with different goal distances suggests that there is some scaling input, making the activity bump move more slowly for longer distances and

more quickly for shorter distances. Joint-area recordings with mFC and areas that strongly input to mFC that are likely to contribute to this process (like ventral hippocampus) will allow us to observe what signals are being sent, and whether they are likely to be contributing to this scaling computation. A strong test would be optogenetic silencing; if goal-progress tuning breaks down, then that region likely contributes to the goal-progress computation.

Altering task and world structure

We already possess data from structural variants (e.g. task variant ABCDE, and adding new locations on to the maze). In a published version of the results in Chapter 3, Chapter 4, and Chapter 5, we report findings from mice that performed interleaved ABCD and ABCDE sequences after being trained on just ABCD. We find, in mFC, both preserved goal-progress tuning across task structures and state tuning for the E state riding on top of that goal progress signal (El-Gaby et al., 2024). Further, we have data from ABCDE + ABCD interleaved recording days for both the MEC and LEC experiments, though these are yet to be analysed. The data from the location addition experiment, once analysed, will test whether the SMB structure can immediately generalise to new spatial locations, suggesting some mechanism for compositionality.

Causal control of buffer dynamics and choice

If activity-bumps on an SMB bias choice, then targeted, timed perturbations should causally steer behaviour. This suggests a difficult, but powerful, set of experiments. All-optical interrogation (simultaneous calcium imaging and holographic photostimulation) of anchored mFC neurons should bias the behaviour of the animal towards anchor revisits, providing strong causal evidence of SMBs on behaviour (Packer et al., 2014; Zhang et al., 2018; Mardinly et al., 2018; Low et al., 2014; Russell et al., 2022).

Learning trajectories and sculpting SMB representation

The strength of goal-progress tuning in our mFC recordings, and the fact that the spatial and state tuning is so frequently conjunctive with it, suggest that goal-progress provides a scaffold onto which other variables are sculpted. This prompts us to ask how these codes form and stabilise over time. With longitudinal imaging, or careful spike-matching with UnitMatch

(van Beest et al., 2025), we can track single units across extended training. Specifically, we can ask if neurons migrate from kinematic responses to goal-progress tuning, and finally to the mature anchor $\times$ lag tuning as SMBs consolidate.

Medial-temporal – frontal interactions

Preliminary EC data suggest strong spatial codes (MEC) and robust goal-progress signals (LEC) without a stable ABCD state code. A next step is multi-region recordings to test directionality. How does frontal cortex learn its spatial model? Do hippocampal replays advance mFC ring dynamics during sleep? Do SMBs receive their anchor inputs from the hippocampus? In the first three mFC cohorts, a subset of the mice had dual mFC-HPC implants but the data were not included due to low individual unit counts. However, questions about the role of theta oscillations or consolidation due to ripple events during sleep can possibly be answered by our local field potential data in the absence of high single-unit yields.

Multiple tasks and interference

Brains must multiplex. What happens when the animal must master two or more concurrently active structures? We can interleave ABCD with an alternative structure (e.g., branching) or impose context cues that switch goals.

Manifold-level detection of abstract codes in EC

A possibility is that abstract ABCD codes exist at the population-manifold level in EC but are not cell-aligned. Future analyses should therefore go beyond single-cell remapping: align across-task EC manifolds (Procrustes/CCA), regress out slow temporal drift, and then test for preserved distances between nominal ABCD states. If an abstract code is present, alignment should reveal a consistent ring even when single neurons remap.

SMBs for planning: predictions from the spacetime attractor model

The spacetime attractors model (Jensen et al., 2025) suggests that having an explicit representation of the world both for *now* and for the *future* permits an attractor-based mechanism for planning-as-inference with static assemblies, rather than relying on, say, hippocampus-like roll-out

mechanisms (Jensen et al., 2024). Such a mechanism provides a cleaner readout for downstream regions (Buzsáki, 2010). As SMBs contain an explicit representation of locations at different time-steps, they provide a plausible neural substrate for planning with assemblies. Future work can test for signatures of planning in SMBs, and look for the world transition structure wiring across activation slots required for such an attractor mechanism to be implemented.

Summary

The programme sketched above raises concrete, testable questions about the origin, composition, and causal role of structured memory buffers in flexible behaviour. Much of this work is already ongoing, so many of these questions will be answered in the near future.

10

Outro

Spend enough time online and you'll probably come across the pop self-help concept of manifestation. That is, holding in mind a desired future state of the world as a means to bring that state about. Its proponents often claim that it works by some alignment with the energy of the universe, but we can imagine a less mystical account. If the paths to get to that desired state are somewhat clear, the prefrontal cortex — with its maps for now, and next, and after that — can load in that end state and some states in between. Then the gears start to turn. They tick-tick away with the odd metres and changing tempos of a free jazz quartet. With their cycles, actions arrive in your mind one by one like items on a supermarket conveyor. If you reach the end of that conveyor, congratulations. You've reached your goal. Synchrony with the universe, then, is not needed. It's just well-tuned neural clockwork.

Bibliography

- Akam, T., Lustig, A., Rowland, J. M., Kapaniaiah, S. K., Esteve-Agraz, J., Panniello, M., Márquez, C., Kohl, M. M., Kätzel, D., Costa, R. M., and Walton, M. E. (2022). Open-source, Python-based, hardware and software for controlling behavioural neuroscience experiments. *eLife*, 11:e67846. Publisher: eLife Sciences Publications, Ltd. pages 39
- Aksay, E., Olasagasti, I., Mensh, B. D., Baker, R., Goldman, M. S., and Tank, D. W. (2007). Functional dissection of circuitry in a neural integrator. *Nature Neuroscience*, 10(4):494–504. pages 203
- Amaral, D. G., Kondo, H., and Lavenex, P. (2014). An analysis of entorhinal cortex projections to the dentate gyrus, hippocampus, and subiculum of the neonatal macaque monkey. *Journal of Comparative Neurology*, 522(7):1485–1505. Author manuscript available in PMC; final edited form published 2014-05-01. pages 32
- Amari, S.-i. (1977a). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27:77–87. pages 202
- Amari, S.-I. (1977b). Neural theory of association and concept-formation. *Biological Cybernetics*, 26(3):175–185. pages 200
- Aronov, D., Nevers, R., and Tank, D. W. (2017). Mapping of a non-spatial dimension by the hippocampal–entorhinal circuit. *Nature*, 543. pages 16, 31, 206
- Bakermans, J. J. W., Warren, J., Whittington, J. C. R., and Behrens, T. E. J. (2025). Constructing future behavior in the hippocampal formation through composition and replay. *Nature Neuroscience*, 28:1061–1072. pages 14
- Banino, A., Barry, C., Uria, B., Blundell, C., Lillicrap, T., Mirowski, P., Pritzel, A., Chadwick, M. J., Degris, T., Modayil, J., Wayne, G., Soyer, H., Viola, F., Zhang, B., Goroshin, R., Rabinowitz, N., Pascanu, R., Beattie, C., Petersen, S., Sadik, A., Gaffney, S., King, H., Kavukcuoglu, K., Hassabis, D., Hadsell, R., and Kumaran, D. (2018). Vector-based navigation using grid-like representations in artificial agents. *Nature*, 557(7705):429–433. pages 14
- Bao, X., Gjorgieva, E., Shanahan, L. K., Howard, J. D., Kahnt, T., and Gottfried, J. A. (2019). Grid-like Neural Representations Support Olfactory Navigation of a Two-Dimensional Odor Space. *Neuron*, 102(5):1066–1075.e5. pages 16, 31, 206
- Baram, A. B., Muller, T. H., Nili, H., Garvert, M. M., and Behrens, T. E. J. (2021). Entorhinal and ventromedial prefrontal cortices abstract and generalize the structure of reinforcement learning problems. *Neuron*, 109. pages 22, 28, 31, 77
- Barone, P. and Joseph, J. . P. (1989). Prefrontal cortex and spatial sequencing in macaque monkey. *Exp. Brain Res.*, 78. pages 22

- Bartlett, F. C. (1933). Remembering: A study in experimental and social psychology. *Philosophy*, 8(31):374–376. pages 2, 10
- Basu, R., Gebauer, R., Herfurth, T., Kolb, S., Golipour, Z., Tchumatchenko, T., and Ito, H. T. (2021). The orbitofrontal cortex maps future navigational goals. *Nature*, 599(7885):449–452. pages 24
- Bateson, G. (1972). *Steps to an Ecology of Mind*. Chandler, San Francisco. pages 9
- Behrens, T. E. J., Muller, T. H., Whittington, J. C. R., Mark, S., Baram, A. B., Stachenfeld, K. L., and Kurth-Nelson, Z. (2018). What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior. *Neuron*, 100(2):490–509. pages 12, 15, 16, 17
- Bellmund, J. L., Deuker, L., and Doeller, C. F. (2019). Mapping sequence structure in the human lateral entorhinal cortex. *eLife*, 8:e45333. Publisher: eLife Sciences Publications, Ltd. pages 32, 33, 150
- Ben-Yishai, R., Bar-Or, R. L., and Sompolinsky, H. (1995). Theory of orientation tuning in visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 92(9):3844–3848. pages 201, 202, 204
- Bernardi, S., Benna, M. K., Rigotti, M., Munuera, J., Fusi, S., and Salzman, C. D. (2020). The Geometry of Abstraction in the Hippocampus and Prefrontal Cortex. *Cell*, 183(4):954–967.e21. pages 22
- Berners-Lee, A., Wu, X., and Foster, D. J. (2021). Prefrontal Cortical Neurons Are Selective for Non-Local Hippocampal Representations during Replay and Behavior. *The Journal of Neuroscience*, 41(27):5894. pages 16
- Birrell, J. M. and Brown, V. J. (2000). Medial frontal cortex mediates perceptual attentional set shifting in the rat. *J. Neurosci.*, 20. pages 17, 19, 172
- Boccaro, C. N., Sargolini, F., Thoresen, V. H., Solstad, T., Witter, M. P., Moser, E. I., and Moser, M.-B. (2010). Grid cells in pre- and parasubiculum. *Nature Neuroscience*, 13(8):987–994. pages 17, 151
- Bongioanni, A., Folloni, D., Verhagen, L., et al. (2021). Activation and disruption of a neural mechanism for novel choice in monkeys. *Nature*, 591:270–274. pages 17, 20, 206
- Borges, J. L. (1999). *Collected Fictions*. Penguin Books, first edition. pages 1
- Bota, A., Goto, A., Tsukamoto, S., Schmidt, A., Wolf, F., Luchetti, A., Nakai, J., Hirase, H., and Hayashi, Y. (2021). Shared and unique properties of place cells in anterior cingulate cortex and hippocampus. *bioRxiv*, page 2021.03.29.437441. pages 15, 135
- Botvinick, M. M. and Plaut, D. C. (2006). Short-term memory for serial order: a recurrent neural network model. *Psychological review*, 113(2):201–233. Place: United States. pages 30, 169
- Broadbent, D. E. (1958). *Perception and Communication*. Pergamon, London. pages 4
- Brodmann, K. (1909). *Vergleichende Lokalisationslehre der Großhirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues*. Johann Ambrosius Barth, Leipzig. Classic cytoarchitectonic atlas; English translation by L. J. Garey, Springer, 2006. pages 26

- Broschard, M. B., Kim, J., Love, B. C., Wasserman, E. A., and Freeman, J. H. (2021). Prelimbic cortex maintains attention to category-relevant information and flexibly updates category representations. *Neurobiology of Learning and Memory*, 185:107524. pages 19
- Brown, G. D., Preece, T., and Hulme, C. (2000). Oscillator-based memory for serial order. *Psychological Review*, 107(1):127–181. pages 169
- Burak, Y. and Fiete, I. R. (2009). Accurate Path Integration in Continuous Attractor Network Models of Grid Cells. *PLOS Computational Biology*, 5(2):e1000291. Publisher: Public Library of Science. pages 14, 205
- Burgess, N. and Hitch, G. J. (1992). Toward a network model of the articulatory loop. *Journal of Memory and Language*, 31(4):429–460. pages 169
- Butler, W. N., Hardcastle, K., and Giocomo, L. M. (2019). Remembered reward locations restructure entorhinal spatial maps. *Science*, 363(6434):1447–1452. pages 17
- Buzsáki, G. (2010). Neural syntax: cell assemblies, synapsembles, and readers. *Neuron*, 68(3):362–385. pages 171, 180
- Cacucci, F., Lever, C., Wills, T. J., Burgess, N., and O'Keefe, J. (2004). Theta-Modulated Place-by-Direction Cells in the Hippocampal Formation in the Rat. *The Journal of Neuroscience*, 24(38):8265. pages 151
- Canto, C. B., Wouterlood, F. G., and Witter, M. P. (2008). What does the anatomical organization of the entorhinal cortex tell us? *Neural Plasticity*, 2008(1):381243. pages 31, 32
- Caruana, R. (1997). Multitask learning. *Machine Learning*, 28(1):41–75. pages 70
- Chaudhuri, R., Gerçek, B., Pandey, B., Peyrache, A., and Fiete, I. (2019). The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nature Neuroscience*, 22(9):1512–1520. pages 204, 205
- Chen, D., Axmacher, N., and Wang, L. (2024a). Grid codes underlie multiple cognitive maps in the human brain. *Progress in Neurobiology*, 233:102569. pages 17, 206
- Chen, J., Zhang, C., Hu, P., Min, B., and Wang, L. (2024b). Flexible control of sequence working memory in the macaque frontal cortex. *Neuron*, 112(20):3502–3514.e6. pages 25, 29
- Chiang, F. . K. and Wallis, J. D. (2018). Spatiotemporal encoding of search strategies by prefrontal neurons. *Proc. Natl Acad. Sci. USA*, 115. pages 170
- Chomsky, N. (1959). A Review of B. F. Skinner's *Verbal Behavior*. *Language*, 35(1):26–58. pages 4
- Cohen, J. D., Dunbar, K., and McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review*, 97(3):332–361. pages 18
- Compte, A., Brunel, N., Goldman-Rakic, P. S., and Wang, X.-J. (2000). Synaptic Mechanisms and Network Dynamics Underlying Spatial Working Memory in a Cortical Network Model. *Cerebral Cortex*, 10(9):910–923. pages 200
- Constantinescu, A. O., O'Reilly, J. X., and Behrens, T. E. J. (2016). Organizing conceptual knowledge in humans with a gridlike code. *Science (New York, N.Y.)*, 352(6292):1464–1468. pages 16, 31, 206

- Dang, W., Jaffe, R. J., Qi, X.-L., and Constantinidis, C. (2021). Emergence of nonlinear mixed selectivity in prefrontal cortex after training. *Journal of Neuroscience*, 41(35):7420–7434. pages 20
- Danjo, T., Toyozumi, T., and Fujisawa, S. (2018). Spatial representations of self and other in the hippocampus. *Science*, 359(6372):213–218. pages 14
- Dannenberg, H., Kelley, C., Hoyland, A., Monaghan, C. K., and Hasselmo, M. E. (2019). The Firing Rate Speed Code of Entorhinal Speed Cells Differs across Behaviorally Relevant Time Scales and Does Not Depend on Medial Septum Inputs. *The Journal of Neuroscience : the official journal of the Society for Neuroscience*, 39(18):3434–3453. Place: United States. pages 157
- den Bakker, H. and Kloosterman, F. (2025). Neurons in the medial prefrontal cortex that are not modulated by hippocampal sharp-wave ripples are involved in spatial tuning and signaling upcoming choice. *eLife*. pages 16
- Deshmukh, S. S. and Knierim, J. J. (2011). Representation of non-spatial and spatial information in the lateral entorhinal cortex. *Frontiers in behavioral neuroscience*, 5:69. Place: Switzerland. pages 32
- DeVito, L. M. and Eichenbaum, H. (2011). Memory for the Order of Events in Specific Sequences: Contributions of the Hippocampus and Medial Prefrontal Cortex. *The Journal of Neuroscience*, 31(9):3169. pages 24
- DeVito, L. M., Kanter, B., and Eichenbaum, H. (2010). The hippocampus contributes to memory expression during transitive inference in mice. *Hippocampus*, 20(1):208–217. pages 13
- Doeller, C. F., Barry, C., and Burgess, N. (2010). Evidence for grid cells in a human memory network. *Nature*, 463(7281):657–661. pages 15, 16, 135
- Dordek, Y., Soudry, D., Meir, R., and Derdikman, D. (2016). Extracting grid cell characteristics from place cell inputs using non-negative principal component analysis. *eLife*, 5:e10094. pages 16
- Dragoi, G. and Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature*, 469. pages 116
- Dupret, D., O'Neill, J., Pleydell-Bouverie, B., and Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nat. Neurosci.*, 13. pages 172
- Dusek, J. A. and Eichenbaum, H. (1997). The hippocampus and memory for orderly stimulus relations. *Proc. Natl Acad. Sci. USA*, 94. pages 13, 16
- Duvelle, É., Grieves, R. M., Liu, A., Jedidi-Ayoub, S., Holeniewska, J., Harris, A., Nyberg, N., Donnarumma, F., Lefort, J. M., Jeffery, K. J., Summerfield, C., Pezzulo, G., and Spiers, H. J. (2021). Hippocampal place cells encode global location but not connectivity in a complex space. *Current Biology*, 31(6):1221–1233.e9. pages 172
- Eichenbaum, H. (2014). Time cells in the hippocampus: a new dimension for mapping memories. *Nature Reviews Neuroscience*, 15:732–744. pages 176
- Eichenbaum, H. (2017a). On the integration of space, time, and memory. *Nature Neuroscience*, 95(5):1007–1018. pages 7

- Eichenbaum, H. (2017b). Prefrontal–hippocampal interactions in episodic memory. *Nat. Rev. Neurosci.*, 18, pages 16
- Eichenbaum, H., Dudchenko, P. A., Wood, E. R., Shapiro, M., and Tanila, H. (1999). The hippocampus, memory, and place cells: Is it spatial memory or a memory space? *Neuron*, 23(2):209–226. pages 7, 16
- Eichenbaum, H., Stewart, C., and Morris, R. G. M. (1990). Hippocampal representation in place learning. *Journal of Neuroscience*, 10(11):3531–3542. pages 12
- El-Gaby, M., Harris, A. L., Whittington, J. C. R., Dorrell, W., Bhomick, A., Walton, M. E., Akam, T., and Behrens, T. E. J. (2024). A cellular basis for mapping behavioural structure. *Nature*, 636(8043):671–680. pages 89, 178, 200
- Ferster, C. B. and Skinner, B. F. (1957). *Schedules of Reinforcement*. Appleton-Century-Crofts, New York. pages 7
- Finn, C., Abbeel, P., and Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pages 1126–1135. pages 9
- Flesch, T., Juechems, K., Dumbalska, T., Saxe, A., and Summerfield, C. (2022). Orthogonal representations for robust context-dependent task performance in brains and neural networks. *Neuron*, 110(7):1258–1270.e11. pages 21
- Frank, L. M., Brown, E. N., and Wilson, M. (2000). Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron*, 27. pages 63, 172
- Gardner, R. J., Hermansen, E., Pachitariu, M., Burak, Y., Baas, N. A., Dunn, B. A., Moser, M.-B., and Moser, E. I. (2022). Toroidal topology of population activity in grid cells. *Nature*, 602(7895):123–128. pages 31, 116, 173, 201, 205, 206
- Garvert, M. M., Dolan, R. J., and Behrens, T. E. (2017). A map of abstract relational knowledge in the human hippocampal–entorhinal cortex. *eLife*, 6:e17086. Publisher: eLife Sciences Publications, Ltd. pages 31
- Gauthier, J. L. and Tank, D. W. (2018). A dedicated population for reward coding in the hippocampus. *Neuron*, 99(1):179–193.e7. pages 17
- George, D., Rikhye, R. V., Gothoskar, N., Guntupalli, J. S., Dedieu, A., and Lázaro-Gredilla, M. (2021). Clone-structured graph representations enable flexible learning and vicarious evaluation of cognitive maps. *Nature Communications*, 12(1):2392. pages 14, 16, 17, 30, 97
- Gershman, S. J. and Niv, Y. (2010). Learning latent structure: carving nature at its joints. *Curr. Opin. Neurobiol.*, 20(2). pages 22
- Gil, M., Ancau, M., Schlesiger, M. I., Neitz, A., Allen, K., De Marco, R. J., and Monyer, H. (2018). Impaired path integration in mice with disrupted grid cell firing. *Nature Neuroscience*, 21(1):81–91. pages 13
- Granon, S. and Poucet, B. (2000). Involvement of the rat prefrontal cortex in cognitive functions: a central role for the prelimbic area. *Psychobiology*, 28(2):229–237. pages 19

- Grosmark, A. D. and Buzsáki, G. (2016). Diversity in neural firing dynamics supports both rigid and learned hippocampal sequences. *Science*, 351. pages 116
- Guldin, W. O., Pritzel, M., and Markowitsch, H. J. (1981). Prefrontal cortex of the mouse defined as cortical projection area of the thalamic mediodorsal nucleus. *Brain, Behavior and Evolution*, 19(3-4):93–107. pages 26
- Guttman, N. and Kalish, H. I. (1956). Discriminability and stimulus generalization. *Journal of Experimental Psychology*, 51(1):79–88. pages 7
- Hafting, T., Fyhn, M., Molden, S., Moser, M. . B., and Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436. pages 13
- Hansel, D. and Sompolinsky, H. (1998). Modeling feature selectivity in local cortical circuits. In Koch, C. and Segev, I., editors, *Methods in Neuronal Modeling: From Synapses to Networks*, pages 499–567. MIT Press, Cambridge, MA, 2 edition. pages 202
- Hardcastle, K., Ganguli, S., and Giocomo, L. M. (2015). Environmental boundaries as an error correction mechanism for grid cells. *Neuron*, 86(3):827–839. pages 202
- Harlow, H. F. (1949). The formation of learning sets. *Psychological Review*, 56(1):51–65. pages 7, 8, 9
- Hinman, J. R., Brandon, M. P., Climer, J. R., Chapman, G. W., and Hasselmo, M. E. (2016). Multiple Running Speed Signals in Medial Entorhinal Cortex. *Neuron*, 91(3):666–679. Publisher: Elsevier. pages 157
- Hisey, E., Purkey, A., Gao, Y., Hossain, K., Soderling, S. H., and Ressler, K. J. (2023). A Ventromedial Prefrontal-to-Lateral Entorhinal Cortex Pathway Modulates the Gain of Behavioral Responding During Threat. *Biological psychiatry*, 94(3):239–248. Place: United States. pages 32, 152
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558. pages 200, 201
- Huang, X., Schlesiger, M. I., Barriuso-Ortega, I., Leibold, C., MacLaren, D. A. A., Bieber, N., and Monyer, H. (2023). Distinct spatial maps and multiple object codes in the lateral entorhinal cortex. *Neuron*, 111(19):3068–3083.e7. pages 32
- Hull, C. (1943). *Principles of Behavior: an introduction to behavior theory*. Appleton-Century-Crofts, New York. pages 4
- Hwang, E., Willis, B. S., and Burwell, R. D. (2017). Prefrontal connections of the perirhinal and postrhinal cortices in the rat. *Behavioural Brain Research*, 354:8–21. Author manuscript available in PMC (2018-11-15); final edited form published 2017-07-29. pages 32
- Høydal, O. A., Skytøen, E. R., Andersson, S. O., Moser, M. B., and Moser, E. I. (2019). Object–vector coding in the medial entorhinal cortex. *Nature*, 568. pages 14
- Inagaki, H. K., Fontolan, L., Romani, S., and Svoboda, K. (2019). Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature*, 566(7743):212–217. pages 201

- Issa, J. B., Radvansky, B. A., Xuan, F., and Dombeck, D. A. (2024). Lateral entorhinal cortex subpopulations represent experiential epochs surrounding reward. *Nature Neuroscience*, 27(3):536–546. pages 32, 33, 150, 158
- Jacobsen, C. F. (1935). Functions of frontal association area in primates. *Archives of Neurology & Psychiatry*, 33(3):558–569. pages 19
- James, W. (1890). *The Principles of Psychology*. Dover Publications, London, England. pages 1
- Jensen, K. T., Doohan, P., Sablé-Meyer, M., Reinert, S., Baram, A., Akam, T., and Behrens, T. E. J. (2025). A mechanistic theory of planning in prefrontal cortex. *bioRxiv*. preprint. pages 170, 179, 204
- Jensen, K. T., Hennequin, G., and Mattar, M. G. (2024). A recurrent network model of planning explains hippocampal replay and human behavior. *Nature Neuroscience*, 27(9):1340–1348. pages 180
- Johnson, A. and Redish, A. D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J. Neurosci.*, 27. pages 172
- Jun, H., Lee, J. Y., Bleza, N. R., Ichii, A., Donohue, J. D., and Igarashi, K. M. (2024). Prefrontal and lateral entorhinal neurons co-dependently learn item-outcome rules. *Nature*, 633(8031):864–871. Place: England. pages 33, 152
- Jun, J. J., Steinmetz, N. A., Siegle, J. H., Denman, D. J., Bauza, M., Barbarits, B., Lee, A. K., Anastassiou, C. A., Andrei, A., Aydın, Ç., Barbic, M., Blanche, T. J., Bonin, V., Couto, J., Dutta, B., Gratiy, S. L., Gutnisky, D. A., Häusser, M., Karsh, B., Ledochowitsch, P., Mora Lopez, C., Mitelut, C., Musa, S., Okun, M., Pachitariu, M., Putzeys, J., Rich, P. D., Rossant, C., Sun, W.-l., Svoboda, K., Carandini, M., Harris, K. D., Koch, C., O’Keefe, J., and Harris, T. D. (2017). Fully integrated silicon probes for high-density recording of neural activity. *Nature*, 551(7679):232–236. pages 151
- Kaefer, K., Nardin, M., Blahna, K., and Csicsvari, J. (2020). Replay of Behavioral Sequences in the Medial Prefrontal Cortex during Rule Switching. *Neuron*, 106(1):154–165.e6. pages 16, 24, 116, 136
- Kanter, B. R., Lykken, C. M., Polti, I., Moser, M.-B., and Moser, E. I. (2025). Event structure sculpts neural population dynamics in the lateral entorhinal cortex. *Science*, 388(6754):eadr0927. _eprint: <https://www.science.org/doi/pdf/10.1126/science.adr0927>. pages 32, 33, 150
- Keene, C. S., Bladon, J., McKenzie, S., Liu, C. D., O’Keefe, J., and Eichenbaum, H. (2016). Complementary Functional Organization of Neuronal Activity Patterns in the Perirhinal, Lateral Entorhinal, and Medial Entorhinal Cortices. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 36(13):3660–3675. Place: United States. pages 32
- Khona, M. and Fiete, I. R. (2022). Attractor and integrator networks in the brain. *Nature Reviews Neuroscience*, 23(12):744–766. pages 27, 200, 205, 206
- Kim, S. S., Rouault, H., Druckmann, S., and Jayaraman, V. (2017). Ring attractor dynamics in the *Drosophila* central brain. *Science*, 356(6340):849–853. Publisher: American Association for the Advancement of Science. pages 204
- Kinnavane, L., Barker, G. R., Banks, P. J., Bashir, Z. I., and Warburton, E. C. (2025). Medial prefrontal cortex input to lateral entorhinal cortex supports both encoding and retrieval of associative recognition memory. *bioRxiv*, page 2025.08.14.670409. pages 32

- Knierim, J. J. and Zhang, K. (2012). Attractor Dynamics of Spatially Correlated Neural Activity in the Limbic System. *Annual Review of Neuroscience*, 35(Volume 35, 2012):267–285. Publisher: Annual Reviews Type: Journal Article. pages 200
- Knudsen, E. B. and Wallis, J. D. (2020). Closed-loop theta stimulation in the orbitofrontal cortex prevents reward-based learning. *Neuron*, 106. pages 20
- Knudsen, E. B. and Wallis, J. D. (2021). Hippocampal neurons construct a map of an abstract value space. *Cell*, 184(18):4640–4650.e10. pages 17
- Kording, K. (2025). Attractors are usually not mechanisms. *Konrad's Substack*. Accessed 2025-09-22. pages 200
- Koscik, T. R. and Tranel, D. (2012). The human ventromedial prefrontal cortex is critical for transitive inference. *Journal of Cognitive Neuroscience*, 24(5):1191–1204. pages 20
- Kriegeskorte, N., Mur, M., and Bandettini, P. (2008). Representational similarity analysis – connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2. pages 141
- Krüger, L., Hohberg, M., Lehmann, W., and Dresing, K. (2018). Assessing the risk for major injuries in equestrian sports. *BMJ Open Sport & Exercise Medicine*, 4(1):e000408. pages 1
- Kuchenhoff, S., Azab, H., and Sheth, S. A. (2025). Personal communication. personal communication. pages 174
- Küchenhoff, S., Baram, A., El-Gaby, M., Sablé-Meyer, M., Harris, A. L., Bakermans, J., Shah, S., Bartoli, E., Watrous, A., Anand, A., Donoghue, T., Maesta-Pereira, S., Topalovic, U., Sakon, J., Smith, E. H., Suthana, N., Fried, I., Jacobs, J., Cowan, R. L., Azab, H., Sheth, S. A., and Behrens, T. (2025). Algorithmic representations in the human brain that underlie schema generalisation. In *Proceedings of the 47th Annual Meeting of the Cognitive Science Society*, volume 47. Cognitive Science Society. pages 174
- Kuruvilla, M. V., Wilson, D. I. G., and Ainge, J. A. (2020). Lateral entorhinal cortex lesions impair both egocentric and allocentric object–place associations. *Brain and Neuroscience Advances*, 4. pages 13
- Laubach, M., Amarante, L. M., Swanson, K., and White, S. R. (2018). What, if anything, is rodent prefrontal cortex? *eNeuro*, 25(5):ENEURO.0315–18.2018. pages 26
- Leonard, C. M. (1969). The prefrontal cortex of the rat. i. cortical projection of the mediodorsal nucleus. ii. efferent connections. *Brain Research*, 12:321–343. pages 26
- Lever, C., Burton, S., Jeewajee, A., O'Keefe, J., and Burgess, N. (2009). Boundary vector cells in the subiculum of the hippocampal formation. *J. Neurosci.*, 29. pages 14
- Lewandowsky, S. and Murdock, Bennett B., J. (1989). Memory for serial order. *Psychological Review*, 96(1):25–57. pages 169
- Liang, Z., Wu, S., Wu, J., Wang, W., Qin, S., and Liu, C. (2024). Distance and grid-like codes support the navigation of abstract social space in the human brain. *eLife*, 12:RP89025. pages 17, 206
- Little, W. A. (1974). The existence of persistent states in the brain. *Mathematical Biosciences*, 19(1):101–120. pages 200

- Low, R. J., Gu, Y., and Tank, D. W. (2014). Cellular resolution optical access to brain regions in fissures: Imaging medial prefrontal cortex and grid cells in entorhinal cortex. *Proceedings of the National Academy of Sciences*, 111(52):18739–18744. pages 178
- MacDonald, C. J., Lepage, K. Q., Eden, U. T., and Eichenbaum, H. (2011). Hippocampal “time cells” bridge the gap in memory for discontinuous events. *Neuron*, 71(4):737–749. pages 176
- Maor, I., Atwell, J., Ascher, I., Zhao, Y., Takahashi, Y. K., Hart, E., Pereira, F., and Schoenbaum, G. (2025). Persistent representation of a prior schema in the orbitofrontal cortex facilitates learning of a conflicting schema. *bioRxiv*. Preprint. pages 22
- Mardinly, A. R., Oldenburg, I. A., Pégard, N. C., Sridharan, S., Lyall, E. H., Chesnov, K., Brohawn, S. G., Waller, L., and Adesnik, H. (2018). Precise multimodal optical control of neural ensemble activity. *Nature Neuroscience*, 21(6):881–893. pages 178
- Mathis, A., Mamidanna, P., Cury, K. M., Abe, T., Murthy, V. N., Mathis, M. W., and Bethge, M. (2018). DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nature Neuroscience*, 21(9):1281–1289. pages 45
- McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., and Moser, M. . B. (2006). Path integration and the neural basis of the ‘cognitive map’. *Nat. Rev. Neurosci.*, 7. pages 14
- Miller, E. K. and Cohen, J. D. (2001). An Integrative Theory of Prefrontal Cortex Function. *Annual Review of Neuroscience*, 24(Volume 24, 2001):167–202. Publisher: Annual Reviews Type: Journal Article. pages 17, 18, 20
- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63(2):81–97. pages 4
- Montchal, M. E., Reagh, Z. M., and Yassa, M. A. (2019). Precise temporal memories are supported by the lateral entorhinal cortex in humans. *Nature Neuroscience*, 22(2):284–288. pages 32
- Morrissey, M. D., Insel, N., and Takehara-Nishiuchi, K. (2017). Generalizable knowledge outweighs incidental details in prefrontal ensemble code over time. *eLife*, 6:e22177. pages 23
- Nair, A., Karigo, T., Yang, B., Ganguli, S., Schnitzer, M. J., Linderman, S. W., Anderson, D. J., and Kennedy, A. (2023). An approximate line attractor in the hypothalamus encodes an aggressive state. *Cell*, 186(1):178–193.e15. pages 203
- Nardin, M., Kaefer, K., and Csicsvari, J. (2021). The generalized spatial representation in the prefrontal cortex is inherited from the hippocampus. *bioRxiv*, page 2021.09.30.462269. pages 16, 136
- Nardin, M., Kaefer, K., Stella, F., and Csicsvari, J. (2023). Theta oscillations as a substrate for medial prefrontal-hippocampal assembly interactions. *Cell Reports*, 42(9). Publisher: Elsevier. pages 16, 136
- Neisser, U. (1967). *Cognitive Psychology*. Appleton-Century-Crofts, New York. pages 4
- Nieh, E. H., Schottdorf, M., Freeman, N. W., Low, R. J., Lewallen, S., Koay, S. A., Pinto, L., Gauthier, J. L., Brody, C. D., and Tank, D. W. (2021). Geometry of abstract learned knowledge in the hippocampus. *Nature*, 595(7865):80–84. pages 17

- Nilssen, E. S., Doan, T. P., Nigro, M. J., Ohara, S., and Witter, M. P. (2019). Neurons and networks in the entorhinal cortex: A reappraisal of the lateral and medial entorhinal subdivisions mediating parallel cortical pathways. *Hippocampus*. Epub 2019-08-13. pages 32
- Ohara, S., Blankvoort, S., Nair, R. R., Nigro, M. J., Nilssen, E. S., Kentros, C., and Witter, M. P. (2021). Local projections of layer Vb-to-Va are more prominent in lateral than in medial entorhinal cortex. *eLife*, 10:e67262. Publisher: eLife Sciences Publications, Ltd. pages 152
- O'Keefe, J. and Dostrovsky, J. (1971). The hippocampus as a spatial map. preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, 34(1):171–175. pages 7, 13
- O'Keefe, J. and Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. Oxford University Press, Oxford. pages 7, 13
- Omer, D. B., Maimon, S. R., Las, L., and Ulanovsky, N. (2018). Social place-cells in the bat hippocampus. *Science*, 359(6372):218–224. pages 14
- Owen, A. M., Roberts, A. C., Polkey, C. E., Sahakian, B. J., and Robbins, T. W. (1991). Extra-dimensional versus intra-dimensional set shifting performance following frontal lobe excisions, temporal lobe excisions or amygdalo-hippocampectomy in man. *Neuropsychologia*, 29. pages 17, 172
- Pachitariu, M., Sridhar, S., Pennington, J., and Stringer, C. (2024). Spike sorting with Kilosort4. *Nature Methods*, 21(5):914–921. pages 43
- Pachitariu, M., Steinmetz, N., Kadir, S., Carandini, M., and D, H. K. (2016). Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. Technical report, UCL. pages 43
- Packer, A. M., Russell, L. E., Dagleish, H. W. P., and Häusser, M. (2014). Simultaneous all-optical manipulation and recording of neural circuit activity with cellular resolution in vivo. *Nature Methods*, 12:140 – 146. pages 178
- Park, S. A., Miller, D. S., and Boorman, E. D. (2021). Inferences on a multidimensional social hierarchy use a grid-like code. *Nature Neuroscience*, 24(9):1292–1301. pages 17, 206
- Passingham, R. E. (2002). The frontal cortex: does size matter? *Nature Neuroscience*, 5(3):190–192. pages 26
- Penfield, W. and Evans, J. P. (1935). The frontal lobe in man: a clinical study of maximum removals. *Brain*, 58(1):115–133. pages 19
- Pereira, T. D., Tabris, N., Matsliah, A., Turner, D. M., Li, J., Ravindranath, S., Papadoyannis, E. S., Normand, E., Deutsch, D. S., Wang, Z. Y., McKenzie-Smith, G. C., Mitelut, C. C., Castro, M. D., D'Uva, J., Kislin, M., Sanes, D. H., Kocher, S. D., Wang, S. S.-H., Falkner, A. L., Shaevitz, J. W., and Murthy, M. (2022). SLEAP: A deep learning system for multi-animal pose tracking. *Nature Methods*, 19(4):486–495. pages 45
- Peyrache, A., Khamassi, M., Benchenane, K., Wiener, S. I., and Battaglia, F. P. (2009). Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nat. Neurosci.*, 12. pages 24, 116
- Peyrache, A., Lacroix, M. M., Petersen, P. C., and Buzsáki, G. (2015). Internally organized mechanisms of the head direction sense. *Nature Neuroscience*, 18(4):569–575. pages 116, 204

- Pfeiffer, B. E. and Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, 497. pages 172
- Piaget, J. (1952). *The origins of intelligence in children*. W. W. Norton & Company, New York, NY. pages 2, 10, 11
- Pilkiw, M., Jarovi, J., and Takehara-Nishiuchi, K. (2022). Lateral Entorhinal Cortex Suppresses Drift in Cortical Memory Representations. *The Journal of Neuroscience*, 42(6):1104. pages 32, 175
- Pratt, W. E. and Mizumori, S. J. Y. (2001). Neurons in rat medial prefrontal cortex show anticipatory rate changes to predictable differential rewards in a spatial memory task. *Behav. Brain Res.*, 123. pages 23
- Preuss, T. M. (1995). Do rats have prefrontal cortex? the rose–woolsey–akert program reconsidered. *Journal of Cognitive Neuroscience*, 7(1):1–24. pages 26
- Reinert, S., Hübener, M., Bonhoeffer, T., and Goltstein, P. M. (2021). Mouse prefrontal cortex represents learned rules for categorization. *Nature*, 593. pages 21
- Rennó-Costa, C., Lisman, J. E., and Verschure, P. F. M. J. (2014). A Signature of Attractor Dynamics in the CA3 Region of the Hippocampus. *PLOS Computational Biology*, 10(5):1–15. Publisher: Public Library of Science. pages 201
- Restle, F. (1957). Discrimination of cues in mazes: a resolution of the place-vs.-response question. *Psychological Review*, 64:217–228. pages 5
- Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., and Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*, 497(7451):585–590. pages 20
- Rolls, E. T. and Treves, A. (1994). Neural networks in the brain involved in memory and recall. In Van Pelt, J., Corner, M., Uylings, H., and Lopes Da Silva, F., editors, *Progress in Brain Research*, volume 102, pages 335–341. Elsevier. pages 201
- Rose, J. E. and Woolsey, C. N. (1948). The orbitofrontal cortex and its connections with the mediodorsal nucleus in rabbit, sheep and cat. *Research Publications of the Association for Research in Nervous and Mental Disease*, 27:210–232. pages 26
- Rubin, A., Sheintuch, L., Brande-Eilat, N., Pinchasof, O., Rechavi, Y., Geva, N., and Ziv, Y. (2019). Revealing neural correlates of behavior without behavioral measurements. *Nature Communications*, 10(1):4745. pages 23
- Russell, L. E., Dalglish, H. W. P., Nutbrown, R., Gauld, O. M., Herrmann, D., Fışek, M., Packer, A. M., Häusser, M., et al. (2022). All-optical interrogation of neural circuits in behaving mice. *Nature Protocols*, 17:1579–1620. pages 178
- Samborska, V., Butler, J. L., Walton, M. E., Behrens, T. E. J., and Akam, T. (2022). Complementary task representations in hippocampus and prefrontal cortex for generalizing the structure of problems. *Nature Neuroscience*, 25(10):1314–1326. pages 24, 28, 78
- Sammons, R. P., Parthier, D., Stumpf, A., and Schmitz, D. (2019). Electrophysiological and Molecular Characterization of the Parasubiculum. *Journal of Neuroscience*, 39(45):8860–8876. Publisher: Society for Neuroscience _eprint: <https://www.jneurosci.org/content/39/45/8860.full.pdf>. pages 151

- Samsonovich, A. and McNaughton, B. L. (1997). Path integration and cognitive mapping in a continuous attractor neural network model. *The Journal of Neuroscience*, 17(15):5900–5920. pages 205
- Sarel, A., Finkelstein, A., Las, L., and Ulanovsky, N. (2017). Vectorial representation of spatial goals in the hippocampus of bats. *Science*, 355(6321):176–180. pages 17
- Sauer, J., Folschweiller, S., and Bartos, M. (2022). Topographically organized representation of space and context in the medial prefrontal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 119(6):e2117300119. pages 15
- Seelig, J. D. and Jayaraman, V. (2015). Neural dynamics for landmark orientation and angular path integration. *Nature*, 521(7551):186–191. pages 204
- Seung, H. (1996). How the brain keeps the eyes still. *Proceedings of the National Academy of Sciences*, 93(23):13339–13344. pages 200, 203
- Shima, K., Isoda, M., Mushiake, H., and Tanji, J. (2007). Categorization of behavioural sequences in the prefrontal cortex. *Nature*, 445(7125):315–318. pages 21
- Skaggs, W., Knierim, J., Kudrimoti, H., and McNaughton, B. (1994). A Model of the Neural Basis of the Rat's Sense of Direction. In Tesauro, G., Touretzky, D., and Leen, T., editors, *Advances in Neural Information Processing Systems*, volume 7. MIT Press. pages 204
- Skinner, B. F. (1938). *The Behavior of Organisms: An Experimental Analysis*. Appleton-Century, New York. pages 4, 7
- Skinner, B. F. (1953). *Science and Human Behavior*. Macmillan, New York. pages 4
- Spence, K. W. (1936). The nature of discrimination learning in animals. *Psychological Review*, 43(5):427–449. pages 7
- Stachenfeld, K. L., Botvinick, M. M., and Gershman, S. J. (2017). The hippocampus as a predictive map. *Nat. Neurosci.*, 20. pages 14
- Steinmetz, N. A., Aydin, C., Lebedeva, A., Okun, M., Pachitariu, M., Bauza, M., Beau, M., Bhagat, J., Böhm, C., Broux, M., Chen, S., Colonell, J., Gardner, R. J., Karsh, B., Kloosterman, F., Kostadinov, D., Mora-Lopez, C., O'Callaghan, J., Park, J., Putzeys, J., Sauerbrei, B., van Daal, R. J. J., Vollan, A. Z., Wang, S., Welkenhuysen, M., Ye, Z., Dudman, J. T., Dutta, B., Hantman, A. W., Harris, K. D., Lee, A. K., Moser, E. I., O'Keefe, J., Renart, A., Svoboda, K., Häusser, M., Haesler, S., Carandini, M., and Harris, T. D. (2021). Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science*, 372(6539):eabf4588. Publisher: American Association for the Advancement of Science. pages 151
- Stensola, H. (2012). The entorhinal grid map is discretized. *Nature*, 492. pages 13, 92
- Stringer, C., Pachitariu, M., Steinmetz, N., Carandini, M., and Harris, K. D. (2019). High-dimensional geometry of population responses in visual cortex. *Nature*, 571(7765):361–365. pages 200
- Sun, C., Yang, W., Martin, J., and Tonegawa, S. (2020). Hippocampal neurons represent events as transferable units of experience. *Nature Neuroscience*, 23(5):651–663. pages 64

- Sun, W., Advani, M., Spruston, N., Saxe, A., and Fitzgerald, J. E. (2023). Organizing memories for generalization in complementary learning systems. *Nature Neuroscience*, 26:1438–1448. pages 23
- Swaminathan, S., Dedieu, A., Vasudeva Raju, R., Shanahan, M., Lázaro-Gredilla, M., and George, D. (2023). Schema-learning and rebinding as mechanisms of in-context learning and emergence. *arXiv*. pages 22
- Taube, J., Muller, R., and Ranck, J. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations. *The Journal of Neuroscience*, 10(2):436. pages 14, 204
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., and Goodman, N. D. (2011). How to grow a mind: statistics, structure, and abstraction. *Science*, 331. pages 24
- Thompson, V. E. and Bucy, P. C. (1969). Learning after cerebral hemidecortication in monkeys. *Physiology & Behavior*, 4(4):455–459. pages 8
- Thorndike, E. L. (1911). *Animal Intelligence: Experimental Studies*. Macmillan, New York. pages 4, 6
- Tian, Z., Chen, J., Zhang, C., Min, B., Xu, B., and Wang, L. (2024). Mental programming of spatial sequences in working memory in the macaque frontal cortex. *Science*, 385(6716):eadp6091. Publisher: American Association for the Advancement of Science. pages 25, 29
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychol. Rev.* pages 6, 7, 12, 16, 39
- Tolman, E. C. and Honzik, C. H. (1930). Introduction and removal of reward, and maze performance in rats. *University of California Publications in Psychology*, 4:257–275. pages 4, 5
- Tolman, E. C., Ritchie, B. F., and Kalish, D. (1946). Studies in spatial learning. ii. place learning versus response learning. *Journal of Experimental Psychology*, 36(3):221. pages 5
- Tomov, M. S., Yagati, S., Kumar, A., Yang, W., and Gershman, S. J. (2020). Discovery of hierarchical representations for efficient planning. *PLoS Comput. Biol.*, 16. pages 24
- Trettel, S. G., Trimper, J. B., Hwaun, E., Fiete, I. R., and Colgin, L. L. (2019). Grid cell co-activity patterns during sleep reflect spatial overlap of grid fields during active behaviors. *Nat. Neurosci.*, 22. pages 116
- Tsao, A., Moser, M.-B., and Moser, E. I. (2013). Traces of Experience in the Lateral Entorhinal Cortex. *Current Biology*, 23(5):399–405. Publisher: Elsevier. pages 32
- Tsao, A., Sugar, J., Lu, L., Wang, C., Knierim, J. J., Moser, M.-B., and Moser, E. I. (2018). Integrating time from experience in the lateral entorhinal cortex. *Nature*, 561(7721):57–62. pages 32, 33, 150
- Tsao, A., Yousefzadeh, S. A., Meck, W. H., Moser, M.-B., and Moser, E. I. (2022). The neural bases for timing of durations. *Nature Reviews Neuroscience*, 23(11):646–665. pages 32
- Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. A., Wood, E. R., Witter, M. P., and Morris, R. G. M. (2007). Schemas and memory consolidation. *Science (New York, N.Y.)*, 316(5821):76–82. pages 22
- Turner-Evans, D., Wegener, S., Rouault, H., Franconville, R., Wolff, T., Seelig, J. D., Druckmann, S., and Jayaraman, V. (2017). Angular velocity integration in a fly heading circuit. *eLife*, 6:e23496. pages 204

- Tye, K. M., Miller, E. K., Taschbach, F. H., Benna, M. K., Rigotti, M., and Fusi, S. (2024). Mixed selectivity: Cellular computations for complexity. *Neuron*, 112(14):2289–2303. pages 20
- van Beest, E. H., Bimbard, C., Fabre, J.-M. J., et al. (2025). Tracking neurons across days with high-density probes. *Nature Methods*, 22:778–787. pages 179
- van Groen, T., Miettinen, P., and Kadish, I. (2003). The entorhinal cortex of the mouse: organization of the projection to the hippocampal formation. *Hippocampus*, 13(1):133–149. Place: United States. pages 31
- van Strien, N. M., Cappaert, N. L. M., and Witter, M. P. (2009). The anatomy of memory: an interactive overview of the parahippocampal–hippocampal network. *Nature Reviews Neuroscience*, 10(4):272–282. pages 31
- Vinograd, A., Nair, A., Kim, J. H., Linderman, S. W., and Anderson, D. J. (2024). Causal evidence of a line attractor encoding an affective state. *Nature*, 634(8035):910–918. pages 203
- Wallis, J. D., Anderson, K. C., and Miller, E. K. (2001). Single neurons in prefrontal cortex encode abstract rules. *Nature*, 411. pages 21
- Walton, M. E., Behrens, T. E., Buckley, M. J., Rudebeck, P. H., and Rushworth, M. F. (2010). Separable learning systems in the macaque brain and the role of orbitofrontal cortex in contingent learning. *Neuron*, 65. pages 17, 19
- Wang, C., Lee, H., Rao, G., and Knierim, J. J. (2024). Multiplexing of temporal and spatial information in the lateral entorhinal cortex. *Nature Communications*, 15(1):10533. pages 32
- Wang, J. X., Kurth-Nelson, Z., Kumaran, D., Tirumala, D., Soyer, H., Leibo, J. Z., Hassabis, D., and Botvinick, M. (2018). Prefrontal cortex as a meta-reinforcement learning system. *Nature Neuroscience*, 21(6):860–868. pages 9, 29, 34
- Wang, J. X., Kurth-Nelson, Z., Tirumala, D., Soyer, H., Leibo, J. Z., Munos, R., Blundell, C., Kumaran, D., and Botvinick, M. M. (2016). Learning to reinforcement learn. In *arXiv preprint arXiv:1611.05763*. pages 9
- Wang, X.-J. (2002). Probabilistic Decision Making by Slow Reverberation in Cortical Circuits. *Neuron*, 36(5):955–968. Publisher: Elsevier. pages 200, 203
- Watson, J. (1913). Psychology as the behaviorist views it. *Psychological Review*, 20:158–177. pages 4
- Whittington, J. C., Dorrell, W., Behrens, T. E., Ganguli, S., and El-Gaby, M. (2025). A tale of two algorithms: Structured slots explain prefrontal sequence memory and are unified with hippocampal cognitive maps. *Neuron*, 113(2):321–333.e6. Publisher: Elsevier. pages 27, 149, 174
- Whittington, J. C. R., McCaffary, D., Bakermans, J. J. W., and Behrens, T. E. J. (2022). How to build a cognitive map. *Nat. Neurosci.*, 25. pages 12, 17, 22, 34, 78
- Whittington, J. C. R., Muller, T. H., Mark, S., Chen, G., Barry, C., Burgess, N., and Behrens, T. E. J. (2020). The Tolman-Eichenbaum Machine: Unifying Space and Relational Memory through Generalization in the Hippocampal Formation. *Cell*, 183(5):1249–1263.e23. pages 14, 16, 17, 27, 30, 33, 34, 78, 97, 149, 150, 176, 202

- Wills, T. J., Lever, C., Cacucci, F., Burgess, N., and O'Keefe, J. (2005). Attractor dynamics in the hippocampal representation of the local environment. *Science*, 308(5723):873–876. pages 201
- Wilson, H. R. and Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, 12(1):1–24. pages 202
- Wilson, R. C., Takahashi, Y. K., Schoenbaum, G., and Niv, Y. (2014). Orbitofrontal cortex as a cognitive map of task space. *Neuron*, 81. pages 22
- Witter, M. P., Doan, T. P., Jacobsen, B., Nilssen, E. S., and Ohara, S. (2017). Architecture of the Entorhinal Cortex A Review of Entorhinal Anatomy in Rodents with Some Comparative Notes. *Frontiers in systems neuroscience*, 11:46. Place: Switzerland. pages 31, 32
- Wood, E. R., Dudchenko, P. A., Robitsek, R. J., and Eichenbaum, H. (2000). Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron*, 27. pages 63, 172
- Wood, R. A., Bauza, M., Krupic, J., Burton, S., Delekate, A., Chan, D., and O'Keefe, J. (2018). The honeycomb maze provides a novel test to study hippocampal-dependent spatial navigation. *Nature*, 554(7690):102–105. pages 13
- Xie, X., Tabuchi, M., Brown, M. P., Mitchell, S. P., Wu, M. N., and Kolodkin, A. L. (2017). The laminar organization of the Drosophila ellipsoid body is semaphorin-dependent and prevents the formation of ectopic synaptic connections. *eLife*, 6:e25328. Publisher: eLife Sciences Publications, Ltd. pages 204
- Xie, Y., Hu, P., Li, J., Chen, J., Song, W., Wang, X.-J., Yang, T., Dehaene, S., Tang, S., Min, B., and Wang, L. (2022). Geometry of sequence working memory in macaque prefrontal cortex. *Science*, 375(6581):632–639. Publisher: American Association for the Advancement of Science. pages 25, 29, 34, 135
- Yu, J. Y., Liu, D. F., Loback, A., Grossrubatscher, I., and Frank, L. M. (2018). Specific hippocampal representations are linked to generalized cortical representations in memory. *Nature Communications*, 9(1):2209. pages 16, 136
- Yu, X. T., Yu, J., Choi, A., and Takehara-Nishiuchi, K. (2021). Lateral entorhinal cortex supports the development of prefrontal network activity that bridges temporally discontinuous stimuli. *Hippocampus*, 31(12):1285–1299. pages 32
- Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *The Journal of Neuroscience*, 16(6):2112. pages 201, 204
- Zhang, Z., Russell, L. E., Packer, A. M., Gauld, O. M., and Häusser, M. (2018). Closed-loop all-optical interrogation of neural circuits in vivo. *Nature Methods*, 15(12):1037–1040. pages 178
- Zhou, J., Jia, C., Montesinos-Cartagena, M., Gardner, M. P. H., Zong, W., and Schoenbaum, G. (2021). Evolving schema representations in orbitofrontal ensembles during learning. *Nature*, 590(7847):606–611. pages 22, 28, 77
- Zhou, J., Montesinos-Cartagena, M., Wikenheiser, A. M., Gardner, M. P. H., Niv, Y., and Schoenbaum, G. (2019). Complementary task structure representations in hippocampus and orbitofrontal cortex during an odor sequence task. *Current Biology*, 29(20):3402–3409. pages 23

- Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., Xiong, H., and He, Q. (2020). A comprehensive survey on transfer learning. *arXiv preprint*. v3, last revised 2020-06-23. pages 70
- Zong, W., Zhou, J., Gardner, M. P. H., Zhang, Z., Costa, K. M., and Schoenbaum, G. (2025). Hippocampal output suppresses orbitofrontal cortex schema cell formation. *Nature Neuroscience*, 28:1048–1060. pages 23

11

Appendix: Mapping variables with attractors

A central aim of this thesis is to identify circuit mechanisms that map the structure of behaviour into neural activity. *Attractor* models offer a useful language for discussing such structure in recurrent networks. To build an intuition for the sorts of circuits that might support the representation of a structured behaviour, we will now briefly review attractor networks in the brain and how they may be used to represent schemas.

11.0.1 What are attractors?

The brain can be viewed as a dynamical system. That is, a set of variables with some rules that govern the evolution of those variables through time. For a network of n neurons, the instantaneous state is a point in $\mathbb{R}^n$ (one axis per unit of activity. Firing rate, calcium-dependent increase in fluorescence, and so on). The state of the network evolves as a function of both its intrinsic dynamical properties (strengths of recurrent synaptic connections, dendritic integration rules, plasticity mechanisms), as well as changes to its inputs. Although, in principle, activity could occupy any position in this high-dimensional state space, many neural circuits appear to constrain their dynamics to a radically smaller set of configurations that neighbouring states will evolve to, and that are stable or only weakly drifting. These constrained sets of fixed points and low-dimensional manifolds, together with their basins of attraction, are the working objects of attractor models. The term carries a precise mathematical meaning, and one can

argue about how strictly biological circuits satisfy the formal criteria (Knierim and Zhang, 2012; Stringer et al., 2019). Even if attractor dynamics are observed in a neural system, the leap from those dynamics to a full *mechanistic* model is hazardous, owing to multiple realisability – see Kording (2025) for a persuasive critique. Indeed, in the published version of this work (El-Gaby et al., 2024), we deliberately avoided using the term attractor at all to avoid taking a side in this debate in a way that would detract from the exposition of our key findings. Here I will not be quite so cautious. To be clear, I am not fully committing to precisely wired attractors as a true mechanistic explanation for either the previous findings discussed in this section, or the results of the work in this thesis. However, I believe the framework is valuable as an intuition pump for some general principles of having an approximately-stable structure to neural representations that mirrors the structure of a real-world variable. Here, we will first consider the appealing properties of attractor networks in general before covering various classes of attractor, examples of them in the brain (or things that look like them), and why they are so suitable for the problems they solve. A much more comprehensive treatment of attractor networks in the brain can be found in Khona and Fiete (2022).

11.0.2 What’s so attractive about attractors?

The broad appeal of attractors is this: attractor dynamics offer a mechanism for achieving robust, perturbation-resistant, structured neural codes. By steering population activity toward a small set of stable states, or along a stable low-dimensional manifold, they automatically correct small perturbations and so preserve representational fidelity in noisy circuits (Hopfield, 1982; Little, 1974; Amari, 1977b). The same recurrent structure allows activity to persist after inputs cease, supporting working memory and temporal integration through self-sustaining patterns of excitation and inhibition (Compte et al., 2000; Wang, 2002). Because the salient dynamics unfold on low-dimensional geometries, downstream readouts can remain simple. For example, the centre of an "activity bump" provides a nearly linear code for a scalar variable (Seung, 1996). Just as importantly, the topology of the manifold can be matched to the topology of the variable to be represented: a discrete set for categorical variables, a line for ordered quantities, a ring for periodic variables, or a torus for two-dimensional periodic lattices. This

alignment makes operations like wrapping-around emerge naturally from local dynamics and facilitates generalisation to new contexts (Ben-Yishai et al., 1995; Zhang, 1996; Gardner et al., 2022). Moreover, an internally structured neural circuit – for example, a ring – in principle allows for the same set of neurons to represent different variables depending on the task demands, merely by changing the inputs and outputs to the circuit. The internal structure of the network itself does not have to change.

11.0.3 Discrete attractors

When the variable that an attractor network represents is categorical, the natural geometry is a set of isolated fixed points. An example of this is auto-associative models like Hopfield networks. In these models, recurrent weights are learned so that previously experienced patterns become attractor basins – points in the high-dimensional state space that neighbouring states will evolve to. This means that partial or corrupted versions converge to the stored pattern, implementing pattern completion and denoising (Hopfield, 1982). This logic has long been linked to hippocampal CA3, whose dense recurrent collaterals endow the CA3 pyramidal network with auto-associative dynamics well suited for discrete attractors (Rolls and Treves, 1994; Rennó-Costa et al., 2014).

A striking example of discrete attractor dynamics occurs in hippocampal place cells. In an experiment in which a square arena is gradually morphed into a circular arena, rats' place cell representation does not change over a continuous gradient mirroring this slow change (Wills et al., 2005). Instead, after a threshold, experimenters observe a global remapping event consistent with the network settling into a new attractor basin when the inputs push the network activity sufficiently far from its fixed point. For an example in frontal circuits, discrete short-term memory attractor dynamics maintain choice-relevant network activity for decision making (Inagaki et al., 2019). Here, once a choice state is entered, recurrent feedback can stabilise network activity against perturbations until an outcome or reset signal arrives.

In the ABCD task, there are (at least) two possible use cases for discrete attractors. First, as mentioned above, discrete attractors can permit the persistence of choice-related activity in frontal circuits. In our frontal cortex recordings from mice performing the ABCD task, it is possible that we will observe discrete attractor-like dynamics for our individual latent task states (sequence positions). That is, the network will have four fixed points, one for A, B, C, and D. These states could then communicate with downstream regions representing specific goal locations to compute the appropriate actions to reach the next goal. The second use case for discrete attractors is hippocampal memory retrieval. In a computational model of the hippocampal-entorhinal network – the Tolman Eichenbaum Machine (Whittington et al., 2020) – a ‘hippocampal’ Hopfield network takes input from an ‘entorhinal’ recurrent neural network that tracks latent position on a graph. Attractor dynamics in the Hopfield network are able to take this position estimate as input and resolve a joint position $\times$ object association to predict what the sensory observation at that position will be. If frontal cortex holds a representation of ABCD sequence position, the hippocampus may use discrete attractor dynamics to convert this abstraction into a representation of a concrete goal.

11.0.4 Continuous attractors

Many variables that the brain needs to represent are continuous, like eye position, head direction, accumulated evidence, position in space. For these variables, continuous attractor networks (CANs) can implement a continuum of stable states that can match the structure of the variable in question. This is typically achieved via translationally symmetric connectivity via local excitation with broader inhibition (Wilson and Cowan, 1972; Amari, 1977a; Hansel and Sompolinsky, 1998; Ben-Yishai et al., 1995). This configuration sustains a stable activity bump that is passed on to neighbouring neurons.

In practice, networks of neurons are often only approximately continuous attractors. Activity bumps drift slowly under noise and are re-anchored by sensory input (Hardcastle et al., 2015). The key property is that the dominant dynamics live on a low-dimensional manifold that preserves neighbourhood relations. We will now cover some classes of continuous attractor.

Line attractors

A line attractor provides a one-dimensional continuum of states and so can store a scalar value or integrate a one-dimensional input. The oculomotor integrator is a key example here. Persistent eye-position signals in brainstem circuits behave like a leaky line attractor whose state is incremented by velocity inputs (Seung, 1996; Aksay et al., 2007). Related mechanisms have been proposed for working memory and evidence accumulation in frontal and parietal cortex, where recurrent excitation and balanced inhibition sustain graded activity or ramping trajectories that reflect the integral of inputs over seconds (Wang, 2002). More recently, technically impressive causal evidence using holographic optogenetics has uncovered line attractor dynamics for affective state in hypothalamic populations (Vinograd et al., 2024; Nair et al., 2023) – Figure 11.1.




Figure 11.1: Line attractor dynamics in a biological circuit

Adapted from Vinograd et al. (2024). Hypothalamic neurons form a line attractor representation for level of aggression. The trajectory of the neural population activity is the black line. The stable area of the neural activity landscape is in blue, the unstable area in red. When a mouse observes aggressive behaviour of an intruder, the population activity is pushed into the unstable area but settles in the blue line attractor basin.

In the ABCD task, the position in task-space (how far the mouse is through the sequence) could be represented through line attractor-like dynamics. Such a mechanism tracks progression through the sequence as a continuous variable with higher granularity than the 4-state discrete attractor option discussed above. It would effectively encode progress between each pair of goals over and above just sequence position (as discussed in the previous section on PFC task representations).

Ring attractors

Periodic one-dimensional variables (like orientation of a visual stimulus or heading direction) naturally map to a circle. For these cases, ring attractors are favoured. In ring attractors, activity forms a localised bump whose angular coordinate encodes the value of the variable.

A ring can be built from lateral-inhibition motifs with rotational symmetry (short-range excitation with broader inhibition) so that all phase-shifted bumps are equivalent fixed points (Ben-Yishai et al., 1995; Zhang, 1996) (Figure 11.2). Beyond providing a stable code, such circuits naturally integrate angular velocity: small, approximately antisymmetric components of recurrent or input drive shift the bump around the ring, implementing path integration of head turns or stimulus rotation (Skaggs et al., 1994; Turner-Evans et al., 2017). Because neurons with similar preferred angles mutually excite, and dissimilar ones inhibit, the network effectively performs inference over noisy cues (sensory input, angular velocity), settling the bump at the angle most compatible with the inputs (Ben-Yishai et al., 1995; Zhang, 1996; Jensen et al., 2025). Experiments across species support ring-like dynamics in head-direction systems (thalamus–presubiculum–retrosplenial networks) and in orientation tuning in visual cortex (Taube et al., 1990; Skaggs et al., 1994; Peyrache et al., 2015). Crucially, ring attractor-like dynamics persist offline during sleep, suggesting genuine internal organisation (Peyrache et al., 2015; Chaudhuri et al., 2019). Because the network connectivity is circular, wrap-around comes for free. If a mouse rotates 360° it ends up where it started; in a ring attractor, the activity bump ends up where it started. This is an appealing match to periodic variables.

Of course, ring attractors do not require the circuit's constituent neurons to be physically arranged in a circle – it is the connectivity profile that is circular. However, in an unusually neat and literal result for biology, neurons in the drosophila head direction circuit are physically arranged in a ring, with a detectable activity bump that follows this physical geometry (Seelig and Jayaraman, 2015; Kim et al., 2017; Xie et al., 2017).




Figure 11.2: A ring attractor network with left and right velocity inputs

(a) Adapted from Khona and Fiete (2022). A ring attractor circuit where neurons (outer circles) have excitatory connections to proximal neighbours but broad inhibition to more distal neurons. An activity bump can be pushed in different directions around the ring by external velocity inputs. (b) Adapted from Chaudhuri et al. (2019). Low-dimensional ring-structured state space of the head-direction system in wake and sleep recovered by dimensionality reduction.

The ABCD task has a periodic structure, with mice having to return to the first goal in the sequence immediately upon collection of the first reward, completing the loop. This lends itself to representation via ring-attractor-like dynamics. In the results of this thesis, we will occasionally use “ring-like” informally to describe circular structure in neural activity, regardless of whether the underlying circuit is a textbook attractor.

Toroidal attractors

Two-dimensional variables with periodic boundaries can be represented on a torus (Figure 11.3). Models of grid cell modules in medial entorhinal cortex (MEC) often invoke 2D continuous attractors whose bump drifts under velocity integration – enabling path integration (Samsonovich and McNaughton, 1997) – and tiles space with a periodic lattice (Burak and Fiete, 2009). Like with line- and ring-structured attractors, this is achieved using mutual excitation between neurons representing nearby locations, with broad inhibition for more distal locations. Recent population recordings demonstrate that grid cell modules occupy a toroidal manifold whose topology is preserved across environments and even in sleep, providing striking neural evidence for a torus *in vivo* (Gardner et al., 2022). As with rings, the appeal is that the

manifold's intrinsic geometry enforces the right boundary conditions and supports accurate path integration with simple local rules.




Figure 11.3: Toroidal topology of the grid cell network

(a, b) Adapted from [Khona and Fiete \(2022\)](#). (a) A 'neural sheet' with centre excitation and local inhibition results in multiple periodically-spaced activity bumps. (b) 2D phase shifts of the multi-peak pattern gives distinct but stable states. Once the translation has reached the lattice periodicity, the pattern repeats, forming a torus. (c) Adapted from [Gardner et al. \(2022\)](#). Top: non-linear low-dimensional embedding of a grid cell module's population activity reveals a toroidal topology. Bottom: firing rate map and 2D autocorrelogram of an example constituent grid cell.

The toroidal activity of grid cells (and the path integration it permits) has primarily been discussed in the spatial domain. However, grid-like codes in entorhinal cortex (and in humans and primates, frontal cortex) have been discovered across species for non-spatial variables ([Aronov et al., 2017](#); [Constantinescu et al., 2016](#); [Bao et al., 2019](#); [Chen et al., 2024a](#); [Liang et al., 2024](#); [Park et al., 2021](#); [Bongioanni et al., 2021](#)). As previously mentioned, an appealing property of internally-structured attractor-like networks is their ability to be reused to map different variables of the same structure by simply altering the inputs and outputs ([Khona and Fiete, 2022](#)). In our recordings from MEC, one possibility is that the 2D periodic toroidal network structure is reused to represent the 1D periodic task structure.

11.0.5 Link to the present work

The attractor perspective will serve as a conceptual scaffold for interpreting our neural findings. In mFC we identify modular, ring-like arrangements of neurons whose dynamics track task

progress relative to behavioural anchors (Chapter 4, Chapter 5, Chapter 6). These features are reminiscent of continuous-attractor intuitions (local bump-like activity, circular topology, offline persistence) without requiring us to assert a literal mathematical attractor.