Exploration of time-warped cortical spike time patterns



James Bryden Isbister
The Queen's College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Michaelmas 2021

For my parents and grandparents,
to whom I owe the world.

Acknowledgements

In the first months of a PhD, a student might pull down a copy of a past thesis from a shelf in their laboratory and open it to find several pages of acknowledgments. The list of gratitudes might seem largely customary, with a few genuine thanks to the friends and family who helped them through the hardest moments and without whom the thesis would not have been possible. The 'personal journey' that the thesis author has been on is impossible to imagine for the new initiate. By the time the PhD student is ready to submit their own thesis, they might think back to this earlier thesis, and are likely to have a much better appreciation for the earlier author's journey. It is probable that they will have learned something about life, friendships, scientific progress and themselves. They'll definitely realise that without the support of mentors, colleagues, partners, friends and family, the outcome would have been very different.

I would first like to thank my parents Barbara and Bryden Isbister who have selflessly given their love, guidance, inspiration and support consistently through my life. The experience of doing a PhD has been a dream come true. This thesis is dedicated to you and your parents, without whom this would not have been possible. It has been a true pleasure to share the journey with you both. Thank you for everything: for taking me to the Natural History Museum and Science Museum as a child, for teaching me to read, write and do maths, our explorations of the world, your enthusiasm for the Arts, scientific advances and wondering how the world works, taking me to school and university, being there for me at every stage of life, sharing interesting and inspiring articles, showing enthusiasm for whatever I chose to pursue, the laughter we have shared, and a million other things, thank you.

I would like to thank my supervisor Dr. Simon M. Stringer for the opportunity to study the brain at one of the most inspiring universities in the world. Your passion to uncover fundamental scientific principles was a great source of inspiration and a privilege to learn from. Your instinctive drive to challenge accepted scientific ideas gave me the confidence to pursue the research questions of the thesis. It was also a pleasure to share a friendship with you through the journey. The environment you created in Oxford was very special and I will cherish my memories from the lab forever.

For taking the time to read the chapters of the thesis and/or the Scientific Reports paper published in July 2021 [136], I would like to thank Dr. Simon Stringer, Dr. Michael Reimann, Dr. Heiko Luhmann, Dr. Vicente Reyes-Puerta, Dr. Jyh-Jang

Sun, Dr. Illia Horenko and the reviewers of the Scientific Reports paper. The insights and guidance they provided were invaluable to the final thesis. Special thanks goes to Dr. Vicente Reyes-Puerta who spent his free time reading and constructively critiquing the Scientific Reports paper, as well as helping with access to the data.

I would like to thank friends from the Oxford Centre for Theoretical Neuroscience and Artificial Intelligence for their support, friendship, patience, understanding and inspiration. Particularly, the thesis author thanks Dr. Harry Jordan, Jonas Hasbach, Dr. Juan Manuel Galeazzi González, Dr. Nasir Ahmad, Dr. Daniel Walters, Niels Leadholm, Dr. Daniel Navarro, Hannah Dury, Toby St Clare Smithe, Gisbert Teepe, Odhran Odonoghue, Elliot Smith, Illenia Salaris, Karen Verstraelen, Jannis Born, Clements Hutter and Dr. Akihiro Eguchi.

I also thank Dr. Michael Reimann, Professor Henry Markram, Dr. Giuseppe Chindemi and the Blue Brain Project for the opportunity to spend a summer internship and the final year of the DPhil at the EPFL Blue Brain Project. The whole Blue Brain Projet team is incredibly supportive and it has been a pleasure to learn through engagement in their ongoing research projects. Special thanks goes to Dr. Michael Reimann, Dr. Srikanth Ramaswamy, Dr. Giuseppe Chindemi, Dr. Christoph Pokorny, Sirio Bolaños Puchet and András Ecker for their teamwork and enjoyable scientific debates.

To the many friends that I have had the privilege of spending time with in Oxford, Switzerland and around the world, for their amazing warmth, care, laughter, discussions and love, including those from The Queen's College, Department of Experimental Psychology, Oxford, and Lausanne. Particular thanks goes to: Patrick Wheatley, Dr. Jakub Vohryzek, Dr. KC Park, Dr. Marten Swaart, Caroline Sarfaty, Dr. Sam Wheeler, Anna Jungbluth, Dr. Vangelis Oikonomou, Dr. Sofia Athanasopoulou, Dr. Nadja Trudel, Dr. Diana Maia Alcaide, Eri Ichijo, Ana Recio de Mur, Gayatri Dangi, Joe Childs, John Vagg, Mairead Healy, Jonas Boze, Rok Medves, Andre Hector, Dr. Romina Mariano, Stephanie Reh, Georgios Iatropoulos, Dr. Michael Doron, Dr. Daril Brown, Dr. Guillaume Bellec, Adrian Clayton, Jakub Mihule, Carlos Mangado, Nils Osseiran, Laurence MG and Matthew Valentine. I couldn't have done this without you.

Exploration of time-warped cortical spike time patterns

James B. Isbister

Copyright © March 2022

Contents

Abstract **xi**

List of Figures **xii**

List of Videos **xlv**

1 Introduction **1**

1.1 Motivating research questions 1

1.1.1 The neural coding problem 3

1.1.1.1 Stimulus-response distributions, non-stationarity and spike train metrics 4

1.1.1.2 Discrimination and neuronal sensitivity 6

1.1.1.3 Tuning curves, decoding algorithms and information analysis 7

1.1.1.4 Neuronal stimulus-response adaptation and other non-stationary processes 8

1.1.1.5 Noise correlations and cortical state 9

1.1.1.6 Decoding 11

1.1.1.7 Summary 12

1.1.2 The feature-binding problem 13

1.1.2.1 Neuronal and modular specialisation 13

1.1.2.2 Correct association of feature-presence and feature-properties 14

1.1.2.3 Joint feature-presence and feature-property encoding 16

1.1.2.4 Overlapping and simultaneously active representations 17

1.1.2.5 Parsing and robustness to false combinations 18

1.1.2.6 Learning and semantic meaning 19

1.1.2.7 Property encoding by (state-dependent?) continua and neural manifolds 19

1.1.2.8 The representation of novel stimuli along continua . . 20
1.1.2.9 Compositionality of continua: property-dependent parsing and holographic projection of information 21
1.1.2.10 Generalisable action 22
1.1.2.11 One-shot or multi-shot learning 23
1.2 Summary of motivating research questions 23
1.3 Multi-neuron first and single spike patterns 24
1.4 Contribution of the thesis 26
1.4.1 Chapter 2: Literature Review 26
1.4.2 Chapter 3: Clustering and control for adaptation uncovers time-warped spike time patterns in cortical networks in vivo (part 1) 27
1.4.3 Chapter 4: Clustering and control for adaptation uncovers time-warped spike time patterns in cortical networks in vivo (part 2) 28
1.4.4 Chapter 5: Comparison of in silico and in vivo responses for a detailed biophysical cortical column model 29
1.4.5 Chapter 6: Hypotheses of time-warped multi-neuron single spike patterns 30
1.4.6 Chapter 7: Co-creation of a spiking neural network simulator and the emergence of relative spike timing continua in a phenomenological spiking neural network 30

2 Literature Review **32**
2.1 Introduction 32
2.1.1 Correlating neural activity with discrimination 32
2.1.2 Fast discrimination and fast cortical processing 33
2.2 Candidate neural code: spike count, mean firing rate and population spike count codes 34
2.2.1 Techniques and considerations 35
2.2.1.1 Low trial counts per unique stimulus 35
2.2.1.2 MUA activity confers higher spike counts 36
2.2.1.3 Non-atomic stimulus presentations and large spike count windows 36
2.2.1.4 Backward masking 37

2.2.1.5 Optogenetic control: multi-neuron single spike patterns are sufficient for visual discrimination 37
2.2.1.6 Multi-neuron single spike pattern represent atomic stimuli in the whisker domain 38
2.2.1.7 Sliding time windows 38
2.2.1.8 Spike count noise correlations and stimulus-response adaptation 39
2.2.1.9 Dimensionality reduction and population trajectories 39
2.2.1.10 Latent state dependent decoding 40
2.2.1.11 Insight from calcium imaging 41
2.2.2 Consideration of spike count encoding w.r.t. the GRQs (visual system) 41
2.2.2.1 Fast spike count / mean firing rate encoding 42
2.2.2.2 Property and continua encoding in the visual system 43
2.2.2.3 Superposition and competition 43
2.3 Candidate neural code: spike timing 44
2.3.1 Temporally precise multi-neuron first and single spike patterns 44
2.3.2 Precise relative spike times, noise correlations and neuronal sensitivity 45
2.3.3 Relative spike time difference continua encoding 46
2.3.4 Possible connectivity-types underlying relative spike time differences 46
2.3.5 Superposition of spike time representations 47
2.3.6 Packet encoding 47
2.3.7 Synfire chainsm synfire braids and polychronisation 47
2.4 Previous theories of feature-binding 48
2.4.1 Binding-by-synchrony 48
2.4.2 Feature-integration theory 48
2.4.3 Population codes 50
2.4.4 Communication through coherence 50
2.4.5 Cortex-wide grid cells 51

3 Clustering and control for adaptation uncovers time-warped spike time patterns in cortical networks in vivo (part 1) 52
3.1 Introduction 52
3.1.1 Chapter overview 55

3.2 The whisker system and barrel cortex . . . 57
3.2.1 The whiskers and whisker follicles . . . 59
3.2.2 The brainstem trigeminal nuclei (TN) . . . 59
3.2.3 The ventro-posterior-medial nucleus . . . 59
3.2.4 The barrel cortex . . . 59
3.2.5 Neuron types . . . 61
3.2.6 Feedforward inhibition . . . 61
3.2.7 Whisker perception and its neural correlates . . . 62
3.2.8 Stimulus representation in the barrel cortex . . . 62
3.2.9 Stimulus-response adaptation and non-stationarity in the barrel cortex . . . 63
3.3 Methods I: Neurophysiology . . . 65
3.4 Results I: Single spike patterns, adaptation and observation of clustered noise correlations . . . 67
3.4.1 Confirmation of single spike patterns following stimulus onset 67
3.4.2 Confirmation of neuronal stimulus-response adaptation for single neurons . . . 69
3.4.3 Observation of clustered spike time noise correlations . . . 70
3.5 Methods II: Two-stage algorithm . . . 73
3.5.1 Stage 1 clusters . . . 73
3.5.2 Correlation surrogate control datasets . . . 74
3.5.3 Stage 2 clusters . . . 74
3.5.4 Comparison to alternative clustering methods . . . 76
3.6 Results II: Clustered spike time noise correlations . . . 78
3.6.1 Clustered spike time noise correlations . . . 78
3.6.1.1 Stage 1 clusters . . . 78
3.6.1.2 Stage 1 correlation controls . . . 79
3.6.1.3 Stage 2 ellipses and correlated unclustered response distributions . . . 79
3.6.1.4 Comparison to alternative clustering algorithms . . . 82
3.7 Methods III: Cluster correlation angle and angle surrogate control datasets . . . 83
3.7.1 Cluster angle . . . 83
3.7.2 Angle surrogate control datasets . . . 83
3.8 Results III: Cluster correlation angles are non-unitary . . . 86
3.8.1 Comparison to Custom GMM . . . 86

3.9 Methods IV: Stationary clusters and predicted encoding capacity . . . 88
3.9.1 Stationary cluster and factor analysis criteria 89
3.10 Results IV: Stationary clusters and state-dependent response distributions 91
3.10.1 Stationary clusters 91
3.10.2 Fixed-state and state-conditioned relative difference response distributions 93
3.10.3 Estimation of state-conditioned response distributions 95
3.11 Methods V: Analysis of non-stationary clusters 98
3.11.1 Differencing 98
3.11.2 ARIMA modelling 101
3.12 Results V: Non-stationary clusters are also modulated by shared trial-to-trial changes in the excitability of neuronal pairs 103
3.13 Results VI: Frequency dependence and spatial distribution of correlated clusters 106
3.14 Discussion 108

4 Clustering and control for adaptation uncovers time-warped spike time patterns in cortical networks in vivo (part 2) 123
4.1 Introduction 123
4.1.1 The problem of low spiking reliability and clustered correlations 124
4.2 Methods I: Pairwise correlation of predicted cortical state 125
4.2.1 Composing the pool 125
4.2.2 Correlations of predicted pairwise states 125
4.2.3 Correlations of predicted states with spike times and spike time differences 126
4.3 Results I: Pairwise correlation of predicted cortical state 126
4.3.1 Composition of the pool 126
4.3.2 Correlations of predicted pairwise states 126
4.3.3 Correlations of predicted states with spike times and spike time differences 126
4.4 Methods II: Spike time co-modulation of larger groups of neurons - the correlation matrix 128
4.4.1 Analysis of first spike time correlation matrices 129
4.4.1.1 Composing the correlation matrix 129

4.5 Results II: Spike time co-modulation of larger groups of neurons - the correlation matrix . . . 131
4.5.1 Example: 1Hz inter trial interval . . . 131
4.5.2 All stimulus conditions . . . 132
4.6 Methods III: Averaged cortical state . . . 132
4.6.1 Improved prediction of cortical state . . . 132
4.6.1.1 The problem with averaging predicted PC states . . . 133
4.6.1.2 Adjustment of cluster offsets and angles . . . 133
4.6.1.3 Predicted states by trial (before and after adjustment) 135
4.6.2 Stationarity of averaged cortical state . . . 136
4.7 Methods IV: Prediction and correlation of spike times and differences from averaged cortical state . . . 137
4.8 Results IV: Prediction and correlation of spike times and differences from averaged cortical state . . . 139
4.9 Methods + Results V: Visualisation of time-warped multi-neuron single spike patterns . . . 139
4.9.1 Without smoothing . . . 139
4.9.2 With smoothing . . . 141
4.10 Discussion . . . 143

5 Comparison of in silico and in vivo responses for a detailed biophysical cortical column model 144
5.1 Introduction . . . 144
5.1.1 Simulations as a validation of biophysical models . . . 145
5.2 The Blue Brain Project (BBP) somatosensory column model [175] . . . 146
5.2.1 Reconstruction . . . 146
5.2.2 Simulations . . . 151
5.2.2.1 Initial thalamic stimulation tests . . . 153
5.3 Other biophysical cortical models . . . 159
5.3.1 Allen Institute mouse primary visual cortex model [23] . . . 159
5.3.2 Responses of morphologically detailed neurons in the barrel cortex [69, 68] . . . 160
5.3.3 Multi-scale EM reconstruction and in vivo functional comparison [277] . . . 160
5.3.4 Null model of the mouse neocortex micro-connectome [222] . . . 161
5.4 Accounting for in vivo sampling bias . . . 161

5.4.1 The sampling problem: recording methods and their biases . . 162
5.4.2 Population distributions of firing rates 163
5.4.3 Long-tailed distributions 163
5.4.4 Layer-specific differences in spontaneous activity 164
5.4.5 Similarity of spontaneous and evoked firing rate distributions . 166
5.4.6 The single neuron's perspective 167
5.5 Methods: comparison to in vivo 167
5.5.1 Model and comparison neuron groups 168
5.5.2 Depolarisation, $[Ca^{2+}]_o$ and input parameter grid search . . . 169
5.5.2.1 Cortical onset 170
5.5.2.2 In vivo spiking activity 171
5.5.2.3 Visual comparison 171
5.5.2.4 Metrics accounting for in vivo sampling bias 172
5.5.2.5 Comparison time windows 172
5.5.2.6 Quantitative comparison framework 173
5.6 Results: comparison to in vivo 173
5.6.1 Visual comparison of stimulus-evoked activity 173
5.6.2 Ratios of stimulus-evoked firing rates between neuronal groups 175
5.6.3 Similarity of PSTH Profiles 179
5.6.3.1 Comparable spike counts and spike time latencies follow from PSTH profile similarity 181
5.6.4 Evoked firing rate similarity to in vivo 183
5.6.4.1 Evoked spontaneous firing rate analysis 186
5.6.4.2 Spontaneous firing rate similarity to in vivo 189
5.6.5 Ratios of spontaneous firing rates between neuronal groups . . 189
5.6.5.1 Ratio between spontaneous activity and stimulus-evoked firing rates 192
5.6.5.2 Selection of model for time-warp test 192
5.7 Methods: in silico test for time-warped spike time noise correlations . 194
5.7.1 Simulation 194
5.7.2 Analysis 194
5.8 Results: in silico test for time-warped spike time noise correlations . . 197
5.9 Discussion 197
5.9.1 Summary 197
5.9.2 Suggested future measures 199
5.9.3 Current use and new analyses 200

5.9.4 In silico test for time-warped spike time noise correlations . . 202

6 Hypotheses of time-warped multi-neuron single spike patterns 205
6.1 General research questions 205
6.2 Hypotheses 206

7 Co-creation of a spiking neural network simulator and the emergence of relative spike timing continua in a phenomenological spiking neural network. 216
7.1 Background: Spiking Neural Network Simulators 218
7.1.1 Modelling of neurons, synapses and plasticity 218
7.1.2 Neuron models 218
7.1.3 Synapse models 218
7.1.3.1 'Voltage-based' synapses 218
7.1.3.2 'Conductance-based' synapses 219
7.1.4 STDP 219
7.1.5 Numerical methods for solving SNN simulations 220
7.1.6 Other simulators 221
7.1.7 GPU programming 221
7.2 Methods: Spike 222
7.2.1 Optimisations 222
7.2.1.1 Synaptic Matrices versus Lists 222
7.2.1.2 Timestep Grouping 223
7.2.1.3 Neuronal-Synaptic Parallelism 223
7.2.1.4 Delay Insensitive 223
7.2.2 Comparison to other simulators 224
7.2.2.1 Comparison hardware 225
7.2.2.2 Benchmark models 225
7.3 Results: Spike 226
7.3.1 Vogels-Abbott Benchmark 226
7.3.1.1 Comparing Simulators (Single Threaded) 226
7.3.1.2 Comparing Simulators (Multi Threaded) 227
7.3.1.3 The Scaling Behaviour of the Spike Simulator 228
7.3.2 Brunel Plasticity Benchmark 228
7.3.2.1 Multithreaded CPU Performance 229
7.3.3 Summary 229
7.4 A model for exploring the emergence of spike timing continua 230

7.4.1 Biological Evidence of STDP . . . 231
7.4.2 Models of STDP . . . 231
7.4.3 STDP Associative Invariance Learning . . . 232
7.5 Methods: the emergence of spike time continua . . . 232
7.5.1 Model architecture and equations . . . 232
7.5.2 Test and training paradigm . . . 234
7.5.2.1 Testing the network . . . 234
7.5.2.2 Training the network . . . 234
7.5.2.3 Test-train-test order . . . 234
7.6 Results: the emergence of spike time continua . . . 235
7.6.1 Individual Simulations . . . 235
7.6.1.1 Number of spiking test trials . . . 235
7.6.1.2 Fully reliable neurons . . . 236
7.6.1.3 Reliable responses over continuous ranges of stimulus transformations . . . 237
7.6.1.4 The effect of location on the number of neurons responding to two stimuli . . . 238
7.6.1.5 Individual Simulation - Space Time Funk Plots . . . 239
7.6.1.6 Medium forward . . . 242
7.6.1.7 Slow fizzle . . . 242
7.6.1.8 Fast forward . . . 242
7.6.1.9 Split response example . . . 243
7.6.1.10 Summary . . . 243
7.6.1.11 Individual Simulation - Space Time Funk Plots (Weight Differences) . . . 243
7.6.1.12 PSTH . . . 244
7.6.1.13 Mean spike counts . . . 244
7.6.1.14 Polynomials fit to chunks . . . 244
7.6.1.15 Weight histograms . . . 245
7.6.2 Effect of parameters . . . 246
7.6.2.1 Mean reliable spike times . . . 246
7.6.2.2 Mean reliable chunk length . . . 246
7.6.2.3 Mean synaptic weight . . . 246
7.6.2.4 Mean 1st lag chunk autocorrelation r-value . . . 246
7.6.2.5 Polynomial residual . . . 248
7.6.2.6 Polynomial 1st-3rd exponents . . . 248

7.7 Discussion 248

Bibliography **270**

Abstract

The thesis explores the nature of information representation by patterns of action potentials in the cortex. The thesis combines analysis of sensory-evoked in vivo cortical spiking activity, large scale biophysical cortical network simulations, phenomenological modelling of representation learning through spike-time-dependent plasticity and the development of a spiking neural network simulator. Together, these approaches engendered novel findings and hypotheses regarding the nature of cortical processing.

The thesis focuses on patterns of spikes in which each neuron spikes once ('multi-neuron single spike patterns') as a possible fundamental form of cortical information encoding. The thesis characterises the temporal structure of stimulus-evoked multi-neuron single spike patterns in response to single whisker deflections in the barrel cortex of the anaesthetised rat. Spike time patterns are found to be stretched and compressed in time (i.e. time-warped) depending on the excitability of the barrel cortex on a single trial. This partly explains why precise spike time patterns are rarely reported in vivo. It is proposed that downstream neurons may decode such representations in an excitability-dependent manner, thus increasing spike time coding capacity.

The characterised form of spike time variability is tested for in a detailed biophysical model of a column of rat somatosensory cortex. It was hoped that the model could be used to understand the phenomena's underlying mechanisms and to make further predictions. A framework was developed for comparing in vivo and in silico spiking activity. Key differences between in vivo and in silico spiking activity were found, however. These differences have been used to guide further model refinement.

Hypotheses of how time-warped multi-neuron patterns of single spikes may facilitate the brain's solution to the Feature Binding Problem are presented, as well as attempts to test these hypotheses in a phenomenological spiking neural network. The latter made up the initial work of the DPhil, which aimed to explore the emergence of spike time continua. Towards this a GPU spiking neural network simulator was co-created, and was demonstrated (by collaborators) to be the fastest of its class.

List of Figures

1.1 (a) Illustration of spike count stimulus-response distribution space for two neurons. Red points illustrate sample stimulus-response distribution. Blue ellipse illustrates the boundary of the true underlying stimulus-response distribution. (b) Illustration of two non-overlapping stimulus-response distributions for (left) spike counts and (right) relative spike time differences. (c) Illustration of shared spike count variability over repeated trials of a condition for a neuron pair (left), for a population (center) and exhibited as alternative single trial trajectories through low-dimensional space (right). 6

1.2 Possible latent processes underlying the noise correlations include adaptative (left) and/or non-adaptative (right) changes in the shared excitability-level of the two neurons from trial-to-trial. 11

1.3 (a) Example of two alternative Necker cubes defined by their properties and constituent low-level features. The top cube has all black edges except for one blue edge, whilst the red cube has all red edges. Each cube has a different orientation. The Necker cube illusion: the perception of a single cube (e.g. the top cube) transiently changes between the perception of two alternative cubes. (b) For two neurons, illustration of two overlapping multi-neuron spike patterns (red and purple). Coincidence detection could play a part in robustness to 'false' spike patterns (black). That is downstream neurons could respond selectively to particular spatiotemporal patterns of spikes. This may be achieved through the spatial organisation of neurons, synaptic innervation patterns and variable conduction delays, such that spike time patterns corresponding to spatially and temporally localised activations of the postsynaptic neurons are more likely to make the post synaptic neurons spike. (c) Illustration of false combinations problem. Particularly, if a neuron or cell assembly represents the presence of a feature over a range of transformations, the false combinations problem asks how a neuron or cell assembly is robust against firing to incorrect combination of constituent low-level features. 15

1.4 (a) Left: Illustration of continua of first spike time distributions, corresponding to first spike times (not relative differences) under fixed-state-decoding (blue) and state-dependent-decoding (green). Right: Illustration of continua of relative spike time difference distributions, which could correspond to fixed-state-decoding and state-dependent-decoding distributions. (b) Illustration of the representation of novel stimulus relative to previously experienced stimuli along relative difference continua. (c) Illustration of possible correspondence between single trial variability in low-dimensional population dynamics and single trial spike times for correlated clusters. 21

1.5 (a) Illustration of multi-neuron first spike patterns (left) and single spike patterns (right). (b) Left: Illustrative (artificial) first spike stimulus-response distribution for two neurons and a single stimulus condition. Each point represents the first spike times of the two neurons on a trial on which both neurons spiked. A spike time noise correlation is observable within the isolated cluster. The vertical distance from a point to the line $y = x$ represents the relative spike time difference for the first spike pair. Right: If the angle of the noise correlation is 45°, the relative spike time difference remains fixed (with random noise) w.r.t. the latent process/processes underlying the correlation. If the angle is different from 45°, the relative spike time difference varies w.r.t. the process/processes. (c) Fixed relative spike time differences vs time-warped spike time difference dependence on non-adaptative single trial shared excitability. (d) Illustration of correlated predicted cortical state for disjoint pairs of neurons. 25

3.1 Explanatory figures. (a) Illustration of shared spike count variability over repeated trials of a condition for a neuron pair (left), for a population (center) and exhibited as alternative single trial trajectories through low-dimensional space (right). (b) Illustration of multi-neuron first spike patterns (left) and single spike patterns (right). (c) Left: Illustrative (artificial) first spike stimulus-response distribution for two neurons and a single stimulus condition. Each point represents the first spike times of the two neurons on a trial on which both neurons spiked. A spike time noise correlation is observable within the isolated cluster. The vertical distance from a point to the line $y = x$ represents the relative spike time difference for the first spike pair. Right: If the angle of the noise correlation is 45°, the relative spike time difference remains fixed (with random noise) w.r.t. the latent process/processes underlying the correlation. If the angle is different from 45°, the relative spike time difference varies w.r.t. the process/processes. (d) The latent process underlying the noise correlation may be due to adaptative (left) and/or non-adaptative (right) changes in the shared excitability-level of the two neurons from trial-to-trial. (e) Fixed relative spike time differences vs time-warped spike time difference dependence on non-adaptative single trial shared excitability. (f) Illustration of correlated predicted cortical state for disjoint pairs of neurons. 53

3.2 Reprinted from [59], with permission from Elsevier. Schematic of the topographic correspondence between whiskers and cortical barrels. The whiskers on the rodent snout are organised in a grid-like layout with each whisker assigned a label based on its row and column. The contralateral barrel cortex is organised topographically in a similar grid-like fashion, with a distinct cortical barrel corresponding to each whisker. 57

3.3 Reprinted from [80], with permission from Elsevier. Original caption: 'Schematic demonstrating the ascending pathways in the rodent whisker-related sensorimotor system, as well as possible connections associating the sensorimotor cortical signals. The ascending pathway starts with the primary afferents in the trigeminal ganglion (TG) transducing whisker vibrations into neuronal signals and projecting to the trigeminal brainstem complex (TN). The TN consist in the principal nucleus (PrV), and the spinal sub-nuclei (interpolaris SpVi; caudalis SpVc; the detailed connectivity of the oralis sub-nucleus is unknown and is omitted in the figure). The SpVi falls into a caudal and rostral part (SpVic and SpVir). The classical mono-whisker lemniscal pathway (lemniscal 1) originates in PrV barrelettes, and projects via VPM barreloid cores to primary somatosensory cortex (S1) barrel columns. A second lemniscal pathway originating from PrV has been recently discovered which carries multi-whisker signals via barreloid heads to septa (and dysgranular zone) of S1. The extra-lemniscal pathway originates in SPVic and carries multi-whisker signals via barreloid tails in VPM to the secondary somatosensory area. Finally the parelemniscal pathway originates in SpVir and carries multi-whisker signals via POm to S1, S2, and primary motor area (M1). The different colors of connections indicate three principal pathways through which associative coupling between the sensorimotor cortical areas may be realized. Black: direct cortico-cotical connections. Blue: cortico-thalamic cascades. Brown: cortico-sub- cortical loops. Projections of S1 and S2 may open or close the lemniscal gate (i.e. gate signal flow through PrV) by modulating intrinsic TN circuitry.' 58

3.4 Onset of cortical activation following whisker stimulation. (a) Histogram of single unit spike times to all stimuli with bins of 0.2 ms and between -50 ms before and +100 ms after stimulus. The histogram was smoothed with a Gaussian filter (width s.d. 1 ms). The mean and standard deviation of the smoothed bin counts was calculated for the period before stimulus onset at 0 ms. The cortical activation onset spiking threshold was determined as this mean value + 3 standard deviations. Cortical onset was determined as the time point when the smoothed histogram went above this threshold (red dot). The left and right figures show the same smoothed histogram between -50 and +100 ms (left) and 0 and +25 ms (right). (b) Table displaying the number of spike sorted neurons recorded over all experimental sessions by layer and E/I type. Note that L2/3 is undersampled (see Fig. 6b) due to sparse activity and responsiveness of L2/3 neurons [153] 65

3.5 Properties of initial responses of cortical neurons to sensory stimulation. (a) Max normalised post-stimulus time histograms (PSTHs) over a 50ms window following stimulus onset (averaged over all trials) for all 8 neuron groups (excitatory in red, inhibitory in blue) in different cortical layers (L) and a 1Hz stimulus. Smoothed with a Gaussian kernel ($\sigma = 1.5ms$) (b) Normalised spike count histogram over all neurons, all trials and stimuli over 50ms following cortical activity onset. Zero spike trials are not included. (c) Histogram of single neuron single stimulus condition did spike once reliabilities. Each value represents the proportion of trials of a stimulus condition that a single neuron spiked at least once. 67

3.6 Confirmation of neural stimulus-response adaptation (a) By stimulation frequency. Left: For each single neuron and stimulus condition combination, spike counts are normalised (combination mean subtracted followed by division by combination standard deviation). Means of these normalised spike counts are calculated for stimulation frequency groups (0-0.2, 1, ..., 10Hz) and 10 trial bins (i.e. trials 0-9, ..., 90-99). These mean values are plotted with a 95% confidence interval by stimulation frequency. Right: Same as left but for normalised first spike times. (b) Same as (a) but by neuron group. Values for excitatory and inhibitory neuron groups shown in red and blue respectively. Note that L2/3 is undersampled (see Fig. 3.4b) due to sparse activity and responsiveness of L2/3 neurons [153]. 68

3.7 Heterogeneous first spike rasters following cortical onset for five single-unit single-stimulus combinations. L4 INH (i) and L5A EXC unit (ii) demonstrate reliable short-latency response between 5 and 10 ms poststimulus with prominent adaptation. L2/3 EXC units (iii and iv) also demonstrate adaptation but responses are more variable. In contrast, L5A EXC unit (v) shows variable responses but not spike-time adaptation. 69

3.8 Single-unit single-stimulus first spike autocorrelations and partial autocorrelations. Autocorrelation and partial autocorrelation analysis allows quantification of further autocorrelative structure in the trial-to-trial sequence of first spike times. This figure considers the first spike times s_t of single neurons on spiking trials $t = 1, \ldots, T$ of single stimulus conditions where T is the number of spiking trials. Lag-n autocorrelations and partial autocorrelations test the correlations and partial autocorrelations between spike times on trials t with spike times on trial t+n. Left and centre plots show the normalized cumulative sums of lag 1-5 autocorrelation (left) and lag (2-5) partial autocorrelation (center) p-values for 2416 single unit single stimulus combinations. Autocorrelations (right) were above chance for lags 1-4 supportive of general adaptative trends. This is demonstrated by the fact that the cumulative sums are slightly above $y = x$, demonstrating non-uniform distributions of p-values. As this is hard to see, accompanying tables show results (χ^2 statistics and p-values) of Fisher-method application to the p-values produced for each lag. This confirms that autocorrelations and partial autocorrelations above chance for several values. 70

3.9 Example clustering of sample first spike response distribution for two units (L5AI, L4I) for single epsilon value over a 50ms window following cortical onset (1 outlier out of view). (a) First spike response distribution for two neurons (L5AI, L4I) over 50ms following cortical onset (1 outlier out of view). Point colours illustrate DBSCAN clustering (epsilon = 1.35; Methods). Green flat ellipse illustrates 4σ confidence interval of Stage 1 cluster. Dark blue angled ellipse illustrates 4σ confidence interval of Stage 2 cluster. Light blue ellipse illustrates 4σ ellipse of all spike pairs under independence and normality assumption. (b) Illustration of detected cluster in first spike pair response distribution (left) and unclustered first spike pair response distribution (right). (c) Point colours illustrate DBSCAN clustering (epsilon = 1.35). Black points are outliers. Dashed ellipses represent 3σ intersection ellipse for each DBSCAN cluster. Setting the DBSCAN 'minimum number of elements in a cluster' parameter to 2 increased the cluster quality control, as intersection ellipses were created for 2 element clusters rather than assigning them as outliers. 71

3.10 Correlation surrogate control dataset generation. Left: Example of original first spike pair response distribution for repetitions of a single stimulus condition and pair of neurons. Right: Generated surrogate dataset examples: (i) original response distribution spike times shuffled and (i) spike times sampled from normal distribution with same single neuron standard deviations as original response distribution. 75
3.11 Histogram of Stage 1 cluster end points (mean + 4σ) for each Stage 2 cluster. 76
3.12 Correlation of Stage 1 clusters (left) and unclustered response distributions (right). Top: Linear regression p-values for Stage 1 clusters (left) and unclustered response distributions (right). Center: Normalised cumulative sums of p-value histograms for (blue) Stage 1 clusters (left) and unclustered response distributions (right), (green) surrogate control dataset 1, (red) surrogate control dataset 2. Surrogate controls were generated separately from correlated response distributions for the clustered and unclustered cases. Bottom: Fisher method confirms that Stage 1 clusters and unclustered correlations are not detected above chance-levels for both correlation surrogate control datasets. . 78
3.13 Stage 2 clustered and unclustered correlations are positively correlated. Top: number of correlations ($p < 0.005$) detected for the clustered (Stage 2 clusters) and unclustered case. Intersection shows the number of response distributions from which a correlation was detected by the clustering algorithm and for the unclustered case. Bottom: linear regression r-values of correlated ($p < 0.005$) Stage 2 clusters (left) and unclustered response distributions (right). 79

3.14 Stage 2 cluster statistics. (a) Left: histogram of variance explained (r^2) by the spike time linear relationship of each Stage 2 cluster (dark blue) and the linear relationship of the unclustered response distributions from which Stage 2 clusters were extracted (grey). Right: histogram of change in linear regression p-values introduced by clustering for response distributions from which Stage 2 clusters were extracted. Particularly, histogram shows Stage 2 cluster linear regression p-values minus originating response distribution linear regression p-values. Clustering reduces p-values in most cases (dark blue). (b) Histogram showing for each Stage 2 cluster, the proportion of conjunctive (both neurons spike) trials contributing a cluster sample (grey) and the proportion of stimulus trials contributing a cluster sample (dark blue). (c) Histogram of r^2-values of linear multi-output spike pair prediction from trial index (grey) and linear correlation (dark blue, blue) for Stage 2 cluster spike times (left) and unclustered response distributions (right). (d) Histograms of linear regression p-values of the first neuron's spike time and the spike time difference. 80

3.15 Comparison to alternative clustering algorithms. (a) Example detection of stationary correlated cluster (dark blue) in shuffled response distribution through fragmentation by the standard GMM algorithm. (b) Cumulative p-value histograms of correlations detected by the Standard GMM algorithm. Fisher method confirms that Stage 1 clusters are detected by the Standard GMM algorithm above chance-levels for both correlation surrogate control datasets. (c) Cumulative p-value histograms of correlations found for custom GMM algorithm. (d) Histogram showing for each Stage 2 cluster found by the custom GMM algorithm the proportion of conjunctive (both neurons spike) trials contributing a cluster sample (blue) and the proportion of stimulus trials contributing a cluster sample (green). 81

3.16 Angle surrogate control dataset. (a) Left: Generation of angle surrogate control dataset (i) example. Stage 2 cluster detected by main clustering algorithm (green ellipse). Points within 6σ circle of cluster mean (where σ is equal to the mean of the cluster spike time standard deviations of the two neurons) are deleted and replaced by an equal number of samples from a 45° Gaussian ellipse (right). New response distribution added to surrogate control dataset (i). (b) Mean θ_{45} angles with 95% empirical confidence intervals for positively correlated Stage 2 clusters detected in angled surrogate control dataset. Blue if significantly greater than 0° ($p < 0.025$) and less than 45° ($p < 0.025$). Red otherwise. Detection of clusters different from 45° appear to be at chance-levels. (c) Histogram of p-values for correlation angles being different from 45° for Stage 2 clusters detected by the main clustering algorithm for the generated angle surrogate control dataset. Correlation angles different from 45° are not found above chance-levels. . . . 84

3.17 Secondary cluster correlation angle test. (a) Example surrogate control response distribution generated by drawing samples from a 2D 45° Gaussian ellipse centred at the original Stage 2 cluster mean (left) and a second uncorrelated response distribution with a mean at a 15° angle to the left cluster (right). (b) Mean θ_{45} angles with 95% empirical confidence intervals for positively correlated Stage 2 clusters detected in secondary cluster correlation angle surrogate control dataset. Blue if significantly greater than 0° ($p < 0.025$) and less than 45° ($p < 0.025$). Red otherwise. Detection of clusters different from 45° appear to be at chance-levels. (c) Histogram of p-values for correlation angles being different from 45° for Stage 2 clusters detected by the main clustering algorithm for the generated angle surrogate control dataset. Correlation angles different from 45° are not found above chance-levels. (d) Mean 45° angles with 95% confidence intervals for positively correlated ($p < 0.005$) unclustered response distributions. As expected, the unclustered approach does not detect the angle of the 45° cluster. Instead, correlation angles are found around 15°. 85

3.18 Correlation angles. (a) Mean θ_{45} angles with 95% empirical confidence intervals for positively correlated Stage 2 clusters (top) and unclustered response distributions (bottom). Green if significantly greater than 0° (p-value ¡ 0.025) and less than 45° ($p < 0.025$). Grey otherwise. (b) Histogram of cluster angle 95% empirical confidence intervals. 86

3.19 Custom GMM correlation angles. (a) Histogram showing for each Stage 2 cluster found by the custom GMM algorithm the proportion of conjunctive (both neurons spike) trials contributing a cluster sample (blue) and the proportion of stimulus trials contributing a cluster sample (green). (b) Angles of Stage 2 clusters found by the custom GMM algorithm. Mean 45° angles with 95% confidence intervals for positively correlated Stage 2 clusters (top) and stationary Stage 2 clusters (bottom) Blue if significantly greater than 0° ($p < 0.025$) and less than 45° ($p < 0.025$). Red otherwise. (c,d) Mean 45 angles with 95% empirical confidence intervals for positively correlated Stage 2 clusters detected by Custom GMM algorithm in secondary cluster correlation angle surrogate control datasets (generated from response distributions in which Custom GMM algorithm originally detected Stage 2 clusters). Blue if significantly greater than 0° ($p < 0.025$) and less than 45° ($p < 0.025$). Red otherwise. (c) Angle correlation datasets (i). (d) Angle correlation datasets (ii). All detected clusters were not significantly different from 45° for both datasets. 88

3.20 Stationarity of single neurons spike times for Stage 2 clusters (top) and correleated unclustered response distributions (bottom). (a) Correlation p-values of Stage 2 cluster single unit spike times with trial index. (b) Correlation p-values of angled cluster single unit spike times with trial index (neuron A vs neuron B). For some clusters, the cluster spike times of neither neuron were correlated with trial index, whereas for other clusters the spike times of at least one neuron in the cluster pair was correlated with trial index. (c) The Kwiatkowski-Phillips-Schmidt-Shin trend stationarity test (KPSS) was used to test whether the cluster first spike times of single neurons were stationary around a deterministic trend. For the 268 Stage 2, 102 of the corresponding 536 (2*268) single unit cluster first spike sequences were determined as not trend stationary (KPSS, $p < 0.05$). Python 'statsmodels' implementation of p-value estimation rounds all $p > 0.1$ to 0.1. 89

3.21 Examples of stationary and non-stationary clusters, and explained variance. (a) Example of stationary (left) and non-stationary cluster (right) determined by stationarity criteria (2)-(5) (see Methods). Top row shows single neuron cluster spike times for each neuron (green and blue) plotted against cluster trial index. Bottom shows the cluster spike times of neuron A and B plotted against each other. Non-stationary cluster shows a clear trial dependence. (b) Top: Histogram of variance explained (r^2) by the spike time linear relationship of each stationary Stage 2 cluster (dark blue) and the linear relationship of the unclustered response distributions from which stationary Stage 2 clusters were extracted (grey). Bottom: r^2 values shown for stationary unclustered response distributions 91

3.22 Stationary Stage 2 cluster single-unit single-stimulus first spike (a) autocorrelations, (b) partial autocorrelations, (c) cross-correlation and (d) partial cross-correlation. Autocorrelation, partial autocorrelation, cross-correlation and partial cross-correlation analysis allows quantification of autocorrelative and cross-correlative structure in the trial-to-trial sequence of first spike times and first spike pairs for the 117 stationary Stage 2 clusters. This figure considers the first spike times $s_{j,t}$ of neurons j=0,1 on cluster trials t=1,. . . ,T where T is the number of cluster trials for a single cluster. Lag-n autocorrelations and partial autocorrelations test the correlations and partial autocorrelations between single neuron cluster spike times on trials t with spike times on trial t+n. Plots (a) and (b) show the normalized cumulative sums of lag 1-5 autocorrelation and lag (2-5) partial autocorrelation p-values for 214 (=2 * 117 stationary Stage 2 clusters) single neuron single cluster spike time sequences. Lag-n cross-correlations and partial cross-correlations were similarly tested between neuron pair cluster spike times on trials t with spike times on trial t+n. Autocorrelations, partial autocorrelations, cross-correlations and partial cross-correlations are at chance-levels for lags 1-5. This is demonstrated by the fact that the cumulative sums lie on y=x, demonstrating uniform distributions of p-values. Accompanying tables show results (χ^2 statistics and p-values) of Fisher-method application to the p-values produced for each lag, which confirms that autocorrelations, partial autocorrelations, cross-correlations and partial cross-correlations are not above chance. 92

3.23 Stationary clusters. Predicted θ_{45} angles with 95% empirical confidence intervals for stationary Stage 2 clusters (top) and stationary correlated unclustered response distributions. Angles significantly different from 0° ($p < 0.025$) and 45° ($p < 0.025$) coloured green. Grey otherwise. . 93

3.24 State-dependent coding predictions and relative difference regimes. (a) Illustration of two non-overlapping stimulus-response distributions for (left) spike counts and (right) relative spike time differences. (b) Sampled first spike response distribution for two neurons (L5B E, L5A E). Point colours illustrate DBSCAN cluster (epsilon = 5.0). Dark blue angled ellipse illustrates 4σ boundary of Stage 2 ellipse (i.e. the area for which the standard deviation of relative differences is calculated under RDR 2). Coloured circles illustrate the 4σ boundary of the state-dependent response distribution predicted by factor analysis for single state values ($< 0, 0, > 0$). The standard deviation of the predicted state-conditioned first spike relative difference distributions (for RDR 3) is calculated based on the size of this area (Methods). (c) The standard deviation of differences of RDR 1 vs RDR 2 (top) and RDR 2 vs RDR 3 (bottom) for criteria fulfilling stationary, differenced and ARIMA modelled clusters. (d) RDR 2 vs RDR 3 (bottom) for criteria fulfilling stationary, differenced and ARIMA modelled unclustered response distributions. 94

3.25 Cluster angle comparisons. (a) PCA angles vs factor analysis angles for the criteria fulfilling clusters. (b) Original PCA angles vs differenced PCA angles for the criteria fulfilling differenced clusters. (c) PCA angles vs factor analysis angles for the criteria fulfilling differenced clusters. (d) Original PCA angles vs ARIMA PCA angles for the criteria fulfilling ARIMA clusters. (e) PCA angles vs factor analysis angles for the criteria fulfilling ARIMA clusters. 96

3.26 Time series modelling techniques. (a) For an example cluster, the original spike times do not fulfil stationarity criteria (2)-(5) (Methods) but do after the application of differencing and ARIMA modelling. The factor analysis predicted state-dependent response distributions are therefore presented for the cluster after differencing and ARIMA modelling (4σ ellipse; bottom centre and bottom right). It is clear that the adaptative trends have been removed by differening and ARIMA but the clusters remain correlated with similar angles to the original cluster. (b) Predicted angles and 95% empirical confidence intervals of differenced and ARIMA modelled criteria fulfilling Stage 2 clusters. Angles significantly different from $0°$ and $45°$ coloured green. Grey otherwise. 99

3.27 Two examples (left and right) of discontinuous and non-monotonic non-stationarity, highlighting the need for differencing over parametric models for the removal of non-stationary trends when testing for non-adaptative pairwise correlations. Top row shows two first spike pair response distributions in which Stage 2 clusters were detected by the main clustering algorithm. Bottom row shows cluster spike times of the two corresponding neurons (red and blue) vs trial index for the original (left), differenced (center) and ARIMA (right) cases. For both of the examples, the original spike times of both neurons show a discontinuity in their relationship with trial index (at around trials 70 and 90 respectively). At these points of discontinuity the latencies of both neurons in both examples become slower, before gradually becoming faster. For the removal of non-stationarities in order to test for pairwise correlations, the non-stationarities of examples such as these could not be modelled with typically used parametric models typically used for adaptative trend removal. Moreover, use of classic parametric models would risk introducing adaptation-independent correlations between the spike times of the two neurons where none existed previously. Alternatively, after the application of differencing and ARIMA, the cluster spike times show no relationship with trial index, and were deemed stationary by the stationarity criteria. This highlights an advantage of differencing based methods over parametric models, for the purposes of this paper. 100

3.28 Negative autocorrelations (left) and negative partial autocorrelations (right) introduced by differencing. Plots for criteria fulfilling differenced clusters (top, centre) and unclustered (bottom) (with the autocorrelation criterion relaxed). As with previous autocorrelation plots (i.e. Fig. 3.22) but p-values instead represent the probability that single neuron spike times are negatively correlated for the criteria fulfilling differenced Stage 2 clusters and unclustered response distributions. Cumulative sum plots and accompanying tables show that spike times are negatively autocorrelated and negatively partially autocorrelated above chance for several lags, as expected by differencing. 101

3.29 Predicted angles and 95% empirical confidence intervals of differenced and ARIMA modelled criteria fulfilling Stage 2 clusters and unclustered response distributions. Angles significantly different from 0° and 45° coloured green. Grey otherwise. 103
3.30 Order of AR and MA models fit to the single neuron spike times of differenced clusters and unclustered response distributions. 104
3.31 Autocorrelations (left) and cross-correlations (right) at chance-levels for criteria fulfilling ARIMA modelled Stage 2 clusters (top, centre) and unclustered response distributions (bottom). As with Fig. 3.22 but for criteria fulfilling ARIMA modelled clusters and unclustered response distributions. Autocorrelations and cross-correlations are also at chance levels confirming that ARIMA modelling successfully removed autocorrelations introduced by differencing and that there is no additional cross-correlative structure between the neurons in ARIMA modelled clusters and unclustered response distributions. 105
3.32 Linear regression r^2 values for criteria fulfilling differenced (left) and ARIMA modelled (right) clusters (top) and unclustered response distributions (bottom). 106
3.33 Histogram showing number of stationary Stage 2 correlations and criteria fulfilling ARIMA modelled correlations (by stimulation frequency) for clusters (left) and unclustered response distributions (right). . . . 107

3.34 Neuron pair spatial separation (unclustered). Neuron pair spatial separation for criteria fulfilling correlated clusters (left) and unclustered response distributions (right) (both stationary and ARIMA modelled). Data from all experimental sessions. Line colour and thickness indicates the number of criteria fulfilling correlations (clusters or unclustered response distributions) found for pairs of neurons in different layers and barrels and of different E/I-type. Red and blue correspond to E and I neuron types respectively. For example, the thick dark green line illustrates that there were 14 criteria fulfilling correlated clusters, for which both neurons were in L5B, 1 neuron was in the principal column and 1 was in a neighbouring column. The numbers in the circles indicate the number of pairs of the same E/I type, layer and barrel. The left most column is shown to illustrate pairs for which neurons are on opposite sides of the principal column. Counts are pooled for stationary clusters and non-stationary clusters made stationary by ARIMA modelling. 32 and 25 neuron pairs were not to added to each plot because at least one of the neurons in the pair lied in a septum (i.e. between barrels). 108

3.35 Discussion figures. (a) Suggested role of relative difference modulation in balancing spike-timing-dependent plasticity. A fixed spike time is shown for the presynaptic neuron, although the figure explains why the presynaptic neuron's spike time should also be modulated. (b) Illustration of possible correspondence between single trial variability in low-dimensional population dynamics and single trial spike times for correlated clusters. (c) Left: Illustration of continua of first spike time distributions, corresponding to first spike times (not relative differences) under RDR 2 (grey) and 3 (dark blue). Right: Illustration of continua of actual relative spike time difference distributions, which could correspond to RDR 2 or 3. 109

3.36 From [29]. Reprinted with permission from AAAS. (a) Uncaging spots (yellow) along the selected dendrite. (b) Somatic responses to IN (red) and OUT (blue) directions at 2.3 mm/ms (averages denoted by bold lines). (c) Relation between peak voltage and input velocity (values normalized to the maximum response in the IN direction for each cell, n = 15). Error bars indicate SEM. (d) Relation between direction selectivity and input velocity (n = 15). (e) Direction selectivity of spike probability (**P = 0.0013, n = 7). (f) Relation between spike probability and velocity (n = 7, average of both directions). 115

3.37 From [29]. Reprinted with permission from AAAS. EPSP supralinearity varies in a graded manner. (a) Spatially clustered activation of 6 spines produced a fast dendritic spike detectable as a notch in the somatic voltage recording, and as a sharp peak in the corresponding dV/dt trace(4, 5, 32, 33) (left; 4 ms interval; all inputs within 15 m); however no fast dendritic spike is produced when the same number of inputs are distributed along the length of a dendritic branch during a direction-selective response (right; 4 ms interval). (b) EPSPs recorded at different membrane potentials produced by sequential activation of 8 dendritic spines at 1 ms intervals; membrane potential was varied with somatic current injection. (c) Supralinear peak voltage (calculated by subtracting the sum of individual spine responses) gradually increases as Vm becomes more depolarized. Arrow indicates resting membrane potential. Pooled data from 5 cells; membrane potential has been normalized to the most depolarized value tested (mean resting potential at arrowhead). (d) Summary plot showing that the supralinear component of the EPSP varies gradually with the velocity and direction of uncaging (normalized to the IN maximum in control cells, P = 0.0410 and P = 0.0348 for IN and OUT respectively, ANOVA test). While at high input velocities (mean 6.3 ± 0.83 m/ms, range indicated by yellow line above X-axis) responses were 44.7 ± 2% of the maximum, at lower input velocities (1.4 ± 0.06 m/ms, green line above X-axis) the supralinearity decreased to 27.8 ± 4% (P ¡ 0.0001; t-test, average of both directions). Also, at the optimal speed for direction sensitivity of each cell, the IN EPSPs were on average 53 ± 3% of the maximum and the OUT responses 33 ± 8% (P ¡ 0.0001; paired t-test). In A and B, traces are averages of 5 sweeps. 116

3.38 From [29]. Reprinted with permission from AAAS. (a) (Left) Activating more synapses (1 to 7) produces EPSPs that are progressively larger than the linear sum of the individual responses. (Right) Summary plot of eight cells. Solid line, experimental data points; dashed line, unity line. Error bars indicate SEM. (b) Velocity and direc- tion sensitivity is abolished in AP5. (c) Comparing peak amplitudes for IN and OUT sequences at the optimal velocity for direction selectivity. (d) Reduced velocity sensitivity in AP5. (e) Random patterns of nine inputs were generated and sorted by their directionality index (35). Four sample patterns (numbers showing activation order) and corresponding directionality index are shown. Patterns were played into a single dendrite in a pyramidal cell model or (f) experimentally in single pyramidal cell dendrites. The peak response is different for different patterns, and responses become similar with hyperpolarization (or without NMDA receptors in the model). (Inset) Dye-filled neuron with dendrite used in the experiment enclosed by a yellow box. Vm, membrane potential. (g) Peak EPSP voltage measured experimentally depends on input-pattern. Asterisks indicate significant differences between EPSP peaks for the example shown; hyperpolarization reduces the ability to discriminate different patterns. Error bars indicate SEM. 118

3.39 From [29]. Reprinted with permission from AAAS. Local impedance differences result in different levels of NMDAR conductance. (A) A cluster of 7 spines was placed either at the tip of the dendrite (green dots) or at the base of the same branch (orange dots), and activated with 1 ms intervals while local Ca2+ influx was monitored by a line scan across the dendrite (indicated by green and orange lines). (B) Activation of distal synapses produces a larger Ca2+ transient than proximal synapses. (C) Across 5 cells, proximal synapses elicited 34 $\pm$ 6% of the peak distal Ca2+ influx (P = 0.0003, paired t-test). 119

3.40 From [29]. Reprinted with permission from AAAS. Effect of inhibition on single dendrite sequence discrimination. Excitatory synapses on a single dendrite were activated in the model with random sequences, and different types of inhibition were evaluated (see Methods for details). (A) Histograms for the peak EPSP amplitude generated under the different conditions (blue - no inhibition; green - somatic inhibition; yellow - dendritic inhibition), showing that strong feedforward dendritic inhibition can shift the EPSP histogram to smaller values, while the variance is maintained. Feedback inhibition has minimal impact on the discrimination of excitatory sequences. (B) Somatic voltage traces in response to the same sequences, while varying the location and type of inhibition. 120

4.1 Neuron pair cortical state correlation. (a) Example of two clusters which do not share a common neuron but have correlated cortical states over trials of a single stimulus. Left and Center: cluster spike times of two clusters coloured by predicted cortical state (value along first principal component). Right: correlated cortical states on trials on which all 4 neurons spike. (b) Right: histograms showing p-values of linear correlation between predicted states of pairs of correlated pairs of neurons, which do not share a common neuron. Inclusion criteria: 15 conjunctive trials. Cases: clustered (orange), unclustered (green) and clustered and unclustered pooled (blue). Left: r-value histograms of linear state correlations (clustered and unclustered pooled). (c) Mean variance explained by each PCA dimension between groups of 4 neurons (on correlated trials that all 4 neurons spiked). 127

4.2 State-conditioned spike times and differences. (a) Left: p-value histogram of correlations between predicted states of single correlated neuron pairs with the spike times of other correlated neuron pairs (clustered and unclustered pooled). Right: r-value histogram of statistically significant (p-value) correlations. (b) Same as (a) but for correlations with spike time differences with non-45° correlation angles. (c) Left: variability comparison. X-axis shows the standard deviation of single neuron spike times within stationary clusters (pooled). Y-axis shows the standard deviations of spike times conditioned on the predicted state of single correlated neuron pairs, which do not share a common neuron. Right: Comparison of predicted state-dependent spike time variability predicted by factor analysis (x-axis) vs spike time variability conditioned on the predicted state of single correlated neuron pairs, which do not share a common neuron (y-axis). (d) Same as (c) but for correlations with spike time differences with non-45° correlation angles. 128

4.3 Example p-value (a) and r-value (b) correlation matrices for an example stimulus condition (1Hz inter trial interval). Each square has either a green or orange border. A green border signifies that the correlation value is for pair of neurons from a single cluster in the pool. An orange border signifies that the correlation was tested for stationary spike times, obtained from other stationary correlated clusters for the same neuron pair. Left: significant correlations ($p < 0.05$) have a green background, with the shade of green corresponding to the p-value (see key). Non-significant correlation ($p >= 0.05$) have a black background. Right: the background colour of each square represents the r-value (blue for a negative correlation and red for positive; see key). (c) Histogram of non-pool correlations p-values. (d) Histogram of r-values of significantly correlated ($p < 0.05$) non pool pairs. 131

4.4 State offsets (left) and state correlation angles (right) before (top) and after (bottom) adjustment for an example stimulus condition. For the example stimulus condition there are 18 stationary correlated clusters. For the state correlation angle plots (right), angle values are only calculated (and shown) where there is a significant correlation ($p < 0.05$). 134

4.5 Histograms of inter cluster state offsets (left) and state correlation angles (right) before (top) and after (bottom) adjustment (meta stats over all stimulus conditions). 136

4.6 (a) Single trial states by trial index before (top) and after (bottom) adjustment procedure for 1Hz stimulus condition. Each coloured dot represents the single trial state predicted for a single stationary correlated cluster. Black dots represent the means of the states predicted from stationary correlated clusters on single trials. Black line traces through the mean values for visualisation purposes. Red, green, blue cycle is purely for visualisation purposes and has no meaning. (b) Original cluster single trial state predictions vs adjusted cluster single trial state predictions. 137
4.7 Left: Histogram of p-values for the linear correlation between trial index and mean population state, Right: Histogram of p-values for the lag-1 autocorrelation of the population mean state. 138
4.8 Population state-conditioned spike times and differences. For each neuron or neuron pair, for which the spike times or spike time differences were to be predicted, mean single trial states were calculated leaving out the clusters to which the neuron or neuron pair belonged. This is referred to as the 'leave one out' population state. (a) Left: p-value histogram of correlations between the leave one out population state and the spike times of stationary correlated clusters and unclustered response distributions. Right: r-value histogram of statistically significant (p-value) correlations. (b) Same as (a) but for correlations with spike time differences with non-45° correlation angles. (c) Left: variability comparison. X-axis shows the standard deviation of single neuron spike times of stationary clusters (pooled). Y-axis shows the standard deviations of spike times conditioned on the leave one out population state. Right: Comparison of predicted state-dependent spike time variability predicted by factor analysis (x-axis) vs spike time variability conditioned on the leave one out population state (y-axis). (d) Same as (c) but for correlations with spike time differences with non-45° correlation angles. (e) Histogram of the number of cluster state predictions available on a single trial, over all leave one out tests of population state correlation with single unit spike times. (f) Histogram of the number of cluster state predictions available on a single trial, over all leave one out tests of population state correlation with pairwise differences. 140

4.9 Time-warp visualisation for example stimulus condition. Top: shows the multi-neuron first spike pattern for each trial of a single stimulus condition. The colour of each multi-neuron single spike pattern represents its index when the trials are ordered by mean predicted state. Red represents trials with high predicted single trial state, and blue low. It can be seen that the spike times of the trials with high (red) single trial state spike earlier than the spike time with low (blue) single trial state. Bottom: Shows the same data after smoothing. 141
4.10 Same as Fig. 4.9 but for a different 1Hz stimulus condition. 142

5.2 Figures used by [175] to compare in vivo ([232]) and in silico [175] stimulus-evoked responses. (a) Top: in vivo (from [232]): 'Comparison of mean FR of EXC (triangles) and INH neurons (circles) before versus after whisker stimulation, plotted in logarithmic scale. Prestimulus time is defined as the 100 ms period before whisker stimulation and post-stimulus time as 0-30 ms afterward. Empty symbols represent neurons showing no significantly different activity in both periods, filled symbols neurons showing significantly different activity ($p < 0.05$). The neurons displayed in panel A are marked by their respective symbols.' Bottom: corresponding in silico plot from [175]. (b) Top: in vivo (from [232]): 'Statistical analysis of first-spike latencies for the different putative neuronal groups. Each box plot represents median, interquartile, and range of latencies across all recorded neurons. Empty circles represent outliers.' Bottom: Corresponding in silico plot from [175]. In vivo figures reprinted form [232] by permission of Oxford University Press. In silico figures reprinted from [175], with permission from Elsevier. 156

5.3 (a) Fig. 3a from [294]. Caption: 'Violin plot of firing rate distributions across cortical regions, in both spontaneous and task-aligned conditions. Width of bar is proportional to the fraction of neurons in that bin. Optimal histogram widths separately derived for each data-set'. Fig. 3b from [294]. Caption: 'Summary of data-sources in panel (a) and best-fit model to each data-set. Exponential, gamma, and log-normal distributions were fitted using maximum likelihood; best-fitting model was chosen using the Bayesian Information Criterion; p(best): the posterior probability of that being the best model within those tested. The Session column indicates whether the data were from a single recording session or pooled across sessions and/or animals.' (c) Fig. 4a from [294]. 'Distribution of spike counts across neurons in the auditory cortex (A1) of awake rats responding to a pure tone pip. Activity distributions are shown before ('spontaneous'), during ('early', 'late') and after ('offset') the tone pip. Best model fits are provided, using respectively an exponential (gray) and a log-normal (black) distribution. The log-normal provides a markedly better fit. Loose patch clamp recordings, from [125]. (d) Fig. 4b from [294]. 'Distribution of spike counts in the layers of barrel cortex (S1) of awake mice involved in a tactile detection task. The mean level of activity is plotted for each neuron against its cortical depth, both during ('localization + response') and between ('intertrials') periods of whisker contact with the stimulus. Loose patch clamp recordings from [203].' 165

5.4 Vertical locations of 300 thalamic fibres within the microcircuit (with respect to the centre of the microcicuit). For an example stimulus in which 10% of fibres are activated, red circles represent fibres which are part of the stimulus, whilst blue circles represent inactive fibres. . . . 169

5.5 In vivo and in silico spike rasters and max normalised PSTH profiles. Spiking activity is shown for the 6 comparison neuron groups for a 25ms period following cortical activity onset. Lines show values of max normalised PSTH profiles for each neuron group (1ms bins). (a) In vivo spike trains of neurons in the principal column in response to principal whisker deflections at 0.066-1.0Hz. For each neuron group a number of spike trains equal to the number of neurons in the corresponding in silico neuron group are plotted. Max normalised PSTH profiles are calculated for all available spike trains for the condition. (b) In silico spike trains for the in silico parameter combination: depolarisation = 100%, $[Ca^{2+}]_{\circ} = 0.9mM$, proportion of input fibers in the stimulus = 0.5 and input pattern width = 5.0ms. The spike trains of all neurons on a single trial are shown for the 6 neuron groups. Spike trains from all 4 trials are used to calculate the max normalised PSTH profiles. . 174

5.6 In vivo and in silico spike rasters and max normalised PSTH profiles. Spiking activity is shown for the 6 comparison neuron groups for a 100ms period following cortical activity onset. Lines show values of max normalised PSTH profiles for each neuron group (1ms bins). (a) In vivo spike trains of neurons in the principal column in response to principal whisker deflections at 0.066-1.0Hz. For each neuron group a number of spike trains equal to the number of neurons in the corresponding in silico neuron group are plotted. Max normalised PSTH profiles are calculated for all available spike trains for the condition. (b) In silico spike trains for the in silico parameter combination: depolarisation = 100%, $[Ca^{2+}]_{\circ} = 0.9mM$, proportion of input fibers in the stimulus = 0.5 and input pattern width = 5.0ms. The spike trains of all neurons on a single trial are shown for the 6 neuron groups. Spike trains from all 4 trials are used to calculate the max normalised PSTH profiles. 175

5.7 Comparison of mean firing rates between pairs of neuron groups. Each green dot represents the values for the in vivo case, whilst each blue dot represents the values for a single in silico parameter combination. Columns show comparison of mean firing rates between pairs of layer-wise inhibitory neuron groups (left column), between pairs of layer-wise excitatory neuron groups (central column) and between inhibitory and excitatory neuron groups within single layers (right column). 176

5.8 Heatmap showing the effect of the 4 parameters on the distance between in vivo and in silico PSTH profiles over 25ms following cortical activity onset. The 6 neuron group max normalised PSTH profile vectors were concatenated. The log of the Euclidean distance between the in vivo vector and the vector of each in silico combination gives a distance (the lower the better) for each in silico combination. Major columns: depolarisation, major rows: $[Ca^{2+}]_{\circ}$, minor columns: input temporal width, minor rows: proportion of fibers. 178
5.9 Heatmap showing the effect of the 4 parameters on the distance between in vivo and in silico PSTH profiles over 100ms following cortical activity onset. The 6 neuron group max normalised PSTH profile vectors were concatenated. The log of the Euclidean distance between the in vivo vector and the vector of each in silico combination gives a distance (the lower the better) for each in silico combination. Major columns: depolarisation, major rows: $[Ca^{2+}]_{\circ}$, minor columns: input temporal width, minor rows: proportion of fibers. 180
5.10 Analysis of spiking activity for in silico parameter combination 1 (depolarisation: 100%, $[Ca^{2+}]_{\circ}$: 0.9mM, fibers: 0.5, input temporal width: 5ms). (a) Lines show values of max normalised PSTH profiles for each comparison neuron group (1ms bins) for in vivo case (green) and in silico parameter combination (blue). (b) Spike count histograms for each comparison neuron group for the in vivo case (green) and in silico parameter combination (blue). (c) Box plots of spike latencies for each comparison neuron group for the in vivo case and in silico parameter combination. 182
5.11 Comparison of mean firing rates between pairs of neuron groups for in silico parameter combination 1 (depolarisation: 100%, $[Ca^{2+}]_{\circ}$: 0.9mM, fibers: 0.5, input temporal width: 5ms). Each green dot represents the values for the in vivo case, whilst each blue dot represents the value for the in silico parameter combination. Columns show comparison of mean firing rates between pairs of layer-wise inhibitory neuron groups (left column) and between pairs of layer-wise excitatory neuron groups (right column). 183

5.12 Heatmap showing the effect of the 4 parameters on the evoked 'mean neuron group firing rate', calculated by averaging the firing rates of the 6 neuron groups. Major columns: depolarisation, major rows: $[Ca^{2+}]_{\circ}$, minor columns: input temporal width, minor rows: proportion of fibers. 184
5.13 Heatmap showing the effect of the 4 parameters on the value of the evoked in vivo mean neuron group firing rate subtracted from the in silico mean neuron group firing rate. Major columns: depolarisation, major rows: $[Ca^{2+}]_{\circ}$, minor columns: input temporal width, minor rows: proportion of fibers. The values of in silico parameter combinations 1 and 2 are highlighted with a yellow box for visualisation purposes. 185
5.14 Heatmap showing the effect of the 4 parameters on the pre-stimulus spontaneous activity 'mean neuron group firing rate', calculated by averaging the firing rates of the 6 neuron groups. Major columns: depolarisation, major rows: $[Ca^{2+}]_{\circ}$, minor columns: input temporal width, minor rows: proportion of fibers. 187
5.15 Line plots showing the effect of each of the 4 parameters on the pre-stimulus spontaneous activity 'mean neuron group firing rate', calculated by averaging the firing rates of the 6 neuron groups. Each line shows the effect of varying the parameter of interest, whilst the other 3 in silico parameters are fixed. The horizontal green line shows the in vivo value for comparison. The four plots show the effect of (a) depolarisation, (b) $[Ca^{2+}]_{\circ}$, (c) the proportion of fibers and (d) the input temporal width. 188
5.16 Heatmap showing the effect of the 4 parameters on the value of the in vivo spontaneous activity mean neuron group firing rate subtracted from the in silico spontaneous activity mean neuron group firing rate. Major columns: depolarisation, major rows: $[Ca^{2+}]_{\circ}$, minor columns: input temporal width, minor rows: proportion of fibers. The values of in silico parameter combinations 1 and 2 are highlighted with a yellow box for visualisation purposes. 190

5.17 Comparison of spontaneous activity mean firing rates between pairs of neuron groups. Each green dot represents the values for the in vivo case, whilst each blue dot represents the values for a single in silico parameter combination. Columns show comparison of spontaneous activity mean firing rates between pairs of layer-wise inhibitory neuron groups (left column), between pairs of layer-wise excitatory neuron groups (central column) and between inhibitory and excitatory neuron groups within single layers (right column). 191
5.18 Heatmap showing the effect of the 4 parameters on the difference between the in silico and in vivo log of the mean neuron group stimulus-evoked/spontaneous firing rate ratio. The values of in silico parameter combinations 1 and 2 are highlighted with a yellow box for visualisation purposes. 193
5.19 Example first spike pair sample response distributions. For each neuron group, sampled first spike pair response distributions are shown for 3 random pairs of neurons in the neuron group. For each pair of neurons, the first spike times of the two neurons are plotted against each other for trials of the stimulus on which both neurons spiked at least once. 195
5.20 Stage 1 cluster correlations for original and surrogate control datasets by neuron group. First column: histograms of the linear correlation p-values of the extracted Stage 1 clusters. Central column: normalised cumulative sum of the linear correlation p-values of the extracted Stage 1 clusters for the original (blue) and surrogate control (red, green) datasets. Right column: the results of Fisher method analysis applied to the 3 datasets. Fisher method p-values show that the distribution of linear correlation p-values are not above chance. 196

6.1 (a) Illustration of possible correspondence between single trial variability in low-dimensional population dynamics and single trial spike times for correlated clusters. (b) Left: Illustration of continua of first spike time distributions, corresponding to first spike times (not relative differences) under non-state-dependent (grey) and state-dependent (dark blue). Right: Illustration of continua of actual relative spike time difference distributions, which could correspond to non-state-dependent or state-dependent continua. (d) For two neurons, illustration of two overlapping multi-neuron spike patterns (red and purple). Coincidence detection could play a part in robustness to 'false' spike patterns (black). That is downstream neurons could respond selectively to particular spatiotemporal patterns of spikes. This may be achieved through the spatial organisation of neurons, synaptic innervation patterns and variable conduction delays, such that spike time patterns corresponding to spatially and temporally localised activations of the postsynaptic neurons are more likely to make the post synaptic neurons spike. 207

6.2 Compostionality explanatory example. Figure illustrates explanatory example. Left: Illustration of the horizontal bar, vertical bar and T is 5 locations. Right: Illustration of spike time continua for the horizontal and vertical bar in the 5 locations (bottom) and the inherited spike time continua for the T in the 5 locations. 212

7.1 The Vogels-Abbott benchmark comparison across simulators. All single threaded or single GPU. 226

7.2 Comparing the multithreaded CPU performance of Auryn against Spike on a single GPU. 227

7.3 Comparing the simulation times of the Spike and Auryn simulators as the Vogels-Abbott Benchmark's network size is scaled. Network size is given by multiplication of the number of neurons and the number of synapses. The x-axis presents the network scale relative to the original benchmark network. On this scale, a value of 1 indicates a network of 4000 neurons, and a value of 100 indicates a network of 400,000 neurons. The smallest network is the size of the previous benchmark with 4000 neurons and 320,000 synaptic connections. The largest network, is 256 time larger with 1,024,000 neurons with 81,920,000 synaptic connections. The y-axis presents the time taken to simulate these networks relative to the time taken by the simulators to simulate a network of scale 1 (a network of 4000 neurons). The gray dashed line indicates linear scaling. 228
7.4 Brunel benchmark comparison across simulators with (black) and without (red) plasticity. 229
7.5 Comparing multithreaded CPU performance with Auryn against Spike with the Brunel non-plastic benchmark. 230
7.6 Heatmaps showing reliable responses by neuron (x), stimulus (y) and percentage of training complete (rows). White represents that a neuron spikes for all test trials of the stimulus; black otherwise. Parameter combination: learning rate: 0.0001, alpha decay: 0.15. 235
7.7 Number of spiking test trials (black = 0, white = all) by neuron (x), stimulus (y) and percentage of training complete (rows). Parameter combination: learning rate: 0.0001, alpha decay: 0.15. 236
7.8 Reliable responses over continuous ranges of stimulus transformations. Histograms showing how chunk lengths change as a function of training time. A chunk is defined as a range of continuous stimulus transforms over which a neuron responds reliably (i.e. on all stimulus trials). Parameter combination: learning rate: 0.0001, alpha decay: 0.15. . . 237
7.9 The effect of stimulus location on the number of neurons responding to both stimuli in a pair of stimuli. Each entry shows the number of neurons responding to both of the stimuli. Black illustrates zero neurons responding to both stimuli, white illustrates that all neurons respond to the pair of stimuli. The figure shows the evolution of the matrix over the course of training. Parameter combination: learning rate: 0.0001, alpha decay: 0.15. 238

7.10 Illustration of how a single neuron learns to respond reliably to the stimulus over a range of transformations. The figure shows how the values of the synaptic weights change over the course of training and the mean spike times that the neuron respond reliably to the stimulus in each location. Firstly the 1D input space is represented by the x-axis (input neurons 0-20,000). For each afferent synapse of the output neuron, there is a circle placed with a y-coordinate equal to the mean (expected) time that the synapse delivers spikes to the postsynaptic neuron. The colour of these circles illustrates the strength of the synapse (navy blue equals 0, yellow equals maximum weight). Red dots are placed when the stimulus responds reliably to a stimulus. The x-location of a red dot represents the central input neuron of the stimulus, and the y-location equals the mean spike time that the neuron responds to the stimulus at. The figure illustrates how the synaptic weights and spike times of the neuron change over the course of training. Parameter combination 'medium forward': learning rate: 0.0001, alpha decay: 0.15. 239
7.11 Figure the same as 7.10 but for the 'slow fizzle' parameter combination: learning rate: 0.0001, alpha decay: 0.25 and a different example neuron. 240
7.12 Figure the same as 7.10 but for the 'fast forward' parameter combination: learning rate: 0.001, alpha decay: 0.05 and a different example neuron. 240
7.13 Figure the same as 7.10 but for the 'medium forward' parameter combination: learning rate: 0.0001, alpha decay: 0.15 and a different example neuron. 241
7.14 Same as 7.10 but instead the colour of the afferent synapses illustrates the weight change over a single training phase (red positive, blue negative). 243
7.15 (a) PSTHs of three parameter combinations at the end of training. (b) Histogram of mean spike counts at the end of training for three parameter combinations. (c) Polynomials fit to chunks at the end of training for three parameter combinations. (d) Histogram of synaptic weights at the end of training for three parameter combinations. . . . 245

7.16 Figure shows the effect of the parameters on (a) mean reliable spike times, (b) mean reliable chunk length, (c) mean synaptic weight and (d) 1st lag autocorrelation r-value. The 1st column shows how these values evolve over training. The 2nd column shows how the standard deviation of these values evolve over training. The 3rd column shows the final values with 3*standard deviation confidence intervals at the end of training. 247
7.17 Figure shows the effect of the parameters on (a) polynomial residual, (b) 1st-3rd polynomial exponents. The 1st column shows how these values evolve over training. The 2nd column shows how the standard deviation of these values evolve over training. The 3rd column shows the final values with 3*standard deviation confidence intervals at the end of training. 249

List of Videos

The following supplementary videos were uploaded with the thesis and are also available on Dropbox at: shorturl.at/quACY

Chapter 3

- **Videos 1 & 2**: First spike response distributions and results of time series analysis for all Stage 2 clusters (Vid. 1) and correlated unclustered response distributions (Vid. 2). Description refers to Stage 2 clusters but same applies to correlated unclustered response distributions (but without clustering steps).

  Top left: Sampled first spike response distribution for two neurons (for all pairs which produced a Stage 2 cluster). Point colours illustrate DBSCAN clusters. Black points are outliers. Solid flat ellipses illustrates 4 boundaries of Stage 1 clusters. Angled ellipses illustrate 4 boundary of Stage 2 clusters. Yellow ellipse illustrates 4 ellipses under the assumption of independence and normality. For criteria fulfilling Stage 2 clusters, purple circle illustrates the 4 boundary of the factor-conditioned response distribution predicted by factor analysis for a single factor value (i.e. 0). For criteria fulfilling Stage 2 clusters, the purple line shows the mean points predicted for the factor values [-4, 4] (i.e. with variance equal to 0).

  Right: This subfigure analyses the spike times of the corresponding Stage 2 clusters. Columns 2 and 3 appear only for differenced Stage 2 clusters. Rows 1 and 2 corresponds with the example in Fig. 6a (see for further details). The columns correspond to (1) the original Stage 2 cluster spike times, (2) the cluster after non-stationary neuron cluster spike times were differenced, and (3) the result of applying AR and/or MA models to the differenced cluster. Row 1 shows the cluster spike times of both neurons. Row 2 shows the corresponding cluster, with the point colour illustrating the trial index (from dark to light).

  If the clusters fulfil the criteria after differencing and/or ARIMA modelling, the purple 4 factor analysis ellipse and line are shown in the corresponding row 2

plot. Rows 3 and 4 present the autocorrelation and partial autocorrelation plots respectively with 95% confidence intervals. Row 5 shows the QQ-plots for each neuron time series.

Chapter 4

- **Videos 3 & 4**: All p-value (Vid. 3) and r-value (Vid. 4) correlation matrices for an example stimulus condition (1Hz inter trial interval). Each square has either a green or orange border. A green border signifies that the correlation value is for a pair of neurons from a single cluster in the pool. An orange border signifies that the correlation was tested for stationary spike times, obtained from other stationary correlated clusters for the same neuron pair. Left: significant correlations ($p < 0.05$) have a green background, with the shade of green corresponding to the p-value (see key). Non-significant correlation ($p >= 0.05$) have a black background. Right: the background colour of each square represents the r-value (blue for a negative correlation and red for positive; see key).

- **Videos 5 & 6**: State predictions by trial index before and after adjustment. Single trial states by trial index before (Vid. 5) and after (Vid. 6) adjustment procedure for 1Hz stimulus condition. Each coloured dot represents the single trial state predicted for a single stationary correlated cluster. Black dots represent the means of the states predicted from stationary correlated clusters on single trials. Black line traces through the mean values for visualisation purposes. Red, green, blue cycle is purely for visualisation purposes and has no meaning.

- **Video 7**: Original cluster single trial state predictions vs adjusted cluster single trial state predictions for each stimulus condition with at least 2 clusters.

- **Videos 8 & 9**: Time-warp visualisations for all stimulus conditions with at least 2 clusters. Vid. 8 shows the multi-neuron first spike pattern for each trial of a single stimulus condition. The colour of each multi-neuron single spike pattern represents its index when the trials are ordered by mean predicted state. Red represents trials with high predicted single trial state, and blue low. It can be seen that the spike times of the trials with high (red) single trial state spike earlier than the spike time with low (blue) single trial state. Vid. 9 shows the same plots after smoothing.

Chapter 5

- **Videos 10 & 11**: In silico spike rasters and max normalised PSTH profiles for all parameter combinations. Spiking activity is shown for the 6 comparison neuron groups for a 25ms (Vid. 10) and 100ms (Vid. 11) periods following cortical activity onset. Lines show values of max normalised PSTH profiles for each neuron group (1ms bins). The spike trains of all neurons on a single trial are shown for the 6 neuron groups. Spike trains from all 4 trials are used to calculate the max normalised PSTH profiles.

Chapter 1

Introduction

The thesis aimed to address the classic neuroscientific problems of feature-binding, neural coding and representation learning. The thesis explored a diverse set of approaches, including analysis of sensory-evoked in vivo cortical spiking activity, large scale biophysical network modelling of cortical processing, and phenomenological modelling of representation learning through spike-time-dependent plasticity. Towards the latter, a spiking neural network simulator was also developed. Together, these approaches engendered novel findings and hypotheses regarding the nature of cortical processing.

Inspired by recent findings, the thesis studies multi-neuron patterns of single spikes (i.e. patterns of spikes in which each neuron spikes only once) as a possible fundamental form of cortical information encoding. A summary of the thesis chapters is given in Section 1.4 at the end of this chapter. The chapter begins, however, by introducing the classical problems of neural coding and feature binding.

1.1 Motivating research questions

> *'If these records give a true measure of the activity in the sensory nerve fibres it is clear that they transmit their messages to the central nervous system in a very simple way. The message consists merely of a series of brief impulses.... In any one fibre the waves are all of the same form.... In fact, the sensory messages are scarcely more complex than a succession of dots in Morse Code.'*
>
> Edgar D. Adrian, 1932

Since Adrian and Zotterman's discovery of discrete action potentials (i.e. spikes) in single sensory nerve fibres [5], the quantity and variety of neuroscience data has increased immeasurably. Despite an abundance of ever larger and more advanced datasets, the neuroscience community continues to yearn for fundamental principles

of information encoding. Perhaps most strikingly, neuroscience still lacks a full understanding of the simplest and most atomic element of the brain's action potential 'Morse Code'.

Specifically, in response to repeated presentations of a single stimulus, variability in the timing of the first action potentials of cortical neurons remains to be fully characterised and explained. This leaves a complete understanding of this most 'atomic' correspondence between extracranial reality and neural activity undetermined at the millisecond scale. It follows that an exact characterisation of the temporal properties of more complex and composite sensory representations, which facilitate action and perception, is incomplete.

To someone outside the field of neuroscience, perhaps nothing would seem more absurd than the continuing mystery of the most atomic form information representation. This, for a field that strives to explain our conscious experience, inspire the creation of artificial general intelligence and produce treatments for neurological diseases. With respect to the 'hard-problem' of experience, would determination of a 'function' mapping neural activity to phenomenological experience not benefit from an exact description of the function's 'inputs' and 'outputs'? For artificial general intelligence, is the implementation of evolution's own 4.5 billion-year-old solution not a likely source of insight into its integrated solution of algorithms and computations? Are the precise temporal specifics of information flow not likely to affect the homeostatic mechanisms of the brain in health and disease?

It is widely accepted that understanding the variability and encoding role of spike times is likely to require consideration of spike times with respect to other neurons or the concerted whole of a local population of neurons. For example, shared response variability across neurons can provide clues towards shared sources of variability [13, 249, 162], which may also affect encoding and decoding.

Furthermore, the ability of a downstream neuron to decode information from the *timing* of spikes requires reference to some other signal, such as the spike times of other neurons. Variability in the spike times of single neurons over repetitions of a single stimulus condition, then, cannot imply that the precise spike times of single neurons are not important to encoding, plasticity or homeostasis. Rather, spike timings should be interpreted with respect to other internal events and processes, including the spike times of other neurons. The nature of such spike time coding, however, remains to be determined.

Beyond the neural representation of single stimuli, the thesis considers the representation of sensory information more broadly. Such consideration offers a constrain-

ing source of requirements for the determination of a sensory neural code. Specifically, the thesis aimed to address the problems of neural coding and feature-binding, which collectively consider the questions of encoding, decoding, sensory perception and sensory-guided action. This chapter describes the neural coding and feature-binding problems; terms which are often used differently by different authors, or are used to describe specific sub-problems.

Classically, the study of neural coding and feature-binding has, in part, been partly motivated by the goal of understanding the neural representation underlying visual experience [170, 272, 171, 259]. The complexity of full visual experience lies beyond the scope of the thesis, however. The thesis, rather, takes a coding-centric perspective, focusing on fundamental aspects of neural representation, which could underlie sensory-guided action and sensory experience.

The thesis particularly considers processes of neural coding and feature-binding during the initial feedforward propagation of activity from sensory receptors to cortical representations. The initial feedforward propagation of activity may offer the simplest and most controlled opportunity to study principles of neural coding. It is also important to characterise the temporal specifics of neural representations and processing in the initial feedforward propagation of activity, as such representations can be viewed as the 'initial condition' of more complex recurrent processing.

The formulation of the neural coding and feature-binding problems motivates a series of scientific questions and requirements of a sensory neural code. These are formalised as a set of explicit research questions in Section 1.2. Section 1.4 then gives an overview of the progress made by the thesis with respect to these research questions. Candidate neural codes and theories of feature-binding are then considered with respect to the described requirements in the literature review chapter (Chapter 2).

1.1.1 The neural coding problem

The neural coding problem addresses the question of how information is encoded by, and decoded from, patterns of action potentials (i.e. spikes) in the nervous system [93, 140, 94]. Whether information is encoded by the number of spikes that neurons emit (i.e. spike count encoding) or by the precise millisecond timing of neuronal spikes (i.e. spike time encoding), remains the subject of considerable debate and study [33, 233, 306]. This section sets out questions that explication of a candidate neural code must address.

The thesis studies on responses following sensory stimulation, which can be analysed to determine principles of neural coding. Precise and reliable temporal spike patterns have been observed over repeated trials of single stimuli [281, 139, 98, 265, 226, 232] and can improve coding (over spike counts) in primary sensory cortices [211, 306]. Nevertheless, a large proportion of spike time variability remains unexplained [270, 251], leaving the definitive form of spike time encoding undetermined.

With the definitive form of spike time encoding undetermined, population spike count (i.e. mean firing rate) encoding is commonly assumed, in which single trial spike count variability averages out over a population [162, 121, 43, 266]. Under this assumption, analyses which smooth spiking data (with a Gaussian filter of 100ms) or which are based on binned spike counts over discrete time windows (usually 2-20ms) are applied. Such analyses, based on binned spike counts, can be used to study the task information or stimulus information encoded by the recorded population [162, 121, 43, 266]. Dimensionality reduction techniques may also be applied to binned spike count representations to reduce the population response space to lower dimensions or study the effect of latent factors on single trials [143, 91, 209]. Such approaches are useful for understanding the dynamics and functional role of a population.

Alternatively, spike time encoding offers the possibility of a faster, more efficient and higher capacity coding form, that would utilise the known temporal sensitivity of neural integration [92, 29, 30] and plasticity [176]. Analysis of continuous spike times (rather than discritised spike counts) could potentially offer improved decoding and inference of single trial dynamics.

1.1.1.1 Stimulus-response distributions, non-stationarity and spike train metrics

It is well known that neurons in sensory cortices exhibit trial-to-trial spike count and spike time variability in response to repeated presentations of single sensory stimuli (e.g. [211, 231]). This implies that neuronal responses to repeated presentations of a single sensory stimulus are drawn from some underlying *'stimulus-response distribution'*. Characterisation of such stimulus-response distributions offers the opportunity to determine reliable response properties that may be relevant to neural processes such as stimulus encoding.

A number of processes are known to affect neuronal responses over time. For example, neuronal stimulus-response adaptation describes how spike counts decrease

and spike time latencies increase for a subset of neurons over rapidly repeated presentations of a stimulus [7, 46, 231, 290], mainly due to short term synaptic depression (i.e. neurotransmitter depletion) [46]. This implies that stimulus-response distributions, for single neurons or groups of neurons, can be *'non-stationary'* (i.e. change over time). Such non-stationarity implies that, without the use of time series analysis or strict experimental control, responses from all trials of a stimulus cannot necessarily be used to precisely characterise the underlying stimulus-response distribution at a moment in time.

For the purposes of understanding, underlying stimulus-response distributions are firstly considered for groups of N neurons under the simplified assumption of stationarity. In the most rudimentary form, the spike train of a neuron i over T milliseconds can be represented as a binary vector $\boldsymbol{v}_i$ of length T, where: $\boldsymbol{v}_i(t) = 1$ if neuron i spiked between *t-1* and t milliseconds (for $t = 1, ..., T$), and 0 otherwise. Concatenating vectors $\boldsymbol{v}_i$ for neurons $i = 1, ..., N$ gives a representation of a stimulus-response sample in $N * T$ dimensional binary space. Changing the temporal resolution (i.e. the size of the time bin, B) changes the dimensionality of binary space to $(N * T)/B$. This demonstrates the ultra-high dimensionality of neuronal responses in this form, especially at high temporal resolution (i.e. for bin sizes smaller than a millisecond).

In this form and at high temporal resolution, very high numbers of samples would be required to estimate the underlying response distribution. Collecting high numbers of response samples per stimulus is made challenging by the limited time of single experimental sessions, in which it is also desirable to characterise responses to multiple stimulus conditions. To characterise the stimulus-response distribution in a form similar to this, spikes are binned with bin sizes on the order of 2-20ms (typically 10ms; e.g. [144]). As bin sizes are increased, however, temporal information is lost, and the code transitions to what is classically considered to be a spike count code.

Fortunately, characterisation of an underlying stimulus response-distribution can be simplified by projecting the sample stimulus-responses to lower-dimensional spaces using spike train metrics (Fig. 1.1a). Examples of spike train metrics include spike counts, first spike times, and the relative difference between the spike times of two neurons. The thesis aims to geometrically characterises the nature of first spike stimulus response distributions, as this constrains the form in which spike times may encode information.

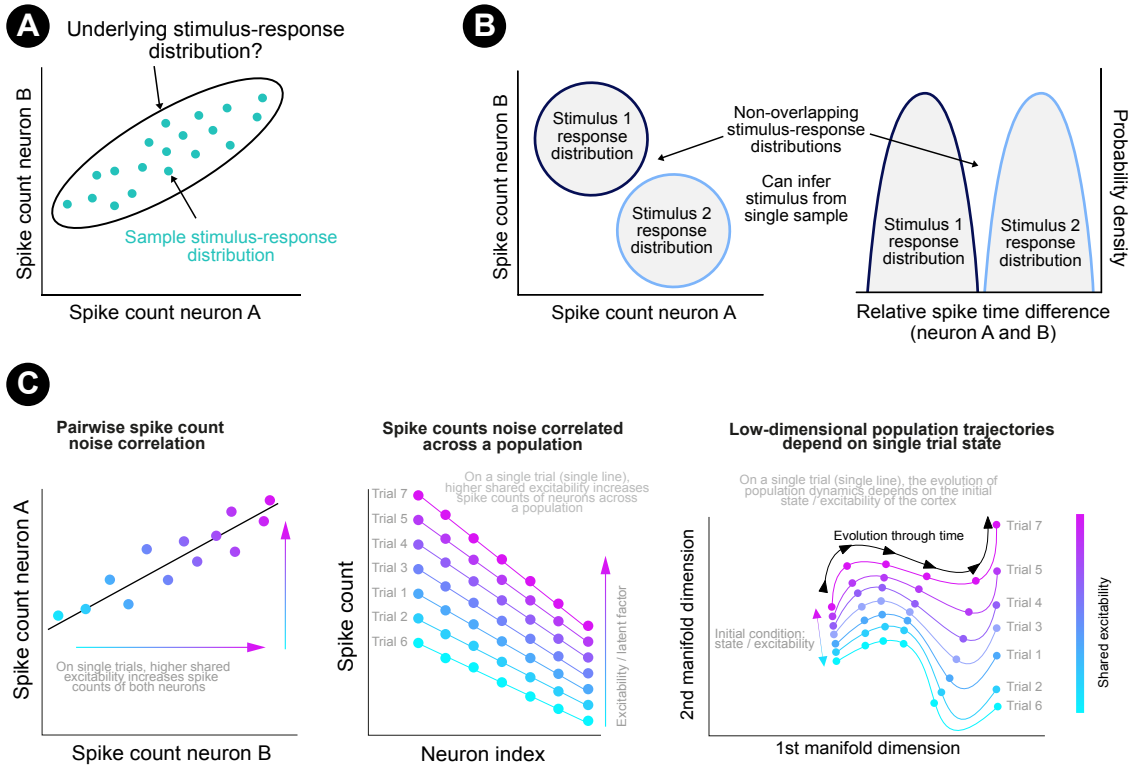


Figure 1.1: (a) Illustration of spike count stimulus-response distribution space for two neurons. Red points illustrate sample stimulus-response distribution. Blue ellipse illustrates the boundary of the true underlying stimulus-response distribution. (b) Illustration of two non-overlapping stimulus-response distributions for (left) spike counts and (right) relative spike time differences. (c) Illustration of shared spike count variability over repeated trials of a condition for a neuron pair (left), for a population (center) and exhibited as alternative single trial trajectories through low-dimensional space (right).

1.1.1.2 Discrimination and neuronal sensitivity

If two sensory stimuli elicit different behavioural responses, downstream neurons must be able to respond differently to samples from the underlying response distribution of either stimulus. This is possible if the response distributions are non-overlapping in dimensions corresponding to spike train metrics that downstream neurons are sensitive to (Fig. 1.1b). If the response distributions are highly non-overlapping for two stimuli which differ by a small value of a stimulus parameter, then the encoding is highly precise, as small changes in the stimulus parameter could (theoretically) be decoded by downstream neurons.

It may not be immediately obvious whether downstream neurons are sensitive to changes in the value of the spike train metric corresponding to a candidate neural

code. For example, decoding of spike times by downstream neurons requires reference to internal events such as the spiking activity of other neurons. For this reason, spike times relative to stimulus onset are considered unlikely as a neural coding candidate. Instead, researchers consider relative differences between the spike times of different neurons, which neurons have been shown to be sensitive to [92, 29, 30]. Whilst the thesis specifically considers relative spike time codes, analysis of first spike time response distributions is still useful for inferring the presence of underlying processes affecting the spike times of multiple neurons.

1.1.1.3 Tuning curves, decoding algorithms and information analysis

Tuning curves plot how the value or mean (over trials) of a spike train metric for a single neuron varies with changes in a stimulus property. Tuning curves have therefore classically been used to study neural encoding of property information. Study of tuning curves does not consider how the coordinated response of the population may encode information collectively.

Classically, decoding algorithms and information analysis are the two main ways in which neural activity is analysed to determine the efficacy of a neural code to encode information in a form that can be discriminated by downstream neurons (see [221, 155] for review). Moreover, these approaches allow for such a quantification for a population of neurons, rather than just single neurons.

Decoding algorithms try to predict the identity of a stimulus on a single trial from the values of spike train metrics. It is assumed that if a decoding algorithm is able to predict the identities of stimuli above chance then the recorded population encodes these stimuli in a form that would allow downstream neurons to discriminate between the stimuli. Decoding accuracy can be calculated for different neural populations to see which population is better at encoding the stimulus information. Decoding accuracy can also be compared for different spike train metrics.

Examples of decoding algorithms include Bayesian decoders, Nearest-neighbour decoders, Fisher linear discriminant decoders, support vector machines and artificial neural networks.

Mutual information and Fisher information are the two main types of information analysis used to study how well populations and spike train metrics encode stimulus information. Mutual information quantifies how much information is encoded in neural responses about each stimulus (over a set of stimuli). Particularly, mutual information calculates the reduction in uncertainty about which stimulus is present,

given neural responses. The lower the degree of overlap between the response distributions of different stimuli, the higher mutual information.

Fisher information offers an alternative to tuning curves for studying how changes in a stimulus parameter affect neural responses. Particularly, Fisher information 'captures the precision with which a stimulus property is encoded in the vicinity of a particular value of the property' [155]. Fischer information offers a quantification of how 'non-overlapping' stimulus response distributions are around a value of a stimulus property.

Whilst, related, these approaches differ from geometrical characterisation of the sample response distribution for single stimuli. Characterisation of sample response distributions may help understand how information is represented by spiking activity.

1.1.1.4 Neuronal stimulus-response adaptation and other non-stationary processes

As discussed, underlying stimulus-response distributions may be non-stationary due to processes such as neuronal stimulus-response adaptation. Neuronal stimulus-response adaptation is characterised by a decrease in spike counts and an increase in spike time latencies over rapidly repeated presentations of a stimulus [7, 46, 231, 290]. Representational drift has also been observed in sensory cortices, showing that underlying stimulus-response distributions can change over the course of minutes, sessions or days [55]. For example, a recent study in mouse visual cortex showed that firing rates of single neurons change over the course of days but that tuning changes to a lesser extent [55]. Importantly, however, population encoding remained stable over time.

In natural sciences there are many examples in multi-scale systems where non-stationarity is explained by low-dimensional latent processes [183]. Due to the multi-scale character of neural processes, it is possible that response non-stationarity in the brain is also affected by the interaction of processes across scales.

Ideally, geometric characterisation of stimulus-response distributions is best suited to experimental conditions in which stimulus-response adaptation is minimised, for example, through the use of long intervals between stimulus presentations. To minimise the effects of adaptation and non-stationary processes, analyses are often applied to low or reduced numbers of trials per unique stimulus. For example, the first 20 stimulus presentations (when the majority of adaptation occurs) might be excluded from analysis [231].

The longer-timescale trend-like effect of neural adaptation (and possibly other non-stationary processes) lends itself to the application of time series analysis to test for and remove trends in neuronal responses over time. This in turn has the potential to allow more samples to be used for estimation of a stimulus-response distribution at a single moment in time. Ideally, time series analysis would determine that no form of non-stationarity was present, such that time series modelling was not required to remove non-stationary trends.

1.1.1.5 Noise correlations and cortical state

As single cortical neurons can generate highly reliable spike trains to direct stimulation in vitro [168], in vivo response variability largely reflects cortical network dynamics [39] and synaptic interactions [31, 46]. Analysis of response variability shared between neurons over repeated trials of a single stimulus ("noise correlations") allows the effect of underlying variability sources to be investigated.

Spike count noise correlations and modulation by latent factors

Shared spike count variability may be studied for pairs of neurons and at the population-level (Fig. 1.1c, left and center [162]). The latter includes analysis of latent factors and trial-to-trial variability in the low-dimensional trajectory of population activity (based on binned spike counts; Fig. 1.1c, right; [143, 209]). For stimulus-evoked activity, noise correlations correspond with correlative structure in a multi-neuron stimulus-response distribution (where each sample represents the response on a single trial). Noise correlations during spontaneous activity, caused by shared fluctuations in excitability, may be studied by applying dimensionality reduction to vectors of binned spike counts [233].

In the visual cortex, shared spike count variability is modulated by shared trial-to-trial variability in the excitability of neurons across a population. Particularly, neurons across the population respond with either higher or lower spike counts depending on the population excitability level on a single trial (Fig. 1.1c, center). Such shared variability is explained by an additive interaction of stimulus-evoked activity with spontaneous background activity [13, 249, 162] and an additional multiplicative interaction [162]. That is, the level of spontaneous background activity and stimulus-evoked activity summate and are modulated by an additional unobserved latent process that varies responses across the population multiplicatively from trial-to-trial. These results show that a population's state of excitability at the point of stimulus presentation affects the representation of a stimulus on a single trial.

[51] explores the efficacy of different features of spontaneous activity on predicting stimulus-evoked activity in the auditory cortex under urethane anaesthesia. More recently, [199] showed that in in vitro recordings of hippocampal slices, responses exhibited state dependence on the time between the last spontaneous burst and the stimulus presentation and that the dependence could be described with a linear model. Moreover, the amount of information encoded by the population showed a state-dependence and that a small number of neurons carrying the most stimulus information also contributed most to the state-dependent information gain.

Analysis of population dynamics in the motor cortex following dimensionality reduction also supports the dependence of neural processing on a population's state at the start of a trial (Fig. 1.1c, right; [143, 209]).

Sources of noise correlations are reviewed in [63]. Subthreshold membrane potential correlations are a prime candidate to underlie noise correlations caused by trial-to-trial changes in shared excitability as they are correlated across areas [220] and affect action potential initiation [219]. The coding implications of spike count noise correlations have been studied theoretically [14, 190] and are discussed in the literature review (Chapter 2).

Neuronal stimulus-response adaptation [7, 46, 231, 290], which describes how spike counts decrease and spike latencies increase for a subset of neurons over rapidly repeated presentations of a stimulus, would give rise to noise correlations in the stimulus response distribution as shown in Fig. 1.2 (left). The thesis explores noise correlations corresponding to non-adaptative trial-to-trial changes in excitability Fig. 1.2 (right). Towards this, characterisation of stimulus-response distributions controls for neuronal stimulus-response adaptation, in order to maximise response distribution stationarity.

Movement is a third source of variability and noise correlations in primary sensory cortices [267]. To simplify investigation of neural coding, however, the thesis mainly considers neural coding when movement is controlled for.

Spike time noise correlations

The thesis considers how spike time representations may be modulated by the same trial-to-trial changes in shared excitability which modulate spike count noise correlations. Particularly, the combination of neuronal stimulus-response adaptation and trial-to-trial changes in shared excitability may obscure spike time encoding and give rise to the apparent complex nature of spike time variability. Such investigation must test and control for neural stimulus-response adaptation, which may also give rise to

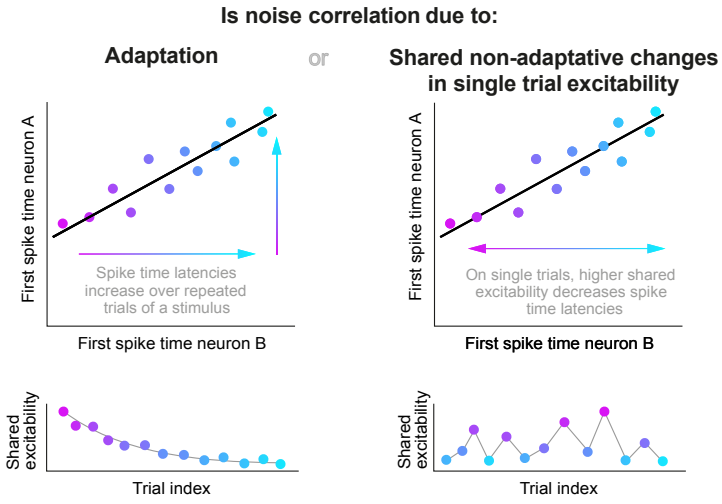


Figure 1.2: Possible latent processes underlying the noise correlations include adaptative (left) and/or non-adaptative (right) changes in the shared excitability-level of the two neurons from trial-to-trial.

noise correlations. Particularly, time series analysis is used to test and control for neural stimulus-response adaptation. Further motivation for the study of multi-neuron single spike patterns is given below.

Single trial variability in population dynamics is thought to reflect differences in population excitability or state at the start of a trial [143]. The thesis proposes that such variability in population dynamics may also be reflected in spike time variability. As a step towards testing this hypothesis, the nature of shared spike-timing variability is characterised.

Continuous cortical state

Throughout the thesis, trial-to-trial changes in shared excitability are referred to as changes in 'continuous cortical state'. This does not refer to major state changes, such as changes between wake and sleep [109], but to trial-to-trial (i.e. between trial) changes in shared excitability. The dimensionality of continuous cortical state changes is discussed in a following section.

1.1.1.6 Decoding

Fixed-state decoding

Classically, the state of downstream decoding neurons is assumed to be fixed. That is, the responses of downstream neurons are assumed to be unaffected by adaptation and changes in excitability. Throughout the thesis, the term 'fixed-state decoders' is used

to refer to this assumption. To the author's knowledge, all previous theoretical and experimental work in neural coding assumes that neurons in the brain are fixed-state decoders.

State-dependent-coding

If the number of latent processes modulating neural responses across a population is low, then it is likely that decoding neurons (within the population and in nearby regions) are also modulated by the same latent processes (i.e. changes in cortical excitability). If this were the case, it could be possible that decoding neurons could decode/respond in a state-dependent manner. [210] suggested that 'downstream areas may extract variables indicating the current state from network activity' ... 'and then use state-dependent decoders to interpret population activity'. This captures the notion of state-dependent coding but does not suggest that shared changes in cortical excitability between areas could enable decoding neurons to respond in a state-dependent manner. As discussed in Section 1.1.1.5, [199] demonstrated an increase in decodable stimulus information under state-dependent decoding for activity from in vitro hippocampal brain slices.

Such state-dependent-coding could offer an increased coding capacity over fixed-state-decoders, as follows. Consider the case in which the responses of two neurons are modulated by fluctuations in a single-dimensional continuous cortical state (i.e. excitability). When the fluctuating state has a value on a single trial, a response comes from a *state-conditioned* stimulus-response distribution, which would occupy a smaller region of response distribution space than the region sampled over multiple trials (Fig. 1.4c, right). Under state-dependent decoding, smaller state-dependent response distributions would offer an improved coding capacity over the classical fixed-state interpretation, as more unique stimuli could be reliably encoded.

1.1.1.7 Summary

The neural coding problem asks how information is encoded by, and decoded from patterns of action potentials, and can be explored through consideration of stimulus-response distributions. Geometrical characterisation of stimulus-response distributions is made non-trivial by non-stationary processes, which change the underlying distribution over time. Through analysis of stimulus-response distributions, the thesis explores how known sources of spike-count variability may interact to affect and obscure spike-time encoding, and to consider the potential implications for encoding capacity.

1.1.2 The feature-binding problem

In sensory processing, the term *'feature'* defines any pattern in the sensory input. The term can refer to sensory primitives (i.e. visual edges and colours) or compositions of sensory primitives (i.e. visual contours and shapes). The term can also refer to more abstract features, such as properties (i.e. orientation), categories (i.e. cat) or object identities (e.g. 'Jerry' the cat).

Any sensory scene is composed of some complex configuration of features at multiple levels of scale. Sensory-guided action and sensory perception relies on a useful representation of such features and their relationships to one another. The classic feature-binding problem (FBP) aims to understand the nature of such a representation and its underlying neural mechanisms [241, 177, 171].

1.1.2.1 Neuronal and modular specialisation

The FBP arose from physiological evidence of neuronal and modular specialisation in the processing of basic visual features. In the primary visual cortex, neurons were found to be organised spatially depending on their selectivity to different features such as colour, stereopsis, movement, and orientation [130, 129, 164]. Early study of the anatomical and functional organisation of the visual cortex demonstrated that regions of the visual cortex were specialised; for example, responding for motion but not colour, and that processing was organised hierarchically [50]. Hubel and Wiesel's classical finding of simple and complex cells [126, 128] additionally supported hierarchical processing.

From such observations of neuronal and modular specialisation, it was thought that a visual image is decomposed into independent representations, such that information regarding the relationships between stimulus features is lost. Under this assumption, the feature-binding problem initially asked how the relationships between features are encoded after such decomposition [241, 275, 177, 274].

A larger hierarchy of visual processing was later characterised for the ventral stream, in which neurons respond to features of increasing complexity with increasing invariance to transformation through successive stages; with highly position and orientation-sensitive units in V1 [194], and position-, size- and pose-invariant units in higher ventral areas [97, 137]. Neurons respond to increasingly complex bar and shape primitives in V2 [127], V4 [213] and posterior inferotemporal cortex (IT) [35], and to complex 3D shape primitives in IT [300, 132, 284]. Categorical object information is encoded in IT and prefrontal cortex [82].

The face patch system identifies six regions in the macaque temporal lobe [83] and three in the prefrontal cortex, with neurons representing: face features in posterior patches; facial feature combinations and spatial relations in middle patches; and identity of faces irrespective of view in anterior patches [83]. These examples demonstrate the highly distributed representation of visual information at multiple spatial scales along the ventral stream that must be associated into coherent object representations.

1.1.2.2 Correct association of feature-presence and feature-properties

The thesis considers how information regarding the presence of a feature is associated with information regarding the feature's properties. Here, 'feature-properties' refers to the lower-level features composing the feature or more global/holistic properties such as orientation, colour and facial expression. Such feature-properties are in fact defined by the composing low-level features but may also be explicitly encoded elsewhere. Encoding of associations between feature-presence and feature-properties has been a major focus of the FBP historically. Such an association is critical to sensory-guided behaviour.

Detection of a feature's *presence* and *properties* signals the *opportunity* for a set of possible action-types. It is the *properties* of the feature, however, that determine the motor specifics of undertaking a particular action-type. The 'combinatorial problem' highlights that a quasi-infinite number of neurons would be required for every possible feature-presence and feature-property combination to be encoded by an individual neuron [241, 287].

Hubel and Wiesel's discovery of complex cells [128] introduced the notion that single cells can encode feature-presence over a range of transformations. This term is similar to the notion of a 'grandmother cell', which describes a cell which responds reliably in the presence of a feature, such as your grandmother. Classically, the term 'transformation-invariant' was also used to describe such a cell but this term additionally introduces the notion that the response of the neuron should be the same (i.e. be invariant) to changes in the preferred feature.

The term 'feature-presence-encoding' is used herein to describe the notion that a neuron or cell assembly responds to the presence of a feature over a range of transformations. The term relaxes the constraint that neurons in the cell assembly should respond invariantly over transformations of it's preferred feature. The term also takes into account that different neurons within the cell assembly are likely to respond with different reliabilities to the preferred feature.

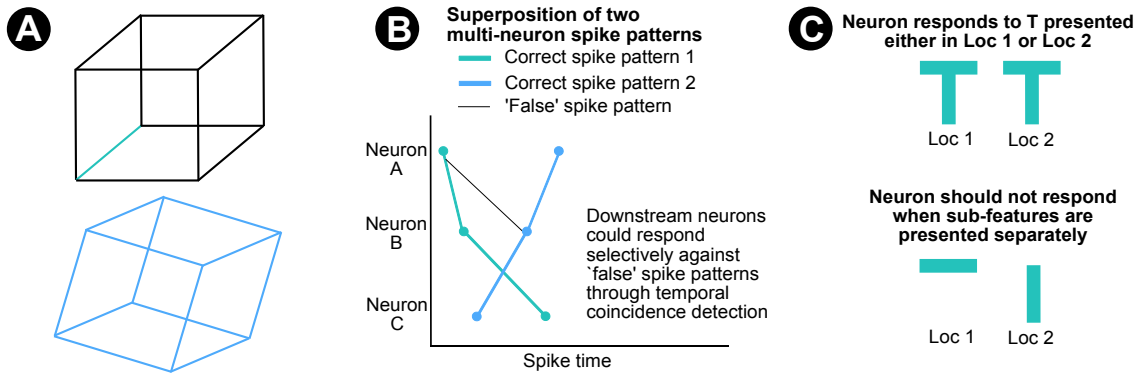


Figure 1.3: (a) Example of two alternative Necker cubes defined by their properties and constituent low-level features. The top cube has all black edges except for one blue edge, whilst the red cube has all red edges. Each cube has a different orientation. The Necker cube illusion: the perception of a single cube (e.g. the top cube) transiently changes between the perception of two alternative cubes. (b) For two neurons, illustration of two overlapping multi-neuron spike patterns (red and purple). Coincidence detection could play a part in robustness to 'false' spike patterns (black). That is downstream neurons could respond selectively to particular spatiotemporal patterns of spikes. This may be achieved through the spatial organisation of neurons, synaptic innervation patterns and variable conduction delays, such that spike time patterns corresponding to spatially and temporally localised activations of the postsynaptic neurons are more likely to make the post synaptic neurons spike. (c) Illustration of false combinations problem. Particularly, if a neuron or cell assembly represents the presence of a feature over a range of transformations, the false combinations problem asks how a neuron or cell assembly is robust against firing to incorrect combination of constituent low-level features.

Early theories of feature-binding proposed how associations might be encoded through *coordination* of spatially distanced representations encoding *feature-presence* and *feature-properties* [185, 275, 177], where feature properties could describe the current feature transform or constituent low-level features, for example. Correct association of feature-presence and feature-property information would be made complicated by multiple simultaneously presented features. This is summarised by the 'superposition catastrophe' problem, which asks how downstream neurons can interpret correct associations when feature representations of multiple objects are co-active [241, 287].

A joint representation of feature-presence and feature-properties additionally accords with our perceptive experience. For example, by switching our experience between two alternative perceptions, the Necker cube [197] confirms that entire features (specifically 'objects' in this case) are perceived beyond the perception of individual photoreceptor activations (Fig. 1.3a). That is, unlike, for example, a digital photo, our visual experience is not explained as a pixel grid of different colours. Instead,

our perceptive experience contains the perception of grouped features and objects. That is, the detection of feature-presence can be considered a requirement of a system underlying perceptive experience.

The Necker cube (or any feature) also demonstrates the obvious, but sometimes forgotten point, that feature-perception is defined not only by feature-presence but also by feature-properties (e.g. rotation, composing low-level features), as perceptions of alternative cubes are determined by their properties and composing low-level features. Fig. 1.3a also demonstrates that our perception includes the simultaneous perception of multiple objects and features, and an understanding of feature relationship at multiple spatial scales. For example, perception of the top cube in Fig. 1.3a includes the understanding that the left bottom edge is the left bottom edge of the top cube.

A joint encoding of feature-presence and feature-properties for multiple simultaneous features, then, can be thought of as a requirement of systems underlying both perceptive experience and sensory-guided action, supporting the likely case that the underlying neural systems are shared. This may seem 'by-definition' but is worthy of further consideration. Particularly, the shared requirements highlight a likely link between the content of experience and the neural processes supporting action. This suggests that research into neural processing can be guided and constrained by research into the content of experience and sensory-guided action (and vice versa). Indeed, it could be said that the FBP aims to determine the relationship between neurophysiology, perceptive experience and sensory-guided action.

Psycho-physical investigation allows the feature-binding content of perceptive experience and the role of attention to be characterised scientifically (e.g. [296, 62]). Consideration of such research was useful for its potential to guide, constrain and motivate investigation, although psycho-physical investigation is not the focus of the thesis. The literature review also includes discussion of early theories aiming to explain how spatially separated representations of feature-presence and feature-properties might be 'coordinated' to create unified object representations.

1.1.2.3 Joint feature-presence and feature-property encoding

The FBP is made simpler by evidence that a cell assembly's activity pattern can encode feature-property information [98, 265, 121, 43]. If activation of a cell assembly signals the presence of a feature, then feature presence and property information could be encoded jointly. Such joint encoding would avoid the need for post-hoc association of representations encoding the opportunity for a particular action-type (i.e. feature

presence and properties) and representations determining the action-specifics (i.e. feature properties). Instead, association of feature presence and property information could be associated during the process of parsing (i.e. the transformation of sensory stimulation to neural representations).

The definitive nature of such property encoding remains an open question. Such encoding was classically explored through analysis of neuronal tuning profiles, which are widely considered to be largely obfuscated by noise at the single cell level, leading to the view that property encoding is achieved at the population level.

IT face-specific neurons are examples of neurons which encode the presence of a feature (i.e. face) through their activation, whilst information regarding feature-properties is encoded through their activity pattern (trial-averaged firing rates [43]). Particularly, neuronal population modelling of neurons across different temporal lobe patches (ML/MF & AM) suggests that, neurons project incoming faces, as vectors, onto specific axes of high dimensional space. Object properties can also be decoded from trial-averaged multi-unit activity (MUA) responses in IT [121]. The possible nature of property encoding is discussed in Section 1.1.2.7.

1.1.2.4 Overlapping and simultaneously active representations

If property information is encoded by the activity of feature-presence-encoding cell assemblies, then an alternative superposition catastrophe problem arises when the feature-presence cell assemblies of two simultaneously presented features overlap with each other. Specifically, how are property encoding activity patterns affected when two features are presented, which both activate a joint subset of neurons? Several firing rate competition schemes have been demonstrated for object representations in the presence of additional objects [253, 305, 15] (see Chapter 2).

Spike times offer a potential solution due to their relatively instantaneous temporal extent. Particularly, as action potentials are short in duration, neuronal spike time patterns could, in theory, 'superimpose' and encode property information of multiple co-active representations quasi-simultaneously (Fig. 1.3b).

Spike coincidence detection by downstream receiving neurons could decrease the occurrence of downstream neurons responding to or 'interpreting' incorrect features/feature-transforms represented by different spatiotemporal patterns. Particularly, downstream neurons could respond selectively to particular spatiotemporal patterns of spikes. As neurons are sensitive to the temporal and spatial proximity of synaptic activations [92, 29, 30], selectivity to particular spatiotemporal patterns could be

achieved through the spatial organisation of neurons, patterns of synaptic innervation, and variable conduction delays. In this way, spike time patterns which provide spatially and temporally localised activations of a postsynaptic neuron would be more likely to make the post synaptic neurons spike compared with patterns of spikes from the same presynaptic neurons which do not deliver spikes to the postsynaptic neuron in close temporal proximity.

1.1.2.5 Parsing and robustness to false combinations

The FBP also concerns how information is correctly parsed into joint feature-presence and feature-property representations, including how neurons respond selectively to a subset of stimuli over ranges of transformations. Such parsing is constrained by neural architecture and the representation of intermediate compositional features and the representation of the final feature into which lower-level representations are to be parsed. Understanding of parsing, therefore requires characterisation of cortical representations.

The DPhil research was initially interested in how the brain learns to represent the presence of a feature over a range of transformations but robustly does not represent the presence of the same feature when *false combinations* of possible constituent low-level features are presented (Fig. 1.3c). This is closely related to the 'adversarial example' problem in machine learning [100], which considers how a machine learning network (usually convolutional neural networks) can be fooled into reporting the presence of an object category after the addition of imperceivable (to humans) noise to an image of a different object category. Put differently, the network reports the presence of an object despite the configuration of constituent low-level features not according with a true exemplar of that object category.

The brain's solution against such adversarial examples and the more general question of how the brain represents and parses information is therefore of interest to machine learning. This is not to say that the brain does not make mistakes in representing the visual world, indeed there is a large history of research studying where human perception breaks down; for example, the study of illusory conjunctions. Instead, it can be said that the types of image that fool the brain and machine learning network are quite different, suggesting that the underlying processing algorithms are also different. Characterising these differences and trying to bring machine learning algorithms in to accordance with the properties of human vision is the subject of continuing research (e.g. [62]).

1.1.2.6 Learning and semantic meaning

The FBP engenders the question of how feature-binding representations emerge through development and experience. Due to activity-dependent-plasticity, learned sensory representations depend on mechanisms of plasticity, the form of past sensory representations and the past history of sensory experience. The thesis considers that sensory representations (and perceptions) are, therefore, an implicit comparison to past experience, and that such comparison to past experience gives semantic meaning to sensory representations. Again, the representation of stimuli, which the thesis explores, constrains the process of experience-guided-learning and comparison to past experience.

1.1.2.7 Property encoding by (state-dependent?) continua and neural manifolds

The thesis asks how property information may be encoded along neural activity continua. That is, small (continuous) changes in feature-transform should correspond with small (continuous) changes in stimulus response-distribution space, such that stimulus-response distributions for separate transforms are regions on a continua in stimulus-response distribution space (Fig. 1.4a). Such continua should lie along dimensions corresponding to spike train metrics that downstream neurons are sensitive to.

The idea of continua is similar to that of a neural manifold (described as follows). Given the recorded activity of a large population of neurons, dimensionaliity reduction techniques are used to convert neural responses to a lower dimensional space, under the assumption that there is a high degree of noise (example reviews: [134, 131]). The neural manifold is the surface traced in the reduced response space, as stimulus parameters are varied.

[131] proposes alternative 'strong' and 'weak' principle of neural manifolds. Under the strong principle, it is thought that all task or stimulus information is encoded in the reduced low-dimensional space. Under this assumption, it is only the coordinated activity of the population that encodes important task or stimulus information with neural activity not captured by the manifold space considered as noise. The strong principle therefore assumes that the dimensionality of the encoding space is much lower than the number of neurons, implying a high degree of redundancy.

Under the weak principle of neural manifolds, it is assumed that neural manifolds are only tools for smoothing out complex high-dimensional responses rather than models which capture all of the important features of neuronal responses. The

weak principle assumes that such smoothing can enable some understanding of neural processing.

Most dimensionality reduction techniques are applied to binned spike counts. Information encoded at a higher temporal resolution could be lost by this process of binning. In relation to neural manifolds and continua, the thesis considers whether the spike timings of individual neurons are structured enough, such that each neuron could encode a separate dimension. If spike timing responses of individual neurons exhibited high independent temporal noise (as is often assumed), then only a signal encoded by spike counts or in low-dimensional (binned spike count) manifold would contain signal.

Alternatively, single trial temporal noise may be coordinated across the population, with respect to the single trial cortical state (i.e. excitability level of the cortex), as suggested in Section 1.1.1.5. If downstream decoding neurons could respond in a state-dependent manner (as described in Section 1.1.1.6), then state-dependent single trial representations would come from a precise area of response distribution space.

It may be expected that such precise state-dependent representations may vary smoothly with changes in stimulus parameters. Moreover, the state-dependent precision of single neurons could enable each neuron (or pairs of neurons for spike time difference encoding) to encode single stimulus parameters. Under such a regime, the dimensionality of the neural manifolds/continua would be closer to the number of neurons, making the encoding capacity of the population much higher than the non-state-dependent regime.

Characterising the effect of latent processes, such as trial-to-trial changes in cortical state, on the spike time representation of single stimuli, is a first step towards understanding the effect of latent processes on spike time continua.

1.1.2.8 The representation of novel stimuli along continua

If continua arose through experience such that some regions on the continua represented the stimulus-response distributions for transforms experienced during learning, the region representing a novel transform on the continua would implicitly encode the transform *relative* to previously experienced feature transforms (Fig. 1.4b). The process by which continua emerge through experience would then be of critical importance to understanding how semantic meaning emerges in the brain. Again, this is constrained by the nature of single stimulus representations, as described.

Continua encoding solves the 'combinatorial problem' of requiring a unique neuron to encode each of a quasi-infinite number of unique feature-presence-feature-property

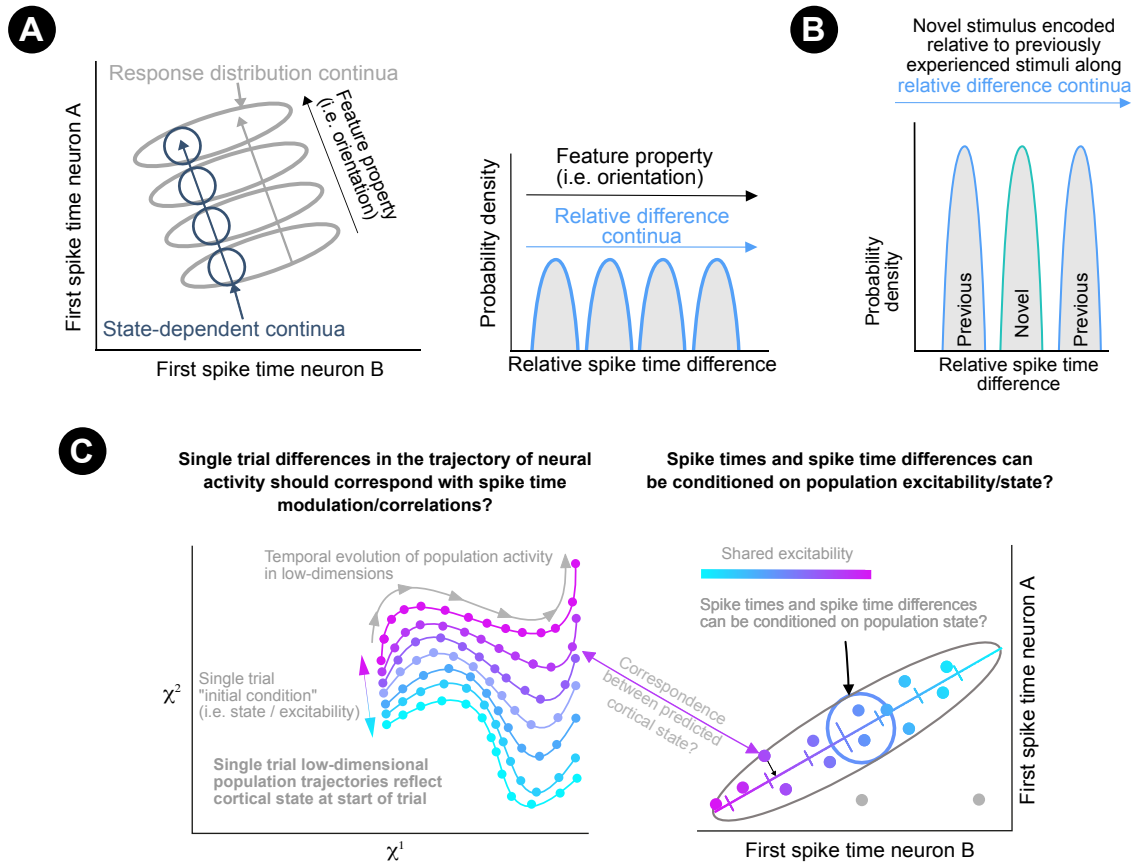


Figure 1.4: (a) Left: Illustration of continua of first spike time distributions, corresponding to first spike times (not relative differences) under fixed-state-decoding (blue) and state-dependent-decoding (green). Right: Illustration of continua of relative spike time difference distributions, which could correspond to fixed-state-decoding and state-dependent-decoding distributions. (b) Illustration of the representation of novel stimulus relative to previously experienced stimuli along relative difference continua. (c) Illustration of possible correspondence between single trial variability in low-dimensional population dynamics and single trial spike times for correlated clusters.

combinations. Consideration of how spike-timing continua-encoding could address other aspects of the FBP and sensory processing are considered in Chapter 6.

1.1.2.9 Compositionality of continua: property-dependent parsing and holographic projection of information

If property-encoding forms were *general* for representations of lower and higher-level features, such representations could be *compositional*. The word 'compositional' here firstly means that property information could be taken into account when parsing a group of low-level features into a high-level feature such that the high-level feature

representation is only instantiated if the properties of the low-level features are configured correctly with respect to one another. This would introduce robustness to false combinations of low-level features.

Secondly, 'compositional' here means that property information would be transferred from the low-level representation into the activity pattern of the representation of the high-level feature. Under a continua-scheme, continuous changes in the activity pattern of a low-level feature representation could translate to continuous changes in the activity pattern of the high-level representation. In this way, some property information could be maintained through multiple levels of a processing hierarchy. The encoding of property information along a continua could then be used in the robust parsing of further higher-level features or to inform the motor-specifics of action.

1.1.2.10 Generalisable action

The DPhil research initially took interest in visually-guided action. Continua encoding offers a natural means by which action can be produced for features in novel configurations. For example, consider a subset of neurons, which respond in the presence of a feature over a range of transformations (i.e. a feature-presence-encoding cell assembly). Also consider that the activity pattern of the cell assembly represents properties of the feature along a continua. Firstly, activation of the cell assembly signals the opportunity for a subset of action-types.

Secondly, consider that there exists a mapping from points on the continua representing previously experienced instantiations of the feature to the specific motor commands required to perform the action on these instantiations. Then, when the feature appears in a novel transformation it would instantiate an activity pattern in the feature-presence-encoding cell assembly that is some interpolation of activity patterns elicited for previous transformations of the feature. This representation would in turn instantiate an interpolation of motor commands for appropriate performance of the action upon the object. Again, the representation of feature transformations would determine the nature of the continua, which would in turn constrain such a mapping to motor commands.

Precise spike times have been demonstrated in the motor cortex [268] and muscle units (the muscle fibers innervated by a single motor neuron) [264] (see [262] for review), supporting the possibility that continua and generalisation are spike timing based.

1.1.2.11 One-shot or multi-shot learning

Finally, a solution to the FBP should also answer how the brain is able to represent the presence and properties of novel features, identities and categories over a range of transformations, after experiencing very few example transformations of that feature. In machine vision, learning to represent the presence of an identity or category over a range of transformations from a single or minimal number of training images is referred to as one-shot or multi-shot learning.

The human brain's solution to one-shot or multi-shot learning (with property encoding) is complex and involves the interaction of many processes, from synaptic mechanisms to patterns of visual saccades. Characterising the detailed nature of feature encoding is a necessary component to understanding how the brain achieves one- or multi-shot learning, as it is the representation underlying such a solution.

1.2 Summary of motivating research questions

For reference throughout the thesis, the discussed motivating research questions, are summarised as follows:

- General Research Question 1 (GRQ 1): How are short sensory stimuli represented atomically in single cortical areas?

  By 'atomic representation', the thesis refers to the minimal activity within a cortical area that critically supports a perceptual discrimination. Such a 'packet' of activity could be considered as the fundamental unit of representation used to transmit behaviourally relevant information through-out the brain.

- General Research Question 2 (GRQ 2): What is the nature of continua encoding for atomic sensory representations?

- General Research Question 3 (GRQ 3): How are short sensory stimuli parsed (i.e. transformed) into neural representations? And how does the process of parsing provide selectivity (including against incorrect combinations of low-level features).

- General Research Question 4 (GRQ 4): How do continua representations emerge through experience?

- General Research Question 5 (GRQ 5): How is feature-property continua encoding (for features at any scale) affected when multiple overlapping cell assemblies are simultaneously activated?
- General Research Question 6 (GRQ 6): How is one-shot or multi-shot learning facilitated by sensory representations in the initial forward pass from sensory stimulation?

It is important to stress that these were *motivating* research questions, and not all are directly addressed. In fact, the thesis directly address GRQ 1 (and partly GRQ 4) and, based on this work, sets out hypotheses for the remaining GRQs.

1.3 Multi-neuron first and single spike patterns

The thesis considers multi-neuron single spike patterns with respect to the GRQs and as a possible atomic and universal coding form. Recent research suggests that multi-neuron first spike patterns are a fundamental form of cortical representation [211, 232, 230], which may generalise to multi-neuron single spike patterns (Fig. 1.5a). For example, the majority of neurons in the barrel field of the rodent somatosensory cortex ("barrel cortex") respond to single whisker deflections with 0 or 1 spike [232] and first spike patterns in primary visual cortex are sufficient to enable visual discrimination [230]. Multi-neuron single spike patterns are well suited to utilise a spike-timing code to address the GRQs (as will be explained). As such, the thesis considers the possible nature of a spike-timing based multi-neuron single spike pattern code.

First spike times following sensory stimulation offer the simplest opportunity to study spike time encoding and are a prime candidate for spike time encoding. With regards to the timing characteristics of multi-neuron single spike patterns, previous analysis of the dataset analysed in Chapters 3 and 4 found millisecond precise multi-neuron spike patterns in the barrel cortex that were organised intra- and inter-laminar, as well intra- and inter-columnar [232], and that stimulus properties (stimulation frequency and location) were decodable from spike time patterns and (weakly noise correlated) spike counts [231]. First spike times improve decoding of whisker stimulus information by 44% over spike counts [211]. Millisecond precise first spike times have also been observed relative to stimulus onset in sensory organs [139, 281, 98] and thalamic areas [226, 265].

As discussed, spike times relative to stimulus onset are considered an unlikely neural coding candidate as they require reference to an internal event for information

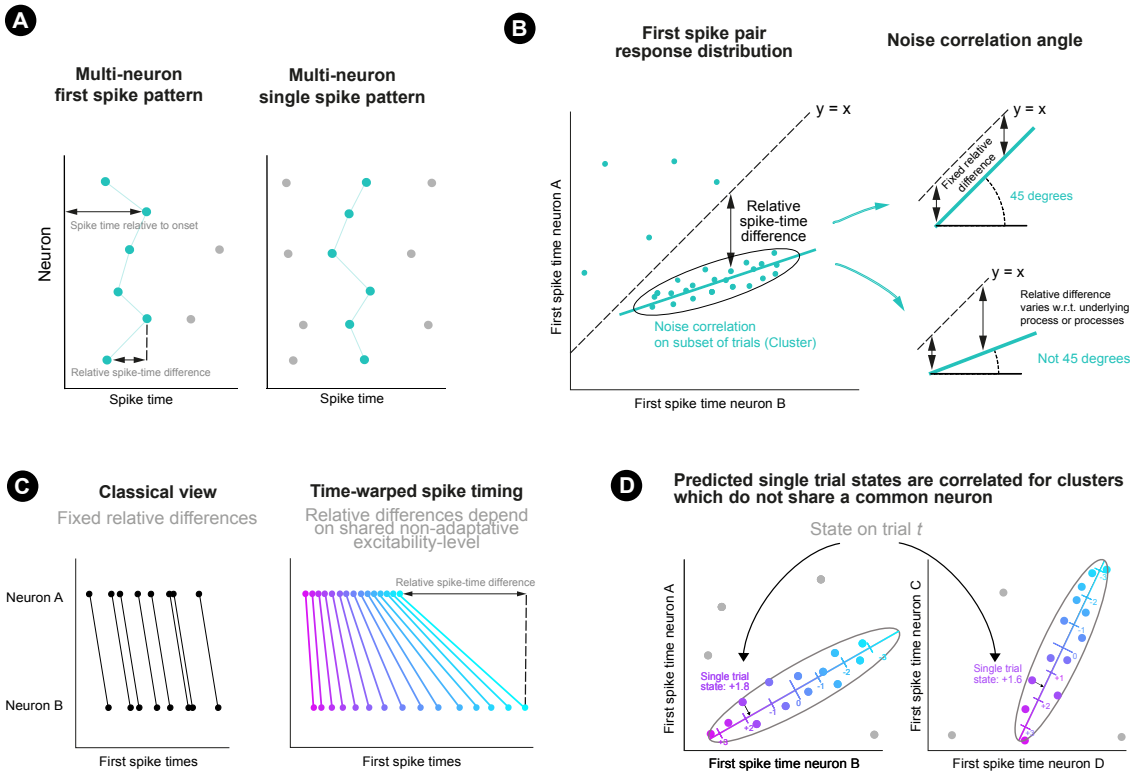


Figure 1.5: (a) Illustration of multi-neuron first spike patterns (left) and single spike patterns (right). (b) Left: Illustrative (artificial) first spike stimulus-response distribution for two neurons and a single stimulus condition. Each point represents the first spike times of the two neurons on a trial on which both neurons spiked. A spike time noise correlation is observable within the isolated cluster. The vertical distance from a point to the line $y = x$ represents the relative spike time difference for the first spike pair. Right: If the angle of the noise correlation is 45°, the relative spike time difference remains fixed (with random noise) w.r.t. the latent process/processes underlying the correlation. If the angle is different from 45°, the relative spike time difference varies w.r.t. the process/processes. (c) Fixed relative spike time differences vs time-warped spike time difference dependence on non-adaptative single trial shared excitability. (d) Illustration of correlated predicted cortical state for disjoint pairs of neurons.

to be decodable by downstream neurons. Precise relative differences between the timings of spikes, which do not require reference to an onset signal to be decoded, have been observed experimentally [139, 98, 232]. Infrequent reporting of precise and reliable relative spike differences, however, is a key argument against spike time coding. Here it is proposed that the temporal variability of spike time patterns requires more detailed characterisation to be well understood.

1.4 Contribution of the thesis

The following subsections summarise the DPhil research. Each of the outlined General Research Questions (GRQs) were addressed to varying degrees, and for some, only hypotheses are formed. As GRQ 1 constrains the remaining GRQs, it is explored through analysis of in vivo spiking activity from adult rat barrel cortex in the first two results chapters (Chapters 3 and 4). The chapters demonstrate how known sources of spike count variability interact to obscure precise spike time patterns between neurons separated over multiple cortical layers and columns.

In the third results chapter (Chapter 5), the characterised form of spike time variability revealed in Chapters 3 and 4 is tested for in a detailed biophysical model of a column of rat somatosensory cortex [175]. It was hoped that the form of spike time variability characterised in Chapters 3 and 4 would be present in the model, in turn allowing the model to be used to make predictions regarding the remaining GRQs. A framework was developed for comparing in vivo and in silico spiking activity. Key differences between the in vivo and in silico spiking activity were found, however.

Chapter 6 then sets out hypotheses regarding how multi-neurons patterns of single spikes (possibly under the characterised form of spike time variability) may facilitate the brain's solution to the FBP; particularly forming hypotheses for the remaining GRQs. Chapter 7 then presents a framework for testing these hypotheses in a phenomenological spiking neural network. This made up the initial work of the DPhil, which aimed to demonstrate the emergence of spike time continua. This included co-creation of a GPU spiking neural network simulator, which was demonstrated (by collaborators) to be the fastest of its class and is also described in Chapter 7.

1.4.1 Chapter 2: Literature Review

The literature review chapter considers candidate neural codes and theories of feature binding with respect to the GRQs. Spike time and spike count codes are assessed as the main alternatives. The chapter supports investigation into the spike timing characteristics of multi-neuron first spike patterns. In general, each results chapter reviews additional literature specific to its domain.

1.4.2 Chapter 3: Clustering and control for adaptation uncovers time-warped spike time patterns in cortical networks in vivo (part 1)

By exploring the nature of spike-timing variability, the first results chapter aims to address the classic dichotomy between spike time and spike count representation. The chapter explores how known sources of spike count variability (adaptation and trial-to-trial changes in shared excitability) affect spike time patterns (in response to single whisker deflections) between neurons separated over multiple layers and columns in the adult rat somatosensory cortex in vivo. The topographic representation of each single whisker on the snout of a rodent in the barrel cortex, provides an excellent model to study responses to controlled sensory stimulation [80]. The results of this chapter were published in Scientific Reports in June 2021 [136].

The chapter analyses first spike stimulus-response distributions for pairs of neurons in response to repeated deflections of a single whisker. Particularly, for each stimulus condition and each pair of neurons, the first spike times of the first and second neuron in the pair were plotted against each other for trials of the stimulus condition on which both neurons spiked (Fig. 1.5b).

Within such 'first spike pair stimulus-response distributions', spike time clusters (isolated regions of high response probability) are found to be positively correlated above chance-levels, representative of spike time noise correlations (i.e. co-modulation of first times by some underlying processes; (Fig. 1.5b)). Moreover the chapter shows that clustering is necessary for the detection of a large proportion of noise correlations, explaining why co-modulation of spike times is rarely detected as a variability source.

To the authors' knowledge, all previous investigations of precise spike time noise correlations have assumed fixed spike time differences, either with a zero [3] or non-zero [138] delay, which would correspond with correlation angles of 45° (Fig. 1.5b) in the stimulus-response distribution. In contrast, the chapter shows that cluster correlation angles are often different from 45° (i.e. non-unitary correlation gradients).

This shows that first spike times of neuron pairs co-vary with respect to some underlying process or processes, such that the relative first-spike time difference varies with respect to this process or processes (Fig. 1.2). Such covariation cannot be solely explained by variability in stimulus onset or cortical onset, which would correspond with a fixed relative difference and 45° correlation angle. The finding of non-unitary noise correlation angles after clustering, explains why millisecond precise spike time noise correlations corresponding to fixed relative spike patterns are seldom reported [139, 98, 232].

The correlations of a subset of clusters are partially modulated by neural stimulus-response adaptation. However, after controlling for adaptation, many clusters remained correlated with non-unitary correlation angles, demonstrating that relative first spike time differences are 'warped' (i.e. stretched/compressed) by trial-to-trial (i.e. between trial) changes in the shared excitability level of neuron pairs (Fig. 1.5c). Detection of adaptation-independent 'time-warped multi-neuron single spike patterns' for neurons separated over multiple cortical columns and layers introduced the possibility of modulation of spike time patterns by population-wide trial-to-trial changes in a population wide cortical excitability (i.e. trial-to-trial changes in continuous cortical state).

1.4.3 Chapter 4: Clustering and control for adaptation uncovers time-warped spike time patterns in cortical networks in vivo (part 2)

Co-modulation of spike times for neurons separated over multiple cortical columns and layers introduces the possibility that spike time patterns are modulated by low-dimensional trial-to-trial population-wide changes in neuronal excitability not caused by adaptation. In support of this (although not conclusive), Chapter 4 shows that predictions of single trial excitability were correlated for a large proportion of pairs of clusters, which did not share a common neuron (Fig. 1.5d).

As discussed, if decoding neurons were also co-modulated by population-wide changes in neuronal excitability, they could integrate spike time patterns in an excitability-dependent (or 'state-dependent') manner. Chapter 4 presents that an improved potential encoding capacity is predicted under the classical and novel assumptions of 'fixed-state' and 'state-dependent' decoding, where state defines the shared population excitability-level at the start of a trial (Fig. 1.4). The results of Chapter 4 described so far were published along side the results of Chapter 3 in Scientific Reports [136].

Chapter 4 presents additional analyses aiming to explore the dimensionality of spike time pattern modulation. Particularly, state predictions from correlated neuron pairs are shown to be correlated for groups of up to 12 neurons, strongly supporting low dimensional spike time pattern modulation. Following this finding a single population state estimate is made for each stimulus trial by averaging over the predicted states of correlated neuron pairs. This population state estimate was used to predict spike times and spike time differences. The single trial population state estimate also

allowed visualisation of spike time pattern modulation by cortical state for up to 12 neurons.

1.4.4 Chapter 5: Comparison of in silico and in vivo responses for a detailed biophysical cortical column model

Following the initial observations of the spike-time phenomena (later characterised in Chapters 3 and 4), the decision was made to test for the same phenomena in a detailed biophysical model of a cortical column. It was hoped that, if the model exhibited the same forms of spike time representations characterised in Chapters 3 and 4, then the model could be used to probe the mechanisms underlying the characterised phenomena and to make further predictions for GRQs 1-5.

For example, the model could be used to probe the nature of spike time modulation by trial-to-trial changes in cortical excitability. The ability to sample all neurons in the model for an unlimited number of stimulus trials could also be used to make predictions about the structure of stimulus response distributions in higher dimensions (i.e. for larger numbers of neurons). It could also be tested whether downstream neurons decode afferent spike time patterns in a state-dependent manner. It was also hoped that the model could be used to make predictions about the nature and emergence of continua encoding, how sensory stimuli are parsed robustly into such representations and how the characterised representations might be affected in the presence of multiple stimuli (i.e. GRQs 2-5).

To undertake this work, the author visited the EPFL Blue Brain Project (BBP) for 3 months in Summer 2019, supported by an academic internship scheme available throught the DPhil grant (ESRC/EPSRC). The author completed the work of the chapter, during the remainder of the DPhil time period, including as a full member of the BBP from November 2020. Over the course of the project, the author developed a tools for comparing in vivo and in silico neuron-group and layer-wise spiking activity, using the same in vivo dataset analysed in Chapter 3 and 4.

The chapter finds and characterises key differences between the in vivo and in silico activity, and the analysis did not detect the phenomena characterised in Chapter 3 and 4. The tools developed for this chapter continue to be used and refined towards validating and and guiding the refinement of future models, in which it is hypothesised time-warped multi-neuron single spike pattern will emerge.

1.4.5 Chapter 6: Hypotheses of time-warped multi-neuron single spike patterns

This chapter sets out hypotheses of how time-warped multi-neuron single spike patterns might address the FBP, as set out by the GRQs. The hypotheses are general across species and modalities and can be tested both in vivo and in silico. The hypotheses were largely inspired by reading during the first years of the DPhil and the modelling work presented in the Chapter 7, which also prompted the analysis of in vivo spiking data (Chapters 3 and 4). The presented hypotheses are updated to account for the results of Chapter 3 and 4.

1.4.6 Chapter 7: Co-creation of a spiking neural network simulator and the emergence of relative spike timing continua in a phenomenological spiking neural network

The final results chapter then presents tools and analysis for testing these hypotheses in a phenomenological spiking neural network. This represents the initial work of the DPhil, which aimed to explore the emergence of spike timing-based feature-binding representations through spike-time-dependent plasticity in phenomenological spiking neural network models. The chapter gives a basic demonstration of the emergence of spike time continua. Particularly, spike times are shown to vary continuously, with continuous changes in the input following 'visually-guided' learning.

Unfortunately, the project was limited by hardware issues that took over a year to resolve, including 4 months of investigation by the manufacturer. During this time, the author began to work on the studies of the other chapters, in order to diversify the thesis. The results and conceptual development of this chapter remained useful in providing the main inspiration for these studies. That is, this work helped foster the hypotheses and investigations of the previous chapters.

The thesis presents the developed tools for demonstrating a proof of principal of the hypotheses of Chapter 6 and exploring these hypotheses further. It was also hoped that such a model could be applied to machine vision tasks. The chapter highlights limitations of the model, which could be investigated using the framework in future research.

To facilitate this work, an advanced GPU spiking neural network simulator was co-created by the author in collaboration with colleagues. The simulator was shown, by others in the laboratory, to be the fastest of its class. A description of the simulator and a demonstration of its performance compared to other spiking neural network

simulators was made in a paper uploaded to the preprint server 'bioRxiv' [8]. The simulator enabled fast investigative simulations for the author and other members of the laboratory in the early stages of the DPhil. The simulator is available for the community on GitHub.

Chapter 2

Literature Review

2.1 Introduction

This chapter reviews the general topics of the thesis, which were not reviewed in Chapter 1. Each results chapter reviews additional literature specific to its domain. Sections 2.2 and 2.3 consider candidate neural codes with respect to the General Research Questions (GRQs) outlined in Chapter 1. This includes discussion relevant to atomic representations of sensory information, including how different experimental approaches and analysis techniques may influence conclusions on neural coding. Analysis techniques for detecting the effect of latent processes on neural representations are also described. Section 2.4 describes several historic theories of feature-binding.

The chapter focuses on research in the visual and whisker domains, but examples from other domains may be given. The visual domain originally inspired the GRQs but they are general across domains. The whisker domain is reviewed in Chapter 3. Chapter 5 reviews large scale biophysical modelling of sensory processing and the first Blue Brain Project cortical column model [175], which is the focus of the chapter.

2.1.1 Correlating neural activity with discrimination

A major aim of neuroscience is to understand the precise correspondence between neural activity and behaviour. GRQ 1 concerns the nature of 'atomic representations' of information, which the thesis defines as the minimal activity within a cortical area that critically supports a perceptual discrimination. Studying stimulus-evoked neural activity offers a method for understanding how sensory information is represented in the cortex. It can also be determined which properties of stimulus-evoked activity correlate with perceptual decisions made about the presented stimuli (e.g. [144]). More recently, the use of optogentics to selectively activate and deactivate neuronal

sub-populations has enabled experimentalists to study how areas are causally involved in perceptual decisions [230, 52].

Humans, non-human primates and rodents (see [40]) can be trained to perform visual discrimination tasks. Go/no-go and forced choice tasks are the simplest visual tasks beyond passive viewing. In go/no-go tasks, the subject can choose either to respond (i.e. go) or not respond (i.e. no-go) with the press of a button, by taking a reward or performing some other action. In a forced choice tasks, the subject must choose between two or more options. Both of these types of task can be used to quantify the subjects perception of a stimulus.

Reaction time, neural processing time and perceptive accuracy vary with the complexity of stimuli and the perceptual task. This must be considered when investigating the atomic representation of sensory information, which enable perceptual decisions. For example, it was shown that reaction time and accuracy of humans and non-human primates in a simple two-way forced choice task depended on image complexity [144].

2.1.2 Fast discrimination and fast cortical processing

In a classic go/no-go task, humans were able to report whether a previously unseen image presented for only 20ms contained an animal [269]. Interestingly, EEG revealed a frontal negativity specific to no-go trials that developed roughly 150ms after stimulus onset, showing that the processing needed to perform the discrimination was achieved in under 150ms. This result has since been proposed as evidence that visual processing is performed quickly in an initial feedforward pass.

As the circuitry of the initial feedforward pass (which passes from the retina to anterior IT cortex) traverses at least eight sets of synapses, it has been proposed by multiple authors (e.g. [269, 283, 133, 60, 230]) that there is only time for each sequential cortical area to process and represent the stimulus rapidly, i.e. over a 10-20ms window in which most neurons spike one or very few times.

Neuroscience has continuously tried to determine the minimal time window of activity required for a cortical area to process and represent a stimulus for purpose of facilitating a perceptual decision (i.e. GRQ 1). This has involved a variety of methods combining pyshcophysics, neurophysiology and information analysis (e.g. [240, 144]), and recently optogenetics [230]. Studies have shown that stimulus information is decodable from short time windows of neural activity; for example, in higher-visual area from spike counts in 10-12.5ms windows [144], and from first spike times in primary visual areas [211, 231].

The thesis aims to improve understanding of the nature of such short stimulus representations, and to consider their role in perceptual decisions. The thesis considers that studying information representation within an initial feedforward pass could uncover principles of information representation that are general to neural processing through-out the brain, and for recurrent processing unrelated to an initial 'feedforward pass'. The following sections consider candidate neural codes for atomic representation, with particular focus on the initial feedforward pass of sensory processing.

2.2 Candidate neural code: spike count, mean firing rate and population spike count codes

Trial-to-trial variability in single neuron spike counts and spike times classically motivated analysis of mean firing rates [270, 251], which are calculated by averaging spike counts; either over trials of an experimental condition, over neurons in a population or over time [94]. Demonstration of significant differences between mean firing rates for different experimental conditions and correlating such differences with behaviour, lead to the common assumption that the brain uses a mean firing rate code, in which spike counts encode information and spike count variability cancels out over a population.

Success in decoding information from mean firing rates and correlating mean firing rates with behaviour lead to the common assumption that the brain uses a mean firing rate code. For example, analysis has been highly successful in determining what types of information are represented by neurons in different areas (e.g. [213, 35]), which has proven useful in the search for spike timing representations (e.g. [211, 232, 231]).

Multi-trial information required to calculate trial-averaged mean firing rates is not available on single trials. As neurons in the brain show high spike count variability, and do not reliably encode stimulus information on single trials, it is now largely accepted that stimulus information is reliably encoded by a population of neurons as a whole (e.g. [93, 221]). Particularly, it is assumed that such "distributed" population codes solve the problem of single neuron variability by recruiting a population of neurons to collectively encode information.

The concept of population spike count encoding is commonly explored through either 'decoding' analysis or 'information analysis' (see [221, 155] for review and Section 1.1.1.3). Both approaches commonly incorporate the notion of spike count encoding by conducting analysis of binned spike counts, where spike counts are binned

either over a single time window or multiple time windows. To characterise the stimulus-response distribution in this form, spikes are binned with bin sizes on the order of 2-20ms (typically 10ms; e.g. [144]). As bin sizes are increased, however, temporal information is lost, making the code what is classically considered to be a spike count code (rather than a precise spike timing code).

The question remains as to whether spike counts over time windows $\sim 10ms$ are fundamental form of representation in the brain, and indeed whether they are the most fundamental form of representation. As downstream neurons are sensitive to the precise timings of their inputs [92, 29, 30], it is possible that observed spike count encoding offers only a coarse interpretation of a precise spike timing regime. On the other hand, it is also possible that a spike count code may be complimentary to spike timing.

2.2.1 Techniques and considerations

Before discussing spike count codes in relation to the GRQs, it is useful to introduce how spike count encoding is often studied. This include several considerations, which influence how neural data is interpreted (especially with respect to the GRQs).

2.2.1.1 Low trial counts per unique stimulus

Historically, a major goal of electrophysiology has been to characterise the stimulus preferences of neurons in different brain areas. Towards this goal, single experiments often present a large battery of stimuli in order to determine neuronal preferences. As the duration of single experimental session is often limited by the subjects attention, the time needed to make a behavioural decision (in a behavioural task) or by the duration that anaesthesia enables reasonable experimentation, the number of presentations per unique stimulus is also limited; often less than 5 [15] and rarely greater than 100.

For low trial counts per unique stimulus, detailed characterisation of underlying multi-neuron spike count and spike time response distributions is often not possible. The impossibility of characterising spike-timing trial-to-trial variability in such studies partly motivates the application of population based approaches. This in turn further reinforces the spike count encoding view, in which neural variability cancels out across the population. Importantly, failure to characterise the true nature of neural variability risks missing both the fundamental form of information representation in the brain and the effect of latent processes on neural variability, which may also be relevant to neural coding.

2.2.1.2 MUA activity confers higher spike counts

Multi-unit activity (MUA) refers to spikes detected from extracellular recordings at single electrode sites without the use of spike sorting to assign action potentials to individual neurons. Spike thresholding is used to detect MUA action potentials; particularly an action potential is assigned to a recording site each time the extracellular potential recorded at the electrode site passes above a certain threshold.

Calculation of mean firing rate statistics and quantification of information content are often made for multi-unit activity (MUA) rather than for spike sorted single neurons (e.g. [121, 144]). MUA responses are therefore likely to represent the responses of multiple neurons such that they are abstracted from single neuron responses. In turn, MUA activity provides higher spike counts than single neurons. It is important to note that such higher spike counts could enable decoding algorithms to perform better than they would for the lower spike count neuron-by-neuron spike trains. On the other hand, the loss of information about which neurons contributed which spikes could in fact disadvantage decoding from MUA activity.

2.2.1.3 Non-atomic stimulus presentations and large spike count windows

Visual stimuli are often presented for periods longer than 100ms in electrophysiology studies [121, 43, 144]; often on the order of several hundred milliseconds or seconds. Whilst stimulus presentations on the order of several hundred milliseconds accord with natural fixation times between saccades, high spike count values may be elicited by the temporally-protracted stimulus presentation and may not correspond to atomic visual stimuli.

Whilst cortical integration of thalamic spike trains over longer time periods is useful for the detection of motion in vision, it is not necessarily the case that cortical activity over several hundred milliseconds represents the presence of the stimulus within a short time window of several tens of milliseconds, say at the start of the stimulus presentation. Indeed, neural activity at 100ms after stimulus onset is likely to largely reflect the stimulus at a different point in time than neural activity at 200ms following stimulus onset. High-spike count values counted over periods of several hundred milliseconds may confer higher encoding capacity than exists for short or atomic stimuli. As such, in the visual domain, it is often difficult to draw general conclusions about the direct relationship between atomic sensory information and neural activity.

2.2.1.4 Backward masking

Even a brief 16ms visual stimulus can trigger a response in the visual cortex that lasts hundreds of milliseconds [240]. This means that the length of the visually-evoked response can not be made short by reducing the length of the stimulus presentation, which would be useful for determining the minimal time window of activity needed for an area to support a visual discrimination.

Backward masks are often noisy stimuli presented shortly after a main stimulus, and have classically been used in vision to disrupt and thus attempt to control the amount of visually-evoked neural activity (e.g. [240]). For example, if a visual mask is presented to humans 20ms after the presentation of a 16ms stimulus, humans are able to identify the stimulus [239]. Under these conditions, neurons in the macaque temporal cortex were shown to encode stimulus information within a short 30ms time window [240], providing further evidence for rapid visual processing in a cortical area.

Still, it is difficult to disambiguate the effect of a backward mask on perception and whether neural activity corresponds to the suppression of the stimulus-evoked response or to the generation of activity in response to the mask [167]. Moreover, backward masks are not area specific and so can not address the minimal duration of activity for a single visual area.

2.2.1.5 Optogenetic control: multi-neuron single spike patterns are sufficient for visual discrimination

As stated, rodents are able to discriminate visual stimuli in both go/no-go and forced choice tasks ([40]), allowing study of the neural correlates of the perceptual decisions. The wide range of mouse genetic models now allows the use of optogenetics to control neural activity within specific populations and study the effects on perceptual decision making. For example, optogenetics can selectively, rapidly and completely silence neuronal activity in a given brain area [161] at any delay following stimulus presentation [227]. This can be used to control the duration of visually-evoked activity in a cortical area during a discrimination task.

In a recent study, mice were trained to perform a go/no-go task, which relied on primary visual cortex to discriminate between gratings that were orientated differently by 45 degrees [230]. The authors silenced cortical activity by optogenetically activating inhibitory neurons in the primary visual cortex (contralateral to the visual stimulus) of transgenic mice, which selectively expressed the microbial light activated cation channel Channelrhodopsin2 (ChR2) in inhibitory neurons.

The authors explored the effect of silencing primary visual cortex activity at different time points following stimulus-evoked cortical onset. The authors found that 40-80ms of activity in primary visual cortex was sufficient to support reliable visual discrimination between the orientated gratings. The authors used additional lesion studies to confirm that the primary visual cortex was necessary for the visual discrimination task.

Most interestingly, within this short time period the majority of recorded neurons in the primary visual cortex spiked 0 or 1 times, providing convincing evidence that 40-80ms multi-neuron single spike patterns critically support visual discrimination, and could be an atomic form of cortical representation. The form in which such short multi-neuron single spike patterns encode information remains to be determined, however.

2.2.1.6 Multi-neuron single spike pattern represent atomic stimuli in the whisker domain

The discrete and short nature of single whisker deflections offer a temporally atomic stimulus ($\sim 2ms$ in the study analysed in Chapters 3 and 4). The rodent whisker system therefore allows study of responses to controlled atomic sensory stimulation [80]. The majority of neurons in the barrel cortex respond to single whisker deflections with 0 or 1 spikes [232], supporting multi-neuron single spike patterns as an atomic encoding form. This provides evidence of a sensory system which does not seem to use high spike counts to encode sensory information. This is supported by evidence that first spike times improve decoding of whisker stimulus information by 44% over spike counts [211].

2.2.1.7 Sliding time windows

Information content is often calculated for sliding time bins, in order to assess how information encoded by the population changes over time [144]. For example, recent studies have applied linear classifiers to calculate how the amount of information linearly decodable from spike counts within 10ms time bins changes as a function of time following cortical activation onset [144]. Importantly, the decoding accuracy of linear decoders increases over time following cortical activation onset. Moreover, the time it takes for a linear decoder to reach a decoding accuracy equal to behavioural accuracy, referred to as the object solution time (OST), can be studied. Evidence that information is decodable from 10ms time bins supports short cortical representation of information. Further consideration of the specifics of such studies is given below.

2.2.1.8 Spike count noise correlations and stimulus-response adaptation

Spike count noise correlations and shared variability were discussed for spike counts in Section 1.1.1.5 of Chapter 1. With regards to interpretation of studies analysing spike count representations, it is important to consider that some studies, which present large numbers of stimuli in rapid succession (i.e. 200ms inter-trial-interval) may not explicitly control for stimulus-response adaptation (e.g. [144]. This complicates interpretation of the results, with respect to consideration of the underlying response distribution.

When searching for noise correlations caused by changes in shared excitability, it is important to design experimental paradigms, which maximise response stationarity. Ideally, stimuli would be presented at large temporal intervals (to reduce the effect of adaptation) but this is often not done, in order to increase the number of stimuli that can be presented in a single experimental session.

Studies may instead randomly interleave stimuli, to avoid the confound of neural stimulus-response adaptation when assessing spike count noise correlations, although statistical checks for adaptation are not always made [249]. Alternatively, if each block only contains presentations of a single stimulus, the first 20 stimulus trials (over which the majority of adaptation occurs) may be discarded [231].

2.2.1.9 Dimensionality reduction and population trajectories

Shared variability can be studied using techniques designed to transform population activity into low-dimensional spaces. Study of activity in motor areas related to movement has been the main focus for the testing and development of such techniques (see [91, 289, 254] for review). Such techniques are used increasingly to analyse neural activity at the population-level.

Characterisation of population activity in low-dimensions can be made using techniques which either do or do not assume a dynamical model relating the population activity of sequential time windows. For example, in the simplest case, linear techniques such as factor analysis (FA) and principle component analysis (PCA) can be applied to population spike count vectors for single time windows. A simple extension to factor analysis is Gaussian Process Factor Analysis, which effectively smooths activity over sequential time windows with a Gaussian kernel.

Following such dimensionality reduction, the progression of neural activity within low-dimensions can be studied. The manifold hypothesis (see [91]) suggests that population activity progresses within low-dimensional 'manifolds' (hyperplanes within

the original spike count response distribution). The hypothesis explicitly suggests that population activity can be reduced due to a 'high degree of correlation and redundancy across individual neurons'. Moreover, it is argued that such activity is suitable for the 'low dimensionality of the stereotypical motor behaviors used in most motor control studies'. It is important to note that the suggestion explicitly does not consider more complex motor tasks beyond current motor control studies. It may be the case that control of more complex motor movements requires information encoded at the single cell level.

Alternatively to simple methods such as FA and PCA, dynamical models can be used to model population activity. For example, linear dynamical models model the progression of low-dimensional population activity from time window to time window [47, 143, 254]. More recently, autoencoders have been successfully used to fit such models [209]. Such modelling of population dynamics has enabled inference of trial-to-trial differences in the progression of population dynamics [143, 209]. As discussed in the Chapter 1, such analysis of the single trial trajectories of population activity in low-dimensional space supports the dependence of population activity on the cortical state at the start of a trial [143, 209].

Condition specific information can be determined from analysis of activity after projection to low-dimensions using a variety of dimensionality reduction techniques. Recently a two step procedure was applied to recordings of spiking activity from the motor cortex of a paralysed patient to successfully decode alphabetic layers, whilst the patient imagined writing the letter with their hand [292].

The first step aligned spike trains to account for fluctuations in writing speed using a time-warp analysis technique [293] (discussed in Chapter 3. The second step applied a nonlinear dimensionality reduction method called t-distributed stochastic neighbour embedding (t-SNE) to binned spike counts. This projected population activity to a 2-dimensional representation. The projection revealed tight clusters for each character, showing that the 'presence' of a (motor) feature (i.e. the motor intention of a particular alphabetic letter) was encoded by a low-dimensional population representation. How the specifics (i.e. properties) of such a motor movement would be encoded remains an open question and is likely to require consideration of neural activity beyond 2-dimensions.

2.2.1.10 Latent state dependent decoding

As discussed in Sections 1.1.1.6 and 1.1.1.5 of Chapter 1, population-wide spike count noise correlations suggest that spike counts are modulated by trial-to-trial changes in

shared excitability and the thesis suggests that such changes could modulate downstream decoder neurons, enabling them to integrate information in a state-dependent manner. Almost all information and decoding analysis has assumed that decoding neurons are fixed state, and thus conclusions about encoding capacity of spike count neural codes currently apply under the fixed-state assumption. Interpretation of neurophysiology results with respect to the GRQs must consider this.

2.2.1.11 Insight from calcium imaging

The use of calcium imaging has become popular in recent years as it allows the activity of many visually observable individual neurons in a population to be tracked over long periods (see [119] for review). Spatially fine calcium activity also allows the activity of sub-parts of neurons such as basal or apical dendrites to be recorded.

Due to mechanical scanning of two-photon rays and the slow decay time of calcium activity induced by neuronal activation, calcium imaging has a temporally coarse resolution; on the order of several hundred milliseconds. Calcium imaging of neuronal responses does not allow the determination of discrete action potentials that is required to explore spike time coding, although techniques are continuously in development to enable reliable inference of spike times from calcium imaging (i.e. [119]).

Calcium imaging better reflects firing rate representations calculated over time windows of several hundred milliseconds, although recording may be non-linear and there is a low signal-to-noise ratio. The popular use of calcium imaging therefore encourages theories of neural coding based on mean firing rates. For example, calcium imaging has been used recently to explore neuronal correlations across a primary visual cortex [266].

2.2.2 Consideration of spike count encoding w.r.t. the GRQs (visual system)

This subsection considers spike count codes with the respect to GRQs 1-3, which are:

- General Research Question 1 (GRQ 1): How are short sensory stimuli represented atomically in single cortical areas?
- General Research Question 2 (GRQ 2): What is the nature of continua encoding for atomic sensory representations?
- General Research Question 3 (GRQ 3): How are short sensory stimuli parsed (i.e. transformed) into neural representations? And how does the process of

parsing provide selectivity (including against incorrect combinations of low-level features).

2.2.2.1 Fast spike count / mean firing rate encoding

Category information is decodable from population MUA trial-averaged firing rates binned over 10-12.5ms windows of activity in macaque inferotemporal cortex (IT), supporting short stimulus representations [133, 144]. In [133], such representations generalised over a range of object positions and scales, including for novel objects. In [144], linear decoding accuracy reached behavioural accuracy, but linear decoders failed to reach behavioural accuracy on a subset of 'challenge' images.

Both of the studies used a one-versus-all approach whereby one binary classifier was trained for each class of stimuli and the overall classifier prediction was given by the binary classifier with the maximum activation. In both studies, the amount of category and identity information decodable from population of MUA firing rates for single 10-12.5ms bins increases from the onset of activity in IT up to a peak at approximately 40ms after the onset of activity in IT. This supports fast encoding at the onset of activity in IT, with improved decoding at a slightly later (40ms) time point. In [133], 12.5ms single neuron bins typically contained 0-2 spikes (0.18 ± 0.26 spikes/bin, mean $\pm$ SD), supporting fast encoding by single spike patterns, rather than high spike counts.

Object category could in fact be decoded with 70% accuracy using spikes in single 12.5ms time bins. The authors suggest this number as a lower bound for the amount of object information that is linearly decodable from the population. Indeed, in a different paper, the same authors suggest that a finer temporal code is likely to encode [121]. Classifier performance increased approximately linearly with the logarithm of the number of MUA sites in [133], indicative of a distributed rather than grandmother-like representation. Interestingly, however, decoders did not generalise to later windows in [144], quashing the notion of sustained firing rate representations in IT.

It has been suggested that firing rate estimation by downstream neurons requires integration over timescales that are discordant with behavioural speed [269] and the fast propagation of activity through cortical layers and areas [218]. This supports the view of a spike timing encoding regime based on single spike patterns in short time windows, rather than a regime based on high spike counts.

2.2.2.2 Property and continua encoding in the visual system

Property and continua encoding was classically explored through analysis of neuronal tuning profiles, which are widely considered to be largely obfuscated by noise at the single cell level, leading to the view that property encoding is achieved at the population level. Classically, in the monkeys, identity-specific information (e.g. category membership) is thought to be represented in the ventral stream, whereas other visual variables (e.g. position) are represented separately, either in the dorsal stream [99, 280] or perhaps directly accessed from V1 [25].

It was recently shown, however, that for trial-averaged MUA firing rates calculated over 100ms, the amount of non-categorical object information (i.e. position, pose and other properties) increases along the ventral stream [121], as it does for categorical information. The authors also found that the accuracy of a decoder applied to activity from IT was consistent with measured human behavioural patterns across tasks (i.e. determination of categorical and non-categorical object information). These results show that information is broadly distributed throughout the IT population, rather than concentrated in task specific specialist units.

The face-patch system, provides the opportunity to study property continua encoding by category (face) encoding cell assemblies. IT face-specific neurons are examples of neurons which encode the presence of a feature (i.e. face) through their activation, whilst information regarding feature-properties is encoded through their activity pattern (trial-averaged firing rates [43]). Particularly, neuronal population modelling of neurons across different temporal lobe patches (ML/MF & AM) suggests that, neurons project incoming faces, as vectors, onto specific axes of high dimensional space.

The results of both [121, 43] were for trial-averaged MUA firing rates calculated over windows $> 100ms$, leaving open the question of how property information might be encoded over shorter time windows in higher visual areas. The authors of [121] found a limit on the spatial resolution of object properties that could be decoded, and highlighted that a finer temporal code is likely to underlie finer property encoding.

2.2.2.3 Superposition and competition

Multiple spike count and firing rate representations cannot superimpose without disrupting one another. For example, if that neuron is activated by two simultaneously presented features and if property information is encoded by the spike count of a neuron, how is the property information of both features encoded simultaneously? If the

firing rate of the neuron equalled the sum of its firing rates in response to each of the features presented in isolation, then the explicit encoding of the property information for the two features would be lost.

Several firing rate competition schemes have been suggested for simultaneously active overlapping representations [15]. Previous research showed that cells reduced their responses in the presence of additional objects [253] or that responses could be explained by a weighted averaging of their responses to two objects in isolation [305]. Winner-take-all and contralateral-take-all rules have also been found in face patches, and depend on the stimulus condition [15].

2.3 Candidate neural code: spike timing

Spike time codes offer the potential of a coding form that would leverage the known temporal sensitivity of neural integration [92, 29, 30] and plasticity [176]. Whilst precise and reliable temporal spike patterns have been observed in response to sensory stimulation [281, 139, 98, 265, 226, 232], a large proportion of spike time variability remains unexplained [251].

Such unexplained variability leads to the common assumption that the brain uses a population code based on spike counts [231, 162] or trial-averaged firing rates [121, 43, 144, 266]. The dichotomy between spike time and rate-based codes remains the subject of considerable debate and study [33, 233, 306]. Chapters 3 and 4 aim to better characterise the nature of spike time variability for pairs of cortical neurons over repeated trials of single stimuli, as this constrains the form in which spike times encode sensory information.

2.3.1 Temporally precise multi-neuron first and single spike patterns

As discussed, multi-neuron single spike patterns are well implicated as a candidate encoding form. Analysis of multi-neuron patterns of first spikes following sensory stimulation allows simplified study of the temporal specifics of such multi-neuron single spike patterns. For example, in the barrel cortex, neurons in all layers demonstrate a clear initial response to single whisker deflections, in which the majority of neurons respond 0 or 1 times [232].

With regards to the timing characteristics of multi-neuron first spike patterns, first spike times can have millisecond order precision relative to stimulus onset in sensory organs [139, 281, 98], thalamic areas [226, 265] and early sensory cortices [232], as

well as relative to oscillations [111]. First spike times in rat barrel cortex improve stimulus encoding by 44% over spike counts [211, 306].

Previous analysis of the dataset analysed in Chapters 3 and 4 found millisecond precise multi-neuron spike patterns were organised intra- and inter-laminar, as well intra- and inter-columnar in the barrel cortex [232], and that stimulus properties (stimulation frequency and location) were decodable from spike time patterns and (weakly noise correlated) spike counts [231].

2.3.2 Precise relative spike times, noise correlations and neuronal sensitivity

Neurons are sensitive to the precise temporal and spatial pattern of synaptic activations [92, 29, 30]. Moreover, synaptic plasticity is dependent on the precise timings of synaptic activations [176]. Precise relative differences [138] between the timings of spikes are therefore an obvious form of spike-time encoding, which do not require reference to an onset signal to be decoded by downstream neurons.

Precise relative spike time differences have been observed experimentally [139, 98, 232] but have only occasionally been shown to encode stimulus information, for example when averaged over trials [98]. Infrequent reporting of precise relative spike differences, however, is a key argument against spike time coding. We propose that the temporal variability of spike time patterns requires more detailed characterisation.

Precise relative spike time differences and spike time noise correlations are classically linked. Particularly, under the assumption that relative spike time differences are precise, it is often (and most often implicitly) assumed that the spike times covary with respect to sources of noise such that relative spike time differences remain precise. Spike time noise correlations refer to such co-variation of spike-times.

Previous studies testing for reliable and millisecond precise spike time patterns (for single stimulus conditions) have assumed that relative differences between the spike times of different neurons are fixed with some trial-to-trial noise [138, 3]. Approaches testing for such fixed precise spike time patterns implicitly assume that noise correlations between the spike times of different neurons have correlation angles of 45° (Fig. 3.1c, upper right) in the stimulus-response distribution. This assumption, that relative differences remain fixed with respect to changes in excitability and other latent processes, may limit the characterisation, detection and estimated coding capacity of relative difference representations.

Study of noise correlations have considered and quantified noise correlations angles but, to the author's knowledge, only for spike count and smoothed firing rate metrics

[162, 244, 282] and not millisecond precise spike times and latencies. Whilst many studies consider delayed noise correlations for various delays (on different timescales), the author is unaware of studies exploring the effect of shared latent processes on precise spike time differences on the millisecond scale.

2.3.3 Relative spike time difference continua encoding

To the author's knowledge there is no known evidence of continua encoding by relative spike time differences in the brain. The author is also not aware of studies, which explore the tuning profiles of relative spike time differences over ranges of stimulus transforms. If relative spike time differences were to be a ubiquitous encoding form in the brain, it might be expected that continuous changes in the transformation of a stimulus, would correspond with continuous changes in the brain's relative difference representation (i.e. continua encoding). Chapters 3 and 4 aims to better characterise the nature of relative spike time difference variability for single stimuli, as a step towards understanding the form of continua encoding in the brain.

2.3.4 Possible connectivity-types underlying relative spike time differences

Precise relative spike time differences between pairs of neurons may be facilitated by three general types of connectivity. For the first type, neuron A may contribute strongly to the activation of neuron B. It is often assumed that this type of connectivity would produce fixed relative differences plus some independent noise (e.g. [138]). However, the thesis argues that this assumption may be incorrect as trial-to-trial variability in shared excitability could cause neuron B to spike either slower or faster after activation from neuron A.

The second type of connectivity describes the case in which neither neuron A or B are strongly involved in activating each other. Instead, neuron A and B are activated by shared or disjoint inputs. Under the common assumption of fixed relative spike time differences, it might be thought that shared or disjoint activations might be uniformly delayed by trial-to-trial differences in shared excitability, such that the relative difference between neuron A or B remains fixed. Alternatively, it might be assumed that the excitability of neuron A and B vary from trial-to-trial, and that this affects the spike times of neuron A and B uniformly, such that the relative difference remains fixed. Again, the thesis suggests that these may be incorrect assumptions,

as it is possible that trial-to-trial differences in excitability vary the time it takes neurons to reach threshold heterogeneously.

The third-type of connectivity captures the intermediate case, in which the neuron A partially activates neuron B, and the comments on fixed relative differences again apply.

2.3.5 Superposition of spike time representations

As discussed in Section 1.1.2.4, the relatively instantaneous extent of action potentials ($\sim 1ms$ means that multiple spike time representations may be able to overlap with minimal disruption. That is, as action potentials are short in duration, neuronal spike time patterns could, in theory, 'superimpose' and encode property information of multiple co-active representations simultaneously (Fig. 1.3b). Spike coincidence detection could decrease the occurrence of downstream neurons responding to or 'interpreting' incorrect features/feature-transforms represented by different spatiotemporal patterns.

As neurons are sensitive to the temporal and spatial proximity of synaptic activations [92, 29, 30], selectivity to particular spatiotemporal patterns could be achieved through organisation of axonal conduction delays, synaptic innervation patterns and synaptic efficacy. Particularly, spike time patterns causing spatially and temporally localised activations of a postsynaptic neuron would be more likely to make postsynaptic neurons spike.

2.3.6 Packet encoding

Representation by multi-neuron single spike patterns is similar but different to the notion of 'packet' encoding [210], which describes the following general notion:

> *'A pattern of firing consisting of a group of neurons firing transiently in a relatively stereotyped sequence. A packet may encode different stimulus features by modulating the relative timing and relative firing rate of the subsets of neurons firing within the sequence.'*
>
> [210]

2.3.7 Synfire chainsm synfire braids and polychronisation

The theoretical notion of 'synfire chains' [4, 2] describes when a volley of spikes in one pool causes a cascade of synchronous volleys in downstream pools. [22] intro-

duced non-uniform delays to the concept to create the concept of 'synfire braids', in which pools of neurons can cascadically activate asynchronously. These notions are different from the notion of precise multi-neuron single spike patterns, in which neurons in a 'pool' can fire precisely and asynchronously. Moreover, no mechanism of self-organisation was given for either synfire chains or synfire braids. Synfire chains were suggested as a solution to the feature-binding problem.

The term 'polychronisation' initially introduced by [138] is a general term for non-synchronous precise spike time pattern, in which a group of neurons spike precisely in a reproducible order.

2.4 Previous theories of feature-binding

2.4.1 Binding-by-synchrony

Binding-by-synchrony (BBS) [185, 170] was the first coordination theory of feature-binding that continues to be considered today. BBS suggests object property representations are associated by firing in synchrony. This requires coordination of distant and spatially-localised feature representations. In support of BBS, proponents showed that a fraction of local MUA pairs oscillate in periodic synchronicity [103, 32]. BBS does not consider that activity patterns can encode property information explicitly. The periodicity of synchrony for BBS is now considered to be in the gamma range, leaving few phases over which spikes could be associated to different features, contradicting our simultaneous perception of many spatial relationships.

2.4.2 Feature-integration theory

Studies by Anne Treisman and colleagues gave rise to one of the first major theories of feature-binding, which continues to be influential to thought on the role of attention in perception [296]. Feature-Integration Theory (FIT) claimed that there was a basic set of features that could be efficiently processed 'pre-attentively' in parallel, whilst feature combinations had to be inefficiently processed serially [275, 274, 272, 273]. FIT is a coordination theory of feature-binding, where 'attention' or re-entrant processing binds spatially distanced feature representations into a single object representation, by preferentially activating feature representations in a single spatial locus, as such:

'... features are registered early, automatically, and in parallel across the visual field, while objects are identified separately and only at a later stage, which requires focused attention. We assume that the visual scene is initially coded along a number of separable dimensions, such as color, orientation, spatial frequency, brightness, direction of movement. In order to recombine these separate representations and to ensure the correct synthesis of features for each object in a complex display, stimulus locations are processed serially with focal attention. Any features which are present in the same central "fixation" of attention are combined to form a single object.'

[275]

Additional theories on the role of attention have also since been proposed, such as Guided-Search, in which pre-attentive feature maps are used to guide the process of search [298, 297]. Guided Search explains many cases where search time is on continuum between efficient and inefficient search. Classically, inefficient serial search requires processing time on the order of 20ms per object, which interestingly accords with the length of representations supported by the thesis.

Following FIT's conception, examples soon emerged showing that serial attention is not always necessary for the detection of conjunctive features [124, 212]. Famously, Houck and Hoffman used the McCollough Effect, an orientation contingent colour after effect, which shows that after adapting to, say, red vertical and green horizontal gratings, black vertical and horizontal gratings appear pink and greenish respectively [181]. Houck and Hoffman showed that the effect could be produced without attention to the adapting stimuli [124], even though the effect requires the association of orientation and colour.

FIT has been refined multiple times since it's initial conception [272, 273], to account for such findings. Treisman now supports that in a:

'first rapid pass through the visual system before attention is brought to bear, or when attention is focused elsewhere' ... 'visual stimuli are initially registered as implicit conjunctions of features by local units in V1' with 'information about their properties implicitly present to be extracted by further processing.'

[273]

Treisman therefore accepts that feature conjunctions and property information are jointly encoded early in the visual system in an initial forward pass, rather than totally independently, as originally suggested by FIT. The thesis considers that representations of features and objects require joint representation of feature-presence and feature-properties, and considers such representations generally for features and ob-

jects in an initial forward pass. That is, the thesis aims to constrain feature-binding by working towards an understanding the *form* in which features and objects are represented.

As discussed, study of the FBP was partly motivated by its potential to describe the richness of our perception beyond single objects, in contrast with the original FIT view of serial processing of conjunctive features. Treisman now supports a dichotomy between focused attention (for objects) and distributed attention (for whole scenes) [273], with distributed attention shown to be useful for a statistical understanding of a scene [45]. Treisman also accepts that even more recent versions of FIT seem to miss the richness of experience:

> *'Is this enough to get us seeing the everyday world? I am not sure. There are some puzzling results in the literature. So far I have discussed results obtained with simple displays, designed to tease apart particular mechanisms underlying performance. However, the world we deal with is normally rich, complex, full of meaning and emotional associations.'*
>
> [273]

2.4.3 Population codes

The high capacity of population codes may be used to argue that they provide a solution to the FBP. That is population codes are proposed to solve *'the combination problem'* [287], which states that a near infinite number of cells would be needed to encode every possible feature-presence/property combination individually. The thesis considers that such an explanation of feature-binding fails to define specific answers to the requirements of the FBP and the GRQs. For example, what is the nature of such a population code, how can it be used to guide 'meaningful' behaviour and how do such representations self-organise to represent novel stimuli?

2.4.4 Communication through coherence

Communication through coherence (CTC) [84, 86, 85] suggests periodic synchronization within pre- and postsynaptic groups creates effective connectivity [299] by focusing excitability to short windows [85], such that presynaptic groups activating a postsynaptic group first or with higher coherence 'beat' other groups, which provide excitation during a subsequent inhibitory phase. Evidence supports modulation of such 'selective coherence' by attention [26, 85]. An updated theory accounting for

delays proposes a mechanism for selective 'top-down attention' through bidirectional coherence, in which the representation with highest coherence is maintained [235, 85].

2.4.5 Cortex-wide grid cells

In the entorhinal cortex, locations are represented by combinations of grid cell activations [107], and environments by unique sets of such location representations. It has recently been suggested that an object may analogously be encoded by 1000s of grid cell systems encoding locations of sensor patches w.r.t. the object [112]. Continuous object-transformations would 'continuously' vary the neurons encoding the object, offering no way of making associations to all transforms. The theory does not address learning, parsing or superposition, and has not been observed in previous studies across the cortex.

Chapter 3

Clustering and control for adaptation uncovers time-warped spike time patterns in cortical networks in vivo (part 1)

3.1 Introduction

The neural coding problem addresses how information is encoded by, and decoded from, patterns of action potentials (i.e. spikes) in the nervous system [93, 140, 94]. Whether information is encoded by spike counts or spike times, remains the subject of debate and study [33, 233, 306]. Responses following sensory stimulation can be analysed to determine principles of neural coding.

Precise and reliable temporal spike patterns have been observed over repeated trials of single stimuli [281, 139, 98, 265, 226, 232] and can improve coding (over spike counts) in primary sensory cortices [211, 306]. Nevertheless, a large proportion of spike time variability remains unexplained [251].

With the definitive form of spike time encoding undetermined, population spike count (i.e. mean firing rate) encoding is commonly assumed, in which spike count variability averages out over a population [162, 121, 43, 266]. Spike time encoding, however, offers the possibility of a fast, efficient and higher capacity coding form, that could leverage the known temporal sensitivity of neural integration [92, 29, 30] and plasticity [176]. This chapter aims to better characterise shared spike time variability in the somatosensory cortex over repeated trials of single sensory stimuli, as this constrains the form in which spike times encode sensory information.

As single cortical neurons can generate highly reliable spike trains to direct stimulation in vitro [168], in vivo response variability largely reflects cortical network dy-

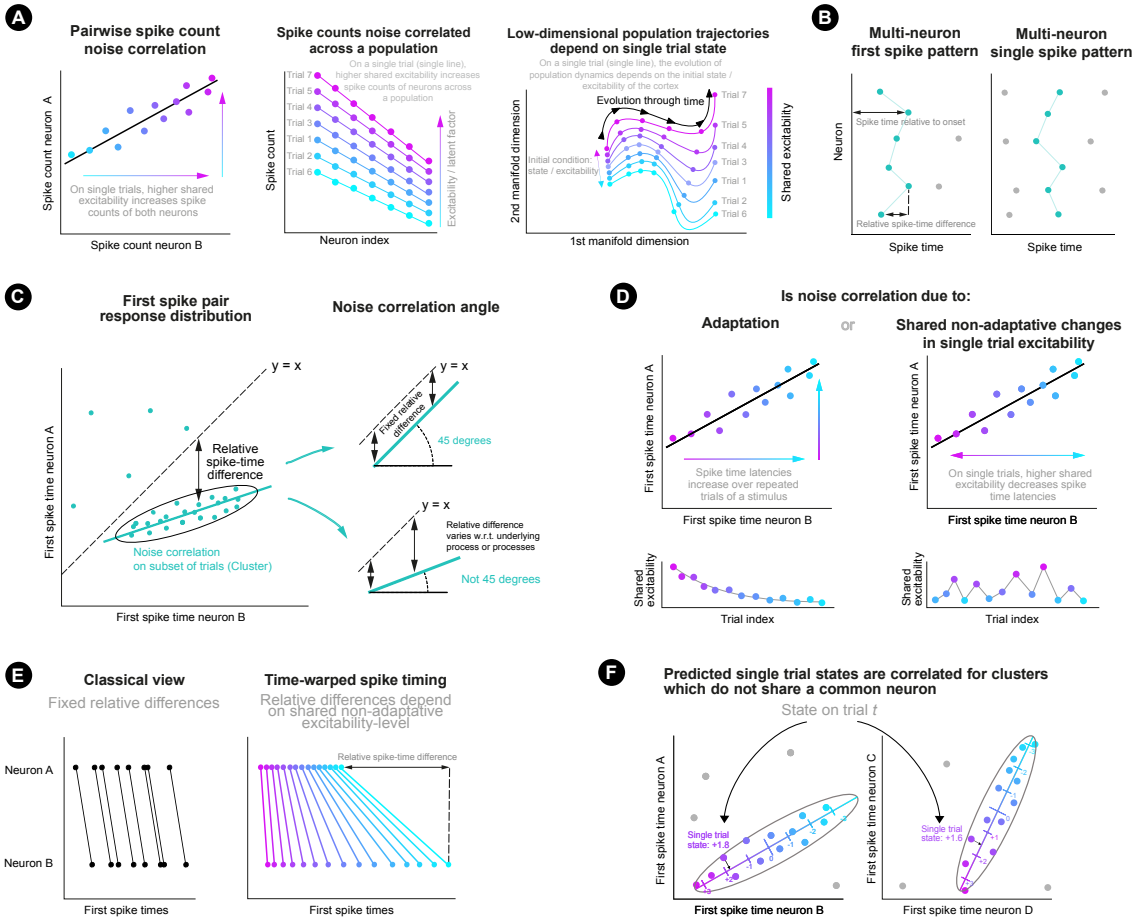


Figure 3.1: Explanatory figures. (a) Illustration of shared spike count variability over repeated trials of a condition for a neuron pair (left), for a population (center) and exhibited as alternative single trial trajectories through low-dimensional space (right). (b) Illustration of multi-neuron first spike patterns (left) and single spike patterns (right). (c) Left: Illustrative (artificial) first spike stimulus-response distribution for two neurons and a single stimulus condition. Each point represents the first spike times of the two neurons on a trial on which both neurons spiked. A spike time noise correlation is observable within the isolated cluster. The vertical distance from a point to the line $y = x$ represents the relative spike time difference for the first spike pair. Right: If the angle of the noise correlation is 45°, the relative spike time difference remains fixed (with random noise) w.r.t. the latent process/processes underlying the correlation. If the angle is different from 45°, the relative spike time difference varies w.r.t. the process/processes. (d) The latent process underlying the noise correlation may be due to adaptative (left) and/or non-adaptative (right) changes in the shared excitability-level of the two neurons from trial-to-trial. (e) Fixed relative spike time differences vs time-warped spike time difference dependence on non-adaptative single trial shared excitability. (f) Illustration of correlated predicted cortical state for disjoint pairs of neurons.

namics [39] and synaptic interactions [31, 46]. Analysis of response variability shared between neurons over repeated trials of a single stimulus ("noise correlations") allows the effect of underlying variability sources to be investigated. For spike counts, shared variability may be studied for pairs of neurons and at the population-level (Fig. 3.1a, left and center [162]). The latter includes analysis, of latent factors and trial-to-trial variability in the low-dimensional trajectory of population activity (Fig. 3.1a, right; [143, 209]).

In the visual cortex, shared spike count variability is modulated by shared trial-to-trial variability in the excitability of neurons across a population. Particularly, neurons across the population respond with either higher or lower spike counts depending on the population excitability level on a single trial (Fig. 3.1a, center). Such shared variability has been explained by an additive interaction of stimulus-evoked activity with spontaneous background activity [13, 249, 162] and an additional multiplicative interaction [162].

A population's state of excitability at the start of a trial therefore affects the representation of a stimulus. Analysis of population dynamics in the motor cortex also supports the dependence of neural processing on a population's state at the start of a trial (Fig. 3.1a, right; [143, 209]). This chapter characterises how precise spike time representations are modulated by the shared excitability-level of spatially separated neurons on single trials.

Multi-neuron first spike patterns following stimulus onset offer the simplest opportunity to study spike time encoding and are supported as a fundamental form of representation [211, 232, 230], which may generalise to multi-neuron single spike patterns (Fig. 3.1b). For example, the majority of neurons in the barrel field of the rodent somatosensory cortex ("barrel cortex") respond to single whisker deflections with 0 or 1 spike [232] and first spike patterns in primary visual cortex are sufficient to enable visual discrimination [230].

With regards to the timing characteristics of multi-neuron single spike patterns, previous analysis of the dataset used here found millisecond precise multi-neuron spike patterns were organised intra- and inter-laminar, as well intra- and inter-columnar in the barrel cortex [232], and that spike time patterns and (weakly noise correlated) spike counts encoded stimulus information (stimulation frequency and location) [231]. It has been shown that first spike times can improve decoding of whisker stimulus information by 44% over spike counts [211].

Precise relative differences between the timings of spikes do not require reference to an onset signal to be decoded by downstream neurons (Fig. 3.1b) and have been

observed experimentally [139, 98], including for this dataset [232]. Infrequent reporting of precise relative spike time differences, however, is a key argument against spike time coding. Here it is proposed that the temporal variability of spike time patterns requires more detailed characterisation.

3.1.1 Chapter overview

The topographic representation of each whisker on the snout of a rodent in the barrel cortex, represents an excellent model to study timing characteristics of multi-neuron single spike patterns in response to controlled sensory stimulation [80]. Using a large dataset of spike sorted recordings in anesthetized adult rat barrel cortex this chapter and Chapter 4 aims to characterise the non-trivial timing variability of multi-neuron first spike patterns for groups of neurons separated across layers and columns of the barrel cortex.

Firstly, for each stimulus condition and each pair of neurons, the first spike times of the first and second neuron in the pair were plotted against each other for trials of the stimulus condition on which both neurons spiked (Fig. 3.1c). Within such 'first spike pair stimulus-response distributions', spike time clusters (isolated regions of high response probability) were found to be positively correlated above chance-levels. Such positively correlated spike time clusters represent spike time noise correlations (i.e. co-modulation of first spike times by some underlying latent process(es); Fig. 3.1c).

The author chose to analyse the correlated and isolated regions of high response probability (i.e. clusters), as they represented reliable regions of the response space. Firstly, it was possible that analysis of these regions, with noise excluded, would provide a better characterisation of the properties of the pairwise distributions and the effects of latent factors (including in relation to larger groups of neurons). Secondly, it was possible that the clusters corresponded to stimulus trials in which the pair of neurons reliably contribute to encoding of the stimulus.

Previous studies testing for reliable and millisecond precise spike time patterns (for single stimulus conditions) have assumed that relative differences between the spike times of different neurons are fixed with some trial-to-trial noise [138, 3]. Approaches testing for such fixed precise spike time patterns implicitly assume that noise correlations between the spike times of different neurons have correlation angles of 45° (Fig. 3.1c, upper right) in the stimulus-response distribution. In contrast, the chapter finds that cluster correlation angles are often different from 45°. This shows that first spike times of neuron pairs co-vary with respect to some underlying latent

process(es), such that the relative first-spike time difference varies with respect to this process(es) (Fig. 3.1c, lower right). This explains why precise fixed relative spike time patterns are seldom reported [139, 98, 232].

After controlling for neural stimulus-response adaptation, which causes spike time latencies to increase over rapidly repeated whisker deflections, many clusters remained positively correlated with non-45° correlation angles. This cannot be explained by variability in stimulus or cortical onset, which would correspond with a fixed relative difference (45° correlation angle; Fig. 3.1d).

Non-45° correlation angles suggest that relative first spike time differences are 'warped' (i.e. stretched/compressed) with respect to an underlying latent factor, which varies from trial-to-trial due to a process other than adaptation. As well as considering other hypotheses, the thesis suggests that the latent factor may correspond with trial-to-trial changes in shared excitability modulating the first spike times of neuron pairs (Fig. 3.1d,e).

Co-modulation of spike times for neurons separated over multiple cortical columns and layers introduces the possibility that spike time patterns are modulated by low-dimensional trial-to-trial population-wide changes in neuronal excitability not caused by adaptation. This hypothesis is explored in Chapter 4. All results in this chapter, as well as part of the results in Chapter 4 were published in Scientific Reports in June 2021 [136].

3.2 The whisker system and barrel cortex

The rodent whisker system offers an excellent model for studying cortical processing in response to controlled sensory stimulation and is one of the most extensively studied models of cortical microcircuitry (see [59, 80] for review). The snout of a rodent is organised in a grid-like layout of about 35 long and thick facial hairs known as whiskers or vibrissae. Fig. 3.2 shows how each whisker is identified by a unique letter-number combination corresponding to its row and arc (e.g., row D, arc 2, or D2).

One of the main advantages of studying the whisker system is the topographic representation of each whisker in the barrel field of the rodent somatosensory cortex ('barrel cortex'; Fig. 3.2). Particularly, in the barrel cortex, each whisker is principally represented by a distinct 'principal barrel' (or 'principal column') of neurons. Neurons within a barrel respond principally to deflections of the corresponding principal whisker, but a proportion of neurons also respond to deflections of whiskers corresponding to neighbouring whiskers. The topographic correspondence between whiskers and barrels allows electrophysiologists to record the activity of neurons which primarily respond to stimulation of a particular whisker.

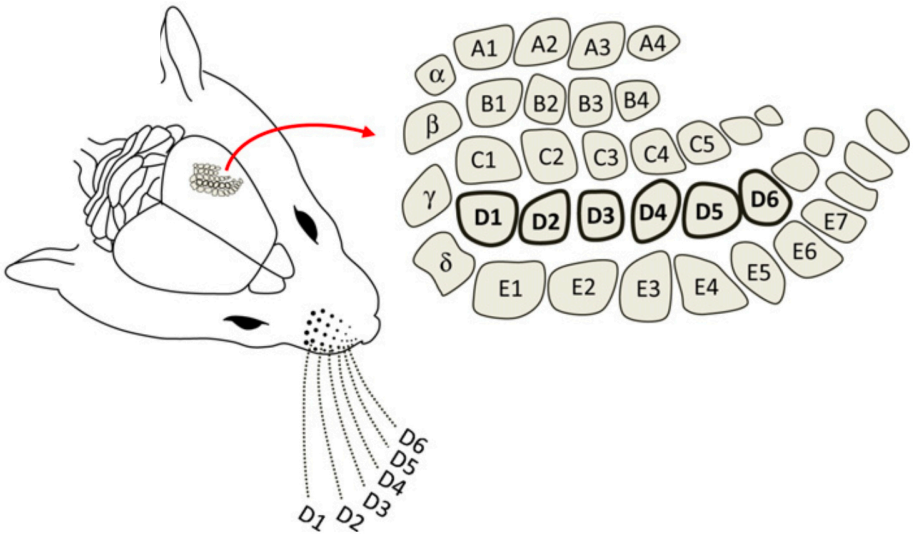


Figure 3.2: Reprinted from [59], with permission from Elsevier. Schematic of the topographic correspondence between whiskers and cortical barrels. The whiskers on the rodent snout are organised in a grid-like layout with each whisker assigned a label based on its row and column. The contralateral barrel cortex is organised topographically in a similar grid-like fashion, with a distinct cortical barrel corresponding to each whisker.

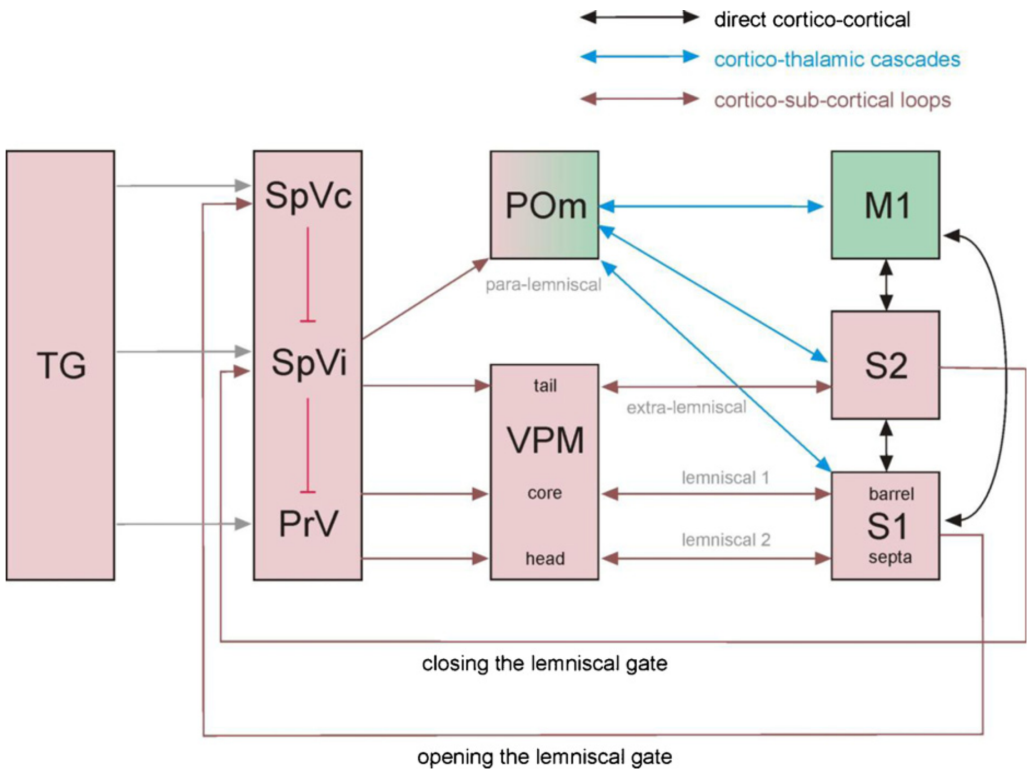


Figure 3.3: Reprinted from [80], with permission from Elsevier. Original caption: 'Schematic demonstrating the ascending pathways in the rodent whisker-related sensorimotor system, as well as possible connections associating the sensorimotor cortical signals. The ascending pathway starts with the primary afferents in the trigeminal ganglion (TG) transducing whisker vibrations into neuronal signals and projecting to the trigeminal brainstem complex (TN). The TN consist in the principal nucleus (PrV), and the spinal sub-nuclei (interpolaris SpVi; caudalis SpVc; the detailed connectivity of the oralis sub-nucleus is unknown and is omitted in the figure). The SpVi falls into a caudal and rostral part (SpVic and SpVir). The classical mono-whisker lemniscal pathway (lemniscal 1) originates in PrV barrelettes, and projects via VPM barreloid cores to primary somatosensory cortex (S1) barrel columns. A second lemniscal pathway originating from PrV has been recently discovered which carries multi-whisker signals via barreloid heads to septa (and dysgranular zone) of S1. The extra-lemniscal pathway originates in SPVic and carries multi-whisker signals via barreloid tails in VPM to the secondary somatosensory area. Finally the parelemniscal pathway originates in SpVir and carries multi-whisker signals via POm to S1, S2, and primary motor area (M1). The different colors of connections indicate three principal pathways through which associative coupling between the sensorimotor cortical areas may be realized. Black: direct cortico-cotical connections. Blue: cortico-thalamic cascades. Brown: cortico-sub- cortical loops. Projections of S1 and S2 may open or close the lemniscal gate (i.e. gate signal flow through PrV) by modulating intrinsic TN circuitry.'

3.2.1 The whiskers and whisker follicles

The whisker system pathway begins when mechanical energy, induced by whisker motion, transmits to the whisker follicle. In the follicle this mechanical energy is transduced into action potentials by neurons in the trigeminal ganglion. These neurons may be sensitive to features of whisker motion such as deflection direction, velocity, acceleration, position, vibration and texture [160, 256, 12] and are either slowly or rapidly adapting [95, 256]. Trigeminal ganglion neurons send mono-whisker signals to the brainstem trigeminal nuclei (TN) [48, 271].

3.2.2 The brainstem trigeminal nuclei (TN)

The TN consists of a principal nucleus (PrV) and three spinal nuclei (oralis, SpVo; interpolaris, SpVi and caudalis, SpVc; Fig. 3.3), which separate processing into four asscending pathways: two 'lemniscal', one 'extralemniscal' and one 'paralemniscal' [58]. Each nuclei is subdivided into histologically defined zones, called 'barrellettes', with each barrellette predominantly representing one whisker but with some neurons integrating information from multiple whiskers. Direction selectivity and response adaptation is also computed by barrellette neurons [186]. Barrellettes are interconnected and receive feedback projections from the somatosensory and motor cortex [90].

3.2.3 The ventro-posterior-medial nucleus

The four ascending pathways pass through dedicated zones in the posterior nuclei (POm) and the ventro-posterior-medial nucleus (VPM; Fig. 3.3), the latter of which is the whisker-related primary somatosensory thalamic nucleus [44]. The VPM contains histological compartments associated with individual whiskers, termed 'barreloids' subdivided into a 'head', 'core', and 'tail' region relaying the different ascending pathways [57, 89].

3.2.4 The barrel cortex

Whisker-related barrel columns are cylindrical structures vertically spanning six layers of the primary somatosensory cortex and receive input predominantly from a single whisker (referred to as the principal whisker). Barrels are arranged topographically in correspondence with the layout of the whiskers on the snout (Fig. 3.2). The borders of barrel columns are defined mainly by dendrites and VPM afferents in L4. Septal columns separate the barrels and are defined by vertically aligned neurons above and

below the septa in layer 4. An S1 dysgranular zone surrounds the barrel field. Mouse barrel columns seem to be organized similarly but with less prominent septal columns.

The four pathways follow different paths to different cortical areas (see [80] for review), including the primary somatosensory cortex (S1), the secondary somatosensory cortex (S2) and the primary motor cortex (M1):

(i) The first lemniscal pathway passes from mono-whisker neurons in the PrV in TN through barreloird cores in VPM, to L4, basal L3, L5B, and 6A of the barrel columns.

(ii) The second lemniscal path passes from PrV multi-whisker neurons to the head of barreloids in VPM and to septal columns in the cortex.

(iii) The extralemniscal path passes from multi-whisker cells in SpVi, via the barreloid tail and terminates in dysgranular S1 and in S2.

(iv) The paralemniscal pathway passes from multi-whisker cells in the SpVi to the anterior POm to barrel cortex L1 and L5a irrespective of barrel and septal columns. L4 innervation from POm is specific for septa. POm also projects to S2 and M1 and several other sensorimotor or somatosensory associative areas.

The internal connections of a single barrel column can be considered with respect to the (disputed) notion of the canonical cortical microcircuit [78] originally proposed by [64]. Indeed, the major excitatory pathways of the canonical microcircuitry defined by Douglas and Martin for the cat primary visual cortex are present in barrel columns (see [79] for discussion). To summarise, the lemniscal pathway passes through the VPM of the thalamus and terminates in L4, as do fibers from LGN in V1. Similarly to V1, the majority of L4 output fibres project to L2/3 which in turn projects to L5. Also as in V1, a thalamo–cortico–thalamic feedback loop exists in which the two major thalamorecipient layers (L4 and L6A) are interconnected and relay activity back to the thalamus.

Importantly, despite such similarities, there are major differences between sensory areas. For example in L4 of the barrel column, unlike in V1, spiny neurons target L5A pyramidal neurons. As another example, L5A pyramidal neurons in the barrel cortex form unusually strong projections to supragranular layers, whilst projections from L6 to L4 have not been characterized in the barrel column.

3.2.5 Neuron types

Through consideration of developmental, molecular, electrophysiological and morphological properties, many studies suggest the classification of inhibitory neurons into at least 5 main types:

(i) Chandelier cells (targeting exclusively axon initial segments)

(ii) Basket cells (believed to target largely somata and proximal dendrites)

(iii) Martinotti cells (targeting distal dendrites, often in L1)

(iv) Neurogliaform cells (thought to signal using volume transmission) and

(v) Bipolar/bitufted cells (often targeting dendritic shafts and spines).

Such subcellular target specificity of the axon of inhibitory neurons is the hallmark of many classifications and has been suggested to be independent of the laminar location of the cell body within a cortical column [174]. See [80] for further classifications of inhibitory neuron types based on connectivity in the barrel cortex. GABAergic interneurons can alternatively be classified based on molecular markers. A recent scheme captures virtually 100% of cortical GABAergic neurons [158], grouping them into the following three non-overlapping groups:

(i) Parvalbumin expressing (PV, a calcium binding protein).

(ii) Somatostatin expressing (SST).

(iii) Ionotropic serotonin receptor 5HT3aR expressing. This is a superset of vasoactive intestinal polypeptide expressing (VIP).

3.2.6 Feedforward inhibition

It is hypothesised that for each major intrinsic excitatory connection, a parallel inhibitory one exists that provides feed-forward inhibition for the pathway (see [80] for review). The thalamocortical synapses from VPM are the best studied case of cortical feedforward inhibition, where the precise connectivity of thalamic fibres with different inhibitory neuron types is thought to set a short window for synaptic integration by excitatory L4 cells (also reviewed by [80]).

Parallel feed-forward inhibition also seems to be implemented in the projections of L4 to L2/3 and L2/3 to L5, such that the three most prominent intrinsic projections

in sensory cortices are controlled by parallel feedforward inhibitory pathways. [80] suggest that *'more work is needed to elucidate the contribution of these trans-laminar double track pathways to the balance of excitation and inhibition and to the restriction of signal flow to narrow time windows.'*

3.2.7 Whisker perception and its neural correlates

Although not the focus of the analyses of this chapter, it is important to consider work related to perception. In 1990, Carvell and Simons [41] first reported that rats can discriminate a smooth surface from one textured with a grid of microscopic groves $30\mu m$ deep at $90\mu m$ distance. As discussed in [80] it may seem unreasonable to consider the relation between perception and neural processing for a single barrel, as decision making is likely to depend on processing by, and interaction with, higher cortical areas. Despite this, it has been shown that disturbing the integrity of a single barrel blocks detection and discrimination of tactile stimuli [188].

The first report of perceptual correlates with barrel column spike trains was performed in 2007 using a choice between drinking from two alternative wells [288]. The authors reported that animals used very short touch times ($\sim 50ms$ per whisker stroke on average) and that increased numbers of touches and higher spike counts correlated with correct discriminations. Stick-and-slip events depending on the roughness of textures were later reported [295], which could contribute significantly to perception.

3.2.8 Stimulus representation in the barrel cortex

Many previous studies have explored stimulus representation in the barrel cortex. This chapter focuses on neural responses evoked by single whisker deflections. Single whisker deflections elicit the largest number of spikes in the principal barrel ([248, 148]. Analysis is complicated by the fact that responses on the fine scale are highly heterogeneous [248, 232], proposed to reflect cell type heterogeneity or sub-networks [247].

This chapter analyses responses in adult rats under anaesthesia. Anaesthesia minimises the impact of motor signals and controls for movement, a well-known source of variability in sensory cortices [267]. Importantly, neural latencies in the principal barrel column are unaffected by the level of anaesthesia [258]. Experimental techniques are continuously refined to allow extracellular recording of many neurons in fully awake head fixed rodents.

The majority of neurons in the barrel cortex respond to single whisker deflections with 0 or 1 spike [232]. This is quantified for the current dataset later in this chapter. As discussed, this implicates multi-neuron first spike patterns as a prime candidate for encoding of single whisker deflections. With regards to the timing characteristics of multi-neuron single spike patterns, previous analysis of the dataset used in this chapter found millisecond precise multi-neuron spike patterns in the barrel cortex that were organised intra- and inter-laminar, as well intra- and inter-columnar [232], and that stimulus properties (stimulation frequency and location) were decodable from spike time patterns [231]. Note again, however, for the spike time pattern analysis, spikes of individual neurons were represented with respect to stimulus onset, as spiking/non-spiking events within 5ms time bins. Moreover, spike counts were found to be weakly noise correlated [231].

In other studies, first spike times have been shown to improve decoding of whisker stimulus information by 44% over spike counts [211]. A two-photon imaging study revealed a pin-wheel like direction map in adult but not juvenile rats [154]. [148] found that pairwise correlations (in calcium transients) between pairs of neurons decreased with inter-neuron distance, including a gradient within a column descending outwards from the centre of the barrel.

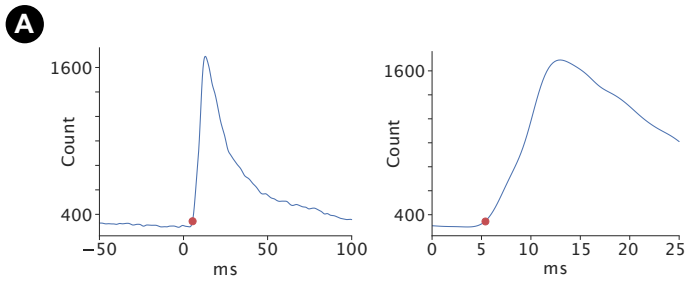
3.2.9 Stimulus-response adaptation and non-stationarity in the barrel cortex

Adaptation describes the 'continuous recalibration of the sensory system to compensate for the changes in the statistics of the input stream' [151], with specific functional roles including gain modulation [6, 114, 149, 189], efficiency [252] and detectability of novel stimuli [65, 279]. The rate and degree of adaptation depends on stimulus parameters, neuron location and neuron properties (see [151] for review).

Adaptation modulates correlations and response heterogeneity (see [151] for review), and depends on brain state and anaesthesia [42, 146]. Adaptation from repetitive stimulation changes excitatory-inhibitory balance [115, 114, 169] thought to be caused by short term synaptic depression (due to neurotransmitter depletion) observed in trigeminal complex [189], thalamus [56] and cortex [115, 114, 169] (mostly at thalamocortical synapses [46]).

In mouse barrel cortex, spiking statistics change through development [27]. Functional imaging in adult mice shows day-to-day maintenance of population activity [172, 180] but variability in the day-to-day stability of single unit representations

[180]. Functional imaging shows that in addition to tuning sensory maps, plasticity and sensory experience affects response gain, variability, firing correlations, and top-down modulation by task context (see [159]). Collectively, these studies characterise the barrel cortex as a highly dynamic system, with representations continuously changing.



	L2/3	L4	L5A	L5B
E	31	58	198	135
I	4	19	16	33

Figure 3.4: Onset of cortical activation following whisker stimulation. (a) Histogram of single unit spike times to all stimuli with bins of 0.2 ms and between -50 ms before and +100 ms after stimulus. The histogram was smoothed with a Gaussian filter (width s.d. 1 ms). The mean and standard deviation of the smoothed bin counts was calculated for the period before stimulus onset at 0 ms. The cortical activation onset spiking threshold was determined as this mean value + 3 standard deviations. Cortical onset was determined as the time point when the smoothed histogram went above this threshold (red dot). The left and right figures show the same smoothed histogram between -50 and +100 ms (left) and 0 and +25 ms (right). (b) Table displaying the number of spike sorted neurons recorded over all experimental sessions by layer and E/I type. Note that L2/3 is undersampled (see Fig. 6b) due to sparse activity and responsiveness of L2/3 neurons [153]

3.3 Methods I: Neurophysiology

The dataset used in this chapter is reused from earlier studies, of which analyses, data collection and spike sorting methods have previously been reported [232, 231, 233, 234]. In summary, extracellular recordings were performed in 18 anaesthetised adult Wistar rats using 16 or 128 electrode silicon probes, inserted perpendicularly into the barrel cortex (details below).

Single whisker deflections were made at various frequencies (0.066-10Hz) in blocks, each containing typically 200 trials (sweeps) with an inter trial interval of 1/f. This was repeated for each of the selected target whiskers (maximum 4). Each block therefore corresponds to repetitions of a unique stimulus condition. Neuronal responses were analysed either following stimulus or cortical activation onset. Most analyses study spiking activity following cortical activation onset, determined as 5.4ms following stimulus onset (Fig. 3.4a).

Experimental sessions varied in terms of electrode placement, stimulated whiskers, stimulation frequencies and number of trials per unique stimulus. 12 experimental sessions stimulated at all frequency groups (0.066-0.2, 1, 2, ..., 10Hz), 2 sessions at a low frequencies and a subset of higher frequencies (0.066-0.2, 1, 2, 4, 7 and 10Hz), and 4 sessions at low frequencies only (0.066-0.2Hz). Data from all experimental sessions

were used in all analyses.

Silicon recording probes consisted of either 1 or 8 electrode 'shanks'. Each shank consisted of 16 electrodes spaced uniformly in a line. Silicon probes were inserted such that the shanks were perpendicular to the cortical surface and spanned multiple layers of the cortex. 8 experimental sessions used single shanks with 16 electrodes, whilst 10 experimental sessions used 8 neighbouring shanks with 16 electrodes per shank covering 2-4 barrel-related columns (median 4).

Spikes were detected using groups of 2 to 4 channels defined as 'virtual tetrodes' [110, 71]. Spike detection was performed using amplitude-thresholding [104]. Spike sorting was performed leveraging waveform feature engineering and open source software [110, 113]. Neuron quality was verified using several additional steps [104]. Cells were classified as putative excitatory (E) and inhibitory (I) neurons based on the mean spike waveform [260, 245, 242].

Voltage-sensitive dye (VSD) imaging, current-source density analyses and histology were used to identify the appropriate whisker-related cortical column and to determine the three-dimensional location of electrodes so that each neuron could be assigned to one cortical layer (L2/3, L4, L5A, L5B/6). Fig. 3.4b shows the number of neurons extracted by spike sorting for each neuron group.

3.4 Results I: Single spike patterns, adaptation and observation of clustered noise correlations

To summarise the experimental paradigm, single whisker deflections (sweeps) were made to single whiskers in blocks of typically 200 trials. For each block, single deflections were made at a certain frequency (between 0.066-10Hz), such that the inter-trial-interval was 1/f.

3.4.1 Confirmation of single spike patterns following stimulus onset

Each neuron-type and layer group demonstrates a clear initial response to single whisker deflections (Fig. 3.5a). Moreover, the majority of responding neurons spike only once (Fig. 3.5b), suggesting low spike count encoding capacity. This also supports previous studies [211] and analyses of the current dataset [232] in implicating multi-neuron first spike patterns as a prime encoding form. This chapter aims to characterise the spike time variability of such multi-neuron single spike patterns.

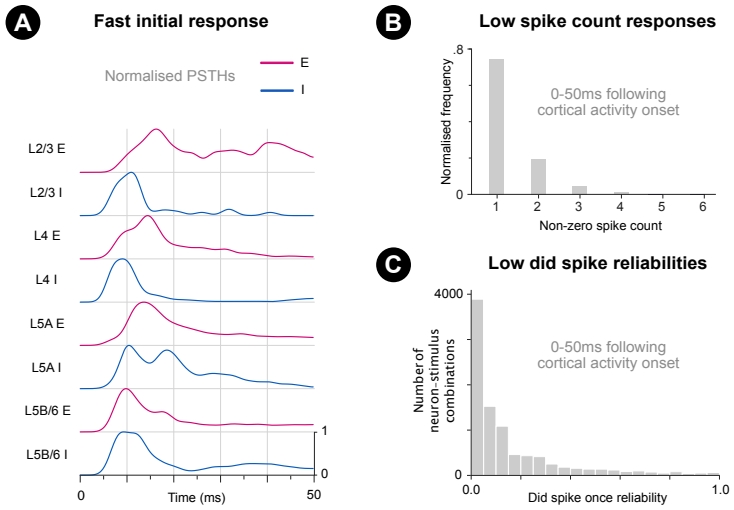


Figure 3.5: Properties of initial responses of cortical neurons to sensory stimulation. (a) Max normalised post-stimulus time histograms (PSTHs) over a 50ms window following stimulus onset (averaged over all trials) for all 8 neuron groups (excitatory in red, inhibitory in blue) in different cortical layers (L) and a 1Hz stimulus. Smoothed with a Gaussian kernel ($\sigma = 1.5ms$) (b) Normalised spike count histogram over all neurons, all trials and stimuli over 50ms following cortical activity onset. Zero spike trials are not included. (c) Histogram of single neuron single stimulus condition did spike once reliabilities. Each value represents the proportion of trials of a stimulus condition that a single neuron spiked at least once.

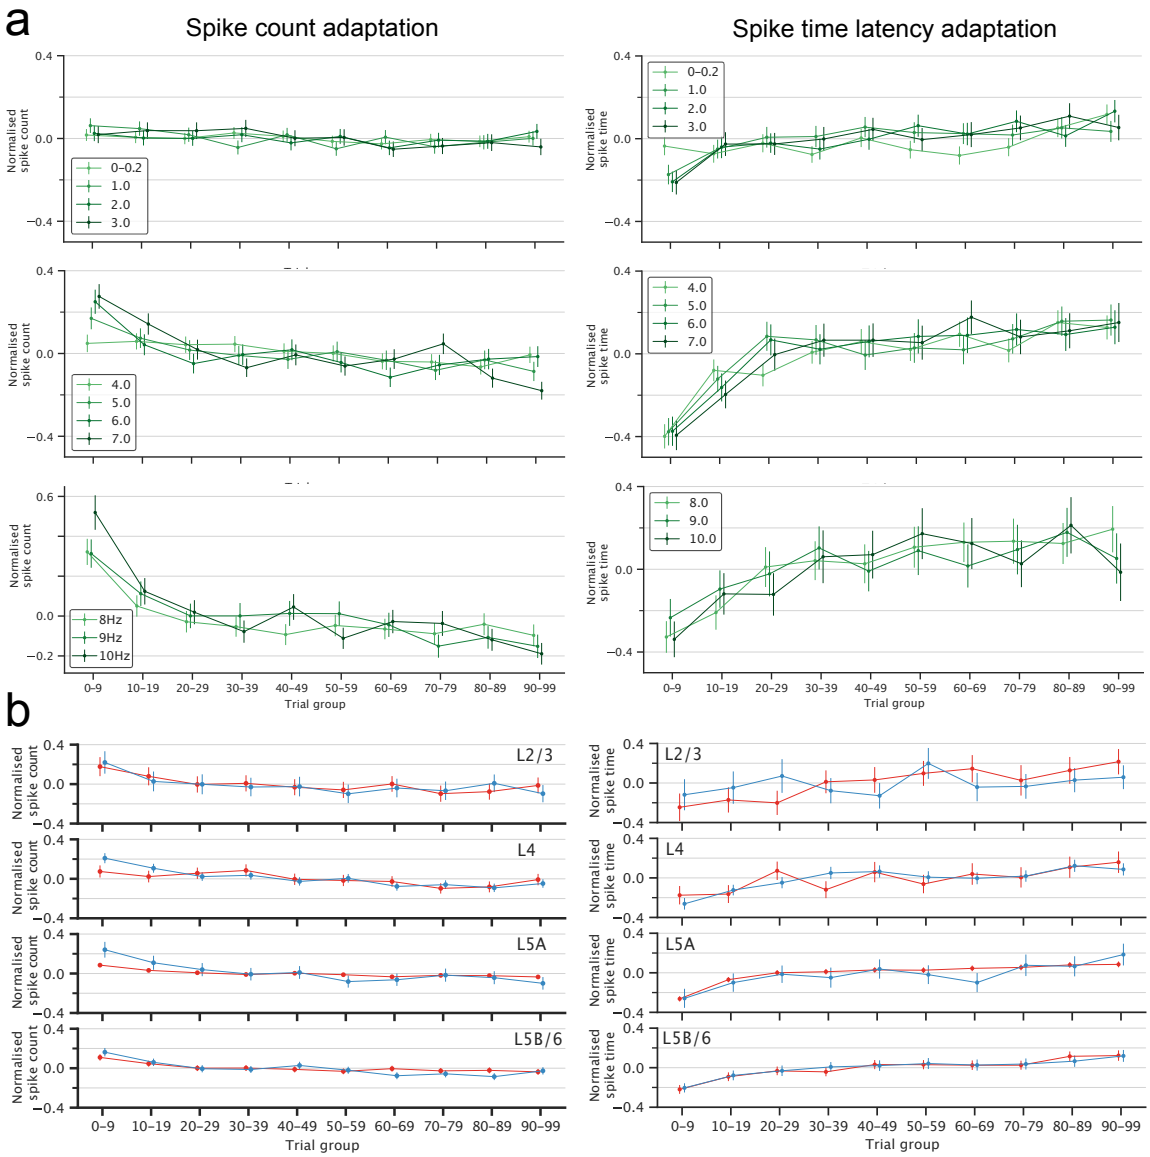


Figure 3.6: Confirmation of neural stimulus-response adaptation (a) By stimulation frequency. Left: For each single neuron and stimulus condition combination, spike counts are normalised (combination mean subtracted followed by division by combination standard deviation). Means of these normalised spike counts are calculated for stimulation frequency groups (0-0.2, 1, ..., 10Hz) and 10 trial bins (i.e. trials 0-9, ..., 90-99). These mean values are plotted with a 95% confidence interval by stimulation frequency. Right: Same as left but for normalised first spike times. (b) Same as (a) but by neuron group. Values for excitatory and inhibitory neuron groups shown in red and blue respectively. Note that L2/3 is undersampled (see Fig. 3.4b) due to sparse activity and responsiveness of L2/3 neurons [153].

Low response reliability (Fig. 3.5c) complicates analysis of multi-neuron single spike patterns. For example, a matrix of first spike times (for a single stimulus condition)

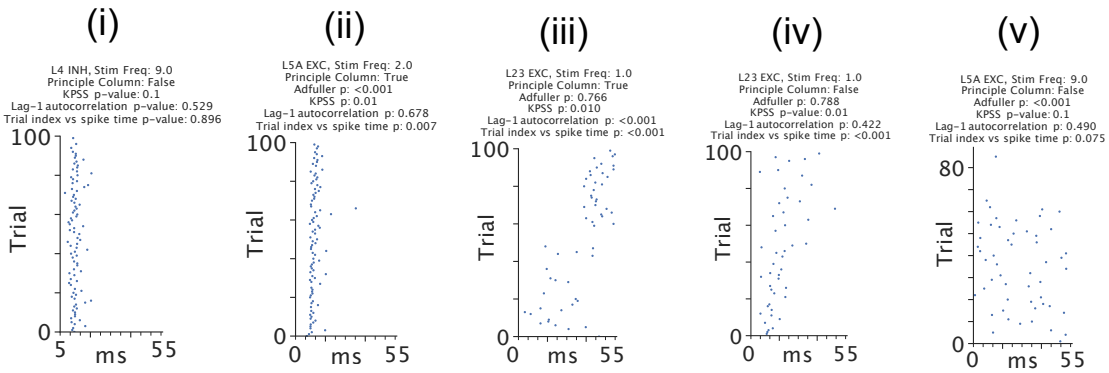


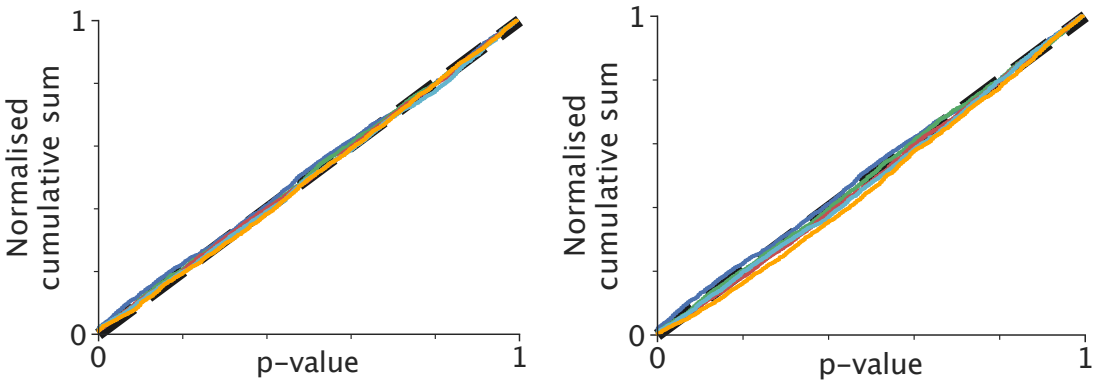
Figure 3.7: Heterogeneous first spike rasters following cortical onset for five single-unit single-stimulus combinations. L4 INH (i) and L5A EXC unit (ii) demonstrate reliable short-latency response between 5 and 10 ms poststimulus with prominent adaptation. L2/3 EXC units (iii and iv) also demonstrate adaptation but responses are more variable. In contrast, L5A EXC unit (v) shows variable responses but not spike-time adaptation.

by neuron and trial index would contain a large number of missing values.

3.4.2 Confirmation of neuronal stimulus-response adaptation for single neurons

Consistent with previous reports [46, 150, 290], including for the current dataset [231], depressive spike count and spike time latency adaptation were confirmed at stimulation frequencies $\geq 5Hz$ and $\geq 1Hz$ respectively, with the majority of adaptation occurring in the first 20 stimulus trials (Fig. 3.6a; Supplementary Fig. 1). Such adaptation was confirmed for all neuron groups, excluding spike time adaptation for the undersampled L2/3 I group (Fig. 3.6b). Such spike time latency adaptation can underlie spike time noise correlations and therefore must be controlled for when testing for noise correlations caused by non-adaptative trial-to-trial changes in shared excitability.

Controlling for adaptative trends must consider the heterogeneity of responses and adaptation over stimulus trials (Fig. 3.7). Autocorrelation and partial autocorrelation analysis of the the trial-to-trial sequence of first spike times of single neurons supported the presence of general adapatative trends without additional autocorrelative structure in the trial-to-trial sequence of single neuron first spike times (Fig. 3.8).



Autocorrelation. Fisher Method (df=4832, N=2416).

Trial-lag	χ^2	p-value
1	5548.9	<0.001
2	5173.6	<0.001
3	5027.3	0.025
4	4954.0	0.108
5	4800.6	0.623

Partial autocorrelation. Fisher Method (df=4832, N=2416).

Trial-lag	χ^2	p-value
2	4806.0	0.602
3	4517.2	>0.999
4	4599.1	0.992
5	4213.8	>0.999

Figure 3.8: Single-unit single-stimulus first spike autocorrelations and partial autocorrelations. Autocorrelation and partial autocorrelation analysis allows quantification of further autocorrelative structure in the trial-to-trial sequence of first spike times. This figure considers the first spike times s_t of single neurons on spiking trials $t = 1, \ldots, T$ of single stimulus conditions where T is the number of spiking trials. Lag-n autocorrelations and partial autocorrelations test the correlations and partial autocorrelations between spike times on trials t with spike times on trial t+n. Left and centre plots show the normalized cumulative sums of lag 1-5 autocorrelation (left) and lag (2-5) partial autocorrelation (center) p-values for 2416 single unit single stimulus combinations. Autocorrelations (right) were above chance for lags 1-4 supportive of general adaptative trends. This is demonstrated by the fact that the cumulative sums are slightly above $y = x$, demonstrating non-uniform distributions of p-values. As this is hard to see, accompanying tables show results (χ^2 statistics and p-values) of Fisher-method application to the p-values produced for each lag. This confirms that autocorrelations and partial autocorrelations above chance for several values.

3.4.3 Observation of clustered spike time noise correlations

Analysing the joint responses of two neurons allows noise correlations caused by adaptation and non-adaptative latent processes such as trial-to-trial changes in shared excitability to be discerned. The first spike times of pairs of neurons were plotted against each other for trials of a single stimulus condition on which they both spiked. For example, for repeated 4Hz stimulation of a single whisker, Fig. 3.9a shows the first spike times of two neurons (L5AI and L4I) in the barrel column corresponding

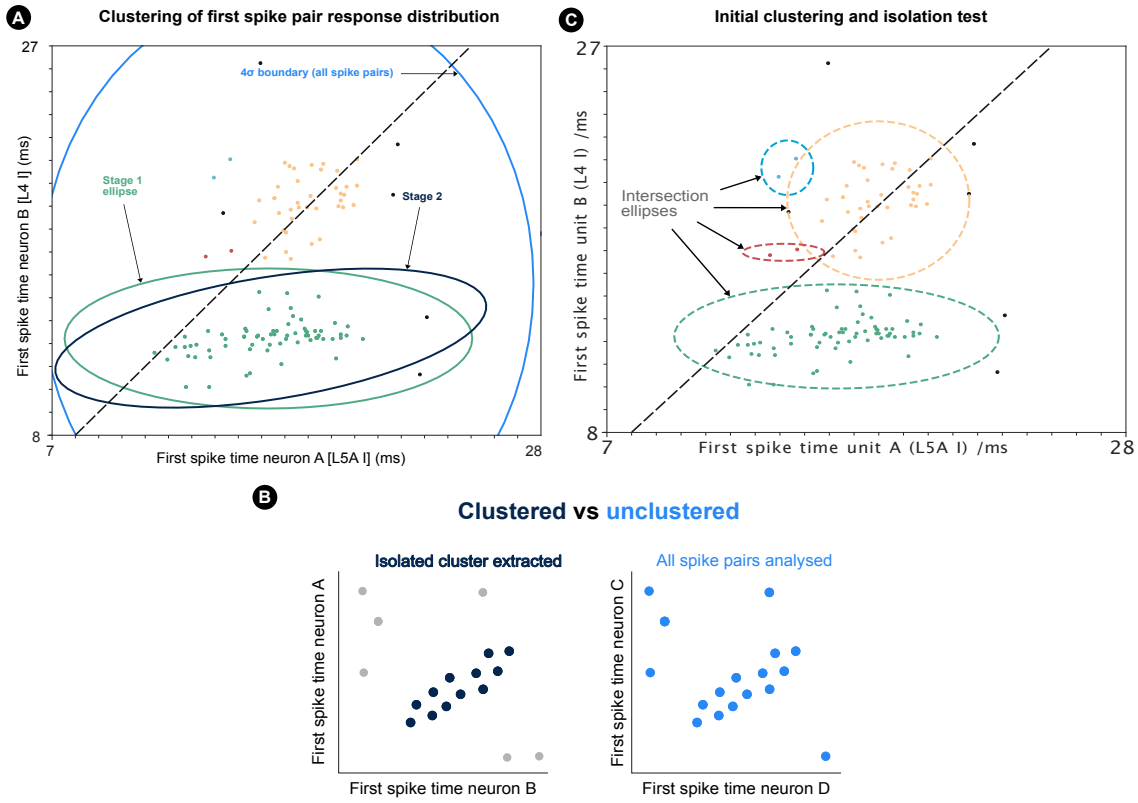


Figure 3.9: Example clustering of sample first spike response distribution for two units (L5AI, L4I) for single epsilon value over a 50ms window following cortical onset (1 outlier out of view). (a) First spike response distribution for two neurons (L5AI, L4I) over 50ms following cortical onset (1 outlier out of view). Point colours illustrate DBSCAN clustering (epsilon = 1.35; Methods). Green flat ellipse illustrates 4σ confidence interval of Stage 1 cluster. Dark blue angled ellipse illustrates 4σ confidence interval of Stage 2 cluster. Light blue ellipse illustrates 4σ ellipse of all spike pairs under independence and normality assumption. (b) Illustration of detected cluster in first spike pair response distribution (left) and unclustered first spike pair response distribution (right). (c) Point colours illustrate DBSCAN clustering (epsilon = 1.35). Black points are outliers. Dashed ellipses represent 3σ intersection ellipse for each DBSCAN cluster. Setting the DBSCAN 'minimum number of elements in a cluster' parameter to 2 increased the cluster quality control, as intersection ellipses were created for 2 element clusters rather than assigning them as outliers.

to the whisker (i.e. principal column) for trials on which both neurons spiked.

First spike pairs appear to come from one of two clusters (isolated regions of high response probability) in the 'first-spike pair stimulus-response distribution'. Moreover, the cluster of green first spike pairs appears to be correlated with a non-45° angle, such that the relative difference between the spike times of the two neurons varies with respect to some underlying latent process(es) (e.g. adaptation or non-

adaptative trial-to-trial changes in shared excitability, for example). Such correlated clusters with apparent non-unitary gradients were observed in many first spike pair distributions (Vid. 1), representing reliable regions of response distribution space, which are noise correlated.

3.5 Methods II: Two-stage algorithm

The occurrence and angle of spike time noise correlations, and the nature of the underlying latent process(es), were quantified both for (i) *isolated clusters* and (ii) the original *unclustered* first spike pair response distributions (Fig. 3.9b). For (i), a two-stage algorithm was developed to quantify the occurrence and angle of correlated clusters observed in first spike response distributions for pairs of neurons, and to allow analysis of the latent processes underlying cluster correlations.

3.5.1 Stage 1 clusters

For the clustered case, a two-stage algorithm was developed to extract isolated clusters which did not contain smaller isolated sub-clusters. Stage 1 of the algorithm utilised 'density-based spatial clustering of applications with noise' (DBSCAN) [74, 238], an established clustering algorithm which detects clusters based on the underlying probability density function. Stage 1 applied two additional steps, firstly discarding DBSCAN clusters, which (like the cluster of orange points in Fig. 3.9c) were likely to be non-isolated sub-parts of larger clusters, and secondly including points within a 4σ ellipse calculated for the DBSCAN cluster points so that correlation tests were unbiased. The details of Stage 1 are described in the following steps:

1. *DBSCAN clustering.* A Python implementation [216] of DBSCAN [74, 238] was used to detect clusters in first spike pair response distributions. Increasing the epsilon parameter increases the resulting cluster sizes. Clustering was performed using the following epsilon values sequentially: [0.4, 0.45, ..., 5.0]. The colour of points in Fig. 3.9c represents the clusters to which points were assigned for epsilon = 1.35. If the previous epsilon value produced the same clustering, then the new clustering was discarded. Stage 1 clusters with less than 30 elements were also discarded. Checks for similar or identical *individual* clusters were not made at this stage to avoid sampling bias, which could affect the uniformity of surrogate control tests.

2. *Isolation test.* This step aimed to ensure that clusters were not sub-parts of larger clusters but were instead well isolated clusters in response distribution space. For each cluster produced in step 1, an 'intersection ellipse' was created (dashed ellipses in Fig. 3.9c), with centre equal to the cluster mean and widths equal to $w_i = 3 * max(0.5, \sigma_i)$ for neurons $i = 0, 1$, where σ_i is the standard deviation of samples in the cluster for neuron i. Any cluster with an intersection

ellipse within or intersecting with another intersection ellipse was excluded from further analysis. In Fig. 3.9c, this corresponded to the orange, red and light blue dashed clusters.

3. *Stage 1 Clusters.* For each remaining cluster, a flat bounding ellipse was created (flat green ellipse in Fig. 3.9a) with centre equal to the cluster mean, and widths equal to $w_i = 4 * \sigma_i$ for neurons $i = 0, 1$, where σ_i is the standard deviation of samples in the cluster for neuron i. All elements within the bounding ellipse were assigned to the Stage 1 cluster. The use of non-angled ellipses was necessary for a controlled test of cluster correlations.

3.5.2 Correlation surrogate control datasets

To test that correlated clusters found in Stage 1 were not artifacts of the clustering algorithm and were above chance levels, the algorithm was applied to the following surrogate control datasets (results shown below):

(i) *All shuffled.* Pairwise shuffling of elements of first spike pair distribution. For each first spike pair stimulus-response distribution, first spike pairs are shuffled so that each spike time of one neuron is matched with a different spike time of the other neuron (Fig. 3.10).

(ii) *Single Gaussian.* Spike pairs drawn from an independent bivariate normal distribution with the same mean and standard deviation as the original spikes. The number of samples drawn from the distribution was equal to the number of original samples (Fig. 3.10).

3.5.3 Stage 2 clusters

An angled 'Stage 2' ellipse was calculated to better estimate the distribution underlying each correlated ($p < 0.005$) Stage 1 cluster. Clusters which were not well isolated or were similar to clusterings produced by lower DBSCAN EPS values were discarded. The details of Stage 2 are described in the following steps:

4. *Angled ellipses.* An angled 'Stage 2' ellipse was calculated to better estimate the distribution underlying each correlated ($p < 0.005$) Stage 1 cluster using a 3-stage bootstrap procedure. In the bootstrap step 1, a number of samples were taken with replacement from the Stage 1 cluster equal to the number of samples

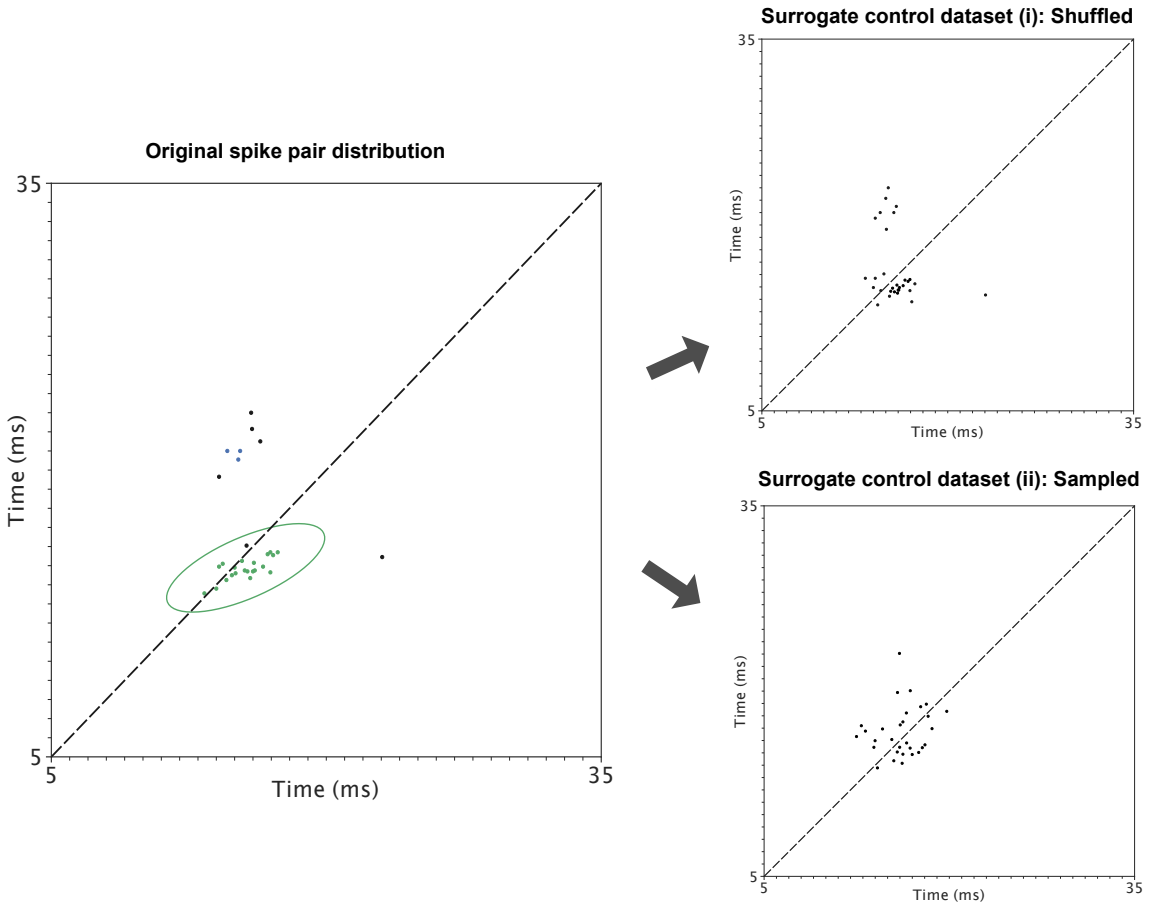


Figure 3.10: Correlation surrogate control dataset generation. Left: Example of original first spike pair response distribution for repetitions of a single stimulus condition and pair of neurons. Right: Generated surrogate dataset examples: (i) original response distribution spike times shuffled and (i) spike times sampled from normal distribution with same single neuron standard deviations as original response distribution.

in the Stage 1 cluster. PCA was performed on the sample to attain a sample mean and covariance matrix. 10,000 such bootstrap iterations were made. The sample means and covariance matrices were then averaged to construct a new bounding ellipse (similar to the rotated dark blue ellipse in Fig. 3.9a) with centre equal to the bootstrapped mean, rotation equal to the angle that the principal component made with the positive x-axis, and widths equal to $4\sigma_i$ where σ_i is the square root of the variance explained by each principal component.

The bootstrap procedure was then repeated by sampling the elements within this new ellipse (bootstrap step 2), with the averaged mean and covariance matrix used to construct a third ellipse for sampling in bootstrap step 3. Elements

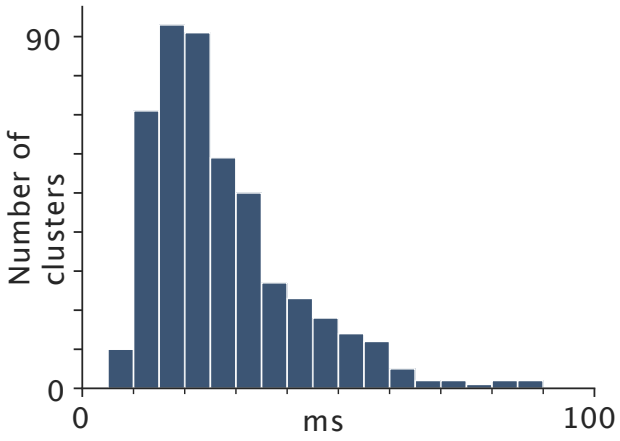


Figure 3.11: Histogram of Stage 1 cluster end points (mean + 4σ) for each Stage 2 cluster.

within the final bounding ellipse defined by the averaged mean and covariance matrix of bootstrap step 3 were assigned to the Stage 2 cluster.

5. *Cluster similarity check.* Angled clusters which shared at least 15 elements to previous angled clusters (produced by lower epsilon values for the same response distribution) were excluded from further analysis. This means that no clusters were included, which contained well isolated 'sub-clusters' produced by lower epsilon values. As a result, all extracted Stage 2 clusters were from unique response distributions, guaranteeing that there were no repeated clusters.

As 94.6% of Stage 2 clusters extracted over 100ms following cortical onset lied within 50ms following cortical onset (Fig. 3.11), analysis was constrained to spikes in the first 50ms following cortical onset.

3.5.4 Comparison to alternative clustering methods

The clustering algorithm was compared to the following two Gaussian Mixture Model approaches:

(i) *Standard GMM.* The first approach uses the classical Bayesian information criterion (BIC) to select an optimal number of Gaussians (between 1 and 10) for the mixture model of each pairwise response distribution. This approach was therefore not constrained to find isolated regions of high response probability.

(ii) *Custom GMM*. An augmented GMM approach was used in which steps 2-4 of the main clustering algorithm were additionally applied after initial BIC based GMM clustering (i). This clustering type therefore included checks for isolation.

3.6 Results II: Clustered spike time noise correlations

3.6.1 Clustered spike time noise correlations

3.6.1.1 Stage 1 clusters

Stage 1 extracted 3929 well isolated clusters defined by the 4σ boundary of a flat 2D Gaussian ellipse (Fig. 3.9b)). As a reminder, Stage 1 did not check for repetitions or similar clusterings produced by different values of the DBSCAN epsilon parameter (to avoid bias). Stage 1 clusters were statistically significantly correlated above chance levels (Fig. 3.12, left; Fisher's method: $\chi^2(7858, N = 3929) = 17308.0$, $p < .001$) as were the unclustered response distributions (Fig. 3.12, right; Fisher's method: $\chi^2(8676, N = 4338) = 14863.4$, $p < .001$).

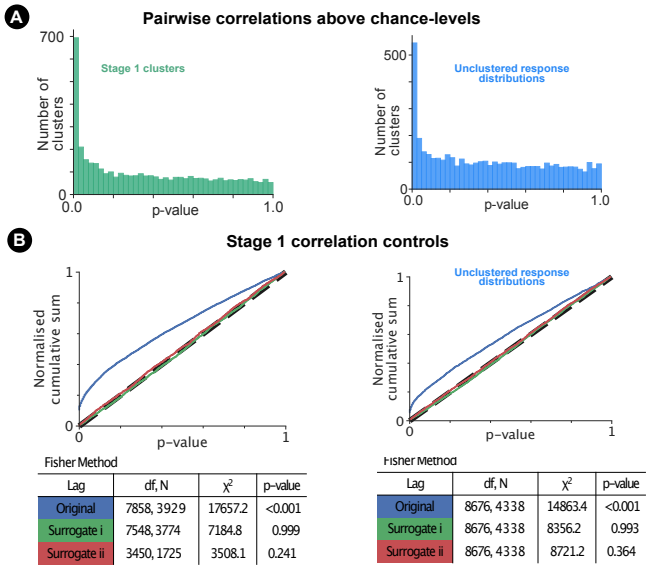


Figure 3.12: Correlation of Stage 1 clusters (left) and unclustered response distributions (right). Top: Linear regression p-values for Stage 1 clusters (left) and unclustered response distributions (right). Center: Normalised cumulative sums of p-value histograms for (blue) Stage 1 clusters (left) and unclustered response distributions (right), (green) surrogate control dataset 1, (red) surrogate control dataset 2. Surrogate controls were generated separately from correlated response distributions for the clustered and unclustered cases. Bottom: Fisher method confirms that Stage 1 clusters and unclustered correlations are not detected above chance-levels for both correlation surrogate control datasets.

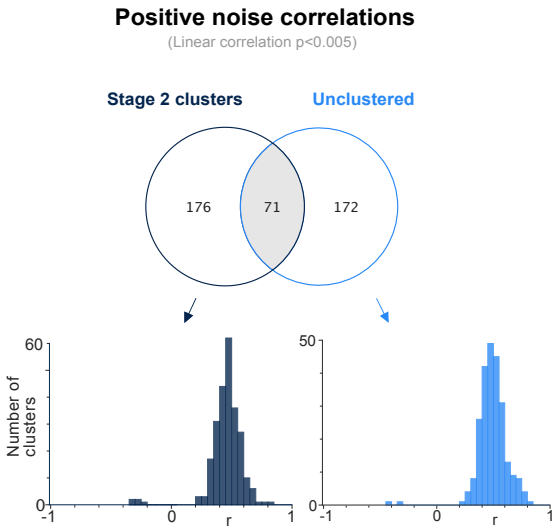


Figure 3.13: Stage 2 clustered and unclustered correlations are positively correlated. Top: number of correlations ($p < 0.005$) detected for the clustered (Stage 2 clusters) and unclustered case. Intersection shows the number of response distributions from which a correlation was detected by the clustering algorithm and for the unclustered case. Bottom: linear regression r-values of correlated ($p < 0.005$) Stage 2 clusters (left) and unclustered response distributions (right).

3.6.1.2 Stage 1 correlation controls

The clustering algorithm found correlated clusters at chance-levels for two surrogate control datasets (Fig. 3.12) confirming that the detection of correlated clusters above chance-levels for the original spike pair distributions were not artifacts of the algorithm.

3.6.1.3 Stage 2 ellipses and correlated unclustered response distributions

An angled 'Stage 2' ellipse was calculated to better estimate the distribution underlying each correlated ($p < 0.005$) Stage 1 cluster (Fig 3.9). Each Stage 2 cluster came from a unique first spike pair response distribution. 247 unique Stage 2 clusters were extracted, 98.0% (242/247) of which remained correlated above chance ($p < 0.005$). Unclustered response distributions were also further analysed, and 243 unclustered response distributions were determined as correlated ($p < 0.005$).

The detected Stage 2 clusters and correlated unclustered response distributions are shown in (Vids. 1 & 2 respectively). In the following sections, the same analyses are performed on the correlated Stage 2 clusters and correlated unclustered response distributions. For a subset of first spike pair response distributions, correlations

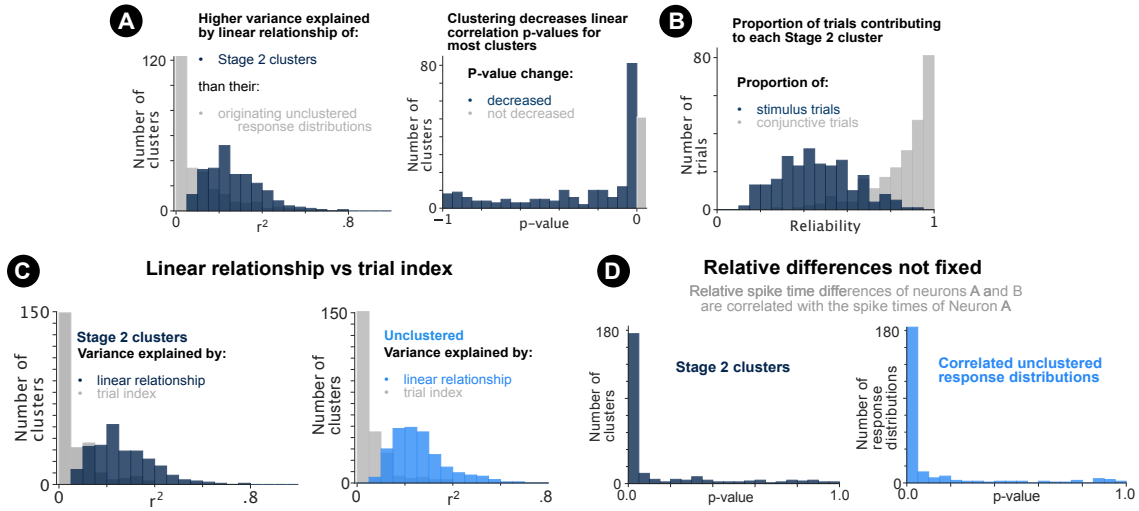


Figure 3.14: Stage 2 cluster statistics. (a) Left: histogram of variance explained (r^2) by the spike time linear relationship of each Stage 2 cluster (dark blue) and the linear relationship of the unclustered response distributions from which Stage 2 clusters were extracted (grey). Right: histogram of change in linear regression p-values introduced by clustering for response distributions from which Stage 2 clusters were extracted. Particularly, histogram shows Stage 2 cluster linear regression p-values minus originating response distribution linear regression p-values. Clustering reduces p-values in most cases (dark blue). (b) Histogram showing for each Stage 2 cluster, the proportion of conjunctive (both neurons spike) trials contributing a cluster sample (grey) and the proportion of stimulus trials contributing a cluster sample (dark blue). (c) Histogram of r^2-values of linear multi-output spike pair prediction from trial index (grey) and linear correlation (dark blue, blue) for Stage 2 cluster spike times (left) and unclustered response distributions (right). (d) Histograms of linear regression p-values of the first neuron's spike time and the spike time difference.

($p < 0.005$) were detected for both the clustered and unclustered case (Fig. 3.13, top).

98.0% (242/247) of Stage 2 clusters and 99.2% (241/243) of correlated unclustered response distributions were positively correlated (Fig. 3.13, bottom) suggestive of modulation by spike time latency adaptation or another latent process such as trial-to-trial changes in pairwise shared excitability.

Linear relationships explained a large proportion of spike time variance within Stage 2 clusters, but not within the response distributions from which Stage 2 clusters were extracted (Fig. 3.14a). In fact, clustering was necessary to uncover a large proportion of spike time noise correlations detected through clustering (Fig. 3.14b). This result partly explains why precise spike time noise correlations are rarely reported. Fig. 3.14c confirms that noise correlated clusters consisted of responses from

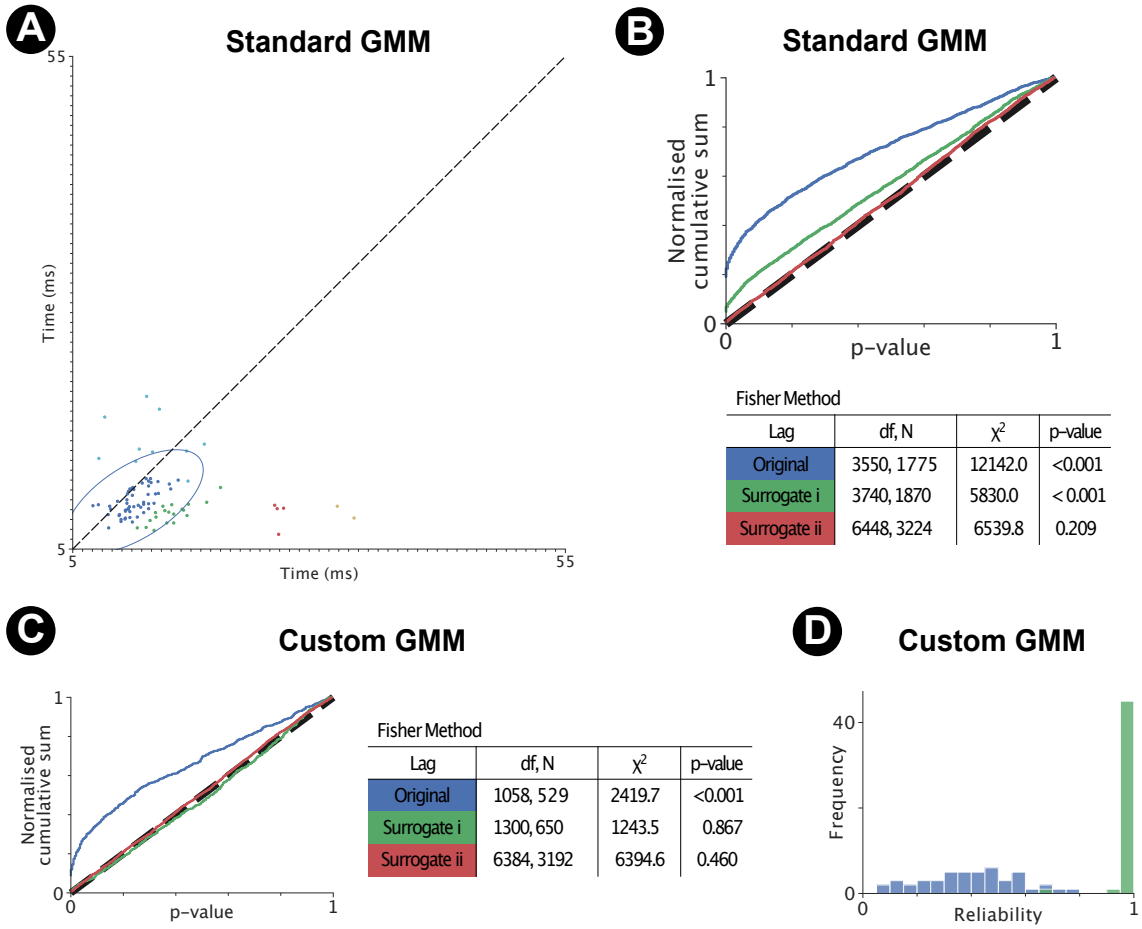


Figure 3.15: Comparison to alternative clustering algorithms. (a) Example detection of stationary correlated cluster (dark blue) in shuffled response distribution through fragmentation by the standard GMM algorithm. (b) Cumulative p-value histograms of correlations detected by the Standard GMM algorithm. Fisher method confirms that Stage 1 clusters are detected by the Standard GMM algorithm above chance-levels for both correlation surrogate control datasets. (c) Cumulative p-value histograms of correlations found for custom GMM algorithm. (d) Histogram showing for each Stage 2 cluster found by the custom GMM algorithm the proportion of conjunctive (both neurons spike) trials contributing a cluster sample (blue) and the proportion of stimulus trials contributing a cluster sample (green).

a subset of stimulus trials.

More variance was explained by the linear relationship between the spike time pairs than by linear models fit to predict the joint spike time pairs from trial number, suggesting that clustered and unclustered noise correlations are modulated more by trial-to-trial changes in a latent factor such as pairwise shared excitability than by adaptation (Fig. 3.14c; tested further below).

For Stage 2 clusters and correlated unclustered response distributions, relative

spike time differences were correlated with the spike time of the first neuron in the pair (Fig. 3.14d). Correlations are therefore not explained by $s_1 = s_0 + d + independent_noise$ where d is a fixed difference and s_0 (the spike time of neuron 0) varies from trial-to-trial.

3.6.1.4 Comparison to alternative clustering algorithms

The clustering algorithm was compared to the two alternative Gaussian Mixture Model (GMM) approaches. As the Standard GMM approach was not constrained to find isolated regions of high response probability, Stage 1 of the Standard GMM algorithm detected clusters, which were fragmentations of regions of high-response probability (e.g. Fig. 3.15a), rather than isolated regions of high response probability (which were the focus of the study). In fact, the algorithm was biased to detect correlations, such that it found above-chance correlations for the two correlation control datasets (Fig. 3.15b). This made it unsuitable for detecting correlated clusters in the main dataset.

For the Custom GMM algorithm, correlations were found at chance-levels (Fig. 3.15c). However, almost all of the detected clusters included all samples from their originating response distribution (Fig. 3.15d), showing that the algorithm did not find isolated regions of high-response probability that contained less than 95% of the samples in the original sample distribution.

3.7 Methods III: Cluster correlation angle and angle surrogate control datasets

To test whether cluster and unclustered correlation angles were different from 45°, the correlation angle θ that each Stage 2 cluster and correlated unclustered response distribution made with the positive x-axis was estimated. $\theta_{45} = min(\theta, 90 - \theta)$ was used, as only the ordering of the neurons determined whether θ was above or below 45°. An additional surrogate control dataset was generated to test that detected differences from 45° were not artifacts of the algorithm.

3.7.1 Cluster angle

To estimate the correlation angle of each Stage 2 cluster and each correlated unclustered response distribution, the 10,000 sample first principal components produced in bootstrap stage 3 were used to estimate the angle that the first principal component makes with the positive x-axis (with empirical confidence intervals). To account for symmetry, sample principal components with negative y-components were replaced by their anti-correlated vector so that all vectors had angles with the positive x-axis between 0 and 180°. To allow averaging and angular difference calculations without the effects of discontinuity, the vectors were rotated anti-clockwise about the origin such that the anticlockwise angles of the vectors with the positive x-axis were between 0 and 360° (double their original angle). A mean first principal component vector was then calculated in this space by averaging over the bootstrap samples. For each cluster, the 250th and 9750th largest bootstrap sample angles were used to define a 95% empirical confidence interval (in this 360° space). The resulting mean vector was then rotated clockwise by half of its anti-clockwise angle with the positive x-axis so that it was between 0 and 180°. This was used as the final θ angle for the cluster. The confidence interval bounds were divided by 2 to create angle confidence intervals for the first principal component on the correct scale.

3.7.2 Angle surrogate control datasets

The following two angle surrogate control datasets were created to verify the main clustering algorithm and compare it to the unclustered case:

(i) 45° *control.* To validate that differences from 45° were not artifacts of the algorithm, a surrogate control dataset was created by replacing each Stage 2 cluster and all samples within a 6σ circle of its mean (where σ was equal to

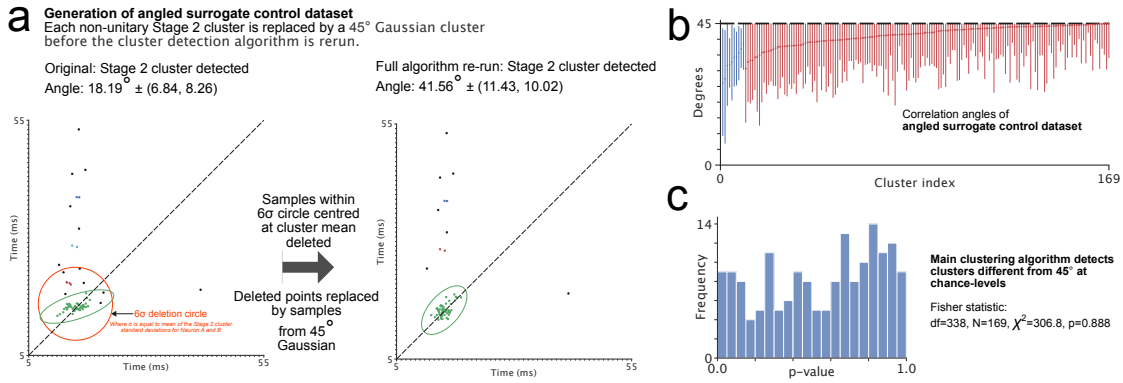


Figure 3.16: Angle surrogate control dataset. (a) Left: Generation of angle surrogate control dataset (i) example. Stage 2 cluster detected by main clustering algorithm (green ellipse). Points within 6σ circle of cluster mean (where σ is equal to the mean of the cluster spike time standard deviations of the two neurons) are deleted and replaced by an equal number of samples from a 45° Gaussian ellipse (right). New response distribution added to surrogate control dataset (i). (b) Mean θ_{45} angles with 95% empirical confidence intervals for positively correlated Stage 2 clusters detected in angled surrogate control dataset. Blue if significantly greater than 0° ($p < 0.025$) and less than 45° ($p < 0.025$). Red otherwise. Detection of clusters different from 45° appear to be at chance-levels. (c) Histogram of p-values for correlation angles being different from 45° for Stage 2 clusters detected by the main clustering algorithm for the generated angle surrogate control dataset. Correlation angles different from 45° are not found above chance-levels.

the mean of the cluster spike time standard deviations of the two neurons) with samples drawn from a 45° 2D Gaussian (Fig. 3.16a). The algorithm found correlated clusters with angles different from 45° at chance-levels (Fig. 3.16b,c).

(ii) *2nd distractor cluster.* An additional surrogate control dataset was created to demonstrate that the clustering algorithm, in the presence of a second uncorrelated cluster, detects the first correlated cluster and calculates its correct correlation angle (Fig. 3.17a). For each response distribution in which a Stage 2 cluster was extracted by the main clustering algorithm, a new response distribution was created with two new clusters. The first cluster was created by sampling a number of samples equal to the number of samples in the original Stage 2 cluster, drawn from a 45° Gaussian centred at the mean of the original Stage 2 cluster, with the new cluster spike time standard deviations of the two neurons equal to the mean of their original standard deviations. The second cluster was created by sampling a number of samples equal to the number of samples in the original Stage 2 cluster, drawn from an uncorrelated 2D Gaussian

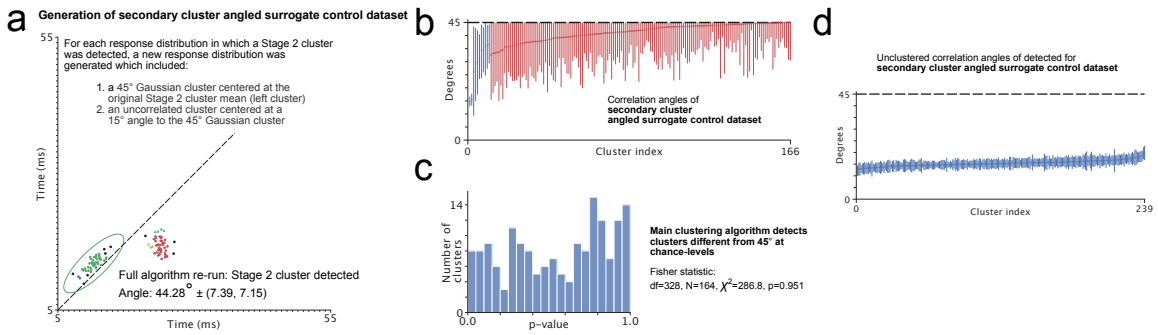


Figure 3.17: Secondary cluster correlation angle test. (a) Example surrogate control response distribution generated by drawing samples from a 2D 45° Gaussian ellipse centred at the original Stage 2 cluster mean (left) and a second uncorrelated response distribution with a mean at a 15° angle to the left cluster (right). (b) Mean θ_{45} angles with 95% empirical confidence intervals for positively correlated Stage 2 clusters detected in secondary cluster correlation angle surrogate control dataset. Blue if significantly greater than 0° ($p < 0.025$) and less than 45° ($p < 0.025$). Red otherwise. Detection of clusters different from 45° appear to be at chance-levels. (c) Histogram of p-values for correlation angles being different from 45° for Stage 2 clusters detected by the main clustering algorithm for the generated angle surrogate control dataset. Correlation angles different from 45° are not found above chance-levels. (d) Mean 45° angles with 95% confidence intervals for positively correlated ($p < 0.005$) unclustered response distributions. As expected, the unclustered approach does not detect the angle of the 45° cluster. Instead, correlation angles are found around 15°.

centred to right at a 15° angle from the first cluster with the cluster spike time standard deviations of the two neurons equal to those of the original Stage 2 cluster. The x coordinate of the mean point was equal to the mean of the original Stage 2 cluster plus 10 standard deviations of the x-axis neuron's cluster spike times.

For this surrogate control dataset, the clustering algorithm detected clusters different from 45° at chance-levels (Fig. 3.17b,c).

The Henze-Zirkler null hypothesis of normality was rejected ($p < 0.05$) for 6.5% (11/169) and 7.2% (12/166) of the Gaussian sampled clusters detected in the two angle control datasets, verifying the algorithm and Henze-Zirkler test could detect normally distributed clusters.

This differed to the unclustered approach, which found correlations with angles $\sim 15°$ (Fig. 3.17d), highlighting an advantage of the clustering algorithm over the unclustered approach.

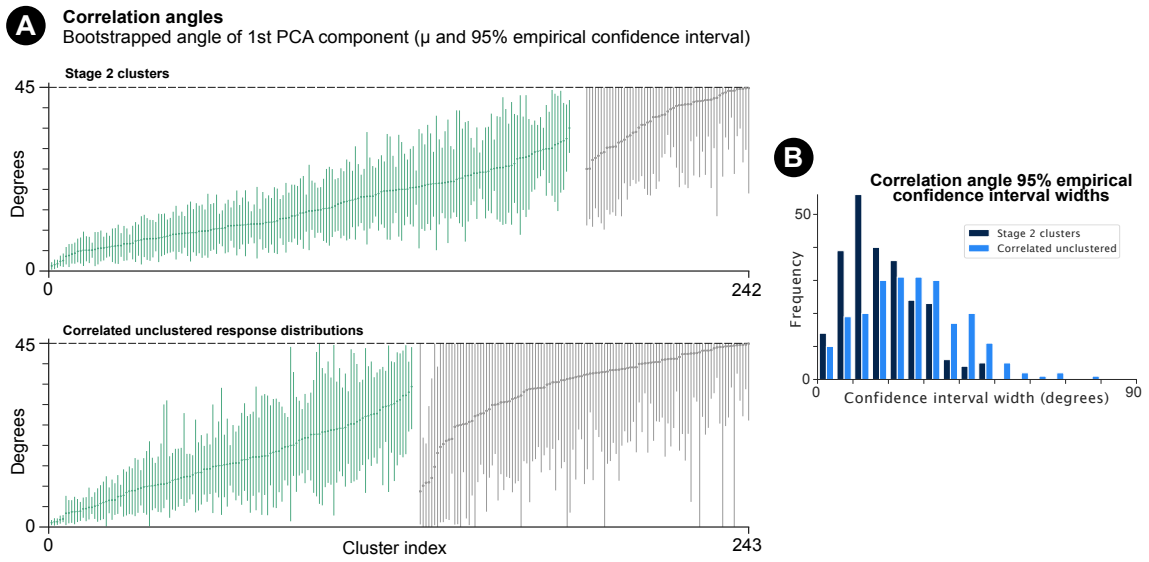


Figure 3.18: Correlation angles. (a) Mean θ_{45} angles with 95% empirical confidence intervals for positively correlated Stage 2 clusters (top) and unclustered response distributions (bottom). Green if significantly greater than 0° (p-value ¡ 0.025) and less than 45° ($p < 0.025$). Grey otherwise. (b) Histogram of cluster angle 95% empirical confidence intervals.

3.8 Results III: Cluster correlation angles are non-unitary

It is important to note that the Henze-Zirkler null hypothesis of normality was rejected ($p < 0.05$) for 71.7% (177/247) of Stage 2 clusters and 76.5% (186/243) of correlated unclustered response distributions.

73.7% (182/247) and 47.7% (116/243) of θ_{45} angles were between and significantly different from 0° ($p < 0.025$) and 45° ($p < 0.025$) for Stage 2 clusters and correlated unclustered response distributions respectively (Fig. 3.18a). This suggests that the relative difference representation of a stimulus is changed by an underlying latent process(es). The clustering algorithm enabled estimations of correlation angles with smaller confidence intervals compared to the unclustered response distributions (Fig. 3.18b).

3.8.1 Comparison to Custom GMM

For the Custom GMM algorithm, almost all of the detected clusters included all samples from their originating response distribution (Fig. 3.19a), showing that the algorithm did not find isolated regions of high-response probability that contained

less than 95% of the samples in the original sample distribution. The Custom GMM algorithm did, however, find clusters with correlation angles different from 45° (Fig. 3.19b), although much fewer than the main clustering algorithm and unclustered approach. The Henze-Zirkler null hypothesis of normality was rejected ($p < 0.05$) for only 31.9% (15/47) of Stage 2 clusters detected by the Custom GMM algorithm, confirming that the Custom GMM algorithm was biased towards detecting clusters likely to be Gaussian distributed.

The Custom GMM algorithm was also verified for the two angle control datasets (Fig. 3.19c,d). Moreover, the Henze-Zirkler null hypothesis of normality was rejected ($p < 0.05$) for 3.8% (1/26) and 3.3% (1/30) of the Gaussian sampled clusters detected in the two angle control datasets, further verifying the Henze-Zirkler test of normality when used in combination with clustering.

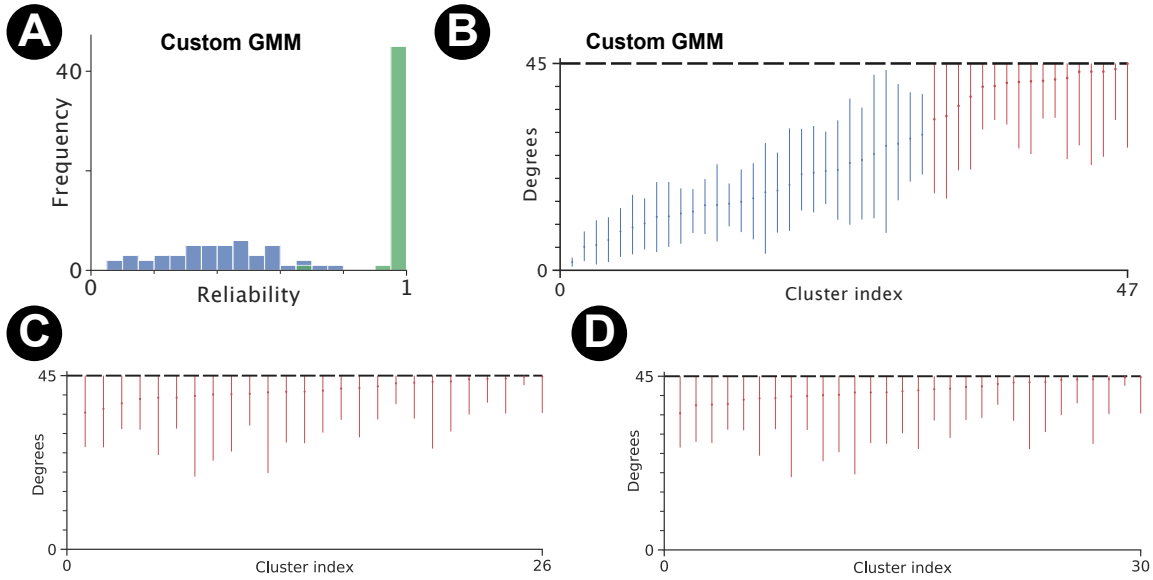


Figure 3.19: Custom GMM correlation angles. (a) Histogram showing for each Stage 2 cluster found by the custom GMM algorithm the proportion of conjunctive (both neurons spike) trials contributing a cluster sample (blue) and the proportion of stimulus trials contributing a cluster sample (green). (b) Angles of Stage 2 clusters found by the custom GMM algorithm. Mean 45° angles with 95% confidence intervals for positively correlated Stage 2 clusters (top) and stationary Stage 2 clusters (bottom) Blue if significantly greater than 0° ($p < 0.025$) and less than 45° ($p < 0.025$). Red otherwise. (c,d) Mean 45 angles with 95% empirical confidence intervals for positively correlated Stage 2 clusters detected by Custom GMM algorithm in secondary cluster correlation angle surrogate control datasets (generated from response distributions in which Custom GMM algorithm originally detected Stage 2 clusters). Blue if significantly greater than 0° ($p < 0.025$) and less than 45° ($p < 0.025$). Red otherwise. (c) Angle correlation datasets (i). (d) Angle correlation datasets (ii). All detected clusters were not significantly different from 45° for both datasets.

3.9 Methods IV: Stationary clusters and predicted encoding capacity

This section tests for neuronal stimulus-response adaptation to determine when non-45° correlation angles are due to trial-to-trial changes in a non-adaptative latent factor such as the shared excitability of neuronal pairs. This section first describes the criteria for a Stage 2 cluster to be determined as stationary. The results of these tests are then presented. The correlation angles of stationary Stage 2 clusters are then analysed, and a prediction is made for their potential encoding capacity under a fixed and state-dependent coding regime.

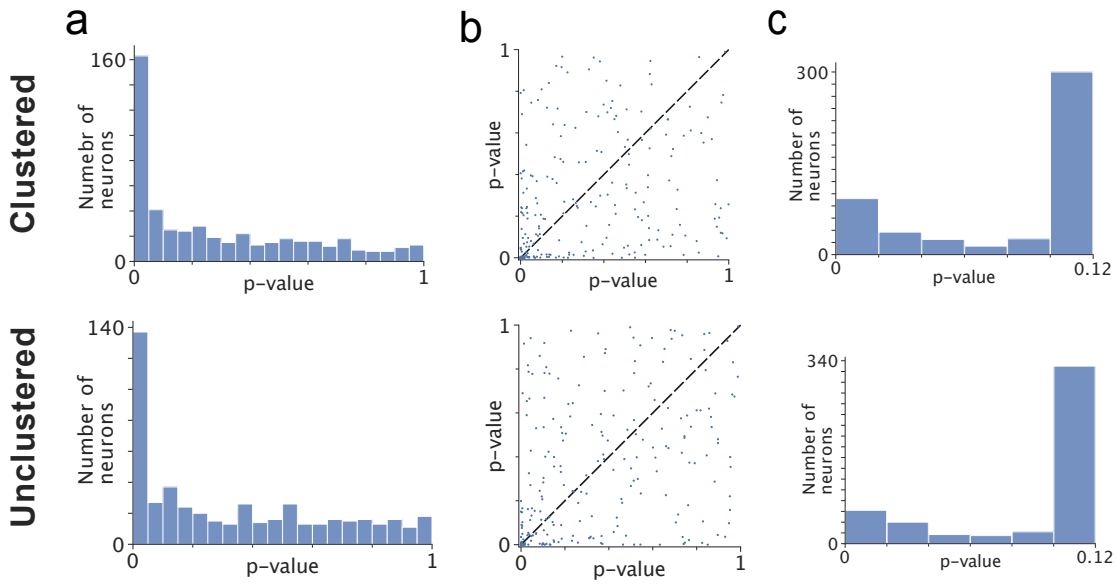


Figure 3.20: Stationarity of single neurons spike times for Stage 2 clusters (top) and correleated unclustered response distributions (bottom). (a) Correlation p-values of Stage 2 cluster single unit spike times with trial index. (b) Correlation p-values of angled cluster single unit spike times with trial index (neuron A vs neuron B). For some clusters, the cluster spike times of neither neuron were correlated with trial index, whereas for other clusters the spike times of at least one neuron in the cluster pair was correlated with trial index. (c) The Kwiatkowski-Phillips-Schmidt-Shin trend stationarity test (KPSS) was used to test whether the cluster first spike times of single neurons were stationary around a deterministic trend. For the 268 Stage 2, 102 of the corresponding 536 (2*268) single unit cluster first spike sequences were determined as not trend stationary (KPSS, $p < 0.05$). Python 'statsmodels' implementation of p-value estimation rounds all $p > 0.1$ to 0.1.

3.9.1 Stationary cluster and factor analysis criteria

Clusters were defined as stationary and sufficiently correlated such that trial-to-trial changes in latent factor such as shared pairwise excitability could underlie the covariation if they passed the following tests, and non-stationary otherwise:

1. Cluster spike times were correlated: linear regression $p < 0.005$, $r > 0.3$.

2. The single neuron cluster spike times of both neurons were not were not linearly correlated with trial index ($p > 0.05$).

3. Cluster correlation was not explained by trial index. Particularly, a linear regressor predicting spike pairs from trial index must have $r^2 < 0.05$.

4. Both neurons KPSS stationary: $p > 0.05$. The Kwiatkowski-Phillips-Schmidt-Shin trend stationarity test (KPSS) tested whether cluster first spike times of single neurons were stationary around a deterministic trend.

5. Cluster spike times of both neurons have a trial-lag-1 autocorrelation $p > 0.05$ (this requirement was relaxed for the differenced clusters but not ARIMA clusters). The trial-lag-1 autocorrelation describes the following. Consider the cluster spike times $s_{j,t}$ of neurons j=0,1 on cluster trials t=1,...,T where T is the number of cluster trials for the cluster. Trial-lag-1 autocorrelations tests the correlations between single neuron cluster spike times on trials t with spike times on trial t+1.

6. Bartlett sphericity $p > 0.05$. Bartlett sphericity tests that the correlation matrix is significantly different from the identity, and is a standard pre-test for factor analysis.

7. The factor correlation matrix had only one eigenvalue greater than 1, suggesting that a single factor underlay each cluster correlation.

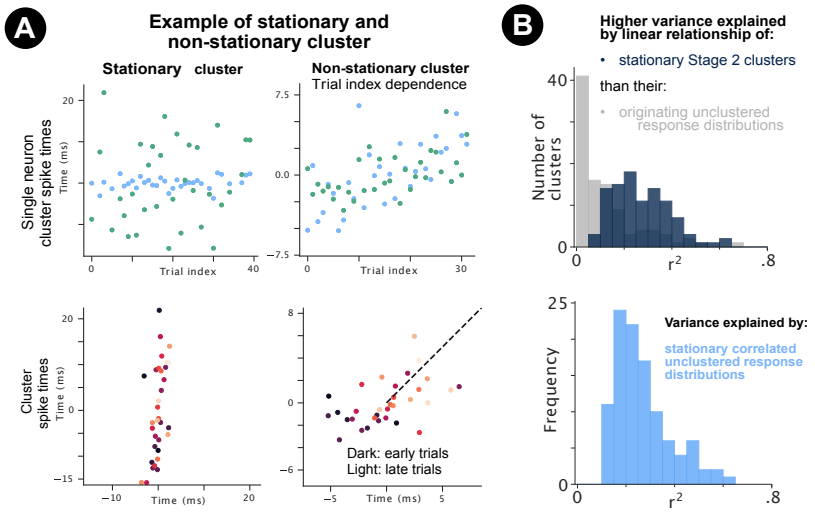
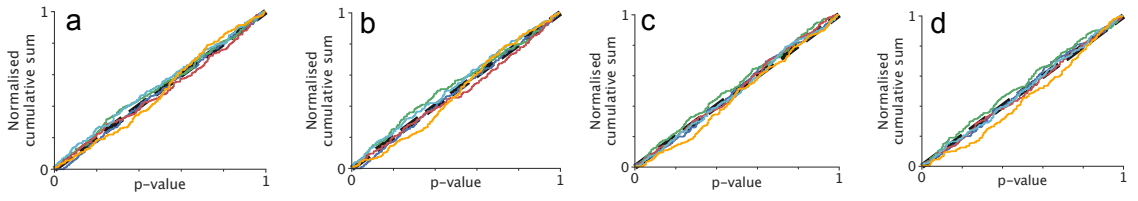


Figure 3.21: Examples of stationary and non-stationary clusters, and explained variance. (a) Example of stationary (left) and non-stationary cluster (right) determined by stationarity criteria (2)-(5) (see Methods). Top row shows single neuron cluster spike times for each neuron (green and blue) plotted against cluster trial index. Bottom shows the cluster spike times of neuron A and B plotted against each other. Non-stationary cluster shows a clear trial dependence. (b) Top: Histogram of variance explained (r^2) by the spike time linear relationship of each stationary Stage 2 cluster (dark blue) and the linear relationship of the unclustered response distributions from which stationary Stage 2 clusters were extracted (grey). Bottom: r^2 values shown for stationary unclustered response distributions

3.10 Results IV: Stationary clusters and state-dependent response distributions

3.10.1 Stationary clusters

For Stage 2 clusters and correlated unclustered response distributions, spike times were either not modulated by adaptation, or the cluster spike times of one or both of the neurons were modulated by adaptation (Fig. 3.20). Clusters or response distributions were defined as 'stationary' if the spike times of both neurons were not modulated by adaptation, and 'non-stationary' otherwise (stationarity criteria in Methods 3.9). Fig. 3.21a shows example stationary and non-stationary clusters. The non-stationary cluster shows a clear trial dependence. 40.5% (100/247) of Stage 2 clusters and 43.2% (105/243) of correlated unclustered response distributions were determined as stationary. Fig. 3.21b shows the proportion of variance explained by the correlations of these stationary Stage 2 clusters and stationary correlated unclustered response distributions.



Autocorrelation.
Fisher Method (df=400,N=200).

Trial-lag	χ^2	p-value
1	347.1	0.973
2	417.0	0.269
3	382.3	0.729
4	447.9	0.049
5	367.1	0.880

Partial autocorrelation.
Fisher Method (df=400,N=200).

Trial-lag	χ^2	p-value
2	437.8	0.093
3	369.7	0.859
4	441.2	0.076
5	351.3	0.962

Cross-correlation.
Fisher Method (df=400,N=200).

Trial-lag	χ^2	p-value
1	389.5	0.638
2	423.3	0.203
3	386.1	0.683
4	376.3	0.797
5	329.0	0.996

Partial cross-correlation.
Fisher Method (df=400,N=200).

Trial-lag	χ^2	p-value
2	427.9	0.168
3	374.4	0.816
4	387.2	0.667
5	315.2	>0.999

Figure 3.22: Stationary Stage 2 cluster single-unit single-stimulus first spike (a) autocorrelations, (b) partial autocorrelations, (c) cross-correlation and (d) partial cross-correlation. Autocorrelation, partial autocorrelation, cross-correlation and partial cross-correlation analysis allows quantification of autocorrelative and cross-correlative structure in the trial-to-trial sequence of first spike times and first spike pairs for the 117 stationary Stage 2 clusters. This figure considers the first spike times $s_{j,t}$ of neurons j=0,1 on cluster trials t=1,...,T where T is the number of cluster trials for a single cluster. Lag-n autocorrelations and partial autocorrelations test the correlations and partial autocorrelations between single neuron cluster spike times on trials t with spike times on trial t+n. Plots (a) and (b) show the normalized cumulative sums of lag 1-5 autocorrelation and lag (2-5) partial autocorrelation p-values for 214 (=2 * 117 stationary Stage 2 clusters) single neuron single cluster spike time sequences. Lag-n cross-correlations and partial cross-correlations were similarly tested between neuron pair cluster spike times on trials t with spike times on trial t+n. Autocorrelations, partial autocorrelations, cross-correlations and partial cross-correlations are at chance-levels for lags 1-5. This is demonstrated by the fact that the cumulative sums lie on y=x, demonstrating uniform distributions of p-values. Accompanying tables show results (χ^2 statistics and p-values) of Fisher-method application to the p-values produced for each lag, which confirms that autocorrelations, partial autocorrelations, cross-correlations and partial cross-correlations are not above chance.

For stationary Stage 2 clusters, cluster spike times of the first neuron on cluster trials t, were not correlated with the cluster spike times of the second neuron on cluster trials $t+1$ (i.e. trial-lag-1 cross-correlation; Fig.3.22). The same was true for stationary correlated unclustered response distributions. This showed that removal of single trial correlations removed the correlations of the stationary clusters and unclustered response distributions, which would not be the case if general adaptative trends underlied the stationary correlations. Moreover, chance-level autocorrelations and cross-correlations for multiple trial-lags confirmed that the stationary correlations were not affected by adaptative changes or autocorrelative structure (Fig.3.22), providing an important control test.

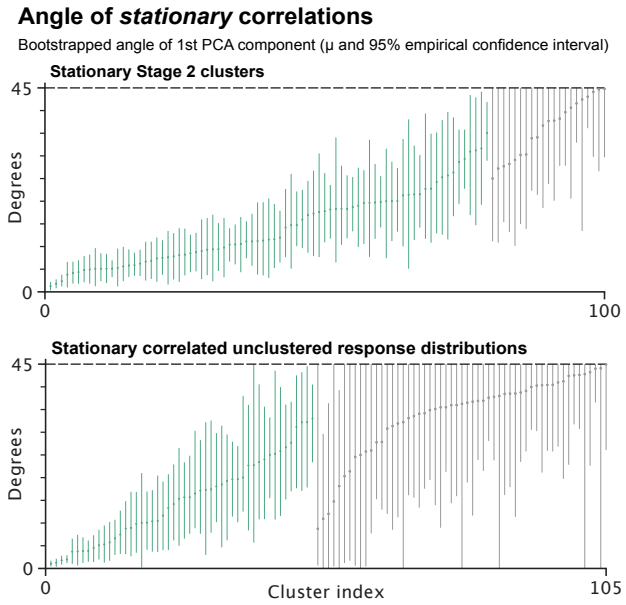


Figure 3.23: Stationary clusters. Predicted θ_{45} angles with 95% empirical confidence intervals for stationary Stage 2 clusters (top) and stationary correlated unclustered response distributions. Angles significantly different from 0° ($p < 0.025$) and 45° ($p < 0.025$) coloured green. Grey otherwise.

77.0% (77/100) of stationary Stage 2 clusters and 43.8% (46/105) of stationary correlated unclustered response distributions had θ_{45} angles between and significantly different from 0° and 45° (Fig. X) showing that relative first spike time differences are warped by a non-adaptative latent factor such as shared trial-to-trial changes in the excitability of neuronal pairs. It is important to note that the Henze-Zirkler null hypothesis of normality was rejected ($p < 0.05$) for 75.0% (75/100) of stationary Stage 2 clusters and 76.2% (80/105) of stationary correlated unclustered response distributions.

3.10.2 Fixed-state and state-conditioned relative difference response distributions

This section compares the potential encoding capacity of the characterised clusters under three alternative relative spike time difference coding regimes named Relative Difference Regimes (RDR) 1-3. It is first necessary to consider how downstream neurons discriminate between stimuli. If two sensory stimuli elicit a behaviourally different response, downstream neurons must be able to respond differently to samples from the response distribution of either stimulus. This is possible if the response distributions of the two stimuli are non-overlapping in dimensions that downstream

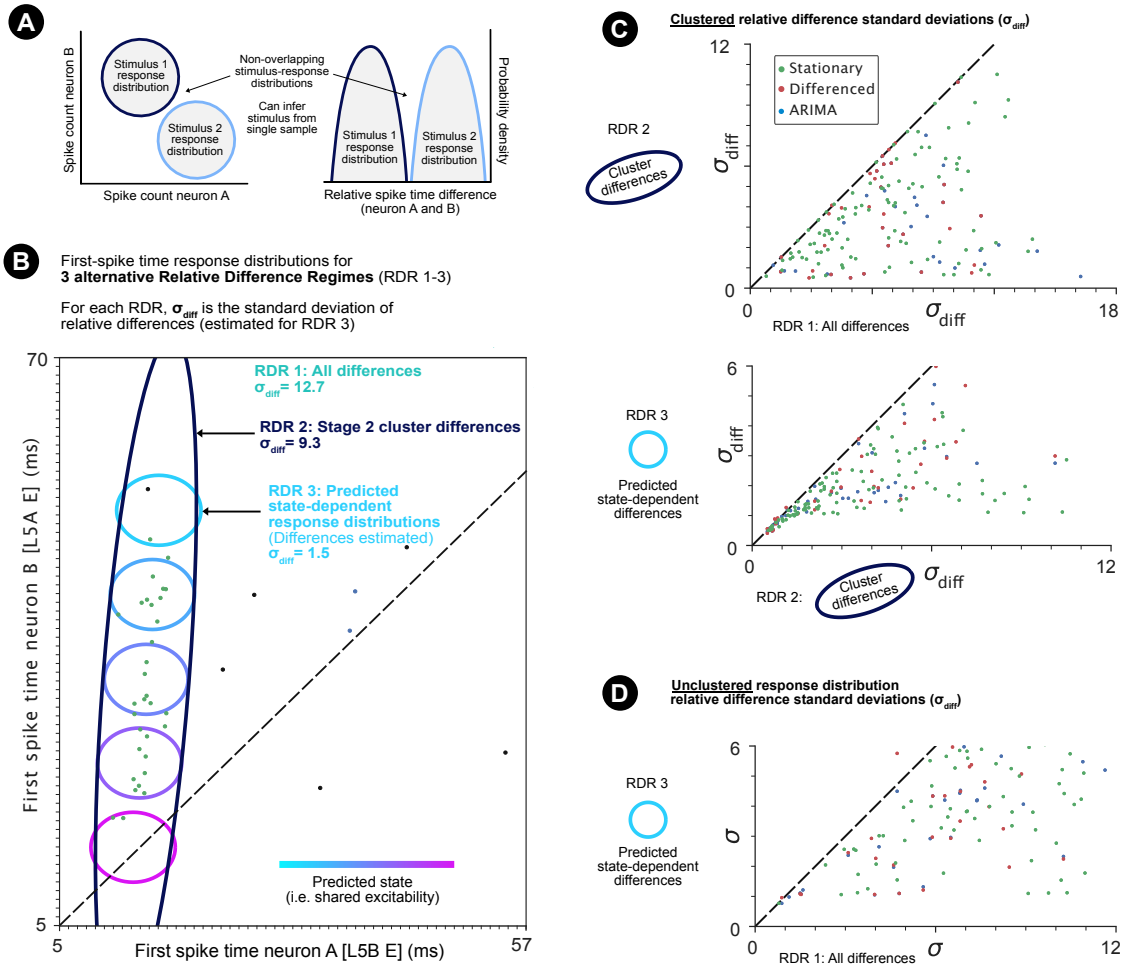


Figure 3.24: State-dependent coding predictions and relative difference regimes. (a) Illustration of two non-overlapping stimulus-response distributions for (left) spike counts and (right) relative spike time differences. (b) Sampled first spike response distribution for two neurons (L5B E, L5A E). Point colours illustrate DBSCAN cluster (epsilon = 5.0). Dark blue angled ellipse illustrates 4σ boundary of Stage 2 ellipse (i.e. the area for which the standard deviation of relative differences is calculated under RDR 2). Coloured circles illustrate the 4σ boundary of the state-dependent response distribution predicted by factor analysis for single state values ($< 0, 0, > 0$). The standard deviation of the predicted state-conditioned first spike relative difference distributions (for RDR 3) is calculated based on the size of this area (Methods). (c) The standard deviation of differences of RDR 1 vs RDR 2 (top) and RDR 2 vs RDR 3 (bottom) for criteria fulfilling stationary, differenced and ARIMA modelled clusters. (d) RDR 2 vs RDR 3 (bottom) for criteria fulfilling stationary, differenced and ARIMA modelled unclustered response distributions.

neurons are sensitive to, such as spike counts or relative spike time differences (Fig. 3.24a). The smaller the region of stimulus-response distribution space that single

stimulus-response distributions occupy, the higher the capacity for encoding the presence of different stimuli reliably.

Two possible relative difference coding regimes are firstly considered under the classical assumption of fixed-state-decoding (i.e. decoding neurons are not modulated by processes of adaptation or shared trial-to-trial changes in the excitability of neuronal pairs). Relative Difference Regime 1 (RDR 1) describes the assumption that pairs of neurons encode information through a fixed relative spike time difference. The standard deviation of relative differences over all spike pairs ($RDR1_\sigma_{diff}$) offers a quantification and prediction of the potential encoding capacity for relative differences under RDR 1. The assumption, however, does not capture the quantified structure of the observed response distributions (Fig. 3.24b; Vid. 1).

Observation of first spike stimulus-response distributions (Fig. 3.24b; Vid. 1) shows that clusters correspond to regions of high response probability, which better estimate the underlying stimulus-response distributions. Relative Difference Regime 2 (RDR 2) describes the notion that correlated clusters may represent trials on which the neuron pair's relative first spike time difference reliably contributes to the representation of the stimulus. Under the classical assumption of fixed-state-decoders, the standard deviation of relative differences ($RDR2_\sigma_{diff}$) calculated for samples in a stationary Stage 2 cluster quantifies the size of the relative difference response distribution and the potential encoding capacity for relative differences under RDR 2. As stationary clusters are not affected by adaptation, all cluster samples can be used in estimating $RDR2_\sigma_{diff}$. The smaller first spike pair relative difference standard deviations (σ_{diff}) under RDR 2 compared to RDR 1 demonstrates an improved potential coding capacity for RDR 2 (Fig. 3.24b,c; Vid. 1).

3.10.3 Estimation of state-conditioned response distributions

RDR 3 considers the hypothesis that the latent factor co-modulating the spike times of neuron pairs corresponds with trial-to-trial changes in excitability that are shared across many neurons in the barrel cortex. Under this assumption, RDR 3 considers that downstream decoding neurons may be modulated by the same shared trial-to-trial changes in the excitability that modulate cluster correlations (i.e. state-dependent-decoding). For example, a higher-level of excitability (i.e. continuous cortical state) on a single trial would decrease the latency of both neurons. Spike times of both neurons would then come from a smaller state-conditioned stimulus-response distribution. If the cluster correlation is non-unitary, as in Fig. 3.24b, different state values would correspond with different distributions of relative differences. A

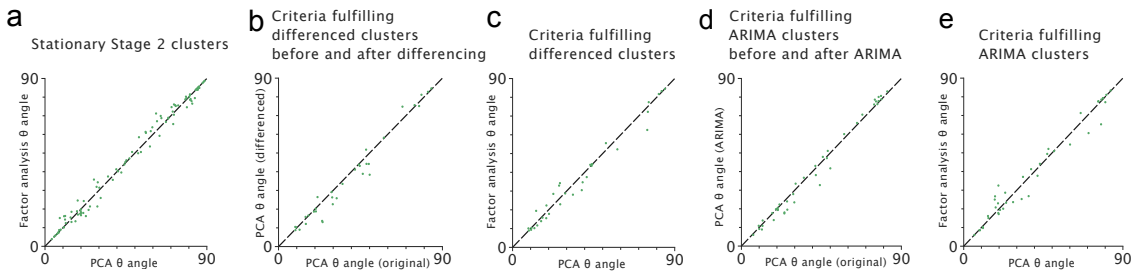


Figure 3.25: Cluster angle comparisons. (a) PCA angles vs factor analysis angles for the criteria fulfilling clusters. (b) Original PCA angles vs differenced PCA angles for the criteria fulfilling differenced clusters. (c) PCA angles vs factor analysis angles for the criteria fulfilling differenced clusters. (d) Original PCA angles vs ARIMA PCA angles for the criteria fulfilling ARIMA clusters. (e) PCA angles vs factor analysis angles for the criteria fulfilling ARIMA clusters.

downstream neuron modulated by the same trial-to-trial changes in excitability (i.e. continuous cortical state) could respond in a state-dependent manner. Such (smaller) state-conditioned first spike difference stimulus-response distributions would provide a higher coding capacity under state-dependent decoding.

To first estimate the size of state-dependent first spike response distributions for each stationary Stage 2 cluster, the following linear single factor model was fit to each zero-centred stationary Stage 2 cluster: $s_i = \lambda_i \eta + \sigma_i$ for $i \in [0, 1]$. η represents the value of the continuous underlying state (i.e. shared level of excitability) with most values between [-3, 3], λ_i is the gradient at which the spike time of neuron i varies w.r.t. the state η, and σ_i is the predicted standard deviation of the spike times s_i given η. Fig. 3.24b illustrates state-dependent first spike response distributions for example values of $\eta = [-2.0, -1.0, 0.0, 1.0, 2.0]$.

The estimated size of state-conditioned distributions of decodable relative spike time differences (RDR 3) was then calculated from the factor analysis model as $\sigma_{diff} = \sqrt{\sigma_0^2 + \sigma_1^2}$ where σ_i is the state-dependent standard deviation of neuron i predicted by factor analysis. As the factor analysis model assumes that noise for each neuron given η is normally distributed and independent between neurons, the state-dependent spike time relative difference standard deviation predicted by the model is equal to $\sigma_{diff} = \sqrt{\sigma_0^2 + \sigma_1^2}$. The standard deviations of state-conditioned distributions of relative differences ($RDR3_\sigma_{diff}$) offers an improved potential coding capacity over RDR 2 (Fig. 3.24b,c; Vid. 1) and thus also over RDR 1, but would require confirmation of modulation by a population-wide state, which varies from trial-to-trial. Fig. 3.24d shows the predicted improvement of $RDR3_\sigma_{diff}$ over RDR 2 for the correlated unclustered response distributions.

θ and θ_{45} angles of factor analysis mean lines were estimated using the same technique as for PCA angles. Factor analysis θ angles were similar to the PCA θ angles for stationary, differenced and ARIMA modelled clusters (Fig. 3.25a,c,e), suggesting that the factor analysis model provided a good fit of the data.

3.11 Methods V: Analysis of non-stationary clusters

To test whether a positive non-adaptative latent factor (which could be trial-to-trial changes in shared excitability) could also modulate the remaining non-stationary Stage 2 clusters and non-stationary correlated unclustered response distributions, and to estimate the size of underlying relative difference response distributions at a single point in time under RDR 2 and 3, single neuron adaptative trends of non-stationary clustered and unclustered spike times were removed using time series modelling. Two techniques with different advantages were applied to non-stationary spike times (Fig. 3.26). For the clustered case, adaptative trends were not removed from spike times before clustering so as to ensure that clusters and spike time correlations were not introduced by such pre-processing. Simple parametric models were not used to remove trends because non-stationarities in the trial-to-trial sequence of single neuron and cluster spike time did not always follow standard adaptative/non-stationary trends (Fig. 3.27). In these cases, standard parameteric models would not remove the non-stationarity and would risk introducing correlations where none were present.

3.11.1 Differencing

Firstly, in order to capture spike time correlations caused by a non-adaptative latent process (such as trial-to-trial changes in shared excitability) and not by non-stationary processes, first order 'differencing' [195] was applied to remove adaptative trends of single neuron cluster spike times determined as non-stationary by stationarity criteria (2)-(5). Differencing a series of first spike times subtracts the first spike time on trial t-1 from the first spike time trial t (for trials that contributed samples to the cluster or unclustered response distribution). Differencing reduces the number of samples by 1. Differencing was applied to single neuron cluster spike times if any of the following stationarity criteria were not met for the cluster spike times of either neuron: (1) Trial index vs spike time linear regression $p > 0.05$, (2) KPSS stationarity $p > 0.05$, (3) Trial-lag-1 autocorrelation $p > 0.05$. For clusters for which the spike times of only one neuron was differenced, the first sample of the undifferenced neuron was removed, so that the number of samples was equal for both neurons. In this case a spike time for an undifferenced trial t was matched with the difference between the spike times on trial t and trial t - 1.

As differencing increases the standard deviation of the time series, the differenced time series were rescaled appropriately. Particularly, differencing increases the stan-

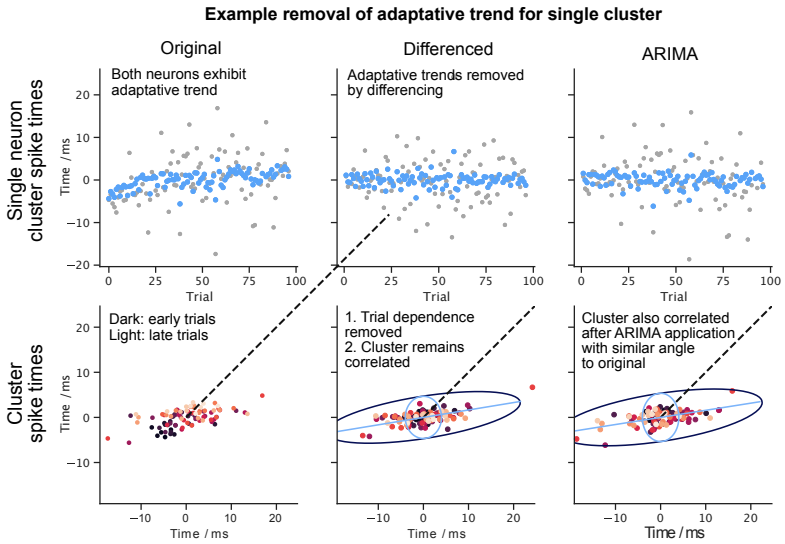


Figure 3.26: Time series modelling techniques. (a) For an example cluster, the original spike times do not fulfil stationarity criteria (2)-(5) (Methods) but do after the application of differencing and ARIMA modelling. The factor analysis predicted state-dependent response distributions are therefore presented for the cluster after differencing and ARIMA modelling (4σ ellipse; bottom centre and bottom right). It is clear that the adaptative trends have been removed by differening and ARIMA but the clusters remain correlated with similar angles to the original cluster. (b) Predicted angles and 95% empirical confidence intervals of differenced and ARIMA modelled criteria fulfilling Stage 2 clusters. Angles significantly different from 0° and 45° coloured green. Grey otherwise.

dard deviation of a time series $S = (s_1, ..., s_t)$ by approximately $\sqrt{2}$. This is shown by considering differencing as approximate to negating two independent random variables $S_{(1,t-1)} = (s_1, ..., s_{t-1})$ and $S_{(2,t)} = (s_2, ..., s_t)$, where t is the number of samples in the original time series. As the standard deviation of $S_{(0,t-1)}$ and $S_{(1,t)}$ are approximate to the standard deviation of the original series S, the standard deviation of the differences is on the order of $\sigma_{differenced} \approx \sqrt{\sigma^2_{(1,t-1)} + \sigma^2_{(2,t)}} \approx \sqrt{2\sigma^2} = \sqrt{2}\sigma$ where σ is the standard deviation of the original time series. Differenced time series were therefore divided by $\sqrt{2}$ to return the differenced time series to the scale of the original time series. This is shown to be a reasonable operation below. The assessment of 'raw' differences allows a direct assessment of cluster angles after trend removal, without further application of time series methods.

Negative autocorrelations known and shown (Fig. 3.28) to be introduced by differencing were ignored, because the resulting cluster spike times allowed direct assessment of cluster angles after trend removal without further processing.

Differencing was applied to the spike times of clusters rather than to all first spike

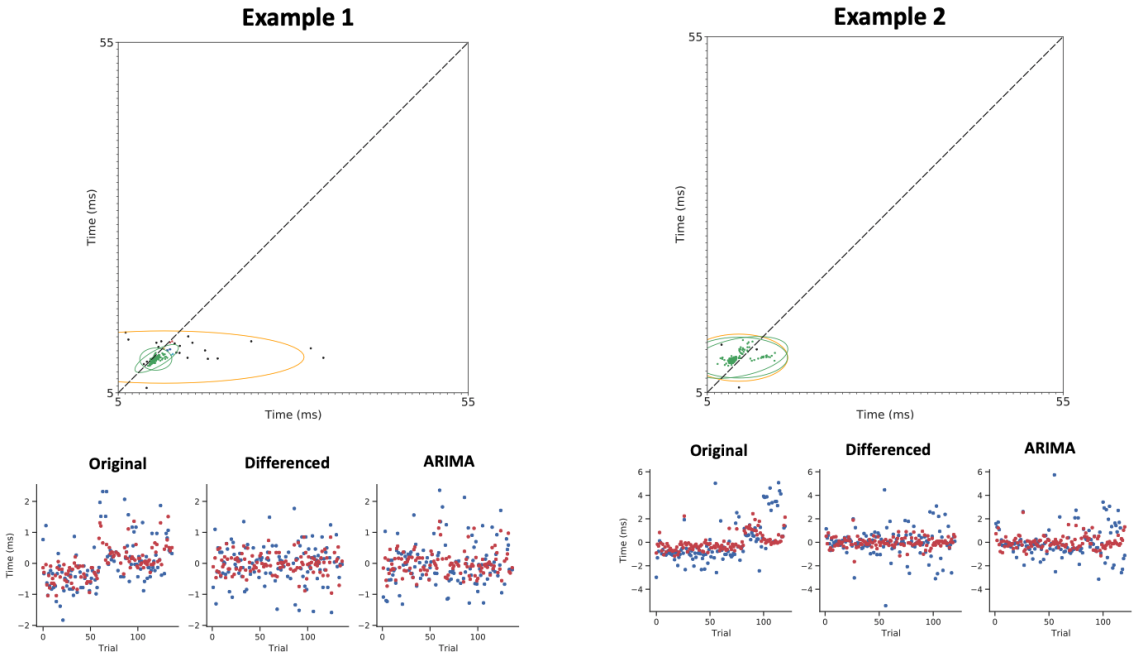


Figure 3.27: Two examples (left and right) of discontinuous and non-monotonic non-stationarity, highlighting the need for differencing over parametric models for the removal of non-stationary trends when testing for non-adaptative pairwise correlations. Top row shows two first spike pair response distributions in which Stage 2 clusters were detected by the main clustering algorithm. Bottom row shows cluster spike times of the two corresponding neurons (red and blue) vs trial index for the original (left), differenced (center) and ARIMA (right) cases. For both of the examples, the original spike times of both neurons show a discontinuity in their relationship with trial index (at around trials 70 and 90 respectively). At these points of discontinuity the latencies of both neurons in both examples become slower, before gradually becoming faster. For the removal of non-stationarities in order to test for pairwise correlations, the non-stationarities of examples such as these could not be modelled with typically used parametric models typically used for adaptative trend removal. Moreover, use of classic parametric models would risk introducing adaptation-independent correlations between the spike times of the two neurons where none existed previously. Alternatively, after the application of differencing and ARIMA, the cluster spike times show no relationship with trial index, and were deemed stationary by the stationarity criteria. This highlights an advantage of differencing based methods over parametric models, for the purposes of this paper.

times of a single neuron in response to a single stimulus. This was to ensure that differences were calculated between spike times of the same trials for both neurons in a pair. That is, the difference between spike times on trial X and spike times on trial Y were compared for neuron pairs. Differencing was not alternatively applied to all samples of a first spike stimulus-response distribution, as this could give rise to corre-

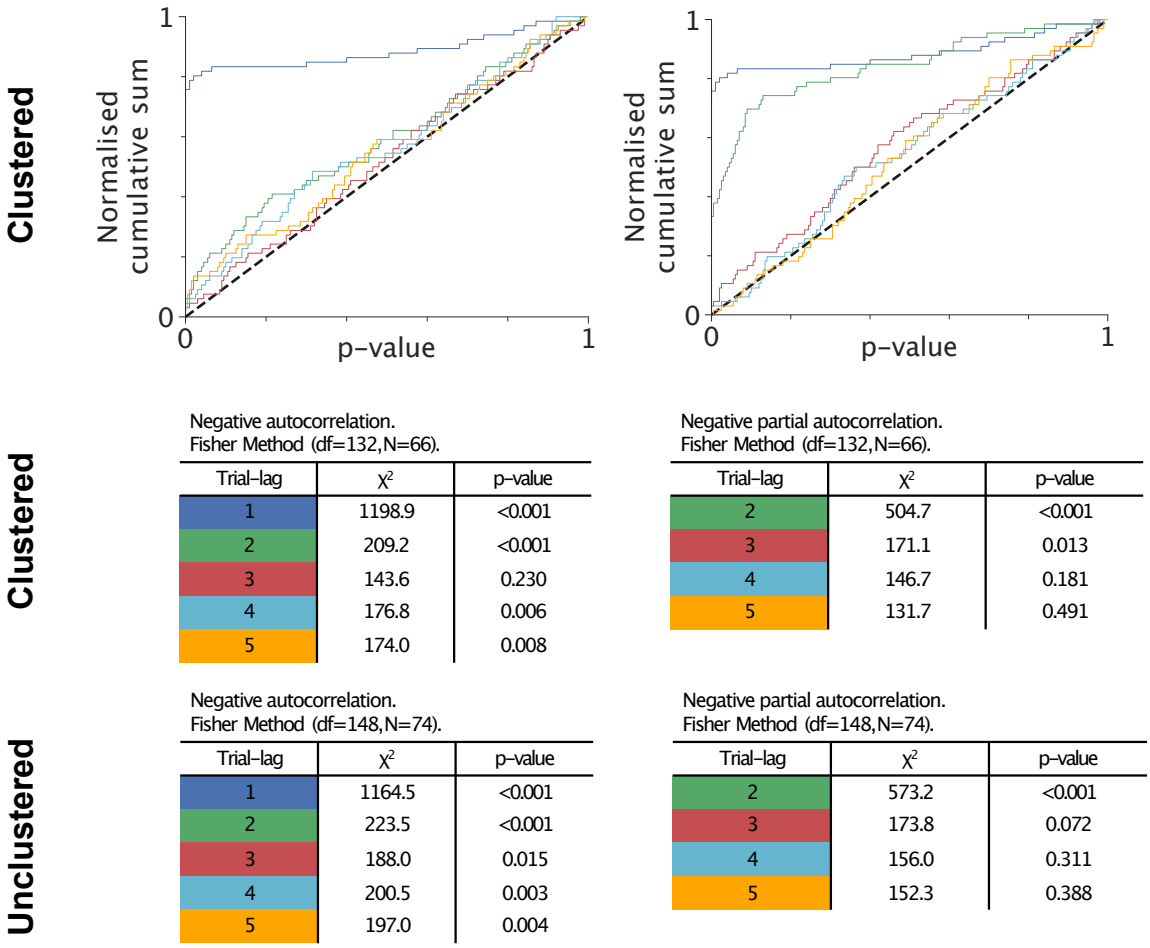


Figure 3.28: Negative autocorrelations (left) and negative partial autocorrelations (right) introduced by differencing. Plots for criteria fulfilling differenced clusters (top, centre) and unclustered (bottom) (with the autocorrelation criterion relaxed). As with previous autocorrelation plots (i.e. Fig. 3.22) but p-values instead represent the probability that single neuron spike times are negatively correlated for the criteria fulfilling differenced Stage 2 clusters and unclustered response distributions. Cumulative sum plots and accompanying tables show that spike times are negatively autocorrelated and negatively partially autocorrelated above chance for several lags, as expected by differencing.

lations caused by switching between clusters rather than the excitability fluctuations of interest.

3.11.2 ARIMA modelling

The second technique fitted autoregressive (AR) and/or moving average (MA) models appropriately to the differenced time series (i.e. 'ARIMA' modelling [118, 28]), in order to remove the autocorrelations introduced by differencing and to return the

time series to the original scale without rescaling. The 'I' in ARIMA refers to the differencing applied previously in the first step. The AR model component regresses the values of the time series on values of the time series on previous trials. The MA model component regresses the value of a time series on a moving average of the previous values. The order of the AR and MA components determines the number of previous values that are used. An order of 0 means that no AR or MA component was applied.

The orders of autoregressive (AR) and moving average (MA) components were determined using a simplified and automated implementation of the Box-Jenkins method, which is classically applied manually based on observation of autocorrelations and partial autocorrelations. This implementation was based on the description of the Box-Jenkins method found at [196], and was particularly designed to account for the negative autocorrelations introduced by differencing and other autocorrelations.

Firstly, for determination of the MA order, if the trial-lag-1 autocorrelation was negative and significant, the series is likely to be slightly overdifferenced. The Box-Jenkins method checks for a sudden cut off in the autocorrelation plot. The lag at which there is a cut-off is used as the MA order. To automate this, the algorithm checked if the trial-lag-1 autocorrelation was negative ($r < -0.1$). If so, the p-values of trial-lagged autocorrelations were tested (trial-lags 1-5). If there was a sudden increase in p-value from below 0.05, from one lag to another, by more than 0.15, then the value of the last significant autocorrelation was used as the MA order.

For determination of the AR order, if the trial-lag-1 autocorrelation is positive ($r > 0.1$) and significant, the series is slightly underdifferenced, and the Box-Jenkins method checks similarly for a sudden cut-off in the partial-autocorrelation plot to determine the AR term.

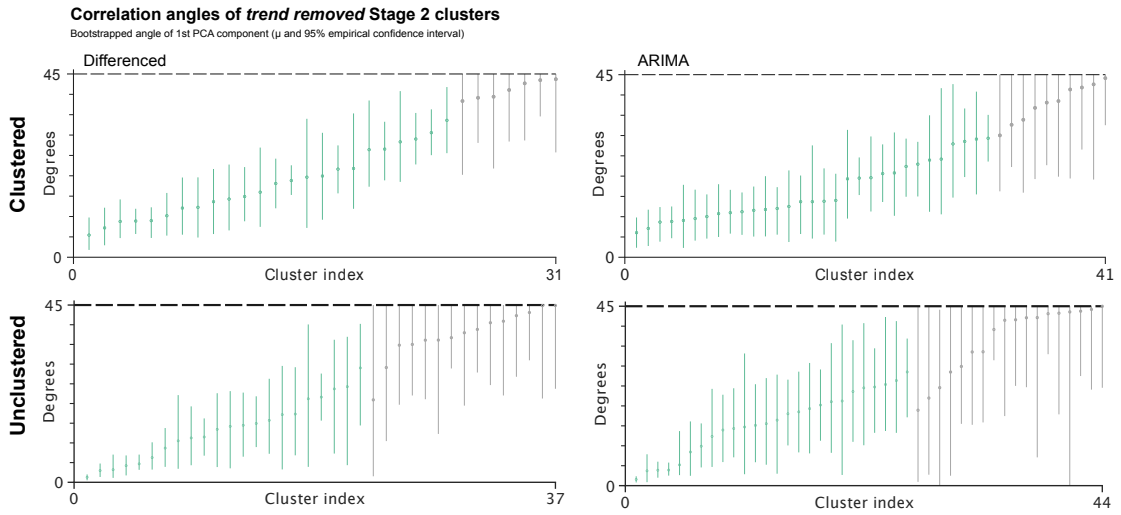


Figure 3.29: Predicted angles and 95% empirical confidence intervals of differenced and ARIMA modelled criteria fulfilling Stage 2 clusters and unclustered response distributions. Angles significantly different from 0° and 45° coloured green. Grey otherwise.

3.12 Results V: Non-stationary clusters are also modulated by shared trial-to-trial changes in the excitability of neuronal pairs

76.2% (112/147) of non-stationary clusters and 80.4% (111/138) of non-stationary correlated unclustered response distributions had differencing applied to at least one neuron. After differencing, 29.5% (33/112) and 33.3% (37/111) of these clusters and unclustered response distributions fulfilled the stationarity and correlation criteria, of which 72.7% (24/33) and 54.1% (20/37) had θ_{45} angles between and significantly different from 0° ($p < 0.025$) and less than 45° ($p < 0.025$; Fig. 3.29). It is important to note that the Henze-Zirkler null hypothesis of normality was rejected ($p < 0.05$) for 21.2% (7/33) and 62.2% (23/37) of the criteria fulfilling differenced Stage 2 clusters and unclustered response distributions.

96.4% (108/112) and 100.0% (111/111) of these differenced clusters and unclustered response distributions had an AR and/or MA model applied to at least one neuron. The majority of fitted AR or MA models were of the MA type with a trial lag of 1 (Fig. 3.30), as expected for removal the negative autocorrelations introduced by differencing. That is, the algorithm automatically selects to remove the dependence of a spike time on the previous spiking trial's spike time that is introduced by the differencing operation.

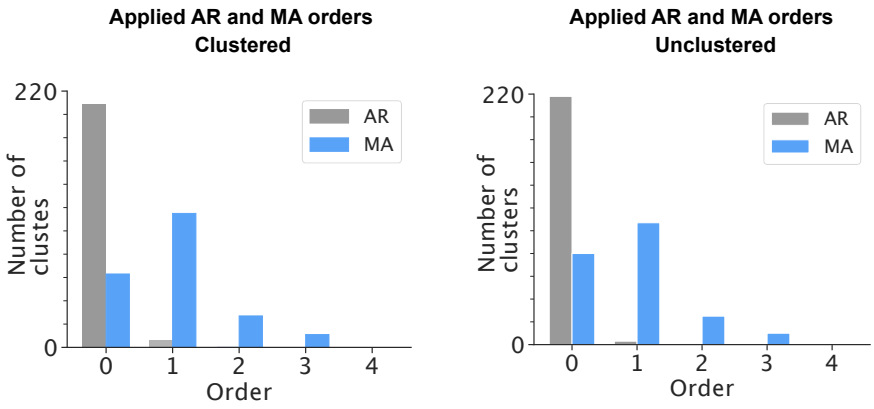
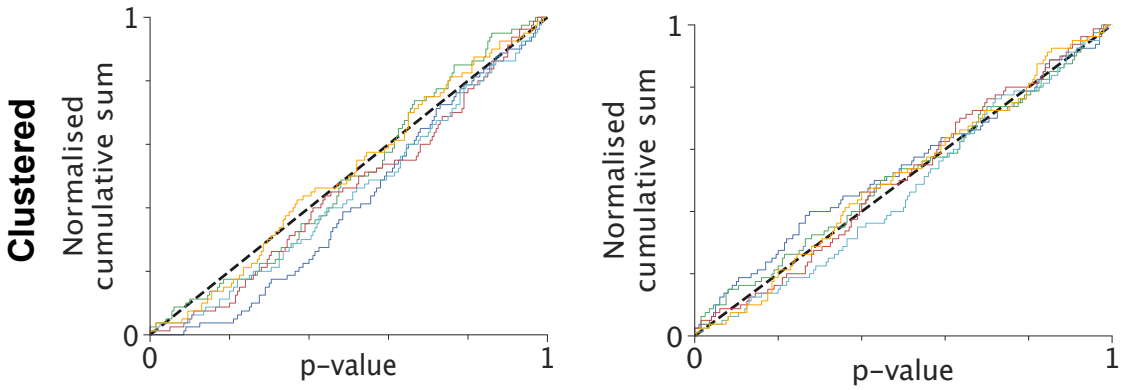


Figure 3.30: Order of AR and MA models fit to the single neuron spike times of differenced clusters and unclustered response distributions.

37.0% (40/108) and 39.6% (44/111) of ARIMA modelled clusters and unclustered response distributions fulfilled the criteria, of which 72.5% (29/40) and 56.8% (25/44) had θ_{45} angles between and significantly different from 0° ($p < 0.025$) and less than 45° ($p < 0.025$; Fig. 3.29). The Henze-Zirkler null hypothesis of normality was rejected ($p < 0.05$) for 21.2% (7/33) and 62.2% (23/37) of the criteria fulfilling ARIMA modelled Stage 2 clusters and unclustered response distributions. Autocorrelations and cross correlations were below chance levels for the criteria fulfilling ARIMA modelled clusters (Fig. 3.31), showing that AR and MA modelling successfully removed the negative autocorrelations introduced by differencing.

These results demonstrated that the spike times and relative differences of a proportion of non-stationary Stage 2 clusters and unclustered response distributions are also modulated by trial-to-trial changes in a non-adaptative latent factor, which could be shared non-adaptative trial-to-trial changes in the excitability of neuronal pairs. Fig. 3.32 show the variance explained by the correlations after trend removal (differencing and ARIMA). The angles of the differenced and ARIMA modelled spike times corresponded with the correlation angles before the application of differencing and ARIMA modelling (Fig. 3.25). This supports co-modulation by trial-to-trial changes in a non-adaptative latent factor and shows that the non-45° correlation angles are not artifacts of time series analysis.

The trend removed clusters allowed all cluster samples to be used in the estimation of the underlying cluster and corresponding distribution of relative differences (RDR 2) at a single point in time. Factor analysis was applied to the trend removed clusters to estimate state-dependent distributions of relative differences at a single point in time (RDR 3). Again, RDR 2 gave an improved potential encoding capacity over



Clustered

Autocorrelation.
Fisher Method (df=160,N=80).

Trial-lag	χ^2	p-value
1	100.8	>0.999
2	166.8	0.340
3	126.0	0.978
4	135.7	0.918
5	150.0	0.704

Cross-correlation.
Fisher Method (df=160,N=80).

Trial-lag	χ^2	p-value
1	186.1	0.077
2	177.5	0.163
3	163.4	0.410
4	137.3	0.903
5	151.1	0.681

Unclustered

Autocorrelation.
Fisher Method (df=176,N=88).

Trial-lag	χ^2	p-value
1	129.7	0.996
2	150.0	0.923
3	172.1	0.570
4	144.8	0.959
5	164.8	0.717

Cross-correlation.
Fisher Method (df=176,N=88).

Trial-lag	χ^2	p-value
1	181.9	0.365
2	169.0	0.634
3	171.1	0.589
4	160.3	0.796
5	192.7	0.185

Figure 3.31: Autocorrelations (left) and cross-correlations (right) at chance-levels for criteria fulfilling ARIMA modelled Stage 2 clusters (top, centre) and unclustered response distributions (bottom). As with Fig. 3.22 but for criteria fulfilling ARIMA modelled clusters and unclustered response distributions. Autocorrelations and cross-correlations are also at chance levels confirming that ARIMA modelling successfully removed autocorrelations introduced by differencing and that there is no additional cross-correlative structure between the neurons in ARIMA modelled clusters and unclustered response distributions.

RDR 1, whilst RDR 3 gave an improved potential encoding capacity over both RDR 1 and 2 (Fig. 3.24c,d).

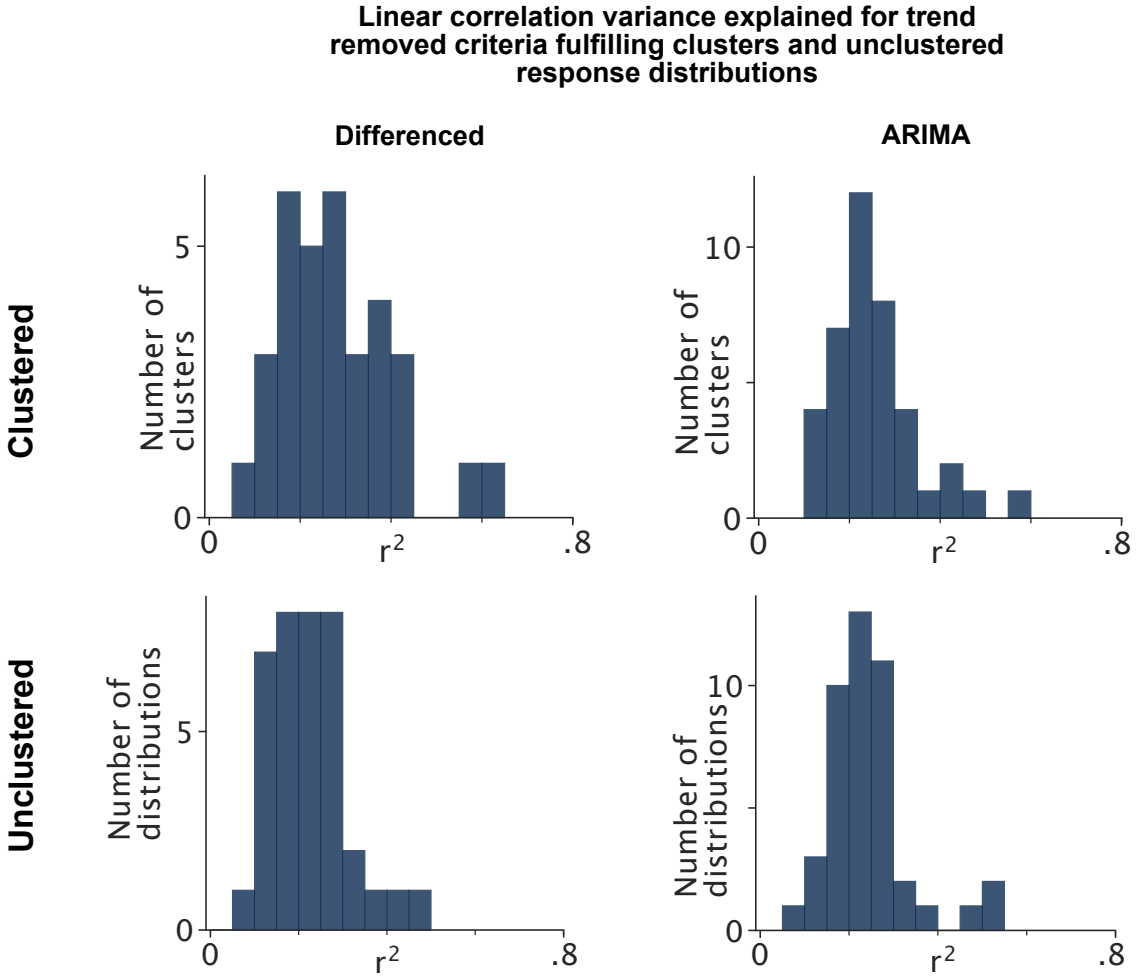


Figure 3.32: Linear regression r^2 values for criteria fulfilling differenced (left) and ARIMA modelled (right) clusters (top) and unclustered response distributions (bottom).

3.13 Results VI: Frequency dependence and spatial distribution of correlated clusters

Stationary correlations were found at all stimulation frequencies, but mostly at lower frequencies (Fig. 3.33), showing that lower frequencies are better suited for stimulus response distribution characterisation as adaptation occurs to a lesser extent. ARIMA modelled clusters were found for all frequencies $\leq 8Hz$ (Fig. 3.33), demonstrating that the techniques can characterise noise correlations caused by a non-adaptative latent process in the presence of adaptation, which occurs both at lower frequencies suited for response distribution characterisation ($0.0 - 0.2Hz$), and behavioural whisking frequencies ($1 - 20Hz$ [41]).

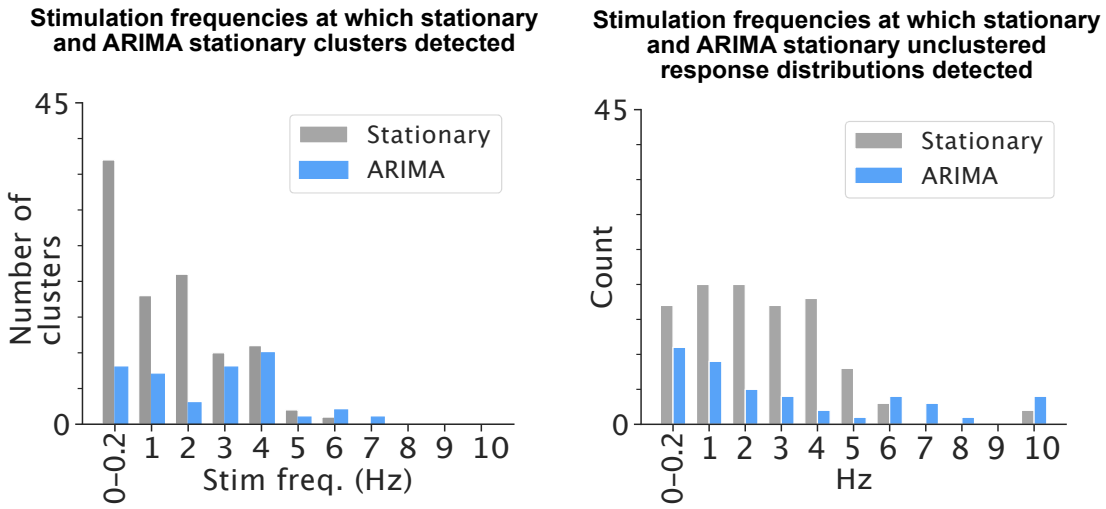


Figure 3.33: Histogram showing number of stationary Stage 2 correlations and criteria fulfilling ARIMA modelled correlations (by stimulation frequency) for clusters (left) and unclustered response distributions (right).

Spike time correlations (stationary and ARIMA trend removed) were found for E-E, I-I and E-I neuron pairs separated over multiple (up to 4) cortical columns and layers (Fig. 3.33), demonstrating positive co-modulation by trial-to-trial changes in a non-adaptative latent factor for neuronal pairs over a large spatial distance. This introduces the possibility that spatially broad population-wide changes in neuronal excitability could modulate decoding neurons and allow them to respond in an state-dependent manner. The most correlations were found between E-type neurons in L5A and L5B/6 of the principal and neighbouring columns, although differences in the number of spike sorted neurons extracted by layer and E/I type must be considered (Fig. 3.4). Low numbers of correlated clusters were found for neuron pairs within the same group, which could be a consequence of the local spike overlapping problem (i.e. spike detection difficulties when two or more neurons fire in close temporal and spatial proximity). For 32 (clustered) and 25 (unclustered) neuron pairs, at least one neuron lied in a septum (i.e. between barrels).

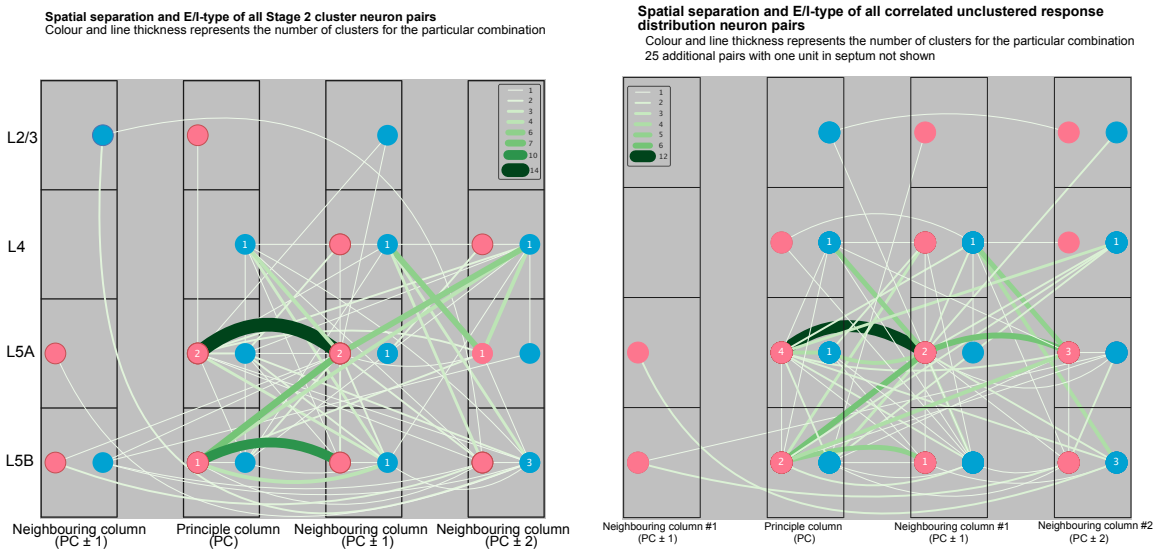


Figure 3.34: Neuron pair spatial separation (unclustered). Neuron pair spatial separation for criteria fulfilling correlated clusters (left) and unclustered response distributions (right) (both stationary and ARIMA modelled). Data from all experimental sessions. Line colour and thickness indicates the number of criteria fulfilling correlations (clusters or unclustered response distributions) found for pairs of neurons in different layers and barrels and of different E/I-type. Red and blue correspond to E and I neuron types respectively. For example, the thick dark green line illustrates that there were 14 criteria fulfilling correlated clusters, for which both neurons were in L5B, 1 neuron was in the principal column and 1 was in a neighbouring column. The numbers in the circles indicate the number of pairs of the same E/I type, layer and barrel. The left most column is shown to illustrate pairs for which neurons are on opposite sides of the principal column. Counts are pooled for stationary clusters and non-stationary clusters made stationary by ARIMA modelling. 32 and 25 neuron pairs were not to added to each plot because at least one of the neurons in the pair lied in a septum (i.e. between barrels).

3.14 Discussion

'Time-warped multi-neuron single spike patterns' describes the novel finding that, after clustering and control for adaptation, pairwise relative spike time differences are 'warped' by positive trial-to-trial changes in a latent factor for pairs of neurons separated over multiple cortical layers and columns. It is proposed the latent factor may correspond with trial-to-trial changes in the shared of excitability of neuron pairs separated over multiple cortical layers and columns. Alternative explanations for the latent factor modulation are discussed below. The results show that clustering, non-45° noise correlations and adaptation, together, explain a large proportion of spike time variability and should be considered collectively when aiming to detect spike

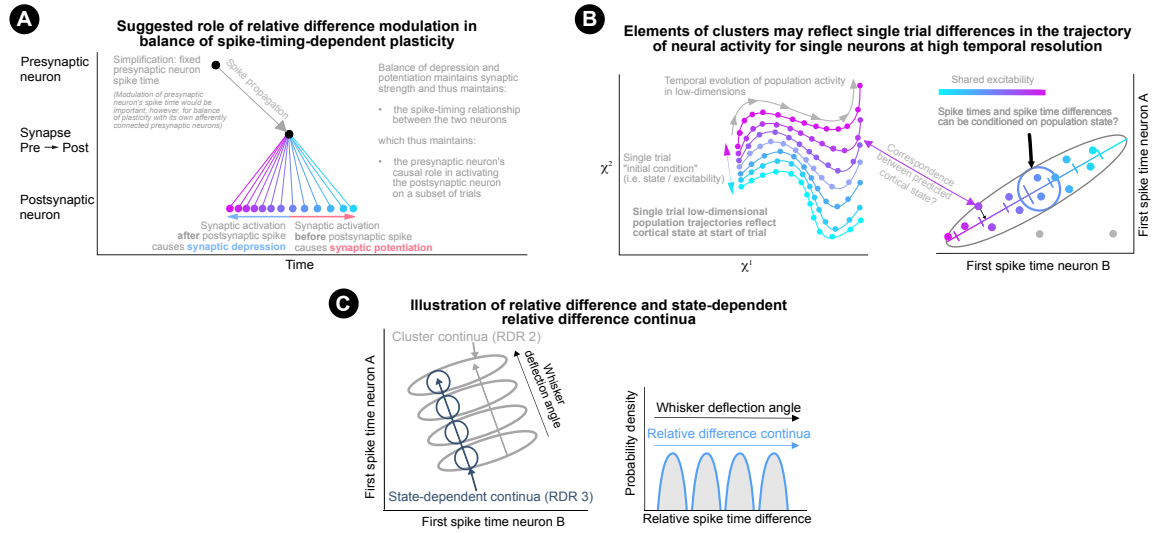


Figure 3.35: Discussion figures. (a) Suggested role of relative difference modulation in balancing spike-timing-dependent plasticity. A fixed spike time is shown for the presynaptic neuron, although the figure explains why the presynaptic neuron's spike time should also be modulated. (b) Illustration of possible correspondence between single trial variability in low-dimensional population dynamics and single trial spike times for correlated clusters. (c) Left: Illustration of continua of first spike time distributions, corresponding to first spike times (not relative differences) under RDR 2 (grey) and 3 (dark blue). Right: Illustration of continua of actual relative spike time difference distributions, which could correspond to RDR 2 or 3.

time relative difference representations.

Co-modulation of spike times for pairs of neurons separated over multiple cortical columns and layers introduced the possibility that spike time patterns are modulated by low-dimensional non-adaptative trial-to-trial population-wide changes in neuronal excitability. This is tested further in Chapter 4. The results presented in this chapter were published in Scientific Reports in July 2021, along with some of the results from Chapter 4.

Other studies

The study differs from revealment of oscillatory activity and spike train alignment to behavioural events through optimisation of population-wide shift and/or stretch functions [293], firing rate time-warp modelling [157] and spike train alignment to behavioural events [257]. This chapter characterises the millisecond and sub-millisecond structure of spike-timing stimulus-response distributions, whilst accounting for clustering and non-stationarity. A novel form of sensory representation (time-warped multi-neuron single spike patterns) is proposed with a potential for sub-millisecond

relative spike time difference encoding. Non-45° correlation angles also demonstrate different time-warp functions for the spike times of different neurons. One previous study modelled time-warp-invariant integration [106], which differs from state-dependent processing.

Membrane potential correlations, spontaneous activity & balance of plasticity

Sources of noise correlations are reviewed in [63]. Subthreshold membrane potential correlations are a prime candidate to underlie spike time noise correlations caused by trial-to-trial changes in shared excitability, as they are correlated across areas [220] and affect action potential initiation [219]. Heterogeneous relationships between membrane potential changes and spike time latency could underlie non-45° correlations.

Patch clamp recordings of single cells have highlighted different states of neuronal membrane potentials under different levels of awake arousal, measured by pupil diameter and task performance [182]. In particular, [182] showed that pupil diameter was correlated with changes in membrane potentials and spiking patterns across neurons, as well as task performance. As pupil diameter is correlated with changes in state across neurons, these recordings could offer the opportunity to study the effect of cortical state changes on spike time patterns and multi-neuron single spike patterns.

The results open the opportunity to study the relationship between stimulus-evoked spike time noise correlations and noise correlations during spontaneous activity (previously characterised for the dataset [231]). The thesis proposes that the purpose of relative spike time difference modulation, may be to maintain spike time patterns/relationships, by balancing spike-timing-dependent synaptic depression and potentiation (Fig. 3.35a).

Population dynamics

Single trial variability in population dynamics is thought to reflect differences in population excitability or state at the start of a trial [143]. As such, population dynamics may offer a proxy for a cortical state, which modulates spike time patterns. Furthermore, the characterised spike timing variability may reflect the relationship between variability in population dynamics and single neuron spike time (and pairwise spike time difference) variability. In another sense, the characterised spike timing variability may reflect variability in the trajectory of population dynamics at a single neuron level and at high temporal resolution (Fig. 3.35b). Indeed, the characterisation

of shared variability on a continuous spike timing scale could offer the opportunity to study population dynamics without loss of information through discretisation of neuronal responses, which typically uses bins of 10-20ms. The precise structure and latent factor dependence of relative difference representations uncovered here suggests that such representations may encode information that is additional to information decodable from classical binned spike count analysis (including classical modelling of population dynamics).

Encoding + Decoding.

The results offer a large potential improvement for discriminable encoding under fixed-state-decoding and state-dependent-decoding regimes. Testing of decoding performance was limited by low cluster counts per stimulus condition and variability in cluster detection between stimulus conditions. For example, most clusters were detected at $0 - 4Hz$ stimulations and clusters were unlikely to be detected for many stimulus conditions for the same neuron pair, likely due to more non-spiking trials for higher stimulation frequencies due to spike count adaptation above $4Hz$ ([46, 231, 290]). This does not imply the characterised form of representation is not utilised at higher frequencies; but means that fewer samples are available for characterisation at higher frequencies.

Dimensionality

The detection of noise correlations over multiple cortical layers and columns suggests that possible trial-to-trial population-wide changes in neuronal excitability could be low-dimensional. This would increase the likelihood that downstream neurons were modulated by the same changes and could respond in a state-dependent manner. Determining the dimensionality of excitability changes requires characterisation of noise correlations over larger groups of neurons (explored further in Chapter 4). Determining the dimensionality of single spike pattern noise correlations is made challenging by the observed clustered nature of pairwise spike time response distributions and low spiking reliability for single neurons. The observed clusters, however, offer the opportunity to to study a reliable part of stimulus representations, and to explore the dimensionality of spike time pattern modulation.

Future experiments.

Future exploration of the characterised phenomena, including determining dimensionality of possible excitability changes and assessing encoding performance, would benefit from high trial counts and experimental paradigms designed to minimise adaptation [108]. Use of dynamic non-stationary and non-parametric clustering may also be beneficial both in the presence and absence of adaptation (Methods).

Clustering + Gaussianity.

As the central goal of the manuscript was to extract isolated regions of high response probability, the author chose to use DBSCAN as this is known to perform more robustly than GMMs at capturing the structure of the underlying probability density function when regions are non-Gaussian [238]. For example, because adaptation was likely to affect the shape of these regions by increasing latencies over trials, it was important to extract regions of high density for further analysis independently of an assumption of their structure. The non-Gaussianity of the isolated regions of high response probability was confirmed for a proportion of the extracted clusters.

The first principal component and its angle offered a first approximation of the relationship of spike times and spike time differences with trial-to-trial changes in shared excitability (for the isolated regions of high response probability). Moreover, this approximation was able to explain the variance in spike times of other neuron pairs. Compared to the unclustered case, clustering better captured isolated regions of high-response probability for application of PCA, leading to novel detection of correlations and more precise cluster angle estimation. Moreover, a large proportion of clusters determined as Gaussian had correlation angles different from 45° .

Future work may aim to model non-linear relationships with underlying state, which in turn could enable more precise estimations of single trial state. Silhouette analysis could be used as a measure of cluster quality and as a criteria for selecting clusters.

The form of continua encoding

The encoding performance of RDR 2 and 3 would be best assessed over a range of similar stimuli. For example, small changes in whisker stimulation angle are likely to elicit slight changes to the response distributions, whilst not drastically changing the number of spiking trials per neuron. This would improve similarity in the number

of clusters detected per stimulus and neuron pair. It can then be tested which response distribution properties improve decoding of transformation and correlate with behavioural performance. It is predicted that cluster (RDR 2) or state-dependent relative difference distributions (RDR 3) may change continuously with continuous stimulus transformations (Fig. 3.35c). More broadly, it is proposed that continua may be encoded in the cortex by time-warped relative spike time differences.

Alternative hypotheses of non-latent factor

Alternative hypotheses of the non-adaptative co-modulating latent factor are:

1. Changes in the afferent signal. This hypothesis suggests that trial-to-trial changes in the afferent stimulus-evoked signal may affect the recorded stimulus-evoked pattern. For example, the latencies of different neurons may respond heterogeneously with respect to the strength of the input pattern. Under this hypothesis, then state would be implicitly encoded in the afferent neural activity. Under this alternative perspective, downstream neurons might be able to infer state from afferent neural activity.

2. Causal neural interactions. This hypothesis would suggest that changes in response latencies of single neurons depend on changes in the spike timings of other neurons. This reduces to a problem of what causes the timings of the afferent spikes to change. Differences in the timing of inhibition, for example, are a possible cause of the observed spike time modulation. The results show that the timing of inhibitory neuron spikes are positively correlated with the timing of excitatory neuron spikes. This is the opposite to the expected effect, in which early inhibition would decrease the latency of other neurons.

Further discussion of methods

To control the possible ill-posedness and the robustness problem of the underlying clustering problem in non-Gaussian situations with strong noise and small data samples, a robust clustering pipeline was designed which took into account the level of noise present in the data and the low number of samples in each response distribution. It follows that surrogate control datasets should match the general statistics of the response distributions. For these reasons surrogate control datasets were created by transforming pre-existing response distributions.

The joint roles of adaptation and trial-to-trial changes in a non-adaptative latent factor promote the use of dynamic clustering techniques in future work where adaptation is present [183]. The multiscale character of the underlying neural processes introduces the possibility of low-dimensional temporally persistent latent regime-switching processes, which are a common source of non-stationarity across nature [183] and are also commonly ignored by analysis techniques [291]. Such processes are likely to contribute to the complexity of the observed time series, including clusters.

Techniques, which address these issues for limited sample sizes (i.e. [122]) should be applied in future research where adaptation is present to (i) characterise the multiscale nature of regime switching processes, (ii) estimate stimulus representation independently of adaptation and invisible biases caused by multi-scale processes beyond depressive spike time latency adaptation [291] and (iii) characterise the dimensionality of trial-to-trial changes in excitability . Understanding of (i) and (ii) will allow stricter control of adaptative and multi-scale processes for the understanding of (iii). This paper aimed to take a first step towards understanding such dynamical processes by characterising a reliable subset of responses (clusters).

Autoencoders could be used to infer the dimensionality of the spiking responses. Particularly, an autoencoder could be designed to reproduce the spike latencies. The number of hidden states needed for the autoencoder to reproduce the spike latencies might give information regarding the dimensionality of latent factor modulation. Focus on spike latencies should be simpler for an autoencoder to reproduce compared to full spike trains. Attention must be paid to the form in which the latencies are represented for the autoencoder.

To test of the hypothesis that trial-to-trial changes in shared excitability underlie the observed phenomena, it could be tested whether the single trial state predictions can predict non-spiking trials using a generalised linear model, for example.

Recent techniques using recurrent neural networks to reproduce spiking activity [18] could be adapted for a latency code with hidden factor variables. As with autoencoders, the number of hidden variables needed to reproduce neural activity could give information regarding the dimensionality of latent factor modulation.

Analyses could have been performed skipping the first 5-20 trials, as this may have increased the number of stationary clusters. Responses could have been pooled over low-frequency stimulus conditions to increase the number of trials per stimulus.

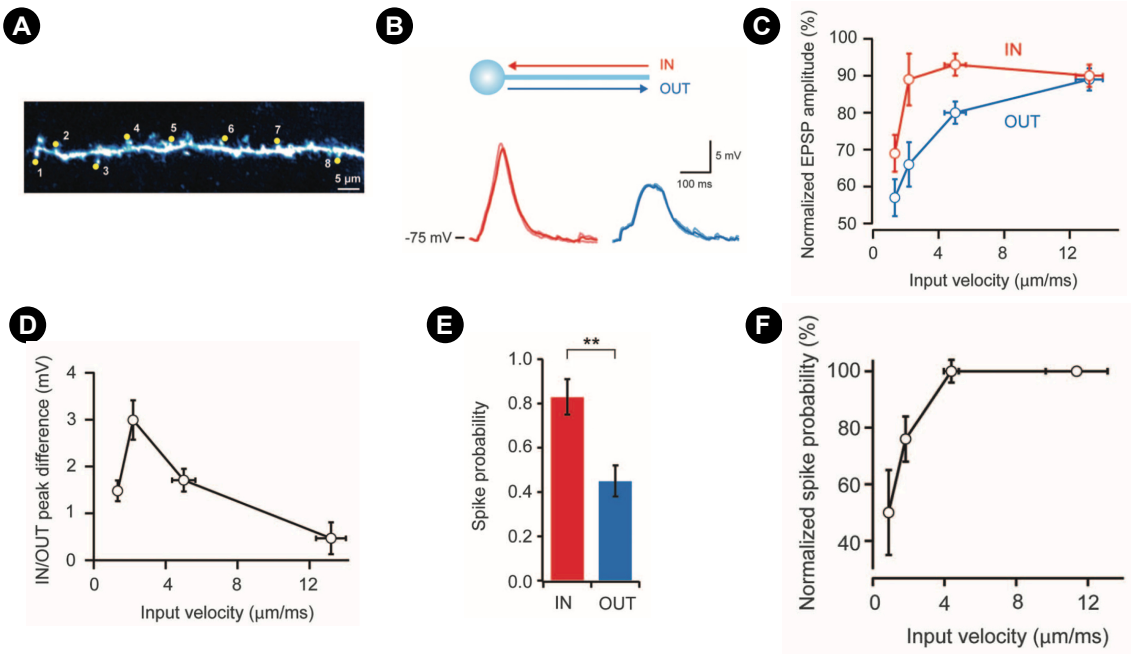


Figure 3.36: From [29]. Reprinted with permission from AAAS. (a) Uncaging spots (yellow) along the selected dendrite. (b) Somatic responses to IN (red) and OUT (blue) directions at 2.3 mm/ms (averages denoted by bold lines). (c) Relation between peak voltage and input velocity (values normalized to the maximum response in the IN direction for each cell, n = 15). Error bars indicate SEM. (d) Relation between direction selectivity and input velocity (n = 15). (e) Direction selectivity of spike probability (**P = 0.0013, n = 7). (f) Relation between spike probability and velocity (n = 7, average of both directions).

Possible state-dependent processing and readout mechanisms

Finally, this section considers how postsynaptic neurons might integrate relative differences in a state-dependent manner. Previous work has explored how neuronal responses are sensitive to the spatial and temporal properties of synaptic patterns of activation [92, 29, 30], as well as how membrane potential affects this sensitivity. The relevant results of [29] characterising mechanisms of dendritic decoding are first outlined. A proposal of how the mechanisms of [29] may combine to support the hypothesised state-dependent decoding of multi-neuron single spike patterns, is then given. This is followed by additional discussion of state-dependent decoding.

[29] studied how temporal input sequences are integrated by single dendrites using two-photon glutamate uncaging. Particularly, synapses on single dendrites were activated in spatial temporal patterns (Fig. 3.36a), and the effect of temporal and spatial properties of such patterns was explored. The effect of membrane potential was also studied.

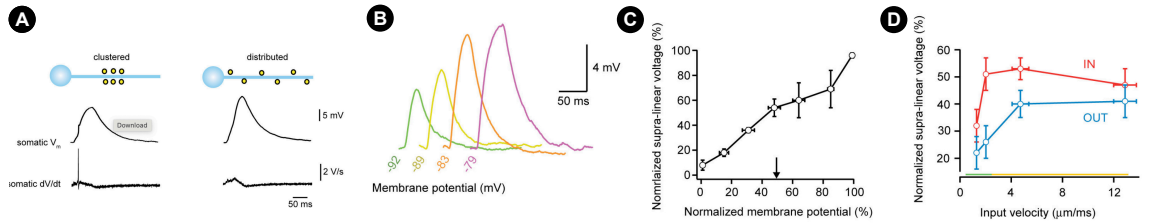


Figure 3.37: From [29]. Reprinted with permission from AAAS. EPSP supralinearity varies in a graded manner. (a) Spatially clustered activation of 6 spines produced a fast dendritic spike detectable as a notch in the somatic voltage recording, and as a sharp peak in the corresponding dV/dt trace(4, 5, 32, 33) (left; 4 ms interval; all inputs within 15 m); however no fast dendritic spike is produced when the same number of inputs are distributed along the length of a dendritic branch during a direction-selective response (right; 4 ms interval). (b) EPSPs recorded at different membrane potentials produced by sequential activation of 8 dendritic spines at 1 ms intervals; membrane potential was varied with somatic current injection. (c) Supralinear peak voltage (calculated by subtracting the sum of individual spine responses) gradually increases as Vm becomes more depolarized. Arrow indicates resting membrane potential. Pooled data from 5 cells; membrane potential has been normalized to the most depolarized value tested (mean resting potential at arrowhead). (d) Summary plot showing that the supralinear component of the EPSP varies gradually with the velocity and direction of uncaging (normalized to the IN maximum in control cells, P = 0.0410 and P = 0.0348 for IN and OUT respectively, ANOVA test). While at high input velocities (mean 6.3 ± 0.83 m/ms, range indicated by yellow line above X-axis) responses were 44.7 ± 2% of the maximum, at lower input velocities (1.4 ± 0.06 m/ms, green line above X-axis) the supralinearity decreased to 27.8 ± 4% (P ¡ 0.0001; t-test, average of both directions). Also, at the optimal speed for direction sensitivity of each cell, the IN EPSPs were on average 53 ± 3% of the maximum and the OUT responses 33 ± 8% (P ¡ 0.0001; paired t-test). In A and B, traces are averages of 5 sweeps.

Firstly, [29] showed that the direction of synaptic activations along a dendrite affects the size and shape of somatic EPSPs, with the inward direction (from branch tip to soma) producing higher EPSPs (Fig. 3.36b). Moreover EPSP magnitudes were shown to depend on the velocity of synaptic activations along the dendrite (Fig. 3.36c,d). Somatic spike probability was higher for the inward direction (Fig. 3.36e) with spike probability also increasing with the velocity of synaptic activations along the dendrite (Fig. 3.36f). Spatially clustered synaptic activations were more likely to produce fast dendritic spikes (Fig. 3.37a).

EPSPs were larger at higher membrane potentials (Fig. 3.37b)), although membrane potentials were controlled with somatic current injection. In reality, EPSPs would cause excitatory synaptic conductances to decrease, and so the background

current flowing into the cell would be smaller during the EPSP than before. As such, one might expect that the effect of membrane potentials on EPSP magnitude might be smaller than that observed in Fig. 3.37b.

Supralinear increases in voltage (likely caused by opening of voltage-gated NMDA receptors) increased at higher base membrane potentials (Fig. 3.37c), although the same caveats about current injection apply. Supralinear increases in voltage also increased with the velocity of synaptic activations in both the inward and outward directions (Fig. 3.37d).

Fig. 3.38a demonstrates the supralinear effect of activating multiple synapses at once. Fig. 3.38b,c,d give example demonstrations of the dependence of dendritic activation direction selectivity on NMDA receptors. In particular, direction selectivity is abolished in the presence of an NMDA receptor blocker.

Peak EPSP amplitudes varied with the directionality of the pattern at a depolarised membrane potential of -66mV but not at a hyperolarised membrane potential of -80mV (Fig. 3.38e,f,g). The reduced EPSPs at lower membrane potentials would reduce discriminability. Interestingly, 'directionless' patterns produced lower EPSP peaks than outward patterns (Fig. 3.38g).

[29] show that activation of distal synapses produces larger influxes of cations (Fig. 3.39), due to increased NMDAR conductance. This is used to explain the differences in EPSPs for the inward and outward directions.

[29] also modelled the effect of both dendritic and backpropagating inhibition on peak EPSPs. The model showed that dendritic 'feedforward' inhibition decreased EPSP magnitudes, whilst the variability in EPSP magnitudes was maintained (Fig. 3.40). The authors use the second finding to suggest that pattern discriminability is also maintained under inhibition.

Proposal of state-dependent coding mechanisms based on the results of [29]

Given the mechanisms outlined in this section, this subsection considers how these mechanisms might combine to produce the hypothesis of state-dependent decoding of multi-neuron single spike patterns. To resummarise, the relevant mechanisms are:

1. The direction of synaptic activations along a dendrite affects the size and shape of somatic EPSPs, with the inward direction (from branch tip to soma) producing higher EPSPs (Fig. 3.36b).

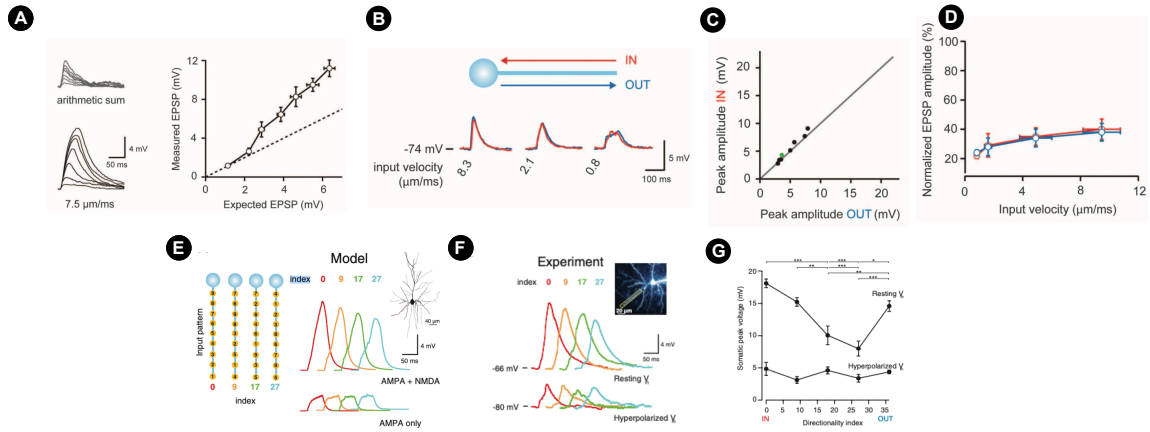


Figure 3.38: From [29]. Reprinted with permission from AAAS. (a) (Left) Activating more synapses (1 to 7) produces EPSPs that are progressively larger than the linear sum of the individual responses. (Right) Summary plot of eight cells. Solid line, experimental data points; dashed line, unity line. Error bars indicate SEM. (b) Velocity and direc- tion sensitivity is abolished in AP5. (c) Comparing peak amplitudes for IN and OUT sequences at the optimal velocity for direction selectivity. (d) Reduced velocity sensitivity in AP5. (e) Random patterns of nine inputs were generated and sorted by their directionality index (35). Four sample patterns (numbers showing activation order) and corresponding directionality index are shown. Patterns were played into a single dendrite in a pyramidal cell model or (f) experimentally in single pyramidal cell dendrites. The peak response is different for different patterns, and responses become similar with hyperpolarization (or without NMDA receptors in the model). (Inset) Dye-filled neuron with dendrite used in the experiment enclosed by a yellow box. Vm, membrane potential. (g) Peak EPSP voltage measured experimentally depends on input-pattern. Asterisks indicate significant differences between EPSP peaks for the example shown; hyperpolarization reduces the ability to discriminate different patterns. Error bars indicate SEM.

2. EPSP magnitudes depend on the velocity of synaptic activations along the dendrite (Fig. 3.36c,d).

3. Somatic spike probability is higher for the inward direction (Fig. 3.36e) and also

4. Somatic spike probability increases with the velocity of synaptic activations along the dendrite (Fig. 3.36f).

5. Spatially clustered synaptic activations are more likely to produce fast dendritic spikes (Fig. 3.37a).

6. EPSPs were larger at higher membrane potentials (Fig. 3.37b)), although membrane potentials were controlled with somatic current injection.

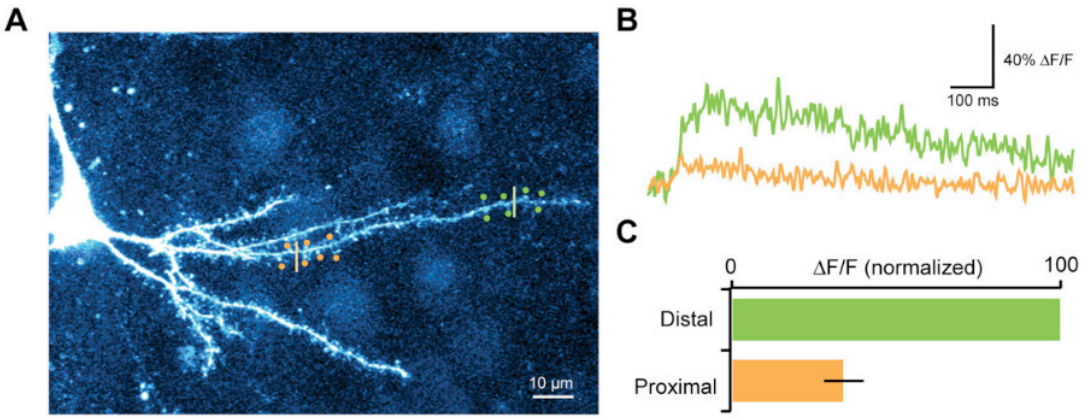


Figure 3.39: From [29]. Reprinted with permission from AAAS. Local impedance differences result in different levels of NMDAR conductance. (A) A cluster of 7 spines was placed either at the tip of the dendrite (green dots) or at the base of the same branch (orange dots), and activated with 1 ms intervals while local Ca2+ influx was monitored by a line scan across the dendrite (indicated by green and orange lines). (B) Activation of distal synapses produces a larger Ca2+ transient than proximal synapses. (C) Across 5 cells, proximal synapses elicited 34 ± 6% of the peak distal Ca2+ influx (P = 0.0003, paired t-test).

7. Supralinear increases in voltage are higher at higher base membrane potentials (Fig. 3.37c), although the caveats about current injection apply.

8. Supralinear increases in voltage also increased with the velocity of synaptic activations in both the inward and outward directions (Fig. 3.37d).

9. Peak amplitudes were shown to vary with the directionality of the pattern at higher membrane potentials (Fig. 3.38e,f,g). EPSPs were reduced at lower depolarisations, however, reducing discriminability. Interestingly, 'directionless' patterns produced lower EPSP peaks than outward patterns (Fig. 3.38g).

10. Modelling showed that dendritic 'feedforward' inhibition decreased EPSP magnitudes, whilst the variability in EPSP magnitudes was maintained (Fig. 3.40). The authors use the second finding to suggest that pattern discriminability is also maintained.

Please note that the following is intended as a speculative first exercise/hypothesis considering the mechanisms in relation to the described form of state-dependent decoding. The proposal is unlikely to be correct, but has the potential to guide future thought and refinement of state-dependent coding hypotheses.

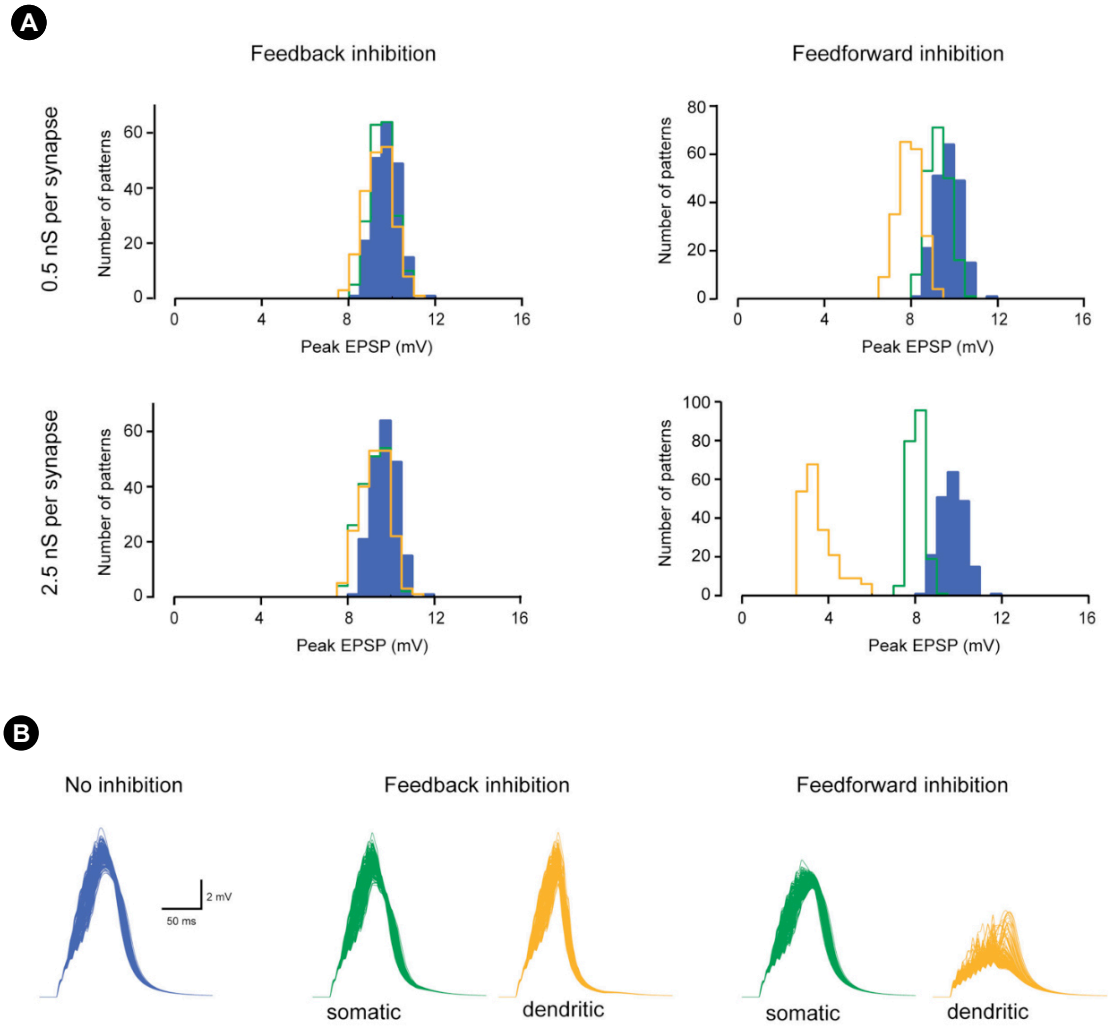


Figure 3.40: From [29]. Reprinted with permission from AAAS. Effect of inhibition on single dendrite sequence discrimination. Excitatory synapses on a single dendrite were activated in the model with random sequences, and different types of inhibition were evaluated (see Methods for details). (A) Histograms for the peak EPSP amplitude generated under the different conditions (blue - no inhibition; green - somatic inhibition; yellow - dendritic inhibition), showing that strong feedforward dendritic inhibition can shift the EPSP histogram to smaller values, while the variance is maintained. Feedback inhibition has minimal impact on the discrimination of excitatory sequences. (B) Somatic voltage traces in response to the same sequences, while varying the location and type of inhibition.

Fig. 3.38g from [29] plotted the directionality of a pattern against the somatic peak voltage at two membrane potentials (resting potential and hyperpolarised). A monotonic relationship between directionality and peak somatic EPSP at resting potential, is assumed. That is, the response to the outward pattern (directionality

index 35) in Fig. 3.38g is ignored. This simplifies the explanation and allows a first hypothesis of a mechanism underlying state-dependent coding. It is possible that such outward pattern do not occur in vivo.

From Fig. 3.38g, and as an exercise, it is assumed that the inward pattern corresponds to the preferred pattern at high excitability. This is because the inward pattern produces the highest amplitude peak somatic EPSP at resting potential. It is important to note that the results of the paper do not explore the effect of inhibition, except through modelling.

Under the superposition that the inward pattern corresponds to the dendrite's preferred pattern in the high excitability state, it would be likely that feedforward inhibition arrives faster when the feedforward pattern is instantiated. This could have the effect of lowering peak somatic EPSPs for inward patterns (relative to observed experimental values which do not consider inhibition). This reduction in EPSPs could be such that there is a positive monotonic relationship between direction index and peak somatic EPSP at lower membrane potentials, whilst the negative monotonic relationship is maintained at higher membrane potentials, with the magnitude of the gradient decreased. This would create opposite relationships between direction selectivity and peak somatic EPSP, at the two membrane potentials, which would support selective state-dependent decoding. This offers a first hypothesis of how state-dependent decoding could be supported.

Further discussion of state-dependent decoding mechanisms

It may be useful to note that the proposed form of state-dependent decoding can be though of as a 2-dimensional XOR problem, where one of the variables is state/membrane potential, and the second could be some property of the pattern of synaptic activations on the dendrite (i.e. directionality/velocity). Considering line plots such as Fig. 3.38g or heatmaps (x: direction/relative difference, y: membrane potential, color: EPSP) may be useful in considering how mechanisms might combine to achieve a form of state-dependent coding.

It is also possible that the main proposal that multi-neuron single spike patterns may be interpreted in a state-dependent manner, may be incorrect. For example, it might be that discriminability is increased at higher membrane potentials, and reduced significantly at lower membrane potentials, such that high membrane potentials offer a more precise decoding regime.

Finally, it is useful to consider that state-dependent decoding in the broadest sense of neural dynamics. The term 'to discriminate' describes a neuron's ability to

produce different output spike trains over a set of input patterns. State-dependent discrimination, then, in the simplest form would mean that a neuron's membrane potential would affect the neuron's output spike train in response to a particular input pattern. In this simplest sense, neurons are natural state-dependent discriminators. That is, neurons can be considered state-dependent discriminators if there is a mapping from input patterns to output spike trains, which depends on the neuron's membrane potential.

In the supplementary materials [29] also provides an analysis of the effect of jitter on the postsynaptic responses, which will be important to consider.

Chapter 4

Clustering and control for adaptation uncovers time-warped spike time patterns in cortical networks in vivo (part 2)

4.1 Introduction

In Chapter 3, the first spike times of pairs of neurons, separated over multiple cortical columns and layers, were shown to be correlated after controlling for stimulus-response adaptation. Such co-modulation of spike times for pairs of neurons separated over multiple cortical columns and layers introduced the possibility that spike time patterns are modulated by low-dimensional non-adaptative trial-to-trial population-wide changes in neuronal excitability. It was proposed that such wide-spread low-dimensional modulation makes it possible that decoding neurons are also modulated by the same changes in excitability, and that this could allow decoding neurons to process spike time patterns in an excitability-dependent (i.e. state-dependent) manner. An improved potential encoding capacity was proposed under the possibility of this regime.

This chapter presents both published and unpublished analyses aiming to explore the dimensionality of spike time modulation for groups of 4 or more neurons. This extends the results of Chapter 3, which explored spike time correlations for pairs of neurons. The chapter first presents the additional analyses presented in the author's publication in Scientific Reports [136]. These analyse explore correlations between the predicted states of disjoint pairs of correlated neuron pairs. Particularly, the spike times of each stationary correlated cluster or unclustered response distribution from Chapter 3 enables a prediction of single trial cortical state. In this chapter it is shown

that the these single trial cortical state predictions are correlated between pairs of neuron pairs.

The second (unpublished) set of analyses presents methods designed to explore the co-modulation and dimensionality of larger groups of neurons, with up to 12 neurons analysed for a single stimulus condition. The analyses presented in this chapter account for the difficulties introduced by non-spiking trials (discussed in Section 4.1.1). The finding that the spike times of a large number of neuron pairs in groups of 3-12 neurons are correlated, strengthens the case that trial-to-trial changes in cortical excitability are low-dimensional. A prediction of population cortical state is made from the spike times of groups of 3-12 neurons. The correlation of the predicted population cortical state with single neuron spike times and neuron pair spike time differences is then made. The prediction of population cortical state also enables a novel visualisation of the effect of population cortical state on the temporal structure of multi-neuron single spike patterns.

4.1.1 The problem of low spiking reliability and clustered correlations

As expected for cortical neurons and as seen in Chapter 3, the recorded neurons exhibited a range of spiking reliabilities (Fig. 3.5), with the majority of recorded neurons having low spiking reliabilities over 50ms following stimulus onset. As discussed briefly in Chapter 3, low spiking reliabilities complicate analysis of multi-neuron single spike patterns. For example, for a single stimulus condition, a matrix of first spike times by neuron and trial index would contain a large number of missing values.

Low spiking reliability complicates analysis of the correlations between the predicted single trial states of pairs of correlated neuron pairs. Firstly, the number of trials for which a pair of neurons both spike is almost always lower than the number of trials of the stimulus condition. If correlations between pairs of neurons are detected for a subset (i.e. a cluster) of trials of a stimulus condition, there is a further reduced number of trials for which a state can be predicted from a pair of neurons. The conjunction of trials on which two pairs of neurons spike, leaves a further reduced subset of trials. This reduces the statistical power of the test for correlation between the predicted states of two disjoint neuron pairs, and in many cases there are insufficient samples to make a useful test of the correlation.

The issue of low spiking reliability also complicates analysis of correlation matrices. If there are insufficient conjunctive trials for a pair of neurons, then the r-value of

the correlation between the pair of neurons may be inaccurate. The analyses of this chapter were implemented to handle missing data due to low-spiking reliability.

4.2 Methods I: Pairwise correlation of predicted cortical state

The spike times of each stationary correlated cluster or unclustered response distribution from Chapter 3 enables a prediction of single trial cortical state for the corresponding neuron pair. This section analyses the correlations between the predicted cortical state of disjoint neuron pairs, as a first step in testing whether non-adaptative trial-to-trial changes in shared excitability were low-dimensional. If the predicted cortical states of pairs of neuron pairs were correlated, this case for low-dimensional population-wide modulation of spike times would be strengthened. These analyses were presented as part of the author's publication in Scientific Reports [136].

4.2.1 Composing the pool

Correlated Stage 2 clusters and correlated unclustered response distributions, which were either determined as stationary or made stationary (by time series modelling) in the previous chapter, were pooled. If a cluster in the pool came from a correlated unclustered response distribution which was also in the pool, the correlated unclustered response distribution (rather than the cluster) was removed from the pool so that the pool contained only one example per response distribution.

For simplicity in this chapter, 'clusters' and 'unclustered response distributions' are herein collectively referred to as 'clusters', unless explicitly stated otherwise. In this chapter, 'stationary' clusters and clusters 'made stationary' by time series modelling are herein referred to as 'stationary'.

4.2.2 Correlations of predicted pairwise states

The position of each spike pair along a cluster's first principal component enabled a prediction of cortical state (i.e. excitability) on a single trial (Fig. 4.1a). As an initial test, the correlations between pairs of correlated clusters in the pool were tested. Only pairs of clusters which did not share a common neuron were included in the analysis. Moreover, more than 15 conjunctive trials between a pair of clusters were needed for the pair of clusters to be included in the analysis.

4.2.3 Correlations of predicted states with spike times and spike time differences

To test the efficacy of single cluster state predictions to predict the spike times and spike differences of other clusters, the correlations between predicted states of correlated pairs with the spike times and spike time differences of other clusters was also tested. Only pairs of clusters which did not share a common neuron were included in the analysis. Again, more than 15 conjunctive trials between a pair of clusters were needed for the pair of clusters to be included in the analysis.

4.3 Results I: Pairwise correlation of predicted cortical state

4.3.1 Composition of the pool

In total, the pool contained 254 elements: 140 correlated clusters and 114 correlated unclustered response distributions (referred to collectively as 'clusters' in this chapter unless stated otherwise). 182 pool clusters were originally stationary, with the remaining 72 made stationary through time series analysis.

4.3.2 Correlations of predicted pairwise states

Predictions of single trial excitability were correlated for a large proportion of disjoint pairs of clusters (i.e. pairs of clusters which do not share a common neuron; Fig. 4.1b, left). Moreover, all statistically significant correlations ($p < 0.05$) were positive (Fig. 4.1b, right), providing evidence that trial-to-trial changes in excitability modulate the spike times of up to 4 spatially separated neurons.

Fig. 4.1c shows that the first PCA dimension explains $\sim 70\%$ of explainable variance between the spike times of groups of 4 neurons (on trials that all 4 neurons spiked). This further strengthens the possibility that spike time patterns are modulated by low-dimensional non-adaptative trial-to-trial population-wide changes in neuronal excitability.

4.3.3 Correlations of predicted states with spike times and spike time differences

Over trials of single stimulus conditions, the predicted state of single correlated neuron pairs is correlated with the spike times (Fig. 4.2a) and spike time differences (Fig.

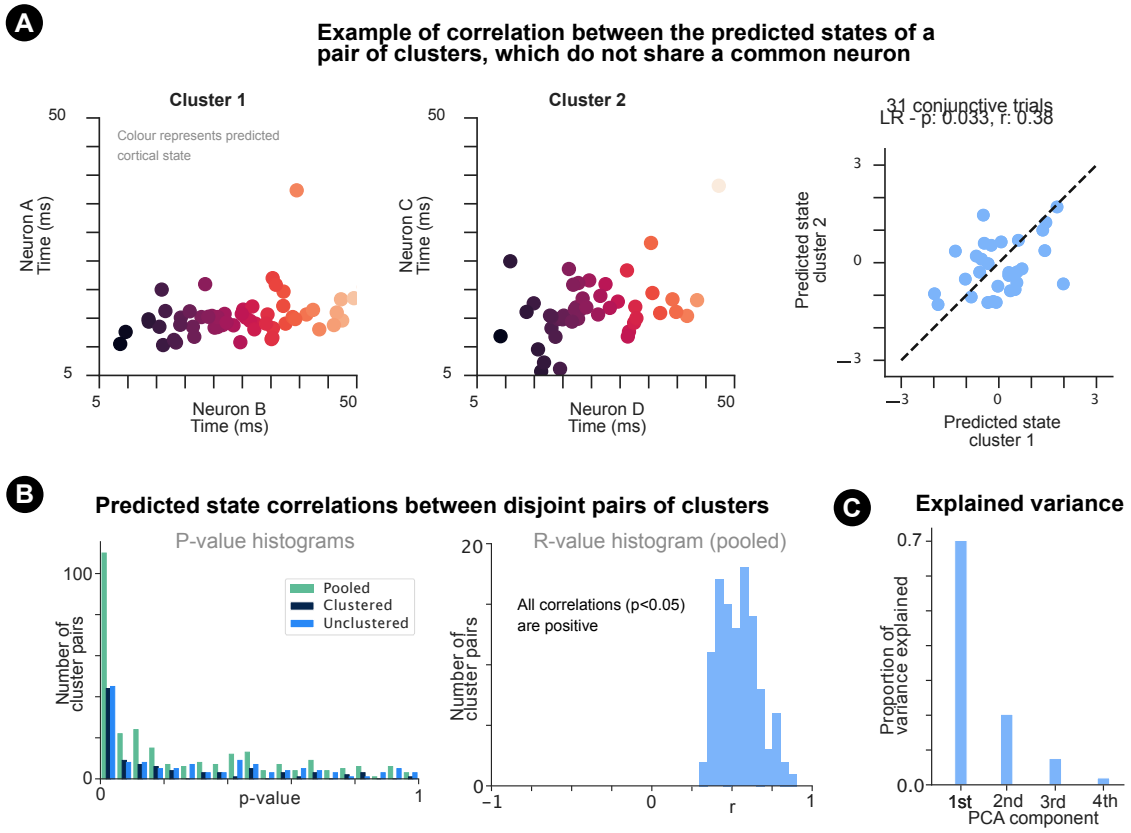


Figure 4.1: Neuron pair cortical state correlation. (a) Example of two clusters which do not share a common neuron but have correlated cortical states over trials of a single stimulus. Left and Center: cluster spike times of two clusters coloured by predicted cortical state (value along first principal component). Right: correlated cortical states on trials on which all 4 neurons spike. (b) Right: histograms showing p-values of linear correlation between predicted states of pairs of correlated pairs of neurons, which do not share a common neuron. Inclusion criteria: 15 conjunctive trials. Cases: clustered (orange), unclustered (green) and clustered and unclustered pooled (blue). Left: r-value histograms of linear state correlations (clustered and unclustered pooled). (c) Mean variance explained by each PCA dimension between groups of 4 neurons (on correlated trials that all 4 neurons spiked).

4.2b) of other correlated neuron pairs, which do not share a common neuron. Fig. 4.2c (left) demonstrates a reduction in variability when spike times are conditioned on the predicted state of other single correlated neuron pairs. In some cases, such state-conditioned spike time variability was on a similar order to that predicted by the factor analysis model above (Fig. 4.2c, right).

A reduction in spike time difference variability is also seen when non-45° correlations are conditioned on the predicted state of other correlated single neuron pairs

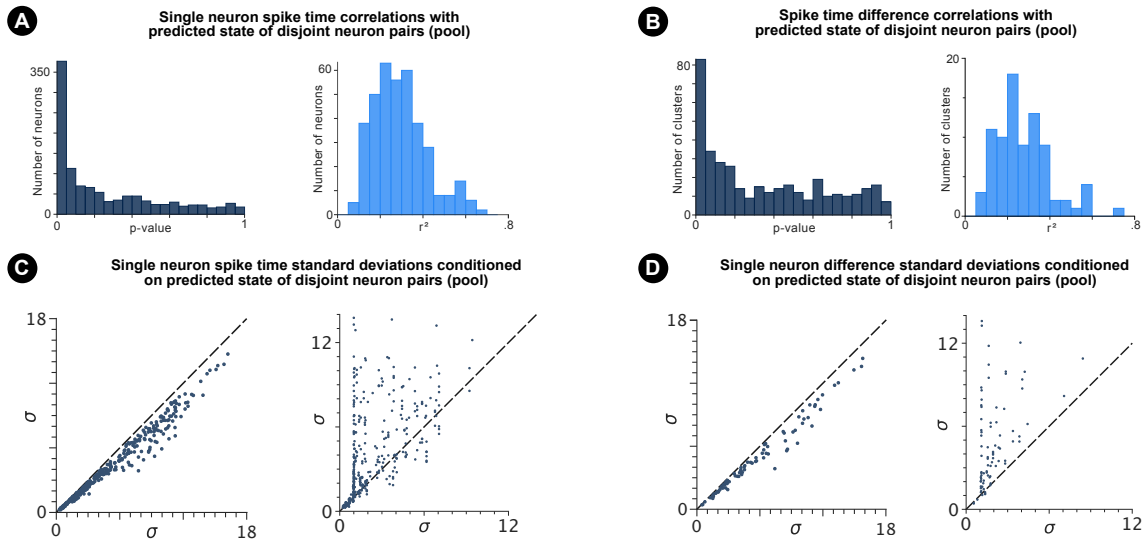


Figure 4.2: State-conditioned spike times and differences. (a) Left: p-value histogram of correlations between predicted states of single correlated neuron pairs with the spike times of other correlated neuron pairs (clustered and unclustered pooled). Right: r-value histogram of statistically significant (p-value) correlations. (b) Same as (a) but for correlations with spike time differences with non-45° correlation angles. (c) Left: variability comparison. X-axis shows the standard deviation of single neuron spike times within stationary clusters (pooled). Y-axis shows the standard deviations of spike times conditioned on the predicted state of single correlated neuron pairs, which do not share a common neuron. Right: Comparison of predicted state-dependent spike time variability predicted by factor analysis (x-axis) vs spike time variability conditioned on the predicted state of single correlated neuron pairs, which do not share a common neuron (y-axis). (d) Same as (c) but for correlations with spike time differences with non-45° correlation angles.

(Fig. 4.2d, left), although variability was mostly higher than the variability predicted by the factor analysis model under RDR 3 (see Chapter 3).

It is important to note that there is inherent noise in the state predicted from single pairs of neurons, which hampers spike time and spike time difference prediction. Future work may explore improved estimates of single trial state.

4.4 Methods II: Spike time co-modulation of larger groups of neurons - the correlation matrix

The remaining analyses of this chapter are currently unpublished and present methods designed to study co-modulation and correlations between larger groups of neurons (4 or greater).

4.4.1 Analysis of first spike time correlation matrices

Analysis of firing rate correlation matrices for repeated trials of single experimental conditions is a typical method for exploring the dimensionality of spiking activity modulation by latent processes. The problem of low spiking reliability (described in Section 4.1.1) complicates such analysis for first spike times, as for some pairs of neurons there are very few conjunctive spiking trials (i.e. trials on which both neurons spike). Estimation of first spike time correlation r-values for such neuronal pairs is likely to have a large error. Moreover, for such pairs, there are insufficient samples to test the significance of the correlation. The correlation matrix is also different to standard correlation matrices, in that each entry in the matrix represents the correlation values for a different subset of trials. This subsection takes a preliminary step in understanding the structure of first spike time correlation matrices through visualisation.

4.4.1.1 Composing the correlation matrix

The following analyses of first spike time correlation matrices are made for stimulus conditions for which at least 3 stationary clusters were detected. Particularly, for each stimulus condition, clusters from the pool of clusters described in 4.2.1 are used. Up to 12 neurons are analysed for a single stimulus conditions.

P-value and r-value matrices were created for each stimulus condition for which more than two correlated clusters were detected. Each row and column of a stimulus condition correlation matrix corresponded to a different neuron which was part of at least one cluster in the pool. For each pair of neurons for which a stationary correlation was detected, the correlation p-value and r-value were inserted in the matrix for the stimulus condition, at the position corresponding to the two neurons.

Fig. 4.3 shows the p-value and r-value matrices for an example stimulus condition (1Hz stimulation frequency). In total, there were 12 neurons for stimulus condition clusters in the pool. For both matrices a green border highlights the neuron pairs for which a stationary correlation was detected (in Chapter 3). This left a number of empty entries in the correlation matrix: for pairs of neurons in the group of 12 neurons where stationary correlations were not detected by the algorithm in Chapter 3. These neuron pairs are highlighted by an orange border in Fig. 4.3).

There were several possibilities for filling the remaining entries of the correlation matrices (i.e. those with orange borders in Fig. 4.3). For example, linear correlation values could simply have been calculated for the conjunctive first spike times of these

additional neuron pairs. This approach had one major disadvantage. Particularly, if analysis of the correlation matrix is for the purpose of considering stationary latent processes, it is important that the spike times of these neuron pairs were stationary. For many of the additional neuron pairs, there were few conjunctive trials, such that tests for non-stationarity would have contained large errors. This would have confounded observation of the correlation matrices when considering the effect of stationary latent processes. It may be beneficial for future work to test this approach, however.

The following describes a first approach to filling the remaining values of the correlation matrix. For each neuron, a vector of first spike times was constructed over all trials of the stimulus condition. This vector was filled with first spike times, which contributed to stationary clusters in the pool. This confirmed that only adaptation controlled spike times were used. For some single neurons, the adaptation controlled correlated first spike times were different between clusters due to different applications of time series modelling to the spike times of different clusters. In such cases, where there were multiple possible values for a single neuron on a particular trial, the spike times of the neuron from the first cluster was used for the single neuron first spike time vector.

At this stage, the spike time vectors were not tested further for non-stationarity and adaptative trends, although this step could be added in future work. It is useful to remember that, the spike times selected for the spike time vectors, had previously been tested for non-stationarity and/or had time series modelling applied. Additional tests for non-stationarity will also made for later analyses in later sections, however.

Following the creation of the first spike time by trial vectors (for each neuron and stimulus condition combination), correlations were then tested between the conjunctive spike times of the additional pairs of neurons (for which spike time correlations were not originally detected). The p and r-values of these correlations were then added to the correlation matrices. As a reminder, these additional entries have an orange border in the correlation matrices shown in Fig. 4.3. All of the p and r-value correlation matrices are shown in Vids. 3 & 4.

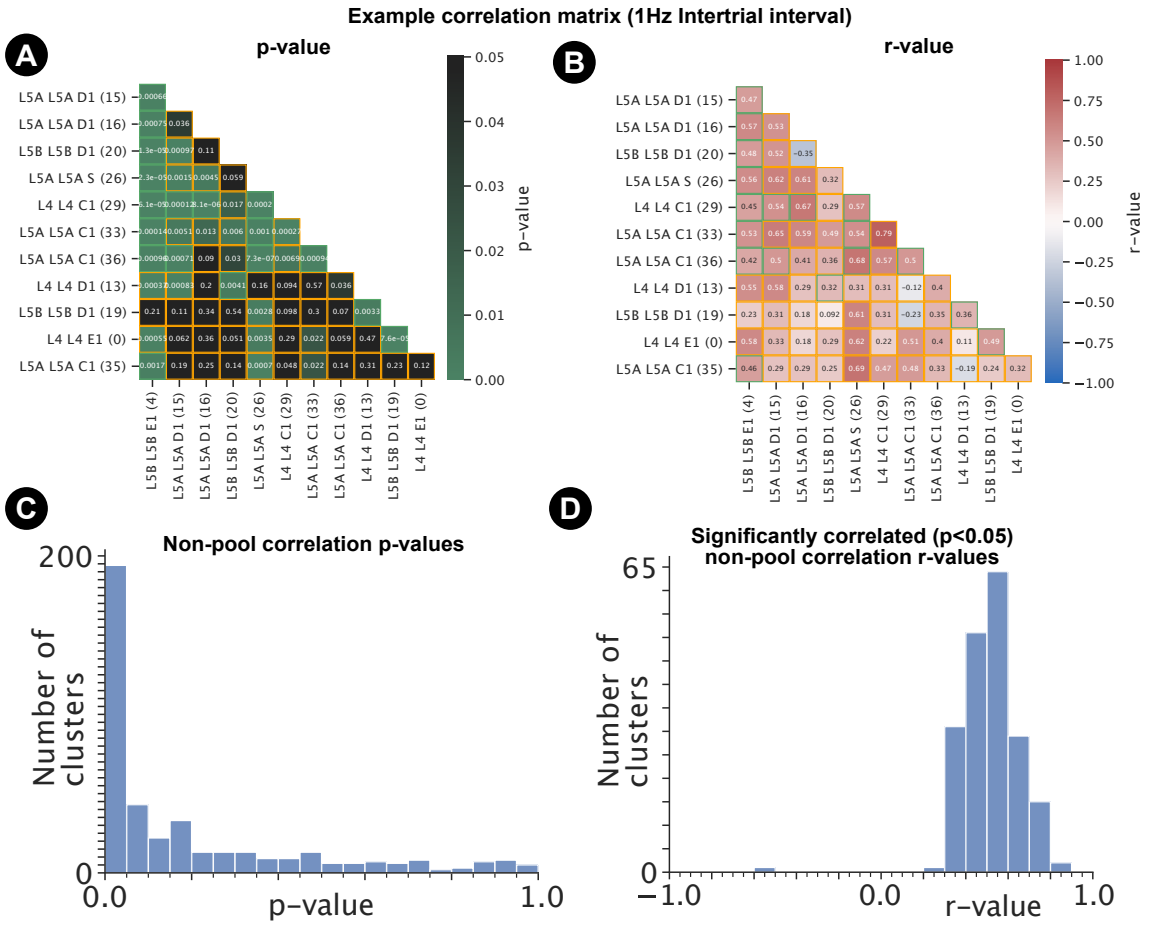


Figure 4.3: Example p-value (a) and r-value (b) correlation matrices for an example stimulus condition (1Hz inter trial interval). Each square has either a green or orange border. A green border signifies that the correlation value is for pair of neurons from a single cluster in the pool. An orange border signifies that the correlation was tested for stationary spike times, obtained from other stationary correlated clusters for the same neuron pair. Left: significant correlations ($p < 0.05$) have a green background, with the shade of green corresponding to the p-value (see key). Non-significant correlation ($p >= 0.05$) have a black background. Right: the background colour of each square represents the r-value (blue for a negative correlation and red for positive; see key). (c) Histogram of non-pool correlations p-values. (d) Histogram of r-values of significantly correlated ($p < 0.05$) non pool pairs.

4.5 Results II: Spike time co-modulation of larger groups of neurons - the correlation matrix

4.5.1 Example: 1Hz inter trial interval

Fig. 4.3a shows example p-value and r-value correlation matrices for a single stimulus condition (1Hz inter trial interval). By definition, all correlations for pairs within the

pool have a significant correlation ($p < 0.05$; shown by the green backgrounds for the green bordered squares). The correlation p-values for these pairs are actually less than 0.005 by definition for inclusion in the pool.

Interestingly, a large proportion (20/48) of the non-pool pairs also have significant correlations ($p < 0.05$). Only 2/48 of the non-pool pairs had a correlation $p > 0.5$ suggesting that correlations of some (or all) non-significant ($p > 0.05$) non-pool pairs may be non-significant due to insufficient samples.

The r-value correlation matrix in Fig. 4.3b shows that only 4/66 pairwise correlations have negative r-values. The correlations of none of these pairs were found to be significant ($p < 0.05$), however. All significant correlations ($p < 0.05$) were in fact positive. The fact that the large majority (62/66) of r-values are in fact positive and all significant correlations are positive, supports the notion that the spike times of all of the 12 neurons are modulated by low dimensional (possibly single dimensional) trial-to-trial changes in shared excitability.

4.5.2 All stimulus conditions

Vids. 3 & 4 show the p-value and r-value correlation matrices for all stimulus conditions for which at least 2 stationary correlated clusters were detected. Observation of Vid. 3 shows that a large proportion of non-pool pairs are significantly correlated ($p < 0.05$). Fig. 4.3c show the distribution of non-pool p-values. Indeed a proportion 42% (194/420) are significantly correlated ($p < 0.05$). Moreover, the histogram is clearly skewed towards low values.

Observation of Vid. 4 shows that for each stimulus condition the vast majority of pairwise r-values (both for pool and non-pool pairs) are positive. It was shown in Chapter 3 that the majority of pool correlations are positive. Fig. 4.3d shows that 99.5% (193/194) of significantly correlated ($p < 0.05$) non-pool pairs are also positively correlated. This supports the notion that the spike times of the *neurons* in the pool are modulated by single dimensional trial-to-trial changes in shared excitability.

4.6 Methods III: Averaged cortical state

4.6.1 Improved prediction of cortical state

The results of the chapter so far suggest that the dimensionality of spike time modulation is low. This section tests whether a one-dimensional 'global' population state can be estimated from the predicted states of multiple correlated pairs for a single

stimulus condition, and whether such a population state can predict the spike times and spike time differences of neurons not used in the calculation of the cortical state. If spike time modulation was indeed largely explained by a single dimension, then it would follow that the an improved estimate of population state could possibly be predicted by averaging over single trials states predicted for neuron pairs.

4.6.1.1 The problem with averaging predicted PC states

This approach has an immediate complication, which is described as follows. Predictions of single trial states were so far predicted for each cluster by taking the value of the single trial spike pair along the cluster's first principal component on a single trial (Fig. 4.1a). A problem arises in that the centres and scales of the first principal component axis for each cluster are likely to be different between clusters.

The problem can be understood by considering a hypothetical case of two clusters. Imagine that spike pairs occur for the first cluster on average at higher levels of excitability than spike pairs of the second cluster. For the first cluster the mean spike time pair (corresponding to a state of 0 on the 1st PC axis) would represent a higher level of excitability then the mean spike time pair of the second cluster. This demonstrates how the 0 state predicted by different clusters may be different between clusters. This complicates a simple averaging of single trial states over clusters, for the purposes of producing a more accurate state prediction.

Even if the 0 states were the same between clusters, it does not follow that the 1st PC axes have the same scale between clusters. The problem can again be understood by considering a hypothetical case of two clusters. Imagine that first spike pairs occur for the first cluster over all levels of excitability but only over a small range of high excitability for the second cluster. For the first cluster, a spike time pair corresponding to low excitability could correspond with a 1st PC value of $\sim +3$. For the second cluster, a spike time pair corresponding to high excitability, but low within the range of excitabilities for which first spike pairs occur, could also correspond with a 1st PC value of $\sim +3$. This demonstrates how the scale of 1st PC axes could be different between clusters. This again complicates a simple averaging of single trial states over clusters, for the purposes of producing a more accurate state prediction.

4.6.1.2 Adjustment of cluster offsets and angles

To test the efficacy of a single dimensional population state in explaining the observed spike time noise correlations, an algorithm was developed to estimate an improved

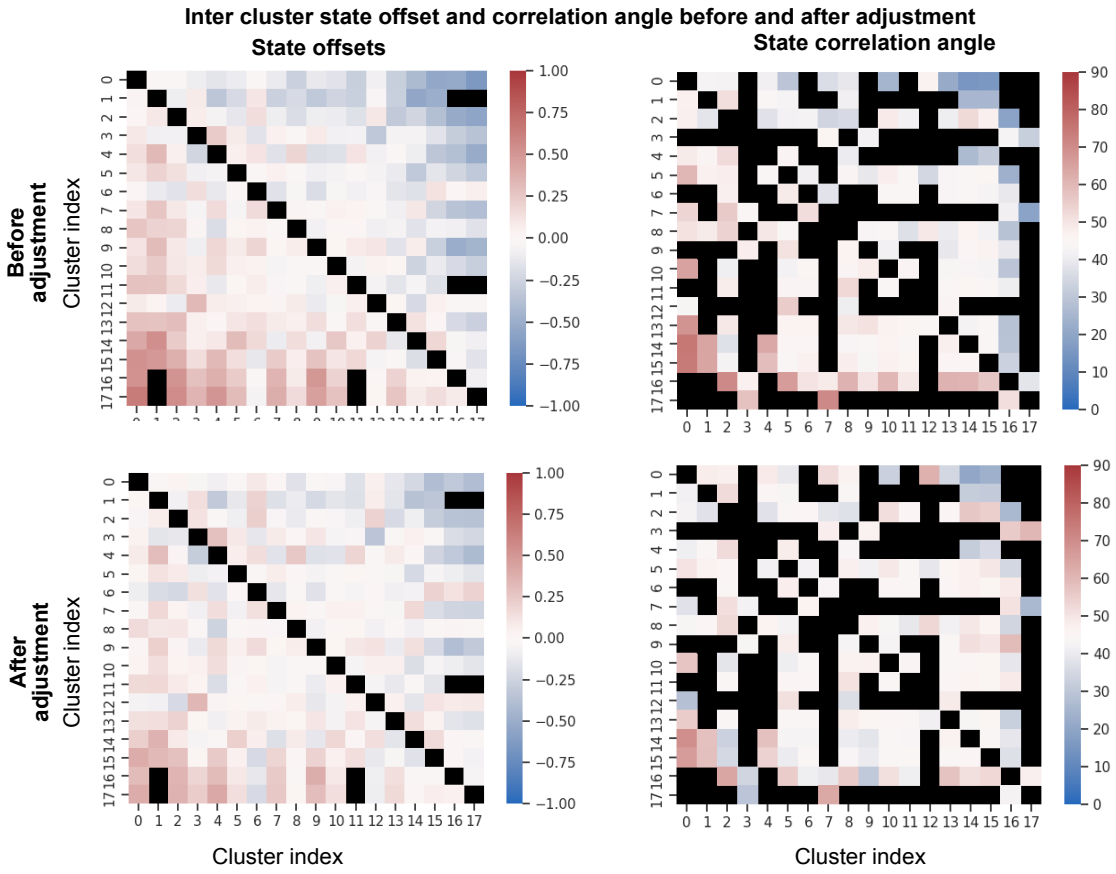


Figure 4.4: State offsets (left) and state correlation angles (right) before (top) and after (bottom) adjustment for an example stimulus condition. For the example stimulus condition there are 18 stationary correlated clusters. For the state correlation angle plots (right), angle values are only calculated (and shown) where there is a significant correlation ($p < 0.05$).

single trial population state by averaging over the predicted states of multiple correlated clusters. The algorithm acts as a first attempt, and tries to account for the problems of different PCA scales (between clusters) described in the previous section (Sec 4.6.1.1)

For a single stimulus condition, the pool of stationary correlated clusters were considered. For each pair of clusters in the pool, the mean state difference was calculated over trials on which there was a state prediction for each cluster. The value of the mean state difference between clusters is visualised in the top left panel of Fig. 4.4. It can be seen that there are clusters (i.e. cluster 17), which have positive state offsets with respect to most of the other clusters, and some with a negative offset.

For pairs of clusters for which the states predictions were significantly correlated

($p < 0.05$), the angle of the state correlation was calculated. These state correlation angles are shown in the top right panel of Fig. 4.4. It can be seen that there are some clusters which have an angle greater than 45° with the states of most other clusters (e.g. cluster 17).

The algorithm aimed to shift and rotate the states of each cluster so that their offsets between clusters were closer to zero and the correlation angles between the states of different clusters were closer to 45°. It was hoped that this would put state predictions from different clusters on a more similar scale, such that averaging would produce a better estimation of single trial state.

For each pool cluster (for a single stimulus condition), the algorithm calculated the mean of the cluster's mean state offsets with other clusters. This value is given by O_diff_i where i indexes the cluster. For the example, this value is the mean of a row in the top left panel of Fig. 4.4.

The mean angle that the cluster's state values made with the state values of other clusters was also calculated. This value is given by A_diff_i where i indexes the cluster. For the example, this value is a mean of a row in the top panel of Fig. 4.4.

The state values of each cluster were then shifted by $-O_diff_i/2$ and rotated by $-A_diff_i/2$. When the angles and offsets are recalculated for the adjusted cluster states (for the example stimulus condition), the state offsets and state correlation angles between clusters are reduced for the example (Fig. 4.4 lower).

Fig. 4.5 shows that, over all stimulus conditions, the reduction on the state offsets between clusters are closer to zero and that the state correlation angles are closer to 45°, confirming that the algorithm does increase the similarity between states.

4.6.1.3 Predicted states by trial (before and after adjustment)

For visualisation purposes, states predicted from all stationary correlated clusters were plotted against trial index before and after the adjustment procedure described in Sec 4.6.1.2. These can be seen in Vids. 5 & 6 for all of the stimulus conditions. Fig. 4.6 shows the values for an example 1Hz stimulus condition. The averaged single trial states are illustrated by the black lines.

In Vids. 5 & 6 and Fig. 4.6a it is unclear whether the mean single state values have changed significantly. Fig. 4.6b plot the mean single trial states before and after clustering for the example stimulus condition, showing that there is not a large change to the mean single trial state values. The small change in mean predictive state was observed for all the stimulus conditions (Vid. 7), showing that adjustment algorithm had little effect on predicted states.

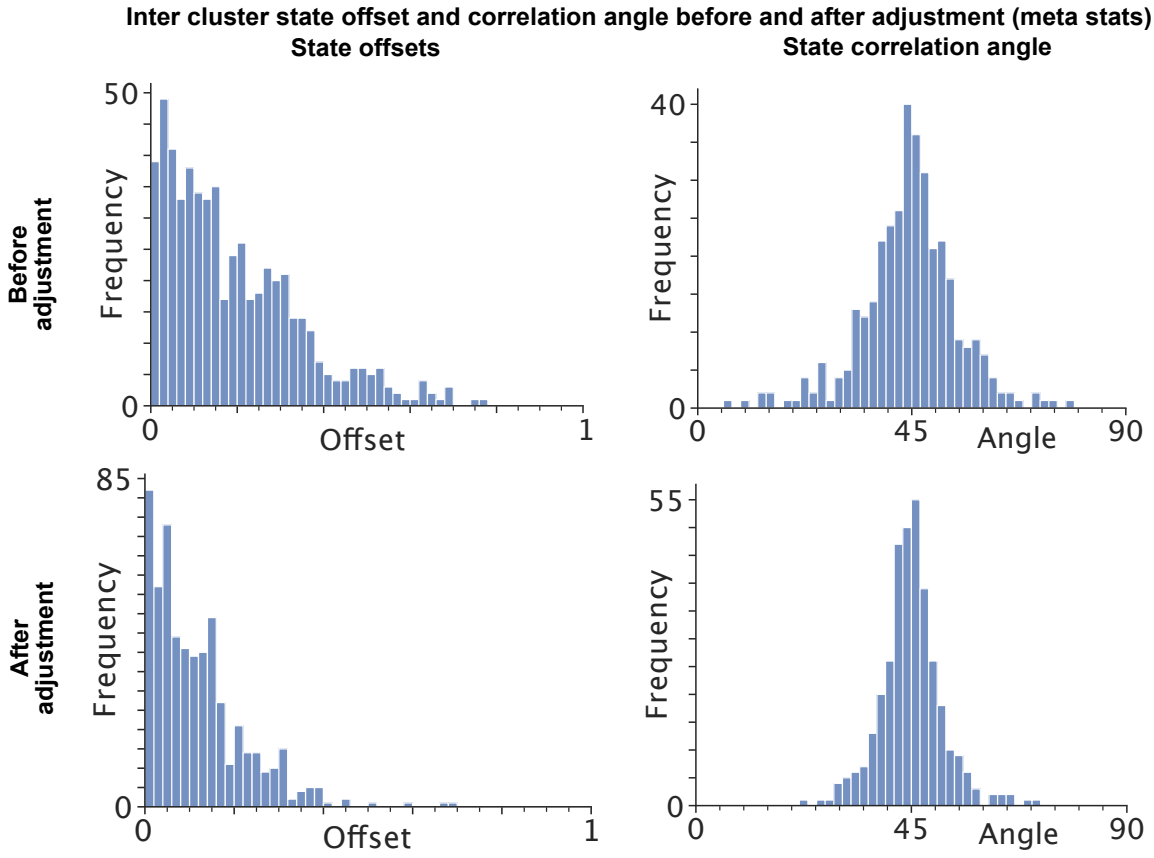


Figure 4.5: Histograms of inter cluster state offsets (left) and state correlation angles (right) before (top) and after (bottom) adjustment (meta stats over all stimulus conditions).

For each trial there are only a small number of stationary correlated clusters for which there is a spike time pair. This will be quantified in a later section but highlights the issue that single trial cortical state predictions calculated through averaging are likely to contain some error due to low number of state predictions available.

4.6.2 Stationarity of averaged cortical state

At this stage it is possible to confirm that the averaged cortical states are indeed stationary, and that the changes in cortical state are not due to adaptation or other non-stationary processes. Fig. 4.7 shows the histogram of p-values for the linear correlation between trial index and mean population state and the histogram of p-values for the lag-1 autocorrelation of the population mean state.

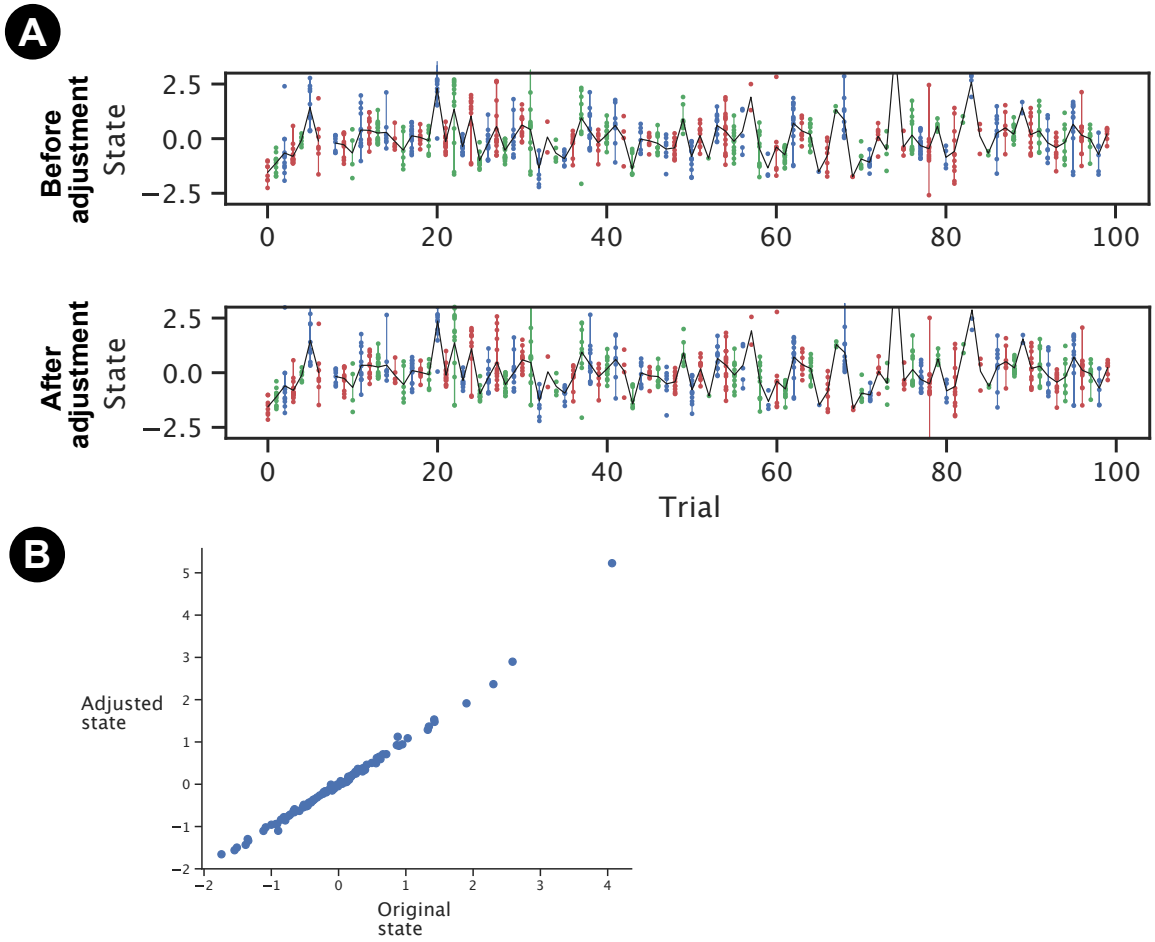


Figure 4.6: (a) Single trial states by trial index before (top) and after (bottom) adjustment procedure for 1Hz stimulus condition. Each coloured dot represents the single trial state predicted for a single stationary correlated cluster. Black dots represent the means of the states predicted from stationary correlated clusters on single trials. Black line traces through the mean values for visualisation purposes. Red, green, blue cycle is purely for visualisation purposes and has no meaning. (b) Original cluster single trial state predictions vs adjusted cluster single trial state predictions.

4.7 Methods IV: Prediction and correlation of spike times and differences from averaged cortical state

This section tests the correlative and predictive relationships between averaged population state and both spike times and spike time differences. To test the correlations of averaged population state with the first spike times of a neuron or spike differences of pair, it was necessary leave out these neurons before making the average cortical state estimations using the procedure outlined in Sec 4.6.1. In particular, for each

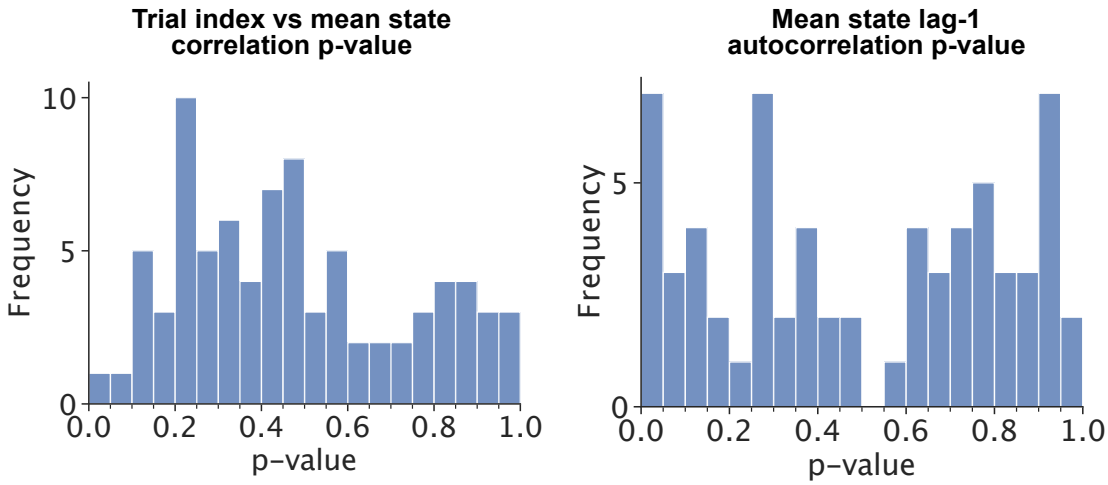


Figure 4.7: Left: Histogram of p-values for the linear correlation between trial index and mean population state, Right: Histogram of p-values for the lag-1 autocorrelation of the population mean state.

neuron or neuron pair, for which the spike times or spike time differences were to be predicted, mean single trial states were calculated by leaving out the clusters to which the neuron or neuron pair belonged. This will be referred to as the 'leave one out' population state.

The results section presents the correlations and predictions of spike times and spike differences with the leave one out population state. The necessity to leave out a number of state predictions, reduces the number of cortical state predictions which contribute to the prediction of the average state. Fig. E, F shows the numbers of state predictions, which were available for averaging on each trial. From the low number of states available, the state prediction is likely to remain noisy, and it was possible that the averaged population state might not majorly improve the predictions from states predicted from neuron pairs in Sec 4.3.

As for the tests of correlations between states, spike times and differences between disjoint pairs of neurons (Sec 4.2 & 4.3), more than 15 conjunctive trials were needed between the averaged population state and the variable of interest.

4.8 Results IV: Prediction and correlation of spike times and differences from averaged cortical state

Over trials of single stimulus conditions, the leave one out population state is correlated with the spike times (Fig. 4.8a) and spike time differences (Fig. 4.8b) of stationary correlated clusters. Fig. 4.8c (left) demonstrates a reduction in variability when spike times are conditioned on the leave one out population state. In some cases, such state-conditioned spike time variability was on a similar order to that predicted by the factor analysis model above (Fig. 4.8c, right).

A reduction in spike time difference variability is also seen when non-45° correlations are conditioned on the the leave one out population state (Fig. 4.8d, left), although variability was mostly higher than the variability predicted by the factor analysis model under RDR 3 of Chapter 4. It is important to note that there is inherent noise in the leave one out population state, which hampers spike time and spike time difference prediction. Future work may explore improved estimates of single trial state.

4.9 Methods + Results V: Visualisation of time-warped multi-neuron single spike patterns

Finally, the calculation of a single dimensional cortical state, allowed visualisation of how the timing properties of multi-neuron single spike patterns (for groups of up to 12 neurons) are affected by such a state. This provides a useful illustration of how the representation a stimulus varies depending on the single dimensional cortical state. Illustrations are presented with and without smoothing.

4.9.1 Without smoothing

The predictions of single trial cortical state allow the multi-neuron single spike pattern elicited on a single trial to be colour coded depending on the value of the cortical state. Such patterns are first be presented in this original form (Figs. 4.9 & 4.10; top). Vid. 8 shows the plots for all the stimulus conditions with at least two clusters. There are gaps in each figure, where neurons do not spike on a particular trial. Moreover, there is noise in the relationship between predicted state (colour) and the spike times of each neuron. This is to be expected from the unexplained noise predicted by the

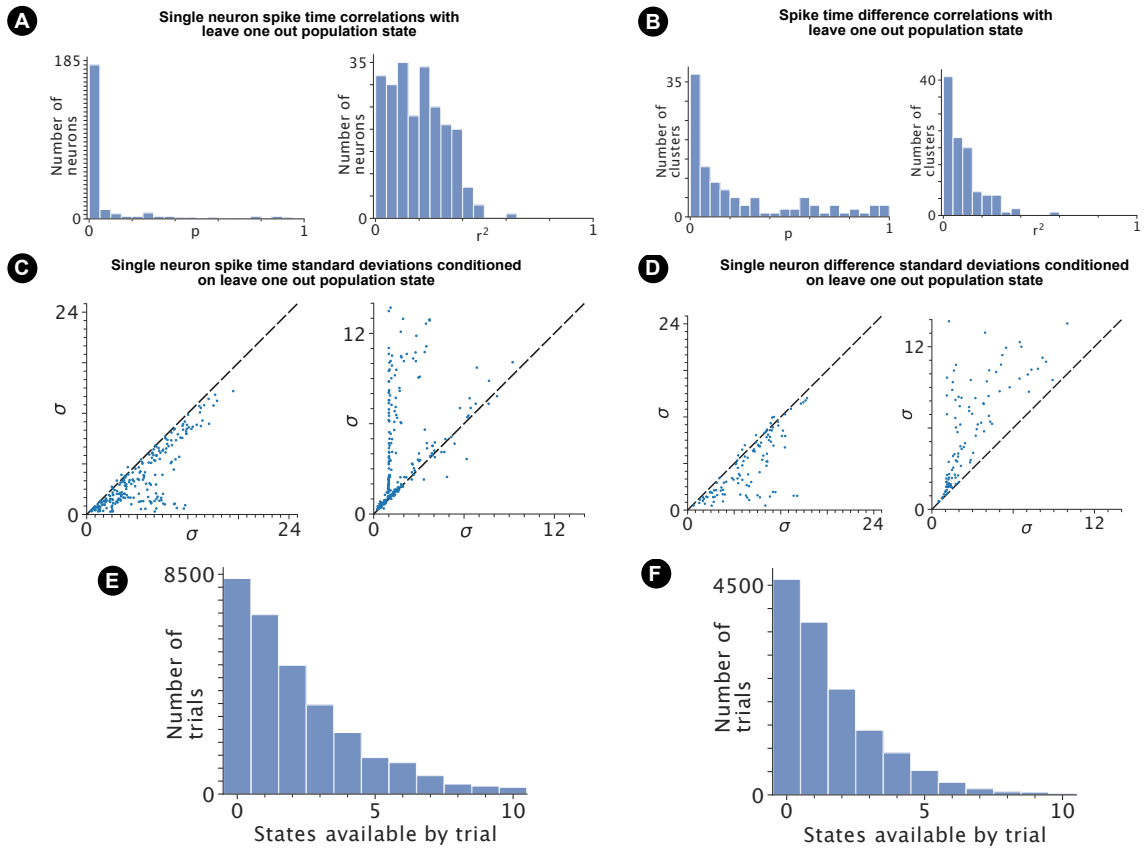


Figure 4.8: Population state-conditioned spike times and differences. For each neuron or neuron pair, for which the spike times or spike time differences were to be predicted, mean single trial states were calculated leaving out the clusters to which the neuron or neuron pair belonged. This is referred to as the 'leave one out' population state. (a) Left: p-value histogram of correlations between the leave one out population state and the spike times of stationary correlated clusters and unclustered response distributions. Right: r-value histogram of statistically significant (p-value) correlations. (b) Same as (a) but for correlations with spike time differences with non-45° correlation angles. (c) Left: variability comparison. X-axis shows the standard deviation of single neuron spike times of stationary clusters (pooled). Y-axis shows the standard deviations of spike times conditioned on the leave one out population state. Right: Comparison of predicted state-dependent spike time variability predicted by factor analysis (x-axis) vs spike time variability conditioned on the leave one out population state (y-axis). (d) Same as (c) but for correlations with spike time differences with non-45° correlation angles. (e) Histogram of the number of cluster state predictions available on a single trial, over all leave one out tests of population state correlation with single unit spike times. (f) Histogram of the number of cluster state predictions available on a single trial, over all leave one out tests of population state correlation with pairwise differences.

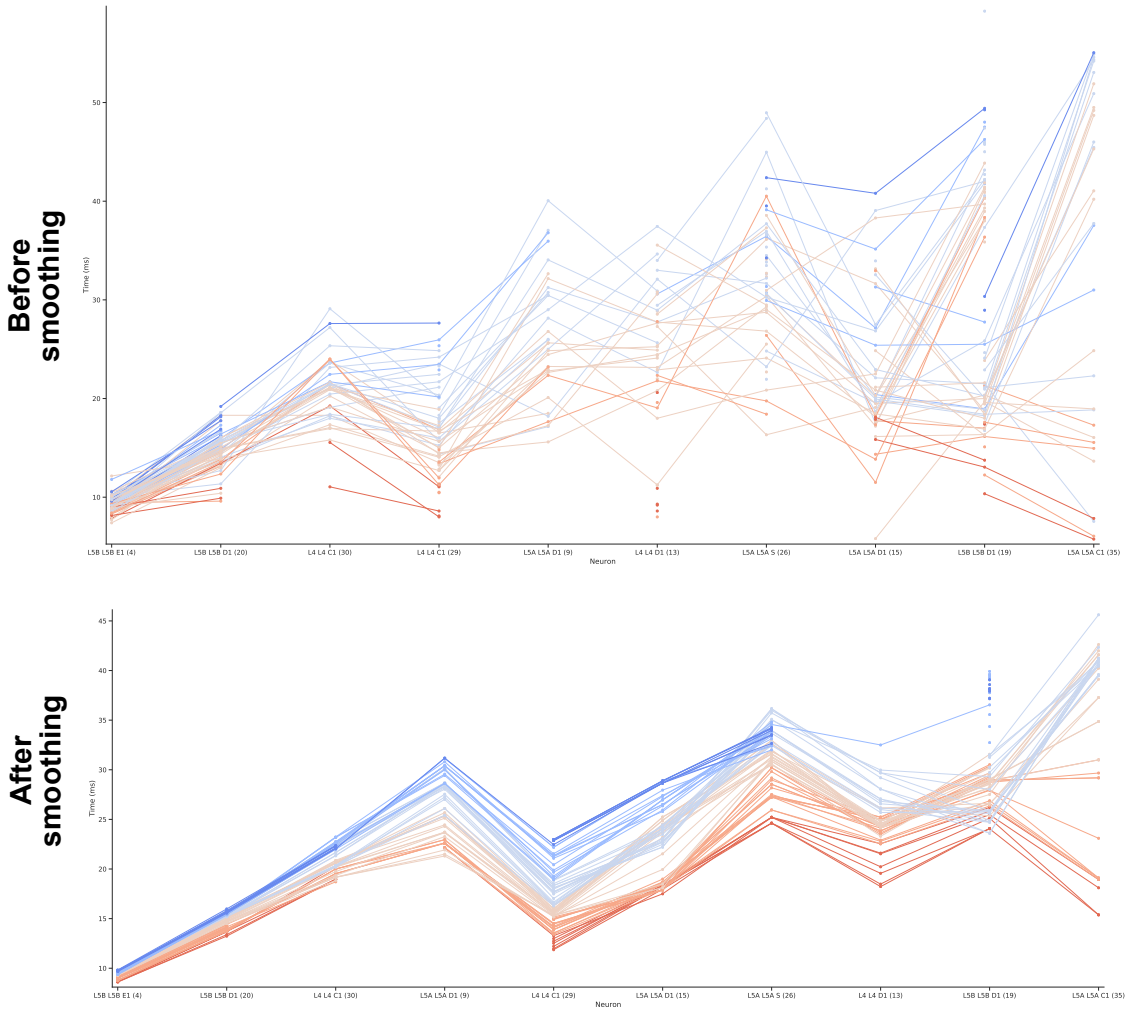


Figure 4.9: Time-warp visualisation for example stimulus condition. Top: shows the multi-neuron first spike pattern for each trial of a single stimulus condition. The colour of each multi-neuron single spike pattern represents its index when the trials are ordered by mean predicted state. Red represents trials with high predicted single trial state, and blue low. It can be seen that the spike times of the trials with high (red) single trial state spike earlier than the spike time with low (blue) single trial state. Bottom: Shows the same data after smoothing.

factor analysis model in Chapter 3. Despite this, there is a clear relationship between state and the spike times of different neurons.

4.9.2 With smoothing

Figs. 4.9 & 4.10 (bottom) show the same data after the following smoothing operation, which calculates an average spike time for each neuron and each single trial state value.

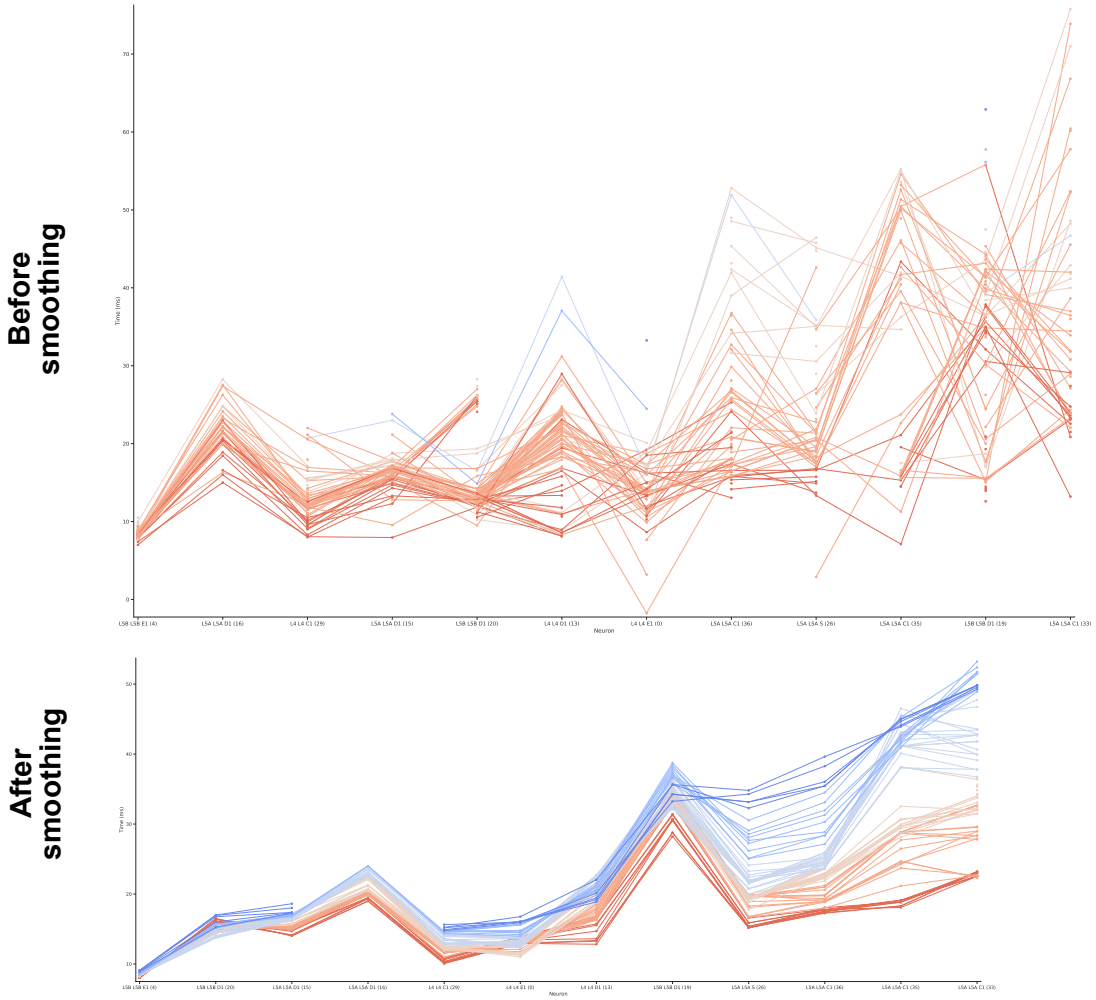


Figure 4.10: Same as Fig. 4.9 but for a different 1Hz stimulus condition.

Vid. 9 shows the plots for all the stimulus conditions with at least two clusters. The smoothing operation aimed to remove unexplained noise for the purpose of visualising state-dependent single spike patterns.

The smoothing operation is described as follows. Firstly, trials of a stimulus condition were sorted based on the predicted cortical state such that trial t_0 had the highest cortical state and t_T had the lowest cortical state (and T is the total number of trials of a stimulus condition). For each neuron, a mean spike time was calculated for each state sorted trial t_i (where i is the trial with the ith highest excitability) by averaging the spike times over 20 trials (i.e. $[i, i + 19]$).

Due to low spiking reliability, it was rare that a neuron spiked on all 20 trials $[i, i + 19]$ for any trial i. A mean spike time was only calculated if at least 7 trials in the range $[i, i + 19]$ had a spike, explaining why some of the visualisations have

some missing values for certain neurons. Future work, could use different smoothing techniques and observe the differences between stationary and made stationary plots.

4.10 Discussion

The results of this chapter further strengthen the case that spike time modulation is low-dimensional and that modulation is caused by trial-to-trial changes in shared excitability. Predictions of spike times and spike time differences from both the predicted state of single clusters and the shared population state provide a demonstration that information about spike time pattern modulation is available in spatially separated neurons. This increases the chance that downstream neurons could respond in a state-dependent manner.

The technique for adjusting the state scales of individual clusters provided limited improvement in predictability, as the state values did not change greatly. This could be due to the low number of single cluster state predictions per trial, however. It would be interesting to see if the adjustment technique improved with greater numbers of single cluster state predictions per trial.

The analyses developed in this chapter are general for the source of the single trial state prediction. In future it would be interesting to use the same analyses with different state prediction techniques, for example, those which estimate the low-dimensional single trial trajectories of population activity.

Chapter 5

Comparison of in silico and in vivo responses for a detailed biophysical cortical column model

5.1 Introduction

The work of this chapter initially aimed to test whether adaptation-independent time-warped spike time patterns, as characterised in Chapters 3 & 4, existed in the Blue Brain Project's (BBP) biophysical model of a small volume of tissue from layers 1-6 of the hind limb somatosensory cortex of juvenile 2-week-old Wistar rat [175]. It was hoped that if the model exhibited the same time-warped spike time patterns, then the model could be used to explore latent processes governing the time-warp effect and to make further predictions, for example about how time-warped spike time patterns vary with changes in the stimulus.

It was first necessary to compare in vivo and in silico neuronal response properties at the population-level (i.e. by layer and E/I-type), to ensure similarity between high-level in vivo and in silico response properties. The spiking activity of the in silico model and in vivo dataset analysed in Chapters 3 & 4 had in fact previously been compared, in the initial publication on the model [175]. Particularly, figures produced from in silico spiking activity were visually compared to published figures for the in vivo dataset [232]. In this chapter, however, direct access to the same in vivo dataset [232] (analysed in Chapters 3 & 4) allowed direct comparison of in vivo and in silico activity.

The chapter outlines a set of measures for directly comparing in vivo and in silico neural activity. The measures consider the possibility of sampling biases likely to be introduced by in vivo extracellular recording and spike sorting. The chapter finds

discrepancies between in vivo and in silico spiking activities under conditions similar to those used in the original model paper [175]. The measures also enable analysis of how in silico parameters affect in silico spiking activity and its similarity to in vivo spiking activity.

A set of model parameters are determined, which produce in silico spiking activity closely resembling in vivo activity (as defined by particular measures). Interestingly, the model produced the most 'in vivo-like' activity when an important biophysical model parameter, extracellular calcium concentration ($[Ca^{2+}]_o$) was brought within its in vivo range, reported to lie in the range 0.9-1.1 mM [142, 179, 11, 24]. This parameter had previously been set outside of its in vivo range for studies of in vivo-like activity [175].

The model was then tested for the time-warp phenomena for this set of model parameters. The time-warp phenomena was not found for this set of model parameters, however, suggesting that the model may need further refinement to reproduce the time-warped spike time phenomena. The methods developed in this chapter, however, can be used to test and validate the spiking activity of future cortical models, and to aid in highlighting where further model refinement may be required.

The presented comparison methods represent the state of the analysis tools in April 2021 when the analysis of the original 2015 model [175] was completed. Since then, the author has further advanced the tools and analyses for a currently unpublished model. A brief summary of these tools and analyses is given in the Discussion of this chapter (Secs 5.9.2 & 5.9.3).

5.1.1 Simulations as a validation of biophysical models

Biophysical computational models provide the opportunity to compose and cross-validate the large quantity of morphological and physiological data produced by modern neuroscience [175, 23]. In doing so, such models can highlight and make predictions where data is missing [223]. In silico simulations of biophysical computational models provide the opportunity to bring such data 'to life'.

Close resemblance between in silico and in vivo spiking activity provides an important validation of such a model, which would strengthen the argument for the model's use as an investigative tool. Initiatives now exist to reconstruct and simulate ever larger parts of the cortex [175, 23, 277], including of the barrel cortex [69, 68]. Section 5.2 reviews the cortical column model studied in this chapter. Relevant other models are discussed in Section 5.3.

In vivo and in silico activity can be compared in the stationary case (i.e. with no mechanisms of adaptation or plasticity). Resemblance of in silico neuron group (i.e. layer and E/I-type) spiking statistics with in vivo offers an important benchmark for in silico models, which large-scale modelling efforts continue to work towards. Beyond stationary response properties, an in silico model implemented with mechanisms of plasticity or adaptation would hope to capture how stimulus-responses change over multiple stimulus presentations or longer time periods.

For example, the in silico model has recently been used to study the dependence of stimulus specific adaptation on synaptic depression and spike frequency adaptation [10]. For such studies, it is important that the initial activity of the model closely resembles in vivo activity at the start of the non-stationary process. The methods described here are designed to quantify resemblance to in vivo spiking activity at a single point in time, for example at the start of a non-stationary processes or over multiple time points of an experiment.

5.2 The Blue Brain Project (BBP) somatosensory column model [175]

This section describes the BBP's model of a small volume of tissue from layers 1-6 of the hind limb somatosensory cortex of juvenile 2-week-old Wistar rat [175], which is the model studied in this chapter. The model was validated against experimental datasets not used in the reconstruction, demonstrating that sparse complimentary datasets can be integrated to obtain dense maps of neural micro-circuitry, without needing to measure every biological parameter.

Model reconstruction and simulations are described in the following subsections. The description of the model's reconstruction, given here, acts as a description of the model, and was also useful to improving the author's understanding of the model. Section 5.2.2 describes the original simulations made using the in silico model, including comparisons of in silico and in vivo spiking activity. Readers who are familiar with the model may wish to proceed to Section 5.3.

5.2.1 Reconstruction

Morphological Diversity of Neocortical Neurons

Over 14,000 cells were recorded in in vitro brain slices using patch-clamp electrodes from all six layers of the somatosensory cortex of P14 (2-week-old) male Wistar (Han)

rats. 2,052 cells were classified into 55 established layer-specific morphological types (m-types) based on their dendritic and axonal arbors. L2/3 was considered as a single layer. An initial pool of 1,009 cells were reconstructed and their classification validated. Where there were no morphological reconstructions for rare layer-specific m-types, exemplars of the same morphology were used from other layers. Arbors damaged by slicing were repaired using other m-type exemplars and validated. A larger pool of exemplars was created by jittering the branch angles and section lengths of m-type clones, and validating distributions of morphometric feature against the original pool.

Reconstructing Neuron Densities, Ratios, and Composition

The average neocortex height in six animals was $2082 \pm 80\mu m$ (mean±SD, $n = 6$). Layer thicknesses were determined by transitions in cell densities and soma sizes in NeuN-stained tissue blocks. Layer-wise percentages of excitatory ($87 \pm 1\%$) and inhibitory ($13\%\pm1\%$) neurons were established by DAI, NeuN and GABA tissue block staining. Layer-wise proportions of m-types within E and I groups were based on the frequencies of classified neurons. E-I and m-type compositions were broadly consistent with previous reports. Cells counts in NeuN-stained tissue blocks determined a mean cell density of $108,662 \pm 2,754\ neurons/mm^3$ (mean ± SEM, n = 6), comparable to rat barrel cortex.

Unlike the barrel cortex, the hind-limb somatosensory cortex has no anatomically defined horizontal columnar organisation [123, 173]. The volume's radius was determined as $210\mu m$ (similar to rodent barrel cortex) by placing reconstructed neurons in a cylindrical volume and determining the radius where dendrite density saturated. The microcircuit was then defined as a hexagonal prism with cross-sectional area equal the calculated radius. 35 instantiations were made using measurements obtained from five animals Bio1-Bio5 (7 instantiations per animal) and 7 instantiations were made using data averaged from the 5 animals (BioM).

Positioning Morphologically Reconstructed Neurons

Neuron positions were drawn from 2D Gaussians around the centers of 310 minicolumns, consistent with reports of weak minicolumn organisation [193]. Vertical positions were chosen randomly within each layer (checking that somata did not overlap). Morphologies were selected for each position from the top 8% of morphologies, scored by their match to typical patterns of arborisation within and between layers. Comparison of in silico and in vitro immunohistochemical staining for 7 markers was

used to validate layer-dependent patterns and the number of neurons at different depths.

Morpho-Electrical Composition

Over 3,900 neurons from all layers were classified into 11 electrical types (e-types; 10 inhibitory, 1 excitatory) according to the Petilla convention [1] using quantified features of the neuronal response to step current pulses and clustering validation. No significant bursting behaviour was observed for excitatory m-types so all excitatory m-types were classified as continuous adapting (cAD) neurons. Consistent with previous observations, most inhibitory m-types expressed multiple e-types [1]. In total there were 207 morpho-electrical types (me-types). The relative proportions of e-types for each inhibitory m-type were used to determine the relative proportion of me-types for each inhibitory m-type (layer-wise when sample size was sufficient). Layer-wise me-type neuron counts were calculated by combining relative proportions of me-types with neuron densities.

Digital Reconstruction of Connectivity

Synaptic connectivity was reconstructed using an approach developed in a companion study [223]. The algorithm yielded plausible multi-synapse connections, reproduced the connectivity (numbers and locations [116]) of all connections that had been studied experimentally, reproduced numerous other experimental findings and made further predictions. Across the 35 Bio1-5 reconstructions, the predicted mean number of synapses per connection was 4.5 ± 0.1 (3.6 for excitatory connections, 13.9 for inhibitory connections). Each neuron innervated 255 ± 13 other neurons, forming an average of $1,145 \pm 75$ synapses with neurons in the microcircuit. An average of 36.7 ± 4.2 million synapses were formed in each microcircuit (27.0 ± 2.9 million excitatory, 9.7 ± 1.5 million inhibitory), made up of 8.1 ± 0.9 million connections.

The algorithm predicted that afferent fibres from beyond the microcircuit form a further 147 ± 4 million synapses (assuming a density of 1.1 excitatory synapses/μm on dendrites, based on observed spine densities), such that $20\% \pm 2\%$ of synapses are formed by neurons in the mircrocircuit, consistent with previous estimates.

Reconstructing Neuronal Physiology

Neurons were simulated as multicompartmental conductance-based models with intracellular Ca^{2+} dynamics. Neurons had separate axon initial segments (AIS), somata, basal dendrites and apical dendrites. Each AIS was represented by two 30mm

sections, with diameters obtained from reconstructed morphologies. Axons had an assumed conduction velocity of $0.3m/s$ such that a delay could be computed based on the path distance between AIS and synapse location on the axon. Interneurons contained only one dendritic region.

Neurons had 13 Hodgkin-Huxley (HH)-type models of known active ion channel types. A multi-objective optimisation algorithm [66] computed ion channel conduction densities that best reproduced features such as spike amplitudes and widths, spike frequency, and changes in frequency. An automated workflow generated a pool of 121,231 neuron models containing exemplars of all 207-me types, with variation comparable to the biological variation. Particularly, the multi-objective algorithm generated a solution for one given morphology, which was then transplanted to all morphologies of the same m-type. The algorithm then accepted or rejected the combination, based on how well it reproduced the spiking features.

Synaptic Physiology Reconstruction

Synaptic physiology was predicted from published paired-recording data and additionally measured synaptic properties (i.e. conductances, postsynaptic potentials, EPSPs/IPSPs, latencies, rise and decay times, failures, release probabilities). Synaptic connections were classified as either facilitating (E1 and I1), depressing (E2 and I2), or pseudo-linear (E3 and I3) s-types (see [175] for original references). S-types for broad classes of connections between me-types were predicted based on five rules (see [175] for supporting references):

1. Pyramidal-to-pyramidal connections are always depressing (E2).
2. Pyramidal-to-interneuron connections are also depressing (E2).
3. Facilitation from inhibitory neurons is around two times stronger than from excitatory neurons.
4. Synaptic dynamics are preserved across layers for all me-type-specific connections.
5. Any remaining connections belong to the most common s-type (type 2; E2 or I2).

Synaptic conductances obtained between specific pairs of m-types, yielded in silico PSPs which were systematically lower than in vitro somatic PSPs. Synaptic conductances were therefore adjusted until in silico PSPs matched experimental levels, suggesting that reported conductances were about 3 and 2-fold too low for excitatory and inhibitory connections respectively. Average corrected synaptic conductances calculated for broader classes of synaptic connections (e.g., E-E, E-I, I-I, I-E) were used where data was missing. Resulting in silico PSP amplitudes were validated experimentally for ten connection types not used to calculate the conductances (regression, r=0.6).

The authors reported that the derived synaptic dynamics and quantal conductances compared well with previous reports, generating a first prediction of mean synaptic conductances for all m-type connection types. Quantal synaptic conductances for single synapses were drawn from truncated normal distributions around these means. In silico paired recordings of m-type-to-m-type connections were comparable to previously published paired recordings in vitro. On average, neurons in the model had excitatory and inhibitory conductances of $\sim 750nS$ and $\sim 250nS$ respectively (total from afferent synapses). Excitatory and inhibitory synapses had a predicted average quantal conductance of $0.85 \pm 0.44nS$ (corresponding to $\sim$150 AMPA and $\sim$20 NMDA receptors) and $0.84 \pm 0.29nS$ for inhibitory synapses (corresponding to $\sim$40 GABA-A receptors).

Synaptic transmission stochasticity and neurotransmitter release probability

Synaptic transmission stochasticity extended a previously reported model [88], and was validated by comparing the coefficient of variation of first PSPs. Neurotransmitter release probability was modified consistent with experimental data for the specific sensitivities of different s-types to changes in extracellular calcium ion concentration $[Ca^{2+}]_o$.

Thalamic stimulation

A 'mesocircuit' consisting of a central microcircuit surrounded by 6 additional microcircuits was used to explore evoked activity. Connectivity was established for the mesocircuit so there were no anatomical borders between microcircuits. Input from the thalamus to the central microcircuit was approximated using experimental data for the ventral posteromedial (VPM) thalamic input to the barrel region of somatosensory cortex [49, 96, 184].

A single fiber was instantiated and centered in each minicolumn with a horizontal spread [184]. The synapses that each thalamic fiber formed onto different m-types was determined by layer-by-layer bouton density profiles [184], measurements of the mean number of synapses per thalamic connection in L4 [96, 9], and the multi-synapse principal described above. This reproduced the number of synapses formed on L4PCs [96, 9] and made predictions for the average number of synapses on other neurons (e.g. ∼12 for L5 pyramidal neurons). This predicted that each thalamic fiber innervates 903±66 neurons (mean ± SD; n=100 fibers; 775±57 excitatory and 83±11 inhibitory neurons). This gave an average of 8.1 ± 4.2 synapses/connection, totalling ∼2 million synapses in the central microcircuit (∼1% of synapses across all layers).

Thalamocortical synaptic transmission was modelled using data of in vitro synaptic dynamics [96, 9], the excitatory-to-excitatory conductances derived above (similar to L4E, E2 s-type) and $[Ca^{2+}]_o$ dependency as for other excitatory connections (described in the previous subsection). In silico PSPs in L4 and L5 PCs in low-$[Ca^{2+}]_o$ conditions were validated against in vivo reports [38, 49].

5.2.2 Simulations

Simulating spontaneous activity

Simulations explored neuronal activity of the microcircuit under both in vitro and in vivo conditions. This chapter focuses on modelling of in vivo conditions but insight from modelling of the in vitro condition is relevant for studying correlated activity. To allow comparison to in vitro experiments, a virtual brain slice was created with width equal to that of 7 microcircuits. Both in the standalone microcircuit and virtual brain slice, tonic depolarisation reproduced oscillatory bursts ($\sim 1Hz$) comparable to previous in vitro experiments [165, 246]

In vitro experiments are typically performed at 2mM $[Ca^{2+}]_o$ but the in vivo level is reported to lie in the range 0.9-1.1mM [142, 179, 11, 24], increasing to 1.2-1.3mM to produce oscillatory cycles during the transition from wakefulness to sleep [11]. As an approximation of in vivo conditions, the behaviour of the circuit at lower $[Ca^{2+}]_o$ levels was explored; slow oscillatory bursting disappeared and the neuronal activity became asynchronous and irregular. This was validated by multi-electrode array recordings in vitro in high and low $[Ca^{2+}]_o$ conditions.

Following a literature search, the authors found that the release probabilities of different synapse-types (particularly inhibitory and excitatory) had different sensitivities to changes in extracellular $[Ca^{2+}]_o$. After implementing these different $[Ca^{2+}]_o$

sensitivities, the authors found a spectrum of network states ranging from largely synchronous to asynchronous neural activity (the synchronous-asynchronous [SA] spectrum), which depended on $[Ca^{2+}]_\circ$ concentration. A decrease of just $\sim 0.15mM$ led to a sharp decrease in correlated spiking and a transition in the SA spectrum. The point of transition varied slightly between model instantiations. These characterisations suggest that the position of the microcircuit on the SA spectrum will be important to the possible occurrence of the time-warped spike time patterns.

The authors suggested that the differential $[Ca^{2+}]_\circ$ sensitivities of inhibitory and excitatory synapse types was likely to underlie the transition; supported by demonstration that a change in $[Ca^{2+}]_\circ$ from 2mM to 1.3mM alters the ratio between excitatory and inhibitory synaptic PSPs by ~ 3.5, in favour of inhibition. Other mechanisms that change the excitatory-inhibitory balance may also affect the SA spectrum transition. In silico knockout experiments indeed showed that different layers, neurons, and connections affect the position of the microcircuit along the SA spectrum. Finally, correlations in neuronal activity were reported slightly more within synaptically coupled assemblies previously found experimentally [217, 223], than in randomly sampled neurons.

Excitatory conductances were highly correlated and are cancelled out by anti-correlated inhibitory conductances, as suggested previously [17, 205]. Spiking in the mesocircuit was correlated with momentary imbalances between excitatory and inhibitory conductances lasting $< 10ms$ and the timing of spikes could be predicted by this imbalance. The precision with which the imbalances drove spiking fell with distance from the SA transition in either direction. The authors reported that correlation was strong but broad in the synchronous state, supporting a temporally imprecise increase in spiking suitable for a rate code, whereas correlation was sharp but too weak to drive spiking in the asynchronous state, supporting a population code based on highly correlated activity.

As reported in vivo [206], for L5 and L6, some neurons, predominantly pyramidal neurons, were tightly correlated to the population (choristers), whilst others, predominantly interneurons, displayed diversity of spiking patterns with lower correlation to the population than expected by chance (soloists). Pyramidal cells tended to be soloists when their spontaneous firing rate is high, the ratio of excitatory to inhibitory synaptic innervation was high, and most of the innervating synapses were close to the soma. Shifting the network into the synchronous regime, increased the number of choristers, whereas shifting the network into the asynchronous regime reduced choristers and soloists.

5.2.2.1 Initial thalamic stimulation tests

Before comparisons were made with stimulus-evoked in vivo neuronal responses, the general properties of in silico responses to thalamic stimulation were tested. With $[Ca^{2+}]_o$ = 1.25 mM and moderate depolarization (referred to as in-vivo-like conditions by the authors), the main response to thalamic stimulation was in L4 to L6, with most m-types responding. Under in-vivo-like conditions, a single synchronous spike was used to activate progressively more thalamic fibers producing a graded response, with 20-30 Hz oscillations emerging in the lower layers for higher numbers of fiber stimulations, whilst activity remained localised. Under in-vitro-like conditions (2.0mM), stimulating as few as four thalamic fibers produced all-or-none behaviour across the whole mesocircuit.

Correlations increased with increasing $[Ca^{2+}]_o$ levels. Over time, the spread of activity accelerated in the synchronous regime, whilst it was constant before falling to zero in the asynchronous regime. A series of simulations varying depolarisation and levels of $[Ca^{2+}]_o$ were used to more completely demarcate between four distinct activity regimes: evoked regenerative (ER), spontaneous regenerative (SR), evoked non-regenerative (EN), and spontaneous non-regenerative (SN). The level of $[Ca^{2+}]_o$ at the transition between the synchronous and asynchronous state for each reconstruction Bio1-Bio5 ranged from 1.23 to 1.31mM $[Ca^{2+}]_o$.

Stimulus response curves measured over 10ms of thalamic stimulation, shifted from a linear to a sharp sigmoidal shape, with increasing $[Ca^{2+}]_o$. The simulations predicted that, at the average $[Ca^{2+}]_o$ reported in the awake state [142, 179, 11, 24], responses will be graded and spatially restricted.

For a thalamic stimulation roughly comparable to a single whisker deflection (stimulation of a cluster of 60 thalamic fibres), trial-to-trial spiking variability was examined across neurons, neuron-types, layers and model instantiations. Neuronal responses showed varying degrees of trial-to-trial variability, as expected by the model's stochastic mechanisms of synaptic transmission, spontaneous release and (some) ion channels. Each neuron-type displayed a characteristic delay to first spike response. Variance in response times decreased with cortical depth for excitatory and inhibitory neurons.

Importantly, in all layers, trial-to-trial variability was lower than the variability between individual neurons of the same type in single trials, demonstrating spatiotemporal structure in the stimulus population response. Setting the $[Ca^{2+}]_o$ level such that each Bio1-Bio5 reconstruction was shifted to the same point along the spectrum

relative to the transition, reduced the response latency variance for specific m-types between Bio1-Bio5 reconstructions.

Precise repeating spike triplet structures, which had previously been found above chance in spontaneous activity in vivo [166], were also found in the analysis of evoked responses to thalamic stimulation of L5 neurons. At lower $[Ca^{2+}]_o$ levels (1.0mM) these triplet structures did not occur above chance.

Comparison to in vivo neuronal responses

The authors next compared in silico and in vivo stimulus-evoked responses reported for previous analysis of the dataset studied in Chapter 3. Particularly, in silico responses to single pulses activating 60 thalamic fibres projecting to the center of the digital microcircuit, were compared to neuronal responses reported for single-whisker deflections in the corresponding principal column of the barrel cortex of anaesthetized adult rats [232]. As the authors did not have direct access to the in vivo dataset, comparisons were made with figures from the original in vivo paper. The authors reported the following comparisons:

1. *'As shown in Figure 17A1, the response to the stimulation displays cell-type-specific diversity that compares reasonably well with the results reported in Figure 3A of the Reyes-Puerta et al. study [232], with the exception of the OFF response, which is not as prominent.'*

   The in vivo and in silico figures referred to in the quote are shown here in Fig. 5.2.2.1.

2. *'The general distribution of responses for excitatory and inhibitory cells is also comparable (Figure 17A2 in silico versus Figure 3B in vivo), though again with fewer OFF responses.'*

   The in vivo and in silico figures referred to in the quote are shown here in Fig. 5.2.

3. *'As in Reyes-Puerta et al. [232], most responses occurred within 10-20 ms of the stimulus and were generally led by inhibitory cells and, more specifically, by inhibitory cells in L4 and L5 (Figure 17A3 in silico versus Figure 4B1 in vivo).'*

   The in vivo and in silico figures referred to in the quote are shown here in Fig. 5.2.

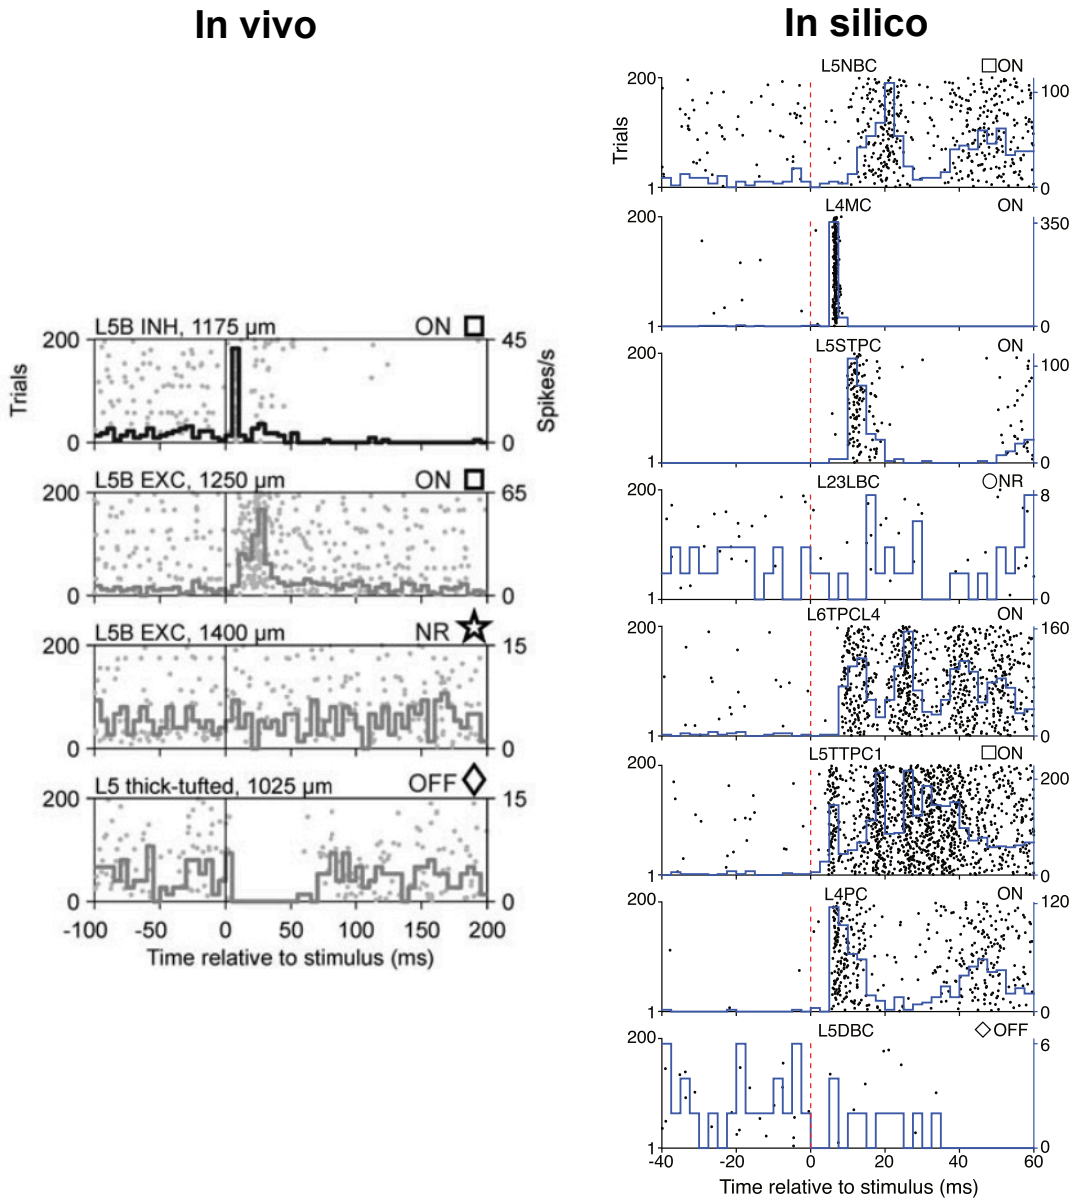


Figure 5.1: Figures used by [175] to compare in vivo ([232]) and in silico [175] stimulus-evoked responses. Left: in vivo rasters of 4 sample spike sorted neurons (from [232]): 'Cellular response types to single-whisker stimulation in barrel cortex. Typical responses of one inhibitory (INH) and 3 excitatory (EXC) neurons recorded in parallel. Each horizontal subplot represents the activity of an individual cell containing the raster plot aligned to whisker stimulation and the PSTH. Trials in the raster plot are presented in the same sequential order as recorded. Note that neurons upon whisker stimulation decreased their firing rate (FR) (ON cells), showed no change in FR (NR) or decreased their FR (OFF cells).' Reprinted form [232] by permission of Oxford University Press. Right: Corresponding in silico plot from [175]. Reprinted from [175], with permission from Elsevier.

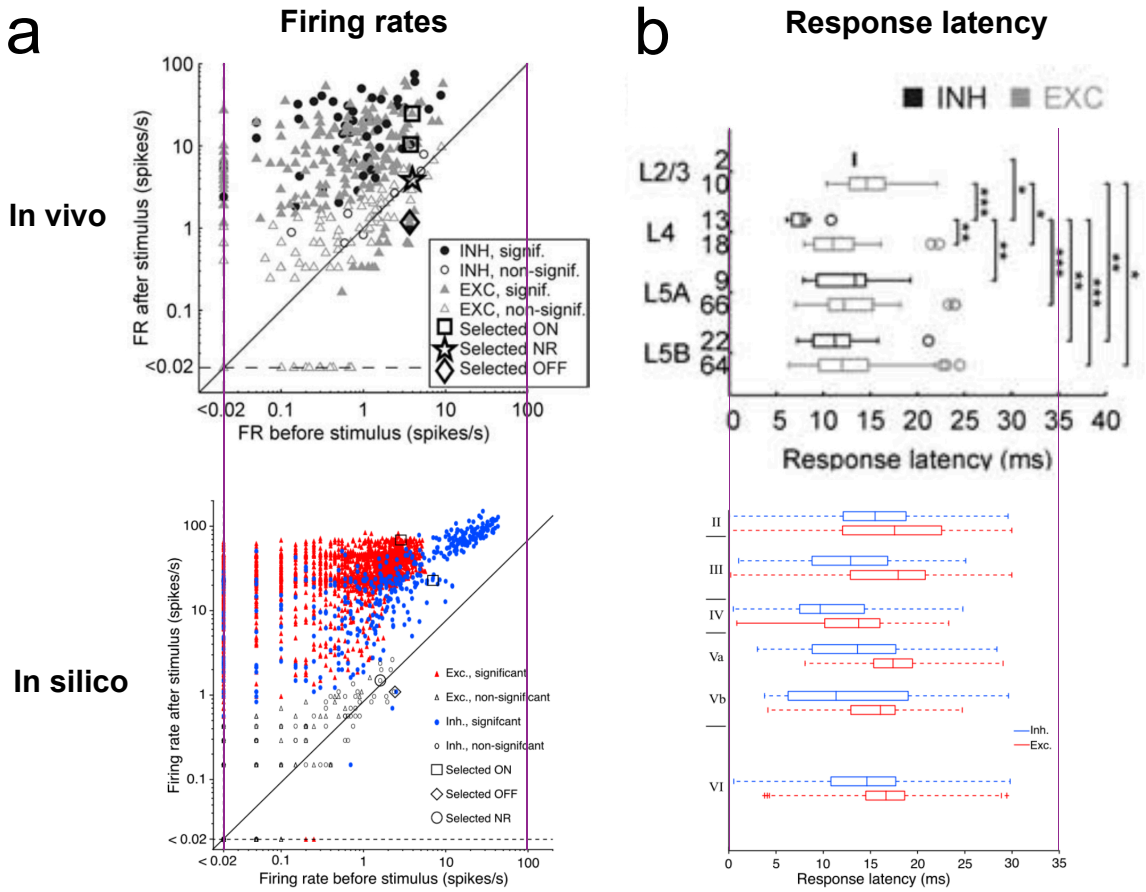


Figure 5.2: Figures used by [175] to compare in vivo ([232]) and in silico [175] stimulus-evoked responses. (a) Top: in vivo (from [232]): 'Comparison of mean FR of EXC (triangles) and INH neurons (circles) before versus after whisker stimulation, plotted in logarithmic scale. Pre-stimulus time is defined as the 100 ms period before whisker stimulation and post-stimulus time as 0-30 ms afterward. Empty symbols represent neurons showing no significantly different activity in both periods, filled symbols neurons showing significantly different activity ($p < 0.05$). The neurons displayed in panel A are marked by their respective symbols.' Bottom: corresponding in silico plot from [175]. (b) Top: in vivo (from [232]): 'Statistical analysis of first-spike latencies for the different putative neuronal groups. Each box plot represents median, interquartile, and range of latencies across all recorded neurons. Empty circles represent outliers.' Bottom: Corresponding in silico plot from [175]. In vivo figures reprinted form [232] by permission of Oxford University Press. In silico figures reprinted from [175], with permission from Elsevier.

The visual comparisons between in vivo and in silico figures shown here in Fig. 5.2 & 5.2.2.1 were given as the main validation of in silico responses in [175]. Comparison of the in vivo and in silico plots in Fig. 5.2, however, shows clear differences between the two cases. As the authors of the in silico model did not have access to the in vivo

dataset, they were unable to compare the two cases directly.

Functional Implications

Many L5 pyramidal neurons could discriminate between inputs separated by $150\mu m$ or more. A few could discriminate between stimuli with separations as small as $50\mu m$ (i.e., approximately two minicolumns apart). Generally, timing-based discrimination was much stronger than rate-based discrimination. Interestingly, discrimination was strongest at $[Ca^{2+}]_\circ$ levels close to the transition between the synchronous and asynchronous regimes. The authors noticed spatial asymmetries in the discriminatory power of neurons in the microcircuit, which was reduced for the mesocircuit. The authors reported strong variation in discrimination power across instantiations.

In reconstructions smaller than the anatomically defined microcircuit, the transition between the synchronous and asynchronous states occurred at high levels of $[Ca^{2+}]_\circ$. The transition fell sharply with increasing size of the reconstruction, reaching a plateau in reconstructions larger than ∼300 minicolumns. In reconstructions with 1,000 minicolumns, the spectrum of states did not change qualitatively.

Similarity in the time course of spontaneous firing rates in clusters of ∼10 minicolumns was high but progressively decreased with increases in the size of the reconstruction, plateauing at distances on the order of the diameter of the anatomically defined microcircuit ($\sim 202\mu m$). Correlated activity within a central set of ∼50 minicolumns decreased exponentially as the surrounding network increased in size, bottoming out in microcircuits of ∼300 minicolumns and larger.

Discussion

The authors suggested the following ways in which the reconstruction could be improved:

> *'Additional cell type markers would improve the accuracy of the morphological composition, saturated EM reconstructions could be used to further validate the derived connectivity, more experiments reporting combined voltage and current measurements for synaptic responses will strengthen conclusions on quantal conductances and connection-specific synaptic dynamics, and further characterization of the sensitivity of different synapses to* $[Ca^{2+}]_\circ$ *may allow more accurate demarcation of the transitions between different activity states and may allow a more precise determination of the role played by each neuron and synapse type in maintaining and shifting regimes.'*

The authors also admit that the model:

'omits many important details of microcircuit structure and function, such as gap junctions, receptors, glia, vasculature, neuromodulation, plasticity, and homeostasis. Furthermore, it represents a snapshot of just one brain region, in one strain of male rat, at a young age. This limits the generality of the conclusions that can be drawn. For instance, in animals of the age used for the study, dendritic morphologies have already matured to adult levels (Larkman, 1991a; Romand et al., 2011), but the ascending axons may not be fully represented and are certainly not completely mature (Romand et al., 2011).

Ensuing studies

Following the initial publication describing the reconstruction and simulation of the model [175], a number of ensuing studies have followed, which extend the model further and explore the model's properties. These studies provide a range of avenues for future research and include:

- Exploring the structure of low-dimensional activity manifolds and their relationship to underlying connectivity [225].
- Demonstrating stimulus-specific adaptation (SSA) [10] under a model of short-term synaptic depression (i.e. neutransmitter depletion) [276].
- Showing that higher-order structure imposed by morphological diversity within neuronal types contributes to a small world topology and impacts emergent cortical activity [201].
- Implementation of in silico voltage sensitive dye imaging (VSDI) to study the relationship between VSDI imaging of subthreshold membrane potential fluctuations and population spiking activity [198] in silico. The authors concluded that VSDI signal is driven primarily by neurons in layers 2/3 and 5. Moreover, membrane potential fluctuations are anticorrelated with population firing rates, suggesting that VSDI primarily reflects inhibitory responses to ongoing excitatory dynamics.
- Demonstration that noisy and chaotic network dynamics caused by stochastic neurotransmitter release can be transiently overcome by weak feed-forward thalamocortical inputs, supporting reliable spike times with millisecond precision [200].

- Demonstration of the relationship between the model's structural and functional algebraic topology [224].

5.3 Other biophysical cortical models

This section describes several biophysical cortical models which are of relevance to the work of this chapter.

5.3.1 Allen Institute mouse primary visual cortex model [23]

This section discusses a new model of the mouse primary visual cortex recently published by the Allen Institute [23]. The model has a radius of $845\mu m$ and contains $\sim 230,000$ model neurons of 17 neuronal types. The model can be simulated either as a biophysically detailed model or point neuron model. Input to the model was simulated using a specialised LGN module that converts stimulus images to realistic firing rate spike trains. This allowed in silico and in vivo spiking activity to be compared in response to the same stimuli.

The work is of relevance because it shows how the spiking activity of another state-of-the-art biophysical cortical model is compared to in vivo. The Allen Institute requires synaptic conductances to be tuned sequentially on a layer-by-layer basis, however. Particularly, the initial ranges of recurrent synaptic weights between neuronal classes were found through a grid search that optimised for (1) spontaneous firing rates, (2) resemblance to in vivo peak firing rates in response to a single trial of a drifting grating (0.5 s long) and (3) non epileptic activity. The 0.5s drifting grating stimulus used for tuning is 250 times longer than the 2ms whisker deflection studied in this chapter. In vivo and in silico responses were also compared for full field flashes, natural movies and looming stimuli.

In vivo and in silico spiking activity were compared for direction selectivity, orientation selectivity, firing rates, time to peak, population sparsity, noise correlations, signal correlations and the correlation between noise and signal correlations. Responses may be compared by layer and 4 neuron types (excitatory, Htr3a, Sst, Pvalb). Similarity between in vivo and in silico varied between measures. Most notably, for full field flashes, the in silico time to first peak statistic shows big differences with in vivo ($\sim$20-40ms for different neuron groups; Figure S8 in [23]), suggesting that the millisecond temporal dynamics of the model are different from in vivo.

The tuning of synaptic conductances differs from the approach of the Blue Brain Project, in which the tuning of model parameters is generally avoided. Indeed, this

chapter explores the effect of 4 parameters on the similarity of spiking activity to in vivo similarity. 2 of these parameters are related to the input pattern and 2 are related to the global excitability of the model. The Allen model uses similar synaptic connectivity principles to the BBP model. In a second version of the model, the paper incorporates recurrent synaptic connectivity probabilities and synaptic weights that are dependent on the similarity of responses of neurons to the angle of drifting gratings. This improved the similarity of orientation selectivity measures with in vivo. This approach could be interesting to studying direction-selectivity in the BBP model in future.

5.3.2 Responses of morphologically detailed neurons in the barrel cortex [69, 68]

Pertinent to the work of this chapter is a model of the barrel cortex [69, 68], which allows comparison of in vivo and in silico whisker evoked responses. Rather than modelling every neuron in a column of the barrel cortex in high morphological detail, however, these studies place morphologically detailed neuronal models within a connectome, in order to infer synaptic innervation patterns for these single neuron models. Activating these neurons based on constrained innervation patterns produces activations similar to neurons in the barrel cortex. As these neuronal models are activated in isolation, these neurons would not exhibit shared changes in excitability, which are of interest to the thesis.

Recent work utilising similar connectivity principles to those used in the original BBP model, has enabled the development of a dense connectome of the barrel cortex [278], which may be used in future for morphologically detailed simulations of the barrel cortex. For future reference, the techniques employed in [68] offer predictions of membrane potential depolarisation levels for the modelled neurons.

5.3.3 Multi-scale EM reconstruction and in vivo functional comparison [277]

A recent study from L2/3 mouse V1 presents a detailed reconstruction of a 250x140x90 μm^3 EM volume. The model is accompanied by 2-photon calcium imaging data for a subset of the reconstructed pyramidal cells. These images were co-registered to the reconstruction, enabling the relationship between structure and function to be studied using directly recorded in vivo responses for the reconstructed neurons. The

reconstruction contains 8 million segmented objects and 3 million synapse predictions. The segmentation also has several types of glia, including astrocytes, microglia, oligodendrocytes and their precursor cells. The authors also performed automated reconstruction of the roughly 2.4 million mitochondria within the volume.

5.3.4 Null model of the mouse neocortex micro-connectome [222]

The BBP has recently published a full statistical micro-connectome of the mouse neocortex [222], opening a path to the first full-scale biophysical reconstruction of the neocortex. Neurons in the model studied in this chapter [175] require artificial somatic current injection to substitute for excitation that would be provided by afferent connections from neurons outside of the reconstructed piece of tissue. Simulation of the full neocortex would reduce the level of artificial current injection required to simulate neural activity.

5.4 Accounting for in vivo sampling bias

Before making visual and quantitative comparisons between in vivo and in silico responses, it is important to the consider the differences in sampling techniques for the in vivo and in silico datasets, and for in vivo recording techniques in general. Firstly, the in silico model allows access to the spike trains of all neurons in the volume. This means that statistics calculated for the in silico model reflect the true spiking activity of the model. Unlike the in silico case, the in vivo dataset is likely to contain some sampling biases; for example caused by (i) biased detection of more frequently occurring and higher amplitude action potential waveforms and (ii) the local spike overlapping problem (i.e. spike detection difficulties when two or more neurons fire in close temporal and spatial proximity).

Moreover, it is well known that a large proportion of cortical neurons are silent or spike very rarely [255]. Biased detection of more frequently occurring and higher amplitude action potential waveforms produces a bias against such non-spiking neurons, and would introduce an overall positive bias. That is, the recorded in vivo population response (for a neuron group for example), would be higher than the true average in vivo population response. Under such an in vivo bias, an silico model acting under the same regime as the true in vivo case, would appear to have a lower activity than the recorded in vivo case. Similarly, an in silico model with activity appearing to match the in vivo case, would have a higher activity than the true in vivo regime.

The local spike overlapping problem (i.e. spike detection difficulties when two or more neurons fire in close temporal and spatial proximity) is likely to introduce a negative bias when firing rates are high, for example, during stimulus-evoked activity. That is, the spike overlap problem may reduce the recorded in vivo firing rate relative to its true value.

The paper titled 'Population-wide distributions of neural activity during perceptual decision-making' [294] gives an excellent review on population-wide distributions of firing rates and the issues of in vivo recording bias. The details of this review provide an important basis for the author's future work, and so it was useful to summarise the key points of the review regarding firing rate distributions. The author plans to use this information as a basis for comparing in vivo and in silico firing rate distributions in future. The following paragraphs summarise the key points of [294] and findings relating to these subjects.

5.4.1 The sampling problem: recording methods and their biases

[294] emphasises that most recording techniques are likely be biased. The paper first discusses the bias of extracellular recording for detecting the most active neurons (citing [207]). Analysis of the expected numbers of neurons within the recording radius of a microelectrode suggests that most neurons are probably silent. This is known as the 'dark neuron' problem; [255]. This 'activity-dependent' bias leads to mean and standard error characterisations of population activity which are higher than the true population values, although the paper suggests that multi-electrode + tetrode spike detection techniques offer an improvement over single electrode recording.

Patch-clamping and calcium-imaging are given as methods, which circumvent the problem of activity bias. Patch-clamping finds neurons by physical contact alone, thus removing the bias towards highly active neurons. Patch-clamping indeed confirms that more neurons are found to be silent than for extracellular techniques (see paper Sec 3.1). Particularly, patch-clamping may bias sampling towards larger neurons or reduce their firing rates, whilst calcium-imaging may bias the sampling towards neurons that favour expression or uptake of fluorescent markers. The paper suggests that these biases are 'likely to be less important than the activity-dependent bias of traditional electrophysiology'. Automated patch clamping techniques are likely to improve the quantity of data available across layers and systems (e.g. [120]).

5.4.2 Population distributions of firing rates

The paper emphasises that it is necessary to characterise the probability distribution of firing rates over a population, although such characterisations have rarely been made. The paper shows that 'mean firing rate tells us little about the population, and that firing rate distributions are clearly not Gaussian'. Interestingly, and perhaps different from the normal conception, the paper states that 'whether characterising spontaneous activity, stimulus-evoked activity, or activity related to the execution of actions, the distributions are remarkably similar and robust'.

5.4.3 Long-tailed distributions

The paper focuses on the observation that distributions of firing rates are dominated by neurons with low firing rates, with a small set of neurons scattered over orders of magnitude higher rates, forming long tail distributions[1].

In the barrel cortex of awake mice using cell attached recordings revealed long-tailed distributions of firing rates when averaged over stimulus trials [203]. [294] tries to fit various probabilistic models to distributions of firing rates from various studies across the cortex, and collates this analysis with previous similar work. The ubiquity of long-tailed distributions is demonstrated, and it is emphasised that the most common model points to non-negligible peaks at low firing rates. The analyses are made for activity pooled during anaesthesia, sleep, quiet wakefulness, "baseline" epochs of recordings before a stimulus presentation or a movement, and "aggregate" recordings analysed across presentations of stimuli or productions of actions. The various datasets use a variety of recording techniques, although consideration of recording bias sometimes seems lost in the paper's discussion of firing rate distributions.

Fig. 5.3a (from Fig. 3a [294]) shows that 'neural populations' distributions of spontaneous and aggregate firing rates are consistently right-skewed and long-tailed in examples from across the whole cortex' with 'consistency across species, recording technique and vigilance state – under anaesthetic, awake and resting, or awake and behaving'. The authors state that 'in every case, the distributions are dominated by neurons firing less than 1 spike/s, with few neurons firing at rates an order of

[1]Useful footnote from [294]: 'In statistical literature, "long tails" often refer to precise mathematical definitions which essentially characterize a slower asymptotic drop-off than the exponential distribution. Since such notions are not applicable to finite, empirical distributions, we use the word "long-tailed" in a more informal meaning, to describe distributions which clearly drop off more slowly than a Gaussian.'

magnitude higher'. The authors fitted exponential, log-normal, and gamma probability distributions, and found that all three distributions occur in the cortical activity data-sets (Fig. 5.3b; from Fig. 3b,c [294]). The authors state that the differences between distributions 'could represent differences between vigilance states (awake vs anaesthetised), of recording type, of the sample size and construction (pooled vs single recording), or of the spatial sampling (see results on layer-specificity below)' or 'different characteristics of the cortical micro-circuits in different areas.'

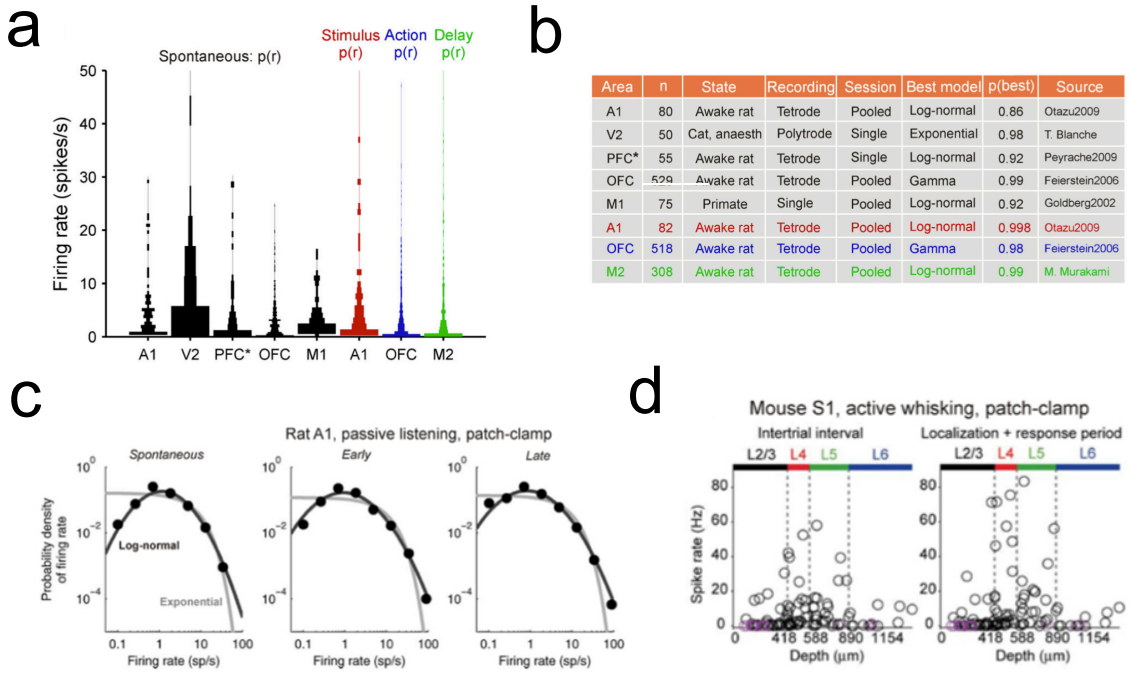
From the analyses, the authors concluded that most cortical neurons fire at sub-1Hz, which fits with the metabolic demands (see [294] for reference) and the likely recording radius of electrode technology [255]. The authors were surprised by the ubiquity of the rejection of an exponential model, suggesting that 'low firing rates dominate, but not arbitrarily low'. The authors do not discuss whether this may have been caused by low detection of non-spiking neurons, however.

5.4.4 Layer-specific differences in spontaneous activity

[294] cites [16] as giving a comprehensive review of differences in average firing rate between layers using different recording techniques. [294] discusses the following examples of between layer differences in average firing rates from [16]:

- [147] was given as a pioneering use of calcium-imaging in vivo, which reported that 'layer 2/3 in both A1 and S1 of the rat had very sparse firing rates under anaesthesia, with a maximum of 0.1 Hz over ten minutes of recording across 212 neurons (their Figure 6G)'. Moreover, 'such low rates were confirmed by simultaneous cell-attached or whole-cell recordings from a sample of those neurons'.

- Other studies confirmed that 'layer 2/3 neurons have substantially lower rates than layer 4 and 5 neurons for both spontaneous and stimulus-evoked firing [54] (barrel) and [245] (A1)'. For future reference [54] provides spontaneous firing rates across layers in barrel cortex using juxtasomatic recordings, which it claims provides an unbiased sampling technique. Particularly, [54] describes the following methodology: 'the electrode resistance of the recording electrode was monitored while advancing in $2\mu m$ steps (conventional method for whole-cell recordings). Using this method, unit isolation was independent of spiking frequency of cortical cells to ensure completely unbiased sampling (48 cells). Upon increase of electrode resistance, APs would generally already appear as negative-orientated waveforms. The electrode was then advanced stepwise at



Area	n	State	Recording	Session	Best model	p(best)	Source
A1	80	Awake rat	Tetrode	Pooled	Log-normal	0.86	Otazu2009
V2	50	Cat, anaesth	Polytrode	Single	Exponential	0.98	T. Blanche
PFC*	55	Awake rat	Tetrode	Single	Log-normal	0.92	Peyrache2009
OFC	529	Awake rat	Tetrode	Pooled	Gamma	0.99	Feierstein2006
M1	75	Primate	Single	Pooled	Log-normal	0.92	Goldberg2002
A1	82	Awake rat	Tetrode	Pooled	Log-normal	0.998	Otazu2009
OFC	518	Awake rat	Tetrode	Pooled	Gamma	0.98	Feierstein2006
M2	308	Awake rat	Tetrode	Pooled	Log-normal	0.99	M. Murakami

Figure 5.3: (a) Fig. 3a from [294]. Caption: 'Violin plot of firing rate distributions across cortical regions, in both spontaneous and task-aligned conditions. Width of bar is proportional to the fraction of neurons in that bin. Optimal histogram widths separately derived for each data-set'. Fig. 3b from [294]. Caption: 'Summary of data-sources in panel (a) and best-fit model to each data-set. Exponential, gamma, and log-normal distributions were fitted using maximum likelihood; best-fitting model was chosen using the Bayesian Information Criterion; p(best): the posterior probability of that being the best model within those tested. The Session column indicates whether the data were from a single recording session or pooled across sessions and/or animals.' (c) Fig. 4a from [294]. 'Distribution of spike counts across neurons in the auditory cortex (A1) of awake rats responding to a pure tone pip. Activity distributions are shown before ('spontaneous'), during ('early', 'late') and after ('offset') the tone pip. Best model fits are provided, using respectively an exponential (gray) and a log-normal (black) distribution. The log-normal provides a markedly better fit. Loose patch clamp recordings, from [125]. (d) Fig. 4b from [294]. 'Distribution of spike counts in the layers of barrel cortex (S1) of awake mice involved in a tactile detection task. The mean level of activity is plotted for each neuron against its cortical depth, both during ('localization + response') and between ('intertrials') periods of whisker contact with the stimulus. Loose patch clamp recordings from [203].'

603 0.1Hz until positive AP waveforms were recorded with amplitudes of 1–3 mV
604 (also additional 3–5 steps).'

605 Beyond average firing rates, [294] highlight the following examples of layer-specific
606 delineation of firing rate distributions:

- [203] reported that 'barrel cortex had very long-tailed distributions in each layer (2/3, 4, 5, and 6); moreover, the distributions were significantly different between layers, with layers 2/3 and 6 showing lower rates than layers 4 and 5 (Fig. 5.3d).'

- [236] 'constructed a comprehensive map of the layer-by-layer differences in spontaneous firing rates of monkey V1 ([294] Figure 4c), showing that, while all layers were dominated by low firing rates of less than 1 Hz, layers 4-6 had markedly more neurons with higher firing rates than layers 2/3.'

- [125] found 'that the firing rate distributions in rat A1 were consistently long-tailed in each layer (2/3, 4, 5 and 6) though not explicitly tested, that the distributions were not the same between layers.'

5.4.5 Similarity of spontaneous and evoked firing rate distributions

The authors also analysed the firing rate distributions for stimulus or event evoked activity in different cortical areas. The authors again found that 'silence reigns'. Log-normal fitted the activity of A1 and M2 best whilst a gamma distribution best fit the activity of OFC. Moreover, [203] reported long-tailed distributions of firing rates (patch clamp) in mouse barrel cortex (S1) aligned to stimulus contact by the whiskers. Interestingly, the authors found that log-normal distributions, rather than exponential, best fit cortical firing rates recorded by two different techniques (single cell patch clamp [125] and multiple cell (tetrode) [208] in a single cortical area (rat A1) during different behavioural states (passive receipt of stimuli vs active engagement in a decision-making task).

[294] state that 'it has generally been thought that population activity in visual cortices changes dramatically when a stimulus is presented'. [294] challenge this by reviewing results which reveal 'that population-wide statistics of task-driven and spontaneous activity may be much closer than previously expected'. Particularly [294] Fig. 4a and b show 'that distributions of mean firing rates in rat A1 [125] and mouse S1 [203] barely changed between baseline and evoked firing'. Similar results were confirmed for action related [76] and stimulus-evoked activity in rat OFC and a different rat A1 study [208, 294]), including a 'small and smooth increase in the probability of neurons appearing in the tail of the distributions of evoked firing rates' ([294] Fig. 4d).

5.4.6 The single neuron's perspective

[294] suggest that 'these small changes in distribution obscure substantial changes at the single-neuron level between spontaneous (baseline) and evoked activity.' Details supporting this statement are given in the paper but importantly they state that 'since firing rate changes of individual neurons can be both positive and negative, the overall effect of these changes on the firing rate distributions is rather small'. [294] considers how these results are compatible with the fluctuations in population-wide activity traditionally observed with imaging methods such as fMRI and EEG.

[294] suggest that there is no fundamental contradiction because overall firing rates are so low that few cells with strong responses may actually be sufficient to significantly increase the overall number of spikes. That is, as the distributions are long-tailed, this overall increase in average firing rate is compatible with only small changes in the distribution itself. The paper gives the examples:

- [125] 'report both that 16% of neurons contributed 50% of spikes and an overall increase of mean firing rate of almost 50% between pre-stimulus and early stimulus responses.'

- 'In mouse barrel cortex (S1) responding to a simple tactile discrimination task, [203] report that around 10% of the cells in a single column account for 50% of all spikes fired' and that 'their data also suggests that around 15% of the cells increased their firing rate by more than 10 spikes/s after the onset of active whisking (their Figure S3, panel A1).'

- Similar results were reported for rat OFC [76].

Section 5.5.2.4 discusses the approaches used in this chapter to account for in vivo sampling bias.

5.5 Methods: comparison to in vivo

The original model paper compared in silico and in vivo stimulus-evoked activity for $[Ca^{2+}]_o$: $\sim 1.25mM$, as this marked a transition point between asynchronous and synchronous spiking regimes. In initial investigations under such conditions, an early version of the methodology described in Chapter 3 did not detect the time-warp phenomena. Furthermore, observation of first spike time plots for pairs of neurons

confirmed a qualitative difference between in vivo and in silico spiking activity. Moreover, in silico stimulus-evoked activity appeared stronger than in vivo activity, with some neurons firing with higher precision, reliability and spike counts.

It was hypothesised that such strong neuronal responses could be the result of heightened excitability in the model, that could possibly be caused by above in-vivo $[Ca^{2+}]_o$-levels. Such heightened excitability could 'override' sources of variability such as trial-to-trial changes in shared excitability and explain why the time-warp phenomena was not detected.

The original paper reports that the true in vivo $[Ca^{2+}]_o$ range lies between 0.9-1.1 mM, citing the following references: [142, 179, 11, 24]. It was hypothesised that the model may exhibit the time-warp phenomena under conditions for which the population statistics of in silico spontaneous and stimulus-evoked activity closely resembled in vivo activity. As a first step, the author aimed to quantify the effect of $[Ca^{2+}]_o$, artificial depolarisation-level and input parameters on in silico neural activity; particularly quantifying the similarity of in vivo and in silico activity under different conditions. The methods for such comparison are described here and attempt to take into account the possibility of sampling biases introduced by extracellular in vivo recording.

5.5.1 Model and comparison neuron groups

This chapter uses the mesocircuit explored in the initial model paper (described above). The mesocircuit consists of a central microcircuit, surrounded by 6 additional microcircuits, which are included to minimise edge effects in the central microcircuit. The chapter analyses activity within the central microcircuit, which is on the order of size of a single barrel column. In vivo and in silico spiking activity was compared by cortical layer and E/I type. 3 layer groups were defined for comparison, giving a total of 6 E/I layer groups. The in vivo layer groups were defined as L2/3, L4 and 'L5', which contained neurons assigned to L5A and L5B/6. These 3 groups were compared directly to in silico layer groups L2/3, L4 and L5. Future work could define in silico L5A and L5B/6 layers, and separate in vivo neurons assigned to L5B and L6.

As described in Section 5.2.1, a single fiber was instantiated and centered in each of 300 minicolumns in the central mesocircuit with a horizontal spread. As such, fibres are placed in a circle of radius $200\mu m$, in which thalamocortical projection axons were spaced uniformly (Fig. 5.4). The number of thalamic projection fibres was consistent with estimations that a barreloid in the rat may contain 250-300 neurons, with 95% of retrogradely labelled neurons in the barreloid that is isomorphic to the injected barrel

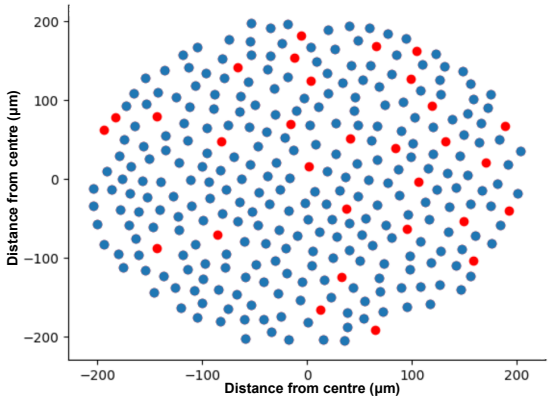


Figure 5.4: Vertical locations of 300 thalamic fibres within the microcircuit (with respect to the centre of the microcicuit). For an example stimulus in which 10% of fibres are activated, red circles represent fibres which are part of the stimulus, whilst blue circles represent inactive fibres.

and up to 5% in a single adjacent barreloid [156]. The number of these fibres that are used to evoke activity varies between experimental conditions (described below).

5.5.2 Depolarisation, $[Ca^{2+}]_o$ and input parameter grid search

To quantify the effect of artificial somatic depolarisation, extracellular calcium concentration ($[Ca^{2+}]_o$) and in silico thalamic input parameters on the resemblance of in silico and in vivo stimulus-evoked activity, simulations were run for all combinations of values of the following 4 parameters:

- *Somatic depolarisation.* Neurons in the model require artificial somatic current injection to compensate for excitation that would be provided by afferent connections from neurons outside of the reconstructed piece of tissue. Somatic current injection was expressed in terms of percentage of the minimum step current injection required for each neuron to spike at least once. Previous work has used a depolarisation of 100%. Artificial current injection is applied with a very small amount of variance so that the variance of the depolarisation is 0.001. This method of artificial depolarisation has been used in all previous investigations, but novel methods of depolarisation are currently being explored (see Discussion). It is hoped that the methods presented here will be useful for the testing of such novel depolarisation techniques. In this chapter, depolarisation is varied over the following values: 80%, 85%, 90%, 95%, 98%, 99%, 100%.

- *Extracellular calcium concentration.* Extracellular calcium concentration ($[Ca^{2+}]_o$) was varied over the values: 0.9mM, 0.95mM, 1.0mM, 1.05mM, 1.1mM, 1.15mM, 1.2mM, 1.25mM. This set of values covered the range of values reported in vivo (0.9-1.1mM), as well as the range previously used to create 'in vivo-like' conditions in in silico simulations for the model (1.15mM-1.25mM). For future

reference, [141] gives a good summary on in vivo extracellular calcium dynamics.

- *Thalamic input pattern density.* This defines the proportion of the 300 thalamic projections that are activated in the input pattern, as shown in Fig. 5.4. Thalamic input pattern density was varied over the values: 0.1, 0.3, 0.5.

- *Thalamic input pattern temporal width.* Each thalamic projection is assigned a unique spike time with respect to stimulus onset. This spike time is fixed over repetitions of the stimulus (i.e. there is no input jitter over multiple trials of the stimulus). This allows study of variability, which is intrinsic to the microcircuit. Such spike times were drawn from uniform distributions of width 0.0ms, 2.5ms and 5.0ms. Whilst there are several characterisations of trial-to-trial temporal jitter of single neurons in the VPM, to the authors knowledge there are no characterisations of millisecond scale trial-to-trial noise correlations between VPM neurons. The addition of independent trial-to-trial jitter to individual neuron spike times could potentially decorrelate (over trials) the input compared to in vivo conditions.

This gave a total of 504 parameter combinations. For each parameter combination, a simulation was run for 4.25 seconds of model time. During this simulation, 5 identical thalamic stimulations (defined by the input parameters of the specific parameter combination) were made at 500ms intervals with the first stimulation starting after 2 seconds of model time. Responses to the first stimulation were ignored, giving 4 trials per unique parameter combination, which was sufficient for calculating the population statistics in this part of the chapter. Each of the 4 stimulations is referred to as a single trial.

5.5.2.1 Cortical onset

As described, the model does not simulate whisker stimulations and the propagation of activity from whisker follicles to the thalamus. Instead sensory stimulation is modelled by assigining spike times to thalamic neurons. Thalamocortical synapses are then activated after a given delay based on their distance from the thalamocortical neuron. This means that in silico and in vivo cortical activity can not be compared with respect to stimulus onset.

So that cortical activity was fairly compared with the in vivo case, a unique cortical onset time was calculated for each parameter combination. Cortical activity onset was

calculated by pooling spike times over all neurons and trials for a single parameter combination, and applying the same method outlined for neurophysiology data in the previous chapter but with different parameters as described here. Particularly, for a single parameter combination, cortical onset was determined by calculating a histogram of single neuron spike times between -10ms before and +100ms after stimulus with bin size equal 0.2 ms. The histogram was smoothed with a Gaussian kernel (σ = 1 ms). The mean and standard deviation of the smoothed bin counts was calculated for the period before stimulus onset. The cortical activation onset spiking threshold was determined as this mean value + 10 standard deviations. 1.5ms was added to this value to account for the width of the Gaussian smoothing kernel. Observation of the spike rasters showed that this approach gave a consistent alignment with the in vivo case. As activity is brought closer to the in vivo case (particularly with more similar spontaneous firing rates), it is likely that the parameters of cortical onset detection can be brought closer to the in vivo case.

5.5.2.2 In vivo spiking activity

The dataset described and analysed in Chapter 3 is used as the source of in vivo spiking data. In silico spiking activity is compared to in vivo spiking activity from the principal column recorded in response to deflections of the principal whisker at frequencies between 0.066-1.0Hz. In vivo neural activity is pooled from all 18 experimental neurophysiology sessions (described in Chapter 3). As described and demonstrated in Chapter 3, low stimulation frequencies limit the effect of adaptation, offering the best example of stationary in vivo activity. The in vivo data recorded for higher stimulation frequencies will allow in vivo comparison for future models equipped with adaptative mechanisms (e.g. [10]).

5.5.2.3 Visual comparison

For each parameter combination, single trial raster plots were plotted for each E/I layer group zeroed by the determined cortical activity onset point. A max normalised PSTH was calculated for each E/I layer group. For the same neuron groups, a similar raster plot was created for in vivo spiking activity in response to deflections of the principal whisker occurring at frequencies of $0.066 - 1Hz$. Particularly, a number of randomly sampled spike trains, equal to the number of in silico neurons in the E/I layer group (of the central column), were plotted. This allowed direct visual comparison between in vivo and in silico activity. Such visual comparison between the

in vivo and in silico raster plots showed clear differences in activity for all parameter combinations. These differences are described and quantified below.

5.5.2.4 Metrics accounting for in vivo sampling bias

The chapter explores several measures, which could reduce the effect of such in vivo sampling biases. For example, comparing the shapes of PSTHs independently of amplitude removes the effect of a positive bias towards larger and frequently occurring neurons, or of negative biases (if they were present). Calculating firing rate ratios provides another useful means for removing in vivo biases. For example, calculating the ratio between the mean firing rates of two neuronal groups would remove a bias that was consistent in both neuron groups. The values of such ratios could then be compared with in silico. Another example is the ratio between spontaneous and stimulus-evoked firing rates. It is important to note, however, that in vivo firing rates could be more negatively biased by the spike overlapping problem for stimulus-evoked activity than for spontaneous activity.

These measures and comparisons to in vivo were intended as a first step towards testing whether an in silico cortical model is close to a true in vivo regime, with the primary aim to provide a network which could be tested for the time-warp phenomena. The following results section also aims to select a model with high similarity to in vivo activity for the test of the time-warp phenomena. The set of comparisons were also chosen to act as a first 'benchmark', which could guide data-driven refinement of the model. Analysis of neuron group mean firing rates for the spontaneous regime is sub-optimal compared to characterisation of firing rate distributions, but the metrics offered a quick and straightforward way of determining appropriate network parameters for the time-warp test. With the large amount of data generated by hundreds of large scale simulations under a variety of network states, the metrics provided an interpretable set of analyses, which can be built upon and improved in future work. The discussion gives a brief summary of the steps which have been taken following the completion of this chapter in April 2021.

5.5.2.5 Comparison time windows

To test for the time-warp phenomena within the model, the thesis primarily considers activity shortly following the onset of cortical activity, as well as the relationship of this activity with pre-stimulus spontaneous activity. It is therefore important that model activity closely resembles in vivo activity shortly after cortical activity onset. A time window of 25ms following cortical onset is chosen for the main in vivo / in

silico comparisons, as this represents a large proportion of the stimulus-evoked rise and decay of spiking activity, as will be shown below.

It is also important that the model exhibits a similar decay of activity over a longer timescale of 100ms, as this provides an important validation that the model is in a similar regime to the in vivo case. Finally, it is important that the model exhibits spontaneous activity that is similar to in vivo. For this, in vivo and in silico spiking activity will be compared over a 50ms time window preceding the stimulus.

5.5.2.6 Quantitative comparison framework

A general comparison framework was developed to compare in vivo and in silico spiking activity on multiple timescales, both by neuron group and for the entire population. The framework currently allows comparison of in silico activity with in vivo responses to principal whisker stimulations made at ranges of defined frequencies (i.e. 0.066-1.0Hz) on certain trials (i.e. trials 0-10). The framework is designed to allow comparison to in vivo activity recorded in barrels neighbouring the principal column in future. The framework enables visual and quantitative comparison of similarity between in vivo and in silico spiking statistics. The framework also enables the effect of model parameters on spiking statistics and similarity to in vivo to be analysed. The general comparison framework allows in vivo / in silico comparison for a range of statistics by neuron group and for the entire population, although not all comparisons are presented here. The comparison framework is also designed to enable the addition of further comparisons in future.

5.6 Results: comparison to in vivo

5.6.1 Visual comparison of stimulus-evoked activity

Before quantifying the similarity of in vivo and in silico activity, it is useful to visually compare the neural responses. Fig. 5.5a shows a spike raster of in vivo spike trains for the 6 comparison neuron groups (L2/3 E, L2/3 I, L4 E, L4 I, L5 E, L5 I) over a 25ms window following cortical activity onset. Particularly, a number of spike trains equal to the number of neurons in the corresponding in silico group (in the central column) are plotted. Fig. 5.5b shows in silico spiking activity for a single in silico parameter combination for the same 6 neuron groups for a single thalamic stimulation. The rasters allow direct visual comparison of in vivo and in silico responses.

Fig. 5.6a,b show spike rasters over a 100ms window following cortical activity onset for the in vivo case and an in silico parameter combination respectively. Note

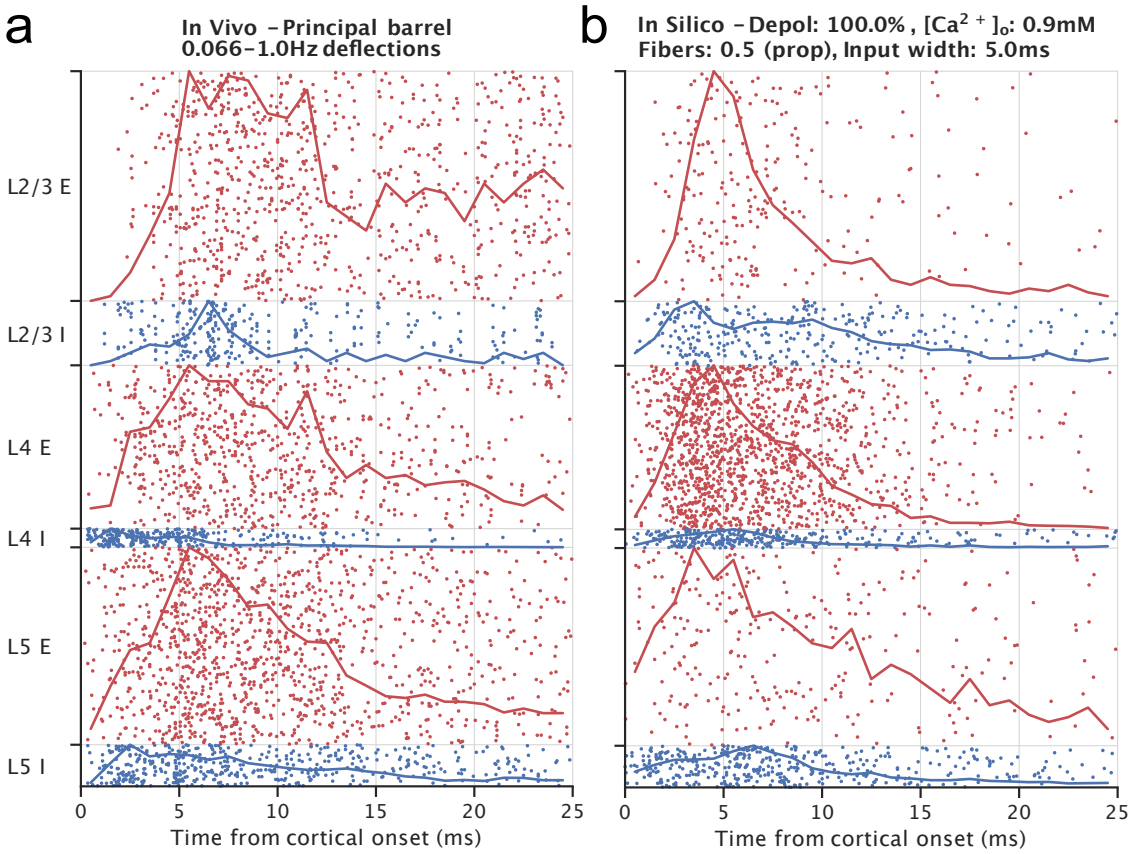


Figure 5.5: In vivo and in silico spike rasters and max normalised PSTH profiles. Spiking activity is shown for the 6 comparison neuron groups for a 25ms period following cortical activity onset. Lines show values of max normalised PSTH profiles for each neuron group (1ms bins). (a) In vivo spike trains of neurons in the principal column in response to principal whisker deflections at 0.066-1.0Hz. For each neuron group a number of spike trains equal to the number of neurons in the corresponding in silico neuron group are plotted. Max normalised PSTH profiles are calculated for all available spike trains for the condition. (b) In silico spike trains for the in silico parameter combination: depolarisation = 100%, $[Ca^{2+}]_o = 0.9mM$, proportion of input fibers in the stimulus = 0.5 and input pattern width = 5.0ms. The spike trains of all neurons on a single trial are shown for the 6 neuron groups. Spike trains from all 4 trials are used to calculate the max normalised PSTH profiles.

that the in silico combination is different to the combination shown in Fig. 5.5b. The chapter is primarily interested in the millisecond scale dynamics over a short time window (i.e. 25ms) following cortical activity onset. Comparison of activity over a longer time period (i.e. 100ms) provides an important indication as to whether the in silico model is in a similar regime to the in vivo case. Vids. 10 & 11 show the in silico spike rasters for all in silico parameter combinations over 25ms and 100ms respectively. The following subsections make a series of quantitative comparisons,

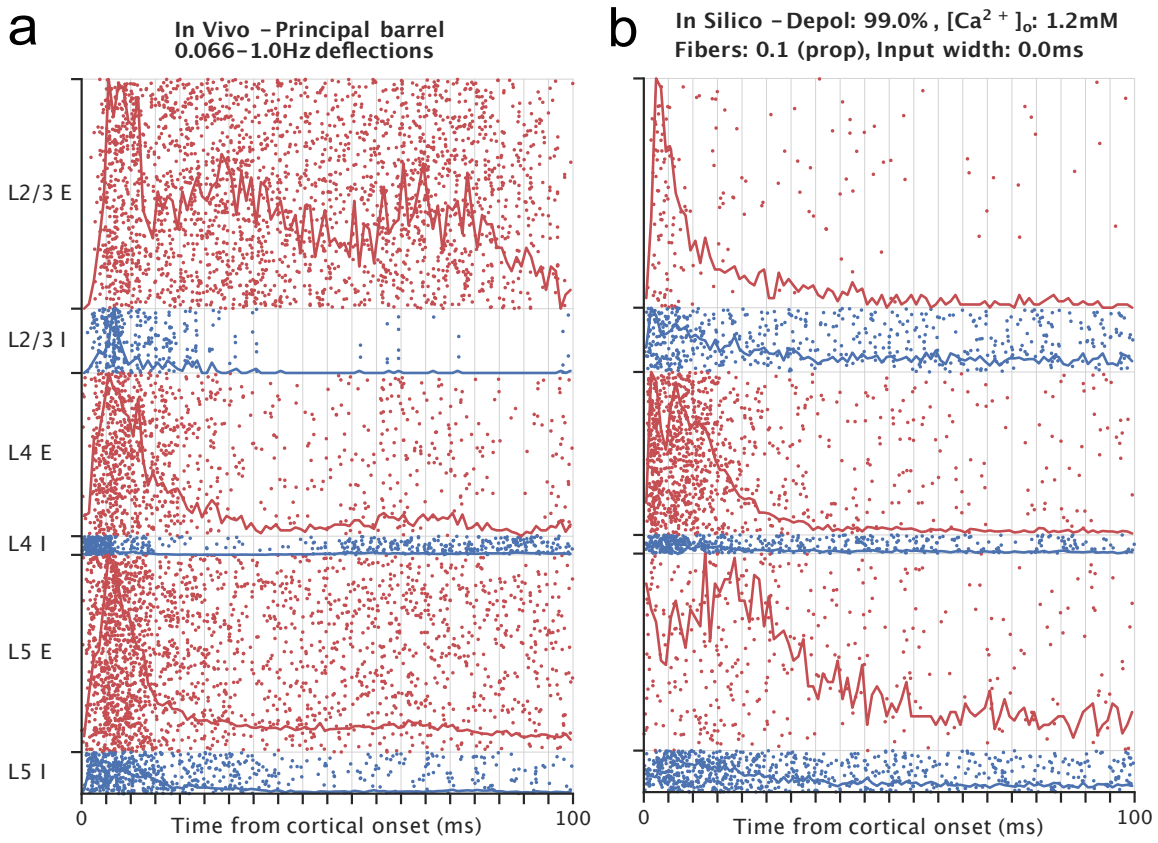


Figure 5.6: In vivo and in silico spike rasters and max normalised PSTH profiles. Spiking activity is shown for the 6 comparison neuron groups for a 100ms period following cortical activity onset. Lines show values of max normalised PSTH profiles for each neuron group (1ms bins). (a) In vivo spike trains of neurons in the principal column in response to principal whisker deflections at 0.066-1.0Hz. For each neuron group a number of spike trains equal to the number of neurons in the corresponding in silico neuron group are plotted. Max normalised PSTH profiles are calculated for all available spike trains for the condition. (b) In silico spike trains for the in silico parameter combination: depolarisation = 100%, $[Ca^{2+}]_o = 0.9mM$, proportion of input fibers in the stimulus = 0.5 and input pattern width = 5.0ms. The spike trains of all neurons on a single trial are shown for the 6 neuron groups. Spike trains from all 4 trials are used to calculate the max normalised PSTH profiles.

some of which support qualitative observations.

5.6.2 Ratios of stimulus-evoked firing rates between neuronal groups

Observation of the in vivo spike raster in Fig. 5.5a suggests that the level of activity is reasonably well balanced between in vivo neuron groups in different layers and between neuron groups within layers. Particularly the density of spikes is seemingly

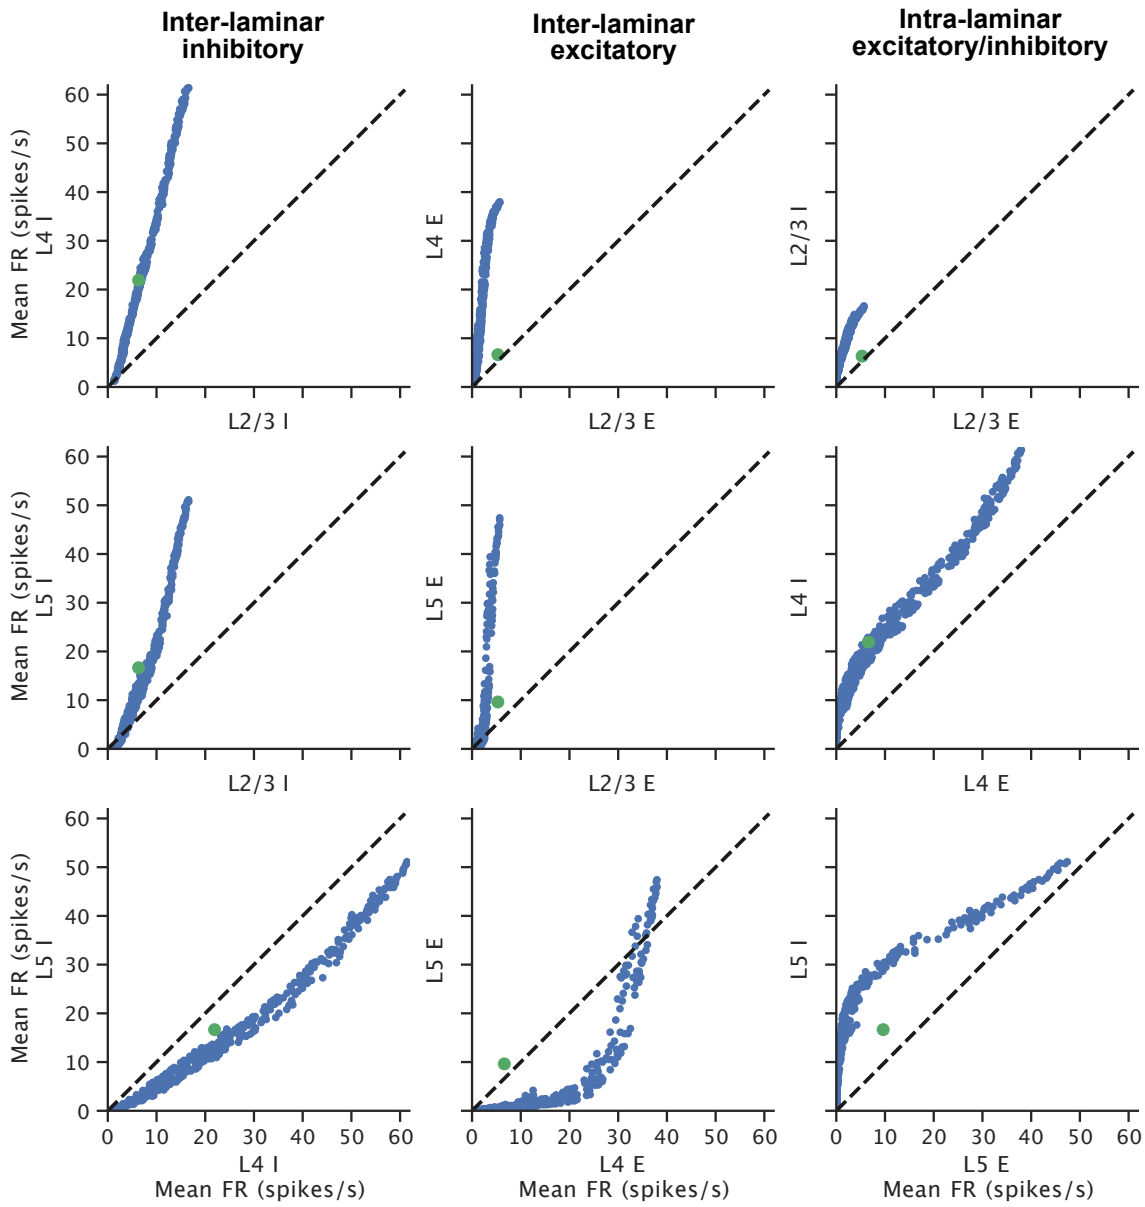


Figure 5.7: Comparison of mean firing rates between pairs of neuron groups. Each green dot represents the values for the in vivo case, whilst each blue dot represents the values for a single in silico parameter combination. Columns show comparison of mean firing rates between pairs of layer-wise inhibitory neuron groups (left column), between pairs of layer-wise excitatory neuron groups (central column) and between inhibitory and excitatory neuron groups within single layers (right column).

similar between neuron groups in different layers and between neuron groups within layers. Observation of the in silico spike rasters in Fig. 5.5b and Vids. 10 & 11 show that the balance of activity between excitatory neuron groups is less balanced for most in silico parameter combination.

To quantify these differences, the mean stimulus-evoked firing rate of each neuron group was calculated (over 25ms following cortical activity onset) for the in vivo and in silico cases. The ratio of stimulus-evoked firing rates between neuron groups offers a measure, which would remove sampling biases that are consistent across neuron groups. The left column of Fig. 5.7 shows that the ratios of firing rates between inhibitory neuron groups of the in silico combinations closely match the in vivo case for all parameter combinations. This is shown by the fact that the values for the in silico parameter combinations trace a linear lines that pass through the origin and the in vivo value. This offers an interesting validation of the model: that the relative stimulus-evoked firing rates between layer-wise inhibitory neuron groups match the in vivo case.

The ratios of stimulus-evoked firing rates between the excitatory groups of the in silico parameter combinations do not match the in vivo case, however. Whilst the stimulus-evoked firing rates of the layer-wise excitatory neuron groups are similar to each other in vivo (shown by the green dots close to $y = x$ in plots of the central column in Fig. 5.7), the in silico combinations trace a line that is further from $y = x$. In fact, for the majority of in silico combinations, L2/3 E and L5 E firing rates are lower than the firing rates of L4 E. This confirms the observation that was made for the raster plots in Figs. 5.5 & 5.6 and Vids. 10 & 11.

The right column of Fig. 5.7) plots the firing rates of the excitatory and inhibitory neuron groups for single layers. This plot essentially contains the same data that is available in the first two columns. The plot confirms that the L2/3 E and L5 E groups are consistently under excited with respect to the corresponding L2/3 I and L5 I groups and the in vivo case. In combination with the results of the first two columns, this suggests that the in silico L2/3 E and L5 E excitatory groups are under excited with respect to the other neuron groups compared to in vivo.

The discrepancies show that tuning of the four parameters cannot produce ratios of interlaminar excitatory group firing rates that match in vivo, and this must be considered when selecting an in silico parameter combination for the in silico test for the time-warp phenomena. Since presenting this comparison to the colleagues at the BBP, efforts have begun to address the imbalance identified here through the integration of additional data from the literature. Particularly, improved thalamic innervation is now being introduced which includes innervation from the POm thalamic area. POm fibres are known to project to L1 and L5 so are likely to improve the balance of stimulus-evoked activity between layers.

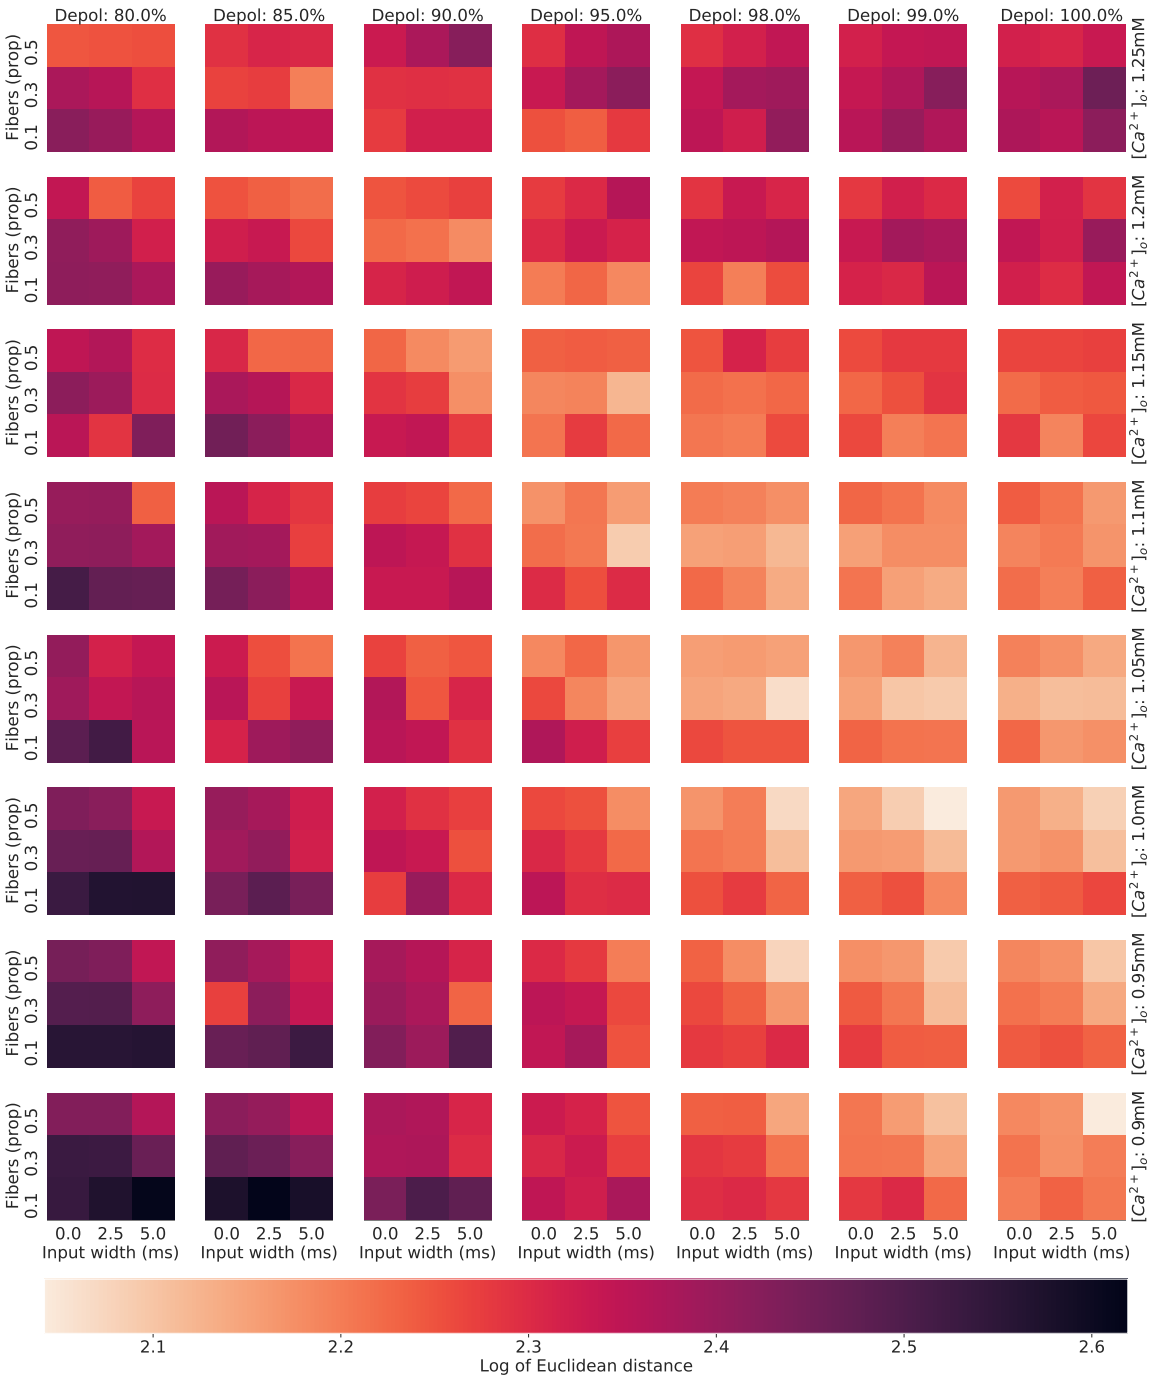


Figure 5.8: Heatmap showing the effect of the 4 parameters on the distance between in vivo and in silico PSTH profiles over 25ms following cortical activity onset. The 6 neuron group max normalised PSTH profile vectors were concatenated. The log of the Euclidean distance between the in vivo vector and the vector of each in silico combination gives a distance (the lower the better) for each in silico combination. Major columns: depolarisation, major rows: $[Ca^{2+}]_o$, minor columns: input temporal width, minor rows: proportion of fibers.

5.6.3 Similarity of PSTH Profiles

The effect of in silico parameters on the absolute similarity of in vivo and in silico firing rates will be analysed in a following subsection. This subsection, however, considers the similarity of in vivo and in silico neuron group PSTH profiles. As shown in the previous subsection, firing rate ratios in silico are dissimilar from the in vivo case for certain pairs of neuron groups. This suggests that there will be differences between in vivo and in silico PSTHs by neuron group for single in silico parameter combinations. Moreover, direct comparison of in vivo and in silico PSTHs does not take into account the possibility of in vivo sampling biases. For example, in vivo sampling biases could increase or decrease the magnitude of PSTHs relative to true in vivo PSTHs.

This section attempts to account for these issues of sampling bias and unbalanced inter-neuronal group firing rates by comparing max normalised in vivo and in silico PSTHs by neuron group. Particularly, the magnitude of each neuron group PSTH is normalised or 'removed' by dividing the values of the PSTH by the maximum (i.e. peak) value of the PSTH. This gives, for each neuron group, a max normalised PSTH with a peak value of 1. Such max normalised PSTHs are referred to as 'PSTH profiles', as they capture only the shape of the PSTH and not the magnitude of the response. The original PSTHs were calculated by binning spike counts in 1ms bins over all trials of the in vivo condition or parameter combination, and dividing by the number of trials and number of neurons in the neuron group (either in silico or in vivo).

The lines in Fig. 5.5a,b show the max normalised PSTH profiles for each of the in vivo and in silico neuron groups (for the example in silico parameter combination). The PSTH profile of each neuron group is therefore defined by 25 values representing the max normalised number of spikes for that neuron group in a single millisecond time bin. Comparison of such max normalised PSTH profiles offer a means of comparing the millisecond scale dynamics of in silico parameter combinations with in vivo dynamics.

To study the effect of the 4 in silico parameters on the similarity of PSTH profiles to in vivo, a PSTH profile distance was calculated for each in silico combination. In particular, for the in vivo case and each in silico parameter combination, the 6 PSTH profile vectors were concatenated to form vectors of length 150 (6 neuron groups x 25 time bins). The Euclidean distance between the in vivo vector and the vector of each in silico combination could then be taken. For presentation purposes, the log of this values was taken to give a final PSTH distance for each in silico combination.

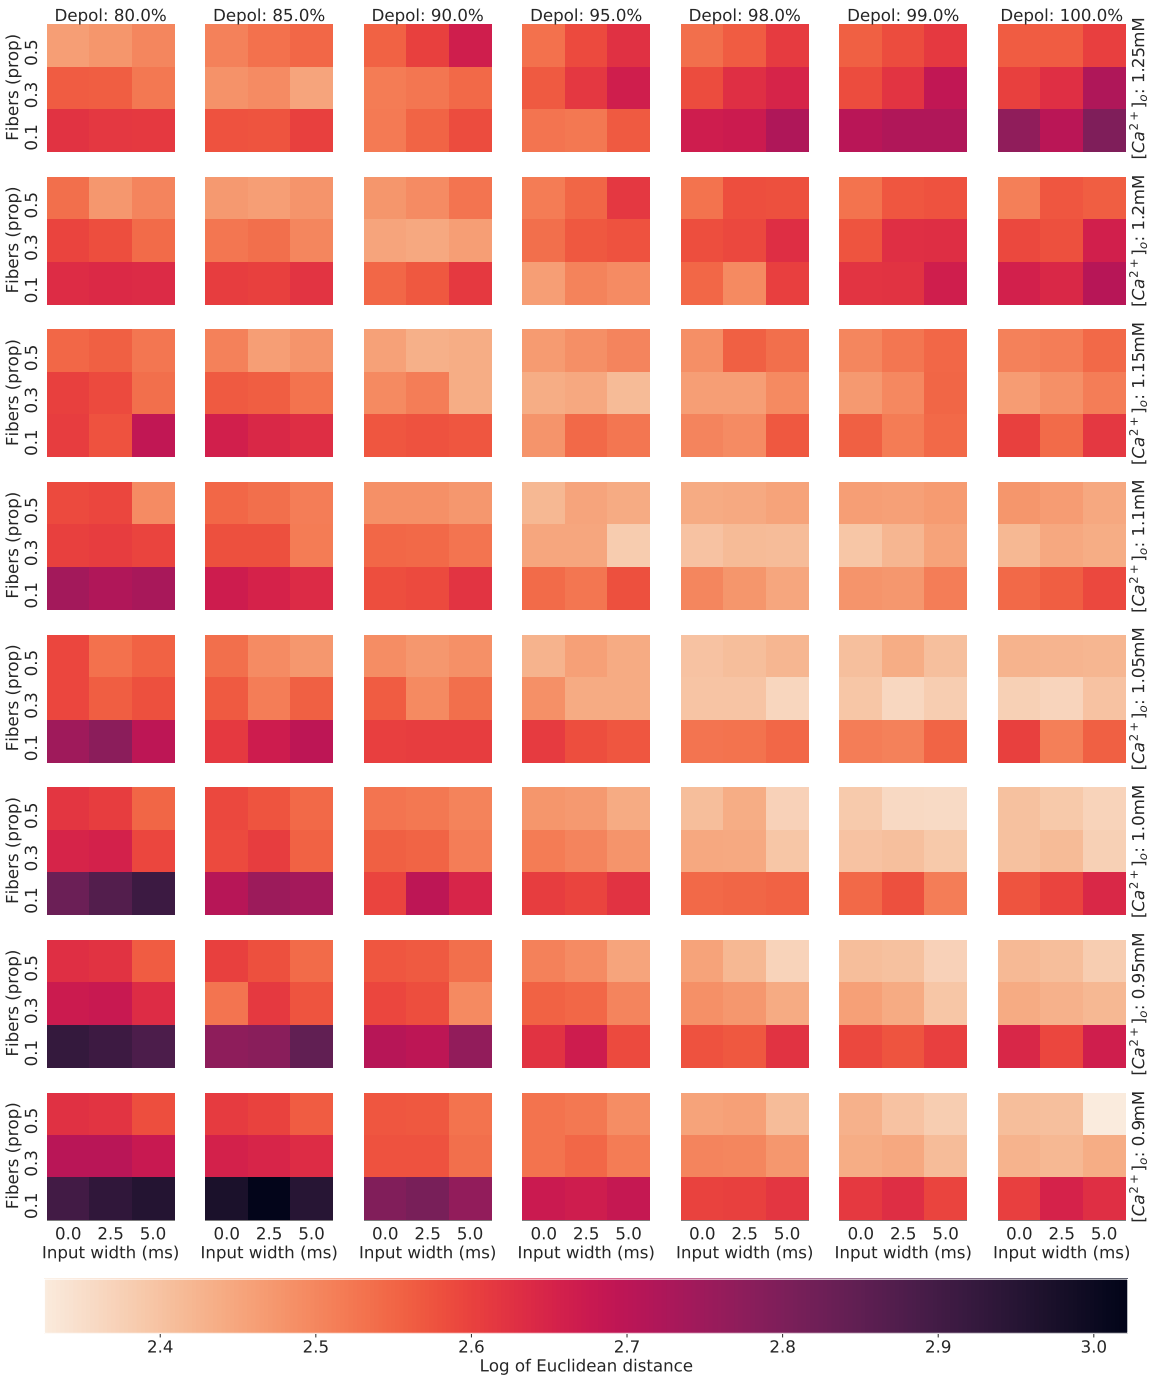


Figure 5.9: Heatmap showing the effect of the 4 parameters on the distance between in vivo and in silico PSTH profiles over 100ms following cortical activity onset. The 6 neuron group max normalised PSTH profile vectors were concatenated. The log of the Euclidean distance between the in vivo vector and the vector of each in silico combination gives a distance (the lower the better) for each in silico combination. Major columns: depolarisation, major rows: $[Ca^{2+}]_o$, minor columns: input temporal width, minor rows: proportion of fibers.

Fig 5.8 shows the effect of the four parameters on the PSTH profile distance. Increasing depolarisation clearly decreases distance (lighter heatmap colours). Models with high depolarisation values (98-100%) and calcium values within the in vivo range (0.9-1.1mM) show the highest similarity (shown by lower distance values). The results are similar over 100ms millisecond (Fig. 5.9), although further analysis of activity over longer timescales is left for future models. The following parameter combinations give the best in vivo similarity for the metric over 25ms (exhibited by the lowest values in Fig 5.8):

1. Depolarisation: 100%, $[Ca^{2+}]_{\circ}$: 0.9mM, fibers: 0.5, input temporal width: 5ms.
2. Depolarisation: 99%, $[Ca^{2+}]_{\circ}$: 1.0mM, fibers: 0.5, input temporal width: 5ms.

Herein, these in silico parameter combinations are referred to as in silico parameter combination 1 and 2. Both lie within the range of in vivo like $[Ca^{2+}]_{\circ}$ levels. The similarity of millisecond scale dynamics to in vivo makes these in silico parameter combinations candidates for the time-warp test. The similarity of in silico parameter combination 1 to in vivo for other metrics will be considered in the following sections.

Fig. 5.5b shows the in silico raster of in silico parameter combination 1. It can be seen that the max normalised PSTH profiles are similar to those of the in vivo case. The in vivo and in silico max normalised PSTH profiles (for in silico parameter combination 1) are plotted together for direct comparison in Fig. 5.10a.

5.6.3.1 Comparable spike counts and spike time latencies follow from PSTH profile similarity

Fig. 5.10b shows the spike count histograms of each neuron group for the in vivo case and in silico parameter combination 1. Interestingly, each in silico neuron group shows comparable activity to the in vivo case for the given metrics (especially for the inhibitory neuron groups). Notably, the L2/3 E and L5 E neuron groups show lower spike counts than in vivo, as expected from the analysis of stimulus-evoked firing rate ratios in Section 5.6.2.

For in silico parameter combination 1, spike time latencies closely match the in vivo regime 5.10c. These results demonstrate that selecting the value of the 4 in silico parameters based on max normalised PSTH profile similarity can give rise to similar spike counts and spike time latencies to in vivo (with the caveat that it is difficult to determine how well the calculated in vivo values reflect the true in vivo values). Finally, 5.11 shows (for in silico parameter combination 1) that the inhibitory firing rates show close similarity with the calculated in vivo firing rates, as expected

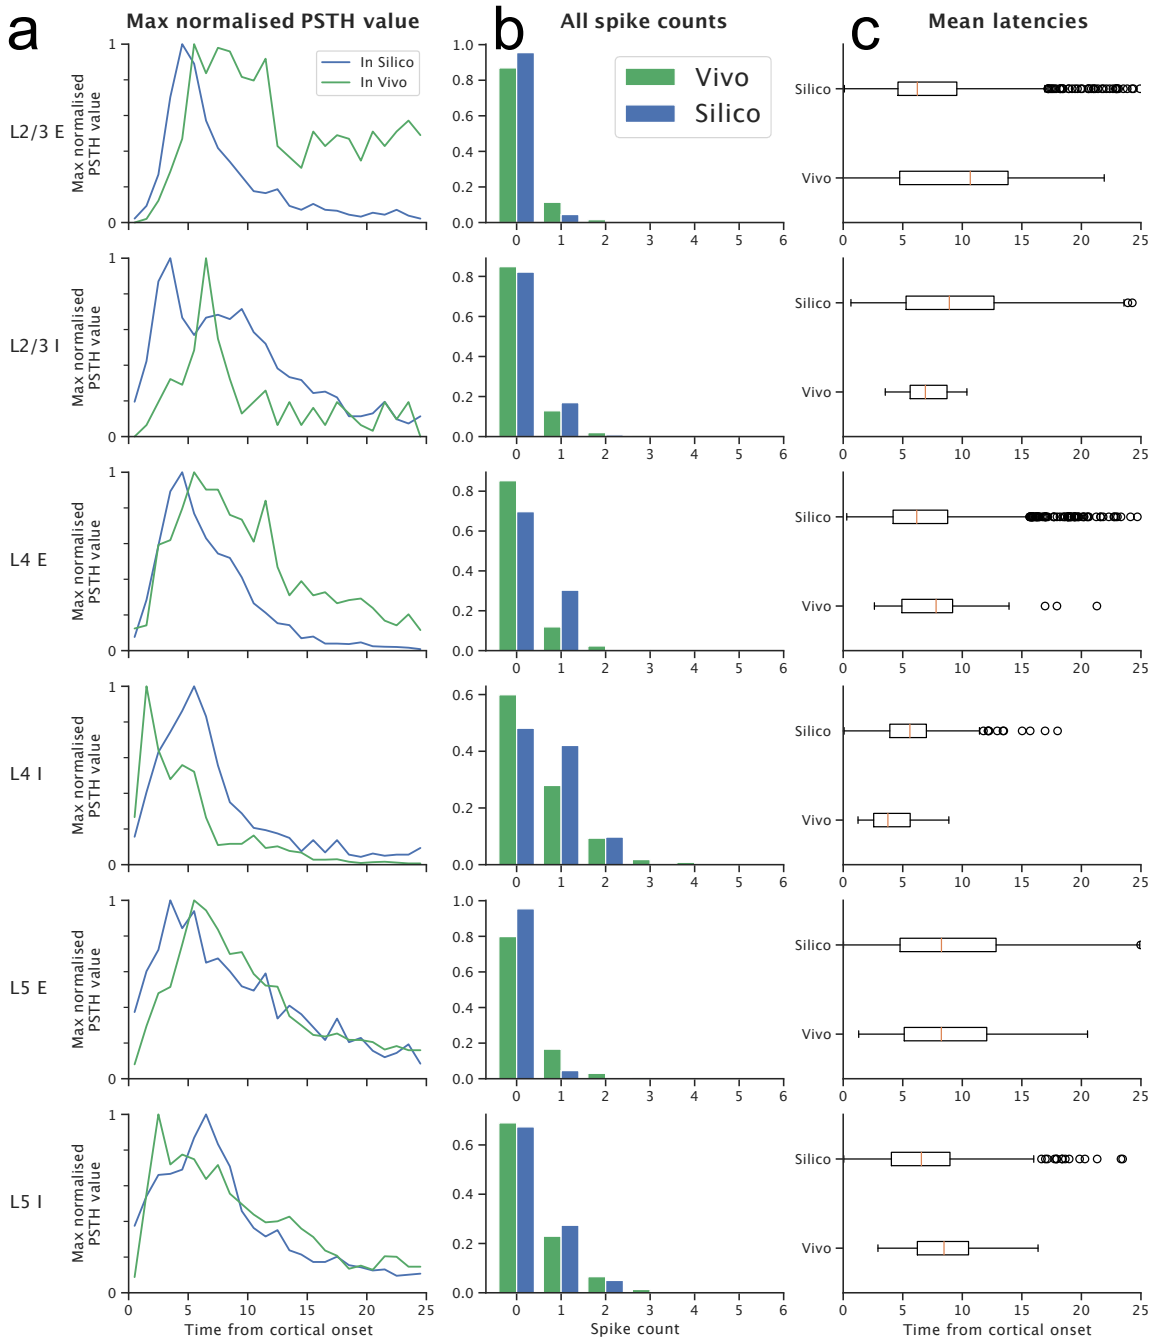


Figure 5.10: Analysis of spiking activity for in silico parameter combination 1 (depolarisation: 100%, $[Ca^{2+}]_o$: 0.9mM, fibers: 0.5, input temporal width: 5ms). (a) Lines show values of max normalised PSTH profiles for each comparison neuron group (1ms bins) for in vivo case (green) and in silico parameter combination (blue). (b) Spike count histograms for each comparison neuron group for the in vivo case (green) and in silico parameter combination (blue). (c) Box plots of spike latencies for each comparison neuron group for the in vivo case and in silico parameter combination.

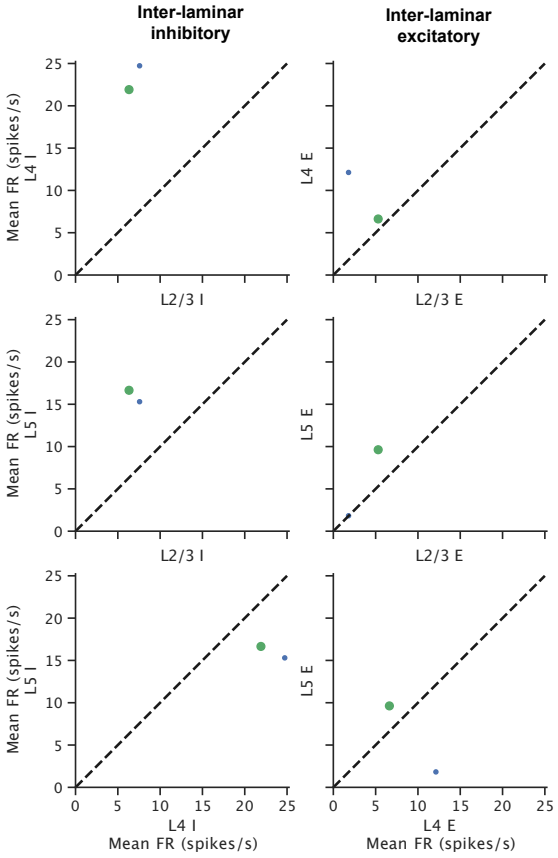


Figure 5.11: Comparison of mean firing rates between pairs of neuron groups for in silico parameter combination 1 (depolarisation: 100%, $[Ca^{2+}]_o$: 0.9mM, fibers: 0.5, input temporal width: 5ms). Each green dot represents the values for the in vivo case, whilst each blue dot represents the value for the in silico parameter combination. Columns show comparison of mean firing rates between pairs of layer-wise inhibitory neuron groups (left column) and between pairs of layer-wise excitatory neuron groups (right column).

from the analysis of spike counts for in silico parameter combination 1 (Fig. 5.10b). Excitatory firing rates show poorer similarity, as expected from the previous analysis of firing rates for all in silico combinations (Fig. 5.7).

5.6.4 Evoked firing rate similarity to in vivo

This subsection studies the effect of the 4 parameters on the similarity of evoked in silico firing rates with in vivo. For this quantification, a simple 'mean neuron group firing rate' is calculated for each parameter combination by averaging the firing rates of the 6 neuron groups to get a single value. This form of average weights the firing rates of each neuron group equally, but other population firing rate metrics could be used. For example, the firing rates of the 6 neuron groups could be represented as a point in 6-dimensional space, where each dimension represents the mean firing rate of one neuron group. The Euclidean distance between in vivo and in silico points in this space could then be calculated as a measure of similarity to in vivo. As expected increasing the depolarisation, $[Ca^{2+}]_o$-level and the proportion of input fibers in the stimulus, increases the current measure (mean neuron group firing rate; Fig. 5.12).

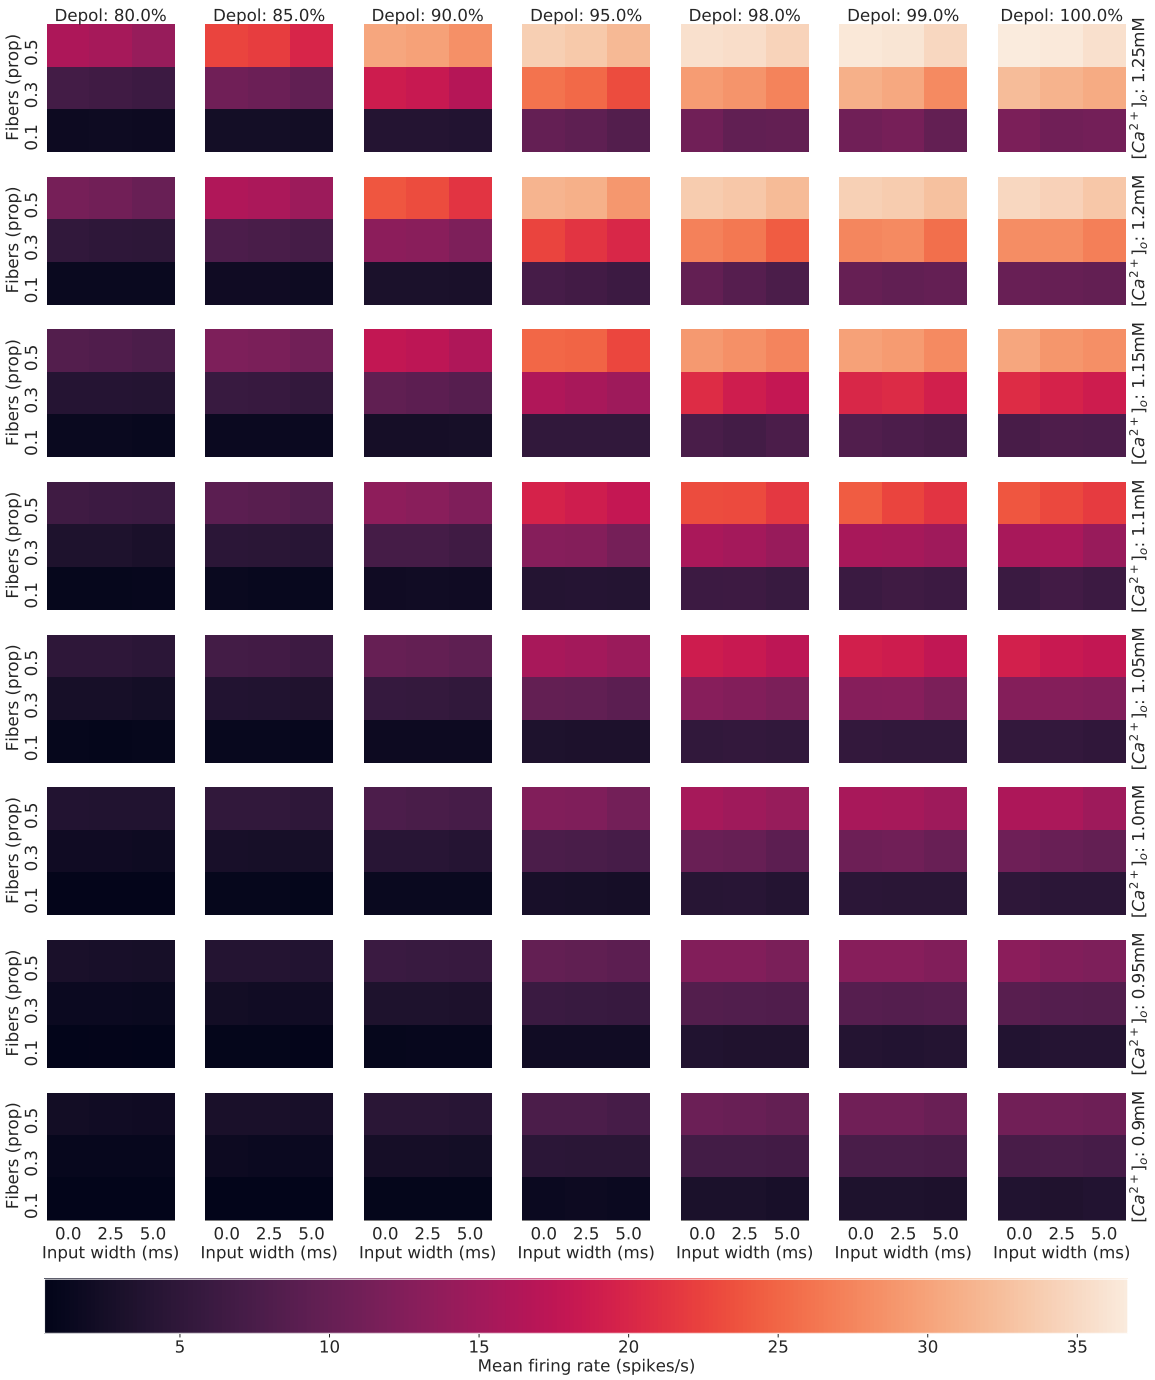


Figure 5.12: Heatmap showing the effect of the 4 parameters on the evoked 'mean neuron group firing rate', calculated by averaging the firing rates of the 6 neuron groups. Major columns: depolarisation, major rows: $[Ca^{2+}]_o$, minor columns: input temporal width, minor rows: proportion of fibers.

Decreasing the temporal width (i.e. increasing the synchronicity) of the input pattern also increases the mean neuron group firing rate.

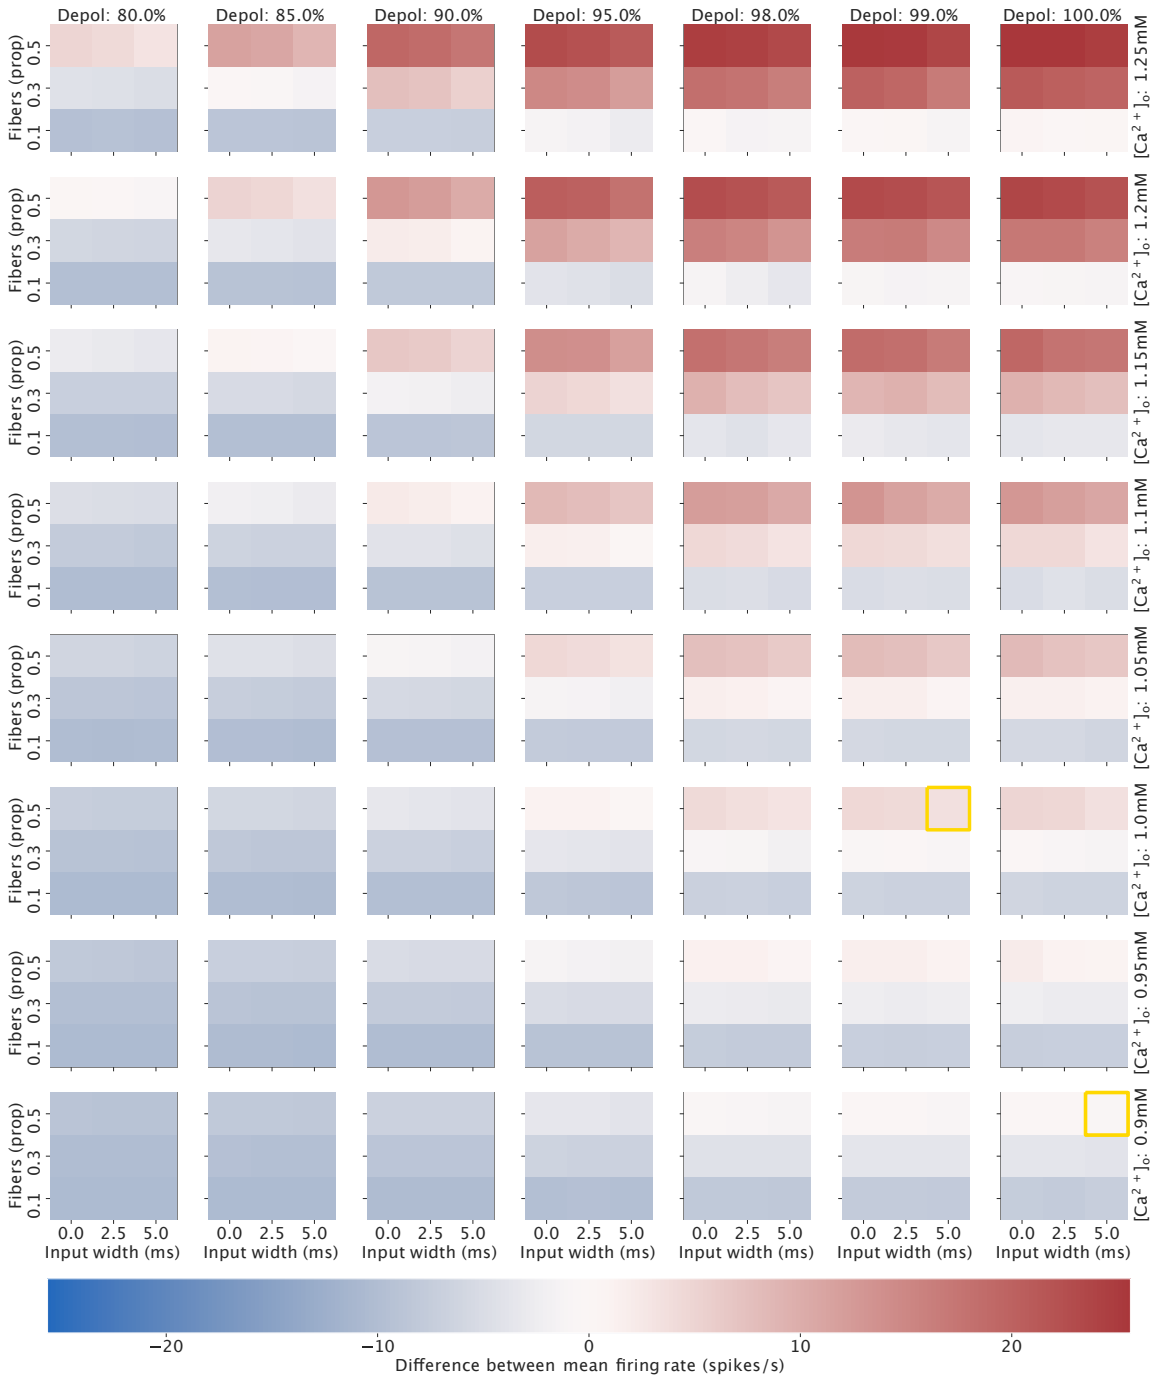


Figure 5.13: Heatmap showing the effect of the 4 parameters on the value of the evoked in vivo mean neuron group firing rate subtracted from the in silico mean neuron group firing rate. Major columns: depolarisation, major rows: $[Ca^{2+}]_{\circ}$, minor columns: input temporal width, minor rows: proportion of fibers. The values of in silico parameter combinations 1 and 2 are highlighted with a yellow box for visualisation purposes.

More informatively, the difference between in vivo and in silico mean neuron group firing rates can be plotted. Fig. 5.13 shows the value of the in vivo mean neuron group firing rate subtracted from the in silico mean neuron group firing rate. As expected from Fig. 5.12, higher values of depolarisation, $[Ca^{2+}]_{\circ}$, the proportion of input fibers in the stimulus and the synchronicity of the input, makes the mean neuron group firing rate higher in silico than in vivo, and vice versa. For all depolarisation and $[Ca^{2+}]_{\circ}$ values it appears that there may be a combination of input fibers within the stimulus and input synchronicity for which the mean neuron group firing rates would match in vivo. Fig. 5.13 shows that in silico parameter combination 1 shows close firing rate similarity to in vivo, as expected from the preceding analysis.

An in vivo recording bias for larger and more reliably firing neurons would lead to a recording of higher in vivo population firing rate values than true population firing rate values. Bias introduced by the spike overlap problem would lead to the opposite effect for high stimulus-evoked activity (i.e. a recording of lower in vivo population firing rates than the true population firing rates). This chapter does not estimate how much both of these biases would affect the recorded firing rate. When interpreting the similarity of in vivo and in silico mean neuron group firing rate similarity, it is important to consider that the in vivo mean neuron group firing rate is likely to contain some recording bias. A following subsection will consider calculation of ratios between spontaneous and evoked firing rates, both as an important metric for validation and as a means to reduce the problem of in vivo recording bias in firing rate comparison metrics.

5.6.4.1 Evoked spontaneous firing rate analysis

So far, the chapter has analysed stimulus-evoked activity. An important validation of the model is that in silico and in vivo spontaneous activity should also show close resemblance (accounting for possible sampling biases). This is especially important for the test of the time-warp phenomena, which is hypothesised to reflect an interaction between sensory stimulation and the trial-to-trial differences in the shared level of cortical excitability. The ratio between spontaneous firing rates and stimulus-evoked firing rates offers another important validation, which as a ratio, may reduce the problem of bias for larger and more reliably spiking neurons. Spontaneous firing rates are analysed first in this section.

As was done for stimulus-evoked firing rates, a single mean neuron group firing rate is calculated for each in silico parameter combination by averaging the mean firing rates of the 6 neuron groups. A 50ms window of spiking activity preceding

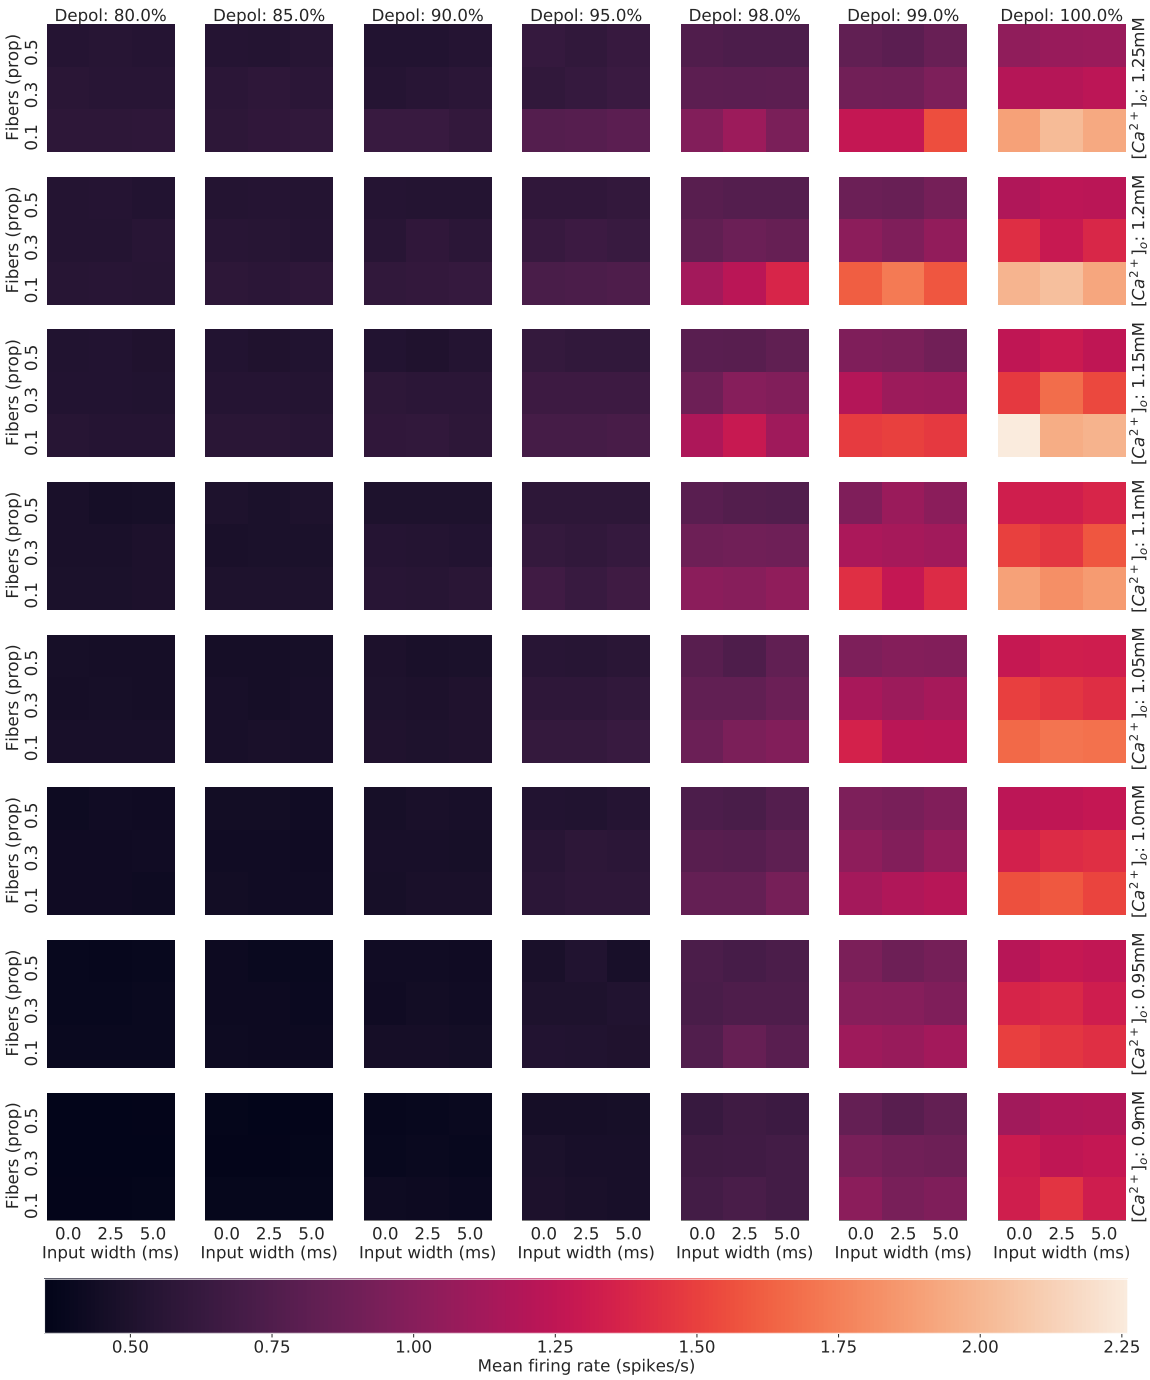


Figure 5.14: Heatmap showing the effect of the 4 parameters on the pre-stimulus spontaneous activity 'mean neuron group firing rate', calculated by averaging the firing rates of the 6 neuron groups. Major columns: depolarisation, major rows: $[Ca^{2+}]_o$, minor columns: input temporal width, minor rows: proportion of fibers.

stimulus-onset is used for the calculation making the activity 'pre-stimulus' spontaneous activity. Fig. 5.14 shows the effect of the four parameters on the in silico

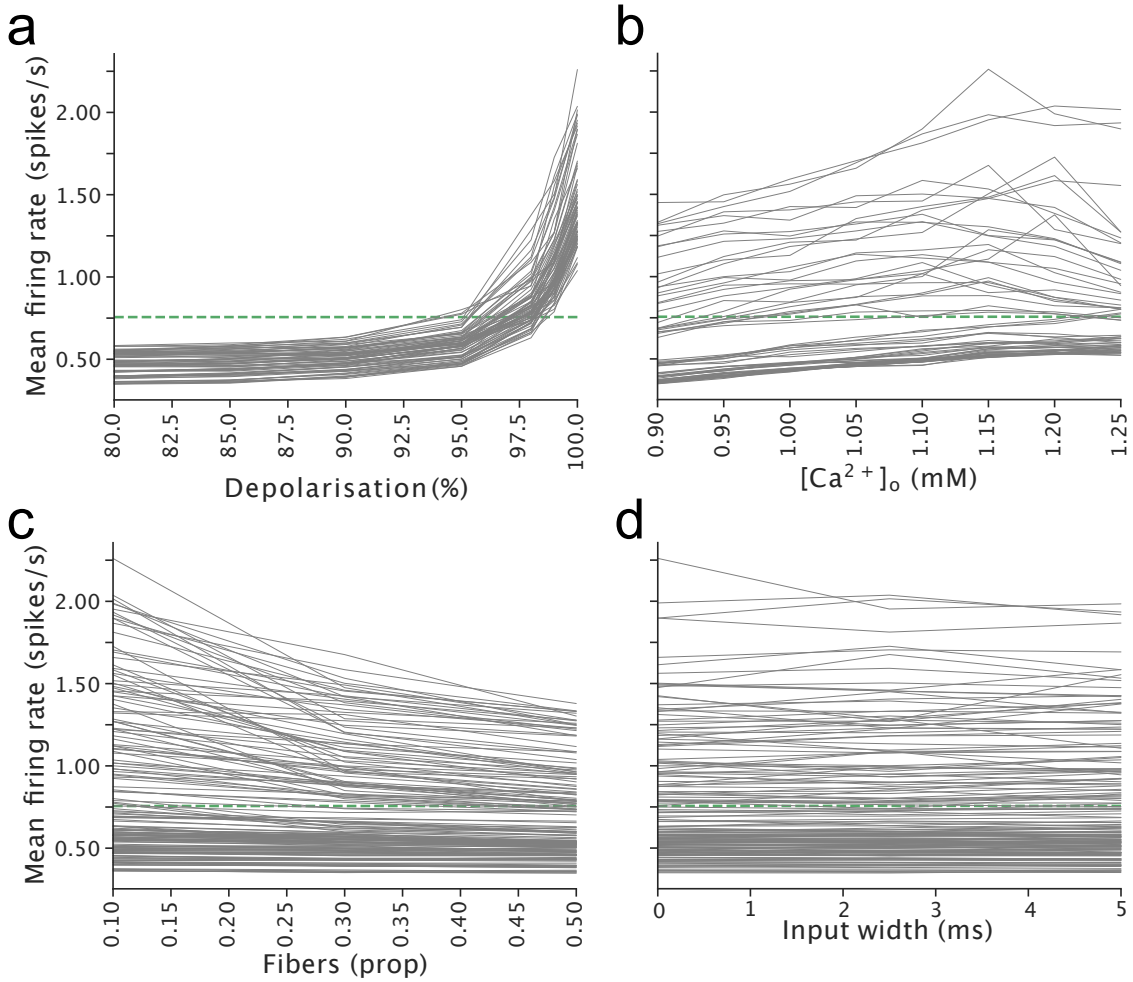


Figure 5.15: Line plots showing the effect of each of the 4 parameters on the pre-stimulus spontaneous activity 'mean neuron group firing rate', calculated by averaging the firing rates of the 6 neuron groups. Each line shows the effect of varying the parameter of interest, whilst the other 3 in silico parameters are fixed. The horizontal green line shows the in vivo value for comparison. The four plots show the effect of (a) depolarisation, (b) $[Ca^{2+}]_o$, (c) the proportion of fibers and (d) the input temporal width.

mean neuron group firing rates. As expected, increasing depolarisation and $[Ca^{2+}]_o$ increases spontaneous firing rates (also shown by Fig. 5.15a,b). Interestingly, the number of fibers in the stimulus decreases the rate of spontaneous firing rates, showing that the stimulus has a negative effect on spontaneous activity 450ms after the stimulus (Fig. 5.15c). This unexpected effect should be checked for networks with different initialisation random seeds but could be due to the higher stimulus-evoked inhibitory activity relative to excitatory activity, as characterised in the preceding sections. Finally, the temporal width of the input has no observable effect (Fig.

5.15d).

5.6.4.2 Spontaneous firing rate similarity to in vivo

Fig. 5.15a shows that a minimum depolarisation of $\sim 93 - 94\%$ is needed for the spontaneous activity mean neuron group firing rate to match in vivo-levels, although depolarisation between $\sim 95 - 98\%$ is needed for most in silico parameter combinations. Fig. 5.16 shows the effect of the 4 parameters on the similarity of in silico mean neuron group firing rates to the in vivo case. An in vivo recording bias for larger and more reliable neurons would make the in vivo spontaneous activity mean neuron group firing rate higher than the true in vivo value. If this were the case, then in silico parameter combinations with spontaneous activity mean neuron group firing rates lower than the in vivo case may better reflect the in vivo case.

Fig. 5.16 shows that in silico parameter combination 1 shows higher spontaneous firing rates than in vivo. The following subsection studies the spontaneous firing rates in more detail (i.e. by neuron group). For the in silico time-warp test it is perhaps preferable to have higher spontaneous activity rather than lower, as the phenomena is hypothesised to reflect an interaction between sensory stimulation and the trial-to-trial level of cortical excitability.

5.6.5 Ratios of spontaneous firing rates between neuronal groups

As was done for stimulus-evoked firing rates in Section 5.6.2, Fig. 5.17 compares spontaneous firing rates between neuron groups with the goal of reducing biases which may be consistent between groups. The best in silico correspondence (with respect to in vivo) between inhibitory groups is for L4 I and L5 I, as all in silico combinations lie close to $y = x$, as for the in vivo case. The correspondence between L2/3 I and the other inhibitory groups is less similar to the in vivo case, possibly due to the over activation of the in silico L2/3 I group with respect to in vivo.

For all in silico parameter combinations, the relationship between the spontaneous activity firing rates of the L2/3 and L5 E groups does have a similar ratio as the in vivo case (central column, central row, Fig 5.17) but are of lower magnitude. This result suggests under activation of these 2 groups, as for the stimulus-evoked case (Section 5.6.2). As for L4 E activity, the firing rates are closer to the in vivo case for some of the in silico parameter combinations. For the firing rate ratios of the L2/3 E and L5 E groups with the L4 E group, under-activation of L2/3 E and L5

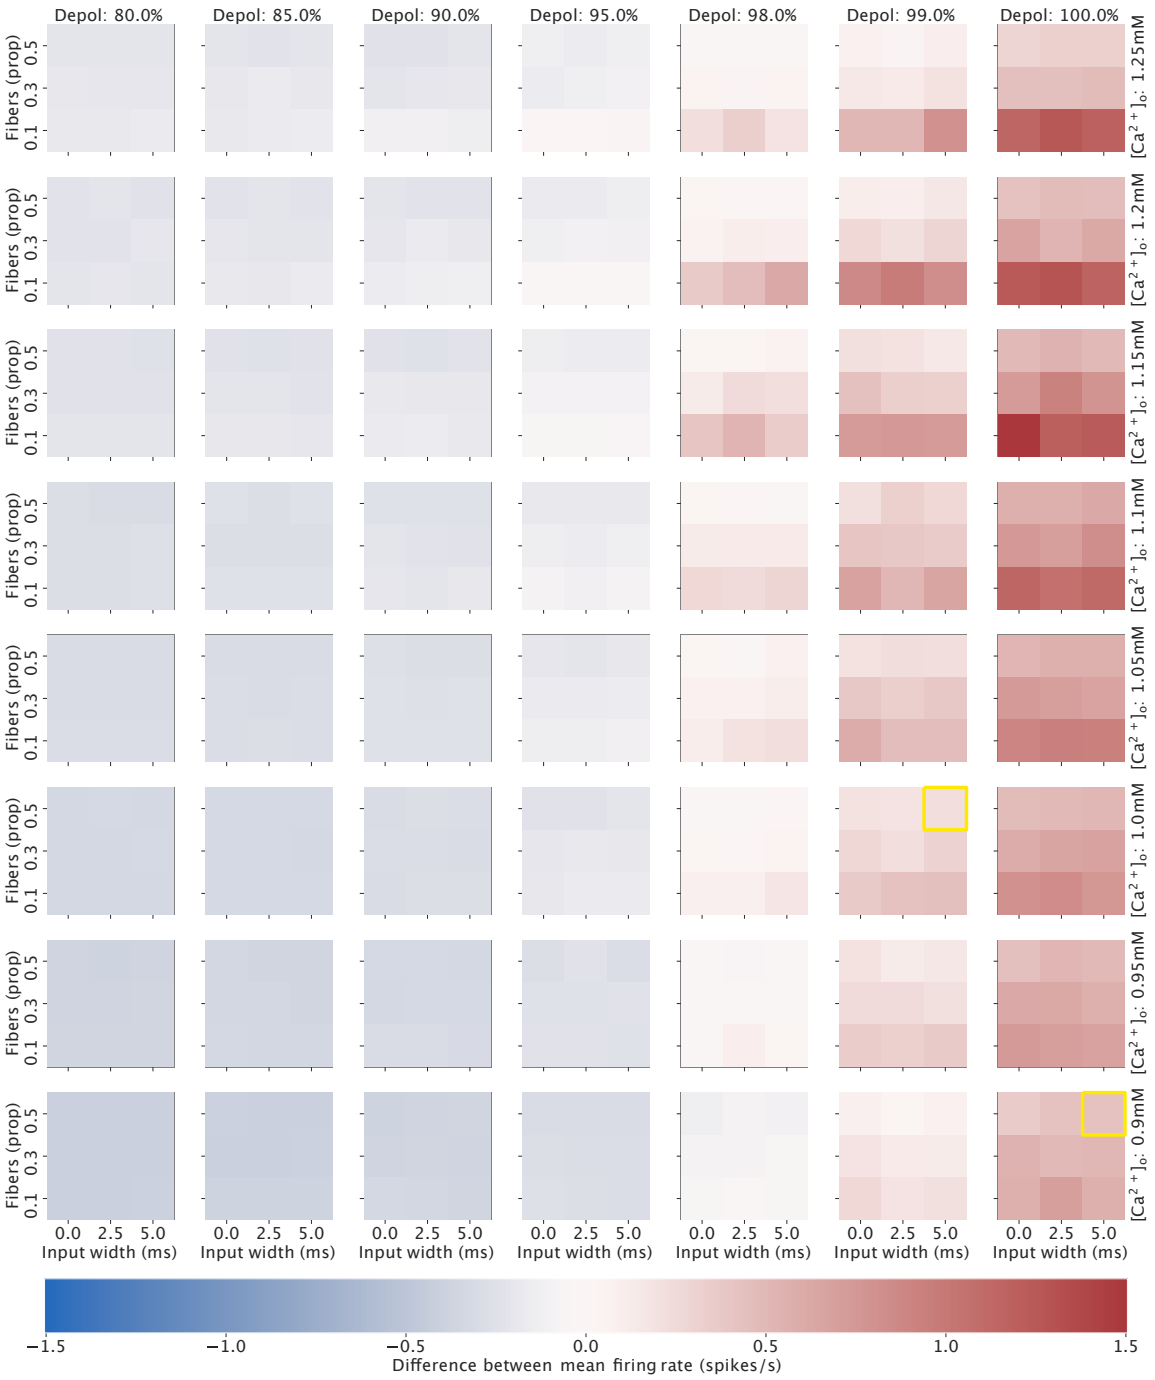


Figure 5.16: Heatmap showing the effect of the 4 parameters on the value of the in vivo spontaneous activity mean neuron group firing rate subtracted from the in silico spontaneous activity mean neuron group firing rate. Major columns: depolarisation, major rows: $[Ca^{2+}]_\circ$, minor columns: input temporal width, minor rows: proportion of fibers. The values of in silico parameter combinations 1 and 2 are highlighted with a yellow box for visualisation purposes.

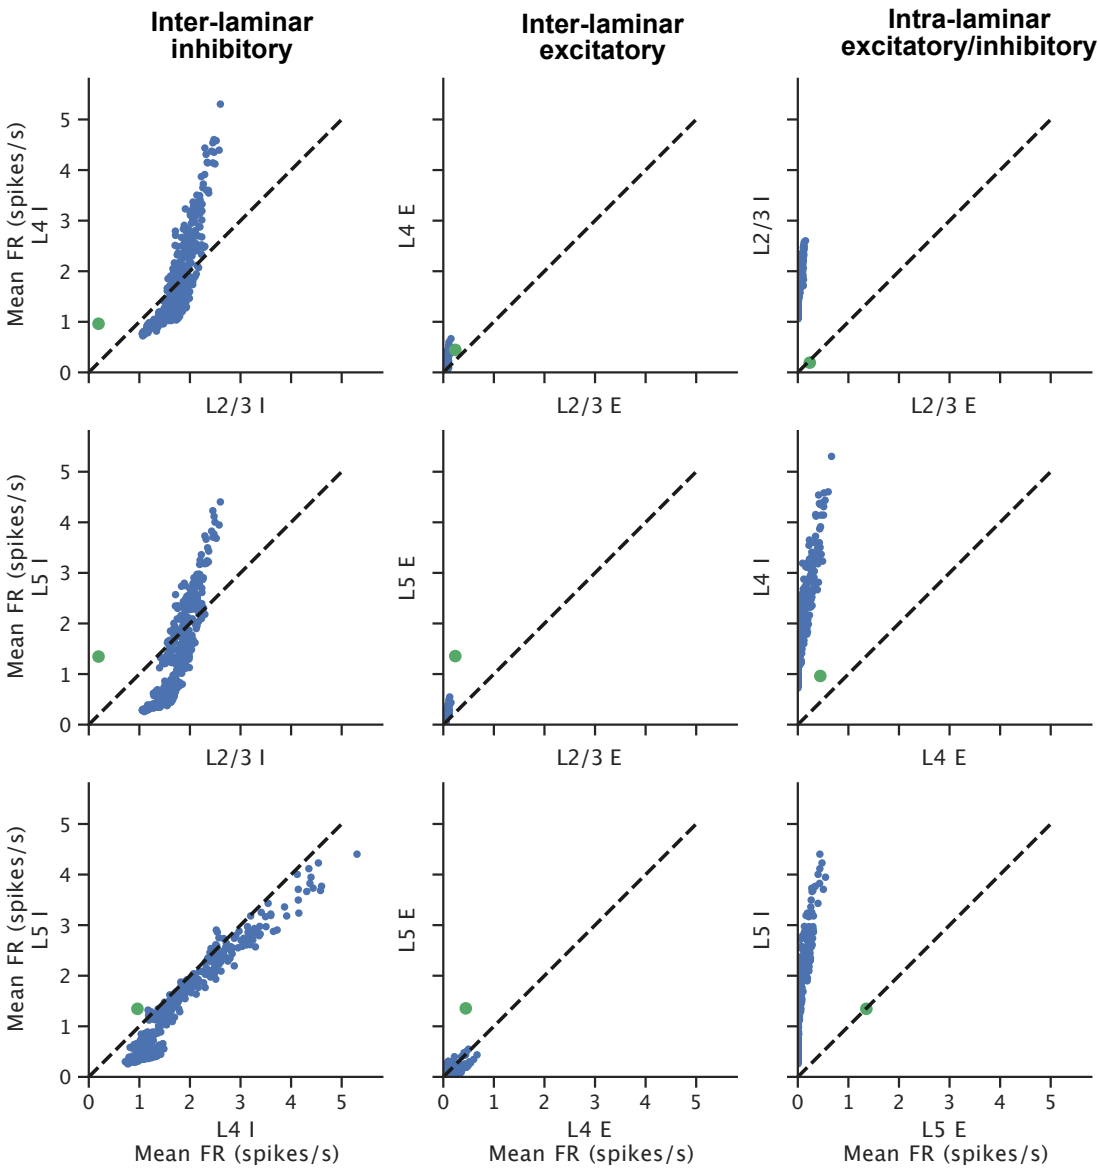


Figure 5.17: Comparison of spontaneous activity mean firing rates between pairs of neuron groups. Each green dot represents the values for the in vivo case, whilst each blue dot represents the values for a single in silico parameter combination. Columns show comparison of spontaneous activity mean firing rates between pairs of layer-wise inhibitory neuron groups (left column), between pairs of layer-wise excitatory neuron groups (central column) and between inhibitory and excitatory neuron groups within single layers (right column).

E again seems to explain the discrepancies between the in vivo and in silico cases. The final column of Fig 5.17 shows that, in silico, the excitatory neuron groups are under activated with respect to the inhibitory neuron groups. This under activation of excitatory groups could be caused by over activation of inhibitory neuron groups.

5.6.5.1 Ratio between spontaneous activity and stimulus-evoked firing rates

Finally, the ratio between spontaneous and stimulus-evoked firing rates offers an additional validation, which, as a ratio, may again partly account for in vivo recording biases for larger and more reliably spiking neurons. It is important to consider that the spike overlap problem is likely to have a larger effect on stimulus-evoked activity, when in vivo firing rates are higher. Under this assumption, the recorded in vivo ratio between stimulus-evoked firing rates and spontaneous firing rates would likely be lower than the true value. An in silico ratio that is slightly higher than the in vivo case may therefore be closer to the true in vivo case.

The ratio is calculated by taking the ratio between the stimulus-evoked and spontaneous firing rates for each neuron group and then taking a mean of this value, to give a 'mean neuron group stimulus-evoked/spontaneous firing rate ratio'. To calculate a measure of difference to in vivo for the purposes of visualisation in Fig. 5.18, the log of the in vivo mean neuron group stimulus-evoked/spontaneous firing rate ratio is calculated and subtracted from the log of the in silico mean neuron group stimulus-evoked/spontaneous firing rate ratio (for each in silico parameter combination). Fig. 5.18 shows the effect of the 4 parameters on the difference between the in silico and in vivo log of the mean neuron group stimulus-evoked/spontaneous firing rate ratio. In silico combinations with a mean neuron group spontaneous firing rate ratio below 0.6 spike/second were greyed out of the visualisation. Fig. 5.18 again shows that in silico parameter combination 1 shows close similarity to in vivo (with the caveat of in vivo recording bias).

5.6.5.2 Selection of model for time-warp test

Based on the similarity of its spiking activity to in vivo, for a range of measures, the decision was made to use in silico parameter combination 1 for the test of the time-warp phenomena. For reference, in silico parameter combination 1 is defined by the following parameters:

- Depolarisation: 100%
- $[Ca^{2+}]_o$: 0.9mM
- Fibers: 0.5 (proportion)
- Input temporal width: 5ms

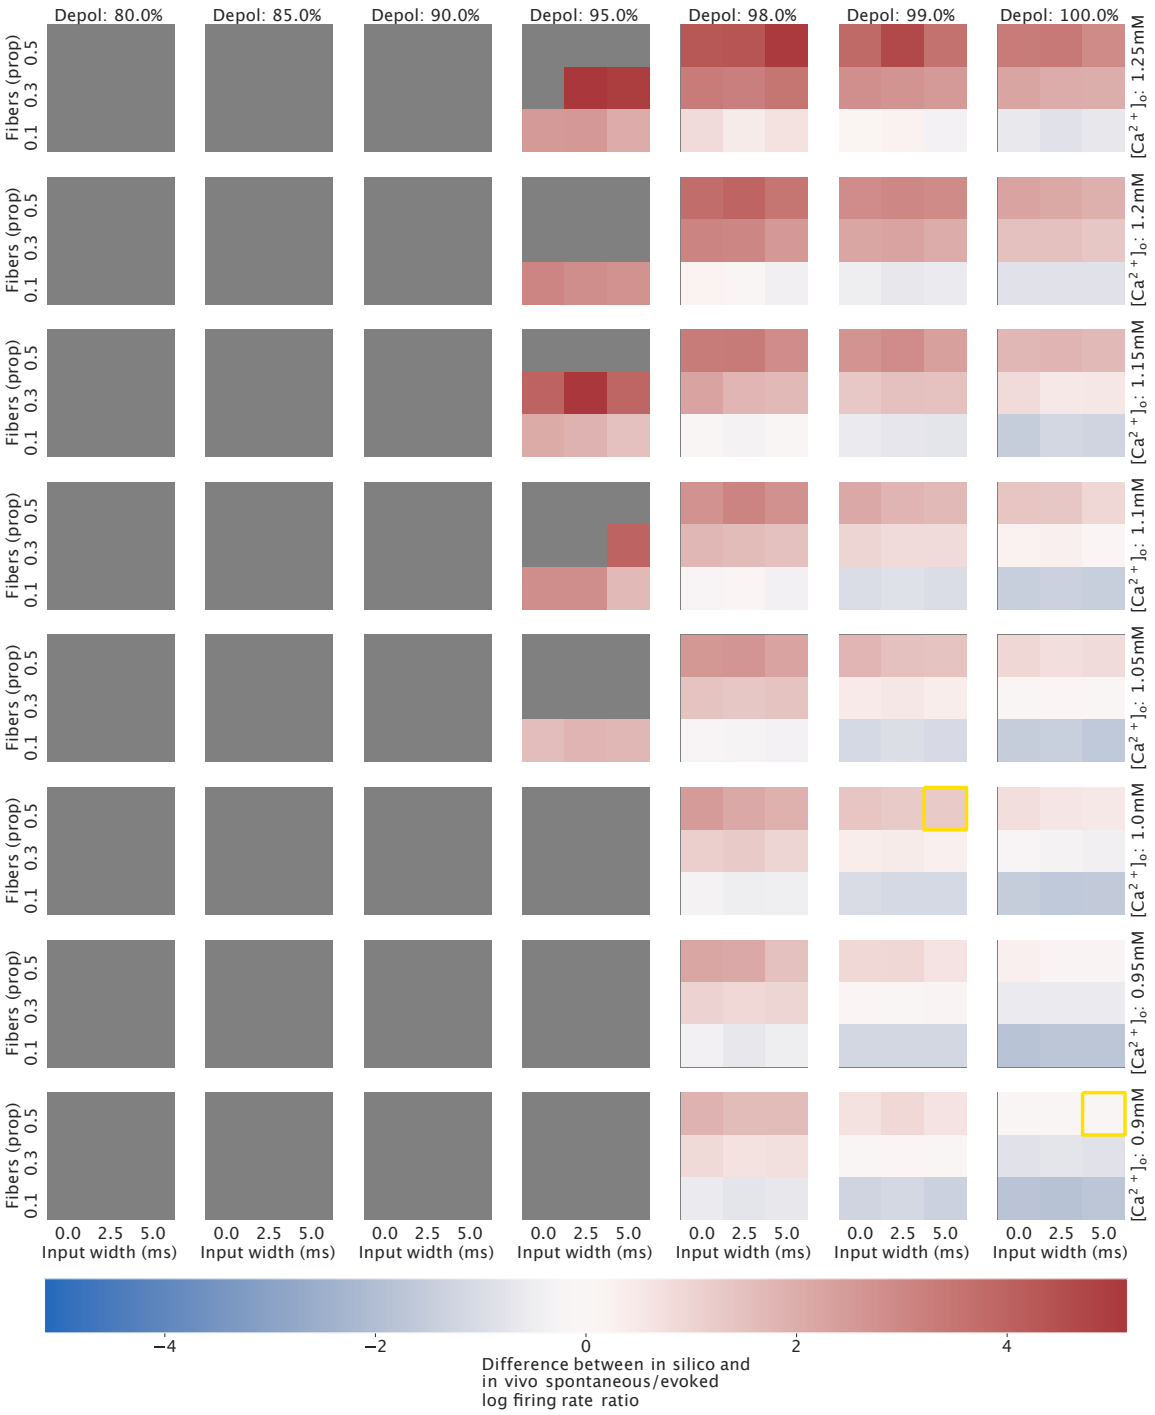


Figure 5.18: Heatmap showing the effect of the 4 parameters on the difference between the in silico and in vivo log of the mean neuron group stimulus-evoked/spontaneous firing rate ratio. The values of in silico parameter combinations 1 and 2 are highlighted with a yellow box for visualisation purposes.

As has been shown, the model represents the best choice in terms of similarity of neuron group PSTH profiles, and shows good similarity to in vivo for neuron group

spike time latencies, neuron group stimulus-evoked spike counts / firing rates, spontaneous firing rates and the ratio of spontaneous and stimulus-evoked firing rates (Figs. 5.10 and 5.11). As highlighted, all parameter combinations shows dissimilarity to in vivo for excitatory neuron group spike counts / firing rates and ratios thereof. Whilst a large proportion of time-warp clusters were found in vivo (Chapter 3) for pairs of excitatory neurons in L5, the decision was made not to select in silico parameters which yielded excitatory firing rates closely resembling in vivo. Instead, the decision was made to select parameters, which overall, demonstrated a first order spiking regime that was most similar to in vivo. In this regard, in silico parameter combination 1 offers a good candidate for the time-warp test for the reasons just outlined. It is planned that the the time-warp detection technique will be applied to future versions of the model, which may show closer resemblance to in vivo.

5.7 Methods: in silico test for time-warped spike time noise correlations

5.7.1 Simulation

To test for the time-warped phenomena, in silico parameter combination 1 was run for 102.25 seconds of model time. During this time, 200 thalamic stimulations were made at 500ms intervals starting after 2 seconds of model time. The thalamic input pattern was the same over all 200 thalamic stimulations as defined by the parameters of in silico parameter combination 1 (proportion of input fibers: 0.5, input pattern temporal width: 5ms).

5.7.2 Analysis

The analysis pipeline from Chapter 3 was adapted and applied to the 200 in silico stimulus trials. The same in vivo analysis parameters used in Chapter 3 were used for the in silico analysis. For example, in the clustering step, the DBSCAN epsilon parameter was increased in 0.05 increments over the range $[0.4, 5.05]$. The time-warp phenomena was tested for 5000 pairs of neurons in each comparison neuron groups (L2/3 E, L2/3 I, L4 E, L4 I, L5 E, L5 I). For example, the time-warp phenomena was tested for 5000 neuron pairs in neuron group L5 E. To provide an additional test of the correlated cluster detection algorithm, correlation surrogate control datasets were generated from the in silico first spike times for pairs of neurons as described

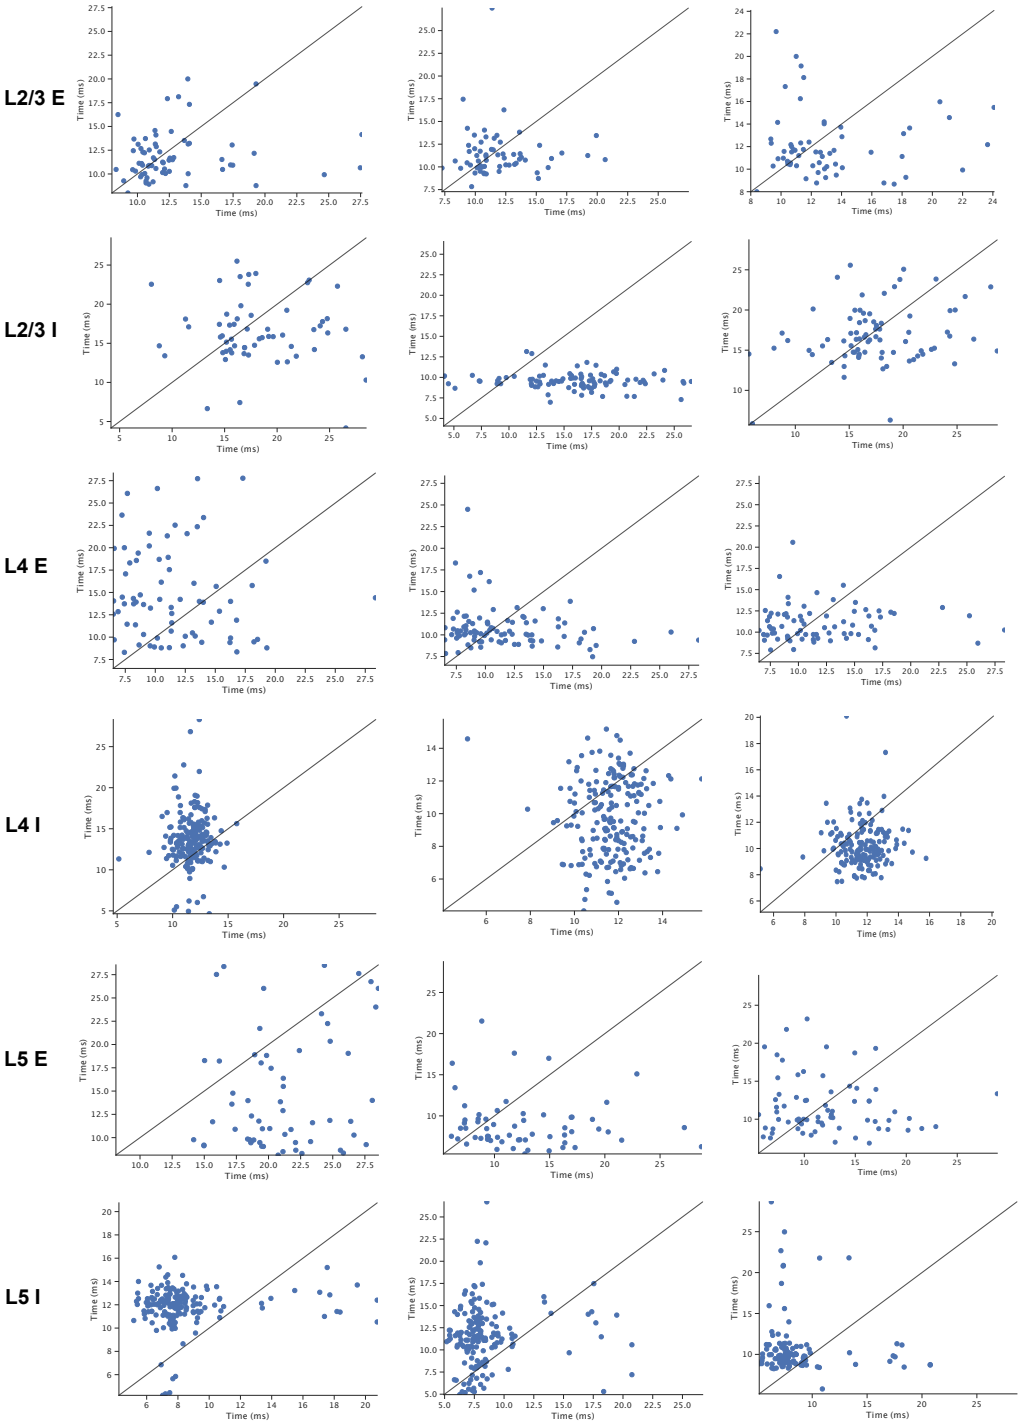


Figure 5.19: Example first spike pair sample response distributions. For each neuron group, sampled first spike pair response distributions are shown for 3 random pairs of neurons in the neuron group. For each pair of neurons, the first spike times of the two neurons are plotted against each other for trials of the stimulus on which both neurons spiked at least once.

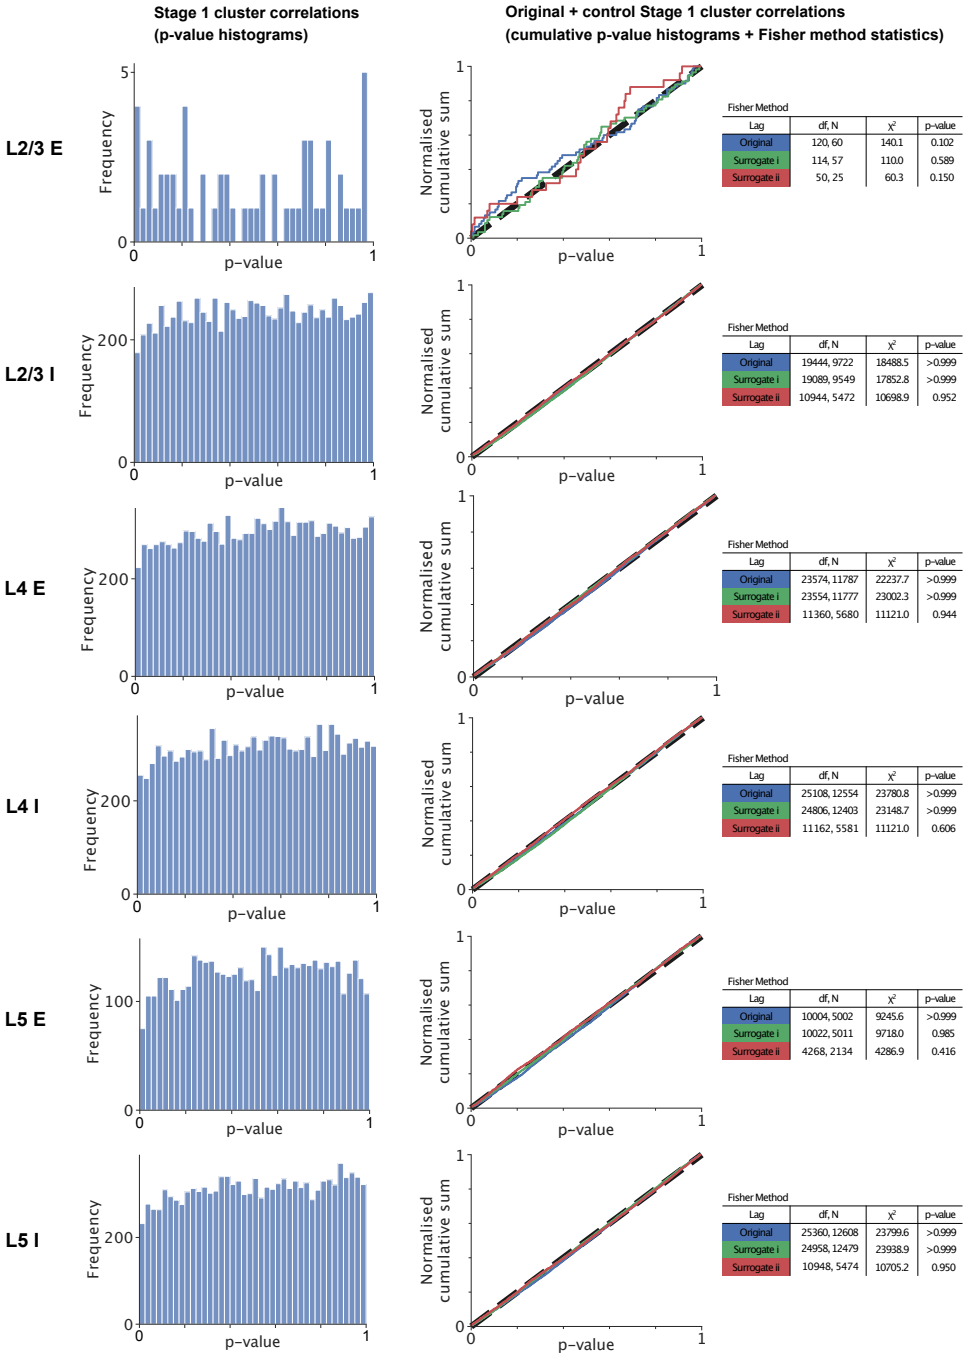


Figure 5.20: Stage 1 cluster correlations for original and surrogate control datasets by neuron group. First column: histograms of the linear correlation p-values of the extracted Stage 1 clusters. Central column: normalised cumulative sum of the linear correlation p-values of the extracted Stage 1 clusters for the original (blue) and surrogate control (red, green) datasets. Right column: the results of Fisher method analysis applied to the 3 datasets. Fisher method p-values show that the distribution of linear correlation p-values are not above chance.

for the in vivo dataset in Section 3.5.2. The surrogate control datasets are described here again for reference:

(i) *All shuffled.* Pairwise shuffling of elements of first spike pair distribution. For each first spike pair stimulus-response distribution, first spike pairs are shuffled so that each spike time of one neuron is matched with a different spike time of the other neuron.

(ii) *Single Gaussian.* Spike pairs drawn from an independent bivariate normal distribution with the same mean and standard deviation as the original spikes. The number of samples drawn from the distribution was equal to the number of original samples.

The correlated cluster detection algorithm was reapplied to these surrogate control datasets to confirm that correlated clusters were only detected at chance levels for these surrogate control datasets.

5.8 Results: in silico test for time-warped spike time noise correlations

First spike time scatter plots for pairs of neurons showed qualitative differences to in vivo (Fig. 5.19). Particularly, there was no obvious appearance of correlated clusters. Application of the analysis pipeline from Chapter 3 confirmed that correlated clusters were only found at chance-levels for each of the comparison neuron groups (Fig. 5.20). This result suggests that in the in silico model, first spike times of pairs of neurons are not linearly modulated by latent processes that change the shared excitability-level. Future work will explore whether further biologically grounded additions to the model lead to the emergence of trial-to-trial changes in shared excitability, which modulate the first spike times of pairs of neurons.

Fig. 5.20 shows that correlated clusters were only found at chance levels for each of the surrogate control datasets generated for each neuron group, providing a further validation of the correlated cluster detection pipeline.

5.9 Discussion

5.9.1 Summary

Sections 5.5 and 5.6 outlined a set of measures for comparing cortical in vivo and in silico spiking activity. These measures aimed to partly account for possible sampling

biases of extracellular in vivo recording and spike sorting. These measures provided an improvement for comparison of the in silico biophysical cortical models with in vivo. For example, previous comparisons of the model with in vivo spiking activity used visual comparisons of figures produced for the in vivo and in silico cases [175]. Analysis of the Allen institute model of the primary visual cortex [23] does not explicitly take into account the possibility of in vivo sampling biases and demonstrates strikingly large differences to in vivo spiking activity on the millisecond and centisecond (1/100th of a second) scales.

The measures and comparisons to in vivo were intended as a first step towards testing whether an in silico cortical model is close to a true in vivo regime, with the primary aim to provide a network which could be tested for the time-warp phenomena. The set of comparisons were also chosen to act as a first 'benchmark', which could guide data-driven refinement of the model. Analysis of neuron group mean firing rates for the spontaneous regime is sub-optimal compared to characterisation of firing rate distributions, but the metrics offered quick, straightforward and often signed measures for understanding and determining appropriate network parameters for the time-warp test. With the large amount of data generated by hundreds of large scale simulations under a variety of network states, the metrics provided an interpretable set of analyses, which have since been built upon and improved.

It is hoped that the methods will be developed further and used towards further model refinement in future. The measures outlined in this chapter offer the possibility of generating an 'in silico model score' based on it's similarity to in vivo. For example, the differences of in silico values to the in vivo values for the various measures could be summed to create a single score. The demonstration that in vivo-like distributions of neuron group spike counts and mean first spike latencies follow for some neuron groups when PSTH profiles closely match in vivo, suggests that improvements to the balance of activity between neuron groups could further improve neuron group mean spike count and mean latency similarity to in vivo.

It is important to note that the BBP cortical column model was described as a first draft by its authors [175]. Given the complexity of the model it is natural to find discrepancies between in silico and in vivo activity. The comparison tools (i.e. implementation of the methodology) were designed with generality in mind so that future iterations of the cortical models could easily be compared to in vivo activity and tested for the time-warp phenomena. The tools are also designed to compare in vivo and in silico at multiple time points of in vivo and in silico experiments. For

example, the tools and data could allow comparison of in vivo and in silico spiking activity under stimulus response adaptation or plasticity.

5.9.2 Suggested future measures

After considering the literature on extracellular recording bias and developing the validations in this chapter, the thesis proposes the following future comparisons of in vivo and in silico spiking activity (as basic validations of cortical activity):

1. *Modelling and comparison of firing rate distributions.* As discussed by [294], the presence of non-Gaussian long-tailed firing rate distributions for both spontaneous and event-related activity in vivo, mean that comparisons of firing rate distributions would be more preferable than comparisons of mean firing rates. Direct access to the barrel cortex dataset means that firing rate distributions can be characterised by neuron group. As done by [294] for in vivo spiking activity, it could be tested whether in silico firing rate distributions are best fit by exponential, log-normal or gamma-distributions.

2. *Layer-wise patch clamp data.* Layer-wise firing rate distributions calculated using patch-clamping would be more optimal for providing non-activity biased spiking measures. Automated patch clamping techniques may provide such data in the near future [120], although limitations on the number of cells which can be patched simultaneously would limit tests for the time-warp phenomena in the same dataset. Such a dataset would allow direct comparison between in vivo and in silico firing rate distributions.

3. *Firing rate distribution sub 1Hz peak.* As highlighted by [294], previous characterisations of firing rate distributions consistently report a peak firing rate below 1Hz during both spontaneous and event-related activity. This provides a simple validation of cortical models for future work.

4. *Similar firing rate distributions for spontaneous and event-related activity (with extension of long-tail for event-related activity).* Contrary to the common conception, [294] reports that the distribution of firing rates for event-related activity is similar to the distribution of firing rates for stimulus-evoked activity. [294] report that long-tailed distributions of firing rates for event-related activity only extend the long-tail of firing rate distributions and that the distributions remain similar. This provides an interesting validation of the model for stimulus-evoked

activity, that would validate that the sparsity of stimulus-evoked activity is partly comparable to in vivo.

5. *Sparsity measures.* It may be possible to create validations which compare the total number of active neurons during spontaneous and stimulus-evoked activity. This would require careful selection of an appropriate in vivo dataset (preferably patch clamp data). The juxtasomally and patch-clamped recordings of [54, 302] by layer and cell-type in barrel cortex may be useful in this regard.

6. *More complex measures relating to the test for the time-warp phenomena* are discussed below. This includes a test for the distance dependence and dimensionality of spontaneous noise correlations that was previously made for the in vivo datasset used here [233].

Since the completion of the chapter, the author has developed several of these analyses (discussed in the following subsection).

5.9.3 Current use and new analyses

The work of this chapter has been used to guide refinement of a new cortical model, which is currently in development. Particularly, where in silico activity has shown large discrepancies with in vivo activity, the validations have given clues as to where the model may be missing important data. This has lead to refinement of thalamic input patterns, thalamocortical innervation patterns and thalamocortical synaptic conductances, all through the integration of additional data in the literature. Moreover, the validataions have inspired the development of a more advanced artificial depolarisation method (compensating for missing synaptic inputs).

Since the completion of the chapter, the author has made the following additions to the in vivo comparison framework. As these have been applied to a currently unpublished model, the results are not presented.

1. *Long-tailed distribution test and sub 1Hz peak test.* 3 long-tailed distributions are now fit to the firing rate distributions of each target neuron group. A boolean test has been added for each target neuron group, which checks whether the peak of the best fit long-tailed distribution is less than 1Hz. This provides 6 (1 for each target neuron group) boolean values which are used in the calculation of the improved model score described below.

2. *Bursting test.* A bursting test has been added by comparing the ratio of the minimum and maximum firing rates of the population during spontaneous activity. This provides a single value for the calculation of the improved model score described below.

3. *Sustained evoked response test.* An automated test was added to test whether evoked response were sustained (or increased) over unnaturally long time periods following sensory stimulation. This provides 6 (1 for each target neuron group) boolean values which are used in the calculation of the improved model score described below.

4. *E/I spiking correlation test.* The correlation between the excitatory and inhibitory population firing rates is now calculated, summarising the correlation between the E/I populations. R-values above .75 correlate with bursting-like activity in spike rasters. Simulations with r-values above .4 and below .75 corresponded with simulations in which the spiking activity of the I population closely followed the E population. For these simulations the firing rates stayed within a controlled range (i.e. not demonstrating bursting-like behaviour). Moreover, the firing rates of the E and I populations varied slowly with time. These simulations are likely to represent simulations with good E/I balance and emergent oscillations due to the E/I balance. These emergent oscillations are likely to correspond with shared fluctuations in membrane potentials, which could lead to the timewarp phenomena. A single boolean test has been added to the model score, particularly testing whether the correlation r-value is above .4 and below .75.

5. *Improvement to model score.*

   The current model score combines the following tests:

   - Sub-1Hz firing rate test per target neuron group (see above).
   - Firing rates less than in vivo. This accounts for extracellular recording bias by checking that firing rates are less than the extracellularly recorded values (see above).
   - Non-bursting (see above).
   - Non-sustained or increasing stimulus-evoked activity (see above).
   - E/I firing rate correlation r-value between 0.4 and 0.75 (see above).

- Euclidean distance between concatenated PSTH profiles.
- Michelson contrast of spontaneous firing rate distances of in silico and in vivo firing rates. Although not described here, this provides an improved metric over the Euclidean distance, which punishes zero firing rates.

The author plans to soon add an explicit metric for the closeness of inter- and intra-layer firing rate ratios to in vivo.

6. *Single simulation summary pages.*

   To better visualise the analyses of each simulation, a single analysis page is now produced for each simulation. The page visualises the spiking rasters for spontaneous and evoked activity, as well as the results of the other important analyses. The pages from a single campaign of simulations are concatenated into a single video and pdf for ease of viewing.

7. *Multiple simulation summary.*

   To accompany the single simulation summaries, a single pdf is generated with the key meta-analyses of the simulation campaign. This includes visualisation of the effect of parameters on the simulation score.

8. *Refactoring to use Pandas dataframes.*

   The comparison framework has been refactored to use Pandas dataframes. This has allowed the output of the analysis to be used more easily by others in the research group.

9. *Standardised simulation test protocol and automated analysis.*

   A standardised protocol has been developed for testing single simulations. The simulation is run for 5 seconds of spontaneous activity followed by a single whisker deflection. The old and new tests are then run automatically for the produced spiking activity. This standardised addition has improved the rate of testing and refining the new cortical model.

5.9.4 In silico test for time-warped spike time noise correlations

Sections 5.7 and 5.8 tested whether the form of precise temporal structure characterised in Chapters 3 & 4 also existed for the in silico spiking activity. It was shown that the spiking activity of the model did not exhibit the precise temporal structure

characterised in Chapters 3 & 4. These differences were quantified by the methodology described in Chapters 3 & 4. Particularly, it was shown that correlated clusters were not detected above chance-levels for the in silico model. This supported qualitative observations for the first spike times of pairs of neurons within each of the 6 comparison neuron groups.

In future, the author hopes to apply more standard spike count noise correlation analysis to the in silico activity, and compare this with in vivo. This would provide an additional and more basic test for the presence of spiking activity modulation by shared latent factors. This should be accompanied by analysis of subthreshold membrane potential correlations between pairs of neurons, as well as analyses of latent modulation, which do not make assumptions about distribution structure (i.e. [237]). Importantly, the dimensionality and distance dependence of trial-to-trial changes in shared excitability have previously been characterised during spontaneous activity for the in vivo dataset used in this paper [233]. As the time-warp phenomena is suggested to depend on the shared excitability-level of neuron pairs at the point of stimulus onset, changes in shared excitability during spontaneous activity provides an important validation of model activity that can be made in future.

Such a set of analyses would offer a set of measures for understanding the conditions necessary for the presence and/or emergence of shared latent modulation of spiking activity, which appears to be missing in the current model. A large number candidates exist for such latent modulation (i.e. trial-to-trial changes in shared excitability). At this point, it is again important to consider the artificial depolarisation of neurons in the model. Particularly neurons are depolarised close to their spiking threshold to compensate for connections, which, in vivo, would originate from outside of the circuit. In the current model, low variance (0.001) is applied to this depolarisation. Under high $[Ca^{2+}]_{\circ}$ levels the model has previously been shown to exhibit epileptic spontaneous activity [175]. Under in vivo-like $[Ca^{2+}]_{\circ}$ levels, the model is unlikely to produce emergent oscillatory activity.

Trial-to-trial changes in excitability may emerge from network effects when larger areas of the cortex are simulated (currently in development). Alternatively, thalamocortical interactions are known to play an important part in cortical rhythmicity (see [204] for review). Current work at the BBP to model thalamo-cortical and cortico-thalamic connectivity is likely to play an important part of such dynamics and may be necessary for the emergence of the time-warp phenomena.

Voltage sensitive dye imaging (VSDI) is likely to be useful in determining shared fluctuations in subthreshold membrane potentials. For example, a new voltage sensitive-

dye was used to show delta (0.5–4 Hz), theta (4–10 Hz), low beta (10–20 Hz), and low gamma (30-40 Hz) oscillations in cortical L1 [53]. The technique is useful for studying oscillatory activity under different states and all oscillations were stronger in awake animals. Moreover, the power of delta, theta, and low beta oscillations increased with depth, the power of low gamma was more constant throughout L1. The authors concluded that the findings identify L1 as an important coordination hub for the dynamic binding process of neurons mediated by oscillations. Implementation of in silico VSDI imaging by the BBP (see [198]; described above) offers the opportunity to compare in silico and in vivo VSDI responses to ensure that fluctuations in subthreshold membrane potentials across the population are similar to the in vivo case.

Additional surrogate control datasets were also generated from the in silico spiking activity. The methodology from Chapters 3 & 4 was then applied to these additional surrogate control datasets. Correlated Stage 1 clusters were not found above chance-levels for the two surrogate control datasets, providing further validation of the analysis. Although, generated from different datasets, the two surrogate control datasets are likely to be quite similar to the 2 corresponding surrogate control datasets used in Chapters 3 & 4. The additional validation of the surrogate control datasets provided in this chapter offers only a limited additional validation of the methodology. Future work should still aim to improve the methodology and test it on a broader range of synthetic datasets.

Chapter 6

Hypotheses of time-warped multi-neuron single spike patterns

This chapter presents hypotheses of how time-warped multi-neuron single spike patterns may underlie the representation of information in the brain. The chapter also presents hypotheses of how such representations may emerge through activity-dependent plasticity. The hypotheses aim to address the GRQs outlined in Chapter 1. The chapter discusses how the outlined hypotheses could be tested in vivo and important considerations for doing so. It is possible that the hypotheses can be further refined through the results of in silico simulations. The hypotheses may also provide a roadmap for the development of a phenomological model that can be applied to sensory machine learning tasks. Such a model would also provide a proof-of-principle for the hypotheses. Chapter 7 describes an early attempt at developing such model.

6.1 General research questions

The General Research Questions outlined in Chapter 1 are given again as reminder:

- General Research Question 1 (GRQ 1): How are short sensory stimuli represented atomically in single cortical areas?
- General Research Question 2 (GRQ 2): What is the nature of continua encoding for atomic sensory representations?
- General Research Question 3 (GRQ 3): How are short sensory stimuli parsed (i.e. transformed) into neural representations? And how does the process of parsing provide selectivity (including against incorrect combinations of low-level features).

- General Research Question 4 (GRQ 4): How do continua representations emerge through experience?
- General Research Question 5 (GRQ 5): How is feature-property continua encoding (for features at any scale) affected when multiple overlapping cell assemblies are simultaneously activated?
- General Research Question 6 (GRQ 6): How is one-shot or multi-shot learning facilitated by sensory representations in the initial forward pass from sensory stimulation?

6.2 Hypotheses

Hypothesis 1: Single atomic entities are represented by time-warped multi-neuron single spike patterns.

It is hypothesised that time-warped multi-neuron single spike patterns are an atomic form of information representation through-out the brain. It is not suggested that time-warped multi-neuron single spike patterns are the only form of spike pattern based information encoding in brain, however. The possibility that multi-neuron single spike patterns are universal form of encoding across the brain is strengthened by evidence of low spike count responses in different brain areas [211, 232, 230]. For example, the study of [230] shows that patterns of largely single spike pattern responses in primary visual cortex are sufficient to enable visual discrimination [230].

The presence of multi-neuron single spike patterns in response to atomic stimuli could be tested for in other brain regions. Stimulus-response paradigms in which movement of the animal is minimised provides the simplest paradigm for testing for the presence of time-warped multi-neuron single spike patterns. The analysis should be applied to data from experiments in which the animal is fully awake, although this may add further noise to stimulus response distributions. The addition of additional sources of noise could obscure the presence of time-warped noise correlations. Separating variance caused by different sources of noise may be critical to uncovering time-warping of spike-time representations by trial-to-trial changes in shared excitability.

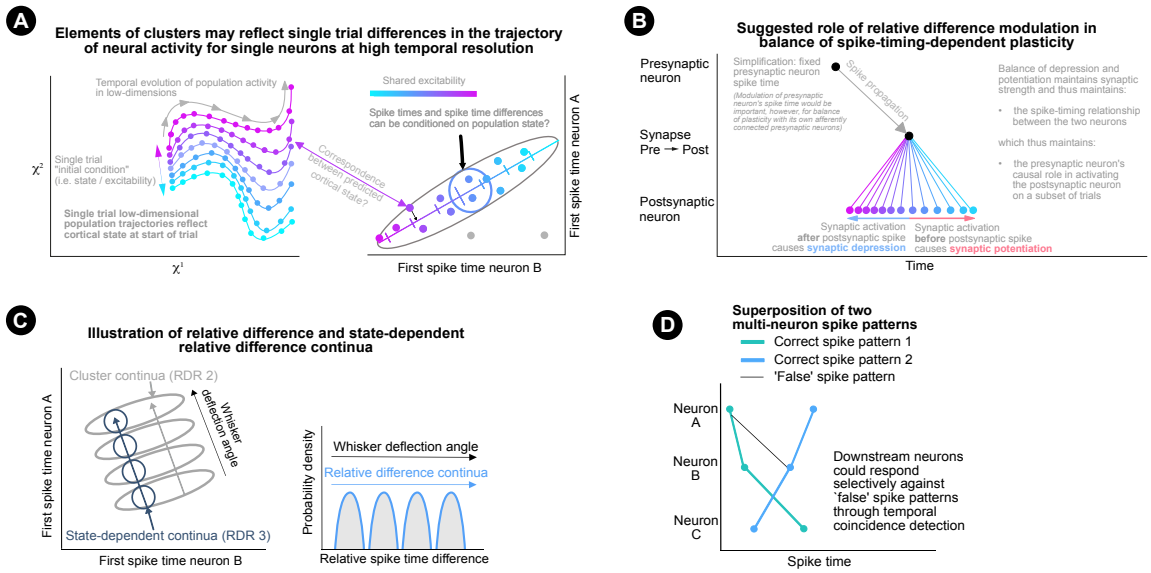


Figure 6.1: (a) Illustration of possible correspondence between single trial variability in low-dimensional population dynamics and single trial spike times for correlated clusters. (b) Left: Illustration of continua of first spike time distributions, corresponding to first spike times (not relative differences) under non-state-dependent (grey) and state-dependent (dark blue). Right: Illustration of continua of actual relative spike time difference distributions, which could correspond to non-state-dependent or state-dependent continua. (d) For two neurons, illustration of two overlapping multi-neuron spike patterns (red and purple). Coincidence detection could play a part in robustness to 'false' spike patterns (black). That is downstream neurons could respond selectively to particular spatiotemporal patterns of spikes. This may be achieved through the spatial organisation of neurons, synaptic innervation patterns and variable conduction delays, such that spike time patterns corresponding to spatially and temporally localised activations of the postsynaptic neurons are more likely to make the post synaptic neurons spike.

Hypothesis 2: Single trial time-warping is correlated with trial-to-trial differences in the trajectory of population activity in low-dimensional space.

Single trial variability in population dynamics is thought to reflect differences in population excitability or state at the start of a trial [143]. As such, population dynamics may offer a proxy for a cortical state, which modulates spike time patterns. Time-warped multi-neuron single spike patterns may demonstrate the effect of population state (normally analysed for population dynamics) for single neuron spike times and pairwise spike time differences. In another sense, the time-warped multi-neuron single spike patterns may reflect variability in the trajectory of population dynamics at a single neuron level and at high temporal resolution (Fig. 6.1a).

The characterisation of shared variability on a continuous spike timing scale could offer the opportunity to study population dynamics without loss of information through discretisation of neuronal responses, which typically uses bins of 10-20ms. The precise structure and latent factor dependence of relative difference representations uncovered in Chapters 3 & 4 suggests that such representations may encode information that is additional to information decodable from classical binned spike count analysis (including classical modelling of population dynamics).

Hypothesis 3: The degree of trial-to-trial time-warping is correlated with shared trial-to-trial changes in neuronal membrane potentials.

This may be simple to study for single neurons using previously collected single-neuron patch clamp recordings. Particularly, the correlation between a neuron's membrane potential before the onset of a stimulus and the latency of the first spike following sensory stimulation could be analysed in vivo. It would be surprising if such a characterisation had not been made previously. The results of Chapters 3 & 4 would predict heterogeneous relationships between the pre-stimulus membrane potential and first spike time latencies across neurons.

The hypothesis also requires that subthreshold membrane potentials are correlated within a population. This has been studied using simulataneous cell recordings of neuron pairs in different areas [220]. Simultaneous patch clamp recordings of 2 or more cells would allow the hypothesis to be tested directly (i.e. that time-warping between neuron pairs is correlated with shared trial-to-trial changes in neuronal membrane potentials).

In future, the author would like to further study what is known from voltage-sensitive-dye (VSD) imaging experiments. Interestingly, a new VSD was used to show delta (0.5-4 Hz), theta (4-10 Hz), low beta (10-20 Hz), and low gamma (30-40 Hz) oscillations in cortical L1 [53], and oscillations were stronger in awake animals. Correlations between VSD imaging an the degree of time-warping could also be explored. It is also possible that the trial-to-trial differences in the shared excitability of neuron pairs (and larger groups of neurons) are governed by an alternative mechanism.

Hypothesis 4: Downstream neurons decode information encoded by time-warped multi-neuron single spike pattern depending on the shared cortical excitability.

As discussed in Chapters 1, 3 & 4, the finding that spike time patterns are modulated by low-dimensional trial-to-trial changes in excitability over multiple layers and columns of cortex, introduces the possibility that downstream decoding neurons are modulated by the same changes in excitability. Downstream neurons may then decode afferent spike time patterns in a state-dependent manner, thus conferring an improved potential coding capacity over classical fixed-state decoding.

Hypothesis 5: Time-warping of multi-neuron single spike patterns balances potentiation and depression of spike-time-dependent plasticity in order to maintain spike time patterns/relationships.

It is hypothesised that the purpose of relative spike time difference modulation may be to maintain spike time patterns/relationships by balancing spike-timing-dependent synaptic depression and potentiation. The reasoning for this hypothesis is given in Fig. 6.1b. This hypothesis was partly inspired by the initial spiking neural network simulations of the DPhil project (Chapter 7), which showed the author that over potentiation by repetitive presentation of the same spike time input pattern can lead to an over potentiation, which leads to the causal role of presynaptic spikes in driving the postsynaptic spike being lost.

Hypothesis 6: Information is represented along time-warped multi-neuron single spike pattern continua, offering a higher encoding capacity.

It is hypothesised that information is encoded along time-warped multi-neuron single spike pattern continua (Fig. 6.1c, left). Particularly, it is hypothesised that continua may be encoded in the cortex by time-warped relative spike time differences (between the single spike times of different neurons in a multi-neuron single spike pattern). In the context of stimulus-encoding this would mean that small changes in the stimulus would correspond with small changes in the distribution form which multi-neuron single spike patterns were drawn (Fig. 6.1c, right). The potential of downstream neurons to discriminate or determine the distribution from which time-warped multi-neuron single spike patterns were drawn for similar stimuli would determine the precision of such an encoding form.

Moreover, decoding of information along such time-warped relative difference continua may be state-dependent. Then, if precise state-dependent representations vary smoothly with changes in stimulus parameters, the state-dependent precision of single neurons could enable each neuron (or pairs of neurons for spike time difference encoding) to encode single stimulus parameters. Under such a regime, the dimensionality of the neural manifold/continua would be closer to the number of neurons, making the encoding capacity of the population much higher than the non-state-dependent regime.

Hypothesis 7: Joint feature-presence and (continua) feature-property encoding.

It is hypothesised that feature-presence and feature-property information are encoded jointly by time-warped multi-neuron single spike patterns. Particularly, it is hypothesised that activation of the neurons in multi-neuron single spike pattern may encode the presence of a feature. In the context of the whisker system, activation of a particular subset of neurons could represent that a particular whisker has been stimulated over a range of deflection angles and speeds. It is hypothesised that the state-dependent temporal pattern of the multi-neuron single spike pattern encodes information regarding the properties of the feature, for example, the exact angle of the whisker deflection and its speed. Moreover, small changes in the transformation (i.e. properties) of the feature correspond with continuous changes in the distribution from which multi-neuron single spike patterns are drawn on single trials. It is necessary to say that the exact neurons which respond for different transformations (i.e. different properties) of the feature over a range of transformations may also change with continuous changes in the stimulus. It is hypothesised that the concept of the feature over a range of transformations is represented by activation of a subset of neurons which respond with some reliability to the presence of the feature.

Hypothesis 8: Overlapping multi-neuron single spike patterns.

It is hypothesised that single neurons can take part in multiple simultaneously active time-warped multi-neuron single spike patterns, such that single neurons can take part in the representation of multiple simultaneously presented features (Fig. 6.1d). For example, consider two features which in isolation would both activate a particular single neuron as part of two separate multi-neuron single spike patterns. Consider that some aspect of property information is encoded for the two features by the time

that the single neuron spikes with respect to other neurons in the multi-neuron single spike pattern.

If the two features were presented simultaneously, then it is possible that the neuron could spike twice at its usual time points in the two multi-neuron single spike patterns. In this way, the representation of the two features would not be lost. As such, it is hypothesised that multi-neuron single spike patterns could facilitate supposition of representations (with property information maintained).

It is possible, however, that the timings of spikes for the one feature could cause a downstream neuron to misinterpret the property information of the other feature. Coincidence detection by downstream neurons could decrease the occurrence of downstream neurons responding to or 'misinterpreting' spikes from different multi-neuron single spike patterns. Particularly, downstream neurons could be tuned to respond selectively to particular multi-neuron single spike patterns.

As neurons are sensitive to the temporal and spatial proximity of synaptic activations [92, 29, 30], selectivity to particular multi-neuron single spike patterns could be achieved through the synaptic plasticity, spatial organisation of neurons, patterns of synaptic innervation, and variable conduction delays. In this way, multi-neuron single spike patterns providing spatially and temporally localised activations of a postsynaptic neuron would be more likely to make a postsynaptic neuron spike compared to multi-neuron single spike patterns from the same presynaptic neurons, which do not deliver spikes to the postsynaptic neuron in close temporal proximity.

Hypothesis 9: Multi-neuron single spike patterns are compositional (property-dependent parsing and holographic projection of property information).

It is hypothesised that continua-encoding multi-neuron single spike patterns continua-encoding are compositional and general for lower- and higher-level features. Particularly, it is hypothesised that:

1. Property information (encoded by the precise temporal pattern of multi-neuron single spike patterns) is taken into account by downstream coincidence detecting neurons when parsing multiple low-level features (each encoded by a multi-neuron single spike pattern) into a high-level feature representation. Particularly, the high-level feature representation (most likely a cell assembly) is only

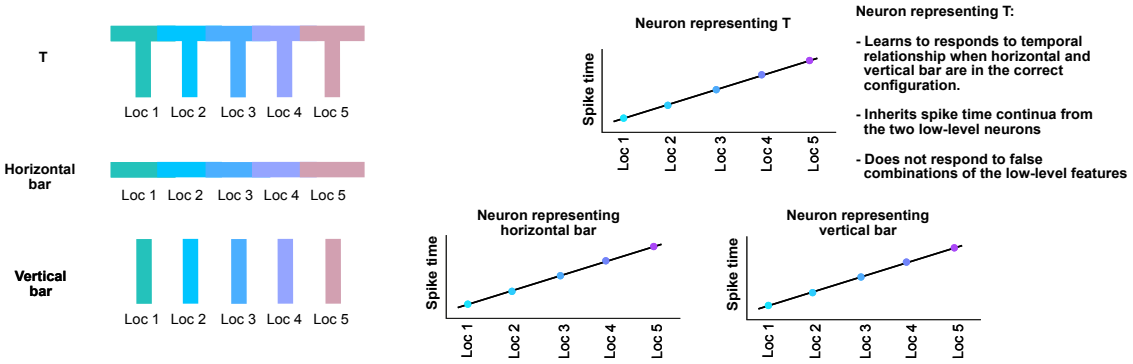


Figure 6.2: Compostionality explanatory example. Figure illustrates explanatory example. Left: Illustration of the horizontal bar, vertical bar and T is 5 locations. Right: Illustration of spike time continua for the horizontal and vertical bar in the 5 locations (bottom) and the inherited spike time continua for the T in the 5 locations.

'fully' instantiated if the properties of the low-level features are configured correctly with respect to one another. It is hypothesised that robustness to false combinations of low-level features is at least partly achieved in this way.

2. Property information is transferred from the low-level representation into the temporal pattern of the multi-neuron single spike pattern representation of the high-level feature. It is hypothesised that continuous changes in the constituent low-level features, translate to continuous changes in the activity pattern of the high-level representation. In this way, the multi-neuron single spike pattern representing the high-level feature represents the properties of the feature a long a continua. The encoding of property information for the high-level feature could then be used in the robust parsing of more complex features or to inform the motor-specifics of action.

Explanatory example and further details

A simplified example is given to illustrate the hypothesis and its further details. We consider a postsynaptic neuron which we would like to represent the presence and position of a T in 5 locations. We carry forward the idea that presence is encoded through activation of the postsynaptic neuron and that position is encoded through the timing of the neuron's spike (with respect to spike timing of other neurons in its local population).

The T in each of the locations shares some overlap with the T in the immediate neighbouring locations but not when the T is in the other locations (Fig. 6.2). In each location the square is represented in the input layer by neurons representing the

horizontal and vertical bars of the T (Fig. 6.2). We consider that the neuron is able to learn to respond the T in each of the 5 locations by potentiating synapses from the input neurons activated when the square is in each of the 5 locations.

Without specificity for the correct component subfeatures (i.e. the horizontal and vertical bar being in the correct location to form a T), the postsynaptic neuron would spike for incorrect combinations of the component subfeatures. For example, the neuron would respond to presentation of the horizontal bar of the T in location 1 and the vertical bar of the T in location 5.

We next suppose that there are neurons (i.e. cell assemblies) which respond in the presence of either the vertical or horizontal bar over the range of locations. For the purposes of explanation we suppose there are several neurons responding in the presence of the vertical bar over the 5 locations and several neuron responding in the presence of the horizontal bar over the 5 locations. Moreover, we suppose that these neurons represent the location of the bar over the 5 locations along spike time continua. These neurons are referred to as the low-level feature encoding neurons.

We next consider a postsynaptic neuron with afferent connections from this pool of low-level feature encoding neurons. We consider a training period, in which the T is presented repeatedly in the 5 locations. We hypothesise that the postsynaptic neuron would learn to potentiate synapses coming from the presynaptic neurons which maintain a relative spike time temporal relationship over the 5 locations (Fig. 6.2).

For example, we consider that there is one low-level neuron representing the location of the horizontal bar with a positive monotonic trend, and another low-level neuron representing the location of the vertical bar also with a positive monotonic trend. If the gradient of the two monotonic trends were similar, then the postsynaptic neuron would consistently receive the same temporal pattern of activation from the two neurons over presentations of the T in the 5 locations. This repetitive stimulation with the same temporal relationship could allow the synapses of the postsynaptic neuron to self-organise to consistently respond to the consistent spike time difference over the 5 locations. Particularly, the postsynaptic neuron may self-organise its afferent synaptic connections such that it spikes when the spikes, emitted in the consistent temporal relationship by the two neurons, arrive synchronously. There are multiple mechanisms through which tuning for certain temporal delays (either axonal or dendritic) could be achieved through plasticity.

Incorrect combinations of the the horizontal and vertical bar primitives would lead to less synchronous arrival of afferent spikes from the two low-level primitive neurons. The membrane potential would decay between the arrival of the first and

second spike, thus reducing the probability that the postsynaptic neuron spikes to the incorrect combinations of primitives.

With respect to compositionality, it can be considered that the postsynaptic neuron automatically encodes the location of the T along a spike time continua. When the two low-level feature encoding neurons spike early, the postsynaptic neuron will also spike early, and when the two low-level feature encoding neurons spike late, the postsynaptic neuron will also spike late. In this way, the postsynaptic neuron becomes a high-level feature encoding neuron, which encodes the location of the full T along a spike time continua. This T encoding neuron could then be used compositionally to create higher-level representations, with information encoded along continua using the same mechanism.

Hypothesis 10: One-shot and multi-shot learning with configurational specicifity and property information.

The above example can be used to highlight an additional requirement of the system. Firstly, a joint representation of a feature's presence and properties would ideally form from presentation of a minimal number of instantiations of the feature (possibly only one). In the example, it would be optimal if the transformation-invariant representation of the T formed from presentation of the T in only one or several of the five locations. This is referred to as 'one-shot' and 'multi-shot' learning in the machine learning literature.

This hypothesis extends the previous hypothesis and states that only one or several presentations of a stimulus may be needed for the postsynaptic neuron to self-organise such that it can respond robustly to the letter T (and not to false combinations of low-level features) and to encode the properties (i.e. location) of the T along a continua. Particularly, the postsynaptic neuron could potentiate synapses from the two low-level feature encoding neurons, through exposure to the stimulus in only one or several locations. That is, the neuron could learn to respond to the correct configuration of low-level features (and the consistent temporal spike relationship) from only one or several presentations of the stimulus.

The hierarchical organisation of sensory processing areas allows us to consider the problem of one shot learning with the given that there are a set of transformation-invariant low-level sensory primitives from which higher-level representations can be learned. It is widely accepted however, that the primary sensory cortices contain cells which respond to simple stimuli over a range of transformations. Furthermore, such primitives are thought to emerge as a result of both spontaneous activity (e.g.

retinal waves) and sensory exposure. Many theoretical studies have also shown how primitive representations can emerge in neural networks through exposure to sensory stimuli.

Hypothesis 11: Generalisation of motor commands

Precise spike times have been demonstrated in the motor cortex [268] and muscle units (the muscle fibers innervated by a single motor neuron) [264] (see [262] for review), supporting the possibility that continua and generalisation are spike timing based. Continua encoding offers a natural means by which action can be produced for features in novel configurations. Firstly, consider that activation of a cell assembly signals the presence of a feature, which in turn signals the opportunity for a subset of action-types. Secondly, consider that there exists a mapping from points on the continua representing previously experienced instantiations of the feature to the specific motor commands required to perform the action on these instantiations. Then, when the feature appears in a novel transformation it would instantiate an activity pattern in the feature-presence-encoding cell assembly that is some interpolation of activity patterns elicited for previous transformations of the feature. This representation would in turn instantiate an interpolation of motor commands for appropriate performance of the action upon the object. It is hypothesised that such continua are facilitated in the brain by time-warped multi-neuron single spike patterns.

Chapter 7

Co-creation of a spiking neural network simulator and the emergence of relative spike timing continua in a phenomenological spiking neural network.

This chapter presents work from the start of the DPhil, which aimed to explore the emergence of spike-timing feature binding representations through spike-time-dependent plasticity (STDP) in a multi-layer phenomenological spiking neural network with ranges of axonal conduction delays. Towards this, the author co-created an open source GPU spiking neural network simulator with other members of the research group, as available spiking neural network simulators were slow when heterogeneous axonal conduction delays and STDP were simulated. The simulator was designed to be 'delay insensitive' and optimised for STDP. That is, the introduction of action potential transmission delays between neurons in a network did not increase simulation time. SNN models using ranges of transmission delays are used to investigate a range of phenomena including sound localization, reservoir computing, auditory processing and visual feature binding [101, 215, 73, 70]

The simulator was shown to be the fastest in its class (as of March 2019), and a paper on this work was uploaded to the preprint server 'bioRxiv' [8]. Other research groups developing more widely used SNN simulators have since begun incorporating some of the simulator's optimisations. The simulator was used by other members of the research group for various research projects.

In the early stages of the DPhil, the simulator was used in several collaborative studies exploring the emergence of spike-timing representations in a multi-layer recur-

rent spiking neural network [70, 135]. The author produced the simulation code and redesigned the simulator for these studies, and also produced some simple analyses for the second publication [135], for which he was the first author. As the majority of the conceptual work and analysis for the two simulation studies was not done by the thesis author, the results of the two papers are not presented here.

This chapter gives a brief overview of the simulator [8] and unpublished simulations which aimed to provide a proof of principle of the hypotheses set out in Chapter 6. The unpublished simulations presented in this chapter explore how leaky integrate-and-fire (LIF) neurons learn to respond to simple one-dimensional stimuli over a range of translations through STDP. The early investigations related to this work and the author's consideration of neural coding and feature-binding, inspired the hypotheses of Chapter 6 and the decision to analyse in vivo spiking activity, as presented in Chapters 3 & 4, and published in Scientific Reports [136].

As stated, the development of the SNN simulator and the simulation work presented in this chapter ultimately aimed to explore whether the hypotheses and principles proposed in Chapter 6 could be fulfilled in a phenomenological SNN model. It was hoped that the simulations would demonstrate a proof of principle of the hypotheses. It was anticipated that complexities would arise but that the model would enable iterative improvements based on biological findings. It was hoped that a complete model could be scaled up, and applied to machine vision problems. These plans proved to be highly ambitious. The project was also postponed by over a year by computer hardware issues, which the author spent a large amount of time trying to solve. After returning the hardware for repair, the issues took the manufacturer 3 months to resolve.

After describing the SNN simulator, the chapter presents the first steps towards producing the proof of principle simulations. Particularly, the chapter presents simulations and analyses exploring how single spiking neurons learn (through STDP) to respond to a spike time encoded stimulus over a range of translations. Moreover, the analyses study how the transformation (i.e. location) of the stimulus becomes encoded by the first spike time of the neuron through learning. The chapter takes particular interest in how the spike time of the neuron varies continuously with continuous changes of stimulus location. Moreover, the chapter explores how the self-organisation of such 'spike time continua' may also facilitate robustness to false combinations of composing features. The investigation presents interesting results but also highlights major challenges, which should be addressed in future work, either through in vivo experiments or in silico modelling (either phenomenological or biophysical).

7.1 Background: Spiking Neural Network Simulators

7.1.1 Modelling of neurons, synapses and plasticity

SNNs consist of neurons, synapses and plasticity rules, which can each be modelled at different levels of detail. For example, neurons can be modelled with detailed morphologies and ion channels, or as abstract point-neuron models (i.e. neurons that have a single or small number of state variables). The SNN simulator described here, 'Spike', is designed to use point-neuron models. Point neuron models can be used in combination with various synapse models, for example, conductance and voltage based synapse models (described below). The dynamics of neurons, synapses and plasticity rules can each be described by Ordinary Differential Equations (ODEs).

7.1.2 Neuron models

For example, the simplest leaky integrate-and-fire (LIF) neuron can be modelled as:

$$\tau_m \frac{dV(t)}{dt} = V_0 - V(t) + I(t) \tag{7.1}$$

where $V(t)$ is the voltage of the cell on timestep t, V_0 is the resting potential of the cell and I(t) is the current injected into the cell on timestep t. If the membrane voltage rises above a threshold V_{thresh}, the neuron emits a spike and the membrane voltage is reset to the cells resting potential V_{rest}. A refractory period is imposed by fixing the membrane voltage at V_{rest} for a set duration.

7.1.3 Synapse models

7.1.3.1 'Voltage-based' synapses

The simplest synaptic model increases the voltage of the postsynaptic cell by the value of the synaptic weight (when the synapse is activated by a spike arriving from the presynaptic cell). The instantaneous activation of the synapse is modelled by the Dirac delta function. The equation of such a voltage based synapse is be given by:

$$I_{\text{syn}}(t) = w_{\text{syn}} \cdot \sum_f \delta(t - t^f_{\text{pre}}), \tag{7.2}$$

where w_{syn} is the synaptic weight, δ is the Dirac delta function, and t^f_{pre} is the set of times at which spikes arrive at the synaptic connection from the presynaptic neuron,

the time of each spike being indexed by f. The total input to the postsynaptic cell from a set S of such voltage-based synapses is then given by:

$$I(t) = \sum_{\text{syn}}^{S} I_{\text{syn}}(t). \tag{7.3}$$

7.1.3.2 'Conductance-based' synapses

Alternatively, the effect of a synaptic activation on the postsynaptic cell can be dynamically modelled. For example, the contribution of a synaptic activation on the postsynaptic cell can be modelled as a decaying process which is also coupled to the postsynaptic cell's membrane potential. The latter allows modelling of synaptic activations as causing changes in relative ion-concentrations. The effect of 'conductance-based' synapses on the membrane potential of the postsynaptic cell depends upon the conductance value of the synapse, $g_{\text{syn}}(t)$, which varies with time, and the reversal potential of the synapse, E_{syn}:

$$I_{\text{syn}}(t) = g_{\text{syn}}(t) \cdot (E_{\text{syn}} - V(t)). \tag{7.4}$$

Conductance-based synapses therefore provide an input which is proportional to the difference between the cell's membrane voltage and the reversal potential of the synapse, such that the synapses drive the membrane voltage towards the synapses reversal potential. As described, the conductance value $g_{\text{syn}}(t)$ can be modelled as a process which decays to zero after increasing by the value of the synaptic weight, w_{syn}, upon the arrival of a presynaptic spike:

$$\tau_{\text{syn}} \cdot \frac{dg_{\text{syn}}(t)}{dt} = -g_{\text{syn}}(t). \tag{7.5}$$

7.1.4 STDP

The simplest form of STDP is based on the original electrophysiological observations that a postsynaptic neuron spiking after the arrival of a presynaptic spike results in a positive synaptic weight change, and that the arrival of a presynaptic neuron spike after the postsynaptic neuron spike resulted in a negative weight change [176, 20, 304]. [20, 176] showed that the magnitude of the synaptic weight is related to the exponential of the pre and post-synaptic neuron spike timings. Such long-term positive and negative changes in the synaptic weight are known as long term potentiation (LTP) and long term depression (LTD) respectively. This dependence

of the synaptic weight change on the difference between the spike times of the pre t_{pre} and postsynaptic t_{post} spike times is captured by the following model:

$$\Delta w = \begin{cases} \alpha_+ \cdot e^{\frac{-(t_{\text{post}} - t_{\text{pre}})}{\tau_+}}, & \text{if } t_{\text{post}} > t_{\text{pre}} \\ -\alpha_- \cdot e^{\frac{(t_{\text{post}} - t_{\text{pre}})}{\tau_-}}, & \text{otherwise.} \end{cases} \tag{7.6}$$

where Δw is the weight change, α_+ and α_- are the magnitudes of potentiation and depression, and τ_+ and τ_- are the decay constants of potentiation and depression. For pre-synaptic neuron, t_{pre} is offset by the synaptic transmission delay. The above STDP rule is an unbounded rule, in that the synaptic weight can increase and decrease unbounded. Hard (explicit) and soft (implicit) rules can be introduced to limit the synaptic weight values. An implicit rule may achieve this by enforcing that the amount of potentiation or depression depends on the current value of the synaptic weight.

7.1.5 Numerical methods for solving SNN simulations

The pulsatile form of spiking activity means that the solving methods for equations of neurons, synapses and plasticity must be stopped upon the occurrence of each spike to update the state variables such as membrane voltage and synaptic conductances. Determination of analytical solutions are further complicated by the common use of temporally stochastic inputs. Furthermore, the dynamics of LIF neuron models can be solved analytically but only approximate analytical solutions have been proposed LIF neurons combined with conductance synapses [243].

As an alternative to analytical solutions, numerical integration methods are used to evolve the state variables incrementally by small steps, known as the timestep. Smaller timesteps lead to more accurate estimations of the state values but lead to longer simulation times. The Forward Euler method uses the gradient of variable to evolve its value. Due to its efficiency and the regular interuptions due to spikes, the Forward Euler method is the most commonly used method for evolving SNN state variables. Particularly, the Forward Euler solver assumes that the value of the function $f(t)$ at some small time δt in the future is approximated by:

$$f(t + \delta t) = f(t) + \delta t \cdot \frac{df(t)}{dt} \tag{7.7}$$

7.1.6 Other simulators

There are a variety of available simulators for SNN research. The best choice of simulator depends on the project and available computing resources. Different simulators vary in terms of the scale and level of detail they are designed to simulate. Different simulators are also designed to run on different computing architectures. Different simulators also offer different levels of flexibility and ease of use.

The most established simulator projects include NEURON [117], Brian/Brian2 [102], and NEST [72]. These benefit from a large support base and development history, as well as a large variety of models and features. More recently, Auryn [303] and ANNarchy [285] have emerged, with Auryn being demonstrated as the fastest SNN simulator for point neuron models [303, 285] using multi-core and multi-node parallelism.

Following the increased availability of General Purpose Graphical Processing Units (GPGPUs, referred to as GPUs herein), a number of GPU based simulators have been developed, due to the well-suitedness of GPUs to point neuron SNN simulations (see Section X). NeMo [81], GeNN [301] and CarlSIM [19] are examples of GPU based simulators. Until recently [8], GPU simulators have not been directly compared to CPU based simulators.

7.1.7 GPU programming

Modern GPU devices allow large scale parallelism as they contain large numbers of cores (on the order of 1000s) [187]. GPU devices use a single instruction multiple data (SIMD) paradigm, in which large numbers of cores execute the same instructions simultaneously but independently on different data. For example, a small piece of code can be exectuted for indices i equal 0 to 100. The code gets executed for each index i on a separate core, with the index i being used to access and operate on different pieces of data.

'Spike' was originally designed to run on NVIDIA GPU devices, although the code was later adapted by a member of the research group to allow processing to be performed using other languages. The language for development on NVIDIA GPUs is called CUDA and is an extension of C++. In CUDA, the pieces of code that are executed in parallel across multiple cores are referred to as kernels [202]. Kernels are launched from the CPU-side using C++. The 'kernel launch' refers to the temporal overhead between the time the kernel is called and the time that it is executed on the GPU. An overview of GPU computing is given by [36]. A number of optimisations

which have previously been implemented in the context of other SNN simulators will be be revisited in the description of the optimisations made for Spike.

7.2 Methods: Spike

As described Spike is a multi-purpose point neuron simulator written in C++ and CUDA. Spike is designed to run on single GPUs and is optimised for STDP and to be delay insensitive. Spike contains a flexible hierarchical class structure for neuron, synapse, and plasticity models.

Spike was in development from 2015-2018 by the thesis author and others in the laboratory, particularly Nasir Ahmad and Toby St Claire Smithe. Nasir Ahmad and the thesis author co-created the project which is available at `www.github.com/oftnai/spike`. Nasir Ahmad started the project by developing a CUDA implementation of an SNN model. The thesis author then converted it into a structured code base which was then further developed by Nasir Ahamad and the thesis author. Toby St Claire Smithe later adjusted the code to allow processing to be performed using languages other than CUDA.

7.2.1 Optimisations

The following subsections describe the optimisations made to Spike.

7.2.1.1 Synaptic Matrices versus Lists

As the connections in modelling studies are often highly sparse, it is efficient (in terms of memory and data storage) to use sparse representations to store synaptic information. Such sparse representations also allow efficient data operations to be made [34] and for multiple synapses between pairs of neurons to be made, as has been observed experimentally [145]. Sparse representations also allow each connection between a single pair of neurons to have a different transmission delay.

Spike uses a list based sparse representation. Particularly, Spike stores lists of presynaptic and postsynaptic neuron IDs, synaptic weights, synaptic delays. In these lists, the value for a synapse s_i, can be indexed using the index i. The lists are ordered by presynaptic neuron index. A look up list identifies the range of postsynaptic synaptic indices for each presynaptic neuron. Storage of efferent synapse information in neighbouring memory locations, allows efficient 'coalesced' memory updates of efferent synapse information upon a spike. It is important to note that early GPU based simulators organised synapses similarly [81].

7.2.1.2 Timestep Grouping

A similar optimisation to that described here has been used in the context of multi-node SNN simulation under the name 'index buffering' [192, 72], as a way of minimising inter node communication. Index buffering takes advantage of the the fact that for any time period of the simulation which is smaller than the minimum axonal delay t_{min}, only spikes which were emitted before that time period can affect neuron dynamics during that time period.

Spike demonstrated the first use of index buffering for GPU SNN simulation, as a way of minimising the number of kernel launches each of which carry a temporal cost. A new name 'timestep grouping' was used to describe the concept in the context of GPU SNN simulation. Particularly, each kernel launch iterates over a number of timesteps equal to the minimum transmission delay divided by the simulation timestep. Spiking activity within this period only affects other neurons after the minimum axonal delay has elapsed, and can used in the subsequent timestep grouping.

7.2.1.3 Neuronal-Synaptic Parallelism

As the number of synapses in SNN models is often large compared to the number of neurons, processing of synaptic dynamics can be expensive. The dynamics of a synapse must be updated when a spike arrives at the synapse. Checking whether a spike has arrived at each synapse on each timestep is inefficient. Spike introduces Neuronal-Synaptic (N-S) parallelism as an alternative. Upon each presynaptic neuron spike, N-S records the index range of the efferent synapses of the presynaptic neuron. These 'active synapses' are then processed.

7.2.1.4 Delay Insensitive

To achieve insensitivity to heterogeneous transmission delays, each neuron is assigned two input buffers (one for each synapse type, e.g. excitatory and inhibitory). The neurons iterate over these buffers to collect incoming synaptic updates. Upon each timestep, a single neuron shifts forward on its buffers by one location and reads the synaptic inputs for its current timestep. The buffer is treated as a ring to conserve memory. The length of the buffer is equal to the maximum possible synaptic delay. Synaptic inputs are placed on the buffer at a distance (corresponding to the delay of that synapse) from where the neuron is currently querying for inputs. This distance corresponds to the synaptic delay in timesteps and thereby ensures that the neuron only reaches this input after the corresponding delay. The circular nature of the buffer

ensures that memory is conserved. Simulation speed is unaffected by transmission delays but memory usage is increased. The technique was first described by [192], and has been used in several other simulators.

7.2.2 Comparison to other simulators

Benchmark comparisons to other actively used simulators were performed by another member the lab (Nasir Ahmad). The code for the performing the benchmarks comparisons for different SNN simulators is publicly available at the URL: https://github.com/nasiryahm/SNNSimulatorComparison. The details of the comparisons are available in the Spike preprint [8] with the most recent benchmark tests available at the URL.

Spike was compared to four CPU based simulators:

- ANNarchy [285].
- Auryn [303].
- Brian2 [102].
- NEST [163].

and the GPU based simulator:

- GeNN [301]

The latest comparisons were made in March 2019 for the master branches of the following simulators (git commit ID also given):

- ANNarchy @ d27f8a4
- Auryn @ 537f293
- Brian2 @ c429811
- GeNN @ 4419d8b
- NEST @ a16b67f
- Spike @ 68c5767

7.2.2.1 Comparison hardware

Single thread benchmarks were produced using a system with the following specifications:

- CPU: Intel i7-4770K
- GPU: NVIDIA GTX 1070 founders edition

Multi thread benchmarks were produced using a system with the following specifications:

- CPU: 2x Intel(R) Xeon(R) CPU E5-2623 v4 @ 2.60GHz

7.2.2.2 Benchmark models

The two benchmark models have previously been published and are referred to as the Vogels-Abbott and Brunel benchmarks [285, 303].

The Vogels-Abbott benchmark is based upon a reduced scale version of the network presented in [286] and consists of:

- 3200 excitatory and 800 inhibitory Leaky Integrate and Fire (LIF) neurons.
- 2% random synaptic connectivity between all populations of neurons.
- Conductance-based synapses.
- Background input maintains a firing rate close to 17 $spikes/second$.
- Network parameters equivalent to equivalent to Zenke at al. [303] and Vitay et al. [285], and are given in the Spike preprint [8].

The second benchmark, referred to as the 'Brunel benchmark', is based upon a network adapted by [303] from [37]. Simulators are compared with the Brunel benchmark both with and without plasticity. The network consists of:

- 8,000 excitatory and 2,000 inhibitory LIF neurons
- 10% random connectivity between them with voltage-based synaptic connections.
- 10,000 excitatory input neurons which project to the LIF with 10% connection probability and spike randomly at time sampled from a Poisson process with an average firing rate of 20Hz.

- Plasticity benchmarks are performed using a weight-dependent STDP rule operating on synapses between the excitatory LIF Neurons. The plasticity rule is expected to lead to a normal weight distribution with a mean of 0.1mV.
- All details of the network parameters are given in the Spike preprint [8].

7.3 Results: Spike

7.3.1 Vogels-Abbott Benchmark

The effect of different synapse update schemes on the performance of Spike are given in the Spike preprint [8]. The combination of timestep-grouping (TG) and Neuron-Synapse Parallelism (N-S) gave the fastest simulation speed and was therefore used for the comparisons to other simulators.

7.3.1.1 Comparing Simulators (Single Threaded)

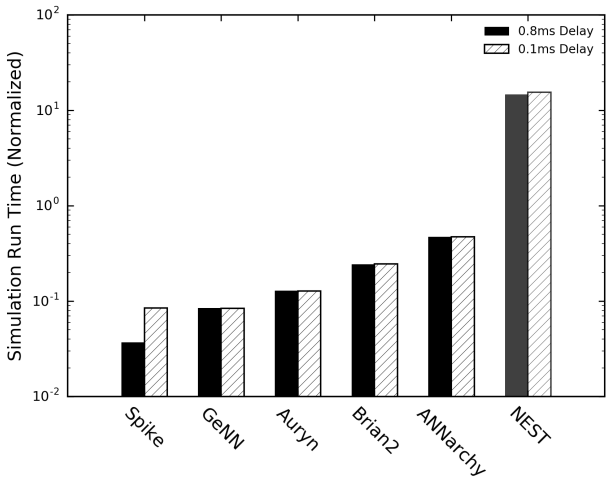


Figure 7.1: The Vogels-Abbott benchmark comparison across simulators. All single threaded or single GPU.
The solid bars show the network simulated with a 0.8ms delay as described in the original benchmark studies [303]. Hatched bars shows the performance of these simulators when the axonal delay is equal to the timestep of the simulation (0.1ms).

Figure 7.1 compares the normalized simulation time required to simulate the Vogels-Abbott benchmark for a set of CPU and GPU based SNN simulators. A delay of 0.8ms was used, as is standard for SNN comparisons [303]. The two GPU based simulators, GeNN and Spike, show a significant speedup compared to competing simulators, with Spike offering the fastest simulaton time. Spike was also the

only simulator which significantly benefits in simulation speed from the introduction of delays in network structure. The fastest CPU simulator is Auryn which has been reported as the fastest CPU based simulator since its release [303, 285]. NEST has incorporated an optimisation based upon delays, but only produces benefits for cluster-based simulations.

7.3.1.2 Comparing Simulators (Multi Threaded)

Auryn is the leading CPU based simulator both in the single and multithreaded use case (based on results presented by [303] and [285]). Spike is therefore compared to Auryn for the multithreaded case for Vogels-Abbott benchmark.

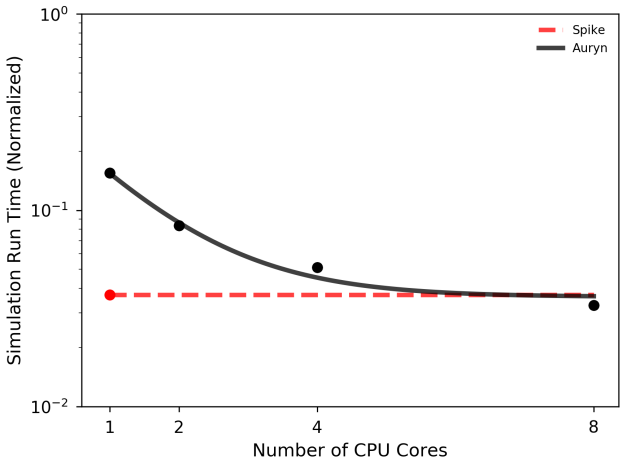


Figure 7.2: Comparing the multithreaded CPU performance of Auryn against Spike on a single GPU.
This plot shows Spike in a single GPU mode only (therefore the dotted line is not a prediction of it's performance over more CPU cores but instead just a baseline). Auryn was benchmarked with 1, 2, 4 and 8 CPU cores. An exponential curve was then fitted to these points to produce the plot shown.

Figure 7.2 compares the speed of Auryn in multithreaded execution mode to the Spike simulator on a single GPU. As can be observed, only under multithreading with eight threads does the Auryn simulator exceed the speed of Spike on a single GPU. Though Auryn does exceed the speed of Spike with eight threads for this network size, it is shown below how the speed of Auryn suffers for larger network sizes. Given that a single system can support multiple GPUs, it is also important to note that with multiple GPUs, the Spike simulator could run multiple models independently (effectively multiplying its simulation throughput), though this is not considered in any benchmarks or plots.

7.3.1.3 The Scaling Behaviour of the Spike Simulator

This section studies how performance of the Spike and Auryn simulators scale with network size. Figure 7.3 shows the scaling of simulation time as network size is increased. As can be seen in Figure 7.3, the Auryn simulator exhibits linear scaling for small network sizes but super-linear scaling at larger network size. The Spike simulator shows sub-linear scaling as the network size is increased and it is assumed that this would continue until GPU memory is exceeded.

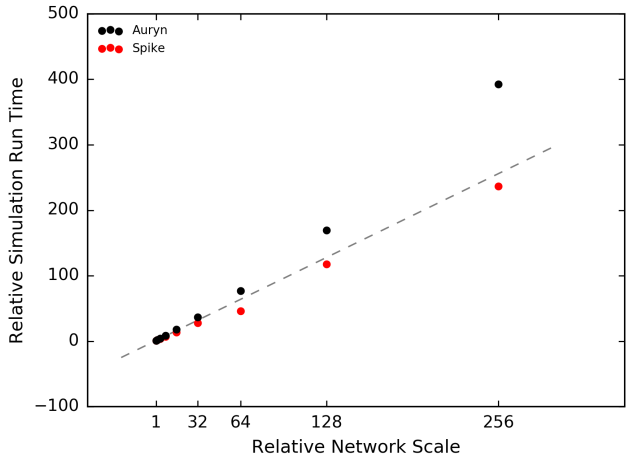


Figure 7.3: Comparing the simulation times of the Spike and Auryn simulators as the Vogels-Abbott Benchmark's network size is scaled. Network size is given by multiplication of the number of neurons and the number of synapses. The x-axis presents the network scale relative to the original benchmark network. On this scale, a value of 1 indicates a network of 4000 neurons, and a value of 100 indicates a network of 400,000 neurons. The smallest network is the size of the previous benchmark with 4000 neurons and 320,000 synaptic connections. The largest network, is 256 time larger with 1,024,000 neurons with 81,920,000 synaptic connections. The y-axis presents the time taken to simulate these networks relative to the time taken by the simulators to simulate a network of scale 1 (a network of 4000 neurons). The gray dashed line indicates linear scaling.

7.3.2 Brunel Plasticity Benchmark

The second SNN benchmark tests the efficiency of the simulators when computing weight updates through synaptic plasticity. Figure 7.4 compares simulation times for the different simulators with and without plasticity. The order of the simulators in terms of speed is the same as for the Vogels-Abbott benchmark but the gap between CPU and GPU simulators is larger. Spike and GeNN are much faster compared to the

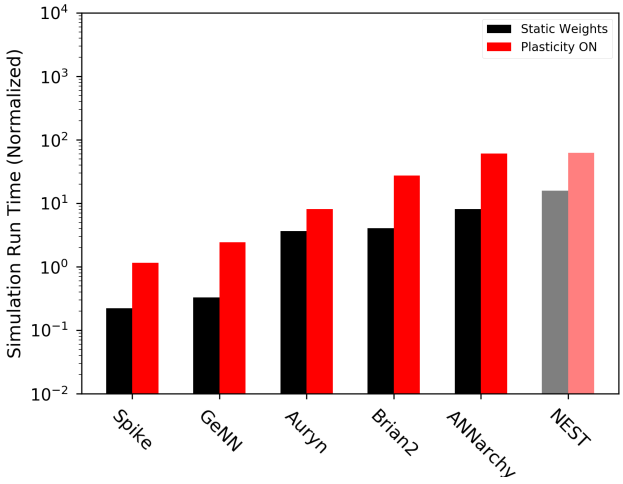


Figure 7.4: Brunel benchmark comparison across simulators with (black) and without (red) plasticity.
All single threaded or single GPU.

CPU simulators in both the plastic and non-plastic. There is an order of magnitude difference between the fastest CPU simulator and Spike. Compared to Auryn, Spike is approximately seven times faster in the plastic case and approximately 15 times faster in the non-plastic case. GeNN simulator remains on the order of two times slower than Spike for both the plastic and non-plastic cases.

7.3.2.1 Multithreaded CPU Performance

As for the Vogels-Abbott benchmark, the performance of the Spike simulator on a single GPU was compared against a multithreaded execution of the Auryn simulator for the non-plastic case of the Brunel benchmark (Figure 7.5). The speed of the Auryn simulator converges towards a longer simulation time for higher numbers of CPU cores.

7.3.3 Summary

This concludes the description of the Spike simulator and its optimisations. Benchmark comparisons to other simulators (made by another member of the research group) were also briefly summarised to show the strengths of the Spike simulator (which the author of the thesis co-created and designed). The Spike simulator was shown to be the fastest in its class for large point neuron simulations with heterogeneous distributions of axonal conduction delays and STDP. This enabled the plastic-

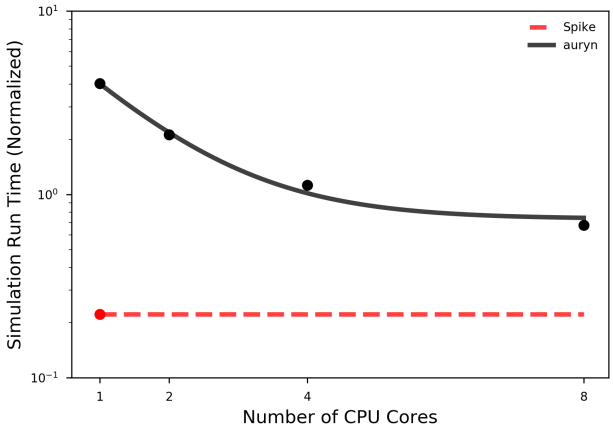


Figure 7.5: Comparing multithreaded CPU performance with Auryn against Spike with the Brunel non-plastic benchmark.
This plot shows Spike in a single GPU mode only (therefore the dotted line is not a prediction of it's performance over more CPU cores but instead just a baseline). Auryn was benchmarked with 1, 2, 4 and 8 CPU cores. An exponential decay curve was then fitted to these points to produce the plot above.

ity simulations in the following section, which in turn lead to the work of all of the preceding chapters.

7.4 A model for exploring the emergence of spike timing continua

The unpublished simulations presented in this chapter explore how leaky integrate-and-fire (LIF) neurons learn to respond to simple one-dimensional stimuli over a range of translations through STDP. The early investigations related to this work inspired the hypotheses of Chapter 6 and the decision to analyse in vivo spiking activity, as presented in Chapters 3 & 4, and published in Scientific Reports [136].

The simulation work presented in this chapter aimed to provide tools and a simplified model, for exploring the hypotheses of Chapter 6, with the hope of providing a proof of principle of the hypotheses in a single model. The chapter presents the first steps towards this. Particularly, the chapter presents simulations and analyses exploring how single spiking neurons learn (through STDP) to respond to a spike time encoded stimulus over a range of translations. Moreover, the analyses study how the transformation (i.e. location) of the stimulus becomes encoded by the first spike time of the neuron through learning.

The chapter takes particular interest in how the spike time of the neuron varies continuously with continuous changes of stimulus location. It was hoped that the chapter would explore how the self-organisation of such spike time continua could facilitate robustness to false combinations of composing features. As described in the Introduction to this chapter, technical and theoretical limitations limited the progression of the project. The following sections presents the work which was completed in the fist years of the DPhil.

7.4.1 Biological Evidence of STDP

Spike-time-dependent plasticity (STDP) describes the finding that changes in synaptic efficacy depend on the relative spike times of the pre and postsynaptic neuron. [176] first showed that the order of pre and postsynaptic spikes affects the sign of synaptic change, and [21] further characterised the dependence. STDP usually requires 60-100 spike pairs. STDP tuning curve shapes vary between synapses and brain areas. STDP also depends on firing rate and postsynaptic voltage [176, 261], the baseline synaptic weight [21] and neuromodulation [250]. Other experiments show non-linear effects for spike triplets [87]. See [77] for review.

7.4.2 Models of STDP

Biophysical models can be computationally expensive with many free parameters. Phenomenological models are useful for analytic studies and larger simulations (see [191] for review), with potential real world applications. Some recent models try to bridge the gap between phenomenological and biophysical models [67] and some models attempt to model non-linear multi-spike learning rules.

Phenomenological models often have both a pre and postsynaptic activity trace, which are updated by pre and postsynaptic spikes respectively. Upon a postsynaptic spike, the synapse is typically potentiated by the presynaptic trace multiplied by a function $F_+(w)$ of the synaptic weight. Upon presynaptic spikes, the synapse is typically depressed by the postsynaptic neuron trace multiplied by a function $F_-(w)$.

A rule is 'additive' if $F_+(w)$ and $F_-(w)$ are equal to constants. For additive rules and Poisson input trains, synaptic efficacies form a bimodal distribution. A rule is 'multiplicative' if $F_+(w)$ and $F_-(w)$ are proportional to the weight and instead leads to a unimodal distribution of the weights, which is more realistic. Thirdly, a rule is a 'power law update' if $F_+(w)$ and $F_-(w)$ are proportional to a power of the weights

between 0 and 1, creating an interpolation between an additive and multiplicative rule [105].

7.4.3 STDP Associative Invariance Learning

Multiple spiking models have attempted to learn invariant visual features through STDP. These can be divided into models with equal or zero delays [178, 75, 263, 61, 152] and those with ranges of delays that act in a polychronous state [138, 214, 228, 229].

7.5 Methods: the emergence of spike time continua

The Spike simulator described in the first part of this chapter was used for all simulations. This section describes the network, as well as the training and testing paradigm used for the simulations presented in this chapter.

7.5.1 Model architecture and equations

Network architecture

The network is composed of a single 1-dimensional input-layer composed of 20,000 input neurons and a 1-dimensional output layer composed of 50 LIF neurons. The input and output layer were connected in a feedforward manner with excitatory voltage-based synapses that had a connection probability of 30%. The code base was designed to facilitate inhibitory synapses, multiple excitatory and inhibitory neuron groups and different synapse types between neuron groups. Many exploratory and more complex simulations were run in the first years of the DPhil. Due to the complexities of such models, a simplified model was instead chosen, in order to better understand the base mechanism of how continua emerged.

Input neurons + stimuli

As described, the input layer consists of 20,000 neurons. Each input neuron was assigned a mean spike time in the range $[0, 50ms]$. Presentation of a stimulus consisted of activating 5,000 contiguous input neurons. Activation of an input neuron consisted of a single spike at its mean spike time with additional temporal jitter drawn from a normal distribution with standard deviation equal to 1. In total there were 200

stimuli spaced uniformly over the 1D input space such that stimuli shared a degree of overlap with their neighbouring stimuli.

Output neurons

Output neurons are modelled as leaky integrate-and-fire neurons using the following differential equation:

$$\tau_m^{\gamma_l} \frac{dV_i(t)}{dt} = V_0^{\gamma_l} - V_i(t) + I_i(t) \tag{7.8}$$

where $V(t)$ is the voltage of the cell on timestep t, V_0 is the resting potential of the cell and I(t) is the current injected into the cell on timestep t. If the membrane voltage rises above a threshold V_{thresh}, the neuron emits a spike and the membrane voltage is reset to the cells resting potential V_{rest}. A refractory period of 2ms is imposed by fixing the membrane voltage at V_{rest} for a set duration.

Synapses

Synapses are modelled as 'voltage-based' synapses, such that the total input current to the neuron at a single timestep is equal to the sum of the weights of synapses which have spikes arriving (from their associated presynaptic neuron) at the given timestep:

$$I_i(t) = \sum_{j \epsilon S_t} w_{ij}(t) \tag{7.9}$$

All synapses in the model are excitatory, in that they contributed a positive value to the net input current. Starting weights are sampled from a uniform distribution defined by the range $[0, \lambda]$, where λ is a biological scaling constant. Over the course of plasticity, weights for excitatory synapse groups are constrained between the initial range, as will be described. The process by which λ is set for different synapse groups will also be described below.

All synapses have an associated axonal propagation delay, which is the time for a presynaptic spike to 'propagate' from the presynaptic neuron to the synapse. Each synapse group has an associated distribution of axonal propagation delays from which axonal transmission delays are drawn. Axonal transmission delays were drawn from a uniform distribution defined by the range $[1, 15]ms$. As such input spikes arrived at the postsynaptic neuron at a time equal to the input neuron's spike time plus the associated axonal delay.

Plasticity rule

Synapses in both model architectures are updated by a model of STDP with weight decay given by the following equations:

$$\Delta w = \begin{cases} l * \alpha_+ e^{-\delta t/\tau_+} & \text{if } \delta t \geq 0 \\ l * w * \alpha_- e^{\delta t/\tau_-} & \text{otherwise} \end{cases} \quad \text{for every pre/post-synaptic spike pair} \tag{7.10}$$

$$\Delta w = w - \alpha_d, \text{for every pre-synaptic spike} \tag{7.11}$$

where l is the learning rate. α_- and α_+ are the classical constants for scaling potentiation and depression. Whilst the effect of these parameters was explored in earlier simulations, they were set to 1 in the analyses presented here. In the simulations, synaptic depression through STDP (rather than decay) was weight dependent in order to slow the convergence of weights to zero. This was a choice made 3 years ago when the simulations were run.

7.5.2 Test and training paradigm

7.5.2.1 Testing the network

The network was tested at different stages of training. A single network test consisted of presenting each stimulus 20 times. This allowed calculation of each neuron's reliability and mean spike time in response to the stimulus in each location. The network was reset between each stimulus presentation.

7.5.2.2 Training the network

A single training iteration consisted of presenting each stimulus once. The order of stimulus presentation was randomised on each training iteration and the network was reset between each stimulus presentation. A single training phase was defined as 5 training iterations. In total there were 500 training iterations, and thus 100 training phases.

7.5.2.3 Test-train-test order

The network was tested before the start of training, and then after each training phase. The synaptic weight values were also stored before the start of training, and then after each training phase. Together, testing the network and storing the synaptic weight values allowed the progression of the networks state and activity to be analysed over the course of training.

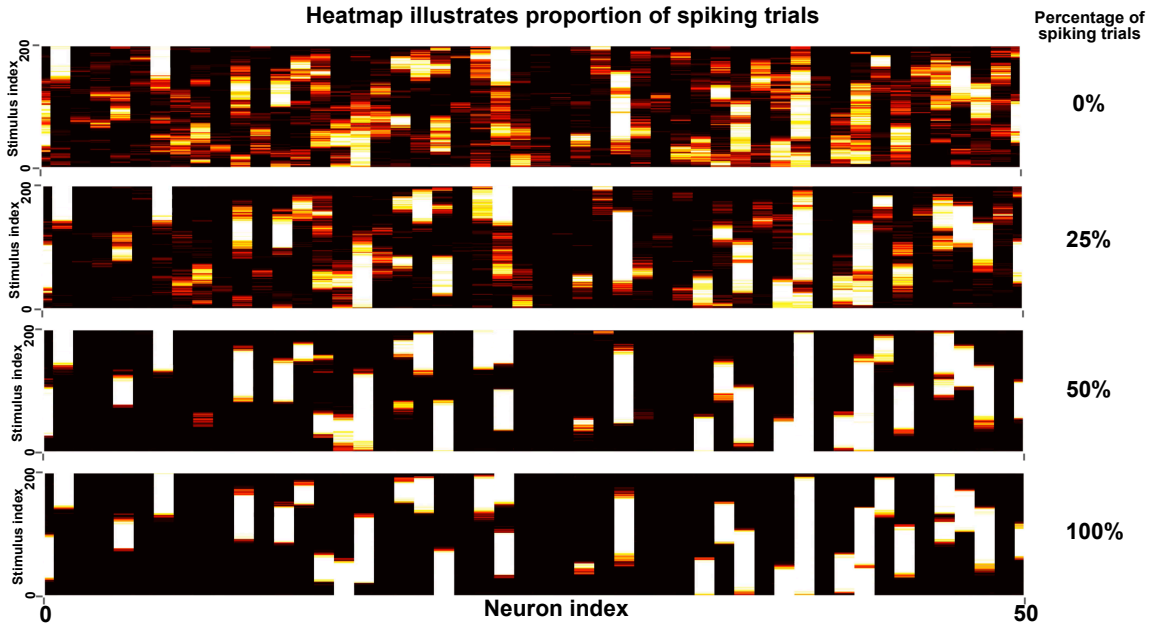


Figure 7.6: Heatmaps showing reliable responses by neuron (x), stimulus (y) and percentage of training complete (rows). White represents that a neuron spikes for all test trials of the stimulus; black otherwise. Parameter combination: learning rate: 0.0001, alpha decay: 0.15.

7.6 Results: the emergence of spike time continua

This section presents the simulation results for the following set of parameters:

- Learning rate l: 0.001, 0.0001
- Weight decay constant α_d: 0.05, 0.15, 0.25

The main purpose of this section is to present the analyses created for the study. The analyses can be divided into analyses for single simulations and analyses studying the effect of parameters over the range of parameters. Explorations were made varying other parameters, such as the input to output connection probability and the number of input neurons in each stimulus, although these are not presented here.

7.6.1 Individual Simulations

7.6.1.1 Number of spiking test trials

This first analysis visualises how the output neurons learn to respond to the the stimulus over a range of transformations. Fig. 7.6 illustrates an example for the parameter combination (learning rate: 0.0001, alpha decay: 0.15) over several training

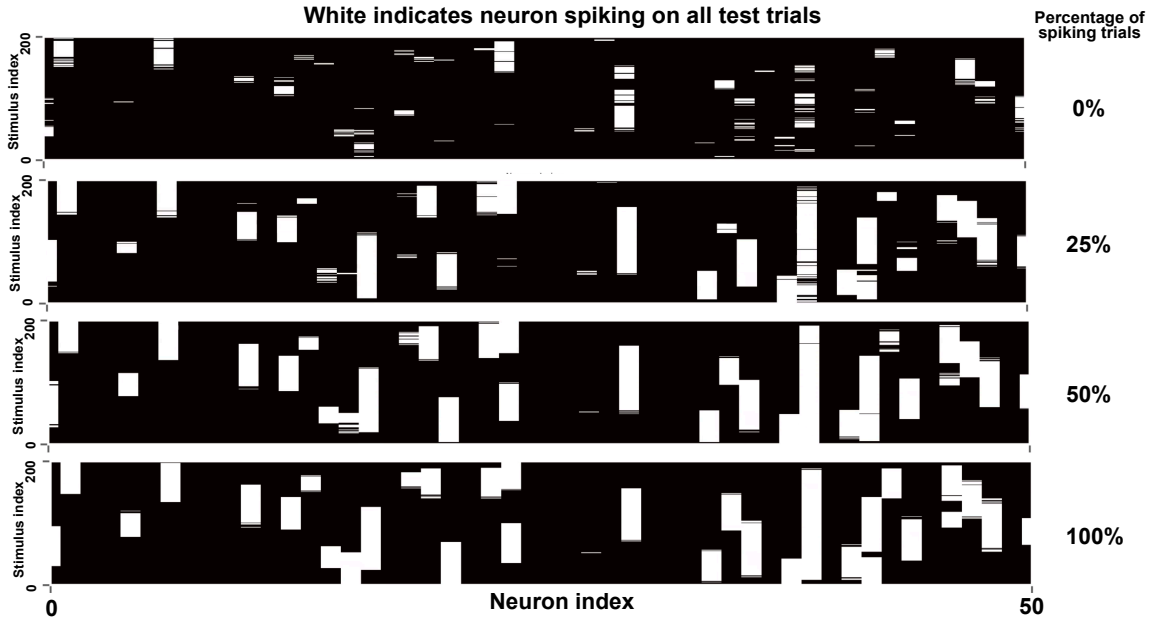


Figure 7.7: Number of spiking test trials (black = 0, white = all) by neuron (x), stimulus (y) and percentage of training complete (rows). Parameter combination: learning rate: 0.0001, alpha decay: 0.15.

phases. The colour illustrates the number of trials on which a neuron spikes to the stimulus in a particular location. A video showing the evolution of this figure over the course of learning was created but is not included in the thesis.

It can be seen that the output neurons learn to respond to the stimulus in more locations and on more trials over the course of training. Moreover the output neurons learn to respond on more trials over larger ranges of continuous locations. This demonstrates the emergence of translation (or transformation) invariance through learning, as the neurons learn to respond to the stimulus over a range of continuous transformations but do not respond to the stimulus over continuous ranges of translations.

7.6.1.2 Fully reliable neurons

This analysis is similar to the previous analysis but instead plots a white bar when a neuron responds reliably to the stimulus in a particular location on all trials. In the example, it can be seen that neurons learn to respond reliably (on all trials) to the stimulus over a range of continuous locations (Fig. 7.7. As with Fig 7.6, this illustrates the emergence of transformation invariance over a range of transformations. The example is again for the parameter combination: learning rate: 0.0001, alpha

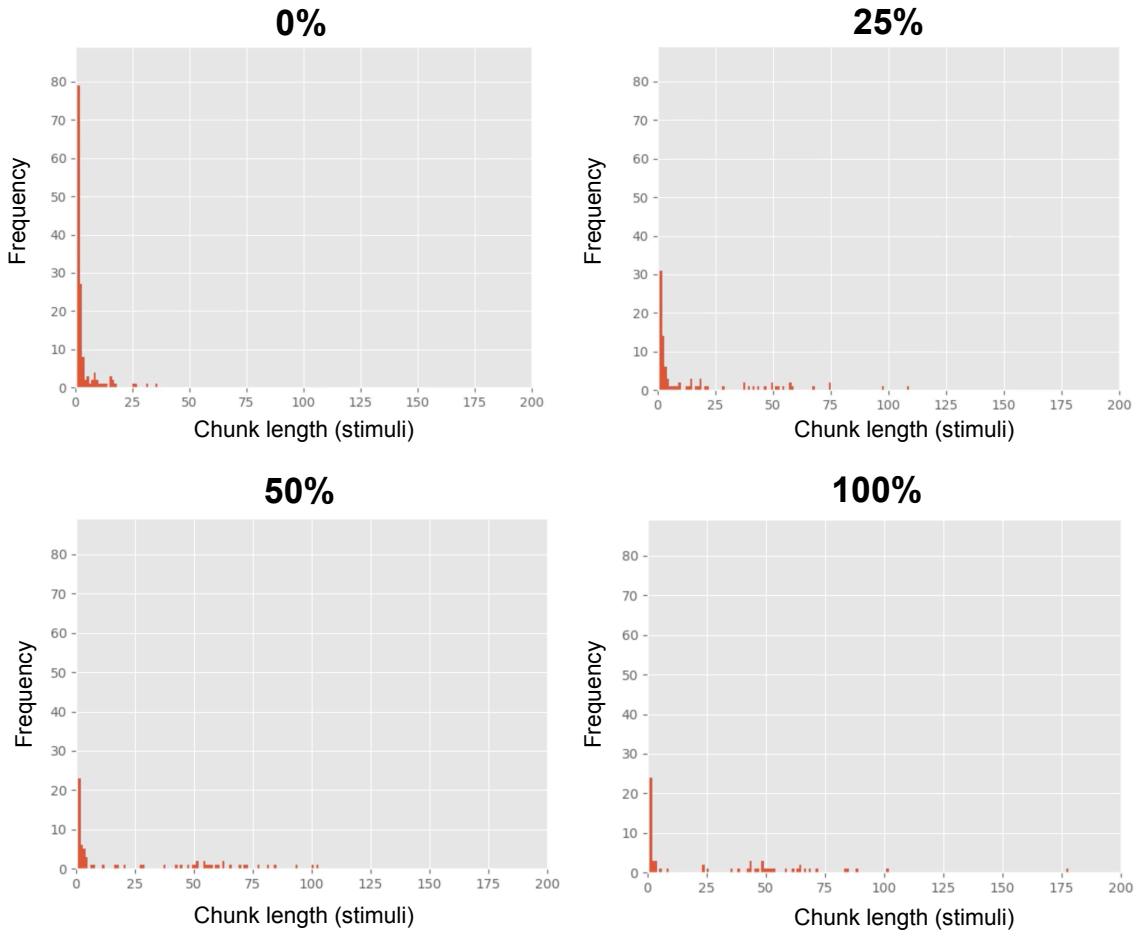


Figure 7.8: Reliable responses over continuous ranges of stimulus transformations. Histograms showing how chunk lengths change as a function of training time. A chunk is defined as a range of continuous stimulus transforms over which a neuron responds reliably (i.e. on all stimulus trials). Parameter combination: learning rate: 0.0001, alpha decay: 0.15.

decay: 0.15.

7.6.1.3 Reliable responses over continuous ranges of stimulus transformations

This analysis visualises how the range of continuous stimulus transforms over which output neurons learn to respond to changes over the course learning. A chunk is defined as a range of continuous stimulus transforms over which a neuron responds reliably (i.e. on all stimulus trials). Fig. 7.8 visualises how the length of these chunks changes over the course of the simulation. It can be seen that the number of larger chunks increases over training.

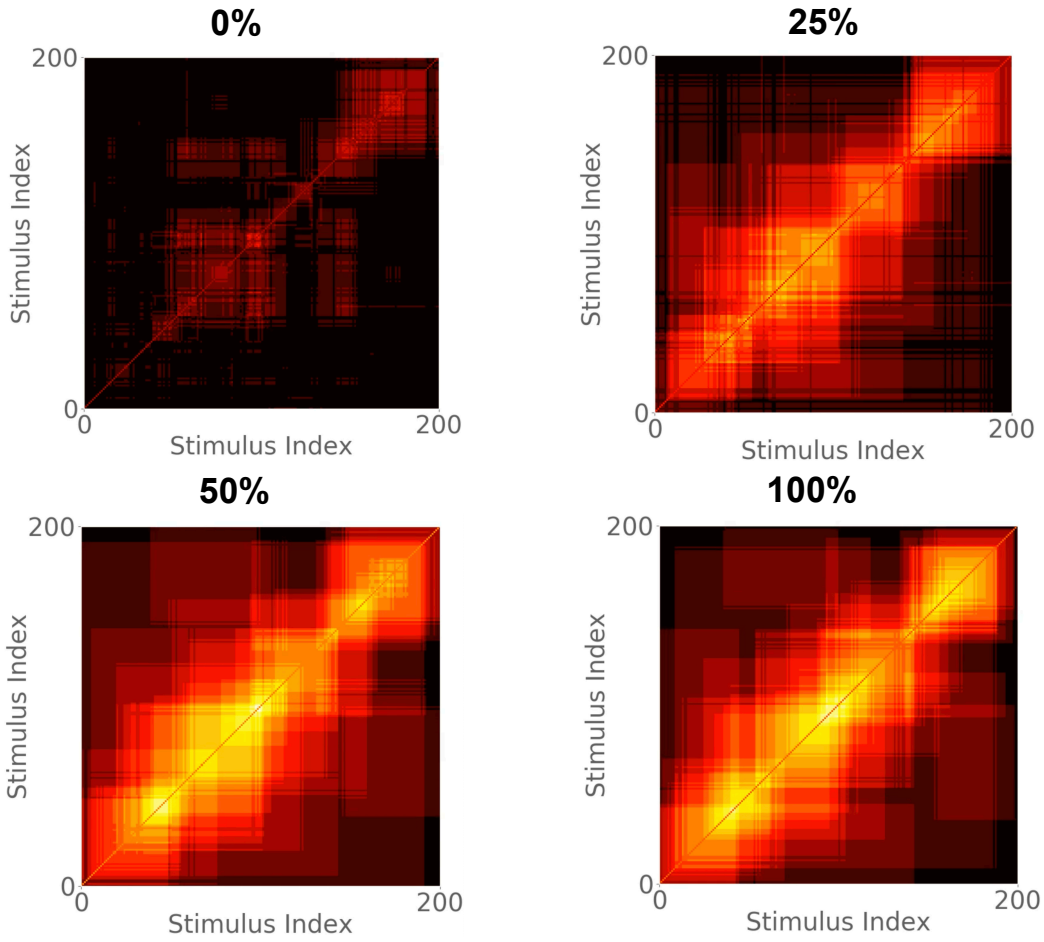


Figure 7.9: The effect of stimulus location on the number of neurons responding to both stimuli in a pair of stimuli. Each entry shows the number of neurons responding to both of the stimuli. Black illustrates zero neurons responding to both stimuli, white illustrates that all neurons respond to the pair of stimuli. The figure shows the evolution of the matrix over the course of training. Parameter combination: learning rate: 0.0001, alpha decay: 0.15.

7.6.1.4 The effect of location on the number of neurons responding to two stimuli

For each pair of stimuli, Fig. 7.9 plots the number of neurons which respond reliably (i.e. on all stimulus trials) to both stimuli. Closer (in translation) stimuli have a higher number of neurons which spike reliably to both stimuli. This is a further demonstration of the emergence of transformation invariance.

The analysis was particularly designed for experiments in which multiple objects were presented. It was hoped that under this case the matrix would show higher similarity for the stimuli which were transformation of the same stimulus.

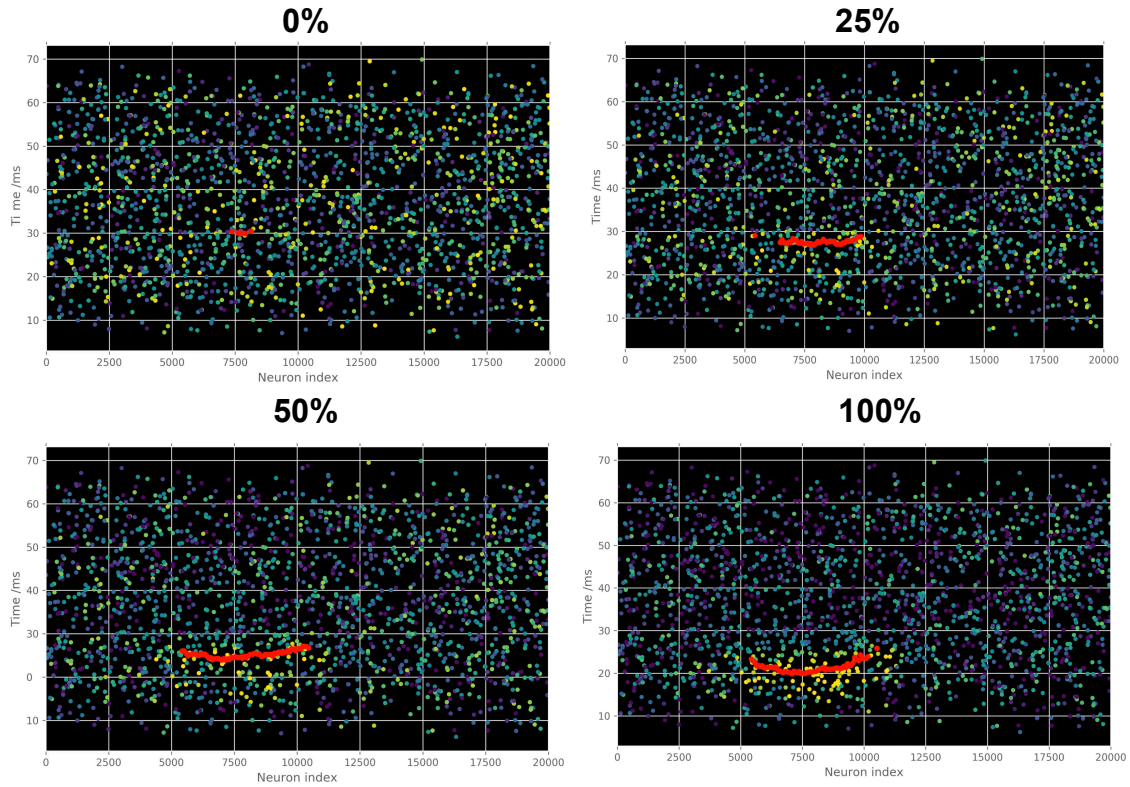


Figure 7.10: Illustration of how a single neuron learns to respond reliably to the stimulus over a range of transformations. The figure shows how the values of the synaptic weights change over the course of training and the mean spike times that the neuron respond reliably to the stimulus in each location. Firstly the 1D input space is represented by the x-axis (input neurons 0-20,000). For each afferent synapse of the output neuron, there is a circle placed with a y-coordinate equal to the mean (expected) time that the synapse delivers spikes to the postsynaptic neuron. The colour of these circles illustrates the strength of the synapse (navy blue equals 0, yellow equals maximum weight). Red dots are placed when the stimulus responds reliably to a stimulus. The x-location of a red dot represents the central input neuron of the stimulus, and the y-location equals the mean spike time that the neuron responds to the stimulus at. The figure illustrates how the synaptic weights and spike times of the neuron change over the course of training. Parameter combination 'medium forward': learning rate: 0.0001, alpha decay: 0.15.

7.6.1.5 Individual Simulation - Space Time Funk Plots

The next step aimed to examine how the first spike times of neurons depended on the stimulus transformation over the course of training. Figs. 7.10, 7.11, 7.12 & 7.13 illustrate how 4 single neurons learn to respond to the stimulus over a range of transformations over the course of training. Particularly, Figs. 7.10, 7.11, 7.12 & 7.13 show how the values of the synaptic weights change over the course of training, as

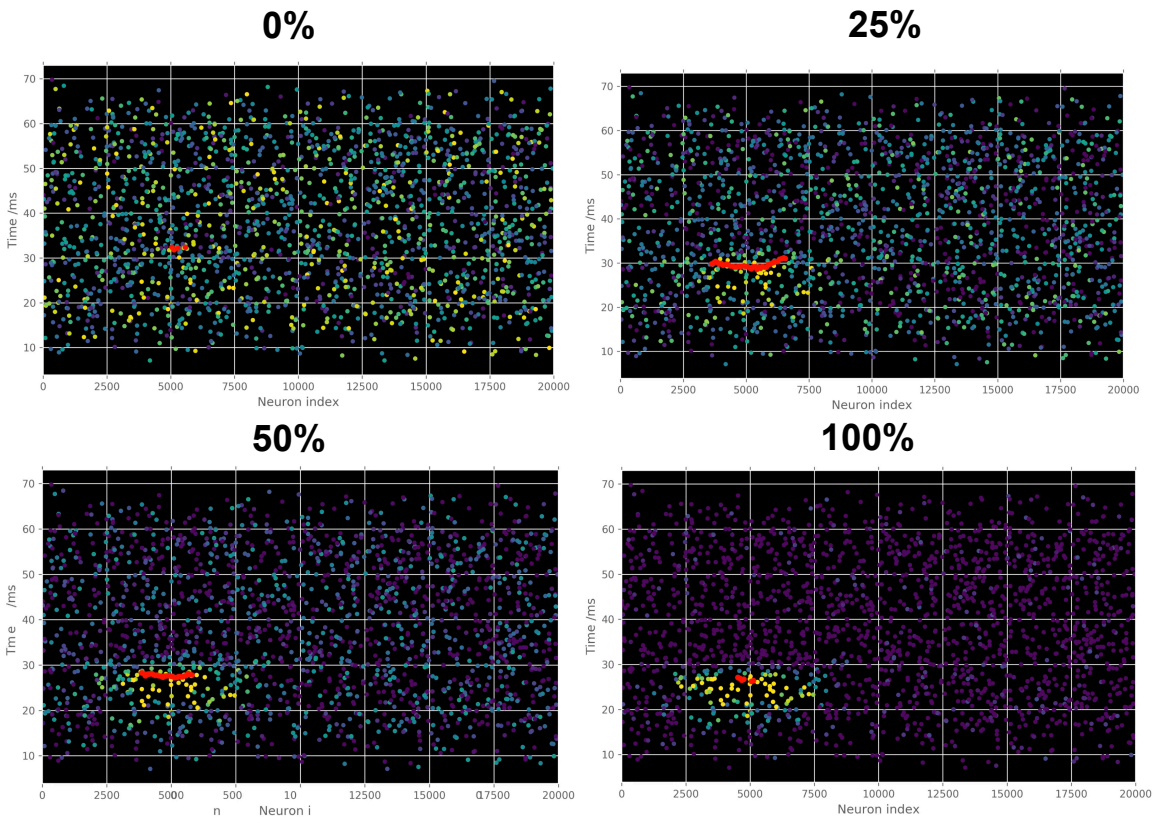


Figure 7.11: Figure the same as 7.10 but for the 'slow fizzle' parameter combination: learning rate: 0.0001, alpha decay: 0.25 and a different example neuron.

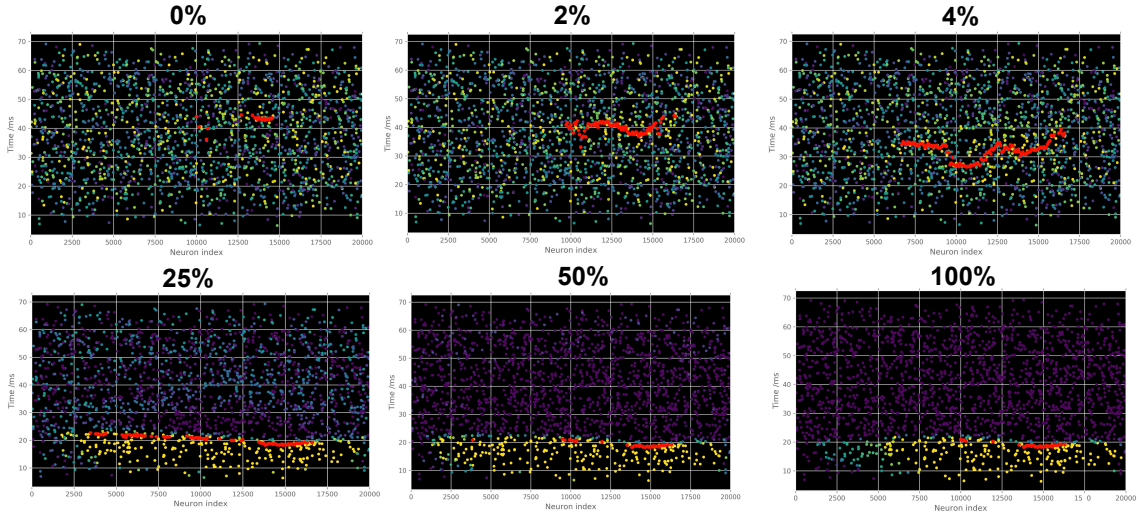


Figure 7.12: Figure the same as 7.10 but for the 'fast forward' parameter combination: learning rate: 0.001, alpha decay: 0.05 and a different example neuron.

well as the mean spike times at which the neuron responds reliably to the stimulus in each location.

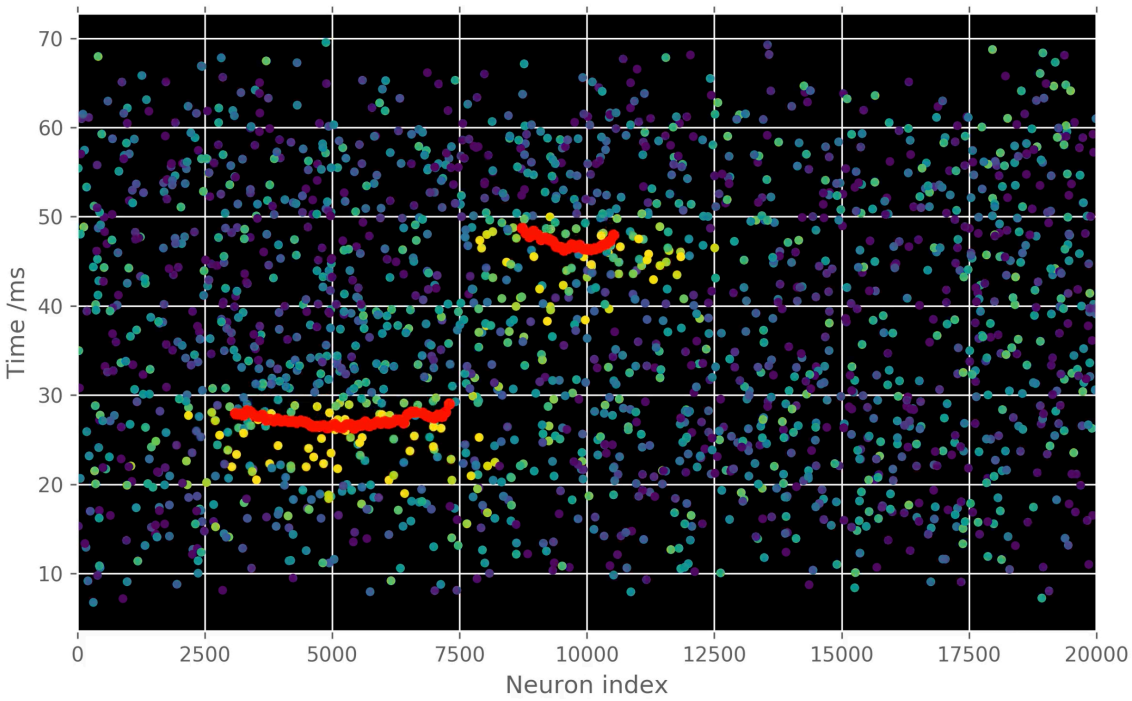


Figure 7.13: Figure the same as 7.10 but for the 'medium forward' parameter combination: learning rate: 0.0001, alpha decay: 0.15 and a different example neuron.

Firstly, the 1D input space is represented by the x-axis (input neurons 0-20,000). For each afferent synapse of the output neuron, there is a circle placed with a y-coordinate equal to the mean (expected) time that the synapse delivers spikes to the postsynaptic neuron. The colour of these circles illustrates the strength of the synapse (navy blue equals 0, yellow equals maximum weight).

Red dots are placed when the stimulus responds reliably to a stimulus. The x-location of a red dot represents the central input neuron of the stimulus, and the y-location equals the mean spike time that the neuron responds to the stimulus at. Figs. 7.10, 7.11, 7.12 & 7.13 thus illustrate how the synaptic weights and spike times of the neuron change over the course of training.

Figs. 7.10, 7.11 & 7.12 illustrate 3 different regimes that were general for the following 3 parameter combinations:

- Medium forward: Learning rate: 0.0001, alpha decay: 0.15 (Fig. 7.10).
- Slow fizzle: Learning rate: 0.0001, alpha decay: 0.25 (Fig. 7.11).
- Fast forward: Learning rate: 0.001, alpha decay: 0.25 (Fig. 7.12).

7.6.1.6 Medium forward

Fig. 7.10 illustrates an example for the 'medium forward' parameter combination in which the neuron gradually learns to respond to the stimulus over a continuous range of transformations. Over time the neuron responds earlier to the stimulus. This is due to the STDP potentiation. As the strength of the synapses which have causally activated the postsynaptic neuron increase in strength, the neuron fires sooner. This leads synapses which deliver spikes earlier to the postsynaptic neuron to also be potentiated. In turn this causes the neuron to spike earlier to the stimulus over the range of transformations. This cycle continues causing the neuron to spike earlier and earlier. This process of 'forward potentiation' can be seen clearly by observing the red mean spike times with respect to the horizontal grid lines over the course of training.

As training progresses, however, weight decay reduces the strength of synapses over the population. This can be seen by the colour of the non-yellow synapses becoming more blue. The process of weight decay slows down the process of potentiation which causes the neuron to spike earlier. It was hoped that the weight decay would stabilise the process of potentiation which causes the neuron to spike earlier.

7.6.1.7 Slow fizzle

Fig. 7.11 illustrates an example for the 'medium forward' parameter combination in which the neuron learns to respond to the stimulus over a continuous range of transformations but the spike times do not move forward as quickly in time over the course of training due to a higher alpha decay constant. Whilst this stabilises the times that the neuron responds, the weight decay eventually causes the neuron to stop firing for some of the stimulus transformations for which it initially learned to respond.

7.6.1.8 Fast forward

Fig. 7.12 illustrates an example for the 'fast forward' parameter combination in which the neuron quickly learns to respond to the stimulus over a range of transformations. The fast potentiation causes the spike times over which the neuron responds to move earlier in time very quickly over the course of training. Eventually, the neuron stops responding over the range of transformations due to weight decay, however.

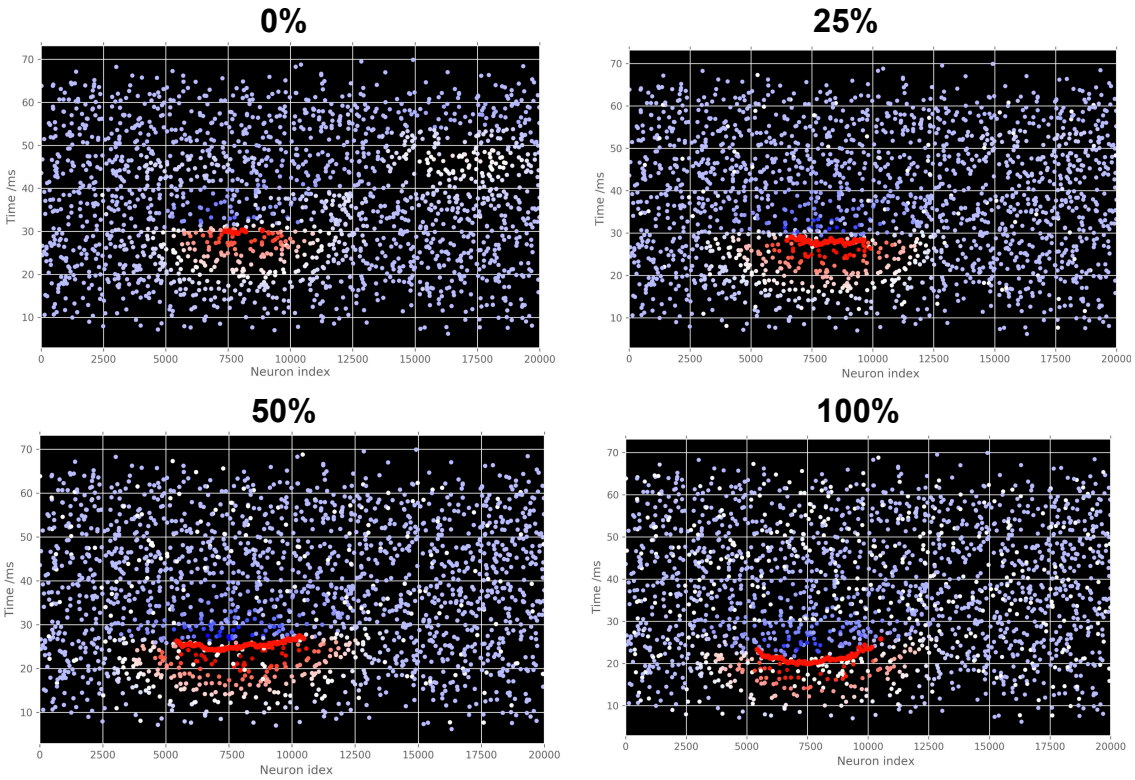


Figure 7.14: Same as 7.10 but instead the colour of the afferent synapses illustrates the weight change over a single training phase (red positive, blue negative).

7.6.1.9 Split response example

Fig. 7.13 illustrates an example for the 'fast forward' parameter combination in which the neuron learns to respond the stimulus over two different continuous ranges of transformations.

7.6.1.10 Summary

To summarise, this section showed that (visually) the emergence of transformation invariance with spike time encoding depends strongly on parameters of plasticity. These were general properties for each of the parameter combinations. These general properties will be quantified below.

7.6.1.11 Individual Simulation - Space Time Funk Plots (Weight Differences)

To better understand the synaptic weight changes and their influence on the spiking activity of the postsynaptic neuron, a similar figure to the previous analysis was made, but with the magnitude of the weight change plotted, rather than the synaptic weight

itself. An example can be seen in Fig. 7.14. The influence of spike timing on the synaptic weight change can clearly be seen, as well as the effect of the weight change on the times that the neuron learns to respond to the stimulus.

7.6.1.12 PSTH

PSTHs were plotted as function of time to better understand how spiking activity changed at the population-level over the course of training. It was important that the PSTH (i) stabilised and (ii) had a broad range of spike times. Both were important to the hypotheses that spike times encoded stimulus information along continua. A suitable PSTH was essential therefore, and so attaining such a PSTH offered a first step requirement.

Fig. 7.15a shows the PSTHs at the end of training for the three example parameter combinations. The 'medium forward' combination provided a broad range of spike times but skewed towards spiking earlier as expected by the previously observed process of forward potentiation. The 'slow fizzle' combination provided a range of spike times but it much fewer neurons spiked at the end of training. The 'fast forward' parameter combination produced only early spike times.

7.6.1.13 Mean spike counts

Fig. 7.15b visualises the proportions of test trials (of a stimulus) that neurons respond to at the end of training. The 'medium forward' parameter combination produces the most fully reliable (i.e. spiking on all test trials) neuron-stimulus combinations at the end of training.

7.6.1.14 Polynomials fit to chunks

For analysis of the properties of the continua, a 3rd order polynomial was fit to each 'chunk'. Particularly, 3rd order polynomials were fit to the mean first spike times over continuous stimulus ranges for which the neuron responded reliably (i.e. on all trials) to all stimuli. Fig. 7.15c shows all the chunk polynomials (for all neurons) at the end of training for the three example parameter combinations.

The 'medium forward' combination provided a polynomials over a broad time window. Most appear to be of 2nd order (i.e. quadratic). This is poor for their ability to encode stimulus location through their spike timing, as a spike time could mean the that the stimulus is on the left or right. The other two example have much smaller chunk polynomials.

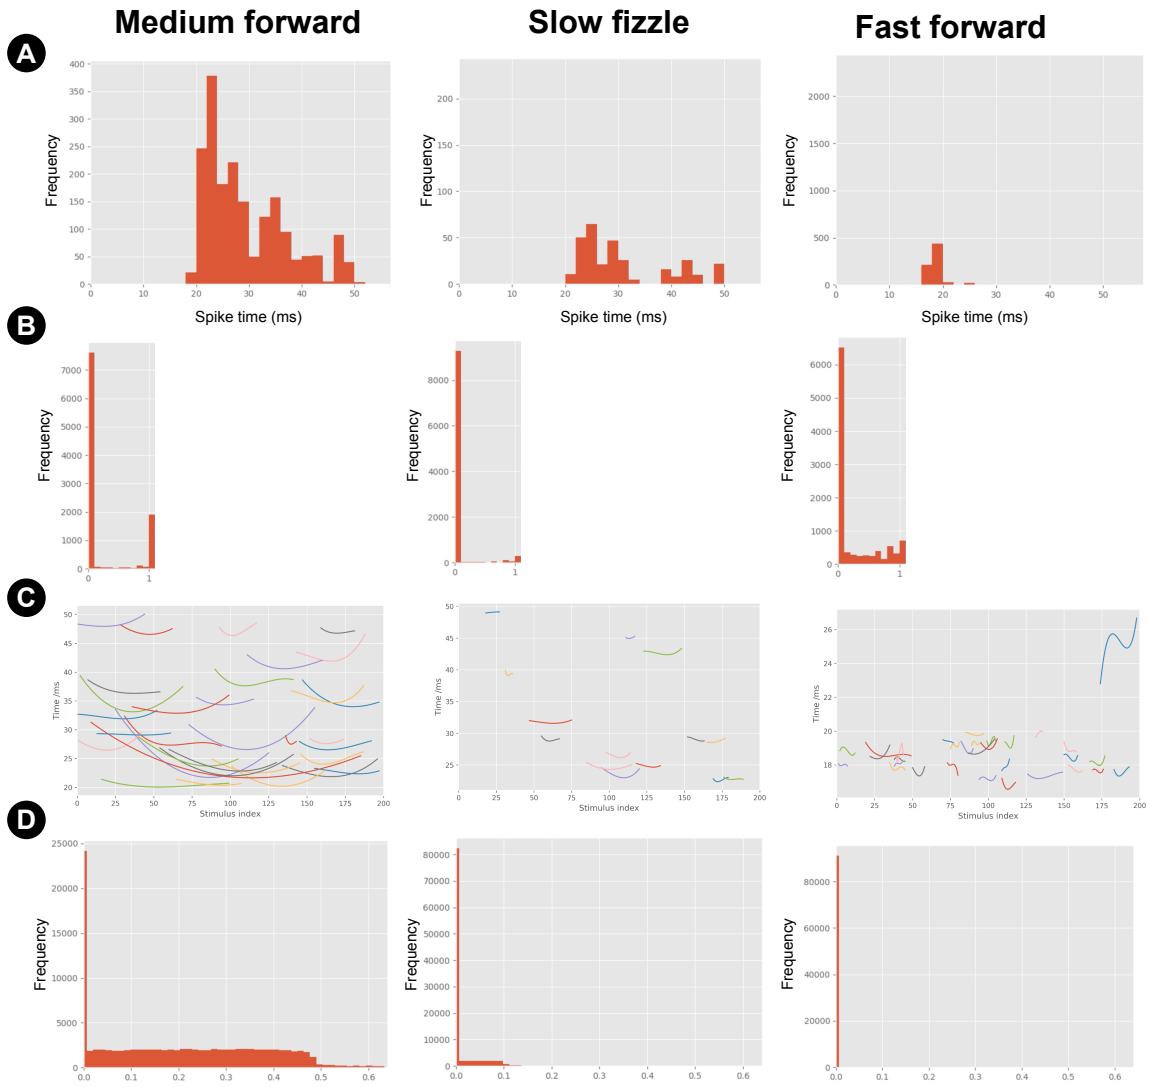


Figure 7.15: (a) PSTHs of three parameter combinations at the end of training. (b) Histogram of mean spike counts at the end of training for three parameter combinations. (c) Polynomials fit to chunks at the end of training for three parameter combinations. (d) Histogram of synaptic weights at the end of training for three parameter combinations.

7.6.1.15 Weight histograms

Finally, to better understand the dynamics of plasticity, a histogram of the weights as a function of training time was plotted. 7.15d shows these histograms at the end of training for the 3 example parameter combinations.

The 'medium forward' combination still has a lot of high valued weight values at the end of training despite the gradual decay of the majority of weights towards zero. The 'slow fizzle' combination has a larger number weights closer to zero at the end

of training. For the 'fast forward' combination, nearly all weights are at zero.

7.6.2 Effect of parameters

The following section plots the effect of the parameters on various metrics.

7.6.2.1 Mean reliable spike times

Fig. 7.16a plots the effect of the parameters on the mean first spike times of reliable (i.e. spiking on all trials) responses change over the course of training. The larger the learning rate and the smaller the learning rate, the faster the mean reliable spike time decreases. The parameter combination with the smallest learning rate and largest alpha decay (green) roughly maintains a constant mean reliable spike time.

7.6.2.2 Mean reliable chunk length

Fig. 7.16b plots the effect of the parameters on the mean reliable chunk length over the course of training. The dynamics of this metric are more complex. For the three parameter combinations with the larger learning rate of 0.001, the metric first increases rapidly and then decreases close to zero. This quantifies the characteristic behaviour of the " fast forward parameter combination, in which chunk length rapidly increases before then decreasing to close to zero. For the three parameter combinations with the smaller learning rate of 0.0001, the metric increases gradually before stabilising for the parameter combinations with alpha decay constant 0.15 and 0.25.

7.6.2.3 Mean synaptic weight

Fig. 7.16c plots the effect of the parameters on the mean synaptic weight over the course of training. The higher the learning rate, the faster the mean synaptic weight decreases. The parameter combination with learning rate: 0.0001 and alpha decay: 0.05 almost has a constant mean weight but is still decreasing.

7.6.2.4 Mean 1st lag chunk autocorrelation r-value

Fig. 7.16d plots the effect of the parameters on the 1st lag autocorrelation r-value over the course of training. This metric aimed to give a quantification of smoothness of chunk spike times over stimulus transforms. Particularly, for a continuous range of stimulus transformations over which the neuron spiked on all trials, the lag-1 autocorrelation r-value was calculated. For each parameter combination at a stage of training, the mean of these values over all chunks and neurons was taken.

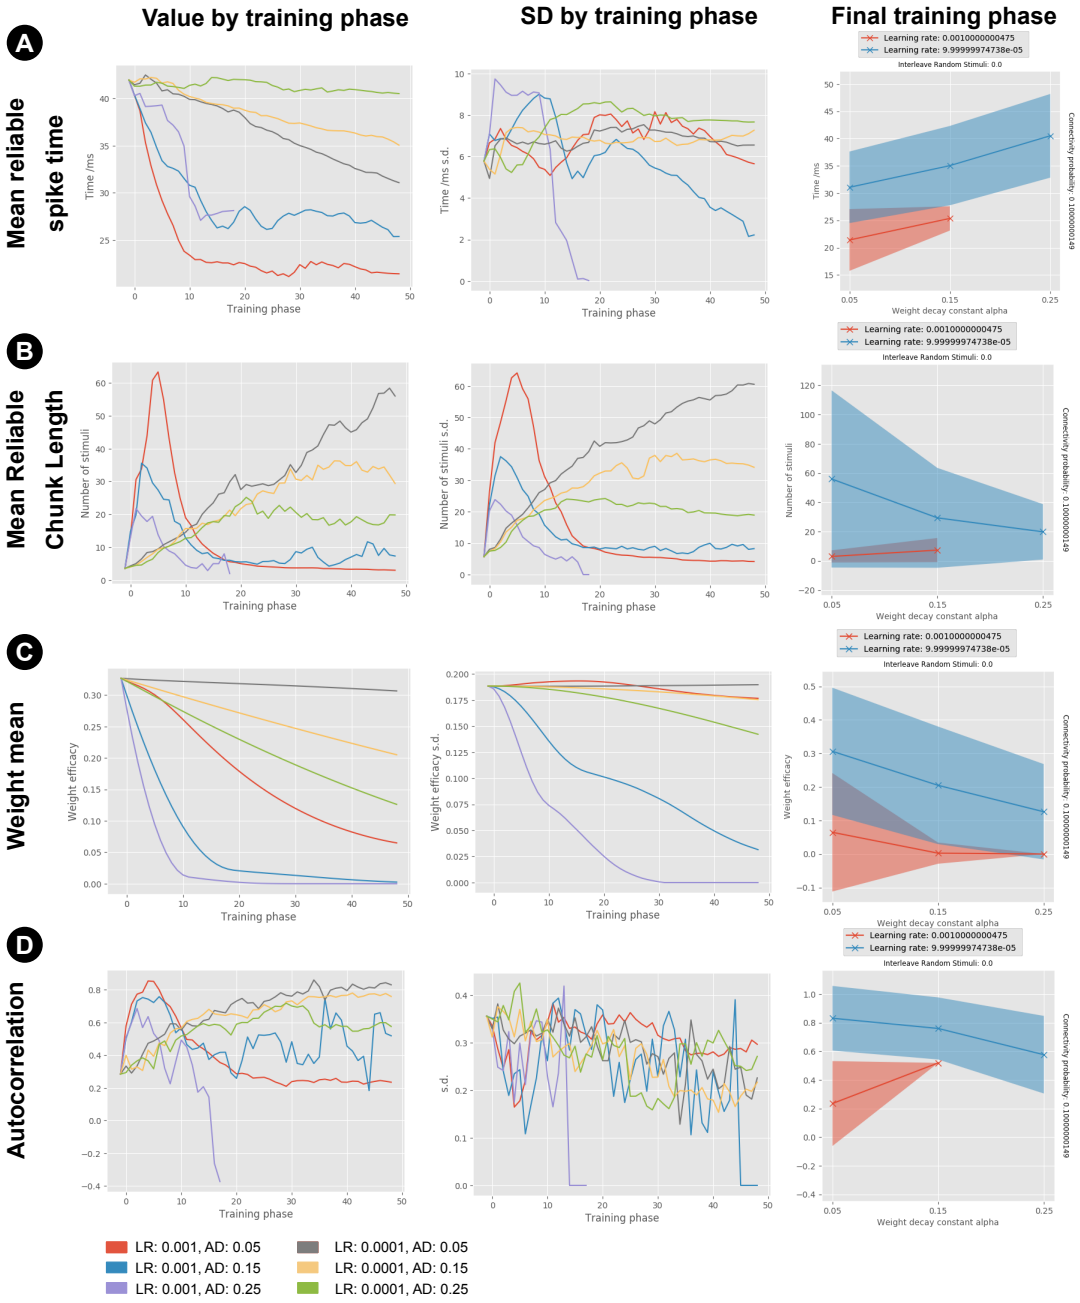


Figure 7.16: Figure shows the effect of the parameters on (a) mean reliable spike times, (b) mean reliable chunk length, (c) mean synaptic weight and (d) 1st lag autocorrelation r-value. The 1st column shows how these values evolve over training. The 2nd column shows how the standard deviation of these values evolve over training. The 3rd column shows the final values with 3*standard deviation confidence intervals at the end of training.

For the three parameter combinations with the larger learning rate of 0.001, the metric first rapidly increases before switching polarity and decreasing. This captures

the fast emergence of smooth continua before a deterioration, which was observed for the 'fast forward' parameter combination example. For the three parameter combinations with the smaller learning rate of 0.0001, the metric increases gradually before stabilising for the parameter combinations with alpha decay constant 0.15 and 0.25.

7.6.2.5 Polynomial residual

To try to characterise the shape of the continua, a 3rd order polynomial was fit to each chunk. Fig. 7.17a plots the effect of the parameters on the residual of the polynomial fit over the course of training. The parameter combination with learning rate: 0.0001 and alpha decay 0.05, is the only one which increases and stabilises.

7.6.2.6 Polynomial 1st-3rd exponents

Fig. 7.17b plots the effect of the parameters on the exponents of the 1st-3rd fitted polynomial components over the course of training. For all combinations the first the 3rd component is close to zero. The first and then second components have the highest values, with the first component having the highest.

7.7 Discussion

The presented analyses and metrics provide a useful set of tools for exploring the emergence of first spike time continua, and are largely independent of the underlying model. The simulation project proved to be highly challenging but provided an excellent opportunity, to consider neural coding, learning and feature binding generally. As stated, the work of the chapter provided inspiration for all of the earlier chapters. This project was limited by the lack of a full theoretical understanding of the learning process, as well as the issues with hardware described in the Introduction to this chapter.

The Spike simulator provided a useful tool for other researchers in the research group and inspired optimisations to other highly used simulators. The simulator required expertise in CUDA and C++ to develop other neuron, synapse and plasticity models, however. The development of the simulator was a necessary for the efficient running of exploratory simulations which lead to the projects and hypotheses of earlier chapters.

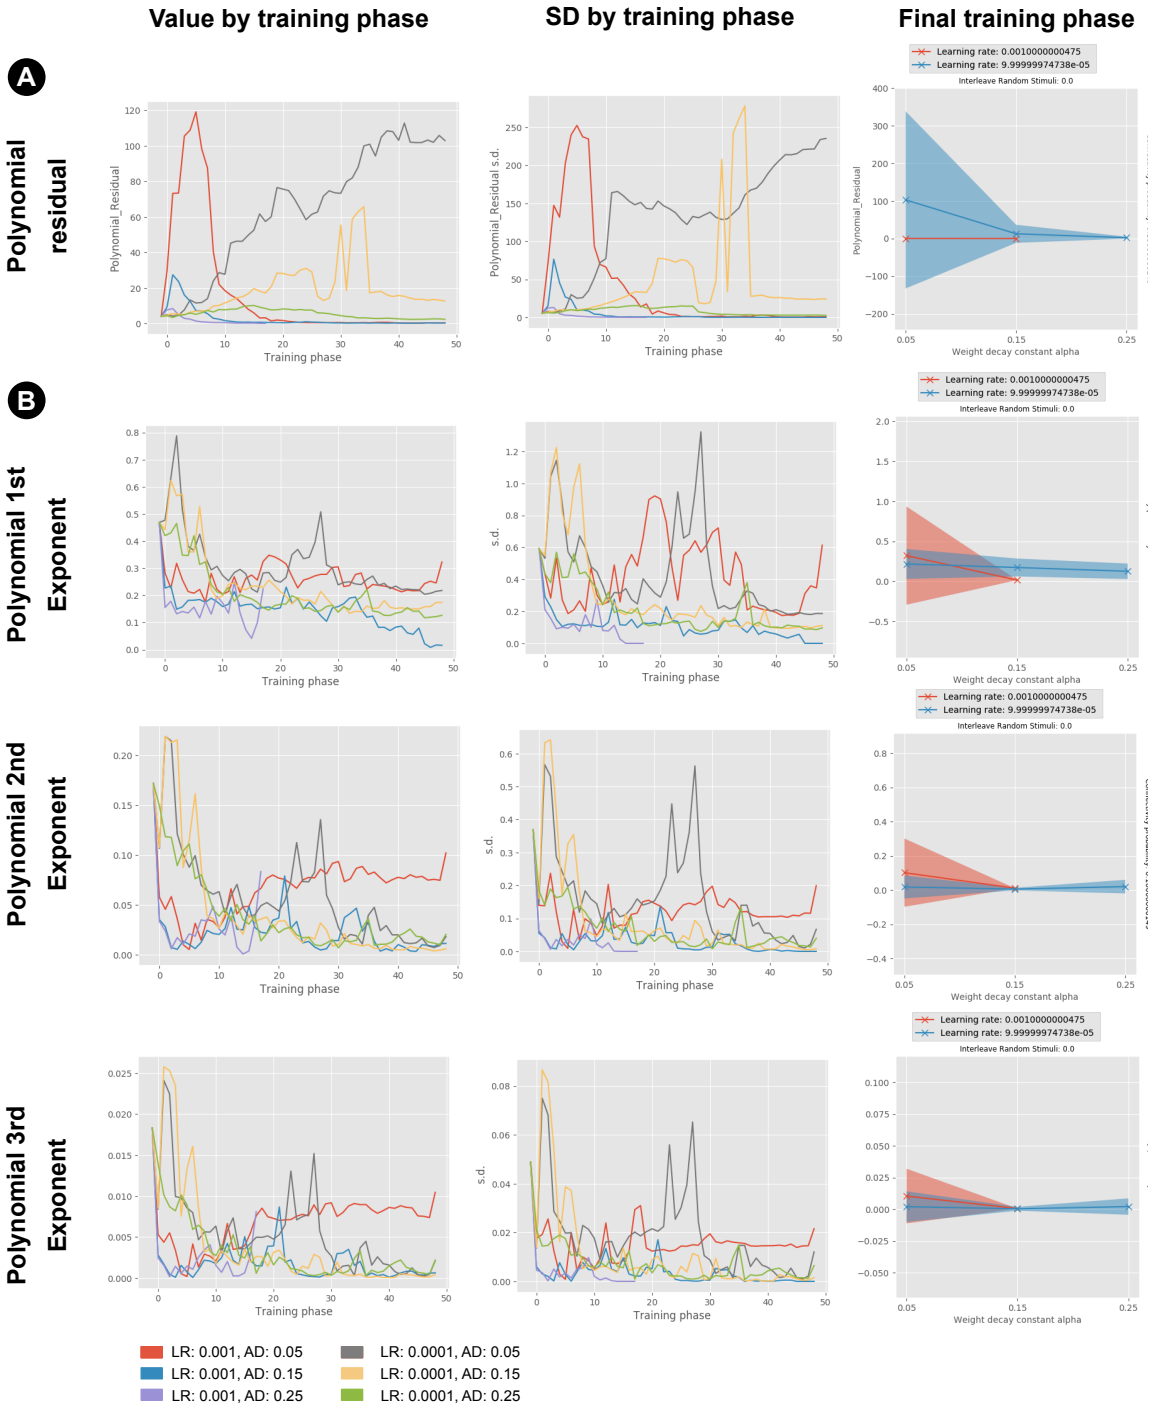


Figure 7.17: Figure shows the effect of the parameters on (a) polynomial residual, (b) 1st-3rd polynomial exponents. The 1st column shows how these values evolve over training. The 2nd column shows how the standard deviation of these values evolve over training. The 3rd column shows the final values with 3*standard deviation confidence intervals at the end of training.

Bibliography

[1] Petilla Interneuron Nomenclature Group (PING et al. "Petilla terminology: nomenclature of features of GABAergic interneurons of the cerebral cortex". In: *Nature reviews. Neuroscience* 9.7 (2008), p. 557.

[2] Moshe Abeles. *Corticonics: Neural circuits of the cerebral cortex*. Cambridge University Press, 1991.

[3] Moshe Abeles. *Local cortical circuits: an electrophysiological study*. Vol. 6. Springer Science & Business Media, 2012.

[4] Moshe Abeles and Moshe Abeles. *Interactions Between Pairs of Cells*. Vol. 6. springer, 1982. ISBN: 9783642817106. DOI: 10.1007/978-3-642-81708-3_4.

[5] Edgar D Adrian and Yngve Zotterman. "The impulses produced by sensory nerve-endings: Part II. The response of a Single End-Organ". In: *The Journal of physiology* 61.2 (1926), p. 151.

[6] Ehud Ahissar, Ronen Sosnik, and Sebastian Haidarliu. "Transformation from temporal to rate coding in a somatosensory thalamocortical pathway". In: *Nature* 406.6793 (2000), pp. 302–306.

[7] Ehud Ahissar et al. "Temporal frequency of whisker movement. II. Laminar organization of cortical representations". In: *Journal of Neurophysiology* 86.1 (2001), pp. 354–367.

[8] Nasir Ahmad et al. "Spike: A GPU Optimised Spiking Neural Network Simulator". In: *bioRxiv* (2018), p. 461160.

[9] Yael Amitai. "Thalamocortical synaptic connections: efficacy, modulation, inhibition and plasticity". In: *Reviews in the neurosciences* 12.2 (2001), pp. 159–174.

[10] Oren Amsalem et al. "Dense Computer Replica of Cortical Microcircuits Unravels Cellular Underpinnings of Auditory Surprise Response". In: *bioRxiv* (2020).

[11] Florin Amzica, Marcello Massimini, and Alfredo Manfridi. "Spatial buffering during slow and paroxysmal sleep oscillations in cortical networks of glial cells in vivo". In: *Journal of Neuroscience* 22.3 (2002), pp. 1042–1053.

[12] Ehsan Arabzadeh, Erik Zorzin, and Mathew E Diamond. "Neuronal encoding of texture in the whisker sensory pathway". In: *PLoS Biol* 3.1 (2005), e17.

[13] Amos Arieli et al. “Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses”. In: *Science* 273.5283 (1996), pp. 1868–1871.

[14] Bruno B Averbeck, Peter E Latham, and Alexandre Pouget. “Neural correlations, population coding and computation”. In: *Nature reviews neuroscience* 7.5 (2006), pp. 358–366.

[15] Pinglei Bao and Doris Y Tsao. “Representation of multiple objects in macaque category-selective areas”. In: *Nature communications* 9.1 (2018), p. 1774.

[16] Alison L Barth and James FA Poulet. “Experimental evidence for sparse firing in the neocortex”. In: *Trends in neurosciences* 35.6 (2012), pp. 345–355.

[17] Michael Beierlein, Jay R Gibson, and Barry W Connors. “A network of electrically coupled interneurons drives synchronized inhibition in neocortex”. In: *Nature neuroscience* 3.9 (2000), pp. 904–910.

[18] Guillaume Bellec et al. “Fitting summary statistics of neural data with a differentiable spiking network simulator”. In: *Advances in Neural Information Processing Systems* 34 (2021).

[19] M Beyeler et al. “CARLsim 3: A user-friendly and highly optimized library for the creation of neurobiologically detailed spiking neural networks”. In: *2015 International Joint Conference on Neural Networks (IJCNN)*. ieeexplore.ieee.org, July 2015, pp. 1–8.

[20] G Q Bi and M M Poo. “Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type”. en. In: *J. Neurosci.* 18.24 (Dec. 1998), pp. 10464–10472.

[21] GQ Bi and MM Poo. “Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type.” In: *J Neurosci Official J Soc Neurosci* 18.24 (1998), pp. 10464–72. ISSN: 0270-6474.

[22] Elie Bienenstock. “A model of neocortex”. In: *Netw Comput Neural Syst* 6.2 (1995), pp. 179–224. ISSN: 0954-898X. DOI: 10.1088/0954-898X_6_2_004.

[23] Yazan N Billeh et al. “Systematic integration of structural and functional data into multi-scale models of mouse primary visual cortex”. In: *Neuron* (2020).

[24] J Borst and Gerard G. “The low synaptic release probability in vivo”. In: *Trends in neurosciences* 33.6 (2010), pp. 259–266.

[25] William H Bosking, Justin C Crowley, and David Fitzpatrick. “Spatial coding of position and orientation in primary visual cortex”. In: *Nature neuroscience* 5.9 (2002), pp. 874–882.

[26] Conrado A Bosman et al. “Attentional stimulus selection through selective synchronization between monkey visual areas”. In: *Neuron* 75.5 (2012), pp. 875–888.

[27] Alexander van der Bourg et al. “Layer-specific refinement of sensory coding in developing mouse barrel cortex”. In: *Cerebral cortex* 27.10 (2017), pp. 4835–4850.

[28] George EP Box et al. *Time series analysis: forecasting and control*. John Wiley & Sons, 2015.

[29] Tiago Branco, Beverley A Clark, and Michael Hausser. “Dendritic discrimination of temporal input sequences in cortical neurons”. In: *Science* 329.5999 (2010), pp. 1671–1675.

[30] Tiago Branco and Michael Hausser. “Synaptic integration gradients in single cortical pyramidal cell dendrites”. In: *Neuron* 69.5 (2011), pp. 885–892.

[31] Tiago Branco and Kevin Staras. “The probability of neurotransmitter release: variability and feedback control at single synapses”. In: *Nature Reviews Neuroscience* 10.5 (2009), pp. 373–383.

[32] Michael Brecht, Wolf Singer, and Andreas K Engel. “Correlation analysis of corticotectal interactions in the cat visual system”. In: *Journal of neurophysiology* 79.5 (1998), pp. 2394–2407.

[33] Romain Brette. “Philosophy of the spike: rate-based vs. spike-based theories of the brain”. In: *Frontiers in systems neuroscience* 9 (2015), p. 151.

[34] Romain Brette and Dan F M Goodman. “Simulating spiking neural networks on GPU”. en. In: *Network* 23.4 (Oct. 2012), pp. 167–182.

[35] Scott L Brincat and Charles E Connor. “Underlying principles of visual shape selectivity in posterior inferotemporal cortex”. en. In: *Nat. Neurosci.* 7.8 (Aug. 2004), pp. 880–886.

[36] Andre R Brodtkorb, Trond R Hagen, and Martin L Saetra. “Graphics processing unit (GPU) programming strategies and trends in GPU computing”. In: *J. Parallel Distrib. Comput.* 73.1 (Jan. 2013), pp. 4–13.

[37] N Brunel. “Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons”. en. In: *J. Comput. Neurosci.* 8.3 (May 2000), pp. 183–208.

[38] Randy M Bruno and Bert Sakmann. “Cortex is driven by weak but synchronously active thalamocortical synapses”. In: *Science* 312.5780 (2006), pp. 1622–1627.

[39] Matteo Carandini. “Amplification of trial-to-trial response variability by neurons in visual cortex”. In: *PLoS Biol* 2.9 (2004), e264.

[40] Matteo Carandini and Anne K Churchland. “Probing perceptual decisions in rodents”. In: *Nature neuroscience* 16.7 (2013), pp. 824–831.

[41] George E Carvell and Daniel J Simons. “Biometric analyses of vibrissal tactile discrimination in the rat”. In: *Journal of Neuroscience* 10.8 (1990), pp. 2638–2648.

[42] Manuel A Castro-Alamancos. "Absence of rapid sensory adaptation in neocortex during information processing states". In: *Neuron* 41.3 (2004), pp. 455–464.

[43] Le Chang and Doris Y Tsao. "The code for facial identity in the primate brain". In: *Cell* 169.6 (2017), pp. 1013–1028.

[44] Jolanta Chmielowska, George E Carvell, and Daniel J Simons. "Spatial organization of thalamocortical and corticothalamic projection systems in the rat SmI barrel cortex". In: *Journal of Comparative Neurology* 285.3 (1989), pp. 325–338.

[45] Sang Chul Chong and Anne Treisman. "Attentional spread in the statistical processing of visual displays". In: *Perception & Psychophysics* 67.1 (2005), pp. 1–13.

[46] Sooyoung Chung, Xiangrui Li, and Sacha B Nelson. "Short-term depression at thalamocortical synapses contributes to rapid adaptation of cortical sensory responses in vivo". In: *Neuron* 34.3 (2002), pp. 437–446.

[47] Mark M Churchland et al. "Neural population dynamics during reaching". In: *Nature* 487.7405 (2012), pp. 51–56.

[48] WB Clarke and David Bowsher. "Terminal distribution of primary afferent trigeminal fibers in the rat". In: *Experimental neurology* 6.5 (1962), pp. 372–383.

[49] Christine M Constantinople and Randy M Bruno. "Deep cortical layers are activated directly by thalamus". In: *Science* 340.6140 (2013), pp. 1591–1594.

[50] Alan Cowey. "Cortical maps and visual perception the grindley memorial lecture". In: *Quarterly Journal of Experimental Psychology* 31.1 (1979), pp. 1–17.

[51] Carina Curto et al. "A simple model of cortical dynamics explains variability and state dependence of sensory responses in urethane-anesthetized auditory cortex". In: *Journal of neuroscience* 29.34 (2009), pp. 10600–10612.

[52] Henry WP Dalgleish et al. "How many neurons are sufficient for perception of cortical activity?" In: *Elife* 9 (2020), e58889.

[53] Neil Dalphin et al. "Voltage imaging of cortical oscillations in layer 1 with two-photon microscopy". In: *Eneuro* 7.3 (2020).

[54] CPJ De Kock et al. "Layer-and cell-type-specific suprathreshold stimulus representation in rat primary somatosensory cortex". In: *The Journal of physiology* 581.1 (2007), pp. 139–154.

[55] Daniel Deitch, Alon Rubin, and Yaniv Ziv. "Representational drift in the mouse visual cortex". In: *bioRxiv* (2020).

[56] Martin Deschenes, Elena Timofeeva, and Philippe Lavallee. "The relay of high-frequency sensory signals in the whisker-to-barreloid pathway". In: *Journal of Neuroscience* 23.17 (2003), pp. 6778–6787.

[57] Martin Deschenes, Pierre Veinante, and Zhong-Wei Zhang. "The organization of corticothalamic projections: reciprocity versus parity". In: *Brain research reviews* 28.3 (1998), pp. 286–308.

[58] Martin Deschenes et al. "The vibrissal system as a model of thalamic operations". In: *Progress in brain research* 149 (2005), pp. 31–40.

[59] Mathew E Diamond and Ehsan Arabzadeh. "Whisker sensory system–from receptor to decision". In: *Progress in neurobiology* 103 (2013), pp. 28–40.

[60] James J DiCarlo, Davide Zoccolan, and Nicole C Rust. "How does the brain solve visual object recognition?" In: *Neuron* 73.3 (2012), pp. 415–434.

[61] Peter Diehl and Matthew Cook. "Unsupervised learning of digit recognition using spike-timing-dependent plasticity". In: *Frontiers in Computational Neuroscience* 9.August (2015), pp. 99–99. DOI: 10.3389/fncom.2015.00099.

[62] Adrien Doerig et al. "Capsule networks as recurrent models of grouping and segmentation". In: *PLoS computational biology* 16.7 (2020), e1008017.

[63] Brent Doiron et al. "The mechanics of state-dependent neural correlations". In: *Nature neuroscience* 19.3 (2016), pp. 383–393.

[64] Rodney J Douglas and Kevan AC Martin. "Neuronal circuits of the neocortex". In: *Annu. Rev. Neurosci.* 27 (2004), pp. 419–451.

[65] Valentin Dragoi et al. "Dynamics of neuronal sensitivity in visual cortex and local feature discrimination". In: *Nature neuroscience* 5.9 (2002), pp. 883–891.

[66] Shaul Druckmann et al. "A novel multiple objective optimization framework for constraining conductance-based neuron models by experimental data". In: *Frontiers in neuroscience* 1 (2007), p. 1.

[67] Rodrigo Echeveste and Claudius Gros. "Two-Trace Model for Spike-Timing-Dependent Synaptic Plasticity". In: *Neural Comput* 27.3 (2015), pp. 672–698. ISSN: 0899-7667. DOI: 10.1162/NECO_a_00707.

[68] Robert Egger et al. "Cortical Output Is Gated by Horizontally Projecting Neurons in the Deep Layers". In: *Neuron* 105.1 (2020), pp. 122–137.

[69] Robert Egger et al. "Robustness of sensory-evoked excitation is increased by inhibitory inputs to distal apical tuft dendrites". In: *Proceedings of the National Academy of Sciences* 112.45 (2015), pp. 14072–14077.

[70] Akihiro Eguchi et al. "The emergence of polychronization and feature binding in a spiking neural network model of the primate ventral visual system". en. In: *Psychol. Rev.* 125.4 (July 2018), pp. 545–571.

[71] Gaute T Einevoll et al. "Towards reliable spike-train recordings from thousands of neurons with multielectrodes". In: *Current opinion in neurobiology* 22.1 (2012), pp. 11–17.

[72] Jochen Martin Eppler et al. "PyNEST: A Convenient Interface to the NEST Simulator". en. In: *Front. Neuroinform.* 2 (2008), p. 12.

[73] Nafise Erfanian Saeedi et al. “Learning Pitch with STDP: A Computational Model of Place and Temporal Pitch Perception Using Spiking Neural Networks”. en. In: *PLoS Comput. Biol.* 12.4 (Apr. 2016), e1004860.

[74] Martin Ester et al. “A density-based algorithm for discovering clusters in large spatial databases with noise.” In: *Kdd*. Vol. 96. 34. 1996, pp. 226–231.

[75] Benjamin D Evans. “Transformation-invariant visual representations in self-organizing spiking neural networks”. In: *Frontiers in Computational Neuroscience* 6.July (2012), pp. 1–19. DOI: 10.3389/fncom.2012.00046.

[76] Claudia E Feierstein et al. “Representation of spatial goals in rat orbitofrontal cortex”. In: *Neuron* 51.4 (2006), pp. 495–507.

[77] Daniel E Feldman. “The Spike-Timing Dependence of Plasticity”. In: *Neuron* 75.4 (2012), pp. 556–571. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2012.08.001.

[78] Jerome Feldman. “The neural binding problem(s)”. In: *Cogn Neurodynamics* 7.1 (2013), pp. 1–11. ISSN: 1871-4080. DOI: 10.1007/s11571-012-9219-8.

[79] Dirk Feldmeyer. “Excitatory neuronal connectivity in the barrel cortex”. In: *Frontiers in neuroanatomy* 6 (2012), p. 24.

[80] Dirk Feldmeyer et al. “Barrel cortex function”. In: *Progress in neurobiology* 103 (2013), pp. 3–27.

[81] A K Fidjeland et al. “NeMo: A Platform for Neural Modelling of Spiking Neurons Using GPUs”. In: *2009 20th IEEE International Conference on Application-specific Systems, Architectures and Processors*. July 2009, pp. 137–144.

[82] David J Freedman et al. “A comparison of primate prefrontal and inferior temporal cortices during visual categorization”. In: *Journal of Neuroscience* 23.12 (2003), pp. 5235–5246.

[83] Winrich A Freiwald and Doris Y Tsao. “Functional compartmentalization and viewpoint generalization within the macaque face-processing system”. In: *Science* 330.6005 (2010), pp. 845–851.

[84] Pascal Fries. “A mechanism for cognitive dynamics: neuronal communication through neuronal coherence”. In: *Trends in cognitive sciences* 9.10 (2005), pp. 474–480.

[85] Pascal Fries. “Rhythms for cognition: communication through coherence”. In: *Neuron* 88.1 (2015), pp. 220–235.

[86] Pascal Fries, Danko Nikolic, and Wolf Singer. “The gamma cycle”. In: *Trends in neurosciences* 30.7 (2007), pp. 309–316.

[87] RC Froemke, IA Tsay, and Raad - M of. “Contribution of individual spikes in burst-induced long-term synaptic modification”. In: *Journal of Neurophysiology* 95.3 (2006), pp. 1620–1629.

[88] Galit Fuhrmann et al. “Coding of temporal information by activity-dependent synapses”. In: *Journal of neurophysiology* (2002).

[89] Takahiro Furuta, Takeshi Kaneko, and Martin Deschenes. "Septal neurons in barrel cortex derive their receptive field input from the lemniscal pathway". In: *Journal of Neuroscience* 29.13 (2009), pp. 4089–4095.

[90] Takahiro Furuta et al. "Corticofugal control of vibrissa-sensitive neurons in the interpolaris nucleus of the trigeminal complex". In: *Journal of Neuroscience* 30.5 (2010), pp. 1832–1838.

[91] Juan A Gallego et al. "Neural manifolds for the control of movement". In: *Neuron* 94.5 (2017), pp. 978–984.

[92] Sonia Gasparini and Jeffrey C Magee. "State-dependent dendritic computation in hippocampal CA1 pyramidal neurons". In: *Journal of Neuroscience* 26.7 (2006), pp. 2088–2100.

[93] Wulfram Gerstner et al. "Neural codes: firing rates and beyond". In: *Proceedings of the National Academy of Sciences* 94.24 (1997), pp. 12740–12741.

[94] Wulfram Gerstner et al. *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press, 2014.

[95] John M Gibson and Wally I Welker. "Quantitative studies of stimulus coding in first-order vibrissa afferents of rats. 2. Adaptation and coding of stimulus parameters". In: *Somatosensory research* 1.2 (1983), pp. 95–117.

[96] Ziv Gil, Barry W Connors, and Yael Amitai. "Efficacy of thalamocortical and intracortical synaptic connections: quanta, innervation, and reliability". In: *Neuron* 23.2 (1999), pp. 385–397.

[97] Paul M Gochin. "The representation of shape in the temporal lobe". In: *Behavioural brain research* 76.1-2 (1996), pp. 99–116.

[98] Tim Gollisch and Markus Meister. "Rapid neural coding in the retina with relative spike latencies". In: *science* 319.5866 (2008), pp. 1108–1111.

[99] Melvyn A Goodale, A David Milner, et al. "Separate visual pathways for perception and action". In: (1992).

[100] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. "Explaining and harnessing adversarial examples". In: *arXiv preprint arXiv:1412.6572* (2014).

[101] Dan F M Goodman and Romain Brette. "Learning to Localise Sounds with Spiking Neural Networks". In: *Proceedings of the 23rd International Conference on Neural Information Processing Systems - Volume 1*. NIPS'10. Vancouver, British Columbia, Canada: Curran Associates Inc., 2010, pp. 784–792.

[102] Dan F M Goodman and Romain Brette. "The brian simulator". en. In: *Front. Neurosci.* 3.2 (Sept. 2009), pp. 192–197.

[103] Charles M Gray et al. "Oscillatory responses in cat visual cortex exhibit intercolumnar synchronization which reflects global stimulus properties". In: *Nature* 338.6213 (1989), p. 334.

[104] Charles M Gray et al. “Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex”. In: *Journal of neuroscience methods* 63.1-2 (1995), pp. 43–54.

[105] R Gutig and R Aharonov. “Learning input correlations through nonlinear temporally asymmetric Hebbian plasticity”. In: *The Journal of Neuroscience* 23.9 (2003), pp. 3699–3714.

[106] Robert Gutig and Haim Sompolinsky. “Time-warp–invariant neuronal processing”. In: *PLoS biology* 7.7 (2009), e1000141.

[107] Torkel Hafting et al. “Microstructure of a spatial map in the entorhinal cortex”. In: *Nature* 436.7052 (2005), p. 801.

[108] Kenneth D Harris. “Nonsense correlations in neuroscience”. In: *bioRxiv* (2020).

[109] Kenneth D Harris and Alexander Thiele. “Cortical state and attention”. In: *Nature reviews neuroscience* 12.9 (2011), pp. 509–523.

[110] Kenneth D Harris et al. “Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements”. In: *Journal of neurophysiology* 84.1 (2000), pp. 401–414.

[111] Martha N Havenith et al. “Synchrony makes neurons fire in sequence, and stimulus properties determine who is ahead”. In: *Journal of neuroscience* 31.23 (2011), pp. 8570–8584.

[112] Jeff Hawkins et al. “A Framework for Intelligence and Cortical Function Based on Grid Cells in the Neocortex”. In: *Frontiers in Neural Circuits* 12 (2019), p. 121. ISSN: 1662-5110. DOI: 10.3389/fncir.2018.00121. URL: https://www.frontiersin.org/article/10.3389/fncir.2018.00121.

[113] Lynn Hazan, Michael Zugaro, and Gyorgy Buzsaki. “Klusters, NeuroScope, NDManager: a free software suite for neurophysiological data processing and visualization”. In: *Journal of neuroscience methods* 155.2 (2006), pp. 207–216.

[114] Jaime E Heiss et al. “Shift in the balance between excitation and inhibition during sensory adaptation of S1 neurons”. In: *Journal of Neuroscience* 28.49 (2008), pp. 13320–13330.

[115] Michael J Higley and Diego Contreras. “Balanced excitation and inhibition determine spike timing during frequency adaptation”. In: *Journal of Neuroscience* 26.2 (2006), pp. 448–457.

[116] Sean L Hill et al. “Statistical connectivity provides a sufficient foundation for specific functional connectivity in neocortical neural microcircuits”. In: *Proceedings of the National Academy of Sciences* 109.42 (2012), E2885–E2894.

[117] M L Hines and N T Carnevale. “The NEURON simulation environment”. en. In: *Neural Comput.* 9.6 (Aug. 1997), pp. 1179–1209.

[118] SL Ho and Min Xie. “The use of ARIMA models for reliability forecasting and analysis”. In: *Computers & industrial engineering* 35.1-2 (1998), pp. 213–216.

[119] Huu Hoang et al. “Improved hyperacuity estimation of spike timing from calcium imaging”. In: *Scientific reports* 10.1 (2020), pp. 1–16.

[120] Gregory L Holst et al. “Autonomous patch-clamp robot for functional characterization of neurons in vivo: development and application to mouse visual cortex”. In: *Journal of neurophysiology* 121.6 (2019), pp. 2341–2357.

[121] Ha Hong et al. “Explicit information for category-orthogonal object properties increases along the ventral stream”. In: *Nature neuroscience* 19.4 (2016), p. 613.

[122] Illia Horenko. “On a Scalable Entropic Breaching of the Overfitting Barrier for Small Data Problems in Machine Learning”. In: *Neural Computation* (2020), pp. 1–17.

[123] Jonathan C Horton and Daniel L Adams. “The cortical column: a structure without a function”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 360.1456 (2005), pp. 837–862.

[124] Michael R Houck and James E Hoffman. “Conjunction of color and form without attention: Evidence from an orientation-contingent color aftereffect.” In: *Journal of Experimental Psychology: Human Perception and Performance* 12.2 (1986), p. 186.

[125] Tomas Hromadka, Michael R DeWeese, and Anthony M Zador. “Sparse representation of sounds in the unanesthetized auditory cortex”. In: *PLoS Biol* 6.1 (2008), e16.

[126] D H Hubel and T N Wiesel. “Integrative action in the cat’s lateral geniculate body”. en. In: *J. Physiol.* 155 (Feb. 1961), pp. 385–398.

[127] D H Hubel and T N Wiesel. “Receptive fields and functional architecture in two non-striate visual areas (18 and 19) of the cat”. en. In: *J. Neurophysiol.* 28 (Mar. 1965), pp. 229–289.

[128] D H Hubel and T N Wiesel. *Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex*. 1962.

[129] DH Hubel and TN Wiesel. “Functional architecture of macaque monkey visual cortex”. In: *Proc. Roy. Soc. Lond. B* 198.1-59 (1977).

[130] DH Hubel and TN Wiesel. “Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex”. In: *J Physiology* 160.1 (1962), pp. 106–154. ISSN: 1469-7793. DOI: 10.1113/jphysiol.1962.sp006837.

[131] Mark D Humphries. “Strong and weak principles of neural dimension reduction”. In: *arXiv preprint arXiv:2011.08088* (2020).

[132] Chia-Chun Hung, Eric T Carlson, and Charles E Connor. “Medial axis shape coding in macaque inferotemporal cortex”. en. In: *Neuron* 74.6 (June 2012), pp. 1099–1113.

[133] Chou P Hung et al. “Fast readout of object identity from macaque inferior temporal cortex”. In: *Science* 310.5749 (2005), pp. 863–866.

[134] Cole Hurwitz et al. “Building population models for large-scale neural recordings: opportunities and pitfalls”. In: *arXiv preprint arXiv:2102.01807* (2021).

[135] James B Isbister et al. “A new approach to solving the feature-binding problem in primate vision”. In: *Interface Focus* 8.4 (2018), p. 20180021.

[136] James B Isbister et al. “Clustering and control for adaptation uncovers time-warped spike time patterns in cortical networks in vivo”. In: *Scientific Reports* 11.1 (2021), pp. 1–20.

[137] Minami Ito et al. “Size and position invariance of neuronal responses in monkey inferotemporal cortex”. In: *Journal of neurophysiology* 73.1 (1995), pp. 218–226.

[138] Eugene M Izhikevich. “Polychronization: computation with spikes”. In: *Neural computation* 18.2 (2006), pp. 245–282.

[139] Roland S Johansson and Ingvars Birznieks. “First spikes in ensembles of human tactile afferents code complex spatial fingertip events”. In: *Nature neuroscience* 7.2 (2004), pp. 170–177.

[140] Kenneth O Johnson. “Neural coding”. In: *Neuron* 26.3 (2000), pp. 563–566.

[141] Brian L Jones and Stephen M Smith. “Calcium-sensing receptor: a key target for extracellular calcium signaling in neurons”. In: *Frontiers in physiology* 7 (2016), p. 116.

[142] HC Jones and RF Keep. “Brain fluid calcium concentration and response to acute hypercalcaemia during development in the rat.” In: *The Journal of Physiology* 402.1 (1988), pp. 579–593.

[143] Jonathan C Kao et al. “Single-trial dynamics of motor cortex and their applications to brain-machine interfaces”. In: *Nature communications* 6.1 (2015), pp. 1–12.

[144] Kohitij Kar et al. “Evidence that recurrent circuits are critical to the ventral stream’s execution of core object recognition behavior”. In: *Nature neuroscience* 22.6 (2019), p. 974.

[145] Narayanan Kasthuri et al. “Saturated Reconstruction of a Volume of Neocortex”. en. In: *Cell* 162.3 (July 2015), pp. 648–661.

[146] Yonatan Katz, Michael Okun, and Ilan Lampl. “Trial-to-trial correlation between thalamic sensory response and global EEG activity”. In: *European Journal of Neuroscience* 35.6 (2012), pp. 826–837.

[147] Jason ND Kerr, David Greenberg, and Fritjof Helmchen. “Imaging input and output of neocortical networks in vivo”. In: *Proceedings of the National Academy of Sciences* 102.39 (2005), pp. 14063–14068.

[148] Jason ND Kerr et al. “Spatial organization of neuronal population responses in layer 2/3 of rat barrel cortex”. In: *Journal of neuroscience* 27.48 (2007), pp. 13316–13328.

[149] Vivek Khatri, Randy M Bruno, and Daniel J Simons. "Stimulus-specific and stimulus-nonspecific firing synchrony and its modulation by sensory adaptation in the whisker-to-barrel pathway". In: *Journal of neurophysiology* 101.5 (2009), pp. 2328–2338.

[150] Vivek Khatri, Jed A Hartings, and Daniel J Simons. "Adaptation in thalamic barreloid and cortical barrel neurons to periodic whisker deflections varying in frequency and velocity". In: *Journal of neurophysiology* 92.6 (2004), pp. 3244–3254.

[151] Ehsan Kheradpezhouh, Mehdi Adibi, and Ehsan Arabzadeh. "Response dynamics of rat barrel cortex neurons to repeated sensory stimulation". In: *Scientific reports* 7.1 (2017), pp. 1–10.

[152] Saeed Kheradpisheh, Mohammad Ganjtabesh, and Timothee Masquelier. "Bio-inspired unsupervised learning of visual features leads to robust invariant object recognition". In: *Neurocomputing* 205 (2016), pp. 382–392. ISSN: 0925-2312. DOI: 10.1016/j.neucom.2016.04.029.

[153] Christiaan PJ de Kock and Bert Sakmann. "Spiking in primary somatosensory cortex during natural whisking in awake head-restrained rats is cell-type specific". In: *Proceedings of the National Academy of Sciences* 106.38 (2009), pp. 16446–16450.

[154] Yves Kremer et al. "Late emergence of the vibrissa direction selectivity map in the rat barrel cortex". In: *Journal of Neuroscience* 31.29 (2011), pp. 10689–10700.

[155] Nikolaus Kriegeskorte and Xue-Xin Wei. "Neural tuning and representational geometry". In: *arXiv preprint arXiv:2104.09743* (2021).

[156] Peter W Land, Sam A Buffer Jr, and Jamie Dawn Yaskosky. "Barreloids in adult rat thalamus: three-dimensional architecture and relationship to somatosensory cortical barrels". In: *Journal of Comparative Neurology* 355.4 (1995), pp. 573–588.

[157] Patrick N Lawlor et al. "Linear-nonlinear-time-warp-poisson models of neural activity". In: *Journal of computational neuroscience* 45.3 (2018), pp. 173–191.

[158] SooHyun Lee et al. "The largest group of superficial neocortical GABAergic interneurons expresses ionotropic serotonin receptors". In: *Journal of Neuroscience* 30.50 (2010), pp. 16796–16808.

[159] Amy M LeMessurier and Daniel E Feldman. "Plasticity of population coding in primary sensory cortex". In: *Current opinion in neurobiology* 53 (2018), pp. 50–56.

[160] SH Lichtenstein, GE Carvell, and DJ Simons. "Responses of rat trigeminal ganglion neurons to movements of vibrissae in different directions". In: *Somatosensory & motor research* 7.1 (1990), pp. 47–65.

[161] Anthony D Lien and Massimo Scanziani. "Tuned thalamic excitation is amplified by visual cortical circuits". In: *Nature neuroscience* 16.9 (2013), pp. 1315–1323.

[162] I-Chun Lin et al. "The nature of shared cortical variability". In: *Neuron* 87.3 (2015), pp. 644–656.

[163] Charl Linssen et al. *NEST 2.16.0*. Aug. 2018. DOI: 10.5281/zenodo.1400175. URL: https://doi.org/10.5281/zenodo.1400175.

[164] Margaret Livingstone and David Hubel. "Segregation of form, color, movement, and depth: anatomy, physiology, and perception". In: *Science* 240.4853 (1988), pp. 740–749.

[165] Magor L Lorincz et al. "A distinct class of slow (0.2–2 Hz) intrinsically bursting layer 5 pyramidal neurons determines UP/DOWN state dynamics in the neocortex". In: *Journal of Neuroscience* 35.14 (2015), pp. 5442–5458.

[166] Artur Luczak et al. "Sequential structure of neocortical spontaneous activity in vivo". In: *Proceedings of the National Academy of Sciences* 104.1 (2007), pp. 347–352.

[167] Stephen L Macknik and Margaret S Livingstone. "Neuronal correlates of visibility and invisibility in the primate visual system". In: *Nature neuroscience* 1.2 (1998), pp. 144–149.

[168] Zachary F Mainen and Terrence J Sejnowski. "Reliability of spike timing in neocortical neurons". In: *Science* 268.5216 (1995), pp. 1503–1506.

[169] Katayun Cohen-Kashi Malina et al. "Imbalance between excitation and inhibition in the somatosensory cortex produces postadaptation facilitation". In: *Journal of Neuroscience* 33.19 (2013), pp. 8463–8471.

[170] Christoph von der Malsburg. "The correlation theory of brain function (Internal Report 81-2)". In: *Goettingen: Department of Neurobiology, Max Planck Intitute for Biophysical Chemistry* (1981).

[171] Christoph von der Malsburg. "The What and Why of Binding". In: *Neuron* 24.1 (1999), pp. 95–104. ISSN: 0896-6273 (Print). DOI: 10.1016/S0896-6273(00)80825-9.

[172] David J Margolis et al. "Reorganization of cortical population activity imaged throughout long-term sensory deprivation". In: *Nature neuroscience* 15.11 (2012), p. 1539.

[173] Henry Markram. "Fixing the location and dimensions of functional neocortical columns". In: (2008).

[174] Henry Markram et al. "Interneurons of the neocortical inhibitory system". In: *Nature reviews neuroscience* 5.10 (2004), pp. 793–807.

[175] Henry Markram et al. "Reconstruction and simulation of neocortical microcircuitry". In: *Cell* 163.2 (2015), pp. 456–492.

[176] Henry Markram et al. “Regulation of synaptic efficacy by coincidence of post-synaptic APs and EPSPs”. In: *Science* 275.5297 (1997), pp. 213–215.

[177] Christoph von der Marlsburg. “The Correlation Theory of Brain Function”. In: *Dept. of Neurobiology, Max-Planck-Institute for Biophysical Chemistry* (1981).

[178] Timothee Masquelier and Simon J Thorpe. “Unsupervised learning of visual features through spike timing dependent plasticity”. In: *PLoS Computational Biology* 3.2 (2007), pp. 0247–0257. ISSN: 1553-7358. DOI: 10.1371/journal.pcbi.0030031.

[179] Marcello Massimini and Florin Amzica. “Extracellular calcium fluctuations and intracellular potentials in the cortex during the slow sleep oscillation”. In: *Journal of neurophysiology* 85.3 (2001), pp. 1346–1350.

[180] Johannes M Mayrhofer et al. “Sparse, reliable, and long-term stable representation of periodic whisker deflections in the mouse barrel cortex”. In: *Neuroimage* 115 (2015), pp. 52–63.

[181] Celeste McCollough. “Color adaptation of edge-detectors in the human visual system”. In: *Science* 149.3688 (1965), pp. 1115–1116.

[182] Matthew J McGinley, Stephen V David, and David A McCormick. “Cortical membrane potential signature of optimal states for sensory signal detection”. In: *Neuron* 87.1 (2015), pp. 179–192.

[183] Philipp Metzner, Lars Putzig, and Illia Horenko. “Analysis of persistent non-stationary time series and applications”. In: *Communications in Applied Mathematics and Computational Science* 7.2 (2012), pp. 175–229.

[184] Hanno S Meyer et al. “Cell type–specific thalamic innervation in a column of rat vibrissal cortex”. In: *Cerebral cortex* 20.10 (2010), pp. 2287–2303.

[185] Peter M Milner. “A model for visual shape recognition.” In: *Psychological review* 81.6 (1974), p. 521.

[186] Brandon S Minnery and Daniel J Simons. “Response properties of whisker-associated trigeminothalamic neurons in rat nucleus principalis”. In: *Journal of neurophysiology* 89.1 (2003), pp. 40–56.

[187] M J Mivsic, DJ M Djurdjevic, and M V Tomav sevic. “Evolution and trends in GPU computing”. In: *2012 Proceedings of the 35th International Convention MIPRO*. May 2012, pp. 289–294.

[188] Toshio Miyashita and Daniel E Feldman. “Behavioral detection of passive whisker stimuli requires somatosensory cortex”. In: *Cerebral cortex* 23.7 (2013), pp. 1655–1662.

[189] Boaz Mohar, Yonatan Katz, and Ilan Lampl. “Opposite adaptive processing of stimulus intensity in two major nuclei of the somatosensory brainstem”. In: *Journal of Neuroscience* 33.39 (2013), pp. 15394–15400.

[190] Ruben Moreno-Bote et al. “Information-limiting correlations”. In: *Nature neuroscience* 17.10 (2014), pp. 1410–1417.

[191] A Morrison and M Diesmann. “Phenomenological models of synaptic plasticity based on spike timing”. In: *Biological cybernetics* 98 (2008), pp. 459–478.

[192] Abigail Morrison et al. “Advancing the boundaries of high-connectivity network simulation with distributed computing”. en. In: *Neural Comput.* 17.8 (Aug. 2005), pp. 1776–1801.

[193] Vernon B Mountcastle. *Perceptual neuroscience: The cerebral cortex*. Harvard University Press, 1998.

[194] JA Movshon, Ian D Thompson, and David J Tolhurst. “Spatial summation in the receptive fields of simple cells in the cat’s striate cortex.” In: *The Journal of physiology* 283.1 (1978), pp. 53–77.

[195] Robert Nau. *Identifying the orders of AR and MA terms in an ARIMA model*. `https://people.duke.edu/~rnau/411diff.htm`. (Accessed on 10/28/2020). Oct. 2020.

[196] Robert Nau. *Identifying the orders of AR and MA terms in an ARIMA model*. `https://people.duke.edu/~rnau/411arim3.htm`. (Accessed on 10/28/2020). Oct. 2020.

[197] Louis A Necker. “LXI. Observations on some remarkable optical phaenomena seen in Switzerland; and on an optical phaenomenon which occurs on viewing a figure of a crystal or geometrical solid”. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 1.5 (1832), pp. 329–337.

[198] Taylor H Newton et al. “Voltage-sensitive dye imaging reveals inhibitory modulation of ongoing cortical activity”. In: *bioRxiv* (2019), p. 812008.

[199] Thierry Nieus et al. “State-dependent representation of stimulus-evoked activity in high-density recordings of neural cultures”. In: *Scientific reports* 8.1 (2018), pp. 1–19.

[200] Max Nolte et al. “Cortical reliability amid noise and chaos”. In: *Nature communications* 10.1 (2019), pp. 1–15.

[201] Max Nolte et al. “Impact of higher order network structure on emergent cortical activity”. In: *Network Neuroscience* 4.1 (2020), pp. 292–314.

[202] Cuda Nvidia. “Cuda C programming guide v8. 0”. In: *Nvidia Corporation* (2017).

[203] Daniel H O’Connor et al. “Neural activity in barrel cortex underlying vibrissa-based object localization in mice”. In: *Neuron* 67.6 (2010), pp. 1048–1061.

[204] Christian O’Reilly et al. “Rodent somatosensory thalamocortical circuitry: neurons, synapses, and connectivity”. In: (2020).

[205] Michael Okun and Ilan Lampl. “Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities”. In: *Nature neuroscience* 11.5 (2008), pp. 535–537.

[206] Michael Okun et al. “Diverse coupling of neurons to populations in sensory cortex”. In: *Nature* 521.7553 (2015), pp. 511–515.

[207] Bruno A Olshausen and David J Field. “What is the other 85 percent of V1 doing”. In: *L. van Hemmen, & T. Sejnowski (Eds.)* 23 (2006), pp. 182–211.

[208] Gonzalo H Otazu et al. “Engaging in an auditory task suppresses responses in auditory cortex”. In: *Nature neuroscience* 12.5 (2009), p. 646.

[209] Chethan Pandarinath et al. “Inferring single-trial neural population dynamics using sequential auto-encoders”. In: *Nature methods* 15.10 (2018), pp. 805–815.

[210] Stefano Panzeri et al. “Neural population coding: combining insights from microscopic and mass signals”. In: *Trends in cognitive sciences* 19.3 (2015), pp. 162–172.

[211] Stefano Panzeri et al. “The role of spike timing in the coding of stimulus location in rat somatosensory cortex”. In: *Neuron* 29.3 (2001), pp. 769–777.

[212] Harold Pashler. “Detecting conjunctions of color and form: Reassessing the serial search hypothesis”. In: *Perception & Psychophysics* 41.3 (1987), pp. 191–201.

[213] A Pasupathy and C E Connor. “Shape representation in area V4: position-specific tuning for boundary conformation”. en. In: *J. Neurophysiol.* 86.5 (Nov. 2001), pp. 2505–2519.

[214] Helene Paugam-Moisy, Regis Martinez, and Samy Bengio. “Delay learning and polychronization for reservoir computing”. In: *Neurocomputing* 71.7-9 (2008), pp. 1143–1158. ISSN: 09252312. DOI: 10.1016/j.neucom.2007.12.027.

[215] Hellne Paugam-Moisy, Regis Martinez, and Samy Bengio. “Delay learning and polychronization for reservoir computing”. In: *Neurocomputing* 71.7–9 (2008), pp. 1143–1158.

[216] F. Pedregosa et al. “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830.

[217] Rodrigo Perin, Thomas K Berger, and Henry Markram. “A synaptic organizing principle for cortical neuronal groups”. In: *Proceedings of the National Academy of Sciences* 108.13 (2011), pp. 5419–5424.

[218] David I Perrett, Edmond T Rolls, and Woody Caan. “Visual neurones responsive to faces in the monkey temporal cortex”. In: *Experimental brain research* 47.3 (1982), pp. 329–342.

[219] Carl CH Petersen et al. “Interaction of sensory responses with spontaneous depolarization in layer 2/3 barrel cortex”. In: *Proceedings of the National Academy of Sciences* 100.23 (2003), pp. 13638–13643.

[220] James FA Poulet and Carl CH Petersen. “Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice”. In: *Nature* 454.7206 (2008), p. 881.

[221] Rodrigo Quian Quiroga and Stefano Panzeri. "Extracting information from neuronal populations: information theory and decoding approaches". In: *Nature Reviews Neuroscience* 10.3 (2009), pp. 173–185.

[222] Michael W Reimann et al. "A null model of the mouse whole-neocortex micro-connectome". In: *Nature communications* 10.1 (2019), pp. 1–16.

[223] Michael W Reimann et al. "An algorithm to predict the connectome of neural microcircuits". In: *Frontiers in computational neuroscience* 9 (2015), p. 28.

[224] Michael W Reimann et al. "Cliques of neurons bound into cavities provide a missing link between structure and function". In: *Frontiers in computational neuroscience* 11 (2017), p. 48.

[225] Michael W Reimann et al. "Topology of synaptic connectivity constrains neuronal stimulus representation, predicting two complementary coding strategies". In: *bioRxiv* (2020).

[226] Pamela Reinagel and R Clay Reid. "Precise firing events are conserved across neurons". In: *Journal of Neuroscience* 22.16 (2002), pp. 6837–6841.

[227] Kimberly Reinhold, Anthony D Lien, and Massimo Scanziani. "Distinct recurrent versus afferent dynamics in cortical visual processing". In: *Nature neuroscience* 18.12 (2015), pp. 1789–1797.

[228] Banafsheh Rekabdar et al. "A Scale and Translation Invariant Approach for Early Classification of Spatio-Temporal Patterns Using Spiking Neural Networks". In: *Neural Process Lett* 43.2 (2015), pp. 327–343. ISSN: 1370-4621. DOI: 10.1007/s11063-015-9436-3.

[229] Banafsheh Rekabdar et al. "Using patterns of firing neurons in spiking neural networks for learning and early recognition of spatio-temporal patterns". In: *Neural Comput Appl* 28.5 (2016), pp. 881–897. ISSN: 0941-0643. DOI: 10.1007/s00521-016-2283-y.

[230] Arbora Resulaj et al. "First spikes in visual cortex enable perceptual discrimination". In: *Elife* 7 (2018), e34044.

[231] Vicente Reyes-Puerta et al. "High stimulus-related information in barrel cortex inhibitory interneurons". In: *PLoS computational biology* 11.6 (2015), e1004121.

[232] Vicente Reyes-Puerta et al. "Laminar and columnar structure of sensory-evoked multineuronal spike sequences in adult rat barrel cortex in vivo". In: *Cerebral Cortex* 25.8 (2014), pp. 2001–2021.

[233] Vicente Reyes-Puerta et al. "Long-range intralaminar noise correlations in the barrel cortex". In: *Journal of neurophysiology* 113.9 (2015), pp. 3410–3420.

[234] Vicente Reyes-Puerta et al. "Propagation of spontaneous slow-wave activity across columns and layers of the adult rat barrel cortex in vivo". In: *Brain Structure and Function* 221.9 (2016), pp. 4429–4449.

[235] John H Reynolds, Leonardo Chelazzi, and Robert Desimone. "Competitive mechanisms subserve attention in macaque areas V2 and V4". In: *Journal of Neuroscience* 19.5 (1999), pp. 1736–1753.

[236] Dario L Ringach. "Spatial structure and symmetry of simple-cell receptive fields in macaque primary visual cortex". In: *Journal of neurophysiology* (2002).

[237] Davi R Rodrigues et al. "A deeper look into natural sciences with physics-based and data-driven measures". In: *iScience* (2021), p. 102171.

[238] Alex Rodriguez and Alessandro Laio. "Clustering by fast search and find of density peaks". In: *Science* 344.6191 (2014), pp. 1492–1496.

[239] Edmund T Rolls and Martin J Tovee. "Processing speed in the cerebral cortex and the neurophysiology of visual masking". In: *Proceedings of the Royal Society of London. Series B: Biological Sciences* 257.1348 (1994), pp. 9–15.

[240] Edmund T Rolls, Martin J Tovee, and Stefano Panzeri. "The neurophysiology of backward visual masking: information analysis". In: *Journal of cognitive neuroscience* 11.3 (1999), pp. 300–311.

[241] Frank Rosenblatt. "Principles of Neurodynamics. Perceptrons and the Theory of Brain Mechanisms." In: (1961). DOI: 10.21236/AD0256582.

[242] Sebastien Royer et al. "Control of timing, rate and bursts of hippocampal place cells by dendritic and somatic inhibition". In: *Nature neuroscience* 15.5 (2012), pp. 769–775.

[243] Michelle Rudolph-Lilith, Mathieu Dubois, and Alain Destexhe. "Analytical integrate-and-fire neuron models with conductance-based dynamics and realistic postsynaptic potential time course for event-driven simulation strategies". en. In: *Neural Comput.* 24.6 (June 2012), pp. 1426–1461.

[244] Caroline A Runyan et al. "Distinct timescales of population coding across cortex". In: *Nature* 548.7665 (2017), pp. 92–96.

[245] Shuzo Sakata and Kenneth D Harris. "Laminar structure of spontaneous and sensory-evoked population activity in auditory cortex". In: *Neuron* 64.3 (2009), pp. 404–418.

[246] Maria V Sanchez-Vives and David A McCormick. "Cellular and network mechanisms of rhythmic recurrent activity in neocortex". In: *Nature neuroscience* 3.10 (2000), pp. 1027–1034.

[247] Takashi R Sato and Karel Svoboda. "The functional properties of barrel cortex neurons projecting to the primary motor cortex". In: *Journal of Neuroscience* 30.12 (2010), pp. 4256–4260.

[248] Takashi R Sato et al. "The functional microarchitecture of the mouse barrel cortex". In: *PLoS Biol* 5.7 (2007), e189.

[249] Marieke L Scholvinck et al. "Cortical state determines global variability and correlations in visual cortex". In: *Journal of Neuroscience* 35.1 (2015), pp. 170–178.

[250] Geun Seol et al. "Neuromodulators control the polarity of spike-timing-dependent synaptic plasticity". In: *Neuron* 55.6 (2007), pp. 919–929.

[251] Michael N Shadlen and William T Newsome. "The variable discharge of cortical neurons: implications for connectivity, computation, and information coding". In: *Journal of neuroscience* 18.10 (1998), pp. 3870–3896.

[252] Tatyana O Sharpee et al. "Adaptive filtering enhances information transmission in visual cortex". In: *Nature* 439.7079 (2006), pp. 936–942.

[253] David L Sheinberg and Nikos K Logothetis. "Noticing familiar objects in real world scenes: the role of temporal cortical neurons in natural vision". In: *Journal of Neuroscience* 21.4 (2001), pp. 1340–1350.

[254] Krishna V Shenoy and Jonathan C Kao. "Measurement, manipulation and modeling of brain-wide neural population dynamics". In: *Nature Communications* 12.1 (2021), pp. 1–5.

[255] Shy Shoham, Daniel H O'Connor, and Ronen Segev. "How silent is the brain: is there a 'dark matter' problem in neuroscience?" In: *Journal of Comparative Physiology A* 192.8 (2006), pp. 777–784.

[256] Donald Shoykhet, Daniel Doherty, and Michael J. Simons. "Coding of deflection velocity and amplitude by whisker primary afferent neurons: implications for higher level processing". In: *Somatosensory & motor research* 17.2 (2000), pp. 171–180.

[257] Roman Shusterman et al. "Sniff invariant odor coding". In: *Eneuro* 5.6 (2018).

[258] Daniel J Simons et al. "Responses of barrel cortex neurons in awake rats and effects of urethane anesthesia". In: *Experimental brain research* 91.2 (1992), pp. 259–272.

[259] W Singer. "Neuronal synchrony: a versatile code for the definition of relations?" In: *Neuron* 24 (1999), pp. 49–65.

[260] Anton Sirota et al. "Entrainment of neocortical neurons and gamma oscillations by the hippocampal theta rhythm". In: *Neuron* 60.4 (2008), pp. 683–697.

[261] PJ Sjostrom, GG Turrigiano, and Nelson SB Neuron. "Rate, timing, and cooperativity jointly determine cortical synaptic plasticity". In: *Neuron* 32 (2001), pp. 1149–1164.

[262] Samuel J Sober et al. "Millisecond spike timing codes for motor control". In: *Trends in neurosciences* 41.10 (2018), pp. 644–648.

[263] Narayan Srinivasa and Qin Jiang. "Stable learning of functional maps in self-organizing spiking neural networks with continuous synaptic plasticity." In: *Frontiers in computational neuroscience* 7.February (2013), p. 10. DOI: 10.3389/fncom.2013.00010.

[264] Kyle H Srivastava et al. "Motor control by precisely timed spike patterns". In: *Proceedings of the National Academy of Sciences* 114.5 (2017), pp. 1171–1176.

[265] Riccardo Storchi et al. *Comparison of latency and rate coding for the direction of whisker deflection in the subcortical somatosensory pathway*. 2012.

[266] Carsen Stringer, Michalis Michaelos, and Marius Pachitariu. "High precision coding in visual cortex". In: *BioRxiv* (2019), p. 679324.

[267] Carsen Stringer et al. "Spontaneous behaviors drive multidimensional, brain-wide activity". In: *Science* 364.6437 (2019).

[268] Claire Tang et al. "Millisecond-scale motor encoding in a cortical vocal area". In: *PLoS Biol* 12.12 (2014), e1002018.

[269] Simon Thorpe, Denis Fize, and Catherine Marlot. "Speed of processing in the human visual system". In: *nature* 381.6582 (1996), p. 520.

[270] David J Tolhurst, J Anthony Movshon, and Andrew F Dean. "The statistical reliability of signals in single neurons in cat and monkey visual cortex". In: *Vision research* 23.8 (1983), pp. 775–785.

[271] Ansgar Torvik. "Afferent connections to the sensory trigeminal nuclei, the nucleus of the solitary tract and adjacent structures. An experimental study in the rat". In: *Journal of Comparative Neurology* 106.1 (1956), pp. 51–141.

[272] Anne Treisman. "Features and objects: The fourteenth Bartlett memorial lecture". In: *The Quarterly Journal of Experimental Psychology Section A* 40.2 (1988), pp. 201–237.

[273] Anne Treisman. "How the deployment of attention determines what we see". In: *Visual cognition* 14.4-8 (2006), pp. 411–443.

[274] Anne Treisman and Hilary Schmidt. "Illusory conjunctions in the perception of objects". In: *Cognitive Psychol* 14.1 (1982), pp. 107–141. ISSN: 0010-0285. DOI: 10.1016/0010-0285(82)90006-8.

[275] Anne M Treisman and Garry Gelade. "A feature-integration theory of attention". In: *Cognitive Psychol* 12.1 (1980), pp. 97–136. ISSN: 0010-0285. DOI: 10.1016/0010-0285(80)90005-5.

[276] Misha V Tsodyks and Henry Markram. "The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability". In: *Proceedings of the national academy of sciences* 94.2 (1997), pp. 719–723.

[277] Nicholas L Turner et al. "Multiscale and multimodal reconstruction of cortical structure and function". In: *bioRxiv* (2020).

[278] Daniel Udvary et al. "A Theory for the Emergence of Neocortical Network Architecture". In: *bioRxiv* (2020).

[279] Nachum Ulanovsky, Liora Las, and Israel Nelken. "Processing of low-probability sounds by cortical neurons". In: *Nature neuroscience* 6.4 (2003), pp. 391–398.

[280] Leslie G Ungerleider and James V Haxby. "'What' and 'where' in the human brain". In: *Current opinion in neurobiology* 4.2 (1994), pp. 157–165.

[281] VJ Uzzell and EJ Chichilnisky. "Precision of spike trains in primate retinal ganglion cells". In: *Journal of Neurophysiology* 92.2 (2004), pp. 780–789.

[282] Martina Valente et al. “Correlations enhance the behavioral readout of neural population activity in association cortex”. In: *Nature neuroscience* 24.7 (2021), pp. 975–986.

[283] Rufin VanRullen and Simon J Thorpe. “Surfing a spike wave down the ventral stream”. In: *Vision research* 42.23 (2002), pp. 2593–2615.

[284] Siavash Vaziri et al. “A channel for 3D environmental shape in anterior inferotemporal cortex”. en. In: *Neuron* 84.1 (Oct. 2014), pp. 55–62.

[285] Julien Vitay, Helge U Dinkelbach, and Fred H Hamker. “ANNarchy: a code generation approach to neural simulations on parallel hardware”. en. In: *Front. Neuroinform.* 9 (July 2015), p. 19.

[286] Tim P Vogels and L F Abbott. “Signal propagation and logic gating in networks of integrate-and-fire neurons”. en. In: *J. Neurosci.* 25.46 (Nov. 2005), pp. 10786–10795.

[287] Christoph Von der Malsburg. “The what and why of binding: the modeler’s perspective”. In: *Neuron* 24.1 (1999), pp. 95–104.

[288] Moritz Von Heimendahl et al. “Neuronal activity in rat barrel cortex underlying texture discrimination”. In: *PLoS Biol* 5.11 (2007), e305.

[289] Saurabh Vyas et al. “Computation through neural population dynamics”. In: *Annual Review of Neuroscience* 43 (2020), pp. 249–275.

[290] Qi Wang, Roxanna M Webber, and Garrett B Stanley. “Thalamic synchrony and the adaptive gating of information flow to cortex”. In: *Nature neuroscience* 13.12 (2010), p. 1534.

[291] Jana de Wiljes, Lars Putzig, and Illia Horenko. “Discrete nonhomogeneous and nonstationary logistic and Markov regression models for spatiotemporal data with unresolved external influences”. In: *Communications in Applied Mathematics and Computational Science* 9.1 (2014), pp. 1–46.

[292] Francis R Willett et al. “High-performance brain-to-text communication via handwriting”. In: *Nature* 593.7858 (2021), pp. 249–254.

[293] Alex H Williams et al. “Discovering precise temporal patterns in large-scale neural recordings through robust and interpretable time warping”. In: *Neuron* 105.2 (2020), pp. 246–259.

[294] Adrien Wohrer, Mark D Humphries, and Christian K Machens. “Population-wide distributions of neural activity during perceptual decision-making”. In: *Progress in neurobiology* 103 (2013), pp. 156–193.

[295] Jason Wolfe et al. “Texture coding in the rat whisker system: slip-stick versus differential resonance”. In: *PLoS Biol* 6.8 (2008), e215.

[296] Jeremy M Wolfe. “Approaches to visual search: Feature integration theory and guided search”. In: *Oxford handbook of attention* (2014), pp. 11–55.

[297] Jeremy M Wolfe. “Guidance of visual search by preattentive information”. In: *Neurobiology of attention*. Elsevier, 2005, pp. 101–104.

[298] Jeremy M Wolfe, Kyle R Cave, and Susan L Franzel. "Guided search: an alternative to the feature integration model for visual search." In: *Journal of Experimental Psychology: Human perception and performance* 15.3 (1989), p. 419.

[299] Thilo Womelsdorf et al. "Modulation of neuronal interactions through neuronal synchronization". In: *science* 316.5831 (2007), pp. 1609–1612.

[300] Yukako Yamane et al. "A neural code for three-dimensional object shape in macaque inferotemporal cortex". en. In: *Nat. Neurosci.* 11.11 (Nov. 2008), pp. 1352–1360.

[301] Esin Yavuz, James Turner, and Thomas Nowotny. "GeNN: a code generation framework for accelerated brain simulations". en. In: *Sci. Rep.* 6 (Jan. 2016), p. 18854.

[302] Jianing Yu et al. "Recruitment of GABAergic interneurons in the barrel cortex during active tactile behavior". In: *Neuron* 104.2 (2019), pp. 412–427.

[303] Friedemann Zenke and Wulfram Gerstner. "Limits to high-speed simulations of spiking neural networks using general-purpose computers". en. In: *Front. Neuroinform.* 8 (Sept. 2014), p. 76.

[304] L I Zhang et al. "A critical window for cooperation and competition among developing retinotectal synapses". en. In: *Nature* 395.6697 (Sept. 1998), pp. 37–44.

[305] Davide Zoccolan, David D Cox, and James J DiCarlo. "Multiple object response normalization in monkey inferotemporal cortex". In: *Journal of Neuroscience* 25.36 (2005), pp. 8150–8164.

[306] Yanfang Zuo et al. "Complementary contributions of spike timing and spike rate to perceptual decisions in rat S1 and S2 cortex". In: *Current Biology* 25.3 (2015), pp. 357–363.