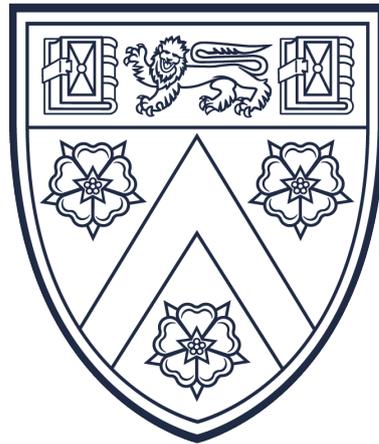


**Statistical analysis of neuronal data: Development
of quantitative frameworks and application to
microelectrode array analysis and cell type
classification**



Ellese Cotterill
Trinity College
University of Cambridge

July 2017

This dissertation is submitted for the degree of *Doctor of Philosophy*

Declaration

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except as specified in the text. It is not substantially the same as any that I have submitted, or is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution. No substantial part of my dissertation has already been submitted, or is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution. My Degree Committee (Mathematics) does not specify a word limit.

Acknowledgements

First and foremost, I would like to thank my supervisor, Stephen Eglén, without whom this thesis would not have been possible. I am incredibly grateful for his support, encouragement and enthusiasm, which has gone above and beyond expectations.

Thanks to my former colleagues, Catherine Cutts and Johannes Hjorth, for their feedback on my work throughout my PhD and to my first year examiners, Julia Gog and John Aston, for their useful comments and advice. A special thanks to John Aston and Julian Budd for their comments on this thesis. Thanks also to our collaborators at the US Environmental Protection Agency, Diana Hall and Timothy Shafer.

I would also like to thank the directors of the Wellcome Trust Mathematical Genomics and Medicine PhD programme for giving me the chance to be a part of such a fantastic programme. Also many thanks to the support and administrative staff of the Centre for Mathematical Sciences and CATO, in particular Joanne Heritage, for all of their assistance and patience at answering my numerous queries and helping make my conference travel throughout my PhD as smooth as possible.

I am greatly appreciative of my family and friends in Australia for supporting my decision to move across the world to pursue this PhD, and for putting up with my four year absence. Special thanks to my dad for proof reading sections of this thesis. Thank you also to my Cambridge friends and my housemates, Katie, Mariel and Nicola, for all of the fun adventures and for their support throughout the more difficult times. Also to Tom, for getting me through this in one piece.

My PhD has been funded by the Wellcome Trust and Cambridge Biomedical Research Centre, who I thank for their kind generosity. I would also like to thank Cambridge Australia Scholarships and Trinity College, for providing me with many academic and cultural opportunities throughout the course of my PhD.

Abstract

With increasing amounts of data being collected in various fields of neuroscience, there is a need for robust techniques for the quantitative analysis of this information. My thesis focuses on the development and assessment of frameworks for the analysis and classification of neuroscience data in a variety of contexts. Firstly, I investigate methods for analysing recordings of the electrophysiological activity of *in vitro* neuronal networks on microelectrode arrays (MEAs). Analysis of this type of activity has applications in many areas, including the study of spontaneous neuronal network activity during development and screening for functional effects of genetic or pharmacological manipulations.

An essential element in characterising the activity of neuronal networks is the accurate identification of episodes of ‘bursting’, which are periods of rapid firing of neurons believed to be associated with a range of physiological processes. Despite this, no single widely used method for the robust identification of bursting activity has been adopted in the field. To address this, I perform an unbiased assessment of the effectiveness of a variety of existing burst detection techniques using both simulated and experimental data. Based on these results, I provide recommendations for the robust analysis of bursting activity in experimental recordings and suggest adaptations to existing burst detection techniques to improve their applicability to a wide range of spike trains. I further apply these techniques to the analysis of novel MEA recordings of networks of human induced pluripotent stem cell-derived neurons, to describe the ontogeny of bursting activity in these networks.

Results from this review of burst detection techniques are then used to inform the development of a broader quantitative framework for the analysis of MEA recordings of spontaneous neuronal network activity. Using this framework, I characterise the developmental profile of spontaneous activity patterns in two sets of experimental recordings of *in vitro* neuronal networks. I show that significant differences in the nature of these activity patterns can be observed across different brain regions and developmental ages, and demonstrate that recordings from these different network types can be reliably differentiated using features of spontaneous activity and standard classification techniques. I further investigate the most appropriate approaches for using MEA devices for the high-throughput screening of neuronal network activity under multiple experimental conditions.

Finally, I consider the classification of neurons into their underlying cell types. Previous studies in this area have generally used only one type of feature of neurons for their classification. I investigate the potential benefits of utilising multiple types of neuronal properties, specifically their morphological and electrophysiological features, in classification studies. In particular, I consider the use of multi-view clustering techniques that are designed to find the shared cluster structure in two or more ‘views’ or feature sets obtained from the same data. By applying a number of multi-view techniques to the clustering of two neuronal data sets, I show that several multi-view clustering algorithms can more reliably differentiate between classes of neurons in these data sets compared to single-view clustering techniques. To close, I examine the properties of the cell types identified by these methods.

Publications

The following journal articles have been published based on work included in this thesis:

Cotterill, E., Charlesworth, P., Thomas, C. W., Paulsen, O. & Eglén, S. J. A comparison of computational methods for detecting bursts in neuronal spike trains and their application to human stem cell-derived neuronal networks. *J. Neurophysiol.* 306-321 (2016). doi:10.1152/jn.00093.2016

Charlesworth, P.*, **Cotterill, E.***, Morton, A., Grant, S. G. & Eglén, S. J. Quantitative differences in developmental profiles of spontaneous activity in cortical and hippocampal cultures. *Neural Dev.* 10, 1-10 (2015). doi:10.1186/s13064-014-0028-0

Cotterill, E.*, Hall, D.*, Wallace, K., Mundy, W. R., Eglén, S. J. & Shafer, T. J. Characterization of Early Cortical Neural Network Development in Multiwell Microelectrode Array Plates. *J. Biomol. Screen.* 21, 510-519 (2016). doi:10.1177/1087057116640520

Contents

List of Tables	xiv
List of Figures	xvi
1 Introduction and biological background	1
1.1 Introduction	1
1.2 Neuron structure and function	2
1.2.1 Basic neuron anatomy	2
1.2.2 Signalling via action potentials	3
1.2.3 Intrinsic neuron properties	5
1.3 The cortex	7
1.3.1 Structure of the cerebral cortex	7
1.4 Cortical cell types	10
1.4.1 Excitatory neurons	10
1.4.2 Inhibitory interneurons	12
1.5 Spontaneous activity in developing neuronal networks	19
1.5.1 Mechanisms of spontaneous activity	20
1.5.2 The role of spontaneous activity	21
1.5.3 Recording of spontaneous activity on microelectrode arrays	23
1.5.4 Properties of cultured neuronal networks on MEAs	25
1.6 Spike trains	27
1.7 Bursting in neuronal networks	28
1.7.1 The functional role of bursting over single spikes	29
1.7.2 The role of bursting in <i>in vivo</i> behaviours	34
1.7.3 Cellular mechanisms underlying bursting	35
1.8 Thesis aims and overview	36
2 Burst analysis	39
2.1 Summary	39

2.2	Introduction	40
2.3	Previous approaches to burst detection	42
2.4	Methods	44
2.4.1	Approach	44
2.4.2	Burst detection methods	44
2.4.3	Analysis of synthetic data	51
2.4.4	Analysis of mouse RGC data	56
2.4.5	Experimental details for hiPSC-derived neuronal network recordings	57
2.5	Results	58
2.5.1	Performance on synthetic data	58
2.5.2	Performance on experimental data	66
2.5.3	Analysis of experimental data using high performing burst detectors	68
2.6	Discussion of existing burst detection techniques	77
2.7	Adapted burst detection methods	80
2.7.1	Adapted CMA algorithm	80
2.7.2	Adapted Poisson Surprise method	84
2.7.3	Performance of adapted burst detectors	86
2.7.4	Discussion of adapted burst detection techniques	94
2.8	Conclusions	94
2.8.1	Limitations and future work	95
3	Spontaneous activity on MEAs	99
3.1	Summary	99
3.2	Comparison of cortical and hippocampal networks	101
3.2.1	Introduction	101
3.2.2	Methods	102
3.2.3	Results	106
3.2.4	Discussion	112
3.3	Ontogeny of spontaneous activity on multi-well MEAs	114
3.3.1	Introduction	114
3.3.2	Methods	115
3.3.3	Results	117
3.3.4	Discussion	126
3.4	Conclusions	128
4	Multi-view neuronal cell type classification	131

4.1	Summary	131
4.2	Introduction	132
4.2.1	Previous work in neuronal classification	133
4.2.2	Neuronal classification using multiple feature types	136
4.2.3	Overview of multi-view clustering	137
4.2.4	Aims of this chapter	142
4.3	Methods	142
4.3.1	Multi-view clustering methods	142
4.3.2	Neuronal data sets	148
4.3.3	Clustering of multi-view neuronal data sets	152
4.4	Results	157
4.4.1	Analysis of the mPFC data set	157
4.4.2	Single-view clustering of the mPFC data set	160
4.4.3	Multi-view clustering of the mPFC data set	164
4.4.4	Cell types from multi-view clustering of the mPFC data set	167
4.4.5	Overview of cell types in the mPFC data set	171
4.4.6	Conclusions from clustering of the mPFC data set	174
4.4.7	Analysis of Allen Cell Types Database data set	176
4.4.8	Single-view clustering of the Allen data set	178
4.4.9	Multi-view clustering of the Allen data set	183
4.4.10	Cell types from multi-view clustering of the Allen data set	186
4.4.11	Overview of cell types in the Allen data set	190
4.4.12	Conclusions from clustering of the Allen data set	194
4.5	Discussion and conclusions	196
4.5.1	Limitations and future work	200
5	Discussion and conclusions	203
5.1	Summary of key contributions	203
5.1.1	Methods of burst detection	205
5.1.2	Frameworks for the analysis and classification of developmental spontaneous activity	206
5.1.3	Multi-view clustering and its application to neuronal cell type classification	206
5.2	Future work	207
5.2.1	Analysis of additional MEA recordings	208
5.2.2	Increased size and diversity of multi-view neuronal data sets	210
5.2.3	Multi-modal neuronal models	211
5.2.4	Concluding remarks	212

A	Supplementary information for Chapter 4	213
A.1	Multi-view clustering of labelled data sets	213
A.1.1	Labelled data sets	213
A.1.2	Results from clustering of labelled data sets	215
A.2	mPFC data set features	218
A.2.1	Morphological features	218
A.2.2	Electrophysiological features	218
A.3	Allen Cell Type Database data set features	219
A.3.1	Morphological features	219
A.3.2	Electrophysiological features	221
A.4	Cell types from multi-view clustering of the Allen data set	226
A.4.1	Cell types from the RCCA method	226
A.4.2	Cell types from the TW- <i>k</i> -means method	227
	Bibliography	229

List of Tables

2.1	Burst detectors and the parameter values used for the implementation of each method on synthetic spike trains	45
2.2	Desirable properties for a burst detector	53
2.3	Models and parameter values used to generate synthetic spike trains .	54
2.4	Summary of the performance of each burst detector on binary desirable properties D1–D4	65
2.5	The relative rank of the performance of each burst detector on desirable properties D5–D11	65
2.6	Parameter values used for burst detection on human induced pluripotent stem cell-derived neuronal networks	71
3.1	MaxInterval method burst detection parameters	103
3.2	Classification performance using a reduced number of features	112
3.3	Classifier performance at predicting age of arrays from the four ages in the study	121
3.4	Classifier performance at predicting the age of arrays for each pairwise combination of ages	122
4.1	Summary of the major features of the mPFC and Allen Cell Types Database data sets	152
4.2	Similarity between the cluster solutions from hierarchical clustering of the morphological, electrophysiological and concatenated mPFC feature sets	161
4.3	Solution from hierarchical clustering using the concatenated feature set of mPFC cells	162
4.4	Solution from hierarchical clustering using the morphological features of mPFC cells	162
4.5	Solution from hierarchical clustering using the electrophysiological features of mPFC cells	163

4.6	Results of the single and multi-view clustering techniques on the mPFC data set	166
4.7	Solution from clustering using the RCCA method on the multi-view mPFC data set	168
4.8	Solution from clustering using the LCCA method on the multi-view mPFC data set	170
4.9	Similarity between the cluster solutions from hierarchical clustering of the morphological, electrophysiological and concatenated feature sets from the Allen Cell Types Database	179
4.10	Solution from hierarchical clustering using the concatenated feature set from the Allen Cell Types Database	182
4.11	Results of the single and multi-view clustering techniques on the Allen Cell Types Database data set	185
4.12	Purity of excitatory clusters by layer and inhibitory clusters by molecular marker from the top performing clustering methods on the Allen Cell Types Database data set	186
4.13	Similarity between the clustering solutions of the Allen Cell Types Database data set	190
A.1	Results of the single and multi-view clustering techniques on the handwritten digits data set	215
A.2	Results of the single and multi-view clustering techniques on the BBC data set	216
A.3	Results of the single and multi-view clustering techniques on the multilingual data set	217
A.4	Solution from clustering of the multi-view Allen Cell Types Database data set using the RCCA method	226
A.5	Solution from clustering of the multi-view Allen Cell Types Database data set using the TW- k -means method	227

List of Figures

1.1	Major features of neuronal anatomy	3
1.2	Representation of the membrane potential of an axon during action potentials	4
1.3	Appearance of the layers of the neocortex	8
1.4	Cell layers in different regions of the cerebral cortex	9
1.5	Major cell types in the monkey cerebral cortex	11
1.6	Classification of interneurons is generally based on a combination of several properties	13
1.7	The firing pattern of neurons is characterised by both their steady state response and onset response to a stepped depolarization	15
1.8	A possible incomplete classification of neocortical inhibitory interneurons	18
1.9	Factors influencing the development of neuronal networks	19
1.10	Microelectrode array from Multi Channel Systems	23
1.11	Example spike trains	27
1.12	Voltage trace from a rhythmically bursting neuron, recorded from cat visual cortex <i>in vivo</i>	29
1.13	Facilitation and depression at synapses plays a role in the filtering of bursts	31
1.14	Spikes recorded from cells in the monkey visual cortex	33
2.1	Example of bursting activity recorded in a mouse retinal ganglion cell	40
2.2	Examples of spike trains recorded from mouse and human neuronal networks	41
2.3	Example of log-adjusted ISI histogram with the threshold for intraburst and interburst intervals found using the logISI method	48
2.4	Example of ISI histogram with the threshold for intraburst and interburst intervals found using the CMA method	49
2.5	Illustration of the parameters used by the MaxInterval method	50

2.6	Examples of one-minute subsets of simulated spike trains for evaluating desirable features D5–D11	55
2.7	Examples of one minute sections of annotated spike trains from retinal ganglion cell recorded at postnatal days 9–15	56
2.8	Results of each burst detector at analysing 100 synthetic spike trains simulated to evaluate desirable properties D5–D8	61
2.9	Results of each burst detector at analysing 100 synthetic spike trains simulated to evaluate desirable properties D9–D11	63
2.10	ROC curves showing the fraction of true positive and false positive spikes identified as being within bursts by each burst detector	67
2.11	Detailed analysis of mouse retinal ganglion cell recordings	70
2.12	Analysis of recordings of networks of human induced pluripotent stem cell-derived neurons	75
2.13	Results of the four burst detectors applied to samples of human induced pluripotent stem cell-derived neuronal network recordings	76
2.14	Histogram of ISIs from example simulated spike trains analysed using the CMA method	81
2.15	Example of ISI histograms with the thresholds for intraburst and interburst intervals found using the thresholded CMA method	83
2.16	Example segmented spike train	85
2.17	Results of each burst detector at analysing 100 synthetic spike trains simulated to evaluate desirable properties D5–8 with adapted burst detectors	87
2.18	Results of each burst detector at analysing 100 synthetic spike trains simulated to evaluate desirable properties D9–D11 with adapted burst detectors	88
2.19	ROC curves showing the fraction of true positive and false positive spikes identified as being within bursts by each adapted and original burst detection method	90
2.20	Analysis of recordings of networks of human induced pluripotent stem cell-derived neurons with adapted and original burst detectors	92
2.21	Results of original and adapted burst detectors applied to samples of human induced pluripotent stem cell-derived neuronal network recordings	93
3.1	Detection of theta bursting on an electrode	105
3.2	Examples of MEA recordings of spontaneous activity in one hippocampal and cortical culture	106

3.3	Features of spontaneous activity in hippocampal and cortical cultured networks	107
3.4	Pairwise electrode correlation values from two MEA recordings	109
3.5	Relative importance of the features used to characterise spontaneous activity patterns	111
3.6	Raster plots of 5 min of activity recorded from one well at each DIV	118
3.7	Features of spontaneous activity over development	119
3.8	Accuracy of classification of wells by their developmental age, using $n \leq 48$ wells per plate	123
3.9	Schematic of the relationship between cultures and plates used for analysis of intra- and inter-culture variability	124
3.10	Firing rate variability between cultures	125
3.11	Network spike variability between cultures	126
4.1	Examples of electrophysiological features of spike trains and single spikes commonly used for neuronal classification.	134
4.2	Examples of morphological features of dendrites calculated using the standardised software L-measure	135
4.3	Schematic of single-view and multi-view clustering	138
4.4	Representation of four approaches to multi-view clustering.	139
4.5	Example morphological reconstructions from the mPFC data set	149
4.6	Example morphological reconstructions and spiking responses from cells in the Allen Cell Types Database data set	151
4.7	The values of AMI, ARI, purity and precision from clusterings of the mPFC data set	155
4.8	Correlation between electrophysiological and morphological features in the mPFC data set	158
4.9	Correlation between electrophysiological and morphological features in the Allen Cell Types Database data set	176
4.10	Principal components analysis of the Allen Cell Types Database feature sets	178
4.11	Several features differentiated the inhibitory and excitatory clusters found using hierarchical clustering of the concatenated Allen Cell Types Database feature sets	180
4.12	Feature weights from TW- k -means clustering of the Allen Cell Types Database data set	189
4.13	Normalised cortical depth of cells by layer in the Allen Cell Types Database data set	191

A.1	Examples of single action potential features	222
A.2	Examples of features calculated using multiple input current sweeps .	225

Abbreviations

ADP	After depolarizing potential
AMI	Adjusted mutual information
ARI	Adjusted Rand Index
CCA	Canonical correlation analysis
CMA	Cumulative moving average
CV	Coefficient of variation
DIV	Days in vitro
EAP	Extracellular action potential
EPSP	Excitatory postsynaptic action potential
FDR	False discovery rate
GABA	γ -aminobutyric acid
GDP	Giant depolarizing potential
HCN	Hyperpolarization-activated cyclic nucleotide-gated
hiPSC	Human induced pluripotent stem cell
HSMM	Hidden semi-Markov model
IBI	Interburst interval
IRT	Interspike interval rank threshold
ISI	Interspike interval
LCCA	Linear canonical correlation analysis
LOOCV	Leave-one-out cross-validation
LTD	Long term depression
LTP	Long term potentiation
MEA	Microelectrode array
MI	MaxInterval
mPFC	Medial prefrontal cortex
PC	Principal component
PLC	Probabilistic latent clustering
PS	Poisson surprise
PV	Parvalbumin
RCCA	Randomised non-linear canonical correlation analysis
RGC	Retinal ganglion cell
RGS	Robust Gaussian surprise
RMP	Resting membrane potential
RS	Rank surprise
SST	Somatostatin
WAP	Weeks after plating
5HT3aR	Ionotropic serotonin receptor 5HT3a

Chapter 1

Introduction and biological background

1.1 Introduction

Modern neuroscience is a diverse field, centred around the goal of increasing knowledge of the structure and function of the brain and wider nervous system and how these regions dictate cognitive function and behaviour. The basic building blocks of this system are neurons, nerve cells that have a variety of distinct anatomical and functional features. A complete understanding of the nervous system requires knowledge of the properties of these cells and the connections between them, how these networks are developed and maintained, as well as how they contribute to the proper functioning of an organism.

Since the early work of Santiago Ramón y Cajal, the ‘father of modern neuroscience’, allowed neurons to be visualised as distinct connected cells (De Carlos and Borrell, 2007), the development of novel and increasingly sophisticated technologies for collecting experimental data has revolutionised the spatial and temporal resolution with which the nervous system can be studied. These techniques demand equally advanced computational methods to ensure that the information and insights extracted from painstakingly collected experimental data are as accurate and comprehensive as possible.

The work in this thesis focuses on the development, examination and critical evaluation of computational analysis techniques in two important areas of neuroscience.

Firstly, I study techniques for analysing the spontaneously generated activity of neurons, which has been observed extensively across many regions of the nervous system during development, and is thought to play an essential role in the proper development of functional circuits. Study of the spontaneous activity of neuronal networks *in vitro* can elucidate important aspects of how these networks functionally and structurally develop, as well as the effect of pharmacological and genetic manipulations to these systems.

Understanding the underlying mechanisms responsible for neuronal activity patterns also requires the cellular components of these networks to be defined. Such classification of neurons is an essential step in understanding the role of different cells within neuronal networks, how these roles are dictated by the intrinsic structural, molecular and functional properties of cells, as well as how each cell type contributes to computation and cognition. The second half of this thesis focuses on computational techniques that can be applied to the task of characterising neuronal cell types from diverse data sources, and in particular, the potential contribution of multi-view clustering methods.

The bulk of this chapter consists of an overview of the biological concepts and experimental techniques that are relevant for the remainder of this thesis. Following this, I outline the overall aims and general structure of this thesis.

1.2 Neuron structure and function

1.2.1 Basic neuron anatomy

Neurons typically consists of four major anatomical regions (Kandel et al., 2000) (Figure 1.1):

- Soma - the cell body, containing the nucleus and endoplasmic reticulum.
- Dendrites - branched projections from the soma that receive and propagate incoming signals to the neuron.
- Axon - a single long tubular projection from the soma responsible for carrying signals to postsynaptic cells. The axon is connected to the soma at the axon hillock.
- Presynaptic terminals - the terminating points of the axonal branches, at which the presynaptic cell communicates with postsynaptic cells across synapses.

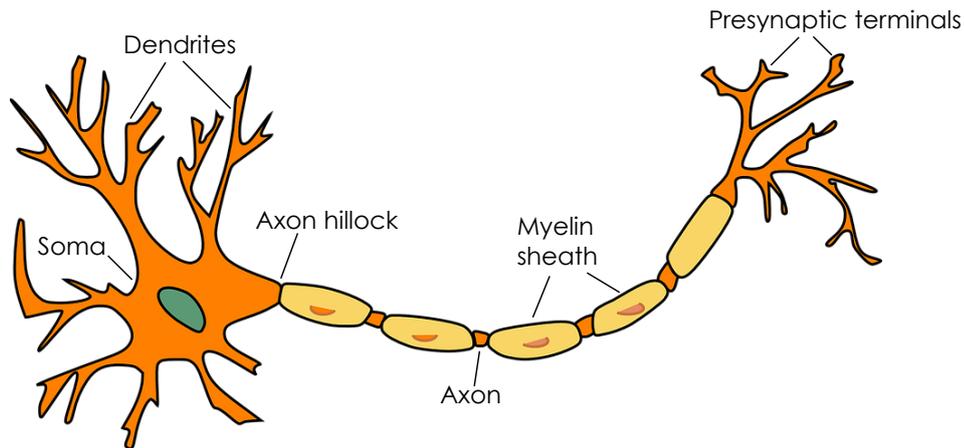


Figure 1.1: Major features of neuronal anatomy. The cell body or soma gives rise to numerous dendrites and generally a single axon that terminates in presynaptic terminals. In many mammalian neuron types, an insulating myelin sheath coats much of the axon.

Within this general structure, different types of neurons differ vastly in their anatomical features, ranging from those with highly branched dendritic trees to those with low number of dendrites, and axons of various sizes and degrees of branching (see Figure 1.5 for examples).

1.2.2 Signalling via action potentials

The primary mode of communication between neurons is in the form of electrical impulses called action potentials, which are sudden depolarizations of the neuronal membrane followed by a swift repolarization (Lodish et al., 2003; Hodgkin and Huxley, 1952). These voltage changes are a result of the rapid influx of Na^+ ions into the neuron through voltage-gated Na^+ channels embedded in the cell membrane, causing a rise in the membrane potential to a certain peak voltage close to the sodium equilibrium potential (Lodish et al., 2003; Alberts et al., 2002). At this point, efflux of K^+ ions from the neuron through voltage-gated K^+ channels returns the membrane to its resting potential (Purves et al., 2001). The entire process of the action potential generally lasts for approximately 1 ms and has a peak voltage of around +50 mV (Alberts et al., 2002) (Figure 1.2).

Action potentials are the only type of membrane potential fluctuations that can travel over long distances in a neuron (Dayan and Abbott, 2001). They are typically generated at the axon hillock or first node of the axon, and travel rapidly along the

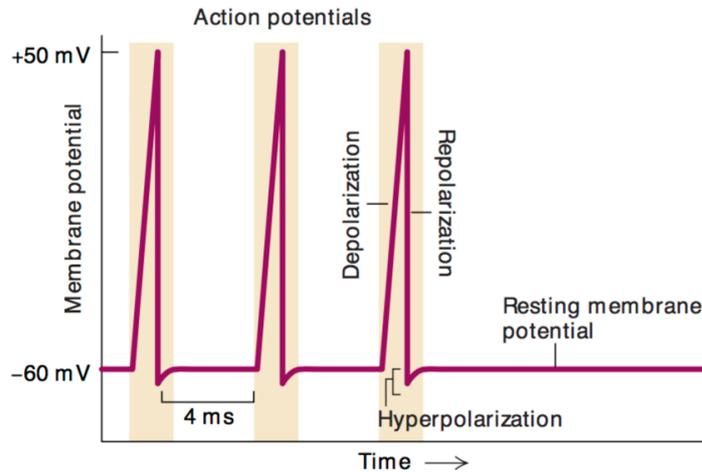


Figure 1.2: Representation of the membrane potential of an axon during repeated action potentials, showing the periods of depolarization and repolarization, as well as the brief period of hyperpolarization that occurs before the neuron returns to its resting potential. Figure reproduced from Lodish et al. (2003).

axon with no decay before terminating in the presynaptic terminals (Kandel et al., 2000). The arrival of an action potential at the presynaptic terminals results in an influx of Ca^{2+} ions and subsequent release of neurotransmitters into the synaptic cleft between the presynaptic and postsynaptic neuron (Lodish et al., 2003). These neurotransmitters diffuse across the synaptic cleft and bind to transmitter-gated ion channels on the postsynaptic neuron (Alberts et al., 2002). Depending on the nature of the postsynaptic receptor to which they bind, the neurotransmitters can produce an excitatory or inhibitory response in the postsynaptic cell (Lodish et al., 2003).

A neuron can simultaneously receive a number of signals from both excitatory and inhibitory synapses, which propagate along the dendrites to the soma and axon hillock (Lodish et al., 2003). At this point, the neuron integrates all of the incoming signal, and produces an action potential if the membrane at this ‘trigger point’ becomes sufficiently depolarized, to a level above the threshold potential (Kandel et al., 2000). An action potential is thus an ‘all-or-none’ signal - a neuron will produce an action potential of approximately the same amplitude and duration in response to any summed stimuli above this threshold, and no response otherwise (Alberts et al., 2002).

The electrical activity of neurons can be measured either intracellularly or using extracellular methods. Intracellular techniques typically involve connecting a hollow glass microelectrode filled with an electrolyte to the neuron and measuring the potential relative to an extracellular reference electrode (Heinricher, 2004). This can

be achieved using sharp electrodes inserted through the cell membrane, or patch pipettes which form a high resistance seal that allows current from the membrane to flow into the pipette and be measured (Hamill et al., 1981; Brette and Destexhe, 2012). Intracellular methods can record both action potentials and subthreshold membrane potential fluctuations. Conversely, extracellular techniques, which involve recording induced extracellular currents using an electrode placed near the neuron, can only detect action potentials and not subthreshold activity (Heinricher, 2004).

1.2.3 Intrinsic neuron properties

Several passive properties of neuronal cells contribute to electrical signalling within a neuron. These include:

- **Resting membrane potential:** In the membrane of a neuron at rest, many K^+ permeable ion channels are open, allowing the flow of K^+ ions out of the cell along their concentration gradient, whereas channels permeable to other ions such as Na^+ are largely closed (Lodish et al., 2003; Kandel et al., 2000). This selective permeability of the membrane, as well as the high intracellular concentration of negatively charged protein ions, results in a difference of charge across the neuronal cell membrane of approximately -60 mV to -70 mV , called the resting membrane potential (Purves et al., 2001).
- **Membrane capacitance:** The ion imbalance across a neuronal membrane at rest means that it acts as a capacitor, with an excess of negative charge along the intracellular surface of the membrane and a corresponding positive charge extracellularly (Dayan and Abbott, 2001). As in all capacitors, the capacitance of the membrane is dependent on the area of the capacitor, the nature of the insulating medium and the distance between capacitor plates. Since cell membranes consist of relatively equally spaced lipid bilayers with similar insulating properties, the specific capacitance per unit area of a biological membrane is a constant value, approximately $1\mu F cm^{-2}$ (Kandel et al., 2000). The total capacitance of a neuron is thus dependent only on its surface area (Dayan and Abbott, 2001).
- **Input resistance:** The input resistance (R_{in}) of a neuron dictates the degree of hyperpolarization or depolarization (ΔV) of the neuron in response to a certain level of injected current (I_e), according to Ohm's law (Dayan and Abbott, 2001):

$$\Delta V = I_e \times R_{in}.$$

Since ion channels are responsible for conducting current across a neuronal membrane, the level of input resistance of a neuron is dependant on the density of resting ion channels in the membrane and the size of the neuron, as neurons with greater membrane surface area generally contain a larger number of ion channels (Kandel et al., 2000).

- **Membrane time constant:** The membrane time constant, τ , is defined as the time taken for the membrane potential of a neuron to increase to $(1 - \frac{1}{e})$ of its resting value, and is the product of the input resistance and membrane capacitance (Kandel et al., 2000):

$$\tau = R_{in} \times C_m.$$

As capacitance and resistance have inverse dependencies on membrane surface area, τ is independent of the surface area of the cell (Dayan and Abbott, 2001).

In addition to voltage-gated K^+ and Na^+ channels, another type of ion channel, called hyperpolarization-activated cyclic nucleotide-gated (HCN) cation channels, are thought to influence both the passive and active properties of neurons (George et al., 2009). In contrast to the majority of voltage-gated ion channels in a neuronal membrane, HCN channels are activated by hyperpolarization below -50 mV to -60 mV (Robinson and Siegelbaum, 2003). Activation of these channels produces an inward cation current, called the hyperpolarization-activated cation current or I_h , that depolarizes the local membrane (Biel et al., 2009).

Because HCN channels are activated at levels at or negative to the typical resting membrane potential of a neuron, many of these channels are open in a neuron at rest. This allows for the inflow of I_h , which has a depolarizing effect of 5 mV to 10 mV on the resting membrane potential and increases resting membrane conductance (George et al., 2009; Doan and Kunze, 1999; Angelo et al., 2007; Pape, 1996). Due to the relationships outlined above, this means that I_h also has the effect of decreasing the input resistance (which is the inverse of membrane conductance) and membrane time constant of a cell (Robinson and Siegelbaum, 2003).

Activation of I_h in response to hyperpolarizing current pulses is also responsible for the depolarizing deflection or ‘sag’ observed in neurons with active HCN channels. Additionally, in neurons such as hippocampal CA1 and pyramidal cells of layer V neocortex, an increased density of I_h has been observed with increasing distance from the soma (Williams and Stuart, 2000; Angelo et al., 2007). This is believed to

influence dendritic integration, by normalising the time course of EPSPs received at different distances from the soma (Robinson and Siegelbaum, 2003).

1.3 The cortex

The largest regions of the human brain are the cerebral hemispheres, which consist of the cerebral cortex, white matter, basal ganglia, amygdala and hippocampus (Kandel et al., 2000). The thin convoluted outer layer of the cerebral hemispheres is made up of a continuous 2 mm to 4 mm thick layered sheet of neurons and other cells, known as the cerebral cortex (Purves et al., 2001). Different regions of the cerebral cortex are responsible for a range of behaviours, including sensory processing, motor functions, problem solving and complex thought (Kandel et al., 2000).

1.3.1 Structure of the cerebral cortex

The majority of the cerebral cortex is neocortex. Neocortex is composed of six layers, each containing distinctive populations of cells (Purves et al., 2001; Kandel et al., 2000) (Figure 1.3):

- Layer I - the outermost layer of the cortex, often called the ‘molecular layer’. This layer contains few neurons, all of which are inhibitory neurons, and consists primarily of the axons and dendrites from neurons in deeper layers.
- Layer II - the ‘external granular layer’, composed mostly of small spherical cells called granule cells and some inhibitory interneurons, as well as the apical dendrites of layer V and VI cells.
- Layer III - the ‘external pyramidal layer’, which contains a variety of cell types, the majority of which are small pyramidal cells.
- Layer IV - the ‘internal granule cell layer’, consists of mostly granule cells and is the principal location of spiny stellate cells. Layer IV is the primary layer for receiving input from the thalamus.
- Layer V - the ‘internal pyramidal layer’, consists primarily of large pyramidal cells and a smaller proportion of inhibitory neurons. Axons of layer V pyramidal cells often make long range projections that leave the cortex.
- Layer VI - the ‘multiform layer’, is a heterogeneous layer consisting of a range of neurons, and blends into the white matter.

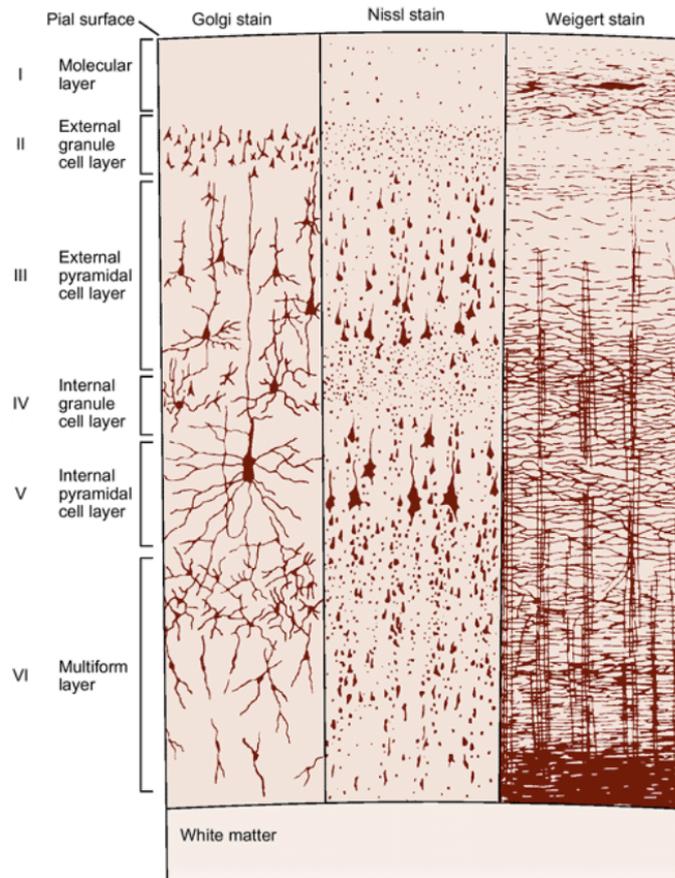


Figure 1.3: The appearance of the layers of the neocortex by different staining methods. The Golgi stain (left) shows the soma and dendritic trees of neurons, while the Nissl stain (centre) reveals cell bodies and proximal dendrites of neurons and the Weigert stain (right) shows axonal distributions. Figure reproduced from Kandel et al. (2000).

Despite this general structure, there are also regional differences in laminar organisation across the cortex which define ‘cytoarchitectonic areas’ that often have distinct functional roles (Purves et al., 2001) (Figure 1.4). For example, sensory areas such as the primary visual cortex have a very prominent layer IV that can further be divided into three or more layers (Kandel et al., 2000). In the motor cortex, layer IV is small or non-existent and layer V is the most prominent (Barbas and Miguel, 2015). This variability in structure can be understood in terms of the inputs and outputs of the cortical layers. Layer IV is the primary target for sensory information from the thalamus, and in primates that rely heavily on visual inputs, this information is topographically mapped into layer IV (Sun et al., 2016). The motor cortex serves as a primarily output region and receives little input from the thalamus, accounting for the small layer IV in this region (Kandel et al., 2000). Most motor signals originate in layer V, explaining the prominence of this layer in the motor cortex (Shipp, 2007).

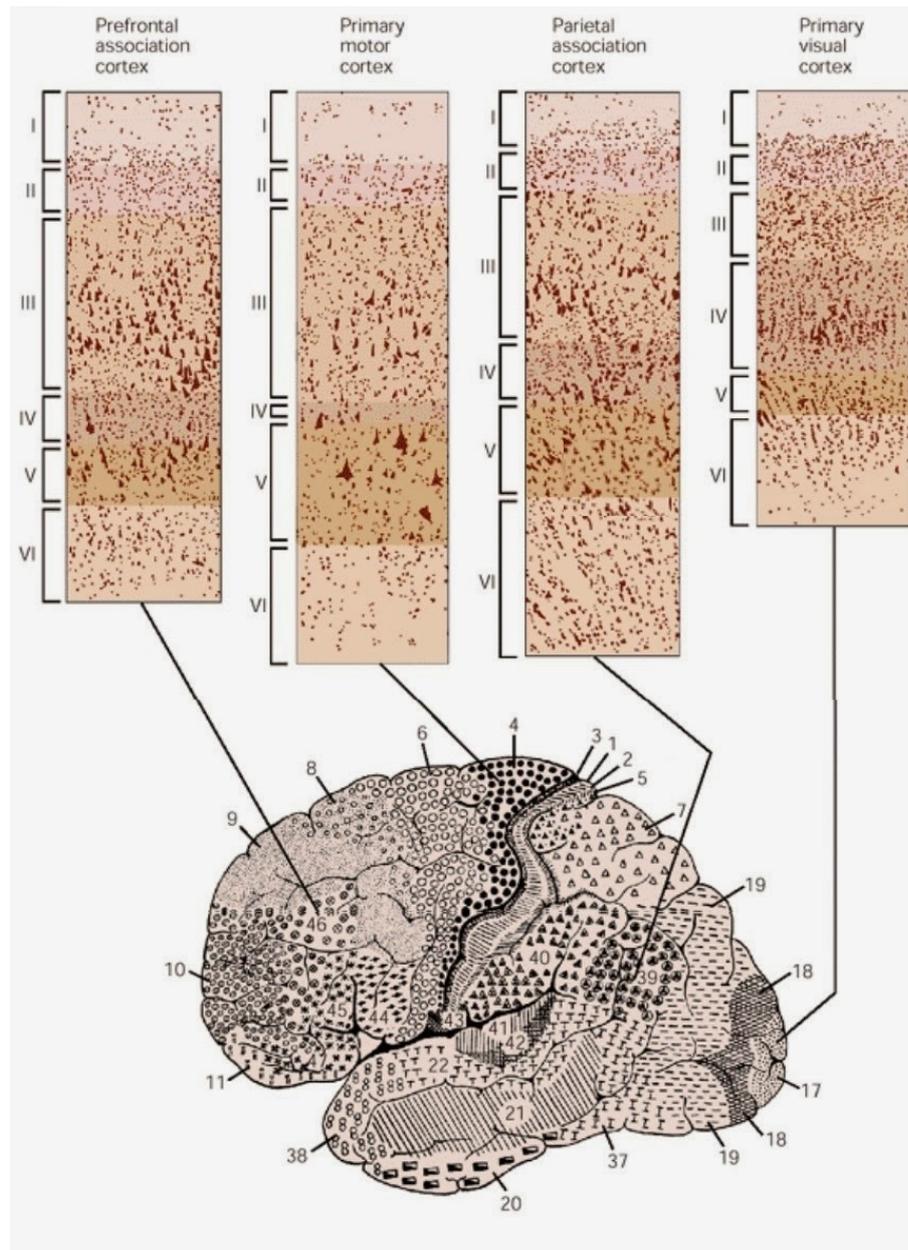


Figure 1.4: Cell layers in different regions of the cerebral cortex. Sensory areas, for example the primary visual cortex, have pronounced layer IV, while motor cortices have small layer IV and prominent output layers, including layer V. Bottom panel shows the lateral view of the brain with regions labelled. Figure adapted from Kandel et al. (2000).

1.4 Cortical cell types

Neurons in the cortex can be broadly grouped into two types, excitatory and inhibitory. Excitatory neurons comprise approximately 70 to 80% of the total neuronal population of the cortex, and use the excitatory amino acid glutamate as their primary neurotransmitter (Kandel et al., 2000). The remaining 20 to 30% of neurons in the cortex are largely inhibitory interneurons which utilise γ -aminobutyric acid (GABA) as their primary neurotransmitter (Harris and Mrsic-Flogel, 2013).

These two major groups of neurons can be subdivided into a number of classes, based on a variety of criteria, including their morphology, physiological properties, gene expression profiles, connectivity patterns and developmental history, amongst others (Tasic et al., 2016). Although there has been some success in developing a top-level classification of neurons based on a subset of these criteria, within each of these top-level classes exists an unknown number of subclasses of neurons, and the development of a definitive classification criteria for neurons that is consistent across the multiple modalities is an ongoing area of investigation (Harris and Shepherd, 2015). In the following sections, I give a brief overview of our current understanding of the major types of excitatory and inhibitory neurons.

1.4.1 Excitatory neurons

Morphologically, excitatory neurons can be divided into two primary cell types (Figure 1.5, left):

- **Pyramidal cells:** Pyramidal cells are the most common type of excitatory neuron, and are characterised by their pyramidal-shaped soma and distinct apical and basal dendritic trees (Toledo-Rodriguez et al., 2003). Although the morphology of pyramidal cells varies between different layers, regions and species, they generally have several relatively short basal dendrites which occupy the same layer as the soma and a single prominent apical dendrite, which branches into a tuft at a distance from the soma (Spruston, 2008).
- **Spiny stellate neurons:** Spiny stellate neurons are generally smaller than pyramidal cells and are characterised by a star-like dendritic arbour, consisting of many short dendrites surrounding a spherical soma (Toledo-Rodriguez et al., 2003; Markram et al., 2004). Spiny stellate cells are almost exclusively found in layer IV of primary sensory cortical areas and project mostly to areas local to their soma (Costa and Müller, 2015).

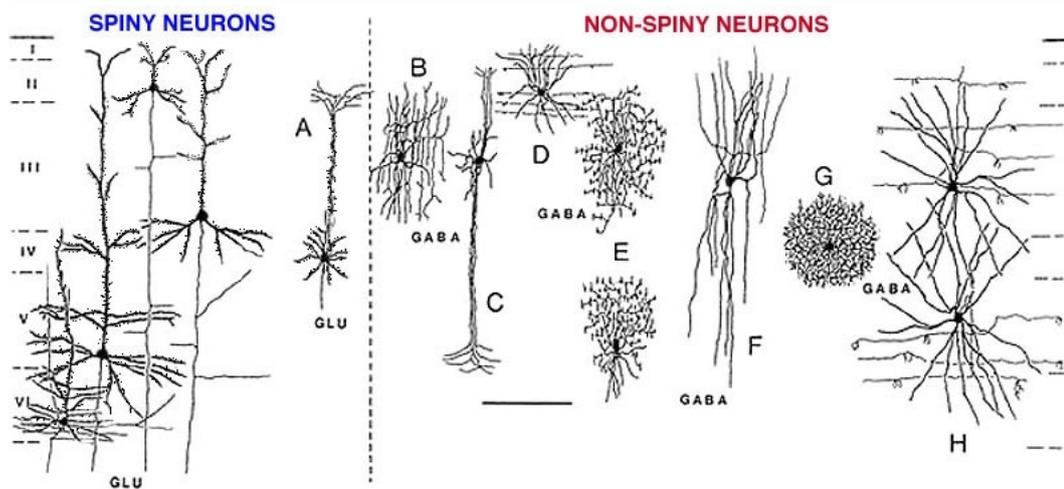


Figure 1.5: Major cell types in the monkey cerebral cortex. These can be divided broadly into spiny (excitatory) and non-spiny (inhibitory) neurons. Spiny neurons include pyramidal cells (left most cells) and **A** spiny stellate neurons, which both use glutamate as their primary neurotransmitter. A variety of smooth inhibitory cells which use GABA as their primary neurotransmitter are also shown, including **B** a cell with local axon arcades, **C** double bouquet cell, **D**, **H** basket cells, **E** chandelier cells, **F** bitufted cell and **G** neurogliaform cell . Figure reproduced from Schmolesky (2005).

Both pyramidal and spiny stellate cells possess dendritic spines, which are small membranous protrusions from the dendrites that act as the postsynaptic sites for the majority of the cell's glutamatergic synapses (Spruston, 2008). The soma of these cells generally only receive inhibitory synapses (Toledo-Rodriguez et al., 2003).

Limited electrophysiological diversity has been observed between excitatory cell types in the cortex. The most common discharge response of excitatory cells to sustained input currents has been labelled 'regular-spiking', in which a neuron fires repeatedly with decreasing frequency, called spike train adaptation or accommodation (Toledo-Rodriguez et al., 2003). Some pyramidal and spiny stellate cells also exhibit intrinsic bursting behaviour, characterised by an initial rapid burst of spikes followed by regular single spikes or bursts (Connors and Gutnick, 1990). Rarer subtypes of pyramidal cells have been observed to display rhythmic firing behaviour without accommodation, or clustered 'chattering' activity patterns (Toledo-Rodriguez et al., 2003).

Excitatory cell types have been shown to have characteristic laminar distributions, and can be divided into layer-specific broad classes within which several subclasses exist (Tasic et al., 2016; Harris and Shepherd, 2015; Zeisel et al., 2015). Layer II/III excitatory neurons have often been thought of as a homogeneous cell type, but evidence suggests they consist of a number of yet-to-be completely classified subtypes, which vary in their firing activity, gene expression and axonal targets (Molyneaux et al., 2009; Harris and Mrsic-Flogel, 2013; Tasic et al., 2016). Layer IV excitatory cells consist of cells from the two major morphological classes, pyramidal and spiny stellate cells. These cells project most strongly to layer II/III, and contain sub-populations of regular spiking and intrinsically bursting cells (Harris and Mrsic-Flogel, 2013; Staiger et al., 2004).

Layer V excitatory neurons can be divided into two major subclasses based on their patterns of axonal projection. Intratelencephalic (IT) layer V neurons are generally found in the upper part of layer V and project axons exclusively within the telencephalon, which consists of the neocortex, striatum and corticoid structures (Reiner et al., 2003). They generally have slender apical dendrites and exhibit adaptive firing activity (Harris and Mrsic-Flogel, 2013; Molnár and Cheung, 2006). Pyramidal tract (PT) neurons, also known as subcerebral projection neurons, are large pyramidal neurons which project to subcerebral locations (Harris and Shepherd, 2015; Molyneaux et al., 2007). These cells generally have thick apical dendrites and prominent dendritic tufts in layer I, and exhibit limited spike train adaption (Harris and Mrsic-Flogel, 2013).

Layer VI contains a morphologically diverse population of excitatory neurons that can be divided into at least two broad classes, within which a number of subclasses exist (Briggs, 2010). Corticothalamic (CT) layer VI neurons project principally to the ipsilateral thalamus as well as layer IV (Harris and Shepherd, 2015; Briggs, 2010). These cells generally have short upright pyramidal morphologies with local axonal arbours (Thomson, 2010). Cortico-cortical layer VI cells project only within the cortex, have small dendritic trees with a range of morphologies and long range horizontally orientated axons (Thomson, 2010; Harris and Mrsic-Flogel, 2013).

1.4.2 Inhibitory interneurons

The vast majority of inhibitory neurons in the cortex are ‘interneurons’, named as such due to the fact that their axonal and dendritic arbours branch only within the

neocortex (Markram et al., 2004). As well as their GABAergic nature, inhibitory interneurons can also be distinguished from excitatory cell types by their absence of an apical dendrite, low spine density and locally projecting axons (Toledo-Rodriguez et al., 2003). Also, unlike excitatory cells, the soma of inhibitory interneurons can receive both excitatory and inhibitory synapses (Toledo-Rodriguez et al., 2003). Inhibitory interneurons are highly heterogeneous in term of their dendritic morphology and physiological and molecular characteristics. Attempts have been made to classify neurons based on morphological, electrophysiological and molecular criteria, however, the overlapping nature of these features in many interneuron subclasses has hindered efforts to develop a unambiguous classification (Ascoli et al., 2008; Kelsom and Lu, 2013) (see Figure 1.6).

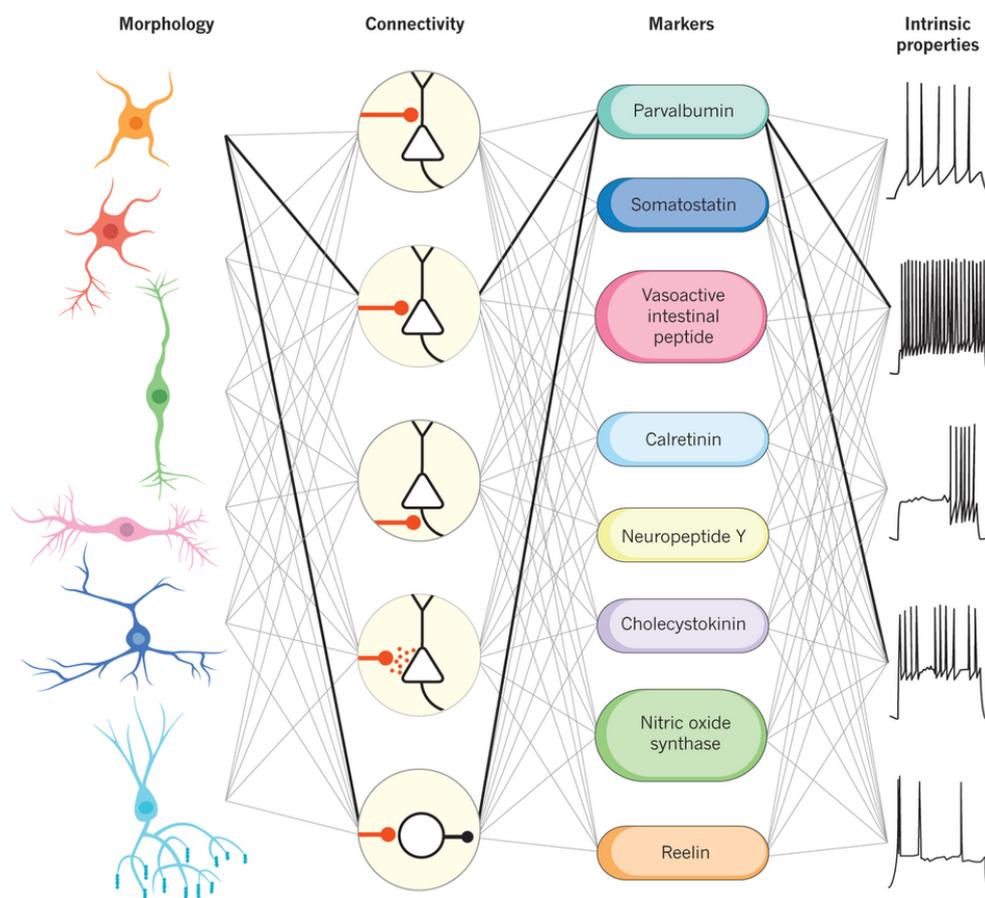


Figure 1.6: Classification of interneurons is generally based on a combination of their morphology, connectivity, synaptic properties, expression of molecular markers and electrophysiological properties. The bold connections represent fast-spiking basket cells, which express parvalbumin and target the soma of pyramidal cells and interneurons. Figure reproduced from Kepecs and Fishell (2014).

GABAergic interneurons can be classified into a number of morphologically defined cell types (see Figure 1.5, right for examples), including:

- **Basket cells:** Basket cells make up $\sim 50\%$ of inhibitory interneurons, and preferentially target the soma and proximal dendrites of both pyramidal cells and interneurons (Markram et al., 2004). Three main subclasses of basket cells have been defined morphologically. These are large basket cells, which have large expansive axonal arbours; small basket cells, which generally have local dense axonal clusters; and nest basket cells, which can have sparse or dense local arbours with infrequent long horizontal axonal branches and simple short dendritic trees (Toledo-Rodriguez et al., 2003; Wang et al., 2002).
- **Bipolar, double bouquet and bitufted cells:** These three cell types generally have ovoid soma and dendrites that emerge from opposite poles of the soma, forming narrow bipolar or bitufted morphologies (Markram et al., 2004). These cell types, however, differ in their axonal morphology. Bipolar cells generally have narrow vertical axonal projections, while double bouquet cells have axons that extend vertically in a tight ‘horse tail’-like bundle (Markram et al., 2004; Schmolesky, 2005) (Figure 1.5C). Bitufted cells generally have wider horizontal axonal spans than bipolar or double bouquet cells, that can extend into neighbouring cortical columns (Toledo-Rodriguez et al., 2003) (Figure 1.5F).
- **Chandelier cells:** Chandelier cells can be multipolar or bitufted and exclusively target the axons of postsynaptic cells (Markram et al., 2004; Kelsom and Lu, 2013). They are characterised by local axonal clusters which branch frequently into vertical sections, resembling a chandelier (Toledo-Rodriguez et al., 2003) (Figure 1.5E).
- **Martinotti cells:** Martinotti cells generally have the most elaborate dendritic projections of all interneuron types (Toledo-Rodriguez et al., 2003). They are characterised by axons that project to layer I, where they branch extensively (Rudy et al., 2011). These cells preferentially target the dendrites and dendritic tufts of postsynaptic cells (Markram et al., 2004).
- **Neurogliaform cells:** Neurogliaform cells are small neurons that have dense local spherical dendritic and axonal fields and primarily target dendritic shafts (Toledo-Rodriguez et al., 2003) (Figure 1.5G).
- **Cajal-Retzius cells:** Cajal-Retzius cells are multipolar cells that are unique to layer I and have long horizontal dendrites and horizontally projecting axons

that are confined to layer I (Markram et al., 2004).

- **Small layer I cells:** A heterogeneous group of small multipolar neurons which vary in their axonal arbourisation patterns (Toledo-Rodriguez et al., 2003).

Inhibitory neurons also exhibit diversity in their electrophysiological properties, and can be classified into several major electrophysiological subclasses based on both their onset and steady state responses to sustained step-current injections to the soma (Toledo-Rodriguez et al., 2003; Ascoli et al., 2008) (see Figure 1.7 for details).

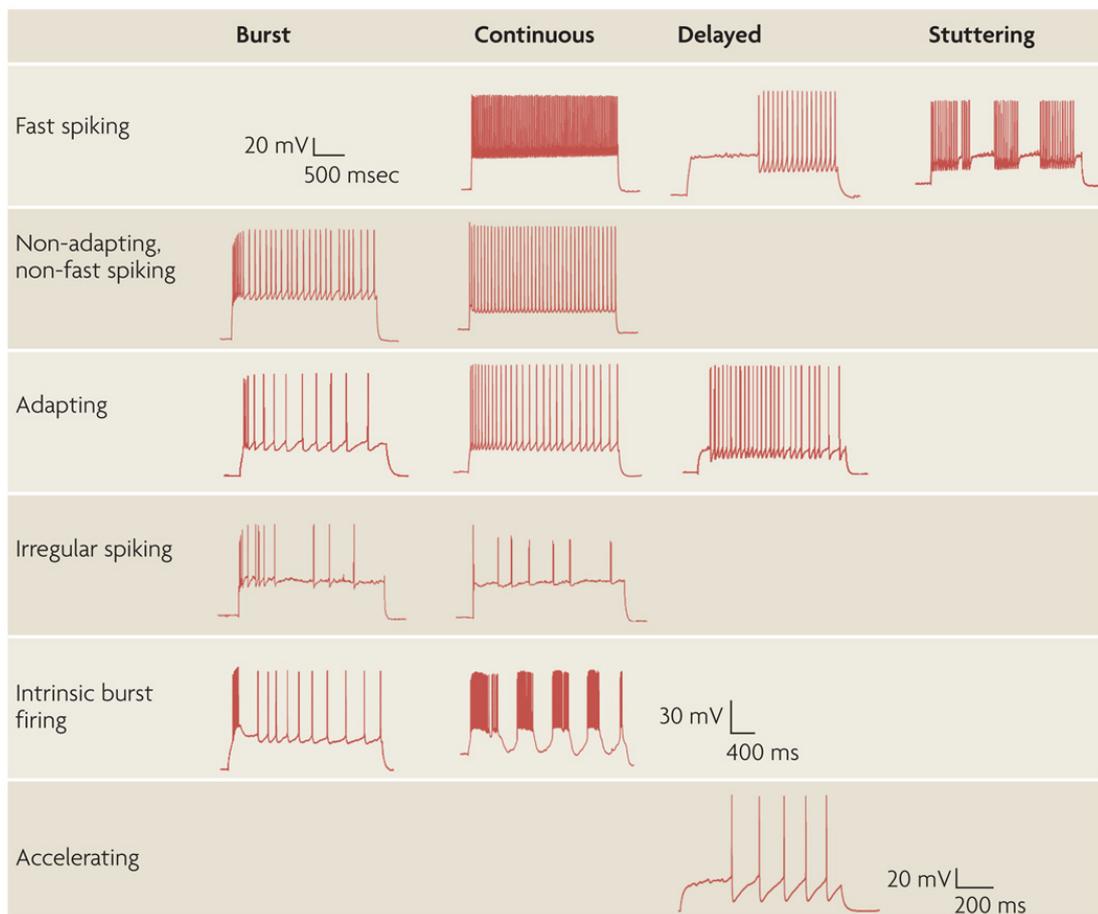


Figure 1.7: The firing pattern of neurons is characterised by both their steady state response (rows) which includes fast spiking, non-adapting non-fast spiking, adapting, irregular spiking, intrinsic bursting and accelerating, as well as their onset response to a stepped depolarization (columns), which can be bursting, continuous, delayed or stuttering. For example, a delayed adapting neuron (row three, column three) would exhibit a delay between the onset of the injected current and the discharge of the cell, followed by action potentials which gradually decrease in frequency (adaptation) (Toledo-Rodriguez et al., 2003). A burst non-adapting interneuron (row two, column one) initially fires a ‘burst’ of spikes in quick succession, followed by repeated spikes that have no frequency adaptation. Figure reproduced from Ascoli et al. (2008).

The cells in the previously detailed morphologically identified classes can fall into a range of these electrophysiologically defined classes. For example, different Martinotti cells have been seen to exhibit continuous non-adapting, bursting adapting, continuous adapting or continuous irregular spiking firing behaviours (Toledo-Rodriguez et al., 2003).

Inhibitory interneurons can also be classified molecularly. Recently, it has been shown that GABAergic interneurons can be divided into three broad subgroups based on their expression of three molecular markers, the calcium binding protein parvalbumin (PV), the neuropeptide somatostatin (SST) and the ionotropic serotonin receptor 5HT3a (5HT3aR) (Lee et al., 2010). These subgroups are non-overlapping, and together account for almost 100% of total neocortical inhibitory interneuron diversity (Rudy et al., 2011). These three broad classes, in particular the 5HT3aR-expressing class, represent heterogeneous populations of cells that express other proteins and neuropeptides, and exhibit morphological and electrophysiological diversity that can be used to further subdivide these classes (Kelsom and Lu, 2013; Kubota et al., 2011).

Approximately 40% of GABAergic neurons express PV (Rudy et al., 2011). These cells have been shown to have low input resistance and fast membrane time constants compared to other interneuron types, and largely exhibit fast spiking firing patterns (Ascoli et al., 2008; Cauli et al., 1997; Kawaguchi and Kubota, 1997; Rudy et al., 2011). PV-expressing interneurons are comprised largely of two morphologically defined cell types, basket and chandelier cells (Ascoli et al., 2008; Harris and Mrsic-Flogel, 2013; Kawaguchi and Kubota, 1997).

Around 30% of inhibitory interneurons in the neocortex are SST-expressing, the largest subtype of which have Martinotti morphologies (Kepecs and Fishell, 2014). SST-expressing Martinotti cells are found in layers II-V, and most commonly exhibit continuous adapting firing patterns (Rudy et al., 2011). Non-Martinotti SST-expressing cells have also been observed, which include cells with bitufted, ‘basket’ and long-projecting morphologies and irregular spiking, stuttering or fast spiking firing patterns (Urban-Ciecko and Barth, 2016; Yavorska and Wehr, 2016). The observed diversity in firing patterns, morphology, as well as the molecular characteristics and connectivity of SST-expressing interneurons suggests that there are many yet unclassified subtypes in this population, which could number as high as one hundred (Ma et al., 2006; McGarry et al., 2010; Rudy et al., 2011; Urban-Ciecko and Barth, 2016).

The remaining inhibitory neurons in the cortex are largely 5HT3aR-expressing interneurons, which form the most populous interneuron type in the superficial layers (layers I–III) (Lee et al., 2010). These cells represent the most heterogeneous class of molecularly defined interneurons, with subclasses that have yet to be completely characterised. A subset of 5HT3aR-expressing neurons have been found to also express the neuropeptide vasoactive intestinal peptide (VIP) (Lee et al., 2010). These VIP-expressing interneurons are thought to be composed of several cell types, including those with bipolar, bitufted and multipolar morphologies and a variety of intrinsic spiking patterns (Rudy et al., 2011; Kelsom and Lu, 2013; Xu et al., 2010; Kawaguchi and Kubota, 1997).

There are also approximately 60% of 5HT3aR-expressing interneurons which are VIP-negative, 80% of which express the interneuronal marker reelin (Lee et al., 2010; Miyoshi et al., 2010). A prominent group of these cells are morphologically defined as neurogliaform cells, which generally have small spherical dendritic fields and late spiking firing patterns (Toledo-Rodriguez et al., 2003; Rudy et al., 2011). Other 5HT3aR-expressing, VIP-negative, reelin-positive cell types have also been identified, which exhibit a variety of electrophysiological characteristics (Rudy et al., 2011; Miyoshi et al., 2010; Lee et al., 2010).

Figure 1.8 shows one possible incomplete classification of cortical inhibitory interneurons from Rudy et al. (2011). Several of these subtypes are still poorly defined, or it is unknown if they represent true interneuron classes. For example, it has been suggested that the SST-expressing X94 cell type shown in Figure 1.8 may not be a true interneuron subtype, but instead represent a subset of a single cell type labelled ‘L4 SST cells’ that includes both X94 and non-X94 cells (Yavorska and Wehr, 2016; Xu et al., 2013).

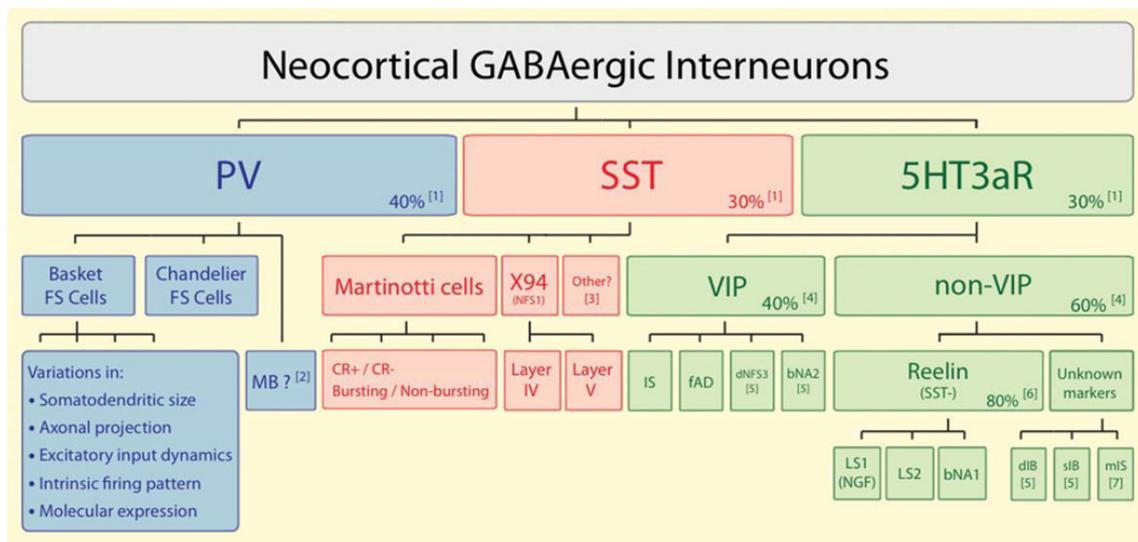


Figure 1.8: A possible incomplete classification of neocortical inhibitory interneurons by Rudy et al. (2011). Cells are divided into three main subclasses, based on their expression of PV, SST or 5HT3aR. Each of these classes can be further subdivided based on a variety of features, including molecular, morphological and electrophysiological properties, however there still exist a number of poorly defined classes. A full description is given in Rudy et al. (2011), from which this figure was reproduced.

1.5 Spontaneous activity in developing neuronal networks

Of equal importance to the types of neurons that make up a neuronal network to its function, are the number and configuration of the synaptic connections between them. Correct wiring of the over 100 trillion connections between neurons in the human nervous system is essential for functions ranging from the processing of sensory information to fine motor movement (Squire et al., 2008). Activity dependent rewiring and synaptic changes, in which the relative activity patterns between pairs of neurons dictates the strength of their connection, are thought to play a critical role in the development and maintenance of these neuronal networks (Butz et al., 2009). Postnatally, this activity is driven by sensory input and experience (Chapman and Stryker, 1993). However, in immature networks during development, it is believed that spontaneous activity assumes this role (Khazipov and Luhmann, 2006) (Figure 1.9).

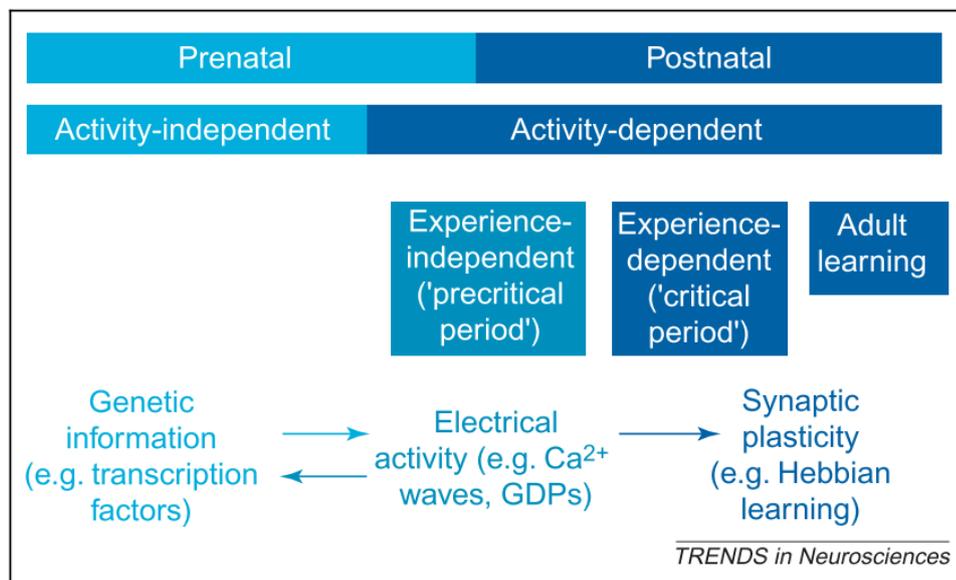


Figure 1.9: Factors influencing the development of neuronal networks. Embryonically, genetic information is used to lay down crude maps. Later (precritical period), mechanisms such as gap-junctions and extrasynaptic neurotransmission result in correlated network activity that shapes the gross and fine structure of neuronal networks. During the critical period, experience-dependent modifications to neuronal networks occur in a Hebbian manner. Figure reproduced from Khazipov and Luhmann (2006).

1.5.1 Mechanisms of spontaneous activity

Spontaneous activity has been observed extensively in numerous parts of the mammalian nervous system, including the hippocampus, cortex, spinal cord, retina, cochlea and higher areas of the visual system (Garaschuk et al., 1998; Hanson and Landmesser, 2003; Meister et al., 1991; Weliky and Katz, 1999; Corlew et al., 2004; Tritsch et al., 2007). Despite significant differences in their structure and function, the nature of the spontaneous activity patterns in these networks and the physiological mechanisms underlying this activity show strong similarities (Blankenship and Feller, 2010).

Although the complete physiological processes underlying the generation of spontaneous activity across the different regions of the nervous system are not well understood, several mechanisms have frequently been observed to be involved. These include the depolarizing activity of GABA and glycine during early development, extrasynaptic neurotransmission, and the presence of gap junctions and transient connections between cells (Ben-Ari, 2001; Blankenship and Feller, 2010; Kerschensteiner, 2014; Syed et al., 2004; Zheng et al., 2006). As the physiological properties of neuronal networks change during development, for example the switch from electrical to chemical synapses, the dominant mechanisms underlying the generation of spontaneous activity are also believed to vary, which is reflected by significant changes in the nature of the spontaneous activity patterns. In most brain regions, this involves a shift from repetitive activity synchronised across small networks to more complex large-scale patterns of activity (Luhmann et al., 2016; Khazipov and Luhmann, 2006; Crépel et al., 2007).

One well studied area that exhibits this phenomenon is the hippocampus. Early during development, the spontaneous activity in the hippocampus takes the form of synchronous plateau assemblies. These consist of small groups of neurons connected via gap junctions, which exhibit bursts of intracellular calcium transients (Crépel et al., 2007; Garaschuk et al., 1998). Later, spontaneous activity in the hippocampus transitions to a form of spreading waves called giant depolarizing potentials (GDPs), which involve slow depolarizations correlated across larger networks of neurons and are dependent on GABAergic and glutamatergic neurotransmission (Leinekugel et al., 1997; Ben-Ari et al., 1989). These GDPs are believed to be initiated primarily in the CA3 region of the hippocampus, where GABAergic interneurons depolarize pyramidal cells, producing periodic bursts of action potentials that propagate throughout the hippocampus (Ben-Ari et al., 1989; Ben-Ari et al., 2007).

Although spontaneous activity in developing neuronal networks has most commonly been studied in rodent, this phenomenon has also been observed in a range of mammalian species, including ferrets, monkeys and humans, with remarkable similarity across species (Khazipov and Luhmann, 2006). In humans, bursting patterns of spontaneous activity have been observed in the cerebral cortex from mid-gestation to just after birth (Tolonen et al., 2007; Moore et al., 2011). Electroencephalogram recordings of human neonates at the later stages of gestation have also shown the presence of delta-brush patterns consisting of rapid bursts, which share several features with the early patterns of spontaneous synchronised activity observed in neonatal rodent cortex (Vanhatalo and Kaila, 2006; Khazipov and Luhmann, 2006).

1.5.2 The role of spontaneous activity

As well as being conserved across numerous brain regions and species, spontaneous activity during development has also been observed to exhibit remarkable robustness to genetic and pharmacological manipulations, suggesting that redundancies exist in these networks to ensure that spontaneous activity is maintained (Chub and O'Donovan, 1998; Charlesworth et al., 2016). When spontaneous activity is successfully disrupted, this has been shown to lead to deficits in network structure (Cang et al., 2005; Clause et al., 2014; Hanson and Landmesser, 2004; Huberman et al., 2006). Together, these observations suggest that spontaneous activity plays a crucial role in neuronal network development.

Although the precise function of developmental spontaneous activity across the nervous system is still an area of investigation, a role for spontaneous activity has been implicated in numerous physiological processes in developing neuronal networks, including:

- **Neurogenesis:** Spontaneous electrical activity has been observed to play a role in regulating the production and proliferation of neuronal cells during early development, both directly as well as indirectly, such as through the regulation of morphogenic factors (Loturco et al., 1995; Weissman et al., 2004; Belgacem and Borodinsky, 2015). Suppression of retinal waves in embryonic mice has been shown to increase neurogenesis and alter the structure of cortical layers (Bonetti and Surace, 2010).
- **Apoptosis:** Bursting patterns of spontaneous activity are also believed to have an important role in regulating cell survival. High frequency bursting has been

shown to increase neuronal survival in cortical cultures, while suppression of spontaneous activity has been observed to greatly increase rates of programmed cell death (Golbs et al., 2011; Heck et al., 2008).

- **Neuronal migration:** Neurons are produced in a limited set of locations in the brain, from which they must migrate to appropriate sites before being incorporated into local networks (Purves et al., 2001). The rate and extent of this migration has been shown to be influenced by temporal changes in spontaneous Ca^{2+} transients and neurotransmitters whose expression is regulated by neuronal activity (Komuro and Kumada, 2005; Luhmann et al., 2015). Additionally, GABAergic interneurons have been observed to cease their migration at a similar time to which they begin to participate in synchronised network activity (de Lima et al., 2009).
- **Neuronal differentiation:** Dendritic differentiation has been shown to be regulated by spontaneous Ca^{2+} transients and activity-dependent release of transcription factors that are essential for dendritic development (Chen and Ghosh, 2005; Bando et al., 2016). Spontaneous electrical activity has also been shown to be involved in axon pathfinding, with disruption of this activity shown to cause pathfinding errors and alterations in molecular cues believed to be involved in axon guidance (Hanson and Landmesser, 2004; Kastanenka and Landmesser, 2013).
- **Specification of neurotransmitters:** Changes in spontaneous activity patterns have been associated with the alteration of neurotransmitter phenotype (Spitzer, 2012). Increasing calcium spike frequency has been observed to increase (decrease) the expression of inhibitory (excitatory) neurotransmitters in embryonic spinal cord (Borodinsky et al., 2004).
- **Network connectivity:** Patterned spontaneous activity has been shown to play an important role in the development of precise connectivity patterns essential for proper functioning of neuronal networks. In particular, spontaneous activity has been observed to be instructive in the formation of topographic maps in sensory systems, and the disruption of this activity has been shown to lead to deficits in both the gross organisation and precise refinement of these networks (Mrsic-Flogel et al., 2005; Cang et al., 2005; Clause et al., 2014; Xu et al., 2011).

1.5.3 Recording of spontaneous activity on microelectrode arrays

As well as being observed extensively *in vivo*, correlated spontaneous activity patterns have frequently been recorded *in vitro*, where they also exhibit increasingly complex activity patterns over developmental time (Sun et al., 2010; Wagenaar et al., 2006a; Baker et al., 2006). A common method for studying spontaneous activity during development is using recordings of the electrical activity of cultured neurons on microelectrode arrays (MEAs, also known as multielectrode arrays) (Sun et al., 2010; Chen et al., 2006; Pasquale et al., 2008; Van Pelt et al., 2005; Chiappalone et al., 2006). These devices consist of planar grids of electrodes, which facilitate the simultaneous and long-term recording of extracellular action potentials and local field potentials from neurons with high temporal resolution (Obien et al., 2015) (Figure 1.10).

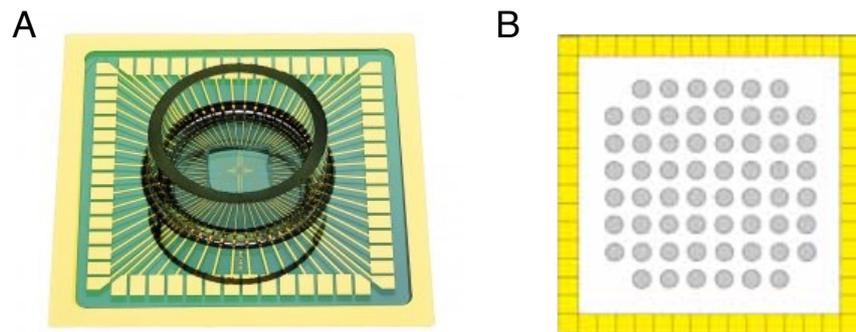


Figure 1.10: **A** A 64-channel MEA from Multi Channel Systems, and **B** configuration of the 8×8 layout of electrodes in a 64-channel array. Standard electrode diameters are either $10 \mu\text{m}$ or $30 \mu\text{m}$ and the spacing of electrodes is either $100 \mu\text{m}$ or $200 \mu\text{m}$. Images obtained from Multi Channel Systems MCS GmbH (2015).

MEAs were initially developed more than four decades ago to allow for non-destructive simultaneous recording of networks of cultured cells (Thomas Jr. et al., 1972; Gross et al., 1977; Pine, 1980). Since this time, improved semiconductor lithographic techniques have led to the development of more reliable MEA devices with increased numbers and densities of electrodes (Obien et al., 2015). Modern high density MEAs can contain over 10 000 electrodes, allowing for the recording of activity at sub-cellular resolution (Spira and Hai, 2013; Frey et al., 2009).

The microelectrodes on an MEA record changes in the extracellular voltage produced by the electrical activity of nearby cells, including both action potentials and local field potentials (Buzsáki et al., 2012). The strongest of these currents are the fast action potentials described in Section 1.2.2, which are recorded extracellularly as ‘spikes’. Due to the rapid influx of Na^+ into the cell during an action potential, a large negative spike in the extracellular potential is observed, followed by a small positive spike as K^+ ions exit the cell (Obien et al., 2015). These extra-cellular action potentials (EAPs) have similar durations to intracellular action potentials, but generally have amplitudes that are several orders of magnitudes lower (Buzsáki et al., 2012) and can be recorded by electrodes placed near ($\sim 100\mu\text{m}$) the origin of the spike (Egert et al., 2002). As well as EAPs, MEAs also record local field potentials, which are the low-frequency component of extracellular voltage, resulting from the dynamic activity of neurons surrounding the electrode, as well as noise (Rey et al., 2015). Noise can arise from biological sources, such as the action potentials of distant cells and the subthreshold events of nearby neurons, as well as experimental factors such as the liquid-metal interface of the MEA (Obien et al., 2015).

Extracting the spiking activity of individual neurons from the raw voltage data recorded by MEA electrodes involves a number of steps. The first is to use band-pass filtering, generally with a band in the range of 300–3000 Hz, to separate EAPs from local field potentials and noise (Quian Quiroga, 2009). Next, spikes are detected, generally using a threshold set at a multiple of the baseline noise level (Rey et al., 2015; Rizk and Wolf, 2009). These spikes consist of the superimposed spiking activity of multiple neurons positioned at different distances from the electrode, and spike sorting techniques can be used to differentiate the spikes originating from each neuron (Quian Quiroga, 2012).

Using dissociated cultures plated on MEAs to study neuronal network development has several advantages. A significant benefit of MEAs is their ability to record simultaneously from multiple neurons in a network (Spira and Hai, 2013). Other techniques, such as simultaneous patch clamp recordings, have been recently developed to allow for simultaneous intracellular recordings of neurons, however, they are limited in the number of simultaneous recordings they allow (~ 12 neurons in vitro), compared to MEAs which can contain several thousand electrodes (Obien et al., 2015; Perin et al., 2011). The non-invasive nature of MEAs also facilitate long term recordings, allowing for repeated incubation and recording from the same cultures (Breckenridge et al., 1995). Cultured networks have been shown to survive and

exhibit reliable spontaneous activity for over a year on MEAs, greatly exceeding the short recording times offered by other experimental techniques (Potter and DeMarse, 2001). The use of glass arrays and the monolayer nature of the cultured neurons also facilitates the imaging of cultured networks on MEAs, allowing for changes in neuronal morphology to be monitored (Potter et al., 2001). MEAs also allow for the manipulation of network activity using electrical stimuli delivered through the microelectrodes, the effect of which can be monitored effectively in the small cultured networks (Wagenaar et al., 2004; Hales et al., 2010).

MEAs, however, also have some limitations. The extracellular nature of the recordings from the microelectrodes results in smaller amplitude recordings and the inability to measure subthreshold variations in the membrane potential of neurons, which can only be detected using intracellular recording techniques such as patch clamp (Henze et al., 2000). The flexibility of MEA recordings is also limited by the fixed location of their electrodes, and MEAs can exhibit greater sensitivity to variations in fluid level than other instruments (Whitson et al., 2006). Despite this, network recordings on *in vitro* MEAs remain a common method for studying the spontaneous electrical activity of neuronal networks.

1.5.4 Properties of cultured neuronal networks on MEAs

In vitro cultured neuronal networks are generally produced from dissociated cells from the brains of early postnatal or embryonic animals (Marom and Shahaf, 2002). Here I will focus on describing the developmental properties of rat cortical cultures, which are most commonly used in experimental studies. Within hours of plating, the dissociated neurons begin producing new neurites and rapidly progress from a population of unconnected independent neurons to a densely connected network (Huettnner and Baughman, 1986; Van Pelt et al., 2005; Wagenaar et al., 2006a). For cortical cultures in the absence of cell purification, cultured networks have been found to contain similar distributions of cell types to that of *in vivo* networks (Nakanishi and Kukita, 2000; Huettnner and Baughman, 1986). From the first week *in vitro*, synapses begin to form, and rapidly increase in density over the following weeks (Habets et al., 1987; Ichikawa et al., 1993). This increase in synaptic density has been shown to be correlated with increasing frequency of spontaneous activity in the cultured networks (Ito et al., 2013; Muramoto et al., 1993).

Spontaneous activity in cultured cortical networks has been observed as early as 3 days *in vitro* (DIV), and typically initially takes the form of uncorrelated firing

of neurons (Kamioka et al., 1996; Wagenaar et al., 2006a). After a week *in vitro*, synchronous regular bursting of neurons, termed ‘network bursts’, become the dominant activity pattern in the cultured networks (Chiappalone et al., 2005; Van Pelt et al., 2004b). Initially, this activity consists of long bursts lasting several hundreds of milliseconds and separated by long periods of quiescence (Kamioka et al., 1996). These bursts increase in frequency and size before reaching a peak around 3 weeks *in vitro* (Van Pelt et al., 2004b; Van Pelt et al., 2004a; Chiappalone et al., 2006). This peak in network bursting activity corresponds to the period in which synaptic density in the network reaches its maximum (Van Huizen et al., 1985; Kamioka et al., 1996; Van Pelt et al., 2004b). This is followed by a period of shortening of network bursts, which coincides with a stage of ‘pruning’ of dendritic spine synapses and maturation of excitatory connections between neurons (Chiappalone et al., 2006; Ichikawa et al., 1993; Van Pelt et al., 2005).

At later ages (> 30 DIV), the networks appear to reach a stable network state, characterised by complex and non-periodic bursting activity with properties that remain stable over long developmental periods (Chiappalone et al., 2006; Kamioka et al., 1996; Van Pelt et al., 2004b). As well as in rat cortical cultures, the presence of network bursting activity has also been observed in a variety of other brain regions and species *in vitro* (Van Den Pol et al., 1996; Ben-Ari, 2001; Rhoades and Gross, 1994; Harris et al., 2002; Meister et al., 1991) and *in vivo* (Chiu and Weliky, 2001; Leinekugel et al., 2002; Weliky and Katz, 1999).

Although animal models provide a useful tool for studying spontaneous developmental activity, human cortical development has been observed to exhibit certain species-specific features, and the translational value of animal models for studying human neuronal development, particularly in diseased states, remains poor (Suzuki and Vanderhaeghen, 2015; Kuijlaars et al., 2016). Recently, *in vitro* studies of the spontaneous activity of developing human neuronal networks on MEAs have been performed using embryonic or induced pluripotent stem-cell derived neurons (Heikkilä et al., 2009; Odawara et al., 2016; Amin et al., 2016). Synchronous bursting resembling that in rat cortical cultures was observed in these networks, generally arising 8 to 12 weeks after differentiation and increasing in frequency over development.

Using calcium imaging, Kirwan et al. (2015) also observed similar developmental progression of spontaneous activity in networks of human pluripotent stem cell-derived neurons to that previously seen in rodent cortical cultures. This began

with synchronised bursting of small networks of neurons after ~ 60 days in culture, progressing to larger bursts with longer interburst periods. The frequency and size of these bursts peaked around 70 days in culture, after which the frequency of correlated bursts decreased and were eventually replaced by more complex activity patterns. These results suggest that human stem cell-derived cultures form functional neuronal networks that exhibit stereotypical patterns of developmental spontaneous activity and have the potential to replace animal models for studying spontaneous activity during development (Kirwan et al., 2015; Amin et al., 2016).

1.6 Spike trains

As mentioned in Section 1.5.3, studying the electrical activity of neurons using MEAs frequently involves reducing the raw voltage measurements recorded by each electrode to a series of binary events or ‘spikes’ that represent the time of the peaks of the recorded action potentials (see Figure 1.11 for examples). This is based on the idea that action potentials convey information primarily through their timing, rather than any feature of the shape of the action potential waveform (Purves et al., 2001).

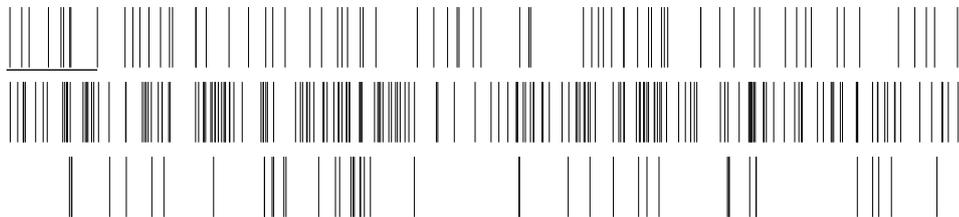


Figure 1.11: Examples of three spike trains from MEA recordings of mouse cortical neurons. Each row shows 10s of spiking activity recorded from one electrode, where each vertical line represents the time of a spike. Scale bar represents 1 s. Experimental recordings by Paul Charlesworth from the Department of Physiology, Development and Neuroscience.

Statistical analysis of both single ‘spike trains’ as well as multiple simultaneously recorded spike trains from several neurons can reveal important details about the mechanisms of neuronal signalling and the interactions between neurons, as well as how neurons encode information (Perkel et al., 1967a). Analysis of a single spike train typically involves examining a variety of first and second order statistics such as the instantaneous mean firing rate, Fano factor and autocorrelation of the spiking activity, as well as features of the interspike interval (ISI) distribution (Perkel et al., 1967a; Purves et al., 2001). If simultaneously recorded spike trains from several neurons are available, this can be expanded to examining features of the correlation

and associations between the spiking activity of the multiple neurons (Perkel et al., 1967b; Brown et al., 2004).

As spikes are regarded as instantaneous and indistinguishable apart from their temporal location, spike trains can be modelled mathematically as stochastic point processes (Perkel et al., 1967a). The simplest model used to approximate the stochastic behaviour of neuronal firing is the homogeneous Poisson process, for which the spike count in any interval on the spike train is assumed to be Poisson distributed with rate λ (Purves et al., 2001; Tuckwell, 1988). This requires the simplifying assumption of a stationary firing rate as well as the independence of individual spike times. Spike trains have also commonly been modelled using inhomogeneous Poisson processes, in which the firing rate is a time dependant function, $\lambda(t)$ (Purves et al., 2001; Perkel et al., 1967a).

Individual spike trains from real neurons rarely correspond to Poisson processes, due to various factors, such as the presence of a refractory period, bursting, and long range temporal correlations between spikes, and several other models of neuronal activity have been developed to incorporate these features (Barbieri et al., 2001; Averbeck and Lee, 2003; Câteau and Reyes, 2006; Pillow, 2009; Kass and Ventura, 2001; Brown et al., 2001). Despite this, the Poisson process is still a useful model for approximating stochastic neuronal spiking activity, and in particular for describing the trial-averaged response of a neuron to repeated presentations of a stimulus (Pouzat and Chaffiol, 2009).

Analysis of such spike trains recorded from neurons in response to certain stimuli are commonly used in studies of neuronal coding, which aim to uncover how neurons use spikes to encode and transmit information (Aljadeff et al., 2016; Awiszus, 1997). An open question in this field is whether it is simply the overall spike rate of a neuron that carries information about the features of stimuli ('rate coding'), or if the precise pattern of spike timing also plays an important role in information transmission ('temporal coding') (Brette, 2015).

1.7 Bursting in neuronal networks

The analysis of periods of 'bursting' in spike trains has been a component of spike train analysis for several decades (Smith and Smith, 1965; Legéndy and Salcman, 1985). Although many studies of developmental spontaneous activity focus on the analysis of network bursts, single cell bursts have also been observed commonly in

MEA recordings from cultures of dissociated neurons, as well as in numerous *in vitro* systems (Wagenaar et al., 2006a; Pasquale et al., 2008; Weyand et al., 2001; Legéndy and Salcman, 1985). These single cell bursts consist of brief periods of high frequency firing of a neuron separated by periods of quiescence (see Figure 1.12 for an example), and can be produced by strong dynamic inputs or generated by intrinsic ionic mechanisms of certain cell types (Lisman, 1997; Izhikevich et al., 2003).

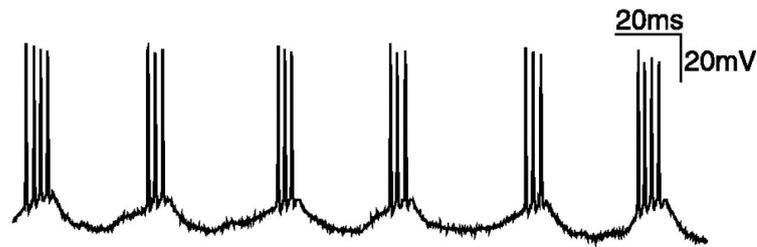


Figure 1.12: Voltage trace from a rhythmically bursting neuron, recorded from cat visual cortex *in vivo*. Figure adapted from Wang (2010).

Although this type of bursting is widely believed to play an important role in information transmission in the brain, the purpose of bursts and the precise mechanisms underlying their generation are only partially understood (Lisman, 1997). In the following sections, I present the current understanding of several burst-related phenomena.

1.7.1 The functional role of bursting over single spikes

Central synapses in various brain regions have been shown to exhibit low probabilities of neurotransmitter release in response to single presynaptic spikes, making information transfer by single spikes unreliable (Borst, 2010; Branco and Staras, 2009; Allen and Stevens, 1994). However, rapid bursts of spikes can lead to a build-up of intracellular Ca^{2+} in the presynaptic terminal, as Ca^{2+} levels have insufficient time to return to their baseline levels between each spike (Thomson, 1997; Wu and Saggau, 1994; Tank et al., 1995). This increases the probability of neurotransmitter release and the resultant production of EPSPs with subsequent spikes (Krahe and Gabbiani, 2004) (Figure 1.13a, left).

This process of ‘facilitation’ has been observed experimentally for a variety of synapses (Snider et al., 1998; Csicsvari et al., 1998; Thomson, 2000a). This includes hippocampal CA1 synapses, for which the probability of postsynaptic response caused by a single spike was found to be 0.2 – 0.6 when a neuron was stimulated with closely

spaced spikes (Stevens and Wang, 1995). In the absence of a response from the first spike, the probability of neurotransmitter release at the time of the second spike exceeded 0.9. However, in the case that a postsynaptic response was induced by the first spike, the probability of a second release of neurotransmitters at the time of the second spike was < 0.1 , suggesting that the synapses were experiencing depression (Figure 1.13a, right).

Depression caused by bursting activity has also been observed in other synapses that have high probabilities of neurotransmitter release in response to single spikes, such as somatosensory thalamocortical synapses (Chung et al., 2002; Gil et al., 1997). It has been suggested that this interaction of synaptic facilitation and depression could serve to produce a deterministic postsynaptic response to bursts (i.e. one single postsynaptic spike produced by each burst), as facilitation increases the probability of at least one neurotransmitter release, however, depression greatly decreases the probability of a second release (Lisman, 1997) (Figure 1.13b). In the case that a second transmitter release does occur during bursting, the evoked postsynaptic potentials are attenuated, due to depletion of readily releasable neurotransmitters and previous desensitisation of postsynaptic receptors by the first release of neurotransmitters (Thomson, 2000b; Clements, 1996).

As well as being involved in these mechanisms of short term plasticity, bursting has also been implicated in long term potentiation (LTP) and depression (LTD). In the hippocampus, postsynaptic bursting at temporally relevant intervals has been shown to be sufficient to produce long term synaptic changes (Pike et al., 2004; Froemke et al., 2006; Thomas et al., 1998). Burst-induced LTP and LTD has also been observed in other brain areas (Birtoli and Ulrich, 2004).

Izhikevich et al. (2003) proposed an additional role for bursting in terms of information transfer between neurons. Short-term synaptic plasticity has been shown to play a role in the temporal filtering of spikes, as the effectiveness of high frequency presynaptic firing at evoking postsynaptic responses is enhanced by facilitation, while depression filters out high-frequency firing (Fortune and Rose, 2001; Stevens and Wang, 1995). Synapses capable of exhibiting both depression and facilitation can thus act as a filter, for which only presynaptic firing in a particular window of resonant frequencies produces large EPSPs (Izhikevich et al., 2003).

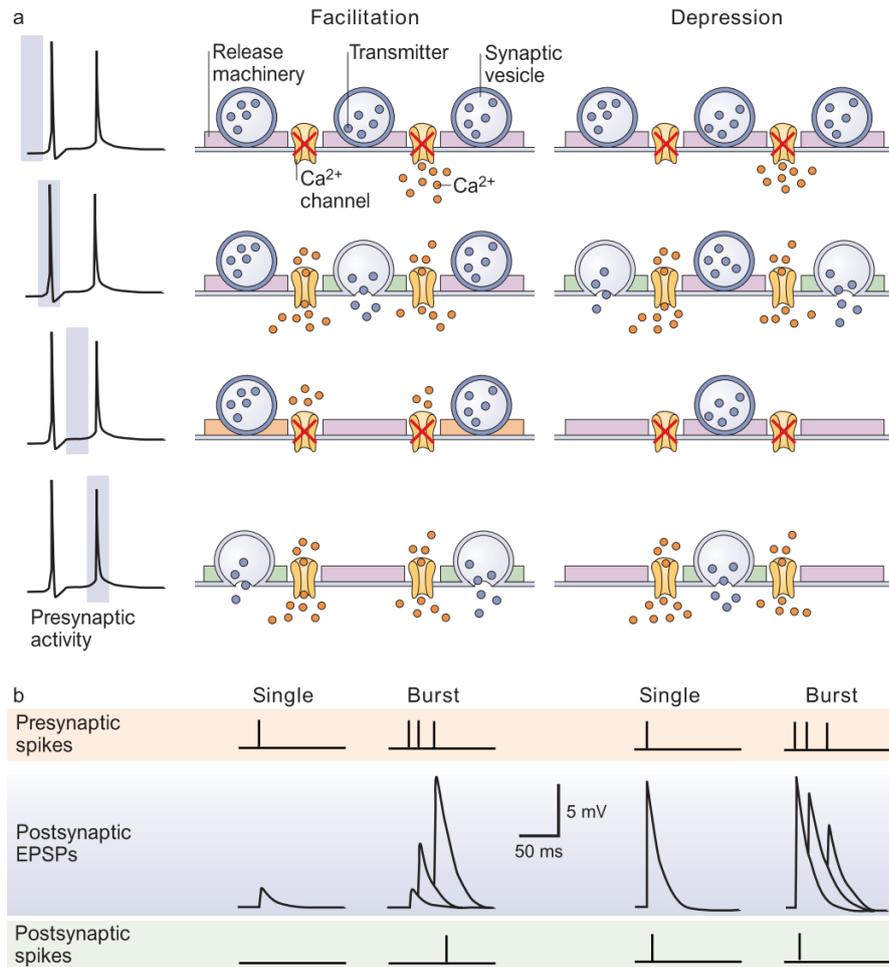


Figure 1.13: Facilitation and depression at synapses plays a role in the filtering of bursts. **(a)** Activity at the presynaptic terminal during the arrival of two action potentials (shown in leftmost column, with the blue bar representing the corresponding time period). Before the arrival of an action potential, ion channels are closed ('Facilitation', top row). Arrival of the action potential results in the opening of these channels, and subsequent Ca^{2+} influx. At facilitating synapses, most neurotransmitter-containing vesicles are not released at this point ('Facilitation', 2nd row), however the intracellular Ca^{2+} from the first action potential can prime the release machinery ('Facilitation', 3rd row) and increase the probability of neurotransmitter release with successive spikes ('Facilitation', 4th row). At synapses with a high release probability, single spikes can cause the release of many vesicles ('Depression', 2nd row). If neurotransmitters are not replenished before arrival of the next spike ('Depression', 3rd row), this spike will induce lower levels of neurotransmitter release ('Depression', 4th row). **(b)** During facilitation (right), EPSP amplitudes increase during a presynaptic bursts until they trigger a postsynaptic action potential. At depressing synapses (left), single spikes can produce postsynaptic action potentials, and EPSP amplitudes tend to decrease over a presynaptic burst. Figure reproduced from Krahe and Gabbiani (2004).

Since the value of these resonant frequencies can differ between the various synapses of the same presynaptic cell, a cell can selectively activate certain postsynaptic cells by producing bursts with specific intra-burst frequencies (Gupta et al., 2000; Izhikevich et al., 2003; Markram et al., 1998). Additionally, at the neuronal level, certain cell types exhibit resonant properties due to submembrane threshold oscillations in membrane potential (Lampl and Yarom, 1997; Leung and Yu, 1998). Bursts at the natural resonant frequency of a cell have the greatest probability of evoking an action potential (Hutcheon and Yosef, 2000). This again allows a presynaptic neuron to selectively communicate with certain postsynaptic cells by producing bursts with frequencies that are resonant only for specific cells (Izhikevich et al., 2003).

It has also been suggested that bursts of spikes transmit information with a higher signal-to-noise ratio than single spikes (Sherman, 2001). Evidence of this has been seen in a variety of brain regions, such as the hippocampus, where place-fields have been shown to be more accurately defined by bursts rather than single spikes (Otto et al., 1991). Bursting has also been shown to produce sharper sensory tuning curves (Cattaneo et al., 1981; Krahe and Gabbiani, 2004) and more reliable feature extraction than single bursts (Gabbiani et al., 1996; Sherman, 2001; Krahe et al., 2002). Livingstone et al. (1996) demonstrated that spikes recorded from the primary visual cortex of freely viewing monkeys weakly correlated with input stimulus, however, when only bursts were considered, the stimulus could be more clearly discerned in the response map of the cells (Figure 1.14).

Some experimental evidence has also suggested that particular features of bursts, such as their duration or temporal pattern of spikes may also transmit specific information about particular features of stimuli (Martinez-Conde et al., 2002; Furukawa and Middlebrooks, 2002; Nelken and Chechik, 2005; Mease et al., 2017), however, more research in this area is required to deduce unequivocally which, if any, parameters of bursting activity in different neuronal networks carry additional information.

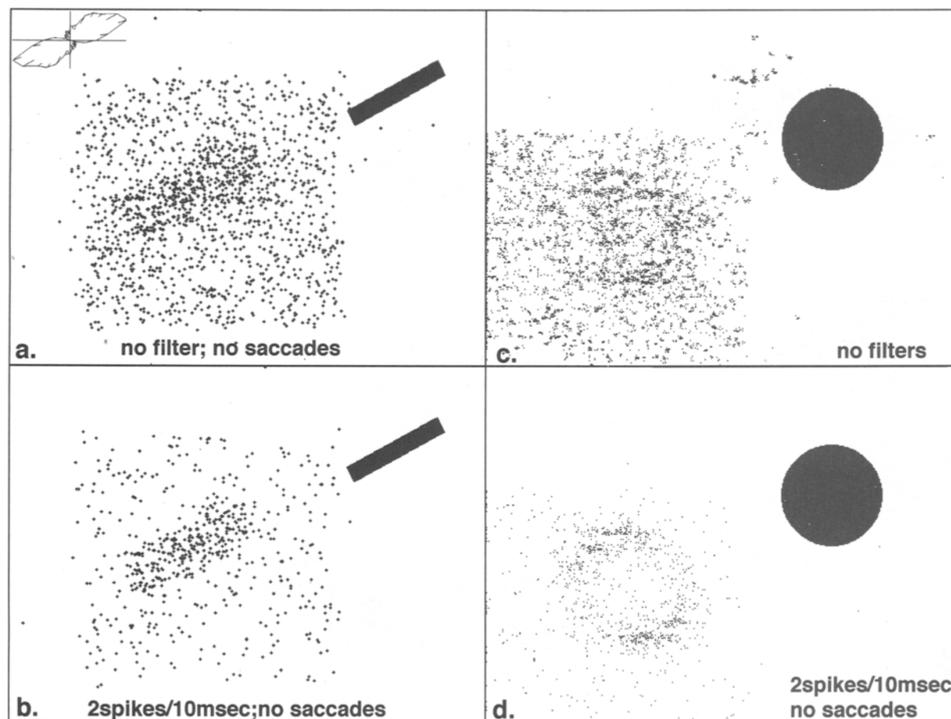


Figure 1.14: Spikes recorded from cells in the monkey visual cortex. (**a**, **b**) Response of cells to a flashing bar stimulus, shown in the top right of each panel. (**a**) All spikes, and (**b**) spikes filtered to show only bursts. (**c**, **d**) Spikes recorded in response to a white disc stimulus with (**a**) no filtering and (**b**) burst filtering. Spikes in all panels are mapped according to the position of the receptive field of the neuron from which the spike was produced. The input stimuli can be more clearly discerned when only bursts are considered (**b**, **d**). Figure reproduced from Livingstone et al. (1996).

1.7.2 The role of bursting in *in vivo* behaviours

The importance of neuronal bursting has also been demonstrated through its association with a variety of behaviours *in vivo*, including visual processing and reward and goal directed behaviour (Cattaneo et al., 1981; Krahe and Gabbiani, 2004; Tobler et al., 2003; Schultz et al., 1997; Schultz, 1998). Bursting activity has long been associated with sleep and resting conditions (Evarts, 1964; Barrionuevo et al., 1981; McCarley et al., 1983; Weyand et al., 2001; Steriade et al., 2001). During slow-wave sleep in rats, sequences of spikes that were observed during awake behaviour have been found to be replayed in a temporally compressed manner during sharp-wave bursts in the hippocampus, suggesting that bursting activity is involved in memory consolidation (Nádasy et al., 1999; Ji and Wilson, 2007).

Bursting of hippocampal place cells has also been observed during exploration of new environments (O’Keefe and Recce, 1993; Epsztein et al., 2011). The presence of bursting in these, as well as other memory-related behaviours (Burgos-Robles et al., 2007; Xu et al., 2012), suggests that bursting plays a specific role in memory and learning in the adult brain (Paulsen and Sejnowski, 2000).

Bursting is also believed to be involved in changes in attention or brain states. In the thalamus, relay neurons display two modes of firing, tonic firing and bursting (Sherman, 2001). Bursting has been shown to be more prevalent in states of sleep or drowsiness, however, irregular bursts have also been observed in awake states (Ramcharan et al., 2000a; Weyand et al., 2001). Evidence has suggested that these thalamic bursts are involved in attentional changes, acting as an alerting signal or ‘wake up call’ to the cortex in response to novel stimuli (Sherman, 2001; Swadlow and Gusev, 2001; Hu and Agmon, 2016; Sherman, 2017). Recently, Li et al. (2009) also showed that bursting of a single cortical neuron could induce a switch between states resembling slow-wave and rapid-eye-movement sleep.

Bursting activity has also been seen to be altered in certain pathological conditions (Walker et al., 2008; Jackson et al., 2004; Miller et al., 2011; Singh et al., 2016). For example, increased bursting activity has been observed in the basal ganglia of Parkinson’s patients, with correlations between the level of bursting activity and the progression of the disease (Lobb, 2014; Ni et al., 2001). This suggests that the study of bursting activity could not only reveal important features of normal brain function, but also how this is altered in diseased states.

1.7.3 Cellular mechanisms underlying bursting

A large variety of physiological mechanisms have been observed to play a role in the generation and maintenance bursting activity in different neuronal networks, with the precise mechanisms underlying bursting across all cell types yet to be fully understood (Lisman, 1997; Jasinski et al., 2013). Despite this diversity, a commonality across cell types is that bursting is generally a result of the interaction of fast membrane conductances responsible for spiking activity and slower subsystems which modulate the spiking activity (Krahe and Gabbiani, 2004). The mechanisms of the slow subsystem vary between different cell types, but generally consist of voltage or Ca^{2+} -mediated activation/inactivation of an outward/inward current that is required to sustain spiking activity (Izhikevich, 2007).

In thalamic relay neurons, the cessation of bursting relies on voltage-gated inactivation of an inward Ca^{2+} current (I_T). When a strong depolarization proceeds a period of hyperpolarization in these cells, I_T is activated, producing a Ca^{2+} spike which generally results in a burst of action potentials. Repetitive spiking during the burst inactivates I_T , and K^+ influx repolarizes the cell (Jahnsen and Llinas, 1984; Sherman, 2001; McCormick, 1991; Ramcharan et al., 2000b).

For CA1 hippocampal neurons, the precise cellular mechanisms of bursting are not well understood. Much of the research in this area has involved investigating the mechanisms underlying the afterdepolarizing potential (ADP) that has been observed in these cells. This is a positive deflection of membrane potential following an action potential that can trigger burst firing (Yue et al., 2005). Numerous studies have suggested a contribution of M-type K^+ channels, persistent Na^+ and transient Ca^{2+} channels to this activity (Metz et al., 2005; Azouz et al., 1996; Jensen et al., 1994; Shuai et al., 2003; Yue et al., 2005; Yue and Yaari, 2004; Chen and Yaari, 2008).

Other mechanisms have also been implicated in burst generation and moderation in different neuronal systems, including alternative activation of persistent Na^+ currents and slow repolarizing K^+ , and ‘ping-pong’ mechanisms caused by interaction of fast Na^+ and K^+ conductances in the soma and slower Ca^{2+} and K^+ dendritic currents (D’Angelo et al., 2001; Lemon and Turner, 2000; Mainen and Sejnowski, 1996; Wang, 1999).

1.8 Thesis aims and overview

As evidenced in the preceding sections, despite great progress in our understanding of the development, structure and function of the nervous system and the neurons from which it is composed, there remain many aspects that are not well understood. In particular, the precise role and mechanisms of the changing patterns of spontaneous activity in neuronal networks over development, and which aspects of these patterns, in particular bursting behaviour, are essential for the proper development of functional neuronal circuits remains unclear. Neuronal cell types are also far from being well-classified, with a lack of agreement regarding how to combine the multiple characterisations of neurons (electrophysiological, morphological, molecular, etc.) into a consensus classification.

This thesis aims to address these issues by identifying, evaluating and developing computational techniques for the analysis of recordings of spontaneous neuronal network activity, as well as experimental characterisations of single neuronal cells. Throughout this thesis, I critically assess existing computational methods for analysing experimental data from these areas and propose improvements to these techniques and new applications of analysis methods adopted from other fields. Quantitative analysis frameworks are also developed and used to extract insights from novel experimental data sources. An overarching goal of this work is to provide insight on the most effective methods for collecting and analysing experimental data from these fields, as well as the limitations of these techniques and possible improvements.

In Chapter 2, I investigate computational techniques for the identification of bursting activity in recordings of neuronal networks from *in vitro* MEAs. I first use simulated data to evaluate the performance of numerous existing burst detection techniques in relation to a list of desirable properties that an ideal burst detector should possess. Experimental data is also used to further evaluate these methods under realistic conditions. The high performing methods from this evaluation are then used to analyse bursting activity in novel recordings of human induced pluripotent stem cell-derived neuronal networks. Following this, I propose adaptations to several existing burst detection techniques, and show that these alterations can improve their performance at identifying bursts under certain conditions.

In Chapter 3, I develop a broader quantitative framework for the analysis of spontaneous activity patterns in developing neuronal networks and use this framework to characterise the developmental profiles of neuronal networks recorded in two experimental contexts. Firstly, I show that by using a number of features to quantify spontaneous activity patterns, significant differences between the activity of developing neuronal networks produced from hippocampal and cortical cells can be observed. These features allow for the activity of hippocampal and cortical networks to be differentiated at a variety of ages with reasonable levels of reliability using standard classification techniques. Following this, I show that a similar framework of quantitative features can be used to characterise and differentiate between the activity of neuronal networks at different developmental ages using recordings from multi-well MEA devices. I further analyse the major sources of variability in these recordings and investigate the most appropriate approaches for using these types of arrays to study multiple experimental conditions in a high-throughput manner.

In Chapter 4, I examine the potential benefits of using multiple characterisations of neurons for the classification of neuronal cell types. In particular, I investigate the impact of utilising multi-view clustering techniques that have been developed for finding the shared cluster structure in data represented by two or more feature sets or ‘views’, in neuronal classification studies. I perform a thorough literature search to identify a range of multi-view clustering methods, and apply these techniques to the analysis of two multi-view data sets consisting of the electrophysiological and morphological properties of neurons. I show that using both the electrophysiological and morphological features of neurons can be advantageous for their classification, compared to the use of a single feature type. I investigate which approaches to multi-view clustering are the most effective for classifying these types of data sets and what can be learnt about the underlying cell types present using these methods.

In Chapter 5, I summarise and discuss the key results from this thesis and suggest directions for future work.

Chapter 2

Burst analysis

The majority of the material in this chapter has been previously published in the article ‘A comparison of computational methods for detecting bursts in neuronal spike trains and their application to human stem cell-derived neuronal networks’ (Cotterill et al., 2016b). Some text has been quoted verbatim from this paper, other sections have been expanded from the article, and Section 2.7 is additional research which was not included in the original paper. The burst detection methods outlined and analysed in Sections 2.4 and 2.5 are the work of others. The experimental data analysed in Sections 2.5.2 and 2.5.3 were provided by Stephen Eglen and Paul Charlesworth, respectively. All analysis of burst detection methods and experimental data presented in this chapter is my own work.

2.1 Summary

Features of neuronal bursts, defined as short periods of high-frequency firing of neurons, are frequently used to characterise activity patterns in a variety of neuronal networks, and identify variations in network state. Despite its importance for classifying network activity, no one technique for identifying the bursting activity of neurons has been widely adopted in the field. Instead, a large variety of burst detection methods have been developed, many of which have only been evaluated by their original authors and under a narrow range of specific experimental conditions.

To provide a more robust evaluation of existing burst detection techniques, in this chapter I conduct an unbiased assessment of eight previously developed burst detectors. I first suggest a list of properties that an ideal burst detection method should

possess, and evaluate each burst detector against these properties using simulated spike trains. Further assessment of the methods is then achieved by examining the coherence of the bursts detected by each method with visually annotated bursts in experimental recordings of mouse retinal ganglion cells. Based on these assessments, I find that several burst detection techniques perform poorly at accurately classifying bursting activity in spike trains with a range of properties.

The four most promising burst detection methods from this initial analysis are then used to detect bursting activity in novel recordings of networks of human induced pluripotent stem cell-derived neurons, provided by Paul Charlesworth. I find that no current method can provide ‘perfect’ burst detection, however two techniques in this study outperform compared to the others. Based on these results, I provide recommendations regarding the optimal approach for robust burst analysis of experimental data using existing techniques. Additionally, in an effort to improve their burst detection performance, I adapt two of the burst detection methods in this study, and examine the impact of the modifications on their ability to accurately identify bursting activity.

2.2 Introduction

Bursting activity, defined as brief episodes of rapid spiking of a neuron separated by periods of relative quiescence (see Figure 2.1), has been observed extensively in both *in vitro* and *in vivo* neuronal networks (Weyand et al., 2001; Chiappalone et al., 2005; Pasquale et al., 2010).



Figure 2.1: Example of bursting activity in a spike train recorded from mouse retinal ganglion cells. Horizontal blue lines show the location of bursts. Scale bar represents 1 s.

Bursting activity is believed to play a role in a range of physiological processes, including synapse formation (Maeda et al., 1995) and long-term potentiation (Lisman, 1997) (see Section 1.7 for details). Analysis of patterns of bursting activity can thus be used as a proxy for studying the underlying physiological processes and structural features of neuronal networks. A common method of studying bursting activity *in*

vitro involves the use of MEA recordings of spontaneous or evoked neuronal network activity. This approach has been employed to study changes in spontaneous network activity over development (Wagenaar et al., 2006a), and the effect of pharmacological or genetic manipulations (Eisenman et al., 2015; Charlesworth et al., 2016).

Despite the prevalence of bursting as a feature used to analyse neuronal network activity, there remains a lack of agreement in the field about the definitive formal definition of a burst (Cocatre-Zilgien and Delcomyn, 1992; Gourévitch and Eggermont, 2007). There is also no single technique that has been widely adopted in the field for identifying the location of bursts in spike trains. Instead, a large variety of burst detection methods have been proposed, many of which have been developed and assessed using specific data sets and single experimental conditions. As the majority of existing studies of bursting activity have been performed on experimental data from recordings of rodent or feline neuronal networks (Charlesworth et al., 2015; Mazzoni et al., 2007), this type of data has most often been used to assess the performance of burst detection techniques (Chiappalone et al., 2005; Mazzoni et al., 2007; Gourévitch and Eggermont, 2007).

Recently, it has been shown that networks of neurons derived from human stem cells can be grown successfully on MEAs and exhibit spontaneous electrical activity, including bursting (Illes et al., 2007; Heikkilä et al., 2009). Human stem-cell derived neuronal cultures have also been demonstrated to be a suitable alternative to rodent neuronal networks in applications such as neurotoxicity testing (Ylä-Outinen et al., 2010). This has led to a demand for a robust method of analysing bursting in these networks, which commonly exhibit more variable and complex patterns of bursting activity than rodent neuronal networks (Kapucu et al., 2012) (see Figure 2.2).

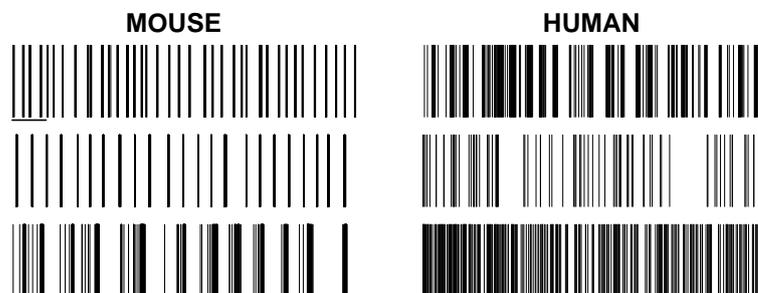


Figure 2.2: Examples of spike trains from mouse and human neuronal networks. Each row represents the spikes recorded from one electrode and the scale bar represents 30 s. Recordings from human neuronal networks often exhibit more variable and complex spontaneous activity patterns.

2.3 Previous approaches to burst detection

Since the development of the first methods to identify bursting in neuronal networks more than three decades ago, a vast number of burst detection techniques have been proposed. These methods take a variety of approaches, the simplest of which impose thresholds on values such as the minimum firing rate or maximum allowed interspike interval (ISI) within a burst, classifying any sequence of consecutive spikes satisfying these thresholds as a burst. In well-ordered spike trains, these thresholds can be set as fixed values by visual inspection (Weyand et al., 2001; Chiappalone et al., 2005). Other methods also incorporate additional thresholds on relevant parameters such as the minimum interval between two bursts and the minimum duration of a burst, to restrict detected bursts to those with biologically realistic properties (Nex Technologies, 2014).

As opposed to having fixed threshold parameters that are chosen by the user, other burst detection algorithms derive the values of their threshold parameters adaptively from properties of the data, such as the mean ISI (Chen et al., 2009a) or total spiking rate (Pimashkin et al., 2011). Commonly, this involves the use of some form of the distribution of ISIs on a spike train. For spike trains containing bursting activity, the smoothed histogram of ISIs on the train should have a peak in the region of short ISIs, which represents within-burst ISIs, and one or more peaks at higher ISI values, representing intra-burst intervals. A threshold for the maximum ISI allowed within a burst can be set at the ISI value representing the turning point in the histogram (Cocatre-Zilgien and Delcomyn, 1992).

Several other adaptive burst detection algorithms also use distributions related to the ISI histogram. Selinger et al. (2007) and Pasquale et al. (2010) argue that the histogram of $\log(\text{ISI})$ s provides a better separation of within and between-burst intervals, and use this histogram to set the threshold for the maximum within-burst ISI at the minimum between the first two well separated peaks. Kaneoke and Vitek (1996) use the histogram of discharge density rather than ISIs for burst detection, while Bakkum et al. (2013) employ the ISI histogram between every n th spike in a network, and Kapucu et al. (2012) derive the threshold parameters for detecting bursts in their algorithm from the cumulative moving average of the ISI histogram.

Another category of burst detection techniques are the surprise-based methods, which use statistical techniques to distinguish periods of bursting from baseline neuronal firing. The earliest of such methods was developed by Legéndy and Salcman (1985), and detects bursts as periods of deviation from an assumed underlying Poisson process of neuronal firing. This method is limited by its assumption of Poisson-distributed spike trains, which has been shown to be inappropriate for many common spike trains, in particular because of the refractory period between spikes (Câteau and Reyes, 2006) (see Section 1.11). Despite this, the Poisson Surprise method has been one of the most widely used burst detection methods since its development over thirty years ago (348 citations as at May 2017) and is still commonly used for analysing bursting activity in experimental studies of numerous neuronal network types (Singh et al., 2016; Pluta et al., 2015; Senn et al., 2014). More recently, other surprise-based burst detectors have been developed that replace the assumption that baseline firing follows a Poisson process with other assumptions about the underlying distribution of spikes (Ko et al., 2012; Gourévitch and Eggermont, 2007).

Other burst detectors take alternative approaches to separate bursting from background spiking activity. Turnbull et al. (2005) examine the slope of the plot of spike time against spike number to detect bursts as periods of high instantaneous slope. Martinson et al. (1997) require bursts to be separated by intervals at least two standard deviations greater than their average within-burst ISIs, while Tam (2002) propose a parameter-free burst detection method, in which sequences of spikes are classified as bursts if the sum of their within-bursts ISIs is less than the ISIs immediately before and after the burst.

Numerous studies have also used various forms of Hidden Markov Models to analyse neuronal activity patterns (Radons et al., 1994; Chen et al., 2009b; Abeles et al., 1995). These methods assume that a neuron stochastically alternates between two or more states, characterised by differences in their levels of activity. Tokdar et al. (2010) apply this idea to burst detection by modelling neuronal activity using hidden semi-Markov models.

2.4 Methods

2.4.1 Approach

An approach similar to Cutts and Eglén (2014) was used to perform an unbiased evaluation of eight burst detection techniques using simulated data. As the burst detection techniques examined in this study were the work of others, in this case a blind approach to evaluation was unnecessary, and the identities of the burst detectors were known throughout. In order to make sure the evaluation was as fair as possible, efforts were made to ensure the implementation of each of the burst detectors was as close as possible to the original authors' implementations. Following this, further assessment of each of the burst detectors was performed using experimental data. In these sections, the parameter values for each burst detector were allowed to vary from their original values to examine the sensitivity and specificity of the methods.

2.4.2 Burst detection methods

Eight burst detection techniques which I believe are representative of the major approaches to burst detection and have sufficient general applicability to allow for their use on a wide range of spike trains, were chosen for analysis. This section provides an outline of each of these methods as applied to a single spike train. A common parameter in several of the burst detection algorithm is the maximum ISI allowed within a burst, which I will refer to as *maxISI*.

Each burst detector was implemented using R statistical software (R Development Core Team, 2011), using code that was either provided by the original authors (HSMM, RGS), adapted from code already available in our group (MI, PS, RS) or developed myself (logISI, CMA, IRT) based on the algorithms or alternative code provided by the authors. All of this code has been made publicly available at <https://github.com/ellesec/burstanalysis>.

To minimise the impact of my implementation, the parameters used for each method were generally left set to the values suggested by the original authors. The exception to this was the minimum number of spikes in a burst, which was set to three for all methods. The minimum surprise value for the three surprise-based burst detectors was also set to $-\log(0.01)$, to ensure consistency across the methods. The parameters used for burst detection on simulated data are shown in Table 2.1.

Method	Parameter	Value	
PS	Poisson Surprise (Legéndy and Salcman, 1985)	Minimum surprise value $-\log(0.01) \approx 4.6$	
	MI	MaxInterval (Nex Technologies, 2014)	Maximum beginning ISI 0.17 s Maximum end ISI 0.3 s Minimum interburst interval 0.2 s Minimum burst duration 0.01 s Minimum spikes in a burst 3
CMA	Cumulative Moving Average (Kapucu et al., 2012)	α_1 (α_2) 1.0 (0.5) if skew < 1 0.7 (0.5) if $1 \leq \text{skew} < 4$ 0.5 (0.9) if $4 \leq \text{skew} < 9$ 0.3 (0.1) if $9 \leq \text{skew}$	
	RS	Rank Surprise (Gourévitch and Eggermont, 2007)	Largest allowed ISI in burst 75th percentile of ISIs Minimum surprise value $-\log(0.01) \approx 4.6$
	IRT	ISI Rank Threshold (Hennig et al., 2011)	Rank threshold, θ_R 0.5
		RGS	Robust Gaussian Surprise (Ko et al., 2012)
logISI	LogISI (Pasquale et al., 2010)	Minimum burst surprise $-\log(0.01) \approx 4.6$ Maximum cutoff value 100 ms	
HSMM	Hidden Semi-Markov Model (Tokdar et al., 2010)	Probabilistic cutoff 0.5 Other parameters (N=23) As per paper*	

Table 2.1: The eight burst detectors and the parameter values used for the implementation of each method on synthetic spike trains. Table reproduced from Cotterill et al. (2016b). *These parameters were left set to the default values provided in the ‘burstHSMM’ R package.

Poisson Surprise method (Legéndy and Salcman, 1985)

The average firing rate, λ , on a spike train is calculated, and the underlying activity on this spike train is assumed to follow a Poisson process with rate λ . The Poisson Surprise (PS) statistic for any period of length T containing N spikes is calculated as

$$S = -\log P$$

where

$$P = \exp\left(-\lambda T \sum_{n=N}^{\infty} \frac{(\lambda T)^n}{n!}\right)$$

is the probability that N or more spikes occur randomly in a period of length T .

A surprise maximization algorithm described in Legéndy and Salcman (1985) is then used to find the set of bursts that maximises the PS statistic across the entire spike train. This involves initially identifying bursts as any sequence of three consecutive spikes separated by ISIs which are less than half of the mean ISI on the spike train. Spikes are then added to the end and removed from the beginning of each of these initial bursts until the sequence of spikes with the maximum PS statistic is found. Finally, any bursts which have a PS statistic below a pre-defined threshold level are discarded.

Rank Surprise method (Gourévitch and Eggermont, 2007)

The Rank Surprise (RS) burst detection algorithm is a non-parametric adaptation of the Poisson Surprise approach. To implement this method, all ISIs on a spike train are ranked by size, with the smallest ISI given a rank of one. Let r_n denote the rank of the n th ISI on the spike train. In the absence of any bursting activity, the ISI ranks should be independently and uniformly distributed. For any period containing N spikes separated by $N - 1$ ISIs with ranks r_n, \dots, r_{n+N-1} , the Rank Surprise statistic is defined as

$$RS = -\log(P(D_N \leq r_n + \dots + r_{n+N-1}))$$

where D_N is the discrete uniform sum distribution between 1 and N .

Bursts are then chosen to maximise the RS statistic across the entire spike train using an exhaustive surprise maximisation algorithm, outlined in Gourévitch and Eggermont (2007). A fixed threshold for *maxISI* is first calculated from the distribution of ISIs on the spike train. The first sequence of at least three spikes with ISIs

less than $maxISI$ are found and an exhaustive search of all of the subsequences of ISIs within this period is performed to find the subsequence with the highest RS value. If this value is above a fixed minimum significance threshold, chosen a priori, it is labelled as a burst. This process is repeated on the remaining ISI subsequences within the period of interest until all significant bursts are found. Following this, the next sequence of spikes with ISIs below $maxISI$ is examined in a similar fashion, and this process is continued until the end of the spike train.

Robust Gaussian Surprise method (Ko et al., 2012)

In the Robust Gaussian Surprise (RGS) method, the distribution of $\log(ISI)$ s on each spike train is found and centred around zero. The normalised $\log(ISI)$ s from each spike train in the study are then pooled and the central distribution of this joint data set is found using a procedure outlined in Ko et al. (2012). A burst detection threshold for $maxISI$ is set at the 0.5 percentile of this central distribution, which is estimated as 2.58 times the median absolute deviation of the distribution.

The Gaussian burst surprise value in any interval on a spike train is defined as

$$GS_B = -\log(P)$$

where P is the probability that the sum of normalised $\log(ISI)$ s in the interval is greater than or equal to the sum of an equal number of i.i.d. Gaussian random variables with mean and variance equal to that of the central distribution.

Any consecutive sequence of spikes separated by intervals less than $maxISI$ are classified as burst cores. These burst cores are then extended by adding intervals to the beginning and end of the burst cores until the sequence with the maximum value of GS_B is found. In the case of overlapping bursts, the burst with the largest GS_B value is retained. Finally, any detected bursts with GS_B below a predefined threshold value are discarded. Ko et al. (2012) also propose a similar method for identifying pauses in spike trains, but for this study, these periods of below average firing rate were not examined.

LogISI method (Pasquale et al., 2010)

The histogram of $\log(ISI)$ s on a spike train is computed, using a bin size of 0.1 in $\log(ISI)$ units. Let C_k denote the ISI count in the k th bin of this histogram, which corresponds to an ISI size of ISI_k , and MCV denote a pre-specified threshold value, known as the maximum cutoff value. The location of the peaks of this histogram

are found using a custom peak finding algorithm described in Pasquale et al. (2010). The largest peak of the histogram corresponding to an ISI less than or equal to MCV is set as the intraburst peak, C_{IBP} . If no peak is found in the histogram with $ISI_k \leq MCV$, the spike train is classified as containing no bursts.

In the case that an intraburst peak is present, the minimum value of the histogram between the intraburst peak and each of the following peaks, C_{p_i} ($i = 1, \dots, N$), is found. For each minimum, a void parameter is calculated that represents how well the corresponding peak is separated from the intraburst peak, as

$$void(i) = 1 - \frac{C_{min_i}}{\sqrt{C_{IBP} \cdot C_{p_i}}}$$

where C_{min_i} is the minimum value of C_k for $IBP < k < p_i$.

The smallest ISI_{min_i} for which $void(i) > 0.7$ is set as the threshold for the maximum ISI in a burst, $maxISI$ (see Figure 2.3). Any series of at least three spikes separated by ISIs less than $maxISI$ are classified as bursts. If no point with a void value above 0.7 is found, or if $maxISI > MCV$, bursts are detected using MCV as the threshold for the maximum ISI in a burst and then extended to include spikes within $maxISI$ of the beginning or end of each of these bursts.

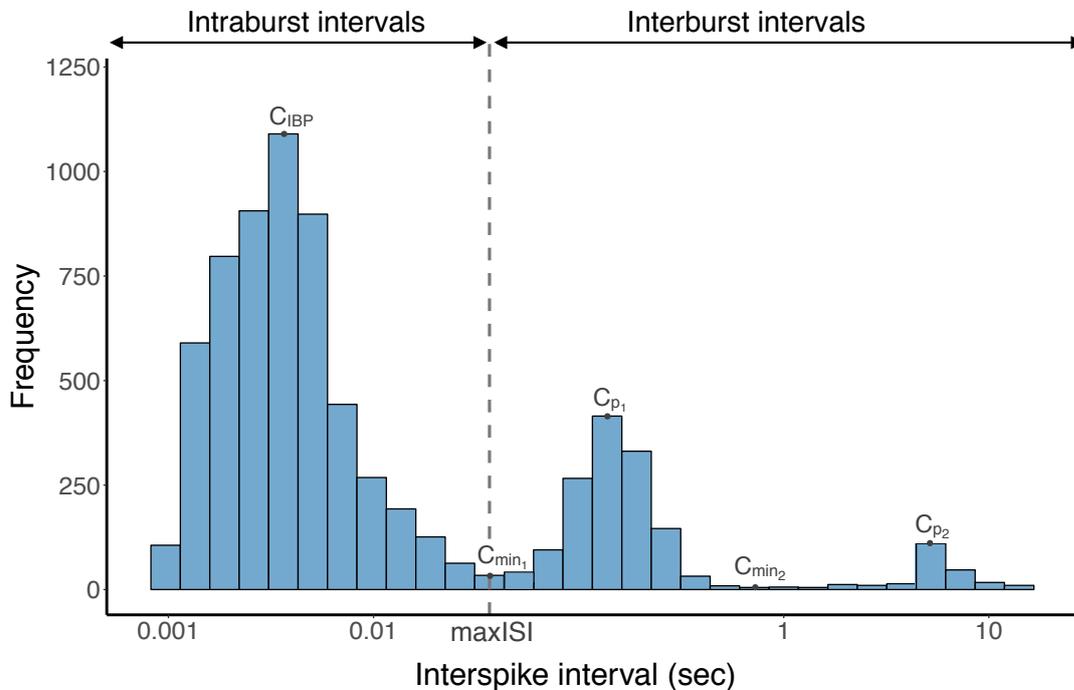


Figure 2.3: Example of log-adjusted ISI histogram with the threshold for intraburst and interburst intervals found using the logISI method.

Cumulative Moving Average (CMA) method (Kapucu et al., 2012)

This method also uses the histogram of ISIs on a spike train. The cumulative moving average (CMA) at each ISI bin of the histogram is calculated. The CMA of the N th ISI bin is defined as

$$CMA_N = \frac{1}{N} \sum_{k=1}^N C_k,$$

where C_k is the ISI count in the k th bin. The skewness of the CMA distribution is used to determine the values of two threshold parameters, α_1 and α_2 , based on the scale given in Kapucu et al. (2012). The maximum of the CMA distribution, CMA_{max} , is found and the value of $maxISI$ is set at the ISI bin at which the CMA is closest in value to $\alpha_1 \cdot CMA_{max}$ (see Figure 2.4). Burst cores are then found as any sequences of at least three spikes separated by ISIs less than $maxISI$.

Kapucu et al. (2012) suggest extending these burst cores to include burst-related spikes. These are found using a second cutoff, set at the value of the ISI bin at which the CMA is closest to $\alpha_2 \cdot CMA_{max}$. Spikes within this cutoff distance from the beginning or end of the existing burst cores are classified as burst-related spikes. For this study, only the burst cores detected by this method were examined, omitting any burst related spikes.

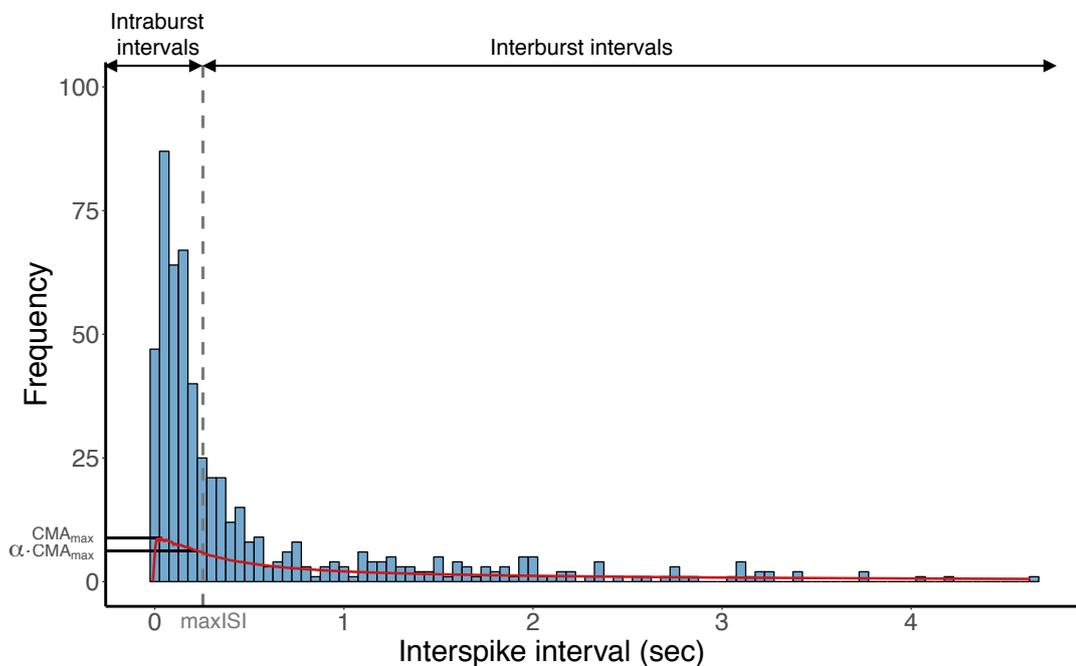


Figure 2.4: Example of ISI histogram with the threshold for intraburst and interburst intervals found using the CMA method. Red line shows the cumulative moving average of the ISI histogram.

ISI rank threshold method (Hennig et al., 2011)

In the ISI rank threshold (IRT) method, the rank of each ISI on a spike train relative to the largest ISI on the train is calculated, with $R(t)$ denoting the rank of the ISI beginning at time t . The probability distribution, $P(C)$, of spike counts in one-second time bins over the spike train is also found. A rank threshold, θ_R , is set to a fixed value, and a spike count threshold, θ_C , is calculated from $P(C)$. A burst is then defined to begin at a spike at time t if the rank of the preceding ISI satisfies $R(t) < \theta_R$ and the spike count in the following second, $C(t, t + 1)$, exceeds θ_C . The burst continues until a spike is found for which $C(t, t + 1) < \frac{\theta_C}{2}$.

Hidden Semi-Markov Model method (Tokdar et al., 2010)

This method is based on the assumption that neurons switch stochastically between two states: ‘non-bursting’ (state 0) and ‘bursting’ (state 1), that can be modelled using a Hidden semi-Markov model. The transition times between the two states are modelled using two gamma distributions, f_0^{ITI} and f_1^{ITI} . Within each of the states, the ISI times are modelled using two additional gamma distributions, f_0^{ISI} and f_1^{ISI} . The parameters of these four distributions are learned from the data. A custom Markov chain Monte-Carlo (MCMC) algorithm described in Tokdar et al. (2010) is then used to compute the posterior probability that a neuron is in a bursting state at any given time. A fixed threshold value is chosen a priori, and any periods during which the posterior probability exceeds this value are classified as bursts.

MaxInterval method (Nex Technologies, 2014)

Bursts are defined using five fixed threshold parameters, shown in Figure 2.5. The value of these parameters are chosen a priori and any series of spikes that satisfy these thresholds is classified as a burst.

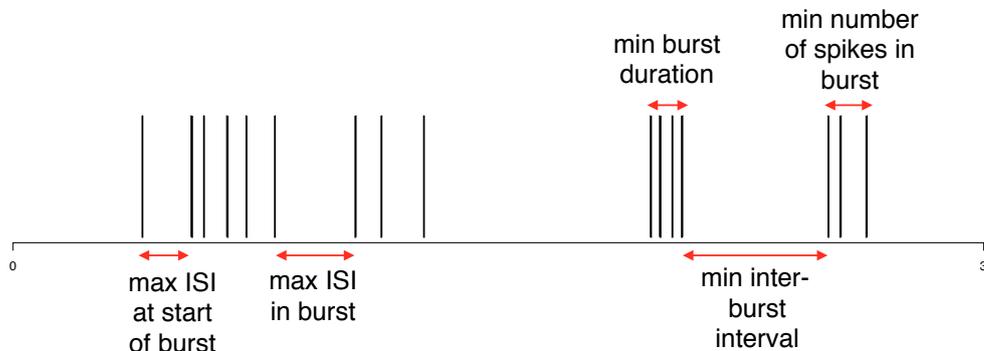


Figure 2.5: Illustration of the parameters used by the MaxInterval method.

Most burst detectors described here separate data sets consisting of multiple spike trains, such as those recorded from the multiple electrodes of an MEA, into individual spike trains, and calculate any associated parameter values, such as $maxISI$, and the location of bursts on each spike train separately. The MaxInterval method is an exception to this, as in the standard implementation of this method, the same values for the threshold parameters are used to detect bursts across all electrodes on an array to simplify the analysis. The RGS method also analyses multiple spike trains from the same array simultaneously by combining the ISIs from all electrodes on an array and using this pooled data to calculate the central distribution of the data and find the cutoff value used to detect the initial burst cores on all spike trains in the study.

2.4.3 Analysis of synthetic data

As an initial assessment of the burst detection methods, the performance of each burst detector was evaluated against a list of eleven properties that I deemed desirable in a burst detection method, shown in Table 2.2. For the desirable properties D1–D3, assessment was made based on the details of the burst detector’s implementation, while testing on simulated data was performed to assess the remaining properties (D4–D11). I used simulated data to generate spike trains with specific properties of interest. The output of each burst detector could then be compared with the ‘ground truth’ bursting activity in each of the simulated spike trains. The simulated spike trains used in this study were generated using the following models, with the parameter values specified in Table 2.3.

Poisson and Gamma distributions

Two different types of non-bursting spike trains were simulated, with ISIs that were distributed according to a homogeneous Poisson process, or gamma distributed ISIs, respectively. For both types of spikes trains, the smallest 10th percentile of ISIs were removed by omitting the corresponding spikes, to eliminate any burst-like activity that randomly arose in the simulated spike trains.

Inhomogeneous Poisson distribution

Spike trains with time-varying firing rates and no bursting activity were simulated using a Poisson process with non-homogeneous intensity, $\lambda(t)$. To eliminate any possible bursting activity in each spike train, the spikes corresponding to the smallest 10th percentiles of ISIs were removed.

Poisson bursting

Spike trains containing bursts were simulated using the Poisson bursting model. The time of the centre of each burst was modelled as a Poisson process with a fixed intensity, λ . The number of spikes in each burst was drawn randomly from a Poisson distribution with mean n . The position of the spikes in each burst relative to the centre of the burst were then drawn from a uniform distribution on $[-\frac{r}{2}, \frac{r}{2}]$. In the case that two simulated bursts overlapped, only the first was kept.

To simulate spike trains with non-stationary bursts, the values of n and r were drawn randomly from uniform distributions for each burst, rather than being held at fixed values. Only the simulated bursts with within-burst firing rates above 5 Hz were included in these spike trains.

To simulate noisy spike trains, between-burst spikes were modelled with gamma-distributed ISIs with the smallest 10th percentile of ISIs removed. These noise spikes were added to the Poisson bursting spike trains, and to avoid any overlap between burst and noise spikes, any noise spikes which overlaid or were within 0.5 s of the beginning or end of any burst were removed.

One hundred spike trains of 300 s duration were simulated for each desirable property and analysed using each of the burst detectors in this study. Examples of simulated spike trains used for evaluating properties D5–11 are shown in Figure 2.6. A comparison of the ‘ground truth’ bursting activity in each spike train and the results from each of the methods was then performed.

For spike trains containing both bursts and noise spikes, this involved calculating the:

1. Fraction of true positive spikes, defined as the proportion of within-burst spikes correctly identified by the burst detector as being within bursts
2. Fraction of false positive spikes, defined as the fraction of all noise spikes which were incorrectly identified as being a part of bursts.

Desirable properties	
D1	Deterministic: The method should detect the same bursts over repeated runs on the same data, to ensure consistency and reproducibility of results
D2	No assumption of spike train distribution: The method should not assume ISIs follow a standard statistical distribution, to ensure wide applicability to a variety of spike trains
D3	Number of parameters: The method should have few parameters, to reduce the variability inherently introduced through parameter choice
D4	Computational time: The method should run in a reasonable amount of time using standard personal computers
D5	Non-bursting trains: The method should detect few spikes as being within bursts in spike trains containing no obvious bursting behaviour
D6	Non-stationary trains: The method should detect few spikes as being within bursts in spike trains with non-stationary firing rates that contain no obvious bursting behaviour
D7	Regular short bursts: The method should detect a high proportion of spikes in bursts in spike trains containing short well-separated bursts
D8	Non-stationary bursts: The method should detect a high proportion of spikes in bursts in spike trains containing bursts with variable durations and numbers of spikes per burst
D8	Regular long bursts: The method should detect a high proportion of spikes in bursts and accurate number of bursts in spike trains containing long bursts with low within-burst firing rates
D10	High frequency bursts: The method should detect a high proportion of spikes in bursts and accurate number of bursts in spike trains containing a large number of short bursts
D11	Noisy train: The method should classify a high number of within-burst spikes as bursting and a low number of interburst spikes as bursting in spike trains containing both bursts and noise spikes

Table 2.2: Desirable properties for a burst detector. Table reproduced from Cotterill et al. (2016b).

Spiking model	Property	Parameters	Mean % spikes in bursts
100 Poisson spiking	Computational time (D4)	$\lambda = 1 \text{ Hz}$	0
50 Poisson spiking	Non-bursting (D5)	$\lambda = 0.5 \text{ Hz}, N = 50$	0
50 Gamma distributed ISIs		$\alpha = 1, \beta = 0.5, N = 50$	
100 Inhomogeneous Poisson	Non-stationary (D6)	$\lambda(t) = 1 + \frac{1}{300}t$	0
100 Poisson bursting	Short bursts (D7)	$\lambda = 0.2 \text{ Hz}, n = 5, r = 0.3 \text{ s}$	100
100 Poisson bursting	Non-stationary bursts (D8)	$\lambda = 0.3 \text{ Hz}, n \sim \mathcal{U}(5, 18), r \sim \mathcal{U}(0.3, 3) \text{ s}$	100
100 Poisson bursting	Long bursts (D9)	$\lambda = 0.1 \text{ Hz}, n = 18, r = 3 \text{ s}$	100
100 Poisson bursting	High frequency (D10)	$\lambda = 1 \text{ Hz}, n = 10, r = 0.5 \text{ s}$	100
100 Poisson bursting with Gamma distributed noise ISIs	Noisy train (D11)	$\lambda = 0.5 \text{ Hz}, n = 8, r = 0.8 \text{ s}$ $\alpha = 1, \beta = 0.5$	91

Table 2.3: Models and parameter values used to generate synthetic spike trains for each desirable property. Each spike train was 300 s duration, and the number, N , of simulated trains was 100, unless otherwise stated. α and β represent the shape and inverse scale parameters of the Gamma distribution, respectively. Table reproduced from Cotterill et al. (2016b).

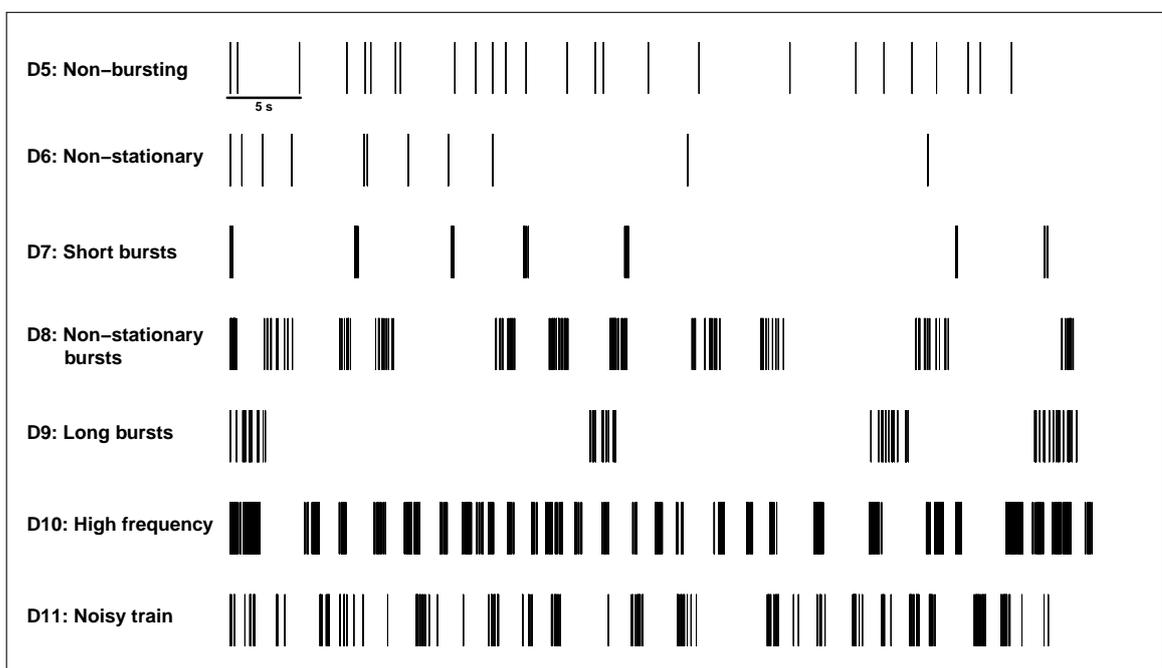


Figure 2.6: Examples of one-minute subsets of simulated spike trains for evaluating desirable features D5–D11. Scale bar represents 5 s. Figure reproduced from Cotterill et al. (2016b).

2.4.4 Analysis of mouse RGC data

MEA recordings of mouse retinal ganglion cells (RGCs) originally analysed in Demas et al. (2003) were re-analysed using the eight burst detectors in this study. These recordings are publicly available at Eglén et al. (2014). Four one hour recordings from mouse retina in the control group at postnatal days (P) 9, 11, 13 and 15 were chosen for re-analysis. In each of these recordings, I manually annotated five randomly chosen spikes trains by visual inspection. Example segments of these annotated spike trains at each age are shown in Figure 2.7.

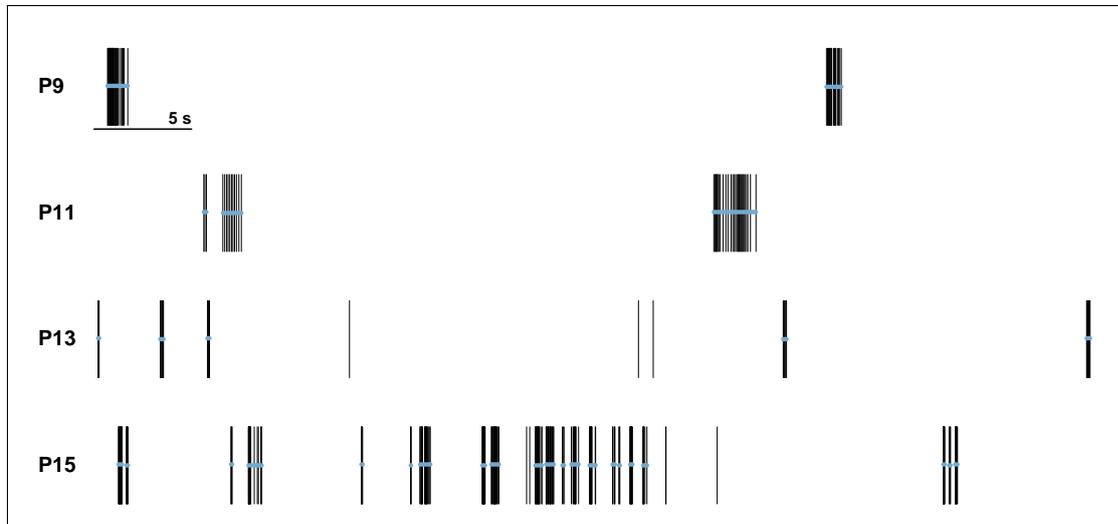


Figure 2.7: Examples of one minute sections of annotated spike trains from retinal ganglion cell recorded at postnatal days 9–15. Blue horizontal lines represent the bursts I annotated by visual inspection. Scale bar represents 5 s. Figure adapted from Cotterill et al. (2016b).

As recordings from MEAs do not reveal the ‘ground truth’ location of bursts in spike trains, these visually annotated bursts were used as a proxy ground truth to assess the performance of each of the burst detectors against. Similar comparison to visually identified bursts has been employed previously to assess the performance of burst detection methods (Chen et al., 2009a; Gourévitch and Eggermont, 2007; Pasquale et al., 2010).

The sensitivity and specificity of each burst detector was then examined by repeatedly analysing the annotated spike trains using each method, while varying the value of a key underlying input parameter. The parameter that was varied in the HSMM and surprise-based methods (PS, RS, RGS) was the probability cutoff. For the IRT method, the value of the spike count cutoff, θ_C , was varied while in

the logISI method, the limit on the maximum allowed ISI cutoff value, MCV, was varied from its original value of 100 ms. For the MI method, the majority of the parameters were kept at the values shown in Table 2.1, with the exception of the maximum beginning and end ISIs, which were varied so that the difference between the maximum beginning and end ISIs was maintained at 0.130 s. The CMA method had no obvious parameters to vary, so for this method only a single value of sensitivity and specificity was found.

Receiver operating characteristic (ROC) curves were generated which plotted 1 – specificity vs. sensitivity for the various input parameter values for each burst detection method. The value of 1 – specificity, or the false positive rate, was taken as the number of spikes falsely detected as being within bursts, as a fraction of the total number of spikes outside of the visually annotated bursts. Sensitivity was defined as the number of spikes correctly identified as being within bursts, as a fraction of the total number of within-burst spikes.

2.4.5 Experimental details for hiPSC-derived neuronal network recordings

These experimental recordings were all performed by Paul Charlesworth in the Paulsen lab of the Department of Physiology, Development and Neuroscience, and do not represent my own work. The following experimental details section (Section 2.4.5) was partially written by Paul Charlesworth and is copied verbatim from Cotterill et al. (2016b).

Neuronal networks were grown from late-stage neuronal precursors differentiated from human induced pluripotent stem cells (hiPSCs) (Axol Bioscience, Cambridge, UK). hiPSCs were generated by reprogramming of embryonic cord blood cells and then differentiated to the neural lineage with protocols based on those in Shi et al. (2012). All of the recordings (447 recordings from 73 MEA platings on 11 plating dates, 4 thawed vials) were obtained with a single line and neural induction (AX0015; <https://www.axolbio.com/page/neural-stem-cells-cerebral-cortex>).

Late-stage neuronal precursors (1×10^6) were thawed and expanded by growing on six-well tissue culture plates coated with polyornithine and laminin (2×10^5 cells/well) in Neural Maintenance Medium (NM), supplemented with $10 \mu\text{M}$ Y-27632 (rho-associated protein kinase inhibitor) for the first 24 h. After 4-5 days cells were

dissociated with Accutase, centrifuged, resuspended in NM, and either plated to MEAs ($2 \times 10^4 - 1 \times 10^5$) or expanded further on six-well plates as above. MEAs (60MEA200/30-Ti; Multi Channel Systems, Reutlingen, Germany) were coated with polylysine followed by laminin as described previously (Charlesworth et al., 2016).

hiPSC-derived neuronal network (hiPSN)-MEA cultures were maintained in NM medium under zero-evaporation lids (Potter and DeMarse, 2001) housed in tissue culture incubators maintained humidified at 37°C and 5% CO₂-10% O₂-85% N₂. Medium was completely exchanged after 24 h to remove Y-27632. Thereafter, MEA cultures were fed by exchanging 40-50% medium with fresh NM three times per week. NM composition was a 1:1 mixture of N2 and B27 supplemented media (N2 medium: DMEM-F-12 + N2 supplement and 5 µg/ml insulin, 1 mM L-glutamine, 100 µM nonessential amino acids; B27 medium: Neurobasal Medium + B27 supplement and 0.5 mM L-glutamine).

Recordings (300 s) of spontaneous extracellular neuronal activity in hiPSN-MEA cultures were made weekly with an MEA system supplied by Multi Channel Systems (MEA 1060INV, with 60MEA200/30-Ti arrays; titanium nitride electrodes, 30 µm diameter, 200 µm spacing, internal reference electrode). The signal was sampled at 25 kHz and stored with a 64-channel data acquisition board (MC Card; Multi Channel Systems) and the acquisition software MCRack (Multi Channel Systems). Action potentials were detected by crossing of a threshold set to a level of 6 standard deviations from the baseline noise level. Record samples (1 ms before and 2 ms after crossing of threshold) confirmed the characteristic action potential waveform. No spike sorting was performed. Action potential timestamps were extracted to text file with batch scripts written for NeuroExplorer (Nex Technologies, Littleton, MA). Recordings made at dates above 16 weeks after plating (WAP) were excluded from the analysis because of the small number of available recordings at these dates, resulting in 424 recordings being analysed.

2.5 Results

2.5.1 Performance on synthetic data

To evaluate each of the burst detectors, they were assessed against each of the eleven desirable properties shown in Table 2.2. An ideal burst detector would possess all of

these desirable properties. For binary properties D1–D4, each burst detector was judged to either possess the property or not. For properties D5–D11, the performance of each burst detector was ranked against all other burst detection methods in the study, based on the median and variance of their performance at analysing one hundred synthetic spike trains, as outlined in Section 2.4.3.

The first desirable property for a burst detection method was that it is deterministic (D1), as this ensures the reproducibility of results and avoids the need to perform repeated trials to find a ‘consensus’ set of bursts from a spike train, which may be computationally expensive. The HSMM method was the only non-deterministic burst detector examined, due to its use of MCMC methods. The bursts detected by this method vary considerably between trials. For example, repeated analysis of a single simulated Poisson bursting spike train with burst frequency of 0.2Hz using the HSMM method found 51 ± 9.75 (mean \pm s.d.) bursts over one hundred trials with the same parameter settings.

The second desirable property for the burst detectors was that they do not assume that the ISIs from a spike train follow a specific statistical distribution (D2). Although ISI distributions in spike trains have often classically been modelled as homogeneous Poisson processes (Legéndy and Salcman, 1985; Gerstein and Mandelbrot, 1964) (see Section 1.11), non-Poisson spiking activity has also been frequently observed (DeWeese et al., 2003; Gur and Snodderly, 2006; Amarasingham, 2006). This can be attributed to various factors, such as the presence of a refractory period or long range temporal correlations between spikes (Averbeck and Lee, 2003). There is thus no consensus about which statistical distribution best represents underlying spiking activity, and imposing an assumption about how ISIs are distributed restricts the applicability of a burst detector to a narrow range of spike trains.

Three burst detectors (PS, RGS, HSMM) assume a fixed statistical distribution for the underlying spike trains. These methods assume ISIs can be modelled using a Poisson process, Gaussian distribution and Gamma distributions respectively. However, for the RGS and PS methods, bursts are detected as deviations from these assumed underlying distributions. These methods thus remain somewhat robust to departures from the distribution assumptions, as ‘surprising’ sequences of spikes from one distribution, will generally have low probability as measured by other distributions commonly used to model spike trains (Legéndy and Salcman, 1985).

A common issue that arises when detecting bursts in sets of spike trains with high variability in their statistical properties, such as those from MEA recordings of hiPSC-derived neuronal networks, is how to choose the optimal parameter values for burst detection. This is further confounded when detecting bursts from recordings spanning a large range of developmental ages or differing experimental conditions. An ideal burst detector should thus have a low number of parameters (D3), to minimise the impact of the choice of parameter values on the resultant detected bursts. The majority of burst detectors examined here required only two or fewer parameters. An exception to this was the MI method, which required five threshold parameters to be set a priori to implement burst detection. The HSMM method required the most ($n = 23$) parameter values to be chosen of all the methods. However, many of these represent initial values which are optimised by the HSMM algorithm. These parameters can thus be left set to the values suggested by the original authors (Tokdar et al., 2010), with little impact on the effectiveness of the method.

Recently, high density MEAs have been developed, which contain up to several thousand electrodes (Berdondini et al., 2009; Eversmann et al., 2003). The increasing use of these arrays for recording of *in vitro* neuronal network activity (Maccione et al., 2014; Lonardoni et al., 2015), as well as the prevalence of multi-well MEAs in applications such as high-throughput neurotoxicity screening (Valdivia et al., 2014; Nicolas et al., 2014) and drug safety testing (Gilchrist et al., 2015), has meant that the computational complexity of burst detectors has become increasingly relevant.

To evaluate the computational time of each burst detection method (D4), each burst detector was used to analyse one hundred simulated non-bursting spike trains of five minutes duration with an average firing rate of 1 Hz. For most methods, the average time to analyse one such spike trains using a standard personal computer was in the order of $10^{-4} - 10^{-1}$ s. The HSMM method was the exception to this, and had an average computational time over 20 times greater than any other burst detection method in this study.

Of equal importance to accurately detecting the presence of bursts in spike trains is a burst detector's ability to identify their absence. Frequently in MEA recordings, particularly those of human neuronal networks, many electrodes record sparse or no bursting activity. An ideal burst detector would find no or few bursts in spike trains containing this type of activity (D5). The majority of burst detectors in this study performed reasonably well at detecting a low amount of bursting activity in spike

trains simulated to contain no bursts (Figure 2.8A). The HSMM and CMA methods, however, greatly overestimated the amount of bursting activity in these spike trains.

When a non-stationary firing rate was incorporated into spike trains containing an absence of bursting activity, the amount of erroneous bursting identified by almost all of the burst detectors increased (Figure 2.8B). The proportion of bursting activity detected by the PS and CMA methods in these spike trains, in particular, significantly increased compared to spike trains with a static mean firing rate. These methods detect bursts as unusual periods of activity relative to baseline spiking activity, and tended to identify bursts during the periods of relatively high firing rate in these non-stationary spike trains.

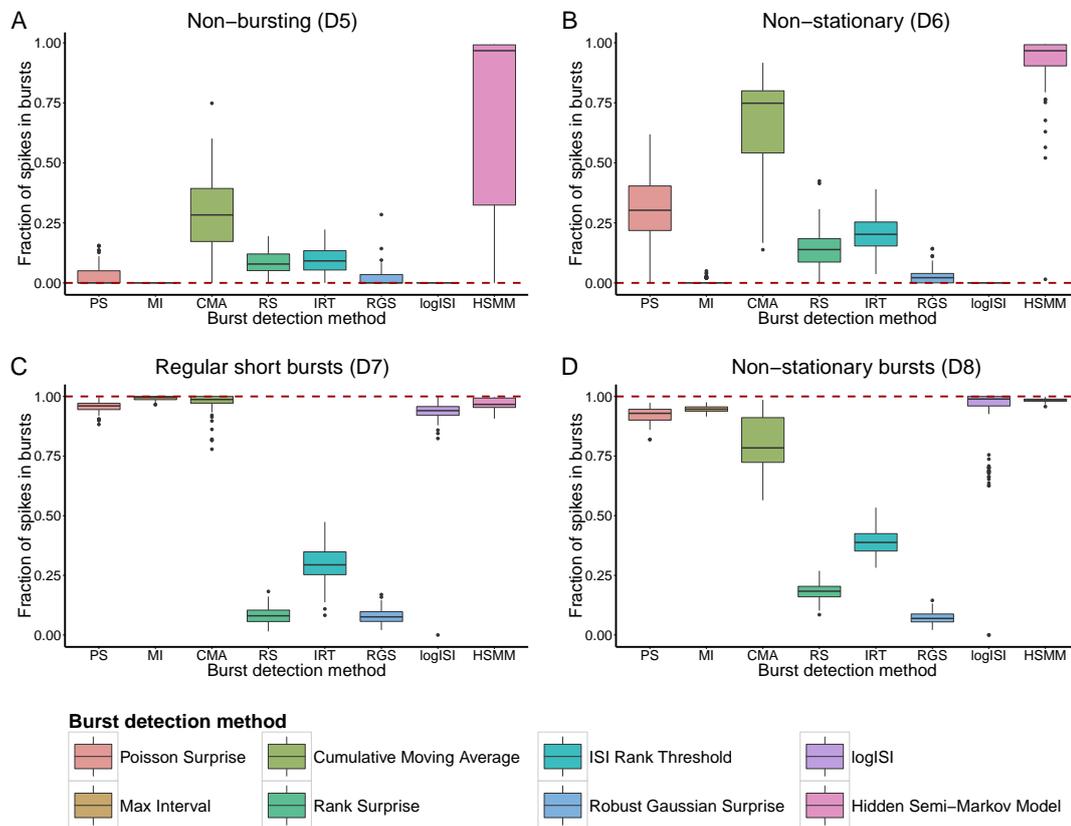


Figure 2.8: Fraction of spikes in bursts found by each burst detector in 100 synthetic trains with **A** no bursting (D5), **B** no bursting and non-stationary firing rate (D6), **C** short regular bursts (D7), **D** bursts with non-stationary burst lengths and durations (D8). Dotted line shows desired result from an ideal burst detector; methods close to this line are deemed to work well. In each ‘box-and-whisker’ plot, boxes show the median \pm inter-quartile range (IQR), and whiskers extend to median $\pm 1.5 \times$ IQR. Outliers are represented as points. Figure reproduced from Cotterill et al. (2016b).

Ideal burst detection methods should also accurately identify bursts in spike trains which consist of only bursting activity, particularly those in which the bursts are regular and well separated (D7). Most burst detectors in the study possessed this property, and could correctly detect over 90% of bursting activity in spike trains simulated to contain only regular bursting (Figure 2.8C). The RS, IRT and RGS methods, however, detected less than half of the bursting activity in these simulated spike trains. A feature of all three of these methods is that they enforce a limit on the maximum proportion of ISIs in a spike train that can be classified as being within bursts. It is thus unsurprising that these methods cannot accurately detect the majority of bursts in spike trains containing no non-bursting activity.

The performance of the burst detectors at analysing spike trains containing irregular bursting activity was also assessed. This included spike trains containing bursts with varying lengths and durations (D8). As these spike trains again only contained bursting activity, the RS, IRT and RGS methods detected only a low proportion of the bursting activity in these synthetic spike trains (Figure 2.8D). Most other methods were able to detect the majority of the bursting activity in these spike trains. However, the proportion of within-burst spikes correctly identified by the CMA method was significantly lower in these spike trains compared to those containing regular bursting activity.

I also assessed the effectiveness of the burst detectors at analysing spike trains containing bursts with long durations and relatively low within-burst firing rates (D9). Of the burst detectors in this study, only the PS and HSMM methods identified both the fraction of spikes within bursts and number of bursts in these spike trains with reasonably high levels of accuracy (Figure 2.9A, B). The MI and CMA methods correctly identified the majority of spikes as being within bursts, but tended to divide the long bursts in these synthetic spike trains into several short bursts. The remaining methods greatly underestimated the proportion of bursting activity in these spike trains.

Another type of non-standard bursting activity observed in recordings of hiPSC-derived neuronal networks is the presence of short, poorly separated bursts occurring at a high frequency. The MI, logISI and HSMM methods were able to correctly detect the majority of within-burst spikes in spike trains containing high-frequency bursting (D10), however all of these methods tended to combine the short bursts in these synthetic spike trains into a smaller number of long bursts (Figure 2.8C, D).

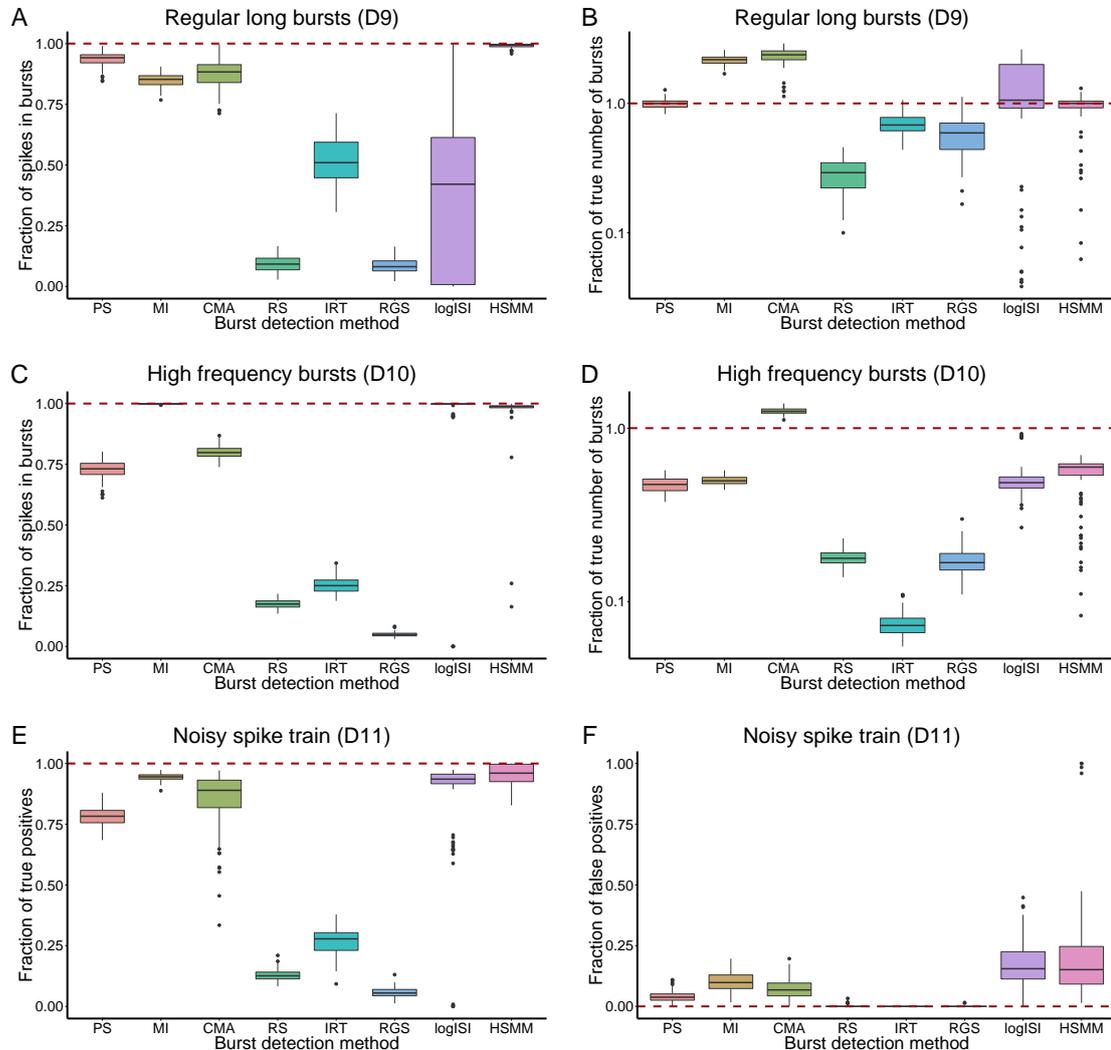


Figure 2.9: Results of each burst detector at analysing 100 synthetic spike trains. **A** Fraction of spikes in bursts, and **B** fraction of true number of bursts in spike trains with regular long bursts (D9); **C** fraction of spikes in bursts, and **D** fraction of true number of bursts in spike trains with high frequency bursting (D10); **E** fraction of true positive, and **F** fraction of false positive spikes in bursts in spike trains containing both bursting and noise (D11). Values calculated as outlined in the methods. Box plots and dotted line as per Figure 2.8 legend. **B** and **D** are presented on a log-scale. Figure reproduced from Cotterill et al. (2016b).

The CMA method was the most accurate of the burst detectors at identifying the correct number of bursts in these spike trains, but tended to underestimate the total proportion of within-burst spikes. The RS, IRT and RGS methods again underestimated the bursting activity in these synthetic spike trains, both in terms of the number of bursts and proportion of spikes in bursts.

Finally, an ideal burst detection method should be able to correctly differentiate between periods of bursting and between-burst spikes in spike trains in which some spiking occurs outside of bursts (D11). To assess this, I compared the bursts detected by each of the methods to the ‘ground truth’ bursting activity in spike trains simulated to contain both bursts and between-burst, or noise, spikes and calculated the fraction of true positive and false positive spikes in bursts found by each method, as defined in Section 2.4.3.

The MI, CMA, logISI and HSMM methods all exhibited reasonably high true positive rates for identifying bursting spikes in the synthetic spike trains, however, of these methods, the logISI and HSMM methods classified a higher number of noise spikes as being within bursts (Figure 2.9E, F). The RS, IRT and RGS methods exhibited low false positive rates, however this came at the expense of very low true positive rates, giving them low overall recall. The performance of the PS method was between the extremes of the other methods, with both lower false positive and true positive rates compared to the highest performing burst detectors on this set of noisy spike trains.

Tables 2.4 and 2.5 show the performance of each of the burst detection methods on desirable properties D1–D4 and D5–D11 respectively. Table 2.5 ranks the burst detectors from best (1), to worst (8), based on the median and range of the box plots in Figures 2.8 and 2.9, with methods with similar results given equal rankings. Based on these rankings, three methods, namely the RS, IRT and RGS methods, clearly underperformed. To further evaluate their performance, each of the burst detectors were employed to analyse experimental recordings from mouse RGCs in the following sections.

Desirable property	Burst detection method							
	PS	MI	CMA	RS	IRT	RGS	logISI	HSMM
D1 Deterministic	✓	✓	✓	✓	✓	✓	✓	×
D2 Distribution assumption	×	✓	✓	✓	✓	×	✓	×
D3 Number of parameters	✓	×	✓	✓	✓	✓	✓	×
D4 Computational time	✓	✓	✓	✓	✓	✓	✓	×

Table 2.4: Summary of the performance of each method on binary desirable properties D1–D4. Ticks represent a burst detector possessing the stated property. Table reproduced from Cotterill et al. (2016b).

Desirable property	Burst detection method							
	PS	MI	CMA	RS	IRT	RGS	logISI	HSMM
D5 Non-bursting	4	1	7	5	6	3	1	8
D6 Non-stationary	6	2	7	4	5	3	1	8
D7 Regular bursting	4	1	2	7	6	7	5	3
D8 Non-stat bursts	4	3	5	7	6	8	2	1
D9 Long bursts	2	4	3	8	5	7	6	1
D10 High frequency	5	1	4	7	6	8	2	3
D11 Noisy bursts	5	1	2	7	6	8	4	2
Total (Relative rank)	30 (4)	13 (1)	30 (4)	45 (8)	40 (6)	44 (7)	21 (2)	26 (3)

Table 2.5: The relative rank of the performance of each method on desirable properties D5–D11 (1=best, 8=worst). Table reproduced from Cotterill et al. (2016b).

2.5.2 Performance on experimental data

Preliminary analysis of mouse RGC data

Demas et al. (2003) used MEA recordings from mouse RGCs to study the changes in spontaneous activity in the retina of control and dark-reared mice over the first six weeks of postnatal development. To examine the performance of the burst detectors at analysing experimental data, four control recordings from Demas et al. (2003) were re-analysed using each of the burst detection methods, as outlined in Section 2.4.4. The sensitivity and specificity of each burst detection method at a range of parameter values was calculated, and averaged across the five annotated spike trains from each recording, to produce the ROC curves in Figure 2.10.

Several of the burst detection algorithms place innate restrictions on the definition of bursts, for example that bursts must contain a minimum of three spikes, or restrict the total proportion of ISIs in a spike train that can be within bursts. This means that many burst detectors do not allocate either no spikes or 100% of spikes as being within bursts for any choice of parameter values, resulting in ROC curves that do not span the entire range of sensitivity and specificity values. I thus assessed each of the burst detection methods based on their minimum distance from the point of perfect classification at $(0, 1)$ for each ROC curve, rather than the area under the curve.

The MI method performed well at analysing the RGC recordings, and there exists specific parameter values at each age for which the MI method exhibited close to perfect performance. The logISI, PS and CMA methods also displayed high performance across most ages, however at P15 the PS and CMA methods exhibited higher false positive rates than other burst detectors (Figure 2.10D). The performance of the RGS method did not vary significantly as its parameter value was altered, and this method did not have high levels of sensitivity or specificity for any choice of the surprise threshold value. The sensitivity and specificity of the RS and IRT methods spanned a large range of values, however, neither of these methods were able to reach the levels of sensitivity of the other burst detectors for any choice of parameter values. The HSMM method reached high levels of sensitivity, but only at points with relatively low levels of specificity, and generally performed poorly at detecting bursts across all four ages of recordings.

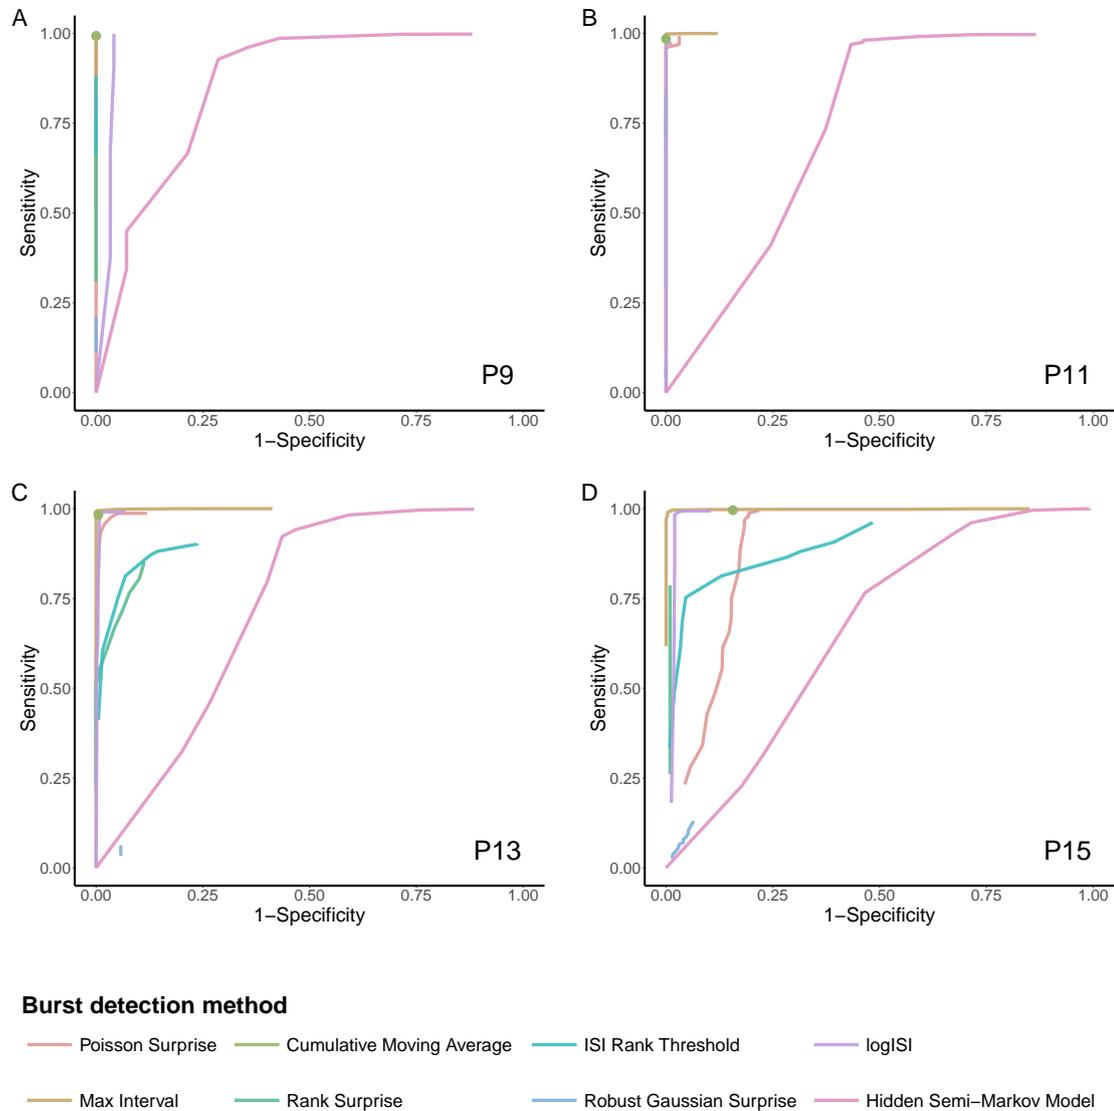


Figure 2.10: ROC curves showing the fraction of true positive (sensitivity) and false positive (1-specificity) spikes identified as being within bursts by each burst detector for a variety of input parameter values, for recordings of mouse retinal ganglion cells at **A** P9, **B** P11, **C** P13 and **D** P15. The ground truth bursts for hour-long recordings of five randomly selected electrodes at each age were determined by visual inspection, examples of which are shown in Figure 2.7, and the mean performance of each burst detector over the five electrodes at each age is shown.

Evaluation of the methods

Based on the performance of the methods in relation to the desirable properties D1-D11, the RS, IRT and RGS methods underperformed compared to the other methods in this study. This was reinforced by their poor performance when detected bursts were compared to visually identified bursts in experimental RGC recordings, where they did not reach the levels of sensitivity of the other methods. These three methods were thus excluded from any further analysis.

The HSMM method had strong performance in terms of several of the desirable properties (D7–D11), however performed poorly at analysing simulated spike trains lacking bursting activity (D5–D6). The complexity of the algorithm’s implementation also meant that it was the lowest performing of the methods on desirable properties D1–D4. The failure of the HSMM method to simultaneously achieve high levels of sensitivity and specificity when analysing experimental RGC recordings at any age cemented the decision to eliminate this method from further consideration.

2.5.3 Analysis of experimental data using high performing burst detectors

Based on the evaluation of burst detectors using simulated and experimental data in the preceding sections, four of the eight burst detectors in this study were seen to be the most promising candidates for providing high quality burst detection in a range of contexts. These were the PS, MI, CMA and logISI methods. In the following sections, these four burst detectors are further used to analyse experimental data. Firstly, the complete set of spike trains from the mouse RGC recordings from Demas et al. (2003) are analysed using these methods, and the results compared to those of the original authors. Following this, the four burst detectors are used to analyse the novel recordings from hiPSC-derived neuronal networks detailed in Section 2.4.5 and describe the ontogeny of bursting activity in these networks.

Further analysis of mouse RGC recordings

The complete set of spike trains from the four control RGC recordings from Demas et al. (2003) were analysed using the four remaining burst detectors in this study. For this analysis, the parameter values used for each burst detector were based on those that resulted in the best performance in the ROC curves in Figure

2.10, as measured by the minimum distance of the curve from the point of perfect classification in the top-left corner. In the original study of this data, Demas et al. (2003) used the autocorrelation of the spike trains to determine the average duration of bursts at each age, rather than explicitly identifying the location of bursts. By explicitly classifying bursting periods in these spike trains using the four remaining burst detectors, I was able to provide a more detailed description of the bursting activity in each recording and compare this with the authors' original results.

The four burst detectors were generally in agreement about the fraction of spikes in bursts across the four ages. Each method found that the proportion of spikes occurring within bursts decreases with increasing developmental age (Figure 2.11A). This is in agreement with the original analysis of Demas et al. (2003), who found that only very few spikes occurred outside of bursts at early ages, while by P15, many cells were active outside of bursts.

All four bursts detectors also detected bursts with similar durations over P9–P13, which resemble the values found by the autocorrelogram method of Demas et al. (2003) (Figure 2.11B). However, at P15, there was a significant deviation between the duration of bursts detected by the PS method and the other burst detectors. The bursts detected by the MI, CMA and logISI methods follow the trend of decreasing burst duration with increasing developmental age, as in the original analysis. The PS method, on the other hand, identified bursts with significantly increased duration at P15. This can be attributed to the fact that at this age, many electrodes recorded regular 'bursting episodes'. These are periods of elevated spiking activity generally spanning around 10–20 s, that consist of a series of shorter bursts (Example in Figure 2.11C). Most burst detection methods identified several short bursts in these bursting episodes, however, the PS method had a tendency to classify these episodes as one long burst, as shown in the example spike train in Figure 2.11D.

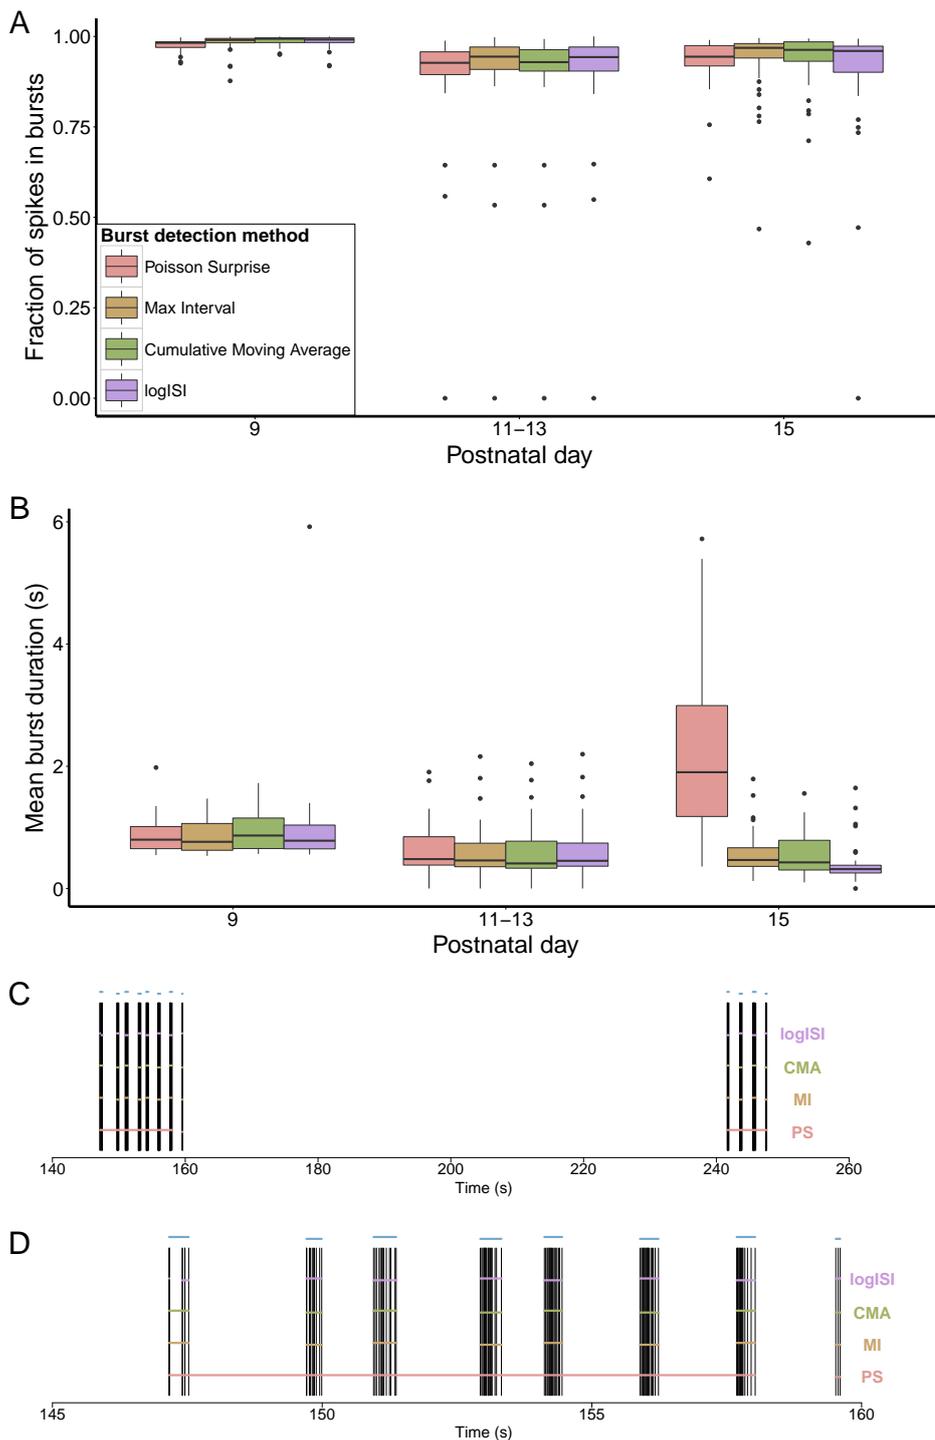


Figure 2.11: Detailed analysis of mouse retinal ganglion cell recordings. **A** Fraction of spikes in bursts, and **B** mean burst duration found by each burst detector. Each electrode was counted as one data point in the box plots. The legend in **A** applies to both **A** and **B**. **C** Bursts detected by each burst detector over a 120s sample of a P15 spike train, and **D** 15s sample showing the first bursting episode from the same spike train. Horizontal bars in **C** and **D** denote the bursts detected by each method. Blue bars above the spike train represent the bursts annotated by a human observer. Figure reproduced from Cotterill et al. (2016b).

Performance at analysing hiPSC-derived neuronal network recordings

A key aim of this study was to identify the most appropriate burst detection methods for analysing recordings of human neuronal networks. I thus next assessed the four remaining burst detectors in the study by employing each method to detect bursts in 424 MEA recordings from 73 cultures of hiPSC-derived neuronal networks. Recordings from these networks were taken at regular intervals from two to sixteen weeks after plating (WAP), as described in Section 2.4.5.

The parameter values used for the MI, PS and logISI methods in this case were chosen by inspection as those that detected bursts which most accurately agreed with visually identified bursts on five randomly chosen spike trains in this data set that had mean firing rates close to 1 Hz, which is the minimum firing rate at which regular bursting activity tends to arise in these recordings. For the CMA method, the original scale for α_1 values from Kapucu et al. (2012) was used, and the author’s suggestions for post hoc screening were also implemented. This meant that any analysed spike trains that were found to have an average burst duration above 5 s or average burst length greater than 50 spikes per burst by the CMA method, were declared as non-bursting. The resultant parameter values used to implement each of the burst detectors are shown in Table 2.6.

Method	Parameter	Value
Poisson Surprise	Minimum surprise value	$-\log(0.0025) \approx 6$
MaxInterval	Maximum beginning ISI	0.2 s
	Maximum end ISI	0.3 s
	Minimum interburst interval	0.2 s
	Minimum burst duration	0.01 s
	Minimum spikes in a burst	3
Cumulative Moving Average	α_1	1.0 if skew < 1
		0.7 if $1 \leq \text{skew} < 4$
		0.5 if $4 \leq \text{skew} < 9$
		0.3 if $9 \leq \text{skew}$
	Maximum mean burst duration	5 s
Maximum mean spikes per burst	50	
LogISI	Maximum cutoff value	150 ms

Table 2.6: Parameter values used for burst detection on human induced pluripotent stem cell-derived neuronal networks. The CMA scale parameters, and most of the MI parameters, with the exception of the maximum beginning ISI were equivalent to the values in the original implementation show in Table 2.1, while the remaining values differ. Table reproduced from Cotterill et al. (2016b).

There were some differences in the absolute level of bursting activity found by each of the burst detectors in this data set. However, the general trend from most methods showed a ‘ramping up’ of bursting, in terms of the fraction of spikes in bursts, with increasing developmental age (Figure 2.12B). The level of bursting activity in these human network recordings, as measured by the fraction of spikes in bursts, tended to be significantly lower than that which has been previously observed in recordings of rodent hippocampal or cortical networks over equivalent developmental timescales (Charlesworth et al., 2015; Chiappalone et al., 2005; Cotterill et al., 2016a). Figure 2.12B also suggests that the proportion of bursting activity in these networks may begin to decrease with age after reaching a peak around 14 WAP. This would be consistent with a previous study using calcium imaging to record the activity of human pluripotent stem cell-derived neuronal networks, which found that bursting activity initially increased with developmental age, before decreasing at later stages of development, when it is replaced by more complex firing patterns (Kirwan et al., 2015). However, additional recordings at time points > 16 WAP would be required to confirm this trend in the data presented here.

Unlike in several previous studies of rodent neuronal networks (Charlesworth et al., 2015; Demas et al., 2003; Cotterill et al., 2016a; Charlesworth et al., 2016), there was no obvious trend relating burst duration and culture age in the hiPSC-derived neuronal network recordings, with bursts detected by most methods remaining short over the entire developmental period examined (Figure 2.12A). Similarly, the degree of regularity of the bursting activity, measured by the coefficient of variation of interburst intervals (CV of IBI), did not appear to vary significantly with increasing developmental age (Figure 2.12C).

To quantify the similarity between the bursts identified by each burst detection method in these recordings, I converted each spike train into a time series by dividing the 300 s recordings into 50 ms time bins. A binary vector was then calculated to represent the bursts detected by each method. This vector takes a value of one if the spike train was in a bursting state during a time bin, or zero otherwise. I then used the Hamming distance to quantify the differences between these binary vectors. The normalised Hamming distance (Hamming, 1950) represents the fraction of points at which two binary strings differ, which for vectors u and v of length n is defined as

$$d(u, v) = \sum_{i=1}^n \frac{1 - \delta(u_i, v_i)}{n},$$

where δ is the Kronecker Delta. The normalised Hamming distance was calculated between each pair of burst detectors for every spike train in which bursting activity was detected by all four methods.

The normalised Hamming distance between the majority of methods was generally low, remaining below 0.05, or 5% at most WAPs. At 12 WAP, however, there was a peak in Hamming distances, in particular those measuring the difference between the bursts identified by the MI method and the other methods in this study. At 12 WAP, the hiPSC-derived neuronal network recordings exhibited higher mean firing rates and lower variability of ISIs compared to recordings at other WAPs, with many electrodes recording tonic spiking or bursting activity at a high frequency (examples in Figure 2.13E, F). As the MI method detects bursts based on the absolute size of ISIs, this method tended to classify a large amount of activity as bursting in these high frequency spike trains. The other burst detectors, which are based on identifying bursts as periods of elevated firing rate relative to background spiking activity, generally found a much lower proportion of bursting activity in these spike trains. Altering the parameter values of the MI method at this WAP, to reduce the maximum allowed beginning and end ISIs in a burst, could bring its results more in line with those of the other burst detectors at this age.

Visual inspection of a subset of spike trains at each WAP was also performed to gather insight on how bursts were detected in these recordings and the differences between the results from each burst detector (Figure 2.13). In several spike trains, the CMA method failed to identify numerous periods which visual inspection and the other burst detectors generally classified as bursting (Figure 2.13B,C). This may explain the lower fraction of spikes in bursts found by this method, when compared to the other burst detectors (Figure 2.12B). The logISI method also found low proportions of spikes in bursts at many WAPs (Figure 2.12B). This may account for the generally low Hamming distances between the CMA and logISI methods (Figure 2.12D), indicating a high degree of similarity between the bursts identified by these two methods.

The PS method exhibited a tendency to combine bursts that other methods identified as multiple separate bursts in these recordings, and extend bursts to incorporate additional spikes that visual inspection suggests should not be included in bursts (Figure 2.13B,D). This may account for the significantly longer burst durations found by this method up to 12 WAP (Figure 2.12A). In these highly variable spike

trains, it is unlikely that two human investigators would agree on how to visually annotate bursts. Similarly, no burst detector agreed perfectly with my own manual identification of bursts in these spike trains. However, based on visual inspection of a large subset of spike trains, I found that on average the MI method corresponded most closely to how I would annotate bursts visually.

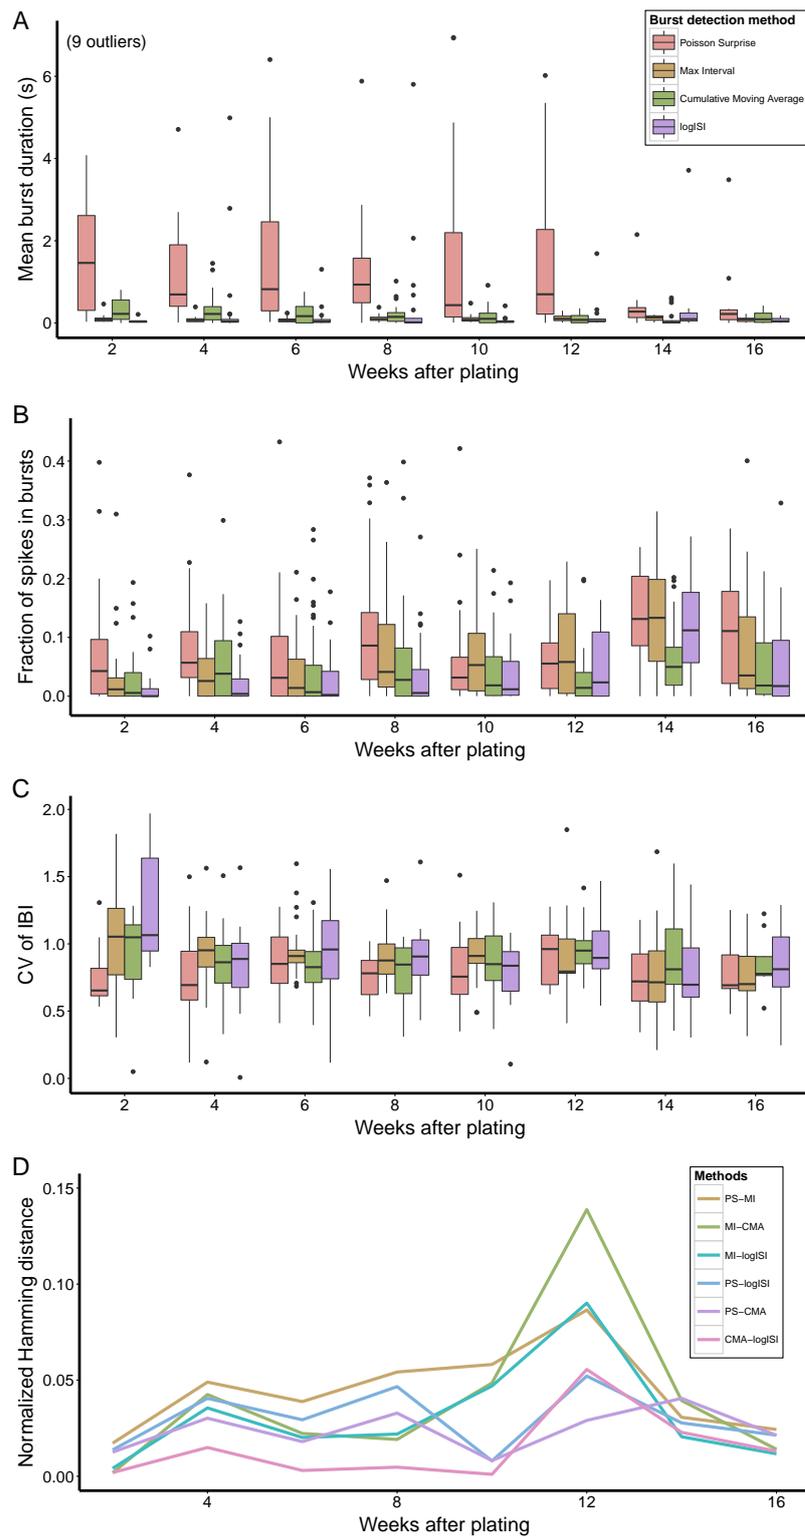


Figure 2.12: Analysis of recordings of networks of human induced pluripotent stem cell-derived neurons. **A** Mean burst duration, **B** fraction of spikes in bursts, and **C** coefficient of variation of interburst intervals (CV of IBI). Each data point in the box plots is the mean value across all electrodes from one recording. **D** Median normalized Hamming distance between each pairwise combination of burst detection methods at each week after plating. Figure reproduced from Cotterill et al. (2016b).

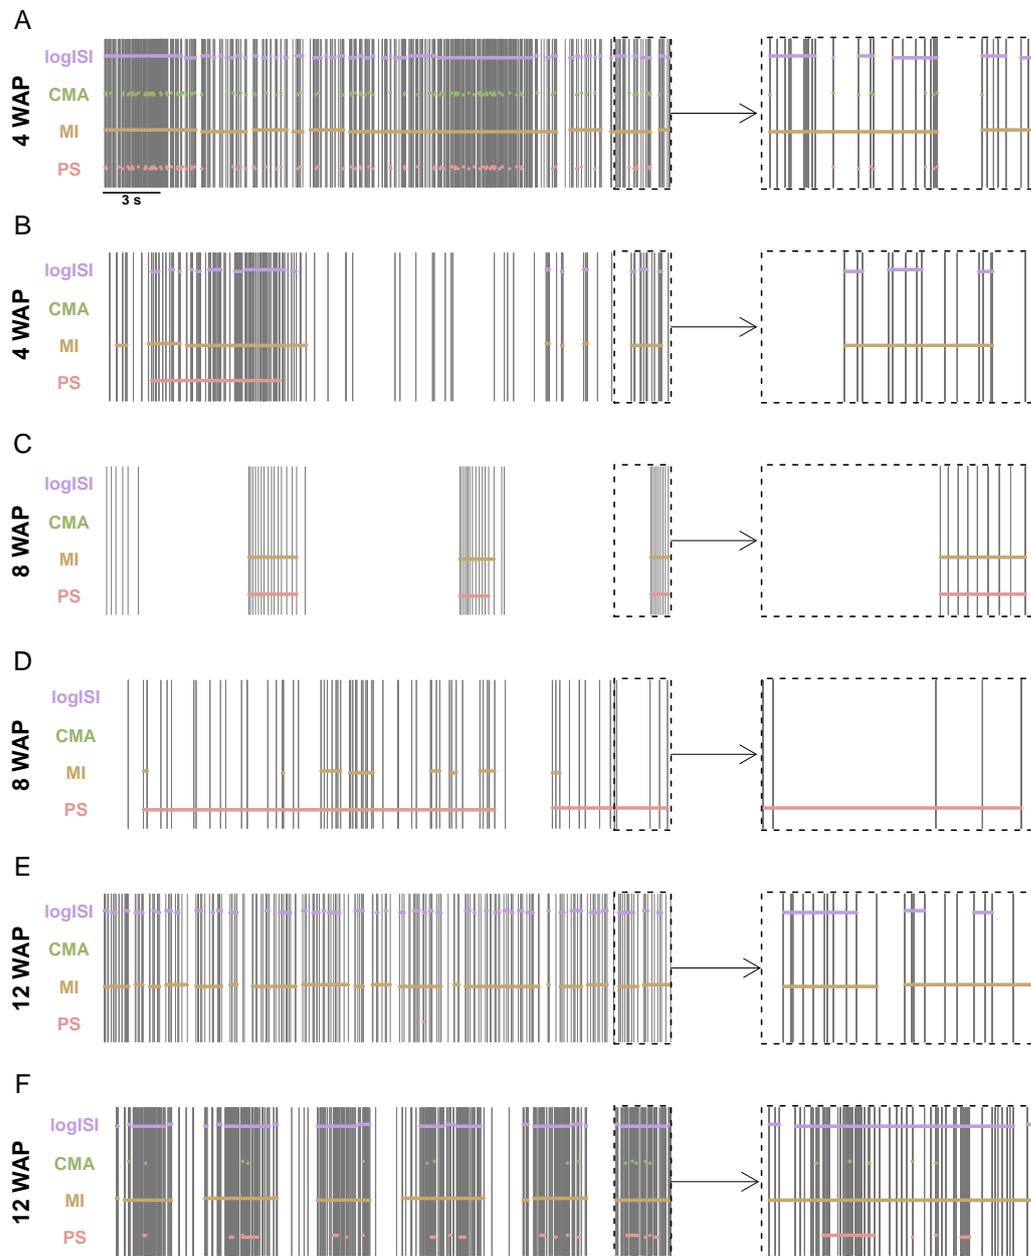


Figure 2.13: Results of the four burst detectors applied to samples of human induced pluripotent stem cell-derived neuronal network recordings at **A, B** 4 weeks after plating (WAP), **C, D** 8 WAP, and **E, F** 12 WAP. Spike trains on the left show 30s of activity, with the scale bar representing 3s. The inset on the right of each spike train is an enlarged version of the last 3s of this activity. Horizontal bars denote the bursts detected by each method. Figure reproduced from Cotterill et al. (2016b).

2.6 Discussion of existing burst detection techniques

Despite the importance of the accurate identification of bursting activity for the analysis and classification of neuronal network activity in a range of contexts, a single consistently widely used method for burst detection has yet to be adopted in the field. By assessing the performance of eight existing burst detection methods at analysing both synthetic and experimental data, I found that a number of these burst detectors perform poorly at identifying bursts in spike trains with a variety of properties.

In the previous sections, I have shown that a number of burst detection methods that were developed based on the analysis of recordings from single experimental conditions do not necessarily generalise to use on a wide range of spike trains. For example, the RGS method, which was originally developed to analyse recordings of dopaminergic neurons, could not detect the majority of bursts in spike trains simulated to contain a range of bursting behaviours. This method also performed poorly at analysing experimental recordings from mouse RGCs, even when its probabilistic threshold parameter was varied over a range of values. Other studies have also found issues using the RGS method to analyse the changes in bursting activity in neuronal networks resulting from treatment with different compounds (Eisenman et al., 2015).

The IRT method also performed poorly at accurately identifying bursts in a range of simulated and experimental spike trains. Unlike the other burst detectors studied, this method was not published in a methods paper, but rather was a heuristic method developed for the analysis of a specific data set of recordings of mouse retinal waves, which was not spike sorted (Hennig et al., 2011). The lack of adaptability of this method is thus not particularly surprising.

I have also shown that the complexity of a burst detection algorithm does not necessarily correlate with its effectiveness. The burst detector with the most complex implementation in this study, the HSMM method, often performed worse than simpler methods, particularly at analysing spike trains with sparse bursting activity. Furthermore, the high computational complexity and non-deterministic nature of this method severely limits its ability to be scaled up for use in high-throughput analysis of MEAs, which is becoming increasingly prevalent in applications such as large-scale neurotoxicity testing (Nicolas et al., 2014). These features also suggest that the HSMM method may perform poorly at analysing recordings from recently developed

high-density MEAs, which have far greater numbers of electrodes than traditional arrays, and a tendency to record few bursts on single electrodes (Lonardoni et al., 2015).

The performance of other burst detectors were hindered by their underlying assumptions, such as the RS method, which had the tendency to assign roughly the same proportion of spikes as being within bursts, regardless of how spikes were distributed in the spike trains. This resulted in a tendency of the RS method to both overestimate bursting activity in non-bursting or sparsely bursting trains and underestimate the proportion of bursting in spike trains in which the majority of spikes occurred within bursts. This makes the RS method unsuitable for analysing MEA recordings in which there is not a consistent level of bursting activity across all electrodes on an array, without manual tuning of the method's parameter values on each electrode.

The CMA method, which was designed specifically for analysing recordings of developing neuronal networks, was a promising burst detector throughout much of this study. The major limitation of this method was its tendency to erroneously detect a large amount of bursting activity in spike trains containing sparse or no bursting activity. This was particularly apparent in spike trains with temporally varying firing rates. Kapucu et al. (2012) suggest post hoc screening for outliers to address this issue. However, this can lead to underestimation of the level of bursting activity in some spike trains, as it does not allow for any shorter bursts to be identified in spike trains in which long erroneous bursts were initially detected by the CMA method. In Section 2.7, I investigate an adaptation to the CMA method aimed at overcoming this limitation.

The PS method had reasonable performance across most of the desirable burst detector properties identified here, as well as at analysing experimental RGC recordings at most ages. The main shortcomings I identified in this method was its propensity to overestimate bursting in non-bursting spike trains with non-stationary firing rates, as well as its tendency to combine temporally close short bursts into single longer bursts. Again, in Section 2.7, I examine the effect of an alteration to this method, which in this case involves applying the PS method to sliding windows across a spike train.

Based on this analysis, two existing burst detectors showed the most promise for robust burst analysis in a range of contexts. These were the MI and logISI methods.

These methods possessed most properties I deemed desirable for a burst detection method and were generally able to achieve high coherence with visually detected bursts in experimental MEA recordings when the values of their input parameters were chosen optimally. These methods, however, still had limitations. The MI method requires the choice of five parameters, the optimal values of which can be challenging to determine, particularly when analysing recordings from a variety of experimental conditions. The logISI method had a tendency to underestimate bursting in some spike trains, particularly those with non-standard bursting activity.

Since I have found no ‘perfect’ burst detector, my advice is to choose a burst detector based on the number of degrees of freedom the user wishes to control. The MI method consistently outperformed throughout this analysis, and would be my first choice recommendation when selecting a robust burst detection method. Although it has a significant number of parameters which must be set by the user, unlike burst detectors with probabilistic thresholds, these parameters are easy to interpret biologically and adjust to achieve the desired burst detection results for the specific situations in which it is utilised. If appropriate parameters cannot be found for the MI method, a high performing alternative is the logISI method, which can be implemented without choosing any input parameters. This method is most effective when there is a clear distinction between the size of within-burst and between-burst intervals on a spike train. In cases when this distinction is not apparent, the PS and CMA methods are reasonably effective alternative burst detection methods. Due to their tendency to overestimate burst durations in some situations, however, post hoc screening for outliers in terms of burst durations is advisable when using either of these methods for burst analysis.

The most robust approach to burst detection would be to use a number of burst detectors to analyse the data of interest, and compare the results of each method using summary statistics such as burst rate or duration, or measures such as the Hamming distance. If the burst detectors are largely in agreement, this provides confidence in the nature of the bursting activity identified in the experimental data. Any major discrepancies between the results from the methods can also be used to identify areas where one or more burst detectors may be performing poorly, which can be further investigated through inspection of the specific spike trains of interest. For example, periods in which the bursts identified by the MI method deviate significantly from those found by other methods may suggest that the MI parameters used for analysis were suboptimal for these spike trains. Generally, I found that for

spike trains in which bursts are easy to annotate by visual inspection, the results from the high performing burst detectors tend to be in close agreement. However, in spike trains for which two human observers may not be able to agree on how to best allocate spikes to bursts, it is likely that the burst detectors will also disagree, and discretion is required in these cases.

By employing this method of applying a number of burst detectors to the analysis of recordings of hiPSC-derived neuronal networks over a range of plating ages, I found that bursting arises in a majority of these networks around eight to ten weeks after differentiation. This is a similar time frame to the findings from some previous studies of human stem cell-derived neuronal network activity (Heikkilä et al., 2009; Kirwan et al., 2015). Although some increase in bursting activity with increasing developmental age was observed in these recordings, the rate of this increase was far lower than that which has been commonly seen in developing rodent neuronal networks (Chiappalone et al., 2005; Wagenaar et al., 2006a), including studies I have been a part of, which are described in Chapter 3.

There are several avenues through which this work could be extended. One area that was not explored here is the possibility of improving the results of burst detection by using a pre-processing step (Martens et al., 2014). Furthermore, to ensure a fair and unbiased assessment of the available burst detectors, in the previous sections, the study was restricted to evaluating the original implementation of the authors' method. In the following section, I implement adaptations to some of the methods in this study in an attempt to enhance their performance.

2.7 Adapted burst detection methods

2.7.1 Adapted CMA algorithm

As discussed in Section 2.6, the CMA method was a promising method throughout much of the evaluation of burst detectors. The CMA method possessed all of the desirable properties D1–D4 which relate to computational complexity and algorithmic assumptions, and performed reasonably well at detecting bursts in simulated spike trains containing both regular and irregular bursting activity. The main drawback of this method was its tendency to overestimate bursting in spike trains in which no bursting activity was present, especially in the case of spike trains with non-stationary

firing rates (Figure 2.8A, B). This can be attributed to the approach employed by the CMA algorithm to calculate the value of the maximum allowed ISI in a burst, $maxISI$, for use in the burst detection step.

The value of $maxISI$ is chosen based on the distribution of the cumulative moving average of the histogram of ISIs (see Section 2.4.2 for details), and is unrestricted in its absolute value. This can lead to large values being chosen for $maxISI$ in cases when the histogram of ISIs is clustered around high ISI values. Figure 2.14 shows ISI histograms for two of the simulated spike trains described in Section 2.4.3. In both of these cases, the value of $maxISI$ found using the CMA algorithm is above 1.6 s. This is far greater than the maximum intra-burst ISI generally allowed by the other burst detectors. For example, the default value of $maxISI$ in the MI method is 0.3 s. Bursts detected with large within-burst intervals > 1 s are unlikely to be biologically plausible.

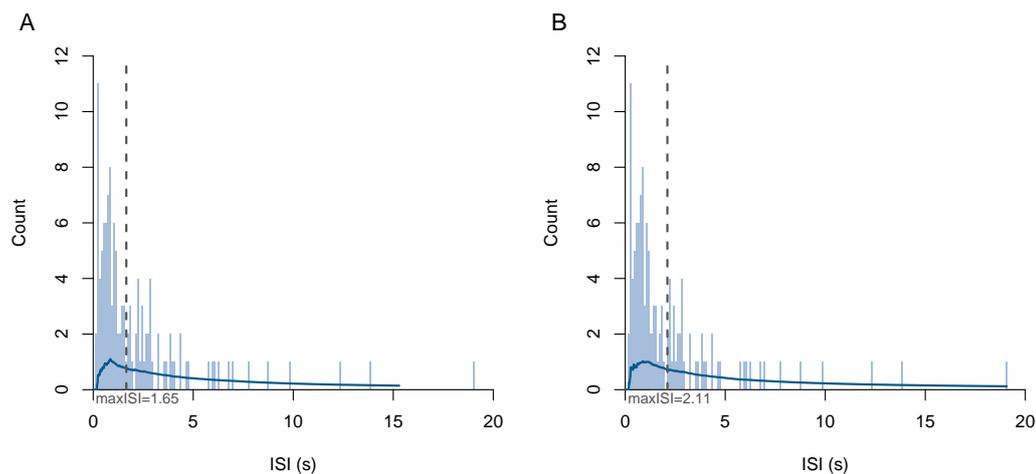


Figure 2.14: Histogram of ISIs from example simulated **A** non-bursting and **B** non-bursting non-stationary spike trains. Dark blue line shows the CMA distribution of the histogram. Vertical dotted line shows the value of $maxISI$ identified by the CMA algorithm.

The original authors recognise this limitation of their CMA algorithm, and suggest a post hoc screening procedure for recordings that are outliers in terms of their bursting activity (Kapucu et al., 2012). This involves identifying any spike trains in which the average number of spikes per burst is above 50 or average burst duration is greater than 5 s, and disregarding any bursting activity identified in these spike trains. This approach is effective at eliminating excessively long bursts identified by the CMA method. However, it can also underestimate bursting activity in some

spike trains, as it does not allow for any bursting activity to be detected in the spike trains identified as outliers. It also increases the computational time of the CMA method by requiring an additional computation step after the initial burst detection.

To improve the performance of the CMA method at performing robust burst detection in a range of contexts, I propose an adaptation to the algorithm which restricts the allowed value of $maxISI$ to a biologically realistic range. This approach is similar to that used in the logISI algorithm, which restricts the value of $maxISI$ to be below a fixed value, set at 100 ms in the original implementation by Pasquale et al. (2010) (see Section 2.4.2).

The adapted CMA algorithm is explicitly outlined below.

Thresholded CMA method

A fixed maximum cutoff value, MCV , is chosen a priori. The CMA of the histogram of ISIs on a spike train is calculated, and the skewness of this distribution is used to determine the values of two threshold parameters, α_1 and α_2 , according to the scale given in Kapucu et al. (2012). The maximum of the CMA distribution, CMA_{max} , is found and the value of $maxISI$ is set at the ISI bin at which the CMA is closest in value to $\alpha_1 \cdot CMA_{max}$. If $maxISI > MCV$, burst cores are found using MCV as the cutoff for the maximum ISI in a burst. The cutoff for burst related spikes is set as $maxISI_{brs} = maxISI$, and any spikes within $maxISI_{brs}$ of the beginning or end of the burst cores are added to the burst cores to define the final bursts (see Figure 2.15, top).

If $maxISI \leq MCV$, burst cores are found using $maxISI$ as the cutoff for the maximum within-burst ISI, and $maxISI_{brs}$ is set at the ISI value at which the CMA is closest to $\alpha_2 \cdot CMA_{max}$. Burst-related spikes within this cutoff distance of the burst cores are added to the bursts detected by this method (see Figure 2.15, bottom).

A range of values between 50 and 300 ms were trialled for MCV . A value of 200 ms was finally chosen, as it was found to provide an adequate trade-off; preventing gross overestimation of bursting activity in non-bursting spike trains, without compromising the ability of the method to accurately detect bursting in spike trains containing primarily bursting activity.

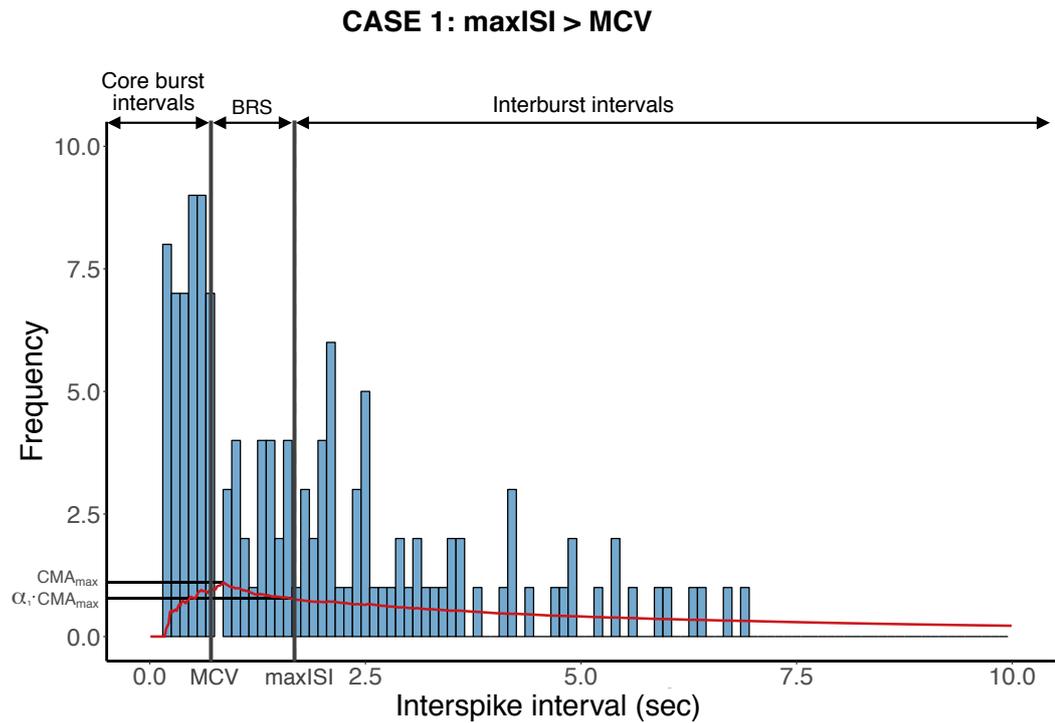


Figure 2.15: Example of ISI histograms with intraburst and interburst intervals found using the thresholded CMA method, in cases when (top) $\max ISI > MCV$ and (bottom) $\max ISI \leq MCV$. In both cases, initial burst cores are found as any series of spikes separated by ISIs in the ‘core burst intervals’ range and extended to included any spikes within ‘BRS’ (burst related spikes) of the initial burst cores.

2.7.2 Adapted Poisson Surprise method

The Poisson Surprise method also exhibited reasonable performance throughout much of the evaluation of burst detectors presented in Section 2.5. In terms of desirable features, the worst performance of the PS method was observed when this method was applied to the analysis of non-stationary spike trains containing an absence of bursting activity (D6). The PS method also performed poorly at analysing experimental spike trains from hiPSC-derived neuronal network recordings with time-varying firing rates (Figure 2.13D), and tended to detect bursts with far greater durations than those found by other burst detectors (Figure 2.12A).

Based on a discussion with Professor Simon Tavaré from Cancer Research UK, an idea was proposed to adapt the PS method by applying the algorithm separately on windows across the spike train, rather than on the spike train as a whole, to allow for a time-varying firing rate. To do so, I implemented an adaptation of the PS method which involves first using a change-point detection algorithm to segment a spike train into K distinct segments, before applying the PS algorithm to each segment individually.

The multiple change-point detection problem assumes that in a set of ordered random variables $Y_t = \{Y_1, \dots, Y_n\}$, there are $K - 1$ distinct changes in the distribution of Y_t (Braun and Müller, 1998). That is,

$$Y_t \sim G(\theta_t, \phi), 1 \leq t \leq n,$$

where G is a probability distribution with parameters ϕ , which is constant, and θ_t , which is piece-wise constant with $K - 1$ changes, one between each segment (Cleyne and Lebarbier, 2014). For the Poisson distribution,

$$G(\theta_t, \phi) = \text{Pois}(\lambda_t).$$

The optimal segmentation is the partition of Y_t into K segments which minimises the negative log-likelihood of the model.

To find the optimal segmentation, I employed the approach of Cleyne and Lebarbier (2014) in the R package `Segmentor3IsBack` (Cleyne et al., 2014). This involves using a pruned dynamic programming algorithm to find the optimal segmentation for each value of K from 1 to K_{max} by minimising the contrast over the K partitions.

The optimal value of K is then found using a penalised maximum-likelihood estimator which satisfies an oracle inequality (Cleynen and Lebarbier, 2014).

Details of the adapted PS method are given below.

Segmented PS method

The underlying activity on a spike train is assumed to follow a Poisson process with a piece-wise constant rate, λ_t , which undergoes $K - 1$ abrupt changes. The spike train is divided into K segments, using the multiple change-point detection approach of Cleynen et al. (2014), with the value of K chosen using the oracle penalty (see Figure 2.16 for example).

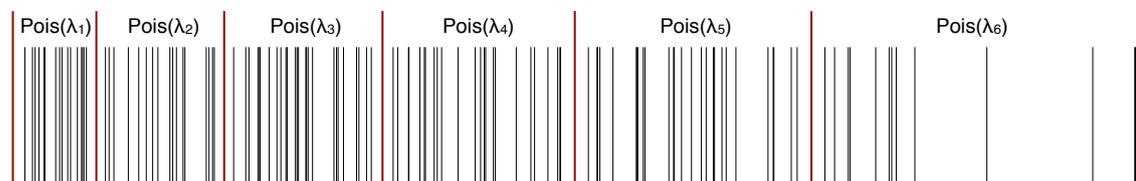


Figure 2.16: Example segmented spike train.

Within each of the K segments, the PS statistic in any period of length T containing N spikes is defined as

$$S = -\log P_t$$

where

$$P_t = \exp\left(-\lambda_t T \sum_{n=N}^{\infty} \frac{(\lambda_t T)^n}{n!}\right)$$

and λ_t is the average spike rate in the t th spike train segment.

Bursts are then detected in each segment using the surprise maximisation algorithm outlined in Legédy and Salcman (1985). First, any sequence of at least three spikes separated by ISIs less than half of the mean ISI in the spike train segment are found. Spikes are then removed from the beginning and added to the end of these initial bursts until the spike sequence with the maximum PS statistic is found, with the restriction that the bursts cannot cross any of the $K - 1$ change points. This is repeated until the end of each segment, and any bursts with a PS statistic below a pre-defined threshold level, which is constant across all segments, are discarded.

2.7.3 Performance of adapted burst detectors

The performance of the adapted burst detectors compared to the standard implementations of the PS and CMA methods was assessed by repeating much of the analysis from Section 2.5.

Results of adapted methods on simulated data

To assess the adapted burst detection methods against the desirable properties in Table 2.2, they were each used to analyse the synthetic spike trains described in Section 2.4.3. For the implementation of the segmented PS method, the minimum surprise value was set equal to that of the original PS method used in this study ($-\log(0.01)$), and the maximum possible number of segments was set at 1/60th of the duration of the spike train in seconds i.e. 5 segments in the 300 s simulated spike trains. The value of MCV in the thresholded CMA method was set at 200 ms and the scale values for α_1 and α_2 proposed by Kapucu et al. (2012) and shown in Table 2.1 were used.

The performance of the burst detectors with regards to the desirable properties D1–D4 did not change significantly with the adaptations presented here. The segmented PS method still makes an assumption about Poisson distributed spikes (D2), however, this assumption is more flexible than in the original PS method, as it allows for the Poisson rate to vary across the spike train. Due to the additional segmentation step, the computational time of the segmented PS method was found to be over 50 times greater than the standard PS method when applied to the analysis of 100 simulated spike trains. The average computational time of this method was still less than one second per 300 s spike train, however, this vast increase in computational complexity may be of relevance when analysing larger data sets e.g. from high-density MEAs. The adaptation of the CMA method did increase the number of parameters in the CMA method from zero to one, however, this is still amongst the lowest number of parameters of all of the burst detectors. The adaptation was not found to increase the computational time of the CMA method.

Figures 2.17 and 2.18 show the performance of the two adapted methods in regards to desirable properties D5–D11. For comparison, the results from the original PS and CMA methods, as well as the CMA method with burst related spikes (BRS) included (see Section 2.4.2 for details), are also shown.

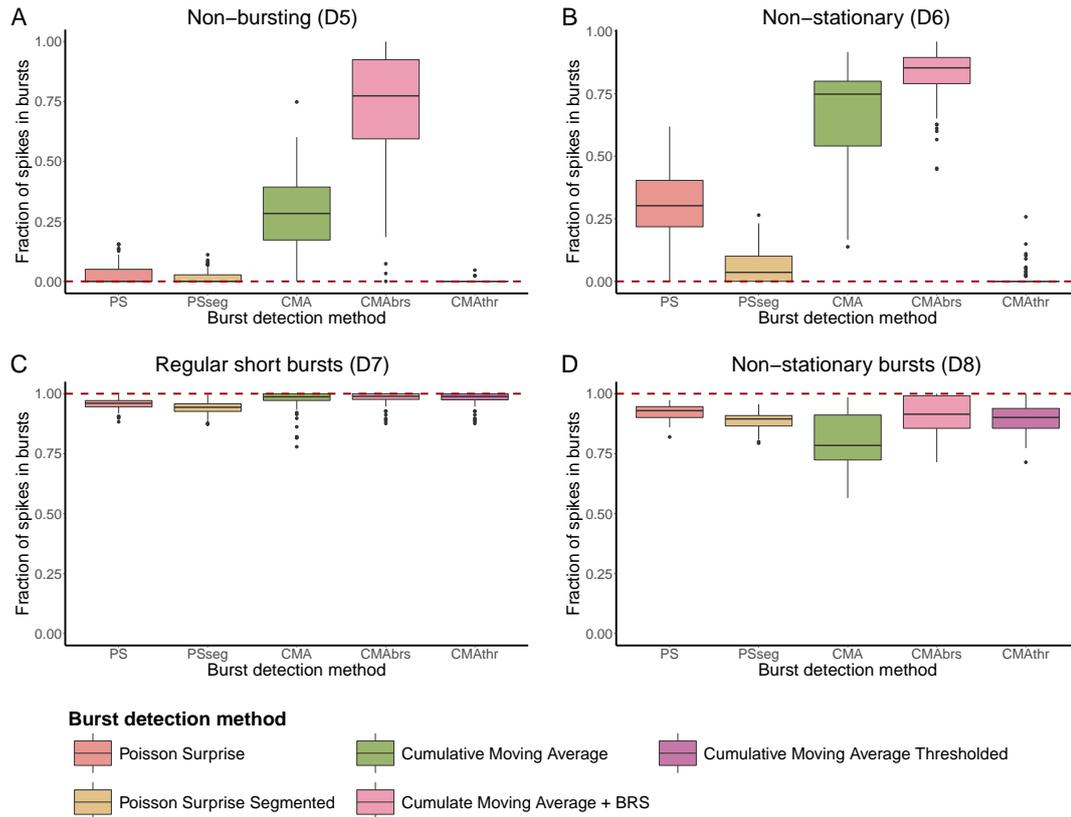


Figure 2.17: Fraction of spikes in bursts found by each burst detector in 100 synthetic trains with **A** no bursting (D5), **B** no bursting and non-stationary firing rate (D6), **C** short regular bursts (D7), **D** bursts with non-stationary burst lengths and durations (D8). Box plots and dotted line as per Figure 2.8 legend.

There was a noticeable reduction in the proportion of erroneous bursting activity detected by the segmented PS method, as compared to the standard PS method, particularly in spike trains with a non-stationary firing rate (Figure 2.17A, B). There was also a slight reduction in the proportion of correctly identified bursting activity by the segmented PS method in spike trains containing only bursting activity (Figure 2.17C, D, Figure 2.18A, C). However, this reduction in accuracy was less than the improvement in accuracy achieved by the adapted PS method when analysing non-stationary non-bursting spike trains.

Adapting the CMA method greatly improved the performance of this method at analysing simulated spike trains lacking bursting activity (Figure 2.17A, B). In particular, introducing a threshold on the value of $maxISI$ reduced the median proportion of spikes in bursts found in non-bursting non-stationary spike trains by the CMA method from almost 75%, or over 85% when burst related spikes are included, to zero (Figure 2.17B).

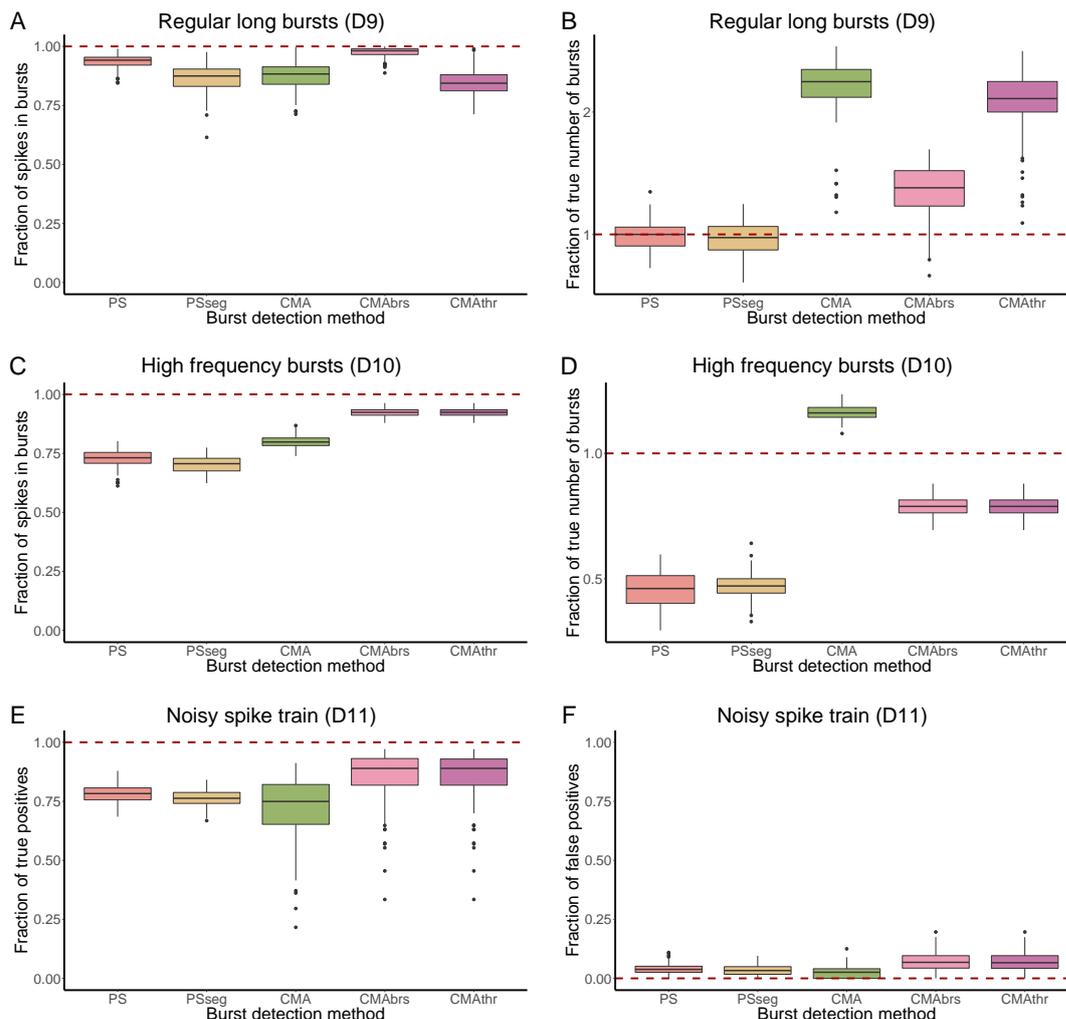


Figure 2.18: Results of each burst detector at analysing 100 synthetic spike trains. **A** Fraction of spikes in bursts, and **B** fraction of true number of bursts in spike trains with regular long bursts (D9); **C** fraction of spikes in bursts, and **D** fraction of true number of bursts in spike trains with high frequency bursting (D10); **E** fraction of true positive, and **F** fraction of false positive spikes in bursts in spike trains containing both bursting and noise (D11). Values calculated as outlined in the methods. Box plots and dotted line as per Figure 2.8 legend. **B** and **D** are presented on a log-scale.

There was a small reduction in the ability of the thresholded CMA method to accurately identify bursting activity in spike trains containing long bursts, compared to the standard CMA method (Figure 2.18A). However, because this method also incorporates burst related spikes into the detected bursts, the thresholded CMA method actually correctly identified a higher proportion of bursting activity in the simulated spike trains containing non-stationary or high frequency bursts, compared to the original method (Figure 2.17D, Figure 2.18C). The thresholded CMA method

also exhibited a small increase in false positive rate when analysing noisy spike trains compared to the original CMA method, however, this was offset by a greater increase in the true positive rate (Figure 2.18E, F). The standard CMA method with the inclusion of burst related spikes slightly outperformed the thresholded method in some cases (Figure 2.17D, Figure 2.18A), however its very poor performance at analysing non-bursting spike trains (Figure 2.17A, B) makes it an unsuitable method for detecting bursts in experimental MEA recordings.

The adaptation of the PS method did not significantly change its ranking compared to other existing burst detectors in Section 2.5.1 (Table 2.5). However, introducing the threshold into the CMA method reduces its total rank from 30 to 22, which places it as the third highest rated burst detection method based on its performance at analysing synthetic data, only one point behind the logISI method (Table 2.5).

Performance of adapted methods on RGC recordings

The adapted CMA and PS methods were further used to analyse the mouse RGC recordings described in Section 2.4.4. In the original analysis in Section 2.5.2, the CMA and PS methods had strong performance at analysing spike trains from P9–P13, however at P15, these two methods exhibited higher false positive rates than many of the other burst detectors (Figure 2.10D). The same spike trains were analysed using the adapted methods, while altering the value of the maximum number of segments in the segmented PS method and the threshold cutoff value, MCV, in the thresholded CMA method. Compared to the standard PS method, the segmented PS method exhibited higher levels of specificity at equivalent levels of sensitivity when analysing the P15 spike trains (Figure 2.19). However, the performance of this method was still lower than that of the logISI and MI methods on this data (Figure 2.10D).

Introducing the threshold into the CMA method allows this method to achieve a range of levels of sensitivity and specificity, compared to the original CMA method which was represented as a single point in the ROC curves in Figure 2.10. However, the minimum distance of the ROC curve of the thresholded CMA method from the point of perfect classification at $(0, 1)$ is only slightly smaller than the distance of the single sensitivity/specificity value of the standard CMA method, indicating that the adapted method does not offer a significant improvement in performance in this case (Figure 2.19).

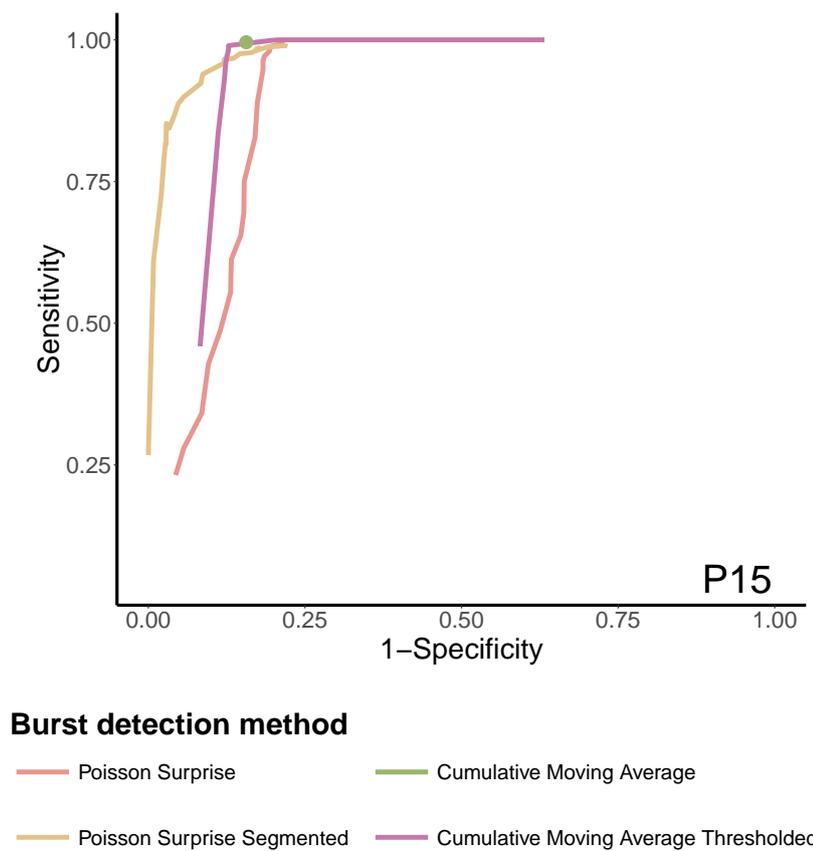


Figure 2.19: ROC curves showing the fraction of true positive (sensitivity) and false positive spikes (1-specificity) identified as being within bursts by each burst detector for a variety of input parameter values, for recordings of mouse retinal ganglion cells at P15.

The original PS method identified bursts with far greater durations than the other burst detectors when used to analyse the full set of spike trains from the mouse RGC recordings at P15 (Section 2.5.3, Figure 2.11). When the segmented PS method was used to analyse these P15 spike trains, the median duration of bursts found was 0.62 s. This is much more in line with that of the other burst detectors, compared to the original PS method, which found bursts with a median duration of 1.90 s. The adaptation of the CMA method did not significantly alter the properties of the bursts identified in the RGC recordings, compared to those found by the original CMA method.

Performance of adapted methods on hiPSC-derived neuronal network recordings

Next, the adapted burst detectors were applied to the analysis of the recordings of hiPSC-derived neuronal networks described in Section 2.4.5. As can be seen in Figure 2.20, the segmented PS method detected bursts with lower median durations and percentage of spikes in bursts in these recordings than the original PS method. This is also evident in the example spike trains shown in Figure 2.21, in which the segmented PS method can be seen to detect a lower amount of overly long bursts than the standard PS method. The results from the segmented PS method are more in line with those of other burst detectors, although the durations of bursts identified at early developmental ages by the segmented method still remained higher than those found by the other high performing burst detection methods examined in Section 2.5.3 (Figure 2.12A).

The adapted CMA method tended to detect lower proportions of spikes in bursts and shorter burst durations than the standard CMA methods in these recordings at all ages except 14 WAP, although the difference between the statistics of these bursts is far lower than those between the PS method and its adaptation. The properties of the bursts identified by the adapted CMA method are more in line with those found by the MI and logISI methods on this data, compared to those from the original CMA method (Section 2.5.3, Figure 2.12). Again, no relationship between CV of IBI and developmental age was evident in these recordings using the adapted burst detection methods (Figure 2.20C).

As can be seen in Figure 2.20D, the Hamming distances between the original PS and CMA methods and their adaptations were generally much lower than those between the different burst detections examined previously (Section 2.5.3, Figure 2.12D), remaining below 2% over the entire developmental period. The Hamming distances between the CMA and thresholded CMA method are particularly low, indicating that there are only small differences in the bursts detected by these two methods. This is also apparent in the example spike trains, in which the only major deviation between the bursts detected by the two methods is in the spike train in Figure 2.21F. In this example spike train, the adapted CMA method detects bursts with significantly longer durations than the original CMA method. This can be attributed to the fact that burst related spikes are included in the bursts detected by the thresholded CMA method, but not in the implementation of the original CMA method used in this analysis.

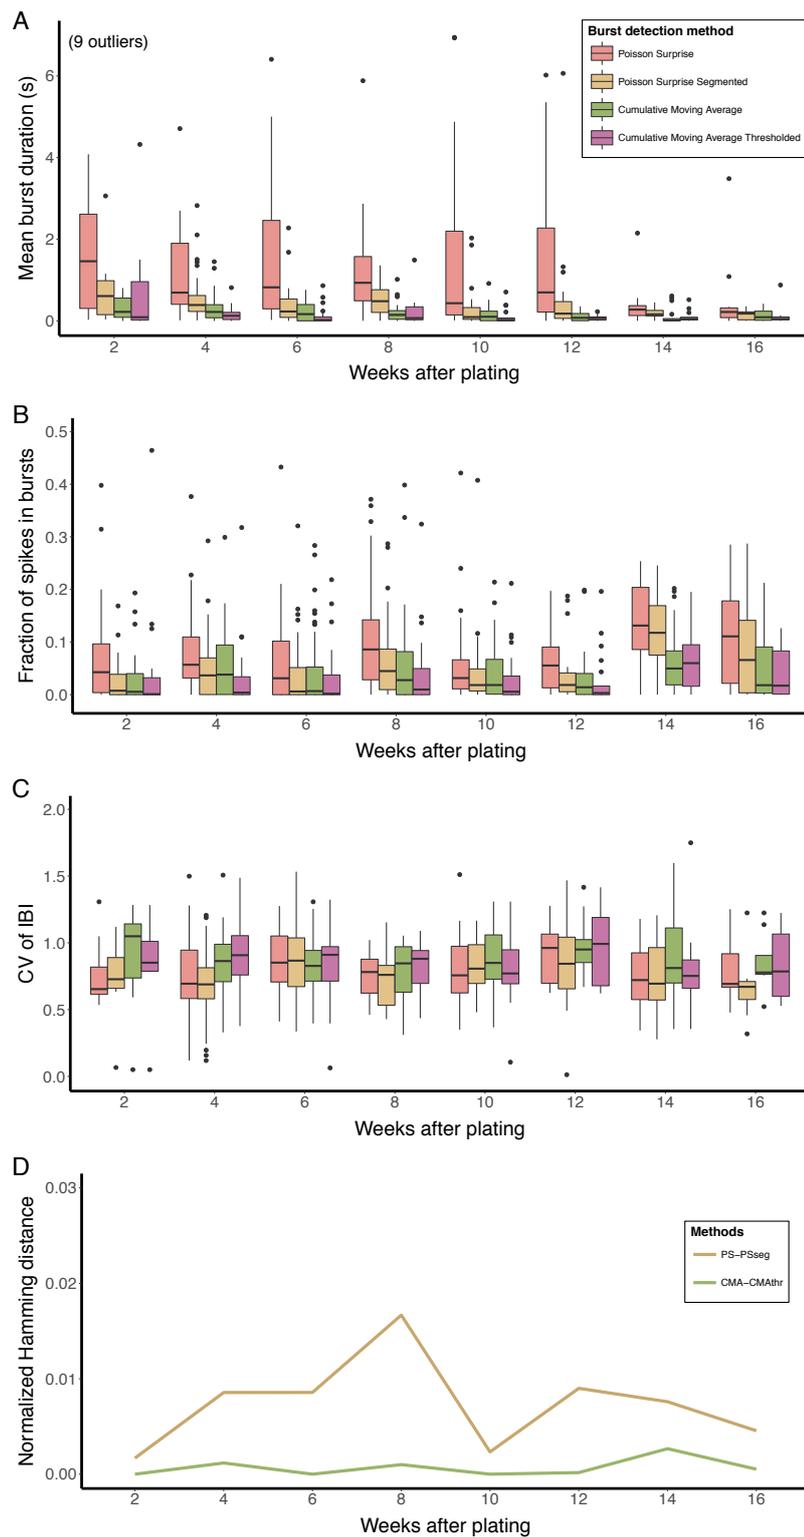


Figure 2.20: Analysis of recordings of networks of human induced pluripotent stem cell-derived neurons. **A** Mean burst duration, **B** fraction of spikes in bursts, and **C** coefficient of variation of interburst intervals (CV of IBI). Each data point in the box plots is the mean value across all electrodes from one recording. **D** Median normalized Hamming distance between each pairwise combination of burst detection methods at each week after plating.

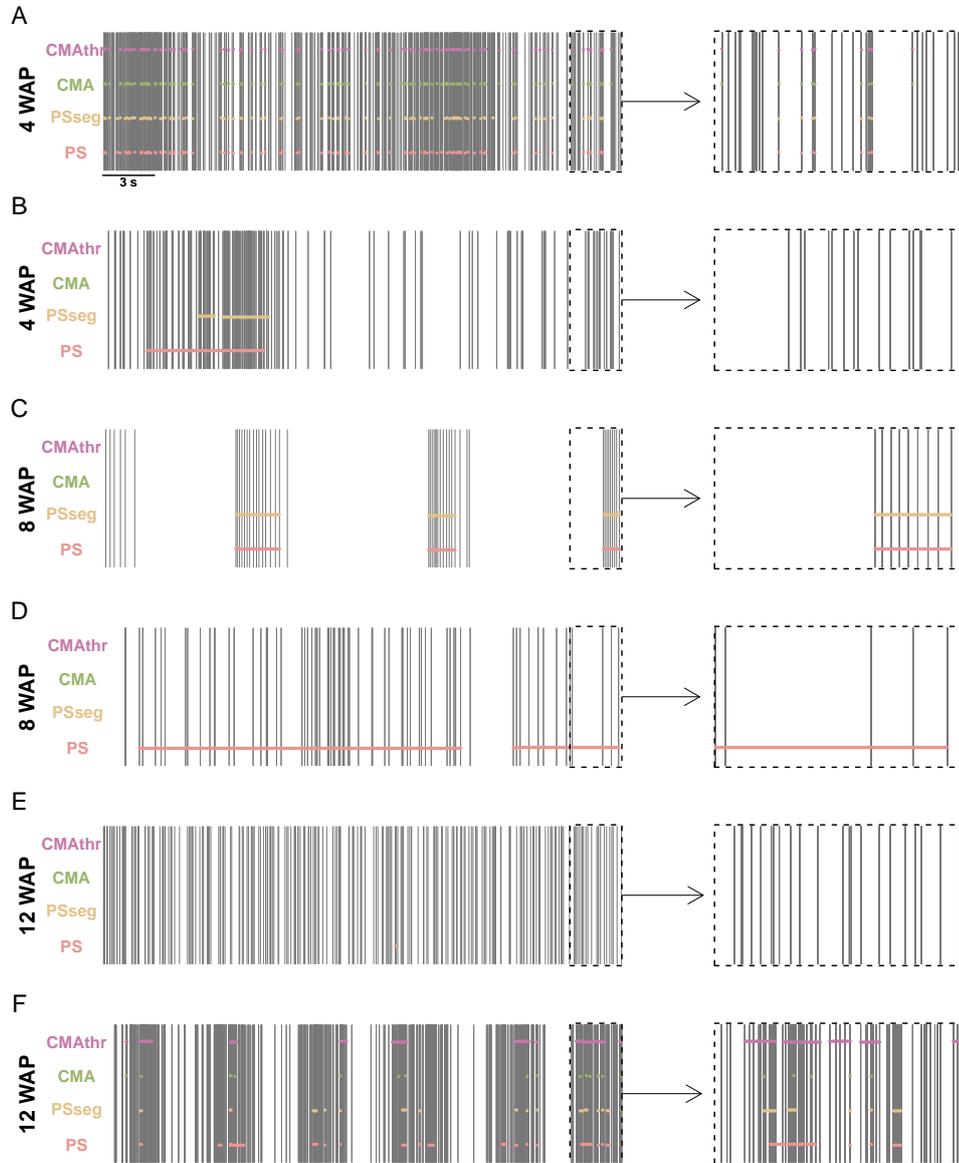


Figure 2.21: Results of the four burst detectors applied to samples of human induced pluripotent stem cell-derived neuronal network recordings at **A**, **B** 4 weeks after plating (WAP), **C**, **D** 8 WAP, and **E**, **F** 12 WAP. Spike trains on the left show 30 s of activity, with the scale bar representing 3 s. The inset on the right of each spike train is an enlarged version of the last 3 s of this activity. Horizontal bars denote the bursts detected by each method. Spike trains are equivalent to those in Figure 2.13.

2.7.4 Discussion of adapted burst detection techniques

In this section, I implemented adaptations to two existing burst detection techniques, namely the PS and CMA methods and assessed their performance using both synthetic and experimental data. Although the adaptation of the PS method presented here reduced its tendency to detect overly long bursts, particularly in spike trains with non-stationary firing rates, it also increased the computational complexity of the method, and decreased its accuracy at detecting bursts in other types of synthetic spike trains. The small improvement offered by this adaptation thus does not justify its inclusion amongst the top methods for burst detection.

The adaptation I developed for the CMA method, however, is conceptually simple and does not increase the computational time of the method. In fact, by removing the need for a post-processing step to eliminate outliers, the adaptation presented here can decrease the computational time required to apply the CMA method. My adaptation significantly improves the performance of this method at analysing spike trains with no or sparse bursting activity. Additionally, restricting the allowed values for $maxISI$ in the CMA algorithm allows for burst related spikes to be incorporated into the detected bursts without causing significant overestimation of the level of bursting activity in non-bursting spike trains. This improves the accuracy of the adapted method at detecting bursts in spike trains containing primarily bursting activity. The introduction of a threshold into the CMA method also increases its flexibility by allowing the user to alter MCV , which provides some control over the sensitivity of the method at detecting bursts. However, as with all burst detectors, there are some limitations of the adapted CMA method, and this method was seen to underestimate bursting in some spike trains with non-standard bursting activity.

2.8 Conclusions

In this chapter, I have evaluated a variety of existing and adapted burst detection methods to assess their performance at performing robust and accurate identification of bursting activity in both synthetic spike trains and experimental recordings containing diverse bursting behaviour. The purpose of this study was to provide implementable advice to researchers for analysing experimental recordings, particularly those from human neuronal networks, which have been observed to exhibit more variable bursting activity than previously studied rodent neuronal networks.

Based on this analysis, I recommend three burst detectors, namely the MaxInterval, logISI and the thresholded CMA methods. These methods were able to detect bursts accurately in both bursting and non-bursting synthetic spike trains, and had strong performance at analysing experimental data. The most appropriate method for a specific analysis situation can be chosen based on the properties of the data to be analysed and the aims of the researcher. The MI method offers the most flexibility, with numerous parameters which can be adjusted by the user to identify bursts with specific properties of interest. The logISI and thresholded CMA methods provide more objective approaches to burst detection which adapt to the properties of the spike trains with little intervention by the user. The logISI method may be preferred when there is a clear separation of within and between burst intervals, while the thresholded CMA method is a high performing alternative when this is not the case. In situations where computational time is not a key concern, the most robust approach to burst detection would be to use all three of these burst detectors to characterise the bursting activity, and compare the results from each method.

2.8.1 Limitations and future work

One limitation of this analysis of burst detection methods was the limited number (eight) of burst detectors examined. This was a deliberate choice, as the extensive number of burst detectors available makes an exhaustive analysis of all methods impossible. Instead of providing a brief analysis of all burst detection methods, the scope of the study was restricted in order to provide a thorough assessment of what I judged as the most promising existing burst detection algorithms, and to offer implementable recommendations and software to researchers working in this area. The eight burst detectors in this study were chosen to represent a broad sample of the major approaches to burst detection, and many of the burst detectors that were not included in this study take similar approaches. For example, there are many variations on the use of the ISI histogram or log-adjusted ISI histogram to determine the value of threshold parameters for burst detection.

The results presented in this chapter were also influenced by how the ‘ground truth’ bursts were chosen by visual inspection in the experimental RGC recordings analysed in Sections 2.5.2 and 2.7.3, which is necessarily a subjective choice. However, the relatively high degree of coherence between the manually annotated bursts and those identified by a number of bursts detectors, suggests that my definition of bursts largely agreed with that of other investigators who have previously developed burst detection techniques.

Given the observed prevalence of ‘network bursts’, or synchronised bursting activity across many electrodes, in developing neuronal networks recorded on MEAs (see Section 1.5.4), an area of future interest would be to investigate if information gathered from observing the array-wide activity could be used to improve the detection of bursts on single channels. Martens et al. (2014) showed that the peaks corresponding to intra and interburst spikes in an ISI histogram were better separated when pooled ISIs from multiple electrodes of an MEA were used, rather than a single spike train, with further separation achievable through the use of return map analysis. Using network-level information to inform the choice of threshold parameter values in single-electrode burst detection techniques thus has the potential to improve their performance, compared to just considering the features of individual spike trains.

The use of autonomous methods, either based on single-electrode or array-wide activity, to set some of the key parameters of the MI method is also an area of future interest. This would reduce the level of manual intervention required to use this method and the subjectivity inherently introduced through manual parameter selection. It would also allow the threshold parameters of this method to be more easily adjusted for each individual electrode, rather than using fixed array-wide values, which may be of particular relevance when analysing recordings from high-density arrays.

Another area for consideration relates to which features of bursts are the most informative to extract from recordings of neuronal network activity. In other studies of rodent neuronal networks, I have shown that the temporal structure of bursting activity, measured by features such as the CV of within-burst or between-burst ISIs, can be important for distinguishing between different types of network activity (see Chapter 3). However, in the recordings from human neuronal networks examined in this chapter, I found no strong relationship between the CV of IBI and the developmental age of arrays. A greater understanding of which are the most distinguishing features of bursting activity in human neuronal networks may inform future approaches to burst detection in these networks. The protocols for recording from hiPSC-derived neuronal networks are also still in their infancy, and the refinement of these experimental methods may lead to more stable recordings with clearer developmental profiles of activity in the future.

As well as its relevance for analysing the activity of neurons, the detection of ‘bursts’, or events with usually high frequency, is also a more general problem in time series analysis, with applications in a variety of fields. For example, identifying bursts of gamma rays can aid in the detection of black holes and other astronomical phenomena and detection of periods of high trading volume of a stock is of relevance to regulators looking for insider trading (Zhu and Shasha, 2003). A recent area of interest in this field is the detection of bursts in streams of information that arrive continuously over time (He and Parker, 2010).

An influential model of bursting in this field was developed by Kleinberg (2002) to identify periods of elevated frequency of certain words in an incoming email stream. This method uses an infinite state automaton, in which each state represents a certain arrival rate and the cost of a change in state is proportional to the increase in state number, to produce a time series of burst ‘weights’ or levels of burstiness by finding the optimal state transition sequence. Versions of this method have been used for detecting bursts in a variety of contexts, such as finding surges of interest in certain topics in scientific publications over time (Boyack et al., 2004; Mane and Börner, 2004) and identifying bursty communities of blogs (Kumar et al., 2003).

Another approach to burst detection in data streams is to detect bursts as periods in which the frequency of an event in a sliding window exceeds a threshold value (Zhu and Shasha, 2003). Zhang and Shasha (2006) developed an elastic burst detection approach based on this idea, which uses specialised data structures called Shifted Aggregation Trees and a heuristic algorithm to efficiently identify bursts in all possible window sizes in a data stream.

Ideas from these burst detection techniques developed in other domains may be useful for informing future approaches to burst detection in a neuroscience context. Although the spike trains analysed in this chapter were fixed and short in duration, the use of such real-time burst detection techniques could be relevant for other neuroscience applications. These include the analysis of neuronal network recordings made over extended continuous time periods, studies of real time learning in embodied cultured networks or applications involving bidirectional communication between biological tissue and computer interfaces (Wagenaar et al., 2005; Bakkum et al., 2004).

Postscript

In July 2017 I became aware of Valkki et al. (2017), published on 31st May 2017. Citing the issues with the CMA method reported in our paper (Cotterill et al., 2016b), which are discussed in Section 2.6, this article presents an adaptation to the CMA method, called multi-CMA, to address these issues. As discussed in this section as a possible future approach to burst detection, the multi-CMA method utilises network-wide activity to set the threshold for *maxISI* used for burst detection. Based on the results in Valkki et al. (2017), the adapted method does appear to address the issues I raised in this chapter, however, further analysis would be required to confirm this and assess the performance of this new burst detection method against the other burst detectors in this study.

Chapter 3

Analysis of spontaneous neuronal activity on microelectrode arrays

The majority of the material in this chapter has been published in two journal articles. The study in Section 3.2 has been previously published in the article ‘Quantitative differences in developmental profiles of spontaneous activity in cortical and hippocampal cultures’ (Charlesworth et al., 2015). All analyses presented in this section are my own work, and the joint first author on this paper, Paul Charlesworth, was involved solely in the collection of experimental data for this study.

The analysis presented in Section 3.3 was performed in collaboration with Diana Hall and Timothy Shafer from the US Environmental Protection Agency and ideas regarding the analysis approach were developed through discussions with these collaborators. This work has been previously published in ‘Characterization of early cortical neural network development in multiwell microelectrode array plates’ (Cotterill et al., 2016a). This article contains analysis which was performed by our collaborators as well as myself. However, any analysis in the article which was not performed by myself personally has not been included in this thesis.

3.1 Summary

The presence of spontaneous neuronal activity during development has been observed in numerous parts of the nervous system, and is thought to play an important role in shaping the structure of neuronal networks (see Section 1.5 for details). Recordings of the spontaneous activity of *in vitro* neuronal networks have been

used extensively to study the time course development of neuronal networks in both control and genetically or chemically manipulated states (Wagenaar et al., 2006a; Charlesworth et al., 2016; Valdivia et al., 2014). A common method of recording this activity *in vitro* is through the use of microelectrode arrays (MEAs), which allow the spiking activity of large networks of neurons to be recorded simultaneously and in a non-invasive manner (see Section 1.5.3 for details). Due to their prevalence, a variety of MEA devices have been developed. These include single-well MEAs, with electrodes numbers ranging from values in the order of tens to thousands, as well as multi-well microelectrode arrays (mwMEAs), which consist of multiple independent wells containing generally a lower number of electrodes per well than their single-well counterparts (Spira and Hai, 2013; Valdivia et al., 2014). In this chapter, I present the analysis of two independent experiments involving recordings of spontaneous neuronal network activity on *in vitro* MEAs.

The first study uses experimental data collected by Paul Charlesworth, formerly of the Department of Physiology, Development and Neuroscience. These recordings are from hippocampal and cortical networks cultured on single-well MEAs, which I employ to characterise and compare the ontogeny of spontaneous activity in these two brain regions. I develop a framework for quantitatively characterising the activity observed in these recordings at a variety of ages using a number of common properties of spontaneous neuronal activity. This includes the analysis of bursting activity, informed by the results from the evaluation of burst detectors in Chapter 2. Using this framework, I show that spontaneous activity in both network types appeared within a few days of plating and strengthened over early development, before reaching a plateau around 21 days *in vitro*. A number of features were shown to differ significantly between hippocampal and cortical networks. In particular, the activity of neurons in hippocampal networks was more correlated, exhibited higher rates of theta-bursting and had lower variability in the size of interburst intervals compared to cortical networks of equivalent age. I further show that the quantitative features used in this study are sufficient to differentiate cortical and hippocampal recordings at each age using common supervised classification methods.

The second study in this chapter employed experimental recordings of neuronal networks on multi-well MEAs (mwMEAs) generated by Timothy Shafer, Kathleen Wallace and William Mundy at the US Environmental Protection Agency. The analysis was conducted in collaboration with Diana Hall and Timothy Shafer. Using a similar quantitative framework as in the first study in this chapter, I examine the

development of spontaneous activity in these mwMEA recordings and find that it resembles the activity patterns of similar cultured networks recorded on single-well MEAs. I also demonstrate that the recordings can be classified by age with reasonable levels of accuracy using the quantitative features in our study. Further, I develop a re-sampling approach for determining the optimal experimental protocol for using mwMEAs to simultaneously study cultured networks under a number of experimental conditions, which has applications in areas such as investigating the effects of chemical manipulations on neuronal network activity in a high-throughput manner.

3.2 Comparison of developmental spontaneous activity in cortical and hippocampal cultures

3.2.1 Introduction

As outlined in Section 1.5, spontaneous activity has been associated with a range of physiological processes in developing neuronal networks, and the study of the nature of these temporally changing activity patterns has the potential to increase our understanding of the structural and functional development of the nervous system. Commonly, MEA recordings of cultured neurons are used to study developmental spontaneous activity *in vitro* (see Section 1.5.3). Although many common features of developmental spontaneous activity patterns in such *in vitro* systems have been documented (Wagenaar et al., 2006a; Chiappalone et al., 2006; Rolston et al., 2007), a systematic understanding of how these properties vary across developmental time and which features are the most informative for describing the activity patterns remains lacking. Greater understanding of these aspects of spontaneous activity could aid in revealing which features of the activity patterns are essential for the proper development of functional neuronal circuits across the nervous system.

To address these issues, I performed an analysis of the developmental profiles of spontaneous activity in cultured networks from two well studied areas of the nervous system, the hippocampus and the cortex. As outlined in Section 1.3.1, the cerebral cortex is the thin outer layer of the cerebral hemispheres, and consists of a variety of regions, responsible for such functions as the processing of sensory information, motor functions, perception and thought. The hippocampus is located underneath the cerebral cortex, inside the temporal lobe and is primarily associated with memory

and spatial navigation (Purves et al., 2001).

Although the presence of spontaneous activity during neuronal development has been previously observed in hippocampal and cortical regions both *in vitro* and *in vivo* (Garaschuk et al., 1998; Leinekugel et al., 2002; Seelke and Blumberg, 2010; Wagenaar et al., 2006a), a comprehensive quantitative description and comparison of the developmental profile of activity in these two brain regions has yet to be performed. In this section, I present the results from a quantitative analysis of MEA recordings of cultured networks composed of dissociated neurons from embryonic mouse hippocampus, as well as networks consisting of cortical neurons obtained using identical experimental protocols. As well as offering insights about the differences between spontaneous activity patterns in hippocampal and cortical regions specifically, this work also provides a framework for examining the differences in the spontaneous activity of cultures under different conditions, for example wild type versus mutant or control versus chemically treated networks.

3.2.2 Methods

MEA recordings

The neuronal networks used in this analysis were prepared from cultures of dissociated hippocampal and cortical neurons at embryonic day 17–18, and plated on 60-channel single well MEAs from MultiChannel Systems. In the four weeks after plating, two 15-minute recordings of each array were made each week. A fixed threshold of $-20 \mu V$ was used for spike detection, and action potential waveforms were confirmed using 1ms pre and 2ms-post threshold record samples. Spikes were not sorted to separate signals produced by individual neurons, and so the spikes from each electrode represent multi-unit activity. A total of 214 recordings from 32 arrays of cortical neurons, and 329 recordings from 61 hippocampal networks at various days from 7 to 28 days *in vitro* (DIV) were produced for use in this analysis. All experimental procedures were performed by Paul Charlesworth and further experimental details are available from Charlesworth et al. (2015).

Data analysis

To quantify the activity in each of the 15-minute MEA recordings, the value of eleven features were calculated. These features were chosen to capture commonly observed properties of the spontaneous electrophysiological activity of developing neuronal networks, such as bursting, correlated network activity and theta bursting. The majority of features were calculated individually for each electrode and a summary array-wide value was computed from all of the electrode values on a single MEA. The eleven features are defined as below:

1. *Firing rate* - The number of spikes on an electrode divided by the length of the recording period. The array value was the median of all electrode values.
2. *Within burst firing rate* - Bursts were detected on each electrode using the MaxInterval Method (Nex Technologies, 2014), with parameter values shown in Table 3.1 (see Section 2.4.2 for details). The firing rate within each burst was calculated, and averaged across all bursts on an electrode. The array value was calculated as the median of within-burst firing rates on all electrodes, excluding those electrodes that recorded no bursting activity.

Parameter	Value
Maximum beginning ISI	0.1 s
Maximum end ISI	0.25 s
Minimum IBI	0.8 s
Minimum burst duration	0.05 s
Minimum number of spikes in a burst	6

Table 3.1: MaxInterval method burst detection parameters.

3. *Burst rate* - The number of bursts on an electrode, divided by the length of the recording period in minutes. The array value was calculated as per feature 2.
4. *Burst duration* - The electrode value was calculated as the mean duration of all bursts on an electrode. The array value was calculated as per feature 2.
5. *Fraction of spikes in bursts* - The electrode value was calculated as the total number of spikes classified as being within a burst, divided by the total number of spikes recorded on the electrode. The array value was calculated as per feature 2.

6. *Coefficient of variation of interburst intervals (CV of IBI)* - The electrode value was the standard deviation of all interburst intervals (IBIs) on an electrode divided by the mean of all IBIs on the electrode. The array value was calculated as per feature 2.
7. *Network spike rate* - Network spikes (Eytan and Marom, 2006) were defined as any period in which 10 or more electrodes were active within a 3 ms time bin. The array value was calculated as the number of network spikes detected on the array divided by the length of the recording period in minutes.
8. *Network spike peak* - The maximum number of active electrodes during each network spike was found. The array value was calculated as the median of the peak values from all network spikes observed on an array.
9. *Network spike duration* - The duration of a network spike was defined as the length of time during which the number of active electrodes exceeded the threshold of 10. The array value was calculated as per feature 8.
10. *Mean correlation* - The correlation between the activity on each pair of electrodes was calculated using the spike time tiling coefficient (Cutts and Eglen, 2014). This is defined as

$$STTC = \frac{1}{2} \left(\frac{P_1 - T_2}{1 - P_1 T_2} + \frac{P_2 - T_1}{1 - P_2 T_1} \right)$$

where P_1 is the proportion of spikes on electrode 1 that occur within $\pm\Delta t$ of a spike on electrode 2, and T_1 is the fraction of the total recording time that lies within $\pm\Delta t$ of the spikes on electrode 1. P_2 and T_2 are defined similarly for electrode 2. A coincidence window of $\Delta t = 5$ ms was used. Windows of $\Delta t = 50$ ms and $\Delta t = 0.5$ ms were also trialled, with qualitatively similar results. The array value was calculated as the mean of the correlation values between each distinct pairwise combination of electrodes on an array.

11. *Fraction of electrodes exhibiting theta bursting* - For each electrode, the histogram of log-adjusted ISIs was calculated and smoothed using a Gaussian kernel density estimation routine (Venables and Ripley, 2002). An electrode was classified as theta bursting if a peak was found in the region of 4 to 10 Hz of this histogram (see Figure 3.1). The array value was calculated as the fraction of electrodes on the array that were classified as exhibiting theta bursting.

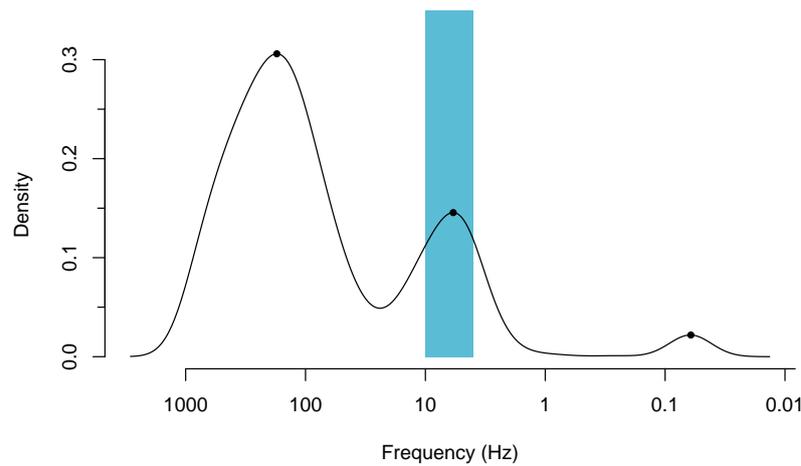


Figure 3.1: Detection of theta bursting on an electrode. Plot shows the smoothed log-adjusted ISI histogram from an electrode with peaks marked. Blue box shows the frequency region representing theta bursting. An electrode was classified as theta bursting if one of the peaks falls within this region. Figure adapted from Charlesworth et al. (2015).

Classification of network types

Classification of the MEA recordings by their culture type was performed using two common supervised learning techniques. For each technique, ten separate binary classifiers were built, one for each age, to predict the brain region of origin of the recordings (cortex or hippocampus) using the eleven features detailed above.

The first classification approach involved the use of bagged classification trees (Breiman, 1996) with 500 bootstrapped training sets. The second approach was to use Support Vector Machines (SVMs) (Boser et al., 1992) with a radial kernel. A separate SVM was built to classify the recordings at each age, using parameter values that were chosen by 10 fold cross-validation on the entire data set. For both types of classifiers, two-thirds of the recordings were used as training data and the remaining one-third used to calculate test accuracy. Five hundred repetitions with random choices of test and training sets were performed, and the mean classification accuracy across these trials is reported.

3.2.3 Results

Developmental profile

Activity rate

Spontaneous spiking activity emerged in both hippocampal and cortical cultures around 7 DIV, and generally increased in frequency as the cultures matured (Figure 3.2). Over early development, the rate of increase in spiking activity was greater in hippocampal cultures compared to cortical cultures. However, by maturity, there was no significant difference between the activity levels in the two types of networks (Figure 3.3A).

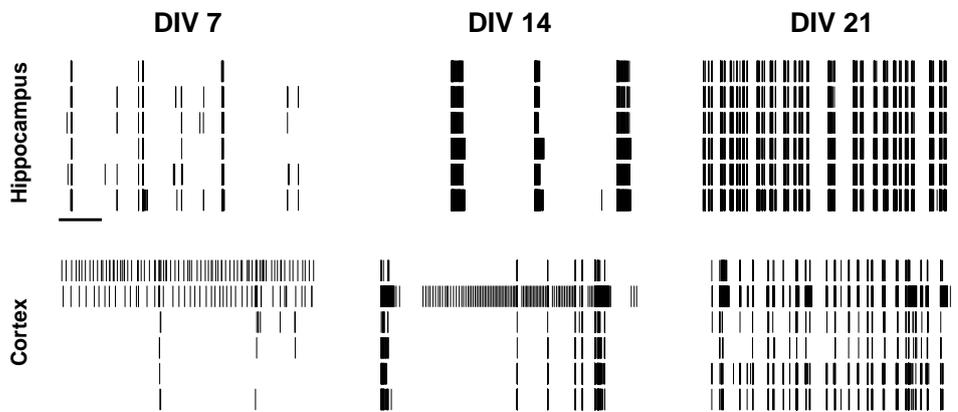


Figure 3.2: Examples of MEA recordings of spontaneous activity in one (top) hippocampal and (bottom) cortical culture at 7, 14 and 21 days in vitro (DIV). In each plot, each row is the spike train from one electrode. Six electrodes are shown from each MEA. The scale bar represents 10 s of recording time. Figure reproduced from Charlesworth et al. (2015).

Bursting

Based on the results from Chapter 2, the MaxInterval method was chosen to detect periods of bursting in these MEA recordings. Using this method, the rate of bursting activity in both network types was found to increase over early development before stabilising around 14–17 DIV (Figure 3.3C). Two properties of bursting activity differed significantly between the two types of cultured networks. Firstly, over the first 24 DIV, the proportion of spikes in bursts was significantly higher in hippocampal cultures compared to cortical cultures of an equivalent age (Figure 3.3E). Secondly, hippocampal networks exhibited more regularly spaced bursting activity than cortical networks, as indicated by the lower CV of IBI in these networks (Figure 3.3F). Weaker differences were observed between the two network types for the other features of bursting examined, namely the burst rate, duration and within burst firing rate, which were significant only at certain DIVs.

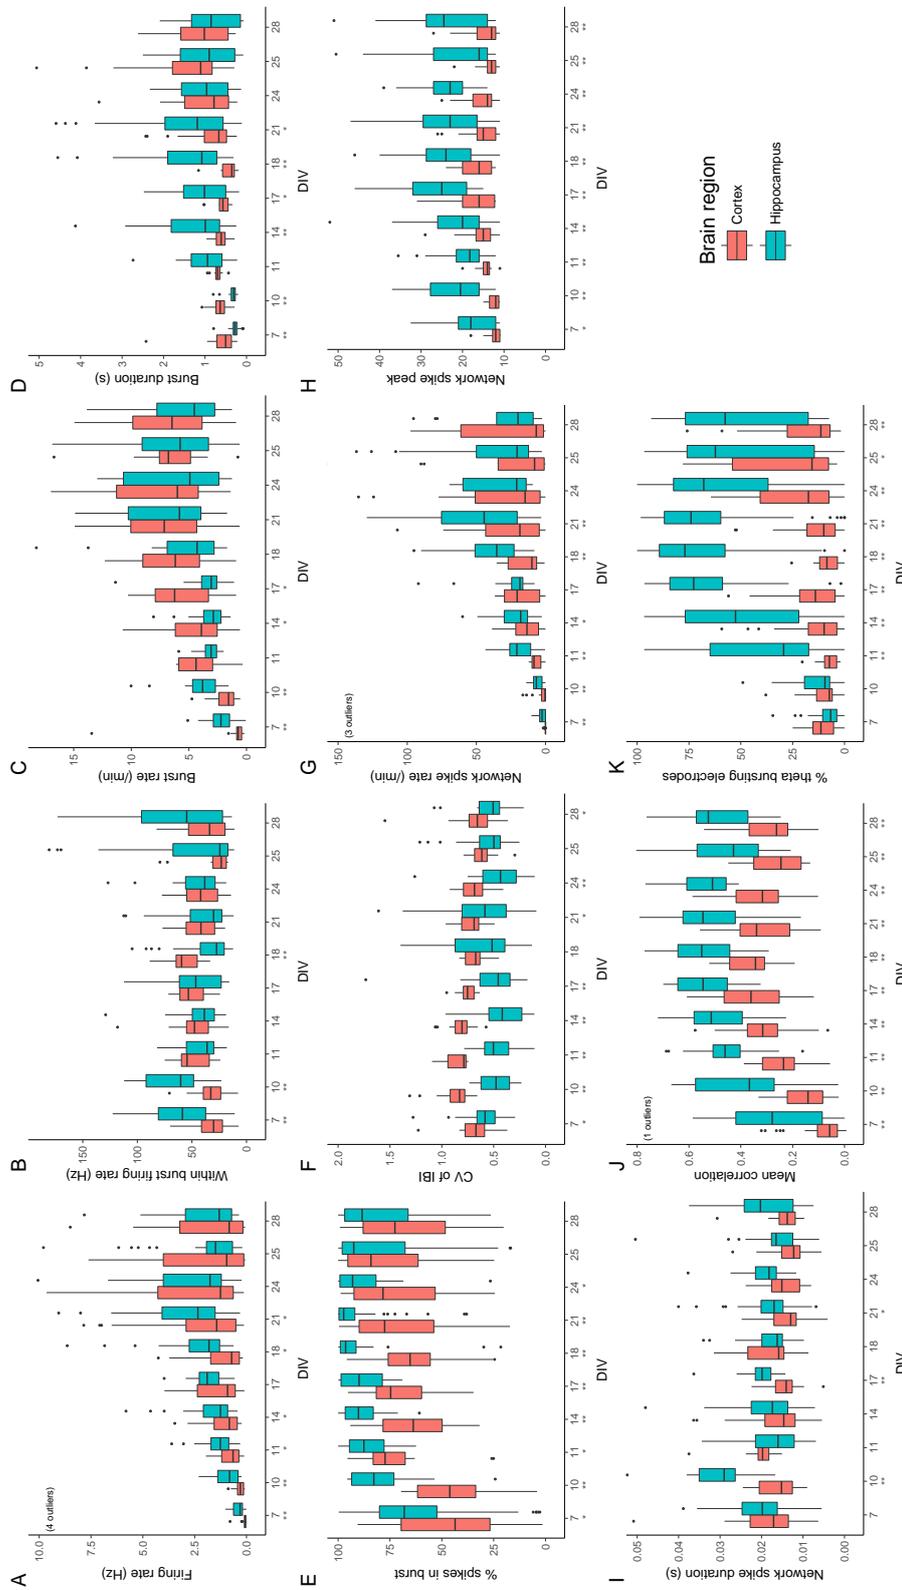


Figure 3.3: **A** Firing rate, **B** within burst firing rate, **C** burst rate, **D** burst duration, **E** percentage of spikes in bursts, **F** coefficient of variation of interburst intervals (CV of IBI), **G** network spike rate, **H** network spike peak, **I** network spike duration, **J** mean correlation, **K** percentage of electrodes exhibiting theta bursting in hippocampal and cortical network cultures. Box plots are of median values from each MEA. Stars underneath each age represent significant difference between the median hippocampal and cortical values at (*) 0.05 and (**) 0.01 significance levels. p values were obtained from a Mann-Whitney test, adjusted for multiple comparisons using the Benjamini & Hochberg procedure. Figure based on Figure 3 from Charlesworth et al. (2015).

Correlated network activity

Two types of features were used to measure the degree of correlated network activity in the cultured networks. Firstly, the presence of network spikes (Eytan and Marom, 2006), which are defined as events in which a number of neurons are simultaneously active, was examined. Network spike activity is thought to represent the activation of neuronal assemblies, which are densely connected groups of neurons believed to be the basic functional unit in the mammalian nervous system (Buzsáki, 2010). *In vivo*, neuronal assemblies have been associated with various physiological processes, including learning and memory encoding and recall (Holtmaat and Caroni, 2016).

In these recordings, network spikes were generally more prevalent in hippocampal networks compared to cortical cultures of the same age (Figure 3.3G). Hippocampal cultures also exhibited larger network spikes than cortical networks, indicated by the value of the network spike peak, which was significantly higher in hippocampal networks over the entire developmental period (Figure 3.3H). Network spikes also tended to be longer in hippocampal networks, however, this difference was not consistent across the entire period studied (Figure 3.3I).

As another measure of correlated network activity, the correlation coefficient between each distinct pair of electrodes on an array was calculated using the spike time tiling coefficient (Cutts and Eglén, 2014). Initially, it was investigated if there was any relationship between the distance separating electrodes and their pairwise correlation. Previous studies in the retina have observed an exponentially decaying relationship between correlation values and the distance between neurons (Wong et al., 1993; Cutts and Eglén, 2014), however no significant distance dependence between correlation values was found in both the hippocampal and cortical network recordings (examples in Figure 3.4).

The mean pairwise correlation of electrodes on an array was instead used as a measure of the overall level of correlation of network activity. Both types of networks exhibited an increase in their level of correlated activity over early development, before stabilising around 17 DIV. However, mean correlation in hippocampal cultures was found to be consistently stronger than in cortical networks over the entire developmental period (Figure 3.3J). Together, the results from the features used to measure network correlation suggest that hippocampal cultures tend to exhibit more highly correlated activity and contain larger networks of simultaneously active neurons compared to cortical cultures of equivalent ages.

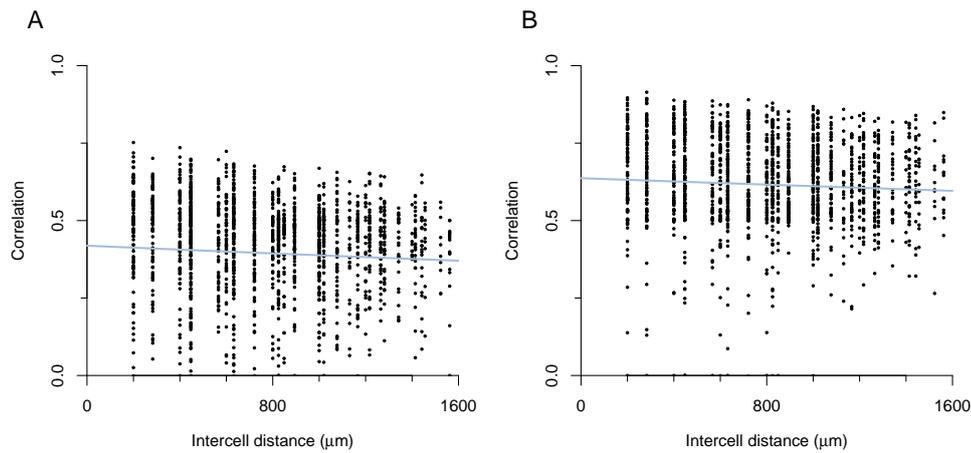


Figure 3.4: Pairwise electrode correlation values from MEA recordings from one **A** cortical and **B** hippocampal culture at 21 DIV. Blue lines show an exponential fit to the data.

Theta bursting

Theta bursting is a phenomenon involving regular oscillations of a neuronal network in a frequency range of 4 to 10Hz. Theta bursting has been observed in various mammalian neuronal networks both *in vivo* and *in vitro* (Bocian and Konopacki, 2004; Cappaert et al., 2009; Konopacki et al., 2006) and has been associated with numerous behaviours, including spatial navigation (Caplan et al., 2003), REM sleep (Vanderwolf, 1969), locomotor activity (Buzsáki, 2002) and learning (Leondopulos et al., 2012).

Here, theta bursting was identified by examining the log-adjusted ISI histogram of a spike train. An electrode was classified as theta bursting if there was a peak in this histogram in the range of 0.1–0.25 s. The level of theta bursting in the network was then measured as the fraction of all electrodes exhibiting theta bursting. Using this definition, a significant level of theta bursting was detected in hippocampal networks from DIV 11 onwards, however, the level of theta bursting in cortical networks remained low across the entire developmental period (Figure 3.3K).

Differentiation of network types

From Figure 3.3, it is evident that there are significant differences between the activity of hippocampal and cortical cultured networks across a range of features. In particular, hippocampal networks exhibit more prevalent and more regular bursting activity, higher levels of correlated network activity and an increased presence of theta bursting compared to cortical networks of equivalent developmental age. To assess which of the features best discriminate between hippocampal and cortical networks, and the degree to which the recordings could be correctly classified by their network types, I employed two common classification techniques, outlined in Section 3.2.2.

Initially, bagged classification trees were used to assign the recordings at each developmental age to either the hippocampal or cortical network type. Using this approach, I found that the recordings could be assigned to their correct network type with above 75% accuracy at each DIV. I then investigated which features were the most dominant in driving the classification using the measure of the mean decrease in Gini Index. The Gini Index is defined as

$$G = \sum_{c=1}^C p_{mc}(1 - p_{mc}),$$

where p_{mc} is the fraction of training observations in the m th node that are from class c (James et al. (2014), p. 312). The mean decrease in Gini index associated with a feature is thus a measure of the degree of the contribution of the feature to the homogeneity of the nodes in the tree-based classification.

At each age, different features were found to dominate the classification, however some features were consistently important across several DIVs. By averaging the relative importance of each feature over the classifiers from the ten separate ages, three features were found to be particularly dominant (Figure 3.5). The most informative feature as identified by this method was the CV of IBI. As seen in Figure 3.3, CV of IBI values were consistently lower in hippocampal cultures across development, indicating that these networks tend to exhibit more regularly spaced bursts than cortical networks. The degree of theta bursting in a network was also an important feature. From 14 DIV, the majority of electrodes in hippocampal cultures recorded the presence of theta bursting, whereas generally less than half of the electrodes in cortical cultures detected theta bursting across the entire developmental period.

Mean correlation, which was significantly higher in hippocampal compared to cortical networks at all DIVs, was the third most informative feature for classification.

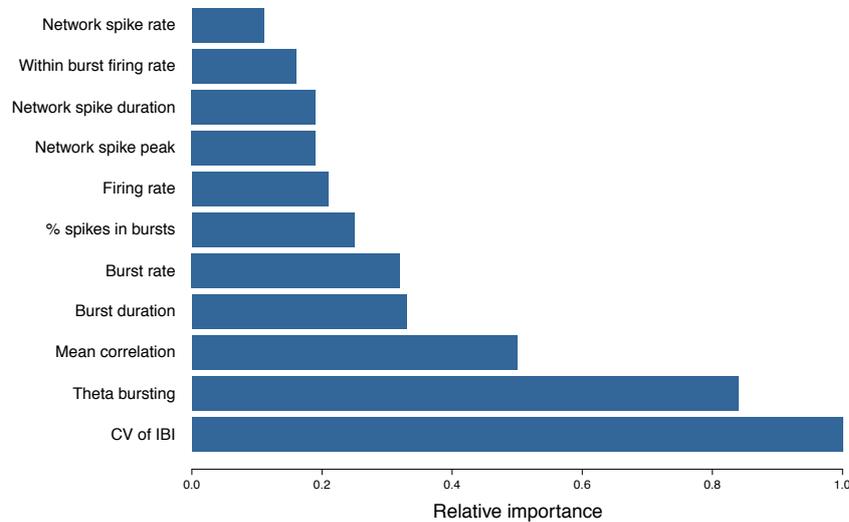


Figure 3.5: Importance of each feature, calculated using the mean decrease in Gini Index, relative to the top score.

This bagged classification tree approach was used because it is easily interpretable and allowed for the most informative features in the classification to be examined. Although bagged classification trees achieved reasonable levels of accuracy, I found that using SVMs to perform the binary classification generally outperformed the tree-based classification, achieving 81–96% classification accuracy over the ten ages studied (Table 3.2, bottom row). The lowest classification accuracies were apparent for the more mature networks in the study. At these ages, the values of many of the features plateau, after initially growing over development, and there are often less significant differences between the values from hippocampal and cortical networks.

The change in classification performance with reduced numbers of features was then examined by removing each feature successively from the feature set in order of their increasing importance and re-examining the accuracy of the SVM classifiers built using this reduced feature set (Table 3.2, columns 3–12). Classification accuracy generally remained high as the size of the feature set was gradually reduced. However, use of multiple features was essential to achieve high classification accuracy at most ages, with four features the minimum number required to achieve above 70% classification accuracy across all DIVs.

Feature	Impt.	Classification accuracy at given DIV									
		7	10	11	14	17	18	21	24	25	28
CV of IBI	1.00	65	87	94	89	82	75	64	68	71	67
Theta bursting	0.84	64	86	95	84	88	88	85	85	75	67
Mean correlation	0.50	67	90	94	91	85	93	88	87	75	78
Burst duration	0.33	73	88	96	92	91	89	87	81	73	79
Burst rate	0.32	77	86	98	91	93	96	89	93	77	82
% spikes in bursts	0.25	80	87	93	93	91	95	89	89	78	82
Firing rate	0.21	84	85	95	94	90	94	91	89	77	81
Network spike peak	0.19	82	86	96	94	91	96	91	90	77	82
Network spike duration	0.19	81	90	97	94	91	95	91	82	75	81
Within burst firing rate	0.16	84	87	97	92	91	96	93	88	80	79
Network spike rate	0.11	85	89	96	92	91	96	93	89	81	81

Table 3.2: Classification performance using a reduced number of features. Features are listed according to their decreasing importance, as judged by the mean decrease in Gini index (shown in Figure 3.5), shown relative to the top feature (Impt., column 2). The value of the classification accuracy in row $n = 1, \dots, 11$ is the mean percentage of correctly classified recordings using the top n features.

3.2.4 Discussion

In this section, I have described the development of a quantitative framework for analysing spontaneous activity in recordings of neuronal cultures on MEAs, and the use of this framework to analyse and compare the changes in activity patterns observed in hippocampal and cortical networks over the first 28 days of development. Using this framework, I identified quantitative differences in the spontaneous activity patterns of these two network types. Specifically, activity in hippocampal networks was shown to be more strongly correlated, with more regular bursting behaviour and a higher prevalence of theta bursting compared to cortical networks at corresponding stages of development.

I also found that the eleven quantitative features employed in this study contained sufficient information to discriminate between cortical and hippocampal networks of equivalent ages with reasonable levels of reliability. In particular, three features were found to dominate the classification, namely the CV of IBI, proportion of theta bursting and mean correlation, which were shown to differ significantly between the two network types at many DIVs.

The choice of simple quantitative features to describe the spontaneous activity in these cultures was deliberate, as it allowed me to examine whether these features were sufficient to differentiate between hippocampal and cortical network activity. The reasonably high classification accuracy achieved using this feature set suggests that there are significant differences between the two network types that can be captured using these simple features. More complex features may also exhibit significant differences between the network types. These could include measures of the functional connectivity of the network in different frequency bands, or the spatiotemporal features of periods of repetitive bursting (Ito et al., 2014; Wagenaar et al., 2006b). However, introducing unnecessary complexity into the feature set may lead to greater variance and increased computational time, the later of which is of particular concern in high-throughput contexts, or when analysing recordings from high-density MEAs.

One unknown factor in this study is the degree to which the *in vitro* activity in these recordings reflects *in vivo* neuronal activity in cortical and hippocampal brain regions. It may be the case that use of a different experimental protocol, such as brain slices rather than cultures of dissociated neurons, may generate recordings which more closely mimic *in vivo* spontaneous activity (Mao et al., 2001). However, some recent studies have suggested that spontaneous activity in cultured networks shows greater similarity to *in vivo* network activity compared to that from acute brain slices (Okamoto et al., 2014). The results from this study are not used to infer any differences between activity in hippocampal and cortical regions *in vivo*. Instead, the focus here is on providing a framework for analysing cultured neuronal network activity, which can also be applied to the study of *in vitro* networks under a range of experimental conditions.

3.3 Analysis of developmental spontaneous activity in cortical networks on multi-well MEAs

The work in this section was performed in collaboration with Diana Hall and Timothy Shafer from the US Environmental Protection Agency (EPA). However, all analysis, figures and text presented here are my own work.

3.3.1 Introduction

As evidenced in Section 3.2, MEAs provide a useful tool for recording and studying the simultaneous electrophysiological activity of networks of cultured neurons. Although single-well MEAs have been used extensively for the analysis of spontaneous and evoked neuronal network activity *in vitro* (Chiappalone et al., 2008; Van Pelt et al., 2005; MacLaren and Charlesworth, 2011), they offer limited throughput, making them inadequate for applications such as large-scale assessment of the effect of pharmacological or genetic manipulations on neuronal network activity.

Recently, multi-well MEAs (mwMEAs) have been developed. These consist of several (6–96) independent wells, with each well generally containing a lower number (3–64) of electrodes compared to single-well MEAs. The cultures on each well behave independently, but can be recorded from simultaneously, allowing for a number of replicates or different experimental conditions to be studied concurrently. The increased throughput provided by these mwMEAs has expanded their use in applications such as large scale toxicity screening (Nicolas et al., 2014; Valdivia et al., 2014), and for studying the impact of various disease states on neuronal network activity (Woodard et al., 2014; Wainger et al., 2014). As well as improving throughput, the use of mwMEAs also has the potential to decrease the variability observed in MEA recordings by increasing the number of recordings that can be made from a single culture (Johnstone et al., 2010).

An area of particular concern to the EPA is the assessment of the effect of potentially neurotoxic compounds on early neuronal network development, given that such chemicals have been implicated in several neurodevelopmental disorders (Grandjean and Landrigan, 2014). mwMEAs have the potential to facilitate high-throughput screening of large numbers of compounds and their impact on developmental network

activity. However, before mwMEAs can be employed for such purposes, the ontogeny of cultured network activity on such arrays and the ability of the functional features of this activity to be reliably recorded by the lower number of electrodes present in mwMEAs must be examined. In this study, analysis of recordings of rat cortical neuronal networks cultured on mwMEAs were used to address these questions, as well as to examine the optimal experimental protocol for using mwMEAs to study multiple experimental conditions simultaneously.

3.3.2 Methods

MEA recordings

The data set used in this study consisted of recordings of mixed primary cortical cultures from newborn rat cortex. The neurons were cultured and recorded on 48-well MEA plates with 16 electrodes per well. These cultures contained both excitatory and inhibitory neurons and glia. Recordings of at least 15 minutes were taken from each plate on DIV 5, 7, 9 and 12 and spikes were detected using a threshold of 8 times the root mean square of noise on each electrode. All experimental procedures were performed by Kathleen Wallace and William R. Mundy, and further experimental details are available from Cotterill et al. (2016a).

Data Analysis

Sixteen MEA plates produced from 15 primary cortical cultures were recorded from at DIV 5, 7, 9 and 12. A recording from one of the plates at DIV 12 was missing from the data set, and one recording at DIV 7 was excluded from the analysis, because its mean firing rate (1.2 Hz) was above 2 SD higher than the mean firing rate of other DIV 7 arrays (mean = 0.6 Hz, SD = 0.3 Hz). This resulted in 62 array recordings, composed of 2976 individual well recordings, being included in the analysis. Because many of these arrays were used at later DIVs for other experimental procedures, short recording times of 15–30 min were used to minimise the impact of repeated removal of cells from the incubator during development. As there was some variation in the length of each recording, to ensure consistency, only the last 15 minutes of each recording was analysed.

Twelve features were used to quantify the activity patterns in each well recording. Ten of the features were equivalent to Features 1-10 from Section 3.2.2, however, as each well on these mwMEAs represent independent networks, the feature values were calculated for each well separately, rather than an array-wide value being found. The feature relating to the presence of theta bursting (Feature 11 in Section 3.2.2) was excluded from this study, as the proportion of theta bursting remains low in cortical networks over development (Section 3.2.3). A low level of theta bursting was again apparent in this cortical data set. Based on discussions with our collaborators at the EPA, two additional features were included in the analysis of this data set. These were:

11. *Fraction of bursting electrodes* - An electrode was classified as bursting if the burst rate over the recording period equalled or exceeded one burst per minute. The well value was calculated as the number of bursting electrodes as a fraction of the total number of active electrodes in the well.
12. *CV of within burst ISIs* - The electrode value was calculated as the ratio of the standard deviation to the mean of all within burst ISIs on the electrode. The well value was the median value from all electrodes in a well, excluding those that recorded no bursting activity.

Again, the MaxInterval method was utilised to detect bursts in these recordings, using equivalent parameters to those in Section 3.2.2 (Table 3.1). Due to the smaller number of electrodes present in each well of the mwMEAs, for these recordings network spikes were defined as any period in which five or more electrodes were active within a 3 ms time bin.

Classification

A similar approach to Section 3.2.2 was used to classify the well recordings using the twelve quantitative features in this study. As all of the cultures were composed of the same type of neurons, the recordings were classified by their developmental age rather than their region of origin. First, bagged classification trees were built and interrogated to investigate the importance of each of the twelve features for classification.

Next, SVMs were used to examine the classification accuracy using reduced subsets of features. Since this classification problem involved classifying each well recording as one of four ages (5, 7, 9 or 12 DIV), a ‘one-against-one’ approach was employed. To implement this, six SVMs, one for each pairwise combination of ages, were built, each of which performed a binary classification of each recording into one of two ages. The class to which a recording was most frequently assigned across these six binary classifications was taken as its correct class.

Radial kernels were used for the SVMs because it was found that they had superior performance at classifying this data set compared to linear or polynomial kernels. The optimal parameter values for the radial kernel ($\gamma = \frac{1}{10}$, $C = 10$) were found using a grid search across a range of parameter values and 10-fold cross validation on the entire data set. For both the bagged classification tree and SVM approaches, two thirds of the well recordings were used as training data, and the remainder as a test set to assess the accuracy of the classification. The average accuracy across 100 repetitions with random choices of training and test sets is presented.

Due to the small number of electrodes in each well, bursting and network spike activity was absent in a number of well recordings, particularly at early stages of development. For any electrodes that recorded no bursting activity, the within burst firing rate and burst duration were set to zero. Similarly, for all wells that recorded no network spike activity, zero values were used for the features of network spike peak and duration. Wells for which the other features, CV of IBI, CV of within burst ISI or mean correlation, had null values were excluded from the data set used for classification. After these adjustments, the data set employed for classification consisted of 2606 well recordings, which represents 87.6% of the original number of recordings.

3.3.3 Results

Developmental profile

Spontaneous spiking activity was observed in the majority of recordings from 5 DIV. As development progressed, the overall spike rate and degree of correlation of activity across the networks tended to increase in most well recordings (example in Figure 3.6).

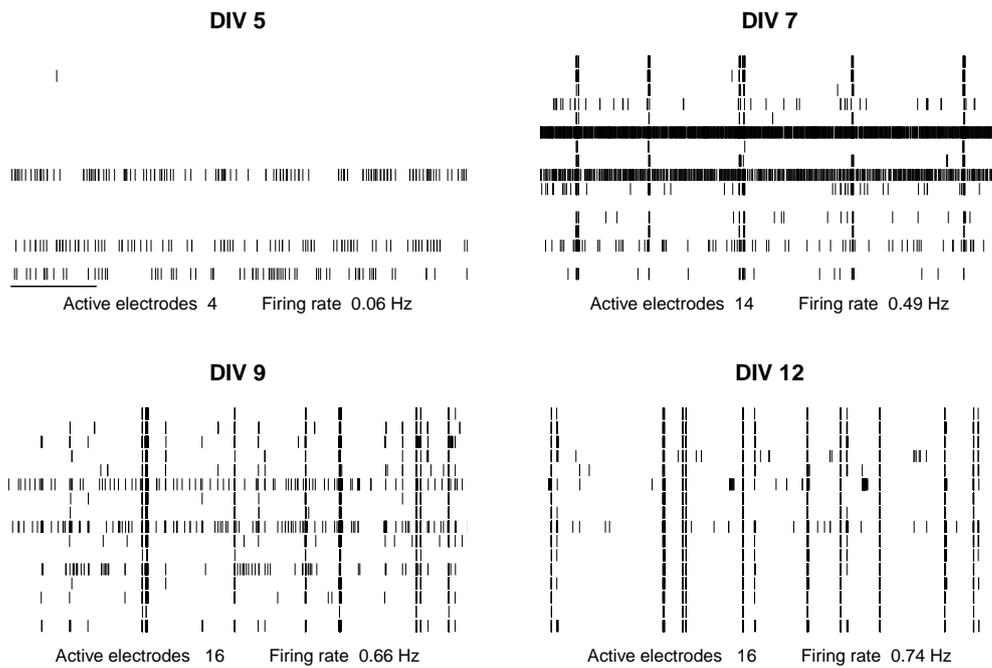


Figure 3.6: Raster plots of 5 min of activity recorded from the same well at each DIV. Each row represents the activity recorded one electrode. The scale bar represents 60 s.

Similar to the activity observed in single well MEA recordings in Section 3.2.3, the average firing rate in these cortical cultures increased monotonically over development (Figure 3.7A). The prevalence of bursting, as measured by the burst rate, fraction of bursting electrodes and proportion of spikes in bursts, also increased with increasing developmental age (Figure 3.7B, D, F). The firing rate within bursts also increased over development, while burst duration did not exhibit strong developmentally related changes (Figure 3.7C, E). This is similar to the relationships observed between culture age and bursting properties over early developmental time periods in my previous analysis of cortical cultures in Section 3.2.3. The regularity of within and between burst intervals both decreased over time, as indicated by the increasing CV of within burst ISIs and IBIs (Figure 3.7G, H).

Similar to the cortical cultures in Section 3.2.3, both the rate of network spikes and mean pairwise correlation between electrodes exhibited strong growth over development (Figure 3.7I, L). The duration and peak number of electrodes involved in network spikes also increased over very early development, but tended to plateau by 7 DIV (Figure 3.7J, K).

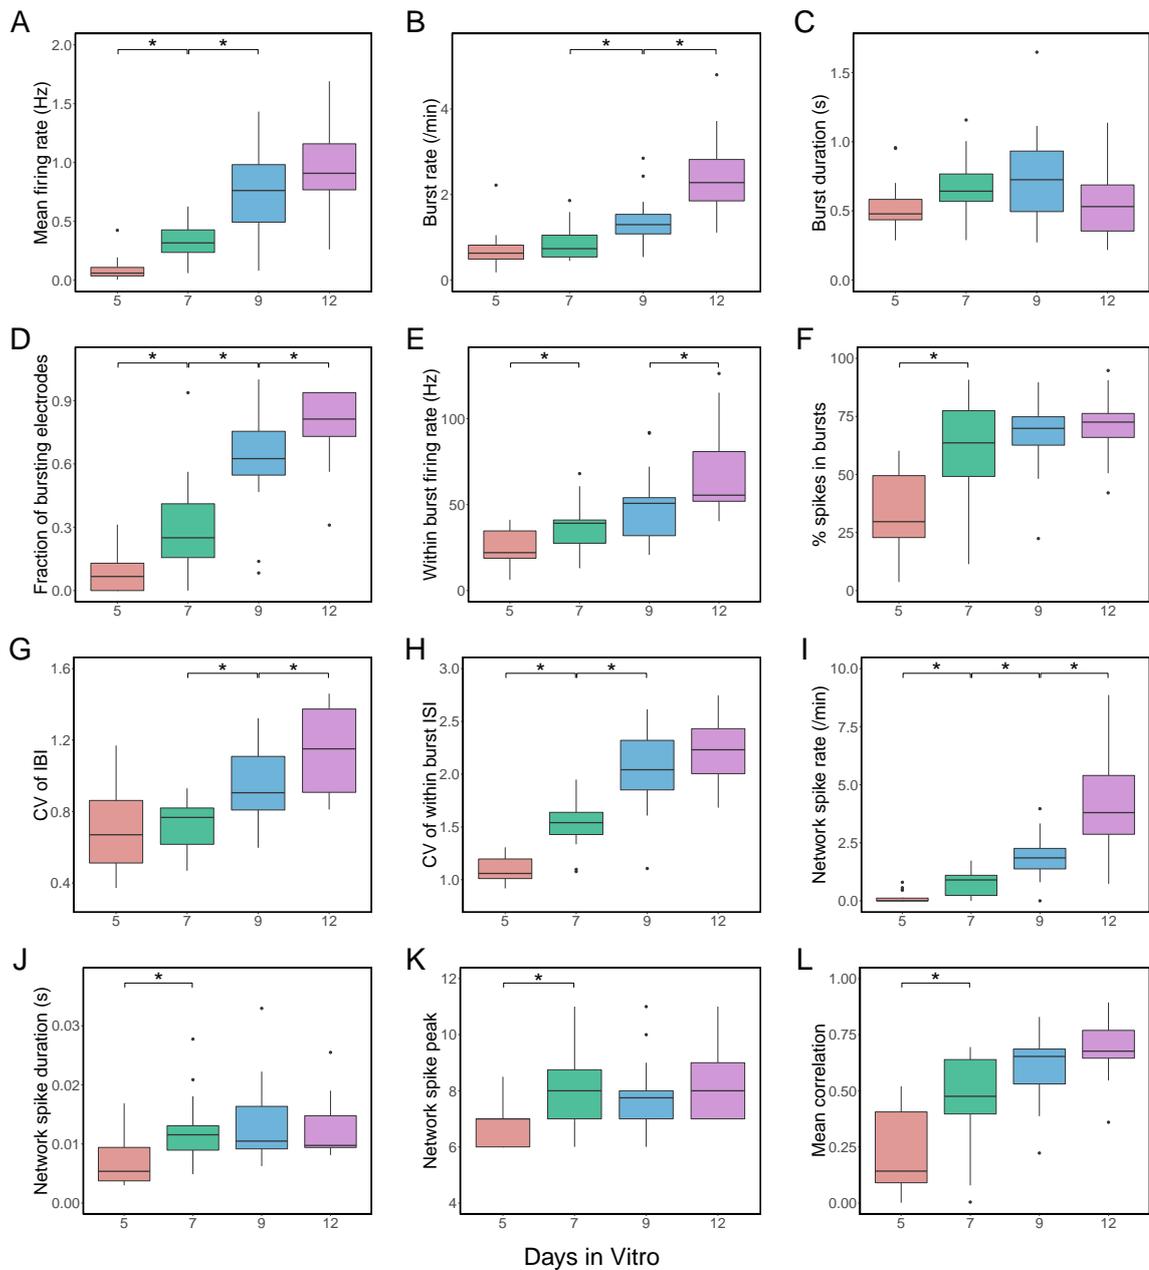


Figure 3.7: Feature values of cortical network recordings over development. **A** Mean firing rate, **B** burst rate, **C** burst duration, **D** fraction of bursting electrodes, **E** within burst firing rate, **F** percentage of spikes in bursts, **G** CV of IBI, **H** CV of within burst ISI, **I** network spike rate, **J** network spike duration, **K** network spike peak, and **L** mean pairwise correlation. Box plots show the median and interquartile range for the $n = 16$ plates at each DIV. Stars show significant differences between median values at consecutive ages (5 vs 7, 7 vs 9, 9 vs 12) at $p < 0.05$. Significance calculated using a Mann-Whitney test, with correction for multiple comparisons using the Benjamini & Hochberg procedure. Figure based on figure from Cotterill et al. (2016a).

Classification by developmental age

The classification approach outlined in Section 3.3.2 was used to assess the extent to which the twelve features in this study could differentiate between recordings of different network types. Since this data set contained only control recordings from cortical networks, developmental age was used as a proxy for network type. Two classification approaches, bagged classification trees and SVMs, were used to examine the accuracy with which the recordings from each well could be classified by their age using the twelve features from Section 3.3.2.

The bagged classification tree approach was able to differentiate between recordings of the four developmental ages with approximately 72% accuracy, which greatly exceeds the 25% classification accuracy that could be expected by random chance. Using the measure of the mean decrease in Gini index, the two features related to the regularity of bursting structure, namely the CV of within burst ISIs and CV of IBI, were found to be the most informative for classifying the cultures according to their developmental age (Table 3.3, column 2). The values of these two features were shown to grow monotonically over development in the recordings (Figure 3.7G, H).

A slightly higher classification accuracy of $\sim 73\%$ was achieved using SVMs and all twelve of the features for classification. Using the ranking of feature importance obtained from the bagged classification trees, the variation in the classification accuracy of SVMs built using only a subset of the twelve features was examined. The classification accuracy generally gradually declined as the number of features included in the feature set was reduced (Table 3.3, column 3). Using just one feature, namely the CV of within burst ISI, well recordings could be differentiated by their developmental age with almost 50% accuracy, while four features was the minimum number required to achieve a classification accuracy of at least 65%.

Next, the accuracy of SVMs at performing a binary classification of the well recordings from each pair of ages was examined. This was achieved by building six individual SVM classifiers, one for each pairwise combination of ages (5 vs 7, 5 vs 9, etc.). Table 3.4 shows the average classification accuracy attained for each pair of ages using all twelve features, as well as subsets of the feature set.

Feature removed	Importance	Accuracy (%)
CV of within burst ISI	1.00	49.2
CV of IBI	0.70	58.3
Network spike rate	0.50	62.0
Burst rate	0.49	65.0
Burst duration	0.44	66.0
% Spikes in bursts	0.39	68.3
Within burst firing rate	0.36	69.5
Firing rate	0.35	71.4
Correlation	0.31	72.7
Bursting electrodes	0.22	73.0
Network spike duration	0.18	73.5
Network spike peak	0.09	73.4

Table 3.3: Classifier performance at predicting age of arrays from the four ages used in the study. Features are listed in decreasing order of importance (score in column 2, normalised to the top score), and the value in each row $n = 1, \dots, 12$ is the mean percentage of correctly classified recordings using the top n features. Table reproduced from Cotterill et al. (2016a).

The classification accuracy was the highest for pairs of recordings with the greatest difference in their ages. For example, the classifier could differentiate between DIV 5 and DIV 12 recordings with over 96% accuracy using all twelve features. Even when only one feature, CV of within burst ISI, was used in the classification, these wells could be separated by age with almost 92% accuracy.

The classification of arrays with small differences in age, particularly those at the later end of the developmental period studied, had the lowest accuracy. The accuracy of the SVM classifier at distinguishing between DIV 9 and DIV 12 well recordings was approximately 82% using all twelve features, and decreased to just above chance when only one feature was used for classification.

Feature	Classification accuracy %					
	5 vs 7	5 vs 9	5 vs 12	7 vs 9	7 vs 12	9 vs 12
CV of within burst ISI	75.0	87.5	91.8	69.9	78.1	57.4
CV of IBI	77.4	89.5	93.6	76.8	85.5	64.7
Network spike rate	79.3	90.3	95.5	79.6	88.0	68.7
Burst rate	79.3	90.3	95.3	81.0	88.4	72.8
Burst duration	79.7	90.8	95.3	81.0	88.6	74.0
% Spikes in bursts	81.6	91.3	95.6	81.5	90.4	76.4
Correlation	82.2	92.1	95.6	82.3	90.9	77.1
Firing rate	82.3	91.7	95.7	82.3	90.9	80.2
Within burst firing rate	84.2	92.7	96.2	82.4	91.3	81.2
Bursting electrodes	83.6	92.5	96.2	82.6	91.8	81.5
Network spike duration	83.7	92.9	96.8	82.9	92.8	82.0
Network spike peak	82.2	92.5	96.8	82.6	93.0	82.1

Table 3.4: Classifier performance at predicting the age of arrays for each pairwise combination of ages. Features are listed in decreasing order of importance, and the value in each row $n = 1, \dots, 12$ is the mean percentage of correct classifications using the top n features, as described in Table 3.3. Table reproduced from Cotterill et al. (2016a).

Classification with reduced numbers of wells

In the recordings analysed in this study, all 48 wells on each array were used as replicates for the control experimental condition. However, mwMEAs can also be employed for high throughput screening by applying different experimental conditions to the individual wells across the array. When using this method, the optimal approach for assigning different experimental conditions to the wells on an array must be considered, taking into consideration the trade-off between the number of experimental conditions that can be tested simultaneously, and the inevitable reduction in power that stems from using a reduced number of replicates of each condition.

To examine the robustness of the classification to variations in replicate number, I simulated an experiment in which only a fraction of wells on each plate are used for each condition by repeating the SVM classification from section 3.3.3 using only the data from $n < 48$ randomly selected wells from each array. This was repeated for a range of values of n , and the mean classification accuracy across 100 repetitions with random choices of wells is shown in Figure 3.8.

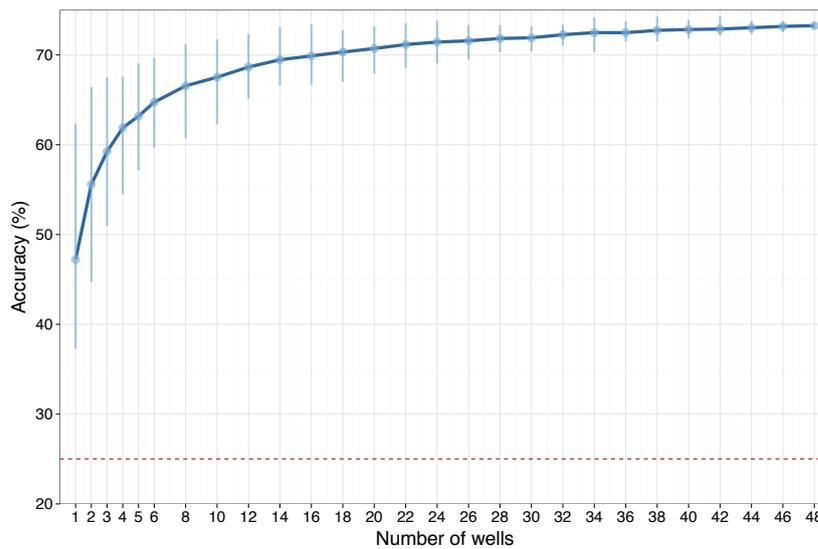


Figure 3.8: Accuracy of classification of wells by their developmental age, using $n \leq 48$ wells per plate. Points show mean classification accuracy and vertical bars show the variability across 100 trials with random choices of wells. Dotted red line shows the baseline accuracy expected by chance.

The classification accuracy remained relatively stable as the number of wells used in the analysis was reduced to a third of the original value (16 wells, $\sim 70\%$ accuracy). When data from eight wells per plate were included in the classification, the mean accuracy remained above 65%, albeit with higher variance. This suggests that for the particular problem of classifying cultures according to their four ages in this data set, $\frac{1}{6}$ to $\frac{1}{3}$ of each mwMEA could be used for each separate condition without a large reduction in the reliability of the results.

Variability across cultures

By allowing for increased numbers of recordings to be made from a single culture, the use of mwMEAs also has the potential to reduce the variability observed across multiple replicate MEA recordings (Johnstone et al., 2010). To assess the degree to which this approach could reduce variability, the major sources of variability in cultured mwMEA recordings was next investigated. In particular, I examined whether there was a significant effect of the culture-of-origin of each network on their observed activity patterns.

In the data set examined in the previous sections, 15 of the 16 arrays were produced from different cultures. A separate data set was thus employed for this analysis, which consisted of recordings from 21 mwMEAs produced from 7 cultures. Each

culture was used to produce three of the arrays, and recordings were made from each array at DIV 5, 7, 9 and 12 DIV. Two recordings were missing from the data set, one from culture four at DIV 5 and another from culture seven at DIV 12. Hence only two recordings from arrays produced from these cultures were available on these days. The mwMEAs in this data set consisted of both control and dosed wells, however, only the six control wells from each plate were used in this analysis.

The inter- and intra-culture variability of the features in Section 3.3.2 were examined by comparing the values of these features between the three plates produced from the same cultures, as well as between the groups of plates produced from each of the seven cultures (see Figure 3.9). Figure 3.10 shows the inter and intra-culture variability of mean firing rate across the recordings at each DIV. The points show the mean firing rate from each of the six control wells on each plate. A plate value was then calculated as the median of these values from all of the wells on a plate. The box plots in Figure 3.10 show the median and interquartile range (IQR) of the plate values from each culture (each culture consisting of 2–3 plate recordings, depending on recording age). It can be seen that the variability of firing rate values between different cultures is generally far greater than between networks produced from the same culture.

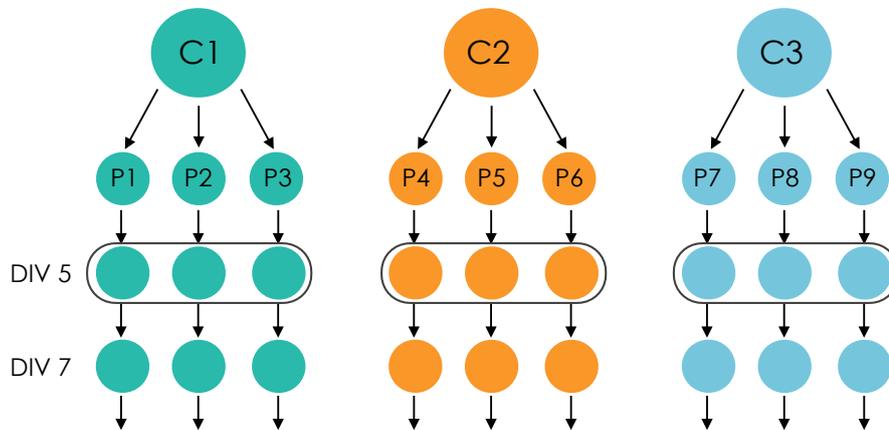


Figure 3.9: Schematic of the relationship between cultures and plates used for analysis of inter- and intra-culture variability. Three plates (P) are produced from each culture (C) and recorded from at each DIV. Intra-culture variability was measured between the plates encircled in the same ring (represented by the box plots in Figures 3.10, 3.11). Inter-culture variability was measured between the groups of ringed plates (represented by the horizontal lines in Figures 3.10, 3.11). This was repeated at each DIV. Only three of the seven cultures used in this analysis are shown in this schematic.

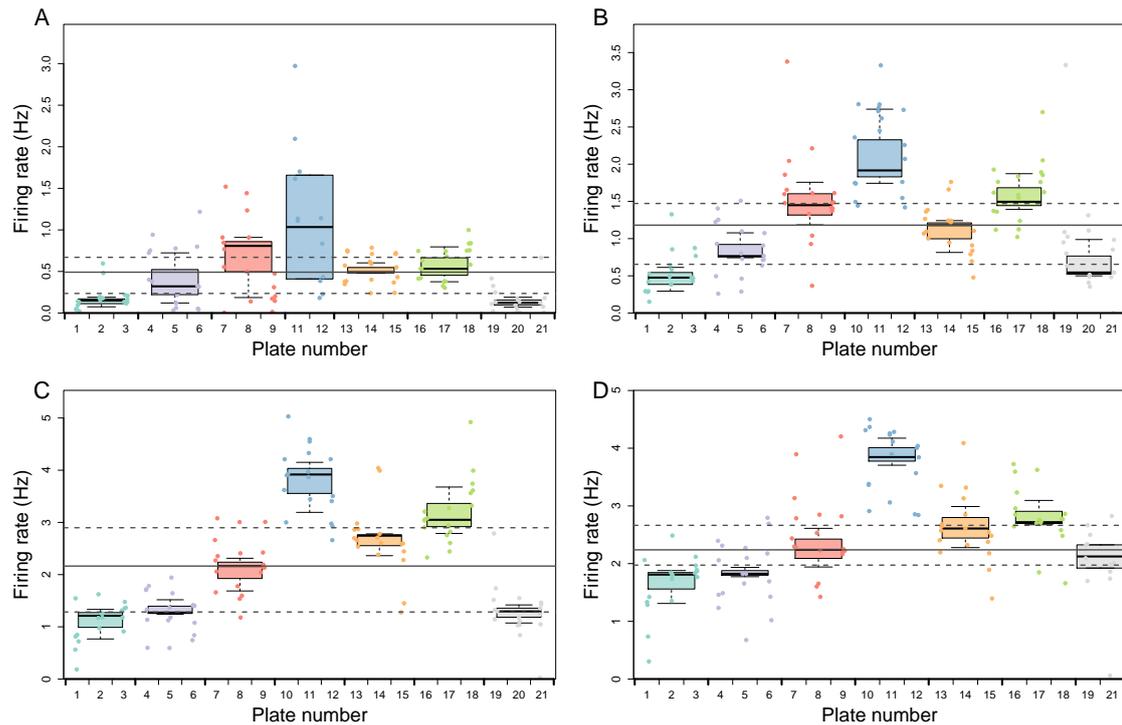


Figure 3.10: Firing rate per well at **A** 5 DIV, **B** 7 DIV, **C** 9 DIV and **D** 12 DIV. Dots show the value of the firing rate on the control wells from each plate. Box plots show the range of median plate values per culture. Horizontal lines show the median and IQR of the median values from all cultures at each DIV.

In this case, culture four appeared to produce the most active networks. At DIV 7–12, recordings from all of the wells from plates produced from culture four (plates 10–12) exhibit higher mean firing rates than almost any well produced from cultures one, two or seven (plates 1–6, 19–21) (Figure 3.10). At DIVs above 5, the IQR across the cultures (dotted lines in Figure 3.10) is greater than the IQR of values from any individual culture.

A similar relationship between inter- and intra-cluster variability was observed for a number of other features. The variability of the network spike rate shown in Figure 3.11 suggests that in this data set, correlated network activity develops at earlier ages in networks produced from certain cultures compared to others. At 5 DIV, only plates produced from three cultures (cultures 4–6, consisting of plates 11–18) exhibited network spiking activity. Network spikes were observed on every plate produced from these three cultures, and not on any plates produced from other cultures. By DIV 7, network spikes are present on all plates, however, network spikes generally remain more prevalent on plates from cultures 4–6 compared to the other cultures over the entire developmental period.

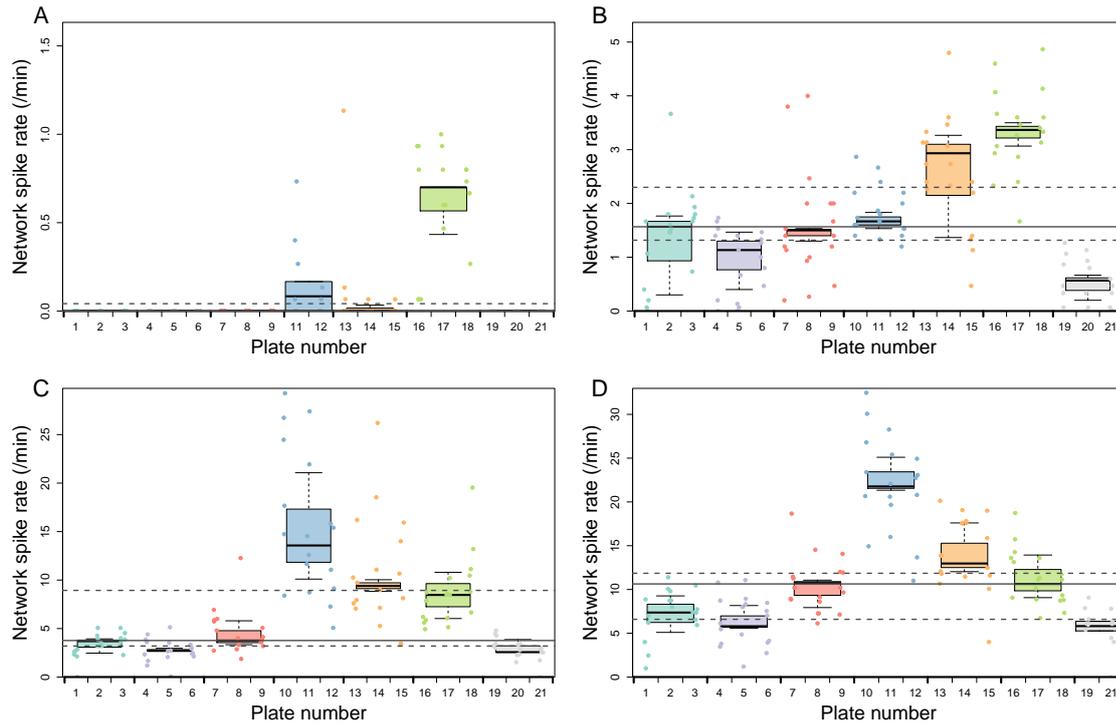


Figure 3.11: Network spike rate per well at **A** 5 DIV, **B** 7 DIV, **C** 9 DIV and **D** 12 DIV. Formatted as per Figure 3.10.

3.3.4 Discussion

In this section, I have presented a quantitative analysis of the ontogeny of spontaneous activity patterns in neuronal networks cultured on 48-well mwMEAs, which showed a rapid increase in overall spiking, bursting and synchronous activity in these networks over the first 12 DIV. The developmental profile of this activity is qualitatively similar to activity observed in single-well MEAs, both from our own experimental studies (Charlesworth et al. (2015), Section 3.2.3) and those of others (Kamioka et al., 1996; Chiappalone et al., 2006; Robinette et al., 2011; Van Pelt et al., 2005). These studies describe a progression of spontaneous activity from random infrequent spiking to more temporally structured and correlated activity, characterised by bursts of synchronised network activity, over the first 4 weeks *in vitro*.

The emergence of spiking activity by DIV 5 in the recordings studied here is consistent with several previous studies of rat cortical cultures, which have detected spiking activity beginning from 3–6 DIV (Wagenaar et al., 2006a; Kamioka et al., 1996). At DIV 5, a low level of synchronous network activity was observed in the recordings, as evidenced by the very low rate of network spikes and low mean pairwise correlation between electrodes. This can be physiologically attributed to the relative

lack of connectivity between neurons at DIV 5 in this culture model, for which synaptogenesis has been shown to begin after approximately 6 DIV and continue to around DIV 12 (Harrill et al., 2011). An increase in the values of the features relating to network connectivity was observed from DIV 7, reflecting this change in the underlying morphology of the network.

Despite exhibiting qualitatively similar developmental trends, spontaneous activity developed more rapidly in these network recordings than that observed in some previous studies of rat cortical neurons recorded on MEAs (Hogberg et al., 2011; Bakkum et al., 2008; Chiappalone et al., 2006). This can be attributed to several factors, including the later age of cells when initially cultured (postnatal day 0 compared to embryonic day 18) and the higher plating density used in this study compared to others (150 000 vs 50 000 cells/well), which has been shown to have an effect on the speed of maturation of cultured networks (Wagenaar et al., 2006a; Chiappalone et al., 2006; Valdivia et al., 2014; Biffi et al., 2013). This earlier development of spontaneous activity in the cultured networks may in fact be beneficial in some applications such as chemical screening, offering increased throughput and reduced assay times (Biffi et al., 2013).

As well as providing a characterisation of the ontogeny of spontaneous activity patterns in cortical networks recorded on mwMEAs, this study also has implications for applications such as drug development or neurotoxicity screening of compounds using MEAs. In the past, researchers have commonly used only one feature, mean firing rate, to describe the effect of different compounds on neuronal network activity (Nicolas et al., 2014; Johnstone et al., 2010). I have shown that the accuracy of the classification of the recordings in this study by developmental age was improved from $\sim 49\%$ using just one feature to 65% when four features were used, or above 70% when a feature set of eight or more features was used. This suggests that using an increased number of features may allow chemical-induced effects on neuronal network activity to be detected with greater sensitivity, and facilitate compound ‘fingerprinting’ approaches (Mack et al., 2014).

I have also shown that network type can be reliably determined using data from a fraction of wells from each mwMEA, suggesting that a single array can be used to record from networks under multiple experimental conditions. In this case, reducing the size of the data set to between $\frac{1}{6}$ and $\frac{1}{3}$ of the number of wells from each array did not greatly decrease the average reliability with which networks could be classified

by age using SVMs, however there was some increase in the variability of results. In other experimental studies, the optimal number of conditions per mwMEA will vary depending on the the gross level of difference between the activity patterns in each of the network types being studied. However, the approach developed in Section 3.3.3 can be applied to any data set to determine the appropriate number of wells per condition that should be employed by experimentalists to ensure the reliability of results. As far as I am aware, this study is the first to address the issue of how to most appropriately use the wells on a mwMEA for a variety of experimental conditions.

Further, by analysing an additional mwMEA data set, I have shown that culture-to-culture variability in network activity tends to be greater than the variability between plates produced from the same culture. This suggests that when using mwMEAs to record from networks under different experimental conditions, maximising the number of wells produced from the same culture may be advantageous compared to using replicates from different cultures.

In summary, as part of this collaborative study, I have shown that the ontogeny of spontaneous activity in neuronal networks cultured on mwMEAs is qualitatively equivalent to network activity in single-well MEAs. Using a number of quantitative features, this activity can be characterised, and different network states reliably differentiated using either data from all or only a fraction of wells from each array. Together, these results suggest that mwMEAs are a suitable tool for studying the effects of chemical or genetic manipulations on neuronal network activity in a high-throughput manner.

3.4 Conclusions

Recordings of the electrophysiological activity of neurons on MEAs have been used extensively to study network activity in a variety of contexts. In this chapter, I have described the development of a quantitative framework for analysing MEA recordings. By applying it to the analysis of two separate data sets, I have shown that the simple features in this framework are sufficient to characterise the network activity recorded on both single-well and multi-well MEAs, as well as to classify recordings according to their network types. In both studies, features related to the regularity of interburst or interspike intervals were shown to be important in

differentiating between network types, suggesting that these features of spontaneous activity patterns may play an important role in neuronal network development.

Using the framework developed in this chapter, I described the ontogeny of spontaneous activity in both cortical and hippocampal networks *in vitro* and showed that the activity patterns in these networks can be reliably differentiated using a small number of quantitative features. I also demonstrated the suitability of multi-well MEAs for the high-throughput study of the activity of networks under multiple experimental conditions, and proposed an approach for determining the most appropriate allocation of conditions to wells in these circumstances. Although in this chapter, only recordings of spontaneous activity in untreated neuronal networks were examined, the frameworks presented here could also be applied to a wide range of applications, such as analysing recordings of stimulated network activity, studying models of disease or predicting the neurotoxicity of compounds (Brown et al., 2016; Jimbo et al., 1999; Chong et al., 2011).

Chapter 4

Multi-view neuronal cell type classification

4.1 Summary

Comprehensive classification of the numerous categories of neuronal cell types is an essential step in understanding and accessing neuronal networks. Previous classification studies have generally employed only one type of feature, such as the anatomical, electrophysiological, molecular or genetic properties of neurons, to identify their underlying cell types. This fails to take advantage of experimental advances which allow for electrophysiological recordings and morphological reconstructions to be obtained from the same cells (Armañanzas and Ascoli, 2015).

In this chapter, I utilise two neuronal data sets to examine the potential benefits of using both the electrophysiological and morphological characterisations of neurons for their classification. In particular, I consider the use of multi-view clustering techniques that are specifically designed to find the shared cluster structure in two or more ‘views’ or feature sets obtained from the same data. I first explore the relationship between the electrophysiological and anatomical properties of the cells in the two multi-view neuronal data sets in this study, before examining the performance of single and multi-view techniques at clustering the data. Several multi-view techniques are shown to outperform hierarchical clustering applied to either an individual view or the concatenated views from these data sets, as judged by standard evaluation measures. The high performing clustering techniques are then used to investigate the cell types present in each neuronal data set. As far as I am aware, this is the first study to use multi-view clustering techniques for the purpose of classifying neuronal cell types.

4.2 Introduction

Classification of neuronal cell types has recently become a major focus of modern neuroscience, as evident from its inclusion as one of the key goals of the BRAIN initiative (BRAIN Working Group, 2014), as well as a priority in the European Human Brain Project. Completing a census of neuronal cell types is important for several reasons. Firstly, identifying the numerous cell types present in a neuronal network is an essential first step to understanding the connectivity of the network, and how each cell type contributes to overall network function (Hosp et al., 2014). From a developmental perspective, being able to identify a cell type before it reaches its final form would also allow researchers to study the functional and structural developmental processes of different neurons (Sanes and Masland, 2015). Additionally, classification of cell types is important for differentiating the effect of certain diseases on specific cell types (Della Santina et al., 2013), as well as in the development of therapeutics (Lagali et al., 2008).

Although there remains lacking a consensus about what constitutes a specific neuronal ‘cell type’ and the level of granularity of cells to which this term should be applied, it is commonly thought that the definition should be multimodal, taking into account the morphological, physiological and molecular features of cells (Ascoli et al., 2008; Seung and Sümbül, 2014; Armañanzas and Ascoli, 2015). However, most studies that have worked on the problem of classifying neuronal cell types have used only a single type of property of neurons for their classification (Karagiannis et al., 2009; Li et al., 2012; Helmstaedter et al., 2009; Druckmann et al., 2013; Cauli et al., 1997; Sümbül et al., 2014; Helmstaedter et al., 2013; Tasic et al., 2016).

Recent experimental advances have increased the prevalence of data sets consisting of multiple modalities of neuronal features, such as electrophysiological and morphological characterisations of the same cells. Several such data sets have been used for the classification of neuronal cell types. However, most studies using such multi-view data sets for neuronal classification have taken the simplistic approach of performing single-view clustering on a concatenation of the feature sets to identify the underlying cell types (Marx and Feldmeyer, 2013; Helmstaedter et al., 2009; McGarry et al., 2010; Hosp et al., 2014; Karagiannis et al., 2009). This approach fails to take into account the unique statistical properties of the different representations as well as the correlations between data modalities.

Data sets consisting of multiple representations or ‘views’ are also common outside of neuroscience, such as images with text captions or web pages represented by their text and in/out links. A number of machine learning techniques have been specifically designed to find the shared cluster structure in such multi-view data sets. These algorithms take a variety of approaches, including using weighted combinations of kernels, joint graphical representations or shared latent subspaces for clustering (Tzortzis and Likas, 2012; de Sa, 2005; Guo, 2013). Multi-view clustering methods have been employed extensively in applications such as image, document and web page clustering (Gao et al., 2005; Rege et al., 2007; Kolda and Bader, 2006; Liu et al., 2009; Cai et al., 2011), however they have rarely been used for the analysis of neuroscientific data.

In this chapter, I conduct a thorough literature search to identify a number of multi-view clustering techniques that can be applied to the classification of data sets consisting of the morphological and electrophysiological features of neurons. These techniques are then used to investigate the cell types present in two such multi-view neuronal data sets. Several multi-view methods are shown to outperform single-view clustering techniques on these data sets, and are used to characterise the cell types present in the data. As well as its relevance for neuronal cell type classification, results from this study may also have implications for the analysis of other biological multi-view data sets, such as joint microRNA and mRNA data (Qin, 2008), genomic and transcriptomic data from tumour samples (Ali et al., 2014) and joint diffusion tensor imaging and fMRI recordings (Zhu et al., 2014).

4.2.1 Previous work in neuronal classification

Classification of neuronal cells has been of interest to neuroscientists for more than a century, since the early work of Golgi and Cajal revealed a diversity in the structure of neurons (De Carlos and Borrell, 2007). A common approach to classifying neuronal cell types since this time has been to manually classify cells based on features of their size, shape and activity, and standardised terminology for describing different cell types has been developed (Ascoli et al., 2008). These types of manual approaches have been used to identify numerous cell types, however much of the diversity of the nervous system remains uncharacterised. Although the number of neuronal cell types is unknown, by extrapolating from classification studies in small brain regions, it has been estimated that the number of distinct cell types in the cortex alone may exceed one thousand (Stevens, 1998), making manual approaches to neuronal classification at this scale untenable.

Most previous studies that have used computational approaches for cell type classification have generally utilised one single type of property of neurons. Commonly these are either electrophysiological (Druckmann et al., 2013; Farrow et al., 2011; Zaitsev et al., 2012; Helm et al., 2013; Baden et al., 2013; Li et al., 2012; Suzuki and Bekkers, 2010), transcriptomic (Cahoy et al., 2004; Sugino et al., 2006; Zeisel et al., 2015; Tasic et al., 2016) or morphological (Karube et al., 2004; Guerra et al., 2011; Lu et al., 2015; Mihaljević et al., 2015) features. The most common approach to identifying the underlying cell types in these studies has been to use hierarchical clustering of the feature set (Karube et al., 2004; Cahoy et al., 2004; Sugino et al., 2006; Zaitsev et al., 2012; Helm et al., 2013; Li et al., 2012).

The electrophysiological features commonly used to classify neurons include both their passive properties, such as resting membrane potential and input resistance, as well as active cell properties, which include features of the spike trains and single spikes invoked by input injections of current (see Figure 4.1 for examples). Cluster analysis using these types of features has been successful at identifying electrophysiological classes of cells in a range of brain regions.

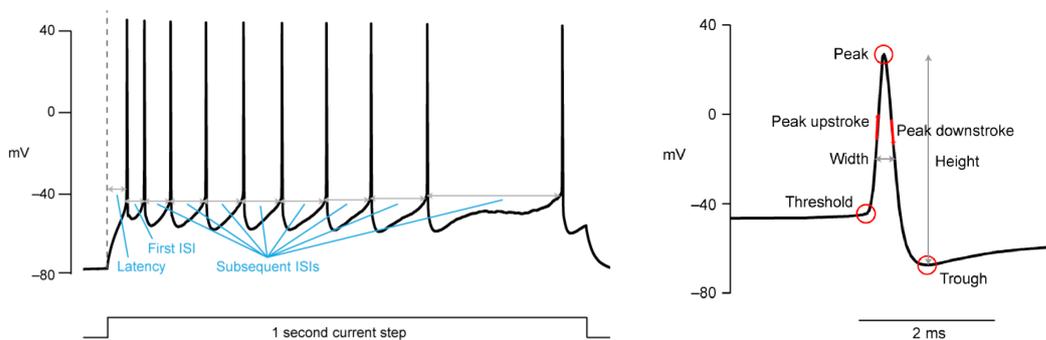


Figure 4.1: Examples of electrophysiological features of (left) spike trains and (right) single spikes commonly used for neuronal classification. Figures adapted from Allen Institute for Brain Science (2016).

Recent advances in single cell transcriptomic profiling have also facilitated neuronal classification studies using information from single-cell RNA-seq (Zeisel et al., 2015; Tasic et al., 2016). Such studies have identified a number of molecularly distinct classes of neuronal cells, however it is still unclear whether some transcriptomic differences between cells simply reflect the state or activity level of the cell, rather than being characteristic of a specific cell type (Tasic et al., 2016; Dehorter et al., 2015).

Morphological features have also commonly been used for classification of neurons. Generally morphological classification is performed using quantitative features extracted from the computational reconstruction of cells, such as the somatic surface area and the length and number of branch points of axons and dendrites. Software tools have been developed that have allowed for the calculation of a comprehensive catalogue of morphological parameters in a standardised fashion across a range of classification studies (Scorcioni et al., 2008) (see Figure 4.2 for examples).

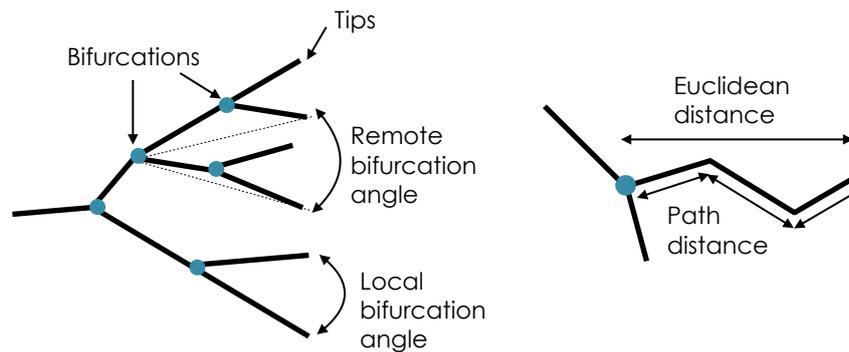


Figure 4.2: Example of morphological features of dendrites calculated using the standardised software, L-measure (Scorcioni et al., 2008).

Recently, alternative approaches to quantifying the morphological similarity of cells have been developed. NBLAST (Costa et al., 2016) quantises the branching patterns of neurons by assigning a coordinate and tangential vector to each small segment of the neuronal morphology. This allows the similarity of two neurons to be measured using a probabilistic score calculated between each pair of nearest neighbour segments of the neurons. Using this approach and a combination of affinity propagation and hierarchical clustering, Costa et al. (2016) identified 1052 clusters in a data set of ~ 16000 *Drosophila* neurons, several of which were previously identified cell types. Alternatively, Smbl et al. (2014) take the approach of representing each neuron according to its voxel-wise arbor density. The dissimilarity between cells is then measured as the Euclidean distance between the vector representations of their arbor densities. This method was utilised to identify 15 cell types in a dataset of ~ 300 mouse RGCs, many of which mapped to molecularly identified cell types.

Rather than taking a purely computational approach, other studies have combined both manual and machine learning methods to classify cells. DeFelipe et al. (2013) proposed an approach in which neurons are manually characterised by experts based on six pre-defined characteristics. The majority vote across the experts is then taken as the ‘ground truth’ label of each cell and used to train supervised algorithms to

classify cells based on their morphological features. López-Cruz et al. (2014) use the same manually classified data set to build a Bayesian multinnet representing the opinion across the group of experts. Inference on this multinnet is then used to find the consensus definition of different GABAergic interneuron types.

4.2.2 Neuronal classification using multiple feature types

A small number of recent studies have used data sets consisting of multiple types of features of neurons for cell type classification. Commonly, these data sets consist of the morphological and electrophysiological properties of cells. Santana et al. (2013) used affinity propagation to cluster a concatenated data set comprised of electrophysiological and morphological features of $n = 50$ GABAergic cortical neurons, previously identified as containing four neuronal cell types. This method was highly effective at distinguishing PV from SST-expressing neurons and achieved 78% accuracy at separating the four cell types in the data set. Dumitriu et al. (2007) used the morphological and electrophysiological features of $n = 18$ cells and hierarchical clustering to identify three interneuron types in the mouse visual cortex, which match molecularly defined cell types. In this case, they found that the morphological properties of axons were the most informative feature type for classification.

McGarry et al. (2010) utilised Ward's hierarchical clustering on a partially overlapping data set consisting of morphological and electrophysiological features from SST-expressing neocortical interneurons. This study found high levels of similarity between the clusters identified using morphological and electrophysiological features alone, with only one of $n = 16$ cells in the joint data set allocated to different groups when clustered using the two separate feature sets. In this case, clustering using the concatenated feature set produced less compact cluster structures than using the morphological or electrophysiological features alone. Conversely, Hosp et al. (2014) found that morphological features were more informative than electrophysiological features for the classification of a data set of $n = 22$ GABAergic neurons in the dentate gyrus, however, both types of features were required to achieve an unequivocal classification of the neurons into five well-separated classes. Other studies have also used hierarchical clustering on concatenated feature sets obtained from whole-cell patch clamp recordings and morphological reconstructions of biocytin stained cells to classify larger data sets of $n > 90$ L6b somatosensory barrel cortical neurons (Marx and Feldmeyer, 2013) and cells from layers 2–6 of the rat medial prefrontal cortex (Van Aerde and Feldmeyer, 2015).

There are also several studies that have used multi-view data sets consisting of different types of features of cells for neuronal classification. Perrenoud et al. (2012) used electrophysiological, morphological and molecular properties and hierarchical clustering to characterise $n = 42$ interneurons from mouse barrel cortex. Karagiannis et al. (2009) employed a combined data set consisting of laminar location, electrophysiological and molecular features of $n = 200$ neurons, as well as morphological features for a subset ($n = 68$) of these neurons. In this case, which studied NPY-expressing neocortical interneurons, the authors found that the morphological features were the least informative and molecular features were the most informative for classification. Subsets of this data have also been used in other classification studies (Battaglia et al., 2013; Pohlkamp et al., 2014).

Despite the diversity of these studies that have used multiple characterisations of neurons for their classification, a commonality is their use of single-view clustering techniques applied to either the individual or concatenated feature sets to identify the underlying cell types present. None of these studies have taken a multi-view approach to clustering of neuronal data sets.

4.2.3 Overview of multi-view clustering

The field of multi-view learning is a growing area of interest in modern machine learning, motivated by a rise in the number of real world applications producing data sets composed of multiple views (Sun, 2013). The term multi-view learning has been used to refer to a variety of situations in the literature. In this thesis, I use the term ‘multi-view’ to denote a data set consisting of two or more distinct feature sets that were produced from the same underlying samples. For example, for a set of video clips, features of the still visual frames could compose one view of the data while features of the audio track could be used as a second view (see Figure 4.3 for schematic of multi-view clustering).

Data sets consisting of multiple representations or ‘views’ arise commonly from real world data sources, such as marketing data from multiple databases, 3D objects captured from several camera angles and documents translated into multiple languages. These views can vary between being completely compatible to having a high level of disagreement, complicating the task of finding a common cluster structure between the views (Greene and Cunningham, 2009).

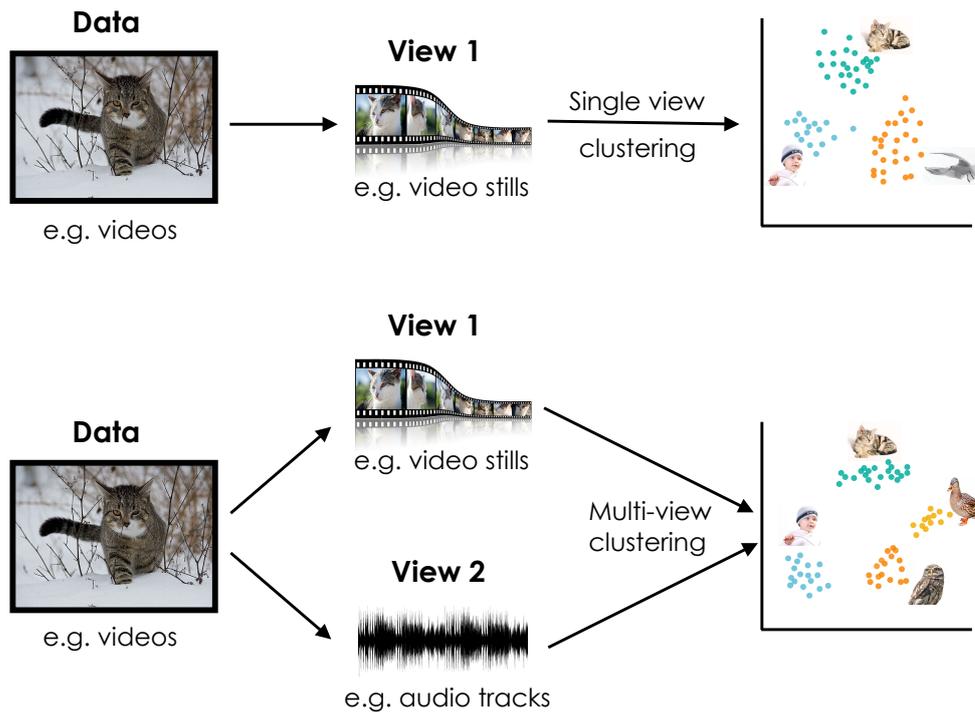


Figure 4.3: Schematic of (top) single-view and (bottom) multi-view clustering. Multi-view clustering techniques can produce different clustering solutions to single-view methods applied to the same data set.

A number of multi-view clustering techniques have been developed to deal with these kinds of representations. These methods take a variety of approaches, with the general aim of reconciling the discordance and maximising the agreement between the partitioning of the multiple views.

To understand the diversity of multi-view clustering techniques, I conducted an extensive literature search of multi-view clustering algorithms. In doing so, I identified four main approaches to clustering that multi-view approaches take, and classified each method into one of these categories. These four approaches are represented in Figure 4.4, and outlined below:

1. **Early integration:** Information from the multiple views is combined before clustering, often by combining kernels or through graphical approaches.
2. **Latent subspace:** A latent subspace is found using the multiple views of the data. Clustering is then performed within this latent subspace.
3. **Bootstrapping:** Information obtained in the process of clustering one view is used to inform clustering in the alternate views, often in an iterative manner.

4. **Post-clustering combination:** The multiple solutions obtained from single-view clustering of each of the individual views are combined post-clustering.

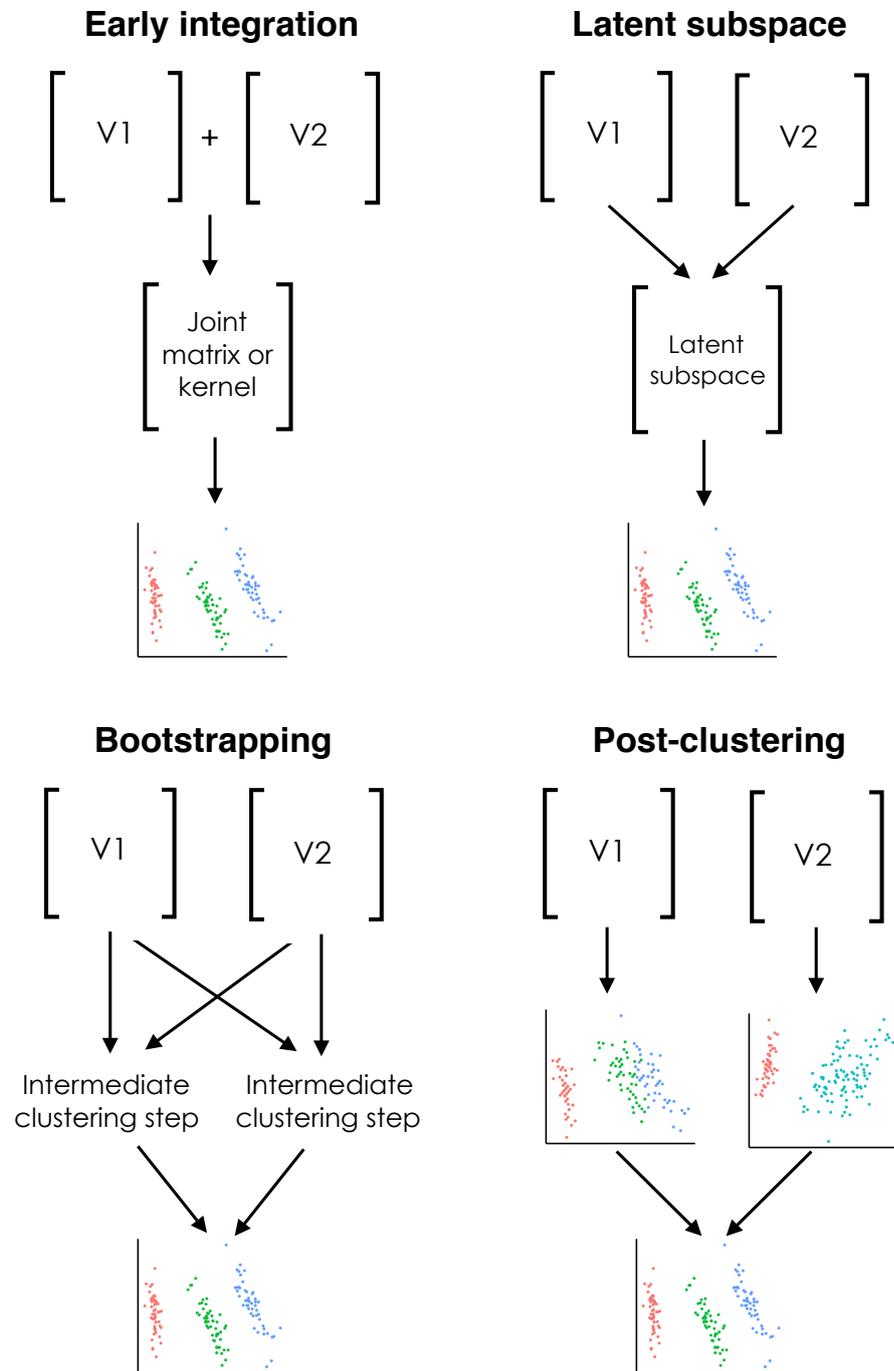


Figure 4.4: Representation of four approaches to multi-view clustering.

Early integration clustering methods use a variety of approaches to integrate information from the multiple views of the data prior to clustering. A common approach is combining the kernels of each view, generally with some weighting of the kernels and/or individual features (Tzortzis and Likas, 2012; Chen et al., 2013; Liu et al., 2009). Graph-based early integration clustering techniques have also been proposed. de Sa et al. (2010) developed a method in which a multipartite graph is constructed that connects nodes representing the same data points in the multiple views, while Zhou and Burges (2007) proposed combining transition probability matrices derived from each view using a Markov mixture. In both cases, spectral clustering techniques are employed to partition these combined graphs.

Latent subspace multi-view clustering methods perform clustering within a subspace found using information from the multiple views. These methods take a variety of approaches to estimate this subspace, including non-negative matrix factorization (Li et al., 2014), linked matrix factorization (Tang et al., 2009), tensor analysis (Liu et al., 2013) and low-rank and sparse decomposition (Xia et al., 2014). Guo (2013) formulate the problem of learning the optimal subspace as a min-max optimisation problem with a sparsity inducing norm, while Mo et al. (2013) extend the latent variable model for finding cancer subtypes originally developed by Shen et al. (2009) to allow it to be applied to variety of data types. This method uses a Monte Carlo algorithm to estimate the values of underlying latent variables, which are then partitioned using k -means clustering.

As an alternative latent subspace approach, Chaudhuri et al. (2009) proposed a method that employs Canonical Correlation Analysis (CCA) to find the subspaces in which the projections of the features from the multiple views are maximally correlated. Clustering is then performed on the projection of one of the views into its corresponding subspace. Numerous variations of clustering using CCA have also been proposed. These include a random non-linear extension of CCA (Lopez-Paz et al., 2014), and the correlational spectral method (Blaschko and Lampert, 2008), which performs k -means clustering in a subspace found by implementing kernel CCA on a training subset of the data, and assigns test points to the closest cluster centroid in the projected space.

The **bootstrapping** clustering techniques were defined as those that use information obtained in clustering of one of the views to influence clustering in the alternative views, generally in an iterative manner across the multiple views. A general expectation-maximisation algorithm for multi-view learning was developed

by Bickel and Scheffer (2004) and uses a maximum likelihood approach to iteratively update the model parameters in one view, taking into account the expected values of the hidden variables in the other view. Several other bootstrapping approaches adapt the spectral clustering algorithm into a multi-view setting. Kumar et al. (2011) and Cai et al. (2011) both developed co-regularization approaches to spectral clustering which utilise an objective function containing the Laplacians from the multiple views. Co-training approaches for two-view data sets have also been proposed by Kumar and Daumé (2011) and Zhao et al. (2014). Kumar and Daumé (2011) developed a method that uses information from spectral clustering in one view to alter the similarity matrix in the alternate view, while in Zhao et al. (2014) discriminant subspaces are learned iteratively in each view using cluster labels obtained from k -means clustering in the other view. Wang et al. (2013) take an alternative approach to clustering with two views, which involves using Pareto optimisation to find the normalized min-cut on two graphs which represent the two views of the data. This technique produces a set of Pareto optimal cuts that represent equally good partitions of the data, which can be consolidated into a single clustering solution using k -means clustering on the Pareto embedding of the data.

A final category of techniques for the clustering of multi-view data are the **post-clustering combination** methods. Many of these methods were not specifically developed for use in multi-view settings, and also have applications in other areas such as distributed clustering, clustering with multiple solutions and knowledge reuse (Li and Ding, 2008; Strehl and Ghosh, 2002). Li and Ding (2008) proposed a non-negative matrix factorisation-based approach to cluster combination, with optimal weights for each clustering solution found using L_1 norm regression. Greene and Cunningham (2009) also employ matrix factorisation to assign data points to one of k meta-clusters. Strehl and Ghosh (2002) proposed multiple post-clustering combination algorithms that combine clusters in such a way as to maximise the average normalised mutual information between the individual cluster solutions. Bruno and Marchand-Maillet (2009) employ Probabilistic Latent Semantic Analysis to combine multiple clustering solutions, while other algorithms use voting mechanisms to find consensus clusterings (Ayad and Kamel, 2008; Fred and Jain, 2005).

4.2.4 Aims of this chapter

As seen in the previous sections, a limited number of studies have utilised joint electrophysiological and morphological feature sets for neuronal cell type classification. Those studies that have used such multi-view feature sets have generally utilised data sets that are small in size and have taken simplistic approaches to classification, leading to conflicting conclusions about the usefulness of multi-view neuronal data sets for cell type classification. In particular, none of these studies have used multi-view clustering techniques for the classification of neuronal cell types. Given this, in this chapter I present a more comprehensive analysis of the use of multi-view neuronal data sets for cell type classification, with the aim of addressing the following questions:

1. Is using both the electrophysiological and morphological features of neurons advantageous for classification, compared to using only a single feature type?
2. Which approach to clustering using multi-view information achieves the greatest accuracy at classifying neuronal cell types? Do multi-view clustering techniques offer an improvement over single-view approaches in this context?
3. Can anything be learnt about the relationship between the electrophysiology and morphology of neurons from this study?

To address these questions, I examined the performance of single-view and several multi-view clustering techniques at identifying the cell types present in two existing data sets consisting of the morphological and electrophysiological features of neurons.

4.3 Methods

4.3.1 Multi-view clustering methods

Twelve methods that are representative of the four approaches to multi-view clustering outlined in Section 4.2.3 were chosen for use in this study. Since the focus of this work was to provide advice and useable tools for the analysis of neuronal data, only methods for which open-source code was available or for which code was provided by the authors upon request were included. A brief description of each of the methods is given here. Although some methods are applicable to data sets with more than two views, here I limit the description to their application to two-view datasets.

Early integration methods

Kernel-based weighted multi-view clustering (Tzortzis and Likas, 2012)

This method performs single-view kernel-based clustering on the weighted sum of the kernels from each individual view, defined as

$$K_{sum} = \sum_{n=1}^N w_n^p K_n,$$

where K_n is the kernel from view n , w_n is the weight of the n th kernel, and p is a fixed sparsity inducing parameter, chosen a priori. Tzortzis and Likas (2012) propose two options for the clustering step, which are kernel k -means and spectral clustering. For each approach, the algorithm iterates between clustering using K_{sum} and updating the kernel weights, w_n , using a closed form update rule that minimises intra-cluster variance. These two steps are repeated until convergence.

TW- k -means (Chen et al., 2013)

This method is an extension of the standard k -means algorithm, with additional steps to calculate the optimal weights for both the views and the individual features in each view. The view weights are chosen according to the compactness of the cluster structures in each view. Within each view, the individual feature weights reflect the importance of each variable in the clustering, and are calculated in an equivalent fashion to a single-view variable weighting method. The view and feature weights are incorporated into the objective function for k -means clustering, and the weights and clusters are iteratively updated until convergence.

Spectral Minimizing-Disagreement algorithm (sMD) (de Sa et al., 2010)

This method constructs a multipartite graph in which each node represents a data point in one view. This graph contains edges only between nodes representing the same data point in the two views, which are weighted to reflect the neighbourhood relationships within each of the views using the product of the Gaussian kernels of the views. The affinity matrix of this graph is then used as input to the single-view spectral clustering algorithm. The final clusters are found using k -means clustering on either the singular values corresponding to the more reliable of the two views, or the average of the singular values from both views.

Latent subspace methods

Linear Canonical Correlation Analysis (LCCA) (Chaudhuri et al., 2009)

This method finds a set of basis vectors for each view, such that the correlation between the views projected onto their corresponding set of basis vectors is mutually maximised. The projection of one of the views into the space spanned by these vectors is then used as input to a single-view clustering algorithm. In the implementation of linear CCA in this study, the partition around medoids (PAM) approach was used to find the final clusters (Kaufman and Rousseeuw, 1990).

Randomised non-linear CCA (RCCA) (Lopez-Paz et al., 2014)

This method involves performing linear CCA on randomised non-linear mappings of the two views. For each view, consisting of feature vectors (x_1, \dots, x_m) , the i components of the random non-linear mapping are calculated as the vectors

$$z_i = [\cos(w_i^T x_1 + b_i), \dots, \cos(w_i^T x_m + b_i)],$$

where w_i are randomly sampled weights $w_i \sim \frac{1}{\sigma_n} \mathcal{N}(0, 1)$, σ_n is the Gaussian kernel width of view n and b_i are random uniformly distributed bias vectors. The linear CCA approach described above is then applied to the mapped feature sets to partition the data.

Partial multi-view clustering (Li et al., 2014)

The partial multi-view clustering algorithm was developed for the clustering of data sets in which information is missing in one or more views, but can also be applied to multi-view data sets with no missing information. This method finds a subspace in which there is a small distance between data points representing the same example in the two different views, and a good grouping of similar data points. This is achieved by using non-negative matrix factorization to estimate the common latent subspace and latent representation of the data in an iteratively updating procedure. k -means clustering is then performed on the latent representation to obtain the final clustering solution.

Robust multi-view spectral clustering (Xia et al., 2014)

For each view, a transition probability matrix is constructed from the Gaussian kernel of the feature matrix. Each of these matrices is assumed to consist of the sum of a shared low-rank latent transition probability matrix, and a sparse error matrix unique to each view. Low-rank and sparse decomposition are used to estimate the

joint transition probability matrix between the views. Clusters are then found using the Markov chain method for spectral clustering (Zhou et al., 2005) with this matrix as input.

iCluster+ (Mo et al., 2013)

The iCluster+ algorithm was originally developed to identify breast cancer subtypes using multi-view genomic data sets. This method involves finding a common set of $k - 1$ latent variables for the views that span a lower dimensional joint subspace. For continuous data, a standard linear regression of the features with respect to the latent variables is performed, with sparsity introduced using the lasso penalty (Tibshirani, 2007). A Monte Carlo Newton-Raphson algorithm is used to maximise the penalised log-likelihood in this regression and find the expected values of the latent variables. k -means clustering is then used to divide the data set into k clusters using the values of the $k - 1$ latent variables.

Bootstrapping methods

Co-regularized spectral clustering (Kumar et al., 2011)

This method modifies single view spectral clustering to include the Laplacians from each view in the clustering objective function, and regularises on the eigenvalues of both Laplacians to enforce consistency in the cluster structures in the two views. Two methods are proposed for co-regularisation:

1. Pairwise co-regularisation: This method minimises the pairwise disagreement between the clustering solutions in the two views. This is achieved by replacing the graph Laplacian, L_v , in the spectral clustering objective of view v ,

$$\max_{U_v} \text{tr}(U_v^T L_v U_v) \quad \text{s.t.} \quad U_v^T U_v = I, \quad (4.1)$$

with $L_v + \lambda U_{v'} U_{v'}^T$, where U_v is the matrix of the eigenvectors of L_v , $U_{v'}$ is the equivalent matrix for the alternate view v' , and λ is a hyperparameter. This is solved by alternatively updating U_v and $U_{v'}$ until convergence. The U from the most important view chosen a priori, or the average of U_v and $U_{v'}$ is then used as input to the k -means algorithm for clustering.

2. Centroid-based co-regularization: In this method, the set of eigenvectors in each view are regularized towards a common consensus set of eigenvectors. This is equivalent to replacing the graph Laplacian, L_v , in the spectral clustering objective of view v (Equation 4.1) with $L_v + \lambda_v U^* U^{*T}$, where U^* is a matrix

consisting of a consensus set of eigenvectors and the λ_v s are hyperparameters, set separately for each view according to its level of noise. The algorithm then proceeds by alternatively updating U_v in each view, followed by optimising U^* , which is repeated until convergence. Clusters are then found using the final U^* matrix as input to the k -means clustering algorithm.

Co-training for spectral clustering (Kumar and Daumé, 2011)

For each view, an initial similarity matrix is calculated as the Gaussian kernel of the feature matrix. Spectral clustering is performed on each similarity matrix to find the discriminative eigenvectors in each view. A new similarity matrix for each view is then calculated by projecting the original similarity matrix onto the top $k \times p$ eigenvectors from the alternate view, and then back-projecting into the original space, where k , the number of clusters, and $p \geq 1$ are constants set a priori. This is repeated for a fixed number of iterations. Clusters are then found using the k -means algorithm, utilising as input the top k eigenvectors of the final similarity matrix from either the most informative of the two views, chosen a priori, or a concatenation of the final similarity matrices from the two views.

Post-clustering combination

These algorithms take as input clustering solutions obtained from single-view clustering in each view individually, and find an optimal combined clustering solution. In this study, both k -means and hierarchical clustering were used to obtain the cluster solutions from the individual views for input to the post-clustering combination algorithms, and the results from both are presented.

Cluster ensembles (Strehl and Ghosh, 2002)

This method seeks the combined clustering solution that has the maximum average normalised mutual information with each of the individual clustering solutions. Three algorithms are proposed to find this solution:

1. Cluster-based Similarity Partitioning Algorithm (CPSA) : This method calculates a binary similarity matrix for each clustering solution, in which points in the same cluster have a similarity of one, and zero otherwise. Re-clustering is then performed on the average of the similarity matrices from all individual clustering solutions.
2. HyperGraph-Partitioning Algorithm (HGPA): This method represents the clustering solutions as a hypergraph, in which each node represents a data

point and each equally-weighted hyperedge represents a cluster in one view. The final clustering solution is chosen as the partition of this graph into k components which cuts the minimal number of hyperedges.

3. Meta-CLustering Algorithm (MCLA): As in the HGPA method, in this approach each cluster from each view is represented as a hyperedge in a hypergraph. Graph-based clustering of the hyperedges is then performed to find k meta-clusters. The set of hyperedges in each meta-cluster is collapsed into a single meta-hyperedge. Each data point is then assigned to the meta-cluster with which it is most strongly associated.

The cluster ensemble method applies each of these three algorithms to the clusters found in the individual views, and outputs the combined clustering solution from the algorithm that achieves the highest average normalised mutual information with the single-view clustering solutions.

Probabilistic Latent Clustering (PLC) (Bruno and Marchand-Maillet, 2009)

This method employs Probabilistic Latent Semantic Analysis (PLSA), which was originally developed for the analysis of document-term matrices. The clustering solution from each view is represented as a $k \times N$ binary matrix. The joint cluster-data point probability is then given by the row-wise concatenation of these matrices, with appropriate normalisation. This is used to derive the joint cluster-cluster probability, $P(c_j, c_{j'})$, defined as the fraction of data points allocated to clusters c_j and $c_{j'}$ in the clustering solutions from the two views. It is assumed that there exist k binary latent variables, z_i , in the process of generating this joint cluster-cluster probability, and clustering of each data point, d , is achieved by assigning to it the latent variable that maximises the posterior probability $P(z_i|d)$.

A preliminary investigation of each of the multi-view clustering techniques was first performed by using each method to cluster well-established labelled data sets from fields outside of neuroscience (see Appendix A.1). This was used to gain a greater understanding of each of the methods, and their performance and scalability when applied to data sets with a variety of properties. The focus of this study was then on the application of these techniques to the clustering of multi-view neuronal data sets.

4.3.2 Neuronal data sets

An ideal approach for assessing each of the multi-view clustering techniques for the purpose of neuronal cell type classification would be to evaluate their performance at clustering a variety of labelled multi-view neuronal data sets. However, as the precise identity of all neuronal cell types is still unknown, multi-view neuronal datasets with ‘ground truth’ labelled cell types are not available. However, there do exist experimental data sets consisting of multiple types of features of unlabelled neurons. These generally contain some basic information about the cell types present, such as their layer of origin or a broad classification of the cells into two or more basic types. This information can be used as a proxy for ‘ground truth’ labels, to provide some indication of the quality of the clustering solutions produced by each clustering technique. In this study, two unlabelled multi-view data sets consisting of the morphological and electrophysiological features of neurons were used to analyse the performance of the clustering techniques at classifying neuronal cell types. These two data sets are detailed below.

Rat medial prefrontal cortex data set

The first neuronal data set studied consisted of features of cells from rat medial prefrontal cortex (mPFC), collected and previously published by Van Aerde and Feldmeyer (2015). Slices were obtained from the prefrontal cortex of P24–46 rats, and select pyramidal cells from these slices were chosen for study based on their morphology and pattern of firing. *In vitro* whole-cell patch clamp recordings were utilised to measure the electrophysiological characteristics of cells, using hyperpolarizations and stepped depolarizations. Cells were then filled with biocytin and imaged, and their three-dimensional morphology was computationally reconstructed using NeuroLucida (MBF Bioscience) (examples in Figure 4.5). Morphological properties were then extracted from the reconstructions using NeuroExplorer (Nex Technologies, Littleton, MA).

The cells in this data set were from cortical layers (L) 2, 3, 5 and 6, with layer borders determined microscopically based on cytoarchitectural features. The feature set used in my analysis consisted of the same features employed for the cluster analysis of pyramidal neurons in L2–L6 in the original article (Table 8 in Van Aerde and Feldmeyer (2015)). These were comprised of thirteen morphological and nine electrophysiological features, detailed in Appendix A.2. All experimental procedures and calculation of feature values were performed by Karlijn I Van Aerde and Dirk

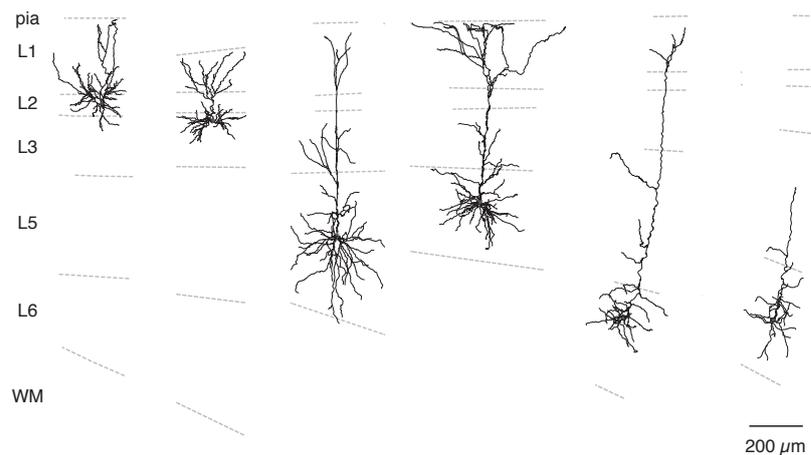


Figure 4.5: Example morphological reconstructions from the mPFC data set, showing cells from various layers. The dotted grey lines show the experimentally defined layer borders. Images obtained from Van Aerde and Feldmeyer (2015).

Feldmeyer, and the processed feature matrix used in my analysis was kindly provided by them. Further experimental details are available from their paper (Van Aerde and Feldmeyer, 2015).

Only cells for which both morphological and electrophysiological properties were available were included in the data set used for clustering, resulting in $n = 103$ neurons. For pyramidal cells in the cortex, each cortical layer is thought to contain distinct populations of neurons (Harris and Shepherd, 2015) (see Section 1.4.1). Cortical layer (2, 3, 4 or 6) was thus used as a proxy ‘ground truth’ label to assess the quality of each clustering of this data set.

Allen Cell Types Database data set

The second neuronal data set examined consisted of electrophysiological and morphological features of cells from mouse visual cortex collected by the Allen Institute and made publicly available through the Allen Brain Atlas Cell Types Database (Allen Cell Types Database (2015)).

Slices were obtained from P45–P70 mice from either interneuron-specific or layer-specific Cre driver lines crossed to an Ai14 tdTomato reporter line, and imaged to allow for brain region identification. Whole cell current clamp recordings were taken from selected neurons, with various stimulation waveforms used to examine a variety of electrophysiological properties of the cells (examples in Figure 4.6). This included injected currents with ramped intensities, short and long square pulses and noise.

High temporal resolution measurements of the membrane potential of the clamped cell were obtained during each stimulation period, and the derivative of the voltage trace was used to detect action potentials. Numerous properties of single action potentials as well as features of the induced spike trains were calculated from this data (Allen Institute for Brain Science, 2016).

Neurons were filled with biocytin during the whole-cell patching procedure (Allen Institute for Brain Science, 2017a). The quality of the possible morphological reconstruction of each cell was first assessed using high resolution image stacks of the filled neurons. For neurons which passed this quality control, the image stack was processed and subjected to multi-scale enhancement. Automatic reconstruction of the morphology of the cells from the enhanced images was then performed using the Neuron Crawler tool (Zhou et al., 2015), followed by extensive manual curation. The dendritic tree and initial part of the axon were reconstructed for all cells, with the full axon reconstruction available for a subset of inhibitory neurons (examples in Figure 4.6).

The dendritic morphology of the cells was labelled as either aspiny, sparsely spiny or spiny by manual inspection of the $63\times$ image. The aspiny and sparsely spiny morphologies were assumed to correspond to interneurons, while cells with spiny morphologies were assumed to be pyramidal or spiny stellate neurons. Cells were assigned to a cortical region and layer using manual annotations of the slices and blockface images. Further experimental details are available from the Allen Cell Types Database Technical White Papers (Allen Institute for Brain Science, 2016; Allen Institute for Brain Science, 2017a; Allen Institute for Brain Science, 2017b).

Electrophysiological features and morphological variables were downloaded for all cells for which both of these properties were available from the Allen Cell Types Database. After processing and removal of uninformative variables, $n = 280$ cells remained that were described by thirty six electrophysiological and twenty morphological features. Explicit descriptions of the features in this data set are given in Appendix A.3. The ‘ground truth’ variables used to assess the quality of the clustering solutions were the classification of the cells into two categories based on their dendrite types, aspiny, which also included the $n = 5$ sparsely spiny cells in the data set, and spiny. As previously mentioned, these two groups are assumed to correspond to predominately inhibitory interneurons and excitatory pyramidal or spiny stellate cells respectively, and so were labelled ‘inhibitory’ and ‘excitatory’ for the rest of this study.

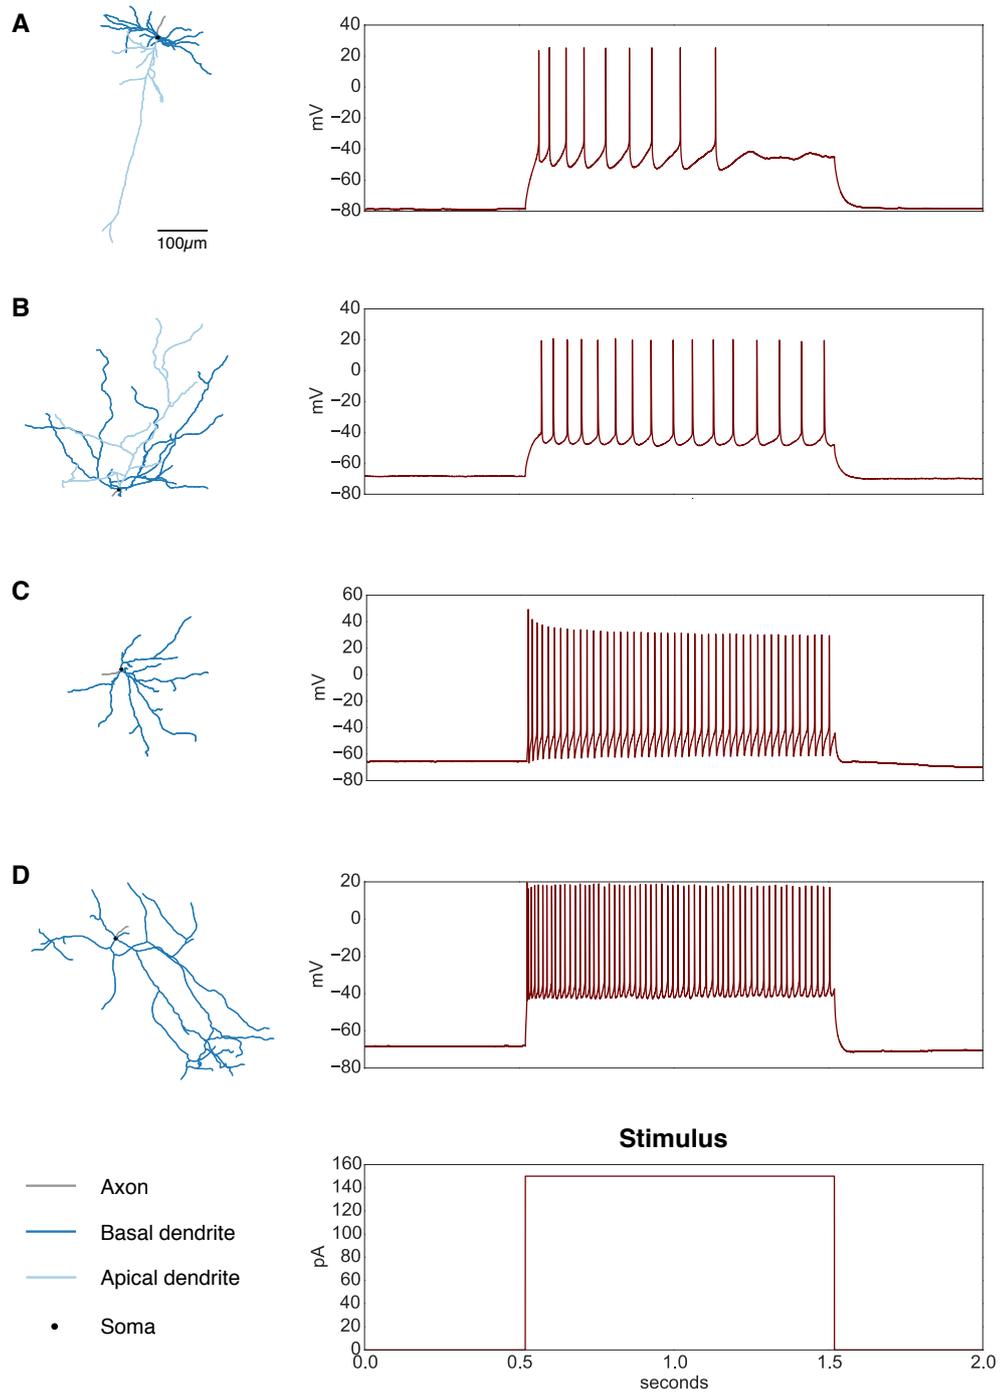


Figure 4.6: Example morphological reconstructions (left) and spiking responses (right) of cells from the Allen Cell Types Database to a 1 s long square stimulus current of 150 pA (shown in bottom right). **A** L5 excitatory neuron, **B** L6b excitatory neuron, **C** L5 SST-expressing interneuron, and **D** L2/3 5HT3aR-expressing interneuron. Scale bar in A and key in bottom left applies to all reconstructions. Figures produced using the Allen Brain Atlas API and Software Development Kit (Allen API (2015); Allen SDK (2015)).

Overview of neuronal data sets

The major features of the two neuronal data sets in this study are summarised below.

Data set	Age range	Species	Region	# cells	# E feats	# M feats
mPFC	P24–46	Rat	mPFC	103	9	13
Allen	P45–70	Mouse	Visual cortex	280	36	20

Table 4.1: Summary of the major features of the mPFC and Allen Cell Types Database data sets.

4.3.3 Clustering of multi-view neuronal data sets

To evaluate the performance of the clustering techniques at analysing multi-view neuronal data sets, the multi-view methods outlined in Section 4.3.1 were applied to the clustering of the two neuronal data sets described in Section 4.3.2. As a baseline, single-view hierarchical clustering using Ward’s method and Euclidean distances was also performed on the individual views as well as a concatenation of the two views from each data set. To examine the robustness of the partitions found by hierarchical clustering, the re-sampling method *BagClust2* (Dudoit and Fridlyand, 2003) was used to repeat hierarchical clustering on a number of bootstrap samples of the data sets, and the mean and standard deviation of several evaluation measures across the bootstrap samples are presented.

As the true number of classes in each of the neuronal data sets is unknown, for each clustering technique the optimal number of clusters was chosen using the gap statistic (Tibshirani et al., 2001). This method estimates the number of clusters in a data set by comparing the pooled within-cluster dispersion of clustering solutions at a range of values of k with the expected dispersion, found using a null reference distribution. For the clustering methods which required the tuning of hyperparameters, this was achieved using a grid search and 10-fold cross validation. The hyperparameter values were chosen as those that minimised the average distance of the left-out data points to their closest cluster centre.

Evaluation measures

In the case of data sets that possess ‘ground truth’ class labels, the performance of clustering techniques can be measured by comparing the similarity of the clusters generated by each method to the correct class structure. Several measures can be used to assess the degree of this similarity, and I employ a number of these in this study. However, as the ground truth cell type labels are unknown for the two neuronal data sets examined, in the place of complete class labels, the labels used were the cortical layer of origin for the cells in the mPFC data set and the classification of the neurons into inhibitory or excitatory type for the Allen Cell Types Database cells.

Firstly, I examined the purity of the split of the known class labels across the clusters identified by each clustering method. Purity is defined as

$$P = \frac{1}{N} \sum_{i=1}^k \max_j |c_i \cap w_j|$$

where N is the number of data points, $C = \{c_1, \dots, c_k\}$ is the set of k predicted clusters and $W = \{w_1, \dots, w_l\}$ is the set of l labelled classes.

However, purity is a biased measure which favours clustering solutions with higher numbers of clusters. In fact, a clustering solution containing N singleton clusters has a purity of one, while providing no useful information about the underlying structure of the data set. As the number correct number of classes in the two neuronal data sets in this study is unknown, the value of k found by each clustering method is a variable that can be optimised, and will generally vary between different clustering techniques. For this reason, two chance-adjusted measures were also used to evaluate the clustering solutions from each of the multi-view clustering techniques. These were the Adjusted Rand Index and Adjusted Mutual Information, both of which are not obviously biased towards any value of k (Vinh et al., 2010)

An adjusted index is defined as

$$\text{Adjusted index} = \frac{\text{Index} - \text{Expected index}}{\text{Max index} - \text{Expected index}}.$$

Adjusted indices have an upper bound of one, and take the value zero when the index is equal to its expected value (Hubert and Arabie, 1985). Two types of adjusted indices were used to evaluate the clustering solutions. The first was the Adjusted

Rand Index (ARI), which was developed by Hubert and Arabie (1985) and based on the Rand Index (Rand, 1971).

The Rand Index (RI) is a measure of the pairwise agreement between two sets of clusters. For two clustering solutions, C and W containing N data points, the RI is defined as

$$RI(C, W) = \frac{N_{11} + N_{00}}{\binom{N}{2}}$$

where N_{11} is the number of pairs of elements that are in the same cluster in both C and W , and N_{00} is the number of pairs of elements that are in different clusters in both C and W . Since the expected value of RI is not constant, the Rand Index has a variable and often high baseline value. By reducing the expected baseline value, the adjusted RI provides a more sensitive measure of cluster similarity.

A permutation model is used to derive the expected value of the Rand Index, based on the assumption that the elements of the contingency table follow a generalised hypergeometric distribution (Hubert and Arabie, 1985). The resultant definition of the ARI between sets of clusters C and W is:

$$ARI(C, W) = \frac{2(N_{11}N_{00} - N_{01}N_{10})}{(N_{11} + N_{01})(N_{00} + N_{01}) + (N_{11} + N_{10})(N_{00} + N_{10})},$$

where N_{01} is the number of pairs of elements in the same cluster in C but different clusters in W , N_{10} is the number of pairs of elements in different clusters in C but the same cluster in W , and N_{11} and N_{00} are defined as previously. Due to its numerous desirable properties, the ARI has become one of the most widely used measures of cluster quality (Vinh et al., 2010).

The third measure used to evaluate clustering quality was a chance-corrected measure of mutual information called Adjusted Mutual Information (AMI). Mutual information measures the shared information between two sets of clusters. The mutual information (Cover and Thomas, 2006) between the set of k clusters, $C = \{c_1, \dots, c_k\}$, and l classes, $W = \{w_1, \dots, w_l\}$, is defined as

$$I(C, W) = \sum_{j=1}^l \sum_{i=1}^k \frac{|c_i \cap w_j|}{N} \log \frac{N|c_i \cap w_j|}{|c_i||w_j|},$$

where N is the number of data points.

The normalised form of this measure, normalised mutual information, is commonly used as a cluster evaluation measure, however it is not unbiased to the value of k , and increases with increasing k . For this reason, the unbiased measure of Adjusted Mutual Information, AMI_{max} , developed by Vinh et al. (2010) was instead employed for cluster evaluation. AMI_{max} is defined as

$$AMI_{max} = \frac{I(C, W) - E\{I(C, W)\}}{\max\{H(C), H(W)\} - E\{I(C, W)\}}$$

where the entropy of C , $H(C)$, is defined as

$$H(C) = - \sum_{i=1}^k \frac{|c_i|}{N} \log \frac{|c_i|}{N}$$

and similarly for W .

Although ARI and AMI are not biased towards higher values of k , they take their maximum value of one only when the clustering solution is identical to the true class structure, and thus penalise solutions in which the ground truth classes are split into a greater number of subclasses. To examine the variability of the three evaluation measures outlined above at different values of k , preliminary hierarchical and k -means clustering was performed on the concatenated mPFC data set described in Section 4.3.2 for a range of values of k . The value of each of the evaluation measures was calculated by comparing the output clusters with the four proxy ‘ground truth’ labels represented by the layer of origin of each cell in this data set (L2, 3, 5 or 6).

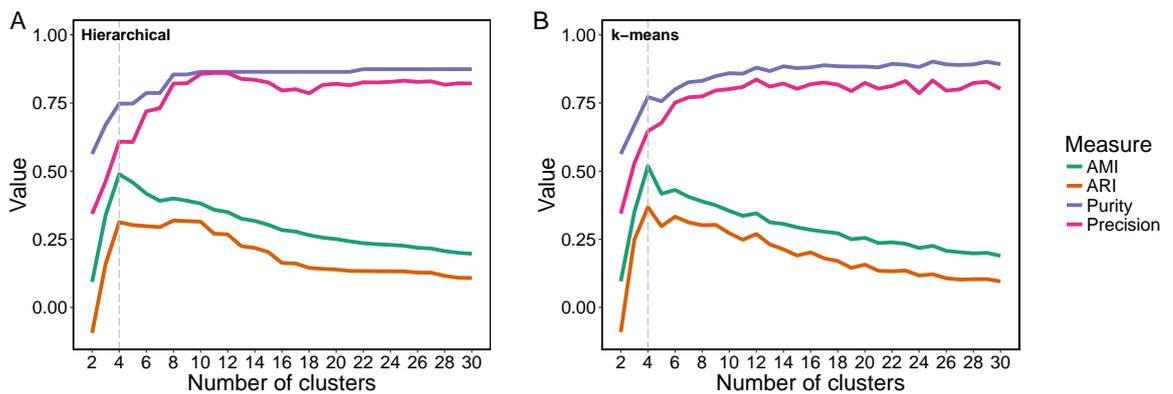


Figure 4.7: The values of AMI, ARI, purity and precision from **A** hierarchical and **B** k -means clustering of the mPFC data set for various values of k . Dotted vertical line shows the number of classes in the ‘ground truth’ labels.

As expected, the value of the purity measure generally increased with increasing k in the clustering results from both hierarchical and k -means clustering (Figure 4.7). For the k -means solution, this was not a strictly monotonic increase, however, this can be attributed to the k -means algorithm converging to different clustering solutions at different values of k . AMI and ARI values were very low for low values of k and initially increased rapidly with increasing k (Figure 4.7). The values of these measures generally peaked at the value of k equal to the number of proxy ‘ground truth’ classes, before decreasing gradually at higher values of k . However, for the hierarchical clustering solution, there was a second peak in ARI at values of k around 8–10 (Figure 4.7A).

For both of the neuronal data sets in this study, the ‘ground truth’ labels take on a smaller number of values than the expected number of cell types in the data. For example, in the mPFC data set, the variable used as the proxy ‘ground truth’ label takes on four values, representing the cortical layer of origin of the cells, however the number of cell types in this data set is expected to exceed four. The clustering methods that identify large numbers of cell types in this data set may thus be disadvantaged by using AMI and ARI to measure the similarity of the clusters found by these methods to the four-class solution provided by the proxy ground truth variables. Due to this, I introduced one additional measure to assess the clustering quality of each method on the multi-view neuronal data sets.

In the case of the neuronal data sets, for a good clustering solution, the number of data points in the same cluster but different ‘ground truth’ classes i.e. N_{10} , should be minimal. For example, for the mPFC data set, N_{10} signifies cells from different layers that are allocated to the same cell type, which should be minimised in a quality clustering solution. But in this case we are unconcerned with the number of data points that are in the same ‘ground truth’ classes but different clusters, i.e. N_{01} , as this simply represents different cell types being present in the same layer. A good clustering solution in this case would have a low value of the following measure:

$$\frac{N_{10}}{N_{10} + N_{11}},$$

which decreases as N_{10} decreases and N_{11} , the number of pairs of data points assigned to the same cluster that have the same ‘ground truth’ labels, increases. This measure is simply $1 - \text{precision}$, so precision was used as the fourth measure to evaluate the clusters identified in the two neuronal data sets, with a higher precision value representing a better clustering solution. As can be seen in Figure 4.7, the value of

precision for the clustering solutions from the mPFC data set initially increased with increasing k . However, this was not a monotonic increase like the purity measure, and further subdivision of the hierarchical clustering solution resulted in a decrease in the precision measure for some values of k , which is a reflection of the underlying structure of the clusters identified (Figure 4.7A).

As most of the clustering methods used in this study were non-deterministic, clustering of each of the data sets was repeated 100 times using each clustering method, and the mean and standard deviation of the four evaluation measures outlined above across the 100 trials were calculated. The median values of the evaluation measures across the 100 trials were also examined, however these were largely similar to the mean values and thus are not presented here. The weighted kernel k -means method is deterministic, as is the linear CCA method using PAM clustering, so only a single value for each of the evaluation measures is presented for these techniques.

4.4 Results

4.4.1 Analysis of the mPFC data set

Relationship between electrophysiological and morphological features

Previous studies have shown that there often exist strong correlations between the electrophysiological and morphological properties of neuronal cells (Hosp et al., 2014; Dumitriu et al., 2007). In their original analysis, Van Aerde and Feldmeyer (2015) describe correlative relationships between some features for subset of cells in the mPFC data set, however, a comprehensive analysis of the correlation between the morphological and electrophysiological variables in the entire data set is not presented. Figure 4.8 shows the Pearson correlation between the 13 morphological and 9 electrophysiological variables in the mPFC data set, as well as their asymptotic p values. Strong correlations are evident between numerous features.

In particular, significant negative correlations were observed between input resistance and several morphological properties, particularly those relating to somatic surface area and the length and number of branches of dendrites. Physiologically, the input resistance of a neuron is inversely proportional to its membrane surface area (see Section 1.2.3), so these correlations are unsurprising from a biological standpoint. The voltage sag, which measures the difference between the initial and sustained response of a cell to a 1 s injection of current that causes hyperpolarization, was

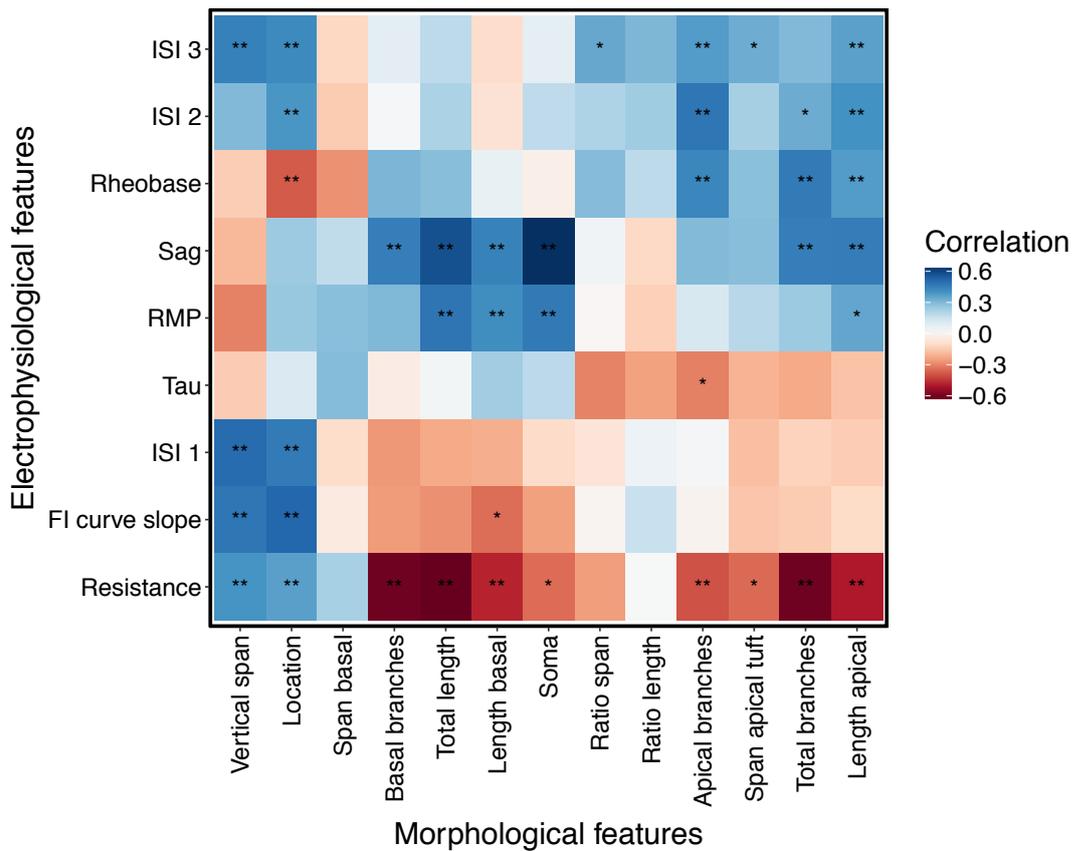


Figure 4.8: Correlation between electrophysiological and morphological features in the mPFC data set. Colours show the sign and magnitude of the Pearson correlation. Stars represent asymptotic p values, at levels $p < 0.005$ (*) and $p < 0.001$ (**), with the null hypothesis of the independence of the feature distributions. p values have been adjusted for multiple comparisons using the Benjamini & Hochberg procedure.

strongly positively correlated with morphological variables relating to dendrite length and somatic area. The value of the voltage sag represents the level of expression of HCN channels in a neuron, which have been shown to increase in density with increasing distance from the soma (see Section 1.2.3). This could explain the larger voltage sags observed in neurons with greater dendritic lengths in this data set. HCN channels are also believed to play a role in regulating the resting membrane potential (RMP) of a neuron (Robinson and Siegelbaum, 2003), which may account for fact that RMP exhibited similar correlations with morphological features related to dendritic length.

The degree of correlation between the morphological and electrophysiological features in this data set varied for cells in different layers. For the cells in L2, there were no significant correlations between their morphological and electrophysiological prop-

erties. For L3 cells, significant positive correlations were observed between apical dendrite length and several electrophysiological parameters such as rheobase, RMP, voltage sag and the membrane time constant at $p=0.01$.

The cells in L5 exhibited the strongest correlations between their morphological and electrophysiological features. Similar to the relationships found in the entire mPFC data set, there were significant negative correlations between input resistance and properties relating to dendritic length and branch numbers, and positive correlations between these dendritic properties and voltage sag and rheobase evident for L5 cells. There were also significant positive correlations between ISI 2 and ISI 3 and features of dendritic morphology. The ISI features measure the rate of spike frequency adaptation of a neuron to an input current. This suggests that the fast-adapting L5 neurons in this data set had larger dendritic arbours than L5 neurons that adapt slowly to input current injections. Considerable morphological diversity existed across the L6 cells in the mPFC data set, and no significant correlations between the morphological and electrophysiological properties of these cells were found.

Clustering of the mPFC data set

The single and multi-view methods outlined in Section 4.3.1 were applied to the clustering of the mPFC data set, and the four evaluation measures described in Section 4.3.3 used to assess the output of each method. In order to most closely replicate the original analysis by Van Aerde and Feldmeyer (2015), the full electrophysiological and morphological feature sets listed in Table 8 of Van Aerde and Feldmeyer (2015) and described in Appendix A.2 were used as input to each of the clustering methods. The impact of performing principal components analysis to reduce the dimensionality of the feature sets before clustering was also investigated, but this produced similar results to the use of the full feature set, so the results are not presented here.

In their original paper, Van Aerde and Feldmeyer (2015) present the results from hierarchical clustering of the concatenated morphological and electrophysiological feature sets from the mPFC cells. In this study, I also perform hierarchical clustering using the same concatenated feature set. However, to ensure consistency in the implementation of all of the clustering methods in this study, the gap statistic was used to find the optimal number of clusters from hierarchical clustering in this case, resulting in some differences in the results presented here and those in the original paper. I also describe the results from hierarchical clustering of

the individual morphological and electrophysiological feature sets, which was not examined in the original analysis of this data set by Van Aerde and Feldmeyer (2015).

Although the expected values of AMI and ARI from a random clustering are zero due to the way in which they are defined, there is no fixed expected value for the purity and precision measures. To investigate this, as well as the variability of each measure, I calculated the 95th percentile value of the four evaluation measures from hierarchical clustering of 100 randomly shuffled versions of the concatenated mPFC feature set. This acts as a baseline, used to assess the significance of the clustering solutions from each method compared to a random clustering of the data set. The 95th percentile values varied for different values of k , and the top line of Table 4.6 shows the maximum 95th percentiles of each measure for $k \geq 5$.

4.4.2 Single-view clustering of the mPFC data set

First, the clustering results from single-view hierarchical clustering of the individual and concatenated mPFC feature sets were examined. Bagged hierarchical clustering of the concatenated feature set identified twelve clusters with a 86.6% purity in the split of cells by cortical layer, as well as reasonably high values of AMI, ARI and precision (Table 4.6). Using only the thirteen morphological variables, eleven clusters were identified, with purity of over 89%, as well as higher AMI and ARI than the clustering solution from the concatenated feature set. Using only the electrophysiological feature set, on the other hand, produced a lower quality clustering solution, with $\sim 74\%$ purity in the split of cells by layer and also low values for AMI, ARI and precision. This suggests that the electrophysiological characteristics are less informative for separating the neurons by cell type in this data set.

Table 4.2 shows the level of similarity between the clusters found using each set of features, as measured by both the AMI and ARI. There was a low level of similarity between the clusters from the morphological feature set and those found using electrophysiological features only. Higher levels of similarity were observed between the clusters from the concatenated feature set and those found using either the morphological or electrophysiological features alone. However, these two values are comparable, suggesting that hierarchical clustering of the concatenated feature set does not just find the dominant clusters in one of the views, but identifies a different set of clusters using both morphological and electrophysiological information.

	AMI			ARI		
	Morph	Ephys	Concat	Morph	Ephys	Concat
Morph	1	-	-	1	-	-
Ephys	0.193	1	-	0.298	1	-
Concat	0.414	0.506	1	0.585	0.523	1

Table 4.2: Similarity between the cluster solutions from hierarchical clustering of the morphological, electrophysiological and concatenated mPFC feature sets, measured by AMI and ARI.

Next, the three explicit hierarchical clustering solutions of this data set are examined. A common notation is used throughout, in which $\mathcal{C}_i^{A,B}$ refers to the i th cluster produced from data set A using clustering method B . i.e. $\mathcal{C}_1^{M,H}$ refers to the first cluster produced from hierarchical clustering (H) of the morphological feature set (M). For each clustering solution, the importance of each of the features for differentiating between the clusters was assessed by performing pairwise Mann-Whitney tests to compare the distribution of each feature between each of the pairs of clusters, with p-values adjusted for multiple comparisons using the Benjamini & Hochberg procedure. The features that exhibited the highest number of pairwise significant differences at $p=0.05$ for each clustering solution are given in Tables 4.3, 4.4 and 4.5, along with the breakdown of the cells by layer in each cluster. Note here that the absolute values of these p values do not represent the true level of statistical significance, as the cluster solutions have been derived using the same feature set. The relative p values, however, are informative for indicating which features differ the most between clusters.

Layer 2

In the solution from hierarchical clustering of the concatenated feature set, L2 cells were split over two clusters, one containing purely L2 cells ($\mathcal{C}_1^{C,H}$), while the other contained a mixture of L2 and L3 cells ($\mathcal{C}_2^{C,H}$), which had lower apical lengths and higher rheobase values compared to the cells in $\mathcal{C}_1^{C,H}$ (Table 4.3). Clustering using the morphological features alone was slightly more effective at separating L2 cells from cells in other layers (Table 4.4), while electrophysiological features alone could not separate L2 and L3 cells effectively (Table 4.5).

Cluster	Cells per layer				Feature values		
	L2	L3	L5	L6	Basal length (μm)	Basal span (μm)	Rheobase (pA)
$\mathcal{C}_1^{C,H}$	6	0	0	0	1527 ± 269	183 ± 12	162 ± 12
$\mathcal{C}_2^{C,H}$	5	6	0	0	2455 ± 203	235 ± 11	232 ± 11
$\mathcal{C}_3^{C,H}$	0	11	0	0	2887 ± 188	275 ± 9	121 ± 8
$\mathcal{C}_4^{C,H}$	0	3	10	0	3629 ± 213	342 ± 20	85 ± 6
$\mathcal{C}_5^{C,H}$	0	1	5	0	4835 ± 335	384 ± 33	177 ± 12
$\mathcal{C}_6^{C,H}$	0	0	19	0	5624 ± 261	406 ± 9	111 ± 6
$\mathcal{C}_7^{C,H}$	0	0	2	5	1933 ± 231	231 ± 15	77 ± 13
$\mathcal{C}_8^{C,H}$	0	0	3	3	3043 ± 346	722 ± 116	55 ± 7
$\mathcal{C}_9^{C,H}$	0	0	19	0	4580 ± 208	323 ± 8	173 ± 11
$\mathcal{C}_{10}^{C,H}$	0	0	0	2	1563 ± 154	173 ± 22	65 ± 5
$\mathcal{C}_{11}^{C,H}$	0	0	0	1	2317	293	70
$\mathcal{C}_{12}^{C,H}$	0	0	0	2	2445 ± 531	304 ± 31	160 ± 20

Table 4.3: Solution from hierarchical clustering using the concatenated feature set from mPFC cells. Columns 2–5 show the number of cells from each layer in each cluster. Bold values show the dominant layer in each cluster. Columns 6–8 show the values of various features in each cluster (mean \pm standard error of the mean).

Cluster	Cells per layer				Feature values		
	L2	L3	L5	L6	Span (μm)	Sum length (μm)	Total branches
$\mathcal{C}_1^{M,H}$	1	0	10	0	298 ± 10	10397 ± 357	74 ± 2
$\mathcal{C}_2^{M,H}$	8	0	0	0	188 ± 10	5047 ± 137	47 ± 2
$\mathcal{C}_3^{M,H}$	2	19	0	0	250 ± 8	5260 ± 176	39 ± 2
$\mathcal{C}_4^{M,H}$	0	2	19	0	390 ± 20	8100 ± 163	47 ± 1
$\mathcal{C}_5^{M,H}$	0	0	8	0	410 ± 20	10052 ± 235	58 ± 2
$\mathcal{C}_6^{M,H}$	0	0	9	2	357 ± 17	6068 ± 348	37 ± 2
$\mathcal{C}_7^{M,H}$	0	0	3	4	232 ± 16	5770 ± 419	44 ± 3
$\mathcal{C}_8^{M,H}$	0	0	7	0	352 ± 15	12454 ± 568	79 ± 4
$\mathcal{C}_9^{M,H}$	0	0	2	2	884 ± 82	4831 ± 358	24 ± 3
$\mathcal{C}_{10}^{M,H}$	0	0	0	4	238 ± 41	3471 ± 562	36 ± 3
$\mathcal{C}_{11}^{M,H}$	0	0	0	1	293	6634	36

Table 4.4: Solution from hierarchical clustering using the morphological features of mPFC cells. Table formatted as per Table 4.3.

Cluster	Cells per layer				Feature values		
	L2	L3	L5	L6	Time constant (ms)	Resistance (M Ω)	Rheobase (pA)
$\mathcal{C}_1^{E,H}$	3	4	1	0	25.0 \pm 1.6	150 \pm 13	154 \pm 8
$\mathcal{C}_2^{E,H}$	7	6	1	0	18.4 \pm 0.5	148 \pm 10	223 \pm 10
$\mathcal{C}_3^{E,H}$	1	0	0	5	17.5 \pm 1.4	264 \pm 34	123 \pm 16
$\mathcal{C}_4^{E,H}$	0	1	5	3	32.0 \pm 0.9	235 \pm 17	79 \pm 6
$\mathcal{C}_5^{E,H}$	0	7	3	0	27.3 \pm 1.4	135 \pm 8	126 \pm 7
$\mathcal{C}_6^{E,H}$	0	1	16	0	33.8 \pm 1.3	156 \pm 9	108 \pm 5
$\mathcal{C}_7^{E,H}$	0	2	10	0	44.2 \pm 1.7	209 \pm 14	71 \pm 5
$\mathcal{C}_8^{E,H}$	0	0	21	0	21.4 \pm 0.9	93 \pm 5	178 \pm 9
$\mathcal{C}_9^{E,H}$	0	0	1	4	33.8 \pm 1.7	396 \pm 32	48 \pm 4
$\mathcal{C}_{10}^{E,H}$	0	0	0	1	27.3	473	70

Table 4.5: Solution from hierarchical clustering using the electrophysiological features of mPFC cells. Table formatted as per Table 4.3.

Layer 3

Clustering using only the electrophysiological feature set also performed very poorly at separating L3 cells from cells in all other layers, with L3 cells present in six of the ten clusters identified in this feature set (Table 4.5). Clustering using the concatenated or morphological feature sets achieved a clearer separation of L3 cells from cells in other layers. These two clustering solutions disagreed on the number of cell types present in L3, with the morphological features identifying one cell type in L3 ($\mathcal{C}_3^{M,H}$), while the concatenated clustering solution split the L3 cells primarily into two main clusters ($\mathcal{C}_{\{2,3\}}^{C,H}$) (Tables 4.3, 4.4).

Layer 5

L5 cells were split over a number of clusters in the solutions from each of the three feature sets. In the predominantly L5 clusters from the concatenated feature set ($\mathcal{C}_{\{4,5,6,8,9\}}^{C,H}$), morphological properties showed the most significant variations (Table 4.3). In particular, the basal dendrite length, total dendrite length and total number of dendritic branches varied greatly between these L5 clusters.

Several L5 cells with low basal lengths ($\mathcal{C}_7^{C,H}$) or high basal span ($\mathcal{C}_8^{C,H}$) were incorporated into clusters with L6 cells with similar properties in the concatenated clustering solution. Similar combined L5 and L6 clusters were also present in the clustering solution from the morphological feature set ($\mathcal{C}_{\{7,9\}}^{M,H}$) (Table 4.4). In the

clusters found using electrophysiological features only, L5 cells were largely split over four clusters which varied significantly in terms of their input resistance, rheobase and time constant (Table 4.5).

Layer 6

L6 cells were also split across numerous clusters in all three clustering solutions. In the solution from the concatenated feature set, L6 cells were commonly clustered with L5 cells with similar features ($\mathcal{C}_{\{7,8\}}^{C,H}$) (Table 4.3). There were also three small clusters containing purely L6 cells in this solution ($\mathcal{C}_{\{10,11,12\}}^{C,H}$), which were well separated by the morphological feature of their apical dendrite length. The pure L6 clusters found in the morphological feature set were similar to those identified using the concatenated feature set, however two of the concatenated clusters ($\mathcal{C}_{\{10,12\}}^{C,H}$) were combined into a single cell type in this solution ($\mathcal{C}_{10}^{M,H}$) (Table 4.4).

One of the L6 clusters identified in the concatenated feature set ($\mathcal{C}_{11}^{C,H}$) contained a single L6 cell, which was identical to cluster $\mathcal{C}_{11}^{M,H}$ found using only morphological features. This cell was an outlier in terms of the span of its apical tuft dendrites, which at $839\ \mu\text{m}$ was the largest of all the cells in the data set, and significantly greater than that of the cell with the next longest span ($642\ \mu\text{m}$). This cell was not an outlier in terms of its electrophysiological features, and when clustered using electrophysiological features only, was incorporated into a cluster with other L3, 5 and 6 cells ($\mathcal{C}_4^{E,H}$) (Table 4.5).

4.4.3 Multi-view clustering of the mPFC data set

The multi-view methods described in Section 4.3.1 were next used to cluster the mPFC data set, and the clustering quality assessed using the four evaluation measures. The evaluation measure scores from all of the clustering methods exceeded the baseline values given by the 95th percentiles from hierarchical clustering of 100 random shuffled versions of the mPFC data set (Table 4.6, top row), indicating that all of the clustering methods considerably outperformed random clustering of the data set.

Excluding the iCluster+ algorithm and cluster ensemble method using k -means clustering, all of the multi-view methods outperformed hierarchical clustering of the electrophysiological feature set according to the majority of the evaluation measures. Additionally, a number of multi-view methods outperformed hierarchical clustering of the morphological or concatenated feature set (Table 4.6).

Two of the highest performing methods belonged to the latent subspace category of multi-view techniques, namely the linear CCA and RCCA methods, which both achieved higher scores than the three hierarchical clustering solutions according to all of the evaluation measures. The RCCA method outperformed the linear CCA method according to the purity and precision measures. This method identified thirteen clusters with the highest purity (0.912) and precision (0.854) of all of the clustering methods, as well as high values for AMI and ARI. The linear CCA method identified nine clusters with the highest AMI and ARI values (0.396 and 0.277 respectively) of any clustering solution. The purity and precision values from this method were also amongst the top three of any clustering method on this data set.

Two of the bootstrapping clustering techniques, the spectral centroid and spectral pairwise methods, produced clustering solutions which scored highly across the evaluation measures. Of these two methods, the spectral centroid method produced the highest quality clusters, and in particular achieved high values of AMI (0.367) and ARI (0.274). The eight clusters found by this method, however, were not amongst the highest performing methods according to their purity and precision scores, which were both below 0.86.

The early integration clustering methods exhibited reasonable performance at clustering the mPFC data set, and often outperformed simple hierarchical clustering of the concatenated feature set, however, their performance was generally not comparable to the highest performing multi-view clustering methods on this data set. Of these methods, the weighted kernel k -means method, which found 14 clusters with purity of 87.4% and a high precision value of 0.846, exhibited the strongest performance.

The post-clustering combination methods were the poorest performing multi-view techniques according to the evaluation measures. The tendency of the cluster ensemble method to find relatively equally sized clusters could account for the poor performance of this technique at clustering the mPFC data set, in which the number of cells are not evenly distributed across layers. The clusters identified by the PLC method using k -means or hierarchical initial clusterings scored reasonably according to the AMI and ARI measures, but had relatively low purity and precision compared to the other multi-view techniques.

Method	k	AMI	ARI	Purity	Precision
Shuffled hierarchical (95th %ile)	≥ 5	0.054	0.047	0.612	0.437
Single					
Hierarchical (Morph)	11	0.348 (0.015)	0.233 (0.021)	0.895 (0.015)	0.828 (0.028)
Hierarchical (Ephys)	10	0.226 (0.015)	0.157 (0.031)	0.744 (0.013)	0.670 (0.037)
Hierarchical (Concatenated)	12	0.321 (0.014)	0.211 (0.023)	0.866 (0.017)	0.832 (0.031)
Early int					
Weighted kernel k -means	14	0.323 (*)	0.242 (*)	0.874 (*)	0.846 (*)
Weighted kernel spectral	10	0.298 (0.023)	0.201 (0.016)	0.833 (0.030)	0.760 (0.043)
TW- k -means	10	0.335 (0.023)	0.243 (0.037)	0.849 (0.026)	0.779 (0.049)
sMD	10	0.315 (0.012)	0.200 (0.011)	0.838 (0.012)	0.749 (0.018)
Subspace					
Linear CCA	9	0.396 (*)	0.277 (*)	0.903 (*)	0.833 (*)
RCCA	13	0.350 (0.020)	0.260 (0.026)	0.912 (0.014)	0.854 (0.029)
Partial multi-view	11	0.286 (0.020)	0.222 (0.025)	0.809 (0.024)	0.786 (0.033)
Robust multi-view spectral	9	0.293 (0.014)	0.193 (0.013)	0.799 (0.015)	0.738 (0.020)
iCluster+	13	0.222 (0.029)	0.132 (0.029)	0.802 (0.036)	0.678 (0.063)
Bootstr.					
Spectral centroid	8	0.367 (0.029)	0.274 (0.030)	0.860 (0.032)	0.800 (0.046)
Spectral pairwise	8	0.361 (0.011)	0.273 (0.011)	0.827 (0.017)	0.782 (0.017)
Spectral co-training	8	0.309 (0.023)	0.216 (0.031)	0.793 (0.024)	0.706 (0.034)
Post clust					
Cluster ensemble k -means	9	0.246 (0.040)	0.143 (0.030)	0.765 (0.045)	0.657 (0.054)
Cluster ensemble hierarchical	10	0.255 (0.023)	0.144 (0.025)	0.803 (0.026)	0.693 (0.034)
PLC k -means	6	0.321 (0.042)	0.255 (0.049)	0.768 (0.047)	0.677 (0.051)
PLC hierarchical	10	0.279 (0.034)	0.200 (0.032)	0.812 (0.039)	0.743 (0.052)

Table 4.6: Results of the single and multi-view clustering techniques on the mPFC data set. The top row shows the 95th percentile values from hierarchical clustering of 100 shuffled versions of the data set. The following rows show the mean and standard deviation of each evaluation measure from 100 repetitions of clustering of the data set using each method. Bold values show the top three performing methods according to each evaluation measure. (* = N/A sd, since method is deterministic)

4.4.4 Cell types from multi-view clustering of the mPFC data set

Based on the results in Table 4.6, two methods stood out as the highest performing multi-view clustering techniques on the mPFC data set. These were the RCCA and linear CCA (LCCA) methods. These two methods identified clusters which had higher scores than hierarchical clustering of any of the single or concatenated feature sets and were also amongst the top four multi-view methods according to all four evaluation measures. The RCCA method found a greater number of clusters ($n = 13$) compared to the LCCA method and outperformed according to the purity and precision measures, while the LCCA method found a smaller number ($n = 9$) of clusters, with the highest AMI and ARI values of all of the clustering methods.

The cell types found by these two high performing methods are next explicitly examined. Due to the use of PAM to partition the projected feature set in the LCCA method, a single clustering solution was produced for LCCA, which is examined below. Because of the random projection element of RCCA, the solutions varied between trials of this method. To find a single clustering solution to use for the analysis of the cell types in the data set, clustering of the data was repeated 100 times using the RCCA method, and the solution with evaluation scores closest to the median values was chosen for examination.

Cell types from the RCCA method

The RCCA method identified thirteen cell types in the mPFC data set (Table 4.7). The clusters found by this method were more similar to those found by hierarchical clustering of the concatenated feature set (AMI=0.679, ARI=0.594) than those from clustering of the morphological feature set alone (AMI=0.553, ARI=0.384) or using electrophysiological features only (AMI=0.382, ARI=0.295).

Layer 2

The RCCA method achieved a clearer separation of L2 cells from cells of other layers than hierarchical clustering of either of the single or concatenated feature sets. This method found one cell type in L2 ($\mathcal{C}_1^{MV,RC}$), which was characterised by cells with low basal dendrite lengths and resting membrane potentials.

Cluster	Cells per layer				Feature values		
	L2	L3	L5	L6	Total length (μm)	Soma area (μm^2)	Total branches
$\mathcal{C}_1^{MV,RC}$	11	3	0	0	5341 ± 229	202 ± 7	49 ± 2
$\mathcal{C}_2^{MV,RC}$	0	10	0	0	5430 ± 265	172 ± 8	38 ± 2
$\mathcal{C}_3^{MV,RC}$	0	5	0	0	5536 ± 413	247 ± 8	37 ± 2
$\mathcal{C}_4^{MV,RC}$	0	2	14	0	6486 ± 306	318 ± 10	41 ± 2
$\mathcal{C}_5^{MV,RC}$	0	1	4	0	7680 ± 285	193 ± 21	45 ± 3
$\mathcal{C}_6^{MV,RC}$	0	0	19	0	9167 ± 255	306 ± 10	52 ± 2
$\mathcal{C}_7^{MV,RC}$	0	0	5	0	13178 ± 349	395 ± 35	82 ± 5
$\mathcal{C}_8^{MV,RC}$	0	0	12	0	10676 ± 230	375 ± 19	74 ± 2
$\mathcal{C}_9^{MV,RC}$	0	0	3	3	4495 ± 330	208 ± 17	26 ± 3
$\mathcal{C}_{10}^{MV,RC}$	0	0	1	5	5876 ± 407	192 ± 18	45 ± 3
$\mathcal{C}_{11}^{MV,RC}$	0	0	0	2	2505 ± 37	167 ± 14	32 ± 4
$\mathcal{C}_{12}^{MV,RC}$	0	0	0	2	4437 ± 169	252 ± 51	40 ± 4
$\mathcal{C}_{13}^{MV,RC}$	0	0	0	1	6634	188	36

Table 4.7: Solution from clustering using the RCCA method on the multi-view mPFC data set. Table formatted as per Table 4.3.

Layer 3

The RCCA method was also able to better separate L3 cells from cells in other layers compared to hierarchical clustering of the concatenated feature set. The RCCA method identified two clusters containing primarily L3 cells ($\mathcal{C}_{\{2,3\}}^{MV,RC}$). These cell types were well separated in terms of their apical length, somatic surface area and voltage sag, with $\mathcal{C}_2^{MV,RC}$ containing cells with greater apical dendrite lengths, lower soma surface areas and lower voltage sag.

Layer 5

The RCCA method identified five clusters containing a majority of L5 cells, which had very high purity in terms of cell layer. Significant variation was observed between both the morphological and electrophysiological features of cells in these L5 clusters. In particular, the total dendrite length, input resistance and rheobase showed high levels of pairwise significant difference between the L5 clusters.

There was also a cluster that contained equal numbers of L5 and L6 cells in the RCCA solution ($\mathcal{C}_9^{MV,RC}$). This cluster was very similar to the clusters $\mathcal{C}_9^{M,H}$ and $\mathcal{C}_8^{C,H}$ from the hierarchical clustering solutions. The majority of cells in this cluster were also combined into a single cell type labelled ‘L5 wide’ in the original analysis by Van Aerde and Feldmeyer (2015).

Layer 6

The RCCA method identified four cell types containing a majority of L6 cells. The vertical span was an important feature for separating these clusters, with one cluster containing cells with long range projections to L1 or L3 ($\mathcal{C}_{10}^{MV,RC}$), while three clusters contained cells which project to local layers (L5 or L6) ($\mathcal{C}_{\{11,12,13\}}^{MV,RC}$).

One of the L6 clusters ($\mathcal{C}_{13}^{MV,RC}$) contained a single cell, which is an outlier in terms of its apical tuft span. This is the same cell that was allocated to a single-cell cluster in the solutions from hierarchical clustering of the morphological and concatenated feature sets ($\mathcal{C}_{11}^{C,H}$, $\mathcal{C}_{11}^{M,H}$). The other two L6 clusters in this solution, which contained cells with short range projections, were well separated in terms of their total dendritic lengths, with $\mathcal{C}_{12}^{MV,RC}$ containing larger cells. Using a similar naming scheme to the original analysis by Van Aerde and Feldmeyer (2015), these cell types ($\mathcal{C}_{\{10,11,12,13\}}^{MV,RC}$) could be labelled as L6 tall, L6 small, L6 short and L6 wide respectively.

Cell types from the LCCA method

The LCCA method identified nine cell types in the mPFC data set (Table 4.8). These clusters had a reasonable level of similarity to the clusters found by the RCCA method (AMI=0.641, ARI=0.576). The LCCA solution also again had higher levels of similarity to the clustering solution from the concatenated feature set (AMI=0.643, ARI = 0.569) than the clusters found using morphological features alone (AMI = 0.581, ARI = 0.441), or just the electrophysiological feature set (AMI=0.419, ARI = 0.347).

Layer 2

The LCCA method produced a clearer separation of L2 cells from cells of other layers than in the RCCA or hierarchical clustering solutions, allocating all L2 cells to one cluster ($\mathcal{C}_1^{MV,LC}$), which contained cells with low basal dendrite lengths and low cortical depth.

Layer 3

The LCCA method found one cell type in L3 ($\mathcal{C}_2^{MV,LC}$). These cells had relatively low voltage sag and low basal spans. The L3 cells in this data set with high voltage sag or larger basal spans tended to be allocated to two clusters containing primarily L5 cells in this solution ($\mathcal{C}_{\{3,4\}}^{MV,LC}$).

Cluster	Cells per layer				Feature values		
	L2	L3	L5	L6	Cortical location (%)	Basal length (μm)	Total length (μm)
$\mathcal{C}_1^{MV,LC}$	11	1	0	0	19.8 ± 1.2	1739 ± 148	5199 ± 241
$\mathcal{C}_2^{MV,LC}$	0	15	1	0	28.1 ± 1.7	2941 ± 145	5514 ± 182
$\mathcal{C}_3^{MV,LC}$	0	4	12	0	42.9 ± 2.0	3593 ± 209	6312 ± 322
$\mathcal{C}_4^{MV,LC}$	0	1	12	0	50.5 ± 2.6	4642 ± 228	8312 ± 300
$\mathcal{C}_5^{MV,LC}$	0	0	13	0	48.6 ± 2.1	6117 ± 260	9173 ± 289
$\mathcal{C}_6^{MV,LC}$	0	0	17	0	51.5 ± 1.9	4622 ± 206	11412 ± 341
$\mathcal{C}_7^{MV,LC}$	0	0	3	8	73.5 ± 2.3	2359 ± 287	5131 ± 334
$\mathcal{C}_8^{MV,LC}$	0	0	0	4	83.3 ± 7.8	1999 ± 338	3471 ± 562
$\mathcal{C}_9^{MV,LC}$	0	0	0	1	75.0	2317	6634

Table 4.8: Solution from clustering using the LCCA method on the multi-view mPFC data set. Table formatted as per Table 4.3.

Layer 5

The LCCA method identified four clusters containing a majority of L5 cells. Both morphological and electrophysiological features varied across the L5 clusters found by the LCCA method. In particular, total dendritic length, basal length, total number of branches, soma surface area and input resistance showed significant variance between the L5 clusters.

Layer 6

The LCCA method found three cell types in L6. One cluster contained the L6 cells with long range projections, as well as three L5 cells ($\mathcal{C}_7^{MV,LC}$). This cluster contained most of the cells allocated to the combined L5/6 clusters $\mathcal{C}_8^{C,H}$ and $\mathcal{C}_9^{M,H}$, plus some additional L6 cells.

The LCCA method identified the same single-cell L6 cluster found by the RCCA and hierarchical clustering methods, containing the cell with the highest apical tuft span ($\mathcal{C}_9^{MV,LC}$). The other short range projection cells in L6 were combined into one cluster in this solution ($\mathcal{C}_8^{MV,LC}$), as opposed to the RCCA solution, in which they were separated into ‘small’ and ‘short’ L6 cell types ($\mathcal{C}_{\{11,12\}}^{MV,RC}$).

Robustness of multi-view clustering solutions

To examine the robustness of the above cluster solutions found by the RCCA and LCCA methods, I took a leave-one-out cross-validation (LOOCV) approach similar to Sümbül et al. (2014). For each method, this involved removing each neuron from the data set in turn, and using the method to cluster the remaining 102 neurons. The removed neuron was then allocated to the closest cluster centre in the projected space. The similarity between the LOOCV solution for each left-out data point and the original cluster solution examined in the preceding sections was measured using the Rand Index (RI).

Using this approach, the RI for LOOCV of the RCCA solution was found to be 0.930 ± 0.015 (mean \pm sd) across the 103 LOOCV trials, and 0.939 ± 0.019 for the LCCA method. This suggests that there is some variability in the clustering solutions from both of these methods caused by small changes to the data set. Given the small size of some of the clusters identified in this data set, some variation caused by leaving out single data points is unsurprising.

4.4.5 Overview of cell types in the mPFC data set

Reclassification of cell layer of origin

Although the layer of origin of the pyramidal cells in the mPFC data set was used to assess the performance each of the clustering techniques in this study, all deviations from this proxy ‘ground truth’ may not necessarily reflect errors in the clustering, but could instead be of biological significance. For example, although pyramidal cell types in the cortex have a general laminar distribution, it has been seen that these distributions do not necessarily precisely match layer borders defined using cytoarchitectonic criteria, and cell types may intermingle at the borders between layers (Oberlaender et al., 2012). In their original analysis, Van Aerde and Feldmeyer (2015) reclassified six cells in this data set into layers different to their experimentally identified layers based on the results from hierarchical clustering of the concatenated feature set.

The clustering solutions from the multi-view techniques presented here also suggest that several neurons in this data set may belong to cell types that are in opposition to their experimentally classified layer. Firstly, of the three L3 cells that were allocated to the predominantly L2 cluster in the RCCA solution ($\mathcal{C}_1^{MV,RC}$), two of these were

located very close (within 8% of the total width of L3) to the layer border with L2. One of these cells, which was 3% from the L2 border, was also classified into the primarily L2 cluster $\mathcal{C}_1^{MV,LC}$ using the LCCA method. This strongly suggests that this cell may actually represent a L2 cell type that was misclassified as a L3 cell based on cytoarchitectonic criteria.

There were also several cells near the L3/L5 border that the clustering results indicate may be experimentally misclassified. The one L3 cell in the predominantly L5 cluster $\mathcal{C}_5^{MV,RC}$ found by the RCCA method was located within 5% of the layer border with L5. This cell was also allocated to a predominantly L5 cluster by the LCCA method and hierarchical clustering of the morphological and concatenated feature sets $(\mathcal{C}_4^{MV,LC}, \mathcal{C}_5^{C,H}, \mathcal{C}_4^{M,H})$, as well as being reclassified in the original study by Van Aerde and Feldmeyer (2015) to an L5 cell type. This strongly suggests that this is in fact a L5 cell type. The other L3 cell that was reclassified as L5 by the original authors was also allocated to a predominantly L5 cluster using the LCCA method and hierarchical clustering of the morphological or concatenated feature sets $(\mathcal{C}_3^{MV,LC}, \mathcal{C}_4^{C,H}, \mathcal{C}_4^{M,H})$, so its reclassification to a L5 cell type is likely correct.

Location in the cortex was included as a parameter in the morphological feature set used to cluster these cells, and thus cortical depth influenced the clusters found using the multi-view methods and hierarchical clustering of the morphological or concatenated feature sets. However, even when clustering was repeated with this feature removed, the results still suggested that the cells identified above belonged to different layers than their experimentally classified layer. This indicates that it is intrinsic properties of these cells, rather than just their laminar location, that has prompted their reclassification.

Van Aerde and Feldmeyer (2015) also proposed the existence of a rare subclass of cells in this data set they label as ‘L5 wide’, which contains cells located in both L5 and L6 that are not close to the layer border, but are sufficiently similar to be classified as a single cell type. The two L6 cells that were allocated to this ‘L5 wide’ subtype were also classified into combined L5/L6 clusters in the results from the LCCA and RCCA methods $(\mathcal{C}_9^{MV,RC}, \mathcal{C}_7^{MV,LC})$. This provides confidence in the existence of a cell type in this data set that contains both L5 and L6 cells.

Other results from the multi-view clustering methods presented here, however, disagree with the reclassification of some cells suggested in the original paper. A neuron that was experimentally labelled as L6 but reclassified as ‘L5 slender high resistance’ by Van Aerde and Feldmeyer (2015) based on their hierarchical clustering results, was allocated to clusters containing primarily L6 cells in the RCCA and LCCA solutions ($\mathcal{C}_{10}^{MV,RC}$, $\mathcal{C}_8^{MV,LC}$). This suggests that in this case, this cell may actually be a L6 cell type that was incorrectly classified by Van Aerde and Feldmeyer (2015) due to their use of single-view clustering techniques. The same is true of an experimentally identified L2 cell that was relabelled as a L3 cell type by Van Aerde and Feldmeyer (2015), but was allocated to clusters containing predominantly L2 cells by both the RCCA and LCCA methods ($\mathcal{C}_1^{MV,RC}$, $\mathcal{C}_1^{MV,LC}$).

Summary of mPFC cell types

Given the innate heterogeneity of biological neurons, and the fact that the layer of origin of these pyramidal cells does not necessarily reflect the absolute ‘ground truth’ of how cells should be separated, a definitive classification of neurons in the rat mPFC into their underlying cell types using the data set given here can not be expected. However, based on the results from the high-performing clustering methods, several conclusions about the identity of the cell types in this data set can be drawn.

Firstly, the mPFC data set appears to contain one type of L2 neurons, which consists of cells with low basal dendrite lengths and resting membrane potentials. Two cell types can be identified in L3, one consisting of neurons with greater apical dendrite lengths and lower voltage sag. L5 contains 4–5 cell types, which vary in both their morphological and electrophysiological properties. Two cells in this data set located in L3 based on cytoarchitectonic criteria but near the border with L5 are most likely L5 cell types. A rare subtype that contains both L5 and L6 cells is also apparent in this data set, and a single cell in L6 with a large apical tuft dendrite span appears to belong to its own class of cells. Three other cell types exist in L6 in this data set, one ‘tall’ cell type that have projections to L1 or L3, and two cell types with projections to local layers that can be classified as ‘short’ or ‘small’ based on the total length of their apical dendrites.

4.4.6 Conclusions from clustering of the mPFC data set

Analysis of this data set of cells from rat mPFC revealed correlations between numerous morphological and electrophysiological features, several of which can be explained by known physiological properties of neurons. Based on hierarchical clustering of the single feature sets, I concluded that the morphological features were the most informative for clustering of this data set, and that cell types could not be identified with a high level of accuracy using the electrophysiological features alone. Although using only the morphological variables can achieve a reasonably good clustering of this data set, as measured by the separation of cells by their cortical layer, the best cluster solutions were produced by using both types of features, in a more informed manner than simple hierarchical clustering of the concatenated feature set.

In particular, the LCCA and RCCA methods, which cluster multi-view data sets by first projecting the features into a latent subspace, achieved the most accurate classification of cell types in the mPFC data set. These two methods produced clustering solutions of higher quality than the single and concatenated hierarchical clustering solutions and were also amongst the top four multi-view methods according to all of the evaluation measures. Analysis of the solution from a single iteration of these methods revealed clusters that were often better separated by cortical layer than those from hierarchical clustering of the single or concatenated feature sets, and generally showed variation in both their morphological and electrophysiological properties.

The RCCA method identified clusters with the highest purity of split of cells by layer, and in particular performed very well at separating L5 and L6 cells from the cells in other layers. With its lower number of clusters, the purity of the LCCA method solution was not as high, and in particular had lower performance at separating L3 cells from other cell types. This method also tended to combine some of the RCCA clusters into a smaller number of cell types, for example combining two L6 cell types identified by the RCCA method into a single cluster.

Some of the cell types in the RCCA and LCCA solutions map well to the cell types found by hierarchical clustering in the original analysis by Van Aerde and Feldmeyer (2015), but there were also differences. For example, the RCCA method identified the ‘small’ and ‘short’ L6 cell type from the original analysis, but had different allocations of L5 cells. There were also several cells that were reclassified

into different layers from their experimentally identified layers by Van Aerde and Feldmeyer (2015) that the RCCA and LCCA methods identified as indeed belonging to their originally experimentally labelled layer. This suggests that these multi-view methods may be able to provide a more accurate classification of cells that represent edge cases than simple hierarchical clustering of the concatenated feature set.

As well as performing hierarchical clustering on the complete mPFC data set, Van Aerde and Feldmeyer (2015) used the results from hierarchical clustering of the cells from each layer independently to obtain the descriptive cell labels in the original study. For each layer, ANOVA tests were used to find the parameters that significantly differed between the clusters, and clustering was repeated using only these parameters. The result of this procedure was that different parameters were used to cluster the cells in each layer. As well as producing cluster solutions with a clearer separation of cells by cortical layer, the LCCA and RCCA methods are also a more automated and unbiased method for classifying cells compared to the approach used in Van Aerde and Feldmeyer (2015), as cells from all layers are clustered simultaneously using the same feature set.

Based on the coherence of the identified clusters with the ‘ground truth’ cell labels, as well as explicit analysis of the cases in which the clustering results deviated from these labels, the RCCA method was judged to produce the most appropriate classification of cell types in the mPFC data set. However, there was some variation in the solution from this method with small variations to the data set, so the cell types identified by this method should be treated with caution and may not generalise to other data sets of mPFC cells.

In conclusion, for the given data set of mPFC cells I have shown that there appears to be no advantage to using both the electrophysiological and morphological properties of cells for their classification if only single-view clustering techniques are considered, as hierarchical clustering of the morphological feature set outperformed that of the concatenated feature set. However, clustering of the electrophysiological feature set alone outperformed what would be expected by random chance, indicating that there is some signal relevant to cell type classification in this feature set. Multi-view techniques allowed the electrophysiological features to inform the clustering in a more effective manner than hierarchical clustering of the concatenated feature set. In particular, the RCCA and LCCA methods outperformed hierarchical clustering of the morphological and concatenated feature sets of this data set, and demonstrate that use of both types of features can improve the clustering results in this case.

4.4.7 Analysis of Allen Cell Types Database data set

Relationship between electrophysiological and morphological features

As in the mPFC data set, correlations were evident between several morphological and electrophysiological variables in the Allen Cell Types Database data set described in Section 4.3.2. In particular, the morphological features relating to dendritic size such as the number of branches, nodes, tips and bifurcations, as well as total branch length were strongly correlated with a number of electrophysiological features, particularly those relating to upstroke-downstroke (UD) ratios and peak and fast trough voltages (Figure 4.9).

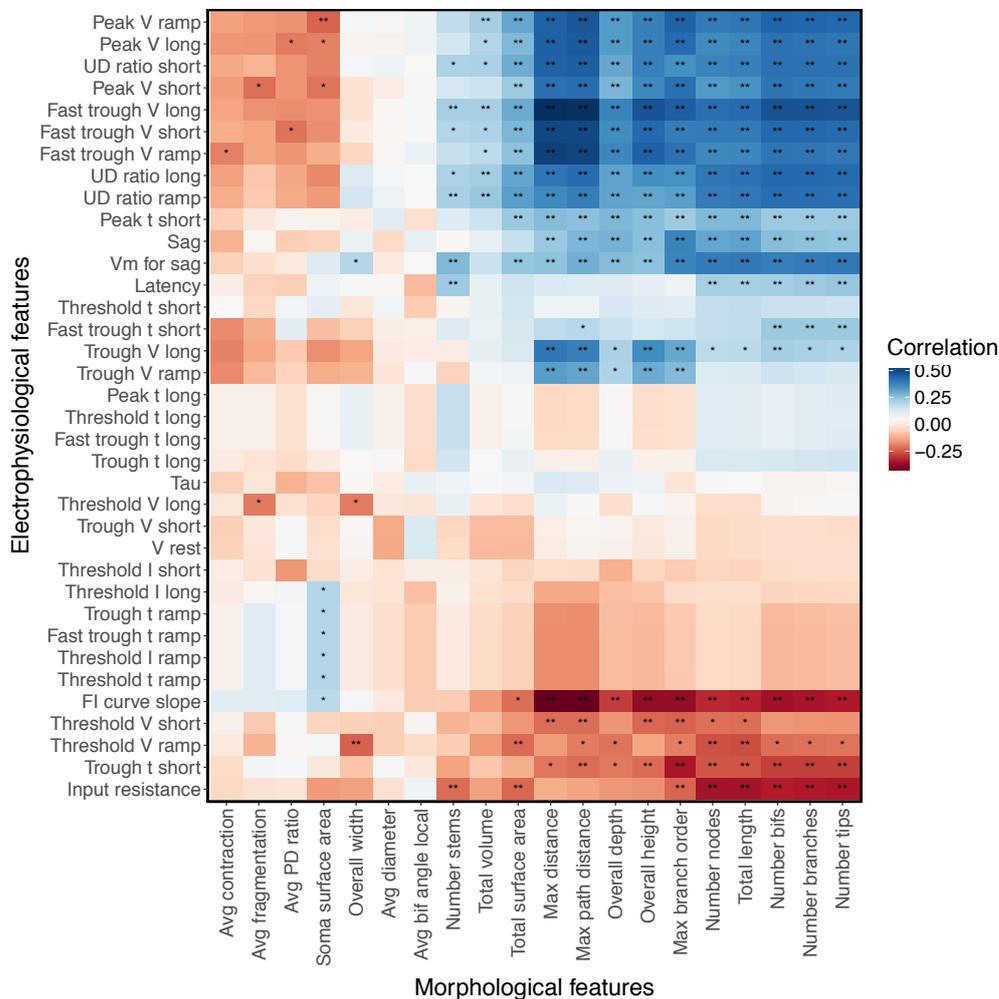


Figure 4.9: Correlation between electrophysiological and morphological features in the Allen Cell Types Database data set. Colours show the level of Pearson correlation. Stars represent asymptotic p values, at levels $p < 0.005$ (*) and $p < 0.001$ (**). p values have been adjusted for multiple comparisons using the Benjamini & Hochberg procedure.

Similar to the mPFC data set, these morphological properties were also significantly positively correlated with voltage sag and negatively correlated with input resistance. Combined, these relationships suggest that the neurons with longer and more branched dendrites in this data set produce larger spikes with faster rates of depolarization, and have lower resistance and more active HCN channels. Other similarities between the correlations in this data set and those in the mPFC data set were also evident. For example, the FI curve slope was negatively correlated with features relating to dendritic size, however the correlation between these features had higher levels of significance in this data set compared to the mPFC data set.

Reduction to principal components

As well as significant correlations between the morphological and electrophysiological features in this data set, there were also strong correlations between the features within each of the two views. Principal components analysis (PCA) was thus used to reduce the dimensionality of the feature sets prior to clustering. Figure 4.10 shows the projection of the data onto the first two principal components of each of the two individual feature sets as well as the concatenation of the feature sets. By visual inspection, there appears to be a clearer division of cells by dendrite type in the concatenated and electrophysiological features sets, compared to the morphological feature set, although there is no linear separation of the cells by dendrite type in any of the feature sets.

Figure 4.10 also shows the cumulative fraction of variance explained by the principal components (PCs) of each of the three feature sets. For the morphological feature set, over 50% of the variance of the data is captured by the first PC. For the electrophysiological feature set, a lower amount of variance is explained by the first principal component, however, there is a steep increase in the cumulative proportion of explained variance as additional components are added. The cutoff for the number of PCs to use for further analysis was chosen as the PC at which at least 90% of cumulative variance was explained, and adding an additional PC explained less than a further 1% of variance. These values were 11, 12 and 16 for the morphological, electrophysiological and concatenated feature sets respectively.

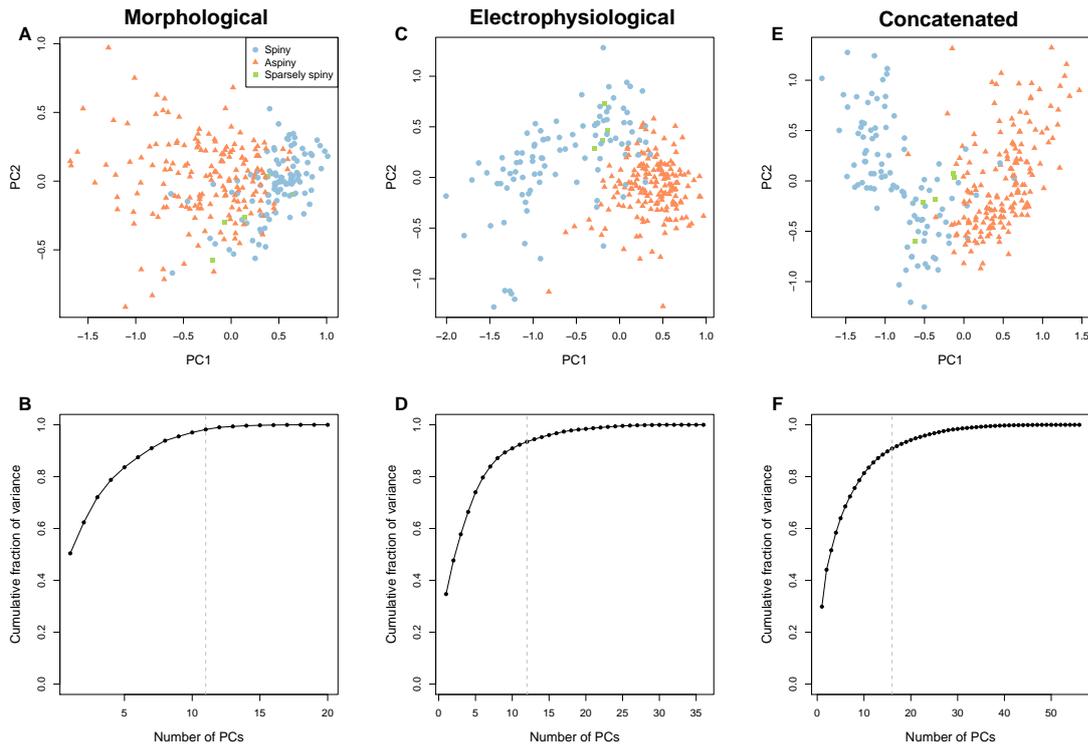


Figure 4.10: Each column represents the PCA from one feature set from the Allen Cell Types Database data set. **A**, **B** Morphological features, **C**, **D** electrophysiological features, and **E**, **F** concatenated feature set. The first row shows the data projected onto the first two PCs of each feature set. Data points are coloured by dendrite type, with the legend in A applicable to all three plots. Spiny dendrites correspond to excitatory cells and sparsely spiny and aspiny dendrites represent inhibitory cell types. The second row shows the cumulative fraction of variance explained by each PC. The grey dotted line shows the number of PCs chosen for use in future analysis.

4.4.8 Single-view clustering of the Allen Cell Types Database data set

First, the results from single-view hierarchical clustering of the individual and concatenated feature sets were examined. Bagged hierarchical clustering was performed on the top PCs from each of the individual views and the concatenated feature set, and the quality of each clustering solution judged by comparing them against the ‘ground truth’ inhibitory or excitatory label of each cell type. Hierarchical clustering of the concatenated feature set identified 21 clusters with an average purity of over 95% in the separation between inhibitory and excitatory cells, as well as reasonably high values of AMI, ARI and precision (Table 4.11). Hierarchical clustering of either of the individual views produced clusters of lower quality than those found by clustering of the concatenated feature set (Table 4.11). This is in contrast to

the previously analysed mPFC data set, for which clustering of the morphological view of the data alone produced clusters that scored higher according to the evaluation measures than the clusters identified in the concatenated data set (Section 4.4.2).

For the Allen Cell Types Database data set, electrophysiological features appeared to be more informative for separating cells by their excitatory and inhibitory cell type compared to their morphological properties. Hierarchical clustering of the electrophysiological features alone identified 21 clusters with $\sim 94\%$ purity, compared to the $\sim 88\%$ purity of the 23 clusters found using hierarchical clustering of the morphological feature set. As seen in Table 4.9, the similarity between the clusters found using the individual feature sets was very low, and the cluster solution from the concatenated feature set had higher a level of similarity to the electrophysiological clusters than the morphological clusters.

	AMI			ARI		
	Morph	Ephys	Concat	Morph	Ephys	Concat
Morph	1	-	-	1	-	-
Ephys	0.155	1	-	0.060	1	-
Concat	0.307	0.495	1	0.150	0.301	1

Table 4.9: Similarity between the cluster solutions from hierarchical clustering of the morphological, electrophysiological and concatenated feature sets from the Allen Cell Types Database, measured by AMI and ARI.

For brevity, here I will focus on examining the cell types found using hierarchical clustering of the concatenated feature set, which produced the highest quality clusters of any of the single-view approaches. Twenty-one clusters were found in the concatenated feature set, thirteen of which consisted primarily of excitatory cells and eight of which contained a majority of inhibitory cells (Table 4.10). The inhibitory and excitatory clusters were well separated by a number of electrophysiological and morphological features. For example, cells in the inhibitory clusters generally had lower UD ratios and were smaller than excitatory cells, with lower maximum Euclidean distances of their dendritic branches (Figure 4.11). As pyramidal cells comprise the majority of excitatory neurons in the cortex (see Section 1.4.1), the bulk of excitatory cells in this data set would be expected to have pyramidal morphologies. The larger maximum Euclidean distances of neurons in the excitatory clusters can thus be attributed to the presence of apical dendrites in these pyramidal cells.

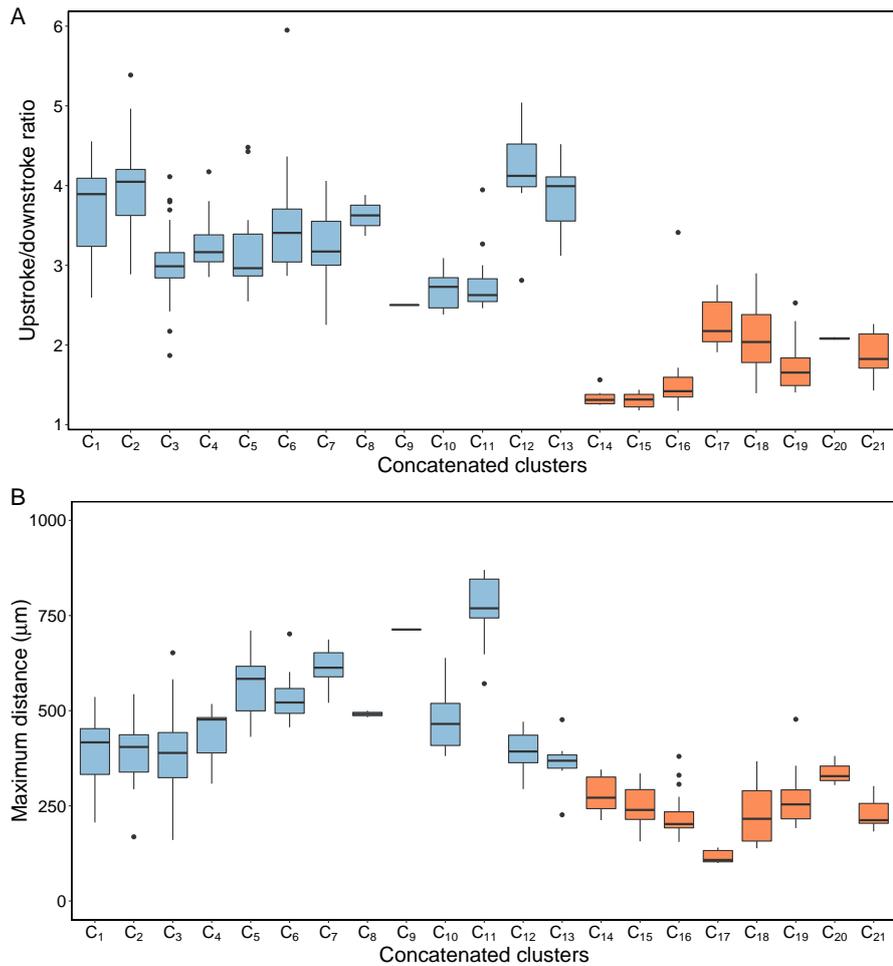


Figure 4.11: Several features differentiated the inhibitory and excitatory clusters found using hierarchical clustering of the concatenated Allen Cell Types Database feature sets. **A** Upstroke/downstroke ratio with a short square input was generally higher in excitatory clusters, as was the **B** maximum Euclidean distance of branches. Clusters are coloured based on the majority cell type in each cluster, either excitatory (blue) or inhibitory (orange).

As in the mPFC data set, the excitatory cell types in this data set would be expected to be well separated by layer. Although some inhibitory cell types display a degree of laminar enrichment, they are generally not as layer-specific as excitatory cells (Tasic et al., 2016). However, previous studies have shown that somatostatin (SST), parvalbumin (PV) and ionotropic serotonin receptor 5HT3a (5HT3aR)-expressing cells divide GABAergic interneurons into three major subtypes with little overlap between classes (see Section 1.4.2). The cells in the Allen Cell Types Database were extracted from a set of Cre recombinase-driver mouse lines, and the majority of inhibitory cells in this data set were Cre-positive for one of these three molecular markers (Allen Institute for Brain Science, 2017b). Assuming that Cre expression reflects

the expression of the corresponding endogenous molecular marker, the inhibitory clusters found in this data set should be well separated by these three markers if they identify true pure cell types.

The excitatory clusters from hierarchical clustering of the concatenated Allen Cell Types Database feature sets exhibited some degree of separation of cells by layer (Table 4.10 rows 1–13, columns 4–9). L2/3 excitatory cells were generally not well separated from cells in other layers and were spread mostly over three clusters, all containing primarily L4 cells ($\mathcal{C}_{\{1,2,3\}}^{C,H}$). Four clusters were identified that consisted of a majority of L4 excitatory cells. These clusters did not have a high level of purity by layer ($\sim 68\%$ average purity across the four clusters), however, they largely contained only cells from neighbouring layers i.e. L2/3–L5.

Five clusters contained primarily excitatory cells from L5. These clusters had reasonably high levels of purity in terms of layer, however, there were also L5 excitatory cells allocated to a number of other excitatory clusters in this solution. One of the L5 clusters ($\mathcal{C}_9^{C,H}$) contained a single cell, which had the highest trough and fast trough voltages evoked by a ramp input current of all of the cells in the Allen Cell Types Database data set.

The L6a excitatory neurons were the most cleanly separated from cells in other layers in this clustering solution. Two clusters containing L6a cells were identified by hierarchical clustering of the concatenated feature set, both of which contained few cells from other layers ($\mathcal{C}_{\{10,11\}}^{C,H}$). The L6b cells in this data set were allocated to two clusters in the hierarchical clustering solution, one which had a high level of purity ($\mathcal{C}_{13}^{C,H}$), while the other contained cells from a range of layers ($\mathcal{C}_{12}^{C,H}$).

The importance of each of the electrophysiological and morphological features for differentiating between the clusters was again assessed by finding the features with the highest number of significant pairwise differences between the clusters. The features that exhibited the greatest number of pairwise significant differences at $p=0.05$ between the excitatory clusters in this solution included both morphological and electrophysiological features, such as the UD ratio with a ramp input, total length, maximum branch order and the number of bifurcations, tips and branches of dendrites.

Cluster	Cells per E/I type		Cells per layer						Cells per molecular marker			
	I	E	L1	L2/3	L4	L5	L6a	L6b	PV	5HT3aR	SST	Other
$\mathcal{C}_1^{C,H}$	0	20	0	5	12	3	0	0	–	–	–	–
$\mathcal{C}_2^{C,H}$	0	30	0	4	23	3	0	0	–	–	–	–
$\mathcal{C}_3^{C,H}$	7	30	0	8	21	6	2	0	–	–	–	–
$\mathcal{C}_4^{C,H}$	0	13	0	0	12	1	0	0	–	–	–	–
$\mathcal{C}_5^{C,H}$	1	11	0	1	3	7	1	0	–	–	–	–
$\mathcal{C}_6^{C,H}$	0	15	0	0	6	9	0	0	–	–	–	–
$\mathcal{C}_7^{C,H}$	0	18	0	0	2	16	0	0	–	–	–	–
$\mathcal{C}_8^{C,H}$	0	2	0	0	1	1	0	0	–	–	–	–
$\mathcal{C}_9^{C,H}$	0	1	0	0	0	1	0	0	–	–	–	–
$\mathcal{C}_{10}^{C,H}$	0	9	0	0	0	2	7	0	–	–	–	–
$\mathcal{C}_{11}^{C,H}$	0	11	0	0	0	2	9	0	–	–	–	–
$\mathcal{C}_{12}^{C,H}$	1	7	0	0	2	1	2	3	–	–	–	–
$\mathcal{C}_{13}^{C,H}$	0	7	0	0	0	0	1	6	–	–	–	–
$\mathcal{C}_{14}^{C,H}$	9	0	0	0	1	6	2	0	7	0	0	2
$\mathcal{C}_{15}^{C,H}$	9	0	0	4	0	4	1	0	5	0	1	3
$\mathcal{C}_{16}^{C,H}$	34	1	0	4	7	18	6	0	20	3	6	6
$\mathcal{C}_{17}^{C,H}$	5	0	0	5	0	0	0	0	0	5	0	0
$\mathcal{C}_{18}^{C,H}$	14	0	1	5	1	7	0	0	0	7	7	0
$\mathcal{C}_{19}^{C,H}$	13	0	0	3	3	6	1	0	2	2	7	2
$\mathcal{C}_{20}^{C,H}$	3	0	0	0	0	2	1	0	0	0	2	1
$\mathcal{C}_{21}^{C,H}$	9	0	0	3	3	3	0	0	0	0	8	1

Table 4.10: Solution from hierarchical clustering using the concatenated feature set from the Allen Cell Types Database. Columns 2–3 show the number of cells that were either inhibitory (I) or excitatory (E) in each cluster. Columns 4–9 show the number of cells per layer in each cluster. Columns 10–13 show the number of cells per molecular marker, only for the clusters containing a majority of inhibitory cells. Bold values show the dominant values in each cluster.

Table 4.10 also shows the split of cells by their expression of the three major inhibitory molecular markers in the data set (PV, 5HT3aR and SST), only for the clusters that contained a majority of inhibitory cells (Table 4.10 rows 14–21, columns 10–13). There was some separation of cells by molecular marker across the nine inhibitory clusters found by hierarchical clustering of the concatenated feature set.

The purest of the inhibitory clusters contained only L2/3 5HT3aR-expressing cells ($\mathcal{C}_{17}^{C,H}$). There were also another three clusters that contained cells expressing only one of the three major molecular markers ($\mathcal{C}_{\{14,20,21\}}^{C,H}$). However, several of the other clusters were quite mixed in terms of molecular marker, with an average purity of $\sim 63\%$ across the eight inhibitory clusters. The features that showed the highest level of pairwise significant difference between the inhibitory clusters in this solution were all electrophysiological features, including the input resistance, trough time with a long square input current and UD ratios with ramp and long square inputs.

4.4.9 Multi-view clustering of the Allen Cell Types Database data set

Next, the multi-view methods from Section 4.3.1 were used to cluster the Allen Cell Types Database data set. For most of the methods, the top PCs from each view identified in Section 4.4.7 were used as input. The exceptions to this were the RCCA, iCluster+ and TW- k -means methods. Since the RCCA method achieves dimensionality reduction by random projection of the features followed by CCA, the full multi-view feature set was used as input to this method. I also used the full feature sets as input to the iCluster+ method, as this method was designed to cluster high-dimensional genomics data sets and uses a Lasso penalty to reduce the contribution of uninformative variables, as well as the TW- k -means method, which performs weighting of individual features as part of the clustering procedure. Using the full feature set for TW- k -means clustering also allowed the contribution of each feature to the final clustering solution to be examined explicitly.

Again, baseline clustering performance was measured by calculating the 95th percentile values of each of the evaluation measures from clustering of the randomly shuffled concatenated data set (Table 4.11, top row). All of the clustering methods outperformed this baseline. Although the clusters found from hierarchical clustering of the concatenated feature set scored reasonably highly across the four evaluation measures, a number of the multi-view clustering techniques outperformed this approach according to several of the evaluation measures (Table 4.11).

Several high performing clustering methods on this data set belonged to the latent subspace category of multi-view techniques. The clusters found by the RCCA method achieved the highest scores of all of the clustering methods across the majority of the evaluation measures. This method identified 19 clusters, with a purity of over 96%, as well as very high values of AMI, ARI and precision. The partial multi-view method also exhibited strong performance at clustering this data set, and was amongst the top three clustering techniques according to most of the evaluation measures. This method identified 20 clusters with the highest ARI (0.107) of any of the clustering methods, as well as high AMI and precision. Additionally, the LCCA method had reasonably high performance at clustering the Allen data set, identifying twenty clusters with comparable evaluation measure scores to hierarchical clustering of the concatenated feature set. The other subspace-based methods in the study exhibited reasonable performance at clustering the Allen data set, but did not consistently outperform single-view hierarchical clustering of the concatenated feature set.

Several early integration multi-view clustering methods also had strong performance at clustering this data set. The highest performing of these methods was the TW- k -means method, which identified 25 clusters with over 96% purity in the split of cells by excitatory and inhibitory type, and outperformed concatenated hierarchical clustering according to most of the evaluation measures. The weighted kernel k -means and sMD methods also exhibited strong performance at clustering this data set, whereas the weighted kernel spectral method generally underperformed hierarchical clustering of the concatenated feature set or either individual view.

Of the bootstrapping approaches to multi-view clustering, the spectral centroid and pairwise methods exhibited the strongest performance at clustering this data set. These methods generally outperformed hierarchical clustering of the individual views of the data, but underperformed hierarchical clustering of the concatenated feature set and the other high performing multi-view methods.

Finally, none of the post-clustering combination methods identified clusters of high quality in the Allen Cell Types Database data set. The clusters found by these methods had purity values generally below 90%, and ARI and AMI values comparable to or lower than those from clustering of the morphological feature set alone.

Method	k	AMI	ARI	Purity	Precision
Shuffled hierarchical (95th %ile)	≥ 5	0.054	0.014	0.686	0.554
Single					
Hierarchical (Morph)	23	0.120 (0.006)	0.055 (0.008)	0.881 (0.011)	0.804 (0.024)
Hierarchical (Ephys)	21	0.162 (0.005)	0.080 (0.008)	0.935 (0.011)	0.892 (0.027)
Hierarchical (Concatenated)	21	0.178 (0.008)	0.083 (0.006)	0.956 (0.008)	0.928 (0.018)
Early int					
Weighted kernel k -means	29	0.150 (*)	0.075 (*)	0.954 (*)	0.934 (*)
Weighted kernel spectral	24	0.105 (0.012)	0.044 (0.005)	0.864 (0.022)	0.807 (0.032)
TW- k -means	25	0.171 (0.006)	0.084 (0.008)	0.961 (0.006)	0.937 (0.014)
sMD	20	0.159 (0.008)	0.073 (0.005)	0.933 (0.014)	0.892 (0.019)
Subspace					
Linear CCA	20	0.166 (*)	0.083 (*)	0.946 (*)	0.914 (*)
RCCA	19	0.194 (0.011)	0.101 (0.006)	0.966 (0.014)	0.953 (0.020)
Partial multi-view	20	0.179 (0.007)	0.107 (0.009)	0.950 (0.009)	0.940 (0.011)
Robust multi-view spectral	17	0.167 (0.009)	0.083 (0.005)	0.931 (0.013)	0.893 (0.017)
iCluster+	21	0.145 (0.011)	0.074 (0.005)	0.909 (0.015)	0.866 (0.020)
Bootstr.					
Spectral centroid	24	0.153 (0.010)	0.060 (0.005)	0.932 (0.019)	0.900 (0.024)
Spectral pairwise	23	0.154 (0.010)	0.063 (0.005)	0.931 (0.017)	0.902 (0.022)
Spectral co-training	18	0.141 (0.011)	0.068 (0.006)	0.899 (0.020)	0.845 (0.025)
Post clust					
Cluster ensemble k -means	20	0.124 (0.015)	0.053 (0.007)	0.884 (0.024)	0.823 (0.033)
Cluster ensemble hierarchical	20	0.136 (0.010)	0.057 (0.004)	0.901 (0.020)	0.849 (0.024)
PLC k -means	28	0.116 (0.013)	0.046 (0.007)	0.893 (0.021)	0.835 (0.035)
PLC hierarchical	17	0.128 (0.016)	0.069 (0.010)	0.870 (0.024)	0.795 (0.036)

Table 4.11: Results of the single and multi-view clustering techniques on the Allen Cell Types Database data set. Table formatted as per Table 4.6.

4.4.10 Cell types from multi-view clustering of the Allen Cell Types Database data set

As discussed in Section 4.4.8, as well as the division of cells by their inhibitory and excitatory type, the split of excitatory cells by layer and inhibitory cells by molecular marker offers another measure of the quality of each clustering solution. Table 4.12 shows the purity levels of these values in the clusters produced from hierarchical clustering of the concatenated feature set, as well as the three multi-view methods that outperformed hierarchical clustering according to at least three of the evaluation measures. These were the RCCA, TW- k -means and partial multi-view methods. The single values from the deterministic hierarchical solution are presented, alongside the mean and standard deviation across 100 runs of the multi-view methods.

Method	Excitatory purity	Inhibitory purity
Hierarchical	0.694	0.629
RCCA	0.700 (0.016)	0.596 (0.025)
TW- k -means	0.700 (0.024)	0.664 (0.034)
Partial multi-view	0.626 (0.018)	0.500 (0.026)

Table 4.12: Purity of excitatory clusters by layer and inhibitory clusters by molecular marker from the top performing clustering methods on the Allen Cell Types Database data set.

The RCCA, TW- k -means and hierarchical clustering solutions achieved comparable levels of purity in the split of excitatory cells by layer, of 69 – 70%, while the partial multi-view solution exhibited lower levels of layer specific separation amongst its excitatory clusters. The TW- k -means solution exhibited the highest purity in the split of inhibitory cells by molecular marker, although a degree of the high purity scores from this method may be attributable to the larger number of clusters found by this method ($n = 25$) compared to the other high performing clustering techniques ($n \leq 21$). The inhibitory clusters from the RCCA method had slightly lower purity compared to the hierarchical clustering solution, while the partial multi-view method achieved a far poorer separation of inhibitory cells by molecular marker than the other high performing clustering techniques.

Based on these results and those in Table 4.11, the TW- k -means and RCCA methods appeared to be the multi-view clustering techniques that achieved the most accurate clustering of the Allen Cell Types Database data set. The explicit cell types identified by a single iteration of these methods are next examined. As these methods are both

non-deterministic, the single solutions chosen for examination were those that had evaluation measure scores closest to the median values from 100 repetitions of each method. The explicit break down of cells by layer and molecular marker in each of these clustering solutions are given in Appendix A.4 (Tables A.4, A.5).

Cell types from the RCCA method

The RCCA method identified 19 cell types in the Allen data set, ten of which consisted primarily of excitatory cells and nine consisting of a majority of inhibitory cells (Table A.4). As in the hierarchical clustering solution, L2/3 excitatory cells were not well separated from cells from other layers in this solution, and were spread over five clusters containing primarily L4 cells ($\mathcal{C}_{\{1,2,3,4,5\}}^{MV,RC}$). Similar to the hierarchical clustering solution, these primarily L4 clusters generally only contained cells from neighbouring layers.

As in the hierarchical clustering solution, L5 excitatory cells were present in a large number of the excitatory clusters, primarily those dominated by L4 cells. However, there were also three clusters in this solution that consisted of a majority of L5 cells, two of which had high levels of purity ($\mathcal{C}_{\{6,8\}}^{MV,RC}$). As opposed to the hierarchical clustering solution, in which two L6a excitatory cell types were identified, excitatory L6a cells were primarily allocated to one cluster, which also contained several L5 cells ($\mathcal{C}_9^{MV,RC}$). L6b cells were more cleanly separated from cells in other layers by the RCCA method compared to the hierarchical clustering solution, with one cluster containing almost all of the L6b excitatory cells in the data set ($\mathcal{C}_{10}^{MV,RC}$). The features with high pairwise significant difference between these excitatory clusters were all morphological features, including the overall height, and number of dendritic branches and tips.

In the inhibitory clusters from the RCCA method, PV-expressing cells were generally not well separated from cells expressing the other molecular markers (Table A.4, rows 11–19, columns 10–13). However, 5HT3aR-expressing cells were generally better separated from other inhibitory cell types than in the hierarchical clustering solution. An almost identical cluster to cluster $\mathcal{C}_{17}^{C,H}$ from the hierarchical clustering solution was identified by the RCCA method ($\mathcal{C}_{15}^{MV,RC}$), which contained the same five L2/3 5HT3aR-expressing cells as in the hierarchical cluster $\mathcal{C}_{17}^{C,H}$, plus the single L1 cell in the data set. As in the hierarchical clustering solution, SST-expressing cells were present in a number of inhibitory clusters. Several electrophysiological features

showed high rates of pairwise significant difference between these inhibitory clusters, including the input resistance, threshold current, threshold time, trough and fast trough time and UD ratio with a ramp input.

Cell types from the TW- k -means method

The TW- k -means method identified a larger number ($n = 25$) of clusters in the Allen Cell Types Database data set than the other high performing clustering methods, eighteen of which consisted of a majority of excitatory cells, and seven of which contained primarily inhibitory cells (Table A.5). L2/3 excitatory cells were again not well separated from cells in other layers in this solution, with L2/3 excitatory cells present in six clusters. Eight clusters in this solution contained a majority of L4 excitatory cells. Although several of these clusters contained only cells from neighbouring layers, there were also two clusters in which L4 excitatory cells were clustered with cells from L6b ($\mathcal{C}_{\{1,7\}}^{MV,TW}$).

There were seven clusters containing a majority of L5 excitatory cells in this solution, which had reasonable levels of purity by layer. L6a and L6b excitatory cells were less well separated from cells from other layers by the TW- k -means method compared to the clustering solutions from the hierarchical and RCCA methods, and were spread across several clusters. The features that showed the highest levels of pairwise significant difference between the excitatory clusters identified by the TW- k -means method were all morphological features, including the overall height, maximum path distance, maximum Euclidean distance and number of dendritic tips.

The TW- k -means method achieved a higher purity of split of cells by molecular marker across its inhibitory clusters than the hierarchical clustering or RCCA solutions, and in particular was more effective at separating PV-expressing neurons from other cells. The TW- k -means method did not find the same pure cluster of L2/3 5HT3aR-expressing cells as was found in the hierarchical clustering solution ($\mathcal{C}_{17}^{C,H}$), however four of the five cells in this cluster were allocated to the same cluster in the TW- k -means method solution ($\mathcal{C}_{23}^{MV,TW}$). The fact that this cluster was mostly conserved across the three clustering solutions provides confidence that this may represent a pure cell type. The features that most frequently significantly differed between the inhibitory clusters from the TW- k -means method were all electrophysiological features, such as the UD ratio with a short square input, threshold current, threshold time and trough time with a ramp input current.

As well as generating a clustering solution, the TW- k -means method also finds optimal weights for each of the views and the individual features as part of the clustering procedure. For the Allen Cell Types Database data set, the view weights found by the TW- k -means method were approximately equal. The distribution of feature weights within each of the two views, however, varied (Figure 4.12). In the electrophysiological view, a small number of features were very highly weighted compared to other features (Figure 4.12A). These included latency, the slope of the FI curve, membrane time constant and fast trough time with a short square input. The top morphological features were more equally weighted, and consisted mostly of features relating to the length and degree of branching of dendrites (Figure 4.12B).

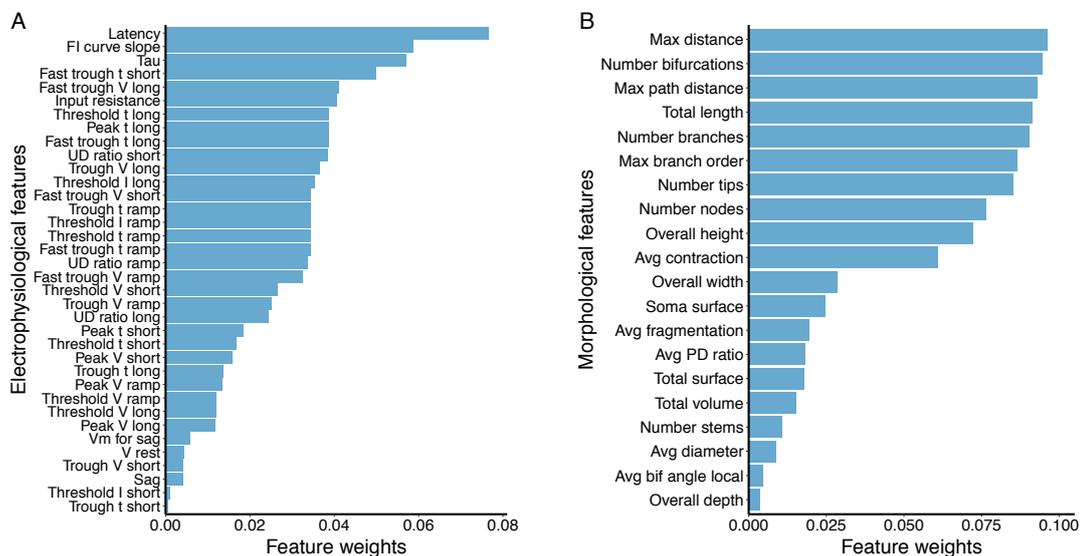


Figure 4.12: Feature weights from TW- k -means clustering of the Allen Cell Types Database data set. **A** Electrophysiological features and **B** morphological features. Within each feature set, the weights sum to one.

Robustness of the multi-view clustering solutions

The LOOCV procedure described in Section 4.4.4 was also performed on the solutions from the RCCA and TW- k -means methods for the Allen Cell Types Database data set. For the RCCA solution, the LOOCV RI was 0.922 ± 0.005 , and 0.936 ± 0.06 for the TW- k -means method. This suggests that both methods show some variability in their clustering output with alterations to the data set, however some of this can be attributed to the non-deterministic nature of these two clustering methods.

4.4.11 Overview of cell types in the Allen Cell Types Database data set

Comparison of Allen Cell Types Database clustering solutions

Table 4.13 shows the similarity between the hierarchical clustering and two multi-view clustering solutions examined previously. The highest levels of similarity were observed between the results from the TW- k -means method and concatenated hierarchical clustering solution, whereas the RCCA and TW- k -means methods solutions were the most dissimilar. Despite not being large in absolute value, the similarities between these solutions were all significantly higher than those between the hierarchical clustering solutions from the individual views of the Allen Cell Types Database data set (Table 4.9).

	AMI			ARI		
	Concat	RCCA	TW- k -means	Concat	RCCA	TW- k -means
Concat	1	-	-	1	-	-
RCCA	0.469	1	-	0.271	1	-
TW- k -means	0.505	0.397	1	0.321	0.215	1

Table 4.13: Similarity between the clustering solutions of the Allen Cell Types Database data set, measured by AMI and ARI.

Reclassification of Allen Cell Types Database cells

As discussed in Section 4.4.5, the points at which the clustering solutions deviate from the proxy ‘ground truth’ labels may be of biological significance. As in the mPFC data set, for the excitatory cells in the Allen Cell Types Database data set, some intermingling of cell types at the layer borders would be expected. However, for this data set, no feature relating to the cortical depth was included in the feature set used for clustering, so the cell types identified were completely independent of their cortical location. The cortical depth of the cells, however, was reported and can be used to analyse which cells were likely to be located close to layer borders. As can be seen in Figure 4.13, there was some overlap between the normalised cortical depth from cells allocated to different layers in this data set.

Based on the results from hierarchical clustering and the high-performing multi-view methods, one neuron in this data set that has a high probability of being experimentally mislabelled is the L6a excitatory cell with the highest cortical depth of all L6a

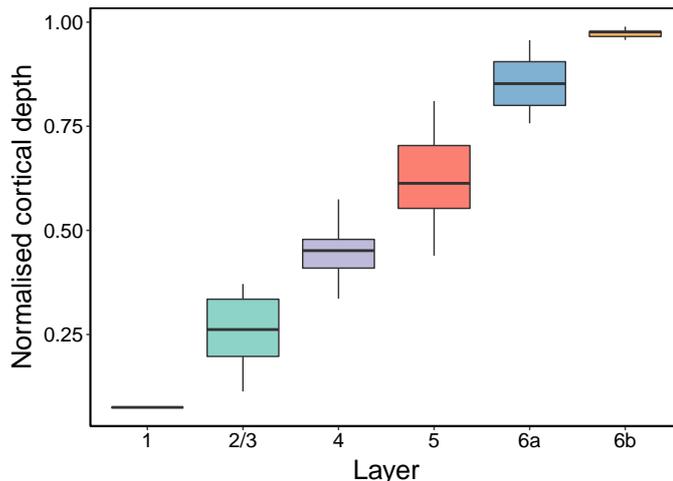


Figure 4.13: Normalised cortical depth of cells by layer in the Allen Cell Types Database data set.

cells in the data set (~ 0.96). This cell was allocated to a cluster containing primarily L6b cells in the solutions from hierarchical clustering of the concatenated feature set and the RCCA and TW- k -means methods ($\mathcal{C}_{13}^{C,H}$, $\mathcal{C}_{10}^{MV,RC}$, $\mathcal{C}_{18}^{MV,TW}$). There is thus a high chance that this cell belongs to a L6b cell type, rather than L6a as it was experimentally labelled. The L5 excitatory cell with the highest cortical depth of all L5 cells (~ 0.81) was also clustered with predominantly L6a cells in all three clustering solutions ($\mathcal{C}_{10}^{C,H}$, $\mathcal{C}_9^{MV,RC}$, $\mathcal{C}_{15}^{MV,TW}$), and so is likely to be a true L6a cell type. Additionally, the two L4 cells in the predominantly L5 cluster $\mathcal{C}_7^{C,H}$ from the hierarchical clustering solution had high cortical depth relative to other L4 cells (> 0.50), and were both allocated to primarily L5 clusters in the multi-view clustering solutions ($\mathcal{C}_6^{MV,RC}$, $\mathcal{C}_{10}^{MV,TW}$). These may also represent neurons that do not correspond to cell types from their experimentally identified layer, and are actually better classified as L5 cell types.

For inhibitory cells, there has been shown to be minimal overlap between cells expressing the three molecular markers PV, 5HT3aR and SST (see Section 1.4.2). For this data set, I have assumed that cells that are Cre-positive for a particular molecular marker endogenously express the corresponding marker. However, it has been observed that Cre lines do not perfectly coincide to the expression of molecular markers and co-expression of these interneuron-specific markers has been detected in some Cre driver lines (Nassar et al., 2015; Hu et al., 2013). For example, SST-IRES-Cre has been seen to be expressed in some PV interneuron types (Tasic et al., 2016). It is thus possible that some of the inhibitory cells in this data set co-express these molecular markers.

Several 5HT3aR expressing cells in the combined SST/5HT3aR cluster found using hierarchical clustering of the concatenated feature set ($\mathcal{C}_{17}^{C,H}$) were also allocated to clusters containing predominantly SST-expressing cells in the two multi-view clustering solutions examined ($\mathcal{C}_{\{19,17\}}^{MV,RC}$, $\mathcal{C}_{24}^{MV,TW}$). There is thus a possibility that these cells form a subclass of inhibitory cells that express both SST and 5HT3aR. Some of the neurons in the combined PV/SST clusters from the three high performing clustering solutions ($\mathcal{C}_{16}^{C,H}$, $\mathcal{C}_{\{12,13,14\}}^{MV,RC}$, $\mathcal{C}_{21}^{MV,TW}$) may also represent double-positive cells, however, additional molecular characterisation of the cells would be required to confirm this. Given the low numbers of cells co-expressing these molecular markers that have been observed in previous studies, it is probably more likely that these clusters primarily contain several similar cell types which the feature set and clustering methods used here are insufficient to separate.

Since the classification of neurons in this data set into inhibitory and excitatory type was performed by visual assessment of the spiny/aspiny nature of the dendrites, there is also a possibility that the inhibitory or excitatory nature of some cells has been mislabelled. The cluster with the lowest purity in terms of its inhibitory/excitatory split amongst the single and multi-view clustering solutions examined was the cluster $\mathcal{C}_3^{C,H}$ found using hierarchical clustering of the concatenated feature set, which contained 7/37 inhibitory cells. Four of these cells were allocated to predominantly excitatory clusters in the TW- k -means clustering solution, however they were all allocated to primarily inhibitory clusters in the RCCA solution ($\mathcal{C}_{\{16,18\}}^{MV,RK}$). These two RCCA clusters, however, did contain a mix of inhibitory and excitatory cells.

Visual inspection of the reconstructions of these neurons revealed that some had more pyramidal-like morphologies than other inhibitory cells. These cells may represent mislabelled excitatory neurons, however, investigation of the experimental microscopy images would be required to more definitely decide if this is the case, or if these cell simply represent a class of inhibitory neurons with similar properties to excitatory cells. The fact that they were not consistently assigned to excitatory clusters across all three clustering solutions suggests that the later may be more likely, and their allocation to an excitatory cluster in the hierarchical clustering solution may be due to the clustering method chosen.

Summary of cell types in the Allen Cell Types Database data set

A definitive classification of the cell types in the Allen Cell Types Database data set was not evident using the available feature sets and clustering methods, however some conclusions can be drawn from the hierarchical and multi-view clustering solutions examined in the previous sections. Firstly, L2/3 excitatory cells were generally not well separated from L4 cells in any clustering solution, suggesting that the features present in the Allen Cell Types Database data set may be insufficient to separate these cell types. There were at least four excitatory cell types present in L4 in this data set, which had reasonable degrees of similarity to L2/3 and L5 cells. L5 contained at least three excitatory cell types, which were morphologically larger than cells in other layers, with greater overall heights, widths and total dendritic lengths. L6b appeared to contain a largely homogeneous group of cells that can be assigned to one class, consisting of cells with higher UD ratios than other cell types. Finally, L6a excitatory neurons in this data set likely consisted of two cell types, one containing cells with greater overall heights.

This data set contained cells from the two major morphological categories of excitatory neurons, pyramidal and spiny stellate. Unlike pyramidal cells, which are found across a range of cortical layers, spiny stellate cells are located primarily in L4 and are characterised by their star-like dendritic arbours and absence of a prominent apical dendrite (see Section 1.4.1). In the three clustering solutions from this data set examined in the previous sections, there was no obvious separation of spiny stellate cells from other cell types. This suggests that the data set is insufficient to distinguish this cell type from other cortical cells. This could be due to the features used, or the possibly small number of spiny stellate cells present in the data set, which represents a random sample of cortical cells for which the experimental reconstructions were determined to be of high quality (see Section 4.3.2).

For the inhibitory cells in this data set, the clustering solutions strongly suggested the existence of a cell type of L2/3 5HT3aR-expressing inhibitory interneurons, which was mostly well conserved across the three clustering solutions. The clustering results also suggested the presence of one other 5HT3aR-expressing inhibitory cell type, two to three SST-expressing interneuron types and at least three cell types consisting of PV-expressing interneurons. Although a small number of cells in this data set may have mislabelled ‘ground truth’ variables, the lack of a clear separation of cells by layer or molecular marker in any of the clustering solutions suggests that this data set is yet insufficient to correctly classify all of the cell types present.

4.4.12 Conclusions from clustering of the Allen Cell Types Database data set

Analysis of the Allen Cell Types Database data set revealed significant correlations between many of the morphological and electrophysiological properties of these neurons, several of which resembled the relationships between features in the mPFC data set. Based on the results from hierarchical clustering, the electrophysiological features of this data set appeared more informative than morphological features for separating the neurons into their excitatory and inhibitory types, however, using the concatenated feature set produced clusters of superior quality than either of the single views. Similar to the results from analysis of the mPFC data set, several multi-view methods exhibited comparable or superior performance to single view hierarchical clustering of the concatenated Allen Cell Type Database feature sets.

The RCCA method exhibited strong performance at clustering the Allen Cell Types Database data set. When judged by its separation of cells into their inhibitory and excitatory type, the clustering solution from the RCCA method scored highly across the four evaluation measures in this study. Analysis of a single iteration of this method revealed that the excitatory cell types found had reasonable levels of separation by cortical depth. The inhibitory cell types identified by the RCCA method exhibited some separation by molecular marker, but did not display the high level of purity that would be expected from previous studies, which have found little overlap between SST, PV and 5HT3aR-expressing cells in the cortex (Urban-Ciecko and Barth, 2016; Lee et al., 2010).

The TW- k -means method was also a high performing method at clustering the Allen Cell Types Database data set. The excitatory clusters identified by this method had a reasonable level of separation by cortical layer, however there was a greater tendency to allocate cells from non-neighbouring layers into the same clusters by this method, compared to the other high performing clustering methods. Despite identifying the lowest number of inhibitory clusters of all the high performing clustering methods on the Allen Cell Types Database data set, this method achieved the highest level of purity in the separation of these cells by molecular marker.

Compared to the similarity between the multiple clustering solutions of the mPFC data set, there were lower levels of similarity between the clusters identified in the Allen Cell Types Database data set by the single-view and high performing multi-view

clustering methods. The cells in this data set were also not purely separated by layer or molecular marker by any of the clustering methods. These factors resulted in a coarser classification of cell types than what was possible for the mPFC data set, and the choice of a single ‘best’ classification of cell types from the high-performing clustering methods would require a more thorough examination of the cells in each cluster and their relationship to previously identified cortical cell types than what was relevant for this study.

The less definitive classification of the Allen Cell Types Database data set can be attributed to the increased complexity of this data set, which contains a wider variety of neurons (both inhibitory and excitatory cells over a number of cortical layers) compared to the previously analysed mPFC data set. Additionally, the neurons for which both morphological and electrophysiological information was available in this data set were those that were judged to have morphological reconstructions of high quality during experimental data collection (Allen Institute for Brain Science, 2017a), and thus represent a random selection of neurons from which electrophysiological recordings were taken.

This randomness may mean that the 280 neurons in the data set used in this study are insufficient to characterise all of the cell types present, and many cortical cell types may be missing from the data. In the results from the three high-performing clustering methods on this data set, the number of excitatory cell types exceeded the number of inhibitory cell types identified in each case. Given that inhibitory interneurons are widely believed to be more diverse than excitatory cells in the cortex (Ascoli et al., 2008; Kelsom and Lu, 2013; Toledo-Rodriguez et al., 2003), this suggests that the full spectrum of cortical neuronal diversity has not yet been captured in this data set.

A recent study of $n = 1654$ neurons in the mouse visual cortex identified seven inhibitory cell types in L2/3 and six inhibitory cell types in L5 using sparse logistic regression classifiers based on both the morphological and electrophysiological properties of the neurons (Jiang et al., 2015). Extrapolating this to include excitatory cell types and the additional layers in this data set, there could be > 50 cell types present in this brain region alone, which are unlikely to be identified correctly using 280 data points. More data points or additional features may lead to an improved classification of the neurons in this region.

One area that was not examined here was the use of partial multi-view information for clustering. Several multi-view clustering techniques in this study, for example the partial multi-view clustering, LCCA and RCCA methods, can be used when both views are available for only a subset of the total data points. As well as morphological and electrophysiological information for close to 300 cells, the Allen Cell Types Database also contains just electrophysiological characterisations of an additional ~ 650 cells, collected using the same experimental protocol. For the RCCA, LCCA and partial multi-view methods, the subspace used for feature projection can be learnt using the data points with both morphological and electrophysiological features in this data set, and clustering performed on the representation of the electrophysiological features of the larger data set in this space. This type of analysis of the larger electrophysiology-only data set may be able to provide a more definitive classification of the cell types present in the Allen Cell Types Database. The performance of partial multi-view clustering techniques on this data set could also inform future experimental approaches, as to whether full multi-view characterisations of cells are necessary for their classification, or if equivalent classification performance can be achieved by collecting multi-view information from a subset of cells.

4.5 Discussion and conclusions

Previous work in neuronal cell type classification has generally been restricted to the use of single types of features, predominantly the electrophysiological, morphological or molecular properties of neurons, for their classification. The few studies that have utilised multi-view feature sets for neuronal cell type classification have generally taken simplistic approaches to the use of this information, applying single-view clustering techniques to the individual or concatenated feature sets.

In this chapter, I have performed a thorough investigation of the application of multi-view clustering techniques to two two-view neuronal data sets, to address the three questions posed in Section 4.2.4:

1. *Is using both the electrophysiological and morphological features of neurons advantageous for classification, compared to using only a single feature type?*

I have shown that utilising both the electrophysiological and morphological characterisations of neurons is advantageous for the classification of cell types in the two data sets in this study, compared to using a single type of feature. In the case of the mPFC data set, superior clustering results could be obtained

from hierarchical clustering of just the morphological features, compared to hierarchical clustering of the concatenated feature set. However, the best classification of cell types was achieved using both feature types and multi-view clustering techniques. For the Allen Cell Types Database data set, hierarchical clustering of the concatenated feature set and several multi-view methods outperformed single-view clustering using either the electrophysiological or morphological features alone.

2. *Which approach to clustering using multi-view information achieves the greatest accuracy at classifying neuronal cell types? Do multi-view clustering techniques offer an improvement over single-view approaches in this context?*

For both neuronal data sets in this study, several multi-view clustering techniques provided an equivalent or more accurate classification of cell types than single-view clustering methods applied to either the individual feature sets or a concatenation of the feature sets, as measured by the separation of cells according to some known ‘ground truth’ variable in each clustering solution. The performance of the specific multi-view techniques varied across the two neuronal data sets in this study, however, the RCCA method exhibited strong performance at clustering both of the data sets.

3. *Can anything be learnt about the relationship between the electrophysiology and morphology of neurons from this study?*

Correlative relationships were evident between numerous morphological and electrophysiological properties of the cells in the two data sets in this study, several of which could be explained by the known physiology of neurons. By examining the features which significantly differed between the cell types identified by clustering of the Allen Cell Types Database data set, I found that in general, electrophysiological features appeared to be more informative for differentiating between inhibitory cells whereas morphological features generally showed higher levels of significant difference between excitatory cell types.

Despite not identifying a specific clustering technique that definitively achieves the best separation of cell types in data sets composed of the electrophysiological and morphological properties of neurons, several insights can be gathered from this study.

Firstly, for both neuronal data sets in this study, relatively low levels of similarity were observed between the clustering solutions produced from the individual electrophysiological and morphological feature sets. Some previous studies have similarly found significant variation between the clusters identified using morphological

and electrophysiological feature sets independently (Hosp et al., 2014). In contrast, others have shown high degrees of agreement between neuronal cell types classified using morphological or electrophysiological features alone (McGarry et al., 2010; Dumitriu et al., 2007), however, these two studies utilised only very small numbers of cells ($n \leq 18$).

One reason for the observed variation between the single-view clustering solutions in this study may be that the electrophysiological and morphological features had different levels of relevance for classifying cells of different types in these data sets. For example, morphological features appeared to be more informative for separating excitatory cell types in the Allen Cell Types Database data set, while electrophysiological features tended to show greater variation between inhibitory cell types. Additionally, in the mPFC data set, significant correlations were evident between several electrophysiological and morphological properties of cells in layers 3 and 5, however, this was not the case for layer 2 and 6 cells. This suggests that it is important to utilise both morphological and electrophysiological information to accurately identify the various cell types in these data sets.

Secondly, the two forms of clustering using canonical correlation analysis in this study, namely the LCCA and RCCA methods, were found to be high performing methods for clustering multi-view electrophysiological and morphological neuronal data sets. For both data sets, the RCCA method was amongst the top three performing clustering methods, and the LCCA method exhibited equivalent or superior performance to hierarchical clustering of the single or concatenated feature sets. These two methods, however, are not without their limitations. For the LCCA method, one disadvantage is that the dimensionality of the CCA subspace found by this method cannot exceed $\min(p_1, \dots, p_n)$, where p_i , $i = 1, \dots, n$ is the number of features in the i th view of the data. This may be an issue when clustering data sets with large differences in the dimensionality of the views, such as those that contain high-dimensional genomic feature sets as well as electrophysiological or morphological views, which are generally more limited in their number of features.

The RCCA method overcomes this disadvantage by randomly mapping the features in each of the views to generate new feature sets with dimensionality chosen by the user. The non-linear nature of this mapping also allows the non-linear relationships between the features in each view to be considered, whereas LCCA can only take into account linear relationships. The RCCA method may thus be preferred for

clustering of more complex data sets with large differences in the dimensionality of the multiple views. Given its strong performance at clustering the two neuronal data sets in this study, the RCCA method would be my current first choice recommendation for a multi-view clustering technique when approached with the task of classifying cell types using data sets consisting of the electrophysiological and morphological properties of neurons.

Other multi-view methods that exhibited reasonable performance at clustering the neuronal data sets in this study included the TW- k -means, weighted kernel k -means and spectral centroid methods, however, these were not amongst the highest performing clustering methods on both of the neuronal data sets. Further analysis of additional multi-view neuronal data sets may provide a more definitive assessment of their performance at classifying neuronal cell types.

At the other end of the scale, the results from analysis of the two multi-view neuronal data sets in this study suggest that post-clustering combination approaches to multi-view clustering, such as the cluster ensemble method, are not appropriate for clustering these types of two-view data sets. These post-clustering combination methods generally only achieved similar or poorer performance than hierarchical clustering of a single view of the data.

The scalability of each of the multi-view clustering techniques in this study was also examined by utilising each method to cluster non-neuronal multi-view data sets with numbers of data points and features in $\mathcal{O}(10^3)$ (see Appendix A.1.1). Apart from the iCluster+ method, which exhibited computational times in the order of hours when analysing such data sets, the other multi-view techniques in this study did not have significantly higher implementation times than simple single-view hierarchical clustering, suggesting that these techniques are also suitable for analysing multi-view neuronal data sets of greater size and complexity than those examined in this chapter. For the RCCA method, the computational time was dependant on the number of random feature projections utilised, and the original authors suggest setting this value based on the computational budget of the project (Lopez-Paz et al., 2014).

4.5.1 Limitations and future work

A limitation of this study was the lack of full ‘ground truth’ cell labels for the neuronal data sets, which introduced a degree of ambiguity into assessing the performance of the clustering methods on these data sets. A number of measures were thus used to assess each of the clustering methods, some of which favoured solutions with high numbers of clusters, and others that tended to punish methods that divided the data into larger numbers of clusters than the proxy ‘ground truth’ labels. The value of these evaluation measures was thus influenced by both the value of k as well as the innate quality of the clustering solutions found by each method. However, the best performing clustering techniques on each of the data sets generally scored highly across all four of the evaluation measures, providing confidence in the quality of these clustering solutions. Having more substantial ground truth labels, particularly for the Allen Cell Types Database data set, in which the binary split into excitatory and inhibitory types was used to assess the clustering of the data set into > 15 cell types, could allow for a clearer assessment of the performance of the multi-view methods in this context. Another possible proxy ‘ground truth’ label that could be used in future studies is the projection target of excitatory cells.

As the actual number of cell types in the cortical regions represented by the two multi-view neuronal data sets in this study are unknown, the gap statistic was used to choose the optimal value of k for each of the clustering methods. This meant that the clustering techniques often identified different numbers of cell types in the neuronal data sets. An area of consideration is thus the level of granularity at which a cell type should be defined, and how to assess whether the appropriate number of cell types have been identified in a given brain region. This problem has somewhat been addressed in the retina using the principle of ‘tiling’ (Wassle et al., 2009). It has been shown that in the retina, the receptive fields of different types of bipolar cells ‘tile’ the retina with minimal overlap of the receptive fields of cells of the same type (Seung and Sümbül, 2014; Wassle et al., 2009). This concept has been used in classification studies to decide the optimal level at which to define cell types, and if further separation of clusters is required (Light et al., 2012; Wassle et al., 2009).

The tiling property, however, does not extend to all types of cells in the retina, such as some classes of retinal ganglion cells, which show extensive overlap of their dendritic fields (Borghuis et al., 2008; Masland, 2012). These cell types have been shown to adhere to a more general mosaic property. In these networks, the soma of

cells of the same type are regularly arranged with even spacing between cells, but are positioned randomly with respect to the soma of cells of other types (Rockhill et al., 2000). There is also some evidence that this mosaic patterning of cells extends beyond the retina, to regions such as the cerebellum and tectum (Cook and Chalupa, 2000). However, it is still unknown whether a similar concept applies to the entire cortex, and more work in this area is required before the actual number of cell types present in cortical regions can be deduced.

Although in this chapter, I have restricted my analysis to multi-view data sets consisting of the electrophysiological and morphological features of cells due to the availability of these data sets, the same frameworks could be applied to other types of multi-view neuronal data. Recently, Patch-seq methodology has been developed, which allows electrophysiological characterisations and transcriptomic data to be obtained from the same cells through a combination of whole cell patch-clamp recordings and single-cell RNA-sequencing (Cadwell et al., 2015; Fuzik et al., 2016). Researchers at the Allen Institute are also developing techniques which allow morphological reconstructions to be obtained from cells that are processed using Patch-seq methodology, producing three-view morphological, electrophysiological and transcriptomic data sets¹. This type of three-view data has the potential to provide a more definitive classification of cell types than what is possible using currently available data sets. Also, although only incremental increases in clustering quality were produced by the multi-view methods for the electrophysiological and morphological data sets examined in this study compared to single-view clustering techniques, the larger variability in the nature and dimensionality of the features in such three view data sets may result in greater benefits from the use of multi-view techniques in such contexts. The high-dimensional nature of the transcriptomic view in this data set may also influence which multi-view techniques are the most appropriate for clustering of these types of data sets.

Although in this study, I have assessed the performance of the clustering techniques based on their ability to produce the ‘best’ separation of neurons by type, the idea of achieving a ‘perfect’ classification of neuronal cell types is problematic in itself. Unlike, for example, the clustering of handwritten digits, for which there exists distinct discrete classes into which the data can be allocated, due to the innate continuous variation that is apparent in all biological systems, there may not be

¹Presentation at Allen Institute Showcase Symposium, video available at <https://www.youtube.com/watch?v=QxTJeSao8Uo&list=PLN-QyZNMh3PvA3M1nTLDDAaBjWm5VYDcc&index=6>

any natural classification of the entire spectrum of neurons in all brain regions into distinct categories or cell types. Given this, the current purpose of neuronal classification is to categorise neurons into cell types that are defined in such a way as to be functionally significant and useful for understanding how different neurons contribute to computation and cognition. Although cell types have been shown to be generally governed by several overarching principles, for example the laminar distribution of pyramidal cell types in the cortex, in rare cases, functionally useful definitions of cell types may not precisely adhere to these rules. As the role of individual cells within neuronal networks is dictated by a variety of their intrinsic properties, including their morphological, electrophysiological and molecular features, utilising multiple characterisations of neurons is important for achieving a functional definition of cell types.

In this chapter, I have shown that multi-view clustering is a useful tool for incorporating multiple characterisations of cells into neuronal classification studies. Several multi-view clustering methods have been shown to allocate neurons to cell types that correspond more closely to their broad neuronal class than what was achievable using only a single type of feature of neurons. These techniques can also suggest cases in which the experimental labels of neurons disagree with their underlying cell types and offer improved classification of neurons compared to clustering approaches that treat multi-view data as a single feature set, particularly for edge cases. Several multi-view techniques also have added benefits over single-view approaches to clustering, such as being able to provide a multi-view classification of neuronal data sets for which multiple characterisations are only available for a subset of cells.

Since it is still unknown whether all neurons can be unambiguously allocated to a unique cell type, the fact that the clustering techniques in this study did not produce a definitive classification of cell types in the analysed data sets is unsurprising. In the future, a clearer understanding of the true identities of neuronal cell types and the level of separation or overlap between them may be possible using improved data sets of increased size and diversity, and I believe that multi-view clustering techniques will have an important role to play in gaining insights from such data.

Chapter 5

Discussion and conclusions

The aim of this thesis has been to investigate and develop techniques and frameworks for the statistical analysis of neuronal data in two main contexts. Firstly, tools for the analysis of the spontaneous electrophysiological activity, and in particular bursting activity, of developing neuronal networks recorded on MEAs. And secondly, methods for classifying neurons into their underlying cell classes based on multiple types of neuronal properties. To this end, throughout this thesis I have evaluated existing analysis algorithms, proposed adaptations to improve existing methods, investigated the applicability of methods from other fields to the analysis of neuronal data sets and developed frameworks that can be applied to the analysis of experimental data and inform experimental practices. In this chapter, I summarise the key contributions of this thesis and provide some suggestions for future work.

5.1 Summary of key contributions

In this section, I summarise the main results from this thesis:

1. A thorough evaluation of existing methods for detecting bursts in MEA recordings of neuronal networks revealed that no existing burst detection technique possesses all of the desired properties for accurate analysis of bursting activity in highly variable network recordings such as those from human stem cell-derived neuronal networks (Section 2.5).
2. Of the existing burst detection techniques, the MaxInterval and logISI methods are the recommended techniques for analysing bursting in neuronal network recordings (Section 2.5).

3. An adaptation of the Cumulative Moving Average method to restrict its maximum allowed ISI cutoff improves its performance and makes it a suitable method for accurately detecting bursts in a range of contexts (Section 2.7)
4. Analysis of novel recordings of hiPSC-derived neuronal networks revealed an increase in levels of bursting activity in these networks over early development, but to a lesser extent than in developing rodent neuronal networks (Section 2.5.3).
5. Using a quantitative framework for analysing spontaneous activity patterns, hippocampal and cortical cultured networks were shown to exhibit differences in their spontaneous activity over development that can be characterised using a number of features. In particular, hippocampal networks were shown to exhibit higher levels of synchronous activity, more regular bursting activity and higher rates of theta bursting than cortical cultures at equivalent ages (Section 3.2.3).
6. Recordings of hippocampal and cortical networks can be differentiated over the first month of development with reasonable levels of reliability using a small number of quantitative features and standard classification techniques (Section 3.2.3).
7. Recordings from multi-well MEAs can be used to characterise the differences in spontaneous neuronal network activity patterns across developmental time, and are a suitable tool for studying a variety of experimental conditions on the same array (Section 3.3.3).
8. Greater levels of variability can be observed between MEA recordings of neuronal networks produced from different cultures compared to the intra-culture variability between networks recorded on different arrays (Section 3.3.3).
9. Clustering of two multi-view neuronal data sets showed that utilising both the electrophysiological and morphological properties of neurons is advantageous for the classification of cell types, compared to use of a single feature type (Section 4.4).
10. Several multi-view clustering techniques can provide a more accurate classification of neuronal cell types when multiple characterisations of cells are available, compared to single-view clustering of individual or concatenated feature sets (Section 4.4).

5.1.1 Methods of burst detection

In Chapter 2, I performed an evaluation of existing techniques for detecting bursting activity in recordings of the spontaneous activity of neuronal networks on MEAs. A particular focus of this study was to identify burst detection techniques that are appropriate for the analysis of recordings of human stem cell-derived neuronal networks, which have been shown to exhibit highly variable firing activity. This was achieved by first developing a list of desirable properties that a burst detection technique should possess in order to accurately detect bursting in spike trains with a range of features commonly observed in recordings of developing human neuronal networks. Simulated data was then used to evaluate the performance of numerous burst detectors against each of these desirable properties. Based on these results, as well as further investigation of the burst detectors using experimental recordings from mouse RGCs, I found that no existing method for detecting bursts possessed all of the desirable properties deemed necessary for accurate burst detection in a wide range of contexts.

Several methods, however, exhibited strong performance across a number of desirable properties and four of these were chosen to analyse bursting activity in novel recordings from hiPSC-derived neuronal networks across a range of developmental ages (Recordings made by Paulsen Lab, Department of Physiology, Development and Neuroscience). Using these techniques, bursting activity in these networks was shown to increase over early development, however, at a far lower rate than what has been commonly observed in rodent neuronal networks over equivalent developmental timescales.

Next, I proposed adaptations to two of the existing burst detectors in this study, the cumulative moving average (CMA) and Poisson surprise (PS) methods (Kapucu et al., 2012; Legéndy and Salcman, 1985), and re-evaluated these methods using both simulated and experimental data. An adaptation of the CMA method that restricts its maximum allowed ISI cutoff was shown to improve the ability of this method to accurately detect bursts in a range of contexts, and bring its performance in line with other high performing burst detectors. Based on these results, I concluded that the existing MaxInterval and logISI methods (Nex Technologies, 2014; Pasquale et al., 2010) as well as the adapted CMA method are the most appropriate currently available tools for analysing bursting activity in a range of spike trains.

5.1.2 Frameworks for the analysis and classification of developmental spontaneous activity

In Chapter 3, I developed a framework for analysing the ontogeny of spontaneous activity patterns recorded from cultured neuronal networks on MEAs. By applying this framework to novel experimental recordings from *in vitro* hippocampal and cortical neuronal networks, I showed that the spontaneous activity patterns of hippocampal and cortical cultured neuronal networks exhibited quantifiable differences over development. In particular, hippocampal networks tended to exhibit higher degrees of correlation, more regular bursting activity and higher rates of theta bursting than cortical networks at equivalent ages. I further demonstrated that these recordings of hippocampal and cortical networks could be differentiated over the first month of development with reasonable levels of reliability using a small number of quantitative features and standard supervised classification techniques.

Next, I applied a similar framework to recordings of rat cortical neurons on multi-well MEAs, obtained through a collaboration with researchers at the US Environmental Protection Agency. By analysing the developmental profile of spontaneous activity in these recordings and applying a similar classification scheme to the previous study, I demonstrated that these multi-well recordings provide sufficient levels of granularity to capture the differences in network activity at various developmental ages. This suggests that multi-well MEAs are a suitable tool for studying a variety of experimental conditions in a high-throughout manner on a single array, and I further presented a sampling method that can be used to deduce the most appropriate allocation of experimental conditions to the multiple wells in such arrays. Using a separate data set, I was also able to investigate the major sources of variability in multi-well MEA recordings of cultured networks, and showed that the level of variability between networks produced from different cultures exceeded the intra-culture variability across different recording plates.

5.1.3 Multi-view clustering and its application to neuronal cell type classification

In Chapter 4, I investigated the potential advantages of using multi-view data sets consisting of electrophysiological and morphological features of neurons for classification of these cells into classes or ‘types’. Previously, combined electrophysiological and morphological neuronal data sets have typically been analysed using single-view

clustering techniques applied to the individual or concatenated feature sets. This fails to take into account the unique statistical representations of the cells in each view. In this study, I examined the impact of utilising multi-view clustering techniques for the analysis of such data sets. I first performed a thorough literature search to identify a range of methods that have been developed for finding cluster structure in data represented by two or more distinct feature sets or ‘views’, before applying these methods to the clustering of two multi-view neuronal data sets. As far as I am aware, this is the first study to use multi-view clustering techniques for the classification of neuronal cell types.

I first examined the relationship between the electrophysiological and morphological features of the cells in the two multi-view neuronal data sets in this study, and showed that several correlative relationships exist. By applying several clustering techniques to these data sets, I found that using both electrophysiological and morphological properties was advantageous for classification of both neuronal data sets, compared to utilising only one of the feature types. I also showed that using multi-view techniques to cluster these data sets can lead to improved classification of cell types, as measured by the separation of cells by some known underlying variable across the clusters. By interrogating the explicit clustering solutions from high performing methods on these data sets, I observed general trends in how cell types were separated. For example, I found that electrophysiological properties were generally the most informative for differentiating between different types of inhibitory neurons, while morphological features tended to differ most significantly between classes of excitatory neurons. The results from this chapter suggest that the use of multi-view clustering techniques may allow for a more comprehensive classification of neuronal cell types in future experimental data sets.

5.2 Future work

As specific suggestions for future work have been proposed in the discussion and conclusions section of each relevant chapter, here I present an overview of the key points and general future directions for the field.

5.2.1 Analysis of additional MEA recordings

The availability of additional recordings of spontaneous neuronal network activity, in particular those recorded on high-density MEAs, over longer developmental timescales or from treated neuronal cultures, would provide several avenues for future research. Chapters 2 and 3 of this thesis involved the analysis of bursting and spontaneous activity patterns of cultured neuronal networks recorded on MEAs with ≤ 64 electrodes per well or array. Although the resolution of these arrays were sufficient to differentiate between recordings made from different network types and ages, recordings from high-density MEAs, which can contain over 4000 electrodes, would allow for the precise dynamics of neuronal network activity and bursting to be studied at higher spatial resolutions (Berdondini et al., 2009).

Recordings of spontaneous neuronal activity from high-density MEAs have been shown to exhibit similar overall dynamics to networks recorded on standard low-density arrays, however in some cases, a lower proportion of bursting activity and higher numbers of random spikes have been observed in these recordings (Maccione et al., 2014; Lonardoni et al., 2015). This could have implications for the performance of burst detection techniques at analysing data recorded from high-density arrays. For example, the presence of a large number of electrodes recording low levels of bursting activity would complicate the task of choosing fixed threshold parameters for burst detection methods such as the MaxInterval method, and reduce the likelihood that single fixed parameters would be appropriate for detecting bursts across all of the electrodes on the array. An extension of the evaluation of burst detection techniques presented in Chapter 2 to spike trains recorded from these high-density arrays would be an interesting area of future research. This could also involve an investigation of the possibility of incorporating methods to automatically tune the parameters of certain burst detection techniques, which could facilitate the setting of parameter values separately for each electrode, rather than on an array-wide basis.

The finer spatial resolution offered by high-density MEAs could also influence the features used to characterise the developmental profile of spontaneous activity patterns, which was the focus of Chapter 3 of this thesis. For example, these types of arrays allow for the source of bursts of network activity to be identified, and their spatial propagation to be tracked (Maccione et al., 2010; Gandolfo et al., 2010). Thus, instead of treating network spikes as a binary event as I did in Chapter 3, additional features such as the average speed of propagation of single spikes within a

network spike as well as their spatial dynamics could be included in the quantitative profiling spontaneous activity patterns. This could reveal additional ways in which the nature of spontaneous activity patterns vary over development.

High-density arrays are also a more effective tool for recording the activity of low density neuronal cultures compared to standard MEA devices (Maccione et al., 2010; Maccione et al., 2012). The sparse nature of these cultures allows for their morphology to be imaged with greater clarity than cultures containing high densities of cells (Ullo et al., 2014). Combined imaging and electrophysiological studies of low density cultures on high-density MEAs thus have the potential to provide an improved tool for studying the correlation between the morphological and electrophysiological development of neuronal networks (Maccione et al., 2012).

The availability of MEA recordings that span greater developmental time frames could also facilitate improved study of developmental spontaneous activity patterns, particular in human derived networks. In Chapter 2, I investigated the developmental patterns of bursting activity in hiPSC-derived neuronal networks over the first 16 weeks after plating. Some studies published since we completed this work have suggested that hiPSC-derived neurons at this age may still be relatively immature, and that longer recording time frames may be required to ensure that the activity of functional and mature human neuronal networks is captured (Amin et al., 2016; Odawara et al., 2016). Future studies could investigate the ontogeny of bursting activity in these networks over longer time scales and address questions such as if the decrease in the proportion of bursting activity observed in the final time point of the data presented in Chapter 2 (Section 2.5.3) is a trend that continues over development or simply an artefact of this data set.

In the second half of Chapter 3, I presented the results of a project conducted in collaboration with the EPA, the purpose of which was to inform their ongoing work in developing a high-throughput framework for screening potential developmental neurotoxins using multi-well MEAs. In the absence of recordings from drug treated neuronal networks, network age was used as a proxy for different experimental conditions in this study. However, it is difficult to know how the gross levels of activity pattern changes between the different network ages would compare to those between networks under different pharmacological conditions. Replicating the analysis from this section on recordings from drug treated networks could address this question, as well as allow for the investigation of other issues, such as how to

deduce the significance of observed changes to developmental spontaneous activity patterns caused by pharmacological treatments, given the innate variability in the activity of control networks.

On a larger timescale, the analysis methods from Chapters 2 and 3 could be expanded to applications with clinical implications. This includes large scale toxicity screening of novel compounds on human neuronal networks as well as personalised optimisation of therapeutic treatments using recordings of patient-specific iPSC-derived neuronal networks (Inoue et al., 2014).

5.2.2 Increased size and diversity of multi-view neuronal data sets

The application of the multi-view clustering techniques examined in Chapter 4 to additional multi-view neuronal data sets is also an area of future interest. The collection of electrophysiological and morphological characterisations of cells in the mouse visual cortex by the Allen Institute is an ongoing project, with regular data releases. Given the vast diversity of cell types in this brain region, larger data sets would be expected to allow for a more definitive classification of cell types than what is currently possible. The release of the Patch-seq data set containing electrophysiological, morphological and transcriptomic characterisations of single cells from the Allen Institute that is expected at the end of the year (personal communication with the Allen team) would also provide interesting data for future study. In particular, the high-dimensional nature of the transcriptomic component of this data set may have an impact upon which multi-view clustering techniques are the most appropriate for the analysis of this data.

Given the experimental effort involved in collecting these multi-view neuronal data sets, another area of future interest is the effectiveness of using partial multi-view data sets for neuronal classification. Several multi-view techniques outlined in Chapter 4 can be applied to data sets in which complete multi-view information is only available for a subset of data points. If the effectiveness of these methods at classifying neuronal cell types were shown to be equivalent to that which can be achieved using full multi-view data sets, it would allow for the cell types in large data sets of neurons for which only one set of features e.g. their electrophysiological properties, have been measured, to be classified by collecting multi-view representations of a smaller subset of cells. This would greatly reduce experimental time and expense compared to generating complete multi-view data sets from neurons.

5.2.3 Multi-modal neuronal models

In Chapter 4, I applied several clustering techniques to the task of identifying cell types in multi-view neuronal data sets. The lack of availability of experimental multi-view neuronal data sets with complete ‘ground truth’ class labels lead to difficulties in assessing which clustering methods achieved a good separation of cell types. A potential future area of investigation would be the use of simulated data, for which the ground truth cell types are known, for this purpose.

Models that can replicate certain neuronal activity dynamics of single cells based on their morphology are available, but these have generally been developed to simulate the activity of specific well studied cell types, such as neocortical layer V pyramidal cells (Hay et al., 2011; Bahl et al., 2012). Biophysical models tuned to replicate experimental data are also available, for example the Allen Institute have built such models for cells in their combined morphological and electrophysiological database¹. However, since these models simply replicate the experimentally observed dynamics, they do not provide a more useful data set for cell type classification compared to simply using the underlying experimental data for clustering. Improved diversity of models that can simulate realistic electrophysiological activity of cells of different types would increase the availability of labelled simulated data that could be used in multi-view neuronal classification studies, to assess the performance and aid in the development of new clustering techniques.

Conversely, a greater understanding of the relationships between the morphological and electrophysiological features of different cell types, such as that provided through multi-view classification studies, would also allow for the development of improved cell type-specific models that can predict and simulate electrophysiology from morphological representations. Increasing availability of experimental data sets which additionally include gene expression profiles of neurons could also allow for the incorporation of these factors into predictive models of neuron behaviour.

These ideas could also be extended to studying the combined morphological and electrophysiological properties of different cell types during their development using modelling approaches. Although there has been some success in tracking and recording the electrophysiological activity of specific single cells over long developmental time courses, it is an experimentally challenging task, leading to low success

¹<http://celltypes.brain-map.org>

rates and sampling densities (Cohen et al., 2013; McMahon et al., 2014). Several computational models have been developed that successfully synthetically replicate dendritic development in numerous cell types (Donohue and Ascoli, 2008; Graham and Ooyen, 2006; Cuntz et al., 2010). A greater understanding of the relationship between morphological properties of different cell types and their electrophysiology would allow these or similar models to be extended to facilitate the simulation of the electrophysiological features of neurons during development alongside their morphological properties. Given the challenges of collecting experimental data in this area, these models would allow for the improved study of the combined morphological and electrophysiological profiles of developing neuronal cells and networks.

5.2.4 Concluding remarks

Progress in neuroscience research requires a close interplay between experimental data collection and the computational techniques required to analyse this data. Throughout my research I have encountered a frequent disconnect between these two areas, resulting in the slow adoption of new computational methods in favour of the ‘tried and true’ techniques. For example, the oldest and most widely used method for burst detection, the Poisson Surprise method (Legéndy and Salcman, 1985), remains the most commonly used method for analysing bursting activity in a range of *in vitro* and *in vivo* experimental studies to this day, despite the development of a large variety of alternative burst detection algorithms in the preceding thirty years. Computational techniques are also often applied to the analysis of data in contexts far removed from those for which they were originally designed, with a lack of awareness of their limitations.

This thesis has attempted to address these issues by providing unbiased investigations and evaluations of computational methods that can be applied to the analysis of experimental neuroscience data in a range of contexts. I believe that the work presented in this thesis has suggested improvements to the way in which experimental data is analysed, and provided practical advice and tools for the analysis of this data. Given the painstaking effort that goes into collecting experimental data in the neurosciences, I hope that the adoption of ideas presented here can lead to improvements in the clarity and completeness of the conclusions drawn from experimental data, and contribute to future progress in this field.

Appendix A

Supplementary information for Chapter 4

A.1 Multi-view clustering of labelled data sets

A.1.1 Labelled data sets

A preliminary investigation of the multi-view clustering techniques described in Section 4.3.1 was performed using the following data sets:

1. **Handwritten digits data set:** Also known as the UCI Multiple features data set¹, this data set consists of handwritten digits ‘0’ to ‘9’ extracted from Dutch Utility maps. The 2000 data points consist of 200 samples per class. Six feature sets are available, two of which were used in this study. The first of these was a feature set consisting of 76 Fourier coefficients of the shapes of the digits, and the second consisted of 240 pixel averages in windows of size 2×3 across the images of the digits.
2. **BBC segmented data set:** This data set consists of 2225 documents from the BBC news website from 2004–2005, partitioned into five topic categories which are used as the ground truth labels (Greene and Cunningham, 2006). This data set was made multi-view synthetically by dividing each document into segments and randomly assigning each segment to a view, with the restriction that at most one segment from a document can be assigned to each view (Greene and Cunningham, 2009). The segmentation of the documents into two

¹<https://archive.ics.uci.edu/ml/datasets/Multiple+Features>

views was used, with only documents present in both views included in this study, resulting in a total of 2012 documents. The features in each view were the term frequencies (6838 in view one, 6790 in view two), after pre-processing. This data set is available online².

- 3. Reuters Multilingual data set:** This is a subset of the Reuters Corpus Volume 1 Multilingual data set, which is a set of Reuters documents from six major categories. A sample consisting of 200 documents from each of the six classes, produced by Clément Grimal, was used³. For the two views, the original English documents and their French translations were utilised in this study. The 2000 features in each view were the term frequency–inverse document frequency (TF-IDF) scores for each document, after standard pre-processing.

Pre-calculated feature matrices provided by the specified authors were used for the clustering of each data set, with no further calculation of features performed by myself. Each data set was clustered using each of the multi-view techniques with the number of partitions, k , set to the correct number of classes. For the clustering methods which required the tuning of hyperparameters, this was achieved using a grid search and 10-fold cross validation. The hyperparameter values were chosen as those that minimised the average distance of the left-out data points to their closest cluster centre. The mean and standard deviation of the AMI, ARI and purity values were calculated across 10 repetitions of clustering of each data set.

²<http://mlg.ucd.ie/datasets/segment.html>

³<http://membres-lig.imag.fr/grimal/data.html>

A.1.2 Results from clustering of labelled data sets

	Method	AMI	ARI	Purity
Single	Hierarchical (View 1)	0.603 (0.009)	0.487 (0.019)	0.646 (0.016)
	Hierarchical (View 2)	0.804 (0.011)	0.728 (0.018)	0.815 (0.011)
	Hierarchical (Concatenated)	0.811 (0.011)	0.737 (0.016)	0.822 (0.010)
Early int	Weighted kernel k -means	0.842 (*)	0.817 (*)	0.911 (*)
	Weighted kernel spectral	0.727 (0.002)	0.658 (0.003)	0.792 (0.002)
	TW- k -means	0.826 (0.010)	0.783 (0.021)	0.889 (0.023)
	sMD	0.756 (0.001)	0.704 (0.001)	0.841 (0.001)
Subspace	Linear CCA	0.669 (*)	0.568 (*)	0.738 (*)
	RCCA	0.760 (0.004)	0.673 (0.005)	0.786 (0.002)
	Partial multi-view	0.482 (0.002)	0.401 (0.002)	0.463 (0.001)
	Robust multi-view spectral	0.792 (0.006)	0.757 (0.005)	0.878 (0.003)
	iCluster+	0.681 (0.020)	0.662 (0.028)	0.701 (0.013)
Bootstr.	Spectral centroid	0.771 (0.019)	0.732 (0.034)	0.861 (0.031)
	Spectral pairwise	0.768 (0.001)	0.725 (0.001)	0.857 (0.001)
	Spectral co-training	0.746 (0.031)	0.694 (0.053)	0.824 (0.052)
Post clust	Cluster ensemble k -means	0.527 (0.036)	0.392 (0.045)	0.560 (0.046)
	Cluster ensemble hierarchical	0.563 (0.046)	0.419 (0.051)	0.565 (0.053)
	PLC k -means	0.612 (0.075)	0.520 (0.090)	0.678 (0.076)
	PLC hierarchical	0.658 (0.032)	0.571 (0.048)	0.734 (0.032)

Table A.1: Results of the single and multi-view clustering techniques on the handwritten digits data set. The mean and standard deviation of each measure from 10 iterations of each method are shown. Bold values show the top three performing methods according to each evaluation measure. (*= N/A sd, since method is deterministic)

	Method	AMI	ARI	Purity
Single	Hierarchical (View 1)	0.408 (0.049)	0.312 (0.056)	0.604 (0.046)
	Hierarchical (View 2)	0.363 (0.056)	0.259 (0.041)	0.564 (0.045)
	Hierarchical (Concatenated)	0.541 (0.039)	0.427 (0.066)	0.700 (0.041)
Early int	Weighted kernel k -means	0.463 (*)	0.309 (*)	0.641 (*)
	Weighted kernel spectral	0.755 (0.000)	0.779 (0.001)	0.906 (0.000)
	TW- k -means	0.353 (0.068)	0.187 (0.096)	0.467 (0.069)
	sMD	0.773 (0.001)	0.804 (0.001)	0.916 (0.001)
Subspace	Linear CCA	0.684 (*)	0.686 (*)	0.857 (*)
	RCCA	0.655 (0.035)	0.652 (0.064)	0.845 (0.032)
	Partial multi-view	0.759 (0.001)	0.792 (0.001)	0.910 (0.001)
	Robust multi-view spectral	0.777 (0.000)	0.806 (0.001)	0.917 (0.001)
	iCluster+	0.001 (0.000)	0.000 (0.001)	0.243 (0.003)
Bootstr.	Spectral centroid	0.721 (0.001)	0.732 (0.000)	0.884 (0.001)
	Spectral pairwise	0.749 (0.001)	0.781 (0.001)	0.907 (0.001)
	Spectral co-training	0.713 (0.001)	0.723 (0.001)	0.879 (0.001)
Post clust	Cluster ensemble k -means	0.247 (0.055)	0.192 (0.069)	0.453 (0.073)
	Cluster ensemble hierarchical	-0.003 (0.000)	-0.002 (0.000)	0.240 (0.000)
	PLC k -means	0.478 (0.087)	0.402 (0.113)	0.671 (0.094)
	PLC hierarchical	0.487 (0.044)	0.452 (0.044)	0.709 (0.054)

Table A.2: Results of the single and multi-view clustering techniques on the BBC data set. Table formatted as per Table A.1.

	Method	AMI	ARI	Purity
Single	Hierarchical (View 1)	0.217 (0.020)	0.135 (0.019)	0.389 (0.022)
	Hierarchical (View 2)	0.251 (0.018)	0.156 (0.008)	0.421 (0.015)
	Hierarchical (Concatenated)	0.255 (0.021)	0.165 (0.014)	0.433 (0.014)
Early int	Weighted kernel k -means	0.217 (*)	0.136 (*)	0.414 (*)
	Weighted kernel spectral	0.252 (0.001)	0.195 (0.001)	0.488 (0.001)
	TW- k -means	0.181 (0.021)	0.129 (0.010)	0.385 (0.024)
	sMD	0.248 (0.004)	0.192 (0.003)	0.487 (0.003)
Subspace	Linear CCA	0.240 (*)	0.168 (*)	0.466 (*)
	RCCA	0.207 (0.012)	0.127 (0.008)	0.392 (0.009)
	Partial multi-view	0.257 (0.001)	0.217 (0.001)	0.508 (0.001)
	Robust multi-view spectral	0.254 (0.004)	0.197 (0.001)	0.490 (0.004)
	iCluster+	-0.001 (0.001)	-0.000 (0.001)	0.197 (0.001)
Bootstr.	Spectral centroid	0.266 (0.001)	0.209 (0.001)	0.504 (0.001)
	Spectral pairwise	0.261 (0.001)	0.203 (0.001)	0.497 (0.000)
	Spectral co-training	0.239 (0.015)	0.192 (0.008)	0.478 (0.016)
Post clust	Cluster ensemble k -means	0.170 (0.031)	0.136 (0.027)	0.399 (0.040)
	Cluster ensemble hierarchical	0.192 (0.003)	0.112 (0.003)	0.400 (0.005)
	PLC k -means	0.210 (0.036)	0.160 (0.028)	0.426 (0.043)
	PLC hierarchical	0.195 (0.027)	0.139 (0.014)	0.418 (0.025)

Table A.3: Results of the single and multi-view clustering techniques on the multilingual data set. Table formatted as per Table A.1.

A.2 mPFC data set features

A.2.1 Morphological features

1. Location in the cortex (%): Location of the soma between the pia (0%) and white matter (100%).
2. Field span tuft apical dendrites (μm): Widest distance between the apical dendrites, measured parallel to the pia.
3. Field span basal dendrites (μm): Widest distance between the basal dendrites, measured parallel to the pia.
4. Vertical span apical dendrite (end layer): Layer of termination of the furthest apical branch from the soma.
5. Ratio field span apical/basal dendrites: Ratio of features 2 and 3.
6. Total length apical dendrites (μm): Total path length of all apical dendrites.
7. Total length basal dendrites (μm): Total path length of all basal dendrites.
8. Ratio length apical/basal dendrites: Ratio of features 6 and 7.
9. Sum length apical + basal dendrites (μm): Sum of features 6 and 7.
10. Soma (μm^2): The maximal soma circumference after projection to the 2D plane.
11. Number of apical branches: The total number of compartments between branch points or between branch and terminal points of the apical dendrites.
12. Number of basal branches: The total number of compartments between branch points or between branch and terminal points of the basal dendrites.
13. Total number of branches: Sum of features 11 and 12.

A.2.2 Electrophysiological features

1. Ratio ISI-1/ISI-9: Ratio of the first interspike interval to the 9th interspike interval of the spiking activity evoked by the current step which produced at least 10 action potentials. This is a measure of spike-time adaptation.
2. Ratio ISI-2/ISI-9: Ratio of the second to the 9th interspike interval, as per feature 1.

3. Ratio ISI-3/ISI-9: Ratio of the third to the 9th interspike interval, as per feature 1.
4. Increase in firing frequency (Hz/100 pA): Slope of the linear fit to the plot of firing rate against injected current for 0 – 300 pA.
5. Resting membrane potential (mV): Pre-stimulus membrane potential of the neuron.
6. Input resistance ($M\Omega$): Slope of the linear fit to the current-voltage curve between -60 mV to -70 mV.
7. Rheobase (pA): The minimum level of current required to produce an action potential. Calculated using a step size of 10 pA.
8. Time constant (ms): Calculated from the exponential fit of the voltage response curve after a -50 pA current step.
9. Voltage sag (%): The difference between the initial and sustained response to a 1 s current injection that produced a hyperpolarization of approximately -7.5 mV.

A.3 Allen Cell Type Database data set features

A.3.1 Morphological features

1. Number of nodes: The total number of nodes in the neuron reconstruction. A node is a single sample point of the reconstruction.
2. Soma surface area: The surface area of the soma, computed either by assuming the soma is spherical (if the soma is one node) or as the sum of the external cylindrical surfaces of the multiple nodes comprising the soma.
3. Number of stems: The total number of stems attached to the soma.
4. Number of bifurcations: The total number of bifurcation points of the neuron. A bifurcation point is defined as any point with two daughter branches.
5. Number of branches: The total number of branches of the neuron. A branch is any compartment between two bifurcation points or between a bifurcation point and a terminal point.
6. Number of tips: Number of terminal end points of the neuron.

7. Overall height: Difference between the minimum and maximum y -coordinates of the morphological reconstruction after removing outer points using the 95% approximation of y -values.
8. Overall width: Calculated as per feature 7, on the x -coordinates of the neuron.
9. Overall depth: Calculated as per feature 7, on the z -coordinates of the neuron.
10. Average diameter: The average of the diameter of all of the compartments of the neuron.
11. Total length: The sum of the length of all branches of the neuron. Branches are defined as per feature 5.
12. Total surface area: The total surface area of the neuron.
13. Total volume: The sum of the volume of all compartments of the neuron.
14. Maximum Euclidean distance: The straight line distance between the soma and the furthest node from the soma.
15. Maximum path distance: The maximum path distance of all nodes of the neuron. The path distance of a node is the sum of the length of all of the connected compartments between the soma and the node.
16. Maximum branch order: The maximum branch order of all branches of the neuron. Branch order is defined based on the number of compartments between the branch and the soma, where the soma has branch order zero, the first bifurcation has branch order 1, the second bifurcation has branch order 2, etc.
17. Average contraction: The average contraction value of all neuronal branches. The contraction of a branch is the ratio of its Euclidean distance from the soma to its path length from the soma.
18. Average fragmentation: The average number of compartments between two bifurcation points or between a bifurcation point and a terminal end point of the neuron.
19. Average parent-daughter ratio: The average parent-daughter ratio of all bifurcation points of the neuron. One value is calculated for each daughter at a bifurcation point, as the ratio between the diameter of the daughter and its parent branch.
20. Average bifurcation angle local: The average angle between the daughter branches at all bifurcation points of the neuron.

A.3.2 Electrophysiological features

1. Threshold voltage short square: The threshold voltage of the first action potential evoked with a ‘short square’ 3 ms injection of current at the lowest amplitude that reliably produced action potentials. The threshold voltage was defined as the voltage at the point at which dV/dt was equal to 5% of the maximal dV/dt across all action potentials in a sweep.
2. Threshold time short square: The time of the threshold defined in feature 1.
3. Threshold current short square: The level of injected current that produced the threshold defined in feature 1.
4. Peak voltage short square: The maximum value of the membrane potential between the time of the first action potential threshold and the time of the next action potential or the end of the recorded response, evoked by the short square injected current described in feature 1.
5. Peak time short square: The time of the peak voltage defined in feature 4.
6. Trough voltage short square: The minimum of the membrane potential in the period between the peak of the first action potential and the time of the next action potential evoked by the short square injected current described in feature 1.
7. Trough time short square: The time of the trough described in feature 6.
8. Fast trough voltage short square: The minimum membrane potential in the 5 ms after the first action potential peak evoked by the short square injected current described in feature 1.
9. Fast trough time short square: The time of the trough described in feature 8.
10. Upstroke/downstroke ratio short square: The upstroke/downstroke ratio of the first action potential produced by the short square current described in feature 1. Upstroke/downstroke ratio is defined as

$$\text{Upstroke/downstroke ratio} = \frac{|\text{Action potential peak upstroke}|}{|\text{Action potential peak downstroke}|}$$

where the action potential peak upstroke is the maximum value of dV/dt between the action potential threshold and peak and the action potential peak downstroke is the minimum value of dV/dt between the action potential peak and trough (See Figure A.1).

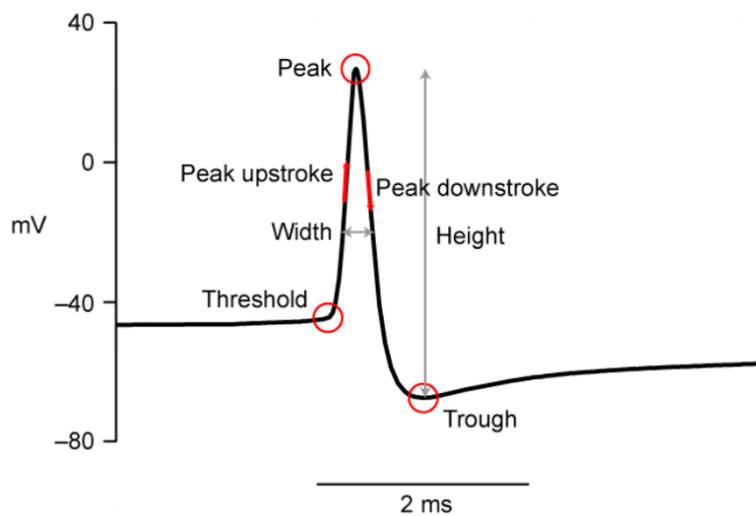


Figure A.1: Examples of single action potential features. Figure reproduced from Allen Institute for Brain Science (2016).

11. Threshold voltage long square: The threshold voltage defined as per feature 1 produced by a 'long square' 1 s current injection at the minimum amplitude which evoked reliably evoked action potentials.
12. Threshold time long square: Defined as per feature 2, with the injected current described in feature 11.
13. Threshold current long square (rheobase): Defined as per feature 3, with the injected current described in feature 11.
14. Peak voltage long square: Defined as per feature 4, with the injected current described in feature 11.
15. Peak time long square: Defined as per feature 5, with the injected current described in feature 11.
16. Trough voltage long square: Defined as per feature 6, with the injected current described in feature 11.
17. Trough time long square: Defined as per feature 7, with the injected current described in feature 11.
18. Fast trough voltage long square: Defined as per feature 8, with the injected current described in feature 11.
19. Fast trough time long square: Defined as per feature 9, with the injected current described in feature 11.

20. Upstroke/downstroke long square: Defined as per feature 10, with the injected current described in feature 11.
21. Threshold voltage ramp: The threshold voltage defined as per feature 1 produced by a ‘ramp’ injected current with an amplitude beginning at 0 pA and increasing linearly at 25 pA/sec, that is terminated after a series of action potentials are evoked.
22. Threshold time ramp: Defined as per feature 2, with the injected current described in feature 21.
23. Threshold current ramp: Defined as per feature 3, with the injected current described in feature 21.
24. Peak voltage ramp: Defined as per feature 4, with the injected current described in feature 21.
25. Trough voltage ramp: Defined as per feature 6, with the injected current described in feature 21.
26. Trough time ramp: Defined as per feature 7, with the injected current described in feature 21.
27. Fast trough voltage ramp: Defined as per feature 8, with the injected current described in feature 21.
28. Fast trough time ramp: Defined as per feature 9, with the injected current described in feature 21.
29. Upstroke/downstroke ratio ramp: Defined as per feature 10, with the injected current described in feature 21.
30. Voltage rest: The average pre-stimulus membrane potential across all of the sweeps using a long square input current pulse described in feature 11.
31. Input resistance: The slope of the linear fit to the plot of the minimum membrane potentials during the response of the neuron to a long square input current pulse of 1 s against the level of injected current, for a range of negative current amplitudes > -100 pA (Figure A.2A, left).
32. F-I curve slope: The F-I curve was found by plotting the average firing rate of the neuron to long square 1 s input current stimuli against the stimulus amplitude. The F-I curve slope was the slope of the straight line fit to the suprathreshold part of this curve (Figure A.2C).

33. Membrane time constant (τ): The voltage response curve of the neuron to an injected long square input current of 1 s was fit with an exponential curve between 10% of the maximum voltage deflection and the minimum membrane potential during the response. This was repeated for a range of input currents and the time constant of the fits to the multiple curves were averaged to find the overall membrane time constant (Figure A.2A, centre and right).
34. Sag: The voltage sag was calculated from the response of the neuron to an injected current that produced a minimum membrane potential closest to -100 mV. The sag was defined as

$$\text{Sag} = \frac{\text{Minimum resting potential} - \text{Steady state resting potential}}{\text{Peak deflection}}$$

during the stimulus (see Figure A.2B)

35. Membrane voltage for sag: The membrane potential at which the sag was calculated.
36. Latency: The time between the beginning of a stimulus and the first action potential evoked by the stimulus at the lowest long square stimulus amplitude that was between 40 pA–60 pA higher than rheobase.

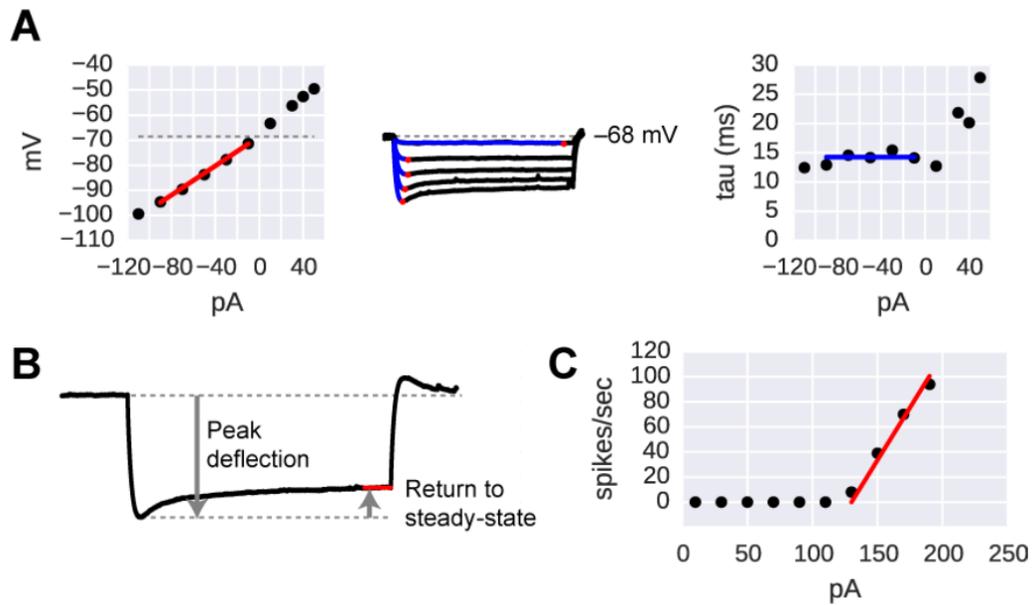


Figure A.2: **A** (left) Plot of minimum membrane potential values (shown as red dots in centre figure) against the level of injected current. Linear fit (red line on left plot) is used to estimate the input resistance of the neuron. Exponential fits are made to the current steps (blue lines on centre plot) and the time constants of these fits are averaged to find the membrane time constant (blue line on right plot). **B** Variables used to calculate the voltage sag. **C** F-I curve with the linear fit to the suprathreshold component that is used to calculate the F-I curve slope shown in red. Figure reproduced from Allen Institute for Brain Science (2016).

A.4 Cell types from multi-view clustering of the Allen Cell Types Database data set

A.4.1 Cell types from the RCCA method

Cluster	Cells per E/I type		Cells per layer						Cells per molecular marker			
	I	E	L1	L2/3	L4	L5	L6a	L6b	PV	5HT3aR	SST	Other
$\mathcal{C}_1^{MV,RC}$	0	29	0	3	25	1	0	0	–	–	–	–
$\mathcal{C}_2^{MV,RC}$	0	26	0	5	18	3	0	0	–	–	–	–
$\mathcal{C}_3^{MV,RC}$	1	11	0	2	7	3	0	0	–	–	–	–
$\mathcal{C}_4^{MV,RC}$	0	27	0	3	13	9	2	0	–	–	–	–
$\mathcal{C}_5^{MV,RC}$	0	18	0	2	12	3	0	1	–	–	–	–
$\mathcal{C}_6^{MV,RC}$	0	12	0	0	2	10	0	0	–	–	–	–
$\mathcal{C}_7^{MV,RC}$	0	7	0	0	2	3	2	0	–	–	–	–
$\mathcal{C}_8^{MV,RC}$	0	9	0	0	0	9	0	0	–	–	–	–
$\mathcal{C}_9^{MV,RC}$	1	21	0	0	1	6	15	0	–	–	–	–
$\mathcal{C}_{10}^{MV,RC}$	0	9	0	0	0	0	1	8	–	–	–	–
$\mathcal{C}_{11}^{MV,RC}$	9	0	0	0	0	7	2	0	7	1	0	1
$\mathcal{C}_{12}^{MV,RC}$	23	1	0	6	5	12	1	0	11	1	4	8
$\mathcal{C}_{13}^{MV,RC}$	18	0	0	1	1	10	6	0	10	1	5	2
$\mathcal{C}_{14}^{MV,RC}$	10	0	0	4	2	3	1	0	6	0	2	2
$\mathcal{C}_{15}^{MV,RC}$	6	0	1	5	0	0	0	0	0	6	0	0
$\mathcal{C}_{16}^{MV,RC}$	11	2	0	6	4	3	0	0	0	11	0	2
$\mathcal{C}_{17}^{MV,RC}$	6	0	0	0	2	4	0	0	0	1	5	0
$\mathcal{C}_{18}^{MV,RC}$	7	3	0	2	0	5	3	0	0	2	5	3
$\mathcal{C}_{19}^{MV,RC}$	13	0	0	3	3	7	0	0	0	3	9	1

Table A.4: Solution from clustering of the multi-view Allen Cell Types Database data set using the RCCA method. Table formatted as per Table 4.10.

A.4.2 Cell types from the TW- k -means method

Cluster	Cells per E/I type		Cells per layer						Cells per molecular marker			
	I	E	L1	L2/3	L4	L5	L6a	L6b	PV	5HT3aR	SST	Other
$\mathcal{C}_1^{MV,TW}$	0	6	0	2	2	0	0	2	–	–	–	–
$\mathcal{C}_2^{MV,TW}$	0	9	0	3	6	0	0	0	–	–	–	–
$\mathcal{C}_3^{MV,TW}$	4	6	0	4	5	0	1	0	–	–	–	–
$\mathcal{C}_4^{MV,TW}$	1	28	0	2	23	4	0	0	–	–	–	–
$\mathcal{C}_5^{MV,TW}$	0	28	0	1	23	3	1	0	–	–	–	–
$\mathcal{C}_6^{MV,TW}$	2	15	0	4	7	5	1	0	–	–	–	–
$\mathcal{C}_7^{MV,TW}$	0	9	0	0	7	0	0	2	–	–	–	–
$\mathcal{C}_8^{MV,TW}$	0	6	0	0	3	3	0	0	–	–	–	–
$\mathcal{C}_9^{MV,TW}$	0	9	0	0	3	6	0	0	–	–	–	–
$\mathcal{C}_{10}^{MV,TW}$	0	18	0	0	2	16	0	0	–	–	–	–
$\mathcal{C}_{11}^{MV,TW}$	1	5	0	0	1	4	1	0	–	–	–	–
$\mathcal{C}_{12}^{MV,TW}$	0	2	0	0	0	1	1	0	–	–	–	–
$\mathcal{C}_{13}^{MV,TW}$	1	3	0	0	0	4	0	0	–	–	–	–
$\mathcal{C}_{14}^{MV,TW}$	0	3	0	0	0	2	1	0	–	–	–	–
$\mathcal{C}_{15}^{MV,TW}$	0	8	0	0	1	1	6	0	–	–	–	–
$\mathcal{C}_{16}^{MV,TW}$	0	10	0	0	0	2	8	0	–	–	–	–
$\mathcal{C}_{17}^{MV,TW}$	0	1	0	0	0	0	1	0	–	–	–	–
$\mathcal{C}_{18}^{MV,TW}$	0	7	0	0	0	1	1	5	–	–	–	–
$\mathcal{C}_{19}^{MV,TW}$	15	0	0	2	5	5	3	0	10	1	2	2
$\mathcal{C}_{20}^{MV,TW}$	26	0	0	5	1	17	3	0	17	2	1	6
$\mathcal{C}_{21}^{MV,TW}$	14	0	0	1	2	7	4	0	7	0	4	3
$\mathcal{C}_{22}^{MV,TW}$	8	2	1	5	2	2	0	0	0	7	1	2
$\mathcal{C}_{23}^{MV,TW}$	8	0	0	8	0	0	0	0	0	5	3	0
$\mathcal{C}_{24}^{MV,TW}$	17	0	0	4	4	9	0	0	0	3	12	2
$\mathcal{C}_{25}^{MV,TW}$	8	0	0	1	0	6	1	0	0	1	6	1

Table A.5: Solution from clustering of the multi-view Allen Cell Types Database data set using the TW- k -means method. Table formatted as per Table 4.10.

Bibliography

- © 2015 Allen Institute for Brain Science. *Allen Brain Atlas API*. Available from: <http://brain-map.org/api/index.html>.
- © 2015 Allen Institute for Brain Science. *Allen Brain Atlas Software Development Kit*. Available from: <http://alleninstitute.github.io/AllenSDK/>.
- © 2015 Allen Institute for Brain Science. *Allen Cell Types Database*. Available from: <http://celltypes.brain-map.org>.
- Abeles, M., Bergman, H., Gat, I., Meilijson, I., Seidemann, E., Tishby, N., and Vaadia, E. (1995). “Cortical activity flips among quasi-stationary states”. *Proc. Natl. Acad. Sci. U.S.A.* 92, 8616–8620.
- Alberts, B., Johnson, A., Lewis, J., Raff, M., Roberts, K., and Walter, P. (2002). *Molecular Biology of the Cell*. 4th edition. New York: Garland Science.
- Ali, H. R., Rueda, O. M., Chin, S.-F., Curtis, C., Dunning, M. J., Aparicio, S. A., and Caldas, C. (2014). “Genome-driven integrated classification of breast cancer validated in over 7,500 samples”. *Genome Biol.* 15, 431.
- Aljadeff, J., Lansdell, B. J., Fairhall, A. L., and Kleinfeld, D. (2016). “Analysis of neuronal spike trains, deconstructed”. *Neuron* 91, 221–259.
- Allen Institute for Brain Science (2016). *Allen Cell Types Database Technical White Paper: Electrophysiology*. Tech. rep., 1–15.
- Allen Institute for Brain Science (2017a). *Allen Cell Types Database Technical White Paper: Morphology*. Tech. rep., 1–12.
- Allen Institute for Brain Science (2017b). *Allen Cell Types Database Technical White Paper: Overview*. Tech. rep., 1–10.
- Allen, C. and Stevens, C. F. (1994). “An evaluation of causes for unreliability of synaptic transmission”. *Proc. Natl. Acad. Sci. U.S.A.* 91, 10380–10383.
- Amarasingham, A. (2006). “Spike count reliability and the Poisson hypothesis”. *J. Neurosci.* 26, 801–809.
- Amin, H., Maccione, A., Marinaro, F., Zordan, S., Nieuw, T., and Berdondini, L. (2016). “Electrical responses and spontaneous activity of human iPS-derived neuronal networks characterized for 3-month culture with 4096-electrode arrays”. *Front. Neurosci.* 10, 1–15.
- Angelo, K., London, M., Christensen, S. R., and Häusser, M. (2007). “Local and global effects of I_h distribution in dendrites of mammalian neurons”. *J. Neurosci.* 27, 8643–8653.
- Armañanzas, R. and Ascoli, G. A. (2015). “Towards the automatic classification of neurons”. *Trends Neurosci.* 38, 307–318.

- Ascoli, G. A., Alonso-Nanclares, L., Anderson, S. A., Barrionuevo, G., Benavides-Piccione, R., Burkhalter, A., Buzsáki, G., Cauli, B., Defelipe, J., et al. (2008). “Petilla terminology: nomenclature of features of GABAergic interneurons of the cerebral cortex.” *Nat. Rev. Neurosci.* 9, 557–568.
- Averbeck, B. B. and Lee, D. (2003). “Neural noise and movement-related codes in the macaque supplementary motor area”. *J. Neurosci.* 23, 7630–7641.
- Awiszus, F. (1997). “Spike train analysis”. *J. Neurosci. Methods* 74, 155–166.
- Ayad, H. G. and Kamel, M. S. (2008). “Cumulative voting consensus method for partitions with variable number of clusters”. *IEEE Trans. Pattern Anal. Mach. Intell.* 30, 160–173.
- Azouz, R., Jensen, M. S., and Yaari, Y. (1996). “Ionic basis of spike after-depolarization and burst generation in adult rat hippocampal CA1 pyramidal cells”. *J. Physiol.* 492, 199–210.
- BRAIN Working Group (2014). *Brain Research through Advancing Innovative Neurotechnologies (BRAIN) Working Group Report to the Advisory Committee to the Director, NIH*. Tech. rep., 146.
- Baden, T., Berens, P., Bethge, M., and Euler, T. (2013). “Spikes in mammalian bipolar cells support temporal layering of the inner retina”. *Curr. Biol.* 23, 48–52.
- Bahl, A., Stemmler, M. B., Herz, A. V. M., and Roth, A. (2012). “Automated optimization of a reduced layer 5 pyramidal cell model based on experimental data”. *J. Neurosci. Methods* 210, 22–34.
- Baker, R. E., Corner, M. A., and van Pelt, J. (2006). “Spontaneous neuronal discharge patterns in developing organotypic mega-co-cultures of neonatal rat cerebral cortex”. *Brain Res.* 1101, 29–35.
- Bakkum, D. J., Shkolnik, A. C., Ben-Ary, G., Gamblen, P., DeMarse, B., and Potter, S. M. (2004). “Removing some ‘A’ from AI : Embodied cultured networks”. *Embodied Artif. Intell.* Ed. by F. Iida, R. Pfeifer, L. Steels, and Y. Kuniyoshi. Springer-Verlag, 130–146.
- Bakkum, D. J., Chao, Z. C., and Potter, S. M. (2008). “Long-term activity-dependent plasticity of action potential propagation delay and amplitude in cortical networks”. *PLoS One* 3, e2088.
- Bakkum, D. J., Radivojevic, M., Frey, U., Franke, F., Hierlemann, A., and Takahashi, H. (2013). “Parameters for burst detection”. *Front. Comput. Neurosci.* 7, 1–12.
- Bando, Y., Irie, K., Shimomura, T., Umeshima, H., Kushida, Y., Kengaku, M., Fujiyoshi, Y., Hirano, T., and Tagawa, Y. (2016). “Control of spontaneous Ca^{2+} transients is critical for neuronal maturation in the developing neocortex”. *Cereb. Cortex* 26, 106–117.
- Barbas, H. and Miguel, A (2015). “Motor cortex layer 4: less is more”. *Cell* 38, 259–261.
- Barbieri, R., Quirk, M. C., Frank, L. M., Wilson, M. A., and Brown, E. N. (2001). “Construction and analysis of non-Poisson stimulus-response models of neural spiking activity”. *J. Neurosci. Methods* 105, 25–37.
- Barrionuevo, G., Benoit, O., and Tempier, P. (1981). “Evidence for two types of firing pattern during the sleep-waking cycle in the reticular thalamic nucleus of the cat”. *Exp. Neurol.* 72, 486–501.

- Battaglia, D., Karagiannis, A., Gallopin, T., Gutch, H. W., and Cauli, B. (2013). “Beyond the frontiers of neuronal types”. *Front. Neural Circuits* 7, 13.
- Belgacem, Y. H. and Borodinsky, L. N. (2015). “Inversion of Sonic hedgehog action on its canonical pathway by electrical activity”. *Proc. Natl. Acad. Sci. U.S.A.* 112, 4140–4145.
- Ben-Ari, Y., Cherubini, E., Corradetti, R., and Gaiarsa, J. L. (1989). “Giant synaptic potentials in immature rat CA3 hippocampal neurones”. *J. Physiol.* 416, 303–325.
- Ben-Ari, Y., Gaiarsa, J.-L., Tyzio, R., and Khazipov, R. (2007). “GABA: A pioneer transmitter that excites immature neurons and generates primitive oscillations”. *Physiol. Rev.* 87, 1215–1284.
- Ben-Ari, Y. (2001). “Developing networks play a similar melody”. *Trends Neurosci.* 24, 353–360.
- Berdondini, L., Imfeld, K., Maccione, A., Tedesco, M., Neukom, S., Koudelka-Hep, M., and Martinoia, S. (2009). “Active pixel sensor array for high spatio-temporal resolution electrophysiological recordings from single cell to large scale neuronal networks”. *Lab Chip* 9, 2644–2651.
- Bickel, S. and Scheffer, T. (2004). “Multi-view clustering”. *Proc. IEEE Int. Conf. Data Min.* 19–26.
- Biel, M., Wahl-schott, C., Michalakis, S., and Zong, X. (2009). “Hyperpolarization-activated cation channels: From genes to function”. *Physiol. Rev.* 80, 847–885.
- Biffi, E., Regalia, G., Menegon, A., Ferrigno, G., and Pedrocchi, A. (2013). “The influence of neuronal density and maturation on network activity of hippocampal cell cultures: A methodological study”. *PLoS One* 8.
- Birtoli, B. and Ulrich, D. (2004). “Firing mode-dependent synaptic plasticity in rat neocortical pyramidal neurons”. *J. Neurosci.* 24, 4935–4940.
- Blankenship, A. G. and Feller, M. B. (2010). “Mechanisms underlying spontaneous patterned activity in developing neural circuits”. *Nat. Rev. Neurosci.* 11, 18–29.
- Blaschko, M. B. and Lampert, C. H. (2008). “Correlational spectral clustering”. *26th IEEE Conf. Comput. Vis. Pattern Recognition*, 1–8.
- Bocian, R. and Konopacki, J. (2004). “The effect of posterior hypothalamic injection of cholinergic agents on hippocampal formation theta in freely moving cat”. *Brain Res. Bull.* 63, 283–294.
- Bonetti, C. and Surace, E. M. (2010). “Mouse embryonic retina delivers information controlling cortical neurogenesis”. *PLoS One* 5, 6–11.
- Borghuis, B. G., Ratliff, C. P., Smith, R. G., Sterling, P., and Balasubramanian, V. (2008). “Design of a neuronal array”. *J. Neurosci.* 28, 3178–3189.
- Borodinsky, L. N., Root, C. M., Cronin, J. A., Sann, S. B., Gu, X., and Spitzer, N. C. (2004). “Activity-dependent homeostatic specification of transmitter expression in embryonic neurons”. *Nature* 429, 523–530.
- Borst, J. G. G. (2010). “The low synaptic release probability in vivo”. *Trends Neurosci.* 33, 259–266.
- Boser, B. E., Guyon, I. M., and Vapnik, V. N. (1992). “A training algorithm for optimal margin Classifiers”. *Proc. 5th Annu. ACM Work. Comput. Learn. Theory*, 144–152.

- Boyack, K. W., Mane, K., and Börner, K. (2004). “Mapping medline papers, genes and proteins related to melanoma research”. *Proc. Eighth Int. Conf. Inf. Vis.* 965–971.
- Branco, T. and Staras, K. (2009). “The probability of neurotransmitter release: variability and feedback control at single synapses”. *Nat. Rev. Neurosci.* 10, 373–383.
- Braun, J. V. and Müller, H.-G. (1998). “Statistical methods for DNA sequence segmentation”. *Stat. Sci.* 13, 142–162.
- Breckenridge, L. J., Wilson, R. J. A., Connolly, P., Curtis, A. S. G., Dow, J. A. T., Blackshaw, S. E., and Wilkinson, C. D. W. (1995). “Advantages of using microfabricated extracellular electrodes for in vitro neuronal recording”. *J. Neurosci. Res.* 42, 266–276.
- Breiman, L. (1996). “Bagging Predictors”. *Mach. Learn.* 24, 123–140.
- Brette, R. (2015). “Philosophy of the spike : Rate-based vs. spike-based theories of the brain”. *Front. Syst. Neurosci.* 9, 1–14.
- Brette, R. and Destexhe, A. (2012). “Intracellular recording”. *Handbook of Neural Activity Measurement*. Ed. by R. Brette and A. Destexhe. Cambridge, UK: Cambridge University Press, 44–91.
- Briggs, F. (2010). “Organizing principles of cortical layer 6”. *Front. Neural Circuits* 4, 1–8.
- Brown, E. N., Kass, R. E., and Frank, L. M. (2001). “The time-rescaling theorem and its application to neural spike train data analysis”. *Neural Comput.* 14, 325–346.
- Brown, E. N., Kass, R. E., and Mitra, P. P. (2004). “Multiple neural spike train data analysis: state-of-the-art and future challenges”. *Nat. Neurosci.* 7, 456–461.
- Brown, J. P., Hall, D., Frank, C. L., Wallace, K., Mundy, W. R., and Shafer, T. J. (2016). “Evaluation of a microelectrode array-based assay for neural network ontogeny using training set chemicals”. *Toxicol. Sci.* 154, 126–139.
- Bruno, E. and Marchand-Maillet, S. (2009). “Multiview clustering: A late fusion approach using latent models categories and subject descriptors”. *Proc. 32nd Int. ACM SIGIR Conf. Res. Dev. Inf. Retr.* 736–737.
- Burgos-Robles, A., Vidal-Gonzalez, I., Santini, E., and Quirk, G. J. (2007). “Consolidation of fear extinction requires NMDA receptor-dependent bursting in the ventromedial prefrontal cortex”. *Neuron* 53, 871–880.
- Butz, M., Wörgötter, F., and van Ooyen, A. (2009). “Activity-dependent structural plasticity”. *Brain Res. Rev.* 60, 287–305.
- Buzsáki, G. (2002). “Theta oscillations in the hippocampus”. *Neuron* 33, 325–340.
- Buzsáki, G. (2010). “Neural syntax: Cell assemblies, synapse ensembles and readers”. *Neuron* 68, 362–385.
- Buzsáki, G., Anastassiou, C. A., and Koch, C. (2012). “The origin of extracellular fields and currents - EEG, ECoG, LFP and spikes”. *Nat. Rev. Neurosci.* 13, 407–420.
- Cadwell, C. R., Palasantza, A., Jiang, X., Berens, P., Deng, Q., Yilmaz, M., Reimer, J., Shen, S., Bethge, M., et al. (2015). “Electrophysiological, transcriptomic and morphologic profiling of single neurons using Patch-seq”. *Nat. Biotechnol.* 34, 1–8.

- Cahoy, J. D., Emery, B., Kaushal, A., Foo, L. C., Zamanian, J. L., Christopherson, K. S., Xing, Y., Lubischer, J. L., Krieg, P. A., et al. (2004). “A transcriptome database for astrocytes, neurons, and oligodendrocytes: A new resource for understanding brain development and function”. *J. Neurosci.* 28, 264–278.
- Cai, X., Nie, F., Huang, H., and Kamangar, F. (2011). “Heterogeneous image feature integration via multi-modal spectral clustering”. *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.* 1977–1984.
- Cang, J., Rentería, R. C., Kaneko, M., Liu, X., Copenhagen, D. R., and Stryker, M. P. (2005). “Development of precise maps in visual cortex requires patterned spontaneous activity in the retina”. *Neuron* 48, 797–809.
- Caplan, J. B., Madsen, J. R., Schulze-Bonhage, A., Aschenbrenner-Scheibe, R., Newman, E. L., and Kahana, M. J. (2003). “Human theta oscillations related to sensorimotor integration and spatial learning”. *J. Neurosci.* 23, 4726–4736.
- Cappaert, N. L. M., Lopes Da Silva, F. H., and Wadman, W. J. (2009). “Spatio-temporal dynamics of theta oscillations in hippocampal-entorhinal slices”. *Hippocampus* 19, 1065–1077.
- Câteau, H. and Reyes, A. D. (2006). “Relation between single neuron and population spiking statistics and effects on network activity”. *Phys. Rev. Lett.* 96, 1–4.
- Cattaneo, A., Maffei, L., and Morrone, C. (1981). “Two firing patterns in the discharge of complex cells encoding different attributes of the visual stimulus”. *Exp. Brain Res.* 43, 115–118.
- Cauli, B., Audinat, E., Lambolez, B., Angulo, M. C., Ropert, N., Tsuzuki, K., Hestrin, S., and Rossier, J. (1997). “Molecular and physiological diversity of cortical nonpyramidal cells”. *J. Neurosci.* 17, 3894–3906.
- Chapman, B. and Stryker, M. P. (1993). “Development of orientation selectivity in ferret visual cortex and effects of deprivation”. *J. Neurosci.* 13, 5251–5262.
- Charlesworth, P., Cotterill, E., Morton, A., Grant, S. G., and Eglen, S. J. (2015). “Quantitative differences in developmental profiles of spontaneous activity in cortical and hippocampal cultures”. *Neural Dev.* 10, 1–10.
- Charlesworth, P., Morton, A., Eglen, S. J., Komiyama, N. H., and Grant, S. G. N. (2016). “Canalization of genetic and pharmacological perturbations in developing primary neuronal activity patterns”. *Neuropharmacology* 100, 47–55.
- Chaudhuri, K., Kakade, S. M., Livescu, K., and Sridharan, K. (2009). “Multi-view clustering via canonical correlation analysis”. *Proc. 26th Int. Conf. Mach. Learn.* 1–8.
- Chen, C., Chen, L., Lin, Y., Zeng, S., and Luo, Q. (2006). “The origin of spontaneous synchronized burst in cultured neuronal networks based on multi-electrode arrays”. *Biosystems.* 85, 137–143.
- Chen, L., Deng, Y., Luo, W., Wang, Z., and Zeng, S. (2009a). “Detection of bursts in neuronal spike trains by the mean inter-spike interval method”. *Prog. Nat. Sci.* 19, 229–235.
- Chen, S. and Yaari, Y. (2008). “Spike Ca^{2+} influx upmodulates the spike afterdepolarization and bursting via intracellular inhibition of K_V7/M channels.” *J. Physiol.* 586, 1351–63.

- Chen, X., Xu, X., Huang, J. Z., and Ye, Y. (2013). “TW-(k)-means: Automated two-level variable weighting clustering algorithm for multiview data”. *IEEE Trans. Knowl. Data Eng.* 25, 932–944.
- Chen, Y. and Ghosh, A. (2005). “Regulation of dendritic development by neuronal activity”. *J. Neurobiol.* 64, 4–10.
- Chen, Z., Vijayan, S., Barbieri, R., Wilson, M. A., and Brown, E. N. (2009b). “Discrete- and continuous-time probabilistic models and algorithms for inferring neuronal UP and DOWN states”. *Neural Comput.* 21, 1797–1862.
- Chiappalone, M., Novellino, A., Vajda, I., Vato, A., Martinoia, S., and van Pelt, J. (2005). “Burst detection algorithms for the analysis of spatio-temporal patterns in cortical networks of neurons”. *Neurocomputing* 65-66, 653–662.
- Chiappalone, M., Bove, M., Vato, A., Tedesco, M., and Martinoia, S. (2006). “Dissociated cortical networks show spontaneously correlated activity patterns during in vitro development”. *Brain Res.* 1093, 41–53.
- Chiappalone, M., Massobrio, P., and Martinoia, S. (2008). “Network plasticity in cortical assemblies”. *Eur. J. Neurosci.* 28, 221–237.
- Chiu, C. and Weliky, M. (2001). “Spontaneous activity in developing ferret visual cortex in vivo.” *J. Neurosci.* 21, 8906–8914.
- Chong, S.-A., Benilova, I., Shaban, H., Strooper, B. D., Devijver, H., Moechars, D., Eberle, W., Bartic, C., Leuven, F. V., and Callewaert, G. (2011). “Synaptic dysfunction in hippocampus of transgenic mouse models of Alzheimer’s disease : A multi-electrode array study”. *Neurobiol. Dis.* 44, 284–291.
- Chub, N. and O’Donovan, M. J. (1998). “Blockade and recovery of spontaneous rhythmic activity after application of neurotransmitter antagonists to spinal networks of the chick embryo”. *J. Neurosci.* 18, 294–306.
- Chung, S., Li, X., and Nelson, S. B. (2002). “Short-term depression at thalamocortical synapses contributes to rapid adaptation of cortical sensory responses in vivo”. *Neuron* 34, 437–446.
- Clause, A., Kim, G., Sonntag, M., Weisz, C. C., Vetter, D., Rubsamen, R., and Kandler, K. (2014). “The precise temporal pattern of prehearing spontaneous activity is necessary for tonotopic map refinement”. *Neuron* 82, 822–835.
- Clements, J. D. (1996). “Transmitter timecourse in the synaptic cleft: Its role in central synaptic function”. *Trends Neurosci.* 19, 163–171.
- Cleynen, A. and Lebarbier, E. (2014). “Segmentation of the Poisson and negative binomial rate models: a penalized estimator”. *ESAIM Probab. Stat.* 18, 750–769.
- Cleynen, A., Koskas, M., Lebarbier, E., Rigai, G., and Robin, S. (2014). “Segmentor3IsBack: an R package for the fast and exact segmentation of Seq-data.” *Algorithms Mol. Biol.* 9, 1–11.
- Cocatre-Zilgien, J. H. and Delcomyn, F. (1992). “Identification of bursts in spike trains”. *J. Neurosci. Methods* 41, 19–30.
- Cohen, L., Koffman, N., Meiri, H., Yarom, Y., Lampl, I., and Mizrahi, A. (2013). “Time-lapse electrical recordings of single neurons from the mouse neocortex”. *Proc. Natl. Acad. Sci. U.S.A.* 110, 5665–5670.
- Connors, B. W. and Gutnick, M. J. (1990). “Intrinsic firing patterns of diverse neocortical neurons”. *TINS* 13, 99–104.

- Cook, J. E. and Chalupa, L. M. (2000). “Retinal mosaics: new insights into an old concept”. *Trends Neurosci.* 23, 26–34.
- Corlew, R., Bosma, M. M., and Moody, W. J. (2004). “Spontaneous, synchronous electrical activity in neonatal mouse cortical neurones”. *J. Physiol.* 560, 377–390.
- Costa, M. R. and Müller, U. (2015). “Specification of excitatory neurons in the developing cerebral cortex: progenitor diversity and environmental influences”. *Front. Cell. Neurosci.* 8, 1–9.
- Costa, M., Manton, J. D., Ostrovsky, A. D., Prohaska, S., and Jefferis, G. S. X. E. (2016). “NBLAST: Rapid, sensitive comparison of neuronal structure and construction of neuron family databases”. *Neuron* 91, 293–311.
- Cotterill, E., Hall, D., Wallace, K., Mundy, W. R., Eglén, S. J., and Shafer, T. J. (2016a). “Characterization of early cortical neural network development in multiwell microelectrode array plates”. *J. Biomol. Screen.* 21, 510–519.
- Cotterill, E., Charlesworth, P., Thomas, C. W., Paulsen, O., and Eglén, S. J. (2016b). “A comparison of computational methods for detecting bursts in neuronal spike trains and their application to human stem cell-derived neuronal networks”. *J. Neurophysiol.* 116, 306–321.
- Cover, T. M. and Thomas, J. A. (2006). *Elements of Information Theory*. 2nd Edition. John Wiley and Sons, Inc.
- Crépel, V., Aronov, D., Jorquera, I., Represa, A., Ben-Ari, Y., and Cossart, R. (2007). “A parturition-associated nonsynaptic coherent activity pattern in the developing hippocampus”. *Neuron* 54, 105–120.
- Csicsvari, J., Hirase, H., Czurko, A., and Buzsáki, G. (1998). “Reliability and state dependence of pyramidal cell-interneuron synapses in the hippocampus: An ensemble approach in the behaving rat”. *Neuron* 21, 179–189.
- Cuntz, H., Forstner, F., Borst, A., and Häusser, M. (2010). “One rule to grow them all: A general theory of neuronal branching and its practical application”. *PLoS Comput. Biol.* 6, e1000877.
- Cutts, C. S. and Eglén, S. J. (2014). “Detecting pairwise correlations in spike trains: An objective comparison of methods and application to the study of retinal waves”. *J. Neurosci.* 34, 14288–14303.
- D’Angelo, E., Nieuws, T., Maffei, A., Armano, S., Rossi, P., Taglietti, V., Fontana, A., Naldi, G., and Angelo, E. D. (2001). “Theta-frequency bursting and resonance in cerebellar granule cells: experimental evidence and modeling of a slow K^+ -dependent mechanism”. *J. Neurosci.* 21, 759–770.
- Dayan, P. and Abbott, L. F. (2001). *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. Cambridge, MA: MIT Press.
- De Carlos, J. A. and Borrell, J. (2007). “A historical reflection of the contributions of Cajal and Golgi to the foundations of neuroscience”. *Brain Res. Rev.* 55, 8–16.
- DeFelipe, J., López-Cruz, P. L., Benavides-Piccione, R., Bielza, C., Larrañaga, P., Anderson, S., Burkhalter, A., Cauli, B., Fairén, A., et al. (2013). “New insights into the classification and nomenclature of cortical GABAergic interneurons”. *Nat. Rev. Neurosci.* 14, 202–216.
- DeWeese, M. R., Wehr, M. S., and Zador, A. M. (2003). “Binary spiking in auditory cortex”. *J. Neurosci.* 23, 7940–7949.

- Dehorter, N., Ciceri, G., Bartolini, G., Lim, L., Pino, I. del, and Marin, O. (2015). “Tuning of fast-spiking interneuron properties by an activity-dependent transcriptional switch”. *Science* 349, 1216–1220.
- Della Santina, L., Inman, D. M., Lupien, C. B., Horner, P. J., and Wong, R. O. L. (2013). “Differential progression of structural and functional alterations in distinct retinal ganglion cell types in a mouse model of glaucoma.” *J. Neurosci.* 33, 17444–57.
- Demas, J., Eglén, S. J., and Wong, R. O. L. (2003). “Developmental loss of synchronous spontaneous activity in the mouse retina is independent of visual experience”. *J. Neurosci.* 23, 2851–2860.
- Doan, T. N. and Kunze, D. L. (1999). “Contribution of the hyperpolarization-activated current to the resting membrane potential of rat nodose sensory neurons”. *J. Physiol.* 514, 125–138.
- Donohue, D. E. and Ascoli, G. A. (2008). “A comparative computer simulation of dendritic morphology”. *PLoS Comput. Biol.* 4, e1000089.
- Druckmann, S., Hill, S., Schürmann, F., Markram, H., and Segev, I. (2013). “A hierarchical structure of cortical interneuron electrical diversity revealed by automated statistical analysis”. *Cereb. Cortex* 23, 2994–3006.
- Dudoit, S. and Fridlyand, J. (2003). “Bagging to improve the accuracy of a clustering procedure”. *Bioinformatics* 19, 1090–1099.
- Dumitriu, D., Cossart, R., Huang, J., and Yuste, R. (2007). “Correlation between axonal morphologies and synaptic input kinetics of interneurons from mouse visual cortex”. *Cereb. Cortex* 17, 81–91.
- Egert, U., Heck, D., and Aertsen, A. (2002). “Two-dimensional monitoring of spiking networks in acute brain slices”. *Exp. Brain Res.* 142, 268–274.
- Eglén, S. J., Weeks, M., Jessop, M., Simonotto, J., Jackson, T., and Sernagor, E. (2014). “A data repository and analysis framework for spontaneous neural activity recordings in developing retina”. *Gigascience* 3.
- Eisenman, L. N., Emnett, C. M., Mohan, J., Zorumski, C. F., and Mennerick, S. (2015). “Quantification of bursting and synchrony in cultured hippocampal neurons”. *J. Neurophysiol.* 114, 1059–1071.
- Epsztein, J., Brecht, M., and Lee, A. K. (2011). “Intracellular determinants of hippocampal CA1 place and silent cell activity in a novel environment”. *Neuron* 70, 109–120.
- Evarts, E. V. (1964). “Temporal patterns of discharge of pyramidal tract neurons during sleep and waking in the monkey”. *J. Neurophysiol.* 27, 152–171.
- Eversmann, B., Jenkner, M., Hofmann, F., Paulus, C., Brederlow, R., Holzapfl, B., Fromherz, P., Merz, M., Brenner, M., et al. (2003). “A 128 x 128 CMOS biosensor array for extracellular recording of neural activity”. *IEEE J. Solid-State Circuits* 38, 2306–2317.
- Eytan, D. and Marom, S. (2006). “Dynamics and effective topology underlying synchronization in networks of cortical neurons”. *J. Neurosci.* 26, 8465–8476.
- Farrow, K., Masland, R. H., Farrow, K., and Masland, R. H. (2011). “Physiological clustering of visual channels in the mouse retina”. *J. Neurosci.* 31, 1516–1530.
- Fortune, E. S. and Rose, G. J. (2001). “Short-term synaptic plasticity as a temporal filter”. *Trends Neurosci.* 24, 381–385.

- Fred, A. L. N. and Jain, A. K. (2005). “Combining multiple clusterings using evidence accumulation”. *IEEE Trans. Pattern Anal. Mach. Intell.* 27, 835–850.
- Frey, U., Egert, U., Heer, F., Hafizovic, S., and Hierlemann, A. (2009). “Microelectronic system for high-resolution mapping of extracellular electric fields applied to brain slices”. *Biosens. Bioelectron.* 24, 2191–2198.
- Froemke, R. C., Ishan, A. T., Raad, M., Long, J. D., and Dan, Y. (2006). “Contribution of individual spikes in burst-induced long-term synaptic Modification”. *J. Neurophysiol.* 95, 1620–1629.
- Furukawa, S. and Middlebrooks, J. C. (2002). “Cortical representation of auditory space: Information-bearing features of spike patterns”. *J. Neurophysiol.* 87, 1749–1762.
- Fuzik, J., Zeisel, A., Máté, Z., Calvigioni, D., Yanagawa, Y., Szabó, G., Linnarsson, S., and Harkany, T. (2016). “Integration of electrophysiological recordings with single-cell RNA-seq data identifies neuronal subtypes”. *Nat. Biotechnol.* 34, 175–183.
- Gabbiani, F., Metzner, W., Wessel, R., and Koch, C. (1996). “From stimulus encoding to feature extraction in weakly electric fish”. *Nature* 384, 563–567.
- Gandolfo, M., Maccione, A., Tedesco, M., Martinoia, S., and Berondini, L. (2010). “Tracking burst patterns in hippocampal cultures with high-density CMOS-MEAs”. *J. Neural Eng.* 7, 1–16.
- Gao, B., Liu, T.-Y., Qin, T., Zheng, X., Cheng, Q.-S., and Ma, W.-Y. (2005). “Web image clustering by consistent utilization of visual features and surrounding texts”. *Proc. 13th Annu. ACM Int. Conf. Multimed.* 112–121.
- Garaschuk, O., Hanse, E., and Konnerth, A. (1998). “Developmental profile and synaptic origin of early network oscillations in the CA1 region of rat neonatal hippocampus”. *J. Physiol.* 507, 219–236.
- George, M. S., Abbott, L. F., and Siegelbaum, S. A. (2009). “HCN hyperpolarization-activated cation channels inhibit EPSPs by interactions with M-type K⁺ channels”. *Nat. Neurosci.* 12, 577–584.
- Gerstein, G. L. and Mandelbrot, B. (1964). “Random walk models for the spike activity of a single neuron.” *Biophys. J.* 4, 41–68.
- Gil, Z., Connors, B. W., and Amitai, Y. (1997). “Differential regulation of neocortical synapses by activity and neuromodulators”. *Neuron* 19, 679–686.
- Gilchrist, K. H., Lewis, G. F., Gay, E. A., Sellgren, K. L., and Grego, S. (2015). “High-throughput cardiac safety evaluation and multi-parameter arrhythmia profiling of cardiomyocytes using microelectrode arrays”. *Toxicol. Appl. Pharmacol.* 288, 249–257.
- Golbs, A., Nimmervoll, B., Sun, J.-J., Sava, I. E., and Luhmann, H. J. (2011). “Control of programmed cell death by distinct electrical activity patterns”. *Cereb. Cortex* 21, 1192–1202.
- Gourévitch, B. and Eggermont, J. J. (2007). “A nonparametric approach for detection of bursts in spike trains”. *J. Neurosci. Methods* 160, 349–358.
- Graham, B. P. and Ooyen, A. van (2006). “Mathematical modelling and numerical simulation of the morphological development of neurons.” *BMC Neurosci.* 7 (Suppl 1), S9.

- Grandjean, P. and Landrigan, P. J. (2014). “Neurobehavioural effects of developmental toxicity”. *Lancet Neurol.* 13, 330–338.
- Greene, D. and Cunningham, P. (2009). “A matrix factorization approach for integrating multiple data views”. *ECML PKDD Jt. Eur. Conf. Mach. Learn. Knowl. Discov. Databases*, 423–438.
- Greene, D. and Cunningham, P. (2006). “Practical solutions to the problem of diagonal dominance in kernel document clustering”. *Proc. 23rd Int. Conf. Mach. Learn.* 377–384.
- Gross, G. W., Rieske, E., Kreutzberg, G. W., and Meyer, A. (1977). “A new fixed-array multi-microelectrode system designed for long-term monitoring of extracellular single unit neuronal activity in vitro”. *Neurosci. Lett.* 6, 101–105.
- Guerra, L., McGarry, L. M., Robles, V., Bielza, C., Larrañaga, P., and Yuste, R. (2011). “Comparison between supervised and unsupervised classifications of neuronal cell types: A case study”. *Dev. Neurobiol.* 71, 71–82.
- Guo, Y. (2013). “Convex subspace representation learning from multi-view data”. *27th AAAI Conf. Artif. Intell.* 387–393.
- Gupta, A., Wang, Y., and Markram, H. (2000). “Organizing principles for a diversity of GABAergic interneurons and synapses in the neocortex”. *Science* 287, 273–278.
- Gur, M. and Snodderly, D. M. (2006). “High response reliability of neurons in primary visual cortex (V1) of alert, trained monkeys”. *Cereb. Cortex* 16, 888–895.
- Habets, A. M. M. C., Van Dongen, A. M. J., Van Huizen, F., and Corner, M. A. (1987). “Spontaneous neuronal firing patterns in fetal rat cortical networks during development in vitro: a quantitative analysis”. *Exp. Brain Res.* 69, 43–52.
- Hales, C. M., Rolston, J. D., and Potter, S. M. (2010). “How to culture, record and stimulate neuronal networks on micro-electrode arrays (MEAs)”. *J. Vis. Exp.* 1–7.
- Hamill, O. P., Marty, A., Neher, E., Sakmann, B., and Sigworth, F. J. (1981). “Improved patch-clamp techniques for high-resolution current recording from cells and cell-free membrane patches”. *Pflugers Arch. Eur. J. Physiol.* Vol. 391, 85–100.
- Hamming, R. W. (1950). “Error detecting and error correcting codes”. *Bell Syst. Tech. J.* 29, 147–160.
- Hanson, M. G. and Landmesser, L. T. (2003). “Characterization of the circuits that generate spontaneous episodes of activity in the early embryonic mouse spinal cord”. *J. Neurosci.* 23, 587–600.
- Hanson, M. G. and Landmesser, L. T. (2004). “Normal patterns of spontaneous activity are required for correct motor axon guidance and the expression of specific guidance molecules”. *Neuron* 43, 687–701.
- Harrill, J. A., Robinette, B. L., and Mundy, W. R. (2011). “Use of high content image analysis to detect chemical-induced changes in synaptogenesis in vitro”. *Toxicol. Vitr.* 25, 368–387.
- Harris, K. D. and Mrsic-Flogel, T. D. (2013). “Cortical connectivity and sensory coding”. *Nature* 503, 51–58.
- Harris, K. D. and Shepherd, G. M. G. (2015). “The neocortical circuit: themes and variations”. *Nat. Neurosci.* 18, 170–181.
- Harris, R. E., Coulombe, M. G., and Feller, M. B. (2002). “Dissociated retinal neurons form periodically active synaptic circuits”. *J. Neurophysiol.* 88, 188–195.

- Hay, E., Hill, S., Schürmann, F., Markram, H., and Segev, I. (2011). “Models of neocortical layer 5b pyramidal cells capturing a wide range of dendritic and perisomatic active properties”. *PLoS Comput. Biol.* 7, e1002107.
- He, D. and Parker, D. S. (2010). “Topic dynamics: An alternative model of ‘bursts’ in streams of topics”. *Proc. 16th ACM SIGKDD Int. Conf. Knowlweledge Discov. Data Min.* 443–452.
- Heck, N., Golbs, A., Riedemann, T., Sun, J.-J., Lessmann, V., and Luhmann, H. J. (2008). “Activity-dependent regulation of neuronal apoptosis in neonatal mouse cerebral cortex”. *Cereb. Cortex* 18, 1335–1349.
- Heikkilä, T. J., Ylä-Outinen, L., Tanskanen Jarno M A Lappalainen, R. S., Skottman, H., Suuronen, R., Mikkonen, J. E., Hyttinen, J. A. K., and Narkilahti, S. (2009). “Human embryonic stem cell-derived neuronal cells form spontaneously active neuronal networks in vitro”. *Exp. Neurol.* 218, 109–116.
- Heinricher, M. M. (2004). “Principles of extracellular single-unit recording”. *Microelectrode Recording in Movement Disorder Surgery*. Ed. by Z. Israel and K. J. Burchiel. 1st ed. Thieme, 8–13.
- Helm, J., Akgul, G., and Wollmuth, L. P. (2013). “Subgroups of parvalbumin-expressing interneurons in layers 2/3 of the visual cortex”. *J. Neurophysiol.* 109, 1600–1613.
- Helmstaedter, M., Sakmann, B., and Feldmeyer, D. (2009). “L2/3 interneuron groups defined by multiparameter analysis of axonal projection, dendritic geometry, and electrical excitability”. *Cereb. Cortex* 19, 951–962.
- Helmstaedter, M., Briggman, K. L., Turaga, S. C., Jain, V., Seung, H. S., and Denk, W. (2013). “Connectomic reconstruction of the inner plexiform layer in the mouse retina”. *Nature* 500, 168–174.
- Hennig, M. H., Grady, J., van Coppenhagen, J., and Sernagor, E. (2011). “Age-dependent homeostatic plasticity of GABAergic signaling in developing retinal networks”. *J. Neurosci.* 31, 12159–12164.
- Henze, D. A., Borhegyi, Z., Csicsvari, J., Mamiya, A., Harris, K. D., and Buzsáki, G. (2000). “Intracellular features predicted by extracellular recordings in the hippocampus in vivo”. *J. Neurophysiol.* 84, 390–400.
- Hodgkin, A. L. and Huxley, A. F. (1952). “A quantitative description of membrane current and its application to conduction and excitation in nerve”. *J. Physiol.* 117, 500–544.
- Hogberg, H. T., Sobanski, T., Novellino, A., Whelan, M., Weiss, D. G., and Bal-Price, A. K. (2011). “Application of micro-electrode arrays (MEAs) as an emerging technology for developmental neurotoxicity: Evaluation of domoic acid-induced effects in primary cultures of rat cortical neurons”. *Neurotoxicology* 32, 158–168.
- Holtmaat, A. and Caroni, P. (2016). “Functional and structural underpinnings of neuronal assembly formation in learning”. *Nat. Neurosci.* 19, 1553–1562.
- Hosp, J. A., Strüber, M., Yanagawa, Y., Obata, K., Vida, I., Jonas, P., and Bartos, M. (2014). “Morpho-physiological criteria divide dentate gyrus interneurons into classes”. *Hippocampus* 24, 189–203.

- Hu, H. and Agmon, A. (2016). “Differential excitation of distally versus proximally targeting cortical interneurons by unitary thalamocortical bursts”. *J. Neurosci.* 36, 6906–6916.
- Hu, H., Cavendish, J. Z., and Agmon, A. (2013). “Not all that glitters is gold: Off-target recombination in the somatostatinIRES-Cre mouse line labels a subset of fast-spiking interneurons”. *Front. Neural Circuits* 7, 1–4.
- Huberman, A. D., Speer, C. M., and Chapman, B. (2006). “Spontaneous retinal activity mediates development of ocular dominance columns and binocular receptive fields in V1”. *Neuron* 52, 247–254.
- Hubert, L. and Arabie, P. (1985). “Comparing partitions”. *J. Classification* 2, 193–218.
- Huettner, J. E. and Baughman, R. W. (1986). “Primary culture of identified neurons from the visual cortex of postnatal rats”. *J. Neurosci.* 6, 3044–3060.
- Hutcheon, B. and Yosef, Y. (2000). “Resonance, oscillation and the intrinsic frequency preferences of neurons”. *Trends Neurosci.* 23, 216–222.
- Ichikawa, M., Muramoto, K., Kobayashi, K., Kawahara, M., and Kuroda, Y. (1993). “Formation and maturation of synapses in primary cultures of rat cerebral cortical cells: an electron microscopic study”. *Neurosci. Res.* 16, 95–103.
- Illes, S., Fleischer, W., Siebler, M., Hartung, H.-P., and Dihn e, M. (2007). “Development and pharmacological modulation of embryonic stem cell-derived neuronal network activity”. *Exp. Neurol.* 207, 171–176.
- Inoue, H., Nagata, N., Kurokawa, H., and Yamanaka, S. (2014). “iPS cells: A game changer for future medicine”. *EMBO J.* 33, 409–417.
- Ito, D., Komatsu, T., and Gohara, K. (2013). “Measurement of saturation processes in glutamatergic and GABAergic synapse densities during long-term development of cultured rat cortical networks”. *Brain Res.* 1534, 22–32.
- Ito, S., Yeh, F.-C., Hiolski, E., Rydygier, P., Gunning, D. E., Hottowy, P., Timme, N., Litke, A. M., and Beggs, J. M. (2014). “Recording depicts functional network differences of cortical and hippocampal cultures”. *PLoS One* 9, e105324.
- Izhikevich, E. M. (2007). *Dynamical Systems in Neuroscience: The Geometry of Excitability and Bursting*. Cambridge, MA: MIT Press, 267–323.
- Izhikevich, E. M., Desai, N. S., Walcott, E. C., and Hoppensteadt, F. C. (2003). “Bursts as a unit of neural information: Selective communication via resonance”. *Trends Neurosci.* 26, 161–167.
- Jackson, M. E., Homayoun, H., and Moghaddam, B. (2004). “NMDA receptor hypofunction produces concomitant firing rate potentiation and burst activity reduction in the prefrontal cortex”. *Proc. Natl. Acad. Sci. U.S.A.* 101, 8467–8472.
- Jahnsen, B. Y. H. and Llinas, R. (1984). “Ionic basis for the electroresponsiveness and oscillatory properties of guinea-pig thalamic neurones in vitro”. *J. Physiol.* 349, 227–247.
- James, G., Witten, D., Hastie, T., and Tibshirani, R. (2014). *An Introduction to Statistical Learning: With Applications in R*. New York: Springer.
- Jasinski, P. E., Molkov, Y. I., Shevtsova, N. A., Smith, J. C., and Ryback, I. A. (2013). “Sodium and calcium mechanisms of rhythmic bursting in excitatory neural networks of the pre-B otzinger complex: A computational modelling study”. *Eur. J. Neurosci.* 37, 212–230.

- Jensen, M. S., Azouz, R., and Yaari, Y. (1994). "Variant firing patterns in rat hippocampal pyramidal cells modulated by extracellular potassium". *J. Neurophysiol.* 71, 831–839.
- Ji, D. and Wilson, M. A. (2007). "Coordinated memory replay in the visual cortex and hippocampus during sleep". *Nat. Neurosci.* 10, 100–107.
- Jiang, X., Shen, S., Cadwell, C. R., Berens, P., Sinz, F., Ecker, A. S., Patel, S., and Tolias, A. S. (2015). "Principles of connectivity among morphologically defined cell types in adult neocortex". *Science* 350, 1055–1065.
- Jimbo, Y., Tateno, T., and Robinson, H. P. C. (1999). "Simultaneous induction of pathway-specific potentiation and depression in networks of cortical neurons". *Biophys. J.* 76, 670–678.
- Johnstone, A. F. M., Gross, G. W., Weiss, D. G., Schroeder, O. H.-U., Gramowski, A., and Shafer, T. J. (2010). "Microelectrode arrays: a physiologically based neurotoxicity testing platform for the 21st century". *Neurotoxicology* 31, 331–350.
- Kamioka, H., Maeda, E., Jimbo, Y., Robinson, H. P. C., and Kawana, A. (1996). "Spontaneous periodic synchronized bursting during formation of mature patterns of connections in cortical cultures". *Neurosci. Lett.* 206, 109–112.
- Kandel, E. R., Schwartz, J. H., and Jessell, T. M. (2000). *Principles of Neural Science*. 4th Edition. New York: McGraw-Hill, Health Professions Division.
- Kaneoke, Y. and Vitek, J. L. (1996). "Burst and oscillation as disparate neuronal properties". *J. Neurosci. Methods* 68, 211–223.
- Kapucu, F. E., Tanskanen, J. M. A., Mikkonen, J. E., Ylä-Outinen, L., Narkilahti, S., and Hyttinen, J. A. K. (2012). "Burst analysis tool for developing neuronal networks exhibiting highly varying action potential dynamics". *Front. Comput. Neurosci.* 6, 38.
- Karagiannis, A., Gallopin, T., Dávid, C., Battaglia, D., Geoffroy, H., Rossier, J., Hillman, E. M. C., Staiger, J. F., and Cauli, B. (2009). "Classification of NPY-expressing neocortical interneurons". *J. Neurosci.* 29, 3642–3659.
- Karube, F., Kubota, Y., and Kawaguchi, Y. (2004). "Axon branching and synaptic bouton phenotypes in GABAergic nonpyramidal cell subtypes". *J. Neurosci.* 24, 2853–2865.
- Kass, R. E. and Ventura, V. (2001). "A spike-train probability model". *Neural Comput.* 13, 1713–1720.
- Kastanenka, K. V. and Landmesser, L. T. (2013). "Optogenetic-mediated increases in in vivo spontaneous activity disrupt pool-specific but not dorsal-ventral motoneuron path finding". *Proc. Natl. Acad. Sci. U.S.A.* 110, 17528–17533.
- Kaufman, L. and Rousseeuw, P. J. (1990). "Partitioning around medoids (Program PAM)". *Finding Groups in Data: An Introduction to Cluster Analysis*. Hoboken, NJ: John Wiley & Sons, Inc, 68–125.
- Kawaguchi, Y. and Kubota, Y. (1997). "GABAergic cell subtypes and their synaptic connections in rat frontal cortex". *Cereb. Cortex* 7, 476–486.
- Kelsom, C. and Lu, W. (2013). "Development and specification of GABAergic cortical interneurons". *Cell Biosci.* 3, 1–19.
- Kepecs, A. and Fishell, G. (2014). "Interneuron cell types are fit to function". *Nature* 505, 318–326.

- Kerschensteiner, D. (2014). “Spontaneous network activity and synaptic development”. *Neurosci.* 20, 272–290.
- Khazipov, R. and Luhmann, H. J. (2006). “Early patterns of electrical activity in the developing cerebral cortex of humans and rodents”. *Trends Neurosci.* 29, 414–418.
- Kirwan, P., Turner-Bridger, B., Peter, M., Momoh, A., Arambepola, D., Robinson, H. P. C., and Livesey, F. J. (2015). “Development and function of human cerebral cortex neural networks from pluripotent stem cells in vitro”. *Development* 142, 3178–3187.
- Kleinberg, J. (2002). “Bursty and hierarchical structure in streams”. *Proc. 8th ACM SIGKDD Int. Conf. Knowl. Discov. Data Min.* 91–101.
- Ko, D., Wilson, C. J., Lobb, C. J., and Paladini, C. A. (2012). “Detection of bursts and pauses in spike trains”. *J. Neurosci. Methods* 211, 145–158.
- Kolda, T. and Bader, B. (2006). “The TOPHITS model for higher-order web link analysis”. *Work. Link Anal. Counterterrorism Secur.*
- Komuro, H. and Kumada, T. (2005). “Ca²⁺ transients control CNS neuronal migration”. *Cell Calcium* 37, 387–393.
- Konopacki, J., Eckersdorf, B., Kowalczyk, T., and Gołbiewski, H. (2006). “Firing cell repertoire during carbachol-induced theta rhythm in rat hippocampal formation slices”. *Eur. J. Neurosci.* 23, 1811–1818.
- Krahe, R. and Gabbiani, F. (2004). “Burst firing in sensory systems”. *Nat. Rev. Neurosci.* 5, 13–23.
- Krahe, R., Kreiman, G., Gabbiani, F., Koch, C., and Metzner, W. (2002). “Stimulus encoding and feature extraction by multiple sensory neurons”. *J. Neurosci.* 22, 2374–2382.
- Kubota, Y., Shigematsu, N., Karube, F., Sekigawa, A., Kato, S., Yamaguchi, N., Hirai, Y., Morishima, M., and Kawaguchi, Y. (2011). “Selective coexpression of multiple chemical markers defines discrete populations of neocortical GABAergic neurons”. *Cerebral Cortex* 21, 1803–1817.
- Kuijlaars, J., Oyelami, T., Diels, A., Rohrbacher, J., Versweyveld, S., Meneghello, G., Tuefferd, M., Verstraelen, P., Detrez, J. R., et al. (2016). “Sustained synchronized neuronal network activity in a human astrocyte co-culture system”. *Sci. Rep.* 6, 1–14.
- Kumar, A. and Daumé, H. (2011). “A co-training approach for multi-view spectral clustering”. *Proc. 28th Int. Conf. Mach. Learn.* 393–400.
- Kumar, A., Rai, P., and Daumé III, H. (2011). “Co-regularized multi-view spectral clustering”. *Neural Inf. Process. Syst.* 1413–1421.
- Kumar, R., Road, H., Jose, S., Road, H., Jose, S., Drive, R., Tomkins, A., and Jose, S. (2003). “On the Bursty Evolution of Blogspace”. *Int. World Wide Web Conf.* 568–576.
- Lagali, P. S., Balya, D., Awatramani, G. B., Münch, T. A., Kim, D. S., Busskamp, V., Cepko, C. L., and Roska, B. (2008). “Light-activated channels targeted to ON bipolar cells restore visual function in retinal degeneration”. *Nat. Neurosci.* 11, 667–675.
- Lampl, I. and Yarom, Y. (1997). “Subthreshold oscillations and resonant behavior: Two manifestations of the same mechanism”. *Neuroscience* 78, 325–341.

- Lee, S., Hjerling-Leffler, J., Zaghera, E., Fishell, G., and Rudy, B. (2010). “The largest group of superficial neocortical GABAergic interneurons expresses ionotropic serotonin receptors”. *J. Neurosci.* 30, 16796–16808.
- Legéndy, C. R. and Salcman, M. (1985). “Bursts and recurrences of bursts in the spike trains of spontaneously active striate cortex neurons”. *J. Neurophysiol.* 53, 926–939.
- Leinekugel, X., Medina, I., Khalilov, I., Ben-Ari, Y., and Khazipov, R. (1997). “Ca²⁺ oscillations mediated by the synergistic excitatory actions of GABA_A and NMDA receptors in the neonatal hippocampus”. *Neuron* 18, 243–255.
- Leinekugel, X., Khazipov, R., Cannon, R., Hirase, H., Ben-Ari, Y., and Buzsáki, G. (2002). “Correlated bursts of activity in the neonatal hippocampus in vivo”. *Science* 296, 2049–2052.
- Lemon, N. and Turner, R. W. (2000). “Conditional spike backpropagation generates burst discharge in a sensory neuron”. *J. Neurophysiol.* 84, 1519–1530.
- Leondopulos, S. S., Boehler, M. D., Wheeler, B. C., and Brewer, G. J. (2012). “Chronic stimulation of cultured neuronal networks boosts low-frequency oscillatory activity at theta and gamma with spikes phase-locked to gamma frequencies”. *J. Neural Eng.* 9, 1–12.
- Leung, L. S. and Yu, H. W. (1998). “Theta-frequency resonance in hippocampal CA1 neurons in vitro demonstrated by sinusoidal current injection”. *J. Neurophysiol.* 79, 1592–1596.
- Li, C.-Y. T., Poo, M.-M., and Dan, Y. (2009). “Burst spiking of a single cortical neuron modifies global brain state”. *Science* 324, 643–646.
- Li, S. Y., Jiang, Y, and Zhou, Z. H. (2014). “Partial multi-view clustering”. *AAAI Conf. Artif. Intell.* 1968–1974.
- Li, T. and Ding, C (2008). “Weighted consensus clustering”. *Proc. Eighth SIAM Int. Conf. Data Min.* 798–809.
- Li, W., Doyon, W. M., and Dani, J. A. (2012). “Quantitative unit classification of ventral tegmental area neurons in vivo”. *J. Neurophysiol.* 107, 2808–2820.
- Light, A. C., Zhu, Y., Shi, J., Sazik, S., Lindstrom, S., Davidsob, L., Li, X., Chiodo, V. A., Hauswirth, W. W., et al. (2012). “Organizational motifs for ground squirrel cone bipolar cells”. *J. Comp. Neurol.* 520, 2864–2887.
- Lisman, J. E. (1997). “Bursts as a unit of neural information: Making unreliable synapses reliable”. *Trends Neurosci.* 20, 38–43.
- Liu, X., Yu, S., Moreau, Y., Moor, B. D., Glänzel, W., and Janssens, F. (2009). “Hybrid clustering of text mining and bibliometrics applied to journal sets”. *SIAM Int. Conf. Data Min.* 49–60.
- Liu, X., Ji, S., Glenzel, W., and De Moor, B. (2013). “Multiview partitioning via tensor methods”. *IEEE Trans. Knowl. Data Eng.* 25, 1056–1069.
- Livingstone, M. S., Freeman, D. C., and Hubel, D. H. (1996). “Visual responses in V1 of freely viewing monkeys”. *Cold Spring Harb. Symp. Quant. Biol.* 61, 27–37.
- Lobb, C. J. (2014). “Abnormal bursting as a pathophysiological mechanism in Parkinson’s disease”. *Basal Ganglia* 3, 187–195.
- Lodish, H., Berk, A., Matsudaira, P., Kaiser, C. A., Krieger, M., Scott, M. P., Zipursky, L., and Darnell, J. (2003). *Molecular Cell Biology*. 5th edition. New York: W. H. Freeman.

- Lonardoni, D., Di Marco, S., Amin, H., Maccione, A., Berdondini, L., and Nieuw, T. (2015). “High-density MEA recordings unveil the dynamics of bursting events in Cell Cultures”. *Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.* 3763–3766.
- López-Cruz, P. L., Larrañaga, P., DeFelipe, J., and Bielza, C. (2014). “Bayesian network modeling of the consensus between experts: An application to neuron classification”. *Int. J. Approx. Reason.* 55, 3–22.
- Lopez-Paz, D., Sra, S., Smola, A. J., Ghahramani, Z., and Schölkopf, B. (2014). “Randomized nonlinear component analysis”. *Proc. Int. Conf. Mach. Learn.* 1359–1367.
- Loturco, J. J., Owens, D. F., Mark, J. S., Davis, M. E., and Kriegsteing, A. R. (1995). “GABA and glutamate depolarize cortical progenitor cells and inhibit DNA synthesis”. *Neuron* 15, 1287–1298.
- Lu, Y., Carin, L., Coifman, R., Shain, W., and Roysam, B. (2015). “Quantitative arbor analytics: Unsupervised harmonic co-clustering of populations of brain cell arbors based on L-Measure”. *Neuroinformatics* 13, 47–63.
- Luhmann, H. J., Fukuda, A., and Kilb, W. (2015). “Control of cortical neuronal migration by glutamate and GABA”. *Front. Cell. Neurosci.* 9, 1–15.
- Luhmann, H. J., Sinning, A., Yang, J.-W., Reyes-Puerta, V., Stüttgen, M. C., Kirischuk, S., and Kilb, W. (2016). “Spontaneous neuronal activity in developing neocortical networks: From single cells to large-scale interactions”. *Front. Neural Circuits* 10, 1–14.
- Ma, Y., Hu, H., Berrebi, A. S., Mathers, P. H., and Agmon, A. (2006). “Distinct subtypes of somatostatin-containing neocortical interneurons revealed in transgenic mice”. *J. Neurosci.* 26, 5069–5082.
- MacLaren, E. and Charlesworth, P. (2011). “Knockdown of mental disorder susceptibility genes disrupts neuronal network physiology in vitro”. *Mol. Cell. ...* 47, 93–99.
- Maccione, A., Gandolfo, M., Tedesco, M., Nieuw, T., and Imfeld, K. (2010). “Experimental investigation on spontaneously active hippocampal cultures recorded by means of high-density MEAs: analysis of the spatial resolution effects”. *Front. Neuroeng.* 3, 1–12.
- Maccione, A., Garofalo, M., Nieuw, T., Tedesco, M., Berdondini, L., and Martinoia, S. (2012). “Multiscale functional connectivity estimation on low-density neuronal cultures recorded by high-density CMOS Micro Electrode Arrays”. *J. Neurosci. Methods* 207, 161–171.
- Maccione, A., Hennig, M. H., Gandolfo, M., Muthmann, O., van Coppenhagen, J., Eglen, S. J., Berdondini, L., and Sernagor, E. (2014). “Following the ontogeny of retinal waves: pan-retinal recordings of population dynamics in the neonatal mouse”. *J. Physiol.* 592, 1545–1563.
- Mack, C. M., Lin, B. J., Turner, J. D., Johnstone, A. F. M., Burgoon, L. D., and Shafer, T. J. (2014). “Burst and principal components analyses of MEA data for 16 chemicals describe at least three effects classes”. *Neurotoxicology* 40, 75–85.
- Maeda, E., Robinson, H. P., and Kawana, A. (1995). “The mechanisms of generation and propagation of synchronized bursting in developing networks of cortical neurons”. *J. Neurosci.* 15, 6834–6845.

- Mainen, Z. F. and Sejnowski, T. J. (1996). “Influence of dendritic structure on firing pattern in model neocortical neurons”. *Nature* 382, 363–382.
- Mane, K. K. and Börner, K. (2004). “Mapping topics and topic bursts in PNAS”. *Proc. Natl. Acad. Sci. U.S.A.* Vol. 101, 5287–5290.
- Mao, B. Q., Hamzei-Sichani, F., Aronov, D., Froemke, R. C., and Yuste, R. (2001). “Dynamics of spontaneous activity in neocortical slices”. *Neuron* 32, 883–898.
- Markram, H., Wang, Y., and Tsodyks, M. (1998). “Differential signaling via the same axon of neocortical pyramidal neurons”. *Proc. Natl. Acad. Sci. U.S.A.* 95, 5323–8.
- Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G., and Wu, C. (2004). “Interneurons of the neocortical inhibitory system”. *Nat. Rev. Neurosci.* 5, 793–807.
- Marom, S. and Shahaf, G. (2002). “Development, learning and memory in large random networks of cortical neurons: Lessons beyond anatomy”. *Q. Rev. Biophys.* 35, 63–87.
- Martens, M. B., Chiappalone, M., Schubert, D., and Tiesinga, P. H. E. (2014). “Separating burst from background spikes in multichannel neuronal recordings using return map analysis”. *Int. J. Neural Syst.* 24, 1450012.
- Martinez-Conde, S., Macknik, S. L., and Hubel, D. H. (2002). “The function of bursts of spikes during visual fixation in the awake primate lateral geniculate nucleus and primary visual cortex”. *Proc. Natl. Acad. Sci. U.S.A.* 99, 13920–13925.
- Martinson, J., Webster, H. H., Myasnikov, A. A., and Dykes, R. W. (1997). “Recognition of temporally structured activity in spontaneously discharging neurons in the somatosensory cortex in waking cats”. *Brain Res.* 750, 129–140.
- Marx, M. and Feldmeyer, D. (2013). “Morphology and physiology of excitatory neurons in layer 6b of the somatosensory rat barrel cortex”. *Cereb. Cortex* 23, 2803–2817.
- Masland, R. H. (2012). “The neuronal organization of the retina”. *Neuron* 76, 266–280.
- Mazzoni, A., Broccard, F. D., Garcia-Perez, E., Bonifazi, P., Ruaro, M. E., and Torre, V. (2007). “On the dynamics of the spontaneous activity in neuronal networks”. *PLoS One* 2, e439.
- McCarley, R. W., Benoit, O., and Barrionuevo, G. (1983). “Lateral geniculate nucleus unitary discharge in sleep and waking: State- and rate-specific aspects”. *J. Neurophysiol.* 50, 798–818.
- McCormick, D. A. (1991). “Functional properties of a slowly inactivating potassium current in guinea pig dorsal lateral geniculate relay neurons”. *J. Neurophysiol.* 66, 1176–1189.
- McGarry, L. M., Packer, A. M., Fino, E., Nikolenko, V., Sippy, T., and Yuste, R. (2010). “Quantitative classification of somatostatin-positive neocortical interneurons identifies three interneuron subtypes”. *Front. Neural Circuits* 4, 1–19.
- McMahon, D. B. T., Bondar, I. V., Afuwape, O. A. T., Ide, D. C., and Leopold, D. A. (2014). “One month in the life of a neuron: Longitudinal single-unit electrophysiology in the monkey visual system”. *J. Neurophysiol.* 112, 1748–1762.

- Mease, R. A., Kuner, T., Fairhall, A. L., and Groh, A. (2017). “Multiplexed spike coding and adaptation in the thalamus”. *Cell Rep.* 19, 1130–1140.
- Meister, M., Wong, R. O. L., Baylor, D. A., and Shatz, C. J. (1991). “Synchronous bursts of action potentials in ganglion cells of the developing mammalian retina”. *Science* 252, 939–943.
- Metz, A. E., Jarsky, T., Martina, M., and Spruston, N. (2005). “R-type calcium channels contribute to afterdepolarization and bursting in hippocampal CA1 pyramidal neurons”. *J. Neurosci.* 25, 5763–5773.
- Mihaljević, B., Benavides-Piccione, R., Bielza, C., DeFelipe, J., and Larrañaga, P. (2015). “Bayesian network classifiers for categorizing cortical GABAergic interneurons”. *Neuroinformatics* 13, 193–208.
- Miller, B. R., Walker, A. G., Barton, S. J., and Rebec, G. V. (2011). “Dysregulated neuronal activity patterns implicate corticostriatal circuit dysfunction in multiple rodent models of Huntington’s disease”. *Front. Syst. Neurosci.* 5, 26.
- Miyoshi, G., Hjerling-leffler, J., Karayannis, T., Sousa, V. H., Butt, S. J. B., Battiste, J., Johnson, J. E., Machold, R. P., and Fishell, G. (2010). “Genetic fate mapping reveals that the caudal ganglionic eminence produces a large and diverse population of superficial cortical interneurons”. *J. Neurosci.* 30, 1582–1594.
- Mo, Q., Wang, S., Seshan, V. E., Olshen, A. B., Schultz, N., Sander, C., Powers, R. S., Ladanyi, M., and Shen, R. (2013). “Pattern discovery and cancer gene identification in integrated cancer genomic data”. *Proc. Natl. Acad. Sci. U.S.A.* 110, 4245–4250.
- Molnár, Z. and Cheung, A. F. P. (2006). “Towards the classification of subpopulations of layer V pyramidal projection neurons”. *Neurosci. Res.* 55, 105–115.
- Molyneaux, B. J., Arlotta, P., Menezes, J. R. L., and Macklis, J. D. (2007). “Neuronal subtype specification in the cerebral cortex”. *Nature* 8, 427–437.
- Molyneaux, B. J., Arlotta, P., Fame, R. M., Macdonald, J. L., MacQuarrie, K. L., and Macklis, J. D. (2009). “Novel subtype-specific genes identify distinct subpopulations of callosal projection neurons”. *J. Neurosci.* 29, 12343–12354.
- Moore, A. R., Zhou, W.-L., Jakovcevski, I., Zecevic, N., and Antic, S. D. (2011). “Spontaneous electrical activity in the human fetal cortex in vitro”. *J. Neurosci.* 31, 2391–2398.
- Mrsic-Flogel, T. D., Hofer, S. B., Creutzfeldt, C., Cloëz-Tayarani, I., Changeux, J.-P., Bonhoeffer, T., and Hübener, M. (2005). “Altered map of visual space in the superior colliculus of mice lacking early retinal waves”. *J. Neurosci.* 25, 6921–6928.
- Multi Channel Systems MCS GmbH (2015). *Extracellular recording with microelectrode arrays for all applications Overview MEA-System*. Tech. rep., 1–16.
- Muramoto, K., Ichikawa, M., Kawahara, M., Kobayashi, K., and Kuroda, Y. (1993). “Frequency of synchronous oscillations of neuronal activity increases during development and is correlated to the number of synapses in cultured cortical neuron networks”. *Neurosci. Lett.* 163, 163–165.
- Nádasy, Z., Hirase, H., Czurkó, A., Csicsvari, J., and Buzsáki, G. (1999). “Replay and time compression of recurring spike sequences in the hippocampus”. *J. Neurosci.* 19, 9497–9507.

- Nakanishi, K. and Kukita, F. (2000). "Intracellular $[Cl^-]$ modulates synchronous electrical activity in rat neocortical neurons in culture by way of GABAergic inputs". *Brain Res.* 863, 192–204.
- Nassar, M., Simonnet, J., Lofredi, R., Cohen, I., Savary, E., Yanagawa, T., Miles, R., and Fricker, D. (2015). "Diversity and overlap of parvalbumin and somatostatin expressing interneurons in mouse presubiculum". *Front. Neural Circuits* 9, 1–19.
- Nelken, I. and Chechik, G. A. L. (2005). "Encoding stimulus information by spike numbers and mean response time in primary auditory cortex". *J. Comp. Neurol.* 19, 199–221.
- Nex Technologies (2014). *NeuroExplorer Manual*. Nex Technologies, 1–371.
- Ni, Z. G., Bouali-Benazzouz, R., Gao, D. M., Benabid, A. L., and Benazzouz, A. (2001). "Time-course of changes in firing rates and firing patterns of subthalamic nucleus neuronal activity after 6-OHDA-induced dopamine depletion in rats". *Brain Res.* 899, 142–147.
- Nicolas, J., Hendriksen, P. J. M., van Kleef, R. G. D. M., de Groot, A., Bovee, T. F. H., Rietjens, I. M. C. M., and Westerink, R. H. S. (2014). "Detection of marine neurotoxins in food safety testing using a multielectrode array". *Mol. Nutr. Food Res.* 58, 2369–2378.
- O'Keefe, J. and Recce, M. L. (1993). "Phase relationship between hippocampal place units and the EEG theta rhythm". *Hippocampus* 3, 317–330.
- Oberlaender, M., Kock, C. P. J. D., Bruno, R. M., Ramirez, A., Meyer, H. S., Derksen, V. J., Helmstaedter, M., and Sakmann, B. (2012). "Cell type-specific three-dimensional structure of thalamocortical circuits in a column of rat vibrissal cortex". *Cereb. Cortex* 22, 2375–2391.
- Obien, M. E. J., Deligkaris, K., Bullmann, T., Bakkum, D. J., and Frey, U. (2015). "Revealing neuronal function through microelectrode array recordings". *Front. Neurosci.* 9, 423.
- Odawara, A., Katoh, H., Matsuda, N., and Suzuki, I. (2016). "Physiological maturation and drug responses of human induced pluripotent stem cell-derived cortical neuronal networks in long-term culture". *Sci. Rep.* 6, 1–14.
- Okamoto, K., Ishikawa, T., Abe, R., Ishikawa, D., Kobayashi, C., Mizunuma, M., Norimoto, H., Matsuki, N., and Ikegaya, Y. (2014). "Ex vivo cultured neuronal networks emit in vivo-like spontaneous activity". *J. Physiol. Sci.* 64, 421–431.
- Otto, T., Eichenbaum, H., Wible, C. G., and Wiener, S. I. (1991). "Learning-related patterns of CA1 spike trains parallel stimulation parameters optimal for inducing hippocampal longterm potentiation". *Hippocampus* 1, 181–192.
- Pape, H.-C. (1996). "Queer current and pacemaker: The hyperpolarization-activated cation current in neurons". *Annu. Rev. Physiol.* 58, 299–327.
- Pasquale, V., Massobrio, P., Bologna, L. L., Chiappalone, M., and Martinoia, S. (2008). "Self-organization and neuronal avalanches in networks of dissociated cortical neurons". *Neuroscience* 153, 1354–1369.
- Pasquale, V., Martinoia, S., and Chiappalone, M. (2010). "A self-adapting approach for the detection of bursts and network bursts in neuronal cultures". *J. Comput. Neurosci.* 29, 213–229.
- Paulsen, O. and Sejnowski, T. J. (2000). "Natural patterns of activity and long-term synaptic plasticity". *Curr. Opin. Neurobiol.* 10, 172–179.

- Perin, R., Berger, T. K., and Markram, H. (2011). “A synaptic organizing principle for cortical neuronal groups”. *Proc. Natl. Acad. Sci. U.S.A.* 108, 5419–5424.
- Perkel, D. H., Gerstein, G. L., and Moore, G. P. (1967a). “Neuronal spike trains and stochastic point processes I. The Single Spike Train”. *Biophys. J.* 7, 391–418.
- Perkel, D. H., Gerstein, G. L., and Moore, G. P. (1967b). “Neuronal spike trains and stochastic point processes II. Simultaneous Spike Trains”. *Biophys. J.* 7, 419–440.
- Perrenoud, Q., Geoffroy, H., Gauthier, B., Rancillac, A., Alfonsi, F., Kessaris, N., Rossier, J., Vitalis, T., and Gallopin, T. (2012). “Characterization of type I and type II nNOS-expressing interneurons in the barrel cortex of mouse”. *Front. Neural Circuits* 6, 36.
- Pike, F. G., Meredith, R. M., Olding, A. W. A., and Paulsen, O. (2004). “Postsynaptic bursting is essential for ‘Hebbian’ induction of associative long-term potentiation at excitatory synapses in rat hippocampus”. *J. Physiol.* 518, 571–576.
- Pillow, J. W. (2009). “Time-rescaling methods for the estimation and assessment of non-Poisson neural encoding models”. *Adv. Neural Information Processing Systems 22*. Ed. by Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, and A. Culotta. MIT Press, 1473–1481.
- Pimashkin, A., Kastalskiy, I., Simonov, A., Koryagina, E., Mukhina, I., and Kazantsev, V. (2011). “Spiking signatures of spontaneous activity bursts in hippocampal cultures”. *Front. Comput. Neurosci.* 5, 1–12.
- Pine, J. (1980). “Recording action potentials from cultured neurons with extracellular microcircuit electrodes”. *J. Neurosci. Methods* 2, 19–31.
- Pluta, S., Naka, A., Veit, J., Telian, G., Yao, L., Hakim, R., Taylor, D., and Adesnik, H. (2015). “A direct translaminar inhibitory circuit tunes cortical output”. *Nat. Neurosci.* 18, 1631–1640.
- Pohlkamp, T., Dávid, C., Cauli, B., Gallopin, T., Bouché, E., Karagiannis, A., May, P., Herz, J., Frotscher, M., et al. (2014). “Characterization and distribution of reelin-positive interneuron subtypes in the rat barrel cortex”. *Cereb. Cortex* 24, 3046–3058.
- Potter, S. M. and DeMarse, T. B. (2001). “A new approach to neural cell culture for long-term studies”. *J. Neurosci. Methods* 110, 17–24.
- Potter, S. M., Lukina, N., Longmuir, K. J., and Wu, Y. (2001). “Multi-site two-photon imaging of neurons on multi-electrode arrays”. *SPIE Proc.* Vol. 4262, 104–110.
- Pouzat, C. and Chaffiol, A. (2009). “Automatic spike train analysis and report generation . An implementation with R , R2HTML and STAR”. *J. Neurosci. Methods* 181, 119–144.
- Purves, D., Augustine, G. J., Fitzpatrick, D., Katz, L. C., LaMantia, A.-S., McNamara, J. O., and Williams, S. M., eds. (2001). *Neuroscience*. Sunderland, MA: Sinauer Associates.
- Qin, L.-X. (2008). “An integrative analysis of microRNA and mRNA expression— A case study”. *Cancer Inform.* 6, 369–379.
- Quian Quiroga, R. (2009). “What is the real shape of extracellular spikes?” *J. Neurosci. Methods* 177, 194–198.
- Quian Quiroga, R. (2012). “Spike sorting”. *Curr. Biol.* 22, R45–R46.

- R Development Core Team (2011). *R: A Language and Environment for Statistical Computing*. Ed. by R Development Core Team.
- Radons, G., Becker, J. D., Dülfer, B., and Krüger, J. (1994). “Analysis, classification, and coding of multielectrode spike trains with hidden Markov models”. *Biol. Cybern.* 71, 359–373.
- Ramcharan, E. J., Gnadt, J. W., and Sherman, S. M. (2000a). “Burst and tonic firing in thalamic cells of unanesthetized, behaving monkeys”. *Vis. Neurosci.* 17, 55–62.
- Ramcharan, E. J., Cox, C. L., Zhan, X. J., Sherman, S. M., and Gnadt, J. W. (2000b). “Cellular mechanisms underlying activity patterns in the monkey thalamus during Visual behavior”. *J. Neurophysiol.* 84, 1982–1987.
- Rand, W. M. (1971). “Objective criteria for the evaluation of clustering methods”. *J. Am. Stat. Assoc.* 66, 846–850.
- Rege, M., Dong, M., and Hua, J. (2007). “Clustering web images with multi-modal features”. *Proc. 15th Int. Conf. Multimed.* 317–320.
- Reiner, A., Jiao, Y., Del Mar, N., Laverghetta, A. V., and Lei, W. L. (2003). “Differential morphology of pyramidal tract-type and intratelencephalically projecting-type corticostriatal neurons and their intrastriatal terminals in rats”. *J. Comp. Neurol.* 457, 420–440.
- Rey, H. G., Pedreira, C., and Quiñero, R. (2015). “Past, present and future of spike sorting techniques”. *Brain Res. Bull.* 119, 106–117.
- Rhoades, B. K. and Gross, G. W. (1994). “Potassium and calcium channel dependence of bursting in cultured neuronal networks”. *Brain Res.* 643, 310–318.
- Rizk, M. and Wolf, P. D. (2009). “Optimizing the automatic selection of spike detection thresholds using a multiple of the noise level”. *Med. Biol. Eng. Comput.* 47, 955–966.
- Robinette, B. L., Harrill, J. A., Mundy, W. R., and Shafer, T. J. (2011). “In vitro assessment of developmental neurotoxicity: use of microelectrode arrays to measure functional changes in neuronal network ontogeny”. *Front. Neuroeng.* 4, 1–9.
- Robinson, R. B. and Siegelbaum, S. A. (2003). “Hyperpolarization-activated cation currents: from molecules to physiological function”. *Annu. Rev. Physiol.* 65, 453–480.
- Rockhill, R. L., Euler, T., and Masland, R. H. (2000). “Spatial order within but not between types of retinal neurons”. *Proc. Natl. Acad. Sci. U.S.A.* 97, 2303–2307.
- Rolston, J. D., Wagenaar, D. A., and Potter, S. M. (2007). “Precisely timed spatiotemporal patterns of neural activity in dissociated cortical cultures”. *Neuroscience* 148, 294–303.
- Rudy, B., Fishell, G., Lee, S., and Hjerling-Leffler, J. (2011). “Three groups of interneurons account for nearly 100% of neocortical GABAergic neurons”. *Dev. Neurobiol.* 71, 45–61.
- Sanes, J. R. and Masland, R. H. (2015). “The types of retinal ganglion cells: Current status and implications for neuronal classification”. *Annu. Rev. Neurosci.* 38, 221–246.

- Santana, R., McGarry, L. M., Bielza, C., Larrañaga, P., and Yuste, R. (2013). “Classification of neocortical interneurons using affinity propagation”. *Front. Neural Circuits* 7, 185.
- Schmolesky, M. (2005). “The primary visual cortex”. *Webvision Organization of the Retina and Visual System*. Ed. by H. Kolb, E. Fernandez, and R. Nelson. Salt Lake City, UT: University of Utah Health Sciences Center.
- Schultz, W., Dayan, P., and Montague, P. R. (1997). “A neural substrate of prediction and reward”. *Science* 275, 1593–1599.
- Schultz, W. (1998). “Predictive reward signal of dopamine neurons”. *J. Neurophysiol.* 80, 1–27.
- Scorcioni, R., Polavaram, S., and Ascoli, G. A. (2008). “L-Measure: A web-accessible tool for the analysis, comparison and search of digital reconstructions of neuronal morphologies”. *Nat. Protoc.* 3, 866–876.
- Seelke, A. M. H. and Blumberg, M. S. (2010). “Developmental appearance and disappearance of cortical events and oscillations in infant rats”. *Brain Res.* 1324, 34–42.
- Selinger, J. V., Kulagina, N. V., O’Shaughnessy, T. J., Ma, W., and Pancrazio, J. J. (2007). “Methods for characterizing interspike intervals and identifying bursts in neuronal activity”. *J. Neurosci. Methods* 162, 64–71.
- Senn, V., Wolff, S. B. E., Herry, C., Grenier, F., Ehrlich, I., Gründemann, J., Fadok, J. P., Müller, C., Letzkus, J. J., and Lüthi, A. (2014). “Long-range connectivity defines behavioral specificity of amygdala neurons”. *Neuron* 81, 428–437.
- Seung, H. S. and Sümbül, U. (2014). “Neuronal cell types and connectivity: Lessons from the retina”. *Neuron* 83, 1262–1272.
- Shen, R., Olshen, A. B., and Ladanyi, M. (2009). “Integrative clustering of multiple genomic data types using a joint latent variable model with application to breast and lung cancer subtype analysis”. *Bioinformatics* 25, 2906–2912.
- Sherman, S. M. (2001). “Tonic and burst firing: Dual modes of thalamocortical relay”. *Trends Neurosci.* 24, 122–126.
- Sherman, S. M. (2017). “Functioning of circuits connecting thalamus and cortex”. *Compr. Physiol.* 7, 713–739.
- Shi, Y., Kirwan, P., and Livesey, F. J. (2012). “Directed differentiation of human pluripotent stem cells to cerebral cortex neurons and neural networks”. *Nat. Protoc.* 7, 1836–1846.
- Shipp, S. (2007). “Structure and function of the cerebral cortex”. *Curr. Biol.* 17, 443–449.
- Shuai, J., Bikson, M., Hahn, P. J., Lian, J., and Durand, D. M. (2003). “Ionic mechanisms underlying spontaneous CA1 neuronal firing in Ca^{2+} -free solution”. *Biophys. J.* 84, 2099–2111.
- Singh, A., Mewes, K., Gross, R. E., DeLong, M. R., Obeso, J. A., and Papa, S. M. (2016). “Human striatal recordings reveal abnormal discharge of projection neurons in Parkinson’s disease”. *Proc. Natl. Acad. Sci. U.S.A.* 113, 9629–9634.
- Smith, D. R. and Smith, G. K. (1965). “A statistical analysis of the continual activity of single cortical neurones in the cat unanaesthetized isolated forebrain”. *Biophys. J.* 5, 47–74.

- Snider, R. K., Kabara, J. F., Roig, B. R., and Bonds, A. B. (1998). “Burst firing and modulation of functional connectivity in cat striate cortex”. *J. Neurophysiol.* 80, 730–744.
- Spira, M. E. and Hai, A. (2013). “Multi-electrode array technologies for neuroscience and cardiology”. *Nat. Nanotechnol.* 8, 83–94.
- Spitzer, N. C. (2012). “Activity-dependent neurotransmitter respecification”. *Nat. Rev. Neurosci.* 13, 94–106.
- Spruston, N. (2008). “Pyramidal neurons: Dendritic structure and synaptic integration”. *Nat. Rev. Neurosci.* 9(3), 206–221.
- Squire, L., Berg, D., Bloom, F., Du Lac, S., Ghosh, A., and Spitzer, N., eds. (2008). *Fundamental Neuroscience*. Burlington, MA: Academic Press.
- Staiger, J. F., Flaggmeyer, I., and Schubert, D. (2004). “Functional diversity of layer IV spiny neurons in rat somatosensory cortex: Quantitative morphology of electrophysiologically characterized and biocytin labeled cells”. *Cereb. Cortex* 14, 690–701.
- Steriade, M., Timofeev, I., and Grenier, F. (2001). “Natural waking and sleep states: a view from inside neocortical neurons”. *J. Neurophysiol.* 85, 1969–1985.
- Stevens, C. F. (1998). “Neuronal diversity: Too many cell types for comfort?” *Curr. Biol.* 8, R708–R710.
- Stevens, C. F. and Wang, Y. (1995). “Facilitation and depression at single central synapses”. *Neuron* 14, 795–802.
- Strehl, A. and Ghosh, J. (2002). “Cluster ensembles – A knowledge reuse framework for combining multiple partitions”. *J. Mach. Learn. Res.* 3, 583–617.
- Sugino, K., Hempel, C. M., Miller, M. N., Hattox, A. M., Shapiro, P., Wu, C., Huang, Z. J., and Nelson, S. B. (2006). “Molecular taxonomy of major neuronal classes in the adult mouse forebrain”. *Nat. Neurosci.* 9, 99–107.
- Sümbül, U., Song, S., McCulloch, K., Becker, M., Lin, B., Sanes, J. R., Masland, R. H., and Seung, H. S. (2014). “A genetic and computational approach to structurally classify neuronal types.” *Nat. Commun.* 5, 3512.
- Sun, J. J., Kilb, W., and Luhmann, H. J. (2010). “Self-organization of repetitive spike patterns in developing neuronal networks in vitro”. *Eur. J. Neurosci.* 32, 1289–1299.
- Sun, S. (2013). “A survey of multi-view machine learning”. *Neural Comput. Appl.* 23, 2031–2038.
- Sun, W., Tan, Z., Mensh, B. D., and Ji, N. (2016). “Thalamus provides layer 4 of primary visual cortex with orientation- and direction-tuned inputs”. *Nat. Neurosci.* 19, 308–315.
- Suzuki, I. K. and Vanderhaeghen, P. (2015). “Is this a brain which I see before me? Modeling human neural development with pluripotent stem cells”. *Co. Biol.* 142, 3138–3150.
- Suzuki, N. and Bekkers, J. M. (2010). “Distinctive classes of GABAergic interneurons provide layer-specific phasic inhibition in the anterior piriform Cortex”. *Cereb. Cortex* 20, 2971–2984.
- Swadlow, H. A. and Gusev, A. G. (2001). “The impact of ‘bursting’ thalamic impulses at a neocortical synapse”. *Nat. Neurosci.* 4, 402–408.

- Syed, M. M., Lee, S., Zheng, J., and Zhou, Z. J. (2004). “Stage-dependent dynamics and modulation of spontaneous waves in the developing rabbit retina”. *J Physiol* 560, 533–549.
- Tam, D. C. (2002). “An alternate burst analysis for detecting intra-burst firings based on inter-burst periods”. *Neurocomputing* 46, 1155–1159.
- Tang, W., Lu, Z., and Dhillon, I. S. (2009). “Clustering with multiple graphs”. *Proc. IEEE Int. Conf. Data Min.* 1016–1021.
- Tank, D. W., Regehr, W. G., and Delaney, K. R. (1995). “A quantitative analysis of presynaptic calcium dynamics that contribute to short-term enhancement”. *J. Neurosci.* 15, 7940–7952.
- Tasic, B., Menon, V., Nguyen, T. N., Kim, T. K., Jarsky, T., Yao, Z., Levi, B., Gray, L. T., Sorensen, S. A., et al. (2016). “Adult mouse cortical cell taxonomy revealed by single cell transcriptomics”. *Nat. Neurosci.* 19, 335–346.
- Thomas Jr., C. A., Springer, P. A., Loeb, G. E., Berwald-Netter, Y., and Okun, L. M. (1972). “A miniature microelectrode array to monitor the bioelectric activity of cultured cells”. *Exp. Cell Res.* 74, 61–66.
- Thomas, M. J., Watabe, A. M., Moody, T. D., Makhinson, M., and O’Dell, T. J. (1998). “Postsynaptic complex spike bursting enables the induction of LTP by theta frequency synaptic stimulation”. *J. Neurosci.* 18, 7118–7126.
- Thomson, A. M. (1997). “Activity-dependent properties of synaptic transmission at two classes of connections made by rat neocortical pyramidal axons in vitro”. *J. Physiol.* 502, 131–147.
- Thomson, A. M. (2000a). “Facilitation, augmentation and potentiation at central synapses”. *Trends Neurosci.* 23, 305–312.
- Thomson, A. M. (2000b). “Molecular frequency filters at central synapses”. *Prog. Neurobiol.* 62, 159–196.
- Thomson, A. M. (2010). “Neocortical layer 6, a review”. *Front. Neuroanat.* 4, 1–14.
- Tibshirani, R., Walther, G., and Hastie, T. (2001). “Estimating the number of clusters in a data set via the gap statistic”. *J. R. Stat. Soc. B* 63, 411–423.
- Tibshirani, R. (2007). “Regression shrinkage and selection via the lasso”. *J. R. Stat. Soc. B* 58, 267–288.
- Tobler, P. N., Dickinson, A., and Schultz, W. (2003). “Coding of predicted reward omission by dopamine neurons in a conditioned inhibition paradigm”. *J. Neurosci.* 23, 10402–10410.
- Tokdar, S., Xi, P., Kelly, R. C., and Kass, R. E. (2010). “Detection of bursts in extracellular spike trains using hidden semi-Markov point process models”. *J. Comput. Neurosci.* 29, 203–212.
- Toledo-Rodriguez, M., Gupta, A., Wang, Y., Wu, C. Z., and Markram, H. (2003). “Neocortex: Basic Neuron Types”. *The Handbook of Brain Theory and Neural Networks*. Ed. by M. A. Arbib. 2nd edition. Cambridge, MA: MIT Press, 719–725.
- Tolonen, M., Palva, J. M., Andersson, S., and Vanhatalo, S. (2007). “Development of the spontaneous activity transients and ongoing cortical activity in human preterm babies”. *Neuroscience* 145, 997–1006.
- Tritsch, N. X., Yi, E., Gale, J. E., Glowatzki, E., and Bergles, D. E. (2007). “The origin of spontaneous activity in the developing auditory system”. *Nature* 450, 50–55.

- Tuckwell, H. C. (1988). *Introduction to Theoretical Neurobiology: Volume 2, Nonlinear and Stochastic Theories*. Cambridge, UK: Cambridge University Press.
- Turnbull, L., Dian, E., and Gross, G. (2005). “The string method of burst identification in neuronal spike trains”. *J. Neurosci. Methods* 145, 23–35.
- Tzortzis, G. and Likas, A. (2012). “Kernel-based weighted multi-view clustering”. *Proc. IEEE Int. Conf. Data Min.* 675–684.
- Ullo, S., Nieuws, T. R., Sona, D., Maccione, A., Berdondini, L., and Murino, V. (2014). “Functional connectivity estimation over large networks at cellular resolution based on electrophysiological recordings and structural prior”. *Front. Neuroanat.* 8, 1–15.
- Urban-Ciecko, J. and Barth, A. L. (2016). “Somatostatin-expressing neurons in cortical networks”. *Nat. Rev. Neurosci.* 17, 401–409.
- Valdivia, P., Martin, M., LeFew, W. R., Ross, J., Houck, K. A., and Shafer, T. J. (2014). “Multi-well microelectrode array recordings detect neuroactivity of ToxCast compounds”. *Neurotoxicology* 44, 204–217.
- Välkki, I. A., Lenk, K., Mikkonen, J. E., and Kapucu, F. E. (2017). “Network-wide adaptive burst detection depicts neuronal activity with improved accuracy”. *Front. Comput. Neurosci.* 11, 1–14.
- Van Aerde, K. I. and Feldmeyer, D. (2015). “Morphological and physiological characterization of pyramidal neuron subtypes in rat medial prefrontal cortex”. *Cereb. Cortex* 25, 788–805.
- Van Den Pol, A. N., Obrietan, K., and Belousov, A. (1996). “Glutamate hyperexcitability and seizure-like activity throughout the brain and spinal cord upon relief from chronic glutamate receptor blockage in culture”. *Neuroscience* 74, 653–674.
- Van Huizen, F., Romijn, H. J., and Habets, A. M. M. C. (1985). “Synaptogenesis in rat cerebral cortex cultures is affected during chronic blockade of spontaneous bioelectric activity by tetrodotoxin”. *Dev. Brain Res.* 19, 67–80.
- Van Pelt, J., Corner, M. A., Wolters, P. S., Rutten, W. L. C., and Ramakers, G. J. A. (2004a). “Longterm stability and developmental changes in spontaneous network burst firing patterns in dissociated rat cerebral cortex cell cultures on multielectrode arrays”. *Neurosci. Lett.* 361, 86–89.
- Van Pelt, J., Wolters, P. S., Corner, M. A., Rutten, W. L. C., and Ramakers, G. J. A. (2004b). “Long-term characterization of firing dynamics of spontaneous bursts in cultured neural networks”. *IEEE Trans. Biomed. Eng.* 51, 2051–2062.
- Van Pelt, J., Vajda, I., Wolters, P. S., Corner, M. A., and Ramakers, G. J. A. (2005). “Dynamics and plasticity in developing neuronal networks in vitro”. *Prog. Brain Res.* 147, 173–188.
- Vanderwolf, C. H. (1969). “Hippocampal electrical activity and voluntary movement in the rat”. *Electroencephalogr. Clin. Neurophysiol.* 26, 407–418.
- Vanhatalo, S. and Kaila, K. (2006). “Development of neonatal EEG activity: From phenomenology to physiology”. *Semin. Fetal Neonatal Med.* 11, 471–478.
- Venables, W. N. and Ripley, B. D. (2002). *Modern Applied Statistics With S*. 4th Edition. Springer, 126.

- Vinh, N. X., Epps, J., and Bailey, J. (2010). “Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance”. *J. Mach. Learn. Res.* 11, 2837–2854.
- Wagenaar, D. A., Pine, J., and Potter, S. M. (2004). “Effective parameters for stimulation of dissociated cultures using multi-electrode arrays”. *J. Neurosci. Methods* 138, 27–37.
- Wagenaar, D. A., Pine, J., and Potter, S. M. (2006a). “An extremely rich repertoire of bursting patterns during the development of cortical cultures.” *BMC Neurosci.* 7, 11.
- Wagenaar, D. A., Nadasdy, Z., and Potter, S. M. (2006b). “Persistent dynamic attractors in activity patterns of cultured neuronal networks”. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* 73, 051907.
- Wagenaar, D., Demarse, T. B., and Potter, S. M. (2005). “MeaBench: A toolset for multi-electrode data acquisition and on-line analysis”. *Proc. 2nd Int. IEEE EMBS Conf. Neural Eng.* 518–521.
- Wainger, B. J., Kiskinis, E., Mellin, C., Wiskow, O., Han, S. S. W., Sandoe, J., Perez, N. P., Williams, L. A., Lee, S., et al. (2014). “Intrinsic membrane hyperexcitability of ALS patient-derived motor neurons”. *Cell Rep.* 7, 1–11.
- Walker, A. G., Miller, B. R., Fritsch, J. N., Barton, S. J., and Rebec, G. V. (2008). “Altered information processing in the prefrontal cortex of Huntington’s disease mouse models”. *J. Neurosci.* 28, 8973–8982.
- Wang, X., Qian, B., Ye, J., and Davidson, I. (2013). “Multi-objective multi-view spectral clustering via Pareto optimization”. *Proc. 13th SIAM Int. Conf. Data Min.* 234–242.
- Wang, X.-J. (1999). “Fast burst firing and short-term synaptic plasticity: A model of neocortical chattering neurons”. *Neuroscience* 89, 347–362.
- Wang, X.-J. (2010). “Neurophysiological and computational principles of cortical rhythms in cognition.” *Physiol. Rev.* 90, 1195–268.
- Wang, Y., Gupta, A., Toledo-Rodriguez, M., Wu, C. Z., and Markram, H. (2002). “Anatomical, physiological, molecular and circuit properties of nest basket cells in the developing somatosensory cortex”. *Cereb. Cortex* 12, 395–410.
- Wassle, H., Puller, C., Muller, F., and Haverkamp, S. (2009). “Cone contacts, mosaics, and territories of bipolar cells in the mouse retina”. *J. Neurosci.* 29, 106–117.
- Weissman, T. A., Riquelme, P. A., Ivic, L., Flint, A. C., and Kriegstein, A. R. (2004). “Calcium waves propagate through radial glial cells and modulate proliferation in the developing neocortex”. *Neuron* 43, 647–661.
- Weliky, M. and Katz, L. C. (1999). “Correlational structure of spontaneous neuronal activity in the developing lateral geniculate nucleus in vivo”. *Science* 285, 599–604.
- Weyand, T. G., Boudreaux, M., and Guido, W. (2001). “Burst and tonic response modes in thalamic neurons during sleep and wakefulness”. *J. Neurophysiol.* 85, 1107–1118.
- Whitson, J., Kubato, D., Shimono, K., Jia, Y., and Taketani, M. (2006). “Multi-electrode arrays: Enhancing traditional methods and enabling network

- physiology". *Advances in Network Electrophysiology using Multi-Electrode Arrays*. Ed. by M. Taketani and M. Baudry. New York: Springer, 38–68.
- Williams, S. R. and Stuart, G. J. (2000). "Site independence of EPSP time course is mediated by dendritic I_h in neocortical pyramidal neurons". *J. Neurophysiol.* 3177–3182.
- Wong, R. O., Meister, M., and Shatz, C. J. (1993). "Transient period of correlated bursting activity during development of the mammalian retina". *Neuron* 11, 923–938.
- Woodard, C. M., Campos, B. A., Kuo, S. H., Nirenberg, M. J., Nestor, M. W., Zimmer, M., Mosharov, E. V., Sulzer, D., Zhou, H., et al. (2014). "iPSC-derived dopamine neurons reveal differences between monozygotic twins discordant for parkinson's disease". *Cell Rep.* 9, 1173–1182.
- Wu, L. G. and Saggau, P. (1994). "Presynaptic calcium is increased during normal synaptic transmission and paired-pulse facilitation, but not in long-term potentiation in area CA1 of hippocampus". *J. Neurosci.* 14, 645–654.
- Xia, R., Pan, Y., Du, L., and Yin, J. (2014). "Robust multi-view spectral clustering via low-rank and sparse decomposition". *Proc. Conf. AAAI Artif. Intell.* 2149–2155.
- Xu, H., Jeong, H.-Y., Tremblay, R., and Rudy, B. (2013). "Neocortical somatostatin-expressing GABAergic interneurons disinhibit the thalamorecipient Layer 4". *Neuron* 77, 155–167.
- Xu, H. P., Furman, M., Mineur, Y. S., Chen, H., King, S. L., Zenisek, D., Zhou, Z. J., Butts, D. A., Tian, N., et al. (2011). "An instructive role for patterned spontaneous retinal activity in mouse visual map development". *Neuron* 70, 1115–1127.
- Xu, W., Morishita, W., Buckmaster, P. S., Pang, Z. P., Malenka, R. C., and Südhof, T. C. (2012). "Distinct neuronal coding schemes in memory revealed by selective erasure of fast synchronous synaptic transmission". *Neuron* 73, 990–1001.
- Xu, X., Roby, K. D., and Callaway, E. M. (2010). "Immunochemical characterization of inhibitory mouse cortical neurons: Three chemically distinct classes of inhibitory cells". *J. Comp. Neurol.* 518, 389–404.
- Yavorska, I. and Wehr, M. (2016). "Somatostatin-expressing inhibitory interneurons in cortical circuits". *Front. Neural Circuits* 10, 1–18.
- Ylä-Outinen, L., Heikkilä, J., Skottman, H., Suuronen, R., Aänismaa, R., and Narkilahti, S. (2010). "Human cell-based micro electrode array platform for studying neurotoxicity". *Front. Neuroeng.* 3, 1–9.
- Yue, C. and Yaari, Y. (2004). "KCNQ/M channels control spike afterdepolarization and burst generation in hippocampal neurons". *J. Neurosci.* 24, 4614–4624.
- Yue, C., Remy, S., Su, H., Beck, H., and Yaari, Y. (2005). "Proximal persistent Na^+ channels drive spike afterdepolarizations and associated bursting in adult CA1 pyramidal cells". *J. Neurosci.* 25, 9704–9720.
- Zaitsev, A. V., Povysheva, N. V., Gonzalez-Burgos, G., and Lewis, D. A. (2012). "Electrophysiological classes of layers 2-3 pyramidal cells in monkey prefrontal cortex." *J. Neurophysiol.* 108, 595–609.

- Zeisel, A., Machado, A. B. M., Codeluppi, S., Lonnerberg, P., La Manno, G., Jureus, A., Marques, S., Munguba, H., He, L., et al. (2015). “Cell types in the mouse cortex and hippocampus revealed by single-cell RNA-seq”. *Science* 347, 1138–1142.
- Zhang, X. and Shasha, D. (2006). “Better burst detection”. *Proc. 22nd Int. Conf. Data Eng.* 146.
- Zhao, X., Evans, N., and Dugelay, J. L. (2014). “A subspace co-training framework for multi-view clustering”. *Pattern Recognit. Lett.* 41, 73–82.
- Zheng, J., Lee, S., and Zhou, Z. J. (2006). “A transient network of intrinsically bursting starburst cells underlies the generation of retinal waves”. *Nat. Neurosci.* 9, 363–371.
- Zhou, D. and Burges, C. J. C. (2007). “Spectral clustering and transductive learning with multiple views”. *Proc. 24th Int. Conf. Mach. Learn.* 1159–1166.
- Zhou, D., Huang, J., and Schölkopf, B. (2005). “Learning from labeled and unlabeled data on a directed graph”. *Proc. 22nd Int. Conf. Mach. Learn.* 1036–1043.
- Zhou, Z., Sorensen, S. A., and Peng, H. (2015). “Neuron crawler: An automatic tracing algorithm for very large neuron images”. *Proc. Int. Symp. Biomed. Imaging*, 870–874.
- Zhu, D., Zhang, T., Jiang, X., Hu, X., Chen, H., Yang, N., Lv, J., Han, J., Guo, L., and Liu, T. (2014). “Fusing DTI and fMRI data: A survey of methods and applications”. *Neuroimage* 102, 184–191.
- Zhu, Y. and Shasha, D. (2003). “Efficient elastic burst detection in data streams”. *Proc. Ninth ACM SIGKDD Int. Conf. Knowl. Discov. data Min.* 336–345.
- de Lima, A. D., Gieseler, A., and Voigt, T. (2009). “Relationship between GABAergic interneurons migration and early neocortical network activity”. *Dev. Neurobiol.* 69, 105–123.
- de Sa, V. R. (2005). “Spectral clustering with two views”. *Proc. Work. Learn. with Mult. Views, 22nd ICML*, 20–27.
- de Sa, V. R., Gallagher, P. W., Lewis, J. M., and Malave, V. L. (2010). “Multi-view kernel construction”. *Mach. Learn.* 79, 47–71.