

Comparative Analysis of Algorithmic Approaches for Mapping Information Flow in the Human Hippocampus

Thomas Gillespie



4th Year Project Report
Computer Science
School of Informatics
University of Edinburgh

2024

Abstract

Significant improvements to electrophysiological technology now enable the recording of activity in large-scale neural networks at the subcellular level. Although such technology has been used extensively to map neural networks in slices both from animals and human brain ‘organoids’, a lack of application to actual human brain tissue limits understanding of connectivity networks in the human brain. This project considers the relative ability of two algorithms to map associative structures in neural data: Sharf et al.’s ‘functional connectivity’ method and an improved ‘spatial activity profiling’ method originally proposed by Muthmann. Each method was tested on synthetically generated neural data, and applied to six datasets of human hippocampal activity, derived from high-density microelectrode array (HD MEA) recordings of tissue slices resected from patients undergoing surgery for intractable epilepsy. Findings reveal that Muthmann’s adapted method outperforms Sharf et al.’s in capturing information flow within the synthetic data. However, when applied to the real hippocampal data, performance reverses, with Muthmann’s adapted method struggling to identify the large-scale structures successfully captured by Sharf et al.’s method. Although Muthmann’s method may benefit from further refinement, overall, the results present a promising case for using these methods to capture associative structures in neural data. ¹

¹In this project, Muthmann’s local structures of neural association, as calculated by the ‘spatial activity profile’ method, and Sharf et al.’s ‘functional connectivity networks’ will be referred to as ‘associative structures’.

Research Ethics Approval

This project was planned in accordance with the Informatics Research Ethics policy. It did not involve any aspects that required approval from the Informatics Research Ethics committee.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

(Thomas Gillespie)

Acknowledgements

First and foremost, I would like to thank Dr David Sterratt for offering to supervise this self-proposed project alongside his other commitments – I feel very grateful for all of the time he dedicated to our weekly meetings. David was incredibly patient, encouraging me to persist when faced with challenges and reminding me to keep the big picture in mind. His feedback was always thoughtful and constructive, which really helped me to develop my ideas.

I would also like to thank Dr David Haussler and his student Matthew Elliott for sparking my interest in computational neuroscience and electrophysiology, and for welcoming me into their research group during my year abroad at the University of California Santa Cruz. I am incredibly grateful for their continued support and feedback on this project, despite the geographical distance.

Finally, I would like to thank my second supervisor, Dr Angus Chadwick, as well as Dr Matthias Hennig and his Neurons and Systems group for encouraging me to explore this topic and offering several exciting ideas to consider.

Table of Contents

1	Introduction	1
2	Background	3
2.1	Microelectrode Array Technology	3
2.2	Three Dimensional Brain Cultures	4
2.3	Functional Neuronal Circuitry	5
2.4	Taxonomy of Functional Connectivity Metrics	6
2.5	Spike Time Tiling Coefficient	7
2.6	The Human Hippocampus	9
3	Design and Implementation	10
3.1	Methods for Mapping Information Flow	10
3.1.1	Functional Connectivity Method of Sharf et al.	10
3.1.2	Muthmann’s Spatial Activity Profiling Method	11
3.1.3	Improved Version of Muthmann’s Method	14
3.2	Synthetic Neuronal Data	15
3.2.1	Randomised Data	16
3.2.2	Plane Wave Data	17
3.2.3	Ring Wave Data	18
3.2.4	Rotating Loop Wave	20
3.3	Implementation	21
3.3.1	Methods	21
3.3.2	Synthetic data	23
4	Testing on Synthetic Data	24
4.1	Performance of Methods on Synthetic Data	24
4.1.1	Sharf et al.’s Functional Connectivity Method	24
4.1.2	Adaptation of Muthmann’s Spatial Activity Profile Method	26
4.2	Evaluation	28
5	Results	29
5.1	Overview of Human Hippocampal Data	29
5.2	Sharf et al.’s Functional Connectivity Method	31
5.3	Adaptation of Muthmann’s Spatial Activity Profile Method	34
6	Discussion	37

6.1	Method Evaluation	37
6.2	Simulation Evaluation	38
7	Conclusions	40
	Bibliography	41
A	First appendix	45
A.1	Full Detail of Sharf et al.'s Functional Connectivity Method	45
A.2	Full Detail of Human Hippocampal Datasets	47
A.3	Relationship between Connected Pairs and Classification of Neurons in Functional Connectivity Networks	53
A.4	Full Size Functional Connectivity Maps	54
A.5	Spatial Activity Profile Parameters for Human Data	57

Chapter 1

Introduction

The human hippocampus is a small, well-connected region of the brain, known to be responsible for learning and memory [3, 23], and more recently, shown to regulate emotional, navigatory, and hypothalamic function [36, 40, 19]. As a highly plastic region, it is especially vulnerable to damage from neurodegenerative diseases like Alzheimer's and Parkinson's, conditions that impair cognitive function by disrupting adult neurogenesis – the process of forming new neurons [38]. Advancements in non-invasive neuroimaging techniques such as functional magnetic resonance imaging (fMRI), and data collected through large-scale initiatives, such as the 'Human Connectome Project' (an effort to map all neuronal pathways underlying brain function) have revealed how the hippocampus connects to wider brain networks (such as the default mode network), and how activity patterns differ within the region [14]. However, understanding of connectivity networks within the hippocampus, and their implications for research into neurodegenerative disease, remains limited.

Recently, high-density microelectrode arrays (HD MEAs) have been used to record extracellular activity of individual neurons in slices of three-dimensional brain cultures derived from stem cells to study network connectivity, and how it is disrupted by disease [27, 33]. Although such cultures, known as 'organoids', can be grown to represent specific brain regions and assemble representative circuits, they are imperfect models of the human brain, lacking a blood-brain barrier and a human-length lifespan. However, HD MEAs have recently been used to record the extracellular activity of neurons from a different source – hippocampal tissue resected from patients undergoing surgery for treatment-resistant epilepsy [4]. Such recordings offer a unique opportunity to explore the hippocampus' internal circuitry and functional connectivity, with significant potential to improve our understanding of brain function and disease mechanisms more broadly.

A variety of algorithmic approaches to exploring connectivity networks exist, and have been applied both to animal brain slices and human organoid derived data. Muthmann, for instance, considers the spatial and temporal proximity of each spike event to calculate 'spatial activity profiles' that capture 'local structures' of neural association in rodent brain slices [24]. A new approach, proposed by Sharf et al. in 2022 [33], computes and visualises 'functional connectivity networks' from spike-sorted time series data derived

from HD MEA recordings of cerebral organoid slices.

This project builds on the work and results of Muthmann and Sharf et al. by applying their methods to both synthetic and human brain slice derived data, attempting where possible to corroborate findings of each method to validate conclusions drawn. The overall goals of this project are to:

- Evaluate and contrast the two algorithms' (Muthmann's and Sharf et al.'s) effectiveness in capturing patterns of neuronal association when applied to spike-sorted time series data derived from HD MEA recordings.
- Refine these methods as necessary to address any problems encountered, ensuring robust validation of any adaptations.
- Apply the improved methods to recordings of extracellular activity from actual human brain slices, visualising the results in interpretable two-dimensional maps.
- Explore patterns of association revealed by the application of each method to the human data.

To this end, the following contributions were made:

- Sharf et al.'s method for computing and visualising 'functional connectivity networks' was implemented from scratch in Python, as outlined in Chapter 3.
- The 'spatial activity profiles' method proposed by Muthmann was implemented as far as possible in Python, with details presented in Chapter 3.
- Several limitations of Muthmann's method were identified and detailed extensively in Chapter 3. Then, a series of conceptual adaptations intended to resolve the issues were proposed and subsequently implemented in an updated version of the method.
- Synthetic neural datasets were generated using Python code written from scratch. Chapter 3 outlines the design, implementation and biological inspiration behind this data.
- In Chapter 4, the synthetic datasets were used to test and evaluate the ability of each method (Sharf et al.'s, and Muthmann's adaptation) to capture embedded associative structures.
- The human hippocampal data were analysed using both methods in Chapter 5 to investigate connectivity patterns and information flow through the six subregions of the hippocampus.
- The ability of each method to capture and map neural patterns was critically examined in Chapter 6. There, further improvements to the method design and the generation of test data were proposed, alongside a discussion on the future applications of the algorithms.

Chapter 2

Background

2.1 Microelectrode Array Technology

Recent advances in microelectrode array (MEA) technology have helped to overcome several hurdles previously faced by electrophysiology researchers. Although MEAs have been used to take extracellular recordings of neural circuits for some time, their capabilities were limited by a relatively poor spatial resolution (usually $\geq 30 \mu\text{m}$) and low number of channels (often less than 300) [5], meaning that mapping large neural circuits at the cellular level was previously impossible. A new type of MEA that utilises complementary metal-oxide semiconductor (CMOS) technology helps to tackle these limitations [18]. These high-density MEAs place a large number of closely-spaced electrodes together, while utilising multiplexing techniques and incorporating CMOS technology to enable large area, high spatial resolution recordings. This has many exciting implications for the field of electrophysiology and the neuroscience community more broadly, namely that the technology could facilitate the full activity of a neuronal population to be captured by growing neurons on the active area of the MEA [25]. It is thought that this could enable a more robust estimation of network parameters, such as functional connectivity – defined as the statistical dependence between time series data of electrical activity within the brain [22].

Maxwell Biosystems' new device, the MaxOne MEA utilises novel CMOS technology to decrease amplifier size, enabling an integrated circuit of 26,400 electrodes in an area $3.85 \text{ mm} \times 2.10 \text{ mm}$ [5]. This spatial density of electrodes – 3265 per mm^2 – is a significant improvement over legacy MEAs. Indeed, at the turn of the 21st century, researchers at ETH Zurich used an MEA with a total of only 20 electrodes and density of 0.2 electrodes per mm^2 to investigate hippocampal organotypic slice cultures [39]. Another major advantage of the MaxOne MEA is its ability to simultaneously record across 1024 channels, meaning that the electrical activity of an entire network of neurons can be analysed. Moreover, where older MEAs would miss signals from the majority of cells in a population, MEAs such as the MaxOne allow recording of individual neurons at the subcellular resolution, meaning it is possible to further investigate the composition and function of individual neurons in detail. One disadvantage of Maxwell's MEAs is that their electrodes only capture signal from the outer layer of tissue, in contrast

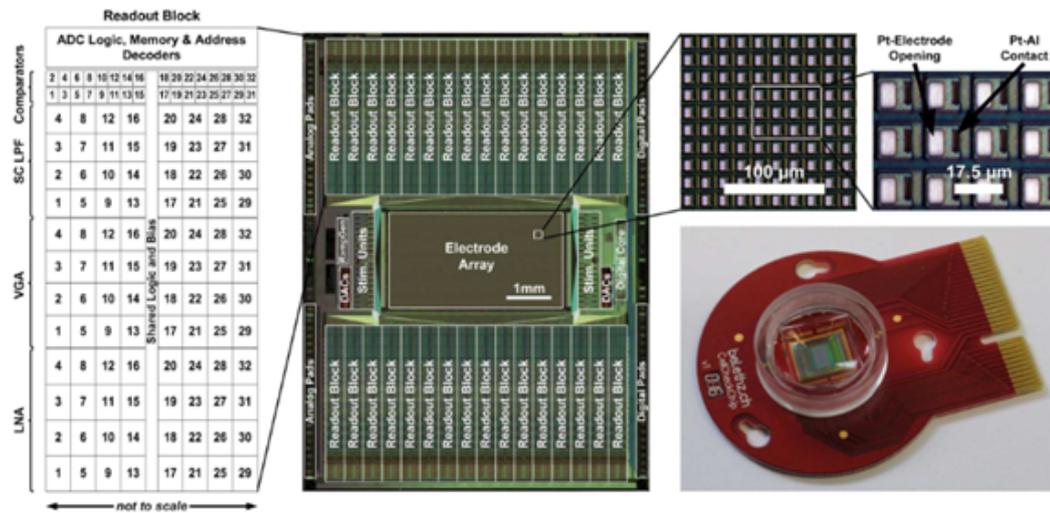


Figure 2.1: A detailed view of Maxwell's MaxOne high-density CMOS microelectrode array. Left to right: Layout of an individual reading block with 32 channels. The chip micrograph with constituent readout blocks above and below central electrode array component, stimulation units to the right and left. Enlarged view of the electrode array, with electrode opening and Pt-Al contact. Packaged device in red. From Ballini et al [5].

to 3Brain's MEAs that utilise needle-like electrodes to penetrate the tissue surface, leading to data with a better signal-to-noise ratio [21, 1]. Modern MEAs including those manufactured both by Maxwell and 3Brain feature numerous stimulation channels that have been shown to reliably stimulate individual neurons, facilitating research into neural plasticity [12, 1].

2.2 Three Dimensional Brain Cultures

To date, there exists no literature detailing experimental procedures in which HD MEAs have been used to take extracellular recordings from human brain slices. However, such devices have been used to record neural activity from a variety of other sources – from animal brain [26] and retinal slices and most notably, from three-dimensional brain cultures derived from animal and human induced pluripotent stem cells (iPSCs) [31]. From a research perspective, the advantage of these brain cultures, also known as cerebral organoids, is that they are simpler and more easily accessible than the animal brain slices that inform much of our current understanding of the human brain. Moreover, cerebral organoids have been shown to model essential features of the developing human brain, and indeed even some mature features, including inhibitory and excitatory synapses and dendritic spines [27].

The accessibility of brain organoids, coupled with advances in HD MEA technology offers a unique opportunity to model disease and brain development – indeed researchers have maintained cultured brain organoids for as long as 20 months [16], providing unique insight into early stage maturation of the human brain and central nervous system. Cerebral organoids are also increasingly being used in stimulation experiments to investigate neural plasticity [45, 12], in research relating to neuronal circuitry [33], and

in many other methods that provide an alternative means of improving understanding of the human brain. However, it is worth noting that understanding of cerebral organoids is still developing – indeed, little is known about the impact of genetic background on iPSC-derived neuronal networks, which have not been benchmarked nearly as extensively as rodent neuronal cultures. Nonetheless, in the absence of access to human primary tissue, cerebral organoids remain a valuable resource to the advancement of neuroscientific knowledge and the development of treatments for neurological disorders. In fact, human organoid derived models or datasets that support the scientific rationale behind new therapies are now often a prerequisite for significant investment by pharmaceutical companies [13].

2.3 Functional Neuronal Circuitry

In their 2022 Nature Neuroscience paper [33], Sharf et al. outline emergent features of human brain organoids that resemble characteristics of functional brain organisation. Using six month old human brain organoids, sliced to 500 μm thickness and plated on a MaxOne HD MEA, Sharf et al. presented analysis on neuronal spiking dynamics, local field potentials and theta oscillations. After selecting the most active 1020 electrodes, spatial maps of extracellular action potential waveforms were constructed which informed the identification of individual neural units by spike-sorting. When the electrodes on such an MEA setup record action potentials, or ‘spikes’ from multiple neurons simultaneously, each channel picks a mixture of signals from different surrounding neurons. Spike-sorting is the process used to categorise these spikes into clusters, each thought to originate from a specific neuron, thereby distinguishing them from background noise. Although Sharf et al. describe the Kilosort2 spike-sorting algorithm as having “optimal accuracy and precision based on ground truth data for arrays with similar electrode densities”, this is not a consensus opinion in the scientific literature, with others arguing that different algorithms such as MountainSort4, might be more reliable, and less prone to producing a large number of false positive units [17]. On the other hand, given that Maxwell’s MaxOne is one of the most dense recording devices, it is arguable that Kilosort2 is suitable in this case because the algorithm is specifically designed for “high-channel count” data – here, 1020 electrodes from a total possible of 26400 were used for recording.

One of the most notable analyses presented in Sharf et al.’s paper is the mapping of information flow through a human brain organoid slice – visualised primarily by a two-dimensional functional connectivity map of the entire region. Constructed by considering pairwise spike correlations and distribution of pairwise mean spike latency, the map is intended to show the patterns of association between different regions in the tissue. The precise components of the method, and their justification for use are outlined in Section 3.1 and Appendix A.1.

The ability to produce two-dimensional visualisations such as the map in Figure 2.2 is a unique characteristic of experiments that rely on HD MEAs to record data. Moreover, a variety of statistics and metrics can be similarly projected onto two dimensions, to create a visual representations of firing rate, covariance or cross-correlation for instance. This novel technique communicates electrophysiological findings in a clear and engaging

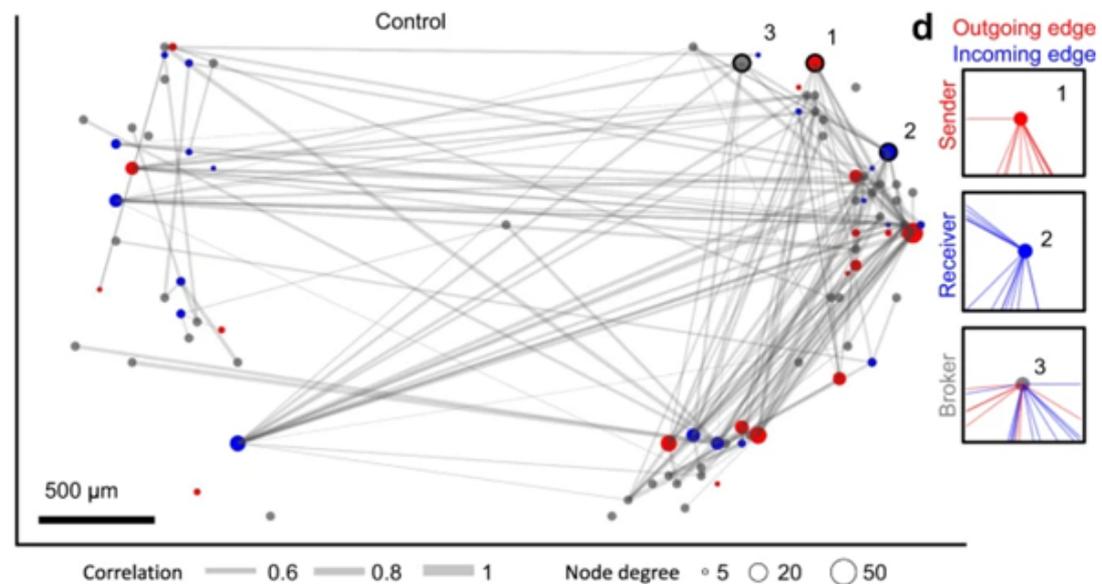


Figure 2.2: A two-dimensional map visualising flow of information through a human brain organoid slice. Left to right: Functional connectivity map showing association between neurons with connecting grey edges, and directionality of neurons by colour. Example sender, receiver and broker nodes and their connections. Taken from Sharf [33].

manner, in contrast to other methods of presentation such as tables, rasters and matrices, which are often difficult to interpret intuitively.

More broadly, the human brain organoids highlighted in Sharf et al.’s paper serve as a useful tool for modeling some of the complex neural activities characteristic of the developing human brain. Improving our understanding of functional connectivity in this way is important, as it will help us understand how different brain regions connect and organise during development, and how they work together to support important functions like memory, language and attention. A deeper understanding of functional connectivity in the context of brain function and development would also be highly valuable to medicine, where it could inform the development of new therapeutic approaches to neurological disorders such as Alzheimer’s disease and epilepsy [37].

2.4 Taxonomy of Functional Connectivity Metrics

A variety of approaches to mapping functional connectivity exist. Bastos et al. offers a comprehensive comparison of many of these methods, describing a “Taxonomy of Functional Connectivity Metrics” [6], in which methods are categorised by their ability to quantify the direction of an interaction between neurons. So-called “directed” measures aim to capture a “statistical causation” from the data, relying on the idea that effects follow their causes, while “non-directed” measures simply capture a degree of interdependence between signals – see Figure 2.3. The authors further differentiate these categories into “model-based” and “model-free” approaches. They explain that model-based approaches such as Pearson’s correlation coefficient assume a linear relationship between two variables, in contrast to metrics that also measure non-linear

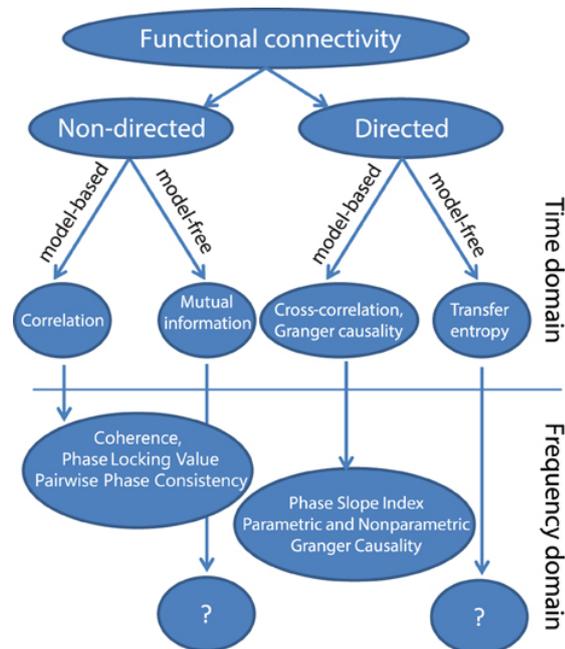


Figure 2.3: Bastos' Classification of Functional Connectivity Techniques [6].

interdependence between variables like “mutual information” or “transfer entropy”.

This project aims to capture some directional element to the interactions between neurons. With this in mind, when using non-directed methods, such as spike time tiling coefficient (STTC), complementary information such as the latency distribution between neural units will be used to estimate the direction of information flow. Although different metrics are used to cross-validate findings when possible, the project is somewhat limited by a lack of model-free approaches to detect directed interactions between neurons. Given more time, it would be interesting to compare model-based findings to results from model-free approaches such as “transfer entropy” [43] that capture non-linear patterns of information flow, as discussed in Section 6.1.

2.5 Spike Time Tiling Coefficient

The spike time tiling coefficient (STTC) is a fundamental component of Sharf et al.'s ‘functional connectivity’ method, which is in turn, central to this project. Proposed by Cutts and Eglen [10], STTC was intended to replace the correlation index as a non-directed estimate of association between pairs of neurons. It is now defined below.

Considering the spike trains of two neurons, A and B , we first compute T_A , defined as the proportion of total recording time which lies within a certain window ($\pm\Delta t$) from any spike in A 's train. T_B is calculated similarly.

Then, we calculate P_A , defined as the proportion of spikes in A 's train that lie within ($\pm\Delta t$) from any spike that occurs in B 's spike train. Similarly, P_B is defined for spikes in B 's train relative to A 's train.

Finally, to compute the coefficient for the example in Figure 2.4 we can use the formula

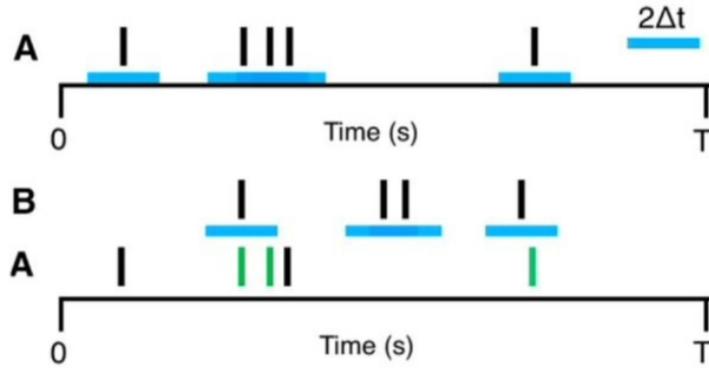


Figure 2.4: T_A is given by the fraction of total recording time (black) which is covered (tiled) by blue bars. Here T_A is $1/3$, and T_B is $1/4$. P_A is the number of green spikes in A (3) divided by the total number of spikes in A (5). Here P_A is $3/5$, and P_B is $2/5$. Cropped from Cutts and Eglen [10].

below, yielding the following result:

$$\begin{aligned}
 STTC &= \frac{1}{2} \left(\frac{P_A - T_B}{1 - P_A T_B} + \frac{P_B - T_A}{1 - P_B T_A} \right) \\
 &= \frac{1}{2} \left(\frac{\frac{3}{5} - \frac{1}{4}}{1 - \frac{3}{5} \cdot \frac{1}{4}} + \frac{\frac{2}{5} - \frac{1}{3}}{1 - \frac{2}{5} \cdot \frac{1}{3}} \right) \\
 &= 0.29
 \end{aligned}$$

STTC values range from -1 to $+1$. A value of $+1$ indicates full autocorrelation, where each spike in train A aligns perfectly with a spike in train B within the specified window. A value of -1 occurs when that no spikes in train A fall within the $(\pm\Delta t)$ window of any of B's spikes, indicating a complete lack of association. Positive values suggest a higher-than-chance level of correlation, while negative values suggest the opposite. A value of 0 suggests no significant alignment or misalignment – it is the correlation expected by chance.

Cutts and Eglen argue that a rival metric, the correlation index – the factor by which the firing rate of neuron A increases over its mean value if measured within a fixed window of spikes from neuron B – has two undesirable properties that mean it is unsuitable for comparing correlations when firing rates differ significantly. They emphasise the importance of quantifying “functionally significant correlations”, citing a study by Demas et al. [11], where correlated activity in spontaneous retinal waves helps neighbouring neurons wire to common targets, contributing to map formation. There is some debate about the what the window $\pm\Delta t$ can take, but Cutts and Eglen conclude that it is often dictated by the specimen investigated and that ultimately, it can take any value.

2.6 The Human Hippocampus

Given the aim of mapping functional networks and pathways of information flow within the human hippocampus, a brief overview of the anatomy, function, and medical significance of the region is set out here. Comprehensive details regarding the origin and methodology used to obtain the extracellular recordings of the human hippocampus examined in this study are available in Section 5.1.

The human hippocampus, named for its resemblance to the seahorse, is an important component of the human brain, often associated with memory consolidation, decision making, and spatial representation [15]. It is located in the medial temporal lobe and is a bilateral structure, with mirror image halves situated in the left and right sides of the brain [41] – see Figure 2.5. Broadly, it can be split into the following sections – the subiculum, the dentate gyrus (DG), and the hippocampus proper. The ‘Cornu Ammonis’, a structure that describes the layers of the hippocampus is split into CA1 through CA4. It is a highly connected region, receiving inputs that arrive in the entorhinal cortex, before complex routing to the CA1 region. Its outputs pass directly from the subiculum to the entorhinal cortex and amygdala, which in turn, connect to many other brain areas.

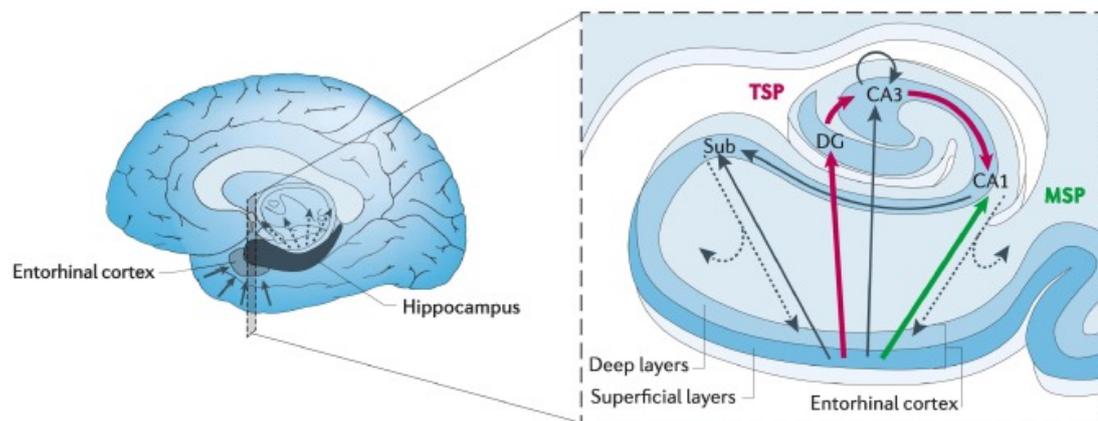


Figure 2.5: Location and anatomical circuit of human hippocampus. Taken from Turk-Browne [41].

Functionally, the hippocampus is generally accepted to be associated with representing environments spatially [20]. It is also understood to be essential for the ability to remember specific events – indeed, ‘H.M.’, a now-famous epileptic patient became severely amnesic and lost all ability to form new memories after having his left hippocampus and temporal lobe removed [35]. Notably, certain parts of the hippocampus – in particular, the entorhinal cortex – are some of the first brain regions to be damaged by Alzheimer’s disease, a degenerative brain condition that primarily affects memory and cognitive function. The region is also thought to be prone to damage by epilepsy and hypoxia, causing further damage to memory function.

Chapter 3

Design and Implementation

3.1 Methods for Mapping Information Flow

3.1.1 Functional Connectivity Method of Sharf et al.

In the 2022 “Functional neuronal circuitry” paper, Sharf et al. present a functional connectivity map to visualise information flow through human brain organoid slices – see Figure 2.2. In essence, their method considers the STTC values and unimodality of the latency distributions between neuron pairs to derive a non-directed functional network. Then, to infer direction, they assess the mean of the latency distributions between neuron pairs, and the number of incoming and outgoing edges for each neuron. The method used to derive the map is very complex and is explained at the high level here and set out with full detail in Appendix A.1.

- **Nodes:** The circles on the graph represent neurons, or “neural units” – depending on the confidence placed in the spike-sorting algorithm. The position of the node on the graph represents its location on the organoid slice.
- **Edges:** The magnitude of the STTC value for each neuron-neuron pair determines the edge strength between neurons. Neuron pairs with more correlated spike trains have a darker, thicker edge connecting them on the functional connectivity map. These edge strengths are determined using a correlation window of 20ms, meaning that spikes that are more than 20ms apart do not contribute to the STTC value, and the corresponding edge weight. Moreover, a minimum STTC threshold is set (in Sharf et al.’s paper at 0.35) to exclude any weak (and therefore less significant) couplings.
- **Directionality:** The authors excluded “bi-directional” connections – those suggested by non-unimodal latency distributions by using Hartigan’s diptest (see Appendix A.1). For each of the neuron pairs, a predominant directionality toward the reference unit was assumed in the case of a negative mean spike time latency to the reference unit, and vice versa. Each neuron-to-neuron edge was then assigned a direction of ‘incoming’ or ‘outgoing’, based on this mean spike time latency. For each neuron in the network, the mean spike timing latency value

to each other neuron it is connected to is computed. Negative values indicate direction towards the reference neuron, and vice versa. Then, for each neuron, the number of ‘incoming edges’ (D_{in}) is set as the number of connected neurons that tend to fire before it (i.e. the mean latency is negative), and its number of ‘outgoing edges’ (D_{out}) as the number of connected neurons that tend to fire after (i.e. mean latency is positive).

- **Classification of nodes:** The nodes are then classified as “Senders” – Red, “Receivers” – Blue, and “Brokers” – Grey, based on their number of incoming and outgoing edges (D_{in} and D_{out}). The threshold is set between 0 and 1 – in Sharf et al.’s paper, 0.8 was selected.

- Senders: $\frac{D_{out}-D_{in}}{D_{out}+D_{in}} > Threshold$

- Receivers: $\frac{D_{in}-D_{out}}{D_{in}+D_{out}} > Threshold$

- Brokers: $\left| \frac{D_{in}-D_{out}}{D_{in}+D_{out}} \right| < Threshold$

3.1.2 Muthmann’s Spatial Activity Profiling Method

Muthmann, supervised by Matthias Hennig, approached the same phenomena from a different angle in his doctoral thesis, “Multiscale Modeling of Homeostatic Processes” [24] by considering the spiking activity of single neurons in relation to network activity over long periods. Specifically, he investigated spatial patterns of the average activity surrounding active neurons at given locations on HD MEAs. Rather than considering pairwise interactions, as Sharf et al.’s functional connectivity method did, Muthmann looked at a “rough profile” of the spatially and temporally surrounding activity whenever a spike occurred.

3.1.2.1 Method Description

Muthmann describes this method with a limited degree of detail. The following section attempts to interpret and explain in full the intended method proposed.

1. **Subset events:** For each spike event in the recording, consider all other events that fall within a certain number of events away from (before or after) the given event. Of these, exclude those that occur outside of a specified spatial range of the given event, leaving a collection of spatially and temporally proximate events.
2. **Angle and distance collection:** Compute the angle and distance from the event of interest to each of these surrounding events, and collect the results in angle and distance histograms, each with six bins.
3. **Spatial averaging:** The angle and distance histograms are “spatially averaged using a spatial sliding window”. This is thought to provide an estimate of average directional and distance activity for locations on the MEA where neurons are not recorded from. Muthmann notes that this estimate leads to a strong bias in the direction of the MEA’s centre, which is where the “average activity is centred”.

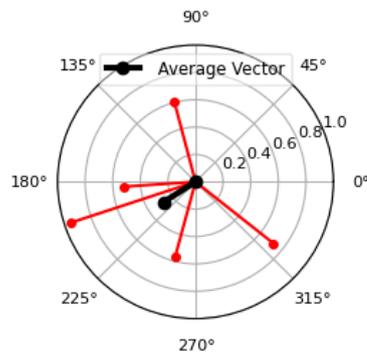


Figure 3.1: Polar plot of 5 random vectors (red) and their average (black).

4. **Local structure identification:** Another average, computed over a larger sliding window is subtracted from the previous density in an attempt to identify local structures which characterise network activity.
5. **Histogram averaging:** The angle and distance histograms for each location were represented in a polar plot. The polar plots then could be averaged to produce representative vectors, as in Figure 3.1.
6. **Normalisation:** Muthmann describes the steps above as corresponding to “a random walk for random spikes”, and proposes normalising the vectors by $\frac{1}{3\sqrt{N}}$, where N is the total number of events.
7. **Colour representation:** The vectors are then represented by colours, in a way that gray suggested the absence of a bias, and a pure colour corresponds to three standard deviations different from random behaviour.

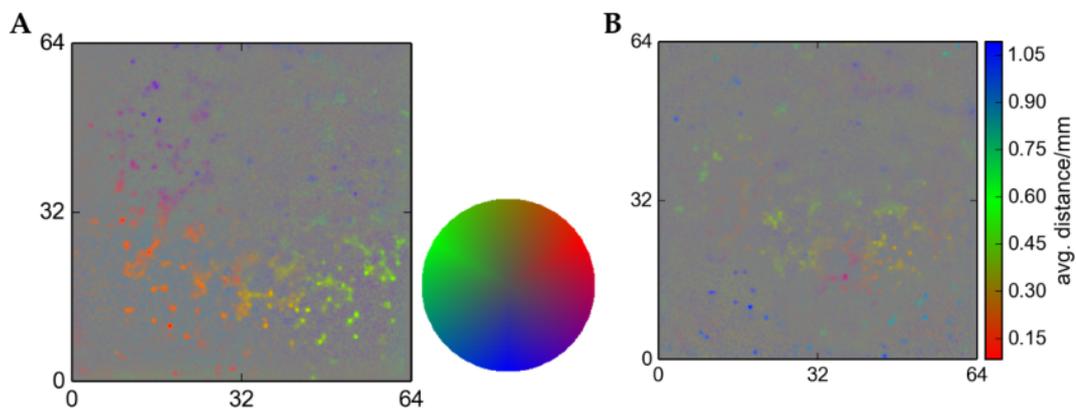


Figure 3.2: Muthmann's analysis of spatial activity in cultured, dissociated hippocampal neurons from mice, conducted using 3Brain's BioCAM4096 platform. **A** The directionality bias of simultaneous neural activity, with various colours on a colour wheel indicating different directions. **B** The distance bias in the activity around the neurons [24].

3.1.2.2 Potential Limitations of the Method

Although Muthmann’s method explores the association between neurons of interest and their surroundings, his approach is limited in three major ways. Firstly, the lack of detail makes his method difficult to interpret, even after sourcing and examining his code. In particular, the multiple stages of spatial averaging are not defined clearly, and his account lacks any justification or explanation for these steps, beyond the intention of “identifying local structures”. This makes it fairly challenging to reproduce his work, and to understand the results obtained.

Furthermore, there are potentially a few structural problems with the method in its current form that make its utility in profiling spatial activity patterns questionable. One problem is that the spatial range used to determine which events are counted as the ‘surrounding activity’ for a given event is very wide, which often results in a directional bias towards the centre of the MEA. For neurons located at the edges of the MEA, this problem is particularly apparent. In Figure 3.3, we see a red X, representing an event occurring at that location. Enclosing the X in a red circle are the surrounding events that are considered as ‘surrounding activity’ by Muthmann’s method – i.e. locations of other events that fall more than 82 μm but less than 1092 μm away from the given event.

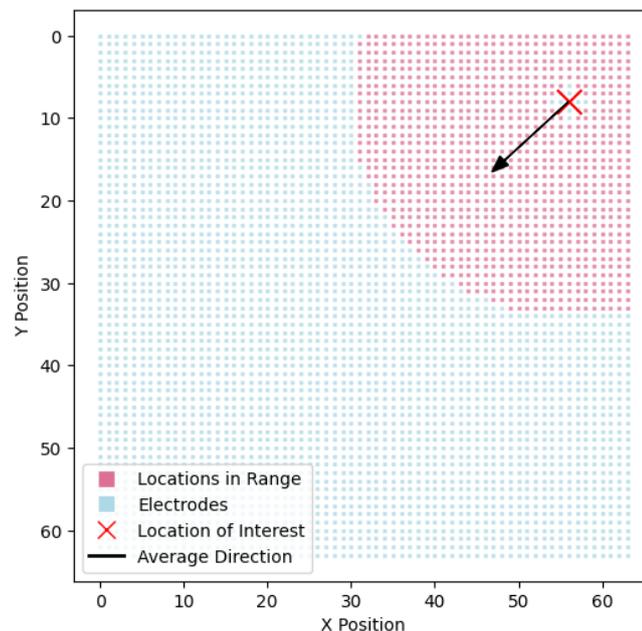


Figure 3.3: Simulated HD MEA displaying 4096 channels (blue squares). Channels ($82 \mu\text{m} < x < 1092 \mu\text{m}$) from ‘X’ are highlighted in red. The black arrow represents the average of the vectors from the red ‘X’ to all locations of interest.

If events were to occur at all of these candidate locations, the average direction of the surrounding activity from the given event would point towards the centre of the MEA. This is seen in Figure 3.3, where the black arrow represents this vector – the average distance and direction of the surrounding activity from the red X. Averaging these average vectors over multiple events throughout the recording that occur at the

same location on the MEA further increases the likelihood that this estimate may be biased towards the centre of the chip. The result of this is that for neurons located on the edges and corners of the MEA, Muthmann's estimate of surrounding activity will almost always be biased towards the centre.

Another limitation of Muthmann's method is the way that it amalgamates activity occurring both before and after an event when considering coincident activity. Consider a scenario where three neurons, A, B, and C, located on a horizontal line, fire in sequence: A, then B, then C, followed by a pause, and further repetition, thereby establishing a clear directional flow from A to B to C. When Muthmann's algorithm evaluates the coincident spatial activity for a particular event at neuron B (referred to as 'B-event'), it examines a certain number X of events before (most of which occur at A), and also X of the events that follow (most of which occur at C). The direction, measured from the positive x -axis of event A from B is 180 degrees, while the direction of event C from B is 0 (or 360) degrees, and the distance between each neuron is equal. This means that the average vector calculated to represent the coincident spatial activity at the 'B-event' would be close to zero, because the vectors to the surrounding activity would effectively cancel one another out when aggregated, telling us nothing of the average direction of signal flow.

The final limitation of Muthmann's method is its approach to data visualisation. To illustrate the directional bias of neurons distributed across the MEA, Muthmann assigns a specific colour to their respective locations. This is visualised using a colour-wheel legend that transitions through blue, red, and green, with grey at its centre. Although these colours may appear distinct to some readers, this specific colour-scheme is not effective for presenting the spatial activity profiles to individuals with red-green colour deficiency – a condition estimated to affect up to 8% of men [9]. Moreover, even for those able to distinguish between these colours, the chosen colours do not intuitively communicate the direction of coincident activity clearly. When attempting to interpret a particular colour on the direction plot, its meaning is unclear without cross-referencing the colour-wheel legend, which itself only vaguely indicates directionality, without specifying precise angles.

3.1.3 Improved Version of Muthmann's Method

Despite the limitations previously outlined, Muthmann's method introduces a valuable approach to exploring neural activity. By considering a spatial profile of surrounding activity for each event instead of looking at pairwise interactions, Muthmann's method has the potential to identify structures in neural data that other approaches, such as Sharf et al.'s functional connectivity method, might miss. With this in mind, several adaptations outlined below were introduced to refine the method.

1. **Temporal Separation:** When considering the 'coincident' activity surrounding a given event, separate it into two distinct categories: surrounding activity that precedes the event, and surrounding activity that follows the event. This resolves the problem described above of opposing vectors cancelling each other when there is a clear, directional flow of information. Adapting the method in this way – to distinguish incoming and outgoing activity patterns for each neuron – enables

a better estimation of functional connectivity within such circuits.

2. **Simplification to Spatial Averaging Steps:** Rather than considering multiple stages of spatial averaging with sliding windows of different widths, the spatial activity profile of activity before and after each event is computed using a single window. This helps to overcome the ambiguity associated with these rather vaguely defined steps, enabling an intuitive approach that is easy to understand and implement.
3. **Narrowed Spatial Range:** The spatial range used to calculate the surrounding activity was narrowed by reducing the upper bound by a factor of two. Previously, Muthmann set the upper bound for this spatial range to approximately 40% of the width of (2.67mm × 2.67mm) MEA – the upper bound should be reduced to 20% of the width of the MEA, or indeed further, if the bulk of the neural population is tightly clustered. This helps avoid the problem described above – where the average direction of surrounding activity for neurons located towards the edge of the MEA is artificially biased towards the centre of the chip.
4. **Vector Visualisation:** To enhance data visualisation and overcome the interpretative challenges posed by the previous colour-coded approach, this method introduces vector-based maps to represent the same underlying statistics. On top of each neuron, an arrow is extended to indicate the average distance and direction of activity preceding or following all events that occurred at the given neuron. This approach offers the reader an intuitive way of understanding the spatial activity profile of neurons recorded using an MEA, negating the necessity of a complex colour scheme that is often inaccessible.

3.2 Synthetic Neuronal Data

Neural data, particularly those obtained from high-density microelectrode array (HD MEA) recordings of brain slices and cerebral organoids are inherently complex. The dynamics underlying such data remain largely unknown, making them prone to misinterpretation by methods that have not been adequately verified. Consequently, it is necessary to validate the methods discussed above against datasets with known and well-understood dynamics to ensure that they can be later applied to actual data with confidence. Although Sharf et al.'s functional connectivity approach is successfully applied to cerebral organoid data, no mention of method validation is made in their paper. Similarly, the adaptation of Muthmann's spatial activity profile method also lacks such validation, highlighting the necessity of testing each method extensively. This validation is not merely a procedural step but a fundamental prerequisite – it gives confidence that the methods are revealing something meaningful in the data.

With this in mind, neural data were simulated to model four well-understood scenarios: randomised data, successive planar waves of propagation, successive ring-like waves of propagation, and successive waves rotating clockwise around a circular pathway. In each case, the synthetic data aimed to mimic spike-sorted neural data obtained from HD MEA recordings. The hypothetical HD MEAs each consisted of 256 channels, evenly spaced on 16 by 16 grids.

3.2.1 Randomised Data

The first scenario simulated was that of randomised data (Figure 3.4). This dataset was designed as a baseline, lacking any of the inherent spatial or temporal patterns that emerge in some of the later more structured synthetic datasets. To simulate the data, spike times were generated for each of the neurons configured in the 16 by 16 grid, intended to mirror the layout of an HD MEA. Each neuron's overall firing rate was determined by a normal distribution with a mean firing rate of 3.45Hz, and a standard deviation of 0.9Hz curtailed below at 0 to avoid a negative firing rate. The simulated recording extended for 60 seconds, a period sufficient to capture representative random activity. By a Poisson process, spikes for each neuron were generated at randomly distributed times within this span while reflecting their assigned firing rates. This process yielded a dataset that was spatially and temporally randomised, lacking any predefined patterns or structure.

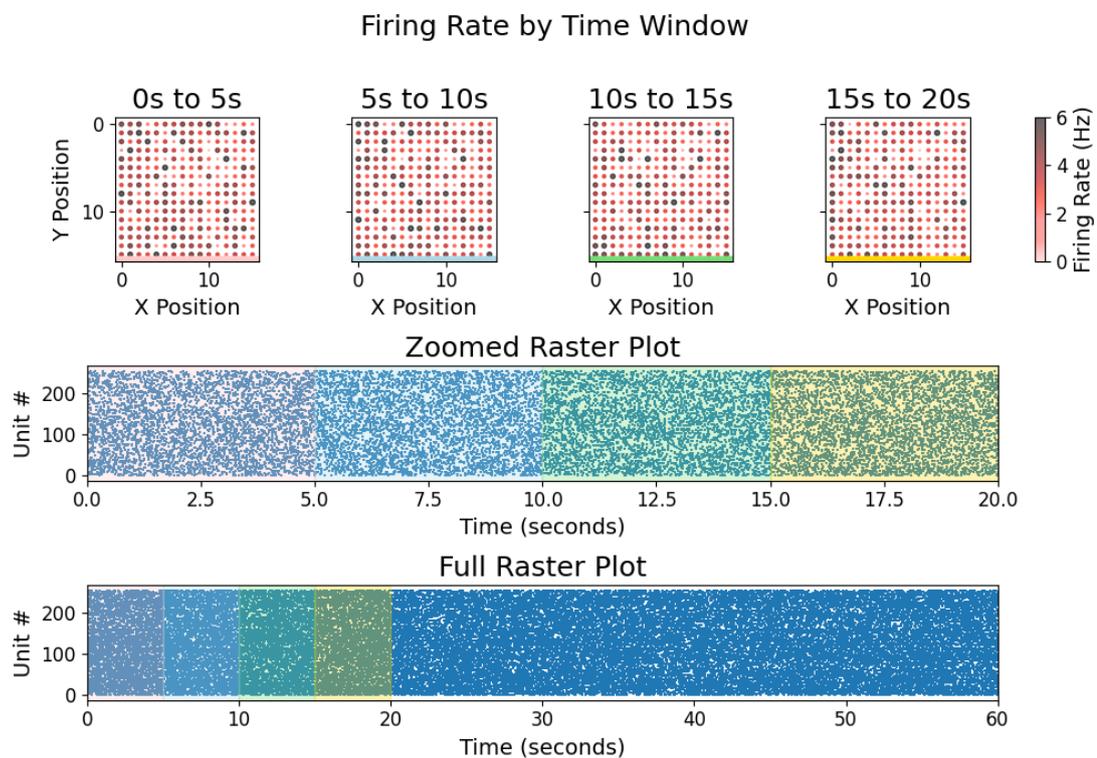


Figure 3.4: Visual representation of simulated neural data reflecting random activity. Top row displays firing rate heat maps by time window, across a 16×16 grid, corresponding to a 256-channel HD MEA. Second row is a zoomed raster plot of neuronal spikes, with each mark representing a spike event, and the coloured bands corresponding to the time windows depicted in the maps above. Third row visualises full dataset using the same colour-coding scheme.

When applied to this dataset, the methods described above are expected to identify no functional connectivity or any other significant patterns, beyond what could be expected by chance. If the methods were to infer strong underlying patterns or functional associations between neurons in this dataset, their utility would be significantly undermined, necessitating further adaptation and re-validation. More broadly, this scenario serves

as a benchmark for further testing of the methods on the following structured datasets, designed to represent more complex neural dynamics.

3.2.2 Plane Wave Data

This structured dataset was designed to reflect coordinated neuronal activity, such as the plane waves in Figure 3.5B, reported by Wu et al. [44]. To achieve this, spikes were generated across the grid in a sequential manner, with each column of neurons in the MEA firing one after the other, creating a plane-wave like pattern of activity across the MEA (Figure 3.6). The simulation ran for a total of 60 seconds, segmented into ten distinct propagation waves. Each slightly slanted wave was structured to traverse the grid from the leftmost to rightmost edge, with a brief inter-wave break period of 2 seconds to demarcate the individual waves. Within each wave, the neurons in a given row were set to fire with an overlap duration that varies around a 100ms baseline by a random amount up to ± 10 ms with the adjacent rows, ensuring a continuous wavelike movement. Within these windows, each neuron's spikes were randomly distributed across the window to introduce further variability.

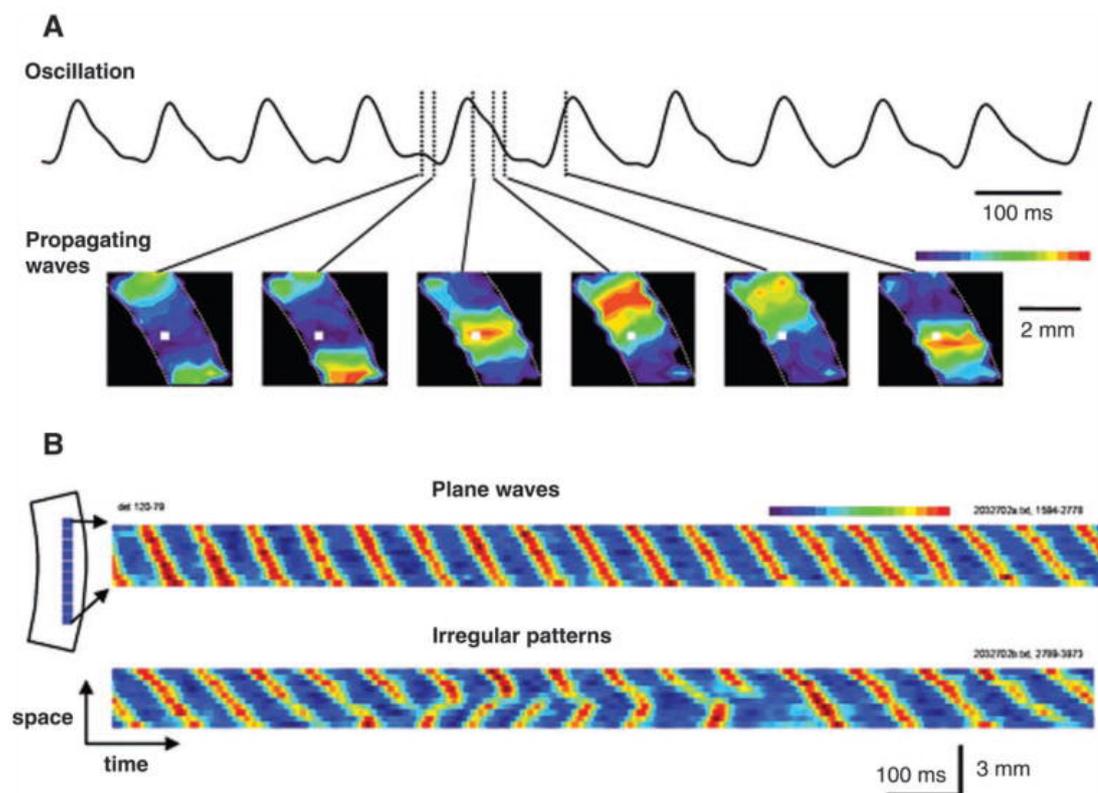


Figure 3.5: In neocortical slices treated with carbachol and bicuculline, spontaneous ~ 10 Hz oscillations occur. **A** Oscillations are captured by voltage-sensitive dyes (VSD), showing a cycle of propagation from bottom to top of the slice, visualised through colour-coded images. **B** A "space-time" map from VSD signals in cortical layers II–III depicts wave interactions during oscillations: initially, waves move downward from the top, later intersecting with waves starting from the bottom, leading to irregular patterns. From Figure 3 of Wu et al. [44]

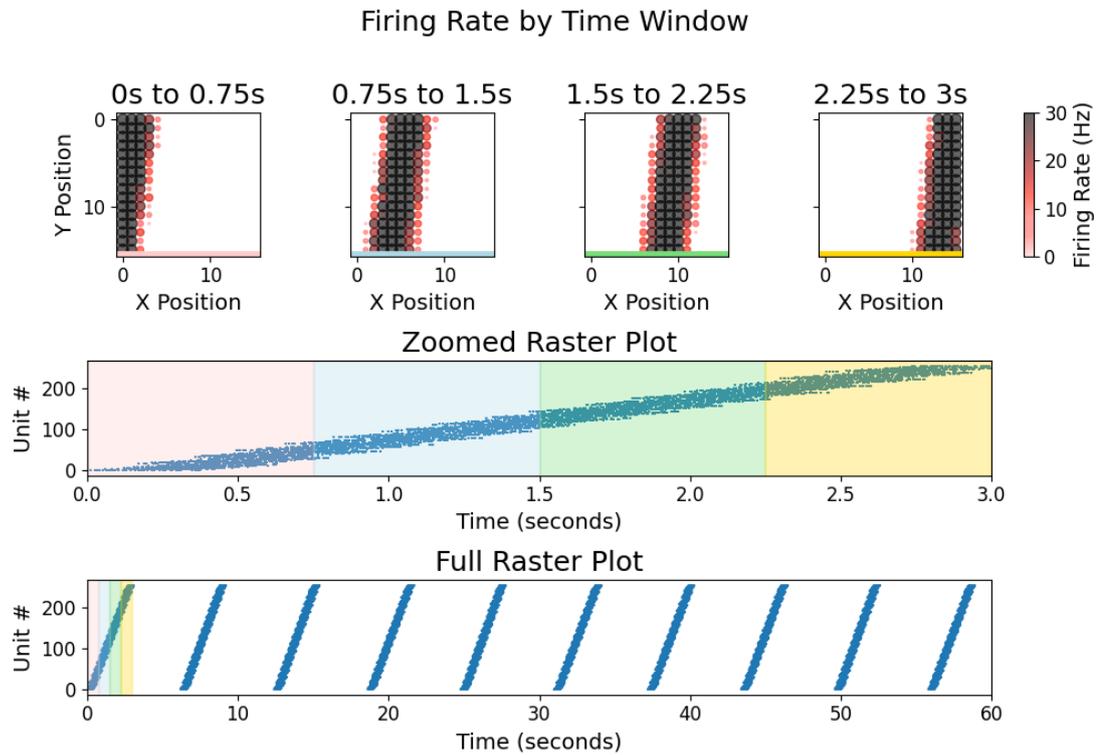


Figure 3.6: Visual representation of simulated neuronal activity featuring repeated waves of linear propagation. Top row displays firing rate heat maps by time window, across a 16×16 grid, corresponding to a 256-channel HD MEA. Second row is a zoomed raster plot of neuronal spikes, with each mark representing a spike event, and the coloured bands corresponding to the time windows depicted in the maps above. Third row visualises full dataset using the same colour-coding scheme.

We would expect the methods described above to pick up on this wave structure in the recording. The adaptation of Muthmann’s method should reveal some amount of directional flow of information from left to right – which should be clear in both the incoming and outgoing spatial activity profile maps, where the average arrow directions should be oppose one another. However, it is possible that the direction of the average preceding and following activity for each neuron might be distorted by the neurons in the same column that fire at the same time. Given the narrowed spatial range considered by this method, it is probable that this effect should not influence the vectors overlaid for the majority of the neurons, apart from a small minority that lie on the far edges of the MEA.

3.2.3 Ring Wave Data

This scenario simulated repeated circular propagations of neural activity, with firing starting in a small circle at the centre of the MEA before widening outward, in a similar pattern to the ripples produced by dropping a stone into a pond (Figure 3.7). Inspired by the “ring wave” pattern observed in rodent hippocampal slices [44], visible in rows 1-2 of Figure 3.8, this dataset is intended to validate the ability of the methods above to

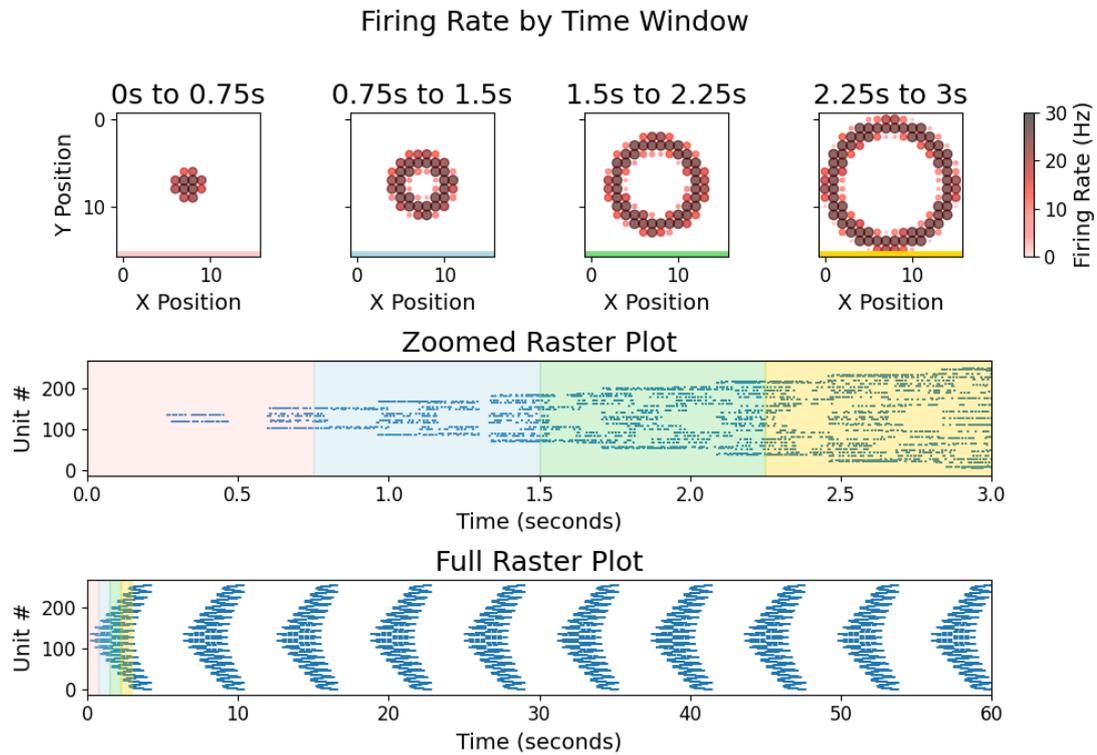


Figure 3.7: Visual representation of simulated neuronal activity featuring repeated waves of circular propagation. Top row displays firing rate heat maps by time window, across a 16×16 grid, corresponding to a 256-channel HD MEA. Second row is a zoomed raster plot of neuronal spikes, with each mark representing a spike event, and the coloured bands corresponding to the time windows depicted in the maps above. Third row visualises full dataset using the same colour-coding scheme.

identify such patterns. Each wave began with a small ring of neurons at the centre of the MEA firing simultaneously, with the firing of subsequent rings of neurons determined by their distance from the centre – neuron rings closer to the centre would fire earlier to create the ripple effect. To smooth the outward-propagating ring wave, there was 200ms (± 10 ms) of overlap between the firing times of neurons adjacent in distance from the centre of the chip. Variability of spike times within these defined windows was also introduced in the same way as above. Again, the simulation ran for a total of 60 seconds, split into ten circular waves separated by a short, 2 second long inter-wave breaks.

It is expected that each of the methods described above will successfully capture to some extent the wave structure and accompanying functional associations between neuron pairs present in the recording. The adaptation of Muthmann's method is expected to visualise directional flow from the centre of the chip to the outer edges, with the average incoming spatial activity profile map pointing towards the centre, and the average outgoing activity map pointing towards away from the centre, to the edges. There is some risk, especially for the smaller rings, that this method mistakenly identifies the average direction of a neuron as being associated with the other neurons in its ring, rather than with the neurons whose firing precedes or follows it. Moreover, if the spatial

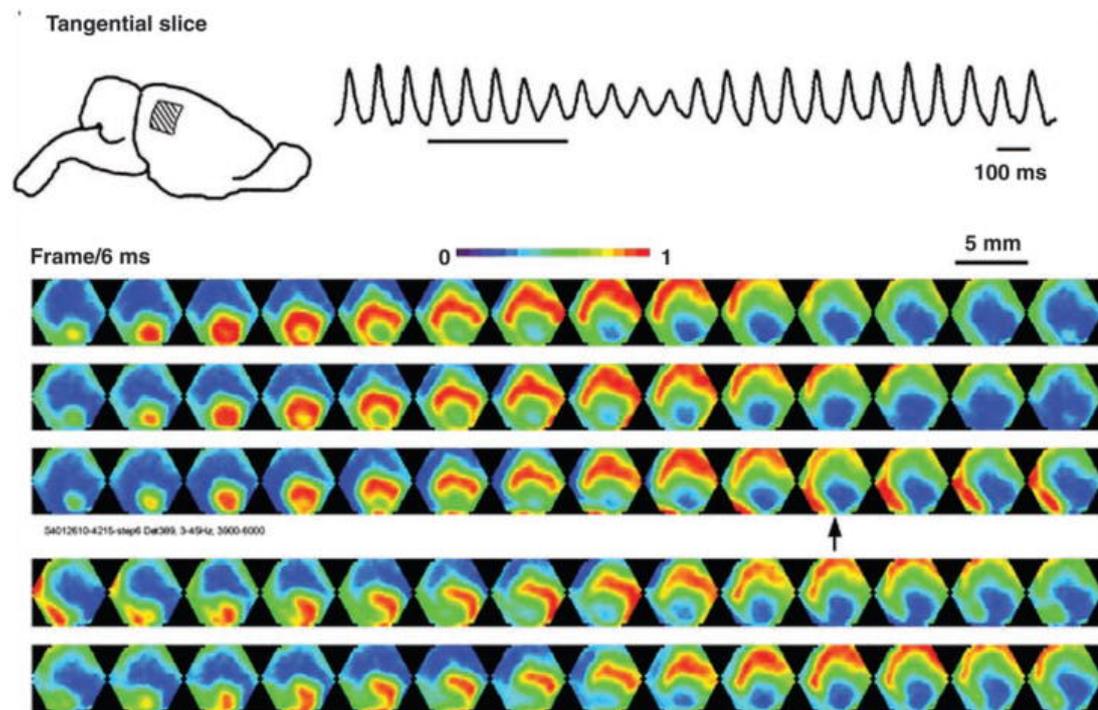


Figure 3.8: In a rat visual cortex slice treated with carbachol and bicuculline, VSD imaging captures wave-to-wave interactions: Rows 1-2 show upward-propagating “ring waves”. The third row marks a collision of two wave fronts, creating a rotating spiral wave. Rows four and five depict the spiral wave’s rotation for 20-30 cycles. From Wu et al. [44].

range window is set large enough, there is a risk that this would cause the direction of an arrow to point toward the centre – repeating the bias towards the centre mentioned in Matthias’ original method. However, it is expected that by setting a narrow enough spatial range, and a wide enough temporal one, that such an effect will be minimised.

3.2.4 Rotating Loop Wave

The final scenario simulated a wave of firing neurons that rotated clockwise around a circle (Figure 3.9). This idea was inspired by the “rotating spiral wave front”, observed by Wu et al., presented in rows 4-5 of Figure 3.8 that occurs after two wave fronts collide. Implementing a spiral-wave rotation proved to be very complex, so a minor simplification was made to simulate a rotating circle of firing neurons instead of a spiral. Indeed, from close inspection of Figure 3.8, we can see that the rotating spiral is not so different from a rotating circle. An overlap of 100ms (± 5 ms) was set between each neuron firing to smooth the rotating wave. As with the two simulations above, a predefined number of spikes were generated within these windows for each neuron. The continuous clockwise rotations ran for a total of 60 seconds.

By offering a simplification of rotating spiral wave-fronts, this simulation is intended to test the ability of the methods to capture repeated rotational information flow. It is expected that the Muthmann’s adapted method will successfully capture this dynamic.

It is also thought that although Sharf et al.’s functional connectivity method will capture the functional associations between the ring of neurons, that it may struggle to offer much value in identifying the direction of these associations.

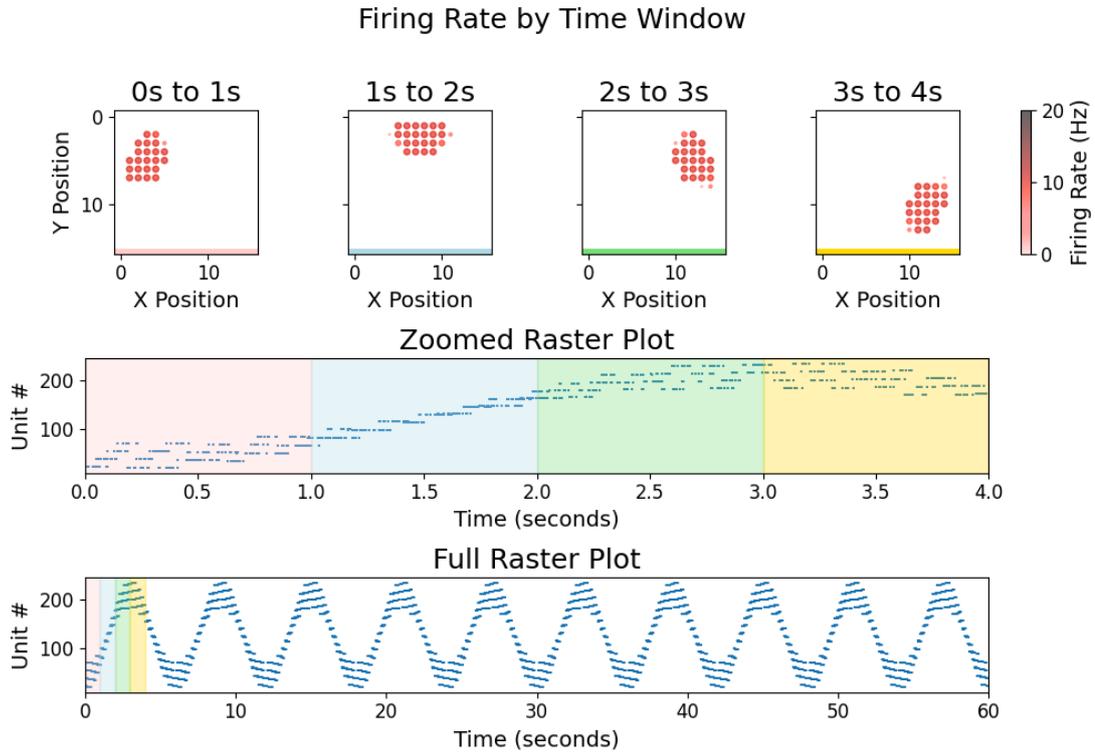


Figure 3.9: Visual representation of simulated neuronal activity featuring a repeated circular loop of firing neurons. Top row displays firing rate heat maps by time window, across a 16×16 grid, corresponding to a 256-channel HD MEA. Second row is a zoomed raster plot of neuronal spikes, with each mark representing a spike event, and the coloured bands corresponding to the time windows depicted in the maps above. Third row visualises full dataset using the same colour-coding scheme.

3.3 Implementation

All code necessary for implementing the exploratory methods and simulating the synthetic data was developed in Python, utilising Jupyter notebooks. Initially, the notebooks were hosted on WetAI, an online platform designed to facilitate collaborative neurobiology research [12], before the code was transferred to GitHub Codespaces, and pushed to the UCSC Genomics Institute’s “braingeneers” repository for future collaboration.

3.3.1 Methods

The methods were implemented in two notebooks – “fcm-figure.ipynb” and “muthmann-figure.ipynb”. At the high level, each file read in simulated or human data, computed

connectivity networks and spatial activity profiles before visualising the results. Although Sharf et al.'s paper provided a link to a GitHub repository containing MATLAB scripts to run the method, its code was largely obsolete and assumed a different data structure to the data available in Table 5.1. As a result, Python code to re-implement Sharf et al.'s functional connectivity method was written almost exclusively from scratch by considering the detailed method description outlined in the paper. When implementing these parts, the code was designed to be modular and reusable – different steps were broken down into reusable functions, with clear and extensive documentation to ensure maintainability.

To implement certain complex details of Sharf et al.'s method – particularly the use of Hartigan's diptest, and full-width-at-half-maximum assessment to assess the modality of spike time latency distributions, Sharf et al.'s publicly available MATLAB script was used as reference [33]. To aid with reading data, calculating STTC and manipulating spike-sorted data, several functions were imported from the University of California Santa Cruz' "braingeneerspy" library [42].

However, in order to analyse the simulated data, it was often necessary to adapt these methods locally. For instance, several of the braingeneerspy functions used to calculate STTC were rewritten locally to bypass division-by-zero issues posed by inactive neurons. On other occasions, braingeneerspy functions were adapted to troubleshoot certain issues – a function was written to calculate the spike time latencies between two individual spike trains to enable the visualisation of spike time latency distributions. This helped to identify an unforeseen problem – that the latency distributions of neuron pairs in the simulated datasets tended to be somewhat discrete (and therefore multimodal) – and fix it by adding randomness to the way spike times were generated for those neurons.

A series of Python scripts to implement some of Muthmann's original analysis were available, but they were difficult to interpret, using outdated libraries and significantly different data structures to what we had been using. Consequently, much of the spatial activity profiles method was rewritten and implemented from scratch, with close reference to the method description in his thesis paper. After evaluating the method, and choosing to adapt certain features, further python scripts were developed to apply the updated method to simulated and human datasets.

When implementing each method, code was designed to maximise computational efficiency. This was an important principle to consider because of the size of the datasets – some of the human datasets for instance, had more than 300,000 spikes, meaning that computationally intensive tasks like ranking events or calculating latency distributions between different spike trains could easily take a long time without careful implementation. A good example of this efficiency is seen in how the angles and distances between events are calculated – the code takes advantage of NumPy's vector operations rather than relying on pure Python, which would be far slower.

3.3.2 Synthetic data

Code to generate synthetic neural data for the four scenarios was generated in the “data-generator.ipynb” notebook. The four distinct functions to generate each type of dataset were written from scratch, with a variety of changeable parameters – such as grid size, recording duration, overall firing rate, and wave overlap duration. This flexibility not only served the immediate objective of generating data to test the methods, but it also paved the way for future explorations and adaptations. Following this generation process, the datasets were serialised using the “pickle” module, ensuring their preservation and accessibility for subsequent analysis.

Further python scripts were developed to visualise the simulated neural data. For each dataset, the visualisation process involved plotting firing rate maps across four time windows and generating two raster plots – one zoomed to the four time windows, and another for the full recording. This code used Matplotlib’s pyplot library to ensure consistent colour schemes across these subplots – with the four time windows being assigned a colour that was then shaded in the raster plots below. This code was designed to be flexible and reusable, accommodating varying numbers of time windows and data with different firing rates.

Chapter 4

Testing on Synthetic Data

4.1 Performance of Methods on Synthetic Data

4.1.1 Sharf et al.'s Functional Connectivity Method

On the whole, the functional connectivity networks depicted in Figure 4.1, which are derived from the four simulated datasets, successfully capture many of the underlying dynamics inherent within these datasets. As expected, application of the functional connectivity method to the (baseline) randomised dataset revealed a complete absence of functional associations or discernible patterns. This is seen in the top left corner of Figure 5.1, where the lack of visible edges connecting any neuron pairs indicates the absence of any underlying associative structure within the data.

By applying the functional connectivity method to the plane wave dataset, the network visualised in the upper right corner of Figure 5.1 was derived. This network successfully captures the underlying structure in the data – the propagation of planar waves from left to right across the MEA with a slight forward slant akin to a '/' symbol. The algorithm correctly classifies the two leftmost columns of neurons – and the upper neurons of the adjacent third column – as the primary wave 'senders', and the two rightmost columns – and the lower neurons of the third column from the right – as 'receivers', positioning the central neurons as 'brokers' in the wave propagation process. This classification is a success – it reflects the underlying pattern of repeated planar wave movements which start at the left of the MEA and end at the right, demarcated by 2 second breaks. However, it is arguable that the network over-captures functional association between these neurons, particularly between neurons in the same column (reflecting their simultaneous firing) – seen in the large number of overlapping edges connecting neuron pairs.

The application of the method to the ring wave dataset derives a network visualised in the lower left corner of Figure 4.1. This network also successfully captures much of the structure present in the ring wave filled data. Specifically, the central neurons are classified as 'senders', while the outer ring of neurons are classified as 'receivers' – clearly representing the outward-propagating ripple-like pattern encoded in the data, where firing begins at the centre of the MEA and ends on the edges.

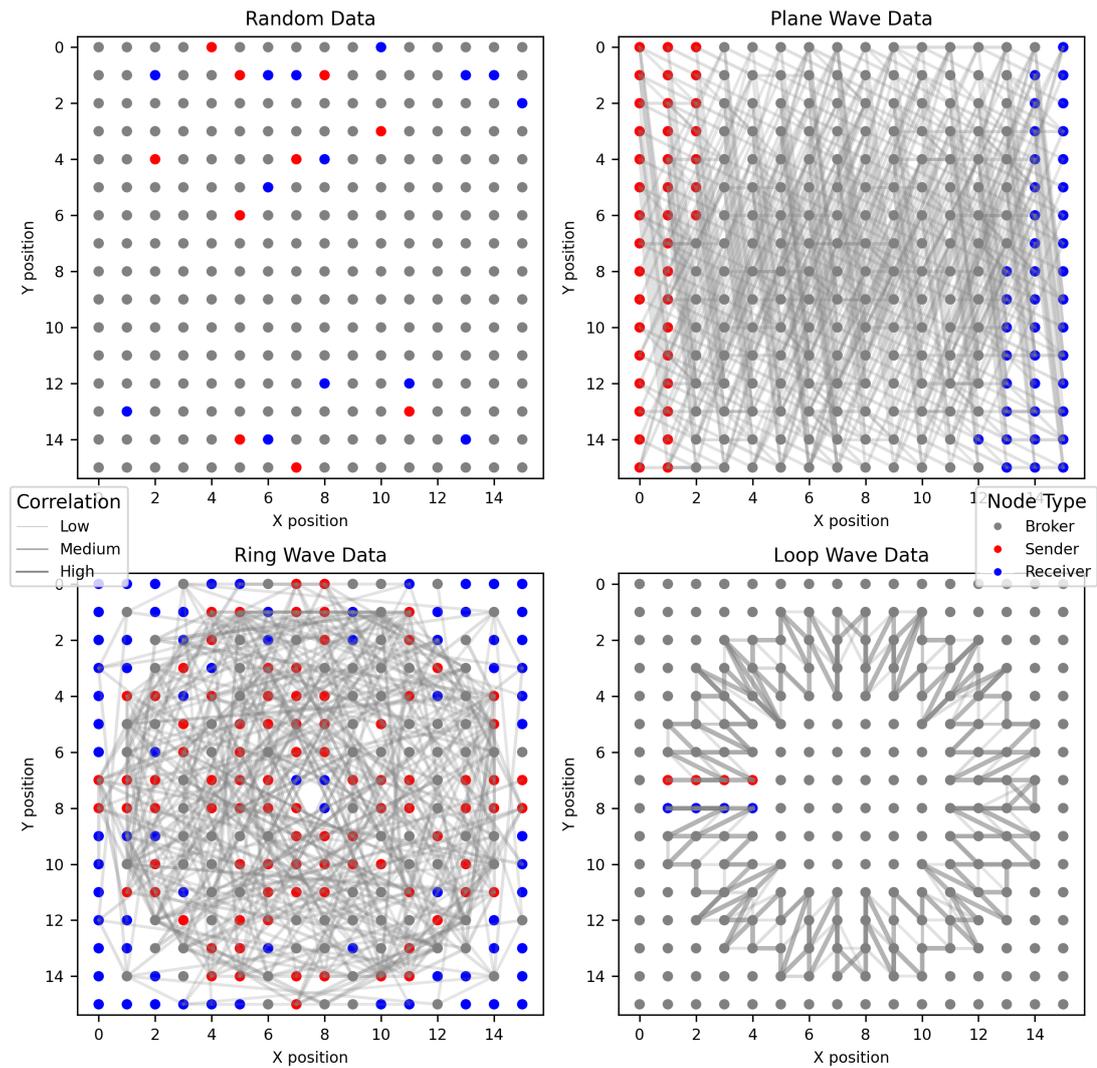


Figure 4.1: A series of two-dimensional maps visualising the functional connectivity networks generated by applying Sharf et al.’s method to the simulated datasets. Colour of neurons indicates their functional classification by the method, with edge opacity indicating association.

Notably, the sender-receiver classification exhibits some spatial variability, with a few of the central neurons are classified as ‘receivers’. However, this can be attributed the deliberate introduction of randomness to the spike times of neurons within their designated firing windows, as detailed above in the methods section. Again, the network likely over-captures functional associations between these neurons – visible through the large number of edges connecting different pairs of neurons.

Finally, applying Sharf et al.’s method to the clockwise-rotating continuous loop wave data derived the network visualised in the lower right of Figure 5.1. This network captured some of the structure in the dataset, but not quite as successfully as it did for the other datasets. In contrast to the other wave-propagation datasets that had inter-wave breaks, the waves in this dataset were continuous – with ten back-to-back full rotations

in a 60-second period. Interestingly, the network still captures the three neurons that the firing begins at, classifying them as ‘senders’, as well as the three neurons that the firing eventually ends at, classified as ‘receivers’. Following the functional connectivity edges from these ‘senders’ to ‘receivers’ (as the red and blue neurons are not directly connected) leads the reader in a clockwise direction – which is indeed present in the data. However, this direction of information flow is not immediately intuitive from looking at the map.

4.1.2 Adaptation of Muthmann’s Spatial Activity Profile Method

When applied to the simulated datasets, the adaptation of Muthmann’s method produces average spatial activity profiles for the incoming and outgoing activity for each neuron (Figure 4.2), and by doing so, successfully reveals much of the structure present in each dataset.

For the randomised dataset (top row of Figure 5.2), the vectors overlaid on each neuron represent the incoming and outgoing spatial activity profiles averaged across all events. As anticipated previously, these vectors are short, and relatively random in direction, indicating that there is no underlying structure present in the data.

When applied to the plane wave data, the adaptation of Muthmann’s method reveals a clear flow of neuron firing from the left of the MEA towards the right. This is evident in the second row of Figure 4.2, where the vectors representing the spatial profile of the average activity before all events at each neuron point towards the left, and slightly upward. This indicates that for events at each neuron, the typical activity that preceded those events came from that direction. In turn, this links clearly to the structure in the data, where waves propagate from left to right, with a slight forward slant. This movement dynamic is further confirmed by the ‘average activity after’ vectors which, on the whole, point to the right. However, although these arrows point slightly downward (see the leftmost column), they do not point down as much as the arrows point upward on the ‘before’ map – perhaps due to the randomness added to the spike times.

Applying the method to the ring wave dataset produces the pair of visualisations present in the third row of Figure 5.2. Both the ‘before’ and ‘after’ activity maps show very clearly the average direction of information flow from the centre of the MEA to the edges. The vectors for each neuron on the ‘before’ map extend to point toward the centre, indicating the average of their where the events that precede their events occur, while the vectors indicating the average direction of following activity point clearly towards the outer edges of the MEA. Again, the vectors in both maps form a slightly imperfect, asymmetric pattern due to the addition of randomness into spike times of the neurons within the simulated data.

The rotating loop wave structure present in the other dataset is clearly captured by the adaptation of Muthmann’s method, and seen in the final row of Figure 5.2. Here, the vector extending from neurons in the ‘average after activity’ map point toward neurons that tend to fire after they do, forming a clear clockwise flow. Similarly, the ‘before activity map’ vectors form an anticlockwise loop, helping to reveal the clockwise propagating flow of neural firing in the simulated dataset.

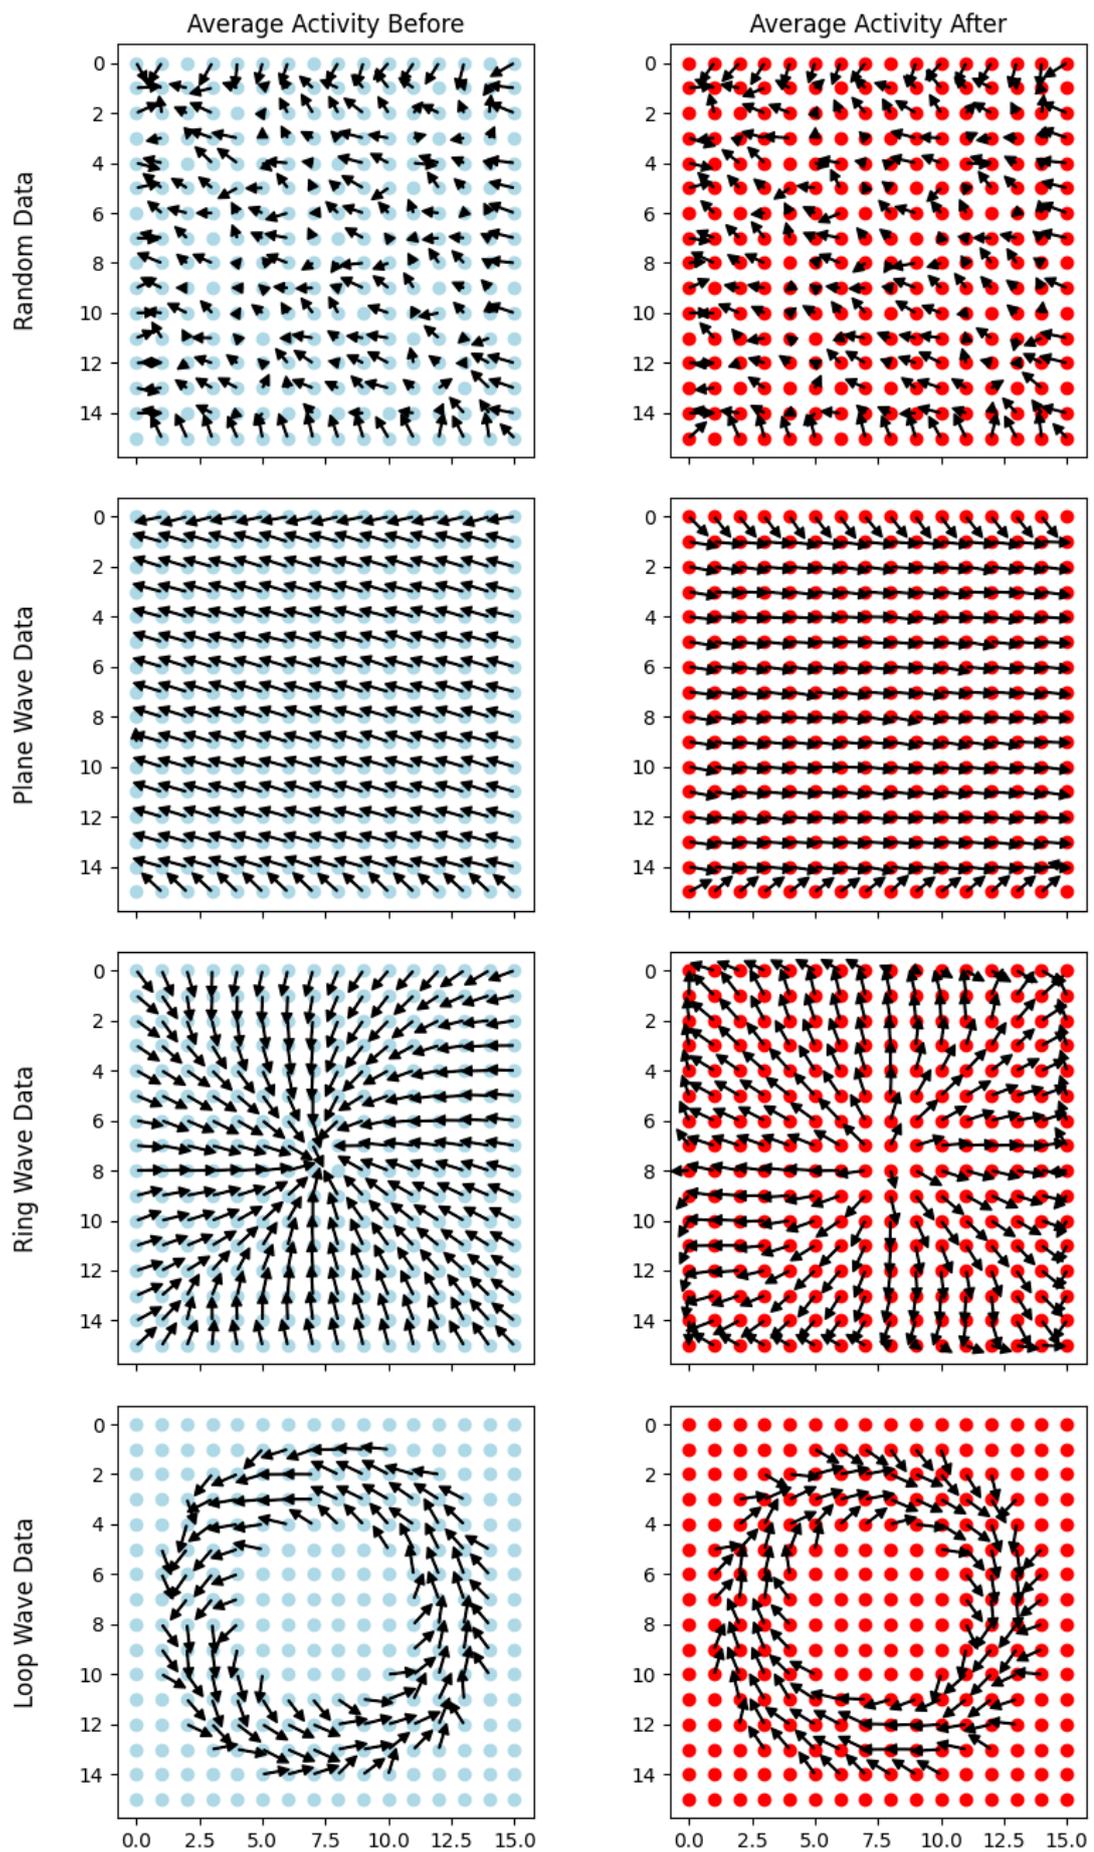


Figure 4.2: In each row, the pair of panels depict the average vector representing to the average incoming and outgoing spatial activity profile for each neuron, as calculated by Muthmann's method. Each row corresponds to one of the four simulated scenarios.

4.2 Evaluation

From this testing, it is clear that both methods are largely successful in identifying the types of coordinated neural activity (or lack thereof) present in each of the simulated datasets. However, from this analysis, it is arguable that the adaptation of Muthmann's method has two desirable properties. Firstly, it is clearly better at visualising representing the structure of the simulated data – rather than having to interpret the meaning of different coloured neurons, and different overlapping edges as with Sharf et al.'s method, the vector arrows used to convey the structure in Muthmann's adapted method are clear and intuitive for the reader. Moreover, it is possible that the refined version of Muthmann's method is more adept than Sharf et al.'s at capturing certain underlying dynamics in the data – particularly those that feature continuous patterns of neural firing. For instance, although Sharf et al.'s method technically captures the direction of the continuously clockwise-rotating loop waves correctly, the distribution of its connectivity edges suggest it is placing unnecessary focus on the associations of neurons that lie on the same radii. This contrasts to the adaptation of Muthmann's method which straightforwardly identifies that rotating pattern, and suggests that when applied to other similar but more sophisticated continuous rotating structures Sharf et al.'s method may fail to correctly identify the underlying pattern.

On the other hand, it is worth pointing out that the utility of the adaptation of Muthmann's method may be limited by the narrow width of the spatial window of activity it considers. While it has been shown to be highly effective at identifying these continuously moving patterns of neural firing, it is possible that for other more complex patterns, perhaps those that involve strongly associated neurons located on opposite sides of the MEA, this narrow width would limit the ability of Muthmann's adapted method to pick up on this structure. In contrast, Sharf et al.'s method ought to be able to pick up on such a dynamic more effectively – as it simply considers the temporal association of the spike trains of each pairs of neurons. Overall, it is clear that although imperfect, each method has distinct advantages – Sharf et al.'s is able to identify large-scale associations, and Muthmann's adaptation can capture local continuous dynamics. Together, these methods could be used in a complementary way to identify structure within real neural data.

Chapter 5

Results

5.1 Overview of Human Hippocampal Data

To explore how information travels through the human brain, organotypic slices from human hippocampal tissues were obtained. These samples, sourced from patients undergoing surgery for treatment-resistant epilepsy, were resected by John Andrews, MD, at the University of California San Francisco's Department of Neurological Surgery. The human tissue samples were collected at UCSF with signed patient consent and full ethical approval from the UCSF Institutional Review Board [4]. After acquisition, the specimens were sliced to 300 μm thickness and cultured in an air-liquid interface within a serum-free medium. For recording, the slices were plated on an HD MEA and oriented to maximise the detection of extracellular activity. Specifically, the CA1 and Dentate Gyrus regions were centred over the MEA's 26,000 channels to capture spontaneous activity occurring after 5 days of culture. However, the MEA surface area was limited to 2100 μm \times 3850 μm , meaning it was not possible to obtain a recording across the entire Dentate Gyrus. Post-recording, the exact anatomical locations of the recorded circuits were confirmed using immunohistochemical staining for neuronal nuclei (NeuN) – see Figure 5.1.d.

The raw activity data from these recordings was processed by a bandpass filter ranging 300-6000Hz and then spike-sorted by Kilosort2 into single unit activity [8]. The resulting spike-sorted datasets were then manually curated using Phy GUI by Jinghui Geng, an electrophysiology researcher at UC Santa Cruz, based on the shape of each unit's quality metrics which include waveform shape, firing rate and interspike interval distribution.

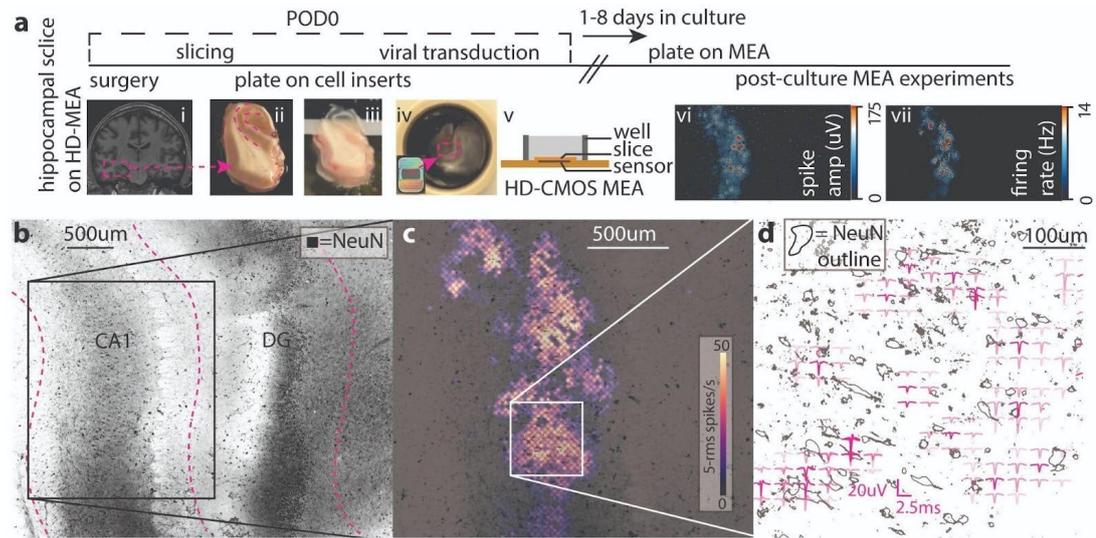


Figure 5.1: **a** Workflow for recording human hippocampal slices using HD MEA. **b** Immunohistochemistry of NeuN marker. **c** Heatmap of spike activity of square box in **b**. **d** NeuN+ neurons represented by black outlines, confirming anatomy of recorded circuits. Taken from John Andrews et al. [4]

Name	Region	# Neurons	Duration (s)	Mean Firing Rate (Hz)
HC_328	CA1/Subiculum	136	398	1.92
HC_52	CA1 Posterior	100	107	2.02
HC_111	DG - Apex	87	418	1.36
HC_120523	DG - Outer Blade 2	51	300	2.38
HC_112823	DG - Outer Blade 1	249	313	4.34
HC_112823_2	DG - Outer Blade 1	351	126	3.46

Table 5.1: Statistical summary of hippocampal datasets.

Table 5.1 provides an overview of the hippocampal location for six of the slices obtained, along with corresponding statistics of their recordings. Of these, two datasets, identified as “HC_328” and “HC_52”, were obtained from the CA1 region, while the remaining four derived from the Dentate Gyrus region. As an illustrative example, Figure 5.2 visualises the neural activity recorded from the outer blade of the Dentate Gyrus (“HC_112823”) through firing maps segmented by time window and corresponding raster plots. Equivalent visualisations for the other five datasets are available in Appendix A.2.

As mentioned in Section 2.2, while HD MEAs have been used to capture extracellular activity from iPSC-derived human neurons [31, 32] and animal tissue slices [26, 24], there is a notable absence of literature on employing HD MEAs to record neural activity from human brain slices. It is in this context that the organotypic human slice datasets presented in Figure 5.1 are valuable – they offer a unique and highly detailed view of the neural circuitry within one of the brain’s most important and connected regions, the hippocampus. By applying the methods described in Section 3.1 to these datasets, functional connectivity networks and spatial activity profiles can be computed for the

sub-regions of the hippocampus present in Table 5.1, to reveal patterns of association within these networks. More broadly, an improved understanding of connectivity within (or in future, across) hippocampal subregions could directly benefit research into neurodegenerative disease – indeed, the functional connectivity of the hippocampus is understood to reduce with Alzheimer’s disease [2].

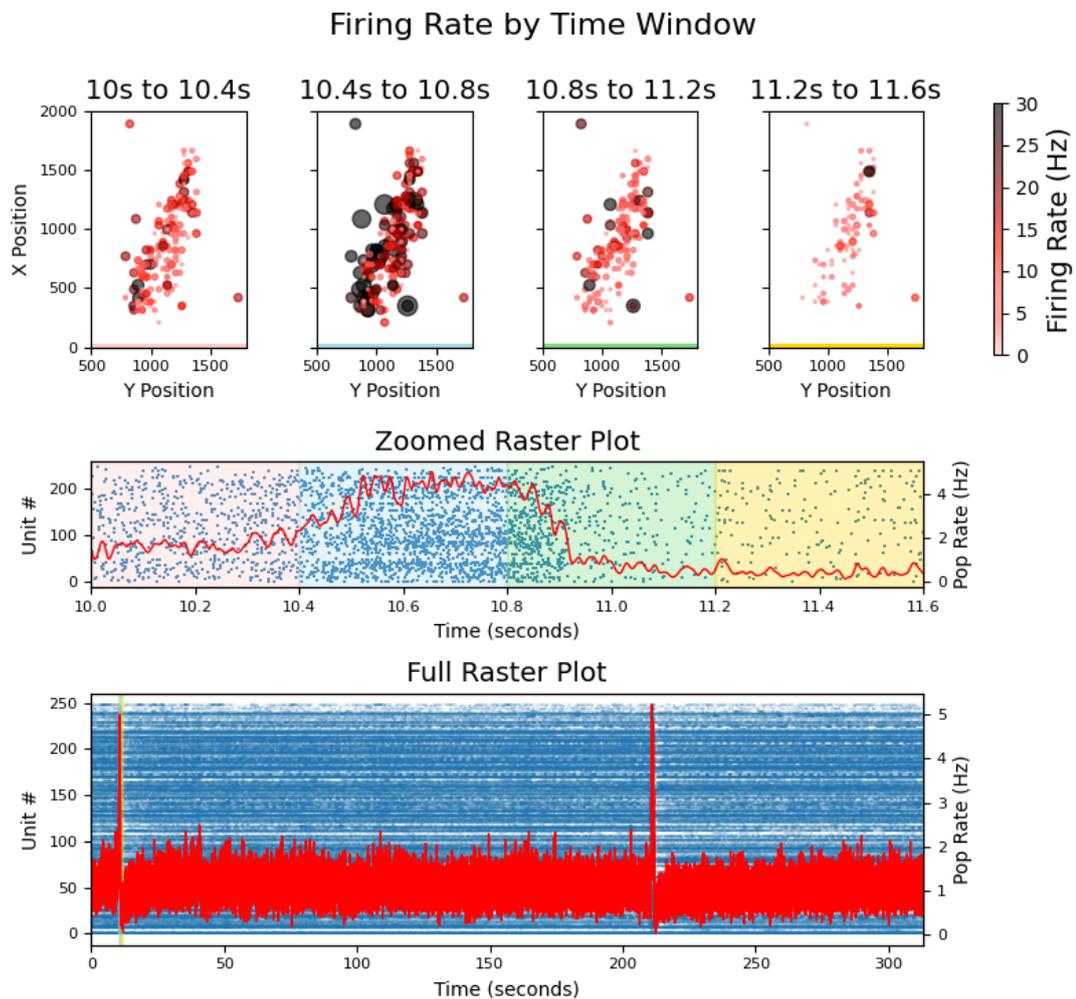


Figure 5.2: Visualisation of spike-sorted extracellular recording ‘HC_112823’ from the outer blade of the dentate gyrus region of the human hippocampus. Top row displays firing rate maps, deliberately rotated. Second row is a zoomed-raster plot of the firing rate windows above. Final row is a raster illustrating full dataset, with same colour-coding scheme.

5.2 Sharf et al.’s Functional Connectivity Method

Applying Sharf et al.’s method to the hippocampal datasets from Table 5.1 produced six functional connectivity networks, visualised in Figure 5.3. This connectivity analysis

revealed variability across slices in the proportion of neuron pairs connected as well as the classification of nodes as ‘brokers’, ‘senders’, and ‘receivers’, detailed in Table 5.2.

Notably, the datasets from the outer blade 1 of the Dentate Gyrus (titled “HC_112823” and “HC_112823_2”) exhibited the lowest percentages of total neurons classified as ‘senders’ and ‘receivers’ (at 0.4% and 0.0% respectively) – a phenomenon clearly visible in the final row of Figure 5.3, in which very few neurons are assigned a colour. This is attributed to the high mean firing rates and neuron count observed in these datasets, as seen in Table 5.1. A high mean firing rate means that neurons both send and receive more signals, thus distributing the number of incoming and outgoing connections more evenly, which likely leads Sharf et al.’s algorithm to classify a greater proportion of neurons as ‘brokers’.

Network analysis of the dataset “HC_52” from the Posterior CA1 region revealed the highest proportion of connected edges at 9.54%. Additionally, this dataset exhibited the highest percentages of neurons as ‘senders’ and ‘receivers’ at 18% and 15% respectively. This observation broadly aligns with the findings of Sharf et al. who classified neurons in fourteen datasets from six iPSC-generated human brain organoids and reported $15.4\% \pm 2.6\%$ senders and $21.9\% \pm 6.5\%$ receivers [33]. It is arguable that this similarity validates the observed connectivity patterns for this Posterior CA1 region, while offering a bridge between iPSC-derived models and actual human brain tissue.

“HC_120523”, the dataset from the outer blade 2 of the Dentate Gyrus, exhibited the next highest proportion of connected edges at 5.54%, with 13.74% of these connected nodes classified as ‘senders’ and ‘receivers’ respectively. These datasets follow the trend that as the percentage of pairs connected increases, the percentage of neurons labelled as ‘brokers’ decrease, with a corresponding increase in the percentage of neurons classified as ‘senders’ and ‘receivers’ – visualised in Figure A.8 in Appendix A.3.

Dataset “HC_328”, from the CA1/Subiculum region presents a distinct functional connectivity network. Figure 5.3 depicts a pronounced linear arrangement of red neurons representing ‘senders’, paralleled by a similarly extended, albeit more fluctuating, line of blue ‘receivers’. This pattern may indicate a preferential signal flow from the red ‘sender’ neurons to the adjacent blue ‘receiver’ neurons. However, despite the apparent organisation in signalling roles, actual connectivity is sparse. As Table 5.2 underscores, only 2.11% of neuron pairs are connected.

Finally, the “HC_111” dataset from the Dentate Gyrus Apex reveals strong connections across the horseshoe-like structure. Here, ‘sender’ and ‘receiver’ neurons on the left side connect strongly to ‘broker’ neurons on the right side. Notably two brokers on the right side establish connections with another broker neuron situated at a significant distance above this structure. However, again, this network is selective – only 2.7% of neuron pairs are connected by the algorithm.

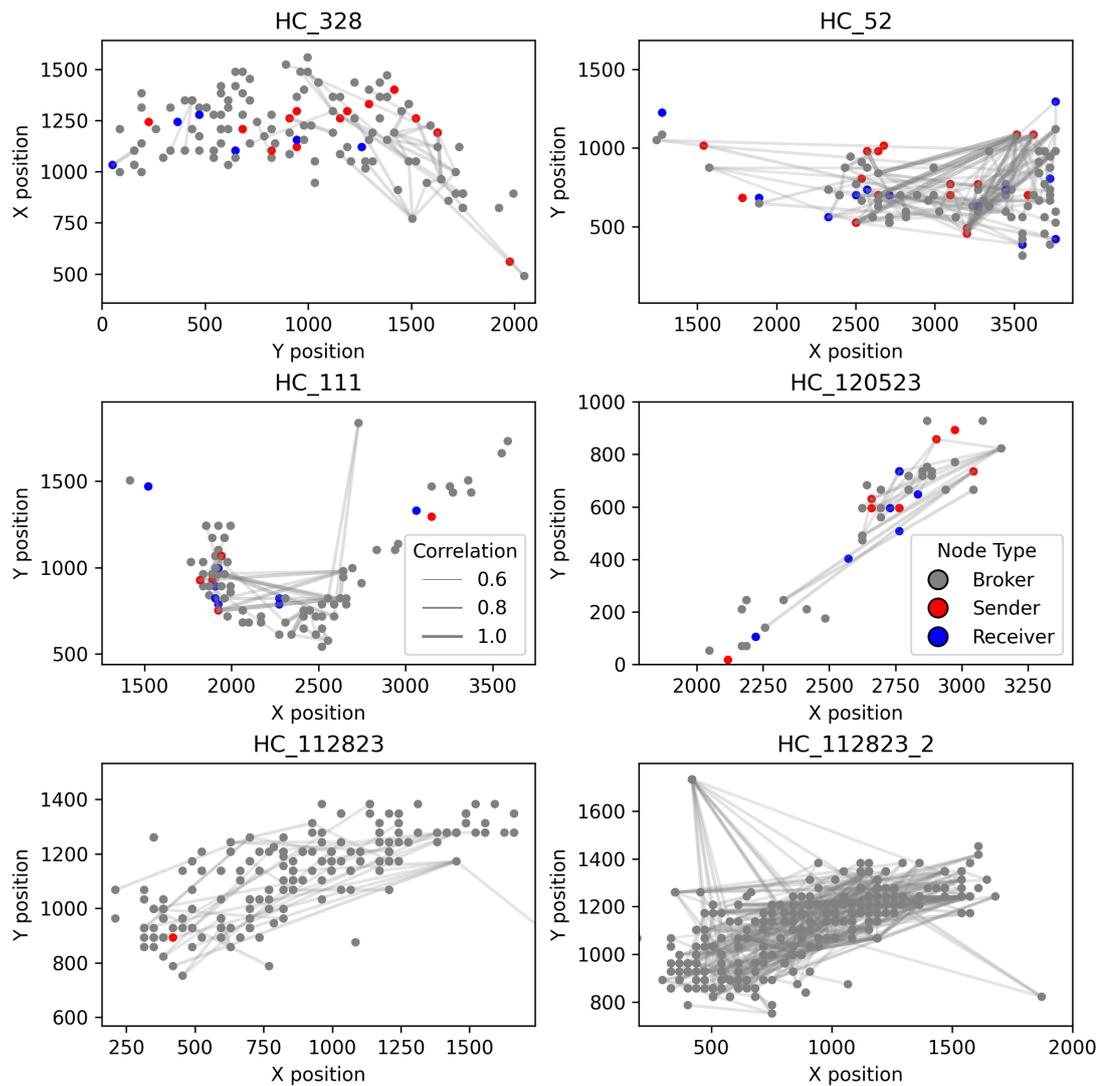


Figure 5.3: Two-dimensional functional connectivity maps to visualise the flow of information through the six of the brain slices recorded from are presented. The top row maps connectivity for the CA1 region, while remaining rows map connectivity in different regions of the Dentate Gyrus. Refer to Table 5.1 for exact locations.

Dataset	% Pairs Connected	% Broker	% Sender	% Receiver
HC_328	2.11	82.35	11.76	5.88
HC_52	9.54	67.00	18.00	15.00
HC_111	2.70	81.61	9.20	9.20
HC_120523	5.54	72.55	13.73	13.73
HC_112823	1.31	99.60	0.40	0.00
HC_112823_2	2.97	100.00	0.00	0.00

Table 5.2: Pair connectivity and neuron classification percentages in hippocampal datasets.

5.3 Adaptation of Muthmann's Spatial Activity Profile Method

By applying the adaptation of Muthmann's method to the hippocampal datasets in Table 5.1, six temporally distinct spatial activity profiles were generated. These profiles provide a comparative analysis of the average neural activity before all events (thought of as the incoming activity), and average activity after all events (outgoing activity), resulting in the visualisations shown in Figure 5.5.

The 'after' spatial activity profile map of "HC_52" displays a distinctive pattern, with arrows on the left side of the horseshoe shape moving downward and then to the right, following the shape's contour. In contrast, the 'average before' map shows a few arrows pointing in the opposite direction, indicating that, on average, the signal in that part of the "HC_52" flows as depicted in the 'after' map. This observation aligns with the functional connectivity map in Figure 5.3, where a cluster of red sender neurons appear at the same location, including one neuron with a strong connection to another neuron positioned below. However, the limited upper spatial bound used to calculate this profile prevents the after map from capturing the robust connection seen in the functional connectivity map between neurons on opposing sides of the horseshoe.

For "HC_120523", derived from the outer blade 2 of the human hippocampus, the spatial activity plots, for both before and after averages, show a significant portion of neurons with overlaid vectors (47.25% and 39.22% respectively). However, this finding is likely explained by the slice's relatively low neuron count of 51. Moreover, despite the application of the updated version of Muthmann's algorithm, the anticipated mitigation of the centre-bias issue, as described in the Section 3.1.2.2 was not achieved. This is evidenced by Figure 5.5, where the polar histogram for "HC_120523" 'average preceding' spatial activity (depicted in blue) closely mirrors the 'average proceeding' activity's outline. This similarity between the profiles of outgoing and incoming activity suggests that the spatial range parameter's upper limit was set too high – meaning that too large an area of activity was considered when computing the spatial activity profile.

However, beyond these observations – particularly that of dataset "HC_52", Muthmann's updated algorithm fails to identify any other distinct dynamics in the data, nor does it corroborate any of the other phenomena described above by Sharf et al.'s method. Generally, either too few neurons are overlaid by vectors to observe any notable patterns, or the spatial range appears to be set too high causing vectors in both the before and after maps to point in the same direction – towards the centre. Different parameter settings for both the upper spatial range and the rank window were tested – both by manually exhausting a range of values, and then by adjusting the spatial range in proportion to the median distance between neurons, and the rank window in inverse proportion to the mean of the median interspike intervals for each unit, but the problem of either too few connections being connected to infer anything, or the centre-bias being present persisted.

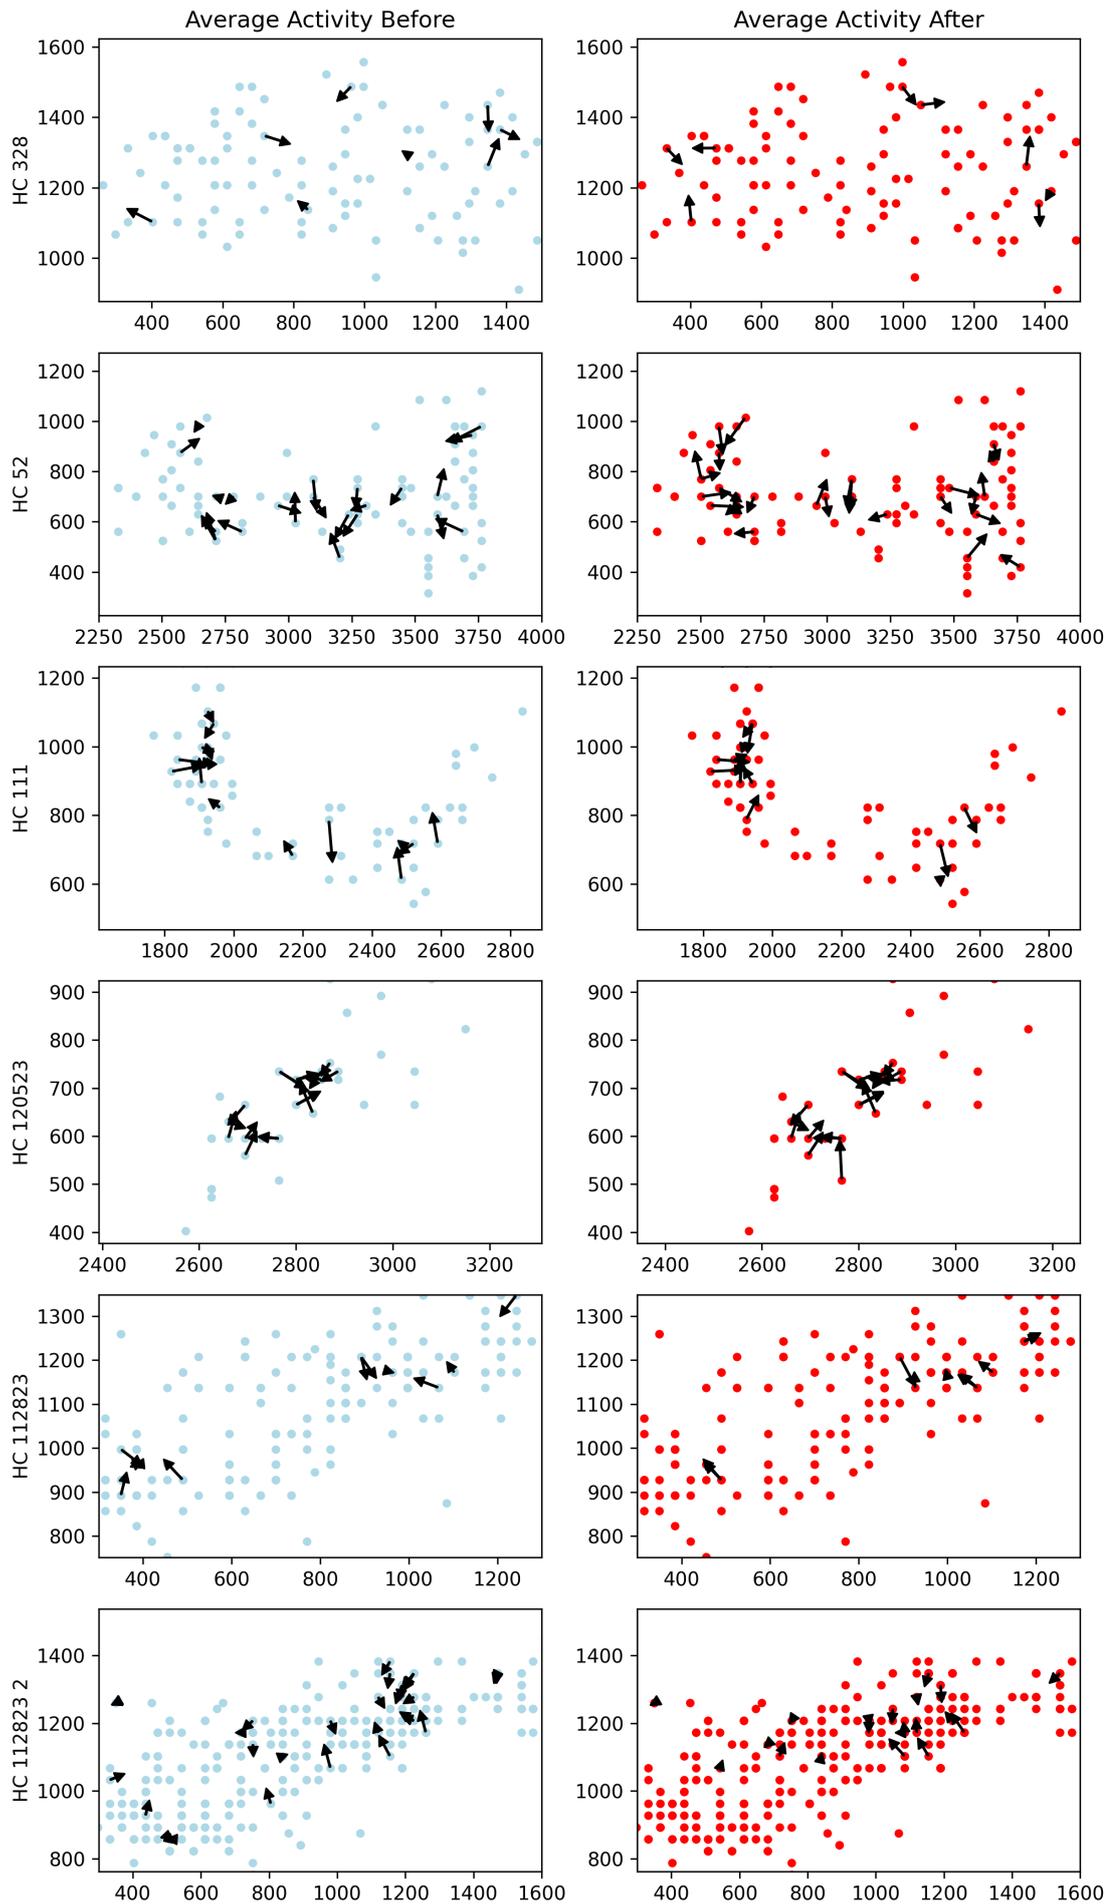


Figure 5.4: In each row, the panel pairs depict the average vector representing the average incoming and outgoing spatial activity profile for each neuron. Each row corresponds to different slice dataset.

Dataset	% Neurons Vectorised Before	% Neurons Vectorised After
HC_328	7.35	5.88
HC_52	23.00	27.00
HC_111	17.24	17.24
HC_120523	37.25	39.22
HC_112823	4.82	4.02
HC_112823_2	9.12	5.98

Table 5.3: Percentage of neurons with a vector in the average spatial activity profile preceding and following the observed events across different datasets.

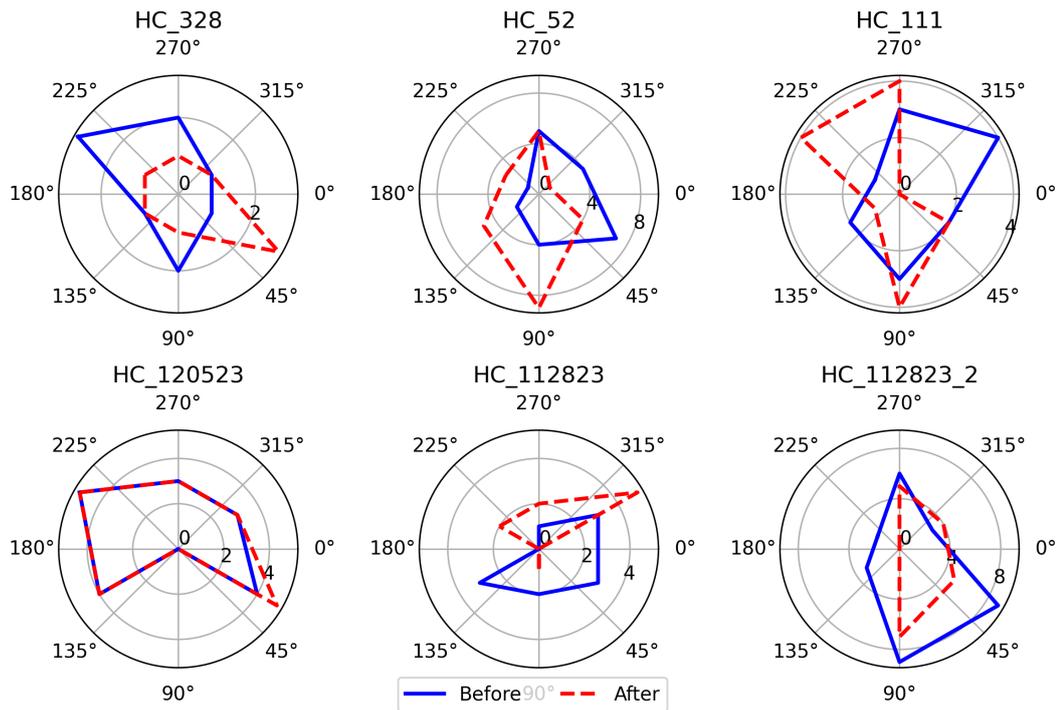


Figure 5.5: Comparative polar histograms representing the directional distribution of the vectors overlaid on the 'before' and 'after' spatial activity plots presented in Figure 5.4.

Chapter 6

Discussion

6.1 Method Evaluation

The ‘functional connectivity method’ developed by Sharf et al. has several clear strengths. Firstly, it has been validated through peer review and used in a publication [33] for its ability to map information flow through cerebral organoids successfully. Here, it has demonstrated its efficacy again by mapping functional associations between neurons in real human brain slices. Moreover, visually, the method is intuitive – a reader can look at the two-dimensional mapping of neurons with connecting lines to indicate association, and relate it to what they imagine neurons in a real organoid or brain slice to look like.

However, one could argue that the classification of neurons into ‘brokers’, ‘senders’ and ‘receivers’ may oversimplify underlying structure in the data. For example, in datasets like “HC_112823” and “HC_112823_2”, the classifications of all neurons as ‘brokers’ seems to overlook more complex interactions potentially present between these neurons. Also, given the method relies heavily on the STTC – a metric that does not estimate direction – significant further analysis is necessary to infer the direction of association between pairs of neurons. An improved approach might employ a directed metric, such as “cross-correlation” [30] to directly infer both association and direction between neuron pairs. This could simplify the method significantly, reducing the need to consider the unimodality of the spike time latency distribution for each neuron pair – a complex process involving Hartigan’s diptest and other width-related assessments. Alternatively, incorporating a model-free approach like “transfer entropy” [43] that could capture more sophisticated, non-linear patterns of information flow might also improve the method. In the future, it would be interesting to adapt the method of Sharf et al. in each of these ways, and compare the results obtained against those from his original method, both for simulated, and real data.

When applied to the simulated datasets, the adaptation of Muthmann’s method appears to fix each of the problems associated with the original version (detailed in Section 3.1.2.2) to a large extent. Indeed, it arguably captures the structure underlying each dataset more successfully than the ‘functional connectivity method’. Moreover, it offers an approach to considering the association between pairs of neurons that does not rely, as Sharf et

al.'s method does, on pairwise spike timing – providing an effective way to corroborate the results of such methods.

However, when applied to the human brain slice datasets, the success of Muthmann's adaptation in revealing patterns of association between neurons is limited. For such data, where there are strong associations between distantly-located neurons (as is evidenced by Figure 5.3), the method is subject to a distinct problem. If its spatial range is too narrow, the method will miss associations between the distant areas, yet if the spatial range is wide enough to capture such long-distance connections, the problem of centre-bias is reintroduced, distorting both the 'before' and 'after' activity profiles towards the recording space's centre.

Despite these limitations, the updated method is effective in capturing localised patterns of association. This is evidenced by its success on the simulated data, and on the "HC_52" dataset, where it identifies a pattern of local activity consistent with the findings of the functional connectivity method (as described above in Section 5.3). Future improvements should focus on refining the spatial averaging process to accommodate data of differing spatial and temporal density. Given more time, it would be useful to further investigate the relationship between the optimal spatial window and the average distance between neurons, as well as exploring how the best rank window relates to firing rate and interspike interval distribution.

6.2 Simulation Evaluation

To more effectively validate and refine the adaptation of Muthmann's method, it could be useful to improve the synthetic data used for testing. While the synthetic datasets presented in Section 3.3.2 successfully represent four distinct scenarios – randomised data, as well as planar, ring and rotating waves – none of the simulations include strong associations between neurons placed wide apart. Incorporating such a dynamic into new synthetic datasets would provide a robust test for future iterations of Muthmann's algorithm, helping it to overcome the spatial range limitation described above. Moreover, creating simulations that reflect other more complex patterns, such as the "irregular wave front", visible in the lower row of Figure 3.5B, or simple models to mimic activity observed with neurodegenerative disease could further improve testing. Aside from additional simulations, validation of the presented methods could also be improved by making the current simulations more realistic, perhaps by adding noise typical of MEA recordings, or by simulating changing connectivity patterns over the course of a single synthetic recording.

Simulating data recorded from MEAs with a range of electrode densities, especially those exceeding the 16×16 channel MEA currently used could also help to examine and address the spatial range challenge in Muthmann's adapted method. To represent the electrode density of MEAs like 3Brain's BioCAM X or Maxwell's MaxOne MEA (used to record the data presented in Table 5.1), it would be necessary to significantly increase the number of simulated channels, perhaps beyond 64×64 . However, not all of these channels would need to be active – after all, the maximum number of neurons curated after spikesorting in any of the hippocampal slice recordings is only 351.

An additional factor to consider – both in designing the methods and synthetic validation data, and interpreting method results – is the thickness of the brain or organoid slices. Typically at least 300 μm thick, these slices may contain neurons that are too far from the surface-based electrodes for their activity to be fully detected [33]. It is possible therefore, that undetected neurons above the slice surface could act as intermediates for the couplings between distant yet strongly associated neurons on the surface. This hypothesis could explain the limited vector output observed in the spatial activity plots of Figure 3.2, despite evidence of strong-cross slice couplings depicted in the functional connectivity networks of Figure 5.3. To improve the ability of future iterations of Muthmann’s method to detect distanced associations facilitated by undetected intermediates, three-dimensional synthetic data could be modelled. Such data might model the entire depth of the slice, including the intermediate neurons, while only the surface-level data would be available to the method. A more direct alternative would be to use MEAs equipped with shank-like electrodes, such as 3Brain’s $\mu\text{Needles}$ [21], to penetrate deeper into the neural tissue, enabling the recording of neural activity throughout the slice. In fact, neuroscientists [7, 34] and other companies, such as Plexon [29, 28], are developing MEAs with multiple electrodes on each penetrating shank (see Figure 6.1). In the future, it is possible that such technology enables high-density, three-dimensional recording of neural tissue. Availability of such data would necessitate adaption of each method – for Sharf et al.’s method to consider functional connectivity networks in three-dimensions, and Muthmann’s adapted method to calculate three-dimensional spatial activity profiles.

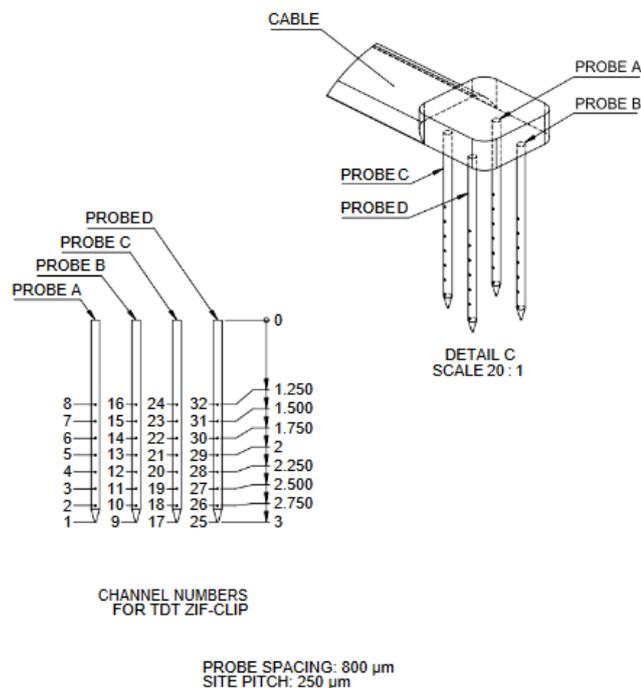


Figure 6.1: Diagram of Plexon’s ‘N-Form’ three dimensional electrode. Taken from Modular Bionics [28].

Chapter 7

Conclusions

This project evaluated the ability of two algorithms – the refined version of Muthmann’s method and Sharf et al.’s approach – to map associative patterns present in spike-sorted time series data from synthetic and biological sources. Through rigorous analysis, several theoretical limitations to Muthmann’s method were identified and detailed extensively, before an improved version designed to mitigate these issues was presented. Synthetic time series data, emulating activity captured by a typical HD MEA were successfully generated for four distinct scenarios, inspired by the biological observations of Wu et al. [44]. Analysis presented in Section 4.1 shows Sharf et al.’s method to be effective in capturing the dynamics within this synthetic testing data. Similarly, Muthmann’s adapted ‘spatial activity profiles’ successfully captured the true structure within the data, corroborating the findings of Sharf et al. Indeed, by the simplicity of its visualisation, it is arguable that Muthmann’s updated method conveyed the identified structure more clearly than the peer-reviewed ‘functional connectivity’ method. However, when applied to the real neuronal data from six human hippocampal slices, it is clear that Sharf et al.’s method outperforms Muthmann’s adaptation in identifying structure in the data. Despite this, there was agreement between some of the findings, with Muthmann’s updated method capturing a local pattern of association also visible in the corresponding ‘functional connectivity’ map.

Overall, this project substantiates the feasibility of using HD MEAs to record extracellular activity from resected brain tissue, and demonstrates that sophisticated algorithms can be used to extract structural information from such data. Although both methods (particularly Muthmann’s) may necessitate further refinement, their potential is evident. Moreover, with the advent of even higher density MEAs, perhaps those featuring hundreds of electrodes per recording shank, the prospect of recording larger resections of the human brain in three dimensions becomes promising. Such comprehensive data, inherently representative of intrinsic neural activity could be coupled with further three-dimensional adaptations to the presented methods, to further improve our understanding of network connectivity and information flow within and across different brain regions. In the long term, such research could be used to map the neural circuits of resections from patients with neurodegenerative disorders, which could improve understanding of disease mechanisms and inform the development of new treatments.

Bibliography

- [1] 3Brain. BioCam Duplex. <https://www.3brain.com/products/single-well/biocam-duplex>, 2024. Accessed: 2024-04-04.
- [2] G. Allen, H. Barnard, R. McColl, A. L. Hester, J. A. Fields, M. F. Weiner, W. K. Ringe, A. M. Lipton, M. Brooker, E. McDonald, C. D. Rubin, and C. M. Cullum. Reduced hippocampal functional connectivity in alzheimer disease. *Archives of Neurology*, 64(10):1482–1487, 2007.
- [3] K. S. Anand and V. Dhikav. Hippocampus in health and disease: An overview. *Annals of Indian Academy of Neurology*, 15(4):239–246, 2012.
- [4] J. Andrews, K. Voitiuk, J. Geng, M. Elliott, D. Shin, A. Robbins, A. Wang, L. Li, A. Spaeth, M. Keefe, J. Sevetson, J. Rivera de Jesus, K. Donahue, H. Larson, D. Ehrlich, K. Auguste, S. Salama, E. Chang, V. Sohal, D. Haussler, C. Cadwell, D. Schaffer, M. Teodorescu, and T. Nowakowski. Optogenetic modulation of epileptiform activity in human brain tissue. Manuscript in Preparation, 2024.
- [5] M. Ballini, J. Müller, P. Livi, Y. Chen, U. Frey, A. Stettler, A. Shadmani, V. Viswam, I. L. Jones, D. Jäckel, M. Radivojevic, M. K. Lewandowska, W. Gong, M. Fiscella, D. J. Bakkum, F. Heer, and A. Hierlemann. A 1024-channel CMOS microelectrode array with 26,400 electrodes for recording and stimulation of electrogenic cells in vitro. *IEEE Journal of Solid-State Circuits*, 49(11):2705–2719, 2014.
- [6] A. M. Bastos and J. M. Schoffelen. A tutorial review of functional connectivity analysis methods and their interpretational pitfalls. *Frontiers in Systems Neuroscience*, 9:175, 2016.
- [7] M. A. Brown, K. M. Zappitelli, L. Singh, R. C. Yuan, M. Bemrose, V. Brogden, D. J. Miller, M. C. Smear, S. F. Cogan, and T. J. Gardner. Direct laser writing of 3D electrodes on flexible substrates. *Nature Communications*, 14(1):3610, 2023.
- [8] A. P. Buccino, S. Garcia, and P. Yger. Spike sorting: new trends and challenges of the era of high-density probes. *Progress in Biomedical Engineering*, 4(2):022005, 2022.
- [9] F. Cramer, G. E. Shephard, and P. J. Heron. The misuse of colour in science communication. *Nature Communications*, 11(1):5444, 2020.
- [10] C. S. Cutts and S. J. Eglén. Detecting pairwise correlations in spike trains: an

- objective comparison of methods and application to the study of retinal waves. *The Journal of Neuroscience*, 34(43):14288–14303, 2014.
- [11] J. Demas, S. J. Eglon, and R. O. Wong. Developmental loss of synchronous spontaneous activity in the mouse retina is independent of visual experience. *The Journal of Neuroscience*, 23(7):2851–2860, 2003.
- [12] M. A. T. Elliott, H. E. Schweiger, A. Robbins, S. Vera-Choqueccota, D. Ehrlich, S. Hernandez, K. Voitiuk, J. Geng, J. L. Severson, Y. M. Rosen, M. Teodorescu, N. O. Wagner, D. Haussler, and M. A. Mostajo-Radji. Internet-connected cortical organoids for project-based stem cell and neuroscience education. *bioRxiv : the preprint server for biology*, 2023.07.13.546418, 2023. Advance online publication.
- [13] S. J. Engle, L. Blaha, and R. J. Kleiman. Best practices for translational disease modeling using human iPSC-derived neurons. *Neuron*, 100(4):783–797, 2018.
- [14] L. Ezama, J. A. Hernández-Cabrera, S. Seoane, E. Pereda, and N. Janssen. Functional connectivity of the hippocampus and its subfields in resting-state networks. *The European Journal of Neuroscience*, 53(10):3378–3393, 2021.
- [15] Lawrence A Fogwe, Veera Reddy, and Fitsum B Mesfin. *Neuroanatomy, Hippocampus*. StatPearls Publishing, 2023.
- [16] A. Gordon, S. J. Yoon, S. S. Tran, C. D. Makinson, J. Y. Park, J. Andersen, A. M. Valencia, S. Horvath, X. Xiao, J. R. Huguenard, S. P. Pasca, and D. H. Geschwind. Long-term maturation of human cortical organoids matches key early postnatal transitions. *Nature Neuroscience*, 24(3):331–342, 2021.
- [17] Robyn Greene. Towards reliable spike sorting evaluation without ground truth data. Master’s thesis, University of Edinburgh, 2021.
- [18] A. Hierlemann, U. Frey, S. Hafizovic, and F. Heer. Growing cells atop microelectronic chips: Interfacing electrogenic cells in vitro with CMOS-based microelectrode arrays. *Proceedings of the IEEE*, 99(2):252–284, 2011.
- [19] M. Koehl and D. N. Abrous. A new chapter in the field of memory: adult hippocampal neurogenesis. *The European Journal of Neuroscience*, 33(6):1101–1114, 2011.
- [20] J. Lisman, G. Buzsáki, H. Eichenbaum, L. Nadel, C. Ranganath, and A. D. Redish. Viewpoints: how the hippocampus contributes to memory, navigation and cognition. *Nature Neuroscience*, 20(11):1434–1447, 2017.
- [21] L. Mapelli, O. Dubochet, M. Tedesco, G. Sciacca, A. Ottaviani, A. Monteverdi, C. Battaglia, S. Tritto, F. Cardot, P. Surbled, J. Schildknecht, M. Gandolfo, K. Imfeld, C. Cervetto, M. Marcoli, E. D’Angelo, and A. Maccione. Design, implementation, and functional validation of a new generation of microneedle 3D high-density CMOS multi-electrode array for brain tissue and spheroids. 2022. Pre-print.
- [22] R. Mohanty, W. A. Sethares, V. A. Nair, et al. Rethinking measures of functional connectivity via feature extraction. *Scientific Reports*, 10:1298, 2020.

- [23] I. Morgado-Bernal. Learning and memory consolidation: linking molecular and behavioral data. *Neuroscience*, 176:12–19, 2011.
- [24] Jens-Oliver Muthmann. *Multiscale modeling of homeostatic processes*. Doctoral dissertation, University of Edinburgh, Manipal University, Manipal, India, 2017. Submitted for the Degree of Doctor of Philosophy (PhD) in the Life Sciences.
- [25] J. Müller, M. Ballini, P. Livi, Y. Chen, M. Radivojevic, A. Shadmani, V. Viswam, I. L. Jones, M. Fiscella, R. Diggelmann, A. Stettler, U. Frey, D. J. Bakkum, and A. Hierlemann. High-resolution CMOS MEA platform to study neurons at subcellular, cellular, and network levels. *Lab on a Chip*, 15(13):2767–2780, 2015.
- [26] M.E.J. Obien, A. Hierlemann, and U. Frey. Accurate signal-source localization in brain slices by means of high-density microelectrode arrays. *Scientific Reports*, 9:788, 2019.
- [27] S. P. Paşca. The rise of three-dimensional human brain cultures. *Nature*, 553(7689):437–445, 2018.
- [28] Plexon. N-Form® 3D Microelectrode Array Chronic Recording Performance. <https://plexon.com/wp-content/uploads/2018/07/N-Form> Accessed: 2024-04-01.
- [29] Plexon. N-form® multi-electrode array. <https://plexon.com/products/n-form/>, 2023. Accessed: 2024-04-01.
- [30] J. Rodu, N. Klein, S. L. Brincat, E. K. Miller, and R. E. Kass. Detecting multivariate cross-correlation between brain regions. *Journal of Neurophysiology*, 120(4):1962–1972, 2018.
- [31] S. Ronchi, A. P. Buccino, G. Prack, S. S. Kumar, M. Schröter, M. Fiscella, and A. Hierlemann. Electrophysiological phenotype characterization of human iPSC-derived neuronal cell lines by means of high-density microelectrode arrays. *Advanced Biology*, 5(3):e2000223, 2021.
- [32] M. Schröter, C. Wang, M. Terrigno, P. Hornauer, Z. Huang, R. Jagasia, and A. Hierlemann. Functional imaging of brain organoids using high-density microelectrode arrays. *MRS Bulletin*, 47(6):530–544, 2022.
- [33] T. Sharf, T. van der Molen, S. M. K. Glasauer, E. Guzman, A. P. Buccino, G. Luna, Z. Cheng, M. Audouard, K. G. Ranasinghe, K. Kudo, S. S. Nagarajan, K. R. Tovar, L. R. Petzold, A. Hierlemann, P. K. Hansma, and K. S. Kosik. Functional neuronal circuitry and oscillatory dynamics in human brain organoids. *Nature Communications*, 13(1):4403, 2022.
- [34] H. Shin, S. Jeong, J. H. Lee, W. Sun, N. Choi, and I. J. Cho. 3D high-density microelectrode array with optical stimulation and drug delivery for investigating neural circuit dynamics. *Nature Communications*, 12(1):492, 2021.
- [35] L. R. Squire. The legacy of patient H.M. for neuroscience. *Neuron*, 61(1):6–9, 2009.
- [36] F. Stella, E. Cerasti, B. Si, K. Jezek, and A. Treves. Self-organization of multiple

- spatial and context memories in the hippocampus. *Neuroscience and Biobehavioral Reviews*, 36(7):1609–1625, 2012.
- [37] C. Stergiadis, D. M. Halliday, and M. A. Klados. Functional connectivity of interictal iieeg and the connectivity of high-frequency components in epilepsy. *Brain Organoid and Systems Neuroscience Journal*, 1:3–12, 2023. Advance online publication.
- [38] J. Terreros-Roncal, E. P. Moreno-Jiménez, M. Flor-García, C. B. Rodríguez-Moreno, M. F. Trinchero, F. Cafini, A. Rábano, and M. Llorens-Martín. Impact of neurodegenerative diseases on human adult hippocampal neurogenesis. *Science*, 374(6571):1106–1113, 2021.
- [39] P. Thiebaud, N.F. de Rooij, M. Koudelka-Hep, and L. Stoppini. Microelectrode arrays for electrophysiological monitoring of hippocampal organotypic slice cultures. *IEEE Transactions on Biomedical Engineering*, 44(11):1159–1163, 1997.
- [40] H. Toyoda, X. Y. Li, L. J. Wu, M. G. Zhao, G. Descalzi, T. Chen, K. Koga, and M. Zhuo. Interplay of amygdala and cingulate plasticity in emotional fear. *Neural Plasticity*, 2011:813749, 2011.
- [41] N. B. Turk-Browne. The hippocampus as a visual area organized by space and time: A spatiotemporal similarity hypothesis. *Vision Research*, 165:123–130, 2019.
- [42] UC Santa Cruz Genomics Institute. Braingeneers GitHub Repository. <https://github.com/braingeneers>. Accessed 3rd April 2024.
- [43] R. Vicente, M. Wibral, M. Lindner, and G. Pipa. Transfer entropy—a model-free measure of effective connectivity for the neurosciences. *Journal of Computational Neuroscience*, 30(1):45–67, 2011.
- [44] J. Y. Wu, Xiaoying Huang, and Chuan Zhang. Propagating waves of activity in the neocortex: what they are, what they do. *The Neuroscientist*, 14(5):487–502, 2008.
- [45] M. P. Zafeiriou, G. Bao, J. Hudson, R. Halder, A. Blenkle, M. K. Schreiber, A. Fischer, D. Schild, and W. H. Zimmermann. Developmental gaba polarity switch and neuronal plasticity in bioengineered neuronal organoids. *Nature Communications*, 11(1):3791, 2020.

Appendix A

First appendix

A.1 Full Detail of Sharf et al.'s Functional Connectivity Method

1. Quantify correlated spiking activity: For each pair of neurons, compute the spike-time-tiling-coefficient of their spike trains using a correlation window of 20ms.
2. Filter by number of spikes: Remove neuron pairs where either neuron has fewer than 5 spikes in its train.
3. Compute latency distribution: For each remaining pair of neurons, compute their spike time latency distributions. Here, latency is defined as being the delay from the time of a spike in neuron A to a spike occurring later in neuron B. To compute the latency distribution from neuron A to neuron B, for each spike in neuron A's train, compute the time difference to every spike in B's train.
4. Filter by unimodality of latency distribution: After removing latencies greater than 20ms, use Hartigan's dip test with a p-value of 0.1 to exclude all electrode pairs with a multimodal latency distribution. This has the effect of isolating the neural pairs with unimodal latency distribution (single peak), enabling the method to focus on simple, direct, one-to-one interactions.

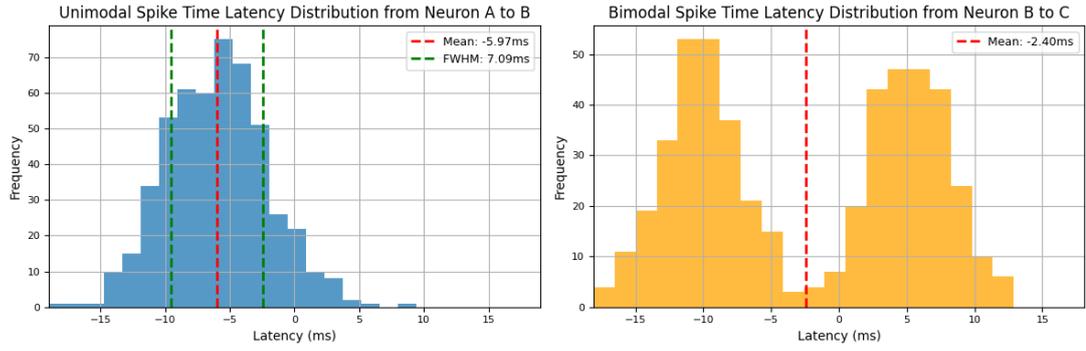


Figure A.1: Example unimodal and bimodal distributions of spike time latencies from neuron A to B and B to C, with means of -5.97ms and -2.4ms, indicating that on average, B fires before A, C before B by these amounts.

5. Filter by width of latency distribution: For all neuron pairs, remove those whose latency distribution has a full width at half the maximum height larger than 15ms. This excludes pairs where the signal transmission is variable.
6. Generate non-directed network: From the remaining neuron pairs, use their STTC values to generate a functional connectivity matrix, where the columns and rows are the indices of the neurons, and the values they access are the STTC values, or 0 for the excluded pairs.
7. Filter by STTC value: All edges of the network with a STTC value less than 0.35 are removed, to minimise the effect of less significant correlations.
8. Assign direction to network edges: For each neuron in the network, compute the mean spike timing latency value to each other neuron it is connected to. Negative values indicate direction towards the reference neuron, and vice versa. Then, for each neuron, set its number of 'incoming edges' (D_{in}) as the number of connected neurons that tend to fire before it (ie, the mean latency is negative), and its number of 'outgoing edges' (D_{out}) as the number of connected neurons that tend to fire after (mean latency is positive).
9. Classify neurons: Classify as neurons as 'senders', 'receivers' and 'brokers', based on their number of incoming and outgoing edges. Precisely:

- Senders: $\frac{D_{out} - D_{in}}{D_{out} + D_{in}} > Threshold$

- Receivers: $\frac{D_{in} - D_{out}}{D_{in} + D_{out}} > Threshold$

- Brokers: $\left| \frac{D_{in} - D_{out}}{D_{in} + D_{out}} \right| < Threshold$

A.2 Full Detail of Human Hippocampal Datasets

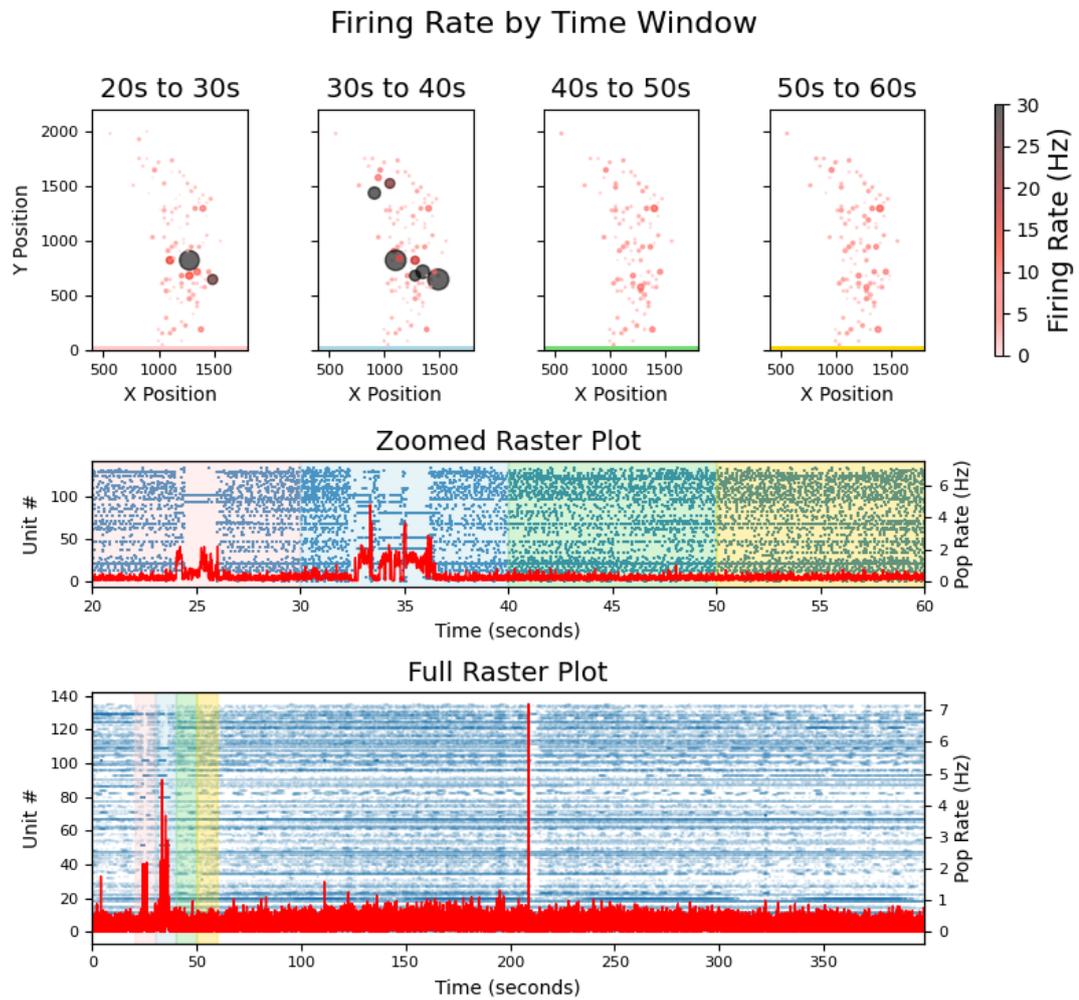


Figure A.2: HC_328. Visualisation of spike-sorted extracellular recording from the outer blade of the dentate gyrus region of the human hippocampus. Top row displays firing rate. Second row is a zoomed-raster plot of the firing rate windows above. Final row is a raster illustrating full dataset, with same colour-coding scheme.

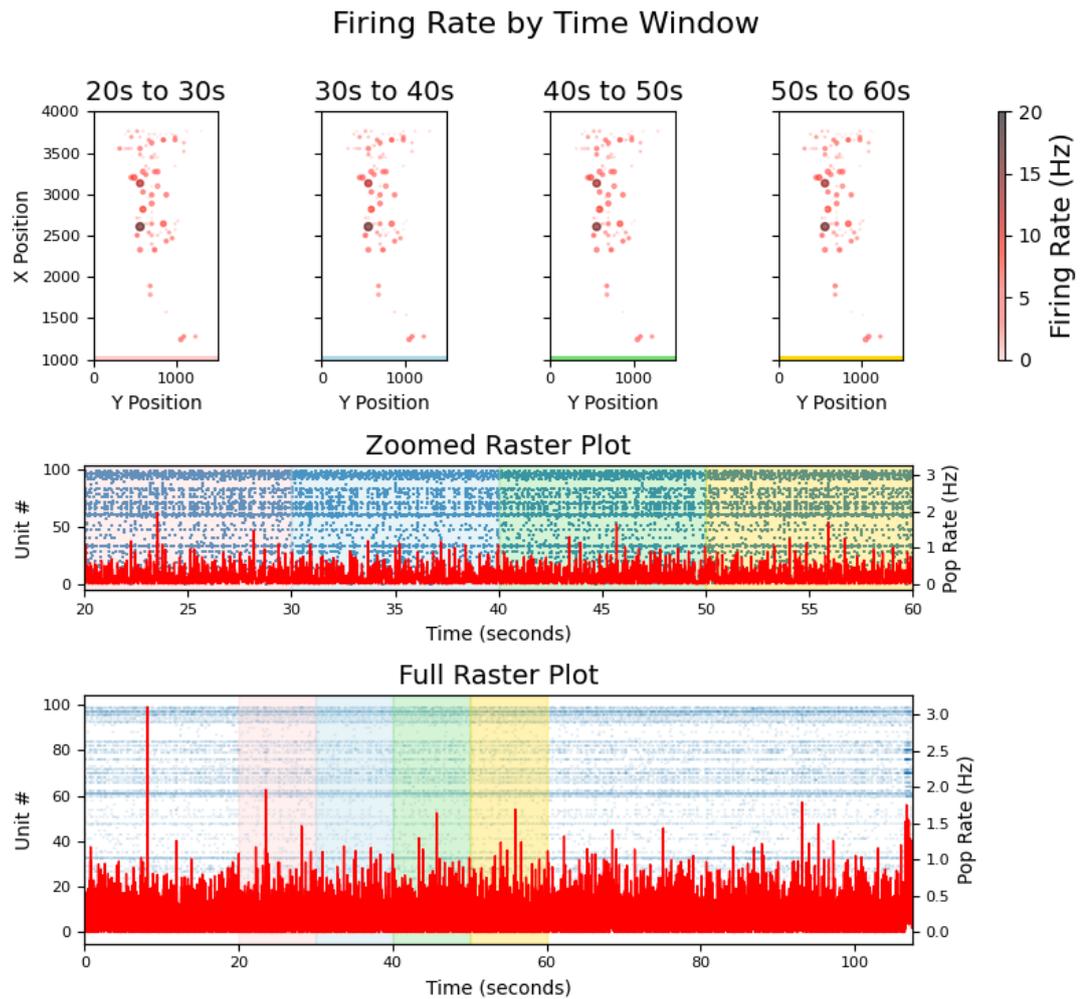


Figure A.3: HC_52. Visualisation of spike-sorted extracellular recording from the outer blade of the dentate gyrus region of the human hippocampus. Top row displays firing rate. Second row is a zoomed-raster plot of the firing rate windows above. Final row is a raster illustrating full dataset, with same colour-coding scheme.

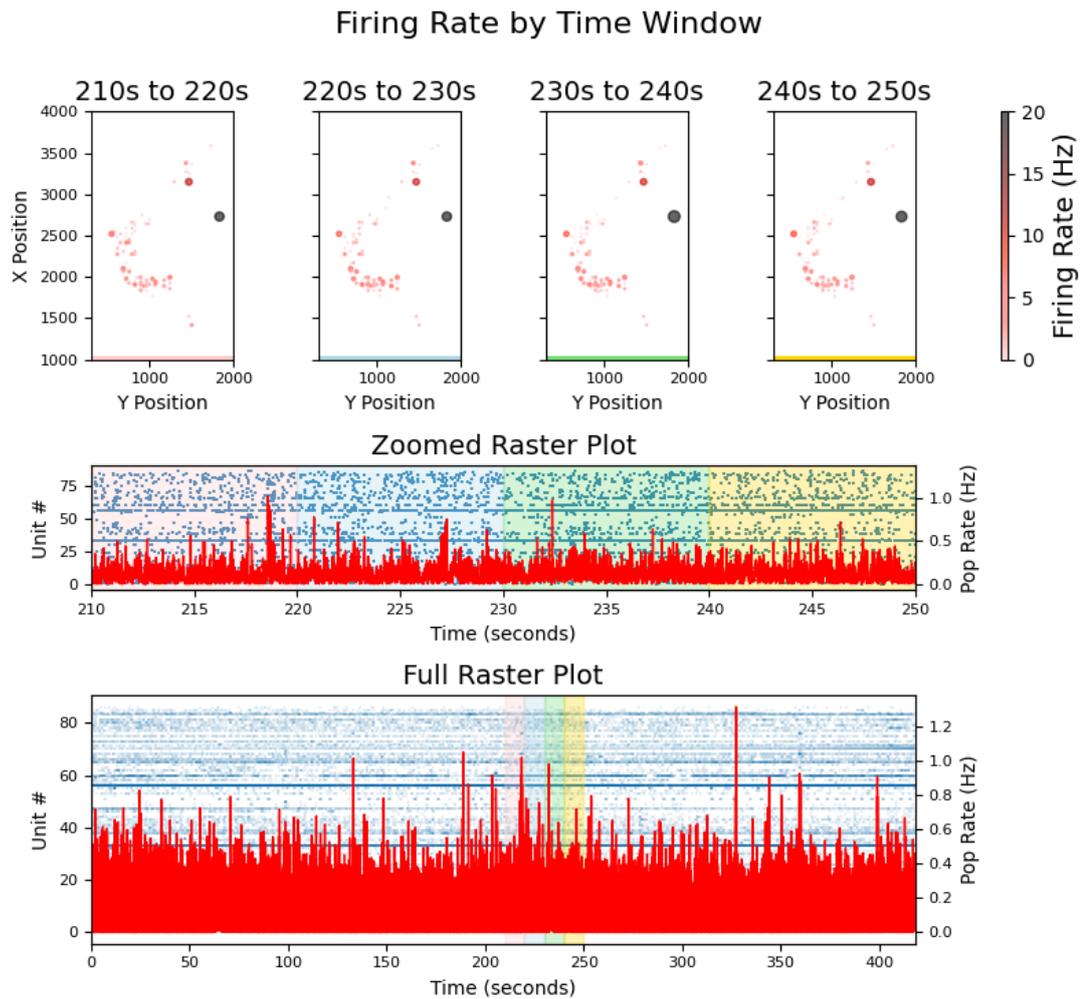


Figure A.4: HC_111. Visualisation of spike-sorted extracellular recording from the outer blade of the dentate gyrus region of the human hippocampus. Top row displays firing rate. Second row is a zoomed-raster plot of the firing rate windows above. Final row is a raster illustrating full dataset, with same colour-coding scheme.

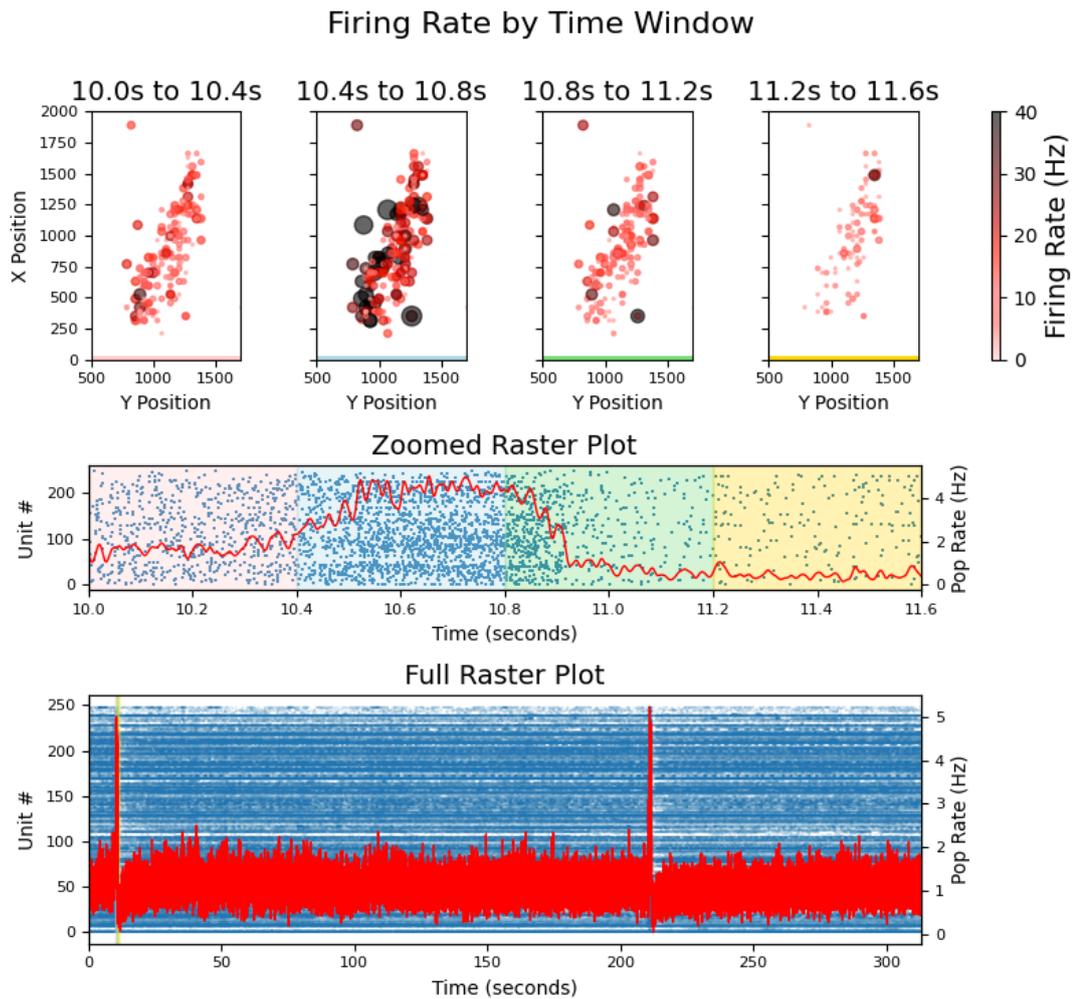


Figure A.5: HC_112823. Visualisation of spike-sorted extracellular recording from the outer blade of the dentate gyrus region of the human hippocampus. Top row displays firing rate. Second row is a zoomed-raster plot of the firing rate windows above. Final row is a raster illustrating full dataset, with same colour-coding scheme.

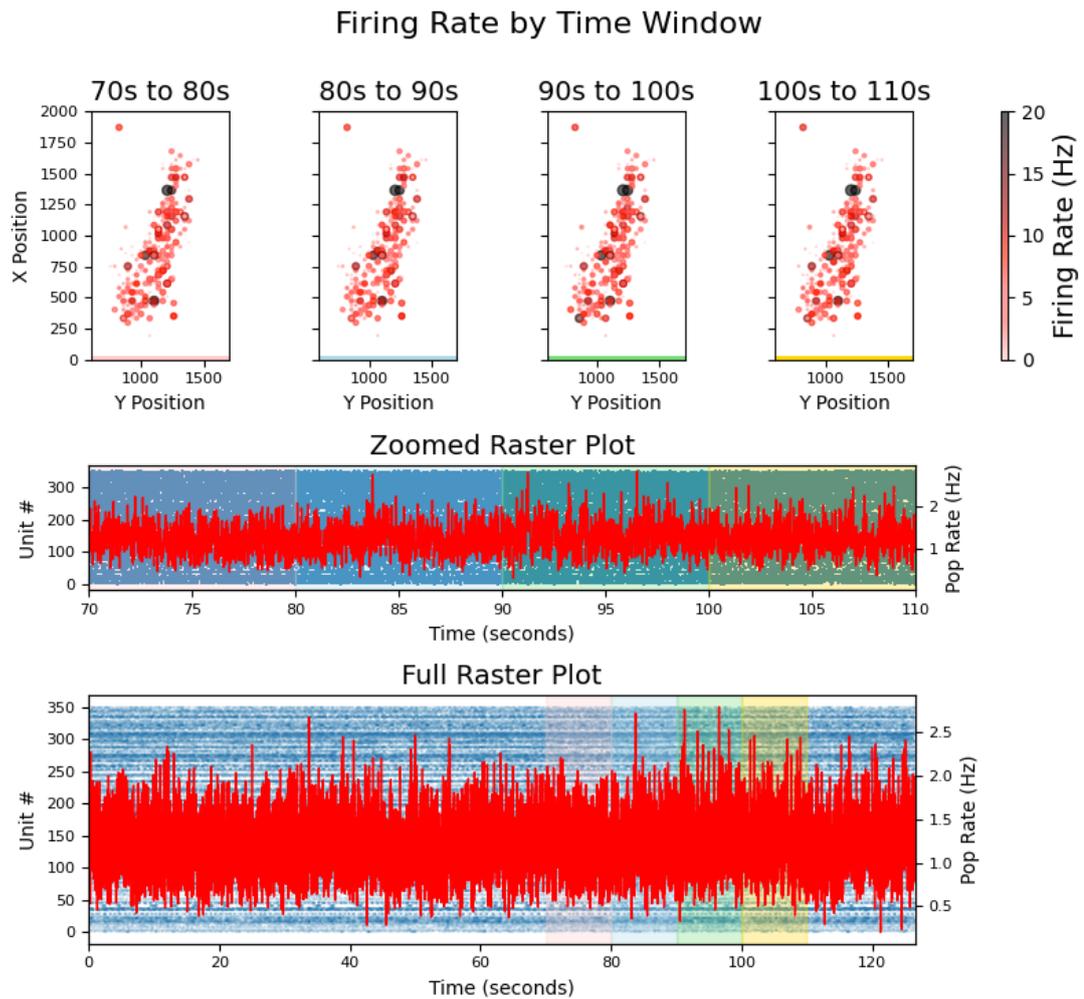


Figure A.6: HC_112823_2. Visualisation of spike-sorted extracellular recording from the outer blade of the dentate gyrus region of the human hippocampus. Top row displays firing rate. Second row is a zoomed-raster plot of the firing rate windows above. Final row is a raster illustrating full dataset, with same colour-coding scheme.

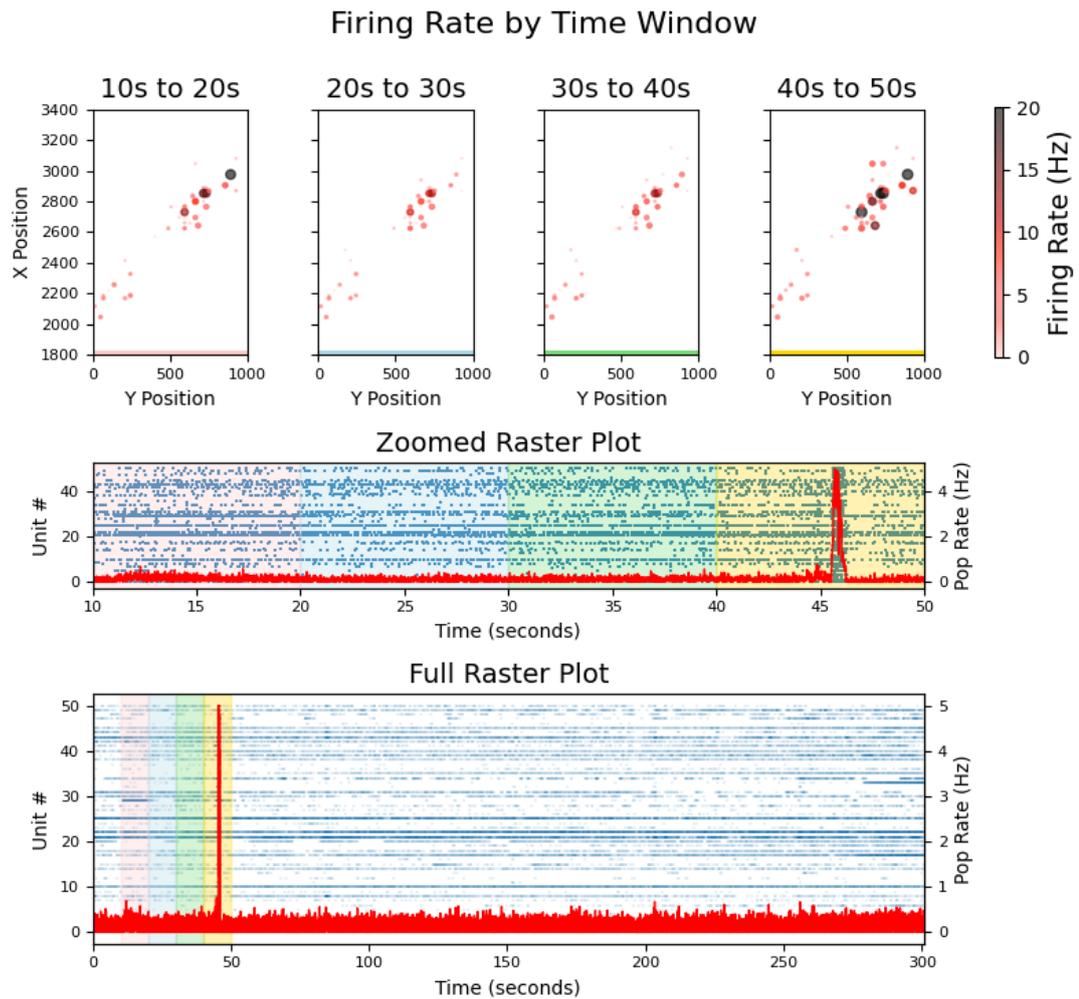


Figure A.7: HC_120523. Visualisation of spike-sorted extracellular recording from the outer blade of the dentate gyrus region of the human hippocampus. Top row displays firing rate. Second row is a zoomed-raster plot of the firing rate windows above. Final row is a raster illustrating full dataset, with same colour-coding scheme.

A.3 Relationship between Connected Pairs and Classification of Neurons in Functional Connectivity Networks

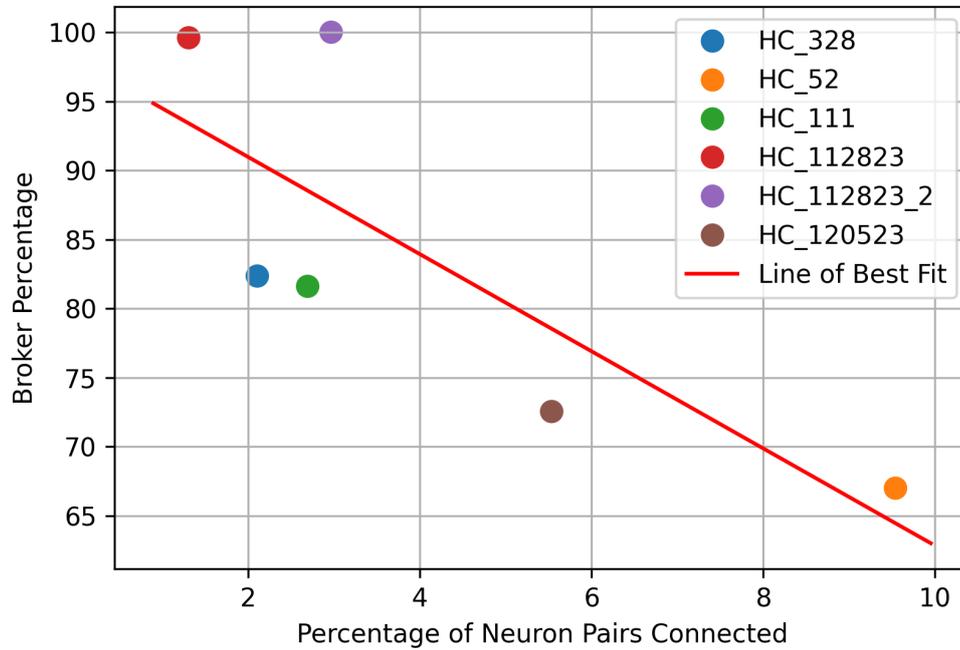


Figure A.8: Relationship between connected pairs and classification of neurons in functional connectivity networks – as the percentage of neuron pairs connected decreases, the percentage of pairs classified as brokers decreases.

A.4 Full Size Functional Connectivity Maps

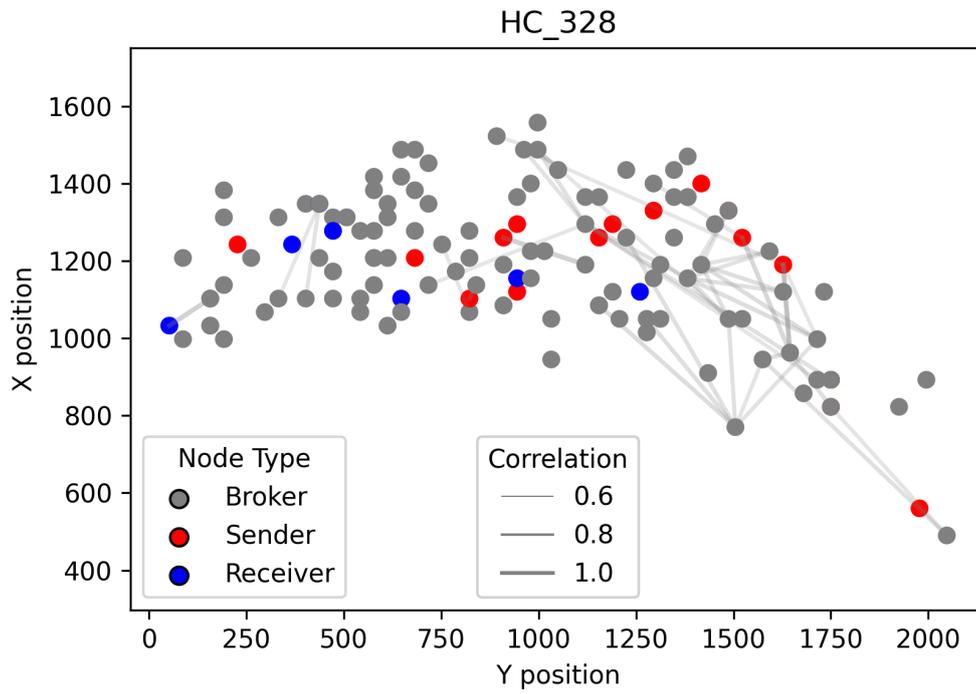


Figure A.9: HC_328 Functional Connectivity Map

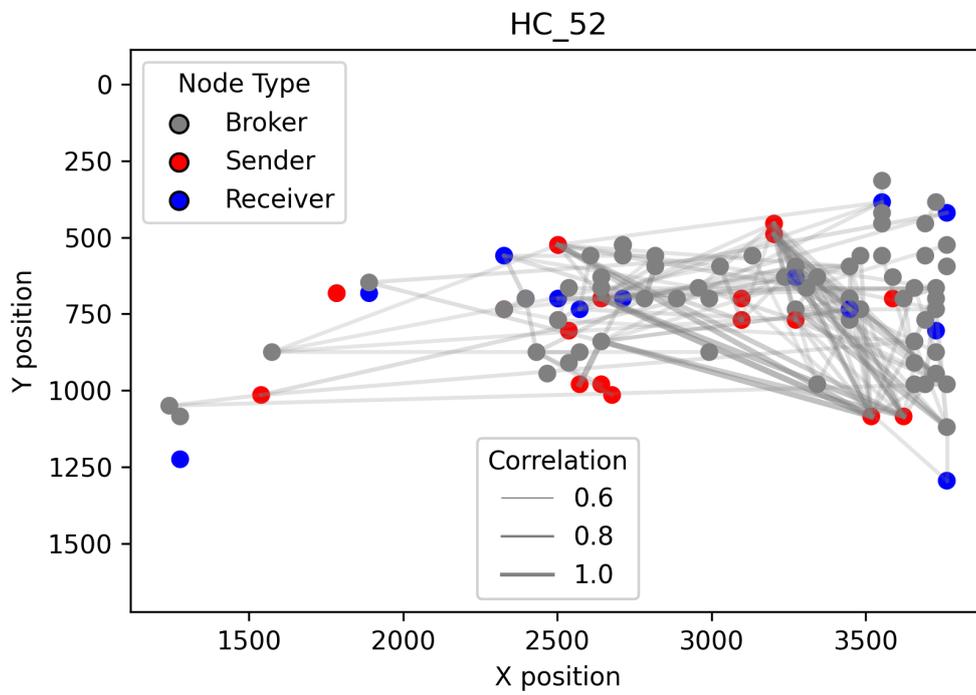


Figure A.10: HC_52 Functional Connectivity Map

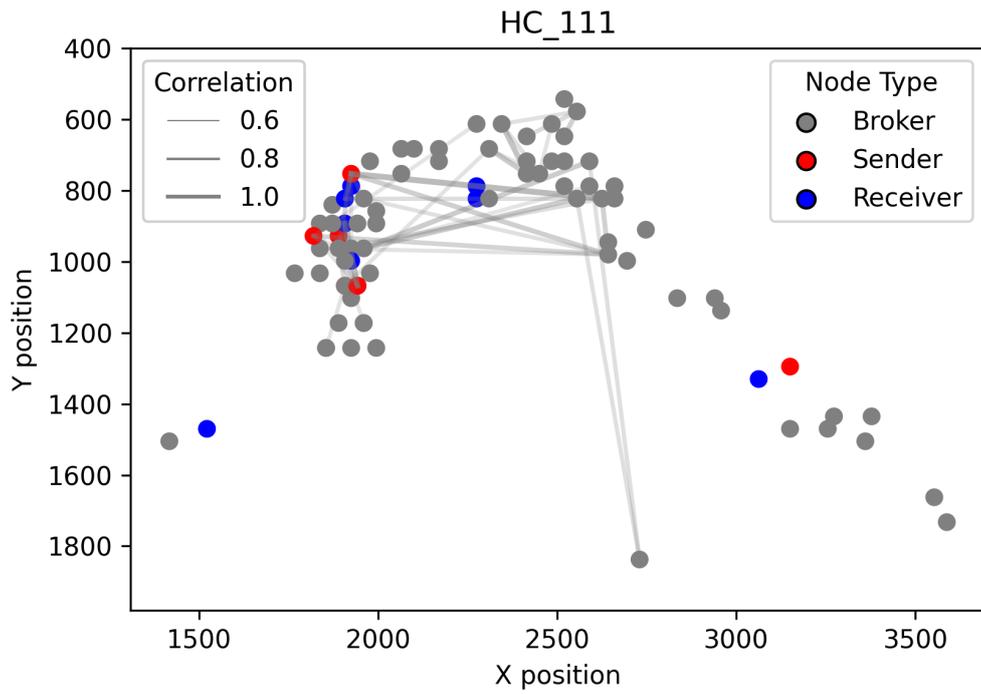


Figure A.11: HC_111 Functional Connectivity Map

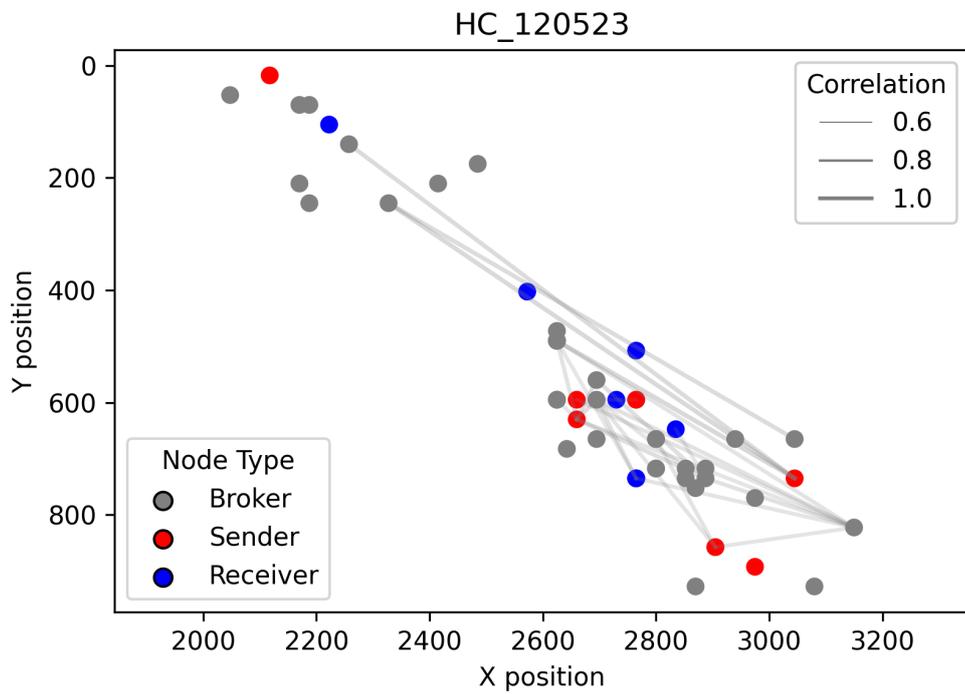


Figure A.12: HC_120523 Functional Connectivity Map

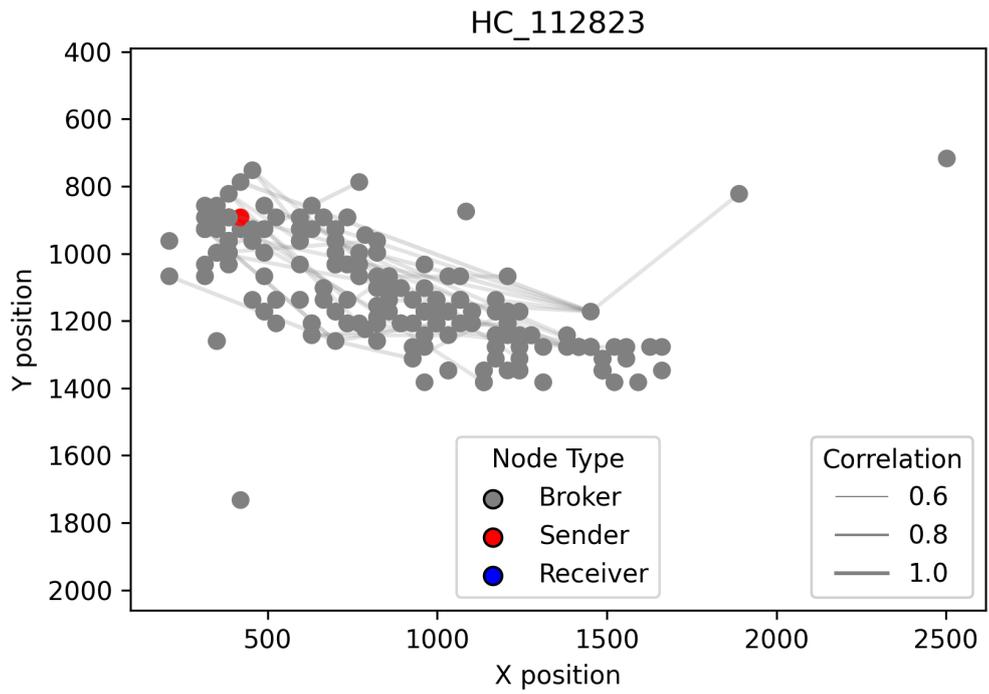


Figure A.13: HC_112823 Functional Connectivity Map

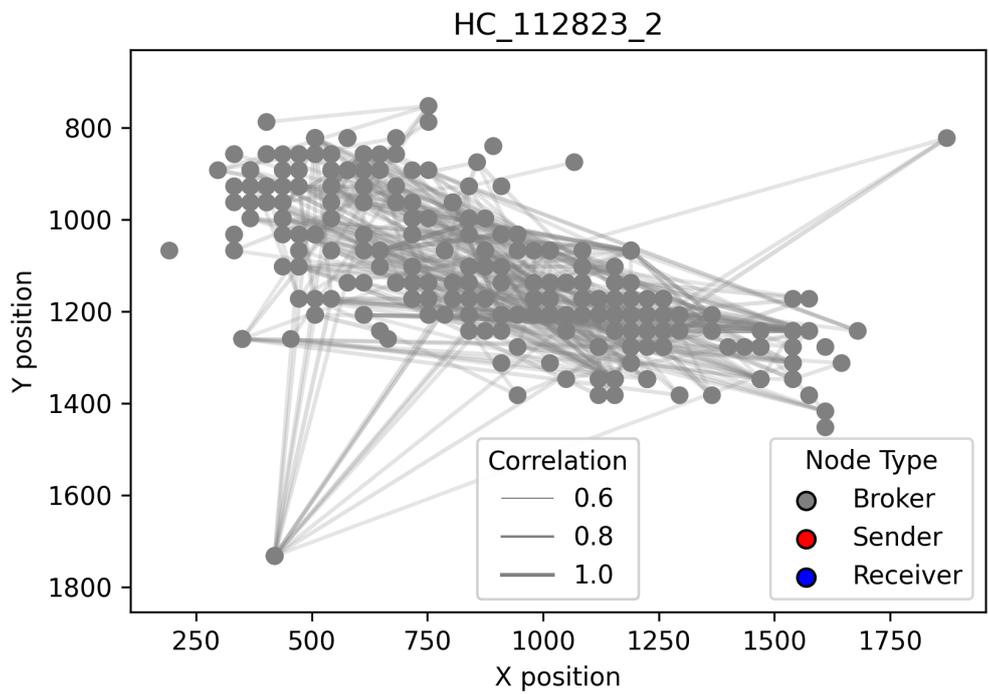


Figure A.14: HC_112823_2 Functional Connectivity Map

A.5 Spatial Activity Profile Parameters for Human Data

The upper spatial bounds used for the computation of the spatial activity profiles presented in Figure 5.4 are presented in Table A.1. The bound was proportional to the calculated median distance between all pairs of neurons. In contrast, the rank window was inversely proportional to the square root of the mean of the median interspike intervals for each unit.

Dataset	Median Distance (μm)	Mean Median ISI (s)	Spatial Bound (μm)	Rank Window (s)
HC_328	560.27	599.23	175	30
HC_52	687.64	1650.96	214	18
HC_111	586.71	1531.33	183	19
HC_120523	315.49	537.61	98	32
HC_112823	409.66	681.23	128	28
HC_112823_2	409.66	1249.27	128	21

Table A.1: Median distance between neuron pairs, mean of the median interspike intervals of each unit, and the accompanying upper spatial bound and rank window used for calculation of the spatial activity profiles presented in Figure 5.4.

Dataset	Mean Lat. (s)	Std Lat. (s)	Mean STTC	Std STTC	Mean ISI (s)	Std ISI (s)
HC_328	9.79	5.88	0.0199	0.0542	41.09	35.82
HC_52	8.24	6.37	0.0396	0.1093	47.47	28.25
HC_111	9.46	6.00	0.0142	0.0662	47.11	25.74
HC_120523	9.85	5.83	0.0678	0.1122	44.24	30.54
HC_112823	9.99	5.79	0.0100	0.0289	48.99	28.66
HC_112823_2	9.97	5.80	0.0031	0.0361	46.96	28.65

Table A.2: Mean and standard deviation of latency, spike time tiling coefficient (STTC), and interspike interval (ISI) in hippocampal datasets.