

**Neurons, networks, and behaviour: Investigating
multiscale interactions between brain regions and
their behavioural correlates**



Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

Seán K. Martin, M.Sc., B.A.

A thesis submitted to the University of Dublin, Trinity College

for the degree of

Doctor of Philosophy

Supervised by

Prof. Shane O'MARA

2024

Declaration

I, Seán Martin, declare that this thesis has not been submitted as an exercise for a degree at this or any other university and it is entirely my own work.

I agree to deposit this thesis in the University's open access institutional repository or allow the Library to do so on my behalf, subject to Irish Copyright Legislation and Trinity College Library conditions of use and acknowledgement.

I consent to the examiner retaining a copy of the thesis beyond the examining period, should they so wish (EU GDPR May 2018).



Seán Martin

March 19th, 2024

Acknowledgements

Thank you to my supervisor, Shane, for providing valuable insight into my chosen topic and offering fantastic suggestions for areas of investigation. Your guidance, comfort, and trust have been instrumental to me on my research journey.

I would also like to thank all my colleagues, including Beth, Charlotte, Emanuela, Eoin, Ham, Jen, John, and Matheus, it has been an absolute pleasure to meet you and learn from you.

I am grateful to the members of the School of Psychology and the Trinity College Institute of Neuroscience for your help and support, particularly Barbara, Ciaran, Elaine, Luisa, and my thesis committee Fiona and Rhodri.

To my partner, Jennifer, I am deeply grateful for your patience, unwavering encouragement, and invaluable advice throughout this process. Your support has meant so much to me.

Finally, I deeply appreciate the support of my family and friends throughout these years, and those that came before – thank you all dearly.

SEÁN MARTIN

University of Dublin, Trinity College
March 19th, 2024

Neurons, networks, and behaviour: Investigating multiscale interactions between brain regions and their behavioural correlates

Seán K. Martin, M.Sc., B.A.
University of Dublin, Trinity College, March 19th, 2024
Supervisor: Shane O'Mara

Understanding the interactions between different brain regions is vital for unravelling the complexities of brain disorders and brain functions such as cognition, behaviour, and memory. While previous investigations of inter-regional communication were primarily performed via neuroimaging and electroencephalography, these are of limited spatial and temporal resolution. Recent technological advances and improved recording techniques have opened new avenues in cross-region intracranial recordings, allowing us to explore the interactions between brain regions at a finer scale than ever before.

This thesis addresses the analysis of inter-regional interactions using a variety of data types obtained from simultaneous recordings of brain activity. To facilitate this analysis, we introduce Simuran, a software tool that enables the analysis of different data formats and promotes interoperability between neural data and analysis tools.

Using Simuran, we investigate the impact of anterior thalamic lesion or inactivation in rats on the hippocampus, subiculum, and retrosplenial cortex. Our results reveal that anterior thalamic lesions disrupt spatial cell coding in the subiculum while leaving CA1 spatial cell coding intact. Interestingly, at the local field potential level, the effect of these lesions is not clear in the subiculum alone; rather, differences emerge when comparing local field potentials across brain regions and examining the relationship between local field potentials and spiking activity.

Next, we explore interactions at the neuron-level and consider the likelihood of recording structurally connected neurons across regions using simulations and statistics. Our findings indicate that modern recording techniques offer a high probability of capturing synaptically connected neurons across well-connected brain areas. Leveraging this connection probability, we turn to analysing large-scale datasets. The results suggest that task performance involves a balance between consistency in the represented information via highly correlated neural activity, while avoiding excessive redundancy. Additionally, patterns in neural ensembles exhibit similarities during both correct and incorrect task performances.

This research builds upon solid foundations to consider cross-region interactions in the brain, shedding light on observed behaviours and highlighting exciting opportunities to leverage modern technology and analytics.

Contents

List of figures	xiii
List of tables	xvii
Acronyms	xix
1 Introduction and thesis outline	1
2 Overview of multi-region neural data analysis	3
2.1 The primary intracranial brain signals recorded <i>in vivo</i>	4
2.1.1 Action potentials from single cells	5
2.1.2 Local field potentials	7
2.2 Recording the brain in the body	9
2.3 Fundamentals concepts for analysing intracranial recordings	10
2.3.1 Types of neural correlation	11
2.3.2 Encoding and decoding of neural signals	11
2.3.3 Approaches to analysing spike trains	12
2.3.4 Drawing insight from the local field potential	13
2.4 Recent techniques for multi-region analysis and future directions	15
2.4.1 Applying machine learning to large-scale neural data	16
2.4.2 Future avenues for cross-region neural ensemble analysis	17
2.5 Conclusion and outlook for this thesis	19

References	20
3 Multi-region neural data analysis and interoperability with Simuran	35
3.1 Challenges in large-scale neural data analysis	36
3.2 An overview of Simuran	37
3.2.1 Primary goals of the software	37
3.2.2 Organisation and design of Simuran	38
3.2.3 Providing a flowchart style user interface	40
3.3 Applications of Simuran to neural data	40
3.3.1 Analysing spikes and local field potentials from the rat limbic system	40
3.3.2 Tracking cells across days in calcium imaging data	43
3.3.3 Performing factor analysis on Allen Institute data	43
3.4 Comparison to other software and conclusion	46
3.5 Evolution of Simuran and user feedback	49
3.6 Instructions for using Simuran and methods	49
3.6.1 Batch data processing with Simuran	50
3.6.2 Software design patterns used in Simuran	50
3.6.3 Creating new functions for the visual editor	52
3.6.4 Processing the visual node editor via graphs	52
3.6.5 Improving performance with Simuran	54
3.6.6 Sharing, documenting, and testing Simuran	54
References	55
4 Anterior thalamic lesions abrogate subicular spatial coding	61
4.1 My contributions to this study	62
4.2 The role of the hippocampus and anterior thalamic nuclei in spatial memory	62

4.3	Abrogation of subicular spatial coding following ATN inactivation . . .	64
4.3.1	Differences in spatial cell recording frequencies between control and lesioned rats	65
4.3.2	Subicular cell firing properties in control and lesioned rats . . .	67
4.3.3	The effects of inactivation via muscimol mimic permanent ATN lesions	71
4.3.4	Place cells are recorded in CA1 despite ATN lesions	73
4.4	Discussing the possibility of multiple spatial circuits and the consequences for diencephalic amnesia	75
4.5	Experimental procedures and analysis methods	77
4.5.1	Summary of experiment details and unit processing	77
4.5.2	Performing downsampling analysis to minimise effects of exploration reduction in muscimol affected rats	78
4.5.3	Computing Bayesian statistics to assess the difference in frequency of spatial cells	78
4.5.4	Calculating bursting cells and spatial coding	79
	References	80
5	The effect of anterior thalamic lesions on local field potentials	85
5.1	My contributions to this study	86
5.2	Using local field potentials to expand upon the anterior thalamic nuclei lesion findings	86
5.3	Results of analyses on local field potentials	88
5.3.1	Rhythms in LFPs are maintained in the subiculum during exploration, but not in retrosplenial cortex	88
5.3.2	Heightened motor activity remains linked to increased theta activity	92
5.3.3	Sharp wave ripples occur during resting periods in both control and lesioned rats	93

5.3.4	Spike-LFP theta coupling is lower for non-spatial units, particularly from lesioned rats	94
5.3.5	Coherence between LFPs from the subiculum and retrosplenial cortex is higher in control rats in the theta and beta bands . . .	95
5.4	Contrasting the LFP results with the abrogation of subicular cells in ATN-lesioned rats	97
5.5	Analysis methods for local field potentials	101
5.5.1	Selecting local field potentials informed by histology	101
5.5.2	Signal filtering and processing of local field potentials	101
5.5.3	Power spectral density analysis	103
5.5.4	Speed theta analysis	104
5.5.5	Coherence analysis	105
5.5.6	Sharp-wave ripple detection during resting	106
5.5.7	Analysing the coupling of spikes to the local field potential . . .	106
5.5.8	Statistical testing	107
5.5.9	Data and code sharing	108
	References	109
6	Recording synaptically connected neurons across brain regions	115
6.1	Neural recording technologies and source transformation rules	116
6.2	Calculating the probability of recording connected neurons across brain regions	117
6.2.1	Model inputs and assumptions	117
6.2.2	Monte Carlo simulation	120
6.2.3	Statistical estimation	120
6.3	Case studies: estimating connection probabilities in the mouse and rat brain	121
6.3.1	Connection probability in the neocortex of the mouse brain . . .	121

6.3.2	Connection probability in the hippocampal formation of the rat brain	125
6.4	Discussion and limitations	128
6.4.1	Implications for experimental design of cross-region analysis . .	128
6.4.2	Limitations of the current model	129
6.4.3	Bias of the model	131
6.4.4	The relationship between functional connectivity and structural connectivity	132
6.4.5	Conclusion	133
6.5	Methods and examples	133
6.5.1	Description of directed graphs and networks of neurons	133
6.5.2	Calculating connection probability via Monte Carlo simulation .	135
6.5.3	Calculating connection probability via statistical methods . . .	137
6.5.4	Technical details of statistical estimation and performance . . .	142
6.5.5	Example of a full probability distribution calculation	145
6.5.6	Calculations in the hippocampal circuit case study	146
6.5.7	Blue Brain Project neocortex connection model and visualisation	148
6.5.8	Power law analysis	150
6.5.9	Code accessibility, sharing and reproducibility	151
	References	152
7	Relating task performance to correlation in multi-region neural data	159
7.1	My contributions to this study	160
7.2	Overview of neural correlation and the correlation-behaviour relationship	160
7.3	Relating neural activity in multiple regions to tasks and behaviour . . .	163
7.3.1	Coherence between the claustrum, RSC, and ACC during operant conditioning	163

7.3.2	Correlation in visual brain areas and limbic system brain areas	165
7.3.3	The shape of neural trajectories in latent space	165
7.4	Describing the balance between information representation and stability in neural activity	172
7.5	Data description and analysis methods	175
7.5.1	Operant conditioning with claustrum, RSC, and ACC recordings in rats	176
7.5.2	Allen Institute for Brain Science and International Brain Labora- tory datasets	178
7.5.3	Coherence analysis of local field potentials	179
7.5.4	Neural trajectory analysis with Gaussian process factor analysis	181
7.5.5	Procrustes analysis and distance measurement to compare neural trajectories	182
7.5.6	Correlation between neural activity with canonical correlation analysis	183
7.5.7	Statistical testing	185
7.5.8	Data processing and software availability	185
	References	185
8	Summary and future directions	191
8.1	Summary of our findings	191
8.2	Future directions and impact	193
8.3	Conclusion	194
8.4	Closing remarks	194
	References	196
A	Appendix	197
A.1	Further details of the anterior thalamic lesion study	197
A.1.1	Behavioural effect of ATN lesions	197

A.1.2	Lesion effectiveness	198
A.1.3	Subject details and experimental design	198
A.1.4	Detailed results from temporary inactivation via muscimol	201
A.1.5	ATN lesions with CA1 recordings	204
A.1.6	Cells in CA1 and possible recording from fibers of passage	206
A.1.7	Ethics	207
A.1.8	Surgical methods	208
A.1.9	Electrophysiological recordings	209
A.1.10	Unit identification and isolation	210
A.1.11	Spatial alternation testing	210
A.1.12	Spatial measures and rate maps	211
A.1.13	Hypothesis testing	212
A.1.14	Perfusion and histology	212
	References	214

List of figures

3.1	The primary components of Simuran’s software architecture	39
3.2	A flowchart in Simuran’s visual interface to load, process, and plot LFPs	41
3.3	Combining Simuran and Snakemake to create an analysis pipeline with branching paths via directed graphs	42
3.4	Using Simuran to track the same excitatory cell across days in calcium imaging data from Allen Brain Institute Visual Behaviour Optical Physiology experiment	44
3.5	A pipeline in Simuran to analyse both calcium imaging data from the Allen Institute for Brain Science and Neuropixels data from the International Brain Laboratory through the same flow	45
3.6	Python code demonstrating the use of Simuran to run factor analysis on the Allen Institute for Brain Science visual behaviour Neuropixels dataset	47
3.7	Using Python to create a new type of node in Simuran’s visual interface that associates single units to the brain region they were recorded from	53
4.1	Limbic system brain regions considered in the ATN lesion study and a brief description of the connections between them	63
4.2	Representative single units recorded in the dorsal subiculum from control and ATN lesioned rats	66
4.3	The posterior distributions of recording spatial and non-spatial cells from the subiculum in control and lesioned rats	68
4.4	Significant differences in bursting and non-spatial cell properties between control and ATN lesioned rats	69
4.5	Comparison of spike widths in control and lesioned rats	72

4.6	Results from spatially-downsampling the subicular spatial units considered in the temporary muscimol ATN lesion experiments	73
4.7	The effect of ATN lesions on CA1 cell recordings	74
5.1	Power spectral densities of subicular and retrosplenial cortex LFPs compared between control and lesion rats during open-field exploration . .	89
5.2	Power in oscillation bands of the subicular and retrosplenial cortex LFPs compared between control and lesion rats during open-field exploration	91
5.3	Comparing running speed and relative theta power during open-field exploration from LFPs recorded in the subiculum and ipsilateral retrosplenial cortex recordings	93
5.4	Sharp-wave ripples in the subiculum and CA1 during resting periods in control and ATN lesioned rats	94
5.5	The subicular LFP to spike relationship measured by spike field coherence and spike-triggered average	96
5.6	Coherence between subicular and retrosplenial cortex LFPs in control and ATN lesioned rats during T-maze spatial alternation and open-field exploration	98
5.7	Examples of retrosplenial cortex bipolar electrode implantation histology	102
6.1	Illustrating the problem of estimating the probability of conjointly recording connected neurons	118
6.2	Probability of recording connected neurons in simultaneous primary visual area and lateral visual area recordings with different probe insertion angles and locations in the Blue Brain model of neocortex	122
6.3	Modelling the probability of simultaneously recording neurons in mouse neocortex	123
6.3	Caption for figure 6.3 continued on this page	124
6.4	Modelling the probability of simultaneously recording directly connected neurons in rat CA3, CA1, and the subiculum	127
6.5	The relationship between the accuracy of probability estimation and maximum allowed synaptic jumps (geodesic distance)	143

6.6	Multiple distributions involved in calculating the probability of recording connected neurons	144
7.1	The length of trials and distributions of lever press times in fixed interval and fixed ratio tasks	164
7.2	The coherence between the claustrum, anterior cingulate cortex, and retrosplenial cortex in fixed interval and fixed ratio tasks	166
7.3	Demonstrating the canonical correlation analysis performed on the multi-region data during change detection trials	168
7.4	Computing the mean correlation between the dimension reduced activity in multiple-brain regions across catch and miss trials	169
7.5	One-second-long latent neural trajectories extracted from visual and limbic areas during a change detection task	171
7.6	The relationship of the similarity of neural trajectory shapes to the distance between the start and end points of trajectories	173
7.7	The log firing rate of the units selected from the Allen Institute’s visual Neuropixels dataset and the International Brain Laboratory’s brain wide survey	180
A.1	Behavioural impacts of ATN lesions on spatial alternation memory performance and novel object recognition	199
A.2	Demonstrating ATN lesion effectiveness from permanent (NMDA) lesions and temporary inactivation via muscimol	200
A.3	Effect of muscimol infusion on unit activity over time and T-maze performance	203
A.4	Spike properties following temporary inactivation of ATN with muscimol	205
A.5	Depicting the maximum and minimum extent of NMDA lesion in ATN lesioned rats with recordings in CA1	206

List of tables

4.1	Contingency table of recordings performed in the subiculum with spatial cells or non-spatial cells recorded	67
4.2	Summary of spike and burst properties of subicular units from control and lesioned rats	70
5.1	The relative bandpowers of local field potentials recorded from control and lesion rats during open-field exploration.	90
6.1	Rat hippocampal formation neuron count estimates, connectivity assumptions, and tetrode sampling assumptions with references	125
7.1	Median and interquartile range of coherence between the claustrum, anterior cingulate cortex, and the retrosplenial cortex along with results from Mann-Whitney U tests	167
7.2	The correlation between the dimension reduced activity in multiple-brain regions compared across catch and miss trials	170
7.3	Number of implantation targets in the claustrum, anterior cingulate cortex, and retrosplenial cortex	177
A.1	Types of recording layouts and the brain regions implanted in along with the number of rats with each type of recording layout in the ATN lesion study	202
A.2	Summary of place and interneuron properties of CA1 units from control and ATN lesioned rats	206

Acronyms

ACC anterior cingulate cortex.

AIS Allen Institute for Brain Science.

ANN artificial neural network.

ANOVA analysis of variance.

AP anterior-posterior.

ATN anterior thalamic nuclei.

ATNx ATN lesion-induced dysfunction.

BMI brain machine interface.

BOLD blood oxygen level dependent.

CA cornu ammonis.

CCA canonical correlation analysis.

CI calcium imaging.

CLA claustrum.

CLES common language effect size.

DAB diaminobenzidine.

DAG directed acyclic graph.

DFT discrete Fourier transform.

DLAG delayed latents across groups.

DREADDs designer receptors exclusively activated by designer drugs.

DV dorsal-ventral.

EMG electromyogram.

FAIR findable, accessible, interoperable, reusable.

FI fixed interval.

FIR finite impulse response.

fMRI functional magnetic resonance imaging.

FR fixed ratio.

GABA *gamma*-aminobutyric acid.

GPFA gaussian process factor analysis.

GUI general user interface.

IBI inter-burst interval.

IBL International Brain Laboratory.

ICA independent component analysis.

IQR inter-quartile range.

ISI inter-spike interval.

LED light emitting diode.

LFADS latent factor analysis via dynamical systems.

LFP local field potential.

LGN lateral geniculate nucleus.

LGv ventral part of the lateral geniculate nucleus.

ML machine learning.

ML medial-lateral.

MOp primary motor area.

NeuN neuronal nuclear protein.

NMDA N-Methyl-D-aspartic acid.

NWB Neurodata Without Borders.

ONE open neurophysiology environment.

PBS phosphate-buffered saline.

PBST phosphate-buffered saline with a low-concentration detergent for protein solubilisation.

PCA principal component analysis.

PFA paraformaldehyde.

PSD power spectral density.

PSTH peristimulus time histogram.

RNN recurrent neural network.

RSC retrosplenial cortex.

SEM standard error of the mean.

SFC spike field coherence.

SSp-II primary somatosensory area associated with lower limbic function.

STA spike-triggered average.

SWR sharp-wave ripple.

UI user interface.

VBN Visual Behaviour Neuropixels.

VISI lateral visual area.

VISp primary visual area.

1 Introduction and thesis outline

In the brain, complex actions require vast amounts of information to be passed along a multitude of short- and long-range connections. For instance, the act of reading provokes an orchestrated response of neurons from distinct cortical areas related to visual processing, memory, imagination, attention, and more. Much of our understanding of this interplay has stemmed from relating the blood oxygen level dependent (BOLD) signal from functional magnetic resonance imaging (fMRI) recordings to cognitive tasks, which provides a coarse-grained view of brain activity. Recent advances in technology have enabled simultaneous recordings of multiple neurons from distinct brain regions, offering new opportunities for studying fine-scale communication in the brain. Despite this, analysing such large-scale recordings poses significant theoretical and computational challenges. Our research tackles the challenge of using this data to improve our understanding of interactions between brain regions and their behavioural correlates.

We begin with an overview of the key topics covered in this thesis:

- To establish the context for our research, we provide a review of brain signals, current analysis methods for multi-region recordings, and our understanding of inter-areal interactions (chapter 2).
- Addressing the need for a flexible tool to facilitate multi-region recording analysis, we developed Simuran, an open-source software package. Simuran enables rapid adoption of new analysis methods and supports interoperability between data formats and analysis tools (chapter 3).
- With a focus on investigating the relationship between brain regions involved in spatial memory and processing, we characterise a spatial network involving the subiculum, hippocampus, thalamus, and retrosplenial cortex. Our findings highlight the importance of the anterior thalamic nuclei in the subiculum's spatial signalling (chapter 4).
- Expanding on this spatial circuit, we explore the dynamics of local field potentials (LFPs), which measure the summed electrical potential in the local area around

the recording site. We observe little effect on anterior thalamic nuclei (ATN) lesions on the LFP within the subiculum alone. The effects only appear when considering the coherence between the subiculum and retrosplenial cortex or spike-LFP relationships (chapter 5).

- Moving from analysing LFPs and lesion studies to neuron-level comparisons, we use simulations and statistics to estimate the probability of recording neurons in one brain region which receive direct synaptic connections from neurons recorded in another anatomically connected brain region. Our results suggest a high likelihood of recording a substantial proportion of connected neurons with current high-fidelity recording devices (chapter 6).
- With connections being likely between well-connected brain areas in large datasets, we investigate the correlation between neural activity and task performance. By analysing visual discrimination tasks, we investigate the complex relationship between correlated activity within and between brain regions and its connection to information representation. Furthermore, this same relationship is explored in LFPs from the claustrum, association cortex, and retrosplenial cortex from rats performing a task that switches between fixed interval and fixed ratio (chapter 7).

This thesis includes work that was performed in collaboration with members of Prof. O'Mara's lab. As such, for chapters that are not solely my contribution, I have included a summary of the work I performed and the work performed by others.

2 Overview of multi-region neural data analysis

We have the tools to look there; Whether we have the intelligence to figure it out, I view that, at least in part, as a theory problem.

Larry Abbott, speaking on understanding the brain

Since as early as the fourth century B.C., Greek scholars suggested that the brain was the organ of sensation (Bear et al., 2020). However, with the discovery of the neuron, and Cajal’s neuron doctrine – the belief that the individual neuron is the structural and functional unit of the nervous system (Cajal & Golgi, 1906), many important questions arose around the function of individual neurons and neuronal circuits. The technique of choice to understand these neuronal circuit dynamics in relation to behaviour continues to be recordings of neurons and associated signals, such as local-field potentials. The captured activity can range wildly in scale, from the blood oxygen level dependent (BOLD) signal in swathes of the brain, to highly precise patch-clamp measurements of activity from individual neurons in the brain.

Considering interactions between neurons in different brain regions has proved particularly challenging. One successful approach has been to use functional magnetic resonance imaging (fMRI) to measure the BOLD signal in the brain, with higher BOLD levels indicating localised active neurons (Ogawa et al., 1990), though see (Westlin et al., 2023). However, such measurements are limited in their temporal and spatial resolution, as the BOLD signal is a coarse-grained measure of neuronal activity, coupled with hundreds of thousands or millions of neurons in each fMRI voxel (Logothetis, 2008). Recent technical advances in neuronal recording methodologies (such as Neuropixels probes; Jun et al., 2017; Steinmetz et al., 2021) allow recording of large numbers of neurons within and between multiple brain regions simultaneously (W. E. Allen et al., 2019; Siegle et al., 2019; Steinmetz et al., 2019; Stringer et al., 2019). These approaches offer an order of magnitude greater neuron yield compared to previous neuron recording technologies, such as tetrodes, and are of remarkably high temporal and spatial resolution (though not without bias). Nonetheless, these recordings remain spatially- and temporally-incomplete samplings and achieving a true population recording is an unrealistic goal (Buzsáki, 2004).

These large-scale, inter-regional, conjoint recordings open the possibility to investigate more granular information processing rules between neurons in connected brain regions. Understanding these kinds of “source transformation rules” (e.g., how neurons transform their synaptic inputs) may enable us to probe how downstream neuronal observers take advantage of upstream signals (Buzsáki & Tingley, 2018). In turn, we might be able to approach how transformations of neural coding within and between brain regions might eventuate in the generation of behaviour and provide a mechanistic basis for understanding cognition.

Though such a goal is still far off, the wheels are certainly in motion. As an example, investigations into the different frequency channels information passes through have suggested that high-frequency signals are correlated with feedforward information while low-frequency signals are correlated with feedback (Bonaiuto et al., 2018; Márton et al., 2019; Vinck et al., 2023). Other studies have investigated topics such as the effect of external stimuli on the communication between neurons (Ruff & Cohen, 2016), the role of correlations between neurons and the action potential of individual neurons (Fournier et al., 2019; Panzeri et al., 1999), and the routing of neuronal activity between cortical areas (Friston, 2011; Semedo et al., 2019).

In this chapter, we will discuss existing approaches to recording from multiple brain regions simultaneously and methods for drawing insights from these massive datasets. We begin with background on recording brain signals, followed by experimental design and analysis methods, closing with the current state of the field of inter-areal interaction studies in the brain.

2.1 The primary intracranial brain signals recorded *in vivo*

Neural dynamics are commonly investigated in living animals, or *in vivo*, to avoid a disconnect from natural behaviours. As it goes, shortly after the discovery of place cells in the hippocampus (cells that activate in specific parts of the environment), Gary Lynch complained to John O’Keefe;

“I’ve tested your theory about these place cells and the spatial function of the hippocampus. I put my slice on wheels, moved it around the lab and it made no difference at all.” (Ferguson et al., 2011; O’Keefe and Dostrovsky, 1971)

Due to the difficulty of achieving intracellular recordings in awake, moving animals (though it has been achieved, e.g., Epsztein et al., 2011; Noguchi et al., 2021) and low cell yield, extracellular recordings are the primary method of choice. These recordings are performed outside the cell body, or soma, where the recording device sits in living

tissue nearby the cell. Traditionally, these recordings have been performed with a set of multiple thin wires to assist with the detection of cell activity (most commonly in a bundle of four wires called tetrodes). More recently, silicon probes with multiple recording sites or tetrode arrays have been used, along with calcium imaging, a technique that uses microscopy to optically measure fluorescent calcium indicators (Grienberger & Konnerth, 2012). Calcium imaging has been applied to record the change in calcium concentration in large populations of neurons in the awake animal (Stringer et al., 2019; Yang et al., 2018) and can be performed in conjunction with silicon probe recordings.

Such recording methods allow access to brain signals such as the electrical action potentials of single cells, calcium content of cells, and the local field potential (LFP) which measures the summed activity of all electrical potentials in the local area around the recording site. These brain measures are often combined with other observable chemical or behavioural aspects, such as running speed or blood sugar levels. However, such recordings do not capture synaptic activity between cells; the direct electrical and chemical messages passed between cells. In other words, though one of our primary goals is to investigate the information transfer between cells, we cannot directly measure this on a large scale currently due to technical limitations (Yoshida et al., 2018). Instead, we must rely on clever analysis methods to synthesise the information from single neuron or neural ensemble responses to infer communication strategies between these neurons. We describe two brain signals that are commonly considered in this research, action potentials and local field potentials.

2.1.1 Action potentials from single cells

Recording *in vivo* action potentials from cells, or “spikes” (so named due to the sharp peak in the characteristic electrical membrane potential waveform produced by an action potential), is generally achieved via extracellular electrical recordings or calcium imaging (CI).

In CI, a fluorescent indicator is loaded into the brain cells (usually chemically or genetically) and a microscope is used to capture a video of cells in the brain over time. The image is then segmented into putative individual cells, and the change in cell fluorescence over time is calculated and normalised by the resting fluorescence level, which is indicative of the calcium concentration in a cell and conventionally labelled $\frac{\Delta F}{F}$. Since the amount of calcium in a cell’s body is proportional to how active the cell is, we can then estimate when a cell is firing spikes according to the frame rate of the video capture – but we do not know exactly how many spikes were fired. As such, the estimation of spiking is not as exact as directly measuring the electrical potential,

which may be solved by voltage imaging (Gong et al., 2015). CI regularly provides the highest output of single-cell measurements of any recording modality, and the cells can be tracked across days since they are classified on morphological features. Additionally, a broad range of coverage is achieved for the cells, and these cells come with cell type information and position in the brain. Despite this, CI can suffer from a low frame rate, difficulty in recording from neurons deep in the brain, and issues around the exact relationship between calcium and spikes.

In extracellular recordings, the raw electrical potential is processed to identify action potentials by a characteristic waveform. These are then sorted to identify putative individual cells producing the action potentials. The putative cells are commonly referred to as “units” because the activity of multiple cells can be inseparable in the sorting process, resulting in a single contaminated “unit”. Once sorted, the action potentials for each unit are identified and represented as a spike train indicating when that unit fired (t_1, t_2, \dots, t_n). Extracellular recordings provide a finer time scale than calcium imaging, capturing short bursts of cell activity ($< 6\text{ms}$) and can more readily reach deeper brain regions, such as the thalamus. However, it is not without its issues, since cells are detected based on electrical activity alone. For instance, this method will miss any cells that are not regularly active, a cell type for which there is abundant evidence in the brain (Shoham et al., 2006). Furthermore, it is generally infeasible to track cells across days if the recording device is removed and re-inserted. It is also difficult to track cells across days even if the recording device is not moved, due to natural movement of the device in the brain.

Historically, both spike sorting and calcium segmentation were performed by hand or in a semi-automated manner, but increased recording yields have necessitated automatic tools (Buccino et al., 2019; Chaure et al., 2018; Pachitariu et al., 2016, 2017; Soltanian-Zadeh et al., 2019). While extracellular recordings and calcium imaging can both be used to infer action potentials (as the electrical action potential occurs along with the chemical influx of calcium into the cell body) it has been shown that there are differences between the action potentials obtained from calcium imaging and extracellular recordings (Huang et al., 2021; Wei et al., 2020). As such, the change in fluorescence detected via calcium imaging can only be seen as a rough estimation of spiking activity, which is partly due to the low sampling rate of calcium imaging.

Overall, the action potentials produced by individual cells are one of the smallest possible units for discussing inter-regional communication in the brain. However, in a cell, the integration of up to hundreds of action potentials along afferent synapses may be required for the cell to produce a single action potential, due to multiple factors, including synaptic failure, excitatory inhibitory balance, and time-synchrony of cell

inputs (Humphries, 2021a). Thus, large-scale recordings are a necessity for investigating inter-regional communication in the brain at the neuron level.

2.1.2 Local field potentials

The local field potential (LFP) is extracted from extracellular recordings by running a low-pass filter over the signal to leave low-frequency information from the raw electrical potential recorded. This field potential depends on the linear sum of potentials from each of the current sources and sinks in the brain, weighted according to distance from the recording site and extracellular conductivity (Pesaran et al., 2018). As such, the LFP captures information from many sources, including neurons, axons, dendritic spines, and glial cells. The LFP is represented as a time series of electrical potential values (V_1, V_2, \dots, V_N) , measured in Volts, at a fixed sampling rate f_s measured in Hertz.

The LFP is an extremely rich signal, which could be used for analysing large-scale brain dynamics (Buzsáki et al., 2012). LFPs are also of interest due to the “coherence as communication” hypothesis, which suggests that LFPs which fluctuate similarly may indicate communication between the regions involved (Fries, 2005; Fries, 2015). However, a recent review of coherence mechanisms suggests that it is communication which causes coherence, as opposed to coherence enabling selective communication (Vinck et al., 2023). Furthermore, if the probability of recording connected neurons between two regions is low, LFPs may allow us to analyse inter-regional communication with low cell counts as LFPs can reflect neural ensemble activity (Gallego-Carracedo et al., 2022). Despite this, LFPs are a difficult signal to work with, as they are sensitive to artefacts from biological activity, such as muscle activity as detected by an electromyogram (EMG), or outside interference such as phone signals or handling (Fabietti et al., 2020). The experimenter can also introduce unwanted artefacts into the signal through signal filtering choices (Widmann & Schröger, 2012).

To remove such artefacts from LFPs, a powerful method is independent component analysis (ICA), which is similar to principal component analysis (PCA) with an added constraint on the reduced dimensions (Jiang et al., 2019). In ICA, the LFPs are broken down into a set of independent components. Then, components that appear across all brain regions, which likely reflect noise or volume conduction, or clear noise such as a saw-like step, can be removed. The algorithm then combines the remaining components to form a cleaner, more region-specific signal. However, given multiple signals from one brain region, ICA can be hard to apply correctly, and one can end up removing components of actual brain signal. An alternative approach in this case is to simply flag

noisy channels, and only consider the rest, preferring to use channels on which spikes were also recorded. Other methods include projection methods, such as signal-space projection, to project the signals into a space where the noise dimension is removed. Automated machine learning methods to remove artefacts from LFPs are also growing in popularity (Fabiatti et al., 2021). In many of these methods, the signals are first normalised by Z-scoring to remove DC and perform scaling.

In addition to artefacts in LFPs, there are other difficulties in analysing LFPs (Herreras, 2016). Many of these difficulties arise from a clear observation – it is impossible to know the exact source of LFPs (Pesaran et al., 2018). One can consider LFPs via forward and inverse models. The forward model, which describes for a given set of neuronal activity what LFP signal is produced, is a well-posed problem. However, the inverse problem, to describe neuronal activity from a given LFP signal, is ill-posed as multiple different configurations can give rise to the same LFP. Since the inverse problem is the general aim of LFP analyses, this creates clear difficulties, especially in any form of causal analysis. Furthermore, effects contribute to LFPs beyond single source and sink potentials. For instance, small-amplitude current sources that are arranged into particular spatial configurations across neurons can give rise to large-amplitude LFPs (Pesaran et al., 2018). This causes many issues, such as unclear sources of LFPs (Einevoll et al., 2013) and concerns over volume conduction, especially in the rodent brain (Lindén et al., 2011).

While investigation into LFPs is ongoing, the common consensus within the field is that LFPs do not exert a causal influence on brain processing. As such, they are considered as a by-product or epiphenomenon of neural activity (Herreras, 2016). Nevertheless, it remains plausible that LFPs could play a functional role by influencing the discharge properties of neurons through alterations in the conductive extracellular space surrounding them (Buzsáki et al., 2012; Fröhlich & McCormick, 2010). Evidence indicates that the extracellular field can impact the membrane potential in neurons through ephaptic coupling (Anastassiou & Koch, 2015; K.-S. Han et al., 2018; Jefferys, 1995; Martinez-Banaclocha, 2018; Su et al., 2012). Consequently, a functional coupling between single-unit activity and LFPs may exist (Pinotsis et al., 2023), necessitating a comprehensive analysis of their interconnected contributions. Whether functionally related or a by-product, LFPs are packed with information, though the impossibility to form an exact inverse back to neuronal activity impedes interpretation.

2.2 Recording the brain in the body

Historically, the focus of electrophysiological experiments was on the function of individual neurons or systems. For instance, Nobel prizes have been awarded for the discovery of single-cell coding in various brain areas, such as the joint award in 2014 for place cells in the hippocampus (O’Keefe & Dostrovsky, 1971) and grid cells in the entorhinal cortex (Hafting et al., 2005). While an understanding at this level is crucial, researchers are gradually considering an integrative approach instead of reductionist methodologies.

Integrative approaches range from considering the relationship between the brain and other organs, or “putting the brain back in the body” (Alonso & Marder, 2020; Tingley, 2021), to considering how the inputs and outputs of multiple brain regions are related (Fournier et al., 2020; Siegle et al., 2021). Related to this, there has been a resurgence in the idea of the neural population doctrine (Saxena & Cunningham, 2019; Yuste, 2015). As opposed to the neuron doctrine, whereby the single neuron is the base unit of computation (Cajal & Golgi, 1906), in the population doctrine, the functional unit of the neural circuit is an ensemble or group of linked neurons. These ensembles can be functional groupings or structural groupings and range from small sets of closely-knit neurons in a single layer of a brain region to large chains of neurons in multiple brain regions (Westlin et al., 2023).

A key factor in the experimental design of inter-areal interaction studies regularly involves linking large-scale multi-region brain recordings to behaviour. An example is fMRI studies where the brain is scanned while subjects are resting or performing tasks, and then the brain activity networks are correlated to the subject’s state (Humphries, 2017; Ju & Bassett, 2020; Kajimura et al., 2023; Koppe et al., 2019; Raichle, 2015). For many, a goal of recording the brain and behaviour is to elucidate how information transfer and interaction between brain regions gives rise to complex behaviours. Simply, neuroscience needs behaviour, as behaviour and neural data are intricately linked (Krakauer et al., 2017; Urai et al., 2021). A strong link to behaviour also allows for top-down modelling, in which the brain circuit under investigation is considered in the context of the problems it is required to solve (Eliasmith & Trujillo, 2014). Current experiments in integrative studies generally follow one of a few paradigms.

In exploratory studies, subjects are recorded in multiple different environments and interesting patterns in brain activity are sought out. Such studies are highly valuable, as a wide range of behavioural tasks can be established, and the data can be publicly shared and re-used to investigate a wide range of topics (Fournier et al., 2020; Siegle et al., 2021; Stringer et al., 2019). However, any causal analysis performed (e.g., Baccalá & Sameshima, 2001; Kim et al., 2011) in such an experiment only indicates

the potential presence of a causal effect and cannot claim to prove a causal effect in the brain. These patterns and any causal analysis findings can then be investigated in a later experimental manipulation study or with computational models (Alonso & Marder, 2020; Dayan & Abbott, 2001; Sussillo et al., 2015).

In experimental manipulation or intervention studies, the natural function of one or more brain regions can be modified while simultaneously recording in multiple brain regions. Two recent developments allow for fine control of experimental manipulations, optogenetics; the modulation of neurons via light (Deisseroth, 2015) and chemogenetics; the modulation of neurons or glia via chemicals (commonly DREADDs – Designer Receptors Exclusively Activated by Designer Drugs; Roth, 2016). For instance, optogenetics allows for rapid on/off stimulation of specific neural populations (Carrillo-Reid et al., 2016), though there are concerns over the unnatural effects of optogenetics (B. D. Allen et al., 2015). Nonetheless, the development of optogenetics has enabled testing certain causal theories, and long-standing models of neural information representation.

More wide-scale manipulations can be performed by permanent lesions, such as N-Methyl-D-aspartic acid infusion, which kills a considerable proportion of cells near the injection site due to excitotoxicity (Salazar et al., 2004). A further viable method is via the control of receptors (protein structures that receive or transmit signals in a biological system) such as *gamma*-aminobutyric acid or N-Methyl-D-aspartic acid (NMDA). Control of these receptors can be performed via drugs, such as D-serine as an NMDA receptor agonist, or muscimol as a *gamma*-aminobutyric acid (GABA) receptor agonist, leading to temporary inhibition or excitation of specific brain regions (Rizzello et al., 2022). Regardless of the experimental design, several viable analysis methods are available for drawing insight from the resulting data.

2.3 Fundamentals concepts for analysing intracranial recordings

To analyse data from neural populations, concepts from statistics, mathematics and systems modelling have been applied. These include dynamical systems modelling, graph theory, dimensionality reduction, and spike correlations. For instance, recent research applied a Wilson-Cowan model to simulate a neuronal population and test theories about communication by comparing simulated data to real data (Ito et al., 2019). Such neural population analysis is not a new idea by any means (Georgopoulos et al., 1986; Nicolelis et al., 1997), but given that the number of simultaneously recordable neurons doubles approximately every 7 years (Urai et al., 2021) we are in a much better position from a data perspective to test theories on ensemble computation and communication.

2.3.1 Types of neural correlation

There are many types of neural correlation that can be considered under the term “correlation” (Kohn et al., 2016), such as:

- Signal correlation, the correlation between neurons’ response tuning curves.
- Pattern correlation, the relation of population level responses from the same neurons, such as the matrix of spike times, to different stimuli.
- “Noise” correlation, the relation of the response of the same neurons to a fixed stimulus across different presentations of that stimulus. This is somewhat of a misnomer, and, as such, other terms are preferred, such as trial-to-trial correlation, or spike count correlations.
- Temporal correlation, the relation of past responses from the same neuron to future responses.

Many methods combine concepts from multiple types of neural correlation. For instance, Gaussian process factor analysis (Yu et al., 2009) considers the response of multiple neurons (pattern correlation), across time (temporal correlation), the result can be averaged in response to the same stimulus (noise or trial-to-trial correlation), and the relation of the responses to different stimuli can be compared (signal correlation).

2.3.2 Encoding and decoding of neural signals

At the highest level, analysis methods fall into two broad categories; encoding and decoding (Dayan & Abbott, 2001; Stevenson & Kording, 2011). To illustrate the difference, imagine sending a secret message via a cypher, which produces a code. You know the true message, or what has been encoded by the cypher to represent your message. But anyone else would have to infer the true message from the code they are provided, i.e., perform decoding. In the brain, the encoder would produce the code that represents sensory signals, while the decoder parses the message from this code to enable successful behaviour (Kriegeskorte & Douglas, 2019).

In terms of analysis, neural decoding involves estimating the value of an external factor, such as the location of the subject, or predicting behaviour from a set of neuronal signals (i.e., predicting external factors from given brain activity). This is essentially what the brain does to draw information from internal brain signals, but performing the calculations ourselves can give insight into the reliability of a certain response pattern in the brain for coding a variable. However, it is possible to decode information that is

completely different from what the brain is encoding – the message that is being stored in the brain. As such, the reverse problem of encoding is also well studied, investigating for a given external stimulus if any neurons (or group of neurons) reliably code for that stimulus (i.e., predicting internal brain activity from given external factors).

We will discuss both of these types of analysis methods based on whether they use information from spikes alone, or if the information from LFPs are incorporated into the analysis. Some analyses use the entire raw electrical potential signal without any processing for decoding purposes (Frey et al., 2019; Trautmann et al., 2019), or use multi-unit activity (Drebitz et al., 2019; S. Han et al., 2017; Smith et al., 2013). These two methods are particularly suited to online processing, such as for brain-machine interfaces, as time-intensive spike sorting is not required. However, raw electrical potentials can be analysed similarly to LFPs, and analysing thresholded multi-unit activity is similar to analysing spike trains. As such, the discussion on spike train analysis and the analysis of LFPs also covers many of the foundations for the analysis of multi-unit activity and the raw electrical potential. This section will focus on the basic ideas shared between many analysis methods and the building blocks of analysis methods, while section 2.4 will give more in-depth examples of currently used methods.

2.3.3 Approaches to analysing spike trains

Recording spikes from many neurons in multiple brain regions simultaneously was not possible until recently. As such, most methods that considered spikes in multiple brain regions involved correlations between the activity of pairs of spiking neurons, or the activity of a group of neurons in one area and a single neuron in another. Most of these ensemble spike train methods involve representing the spike train(s) as a matrix or vector and relating this to another matrix or vector. Many of these methods are easier to perform in controlled trials with a timed stimulus, but not all experiments are amenable in this manner. Nonetheless, we will discuss the basis of inter-areal data analysis methods.

One general approach is to directly use tuning curves or to modify the concept of tuning curves. Given a variable (e.g., running speed) and the spike train of a neuron, a tuning curve plots the value of the variable on the x-axis against the firing rate of the neuron on the y-axis. In a multi-neuron recording, these tuning curves can be calculated independently per neuron. In another similar independent analysis, a histogram is plotted around the time of stimulus presentation to estimate the response rate of a neuron over time in relation to a stimulus, called a peristimulus time histogram (PSTH). These ideas facilitate computationally feasible forms of multi-neuron analysis,

as the tuning curves or PSTH are readily calculated independently per neuron, and then multiple tuning curves can be combined or compared. This idea led to many popular analysis methods, such as the population vector method (Georgopoulos et al., 1986; Georgopoulos, 2014) to combine the preferred tuning of multiple neurons into a system-level approximation of coding, the joint PSTH (Gerstein, 2013) and the population distance index (Nicoletis, 2008) which both compare the response of two neurons in relation to a stimulus.

However, it has been shown that these methods built around tuning curves do not scale accurately to large-scale recordings as the independence assumption leads to a constant accuracy rate despite increasing neuron numbers (Stevenson & Kording, 2011). Furthermore, where trial-averaged metrics are considered, there is question over how relevant trial-averaged metrics are to neuronal computation (Tlaie et al., 2021). As such, methods which account for the more real-world situation where neuron responses are highly correlated and aim to account for these interactions are important. These methods commonly act on the binned spike train matrix for tractability, which is a matrix where each row represents a neuron and each column represents a time bin, and the value of each element is the number of spikes in that time bin for that neuron. An early example of this was the gravity clustering system whereby neurons were represented as particles in space and neuronal responses acted as gravitational forces to move the particles in space (Abeles & Gerstein, 1988; Egert et al., 2002). Though a significant drawback of this method is that it is limited to identifying small chains of interactions.

2.3.4 Drawing insight from the local field potential

To investigate brain dynamics, a promising approach is to record neurons from many distinct brain areas, but these patterns in neuronal activity have redundancies (Cunningham & Yu, 2014; Gallego-Carracedo et al., 2022) and as such LFPs might be used in place of large-scale neuronal recordings. Furthermore, by inferring LFPs from neural sources (forward modelling), possible neuronal correlates for patterns observed in LFPs may be suggested (Einevoll et al., 2013). For example, could highly correlated LFPs in separated regions be observed merely by the volume conduction of uncorrelated synaptic inputs? Regardless of the method of choice, and whether pure relations between LFPs are considered or the relationship between LFPs and spikes, two main categories exist. LFPs can be considered in the time domain (the standard x -axis time vs y -axis electrical potential representation), or the frequency domain via a z -transform, or more commonly a discrete Fourier transform (DFT).

In the frequency domain, the component of the signal at a given frequency can be thought of as having an amplitude and relative phase. The frequency domain representation is particularly powerful and has led to methods such as partial directed coherence (Baccalá & Sameshima, 2001), a frequency domain form of Granger causality, phase-amplitude coupling of LFPs (Hülsemann et al., 2019), and wavelet coherence (Li et al., 2007). One drawback of coherence-based measures is that volume conduction in the rodent brain can cause issues for analysis (Herreras, 2016; Kajikawa & Schroeder, 2011). Recent studies have also suggested that such coherence is a mere by-product of connectivity and power and is produced by communication as opposed to enabling it (Schneider et al., 2020; Vinck et al., 2023). As for partial directed coherence, as with any causality analysis, it can miss the true driver. For instance, if $X \rightarrow Y, X \rightarrow Z$ (where \rightarrow denotes a causal relationship) but X is not recorded yet Y and Z are recorded, then causality analysis will generally give the incorrect result that $Y \rightarrow Z$. Additionally, it should be noted that Granger mentioned the possibility of ridiculous results from applying Granger causality outside of the financial domain in their Nobel prize acceptance speech, though pointed out that a definition of “true causation” was never provided by critics (Granger, 2003).

In the time domain, two popular methods are power analysis, usually performed by integrating the periodogram or predicting a variable from LFPs, called signal decoding (Pesaran et al., 2018). Furthermore, spikes can be considered with LFPs in the time domain very naturally, as they are both time series. This leads to a common measurement of spike-LFP relationship; spike-triggered average, the average LFP around the time of a spike. Spike-LFP causality is a lucrative topic and commonly assessed by Granger causality, though such analysis is difficult, and perhaps intractable, due to the aforementioned ill-posed inverse model problem for LFPs. While the time-domain spike-LFP comparison is quite natural, LFPs can also be related to spikes in the frequency domain by convolving a sum of Dirac deltas with a sine wave kernel (Fanselow et al., 2001), or considering the spikes as a point process, such as in spike field coherence calculations and field-triggered average, which measure how well the spikes line up with phasic oscillations in LFPs.

Rhythmicity in the local field potential

Neuronal oscillations are well known to show distinct frequency patterns in LFPs (Buzsáki & Draguhn, 2004; Lopes da Silva, 1991), particularly in the delta (1 – 4Hz), theta (6 – 12Hz), beta (15 – 30Hz), and gamma (40 – 120Hz) bands. Note that bands and ranges differ per species, and these bands and ranges are given for rodents (for instance, alpha activity from 8 – 12Hz is more prevalent in the human brain).

Each of these oscillation frequencies are related to different cognitive processes. Delta activity is primarily increased during deep sleep (Amzica & Steriade, 1998; Long et al., 2021), though, during wakefulness, delta waves may contribute to long-scale decision-making and concentration (Harmony, 2013; Schultheiss et al., 2020). Theta waves are linked with spatial signalling and increased motor activity, potentially enabling memory consolidation and decision-making (Buzsaki, 2002; Cacucci et al., 2004; Colgin, 2013; Geisler et al., 2007; Kay et al., 2016; Richard et al., 2013; Solomon et al., 2019; Womelsdorf et al., 2010). Beta waves in the hippocampus are related to object-location memory performance, learning, and novelty detection (Berke et al., 2008; França et al., 2014; Grossberg, 2009; Iwasaki et al., 2021). Gamma waves are implicated in decision-making processes (Buzsáki & Wang, 2012; Colgin et al., 2009; Montgomery & Buzsáki, 2007; Trimper et al., 2017) and a relationship between theta phase and gamma amplitude has been demonstrated (Mysin & Shubina, 2022; Scheffer-Teixeira & Tort, 2016). Further afield, neocortical neurons and gamma oscillations relate to theta rhythms (Sirota et al., 2008). As such, each of these oscillation frequencies can be of interest in different contexts, and the complex relationship between them is not fully understood.

2.4 Recent techniques for multi-region analysis and future directions

While the topic of inter-areal interactions and considering the relationship between the brain and behaviour is gaining traction, there are significant challenges. A number of these challenges involve the possibility of identifying spurious interactions or misusing statistics. Other challenges question the general form of investigation used to try and understand the brain. Even the term *understanding the brain* requires thought; does this mean being able to predict future neuron-level brain activity from a given state? Would accurately describing the activity of small circuits in the brain or the evolution of patterns of brain activity in response to stimuli suffice? Perhaps we cannot describe *understanding* concretely and are working towards personal and subjective goals. Nonetheless, remarkable applications of neural ensemble research exist, and a few of these applications involve multi-region communication research. We describe these topics here and raise several major outstanding questions and avenues that may be interesting to pursue.

2.4.1 Applying machine learning to large-scale neural data

The field of metascience has brought to light many difficulties surrounding scientific research practices, and neuroscience is no exception (Button et al., 2013). Many published research findings are likely to be false (more than would be expected by chance; Ioannidis, 2005) and science as a whole faces reproducibility issues (Munafò et al., 2017). These issues arise from reliance on null hypothesis significance testing and an over-emphasis on achieving statistically significant p -values, as opposed to considering effect size, statistical power, or applying the appropriate statistical tests. This has led to a shift in how many experiments are performed and analysed. Of note is the prevalence of machine learning (ML).

Machine learning is a form of computer algorithm designed to recognise patterns in large sets of training data, and then use these patterns to make predictions on unseen data. Machine learning models can find patterns in complex datasets, but the resulting models are notoriously difficult to inspect, often labelled as a black box (though see explainable AI; Barredo Arrieta et al., 2020). A further issue with model complexity in machine learning relates to bias and over-optimism in science. Machine learning can find patterns in anything. Given random data with low samples and a high number of variables, a logistic regression model regularly achieves a misleadingly high accuracy for within training samples (Whelan & Garavan, 2014). Similarly, if machine learning identifies a pattern in data, it must be asked if this should simply be expected by chance, or because of the data structure (Elsayed & Cunningham, 2017). In this manner, formal evaluations of different machine learning methods in neuroscience are of high importance (such as in neuroimaging, Jollans et al., 2019).

Encompassed in the umbrella of machine learning methods are artificial neural networks (ANNs). ANNs drew inspiration from the brain (Richards et al., 2019), forming nodes (like neurons) that were connected (like synapses) and the weights updated via learning (like synaptic plasticity, e.g., Hebbian learning; Hebb, 1949). Trained ANNs have been used to decode images from fMRI activity that resemble images presented to the fMRI subject (Shen et al., 2019). ANNs have limited biological plausibility, as they, for instance, apply backpropagation of errors through the network to update synaptic weights, which there is little evidence in favour of in the brain (Macpherson et al., 2021). Furthermore, as the accuracy of ANNs increases, they begin to resemble the brain less and less (Nonaka et al., 2020). Considerations of this matter are increasing the development of spiking neural networks, which more closely mimic physical biological networks, like the brain (Macpherson et al., 2021; Pfeiffer & Pfeil, 2018; Taherkhani et al., 2020).

Despite difficulties in applying machine learning, the ability to handle complexity has opened many avenues in electrophysiological data analysis (Cunningham & Yu, 2014; Gokcen et al., 2021; Mastrogiuseppe & Ostojic, 2018; Pandarinath et al., 2018; Viejo et al., 2018; Vyas et al., 2020). Outside of using ANNs or spiking neural networks to create models to compare brain function and circuitry, the primary form of machine learning applied to electrophysiological data is dimensionality reduction. The application of machine learning has suggested that neurons communicate in groups to solve tasks in the Visual Cortex (Semedo et al., 2019, 2022) and has also been applied to estimate latent (hidden) brain states (Gallego et al., 2018; Semedo et al., 2014). A particularly prominent example of machine learning applied in electrophysiology is latent factor analysis via dynamical systems (LFADS), which has been trained to accurately predict rhesus macaque monkeys' arm movements in a reaching task using neuronal data recorded from the motor and dorsal premotor cortices (Pandarinath et al., 2018) and has been expanded to calcium imaging (Zhu et al., 2021). This emerging view of the neural circuit as a dynamical system is fascinating, suggesting that the dynamics of neural circuits can be described by an initial state and a vector field describing the direction and magnitude of neural activity over the state space and has been assessed using different forms of dimensionality reduction (Cunningham & Yu, 2014; Glaser et al., 2020; Humphries, 2021b; Pang et al., 2016; Umakantha et al., 2021; Yu et al., 2009). Motor cortex is currently the most promising area for discovering dynamical system properties.

Such studies are exciting and offer promise for the future in many areas, such as human-computer interaction using data from neural ensembles (Chapin et al., 1999; Nicolelis, 2008). Overall, this has contributed to a new view on the role of cell assemblies, and a resurgence in the neural population doctrine (Eichenbaum, 2018; Machado et al., 2022; Nicolelis et al., 1997; Saxena & Cunningham, 2019; Yuste, 2015).

2.4.2 Future avenues for cross-region neural ensemble analysis

Machine learning techniques are promising for investigating the relationship between many variables in multiple brain regions (Gokcen et al., 2021; Semedo et al., 2020). Here we list some findings that have resulted from these techniques and raise related outstanding questions. Neural activity has redundancies, which has allowed for accurate prediction of neural firing rates in many brain regions (Pang et al., 2016; Perich et al., 2018; Pesaran et al., 2018; Williams & Linderman, 2021). But what does this redundancy or correlation represent, and given it takes energy to fire cells, why is this redundancy needed? Perhaps it enables large-scale routing of neural information (Semedo et al., 2019), or represents robust features and a stable and confident representation of the

environment (Cayco-Gajic & Zylberberg, 2021). If this is the case, we might expect this correlation to be linked to behaviour (Ito et al., 2019; Larry & Joshua, 2021). Alternatively, it could reflect synaptic failure, connection strength, and the requisite integration of multiple presynaptic action potentials at the soma to regulate the cell’s activity (Humphries, 2021a). Similarly, correlation over random tensors might display similar features to neural data correlations which would suggest this redundancy is a by-product of such a high-dimensional system (Elsayed & Cunningham, 2017). This kind of redundancy has, for example, led to the previously discussed consideration of the brain as a dynamical system or viewing of neural activity as lying in a low-dimensional manifold. These are exciting, but many outstanding questions remain, and control methods are necessary to avoid detecting spurious patterns (Vyas et al., 2020).

We also know that LFPs have reliable links to neural activity both within and between regions. One of the most well-known examples within a region is phase precession in the hippocampus, which is a link between the phase of the slow-wave theta band (around 8 Hz) signal and place cell firing (O’Keefe & Recce, 1993). In a more recent study, the firing rate of one-quarter of recorded cells in V1 was found to be influenced by the phase of theta oscillations from the hippocampus (namely CA1), indicating a further link between spatial coding and visual processing (Fournier et al., 2020). Another study points to the importance of the theta band oscillations in the anterior cingulate cortex (ACC) for hippocampal processing (Wirt & Hyman, 2019). These observations line up with the suggestion of a form of fundamental neural clock via brain oscillations (Buzsáki & Draguhn, 2004; Buzsáki & Tingley, 2018). Overall, this points to the importance of LFPs, which may turn out to be one of the most important brain signals for understanding neuronal computations (Buzsáki et al., 2012).

Many neural ensemble recordings have been performed concerning brain-machine interface research. The aim is to identify reliable patterns of brain activity and use this activity to control a machine using the mind (Mehonic & Kenyon, 2022). This has resulted in amazing achievements, such as Juliano Pinto, a 29-year-old paraplegic man, kicking off the 2014 World Cup in Brazil (Bates, 2017; Donati et al., 2016). In non-human primates, similar experiments have allowed for robotic arm control (O’Doherty et al., 2011), full virtual-reality arm control with feedback (Ifft et al., 2013), and playing pong (Musk & Neuralink, 2019), all using the mind alone. Other tentative applications exist which use communication rules for multiple brain regions, such as using the transformation from the retina to six interconnected brain areas called the ventral visual stream to activate desired neurons based on sets of *controller images* (Bashivan et al., 2019). Applications of this encoding and decoding work from multiple brain regions will continue to be a primary area of interest for future studies, as they are the most readily available commercial applications of the work.

Despite this, it has been suggested that using recent techniques to investigate the brain, a neuroscientist could not even understand the inner workings of a microprocessor (Jonas & Kording, 2017). Similar problems are faced in breaking apart the computations performed by large ANNs (Barredo Arrieta et al., 2020). This suggests that we may need an overhaul in our approach to brain investigation. A potentially fruitful avenue is to develop interpretable general methods, like Skaggs information content, which applies to any spike train and another variable (Skaggs et al., 1993). Another line of thinking which could prove vital is to avoid relating neuronal activity to external variables like space or time and to rethink basic assumptions and metaphors in neuroscience (Buzsáki, 2020; Buzsáki & Tingley, 2018; Westlin et al., 2023). The idea is, the brain has no direct representation of these, they are an abstract term that we have created as humans and, instead, we should more accurately describe neuronal activity in terms of sequential activity or internal rate of change.

2.5 Conclusion and outlook for this thesis

Hopefully this review provides some insight into the multi-region neural analyses applied to data recorded *in vivo* and a motivation for several of the challenging but fascinating open questions in the field. To give context, we will close with a commentary on how this review informs the later chapters of this thesis:

- It is evident from our review that many methods used in neural data analysis are borrowed from other domains like signal processing, machine learning, and mathematics. To leverage these algorithms, we develop software that can convert neuroscience-specific data formats into more general formats, representing the data fundamentally as vectors, matrices, and time-series (chapter 3).
- To explore the effects of manipulations on a spatial memory circuit in the brain, we focused on temporary inactivation and permanent lesions of the anterior thalamic nuclei. By employing spike sorting and information theoretic methods, we examined the impact of these manipulations at the single neuron level in the subiculum and CA1 regions (chapter 4). Furthermore, we employed time- and frequency-domain methods to study the influence of these manipulations on local field potentials in the subiculum, CA1, and retrosplenial cortex (chapter 5).
- This review emphasises the advances in high yield recording techniques such as calcium imaging and silicon probes and the potential impact for assessing neural communication. We explore the concept of source transformation rules prompted by this review and theorise the feasibility of investigating these rules at the single

neuron level by estimating the number of cross-region synaptically connected neurons that could potentially be recorded in a single session (chapter 6).

- Finally, we investigate the role of neural redundancy and correlation on large scale open datasets with machine learning. Through dimensionality reduction techniques and coherence measures, we examine this relationship in local field potentials from the claustrum, retrosplenial cortex, and anterior cingulate cortex during switching fixed interval and fixed ratio tasks. Furthermore, we explore this relationship using single unit data from the visual and limbic system areas during a visual discrimination task (chapter 7).

References

- Abeles, M., & Gerstein, G. L. (1988). Detecting spatiotemporal firing patterns among simultaneously recorded single neurons. *Journal of Neurophysiology*, *60*(3), 909–924.
- Allen, B. D., Singer, A. C., & Boyden, E. S. (2015). Principles of designing interpretable optogenetic behavior experiments. *Learning & Memory*, *22*(4), 232–238.
- Allen, W. E., Chen, M. Z., Pichamoorthy, N., Tien, R. H., Pachitariu, M., Luo, L., & Deisseroth, K. (2019). Thirst regulates motivated behavior through modulation of brainwide neural population dynamics. *Science*, *364*(6437), eaav3932.
- Alonso, L. M., & Marder, E. (2020). Temperature compensation in a small rhythmic circuit (F. K. Skinner, G. L. Westbrook, F. Nadim, & A. Guet-McCreight, Eds.). *eLife*, *9*, e55470.
- Amzica, F., & Steriade, M. (1998). Electrophysiological correlates of sleep delta waves1correspondence and reprint requests may be addressed to either author.1. *Electroencephalography and Clinical Neurophysiology*, *107*(2), 69–83.
- Anastassiou, C. A., & Koch, C. (2015). Ephaptic coupling to endogenous electric field activity: Why bother? *Current Opinion in Neurobiology*, *31*, 95–103.
- Baccalá, L. A., & Sameshima, K. (2001). Partial directed coherence: A new concept in neural structure determination. *Biological Cybernetics*, *84*(6), 463–474.
- Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, *58*, 82–115.
- Bashivan, P., Kar, K., & DiCarlo, J. J. (2019). Neural population control via deep image synthesis. *Science*, *364*(6439), eaav9436.

- Bates, M. (2017). From brain to body: New technologies improve paralyzed patients? quality of life. *IEEE Pulse*, 8(5), 22–26.
- Bear, M., Connors, B., & Paradiso, M. A. (2020). *Neuroscience: Exploring the brain*. Jones & Bartlett Learning, LLC.
- Berke, J. D., Hetrick, V., Breck, J., & Greene, R. W. (2008). Transient 23–30 hz oscillations in mouse hippocampus during exploration of novel environments. *Hippocampus*, 18(5), 519–529.
- Bonaiuto, J. J., Meyer, S. S., Little, S., Rossiter, H., Callaghan, M. F., Dick, F., Barnes, G. R., & Bestmann, S. (2018). Lamina-specific cortical dynamics in human visual and sensorimotor cortices. *eLife*, 7, e33977.
- Buccino, A. P., Hurwitz, C. L., Magland, J., Garcia, S., Siegle, J. H., Hurwitz, R., & Hennig, M. H. (2019). SpikeInterface, a unified framework for spike sorting. *bioRxiv*, 796599.
- Button, K. S., Ioannidis, J. P. A., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S. J., & Munafò, M. R. (2013). Power failure: Why small sample size undermines the reliability of neuroscience. *Nature Reviews. Neuroscience*, 14(5), 365–376.
- Buzsáki, G. (2002). Theta oscillations in the hippocampus. *Neuron*, 33(3), 325–40.
- Buzsáki, G. (2004). Large-scale recording of neuronal ensembles. *Nature Neuroscience*, 7(5), 446–451.
- Buzsáki, G. (2020). The brain–cognitive behavior problem: A retrospective. *eNeuro*, 7(4).
- Buzsáki, G., Anastassiou, C. A., & Koch, C. (2012). The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, 13(6), 407–420.
- Buzsáki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science*, 304(5679), 1926–1929.
- Buzsáki, G., & Tingley, D. (2018). Space and time: The hippocampus as a sequence generator. *Trends in Cognitive Sciences*, 22(10), 853–869.
- Buzsáki, G., & Wang, X.-J. (2012). Mechanisms of gamma oscillations. *Annual review of neuroscience*, 35, 203.
- Cacucci, F., Lever, C., Wills, T. J., Burgess, N., & O’Keefe, J. (2004). Theta-modulated place-by-direction cells in the hippocampal formation in the rat. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 24(38), 8265–8277.
- Cajal, S. R. y., & Golgi, C. (1906). The neuron doctrine: Theory and facts. *Nobel lecture, 1921*, 190–217.
- Carrillo-Reid, L., Yang, W., Bando, Y., Peterka, D. S., & Yuste, R. (2016). Imprinting and recalling cortical ensembles. *Science*, 353(6300), 691–694.

- Cayco-Gajic, N. A., & Zylberberg, J. (2021). Good decisions require more than information. *Nature Neuroscience*, *24*(7), 903–904.
- Chapin, J. K., Moxon, K. A., Markowitz, R. S., & Nicolelis, M. A. L. (1999). Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nature Neuroscience*, *2*(7), 664–670.
- Chaure, F. J., Rey, H. G., & Quian Quiroga, R. (2018). A novel and fully automatic spike-sorting implementation with variable number of features. *Journal of Neurophysiology*, *120*(4), 1859–1871.
- Colgin, L. L. (2013). Mechanisms and functions of theta rhythms. *Annual Review of Neuroscience*, *36*, 295–312.
- Colgin, L. L., Denninger, T., Fyhn, M., Hafting, T., Bonnevie, T., Jensen, O., Moser, M.-B., & Moser, E. I. (2009). Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature*, *462*(7271), 353–357.
- Cunningham, J. P., & Yu, B. M. (2014). Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, *17*(11), 1500–1509.
- Dayan, P., & Abbott, L. F. (2001). *Theoretical neuroscience: Computational and mathematical modeling of neural systems*. Massachusetts Institute of Technology Press.
- Deisseroth, K. (2015). Optogenetics: 10 years of microbial opsins in neuroscience. *Nature neuroscience*, *18*(9), 1213–1225.
- Donati, A. R. C., Shokur, S., Morya, E., Campos, D. S. F., Muioli, R. C., Gitti, C. M., Augusto, P. B., Tripodi, S., Pires, C. G., Pereira, G. A., Brasil, F. L., Gallo, S., Lin, A. A., Takigami, A. K., Aratanha, M. A., Joshi, S., Bleuler, H., Cheng, G., Rudolph, A., & Nicolelis, M. A. L. (2016). Long-term training with a brain-machine interface-based gait protocol induces partial neurological recovery in paraplegic patients. *Scientific Reports*, *6*(1), 30383.
- Drebitz, E., Schledde, B., Kreiter, A. K., & Wegener, D. (2019). Optimizing the yield of multi-unit activity by including the entire spiking activity. *Frontiers in Neuroscience*, *13*.
- Egert, U., Knott, T., Schwarz, C., Nawrot, M., Brandt, A., Rotter, S., & Diesmann, M. (2002). MEA-tools: An open source toolbox for the analysis of multi-electrode data with matlab. *Journal of Neuroscience Methods*, *117*(1), 33–42.
- Eichenbaum, H. (2018). Barlow versus hebb: When is it time to abandon the notion of feature detectors and adopt the cell assembly as the unit of cognition? *Neuroscience Letters*, *680*, 88–93.
- Einevoll, G. T., Kayser, C., Logothetis, N. K., & Panzeri, S. (2013). Modelling and analysis of local field potentials for studying the function of cortical circuits. *Nature Reviews Neuroscience*, *14*(11), 770–785.

- Eliasmith, C., & Trujillo, O. (2014). The use and abuse of large-scale brain models. *Current Opinion in Neurobiology*, *25*, 1–6.
- Elsayed, G. F., & Cunningham, J. P. (2017). Structure in neural population recordings: An expected byproduct of simpler phenomena? *Nature Neuroscience*, *20*(9), 1310–1318.
- Epszstein, J., Brecht, M., & Lee, A. K. (2011). Intracellular determinants of hippocampal CA1 place and silent cell activity in a novel environment. *Neuron*, *70*(1), 109–120.
- Fabietti, M., Mahmud, M., Lotfi, A., Averna, A., Gugganmos, D., Nudo, R., & Chiappalone, M. (2020). Neural network-based artifact detection in local field potentials recorded from chronically implanted neural probes. *2020 International Joint Conference on Neural Networks (IJCNN)*, 1–8.
- Fabietti, M., Mahmud, M., Lotfi, A., Kaiser, M. S., Averna, A., Guggenmos, D. J., Nudo, R. J., Chiappalone, M., & Chen, J. (2021). SANTI: A matlab-based open-source toolbox for artifact detection and removal from extracellular neuronal signals. *Brain Informatics*, *8*(1), 14.
- Fanselow, E. E., Sameshima, K., Baccala, L. A., & Nicolelis, M. A. L. (2001). Thalamic bursting in rats during different awake behavioral states. *Proceedings of the National Academy of Sciences of the United States of America*, *98*(26), 15330–15335.
- Ferguson, J. E., Jackson, J. C., & Redish, A. D. (2011). An inside look at hippocampal silent cells. *Neuron*, *70*(1), 3–5.
- Fournier, J., Saleem, A. B., Diamanti, E. M., Wells, M. J., Harris, K. D., & Carandini, M. (2019). Modulation of visual cortex by hippocampal signals. *bioRxiv*, 586917.
- Fournier, J., Saleem, A. B., Diamanti, E. M., Wells, M. J., Harris, K. D., & Carandini, M. (2020). Mouse visual cortex is modulated by distance traveled and by theta oscillations. *Current Biology*, *30*(19), 3811–3817.e6.
- França, A. S. C., do Nascimento, G. C., Lopes-dos-Santos, V., Muratori, L., Ribeiro, S., Lobão-Soares, B., & Tort, A. B. L. (2014). Beta2 oscillations (23–30 hz) in the mouse hippocampus during novel object recognition. *European Journal of Neuroscience*, *40*(11), 3693–3703.
- Frey, M., Tanni, S., Perrodin, C., O’Leary, A., Nau, M., Kelly, J., Banino, A., Doeller, C. F., & Barry, C. (2019). Deepinsight: A general framework for interpreting wide-band neural activity. *bioRxiv*, 871848.
- Fries, P. (2005). A mechanism for cognitive dynamics: Neuronal communication through neuronal coherence. *Trends Cogn Sci*, *9*(10), 474–80.
- Fries, P. (2015). Rhythms for cognition: Communication through coherence. *Neuron*, *88*(1), 220–235.

- Friston, K. J. (2011). Functional and effective connectivity: A review. *Brain Connectivity*, *1*(1), 13–36.
- Fröhlich, F., & McCormick, D. A. (2010). Endogenous electric fields may guide neocortical network activity. *Neuron*, *67*(1), 129–143.
- Gallego, J. A., Perich, M. G., Chowdhury, R. H., Solla, S. A., & Miller, L. E. (2018). A stable, long-term cortical signature underlying consistent behavior. *bioRxiv*, 447441.
- Gallego-Carracedo, C., Perich, M. G., Chowdhury, R. H., Miller, L. E., & Gallego, J. Á. (2022). Local field potentials reflect cortical population dynamics in a region-specific and frequency-dependent manner (S. Ostojski & R. L. Calabrese, Eds.). *eLife*, *11*, e73155.
- Geisler, C., Robbe, D., Zugaro, M., Sirota, A., & Buzsáki, G. (2007). Hippocampal place cell assemblies are speed-controlled oscillators. *Proceedings of the National Academy of Sciences*, *104*(19), 8149–8154.
- Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science*, *233*(4771), 1416–1419.
- Georgopoulos, A. P. (2014). Cell directional spread determines accuracy, precision, and length of the neuronal population vector. *Experimental Brain Research*, *232*(7), 2391–2405.
- Gerstein, G. L. (2013). Joint peri stimulus time histogram (JPSTH). In D. Jaeger & R. Jung (Eds.), *Encyclopedia of computational neuroscience* (pp. 1–4). Springer New York.
- Glaser, J. I., Benjamin, A. S., Chowdhury, R. H., Perich, M. G., Miller, L. E., & Kording, K. P. (2020). Machine learning for neural decoding. *arXiv:1708.00909 [cs, q-bio, stat]*.
- Gokcen, E., Jasper, A. I., Semedo, J. D., Zandvakili, A., Kohn, A., Machens, C. K., & Yu, B. M. (2021, September 1). *Disentangling the flow of signals between populations of neurons* (preprint). Neuroscience.
- Gong, Y., Huang, C., Li, J. Z., Grewe, B. F., Zhang, Y., Eismann, S., & Schnitzer, M. J. (2015). High-speed recording of neural spikes in awake mice and flies with a fluorescent voltage sensor. *Neuronal Dynamics*, *7*.
- Granger, C. (2003, December 8). *Granger nobel prize speech*. Retrieved February 3, 2020, from <https://www.nobelprize.org/uploads/2018/06/granger-lecture.pdf>
- Grienberger, C., & Konnerth, A. (2012). Imaging calcium in neurons. *Neuron*, *73*(5), 862–885.
- Grossberg, S. (2009). Beta oscillations and hippocampal place cell learning during exploration of novel environments. *Hippocampus*, *19*(9), 881–885.
- Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, *436*(7052), 801–806.

- Han, K.-S., Guo, C., Chen, C. H., Witter, L., Osorno, T., & Regehr, W. G. (2018). Ephaptic Coupling Promotes Synchronous Firing of Cerebellar Purkinje Cells. *Neuron*, *100*(3), 564–578.e3.
- Han, S., Chu, J.-U., Kim, H., Park, J. W., & Youn, I. (2017). Multiunit activity-based real-time limb-state estimation from dorsal root ganglion recordings. *Scientific Reports*, *7*(1), 44197.
- Harmony, T. (2013). The functional significance of delta oscillations in cognitive processing. *Front Integr Neurosci*, *7*, 83.
- Hebb, D. O. (1949). *The organization of behavior; a neuropsychological theory*. Wiley.
- Herreras, O. (2016). Local field potentials: Myths and misunderstandings. *Frontiers in Neural Circuits*, *10*.
- Huang, L., Ledochowitsch, P., Knoblich, U., Lecoq, J., Murphy, G. J., Reid, C., de Vries, S. E. J., Koch, C., Zeng, H., Buice, M. A., Waters, J., & Li, L. (2021). Relationship between simultaneously recorded spiking activity and fluorescence signal in GCaMP6 transgenic mice (G. L. Westbrook, Ed.). *eLife*, *10*, e51675.
- Hülsemann, M. J., Naumann, E., & Rasch, B. (2019). Quantification of phase-amplitude coupling in neuronal oscillations: Comparison of phase-locking value, mean vector length, modulation index, and generalized-linear-modeling-cross-frequency-coupling. *Frontiers in Neuroscience*, *13*.
- Humphries, M. D. (2017). Dynamical networks: Finding, measuring, and tracking neural population activity using network science. *Network Neuroscience*, *1*(4), 324–338.
- Humphries, M. D. (2021a, March 9). *The spike*.
- Humphries, M. D. (2021b). Strong and weak principles of neural dimension reduction. *arXiv:2011.08088 [q-bio]*.
- Ifft, P. J., Shokur, S., Li, Z., Lebedev, M. A., & Nicolelis, M. A. L. (2013). A brain-machine interface enables bimanual arm movements in monkeys. *Science Translational Medicine*, *5*(210), 210ra154–210ra154.
- Ioannidis, J. P. A. (2005). Why most published research findings are false. *PLOS Medicine*, *2*(8), e124.
- Ito, T., Brincat, S. L., Siegel, M., Mill, R. D., He, B. J., Miller, E. K., Rotstein, H. G., & Cole, M. W. (2019). Task-evoked activity quenches neural correlations and variability in large-scale brain systems. *bioRxiv*, 560730.
- Iwasaki, S., Sasaki, T., & Ikegaya, Y. (2021). Hippocampal beta oscillations predict mouse object-location associative memory performance. *Hippocampus*, *31*(5), 503–511.
- Jefferys, J. G. (1995). Nonsynaptic modulation of neuronal activity in the brain: Electric currents and extracellular ions. *Physiological Reviews*, *75*(4), 689–723.

- Jiang, X., Bian, G.-B., & Tian, Z. (2019). Removal of artifacts from EEG signals: A review. *Sensors (Basel, Switzerland)*, *19*(5).
- Jollans, L., Boyle, R., Artiges, E., Banaschewski, T., Desrivières, S., Grigis, A., Martinot, J.-L., Paus, T., Smolka, M. N., Walter, H., Schumann, G., Garavan, H., & Whelan, R. (2019). Quantifying performance of machine learning methods for neuroimaging data. *NeuroImage*, *199*, 351–365.
- Jonas, E., & Kording, K. P. (2017). Could a neuroscientist understand a microprocessor? *PLOS Computational Biology*, *13*(1), e1005268.
- Ju, H., & Bassett, D. S. (2020). Dynamic representations in networked neural systems. *Nature Neuroscience*, *23*(8), 908–917.
- 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., . . . Harris, T. D. (2017). Fully integrated silicon probes for high-density recording of neural activity. *Nature*, *551*(7679), 232–236.
- Kajikawa, Y., & Schroeder, C. E. (2011). How local is the local field potential? *Neuron*, *72*(5), 847–858.
- Kajimura, S., Margulies, D., & Smallwood, J. (2023). Frequency-specific brain network architecture in resting-state fMRI. *Scientific Reports*, *13*(1), 2964.
- Kay, K., Sosa, M., Chung, J. E., Karlsson, M. P., Larkin, M. C., & Frank, L. M. (2016). A hippocampal network for spatial coding during immobility and sleep. *Nature*, *531*(7593), 185–190.
- Kim, S., Putrino, D., Ghosh, S., & Brown, E. N. (2011). A granger causality measure for point process models of ensemble neural spiking activity. *PLOS Computational Biology*, *7*(3), e1001110.
- Kohn, A., Coen-Cagli, R., Kanitscheider, I., & Pouget, A. (2016). Correlations and neuronal population information. *Annual Review of Neuroscience*, *39*, 237–256.
- Koppe, G., Toutounji, H., Kirsch, P., Lis, S., & Durstewitz, D. (2019). Identifying nonlinear dynamical systems via generative recurrent neural networks with applications to fMRI. *PLOS Computational Biology*, *15*(8), e1007263.
- Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., MacIver, M. A., & Poeppel, D. (2017). Neuroscience needs behavior: Correcting a reductionist bias. *Neuron*, *93*(3), 480–490.
- Kriegeskorte, N., & Douglas, P. K. (2019). Interpreting encoding and decoding models. *Current Opinion in Neurobiology*, *55*, 167–179.
- Larry, N., & Joshua, M. (2021, April 30). *Neural correlations are attenuated in the transformation from the frontal cortex to movement* (preprint). Neuroscience.

- Li, X., Yao, X., Fox, J., & Jefferys, J. G. (2007). Interaction dynamics of neuronal oscillations analysed using wavelet transforms. *Journal of Neuroscience Methods*, *160*(1), 178–185.
- Lindén, H., Tetzlaff, T., Potjans, T. C., Pettersen, K. H., Grün, S., Diesmann, M., & Einevoll, G. T. (2011). Modeling the spatial reach of the LFP. *Neuron*, *72*(5), 859–872.
- Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. *Nature*, *453*(7197), 869–878.
- Long, S., Ding, R., Wang, J., Yu, Y., Lu, J., & Yao, D. (2021). Sleep quality and electroencephalogram delta power. *Frontiers in Neuroscience*, *15*.
- Lopes da Silva, F. (1991). Neural mechanisms underlying brain waves: From neural membranes to networks. *Electroencephalography and Clinical Neurophysiology*, *79*(2), 81–93.
- Machado, T. A., Kauvar, I. V., & Deisseroth, K. (2022). Multiregion neuronal activity: The forest and the trees. *Nature Reviews Neuroscience*, *23*(11), 683–704.
- Macpherson, T., Churchland, A., Sejnowski, T., DiCarlo, J., Kamitani, Y., Takahashi, H., & Hikida, T. (2021). Natural and artificial intelligence: A brief introduction to the interplay between AI and neuroscience research. *Neural Networks*, *144*, 603–613.
- Martinez-Banaclocha, M. (2018). Ephaptic coupling of cortical neurons: Possible contribution of astroglial magnetic fields? *Neuroscience*, *370*, 37–45.
- Márton, C. D., Fukushima, M., Camalier, C. R., Schultz, S. R., & Averbach, B. B. (2019). Signature patterns for top-down and bottom-up information processing via cross-frequency coupling in macaque auditory cortex. *eNeuro*, ENEURO.0467–18.2019.
- Mastrogiuseppe, F., & Ostojic, S. (2018). Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. *Neuron*, *99*(3), 609–623.e29.
- Mehonic, A., & Kenyon, A. J. (2022). Brain-inspired computing needs a master plan. *Nature*, *604*(7905), 255–260.
- Montgomery, S. M., & Buzsáki, G. (2007). Gamma oscillations dynamically couple hippocampal CA3 and CA1 regions during memory task performance. *Proceedings of the National Academy of Sciences*, *104*(36), 14495–14500.
- Munafò, M. R., Nosek, B. A., Bishop, D. V. M., Button, K. S., Chambers, C. D., Percie du Sert, N., Simonsohn, U., Wagenmakers, E.-J., Ware, J. J., & Ioannidis, J. P. A. (2017). A manifesto for reproducible science. *Nature Human Behaviour*, *1*(1), 0021.
- Musk, E., & Neuralink. (2019). An integrated brain-machine interface platform with thousands of channels. *bioRxiv*, 703801.

- Mysin, I., & Shubina, L. (2022). From mechanisms to functions: The role of theta and gamma coherence in the intrahippocampal circuits. *Hippocampus*, *32*(5), 342–358.
- Nicolelis, M. A. L., Fanselow, E. E., & Ghazanfar, A. A. (1997). Hebb’s dream: The resurgence of cell assemblies. *Neuron*, *19*(2), 219–221.
- Nicolelis, M. A. (Ed.). (2008). *Methods for neural ensemble recordings* (2nd). CRC Press/Taylor & Francis.
- Noguchi, A., Ikegaya, Y., & Matsumoto, N. (2021). In vivo whole-cell patch-clamp methods: Recent technical progress and future perspectives. *Sensors (Basel, Switzerland)*, *21*(4), 1448.
- Nonaka, S., Majima, K., Aoki, S. C., & Kamitani, Y. (2020, September 1). *Brain hierarchy score: Which deep neural networks are hierarchically brain-like?* (SSRN Scholarly Paper No. ID 3664362). Social Science Research Network. Rochester, NY.
- O’Doherty, J. E., Lebedev, M. A., Ifft, P. J., Zhuang, K. Z., Shokur, S., Bleuler, H., & Nicolelis, M. A. L. (2011). Active tactile exploration using a brain–machine–brain interface. *Nature*, *479*(7372), 228–231.
- Ogawa, S., Lee, T. M., Kay, A. R., & Tank, D. W. (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences of the United States of America*, *87*(24), 9868–9872.
- O’Keefe, J., & 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.
- O’Keefe, J., & Recce, M. L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, *3*(3), 317–330.
- Pachitariu, M., Steinmetz, N., Kadir, S., Carandini, M., & D, H. K. (2016). Kilo-sort: Realtime spike-sorting for extracellular electrophysiology with hundreds of channels. *bioRxiv*, 061481.
- Pachitariu, M., Stringer, C., Dipoppa, M., Schröder, S., Rossi, L. F., Dalgleish, H., Carandini, M., & Harris, K. D. (2017). Suite2p: Beyond 10,000 neurons with standard two-photon microscopy. *bioRxiv*, 061507.
- Pandarínath, C., O’Shea, D. J., Collins, J., Jozefowicz, R., Stavisky, S. D., Kao, J. C., Trautmann, E. M., Kaufman, M. T., Ryu, S. I., Hochberg, L. R., Henderson, J. M., Shenoy, K. V., Abbott, L. F., & Sussillo, D. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods*, *15*(10), 805–815.
- Pang, R., Lansdell, B. J., & Fairhall, A. L. (2016). Dimensionality reduction in neuroscience. *Current Biology*, *26*(14), R656–R660.

- Panzeri, S., Schultz, S. R., Treves, A., & Rolls, E. T. (1999). Correlations and the encoding of information in the nervous system. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, *266*(1423), 1001–1012.
- Perich, M. G., Gallego, J. A., & Miller, L. E. (2018). A neural population mechanism for rapid learning. *Neuron*, *100*(4), 964–976.e7.
- Pesaran, B., Vinck, M., Einevoll, G. T., Sirota, A., Fries, P., Siegel, M., Truccolo, W., Schroeder, C. E., & Srinivasan, R. (2018). Investigating large-scale brain dynamics using field potential recordings: Analysis and interpretation. *Nature Neuroscience*, *21*(7), 903–919.
- Pfeiffer, M., & Pfeil, T. (2018). Deep learning with spiking neurons: Opportunities and challenges. *Frontiers in Neuroscience*, *12*, 774.
- Pinotsis, D. A., Fridman, G., & Miller, E. K. (2023). Cytoelectric coupling: Electric fields sculpt neural activity and “tune” the brain’s infrastructure. *Progress in Neurobiology*, *226*, 102465.
- Raichle, M. E. (2015). The brain’s default mode network. *Annual Review of Neuroscience*, *38*(1), 433–447.
- Richard, G. R., Titiz, A., Tyler, A., Holmes, G. L., Scott, R. C., & Lenck-Santini, P.-P. (2013). Speed modulation of hippocampal theta frequency correlates with spatial memory performance. *Hippocampus*, *23*(12), 1269–1279.
- Richards, B. A., Lillicrap, T. P., Beaudoin, P., Bengio, Y., Bogacz, R., Christensen, A., Clopath, C., Costa, R. P., de Berker, A., Ganguli, S., Gillon, C. J., Hafner, D., Kepecs, A., Kriegeskorte, N., Latham, P., Lindsay, G. W., Miller, K. D., Naud, R., Pack, C. C., . . . Kording, K. P. (2019). A deep learning framework for neuroscience. *Nature Neuroscience*, *22*(11), 1761–1770.
- Rizzello, E., Martin, S. K., Rouine, J., Callaghan, C., Mathiasen, M. L., & O’Mara, S. M. (2022). Place cells in the claustrum remap under NMDA receptor control. *European Journal of Neuroscience*.
- Roth, B. L. (2016). DREADDs for neuroscientists. *Neuron*, *89*(4), 683–694.
- Ruff, D. A., & Cohen, M. R. (2016). Attention increases spike count correlations between visual cortical areas. *Journal of Neuroscience*, *36*(28), 7523–7534.
- Salazar, R. F., White, W., Lacroix, L., Feldon, J., & White, I. M. (2004). NMDA lesions in the medial prefrontal cortex impair the ability to inhibit responses during reversal of a simple spatial discrimination. *Behavioural Brain Research*, *152*(2), 413–424.
- Saxena, S., & Cunningham, J. P. (2019). Towards the neural population doctrine. *Current Opinion in Neurobiology*, *55*, 103–111.
- Scheffer-Teixeira, R., & Tort, A. B. (2016). On cross-frequency phase-phase coupling between theta and gamma oscillations in the hippocampus (F. K. Skinner, Ed.). *eLife*, *5*, e20515.

- Schneider, M., Dann, B., Sheshadri, S., Scherberger, H., & Vinck, M. (2020, August 10). *A general theory of coherence between brain areas*.
- Schultheiss, N. W., Schlecht, M., Jayachandran, M., Brooks, D. R., McGlothan, J. L., Guilarte, T. R., & Allen, T. A. (2020). Awake delta and theta-rhythmic hippocampal network modes during intermittent locomotor behaviors in the rat. *Behavioral Neuroscience*, *134*(6), 529.
- Semedo, J. D., Gokcen, E., Machens, C. K., Kohn, A., & Yu, B. M. (2020). Statistical methods for dissecting interactions between brain areas. *Current Opinion in Neurobiology*, *65*, 59–69.
- Semedo, J. D., Jasper, A. I., Zandvakili, A., Krishna, A., Aschner, A., Machens, C. K., Kohn, A., & Yu, B. M. (2022). Feedforward and feedback interactions between visual cortical areas use different population activity patterns. *Nature Communications*, *13*(1), 1099.
- Semedo, J. D., Zandvakili, A., Kohn, A., Machens, C. K., & Yu, B. M. (2014). Extracting latent structure from multiple interacting neural populations. *Advances in Neural Information Processing Systems*, *9*.
- Semedo, J. D., Zandvakili, A., Machens, C. K., Yu, B. M., & Kohn, A. (2019). Cortical areas interact through a communication subspace. *Neuron*, *102*(1), 249–259.e4.
- Shen, G., Dwivedi, K., Majima, K., Horikawa, T., & Kamitani, Y. (2019). End-to-end deep image reconstruction from human brain activity. *Frontiers in Computational Neuroscience*, *13*, 21.
- Shoham, S., O’Connor, D. H., & Segev, R. (2006). How silent is the brain: Is there a “dark matter” problem in neuroscience? *Journal of Comparative Physiology A*, *192*(8), 777–784.
- Siegle, J. H., Jia, X., Durand, S., Gale, S., Bennett, C., Graddis, N., Heller, G., Ramirez, T. K., Choi, H., Luviano, J. A., Groblewski, P. A., Ahmed, R., Arkhipov, A., Bernard, A., Billeh, Y. N., Brown, D., Buice, M. A., Cain, N., Caldejon, S., . . . Koch, C. (2019). A survey of spiking activity reveals a functional hierarchy of mouse corticothalamic visual areas. *bioRxiv*, 805010.
- Siegle, J. H., Jia, X., Durand, S., Gale, S., Bennett, C., Graddis, N., Heller, G., Ramirez, T. K., Choi, H., Luviano, J. A., Groblewski, P. A., Ahmed, R., Arkhipov, A., Bernard, A., Billeh, Y. N., Brown, D., Buice, M. A., Cain, N., Caldejon, S., . . . Koch, C. (2021). Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature*, *592*(7852), 86–92.
- Sirota, A., Montgomery, S., Fujisawa, S., Isomura, Y., Zugaro, M., & Buzsaki, G. (2008). Entrainment of neocortical neurons and gamma oscillations by the hippocampal theta rhythm. *Neuron*, *60*(4), 683–97.
- Skaggs, W. E., McNaughton, B. L., & Gothard, K. M. (1993). An information-theoretic approach to deciphering the hippocampal code. In S. J. Hanson, J. D. Cowan,

- & C. L. Giles (Eds.), *Advances in neural information processing systems* 5 (pp. 1030–1037). Morgan-Kaufmann.
- Smith, E., Kellis, S., House, P., & Greger, B. (2013). Decoding stimulus identity from multi-unit activity and local field potentials along the ventral auditory stream in the awake primate: Implications for cortical neural prostheses. *Journal of neural engineering*, *10*(1), 016010.
- Solomon, E. A., Stein, J. M., Das, S., Gorniak, R., Sperling, M. R., Worrell, G., Inman, C. S., Tan, R. J., Jobst, B. C., Rizzuto, D. S., & Kahana, M. J. (2019). Dynamic theta networks in the human medial temporal lobe support episodic memory. *Current Biology*, *29*(7), 1100–1111.e4.
- Soltanian-Zadeh, S., Sahingur, K., Blau, S., Gong, Y., & Farsiu, S. (2019). Fast and robust active neuron segmentation in two-photon calcium imaging using spatiotemporal deep learning. *Proceedings of the National Academy of Sciences*, *116*(17), 8554–8563.
- 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., ... Harris, T. D. (2021). Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science*, *372*(6539).
- Steinmetz, N. A., Zatka-Haas, P., Carandini, M., & Harris, K. D. (2019). Distributed coding of choice, action and engagement across the mouse brain. *Nature*, *576*(7786), 266–273.
- Stevenson, I. H., & Kording, K. P. (2011). How advances in neural recording affect data analysis. *Nature Neuroscience*, *14*(2), 139–142.
- Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C. B., Carandini, M., & Harris, K. D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, *364*(6437), eaav7893.
- Su, C.-Y., Menuz, K., Reiser, J., & Carlson, J. R. (2012). Non-synaptic inhibition between grouped neurons in an olfactory circuit. *Nature*, *492*(7427), 66–71.
- Sussillo, D., Churchland, M. M., Kaufman, M. T., & Shenoy, K. V. (2015). A neural network that finds a naturalistic solution for the production of muscle activity. *Nature Neuroscience*, *18*(7), 1025–1033.
- Taherkhani, A., Belatreche, A., Li, Y., Cosma, G., Maguire, L. P., & McGinnity, T. M. (2020). A review of learning in biologically plausible spiking neural networks. *Neural Networks*, *122*, 253–272.
- Tingley, D. (2021). *Network computation across brain states* [Doctoral dissertation, New York University].
- Tlaie, A., Shapcott, K. A., Tiesinga, P., Schölvinck, M., & Havenith, M. N. (2021, November 29). *Does the brain care about averages? a simple test*.

- Trautmann, E. M., Stavisky, S. D., Lahiri, S., Ames, K. C., Kaufman, M. T., O'Shea, D. J., Vyas, S., Sun, X., Ryu, S. I., Ganguli, S., & Shenoy, K. V. (2019). Accurate estimation of neural population dynamics without spike sorting. *Neuron*, *103*(2), 292–308.e4.
- Trimper, J. B., Galloway, C. R., Jones, A. C., Mandi, K., & Manns, J. R. (2017). Gamma oscillations in rat hippocampal subregions dentate gyrus, CA3, CA1, and subiculum underlie associative memory encoding. *Cell Reports*, *21*(9), 2419–2432.
- Umakantha, A., Morina, R., Cowley, B. R., Snyder, A. C., Smith, M. A., & Yu, B. M. (2021). Bridging neuronal correlations and dimensionality reduction. *Neuron*, *109*(17), 2740–2754.e12.
- Urai, A. E., Doiron, B., Leifer, A. M., & Churchland, A. K. (2021). Large-scale neural recordings call for new insights to link brain and behavior, 24.
- Viejo, G., Cortier, T., & Peyrache, A. (2018). Brain-state invariant thalamo-cortical coordination revealed by non-linear encoders. *PLoS computational biology*, *14*(3), e1006041.
- Vinck, M., Uran, C., Spyropoulos, G., Onorato, I., Broggin, A. C., Schneider, M., & Canales-Johnson, A. (2023). Principles of large-scale neural interactions. *Neuron*, *111*(7), 987–1002.
- Vyas, S., Golub, M. D., Sussillo, D., & Shenoy, K. V. (2020). Computation through neural population dynamics. *Annual Review of Neuroscience*, *43*(1), 249–275.
- Wei, Z., Lin, B.-J., Chen, T.-W., Daie, K., Svoboda, K., & Druckmann, S. (2020). A comparison of neuronal population dynamics measured with calcium imaging and electrophysiology. *PLOS Computational Biology*, *16*(9), e1008198.
- Westlin, C., Theriault, J. E., Katsumi, Y., Nieto-Castanon, A., Kucyi, A., Ruf, S. F., Brown, S. M., Pavel, M., Erdogmus, D., Brooks, D. H., Quigley, K. S., Whitfield-Gabrieli, S., & Barrett, L. F. (2023). Improving the study of brain-behavior relationships by revisiting basic assumptions. *Trends in Cognitive Sciences*, *27*(3), 246–257.
- Whelan, R., & Garavan, H. (2014). When optimism hurts: Inflated predictions in psychiatric neuroimaging. *Biological Psychiatry*, *75*(9), 746–748.
- Widmann, A., & Schröger, E. (2012). Filter effects and filter artifacts in the analysis of electrophysiological data. *Frontiers in Psychology*, *3*.
- Williams, A. H., & Linderman, S. W. (2021). Statistical neuroscience in the single trial limit. *arXiv:2103.05075 [q-bio, stat]*.
- Wirt, R. A., & Hyman, J. M. (2019). ACC theta improves hippocampal contextual processing during remote recall. *Cell Reports*, *27*(8), 2313–2327.e4.

- Womelsdorf, T., Vinck, M., Leung, S., & Everling, S. (2010). Selective theta-synchronization of choice-relevant information subserves goal-directed behavior. *Frontiers in Human Neuroscience*, *4*.
- Yang, W., Carrillo-Reid, L., Bando, Y., Peterka, D. S., & Yuste, R. (2018). Simultaneous two-photon imaging and two-photon optogenetics of cortical circuits in three dimensions (K. Svoboda, Ed.). *eLife*, *7*, e32671.
- Yoshida, E., Terada, S.-I., Tanaka, Y. H., Kobayashi, K., Ohkura, M., Nakai, J., & Matsuzaki, M. (2018). In vivo wide-field calcium imaging of mouse thalamocortical synapses with an 8 k ultra-high-definition camera. *Scientific Reports*, *8*(1), 8324.
- Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. V., & Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Journal of Neurophysiology*, *102*(1), 614–635.
- Yuste, R. (2015). From the neuron doctrine to neural networks. *Nature Reviews Neuroscience*, *16*(8), 487–497.
- Zhu, F., Grier, H. A., Tandon, R., Cai, C., Giovannucci, A., Kaufman, M. T., & Pandarinath, C. (2021, November 21). *A deep learning framework for inference of single-trial neural population activity from calcium imaging with sub-frame temporal resolution*.

3 Multi-region neural data analysis and interoperability with Simuran

Introduction: Recent advances in neural recording technologies, computing power, and experimental methods have made it possible to simultaneously investigate neural recordings from multiple brain regions. However, the availability and continued development of numerous software tools and data formats have presented a challenge for neuroscientists analysing this type of data.

Methods: To address this challenge, we developed Simuran, a Python library that simplifies the process of analysing large neural data recordings and facilitates interfacing with other analysis software. Simuran supports a variety of data formats and sources, including Axona, Neurodata Without Borders, the Allen Institute, and the International Brain Laboratory, in addition to neuroscience analysis libraries such as NeuroChaT, MNE, Elephant, and Neo. Simuran also provides a node-based user interface to streamline the analysis process for inexperienced coders.

Results: We have successfully used Simuran to analyse data from multiple brain regions simultaneously, including handling large experiments and analysing open-source datasets. Simuran is used to convert Axona experiments to common format and analyse spiking activity and local field potentials from the subiculum, retrosplenial cortex, and CA1. We also demonstrate using Simuran to analyse data from the Allen Institute and the International Brain Laboratory.

Conclusions: In summary, we have developed a flexible tool that enables the analysis of neural data from multiple brain regions simultaneously. Simuran is modular and designed to support new tools and data formats, making it a valuable resource for the neuroscience community. We close by discussing future directions for Simuran, and the promise such a tool holds for the analysis of large-scale neural data.

3.1 Challenges in large-scale neural data analysis

A goal for many neuroscientists is to elucidate how the orchestration between multiple neurons enables complex processes and behaviours, such as memory, consciousness, and decision-making. Connected with this goal, neural recording technologies have rapidly improved, with advances in silicon probes, calcium imaging, optogenetics, designer receptors exclusively activated by designer drugs (DREADDs), and anatomical reconstruction of neural connections (Deisseroth, 2015; Gao et al., 2022; Grienberger & Konnerth, 2012; Roth, 2016; Steinmetz et al., 2021; Yang et al., 2018). Large-scale datasets which obtain neural activity in multiple brain regions in the living participant have been captured in humans and other animals (e.g., Paulk et al., 2022; Siegle et al., 2021; Stringer et al., 2019; The International Brain Laboratory et al., 2022). These offer exciting opportunities to relate neural responses across multiple regions to complex behaviours (Machado et al., 2022; Saxena & Cunningham, 2019; Urai et al., 2021).

Despite steady progress in drawing insights from such large-scale neural data (Allen et al., 2019; Carrillo-Reid et al., 2019; Kaefer et al., 2020; Semedo et al., 2019, 2022; Stringer et al., 2019) there remain significant conceptual difficulties (K. D. Harris, 2020; Jonas & Kording, 2017; Jones & Kording, 2021; Stevenson & Kording, 2011) and technical issues in handling and analysing from these huge datasets. To solve these issues, there have been efforts across multiple intermingled strands. One core effort involves standardising data sharing formats with Neurodata Without Borders (NWB), open neurophysiology environment (ONE), Nix, and Neo (Adrian et al., 2014; Garcia et al., 2014; Rübél et al., 2022) and tools for storing, viewing, and retrieving these datasets with DANDI, GIN, Alyx and Allen SDK (Rübél et al., 2022; Siegle et al., 2019; The International Brain Laboratory et al., 2022). Another line of progress has been to provide neuroscience targeted analysis methods focused on dimensionality reduction, including the latent factor analysis via dynamical systems (LFADS), delayed latents across groups (DLAG), and gaussian process factor analysis (GPFA) approaches (Cunningham & Yu, 2014; Gokcen et al., 2021; Pandarinath et al., 2018; Pang et al., 2016; Yu et al., 2009).

Tangential to this, general-purpose Python libraries provide analysis methods which are highly relevant for dimensionality reduction, machine learning, and signal processing such as SciPy, scikit-learn, NumPy, and pandas (C. R. Harris et al., 2020; McKinney, 2010; Pedregosa et al., 2011; Virtanen et al., 2020). Additionally, toolbox style packages have been created which aim to provide an all-in-one solution to specific neuroscience challenges, including Pynapple, NeuroChaT, Elephant, MNE, and Ghostipy (Chu & Kemere, 2021; Denker et al., 2018; Gramfort et al., 2013; Islam et al., 2019; Viejo

et al., 2022). Finally, keeping track of all this experimental data and with flexibility to scale analyses from investigating some tens of neurons on a single lab computer to thousands of neurons on computing clusters causes huge difficulty, which has led to the development of experiment management tools and version control tools, including Snakemake, Sumatra, and Nipype (Davison, 2012; Gorgolewski et al., 2011; Mölder et al., 2021). This overview leaves out many packages, and other branches of tools, such as spiking neural network simulators and modelling of LFPs and neurons.

Creating a fully-fledged analysis in this space requires multiple moving parts and interaction between different packages, leading to increased risk of mistakes, time burden on the programmer, and difficulty in creating a reproducible environment. To tackle these issues, we developed Simuran, an open-source Python library that is designed specifically to enable researchers to analyse large-scale neural datasets. Researchers can use Simuran to overcome file-format incompatibilities, difficulties in software interoperability, and the burden of experiment management. Simuran is aimed primarily towards developers that can use Python to program such analyses but provides a general user interface (GUI) so that non-programmers can perform a subset of the analysis by linking up steps in the pipeline like a flowchart. We show the application of Simuran in analysing tetrode and behavioural data in our lab experiments and analysing openly shared large-scale neural datasets, demonstrating the advantages of our unified approach to neural data analysis.

3.2 An overview of Simuran

3.2.1 Primary goals of the software

Simuran was designed with three main goals:

1. Facilitate using different neuroscience tools and interacting with many data formats through a common framework.
2. Provide a visual node-based interface for program flow management.
3. Allow for fast iteration to readily support new tools via modern object-orientated software design practices.

The first goal is to enable common interfaces to neuroscience software to reduce the burden on programmers and increase benchmarking and reproducibility of work. The most readily comparable situation is in spike sorting, where the abundance of

spike sorting algorithms and different data sources (Pachitariu et al., 2016; Quiroga, 2012; Rossant et al., 2016) led to the necessity of a common entry point to the tools. SpikeInterface was developed to allow for a unified user interface to access spike sorting algorithms and data formats and SpikeForest (Magland et al., 2020) was created to maintain regular benchmarks of the performance of ten of the most popular spike sorters by wrapping them into a unified interface. Simuran aims to provide a similar interface for other neuroscience tools, such as dimensionality reduction, machine learning, and signal processing, and to provide a common interface to data formats.

A visual node-based interface provides an efficient and intuitive GUI for programs, such as Inviwo (Jönsson et al., 2020) for volume rendering, Bonsai (Lopes et al., 2015) for controlling neural hardware, and general-purpose Python visual scripting programs such as Ryven, and PyFlow. These are similar to creating flowcharts and are intuitive to the user and effective at displaying temporal dependencies. In addition, an efficient method of processing node-based visual editors is through directed acyclic graphs (DAGs) which can be linked to workflow management tools such as Snakemake (Mölder et al., 2021) which also process networks via DAGs. Providing a visual wrapper around many common tools along with installation instructions can facilitate a better understanding of the program execution and ease the barrier to entry in scientific programming.

Finally, the neuroscience data analysis field is rapidly evolving, coupled with the constantly increasing limits of data collection (Urai et al., 2021). As such, designing the code to only support a fixed set of data formats and tools would be ill-thought-out. Instead, our third goal is to follow modern object-oriented software design principles (which have built upon Gamma et al., 1995). For instance, the protocol pattern and factory pattern will allow us to provide flexible templates for supporting new software tools or data formats. This will ensure that the software is usable for several tasks related to neural data analysis.

3.2.2 Organisation and design of Simuran

Simuran is organised into components which fall into three categories: creating the foundation of the software, loading and saving data, and performing analysis on data (Figure 3.1). Using these components, the main steps involved to analyse a single session of neural recordings in Simuran are as follows:

1. **Data format specification:** The user creates a *Recording* object which will store all the data, metadata, and analysis results, and then associates a *Loader* object with that *Recording*, indicating the data format they wish to be loaded and how to load that data.

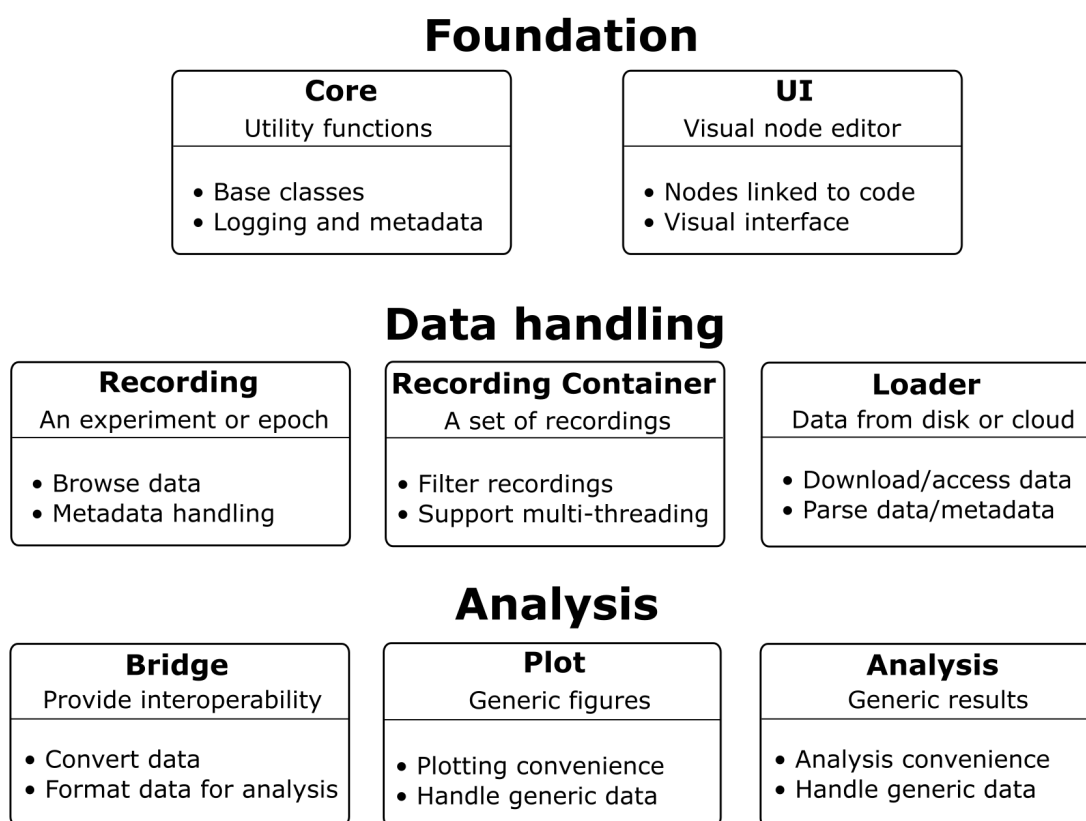


FIGURE 3.1: The primary components of Simuran’s software architecture.

Foundation: The foundation component is the base of the software and is comprised of the *Core* and *UI* modules to provide base classes and a visual interface.

Data handling: The data handling component is comprised of the *Recording* and *Recording container* and *Loader* classes. These provide methods to load data, save data, and browse data.

Analysis: The analysis component is comprised of the *Analysis*, *Bridge*, and *Plot* modules. These provide methods to convert between different data formats, run analyses with multi-processing and save analysis results.

2. **Metadata and data source specification:** The user provides data source information (e.g., the paths to the data files, or a URL to download data from) to the *Recording* object and any optional metadata.
3. **Data access and preprocessing:** The user calls the *load* method on the *Recording* to load the data into computer memory and perform any preprocessing.
4. **Data conversion, analysis, and plotting:** The user can then combine *Bridge* objects, which allow conversion between data formats, with the *Analysis* and *Plot* modules to operate on the data using many different software packages.

3.2.3 Providing a flowchart style user interface

Similar to a flow chart, the visual node editor in Simuran is a GUI that allows users to create a program flow by dragging and dropping nodes onto a canvas and linking them together (Figure 3.2). The nodes represent different actions that can be performed to load, convert, analyse, and plot data, while the connections between the nodes represent the order in which the actions are performed. After linking up the nodes and setting the parameters, the user can run the program flow by clicking the run button. The program flow can be saved (in JSON format) and loaded for later use, and the nodes can be rearranged and edited. In this way, such a visual node-based editor provides a simple and intuitive way to perform complex data analysis with little or no coding knowledge.

3.3 Applications of Simuran to neural data

Simuran was used throughout this thesis to aid with data analysis. As such, many applications can be seen in the results sections of this thesis and the corresponding shared code. In this section, we provide further details and examples of applying Simuran.

3.3.1 Analysing spikes and local field potentials from the rat limbic system

Simuran was used to analyse spikes and LFPs from the subiculum, hippocampus, and retrosplenial cortex in the rat. Rats were implanted with 32-channel tetrode arrays with varying placements of wires in the subiculum, hippocampus, and retrosplenial cortex (Table A.1) in addition to recording those rats performing numerous tasks in many different environments. Each recording session, such as a full free-foraging session, was

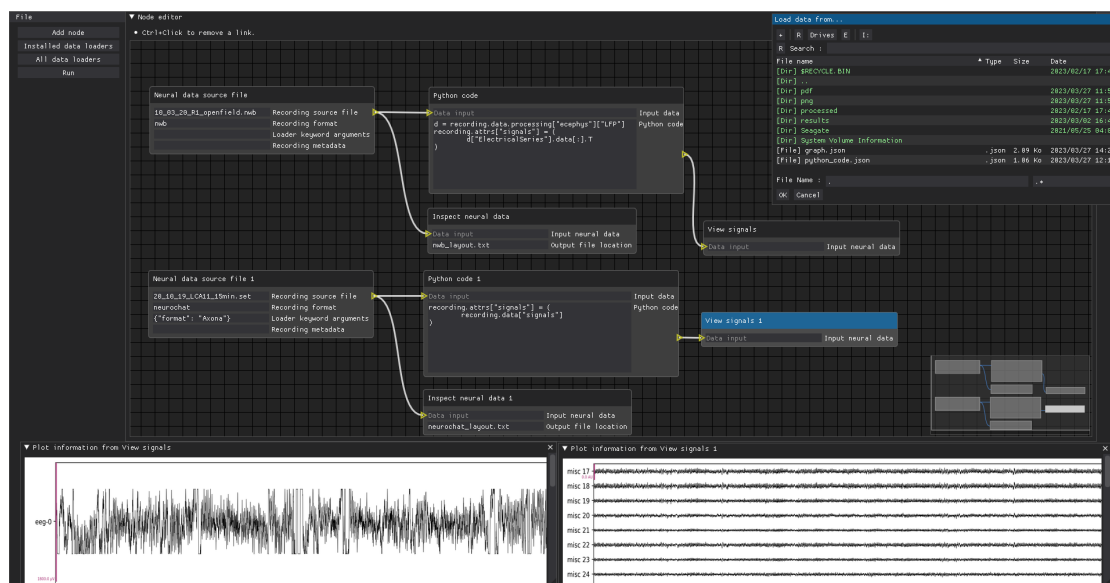


FIGURE 3.2: A flowchart in Simuran’s visual interface to load, process, and plot LFPs.

The example shows how to load data from various sources, and then plot the data for visual inspection. In this case, the data was loaded from the Axona file system using NeuroChaT and from NWB format in the neural data source file nodes. The files were selected using the window shown in the top right of the figure. The loaded data is then passed from these source nodes into recording inspection nodes (which show the data layout) and through a Python code block (which can perform any function). In this case, the Python code block is used to indicate which signals to plot in the GUI. The output of the Python code block is then passed to a plot node, which plots the data for visual inspection in the GUI. Shown at the bottom of the figure are the LFPs in the two datasets, viewed from within the GUI.



FIGURE 3.3: Combining Simuran and Snakemake to create an analysis pipeline with branching paths via directed graphs.

The graph shows the order of execution of the pipeline, and the nodes are the different steps. The nodes are labelled with the name of the node, indicating the operation performed, and edges show the dependencies between the nodes.

then loaded into a *Recording* object, which was then used to store the data, metadata, and analysis results. Using Simuran, we loaded the captured Axona data from the different recording layouts into a generic format by using *Loaders* to read linked mapping files. The data was then pre-processed and converted to the standardised NWB format for further processing and data sharing.

A table of each *Recording*'s metadata was created, to allow filtering on properties, such as, the maze type, rat identifier, and experimental condition (see Figure 3.2 for an example of loading an NWB file and Axona file in the node editor). With the data pre-processed and readily filterable, the *Bridges* in Simuran enabled conversion of the NWB data to analyse LFPs and spikes with a combination of tools, including NeuroChaT (Islam et al., 2019) MNE (Gramfort et al., 2013), and SciPy (Virtanen et al., 2020). Finally, Simuran was used in conjunction with Snakemake (Mölder et al., 2021) to perform batch processing of the data and handle the various configuration files for parameter setting (<https://github.com/seankmartin/atn-sub-lfp-workflow>). The full workflow for this pipeline is shown in Figure 3.3 and the results of this analysis are presented in chapter 4 and chapter 5.

3.3.2 Tracking cells across days in calcium imaging data

We demonstrate using Simuran to track cells across days in calcium imaging data from mice imaged in the Allen Institute Visual Behaviour Optical Physiology dataset (Allen Institute for Brain Science, 2021). Simuran was used to filter the data to only include two-photon imaging data that was captured along with behaviour in mice that respond to GCaMP6f where the imaging data contains at least one cell which was tracked across at least five imaging sessions. Then, Simuran was used to load this data across multiple cores and interface with matplotlib and MNE to plot the maximum intensity projection of the calcium imaging data and the response of the cell in relation to stimuli and behavioural state. The same excitatory cell captured from the primary visual cortex by multi-plane two-photon imaging was tracked and analysed across days (excitatory driver line *Slc17a7-IRES2-Cre*, reporter line GCaMP6; Madisen et al., 2015). In the window of data shown, the cell appears to respond to presentations of a particular image from the novel image set while the mouse is running, but not to images from the familiar image set or when the mouse is stationary (Figure 3.4).

3.3.3 Performing factor analysis on Allen Institute data

We describe using Simuran to compare dimensionality reduction with gaussian process factor analysis (GPFA) performed on calcium imaging data and Neuropixels data that are recorded from the same brain areas. In Simuran, this involves specifying the two data types to be loaded, loading the data, and then converting the data from their specific formats to generic NumPy arrays or lists following pre-processing. Finally, the data is converted from generic format to toolbox specific formats for analysis and plotting of spikes, LFPs, and fluorescence (see the full pipeline in Figure 3.5). To demonstrate Python code using Simuran, we implement a subset of the pipeline, loading visual cortex data from Allen Institute for Brain Science (AIS) Neuropixels recordings (Allen Institute for Brain Science, 2022) and performing GPFA (see results in chapter 7). Elephant (Denker et al., 2018) is used to perform the GPFA analysis and compared with factor analysis performed in scikit-learn (Pedregosa et al., 2011). Using Simuran, the required steps can be created with a relatively small amount of Python code (Figure 3.6):

1. Identify a subset of experimental sessions that contain data from the visual cortex. Then, for each experiment (optionally across multiple cores):
2. Download the corresponding visual cortex data in NWB format from the Allen Institute Visual Behaviour Neuropixels (VBN) dataset.

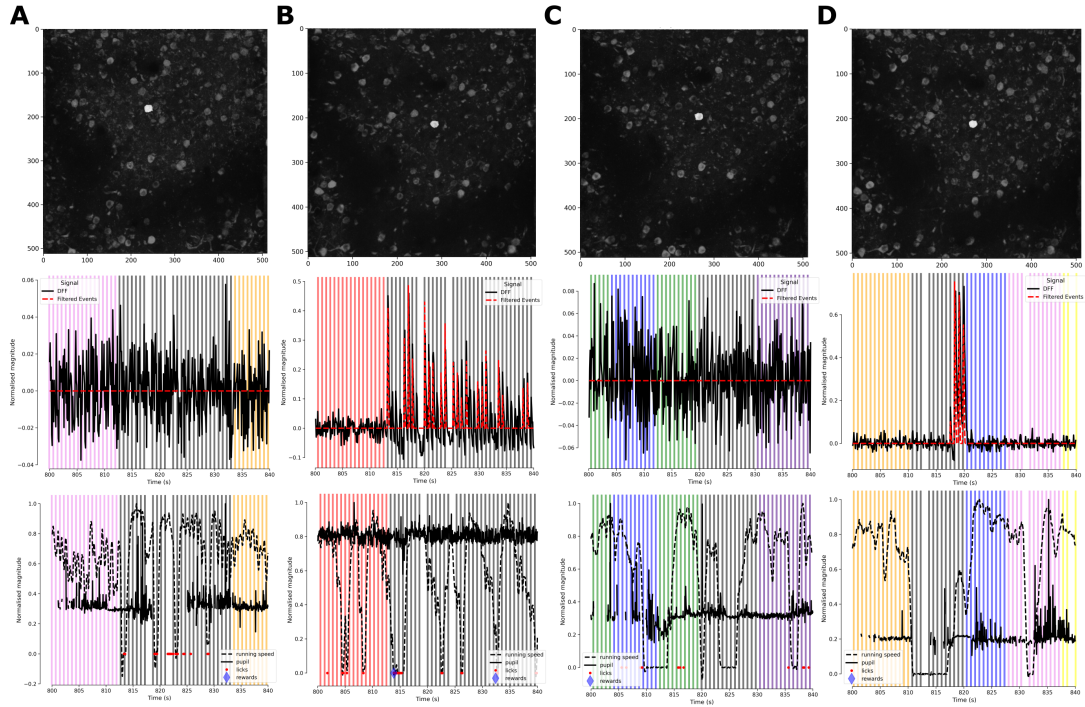


FIGURE 3.4: Using Simuran to track the same excitatory cell across days in calcium imaging data from Allen Brain Institute Visual Behaviour Optical Physiology experiment.

(A) The response of the cell during active go/no-go change detection task performance while the mouse was presented with the set of training images.

(B) The response of the cell during active go/no-go change detection task performance while the mouse was presented with the set of novel images that had only been seen twice before.

(C) The response of the cell during active go/no-go change detection task performance while the mouse was presented with the set of training images.

(D) The response of the cell during passive viewing of stimuli while the mouse was presented with the set of novel images that had only been seen four times before.

(Key) Top – maximum projection image of the imaging plane, with the tracked cell shown in white. Middle – the change in fluorescence of the cell over time and detected likely cell activations (filtered events), with stimulus onset and offset times shown by coloured vertical bars. Bottom – the normalised running speed, pupil activity, licks, and rewards, with stimulus times shown by vertical bars.

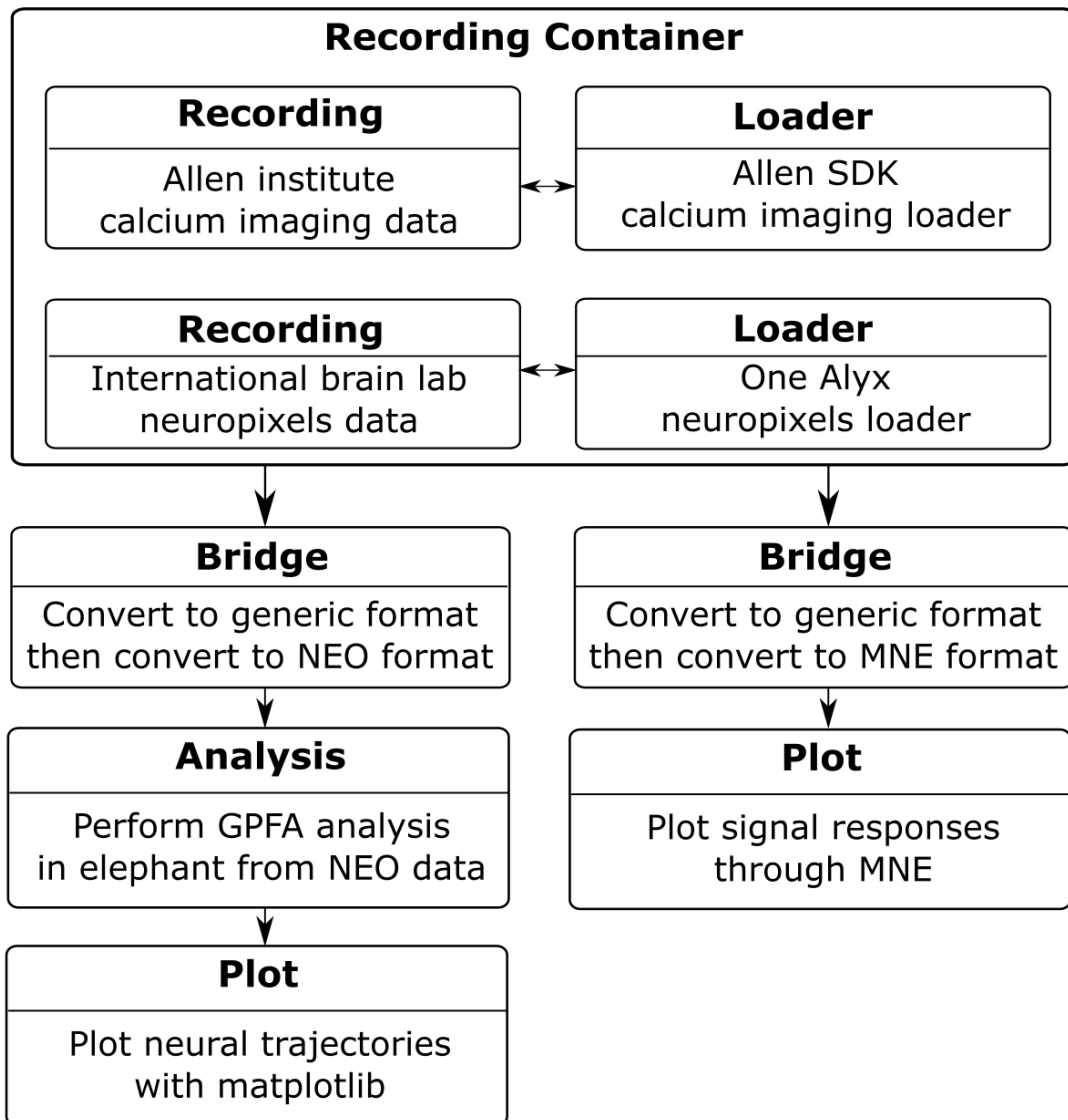


FIGURE 3.5: A pipeline in Simuran to analyse both calcium imaging data from the Allen Institute for Brain Science and Neuropixels data from the International Brain Laboratory through the same flow.

Data loading: The *RecordingContainer* is set up to contain the two distinct types of data that will be analysed. Then, the data is loaded using the AllenSDK calcium imaging *Loader* and the OneAlyx Neuropixels *Loader*.

GPFA analysis: The data is piped into a *Bridge* which converts the data into generic python objects (lists and NumPy arrays), and then into a *Bridge* which converts to NEO format that can be used by the GPFA analysis in Elephant. For Neuropixels data, this involves using the loaded spike sorting data and converting it to a NEO spike train. For calcium imaging data, this involves estimating spikes from the calcium imaging data and then converting this to a NEO spike train. The analysis is run using Elephant and the results are saved.

Signal plotting: The fluorescence signal and the LFP are plotted using MNE after converting to generic NumPy arrays using a *Bridge*, and then another *Bridge* to convert to MNE style data.

3. Load the downloaded data into memory as it is required.
4. Convert the spike data to generic Python datatypes for compatibility with scikit-learn, and then from generic data into the Neo format for compatibility with the elephant toolbox.
5. Compare the dimensionality reduction of the spiking data from the elephant toolbox against factor analysis from the scikit-learn toolbox.

3.4 Comparison to other software and conclusion

We have introduced Simuran, an open-source software package that provides a common interface to many different neuroscience related software packages, allowing for easy conversion between data formats and analysis tools. Due to numerous neuroscience data formats and analysis tools, working in this space can come with significant cognitive and technical burdens, creating difficulty in sharing data, results, and analysis pipelines between researchers. Simuran enables open and findable, accessible, interoperable, reusable (FAIR) neuroscience research by supporting the International Neuroinformatics Coordinating Facility (INCF) standard (Abrams et al., 2022) approved NWB 2.0 (Rübel et al., 2022) and NEO (Garcia et al., 2014) data formats. In addition to general format datasets, such as those on DANDI (DANDI team, 2022), Simuran supports using a subset of bespoke data storage solutions such as the AllenSDK and ONEAlyx (Siegle et al., 2019; The International Brain Laboratory et al., 2022). Furthermore, supporting toolboxes such as MNE and NeuroChaT facilitates access to multiple analysis methods and data formats (Gramfort et al., 2013; Islam et al., 2019), such as SpikeInterface (Buccino et al., 2019) and proprietary Axona Ltd data formats. Finally, the GUI flowchart-like visual node editor in Simuran provides an intuitive way to perform complex data analysis with little to no coding knowledge.

Simuran differs from other analysis software or pipelines (e.g., CalmAn, Minian, and Elephant – Denker et al., 2018; Dong et al., 2022; Giovannucci et al., 2019) as it is aimed towards interoperability of *in vivo* neuroscience analysis related tools. Here, we compare Simuran to other similar software tools aimed towards interoperability. Neuroconv uses a similar concept to the *Bridges* in Simuran to convert different data formats into NWB format via *DataInterface* objects, though the scope of Neuroconv is limited to converting data formats to NWB. Open Source Brain (Gleeson et al., 2019) provides improved accessibility to models of neural circuits by interacting with other modelling tools in addition to providing a free online platform to store models and perform analysis. This allows for flexible handling of neural circuit modelling. In addition, direct integration of

```

1 # Setup - import the necessary packages
2 import simuran, elephant, sklearn
3 from simuran.bridges.allen_vbn_bridge import AllenVBNBridge
4 from simuran.bridges.neo_bridge import NeoBridge
5 from simuran.analysis.unit import bin_spike_train
6
7 # Setup - create a loader for the AIS VBN data and a bridge between software
8 allen_vbn_data_loader = simuran.loader("allen_vbn", cache_directory="tmp")
9 allen_vbn_bridge = AllenVBNBridge()
10
11 # Setup - load the metadata table for AIS VBN and create a recording container
12 allen_vbn_session_metadata = allen_vbn_data_loader.get_sessions_table()
13 allen_vbn_experiment = simuran.RecordingContainer.from_table(
14     table=allen_vbn_session_metadata, loader=allen_vbn_data_loader
15 )
16
17 # Step 1 - Filter the recordings to only those in visual cortex
18 allen_vbn_cortex_recordings = allen_vbn_experiment.filter_table(
19     index
20     for index, recording in enumerate(allen_vbn_experiment)
21     if any(
22         region in allen_vbn_bridge.brain_regions_in_structure("cortex")
23         for region in allen_vbn_bridge.recorded_regions(recording)
24     )
25 )
26
27 # Loop over each allen VBN session with visual cortex data
28 for i, recording in enumerate(allen_vbn_cortex_recordings):
29     # Step 2 - download the data for this recording to disk
30     allen_vbn_data_loader.download_data(recording.attrs["eophys_session_id"])
31
32 # Step 3 - load the data for this recording
33 recording = allen_vbn_cortex_recordings.load(i)
34
35 # Step 4 - get the spike data for this recording
36 unit_table, spike_train = allen_vbn_bridge.spike_train(
37     recording, brain_regions=allen_vbn_bridge.brain_regions_in_structure("cortex")
38 )
39
40 # Step 5 - Convert the spiking data to NEO format
41 neo_spike_train = [NeoBridge.convert_spikes(spikes=spike_train)]
42
43 # Step 6 - pass this data to a GPFA analysis in elephant
44 elephant_gpfa_result = elephant.gpfa.GPFA().fit_transform(neo_spike_train)
45
46 # Step 7 - pass the data to a GPFA analysis in sci-kit learn
47 sklearn_gpfa_result = sklearn.decomposition.FactorAnalysis().fit_transform(
48     bin_spike_train(spike_train=spike_train)
49 )
50
51 # Step 8 - compare the results
52 compare_factor_analysis(elephant_gpfa_result, sklearn_gpfa_result)
53

```

FIGURE 3.6: Python code demonstrating the use of Simuran to run factor analysis on the Allen Institute for Brain Science visual behaviour Neuropixels dataset. Simuran is used to download the data from the Allen Institute, load the data into memory, filter the data to the visual cortex, convert data to the correct format, and then run the analysis.

some neural data analysis software is included, such as NWBExplorer and NWBWidget which allow for detailed inspection of NWB files. Similarly, SpikeInterface (Buccino et al., 2019) provides interoperability between different neural data formats and spike sorting algorithms but does not handle analysis outside of spike sorting.

In a different approach to interoperability than integrating tools, Pynapple (Viejo et al., 2022) offers an abstraction of neural data into the fixed types of *Time Series*, *Interval Sets*, and *Time Series Groups*. As such, any data format which does not naturally fit into one of these groups can be hard to handle, and the user is required to convert their data to fit these groups or create their own subclasses. Likewise, Geppetto (Cantarelli et al., 2018) provides interoperability by abstracting neural data and modelling to enable building tools for visualising neuroscience models, data, and managing simulations (e.g., Open Source Brain is built on top of Geppetto). These forms of abstraction can be powerful for data that fits into the provided categories but can be confusing for data that do not naturally fall into these categories. Simuran avoids this data abstraction by instead providing bridges to convert between different data formats and analysis tools. On the whole, these software tools are valuable in enhancing accessibility and flexibility in neuroscience research, providing an opportunity for researchers to be agnostic to the data format or analysis tool they use.

To continue to improve neuroscience software interoperability, future development of Simuran will support more data formats and analysis tools. As Simuran is written in Python, it can be combined with other biology related workflow management tools to create a full analysis pipeline, such as Snakemake (Mölder et al., 2021) or Nipype (Gorgolewski et al., 2011). Simuran can also be used with tools designed for reproducible research such as Binder (The Jupyter Project et al., 2018) and Sumatra (Davison, 2012). This offers an exciting opportunity to provide free online platforms where anybody can run analysis on their confidential data, or on public data, and share their results with the community. For instance, a user could load up a cloud platform with packages such as Simuran, SpikeInterface, and Open Source Brain preinstalled for spike sorting, neural data analysis, or neural circuit modelling, with the ability to share their data, code, and results in their exacts.

Overall, Simuran enables interoperability between neuroscience software and data formats, allowing for increased data sharing and reproducibility of results. Its platform provides a space for further development of analysis tools, as shown by our related contributions and interactions with open-source neural data analysis packages such as Brainrender, Mountainsort, MNE, and Neo in both code and documentation. The collaboration that tools like Simuran encourage provides a promising future for free online platforms where users can run their analysis and share results.

3.5 Evolution of Simuran and user feedback

We used Simuran for data analysis on projects in our lab. In our first applications, the main aims were to assess how easy-to-use Simuran was, and confirm accuracy of results via tests and comparison to other analysis tools. While we found that results from Simuran and other software were consistent, it was apparent that Simuran was not user friendly. This was caused by two main issues: batch data processing in Simuran was too complex and the Python code was challenging to learn. Batch data processing involved processing a large number of recordings with different recording layouts and experimental conditions. Simuran provided a command line tool designed to assist batch processing, but the time spent learning how to use the tool was too much of an investment. To fix this, we implemented batch data processing in Simuran through linked CSV files and other metadata files. Later feedback on this change was positive, and the new batch data processing system was more intuitive, familiar, and maintainable.

The issue with the barrier to using the Python code was more challenging to solve as it required the user to have strong Python programming knowledge. Simplifying the Python code would only go so far, and it was clear that the best solution would be to provide a visual interface to Simuran. A visual interface would allow researchers to create processing pipelines for their data without needing to be proficient in programming. Through research of other software tools, a node-based design, which is a kind of visual flowchart of program execution, proved intuitive. This visual interface was clearer to understand than the Python code, and in the future could also be used to create Python code that matched the established pipeline - bridging the gap between the user interface (UI) and Python code. While the current visual interface has limited functionality compared to the Python code, Simuran provides templates to make new UI elements that implement missing functionality. Overall, the in-lab use of Simuran was highly beneficial to the development of the software, and the feedback we received from lab members was invaluable. Simuran was simpler to understand after development of the visual interface, and the addition of data tables for batch processing made it easier to manage large experiments.

3.6 Instructions for using Simuran and methods

We plan to extend the instructions for using Simuran on GitHub (<https://github.com/seankmartin/SIMURAN>) and our documentation website (<https://simuran.readthedocs.io/en/latest>). This will allow the instructions to expand and evolve over time. Simuran

can be installed using pip (<https://pypi.org/project/simuran/>) and the instructions on the GitHub page will help to get started with the software. The details that follow here are intended to provide a deeper understanding of the software, including the software engineering approach.

3.6.1 Batch data processing with Simuran

For performing batch analysis, such as analysing all calcium imaging data captured from mice of a certain phenotype, the user can use the *RecordingContainer* class (**Figure 3.1**) to handle a collection of *Recordings*. The following is a brief description of how to perform batch analysis in Simuran:

1. **Experiment description:** Create a table, such as a CSV file, which describes the data to be analysed (e.g., the file paths, the experimental conditions, the stimuli, and other metadata).
2. **Data format specification:** Link the table to a Simuran *RecordingContainer* object and provide a set of *Loaders* that are associated with each *Recording* in the *RecordingContainer*.
3. **Data access and preprocessing:** As data in the *RecordingContainer* is required, it is loaded from disk. By default, this is performed on demand to enable working on a subset of the data without loading the entire dataset.
4. **Data conversion, analysis, and plotting:** Combine *Bridge* objects with the *Analysis* and *Plot* modules to operate on the data using many different software packages. The same analysis can be performed on multiple *Recordings* at once, allowing for multi-processing and reduced user input. For example, one could perform a multi-processing operation to compare data acquired from the same brain region in different mice.
5. **Saving outputs:** Save all analysis results to disk using convenience methods provided by the *RecordingContainer* class and the *Analysis* and *Plot* modules.

3.6.2 Software design patterns used in Simuran

In object-oriented programming, multiple design patterns have been described (Gamma et al., 1995) to enable programmers to effectively architect software that is extendable and maintainable. As programming languages have evolved, several of these design patterns have become obsolete, but many remain highly relevant. In addition to taking general inspiration from design patterns, the following patterns were implemented:

- **Factory pattern:** Node factories create new nodes in the visual editor, while generating unique identifiers for the nodes and registering the nodes with the backend UI controller.
- **Protocol pattern:** Facilitate interaction between different data formats by using *Bridges*. *Bridges* are used to convert data from one format to another, and to provide a common interface for interacting with the data. However, the *Bridges* only expect the data to follow certain conventions (e.g., having an *id* and *label*), and not necessarily to be in a specific format. This allows for the same analysis to be performed on different data formats, and for new data formats to be supported with minimal effort, as *Bridges* can often be re-used.
- **Composite pattern:** The main *Recording* object is composed of other objects as opposed to using inheritance for ease of extension. This allows for the *Recording* object to be extended with new types of data, such as a list of *SpikeTrains*, without having to extensively modify the *Recording* object.
- **Lazy initialisation:** In the *RecordingContainer* object, data is not loaded until the first time it is needed. This allows the user to work with a subset of the data without loading the entire dataset, avoiding filling computer memory.

The protocol design pattern

We will further expand upon how Simuran uses the protocol pattern to interact with different tools and data formats while supporting the addition of new features. In this pattern, one requests that any object that will be supported has certain features. To demonstrate this, consider creating a function to analyse the relationship between all recorded units with a subset of captured LFPs. Creating a function of the signature `def compare_lfp_to_spikes(spike_times_array: np.array, lfp_signals: np.array)` is desirable because of the generic nature of the function with NumPy arrays. However, when the user attempts to load their data for this function, it is possible that this data will not fit in memory simultaneously. Furthermore, the user has to have knowledge of NumPy array slicing to work with only a subset of the LFPs. Instead, by using protocols, we can simply take in one data object and a selector `def compare_lfp_to_spikes(recording: simuran.Recording, desired_channels: List[str])` and any data object which satisfies the protocol will suffice. Here, the protocol is that the given recording object must support accessing its spikes and LFPs like so:

```
1 recording.spike_table
2 recording.lfp_data[desired_channels]
```

To coerce an object that does not directly support this kind of data access into the required format for the protocol, all one needs to do is provide a small mapping, similar to this example. Under the hood, these data accesses could be from Dask arrays (Dask Development Team, 2016), solving the memory issue and hiding the complexity of the data access from the user.

```

1  @dataclass
2  class MyCustomRecording:
3      data : Any
4
5  @property
6  def spike_table():
7      return self.data.spike_table
8
9  @property
10 def lfp_data(slc : slice = None):
11     if slc is not None:
12         return self.data.lfp_data[slc]
13     else:
14         return self.data.lfp_data

```

3.6.3 Creating new functions for the visual editor

Creating a new node type for the UI requires the creation of a *NodeFactory* subclass and a related *Node* subclass. The *NodeFactory* is responsible for creating the *Node* object during runtime and creating the UI elements, while the *Node* is responsible for using the parameters set in the UI to perform operations. To create a new node type, first, create a new *NodeFactory* subclass and give the factory a name and a categorisation. Then, list the parameters modifiable from the UI, establish any input or outputs slots from nodes created by this factory, and optionally add tooltips for the node. Next, create a new *Node* subclass and link it to the factory. In the *Node* subclass, implement the *process* method, which will be called when the node is run to extract information from the parameters set in the UI and use the parameters in an operation. Finally, register the factory with the UI by calling `simuran.register_node_factory(NewNodeFactory)`. Figure 3.7 shows Python code to create a new node that links units to the brain regions they were recorded in (the node is also shown in the UI).

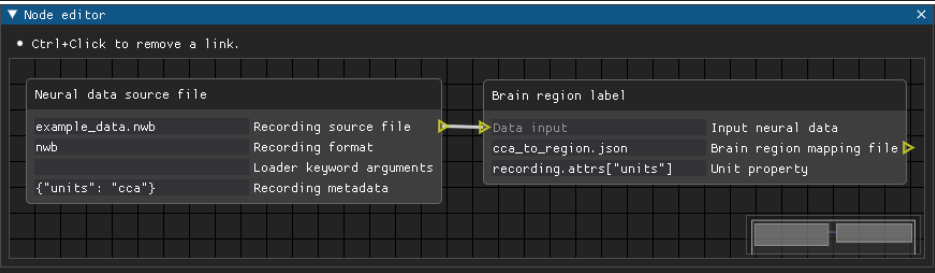
3.6.4 Processing the visual node editor via graphs

To process the visual node editor, Simuran converts the underlying graph into a directed acyclic graph (DAG) and then processes the DAG. This processing is performed by starting at the input nodes and working through the graph until all nodes have been processed. The processing of each node is performed by calling the *process* method of the node, which is responsible for extracting the parameters set in the UI and using

```

1 import simuran
2 from simuran.ui.node import BaseNode, NodeFactory
3 from simuran.ui.node_elements import create_file_select, create_parameter, create_input
4 from simuran.ui.simuran_ui import main_ui as node_editor
5
6
7 class BrainRegionLabelNodeFactory(NodeFactory):
8     def __init__(self, **kwargs):
9         super().__init__(**kwargs)
10        self.label = kwargs.get("label", "Brain region label")
11        self.node_class = BrainRegionLabelNode
12        self.category = "Processing"
13
14        self.attributes = [
15            create_input("Input neural data", "Data input"),
16            create_file_select("Brain region mapping file", output=True),
17            create_parameter("Unit property", tooltip="The unit property to use"),
18        ]
19
20
21 class BrainRegionLabelNode(BaseNode):
22     def process(self, nodes):
23         super().process(nodes)
24
25         # Parse out parameters from the UI
26         self.source_file = self.get_value_of_label(label="Brain region mapping file")
27         self.unit_property = self.get_value_of_label(label="Unit property")
28         self.input_recording_node = self.find_matching_input_node(
29             receiver_label="Input neural data"
30         )
31
32         # Load the brain region mapping onto the input neural data
33         recording = self.input_recording_node.recording
34         recording.attrs["brain_region_mapping"] = simuran.ParamHandler(
35             source_file=self.source_file
36         )
37         recording.attrs["units_to_brain_region"] = {
38             unit: self.brain_region_mapping[unit.getattr(self.unit_property)]
39             for unit in recording.units
40         }
41
42 simuran.register_node_factory(BrainRegionLabelNodeFactory)
43 node_editor()
44

```



The screenshot shows the 'Node editor' window with a grid background. On the left, a 'Neural data source file' node is shown with parameters: 'example_data.nwb' (Recording source file), 'nwb' (Recording format), and '{\"units\": \"cca\"}' (Recording metadata). An arrow points from the 'Recording source file' parameter to the 'Data input' parameter of the 'Brain region label' node on the right. The 'Brain region label' node has parameters: 'cca_to_region.json' (Brain region mapping file) and 'recording.attrs[\"units\"]' (Unit property). The 'Data input' parameter of the 'Brain region label' node is also linked to the 'Input neural data' parameter of the 'Neural data source file' node.

FIGURE 3.7: Using Python to create a new type of node in Simuran’s visual interface that associates single units to the brain region they were recorded from.

The node takes in a brain region mapping file and a unit property. Then the node adds a new attribute to the input neural data (a *Recording*) called `units_to_brain_region` which maps each single-unit to a brain region. The new node is then registered with Simuran and the node editor is opened. A data source node is added and linked to the newly created node type shown on the right (Brain region label node), and parameters are filled out.

these to perform the intended function of the node. The output of each node is then passed to the next node in the graph. This process is repeated until all nodes have been processed. In this way, the processing of the graph is performed in a way that is easy to extend and maintain. Furthermore, Simuran can handle branching paths and multiple inputs or outputs from a single node.

3.6.5 Improving performance with Simuran

Using Simuran can significantly reduce the time to create analysis pipelines, without introducing any significant performance overhead. Furthermore, the software can have significant performance benefits if multi-processing is used, as it is designed to be able to run on a cluster or a single machine with multiple cores. As such, in most cases, users will be able to run their analysis pipelines faster than they would have been able to without Simuran. For instance, in testing on a Windows 10 machine with an Intel 15-6500 4-core CPU 3.20GHz and 16GB RAM, Simuran loaded sets of calcium imaging data from AIS in less than half the time than it took using the AllenSDK alone, which would scale with the number of available cores.

3.6.6 Sharing, documenting, and testing Simuran

Simuran is openly hosted at <https://github.com/seankmartin/SIMURAN> under the GPL-3.0 license, which allows redistribution of the free software, and points out that the software is provided without any warranty. Simuran's documentation is hosted on Read the docs at <https://simuran.readthedocs.io/en/latest/>. This involves a full description of the Python package, examples of usage, and general information about Simuran. Simuran follows many of the best practices for open-source software development, such as continuous integration and automated testing. In addition to this, Simuran openly displays code metrics which indicate the quality of the codebase and how maintainable it is.

To ensure accuracy of the code, Simuran has been manually tested, and the core codebase has been automatically tested. Automatic tests have been performed via `pytest` and the percentage of the codebase that has been tested has been checked via `coverage`. These automatic tests are performed every time Simuran is updated, ensuring that there is no regression in the codebase. For test coverage, at least 90% coverage of the core codebase was achieved, ensuring that the critical code paths were covered. Details of the tests performed are all publicly available. Currently, the custom *Bridges* and *Loaders* for certain data formats have primarily been manually tested, and the automatic tests for these are not as detailed as the core codebase. In the future, the

same test coverage for the core codebase can be enforced for the different *Loaders* and *Bridges* subclasses, along with custom *Analysis* module members.

References

- Abrams, M. B., Bjaalie, J. G., Das, S., Egan, G. F., Ghosh, S. S., Goscinski, W. J., Grethe, J. S., Kotaleski, J. H., Ho, E. T. W., Kennedy, D. N., Lanyon, L. J., Leergaard, T. B., Mayberg, H. S., Milanese, L., Mouček, R., Poline, J. B., Roy, P. K., Strother, S. C., Tang, T. B., . . . Martone, M. E. (2022). A standards organization for open and FAIR neuroscience: The international neuroinformatics coordinating facility. *Neuroinformatics*, *20*(1).
- Adrian, S., Christian, K., Jan, B., Thomas, W., & Jan, G. (2014). File format and library for neuroscience data and metadata. *Frontiers in Neuroinformatics*, *8*.
- Allen, W. E., Chen, M. Z., Pichamoorthy, N., Tien, R. H., Pachitariu, M., Luo, L., & Deisseroth, K. (2019). Thirst regulates motivated behavior through modulation of brainwide neural population dynamics. *Science*, *364*(6437), eaav3932.
- Allen Institute for Brain Science. (2021). Allen brain observatory: Visual behavior optical physiology technical whitepaper.
- Allen Institute for Brain Science. (2022). Allen brain observatory: Visual behavior neuropixels technical whitepaper.
- Buccino, A. P., Hurwitz, C. L., Magland, J., Garcia, S., Siegle, J. H., Hurwitz, R., & Hennig, M. H. (2019). SpikeInterface, a unified framework for spike sorting. *bioRxiv*, 796599.
- Cantarelli, M., Marin, B., Quintana, A., Earnshaw, M., Court, R., Gleeson, P., Dura-Bernal, S., Silver, R. A., & Idili, G. (2018). Geppetto: A reusable modular open platform for exploring neuroscience data and models. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *373*(1758), 20170380.
- Carrillo-Reid, L., Han, S., Yang, W., Akrouh, A., & Yuste, R. (2019). Controlling visually guided behavior by holographic recalling of cortical ensembles. *Cell*, *178*(2), 447–457.e5.
- Chu, J. P., & Kemere, C. T. (2021). GhostiPy: An efficient signal processing and spectral analysis toolbox for large data. *eNeuro*, *8*(6).
- Cunningham, J. P., & Yu, B. M. (2014). Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, *17*(11), 1500–1509.
- DANDI team. (2022). *DANDI: Distributed archives for neurophysiology data integration*.
- Dask Development Team. (2016). *Dask: Library for dynamic task scheduling*.
- Davison, A. (2012). Automated capture of experiment context for easier reproducibility in computational research. *Computing in Science & Engineering*, *14*(4), 48–56.

- Deisseroth, K. (2015). Optogenetics: 10 years of microbial opsins in neuroscience. *Nature neuroscience*, *18*(9), 1213–1225.
- Denker, M., Yegenoglu, A., & Grün, S. Collaborative HPC-enabled workflows on the HBP Collaboratory using the Elephant framework. In: *Neuroinformatics 2018*. 2018, P19.
- Dong, Z., Mau, W., Feng, Y., Pennington, Z. T., Chen, L., Zaki, Y., Rajan, K., Shuman, T., Aharoni, D., & Cai, D. J. (2022). Minian, an open-source miniscope analysis pipeline (C. Kemere, L. L. Colgin, & C. Kemere, Eds.). *eLife*, *11*, e70661.
- Gamma, E., Helm, R., Johnson, R., & Vlissides, J. (1995). *Design patterns: Elements of reusable object-oriented software*. Addison-Wesley Longman Publishing Co., Inc.
- Gao, L., Liu, S., Gou, L., Hu, Y., Liu, Y., Deng, L., Ma, D., Wang, H., Yang, Q., Chen, Z., Liu, D., Qiu, S., Wang, X., Wang, D., Wang, X., Ren, B., Liu, Q., Chen, T., Shi, X., . . . Yan, J. (2022). Single-neuron projectome of mouse prefrontal cortex. *Nature Neuroscience*, *25*(4), 515–529.
- Garcia, S., Guarino, D., Jaillet, F., Jennings, T., Pröpper, R., Rautenberg, P., Rodgers, C., Sobolev, A., Wachtler, T., Yger, P., & Davison, A. (2014). Neo: An object model for handling electrophysiology data in multiple formats. *Frontiers in Neuroinformatics*, *8*.
- Giovannucci, A., Friedrich, J., Gunn, P., Kalfon, J., Brown, B. L., Koay, S. A., Taxidis, J., Najafi, F., Gauthier, J. L., Zhou, P., Khakh, B. S., Tank, D. W., Chklovskii, D. B., & Pnevmatikakis, E. A. (2019). CaImAn an open source tool for scalable calcium imaging data analysis (D. Kleinfeld & A. J. King, Eds.). *eLife*, *8*, e38173.
- Gleeson, P., Cantarelli, M., Marin, B., Quintana, A., Earnshaw, M., Sadeh, S., Piasini, E., Birgiolas, J., Cannon, R. C., Cayco-Gajic, N. A., Crook, S., Davison, A. P., Dura-Bernal, S., Ecker, A., Hines, M. L., Idili, G., Lanore, F., Larson, S. D., Lytton, W. W., . . . Silver, R. A. (2019). Open source brain: A collaborative resource for visualizing, analyzing, simulating, and developing standardized models of neurons and circuits. *Neuron*, *103*(3), 395–411.e5.
- Gokcen, E., Jasper, A. I., Semedo, J. D., Zandvakili, A., Kohn, A., Machens, C. K., & Yu, B. M. (2021, September 1). *Disentangling the flow of signals between populations of neurons* (preprint). Neuroscience.
- Gorgolewski, K., Burns, C., Madison, C., Clark, D., Halchenko, Y., Waskom, M., & Ghosh, S. (2011). Nipype: A flexible, lightweight and extensible neuroimaging data processing framework in python. *Frontiers in Neuroinformatics*, *5*.
- Gramfort, A., Luessi, M., Larson, E., Engemann, D., Strohmeier, D., Brodbeck, C., Goj, R., Jas, M., Brooks, T., Parkkonen, L., & Hämäläinen, M. (2013). MEG and EEG data analysis with MNE-python. *Frontiers in Neuroscience*, *7*.

- Grienberger, C., & Konnerth, A. (2012). Imaging calcium in neurons. *Neuron*, *73*(5), 862–885.
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., Fernández del Río, J., Wiebe, M., Peterson, P., ... Oliphant, T. E. (2020). Array programming with NumPy. *Nature*, *585*, 357–362.
- Harris, K. D. (2020). Nonsense correlations in neuroscience. *bioRxiv*, 2020.11.29.402719.
- Islam, M. N., Martin, S. K., Aggleton, J. P., & O'Mara, S. M. (2019). NeuroChaT: A toolbox to analyse the dynamics of neuronal encoding in freely-behaving rodents in vivo. *Wellcome Open Research*, *4*, 196.
- Jonas, E., & Kording, K. P. (2017). Could a neuroscientist understand a microprocessor? *PLOS Computational Biology*, *13*(1), e1005268.
- Jones, I. S., & Kording, K. P. (2021). Might a single neuron solve interesting machine learning problems through successive computations on its dendritic tree? *Neural Computation*, *33*(6), 1554–1571.
- Jönsson, D., Steneteg, P., Sundén, E., Englund, R., Kottraval, S., Falk, M., Ynnerman, A., Hotz, I., & Ropinski, T. (2020). Inviwo — a visualization system with usage abstraction levels. *IEEE Transactions on Visualization and Computer Graphics*, *26*(11), 3241–3254.
- Kaefer, K., Nardin, M., Blahna, K., & Csicsvari, J. (2020). Replay of behavioral sequences in the medial prefrontal cortex during rule switching. *Neuron*, *0*(0).
- Lopes, G., Bonacchi, N., Frazão, J., Neto, J. P., Atallah, B. V., Soares, S., Moreira, L., Matias, S., Itskov, P. M., Correia, P. A., Medina, R. E., Calcaterra, L., Dreosti, E., Paton, J. J., & Kampff, A. R. (2015). Bonsai: An event-based framework for processing and controlling data streams. *Frontiers in Neuroinformatics*, *9*.
- Machado, T. A., Kauvar, I. V., & Deisseroth, K. (2022). Multiregion neuronal activity: The forest and the trees. *Nature Reviews Neuroscience*, *23*(11), 683–704.
- Madisen, L., Garner, A. R., Shimaoka, D., Chuong, A. S., Klapoetke, N. C., Li, L., van der Bourg, A., Niino, Y., Egolf, L., Monetti, C., Gu, H., Mills, M., Cheng, A., Tasic, B., Nguyen, T. N., Sunkin, S. M., Benucci, A., Nagy, A., Miyawaki, A., ... Zeng, H. (2015). Transgenic mice for intersectional targeting of neural sensors and effectors with high specificity and performance. *Neuron*, *85*(5), 942–958.
- Magland, J., Jun, J. J., Lovero, E., Morley, A. J., Hurwitz, C. L., Buccino, A. P., Garcia, S., & Barnett, A. H. (2020). SpikeForest, reproducible web-facing ground-truth validation of automated neural spike sorters (M. Meister, R. L. Calabrese, & M. Meister, Eds.). *eLife*, *9*, e55167.

- McKinney, W. (2010). Data Structures for Statistical Computing in Python. In S. van der Walt & J. Millman (Eds.), *Proceedings of the 9th Python in Science Conference* (pp. 56–61).
- Mölder, F., Jablonski, K. P., Letcher, B., Hall, M. B., Tomkins-Tinch, C. H., Sochat, V., Forster, J., Lee, S., Twardziok, S. O., Kanitz, A., Wilm, A., Holtgrewe, M., Rahmann, S., Nahnsen, S., & Köster, J. (2021, April 19). *Sustainable data analysis with snakemake* (10:33). F1000Research.
- Pachitariu, M., Steinmetz, N., Kadir, S., Carandini, M., & D, H. K. (2016). Kilo-sort: Realtime spike-sorting for extracellular electrophysiology with hundreds of channels. *bioRxiv*, 061481.
- Pandarinaath, C., O’Shea, D. J., Collins, J., Jozefowicz, R., Stavisky, S. D., Kao, J. C., Trautmann, E. M., Kaufman, M. T., Ryu, S. I., Hochberg, L. R., Henderson, J. M., Shenoy, K. V., Abbott, L. F., & Sussillo, D. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods*, *15*(10), 805–815.
- Pang, R., Lansdell, B. J., & Fairhall, A. L. (2016). Dimensionality reduction in neuroscience. *Current Biology*, *26*(14), R656–R660.
- Paulk, A. C., Kfir, Y., Khanna, A. R., Mustroph, M. L., Trautmann, E. M., Soper, D. J., Stavisky, S. D., Welkenhuysen, M., Dutta, B., Shenoy, K. V., Hochberg, L. R., Richardson, R. M., Williams, Z. M., & Cash, S. S. (2022). Large-scale neural recordings with single neuron resolution using neuropixels probes in human cortex. *Nature Neuroscience*, 1–12.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, *12*, 2825–2830.
- Quiroga, R. Q. (2012). Spike sorting. *Current biology: CB*, *22*(2), R45–46.
- Rossant, C., Kadir, S. N., Goodman, D. F. M., Schulman, J., Hunter, M. L., Saleem, A. B., Grosmark, A., Belluscio, M., Denfield, G. H., Ecker, A. S., Tolias, A. S., Solomon, S., Buzsaki, G., Carandini, M., & Harris, K. D. (2016). Spike sorting for large, dense electrode arrays. *Nature neuroscience*, *19*(4), 634–641.
- Roth, B. L. (2016). DREADDs for neuroscientists. *Neuron*, *89*(4), 683–694.
- Rübel, O., Tritt, A., Ly, R., Dichter, B. K., Ghosh, S., Niu, L., Baker, P., Soltesz, I., Ng, L., Svoboda, K., Frank, L., & Bouchard, K. E. (2022). The neurodata without borders ecosystem for neurophysiological data science (L. L. Colgin & S. P. Jadhav, Eds.). *eLife*, *11*, e78362.
- Saxena, S., & Cunningham, J. P. (2019). Towards the neural population doctrine. *Current Opinion in Neurobiology*, *55*, 103–111.

- Semedo, J. D., Jasper, A. I., Zandvakili, A., Krishna, A., Aschner, A., Machens, C. K., Kohn, A., & Yu, B. M. (2022). Feedforward and feedback interactions between visual cortical areas use different population activity patterns. *Nature Communications*, *13*(1), 1099.
- Semedo, J. D., Zandvakili, A., Machens, C. K., Yu, B. M., & Kohn, A. (2019). Cortical areas interact through a communication subspace. *Neuron*, *102*(1), 249–259.e4.
- Siegle, J. H., Jia, X., Durand, S., Gale, S., Bennett, C., Graddis, N., Heller, G., Ramirez, T. K., Choi, H., Luviano, J. A., Groblewski, P. A., Ahmed, R., Arkhipov, A., Bernard, A., Billeh, Y. N., Brown, D., Buice, M. A., Cain, N., Caldejon, S., ... Koch, C. (2019). A survey of spiking activity reveals a functional hierarchy of mouse corticothalamic visual areas. *bioRxiv*, 805010.
- Siegle, J. H., Jia, X., Durand, S., Gale, S., Bennett, C., Graddis, N., Heller, G., Ramirez, T. K., Choi, H., Luviano, J. A., Groblewski, P. A., Ahmed, R., Arkhipov, A., Bernard, A., Billeh, Y. N., Brown, D., Buice, M. A., Cain, N., Caldejon, S., ... Koch, C. (2021). Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature*, *592*(7852), 86–92.
- 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., ... Harris, T. D. (2021). Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science*, *372*(6539).
- Stevenson, I. H., & Kording, K. P. (2011). How advances in neural recording affect data analysis. *Nature Neuroscience*, *14*(2), 139–142.
- Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C. B., Carandini, M., & Harris, K. D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, *364*(6437), eaav7893.
- The International Brain Laboratory, Banga, K., Benson, J., Bonacchi, N., Bruijns, S. A., Campbell, R., Chapuis, G. A., Churchland, A. K., Davatolhagh, M. F., Lee, H. D., Faulkner, M., Hu, F., Hunterberg, J., Khanal, A., Krasniak, C., Meijer, G. T., Miska, N. J., Mohammadi, Z., Noel, J.-P., ... Winter, O. (2022, May 9). *Reproducibility of in-vivo electrophysiological measurements in mice*. bioRxiv.
- The Jupyter Project, Bussonnier, M., Forde, J., Freeman, J., Granger, B., Head, T., Holdgraf, C., Kelley, K., Nalvarte, G., Osheroff, A., Pacer, M., Panda, Y., Perez, F., Ragan-Kelley, B., & Willing, C. (2018). Binder 2.0 - reproducible, interactive, sharable environments for science at scale. *Proceedings of the 17th Python in Science Conference*, 113–120.
- Urai, A. E., Doiron, B., Leifer, A. M., & Churchland, A. K. (2021). Large-scale neural recordings call for new insights to link brain and behavior, 24.

- Viejo, G., Levenstein, D., Carrasco, S. S., Mehrotra, D., Mahallati, S., Vite, G. R., Denny, H., Sjulson, L., Battaglia, F. P., & Peyrache, A. (2022, December 7). Pynapple: A toolbox for data analysis in neuroscience.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., . . . SciPy 1.0 Contributors. (2020). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, *17*, 261–272.
- Yang, W., Carrillo-Reid, L., Bando, Y., Peterka, D. S., & Yuste, R. (2018). Simultaneous two-photon imaging and two-photon optogenetics of cortical circuits in three dimensions (K. Svoboda, Ed.). *eLife*, *7*, e32671.
- Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. V., & Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Journal of Neurophysiology*, *102*(1), 614–635.

4 Anterior thalamic lesions abrogate subicular spatial coding

Introduction: In this study, we investigated a set of interconnected brain regions within the rodent limbic system, consisting of the subiculum, CA1, retrosplenial cortex, and anterior thalamic nuclei. We focused on the role of the anterior thalamic nuclei for spatial coding in the subiculum and CA1. By understanding the role of the ATN in this circuit, we aimed to shed light on the mechanisms underlying spatial coding and implications for memory and amnesia.

Methods: Electrophysiological recordings were performed in freely moving male rats with permanent anterior thalamic nuclei lesions or temporary inhibition of the anterior thalamic nuclei to obtain single-unit measurements. Recordings of units were obtained from the dorsal subiculum and CA1 of lesioned rats, and the dorsal subiculum of control rats. Control and lesioned rats were tested in behavioural tasks, including spatial alternation tasks and open-field exploration.

Results: Our findings revealed that both permanent and reversible lesions of the anterior thalamic nuclei led to a complete cessation of subicular spatial signalling. As a result, spatial memory performance was reduced to chance levels. Interestingly, place cells in the hippocampal CA1 region remained largely unaffected by the lesions.

Conclusions: From our results, we can conclude that the outputs from CA1 alone are insufficient for subicular spatial coding. Instead, the anterior thalamic nuclei play a critical role in providing key inputs or modulating the subiculum to support spatial coding. These findings deepen our understanding of the intricate mechanisms underlying spatial coding and memory formation in the limbic system.

4.1 My contributions to this study

As this is a multi-author study, I outline my role in this study. I did not perform any of the electrophysiology or histology performed in this study, which was led by Dr. Frost and Dr. Cafalchio. Additionally, some of the data analysis was performed collaboratively with Dr. Frost and Dr. Cafalchio. Figure 4.2, Figure 4.4, and Figure 4.7 were created from this joint analysis. My main contributions were expanding upon the initial finding that subicular spatial cells appeared to be abrogated post ATN lesion. Firstly, I expanded NeuroChaT (Islam et al., 2019) with additional plotting and batching functionality for further analysis in this study. Then I assisted with the analysis of the CA1 data and presented preliminary results of the study at a poster session at the Bernstein conference. Finally, I helped expand upon the initial results with a Bayesian style interpretation of the findings, performed downsampling analysis to account for the lack of movement in muscimol recordings, and investigated if the spike width differences in control and lesion rats suggested recording from two different neural populations.

4.2 The role of the hippocampus and anterior thalamic nuclei in spatial memory

We focused on a circuit in the rodent limbic system involving the subiculum, retrosplenial cortex, anterior thalamic nuclei, and CA1 (Figure 4.1). Our motivation for investigating such a system was three-fold.

Firstly, the subiculum and CA1 are part of a larger hippocampal circuit consisting of the subicular complex (containing the subiculum, the parasubiculum and the presubiculum), hippocampus proper (composed of subfields CA1, CA2, and CA3), the dentate gyrus, and the entorhinal cortex (Amaral & Witter, 1989). The anatomical connections in this circuit have been well studied, and information on projection density is available for the anterior thalamic nuclei (ATN), cornu ammonis (CA) fields, subiculum, and retrosplenial cortex (Jankowski et al., 2013; Sugar et al., 2011). For instance, the projections from CA1 to its major output, the subiculum proper, are organised such that the projections from proximal CA1 terminate in distal subiculum, and the projections from distal CA1 terminate in proximal subiculum (where proximal is relative to the CA1 subiculum border, so distal CA1 borders CA2 while distal subiculum borders the presubiculum, Amaral et al., 1991). Further, the ATN have direct, reciprocal connections with the subiculum (Shibata, 1993) and have further indirect connections via the entorhinal cortex, presubiculum and retrosplenial cortex (see Sugar et al., 2011,

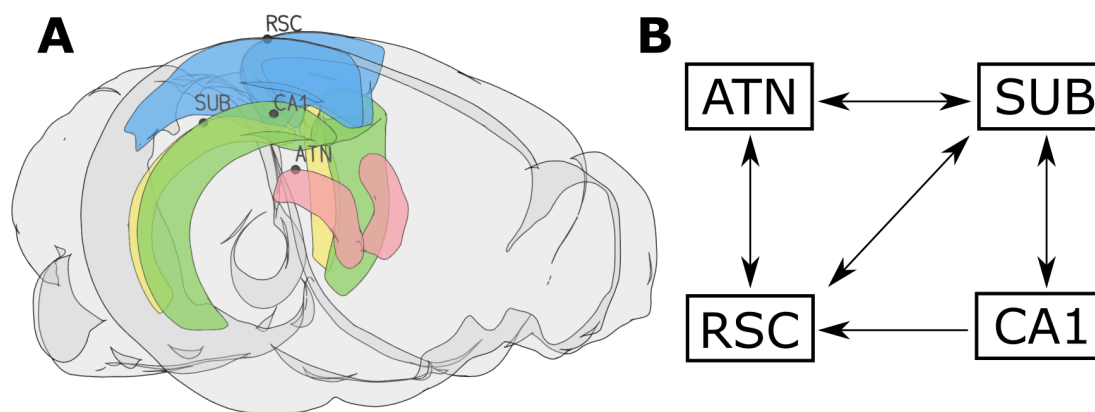


FIGURE 4.1: Limbic system brain regions considered in the ATN lesion study and a brief description of the connections between them.

(A) The four brain regions recorded from or manipulated in this study, shown using the Allen mouse brain atlas in *Brainrender*.

(B) Depicting the primary connections between the brain regions in this study. The arrows indicate the direction of the connection from one region to another as efferent, afferent, or reciprocal.

Key: SUB – subiculum, CA1 – field CA1 of the hippocampus, ATN – anterior thalamic nuclei, RSC – retrosplenial cortex.

for information on retrosplenial connectivity).

Secondly, the subiculum and CA1 both contain cells which are tuned to elements of the spatial environment, but the types of cells found in these regions differ. Place cells are abundant in CA1, while the subiculum displays a greater variety of spatially tuned cells (place, grid, head-direction, and boundary vector cells), but in lower frequency (Brotons-Mas et al., 2017; Brotons-Mas et al., 2010; Hafting et al., 2005; Moser et al., 2008; O’Keefe & Burgess, 1996). These spatially tuned subicular cells tend to produce somewhat less “clean” response maps compared to the model examples of place cells in CA1, or grid cells in entorhinal cortex (Sharp & Green, 1994) – this may be related to the different properties of the pyramidal neurons in CA1 and subiculum (Staff et al., 2000). The ATN, which contains head direction cells (Taube, 1995; Tsanov et al., 2011) may provide key information input to the subiculum for spatial coding. The ATN comprises the anteroventral nucleus, the anteromedial nucleus, and the anterodorsal nucleus of the thalamus, which are associated with different functions (Aggleton et al., 2022). For instance, the anterodorsal nucleus is particularly key to head directional information input to the hippocampal formation (Viejo & Peyrache, 2019). In this study, we inhibited all of the ATN simultaneously due to the difficulty in achieving individual lesions of ATN during *in vivo* studies.

Finally, temporal lobe amnesia (involving damage to the medial temporal lobe, primarily hippocampal) and diencephalic amnesia (involving damage to the diencephalon, primarily the thalamus) are widespread neurological disorders. Amnesia is a form of

memory deficiency involving a combination of retrograde memory impairment (the inability to recall memories before a certain point in time) and anterograde memory impairment (the degradation of short-term memory and learning Smith et al., 2013). With our discussion above in mind, the thalamus and the ATN could play a direct role in diencephalic amnesia, and through connections and information transfer to the hippocampal formation, an indirect role in temporal lobe amnesia. This is further emphasised if we consider that there is a strong link between spatial tuning and memory in the limbic system, as spatially tuned cells enable memories related to space and time. Despite this, the hippocampus has largely been the focus of amnesia studies, even though it is just one part of a larger episodic memory system (O’Mara & Aggleton, 2019).

Combining these facets, we have a well-connected network involving heterogeneous cell types that is heavily implicated to be involved in memory and learning, in which the role of the ATN may be under-represented compared to the role of the hippocampus (Aggleton & O’Mara, 2022). We focused on the role of the ATN and its relationship to the other regions in this system via simultaneous recordings from the dorsal subiculum and retrosplenial cortex with permanent N-Methyl-D-aspartic acid (NMDA) lesions to the ATN (which kills cells due to excitotoxicity) as well as temporary inhibition of the ATN via muscimol (a *gamma*-aminobutyric acid (GABA) receptor agonist). Separate CA1 recordings were also performed to contrast the effect of ATN damage on the subiculum and CA1, as CA1 does not receive direct projections from ATN but displays spatial tuning and projects heavily to the subiculum.

Overall, we found that ATN inactivation abolished the presence of spatially tuned cells in our subicular tetrode recordings, whilst leaving CA1 place cells intact, and additionally degraded the performance to chance level on a spatial alteration task.

4.3 Abrogation of subicular spatial coding following ATN inactivation

Overall, the bilateral ATN NMDA lesions were highly effective, significantly reducing cell counts in the ATN while sparing nucleus reuniens (see subsection A.1.2 for further detail of lesion effectiveness). These lesions resulted in a reduction of spatial alteration task performance to chance level (Figure A.1). Spatial and non-spatial single units were recorded in the dorsal subiculum of all control rats (Figure 4.2). Of 82 single units recorded in the control rats, 47 (57%) were considered spatial units. These units consisted of place ($n = 11$; 23%), head-direction ($n = 20$; 43%), border ($n = 5$; 11%), and grid ($n = 11$; 23%) cells. A further 35 (43%) did not show obvious spatial properties

(e.g., no clear place field or preferred head direction). Strikingly, no spatial units were recorded in the ATN lesion-induced dysfunction (ATNx) rats although 21 non-spatial units, that is, those showing no preferential firing in a specific place or head orientation, were recorded (Figure 4.2). Electrode tracks are reconstructed in Figure A.2F.

4.3.1 Differences in spatial cell recording frequencies between control and lesioned rats

Despite this finding, the presence of spatially tuned cells in ATNx rats may have been missed, despite there being no actual change in the response of spatially tuned cells post-lesion. Here, we investigate whether the absence of spatially tuned cells in ATNx rats could be an artifact stemming from a limited number of recordings performed in different rats or following tetrode manipulation. Consequently, our analysis does not primarily revolve around the proportion of spatial and non-spatial cells. Instead, our aim is to analyse if the data indicates that we have performed a sufficient number of independent recordings in both control and ATNx rats. To assess this absence of spatially tuned cells in ATNx rats and quantify the change in recordings, the presence or absence of spatial cells and non-spatial cells in subicular recordings was noted (Table 4.1). Our null hypothesis was that the proportion of recordings performed with spatial or non-spatial cells recorded would be independent of the rat being a control rat or a lesioned rat. The results from running a two-sided Barnard's unconditional exact test (Barnard, 1947) for a binomial model on this 2×2 contingency table consisting of the observed frequencies from the control and ATNx recordings are as follows:

- **Recordings with spatial cells:** Control rats had 15 out of 53 recordings with at least one spatial cell, while ATNx rats had 0 out of 47 recordings with a spatial cell. The data strongly support rejecting the null hypothesis that performing a recording with at least one spatial cell captured is independent of the rat being a control rat or lesioned rat (Wald statistic 4.574, difference in proportion 0.28, $p < 0.0001$).
- **Non-spatial cell recordings:** Control rats had 16 out of 53 recordings with at least one non-spatial cell, while ATNx rats had 13 out of 47 recordings with at least one non-spatial cell. The data fails to reject the null hypothesis that performing a recording with at least one non-spatial cell captured is independent of the rat being a control rat or lesioned rat (Wald statistic 0.28, difference in proportion 0.0259, $p = 0.81$).

To further investigate how likely performing recording sessions that capture spatial

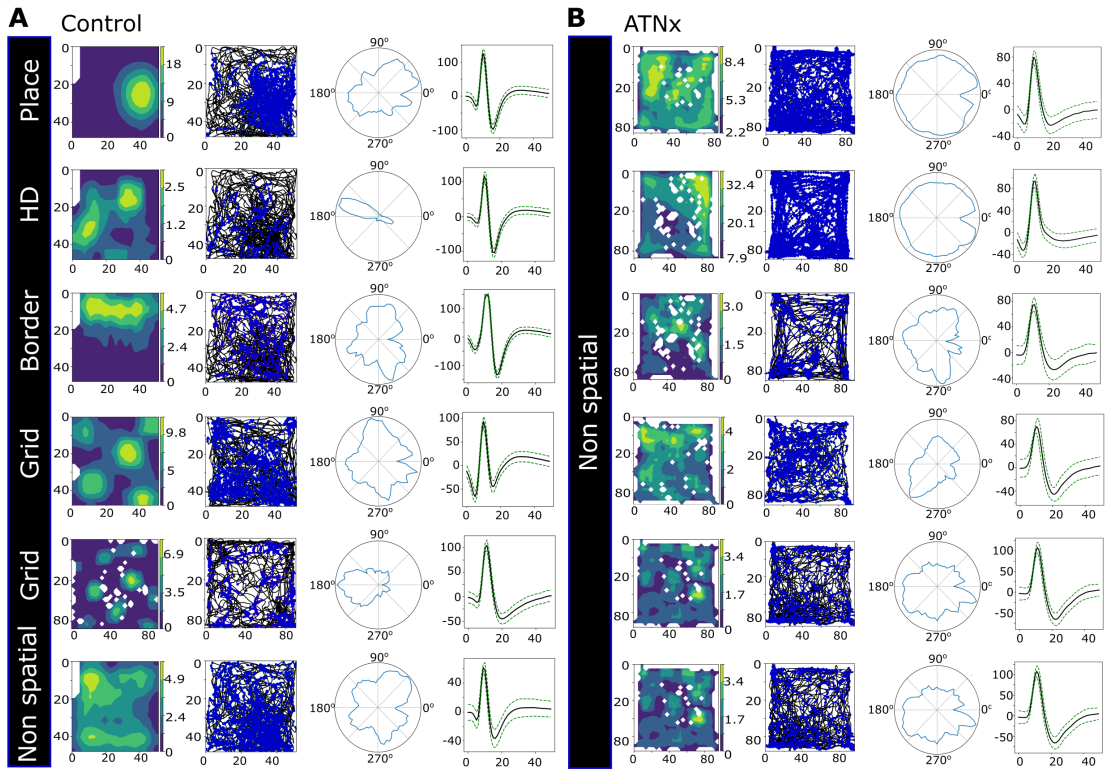


FIGURE 4.2: Representative single units recorded in the dorsal subiculum from control (Control; **A**) and ATN lesioned (ATNx; **B**) rats. For individual units, the figures illustrate the following (from left to right in each row):

1. The smoothing firing rate map, produced by dividing the spatially binned position of the rat when cells fired by the bin occupancy count.
2. The spike position map; with each the position of the animal when the cell spiked plotted as a blue dot over the animal's exploration path (in black).
3. The head direction (HD) map, representing cell activity when the animal faces particular directions.
4. The mean waveform of the unit in μV with error ranges representing 95% confidence intervals.

Observed frequencies for recordings	Control	Lesion	Total
Spatial cells			
At least one spatial cell was recorded	15	0	15
Spatial cells were not recorded	38	47	85
Total recordings performed	53	47	100
Non-spatial cells			
At least one non-spatial cell was recorded	16	13	29
Non-spatial cells were not recorded	37	34	61
Total recordings performed	53	47	100

TABLE 4.1: Recording counts from control and ATNx animals. The recordings (where a “recording” is defined as a single trial of an open field recording session) were selected such that recordings were independent. Recordings were deemed to be independent if they were performed before and after adjustment of electrode position, or if they were several days apart, as the electrodes had likely shifted naturally. For each of these recordings, the presence or absence of both spatial cells and non-spatial cells was marked, and the observed frequencies are presented. Recordings were considered to contain spatial cells or non-spatial cells (regardless of how many of those cells were recorded, or if the same cell had been previously recorded).

or non-spatial cells are in control and ATNx rats, a Bayesian analysis was performed (MacKay, 2003). First, we calculated the joint posterior distribution of performing a recording with at least one spatial or non-spatial cell given the data observed in the contingency table Table 4.1 (Figure 4.3). Integrating over the joint posterior distribution, we find that, at 99.99% chance, it is more likely that a recording would capture a spatial cell in a control recording compared with an ATNx recording. Similarly, we find that at 93.28% chance, it is five times more likely that a recording would capture a spatial cell in a control recording compared with an ATNx recording. While this does not eliminate the possibility of recording a spatial cell from the subiculum of an ATNx rat, it would seem unlikely that we missed the presence of spatially tuned cells in ATNx rats simply due to a lack of recordings performed. As such, the observed difference may stem from a change in the proportion of response of the spatially tuned in the subiculum post ATN-lesion.

4.3.2 Subicular cell firing properties in control and lesioned rats

We compare subicular cell firing properties between control and lesioned rats in Table 4.2 and the main significant findings are summarised in Figure 4.4 .

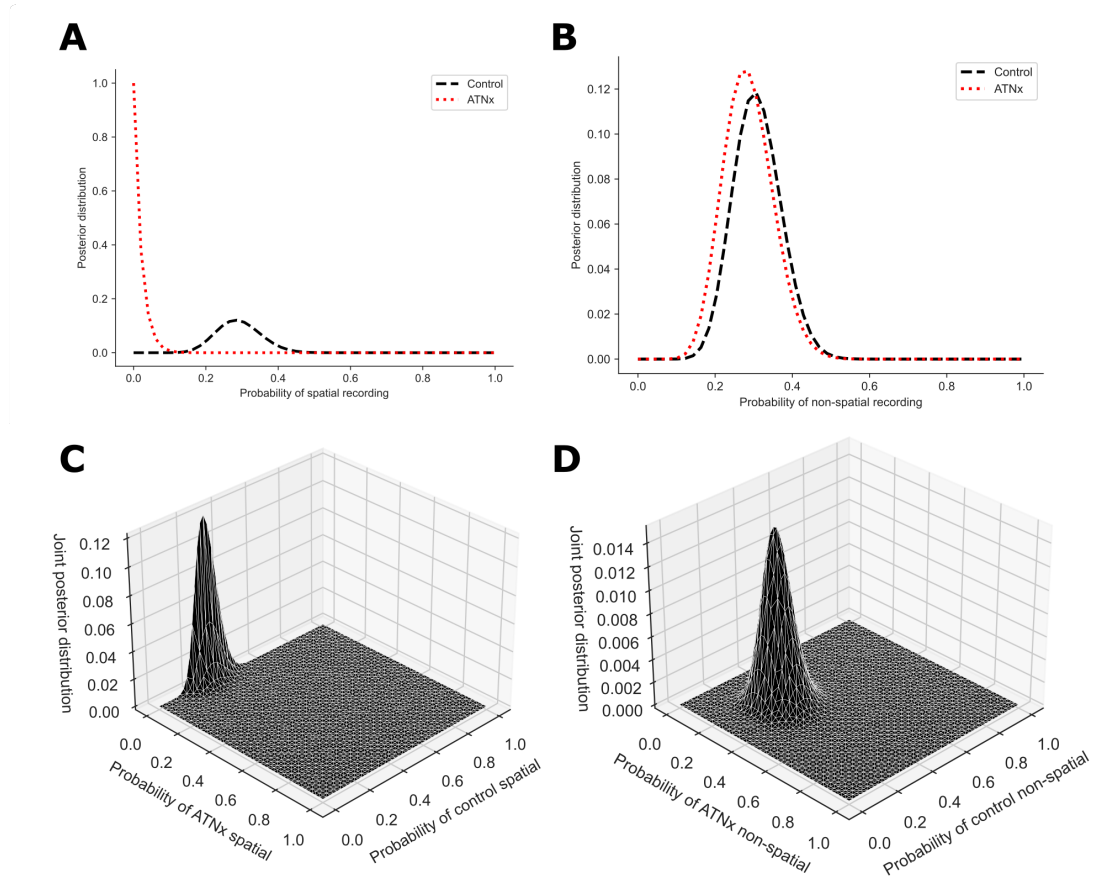


FIGURE 4.3: Exploring the probability of performing a recording with at least one spatial or non-spatial cell from the subiculum in control and lesioned rats following our methods.

(A, B) The posterior distributions of performing a recording in dorsal subiculum and obtaining at least one spatial cell A, and the same, but for at least one non-spatial cell B, given the observed data.

(C, D) Combining the above to form the joint posterior distributions. Integrating over this joint posterior distribution allows for comparison of recording probability between control and lesion rats.

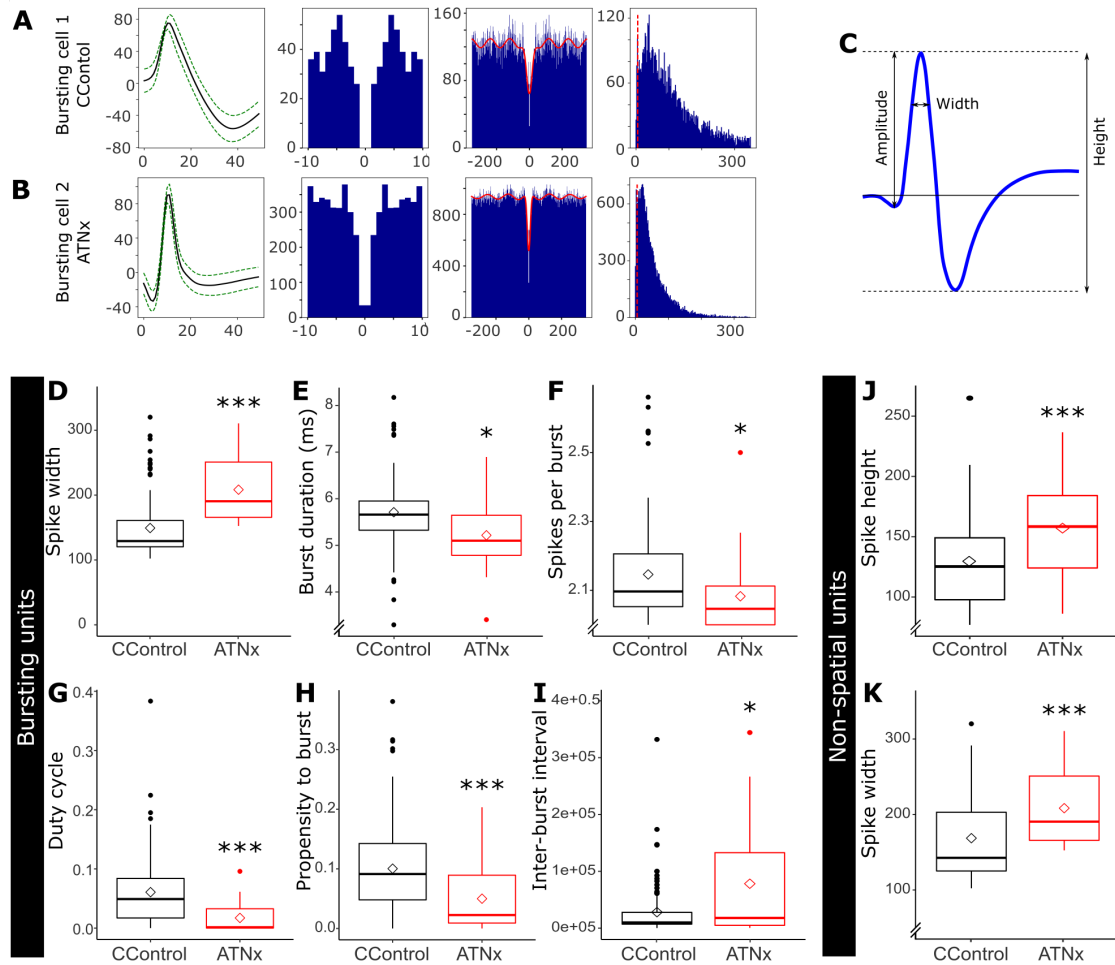


FIGURE 4.4: Significant differences in bursting and nonspatial cell properties between control and ATNx rats.

(A, B) Waveforms and autocorrelation histograms were used for cell classification.

(C) Diagram of waveform properties shown in the boxplots.

(D – I) Boxplots of significant differences between bursting units recorded from control and ATNx rats.

(J, K) Boxplots of significant spike width and height differences between non-spatial units recorded from control and ATNx rats.

Key: In the boxplot, the midline represents the median, while the outer lines represent the inter-quartile range (IQR), with the whiskers extending up/down to the furthest data point less than within $1.5 \times (IQR)$ from the third/first quartile. Points outside the whiskers are deemed outliers and are shown as filled circles. Unfilled diamonds represent the mean. * indicates $p < 0.05$, *** indicates $p < 0.001$; Welch's two-sample t-test for normal data or the Mann-Whitney U test otherwise.

Measurement	Control	Lesion
Spike rate (spikes/s)	2.93 ± 4.31	5.42 ± 7.07
Spike amplitude (μV)	106.91 ± 31.15	104.33 ± 19.20
Spike height (μV)	146.91 ± 53.96	161.18 ± 39.32
Spike width (μs) ^{***}	155.26 ± 50.43	213.88 ± 53.87
Inter-spike interval (s)	1.01 ± 1.38	2.25 ± 3.45
Mean bursting inter-spike interval (ISI) (ms)	4.08 ± 0.32	3.88 ± 0.28
Spikes per burst [*]	2.13 ± 0.11	2.05 ± 0.05
Mean burst duration (ms) [*]	5.61 ± 0.64	5.07 ± 0.45
Inter-burst interval (s) [*]	11.14 ± 11.10	31.83 ± 39.50
Mean duty cycle	0.05 ± 0.05	0.02 ± 0.02
Propensity to burst ^{***}	0.12 ± 0.09	0.06 ± 0.03

TABLE 4.2: Summary of spike and burst properties of subicular units from control and lesioned rats. Values indicate mean plus/minus standard deviation. The top section of the table represents results for all cells recorded, while the bottom section of the table represents results for bursting cells recorded.

key: ^{*} $p < 0.05$, ^{***} $p < 0.001$; Welch's two-sample t test. Propensity to burst is the number of bursting spikes divided by the total number of spikes. Duty cycle is the ratio between the burst duration and the Inter-burst interval.

Comparing spike properties and inter-spike interval distributions

First, we consider all recorded subicular cells pooled into a single group (Table 4.2, top section). The mean spike width was greater in ATNx rats than control rats (control $155.26 \pm 50.43\mu\text{s}$, ATNx $213.88 \pm 53.87\mu\text{s}$; $W = 310$, $p < 0.001$, Mann-Whitney U test). Other spike properties showed no differences between groups, including spike rate (control $2.93 \pm 4.31\text{spikes/s}$, ATNx $5.42 \pm 7.07\text{spikes/s}$; $W = 809$, $p = 0.673$, Mann-Whitney U test), spike amplitude (control $107 \pm 31\mu\text{V}$, ATNx $104 \pm 19\mu\text{V}$; $t_{(49.3)} = 0.47$, $p = 0.642$), and spike height (control $147 \pm 54\mu\text{V}$, ATNx $161 \pm 39\mu\text{V}$; $t_{(40.74)} = -1.34$, $p = 0.187$). However, while the mean ISI was not statistically different between the groups, the median ISI can be a fairer comparison due to the skewed nature of ISI distributions. We tested the ISI distributions for normality using the Shapiro-Wilk test, which indicated these data are not normally distributed, supporting the notion that they may be skewed. The median ISI in control had a mean value of 153ms versus ATNx of 547ms, respectively (Mann-Whitney $W = 386$, $p = 0.010$, two-tailed), which was significantly lower. Both the median ISI and mean ISI are lower in control rats than ATNx rats, despite the mean spike rate being (non-significantly) lower in control rats. This is likely due to a combination of the presence of a small number of cells with very high firing rates in ATNx rats and a more consistent firing rate in cells from control rats.

Collating spike-width distributions across non-spatial units

When non-spatial cells in control rats were compared with putative non-spatial cells in ATNx rats, spike width remained significantly different (control $173 \pm 60\mu\text{s}$, ATNx $213 \pm 54\mu\text{s}$; $t_{(45.60)} = -2.56$, $p = 0.014$, Welch's two-sample t-test, Figure 4.4K), but spike height also differed (control $133 \pm 40\mu\text{V}$, ATNx $161 \pm 39\mu\text{V}$; $t_{(42.11)} = -2.57$, $p = 0.014$, Welch's two-sample t-test, Figure 4.4J). Spike width was assessed further using histograms to determine whether the difference observed between cells from control rats and ATNx rats reflected recording two different neuron populations (Figure 4.5). More cells with short-duration waveforms were recorded from control rats than ATNx, but the samples obtained from the two groups were not distinct (Figure 4.5A). Cells with narrow spike widths were more often spatial, but the classification of narrow and wide waveform cells was mixed in both control and ATNx combined (Figure 4.5D) and control only (Figure 4.5C).

Increased rate of burst firing in control rats

As no spatial units were recorded in ATNx rats, it was unclear whether spatial cells were present but inhibited (and therefore not recorded) or whether the units that were recorded were latent spatial units that were now not responding to "spatial" inputs. Because of the difficulties in comparing non-spatial units in the control to unknown spatial or non-spatial units in ATNx rats, subicular cells were classified according to their spike properties into bursting, fast-spiking, and theta-entrained cells (Anderson & O'Mara, 2003). While the percentage of subicular bursting cells in control (53% of 82 units) and ATNx (57% of 21 units) rats was essentially equivalent, other properties differed (Table 4.2, bottom section). Cells in control rats showed a greater propensity to burst than those in ATNx rats (control 0.12 ± 0.09 , ATNx 0.06 ± 0.03 ; $t_{(35.46)} = 3.69$, $p < 0.001$). Additionally, bursting cells in control rats showed more spikes per burst (control 2.13 ± 0.11 , ATNx 2.05 ± 0.05 ; $W = 1005$, $p = 0.004$, Mann-Whitney U test) and greater burst duration (control $5.61 \pm 0.64\text{ms}$, ATNx $5.07 \pm 0.45\text{ms}$; $t_{(18.21)} = 3.04$, $p = 0.007$).

4.3.3 The effects of inactivation via muscimol mimic permanent ATN lesions

Our findings from permanent NMDA lesions (ATNx) were further reflected in the muscimol temporary inhibition experiments. Inactivation of the ATN via muscimol infusion into the ATN resulted in a significant reduction in the spatial responsivity of subicular spatial cells (Figure A.3) despite no change in firing rate. However, these rats

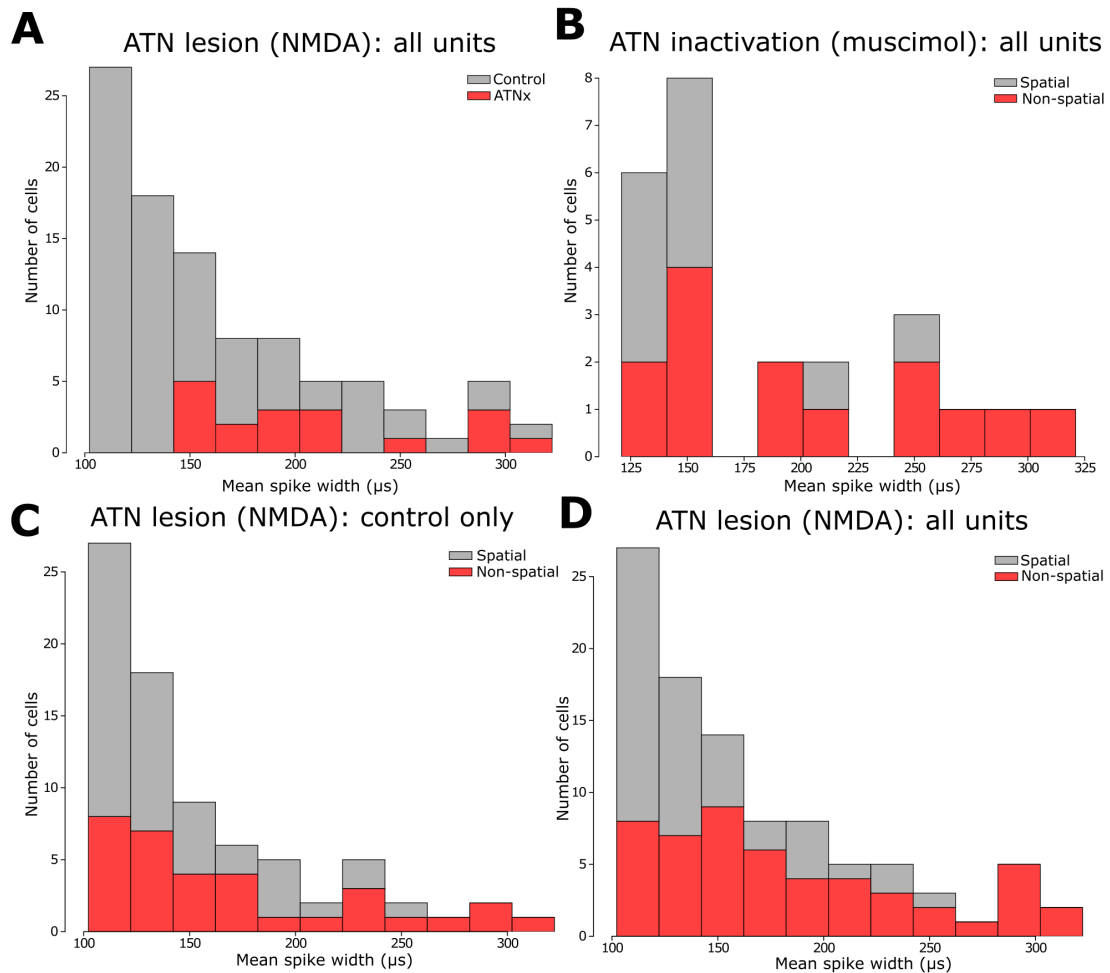


FIGURE 4.5: Comparison of spike widths in control and lesioned rats.

(A) When control rats are compared with rats with permanent NMDA lesions of the ATN (ATNx), more short-duration waveforms were recorded in controls, but the samples obtained from the two groups of rats are not distinct.

(B) Similarly, in rats where the ATN was inactivated with muscimol, narrow waveform cells were more often spatial, but the groups were mixed.

(C, D) In the permanent lesion study, narrow waveform cells were more often spatial, but the classification of wide and narrow waveform cells was mixed. C shows data for only control rats, while D shows combined data for the control and ATNx rats.

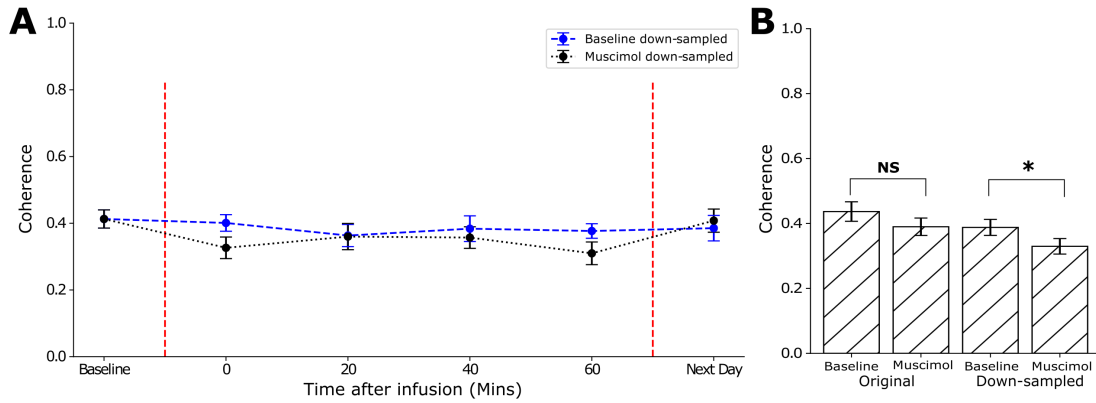


FIGURE 4.6: Results from spatially-downsampling the subicular spatial units considered in the temporary muscimol ATN lesion experiments.

(A) The blue dashed line shows average coherence in the baseline recording spatial/spike data, but spatially downsampled to match the coverage of the arena in the corresponding muscimol recording at that time. The black dotted line shows average coherence for the corresponding muscimol recording, spatially downsampled to match the baseline recording arena coverage. Error bars represent standard error of the mean.

(B) Bar chart representing the difference in coherence with and without downsampling. The bars on the left show the average coherence values from baseline recordings of subicular spatial units and muscimol recordings without any downsampling. The bars on the right show the average coherence values after the downsampling procedure. NS indicates $p \geq 0.05$, * indicates $p < 0.05$ (Welch's Two Sample t-test).

had reduced levels of movement around the time of heightened muscimol effects, which may have contributed to the reduction in spatial responsivity. We thus performed a control measure by spatially downsampling the firing maps before and after muscimol injection to match occupancy in each spatial bin (Figure 4.6). Furthermore, a histogram of spike widths showed that, while cells with narrow waveforms were more often spatial, the spatial and non-spatial groups remained mixed (Figure 4.5B). This suggests that the differences observed were not due to recording different neuron populations in the two groups. Overall, these results suggest that temporary inactivation of the ATN via muscimol infusion into the ATN mimics the effects of permanent ATN lesions on spatial processing in the subiculum (see subsection A.1.4 for detailed muscimol results).

4.3.4 Place cells are recorded in CA1 despite ATN lesions

To explore the effects of anterior thalamic lesions on spatial processing in the dorsal hippocampus, a further 3 rats were implanted with a microdrive apparatus of 8 recording tetrodes into the dorsal CA1 and received bilateral, permanent (NMDA) lesions of the ATN (see Figure 4.7A, B for lesion extent). Behavioural effects of the lesions were similar to those observed in rats with lesions of the entire ATN with subiculum

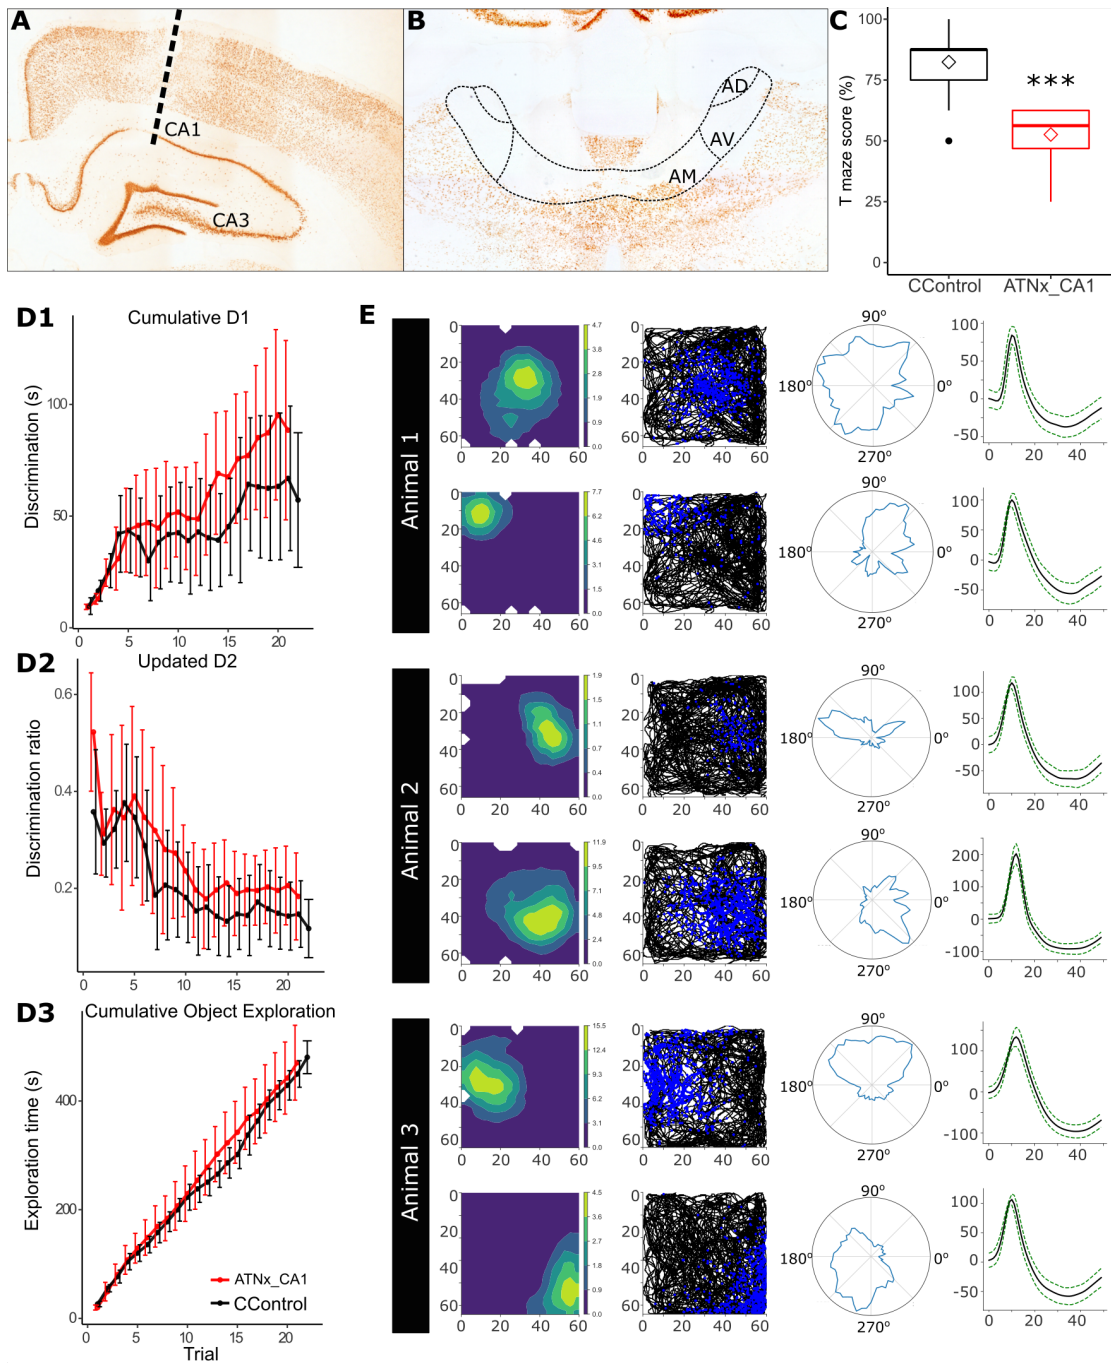


FIGURE 4.7: The effect of ATN lesions on CA1 cell recordings.

(A) Representative electrode placement in CA1.

(B) Representative ATN lesion in NeuN-reacted sections.

(C) Animals with ATN lesions and electrodes implanted in CA1 showed a significant deficit in spatial alternation tasks compared with control animals (control data is reused from the ATN lesion data with subicular recordings).

(D1 – D3) The same cohort of ATNx animals showed no deficit in object recognition in the bow-tie maze.

(E) Representative place cells recorded from CA1 in 3 ATNx animals.

Key: *** indicates $p < 0.001$, Welch's two-sample t-test.

recordings, namely reduced spatial alternation performance, but the same exploration rate (Figure 4.7C, D). Daily recordings were conducted to measure unit activity in the CA1 of ATNx rats performing a pellet-chasing task in an arena. From these rats, 203 well-isolated units were recorded from dorsal CA1. Units were further classified into 107 spatial units (following Schoenenberger et al., 2016) using sparsity and coherence criteria (Rat 1, $n = 12$; Rat 2, $n = 67$; Rat 3, $n = 28$). Putative interneurons were considered to have been recorded from dorsal CA1 if they were recorded on the same tetrode and in the same recording session as a spatial unit. Despite the absence of spatial signals in the subiculum, hippocampal (CA1) place cells appeared intact (Figure 4.7E), and place cell and interneuron properties are shown in Table A.2.

4.4 Discussing the possibility of multiple spatial circuits and the consequences for diencephalic amnesia

The usual heterogeneous population of subicular cells, including place, head-direction, grid, and border cells, is strikingly absent after permanent ATN lesions. Furthermore, the effects of these lesions are specific rather than generalised, as performance on a spatial memory task declines significantly while non-hippocampal object recognition memory remains intact. Remarkably, infusion of muscimol into the ATN of healthy rats causes similar, but temporary effects on behavioral performance and cellular firing. This provides strong evidence for the crucial role of ATN projections, both direct and indirect, in subicular spatial cellular discharge and spatial alternation memory.

Despite the extensive ATN lesions, the place cells in CA1 of the hippocampus remained intact (see also, Calton et al., 2003). While it might be assumed that subicular spatial cells are primarily driven by CA1 inputs due to the heavy innervation of the subiculum by CA1 and the activity-dependent plasticity observed in this projection (Amaral et al., 1991; Cembrowski et al., 2018; Commins & O'Mara, 2000; Commins et al., 1998a, 1998b; O'Mara et al., 2001, 2009), our results suggest otherwise. CA1 projections alone are not sufficient to support subicular spatial firing or spatial alternation memory. Instead, ATN projections, both direct and indirect, play a crucial role in these processes.

This raises an intriguing question: why is spatial coding in CA1 sufficiently independent from direct and indirect ATN inputs, while subicular spatial coding is not? The subiculum, with its dense direct and indirect projections to the ATN, is part of a complex reciprocal network of pathways (Bubb et al., 2017). Perhaps the ATN normally exert a direct modulatory, including oscillatory (Vertes et al., 2001), influence on the subiculum that, when removed, leads to changes in gain control, disrupting information processing. Considering the range of spatial cells in the ATN (Jankowski et al., 2015; Matulewicz

et al., 2019; Taube, 1995), such an input might, for example, help hippocampal and parahippocampal regions coregister their various spatial signals, including those from CA1. This might explain why ATN lesions increased dorsal subiculum spike width and reduced bursting properties, but left other features seemingly intact, such as overall firing rate and spike amplitude.

Subicular neurons are well separated into bursting, non-bursting, and fast-spiking groups, each associated with distinct burst firing patterns (Anderson & O'Mara, 2003; Sharp & Green, 1994; Stewart & Wong, 1993). These patterns may reflect functional differences among the cells, such as sparsely bursting subicular cells potentially carrying more spatial information (Simonnet & Brecht, 2019). Evidence also shows that subicular neurons switching between bursting and single spiking modes are related to the inactivation of sodium channels in these pyramidal neurons, contributing to further network dynamics (Cooper et al., 2005). Here, although the proportion of bursting cells remained equivalent between control and ATNx groups, several properties differed in line with a reduction in spatial coding as cells from controls showed a greater burst propensity, more spikes per burst, and longer burst duration. This alteration in bursting may affect the fidelity of subicular-retrosplenial transmission (see further investigation of this in chapter 5) as bursting cells in dorsal subiculum with direct connections to granular RSC directly impact sharp-wave ripples in RSC (Nitzan et al., 2020). The ATN may therefore be critical for the maintenance of subicular bursting properties, which in turn may be important for the transmission of spatial information to the RSC.

While the origins of diencephalic amnesia have remained somewhat mysterious for over a century, extensive research has been conducted on hippocampal amnesia since the 1950s (Arts et al., 2017; Scoville & Milner, 1957). Clinically, diencephalic amnesia presents as a dense anterograde amnesic syndrome similar to that associated with bilateral hippocampal damage (Arts et al., 2017). The ATN emerges as a critical structure in diencephalic amnesia (Harding et al., 2000; Kril & Harper, 2012), and our study supports a wider network account of ATN loss involving direct efferent actions on the subiculum and indirect actions via parahippocampal and retrosplenial targets. The reciprocal relationship between the ATN and the subiculum, as well as their connections with other regions like the dorsal hippocampus and parahippocampal areas, including the entorhinal cortex, is key to understanding the impact of ATN damage on diencephalic amnesia. Moving forward, it is crucial to delve deeper into the computations performed by the subiculum on the inputs it receives from the ATN to support spatial alternation performance. Additionally, understanding how the substantial input from CA1 is gated by ATN inputs reaching the subiculum is of utmost importance.

Our study highlights the critical roles of the ATN in navigation and spatial learning,

with a particular focus on the pivotal role of the subiculum. The severity of the spatial memory/alternation deficit despite preservation of CA1 place cells underscore the importance of the subiculum as a principal source of hippocampal projections beyond the temporal lobe. These findings shed light on how damage to the anterior thalamic nuclei can indirectly affect other brain regions involved in spatial processing, such as the mammillary bodies, ventral striatum, and medial and orbital prefrontal cortices (Aggleton & Nelson, 2015; Ferguson et al., 2019; Haugland et al., 2019). The data also suggest that subicular spatial signals arise from the convergence of inputs from CA1, the ATN, and parahippocampal areas, including the entorhinal cortex. In conclusion, the loss of anterior thalamic influence on the output regions of the hippocampal formation, particularly the subiculum, is a key factor in diencephalic amnesia. Further research is necessary to unravel the intricacies of this complex network and its implications for spatial memory and navigation.

4.5 Experimental procedures and analysis methods

4.5.1 Summary of experiment details and unit processing

The appendix (section A.1) details the ethics, subject details, experimental design, and histology performed. To briefly summarise here, 28 male lister hooded rats were implanted with 32 electrodes unilaterally, with electrodes targeting a combination of the subiculum, retrosplenial cortex, and CA1. Of the rats implanted, 1 ATNx rat was subsequently excluded because of electrode malfunction, and 1 normal control was excluded because of postsurgical complications. Furthermore, comparisons between sham control groups and normal control groups consistently failed to reveal group differences, and for this reason, they were typically combined. Rats were lesioned permanently in the ATN via NMDA excitotoxicity, or temporarily inactivated in the ATN via muscimol infusion, or kept as sham or normal controls. Rats were then recorded during free-moving exploration, T-maze spatial alternation, and novel object interaction in a bow-tie maze. Table A.1 shows the number of rats in each type of recording layout.

See subsection A.1.10 for a full description of the single-unit sorting and selection process from putative action potentials recorded from these rats. Standard neural data analysis and plots on single-units was performed using NeuroChaT, which was expanded in functionality for this study (Islam et al., 2019). Datasets generated during this study are available at OSF <https://osf.io/vdakx>, and code is available at GitHub <https://github.com/shanemomara/omaraneurolab>. An updated version of this dataset

with LFPs in addition to spike data is described in chapter 5.

4.5.2 Performing downsampling analysis to minimise effects of exploration reduction in muscimol affected rats

We performed a spatial downsampling procedure (based on Boccarda et al., 2019) as a control method to rule out the possibility that the disruption in spatial firing merely reflects a lack of sampling of the environment, given that the rats move less distance and cover less of the environment while under the effects of muscimol. For each cell, we produce a list $L = (x_t, y_t, spike_t)$ where x_t, y_t is the position of the animal at time t and $spike_t$ is the number of spikes the cell emitted in that time bin. The procedure to spatially downsample data from recording A to match the exploration in recording B is as follows. First, the list L from B is binned into 3cm squares based on x_t, y_t and the list L from A is binned using the same number of bins as was used for B . Second, several random samples are drawn (with repetition) from the list L in A such that the number of samples drawn from an individual bin is the minimum of the number of data points falling in that bin between recordings A and B . This procedure is performed 200 times for each cell.

Using the method above, recordings were spatially downsampled to match their own spatial occupancy as an additional control since the random sampling procedure should cause minor changes. After testing spatial coherence, spatial information content, and spatial sparsity, coherence was the only measure tested that was resistant to downsampling against self (mean decrease of 0.015) and downsampling to other randomly generated spatial data with high occupancy of the arena (mean decrease of 0.05). For each cell considered in the muscimol experiments, the baseline recording was downsampled to match the muscimol recordings, and the muscimol recordings were downsampled to match the baseline recording. These resulting firing maps had closely matching occupancy of the environment, and coherence was computed on these maps and compared (Figure 4.6).

4.5.3 Computing Bayesian statistics to assess the difference in frequency of spatial cells

To investigate whether the apparent absence of spatial signal in ATNx rats was significant, a Bayesian approach was applied as follows (MacKay, 2003). Let PC denote the probability that a control recording will contain at least one spatial cell, and PL denote the probability that a lesion recording will contain at least one spatial cell. Let D denote the recorded data.

By Bayes' theorem, $P(PC, PL|D) = P(D|PC, PL) \cdot P(PC, PL)/P(D)$ (i.e., posterior = likelihood \times prior / probability of the evidence). We use a uniform prior distribution, as we have no prior belief of how common recordings in the subiculum with spatial cells are, so $P(PC, PL) = 1$. The probability of the evidence $P(D)$ is a constant normalising parameter, which can be calculated by ensuring the posterior distribution sums to 1. By the independence assumption of selected recordings, the likelihood function can be modelled as the product of evaluating the probability mass functions arising from two binomial distributions as follows:

$$P(D|PC, PL) = \binom{CR}{CS} PC^{CS} (1 - PC)^{CR-CS} \binom{LR}{LS} PL^{LS} (1 - PL)^{LR-LS} \quad (4.1)$$

where:

- D = the data that is observed
- CR = the number of control recordings performed
- CS = the number of control recordings with a spatial cell
- LR = the number of lesion recordings
- LS = the number of lesion recordings with a spatial cell

We then apply this with our recorded data D and note the observations in a contingency table, which allows us to make comparisons between PC and PL , and contrast this with PC^* and PL^* , the probability of recording non-spatial cells in control and lesion Table 4.1. For example, using the contingency table:

$$P(D|PC, PL) = \binom{53}{15} PC^{15} (1 - PC)^{53-15} \binom{47}{0} PL^0 (1 - PL)^{47-0} \quad (4.2)$$

4.5.4 Calculating bursting cells and spatial coding

A burst was defined as a series of at least spikes in which each ISI was a maximum of 6ms. Bursts were excluded if there was not a minimum of 50ms since the last bursting event (Anderson & O'Mara, 2003). Further bursting analyses examined the total number of bursts during a recording session, the number of spikes in the bursting cluster, mean ISI during the burst cluster, the number of spikes per burst, burst duration, duty cycle (the portion of an inter-burst interval (IBI) during which a burst fires), the IBI, and propensity to burst, calculated by dividing the number of bursting spikes by the total number of spikes in a recording.

Multiple indices were used to analyse the spatial properties of unit activity; spatial coherence, spatial information content, and spatial sparsity. Spatial coherence represents

the similarity between the obtained firing map and a smoothed version of the firing map (boxcar smoothing) with a unitless value between 0 and 1. A high coherence indicates that the firing rate of each pixel of the rate map can be well predicted from its neighbours, thus indicating consistency in the response. Spatial information content was expressed in bits per spike (Skaggs et al., 1993) and represents how much bits of information each spike provides about the location of the rat in the arena. Spatial sparsity calculates the fraction of the environment in which a unit is active. Mean spiking frequency is the total number of spikes divided by the total recording time and is expressed in spikes per second. A place field was identified if nine neighbouring pixels (sharing a side) were $\geq 20\%$ of the peak firing rate. Additionally, to be regarded as place cells, the following criteria had to be met: a spatial information content index of ≥ 0.5 , a spatial coherence of ≥ 0.25 , and a mean firing rate of ≥ 0.25 .

References

- Aggleton, J. P., & Nelson, A. J. D. (2015). Why do lesions in the rodent anterior thalamic nuclei cause such severe spatial deficits? *Neuroscience and Biobehavioral Reviews*, *54*, 131–144.
- Aggleton, J. P., Nelson, A. J. D., & O'Mara, S. M. (2022). Time to retire the serial papez circuit: Implications for space, memory, and attention. *Neuroscience & Biobehavioral Reviews*, *140*, 104813.
- Aggleton, J. P., & O'Mara, S. M. (2022). The anterior thalamic nuclei: Core components of a tripartite episodic memory system. *Nature Reviews Neuroscience*, *23*(8), 505–516.
- Amaral, D. G., & Witter, M. P. (1989). The three-dimensional organization of the hippocampal formation: A review of anatomical data. *Neuroscience*, *31*(3), 571–591.
- Amaral, D. G., Dolorfo, C., & Alvarez-Royo, P. (1991). Organization of CA1 projections to the subiculum: A PHA-l analysis in the rat. *Hippocampus*, *1*(4), 415–435.
- Anderson, M. I., & O'Mara, S. M. (2003). Analysis of recordings of single-unit firing and population activity in the dorsal subiculum of unrestrained, freely moving rats. *Journal of Neurophysiology*, *90*(2), 655–665.
- Arts, N. J. M., Walvoort, S. J. W., & Kessels, R. P. C. (2017). Korsakoff's syndrome: A critical review. *Neuropsychiatric Disease and Treatment*, *13*.
- Barnard, G. (1947). Significance tests for 2x2 tables. *Biometrika*, *34*(1), 123–138.
- Boccaro, C. N., Nardin, M., Stella, F., O'Neill, J., & Csicsvari, J. (2019). The entorhinal cognitive map is attracted to goals. *Science*, *363*(6434), 1443–1447.

- Brotons-Mas, J. R., Schaffelhofer, S., Guger, C., O'Mara, S. M., & Sanchez-Vives, M. V. (2017). Heterogeneous spatial representation by different subpopulations of neurons in the subiculum. *Neuroscience*, *343*, 174–189.
- Brotons-Mas, J. R., Montejo, N., O'Mara, S. M., & Sanchez-Vives, M. V. (2010). Stability of subicular place fields across multiple light and dark transitions. *European Journal of Neuroscience*, *32*(4), 648–658.
- Bubb, E. J., Kinnavane, L., & Aggleton, J. P. (2017). Hippocampal-diencephalic-cingulate networks for memory and emotion: An anatomical guide. *Brain and neuroscience advances*, *1*, 2398212817723443.
- Calton, J. L., Stackman, R. W., Goodridge, J. P., Archey, W. B., Dudchenko, P. A., & Taube, J. S. (2003). Hippocampal place cell instability after lesions of the head direction cell network. *Journal of Neuroscience*, *23*(30), 9719–9731.
- Cembrowski, M. S., Wang, L., Lemire, A. L., Copeland, M., DiLisio, S. F., Clements, J., & Spruston, N. (2018). The subiculum is a patchwork of discrete subregions (L. Colgin & G. L. Westbrook, Eds.). *eLife*, *7*, e37701.
- Commins, S., Gigg, J., Anderson, M., & O'Mara, S. M. (1998a). Interaction between paired-pulse facilitation and long-term potentiation in the projection from hippocampal area CA1 to the subiculum. *Neuroreport*, *9*(18), 4109–4113.
- Commins, S., Gigg, J., Anderson, M., & O'Mara, S. M. (1998b). The projection from hippocampal area CA1 to the subiculum sustains long-term potentiation. *Neuroreport*, *9*(5), 847–850.
- Commins, S., & O'Mara, S. M. (2000). Interactions between paired-pulse facilitation, low-frequency stimulation, and behavioral stress in the pathway from hippocampal area CA1 to the subiculum: Dissociation of baseline synaptic transmission from paired-pulse facilitation and depression of the same pathway. *Psychobiology*, *28*(1), 1–11.
- Cooper, D. C., Chung, S., & Spruston, N. (2005). Output-mode transitions are controlled by prolonged inactivation of sodium channels in pyramidal neurons of subiculum. *PLOS Biology*, *3*(6), e175.
- Ferguson, M. A., Lim, C., Cooke, D., Darby, R. R., Wu, O., Rost, N. S., Corbetta, M., Grafman, J., & Fox, M. D. (2019). A human memory circuit derived from brain lesions causing amnesia. *Nature communications*, *10*(1), 1–9.
- Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, *436*(7052), 801–806.
- Harding, A., Halliday, G., Caine, D., & Kril, J. (2000). Degeneration of anterior thalamic nuclei differentiates alcoholics with amnesia. *Brain*, *123*(1), 141–154.
- Haugland, K. G., Sugar, J., & Witter, M. P. (2019). Development and topographical organization of projections from the hippocampus and parahippocampus to the retrosplenial cortex. *European Journal of Neuroscience*, *50*(1), 1799–1819.

- Islam, M. N., Martin, S. K., Aggleton, J. P., & O'Mara, S. M. (2019). NeuroChaT: A toolbox to analyse the dynamics of neuronal encoding in freely-behaving rodents in vivo. *Wellcome Open Research*, *4*, 196.
- Jankowski, M. M., Passecker, J., Islam, M. N., Vann, S., Erichsen, J. T., Aggleton, J. P., & O'Mara, S. M. (2015). Evidence for spatially-responsive neurons in the rostral thalamus. *Frontiers in Behavioral Neuroscience*, *9*, 256.
- Jankowski, M. M., Ronnqvist, K. C., Tsanov, M., Vann, S. D., Wright, N. F., Erichsen, J. T., Aggleton, J. P., & O'Mara, S. M. (2013). The anterior thalamus provides a subcortical circuit supporting memory and spatial navigation. *Frontiers in Systems Neuroscience*, *7*, 45.
- Kril, J. J., & Harper, C. G. (2012). Neuroanatomy and neuropathology associated with korsakoff's syndrome. *Neuropsychology review*, *22*(2), 72–80.
- MacKay, D. J. (2003). *Information theory, inference and learning algorithms*. Cambridge university press.
- Matulewicz, P., Ulrich, K., Islam, M. N., Mathiasen, M. L., Aggleton, J. P., & O'Mara, S. M. (2019). Proximal perimeter encoding in the rat rostral thalamus. *Scientific Reports*, *9*(1), 2865.
- Moser, E. I., Kropff, E., & Moser, M.-B. (2008). Place cells, grid cells, and the brain's spatial representation system. *Annual Review of Neuroscience*, *31*, 69–89.
- Nitzan, N., McKenzie, S., Beed, P., English, D. F., Oldani, S., Tukker, J. J., Buzsáki, G., & Schmitz, D. (2020). Propagation of hippocampal ripples to the neocortex by way of a subiculum-retrosplenial pathway. *bioRxiv*, 2020.02.27.966770.
- O'Keefe, J., & Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature*, *381*(6581), 425–8.
- O'Mara, S. M., & Aggleton, J. P. (2019). Space and memory (far) beyond the hippocampus: Many subcortical structures also support cognitive mapping and mnemonic processing. *Frontiers in Neural Circuits*, *13*.
- O'Mara, S. M., Commins, S., Anderson, M., & Gigg, J. (2001). The subiculum: A review of form, physiology and function. *Progress in Neurobiology*, *64*(2), 129–155.
- O'Mara, S. M., Sanchez-Vives, M. V., Brotons-Mas, J. R., & O'Hare, E. (2009). Roles for the subiculum in spatial information processing, memory, motivation and the temporal control of behaviour. *Progress in Neuro-Psychopharmacology and Biological Psychiatry*, *33*(5), 782–790.
- Schoenenberger, P., O'Neill, J., & Csicsvari, J. (2016). Activity-dependent plasticity of hippocampal place maps. *Nature Communications*, *7*(1), 11824.
- Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of neurology, neurosurgery, and psychiatry*, *20*(1), 11.
- Sharp, P. E., & Green, C. (1994). Spatial correlates of firing patterns of single cells in the subiculum of the freely moving rat. *Journal of Neuroscience*, *14*(4), 2339–2356.

- Shibata, H. (1993). Direct projections from the anterior thalamic nuclei to the retrohippocampal region in the rat. *Journal of Comparative Neurology*, *337*(3), 431–445.
- Simonnet, J., & Brecht, M. (2019). Burst firing and spatial coding in subicular principal cells. *Journal of Neuroscience*, *39*(19), 3651–3662.
- Skaggs, W. E., McNaughton, B. L., & Gothard, K. M. (1993). An information-theoretic approach to deciphering the hippocampal code. In S. J. Hanson, J. D. Cowan, & C. L. Giles (Eds.), *Advances in neural information processing systems 5* (pp. 1030–1037). Morgan-Kaufmann.
- Smith, C. N., Frascino, J. C., Hopkins, R. O., & Squire, L. R. (2013). The nature of anterograde and retrograde memory impairment after damage to the medial temporal lobe. *Neuropsychologia*, *51*(13), 10.1016/j.neuropsychologia.2013.09.015.
- Staff, N. P., Jung, H.-Y., Thiagarajan, T., Yao, M., & Spruston, N. (2000). Resting and active properties of pyramidal neurons in subiculum and CA1 of rat hippocampus. *Journal of Neurophysiology*, *84*(5), 2398–2408.
- Stewart, M., & Wong, R. K. (1993). Intrinsic properties and evoked responses of guinea pig subicular neurons in vitro. *Journal of Neurophysiology*, *70*(1), 232–245.
- Sugar, J., Witter, M., van Strien, N., & Cappaert, N. (2011). The retrosplenial cortex: Intrinsic connectivity and connections with the (para)hippocampal region in the rat. an interactive connectome. *Frontiers in Neuroinformatics*, *5*.
- Taube, J. S. (1995). Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *15*(1), 70–86.
- Tsanov, M., Chah, E., Vann, S. D., Reilly, R. B., Erichsen, J. T., Aggleton, J. P., & O'Mara, S. M. (2011). Theta-modulated head direction cells in the rat anterior thalamus. *The Journal of Neuroscience*, *31*(26), 9489.
- Vertes, R., Albo, Z., & Di Prisco, G. V. (2001). Theta-rhythmically firing neurons in the anterior thalamus: Implications for mnemonic functions of papez's circuit. *Neuroscience*, *104*(3), 619–625.
- Viejo, G., & Peyrache, A. (2019). Precise coupling of the thalamic head-direction system to hippocampal ripples. *bioRxiv*, 809657.

5 The effect of anterior thalamic lesions on local field potentials

Introduction: The anterior thalamic nuclei (ATN) are vital for spatial memory in rodents and appear necessary for episodic memory in humans. We previously showed how ATN lesions abrogate spatial coding in subicular cells. Here, we further elucidate these findings by investigating the relationship between the local field potential (LFP) and single-unit activity in the dorsal subiculum and the LFP in ipsilateral retrosplenial cortex (RSC) post-lesion of the ATN.

Methods: Electrophysiological recordings were performed in rats during free-exploration and spatial alternation tasks with permanent ATN lesions via N-Methyl-D-aspartic acid (NMDA) and temporary ATN inhibition via muscimol to obtain single-unit measurements and LFPs. We primarily focused on analysing theta rhythm activity (related to spatial coding in single units), beta rhythm activity (linked to learning and novelty detection), and sharp-wave ripples during sleep (linked to memory and replay) from LFPs.

Results: Lesion effects on subicular LFPs were unexpectedly subtle, given the gross behavioural impairments present post-lesion. Theta and beta rhythms, sharp-wave ripples during sleep, and the relationship between motor activity and theta rhythm activity were all maintained in the subiculum post-lesion. However, power in RSC LFPs was reduced outside of the theta band post-lesion, the coherence between LFPs from RSC and subiculum was reduced in T-maze trials post-lesion (but less so in open-field exploration), and the coupling between the subicular LFP theta phase and single unit-spiking activity was reduced post-lesion.

Conclusions: We suggest that a key contributor to the behavioural impairments post-lesion may stem from a change in the transfer or representation of information between regions in the hippocampal circuit following anterior thalamic impairment.

5.1 My contributions to this study

This is a multi-author study and I did not perform any of the electrophysiology or histology in this study, which were performed by Dr. Frost and Dr. Cafalchio. The software development and analysis was performed collaboratively with Dr. Cafalchio, who helped to form a data table describing all of the experiments, and assisted in manually splitting out the t-maze trials into blocks. I led the dry-lab aspects of this study, including data processing, analysis, results interpretation, and direction.

5.2 Using local field potentials to expand upon the anterior thalamic nuclei lesion findings

The anterior thalamic nuclei (ATN) are a vital part of the extended neural system for memory, spatial learning, and navigation. ATN lesions in rodents severely impair a wide range of spatial tasks (Aggleton & Nelson, 2015). At the same time, electrophysiological studies reveal a range of spatial cells in the ATN, including head-direction cells (Taube, 1995; Taube et al., 1990) place, and perimeter cells (Jankowski et al., 2015; Matulewicz et al., 2019; O'Mara & Aggleton, 2019). In addition, ATN dysfunction is believed to be a core feature of diencephalic amnesia in humans (Aggleton & Brown, 1999; Carlesimo et al., 2011; Danet et al., 2015) highlighting its potential significance for episodic memory. The observed lesion effects and electrophysiological properties bring into question the interactions of the ATN with the hippocampal formation and retrosplenial cortex (RSC); two structures also vital for rodent spatial processing (Keene & Bucci, 2008; Todd & Bucci, 2015), with the hippocampal formation containing an array of spatial cells, and the presence of head direction cells in the retrosplenial cortex (Cho & Sharp, 2001; Moser et al., 2017). Additionally, both structures have close links with anterograde amnesic syndromes (Maguire, 2001; Spiers et al., 2001; Vann et al., 2009).

In this study, we wish to further elucidate these findings by investigating the relationship between local field potentials (LFPs) in the dorsal subiculum, ipsilateral RSC, and ipsilateral CA1 (relative to the subicular recording site), and single-unit activity (termed, “spikes”, meaning putative action potentials detected from a single or multiple cells during electrophysiological recordings) in the dorsal subiculum. These are investigated in rats that were controls, had ATN lesions via NMDA, or had temporary ATN inhibition via muscimol infusion.

We have previously described our electrophysiological recordings performed in a circuit involving the ATN, CA1, RSC, and the subiculum in chapter 4. To recap the

connections in this circuit, the subiculum is reciprocally connected as the primary output of CA1, and projections are organised such that proximal CA1 projects to distal subiculum (Amaral et al., 1991). The subiculum sends direct projections to both granular and agranular RSC, while CA1 only projects to granular RSC (Sugar et al., 2011). In return, the subiculum receives rather scant afferent projections from granular RSC, while CA1 does not receive any direct connections from RSC. The ATN projects to the dorsal subiculum and are densely reciprocally connected with the RSC (Bubb et al., 2018; Shibata, 1993; Sripanidkulchai & Wyss, 1986). The present study, therefore, examined the impact of both permanent and transient ATN lesions in adult male rats on LFPs simultaneously recorded from the retrosplenial cortex, subiculum and CA1, to examine interactions between these anatomically connected and functionally linked regions.

To draw out the effect of ATN lesions on this circuit, oscillations in LFPs and coherence between LFPs were studied during open-field exploration, sleep, and T-maze trials in control, ATN lesion, and ATN inactivated rats. We begin by investigating the impact of ATN lesions on rhythms in LFPs during open-field exploration. Then we shift focus to investigate potential correlates with the behaviour change in ATNx rats, such as increased motor activity and performing close to chance levels in spatial alternation T-maze tasks. Sharp wave ripples (150 – 210 Hz) were investigated during sleep, as hippocampal ripples are highly linked to memory consolidation and spatial alternation is a memory task (Karlsson & Frank, 2009; Kay et al., 2016; Muessig et al., 2019; O'Neill et al., 2008). We expected some combination of a difference in oscillations in LFPs and coherence between LFPs during T-maze trials, along with a change in ripple activity during sleep, would contribute to the spatial alternation performance decrease.

In terms of oscillatory activity, we primarily focus on the analysis of theta rhythm activity (6 – 12Hz; related to spatial coding in single units and motor activity) and beta rhythm activity (15 – 30Hz; linked to learning and novelty detection) from LFPs (Berke et al., 2008; Buzsáki, 2002; Cacucci et al., 2004; Colgin, 2013; França et al., 2014; Geisler et al., 2007; Grossberg, 2009; Iwasaki et al., 2021; Kay et al., 2016; Richard et al., 2013; Solomon et al., 2019; Womelsdorf et al., 2010). However, delta (1 – 4Hz; linked to concentration) and gamma (40 – 120Hz; related to decision making) rhythms will also be examined (Amzica & Steriade, 1998; Buzsáki & Wang, 2012; Colgin et al., 2009; Harmony, 2013; Long et al., 2021; Montgomery & Buzsáki, 2007; Schultheiss et al., 2020; Trimper et al., 2017). The subiculum is capable of independently generating theta and gamma rhythms, as well as generating theta-gamma coupling based on *in vitro* studies of the subicular structure (Jackson et al., 2011).

Overall, the ATN lesions had a surprisingly subtle impact on oscillations in LFPs and coherence between LFPs given the absence of spatially tuned cells in the subiculum and the behavioural impairments post-lesion. Instead, differences emerge when comparing LFPs across brain regions and examining the relationship between LFPs and spiking activity.

5.3 Results of analyses on local field potentials

5.3.1 Rhythms in LFPs are maintained in the subiculum during exploration, but not in retrosplenial cortex

The average power spectral densities of LFPs in both the subiculum and retrosplenial cortex followed the expected $1/f$ decay of power with frequency (Figure 5.1A, B). Clear peaks (accounting for the $1/f$ decay of power) can be seen in the power spectra in the theta band primarily, followed by the beta band, with a small number of peaks in the delta and gamma bands (Figure 5.1C, D). This general power spectra structure and peaks in oscillatory activity were visible across both the control rats and ATNx (ATN-lesioned) rats, demonstrating that fundamental oscillations in LFPs were not abolished in either of these two areas.

The bandpower in each of the oscillatory bands was calculated from the power spectra (Figure 5.2A, B). Theta was the most dominant bandpower in both the subiculum and RSC, even when accounting for the $1/f$ decay of power (Figure 5.2C, D). Beta band activity was visually higher in control rats compared to ATNx rats, though the effect size was quite small (common language effect size (CLES) of 0.38) and was not significant when correcting for multiple comparisons, while activity in other bands was consistent between the two groups (Table 5.1). In the RSC, delta and theta band activity also remained consistent across the two groups, though beta and gamma oscillatory power were higher in control rats compared to ATNx rats (Table 5.1).

As Table 5.1 reports relative power and has a lower sample size for the ipsilateral RSC LFPs ($n_{\text{control}} = 38$, $n_{\text{lesion}} = 19$) than the subicular data, further tests were considered for the absolute bandpower in the subiculum including all RSC recordings (including ipsilateral and contralateral to the subicular recording site, and the midline; as there is a bilateral lesion) with $n_{\text{control}} = 97$, $n_{\text{lesion}} = 57$ in both cases. These statistical tests were consistent with the findings from the relative bandpower tests in both the subiculum recordings and the ipsilateral RSC recordings.

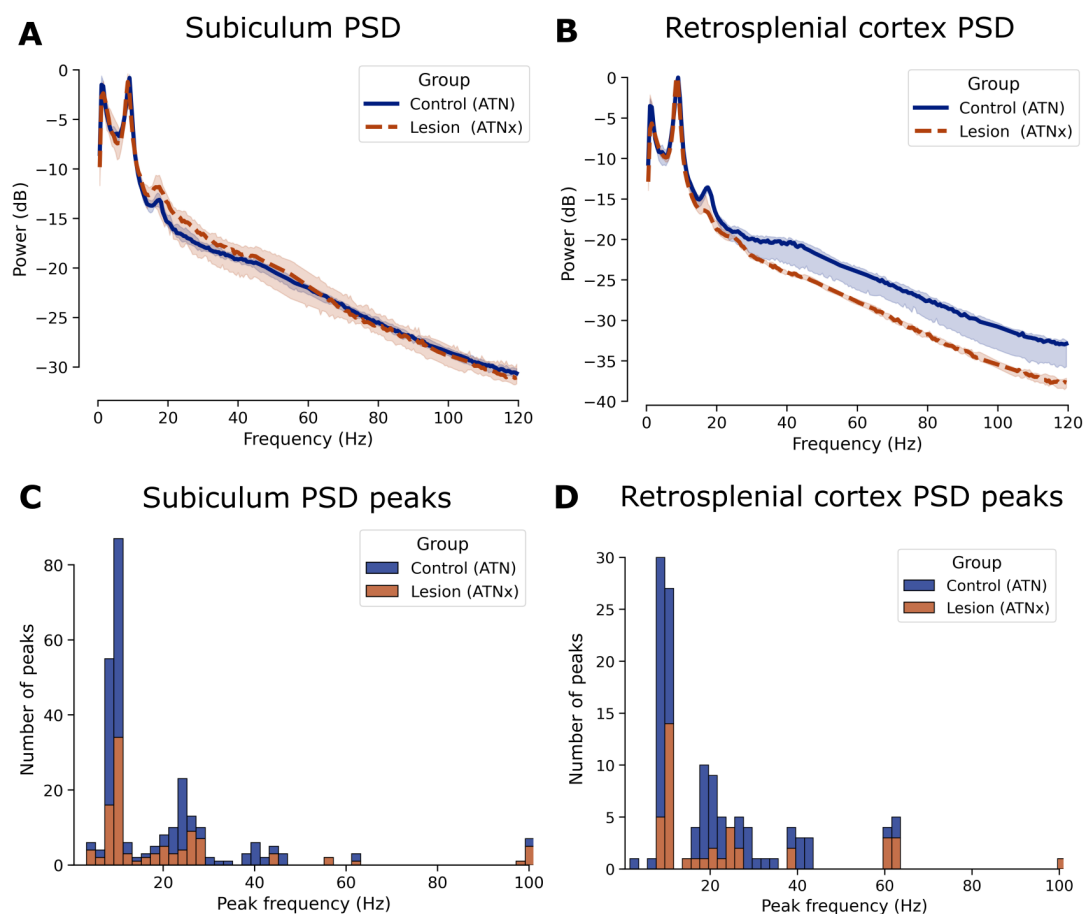


FIGURE 5.1: Power spectral densities of subicular and retrosplenial cortex LFPs compared between control and lesion rats during open-field exploration.

(A) The median power spectral densities of LFPs from electrodes implanted in the subiculum across control and lesioned rats recorded during open-field recordings. The line represents the median power across recordings per frequency, and the error spread represents a 95% confidence interval.

(B) As above, but the median power spectral densities from the retrosplenial cortex.

(C) Peaks in the power spectral densities of LFPs from electrodes implanted in the subiculum across control and lesioned rats recorded during open-field recordings.

(D) As above, but the peaks in the PSD from the retrosplenial cortex.

			Median	IQR	U	CLES	p	BF_{10}
Delta	SUB	Control	0.20	0.06	3023	0.56	0.25	0.34
		Lesion	0.19	0.13				
	RSC	Control	0.14	0.04	246	0.34	0.053	1.00
		Lesion	0.16	0.03				
Theta	SUB	Control	0.39	0.17	2824	0.52	0.69	0.18
		Lesion	0.34	0.25				
	RSC	Control	0.49	0.10	246	0.34	0.053	0.7
		Lesion	0.56	0.11				
Beta	SUB	Control	0.08	0.04	2045	0.38	0.011	1.9
		Lesion	0.11	0.06				
	RSC	Control	0.08	0.009	625	0.87	< 0.001*	18.6
		Lesion	0.06	0.005				
Low Gamma	SUB	Control	0.044	0.014	2520	0.46	0.46	0.21
		Lesion	0.049	0.034				
	RSC	Control	0.038	0.021	676	0.93	< 0.001*	182.3
		Lesion	0.021	0.005				
High Gamma	SUB	Control	0.012	0.006	2270	0.42	0.09	0.50
		Lesion	0.014	0.001				
	RSC	Control	0.009	0.006	656	0.91	< 0.001*	104.5
		Lesion	0.004	0.001				

TABLE 5.1: The relative bandpowers of local field potentials recorded from control and lesion rats during open-field exploration. For LFPs in SUB, $n_1 = 97$, $n_2 = 56$ for the Mann-Whitney U test, representing recordings from 8 control rats and recordings from 6 lesioned rats. For LFPs in RSC, only ipsilateral recordings to the SUB recordings were considered, so $n_1 = 38$, $n_2 = 19$ representing recordings from 4 control and 3 lesioned rats. Multiple tests were corrected for via Bonferroni correction.

Key: SUB – subiculum, RSC – retrosplenial cortex, IQR – inter-quartile range, CLES – common language effect size, BF_{10} – Bayes factor in support of the alternative hypothesis, * $p < 0.01$.

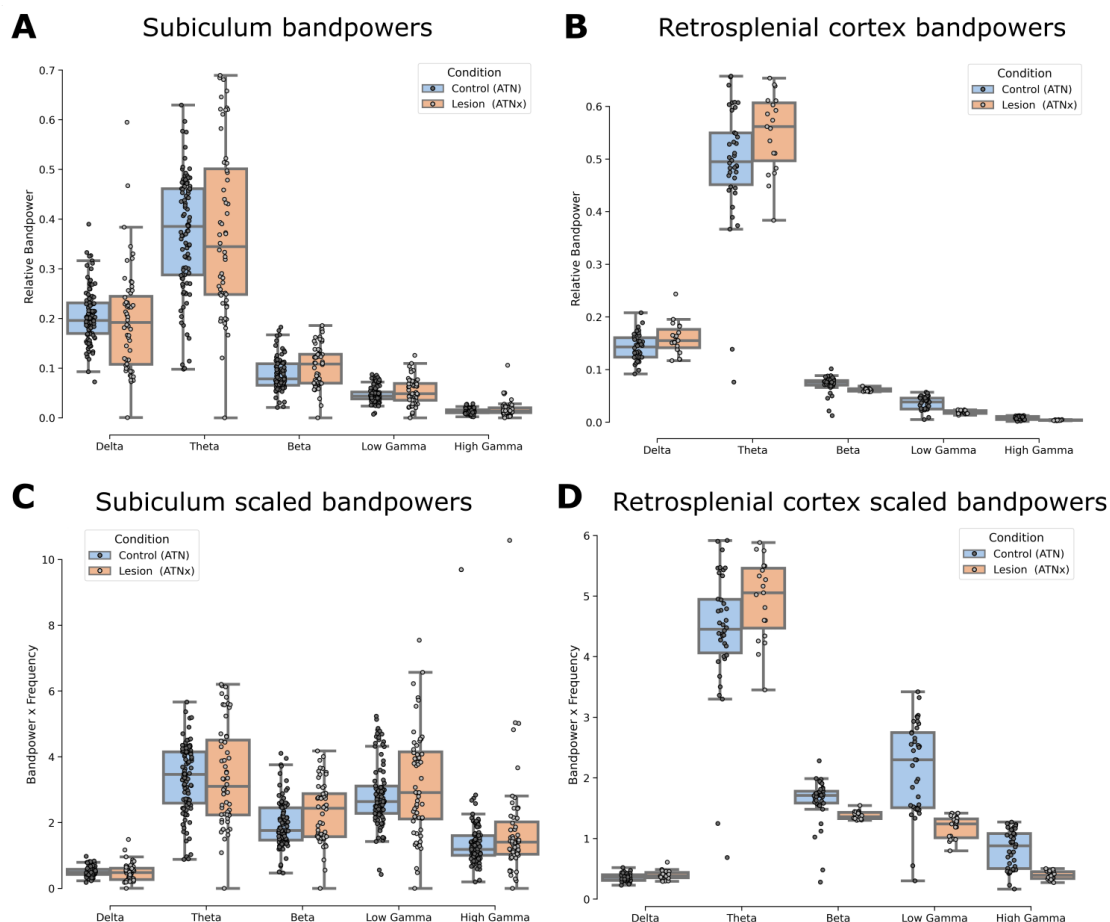


FIGURE 5.2: Power in oscillation bands of the subicular and retrosplenial cortex LFPs compared between control and lesion rats during open-field exploration.

(A) Relative bandpowers in LFPs from the subiculum of control and lesioned rats during open-field exploration. Relative bandpower represents the ratio between the power in that frequency band to the power in the total frequency band (0.5 to 120Hz).

(B) As above, but in LFPs recorded from the ipsilateral retrosplenial cortex.

(C) Scaled bandpowers from LFPs recorded from the subiculum during open-field exploration as relative bandpower \times mean band frequency (e.g., mean of 9Hz for Theta) to account for $\frac{1}{f}$ decay of power with frequency.

(D) As above, but from LFPs recorded from the ipsilateral retrosplenial cortex.

Key: Boxplots represent the median at the line in the box, and the first and third quartiles on the lower and upper edges of the box, respectively. Whiskers extend up/down to the furthest data point less than within $1.5 \times$ IQR from the third/first quartile. Individual data points are overlaid over the boxplots (points outside the whiskers can be considered outliers).

5.3.2 Heightened motor activity remains linked to increased theta activity

Theta waves are regularly present in the hippocampus during various types of locomotor activities, such as voluntary exploration (Buzsaki, 2002). Despite CA1 having the most regular and high amplitude theta oscillations, they are also seen in the further hippocampal formation (including subiculum) and other nearby brain structures such as the perirhinal area and RSC. As theta oscillations are generally absent or weaker during non-moving awake states, we investigated the link between rat movement speed and theta power in control and lesioned rats. Consistent with theta being present and dominant during open-field exploration (Figure 5.1), the expected correlation between increased motor activity and increased theta-band activity was maintained in lesion rats, though at a slightly decreased rate (Figure 5.3). There was significant positive Pearson correlation between rat running speed (cm/s) and relative theta power in LFPs from the subiculum and RSC in both control and lesion rats. The correlation was strongest in the RSC of control animals, with a correlation coefficient of $r = 0.64$. In the below tests N is the number of data points used in the correlation, while n is the number of trials analysed.

- Subiculum control, $r = 0.37$, $N = 3685$, $n = 97$, $p < 0.001$, $\text{BF}_{10} > 100$.
- Subiculum lesion $r = 0.22$, $N = 2162$, $n = 56$, $p < 0.001$, $\text{BF}_{10} > 100$.
- RSC control $r = 0.73$, $N = 1407$, $n = 38$, $p < 0.001$, $\text{BF}_{10} > 100$.
- RSC lesion $r = 0.58$, $N = 729$, $n = 19$, $p < 0.001$, $\text{BF}_{10} > 100$.

The relationship between speed and theta band LFP power exhibits a non-linear trend with a discernible linear component (Figure 5.1). Generally, a second or third degree polynomial aptly captures this relationship. The model, along with the correlation, reveals two key findings. Firstly, the model's intercept is higher in ATN lesion-induced dysfunction (ATNx) rats than in control rats in the retrosplenial cortex, but this is reversed in the subiculum (higher in control rats than lesioned rats). This discrepancy aligns with the data of non-speed related power (Figure 5.2A, B). Secondly, the linear growth of theta power with speed shows a slight reduction in ATNx rats compared to control rats. Consequently, the speed-power (Figure 5.1) relationship shows a non-speed related component and a speed-related component (note that we also repeated the analysis with absolute bandpower to confirm that the difference in relative theta power is not attributable to power changes in other bands).

In examining the non-speed-related theta power component (Table 5.1) no significant differences were observed. However, subicular theta power was lower in ATNx rats

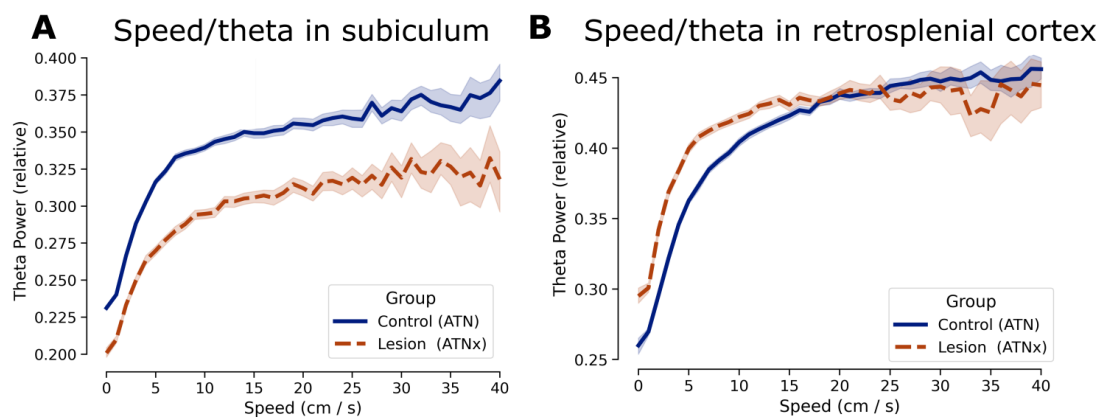


FIGURE 5.3: Comparing running speed and relative theta power during open-field exploration from LFPs recorded in the subiculum (A) and ipsilateral retrosplenial cortex recordings (B). Speed and theta power in both subiculum and retrosplenial cortex are positively correlated, as would be expected in healthy rats. This correlation still holds in ATN lesioned rats.

compared to control rats, while retrosplenial cortex theta power was higher in ATNx rats than in control rats. Given the direct reciprocal connections between the ATN and both regions, this observation may be attributed to several factors. Among potential explanations, a compensatory mechanism is plausible. Moreover, the increased power in the retrosplenial cortex post-lesion, coupled with a concurrent power decrease in the subiculum post-lesion, suggests a form of circuit balancing. This non-speed related theta power change in lesioned rats may compound during movement as there is a difference in theta power growth rate with speed in lesioned rats.

5.3.3 Sharp wave ripples occur during resting periods in both control and lesioned rats

During rest periods (see subsection 5.5.6 for resting detection details), sharp-wave ripples were detected in control rats, lesioned rats, and rats with cannulae for muscimol infusion. Overall, there was no significant difference in the number of sharp-wave ripples per minute between experimental manipulation groups (control vs NMDA lesion and cannulated control vs muscimol infusion) in both subiculum and CA1 (Figure 5.4A). The visual difference between control and cannulated control (CanControl) animals may be due to a difference in the experiments performed before resting, potentially leading to increased memory consolidation. For instance, in cannulated rats, resting recordings were often performed after many T-maze trials, while in control rats, resting recordings were often performed on isolated days or after free exploration. Subicular ripples in control rats were detected at a rate of (mean \pm SE) 13.89 ± 2.12 ripples per minute, while ripples in lesioned rats were detected at a rate of 18.64 ± 0.42 ripples per

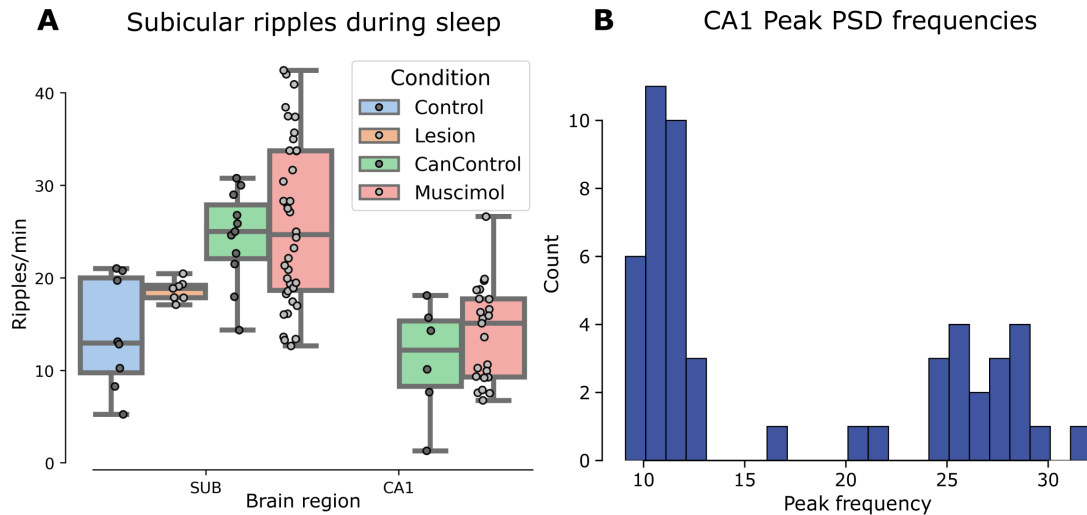


FIGURE 5.4: Sharp-wave ripples in the subiculum and CA1 during resting periods in control and ATN lesioned rats.

(A) The average number of sharp-wave ripples detected in the subiculum and CA1 per minute during resting recordings from control, ATN lesioned, cannulated rats (CanControl), and cannulated rats with muscimol infusion. CA1 data is available in cannulated rats due to a different recording layout with the canulae compared to regular control rats.

(B) The peak frequencies in the power spectra of LFPs recorded from the CA1 of 3 permanent NMDA ATN lesioned rats while rats performed open-field exploration. This further suggests a lack of impact of ATN lesions on the CA1 as regular theta and beta peaks are present.

minute ($t = -2.18$, $df = 7.5$, $p = 0.06$, $CLES = 0.48$, $n_1 = 8$, $n_2 = 7$, $BF_{10} = 1.55$, Welch's t-test), which was not significant. In cannulated control rats, subicular ripples were detected at a rate of 24 ± 1.52 ripples per minute, while ripples were detected at a rate of 25.86 ± 1.45 per minute after muscimol infusion ($t = -0.69$, $df = 29.73$, $p = 0.49$, $CLES = 0.39$, $n_1 = 11$, $n_2 = 38$, $BF_{10} = 0.36$), which was not significant. Ripples in CA1 were detected in cannulated control rats at a rate of 11.18 ± 2.51 ripples per minute, while in muscimol infused rats ripples were detected at a rate of 13.93 ± 1.08 per minute ($t = -1.0$, $df = 7.0$, $p = 0.34$, $CLES = 0.38$, $n_1 = 6$, $n_2 = 23$, $BF_{10} = 0.61$), which was not significant. Additionally, in rats with ATN lesions, the peaks in the power spectral density from CA1 LFPs were as expected (primarily in the theta and beta bands) during exploration (Figure 5.4B).

5.3.4 Spike-LFP theta coupling is lower for non-spatial units, particularly from lesioned rats

In control rats exploring the open-field, spike field coherence (SFC), representing the amount of phase synchronisation between the LFP and spikes, was higher in the theta

band for spatial units than for non-spatially tuned units (Figure 5.5A; median theta SFC for spatial units 0.056, inter-quartile range (IQR) 0.096, non-spatial units median 0.028, IQR 0.046, Mann-Whitney $U = 1095$, CLES = 0.67, $n_1 = 47$, $n_2 = 35$, $p = 0.011$, $BF_{10} = 6.11$, where n_1, n_2 indicate the number of spatial/non-spatial units). The SFC in the delta band was higher for spatial units than non-spatial units, though this was not significant when correcting for multiple comparisons (Figure 5.5A; median delta SFC for spatial units 0.022, IQR 0.002, non-spatial units median 0.020, IQR 0.003, Mann-Whitney $U = 1057$, CLES = 0.64, $n_1 = 47$, $n_2 = 35$, $p = 0.028$, $BF_{10} = 1.47$). Only non-spatial units could be compared between control and lesioned rats as no spatially tuned units were recorded from lesioned rats. Non-spatial units in control rats had a higher SFC in the theta band than non-spatial units from lesioned rats (median theta SFC in controls 0.028, IQR 0.046, lesion median 0.021, IQR 0.008, Mann-Whitney $U = 498$, CLES = 0.71, $n_1 = 35$, $n_2 = 20$, $p = 0.01$, $BF_{10} = 4.51$). Additionally, the SFC was higher in the delta band for non-spatial units from lesioned rats than non-spatial units from control rats, but was not significantly different when correcting for multiple comparisons (median delta SFC in controls, 0.020, IQR 0.001, lesion median 0.023, IQR 0.005, Mann-Whitney $U = 218$, CLES = 0.31, $n_1 = 35$, $n_2 = 20$, $p = 0.021$, $BF_{10} = 2.90$).

The theta SFC from spatially responsive units was compared before muscimol infusion and at 60-80 minutes post-infusion. While there was a visual decrease post-infusion (Figure 5.5C), this was not significant (pre-infusion median theta SFC 0.024, IQR 0.009, post-infusion median 0.021, IQR 0.007, Wilcoxon paired signed-rank $W = 16$, CLES = 0.63, $n_1 = n_2 = 10$, $p = 0.275$, $BF_{10} = 0.67$). There was an increase in the average delta SFC that was not significant when correcting for multiple comparisons post-infusion, though Bayesian tests indicate support of the difference (pre-infusion median delta SFC 0.020, IQR 0.002, post-infusion median 0.027, IQR 0.020, Wilcoxon paired signed-rank $W = 6$, CLES = 0.1, $n_1 = n_2 = 10$, $p = 0.027$, $BF_{10} = 4.6$). There was an associated visual decrease in the theta rhythmic activity in the spike-triggered average (STA), representing the average LFP around spike occurrences, in lesioned rats and post muscimol infusion, with increased delta activity visible (Figure 5.5B, D).

5.3.5 Coherence between LFPs from the subiculum and retrosplenial cortex is higher in control rats in the theta and beta bands

In open-field recordings, there was a high theta peak in coherence in both control rats and ATN-lesioned rats (Figure 5.6A), though the peak was more pronounced in control rats (control median theta coherence value 0.47, IQR 0.21, lesion 0.29, IQR 0.14; Mann-Whitney $U = 687$, CLES = 0.88, $n_1 = 38$, $n_2 = 19$, $p < 0.001$, $BF_{10} = 80.1$).

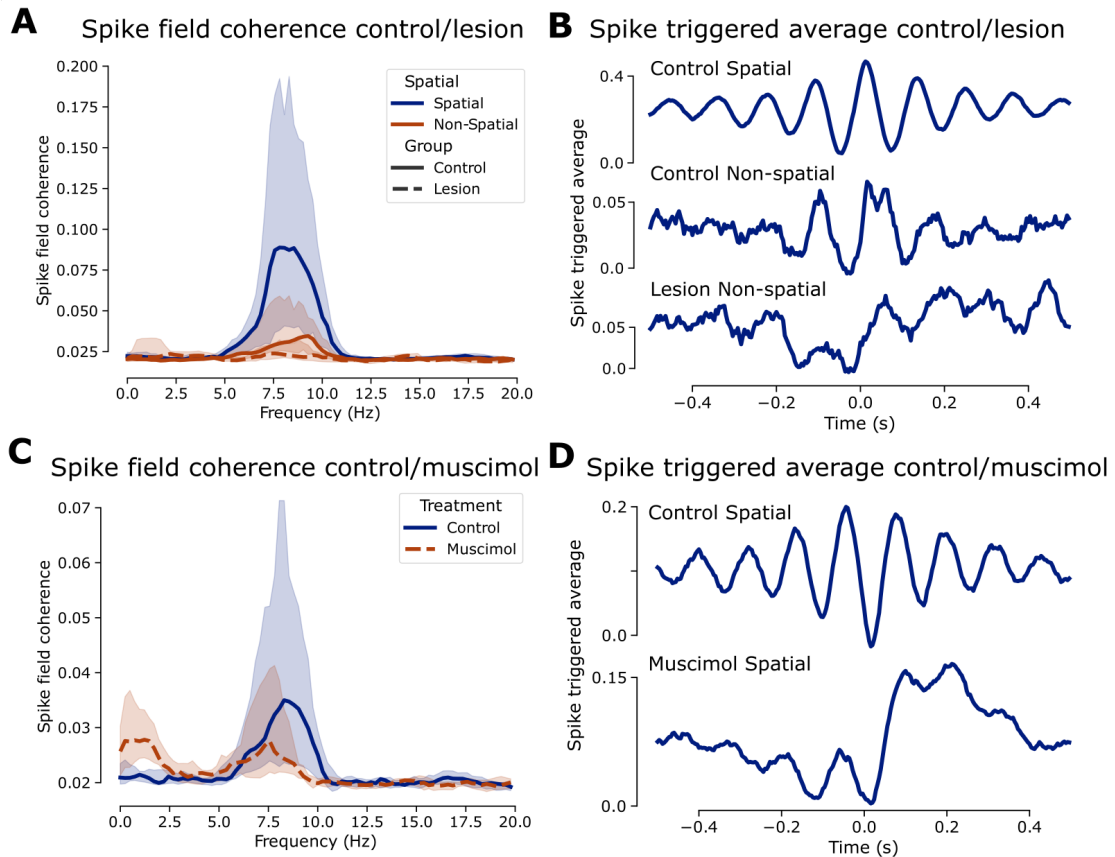


FIGURE 5.5: The subicular LFP to spike relationship measured by spike field coherence and spike-triggered average.

(A) The coherence between subicular units and the LFP in control and lesion rats, split between spatial and non-spatial units (there were no recorded spatially tuned units from lesioned animals).

(B) The average subicular LFP around a spike for units recorded from the subiculum.

(C) The coherence between subicular spatially-tuned units and the subicular LFP recorded before and 60 – 80 minutes after the infusion of muscimol.

(D) The average subicular LFP around a spike for units recorded from the subiculum before and after the infusion of muscimol.

Additionally, there was a peak in beta coherence in control rats that was not present in the ATN-lesioned rats (control median beta coherence 0.13, IQR 0.27, lesion 0.04, IQR 0.03; Mann-Whitney $U = 660$, CLES = 0.91, $n_1 = 38$, $n_2 = 19$, $p < 0.001$, $BF_{10} > 100$). Other coherence band values were not significantly different between control and lesioned rats when correcting for multiple comparisons via Bonferroni correction (control median delta coherence 0.17, IQR 0.3, lesion 0.13, IQR 0.09; Mann-Whitney $U = 487$, CLES = 0.67, $n_1 = 38$, $n_2 = 19$, $p = 0.034$, $BF_{10} = 1.01$. Control median low gamma coherence 0.14, IQR 0.29, lesion 0.08, IQR 0.07; Mann-Whitney $U = 420$, CLES = 0.58, $n_1 = 38$, $n_2 = 19$, $p = 0.322$, $BF_{10} = 0.35$. Control median high gamma coherence 0.25, IQR 0.25, lesion 0.26, IQR 0.07; Mann-Whitney $U = 303$, CLES = 0.52, $n_1 = 38$, $n_2 = 19$, $p = 0.5$, $BF_{10} = 0.32$).

Given these theta and beta peaks in coherence in open-field recordings, the same bands were investigated in a shorter timespan during T-maze spatial alternation (Figure 5.6). Theta coherence was significantly higher during T-maze choice alternation in control rats than in lesioned rats in both correct trials (control median theta coherence 0.5, IQR 0.54, lesion 0.24, IQR 0.18; Mann-Whitney $U = 6963$, CLES = 0.67, $n_1 = 173$, $n_2 = 60$, $p < 0.001$, $BF_{10} = 6.68$) and incorrect trials (control median theta coherence 0.49, IQR 0.58, lesion 0.22, IQR 0.21; Mann-Whitney $U = 1498$, CLES = 0.67, $n_1 = 35$, $n_2 = 64$, $p = 0.006$, $BF_{10} = 2.92$). Beta coherence was also higher during T-maze choice alternation in control rats than in lesioned rats in both correct trials (control median beta coherence 0.34, IQR 0.37, lesion 0.16, IQR 0.29; Mann-Whitney $U = 6550$, CLES = 0.63, $n_1 = 173$, $n_2 = 60$, $p = 0.003$, $BF_{10} = 1.73$) and incorrect trials (control median beta coherence 0.29, IQR 0.40, lesion 0.17, IQR 0.24; Mann-Whitney $U = 1405$, CLES = 0.63, $n_1 = 35$, $n_2 = 64$, $p = 0.037$, $BF_{10} = 1.83$).

5.4 Contrasting the LFP results with the abrogation of subicular cells in ATN-lesioned rats

The anterior thalamic nuclei (ATN) play a critical role in memory, spatial learning, and navigation. ATN lesions in rodents severely impair a wide range of spatial tasks (Aggleton & Nelson, 2015). In our previous chapter, we demonstrated that ATN lesions cause deficits in subicular spatial signalling while keeping CA1 intact (Frost et al., 2021). Additionally, we found that ATN lesions impact spatial alternation memory, which was also corroborated by a complementary study involving DREADDs (Nelson et al., 2020). In this study, we delve into the effects of ATN lesions on subicular local field potentials (LFPs) and investigate the relationship between the subiculum and the retrosplenial cortex — an area crucial for spatial memory and navigation (Aggleton & Nelson, 2015;

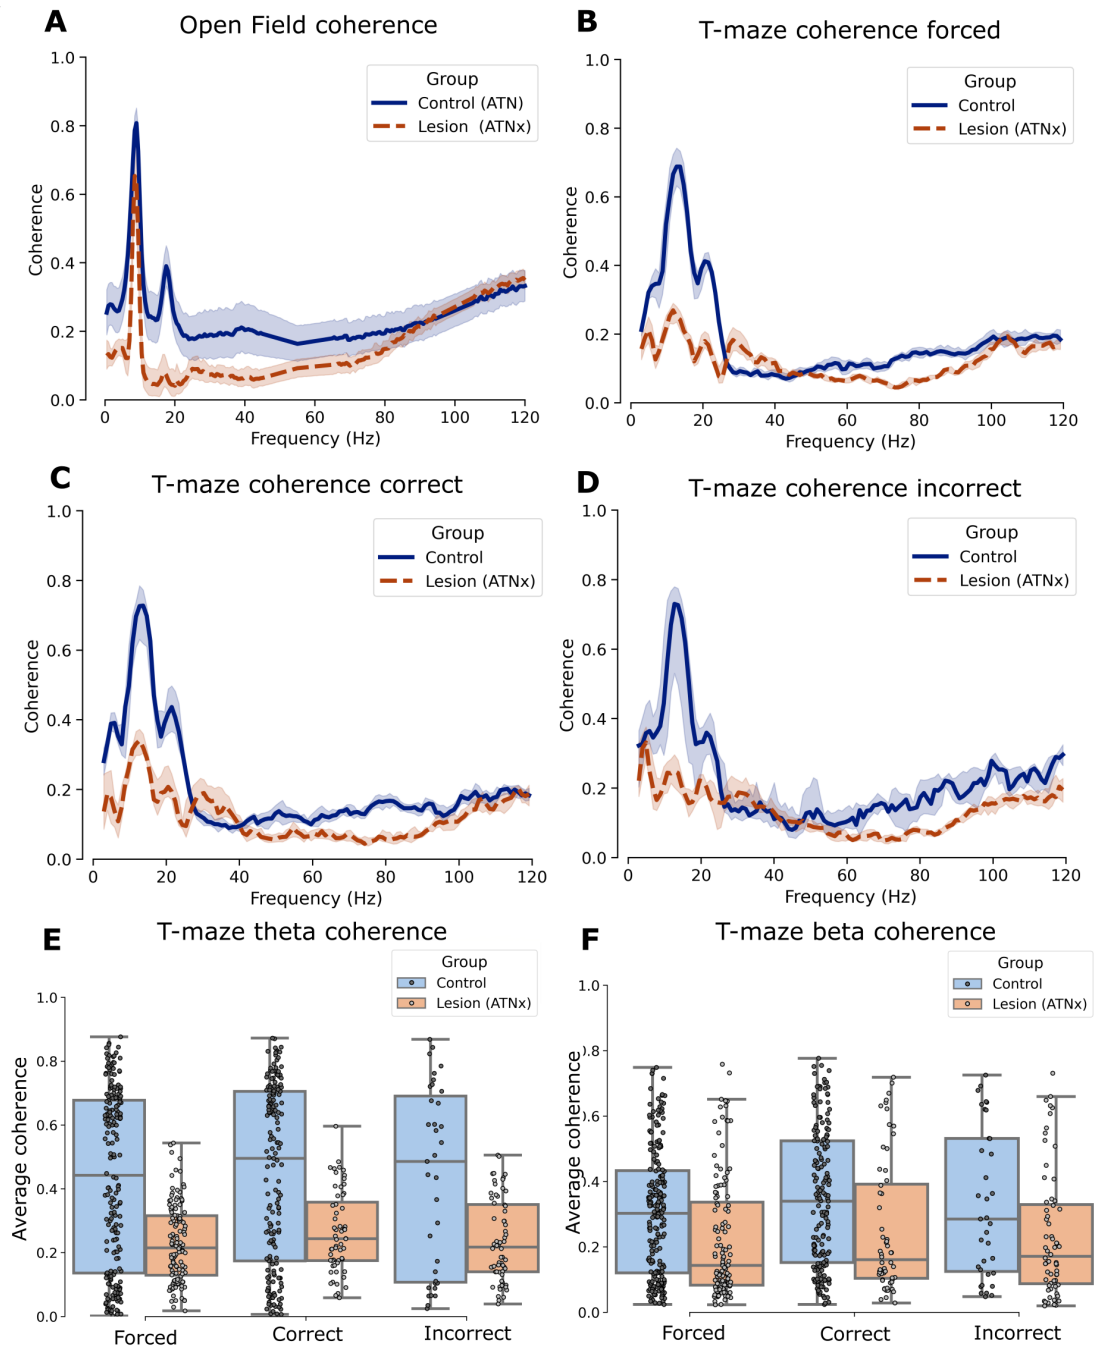


FIGURE 5.6: Coherence between subicular and retrosplenial cortex LFPs in control and ATN lesioned rats during T-maze spatial alternation and open-field exploration. (A) The coherence between LFPs recorded from the subiculum and ipsilateral retrosplenial cortex in open-field recordings from control and lesioned rats. (B, C, D) The coherence between LFPs recorded from the subiculum and ipsilateral retrosplenial cortex in T-maze trial recordings, split based on the trial result. (E, F) The average coherence in the theta (E) and beta (F) bands between LFPs recorded from the subiculum and ipsilateral retrosplenial cortex in T-maze trial recordings.

Key: Forced indicates the starting trial, where the choice is forced, and correct or incorrect indicate the ending trial, based on whether the rat chose the alternative direction in the maze. LFPs are considered between the time when the door is lifted for the rat to enter the maze, and the time the rat completes the maze.

Keene & Bucci, 2008; Todd & Bucci, 2015).

Surprisingly, ATN lesions had only subtle effects on subicular LFPs, despite the significant behavioural impairments observed in spatial alternation memory and the loss of spatial cell firing in the subiculum post-lesion. Rhythms in subicular LFPs remained largely intact during open-field exploration, and the link between increased motor activity and heightened theta band oscillation activity was maintained post-lesion. Moreover, the memory consolidation process during sleep, as indicated by sharp-wave ripples, remained unaffected in lesioned rats. This suggests that the observed deficits in choice alternation performance are likely attributed to factors other than sleep-dependent memory consolidation.

Although many aspects remained stable post-lesion, the relationship between signals within and across brain regions was indeed affected. The coherence between the RSC and subicular LFPs was reduced in the T-maze task after ATN lesion, particularly around the decision point in the maze, aligning with the observed spatial alternation memory deficits. Furthermore, the RSC activity was reduced outside the theta band post-lesion of the ATN, suggesting that changes in the RSC may impact the subiculum. While volume conduction in the rodent brain (Pesaran et al., 2018) can sometimes cause LFPs to reflect large-scale brain dynamics rather than local changes, this possibility is less likely here as subicular LFPs remained similar post-lesion while RSC LFPs exhibited differences, indicating that these LFPs reflect local changes in these regions. Additionally, the time-coupling between subicular the LFP theta phase and spiking activity was reduced in the lesioned rats compared to controls, with irregular spikes occurring in conjunction with theta peaks post-lesion. This alteration in spike-LFP coupling, particularly in the theta range, was further exacerbated during muscimol inactivation, as previously spatially responsive units reduced their theta range coupling after ATN inhibition.

Now, the question arises: What could be the cause of the changes in coherence across regions and spike-LFP coupling, while the rhythms in subicular LFPs themselves remain relatively unaffected? The power in LFPs reflects the coherent spatially oriented activity of pyramidal neurons but is influenced by various factors (Vinck et al., 2023). Since our recording sessions included only a small number of units and the number of units recorded per animal was limited, we primarily analysed LFPs to capture activity from multiple sources in the brain. The spiking activity of cortical neuron populations can often be described by a small set of “latent dynamics” that represent population-wide covariance features, which are strongly associated with LFPs (Gallego-Carracedo et al., 2022). LFPs are influenced by activity in neurons, glia, synapses, and any excitable membrane, but primarily reflect correlated synaptic patterns across the local circuit

near the recording device (Buzsáki et al., 2012). With this in mind, it is plausible that subicular LFPs remain unaffected by ATN lesions because they are driven by a different set of cell populations than those impacted by the lesions in the ATN.

Furthermore, it is unlikely that ATN lesions had a knock-on effect of inhibiting the medial septum, which serves as the primary generator of hippocampal theta rhythm (Nuñez & Buño, 2021). The behavioural and cellular impacts observed in ATN-lesioned rats may, therefore, result from intricate interactions. Perhaps, damage to a specific type of ATN neuron or a set of neural circuits within the ATN is responsible for the change in subicular cell tuning, rather than requiring complete inhibition of the ATN, which could affect other brain areas as well. This opens up the possibility of replicating similar spatial coding deficits using more targeted inactivation methods like optogenetic inhibition or designer receptors exclusively activated by designer drugs (DREADDs).

Previous research has shown that bursting cells in the dorsal subiculum with direct connections to the granular RSC directly impact sharp-wave ripples in the RSC (Nitzan et al., 2020). Thus, the changes in coherence between RSC and subicular LFPs, along with the alterations in RSC LFPs, could partially reflect the modifications in bursting cell firing observed after ATN lesions. Post-lesion, there was a reduction in spike-LFP coupling in the theta range, leading to irregular spikes occurring in conjunction with theta peaks. When spikes phase-lock with LFPs, it indicates increased coordination and coherence, with many spikes occurring at a regular rhythm due to the high influence of correlated synaptic inputs to pyramidal neurons (Schneider et al., 2021). This suggests that the ATN plays a crucial role in the temporal coordination of subicular cell firing, which is disrupted after the lesions and likely contributes to the spatial alternation deficits observed in ATN-lesioned rats.

The coherence changes observed in the subicular and retrosplenial cortex LFPs primarily occurred in the slow-wave theta and beta bands, particularly during the T-maze spatial alternation task. Increased theta band activity is highly associated with memory-related processing and decision-making (Colgin, 2013; Richard et al., 2013; Solomon et al., 2019; Womelsdorf et al., 2010), while beta waves in the hippocampus are related to object-location memory performance, learning, and novelty detection (Berke et al., 2008; França et al., 2014; Grossberg, 2009; Iwasaki et al., 2021). Coherence between two brain regions likely reflects the information transfer between them at a particular frequency. Therefore, theta and beta coherence may reflect longer-term decision-making processes due to their slower speed and the transfer of both long-range and intermediate feedback (Vinck et al., 2023). The decrease in theta coherence could indicate a change in information transfer between neurons involved in spatial alternation memory. Additionally, theta oscillations are linked to gamma amplitude, which may

explain the reduction in gamma power observed in the RSC after ATN lesions.

In summary, this study emphasises the essential role of the ATN in a larger memory system involving the subiculum, retrosplenial cortex, and hippocampus. Permanent or transient ATN lesions result in behavioural impacts on spatial alternation memory and an increase in motor activity, while abolishing spatial coding in subicular neurons. However, it is unlikely that these changes are solely due to alterations in the direct interaction between the ATN and the subiculum, as subicular LFPs remain largely unchanged. The observed changes in coherence and spike-LFP coupling indicate broader circuit-level alterations underlying the behavioural and cellular differences observed.

5.5 Analysis methods for local field potentials

5.5.1 Selecting local field potentials informed by histology

All electrodes targeting the CA1 and subiculum were on target. However, the electrodes targeting the retrosplenial cortex were not always in the ipsilateral retrosplenial cortex to the dorsal subicular recordings. In these cases, the wires remained at the midline, or crossed the midline and were implanted in the contralateral retrosplenial cortex. For analyses involving the retrosplenial cortex, only data from the ipsilateral retrosplenial cortex to the dorsal subicular recordings were included. As rats were bilaterally lesioned, data from the retrosplenial cortex of the contralateral hemisphere could not be used as control data. Figure 5.7 demonstrates example tracking of the electrode trajectories in the retrosplenial cortex. In control rats, 3 midline RSC recordings, 1 contralateral RSC recording, and 4 ipsilateral RSC recordings were performed. In lesion rats, 5 contralateral RSC recordings and 4 ipsilateral RSC recordings were performed, though the data from one ipsilateral recording was corrupted, so three RSC recordings were used. In cannulated rats, 3 contralateral RSC recordings and 4 ipsilateral RSC recordings were performed, and a further 3 recordings were not confirmed as to their RSC location, so were excluded from RSC analyses.

5.5.2 Signal filtering and processing of local field potentials

All LFPs were notch-filtered at 50Hz due to mains power frequency, and then filtered with a one-pass, zero-phase, non-causal band-pass from 0.5 to 120Hz finite impulse response (FIR) filter using MNE (Gramfort et al., 2013). The filter length was 6.6 seconds, the cutoff frequencies were 0.25 and 122.50Hz and a hamming window with 0.0194 passband ripple and 53dB stopband attenuation was used.




FIGURE 5.7: Examples of retrosplenial cortex bipolar electrode implantation histology.

(A) Implantation of a bipolar electrode in a control rat that was on target in the ipsilateral retrosplenial cortex to the dorsal subiculum that tetrodes were implanted in.

(B) Implantation of a bipolar electrode in a lesioned rat that was on target in the ipsilateral retrosplenial cortex to the dorsal subiculum that tetrodes were implanted in.

(C) Implantation of a bipolar electrode in a control rat that was off target in the midline.

(D) Implantation of a bipolar electrode in a control rat that was off target in the contralateral retrosplenial cortex to the dorsal subiculum that tetrodes were implanted in.

For control rats and lesion rats, the two LFPs on the bipolar electrodes in each brain region were considered as these were the only LFPs available in many of these rats and thus the comparison was kept the same even for those rats with 32 LFPs. However, for cannulated rats, 32 LFPs were available from all rats, so all LFPs were considered that were not noisy. To determine noisy LFPs, first, all signals were z-score normalised for comparison as the bipolar electrodes were a different thickness to the tetrode wires (the resistance of a wire is inversely proportional to the diameter of the wire, and thus the absolute power recorded on that wire would differ). This is a linear transformation, and as such, does not impact any of the results, including normalised power spectral density (PSD) in dB, coherence, FOOOF, or spike-LFP (bar the STA, which is scaled, and can be returned to the original by inverting the linear transformation). Then a normalised signal was excluded if it differed from the average normalised signal by more than 10% of the median difference between all signals and the average normalised signal. Then, the remaining non-noisy signals were subsequently averaged to produce one synthetic LFP signal from each brain area in each recording session for further analysis (for bipolar electrode LFP recordings, the recording was discarded if the LFPs on each wire were not in agreement). For ripple detection, however, this was performed on all recorded non-noisy LFPs instead of using a synthetic LFP signal, and then the results were merged.

However, other methods were applied to ensure the noise detection and averaging LFPs into a single synthetic signal did not influence the results, including; two-random LFPs were selected from the non-outlier signals for each brain area; a single random LFP was selected from all LFPs for each brain area; and only LFPs recorded from bipolar electrodes were used. While there were per-recording differences in these methods, it was found that the method of selection of a subset of LFPs did not significantly affect the outcome when summarising across all recordings.

5.5.3 Power spectral density analysis

Power spectral density was calculated from the LFP using Welch's periodogram (Welch, 1967); the modified periodogram was calculated from the LFP signal in 2-second segments with 1-second overlaps, where a Hann window was used. The resulting periodograms from each segment were combined by taking the mean. Since z-score normalised signals were often used, absolute power was disregarded. As such, the power values were converted to decibels, using the maximum power value in the periodogram as the reference point, $p_{\text{dB}} = \log_{10}(p_{\text{V}}/\max(p_{\text{V}}))$, where p_{dB} is the power in decibels, and p_{V} is the power volts. Oscillations in the LFP power were investigated in the following ranges:

- Delta, 1 to 4Hz.
- Theta, 6 to 12Hz.
- Beta, 15 to 30Hz.
- Low-gamma, 40 to 80Hz.
- High-gamma, 80 to 120Hz.
- Ripple waves, 150 to 210Hz.

Bandpower was calculated using Simpson's integration on the discrete samples lying within the band from Welch's periodogram calculation. Bandpowers are discussed relatively, whereby the relative power in each band is considered to be the unitless quantity $bp_r = bp/tp$ where bp is the bandpower, and tp is the total power of the LFP (Cardin et al., 2009). As such, the relative bandpower has a value between 0 and 1, where a value of 1 would indicate all the power in the LFP occurs within the given band. Relative power does not reflect overall increases in the LFP power across all bands (which can occur due to a variety of reasons), but instead reflects an increase/decrease of the power in the band of interest relative to the full bandpower. For plotting the median PSD over all animals during free exploration, the periodogram information is saved per recording, and then the median is taken over the mean PSDs estimates per recording.

One issue with relative power, however, is that a notable change of absolute power in a single band can cause significant knock-on changes to relative power in other bands. Thus, peak frequencies in the power spectra of the LFP signals were identified using a previously described method (Donoghue et al., 2020) to evaluate peaks in the PSD while taking into account the $1/f$ decay of power without resorting to relative powers. The foof algorithm (version 1.0.0) was used to parameterise neural power spectra. Settings for the algorithm were set as: peak width limits: 1 – 8; max number of peaks: 3; minimum peak height: 0.2; peak threshold: 2.0; and aperiodic mode: fixed. Power spectra were parameterised across the frequency range 1 to 90Hz.

5.5.4 Speed theta analysis

The relationship of running speed to the LFP was analysed based on a previously described method (Ahmed & Mehta, 2012). To briefly summarise, speed was calculated from the smoothed position data of the rat, which is tracked by the light emitting diodes (LEDs) on the rat's head. Then, the recording sessions were binned into sequential

200ms bins. The average speed of the rat in these bins was calculated, along with the power spectrum. The relative theta bandpower is computed by dividing the power in the theta band in that bin by the total power of the LFP. Then, to compute the average power at a given speed, all time windows with that average speed are used – and the power is averaged over these time windows.

5.5.5 Coherence analysis

Coherence was considered between LFPs recorded from the subiculum, and ipsilateral RSC. Magnitude squared coherence was calculated by dividing the squared cross-spectral density estimate by the individual power spectral density estimates:

$$C_{xy}(\omega) = \frac{|S_{xy}(\omega)|^2}{S_{xx}(\omega)S_{yy}(\omega)} \quad (5.1)$$

where S_{xy} is the cross-spectral density estimate, and S_{xx} and S_{yy} are the individual power spectral density estimates. In open-field recordings, the coherence was calculated in 2-second segments with 1-second overlaps, where a Hann window was used, and then averaged over all segments.

T-maze activity was manually split into the start time for a trial, the time at which the decision of left or right was made (which was defined as the time at which the rat committed to the direction left or right, without turning back), and the ending time of trial. Coherence around the choice time was evaluated (similar to, Jones & Wilson, 2005), splitting the LFPs recorded during T-maze tasks into three groups: start, choice, and end. For analysis of coherence during T-maze tasks, these three splits were computed two ways for comparison. In both cases, choice was a 1 second chunk of the LFPs around the decision time (where the rat physically committed to turning left or right), with 0.8s before the decision time and 0.2s after the decision time. In the first method, start was delineated as when the rat entered the maze to the beginning of the choice data, and end taken as the termination of the choice data to when the rat had reached the end of the maze. In the second method, the data were analysed in a maximum 3-second-long period with 1 second of data taken at both sides of the central 1 second of choice data. The results from these methods were consistent as most trials completed quickly, thus, only the results from the first method were presented. In T-maze analysis, due to smaller windows of data, 0.25-second segments were instead used that overlap by 0.125 seconds with a Hann window, and then averaged over all segments.

5.5.6 Sharp-wave ripple detection during resting

In the absence of electromyogram (EMG) recordings, we could not unquestionably determine whether the rat was in a slow-wave sleep state, or a quiet and restful state. As such, we restricted all sleep analyses to bouts of rest. Rest periods were defined as being at least 60 continuous seconds during which the rat had a theta/delta power ratio < 2 (indicative of non-REM sleep in hippocampal signals, Louie & Wilson, 2001) and a speed of less than 2.5cm/s (following Muessig et al., 2019). As there were different recording layouts for different rats, the theta/delta power ratio was calculated for either the subiculum or CA1 dependent on which area had the most tetrodes implanted in it (favouring subiculum in the case of a tie). In addition to the automatic detection above, the experimenter additionally marked recordings to indicate the visual appearance of the rat (e.g., sleeping, resting, awake). The automatic detection closely corresponded with the manual markings.

Two different methods were applied to detect sharp-wave ripples (SWRs). In both methods, the 4800Hz LFPs were first restricted to a period of rest, and then filtered in the 150 to 250Hz band using a FIR filter with a lower transition bandwidth of 37.50Hz and an upper transition bandwidth of 62.50Hz. For the first method, LFPs from all tetrodes in a single brain region were squared and summed, then smoothed by a Gaussian kernel ($\sigma = 4\text{ms}$) and the square root of the smoothed sum was analysed (Kay et al., 2016). SWRs were detected when the smoothed sum exceeded 3 standard deviations of the mean in that rest period for at least 15ms. The second method is similar, except a smooth envelope was calculated for each of the LFP using the Hilbert transform instead of squaring and summing the signals (Karlsson & Frank, 2009). Similarly, SWRs were detected when the smooth envelope exceeded 4 standard deviations of the rest period mean for at least 15ms. The results from both methods were well matched, though the Kay et al. method was more conservative than the Karlsson and Frank method and slightly more accurate based on visual inspection and, thus, will be the reported in the plots and statistical tests. However, statistical tests from the Karlsson and Frank method were in agreement.

5.5.7 Analysing the coupling of spikes to the local field potential

To consider the relationship between spikes from single units in the dorsal subiculum and LFPs recorded from the dorsal subiculum, the STA and SFC were calculated (Fries et al., 1997). LFPs were considered in one-second intervals around the spike time. Since the STA and SFC are both dependent on the spike rate (with a lower spike rate leading to a lower SFC, Lepage et al., 2011) and the cells recorded from ATNx rats

generally had a higher spike rate than from control rats (Table 4.2), we controlled for the spike rate in spike-LFP analyses to compare SFC between control and ATNx rats. We use the method presented in NeuroChaT (Islam et al., 2019) to randomly select at most 50 spikes from the full spike train and calculate the STA and SFC in a bootstrapping procedure repeated 500 times for each cell. The resulting STA and SFC are then averaged per cell by taking the mean of these 500 bootstrapped samples.

To further test the STA and SFC measures, both measures are also calculated on randomly shifted spike trains. A high SFC value from a shifted spike train would suggest that a high SFC for this cell and the LFP occurred by chance, acting as a control procedure (Destexhe et al., 1999). This shuffling procedure was performed 50 times per cell, though the number of bootstrapping procedures performed was reduced to 20. Thus, there were $50 \cdot 20 = 1000$ STAs and SFCs computed overall per cell for the shuffling analysis. The shuffling procedure indicated that the observed SFC and STA did not occur from an intrinsic factor of the data.

The spike-LFP relationship was investigated during open-field recordings in control and ATNx animals and compared between these two conditions using a Mann-Whitney U test to compare the peak values of SFC within the theta band. The theta band was chosen for comparison based on the strong relationship between theta and spiking activity in the hippocampus, and the reliable theta peak in the power spectrum identified with FOOOF. Similarly, the spike-LFP relationship was compared before and after the effects of muscimol.

5.5.8 Statistical testing

To select the appropriate manner of comparing the distribution of independent or paired data these steps were followed. Firstly, the distributions in both groups were plotted and visually inspected for skewness and assessed for normality using the Shapiro-Wilk test. Additionally, the two distributions were tested for equal variance using Levene's test.

- For independent samples (e.g., control vs lesion rats) with normal distributions and equal variance, an unpaired student's t -test was applied.
- For independent samples with normal distributions and non-equal variance, Welch's t -test was applied.
- For independent samples with non-normal distributions, a Mann-Whitney U test was applied.

- For paired samples with normal distributions, a student's t -test was applied to the paired differences.
- For paired samples with non-normal distributions, a Wilcoxon signed-rank test was applied to the paired differences.

Tests were performed in both Python using Pingouin v0.5.2 (Vallat, 2018) and Jasp v.0.16.2 (JASP Team, 2022) for comparison, except for the Bayesian tests which were exclusively performed in Jasp. Where possible, a related Bayesian test was also applied in addition to the classical frequentist tests and the Bayes factor reported. The Bayes factor represents the ratio:

$$\text{BF}_{10} = \frac{p(d|h_1)}{p(d|h_0)} \quad (5.2)$$

this value $\text{BF}_{10} > 1$ reflects how much more strongly the data support the alternative hypothesis (H_1) over the null hypothesis (H_0), while a value less than 1 prefers the null hypothesis. A value over 4 indicates substantial support for the alternative hypothesis over the null hypothesis. As well as reporting the test statistic and p -value, the CLES is regularly reported. The CLES is the probability that a score sampled at random from one distribution will be greater than a score sampled from another distribution (McGraw & Wong, 1992). For a paired samples test, this represents the proportion of pairs in which $x > y$. This was calculated with the brute-force version of the formula using Pingouin (Vallat, 2018; Vargha & Delaney, 2000).

5.5.9 Data and code sharing

All code used to analyse this data and produce the figures is freely available on GitHub at <https://github.com/seankmartin/atn-sub-lfp-workflow>. The intermediate results (statistical tests, plots, etc.) resulting from this work are available on the Open Science Framework at <https://osf.io/ytjmd/>. The data from all rats, which has been converted from Axona format to NWB format, is fully described in this documentation (<http://bit.ly/sub-lfp-data>), which includes information on how to access the data. This data includes metadata and intermediate processed results to ease the analysis of the data. To inspect the contents of one of these NWB files, we recommend using Simuran or NWB Explorer. Downloading the data and running the shared code can be used to reproduce the results presented.

References

- Aggleton, J. P., & Brown, M. W. (1999). Episodic memory, amnesia, and the hippocampal-anterior thalamic axis. *The Behavioral and Brain Sciences*, *22*(3), 425–444, discussion 444–489.
- Aggleton, J. P., & Nelson, A. J. D. (2015). Why do lesions in the rodent anterior thalamic nuclei cause such severe spatial deficits? *Neuroscience and Biobehavioral Reviews*, *54*, 131–144.
- Ahmed, O. J., & Mehta, M. R. (2012). Running speed alters the frequency of hippocampal gamma oscillations. *The Journal of Neuroscience*, *32*(21), 7373.
- Amaral, D. G., Dolorfo, C., & Alvarez-Royo, P. (1991). Organization of CA1 projections to the subiculum: A PHA-l analysis in the rat. *Hippocampus*, *1*(4), 415–435.
- Amzica, F., & Steriade, M. (1998). Electrophysiological correlates of sleep delta waves1correspondence and reprint requests may be addressed to either author.1. *Electroencephalography and Clinical Neurophysiology*, *107*(2), 69–83.
- Berke, J. D., Hetrick, V., Breck, J., & Greene, R. W. (2008). Transient 23–30 hz oscillations in mouse hippocampus during exploration of novel environments. *Hippocampus*, *18*(5), 519–529.
- Bubb, E. J., Metzler-Baddeley, C., & Aggleton, J. P. (2018). The cingulum bundle: Anatomy, function, and dysfunction. *Neuroscience and Biobehavioral Reviews*, *92*, 104–127.
- Buzsáki, G. (2002). Theta oscillations in the hippocampus. *Neuron*, *33*(3), 325–40.
- Buzsáki, G., Anastassiou, C. A., & Koch, C. (2012). The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, *13*(6), 407–420.
- Buzsáki, G., & Wang, X.-J. (2012). Mechanisms of gamma oscillations. *Annual review of neuroscience*, *35*, 203.
- Cacucci, F., Lever, C., Wills, T. J., Burgess, N., & O’Keefe, J. (2004). Theta-modulated place-by-direction cells in the hippocampal formation in the rat. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *24*(38), 8265–8277.
- Cardin, J. A., Carlén, M., Meletis, K., Knoblich, U., Zhang, F., Deisseroth, K., Tsai, L.-H., & Moore, C. I. (2009). Driving fast-spiking cells induces gamma rhythm and controls sensory responses. *Nature*, *459*(7247), 663–667.
- Carlesimo, G. A., Lombardi, M. G., & Caltagirone, C. (2011). Vascular thalamic amnesia: A reappraisal. *Neuropsychologia*, *49*(5), 777–789.
- Cho, J., & Sharp, P. E. (2001). Head direction, place, and movement correlates for cells in the rat retrosplenial cortex. *Behavioral Neuroscience*, *115*, 3–25.

- Colgin, L. L. (2013). Mechanisms and functions of theta rhythms. *Annual Review of Neuroscience*, *36*, 295–312.
- Colgin, L. L., Denninger, T., Fyhn, M., Hafting, T., Bonnevie, T., Jensen, O., Moser, M.-B., & Moser, E. I. (2009). Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature*, *462*(7271), 353–357.
- Danet, L., Barbeau, E. J., Eustache, P., Planton, M., Raposo, N., Sibon, I., Albucher, J.-F., Bonneville, F., Peran, P., & Pariente, J. (2015). Thalamic amnesia after infarct: The role of the mammillothalamic tract and mediodorsal nucleus. *Neurology*, *85*(24), 2107–2115.
- Destexhe, A., Contreras, D., & Steriade, M. (1999). Spatiotemporal analysis of local field potentials and unit discharges in cat cerebral cortex during natural wake and sleep states. *Journal of Neuroscience*, *19*(11), 4595–4608.
- Donoghue, T., Haller, M., Peterson, E. J., Varma, P., Sebastian, P., Gao, R., Noto, T., Lara, A. H., Wallis, J. D., Knight, R. T., Shestyuk, A., & Voytek, B. (2020). Parameterizing neural power spectra into periodic and aperiodic components. *Nature Neuroscience*, *23*(12), 1655–1665.
- França, A. S. C., do Nascimento, G. C., Lopes-dos-Santos, V., Muratori, L., Ribeiro, S., Lobão-Soares, B., & Tort, A. B. L. (2014). Beta2 oscillations (23–30 Hz) in the mouse hippocampus during novel object recognition. *European Journal of Neuroscience*, *40*(11), 3693–3703.
- Fries, P., Roelfsema, P. R., Engel, A. K., Konig, P., & Singer, W. (1997). Synchronization of oscillatory responses in visual cortex correlates with perception in interocular rivalry. *Proceedings of the National Academy of Sciences*, *94*(23), 12699–12704.
- Frost, B. E., Martin, S. K., Cafalchio, M., Islam, M. N., Aggleton, J. P., & O’Mara, S. M. (2021). Anterior thalamic inputs are required for subiculum spatial coding, with associated consequences for hippocampal spatial memory. *Journal of Neuroscience*, *41*(30), 6511–6525.
- Gallego-Carracedo, C., Perich, M. G., Chowdhury, R. H., Miller, L. E., & Gallego, J. Á. (2022). Local field potentials reflect cortical population dynamics in a region-specific and frequency-dependent manner (S. Ostoic & R. L. Calabrese, Eds.). *eLife*, *11*, e73155.
- Geisler, C., Robbe, D., Zugaro, M., Sirota, A., & Buzsáki, G. (2007). Hippocampal place cell assemblies are speed-controlled oscillators. *Proceedings of the National Academy of Sciences*, *104*(19), 8149–8154.
- Gramfort, A., Luessi, M., Larson, E., Engemann, D., Strohmeier, D., Brodbeck, C., Goj, R., Jas, M., Brooks, T., Parkkonen, L., & Hämäläinen, M. (2013). MEG and EEG data analysis with MNE-python. *Frontiers in Neuroscience*, *7*.
- Grossberg, S. (2009). Beta oscillations and hippocampal place cell learning during exploration of novel environments. *Hippocampus*, *19*(9), 881–885.

- Harmony, T. (2013). The functional significance of delta oscillations in cognitive processing. *Front Integr Neurosci*, 7, 83.
- Islam, M. N., Martin, S. K., Aggleton, J. P., & O'Mara, S. M. (2019). NeuroChaT: A toolbox to analyse the dynamics of neuronal encoding in freely-behaving rodents in vivo. *Wellcome Open Research*, 4, 196.
- Iwasaki, S., Sasaki, T., & Ikegaya, Y. (2021). Hippocampal beta oscillations predict mouse object-location associative memory performance. *Hippocampus*, 31(5), 503–511.
- Jackson, J., Goutagny, R., & Williams, S. (2011). Fast and slow gamma rhythms are intrinsically and independently generated in the subiculum. *Journal of Neuroscience*, 31(34), 12104–12117.
- Jankowski, M. M., Passecker, J., Islam, M. N., Vann, S., Erichsen, J. T., Aggleton, J. P., & O'Mara, S. M. (2015). Evidence for spatially-responsive neurons in the rostral thalamus. *Frontiers in Behavioral Neuroscience*, 9, 256.
- JASP Team. (2022). JASP (Version 0.16.4)[Computer software].
- Jones, M. W., & Wilson, M. A. (2005). Theta rhythms coordinate hippocampal-prefrontal interactions in a spatial memory task. *PLoS Biol*, 3(12), e402.
- Karlsson, M. P., & Frank, L. M. (2009). Awake replay of remote experiences in the hippocampus. *Nature Neuroscience*, 12(7), 913–918.
- Kay, K., Sosa, M., Chung, J. E., Karlsson, M. P., Larkin, M. C., & Frank, L. M. (2016). A hippocampal network for spatial coding during immobility and sleep. *Nature*, 531(7593), 185–190.
- Keene, C. S., & Bucci, D. J. (2008). Neurotoxic lesions of retrosplenial cortex disrupt signaled and unsignaled contextual fear conditioning. *Behavioral Neuroscience*, 122(5), 1070–1077.
- Lepage, K. Q., Kramer, M. A., & Eden, U. T. (2011). The dependence of spike field coherence on expected intensity. *Neural Computation*, 23(9), 2209–2241.
- Long, S., Ding, R., Wang, J., Yu, Y., Lu, J., & Yao, D. (2021). Sleep quality and electroencephalogram delta power. *Frontiers in Neuroscience*, 15.
- Louie, K., & Wilson, M. A. (2001). Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron*, 29(1), 145–156.
- Maguire, E. (2001). The retrosplenial contribution to human navigation: A review of lesion and neuroimaging findings. *Scandinavian Journal of Psychology*, 42(3), 225–238.
- Matulewicz, P., Ulrich, K., Islam, M. N., Mathiasen, M. L., Aggleton, J. P., & O'Mara, S. M. (2019). Proximal perimeter encoding in the rat rostral thalamus. *Scientific Reports*, 9(1), 2865.
- McGraw, K. O., & Wong, S. P. (1992). A common language effect size statistic. *Psychological bulletin*, 111(2), 361.

- Montgomery, S. M., & Buzsáki, G. (2007). Gamma oscillations dynamically couple hippocampal CA3 and CA1 regions during memory task performance. *Proceedings of the National Academy of Sciences*, *104*(36), 14495–14500.
- Moser, E. I., Moser, M.-B., & McNaughton, B. L. (2017). Spatial representation in the hippocampal formation: A history. *Nature neuroscience*, *20*(11), 1448–1464.
- Muessig, L., Lasek, M., Varsavsky, I., Cacucci, F., & Wills, T. J. (2019). Coordinated emergence of hippocampal replay and theta sequences during post-natal development. *Current Biology*, *29*(5), 834–840.e4.
- Nelson, A. J. D., Kinnavane, L., Amin, E., O'Mara, S. M., & Aggleton, J. P. (2020). Deconstructing the direct reciprocal hippocampal-anterior thalamic pathways for spatial learning. *Journal of Neuroscience*, *40*(36), 6978–6990.
- Nitzan, N., McKenzie, S., Beed, P., English, D. F., Oldani, S., Tukker, J. J., Buzsáki, G., & Schmitz, D. (2020). Propagation of hippocampal ripples to the neocortex by way of a subiculum-retrosplenial pathway. *bioRxiv*, 2020.02.27.966770.
- Nuñez, A., & Buño, W. (2021). The theta rhythm of the hippocampus: From neuronal and circuit mechanisms to behavior. *Frontiers in Cellular Neuroscience*, *15*.
- O'Mara, S. M., & Aggleton, J. P. (2019). Space and memory (far) beyond the hippocampus: Many subcortical structures also support cognitive mapping and mnemonic processing. *Frontiers in Neural Circuits*, *13*.
- O'Neill, J., Senior, T. J., Allen, K., Huxter, J. R., & Csicsvari, J. (2008). Reactivation of experience-dependent cell assembly patterns in the hippocampus. *Nature Neuroscience*, *11*(2), 209–215.
- Pesaran, B., Vinck, M., Einevoll, G. T., Sirota, A., Fries, P., Siegel, M., Truccolo, W., Schroeder, C. E., & Srinivasan, R. (2018). Investigating large-scale brain dynamics using field potential recordings: Analysis and interpretation. *Nature Neuroscience*, *21*(7), 903–919.
- Richard, G. R., Titiz, A., Tyler, A., Holmes, G. L., Scott, R. C., & Lenck-Santini, P.-P. (2013). Speed modulation of hippocampal theta frequency correlates with spatial memory performance. *Hippocampus*, *23*(12), 1269–1279.
- Schneider, M., Brogini, A. C., Dann, B., Tzanou, A., Uran, C., Sheshadri, S., Scherberger, H., & Vinck, M. (2021). A mechanism for inter-areal coherence through communication based on connectivity and oscillatory power. *Neuron*, *109*(24), 4050–4067.e12.
- Schultheiss, N. W., Schlecht, M., Jayachandran, M., Brooks, D. R., McGlothan, J. L., Guilarte, T. R., & Allen, T. A. (2020). Awake delta and theta-rhythmic hippocampal network modes during intermittent locomotor behaviors in the rat. *Behavioral Neuroscience*, *134*(6), 529.

- Shibata, H. (1993). Direct projections from the anterior thalamic nuclei to the retrohippocampal region in the rat. *Journal of Comparative Neurology*, *337*(3), 431–445.
- Solomon, E. A., Stein, J. M., Das, S., Gorniak, R., Sperling, M. R., Worrell, G., Inman, C. S., Tan, R. J., Jobst, B. C., Rizzuto, D. S., & Kahana, M. J. (2019). Dynamic theta networks in the human medial temporal lobe support episodic memory. *Current Biology*, *29*(7), 1100–1111.e4.
- Spiers, H. J., Maguire, E. A., & Burgess, N. (2001). Hippocampal amnesia. *Neurocase*, *7*(5), 357–382.
- Sripanidkulchai, K., & Wyss, J. M. (1986). Thalamic projections to retrosplenial cortex in the rat. *Journal of Comparative Neurology*, *254*(2), 143–165.
- Sugar, J., Witter, M., van Strien, N., & Cappaert, N. (2011). The retrosplenial cortex: Intrinsic connectivity and connections with the (para)hippocampal region in the rat. an interactive connectome. *Frontiers in Neuroinformatics*, *5*.
- Taube, J. S. (1995). Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *15*(1), 70–86.
- Taube, J. S., Muller, R. U., & Ranck, J. B. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. i. description and quantitative analysis. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *10*(2), 420–435.
- Todd, T. P., & Bucci, D. J. (2015). Retrosplenial cortex and long-term memory: Molecules to behavior. *Neural Plasticity*, *2015*, 414173.
- Trimper, J. B., Galloway, C. R., Jones, A. C., Mandi, K., & Manns, J. R. (2017). Gamma oscillations in rat hippocampal subregions dentate gyrus, CA3, CA1, and subiculum underlie associative memory encoding. *Cell Reports*, *21*(9), 2419–2432.
- Vallat, R. (2018). Pingouin: Statistics in Python. *Journal of Open Source Software*, *3*(31), 1026.
- Vann, S. D., Aggleton, J. P., & Maguire, E. A. (2009). What does the retrosplenial cortex do? *Nature Reviews. Neuroscience*, *10*(11), 792–802.
- Vargha, A., & Delaney, H. D. (2000). A critique and improvement of the common language effect size statistics of mcgraw and wong. *Journal of Educational and Behavioral Statistics*, *25*(2), 101–132.
- Vinck, M., Uran, C., Spyropoulos, G., Onorato, I., Broggin, A. C., Schneider, M., & Canales-Johnson, A. (2023). Principles of large-scale neural interactions. *Neuron*, *111*(7), 987–1002.

- Welch, P. (1967). The use of fast fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Transactions on Audio and Electroacoustics*, *15*(2), 70–73.
- Womelsdorf, T., Vinck, M., Leung, S., & Everling, S. (2010). Selective theta-synchronization of choice-relevant information subserves goal-directed behavior. *Frontiers in Human Neuroscience*, *4*.

6 Recording synaptically connected neurons across brain regions

Introduction: Large-scale simultaneous *in vivo* recordings of neurons in multiple brain regions prompt the question of whether we can capture direct interactions between these neurons. This, in turn, opens the possibility of identifying inter-regional communication rules during behavioural tasks, assuming we can detect coordinated activity among neurons in connected brain regions. Here, we take a first step towards addressing this issue by estimating the number of connected neurons that can be recorded through simultaneous recordings in anatomically connected brain regions.

Methods: To investigate this, we construct a directed acyclic graph that represents the chemical synapses between neurons in the specific brain regions under study. Through Monte Carlo simulations and statistical analyses, we develop a method for calculating the probability distribution of “recordable” neurons that are connected. We apply this method to neurons in the neocortex, following the Blue Brain model of the mouse neocortex connectome, as well as to neurons in the hippocampus and subiculum, based on statistical data from anatomical studies.

Results: Our findings reveal that the probability of recording a connected pair of neurons is notably high in well-connected brain regions like the hippocampus and subiculum. In the neocortex, the probability is relatively lower but still significant. Furthermore, along chains of synapses, such as disynaptic or trisynaptic connections, the probability of recording connected neurons increases dramatically.

Conclusions: By employing high-throughput neural recording technologies in well-connected brain areas, such as specific regions within the auditory cortex, we can enhance the likelihood of capturing a substantial number of neuron pairs that are directly connected across brain regions. Consequently, these techniques hold promise for investigating and validating hypotheses related to interregional communication and source transformation rules.

6.1 Neural recording technologies and source transformation rules

In the previous chapter, when analysing interactions between brain areas, we had an intuitive understanding that the number of cells obtained and the overall recording fidelity would not allow for direct spike-to-spike comparisons between different brain regions (chapter 5). But what recording fidelity would be required to achieve such a comparison with a high level of confidence? How many cells would we need to record, and how would that number of cells change with the recording modality used and the regions being considered?

The technique of choice to understand neuronal circuit dynamics *in vivo* continues to be intracranial recordings of single and multiple neurons (and of associated signals, such as local-field potentials). Recent technical advances in neuronal recording methodologies (such as Neuropixels probes; Jun et al., 2017; Steinmetz et al., 2021) allow recording of large numbers of single neurons within and between multiple brain regions simultaneously (Allen et al., 2019; Siegle et al., 2021; Steinmetz et al., 2019). These approaches offer a level of analysis previously impossible using tetrode-based recordings. For example, simultaneously recording c. 3000 neurons across the brain using Neuropixels probes and c. 10,000 neurons in visual cortex using calcium imaging (Stringer et al., 2019).

These large-scale, inter-regional, conjoint recordings leave the problem of understanding granular information processing rules between neurons in connected brain regions. Understanding these information transformation rules may enable us to probe how downstream neuronal observers take advantage of neural coding (Buzsáki, 2004; Buzsáki & Tingley, 2018). But what are the source transformation rules applied to given inputs from one area to another, and can we expect to infer them, given the neuronal sampling current methodologies afford? Here, we provide a first step in approaching the issue of source transformation rules *in vivo*. We focus on connections between neurons, as opposed to parameters such as synaptic weights and firing patterns, because structural connectivity provides the basis for much of the communication between connected brain regions (Goulas et al., 2017), though not all (e.g., ephaptic communication, Anastassiou & Koch, 2015; Jefferys, 1995; Martinez-Banaclocha, 2018).

Many previous studies of inter-areal interactions have analysed the spiking activity of pairs of neurons in different areas (Semedo et al., 2019). Given the importance of structural connectivity to inter-areal interactions, estimating the proportion of neurons that might be connected in a conjoint recording (and how direct the connections are) is key to understanding inter-areal interactions at the neuronal level. However, neuronal connection probabilities on a large scale have received little attention. The probability

of connection between pairs of neurons at a small distance from each other (usually on a scale of a few hundred micrometres), has been well studied (Braitenberg & Schüz, 1998; Holmgren et al., 2003; Kalisman et al., 2003; Liley & Wright, 1994). Here, we provide a method for estimating the proportion of neurons that might be connected in conjoint recordings from anatomically-connected brain regions.

We consider the following question: when conjointly recording neurons from anatomically connected brain areas *in vivo*, how likely are we to record neurons from one brain area that receive chemical connections from recorded neurons in the other brain area (Figure 6.1)? We study several neurons in multiple brain regions, with simplified connections between neurons in these regions and derive the probability distribution for the number of connected neurons that might be recorded between two populations of neurons. This method is applied to network models of the rat hippocampal formation (hippocampus proper and subiculum), and areas of a pre-existing mouse neocortex model (Reimann et al., 2019).

Our findings demonstrate areas in which current neural recording capabilities allow for recording a sample of neurons such that roughly $\geq 50\%$ of the recorded downstream neurons receive direct connections from the recorded upstream neurons. This is primarily for well-connected areas and high yield recording devices with close to a hundred neurons in each brain region. However, when considering chains of connections, such as disynaptic or trisynaptic connections (those involving one or two intermediate neurons, respectively), it is highly likely that only tens of neurons need to be recorded to investigate such effects on downstream neurons. We conclude with an evaluation of the model assumptions, the next steps for the method, and implications for the analysis of inter-regional neural processing.

6.2 Calculating the probability of recording connected neurons across brain regions

Given a model M for the direct synaptic connections within and between two brain regions A and B , we estimate the probability distribution of the number of neurons recorded in B that receive a connection from a neuron recorded in A , when randomly sampling k neurons from A and m neurons from B .

6.2.1 Model inputs and assumptions

The following information is required in our method:

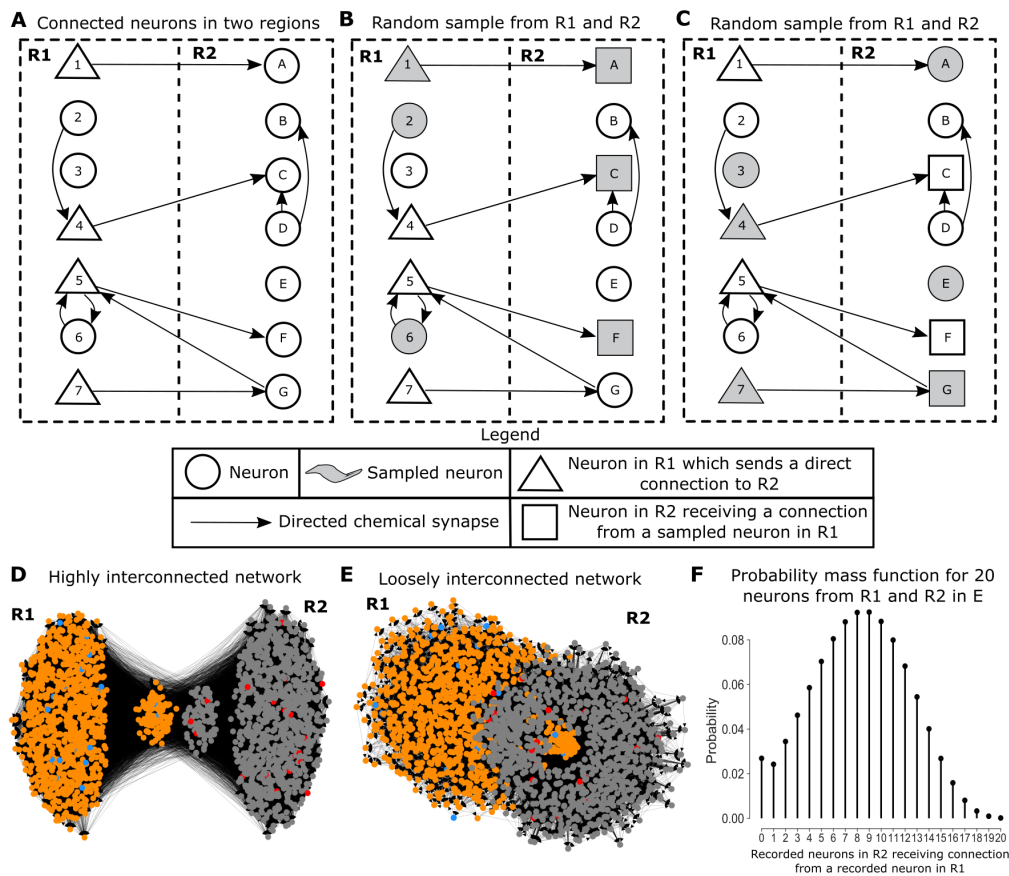


FIGURE 6.1: Illustrating the problem of estimating the probability of conjointly recording connected neurons.

(A) Connections between neurons in two regions, where the neurons in R1 which send direct connections to R2 are marked as triangles and connections indicated by arrows.

(B) Randomly recording three neurons from R1 and three neurons from R2. All three of the sampled neurons in R2 receive a connection from a neuron sampled in R1, but two are indirect connections ($1 \rightarrow A, 2 \rightarrow 4 \rightarrow C, 6 \rightarrow 5 \rightarrow F$).

(C) As in B, but a new recording results in a different number of sampled neurons that are connected. Only one of the sampled neurons receives a connection from a neuron sampled in R1 ($7 \rightarrow G$).

(D, E) Examples of large-scale networks of neurons with one thousand neurons in each region. Blue and red dots indicate representative recorded neurons in R1 and R2.

(F) Using our method on the network in E, we derive the probability distribution of the number of recorded neurons in R2 which receive a direct connection from a neuron recorded in R1 when randomly recording 20 neurons from R1 and 20 neurons from R2.

- A, B : The two brain regions to consider.
- M : A model which describes the connections between the neurons in regions A and B . This is either a full connection matrix, such that $M(i, j) = 1$ indicates a connection between neurons i and j , or a statistical model indicating the distribution of connections related to the underlying topography (e.g., Reimann et al., 2019).
- D : The maximum geodesic distance between connected nodes. In other words, two neurons are considered connected if at most D synapses are needed to connect them. For instance, $D = 2$ would allow for direct connections and disynaptic connections (those involving an intermediate neuron).
- R_1, R_2, \dots, R_n : Information about the n recording devices used, such as target brain location and the 3D shape of the recording area. For example, a tetrode could be assumed to sample neurons in a $130\mu\text{m}$ radius sphere (Mechler et al., 2011).
- Ω : An atlas describing the 3D shape of the brain regions A and B (e.g., Lein et al., 2007).
- Γ : A model of the 3D location of each cell in A and B such that $\Gamma(i) \mapsto (x_i, y_i, z_i)$ (e.g., Erö et al., 2018).

Model assumptions

1. Synaptic connection strength is not considered, just the presence or absence of a synapse.
2. Multiple connections between two specific neurons are considered as one connection.
3. Neurons are randomly sampled from the overall population of neurons that are within range of the recording device(s).
4. The problem is considered as a fixed state, or point estimate – so, for instance, the dynamic nature of brain connectivity over time during chronic recordings is not considered.
5. We do not distinguish between inhibitory and excitatory connections. Setting up the model with statistics of, say, only inhibitory connections would then give probabilities for that type of connection.

6. Only chemical synapses are considered and not electrical synapses (gap junctions), because electric synapses could be considered as undirected connections, while chemical synapses are directed (Reimann et al., 2017).
7. We do not consider the issue of the probability that spiking from a neuron in area A causes spiking in a neuron in area B ; we just consider that they are anatomically connected.
8. Autapses are not considered, as they do not influence whether two neurons are synaptically connected.

6.2.2 Monte Carlo simulation

In the Monte Carlo simulation approach, the problem is modelled in three-dimensional space. Neurons are placed in x, y, z positions and associated with a directed graph that represents the underlying chemical synaptic connections between the neurons in A and B along with the local connections between neurons in the same brain area. Recording devices are simulated in these brain regions, and the neurons that could feasibly be recorded by the devices are noted. Finally, neurons are randomly sampled from within the proximity of the recording devices over many iterations and the neurons that are connected along chains of D or less synapses (e.g., $D = 1$ for direct connections) are noted. After many iterations (roughly 50,000), a stable probability distribution of the number of neurons recorded in B that receive a connection from A along D or fewer synapses is formed. However, this method relies on simulating recording from this graph many times and is slow to calculate in addition to requiring substantial amounts of computer memory to store the graph (see subsection 6.5.2 for more details).

6.2.3 Statistical estimation

In the statistical estimation approach, the distribution of connections between the neurons in A and B both inside and outside the proximity of the recording devices is required, along with the number of neurons that send connections. For instance, when recording from proximal CA1 and distal subiculum in the rat, the parameters can be set to describe how these parts of CA1 and subiculum are well connected, while connections outside of proximal CA1 and outside of distal subiculum are of a different nature. These distributions are then used to estimate $P(X_B = l \mid X_A = k)$, the probability of l neurons in B receiving at least one connection from a neuron recorded in A along D or fewer synapses, given that k neurons are recorded in A which send direct connections to B . Finally, $P(X_B = l \mid X_A = k)$ is used in combination with the

chain rule of probability and the hypergeometric distribution to estimate the probability distribution of the number of cells recorded in B that receive a connection from A . This statistically-based method is significantly faster to compute and has a lower memory and energy footprint, but is not as accurate as the Monte Carlo simulation method due to the necessity to collapse the full graph information into statistical connection distributions (see subsection 6.5.3 for more details).

6.3 Case studies: estimating connection probabilities in the mouse and rat brain

6.3.1 Connection probability in the neocortex of the mouse brain

A recent mouse neocortex model provides parametric rules for short- and long-range stochastic connections between neurons in a three-dimensional volume (Reimann et al., 2019). Such a model opens an opportunity for investigating candidates for simultaneous recordings in the neocortex; namely, those with a high connection probability. We apply our method to large graphs of neurons ($> 400,000$) with highly varying forms of connectivity from this neocortex model. We first detail the connection probability from primary visual area to lateral visual area (Figure 6.2). The influence of the recording device placement on the expected number of connected neurons is demonstrated with Neuropixel probes with an assumed yield of 79 neurons (the average of a Neuropixels probe in a single brain region, where only a portion of the probe is in that brain region, Jun et al., 2017). Placing the probes as in a recording following a previous study (Steinmetz et al., 2019) results in an expected percentage of neurons recorded in lateral visual area (VISl) receiving a connection from a neuron recorded primary visual area (VISp) equal to 10.08% (2.5th percentile = 2.5%, 97.5th percentile = 20.25%) while shifting the probes laterally results in an expected percentage of 8.35% (2.5th percentile = 1.26%, 97.5th percentile = 16.45%).

We further compute the expected number of connected neurons in different pairs of neocortical brain regions via Neuropixel probe recordings directed towards the centre of the two brain regions (Figure 6.3). For instance, 30% of neurons sampled from the posterior auditory cortex would be expected to receive input from at least one of the neurons sampled from the primary auditory cortex in a single Neuropixels probe recording (Figure 6.3C). However, even for regions with lower connection probability, such as primary motor area (MOp) to primary somatosensory area associated with lower limbic function (SSp-II), neurons in SSp-II receive many indirect connections from recorded MOp neurons along disynaptic or trisynaptic paths. With even a single neuron

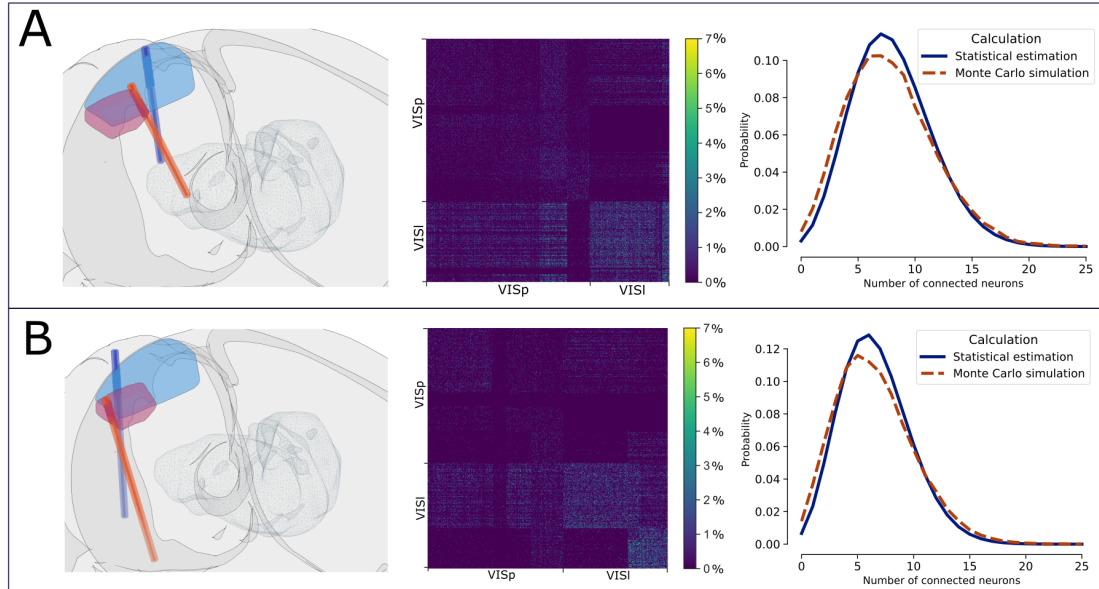


FIGURE 6.2: Probability of recording connected neurons in simultaneous primary visual area and lateral visual area recordings with different probe insertion angles and locations in the Blue Brain model of neocortex. The two recording locations result in different connectivity matrices between cells in proximity of the recording devices, and thus different probabilities of recording connected neurons. **(A)** places two Neuropixel probes used in these brain areas following the exact placement of two probes in a previous study (Steinmetz et al., 2019). **(B)** presents results from two simulated Neuropixel probes that have been moved in a lateral direction from **(A)** and rotated. In each panel:

(left) Visualising the recording span around 384 recording sites of two Neuropixel probes in VISp (primary visual area) and VISl (lateral visual area). Coloured points indicate cells in proximity to the recording sites.

(middle) Visualising the connectivity matrix between VISp and VISl in 10×10 blocks, where $x\%$ indicates that x out of 100 possible connections between these two groups of 10 neurons are formed.

(right) Calculating the expected number of neurons recorded in VISl that receive direct afferent connections from at least one neuron recorded in VISp, with 79 neurons recorded in VISl and VISp. Due to the similarity of efferent direct connections from VISp to VISl between the neural probes in cases (A) and (B) the resulting probability distributions are also similar.

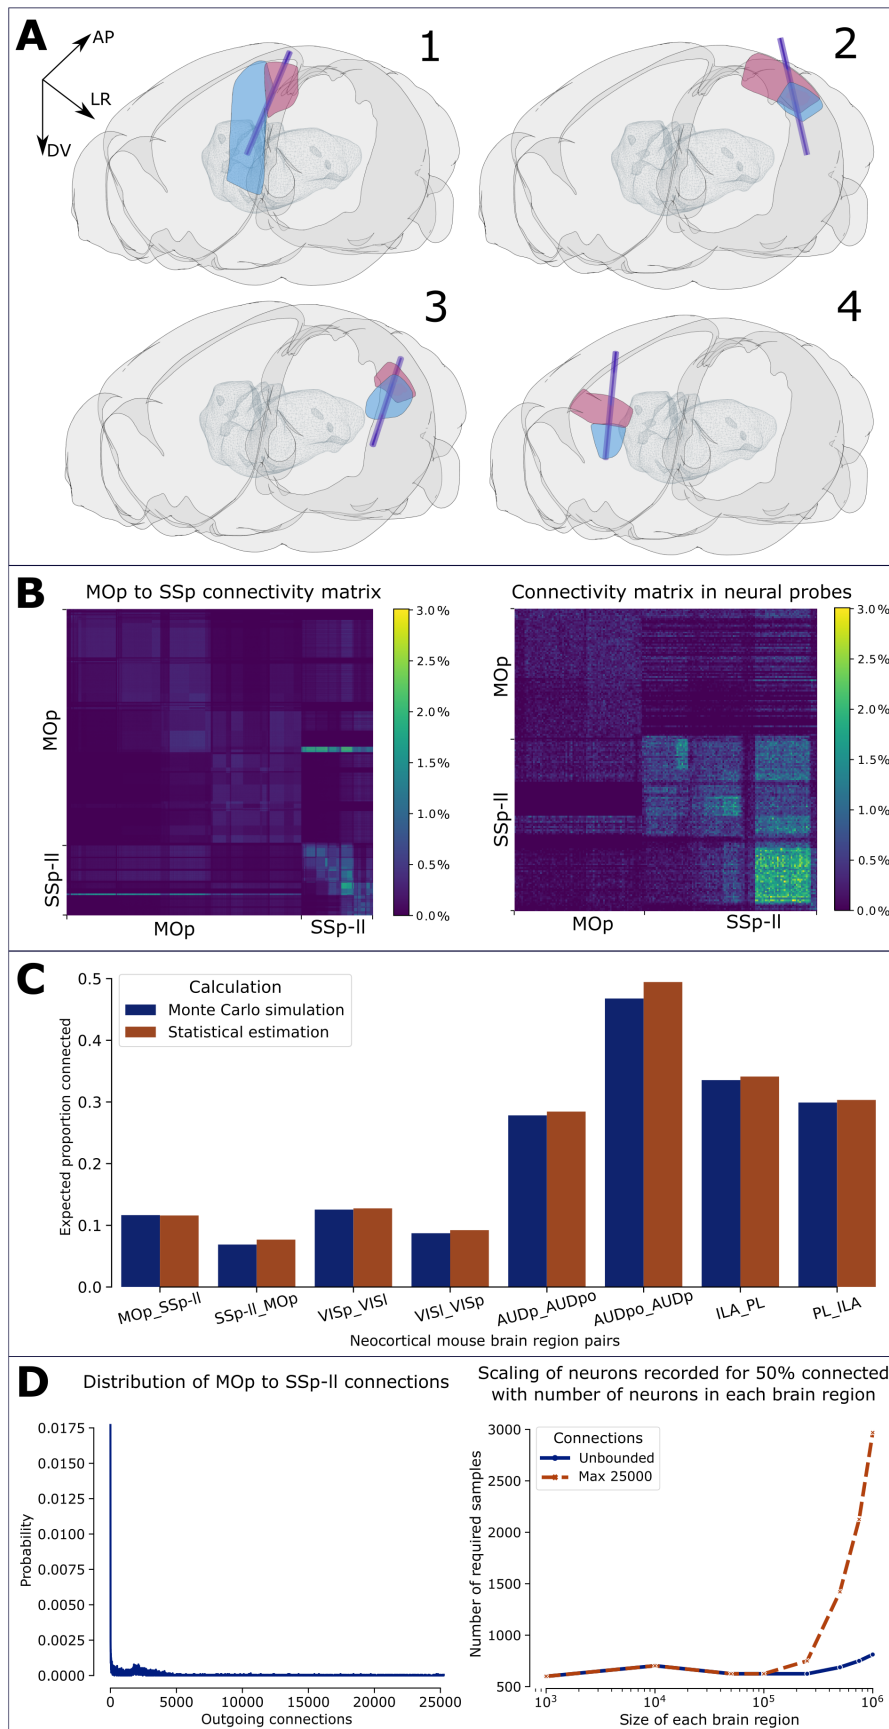


FIGURE 6.3: Caption continued on next page...

FIGURE 6.3: Modelling the probability of simultaneously recording neurons in mouse neocortex.

(A) Recording span of Neuropixel probes in different neocortical areas. Areas: 1 – MOp and SSp-ll, 2 – VISp and VISl, 3 – AUDp and AUDpo, 4 – ILA and PL. Axes: AP – anterior to posterior, LR – left to right, DV – dorsal to ventral.

(B) *left*: Connectivity matrix for all cells in MOp and SSp-ll, in 150×150 blocks (346366 cells in MOp and 107548 cells in SSp-ll). *right*: Connectivity matrix for only cells in proximity of the neural probe between MOp and SSp-ll in (A), in 40×40 blocks (2922 cells in MOp and 4100 cells in SSp-ll). Here, $x\%$ indicates that $x\%$ of the possible connections between two groups of 150/40 neurons are formed.

(C) Expected proportion of directly connected neurons with 79 samples in each pair of brain regions with neural probes as depicted in (A). R1_R2 indicates direct connections from R1 to R2.

(D) *left*: MOp to SSp-ll efferent connection distribution. *right*: Estimating the number of neurons that would need to be recorded from MOp and SSp-ll for 50% of the recorded neurons to be connected if both had the number of neurons shown on the x-axis. The unbounded line indicates the efferent connection distribution is scaled with the size of the brain regions, while the Max 25000 line indicates that no neuron can project to more than 25000 neurons.

Key: MOp, SSp-ll, primary motor area and primary somatosensory area associated with lower limbic function; VISp, VISl, primary and lateral visual areas; AUDp, AUDpo – primary and posterior auditory areas; ILA, PL – infralimbic and prelimbic areas.

sampled from each region they are likely connected by a trisynaptic connection, and with just 7 sampled neurons in each region, all 7 neurons in SSp-ll are likely to receive a disynaptic connection from a sampled neuron in MOp.

By simulating a network of neurons between two regions based on MOp and SSp-ll connectivity statistics (5% of MOp neurons project to SSp-ll neurons and those neurons project to a mean of 2% of the total SSp-ll population), but varying the size of the two brain areas, we can estimate how these connection probabilities scale to larger brains than mice brains. With 100,000 total neurons in both regions (Figure 6.3D), approximately 600 neurons would need to be recorded from both regions for 50% of the recorded neurons in the SSp-ll like region to receive connections from the neurons recorded in the MOp like region. 1,000,000 neurons in each region, this increases marginally, up to the need to sample about 800 neurons from each region. Alternatively, if we scale up the size of the brain areas to 1,000,000 neurons in each region, but fix the distribution such that the maximum outgoing connections from a single neuron is 25,000 (supposing the existence of a fundamental outgoing connection upper bound), then close to 3000 neurons would need to be recorded from each brain area to expect about 50% of the recorded neurons in the SSp-ll like region to receive connections from the neurons recorded in the MOp like region.

Source	Rat hippocampal formation statistic
Andersen et al. (2006)	390,000 and 290,000 principal neurons in CA1 and subiculum, respectively.
Amaral et al. (1990)	CA3 contains 303,930 pyramidal cells.
Bezaire and Soltesz (2013)	89% (347,100) of neurons in CA1 are pyramidal cells.
Ropireddy et al. (2012)	CA1 occupies a volume of 18.8mm^3 , CA3 12.2mm^3 .
West et al. (1978)	The subiculum occupies a volume of 10mm^3 .
Amaral et al. (1990)	A CA1 pyramidal cell might receive input from about 1.8% of the CA3 pyramidal cell population.
Arszovszki et al. (2014)	Pyramidal cells in ventral CA1 commonly have terminals in ventral subiculum, and some projections bifurcate (the apical dendrite of ventral CA1 cells bifurcated in 3 out of 12 pyramidal cells; this would possibly affect statistics of outputs from ventral CA1; the main axon of a ventral CA1 pyramidal cell branches into three different routes towards lateral septum, the amygdalar complex/lateral entorhinal region, and subiculum/entorhinal cortex).
Reimann et al. (2019)	The number of synaptic projections is not identical across all neurons.
Mechler et al. (2011)	A tetrode samples neurons in a $130\mu\text{m}$ radius sphere (volume 0.0092mm^3).
Gray et al. (1995)	A tetrode (e.g., in visual cortex) has a mean yield of about 5 cells.

TABLE 6.1: Rat hippocampal formation neuron count estimates, connectivity assumptions, and tetrode sampling assumptions with references.

6.3.2 Connection probability in the hippocampal formation of the rat brain

Here, we draw on previously published estimates of anatomical statistics in regions of the rat hippocampal formation (Table 6.1), a well-studied brain circuit with implications in critical functions, such as spatial processing and memory processes (Grieves & Jeffery, 2017; O’Keefe & Dostrovsky, 1971; O’Mara & Aggleton, 2019; Sutherland & Rudy, 1989). We focus on areas CA3, CA1 and subiculum, as these, respectively, receive large, well-defined, anatomical projections from the preceding area (Amaral et al., 1990). Subicular projections from CA1 are our primary focus, as CA1 neurons display a relatively homogenous place code (Colgin, 2020), whereas subicular neurons have a very heterogenous spatial code (Brotons-Mas et al., 2017; Frost et al., 2021). We are uncertain about certain key numbers (e.g., distribution of projections, number of senders, proportion of pyramidal cells), as the relevant data do not yet exist in the literature. As such, we first present the case of CA3 to CA1 informed by neuroanatomical studies, and then perform a parameter exploration for proximal CA1 (near CA3) connections to

the distal subiculum (the predominant termination site for proximal CA1).

We consider both tetrode recordings and Neuropixel probe recordings in proximal CA3 and distal CA1 (as the preferential direction of the projections relates to the transverse level of origin, Witter, 2007). We assume that one-third of CA3 pyramidal cells lie in proximal CA3 and one-third of CA1 pyramidal cells lie in distal CA1. A tetrode is assumed to sample from a uniform 130 μm radius sphere of volume 0.0092 mm^3 (Mechler et al., 2011), with an expected yield of 5 neurons (Gray et al., 1995). A Neuropixels probe is assumed to sample from a cylinder with height 3.8mm (the full probe is shank is 10mm) and radius 0.1mm (Buccino et al., 2018; Buccino et al., 2019; Steinmetz et al., 2018) of volume 0.1193 mm^3 , with an expected yield of 79 neurons (Jun et al., 2017). With a tetrode (5 neurons) in both rat CA3 and CA1, there is only a 35% chance of simultaneously recording any directly connected neurons, while for a probe (79 neurons), we would expect 74% (2.5th percentile = 57%, 97.5th percentile = 89%) of neurons recorded in CA3 to receive a direct connection from at least one neuron being recorded in CA1, assuming that the recording devices are appropriately targeted towards proximal CA3 and distal CA1 (Figure 6.4A, B).

Consider recording pyramidal cells in the rat proximal CA1 (near CA3), which project in a columnar fashion to distal subiculum (Amaral et al., 1991). We assume that one-third of subicular pyramidal cells lie in distal subiculum and one-third of CA1 pyramidal cells lie in proximal CA1. Further, we assume there is a spread of projections around the mean proportion of projections, assuming a uniform distribution. We perform a parameter exploration that represents the unknown underlying anatomy and demonstrate how the underlying connectivity affects the probability of recording connected neurons in three different cases:

1. 36% of proximal CA1 pyramidal cells project to 5% of distal subicular pyramidal cells.
2. 10% of proximal CA1 pyramidal cells project to 5% of distal subicular pyramidal cells.
3. 16% of proximal CA1 pyramidal cells project to 1.1% of distal subicular pyramidal cells.

For instance, in a dual Neuropixel probe recording capturing 79 neurons from each brain area, if 36% of proximal CA1 pyramidal cells project to 5% of the distal subicular pyramidal cell population (case 1), c. 60 recorded subicular neurons are expected to receive a connection from at least one of the recorded CA1 neurons; but if instead, 10% project to 5% (case 2), then c. 25 recorded subicular neurons are expected to receive a

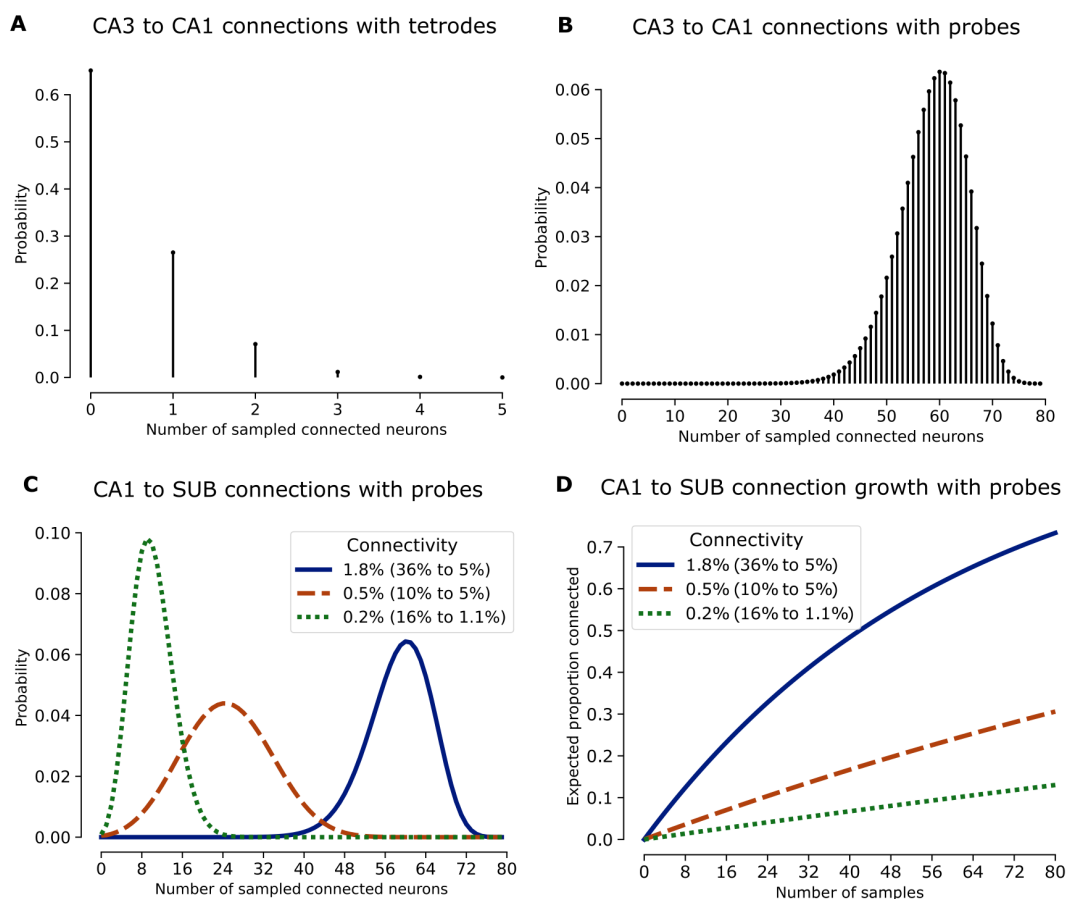


FIGURE 6.4: Modelling the probability of simultaneously recording directly connected neurons in rat CA3, CA1, and the subiculum.

(A) With a tetrode (5 neurons) in both rat CA3 and CA1, there is only a 35% chance of simultaneously recording any directly connected neurons.

(B) Recording directly connected neurons in rat CA3 and CA1 with a Neuropixels probe (79 neurons) is more likely, with a 95% chance of between 57% and 89% of neurons recorded in CA3 receiving a direct connection from at least one neuron in recorded in CA1.

(C) With 79 neurons simultaneously recorded from neural probes in proximal CA1 and distal subiculum, the probability of recording connected neurons is highly dependent on the underlying connectivity.

(D) Similarly to (C), the growth rate of connections is highly dependent on the underlying connectivity between rat proximal CA1 and distal subiculum.

connection (Figure 6.4C). Consider instead simultaneously recording 25 neurons from rat proximal CA1 and distal subiculum in a dual Neuropixel probe recording; in case 1 (36% to 5%), then 8 subicular neurons that receive connections are expected, but in case 3 (16% to 1.1%), this falls to approximately a single subicular neuron that receives connections (Figure 6.4D). Further, we estimate that roughly 40 simultaneously recorded neurons are required from rat proximal CA1 and distal subiculum to expect close to 50% of the recorded population in the subiculum to directly receive a connection from at least one of the recorded CA1 neurons in case 1, which is most akin to CA3 to CA1 connectivity.

6.4 Discussion and limitations

6.4.1 Implications for experimental design of cross-region analysis

An important and unresolved issue in neuroscience is understanding how different brain regions interact and they process inputs from other regions, termed “source transformation rules”. Intracranial recordings of neuronal activity have provided valuable insights into coding principles within specific brain regions, such as the hippocampal formation, where place cells and other spatially responsive cells have been identified (Moser et al., 2008; O’Keefe & Dostrovsky, 1971; Taube et al., 1990). At the limit, recording many connected neurons in multiple brain regions may prove to be a key component to inferring neural source transformation rules.

In this study, we aim to shed light on this issue by estimating the number of connected neurons that can be recorded simultaneously from anatomically connected brain regions. Previous multi-region recordings have often been limited to small numbers of electrodes or tetrodes, resulting in recordings from only a few tens of neurons at best. However, recent advances in recording technology, such as Neuropixels probes, offer the potential to simultaneously record from hundreds of neurons within and between different brain regions. This technological progress intuitively increases the likelihood of capturing neurons that interact directly across brain regions, thus making the detection and isolation of transformation rules between regions more feasible.

To quantify our intuition and assess the tractability of recording directly connected neurons *in vivo*, we develop a method based on directed acyclic graphs that represent the chemical synapses between neurons in the studied brain regions. Through Monte Carlo simulations and statistical analyses, we calculate the probability distribution of “recordable” neurons that are connected. We apply this method to neurons in the neocortex, using the Blue Brain model of the mouse neocortex connectome, as well as

to neurons in the hippocampus and subiculum, using data from anatomical studies.

In parts of the neocortex, in large, sparsely connected areas such as motor or somatosensory regions, capturing a high rate of directly connected neurons may require recording thousands of neurons, which is currently more feasible with optical methods like calcium imaging. Moreover, we observe a dramatic increase in the probability of recording connected neurons along chains of synapses, such as disynaptic or trisynaptic connections. Despite this, if effects are studied that may pass through intermediate neurons (e.g., through disynaptic or trisynaptic connections), then remarkably high rates of connected neurons are achieved with just tens of neurons. This finding shows the effect that perturbations to a small ensemble of neurons can have on local and distant structures. Such spread of neural connections likely echoes why techniques such as optogenetics can be so impactful, even if they act on a small subset of the full neuronal population (Carrillo-Reid et al., 2016, 2019; London et al., 2010).

As an example of a potential investigation of source transformation rules, consider CA1 and subiculum. Most cells in rodent hippocampal area CA1 are place cells and CA1 sends the vast proportion of its efferents to the adjacent subiculum according to a reliable and ordered set of anatomical rules (Amaral & Witter, 1989; Amaral et al., 1991; Cembrowski et al., 2018; O'Mara et al., 2001; Witter & Groenewegen, 1990). Yet, subicular neurons display very heterogenous responses (Brotons-Mas et al., 2017; Frost et al., 2021), from no spatial signal whatsoever, to place, grid, boundary-vector, and other cell classes. One approach to understanding the transformation rules the subiculum applies to its principal CA1 inputs is to record from CA1 and subiculum simultaneously. Our calculations suggest that there is a high probability of recording multiple, directly connected, neurons in these kinds of anatomically densely-connected regions with current high site count neural recording technologies.

6.4.2 Limitations of the current model

Despite the insight into connection probability in conjoint recordings, there are limitations to this work. For instance, the current model is focused on comparing two brain regions at a time. It can be adapted to more brain regions using the Monte Carlo simulation method, but the statistical estimation method may fail due to connection probability dependencies. For instance, if recording in the anterior thalamic nuclei, the subiculum, and the hippocampus, there are multiple pathways between these areas (Aggleton, 2012; Jay & Witter, 1991; Shibata, 1993). Thus, recording a set of anteromedial neurons that highly project to a set of temporal subiculum neurons may affect how likely a simultaneously recorded set of CA1 neurons are to send projections to that set

of temporal subiculum neurons. Further limitations are reflected in the assumptions made about the model (subsection 6.2.1). Here we will address these assumptions in turn, reflecting on how limiting the assumption is, how approximate the assumption is, and what impact deviation from the core assumption has on the results. A hard model limitation cannot be incorporated into the model without notable change. For a soft model limitation, the current model can be used to account for this information and overcome the assumption, but it may require multiple runs of the model, or setting up the model parameters in non-intuitive manners and would perhaps be desirable to be explicitly handled in future endeavours.

1. **Synaptic connection strength is not considered (hard model limitation):** Synapses do have different strengths, and the strength agnostic approach of this model impacts the interpretation of functional interaction between the neurons recorded.
2. **Multiple synaptic connections are treated as one connection (hard model limitation):** This is approximately true, as the direct influence of one neuron onto another is typically weak, being mostly represented by a single synapse (Braitenberg & Schüz, 1998). Furthermore, this assumption has no impact on which neurons receive connections, but does highly impact the interpretation of the strength of connections between neurons. Related to assumption 1, as incorporating synaptic connection strength is linked to the number of synaptic connections.
3. **Randomly sampling neurons from within the recording devices (hard model limitation):** The types of neurons that are favoured by the recording modality is not accounted for (e.g., electrophysiological recordings are biased towards large neurons with high firing rate Wei et al., 2020). This impacts the results as the true probability of connection may be different for the cells the recording devices are biased towards.
4. **Considering the brain at a fixed state (soft model limitation):** The brain is plastic and synaptogenesis occurs naturally throughout a lifespan, especially in early development. The current model can account for this but would need information on synaptic changes and be run at multiple time points.
5. **Only one connection type is considered (soft model limitation):** The model can account for different connection types, but to have a probability of inhibitory and excitatory connections, the model must be set up appropriately and run separately for inhibitory and excitatory connections.

6. **No electrical synapses (hard model limitation):** Electrical synapses exist, but neurons connected in this way lose much of their individuality and such connections are rare in the mammalian brain. The lack of incorporation of these into the model affects the interpretation of functional connectivity.
7. **The probability of spiking between connected neurons is not considered (hard model limitation):** This is currently out of the scope of this work. Significantly affects functional connectivity interpretation.
8. **Autapses are not considered (hard model limitation):** Autapses (synaptic connections from a neuron to itself) do happen, so it is an invalid assumption. However, the existence of such a connection does not determine if two neurons are chemically synaptically connected and does not influence the shortest path of synapses between two neurons. However, ignoring these connections may again impact the interpretation of the strength of connections between neurons.

Two immediate next steps to overcome these limitations would be to associate properties with each neuron such as size, genetic/morphological cell type, position, and firing properties. The second would be to represent connection strength and the number of synapses for each pair of connected neurons. The Monte Carlo method can handle these extensions quite naturally, but incorporating them into the statistical estimation method is challenging. As computers get more powerful, perhaps the need for the statistical estimation case subsides, but it is much easier to parameterise and execute than the Monte Carlo method and faster to run.

6.4.3 Bias of the model

The model incorporates spatial information by assigning a 3D position to each neuron and predicting the potential neurons in proximity of planned recording devices. However, two primary challenges arise in this approach. The first challenge is in obtaining topological information such as likely neuron positions throughout the brain and their corresponding projection locations. Secondly, the statistical estimation method necessitates a degree of averaging, introducing a challenge in capturing the variability of the initially complex topological information. Despite the inherent difficulty in resolving these issues, a question emerges: given the inherent spatial disorganisation of this model, could there be biases in the estimates, either overestimating or underestimating compared to real brain connectivity that tends to exhibit more structure? This question prompts a two-fold investigation.

From a theoretical standpoint, the absence of topological information, including

neuron positions and the distribution of projections from source to target regions, tends to homogenise the situation. However, the reality is more intricate, where connections from single neurons exhibit a skewed distribution with a lengthy tail, and projections seldom span the entire output region (Reimann et al., 2019). Consequently, we estimate that the model tends to overestimate connection numbers with reduced variance in the spatial agnostic case. This is because, in actuality, sampling from neurons that send few projections to a small area of the target region is more likely than predicted by the uniform spatial model.

Taking an experimental perspective, our theoretical view of overestimating with reduced variance finds support in two sources. Firstly, the initial iteration of this model lacked the capability to handle any spatial information (refer to our preprint, Martin et al., 2020). In the updated version, equipped with improved topological information handling, the estimated mean number of connections is consistently decreased with increased variance, compared to the preprint model. Furthermore, the statistical estimation method, which compresses topological information more than the Monte Carlo method, yields a higher estimate of the number of connected neurons (Figure 6.2).

6.4.4 The relationship between functional connectivity and structural connectivity

Approaching the problem of predicting functional connectivity is becoming more tractable as modern technologies get more powerful. Many studies will be required, combining a myriad of experiments, such as tracing, molecular analysis, two-photon optogenetics, electron-microscopy, connectomics, long-term potentiation/depression, genetic studies, RNA sequencing, and synapse-level recordings. For instance, to investigate functional connectivity in a micro-circuit, one could consider using two-photon optogenetics to combine electron-microscopy connectomics with optogenetics and two-photon calcium imaging in a single block of cortex (as in Yang et al., 2018). Predicting functional connectivity would likely require at least the following information:

- A definition of functional connectivity.
- The chemical synapses connecting the cells and the number of those synapses.
- Synaptic strength, i.e., the amplitude of excitatory/inhibitory postsynaptic potentials.
- Gap junctions or electrical synapses.
- Non-structural connectivity, such as ephaptic coupling.

- Local field potential dynamics.
- Synaptic plasticity.
- Genetic and morphological cell types.
- Placement of cells in the cell atlas.
- The recording modalities of choice and the bias of those modalities.

The functional connectivity as estimated from functional magnetic resonance imaging (fMRI) has been related to activity captured from wide-field calcium imaging (Shah-savarani et al., 2023), but the relationship between functional connectivity and the underlying synaptic connectivity is still unclear. In future work, it would be interesting to relate these measures of structural connectivity to the functional connectivity as estimated from fMRI. A possibility is to combine the anatomical connectivity with geometry, which influences functional connectivity like a neural field (Pang et al., 2023).

6.4.5 Conclusion

Overall, we find the probability of simultaneously recording directly connected pairs of neurons *in vivo* in anatomically connected regions is high for regions that are well connected if high site count recording techniques are used, and the recording devices are well placed. This suggests we may be able to infer source transformation rules from multiple, paired, inter-regional recordings *in vivo*, in combination with techniques summing activity across many neurons. The findings presented here provide a foundation for considering the probability of simultaneously recording connected neurons in multiple, anatomically connected brain regions, and considering pathways to understanding neuronal source transformation rules. Furthermore, they highlight the need for quantitative digital neuroanatomical techniques and analytic techniques to further constrain functional theories of neural coding.

6.5 Methods and examples

6.5.1 Description of directed graphs and networks of neurons

A directed graph $G = (V, E)$ represents a set whose elements are related via a series of directed edges, where:

- V is a set, whose elements are called nodes.

- E is a set of ordered pairs, whose elements are called arrows, such that $(a, b) \in E$ represents a connection going from $a \in V$ to $b \in V$ in the graph.

When we consider neurons and synaptic connections, V represents the set of neurons, while E represents the set of directed chemical synaptic connections between these neurons. A node $b \in V$ is reachable from $a \in V$ if $(a, b) \in E$ or there is a path (sequence of edges):

$$\{(a, v_i), (v_i, v_{i+1}), \dots, (v_{n-1}, v_n), (v_n, b)\} \subset E$$

Such a path has length 1 if $(a, b) \in E$ or $n + 1$ otherwise. In terms of neurons, this would be equivalent to the existence of a set of chemical synapses starting from a and linking to b . The geodesic distance between any two nodes in the graph is the length of the shortest path between these nodes. As such, a geodesic distance of 2 between two ordered neurons $\{a, b\} \subset V$ implies that there is no direct chemical synapse from a to b ($(a, b) \notin E$), but that at least one disynaptic connection exists between the two neurons ($\exists v_1 \in V$, such that $\{(a, v_1), (v_1, b)\} \subset E$).

Considering neuronal connections as a directed graph

Consider recording k and m neurons respectively from two brain regions or areas A and B , each containing N and M neurons overall. The regions A, B are two disjoint sets of vertices, representing neurons, in a directed graph $G = (V, E)$, describing the physical connectivity of the neurons through directed chemical synapses. E , the set of arrows connecting the vertices, indicates a direct connection between two neurons in V (Reimann et al., 2017), such that $v_i, v_j \in E$ indicates a synaptic connection between v_i and v_j where v_i is the presynaptic neuron, and v_j is the postsynaptic neuron. Our assumptions (subsection 6.2.1) impose that such a graph is simple directed, as:

- The graph is directed (6 – only chemical synapses)
- The graph has no loops (8 – a neuron is not connected to itself)
- At most, one arrow exists in the graph with the same source and target nodes (2 – multiple synaptic connections between two neurons are bundled into one connection).
- The graph is static (4 – nodes and edges are fixed)
- All nodes and edges in the graph are of a single type (1, 3, 5 – no synaptic weights, neurons sampled with equal weight, no inhibitory/excitatory split).

Our primary question can be restated: *what is the probability that a random sample of vertices $V_S \cup V_E$ from disjoint regions ($V_S \subset A \subset V$, $V_E \subset B \subset V$, $A \cap B = \emptyset$) contains directed paths in which x vertices in V_E are reachable from V_S .*

6.5.2 Calculating connection probability via Monte Carlo simulation

Firstly, a directed graph is generated representing the direct synaptic connections between all neurons in two brain regions A, B following the assumptions (subsection 6.2.1) and a model M to build the connections (see section 6.2 for a summary of the parameters). Then, the problem is modelled in 3D space using Ω , the brain atlas, Γ , the location of cells in Ω , and A, B , the two brain regions in which to place cells with 3D positions. Then the recording devices R_1, R_2, \dots, R_n are simulated as being placed into the brain, and the cells (nodes) in proximity of these devices are noted. Note that if any of M , Ω , or Γ are stochastic models, then the directed graph and cell positioning should be regenerated on a regular basis. For these directed graphs, random sets of source and target nodes are sampled from those nodes in proximity of the recording devices and the target nodes reachable from the source nodes determined. This operation is repeated many times, and the number of reachable target nodes recorded. Such a simulation generates draws from the probability distribution of the number of recorded downstream neurons that receive connections from recorded upstream neurons. Thus, these simulations can be used to estimate a probability distribution in Monte Carlo fashion.

Additional calculations were performed to estimate the convergence rate of the estimated distribution from these simulations to the true underlying distribution. This was evaluated by checking the smoothness and variance of the simulation result with varying numbers of iterations. This was performed on network graphs, as well as on a simulated problem, whereby there are two bags of balls being drawn from. The first bag contains red and black balls, and the number of red balls drawn from the first bag determines the distribution of red and black balls in the second bag. The result is the number of red balls drawn from the second bag. The statistics of this problem mimic the way the problem is to be solved, but without the complications of directed graphs. In all cases, at least 50,000 samples are needed for a steady level of convergence to the full probability distribution in Monte Carlo simulations. If the objective is to obtain summary statistics, such as the mean and variance of the distribution, convergence is faster and can be achieved with a few thousand simulations.

Comparing the statistical estimation to the Monte Carlo simulation

The Monte Carlo simulation method is slow and computationally expensive, but can be considered a form of ground truth for this problem. As such, it was used to evaluate the statistical calculations by comparing the distributions from the estimate to those obtained from Monte Carlo simulations via the Bhattacharyya distance.

The Bhattacharyya distance is defined as follows: consider probability distributions p and q over the same domain X , and let $BC(p, q)$ denote the Bhattacharyya coefficient between p and q , and let $D_B(p, q)$ denote the Bhattacharyya distance between p and q . Then:

$$\begin{aligned} BC(p, q) &= \sum_{x \in X} \sqrt{p(x) \cdot q(x)} \\ D_B(p, q) &= -\ln(BC(p, q)) \end{aligned} \tag{6.1}$$

The lower the Bhattacharyya distance the better (minimum of 0), indicating greater similarity of the probability distributions being evaluated. For instance, the mean Bhattacharyya distance between the distributions for which the expected value is shown in Figure 6.3C was 0.024 with a maximum of 0.088 for the prelimbic area to the infralimbic area.

Representing the directed graphs computationally

A directed graph is represented by a sparse matrix indicating the presence of edges in the graph. As such, directed graphs are represented as lists of target nodes, similar to how a sparse matrix can be efficiently stored as a list of lists. For instance, the representation $[2, 3], [], [0], [1, 2]$ describes a graph with four vertices 0, 1, 2, 3 and five edges $0 \rightarrow 2$, $0 \rightarrow 3$, $2 \rightarrow 0$, $3 \rightarrow 1$, $3 \rightarrow 2$. A connection matrix representation of a directed graph can be converted into this format by taking non-zero entries as (r, c) and inserting the value c into the list r .

Directed graphs were randomly generated in this format in Python, while adhering to configured forms of connectivity. To generate a directed graph representing connections between two brain regions, the number of neurons in each region is provided, as well as distributions of forward, recurrent, and local connections in each region. Connections were then randomly generated to adhere to these provided distributions. Connections were sampled with replacement, representing the multiple synapses that could be formed between the same pairs of neurons (which are then bundled into one single connection as per the model assumptions).

Given a list of source and target nodes in a directed graph, the number of target nodes reachable from the source nodes was determined by running an iterative deepening, depth-limited, depth-first search on the transpose graph (in which the direction of arrows are reversed) starting from each target node, to see if any source node was reachable from the target node. This is equivalent to running a search operation on the original graph, and starting from source nodes to see which target nodes are reachable, but speeds up computation by finding predecessors of target nodes as opposed to the successors of source nodes. The result is a list of target nodes reachable from one or more source nodes. An iterative deepening operation is performed to allow one to consider not only direct connections between source and target nodes, but also connections that require pass-through nodes. With neurons, this allows one to consider indirect connections such as disynaptic or trisynaptic connections.

Directed graph visualisation

For small directed graphs, visualisations were performed using Networkx (Hagberg et al., 2008). Nodes in the graph are positioned using the Fruchterman-Reingold force-directed algorithm, whereby the edges in the graph are considered as springs holding the nodes close, and the nodes apply a repelling force. The system is then simulated until equilibrium is reached, and the graph is plotted at this point. As such, well-connected groups of neurons would be close to each other in the visualisation.

For large directed graphs, visualisations were performed by considering the graph as a connection matrix. In this connection matrix C , row r and column c having the value 1 indicates the existence of a directed connection between node r and c , while the value 0 indicates no connection. The connection matrix is considered in $n \times n$ sized windows (where n is an arbitrarily chosen window size for smoothing), and the percentage of ones in these squares is calculated. The result is then visualisation using a colourmap, from the minimum of 0% representing no connections between the n^2 pairs of neurons, to the maximum representing the most highly connected block of neurons in the network.

6.5.3 Calculating connection probability via statistical methods

Let X be the random variable whose value is the number of neurons recorded in B which receive at least one connection from a neuron recorded in A . Furthermore, let X_A be the random variable whose value is the number of neurons recorded in A sending a connection to B , and X_B be the random variable whose value is the number of neurons in B receiving a connection from a neuron recorded in A . By the chain rule of probability, and taking the marginal distribution over X_A and X_B , the probability of

recording x such neurons in B is then

$$\begin{aligned}
 P(X = x) &= \sum_{k=0}^n \sum_{l=0}^M P(X_A = k, X_B = l, X = x) \\
 &= \sum_{k=0}^n \sum_{l=0}^M P(X_A = k) \cdot P(X_B = l \mid X_A = k) \cdot P(X = x \mid (X_B = l), (X_A = k)) \\
 &= \sum_{k=0}^n \sum_{l=0}^M \left(\frac{\binom{K}{k} \cdot \binom{N-K}{n-k}}{\binom{N}{n}} \right) \cdot P(X_B = l \mid X_A = k) \cdot \left(\frac{\binom{l}{x} \cdot \binom{M-l}{m-x}}{\binom{M}{m}} \right)
 \end{aligned} \tag{6.2}$$

where:

N = number of neurons in A

n = number of neurons recorded in A

K = number of neurons in A which send connections to B

k = number of recorded neurons in A which send connections to B

M = number of neurons in B

m = number of neurons recorded in B

l = number of neurons in B which receive connections from neurons recorded in A

In the above formula, two things are of note. Firstly

$$P(X = x \mid (X_B = l), (X_A = k)) = P(X = x \mid (X_B = l)) \tag{6.3}$$

since knowing how many sampled neurons in A send connections to B is redundant to calculating $P(X = x)$ if the number of neurons in B which receive connections from the neurons sampled in A is known. Secondly, $P(X_A = k)$ and $P(X = x \mid (X_B = l))$ are both calculated from the hypergeometric distribution, which describes the probability of recording distinct neurons sending or receiving connections from a population of neurons.

The hypergeometric distribution

The hypergeometric distribution describes the probability of obtaining k successes in n draws without replacement from a population of size N with K objects that would be considered a success. Let X represent the random variable that takes the value of the number of successes drawn from the population, then

$$P(X = k) = \frac{\binom{K}{k} \cdot \binom{N-K}{n-k}}{\binom{N}{n}} \tag{6.4}$$

If N is much larger than n , the binomial distribution can be used instead as a faster approximation of the hypergeometric distribution. In this case, there would be n draws from a population of N with replacement with $p = K/N$ as the probability of drawing a success. Then

$$P(X = k) = \binom{n}{k} p^k \cdot (1 - p)^{n-k} \quad (6.5)$$

Estimating the proportion of neurons in the target region receiving a connection

We describe calculating $P(X_B = l \mid X_A = k)$, the probability of l neurons in B receiving at least one connection from a neuron recorded in A , given that k neurons are recorded in A which send direct connections to B .

First, consider computing $P(X_B = l \mid X_A = k)$ for direct connections between neurons in A and B . Let the random variables R_1, R_2, \dots, R_k , represent the number of synapses that each sampled neuron in A that projects to B sends to B . We assume R_1, R_2, \dots, R_k are independent and identically distributed with mean μ and variance σ^2 (the distribution is not assumed, it could be uniform, exponential, gaussian, etc. but it is a fixed distribution). Then $S_k = R_1 + R_2 + \dots + R_k$ represents the number of outgoing synaptic connections from A to B . By the central limit theorem $S_k = R_1 + R_2 + \dots + R_k \approx N(k\mu, k\sigma^2)$ for high values of k (usually $k > 10$ is enough for a good approximation, but $k > 30$ is recommended). For k below these values, a recursive convolution operation is applied to the distributions of the random variables (Grinstead & Snell, 2012). If $m_1(x)$ and $m_2(x)$ are the distribution functions of X and Y , then the convolution of $m_1(x)$ and $m_2(x)$ is the distribution function $m_3 = m_1 * m_2$ given by

$$m_3(j) = \sum_i m_1(i) \cdot m_2(j - i)$$

With S_k computed, we then apply a function f to S_k to produce a random variable $Y = f(S_k)$, where this function $f(s)$ is defined to compute the expected number of unique neurons in B with at least one incoming synapse from a neuron in A with s synapses going from A to B . Thus, given M neurons in B :

$$f(s) = M \left(1 - \left(\frac{M-1}{M} \right)^s \right)$$

Finally, $P(Y = l) = P(X_B = l \mid X_A = k)$, the original distribution we aimed to compute.

The central limit theorem will introduce small errors in the case of large graphs with

skewed node-link distributions. This is because, unlike dense graphs, in which large enough subgraphs closely follow the properties of the full graph, sparse graphs with skewed distributions have no such guarantee (Lovász, 2012). As such, the statistical estimation method is prone to a slight overestimation of the expected number of recorded neurons that are connected when the forward chemical connection distribution is right-skewed (see Figure 6.2).

Simplifying the full connection distribution with a mean

Calculating $P(X = x)$ using the full distributions of $P(X_B = l \mid X_A = k)$ for each k can be too time-consuming, so instead, we can consider the random variable X_B to be distributed as a single point at the mean of the true distribution with 0 variance, whereby $P(X_B = l \mid X_A = k) = 1$ if $l = \mu$, and 0 otherwise. This reduces the accuracy, but increases the speed of calculating $P(X = x)$. This approach is often necessary for large graphs ($> 100,000$ nodes) with a maximum geodesic distance greater than 1.

Estimating the proportion of neurons in the target region receiving a connection for different geodesic distances

In general, to compute $P(X_B = l \mid X_A = k)$, we limit the maximum geodesic distance considered between nodes A and B in the graph, or the number of synapses connecting the neurons in A and B , to a value D . Then, using the connectivity statistics between the two regions, the interactions of outgoing (AB), recurrent (BA), and local (AA , BB) connections are described, and used as building blocks, since these will comprise any paths between A and B . For a maximum geodesic distance of D , the sum over $2^D - 2$ variables must be taken to compute the marginal distribution $P(X_B = l \mid X_A = k)$.

For the max geodesic distance $D = 1$ (direct connections only), let X_{AB} be the random variable denoting the number of neurons recorded in B directly connected to a neuron recorded in A , then:

$$P(X_B = l \mid X_A = k) = P(X_{AB} = l \mid X_A = k). \quad (6.6)$$

While for the max geodesic distance $D = 2$ (disynaptic connections and direct connections are permitted), also let X_{AAB} denote the number of neurons in B not receiving a direct connection from a neuron recorded in A , but receiving a connection from a neuron recorded in A via an intermediate neuron, where the intermediate neuron remains in A . Similarly, X_{ABB} denotes the number of neurons in B not receiving a direct connection from a neuron recorded in A , but receiving a connection from a neuron recorded in B via an intermediate neuron, where the intermediate neuron lies in B . Then, the sum is

taken over two variables

$$\begin{aligned}
 P(X_B = l \mid X_A = k) &= \sum_{j,h} P(X_{AB} = j \mid X_A = k) \cdot \\
 &\quad P(X_{AAB} = h \mid X_A = k, X_{AB} = j) \cdot \\
 &\quad P(X_{ABB} = l - j - h \mid X_A = k, X_{AB} = j, X_{AAB} = h)
 \end{aligned} \tag{6.7}$$

Connectivity statistics with respect to neural recording devices

When considering neural recording device(s) placed within the brain regions A and B , the following parameters are required to describe direct connections:

1. N_{AR} , N_{BR} – the number of neurons that could potentially be sampled by the recording device (which will be referred to as; *in the proximity of the recording device*). For instance, a tetrode could only record from a small set of neurons that lie in a sphere with volume 0.0092mm^3 (Mechler et al., 2011). So, in CA1 which has a volume of 18.8mm^3 (Ropireddy et al., 2012) approximately 0.05% of the cell population would be in the proximity of a tetrode.
2. N_{ARBR} – the number of neurons in proximity of the recording device in A which send connections to neurons in proximity of the recording device in B .
3. P_{ARBR} – the probability distribution of efferent connections from neurons in proximity of the recording device in A to neurons in proximity of the recording device in B .

While the following additional parameters are required to describe indirect connections:

1. N_A , N_B – the number of neurons in regions A and B .
2. P_{ARB} – the probability distribution of efferent connections from neurons in proximity of the recording device in A to neurons outside the proximity of the recording device in B .
3. P_{ABR} – the probability distribution of efferent connections from neurons outside the proximity of the recording device in A to neurons in proximity of the recording device in B .

4. P_{AB} – the probability distribution of efferent connections from neurons outside the proximity of the recording device in A to neurons outside the proximity of the recording device in B .
5. P_{BA} – the probability distribution of afferent connections to neurons in A from neurons in B .
6. P_{AA} – the probability distribution of local connections in A .
7. P_{ARA} – the probability distribution of local connections from neurons in proximity of the recording device in A to the rest of the neurons in A .
8. P_{BB} – the probability distribution of local connections in B .
9. P_{BBR} – the probability distribution of local connections from neurons in B to neurons in proximity of the recording device in B .
10. N_{AB} – the number of neurons in A which send connections to neurons in B .
11. N_{ABR} – the number of neurons outside the proximity of the recording device in A which send connections to neurons in proximity of the recording device in B .
12. N_{ARB} – the number of neurons in proximity of the recording device in A which send connections to B .
13. N_{BA} – the number of neurons in B which send connections to neurons in A .

Though many parameters can be set for indirect connections, they are primarily used to describe topography. For instance, when recording from proximal CA1 and distal subiculum in the rat, P_{ARB} can be set to describe how these two parts of CA1 and subiculum are well connected, while P_{AB} can be set to describe how connections outside of proximal CA1 and outside of distal subiculum are formed.

6.5.4 Technical details of statistical estimation and performance

Accuracy of mean estimation versus full distribution estimation

Using the mean of $P(X_B = l | X_A = k)$ is not as accurate as using the full distribution, but may be required in certain cases (e.g., large graphs and indirect connections) for computational performance reasons. Furthermore, in general, the statistical estimation is not as accurate at larger maximum geodesic distances, i.e., indirect connections (see Figure 6.5).

Probability distributions at increased max geodesic distance

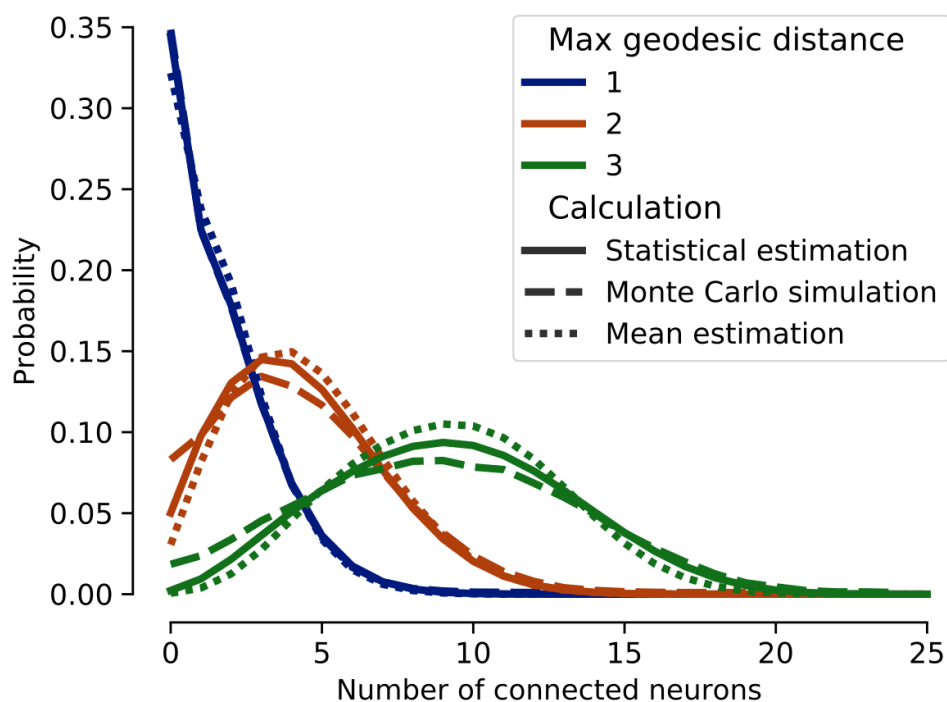


FIGURE 6.5: The relationship between the accuracy of probability estimation and maximum allowed synaptic jumps (geodesic distance).

As the maximum geodesic distance is increased the accuracy in estimating the probability mass function reduces. However, the expectation of the distribution is a better match. These expectations are: max distance 1: 1.59, 1.59, 1.59, max distance 2: 4.29, 4.33, 4.55, max distance 3: 8.86, 9.27, 9.22, in the order; Monte Carlo simulation, statistical estimation; mean estimation.

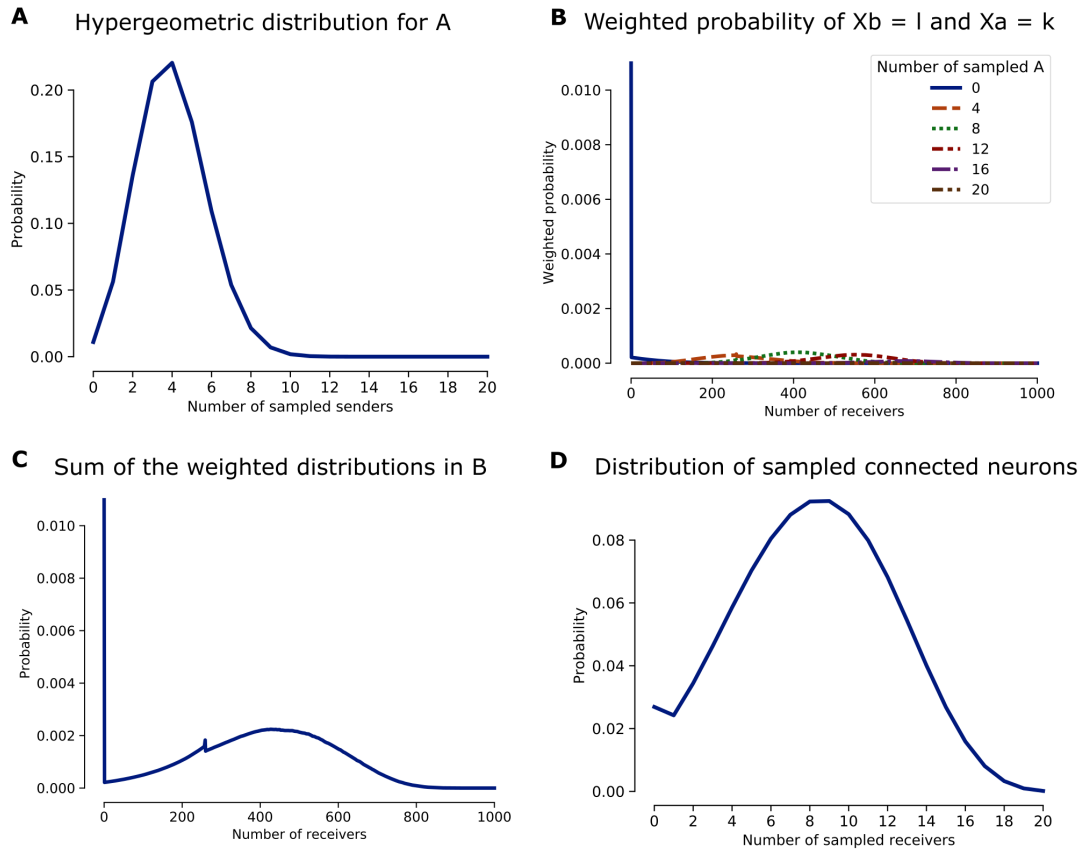


FIGURE 6.6: Multiple distributions involved in calculating $P(X = x)$ for Figure 6.1F with 20 samples in each region.

(A) Plotting k against $P(X_A = k)$, which is hypergeometrically distributed.

(B) Plotting l against $P(X_A = k) \cdot P(X_B = l | X_A = k)$ for different values of k .

(C) Summing the functions in B gives the marginal distribution of $P(X_B = l)$. Note that the blip on the graph around 260 receivers is due to a spike in $P(X_A = 2) \cdot P(X_B = l | X_A = 2)$ around $l = 260$, followed by a sharp decrease.

(D) The final result of x against $P(X = x)$ by using the distribution in C to calculate $P(X = x | X_B = l) \cdot P(X_B = l)$ and then summing over l to obtain the final distribution.

Interpolation accuracy improvement

Figure 6.6 shows the distributions involved in the statistical estimation. Since the numbers involved in calculating $P(X = x)$ with the statistical estimation can be large, the distributions are sometimes taken at sample points, and linearly interpolated in between these points to avoid calculating the full distribution. This allows the method to be run on larger graphs, without having to use the more inaccurate mean estimation method. To help ensure that valuable information is not missed, the second derivative is taken in between the points, and a substantial change in this value suggests missing information between sample points.

6.5.5 Example of a full probability distribution calculation

Consider the situation described in Figure 6.1. Here, there are 7 neurons in each region; there are 4 neurons which send exactly one direct connection from the first region, and we randomly sample 3 neurons in each region. Let X denote the random variable whose value is the number of sampled neurons in the second region receiving a direct connection from a neuron sampled in the first region. A full calculation of $P(X = x)$ is

below for this case:

Let

$$\delta(k) = P(X_A = k) = \left(\frac{\binom{4}{k} \cdot \binom{3}{3-k}}{\binom{7}{3}} \right)$$

$$\gamma(x, l) = P(X = x \mid X_B = l) = \left(\frac{\binom{l}{x} \cdot \binom{7-l}{3-x}}{\binom{7}{3}} \right)$$

Then

$$P(X = x) = \sum_{k=0}^3 \sum_{l=0}^7 P(X_A = k, X_B = l, X = x)$$

$$P(X = x) = \sum_{k=0}^3 \sum_{l=0}^7 P(X_A = k) \cdot P(X_B = l \mid X_A = k) \cdot P(X = x \mid X_B = l) \quad (6.8)$$

$$P(X = x) = \sum_{k=0}^3 \sum_{l=0}^7 \delta(k) \cdot P(X_B = l \mid X_A = k) \cdot \gamma(x, l)$$

With

$$P(X_B = l \mid X_A = k) = 1 \text{ if } l = k, \text{ else } 0$$

Finally

$$P(X = 0) = \delta(0) \cdot \gamma(0, 0) + \delta(1) \cdot \gamma(0, 1) + \delta(2) \cdot \gamma(0, 2) + \delta(3) \cdot \gamma(0, 3) = 0.384$$

$$P(X = 1) = \delta(1) \cdot \gamma(1, 1) + \delta(2) \cdot \gamma(1, 2) + \delta(3) \cdot \gamma(1, 3) = 0.500$$

$$P(X = 2) = \delta(2) \cdot \gamma(2, 2) + \delta(3) \cdot \gamma(2, 3) = 0.113$$

$$P(X = 3) = \delta(3) \cdot \gamma(3, 3) = 0.003$$

6.5.6 Calculations in the hippocampal circuit case study

Here, we demonstrate a set of model parameters for direct connections from CA3 to CA1, (see Table 6.1 for reference, note: the 36% of CA3 pyramidal cells projecting to 5% of the CA1 pyramidal cell population are chosen such that overall, a single CA1 pyramidal cell might receive input from about 1.8% of the CA3 pyramidal cell population):

1. $N_{CA3} = 303930 \cdot \frac{1}{3} = 101310$. $N_{CA1} = 347100 \cdot \frac{1}{3} = 115700$ (total number of pyramidal cells in proximal CA3 and distal CA1, assuming these areas both contain one-third of the total pyramidal cell population in CA3 and CA1).
2. $N_{CA3R} = N_{CA3} \cdot \frac{V_{\text{probe}}}{V_{CA3}} = 101310 \cdot \left(\frac{0.1193}{4.066} \right) = 2972$. $N_{CA1R} = N_{CA1} \cdot \frac{V_{\text{probe}}}{V_{CA1}} = 115700 \cdot \left(\frac{0.1193}{6.266} \right) = 2202$ (number of pyramidal cells that are in proximity of

the recording devices, assuming uniformly distributed cells and probe volume 0.1193mm^3).

3. $N_{CA3_R CA1_R} = 0.36 \cdot 2972 = 1070$ (36% of pyramidal cells in proximity of the recording device send projections to CA1 pyramidal cells in proximity of the recording device).
4. $P_{CA3_R CA1_R} = \text{unif}\{2, 218\}$ (the probability distribution of connections that CA3 cells in proximity of the probe send to CA1 cells in proximity of the probe. Here, we assume a uniform distribution with such that a single cell in CA3 projects to a mean of 5% of the 2202 CA1 pyramidal cells in proximity of the probe).

Calculations of model parameters hippocampal formation case study

The calculations involved in determining the parameters of the statistical model in the hippocampal formation case study are here. For CA3 to CA1, we calculate, for a tetrode in each region:

$$\begin{aligned}
 N &= \frac{0.0092}{12.2} \cdot 303,930 = 229 \text{ (CA3 pyramidal cells in sphere around tetrode)} \\
 n &= 5 \text{ (assumed tetrode yield)} \\
 K &= 229 \text{ (no knowledge of sender proportion)} \\
 M &= \frac{0.0092}{18.8} \cdot 0.89 \cdot 390,000 = 170 \text{ (CA1 pyramidal cells in sphere around tetrode)} \\
 m &= 5 \text{ (assumed tetrode yield)} \\
 E(X_B | X_A = 1) &= \frac{0.0092}{18.8} \cdot 0.89 \cdot 390,000 \cdot 0.018 = 3 \text{ (average CA3 projections)}
 \end{aligned} \tag{6.9}$$

While for a Neuropixels probe in each region:

$$\begin{aligned}
 N &= 303,930 \text{ (CA3 pyramidal cells)} \\
 n &= 79 \text{ (assumed Neuropixels probe yield)} \\
 K &= 303,930 \text{ (no knowledge of sender proportion)} \\
 M &= 390,000 \cdot 0.89 = 347,100 \text{ (CA1 pyramidal cells)} \\
 m &= 79 \text{ (assumed Neuropixels probe yield)} \\
 E(X_B | X_A = 1) &= 390,000 \cdot 0.89 \cdot 0.018 = 6248
 \end{aligned} \tag{6.10}$$

For proximal CA1 to distal subiculum, we calculate, for the high projection rate:

$$\begin{aligned}
 N &= \frac{1}{3} \cdot 0.89 \cdot 390,000 = 115,700 \text{ (pyramidal cells in proximal CA1)} \\
 K &= 0.9 \cdot 115,700 = 104,130 \text{ (assumed projection rate from proximal CA1)} \\
 M &= \frac{1}{3} \cdot 0.9 \cdot 290,000 = 87,000 \text{ (pyramidal cells in distal subiculum)} \\
 E(X_B | X_A = 1) &= 0.02 \cdot 87,000 = 1740
 \end{aligned} \tag{6.11}$$

While for the low projection rate:

$$\begin{aligned}
 N &= \frac{1}{3} \cdot 0.89 \cdot 390,000 = 115,700 \text{ (pyramidal cells in proximal CA1)} \\
 K &= 0.7 \cdot 115,700 = 80,990 \text{ (assumed projection rate from proximal CA1)} \\
 M &= \frac{1}{3} \cdot 0.8 \cdot 290,000 = 77,332 \text{ (pyramidal cells in distal subiculum)} \\
 E(X_B | X_A = 1) &= 0.01 \cdot 87,000 = 777
 \end{aligned} \tag{6.12}$$

6.5.7 Blue Brain Project neocortex connection model and visualisation

Connection matrices were downloaded from the Blue Brain repository (Reimann et al., 2019), version 1.15. Instance 4 ipsilateral and local connections in the right hemisphere of the brain were downloaded for

- The primary motor cortex (MOp).
- The primary somatosensory area associated with lower limbic function (SSp-ll).
- The primary visual area (VISp).
- The lateral visual area (VISl).
- The primary auditory area (AUDp).
- The posterior auditory area (AUDpo).
- The infralimbic area (ILA).
- The prelimbic area (PL).

For the left hemisphere, instance 4 ipsilateral and local connections were downloaded for the primary and lateral visual areas. The data were processed into smaller sparse matrices of connections and are stored on the Open Science Framework at <https://osf.io/u396f/>. These can be loaded using SciPy to handle the npz format files.

Cells were distributed uniformly throughout these brain regions according to the number of cells in the local connection matrix for that brain region. To maintain a structure of the connectivity matrices for these cells, the rows and columns of the connectivity matrices were sorted by the distance of the cells in 3D space to a Hilbert curve (a form of space-filling curve) such that rows that were close together in the connection matrix would also be spatially close to each other in 3D.

Mouse brain atlas and visualisation

The Allen mouse brain atlas (Lein et al., 2007) at 25 μm resolution was used in this project for probe visualisation and localisation. The atlas was loaded using the `bg-atlasapi` package (<https://github.com/brainglobe/bg-atlasapi>) and `Brainrender` (Claudi et al., 2021). For Neuropixel probe visualisation in primary visual area and lateral visual area, probe implant locations in the mouse brain were extracted from (Steinmetz et al., 2019) and registered in common coordinates in the Allen mouse brain atlas. For theoretical Neuropixel probe recordings with one probe implanted between the two areas (such as primary motor cortex to the primary somatosensory area associated with lower limbic function), the probes were angled such that they would pass close by the centre of both brain areas, while obeying a realistic implantation angle.

Neuropixel probes were considered to record from a cylinder with a recording span of height 3800 μm (the full length of the probe is 10,000 μm) and a 100 μm radius (Steinmetz et al., 2018). The 100 μm recording radius of a probe was chosen based on the estimated 130 μm recording radius of a tetrode (Mechler et al., 2011), the maximum distance of 80 μm used in (Buccino et al., 2018) to model extracellular neuronal recordings with neural probes and multielectrode arrays, and the distances considered between a neuron soma and a Neuropixel probe tip in (Buccino et al., 2019) to model the impact of neural probes on extracellular potentials.

Connectivity distribution calculation from connection matrices

Given a connection matrix C representing the connections between two brain regions, we calculate the distribution of forward, recurrent, and local connections. Firstly, it is assumed that brain region A has N neurons, B has M neurons, and the matrix C has size $(N + M) \times (N + M)$. Then, the top left $N \times N$ part of C represents the local connections in A , the top right $N \times M$ part represents the forward connections from A to B , the bottom left $M \times N$ part represents the recurrent connections from B to A , and the bottom right $M \times M$ part represents the local connections in B . The distribution of connections is then calculated by summing the number of ones appearing along rows and

then placing the result into a frequency table (excluding 0 for non-local connections). The frequencies divided by the number of (non-zero for non-local connections) rows indicate the distribution. This is performed for each of these four matrices.

To obtain the statistics in the presence of recording devices, the cells which are in the proximity of the recording devices are noted. Then the connection matrices can be subsampled to include only the row/column indices corresponding to these cells in proximity of the recording devices. For instance, subsampling the connection matrix to only have rows and columns corresponding to those in the proximity of the recording devices in both A and B would give the statistics of connectivity within the local area of all the recording devices. While subsampling to just those cells in the proximity of the recording devices in region A while leaving all cells in B would give connectivity statistics from the local area in the recording devices in A to the whole of B (or the area outside the recording devices in B if the matrix is subsampled to only include cells outside the proximity of the recording devices in B).

6.5.8 Power law analysis

In the power law style analysis, we scale the size of the brain areas which are connected along with the connection distribution between the brain areas. In this, $x\%$ of the neurons in region A project to neurons in region B , regardless of the size of A and B . Similarly, the distribution of efferent connections from A to B is scaled such that the mean of the distribution is fixed in terms of percentage, i.e., a single neuron in A that projects to B projects to $y\%$ of neurons in B on average, regardless of the size of A and B . In addition to a scaled version of the distribution, we can also investigate what number of neurons would need to be sampled if the distribution was fixed, but the number of projecting neurons scales with the size of the brain region.

With these parameters set, we can then set a desired expected rate of connections (e.g., we desire that 50% of the neurons recorded from B receive direct connections from the neurons recorded from A). Finally, we can search the solution space for the number of recorded neurons that would need to be recorded to achieve the desired rate of connections. There are many methods to perform this search (such as gradient descent), but a simple method is to set a minimum number of recorded neurons (1), and a maximum number of recorded neurons (the most the recording device could record, say 1000), and then find the outcome via half-interval search (iteratively halving the search space) over this space. This is an appropriate approach since the solution space is unimodal.

6.5.9 Code accessibility, sharing and reproducibility

The code described in the paper is written in Python 3.8 and freely available online at <https://github.com/seankmartin/NeuralConnections> and an archived version is available at <https://doi.org/10.5281/zenodo.6401535>. The connectivity matrices from the blue brain project used in this study have been archived at <https://osf.io/u396f/>. Instructions to reproduce this work in full are provided and full code documentation is available, along with a command-line and general-user interface to run further experiments.

Python packages used

Many open-source Python packages were vital to this work, and are listed here:

- Mpmath (Johansson & The mpmath development team, 2018) – large floating-point operations at arbitrary precision and handle cases of non-integer binomial co-efficient calculations, by using the gamma function.
- NumPy (Harris et al., 2020) – random number generation, storage of large arrays, vectorised operations, and gradient estimation.
- Networkx (Hagberg et al., 2008) – infer numerical properties of graphs, and for graph visualisation.
- SciPy (Virtanen et al., 2020) – sparse matrix operations and linear interpolation.
- Matplotlib (Hunter, 2007) – plotting colourmaps and graphs.
- Seaborn (Waskom & The seaborn development team, 2020) – further plotting operations.
- Pandas (McKinney, 2010) – storing and retrieving the results of Monte Carlo simulations.
- Brainrender (Claudi et al., 2021) – visualising recording device placement in anatomically registered data.

Computer specifications

A custom-built Windows 10 computer with 16GB DDR4 RAM, an Intel Core i5-6500 CPU @ 3.20GHz with 4 cores, and a NVIDIA GeForce RTX 2060 GPU was used to run the code and produce the results.

References

- Aggleton, J. P. (2012). Multiple anatomical systems embedded within the primate medial temporal lobe: Implications for hippocampal function. *Neuroscience & Biobehavioral Reviews*, *36*(7), 1579–1596.
- Allen, W. E., Chen, M. Z., Pichamoorthy, N., Tien, R. H., Pachitariu, M., Luo, L., & Deisseroth, K. (2019). Thirst regulates motivated behavior through modulation of brainwide neural population dynamics. *Science*, *364*(6437), eaav3932.
- Amaral, D. G., Ishizuka, N., & Claiborne, B. (1990). Neurons, numbers and the hippocampal network. *Progress in Brain Research*, *83*, 1–11.
- Amaral, D. G., & Witter, M. P. (1989). The three-dimensional organization of the hippocampal formation: A review of anatomical data. *Neuroscience*, *31*(3), 571–591.
- Amaral, D. G., Dolorfo, C., & Alvarez-Royo, P. (1991). Organization of CA1 projections to the subiculum: A PHA-l analysis in the rat. *Hippocampus*, *1*(4), 415–435.
- Anastassiou, C. A., & Koch, C. (2015). Ephaptic coupling to endogenous electric field activity: Why bother? *Current Opinion in Neurobiology*, *31*, 95–103.
- Andersen, P., Morris, R., Amaral, D., Bliss, T., & O'Keefe, J. (2006). *The hippocampus book*. Oxford university press.
- Arszovszki, A., Borhegyi, Z., & Klausberger, T. (2014). Three axonal projection routes of individual pyramidal cells in the ventral CA1 hippocampus. *Frontiers in Neuroanatomy*, *8*.
- Bezaire, M. J., & Soltesz, I. (2013). Quantitative assessment of CA1 local circuits: Knowledge base for interneuron-pyramidal cell connectivity. *Hippocampus*, *23*(9), 751–785.
- Braitenberg, V., & Schüz, A. (1998). *Cortex: Statistics and geometry of neuronal connectivity*. Springer Science & Business Media.
- Brotons-Mas, J. R., Schaffelhofer, S., Guger, C., O'Mara, S. M., & Sanchez-Vives, M. V. (2017). Heterogeneous spatial representation by different subpopulations of neurons in the subiculum. *Neuroscience*, *343*, 174–189.
- Buccino, A. P., Kordovan, M., Ness, T. V., Merkt, B., Häfliger, P. D., Fyhn, M., Cauwenberghs, G., Rotter, S., & Einevoll, G. T. (2018). Combining biophysical modeling and deep learning for multielectrode array neuron localization and classification. *Journal of Neurophysiology*, *120*(3), 1212–1232.
- Buccino, A. P., Kuchta, M., Jæger, K. H., Ness, T. V., Berthet, P., Mardal, K.-A., Cauwenberghs, G., & Tveito, A. (2019). How does the presence of neural probes affect extracellular potentials? *Journal of Neural Engineering*, *16*(2), 026030.
- Buzsáki, G. (2004). Large-scale recording of neuronal ensembles. *Nature Neuroscience*, *7*(5), 446–451.

- Buzsáki, G., & Tingley, D. (2018). Space and time: The hippocampus as a sequence generator. *Trends in Cognitive Sciences*, *22*(10), 853–869.
- Carrillo-Reid, L., Han, S., Yang, W., Akrouh, A., & Yuste, R. (2019). Controlling visually guided behavior by holographic recalling of cortical ensembles. *Cell*, *178*(2), 447–457.e5.
- Carrillo-Reid, L., Yang, W., Bando, Y., Peterka, D. S., & Yuste, R. (2016). Imprinting and recalling cortical ensembles. *Science*, *353*(6300), 691–694.
- Cembrowski, M. S., Wang, L., Lemire, A. L., Copeland, M., DiLisio, S. F., Clements, J., & Spruston, N. (2018). The subiculum is a patchwork of discrete subregions (L. Colgin & G. L. Westbrook, Eds.). *eLife*, *7*, e37701.
- Claudi, F., Tyson, A. L., Petrucco, L., Margrie, T. W., Portugues, R., & Branco, T. (2021). Visualizing anatomically registered data with brainrender (M. W. Mathis, K. M. Wassum, & J. Nunez-Iglesias, Eds.). *eLife*, *10*, e65751.
- Colgin, L. L. (2020). Five decades of hippocampal place cells and EEG rhythms in behaving rats. *Journal of Neuroscience*, *40*(1), 54–60.
- Erö, C., Gewaltig, M.-O., Keller, D., & Markram, H. (2018). A cell atlas for the mouse brain. *Frontiers in Neuroinformatics*, *12*, 84.
- Frost, B. E., Martin, S. K., Cafalchio, M., Islam, M. N., Aggleton, J. P., & O'Mara, S. M. (2021). Anterior thalamic inputs are required for subiculum spatial coding, with associated consequences for hippocampal spatial memory. *Journal of Neuroscience*, *41*(30), 6511–6525.
- Goulas, A., Uylings, H. B. M., & Hilgetag, C. C. (2017). Principles of ipsilateral and contralateral cortico-cortical connectivity in the mouse. *Brain Structure and Function*, *222*(3), 1281–1295.
- Gray, C. M., Maldonado, P. E., Wilson, M., & McNaughton, B. (1995). Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex. *Journal of Neuroscience Methods*, *63*(1), 43–54.
- Grieves, R. M., & Jeffery, K. J. (2017). The representation of space in the brain. *Behavioural Processes*, *135*, 113–131.
- Grinstead, C. M., & Snell, J. L. (2012). *Introduction to probability*. American Mathematical Soc.
- Hagberg, A. A., Schult, D. A., & Swart, P. J. (2008). Exploring network structure, dynamics, and function using NetworkX. In G. Varoquaux, T. Vaught, & J. Millman (Eds.), *Proceedings of the 7th python in science conference* (pp. 11–15).
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., Fernández del Río, J.,

- Wiebe, M., Peterson, P., . . . Oliphant, T. E. (2020). Array programming with NumPy. *Nature*, *585*, 357–362.
- Holmgren, C., Harkany, T., Svennenfors, B., & Zilberter, Y. (2003). Pyramidal cell communication within local networks in layer 2/3 of rat neocortex. *The Journal of physiology*, *551*(1), 139–153.
- Hunter, J. D. (2007). Matplotlib: A 2d graphics environment. *Computing in Science Engineering*, *9*(3), 90–95.
- Jay, T. M., & Witter, M. P. (1991). Distribution of hippocampal CA1 and subicular efferents in the prefrontal cortex of the rat studied by means of anterograde transport of phaseolus vulgaris-leucoagglutinin. *Journal of Comparative Neurology*, *313*(4), 574–586.
- Jefferys, J. G. (1995). Nonsynaptic modulation of neuronal activity in the brain: Electric currents and extracellular ions. *Physiological Reviews*, *75*(4), 689–723.
- Johansson, F., & The mpmath development team. (2018, December). *Mpmath: A python library for arbitrary-precision floating-point arithmetic (version 1.1.0)*.
- 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., . . . Harris, T. D. (2017). Fully integrated silicon probes for high-density recording of neural activity. *Nature*, *551*(7679), 232–236.
- Kalisman, N., Silberberg, G., & Markram, H. (2003). Deriving physical connectivity from neuronal morphology. *Biological Cybernetics*, *88*(3), 210–218.
- Lein, E. S., Hawrylycz, M. J., Ao, N., Ayres, M., Bensinger, A., Bernard, A., Boe, A. F., Boguski, M. S., Brockway, K. S., Byrnes, E. J., et al. (2007). Genome-wide atlas of gene expression in the adult mouse brain. *Nature*, *445*(7124), 168–176.
- Liley, D. T. J., & Wright, J. J. (1994). Intracortical connectivity of pyramidal and stellate cells: Estimates of synaptic densities and coupling symmetry. *Network: Computation in Neural Systems*, *5*(2), 175–189.
- London, M., Roth, A., Beeren, L., Häusser, M., & Latham, P. E. (2010). Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature*, *466*(7302), 123–127.
- Lovász, L. (2012). *Large networks and graph limits* (Vol. 60). American Mathematical Soc.
- Martin, S. K., Aggleton, J. P., & O’Mara, S. M. (2020). Can we conjointly record direct interactions between neurons in vivo in anatomically-connected brain areas? probabilistic analyses and further implications. *bioRxiv*, 2020.12.07.415125.
- Martinez-Banaclocha, M. (2018). Ephaptic coupling of cortical neurons: Possible contribution of astroglial magnetic fields? *Neuroscience*, *370*, 37–45.

- McKinney, W. (2010). Data structures for statistical computing in python. In S. v. d. Walt & J. Millman (Eds.), *Proceedings of the 9th python in science conference* (pp. 56–61).
- Mechler, F., Victor, J. D., Ohiorhenuan, I., Schmid, A. M., & Hu, Q. (2011). Three-dimensional localization of neurons in cortical tetrode recordings. *Journal of Neurophysiology*, *106*(2), 828–848.
- Moser, E. I., Kropff, E., & Moser, M.-B. (2008). Place cells, grid cells, and the brain's spatial representation system. *Annual Review of Neuroscience*, *31*, 69–89.
- O'Keefe, J., & 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.
- O'Mara, S. M., & Aggleton, J. P. (2019). Space and memory (far) beyond the hippocampus: Many subcortical structures also support cognitive mapping and mnemonic processing. *Frontiers in Neural Circuits*, *13*.
- O'Mara, S. M., Commins, S., Anderson, M., & Gigg, J. (2001). The subiculum: A review of form, physiology and function. *Progress in Neurobiology*, *64*(2), 129–155.
- Pang, J. C., Aquino, K. M., Oldehinkel, M., Robinson, P. A., Fulcher, B. D., Breakspear, M., & Fornito, A. (2023). Geometric constraints on human brain function. *Nature*, *618*(7965), 566–574.
- Reimann, M. W., Gevaert, M., Shi, Y., Lu, H., Markram, H., & Muller, E. (2019). A null model of the mouse whole-neocortex micro-connectome. *Nature Communications*, *10*(1), 3903.
- Reimann, M. W., Nolte, M., Scolamiero, M., Turner, K., Perin, R., Chindemi, G., Dłotko, P., Levi, R., Hess, K., & Markram, H. (2017). Cliques of neurons bound into cavities provide a missing link between structure and function. *Frontiers in Computational Neuroscience*, *11*.
- Ropireddy, D., Bachus, S. E., & Ascoli, G. A. (2012). Non-homogeneous stereological properties of the rat hippocampus from high-resolution 3d serial reconstruction of thin histological sections. *Neuroscience*, *205*, 91–111.
- Semedo, J. D., Zandvakili, A., Machens, C. K., Yu, B. M., & Kohn, A. (2019). Cortical areas interact through a communication subspace. *Neuron*, *102*(1), 249–259.e4.
- Shahsavarani, S., Thibodeaux, D. N., Xu, W., Kim, S. H., Lodgher, F., Nwokeabia, C., Cambareri, M., Yagielski, A. J., Zhao, H. T., Handwerker, D. A., Gonzalez-Castillo, J., Bandettini, P. A., & Hillman, E. M. (2023). Cortex-wide neural dynamics predict behavioral states and provide a neural basis for resting-state dynamic functional connectivity. *Cell Reports*, *42*(6), 112527.
- Shibata, H. (1993). Direct projections from the anterior thalamic nuclei to the retrohippocampal region in the rat. *Journal of Comparative Neurology*, *337*(3), 431–445.

- Siegle, J. H., Jia, X., Durand, S., Gale, S., Bennett, C., Graddis, N., Heller, G., Ramirez, T. K., Choi, H., Luviano, J. A., Groblewski, P. A., Ahmed, R., Arkhipov, A., Bernard, A., Billeh, Y. N., Brown, D., Buice, M. A., Cain, N., Caldejon, S., . . . Koch, C. (2021). Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature*, *592*(7852), 86–92.
- 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., . . . Harris, T. D. (2021). Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science*, *372*(6539).
- Steinmetz, N. A., Koch, C., Harris, K. D., & Carandini, M. (2018). Challenges and opportunities for large-scale electrophysiology with neuropixels probes. *Current Opinion in Neurobiology*, *50*, 92–100.
- Steinmetz, N. A., Zatka-Haas, P., Carandini, M., & Harris, K. D. (2019). Distributed coding of choice, action and engagement across the mouse brain. *Nature*, *576*(7786), 266–273.
- Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C. B., Carandini, M., & Harris, K. D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, *364*(6437), eaav7893.
- Sutherland, R. J., & Rudy, J. W. (1989). Configural association theory: The role of the hippocampal formation in learning, memory, and amnesia. *Psychobiology*, *17*(2), 129–144.
- Taube, J. S., Muller, R. U., & Ranck, J. B. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. i. description and quantitative analysis. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *10*(2), 420–435.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., . . . SciPy 1.0 Contributors. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in python. *Nature Methods*, *17*, 261–272.
- Waskom, M., & The seaborn development team. (2020, September). *Mwaskom/seaborn* (Version latest). Zenodo.
- Wei, Z., Lin, B.-J., Chen, T.-W., Daie, K., Svoboda, K., & Druckmann, S. (2020). A comparison of neuronal population dynamics measured with calcium imaging and electrophysiology. *PLOS Computational Biology*, *16*(9), e1008198.
- West, M. J., Danscher, G., & Gydesen, H. (1978). A determination of the volumes of the layers of the rat hippocampal region. *Cell and Tissue Research*, *188*(3), 345–359.

- Witter, M. P., & Groenewegen, H. J. (1990). The subiculum: Cytoarchitectonically a simple structure, but hodologically complex. *Progress in Brain Research*, *83*, 47–58.
- Witter, M. P. (2007). Intrinsic and extrinsic wiring of CA3: Indications for connectional heterogeneity. *Learning & Memory*, *14*(11), 705–713.
- Yang, W., Carrillo-Reid, L., Bando, Y., Peterka, D. S., & Yuste, R. (2018). Simultaneous two-photon imaging and two-photon optogenetics of cortical circuits in three dimensions (K. Svoboda, Ed.). *eLife*, *7*, e32671.

7 Relating task performance to correlation in multi-region neural data

Introduction: The relationship between the timing of action potentials from multiple neurons with stimuli and observed behaviours has been a long-standing topic of interest. From a theoretical view, increased neural correlation implies reduced information transmission. Despite this, evidence suggests that increased neural correlation could enhance task performance. This suggests a delicate balance between neural information efficiency and stability, with task performance potentially influencing this balance. In this study, we investigate this relationship and its dependence on task performance.

Methods: We compare coherence from local field potentials in the claustrum, anterior cingulate cortex, and retrosplenial cortex across fixed interval and fixed ratio tasks. We then compare correlations between a reduced dimension representation of neural firing in catch and miss trials during a visual change detection task. To identify these low-dimensional representations of neural firing, we employ canonical correlation analysis. Additionally, we use Gaussian process factor analysis to generate latent neural trajectories in visual and limbic system brain areas, allowing us to assess the similarity of computations between catch and miss trials in the same visual change detection task.

Results: Coherence differed in a region-specific manner between fixed interval and fixed ratio trials. Latent neural trajectories were distinctly separable between catch and miss trials. However, the shape of the neural trajectories often appeared similar for catch and miss trials, suggesting that related computations were being performed. The correlation between the reduced dimensional activity differed between catch and miss trials, with higher correlation in catch trials for visual areas and lower correlation in catch trials for limbic system areas.

Conclusions: Our findings demonstrate that trial-to-trial correlations vary across similar tasks and task performance, but in a region-specific manner. Thus, the correlation-task relationship appears to be highly task- and brain region-dependent, lacking a general rule. These results offer insights that can help explain the conflicting reports on the correlation-task relationship.

7.1 My contributions to this study

In this multi-author study I did not perform any of the electrophysiology in the O'Mara lab, which was performed by Dr. Ham. The initial data analysis and data processing on our in-lab data was a joint contribution with Dr. Ham, which I continued after COVID-19 complications. Additionally, I wrote the behavioural training program logic and helped set up the hardware for the experiment. Recordings in mice were gratefully received from the Allen Institute and International Brain Laboratory. However, the analysis of this open-data was solely my contribution.

7.2 Overview of neural correlation and the correlation-behaviour relationship

Understanding the intricate relationship between neural activity and behaviour is a fundamental pillar of neuroscience research. A common approach to this problem is to consider correlations in the neural activity, which can be measured in several ways (Averbeck et al., 2006; Kohn et al., 2016). About a century ago, experimental practices and theoretical considerations gave rise to a prevailing belief that noise correlations, which represent trial-to-trial variations in stimulus responses, inherently compromised the information conveyed about the environment and stimuli (Barlow, 2001). In contrast, signal correlations, reflecting the similarity of response tuning curves between pairs of neurons, were frequently employed as a metric for evaluating neural activity. Anchored in the efficient coding hypothesis, which posits that the brain minimises spike count to transmit a given signal, both noise and signal correlations were commonly regarded as undesirable, meriting minimisation. This perspective was heavily influenced by Shannon's information theory (Shannon, 1948).

However, as larger-scale recordings were performed and theoretical understanding advanced, the interpretation of neural correlations became more nuanced (Averbeck et al., 2006; Panzeri et al., 1999). Theoretical relationships between signal and noise correlations suggested that simultaneous increases or decreases in both measures would likely reduce information representation, while movement in opposite directions could increase information representation. Specific types of correlations that limited information were also identified, such as those proportional to the product of the derivatives of the tuning curves (Moreno-Bote et al., 2014). The concept of a "noise" correlation itself was recognised as a misnomer (Katz et al., 2021), leading to adoption of alternative terms such as trial-to-trial correlation or spike count correlation.

These advances in understanding correlation dynamics led to significant discoveries, such as the default mode network (Greicius et al., 2003; Raichle et al., 2001), which represents a network that is “task negative” and can exhibit reduced activity and correlation during task performance. Additionally, the concept of mixed selectivity emerged, demonstrating that neurons can exhibit different activity patterns in different contexts, resembling members of different ensembles (Fusi et al., 2016; Rigotti et al., 2013). Given the interplay of various concepts such as correlation, information representation, redundant coding, dimensionality, and mixed selectivity, researchers have studied their relationship with behaviour. Studies have explored the relationship between correlation and learning (Failor et al., 2021), correlation and task performance (Cayco-Gajic & Zylberberg, 2021; Valente et al., 2021), correlation and information transformation in the brain (Ito et al., 2019; Larry & Joshua, 2021), and correlation and coding properties (Franke et al., 2016). Recent investigations have applied these concepts to analyse decision making tasks, such as image change detection (Ni et al., 2018; Siegle et al., 2021), visual contrast discrimination (Steinmetz et al., 2019; Zatka-Haas et al., 2021), and virtual reality tasks (Lee et al., 2020).

Considering the growing body of literature on neural correlation, recent findings have prompted a shift in perspective towards a nuanced understanding of the functional implications of changes in correlation between the responses of multiple neurons. For instance, studies have revealed distinct activity patterns for feedforward and feedback information processing, with the correlation between visual areas often highest during spontaneous activity rather than in response to a stimulus (Semedo et al., 2022). Additional support for this perspective comes from Rhodri Cusack (personal communication), where the functional connectivity between the frontal and auditory cortices decreased during wakefulness, likely due to tangential processing occurring in those areas or non-linear transformations between the two regions. However, it has been shown that spike count correlation in the posterior parietal cortex was significantly higher when mice correctly navigated a T-maze to an auditory cue (Cayco-Gajic & Zylberberg, 2021; Valente et al., 2021). These observations suggest that high correlation between brain regions may not always indicate informative communication – yet there seems to be a contradiction – neural correlation may improve task performance.

In this study, our focus is on the relationship between correlation and task performance. We consider various forms of correlation combined via canonical correlation analysis (CCA) and gaussian process factor analysis (GPFA). Based on experimental findings linking increased correlation between neural activity with task performance, our primary hypothesis posits that the correlation between multiple brain regions implicated in a given task will differ between correctly performed tasks and incorrectly performed tasks. We extend this hypothesis to encompass not only binary outcomes of pass or fail

tasks, but also scenarios involving swapping between two similar tasks. Nevertheless, we must consider that increased correlation may not signify informative communication, and if so, why might a change in correlation between neural activity be important for task performance? Our proposition here is that there exists a delicate balance between brain areas' correlation (perhaps an equilibrium state), enabling a stable representation of the environment while avoiding excessive hindrance to information transmission. An underlying mechanism for this system could involve the initial state of the system serving as a critical task input, potentially more influential than a change in the computational processes performed. This proposition finds some support in recent observations, particularly in the motor cortex, where brain dynamics exhibit similarities to a dynamical system (Keshtkaran et al., 2022; Pandarinath et al., 2018; Rodriguez et al., 2023). Consequently, we hypothesise that neural trajectories will exhibit similar shapes between correct and error trials, but differ in their starting and ending points.

To address these hypotheses, we aim to answer two main questions:

1. How consistent are the interactions between multiple brain regions during successful and failed task performance?
2. How correlated are the responses between two brain areas when executing distinct yet related tasks or during successful and failed task performance?

We will explore these dynamics in two distinct sets of brain regions and tasks. In the first set, we will analyse data from brain regions associated with memory and decision-making during a task in which a rat monitors its surroundings, adjusting behavior based on environmental cues and feedback. Here, we can test the hypothesis of correlation between neural activity differing in two similar tasks. In the second set, we will analyse data from visual and limbic system brain regions during which a mouse is engaged in a visual stimulus change detection task followed by a motor response. Here, the visual areas may be responsible for forming a stable representation of the environment, which is subsequently transmitted to the limbic system for higher-level processing and memory of previous stimuli. This allows us to investigate a situation where high correlation between neural activity could prove to be beneficial or detrimental. Here we will test our hypothesis of a difference in correlation between neural activity across binary task performance, along with the similarity of neural trajectories during correct and incorrect trials. By exploring these inquiries, we aim to enhance our understanding of how correlation dynamics influence task performance, information representation, and the interplay between different brain regions.

7.3 Relating neural activity in multiple regions to tasks and behaviour

7.3.1 Coherence between the claustrum, RSC, and ACC during operant conditioning

We designed a task whereby rats would perform either fixed ratio or fixed interval tasks, but have to switch between them to receive rewards. In the fixed ratio mode, rats would be required to press a lever six times to get a reward. This encourages fast lever pressing as the optimal strategy. However, in the fixed interval task, rats had to wait 30 seconds between lever presses to get a reward. To discourage lever presses during this interval, a double reward was offered for successful waiting, encouraging a more spaced-out lever pressing strategy. The total testing duration was set at 30 minutes, during which rats encountered three blocks of five-minute-long paradigms for both fixed interval and fixed ratio tasks. The order of these blocks were randomised. Rats had three primary cues to determine which paradigm they should be performing. The first cue was a light above the designated lever for pressing. Rats, having been trained to associate each of two levers with a different task, could determine the corresponding task based on the illuminated lever. Additionally, when a new block was starting (indicating a possible paradigm shift, not a guaranteed shift, as the same task could be repeated) the levers would briefly retract and the main buzzer would sound. The third cue was, of course, whether they were receiving any reward. The rats learned to differentiate between the two tasks, and displayed different behaviours and lever pressing strategies for each task (Figure 7.1). This suggested that the rats were able to learn to switch between the two tasks when the underlying rules changed. See subsection 7.5.1 for more details on the operant conditioning data.

Vast anatomical connections stem from the claustrum (CLA), though the functional role of the CLA is not well known (Dillingham et al., 2017; Ham & Augustine, 2022; White et al., 2018). Motivated by this, we aimed to investigate if the CLA might play a role in orchestrating information between brain areas and influencing decision-making in a task that alternates between fixed ratio and fixed interval paradigms. We anticipated that this task would involve spatial and memory components, with the recollection of the task itself and cues to switch between fixed ratio and fixed interval paradigms. Additionally, we expected the task to include decision-making, reward maximisation, and learning aspects. Functionally, the retrosplenial cortex (RSC) is posited to contribute to the spatial and memory processes (Todd & Bucci, 2015; Vann et al., 2009), while the anterior cingulate cortex (ACC) may play a role in reward maximisation and decision-making processes (Bush et al., 2002; Hart et al., 2019;

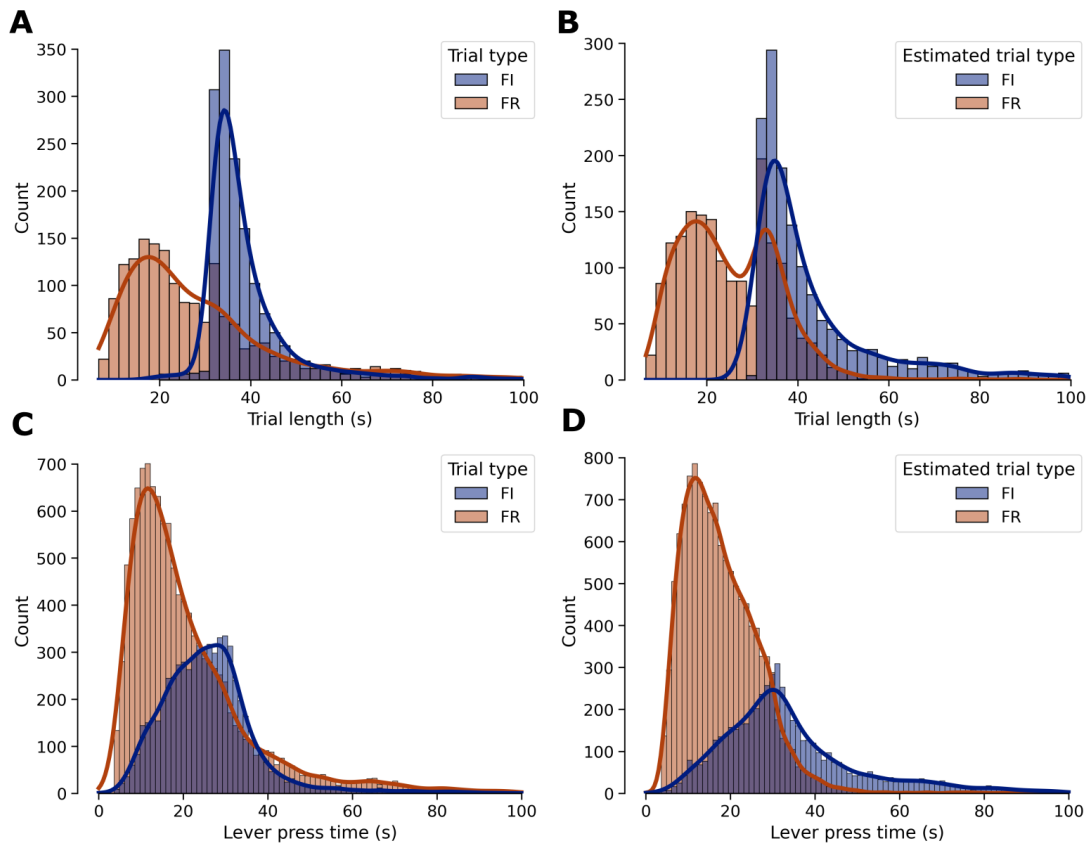


FIGURE 7.1: The length of trials and distributions of lever press times in fixed interval and fixed ratio tasks.

(A) The distribution of the length of trials (from when the trial starts to reward collection) in switching blocks of fixed interval and fixed ratio tasks.

(B) The distribution of the times at which rats pressed the lever during switching blocks of fixed interval and fixed ratio tasks.

(C, D) The same as in A, B, but the trial type is estimated based on the lever presses during the trial. For instance, if the rat pressed the lever in quick succession repeatedly from the start of the trial, that would be estimated as a fixed ratio trial (even if it was during an intended fixed interval block).

Heilbronner & Hayden, 2016; Kolling et al., 2016). Anatomically, the CLA has known projections to the RSC (Ham & Augustine, 2022), and reciprocal connections with the cingulate cortex (Q. Wang et al., 2017). In the context of this task, a possible function of the CLA would then involve assisting in orchestrating or controlling interplay between the RSC and CLA at a higher level.

The coherence between the claustrum (CLA), anterior cingulate cortex (ACC), and retrosplenial cortex (RSC) was compared between the estimated fixed ratio and fixed interval tasks based on the behaviour (see Table 7.5.1 for details of the behaviour estimation procedure). All tests were performed with the Mann-Whitney U test, and the number of coherence band values for each group were $N_1 = 415$ FI trials and $N_2 = 375$ FR trials (see Table 7.1 for full results). The coherence between the CLA and the ACC

(Figure 7.2) was significantly greater in the low and high gamma bands ($p < 0.001$) for the fixed interval (FI) task compared to the fixed ratio (FR) task (for the other bands, the results were: Delta $p = 0.24$, $BF_{10} = 0.10$, Theta $p = 0.10$, $BF_{10} = 0.15$, Beta $p = 0.10$, $BF_{10} = 0.42$). The coherence between the CLA and the RSC was significantly higher ($p < 0.001$) for the FR task compared to the FI task in all bands except the theta band ($p = 0.22$, $BF_{10} = 0.179$). The coherence between the RSC and the ACC was significantly higher ($p < 0.001$) for the FR task compared to the FI task in all bands. Overall, these differences were small to medium in size (Rank-biserial correlation $|r| \in (0.15, 0.4)$, see Table 7.1).

7.3.2 Correlation in visual brain areas and limbic system brain areas

Using canonical correlation analysis (CCA), a linear combination of the binned firing rates with maximal correlation was calculated for multiple brain regions, two regions at a time, with at least 15 units in each region (see an example in Figure 7.3). CCA was performed for visual and limbic system related brain regions from data captured in the Allen Institute for Brain Science (AIS) Visual Behaviour Neuropixels (VBN) dataset. The mice recorded from were performing a visual change detection task, and sessions had a mean of 250 trials, with an average catch rate to close to 80%. See more details on this dataset in subsection 7.5.2.

This canonical correlation was found to be separable based on the task performance, however, whether the correlation increased or decreased was dependent on the brain regions being compared (Figure 7.4). Canonical correlations between the primary visual area (VISp) and lateral visual area (VISl), between VISp and the ventral part of the lateral geniculate nucleus (LGv), and between VISp and the thalamus were found to be significantly higher for catch trials than for miss trials (Table 7.2). However, the effect size was small to medium, with Cohen's d values of 0.24, 0.36, and 0.49 respectively. For limbic system areas, canonical correlations between CA1 and subiculum and between CA1 and thalamus were found to be significantly lower for catch trials than for miss trials (Table 7.2). The effect size for this difference was very small with Cohen's d values of -0.1 and -0.07 respectively. The canonical correlation for the subiculum and thalamus pair was not significantly different when comparing catch trials against miss trials ($p = 0.9$).

7.3.3 The shape of neural trajectories in latent space

Using gaussian process factor analysis (GPFA), neural trajectories were calculated in latent space for the AIS VBN dataset. This was performed using the data from

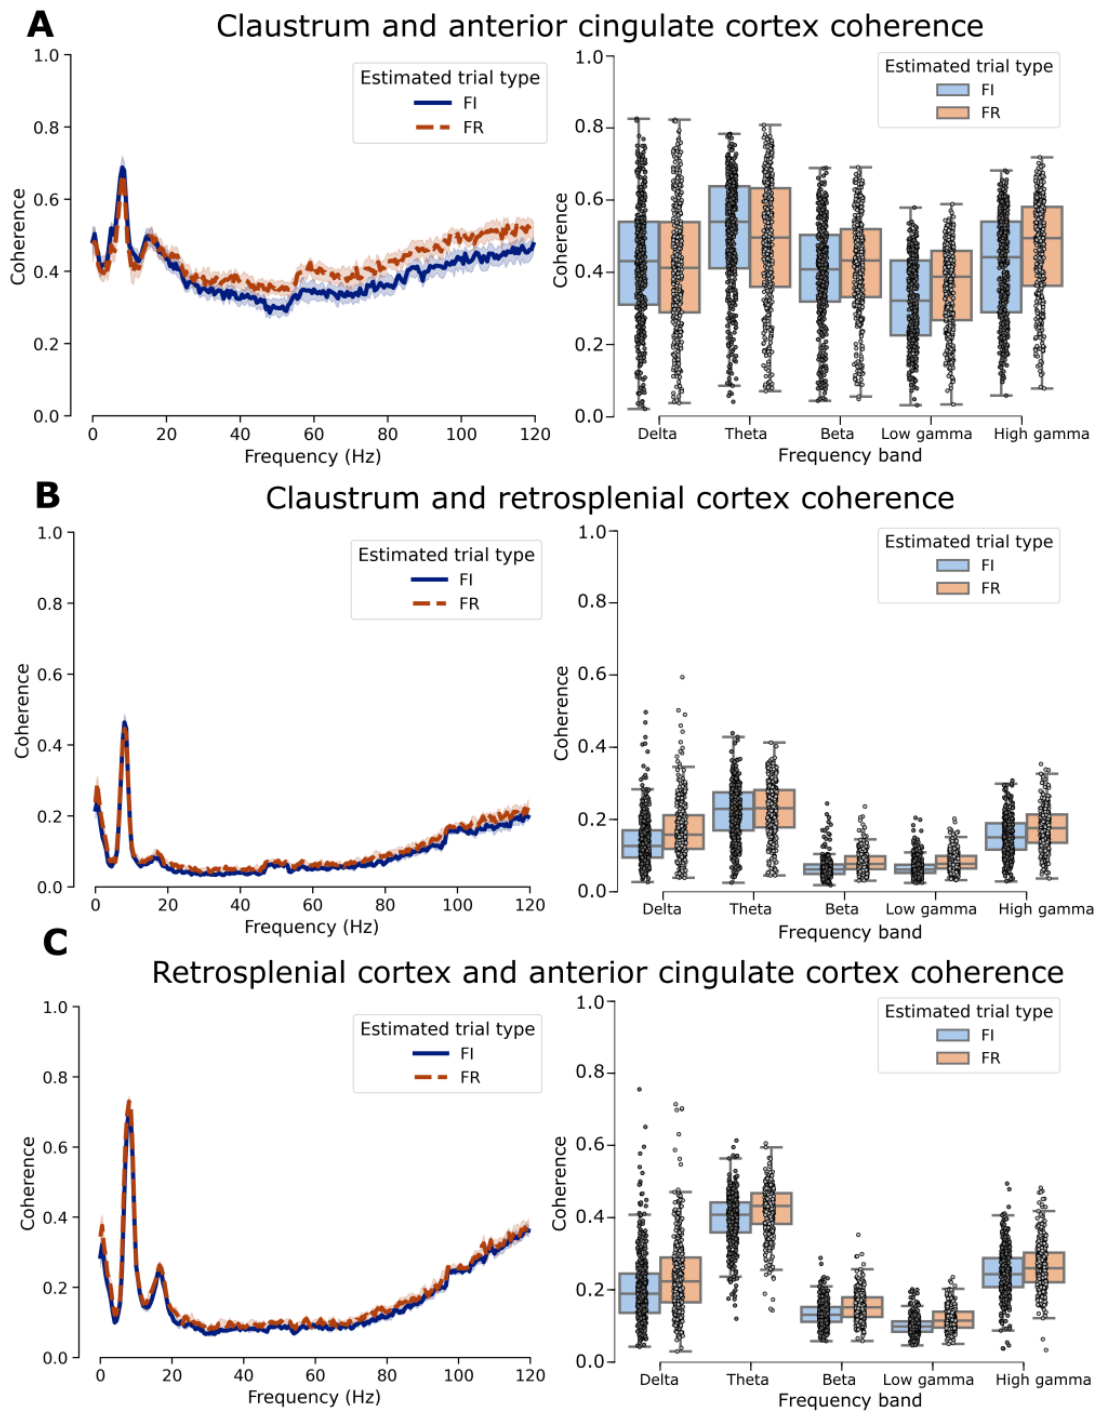


FIGURE 7.2: The coherence between the claustrum, anterior cingulate cortex, and retrosplenial cortex in fixed interval and fixed ratio tasks.

(A) Coherence between the claustrum and anterior cingulate cortex.

(B) Coherence between the claustrum and retrosplenial cortex.

(C) Coherence between the anterior cingulate cortex and retrosplenial cortex.

Key: Panels on the left are line plots of the average coherence with 95% confidence intervals. Panels on the right are boxplots of the band-limited coherence, with individual data points overlaid.

	Delta		Theta		Beta		Low gamma		High gamma	
	FI	FR	FI	FR	FI	FR	FI	FR	FI	FR
Median	0.431	0.412	0.540	0.496	0.408	0.432	0.320	0.387	0.442	0.494
IQR	0.230	0.251	0.228	0.274	0.185	0.189	0.208	0.192	0.252	0.219
W	81557		83133		72579		63776		62184	
p	0.24		0.10		0.10		< 0.001		< 0.001	
BF_{10}	0.10		0.15		0.42		211.6		474.04	
r_B	0.05		0.07		-0.07		-0.18		-0.2	
Claustrum and anterior cingulate cortex coherence										
	Delta		Theta		Beta		Low gamma		High gamma	
	FI	FR	FI	FR	FI	FR	FI	FR	FI	FR
Median	0.127	0.158	0.229	0.232	0.062	0.077	0.062	0.078	0.151	0.176
IQR	0.076	0.094	0.106	0.104	0.026	0.036	0.023	0.035	0.073	0.078
W	57528		73905		47546		47213		59628	
p	0.003		0.22		< 0.001		< 0.001		< 0.001	
BF_{10}	> 10000		0.179		> 10000		> 10000		2150	
r_B	-0.26		-0.05		-0.39		-0.39		-0.23	
Claustrum and retrosplenial cortex coherence										
	Delta		Theta		Beta		Low gamma		High gamma	
	FI	FR	FI	FR	FI	FR	FI	FR	FI	FR
Median	0.190	0.224	0.409	0.433	0.131	0.152	0.099	0.115	0.244	0.260
IQR	0.109	0.124	0.083	0.085	0.041	0.054	0.029	0.044	0.081	0.082
W	62456		60743		52964		50318		65774	
p	< 0.001		< 0.001		< 0.001		< 0.001		< 0.001	
BF_{10}	294.7		1055.5		> 10000		> 10000		41.7	
r_B	-0.2		-0.22		-0.32		-0.35		-0.15	
Retrosplenial cortex and anterior cingulate cortex coherence										

TABLE 7.1: Median and interquartile range of coherence between the claustrum, anterior cingulate cortex, and the retrosplenial cortex. Statistical results are from the Mann-Whitney U test, W is the test statistic, p is the p -value, BF_{10} is the Bayesian factor in support of the alternative hypothesis, and r_B is the rank-biserial correlation coefficient indicating effect size. FI indicates the fixed interval trials, and FR indicates the fixed ratio trials. Recordings from 5 rats are included in the comparisons with 415 FI and 375 FR trials overall (the trial type was estimated based on behaviour).

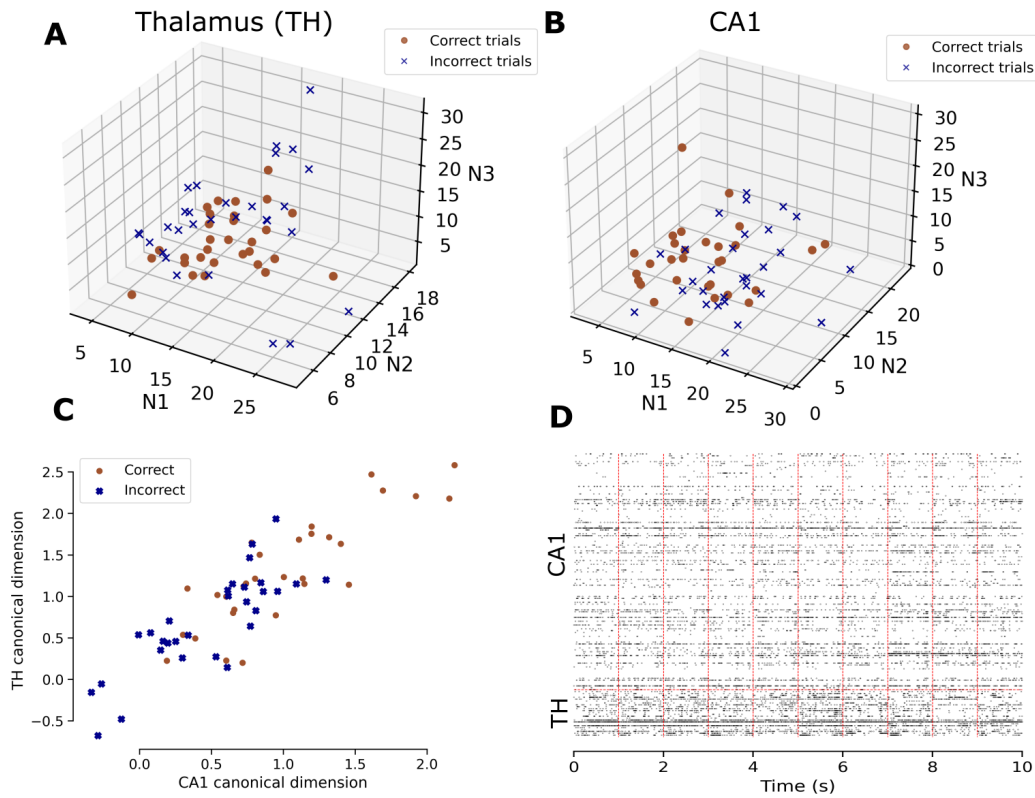


FIGURE 7.3: Demonstrating the canonical correlation analysis performed on the multi-region data during change detection trials.

(A, B) The firing rate of three units from the thalamus (A) and CA1 (B) over 60 trials in the one second window after stimulus presentation. Each dot in the scatter plot represents a change-detection catch (correct) or miss (incorrect) trial.

(C) A plot of those same 60 trials as in (A, B), but after reducing the firing rates of each neuron to a single latent variable via CCA for each brain region. The activity along the canonical dimensions is highly correlated.

(D) The spike raster plot for the first 10 trials plotted in (A, B). The dotted red line splits the trials into one-second-long windows.

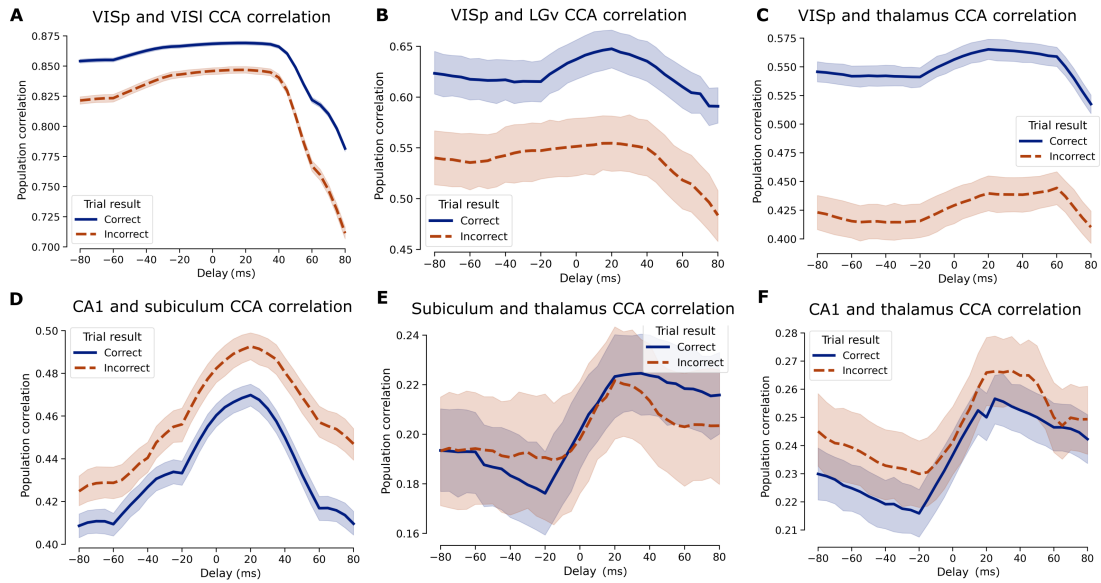


FIGURE 7.4: Computing the mean correlation between the dimension reduced activity in multiple-brain regions across catch (correct) and miss (incorrect) trials. The x -axis delay represents a comparison between shifted versions of the spike trains (e.g., a delay of d ms for VISp and VISl indicates that the spike train from VISp at $[t, t + 1]$ is compared against the spike train from VISl at $[t + d, t + 1 + d]$). The number of trials used to calculate the correlation between the latent activity in each brain region is as follows:

A VISp and VISl: 13500 catch trials and 6883 miss trials.

B VISp and LGv: 620 catch trials and 382 miss trials.

C VISp and thalamus: 1994 catch trials and 1174 miss trials.

D CA1 and subiculum: 7375 catch trials and 3632 miss trials.

E Subiculum and thalamus: 983 catch trials and 407 miss trials.

F CA1 and thalamus: 490 catch trials and 135 miss trials.

Key: VISp – primary visual area (V1), VISl – lateral visual area, LGv – ventral part of the lateral geniculate nucleus.

Brain regions	Trial result	N	Mean	Std	t	p	Cohen's d																																																								
VISp and VISl	Catch	13500	0.866	0.092	15.34	< 0.001	0.24																																																								
	Miss	6883	0.84	0.12				VISp and LGv	Catch	620	0.647	0.24	5.44	< 0.001	0.36	Miss	382	0.555	0.28	VISp and TH	Catch	2994	0.565	0.25	14.23	< 0.001	0.49	Miss	1174	0.44	0.26	CA1 and SUB	Catch	7375	0.47	0.23	-5.2	< 0.001	-0.1	Miss	3632	0.49	0.21	SUB and TH	Catch	983	0.223	0.26	0.119	0.905	0.007	Miss	407	0.222	0.23	CA1 and TH	Catch	3417	0.25	0.26	-2.1	0.038	-0.064
VISp and LGv	Catch	620	0.647	0.24	5.44	< 0.001	0.36																																																								
	Miss	382	0.555	0.28				VISp and TH	Catch	2994	0.565	0.25	14.23	< 0.001	0.49	Miss	1174	0.44	0.26	CA1 and SUB	Catch	7375	0.47	0.23	-5.2	< 0.001	-0.1	Miss	3632	0.49	0.21	SUB and TH	Catch	983	0.223	0.26	0.119	0.905	0.007	Miss	407	0.222	0.23	CA1 and TH	Catch	3417	0.25	0.26	-2.1	0.038	-0.064	Miss	1419	0.266	0.24								
VISp and TH	Catch	2994	0.565	0.25	14.23	< 0.001	0.49																																																								
	Miss	1174	0.44	0.26				CA1 and SUB	Catch	7375	0.47	0.23	-5.2	< 0.001	-0.1	Miss	3632	0.49	0.21	SUB and TH	Catch	983	0.223	0.26	0.119	0.905	0.007	Miss	407	0.222	0.23	CA1 and TH	Catch	3417	0.25	0.26	-2.1	0.038	-0.064	Miss	1419	0.266	0.24																				
CA1 and SUB	Catch	7375	0.47	0.23	-5.2	< 0.001	-0.1																																																								
	Miss	3632	0.49	0.21				SUB and TH	Catch	983	0.223	0.26	0.119	0.905	0.007	Miss	407	0.222	0.23	CA1 and TH	Catch	3417	0.25	0.26	-2.1	0.038	-0.064	Miss	1419	0.266	0.24																																
SUB and TH	Catch	983	0.223	0.26	0.119	0.905	0.007																																																								
	Miss	407	0.222	0.23				CA1 and TH	Catch	3417	0.25	0.26	-2.1	0.038	-0.064	Miss	1419	0.266	0.24																																												
CA1 and TH	Catch	3417	0.25	0.26	-2.1	0.038	-0.064																																																								
	Miss	1419	0.266	0.24																																																											

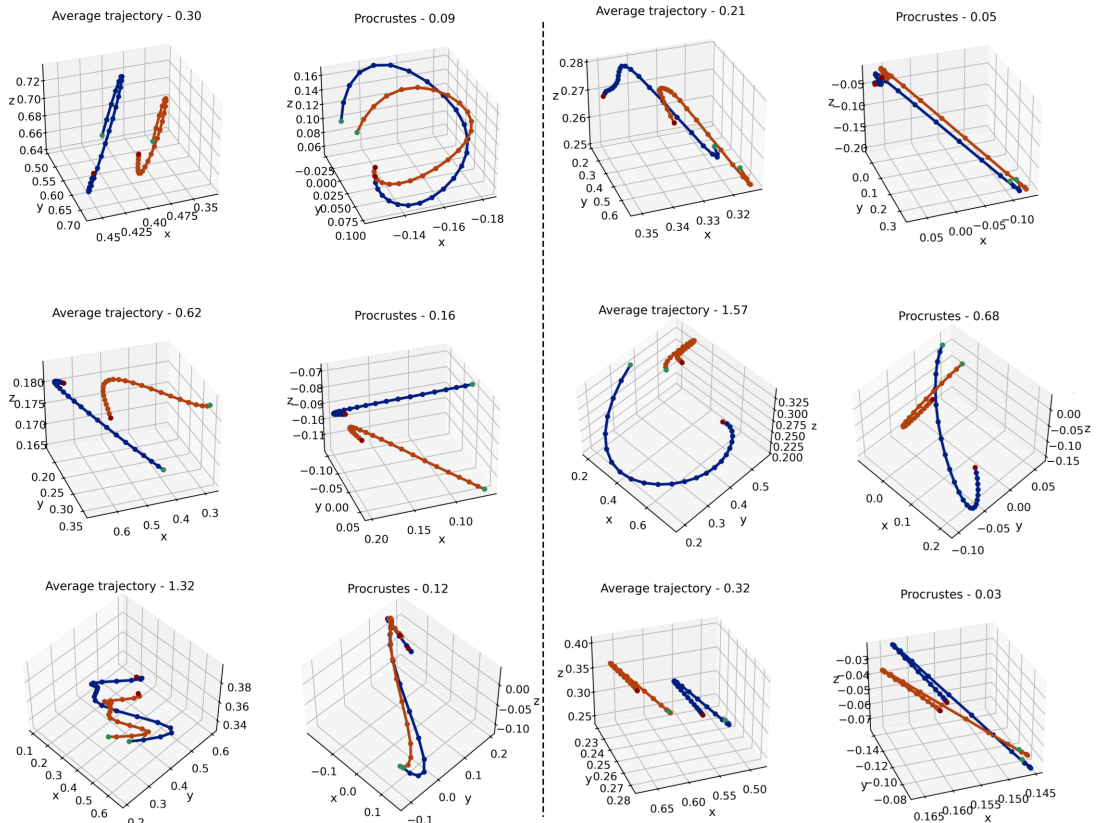
TABLE 7.2: The correlation between the dimension reduced activity in multiple-brain regions compared across catch and miss trials. The results are presented for the delay at which the correlation was maximal. Statistical tests were performed using a two-tailed Welch's t -test. The Cohen's d represents the difference in means between the two groups divided by the pooled standard deviation.

three sets of brain regions at a time if at least 10 units were present in each of the brain regions. The trajectories commonly take the form of a discrete version of a smooth continuous curve without sharp cusps, and are well separated between catch and miss trials (Figure 7.5). Additionally, Procrustes analysis (a form of curve alignment) suggests that the shapes of these trajectories are similar.

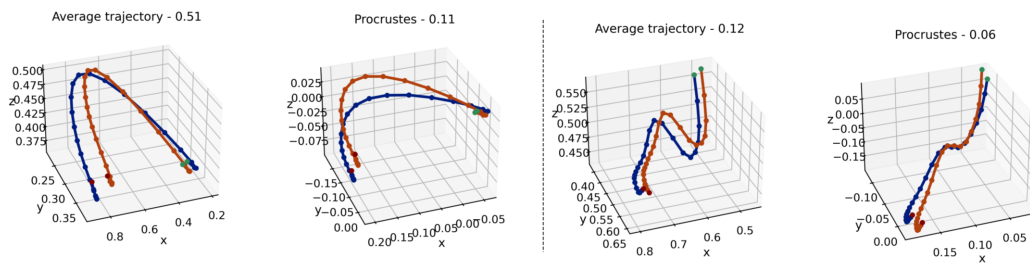
The neural trajectories were found to be similar in shape between correct and error trials, but the start and end points of the trajectories were well distanced (Figure 7.6). For GPFA performed on units from the subiculum, the thalamus, and CA1, the difference in shape between the curves in correct trials and incorrect trials was significantly positively correlated with the difference in the start and end points of the curves ($r = 0.84$, $p < 0.001$, $n = 18$, Figure 7.6A). Additionally, the variance in the neural trajectories was not correlated with the difference in shape between the curves ($r = -0.13$, $p = 0.60$, $n = 18$ Figure 7.6B). The distribution of curve shape similarities does not appear to be a discrete sample from a normal distribution, possibly representing two distinct populations of trials (Figure 7.6E). There was also no difference between the variance in neural trajectories between correct and error trials (Mann Whitney U test, $W = 172$, $p = 0.76$, $n = 18$ Figure 7.6F).

For GPFA performed on units from the primary visual area (VISp), lateral visual

CA1, Thalamus, and Subiculum



Primary visual area, LGv, and Thalamus



Primary and lateral visual areas and LGv

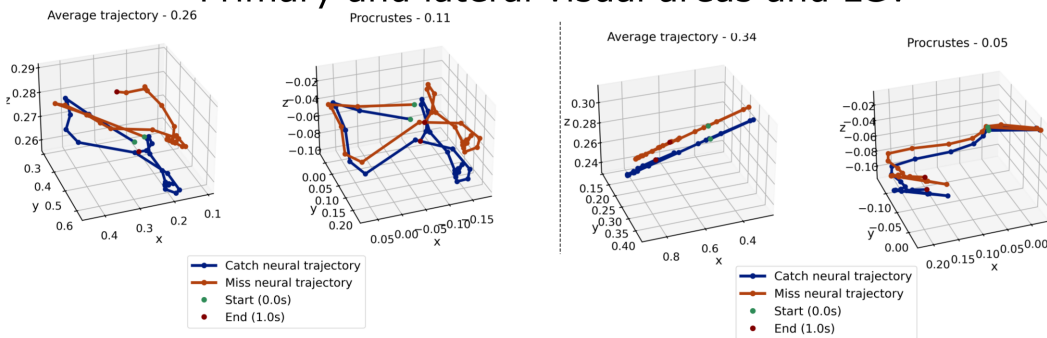


FIGURE 7.5: One-second-long latent neural trajectories extracted from visual and limbic areas during a change detection task. The dotted line separates examples into columns. Panels on the left of each column show average neural trajectories for correct change detections and misses and their Fréchet distance. Panels on the right of each column indicate the same after performing Procrustes analysis to align the two curves.

area (VISl), and the ventral part of the lateral geniculate nucleus (LGv), the difference in shape between the curves in correct trials and incorrect trials was insignificantly positively correlated with the difference in the start and end points of the curves ($r = 0.71$, $p = 0.12$, $n = 6$, Figure 7.6C). The sample size for the visual areas was small, and therefore the difference in variance between correct and incorrect trials was not statistically tested ($n = 3$ for each group).

7.4 Describing the balance between information representation and stability in neural activity

Coherence, which is a measure of linear correlation between two signals at a particular frequency level, may reflect important communication between brain regions. Two completely non-coherent signals represent the maximum amount of information at that frequency, but such a complete lack of correlation can lead to widely varying signals in the brain and a lack of redundancy in neural coding, resulting in no consensus between many pairs of signals. Given the brain's plastic nature, redundant coding is certainly present (Humphries, 2021), and the strengthening/weakening of neural pathways is related to correlation. As such, there is likely a delicately balanced system of highly correlated signals that indicate reliable evidence for stimuli and facilitate redundant coding, without being too correlated such that little information can be transferred. In this study, we aimed to investigate how correlation between neural action potential timing relates to task performance. Our results provide insights into this complex relationship and the dynamics of neural activity between different brain regions during varying task demands.

In the operant conditioning task, rats exhibited distinct behavioural patterns and lever pressing strategies for FI and FR tasks. This suggests that the rats were able to learn and adapt to the changing task requirements. Examining the coherence between brain regions during the operant conditioning task, we observed region-specific differences. The coherence between the claustrum and anterior cingulate cortex was high, with significantly higher coherence in the gamma band during the FI task compared to the FR task. In contrast, the coherence involving the retrosplenial cortex was lower overall and peaked only in the theta band. The coherence between the retrosplenial cortex and claustrum was significantly higher during the FR task in all bands except theta. Similarly, the coherence between the retrosplenial cortex and anterior cingulate cortex was significantly higher during the FR task across all bands. These findings suggest neural coherence is separable for two similar tasks, with different patterns emerging in different brain regions. In future studies, it would be interesting to

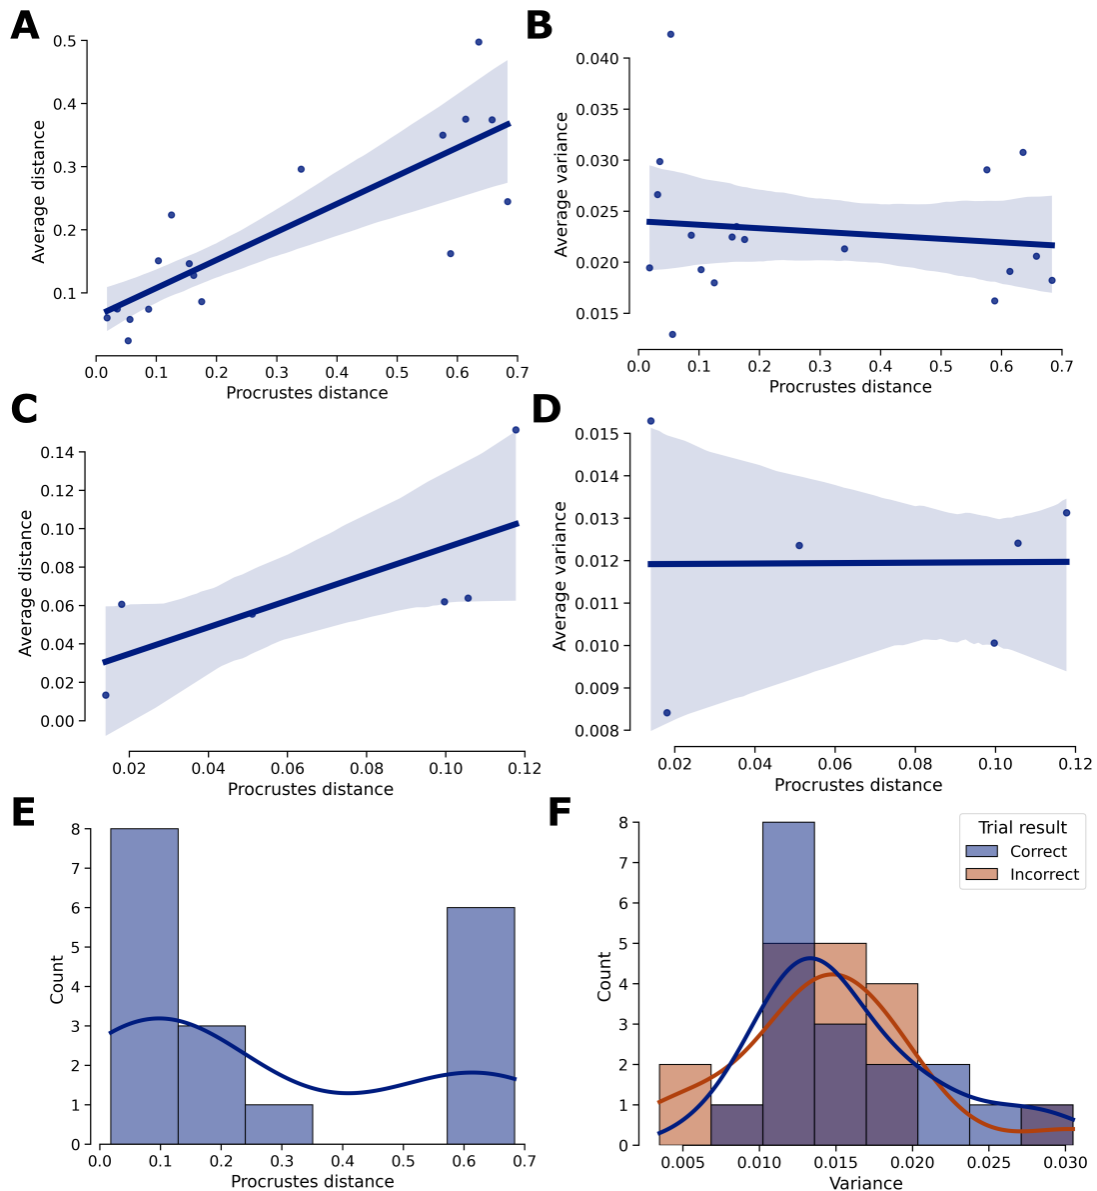


FIGURE 7.6: The relationship of the similarity of neural trajectory shapes to the distance between the start and end points of trajectories. Neural trajectories are estimated with GPFA and are compared across catch (correct) and miss (incorrect) trials to estimate separation of trajectories and shape similarity.

(A, C) The linear correlation of the mean Euclidean distance between the two starting and ending neural trajectory points (Average distance; y -axis) and the Fréchet distance between the Procrustes aligned trajectories (Procrustes distance; x -axis).

(B, D) The linear correlation of the variance of the neural trajectories and the Fréchet distance between the Procrustes aligned trajectories.

(E) The distribution of neural trajectory shape similarities measured by the Fréchet distance between trajectories.

(F) The distribution of neural trajectory variances across trials split.

Key: A, B, E, F data is from the subiculum, CA1, and thalamus. C, D data is from the primary visual area, lateral visual area, and lateral geniculate nucleus. Error shading indicates a 95% confidence interval.

investigate how dynamics in local field potentials (LFPs) relate to neural population activity.

When analysing single-unit activity, we applied CCA to pairs of brain regions at different delays to investigate correlated neural activity. We found that the correlation patterns between brain regions differed with task performance in a change detection task. In visual areas, including VISp, VISl, and ventral part of the lateral geniculate nucleus (LGN) canonical correlations were significantly higher in catch trials compared to miss trials. Conversely, in limbic system areas, including CA1, the subiculum, and the thalamus, canonical correlations were significantly lower or the same in catch trials compared to miss trials. This may be due to the distinct roles of these brain regions in the task. Perhaps, a high correlation between neural activity in visual areas is beneficial to the change detection task; indicating a stable representation of the visual stimulus with high confidence. On the other hand, perhaps, a high correlation between limbic system areas is detrimental to the change detection task; due to decreased information representation efficiency in decision making and memory related brain areas. These findings indicate that correlated activity between brain regions can be related with increased task performance, supporting the notion of a balance between representation stability and information representation efficiency.

Using GPFA to extract neural trajectories, we found clear separability between catch and miss trials in a change detection task. However, the shape of the neural trajectories was similar, though the start and end points of the curves differed between correct and error trials. This suggests that comparable computations are being performed, but with different inputs, resulting in different outputs. Additionally, the distance between the start and end points of the curves was correlated with the difference in shape between the curves. Interestingly, the variance in neural trajectories did not correlate with the shape differences, suggesting that the separability of the neural trajectories is not caused by uncertainty. Relatedly, there appeared to be two distinct populations of trials, one with similar neural activity close to the decision boundary, and another with distinct neural activity far from the decision boundary. This raises the possibility that there are two forms of circuits, one which informs catch trials via minor changes, like a dynamical system, and another circuit which is vastly different in catch and miss trials.

However, it is important to acknowledge some limitations of our study. First, the sample sizes in certain analyses were small due to a limited number of recordings performed in those regions, limiting the generalisability of the findings. Second, our analysis focused on a subset of brain regions, and other regions not included in the analysis might exhibit different correlation-task performance relationships. Finally, a similar dataset to the AIS VBN dataset from the International Brain Laboratory (IBL)

could not be considered in the CCA and GPFA analyses due to the low number of units identified in the brain regions being considered. Perhaps the unit selection process could be more lenient to allow for less clean units to be included in the analysis, as population activity may be more important than individual units. Despite this, one of the most restrictive conditions was that the selected neurons had to be tracked for most of the task, which is important for the analysis of neural trajectories and this likely cannot be avoided. Further studies with larger sample sizes and broader brain coverage are warranted to confirm and expand upon our findings.

Overall, the observed differences in coherence, neural trajectories, and correlated activity between catch and miss trials shed light on task-related modulation of neural dynamics in the studied brain regions. Our results suggest that the relationship between task performance and the correlation in neural activity from multiple brain regions is complex and dependent on the specific task and brain regions investigated. This suggests the lack of a universal rule governing the correlation-task relationship and highlights the need for a task-specific and region-specific approach when investigating neural interactions. One possibility is that the correlation-task relationship is governed by a delicate balance between representation stability and information representation efficiency. This balance may be achieved by a dynamic equilibrium between correlated and uncorrelated activity, which may be modulated by task demands. Complementary to this, the structural connectivity between the regions considered may contribute to the correlation-task relationship. Another possibility is that there are intrinsic components of the neural activity which could be causing misinterpretation in this case, and would need to be separated to fully interpret the neural dynamics resulting from sensory inputs and inputs from other brain regions (Vahidi et al., 2023).

In conclusion, our study provides new insights into the intricate relationship between neural correlation and task performance. We observe task-dependent variations in coherence, neural trajectories, and correlated activity in different brain regions. These findings contribute to our understanding of how neural dynamics are shaped by task demands and provide a foundation for future investigations into the mechanisms underlying flexible behaviour and cognitive processes.

7.5 Data description and analysis methods

Three datasets were used to investigate this project. The first dataset was recorded in our lab, and involved tetrode recordings from the claustrum, retrosplenial cortex, and association cortex in rats performing an alternating fixed interval to fixed ratio task in an operant conditioning chamber. Due to the low number of cells identified in

our within lab claustrum recordings, local field potentials (LFPs) were analysed from these datasets. The other two datasets were publicly available releases from the Allen Institute for Brain Science (AIS) and International Brain Laboratory (IBL).

7.5.1 Operant conditioning with claustrum, RSC, and ACC recordings in rats

Recordings with 16 tetrodes were performed in 14 rats that were completing an alternating fixed interval to fixed ratio task regime in an operant conditioning chamber. The rats were first trained to perform the task for a reward on a regime coded in an assembly style language for MED-PC controlled boxes for operant conditioning without the headstage attached, and then in a BASIC style language for the headstage attached for use with the Axona recording system. Once rats learnt the task and the headstage was attached for recording, the alternating fixed interval to fixed ratio task was performed for 30 minutes, with 6 blocks of 5 minutes. There were three blocks of each task (FI or FR), and the order of the tasks were randomised. The data recorded from the first batch of 6 rats was used to ensure rats could learn the operant conditioning task, and to fine-tune the training regime. Additionally, this data was used to test the recording setup, isolate noise sources, and test video quality and video control. With the recording regime well established, the remaining 8 rats were used for the final data collection and are included in the LFP analyses (though one rat was excluded due to off target claustrum implantation).

Histology was not performed on the rats due to complications with COVID-19 shutting down the lab before the histology could be performed. Recording targets varied between rats, but implants were primarily targeted at the claustrum, retrosplenial cortex (RSC), and anterior cingulate cortex (ACC). As these three areas were the focus of the analysis, the number of tetrodes targeted to these areas per rat is listed in Table 7.3.

The captured electrophysiological data was processed from the raw electrical signal in Axona format (stored in a binary file) by first filtering into low-frequency LFPs and high-frequency multi-unit data, and then thresholding the multi-unit data to extract spikes using the BinConverter package (developed by Geoffrey Barrett in the Hussaini lab <https://github.com/seankmartin/BinConverter>). After filtering and thresholding, this data was processed with SIMURAN into NWB files for easy data sharing following the procedure in chapter 5. Multi-unit data and LFPs along with spatial information, operant chamber responses, and video of the rats in the operant conditioning chamber are available. A blinking red light is present in the video for synchronisation with the electrical signal.

Rat identifier	Clastrum tetrodes	ACC tetrodes	RSC tetrodes
CAR-SA1	8	4	4
CAR-SA2	2	4	4
CAR-SA3	2	2	4
CAR-SA4	4 (uncertain)	4	4
CAR-SA5	2	4	4
CAR-SA6	0	2	4
CAR-R2	7	1	1
CAR-R3	0	1	1

TABLE 7.3: Number of implantation targets in the claustrum, anterior cingulate cortex (ACC), and retrosplenial cortex (RSC). The number of tetrodes targeted at each brain region is shown for each rat. For CAR-SA4, the number of tetrodes implanted in the claustrum is uncertain because the tetrodes intended for the claustrum were angled such that they likely terminated in the dysgranular insular cortex, not the claustrum.

Estimation of trial behaviour as fixed interval or fixed ratio

For each trial performed in the claustrum dataset, the behaviour of the rat in that trial was classified as indicative of a fixed interval (FI) or fixed ratio (FR) task. First, an unsupervised classification was performed on the data to determine if this task was tenable, and that the data was well separable. The results indicated that fixed interval and fixed ratio trials were well separable, and that the data was well suited to this classification. The following criteria were used to classify the trials (in order of precedence):

1. If there were no lever presses during the trial, the trial was discarded.
2. If there were less than 6 lever presses during the trial, the trial was classified as a FI trial.
3. If the first lever press occurred more than 25 seconds after the start of the trial, the trial was classified as a FI trial.
4. If all lever presses during the first 25 seconds of the trial were at least 8 seconds apart of each other, the trial was classified as a FI trial.
5. If the first lever press that was less than 8 seconds after the last lever press was within 0.5 seconds of the end of the trial, the trial was classified as a FI trial.
6. If the average press rate of the lever was 1 press every 5 seconds, the trial was classified as a FR trial.

7. If there were more than 7 lever presses, the trial was classified as a FR trial.
8. In the rare situation that none of these criteria were met, the trial was classified as a FI trial.

Applying the criteria to classify the behaviour in the trial as FI or FR, the distribution of the lever press times remained similar between the two tasks but discouraged early lever presses in the FI task and late lever presses in the FR task (Figure 7.1 C, D). Furthermore, while statistics are reported for the estimated trial type, the results for the actual trial type (irrespective of behaviour) are similar.

7.5.2 Allen Institute for Brain Science and International Brain Laboratory datasets

The AIS data involved Neuropixels recordings from mice performing a change detection task with sets of natural images (Allen Institute for Brain Science, 2022). These Neuropixels recordings were primarily performed with the probes angled through the visual cortex and down into the limbic system. The IBL data involved Neuropixels recordings from mice that spun a wheel to move a single grating stimulus to the centre of the screen (Bonacchi et al., 2020). However, the visual stimulus strength changed over time, being barely visible at points, forcing the mouse to rely on experience to predict when the stimulus would appear and where it would appear on the screen.

To maintain consistency between these datasets, we applied the same unit selection criteria and the same brain regions were analysed. The brain regions analysed were, CA1, the thalamus, the subiculum, all layers of the primary visual cortex (V1), V2, and the ventral part of the LGN. In the AIS dataset, only trials where the stimulus changed were considered, and they were marked as correct, or catch trials, if the mouse correctly identified and responded to the stimulus change and incorrect, or miss trials, otherwise. In the IBL dataset, all trials were considered, and they were marked as correct if the mouse correctly identified and responded to the stimulus location. All units were selected from those recorded such that the following properties were satisfied:

1. The unit had a presence ratio of at least 90%, indicating that it was well tracked and isolated throughout the recording session. This ensures that the unit is present in the majority of trials.
2. The unit had a minimum firing rate of 0.5Hz, as units with a lower firing rate would contribute little to a GPFA or CCA analysis which both begin by binning spikes.

3. The unit had a maximum of 40% of its spikes as inter-spike interval (ISI) violations based on the refractory period of 2ms. An elevated level of ISI violations indicated contamination from multi-unit activity. However, this threshold was left quite high as some multi-unit activity is not expected to severely affect the results of the GPFA or CCA analysis.
4. The fraction of spikes below the spike detection threshold was less than 10%. Multiple spikes below the threshold indicated an incomplete unit was recorded.

Units from the International Brain Laboratory dataset were chosen with the additional conditions that the median amplitude was greater than 40 μ V and that the amplitude was not cutoff.

The resulting firing rate distributions (Figure 7.7) are in line with what would be expected from the literature, an approximately lognormal distribution due to the mix of cell types recorded (Roxin et al., 2011). Units chosen in this manner were not plentiful enough in the IBL dataset for the regions of interest to be analysed for the CCA and GPFA, frequently having less than 5 units in at least one of the brain regions of interest. As such, only the AIS dataset was used for the remainder of the analysis.

7.5.3 Coherence analysis of local field potentials

Coherence was considered between LFPs recorded from the claustrum, retrosplenial cortex, and association cortex in rats performing an alternating fixed interval to fixed ratio task in an operant conditioning chamber. Similarly to chapter 5, a single synthetic LFP was created for each brain region by normalising the signals, excluding noisy outliers, and then averaging LFPs from all remaining channels in the recording from that brain region. Coherence was considered over the entire FI or FR trials, the end of which was denoted by the collection of a reward. Magnitude squared coherence was calculated by dividing the squared cross-spectral density estimate by the individual power spectral density estimates:

$$C_{xy}(\omega) = \frac{|S_{xy}(\omega)|^2}{S_{xx}(\omega)S_{yy}(\omega)} \quad (7.1)$$

where S_{xy} is the cross-spectral density estimate, and S_{xx} and S_{yy} are the individual power spectral density estimates. The coherence was calculated in 2-second segments with 1-second overlaps, where a Hann window was used, and then averaged over all segments.

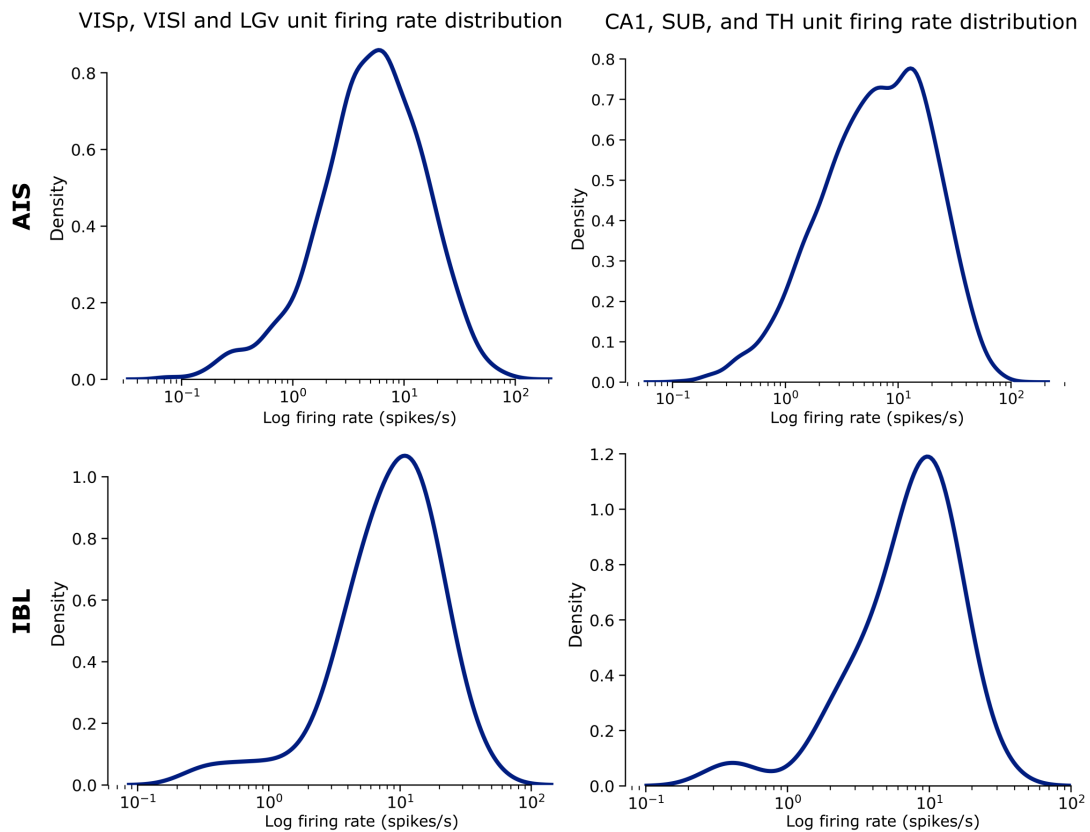


FIGURE 7.7: The log firing rate of the units selected from the Allen Institute’s visual Neuropixels dataset and the International Brain Laboratory’s brain wide survey. Rows separate the data into datasets, while columns separate the data into sets of brain regions.

Key: AIS – Allen Institute visual Neuropixels dataset, IBL – International Brain Laboratory brain wide survey, VISp – primary visual area (V1), VISl – lateral visual area, LGv – ventral part of the lateral geniculate nucleus, SUB – subiculum, TH – thalamus.

7.5.4 Neural trajectory analysis with Gaussian process factor analysis

The overall goal of the neural trajectory analysis was to compare the dynamics of the neural activity between correct and error trials. To do so, the neural activity was projected into a low dimensional latent space using gaussian process factor analysis (GPFA). The total shape of the latent neural trajectories and the Euclidean distance between the start and end points of the latent neural trajectories were compared between correct and error trials. This was to investigate if the neural trajectories were comparable to a dynamical system, whereby correct and incorrect trials would start in separated regions of the latent space but evolve along similarly shaped trajectories – suggesting a form of similarity in information processing but a difference in the initial conditions and the outcome.

GPFA is a method to extract low dimensional latent variables from noisy high dimensional spike train data (Yu et al., 2009). It combines a linear dimensionality reduction step via factor analysis with a non-linear smoothing step via Gaussian processes. In this section, GPFA was performed using the Elephant package (Denker et al., 2018) in Python. To ensure that the dynamics of the latent variables were not dominated by the neural responses in a single brain area, the same number of units were randomly selected from each brain area for the GPFA (to match the recorded area with the minimum number of units). Additionally, brain regions were considered in sets of three, and any recording in which less than 10 units were recorded in one of those brain area was excluded from the analysis. Furthermore, the GPFA was run on the data from each brain area separately and on the data from all brain areas together. The results of the GPFA were then compared between the two cases to ensure that the observed dynamics were reflecting cross-regional information. For spike count binning, a bin size of 20ms was used to bin the spike train observed in the one second interval following the trial start (the time at which the active stimulus was presented). A one second interval was chosen as it was the longest it took for the mouse to respond to the stimuli (mean response time was half a second), and stimuli were presented for 250ms. Post GPFA the latent neural trajectories were split into correct or error trials for comparison (where a correct trial involved correctly detecting a change in stimulus presentation and receiving a reward, while an incorrect trial involved missing a change in stimulus presentation).

As a prerequisite step for GPFA, spike counts were calculated from the observed spike trains in a series of nonoverlapping time bins and then square-rooted. Let us define this matrix of binned spike count observations as $Y \in \mathbb{R}^{N \times T}$, where N is the number of neurons and T is the number of time bins. To perform GPFA, we wish to find a low dimensional representation of the data (the neural states or latent variables),

$X \in \mathbb{R}^{D \times T}$, where D is the number of latent dimensions. To do so, we assume a linear-Gaussian relationship between the latent variables and the observations, such that

$$Y|X \sim CX + \mathbf{d} + \mathbb{N}(0, R) \quad (7.2)$$

where $C \in \mathbb{R}^{N \times D}$ is the transformation matrix which maps the neural trajectory into the space of recorded activity, $\mathbf{d} \in \mathbb{R}^N$ is the offset vector or bias, and $R \in \mathbb{R}^{N \times N}$ is the covariance matrix of the noise – which are the model parameters to be learned. The covariance matrix is constrained to be diagonal, such that $R_{ij} = 0$ for $i \neq j$, and the diagonal elements are the independent noise variances of each neuron. Additionally, each row of X is assumed to be a sample from a Gaussian process, which embodies the notion that the latent neural trajectories should be smooth and continuous. The Gaussian processes for each dimension of the latent space $i = 1 \dots D$ are distributed as $\mathbb{N}(0, K_i)$ where $K_i \in \mathbb{R}^{T \times T}$ following the squared exponential covariance function with characteristic timescale τ_i , signal variance $\sigma_{f,i}^2$, and noise variance $\sigma_{n,i}^2$:

$$K_i(t, t') = \sigma_{f,i}^2 \exp\left(-\frac{(t-t')^2}{2\tau_i^2}\right) + \sigma_{n,i}^2 \delta_{t,t'} \quad (7.3)$$

where δ is the Kronecker delta function (equals 1 if $t = t'$ and 0 otherwise). The model parameters ($C, \mathbf{d}, R, \tau_1, \dots, \tau_D, \sigma_{f,1}, \dots, \sigma_{f,D}, \sigma_{n,1}, \dots, \sigma_{n,D}$) are learned by maximising the log-likelihood of the data given the model via the expectation maximisation algorithm.

7.5.5 Procrustes analysis and distance measurement to compare neural trajectories

Procrustes analysis was performed to compare the shapes of the latent neural trajectories between correct and error trials. Procrustes analysis involves finding the optimal transformation between two sets of points (Green, 1952; Schönemann, 1966). In this case, the two sets of points are the latent neural trajectories when the mouse correctly identified a change in stimulus presentation and the latent neural trajectories when the mouse failed to respond to change in stimulus presentation. By constraining the solution to the orthogonal Procrustes problem, the transformation performed is a rotation, translation, and scaling of the latent neural trajectories. Let us define the latent neural trajectories for correct trials as $X_c \in \mathbb{R}^{D \times T}$ and the latent neural trajectories for error trials as $X_e \in \mathbb{R}^{D \times T}$, where D is the number of latent dimensions, and T the number of time bins. The orthogonal Procrustes problem is then to find the orthogonal matrix $R \in \mathbb{R}^{D \times D}$ which minimises the Frobenius norm of the difference between the two sets

of points:

$$\min_R \|X_c - RX_e\|_F^2, \quad \text{subject to } R^T R = I \quad (7.4)$$

This is equivalent to finding the singular value decomposition of the matrix product $X_c X_e^T = U \Sigma V^T$ and setting $R = UV^T$. This was performed in Python using the SciPy package (Virtanen et al., 2020).

Distance measurement between curves

To quantify the similarity between the latent neural trajectories for correct and error trials, the Euclidean distance between the start and end points of the latent neural trajectories was calculated. Additionally, the shape of the latent neural trajectories were compared using the discrete Fréchet distance. Fréchet distance (δ_F) is a measure of similarity between two curves that considers the location and ordering of the points along the curves (Eiter & Mannila, 1994; Fréchet, 1906). The intuitive idea behind the Fréchet distance is that it is the minimum length of a leash that would allow a dog and its owner to walk along their respective paths at the same time without backtracking. The Fréchet distance is defined as follows; let $f : [a, b] \rightarrow V$ and $g : [a', b'] \rightarrow V$ be two curves that map to V where (V, d) is a metric space.

$$\delta_F(X, Y) = \inf_{\alpha, \beta} \max_{t \in [0, 1]} d(f(\alpha(t)), g(\beta(t))) \quad (7.5)$$

where α and β are arbitrary continuous non-decreasing functions from $[0, 1]$ onto $[a, b]$ and $[a', b']$ respectively. The discrete Fréchet distance is a variant of the Fréchet distance which is defined for curves with a finite number of points, which is the case for the latent neural trajectories. This was calculated using the `similaritymeasures` Python package (Jekel et al., 2019).

7.5.6 Correlation between neural activity with canonical correlation analysis

Correlation between neural activity was measured by projecting the neural activity in each brain area into lower dimensional spaces using canonical correlation analysis (CCA) cross decomposition (Hotelling, 1936; H.-T. Wang et al., 2020), and then calculating the correlation between the projections. CCA cross decomposition can be used to project two matrices into lower dimensional spaces by identifying the sources of common variation in the high-dimensional matrices (H.-T. Wang et al., 2020). CCA is a linear transformation of two sets of variables such that the correlation between the two sets of variables is maximised (Hotelling, 1936). By stacking the binned spike counts from one brain area into a matrix $X \in \mathbb{R}^{N \times T}$ and the binned spike counts from another

brain area into a matrix $Y \in \mathbb{R}^{M \times T}$, where N and M are the number of neurons in each brain area respectively, and T is the number of time bins, we can find the linear combinations of the two sets of variables which have maximum correlation. Define the canonical variates U and V as follows:

$$U = \mathbf{a}^T X, \quad \mathbf{a} \in \mathbb{R}^N, \quad V = \mathbf{b}^T Y, \quad \mathbf{b} \in \mathbb{R}^M \quad (7.6)$$

Then CCA is defined as the following optimisation problem:

$$\max_{\mathbf{a}, \mathbf{b}} \text{corr}(U, V) = \max_{\mathbf{a}, \mathbf{b}} \frac{\text{cov}(U, V)}{\sqrt{\text{var}(U)}\sqrt{\text{var}(V)}} \quad (7.7)$$

Here \mathbf{a} and \mathbf{b} would be the first set of canonical vectors, and further sets of canonical vectors can be found by constraining the solution to be orthogonal to the previous canonical vectors. The order of X and Y is arbitrary, as CCA is symmetric with respect to the two matrices. This was performed in Python using the scikit-learn package (Pedregosa et al., 2018), and the data in X and Y were standardised by subtracting the mean and scaling by the standard deviation before performing CCA.

Binning spike trains for CCA

In our case, CCA was performed with two different methods of binning the spike trains. In the primary method, the spike trains were binned into 20ms bins and the binned spike counts were stacked into the matrices X_1, X_2, \dots, X_K and Y_1, Y_2, \dots, Y_K , where K is the total number of trials, thus forming matrices of binned spike counts for each trial. Then these matrices were stacked across trials to form the matrices $X_t \in \mathbb{R}^{N \times (T \cdot K)}$ and $Y_t \in \mathbb{R}^{M \times (T \cdot K)}$, where K is the number of catch and miss trials in the session. This leads to finding a consistent way to collapse neural activity (via a linear transformation with a and b) into a latent representation across trials, and then finding the correlation between the latent representations of the two brain areas. In the second method, each catch or miss trial of the change-detection task for a single session was considered as one row of the matrix X or Y , boiling down the problem to finding the correlation between the average firing rate of the neurons in each trial. The final matrices were $X \in \mathbb{R}^{(N \times K)}$ and $Y \in \mathbb{R}^{(M \times K)}$, where K is the number of catch and miss trials in the session. Overall, the second method provided little extra over the correlation between the average firing rate of all neurons in each brain area per trial (in other words, $a \approx (1/N, 1/N, \dots, 1/N)$, $b \approx (1/M, 1/M, \dots, 1/M)$). As such, the first method was preferred and reported in the results.

7.5.7 Statistical testing

For normally distributed data with unequal variances, Welch’s t-test was used to test for significant differences between two groups. In these cases, Cohen’s d was reported as a measure of effect size. Cohen’s d was calculated as:

$$d = \frac{\mu_1 - \mu_2}{\sqrt{\frac{(n_1-1)s_1^2 + (n_2-1)s_2^2}{n_1+n_2-2}}} \quad (7.8)$$

where μ_1 and μ_2 are the means of the two groups, s_1 and s_2 are the standard deviations of the two groups, and n_1 and n_2 are the number of samples in the two groups.

For non-normally distributed data, the Mann-Whitney U test was used to test for significant differences between two groups. In these cases, the effect size was reported as the rank-biserial correlation. This was calculated as:

$$r_B = f - u \quad (7.9)$$

Where, f is the proportion of pairs where the value in the first group is greater than the value in the second group, and u is the proportion of pairs where the value in the second group is greater than the value in the first group. This is the same as the difference in common language effect size computed for each group.

7.5.8 Data processing and software availability

All data processing and analysis was performed in Python. For the data from the claustrum, see this documentation (<https://bit.ly/cla-operant-data>) for full details on the data processing pipeline and the GitHub code (https://github.com/seankmartin/Claustrum_Experiment) for implementation. For the analysis of data from the AIS and IBL, the code is available at <https://github.com/seankmartin/task-related-neural-activity> to ensure reproducibility. Associated data can be downloaded via the AllenSDK, and OneAlyx. Convenience functions for downloading the data from the AllenSDK and OneAlyx are available in Simuran under the *loaders* module.

References

Allen Institute for Brain Science. (2022). Allen brain observatory: Visual behavior neuropixels technical whitepaper.

- Averbeck, B. B., Latham, P. E., & Pouget, A. (2006). Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, 7(5), 358–366.
- Barlow, H. (2001). Redundancy reduction revisited. *Network: Computation in Neural Systems*, 12(3), 241–253.
- Bonacchi, N., Chapuis, G., Churchland, A., Harris, K. D., Hunter, M., Rossant, C., Sasaki, M., Shen, S., Steinmetz, N. A., Walker, E. Y., Winter, O., & Wells, T. I. B., Miles Laboratory. (2020, February 23). *Data architecture for a large-scale neuroscience collaboration*.
- Bush, G., Vogt, B. A., Holmes, J., Dale, A. M., Greve, D., Jenike, M. A., & Rosen, B. R. (2002). Dorsal anterior cingulate cortex: A role in reward-based decision making. *Proceedings of the National Academy of Sciences*, 99(1), 523–528.
- Cayco-Gajic, N. A., & Zylberberg, J. (2021). Good decisions require more than information. *Nature Neuroscience*, 24(7), 903–904.
- Denker, M., Yegenoglu, A., & Grün, S. Collaborative HPC-enabled workflows on the HBP Collaboratory using the Elephant framework. In: *Neuroinformatics 2018*. 2018, P19.
- Dillingham, C. M., Jankowski, M. M., Chandra, R., Frost, B. E., & O'Mara, S. M. (2017). The claustrum: Considerations regarding its anatomy, functions and a programme for research. *Brain and Neuroscience Advances*, 1, 2398212817718962.
- Eiter, T., & Mannila, H. (1994). Computing discrete fréchet distance.
- Failor, S. W., Carandini, M., & Harris, K. D. (2021, May 23). *Learning orthogonalizes visual cortical population codes*.
- Franke, F., Fiscella, M., Sevelev, M., Roska, B., Hierlemann, A., & Azeredo da Silveira, R. (2016). Structures of neural correlation and how they favor coding. *Neuron*, 89(2), 409–422.
- Fréchet, M. M. (1906). Sur quelques points du calcul fonctionnel. *Rendiconti del Circolo Matematico di Palermo (1884-1940)*, 22(1), 1–72.
- Fusi, S., Miller, E. K., & Rigotti, M. (2016). Why neurons mix: High dimensionality for higher cognition. *Current Opinion in Neurobiology*, 37, 66–74.
- Green, B. F. (1952). The orthogonal approximation of an oblique structure in factor analysis. *Psychometrika*, 17(4), 429–440.
- Greicius, M. D., Krasnow, B., Reiss, A. L., & Menon, V. (2003). Functional connectivity in the resting brain: A network analysis of the default mode hypothesis. *Proceedings of the National Academy of Sciences of the United States of America*, 100(1), 253–258.
- Ham, G. X., & Augustine, G. J. (2022). Topologically Organized Networks in the Claustrum Reflect Functional Modularization. *Frontiers in Neuroanatomy*, 16.
- Hart, E. E., Blair, G. J., O'Dell, T. J., Blair, H. T., & Izquierdo, A. (2019). Anterior cingulate cortex activity regulates effort-based decision making. *bioRxiv*, 792069.

- Heilbronner, S. R., & Hayden, B. Y. (2016). Dorsal anterior cingulate cortex: A bottom-up view. *Annual Review of Neuroscience*, *39*(1), 149–170.
- Hotelling, H. (1936). Relations between two sets of variates. *Biometrika*, *28*(3), 321–377.
- Humphries, M. D. (2021, March 9). *The spike*.
- Ito, T., Brincat, S. L., Siegel, M., Mill, R. D., He, B. J., Miller, E. K., Rotstein, H. G., & Cole, M. W. (2019). Task-evoked activity quenches neural correlations and variability in large-scale brain systems. *bioRxiv*, 560730.
- Jekel, C. F., Venter, G., Venter, M. P., Stander, N., & Haftka, R. T. (2019). Similarity measures for identifying material parameters from hysteresis loops using inverse analysis. *International Journal of Material Forming*.
- Katz, L. N., Yu, G., Herman, J. P., & Krauzlis, R. J. (2021, September 17). *Correlated variability in primate superior colliculus depends on functional class*.
- Keshtkaran, M. R., Sedler, A. R., Chowdhury, R. H., Tandon, R., Basrai, D., Nguyen, S. L., Sohn, H., Jazayeri, M., Miller, L. E., & Pandarinath, C. (2022). A large-scale neural network training framework for generalized estimation of single-trial population dynamics. *Nature Methods*, *19*(12), 1572–1577.
- Kohn, A., Coen-Cagli, R., Kanitscheider, I., & Pouget, A. (2016). Correlations and neuronal population information. *Annual Review of Neuroscience*, *39*, 237–256.
- Kolling, N., Wittmann, M. K., Behrens, T. E. J., Boorman, E. D., Mars, R. B., & Rushworth, M. F. S. (2016). Value, search, persistence and model updating in anterior cingulate cortex. *Nature Neuroscience*, *19*(10), 1280–1285.
- Larry, N., & Joshua, M. (2021, April 30). *Neural correlations are attenuated in the transformation from the frontal cortex to movement* (preprint). Neuroscience.
- Lee, J. J., Krumin, M., Harris, K. D., & Carandini, M. (2020). Task specificity in mouse parietal cortex. *bioRxiv*, 2020.12.18.423543.
- Moreno-Bote, R., Beck, J., Kanitscheider, I., Pitkow, X., Latham, P., & Pouget, A. (2014). Information-limiting correlations. *Nature Neuroscience*, *17*(10), 1410–1417.
- Ni, A. M., Ruff, D. A., Alberts, J. J., Symmonds, J., & Cohen, M. R. (2018). Learning and attention reveal a general relationship between population activity and behavior. *Science (New York, N.Y.)*, *359*(6374), 463–465.
- Pandarinath, C., O’Shea, D. J., Collins, J., Jozefowicz, R., Stavisky, S. D., Kao, J. C., Trautmann, E. M., Kaufman, M. T., Ryu, S. I., Hochberg, L. R., Henderson, J. M., Shenoy, K. V., Abbott, L. F., & Sussillo, D. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods*, *15*(10), 805–815.
- Panzeri, S., Schultz, S. R., Treves, A., & Rolls, E. T. (1999). Correlations and the encoding of information in the nervous system. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, *266*(1423), 1001–1012.

- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Müller, A., Nothman, J., Louppe, G., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, É. (2018, June 5). Scikit-learn: Machine learning in python.
- Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., & Shulman, G. L. (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences*, *98*(2), 676–682.
- Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., & Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*, *497*(7451), 585–590.
- Rodriguez, A. C., Perich, M. G., Miller, L., & Humphries, M. D. (2023). Motor cortex latent dynamics encode arm movement direction and urgency independently. *bioRxiv*, 2023.05.26.542452.
- Roxin, A., Brunel, N., Hansel, D., Mongillo, G., & Vreeswijk, C. v. (2011). On the distribution of firing rates in networks of cortical neurons. *Journal of Neuroscience*, *31*(45), 16217–16226.
- Schönemann, P. H. (1966). A generalized solution of the orthogonal procrustes problem. *Psychometrika*, *31*(1), 1–10.
- Semedo, J. D., Jasper, A. I., Zandvakili, A., Krishna, A., Aschner, A., Machens, C. K., Kohn, A., & Yu, B. M. (2022). Feedforward and feedback interactions between visual cortical areas use different population activity patterns. *Nature Communications*, *13*(1), 1099.
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell system technical journal*, *27*(3), 379–423.
- Siegle, J. H., Jia, X., Durand, S., Gale, S., Bennett, C., Graddis, N., Heller, G., Ramirez, T. K., Choi, H., Luviano, J. A., Groblewski, P. A., Ahmed, R., Arkhipov, A., Bernard, A., Billeh, Y. N., Brown, D., Buice, M. A., Cain, N., Caldejon, S., . . . Koch, C. (2021). Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature*, *592*(7852), 86–92.
- Steinmetz, N. A., Zatzka-Haas, P., Carandini, M., & Harris, K. D. (2019). Distributed coding of choice, action and engagement across the mouse brain. *Nature*, *576*(7786), 266–273.
- Todd, T. P., & Bucci, D. J. (2015). Retrosplenial cortex and long-term memory: Molecules to behavior. *Neural Plasticity*, *2015*, 414173.
- Vahidi, P., Sani, O. G., & Shanechi, M. M. (2023, March 14). *Modeling and dissociation of intrinsic and input-driven neural population dynamics underlying behavior* (preprint). Neuroscience.
- Valente, M., Pica, G., Bondanelli, G., Moroni, M., Runyan, C. A., Morcos, A. S., Harvey, C. D., & Panzeri, S. (2021). Correlations enhance the behavioral readout of

- neural population activity in association cortex. *Nature Neuroscience*, *24*(7), 975–986.
- Vann, S. D., Aggleton, J. P., & Maguire, E. A. (2009). What does the retrosplenial cortex do? *Nature Reviews. Neuroscience*, *10*(11), 792–802.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., . . . SciPy 1.0 Contributors. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in python. *Nature Methods*, *17*, 261–272.
- Wang, H.-T., Smallwood, J., Mourao-Miranda, J., Xia, C. H., Satterthwaite, T. D., Bassett, D. S., & Bzdok, D. (2020). Finding the needle in a high-dimensional haystack: Canonical correlation analysis for neuroscientists. *NeuroImage*, *216*, 116745.
- Wang, Q., Ng, L., Harris, J. A., Feng, D., Li, Y., Royall, J. J., Oh, S. W., Bernard, A., Sunkin, S. M., Koch, C., & Zeng, H. (2017). Organization of the connections between claustrum and cortex in the mouse. *J Comp Neurol*, *525*(6), 1317–1346.
- White, M. G., Panicker, M., Mu, C., Carter, A. M., Roberts, B. M., Dharmasri, P. A., & Mathur, B. N. (2018). Anterior cingulate cortex input to the claustrum is required for top-down action control. *Cell Reports*, *22*(1), 84–95.
- Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. V., & Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Journal of Neurophysiology*, *102*(1), 614–635.
- Zatka-Haas, P., Steinmetz, N. A., Carandini, M., & Harris, K. D. (2021). Sensory coding and the causal impact of mouse cortex in a visual decision (M. Vinck, J. I. Gold, & K. Svoboda, Eds.). *eLife*, *10*, e63163.

8 Summary and future directions

8.1 Summary of our findings

In this thesis, we aimed to investigate the relationship between neural activity in multiple brain regions and how this changes with observed behaviours and experimental manipulations. Through analyses supported by our software package Simuran, we made several significant findings and contributions to the field.

Firstly, we provided a review of existing approaches to the analysis of *in vivo* neural activity recorded from multiple brain regions simultaneously, establishing the foundation for our study. Additionally, we developed the Simuran software package, which enables the analysis of diverse datasets from multiple laboratories and recording modalities. Using Simuran, we converted our in-lab recordings from proprietary data formats to the open-source Neurodata Without Borders (NWB) format, promoting data accessibility and interoperability in the field.

Regarding the specific research questions addressed in this thesis, we made important discoveries. We demonstrated that lesions of the anterior thalamic nuclei (ATN) abrogate spatial coding in the subiculum and reduce spatial alteration performance to chance level, without effects on spatial coding in CA1. This raises the possibility that there are multiple pathways for spatial information to reach the subiculum. A potential mechanism for this is via the suppression of bursting in the subiculum, which is known to be important for spatial coding.

Moreover, we found that permanent and temporary lesions of the ATN had insignificant effect on the oscillatory power of local field potentials (LFPs) in the subiculum, sharp wave ripple rate during sleep, or the relationship between speed and theta oscillatory power in subicular LFPs. However, lesions of the ATN had a significant effect on the coherence between subicular and retrosplenial cortex LFPs, and spike-LFP coupling in the subiculum within the theta band. This suggests that the behavioural and coding differences observed post-lesion may be related to cross-region communication, rather than reflecting primarily local changes in neural activity within the subiculum.

Additionally, given the link between the theta rhythm and spatial coding, the observed changes in spike-LFP coupling may be related to the observed changes in spatial coding.

Furthermore, we presented two novel methods for estimating the proportion of neurons in one brain region receiving structural chemical synaptic connections from another brain region, enhancing the interpretation of functional connectivity. Our results suggest that recording from densely connected regions with current high-yield technologies is likely to capture a considerable proportion of neurons that are directly connected. When connections along disynaptic or trisynaptic pathways are considered for densely anatomically connected brain regions, almost all neurons recorded in the downstream region are likely to receive a connection from a neuron recorded in an upstream brain region.

Finally, we found that the correlation between neural data in multiple brain regions is related to task performance in a region-specific manner. Analysing data from mice performing a change detection task revealed a task- and region- dependent amount of correlation between spiking activity from pairs of brain regions, whereby correlation between visual areas increased during catch trials, yet correlation between limbic system areas decreased or was unchanged. This demonstrates that the balance between information representation efficiency and the consistency in neural representation may be determined in a region and task specific manner. From the perspective of using LFPs, this may be a valid proxy for spiking activity in this context, as it also demonstrated region- and task-specific modulation in rats that were performing a switching fixed interval to fixed ratio task. Additionally, we found that the neural dynamics for catch and miss trials during the change detection task appear to split into two populations, with one population being more consistent with a dynamical system and having both trajectories close to the decision boundary. The other population was much further from the decision boundary, suggesting that this neural ensemble may be performing quite different computations during catch and miss trials.

The findings of this thesis contribute to our understanding of the neural circuitry underlying spatial navigation, the importance of brain rhythms, and the relationship between neural activity and task performance. Moreover, the development of the Simuran software package enhances the accessibility and interoperability of neural data analysis.

8.2 Future directions and impact

Overall, the ongoing development and improvement of open-source software packages for neural data analysis will facilitate further investigations in the field. With increased data sharing and collaboration, opportunities for large-scale analyses and meta-analyses will arise, enabling the discovery of new insights and the validation of existing findings. In addition, the potential to combine multiple neuroscience related technologies is exciting. For instance, combining connectomics, transcriptomics, and behavioural data could provide a more comprehensive understanding of the brain and its functions.

There are several promising avenues for future research that could take advantage of these efforts. For instance, understanding if there are general methods to analyse and describe neural activity, or if it is region- and task-dependent remains a key question. Existing studies suggest that a combination of general rules and region-specific knowledge may be necessary. This highlights the challenge of modelling these processes and raises the possibility that we can only expect to achieve a low-level or circuit-level understanding. The inference of source transformation rules or causal flow within these circuits could be beneficial in this regard (Camps-Valls et al., 2023; Gokcen et al., 2021).

Exploring the relationship between functional and structural connectivity and understanding source transformation rules in neural systems holds exciting potential for advancing our knowledge of neural circuitry and information processing. With such source transformation rules, we could probe how downstream observers respond to information from upstream sources and how that enables flexible behaviour. There is the potential for the use of expansion microscopy and electron microscopy to observe the structural connectivity of the brain in great detail, aligning with this goal (Kubota et al., 2018; Wassie et al., 2019). Additionally, voltage imaging could be used to observe this flow of information by resolving sub-threshold activity to estimate synaptic inputs, along with manipulation of the circuit via optogenetics (Fan et al., 2020).

Exploring the relationships between different brain signals, such as brain rhythms and spiking activity, and understanding how structural connectivity and network influence contribute to these processes are important areas of inquiry (Barabási et al., 2023; Buzsáki & Vöröslakos, 2023). In addition to structural connectivity, the geometry of the brain is impactful (Chen et al., 2023; Iyer et al., 2021; Pang et al., 2023). As such, the joint combination of structural connectivity and geometry may be necessary to understand the flow of information in the brain. By integrating various aspects of neural activity and connectivity, we can gain a more comprehensive understanding of brain function.

Moreover, this research into multi-region neural recording analysis has the potential to make a real-world impact. For example, in cases of drug-resistant epilepsy, implanting a deep brain stimulator in the ATN has shown positive outcomes (Zangiabadi et al., 2019). However, manually adjusting the stimulation strength is a labour-intensive process for patients and medical practitioners. By leveraging the information transfer process from the ATN to its connected regions, it may be possible to create a closed-loop system where a device monitors ATN signals and automatically adjusts the stimulator strength, eliminating the need for manual intervention.

Notably, machine learning algorithms, such as recurrent neural networks (RNNs), have proven effective for brain machine interfaces (BMIs) by learning patterns and enabling efficient control. This raises the question of whether there are more structured and transferable approaches to improve BMIs. If we had a comprehensive understanding of source transformation rules and brain region connections, could we enhance BMIs? Or will machine learning always be necessary to learn these patterns? It is possible that the rules could be used to seed machine learning algorithms, reducing training time or aiding in result interpretation.

In summary, future research should focus on refining data analysis tools, investigating functional and structural connectivity, and exploring the applicability of source transformation rules. Additionally, understanding neural ensembles, general methods for describing neural activity, and the relationships between different modalities are crucial for advancing our knowledge of the brain and its functions.

8.3 Conclusion

In conclusion, this thesis provides valuable insights into the intricate relationship between neural activity in multiple brain regions and how that changes to enable flexible behaviour. The findings and methodologies presented here contribute to the broader field of neuroscience and lay the groundwork for further advances in our understanding of brain function and its implications in health and disease.

8.4 Closing remarks

I am truly grateful to my supervisor, colleagues, and all those who supported me throughout this journey. Their guidance, expertise, and feedback have been instrumental in shaping this work.

I would also like to acknowledge the funding from the School of Psychology in Trinity

College Dublin to support this thesis, and the wider support from the Wellcome Trust and Science Foundation Ireland to my colleagues.

Overall, I am optimistic about the future of neuroscience and the potential it holds for improving human health and well-being. Undoubtedly, we face significant challenges, but there are amazing people working in the field and fascinating ideas surfacing regularly. Hopefully, this body of work provides helpful insight into this incredible research community, and truly:

We are in one of the golden ages of neuroscience. Never before have so many people thought so hard about the brain – it's a fantastic time to be alive.

Christof Koch, Neuromatch conference 4.0.

References

- Barabási, D. L., Bianconi, G., Bullmore, E., Burgess, M., Chung, S., Eliassi-Rad, T., George, D., Kovács, I. A., Makse, H., Papadimitriou, C., Nichols, T. E., Sporns, O., Stachenfeld, K., Toroczkai, Z., Towlson, E. K., Zador, A. M., Zeng, H., Barabási, A.-L., Bernard, A., & Buzsáki, G. (2023). Neuroscience needs network science.
- Buzsáki, G., & Vöröslakos, M. (2023). Brain rhythms have come of age. *Neuron*, *111*(7), 922–926.
- Camps-Valls, G., Gerhardus, A., Ninad, U., Varando, G., Martius, G., Balaguer-Ballester, E., Vinuesa, R., Diaz, E., Zanna, L., & Runge, J. (2023). Discovering causal relations and equations from data [arXiv.org].
- Chen, L., Lin, X., Ye, Q., Nenadic, Z., Holmes, T. C., Nitz, D. A., & Xu, X. (2023). Anatomical organization of temporally correlated neural calcium activity in the hippocampal CA1 region. *iScience*, *26*(5), 106703.
- Fan, L. Z., Kheifets, S., Böhm, U. L., Wu, H., Piatkevich, K. D., Xie, M. E., Parot, V., Ha, Y., Evans, K. E., Boyden, E. S., Takesian, A. E., & Cohen, A. E. (2020). All-optical electrophysiology reveals the role of lateral inhibition in sensory processing in cortical layer 1. *Cell*, *180*(3), 521–535.e18.
- Gokcen, E., Jasper, A. I., Semedo, J. D., Zandvakili, A., Kohn, A., Machens, C. K., & Yu, B. M. (2021, September 1). *Disentangling the flow of signals between populations of neurons* (preprint). Neuroscience.
- Iyer, R., Siegle, J. H., Mahalingam, G., Olsen, S., & Mihalas, S. (2021, June 23). *Geometry of inter-areal interactions in mouse visual cortex*. bioRxiv.
- Kubota, Y., Sohn, J., & Kawaguchi, Y. (2018). Large volume electron microscopy and neural microcircuit analysis. *Frontiers in Neural Circuits*, *12*.
- Pang, J. C., Aquino, K. M., Oldehinkel, M., Robinson, P. A., Fulcher, B. D., Breakspear, M., & Fornito, A. (2023). Geometric constraints on human brain function. *Nature*, *618*(7965), 566–574.
- Wassie, A. T., Zhao, Y., & Boyden, E. S. (2019). Expansion microscopy: Principles and uses in biological research. *Nature Methods*, *16*(1), 33–41.
- Zangiabadi, N., Ladino, L. D., Sina, F., Orozco-Hernández, J. P., Carter, A., & Téllez-Zenteno, J. F. (2019). Deep brain stimulation and drug-resistant epilepsy: A review of the literature. *Frontiers in Neurology*, *10*.

A Appendix

The appendix describes the technical details of the wet lab work performed by my colleagues that supports the work in this thesis. It is included here for completeness, and to understand the full context of the work presented in this thesis. However, my own contributions to this thesis are strictly computational, theoretical, or hardware-related, so I did not perform any of the wet lab work described here. To help understand the process and potential data capture issues, I have attended multiple surgeries where craniotomies and electrode implantations were performed, but I have not performed any surgeries myself or handled the animals.

A.1 Further details of the anterior thalamic lesion study

A.1.1 Behavioural effect of ATN lesions

Consistent with previous lesion studies (Aggleton & Nelson, 2015) ATN lesion-induced dysfunction (ATNx) rats had considerably lower spatial alternation scores ($52.23 \pm 14.58\%$), compared with normal ($83.10 \pm 10.96\%$) and sham ($81.25 \pm 11.41\%$) controls (ANOVA, $F_{(2,114)} = 77.8$, $p < 0.001$; Tukey *post hoc* $p < 0.001$; Figure A.1B), confirming their severely impaired spatial working memory. In contrast, ATNx rats showed no impairment in novel object discrimination compared with the control rats (cumulative D1 $t_{(142.69)} = 0.99$, $p = 0.324$; updated D2 $t_{(172.25)} = 1.42$, $p = 0.159$; Welch's two-sample t test), indicating that recognition memory under these conditions remains intact (Figure A.1D-F). During free exploration in the square arenas with electrophysiological recordings, the ATNx rats travelled greater distances than the control rats both during habituation (control $85.45 \pm 14.21\text{m}$, ATNx $162.14 \pm 17.43\text{m}$; $t_{(16.14)} = 3.55$, $p = 0.003$; Welch's two-sample t test) and during subsequent recordings (control $120.66 \pm 15.19\text{m}$, ATNx $142.41 \pm 10.89\text{m}$; $t_{(14.55)} = 2.98$, $p = 0.096$), suggesting slightly increased motor activity after anterior thalamic nuclei (ATN) lesion.

As for ATNx rats with CA1 recordings, those rats also showed a deficit in spatial alternation on the T-maze, equivalent to ATNx (SUB) rats (control 82.38 ± 11.18 ,

52.60 ± 11.40 , $t_{(40.93)} = -10.71$, $p < 0.001$; Figure 4.7C), confirming an impairment in spatial working memory. These rats also showed no deficit in object recognition and no difference in overall exploration time, compared with control ($p > 0.05$; Figure 4.7D1-D3).

A.1.2 Lesion effectiveness

Rats received permanent neurotoxic lesions of the ATN (termed ATNx) following injections of N-Methyl-D-aspartic acid (NMDA) or transient ATN lesions following infusion of muscimol. Other rats served as controls. Permanent neurotoxic lesion effectiveness was quantified by comparing the total anti-NeuN-reacted cell counts in 40 μ m thick brain sections containing the anteromedial, anterodorsal, and anteroventral thalamic nuclei in ATNx rats with the corresponding control values (Figure A.2A, B). The ATNx rats had markedly reduced cell counts (control $16,209 \pm 2507$, ATNx 3497 ± 1528 ; $t_{(8.26)} = -9.68$, $p < 0.001$, Welch's two-sample t test; Figure A.2E), while the calbindin-reacted sections helped to confirm that nucleus reuniens remained intact (Figure A.2C, D). The sparing of nucleus reuniens indicates that, despite a significant role in spatial working memory (Griffin, 2015), nucleus reuniens damage was not a contributing factor to differences in either behaviour or cell firing properties in ATNx rats.

A.1.3 Subject details and experimental design

Experiments were conducted on 28 male Lister Hooded rats with preprocedural weights of between 306g and 356g. Upon arrival, rats were cohoused on a 12-hour day/night cycle and handled daily by the experimenter for a week before surgical procedure. Before surgery and during recovery, rats had free access to food and water; during behavioural testing, food was restricted but ensured the rats did not fall below 85% of the animal's free feeding weight. All rats were naive before the present study. Selection of rats between lesion and control groups was alternated according to body weight before surgery so that preprocedural weights were matched, and the groups balanced.

These rats were implanted with 32 electrodes unilaterally, with electrodes targeting a combination of the subiculum, retrosplenial cortex, and CA1. Some rats were lesioned permanently in the ATN, temporarily inactivated in the ATN via muscimol, or kept as sham or normal controls. Rats were then recorded during free-moving exploration, T-maze spatial alternation, and novel object interaction in a bow-tie maze. Table A.1 shows the number of rats in each type of recording layout. One rat was excluded due to a blockage in one canula, resulting in unilateral muscimol infusion as opposed

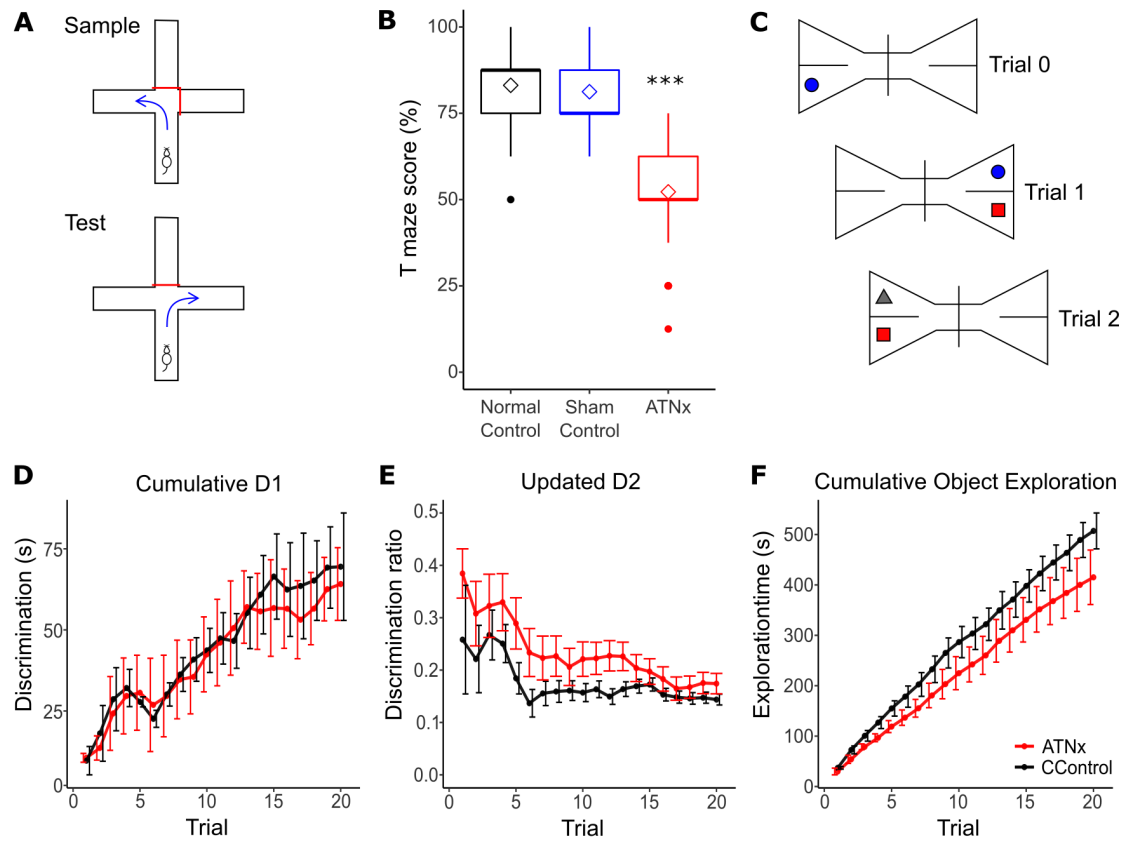


FIGURE A.1: Behavioural impacts of ATN lesions on spatial alternation memory performance and novel object recognition.

(A) Schematic diagram of the spatial alternation T-maze task, where the rat is forced to go one direction in the first trial, and then should go the opposite direction in the second trial.

(B) ATN-lesioned animals (ATNx) showed a significant deficit in spatial alternation ability compared with both control and sham animals.

(C) Schematic of the novel object recognition task, where the rat is shown two objects in one trial, and then in the next trial one of the objects is changed. The rat is expected to investigate the novel object.

(D, E, F) In novel object exploration, there was no difference between control and lesion rats in cumulative D1 (D), updated D2 (E), or total exploration time (F).

Key: *** indicates $p < 0.001$, ANOVA with Tukey *post hoc*. Error bars indicate standard error of the mean (SEM).

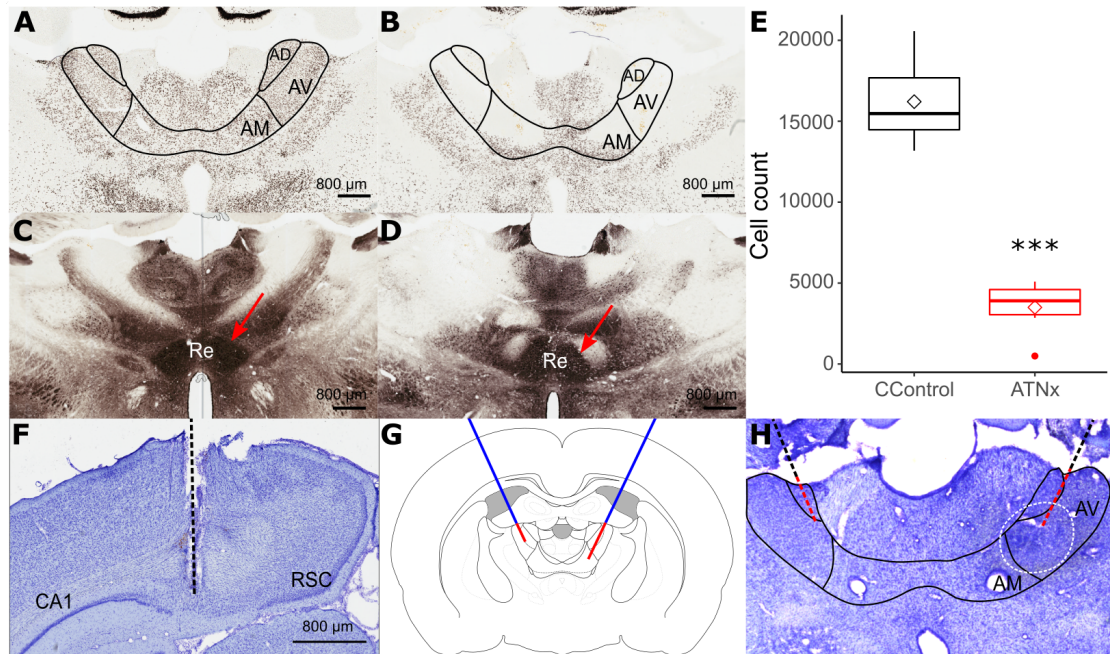


FIGURE A.2: Demonstrating ATN lesion effectiveness from permanent (NMDA) lesions and temporary inactivation via muscimol.

(A, B) NeuN-reacted coronal sections showing the status of the AV (anteroventral), AM (anteromedial), and AD (anterodorsal) thalamic nuclei for control (A) animals and NMDA lesioned rats (B).

(C, D) Calbindin-reacted coronal section showing that the nucleus reuniens (Re – marked with an arrow) was intact in both control (C) and lesion (D) rats.

(E) The ATN cell count was significantly reduced in lesioned animals (ATNx) compared with controls. *** indicates $p < 0.001$, Welch's two-sample t-test.

(F) Nissl-stained coronal sections helped to confirm electrode placement in dorsal subiculum, with the electrode path indicated.

(G) Schematic representing cannula placement (blue) and the two ATN targets of the infusion needle (red).

(H) Cresyl violet-stained section indicating cannula placement, with diaminobenzidine (DAB)-reacted Fluorogold infused to indicate the spread of muscimol.

to bilateral muscimol infusion. The following information in the methods describes processing that is specific to this chapter as opposed to the information in chapter 4 involving these same rats.

A.1.4 Detailed results from temporary inactivation via muscimol

When the ATN was temporarily inactivated with muscimol, spatial alternation percentage dropped to chance levels (before muscimol infusion $87.27 \pm 7.41\%$; after muscimol infusion $50.00 \pm 13.97\%$; Figure A.3B). Two rats also received bilateral saline infusions as a control, showing no deficit in spatial alternation with an 87.5% success rate. The muscimol inactivation caused a significant deficit in spatial alternation compared with both before muscimol and saline infusion (ANOVA, $F_{(3,62)} = 49.69$, $p < 0.001$, Tukey *post hoc* $p < 0.001$).

Before the muscimol infusion sessions, electrophysiological recordings were conducted on cannulated rats to allow for electrode adjustment and habituation to recording equipment. Of 35 units recorded during inactivation experiments, 29 were recorded at baseline (i.e., immediately before muscimol infusion) and so included in the study (see examples in Figure A.3, boxes A1-A3, C1-C3). Of the 29 units recorded at baseline, only one unit was recorded before infusion and not afterwards. For analysis, unless otherwise stated, Welch's two-sample t-tests were performed to compare unit properties with baseline data, with Bonferroni post hoc correction to account for multiple comparisons of time points to baseline. When all cells were grouped, there was no significant change in mean spiking frequency following ATN inactivation (baseline 5.07 ± 9.10 spikes/s, 40-60min after infusion 2.90 ± 3.84 spikes/s; $t_{(45.21)} = 0.21$, $p = 0.230$). Mean spike width (baseline 180.32 ± 55.20 ms, 40-60min after infusion 185.81 ± 50.71 ms, $t_{(55.53)} = 0.39$, $p = 0.698$) and spike amplitude (baseline 113.01 ± 41.43 mV, 40-60min after infusion 115.22 ± 45.48 mV, $t_{(51.13)} = 0.19$, $p = 0.851$) showed no change following ATN inactivation when all cells were grouped. A histogram of spike widths showed that, while cells with narrow waveforms were more often spatial, the spatial and non-spatial groups remained mixed (Figure 4.5B).

Following muscimol infusion, spatial properties of subicular cells declined, despite no decrease in firing frequency. For designated place cells before infusion, the place field became disrupted after ATN inactivation (Figure A.3A1). Head directionality also became disrupted without ATN input (Figure A.3, box A2), and grid cells did not fire in a grid-like pattern (Figure A.3, box A3). In most cases, these spatial properties were recovered by the following day, although in rare cases, the cells could no longer be recorded (Figure A.3, box C3). Interestingly, the subiculum grid cells appeared to

Experimental condition	Electrode positions
Normal control ($n = 3$)	7 tetrodes targeting dorsal subiculum; bipolar electrode targeting RSC and subiculum
Sham control ($n = 3$)	7 tetrodes targeting dorsal subiculum; bipolar electrode targeting RSC and subiculum
ATNx ($n = 5$)	7 tetrodes targeting dorsal subiculum; bipolar electrode targeting RSC and subiculum
ATN muscimol ($n = 6$)	<ul style="list-style-type: none"> • 7 tetrodes targeting dorsal subiculum; bipolar electrode targeting RSC ($n = 2$) • 7 tetrodes targeting dorsal subiculum; bipolar electrode targeting CA1 ($n = 1$) • 4 tetrodes targeting dorsal subiculum; 3 tetrodes targeting CA1; 1 tetrode targeting RSC ($n = 2$) • 4 tetrodes targeting dorsal subiculum; 3 tetrodes targeting CA1; bipolar electrode targeting RSC ($n = 1$)
ATN muscimol control ($n = 5$)	<ul style="list-style-type: none"> • 4 tetrodes targeting dorsal subiculum; 3 tetrodes targeting CA1; 1 tetrode targeting RSC ($n = 3$) • 7 tetrodes targeting CA1; bipolar electrode targeting RSC ($n = 2$)
Control RSC ($n = 2$)	7 tetrodes targeting RSC; bipolar electrode targeting subiculum and RSC
ATNx RSC ($n = 1$)	7 tetrodes targeting RSC; bipolar electrode targeting subiculum and RSC
ATNx CA1 ($n = 3$)	8 tetrodes targeting CA1

TABLE A.1: Types of recording layouts and the brain regions implanted in along with the number of rats with each type of recording layout in the ATN lesion study. Sham indicates a phosphate-buffered saline (PBS) lesion. ATNx indicates Permanent bilateral cytotoxic lesion via NMDA targeting ATN. Bilateral cannulae were targeting the ATN for temporary inactivation via muscimol for the muscimol experiments. For these experiments, muscimol control had the cannulae implanted but no muscimol was infused. RSC indicates the retrosplenial cortex. CA1 indicates the CA1 subfield of the hippocampus. ATN indicates the anterior thalamic nuclei.

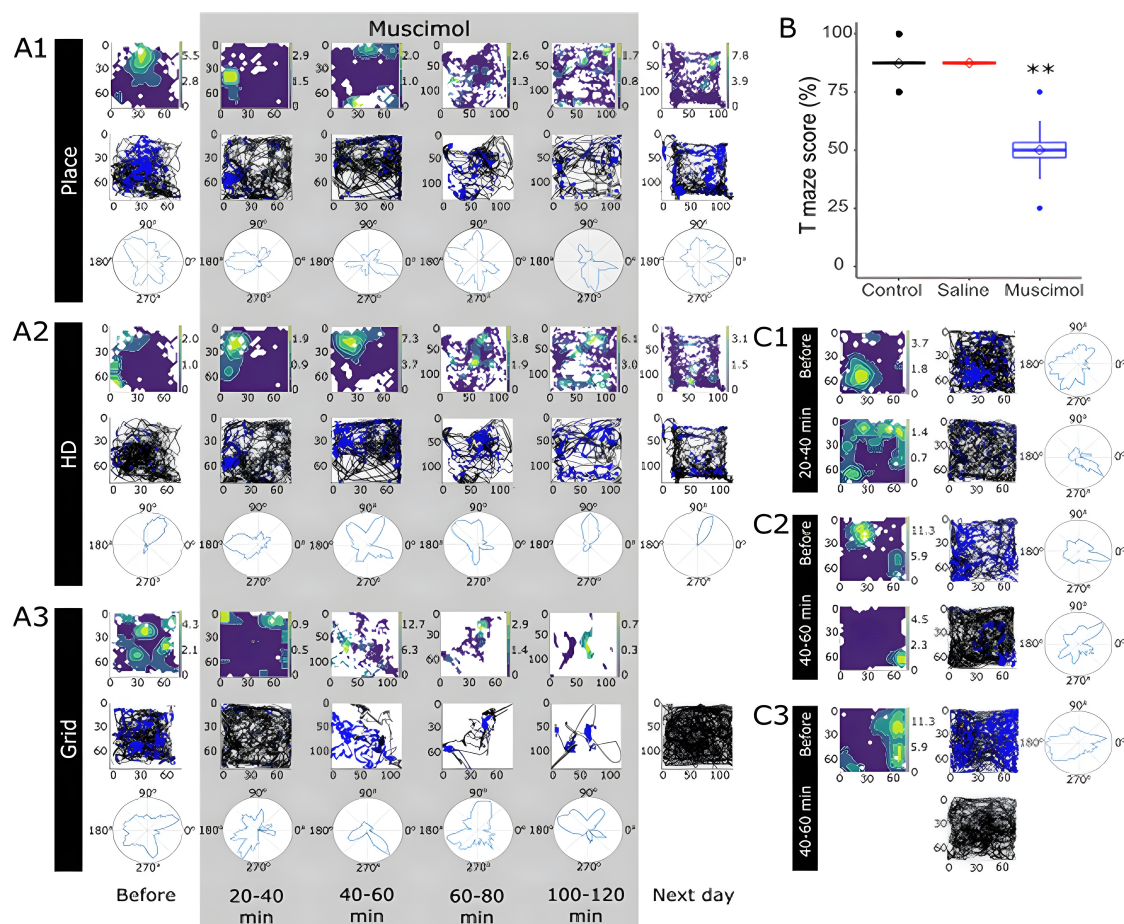


FIGURE A.3: Effect of muscimol infusion on unit activity over time and T-maze performance.

(A1-A3) Examples of spatial units before, during, and after muscimol infusion. Spatial properties of single units decreased when ATN was inactivated. In (A3), the rat was relatively inactive after 100-120 min, and the cell was not recorded the next day.

(B) Spatial alternation dropped to chance levels when the ATN were temporarily inactivated with muscimol, compared with the same animals before infusion. When saline was infused in place of muscimol, no deficit was present.

(C1-C3) Further examples of spatial units before and after ATN inactivation via muscimol. (C1, C2) show disruption of place field shortly after muscimol infusion, while no firing was detected after muscimol infusion in (C3).

Key: ** indicates $p < 0.01$, ANOVA with Tukey post hoc.

lose their grid field initially but retained some head directionality, before this too was disrupted (Figure A.3, boxes A3, C1, C2). However, muscimol is a powerful substance, and as a result, the rats had reduced levels of movement around the time of most heightened muscimol inactivation of the ATN.

Of the 29 units included in the study, at baseline, 20 were classified as bursting, 6 as fast-spiking, and 2 as theta-modulated (before muscimol infusion). For unit property analysis, theta-modulated units were excluded because of their high firing rate and insufficient numbers to perform further statistical analysis. When fast-spiking and bursting units were combined, there was no difference in firing frequency compared with baseline, following ATN inactivation with muscimol (5.07 ± 9.10 spikes/s, 0-20min after infusion 3.20 ± 4.75 spikes/s; $t_{(52.31)} = 0.87$, $p = 0.385$; Figure A.4A). All rats travelled less distance (i.e., showed less activity) as time increased across the experiment (baseline, 115.97 ± 33.12 m; 60-80 min, 58.17 ± 37.41 m; Figure A.4B), although there was no correlation between firing frequency and distance travelled ($r^2 = 0.011$; Figure A.4C). Figure A.4D illustrates the firing frequency of each cell at baseline, then its frequency in 5min intervals following muscimol infusion and again the following day, in those instances where the cell was recorded again. Bursting units did not show any changes in burst properties following inactivation of the ATN, including the number of bursts, spikes per burst, and burst duration ($p > 0.05$). Cells that had been designated as spatial units before infusion showed no changes in firing frequency, indicating that these spatial cells continue to fire following ATN inactivation but no longer displayed spatial properties. Spatial cells displayed a visual decrease in the number of bursts in recordings after muscimol infusion (Figure A.4), though this was not significant.

A.1.5 ATN lesions with CA1 recordings

In rats with only CA1 recordings, ATN lesions were also quantified by comparing anti-NeuN reacted cell counts in the ATN with those from control. In all ATNx (CA1) rats, the surgery consistently produced a marked cell loss throughout almost the entire ATN (see Figure 4.7A, B). An independent-sample t-test showed a significant difference in anti-NeuN cell counts between (4166 ± 1003) and control ($16,209 \pm 2507$, $t_{(7)} = 9.08$, $p < 0.001$) rats (see Figure A.5 for lesion extent). As expected, these ATNx (CA1) rats showed a deficit in spatial alternation on the T-maze, equivalent to ATNx (SUB) rats (control 82.38 ± 11.18 , 52.60 ± 11.40 , $t_{(40.93)} = -10.71$, $p < 0.001$; Figure 4.7C), confirming an impairment in spatial working memory. These rats also showed no deficit in object recognition and no difference in overall exploration time, compared with control ($p > 0.05$; Figure 4.7D1-D3). Unlike subicular recordings with ATN lesions, spatial cells were still present in CA1 recordings with ATN lesions (Figure 4.7E). Full

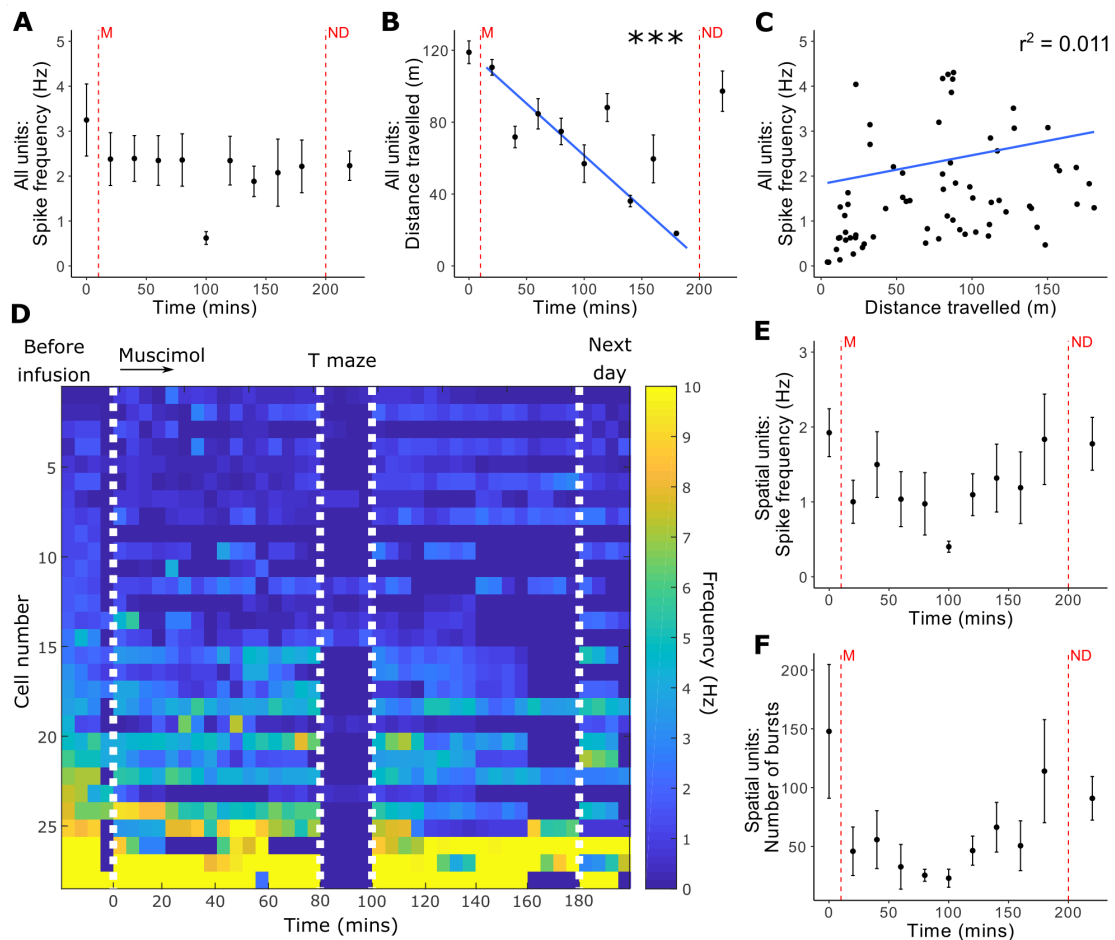


FIGURE A.4: Spike properties following temporary inactivation of ATN with muscimol.

(A) Inactivation of ATN caused no significant decrease in single-unit firing frequency.

(B) Animals showed decreasing levels of activity throughout the experiment.

(C) Despite a decrease in activity, there was no significant correlation between distance travelled and spike frequency.

(D) Firing frequency of each cell recorded at baseline (left), in 5 min bins throughout the experiment. The first white line indicates ATN inactivation with muscimol, after 15-20 min of baseline recording before infusion. In most cases, electrophysiological recording was paused for T-maze testing between 80 and 100 min (second and third white lines), then continued. Recordings in which the animal was largely inactive or asleep were excluded. The final white line indicates data from the day after infusion.

(E, F) There were no significant changes in spike firing in spatial units as a result of ATN inactivation and burst properties remained consistent, including the number of bursts.

Key: (A, B, E, F), the first red vertical line (M) indicates the infusion of muscimol, while the second red vertical line (ND) indicates the recording taken the next day. Each bin represents 20 min of recording. Data are compared with baseline, immediately before inactivation. Error bars on the plots indicate SEM. Theta-entrained cells are removed from (A, B, C) due to high firing frequency compared with other cell classes. *** indicates $p < 0.001$, Welch's two-sample t-test with Bonferroni correction.

Measurement	Place cell	Interneuron
Spike rate (spikes/s)	0.92 ± 0.68	19.71 ± 16.83
Spike amplitude (μV)	134.43 ± 41.38	134.55 ± 42.58
Spike height (μV)	213.64 ± 59.32	217.45 ± 69.76
Spike width (μs)	231.99 ± 29.34	176.29 ± 35.94
Inter-spike interval (s)	1.54 ± 0.84	0.15 ± 0.16
Mean bursting inter-spike interval (ISI) (ms)	4.30 ± 0.34	4.30 ± 0.33
Spikes per burst	2.15 ± 0.11	2.15 ± 0.15
Mean burst duration (ms)	5.96 ± 0.47	5.98 ± 0.71
Inter-burst interval (second)	15.89 ± 13.70	11.69 ± 21.24
Propensity to burst	0.28 ± 0.15	0.125 ± 0.11
Spatial coherence	0.72 ± 0.10	0.45 ± 0.20
Spatial information content	2.12 ± 0.58	0.54 ± 0.29
Spatial sparsity	0.18 ± 0.07	0.61 ± 0.14
Peak spike rate in $3 \times 3\text{cm}$ bins (spikes/s)	18.26 ± 10.63	38.5 ± 23.00

TABLE A.2: Summary of place and interneuron properties of CA1 units from control and ATN lesioned rats over three rats. Values indicate mean plus/minus standard deviation. Propensity to burst represents bursting spikes divided by total spikes.




FIGURE A.5: Depicting the maximum (A) and minimum (B) extent of NMDA lesions in ATN lesioned rats with recordings in CA1 (ATNx-CA1), shown by grey shading superimposed over the image.

Key: AV (anteroventral), AM (anteromedial), and AD (anterodorsal) thalamic nuclei.

place cell and interneuron properties are shown in Table A.2.

A.1.6 Cells in CA1 and possible recording from fibers of passage

Notably, the proportions of CA1 place cells in ATN lesioned rats are in line with our previous experiences in normal rats (Hok et al., 2012; Wang et al., 2012) and that of others (Broussard et al., 2020; Duvelle et al., 2019; Meshulam et al., 2017; Thompson & Best, 1989), contrasting with the more disruptive effects of medial entorhinal cortex lesions on CA1 place cell activity (Hales et al., 2014).

Consistent with previous reports (Brotons-Mas et al., 2017; Brotons-Mas et al.,

2010; Sharp & Green, 1994), we found spatial cells (including grid-like cells) in the dorsal subiculum of control rats. Although grid cells have been widely studied in the hippocampal formation since their discovery (Hafting et al., 2005), the presence of grid cells in the dorsal subiculum is still a relatively new finding (Brotons-Mas et al., 2017). Grid cells have been reported in the presubiculum and parasubiculum (Boccaro et al., 2010), both parahippocampal areas receiving anterior thalamic inputs (van Groen & Wyss, 1990). Anterior thalamic lesions disrupt grid cell (and head-direction) activity in these same parahippocampal areas (Goodridge & Taube, 1997; Winter et al., 2015). Of the ATN, the anteroventral nucleus possesses theta-modulated head-direction cells (Tsanov et al., 2011), and the anteromedial nucleus contains place cells and perimeter/boundary cells (Jankowski et al., 2014; Matulewicz et al., 2019). The grid-like signal found in the presubiculum and subiculum might have thalamic (Goodridge & Taube, 1997) and entorhinal (Hafting et al., 2005) components.

Given the routes of ATN fibers reaching the medial temporal lobe, it is worth considering whether we have recorded from fibers of passage. *In vitro*, “direct recording of single AP (action potential) transmission is challenging” because of the small diameters of axons and recording instability (Radivojevic et al., 2017). However, *in vivo*, during hippocampal and alvear recordings, the presence of short-duration triphasic waveforms having a peak-trough length of 179 μ s have been reported, recorded mostly on only one wire of a tetrode (Robbins et al., 2013). On occasion, we might have recorded from such fibers (e.g., Figure 4.2A, top grid cell; but compare Brotons-Mas et al. (2017); Figures 3B, 4C). The possibility of recordings of fibers of passage in the subiculum needs further investigation. A different concern relates to the disruption of thalamic fibers of passage. Here, muscimol helped to confirm the NMDA lesions targeted ATN neurons, and not fibers of passage (Winter et al., 2015).

A.1.7 Ethics

All experimental procedures were in accordance with the ethical, welfare, legal, and other requirements of the Healthy Products Regulatory Authority regulations and were compliant with the Health Products Regulatory Authority (Irish Medicines Board Acts, 1995 and 2006) and European Union directives on Animal Experimentation (86/609/EEC and Part 8 of the EU Regulations 2012, SI 543). All experimental procedures were approved by the Comparative Medicine/Bioresources Ethics Committee, Trinity College Dublin, Ireland before conduct, and were conducted following the LAST Ireland and international guidelines of good practice.

A.1.8 Surgical methods

Rats were first anaesthetised with isoflurane (5% to induce anaesthesia, 1% – 2% to maintain) combined with oxygen at a rate of 2L/min. After being placed in a stereotaxic frame, chloramphenicol 0.5% eye gel, preoperative antibiotics (Enrocare, 0.1mL in 0.5mL saline) and analgesia (Metacam, 0.1mL) were administered. The skull was then exposed, and connective tissue was removed.

For the permanent ATN lesion cohort (ATN_x, $n = 7$), bilateral neurotoxic lesions targeting the ATN were performed using slow infusions of 0.12mol/L NMDA dissolved in PBS solution (pH 7.35). NMDA was infused over 5 minutes (0.22 or 0.24mL per site) via a 0.5mL Hamilton syringe (25 gauge), with the syringe left in position for a further 5 minutes at each of the four target sites before slow retraction. The craniotomies were then sealed using bone wax. The ATN lesion coordinates, with the skull flat, were as follows from bregma: anterior-posterior (AP) -1.7mm , medial-lateral (ML) $\pm 0.8\text{mm}$, dorsal-ventral (DV) -5.7mm from top of cortex; AP -1.7mm , ML $\pm 1.6\text{mm}$, DV -4.9mm from top of cortex. 3 sham control rats underwent four equivalent infusions of PBS only. A further 4 control rats had no sham lesion procedure and just had electrodes implanted.

During stereotaxic surgery, rats were implanted unilaterally with 28 electrodes of 25 μm diameter platinum-iridium wires arranged in a tetrode formation. Tetrodes were implanted aimed at the dorsal subiculum (AP -5.6mm , ML 2.5mm, DV -2.7mm from top of cortex), CA1 (AP -3.8mm , ML 2.5mm, DV -1.4mm from top of cortex), or both. Electrodes were stabilised with dental cement attached to the screws implanted into the skull. An additional stainless steel bipolar electrode of μm diameter targeting the ipsilateral retrosplenial cortex was also implanted in all but 2 cases. All electrodes were connected to a 32-channel Axona microdrive.

For the ATN inactivation experiment, a further 6 rats (ATN_{muscle}) were implanted with tetrode and bipolar electrode configurations alongside bilateral infusion cannulae (26 gauge, 4mm length). Cannulae were placed targeting ATN (AP -1.7mm , ML $\pm 3.8\text{mm}$, DV -4.0mm from top of cortex, at angle 28.6° toward centre), then fixed in position using dental cement and dummy cannulae (0.203mm diameter, 4mm length) inserted to prevent blockage. Electrodes were positioned as above, targeting dorsal subiculum and RSC, with two exceptions; 1 animal was implanted with four tetrodes targeting the subiculum, three targeting CA1 and one targeting retrosplenial cortex, with no additional bipolar electrode; and 1 rat was implanted with tetrodes targeting the subiculum and the bipolar electrode targeting CA1.

Finally, for the CA1 experiment, a further 3 rats received permanent bilateral

cytotoxic ATN lesions followed by electrode implantation into CA1. In all cases, glucosaline (5-10 mL) was administered subcutaneously postoperatively and the animal allowed to recover. Animal weight, activity, and hydration were closely monitored daily for a minimum of 7 days before beginning electrophysiological recordings.

A.1.9 Electrophysiological recordings

Recordings were obtained using an Axona Ltd 64-channel system, allowing dual recordings of single units and local field potentials from each electrode. Initial habituation recordings were conducted in a 60×60 cm square, walled arena (height 42cm). Later testing involved a larger arena (105×105 cm, 25cm). For both arenas, the walls and floors were made of wood painted matt black. A black curtain could be closed around the arena to remove distal spatial cues, and visual cues could be attached to the curtain as required. The habituation sessions in the small arena allowed the animal to acclimatise to the recording procedure and the experimenter to adjust electrode locations until the optimal recording depth was reached. After the habituation period (usually 3 to 7 days), rats were first trained on the behavioural tasks (T-maze then bow-tie maze) with up to 1 hour of free exploration with pellet-chasing tasks typically recorded afterwards on the same day. Pellet-chasing involved exploration in the small, and then large, arena to assess properties in both, and then consecutive recordings to examine sleep properties.

In the cohort for temporary inhibition of the ATN via muscimol rats were trained and habituated for at least one week before commencing experiments involving muscimol. The rats were trained to get used to the dummy cannulae. Further, the rats were also trained in the spatial alternation task, and electrophysiological activity was recorded during free exploration and pellet-chasing in the large and small arenas.

On the day of experimentation, electrophysiological recordings of exploration and pellet-chasing were first performed before infusion for 20-40 min to establish a baseline. The rat was then lightly restrained and muscimol (concentration 0.5mg/mL saline) infused through a 33 gauge infusion needle with a 1 or 2mm projection past the length of the implanted guide cannula, targeting the ATN. Muscimol was infused over 90 seconds using a 0.5 μ L Hamilton syringe and infusion pump. The infusion needle was retained in position for a further 60 seconds before it was removed and replaced with the dummy cannula. Rats received a 0.2 μ L infusion at each of two locations per hemisphere using both a 5 and 6mm length infusion needle to target the whole ATN.

Following infusion, electrophysiological recordings during pellet-chasing and exploration were conducted in consecutive 20-minute sessions in the small and large square arena. Arenas were swapped between recordings to maintain the animal's interest in

the surroundings and to assess whether spatial remapping occurred. Between 90 and 120 minutes into the experiment, the T-maze task was performed (8 trials, pseudo-randomised starting points). rats were then returned to the square arena and recorded for a further 2-3 hours, including during sleep. Regular diet and water were freely available in the recording arena after the T-maze test. The following day, further electrophysiological recordings were undertaken to determine whether the effects of muscimol had ceased and cell activity had returned to baseline. Muscimol infusions were repeated 1-2 weeks after the initial experiment. In 2 cases, an additional control infusion of saline was given. To visualise the location of the muscimol infusion, the tracer fluorogold was infused 1 day prior to perfusion.

A.1.10 Unit identification and isolation

Spike sorting was performed automatically in Tint using k-means (Axona) and identified k-means clusters refined manually. Unit identification used the following criteria: units had to be active and show consistent waveform characteristics (amplitude, height, and duration) during recording, as well as a clean refractory period (≥ 2 ms) in the ISI histogram. Spike amplitude was measured as the difference between the positive peak and the first negative peak before the positive peak, if present, or zero. Spike height was the difference between the spike peak to the minimum value of the spike waveform. Spike width was the distance in microseconds beyond which the waveform drops below 25% of its peak value. Histograms were used to assess spike width and determine whether qualitatively different neuron populations were being recorded following anterior thalamic lesions.

Units were sometimes seemingly recorded for more than one day, despite electrode lowering. For these cases, cells were monitored on the relevant tetrodes from day to day; for analysis, only clean recordings with the largest sample size and spikes of the highest amplitude were chosen. To avoid double-counting cells, care was taken to exclude seemingly related samples from analysis. During spike sorting, the signals from each cell were carefully followed from first appearance to complete loss, to avoid overestimation of cell counts. Once well-defined neuronal signals were isolated, recording commenced. For permanent lesion experiments, rats had to explore at least 90% of the arena in a session to be included in analyses to allow reliable calculation of spatial characteristics.

A.1.11 Spatial alternation testing

A four-arm cross-shaped wooden maze with raised sides (119 × 119cm full area; each arm 48 × 23cm; wall height 30cm) was used for the spatial alternation task, allowing

the rotation of start points. Each arm could be blocked close to the centre to form a T-maze. In addition, a barrier could be placed within an arm to form a holding area for the start position. Distal spatial cues were available in the recording room, including the pulled-back curtain, electrophysiological recording equipment set on wall-mounted shelves, a desk, and a computer.

Rats were first habituated in pairs to the maze and allowed to freely explore for 10 minutes. Rats then individually had two pre-training days, each consisting of four 5-minute sessions in which they were first placed behind a barrier at the start position at one end of the maze, retained for 10 seconds, then the door removed and the rat allowed to explore the maze. The maze arm opposite the start was blocked, and sucrose pellets were placed in a shallow dish at the end of each open arm so that they were not visible from the centre of the maze. Rats learned to run to an arm of the maze to obtain sucrose pellets, which were replaced once they had been consumed, allowing the animal to continuously run from arm to arm during the 5-minute session. The maze was then altered so that a different start arm and blocked arm were used, and another training session run. Each rat had four 5-minute training sessions per day such that each start/blocked configuration was experienced (opposite, adjacent, opposite) for 2 days initially, with an additional day if required.

For the experiment, rats were placed in the start position for 10 seconds. In each trial, there was an initial forced run (“sample”), in which two arms of the cross-shaped maze were blocked, forcing the animal into either the left or right arm to obtain two sucrose pellets from the end of the arm. The animal was then picked up and returned to the start position and confined for roughly 20 seconds before being released for the choice run (“test”), in which one of the barriers was removed from the maze so that the animal had the choice of either the left or right arm. Sucrose pellets were only available in the arm opposite to the sample run, so the rat was rewarded if it alternated. A choice was determined when the back paws of a rat had entered the arm. The rat was then removed from the maze and placed in an open holding cage for 2-3 minutes while the maze was rearranged. Each animal had eight sessions consisting of eight trials each, with both the start position and forced turn pseudorandomised so that the same arrangement did not occur more than twice consecutively. Rats were connected to the electrophysiological recording equipment throughout pretraining and each test session.

A.1.12 Spatial measures and rate maps

Standard statistical testing was performed using an open-source custom-written suite in Python, NeuroChaT (Islam et al., 2019) and additional custom code in R. Units were

classified based on the spatio-temporal features of their activity in the arena during pellet-chasing. A firing map of a cell was obtained by binning the location the rat was in when the cell fired and dividing this by the binned overall location of the rat. To analyse head direction cells, the animal's head direction was calculated using the relative position of two tracked light emitting diodes (LEDs) on a bar attached to the microdrive. The directional tuning function was determined by plotting the firing rate as a function of the head direction divided into 5° bins, and the firing rate calculated by the total number of spikes divided by time spent in each bin. To determine the existence of a hexagonal grid firing structure, grid index, size, and orientation were calculated. The grid cell analysis included calculating the spatial autocorrelation of the firing rate map and assessing the shape formed by the peaks in autocorrelation. For border cell analyses, cells with a firing profile that was parallel to the border of the arena were selected by plotting the positional firing pattern.

A.1.13 Hypothesis testing

Statistical analyses were performed in R Statistics. No statistical difference was noted between the normal control and sham control groups, so these were combined into a single control group unless otherwise stated. The distribution of variables was assessed using histograms; and where a normal distribution was determined, differences between groups were examined using Welch's two-sample t-test, an adaptation of Student's t-test that is more reliable when samples have unequal variances or sample sizes. In cases where the distribution of a variable was non-parametric, differences between groups were assessed using the Mann-Whitney U test. Where three groups were compared, analysis of variance (ANOVA) with a Tukey post hoc test was applied.

For temporary inactivation experiments, data were compared with baseline (immediately before muscimol infusion) values using Welch's two-sample t-test with a Bonferroni correction to account for multiple comparisons and reduce the likelihood of Type 1 errors. Statistical tests are reported in the text and appropriate figure legends. Boxplots represent median, 25th, and 75th percentiles. Boxplot tails represent the smallest and largest value within 1.5 times the interquartile range, and outliers are defined as values that are ≥ 1.5 times and ≤ 3 times the interquartile range. Data means are represented with a diamond-shaped marker. Error bars on scatter plots indicate SEM.

A.1.14 Perfusion and histology

After completion of experiments, rats were killed and perfused transcardially with 0.1mol/L PBS and then 2.5% paraformaldehyde (PFA) in 0.1mol/L PBS. Brains were

removed and postfixed in PFA for 24 hours and then transferred to 30% sucrose in 0.1mol/L PBS solution for 2 days. A cryostat was used to cut 40 μ m sections in a 1:4 series. One series was mounted onto double gelatin-subbed microscope slides and, once completely dry, washed in decreasing concentrations of alcohol (100%, 90%, 70%) before being stained with cresyl violet, a Nissl stain. Sections were then dehydrated with increasing alcohol concentrations, washed in xylene, and coverslipped. Of the other series, one was reacted against anti-calbindin antibody raised in mouse and another against anti-neuronal nuclear protein (NeuN) antibody raised in mouse. The remaining series was reacted with either anti-parvalbumin antibody raised in mouse or, in the temporary inactivation cohort, with anti-fluorogold raised in rabbit.

Sections were washed in a quench solution (10% methanol and 0.3% hydrogen peroxide in distilled water) before 0.1mol/L PBS, then phosphate-buffered saline with a low-concentration detergent for protein solubilisation (PBST) (2mL Triton X-100 in 1L 0.1mol/L PBS; pH 7.35) washes. Sections were stirred for 1 hour in 4% normal horse serum in PBST before the primary antibody was added (1:5000 dilution in PBST for calbindin, parvalbumin, and fluorogold; 1:10,000 for NeuN) and stirred at 4° overnight. Sections were then washed in PBST before being incubated for 2 hours in 1:250 dilution of horse-anti-mouse or, in the case of fluorogold, horse-anti-rabbit, in PBST. After further PBST washes, sections were incubated at room temperature in Vectastain Elite ABC Solution (Vector Labs) before further PBST and PBS washes. Sections were then reacted with DAB solution and washed in PBS.

All sections were mounted on double-subbed slides, and fluorogold-reacted slides were lightly stained with cresyl violet for improved tissue visualisation before coverslipping. Sections were imaged using either an Olympus BX51 upright microscope or Leica Microsystems Aperio AT2 slidescanner.

Individual images of sections were aligned and tiled using Inkscape (<http://inkscape.org>) and FIJI (Preibisch et al., 2009; Schindelin et al., 2012). Cresyl violet-stained sections helped to confirm electrode placement. To assess lesion success, the ATN was segmented from photomicrographs of anti-NeuN-reacted sections using FIJI and the resulting image thresholded to show nuclei separation, before using the inbuilt Analyze Particles plugin to obtain a cell count. The ATN cell counts were compared between the ATNx and control groups to determine lesion effectiveness. Calbindin-reacted sections helped to determine the status of nucleus reuniens.

References

- Aggleton, J. P., & Nelson, A. J. D. (2015). Why do lesions in the rodent anterior thalamic nuclei cause such severe spatial deficits? *Neuroscience and Biobehavioral Reviews*, *54*, 131–144.
- Boccaro, C. N., Sargolini, F., Thoresen, V. H., Solstad, T., Witter, M. P., Moser, E. I., & Moser, M.-B. (2010). Grid cells in pre-and parasubiculum. *Nature neuroscience*, *13*(8), 987–994.
- Brotons-Mas, J. R., Schaffelhofer, S., Guger, C., O'Mara, S. M., & Sanchez-Vives, M. V. (2017). Heterogeneous spatial representation by different subpopulations of neurons in the subiculum. *Neuroscience*, *343*, 174–189.
- Brotons-Mas, J. R., Montejo, N., O'Mara, S. M., & Sanchez-Vives, M. V. (2010). Stability of subicular place fields across multiple light and dark transitions. *European Journal of Neuroscience*, *32*(4), 648–658.
- Broussard, J. I., Redell, J. B., Zhao, J., Maynard, M. E., Kobori, N., Perez, A., Hood, K. N., Zhang, X. O., Moore, A. N., & Dash, P. K. (2020). Mild traumatic brain injury decreases spatial information content and reduces place field stability of hippocampal CA1 neurons. *Journal of neurotrauma*, *37*(2), 227–235.
- Duvelle, É., Grieves, R. M., Hok, V., Poucet, B., Arleo, A., Jeffery, K. J., & Save, E. (2019). Insensitivity of place cells to the value of spatial goals in a two-choice flexible navigation task. *Journal of Neuroscience*, *39*(13), 2522–2541.
- Goodridge, J. P., & Taube, J. S. (1997). Interaction between the postsubiculum and anterior thalamus in the generation of head direction cell activity. *Journal of Neuroscience*, *17*(23), 9315–9330.
- Griffin, A. L. (2015). Role of the thalamic nucleus reuniens in mediating interactions between the hippocampus and medial prefrontal cortex during spatial working memory. *Frontiers in Systems Neuroscience*, *9*.
- Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, *436*(7052), 801–806.
- Hales, J. B., Schlesiger, M. I., Leutgeb, J. K., Squire, L. R., Leutgeb, S., & Clark, R. E. (2014). Medial entorhinal cortex lesions only partially disrupt hippocampal place cells and hippocampus-dependent place memory. *Cell reports*, *9*(3), 893–901.
- Hok, V., Chah, E., Reilly, R. B., & O'Mara, S. M. (2012). Hippocampal dynamics predict interindividual cognitive differences in rats. *Journal of Neuroscience*, *32*(10), 3540–3551.
- Islam, M. N., Martin, S. K., Aggleton, J. P., & O'Mara, S. M. (2019). NeuroChaT: A toolbox to analyse the dynamics of neuronal encoding in freely-behaving rodents in vivo. *Wellcome Open Research*, *4*, 196.

- Jankowski, M. M., Islam, M. N., Wright, N. F., Vann, S. D., Erichsen, J. T., Aggleton, J. P., & O'Mara, S. M. (2014). Nucleus reuniens of the thalamus contains head direction cells (H. Eichenbaum, Ed.). *eLife*, *3*, e03075.
- Matulewicz, P., Ulrich, K., Islam, M. N., Mathiasen, M. L., Aggleton, J. P., & O'Mara, S. M. (2019). Proximal perimeter encoding in the rat rostral thalamus. *Scientific Reports*, *9*(1), 2865.
- Meshulam, L., Gauthier, J. L., Brody, C. D., Tank, D. W., & Bialek, W. (2017). Collective behavior of place and non-place neurons in the hippocampal network. *Neuron*, *96*(5), 1178–1191.
- Preibisch, S., Saalfeld, S., & Tomancak, P. (2009). Globally optimal stitching of tiled 3d microscopic image acquisitions. *Bioinformatics*, *25*(11), 1463–1465.
- Radivojevic, M., Franke, F., Altermatt, M., Müller, J., Hierlemann, A., & Bakkum, D. J. (2017). Tracking individual action potentials throughout mammalian axonal arbors. *Elife*, *6*, e30198.
- Robbins, A. A., Fox, S. E., Holmes, G. L., Scott, R. C., & Barry, J. M. (2013). Short duration waveforms recorded extracellularly from freely moving rats are representative of axonal activity. *Frontiers in Neural Circuits*, *7*.
- Schindelin, J., Arganda-Carreras, I., Frise, E., Kaynig, V., Longair, M., Pietzsch, T., Preibisch, S., Rueden, C., Saalfeld, S., Schmid, B., Tinevez, J.-Y., White, D. J., Hartenstein, V., Eliceiri, K., Tomancak, P., & Cardona, A. (2012). Fiji: An open-source platform for biological-image analysis. *Nature Methods*, *9*(7), 676–682.
- Sharp, P. E., & Green, C. (1994). Spatial correlates of firing patterns of single cells in the subiculum of the freely moving rat. *Journal of Neuroscience*, *14*(4), 2339–2356.
- Thompson, L., & Best, P. (1989). Place cells and silent cells in the hippocampus of freely-behaving rats. *Journal of Neuroscience*, *9*(7), 2382–2390.
- Tsanov, M., Chah, E., Vann, S. D., Reilly, R. B., Erichsen, J. T., Aggleton, J. P., & O'Mara, S. M. (2011). Theta-modulated head direction cells in the rat anterior thalamus. *The Journal of Neuroscience*, *31*(26), 9489.
- van Groen, T., & Wyss, J. M. (1990). The connections of presubiculum and parasubiculum in the rat. *Brain research*, *518*(1), 227–243.
- Wang, B. W., Hok, V., Della-Chiesa, A., Callaghan, C., Barlow, S., Tsanov, M., Bechara, R., Irving, E., Virley, D. J., Upton, N., et al. (2012). Rosiglitazone enhances learning, place cell activity, and synaptic plasticity in middle-aged rats. *Neurobiology of aging*, *33*(4), 835–e13.
- Winter, S. S., Clark, B. J., & Taube, J. S. (2015). Disruption of the head direction cell network impairs the parahippocampal grid cell signal. *Science*, *347*(6224), 870–874.