Characterizing patches of primary visual cortex with minimal bias

by

Martin A. Spacek

BSc Engineering Physics, University of Alberta, 2001

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL STUDIES
(Neuroscience)

The University of British Columbia
(Vancouver)

June 2015

Abstract

The brain is highly complex, and studying it requires simplifying experiments, analyses, and theories. New techniques can capture more of the brain's complexity while reducing biases in our understanding of how it works.

This thesis describes experiments in primary visual cortex of anesthetized cat, using high-density silicon multisite electrodes to simultaneously record from as many neurons as possible across all cortical layers, thereby characterizing local cortical populations with minimal bias. Recordings were maintained for many hours at a time, and included both spontaneous and stimulus-evoked periods, with a wide variety of naturalistic and artificial visual stimuli. A new "divide-and-conquer" spike sorting method translated correlated multisite voltages into action potentials of spatially localized, isolated neurons. This method tracked neurons over periods of many hours despite drift, and distinguished neurons with firing rates < 0.05 Hz.

Neuron physiology was reasonably normal and mostly agreed with accepted principles of visual cortex, but there were exceptions. Surprisingly, firing rates across the population followed a lognormal distribution, and 82% of neurons had mean firing rates < 1 Hz. Also surprisingly, orientation tuning strength across the population was inversely correlated with log firing rate. Finally, there was evidence for neural shift work: over long durations, as some neurons became silent, others became active. To break down analyses by cell type, neurons were classified by their temporal spike shape and receptive field.

Responses to repeated natural scene movie clips consisted of unique patterns of remarkably sparse, temporally precise (20 ms wide), reliable events. Mean pairwise correlations between neurons, as measured between trial-averaged responses to natural scene movies, were weakly positive. Correlations between simple and complex cells were lower — and between complex cells were higher — than expected, challenging the hierarchical model of complex cells. Cortical state was classified according to the local field potential, revealing greater natural scene movie response precision, sparseness, and reliability during the synchronized than desynchronized cortical state, contrary to reports in rodents.

The open-ended, inclusive, high-dimensional experiments and analyses described here make few assumptions, potentially leading to more insightful theories of brain function than hypothesis-driven research alone.

Preface

Some of the methods described in Sections 2.1–2.3 and Appendix C were published in Blanche et al. (2005) and Spacek et al. (2009), respectively. The clustering technique described in Section 3.9.1 and the overlap index described in Section 3.10.2 were devised by Nicholas Swindale and published in Swindale and Spacek (2014). Experiments on one animal (ptc15, Table 2.1) were done jointly with Tim Blanche, while the rest were done jointly with Nicholas Swindale. Experimental work was covered by UBC Ethics Certificates A04-0098 and A11-0280. The author declares no conflicts of interest.

Table of Contents

Abstract . . . . **ii**

Preface . . . . **iii**

Table of Contents . . . . **iv**

List of Tables . . . . **vii**

List of Figures . . . . **viii**

List of Abbreviations . . . . **x**

Note to Reader . . . . **xiii**

Acknowledgements . . . . **xiv**

1 Introduction **1**
- 1.1 Background 1
- 1.2 Summary 7

2 Experimental Methods **10**
- 2.1 General setup 10
- 2.2 Polytrodes 14
- 2.3 Data acquisition 14
- 2.4 Visual stimulation 17
- 2.5 Histology 21

3 Spike Sorting **22**
- 3.1 Background 22
  - 3.1.1 Existing methods 24
  - 3.1.2 Clustering methods 26
- 3.2 Overview 28
  - 3.2.1 Software 29
- 3.3 Preprocessing 29
- 3.4 Spike detection 30
- 3.5 Spatial localization 35
- 3.6 Initial channel split 36

3.7 Alignment 36
3.8 Dimension reduction 40
3.9 Clustering 45
  3.9.1 Gradient ascent clustering 45
3.10 Cluster verification 48
  3.10.1 Undersplitting 49
  3.10.2 Oversplitting 51
  3.10.3 Misassignment 53
  3.10.4 Duplicate spikes 54
3.11 Autocorrelograms & refractory periods 54
3.12 Simulation 55
3.13 Discussion 57
  3.13.1 Spike detection 57
  3.13.2 Dimension reduction 59
  3.13.3 Clustering 60
  3.13.4 Autocorrelograms & refractory periods 62
  3.13.5 Drift 62

4 Basic Physiology 63
4.1 Summary 63
4.2 Neuron yields 63
4.3 Firing rates 65
4.4 Templates & positions 69
4.5 Orientation tuning 72
4.6 Discussion 78
  4.6.1 Neuron yields 78
  4.6.2 Firing rates 81
  4.6.3 Templates & positions 84
  4.6.4 Orientation tuning 87

5 Cell Type 89
5.1 Introduction 89
5.2 Spike shape 91
5.3 Spatial extent 95
5.4 Receptive field type 95
5.5 Cell type comparison 101
5.6 Discussion 103
  5.6.1 Spike shape 103
  5.6.2 Spatial extent 106

5.6.3 Receptive field type 106
5.6.4 Cell type comparison 109

6 Natural scenes & cortical states 110
6.1 Introduction 110
6.2 Natural scene movie responses 111
6.3 Natural scene movie response correlations 115
6.4 Cortical states 121
6.5 UP/DOWN phases 126
6.6 Natural scene movie responses vs. cortical state 129
6.7 Discussion 137
6.7.1 Natural scene movie responses 137
6.7.2 Natural scene movie response correlations 141
6.7.3 Cortical states 144
6.7.4 UP/DOWN phases 146
6.7.5 Natural scene movie responses vs. cortical state 148
6.7.6 Clustering cortical states 150

7 Conclusion 152

Bibliography 156

Appendix A Receptive field stability 177

Appendix B Spike correlations 178

Appendix C Software 179
C.1 dimstim 179
C.2 spyke 181
C.3 neuropy 184

Appendix D Impedance meter 186

Appendix E ACSF recipe 188

List of Tables

Table 2.1 Cats used in this study 10
Table 2.2 Tracks recorded in this study 12
Table 2.3 54-channel polytrode designs 14
Table 2.4 Stimulus sets 20

Table 4.1 Neuron yield. 64

Table 5.1 Spike & receptive field type counts 95

Table 6.1 Counts of cells responsive to natural scene movies 113

Table E.1 ACSF ingredients 188

List of Figures

Figure 2.1 General experimental setup . . . 11
Figure 2.2 Silicon polytrodes . . . 15
Figure 2.3 50 ms of high-pass voltage waveform data . . . 16
Figure 2.4 Phase locking of V1 to a low refresh CRT monitor . . . 18
Figure 2.5 Example frames from natural scene movie sets . . . 20

Figure 3.1 Motivation for careful spike sorting . . . 23
Figure 3.2 Multichannel template matching . . . 25
Figure 3.3 Example spikes . . . 31
Figure 3.4 Spike detection . . . 33
Figure 3.5 Spike propagation in space . . . 35
Figure 3.6 Spatial locations & V_{pp} of spikes . . . 37
Figure 3.7 Spike misalignment can result in artifactual clusters . . . 39
Figure 3.8 Spike misalignment can result in non-Gaussian clusters . . . 40
Figure 3.9 Channel selection affects clusterability . . . 42
Figure 3.10 ICA separates clusters with large size ratios better than PCA . . . 44
Figure 3.11 GAC dependence on σ_c . . . 47
Figure 3.12 Plotting clusters in time is necessary to check for drift. . . . 51
Figure 3.13 Examining autocorrelograms for refractory period violations . . . 56

Figure 4.1 Mean firing rates had a lognormal distribution . . . 66
Figure 4.2 Coarse firing rates vs. time . . . 68
Figure 4.3 Single channel & multichannel templates . . . 70
Figure 4.4 Cell positions . . . 71
Figure 4.5 Spike parameters vs. time . . . 73
Figure 4.6 Smoothed spike parameters vs. time . . . 74
Figure 4.7 Distributions of changes in y position over time for all cells . . . 75
Figure 4.8 Orientation tuning curves . . . 77
Figure 4.9 Orientation tuning strength vs. mean firing rate . . . 78
Figure 4.10 Orientation preference, tuning strength & normalized depth . . . 80

Figure 5.1 A demonstration of the potential dangers of threshold-based metrics . . . 93
Figure 5.2 Temporal spike shape measures & classification . . . 94
Figure 5.3 Distributions of spatial extent . . . 96

Figure 5.4 Example drifting bar & grating raster plots of simple & complex cells 97
Figure 5.5 Spike triggered average to m-sequence 98
Figure 5.6 Spatiotemporal receptive fields 100
Figure 5.7 Spatial distribution of cell types 102
Figure 5.8 Comparison of receptive field type vs. spike type 103
Figure 5.9 Mean firing rate & activity duration distributions vs. cell type 104

Figure 6.1 Natural scene movie responses (1) 113
Figure 6.2 Venn diagram of active & responsive cells 113
Figure 6.3 Responsive cell counts vs. cell type 114
Figure 6.4 Natural scene movie responses (2) 117
Figure 6.5 Temporally precise response events can be shared by disparate neurons 118
Figure 6.6 Natural scene movie PSTH correlations 119
Figure 6.7 Matrices of natural scene movie mean PSTH correlations vs. cell type 120
Figure 6.8 Power spectral density and amplitude of deep layer LFP 122
Figure 6.9 Cortical state (1) 124
Figure 6.10 Cortical state (2) 126
Figure 6.11 Synchrony index distributions 127
Figure 6.12 Scatter plots of multiunit activity vs. synchrony index 128
Figure 6.13 UP & DOWN phases 130
Figure 6.14 Natural scene movie responses vs. cortical state (1) 131
Figure 6.15 Natural scene movie responses vs. cortical state (2) 132
Figure 6.16 Responsive inactive cells 134
Figure 6.17 Response precision, sparseness & reliability vs. cortical state 135
Figure 6.18 PSTH correlations vs. cortical state 138
Figure 6.19 PSTH correlation differences between cortical states 139
Figure 6.20 Complex cell models 144

Figure A.1 Long duration receptive field stability 177

Figure B.1 Spike train binary code 178

Figure C.1 Example dimstim script 180
Figure C.2 Visualizing time series voltage waveforms with spyke 182
Figure C.3 Sorting detected spikes with spyke 183
Figure C.4 GAC run time vs. σ_c & N 184
Figure C.5 Neuropy 185

Figure D.1 Impedance meter circuit 186

List of Abbreviations

1Dsep	one-dimensional separation metric
A1	primary auditory cortex
α-**BTX**	α-bungarotoxin
ACSF	artificial cerebrospinal fluid
AC	area centralis
ADC	analog-to-digital converter
AHP	afterhyperpolarization
CRF	classical receptive field
CRI	constant rate infusion
CRT	cathode ray tube
CSF	cerebrospinal fluid
CSD	current source density
dH_2O	deionized water
D_{JS}	Jensen-Shannon divergence
D_{KL}	Kullback-Leibler divergence
EEG	electroencephalogram
EM	expectation-maximization
FFT	fast Fourier transform
fMRI	functional magnetic resonance imaging
fps	frames per second
FWHM	full width half maximum
GAC	gradient ascent clustering
GMM	Gaussian mixture model

GUI	graphical user interface
HEPES	4-(2-hydroxyethyl)piperazine-1-ethanesulfonic acid ($C_8H_18N_2O_4S$)
HC	hierarchical clustering
IC	independent component
ICA	independent component analysis
IM	intramuscular
ISI	inter-spike interval
IP	intraperitoneal
IV	intravenous
LGN	lateral geniculate nucleus
LFP	local field potential
LM	Levenberg-Marquardt
M1	primary motor cortex
M2	secondary motor cortex
MEA	multichannel electrode array
MT	middle temporal cortex
MUA	multiunit activity
NDsep	N-dimensional separation metric
PC	principal component
PCA	principal component analysis
PFC	prefrontal cortex
PPT	pedunculopontine tegmenta
PSTH	peristimulus time histogram
PSD	power spectral density
PSP	postsynaptic potential

RAM	random-access memory
REM	rapid eye movement
RF	receptive field
RGC	retinal ganglion cell
RGP	rigid gas permeable
RMS	root mean square
RPV	refractory period violation
S1	primary somatosensory cortex
SI	synchrony index
SPC	superparamagnetic clustering
STA	spike-triggered average
STC	spike-triggered covariance
STRF	spatiotemporal receptive field
V1	primary visual cortex

Note to Reader

Given the inherent limitations of extracellular electrophysiology, the terms "cell" and "neuron" should be considered synonymous with "unit" throughout this thesis.

For Chapters 4–6, methods and results are combined within each main section, followed by a single discussion at the end of each chapter.

Chapter, section, appendix, figure, table, and equation references in this document are active hyperlinks, as are citations and abbreviations. The page numbers in the table of contents, lists of figures and tables, and bibliography are also hyperlinked. Chapters, sections, and appendices are bookmarked in the PDF.

A digital version of this document will be made available at http://mspacek.github.io.

Acknowledgements

My parents, for their constant support over all the years.

My supervisor, Nicholas Swindale, whose door has always been open, and for generously supporting me for so many years. Nick provided extensive training, as well as critical support during animal experiments. Most of the ideas, methods, and results in this thesis are a direct result of interacting with Nick.

Tim Blanche for introducing me to the use of polytrodes, and for setting up the recording system and writing the acquisition software "Surf".

Catalin Mitelut for his work on spike shape clustering, showing that the width of the second peak is more informative than that of the first peak, and for preliminary spike sorting of data briefly mentioned in Section 6.7.5.

My girlfriend, Annay, for putting up with me during thesis writing.

This thesis depended heavily on free, open source software including Linux, Python, Geany, Git, LaTeX, BibTeX, Inkscape, GIMP, and qpdfview. It also benefited from a steady stream of electronic music.

This work was supported by funding from the Canadian Institutes of Health Research (CIHR) and the Natural Sciences and Engineering Research Council of Canada (NSERC).

1 Introduction

> *"Brain cells fire in patterns."* — How the brain works, in 5 words or less.
> Steven Pinker, *The Colbert Report* (2007)

From insects to humans, brains are incredibly complicated computational machines. They may be the most complex systems in the known universe, and studying such systems is an exercise in managing complexity. For most of neurophysiological history, much of this complexity has been dealt with by, in essence, ignoring it. The technology to simultaneously monitor the activity of multiple neurons is relatively new, so by necessity, the interactions between neurons — perhaps the brain's most powerful feature — were long unexamined. Low density short duration extracellular recordings cannot distinguish between missing neurons and silent neurons, so again, partly by necessity, low firing rate neurons have mostly been ignored. High dimensional naturalistic stimuli are more complicated than low dimensional artificial stimuli, so the latter were, until recently, overwhelmingly used, and the responses to the former remain underappreciated. And finally, even when asleep or lightly anesthetized, the brain is likely performing a multitude of different simultaneous tasks. Yet when characterizing neural activity in a supposedly controlled way, over seemingly identical trials with identical stimuli or tasks, neural activity can vary substantially from one trial to the next. This variability is often labelled "noise", and is dealt with by averaging across trials. Although convenient for the neurophysiologist, doing so ignores the reality that the brain functions on a continuous basis. A brain's life is one continuous single trial, and this thesis is an attempt to treat it that way.

1.1 Background

Action potentials (spikes) are the currency of the brain. A spike allows one neuron to communicate via its axon to thousands of downstream neurons the precise moment at which its membrane potential threshold has been exceeded. Spikes are significant and energetically costly events (Attwell and Laughlin, 2001; Lennie, 2003). They can be recorded intracellularly with an electrode placed directly on or within the neuron's cell membrane, or extracellularly with an electrode placed outside but very near the membrane. Perhaps the first to isolate single spikes from single neurons (units) were Adrian and Bronk (1928), who dissected away the phrenic nerve in anesthetized rabbit until only one axon within the nerve remained intact. Removing all other possible transmitting axons within the nerve allowed for unambiguous extracellular recordings from single motoneuron axons using simple, relatively large brush electrodes placed on the nerve at either end of the dissection. This revealed that single-unit firing rates of motoneurons in the spinal cord control the force of diaphragm contraction critical to breathing. Hodgkin and Huxley (1939) later performed the first *in vitro* intracellular recordings (which guarantee single-unit isolation) using a blunt 100 μm wide wire electrode inserted longitudinally into 500 μm diameter squid giant axon. This led to their widely-used mathematical model of spike initiation and propagation (Hodgkin and Huxley, 1952).

However, if the goal is to record isolated spikes from neurons within the intact brain, removing all other units to gain single-unit resolution or limiting oneself to squid giant axon is not viable. Recording extracellularly but very close to the neuron's soma or axon hillock, where its spikes are generated, minimizes contamination by spikes from other nearby cells or axons. To further improve single-unit resolution, extracellular recordings can be filtered into high- and low-pass frequency bands. High-pass frequencies ($>$ 300 Hz) capture most of a spike's signal. Low-pass frequencies ($<$ 100 Hz), referred to as the local field potential (LFP), are the sum of synaptic (Lindén et al., 2011) or spiking (Buzsáki et al., 2012; Reimann et al., 2013) activity of the wider neural population. While not useful for single-unit isolation, the LFP can be used to monitor the overall activity of a region of cortex (Katzner et al., 2009; Kelly et al., 2010; Harris and Thiele, 2011).

Recording either intracellularly or extracellularly directly from intact neurons requires very fine micron-scale electrodes with high input impedance, such that they do minimal damage and do not perturb the very currents and potentials they measure. For extracellular electrodes, high impedance also rejects the noisy extracellular signal from the multitude of neurons located further away from the electrode. The first such microelectrodes were fine (25 μm tip) glass micropipettes, used by Renshaw et al. (1940) to extracellularly isolate single units in intact, anesthetized cat hippocampus. Ling and Gerard (1949) used even finer tapered ($<$ 1 μm tip) glass micropipettes intracellularly. These were fine enough for the cell membrane to seal around them after insertion, minimizing damage.

Glass micropipettes have the advantage of being easily shaped under heat and tension (pulling) to form very fine tips, but they are brittle and delicate. Wire electrodes are more durable, but were difficult to sharpen. Hubel (1957) developed a way of electrolytically sharpening the tips of insulated tungsten wire down to $<$ 0.5 μm, and Green (1958) did the same with stainless steel wire. Sharp wire electrodes allowed for larger scale single-unit investigations, revealing retinotopically organized, localized, oriented, direction selective receptive fields (RFs) of single units in cat and monkey primary visual cortex (V1), as well as the existence of visual stimulus feature maps across the surface of V1 (Hubel and Wiesel, 1959, 1962, 1968).

Hubel and Wiesel also found that single-unit RFs in V1 could be categorized into two types: simple and complex (Hubel and Wiesel, 1962). Simple cells have oriented RFs with alternating, spatially distinct ON and OFF subregions. This gives them spatial phase preference for light and dark oriented stimuli. Complex cells have spatially overlapping subregions, and are therefore invariant to spatial phase, but they still maintain orientation preference, making them slightly more general than simple cells. To explain this, Hubel and Wiesel proposed a hierarchical model in which complex cells pool from many different simple cells with similar orientation preference but different spatial phase preferences, resulting in spatial phase invariance (Figure 6.20A). This has been the dominant model of complex cells ever since.

The next major advances in recording technology were stereotrodes and tetrodes, which consist of twisted pairs or tetrads of insulated sharp wire electrodes, respectively, with tips spaced about 25 μm apart (McNaughton et al., 1983; O'Keefe and Recce, 1993; Wilson and McNaughton, 1993). Increasing the number of closely spaced electrodes improves single-unit isolation by sampling the extracellular potential at multiple points in space, thereby increasing the chance that spike-induced extracellular potentials from even the most closely spaced neurons will appear sufficiently different on at least one electrode channel to separate their (multichannel) spike waveforms (Gray et al., 1995; Harris et al., 2000). However, this benefit comes at the cost of added complexity. Detecting extracellular spikes above noise and separating spikes from multiple neurons (spike sorting) is challenging enough. Doing so for many closely spaced electrode sites with correlated voltages is even more difficult.

Contemporary use of tetrodes often combines many of them into a circular or conical array, such that they all point toward the same region of neural tissue (Wilson and McNaughton, 1993; Jog et al., 2002; Battaglia et al., 2009; Voigts et al., 2013). The depth of each tetrode is individually controlled via independent screw drives, each of which is adjusted to maximize its single-unit yield. Tetrode arrays increase yield by simultaneously recording from many more neurons than a single tetrode. However, tetrodes in an array rarely, if ever, come close enough to one another to isolate spikes from the same or immediately neighbouring neurons, and their relative positions are difficult to infer. Therefore, tetrode arrays cannot fully sample a contiguous volume of brain. Moreover, because neurons that fire at a high rate are easier to isolate than those that fire at a low rate, and due to the desire to maximize neuron yield, arrays of individually adjustable tetrodes may be biased towards high rate neurons, and may neglect the majority of neurons which fire at very low average rates (Lennie, 2003; Olshausen and Field, 2005; Shoham et al., 2006).

A more recent advance in extracellular single-unit electrophysiology is high-density silicon multisite electrodes (polytrodes) (Drake et al., 1988; Henze et al., 2000; Harris et al., 2000; Csicsvari et al., 2003; Blanche et al., 2005; Berényi et al., 2014). These are arrays of ~ 15 μm diameter metal (iridium, platinum or gold) electrode sites on silicon shanks, manufactured using micron-scale deposition and etching technology adapted from the integrated circuits industry. The silicon shanks are often thin (~ 15 μm) to minimize damage, but their widths and lengths can range widely (20–200 μm & 500–6000 μm respectively). Sites are arranged in a 1D, 2D, or even 3D (Du et al., 2009) array in almost any conceivable spatial pattern. Site spacing can be < 20 μm, but is more commonly > 30 μm. Both single- and multi-shank designs exist with 1–64 electrode sites per shank, and the potential for many more.

One major benefit of polytrodes over tetrode arrays is that, like tetrodes, the electrode site density is high, but unlike tetrode arrays, polytrodes can maintain consistently high site density over many hundreds of microns, even millimeters. Another related benefit of polytrodes is that the precise relative location of each electrode site is fixed and known. As a result, correlated spike signals from neighbouring electrode sites can be used to help localize the isolated neuron, and the relative locations of all neurons isolated by the polytrode can therefore be inferred (Henze et al., 2000; Csicsvari et al., 2003; Blanche et al., 2005; Blanche, 2005; Du et al., 2009). The cost of having a fixed array of electrode sites is that neuron yield cannot be maximized by individually positioning each site, or group of sites, as it can with tetrode arrays. Only the entire array can be moved up or down. However, this cost can be considered a benefit because it may reduce bias towards high firing rate neurons.

All of the above recording methods involve inserting a device into the brain, thereby causing at least some micro-scale tissue damage. The most recent advance in large-scale recording of neural populations, two-photon calcium imaging (Denk et al., 1990; Svoboda et al., 1997; Helmchen and Denk, 2005; Kerr et al., 2005; Ohki et al., 2005, 2006), provides an alternative in which the brain remains untouched and therefore undamaged. Neurons are either loaded with a photosensitive calcium indicator, or are genetically engineered to express it endogenously (Wallace et al., 2008; Lütcke et al., 2010). When stimulated with two-photon laser scanning, the calcium indicator fluoresces as a function of intracellular calcium concentration, which itself is a function of spiking activity, allowing for large-scale 3D optical readout of spiking activity of a complete neural population. The laser scanning used for readout can image a volume of superficial cortex $\sim 300\ \mu$m on a side (Ohki et al., 2006), and several hundred microns deep. Imaging depth has steadily improved over time, and has now reached as much as $800\ \mu$m (Mank et al., 2008; Mittmann et al., 2011). Due to differences in cortical thickness and transparency, two-photon calcium imaging works best in rodents or immature cats, and is perhaps not as well suited for adult cats whose primary visual cortex is $\sim 1500\ \mu$m thick (Beaulieu and Colonnier, 1985; Jones and Peters, 1987; Payne and Peters, 2001). More importantly, intracellular calcium signals are slow, with decay time constants on the order of dozens or hundreds of milliseconds (Koester and Sakmann, 2000; Stosiek et al., 2003; Kerr et al., 2005; Wallace et al., 2008; Sasaki et al., 2008), making detection of precise spike times in arbitrary spike patterns (including bursts) difficult, if not impossible. Finally, the equipment required for two-photon calcium imaging (laser, scanning mechanism, optics, detectors) is currently much more expensive than the equipment required for polytrode recordings.

Why is recording precise spike times so important? A long-standing hypothesis in neuroscience is Hebb's cell assembly hypothesis (Hebb, 1949). A cell assembly is a spatially distributed set of neurons whose transient, temporal pattern of activation might serve as the neural substrate for stimuli, memories, concepts, decision making, and motor actions. The appeal of cell assemblies lies in two main attributes. First, their spatial distribution means that they degrade gracefully upon damage to a specific location in the brain. Second, since a neuron can participate in one of many different cell assemblies, the combinatorics allow for nearly unlimited representational capacity, despite a limited number of neurons. Definitively proving the cell assembly hypothesis will require simultaneously recording precise spike times of many more neurons than is currently possible.

The appeal of cell assemblies is closely related to the appeal of spike timing as a coding mechanism. Precise spatiotemporal spike patterns in neuronal recordings could be signatures of cell assemblies at work (Hopfield, 1995; Engel et al., 2001; Harris, 2005). Furthermore, the psychophysically measured speed of object categorization in primates and the conduction velocity of spikes from one cortical area to the next put a hard upper limit on the number of spikes ($\sim$ 1–2) that neurons in each area have time to send to the next area (Thorpe et al., 1996; Kirchner and Thorpe, 2006). VanRullen and Thorpe (2002) go on to argue that this is insufficient for downstream areas to calculate the firing rates of their inputs, and that instead coding must be based on spike arrival time. In addition, sparse coding theory (Olshausen and Field, 1996, 2004) suggests that for various reasons including energy efficiency and representational capacity, only a small fraction of neurons are active at any time, and that mean firing rates are very low. Sparse coding therefore implicitly favours spike time coding over spike rate coding.

Experimental evidence in support of precise spike timing is accumulating (Burr and Ross, 1979; Gray et al., 1989; Mainen and Sejnowski, 1995; Vaadia et al., 1995; Bair and Koch, 1996; Harris et al., 2002, 2003; Salami et al., 2003; Wehr and Zador, 2003; Ikegaya et al., 2004; Johansson and Birznieks, 2004; Luczak et al., 2007; Gollisch and Meister, 2008; Havenith et al., 2011; Mackevicius et al., 2012; Luczak et al., 2013; Garcia-Lazaro et al., 2013), but the debate over spike rate coding vs. spike time coding is far from settled (Softky, 1995; Shadlen and Newsome, 1998; Oram et al., 1999; Baker and Lemon, 2000; Latham et al., 2006; London et al., 2010), making this a fertile area of study. Both types of coding are likely used for different purposes, and there is no hard line separating the two. The debate is mostly about just how fine a temporal scale of integration the brain might use for coding.

Why focus on extracellular and not intracellular electrophysiology for large scale recordings? The benefit of intracellular recording is that it directly measures cell membrane potential, and allows not only detection of spikes with high signal to noise ratio and no risk of spike contamination from neighbouring cells, but also detection of excitatory and inhibitory postsynaptic potentials (PSPs). But since intracellular electrodes touch or pierce the cell membrane and are very sensitive to motion, successfully recording ("holding") one cell intracellularly *in vivo* for hours at a time is challenging. Holding multiple cells simultaneously is that many times more difficult. Ko et al. (2011) recorded intracellularly from 4 cells and Perin et al. (2011) from 12 cells simultaneously, yet these impressive numbers are much lower than for large-scale extracellular recordings, which currently record from dozens or hundreds of neurons at once (Csicsvari et al., 2003; Segev et al., 2004; Buzsáki, 2004; Blanche et al., 2005; Berényi et al., 2014). Moreover, breaching the cell membrane can alter the cell's electrophysiology, and ultimately reduce its lifespan. Therefore, large scale single-unit multielectrode recordings have been mostly limited to extracellular recordings.

Why perform acute recordings in anesthetized animals instead of chronic recordings in awake animals? Acute recordings are generally simpler and require fewer resources than chronic recordings. The same holds for anesthetized vs. awake animals. Acute experiments face no ethical concerns about subject pain upon waking, and have less need for stringent sterile techniques during surgery.

In acute recordings, anesthesia is required for two reasons: ethical, to eliminate pain and stress; and practical, to keep the animal still during recordings. The downside of anesthesia is that brain activity may not be as naturalistic as in an awake head-restrained or freely behaving animal, and may therefore lead to biased data and conclusions. Although it is possible that precise neuronal spike timing may be disrupted during anesthesia (Swindale, 2003), V1 is an early cortical area and may not be as affected by anesthesia as higher level areas. A study by Vincent et al. (2007) found that, at least at the spatial and temporal resolution of functional magnetic resonance imaging (fMRI), patterns of spontaneous activity in monkey cortex, including V1, were independent of depth of anesthesia. Furthermore, there is strong evidence that cortical state (typically measured by LFP frequency content) can vary considerably during both anesthesia and wakefulness (Arieli et al., 1996; Petersen et al., 2003; Fiser et al., 2004; Harris and Thiele, 2011; Marguet and Harris, 2011; Sakata and Harris, 2012; Xu et al., 2012). Considering cortical state may therefore be even more important than considering whether the animal is awake or anesthetized.

There are at least two types of cortical state (Berger, 1929; Harris and Thiele, 2011). The first is the synchronized state, in which large groups of neurons tend to fire in synchrony, resulting in large low-frequency extracellular fluctuations as measured by the electroencephalogram (EEG) or LFP. The second is the desynchronized state, in which neurons fire more independently, and the resulting extracellular fluctuations are lower in amplitude and higher in frequency. The synchronized state occurs during deep anesthesia, slow-wave sleep, and quiet wakefulness, while the desynchronized state occurs during light anesthesia, rapid eye movement (REM) sleep, and alert wakefulness. Because both cortical states can occur during both wakefulness and anesthesia, controlling for cortical state may allow results from anesthetized animals to be directly applicable to those from awake animals (Arieli et al., 1996; Petersen et al., 2003; Fiser et al., 2004; Harris and Thiele, 2011).

Why study V1, and why use cats? Early sensory areas are the information gateways to the brain. This may make them easier to study and understand than higher-level multimodal and cognitive areas. Vision is the richest and most complex sensory modality, with over 100 million photoreceptors alone in human retina dedicated to detecting incident rays of light (Kandel et al., 2012). Despite decades of work, vision research is full of unsolved computational problems. Replicating the brain's natural visual capabilities in machines has proven to be very difficult, although progress is being made (Lowe, 2004; Krizhevsky et al., 2012). Cats have a long history of study in vision. They have good acuity, and as predators have forward facing eyes like humans. Cats are also widely available. However, long-duration large scale recordings combined with a rich set of visual stimuli can both help to reduce the number of animals needed.

V1 has mostly been investigated using simple artificial stimuli, such as drifting and flashed spots, bars and gratings, as well as spatiotemporal white noise, all of which help to characterize the RFs of neurons in V1. Each neuron has its own combination of stimulus preferences for retinotopic position, orientation, spatial and temporal frequency, and even direction of motion. These preferences are quantified using tuning curves, which plot trial-averaged firing rate as a function of stimulus parameter. Although convenient for characterizing such preferences, these artificial stimuli are not behaviourally relevant. The brain did not evolve in an environment of bars and gratings. It evolved in a 3D environment full of physical objects with texture, occlusion, perspective, lighting, shadow, motion, and parallax. The statistics of natural scene movies are very different from those of artificial stimuli (Olshausen and Field, 1996, 2000). Furthermore, natural vision involves full field stimulation, and is not restricted to a small portion of visual space as is often the case with artificial stimuli.

It should therefore not come as a great surprise that the spiking responses of neurons in V1 to full field naturalistic stimuli are very different from responses to spatially localized artificial stimuli (Vinje and Gallant, 2000, 2002; Yen et al., 2007; Haider et al., 2010; Herikstad et al., 2011), or that response models based on artificial stimuli do a poor job of predicting responses to naturalistic stimuli (Olshausen and Field, 2005; Carandini et al., 2005). Specifically, responses to naturalistic stimuli are sparser and more temporally precise and reliable than responses to artificial stimuli. These are still recent revelations that remain underappreciated.

1.2 Summary

This thesis characterizes local neuronal populations in anesthetized cat V1 in a wider manner than usual. Extensive efforts were made to minimize bias at every step of the way, from stimulation, recording and spike sorting, to response characterization, cell typing, and cortical states.

Chapter 2 describes the experimental methodology. Various strategies were used to help minimize bias during recording. Extracellular waveform data were recorded using single shank polytrodes spanning most or all layers of cat visual cortex. Compared to wire or tetrode arrays with individually positionable recording elements, the fixed relative electrode site positions of polytrodes may reduce potential biases towards specific cell types, including high firing rate cells. Recording durations were many hours long in a given position (track). For each track, a host of both artificial and naturalistic stimuli were used, in addition to spontaneous recordings. For artificial stimuli with tunable parameters, a full range of stimulus parameters (such as orientation) were used to minimize bias. Both artificial and natural scene movie stimuli spanned as much of the visual field as possible, well outside the classical RFs of all recorded neurons.

Chapter 3 tackles the methodological challenge of high-density multichannel spike sorting. Before performing any further analysis, spikes must be detected and sorted from the raw extracellular voltage waveform data acquired during the experiment. Spike sorting can lead to a range of biases, such as exclusion of low firing rate cells, or incorrect grouping together of low and high firing rate cells. Most critically, like tetrodes, polytrodes with closely spaced electrode sites present a challenge for spike sorting in that spikes tend to appear on several channels at a time. Yet, unlike tetrodes, polytrodes sample a long and potentially contiguous volume of brain. Existing spike sorting methods cannot deal with all of these challenges, so a novel spike sorting technique was developed. Inspired by existing methods, this “divide-and-conquer” spike sorting method addresses these challenges (Swindale and Spacek, 2014), and is implemented in freely available software. This method sorted millions of spikes from up to 93 simultaneously isolated units. Units were tracked over many hours despite some of them exhibiting significant drift and very low mean firing rates (< 0.05 Hz).

Data were acquired from a total of 15 tracks in 10 hemispheres in 6 cats. Due to time and space constraints, only 3 tracks were fully spike sorted (Section 4.2), resulting in 245 single units which were analyzed in Chapters 4–6.

Chapter 4 describes the basic physiology of these units. Across all spontaneous and stimulus-evoked recording periods, mean firing rates were surprisingly low. Mean firing rates were not normally distributed, but rather lognormally, with a geometric mean of 0.11 Hz. 82% of neurons had mean firing rates below 1 Hz. These results support the theory of sparse coding (Olshausen and Field, 1996, 2004). Evidence is also presented that neuronal populations in V1 perform a kind of shift work. Over the course of minutes and hours, as some cells stopped firing, others started firing, keeping the geometric mean firing rate reasonably constant over time. Shift work could be a useful mechanism for maintaining cell physiology, or even network stability. Neurons were localized relative to the polytrode by their multichannel waveforms. Overall, they were well distributed along its length, but at a finer scale their positions were biased towards the positions of the electrode sites, suggesting that the 65 μm hexagonal electrode site spacing of the polytrodes used here is not dense enough to fully capture the local neural population. Orientation tuning curves were calculated for each cell, and 61% of all cells were significantly tuned, but 87% were tuned if only active neurons (firing rates ≥ 0.05 Hz) were considered. Surprisingly, orientation tuning strength was significantly but inversely correlated with log firing rate. This is incongruent with the notion that higher firing rates result in better stimulus encoding, and further strengthens the relevance of low firing rate cells. The issue of polytrode tissue damage is discussed, and the use of narrower polytrodes is suggested to reduce damage.

Chapter 5 examines the classification of cells according to temporal spike shape, multichannel spatial extent, and RF. As many as 4 different types of spike shapes, and 4 different types of RFs were found. Cells were not found to cluster according to the 2D spatial extent of their multichannel templates. Spike shape and RF types were compared by spatial location and mean firing rate. These classifications were used later in Chapter 6.

Chapter 6 examines responses to natural scene movies, and the dependence of those responses on cortical state as measured from the LFP. Short (5 s) natural scene movie clips were presented hundreds of times each. Trial-aligned spike raster plots had short, temporally precise and reliable response events consisting of as little as one spike each. These resulted in peristimulus time histogram (PSTH) peaks as little as 20 ms wide, showing that responses in V1 are more temporally precise and reliable than widely assumed. Moreover, very low stimulus-evoked firing rates did not preclude cells from having temporally precise and reliable response events. Response event precision and reliability varied from one cell to the next, and there was great diversity in the pattern of response events between cells and across movies. Response correlations between all cell pairs were measured at a fine (20 ms) time scale, and were mostly very weak, even between simple and complex cells. However, between complex cells, response correlations were stronger than expected, challenging the hierarchical model of complex cells and providing evidence for an alternative recurrent model. Cortical state switched spontaneously between synchronized and desynchronized. Stimulus-evoked superficial cell firing rates were higher in the desynchronized than synchronized state, while the reverse was true for spontaneous activity. Deep layer cells showed a more heterogenous relationship. Contrary to other reports, response event precision, sparseness, and reliability for almost all neurons was higher during the synchronized than desynchronized state. Finally, cortical state had a greater influence on response correlations than did the particular movie clip presented, demonstrating the importance of taking cortical state into account.

By minimizing bias during stimulation, recording, spike sorting, and analysis of multiple simultaneously recorded single units from patches of V1, the conclusions from many different studies, each of which may focus on only a subset of layers or cell or stimulus types, can be replicated or challenged. This manner of experimentation is less hypothesis-driven than most neurophysiology experiments, which might be a weakness in the short term. But in the long term, collecting data in such a hypothesis- and analysis-agnostic way is a strength. It allows analyses to be done in a richer, higher dimensional space, and enables examination of neural responses and properties in various combinations (subspaces) that more hypothesis-driven experiments might not be capable of. Besides adding flexibility, such an approach may also help explain apparent neural response and property variability as an artifact of inappropriate pooling of results over many of these dimensions. This approach therefore has the potential to reveal greater reliability of neural responses and properties, and to show that the brain is less noisy than widely believed.

2 Experimental Methods

2.1 General setup

Extracellular recordings were made from cortical areas 17 and 18 of anesthetized domestic cat. The general experimental setup is shown in Figure 2.1. Animal experiments followed the guidelines of the Canadian Council for Animal Care and the Animal Care Committee of the University of British Columbia. Six cats were used in total: four were normal domestic cats, while two were heterozygous lipoprotein lipase deficient, left over from an unrelated study (Table 2.1). Initial stages of each animal experiment were performed with the aid of a veterinarian. For 3 of the 6 cats (ptc20–22), initial sedation was by intramuscular (IM) injection of dexmedetomidine (25 μg/kg) and initial analgesia by IM injection of butorphanol (0.3 mg/kg). The other 3 cats may have been sedated using different drugs as seen fit by the veterinarian. An intravenous (IV) catheter was inserted, and initial anesthesia was induced by IV injection of sodium thiopental or propofol. An endotracheal tube was then inserted and a catheter placed in the urethra. The animal was placed in a stereotaxic frame and its head fixed in place with ear bars coated in topical anesthetic (5% lidocaine). The stereotaxic frame was mounted on an air table which was floated prior to polytrode insertion to minimize vibrations.

For 4 cats, anesthesia was maintained by inhalation of 0.5–1.5% isoflurane with 70% N_2O in O_2. During surgical procedures and euthanization, up to 3% isoflurane was used. For the 2 other cats, constant rate infusion (CRI) of propofol (5–10 mg/kg/h) and fentanyl (5–7 μg/kg/h) was used instead of isoflurane and N_2O. In two cats, the opioid buprenorphine (0.01 mg/kg) was injected subcutaneously every 12 hours as an analgesic. In one cat (ptc18), xylazine (2 mg/kg) was injected

Animal ID	Sex	Age (y)	Weight (kg)	Source	Stimulus set	Drugs		
						Anesthetic	Paralytic	Other
ptc15	F	2	2.5	UV	A	iso + N_2O	PB	
ptc17	M	7	7.0	HLLD	B	iso + N_2O	α-BTX	bupr
ptc18	M	5	5.1	HLLD	B	iso + N_2O	α-BTX	bupr, atrop, xyla
ptc20	F	1	3.4	UCD	C	prop + fent	PB	bupivicaine
ptc21	F	1	3.5	UCD	C	prop + fent	PB	
ptc22	F	1	3.2	UCD	C	iso + N_2O	PB	bupr, dobut

Table 2.1: Cats used in this study. Stimulus sets are summarized in Table 2.4. Drugs mentioned in the text but not listed here were administered to all cats. **UV**: Unique Ventures, Balmoral, MB. **HLLD**: heterozygous lipoprotein lipase deficient cats, UC Davis. **UCD**: UC Davis. **iso**: isoflurane. **prop**: propofol. **fent**: fentanyl. **PB**: pancuronium bromide. **α-BTX**: α-bungarotoxin. **bupr**: buprenorphine. **atrop**: atropine. **xyla**: xylazine. **dobut**: dobutamine.

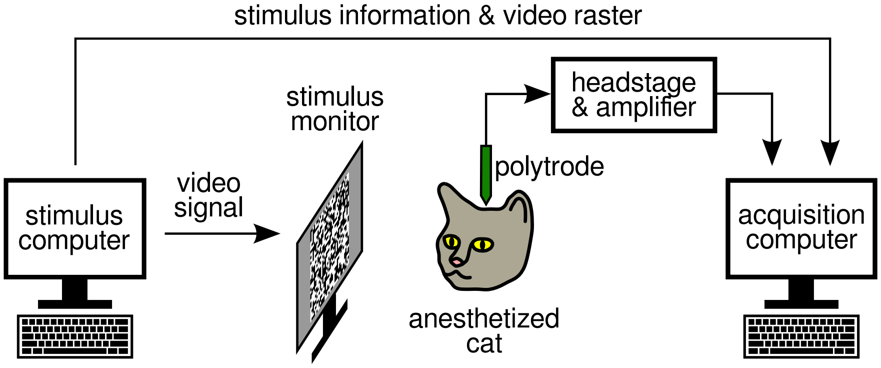

Figure 2.1: General experimental setup. The stimulus computer was the master and the acquisition computer was the slave. To begin recording, the user first armed the acquisition computer to save both spike and stimulus data to disk. Then the user ran the desired stimulus on the stimulus computer, which subsequently triggered the acquisition computer to begin saving to disk. Stimulus information, represented by a 16 bit integer, was acquired on every (5 ms) raster of the stimulus monitor. Adapted from Spacek et al. (2009).

IM as a preanesthetic, and atropine (0.02 mg/kg) was injected IM as needed to decrease salivation and increase heart rate.

Dobutamine (0.25 mg initial bolus, then CRI of 0.15–0.3 mg/kg/h) was administered IV in one cat (ptc22) to increase blood pressure. The antibiotic bupivicaine was injected subcutaneously (0.6 mL) around the scalp wound in one cat (ptc20). Table 2.1 summarizes the drugs administered uniquely for each cat.

Animals were mechanically ventilated (Harvard Apparatus, Holliston, MA) at ~ 20 breaths/min to maintain end-tidal CO_2 of 30–40 mmHg. This was especially necessary during systemic paralysis (see below). Blink and pinna (ear) reflexes and toe pinch were used to ensure sufficient anesthetic depth. Dexamethasone (1 mg/kg) was injected IM to reduce swelling and salivation. The cat was kept hydrated on a mixture of lactated Ringer's salt solution (10–20 mL/h), sometimes with added potassium chloride (20 mEq/L) and dextrose (2.5%). Heart rate and blood oxygenation were monitored with a pulse-oximeter (Nonin 8600V), with the sensor placed on the tongue or a shaved part of the tail. Mean arterial blood pressure was monitored with a doppler blood pressure monitor (Parks Medical 811-B) on a shaved section of hind leg. End-tidal CO_2 and respiration rate was monitored with a capnographer (Hewlett-Packard HP47210A) on the gas exhaust close to the endotracheal tube. Body temperature was monitored with a rectal probe and maintained via closed-loop control with a homeothermic blanket (Harvard Apparatus). All vital signs were logged during the course of each experiment.

Local anesthetic (bupivacaine) was injected subcutaneously around the top of the skull and into the ear muscles before cutting the skin to expose the skull. A roughly 4 × 6 mm craniotomy (1–5 mm lateral and 3–9 mm posterior relative to the centerline and earbar zero, respectively)

Track ID	Duration (hours)	Position			Angle		Area	Polytrode Design	Agar
		L	P	D	ML	AP			
ptc15.tr7c	17.6	3	5	700	0	20	17	2a	4% in saline
ptc17.tr1	17.1	-3	6	0	10	10	17	1a	4% LTA
ptc17.tr2b	12	2	4	300	10	10	17	1a	4% LTA
ptc18.tr1	14.9	-3	6	0	10	10	17	2b	3% LTA
ptc18.tr2c	13.9	3	5	600	15	15	17	2b	3% LTA
ptc20.tr1	15.7	-3	4	0	10	10	17	2b	2.5% LTA
ptc20.tr2	14.2	3	3	0	10	10	18	2b	2.5% LTA
ptc20.tr3	5.1	3	4	0	10	10	17	2a	2.5% LTA
ptc21.tr2	7.5	-3.5	6	0	5	10	17	2a	2.5% LTA
ptc21.tr5c	9.5	-4	5.5	400	5	10	18	2a	2.5% LTA
ptc21.tr6b	6.8	-3	3.5	150	2	10	18	2a	2.5% LTA
ptc22.tr1	9.5	-3.5	5.5	0	0	20	17	1a	2.5% LTA
ptc22.tr2	7.3	3	4	0	0	20	17	1a	2.5% LTA
ptc22.tr4b	8.5	-2.5	4.5	150	0	20	17	1a	2.5% LTA
ptc22.tr5b	6.2	-2	3.5	150	0	16	17	1a	2.5% LTA

Table 2.2: Tracks recorded in this study. Tracks fully sorted and analyzed are highlighted in grey. The first part of each track ID is the associated animal ID from Table 2.1. The second part is the track number within that animal, with a letter suffix to distinguish different depths within the same insertion. **Duration**: track duration, i.e., the total amount of time spent in each track, including any gaps in recording. **L**: lateral distance (mm) from midline. Negative values denote left hemisphere, positive values denote right hemisphere. **P**: posterior distance (mm) from ear bar zero. **D**: depth (μm), a rough estimate of how far, along the path of the track, the topmost sites of the polytrode were beneath the cortical surface. **ML**: mediolateral angle (°): lateral tilt of polytrode from vertical. **AP**: anteroposterior angle (°): posterior tilt of polytrode from stereotaxic vertical. **Area**: Brodmann's area (not definitive) as defined solely by stereotaxic coordinates. Polytrode designs are described in Table 2.3. **LTA**: low-temperature high purity agarose in ACSF.

was drilled with a dental drill (Midwest Stylus, DENTSPLY Professional, Des Plaines, IL) over Brodmann's area 17 and 18. A stereo surgical microscope was used during drilling, removal of meninges, and polytrode insertion. Artificial cerebrospinal fluid (ACSF) (Appendix E) was used to flush away blood and other detritus from the meninges, and to keep them moist. Ophthalmic surgical sponges (Ultracell Eye Spears, Aspen Surgical, Caledonia, MI) were used to wick blood and excess fluid away. Care was taken to not apply pressure to the brain. Sometimes the dura mater was completely reflected to expose the underlying pia mater, but more often only a small area of dura was dissected away one layer at a time with an ophthalmic slit knife (Beaver Optimum 15°, BD Medical, Le Pont-de-Claix, France; or ClearCut 3.2 mm, Alcon, Mississauga, ON). A small nick in the pia was then made with the ophthalmic slit knife to allow for polytrode insertion. Prior to insertion, cerebrospinal fluid (CSF) was wicked away from the point of insertion using an ophthalmic surgical sponge. This seemed to improve unit isolation, perhaps because the CSF acts as a low impedance path to ground. Immediately before or after insertion, high purity low temperature agarose (Type III-A, Sigma-Aldrich, St. Louis, MO) dissolved in ACSF at a concentration of 2.5–4% was applied in liquid form at 38–40°C to the craniotomy. This quickly set and greatly reduced or eliminated brain motion due to respiration and heart beats. Sometimes, if the dura and pia were insufficiently dissected away, the polytrode would cause the brain to dimple during insertion before finally breaking through. Applying agar first helped reduce dimpling, and thereby increased the chance of a successful insertion. The polytrode was advanced through the tissue using a manual micromanipulator (Model 1460 Electrode Manipulator, David Kopf Instruments, Tujunga, CA) under visual control until the topmost electrode sites disappeared below the surface of the cortex. Any subsequent advancement through the tissue was made with a hydraulic micromanipulator (Narishige MHW-4, East Meadow, NY), typically 150–300 μm at a time.

Nictitating membranes were retracted with phenylephrine (10%, 1–2 drops/eye), and pupils were dilated with tropicamide (0.5%, 1–2 drops/eye). Custom-made rigid gas permeable (RGP) contact lenses (14 mm diameter, 7.8–8.7 mm base curvature, +2.00 to +4.00 diopter, Harbour City Contact Lens Service, Nanaimo, BC) protected the eyes and refracted the cat's vision to the distance of the stimulus display monitor. To improve focus, 3 mm diameter artificial pupils were placed directly in front of the lenses. To prevent eye drift, the cat was given an initial IV bolus of the systemic paralytic pancuronium bromide (1 mg/kg), and paralysis was maintained by CRI (0.2 mg/kg/h). For some animals, pancuronium bromide was not used (Table 2.1). Instead, α-bungarotoxin (α-BTX) was injected retrobulbarly (125 μM, 0.5 mL per eye), acting as a local paralytic to prevent eye drift. Especially for α-BTX animals, eye position was closely monitored by reverse ophthalmoscopy (Section 2.4) to ensure stability, using fine blood vessels as landmarks. α-BTX injections were found to be effective at preventing the eyes from drifting.

Experiments lasted up to 3 days. At the end of the experiment, the anesthetic level was increased and the cat was killed with an IV injection of euthanyl. The animal was then perfused with saline followed by 10% paraformaldehyde, and the brain was removed for subsequent histology. Given the long duration of recordings from each insertion in each animal, and the hypothesis-agnostic experimental design, only 6 cats were required. A 'track' was defined as a single polytrode insertion at a single fixed depth. Data were acquired from 15 tracks from 10 hemispheres, for a total track duration of 166 hours (total amount of time in each track, including gaps in recording). However, only 3 of these tracks were fully spike sorted (Section 4.2), with a total track duration of 34.4 hours. Table 2.2 provides details about each of the tracks recorded from in this study.

A separate set of recordings were also performed in V1 of 13 anesthetized rats. These were mostly pilot experiments to test the recording and stimulus systems and to practice experimental techniques. Procedures were similar to those described above for cats, with a few differences. 230–570 g male and female Long Evans rats were anesthetized by intraperitoneal (IP) injection of urethane (1300–1500 mg/kg), followed by maintenance injections as needed (50–100 mg each). Rats were held in place with ear bars in a stereotaxic frame and were not paralyzed or mechanically respirated, but supplementary oxygen was often provided. As above, heart rate and blood

Design	Columns	Layout	Spacing (μm)	Width (μm)	Length (μm)
1a	3	hexagonal	65	207	1138
1b	3	collinear	43 × 50	210	850
1c	3	hexagonal	75	208	1313
2a	2	hexagonal	65	200	1723
2b	2	hexagonal	50	207	1325

Table 2.3: 54-channel polytrode designs. Width is the shank width, and length is the distance between the most vertically distant sites. Collinear design 1b had tighter horizontal than vertical spacing.

oxygenation were continuously monitored and logged, and blink and pinna reflexes and toe pinch were used to ensure sufficient anesthetic depth. A heat pad was used to maintain constant body temperature. Experiments lasted up to 16 hours each. Unfortunately, very few units were isolated during these recordings, and they are therefore only briefly discussed in Section 4.6.1.

2.2 Polytrodes

Extracellular voltage data (Figure 2.3) were recorded using multichannel silicon polytrodes (Blanche et al., 2005) manufactured by the University of Michigan's Center for Neural Communication Technology, and NeuroNexus (Ann Arbor, MI). The polytrodes used here had a single shank, 15 μm thick and ~ 200 μm wide, with an array of platinum-iridium electrode sites, each 15 μm in diameter (Figure 2.2, middle). Designs had 54 channels in two or three column collinear or hexagonal configurations (Table 2.3). Electrode sites were closely-spaced (43–75 μm), such that a spike from a given neuron would usually generate signal on several neighbouring channels (Figure 2.2, right). This provided more information for spike sorting, and also allowed for 2D triangulation of neuronal position with respect to the plane of the polytrode. By using a model of extracellular signal decay in the neuropil, it may also be possible to estimate the distance between each neuron and the polytrode, and hence each neuron's 3D position relative to the polytrode (Section 4.6.3; Blanche, 2005).

Before use, electrode site impedances were tested using a custom-made low current impedance meter (Appendix D). If a polytrode had more than 5 faulty sites, it was rejected. A faulty site could be the result of a poor wire bond between the polytrode shank and its electrode interface board, or poor etching of the electrode site itself. Faulty sites were often noisy and were grounded out so that they would not induce unnecessary noise onto neighbouring sites or their conductors.

2.3 Data acquisition

Extracellular waveforms from all 54 electrode sites were unity-gain buffered by a pair of 27-channel headstages (HS-27, Neuralynx, Tucson, AZ), and then amplified by a 64-channel 5000× amplifier with fixed analog filters (FA-I-64, Multichannel Systems, Reutlingen, Germany). The first 54

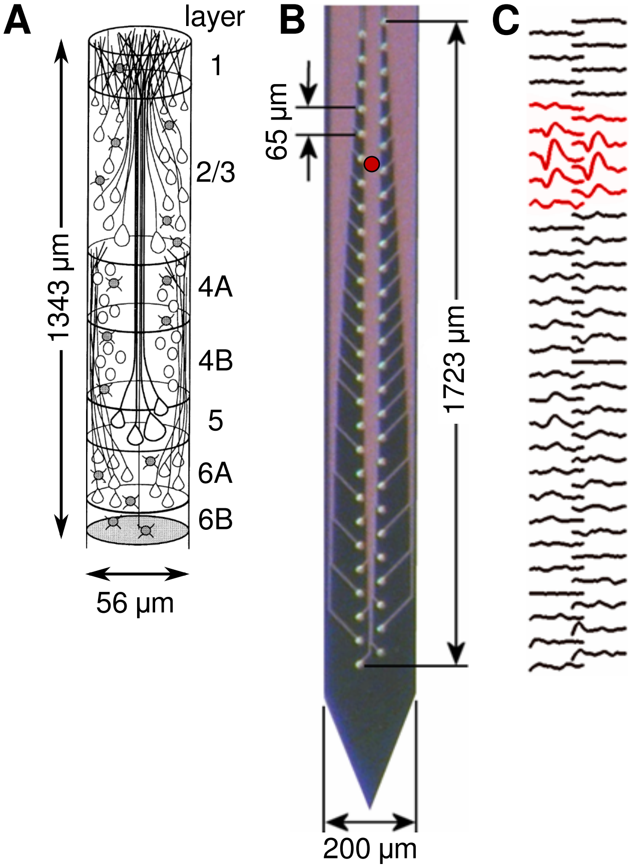

Figure 2.2: Silicon polytrodes. ***A***: A cartoon of a small patch of area 17 of cat primary visual cortex, depicting cells in different layers. Polytrodes are designed to record from all cortical layers. White cell bodies represent pyramidal cells, shaded cell bodies represent non-pyramidal cells. Horizontal and vertical dimensions are not to scale, but ***A*** and ***B*** share the same vertical scale. ***B***: A polytrode ('2a' design) with 54 equally-spaced recording sites arranged in a hexagonal two column layout. Polytrode designs are listed in Table 2.3. ***C***: A spike from a neuron close to the polytrode (red dot in ***B*** shows its estimated 2D position) will typically generate signals on several channels (highlighted in red). ***A*** adapted from Payne and Peters (2001), ***B*** adapted from Blanche (2005).

channels of the amplifier were high-pass filtered (0.5–6 kHz) for use as spike channels (Figure 2.3). Data from a subset of 10 of the 54 electrode sites, evenly distributed along the length of the polytrode, were also separately low-pass filtered (0.1–150 Hz) for use as LFP channels. All 64 channels were then digitally sampled (25 kHz for the high-pass channels, 1 kHz for the low-pass channels) by a pair of 12-bit 32-channel acquisition boards with an internal gain of 1–8× (DT3010, Data Translations, Marlboro, MA). To reduce costs, each acquisition board had only a single analog-to-digital converter (ADC), and could only sample one channel at a time. To closely approximate simultaneous sampling of all channels, each ADC used a sample-and-hold technique, in which its channels were sampled very quickly in succession at 1 MHz, followed by a pause until it was time for the next acquisition timepoint (every 40 μs for high-pass channels, every millisecond for the low-pass channels).

Acquisition was controlled by a custom program "Surf" written in the Delphi programming language (Blanche et al., 2005; Blanche, 2005). A channel of interest could be selected for viewing on-screen (1 ms of data displayed every 100 ms) or directly on an analog oscilloscope. The selected signal was played concurrently through an audio monitor. Data were saved to `.srf` files at a rate of $\sim$ 2.7 MB/s. A single continuous recording consisted of one type of visual stimulus (Section 2.4) and usually lasted no more than 45 minutes, resulting in individual files up to $\sim$ 7 GB in size. Other

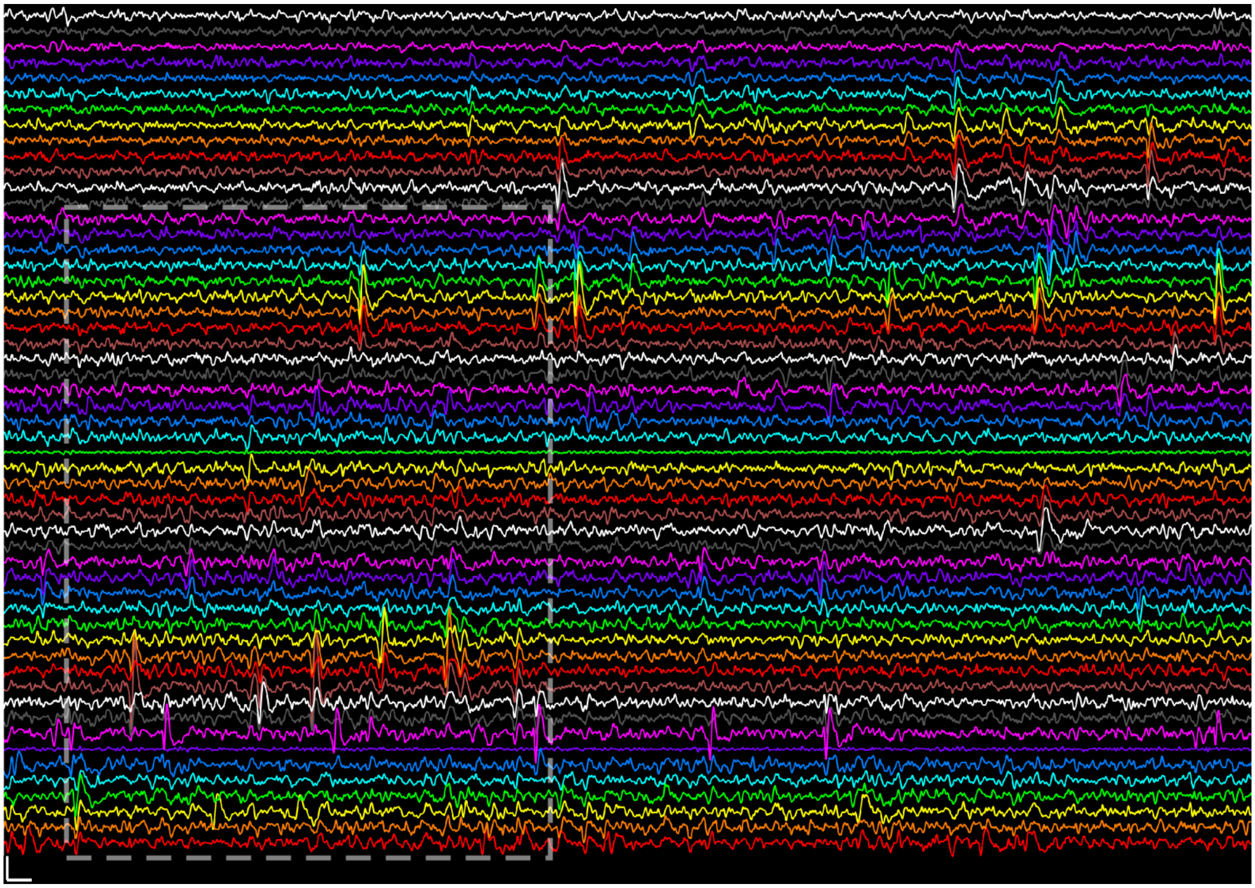

Figure 2.3: 50 ms of high-pass (0.5–6 kHz) voltage waveform data, Nyquist interpolated to 50 kHz (Section 3.3), for all 54 channels of a polytrode recording of spontaneous activity in track ptc15.tr7c (Table 2.2). The channels of this 2 column hexagonal polytrode are arranged here in vertical spatial order, from superficial (top) to deep (bottom). Colours cycle to distinguish neighbouring channels. Two faulty channels (green and violet) were grounded out. Spikes are visible at various times and cortical depths, and most spike waveforms span multiple neighbouring channels. This segment is from a longer sample of spontaneous data available online (Section C.2). Dashed box denotes a segment shown in greater detail in Figure 3.4. Scale bar: 1 ms, 100 μV. Channel spacing: 65 μm.

information was also saved to the `.srf` file, such as polytrode layout, stimulus information (with a microsecond precision time stamp for screen refresh), and the absolute start time of the recording. Saving the absolute start time allowed for later determination of the intervals between recordings, which was important for dealing with drift while spike sorting multiple recordings from the same track all at once (Section 3.2). Online current source density (CSD) analysis (Section 4.6.3) was used as a rough indicator of polytrode depth, and was saved to a separate file. A short period of extracellular multichannel waveform data is shown in Figure 2.3.

2.4 Visual stimulation

Stimuli were displayed on a flat 19" (36 × 27 cm) CRT monitor (Iiyama HM903DTB) at 800×600 resolution and 200 Hz refresh rate. A high refresh rate was required to prevent artifactual phase locking of neurons in V1 to the screen raster (Figure 2.4; Williams et al., 2004). The monitor was placed 57 cm in front of the cat's eyes. At this distance, the monitor subtended horizontal and vertical angles of $\sim 36^\circ$ and 27° respectively, and 1 cm on the screen subtended 1° of visual angle. The monitor was gamma corrected, with a maximum luminance of 116 cd/m^2. This ensured a linear relationship between input pixel values and output luminance. Mouse-controlled stimuli, such as an oriented bar, were displayed simultaneously on two monitors: one for the cat, the other for the experimenter. Both manual and automated stimuli were generated using "dimstim", a freely available custom software package written in Python (Section C.1, Spacek et al., 2009), based on the Vision Egg stimulus library (http://www.visionegg.org; Straw et al., 2006; Straw, 2008) which itself depends on OpenGL (http://opengl.org). Graphics were double-buffered to minimize latency (< 1 ms) and ensure that no frames (5 ms each at 200 Hz refresh rate) were ever dropped. Every screen raster was detected by the acquisition computer directly from the analog video signal cable, and its microsecond precision time stamp was saved concurrently with the spike waveforms. For each type of stimulus, a table of stimulus parameters was generated. On each raster, the stimulus computer sent a stimulus table index value to the acquisition computer which was stored along with the raster time stamp. Thus, there was a record of exactly what was displayed on every screen raster, allowing for precise temporal analysis of the relationship between stimulus and response.

Three different sets of stimuli (designated by letter in Table 2.4) were used depending on the animal (Table 2.1). For all sets (A–C), a mouse-controlled oriented bar was used to map out the average RF position of the population of cells, on which all subsequent stimuli were then centered. Artificial stimuli were used to characterize RF properties of each of the cells in the population, all of which had similar retinotopic positions. Together, the artificial stimuli allowed for mapping of simple and complex cell RFs, as well as orientation, spatial frequency, temporal frequency, phase, and contrast tuning curves. Natural scene movie stimuli were used to examine how responses might differ from artificial stimuli.

Running a full stimulus set took several hours, and each type of stimulus (see below) was usually presented at least once per track. Each stimulus resulted in a single continuous recording of up to 1 h in duration, and an entire track's recordings spanned 5–18 h, including brief gaps in recording (Table 2.2).

Artificial stimuli included: drifting bars at various orientations; flashed gratings at various orientations, spatial frequencies, and phases; drifting gratings at various orientations, spatial and temporal frequencies, phases, and contrasts; and m-sequence white noise movies. Additionally, single raster (5 ms) uniform full-screen white flashes were presented at 1.3 Hz to generate strong synchronous feedforward stimulation of V1 for calculating CSDs (Section 4.6.3). Although not

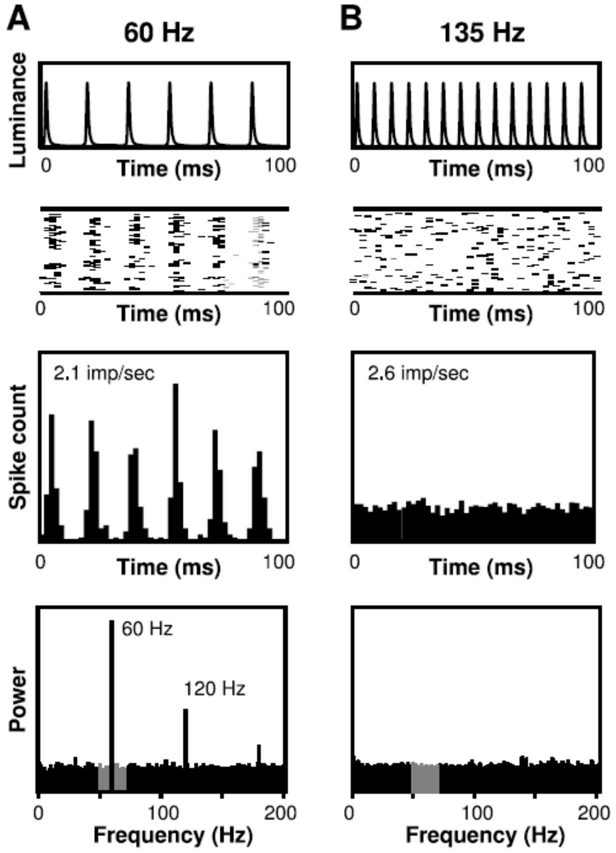

Figure 2.4: Phase locking of V1 to a low refresh CRT monitor. ***A***: When stimulated with a uniform display at 60 Hz refresh rate (***top***), the spike trains of a neuron in macaque V1 phase locked to the display raster (***middle***). A strong 60 Hz component (and its 120 Hz harmonic) is apparent in the power spectrum of its responses (***bottom***). ***B***: Presenting the same uniform display to the same neuron at 135 Hz eliminated phase locking to the display raster. Taken from Williams et al. (2004).

analyzed in this study, sets A and B also included flashed bars at various locations, orientations, and contrasts, and set C included a full-screen sinusoidal grating that progressed over a range of temporal frequencies (‘freqsweep’ column in Table 2.4). For drifting and flashed bars and gratings, all stimulus conditions were presented an equal number of times, and in pseudorandom order to minimize response adaptation (Maffei et al., 1973).

Drifting bars consisted of white and/or black bars on either a grey background or a background of opposite luminance. Bars were 10 or 6° long and 0.5 or 0.3° wide, and drifted at 2.5 or 5°/s for 4 s, for a total of 10 or 20°. Each condition was followed by a 1 or 0.5 s blank period. Steps in orientation were 20 or 30°. Each stimulus condition was presented 8 or 12 times.

Flashed grating stimuli consisted of a rapid succession of stationary, spatially sinusoidal gratings, each displayed for only 40 ms. Flashed gratings can be used to rapidly characterize the RF properties of an entire population of cells (Ringach et al., 1997a,b). Flashed gratings were either presented within a circular aperture 10° in diameter on a mid-grey (50% luminance, 58 cd/m^2) background, or they were presented full screen. Steps in orientation were 10 or 15°. Spatial frequencies ranged roughly logarithmically from 0.05 to 1.6 cycles/°. Steps in spatial phase were 90 or 60°. A Michelson contrast (the luminance difference between bright and dark extremes of the grating, divided by the sum) of 1 was used (100% of the screen’s maximum capability). Mean grating luminance was mid-grey. There was no blank period following each trial, although a blank screen was presented for 40 ms or 2 s every 721 or 6000 trials (29 s or 4 min) as a control to allow estimation of baseline responses. Each stimulus condition was presented 40 or 100 times.

Drifting gratings had a sinusoidal brightness profile in both space and time, and were presented either in a circular aperture 8° in diameter on a mid-grey background, or full screen. Steps in orientation were 45 or 30°. Spatial frequency ranged a roughly logarithmic scale from 0.05 to 5 cycles/°, but was usually kept below 2 cycles/°. Temporal frequency also ranged a roughly logarithmic scale from 0.5 to 20 Hz, but was usually kept below 5 Hz. Trial duration was either adjusted according to temporal frequency, such that each trial consisted of exactly 6 full temporal cycles of the grating, or trial duration was fixed at 6 s. Like the flashed gratings, there was no blank period following each trial, although a blanks screen was presented for 5 or 2s every 24 or 20 trials (120 or 40 s). Contrast ranged roughly logarithmically from 0.016 to 1, but was usually kept at 0.5 or 1. Mean grating luminance was mid-grey. Each stimulus condition was presented 8 times, or only once.

Drifting gratings generally had coarser steps in orientation because more non-orientation stimulus dimensions were included than either drifting bars or flashed gratings (spatial frequency, temporal frequency, and/or contrast). To keep the run time of each stimulus type roughly the same, the number of steps in orientation was therefore decreased.

An m-sequence (Shapley et al., 1991; Reid et al., 1997) was used to generate a white noise movie. The movie had 65535 frames, each 32×32 pixels and 5–12.7° on a side, with an aspect ratio of 1. Each frame was presented for 20 or 40 ms. Every pixel in the m-sequence movie was white and black for an equal amount of time, and each frame had an equal number of white and black pixels. Hence, the first order statistics of the movie were uniform, and summing all the frames from the m-sequence movie resulted in a mid-grey image. The second-order spatial statistics were similarly uniform: all pairwise correlations between pixels were zero when calculated across all movie frames.

Natural scene movies came from two different sources: ‘old’ and ‘new’ (Figure 2.5 & Table 2.4). The ‘old’ set of movies were courtesy of the Peter König lab. They were filmed by attaching a camera to a cat’s head and allowing it to wander through natural environments (Kayser et al., 2003). 64×64 pixel subsets from these movies were displayed on the stimulus monitor at a subtended visual angle of 12.7°. However, these movies had some practical problems, including VHS tape transfer artifacts, low resolution, and a low frame rate necessitating interpolation between frames (courtesy of Nick Lesica). The ‘new’ set of natural scene movies were filmed by the author with a digital camera (Canon PowerShot SD200) at 320×240 pixel resolution and 60 frames per second (fps), low to the ground, with movements approximating that of a cat exploring its environment, as well as sudden saccade-like movements. A total of 52 min were filmed outdoors over 3 days in a variety of wooded or grassy locations in Vancouver, BC. Footage consisted mostly of dense grass and foliage with a wide variety of oriented edges. Focus was kept within 2 m and exposure settings were set to automatic. The horizontal angle subtended by the camera lens (51.6°) was measured for proper

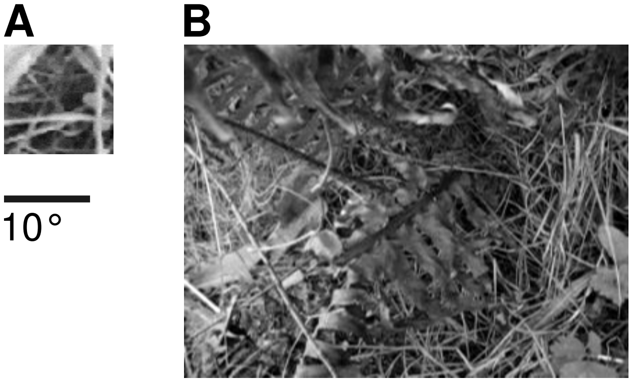

Figure 2.5: Example frames from 'old' (***A***) and 'new' (***B***) natural scene movie sets. Frames are shown at the same relative size as when they were displayed on-screen during experiments. Both sets exceeded the classical receptive field sizes of cells in this study. Individual movie pixels were roughly the same size in both sets (0.2 and 0.16°, respectively). Scale bar applies to both panels. 10° is approximately the width of a closed fist at arm's length.

Set	Movies	Repeats	Blank screen	Refresh rate	Flashed bars	Freqsweep
A	old	25	no	no	yes	no
B	new	200–400	yes	yes	yes	no
C	new	400	yes	yes	no	yes

Table 2.4: Stimulus sets. The set used for each animal is shown in Table 2.1. See text for details.

scaling to match the visual angle subtended by the movie on the stimulus monitor. These movies are available upon request. An example movie is available at http://dimstim.github.io.

Both movie sources resulted in clips of 1–5 min in duration, presented up to 8 times in succession. Shorter clips (4.5–5 s in duration) were also presented, but much more repeatedly. Short clips from the 'old' movies in stimulus set A were repeated 25 times, while short clips in the 'new' movies in sets B and C were repeated 200–400 times ('repeats' column in Table 2.4). These were used to study the reliability of neuronal responses to repeated naturalistic stimuli. They were also intended for studying how responses might change over time as a result of plasticity (Yao et al., 2007).

In stimulus set A, except for the occasional brief periods in between artificial stimulus trials (see above), spontaneous activity was acquired without the screen. In sets B and C, spontaneous activity was acquired while presenting a blank grey screen. Finally, sets B and C included natural scene movies and blank screen stimuli which were run at a 66 Hz refresh rate in addition to the usual 200 Hz. These may be used to investigate in greater detail how a low refresh rate screen might affect responses.

Other than the full-screen flashes and some of the blank screen stimuli, stimuli were presented monocularly to the eye that was dominant for the majority of the recorded neural population. This was judged by listening to visually evoked responses on the audio monitor. Monocular stimulation avoided the difficulty of keeping the eyes accurately converged in an anesthetized animal. Although natural scene movies were more naturalistic than the artificial stimuli, they excluded colour, and because they were presented monocularly they also lacked depth. At least once per track, the area centralis (AC) of the open eye was mapped by reverse ophthalmoscopy onto a sheet of paper on the stimulus monitor, and the center position of the stimuli was marked as well. This confirmed that the neural population had RFs at low eccentricity, within a few degrees of the AC. Classical

RF sizes were usually $< 5°$ at the low eccentricities in this study.

2.5 Histology

At the end of the experiment, all cats except ptc15 were perfused with saline followed by 10% formalin to fix the tissue. The visual cortex was blocked and sectioned coronally into $\sim 50\ \mu$m thick slices on a freezing microtome. The polytrode track could sometimes be localized to less than a handful of slices. This potentially allowed determination of track depth, and therefore roughly which cortical layers the polytrode spanned. Histology also had the potential to roughly determine the angle between the polytrode track and the vector normal to the cortical surface, and therefore how inter- or intra-columnar the track may have been. However, the histology work remains incomplete and is not presented here. Future efforts may yet recover useful histological information from the stored samples.

3 Spike Sorting

Before any spike analysis can be performed, single unit spikes must be extracted from the extracellular multichannel waveform data. Spike sorting has a long history, and can be a very difficult process (see Lewicki (1998) for a review). Not only must spikes be separated from noise, but they must also be separated from each other. A further complication is that the number of neurons captured by a given recording and their specific extracellular multichannel waveform shapes are unknown *a priori*. Accurate spike sorting is essential for later analyses, especially those that depend on precise timing relationships such as spike correlations (Ventura and Gerkin, 2012; Figure 3.1; Appendix B). Spike sorting therefore deserves considerable attention.

3.1 Background

There are a wide variety of problems that may be encountered during spike sorting. During spike detection, noise events may be mistaken for spikes (false positives), or spikes may be mistaken for noise (false negatives). Once spikes have been detected, they need to be clustered. A cluster may be mistakenly divided into two or more clusters (oversplitting), or it may mistakenly combine spikes from multiple neurons (undersplitting) resulting in a multiunit cluster of single unit spikes. Conversely, instead of being discarded, non-isolatable multiunit spikes may be mistaken for isolated single unit spikes. Spike shape variability is another complication. During longer recordings, electrode and tissue drift can result in non-stationary spike waveforms, and during bursts, spike amplitudes can diminish (Buzsáki, 2004). Neurons can also exhibit great variability in their responsivity, sometimes resulting in complete silence for long periods of time (Henze et al., 2000; Lennie, 2003; Buzsáki, 2004; Shoham et al., 2006; Mizuseki and Buzsáki, 2013). All these problems are a potential source of bias.

Electrode arrays with many closely spaced recording sites have the potential to improve spike sorting quality. Traditionally, most electrophysiology has been done using single channel electrodes (either tungsten in glass or single wire), stereotrodes (McNaughton et al., 1983), or tetrodes (Gray et al., 1995). More recently, multichannel electrodes have come into use (Section 1.1). As the number of recording sites increases, so can the complexity of spike sorting. As long as the spacing between each site (or each tetrode) is great enough such that a single spike shows up on no more than one channel (or tetrode), single channel spike sorting methods can easily scale up to any number of channels. However, for electrodes with sufficiently low site spacing (< 100 μm and ~ 1 MΩ site impedances), a single spike may be picked up by multiple electrode sites (Figure 2.2). This has the benefit of potentially increasing the spike signal to background noise ratio, tracking neurons over time despite drift, improving the ability to sort spikes from neurons very close to each other, and providing a better estimate of the locations of recorded neurons relative to the electrode

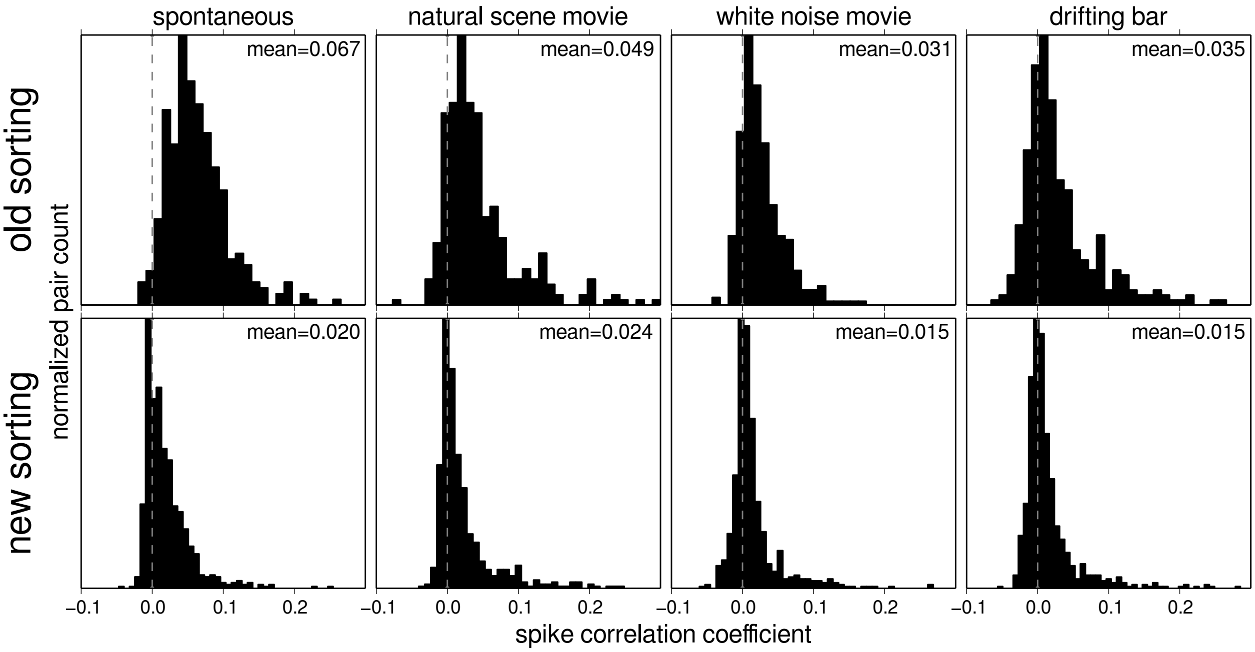

Figure 3.1: Motivation for careful spike sorting. Each plot is a distribution of spike correlations (Appendix B) between all active neuron pairs in a given recording. Four different types of recordings from track ptc15.tr7c are shown, one in each column (3, 15, 22 and 12 min in duration, left to right). The top row used spikes sorted by the previous multichannel template matching method (Section 3.1.1) while the bottom row used spikes sorted by the method presented here. Mean spike correlations between the (fewer) neurons measured from spikes sorted with the previous method were artifactually higher than those measured from spikes sorted with the method presented here (0.031–0.067 vs. 0.015–0.024). One reason for this may have been an implementation problem with the old method: it was possible for a single spike to be assigned to more than one cluster.

and to each other (Drake et al., 1988; Harris et al., 2000; Csicsvari et al., 2003; Blanche et al., 2005; Du et al., 2009; Figures 3.6, 3.12 & 4.3).

However, these potential benefits come at the cost of further complications. One of these is in spike detection. A single spike that generates signal on more than one channel can trigger multiple threshold crossing events. Deciding which threshold crossing events are close enough together in space and time to be considered part of the same spike is not straightforward. It may be that multiple adjacent events belong to the same single spike, or they may belong to multiple spatiotemporally adjacent, yet distinct, spikes. Spike detection therefore becomes a problem not just of detecting threshold crossing events, but of clustering them into distinct spikes (Swindale and Spacek, 2015). Only then can these spikes be clustered into different single units.

Once spikes have been detected, another complication is consistently aligning spikes in space and time. If a neuron's physical location falls in between two or more electrode sites, the channel on which its spike waveforms have their maximum amplitude may vary as a result of multichannel waveform shape variation, whether systematically due to drift or randomly due to noise. This can cause the channels that are selected for inclusion for spike sorting for one spike to differ from those of other spikes from the same neuron. A method is therefore required for choosing a common set of channels and timepoints for all spikes that might conceivably be from the same neuron. A common set of channels and timepoints is necessary for the dimension reduction and clustering steps that follow spike alignment.

Perhaps the greatest difficulty is the current lack of sufficiently realistic ground truth data, whether simulated or experimental, of many simultaneously spiking neurons recorded by closely spaced cortical electrode arrays. Unfortunately, until this situation changes, objective performance comparisons of different spike sorting methods will remain very difficult to make (Section 3.12; Einevoll et al., 2012).

Efforts were made here to deal with all of the above spike sorting challenges, but one further difficulty was not dealt with: spatially adjacent spikes overlapping in time. The solution to this is spike overlap decomposition, in which each spike overlap event is treated as a linear superposition of spikes from different neurons at different temporal shifts (Lewicki, 1994; Segev et al., 2004; Franke et al., 2010; Prentice et al., 2011; Jäckel et al., 2012; Marre et al., 2012; Pillow et al., 2013). Spike overlap decomposition was considered too difficult here in the context of long duration recordings with non-stationary waveforms. Fortunately, given the low mean firing rates of neurons in cat V1 (Section 4.3), spike overlap is much less of a concern here than in retinal ganglion cell (RGC) recordings (Segev et al., 2004; Prentice et al., 2011; Jäckel et al., 2012; Marre et al., 2012; Pillow et al., 2013).

3.1.1 Existing methods

There are a wide variety of spike sorting methods (Lewicki, 1998), but no single existing method addresses all of the problems associated with spike sorting of long duration cortical polytrode recordings. Most methods divide the process into separate spike detection, dimension reduction, and clustering steps. Some methods extract simple spike features, such as spike height or width, as a form of dimension reduction, but may miss other subtle but important differences in spike shape (Gray et al., 1995; Nguyen et al., 2003). Some only handle single channel or independent channel recordings (Lewicki, 1994; Zouridakis and Tam, 1997, 2000; Quian Quiroga et al., 2004; Wood and Black, 2008; Wolf and Burdick, 2009) while others are tailored to stereotrode or tetrode data, but may not scale well to higher numbers of non-independent channels (Gray et al., 1995; Fee et al., 1996; Harris et al., 2000; Nguyen et al., 2003; Gasthaus et al., 2009; Calabrese and Paninski, 2011). Some methods require Gaussian-shaped clusters, and therefore assume stationarity of spike waveforms over time (Harris et al., 2000; Litke et al., 2004; Hazan et al., 2006). Others explicitly deal with non-stationarity (Bar-Hillel et al., 2006; Wolf and Burdick, 2009; Calabrese and Paninski, 2011) but may not scale well to large numbers of spikes and clusters (Hulata et al., 2002; Shoham et al., 2003; Pouzat et al., 2004; Gasthaus et al., 2009).

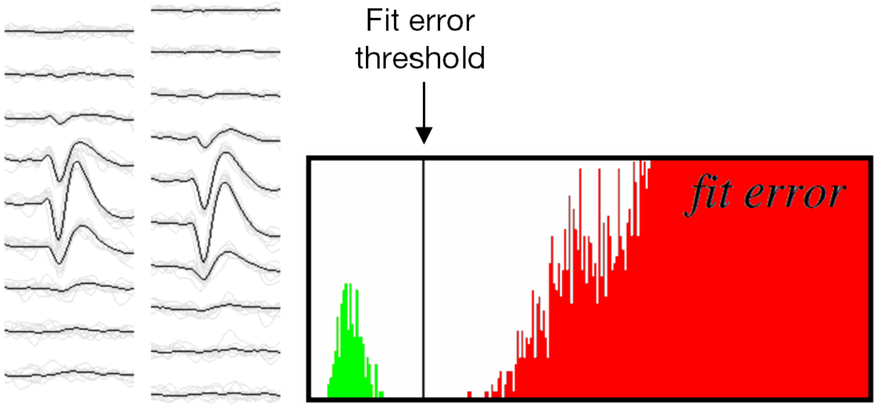

Figure 3.2: Multichannel template matching. ***Left***: A multichannel template with member spikes (grey) and mean waveform (black). ***Right***: The distribution of difference values between the multichannel template and all points in the recording. Ideally, this was a bimodal distribution, allowing a threshold to be set to separate spikes (green) from noise (red). In practice, this distribution was often unimodal, with no obvious place to set the threshold. Adapted from Blanche et al. (2008).

Many of the problems associated with polytrode spike sorting were described by Blanche (2005), but some of that method was never fully implemented in software. More fundamentally, the method of Blanche (2005) had a number of problems, most of which were a result of using multichannel template matching for both spike detection and clustering. Briefly, initial spike detection was performed on short random sections of data, resulting in a few thousand spikes. These spikes were then clustered using a binary split algorithm, followed by k-means clustering. The intent was to sample enough data that all neurons would be represented by a cluster of at least a few spikes each. The multichannel mean waveform, or template, was calculated for each cluster and compared to the data at every timepoint. Temporally local minima in root mean square (RMS) difference values between each template and every timepoint in the data were found, and their distribution plotted (Figure 3.2). Ideally, the distribution of difference values would be bimodal, the lower mode being true matches of the template to its spikes in the data, and the upper mode being coincidental matches to noise. This would give the user a clear indication of where to set the RMS difference threshold between matches to spikes and to noise. However, in practice the distribution was often unimodal, with no clear distinction between matches to spikes and to noise, and therefore no principled way for the user to set a difference threshold.

Another problem was that template matching cannot typically handle large systematic waveform variability, such as that due to drift (Section 4.4). The method of Blanche (2005) assumed that templates derived from just a portion of a track's recordings were accurate across all of its recordings. These could be separated by several hours, and therefore could be subject to substantial amounts of drift or other systematic sources of waveform variability. As a result, templates partly derived from earlier sections of data could generate substantial false positive or false negative matches to later sections of data, and vice versa. A related problem with template matching was that the initial random sampling meant that a rarely firing neuron might never result in a template, and therefore might be missed altogether. This is a significant issue, given that most neurons in cat V1 have very low firing rates (Section 4.3). Finally, an implementation issue allowed some spikes to match multiple templates and therefore belong to multiple neurons at the same time. This likely resulted in artifactually strong spike correlations between neurons (Figure 3.1).

The above arguments against template matching for long duration cortical polytrode recordings also hold for many other spike sorting methods that are based on template matching (Segev et al., 2004; Franke et al., 2010; Prentice et al., 2011; Jäckel et al., 2012; Marre et al., 2012). However, template matching may be more suitable for other kinds of data, such as shorter *in vitro* multichannel electrode array (MEA) recordings of RGCs (Segev et al., 2004; Prentice et al., 2011; Jäckel et al., 2012; Marre et al., 2012). Such recordings may be less vulnerable to drift than the longer *in vivo* cortical polytrode recordings reported here. RGCs also have higher firing rates, thereby mitigating the above mentioned sampling problem. Higher firing rates come with greater severity of spike overlap from spatially adjacent neurons. Fortunately for RGC recordings, template matching is well-suited for spike overlap decomposition of recordings with stationary waveforms.

3.1.2 Clustering methods

Compared to template matching, the use of separate spike detection, dimension reduction, and clustering steps is the more common approach to spike sorting. For the clustering step, a variety of different clustering methods have been used, including manual cluster cutting (Gray et al., 1995), k-means, hierarchical clustering (HC), expectation-maximization (EM) of a Gaussian mixture model (GMM), and superparamagnetic clustering (SPC). All were considered for use here, but all were rejected.

K-means (Forgy, 1965; MacKay, 2003) is a simple, commonly used, two step iterative algorithm. Before iteration begins, a set of cluster centers (or means) is initialized in some manner, usually randomly. First, each data point is assigned to the nearest cluster. Second, the position of each cluster is updated by taking the mean of the positions of all its member points. Given enough iterations, k-means is guaranteed to converge (MacKay, 2003). However, it has two major problems. First, because it uses Dirichilet domains, which assign each point to a cluster based solely on which cluster's centre is nearest, highly elongated and non-convex clusters are often poorly separated. Second, the number of clusters k must be known *a priori*, which it is not during spike sorting. This can be dealt with by running k-means multiple times, each with a different value of k, but then the best solution must be evaluated somehow. Doing so is complicated by the fact that k-means is not deterministic, due to its random initialization which is important to test robustness.

Hierarchichal clustering organizes data points into a hierarchy, with individual points at the bottom, and groups of increasing size further up the hierarchy, culminating in a single group for all the points at the very top. There are two types of HC: agglomerative, which combines points into increasingly larger and fewer clusters; and divisive, which iteratively divides the data into increasingly smaller and more numerous clusters. Both require a pairwise similarity matrix (e.g., distance matrix) of all points. Murtagh and Contreras (2012) provide a recent overview of the field of HC, while Fee et al. (1996) apply it to the problem of spike sorting. An interesting feature of HC is that each level in the hierarchy can be considered a valid clustering. Moving up or down the hierarchy allows one to track how different points are related to each other at different spatial scales. One challenge then is deciding at what level to "slice" the hierarchy. There are many different distance metrics and group linkage criteria available for HC, and deciding which to use is another challenge. In general, for agglomerative clustering, at each level in the hierarchy the two points closest to each other are merged, and for divisive clustering the two points furthest from each other in a group are split. This may make traditional HC vulnerable to the same problem as k-means: if each cluster's position is based solely on its centre, or its edge, or some other single feature that does not fully describe cluster shape and point density, clusters can become poorly separated.

EM is a method of fitting a model to data, in this case, models of clusters to points in cluster space. EM searches for a set of model parameters that maximize the likelihood that the data were generated from the model (Bilmes, 1998). A GMM is one such model, where the number of Gaussians itself is also a parameter. Like k-means, EM is a two step iterative process. First, in the expectation step, it guesses a set of parameter values, and calculates the expectation that the model fits the data. Then, in the maximization step, it chooses parameter values that maximize the expectation. On each iteration, the likelihood of the chosen parameter values is guaranteed to increase and eventually reach a local maximum. KlustaKwik (Harris et al., 2000; Hazan et al., 2006) is a widely used clustering program tailored for spike sorting, which implements EM of a GMM (EM+GMM). However, use of a GMM assumes that clusters are Gaussian in shape, which is not always the case. For example, neurons drifting over time can result in highly non-Gaussian clusters (Figure 3.12). Bar-Hillel et al. (2006) clustered spikes using EM+GMM but allowed for drift by splitting the data into time chunks, and then linking the clusters in consecutive chunks. Wolf and Burdick (2009) and Calabrese and Paninski (2011) both employed similar strategies. However, EM can be computationally expensive, and is non-deterministic, again due to random initialization.

SPC (Blatt et al., 1996, 1997; Quian Quiroga et al., 2004) represents each point as a magnetic domain that can have one of a number of possible spin values. Interactions between points are modelled as an exponential function of squared distance between points. Points are assigned random spins, and a Monte-Carlo algorithm is run at different simulated temperatures, allowing the points to affect each other's spins in a pairwise manner. At each temperature, a point's spin is perturbed, and neighbouring points correspondingly change their spin with some probability. Those that do so with a high enough probability are classified as belonging to the same cluster. This is repeated for all points and many different temperatures. The ideal temperature is picked from a (superparamagnetic) range of temperatures that exhibits cluster stability. The benefit of SPC is that it makes no assumptions about cluster shape, does not require prior knowledge of the number of clusters, and does not necessarily assign points to the nearest cluster mean. However, SPC can be slow if many temperature steps are used, and choosing an ideal temperature can be difficult.

3.2 Overview

The spike sorting procedure used here was similar to that described in Swindale and Spacek (2014). First, raw spike data (0.5–6 kHz) were preprocessed (Section 3.3) by Nyquist interpolation to increase the effective sampling rate, and to correct for sample-and-hold delays in the acquisition hardware. Next, spike detection (Section 3.4) designated plausible neuronal spikes in the raw data. These were then initially partitioned (Section 3.6) into up to one cluster per channel, according to the channel with the sharpest voltage peaks for each spike. The multichannel spikes in each of these initial clusters were then aligned in time (Section 3.7). Next, dimension reduction (Section 3.8) was performed on each cluster using principal component analysis (PCA) or independent component analysis (ICA) of channels with significant signal. The 3 most significant components were plotted in 3D. To check for drift, spike time could be chosen as one of the 3 dimensions. Spikes were then clustered (Section 3.9) in the chosen 3D space using the gradient ascent clustering (GAC) algorithm (Swindale and Spacek, 2014). These more refined clusters were cleaned (usually by further clustering in ICA space) and compared and realigned to similar clusters to check for significant separation (Section 3.10). Sufficiently similar clusters were merged, and sufficiently multimodal clusters were split. Every merge or split required reverification of nearby clusters. All remaining unclustered spikes that had been discarded could be compared to existing clusters for potential incorporation. Finally, each cluster's autocorrelogram could be checked for a satisfactory refractory period, although this was found to be of dubious utility (Section 3.11).

Within a given track, sorting each recording independently would result in a correspondence problem: how to determine cluster correspondence across recordings adjacent in time? Track-wide sorting was done to avoid this correspondence problem. It had the added benefit of potentially increasing the size of clusters that would otherwise be too small to consider on a single recording basis (i.e., when considering a shorter stretch of data), and would therefore be erroneously discarded. While this made small clusters bigger and denser, it also made big clusters even bigger. This was a disadvantage because in PCA space these could then encroach on each other, and on smaller clusters. However, ICA helped get around this problem (Section 3.8). Because it was desirable to conserve time gaps between recordings to keep track of drifting units over time, null data of the appropriate length was inserted between recordings to represent the time gaps. This also made it clear to the user where and how long the gaps were when scrolling through the continuous data. Time gaps would then also show up in the 3D cluster plots of detected spikes when plotted against time (Figure 3.12, right column).

3.2.1 Software

To implement the methods described here, new spike sorting software, called "spyke" (http://spyke.github.io), was written in the Python programming language (http://python.org; Spacek et al., 2009; Spacek and Swindale, 2009). For spike sorting large data sets such as an entire track, spyke requires a 64-bit computer with at least 4 GB of random-access memory (RAM) and a 64-bit operating system. In principle, spyke can run on Linux, Windows, or Mac OS X, but most recently has only been tested in Linux (Xubuntu 14.04). Spyke can deal with millions of spikes at a time ($\sim$ 1 million spikes per GB of RAM). The longest sorted track (ptc15.tr7c) was 11.8 hours long, excluding recording time gaps. There was no limitation in the recording duration that could be sorted, only the memory limitation as a function of the number of detected spikes. A spike sorting session could be saved and resumed at any time. For a track-wide sorting session, the saved data, including waveforms of detected spikes, took $\sim$ 1 GB of disk space per million spikes (7 GB for ptc15.tr7c). This data could be further sorted even in the absence of the original `.srf` raw data files (107 GB for ptc15.tr7c) saved by the acquisition software "Surf" (Blanche et al., 2005). All spike sorting methods described here were implemented in spyke. Most of the methods here were also independently implemented in Fortran (Swindale and Spacek, 2014, 2015), which allowed for cross-validation of methods. Track-wide sorting with spyke was used to sort all spike data for analysis in subsequent chapters. Further details about spyke are in Section C.2.

3.3 Preprocessing

The data acquisition boards acquired raw spike data at 25 kHz using a sample-and-hold technique in which each ADC cycled between its channels at a rate of 1 MHz, taking one sample at a time (Section 2.3). This resulted in an artifactual delay of 1 μs between consecutive channels, and up to 31 μs between channels at opposite ends of a single 32 channel ADC board. Furthermore, spike detection and alignment at 25 kHz is inferior to that at 50 kHz or 100 kHz, because positive and negative peaks in spikes happen quickly, and insufficiently fast sampling can cause these peaks to artifactually fall below spike detection threshold (Blanche and Swindale, 2006). Nyquist interpolation was used to correct for both of these artifacts. Nyquist interpolation to 50 kHz was chosen because there was no appreciable improvement in going any higher (Blanche and Swindale, 2006). Sample-and-hold correction and interpolation were performed in a single step on the fly every time raw data were called for, such as when scrolling through the data or during spike detection. Data were Nyquist interpolated by convolution (`numpy.convolve`) with a bandpass sinc function. The sinc function used for each channel was offset by the appropriate amount to correct for its sample-and-hold delay.

Because the frequency content of extracellular spikes spanned most of the recorded bandwidth (0.5–6 kHz), any further digital filtering to reduce noise would have also reduced signal, and the signal to noise ratio would have remained unchanged (Swindale and Spacek, 2014). Therefore, no filtering was performed on the raw data.

3.4 Spike detection

Next, spike detection was performed on the preprocessed data. Data were divided into slightly overlapping blocks (10 s long, 2 ms overlap), and preprocessing and detection were performed on each block independently. This allowed time dependent voltage thresholds to be calculated and applied, with a time resolution of one block. It also allowed detection to be split up into parallel independent tasks, which made for faster execution on a multicore CPU. Working on blocks was also much more memory efficient than trying to preprocess an entire track's raw data (107 GB for ptc15.tr7c) at once. Blocks overlapped slightly to accommodate spikes that straddled block borders. Channel and time dependent voltage thresholds were calculated by estimating the noise on each channel within each block. Noise levels can be estimated by taking the standard deviation of each channel's voltage data, but estimation by standard deviation is biased by the presence of spikes in the data. Higher spike rates and higher amplitude spikes lead to higher standard deviation, even when background noise levels remain the same. The median is a better method for estimating noise in the presence of spikes (Quian Quiroga et al., 2004), and was used instead. Voltage thresholds were calculated according to

$$V_t(i,j) = \max\left\{A\,\frac{\mathrm{median}(|V(i,j)|)}{0.6745},\ V_{min}\right\} \tag{3.1}$$

(Quian Quiroga et al., 2004), where A is a noise multiplier, V is the preprocessed voltage data, i and j are channel and time block indices respectively, and V_{min} is a minimum threshold. $A = 6$ was applied, and gave good discrimination between noise and spike-like peaks. $V_{min} = 40\ \mu\mathrm{V}$ was also applied: events below that were deemed too difficult to reliably discern from noise events, even on channels and during periods of very low noise.

A peak was defined as the largest extremum between each pair of zero crossings on a given channel. A subset of these peaks had amplitudes exceeding V_t, referred to here as trigger peaks. Trigger peaks were examined one at a time in temporal order. Each was tested to decide if it was part of a spike, and also if it was the primary peak of that spike, i.e., the best peak to align to in space and time. Positive and negative peaks were given equal consideration. Because a spike can generate signal on several channels, and because spikes generally have two or more peaks of opposite sign (Figure 3.3), a local spatiotemporal search was performed around each trigger peak, constrained to channels within a radius of $r_{lockout} = 150\ \mu\mathrm{m}$ and to timepoints within $\pm\ dt_{max} = 0.4$ ms. All peaks within the search window were compared to the trigger peak. Peaks were characterized by their "sharpness", defined as

$$S_p = \frac{{V_p}^2}{dt_z} \tag{3.2}$$

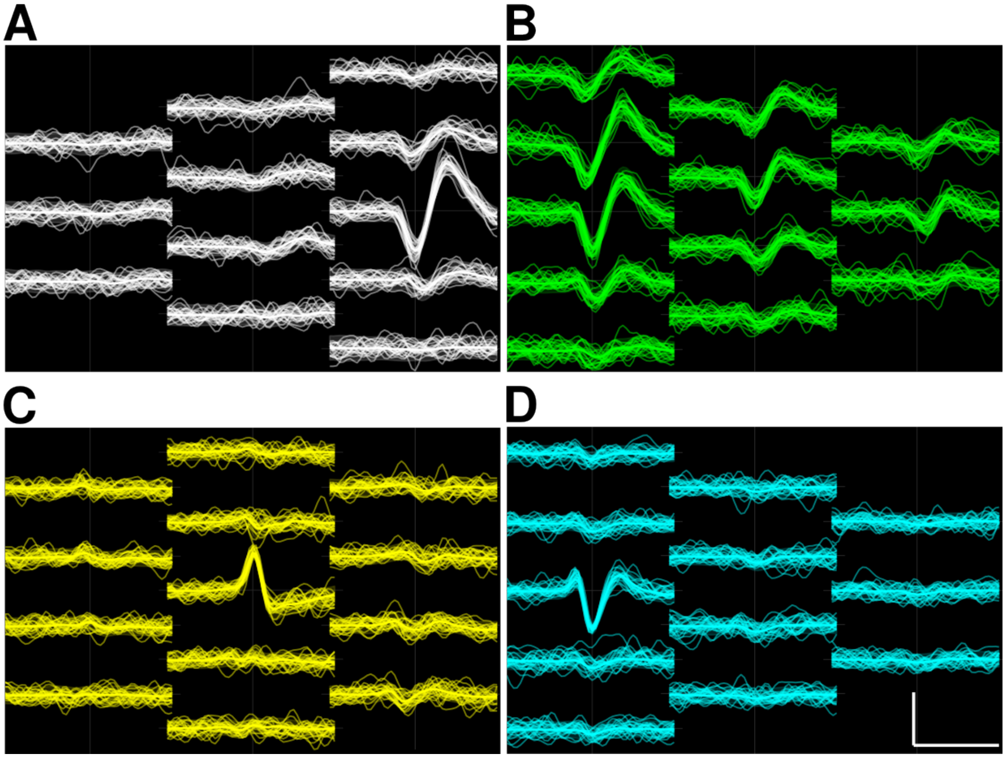

Figure 3.3: Examples of 4 types of spikes, shown clustered, with 20 randomly chosen spikes plotted in each cluster. Mean and standard deviations of waveforms (thick lines and translucent bands, respectively) are mostly obscured by the spikes. Upper panels show larger slower spikes, while lower panels show smaller faster ones. ***A***: Spikes with most of their signal on one channel and much less on surrounding channels. These had a prominent negative and positive peak. ***B***: A cluster with similar spike shape, but with a broader distribution in space. The neuron generating these spikes likely fell in between the two channels with the largest amplitude. ***C***: Rarer, mostly positive-going spikes, with a very slow, low amplitude negative peak. ***D***: Spikes with 3 peaks, in this case 2 smaller positive peaks flanking a larger negative peak. Scale bar: 0.5 ms, 100 μV. Site spacing: 65 μm.

where V_p is the amplitude of the peak and dt_z is the time between zero crossings on either side of the peak. As a result of the squaring, for large and small peaks that were proportionally shaped, larger peaks were defined as sharper.

On each channel within the search window, the sharpest peak was found and paired with its sharpest temporally adjacent peak, either the one immediately to its left or its right, if any. Peak-to-peak sharpness S_{pp} was defined as the sum of the sharpness measures of the adjacent pair. The channel within the search window with maximum S_{pp} was labelled the primary channel of the putative spike. This channel sometimes differed from the channel of the original trigger peak around which the search window was centered. In other words, sometimes the trigger channel and the putative spike's primary channel were not the same. When this was the case, the current trigger peak was rejected, and eventually the sharpest peak on the primary channel would be encountered as a threshold crossing trigger peak in its own right. This heuristic helped prevent excessive shifting around in space from trigger channel to primary channel, which qualitatively helped reduce migration to the spatial edge of the search window, and the associated edge effects. It also prevented low amplitude spikes from locking out (see below) neighbouring larger amplitude spikes that immediately followed.

A second heuristic was also applied in time: if the trigger peak came before the sharpest peak, the trigger peak and its associated search window was again rejected. This prevented prematurely setting the temporal detection lockout (see below) on the primary channel, which was important because sometimes the early peak on the primary channel was part of a different yet to be detected spike on a nearby channel. Both the spatial and temporal heuristics were used to prevent an incorrectly centered detection lockout in space and time. In this way, the search window was used to check if the trigger peak and the sharpest peak were identical.

Next, the adjacent peak pair with maximal S_{pp} was tested to ensure that its peak-to-peak voltage V_{pp} exceeded $V_{ppt} = 1.5V_t$. If so, the adjacent peak pair was considered to be the center of a plausible spike. Together, the V_t and V_{ppt} absolute value voltage thresholds, and the dt_{max} temporal threshold constituted a very general expectation of what a spike should look like. This allowed for detection of both fast and slow spikes with two or more peaks of alternating sign, and no restriction on the relative amplitudes of the peaks, nor their absolute sign. However, monophasic spikes were excluded (see Section 3.13.1).

Unfortunately, there was no parametric measure that could be used to align a spike to one of its peaks. A small change in spike shape due to noise could cause a sudden shift in the sharpness rank of peaks belonging to the spike. Therefore, a new inclusion window was formed aligned in time to the negative peak of the pair (Figure 3.4), and centered in space on the primary channel, as before. The negative peak was chosen because extracellular spikes are mostly negative-going, and prior to clustering with many other similar spikes, there was no way to otherwise determine which of the two sharpest peaks, positive or negative, was more appropriate for alignment. Subjectively evaluated alignment errors occurred mostly on rarer, positive-going spikes (Figure 3.3C). After initial clustering (Section 3.6), such errors were later corrected by best fit realignment, or by forcing realignment to the biggest positive peak (Section 3.7).

The new spatiotemporal window was constrained to an inclusion radius $r_{include}$ which, like $r_{lockout}$, was also set to 150 μm (but could be set differently if desired). For a 3 column polytrode with 65 μm hexagonal spacing, $r_{include} = 150$ μm resulted in about 12 included channels per spike. One millisecond of data was saved: 0.4 ms before the earlier of the two adjacent peaks, and 0.6 ms after (although figures in this chapter show ± 0.5 ms, this was later changed to capture more of the slower afterhyperpolarization (AHP) of each spike). At a sampling rate of 50 kHz, this gave 50 timepoints per channel, for a total of ~ 600 voltage values per spike. These were saved to disk along with various other parameters characterizing each spike. Saved waveforms remained interpolated

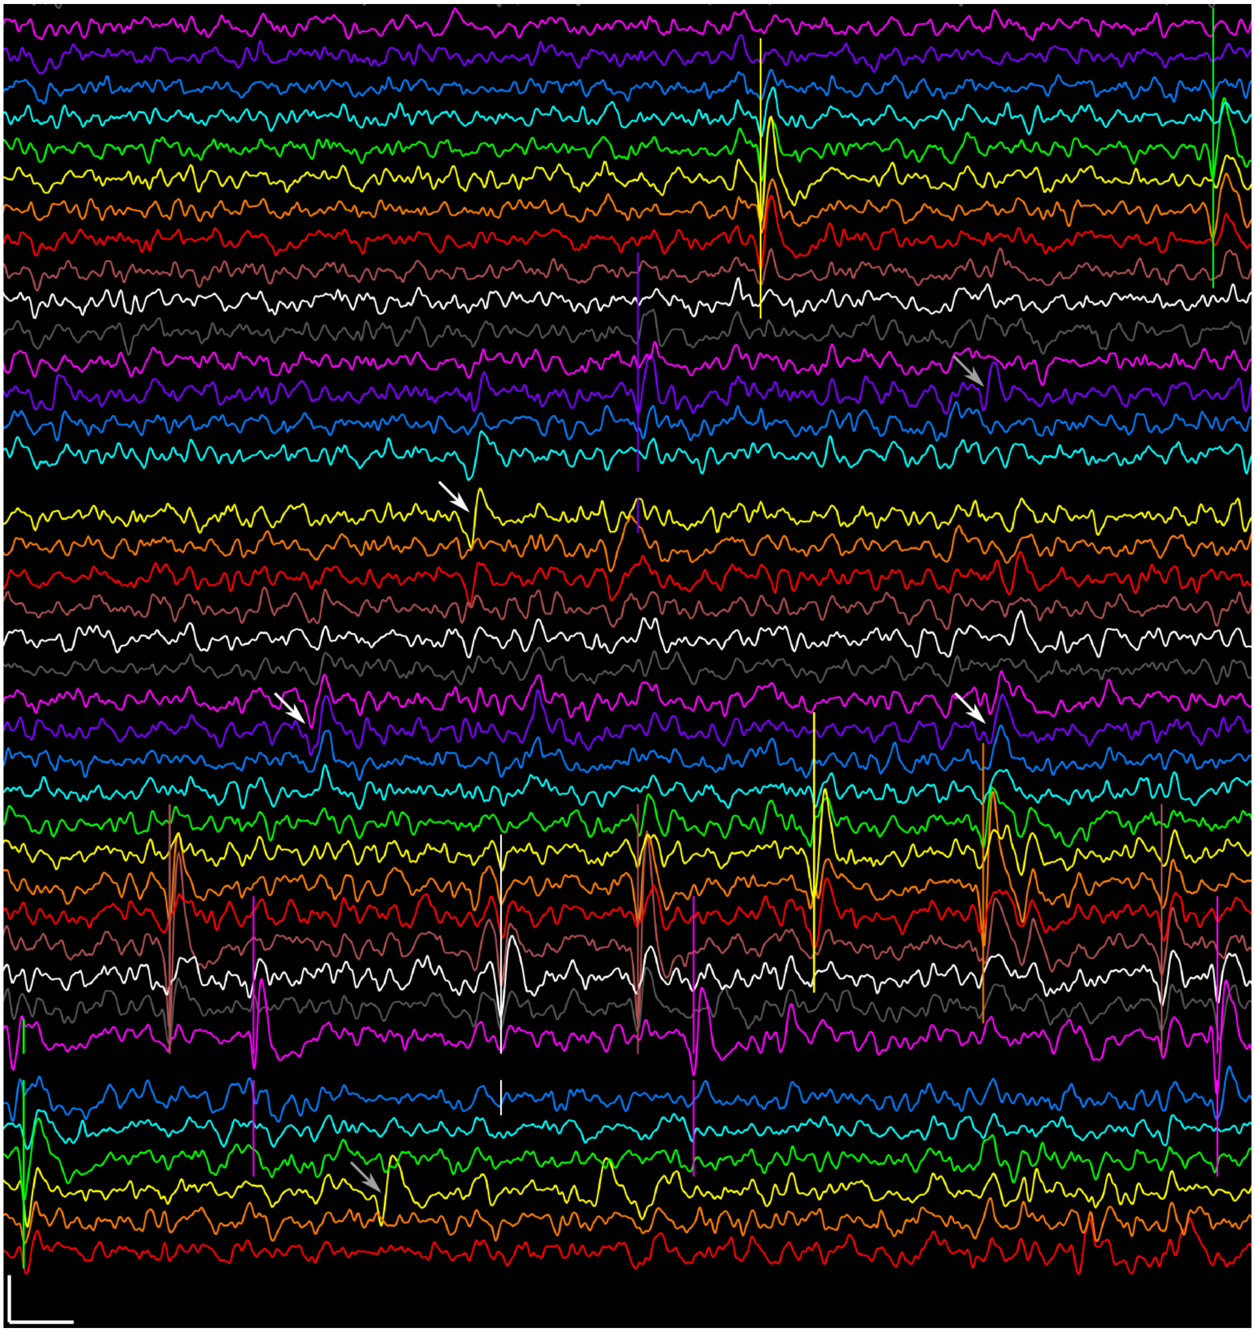

Figure 3.4: Spike detection of data in the dashed box in Figure 2.3. Detected spikes are indicated with a vertical raster line at the negative peak of their primary channel, spanning all the channels that fell within a 150 μm radius — in this case, 4 channels above and below each primary channel. Raster lines are coloured according to the primary channel of each spike. Low amplitude single channel (grey arrows) and multichannel (white arrows) spike-like events were not detected because they were subthreshold. Each channel had a different threshold based on its noise level (Equation 3.1). Grounded channels were disabled for spike detection. Scale bar: 1 ms, 100 μV. Site spacing: 65 μm.

and sample-and-hold corrected. Saved data were approximately 20× smaller than the original `.srf` data acquisition files, and could be relied on almost exclusively for subsequent spike sorting (`.srf` files were still required for substantial spike realignment, see Section 3.7). This allowed the rest of the spike sorting process to be much more portable.

During all of the preceding spike detection steps, a channel dependent temporal lockout was constantly enforced and updated. This lockout prevented multiple registrations of the same spike. For example, when testing a putative spike against the lockout, if a previously detected spike on a proximal channel extended up until or past the primary peak of the putative spike, the putative spike was locked out and hence rejected. Most likely, this putative spike consisted of one or more peaks belonging to the previously detected spike. However, keeping the lockout highly localized in space and time was also desirable to prevent the unintended lockout of distinct spikes. Therefore, a balance was struck in the spatiotemporal extent of the lockout, between minimizing multiple registrations of the same spike, and minimizing the loss of distinct spikes.

The lockout was updated after each detection of a spike. The channels included in the lockout were those that fell within $r_{lockout}$ of the spike's primary channel. The lockout time for each spike was different on each locked out channel. On the primary channel, the lockout was set to the time of the later of the spike's two adjacent peaks. The same held for the corresponding peaks on other channels within $r_{lockout}$. However, peaks on nearby channels were not always simultaneous in time (Figure 3.5). To find the corresponding peaks on nearby channels, each channel within $r_{lockout}$ was searched $\pm\, dt/2$ around the time of each peak on the primary channel, where $dt \leq dt_{max}$ is the time between the spike's adjacent peaks on its primary channel. If more than one peak was found on a given nearby channel within the given time range, the peak closest in time to that on the primary channel was chosen as the corresponding peak. If no peak was found on a given nearby channel within the required time range, the corresponding peak was declared to be at the same timepoint as that on the primary channel. When searching for corresponding peaks, sign was ignored. This allowed for polarity inversion across channels, which could occasional occur (e.g., green unit in Figure 3.10B).

Since sufficiently realistic simulations of extracellular cortical spike data are currently lacking (Section 3.12), manual inspection of hundreds of detected spikes and minutes worth of data was used to evaluate spike detection performance. No obvious spike detection errors were found, whether false positive or negative (but see arrows in Figure 3.4 for less obvious false negative candidates). Errors were generally only found during clustering (Section 3.9), in which many thousands of spikes were considered at a time. By nature, only false positives could be found during clustering. False positive spike detection error rates were $\sim 1\%$ of detected spikes. Many other apparent detection errors were simply spike alignment errors that were easily corrected (Section 3.7).

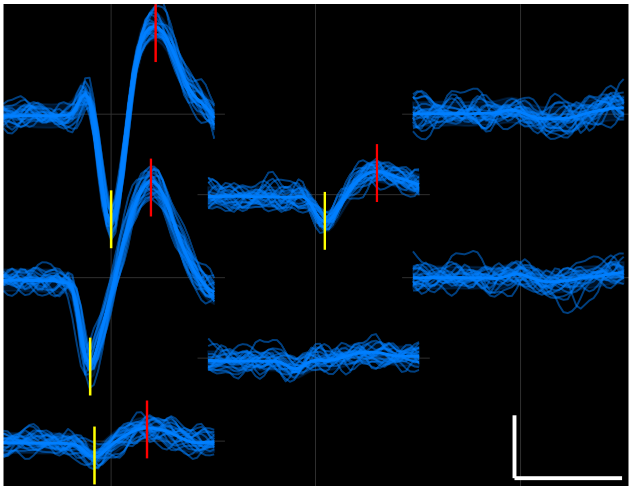

Figure 3.5: Spike propogation in space can have noticeable delays, resulting in non-simultaneous peaks on nearby channels. 20 spikes are shown from a cluster in ptc22.tr1. The top left channel is the primary channel (with maximum S_{pp}), and is also the top left site of the polytrode. Yellow markers denote the primary peak on each channel, and red markers denote the later secondary peak. Markers were aligned to the mean, but each spike deviated slightly. All channels within 150 μm of the primary channel are shown. The peaks of channels without markers were too small and inconsistent to mark in this display. Scale bar: 0.5 ms, 100 μV. Site spacing: 65 μm.

3.5 Spatial localization

Once detected, each spike was localized to a 2D position along the plane of the polytrode. A 2D Gaussian distribution was fitted to the spatial voltage distribution of each spike, namely V_{pp} calculated from the adjacent peak pair on each of the spike's included channels (Section 3.4). The Gaussian model's position was initialized to the V_{pp}-weighted mean position of the spike's channels. Although some cells' voltage distributions were clearly elliptical (e.g., Figure 3.13B, magenta and darker blue units), a circularly symmetric Gaussian distribution with a spatial extent of σ (with $\sigma \equiv \sigma_x \equiv \sigma_y$) was used to reduce the number of free parameters. Circular symmetry was imposed for model stability given the limited amount of spatial data provided by each spike, especially in the x direction which had at most only 3 values for a 3 column polytrode. σ was initialized to 50 μm, approximately the average spatial extent across all spikes (Figure 5.3). V_{pp} at any point (x, y) along the plane of the polytrode was modelled by

$$\hat{V}_{pp}(x, y) = Ae^{-\frac{(x-x_0)^2+(y-y_0)^2}{2\sigma^2}} \tag{3.3}$$

where (x_0, y_0) is the distribution's spatial origin. The amplitude A was fixed to V_{pp} of the primary channel, which was usually also the maximum V_{pp} of all the included channels. Keeping A fixed reduced the number of free parameters to 3.

The Levenberg-Marquardt (LM) least squares algorithm (`scipy.optimize.leastsq`) was used to fit the model to each spike. The algorithm was given V_{pp} for each of the spike's channels, along with the (x, y) coordinates of each channel. On every iteration, the free parameters x_0, y_0, and σ were modified to minimize the squared error between the model ($\hat{V}_{pp}(x, y)$) and the data ($V_{pp}(x, y)$). The final values of x_0, y_0, and σ were saved with the rest of the spike's information.

Spike location and amplitude is one way to reduce a $\sim$ 600 dimensional voltage space (Section 3.4) down to only three dimensions (x, y, V_{pp}), and initial attempts were made to cluster spikes in this space. However, this space on its own ignores all temporal information about spike shape. Spikes with very distinct shapes could share the same position in this space, and conversely, due to spatial drift over time, spikes with very different positions in this space could still belong to the same cluster (Figures 3.6 & 3.12). Therefore, the (x, y, V_{pp}) space was poorly suited for unit clustering. Nevertheless, spatial localization was useful for visualizing the spatial position and amplitude of spikes along the polytrode (Figure 3.6). It was also useful for estimating the spatial position of clusters based on the median of the spatial positions of their constituent spikes (Section 4.4), allowing clusters to be arranged in vertical spatial order.

3.6 Initial channel split

To create an initial set of clusters, the set of detected spikes was split according to each spike's primary channel. This resulted in up to 54 initial clusters, one per polytrode channel, numbered in order from top to bottom of the polytrode. This "divide-and-conquer" approach (Swindale and Spacek, 2014) to splitting the full set of spikes into a few dozen smaller ones was desirable for two reasons. First, a common set of channels within a given set of spikes was required before any dimension reduction (Section 3.8) could be applied to them. Second, clustering with the GAC algorithm (Section 3.9) takes up to $O(N^2)$ computational time, where N is the number of spikes. Splitting up the large initial group of unclustered spikes to reduce N within each group was therefore very computationally advantageous for clustering.

Splitting spikes by their primary channel resulted in only a very preliminary set of clusters. These were then further split into single unit, multiunit, and noise clusters, some of which were later merged (Section 3.10.2). For example, some units fell half-way between two channels, resulting in spikes whose primary channel alternated between two channels (Figure 3.3B). Also, clusters belonging to neurons drifting in time (Figure 3.12) were often initially oversplit, and were merged later on.

The alignment, dimension reduction, clustering, and verification steps that follow were performed iteratively on clusters ordered vertically in space, from top to bottom of the polytrode. Once a cluster at a given position on the polytrode was judged sufficiently clean and well clustered (Section 3.10), it was marked as "good", which highlighted it in green in the cluster list (Figure C.3). Attention was then shifted to the next cluster down the polytrode. This way, progress was made visually obvious to the user. Occasionally, after a few rounds of merging and splitting, clusters were renumbered to keep them spatially ordered from top to bottom.

3.7 Alignment

Before applying any dimension reduction to a group of spikes, spikes had to be properly aligned, otherwise subsequent clustering steps could split groups of spikes based on spike misalignment

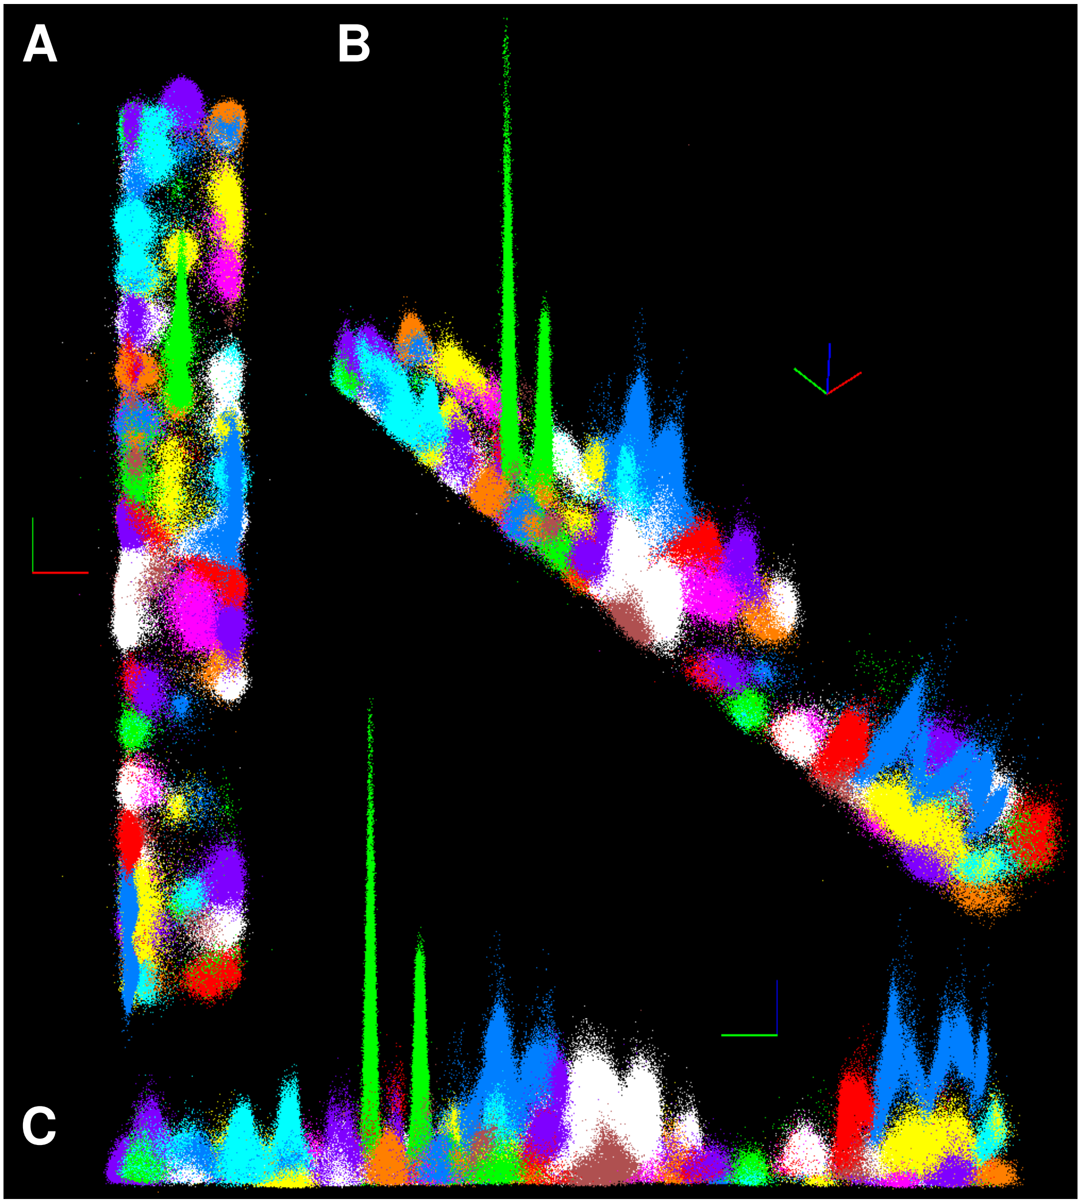

Figure 3.6: Spatial locations and V_{pp} of $\sim$ 1.9M spikes from track ptc22.tr1 after clustering (Section 3.9) and verification (Section 3.10). Spikes are plotted as points in a 3D (x, y, V_{pp}) space, represented by red, green and blue axes respectively. Each dimension is normalized by its standard deviation, except for y, which was normalized by the standard deviation of x to maintain aspect ratio. Positive y points down the polytrode, and V_{pp} of each spike is that of its primary channel. Points are plotted with perspective, from different views: ***A***: above the plane of the polytrode; ***B***: obliquely; and ***C***: from the side. Colours represent distinct clusters, with 10 colours cycling along the y axis. In this space, what appears to be a single cluster can sometimes be multiple clusters when plotted in a more discriminating space, such as PCA (Section 3.8). Alternately, distinct clusters in this space (e.g., the high amplitude bimodal green cluster) can in fact be a single cluster that has drifted in position (Figure 3.12A). Gaps of points are due to faulty sites or an excess of noise events or multiunit spikes. Axes bars represent $\sigma_c = 0.43$ in each dimension (see Section 3.9.1).

instead of genuine differences in spike shape (Figure 3.7). Since all spikes were aligned to their negative peak during detection, and since most spikes consist of a sharp negative peak followed by a wider positive peak, most spikes within a given cluster were already reasonably well aligned to one another. However, some cells had spikes whose primary channel alternated between two or more channels, due to noise or drift. Some spikes had two similarly sharp negative peaks, while other spikes were predominantly positive-going. Any of these cases could lead to alignment errors. Even modest misalignment could result in artifactual clusters (Figure 3.7). Larger misalignments could distort cluster distributions (Figure 3.8), which could further complicate downstream clustering steps.

With spikes split into an initial set of clusters (Section 3.6), they could now be aligned to one another by best fit, instead of by the (often noisy) features of individual spikes. To realign spikes by best fit, the desired cluster was selected, as were the channels and time range on which to calculate waveform differences. Usually, only the top one or two largest amplitude channels were selected. On the selected channels, the mean waveform of all of the cluster's spikes was calculated. Then, for each spike, the sum of squared differences was calculated between the spike's waveform and the mean waveform. The sum of squared differences was used instead of the RMS difference in order to decrease calculation time. Each spike was then shifted up to 2 interpolated timepoints in either direction ($\pm$ 40 μs), for a total of 5 possible positions (including no shift). The sum of squared differences was calculated at each position, and the spike was shifted to the position that gave rise to the smallest sum of squared differences. The user typically executed this process 2 or 3 times, the mean waveform updating each time. After a few best fit realignments, the number of realigned spikes per realignment approached zero, and the standard deviation of the spike waveforms settled at a minimum value.

Note that the above best fit realignment process could be performed on any desired selection of spikes, whether selected individually or by selecting one or more clusters. The only requirement was that all selected spikes shared at least one common channel. However, realignment was most often performed on one single cluster at a time.

When shifting spike waveforms, missing end points were replaced automatically, such that the last real value at a given edge was repeated the appropriate number of times to pad the waveform to the proper width (50 timepoints for the usual case of 1 ms waveforms interpolated to 50 kHz). Though not often necessary, this edge data could be replaced by reloading the actual data from disk. For the occasional spikes that were realigned many times and were shifted by a large amount, this was an important feature. Padding with edge data and only occasionally reloading made for much faster realignment and decreased the need for access to the continuous waveform data on disk.

After spike detection was complete, the spatiotemporal lockout for each spike (which only made sense in continuous time) could no longer be applied. Yet, given enough best fit realignments, it was possible to violate previously observed lockouts and shift a detected spike far enough that

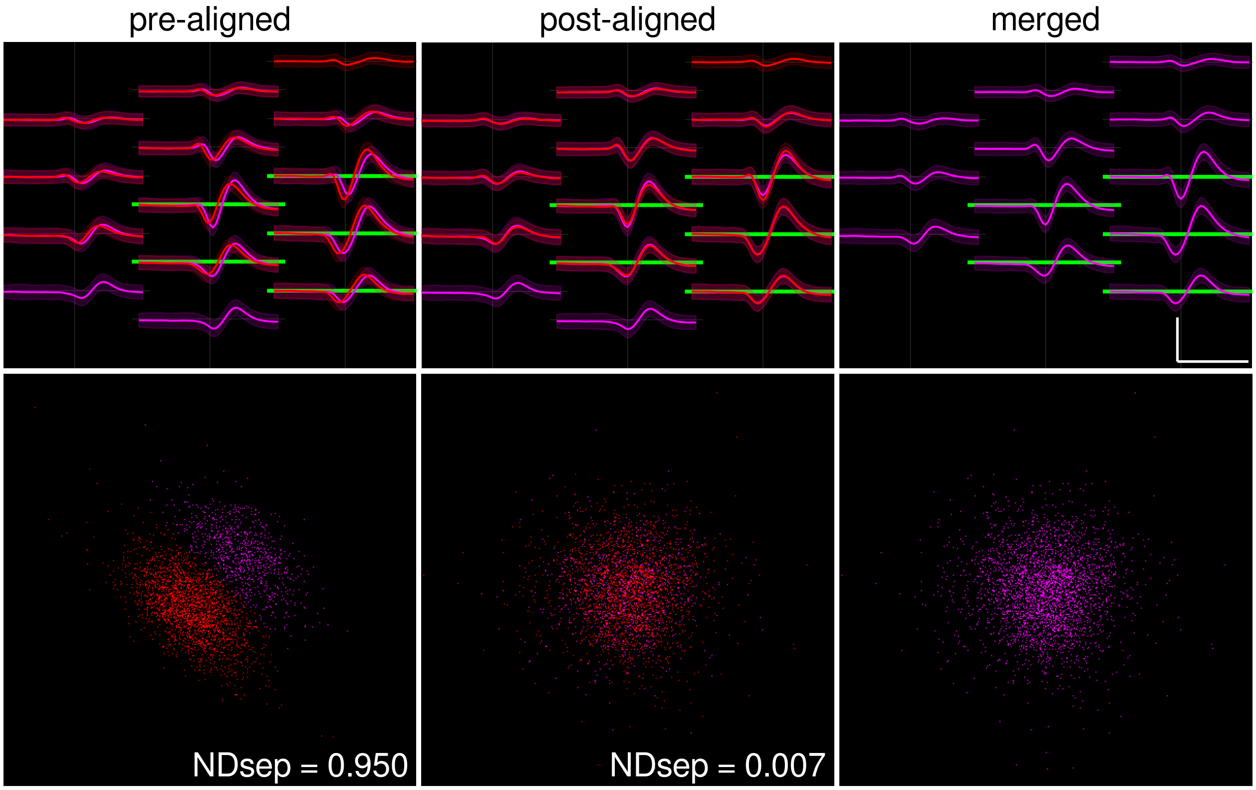

Figure 3.7: Spike misalignment can result in artifactual clusters. ***Left***: Two clusters are plotted, with mean waveforms above and cluster plot below. Cluster plots are based on the first 3 principal components (PCs) (Section 3.8) calculated from the 5 selected green channels. The first cluster plot is rotated to show maximum cluster separation, and the rest maintain the same orientation. In cluster space, the pair are distinct as indicated by their high NDsep value (Equation 3.9), but the similarity of their mean waveforms suggests that they may instead be a single cluster split into two clusters of slightly different alignments. ***Middle***: The two clusters realigned to each other, using the sum of squares best fit of each spike's waveform to the mean waveform of all the spikes from both clusters, measured on only the 5 selected channels (highlighted in green). Realignment was performed 3 times, with fewer spikes shifted each time. Separation in cluster space disappeared, suggesting the pair were indeed oversplit. ***Right***: The two realigned clusters merged into a single cluster. Scale bar: 0.5 ms, 100 μV.

it erroneously merged with another one nearby in space and time. This could happen due to high initial misalignment resulting in the lack of a local minimum in the sum of squared differences. During best fit realignment, spikes were shifted by no more than 40 μs at a time to help prevent such spike mergers. Given enough realignments, even this did not provide complete protection. However, it was generally obvious from the plotted waveforms and the displayed realignment statistics when spikes were shifting too far or in the wrong direction. In such a case, the user would stop any further realignment, and upon clustering, the now grossly misaligned spikes would split off as a cluster, which could be dealt with separately (Figure 3.8, middle). It was also possible to realign spikes to

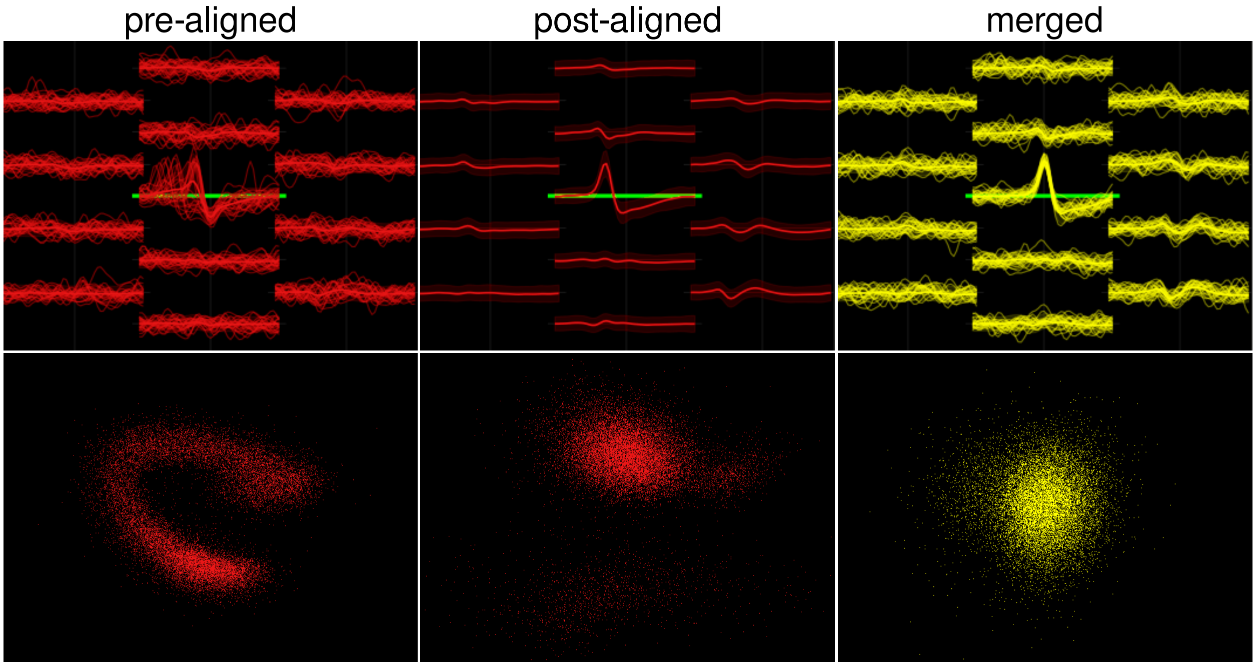

Figure 3.8: Spike misalignment can result in non-Gaussian clusters. ***Left***: Above, 20 randomly selected spikes from a cluster are plotted over top of the cluster's mean waveform (thicker red line). Despite all being aligned to their negative peak by default, many of the cluster's spikes were misaligned due to being predominantly positive-going (Figure 3.3C). Below, the resulting non-Gaussian distribution of the cluster's points, based on the top 3 PCs (Section 3.8) of the selected channel. ***Middle***: After best fit realignment of waveforms on the single selected channel (green), the mean waveform increased in amplitude and the main cluster became much more Gaussian. A second sparser cluster below the main one was the result of spikes that began so far out of alignment that best fit realignment failed to properly realign them. This sparser cluster was split off and realigned with some manual intervention (alignment of each spike to its positive peak, then best fit of spikes within the cluster, then best fit with the main cluster). ***Right***: The sparser cluster was then merged back into the main cluster. The final cluster distribution is much more Gaussian than the original.

either of their two original peaks (positive or negative), or to manually shift a selection of spikes left or right by a desired amount. Once spikes were manually shifted closer to their (subjectively) proper alignment, best fit realignment could once again be used.

3.8 Dimension reduction

Given that each spike inhabited an impractically large ~ 600 dimensional space (Section 3.4), the number of dimensions had to be greatly reduced before clustering. Most spike sorting methods that rely on clustering in a dimension reduced space use only 2 dimensions (Lewicki, 1998). Since 3 dimensions are the most that can be easily visualized, and since any more than the top 3 dimensions from either of the reduction methods below (PCA or ICA) are unlikely to improve clusterability, spikes were reduced to a 3D space before clustering.

One way to reduce dimensionality is with PCA (Shlens, 2009). Neighbouring waveform data values are significantly correlated in space and time. PCA can take such a correlated data set and represent it in terms of orthogonal (principal) components that are uncorrelated. These orthogonal components can then be ranked according to the amount of variance of the original data that they explain. For Gaussian distributed data, typically only the top few components are needed to accurately reconstruct the data (Shlens, 2009). This allows for a great reduction in dimensionality. Note however that the goal here is not to reconstruct the data in a low dimensional space with maximal accuracy, but rather to reduce the data to a space that maximally reveals clusters.

First, for a given selection of spikes, such as those belonging to the initial channel-based clusters (Section 3.6), the relevant channels were selected either automatically or manually. For automatic channel selection, the median of the modelled spatial positions of the selected spikes (Section 3.5) was calculated. Then, a radius r was calculated around this median position such that r encompassed 95% of the modelled spatial σ values of the selected spikes. All channels that fell within r of the median spatial position were then automatically selected. Channels could also be manually toggled by simply clicking on them.

The relevant time range on the selected channels could be left at the full 1 ms width, or manually reduced in 0.1 ms increments to include only the higher amplitude parts of the spike waveforms, around the peaks. In either case, the chosen time range was always centered on the spike waveforms (i.e., at $t = 0.5$ ms of each spike's saved waveform data). Generally, the most relevant channels and time ranges were those with the greatest signal and/or the greatest signal variance (Figure 3.9). Discriminative channels and time ranges could also be revealed by plotting random subsets (e.g., 20) of the selected spikes at a time, and watching which channels and time ranges had signals that varied the most between subsequent random samplings. Selected channels and time ranges were highlighted in green (Figures 3.7, 3.8 & 3.9, all showing full 1 ms width time range selection). Ideally, channel and time range selection excluded waveform data that were not useful for distinguishing spikes, such as channels and times during which spike waveforms had little or no signal. Restricted channel and time range selection also reduced the dimensionality of the inputs to PCA and ICA, allowing them to run faster.

Next, for each spike, the waveforms on the selected channels in the selected time range were concatenated into a single n_{dim} long row vector, where n_{dim} is the product of the number of selected channels and timepoints. Row vectors from all selected spikes were combined to form an $n_{spikes} \times n_{dim}$ input data matrix. PCA was performed on this matrix, and the scores of the top 3 PCs were used to characterize each spike.

Another way to reduce data dimensionality is with ICA (Comon, 1994; Bell and Sejnowski, 1995; Hyvärinen et al., 2001). Unlike PCA, ICA considers higher-order statistics instead of just second order statistics. ICA does a better job of separating out components of non-Gaussian distributed data by searching for components of maximum independence. These components are

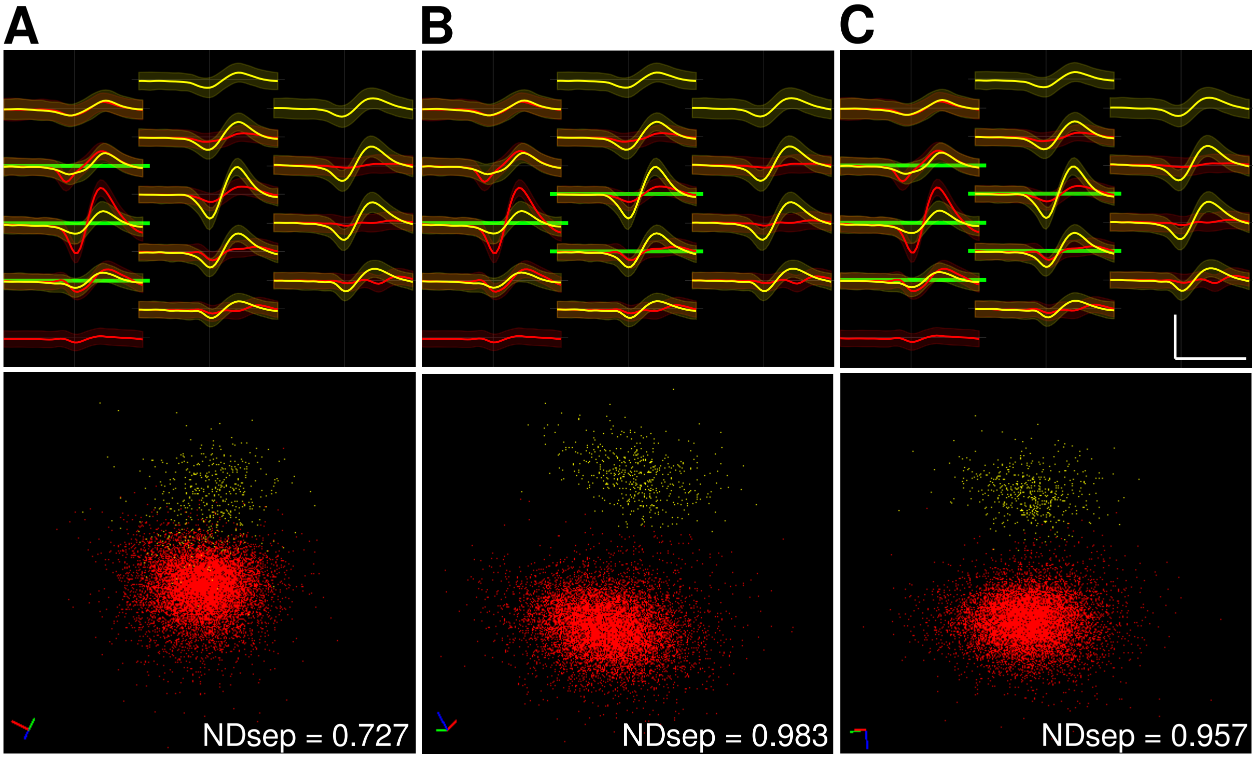

Figure 3.9: Channel selection affects clusterability in a dimension reduced space. This is shown here by comparing two existing clusters (red and yellow), but applies equally well to unclustered spikes. ***Top***: Mean cluster waveforms. Channels and time ranges selected for dimension reduction are highlighted in green. In this example, the full 1 ms time range was selected on each channel, represented by the full length of each green line. Shorter selected time ranges, symmetric about $t = 0.5$ ms (vertical lines), were represented by proportionally shorter green lines. ***Bottom***: 3D PCA cluster plots. PC0, PC1 and PC2 are along the red, green and blue axes respectively. Each cluster plot is rotated to show maximum cluster separation. The optimal manual channel selection in ***B*** resulted in a much more clusterable space than in ***A***, as denoted by the NDsep values (Equation 3.9). The optimal channels were usually those that showed the greatest signal variance. ***C***: Automatic channel selection resulted in separation nearly as good as the optimal manual channel selection in ***B***. Note that even though the channels selected in ***C*** are the union of those in ***A*** and ***B***, separation in ***B*** is better. This shows that, at least in a dimension reduced cluster space, some channels can worsen clusterability. Scale bar: 0.5 ms, 100 μV.

not necessarily orthogonal to each other, as they are in PCA. Whereas PCA can only rotate its entire set of components together, ICA can also change the angles between its components. Various flavours of ICA exist. The method used here was FastICA (Hyvärinen, 1999; Hyvärinen and Oja, 2000), implemented in the MDP Python library (`mdp.fastica`).

Prior to performing ICA, the $n_{spikes} \times n_{dim}$ input data matrix was prepared the same way as described above. But because ICA is much more computationally expensive than PCA, PCA was first applied to partially reduce the dimensionality (Scholz et al., 2004). The top $7 \times n_{chans}$ PCs, where n_{chans} was the number of channels selected, were kept to create a reduced input data matrix. ICA was then performed on this reduced matrix. A factor of 7 was chosen as a good trade-off between speed and cluster separability. Unlike PCs, independent components (ICs) do not come with associated scores. However, components can be scored *post-hoc* by various means. In this case, they were scored by their kurtosis (`scipy.stats.kurtosis`), a 4th order statistical measure of the "peakiness" of a distribution. Kurtosis is one way of measuring how non-Gaussian a distribution is. By sorting by kurtosis, the components were sorted by how non-Gaussian they were, and therefore by their statistical independence from each other (Scholz et al., 2004).

ICA was especially useful for separating clusters of vastly different numbers of points (size ratios > 50, Figure 3.10A & B). The small clusters in such pairs could consist of spikes, often belonging to some other existing larger cluster, or noise events, which were discarded. Conversely, and surprisingly, ICA was dramatically worse than PCA at separating similarly sized clusters (size ratios < 10, Figure 3.10C, see Section 3.13.2 for discussion). Therefore, PCA was generally used first to reveal large clusters within a selection of spikes, and ICA was run later on each of those clusters to check for contamination (Section 3.10.1).

In addition to the top ranked PCs or ICs, other parameters could also be chosen as one or more dimensions to plot in 3D: spike spatial parameters x_0, y_0 and σ; peak-to-peak voltage V_{pp}; the time dt between the peaks of a spike; and spike time t. Of these, only t was found to be useful for clustering. The rest were all encapsulated by PCA and ICA of spike waveforms. Though not easily visualized, clustering in more than 3 dimensions was explored but did not show any practical benefit. Wavelet decomposition (Quian Quiroga et al., 2004) and wavelet packet decomposition (Hulata et al., 2002), extensions of Fourier decomposition which consider both signal frequency and time, were also explored (using the PyWavelets library, http://www.pybytes.com/pywavelets) but did not prove very useful (not shown). Although any 3 parameters could be plotted in any combination, all clustering was done either in pure PCA or ICA space, or the top two PCs or ICs were plotted vs. t to handle drift. In practice, all clusters were examined in four different 3D spaces: pure PCA, pure ICA, the top two PCs vs. t, and the top two ICs vs. t. Each plotted dimension was rescaled such that it had zero mean and unit variance. Spatial parameters x and y were both normalized by the variance of x to maintain a constant spatial aspect ratio.

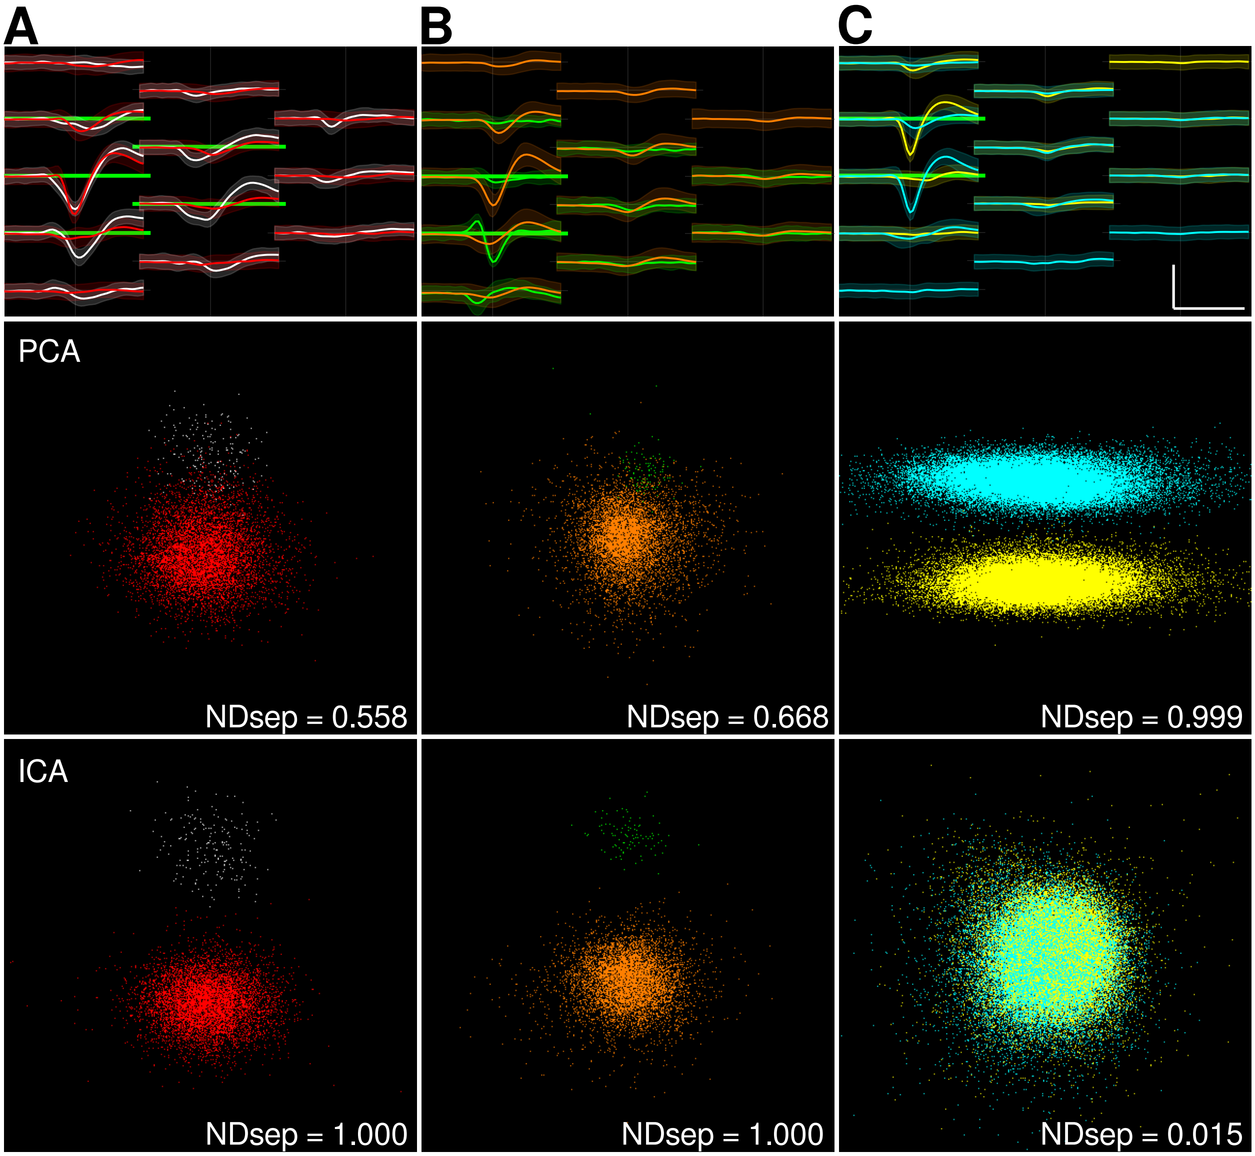

Figure 3.10: ICA separates clusters with large size ratios better than PCA. ***Top***: Mean cluster waveforms, with channels selected for dimension reduction highlighted in green. Clusters are plotted in PCA space (***middle***) and ICA space (***bottom***). Each plot is rotated to show maximum cluster separation. NDsep values (Equation 3.9) quantify the separation of each cluster pair. ***A***: A pair of clusters with 6970 (red) and 190 (white) spikes each. Because there were so few points in the white cluster, the mean waveform of the red cluster was essentially identical to that of all the spikes from both clusters. In PCA space, the smaller white cluster is encompassed by the tail of the larger red one. Their distinctness is revealed only in ICA space. ***B***: Another example cluster pair with a large size ratio (orange: 6772 spikes, green: 101 spikes). ***C***: A cluster pair (same as in Figure 3.12D) with similar numbers of spikes (yellow: 41376; cyan: 36633). In this case, because the clusters were not drastically different in size, PCA did a much better job of separating them than ICA (but see Section 3.13.2). Scale bar: 0.5 ms, 100 μV.

3.9 Clustering

With an appropriately chosen low dimensional space, spikes were ready to be partitioned into clusters representing putative neurons. Initial attempts at manual clustering involved drawing 3D ellipsoids. This generally allowed the user to visually cluster the data as desired, but it incorrectly imposed ellipsoidal shapes with hard boundaries. Just as critically, it proved far too laborious and slow, and was subject to more user bias than an automated algorithm. However, occasional manual "painting" of points in 3D with the mouse was used to quickly examine and potentially split off small numbers of outlier points from a large cluster.

All of the clustering algorithms described in Section 3.1.2 were attempted, but with limited success. HC was tested using the `scipy.cluster.hierarchy` and `hcluster` Python libraries, but was eventually rejected due to poor results. EM+GMM was tested using the method and MATLAB code of Bar-Hillel et al. (2006), but was found to be very slow, especially for more than a few thousand spikes and a few clusters. Also, splitting the data into time chunks resulted in the loss of neurons that fired very sparsely. SPC was tested using software from Blatt et al. (1996) and Quian Quiroga et al. (2004), but often an ideal temperature could not be found to satisfactorily cluster the data presented here (e.g., Figure 3.11). Clusters were either oversplit or undersplit. Also, SPC was sensitive to narrow, low density bridges of points between clusters, and would incorrectly combine such clusters (Swindale and Spacek, 2014). SPC is somewhat insensitive to density variation, and this was found to be a drawback rather than a benefit.

With all existing methods, it was often difficult or impossible to cluster the points in the way the user desired. In 3 dimensions or less, obvious (and often non-Gaussian) clusters were often visible in the data that none of the algorithms could easily find.

3.9.1 Gradient ascent clustering

Instead, a new clustering algorithm, GAC, was developed to mimic the clustering a human user does by eye. Visually we look for local peaks in the density of points, at the appropriate spatial scale. GAC iteratively clusters data according to local density peaks, at the chosen spatial scale. Before iteration began, all points were first duplicated into a set of "scout" points. During iteration, the original data points remained motionless, while the scout points were allowed to move and merge with one other.

In the first step of each iteration, the local density gradient of data points around each scout point was calculated, and each scout point was moved up its gradient by an amount relative to the gradient's magnitude. The update step vector $\Delta\mathbf{s}$ of scout point i's n_{dim}-dimensional position $\mathbf{s}_i$

was

$$\Delta \mathbf{s}_i = \frac{\alpha \sum_{j=1}^{n} \mathbf{d}_{ij} f_{ij}}{\sum_{j=1}^{n} f_{ij}} \tag{3.4}$$

where α is the gain of the step vector ($\alpha = 2.0$ was used), $\mathbf{d}_{ij}$ is the distance vector between scout point i and data point j, f is the spatial kernel used to define locality, and n is the number of data points in the neighbourhood of the scout point. The summation in the numerator represents the local density gradient, normalized by the summation in the denominator. Although n could potentially include all N data points in the cluster space, in practice given a kernel characteristic scale σ_c, there was no benefit in setting $n = N$. Instead, for computational efficiency n was limited to data points that fell within a radius $r_{neigh} = 4\sigma_c$ of the scout point. The chosen spatial kernel was an n_{dim}-dimensional Gaussian distribution,

$$f_{ij} = e^{-\frac{d_{ij}^2}{2\sigma_c^2}} \tag{3.5}$$

where d_{ij} is the Euclidean distance between scout point i and data point j. As mentioned previously (Section 3.8), $n_{dim} = 3$ was used.

The second step of each iteration was the merge step, during which the actual clustering occurred. During the merge step, a search was performed around each scout point for other nearby scout points, merging those that were sufficiently close to one another. The threshold for this decision was a merge radius set to a fraction of σ_c, $r_{merge} = 0.25\sigma_c$. For ease of bookkeeping, for any pair of scout points to be merged, the scout point with the higher index was always merged into the scout point with the lower index. On each merge, the original index of each scout point and the new scout point it had been merged into were both kept track of, such that when the algorithm exited, each original data point (from which all scout points originated) had an index that assigned it to one of the remaining scout points.

The completion of both steps, gradient ascent and merge, constituted a single iteration of the algorithm. On each iteration, the number of scout points generally decreased, while each remaining scout point neared its local density peak. There were two convergence criteria, either of which was sufficient to signify convergence and cause the algorithm to exit. The first kept track of how many iterations in a row had occurred without a single merger of any scout points. If this reached 1000, the algorithm would exit. The second convergence criterion kept track of how far each scout point moved during the update step. Scout points that moved less than $10^{-5}\sigma_c$ on any iteration were deemed to have stopped moving, and were marked as so. Scout points marked as stationary could be skipped during the update step, and pairs of stationary scout points could be skipped during the merge step, reducing computational time. If at any point all of the scout points stopped moving,

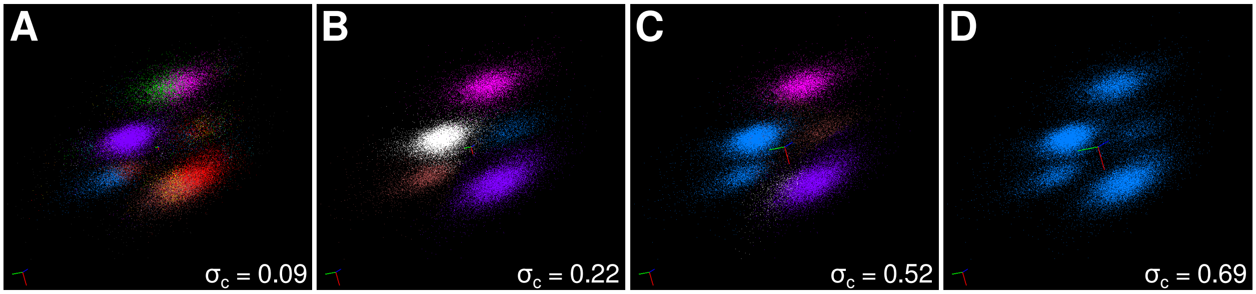

Figure 3.11: GAC dependence on σ_c. In this example 50,000 points are plotted in 3D PCA space, rotated to show maximum separation. Five distinct clusters are visible to the eye. Clustering results from left to right ranged from oversplit to undersplit, depending on the user-set value of σ_c. Of the four values, $\sigma_c = 0.22$ ($\boldsymbol{B}$) gave the best result in this example. Results were consistent within about ± 25% of that value. Two axes are visible in each panel. The size of the central axes corresponded to σ_c, allowing the user to better gauge its optimal value according to the separation of visible clusters.

the algorithm would exit. Upon convergence, each remaining scout point represented the center of a cluster. All scout points that had been merged into a particular remaining scout point shared the same index and belonged to the same cluster. Any clusters with less than 5 points were marked as unclustered outliers by assigning them a cluster index of 0.

The only parameter adjusted by the user on a regular basis was σ_c, which defined the spatial scale of cluster separation judged to be relevant to the data. Low σ_c values yielded many local density peaks, corresponding to many small clusters. High σ_c values yielded fewer local density peaks, corresponding to only a few large clusters (Figure 3.11). A single σ_c value was used uniformly in all dimensions. This was feasible because the data were normalized in all dimensions during dimension reduction (Section 3.8). Importantly, this limited the number of parameters requiring user adjustment to only one.

The normal range of values for σ_c was $0.1 < \sigma_c < 0.9$, and the most common value was $\sigma_c \approx 0.4$. The value of σ_c chosen by the user was visually guided by the spatial structure of the data plotted in the 3D cluster space. A set of axes whose size represented σ_c was scaled until it was roughly the size of the minimum separation between clusters. Using the mouse or keyboard, the cluster space could be zoomed, translated, and rotated at will, and the axes could be independently positioned and sized relative to the data (Figure C.3). The desired clustering, as evaluated by eye, was usually accomplished on the first attempt, but if need be σ_c adjustment and execution of GAC could be done iteratively until the desired clustering result was obtained (Figure 3.11B). Undersplitting was corrected by decreasing σ_c, and oversplitting by increasing it. The user could undo or redo any number of clustering operations, including cluster splitting, merging and deletion. This allowed for easy before and after comparison of GAC results. GAC performed fast enough (Figure C.4) for interactive manipulation of σ_c.

Clustering results were usually consistent within a range of σ_c values. For example, the clustering result in Figure 3.11B was tolerant to within about ± 25% of its optimum value. This meant that for a given set of points in a given cluster space, there was rarely a need to finely optimize σ_c. An automated method of choosing σ_c, which searches for a range of values over which the number and size of found clusters remain invariant, is described in Swindale and Spacek (2014). However, because spyke is heavily focused on user interactivity, this automated σ_c selection method was not implemented here.

The α parameter controlled the rate of gradient ascent. If set too high, scout points could overshoot and then oscillate around the local density peak. If set too low, the algorithm would take excessively long to converge and might even exit prematurely. A good compromise was found by setting $\alpha = 2.0$.

Although devised independently, it was later found that GAC is a modification of the mean-shift clustering algorithm (Fukunaga and Hostetler, 1975). The chief difference between the two is that in GAC, scout points are merged on each iteration, greatly speeding up the algorithm by rapidly decreasing the number of scout points to iterate over.

3.10 Cluster verification

Performing dimension reduction on a channel-based cluster and running GAC on the resulting set of points was only the first step. Despite appearances, there was no guarantee that the newly found clusters were indeed unimodal, distinct, and free of noise events and misassignments. Cluster verification proved to be the most lengthy and laborious step.

There were three major issues to deal with during cluster verification: undersplitting, oversplitting, and misassignment. The outcome of cluster verification was the classification of each cluster as either single unit or multiunit. This classification was made by the user, and was based on both subjective and objective measures. Classification of each cluster depended on its unimodality, its distinctness from other clusters in 3D cluster space, and the consistency of its plotted spike waveforms. If a cluster was mostly composed of what were obviously noise events, or was an unsplittable collection of noise and spike-like events, a third option was also available: deletion. However, no detected spikes (whether genuine or noise events) were ever truly deleted. Instead, events were discarded by assigning them to a special cluster with index 0. Single unit and multiunit clusters were assigned positive and negative cluster indices, respectively. Given sufficient user confidence in the isolation quality of a single unit cluster, it would also be marked as “good” for export for subsequent spike train analysis (Chapters 4–6). Any time such a cluster’s list of member spikes changed due to splitting or merging, its “good” flag was cleared to indicate that it needed to be re-verified. In practice only a subset of single unit clusters were judged to be adequately isolated and were exported. All other clusters, both single unit and multiunit, were kept for potential future refinement but were not exported. Only about half of all threshold crossing events accepted as plausible spikes during spike detection (Section 3.4) were exported (Table 4.1).

3.10.1 Undersplitting

Each cluster was tested for undersplitting to ensure that it was unimodal and consisted of spikes of uniform shape, with as few noise events as possible. Upon choosing the most relevant channels and time range (Section 3.8), each cluster was examined for multimodalities in all four of the critical 3D cluster spaces: pure PCA, pure ICA, and the top two PCs and ICs vs. spike time t. Multimodality in any one of these spaces suggested that the cluster was undersplit, but was not necessarily definitive. For example, a rapid but smooth and continuous shift in waveform shape could result in a multimodal cluster in pure PCA space, but when plotted against time, was revealed to much more likely be the same unit undergoing rapid drift (Figure 3.12A–B). Unfortunately, such situations required user judgement, precluding automation.

Besides visually searching for multimodalities and running GAC to see if it could split off clusters in the chosen 3D space, cluster quality could be further subjectively evaluated by plotting subsets of a cluster's spikes (e.g., 20 at a time), and progressing from one overlapping subset to the next (or previous), in temporal order. This sliding spike selection window allowed the user to observe how the waveform shapes changed (or did not) over time, and whether those changes were gradual enough to suggest that they were due to long-term spike shape variability, such as from drift, and not due to the erroneous merger of separate units.

In a given 3D cluster space, if multiple clusters were visible, it was not necessary to separate them all in a single run of GAC (although this is shown in Figure 3.11B). Splitting off only one cluster at a time was a more effective strategy, and was often (but not always) possible by choosing an appropriate value for σ_c. Each time a cluster was split off, dimension reduction was separately reapplied to not only the remaining points, but also to the cluster that had been split off. This was repeated until only unimodal clusters remained. Splitting off only one cluster at a time allowed PCA and ICA to maximally separate the remaining points, making it easier to subsequently split off further clusters.

Because ICA was better than PCA at separating clusters of greatly differing population sizes (Section 3.8), and because within a given cluster, noise events were typically much fewer in number than actual spikes, ICA was better than PCA at separating noise events from spikes. Therefore, most cleaning (i.e., removal of noise events) was done in ICA space. After splitting off a suspected noise cluster, like any other cluster it was examined in 3D cluster space, and its mean waveform and random samples of its spike waveforms were plotted. If the noise cluster was clearly unimodal, and was composed of inconsistent and clearly not spike-like events, it was deleted. If it was unimodal with inconsistent waveforms, but had at least some spike-like events, it was labelled multiunit. Sometimes what seemed like a noise cluster was instead clearly a set of spikes that had simply been misaligned (Figure 3.8, middle). Instead of discarding them as noise events, these spikes could be realigned to one another (Section 3.7) and then potentially merged back into an existing cluster (Section 3.10.2).

When splitting off cluster outliers, instead of always having to run GAC and find an appropriate

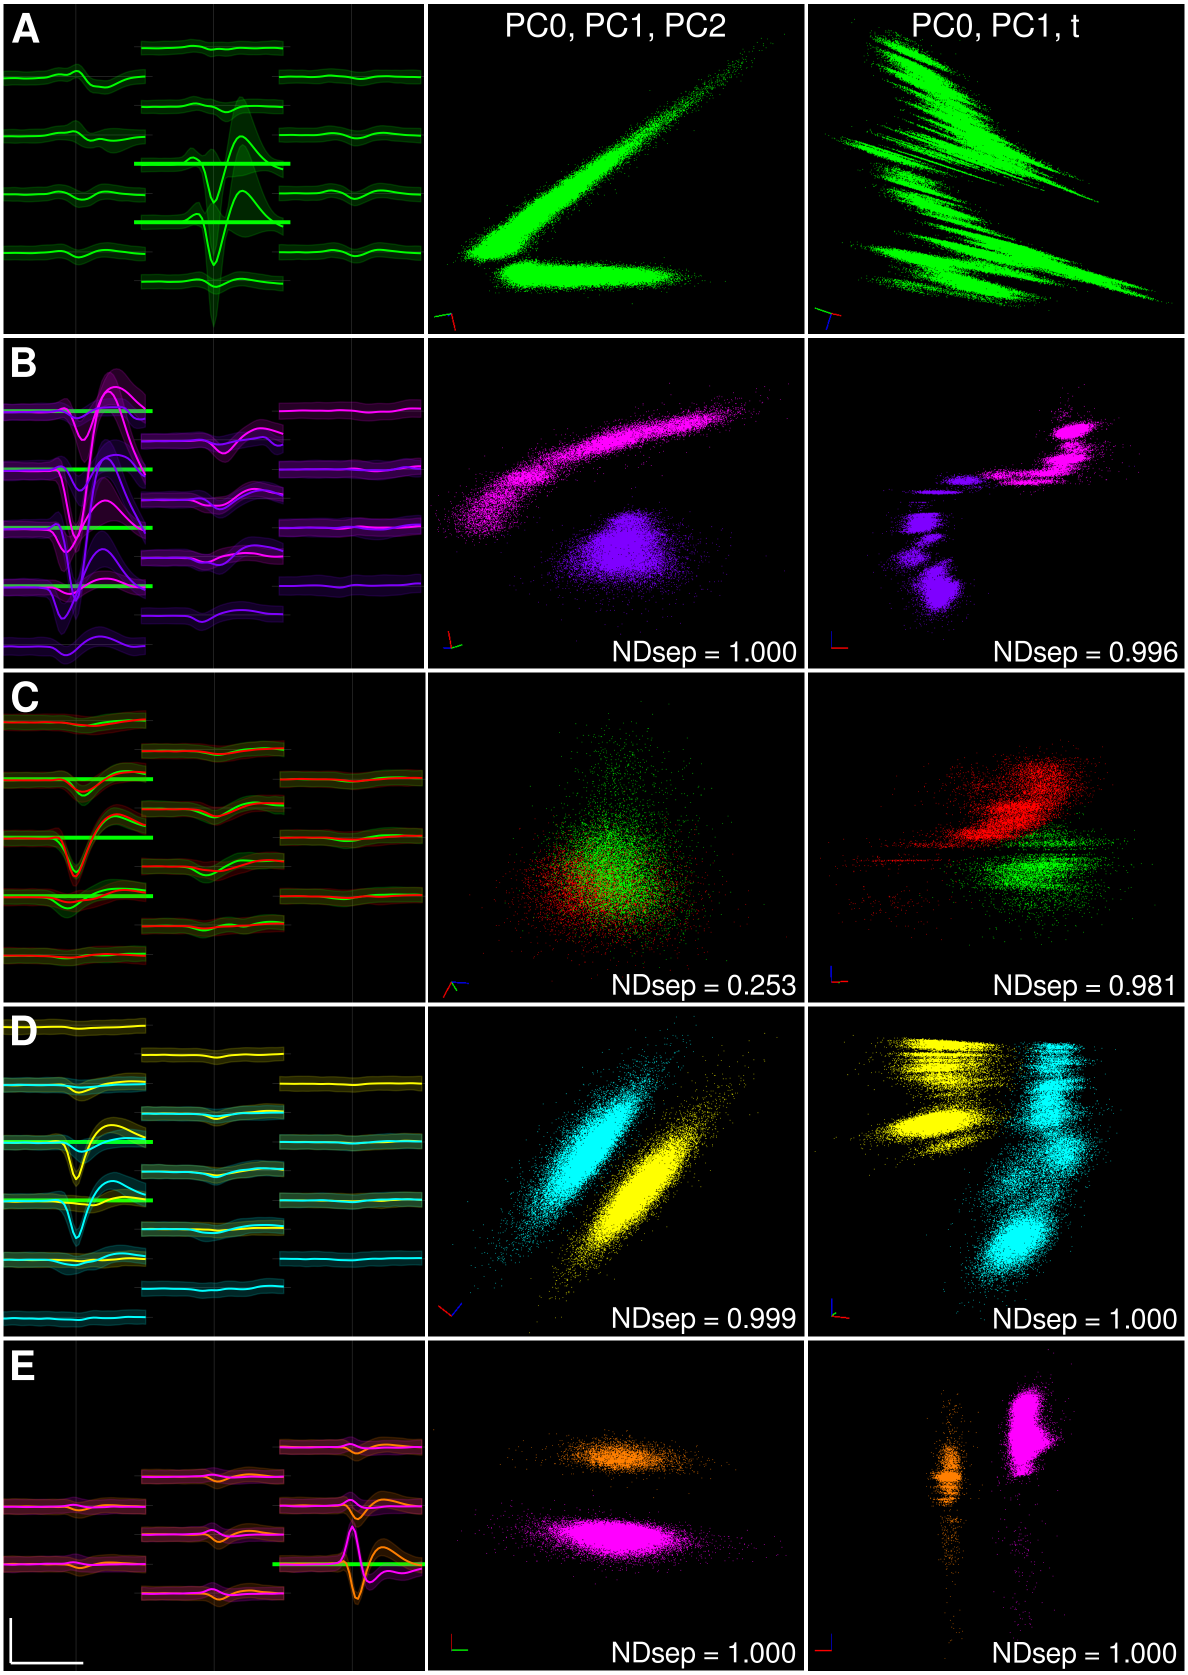
A
PC0, PC1, PC2
PC0, PC1, t
B
NDsep = 1.000
NDsep = 0.996
C
NDsep = 0.253
NDsep = 0.981
D
NDsep = 0.999
NDsep = 1.000
E
NDsep = 1.000
NDsep = 1.000

Figure 3.12: (*Previous page.*) Plotting clusters in time is necessary to check for drift. Sometimes what appeared to be two distinct clusters (***A***–***B***) were only a single cluster that had a relatively sudden drift event. Conversely, sometimes what looked like a single cluster in pure PCA space was in fact two neighbouring clusters drifting together in time (***C***). The channels selected for PCA are highlighted in green (***left***). For each example, cluster plots are shown in pure PCA space (***middle***: PC0, PC1, PC2) and PCA vs. time (***right***: PC0, PC1, t). Red, green and blue axes point along the first, second and third cluster dimensions respectively. Stratification in time (vertical blue axis, ***right***) was visible because of periods of higher or lower activity or brief periods in between recordings making up the track. ***A***: This cluster corresponds to the large amplitude green double-peaked cluster in Figure 3.6. High amplitude variability on the two selected channels was apparent in the large transparent ± 2 standard deviation limits. In pure PCA space, the cluster appeared undersplit, but plotting in time strongly suggested that it was indeed a single cluster, with high variability superimposed on slow drift. ***B***: An erroneously split pair of clusters. In pure PCA space, they were completely distinct, but when plotted against time, the two were obviously a single cluster that had drifted abruptly over a relatively short period of time. The two were later merged. ***C***: Two simultaneously drifting clusters that could only be distinguished when plotted against time. Their average waveforms looked nearly identical. NDsep values (Equation 3.9) quantified the difference in separability. ***D***: A drifting pair (same as in Figure 3.10C) of completely distinct neighbouring clusters. ***E***: A non-drifting pair of completely distinct neighbouring clusters.

value for σ_c, sometimes it was quicker and simpler to manually select points in the 3D cluster space by "painting" them with the mouse. The waveforms of all selected spikes were automatically plotted, and could be inspected before being split off from the parent cluster. If split off, they could then be further inspected, realigned, and cleaned. Points at the fringes of a cluster were routinely selected by painting them, and their waveforms were examined to ensure that they were not notably different from the cluster's mean waveform.

3.10.2 Oversplitting

Every cluster was also tested to ensure that it was measurably distinct from all other clusters, and was therefore not oversplit. The user could quickly compare a cluster to all others near enough in space to share at least one channel with the cluster, in order of increasing mean waveform RMS difference. The mean waveforms of each pair were overplotted, and if a pair looked suspiciously similar, or their RMS difference on the selected channels was $< 10\ \mu$V, or they scored poorly on one of the cluster pair separation metrics (see below), the pair were potentially oversplit. To investigate further, the relevant channels and time range were selected, either manually or automatically (Section 3.8), for realignment and dimension reduction. If after realignment, no combination of channel and time range selection and dimension reduction clearly separated the two clusters, they were merged. This new cluster in turn also had to be compared to all other nearby clusters. Checking for cluster distinctness could therefore be a laborious task.

In addition to mean waveform RMS difference, cluster pair separation was quantified in four ways: Jensen-Shannon divergence ($\mathrm{D_{JS}}$), 1D overlap area ratio, one-dimensional separation metric (1Dsep) and N-dimensional separation metric (NDsep). The first three required projecting the cluster pair onto a line. The line onto which their points were projected was formed by connecting the centers of the two clusters, as calculated by the median positions of their member spikes in the chosen cluster space. Another 1D measure of separation is Fisher's linear discriminant (Hill et al., 2011). It measures cluster separation by taking the ratio of the intercluster variance to the sum of the intracluster variances, along a projection line that considers both the means and the covariance matrices of the two clusters. However, this measure assumes Gaussian clusters, and was therefore not used.

$\mathrm{D_{JS}}$ (Schneidman et al., 2006) is an information theoretic measure (in bits) of the separation of two 1D distributions, Q and P, defined by

$$D_{JS}(P,Q) = \frac{D_{KL}(P,M) + D_{KL}(Q,M)}{2} \tag{3.6}$$

where

$$M = \frac{P+Q}{2} \ . \tag{3.7}$$

$\mathrm{D_{JS}}$ is a symmetric ($D_{JS}(P,Q) \equiv D_{JS}(Q,P)$) version of the Kullback-Leibler divergence ($\mathrm{D_{KL}}$), also measured in bits, defined by

$$D_{KL}(P,Q) = \sum_x P(x) \log_2 \left(\frac{P(x)}{Q(x)} \right) \ . \tag{3.8}$$

$\mathrm{D_{KL}}$ is the extra number of bits required to encode distribution P based on Q versus if it were based on P. The closer $\mathrm{D_{JS}}$ is to 1, the better the separation of the two distributions. The closer it is to 0, the more similar the two distributions. Though there was no hard threshold, cluster pairs were generally judged to have good separation when $\mathrm{D_{JS}} \geq 0.95$.

The 1D overlap area ratio is a simpler measure which calculates the fraction of area over which the two distributions overlapped. At each bin in the common histogram of the two distributions, the lesser of the two values was taken. These lesser values were summed up for all bins, and then divided by the sum of the histogram values of the smaller distribution to get a (worst case) ratio. Again, there was no hard threshold, but values < 0.05 generally signified good separation.

The 1Dsep measure is calculated by normalizing the distance between the medians of the two 1D distributions by 3 times the standard deviation, σ, of the distribution with the greater σ. Generally, a value > 1 was required for clusters to be considered well separated. Like Fisher's linear discriminant, this measure was less meaningful for non-Gaussian distributions, for which σ poorly captures the true spread of the distribution.

NDsep was the preferred measure of cluster separation because it does not suffer from the loss of information inherent in projecting 3D clusters onto a single dimension. Like $\mathrm{D_{JS}}$ and the 1D overlap area ratio, it also does not assume Gaussian distributions. It was therefore relied on much more heavily than the other 3 measures. For a pair of clusters i and j, where the number of points in each is $N_i < N_j$,

$$\text{NDsep}(i,j) = 1 - \frac{1 - N_{nni}/N_i}{1 - N_i/(N_i + N_j)}\ , \quad N_i < N_j \tag{3.9}$$

where N_{nni} is the number of points in cluster i whose nearest neighbour is also in cluster i. This is exactly 1 minus the overlap index described in Swindale and Spacek (2014). The value ranges from 1 for completely separate clusters, to 0 for completely mixed clusters. The numerator is 1 minus the fraction of points in cluster i whose nearest neighbour is in cluster i. The denominator is 1 minus the probability that the nearest neighbour of any point in cluster i was also from cluster i, if the two clusters i and j were completely mixed. Like GAC, NDsep has $O(N^2)$ computational cost because distances between pairs of points must be calculated. Also like GAC, it was implemented in fast Cython code with multithreading (Section C.2). However, unlike the merge step in GAC, the input could be subsampled to limit the computational cost. For each cluster with N > 20,000 points, 20,000 points were randomly sampled from the cluster before calculating NDsep. As a result, the calculation took only a fraction of a second per cluster pair, no matter their size.

Because channel and time range selections (Figure 3.9) and dimension reduction methods (Figure 3.10) could greatly affect the separation of clusters, multiple combinations of channels, time ranges, and dimension reduction methods were tested for every pair of plausibly oversplit clusters. Clusters that remained poorly separated (as measured by the separation metrics and visualized by the user) no matter the cluster space were merged. Plotting against time in cluster space allowed the user to check if one cluster was "drifting into" another. If a pair of clusters showed an uninterrupted transition over time from one to the other, they were likely oversplit and were therefore merged (Figure 3.12B). Another way of checking for oversplitting in time was to use the sliding spike selection window (Section 3.10.1) on both clusters simultaneously. This allowed the user to see exactly how the spike waveforms were changing with time, and whether or not one cluster's spike waveforms gradually morphed into those of the other.

3.10.3 Misassignment

Sometimes when comparing pairs of clusters in cluster space, even if a pair remained distinct, some points would be obviously misassigned. Most often, this was only a small fraction of points, and that fraction might depend on which specific cluster space the points were plotted in. If points remained misassigned regardless of the cluster space, they were easily corrected by simply running GAC on the cluster pair with an appropriate σ_c, or by manually painting the points to be reassigned.

Discarded (unclustered) spikes were tested for potential incorporation into existing clusters. The waveforms of all discarded spikes were exhaustively compared to the mean waveforms of all clusters by calculating the RMS difference between each possible pair (the calculation was skipped for cluster-discarded spike pairs that had no channels in common). The user then selected one target cluster at a time, and the distribution of RMS differences was plotted for all discarded spikes that fit the target cluster better than any other. As in multichannel template matching (Section 3.1.1), an RMS difference threshold was set. Discarded spikes that fell below the difference threshold were automatically selected, and were then examined in the usual ways. If deemed similar enough to the target cluster, they could be merged into it, or GAC could be run on the selected discarded spikes and the target cluster's spikes, potentially merging some of the discarded spikes into the target cluster. Once complete, the user would select the next target cluster and repeat the process.

However, testing discarded spikes by template matching in this way was risky. Discarded spikes could not be optimally aligned to all clusters simultaneously, and therefore the exhaustive RMS difference calculation in the very first step suffered from some inaccuracy. Also, as with template matching, there was often no clear bimodality in the RMS difference distribution, and therefore no guide for where to place the threshold. Finally, merging formerly discarded spikes into existing clusters altered those clusters, and required re-testing them against all other nearby clusters to ensure they remained distinct, a laborious task (Section 3.10.2). This test for discarded spikes was therefore used only rarely and conservatively. Erroneously discarding some spikes was considered better than potentially contaminating existing clusters.

3.10.4 Duplicate spikes

For each cluster, duplicate spikes, i.e., spikes with identical spike times belonging to the same cluster, were searched for and removed. Duplicate spikes were very rare, occurring at a rate of $<$ 0.02% of spikes. For tracks ptc15.tr7c, ptc22.tr1, and ptc22.tr2 respectively, 1154/7.3M, 277/2.05M, and 48/1.4M duplicate spikes were found. This was an overestimate because most clustered units were not considered "good" for export, yet their duplicate spikes were included in the counts. These duplicate spikes may have been due to temporally broad spikes with spike-like peaks just outside their spatial lockout, which were then shifted towards the true spike's primary peak due to excessive realignment, i.e., excessive realignment may have undone the effects of the spatial lockout (Section 3.7).

3.11 Autocorrelograms & refractory periods

Autocorrelograms could be plotted for any cluster and cross-correlograms could be plotted for any pair of clusters. Cluster pairs with very similar autocorrelograms, and whose cross-correlograms were similar to their autocorrelograms, were considered potential candidates for merging (Hazan et al., 2006), but such evidence alone was not considered enough to warrant a merge. Autocorrelograms were also inspected for a minimum refractory period gap in spike intervals of $\Delta t = \pm 0.75$ ms (Figure 3.13A). Although the temporal lockout during spike detection prevented some refractory period violations (RPVs), its duration varied from one spike to the next, and was never greater than 0.4 ms (Section 3.4). Moreover, the temporal lockout was not enforced during later spike realignment (Section 3.7), allowing for further potential RPVs. RPVs were therefore quantified from the autocorrelogram using a fractional measure f, defined for cluster i as

$$f_i = \frac{R_i}{N_i} \tag{3.10}$$

where R_i is the total number of spike intervals within the designated refractory period (0.75 ms) and N_i is the number of consecutive spike intervals (number of spikes minus 1) for cluster i. A value of 0 indicates that no spike intervals fell within the refractory period, and a value of 1 indicates that all of them did (very rare non-consecutive spike interval RPVs may have caused a negligible overestimate of this fraction). As expected for well-sorted clusters (those marked as "good" for export), all had a very small value of f (Figure 3.13A, top two rows). Out of all 3 sorted tracks, the cluster with the highest value had $f = 0.004$.

To show how effective or ineffective autocorrelograms are at revealing cluster contamination, pairs of well-sorted clusters were merged, and the autocorrelograms of the resulting merged spike trains were plotted (Figure 3.13A, bottom row). The RPVs of the merged cluster m were compared to those of the original clusters i and j, and the change in fractional RPVs as a result of the merge was calculated as

$$df_{ij} = \frac{R_m - (R_i + R_j)}{N_i + N_j}\ . \tag{3.11}$$

Because $R_m \geq R_i + R_j$, df also ranged from 0 to 1, with 0 indicating no change in RPVs and 1 indicating the largest possible change (from no spike intervals to all spike intervals falling within the refractory period). This was measured for all possible pairs of neurons in each track. For the vast majority of pairs, df was very small at less than 0.001 (Figure 3.13C). Results were similar for longer designated refractory periods of 1 or 2 ms (not shown). This demonstrates that most of the time it was virtually impossible to detect even blatant cluster contamination by examining autocorrelograms for RPVs.

3.12 Simulation

Spike sorting methodologies may be tested against data with a ground truth of known spike times. Such data either come from simultaneous intracellular and extracellular recordings (Harris et al., 2000), or from simulated extracellular data. Relatively short simulated polytrode recordings were generated by taking templates from real recordings, and adding them at random times to artificial noise generated independently per channel. The autocorrelation of the generated noise was matched to that of spikeless segments of real recordings (Swindale and Spacek, 2015). However, such simulated recordings were trivial to sort compared to real data, and could be sorted almost flawlessly even in a single big (x, y, V_{pp}) cluster space. Few of the methods described here were necessary for sorting such simulated data. A more realistic simulation with the same spike sorting challenges of a multi-hour track of real data would be much more complicated to generate. It would require noise events correlated over multiple channels, variable noise levels, and above all templates

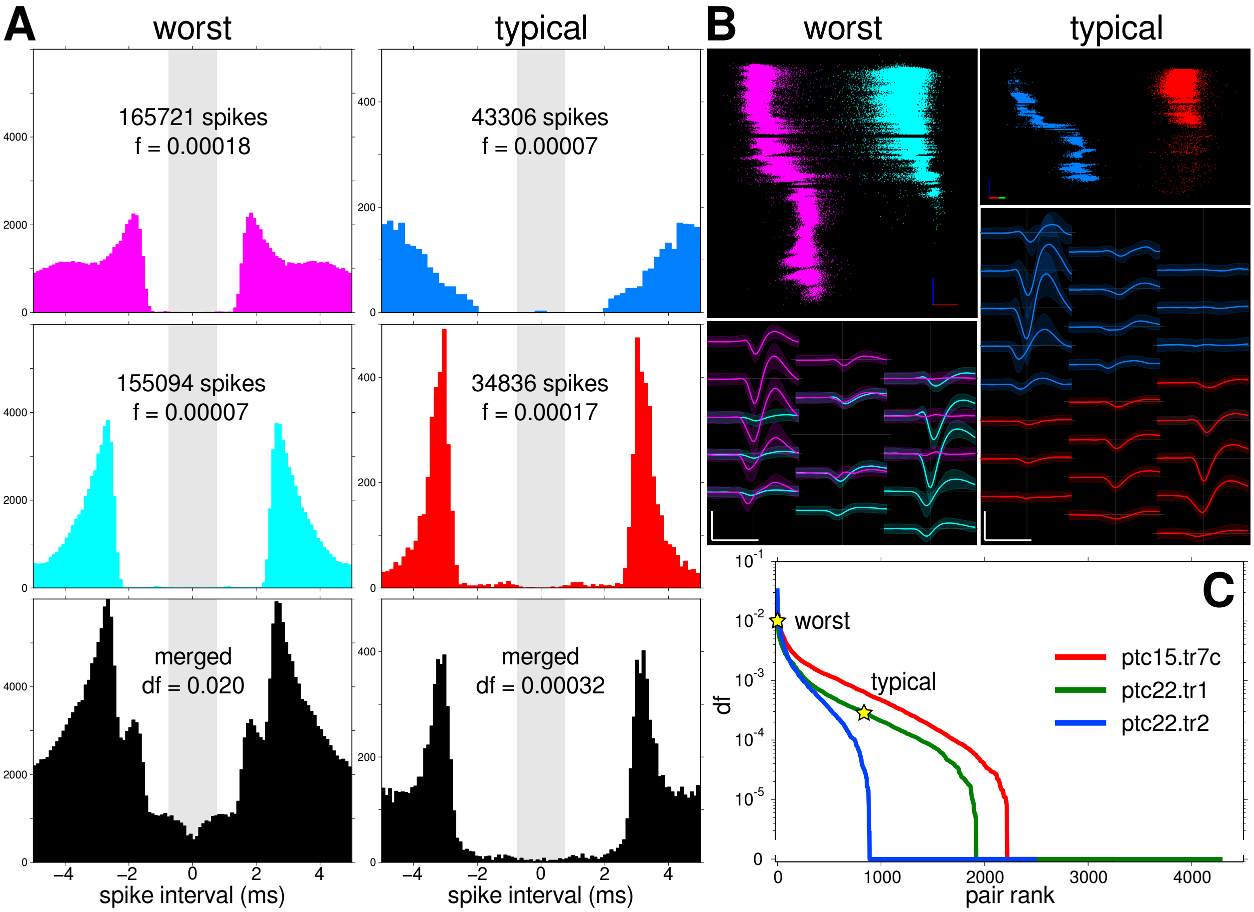

Figure 3.13: Examining autocorrelograms for RPVs is not useful for diagnosing cluster contamination. ***A***: When two distinct clusters (first two rows) were merged, there was typically little increase in the number of spike intervals that fell within the refractory period of the merged cluster, as shown by its autocorrelogram. The worst case scenario is shown in the left column, and the more typical scenario in the right column. Shaded grey regions represent the designated refractory period ($\leq$ 0.75 ms). ***B***: Colour coded 3D cluster plots (first two PCs vs. time along the vertical axis) and mean waveforms of the neurons in ***A***, showing that for each of the two example neuron pairs, the cells were simultaneously active and in distinct spatial locations. White scale bars: 0.5 ms, 100 μV. ***C***: Simultaneously recorded pairs in each of the 3 tracks were sorted by decreasing df (Equation 3.11). Note the vertical log scale. The vast majority of merged pairs had very small df. About half had $df = 0$. The positions of the worst case and typical examples from ***A*** and ***B*** are shown. Results were similar for longer designated refractory periods of 1 or 2 ms.

drifting independently of each other and at non-constant rates. Finally, metrics would be required to verify that any simulated data were comparable in complexity to real data. Simulated data were therefore not extensively used to test the spike sorting method presented here, but would be of great utility in evaluating and improving the performance of all spike sorting techniques (Einevoll et al., 2012).

3.13 Discussion

The spike sorting method described here used a divide-and-conquer approach by splitting polytrode data into multiple smaller overlapping spatial domains (Swindale and Spacek, 2014). As a result, computational complexity and overall spike sorting time (including user time) varied roughly linearly with the number of electrode sites and the number of isolatable units. This is a desirable property, especially if future recording technologies with thousands of sites and neurons are to be made practical (Einevoll et al., 2012; Kording, 2011; Stevenson and Kording, 2011).

Other than during spike detection, there was little automation in the spike sorting procedure described here. This was intentional. There were too many variables, including alignment, channel and time range selection, dimension reduction, and cluster comparison and verification, for an automated procedure to fully explore and safely make decisions about. Instead, the focus was placed on making the software fast to use, in both the time taken to perform each step, as well as in its graphical user interface (GUI) (Section C.2) and keyboard shortcuts. Unfortunately, when spike sorting long duration cortical recordings with many adjacent cells and significant drift, user intervention and experience remain key. Presenting the user with all relevant options with which to quickly examine and compare clusters in a variety of different spaces allowed for greater confidence in the final result.

3.13.1 Spike detection

Putative monophasic spikes (i.e., spikes consisting of only a single peak on the primary channel, with no adjacent peaks within $\pm\, dt_{max} = 0.4$ ms) were ignored during spike detection (Section 3.4), and therefore excluded from all subsequent analysis in this study. The spike detection algorithm was later modified to allow for monophasic spikes, with the requirement that the single peak exceed the peak-to-peak threshold V_{ppt} and that it be bounded by zero-crossings no more than $dt_z = dt_{max}$ from each other (effectively setting a minimum sharpness threshold). Examination of the specific differences in spike detection resulting from these changes revealed that most trigger single peaks were very broad, slow and negative, and were generally part of the AHP of a larger preceding spike. In a recording in ptc15.tr7c (Figure 2.3), only 3/283 ($\sim$ 1%) of trigger single peaks passed all of the above criteria and were detected as monophasic spikes (out of a total of 4491 detected spikes, i.e., 0.07%). By visual inspection, all 3 of these were indeed plausible spikes, but were monophasic only due to distortion from partial spike overlap with a preceding spike that was larger in amplitude (not shown). Moreover, since intracellular spikes must always have at least two phases (depolarization and repolarization), and since extracellular spikes are approximately the negative of the first derivative of intracellular spikes (McCormick et al., 1985; Henze et al., 2000), non-overlapping extracellular somatic spikes should by necessity have at least two peaks of alternating sign. Therefore, excluding monophasic spikes during extracellular spike sorting does not introduce a bias against them, simply because they do not (and should not) exist.

The temporal lockout was different for each detected spike, but the spatial lockout was the same for all spikes. The 150 μm lockout radius was set for the spatially broadest of spikes. Therefore, spikes that were more spatially focal resulted in an unnecessarily broad spatial lockout, which may have resulted in more missed spikes than necessary. A spatially adaptive lockout for each spike based on the modelled spatial position and extent (Section 3.5) would fix this oversight, while still minimizing multiple registrations of the same spike. GAC-based spike detection (Swindale and Spacek, 2015) (see below) avoids this problem altogether by not using a lockout.

The spike detection method described in Section 3.4 is a way of combining large numbers of candidate events into a smaller number of true spikes. It was later found that detecting spikes on many closely spaced channels is fundamentally a clustering problem (Swindale and Spacek, 2015), which was solved here by rectangular spatiotemporal search and lockout windows. Though somewhat crude as a clustering method, subjectively it was good at detecting events that were clearly spikes, rejecting most noise events, and preventing multiple registrations of the same spike. A more sophisticated method now exists that fully treats spike detection as a threshold event clustering problem (Swindale and Spacek, 2015). It uses a modified version of the GAC algorithm (Section 3.9), in which each threshold crossing event is a point in space and time to be clustered, with scout points moving up a spatiotemporal voltage-weighted gradient to find the highest local peak, corresponding to the likely origin of a spike in space and time. This method eliminates the need for search and lockout windows. Besides being more elegant and principled, GAC-based spike detection can resolve spikes slightly closer in space and time than the spatiotemporal lockout method described in Section 3.4, with little or no extra computational cost (Swindale and Spacek, 2015), and should therefore be implemented in spyke in the future.

Spike detection was based on independent single channel thresholds. Visual inspection of the voltage data revealed instances when several neighbouring channels had similar, simultaneous, but subthreshold spike-like events (Figure 3.4, white arrows). These were reasonably distinct from the independent noise on each channel, suggesting a spike. However, because each channel was tested independently for threshold crossings, and because the events were subthreshold on all channels, the spike was not detected. One solution to this problem is to simply lower the threshold on all channels (e.g., by setting $A = 5$ instead of 6, Equation 3.1), but doing so would also greatly increase the number of false positive spike detections, making later clustering more difficult. A better strategy would be to weigh the evidence for a spike by looking across multiple neighbouring channels at similar timepoints to see if several of them show signs of a spike, even if they are all low amplitude. One might consider a low amplitude multichannel event to be as likely a spike as a large amplitude but purely single channel event. A modification of the GAC-based spike detection method (see above) might be a solution. The event detection threshold could be set lower (perhaps to $A = 4$) to detect even the smallest plausible multichannel events using a single channel threshold. Then, GAC-based spike detection would be run on all of these low amplitude events to cluster redundant events close in space and time into putative spikes. Each putative spike could then be tested using a different kind of threshold, based on the height of the "hill" climbed by GAC, i.e., the density of events around the location and time of the putative spike, each event weighted by its voltage amplitude. This way, a large number of small-amplitude events from several nearby channels and timepoints would be about as likely to cross threshold and be accepted as a spike as a large amplitude single channel event.

3.13.2 Dimension reduction

ICA was most useful for separating clusters of vastly different numbers of spikes (size ratios $>$ 50, Figure 3.10A & B). This was initially surprising, but can be explained as follows: PCA seeks to explain variance. The variance of the combination of a very small cluster and a large one will be dominated by the variance of the large cluster (Pedreira et al., 2012). ICA on the other hand considers higher order correlations than just variance, and seeks dimensions of maximum independence instead of just orthogonality. If two clusters are truly statistically independent from each other, ICA can reveal them regardless of their size ratio. These small clusters could be true spikes (often found to belong to some other existing larger cluster) or noise events, which were discarded.

Conversely, ICA was worse than PCA at separating similarly sized clusters (size ratios $<$ 10, Figure 3.10C), sometimes dramatically so. This was also surprising, and its cause remains unclear. One possible reason is that ICs were sorted by signed kurtosis (decreasing from positive to negative) instead of the absolute value of kurtosis (A. Hyvärinen, personal communication). Another possible reason may be that kurtosis is poorly suited to sorting the ICs from similarly sized clusters, and that negative entropy (negentropy) (Hyvärinen, 1998) is a better choice (A. Hyvärinen, personal communication). Both of these possibilities were briefly explored, but only after spike sorting had been completed. Sorting ICs by decreasing absolute value of kurtosis did improve the ability of ICA to separate similarly sized clusters, but only occasionally. Sorting ICs by negentropy separated similarly sized clusters even better, but there were still cases where it performed worse than PCA. Neither modification appeared to degrade the ability of ICA to sort very differently sized clusters. Further work in this regard is required. It may be that sometimes the sources of extracellular signals are not independent, and that therefore ICA is inappropriate to use in such cases (Jäckel et al., 2012).

A method was devised to automatically select channels for dimension reduction and clustering, based on the median spatial extent of the selected spikes (Section 3.8). However, this automated method was often less effective at separating clusters than was manual channel selection (Figure 3.9), hindering downstream clustering. Moreover, time range selection was completely manual, and was identical for all selected channels. An automated method of selecting both channels and time ranges that is at least as effective at separating clusters as manual channel and time range selection could accelerate the entire spike sorting procedure, and would also make it more objective. One such method might choose channels and time ranges (or individual timepoints) based on the mean signal amplitude or variance of the selected spikes. More fundamentally, quantitative comparison of any method of dimension reduction requires a robust way of measuring the clusterability of the resulting distribution of points, i.e., multimodality: the number of modes and their separation. There are many potential ways of measuring multimodality (Hartigan and Hartigan, 1985; Nason and Sibson, 1992), but a method that considers a specific spatial scale of interest might be simplest and most effective in this context. Finding and/or developing such a measure of multimodality requires further work.

3.13.3 Clustering

As a variant of the mean-shift algorithm (Fukunaga and Hostetler, 1975), GAC has several benefits. First, it does not assume clusters are of any particular distribution, such as Gaussian. This is especially beneficial for dealing with clusters with distorted shapes, due to drift or otherwise. Second, GAC clusters points in a way that matches human visual intuition. Third, GAC makes no assumptions about the number of clusters in the data. Fourth, GAC requires the manipulation of only a single parameter, σ_c, describing the spatial scale of features of interest. The choice of σ_c was greatly aided by onscreen user measurement of such features.

A moderate disadvantage of GAC is that it is of $O(N^2)$ computational complexity, where N is the number of points. The $O(N^2)$ complexity comes from the need to measure the distance between all pairs of points at least once. Although the gradient calculation step can subsample the data that fall within r_{neigh} to reduce dependence on N, the merge step cannot. With extensive optimizations in the implementation of GAC (Section C.2), clustering $N = 25{,}000$ data points with $\sigma_c = 0.4$ took ~ 0.5 s on a first generation quad-core i7 laptop. Smaller values of σ_c slowed the algorithm down (Figure C.4), mostly because the merge radius was proportionally smaller.

In practice, Manhattan distance (i.e., city block or rectilinear distance) was used instead of Euclidean distance (Swindale and Spacek, 2014; Equations 3.4 & 3.5) to calculate density gradients for GAC. This was unintentional and went unnoticed until after all spike sorting had been completed. Euclidean distance is more principled and has no computational penalty, and should therefore have been used instead. The clustering performance of the two distance metrics was compared for both real clusters and simulated Gaussian clusters by examining animations of scout point movement, as well as the final cluster counts and positions (not shown). Fortunately, the clustering performance of the two metrics was very similar, and any differences could be attributed to slightly different sensitivities to σ_c, which was manually set to effect by the user anyway. The similarity in clustering performance may be due to the small step sizes used during the gradient ascent step. The Manhattan distance somewhat favoured clustering points that were distributed parallel to any one of the axes of the cluster space over those that were not, such as diagonally distributed points. The Euclidean distance had no such directional bias. Since clusters were most often aligned with one or more of the clustering axes (such as PC and IC axes), it is possible that this directional bias may have even been beneficial for clustering, but this was not explored. The spike sorting results presented in the following chapters used Manhattan instead of Euclidean distance to calculate density gradients, but there is no reason to believe that this made any practical difference.

GAC has some similarity to HC in that during the course of the algorithm, each data point has a history of membership of increasingly populated clusters that could be represented in a hierarchy. However, unlike HC, GAC exits automatically at the appropriate level in the hierarchy. GAC is also similar to gravitational clustering methods (Wright, 1977) which allow points to attract one another and form clusters. The difference is that the density landscape for gravitational algorithms is constantly changing, while for GAC it remains constant.

Several other density-based clustering algorithms exist (Kowalewski, 1995; Ester et al., 1996; Klusch et al., 2003; Wang et al., 2004). All differ from GAC in subtle ways, and like the mean-shift algorithm, were found in the literature only after development of GAC. None were investigated to any great extent here. Ester et al. (1996) suggest an interesting way to quickly determine the optimal value of σ_c. For each point, the distance of its k^{th} nearest neighbour (say, $k = 4$) is found. This is called the "k-dist" of each point. All points are then sorted by their k-dist, in decreasing order, and plotted. Plateaus in such a plot, where many points have the same k-dist value, signify significant structure in the data at that spatial scale. Such plateaus are apparently robust to the choice of k, and points to the left of the first significant plateau are likely outlying noise points. Such a pre-clustering step could run much faster than even a single run of GAC, especially for large N, and may be implemented in the future.

A new density-based clustering algorithm by Rodriguez and Laio (2014) may represent a significant improvement over the mean-shift algorithm and GAC. It preserves the ability to handle non-Gaussian clusters and to automatically discover their number. For the purposes of spike sorting, the algorithm's major advantages are that *a*) it is a fast, simple process, requiring no computationally expensive iteration, and *b*) it has essentially no user-set parameters. In the first step, the local density ρ around each point is calculated, potentially using a computationally cheap step function kernel that simply counts the number of points within a critical distance d_c, to which the algorithm is highly insensitive. In the second step, the distance δ from each point to the nearest point with higher ρ is found for each point. Finally, the product $\gamma = \rho\delta$ is calculated for each point, and the points are ranked in decreasing order. As determined by being above the knee of the γ rank order plot, points with high values of both ρ and δ represent cluster centers. In practice, when splitting off only one cluster at a time (Section 3.10.1), only the point with the highest γ need be found.

Other points are assigned to a cluster by following the linked list of nearest neighbours of higher local density, generated while calculating δ. Although still of $O(N^2)$ computational complexity due to its need to calculate distances between all pairs of points once, the single-shot nature of the algorithm of Rodriguez and Laio (2014) likely makes the gain of the $O(N^2)$ dependence very low. Incorporation of this clustering method into spyke and other existing spike sorting software may therefore help speed up the clustering step.

3.13.4 Autocorrelograms & refractory periods

As an experiment, distinct, simultaneously active clusters were intentionally merged to create blatant spike contamination (Section 3.11). Tests on these merged clusters showed that RPVs, typically visualized with autocorrelograms, were a poor indicator of cluster contamination (Figure 3.13). This is perhaps unsurprising in retrospect. Given the low firing rates (Section 4.3) and weak spike correlations (Section 6.3 & Appendix B) reported here in anesthetized cat V1, the probability that a given pair of cells will both fire a spike within one millisecond of each other is very low. When merging completely independent spike trains each with a mean firing rate of 0.1 Hz (Figure 4.1), one would expect a 1 ms RPV only once every 100,000 s. For mean firing rates of 1 Hz, RPVs would still arrive at a rate of only one every 1000 s. This conclusion, that autocorrelogram RPVs are a potentially unreliable indicator of spike sorting quality, has been briefly mentioned in the literature (Gray et al., 1995; Harris et al., 2000). Yet this is an underappreciated finding, as autocorrelograms and RPVs are widely perceived to be a key tool for testing spike sorting quality (Fee et al., 1996; Alonso et al., 1996; Nirenberg et al., 2001; Litke et al., 2004; Segev et al., 2004; Hill et al., 2011; Prentice et al., 2011; Marre et al., 2012; Pillow et al., 2013). At best, they are necessary but insufficient for the task. At worst, they may provide a false sense of security.

3.13.5 Drift

The greatest difficulty in spike sorting multi-hour polytrode recordings was drift (Figure 3.12). Strangely, not all units drifted by the same amount, although generally drift was in the upward direction. Certainly, polytrode drift must play some role in this, but on its own cannot account for the variability of drift across units. Whether this extra variability is due to a physical property of cortical layers, of glial activity, or of some other source, is unknown. Some tracks had more drift than others (see Figures 4.5 & 4.6 in the next chapter). Efforts to minimize drift should be a priority during future long duration recordings, in order to ease subsequent spike sorting.

4 Basic Physiology

With spike sorting completed, the next step was to examine the basic neurophysiology of the sorted single units, again with the aim of minimizing bias as much as possible. Neurophysiological traits examined in this chapter include neuron yields, firing rates, cell templates and spatial positions, and orientation tuning.

4.1 Summary

Spike sorting of 3 tracks in 3 hemispheres in 2 cats resulted in 245 single units (Section 4.2). Surprisingly, mean firing rates followed a lognormal distribution instead of a normal one, and ranged 5 orders of magnitude, from ~ 0.0001 Hz to ~ 10 Hz (Section 4.3). The geometric mean was 0.11 Hz, and 82% of cells had mean rates below 1 Hz. Many studies exclude such low rate cells from analysis, potentially biasing their results towards the properties of high rate cells. All cells were included for most of the analyses in this study, regardless of firing rates. Surprisingly, initial results suggest that local neural populations perform a kind of shift work: as some cells stop firing, others start firing, keeping the overall geometric mean firing rate relatively constant despite orders of magnitude fluctuations in individual firing rates over time.

Multiple neighbouring channels were crucial for distinguishing neighbouring cells with very similar spike shapes on their primary channel (Section 4.4). Neurons were well distributed across the length of their polytrode, but their spatial positions were biased toward the electrode site locations, suggesting that 65 μm site spacing is not fine enough to fully sample the local population. Spatial positions varied as a function of time, suggesting polytrode-tissue drift, yet this drift was not consistent across the length of the polytrode, across time, or even across neighbouring neurons, and its source is unknown.

Orientation tuning curves were calculated for all neurons. Regardless of firing rate, 61% of cells were significantly tuned, but when restricted to active neurons (trial-averaged rates ≥ 0.05 Hz), 87% of cell were significantly tuned (Section 4.5). As expected for generally vertical polytrode insertions, certain orientations were preferred within each track, but some tracks had a wider range of orientation preferences than others, indicative of a more transcolumnar insertion. Surprisingly, orientation tuning strength was inversely correlated with log firing rate. Cells with trial-averaged rates < 0.01 Hz had the strongest orientation tuning.

4.2 Neuron yields

Initially, rough spike sorting was performed on small subsets of data from all tracks. Tuning curves (Section 4.5) and RFs (Section 5.4) of the most easily sortable units in each track were examined,

track	units	sorted spikes	detected events	%	duration	spike rate
ptc15.tr7c	81	2,661,710	7,287,225	36.5	11.8 h	62.6 Hz
ptc22.tr1	93	1,664,985	2,054,825	81.0	8.3 h	55.5 Hz
ptc22.tr2	71	797,641	1,399,230	57.0	6.2 h	35.6 Hz
total	245	5,124,336	10,741,280	47.7	26.4 h	54.0 Hz

Table 4.1: Neuron yield. For each track and in total, columns show single unit yield, sorted spike and detected event count, percentage of events that were classified as single unit spikes, recording duration, and mean population sorted spike rate. Recording duration excludes recording time gaps within each track (Figure 4.2).

which allowed for a qualitative estimate of the physiological normality of each track. Subjectively, the best 3 tracks (ptc15.tr7c, ptc22.tr1, and ptc22.tr2, highlighted in Table 2.2) were picked for full sorting and subsequent analysis. These 3 tracks were 11.8, 8.3, and 6.2 hours in recording duration (track duration minus any recording time gaps), respectively, for a total of 26.4 hours of recording from two different cats. The two tracks that came from the same cat (ptc22) were from opposite hemispheres. Once the spike sorting method was fully developed, sorting a track took about a month of concerted user effort.

Recording conditions for the first few minutes after polytrode insertion were very different from more stable conditions starting $\sim$ 30 min later (not shown). Fast, purely single channel spikes were prevalent, while larger, slower multichannel spikes were mostly absent. The fast single channel spikes persisted for only a few minutes after insertion, followed by a generally quiet period, and then a slow onset of more typical multichannel spikes. Some of these initial periods were saved to disk, but most recordings saved to disk did not begin until at least 30 min after insertion.

Single unit yields, spike counts, and mean population spike rates are shown in Table 4.1. Single unit yields per track were 81, 93, and 71 respectively, for a total of 245, and a mean of 82 per track. Across all tracks, over 10 million spike-like events were detected (Section 3.4), and about half of these were classified as single unit spikes (i.e., belonging to clusters marked as “good” for export, Section 3.10). The rest were classified as either noise events or multiunit spikes. Single unit spikes arrived at a grand average rate of 54 Hz.

Single unit yields were greater across a track than for any single recording from that track. For example, the first drifting bar recording in ptc22.tr1 lasted 22 minutes. During that time, only 66 cells fired at least one spike, out of a total of 93 cells sorted in that track. Relative to that one recording, track-wide sorting resulted in a 48% increase in yield. If restricted to cells with a mean firing rate of at least 0.05 Hz, the numbers were 30 and 58 cells respectively, a near doubling of yield. Therefore, in addition to likely improving detection of very low firing rate cells (Section 3.2), track-wide sorting increased overall neuron yield.

4.3 Firing rates

To examine the mean firing rates of the neural population, the mean firing rate of each neuron was calculated over the duration of its track. Neurons were plotted in decreasing order of mean firing rate on a log-linear scale (Figure 4.1A). The plots of all 3 tracks (red, green and blue lines) resembled a rotated sigmoid function, as did their mean (black). The same data were also presented as a distribution of mean firing rates on a linear-log plot (Figure 4.1B) of all 245 neurons from all 3 tracks. The log-average (the mean of the base-10 log, i.e., the geometric mean) of the firing rate across all sorted units was -0.97 log units, or 0.11 Hz (black arrow). The log-average of the mean firing rates was calculated instead of the arithmetic mean because on a linear scale the mean firing rates had a very non-normal distribution. Instead, the distribution of mean rates resembled a lognormal distribution, which was fit to the data using the least-squares LM algorithm (Figure 4.1B, magenta curve). The fitted lognormal curve had a mean of -0.82 log units, or 0.15 Hz (magenta arrow), and a standard deviation of 1.05 log units, or roughly an order of magnitude on either side of 0.15 Hz. Mean firing rates of all cells from all tracks ranged 5 orders of magnitude, from $\sim$ 0.0001 Hz to $\sim$ 10 Hz.

The mean trace in black and the lognormal fitted trace in magenta in Figure 4.1A were approximately linear over 2 orders of magnitude of mean firing rates, from $\sim$ 0.01 Hz to $\sim$ 1 Hz. A linear range in the rank-order plot would be equivalent to a decreasing exponential firing rate distribution, which would result in a decreasing sigmoid over the same range of firing rates in Figure 4.1B. Such a sigmoid was not apparent in Figure 4.1B, and the firing rate distribution was therefore unlikely to be exponential instead of lognormal.

Mean firing rates were $<$ 1 Hz for 82% of neurons (74%, 83%, and 89% per respective track), and $<$ 0.05 Hz (dashed grey lines in Figure 4.1) for 35% of neurons (24%, 38%, and 44% per respective track). A 0.05 Hz mean firing rate threshold was used to classify cells as active or inactive. This choice of threshold was somewhat arbitrary, but very low, corresponding to one spike every 20 s on average. Mean firing rates differed for different analyses, depending on the time span encompassed by the analysis. Therefore, even though the mean firing rate threshold was kept constant, the set of active and inactive neurons for an analysis depended on whether that analysis was specific to one or more recordings, or to an entire track.

In addition to track-wide means, firing rates were was also examined over time. To estimate even very low rate time series, coarse 5 min wide overlapping time bins at 1 min resolution were used for all neurons. To prevent time gaps in between recordings from causing underestimation of firing rates, gaps were ignored and time bins were aligned to the start of each recording. Rate time series were plotted on a vertical log scale for the first 5 hours of each track (Figure 4.2A). Some cells had rates that ranged over 2 orders of magnitude. For example, the cyan trace for one cell during a blank screen recording in ptc22.tr2 (arrow, bottom panel, Figure 4.2A) began at 0.02 Hz at t = 1.6 h, increased to 6 Hz 18 min later at t = 1.9 h, and then decreased again to 0.02

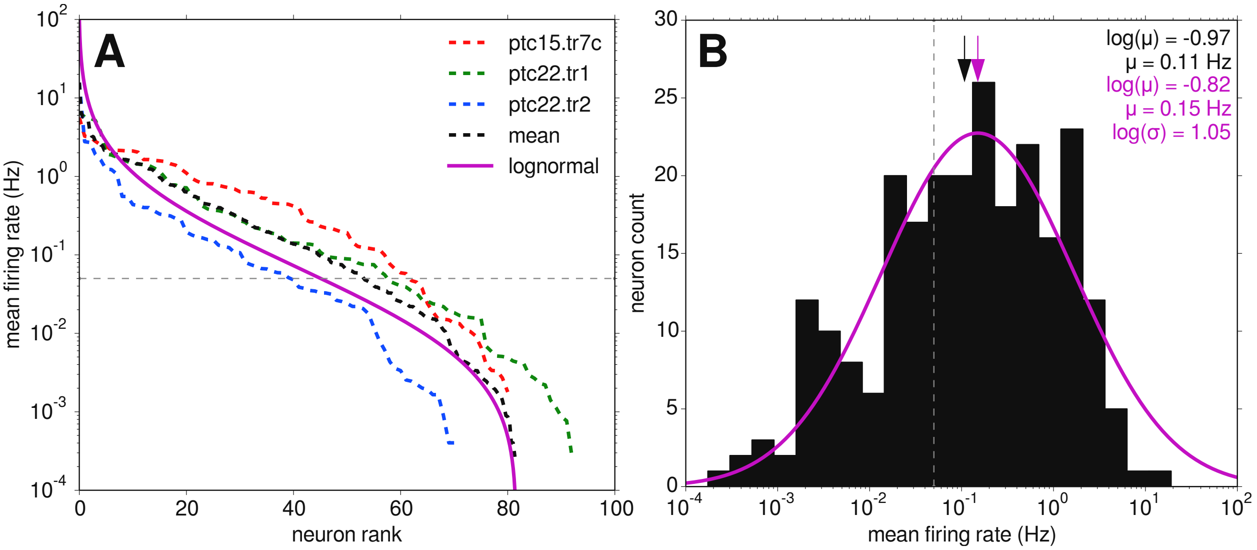

Figure 4.1: Mean firing rates had a lognormal distribution. ***A***: Mean firing rates of all 245 cells of all 3 tracks, ranked in descending order within each track. Note the vertical log scale. The black dashed line ("mean") ranks all 245 cells but is compressed horizontally by a factor of 3 for display purposes. On this scale, a straight line with negative slope corresponds to a decaying exponential firing rate distribution, while a rotated sigmoid corresponds to a lognormal distribution. The solid magenta line corresponds to the lognormal fit in ***B***. ***B***: Mean firing rate distribution of all cells corresponding to the black dashed line in ***A***. Note the horizontal log scale. The magenta line is the LM least-squares best fit lognormal distribution, with fit parameters shown in the top right. Black and magenta arrows denote geometric means of the distribution and fit, respectively. Both ***A*** and ***B*** show that the majority of cells had mean firing rates below 1 Hz. The grey dashed horizontal and vertical lines at 0.05 Hz in both plots denote the minimum firing rate threshold for cell inclusion for some subsequent analyses. 35% of cells fell below this threshold.

Hz another 15 min later at $t = 2.15$ h. Visual inspection showed that deep and middle layer cells generally had higher firing rates than superficial cells, and that some cells covaried in their rates, at least at this coarse time resolution. The firing rates of some cells covaried positively (especially evident for high rate cells in ptc15.tr7c), and others covaried negatively.

How might the firing rate of the entire population fluctuate over time? Since the mean rates of all neurons over all time were lognormally distributed (Figure 4.1B), instead of the arithmetic mean rate, the log-average (geometric mean) rate was calculated for all neurons as a function of time (thick grey transparent line in Figure 4.2A). Compared to the wide range of rates for each neuron across time and across neurons at any given point in time, the log-average remained relatively constant over the duration of each track, even during some periods of positive covariation in high rate neurons. The distributions of log-average rates are shown for the entire duration of each track in Figure 4.2B. Each distribution spanned no more than about 1 order of magnitude, and all were well below 1 Hz. Log-average distribution means for the three tracks were ~ 0.3 Hz, 0.09 Hz, and 0.15 Hz respectively, in line with the 0.11 Hz geometric mean rate shown in Figure 4.1B.

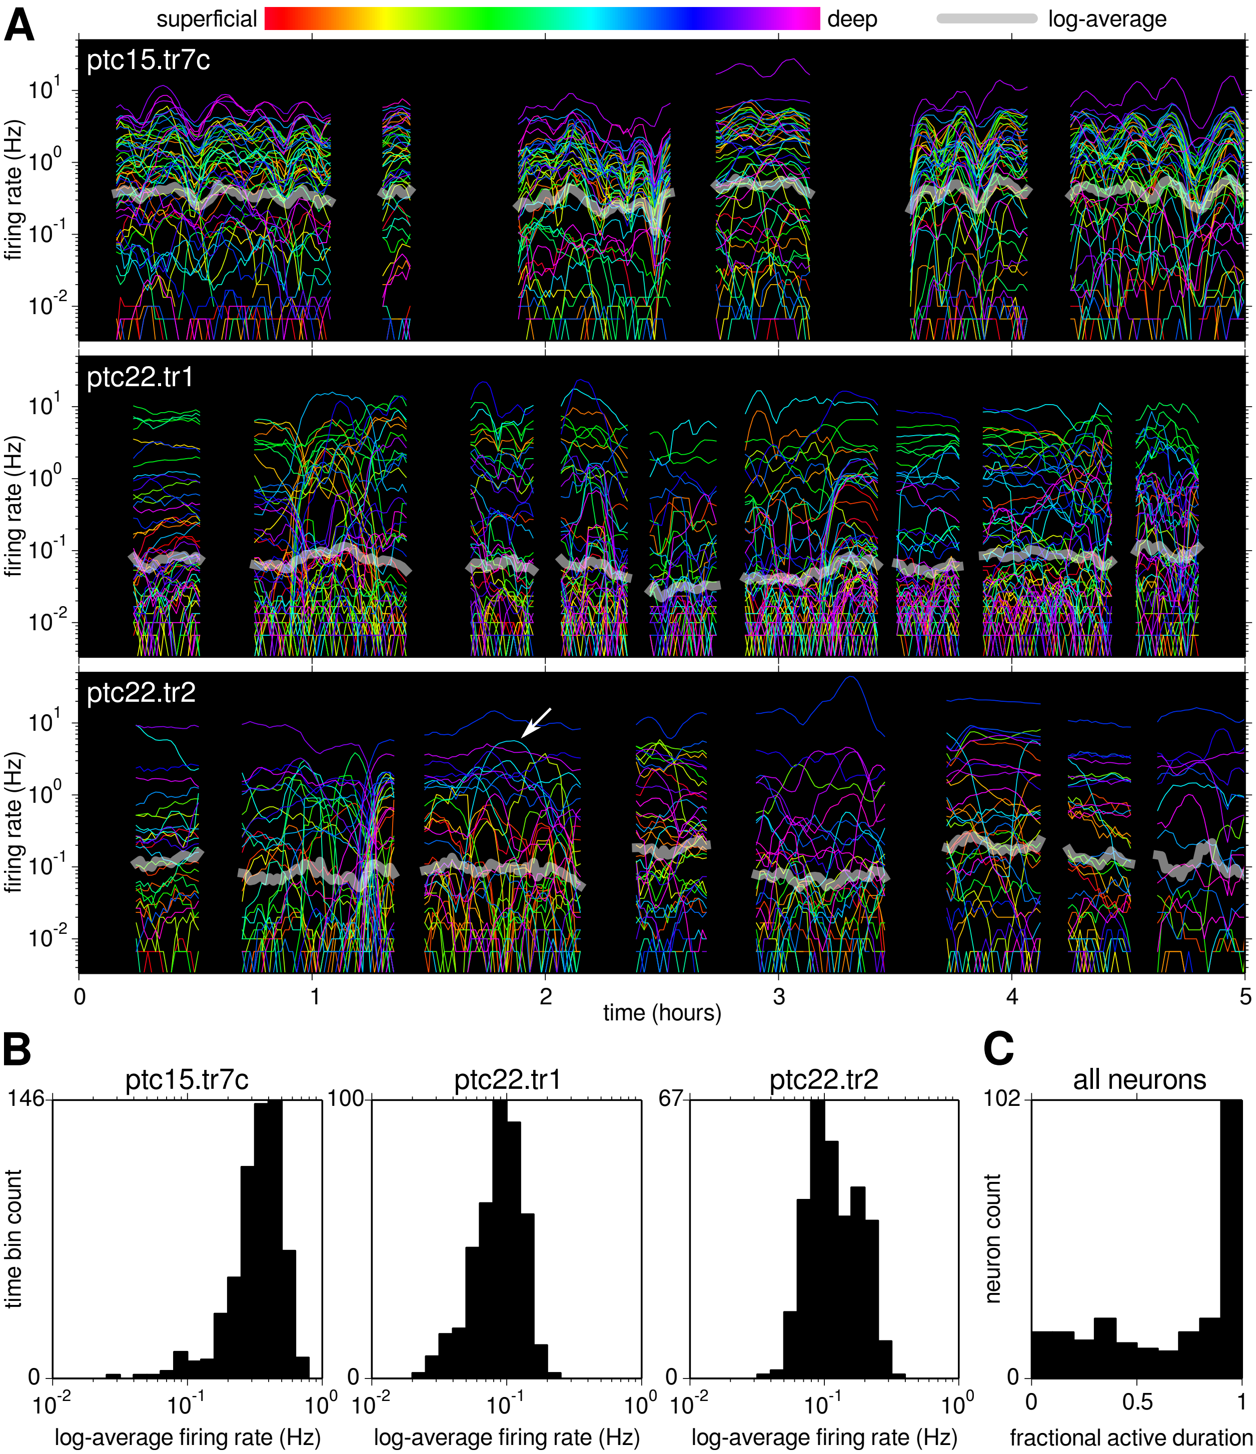
A
superficial
deep
log-average
ptc15.tr7c
ptc22.tr1
ptc22.tr2
firing rate (Hz)
10^1
10^0
10^{-1}
10^{-2}
0
1
2
3
4
5
time (hours)
B
ptc15.tr7c
ptc22.tr1
ptc22.tr2
146
100
67
0
time bin count
10^{-2}
10^{-1}
10^0
log-average firing rate (Hz)
C
all neurons
102
0
neuron count
0
0.5
1
fractional active duration

Figure 4.2: (*Previous page.*) ***A***: Coarse firing rates plotted as a function of time for each neuron in each track. Each neuron's trace is coloured according to its depth rank along the length of the polytrode. Overlapping time bins were 5 min wide at 1 min resolution. Because of the vertical log scale, 0 Hz bins were excluded. For greater visibility, only the first 5 hours of each track are shown. Each contiguous block of traces represents one recording, separated by gaps from neighbouring recordings in that track. The stimulus varied from one recording to the next. Middle and deep layer neurons generally had the highest overall rates. Some neurons had rates that ranged > 2 orders of magnitude (arrow denotes an example). Some rates covaried positively (especially in ptc15.tr7c), some negatively, and others not at all. The thick grey transparent line is the log-average. Although individual rates varied widely, the log-average remained relatively constant. ***B***: Distributions of log-average rates for each track, collected over the entire duration of each track (not just the first 5 hours). Log-averages ranged no more than about an order of magnitude around their mean, but this was an overestimate (see text). ***C***: The distribution of activity duration (time between first and last spike of each neuron) as a fraction of track duration, for all neurons. Other than the $\sim 40\%$ of neurons that were active for their entire track (peak at 1), the distribution was reasonably uniform.

Note that the variance in log-averages (thick grey traces in Figure 4.2A) was overestimated because of the necessary exclusion of many cells during those particular time bins in which they fired zero spikes. As the number of such excluded cells at any point in time was a substantial fraction of the total number of cells, and as it varied substantially over time, this inflated the variance in the calculated log-average distributions. The true log-averages were therefore likely to be even more constant over time than shown here. Wider time bins (20 min instead of 5 min, both at 1 min resolution) somewhat reduced this tendency by reducing the number of 0 Hz bins, thereby decreasing log-average means and variances (not shown).

To further characterize how neural activity fluctuates over broad time scales, the duration of activity of each neuron was calculated as the time between its first and last spike. This was then normalized by total track duration to give the fractional duration that each neuron was active for. The distribution of fractional active durations of all neurons is shown in Figure 4.2C. About 40% of neurons had active durations close to 1, i.e., they were active for their entire track. The other 60% of neurons were active for some shorter fraction of their track, and the distribution of these fractions was surprisingly uniform. There was no characteristic activity duration timescale for these 60% of neurons. Note that some of the 40% of neurons making up the peak at 1 in Figure 4.2C might be there only due to the experimentally limited duration of each track. Had they been recorded from for longer, there would have been greater opportunity for them to stop firing, and they might therefore have moved out of the peak at 1 and into the uniform part of the distribution below 1. Analogous plots of first and last spike times as a fraction of track duration (not shown) were consistent with the above result. First spike times peaked close to 0 but were otherwise uniform, and last spike times peaked close to 1 but were also otherwise uniform.

4.4 Templates & positions

All spikes from a sorted unit were used to calculate its average multichannel waveform, or "template". Templates of all 245 sorted cells from all 3 tracks are shown in Figure 4.3. Many channels had multiple cells centered on them (left panels). Furthermore, many cells had similar, even identical, templates on their primary channel, and could only be distinguished from each other by considering their template differences on neighbouring channels with lower amplitude signal (right panels). For example, one primary channel in ptc15.tr7c was shared by three cells (yellow, red, brown) with very similar primary channel templates (Figure 4.3, inset, left panel, white arrow). However, all three neurons were distinct from one another on 2 of the 4 neighbouring channels (right panel, grey arrows). This demonstrates the necessity of having many closely spaced sites to help distinguish neighbouring cells (Blanche et al., 2005).

The spatial position of each neuron, shown in Figure 4.4A, was taken as the median spatial position of all of its member spikes (see Section 3.5 for spatial localization of spikes). Cells were reasonably well distributed across the length of their respective polytrode, with no particular bias towards any one part of the polytrode. However, x coordinates of cells were strongly biased toward the x coordinates of the electrode sites (Figure 4.4A, top). There was a similar but weaker bias of cell y coordinates to electrode site y coordinates (Figure 4.4A, left). For the y position histogram, bins were placed such that there were equal numbers of bins aligned with and in between electrode site y coordinates. There were 26 local peaks in the y position histogram that aligned with electrode sites, and only 3 local peaks in between electrode sites (90% and 10%, respectively). 154 cells fell in site-aligned bins, while only 91 fell in inter-site bins (63% and 37%, respectively). The bias in x and y cell coordinates to x and y electrode site coordinates suggests that cell positions were biased towards their nearest electrode site.

To further examine this, the distribution of 2D distances between each neuron and its nearest electrode site was measured (Figure 4.4B). Cells were most likely to be found close to electrode sites, with decreasing probability at greater distances. The expected distribution for randomly positioned cells was analytically estimated by approximating the hexagonal domain surrounding each electrode site as a circle. The number of points N at a distance r from the site is the area of an infinitesimally thin concentric ring, times the average areal cell density ρ measured from the data:

$$N = 2\pi r\rho \, dr \tag{4.1}$$

This analytical estimate is shown as the dashed blue line in Figure 4.4B, and ignores the edge effects of a finite polytrode. A uniform random distribution of 100,000 points was also numerically simulated for the 3 column probe (1a design), resulting in the dashed red line in Figure 4.4B. Both expected distributions for randomly positioned neurons increased linearly as a function of nearest electrode site distance, while the measured data showed the opposite. This again suggests that cell positions were not random, but rather were biased to the nearest electrode site. Moreover,

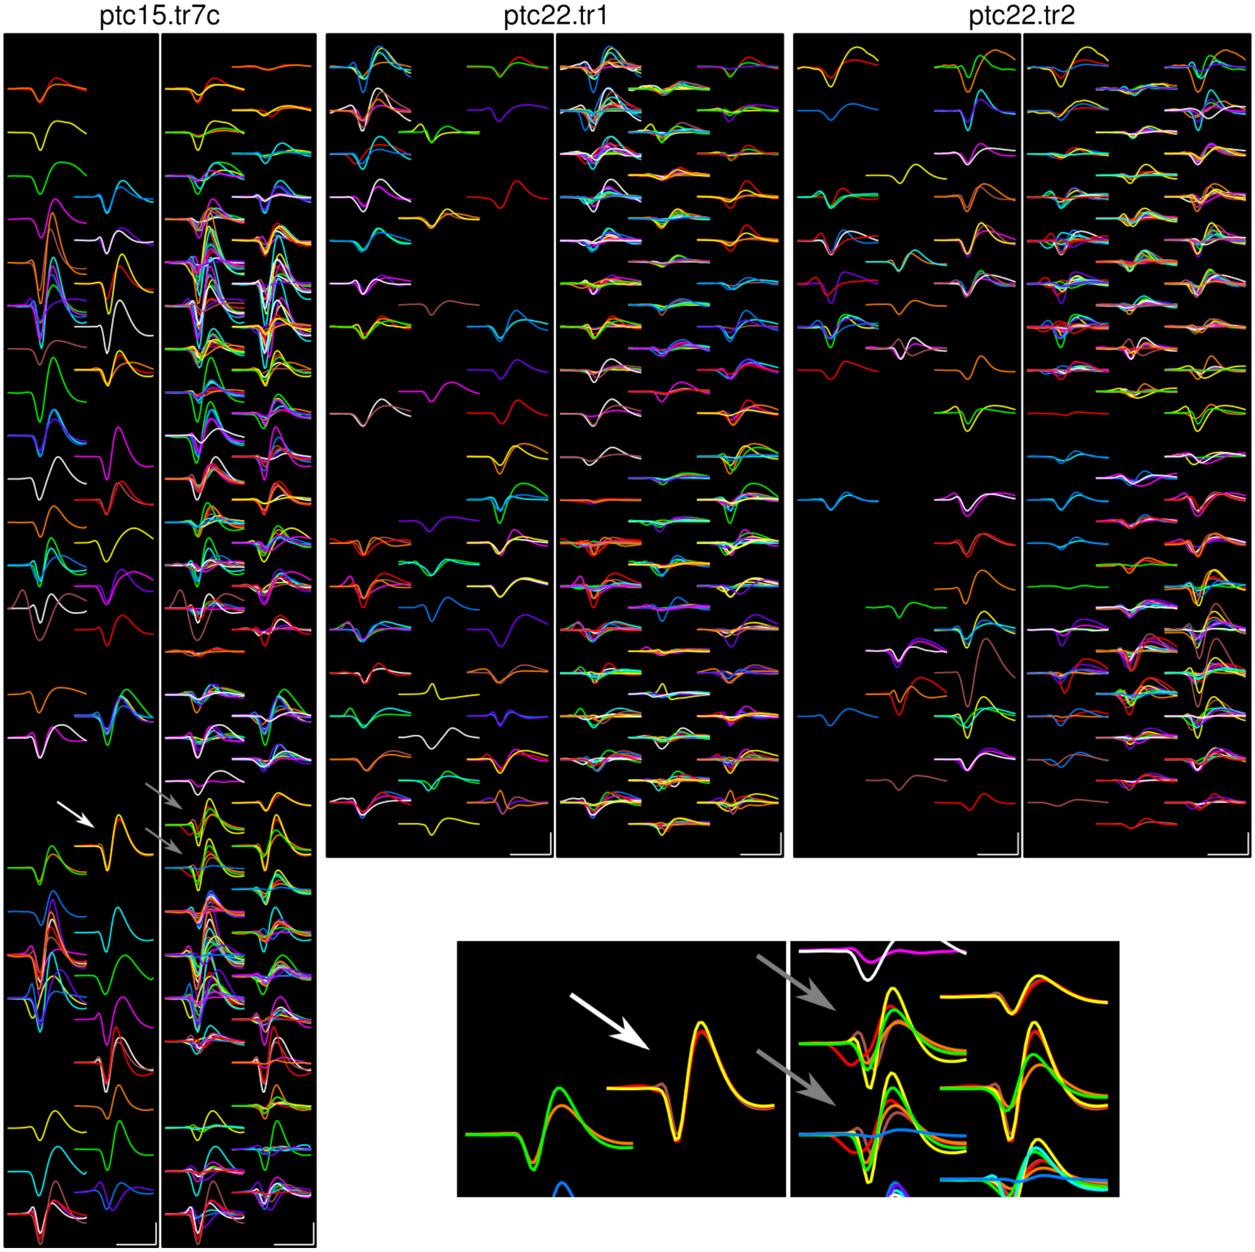

Figure 4.3: Single (primary) channel templates (***left*** in each pair) and primary channel plus nearest neighbour multichannel templates (***right*** in each pair) of all sorted cells from all three tracks. Surrounding channels helped distinguish cells when two or more shared primary channels with very similar signal. Each template was aligned to its primary channel. For nearest neighbour multichannel templates, included surrounding channels did not necessarily align the same way. Blank channels either had no cells for which they were best positioned, or were disabled during recording. A limited set of 10 colours was cycled through in vertical spatial order to help distinguish cells. The white arrow highlights a channel on which three cells (yellow, red, brown; enlarged in inset) had nearly identical primary channel templates, and were only distinguishable on two neighbouring channels (grey arrows). Scale bars: 0.5 ms, 100 μV.

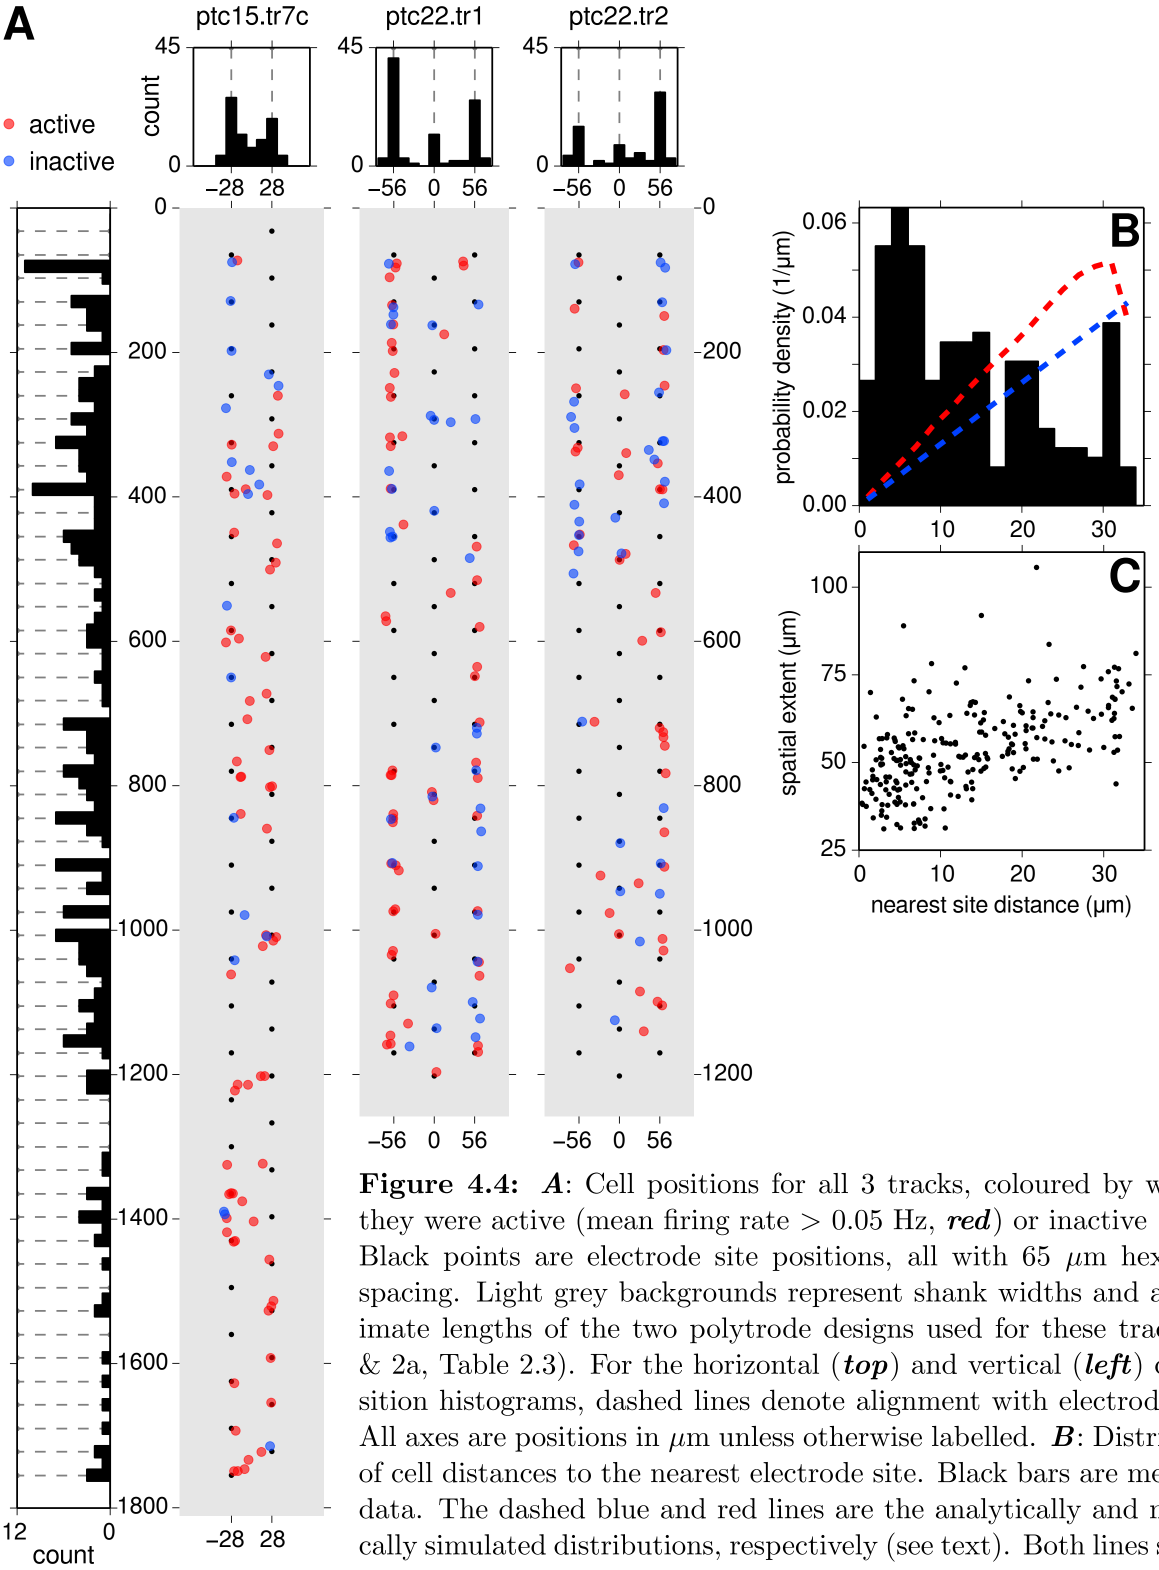

Figure 4.4: ***A***: Cell positions for all 3 tracks, coloured by whether they were active (mean firing rate > 0.05 Hz, ***red***) or inactive (***blue***). Black points are electrode site positions, all with 65 μm hexagonal spacing. Light grey backgrounds represent shank widths and approximate lengths of the two polytrode designs used for these tracks (1a & 2a, Table 2.3). For the horizontal (***top***) and vertical (***left***) cell position histograms, dashed lines denote alignment with electrode sites. All axes are positions in μm unless otherwise labelled. ***B***: Distribution of cell distances to the nearest electrode site. Black bars are measured data. The dashed blue and red lines are the analytically and numerically simulated distributions, respectively (see text). Both lines showed a trend opposite that of the data. ***C***: Plotting each cell's spatial extent vs. distance to the nearest electrode site showed a correlation, suggesting that cells with small closed fields were missed in between sites.

the spatial extent of cell templates (each the median spatial extent of its spikes, Section 3.5) was positively correlated with the distance to the nearest electrode site (Figure 4.4C), suggesting that the positional bias was due to cells with small closed extracellular fields that were too distant from an electrode site to be detected.

There was no strong spatial segregation of active cells (mean rates > 0.05 Hz, Figure 4.4, red dots) and inactive cells (mean rates < 0.05 Hz, blue dots) as a function of position along the polytrode. However, there was a weak tendency for inactive cells to be more superficial.

Cell position was also examined as a function of time, as were other spike parameters. The x and y coordinates, spatial extent (σ), and peak to peak amplitude (V_{pp}) of all 5.1 million spikes from all 3 tracks are shown in Figure 4.5. For better visibility, Figure 4.6 shows the same data, with spike parameters averaged over 10 min wide non-overlapping time bins, resulting in lines representing one cell each. Track ptc15.tr7c showed the most stability, especially in y position. Tracks ptc22.tr1 and ptc22.tr2 were both less stable in y. Many cells in those two tracks appeared to drift upward, in tandem, over a span of hours, suggesting polytrode drift relative to the tissue. A few cells drifted upward by as much as ~ 150 μm. However, many other cells in those two tracks had no upward drift, or had periods of downward drift, and some of these were positioned in between cells that did have upward drift. This cannot be explained by global, or even local, polytrode-tissue drift. In comparison to y position, x position, σ and V_{pp} were all reasonably stable across all tracks.

To further investigate stability in y, dy/dt was calculated from the binned values in Figure 4.6. The distribution of dy/dt values across all time bins and cells within each track are shown in Figure 4.7A. Contrary to the visual impressions from Figures 4.5 & 4.6 of consistent upward drift for many cells in ptc22.tr1 and ptc22.tr2, the distributions of dy/dt for all 3 tracks were nearly symmetric around $dy/dt = 0$, suggesting that at least at the timescale of 10 min wide bins, there was little systematic upward drift. Mean dy/dt was positive (upward) but very low for all three tracks, ranging 0.7–3.4 μm/h. However, dy/dt distributions were broader for the latter two tracks ($\sigma = 64$–78 μm/h), indicating greater vertical positional variability than for ptc15.tr7c ($\sigma = 31$ μm/h). Net rates of change in y position between the first and last spike of each cell ($\Delta y/t$) were also calculated, and their distributions are shown in Figure 4.7B. Results were similar to dy/dt: the means were positive but very low for all three tracks, ranging 2–2.9 μm/h, and the distributions for ptc22.tr1 and ptc22.tr2 were broader than for ptc15.tr7c.

4.5 Orientation tuning

Oriented responses are a characteristic feature of neurons in V1 (Hubel and Wiesel, 1959), and measuring orientation tuning is a standard procedure. Three different types of artificial stimuli were used to measure orientation tuning of neurons: drifting bars, flashed gratings, and drifting gratings (see Section 2.4 for stimulus details).

Depending on the stimulus type, steps in orientation ranged from as fine as 10° to as coarse

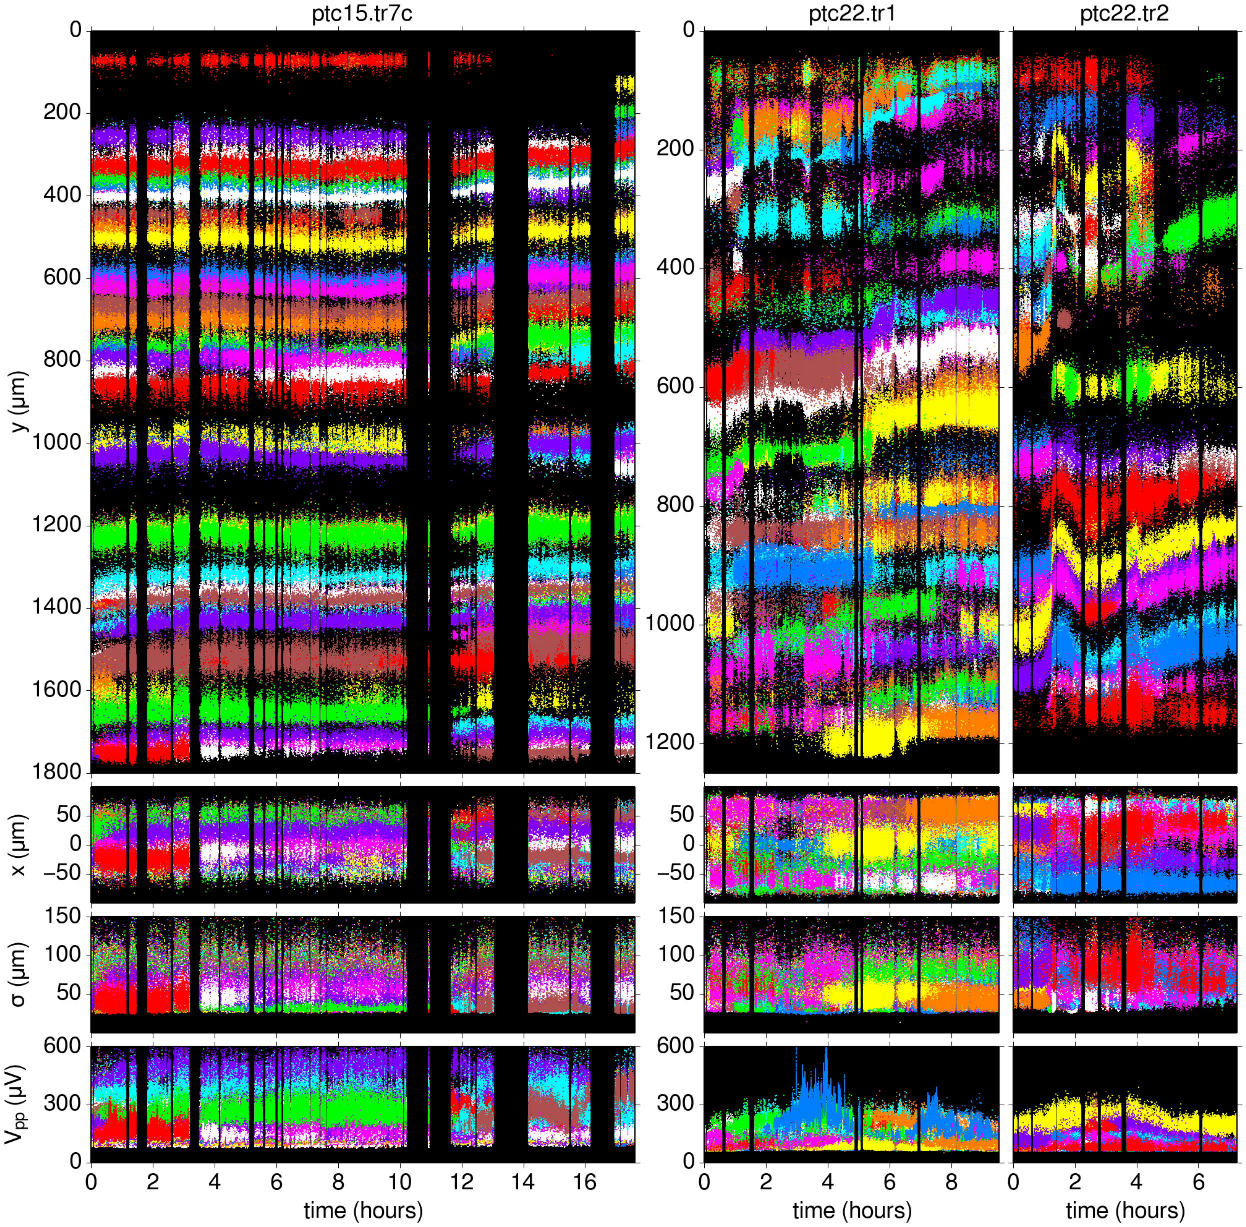

Figure 4.5: Spike parameters as a function of time for all 5.1M spikes from all cells of each track. Each point represents a spike, coloured by cell, with colours corresponding to those in Figure 4.3. From top, plotted spike parameters are vertical (y) and horizontal (x) position, spatial spread (σ), and primary channel peak-to-peak voltage (V_{pp}). Positions are relative to the polytrode used for each respective track. Vertical black lines represent recording time gaps. Many spikes are obscured, especially those of low firing rate cells. To maximize spike visibility, different scales were used for x and y, and a different y scale was used for ptc15.tr7c. y roughly corresponded to cortical depth, but less so in ptc15.tr7c due to its angled insertion (Figures 5.6 & 4.10C). Depending on the track, dy/dt was moderate to high, but dx/dt was relatively low. Temporal variation in σ remained fairly low, while V_{pp} varied moderately.

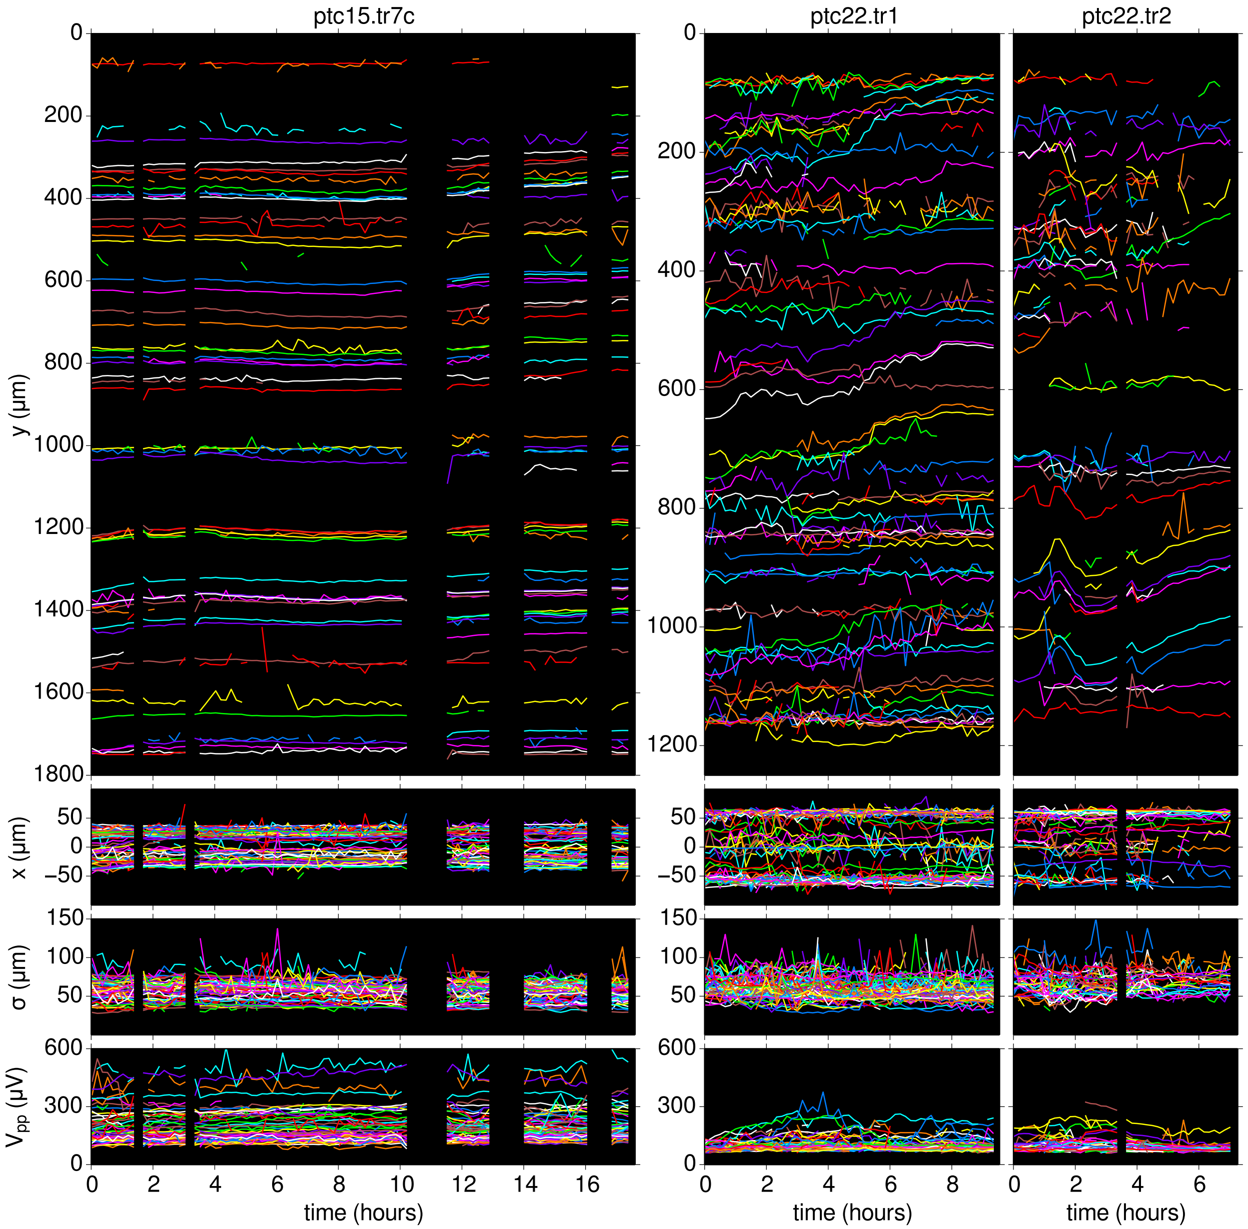

Figure 4.6: Smoothed spike parameters vs. time. Same as Figure 4.5, but with points averaged by binning into 10 min wide non-overlapping bins and plotting the mean value within each bin. Consecutive bin values of each neuron are represented by connected lines. This results in a smoother signal with better visibility of all cells, at the cost of a loss of representation of parameter variance and cell firing rate. Disconnected lines are either due to gaps in recordings or time bins in which the cell did not fire at all.

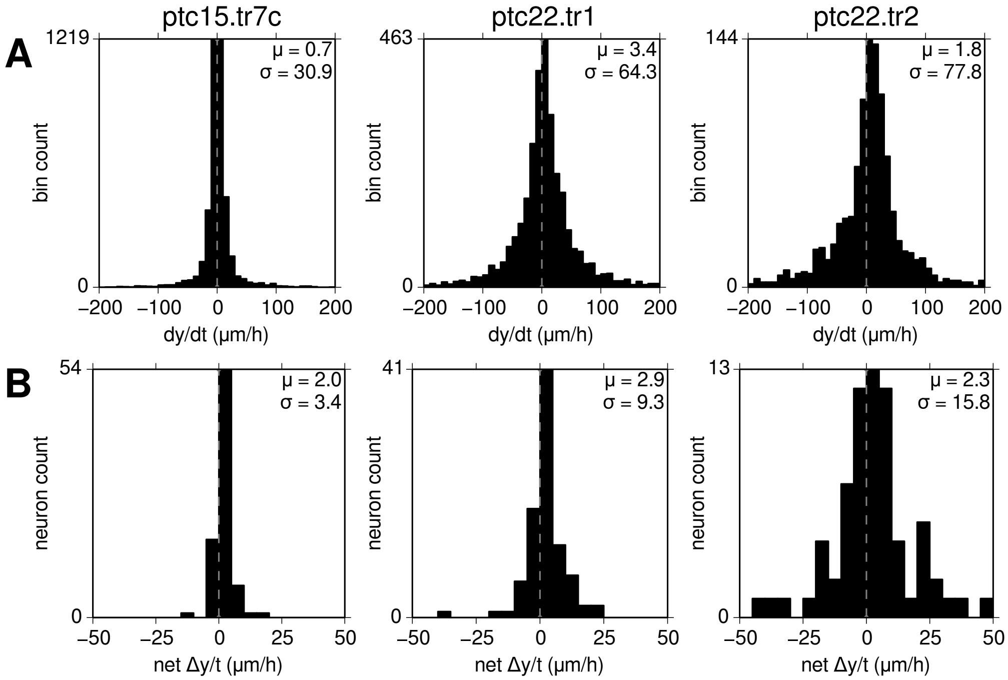

Figure 4.7: Distributions of changes in y position over time for all cells in each track. ***A***: Distributions calculated from the uppermost 3 panels in Figure 4.6, using the same 10 min wide non-overlapping time bins. Positive dy/dt corresponds here to movement up the polytrode, so in this case, y represents distance from the bottom of the polytrode instead of the top. Means and standard deviations are shown for each track. Although many cells in ptc22.tr1 and ptc22.tr2 appeared to drift upwards in Figure 4.6, surprisingly their slope distributions were highly symmetric about $dy/dt = 0$, with little positive skew. While some cells moved up the polytrode at times, they were mostly balanced out by other cells moving down the polytrode. ***B***: Distributions of net rates of change in y position between the first and last time bin of each cell. Note the finer scale on the x axis. Although standard deviations were lower than in ***A***, the means were similar and the same conclusions held.

as 45°. When calculating orientation tuning curves (Figure 4.8, bottom), spikes were counted across all stimulus conditions, but the stimulus values of non-orientation stimulus dimensions were ignored. This means that tuning curves were composed of responses collapsed across all other stimulus dimensions including brightness, contrast, spatial and temporal frequencies, and phase.

Tuning curves for flashed gratings were calculated by forward correlating each 40 ms duration flashed grating with the spikes that occurred 40 to 80 ms after stimulus onset. Although different cells may have different response delays to stimulus onset (Figure 5.6), 40 ms was a good compromise across the population, and was therefore chosen as both the flash duration and the reverse correlation delay. Flashed gratings allow for rapid simultaneous characterization of many stimulus dimensions for each cell, much faster than drifting gratings (Ringach et al., 1997a,b). However, temporal frequency preferences cannot be calculated because the stimulus contains no temporal frequencies (apart from the 25 Hz frame frequency from the 40 ms frame time).

Orientation preference was calculated from the tuning curve of each cell to each stimulus type by taking the vector mean of spike counts at each evenly spaced stimulus orientation (Swindale, 1998). Specifically, this required doubling all the stimulus orientation angles, performing the vector sum, and then dividing the resulting length by the total number of spikes, and dividing the resulting angle by 2. Finally, modulus 180 of the angle was taken to constrain it to between 0° and 180° for motion stimuli that go around the clock (drifting bars and drifting gratings). The resulting vector length, r, represented tuning strength, and ranged from 0 (low tuning) to 1 (high tuning). Cells were tested for orientation tuning significance using Rayleigh's test for circular uniformity (Wilkie, 1983; Fisher, 1995). For each significantly tuned cell, the stimulus that resulted in the sharpest tuning curve was used to characterize the orientation tuning of that cell. The stimulus type that resulted in the strongest orientation tuning varied from cell to cell, with no apparent regularity.

Overall, 61% (150/245) of cells were significantly tuned to orientation ($p < 0.01$). The numbers per track were 62% (50/81), 57% (53/93), and 66% (47/71) for ptc15.tr7c, ptc22.tr1, and ptc22.tr2 respectively. However, when restricted only to cells that were active (with mean firing rates $\geq$ 0.05 Hz) during at least one oriented stimulus, 87% (131/150) of cells were significantly tuned to orientation ($p < 0.01$), with 98% (46/47), 77% (48/62), and 90% (37/41) per respective track.

Spiking responses to drifting gratings and the associated tuning curves of 3 example neurons are shown in Figure 4.8. Some cells exhibited strong and significant orientation tuning despite extremely low firing rates, with orientation preferences in line with those of their higher firing rate neighbours. The example neuron in Figure 4.8C fired only 8 spikes during the entire 22 min drifting bar recording, with a mean firing rate of 0.006 Hz, yet it was significantly orientation tuned ($p = 0.0057$) with a tuning strength of $r = 0.76$. All 8 spikes occurred near the middle of the relevant trials when the bar was at the same position, presumably over the cell's RF. Like other well tuned moderately direction selective cells, it had two peaks in its tuning curve 180° apart. Finally, this cell's preferred orientation of 30.6° was similar to that of its neighbours (Figure 4.10C, middle), providing further evidence of significant tuning (not considered in p value calculation).

There were other similarly low firing rate neurons that were significantly tuned, with orientation preferences similar to their neighbours. Figure 4.9A shows how tuning strength varied with mean firing rate, which was calculated for each of the 150 significantly tuned cells over the duration of the stimulus that resulted in that cell's strongest orientation tuning. Surprisingly, on a linear-log plot of tuning strength vs. mean firing rate, the relationship was inverse: low mean firing rate cells generally had the highest tuning strength, and vice versa. This inverse and approximately linear relationship ($r = -0.5$) held regardless of the significance threshold. For significance thresholds of 0.05, 0.01, 0.001 and 0.0001, the number of significantly tuned cells was 164, 150, 142 and 134, respectively, with r values of -0.56, -0.50, -0.51 and -0.53.

Figure 4.9B shows that the distribution of mean firing rates of the 150 significantly tuned cells was roughly a lognormal distribution, similar to that of the full population shown in Figure 4.1B,

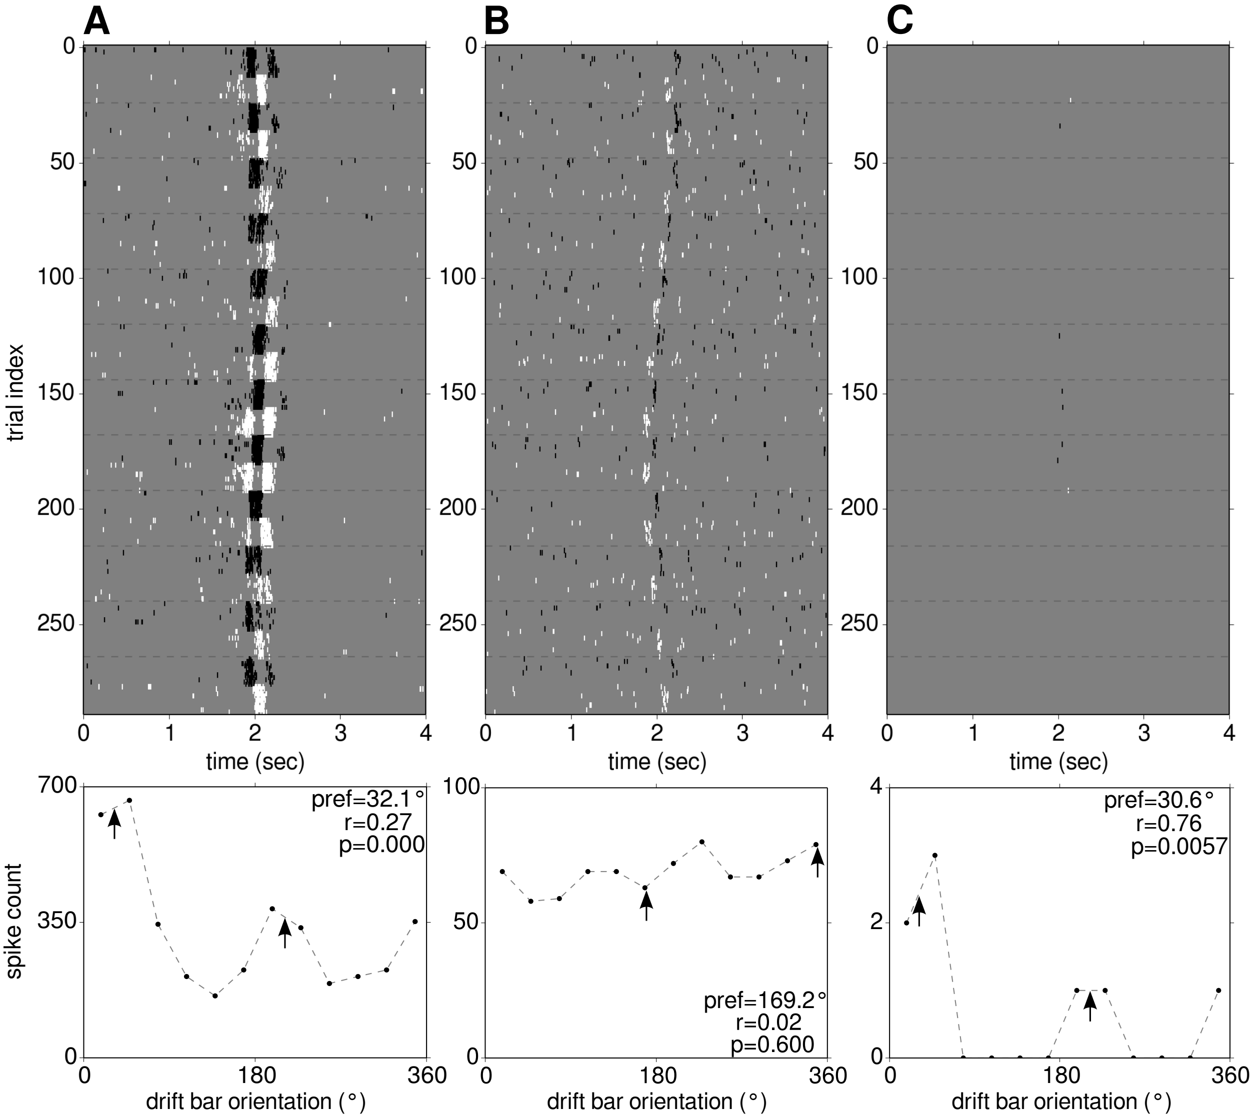

Figure 4.8: Orientation tuning curves calculated from the responses of 3 simultaneously recorded example neurons in track ptc22.tr1 to 288 trials of drifting bar stimulus. ***Top***: Raster plots of all trials, in order of stimulus parameter combination (not temporal order). Responses to white and black drifting bars are represented by white and black ticks respectively, where each tick represents one spike. Trial 1 started at 198° and orientation incremented by 30° every 24 trial indices. Horizontal dashed lines denote orientation increments. ON and OFF subfields (Section 5.4) were discernable in ***A*** and ***B*** as bar brightness alternated every 12th trial. ***Bottom***: Orientation tuning curves with orientation preferences (arrows) calculated using the vector mean of trial spike counts. Tuning strength, represented by normalized vector length r, and p values (Rayleigh test for circular uniformity) are shown. ***A*** was moderately tuned with high significance and ***B*** was very weakly tuned with very low significance. Despite firing only 8 spikes, ***C*** was well tuned and significant ($p < 0.01$). Like ***A***, ***C***'s orientation preference of $\sim 30°$ was similar to that of its neighbours.

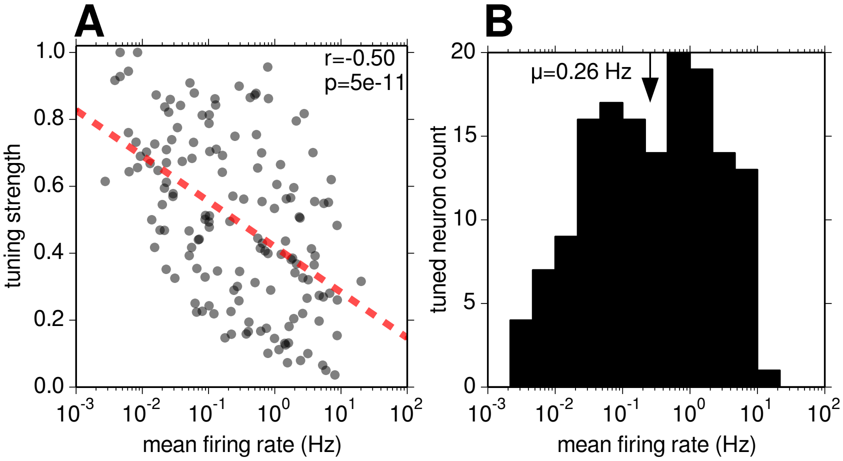

Figure 4.9: ***A***: Orientation tuning strength vs. mean firing rate for the 150 significantly tuned cells. The red dashed line is the least-squares linear regression (r and p values are shown). Surprisingly, tuning strength was inversely correlated with log mean firing rate. ***B***: The mean firing rate distribution for the same set of cells showed weak bimodality. The arrow is the geometric mean.

though with a slightly higher geometric mean of 0.26 Hz vs. 0.11 Hz for the full population. Also, restricting the mean firing rate distribution to only tuned cells revealed a weak bimodality, suggesting two different orientation tuned cell types. This weak bimodality persisted for significance thresholds of 0.05, 0.01 and 0.0001, but not 0.001 (not shown).

The orientation preferences of all 150 significantly tuned cells are shown in Figure 4.10. The full range of possible orientation preferences and tuning strengths was represented across the population of significantly oriented cells, with no particular favoured orientation (Figure 4.10A). To allow for a very rough comparison of cell depths along polytrodes of different lengths, normalized depth was calculated for each cell as the fractional distance down the length of the polytrode. Normalized depth is represented by point darkness in Figure 4.10A, while Figure 4.10B shows normalized depth as a function of orientation preference, with points coloured according to tuning strength. As expected, when pooled across all 3 tracks, there was no relationship between normalized depth and orientation preference. However, when taken individually, each track showed relationships between orientation preference, tuning strength, and normalized depth (Figure 4.10C). For ptc15.tr7c orientation preference varied strongly with normalized depth, at a rate of 104°/mm (counterclockwise), suggesting (as in Figure 5.6) that this was a transcolumnar track. For track ptc22.tr2, orientation preference also varied with depth, at a rate of 79°/mm (clockwise), but this relationship was weaker (statistical significance not calculated).

4.6 Discussion

4.6.1 Neuron yields

Track-wide sorting increased overall neuron yield (Section 4.2). This is unsurprising, given that low firing rate cells (Section 4.3) take a long time to accumulate enough spikes to be detected as clusters. However, further analysis is required to determine if most of the additional units gleaned

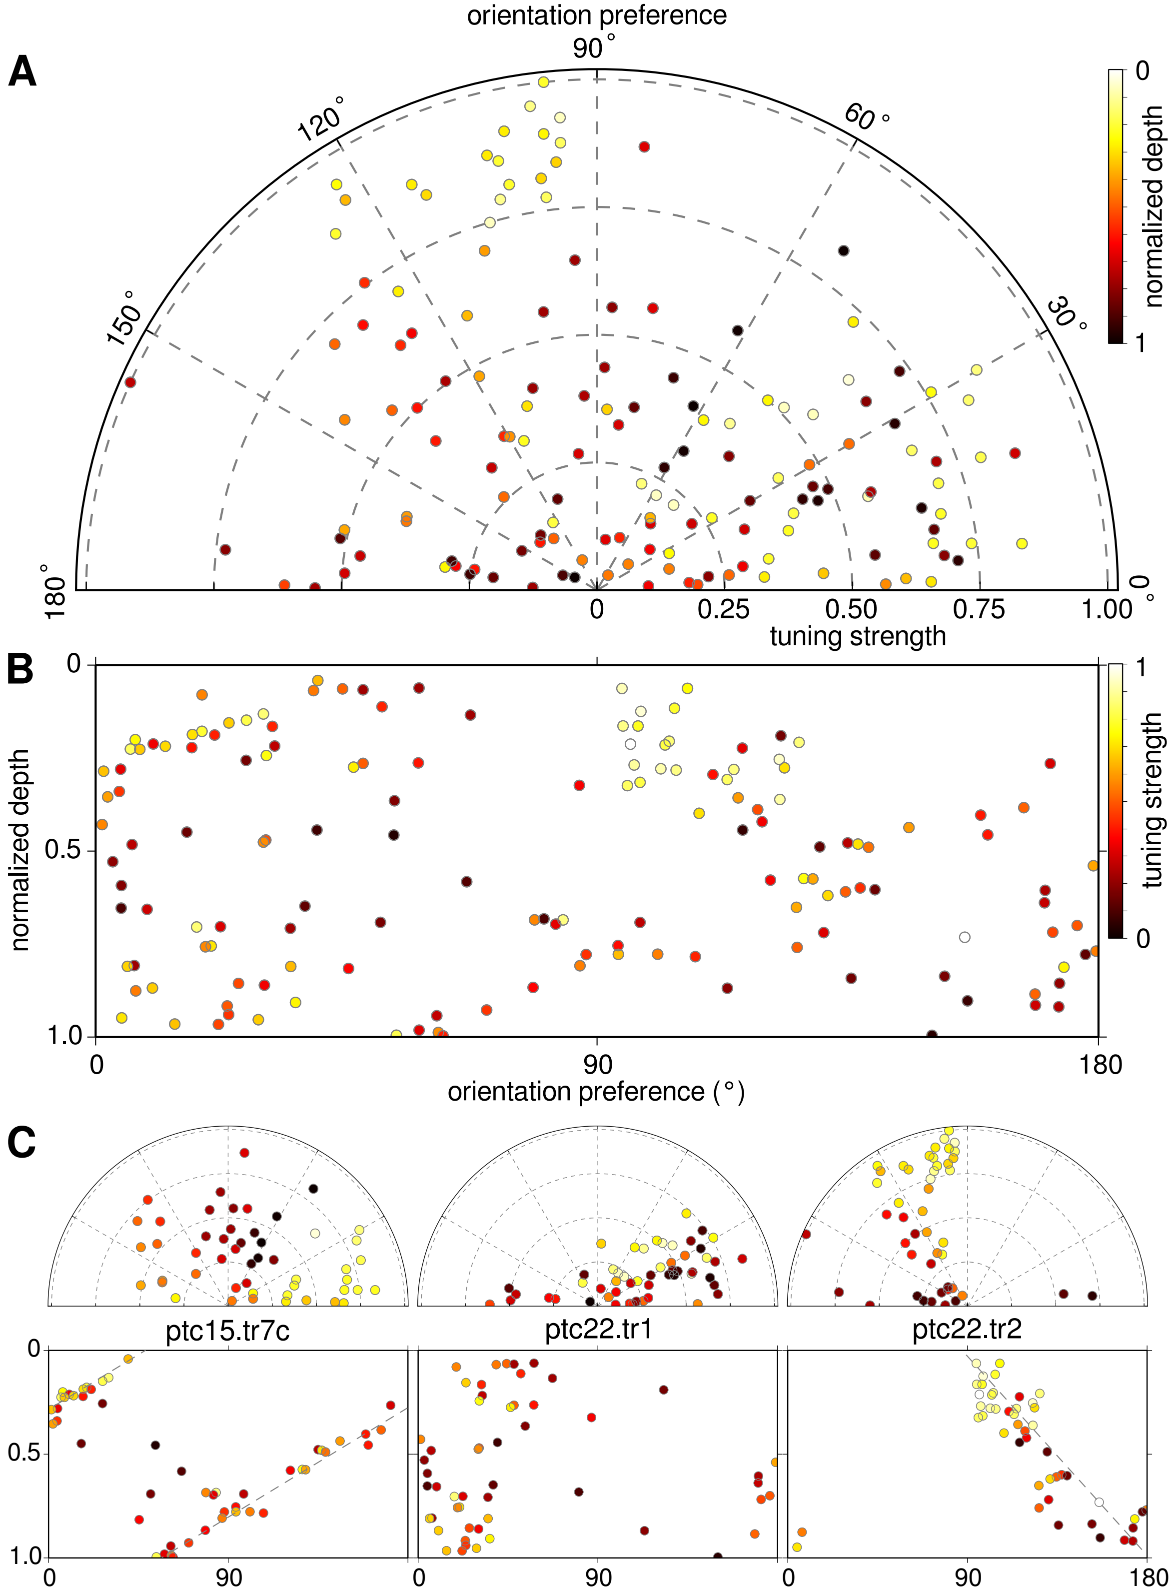
A
orientation preference
90°
120°
60°
150°
30°
180°
0°
0
0.25
0.50
0.75
1.00
tuning strength
0
normalized depth
1
B
0
0.5
1.0
normalized depth
0
90
180
orientation preference (°)
1
tuning strength
0
C
ptc15.tr7c
ptc22.tr1
ptc22.tr2
0
0.5
1.0
0
90
0
90
0
90
180

Figure 4.10: (*Previous page.*) ***A***: Orientation preference, tuning strength, and normalized depth of the 150 cells with significant tuning ($p < 0.01$, Rayleigh test) to at least one drifting bar, flashed grating, or drifting grating experiment. Point darkness represents normalized cell depth along the length of the polytrode. ***B***: Orientation preference as a function of depth along the polytrode. In this case, point brightness represents tuning strength. In both ***A*** and ***B***, there were no obvious relationships between orientation preference, tuning strength, and cell depth, and all were reasonably uniform across the population. ***C***: Same as ***A*** and ***B***, but broken down according to track. ptc15.tr7c and ptc22.tr2 showed a linear relationship between orientation preference and depth (highlighted by manually positioned dashed lines), with strong and weak relationships, respectively, suggesting that these tracks were strongly and weakly transcolumnar, respectively.

from track-wide sorting are low rate cells, cells that respond only to certain stimuli, or cells that are only active at certain times for other reasons, such as a change in cortical state (Chapter 6).

Although not investigated in any detail, perhaps the fast single channel spikes found during the first few minutes after polytrode insertion were inhibitory interneurons (Section 5.1) whose activity was inhibiting that of pyramidal cells. Or perhaps more likely, the fast single channel spikes were axonal spikes from cut axons whose cross-section happened to fall near an electrode site.

Given known cortical neuronal densities and estimates of recording volumes of extracellular electrodes, only about 10% of the population of cells that must be physically present and within recording range of extracellular electrodes are typically detected (Henze et al., 2000; Shoham et al., 2006). This fraction must depend to some extent on spike sorting quality. Shoham et al. (2006) cited Blanche et al. (2005) as having found 60 neurons out of potentially 700, i.e., 8.6%, estimated from the polytrode recording volume and neocortical neuron density. With 82 units reported here for that same track (ptc15.tr7c), that percentage increased, but is still only 11.7% of the cells that should be there. This “dark neuron” problem might be due to most neurons remaining silent, and therefore undetectable in extracellular recordings, for long periods of time (Thompson and Best, 1989; Henze et al., 2000; Hahnloser et al., 2002; Kerr et al., 2005). Stable continuous recording for even longer periods of time (days) and under as wide a variety of stimulus and animal conditions as possible (awake, asleep, anesthetized, behaving, learning) may help resolve this problem.

In addition to cats, polytrode recordings were also made in V1 of 13 urethane-anesthetized rats (Section 2.1). However, very few units were isolated, and even fewer were visually responsive (not shown). Given that many other groups successfully record from rat brain using polytrodes with narrower shanks ($< 100\ \mu$m) (Harris et al., 2003; Csicsvari et al., 2003; Buzsáki, 2004; Sirota et al., 2008; Goard and Dan, 2009; Schjetnan and Luczak, 2011; Luczak et al., 2013; Berényi et al., 2014), this suggests the interesting possibility that the wide ($\sim 200\ \mu$m) polytrodes used here do more damage in smaller brains. Further investigation of this hypothesis would require modelling the number of dendrites and axons that might be cut by polytrodes with different numbers of shanks of different widths and spacing, taking into account the density of cells and the size and shape of their dendritic and axonal fields in different species. This also suggests that using narrower polytrodes

(Section 4.6.3), with either narrower or layered conductors, would further decrease damage and increase neuron yield in cat cortex.

4.6.2 Firing rates

Mean firing rates in anesthetized cat V1 were lognormally distributed with a log-average firing rate of ~ 0.11 Hz and a standard deviation of about one order of magnitude (Figure 4.1B). Although much lower than reported in most studies (see Olshausen and Field (2005) & Carandini et al. (2005) for reviews), this value and the lognormal distribution around it is nevertheless in line with a handful of reports in other species and cortical areas. Using energy analysis, Lennie (2003) estimated that the average firing rate of human cortical cells is 0.16 Hz. Brecht et al. (2003) found that average spontaneous firing rates were 0.068 Hz in superficial layer pyramidal cells in urethane-anesthetized rat barrel cortex. That study used whole-cell recordings whose electrode impedance changes upon contact with a neuron, providing confirmation of neuron isolation even in the absence of spiking activity. This reduces the probable bias of individually maneuverable extracellular single-wire electrodes or tetrodes towards high rate neurons. Kerr et al. (2005) found a mean spontaneous firing rate of 0.05 Hz in superficial layers in urethane-anesthetized rat primary somatosensory cortex (S1) and primary motor cortex (M1). That study used both two-photon calcium imaging and cell-attached recordings, with bias reduction benefits similar to those of whole-cell recordings. Hromádka et al. (2008) found a lognormal distribution of firing rates in primary auditory cortex (A1) of awake rat, using cell-attached recordings. That study used a wide range of auditory stimuli, including naturalistic sounds. Their log-average was higher (~ 3 Hz) than reported here, but that should be expected given that firing rates are higher in auditory cortex. Sakata and Harris (2009) also reported a lognormal distribution of mean spontaneous firing rates across superficial and deep layers of awake and urethane-anesthetized rat A1 (their Figure S14B). Using silicon polytrodes, they found log-averages of 0.5–3 Hz. Finally, Mizuseki and Buzsáki (2013) reported lognormal firing rate distributions in awake and asleep rat hippocampus and entorhinal cortex, with log-averages of 0.5–1 Hz. In that study, rate distributions were best fit by a lognormal curve during slow wave sleep, but were somewhat skewed during REM sleep and awake states.

Given the increasing evidence for lognormal (as opposed to normal) distributions of mean firing rates across species and cortical areas, the geometric mean may be more appropriate than the arithmetic mean for describing the central firing rate tendency of populations of cortical neurons. Furthermore, lognormal rate distributions are consistent with the concept of sparse coding, in which only a small fraction of the population need be active at any given time to encode a stimulus, memory, or action (Olshausen and Field, 1996). The benefits of sparse coding include computationally efficient memory storage, efficient extraction and representation of structure within natural stimuli, and energy efficiency (Olshausen and Field, 2004; Attwell and Laughlin, 2001).

Wohrer et al. (2013) extensively reviewed population distributions of firing rates in various species and neocortical regions. That study made conclusions about publicly available sorted data from cat ptc15, which was sorted using the previous spike sorting technique based on multichannel template matching (Section 3.1.1). They concluded that population mean rates from those data are best fit by a decaying exponential distribution. In linear-log space, such a distribution would have the shape of a decreasing sigmoid, which clearly does not fit the distribution shown in Figure 4.1B. In contrast, the LM least-squares best fit lognormal distribution shown in that figure does provide a reasonable approximation. This is consistent with lognormal distributions described by Wohrer et al. (2013) in other species (rat and primate) and cortical areas (A1, prefrontal cortex (PFC), M1 and secondary motor cortex (M2)). Spike sorting quality may therefore affect conclusions about firing rate distributions.

The majority of cells, 82%, had mean firing rates below 1 Hz (Figure 4.1). This is significant because for many studies, ~ 1 Hz is considered the minimum mean firing rate for unit inclusion for analysis (Olshausen and Field, 2005). Depending on the area of recording, such studies may be inadvertently biased towards the properties of high firing rate cells. Perhaps the biggest reason for setting such a threshold is that sorting very low firing rate units from those that fire at higher rates is difficult (Section 3.8). The task was made easier in this case by the use of ICA instead of PCA, when appropriate.

Some of the lowest mean firing rates might be from cells that were active for only a short period of time. Several such cells are apparent in Figure 4.6. While it is possible that these were cells whose single unit isolation was gained or lost part way through a track, this seems less likely than such cells simply starting or stopping their firing, whether due to a change in stimulus, cortical state (Section 6.4), or some other unknown cell-specific reason. Barring sudden electrode movement, loss of cell isolation is usually a gradual process due to gliosis or tissue swelling (Liu et al., 1999; Szarowski et al., 2003; Biran et al., 2005), accompanied by a gradual decrease in spike amplitude (see reviews by Polikov et al. (2005) & Leach et al. (2010)). No such gradual spike attenuation was encountered during spike sorting, and none was apparent in any of the sorted units (Figure 4.6, bottom row). Since the polytrode is a single rigid structure with fixed electrode site locations, if it were to shift suddenly, simultaneous isolation loss of many cells would be expected. This was not observed in any of the 3 sorted tracks.

At a time scale of 10–20 min, firing rates of individual cells could fluctuate by orders of magnitude, while the population geometric mean rate remained fairly stable (Figure 4.2). This suggests the possibility that neuronal populations are engaged in a kind of shift work, where some cells are active for a period of time, while others are silent, with group membership turning over as a function of time, and network functionality remaining intact throughout. This could perhaps allow for offline physiological cell maintenance, independent of sleep. It may also have implications for network stability, such as preventing synaptic weights from growing out of control and giving too much weight to too few cells (Markram and Tsodyks, 1996). Kerr et al. (2005) provide similar evidence for shift work using two-photon calcium imaging in urethane-anesthetized rat cortex. However, further work is required to demonstrate neural shift work more definitively. One way would be to measure the (log-scale) variance in multiunit activity (MUA) of spike trains randomly shifted in time by different amounts for each neuron, and compare that to the MUA variance of the original spike trains. If temporal shifting increases MUA variance, that would strengthen the evidence for neural shift work.

It is possible that the wide fluctuations in firing rates of individual neurons was due to rapid gain or loss of single unit isolation due to drift. Cells with small closed extracellular fields (single channel cells) might be especially prone to sudden changes in unit isolation. However, while the traces in Figure 4.2 show that individual firing rates fluctuated greatly at different depths and times, these fluctuations did not seem to coincide with sudden changes in y position in Figure 4.6. Further analysis is required to explicitly show this, such as plotting each cell's firing rate as a function of V_{pp} amplitude or y position, to confirm that firing rate is independent of both of these variables. However, simultaneous gain and loss of single unit isolation of many different cells across the polytrode seems an unlikely explanation for the wide fluctuations in firing rate of individual cells shown in Figure 4.2.

Anecdotally, conventional single-wire extracellular electrophysiology can suffer problems "holding" cells for no more than a few hours, followed by their sudden isolation loss. This is generally attributed to the electrode suddenly shifting its position relative to the neuron, even if only by a small amount. However, when slowly approaching or retreating from a neuron using a microdrive, a neuron may remain isolatable over the span of 100 μm or more, contradicting the assumption that isolation loss is due to fine sensitivity to electrode position relative to the neuron. An alternative explanation that resolves this apparent contradiction is that this sudden apparent isolation loss instead reflects dramatic changes in firing rate (Swindale and Spacek, 2012).

Extremely low firing rates could plausibly be an artifact of high drift, causing a single unit's spikes to be split over time into two (or more) clusters, with insufficient evidence to combine them during spike sorting. For example, if the spike counts in the two resulting clusters were equal, this would create two units with exactly half the mean firing rate of the original unit (half the spikes divided by the same full track duration). This would increment that particular half mean firing rate bin in Figure 4.1B twice instead of once at the original mean rate. Since low firing rate cells have fewer spikes (and therefore less evidence for merging) to begin with, perhaps such splitting occurs more often for low firing rate cells, thereby overestimating their numbers. Although one might suspect a few such candidate oversplit pairs while examining changes in unit position as a function of time (Figures 4.5 & 4.6), those plots represent only one narrow view of the sorted units. Many other dimensions were inspected during spike sorting before a decision was made to merge or split (Section 3.8). Even if some clusters were oversplit, it seems safer to oversplit cells and thereby underestimate their mean firing rates than to undersplit them and mistake multiunits for single units in later analyses. Nevertheless, the approximate symmetry in log space of both the high and low end of the mean firing rate distribution in Figure 4.1B suggests that oversplitting was not a common occurrence.

Time gaps between some recordings were up to 30 min long, and should be minimized in the future, both to improve tracking of physiological properties such as firing rate, cell position, and spike shape, as well as to increase confidence in spike sorting of cells that span these time gaps.

4.6.3 Templates & positions

Even when the waveforms of spikes from different cells were indistinguishable on their primary channel, high-density polytrodes allowed cells to be distinguished on neighbouring channels (Figure 4.3). Arguably, the greater the number of electrode sites and the higher their density, the higher the probability that at least some channels will separate cells that would otherwise be indistinguishable with a single wire electrode or tetrode.

At a broad scale, the 2D spatial localization of cells showed a reasonably uniform spatial distribution for each of the 3 tracks (Figure 4.4A). Given that the greatest amount of damage should be expected in superficial layers, past which the greatest fraction of polytrode length must slide to access deeper layers, there was no obvious lack of superficial cells. This suggests that there was no gross anatomical damage from polytrode insertion (Blanche et al., 2005).

However, at a finer scale, estimated cell positions were biased toward electrode sites (Figure 4.4A & B), and cells with small closed extracellular fields were more biased toward electrode site positions than cells with large open fields (Figure 4.4C). This could be due to spatial undersampling of the neuronal population, with electrode sites spaced further apart than optimal. 65 μm hexagonal electrode site spacing may be insufficiently dense to fully capture the local neural population. This is consistent with Du et al. (2011) who reported that about half of all cells recorded using their custom polytrodes ($\sim$ 40 μm site spacing) were detectable no further than 60 μm away from their apparent origin. Although most of the sorted cells here had templates spanning several channels, those may have been prevalent only because at 65 μm spacing, cells with small closed fields were more likely to fall undetected between electrode sites, and were therefore underrepresented.

One way to further test this hypothesis would be to divide all the cells into two or more groups according to spatial extent (σ, Section 3.5). The distribution of cell distances to the nearest site (Figure 4.4B) could be plotted separately for each group. If the hypothesis is correct, the distribution of nearest site distances should fall off more quickly for cells with small fields, and more slowly (or not at all) for cells with large fields. Another way to test this hypothesis would be to sort spikes from other existing data from polytrodes with site spacing less than 65 μm, such as 2b polytrodes (50 μm hexagonal spacing) or 1b polytrodes (collinear, with 43 μm horizontal and 50 μm vertical spacing, from older tracks omitted from Table 2.2). The hypothesis would predict that the distribution of cells to the nearest site should be closer to analytical and numerical simulations (Figure 4.4B) for polytrodes with lower site spacing.

Resolving this issue may require polytrodes with greater site density. At least one suitable single-shank high-density design is already commercially available, with $<$ 38 μm hexagonal spacing of two columns of electrode sites, and a narrower shank width ranging 30–115 μm (A1x64-Poly2-6mm-

23s-160, Anton Sirota lab, manufactured by NeuroNexus, Ann Arbor, MI). Even at this higher site density, with 64 electrode sites the single shank still spans ~ 1.5 mm, roughly the thickness of cat primary visual cortex.

Another cause for this spatial bias could be the spatial localization method itself (Section 3.5). As an alternative to taking the average localized position of each member spike, spatial localization could be performed on the neuronal template, which is smoother and might therefore allow for a better spatial Gaussian fit. However, templates may be subject to distortion due to drift (see below), whereas individual spikes are not. Such distortion could introduce errors in template-based spatial localization. Another possibility is that there may be insufficient unique x values (columns) to properly constrain localization along the x axis. This might require a polytrode with more than only 2 or 3 columns. A computational workaround to this physical limitation could be rotating the coordinate system during localization so that each site has a unique x and y value in the new rotated space. This could trade off the effective number of rows for an increased effective number of columns.

Blanche (2005) reported that neuron distance from the polytrode (z position) could be inferred, in addition to x and y position along the surface of the polytrode, thereby allowing for 3D localization within the tissue instead of only 2D. This was done by modelling the spatial distribution of each neuron's multichannel peak spike amplitude as a mixed monopole/dipole current source, with voltage decay in extracellular space. The extracellular space was treated as a medium with uniform conductivity in the cortical laminar plane, but differing conductivity orthogonal to the laminar plane. Parameter convergence of this model was reported for smooth mean templates, but was not possible for individual spike waveforms due to excess noise. This meant that 3D localization of spikes from drifting neurons was not possible, although it would be possible to generate multiple templates per neuron as a function of time, and then fit the 3D model to each such template. The low number of electrode site columns in current polytrodes (2 or 3, though see above) also made model parameter estimation especially difficult. Although promising, no attempt to use that model or otherwise perform 3D localization was made here.

The spike sorting method described in Chapter 3 (see especially Figure 3.12) tracked vertical changes in cell position ("drift") up to a net total of $\sim 150\ \mu$m (Figure 4.6). Continuous changes in localized spike position were accompanied by continuous changes in spike shape, such as V_{pp} and σ, but also other spike shape features captured by PCA and ICA. Drift rates varied greatly, ranging 0–200 μm/h (Figure 4.7) depending on time, cell, track, and method of calculation. Despite the appearance of predominantly upward drift in the vertical position plots of cell populations (Figure 4.6, upper panels), there was nearly as much downward drift as upward, with a mean upward drift rate across cells and tracks of $\sim 2\ \mu$m/h. There was less drift in ptc15.tr7c than in the other tracks, though it remains unknown why. One possibility is that longer polytrodes (2a vs. 1a designs) exhibit less drift. Another is that transcolumnar tracks exhibit less drift. Track ptc15.tr7c also had more consistent drift across cells and time, suggesting that lower drift rate might correspond to lower variance in drift rate.

Why some cells drifted much more than others is a mystery, and cannot be easily explained by systematic polytrode-tissue movement. Even more perplexing is that cells that drifted at high rates often did so in unison (Figure 4.6, top right panels), even when at opposite ends of the polytrode, while others in between drifted little, not at all, or in the opposite direction. Perhaps some cells adhered to the polytrode, thereby appearing stationary, while others did not and therefore appeared to drift. Another explanation may be the use of 2D instead of 3D localization. With 2D localization, cells closer to the polytrode might appear to drift more than more distant cells. Although this does not explain why some some cells didn't drift at all, or drifted in opposite directions, using a 3D instead of a 2D localization model might help explain some of the strange drift observed here.

Buzsáki (2004) argued that the apparent location of a neuron may be confounded by different possible spike generating mechanisms outside of the usual somatic or axon hillock locations. Perhaps some of what appear to be conventional neuronal spikes are instead spikes from the axons (Section 5.6.1) or dendrites of more distant neurons, whose apparent origination point may be more greatly disturbed by the polytrode, and which may move in different ways from conventional somatic or axon hillock spikes.

Remarkably, a recent study by Xie et al. (2013) showed that the extracellular space in mouse brain during both sleep and anesthesia increases by 60% compared to the awake state. This large change in volume may affect relative neuronal distances and extracellular spike propagation, and could plausibly fluctuate with depth of anesthesia. Extracellular volume fluctuations could be a significant source of the apparent drift found here. If so, long-duration recordings in awake animals might paradoxically exhibit less drift than in anesthetized animals, whose anesthetic depth and cortical state can vary over time (Sections 6.4 & 6.7.3).

After several hours of recording, sometimes the agar surrounding the craniotomy and the polytrode (Section 2.1) would begin to dry up. This may have led to changes in the polytrode-tissue interface, and potentially contributed to some of the drift described here. Future experiments may add a layer of silicone oil on top of the agar to prevent drying.

To address potential template distortion due to drift, future work could calculate templates over shorter periods of time, perhaps every 30 min of recording, or every certain number of spikes. This would also remove the influence of drift on conclusions about whether the discretized neuronal positions are due to the localization method, or due to polytrode geometry.

A strobe stimulus such as a bright full screen flash generates transient synchronized input into cortical layer 4. By calculating the second spatial derivative of the LFP time-locked to this strobe stimulus, a characteristic laminar pattern of current sources and sinks can be revealed. This pattern is called the CSD, and may be used to determine the position of a vertical array of electrode sites with respect to cortical layers (Mitzdorf, 1985; Sakata and Harris, 2009). The CSD therefore allows estimated neuron positions to be approximately localized by cortical layer. After polytrode insertion, the CSD pattern was used here only as a rough indicator of how deep the polytrode was positioned relative to the cortical layers, and therefore how much further to drive it down for optimal coverage of all layers. CSD analysis was not used to accurately determine the laminar position of each localized neuron. Future work should do so, and may not even require a strobe stimulus (Section 6.7.4), which might therefore allow continuous tracking of polytrode-tissue drift over time.

4.6.4 Orientation tuning

Orientation tuning curves were calculated for all isolated cells from responses to orientation stimuli. 61% of all cells, including those that were inactive (firing rates < 0.05 Hz) during all orientation stimuli, were significantly tuned, while 87% of cells that were active during at least one orientation stimulus were significantly tuned (Section 4.5). Most extracellular electrophysiology studies report a high fraction (70–90%) of orientation tuned cells in primary visual cortex. Schiller et al. (1976b) found that 87% (572/654) of cells in rhesus monkey V1 were well tuned, with the disclaimer that the reported proportion of tuned cells depends on recording methods. Hammond and Andrews (1978) found 70% (68/97) of cells in cat areas 17 and 18 were tuned to orientation. De Valois et al. (1982) and Ringach et al. (2002b) reported that 86% (190/222) and $\sim$ 89% ($\sim$ 274/308) respectively of cells in macaque V1 were well tuned (bandwidth $< 90°$).

All of these reports likely excluded very low firing rate cells, cells that responded poorly to artificial stimuli, and/or cells whose activity may have only been detectable over long recording periods. The proportion of tuned cells in these studies may therefore be most directly comparable to the 87% of active cells reported here. In contrast, an extensive *in-vivo* two-photon calcium imaging (Section 1.1) study in isoflurane-anesthetized young (P19–40) cat primary visual cortex (area 18) by Ohki et al. (2005) found 61% (97% of the 63% that were responsive) of 6734 cells in layers 2/3 were significantly tuned to orientation. This corresponds well to the 61% of all isolated cells reported here.

Damage from the wide shank polytrodes used here may be responsible for some of the 39% of all isolated cells that were not significantly orientation tuned to even a single orientation stimulus. Switching to narrower shank polytrodes might increase the fraction of significantly tuned cells. Besides damage, there may be other reasons for non-oriented cells. Perhaps some cells respond only to stimuli more naturalistic stimuli than bars or grating. Although technically challenging, perhaps their orientation tuning could be characterized from responses to naturalistic movies instead. Monocular stimulation instead of binocular stimulation may be insufficient to adequately stimulate some cells. The shift work hypothesis (Section 4.6.2) or changes in cortical state (Chapter 6) might explain the inactivity of cells at certain times, which for some cells might by chance have been during the relatively brief periods of presentation of artificial orientation stimuli. And finally, there may very well be a wide range of degrees of orientation selectivity in visual cortex (De Valois et al., 1982), with some cells lacking significant orientation tuning altogether.

Cells with very sparse firing were still highly selective for stimulus orientation (Figures 4.8C &

4.9). Though often discarded in studies, such low firing rate cells (< 0.1 Hz) may nevertheless be useful for perception. Surprisingly, orientation tuning strength was inversely and linearly correlated with geometric mean firing rate (Figure 4.9A). The author is unaware of any other such report. This contradicts the notion that higher firing rates generally lead to better stimulus encoding (Gershon et al., 1998), and bolsters the importance of low firing rate cells.

Across the population of tuned cells from all 3 tracks, there were no obvious relationships between orientation preference, tuning strength, and normalized cell depth (Figure 4.10A), nor was there any reason to expect otherwise. However, individual tracks exhibited varying degrees of dependence between cell orientation preference and depth, likely due to their vertical angle of insertion (Figure 4.10B). Track ptc15.tr7c had orientation preferences that varied strongly with depth, suggesting that it was a highly transcolumnar track. Direction of motion preferences were not examined here, but future work should show that, like orientation preference, over multiple tracks there was no bias towards any particular preferred direction of motion.

5 Cell Type

5.1 Introduction

Functional variation among neurons may be caused by a variety of factors, including differences in cellular morphology, physiology, and connectivity. Some of the functional variation in extracellularly recorded populations may be explainable by classifying neurons into different types. There are several possible ways to do so, five of which are described here, and the first three of which are used.

One method is to classify neurons according to their temporal waveform shapes, of which there are at least two types: longer waveform duration excitatory pyramidal cells and shorter waveform duration inhibitory interneurons (Csicsvari et al., 1998; Barthó et al., 2004; Blanche, 2005; Luczak et al., 2007; Sirota et al., 2008; Niell and Stryker, 2008; Mizuseki et al., 2009; Sakata and Harris, 2009; Benchenane et al., 2010; Wilson et al., 1994; Gur et al., 1999; Mitchell et al., 2007). Both types are usually labelled as "putative" due to the inability of most extracellular experimental setups (including that used here) to directly visualize and measure the intracellular currents of each extracellularly recorded neuron.

Another way to classify cells may be according to the spatial extent of their waveforms, as measured across multiple channels. Like temporal waveform shape, the spatial extent of the extracellular field might cluster into two groups: pyramidal cells with large open dipolar extracellular fields, and interneurons with small closed monopolar fields (Blanche, 2005). However, there is currently no evidence in the literature for such a bimodality in spatial extent.

Third, in primary visual cortex, neurons can be classified into simple and complex cells according to the spatial separation of the oriented ON and OFF subfields of their RFs (Hubel and Wiesel, 1962). Simple cells have well separated ON and OFF subfields, while complex cells have overlapping subfields (Section 1.1). A third possibility are afferent axons from the lateral geniculate nucleus (LGN), which like simple cells have non-overlapping ON and OFF subfields, but with a centre-surround spatial organization which makes them untuned to orientation. Classifying cells by RF type is a standard procedure (Schiller et al., 1976a; Gilbert, 1977; DeAngelis et al., 1993; Kagan et al., 2002), but the ease with which it can be done depends on the stimulus used. Classification is straightforward with white and black drifting bars or gratings, both of which have large, oriented, spatially separated light and dark regions for independently stimulating potentially separated subfields. Reverse correlation to an m-sequence white noise movie (Section 2.4) to calculate the spike-triggered average (STA) can also be used to extract non-overlapping subfields, and to therefore check if the RF is simple (Ringach and Shapley, 2004). However, the lack of a discernable RF from the m-sequence does not necessarily mean the cell is complex, and may merely mean that the cell was not sufficiently stimulated by the white noise stimulus. Spike-triggered covariance (STC) of m-sequence responses can be used to calculate overlapping subfields, and hence detect complex cells (Schwartz et al., 2006), but calculating the requisite covariance matrix for the STC requires many more spikes than the STA, and was considered impractical here. In comparison, naturalistic stimuli are desirable for evoking greater and more realistic responses, but mapping RFs (and especially complex cell RFs) with naturalistic movies is more complicated than reverse correlation to a linear stimulus with only first order statistics, such as the m-sequence (Ringach et al., 2002a; Smyth et al., 2003; Willmore and Smyth, 2003; Sharpee et al., 2004; Touryan et al., 2005). Also, due to the high spatial correlations in naturalistic movies, very long duration recordings and/or high firing rates may be required to extract reasonable estimates of RFs from naturalistic movie responses. However, the multi-hour recordings described here span both artificial and naturalistic stimuli, with single unit identification across these different stimuli, potentially allowing for the best of both worlds. Note that the true discreteness of simple and complex cell classes has been the subject of debate (Dean and Tolhurst, 1983; Chance et al., 1999; Mechler and Ringach, 2002; Abbott and Chance, 2002; Kagan et al., 2002; Priebe et al., 2004; Mata and Ringach, 2005).

Fourth, neurons may also be classified by their firing patterns as regular spiking (RS), fast spiking (FS), chattering (CH) or intrinsically bursting (IB) (Gray and McCormick, 1996; Nowak et al., 2003; Barthó et al., 2004; Blanche, 2005; Herikstad et al., 2011). This is done by examining each cell's autocorrelogram or inter-spike interval (ISI) histogram. However, cells with very low firing rates, such as those reported here (Section 4.3), have mostly empty autocorrelograms and ISI histograms at the timescales of interest (ISIs < 50 ms), precluding such classification. Note that Herikstad et al. (2011) showed that it may nevertheless be possible to classify the firing patterns of very low rate cells by plotting their ISI histograms on a log timescale instead of a linear one.

Fifth, neurons may be explicitly classified as excitatory or inhibitory by examining their cross-correlograms for evidence of monosynaptic connections. Given enough spikes, excitatory monosynaptic connections result in peaks in the cross-correlogram, and inhibitory connections result in troughs (Perkel et al., 1967). However, monosynaptic connections are very rare, even within a local population. Barthó et al. (2004) found that only 0.2% of local cell pairs recorded in layer 5 of rat somatosenory and prefrontal cortex had cross-correlogram peaks or troughs that suggested monosynaptic connectivity. Furthermore, less than 8% of neurons induced such peaks or troughs in the cross-correlograms of other neurons (Barthó et al. (2004), Table 1, $(72 + 5 + 28)/1414 = 7.4\%$) and could therefore be classified as excitatory or inhibitory. Qualitative examination of cross-correlograms in the data presented here showed similarly low percentages. Therefore, only a small fraction of recorded neurons can be classified in this manner, making it a seemingly impractical method for classifying an entire population. Fortunately, the same study showed that spike waveform duration correlated with whether a cell was excitatory (long duration) or inhibitory (short duration), allowing waveform duration to be used as a reasonable proxy for the excitatory/inhibitory classification.

In this study, cells were classified into types according to their spike shape (Section 5.2) and RFs (Section 5.4). The width of the secondary spike peak, and the amplitude asymmetry of the two peaks were used to cluster cells into 4 spike types: fast, slow, fast asymmetric and slow asymmetric. Surprisingly, the width of the primary spike peak was not useful for cell typing. The spatial extent of spikes was also unimodal and not useful for cell typing (Section 5.3), although this may have been an artifact of using 2D instead of 3D spatial localization. A variety of methods were used to characterize RFs, and their consensus was used to classify each cell as simple, complex, putative LGN afferent, or unknown RF type. There were roughly equal numbers of cells classified as simple, complex, and unknown RF type, with only 5% classified as putative LGN afferents. Cells classified as unknown RF type may have been mostly complex cells damaged by polytrode insertion.

5.2 Spike shape

Six different measures were used to characterize the shape of the mean waveform of the primary channel (template) of each neuron: maximum slope, slope-based spike duration, interpeak interval, peak temporal asymmetry, full width half maximum (FWHM) of primary and secondary peaks, and peak amplitude asymmetry. To prevent quantization of temporal measures, templates were spline interpolated to 1 μs resolution.

The first measure, maximum slope, was simply the maximum absolute value slope of each cell's template (in μV/μs). The second measure, slope-based spike duration, was defined as the interval between the first and last timepoint in the template which exceeded an absolute value slope threshold of 0.4 μV/μs. All cells had a maximum absolute value slope of at least 0.49 (Figure 5.1A), so picking a threshold of 0.4 did not exclude any cells. Rather, this threshold value seemed to correspond to qualitative judgements of when each template first departed significantly from zero, and when it last returned to zero.

The third measure, interpeak interval, was defined as the time between the primary and secondary peaks in the template. Peaks were defined as the extrema between template zero-crossings and edges. The primary peak was considered to be the extremum closest in time to the $t = 400$ ms general alignment point used during spike detection (Section 3.4). The secondary peak was considered to be the extremum immediately following the primary peak, and by definition, of opposite sign. If however the primary peak was the last extremum in the template, it was labelled the secondary peak, and the one to its immediate left became the primary peak. Only 2 cells out of the 245 required such a swap of peak labels.

The fourth measure, peak temporal asymmetry, was an estimate of the temporal skew of each peak around its mode. It was calculated by finding the time between the mode and the median of each peak, normalized by the width of the peak. Peak limits were determined using the fifth measure, FWHM, i.e., the width of the peak at one half its maximum amplitude with respect to 0 volts. FWHM of the primary and secondary peak was designated as FWHM1 and FWHM2, respectively.

The sixth and final measure was peak amplitude asymmetry, inspired by the asymmetry index of Sirota et al. (2008) and Sakata and Harris (2009), and by the slope ratio of the depolarization and repolarization phases of intracellular spikes (McCormick et al., 1985). Sirota et al. (2008) and Sakata and Harris (2009) compared the amplitudes of the first and third peaks of tri-phasic waveforms. Given that most of the waveforms collected here were only biphasic (Figure 4.3), the index was modified to compare the primary and secondary peaks instead. Peak amplitude asymmetry was thus defined as $(V_1 - V_2)/(V_1 + V_2)$ where V_1 and V_2 are the absolute values of amplitudes of the primary and secondary peaks, respectively. Its value ranged from -1 to 1, with 0 representing equal peak amplitudes.

For the purposes of measuring spike shape, template amplitudes were normalized because amplitude as measured at a given electrode site is mostly not an inherent property of the cell. Rather, it is mostly a function of distance between cell and electrode site (Henze et al., 2000).

Distributions of the first three measures are shown in Figure 5.1A–C. Maximum slope was unimodal, while slope-based duration showed a clear bimodality and interpeak interval showed a hint of multimodality. Two clear clusters emerged when slope-based duration was plotted against interpeak interval (Figure 5.1D). Points were divided by a manually placed line separating fast (red) and slow (blue) spike types. Overplotted templates of the two clusters (Figure 5.1E–F) looked convincingly different. However, the bimodality of the slope-based duration metric was merely an artifact of thresholding template slope. For a given slope threshold (0.4 μV/μs in this case), either only one peak or both peaks of a given template would exceed threshold, resulting in spike duration values that were artifactually bimodal. Changing the slope threshold made points jump from one cluster to another, and cluster sizes were very sensitive to the precise threshold used. Slope-based duration was therefore abandoned as a measure of spike shape, but is left here as a demonstration of the dangers of threshold-based metrics. Thresholding can be dangerous because it imposes a nonlinearity, something best avoided before clustering.

The fourth measure, peak temporal asymmetry, was unimodal for both primary and secondary peaks, and its use did not improve clusterability (not shown).

Ultimately, only the fifth and sixth measures, FWHM and amplitude asymmetry, were found to be useful in classifying cells according to spike shape. Their distributions are shown in Figure 5.2A–C. FWHM1 was unimodal, but FWHM2 was bimodal, while amplitude asymmetry was somewhat multimodal. Therefore, FWHM2 vs. amplitude asymmetry was chosen as the space of spike shape measures in which to cluster cells into different spike types (Figure 5.2D). This decision was further supported by examining 2D and 3D scatter plots of all combinations of the spike shape measures described here (excluding slope-based duration). Of all the possible spaces, the FWHM2 vs. amplitude asymmetry space revealed the clearest and greatest number of clusters. For example, interpeak interval was highly correlated with FWHM2, and therefore provided no additional clustering benefit. A second measure of spike duration based on FWHM (measured from the start of

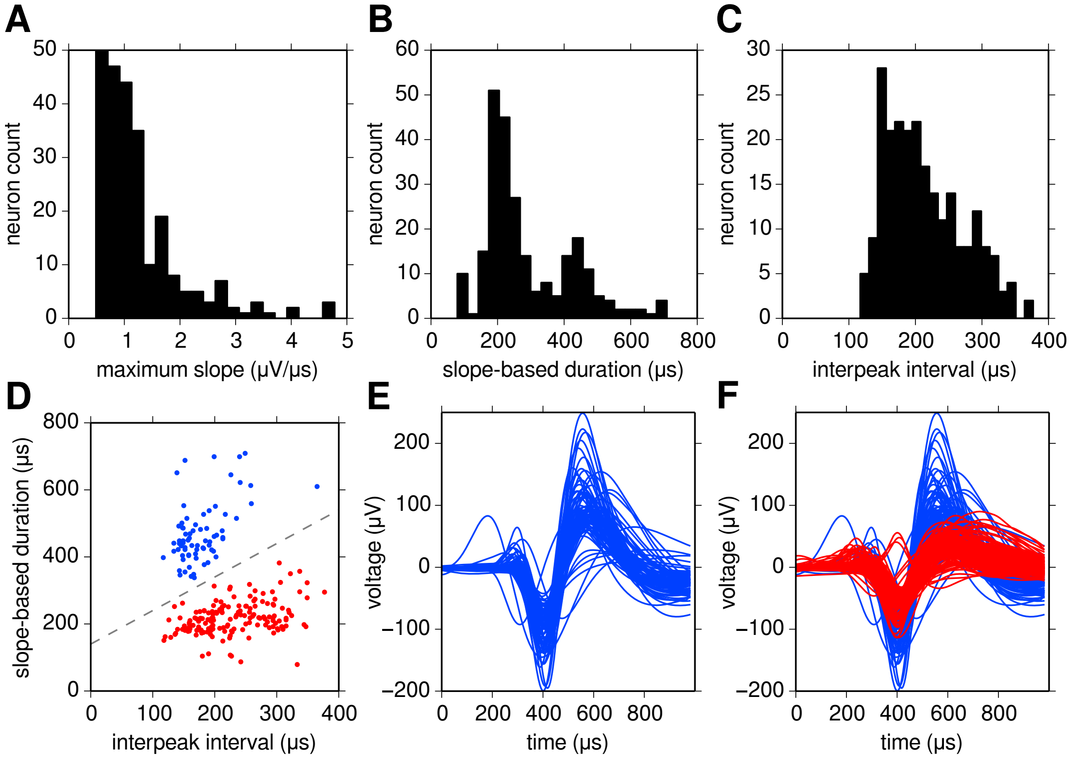

Figure 5.1: A demonstration of the potential dangers of threshold-based metrics. ***A–C***: Distributions of three different temporal waveform measures for all cells (see text for details). Of the three, only slope-based duration (the interval between the first and last timepoints which exceeded an absolute value slope threshold of 0.4 μV/μs) showed clear bimodality. ***D***: Scatter plot of spike duration vs. inter-peak interval showing two well-separated clusters, divided manually via the dashed line into fast (red) and slow (blue) spike types. The clusters were an artifact of thresholding template slope to calculate spike duration. ***E–F***: Overplotted templates, coloured according to the artifactual clustering in ***D***. Slow templates are shown on their own in ***E***, and fast templates are plotted over top of slow in ***F***. Despite being artifactual clusters, their differences in template shape looked convincing.

the primary peak to the end of the secondary peak) was also highly correlated with FWHM2, and provided no extra clustering information. Finally, full width at different fractions of peak maximum (from 0.125 to 0.9) were also tested, but were all less clusterable than FWHM.

Points in FWHM2 vs amplitude asymmetry cluster space were split into 4 clusters using manually placed straight lines (Figure 5.2D). No points were excluded. Clusters were labelled fast (red), slow (blue), fast asymmetric (green), and slow asymmetric (grey). The first 2 clusters were very clear, the third somewhat less so, and the fourth even less. Overplotted unnormalized templates are shown separately for each cluster (Figure 5.2E–H), and simultaneously (Figure 5.2I).

Cell counts according to spike type and track are shown in Table 5.1. Across all tracks, 52% (128) of cells were fast, 30% (73) were slow, 13% (32) were fast asymmetric, and 5% (12) were

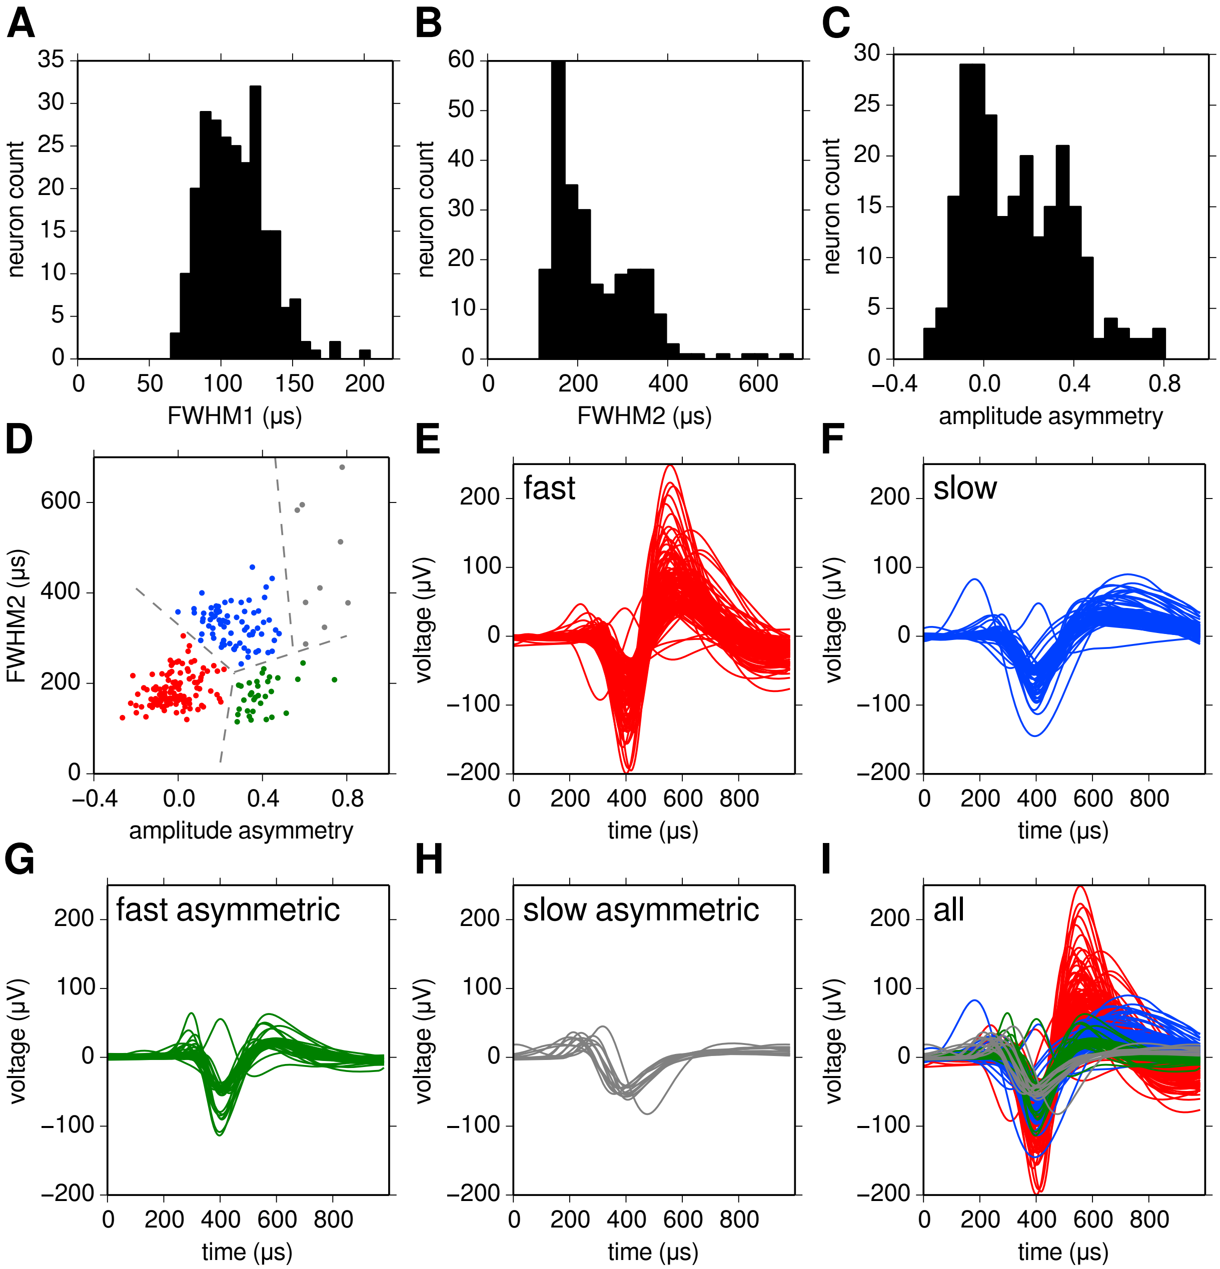

Figure 5.2: Temporal spike shape measures of primary channel templates were used to classify neurons into two clear clusters (fast and slow), plus another two potential clusters (fast asymmetric and slow asymmetric) that were less clear. ***A*–*C***: Distributions of three different temporal waveform measures (see text for details). FWHM1 showed no mulitmodality (***A***), but FWHM2 and amplitude asymmetry did (***B*–*C***). ***D***: Scatter plot of FWHM2 vs. amplitude asymmetry. Points were manually clustered according to the dashed lines. To maximize cluster visibility, the y axis limit was set such that 3 grey points were excluded. ***E*–*H***: Overplotted templates of each cluster, coloured according to ***D***. Although templates were roughly aligned at 400 μs, they were not explicitly aligned with one another due to the many realignment operations performed during spike sorting. Overplotted templates of all four clusters are shown in ***I*** for comparison.

track	fast	slow	fast asym	slow asym	simple	complex	LGN aff	unknown
ptc15.tr7c	61(75)	12(15)	7(9)	1(1)	53(65)	12(15)	0(0)	16(20)
ptc22.tr1	32(34)	36(39)	19(20)	6(6)	18(19)	31(33)	9(10)	35(38)
ptc22.tr2	35(49)	25(35)	6(8)	5(7)	14(20)	32(45)	3(4)	22(31)
total	128(52)	73(30)	32(13)	12(5)	85(35)	75(31)	12(5)	73(30)

Table 5.1: Spike and RF type counts (with percentages in parentheses), for each track and in total. Compared to the other two tracks, ptc15.tr7c had an excess of fast and simple cells, and a lack of slow, slow asymmetric, complex, and putative LGN afferent cells. **asym**: asymmetric. **aff**: afferent.

slow asymmetric. Compared to the other two tracks, ptc15.tr7c was an outlier, with an excess of fast cells, and fewer slow and slow asymmetric cells. Fast cells had the largest peak-to-peak amplitudes, and asymmetric cells had the smallest (Figure 5.2I). The mostly positive values in amplitude asymmetry (Figure 5.2C & D) indicate that the primary peak (the one at ~ 400 μs) was generally greater in amplitude than the secondary peak (the one that followed). The disparity in amplitude between the primary and secondary peaks was greatest for the cell types labelled asymmetric.

5.3 Spatial extent

For each neuron, the median of the spatial extent (σ) of its member spikes (Section 3.5) was taken as the spatial extent of the neuron (analogous to how per-neuron x and y positions were obtained in Section 4.4.) Figure 5.3A shows the distribution of σ for cells in each track, as well as for all cells across all tracks. The distribution across all tracks peaked at 50 μm and showed no multimodality. Save for perhaps a hint of bimodality in ptc15.tr7c, the individual tracks also showed no multimodality. Also, σ did not correlate with any temporal measures (the best two, FWHM2 and amplitude asymmetry are shown in Figure 5.3B), and did not improve clusterability of cells into different types. Spatial extent from 2D spatial localization was therefore not useful for cell typing.

5.4 Receptive field type

Classification of cells into simple, complex and putative LGN afferent RF types was performed using three different types of stimuli: white and black drifting bars at various orientations, drifting gratings at various spatial frequencies and orientations, and m-sequence noise stimuli. Given experimental time constraints and the history of stimulus development associated with each track, only two of the three stimulus types were available for each track. All 3 tracks had m-sequence stimuli, but only ptc15.tr7c had drifting gratings with a sufficient number of presentations of each stimulus condition to allow determination of RF type, and only ptc22.tr1 and ptc22.tr2 had white and black drifting bar stimuli. Results from all relevant recordings were considered when classifying a cell's

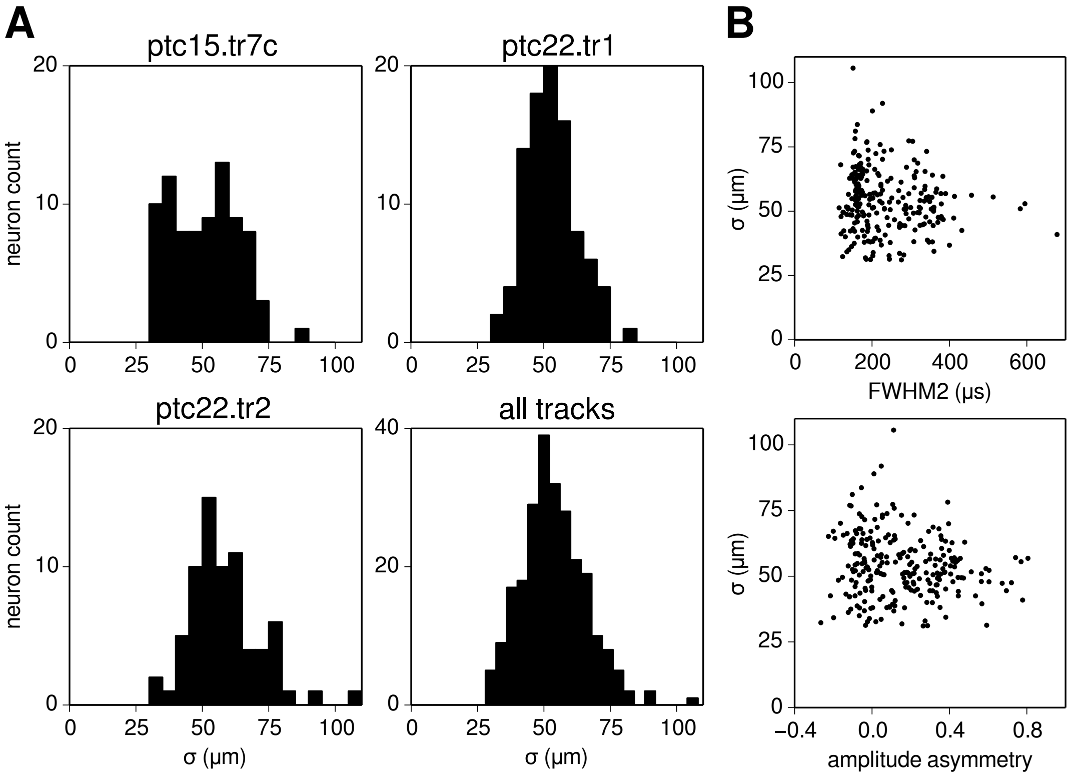

Figure 5.3: ***A***: Distributions of spatial extent (σ) of neurons in individual tracks, as well as across all tracks. Overall, σ had a unimodal distribution. ***B***: σ plotted against the two best temporal waveform measures (FWHM2 and amplitude asymmetry) from Section 5.2, showing that σ had little to no correlation with either of them, and did not improve clusterability.

RF type.

Drifting bar raster plots (Figure 5.4A) were visually inspected for spatial overlap of responses to white and black bars at the preferred orientation. If white and black bar responses were obviously non-overlapping, the cell was classified as simple for that recording. Classifying it as simple did not rule out later classification as a putative LGN afferent based on the characteristics of its STA (see below). Poor orientation tuning in conjunction with non-overlapping ON and OFF responses (Figure 4.8B) was suggestive of an LGN afferent, but was not considered definitive. If the responses were obviously overlapping, the cell was classified as complex for that recording. If the degree of ON and OFF subfield overlap was too difficult to discern by visual inspection of the raster plot, the RF type was classified as "unknown" for that recording.

Drifting grating raster plots (Figure 5.4B) were also visually inspected, this time for modulation at the temporal frequency of the grating (2 Hz), at the preferred orientation and spatial frequency. If there was obvious temporal modulation above the baseline firing rate at the expected temporal frequency, the cell was considered simple. If the rate was elevated (i.e., the cell was tuned) but showed little to no temporal modulation, the cell was considered complex. Otherwise, for that

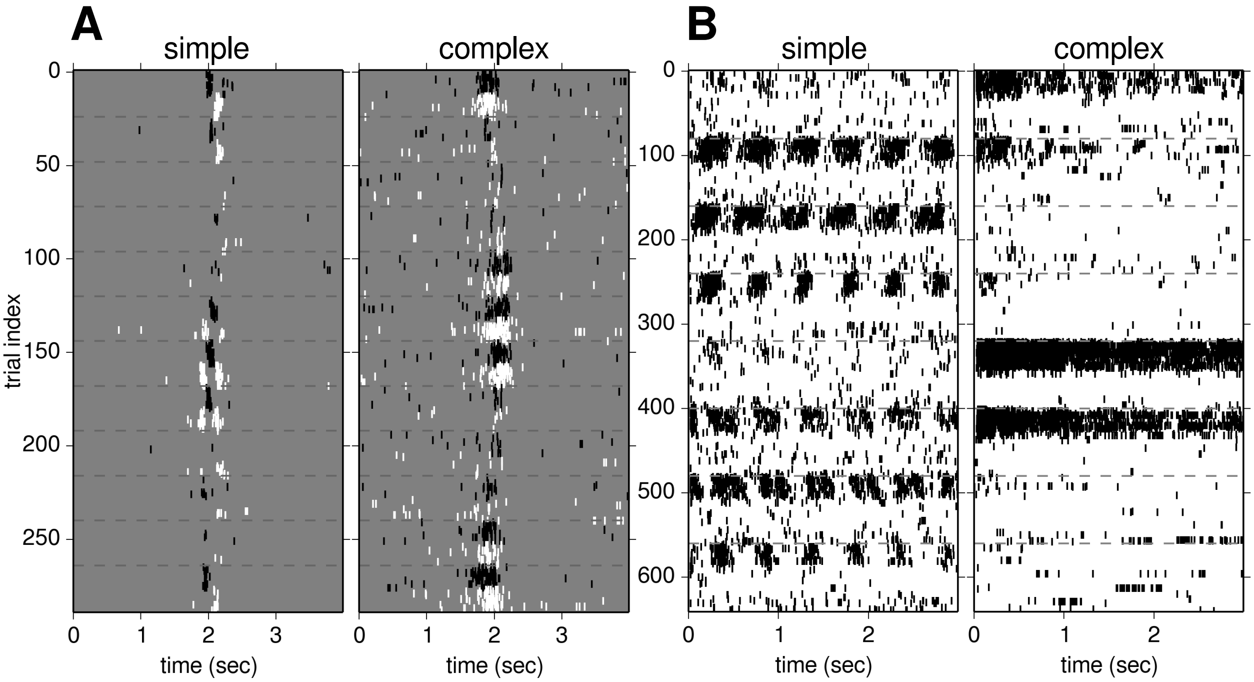

Figure 5.4: Example trial raster plots of drifting bar (***A***) and grating (***B***) responses of simple and complex cells. For all panels, visible rows of responses (8–12 trials each) are in order of stimulus parameter combination (not temporal order). Horizontal dashed lines denote orientation increments. ***A***: Drifting bar raster plots. Responses to white and black drifting bars are represented by white and black ticks respectively, where each tick represents one spike. Stimulus parameters were bar brightness (white or black) and orientation (12 steps of 30°). All possible parameter combinations were presented in pseudorandom order, with 12 presentations per combination. Trial index increased first by presentation number, then bar brightness, and then orientation. Spatial overlap of responses to white and black drifting bars distinguished simple cells from complex. Variation in response timing over the range of orientations was due to slight misalignment of the center of the stimulus and the RF centers of some cells. ***B***: Drifting grating raster plots. Stimulus parameters were spatial frequency (10 values from 0.15–5 cycles/°) and orientation (8 steps of 45°). All possible parameter combinations were presented in pseudorandom order, with 8 presentations per combination. Trial index increased first by presentation number, then spatial frequency, and then orientation. Response modulation at the 2 Hz temporal frequency of the grating relative to baseline firing was obvious for the simple cell, but absent for the complex cell. The raster plots of all four cells showed strong orientation tuning.

recording, the cell was classified as "unknown".

Reverse correlation of spikes to the m-sequence movie to construct the STA is illustrated in Figure 5.5. For a fixed reverse correlation time range of say, 40 ms, each spike's contribution to the STA is found by looking backwards 40 ms from the time of each spike to determine which frame of the movie was on display. That frame's pixels are added to the running total. Once all spikes are reverse correlated in such a manner, the sum of the frames is normalized by the number of spikes to get the STA, representing the average movie frame that triggered a spike. This

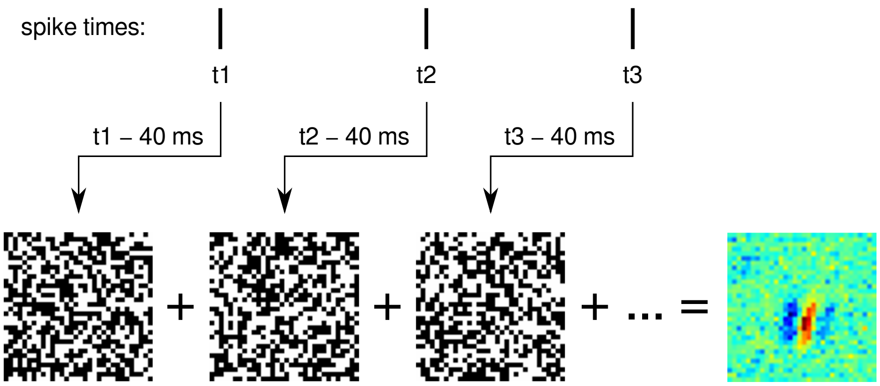

Figure 5.5: An illustration of the spike triggered average, calculated by reverse correlating the spike train of a cell in response to an m-sequence noise movie to estimate the cell's RF 0 to 40 ms following the presentation of each m-sequence frame. A simple cell is illustrated here. Red are ON subregions (stimulated by light pixels, inhibited by dark pixels), blue are OFF subregions (stimulated by dark pixels, inhibited by light pixels).

can be repeated for different time ranges (each typically a multiple of the frame display time) to calculate the spatiotemporal STA, estimating the cell's spatially non-overlapping spatiotemporal receptive field (STRF) (Figure 5.6). Cells with a non-zero (non uniform) STA were considered to be simple cells or putative LGN afferents, with the STA approximating the cell's spatially non-overlapping STRF. Oriented STRFs were considered to be simple cells, while small, unoriented, quickly emerging STRFs were considered to be putative LGN afferents. An absence of an STA at all reverse correlation time ranges suggests that the cell is complex, but due to spatially overlapping ON and OFF subfields, the STA cannot reveal the cell's STRF. A different method, the STC, may be used to reveal spatially overlapping subfields. However, the STC requires roughly an order of magnitude more spikes than the STA to accurately calculate the required covariance matrix, and was not used here to detect complex cells.

Note that reverse correlating the spikes to the movie is equivalent to forward correlating the movie to the spikes. The latter was used here because if was more computationally efficient. Also note that reverse correlation timepoints (e.g., 40 ms, 80 ms, 120 ms, etc.) are referred to here as time ranges (0–40 ms, 40–80 ms, 80–120 ms, etc., see time labels in Figure 5.6), because when reverse correlating from a given spike to a movie frame *t* ms ago, that frame will have been on-screen for anywhere from 0 to *dt* ms, where *dt* is the duration each frame is displayed. Similarly, when forward correlating from a movie frame, a time range of *t* ms after the start of the frame to *t* ms after the end of the frame is used to determine which spikes should accumulate that frame's contents.

Cells with non-zero STAs are shown in Figure 5.6, organized vertically according to depth along the length of the polytrode. Non-zero STAs are shown for 63 cells, from one representative m-sequence recording per track. For all m-sequence recordings from all 3 tracks, there were 93 cells with non-zero STAs. In agreement with DeAngelis et al. (1993), most responses lasted over 100 ms, and RF subfields generally reversed polarity 40–80 ms after their initial appearance, itself 20–40 ms after stimulus onset. RFs appeared normal, most having one or more elongated subfields of alternating sign. There were few obvious holes in RF subfields that might indicate extensive damage to afferents. RFs calculated from multiple m-sequence recordings hours apart were either quite similar (e.g., Figure A.1), or were missing altogether, possibly due to either a gain or loss of responsivity. The first track, ptc15.tr7c, had many more cells with non-zero STAs than did the other two tracks, and their subfield orientation varied systematically as a function of depth along the polytrode, suggesting a transcolumnar insertion. Though difficult to see (due to a conservatively large choice of m-sequence frame size), the subfield orientations of the other two tracks were more constant as a function of depth. Results from the calculation of orientation tuning from oriented stimuli (Section 4.5) agree with the conclusion that track ptc15.tr7c was strongly transcolumnar.

Subjective classification of a unit as a putative LGN afferent was mostly based on the m-sequence STA, but also on responses to oriented stimuli. Small, circular fields that emerged early (within 40 ms of stimulation) suggested an LGN afferent, as did poor orientation tuning and high firing rates in response to drifting bars (Figure 4.8B) or gratings.

All relevant recordings (4 or 5, depending on the track) were used to classify each cell as simple, complex, putative LGN afferent, or unknown. Usually the two stimulus types (either white and black drifting bars and m-sequence, or drifting gratings and m-sequence) agreed on the classification. Since the responsivity of cells could vary substantially over hours of recording (Section 4.3, Swindale and Spacek (2012)), using as many recordings as possible to classify RF type increased the chance of a successful classification (something other than “unknown”). It also helped increase confidence when responses from multiple recordings of different stimulus types gave the same classification. For 11% (26/245) of cells, the responses from two or more recordings conflicted over whether the cell was simple or complex. In these cases, a voting system was used to resolve the conflict, with one vote per recording. There were no recordings that conflicted over the distinction between simple cells and LGN afferents. Appendix A demonstrates simple cell RF stability over a 14 hour period.

Cell counts and percentages according to RF type and track are shown in Table 5.1. Across all 3 tracks, there were roughly equal numbers of simple, complex, and unknown cells (about one third each), with only 5% of cells classified as putative LGN afferents. Of the 12 units classified as LGN afferents, only 3 were significantly ($p < 0.01$) tuned to orientation (Section 4.5), and only weakly so ($\bar{r} = 0.13$). Of the 93 cells with non-zero m-sequence STAs, 73% (68/93) were classified as simple cells, 12% (11/93) as complex cells, 13% (12/93) as LGN afferents, and 2% (2/93) as unknown. As was the case for spike type, ptc15.tr7c was an outlier in terms of its RF type distribution (Table 5.1), with an excess of simple cells and a lack of complex cells and LGN afferents compared to the other two tracks.

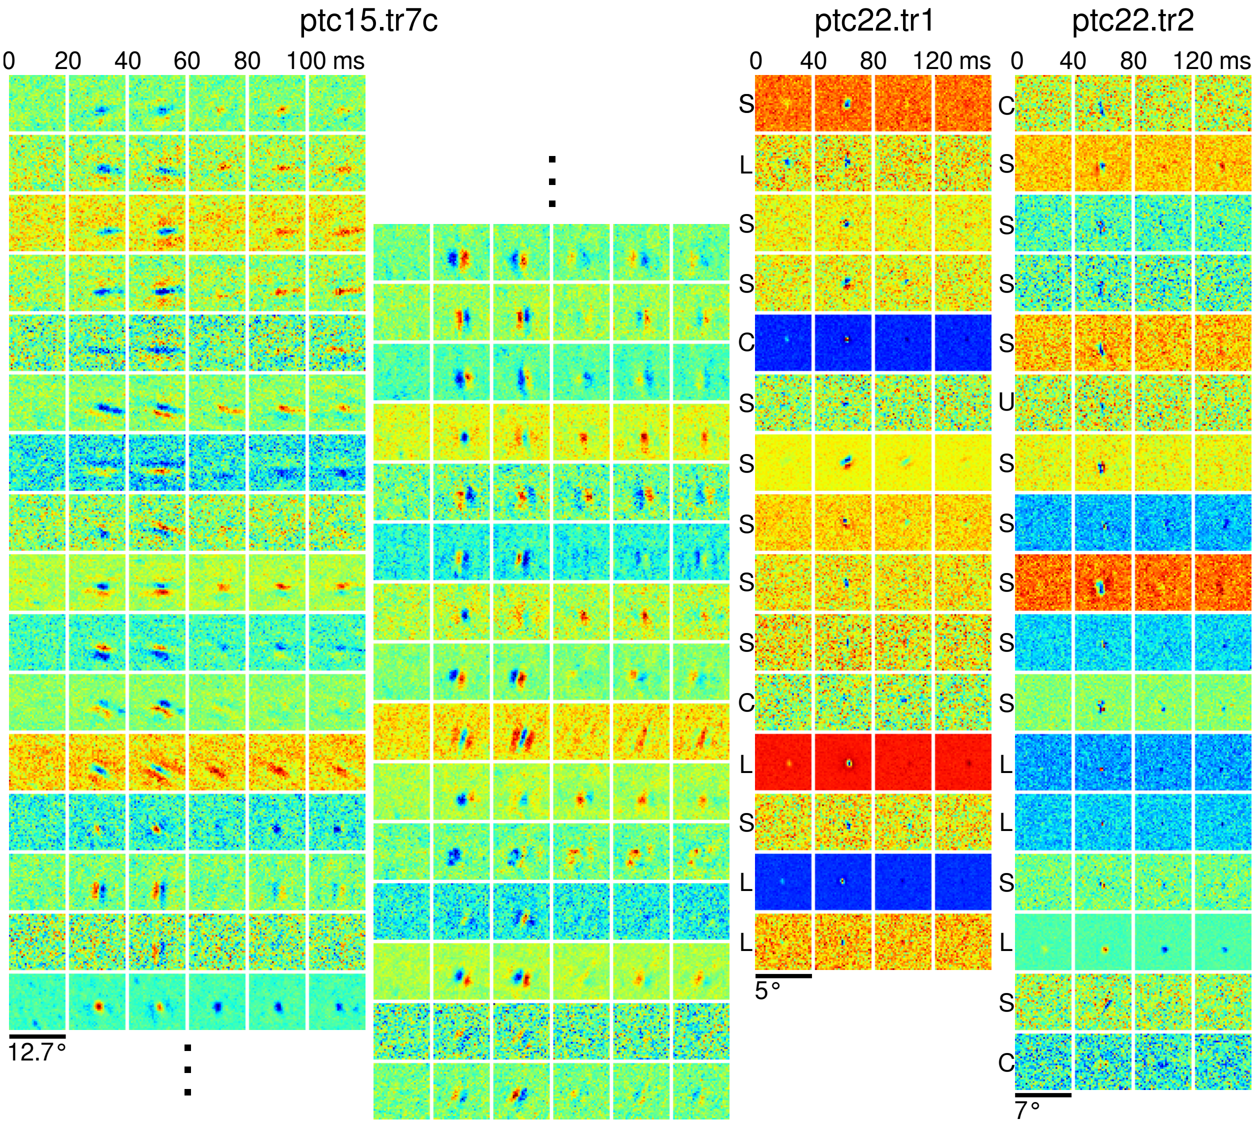

Figure 5.6: Spatiotemporal RFs from STAs to multiple reverse correlation time ranges of an m-sequence noise movie. Each row represents a unit, and each column a reverse correlation time range. Within a row, the same normalization is used for all time ranges, but each row is independently normalized to maximize visibility of RF structure. Units were sorted top to bottom by depth along the polytrode and grouped according to track (ptc15.tr7c split vertically), with one representative m-sequence recording chosen for each track. Only units with non-zero STAs are shown. Other m-sequence recordings in each track revealed other units with non-zero STAs. For ptc15.tr7c (***left***), m-sequence frames were presented for 20 ms each. Orientation tuning changed in a regular clockwise fashion, strongly suggesting the track was transcolumnar. All cells were classified as simple. For ptc22.tr1 and ptc22.tr2 (***middle, right***), far fewer units with non-zero STAs were present. Letters denote RF type: **S**: simple, **C**: complex, **L**: putative LGN afferent, **U**: unknown. Each classification took all relevant recordings into consideration. M-sequence frames were presented for 40 ms each. RFs appear smaller because the absolute movie size in degrees of visual angle was set higher than necessary, conservatively ensuring capture of the full RFs of the entire population of cells.

5.5 Cell type comparison

The spatial distributions of cells by both spike type and RF type are shown in Figure 5.7. For all tracks, slow, complex and unknown cells were fairly evenly distributed along the length of the polytrodes. For track ptc15.tr7c, both fast and simple cells were also evenly distributed across the length of the polytrode, but this result may have been because that track was likely transcolumnar (Sections 4.5 & 5.4). In contrast, for the other two tracks, both fast and simple cells were somewhat more likely to be found higher up the polytrode, presumably in more superficial layers. There were too few fast asymmetric, slow asymmetric, and LGN afferent units to make conclusions about their spatial distributions.

The relationship between cell position and spatial extent was also examined, with points representing cell position (as in Figures 4.4A & 5.7) coloured by each cell's spatial extent (not shown). There was no obvious relationship between cell position and spatial extent for any of the 3 tracks.

To see if there was a relationship between spike type and RF type, various comparisons were made. Figure 5.8A–E shows overplotted templates of cells grouped according to RF type. Examination suggested that simple cells had the highest peak-to-peak amplitude of the four RF types, and were perhaps slightly faster than complex cells (width of secondary peaks at $t \approx 550$ μs of red and blue traces in Figure 5.8E). Putative LGN afferents had the smallest amplitude. Distributions of peak-to-peak amplitudes in Figure 5.8F confirmed that RF type amplitudes from largest to smallest were simple, complex, and LGN afferent. Unknown cells had roughly the same mean amplitude as complex cells, suggesting that unknown cells might nominally be mostly complex cells.

Plotting cells in the spike type classification space (FWHM2 vs amplitude asymmetry) but colouring them according to RF type (Figure 5.8G) did not reveal any obvious relationship, save for perhaps an excess of simple cells (red points) in the fast cell cluster (bottom left). A matrix of cell counts of RF type vs spike type (Figure 5.8H) confirmed the above, showing that simple cells were about 4 times more likely to be fast cells than slow cells, even though overall there were only about 1.75 times as many fast cells as slow cells (Table 5.1, statistical tests for significance were not performed). Furthermore, both complex and unknown cells were classified as fast and slow spike types in roughly equal numbers, suggesting again that unknown cells may nominally be mostly complex cells.

How might cell types differ in their mean track-wide firing rates and fractional activity durations? Both measures were calculated as previously described (Section 4.3), and the distributions for all four spike types and all four RF types are shown in Figure 5.9. Fast cells had somewhat higher mean rates than slow, fast asymmetric, and slow asymmetric cells (Figure 5.9A), with geometric means of 0.17 Hz, 0.078 Hz, 0.053 Hz and 0.054 Hz, respectively. This was despite fast cells having the same mean fractional activity duration as slow cells (Figure 5.9B). Cells with unknown RFs had lower mean firing rates and were active for shorter periods of time than simple, complex, and LGN afferent cells, all of which were very similar in both respects (Figure 5.9C & D). Geomet-

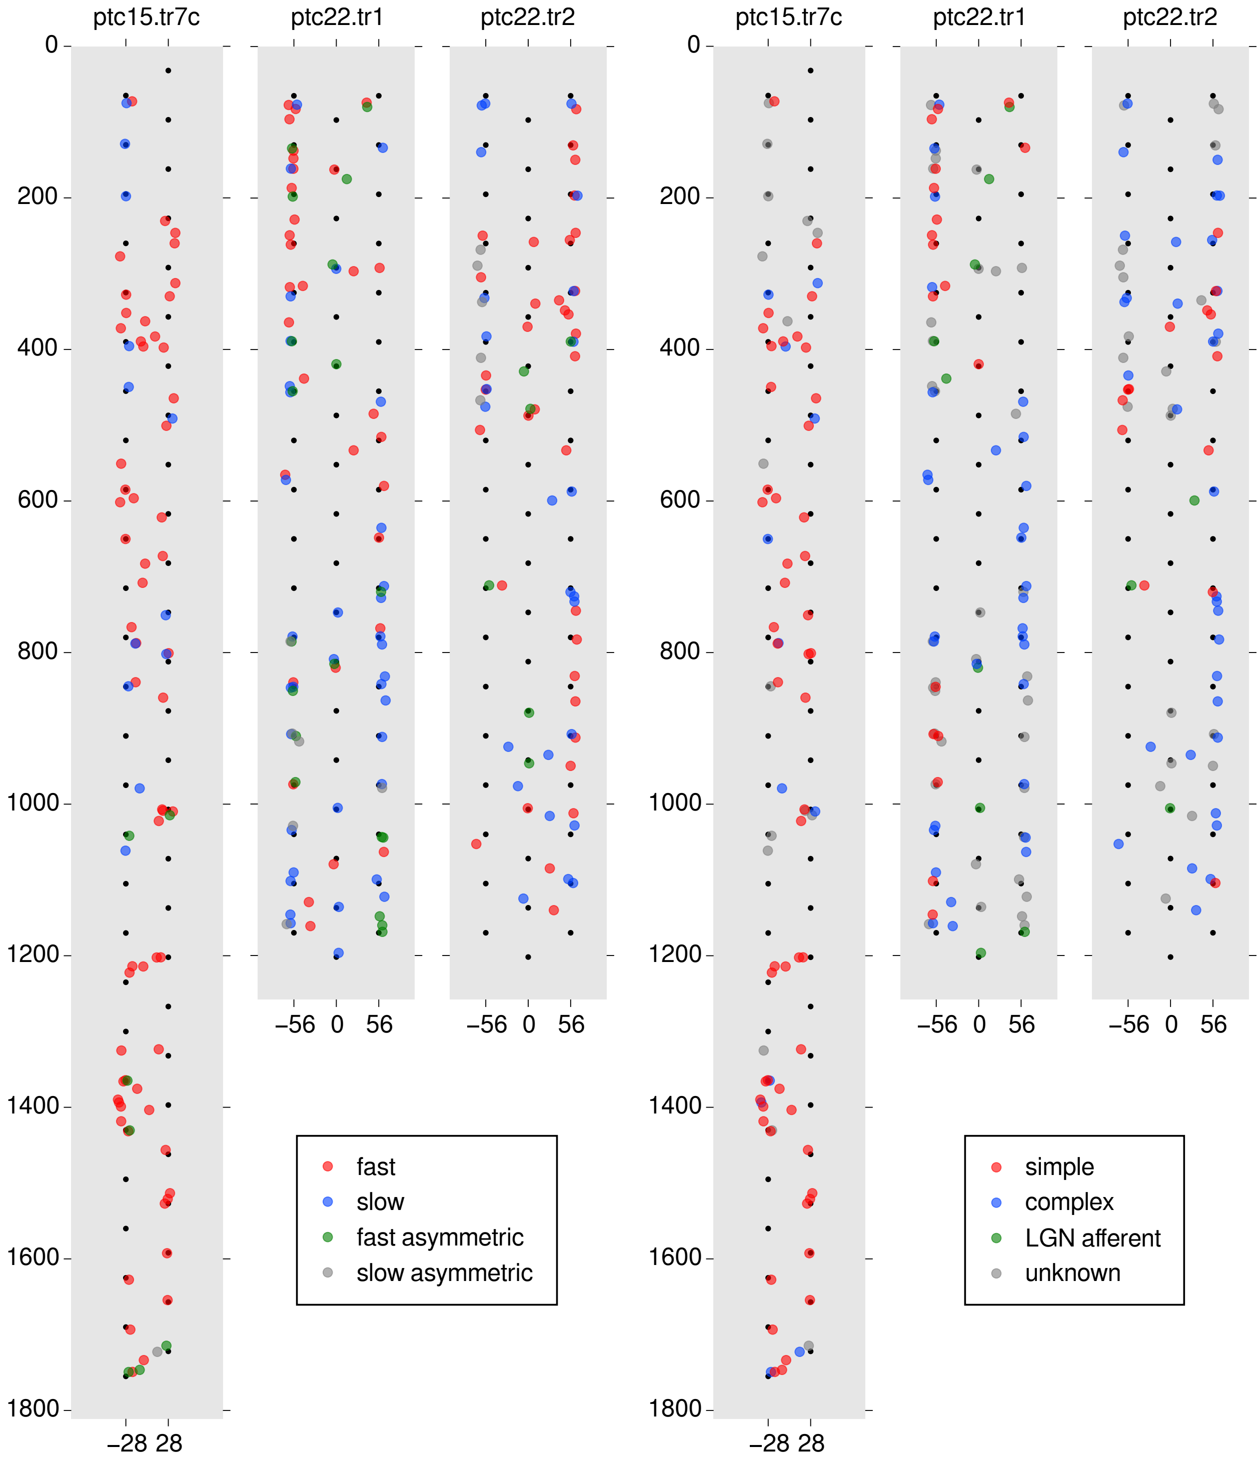

Figure 5.7: Spatial distribution of cell types, per track. ***Left***: Temporal spike shape classification. ***Right***: RF classification. Same layout as in Figure 4.4A. Units are in μm.

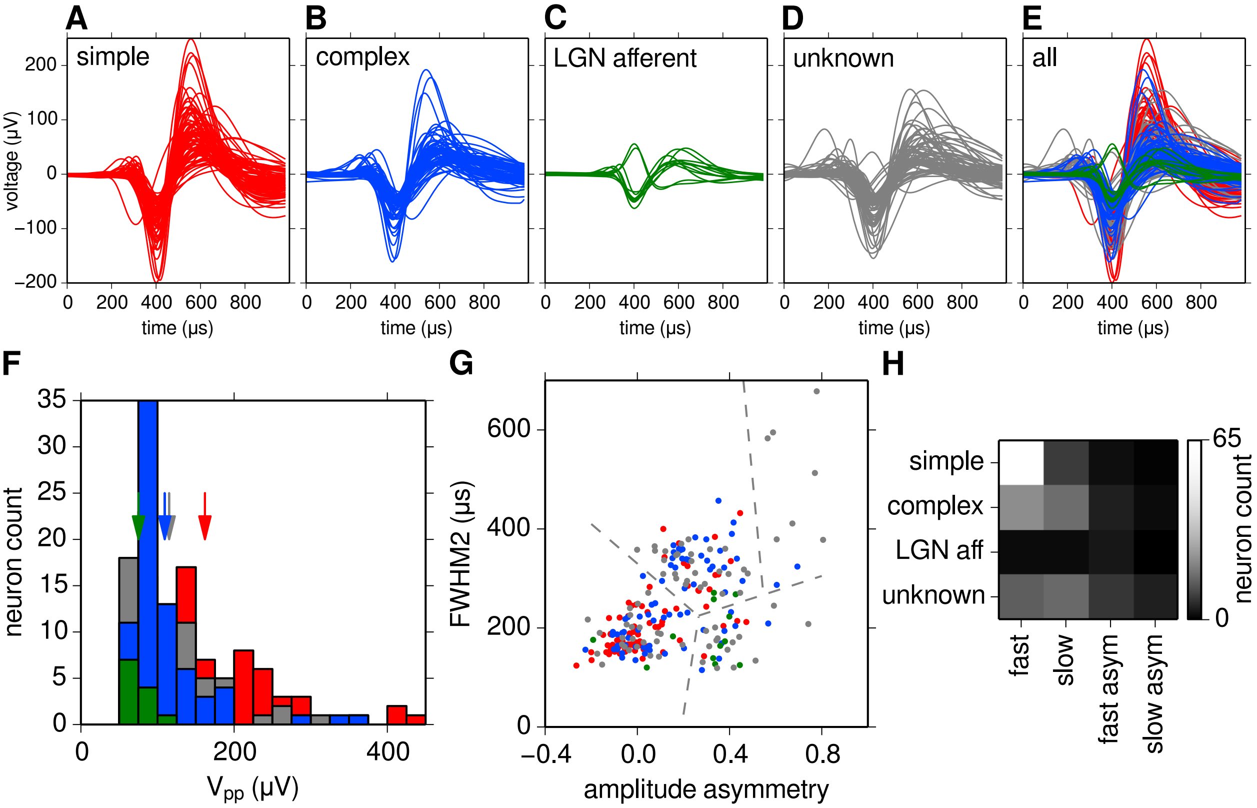

Figure 5.8: Comparison of RF type vs. spike type. ***A*–*D***: Overplotted templates of simple, complex, putative LGN afferent, and unknown RF cell types. ***E***: All RF cell types overplotted for comparison. ***F***: Overlapping distributions of primary channel peak-to-peak voltages of all four RF cell types. Coloured arrows indicate means of their respective distributions. ***G***: Same spike type cluster space as in Figure 5.2D, but with points coloured according to RF type instead of spike type. ***H***: Cell count matrix of RF type (rows) vs. spike type (columns) for all 245 cells. **aff**: afferent. **asym**: asymmetric.

ric means of simple, complex, LGN afferent and unknown cells were 0.33 Hz, 0.26 Hz, 0.23 Hz and 0.011 Hz, respectively. Given that mean rates were calculated track-wide (Section 4.3), the lower mean rates of cells with unknown RFs may have simply been a result of their shorter durations of activity.

5.6 Discussion

5.6.1 Spike shape

Although cell typing by intracellular spike shape has been previously reported in cat visual cortex (McCormick et al., 1985; Nowak et al., 2003; Haider et al., 2010), most published cell typing by extracellular spike shape has been done in rodent cortex and hippocampus, and macaque cortex

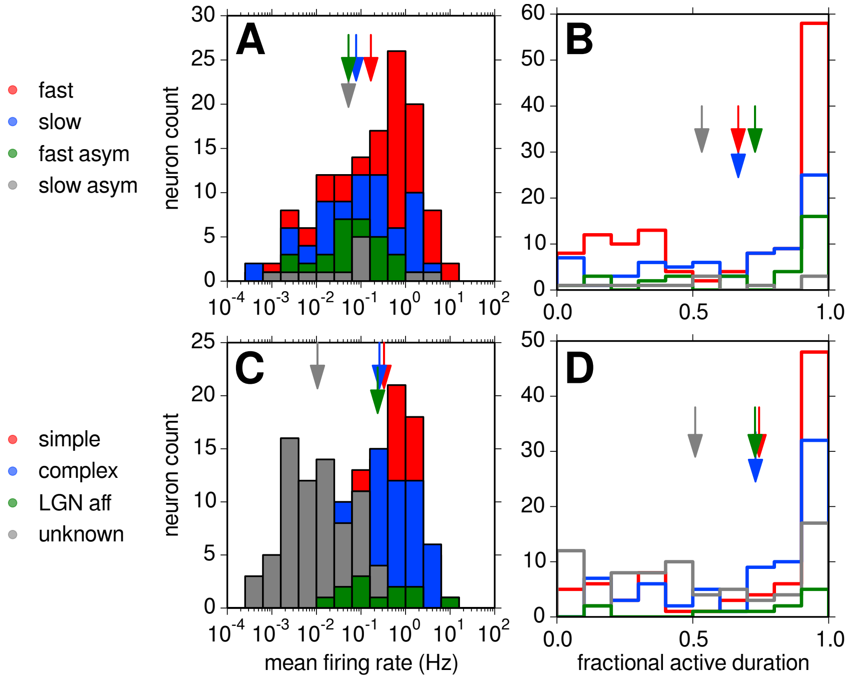

Figure 5.9: Mean firing rate distributions (***left***) and activity duration distributions (***right***) broken down by spike type (***top***) and RF type (***bottom***). Arrows indicate geometric (***left***) and arithmetic (***right***) means separately for each cell type. Plots in ***A*** and ***C*** are as in Figure 4.1B. Plots in ***B*** and ***D*** are as in Figure 4.2C, with bars left unfilled to better distinguish overlapping distributions. **aff**: afferent. **asym**: asymmetric.

(Csicsvari et al., 1998; Barthó et al., 2004; Luczak et al., 2007; Sirota et al., 2008; Niell and Stryker, 2008; Mizuseki et al., 2009; Sakata and Harris, 2009; Benchenane et al., 2010; Wilson et al., 1994; Gur et al., 1999; Mitchell et al., 2007). Showing that extracellular spike shape also clusters into two or more classes in cat visual cortex (Figure 5.2D) helps generalize extracellular spike shape (in this case, FWHM2 and amplitude asymmetry) as a useful cell typing property.

Given that waveform duration has been shown to be a reasonable proxy for physiological cell type (Barthó et al., 2004), slow cells may be considered putative excitatory pyramidal cells, and fast cells putative inhibitory interneurons. However, Robbins et al. (2013) discuss the caveat that axonal spikes (such as LGN afferents in primary visual cortex) are of short duration, and can be easily mistaken for interneuron spikes. A small fraction of units here (5%) were classified as putative LGN afferents (Table 5.1), based on their very small circular RFs as well as their lack of orientation tuning and spatially segregated ON and OFF responses, all of which were indicative of LGN cells. These units also tended to have short duration spikes (Figure 5.8C & green points in G), enforcing the point of Robbins et al. (2013) of the dangers of assuming that fast spiking units are necessarily inhibitory interneurons. Vigneswaran et al. (2011) showed that large pyramidal tract neurons (Betz cells) in macaque layer 5 M1 have extracellular spikes as short in duration as fast-spiking inhibitory interneurons, providing a further caveat to classifying cell type by extracellular spike shape alone.

Unfortunately, despite extensive efforts to minimize sampling bias in these experiments, the greater proportion of fast cells relative to slow cells (52% vs. 30%, see Table 5.1) is consistent with the observation of Henze et al. (2000) that slow, low firing rate pyramidal cells tend to be underrepresented in extracellular recordings. However, if track ptc15.tr7c is excluded due to it being so transcolumnar, the proportion of fast and slow cells from the remaining two tracks is roughly equal, at about 40% each.

Surprisingly, FWHM1, the width of the primary peak, was unimodal and did not cluster, while FWHM2, the width of the secondary peak, was bimodal and did cluster (Figure 5.2A–B). This agrees with observations by McCormick et al. (1985) that the slope of the repolarization phase of an intracellular spike clusters into high and low values, corresponding to narrow and wide secondary peaks in the extracellular spike. In comparison, other groups have found FWHM1 to be multimodal and useful for cell typing in rat prefrontal, motor, somatosensory, auditory, and secondary visual cortices, as well as hippocampus (Barthó et al., 2004; Luczak et al., 2007; Sirota et al., 2008; Sakata and Harris, 2009; Benchenane et al., 2010). None report on the clustering of FWHM2. Those studies also reported bimodality in interpeak interval (or trough to peak time), which was absent here (Figure 5.1C). Furthermore, the fast cells shown here were larger in amplitude than slow cells (Figure 5.2E–F). This is the opposite of the findings of the above studies, which find fast cells (i.e., narrow spiking cells, attributed to interneurons) to be of lower amplitude than slow cells (i.e., wide spiking cells, attributed to pyramidal cells). These differences may be due to differences in species and cortical area. They may also be due to extra damage from the relatively wide (~ 200 μm) polytrodes used here, compared to those used by the above studies (~ 100 μm). Differences in analog and/or digital filter settings can also affect spike shape, sometimes dramatically (Wiltschko et al., 2008). Finally, averaging waveforms across periods of recording with drift may distort templates, thereby affecting temporal or spatial waveform measures.

These differences in spike shape may be explored in the future by using narrower polytrode designs (Section 4.6.3), a new recording system with greater ADC dynamic range (16 bits instead of 12) allowing for open analog filters, and careful digital filtering of LFP from spike waveforms (Wiltschko et al., 2008). As described in Section 4.6.3, the spike shapes reported here may be somewhat distorted due to drift, which could be corrected for by splitting up each cell's spike train in time and averaging the spike waveforms within each time period. Finally, many more than the 245 neurons shown here should be examined to increase confidence in conclusions about spike shape. If differences in spike shape clustering remain, recordings in other species and cortical areas with the same recording system would help clarify whether or not those differences are indeed intrinsic to species and cortical area.

Interestingly, Sirota et al. (2008) recorded 767 cells from a wide variety of cortical areas in rat. They classified their plot of amplitude asymmetry vs. interpeak interval into two clusters: wide (excitatory) and narrow (inhibitory). However, that plot (their Figure 1C) also suggested another two or three clusters. Though not mentioned in the paper, perhaps these were truly distinct clusters as a result of combining recordings from such a wide range of cortical areas (6 in total, from V2 to M1), and therefore perhaps different cortical areas do indeed consist of cells of differing spike shape, even within the same species.

5.6.2 Spatial extent

In cat primary visual cortex, spatial extent (σ) of multichannel neuron waveforms was not a useful metric for cell typing (Section 5.3). Cells did not cluster according to spatial extent (Figure 5.3), spatial extent was not correlated with either spike shape measure, and spatial extent did not show any particular spatial distribution along the polytrode (Section 5.5). However, these null results may have been partly an artifact of calculating each cell's σ as the average σ of all of its member spikes. Since individual spikes are noisier than their cell's template, calculating σ from each spike and then averaging them may have obscured multimodalities in template spatial extent. Future work should also measure spatial extent directly from templates by fitting a 2D Gaussian (Section 3.5) to the template (or multiple templates over time, to account for drift), instead of to the individual spikes.

Another potential artifact is that cell distance from the polytrode affects signal amplitude: the greater the distance, the smaller the signal. This in turn may have an effect on the apparent spatial extent of the cell, as measured by fitting a 2D Gaussian to its spatial amplitude distribution. Since isolatable cell distances likely form a continuous distribution, their influence may smear out any multimodalities in the true 3D spatial extents of the population of cells. As described in Section 4.6.3, a 3D monopole/dipole current source model of extracellular spike potentials (Blanche, 2005) may reveal multimodalities in spatial extent, such as a monopole/dipole distinction, that may be useful for cell classification.

5.6.3 Receptive field type

On and off response overlap to drifting bar stimuli, temporal phase locking to sinusoidal grating stimuli, and STAs from m-sequence movies were all used to classify each cell's RF type as simple, complex, putative LGN afferent, or unknown (Figures 5.4 & 5.6). Each drifting bar and grating raster plot, and each m-sequence STA was visually inspected to subjectively classify each cell's RF type. Future work should employ a more quantitative method of determining on and off subfield overlap (Dean and Tolhurst, 1983; Kagan et al., 2002; Mata and Ringach, 2005), temporal phase-locking (Skottun et al., 1991), and the number, size, and shape of non-overlapping on and off subfields in the m-sequence STA. Though such quantification might not change the classification of many cells, automated methods would be faster and would increase confidence in RF classification.

Although on and off subfield overlap was Hubel and Wiesel's original method of distinguishing simple cells from complex cells (Hubel and Wiesel, 1959, 1962, 1968), a more recent and common method is to measure response phase locking to the temporal frequency of a drifting sinusoidal grating. This involves creating a PSTH of the repeated trials of the stimulus parameter combination with the greatest response, and then taking its Fourier transform to find the power at the grating's temporal frequency (F1). This is then compared to the DC power (the average firing rate), F0. If the F1/F0 ratio (i.e., modulation ratio) is greater than 1, the cell is considered simple, and if it is less than 1, the cell is considered complex (Skottun et al., 1991). However, generating the PSTH can be complicated by the fact that the eyes may slowly drift over the course of a long duration stimulus. This might be corrected for in an automated way, but visually inspecting the spike rasters for noticeable amounts of the correct temporal modulation was deemed here to be simpler. Further, Kagan et al. (2002) found that F1/F0 correlated poorly with the degree of RF subfield overlap in awake monkey V1, with some cells classified as complex according to their subfield overlap, yet classified as simple according to their F1/F0 ratio (but see Dean and Tolhurst (1983) and Mata and Ringach (2005)). Chance et al. (1999) & Kagan et al. (2002) argue that the key distinction is the degree of spatial phase invariance: complex cells have high phase invariance, while simple cells have low phase invariance. On and off subfield overlap may relate more directly to spatial phase invariance than does the F1/F0 ratio, and subfield overlap may therefore be the preferred metric.

Therefore, although F1/F0 has been a standard way to distinguish simple cells from complex, its true effectiveness in doing so is unclear, and relying instead on the subfield overlap may be a better strategy. Note that a combination of the two methods (subfield overlap estimated from drifting light and dark bars and m-sequence STA, and F1/F0 from drifting gratings) was used here.

Simple, complex and unknown RF types each made up about one third of all cells reported here, with only 5% of cells classified as putative LGN afferents. In comparison, Gilbert (1977) reported that most cells were complex (24% (44) simple, 68% (125) complex, and 9% (16) unclassified cells out of 185 cells from all layers of area 17 in 21 cats). Those recordings were performed with single-channel tungsten-in-glass electrodes, with multiple cells characterized within each penetration. However, a more recent review by Carandini et al. (2005) argues that the general consensus is that there are equal numbers of simple and complex cells in V1. Tracks ptc22.tr1 and ptc22.tr2 were somewhat consistent with Gilbert (1977) in that about twice as many cells were complex as simple (Table 5.1). Yet compared to that study, all 3 tracks had a higher fraction of cells with unknown RF type. This could be because wide polytrodes may cause more damage than narrow tungsten-in-glass electrodes.

Similar to Gilbert (1977), Kagan et al. (2002) reported 14% simple, 78% complex, and 8% unclassified cells out of 228 cells in V1 of 5 awake macaques, using single channel platinum-tungsten in glass or quartz electrodes. Kagan et al. (2002) also used spatial overlap instead of F1/F0 to distinguish simple cells from complex, and argued that use of F1/F0 overestimates the fraction of simple cells.

The precise distinction between simple and complex cells is somewhat controversial, as is the methodology used to distinguish them (Dean and Tolhurst, 1983; Carandini and Ferster, 2000; Kagan et al., 2002; Priebe et al., 2004; Mata and Ringach, 2005). Methodology and species may account for some of the variability in these relative proportions across studies. Ideally, natural scene movies should be used to calculate the RFs of both simple and complex cells, but doing so is, at least for now, complicated (Ringach et al., 2002a; Smyth et al., 2003; Willmore and Smyth, 2003; Sharpee et al., 2004; Touryan et al., 2005). What may be simpler is to use natural scene movie responses purely to distinguish simple and complex cells (see Section 6.7.2).

How might polytrode damage affect the relative numbers of simple and complex cells? Damage might make complex cells look more like simple cells, but not vice versa. If complex cells pool from many spatially overlapping simple cells (Hubel and Wiesel, 1962; Alonso and Martinez, 1998; Figure 6.20A), they should become increasingly simple the fewer simple cells they pool from due to cut afferents. If instead like simple cells they receive direct LGN input, but gain their spatial phase invariance from horizontal connectivity from other cells with different spatial phase LGN inputs (Chance et al., 1999; Tao et al., 2004; Figure 6.20B), then damage to their afferents might also make them more simple than complex. Either way, polytrode damage may bias the population toward more simple-like cells and fewer complex cells than would normally be the case.

Unknown RF cells had lower mean firing rates than all other RF types (Figure 5.9C), suggesting that they may have been damaged cells that had lost some of their afferents. There are several weak but independent lines of evidence that unknown RF type cells may have been composed mostly of damaged complex cells in particular. First, unknown RF cells and complex cells shared similar spatially uniform distributions across all three tracks (Figure 5.7). Second, they had similar peak-to-peak spike amplitudes (Figure 5.8F). Third, they were composed of a similar mix of slow and fast cells (Figure 5.8H). Finally, given reports that complex cells outnumber simple cells 3:1 or more (Gilbert, 1977; Kagan et al., 2002), all else being equal, unknown RF cells are more likely to be damaged complex cells than damaged simple cells.

Why might the cells labelled as LGN afferents show up as a distinct RF type? Are these truly LGN afferents, with an axon passing close to the polytrode, or are they instead a distinct cortical cell type that has not been categorized before? The characteristics of the putative LGN afferent RF type described here are similar to what Schiller et al. (1976a) label as T cells, or what Kagan et al. (2002) label as “monocontrast” cells, which they suggest are a third class of RF. With only 12 cells labelled as LGN afferent in the present data set, these questions remain open for the time being.

In retrospect, the use of 40 ms m-sequence frame times and reverse correlation timepoints for tracks ptc22.tr1 and ptc22.tr2 may have been a poor choice. Examination of the finer 20 ms frame times in ptc15.tr7c (Figure 5.6) reveals that averaging the 0–20 ms and 20–40 ms STAs would result in weaker RFs for nearly every cell, since the 20–40 ms STA would be washed out by noise in the 0–20 ms STA. The 0–40 ms STAs do exactly such averaging. Worse yet, there is a subfield sign inversion at around 60 ms, and therefore the STA calculated over 40–80 ms might result in near perfect cancellation of the fields. Receptive field dynamics studied by reverse correlation in macaque and cat V1 also suggest that 20 ms frame times are more suitable than 40 ms (Ringach et al., 1997a,b, 2003).

The switch to 40 ms m-sequence frame times was originally made to increase the duration of the m-sequence recordings, increase the number of spikes, and therefore improve the STA signal-to-noise ratio. Although more spikes were undoubtedly acquired, the STAs may have been inadvertently degraded by averaging over exactly the least complementary response time periods. For future experiments, m-sequence frame times should be restored to 20 ms, and to maintain the desired recording duration and number spikes, the entire m-sequence movie should be presented twice instead of just once.

5.6.4 Cell type comparison

Given the relative proportions of fast and slow cells, simple cells were more likely than expected to be fast cells than slow cells (Figure 5.8H). However, more cells are needed from more tracks and cats to make stronger conclusions about the relationships between spike type and RF type, as well as the spatial distributions of different cell types. Given the uncertainties in the relative laminar position of each track, it was difficult to align spatial distributions of different cell types across tracks, and therefore difficult to make conclusions about the true spatial distribution of cell types across cortical layers. Histology and CSD analysis should greatly reduce these uncertainties in future experiments. Future work could also examine how orientation tuning strength and significance (Section 4.5) might relate to spike and RF type.

6 Natural scenes & cortical states

6.1 Introduction

Visual neuroscience has traditionally relied on artificial stimuli such as drifting bars and gratings to characterize neuron response properties. However, more naturalistic stimuli elicit responses that are poorly predicted from responses to artificial stimuli (David et al., 2004; Olshausen and Field, 2005). Although artificial stimuli are easier to characterize and are of much lower dimensionality than naturalistic stimuli, relying too heavily on artificial stimuli may obscure insights into how the brain processes visual information. For a minimally biased characterization of a neural population in visual cortex, it is therefore important to consider responses to naturalistic stimuli in addition to artificial stimuli. Although sequences of natural images are spatially naturalistic, the gold standard is natural scene movies (Carandini et al., 2005), which are both spatially and temporally naturalistic.

As a complex dynamic system, the brain is never in exactly the same state twice. Yet experiments involving repeated stimulus presentation implicitly assume that it is, or at least that any resulting variability will average out over enough trials. There is increasing evidence that this may not always be the case, and that perhaps some of the response variability due to brain state can be controlled for (Arieli et al., 1996; Petersen et al., 2003; Fiser et al., 2004). There are many non-perceptual tasks that even primary sensory cortices might be engaged in during stimulus presentation. Such tasks might include attention (Roelfsema et al., 1998; Harris and Thiele, 2011), memory formation and recall (Ji and Wilson, 2007), reward encoding (Shuler and Bear, 2006; Chubykin et al., 2013), locomotion (Saleem et al., 2013), visualization, synaptic normalization, and low level cellular maintenance (Vyazovskiy and Harris, 2013; Xie et al., 2013). None of these tasks have much to do with encoding the currently presented stimulus. To deal with this multitude of tasks, cortex may need to engage in some kind of task switching, which could be reflected in cortical state changes. There are two broad categorizations of cortical state: synchronized and desynchronized (Berger, 1929; Destexhe et al., 1999; Harris and Thiele, 2011). The synchronized state is characterized by large amplitude low frequency fluctuations, and occurs during deep anesthesia, slow-wave sleep, and awake quiescent periods (quiet wakefulness). The synchronized state can be further subdivided into UP and DOWN phases (Destexhe et al., 1999; Harris and Thiele, 2011), corresponding to periods of higher and lower resting membrane potentials (Steriade et al., 1993a; Anderson et al., 2000; Sanchez-Vives and McCormick, 2000). The desynchronized state is characterized by low amplitude high frequency fluctuations, and occurs during light anesthesia, REM sleep, and awake attending behaviour. To minimize bias in the characterization of cortex, taking cortical state into account is a necessity.

Natural scene movie responses and the effects of cortical state are both examined in this chapter. Responses to natural scene movies in anesthetized cat V1 consisted of temporally precise and reliable events, with some cells responding to one movie clip but not at all to another similar one. Response correlations were very weak for most cell pairs but were significantly stronger between fast cells and between complex cells. The weakness of simple-complex and the strength of complex-complex pairs challenges the hierarchical model of complex cells. Synchronized and desynchronized cortical states were distinguished using the frequency content of deep-layer LFP, and switched spontaneously. Regardless of cell type, stimulus-evoked superficial (upper layer) cell firing rates were higher in the desynchronized than synchronized state, while the reverse was true for spontaneous activity. Deep layer cells showed a more heterogenous relationship. Contrary to other reports (Goard and Dan, 2009; Marguet and Harris, 2011; Zagha et al., 2013; Pachitariu et al., 2015) natural scene movie responses were more precise, sparse, and reliable during the synchronized state than the desynchronized state. Finally, response correlations between cell pairs were more influenced by cortical state than by the particular movie clip presented.

6.2 Natural scene movie responses

Two different sets of natural scene movies were used, as described in Section 2.4. Each set consisted of both long (several minutes) and short (4.5–5 s) clips, repeated anywhere from 5 to 400 times, usually inversely proportional to their length. Except where otherwise stated, this chapter focuses on the responses to short clips (with many repeats). Trial raster plots of responses to short clips are shown for a few example cells in Figures 6.1, 6.4 & 6.5. A given neuron's responses were often composed of short bursts of spikes that occurred reliably across trials at specific times during each trial, with high temporal precision. These temporally precise features are referred to here as response "events", and show up as vertical lines in the trial raster plots. Responses to repeated natural scene movies were thus very different from responses to repeated artificial stimuli, which elicited responses with much less temporal precision across trials (Figures 4.8 & 5.4).

Though temporally precise, natural scene movie response events were not always reliable, with events missing on some trials, sometimes for many consecutive trials. This is most noticeable in Figures 6.1 & 6.4, panels A–C, as well as Figure 6.5B. One explanation for this variability is changes in cortical state, examined in Section 6.6.

Each cell's natural scene movie raster plots were visually inspected for signs of response events. If a cell showed a reasonable degree of temporal precision and reliability across trials during at least one natural scene movie recording, it was considered "responsive". If its mean firing rate was ≥ 0.05 Hz, it was considered "active". Not all responsive cells were active, and not all active cells were responsive, although the two mostly overlapped (Figure 6.2). Table 6.1 shows the counts and percentages of responsive cells for each track and overall. The percentage of responsive cells was fairly consistent across all three tracks: 47% of all sorted cells and 87% of active cells were

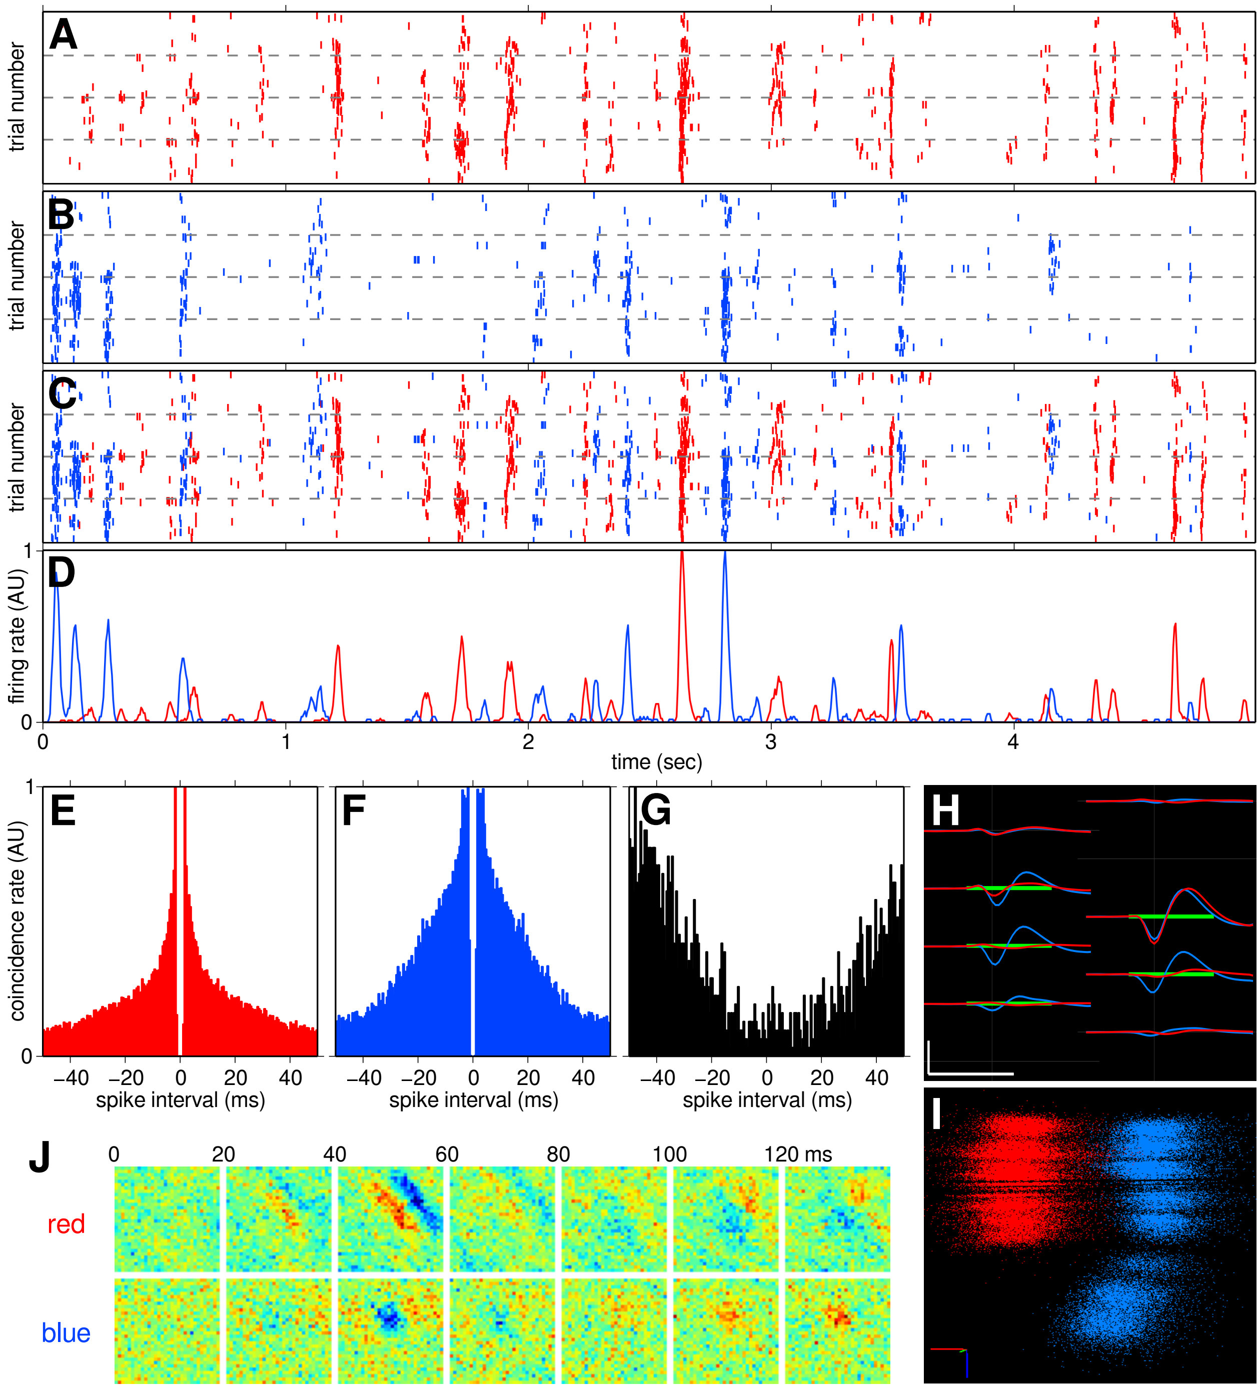
A
B
C
trial number
trial number
trial number
D
firing rate (AU)
1
0
0
1
2
3
4
time (sec)
E
F
G
coincidence rate (AU)
1
0
−40 −20 0 20 40
spike interval (ms)
−40 −20 0 20 40
spike interval (ms)
−40 −20 0 20 40
spike interval (ms)
H
I
J
0 20 40 60 80 100 120 ms
red
blue

Figure 6.1: (*Previous page.*) Natural scene movie responses of a pair of cells (red and blue) in ptc15.tr7c to 96 presentations of a 5 s clip from stimulus set A (Section 2.4). ***A*–*B***: Raster plots of both cells. Responses consisted of distinct temporally precise events visible as vertical lines in the raster plots. Horizontal dashed lines demarcate 4 different versions of the same movie, each at increasing contrast, separated by time gaps of 0–10 min and interleaved with unrelated movie stimuli. ***C***: Overlaid raster plots show that the responses for this pair were very distinct, with no major temporally precise events shared between them. ***D***: Normalized PSTH of both cells, calculated using 20 ms overlapping bins at 5 ms resolution, again showing the distinctness in their responses. ***E*–*F***: Normalized autocorrelograms of the two cells were noticeably different. ***G***: The cross-correlogram of the cell pair had a trough at $t = 0$, indicating that the cells tended to avoid firing together. Autocorrelograms and cross-correlograms were calculated using 0.5 ms wide bins. ***H***: Template waveforms of the cell pair, with channels and timepoints for PCA highlighted in green. Cells were only 16 μm apart, yet their template waveforms were noticeably different. Scale bar: 0.5 ms, 100 μV. ***I***: Cluster plot of the cell pair. The red horizontal axis is PC1, the green depth axis is PC2, and the blue vertical axis is spike time. The clusters remained distinct over time. Note that the autocorrelograms, cross-correlograms, templates and cluster plots were calculated from all recordings in the track, not only from the natural scene movie recording. ***J***: M-sequence STAs (see Figure 5.6) of the cell pair showed marked differences. STAs are 6.4° across. **AU**: arbitrary units.

track	all	active
ptc15.tr7c	47/81 (58%)	47/53 (87%)
ptc22.tr1	39/93 (42%)	34/41 (83%)
ptc22.tr2	29/71 (41%)	29/33 (88%)
total	115/245 (47%)	110/127 (87%)

Table 6.1: Counts of cells responsive (with temporally precise, reliable events in their raster plots) to natural scene movies, for each track and in total, listed in the numerators (percentages in parentheses). Counts in the 'all' column include all cells, regardless of their firing rate. Counts in the 'active' column include only those cells whose mean firing rate was $\geq$ 0.05 Hz during at least one natural scene movie recording.

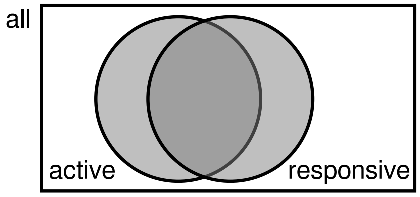

Figure 6.2: Venn diagram of active and responsive cells. Cells could be active (with mean firing rates $\geq$ 0.05 Hz), and/or responsive (with temporally precise, reliable events in their natural scene movie raster plots). Not all active cells were responsive, and not all responsive cells were active, although the two mostly overlapped. Set membership could change from one movie to the next (Figure 6.15).

responsive to natural scenes. Note that ptc22.tr1 had 5 cells that did not qualify as active (39 in the "all" column vs. 34 in the "active" column), but were nevertheless responsive to natural scene movies (3 of these 5 are shown in Figure 6.16).

To see if specific cell types (Chapter 5) were more or less likely to respond to natural scene movies, a matrix of responsivity vs. cell type was constructed (Figure 6.3). Pooling over all cells in all tracks, the only cell type that was significantly more likely than not to respond to natural scene movies was simple cells. In contrast, both fast asymmetric and unknown cell types were significantly more likely to not respond than to respond. However, distributions varied substantially from one

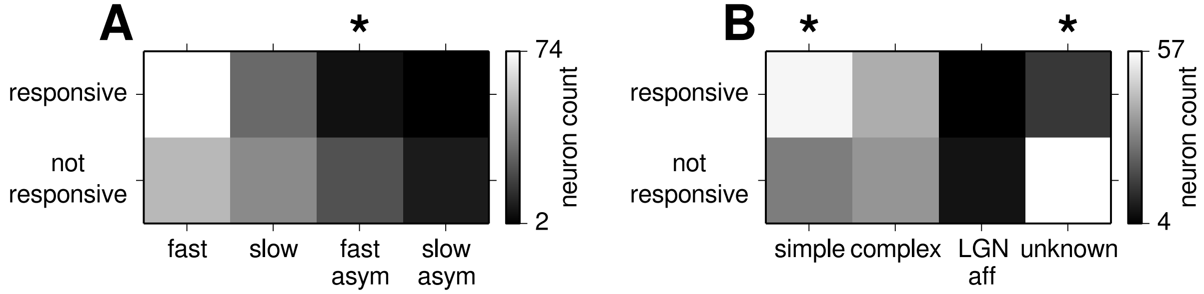

Figure 6.3: Responsive cell counts vs. cell type. For each cell type, matrices show the number of cells that were classified as responsive (top rows) or not responsive (bottom rows) with a qualitative degree of temporal precision and reliability to short natural scene movie clips. Asterisks indicate cell types with significantly unequal numbers of responsive and non-responsive cell counts (χ^2 test applied independently to each column, $p < 0.01$). ***A***: Natural scene movie responsiveness vs. spike type. ***B***: Natural scene movie responsiveness vs. RF type. Simple cells were the only cell type that was significantly more likely than not to respond reliably to natural scene movies. All 245 cells from all 3 tracks are included.

track to the next (not shown), so the overall distributions shown in Figure 6.3 may not be very conclusive.

To quantify natural scene movie response precision, PSTHs were constructed by binning spikes in the trial-triggered raster plots into 20 ms wide overlapping time bins at 5 ms resolution (Figures 6.1D, 6.4D & 6.5) or 0.1 ms resolution (6.17A). This bin width was chosen because 20 ms is roughly the membrane time constant of neocortical layer 5 (Mainen and Sejnowski, 1995) and hippocampal CA1 (Spruston and Johnston, 1992) pyramidal neurons. This is also the timescale at which hippocampal pyramidal cell spike times are best predicted by the activity of peer neurons, and therefore may be the timescale most relevant for cell assemblies (Harris et al., 2003). The widths of peaks in the PSTHs were measured by their FWHM. For those cells classified as responsive to natural scene movies, responses were remarkably sparse, with peaks in the PSTHs ranging 20–100 ms in width, separated by periods of silence (Figures 6.1D, 6.4D, 6.5 & 6.17A).

Each cell's pattern of response events was usually distinct. The cells in the pair shown in Figure 6.1 were separated (Section 4.4) by only 16 μm, yet their response patterns were very distinct. Their cross-correlogram (Figure 6.1G) had a trough at $t = 0$, indicating that they rarely fired together. Both cells were fast and simple with RFs that were close to 180° out of phase (Figure 6.1J). As measured by a drifting bar stimulus, their orientation preference (Section 4.5) differed by 60° and their tuning strength differed by 0.64 (not shown). The marked difference of these cells' responses to natural scene movies despite their physical proximity is further evidence of valid spike sorting.

The cell pair in Figure 6.4 are a counterexample. Though separated by a greater distance (31 μm), the pair's response patterns were remarkably similar. All significant response events were shared by the pair. However, PSTH peak amplitudes (Figure 6.4D) differed between the two in a nonlinear way. Furthermore, a cluster plot of the two cells in Figure 6.4I shows that they were well separated during spike sorting, suggesting that their similarity in event timing was not an artifact of spike sorting errors. Again, both cells were fast and simple. Although the cells' RFs had only minor differences (Figure 6.4J), their orientation preference differed by 10° and their tuning strength differed by 0.48 (not shown).

Temporally precise response events could be shared by more than just two cells. Figure 6.5 shows two groups of cells from a recording in ptc22.tr2. One group was a triplet of cells (Figure 6.5A) and the other a quintuplet (Figure 6.5B). All cells in each group shared at least one obvious response event (black arrows), despite the maximum cell separation of 204 and 277 μm in each group, respectively. Additionally, two events seemed to be shared across the two groups (grey arrows). The separation of the mean positions of the two groups was 509 μm, while the most separated pair of cells across the two groups was 715 μm apart. The triplet of cells was in the more superficial cortical layers, and the quintuplet in deeper layers. Note that in contrast with the cells shown in the transcolumnar ptc15.tr7c recording in Figure 6.4, this recording was from a mostly columnar track (Figure 4.10C, right panel, Figure 5.6, right panel), suggesting that sharing of response events between cells is possible for both columnar and transcolumnar groups of cells. The two groups of cells in Figure 6.5 comprised 18% and 29% of only 17 active cells in this particular natural scene movie recording, and 4% and 7% of the 71 total cells in this track, respectively.

Figure 6.16 shows another example of shared response events. Examination of responses across all natural scene movie recordings suggested that events shared across many cells tended to be broader in time. Those shared by few cells seemed to be more temporally precise. However, this observation was not quantified.

6.3 Natural scene movie response correlations

To visualize and quantify the similarity of natural scene movie responses across an entire population of cells, correlations of PSTHs were calculated for all active cell pairs within a given movie recording, or segment of a recording. Pearson's correlation (ρ) was calculated from the PSTHs of each pair of active cells by

$$\rho = \frac{\overline{xy} - \bar{x}\bar{y}}{\sigma_x \sigma_y} \,, \tag{6.1}$$

where x and y are the PSTHs of each cell, σ_x and σ_y are their standard deviations, and the overline indicates the mean. Calculating ρ for all cell pairs resulted in a matrix, which was ordered top to bottom and left to right by active cell depth along the length of the polytrode.

PSTH correlation matrices from two recordings from the transcolumnar track ptc15.tr7c are shown in Figure 6.6A & B (left panels). For each cell pair, colour represents ρ, and blank rows and columns represent inactive cells. In both recordings, a small minority of cell pairs had ρ values as high as 0.9. For comparison, the example cell pairs in Figures 6.1 and Figure 6.4 from the same

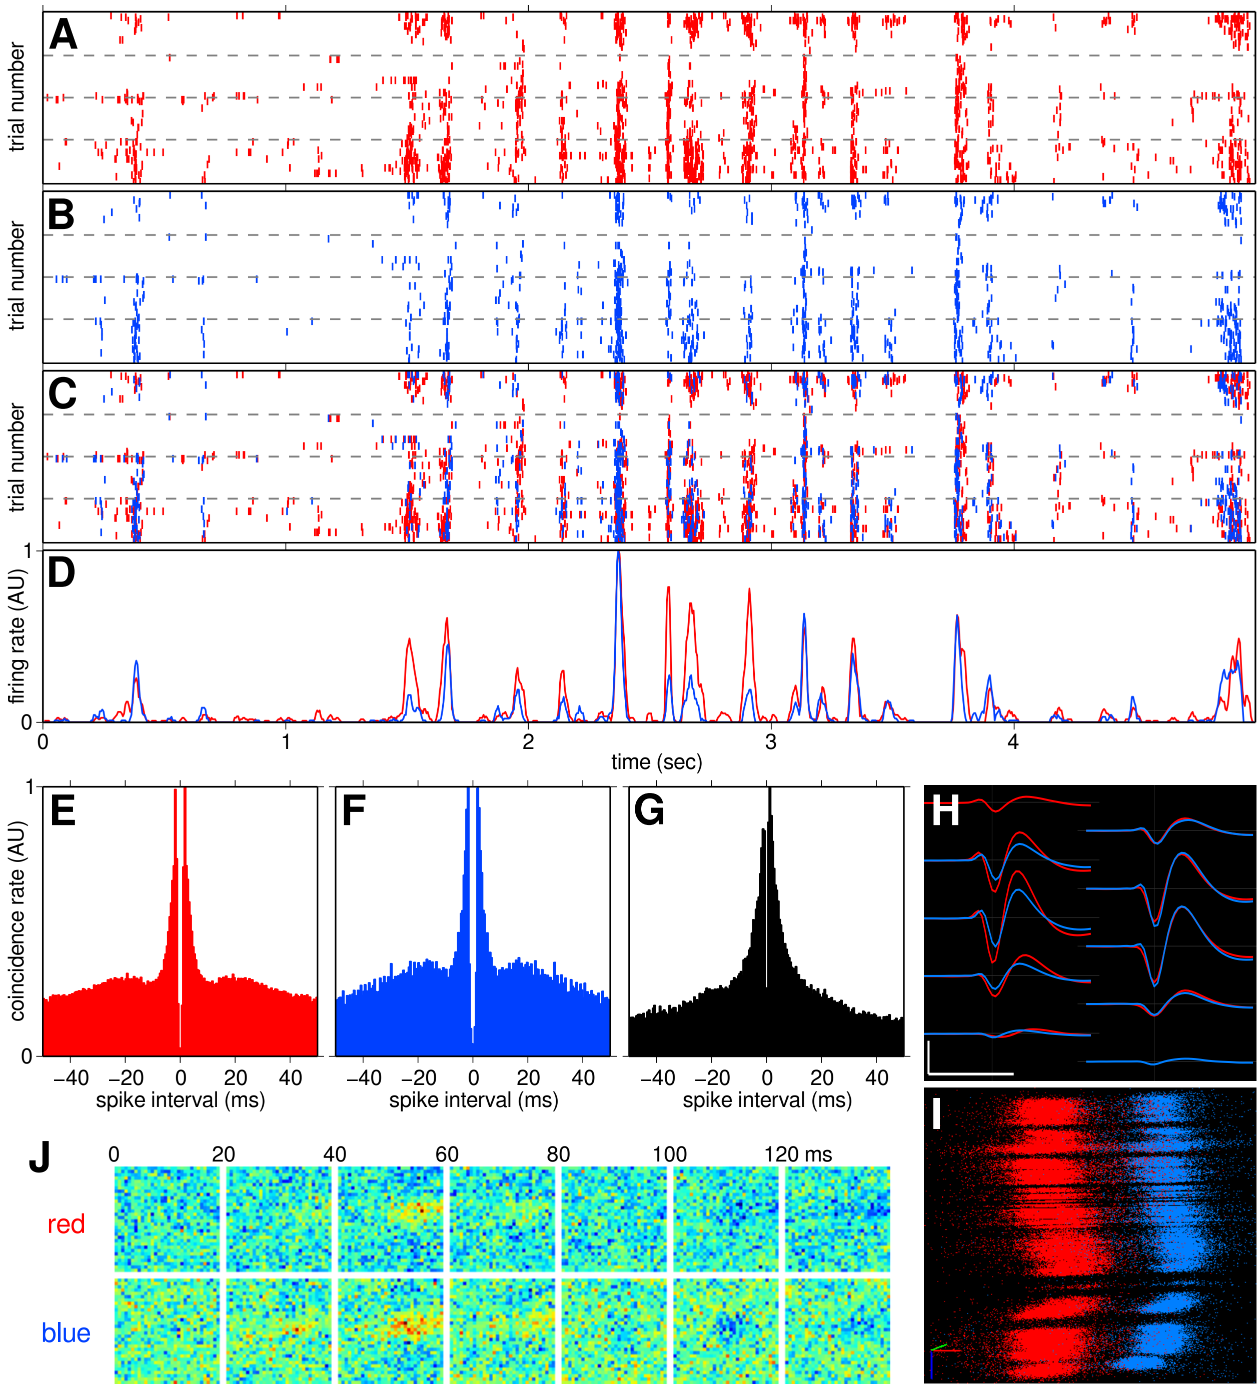
A
B
C
D
trial number
trial number
trial number
firing rate (AU)
1
0
0
1
2
3
4
time (sec)
E
F
G
H
I
coincidence rate (AU)
1
0
−40
−20
0
20
40
spike interval (ms)
spike interval (ms)
spike interval (ms)
J
0
20
40
60
80
100
120 ms
red
blue

Figure 6.4: (*Previous page.*) Same as Figure 6.1 but with a different cell pair from the same recording, located 31 μm apart. ***A*–*D***: Natural scene movie responses for this pair were very similar, with all major temporally precise events shared between them. ***E*–*F***: Normalized autocorrelograms of the two cells had only slight differences. ***G***: The cross-correlogram of the cell pair was peaked around $t = 0$ and reasonably symmetric, indicating that the cells tended to fire together, but not in any consistent order. ***H***: Template waveforms of the cell pair, though similar, had different spatial distributions. Scale bar: 0.5 ms, 100 μV. ***I***: Cluster plot of the cell pair. The red horizontal axis is horizontal position on the polytrode, the green depth axis is vertical position on the polytrode, and the blue vertical axis is spike time. The clusters remained distinct over time. Clusters were also distinct in PCA space (not shown). ***J***: M-sequence STAs of the cell pair showed minor differences.

recording as in Figure 6.6A had $\rho = -0.12$ and $\rho = 0.83$, respectively. The small minority of strongly correlated pairs tended to be close to the diagonal, i.e., the cells in such pairs tended to be physically close to one another. However, the vast majority of cell pairs, some of which were also close to the diagonal, had ρ close to 0. This was also visible in the overall distributions of ρ values (middle panels), which were peaked very close to 0. Mean ρ (red vertical dashed line) was slightly but significantly positive for both recordings, at 0.032–0.05 (two-sided Student's t-test, $p < 5 \times 10^{-14}$). Both recordings had a U-shaped dependence of ρ on cell pair separation (right panels). This was consistent with a transcolumnar track in which the cells near the top of the polytrode had the same orientation preference as those near the bottom, turning through 180° of orientation preference along the way (Figure 4.10C, left).

Despite stimulation by different movie clips and despite being separated by 12 hours, the correlation matrices in Figure 6.6A and B were similar for the set of active cells common to both recordings. Their difference is shown in Figure 6.6C. The difference ($\Delta\rho$'s) had even fewer strongly valued pairs (left panel), indicating that most of the correlations were similar in the two recordings. The overall $\Delta\rho$ distribution was symmetric with a mean of 0.01 (middle panel) that was not significantly different from 0 (two-sided Student's t-test, $p = 0.076$). There was no apparent dependence of $\Delta\rho$ on pair separation (right panel).

To examine how PSTH correlations might relate to cell types, ρ values were calculated for active cells during natural scene movie recordings (both short movies with many repeats and long movies with few repeats) in all 3 tracks, and binned according to the spike and RF cell type combination of each pair. For this analysis, to simultaneously maximize the total duration of recordings examined while also allowing for at least a rough estimate of PSTHs during long movie clips with few trials, all natural scene movie recordings with at least 5 trials per movie clip were included. This resulted in 16, 104 and 53 min of total movie presentation time for tracks ptc15.tr7c, ptc22.tr1 and ptc22.tr2, respectively. If a cell pair was active in more than one natural scene movie recording, only its maximum ρ was kept. Results were divided up into transcolumnar (ptc15.tr7c) and columnar (ptc22.tr1 & ptc22.tr2) tracks. The mean correlation (ρ) of each cell type pair was represented by colour in symmetric matrices (Figure 6.7). Those cell pair types whose $\bar{\rho}$ was not significantly

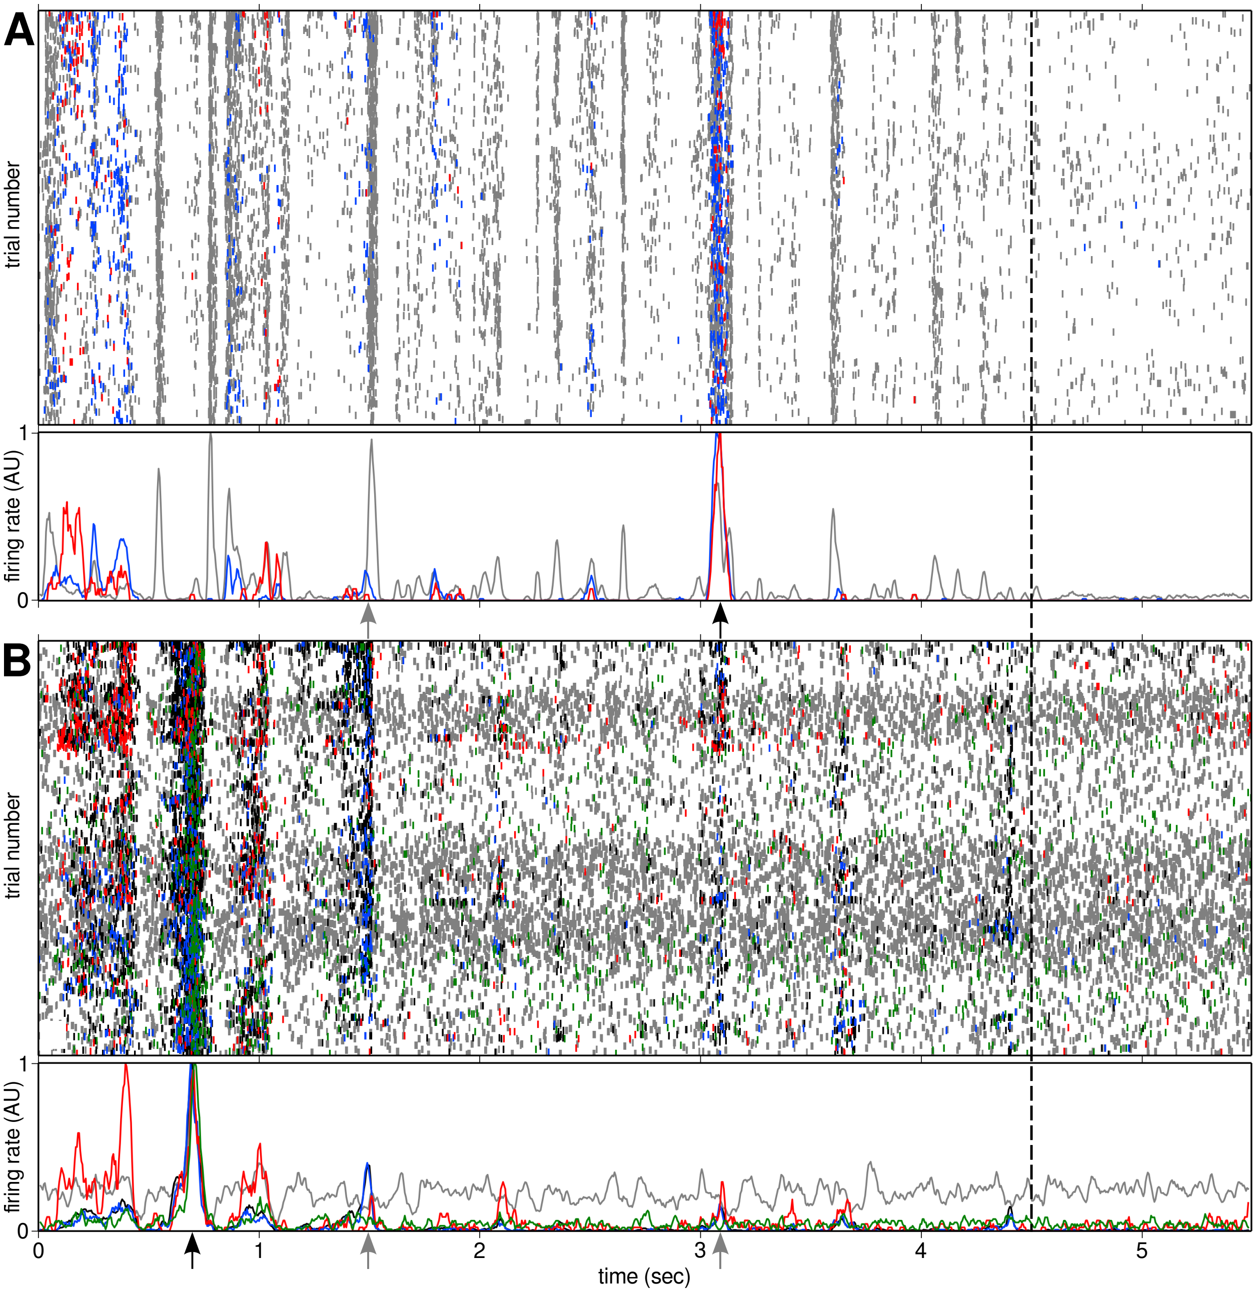

Figure 6.5: Temporally precise response events can be shared by larger more disparate groups of neurons. Raster plots and normalized PSTHs (20 ms overlapping bins at 5 ms resolution) of a superficial layer triplet (***A***) and a deep layer quintuplet (***B***) of cells from ptc22.tr2, responding simultaneously to 400 presentations of a 4.5 s natural scene movie clip from the 'new' movie set (Figure 2.5). Cells are distinguished by colour within each group separately. Trials were separated by 1 s blank periods (black dashed line). Each group exhibited a prominent event shared by all cells within the group (black arrows). Weaker events were shared between groups (grey arrows). Maximum separation of cells within ***A*** and ***B*** was 204 and 277 μm respectively, and across both groups was 715 μm. Polytrode depths of all 3 cells in ***A*** were < 400 μm, while 4/5 cells in ***B*** had depths > 700 μm (one was at 587 μm).

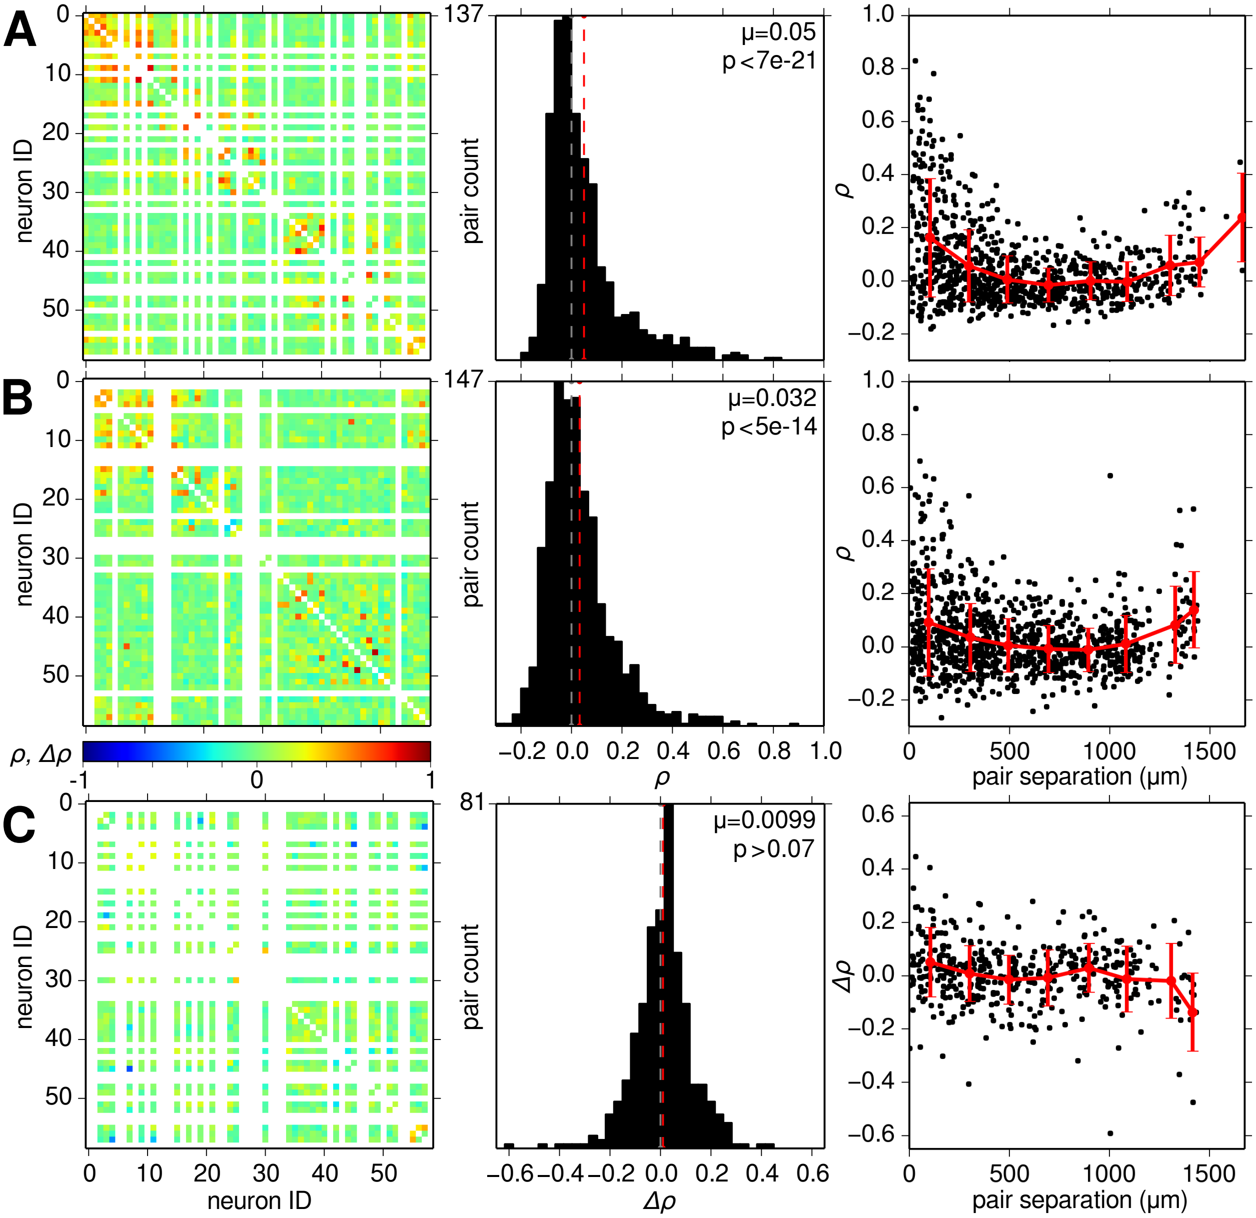

Figure 6.6: Natural scene movie PSTH correlation matrices (***left***), distributions (***middle***), and separation dependency (***right***) for all active cell pairs, for two recordings (***A*** & ***B***) in ptc15.tr7c with two different 5 s movie clips, separated by 12 hours. Correlation matrices were ordered top to bottom and left to right by active cell depth along the length of the polytrode. ***A*** corresponds to the recording in Figures 6.1 & 6.4. ***Left***: White gaps in the matrices represent cells that were active in one recording but not the other. Colour scale bar applies to all 3 matrices. ***Middle***: Red vertical dashed line denotes the mean of each distribution. ***Right***: Red markers are the mean ±1 standard deviation of each pair separation bin. PSTH correlations had a weak U-shaped dependency on cell separation, as might be expected for this transcolumnar track. ***C***: The difference between correlation matrices in ***A*** and ***B***, along with the distributions of the difference values and their separation dependency. PSTH correlations were consistent across the two recordings, for specific cell pairs (***left***), in their distributions (***middle***), and in their dependence on cell separation (***right***).

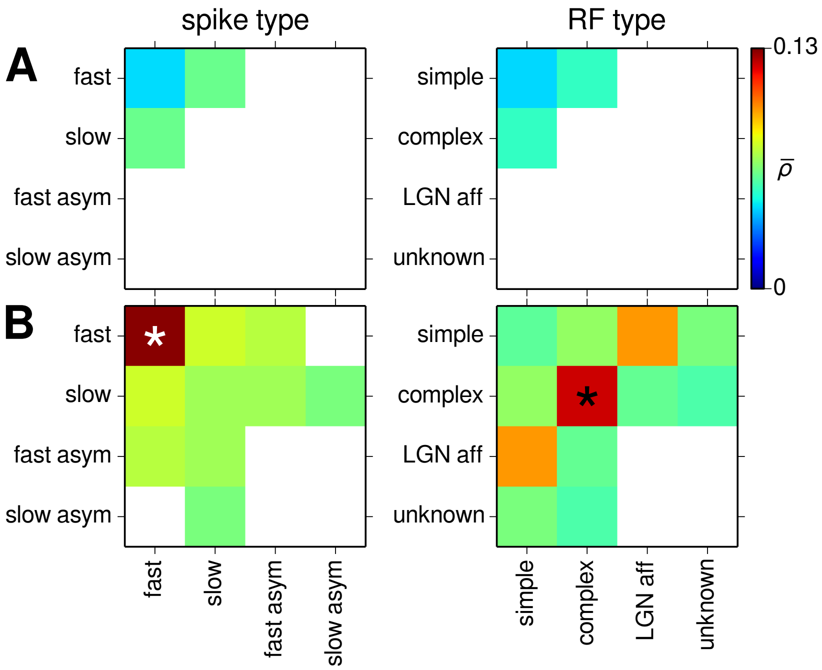

Figure 6.7: Matrices of natural scene movie mean PSTH correlation ($\bar{\rho}$) of active cells, indexed by spike type (***left***) and RF type (***right***). Blank entries had $\bar{\rho}$ that was not significantly different from 0 (two-sided Student's t-test, $p < 0.0005$). ***A***: $\bar{\rho}$ matrices for transcolumnar track ptc15.tr7c (1238 cell pairs). $\bar{\rho}$ was very low and was significant for only 4 cell type combinations: fast-fast, fast-slow, simple-simple, and simple-complex. None were significantly different from one another ($p > 0.35$). ***B***: $\bar{\rho}$ matrices for the two columnar tracks in ptc22 (1976 cell pairs). $\bar{\rho}$ was significant for 13 cell type combinations, and 3 of them were noticeably higher than the rest: fast-fast, simple-LGN afferent, and complex-complex. Fast-fast was significantly different from all other spike type combinations (white asterisk, $p < 0.01$, 268 pairs), and complex-complex was significantly different from all other RF type combinations, save for simple-LGN afferent (black asterisk, $p < 0.005$, 549 pairs). Simple-LGN afferent was not significantly different from any other RF type combinations, likely because of low N (37 pairs). Colour bar applies to all panels.

different from 0 (two-sided Student's t-test, $p < 0.0005$) were shown as blank.

Consistent with Figure 6.6, the transcolumnar track (Figure 6.7A) had very low $\bar{\rho}$ values for all pairwise combinations of cell types. Only 4 out of a total possible 20 spike and RF cell type combinations were significantly different from 0. However, cell pairs from the two columnar tracks resulted in higher $\bar{\rho}$ values, with a total of 13 significant cell type combinations (Figure 6.7B). Pooling was justified across the two tracks in ptc22 not only because both were columnar, but also because their separate $\bar{\rho}$ matrices were similar (not shown). Fast-fast, simple-LGN afferent, and complex-complex cell type pairs had visibly higher $\bar{\rho}$ than the rest. Comparing the distribution of each cell type combination to all others revealed that in the columnar tracks only fast-fast was significantly greater than all other spike type combinations (two-sided Student's t-test, $p < 0.01$), while only complex-complex was significantly greater than all other RF type combinations, save for simple-LGN afferent ($p < 0.005$). Likely due to low N (37 cell pairs), simple-LGN afferent itself was not significantly different from any other RF type combinations. In the transcolumnar track (Figure 6.7A), none of the 4 cell type combinations had significantly different $\bar{\rho}$ from one another ($p > 0.35$).

6.4 Cortical states

Cortical state, as classified by the degree of synchronization or desynchronization of a region of cortex, can be measured in a variety of ways. Measures can be based on the EEG, LFP, MUA, or even membrane potential (Poulet and Petersen, 2008; Li et al., 2009; Sakata and Harris, 2009; Renart et al., 2010; Saleem et al., 2010; Harris and Thiele, 2011; Okun et al., 2012; Poulet et al., 2012). The measure used here was the deep layer LFP L/(L+H) ratio (Saleem et al., 2010), a variant of the L/H ratio (Li et al., 2009), and is referred to here as the synchrony index (SI). Deep layer LFP is considered a better indicator of cortical state than superficial layer LFP (Saleem et al., 2010). L and H represent lowpass and highpass LFP bands, taken here to be 0.5–7 Hz and 15–100 Hz respectively (Figure 6.8A). To calculate the L/(L+H) ratio, first an LFP spectrogram was constructed by taking the fast Fourier transform (FFT) of the LFP within overlapping time bins (30 s wide at 5 s resolution) to find the power spectrum within each time bin. Next, the power of the L and H bands was summed within each time bin, and the L/(L+H) ratio was calculated, resulting in a series of SI values representing cortical state as a function of time.

SI ranges from 0 to 1. Periods dominated by low frequency LFP have high SI, and periods dominated by high frequency LFP have low SI. High SI also corresponded to high LFP amplitude, and low SI to low LFP amplitude (Figure 6.8B & C). Synchronized and desynchronized states are two extremes of a continuous spectrum, with no sharp division between them, and cortical states may even be multidimensional (Harris and Thiele, 2011). Nevertheless, recording segments were classified here as either synchronized or desynchronized by applying a pair of SI thresholds: SI $<$ 0.8 was considered desynchronized, and SI $>$ 0.9 was considered synchronized. Periods that fell in between the two thresholds were left undefined.

Figure 6.9 shows a spontaneous change in cortical state during a 39 min recording in track ptc22.tr1, consisting of 400 presentations of a 4.5 s natural scene movie clip. Superficial and deep layer LFP, the deep layer LFP spectrogram, SI, and the population spike raster plot are all shown on the same timescale. At 2/3 of the way through the recording, the LFP spontaneously switched from low amplitude high frequency activity to high amplitude low frequency activity. As a result, SI switched from low to high, representing a switch in cortical state from desynchronized to synchronized. In this example, firing rates of most cells were higher during the synchronized period. Note that even while stimulus and cortical state remained constant, near the end of the recording, superficial cells spontaneously became silent (grey dashed line). Figure 6.10 shows the opposite:

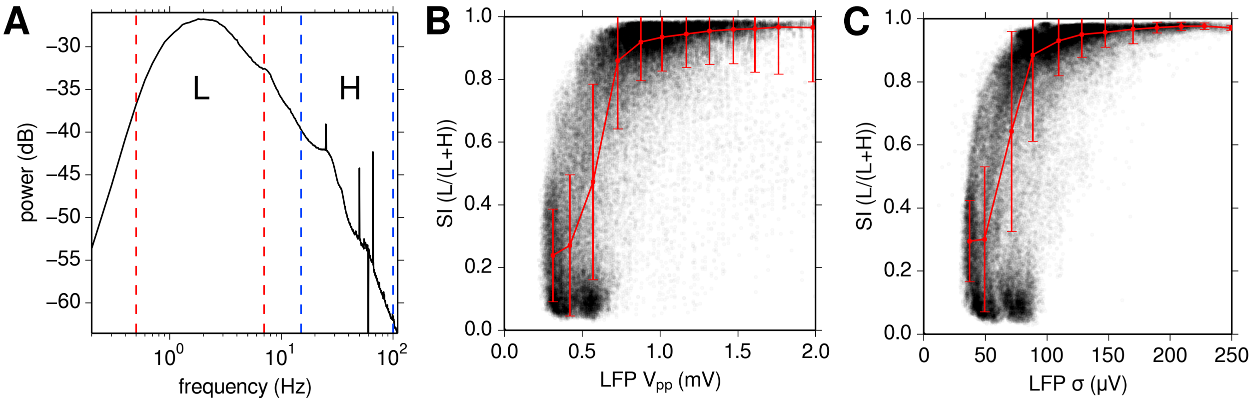

Figure 6.8: Power spectral density (PSD) and amplitude of deep layer LFP. ***A***: PSD of all 15 tracks in Table 2.2. Power is in decibels relative to 1 mV^2. Red and blue vertical dashed lines mark the limits of the low (L) and high (H) bands, respectively, used to calculate the synchrony index, SI. On this log-log scale, the low band is roughly centered on the broad peak at ~ 2 Hz. Some of the attenuation below 1 Hz is due to analog filtering of the LFP during acquisition. The narrow positive peaks at 25, 50 and 66 Hz are stimulus induced. The narrow negative peak at 60 Hz is from filtering out mains interference using a 0.5 Hz wide elliptic notch filter. ***B***: SI vs. peak-to-peak LFP amplitude (V_{pp}), measured from 30 s wide overlapping bins at 5 s resolution. ***C***: SI vs. LFP standard deviation (σ). Red markers are the median ± 1 standard deviation. Data from cat ptc21 caused a small dip in ***B*** & ***C*** at SI > 0.8 (not shown). Excluding ptc21, SI was a monotonically increasing function of LFP amplitude, as measured by both V_{pp} and σ.

cortical state spontaneously switched from synchronized back to desynchronized. Although the two states were distinct, the spectrograms and SI plots during the desynchronized periods in both recordings were far from homogeneous, suggesting other potential categorizations within that state. These were left unexplored.

Consistent with the literature, cortical state in cat primary visual cortex spanned a continuum from large amplitude low frequency fluctuations (synchronized state) to small amplitude high frequency fluctuations (desynchronized state). Also consistent, state changes were more noticeable in deep layer LFP than in superficial layer LFP (not shown, although Figures 6.9A & 6.10A give an indication).

Figure 6.11 shows the distribution of SI for each acquired track (including those that were not spike sorted, see Table 2.2), and overall. The distributions varied widely across tracks. Some were unimodal, others bi- and even trimodal. Some were dominated by low SI, but most were dominated by high SI. Unfortunately, the desynchronized state was therefore poorly represented, especially in those tracks that were spike sorted (labelled in red), resulting in fewer opportunities to compare the effects of cortical state on spiking. Experiment logs were examined for the cause of these wildly different distributions in cortical state. Tracks with greater amounts of synchronization were associated with application of various drugs, most especially atropine (to reduce secretions and/or increase heart rate). Other implicated drugs were tropicamide (to dilate pupils), dobutamine (to

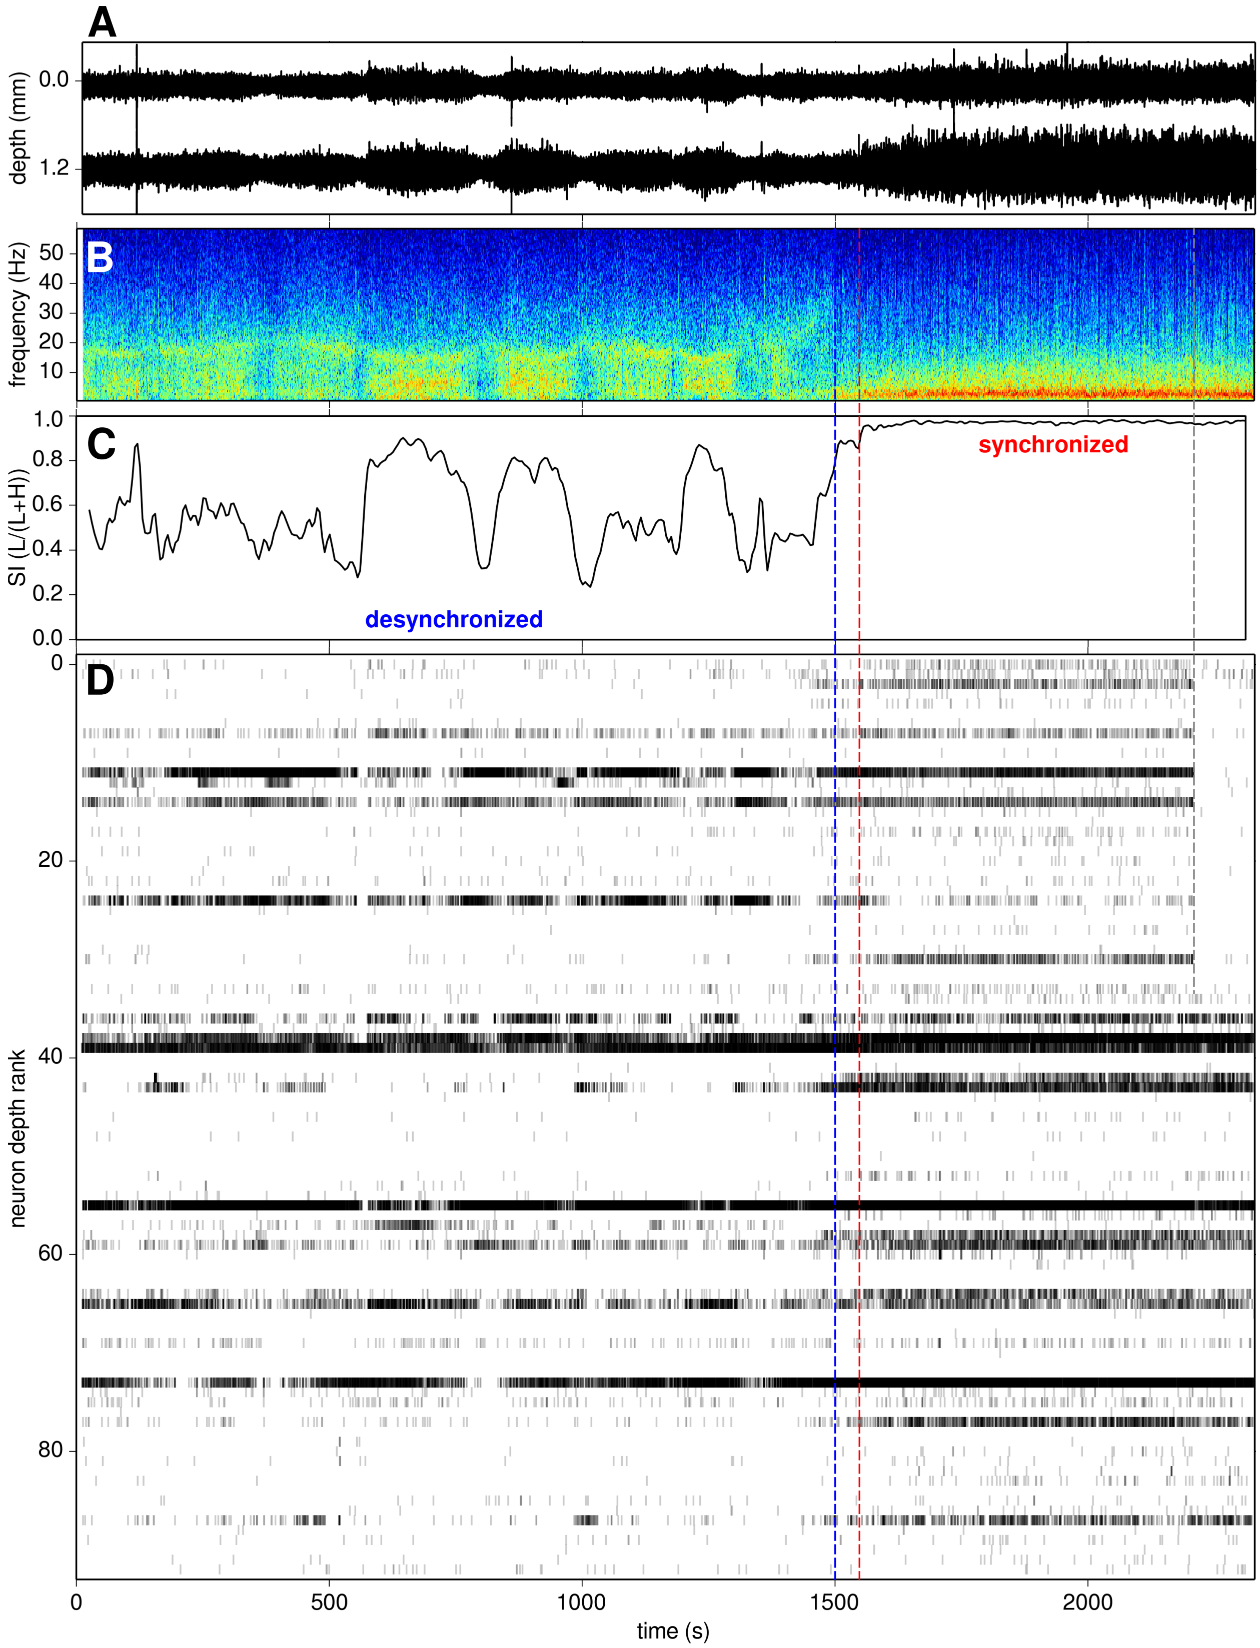
A
depth (mm)
0.0
1.2
B
frequency (Hz)
50
40
30
20
10
C
SI (L/(L+H))
1.0
0.8
0.6
0.4
0.2
0.0
desynchronized
synchronized
D
neuron depth rank
0
20
40
60
80
0
500
1000
1500
2000
time (s)

Figure 6.9: (*Previous page.*) Cortical state during 400 identical presentations of a 4.5 s natural scene movie clip in track ptc22.tr1. Over the course of the recording, cortical state spontaneously changed from desynchronized to synchronized. ***A***: LFP signal of the most superficial and deepest channels. ***B***: Spectrogram (frequency histogram) of the deepest channel. Red and blue represent high and low power, respectively. For display, this spectrogram was constructed using 2 s wide overlapping time bins at 0.5 s resolution. ***C***: Synchrony index (SI) calculated from the L/(L+H) ratio of the spectrogram (which itself was constructed using wider time bins than in ***B***, see text). ***D***: Population raster plot, showing all 93 cells in this track, including those that fired 0 spikes during this movie. Rasters are vertically ordered by cell depth. Transparency is used to increase detail. The blue vertical dashed line marks the end of the desynchronized state, and the red vertical dashed line marks the beginning of the synchronized state. UP and DOWN phases (Section 6.5) are visible in the synchronized state as broadband vertical lines in ***B***. The vertical grey line marks a sudden drop in superficial layer firing, with no apparent change in deep layer (or superficial layer, not shown) LFP frequency content.

increase blood pressure), and buprenorphine (analgesia). Atropine and tropicamide are both anticholinergics, and buprenorphine is known to have anticholinergic-like effects (Wood and Rackham, 1981; Sakuraba et al., 2009). Transitions to greater amounts of synchronization were also found in individual recordings within some tracks soon after application of a suspect drug (not shown). Subsequent recordings tended to remain more synchronized. Cats ptc17 and ptc18 had no pancuronium bromide (local retrobulbar injection of α-BTX was used instead of systemic paralysis, see Table 2.1), and their four tracks had the greatest amount of desynchronization (Figure 6.11). Note that pancuronium bromide is also an anticholinergic drug (Garland et al., 1998). Finally, higher doses of isoflurane and resulting greater depth of anesthesia were also associated with more synchronized cortical states. Slight changes in isoflurane dosage ($\pm 0.25\%$) sometimes resulted in a shift in cortical state (not shown). The two animals that used propofol & fentanyl instead of isoflurane & N_2O, ptc20 and ptc21, spent almost all of their time in the synchronized state (Figure 6.11).

To investigate if a relationship existed between overall firing rates and cortical state, MUA was calculated by combining the spike trains of specific subsets of active cells, and then binning them at the same time resolution as the SI (again 30 s wide bins, but now at 10 s resolution instead of 5 s). MUA was normalized by the number of included active cells. This was done separately for all spontaneous, natural scene movie, and artificial stimulus recordings (m-sequence white noise, drifting bars and gratings, and flashed gratings), as well as for fast, slow, simple, and complex cell types. Additionally, in view of a report that superficial and deep layer cells respond differently to changes in cortical state (Sakata and Harris, 2012), cells were roughly divided into superficial, middle, and deep layers according to their position along the length of the polytrode. Adjustments were made according to how transcolumnar each track was estimated to be. The upper/middle and middle/lower boundaries used for tracks ptc15.tr7c, ptc22.tr1, and ptc22.tr2 were 900 & 1100 μm, 500 & 700 μm, and 550 & 700 μm respectively. To reduce errors as a result of this rough

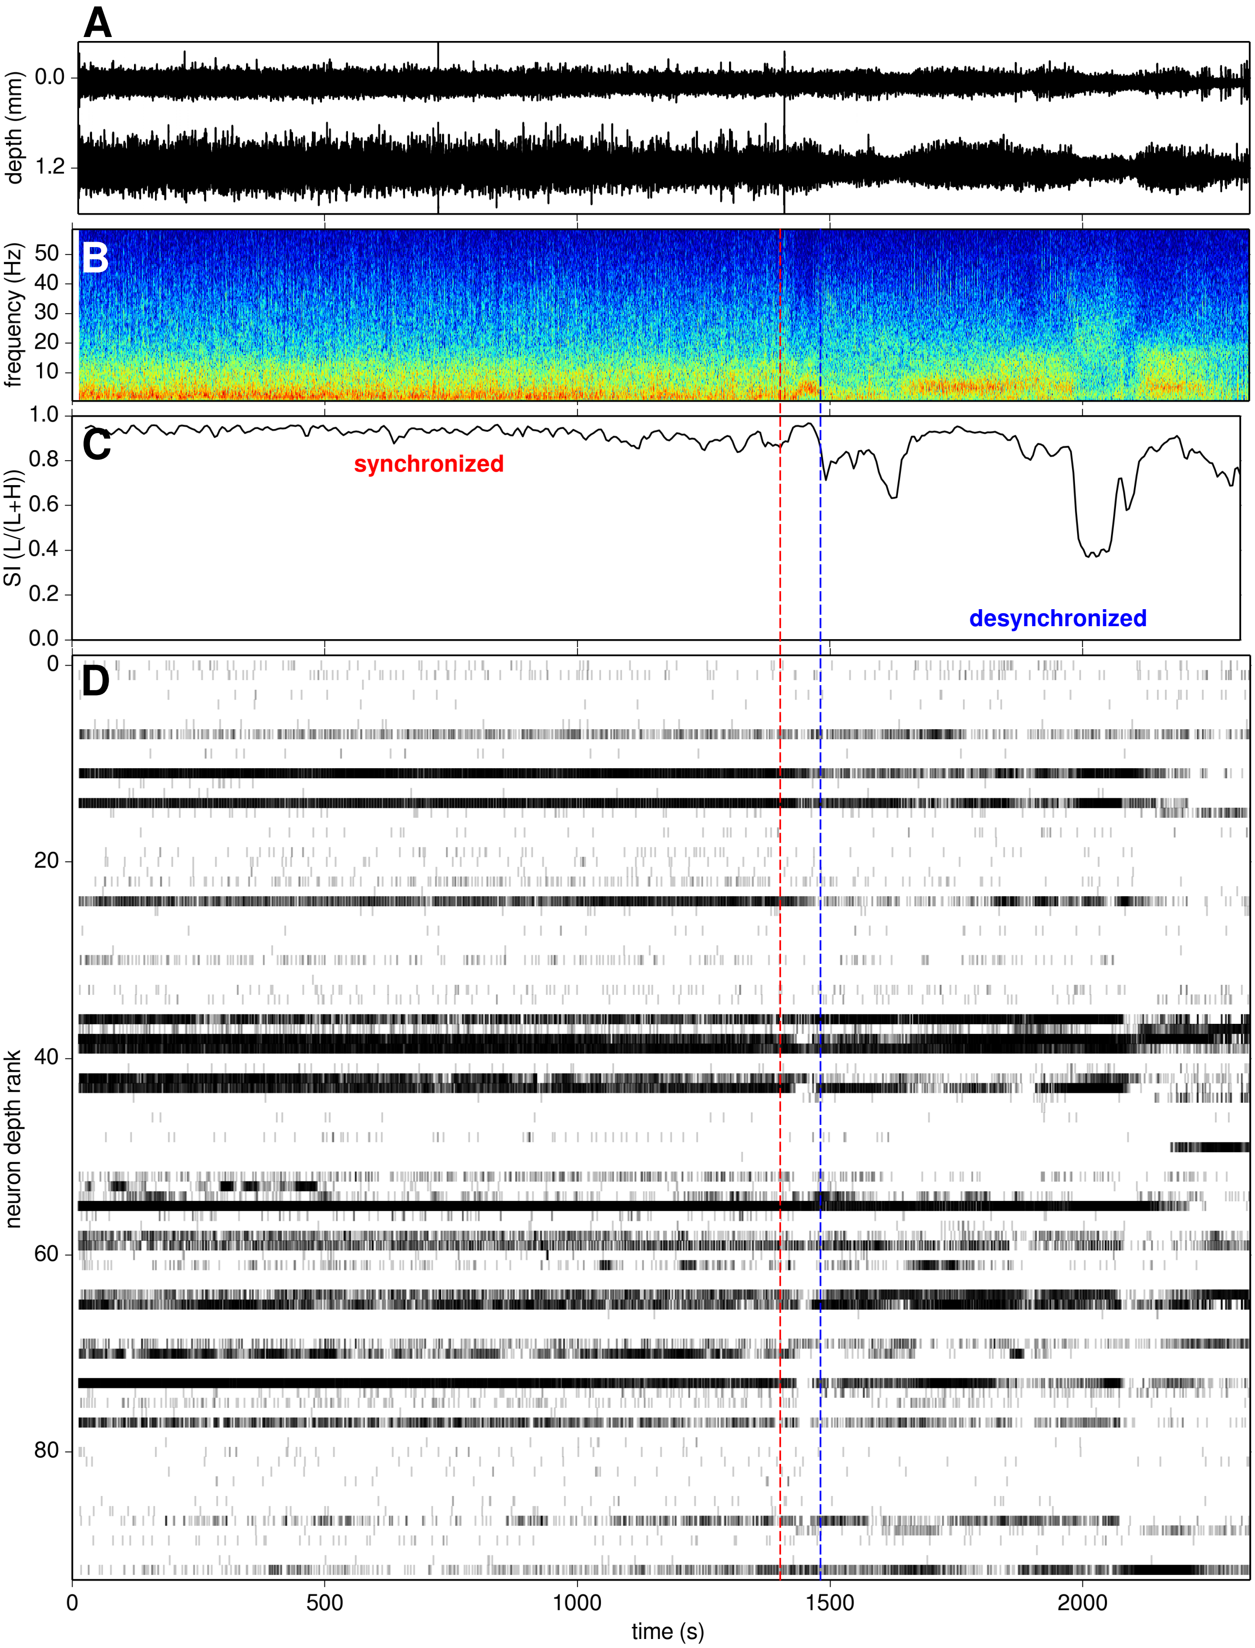
A
depth (mm)
0.0
1.2
B
frequency (Hz)
50
40
30
20
10
C
SI (L/(L+H))
1.0
0.8
0.6
0.4
0.2
0.0
synchronized
desynchronized
D
neuron depth rank
0
20
40
60
80
0
500
1000
1500
2000
time (s)

Figure 6.10: (*Previous page.*) Same as Figure 6.9, but for a different 4.5 s natural scene movie clip, separated from the previous one by $\sim$ 20 min of blank grey screen stimulus. During this recording, the opposite spontaneous change in cortical state occurred: from synchronized back to desynchronized. Neurons in the rows in ***D*** correspond to those in Figure 6.9.

delineation, the middle layer cells were excluded from further analysis.

Scatter plots and least-squares linear regressions of MUA vs. SI are shown in Figure 6.12. Each point represents 30 s of recording time at 10 s resolution. Insignificant relationships are shown faded out. The relationship between SI and MUA was complex, and varied depending on stimulus type, cell type, and layer. Correlation coefficients were low ($|r| < 0.5$), and conclusions were complicated by a shortage of desynchronized periods (low SI) and seemingly nonlinear relationships between MUA and SI. However, for most combinations of stimulus and cell type, MUA and SI had noticeably different relationships for superficial and deep layer cells.

Across stimulus type, layer type, and cortical state, fast cells had higher firing rates than slow cells (Figure 6.12, regression lines in 2 leftmost columns). Given that fast-spiking interneurons are known to have higher firing rates than pyramidal cells, this result adds confidence to the hypothesis that fast and slow cells, as defined here by their spike shape, correspond to fast-spiking interneurons and pyramidal cells, respectively.

Overall, regardless of cortical state or cell type, superficial layer firing rates were lower during spontaneous activity (Figure 6.12, red lines, top row) than during stimulus-evoked activity (middle rows). This was not the case for deep layers (blue lines), which sometimes had higher firing rates during spontaneous activity than during stimulus-evoked activity, especially during the synchronized state (high SI). Across cell types, spontaneously active superficial and deep layer cells had $r > 0$, i.e., on average, spontaneously active cells were more active in the synchronized than desynchronized state (red & blue lines, top row). For superficial layer cells, this relationship was reversed for both natural scene movies and artificial stimuli, which both had $r < 0$ across almost all cell types (red lines, middle rows, 5/6 significant combinations). Superficial layer slow cells during natural scene movies were the exception. Complementary to superficial layers, among the significant deep layer regressions (blue lines), there was greater variety across cell type, and less variety across stimulus type. Across stimulus type, deep layer simple cells had $r < 0$, while deep layer complex cells had $r > 0$. Also across stimulus type, deep layer fast and slow cells both had $r > 0$.

6.5 UP/DOWN phases

UP and DOWN phases are short alternating periods (< 1 Hz) of membrane potential depolarization and hyperpolarization in cortex (Steriade et al., 1993a; Anderson et al., 2000; Sanchez-Vives and McCormick, 2000), and are present during the synchronized cortical state, but absent during the

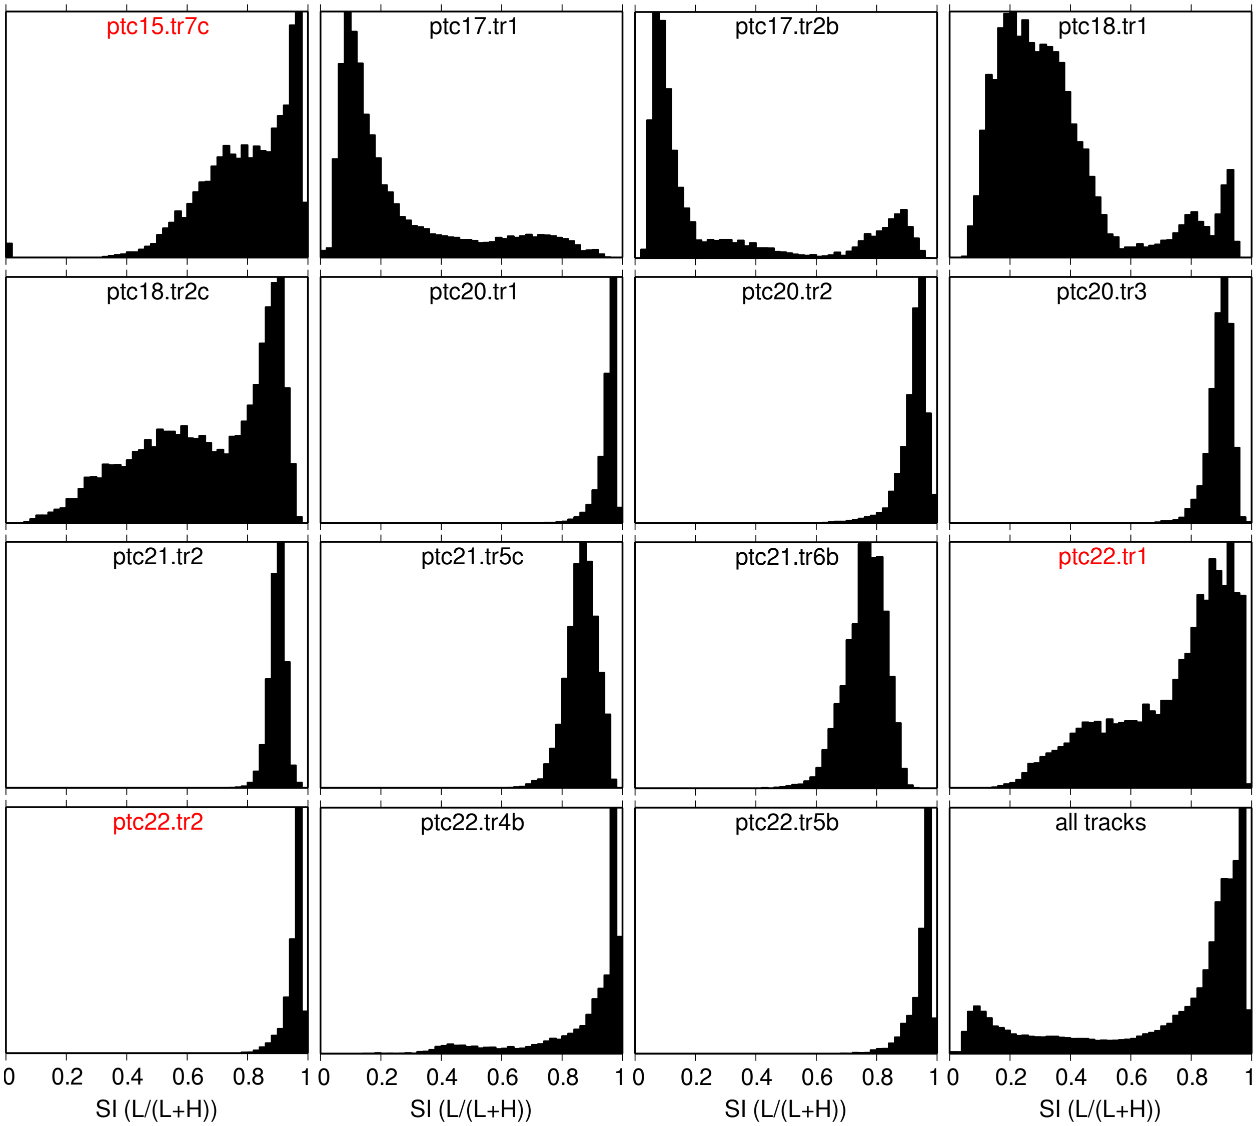

Figure 6.11: SI distributions of all 15 tracks in Table 2.2 (only 3 of which were spike sorted, labelled in red). The bottom right panel shows the distribution over all tracks, representing 137 hours of recording. Cortical state distributions varied widely across tracks, even within the same animal. However, for the sorted tracks, cortical state was mostly synchronized.

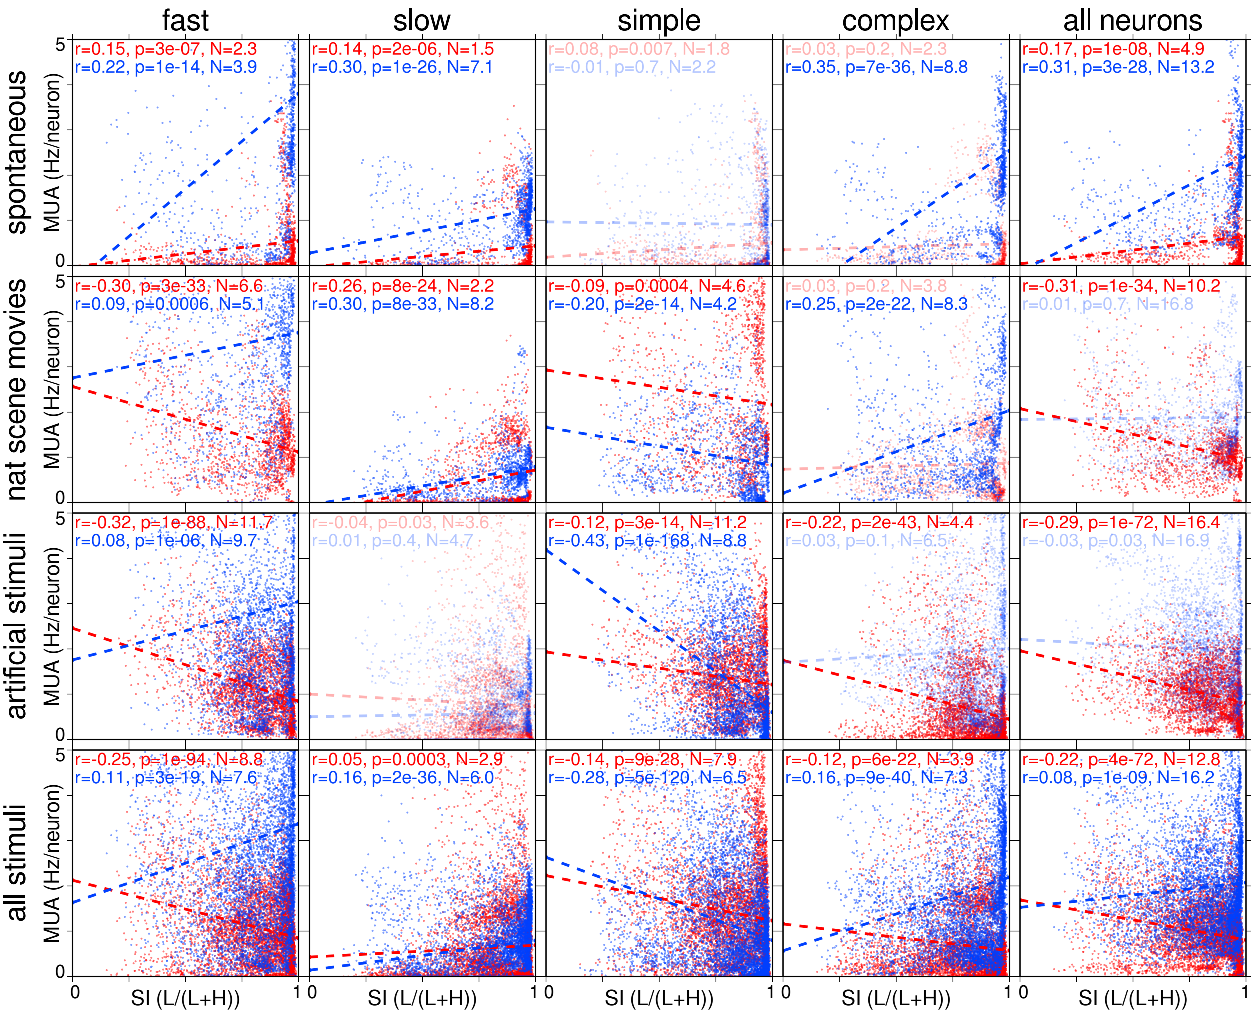

Figure 6.12: Scatter plots of MUA vs. SI for active neurons (activity determined separately for each relevant recording), classified by stimulus type (***rows***), cell type (***columns***), and layer type (***red***: superficial; ***blue***: deep). Artificial stimuli included m-sequence white noise movies, drifting bars and gratings, and flashed gratings. MUA and SI were both calculated using the same 30 s wide overlapping time bins at 10 s resolution. Each point represents one time bin. MUA was calculated separately by layer, while SI was calculated only from deep layer LFP. Least-squares linear regression was performed for each of the 24 classifications (3 stimulus types × 4 cell types × 2 layers) plus 16 more in the 'all' row and column. Regression r and p values are shown. N is the mean number of active cells per point from which MUA was calculated (note that this is less than the total number of unique cells that contributed to the respective regression). Lines, points, and statistics were shown faded if r was not significantly different from 0 (two-sided Student's t-test, $p < 0.001$). SI was mostly dominated by high values (see also Figure 6.11). Correlation coefficients were low ($|r| < 0.5$), but for most combinations of stimulus and cell type, significant regression lines differed noticeably between superficial and deep layer cells. Across cell types, superficial cells during spontaneous periods had low MUA and $r > 0$ (top row), and during stimulus-evoked periods had higher MUA and $r < 0$ (middle rows). Deep layer cells were more heterogenous across stimulus and cell type, and generally had higher rates than superficial cells.

desynchronized cortical state (Saleem et al., 2010; Harris and Thiele, 2011; Sakata and Harris, 2012). In fact, synchronized and desynchronized states are sometimes defined by the presence or absence of UP and DOWN phases, as measured by the coefficient of variation (standard deviation divided by the mean) of the LFP broadband power (Okun et al., 2012) or of MUA (Renart et al., 2010) as a function of time.

To better resolve UP and DOWN phases, the LFP spectrogram was constructed using finer overlapping time bins (1 s wide at 0.25 s resolution) than those used to calculate SI (Section 6.4). Periods of higher SI showed regular increases and decreases in broadband LFP power, consistent with UP/DOWN phases (Figure 6.13A, see also synchronized periods in Figures 6.9B & 6.10B). Population bursts of spiking activity (Figure 6.13C) resulted in peaks in the MUA (Figure 6.13D) which coincided with peaks in broadband LFP power. During spontaneous activity, a full cycle of one UP and one DOWN phase typically took ~ 2.5 s (Figure 6.13). UP phases had MUA peaks (Figure 6.13D) ranging 0.25–1 s in width, as measured by FWHM. DOWN phases lasted somewhat longer, at 0.5–2.5 s each. UP/DOWN phases were also present during stimulus-evoked activity, but were difficult to distinguish from responses to frequencies intrinsic to the stimuli, and were therefore left unexplored.

6.6 Natural scene movie responses vs. cortical state

How might cortical state influence natural scene movie responses? Figure 6.14 shows trial raster plots from 3 example cells to 2 different natural scene movie clips, during which two cortical state transitions occurred: from desynchronized to synchronized (Figure 6.14A–D), and from synchronized to desynchronized (Figure 6.14E–H). Examination of the trial raster plots showed that there was greater temporal precision and reliability in all 3 cells during the synchronized state than during the desynchronized state. Trial raster plots of all responsive cells were visually inspected during both of the natural scene movie clips in Figure 6.14. For the first movie (panel A), of the 27 responsive cells, the temporal precision and reliability of response events increased for 22 cells (82%), decreased for 3 cells (11%), and remained unchanged for 2 cells (7%). For the second movie (panel E), the temporal precision and reliability of response events decreased for all 31 responsive cells. Overall, for these two natural scene movie recordings (totalling 78 min of recording and separated by 20 min of blank grey screen), the vast majority of responsive cell firing pattern transitions (53/58, 91%) were consistent with greater temporal precision and reliability of response events during the synchronized state than during the desynchronized state. Only 5% (3/58) of responsive cell firing pattern transitions showed the opposite, while 3% (2/58) showed no change as a function of cortical state.

Although the 3 example cells shown in Figure 6.14 were responsive to both natural scene movie clips, some cells were responsive to only one movie and not the other. Figures 6.15 & 6.16 show 6 such example cells. Across the two natural scene movie recordings shown in Figures 6.14 &

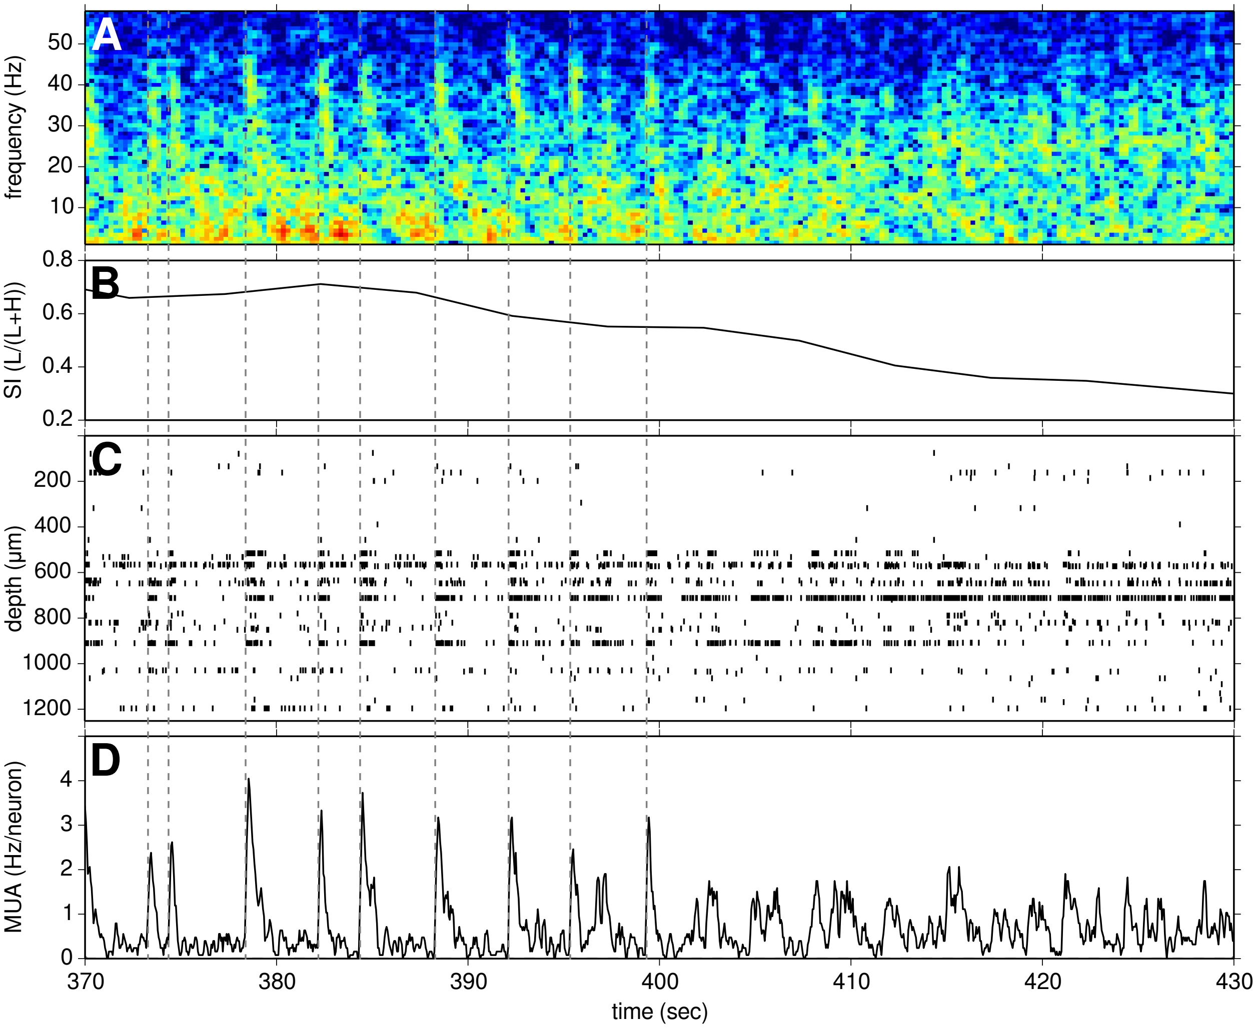

Figure 6.13: UP & DOWN phases are visible in spiking activity when cortical state is sufficiently synchronized. 60 s of example spontaneous activity (blank grey screen at 200 Hz refresh rate) is shown from ptc22.tr1. The deep channel LFP spectrogram (***A***) and resulting synchrony index (***B***) show that the local population became increasingly desynchronized over the time period shown. ***C***: Spatial population raster plot. Vertical axis is cell depth along the length of the polytrode, and spike rasters of cells at very similar depths necessarily overlap. 31 out of 93 neurons were active (with mean firing rates ≥ 0.05 Hz) during this recording. 63 cells fired at least one spike and were therefore used to normalize MUA in ***D***. Superficial layer cells were mostly silent. ***D***: Normalized MUA calculated from ***C*** using 200 ms overlapping time bins at 50 ms resolution. Vertical dashed lines denote the beginning of UP phases, as seen in peaks in the MUA, synchronized spiking in the spatial population raster plot, and broadband power peaks in the LFP spectrogram.

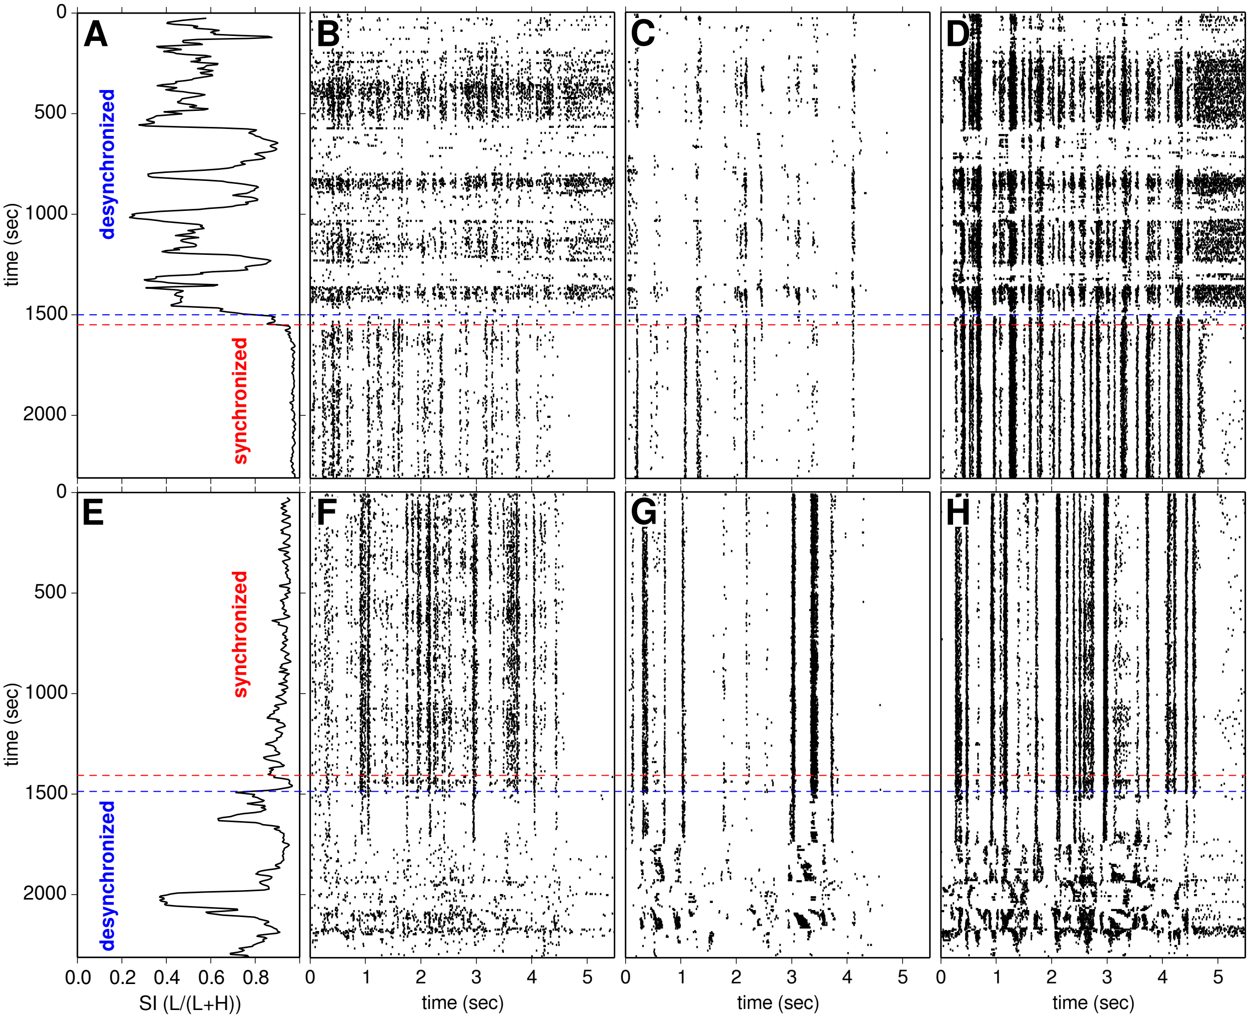

Figure 6.14: Cortical state affects precision and reliability of natural scene movie response events. During repeated presentation of two different 4.5 s natural scene movie clips in ptc22.tr1, two spontaneous cortical state transitions occurred: from desynchronized to synchronized (***A***, same recording as in Figure 6.9), and from synchronized back to desynchronized (***E***, same recording as in Figure 6.10). Horizontal dashed lines indicate transitions. SI is the L/(L+H) power ratio from the deep channel LFP. ***B–D, F–H***: Trial raster plots of natural scene movie responses of three example cells, left to right in order of increasing depth along the polytrode (161, 186 and 820 μm). Cell RF types (Section 5.4) were classified as simple, simple, and LGN afferent, and spike types (Section 5.2) were all fast. Each raster plot consisted of 400 presentations of a 4.5 s movie clip, each presentation separated by 1 s of blank screen. For both recordings, responses were more precise and reliable during the synchronized state than the desynchronized state. There was a gap of $\sim$ 20 minutes of blank grey screen separating the end of the first recording (***A***) from the start of the second (***E***). Responses were distinct for each cell, even for the first two whose physical separation was only $\sim$ 25 μm.

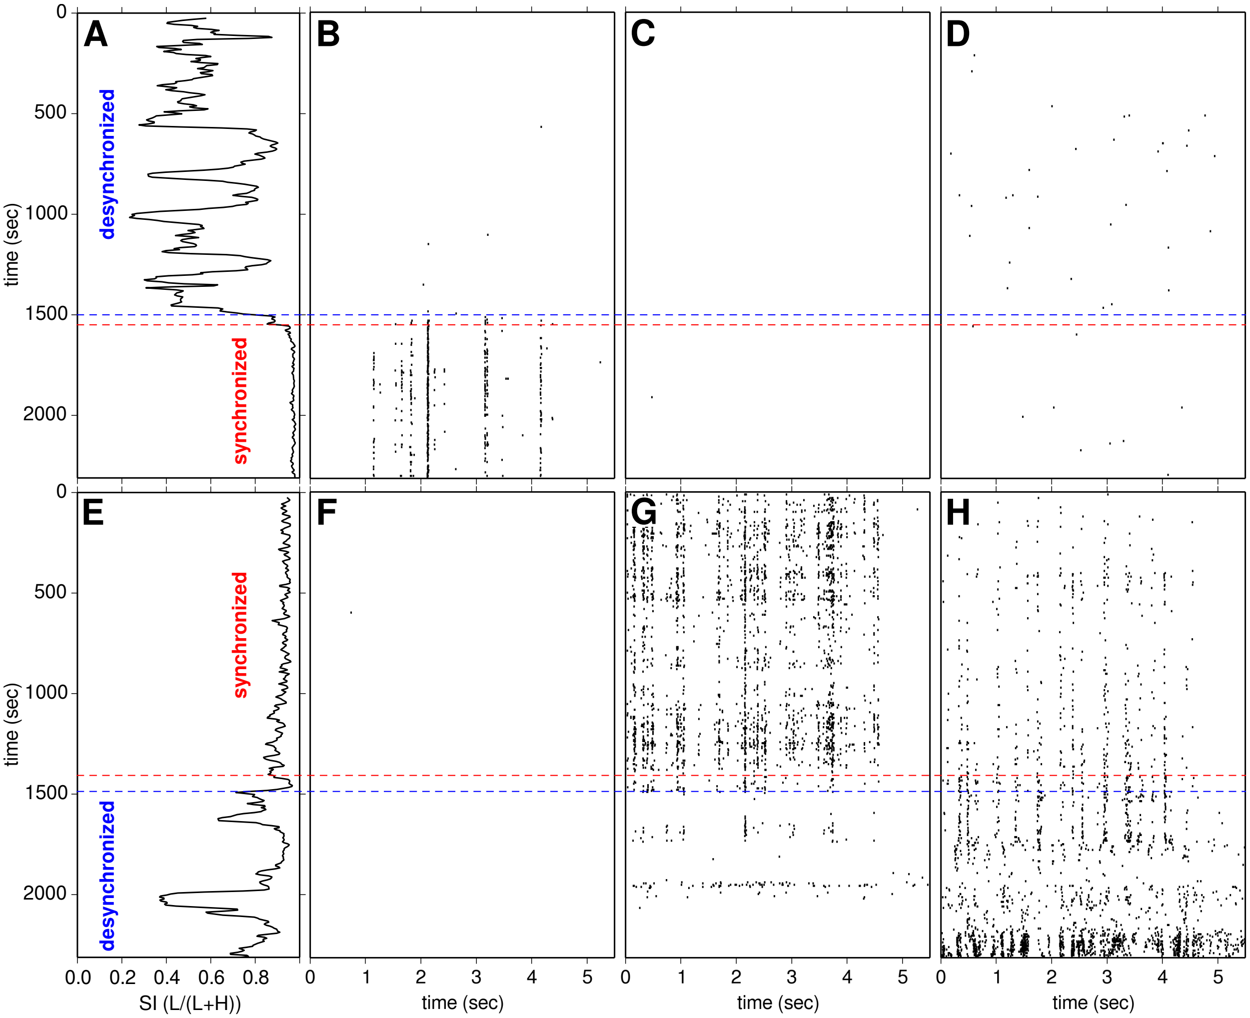

Figure 6.15: Same as Figure 6.14 but with 3 more example neurons, each of which had temporally precise and reliable response events during one movie but not the other. Panels ***C*** & ***F*** had exactly one spike each. Left to right, cells were in order of increasing depth along the polytrode (77, 974 and 1197 μm). Cell RF types were classified as unknown, unknown and LGN afferent, and spike types were fast, fast and slow. Although difficult to see in this layout, the last two cells shared several response events in the second recording. Note that the most superficial cell was most strongly modulated by cortical state, at least during the movie to which it responded (***B***). This was part of a trend (see text), in which responsive superficial layer cells were more likely than responsive deep layer cells to have a complete absence of response events in the desynchronized state, distinguishing them from the types of cells in Figure 6.14, which had response events in both states.

6.15, there were 20 cells that responded to only one movie, out of a total of 39 responsive cells: 8 responded only to the first movie, and 12 responded only to the second. However, as shown in Table 6.1, most cells isolated in that track (54/93) did not respond to either movie.

Another set of cells were those that were responsive in one cortical state but nonresponsive in the other (Figures 6.15B & 6.16). All such cells were responsive during the synchronized state in at least one movie, and nonresponsive during the desynchronized state in both movies. In addition, these cells were more likely to be superficial than deep layer cells, as suggested by the examples in Figures 6.15B & 6.16, which ranged 77–262 μm in depth along the polytrode. The layer boundaries described in Section 6.4 were again applied, and only those cells that were responsive at some point were included. Of such cells, for the first movie 7/12 (58%) superficial and 5/11 (45%) deep layer cells switched from nonresponsive to responsive. During the second movie, 7/9 (78%) superficial and 5/17 (29%) deep layer cells switched from responsive to nonresponsive. There were no cells during either movie whose responsivity switched in the opposite direction. The cells missing from the above fractions were those that still had discernable response events in the desynchronized state.

To quantify the change in temporal precision of response events between cortical states in the two movies shown in Figures 6.14 & 6.15, a peak detection algorithm was applied to the PSTH of every cell, and peak widths and heights were measured separately for all 4 recording periods (2 synchronized and 2 desynchronized). PSTHs were normalized by the number of trials in each recording period, and by the 20 ms bin width, yielding a signal in units of average instantaneous firing rate. For each PSTH, twice the median signal was designated the baseline level, and the peak detection threshold was set to 3 Hz above the baseline. Candidate peaks were those where the local maximum of the PSTH exceeded threshold and then fell below baseline on both sides. For each peak, a search was performed for the left and right FWHM timepoints on either side of the peak, from the baseline crossing on either side of the peak inward towards the peak. If the measured FWHM was greater than 200 ms, the peak was discarded. A range of peak detection parameters were tested, and results were generally insensitive to the particular parameters chosen. PSTHs were spot-checked to ensure that peaks were adequately detected and measured. Automated peak detection is demonstrated for 3 cells in Figure 6.16.

Out of a total possible 372 PSTHs (4 recording periods × 93 neurons in track ptc22.tr1), 283 PSTHs had at least one spike, 120 were active, and 79 were responsive (i.e., had at least one detected peak). There were significantly more PSTH peaks detected in the synchronized state than in the desynchronized state: 386 and 143, respectively (χ^2 test, $p < 5 \times 10^{-26}$). Logarithmic distributions of peak widths are shown in Figure 6.17A, coloured red for the synchronized state and blue for the desynchronized state. Peaks in the synchronized state were significantly narrower than in the desynchronized state, with FWHM geometric means of 36 and 51 ms, respectively (one-sided Mann-Whitney U test, $p < 7 \times 10^{-13}$). Peak amplitudes were also significantly greater in the synchronized state than in the desynchronized state (Figure 6.17B), with geometric means

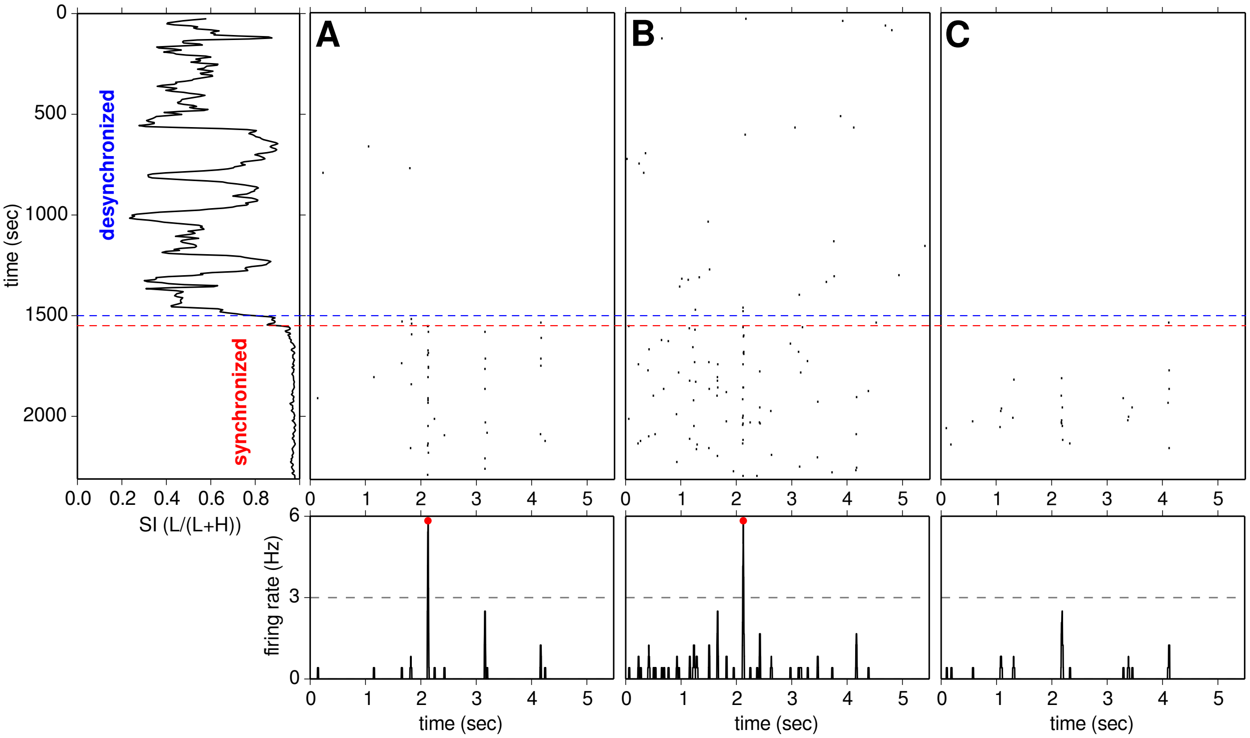

Figure 6.16: Responsive inactive cells. ***Top***: Same as the upper panels in Figures 6.14 & 6.15 but with 3 more example neurons, each of which was responsive during the synchronized state, but inactive overall. Cells ***A***–***C*** were in order of increasing depth along the polytrode (82, 249 and 262 μm). Mean firing rates during this recording were 0.02, 0.047, and 0.01 Hz, respectively. Cell RF types were all classified as simple, and spike types were all classified as fast. None of the 3 cells were responsive or active in the subsequent natural scene movie recording. ***Bottom***: All 3 PSTHs had a 0 Hz baseline firing rate, and therefore all had the same peak detection threshold of 3 Hz (grey horizontal dashed line). Red dots denote detected peaks. The last cell (***C***) had no PSTH peaks that exceeded the automated detection threshold, but for the purposes of Table 6.1 and Figure 6.3 was still classified by visual inspection as responsive.

of 11.2 and 7.5 Hz, respectively ($p < 2 \times 10^{-5}$).

To quantify changes in the sparseness of responses, the sparseness measure of Vinje and Gallant (2000) was applied to each responsive PSTH in each of the 4 recording periods. Sparseness was defined by

$$S = \left(1 - \frac{\left(\sum_{i=1}^{n} r_i/n\right)^2}{\sum_{i=1}^{n} r_i^2/n}\right)\left(\frac{1}{1-1/n}\right) \qquad (6.2)$$

where r_i is the PSTH value in the i^{th} time bin, and n is the number of time bins. Sparseness ranges from 0 to 1, with 0 corresponding to a uniform signal, and 1 corresponding to a signal with all of

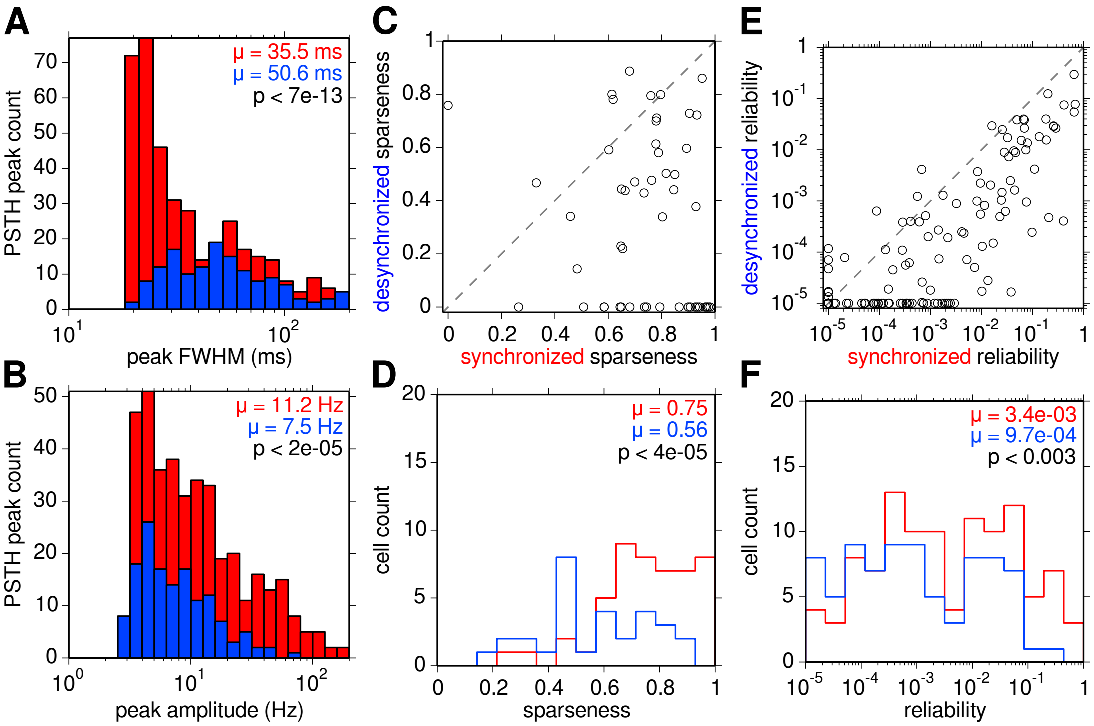

Figure 6.17: Response precision, sparseness & reliability vs. cortical state. ***A***: Distributions of response event widths, measured by FWHM of PSTH peaks, during the synchronized (***red***) and desynchronized (***blue***) periods of the two recordings shown in Figures 6.14 & 6.15. ***B***: Distributions of peak amplitudes relative to baseline. ***C***: Scatter plot of response sparseness in the two cortical states for cells with at least one detected PSTH peak. Cells with no peaks in one of the two states were assigned a sparseness of 0 in that state. 86% of cells fell below the dashed $y = x$ line. ***D***: Sparseness distributions for cells with at least one detected peak. ***E***: Scatter plot of response reliability in the two cortical states for all cells. Cells with no spikes during a cortical state were assigned a reliability of 10^{-5} in that state. 89% of cells fell below the dashed $y = x$ line. ***F***: Response reliability distributions for all cells. Geometric means are shown in ***A***, ***B*** & ***F***, and arithmetic means in ***D***. PSTH peaks were significantly narrower and higher, PSTHs were significantly sparser, and responses were significantly more reliable in the synchronized than desynchronized state (one-sided Mann-Whitney U test, p values shown in each panel).

its energy in a single time bin.

There were significantly more responsive PSTHs in the synchronized state than in the desynchronized state, 49 and 30 respectively (χ^2 test, $p = 0.033$). Figure 6.17C shows a scatter plot of PSTH sparseness values for all responsive cells in the two states. Cells that were responsive in only one state in each synchronized/desynchronized pair of states were assigned a sparseness value of 0 in the nonresponsive state. 86% of cells fell below the $y = x$ line, showing that most cells had sparser responses in the synchronized than desynchronized state. The corresponding distributions of sparseness values are shown in Figure 6.17D. Responses were significantly sparser in the synchronized state than in the desynchronized state, with mean sparseness values of 0.75 and 0.56, respectively (one-sided Mann-Whitney U test, $p < 4 \times 10^{-5}$).

To quantify the reliability of responses across trials, the mean pairwise correlations between all trials were calculated for each cell during each state (Goard and Dan, 2009). Single trial spike trains were divided into 20 ms wide overlapping time bins at 0.1 ms resolution, and the number of spikes in each bin was counted, resulting in a matrix of integer values as a function of time, with one row per trial (this is a modification of the method shown in Figure B.1). Pearson's correlation (Equation 6.1) was calculated between all possible pairs of trials. For trial pairs in which one or both trials had no spikes, their correlation was set to 0. The reliability of each cell during each cortical state was defined as the mean of all of the pairwise correlations of the trials during that state. Alternative methods were tested, including taking the median instead of the mean, or the mean weighted by the number of spikes in each trial pair, but the results were similar to the simple mean. Response reliability could range from -1 to 1, but was mostly positive. For logarithmic plotting, values that fell below 10^{-5} were assigned a value of 10^{-5}.

Figure 6.17E shows a logarithmic scatter plot of response reliability for all cells in the two cortical states. Cells that fired no spikes in one state in each synchronized/desynchronized pair of states were assigned a reliability value of 10^{-5} in that state. 89% of cells fell below the $y = x$ line, showing that most cells responded more reliably in the synchronized than desynchronized state. The corresponding logarithmic reliability distributions are shown in Figure 6.17F. Responses were significantly more reliable in the synchronized than desynchronized state, with mean reliability values of 3.4×10^{-3} and 9.7×10^{-4} respectively (one-sided Mann-Whitney U test, $p < 0.003$).

For the recording shown in Figure 6.5, which came from a different track (ptc22.tr2), cortical state was synchronized the entire time as measured by SI (Figure 6.11, bottom left panel). Consistent with the pair of recordings shown in Figures 6.14 & 6.15, responsive cells generally remained responsive for the duration of the recording, with generally consistent precision and reliability. There were some exceptions however. Inspection of the LFP spectrogram showed changes in low-frequency power that SI was not able to resolve (not shown). These changes were complex, but corresponded in time with changes in the response events of some cells.

Although cortical state was examined in the transcolumnar track (ptc15.tr7c), there was again a lack of desynchronized periods during natural scene movie presentation. However, there was no evidence from that track to contradict the findings from ptc22.tr1 in Figures 6.14–6.17. Furthermore, using LFP at a specific depth to calculate SI in a transcolumnar track may create complications, due to the potential for differences in cortical state between columns (Katzner et al., 2009). Therefore little could be concluded about the relationship between measured cortical state and natural scene movie responses in track ptc15.tr7c.

Given that the temporal precision and reliability of response events was higher in the synchronized state than in the desynchronized state, how might cortical state affect PSTH correlations between cell pairs? To answer this question, the two natural scene recordings in Figures 6.14 &

6.15 were split, as before, into synchronized and desynchronized periods, and PSTHs were calculated separately for active cells in each of the four resulting periods. PSTH correlations were then calculated for all active pairs in each recording period. The results are shown in Figure 6.18, in the same format as Figure 6.6.

Qualitatively, PSTH correlation matrices appeared more similar within cortical state and across movies (Figure 6.18B vs. C, left column) than across cortical state and within movies (panels A vs. B, C vs. D). Mean PSTH correlations ($\bar{\rho}$, red vertical lines, middle column) were higher in the two synchronized periods ($\bar{\rho} = 0.13$) than in the two desynchronized periods ($\bar{\rho} = 0.075$–0.1). As shown for the natural scene movie in Figure 6.6, ρ was independent of cell pair separation, even when considering cortical states individually (Figure 6.18, right column).

To more directly compare the changes in PSTH correlations from one period to the next, $\Delta\rho$ matrices were calculated for neighbouring periods in Figures 6.14 & 6.15. The results are shown in Figure 6.19. Although little was concluded from the spatial patterns of the $\Delta\rho$ values, their means corresponded with the conclusions from Figure 6.18: mean changes in ρ were significantly greater than 0 during the transition from the desynchronized to synchronized state (B-A, $p = 0.0052$), not significantly different within the synchronized state across movies (C-B, $p = 0.63$), and significantly less than 0 during the transition from the synchronized to desynchronized state (D-C, $p = 4.8 \times 10^{-8}$). $\Delta\rho$ values showed no apparent dependence on cell pair separation.

6.7 Discussion

6.7.1 Natural scene movie responses

The majority of active cells (87%) showed temporally precise and reliable response events to natural scene movie stimulation. Response events were as little as 20 ms wide (FWHM). There was great diversity in event patterns, even among cells very close to one another (Figure 6.1). Although only about half of all cells responded to natural scene movies, many of those that did not respond may have simply been inactive at the time, regardless of stimulus. One way to resolve this might be to tightly interleave trials of different stimulus types (including blank screen for spontaneous activity) to determine if some cells expressly do not respond to naturalistic stimuli.

There are a handful of reports of such sparse, temporally precise, and reliable responses to natural scene movies in V1: in awake behaving macaque (Vinje and Gallant, 2000, 2002), and in anesthetized cat, both extracellularly (Yen et al., 2007; Herikstad et al., 2011) and intracellularly (Haider et al., 2010). Yen et al. (2007) also showed that precisely timed events can be shared between nearby cells, as in Figures 6.4–6.5. Bair and Koch (1996) found similar precision and reliability in awake behaving macaque middle temporal cortex (MT) during random dot stimulation with low motion coherence. There have been more reports of even sparser, more temporally precise (as little as ~ 1 ms wide) and more reliable response events to high-entropy (though not necessarily

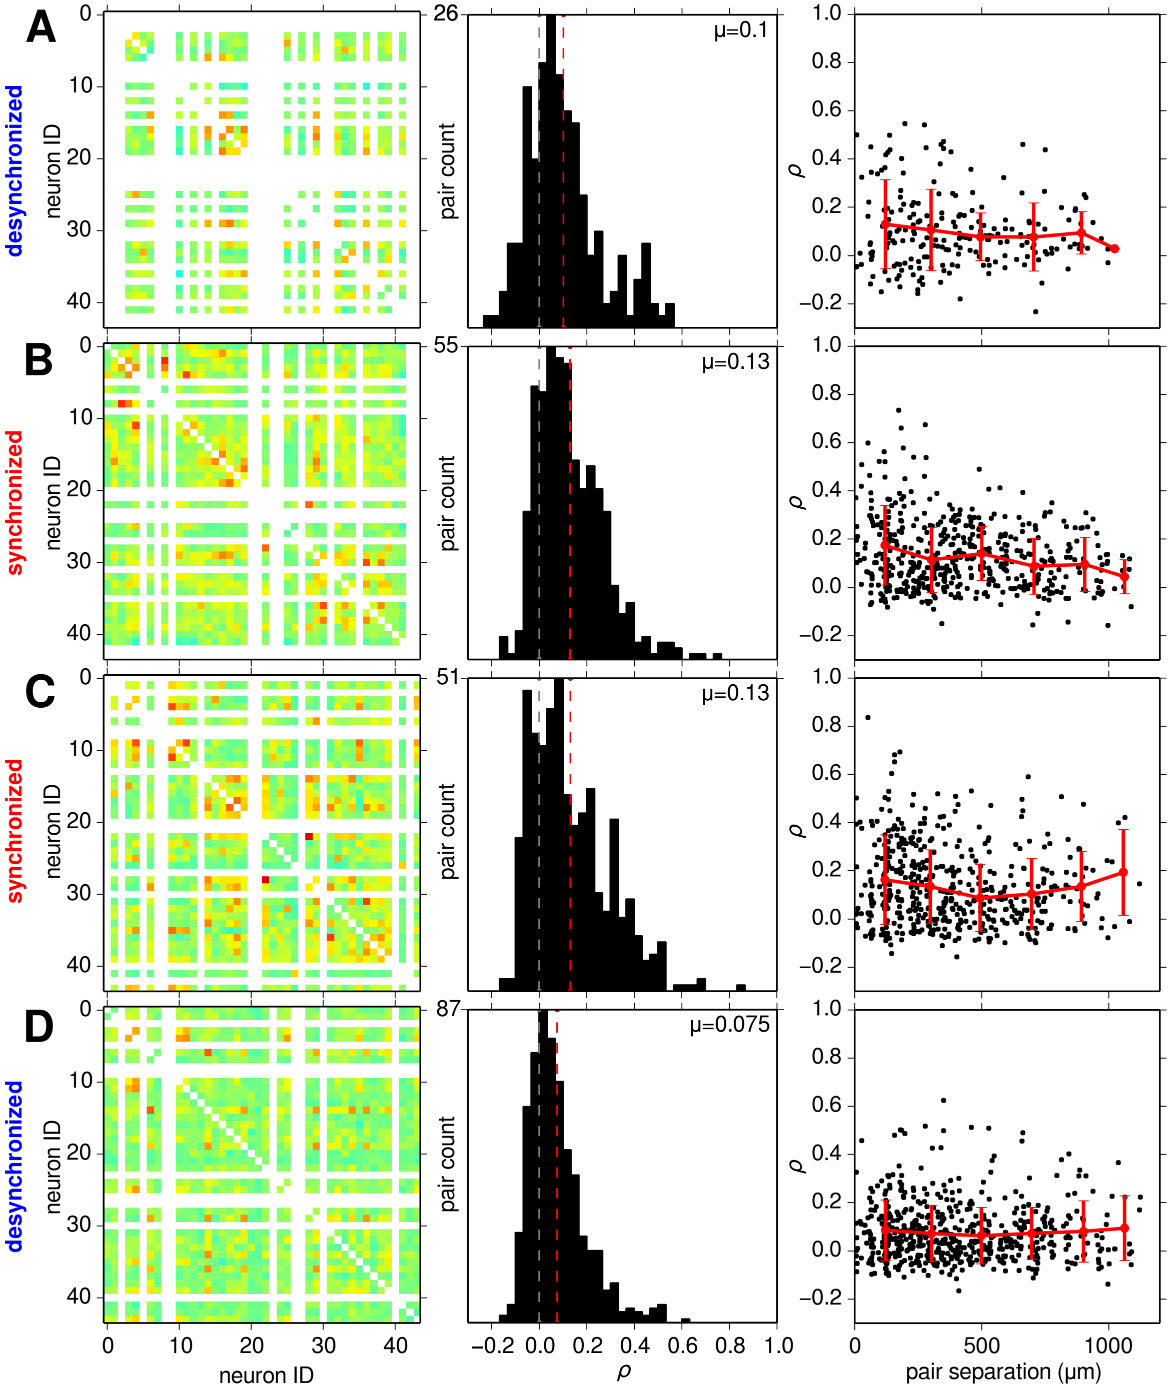

Figure 6.18: PSTH correlations vary as a function of cortical state. Same style of plots and same colour scale bar as in Figure 6.6A–B, but this time for the desynchronized (***A*** & ***D***) and synchronized (***B*** & ***C***) periods of the two recordings in Figure 6.14. ***A*** & ***B*** correspond to the recording in Figure 6.14A–D, while ***C*** & ***D*** correspond to the recording in Figure 6.14E–H. ***Left***: PSTH correlation matrices appeared more similar within state and across recordings (and movies) than within recordings and across state. ***Middle***: Synchronized periods had slightly higher mean PSTH correlations than desynchronized periods. ***Right***: PSTH correlations were independent of cell pair separation.

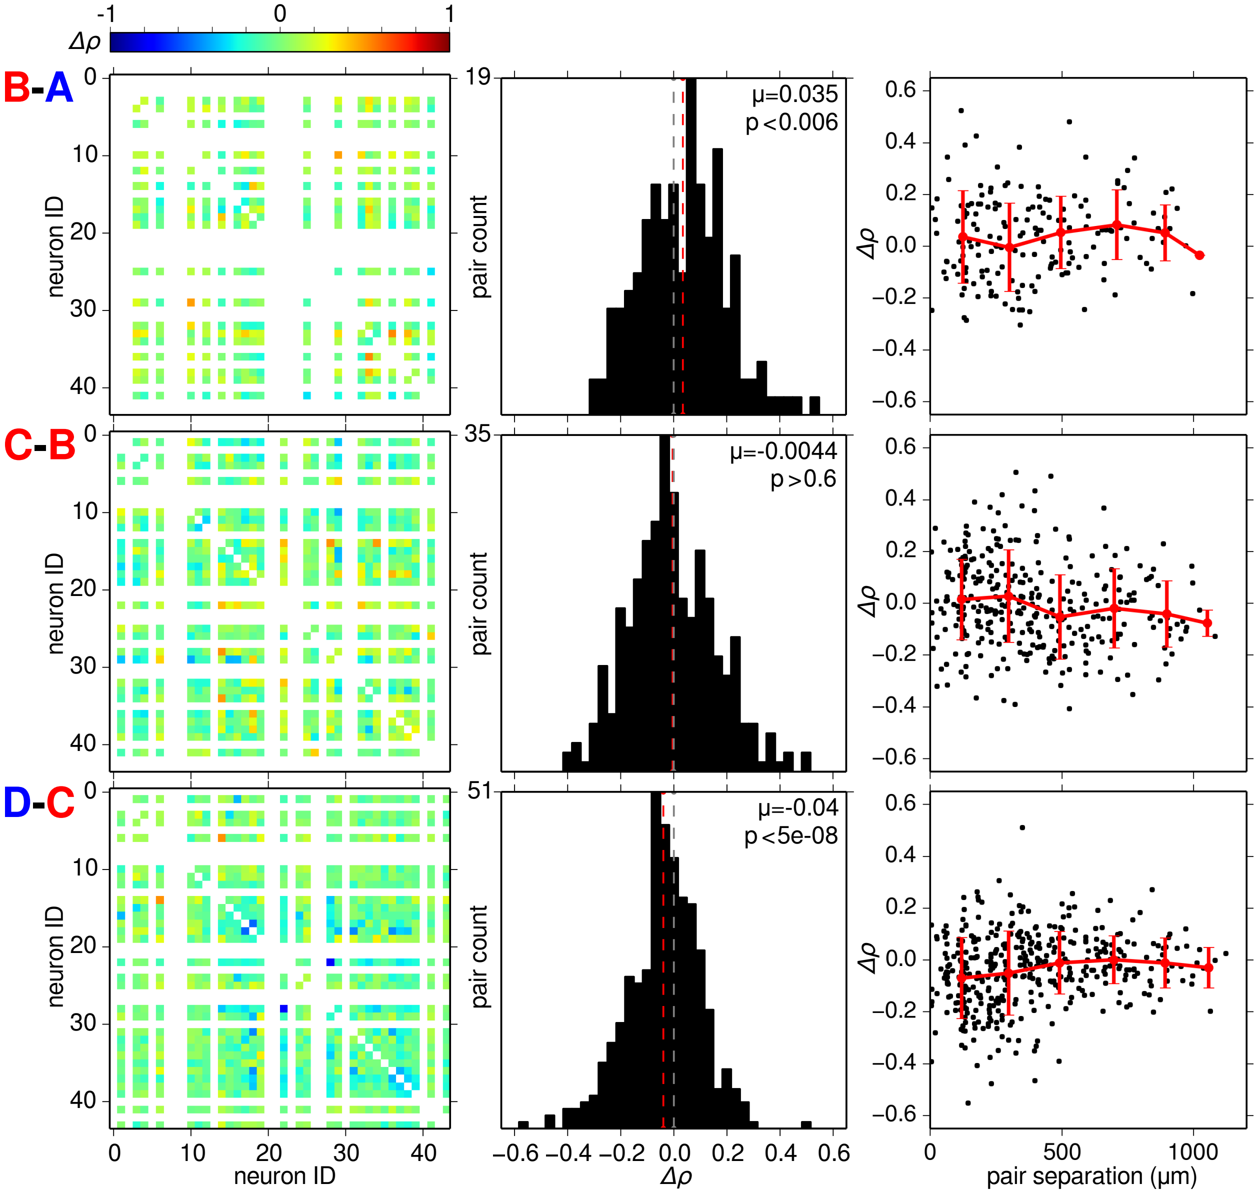

Figure 6.19: Cortical state has a greater influence on natural scene movie PSTH correlations than does the specific natural scene movie being presented. The three rows of panels here correspond to the differences in ρ values between the four recording periods labelled in Figure 6.18. For example, the panel labelled ***B-A*** represents the $\Delta\rho$ values between Figure 6.18B & A. Red and blue panel labels represent synchronized and desynchronized periods, respectively. Only cells active in both recordings were considered. ***C-B***: The mean $\Delta\rho$ for the two synchronized recording periods was not significantly different from zero (middle panel, vertical red dotted line, two-sided Student's t-test), despite being from different movies. For ***B-A*** and ***D-C***, mean $\Delta\rho$ values were significantly positive and negative respectively, congruent with synchronized periods having greater mean ρ than desynchronized periods (Figure 6.18). ***Right***: $\Delta\rho$ was independent of cell pair separation.

naturalistic) stimuli in RGCs in salamander, rabbit, and cat (Berry et al., 1997; Reich et al., 1997; Gollisch and Meister, 2008), and in LGN in anesthetized cat (Dan et al., 1996; Alonso et al., 1996; Reich et al., 1997; Reinagel and Reid, 2000, 2002).

It seems that as visual information propagates from RGCs to LGN to V1, response event precision and reliability decrease (Kara et al., 2000). It is interesting to consider that this precision is retained at all. LGN inputs constitute only a small fraction of synapses onto (mostly layer 4) cortical cells, yet these inputs are very effective at driving the cortex (Ahmed et al., 1994; Binzegger et al., 2004). In addition to the high effectiveness of LGN-V1 synapses, convergent event-like input from LGN cells in response to naturalistic stimuli may be one reason for this strong drive (Alonso et al., 1996; Wang et al., 2010). Clearly, there must be some evolutionary benefit in maintaining, to some extent, these temporally precise response events in V1. Sparse coding, and the energy efficiency that comes with it (Olshausen and Field, 1996; Attwell and Laughlin, 2001; Lennie, 2003) may be one such reason. Another reason, given by Hopfield (1995), is that precise relatively-timed spikes allow for simple scale-invariant representations of stimuli, and are potentially a common theme used throughout cortex. The delay line coding theory presented in that paper is strengthened by evidence, here and elsewhere, that at least for naturalistic stimuli, cortical cells can have responses that are both temporally precise and reliable relative to the stimulus, and therefore also temporally precise and reliable relative to each other.

The natural scene movies used here spanned visual angles of 12.7° for the 'old' set of movies and 51.6° for the 'new' set (Figure 2.5). Even the smaller of these was several times the size of the largest classical RFs shown here. Extra-classical RF stimulation has been found to increase response sparseness, for both artificial and naturalistic stimuli (Vinje and Gallant, 2000, 2002; Yen et al., 2007; Haider et al., 2010; Herikstad et al., 2011). Haider et al. (2010) showed that for regular spiking cells, this is due to more extensive inhibitory barrages by spatially unselective fast-spiking interneurons. The firing of the excitatory cells is thereby sculpted into temporally precise and reliable events. The temporal precision and reliability of responses reported here almost certainly required such wide field stimulation.

An *in-vitro* study by Mainen and Sejnowski (1995) in pyramidal layer 5 cells found sub-millisecond spike precision and high ($\sim$ 95%) reliability across trials in response to naturalistic, yet identical, input current traces. This is much better than the finest $\sim$ 20 ms precision found here. The main difference is that cells *in-vivo* have much greater feedforward, feedback, and horizontal connectivity than do *in-vitro* slices. Cell types are also more heterogenous *in-vivo* than in selective *in-vitro* recordings. Mainen and Sejnowski (1995) also found that a more artificial step input current resulted in much poorer temporal precision. An analogous result may be argued for *in-vivo* responses to artificial visual stimuli (Figures 4.8 & 5.4).

Precise, reliable, and unique responses to natural scene movie clips of 47% of all sorted cells suggests a high quality of spike sorting for those cells, both from noise and from each other. Of course, this says nothing about the spike sorting of the other 53% of cells that, for whatever reason, did not respond to natural scene movies, or for cells that were responsive to some movies but not others. Also, since SI derives from the low-pass LFP data which is methodologically independent of the spike sorted high-pass data, the correspondence in time of sudden changes in both SI and cell firing patterns also increases confidence in the quality of the spike sorting.

Five cells in track ptc22.tr1 were responsive to natural scene movies despite having mean firing rates < 0.05 Hz (Table 6.1, three shown in Figure 6.16). At an average of less than one spike every 20 s, these cells demonstrate that, as for artificial oriented stimuli (Figures 4.8C & 4.9), cells with extremely low firing rates in V1 can still meaningfully represent naturalistic stimuli via sparse responses. These cells also serve as an example of what can be lost when applying firing rate thresholds for unit inclusion and exclusion, even if those thresholds are set very low. Firing rate thresholds are certainly necessary for many analyses, but they should be used as sparingly as possible.

Half of the responsive cells in Figures 6.14 & 6.15 (20/39, 51%) were responsive to one natural scene movie but not the other, even though both recordings occurred closely in time, and even during periods when cortical state was apparently the same. Although loss of unit isolation is always a possibility and is difficult to explicitly distinguish from a natural change in responsivity, the results presented here suggest the possibility that some cells in V1 are extremely selective to specific naturalistic stimuli, with nearly all or nothing responses (Figure 6.15), even when the low level stimulus statistics do not change (not shown). Such extreme selectivity challenges the standard model of simple and complex cells in V1 as fairly simple oriented spatiotemporal filters (Hubel and Wiesel, 1962; Carandini et al., 1997), and may partially explain why such models do so poorly at predicting responses to natural scene movies (David et al., 2004; Olshausen and Field, 2005; Carandini et al., 2005). Another possibility is that individual V1 cells, or perhaps even cell assemblies, are engaged in a kind of shift work (Figures 4.2C, 5.9B & D), and alternate between various computational and biophysical tasks besides encoding stimulus. Such tasks might include visualization, memory formation and recall, reward encoding, synaptic renormalization, network stability, low level cellular maintenance, and energy conservation. Assuming that most of these other tasks take longer than stimulus encoding of a short movie clip, it may be possible to disentangle their potential effects on spiking activity by presenting several different natural scene movie clips with their trials randomly interleaved. Reconstruction of raster plots for each movie would reveal if similar fractions of neurons as reported here are found to respond to only some specific clips but not others, despite their tight interleaving during presentation. Or it may instead show that cells simply become responsive or nonresponsive and then stay that way for an extended period of time, regardless of stimulus. Either result would be interesting.

6.7.2 Natural scene movie response correlations

On average, natural scene PSTH correlations in cat V1 were positive, but very weak ($\bar{\rho} < 0.15$, Figures 6.6, 6.7 & 6.18). This is congruent with experimental and theoretical findings of near-zero spike count correlations in macaque and rat cortex (Ecker et al., 2010; Renart et al., 2010). For the transcolumnar track, PSTH correlations decreased with cell pair separation, and then slightly increased again, as might be expected with orientation preference changing through 180° as a function of position along the length of the polytrode (Figure 6.6). Surprisingly, for the two columnar tracks, PSTH correlations were not dependent on cell separation along the length of the polytrode (Figure 6.18). Considering only cell pairs with strongly correlated responses (say $\rho > 0.4$), such pairs were mostly closely spaced in the transcolumnar track (Figure 6.6), but had a wide range of spacing in the columnar tracks (e.g., Figure 6.18). This distance independence within a functional cortical column suggests that cells that respond similarly to natural scene movies are fairly well distributed across layers within a cortical column, and seems incongruent with functional specialization of cortical layers. A stronger conclusion regarding this will certainly require more cell pairs in more columnar tracks in more animals, with histological verification of polytrode position and orientation.

A recent report found that the more similar the responses of pyramidal cell pairs in layer 2/3 mouse V1 to natural scene movies, the greater the probability of direct synaptic connectivity between the two cells (Ko et al., 2011). That study also concluded that most connected pyramidal cell pairs had bidirectional synaptic connectivity. That study used 2-photon calcium imaging to record responses *in vivo* followed by aligned *in vitro* multiple (up to 4) cell patch-clamp recordings with electrical stimulation to directly gauge synaptic connectivity. Natural scene movie response correlation may therefore be useful as an indirect measure of direct synaptic connectivity between pyramidal cells, and the natural scene movie PSTH correlation matrices presented in this chapter (Figures 6.6 & 6.18) may be reasonable approximations of the synaptic connectivity between pyramidal cells.

When broken down by cell type, the mean natural scene movie PSTH correlations of three cell type combinations were notably higher than all other combinations, but only for columnar recordings: fast-fast, simple-LGN afferent, and complex-complex (Figure 6.7). Assuming that fast cells are indeed fast-spiking inhibitory interneurons (Section 5.6.1), strong fast-fast PSTH correlations (significantly stronger than all other spike type combinations) within a column are consistent with a calcium imaging study in layer 2/3 mouse V1 that found neighbouring parvalbumin-expressing fast-spiking interneurons had more strongly correlated responses to naturalistic movies than did pyramidal cells (Hofer et al., 2011). Strong fast-fast PSTH correlations are also consistent with studies showing that neighbouring fast-spiking inhibitory interneurons are closely coupled via electrical gap junctions (Galarreta and Hestrin, 1999; Gibson et al., 1999), and that natural scene movie responses are sculpted by fast-spiking interneurons (Haider et al. (2010), see previous section). Simple-LGN afferent correlations were higher (though not significantly) than most other RF type combinations. This is consistent with the Hubel and Wiesel hierarchical model of simple cell RFs, and the findings of Reid and Alonso (1995). More interestingly, simple-complex correlations were weak, and complex-complex correlations were significantly stronger than all other RF type combinations (except simple-LGN afferent). This conflicts with the Hubel and Wiesel hierarchical model of complex cell RFs (Figure 6.20A), i.e., that they are the result of monosynaptic input from multiple simple cells.

Perhaps the best evidence for the hierarchical complex cell RF model is a study by Alonso and Martinez (1998). That study found that extracellular spike train cross-correlograms between layer 4 simple cells and layer 2/3 complex cells in anesthetized cat V1 showed monosynaptic simple-complex connectivity, but not vice versa. However, that study may have been biased in its use of a handful of electrodes that were independently positioned to maximize cell isolation, and which recorded no more than a pair of cells at a time. Those results conflict with earlier studies (Toyama et al., 1981a,b; Ghose et al., 1994) which showed, using similar methodology, a lack of simple-complex monosynaptic cross-correlogram peaks, but did show some “antihierarchical” complex-simple monosynaptic cross-correlogram peaks. A more recent intracellular study in anesthetized cat V1 by Yu and Ferster (2013) found that although complex cell membrane potentials triggered off of simple cell spikes did indeed show a significant peak during visual stimulation, that peak disappeared when visual stimulation was replaced with electrical stimulation, suggesting that coordinated activity is what drives complex cell responses, not direct monosynaptic connectivity from simple cells.

These conflicting results suggest that the hierarchical complex cell model may need refinement, as might the distinction between the two cell types (including that made in Section 5.4). These are two independent issues, and there is more support for the first than the second (Toyama et al., 1981a,b; Ghose et al., 1994; Mechler and Ringach, 2002; Abbott and Chance, 2002; Priebe et al., 2004; Mata and Ringach, 2005; Yu and Ferster, 2013). Alternative models (Chance et al., 1999; Tao et al., 2004) suggest that recurrent connectivity between cells in V1 could determine how simple or complex a cell is (Figure 6.20B). In these models, more complex-like cells have more recurrent (bidirectional) connectivity with other nearby cells with similar orientation and spatial frequency preferences, but different spatial phase preferences. These models predict a spectrum of recurrent connectivity from simple to complex, with strong coupling between highly complex cells and weak coupling between highly simple cells. The results shown in Figure 6.7 support these recurrent connectivity models of complex cells. Note that these models do not invalidate the bimodal distribution of spiking responses that have long distinguished simple and complex cells. Although membrane potential responses are unimodal across the population, the spike threshold nonlinearity transforms this unimodal distribution into a bimodal one which distinguishes simple and complex cells (Mechler and Ringach, 2002; Priebe et al., 2004). Therefore, although recurrent synaptic connectivity may vary gradually from simple to complex cells, spiking responses, which are what downstream targets receive, are indeed bimodal.

Unfortunately, traditional methods of distinguishing simple and complex cells rely on artificial stimuli: drifting bars and gratings, light and dark flashed bars and spots, flashed gratings, and counterphase gratings. Ideally, responses from natural scene movies alone should be enough to determine which cells are simple, which are complex, and which (if any) are somewhere in between.

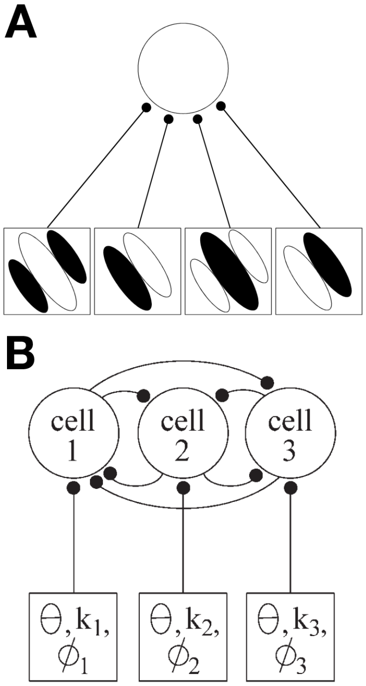

Figure 6.20: Complex cell models. ***A***: Hierarchical model. A complex cell (circle) pools inputs from simple cells (boxes) with similar orientation and spatial frequency preferences, but different spatial phase preferences, making it invariant to spatial phase. ***B***: Recurrent model. Each cortical cell (circles) could be either simple or complex, depending on the strengths of its recurrent connections with others. Each cell receives a weak set of feedforward inputs from LGN (boxes), which determines its orientation, spatial frequency, and spatial phase preference (θ, k, and ϕ respectively). In the absence of strong recurrent connections, a cell retains these preferences and behaves like a simple cell. With increasing recurrent connectivity with other cells of similar orientation and spatial frequency preference, but differing spatial phase preference, the cell becomes more invariant to spatial phase, and hence more like a complex cell. Taken from Chance et al. (1999).

As of yet, no method for doing so exists. Until such a method is devised and applied, the physiological relevance of the artificial stimulus-derived simple/complex cell distinction to natural vision should be considered an open question.

Given the ubiquity of temporally precise and reliable response events during natural scene movie stimulation (at least during the synchronized state), response events could be used as a tool to distinguish simple from complex cells. One method might be to present two versions of the same natural scene movie: the original, and the contrast inverted version. Since simple cells are sensitive to spatial phase, their response events should shift in time between the two movies, or rearrange completely. However, the response events of complex cells should not shift, because ideally complex cells are phase invariant. Increasing degrees of temporal shift of response events could signify increasing simple cell-like properties. It would be interesting to see if these temporal shifts clustered into two groups, or if instead they formed a continuous distribution.

6.7.3 Cortical states

Deep layer LFP spectrograms revealed spontaneous changes in frequency content (Figures 6.9 & 6.10), which roughly clustered into two states: synchronized (slow, large amplitude fluctuations) and desynchronized (fast, low amplitude fluctuations). LFP frequency content was quantified in a more graded way using an index (SI) which describes the degree of synchronization of the local population as a function of time. This index was then correlated with mean population firing rates (MUA) for a variety of cell, layer and stimulus types (Figure 6.12).

Across cell type, population firing rates during spontaneous activity in superficial cells were higher in the synchronized state and lower in the desynchronized state, while the reverse was true for stimulus-evoked activity (Figure 6.12). Deep layer cells showed a more heterogenous relationship across stimulus and cell type. These results agree with those of Sakata and Harris (2012). That study focused on spontaneous activity in awake and urethane-anesthetized rat A1, and found that for both awake and anesthetized animals, superficial layer cells had higher firing rates in the synchronized state than in the desynchronized state. That study also found that the relationship for deep layer cells was more heterogenous and depended on cell type: some pyramidal cell types had higher rates during desynchronization, but fast-spiking interneurons did not. Assuming that the slow and fast cell types described here correspond to pyramidal cells and fast-spiking interneurons respectively (Section 5.6.1), the cell type-specific results presented here for deep layer cells during spontaneous activity are consistent with that study (Figure 6.12, top row, first two columns, blue lines).

Although the laminar and cell type analysis performed here was less reliable than that of Sakata and Harris (2012), which used juxtacellular recording combined with histology and spike shape analysis, the results presented here expand on that study by including stimulus-evoked activity with both natural scene movies and artificial stimuli, in addition to spontaneous activity. The relative differences in stimulus-evoked firing rates of slow and fast cells in both superficial and deep layers in Figure 6.12 (rows 2–4, columns 1–2, red & blue lines, high SI) generally agree with those of an earlier study by the same group (Sakata and Harris (2009), their Figure 2A) which examined laminar and cell type dependence of spontaneous and stimulus-evoked activity during the synchronized state.

This correspondence with other studies in different species and modalities with different anesthetics suggests that the results in Figure 6.12 are not species, modality, or anesthesia specific. Rather, they may be general properties of primary cortical areas and perhaps other areas. What is their functional relevance? Sakata and Harris (2012) suggest that lower superficial firing rates in the desynchronized state during spontaneous activity may be a way of lowering neural noise, perhaps via attention at a specific retinotopic or tonotopic location, to allow detection of stimuli at lower threshold. Although attention is presumably absent during anesthesia, the broad mechanism of global synchronization and desynchronization remains, suggesting that aspects related to attention can be studied in an anesthetized animal (Harris and Thiele, 2011).

For the 3 spike sorted tracks, cortical state was mostly synchronized, and the desynchronized state was infrequent (Figure 6.11). This limited confidence in conclusions about how V1 differs between the two extremes of cortical state. Several steps could be taken to better balance, and even control, the time spent in both states. Since anticholinergic and anticholinergic-like drugs have a synchronizing effect on cortex (Herrero et al., 2008; Harris and Thiele, 2011), they should be used with greater care. Atropine, tropicamide, dobutamine, buprenorphine, and pancuronium bromide were all implicated here in promoting synchronization and preventing desynchronization. Some are essential, but reduced dosage or complete elimination of others may be possible. Although not shown here, sometimes a loud noise, such as a hand clap, was enough to desynchronize V1.

It may also be wise to avoid propofol & fentanyl anesthesia in the future and use isoflurane (and potentially N_2O) instead. The greater synchronization under propofol & fentanyl (ptc20 and ptc21, Figure 6.11) may have been due to the greater difficulty in regulating its intravenous infusion and its longer half life compared to inhaled isoflurane.

Slight changes in isoflurane dosage ($\pm 0.25\%$) sometimes affected cortical state (not shown). Long after completion of surgery and after many consecutive hours of isoflurane administration, it may be possible to temporarily reduce isoflurane dosage to very low levels while still maintaining sufficient anesthesia (Eger II and Johnson, 1987; Pascoe et al., 2006). Doing so should result in more desynchronized periods. Given the long recovery times from long duration anesthesia, it may even be possible to completely turn off isoflurane for short periods, perhaps 10 s at a time, to induce a shift from the synchronized to desynchronized state. Even greater control is possible by direct stimulation of the nucleus basalis (Goard and Dan, 2009) or the pedunculopontine tegmenta (PPT) nucleus (Curto et al., 2009; Sakata and Harris, 2012), at least for V1 and A1 respectively, in urethane-anesthetized rat. However, such stimulation involves greater complexity and risk, and requires further surgery. A combination of some of the above methods should be sufficient to ensure a balance between synchronized and desynchronized cortical states. Until then, spike sorting of existing data that already show greater desynchronization (ptc17 & ptc18, Figure 6.11) is a top priority.

A study by Wörgötter et al. (1998) found that RF size in ketamine- and halothane-anesthetized cat V1 is greater in the synchronized than desynchronized state. A similar report by Castro-Alamancos (2002) found that whisker fields in urethane-anesthetized rat barrel cortex are larger in the synchronized than desynchronized state. Both studies relied primarily on single channel electrodes, and their results could be tested using the data already collected here. The STAs of simple cells from m-sequence noise movies (Figure 5.5) could be separately calculated during synchronized and desynchronized states, and their sizes measured. Alternatively, RF size of both simple and complex cells during the two cortical states could be measured using light and dark drifting bars (Figure 5.4A). Dependence of RF size on cortical state could then be examined according to spike type, RF type, and layer.

6.7.4 UP/DOWN phases

Up and down phases are present only in the synchronized cortical state and are thought to represent two alternating modes of cortical function (Harris and Thiele, 2011). Up phases are brief periods of depolarized membrane potential and desynchronization during which information can flow more easily from one brain area to another (Luczak et al., 2007; Hoffman et al., 2007), akin to the temporary opening of a gate (Luczak et al., 2013), and reminiscent of the data packets and time domain multiplexing used in digital communication systems. Down phases are somewhat longer periods between up phases, with hyperpolarized membrane potential and high synchronization during which information flows more poorly.

As in many other species and cortical regions, in this study UP and DOWN phases were found in the synchronized cortical state but were absent in the desynchronized state (Figure 6.13). During spontaneous activity, one full up/down cycle lasted ~ 2.5 s on average, corresponding well to the ~ 0.3–0.4 Hz frequency of UP/DOWN phase cycling originally reported intracellularly *in vivo* in urethane-anesthetized cat cortex (Steriade et al., 1993a). UP phases lasted 0.25–1 s, and DOWN phases lasted somewhat longer at 0.5–2.5 s, similar to findings during slow wave sleep in cat cortex (Destexhe et al., 1999). UP phase duration in cat V1 was similar to that in urethane-anesthetized rat barrel and somatosensory cortices (Petersen et al., 2003; Hasenstaub et al., 2007; Luczak et al., 2007), but was longer than in awake rat (Petersen et al., 2003; Luczak et al., 2007; Sakata and Harris, 2009; Luczak et al., 2013), and DOWN phase duration was longer than in both anesthetized and awake rat, as well as ketamine/xylazine anesthetized ferret PFC (Haider et al., 2006).

Interestingly, Anderson et al. (2000) found that complex cells in anesthetized cat V1 had bigger amplitude UP and DOWN phases than did simple cells. Contrary to the overall finding here, they also reported that UP phases lasted longer than DOWN phases, especially during visual stimulation. That study recorded intracellularly from one cell at a time. Characterization of UP and DOWN phases according to RF, spike, or layer type was not attempted here, and has yet to be reported using extracellular polytrode recordings. Even though single unit spike trains on their own do not provide the same UP/DOWN phase detection fidelity as single cell intracellular recordings, with enough simultaneously recorded cells of each type, MUA pooled from those cells should reveal UP/DOWN phases specific to that type. Simultaneous recording might also reveal how UP/DOWN phases interact between different RF, spike, or layer types.

Luczak et al. (2007) found that spike latencies relative to UP phase onset were constant, and were different for each neuron. This allowed neurons to be ranked according to their UP phase spike latency during spontaneous activity, without the need for any external manipulation of trial start and end. Using an internal trigger for designating trials may be a more natural and less biased way of denoting trial start than an external trigger. It seems that UP phase spike latency has not yet been investigated in cat V1. Although methodologically difficult, future experiments might also attempt to time the start of stimulus presentation trials to UP phase onset, which might increase response precision and reliability.

UP phase onset may also serve as a valuable tool for determining polytrode laminar position. Sakata and Harris (2009) used UP phase onset as a trigger for CSD (Section 4.6.3) analysis, without the need for an external stimulus trigger. This means that experimental time need not be spent presenting a transiently synchronizing stimulus, such as a full screen flash or direct electrical thalamic stimulation, solely for the purpose of constructing the CSD. It would also provide many more triggers for calculating the CSD with greater fidelity, and allow for continuous tracking of laminar position. Again, using an internal trigger rather than an external one is arguably more naturalistic.

6.7.5 Natural scene movie responses vs. cortical state

Natural scene movie responses were more temporally precise and reliable during the synchronized state than the desynchronized state (Section 6.6). For the two natural scene movie recordings during which a spontaneous state change occurred (Figures 6.14 & 6.15) this conclusion held for 91% of responsive cells. Only 5% of responsive cells showed the opposite, while 3% showed no change. These qualitative differences in precision and reliability between the two states were supported by quantitative measures of response precision, sparseness, and reliability, all of which were significantly greater in the synchronized state than in the desynchronized state (Figure 6.17).

Although other studies have reported temporally precise and reliable responses to naturalistic stimuli in V1 (Vinje and Gallant, 2000, 2002; Yen et al., 2007; Haider et al., 2010; Herikstad et al., 2011), perhaps only one has considered the influence of cortical state on temporal precision and reliability in V1. Contrary to the results shown in Section 6.6, Goard and Dan (2009) reported that responses to natural scene movies in urethane-anesthetized rat V1 were *more* reliable in the desynchronized state than in the synchronized state. Desynchronization was triggered by electrical stimulation of the nucleus basalis in the basal forebrain. Another study by Marguet and Harris (2011) came to the same conclusion, using an amplitude-modulated frozen noise stimulus in urethane-anesthetized rat A1. In that study, desynchronization was induced via tail pinch. A third study by Zagha et al. (2013) reported that tactile stimulus encoding in urethane-anesthetized mouse barrel cortex was also more reliable in the desynchronized state, which was induced by stimulation of M1. In the above three studies, comparisons of trial raster plots or MUA PSTHs between the two cortical states suggest modest changes in temporal precision and reliability compared to the more dramatic and opposite changes shown here in Figures 6.14 & 6.15. A fourth study by Pachitariu et al. (2015) in variously anesthetized gerbil A1 also found greater precision and reliability of responses in the desynchronized state than in the synchronized state, this time to frequency-modulated tones and speech stimuli. Trial raster plot differences between states were more dramatic in that study than in the other three.

There are many experimental differences that might explain the opposing results presented here, including differences in species (cat vs. rodent), anesthetic (isoflurane vs. urethane, ketamine/xylazine, and fentanyl/medetomidine/midazolam), desynchronization method (spontaneous vs. evoked), cortical area (V1 vs. A1 and barrel cortex), stimulus modality (visual vs. auditory and tactile), and stimulus type (naturalistic vs. artificial). Since cortical state is likely multidimensional and SI measures only one such dimension (Harris and Thiele, 2011), it is also possible that there were other undetected changes in cortical state in the two recordings in Figures 6.14–6.15 but not in those reported in the literature (or vice versa). Such undetected changes might account for some of these opposing results.

The species difference may be the most important. Cats have greater columnar organization of stimulus features in V1 than do rodents: cats have ocular dominance and orientation columns that rodents lack (Horton and Adams, 2005). Up phases in the synchronized state can manifest as waves of activity travelling across the cortical surface (Petersen et al., 2003; Massimini et al., 2004; Benucci et al., 2007; Xu et al., 2007; Luczak et al., 2007; Mohajerani et al., 2010), while oriented visual stimuli can evoke standing waves of activity aligned to orientation columns (Benucci et al., 2007). Presumably, stimulus-evoked standing waves are absent in species that lack orientation columns, including rodents. Perhaps an interaction between these travelling and standing waves of activity in the synchronized state increases the temporal precision and reliability of stimulus-evoked responses in cat V1 relative to rat V1. This hypothesis predicts that responses in the synchronized state of anesthetized ferret and primate V1, which also have orientation columns, should also be more precise and reliable than in the desynchronized state.

These opposing results are surprising because neural responses in awake animals are known to be stronger and more synchronized in alert than in quiescent animals, and to attended than to unattended stimuli (Roelfsema et al., 1998; Fries et al., 2001; Cohen and Maunsell, 2009; Mitchell et al., 2009; Chalk et al., 2010). The result presented here therefore conflicts with the hypothesis that the synchronized and desynchronized cortical states in anesthetized animals are respectively analogous to quiescent and attending periods in awake animals (Harris and Thiele, 2011). Perhaps the relationship is more complex than previously thought.

The present finding that responses are less precise and reliable during the desynchronized state is consistent with Fiser et al. (2004). In awake freely viewing (supposedly highly desynchronized) ferret V1, that study concluded that superficial layer multiunit firing patterns were dominated by cortical state rather than stimulus. In other words, stimulus did little to affect multiunit responses in the desynchronized state. However, during anesthesia in the same animals (supposedly in a more synchronized state), stimulus had much more of an effect on multiunit responses. This only indirectly agrees with Figures 6.14–6.15, because Fiser et al. (2004) did not show any trial raster plots from which direct conclusions about precision and reliability might be made.

Marguet and Harris (2011) also found trial-averaged LFP to be more reliable in the desynchronized state than the synchronized state. Although trial-averaged LFP was not calculated here, doing so would be simple, and given the conclusions here on spike train reliability, it is likely that for anesthetized cat V1, the trial-averaged LFP will also be more reliable in the synchronized state.

Almost by definition, PSTH correlations between cell pairs were higher on average in the synchronized than desynchronized state, but only modestly (Figure 6.18). There is wide agreement on this in the literature (Poulet and Petersen, 2008; Goard and Dan, 2009; Curto et al., 2009; Renart et al., 2010; Marguet and Harris, 2011; Schölvinck et al., 2015; Pachitariu et al., 2015). Even though PSTH correlations were only modestly influenced by cortical state, that influence was greater than the particular choice of natural scene movie clip that was presented (Figure 6.19). Though perhaps somewhat controversial, this conclusion agrees with a number of other studies (Arieli et al., 1996; Fiser et al., 2004; Marguet and Harris, 2011; Okun et al., 2012). (Okun et al. (2012) is in a somewhat different context of spike correlations and multineuron words over the course of ferret visual cortex development, but is still relevant. See Appendix B.)

The conclusions made here are based on only 78 min of data from two recordings from the same track. This was due to the paucity of desynchronized periods in the three tracks chosen for careful spike sorting and detailed analysis (ptc15.tr7c, ptc22.tr1, ptc22.tr2). Other movie recordings from these tracks had less dramatic cortical state changes, as exhibited by their LFP. They also seemed to exhibit the same relationship between level of synchronization and temporal precision of response events (not shown), but that relationship was less obvious. The surprising conclusion that responses are more precise and reliable in the synchronized state should be strengthened by spike sorting more tracks with greater periods of desynchronization (Figure 6.11), and therefore a greater chance of spontaneously switching between synchronized and desynchronized states. Indeed, another 4 natural scene movie recordings showed sudden switches in cortical state (not shown), similar to those shown in Figures 6.9 & 6.10. These totalled an additional 130 min of recording, from 4 more tracks from 2 more cats (ptc17.tr2b, ptc18.tr1, ptc18.tr2c, ptc22.tr4b). Preliminary spike sorting of these additional recordings (courtesy of Catalin Mitelut) showed the same relationship between cortical state and spiking response precision and reliability as in Section 6.6 (not shown), thereby strengthening the above conclusion.

The third cell in Figure 6.14 (rightmost column) was classified as an LGN afferent, yet it was clearly modulated by cortical state. Two other such cells were found. As part of the thalamus, the LGN can drive cortical state changes (Steriade et al., 1991, 1993c; Hirata and Castro-Alamancos, 2010), but cortical state changes are not thought to directly affect the LGN, although indirect influence via the thalamic reticular nucleus is possible (Steriade et al., 1993c). The cells classified here as LGN afferents may therefore reflect state changes in LGN which subsequently drive changes in cortical state.

The influence of cortical state on precision and reliability during presentation of artificial stimuli was not examined here, partly because artificial stimuli such as m-sequence white noise movies, drifting bars, and drifting gratings do not induce very precise or reliable responses compared to natural scene movies. Nevertheless, this is something that could be examined in the future.

6.7.6 Clustering cortical states

Labelling cortical state as synchronized and desynchronized is a start, but is almost certainly overly simplistic (Harris and Thiele, 2011). Future work should also partition especially the desynchronized cortical state into further categories, given the large, sudden changes in LFP spectrograms during desynchronized periods (Figures 6.9B & 6.10B), and the trimodality of the SI plots in some tracks (Figure 6.11). Instead of using a simple power ratio of somewhat arbitrarily chosen high and low frequency bands, the full distribution of power across frequencies, i.e., the PSD (Figure 6.8A), at each point in time could be characterized in a more sophisticated way. For example, a high dimensional data set could be generated by collecting LFP PSDs at all points in time in all recordings in all tracks (or perhaps separately for each track). If 1 Hz wide frequency bands were used from 0 to 100 Hz, this would create a 100 dimensional space varying over time. PCA or ICA or some other dimension reduction method could be applied to this space to extract the top 2 or 3 most interesting or descriptive components. Upon mapping each point in time into this low dimensional space, multiple clusters might emerge. Upon clustering of this dimension reduced data (perhaps using GAC, Section 3.9.1), dimension reduction could be redone separately for each cluster, creating a new space with perhaps further subclusters (akin to checking for undersplit clusters in Section 3.10.1). One such expected subclustering is the distinction between UP and DOWN phases within the synchronized state, but there may be many more. Differences in PSDs across channels could also be exploited to characterize cortical state as a function of depth. All of this could allow for a (potentially nested) multidimensional metric of cortical state as a function of time, and a simultaneous increase in the number of labelled states, each potentially having a different computational role.

7 Conclusion

This thesis examined the use of silicon polytrodes for large-scale extracellular recordings of local populations of neurons in anesthetized cat V1. A wide variety of both artificial and naturalistic visual stimuli were used (Section 2.4), and a new divide-and-conquer spike sorting method was devised to translate correlated multisite extracellular potential waveforms into sorted spikes from isolated single units (Chapter 3). From stimulus and recording, to spike sorting and analysis, many steps were taken to minimize bias in the characterization of patches of V1. Many results concurred with previous studies, but some results were surprising.

Mean firing rates were surprisingly low and lognormally distributed, supporting the theory of sparse coding (Section 4.3). There was also evidence for neural shift work, which could be a useful strategy for physiological maintenance or network stability (Section 4.6.2). 65 μm site spacing was found to be too great to fully capture the local neuronal population (Section 4.4), suggesting that future polytrodes should use a higher density of electrode sites to decrease the chance of missed cells. Surprisingly, orientation tuning strength was inversely correlated with stimulus-evoked log firing rates (Section 4.5), demonstrating the utility of low-rate cells.

Cells could be classified into at least 2 and possibly up to 4 different types based on their spike shape (fast, slow, fast asymmetric, slow asymmetric; Section 5.2), and into 4 types based on their spatiotemporal receptive fields (simple, complex, LGN afferent, unknown; Section 5.4). Spatial extent of 2D spatially localized multichannel spike waveforms was not useful for classifying cells (Section 5.3). Approximately equal numbers of cells were classified as simple, complex, and unknown RF type, with only 5% classified as LGN afferents. Cells classified as unknown RF type may have been mostly complex cells damaged by polytrode insertion.

Natural scene movies evoked sparse and reliable spike patterns in most cells that were active, with temporal precision measured in milliseconds (Section 6.2), lending support to the importance of spike timing in cortical neural coding. Spontaneous changes in cortical state (Section 6.4) played a major role in these responses. Consistent with other reports, UP/DOWN phases were present during the synchronized state, and absent during the desynchronized state (Section 6.5). Contrary to reports in other preparations, responses were more precise, sparse, and reliable in the synchronized state than in the desynchronized state (Section 6.6). Pairwise correlations of natural scene movie responses between most cells were close to zero, but mean correlations across the population were slightly positive. Correlations were lower than expected between simple and complex cells, and higher than expected between complex cells (Section 6.3), challenging the hierarchical model of complex cells and supporting a recurrent model instead (Section 6.7.2). Response correlations were influenced more by cortical state than by the specific natural scene movie presented (Section 6.6), providing further evidence of the importance of cortical state.

Although bias was minimized in as many ways as practically possible, some bias inevitably remained. For example, due to time and space limitations, this study focused on only 3 tracks from 3 hemispheres in 2 cats, resulting in 245 isolated cells. Increasing the number of cells, tracks, and cats requires further spike sorting, but is a top priority as it would strengthen many of the conclusions in Chapters 4–6.

Despite long-duration recordings at fixed positions and substantial effort dedicated to detecting low-rate cells, neuronal yields were still $< 15\%$ of the known anatomical number of neurons within isolatable recording range (Section 4.6.1). One possible source of remaining bias was that 65 μm electrode site spacing selected for cells with large open extracellular potential fields, and against those with small closed fields. Fortunately, the solution to this is relatively simple: use a polytrode design with lower electrode site spacing. Balancing this against the desire to simultaneously record from all cortical layers will require more electrode sites in future experiments. Since polytrodes (e.g., Buzsaki256 probe, NeuroNexus, Ann Arbor, MI) and neurophysiological recording systems (e.g., RHD2000 system, Intan Technologies, Los Angeles, CA) with up to 256 sites and channels are now easily and relatively cheaply available, this should not be a major obstacle. Higher channel counts may increase the overall amount of time required for spike sorting, but greater channel densities may make it easier to distinguish neighbouring neurons and may therefore decrease the amount of spike sorting time required per neuron.

Polytrode tissue damage was likely another big source of remaining bias, one that is more difficult to overcome. Damage could be mitigated in a variety of ways. One way is to insert the polytrode at a much slower rate, on the order of 10 μm/min or less (Schjetnan and Luczak, 2011; Berényi et al., 2014), which may reduce damage associated with dimpling. Another is to use narrower polytrodes. The ~ 200 μm shank width of the polytrodes used here is at least twice that of most other designs, and failed to isolate cells in several attempted experiments in urethane-anesthetized rat V1. In comparison, many other groups have had success with narrower polytrodes in rat cortex (Section 4.6.1). If the 200 μm shank width is responsible for such a radical difference in neuron yield in rat cortex, it could also easily affect yield in cat V1, but perhaps less radically. At least one sufficiently long high-density polytrode design of narrower width is available (A1x64-Poly2-6mm-23s-160, Anton Sirota lab, manufactured by NeuroNexus, Ann Arbor, MI) which tapers to a maximum width of 115 μm.

Most of the width of a polytrode is to provide space for non-overlapping conductors from each of the electrode sites up to the headstage connector. Ideally, conductor requirements should not dictate shank width, and polytrode shanks should be no wider than the desired site spacing ($<$ 38 μm in the above design). This would require either much narrower conductors, or a multilayer design, or some combination of the two. Given that these are passive low current conductors, such changes could lead to undesired coupling between neighbouring conductors (but see Du et al. (2011)). The ultimate solution to this problem might require active electronics, perhaps a current buffer stage located immediately adjacent to each electrode site. Unfortunately, given the specialized nature of polytrodes and their low manufacturing volumes, this would likely result in greater costs per polytrode, which are fragile and easily broken during handling.

Given the dramatic influence of cortical states on spiking responses (Section 6.6), further work should carefully control for cortical states, manipulate them experimentally, and classify them with greater fidelity than simply synchronized vs. desynchronized (Section 6.7.6). Spike correlations (Appendix B) are widely used to infer functional connectivity (Barthó et al., 2004; Du et al., 2011) and to explain population spike patterns (Schneidman et al., 2006; Shlens et al., 2006; Yu et al., 2008; Ohiorhenuan et al., 2010; Berkes et al., 2011), yet they are very sensitive to changes in cortical state (Goard and Dan, 2009; Harris and Thiele, 2011; Okun et al., 2012). Cortical state should therefore be considered during spike correlation analysis. Great care should also be taken in the application of anticholinergics, which induce extended periods of synchronization (Section 6.4). During long duration anesthesia, short-term reduction of isoflurane levels may allow for controlled periods of desynchronization without risk of loss of anesthesia (Section 6.7.3). More sophisticated chronic recordings in unanesthetized animals in various waking and sleep states (awake behaving, quiet resting, REM sleep, slow-wave sleep) might allow for a broader characterization of cortical states, and more direct comparisons to cortical states from the same cells in the same animal under anesthesia. This would help strengthen the argument that cortical state in anesthetized animals is closely related to that in awake behaving animals, and that, for example, the desynchronized state is closely related to attention (Harris and Thiele, 2011). Unfortunately, precisely controlling the projection of visual stimuli onto the retina of awake or sleeping animals is very difficult, yet doing so is necessary to gauge the temporal precision and reliability of spiking responses to repeated natural scene movies in such animals.

Moore's law describes the exponential growth of the number of transistors that can be packed into a single integrated circuit, with a doubling time of about 2 years. Although the growth rate of neuronal recording technology, as measured by the maximum number of simultaneously recordable neurons, is slower than Moore's law, it too is exponential, with a doubling period of roughly 7 years (Stevenson and Kording, 2011). Within just a few years, polytrodes could consist of hundreds of electrode sites and dozens of shanks, perhaps multisided or circular, allowing simultaneous monitoring of thousands of neurons (Einevoll et al., 2012) from ever larger but mostly contiguous volumes of grey matter.

Like most technologies, the silicon polytrodes used here for large-scale neuronal recording will eventually be supplanted by a superior technology. Two-photon calcium imaging is currently limited by both its recording depth and its reliance on the slow calcium signal which prevents detection of temporally precise spike patterns (Section 1.1). While the former limitation is improving, the latter remains a greater challenge. Temporal resolution might be improved by developing a genetically engineered fluorophore whose change in fluorescence directly corresponds to change in membrane potential, analogous to voltage-sensitive dye recordings used at coarser spatial scales (Arieli et al., 1995, 1996; Mohajerani et al., 2010). Such a fluorophore would not only allow detection of arbitrarily fine spike patterns, but might also provide optical access to the subthreshold membrane potential, combining the best of both intracellular and extracellular electrophysiology. Such an advance, combined with its lack of tissue damage, would allow two-photon imaging to leapfrog polytrodes for large-scale monitoring of neuronal populations.

A more exotic proposal for large-scale recording might involve molecular “ticker tape”, in which a constant rate of copying of DNA by polymerase in a neuron is interrupted by copying anomalies on every occurrence of a spike (Kording, 2011). This nucleotide signal could then be read out offline by independently sequencing the DNA of each neuron to obtain its temporal pattern of spikes. Another exotic possibility is nano-scale magnetic resonance imaging (McGuinness et al., 2011; Hemmer, 2013), in which tiny ultra-pure diamonds with engineered defects can be used as highly sensitive room temperature magnetometers, capable of detecting changes in the magnetic spin of small numbers of nearby electrons or protons using pulsed optical and radio frequency stimulation and readout. These nanodiamonds might be embedded in cell membranes to provide high-speed non-invasive readout of membrane potential from many neurons at a time.

While ever-improving recording technology will help reduce bias in our understanding of the brain, care must also be taken with all experimental and analytical methods to do the same. As the number of simultaneously recorded cells increases, the number of pairwise interactions to analyze increases with its square. The combinatorial explosion involved in the analysis of such data could prove to be more challenging than the technological developments or the experiments themselves. However these problems are tackled, the study of a system as complex as the brain demands the utmost care and attention towards minimizing bias wherever possible.

Bibliography

Abbott LF, Chance FS (2002) Rethinking the taxonomy of visual neurons. *Nat Neurosci* 5:391–392. → page 90, 143

Adrian ED, Bronk DW (1928) The discharge of impulses in motor nerve fibres: Part I. Impulses in single fibres of the phrenic nerve. *J Physiol* 66:81–101. → page 1

Ahmed B, Anderson JC, Douglas RJ, Martin KAC, Nelson JC (1994) Polyneuronal innervation of spiny stellate neurons in cat visual cortex. *J Comp Neurol* 341:39–49. → page 140

Alonso JM, Martinez LM (1998) Functional connectivity between simple cells and complex cells in cat striate cortex. *Nat Neurosci* 1:395–403. → page 108, 143

Alonso JM, Usrey WM, Reid RC (1996) Precisely correlated firing in cells of the lateral geniculate nucleus. *Nature* 383:815–819. → page 62, 140

Anderson J, Lampl I, Reichova I, Carandini M, Ferster D (2000) Stimulus dependence of two-state fluctuations of membrane potential in cat visual cortex. *Nat Neurosci* 3:617–621. → page 110, 126, 147

Arieli A, Shoham D, Hildesheim R, Grinvald A (1995) Coherent spatiotemporal patterns of ongoing activity revealed by real-time optical imaging coupled with single-unit recording in the cat visual cortex. *J Neurophysiol* 73:2072–2093. → page 154

Arieli A, Sterkin A, Grinvald A, Aertsen A (1996) Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* 273:1868–1871. → page 6, 110, 149, 154

Attwell D, Laughlin SB (2001) An energy budget for signaling in the grey matter of the brain. *J Cereb Blood Flow Metab* 21:1133–1145. → page 1, 81, 140

Bair W, Koch C (1996) Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Comput* 8:1185–1202. → page 5, 137

Baker SN, Lemon RN (2000) Precise spatiotemporal repeating patterns in monkey primary and supplementary motor areas occur at chance levels. *J Neurophysiol* 84:1770–1780. → page 5

Bar-Hillel A, Spiro A, Stark E (2006) Spike sorting: Bayesian clustering of non-stationary data. *J Neurosci Meth* 157:303–316. → page 24, 27, 45

Barthó P, Hirase H, Monconduit L, Zugaro M, Harris KD, Buzsáki G (2004) Characterization of neocortical principal cells and interneurons by network interactions and extracellular features. *J Neurophysiol* 92:600–608. → page 89, 90, 104, 105, 154

Battaglia FP, Kalenscher T, Cabral H, Winkel J, Bos J, Manuputy R, van Lieshout T, Pinkse F, Beukers H, Pennartz C (2009) The lantern: an ultra-light micro-drive for multi-tetrode recordings in mice and other small animals. *J Neurosci Meth* 178:291–300. → page 3

Beaulieu C, Colonnier M (1985) A comparison of the number of neurons in individual laminae of cortical areas 17, 18 and posteromedial suprasylvian (PMLS) area in the cat. *Brain Res* 339:166–170. → page 4

Behnel S, Bradshaw R, Citro C, Dalcin L, Seljebotn DS, Smith K (2011) Cython: The best of both worlds. *IEEE Comput Sci Eng Mag* 13:31–39. → page 181

Bell AJ, Sejnowski TJ (1995) An information-maximization approach to blind separation and blind deconvolution. *Neural Comput* 7:1129–1159. → page 41

Benchenane K, Peyrache A, Khamassi M, Tierney PL, Gioanni Y, Battaglia FP, Wiener SI (2010) Coherent theta oscillations and reorganization of spike timing in the hippocampal-prefrontal network upon learning. *Neuron* 66:921–936. → page 89, 104, 105

Benucci A, Frazor RA, Carandini M (2007) Standing waves and traveling waves distinguish two circuits in visual cortex. *Neuron* 55:103–117. → page 149

Berényi A, Somogyvári Z, Nagy AJ, Roux L, Long JD, Fujisawa S, Stark E, Leonardo A, Harris TD, Buzsáki G (2014) Large-scale, high-density (up to 512 channels) recording of local circuits in behaving animals. *J Neurophysiol* 111:1132–1149. → page 3, 5, 80, 153

Berger H (1929) Über das elektrenkephalogramm des Menschen. *Arch Psychiatr Nervenkr* 87:527–570. → page 6, 110

Berkes P, Orbán G, Lengyel M, Fiser J (2011) Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science* 331:83–87. → page 154

Berry MJ, Warland DK, Meister M (1997) The structure and precision of retinal spike trains. *PNAS* 94:5411–5416. → page 140

Bilmes JA (1998) A gentle tutorial of the EM algorithm and its application to parameter estimation for Gaussian mixture and hidden Markov models . → page 27

Binzegger T, Douglas RJ, Martin KAC (2004) A quantitative map of the circuit of cat primary visual cortex. *J Neurosci* 24:8441–8453. → page 140

Biran R, Martin DC, Tresco PA (2005) Neuronal cell loss accompanies the brain tissue response to chronically implanted silicon microelectrode arrays. *Exp Neurol* 195:115–126. → page 82

Blanche TJ (2005) Large scale neuronal recording. Ph.D. diss., University of British Columbia, Vancouver, British Columbia, Canada. → page 4, 14, 15, 25, 85, 89, 90, 106, 179

Blanche TJ, Godfrey K, Douglas RM, Swindale NV (2008) Spike detection and sorting algorithms for polytrodes. *(unpublished)* . → page 25

Blanche TJ, Spacek MA, Hetke JF, Swindale NV (2005) Polytrodes: high-density silicon electrode arrays for large-scale multiunit recording. *J Neurophysiol* 93:2987–3000. → page iii, 3, 4, 5, 14, 15, 23, 29, 69, 80, 84, 186

Blanche TJ, Swindale NV (2006) Nyquist interpolation improves neuron yield in multiunit recordings. *J Neurosci Meth* 155:81–91. → page 29

Blatt M, Wiseman S, Domany E (1996) Superparamagnetic clustering of data. *Phys Rev Lett* 76:3251–3254. → page 27, 45

Blatt M, Wiseman S, Domany E (1997) Data clustering using a model granular magnet. *Neural Comput* 9:1805–1842. → page 27

Brecht M, Roth A, Sakmann B (2003) Dynamic receptive fields of reconstructed pyramidal cells in layers 3 and 2 of rat somatosensory barrel cortex. *J Physiol* 553:243–265. → page 81

Burr DC, Ross J (1979) How does binocular delay give information about depth? *Vision Res* 19:523–532. → page 5

Buzsáki G (2004) Large-scale recording of neuronal ensembles. *Nat Neurosci* 7:446–451. → page 5, 22, 80, 86

Buzsáki G, Anastassiou CA, Koch C (2012) The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nat Rev Neurosci* 13:407–420. → page 2

Calabrese A, Paninski L (2011) Kalman filter mixture model for spike sorting of non-stationary data. *J Neurosci Meth* 196:159–169. → page 24, 27

Carandini M, Demb JB, Mante V, Tolhurst DJ, Dan Y, Olshausen BA, Gallant J, Rust NC (2005) Do we know what the early visual system does? *J Neurosci* 25:10577–10597. → page 7, 81, 107, 110, 141

Carandini M, Ferster D (2000) Membrane potential and firing rate in cat primary visual cortex. *J Neurosci* 20:470–484. → page 107

Carandini M, Heeger DJ, Movshon JA (1997) Linearity and normalization in simple cells of the macaque primary visual cortex. *J Neurosci* 17:8621–8644. → page 141

Castro-Alamancos MA (2002) Role of thalamocortical sensory suppression during arousal: focusing sensory inputs in neocortex. *J Neurosci* 22:9651–9655. → page 146

Chalk M, Herrero JL, Gieselmann MA, Delicato LS, Gotthardt S, Thiele A (2010) Attention reduces stimulus-driven gamma frequency oscillations and spike field coherence in V1. *Neuron* 66:114–125. → page 149

Chance FS, Nelson SB, Abbott LF (1999) Complex cells as cortically amplified simple cells. *Nat Neurosci* 2:277–282. → page 90, 107, 108, 143, 144

Chubykin AA, Roach EB, Bear MF, Shuler MGH (2013) A cholinergic mechanism for reward timing within primary visual cortex. *Neuron* 77:723–735. → page 110

Cohen MR, Maunsell JHR (2009) Attention improves performance primarily by reducing interneuronal correlations. *Nat Neurosci* 12:1594–1600. → page 149

Comon P (1994) Independent component analysis, a new concept? *Signal Processing* 36:287–314. → page 41

Csicsvari J, Henze DA, Jamieson B, Harris KD, Sirota A, Barthó P, Wise KD, Buzsáki G (2003) Massively parallel recording of unit and local field potentials with silicon-based electrodes. *J Neurophysiol* 90:1314–1323. → page 3, 4, 5, 23, 80

Csicsvari J, Hirase H, Czurko A, Buzsáki G (1998) Reliability and state dependence of pyramidal cell–interneuron synapses in the hippocampus: an ensemble approach in the behaving rat. *Neuron* 21:179–189. → page 89, 104

Curto C, Sakata S, Marguet S, Itskov V, Harris KD (2009) A simple model of cortical dynamics explains variability and state dependence of sensory responses in urethane-anesthetized auditory cortex. *J Neurosci* 29:10600–10612. → page 146, 149

Dagum L, Menon R (1998) OpenMP: an industry standard API for shared-memory programming. *IEEE Comput Sci Eng Mag* 5:46–55. → page 181

Dan Y, Atick JJ, Reid RC (1996) Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *J Neurosci* 16:3351–3362. → page 140

David SV, Vinje WE, Gallant JL (2004) Natural stimulus statistics alter the receptive field structure of V1 neurons. *J Neurosci* 24:6991–7006. → page 110, 141

De Valois RL, Yund EW, Hepler N (1982) The orientation and direction selectivity of cells in macaque visual cortex. *Vision Res* 22:531–544. → page 87

Dean AF, Tolhurst DJ (1983) On the distinctness of simple and complex cells in the visual cortex of the cat. *J Physiol* 344:305–325. → page 90, 106, 107

DeAngelis GC, Ohzawa I, Freeman RD (1993) Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. I. General characteristics and postnatal development. *J Neurophysiol* 69:1091–1117. → page 89, 99

Denk W, Strickler JH, Webb WW (1990) Two-photon laser scanning fluorescence microscopy. *Science* 248:73–76. → page 4

Destexhe A, Contreras D, Steriade M (1999) Spatiotemporal analysis of local field potentials and unit discharges in cat cerebral cortex during natural wake and sleep states. *J Neurosci* 19:4595–4608. → page 110, 147

Drake KL, Wise KD, Farraye J, Anderson DJ, BeMent SL (1988) Performance of planar multisite microprobes in recording extracellular single-unit intracortical activity. *IEEE Trans Biomed Eng* 35:719–732. → page 3, 23

Du J, Blanche TJ, Harrison RR, Lester HA, Masmanidis SC (2011) Multiplexed, high density electrophysiology with nanofabricated neural probes. *PLoS ONE* 6:e26204. → page 84, 153, 154

Du J, Riedel-Kruse IH, Nawroth JC, Roukes ML, Laurent G, Masmanidis SC (2009) High-resolution three-dimensional extracellular recording of neuronal activity with microfabricated electrode arrays. *J Neurophysiol* 101:1671–1678. → page 3, 4, 23

Ecker AS, Berens P, Keliris GA, Bethge M, Logothetis NK, Tolias AS (2010) Decorrelated neuronal firing in cortical microcircuits. *Science* 327:584–587. → page 142

Eger II EI, Johnson BH (1987) Rates of awakening from anesthesia with I-653, halothane, isoflurane, and sevoflurane: a test of the effect of anesthetic concentration and duration in rats. *Anesth Analg* 66:977–982. → page 146

Einevoll GT, Franke F, Hagen E, Pouzat C, Harris KD (2012) Towards reliable spike-train recordings from thousands of neurons with multielectrodes. *Curr Opin Neurobiol* 22:11–17. → page 24, 57, 154

Engel AK, Fries P, Singer W (2001) Dynamic predictions: oscillations and synchrony in top–down processing. *Nature Rev Neurosci* 2:704–716. → page 5

Ester M, Kriegel HP, Sander J, Xu X (1996) A density-based algorithm for discovering clusters in large spatial databases with noise In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, pp. 226–231, Portland, OR. AAAI Press. → page 61

Fee MS, Mitra PP, Kleinfeld D (1996) Automatic sorting of multiple unit neuronal signals in the presence of anisotropic and non-gaussian variability. *J Neurosci Meth* 69:175–188. → page 24, 27, 62

Fiser J, Chiu C, Weliky M (2004) Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature* 431:573–578. → page 6, 110, 149

Fisher NI (1995) *Statistical analysis of circular data.* Cambridge University Press. → page 76

Forgy EW (1965) Cluster analysis of multivariate data: efficiency versus interpretability of classifications. *Biometrics* 21:768–769. → page 26

Franke F, Natora M, Boucsein C, Munk MHJ, Obermayer K (2010) An online spike detection and spike classification algorithm capable of instantaneous resolution of overlapping spikes. *J Comput Neurosci* 29:127–148. → page 24, 26

Fries P, Reynolds JH, Rorie AE, Desimone R (2001) Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291:1560–1563. → page 149

Fukunaga K, Hostetler L (1975) The estimation of the gradient of a density function, with applications in pattern recognition. *IEEE Trans Inf Theory* 21:32–40. → page 48, 60

Galarreta M, Hestrin S (1999) A network of fast-spiking cells in the neocortex connected by electrical synapses. *Nature* 402:72–75. → page 142

Garcia-Lazaro JA, Belliveau LA, Lesica NA (2013) Independent population coding of speech with sub-millisecond precision. *J Neurosci* 33:19362–19372. → page 5

Garland CM, Foreman RC, Chad JE, Holden-Dye L, Walker RJ (1998) The actions of muscle relaxants at nicotinic acetylcholine receptor isoforms. *Eur J Pharmacol* 357:83–92. → page 124

Gasthaus J, Wood F, Görür D, Teh YW (2009) Dependent Dirichlet process spike sorting In *Adv Neural Inf Process Syst*, Vol. 21, pp. 497–504. Curran Associates, Inc. → page 24

Gershon ED, Wiener MC, Latham PE, Richmond BJ (1998) Coding strategies in monkey V1 and inferior temporal cortices. *J Neurophysiol* 79:1135–1144. → page 88

Ghose GM, Freeman RD, Ohzawa I (1994) Local intracortical connections in the cat's visual cortex: postnatal development and plasticity. *J Neurophysiol* 72:1290–1303. → page 143

Gibson JR, Beierlein M, Connors BW (1999) Two networks of electrically coupled inhibitory neurons in neocortex. *Nature* 402:75–79. → page 142

Gilbert CD (1977) Laminar differences in receptive field properties of cells in cat primary visual cortex. *J Physiol* 268:391–421. → page 89, 107, 108

Goard M, Dan Y (2009) Basal forebrain activation enhances cortical coding of natural scenes. *Nat Neurosci* 12:1444–1449. → page 80, 111, 136, 146, 148, 149, 154

Gollisch T, Meister M (2008) Rapid neural coding in the retina with relative spike latencies. *Science* 319:1108–1111. → page 5, 140

Gray CM, König P, Engel AK, Singer W (1989) Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature* 338:334–337. → page 5

Gray CM, Maldonado PE, Wilson M, McNaughton B (1995) Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex. *J Neurosci Meth* 63:43–54. → page 3, 22, 24, 26, 62

Gray CM, McCormick DA (1996) Chattering cells: superficial pyramidal neurons contributing to the generation of synchronous oscillations in the visual cortex. *Science* 274:109–113. → page 90

Green J (1958) A simple microelectrode for recording from the central nervous system. *Nature* 182:962. → page 2

Gur M, Beylin A, Snodderly DM (1999) Physiological properties of macaque V1 neurons are correlated with extracellular spike amplitude, duration, and polarity. *J Neurophysiol* 82:1451–1464. → page 89, 104

Hahnloser RHR, Kozhevnikov AA, Fee MS (2002) An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature* 419:65–70. → page 80

Haider B, Duque A, Hasenstaub AR, McCormick DA (2006) Neocortical network activity in vivo is generated through a dynamic balance of excitation and inhibition. *J Neurosci* 26:4535–4545. → page 147

Haider B, Krause MR, Duque A, Yu Y, Touryan J, Mazer JA, McCormick DA (2010) Synaptic and network mechanisms of sparse and reliable visual cortical activity during nonclassical receptive field stimulation. *Neuron* 65:107–121. → page 7, 103, 137, 140, 142, 148

Hammond P, Andrews DP (1978) Orientation tuning of cells in areas 17 and 18 of the cat's visual cortex. *Exp Brain Res* 31:341–351. → page 87

Harris KD (2005) Neural signatures of cell assembly organization. *Nature Rev Neurosci* 6:399–407. → page 5

Harris KD, Csicsvari J, Hirase H, Dragoi G, Buzsáki G (2003) Organization of cell assemblies in the hippocampus. *Nature* 424:552–6. → page 5, 80, 114

Harris KD, Henze DA, Csicsvari J, Hirase H, Buzsáki G (2000) Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *J Neurophysiol* 84:401–414. → page 3, 23, 24, 27, 55, 62

Harris KD, Henze DA, Hirase H, Leinekugel X, Dragoi G, Czurkó A, Buzsáki G (2002) Spike train dynamics predicts theta-related phase precession in hippocampal pyramidal cells. *Nature* 417:738–741. → page 5

Harris KD, Thiele A (2011) Cortical state and attention. *Nat Rev Neurosci* 12:509–523. → page 2, 6, 110, 121, 129, 145, 146, 148, 149, 150, 154, 178

Hartigan JA, Hartigan PM (1985) The dip test of unimodality. *Ann Stat* 13:70–84. → page 60

Hasenstaub A, Sachdev RNS, McCormick DA (2007) State changes rapidly modulate cortical neuronal responsiveness. *J Neurosci* 27:9607–9622. → page 147

Havenith MN, Yu S, Biederlack J, Chen NH, Singer W, Nikolić D (2011) Synchrony makes neurons fire in sequence, and stimulus properties determine who is ahead. *J Neurosci* 31:8570–8584. → page 5

Hazan L, Zugaro M, Buzsáki G (2006) Klusters, NeuroScope, NDManager: a free software suite for neurophysiological data processing and visualization. *J Neurosci Meth* 155:207–216. → page 24, 27, 54, 181

Hebb DO (1949) *The Organization of Behavior: A Neuropsychological Theory.* John Wiley & Sons, New York. → page 4

Helmchen F, Denk W (2005) Deep tissue two-photon microscopy. *Nat Methods* 2:932–940. → page 4

Hemmer P (2013) Toward molecular-scale MRI. *Science* 339:529–530. → page 155

Henze DA, Borhegyi Z, Csicsvari J, Mamiya A, Harris KD, Buzsáki G (2000) Intracellular features predicted by extracellular recordings in the hippocampus in vivo. *J Neurophysiol* 84:390–400. → page 3, 4, 22, 58, 80, 92, 104

Herikstad R, Baker J, Lachaux JP, Gray CM, Yen SC (2011) Natural movies evoke spike trains with low spike time variability in cat primary visual cortex. *J Neurosci* 31:15844–15860. → page 7, 90, 137, 140, 148

Herrero JL, Roberts MJ, Delicato LS, Gieselmann MA, Dayan P, Thiele A (2008) Acetylcholine contributes through muscarinic receptors to attentional modulation in V1. *Nature* 454:1110–1114. → page 145

Hill DN, Mehta SB, Kleinfeld D (2011) Quality metrics to accompany spike sorting of extracellular signals. *J Neurosci* 31:8699–8705. → page 52, 62

Hirata A, Castro-Alamancos MA (2010) Neocortex network activation and deactivation states controlled by the thalamus. *J Neurophysiol* 103:1147–1157. → page 150

Hodgkin AL, Huxley AF (1939) Action potentials recorded from inside a nerve fibre. *Nature* 144:710–711. → page 1

Hodgkin AL, Huxley AF (1952) A quantitative description of membrane current and its application to conduction and excitation in nerve. *J Physiol* 117:500–544. → page 2

Hofer SB, Ko H, Pichler B, Vogelstein J, Ros H, Zeng H, Lein E, Lesica NA, Mrsic-Flogel TD (2011) Differential connectivity and response dynamics of excitatory and inhibitory neurons in visual cortex. *Nature Neurosci* 14:1045–1052. → page 142

Hoffman KL, Battaglia FP, Harris K, MacLean JN, Marshall L, Mehta MR (2007) The upshot of up states in the neocortex: from slow oscillations to memory formation. *J Neurosci* 27:11838–11841. → page 146

Hopfield JJ (1995) Pattern recognition computation using action potential timing for stimulus representation. *Nature* 376:33–36. → page 5, 140

Horton JC, Adams DL (2005) The cortical column: a structure without a function. *Phil Trans R Soc B* 360:837–862. → page 148

Hromádka T, DeWeese MR, Zador AM (2008) Sparse representation of sounds in the unanesthetized auditory cortex. *PLoS Biol* 6:e16. → page 81

Hubel DH (1957) Tungsten microelectrode for recording from single units. *Science* 125:549–550. → page 2

Hubel DH, Wiesel TN (1959) Receptive fields of single neurones in the cat's striate cortex. *J Physiol* 148:574–591. → page 2, 72, 106

Hubel DH, Wiesel TN (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Physiol* 160:106–154. → page 2, 89, 106, 108, 141

Hubel DH, Wiesel TN (1968) Receptive fields and functional architecture of monkey striate cortex. *J Physiol* 195:215–243. → page 2, 106

Hulata E, Segev R, Ben-Jacob E (2002) A method for spike sorting and detection based on wavelet packets and Shannon's mutual information. *J Neurosci Meth* 117:1–12. → page 24, 43

Hunter JD (2007) Matplotlib: A 2D graphics environment. *IEEE Comput Sci Eng Mag* 9:90–95. → page 181

Hyvärinen A (1998) New approximations of differential entropy for independent component analysis and projection pursuit In *Adv Neural Inf Process Syst*, Vol. 10, pp. 273–279, Cambridge, MA. MIT Press. → page 59

Hyvärinen A (1999) Fast and robust fixed-point algorithms for independent component analysis. *IEEE Trans Neural Networks* 10:626–634. → page 43

Hyvärinen A, Karhunen J, Oja E (2001) *Independent Component Analysis*. Wiley-Interscience. → page 41

Hyvärinen A, Oja E (2000) Independent component analysis: algorithms and applications. *Neural Networks* 13:411–430. → page 43

Ikegaya Y, Aaron G, Cossart R, Aronov D, Lampl I, Ferster D, Yuste R (2004) Synfire chains and cortical songs: Temporal modules of cortical activity. *Science* 304:559–564. → page 5

Jäckel D, Frey U, Fiscella M, Franke F, Hierlemann A (2012) Applicability of independent component analysis on high-density microelectrode array recordings. *J Neurophysiol* 108:334–348. → page 24, 26, 59

Ji D, Wilson MA (2007) Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nat Neurosci* 10:100–107. → page 110

Jog MS, Connolly CI, Kubota Y, Iyengar DR, Garrido L, Harlan R, Graybiel AM (2002) Tetrode technology: advances in implantable hardware, neuroimaging, and data analysis techniques. *J Neurosci Meth* 117:141–152. → page 3

Johansson RS, Birznieks I (2004) First spikes in ensembles of human tactile afferents code complex spatial fingertip events. *Nat Neurosci* 7:170–177. → page 5

Jones E, Oliphant T, Peterson P et al. (2001) SciPy: Open source scientific tools for Python. → page 181

Jones EG, Peters A (1987) *Cerebral Cortex – Further Aspects of Cortical Function, Including Hippocampus*, Vol. 6, chapter 7: Number of neurons and synapses in primary visual cortex, pp. 267–294. Plenum Press, New York. → page 4

Kagan I, Gur M, Snodderly DM (2002) Spatial organization of receptive fields of V1 neurons of alert monkeys: comparison with responses to gratings. *J Neurophysiol* 88:2557–2574. → page 89, 90, 106, 107, 108

Kandel ER, Schwartz JH, Jessell TM, Siegelbaum SA, Hudspeth AJ (2012) *Principles of Neural Science*, Vol. 5. McGraw-Hill, New York. → page 6

Kara P, Reinagel P, Reid RC (2000) Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron* 27:635–646. → page 140

Katzner S, Nauhaus I, Benucci A, Bonin V, Ringach DL, Carandini M (2009) Local origin of field potentials in visual cortex. *Neuron* 61:35–41. → page 2, 136

Kayser C, Einhäuser W, König P (2003) Temporal correlations of orientations in natural scenes. *Neurocomputing* 52:117–123. → page 19

Kelly RC, Smith MA, Kass RE, Lee TS (2010) Local field potentials indicate network state and account for neuronal response variability. *J Comput Neurosci* 29:567–579. → page 2

Kerr JND, Greenberg D, Helmchen F (2005) Imaging input and output of neocortical networks in vivo. *PNAS* 102:14063–14068. → page 4, 80, 81, 82

Kirchner H, Thorpe SJ (2006) Ultra-rapid object detection with saccadic eye movements: Visual processing speed revisited. *Vision Res* 46:1762–1776. → page 5

Klusch M, Lodi S, Moro G (2003) Distributed clustering based on sampling local density estimates In *International Joint Conference on Artificial Intelligence*, Vol. 18, pp. 485–490. Lawrence Erlbaum Associates Ltd. → page 61

Ko H, Hofer SB, Pichler B, Buchanan KA, Sjöström PJ, Mrsic-Flogel TD (2011) Functional specificity of local synaptic connections in neocortical networks. *Nature* 473:87–91. → page 5, 142

Koester HJ, Sakmann B (2000) Calcium dynamics associated with action potentials in single nerve terminals of pyramidal cells in layer 2/3 of the young rat neocortex. *J Physiol* 529:625–646. → page 4

Kording KP (2011) Of toasters and molecular ticker tapes. *PLoS Comput Biol* 7:e1002291. → page 57, 155

Kowalewski F (1995) A gradient procedure for determining clusters of relatively high point density. *Pattern Recognition* 28:1973–1984. → page 61

Krizhevsky A, Sutskever I, Hinton GE (2012) ImageNet classification with deep convolutional neural networks In *Adv Neural Inf Process Syst*, Vol. 25, pp. 1097–1105. Curran Associates, Inc. → page 6

Latham PE, Roth A, Hausser M, London M (2006) Requiem for the spike? *Soc Neurosci Abstr* 36:432.12. → page 5

Leach JB, Achyuta AKH, Murthy SK (2010) Bridging the divide between neuroprosthetic design, tissue engineering and neurobiology. *Front Neuroeng* 2:18. → page 82

Lennie P (2003) The cost of cortical computation. *Curr Biol* 13:493–497. → page 1, 3, 22, 81, 140

Lewicki MS (1994) Bayesian modeling and classification of neural signals. *Neural Comput* 6:1005–1030. → page 24

Lewicki MS (1998) A review of methods for spike sorting: the detection and classification of neural action potentials. *Network* 9:R53–R78. → page 22, 24, 40

Li CY, Poo MM, Dan Y (2009) Burst spiking of a single cortical neuron modifies global brain state. *Science* 324:643–646. → page 121

Lindén H, Tetzlaff T, Potjans TC, Pettersen KH, Grün S, Diesmann M, Einevoll GT (2011) Modeling the spatial reach of the LFP. *Neuron* 72:859–872. → page 2

Ling G, Gerard RW (1949) The normal membrane potential of frog sartorius fibers. *J Cell Compar Physl* 34:383–396. → page 2

Litke AM, Bezayiff N, Chichilnisky EJ, Cunningham W, Dabrowski W, Grillo AA, Grivich M, Grybos P, Hottowy P, Kachiguine S, Kalmar RS, Mathieson K, Petrusca D, Rahman M, Sher A (2004) What does the eye tell the brain?: Development of a system for the large-scale recording of retinal output activity. *IEEE Trans Nucl Sci* 51:1434–1440. → page 24, 62

Liu X, McCreery DB, Carter RR, Bullara LA, Yuen TGH, Agnew WF (1999) Stability of the interface between neural tissue and chronically implanted intracortical microelectrodes. *IEEE Trans Rehabil Eng* 7:315–326. → page 82

London M, Roth A, Beeren L, Häusser M, Latham PE (2010) Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature* 466:123–127. → page 5

Lowe DG (2004) Distinctive image features from scale-invariant keypoints. *Int J Comput Vision* 60:91–110. → page 6

Luczak A, Bartho P, Harris KD (2013) Gating of sensory input by spontaneous cortical activity. *J Neurosci* 33:1684–1695. → page 5, 80, 146, 147

Luczak A, Barthó P, Marguet SL, Buzsáki G, Harris KD (2007) Sequential structure of neocortical spontaneous activity in vivo. *PNAS* 104:347. → page 5, 89, 104, 105, 146, 147, 149

Lütcke H, Murayama M, Hahn T, Margolis DJ, Astori S, zum Alten Borgloh SM, Göbel W, Yang Y, Tang W, Kügler S, Sprengel R, Takeharu N, Miyawaki A, Larkum ME, Hlemchen F, Hasan MT (2010) Optical recording of neuronal activity with a genetically-encoded calcium indicator in anesthetized and freely moving mice. *Front Neural Circuits* 4:9. → page 4

MacKay DJC (2003) *Information Theory, Inference, and Learning Algorithms*, chapter 20: An Example Inference Task: Clustering, pp. 284–292. Cambridge Univ Pr. → page 26

Mackevicius EL, Best MD, Saal HP, Bensmaia SJ (2012) Millisecond precision spike timing shapes tactile perception. *J Neurosci* 32:15309–15317. → page 5

Maffei L, Fiorentini A, Bisti S (1973) Neural correlate of perceptual adaptation to gratings. *Science* 182:1036–1038. → page 18, 179

Mainen ZF, Sejnowski TJ (1995) Reliability of spike timing in neocortical neurons. *Science* 268:1503–1506. → page 5, 114, 140

Mank M, Santos AF, Direnberger S, Mrsic-Flogel TD, Hofer SB, Stein V, Hendel T, Reiff DF, Levelt C, Borst A, Bonhoeffer T, Hübener M, Griesbeck O (2008) A genetically encoded calcium indicator for chronic in vivo two-photon imaging. *Nat Methods* 5:805–811. → page 4

Marguet SL, Harris KD (2011) State-dependent representation of amplitude-modulated noise stimuli in rat auditory cortex. *J Neurosci* 31:6414–6420. → page 6, 111, 148, 149

Markram H, Tsodyks M (1996) Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature* 382:807–810. → page 82

Marre O, Amodei D, Deshmukh N, Sadeghi K, Soo F, Holy TE, Berry II MJ (2012) Mapping a complete neural population in the retina. *J Neurosci* 32:14859–14873. → page 24, 26, 62

Massimini M, Huber R, Ferrarelli F, Hill S, Tononi G (2004) The sleep slow oscillation as a traveling wave. *J Neurosci* 24:6862–6870. → page 149

Mata ML, Ringach DL (2005) Spatial overlap of ON and OFF subregions and its relation to response modulation ratio in macaque primary visual cortex. *J Neurophysiol* 93:919–928. → page 90, 106, 107, 143

McCormick DA, Connors BW, Lighthall JW, Prince DA (1985) Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *J Neurophysiol* 54:782–806. → page 58, 92, 103, 105

McGuinness LP, Yan Y, Stacey A, Simpson DA, Hall LT, Maclaurin D, Prawer S, Mulvaney P, Wrachtrup J, Caruso F, Scholten RE, Hollenberg LCL (2011) Quantum measurement and orientation tracking of fluorescent nanodiamonds inside living cells. *Nature Nanotechnol* 6:358–363. → page 155

McNaughton BL, O'Keefe J, Barnes CA (1983) The stereotrode: a new technique for simultaneous isolation of several single units in the central nervous system from multiple unit records. *J Neurosci Meth* 8:391–397. → page 3, 22

Mechler F, Ringach DL (2002) On the classification of simple and complex cells. *Vision Res* 42:1017–1033. → page 90, 143

Mitchell JF, Sundberg KA, Reynolds JH (2007) Differential attention-dependent response modulation across cell classes in macaque visual area V4. *Neuron* 55:131–141. → page 89, 104

Mitchell JF, Sundberg KA, Reynolds JH (2009) Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* 63:879–888. → page 149

Mittmann W, Wallace DJ, Czubayko U, Herb JT, Schaefer AT, Looger LL, Denk W, Kerr JND (2011) Two-photon calcium imaging of evoked activity from l5 somatosensory neurons in vivo. *Nat Neurosci* 14:1089–1093. → page 4

Mitzdorf U (1985) Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and EEG phenomena. *Physiol Rev* 65:37–100. → page 86

Mizuseki KS A, Pastalkova E, Buzsáki G (2009) Theta oscillations provide temporal windows for local circuit computation in the entorhinal-hippocampal loop. *Neuron* 64:267–280. → page 89, 104

Mizuseki K, Buzsáki G (2013) Preconfigured, skewed distribution of firing rates in the hippocampus and entorhinal cortex. *Cell Reports* 4:1010–1021. → page 22, 81

Mohajerani MH, McVea DA, Fingas M, Murphy TH (2010) Mirrored bilateral slow-wave cortical activity within local circuits revealed by fast bihemispheric voltage-sensitive dye imaging in anesthetized and awake mice. *J Neurosci* 30:3745–3751. → page 149, 154

Murtagh F, Contreras P (2012) Algorithms for hierarchical clustering: an overview. *WIREs Data Mining Knowl Discov* 2:86–97. → page 27

Nason GP, Sibson R (1992) Measuring multimodality. *Stat Comput* 2:153–160. → page 60

Nguyen DP, Frank LM, Brown EN (2003) An application of reversible-jump markov chain monte carlo to spike classification of multi-unit extracellular recordings. *Network* 14:61–82. → page 24

Niell CM, Stryker MP (2008) Highly selective receptive fields in mouse visual cortex. *J Neurosci* 28:7520–7536. → page 89, 104

Nirenberg S, Carcieri SM, Jacobs AL, Latham PE (2001) Retinal ganglion cells act largely as independent encoders. *Nature* 411:698–701. → page 62

Nowak LG, Azouz R, Sanchez-Vives MV, Gray CM, McCormick DA (2003) Electrophysiological classes of cat primary visual cortical neurons in vivo as revealed by quantitative analyses. *J Neurophysiol* 89:1541–1566. → page 90, 103

Ohiorhenuan IE, Mechler F, Purpura KP, Schmid AM, Hu Q, Victor JD (2010) Sparse coding and high-order correlations in fine-scale cortical networks. *Nature* 466:617–621. → page 154, 178

Ohki K, Chung S, Ch'ng YH, Kara P, Reid RC (2005) Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. *Nature* 433:597–603. → page 4, 87

Ohki K, Chung S, Kara P, Hübener M, Bonhoeffer T, Reid RC (2006) Highly ordered arrangement of single neurons in orientation pinwheels. *Nature* 442:925–928. → page 4

O'Keefe J, Recce ML (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3:317–330. → page 3

Okun M, Yger P, Marguet SL, Gerard-Mercier F, Benucci A, Katzner S, Busse L, Carandini M, Harris KD (2012) Population rate dynamics and multineuron firing patterns in sensory cortex. *J Neurosci* 32:17108–17119. → page 121, 129, 149, 154

Oliphant TE (2006) *Guide to NumPy* Provo, UT. → page 181

Olshausen BA, Field DJ (1996) Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381:607–609. → page 5, 7, 8, 81, 140

Olshausen BA, Field DJ (2000) Vision and the coding of natural images. *Am Sci* 88:238–245. → page 7

Olshausen BA, Field DJ (2004) Sparse coding of sensory inputs. *Curr Opin Neurobiol* 14:481–487. → page 5, 8, 81

Olshausen BA, Field DJ (2005) How close are we to understanding V1? *Neural Comput* 17:1665–1699. → page 3, 7, 81, 82, 110, 141

Oram MW, Wiener MC, Lestienne R, Richmond BJ (1999) Stochastic nature of precisely timed spike patterns in visual system neuronal responses. *J Neurophysiol* 81:3021–3033. → page 5

Pachitariu M, Lyamzin DR, Sahani M, Lesica NA (2015) State-dependent population coding in primary auditory cortex. *J Neurosci* 35:2058–2073. → page 111, 148, 149

Pascoe PJ, Ilkiw JE, Frischmeyer KJ (2006) The effect of the duration of propofol administration on recovery from anesthesia in cats. *Vet Anaesth Analg* 33:2–7. → page 146

Payne B, Peters A (2001) *The Cat Primary Visual Cortex.* Academic Press, San Diego, CA. → page 4, 15

Pedreira C, Martinez J, Ison MJ, Quian Quiroga R (2012) How many neurons can we see with current spike sorting algorithms? *J Neurosci Meth* 211:58–65. → page 59

Pérez F, Granger BE (2007) IPython: a System for Interactive Scientific Computing. *IEEE Comput Sci Eng Mag* 9:21–29. → page 181

Perin R, Berger TK, Markram H (2011) A synaptic organizing principle for cortical neuronal groups. *PNAS* 108:5419–5424. → page 5

Perkel DH, Gerstein GL, Moore GP (1967) Neuronal spike trains and stochastic point processes: II. Simultaneous spike trains. *Biophys J* 7:419–440. → page 90

Petersen CCH, Hahn TTG, Mehta M, Grinvald A, Sakmann B (2003) Interaction of sensory responses with spontaneous depolarization in layer 2/3 barrel cortex. *PNAS* 100:13638–13643. → page 6, 110, 147, 149

Pillow JW, Shlens J, Chichilnisky EJ, Simoncelli EP (2013) A model-based spike sorting algorithm for removing correlation artifacts in multi-neuron recordings. *PLoS ONE* 8:e62123. → page 24, 62

Polikov VS, Tresco PA, Reichert WM (2005) Response of brain tissue to chronically implanted neural electrodes. *J Neurosci Meth* 148:1–18. → page 82

Poulet JFA, Frenandez LMJ, Crochet S, Petersen CCH (2012) Thalamic control of cortical states. *Nat Neurosci* 15:370–372. → page 121

Poulet JFA, Petersen CCH (2008) Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature* 454:881–885. → page 121, 149

Pouzat C, Delescluse M, Viot P, Diebolt J (2004) Improved spike-sorting by modeling firing statistics and burst-dependent spike amplitude attenuation: a markov chain monte carlo approach. *J Neurophysiol* 91:2910–2928. → page 24

Prentice JS, Homann J, Simmons KD, Tkačik G, Balasubramanian V, Nelson PC (2011) Fast, scalable, Bayesian spike identification for multi-electrode arrays. *PLoS ONE* 6:e19884. → page 24, 26, 62

Priebe NJ, Mechler F, Carandini M, Ferster D (2004) The contribution of spike threshold to the dichotomy of cortical simple and complex cells. *Nat Neurosci* 7:1113–1122. → page 90, 107, 143

Quian Quiroga R, Nadasdy Z, Ben-Shaul Y (2004) Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural Comput* 16:1661–1687. → page 24, 27, 30, 43, 45

Reich DS, Victor JD, Knight BW, Ozaki T, Kaplan E (1997) Response variability and timing precision of neuronal spike trains in vivo. *J Neurophysiol* 77:2836–2841. → page 140

Reid RC, Alonso JM (1995) Specificity of monosynaptic connections from thalamus to visual cortex. *Nature* 378:281–283. → page 142

Reid RC, Victor JD, Shapley RM (1997) The use of m-sequences in the analysis of visual neurons: Linear receptive field properties. *Visual Neurosci* 14:1015–1027. → page 19

Reimann MW, Anastassiou CA, Perin R, Hill SL, Markram H, Koch C (2013) A biophysically detailed model of neocortical local field potentials predicts the critical role of active membrane currents. *Neuron* 79:375–390. → page 2

Reinagel P, Reid RC (2000) Temporal coding of visual information in the thalamus. *J Neurosci* 20:5392–400. → page 140

Reinagel P, Reid RC (2002) Precise firing events are conserved across neurons. *J Neurosci* 22:6837–6841. → page 140

Renart A, de la Rocha J, Bartho P, Hollender L, Parga N, Reyes A, Harris KD (2010) The asynchronous state in cortical circuits. *Science* 327:587–590. → page 121, 129, 142, 149, 178

Renshaw B, Forbes A, Morison BR (1940) Activity of isocortex and hippocampus: electrical studies with microelectrodes. *J Neurophysiol* 3:74–105. → page 2

Ringach D, Shapley R (2004) Reverse correlation in neurophysiology. *Cognitive Sci* 28:147–166. → page 89

Ringach DL, Hawken MJ, Shapley R (1997a) Dynamics of orientation tuning in macaque primary visual cortex. *Nature* 387:281–284. → page 18, 75, 108

Ringach DL, Hawken MJ, Shapley R (2002a) Receptive field structure of neurons in monkey primary visual cortex revealed by stimulation with natural image sequences. *J Vision* 2:12–24. → page 90, 107

Ringach DL, Hawken MJ, Shapley R (2003) Dynamics of orientation tuning in macaque V1: the role of global and tuned suppression. *J Neurophysiol* 90:342–352. → page 108

Ringach DL, Sapiro G, Shapley R (1997b) A subspace reverse-correlation technique for the study of visual neurons. *Vision Res* 37:2455–2464. → page 18, 75, 108

Ringach DL, Shapley RM, Hawken MJ (2002b) Orientation selectivity in macaque V1: diversity and laminar dependence. *J Neurosci* 22:5639–5651. → page 87

Robbins AA, Fox SE, Holmes GL, Scott RC, Barry JM (2013) Short duration waveforms recorded extracellularly from freely moving rats are representative of axonal activity. *Front Neural Circuits* 7:181. → page 104

Rodriguez A, Laio A (2014) Clustering by fast search and find of density peaks. *Science* 344:1492–1496. → page 61, 62

Roelfsema PR, Lamme VAF, Spekreijse H (1998) Object-based attention in the primary visual cortex of the macaque monkey. *Nature* 395:376–381. → page 110, 149

Sakata S, Harris KD (2009) Laminar structure of spontaneous and sensory-evoked population activity in auditory cortex. *Neuron* 64:404–418. → page 81, 86, 89, 92, 104, 105, 121, 145, 147

Sakata S, Harris KD (2012) Laminar-dependent effects of cortical state on auditory cortical spontaneous activity. *Front Neural Circuits* 6:109. → page 6, 124, 129, 145, 146

Sakuraba S, Tsujita M, Arisaka H, Takeda J, Yoshida K, Kuwana SI (2009) Donepezil reverses buprenorphine-induced central respiratory depression in anesthetized rabbits. *Biol Res* 42:469–475. → page 124

Salami M, Itami C, Tsumoto T, Kimura F (2003) Change of conduction velocity by regional myelination yields constant latency irrespective of distance between thalamus and cortex. *PNAS* 100:6174–6179. → page 5

Saleem AB, Ayaz A, Jeffery KJ, Harris KD, Carandini M (2013) Integration of visual motion and locomotion in mouse visual cortex. *Nat Neurosci* 16:1864–1869. → page 110

Saleem AB, Chadderton P, Apergis-Schoute J, Harris KD, Schultz SR (2010) Methods for predicting cortical UP and DOWN states from the phase of deep layer local field potentials. *J Comput Neurosci* 29:49–62. → page 121, 129

Sanchez-Vives MV, McCormick DA (2000) Cellular and network mechanisms of rhythmic recurrent activity in neocortex. *Nat Neurosci* 3:1027–1034. → page 110, 126

Sasaki T, Takahashi N, Matsuki N, Ikegaya Y (2008) Fast and accurate detection of action potentials from somatic calcium fluctuations. *J Neurophysiol* 100:1668–1676. → page 4

Schiller PH, Finlay BL, Volman SF (1976a) Quantitative studies of single-cell properties in monkey striate cortex. I. Spatiotemporal organization of receptive fields. *J Neurophysiol* 39:1288–1319. → page 89, 108

Schiller PH, Finlay BL, Volman SF (1976b) Quantitative studies of single-cell properties in monkey striate cortex. II. Orientation specificity and ocular dominance. *J Neurophysiol* 39:1320–1333. → page 87

Schjetnan AGP, Luczak A (2011) Recording large-scale neuronal ensembles with silicon probes in the anesthetized rat. *J Vis Exp* 56:e3282. → page 80, 153

Schneidman E, Berry II MJ, Segev R, Bialek W (2006) Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* 440:1007–1012. → page 52, 154, 178

Schölvinck ML, Saleem AB, Benucci A, Harris KD, Carandini M (2015) Cortical state determines global variability and correlations in visual cortex. *J Neurosci* 35:170–178. → page 149, 178

Scholz M, Gibon Y, Stitt M, Selbig J (2004) Independent component analysis of starch deficient pgm mutants In *Proceedings of the German Conference on Bioinformatics*, pp. 95–104. Gesellschaft für Informatik, Bielefeld, Germany. → page 43

Schwartz O, Pillow JW, Rust NC, Simoncelli EP (2006) Spike-triggered neural characterization. *J Vision* 6:484–507. → page 90

Segev R, Goodhouse J, Puchalla J, Berry II MJ (2004) Recording spikes from a large fraction of the ganglion cells in a retinal patch. *Nat Neurosci* 7:1155–1162. → page 5, 24, 26, 62

Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *J Neurosci* 18:3870–3896. → page 5

Shapley R, Reid RC, Soodak R (1991) Spatiotemporal receptive fields and direction selectivity In Landy MS, Movshon JA, editors, *Computational Models of Visual Processing*, pp. 109–118. MIT Press, Cambridge, MA. → page 19

Sharpee T, Rust NC, Bialek W (2004) Analyzing neural responses to natural signals: maximally informative dimensions. *Neural Comput* 16:223–250. → page 90, 107

Shlens J (2009) A tutorial on principal component analysis. *Salk Insitute for Biological Studies* . → page 41

Shlens J, Field GD, Gauthier JL, Grivich MI, Petrusca D, Sher A, Litke AM, Chichilnisky EJ (2006) The structure of multi-neuron firing patterns in primate retina. *J Neurosci* 26:8254–66. → page 154, 178

Shoham S, Fellows MR, Normann RA (2003) Robust, automatic spike sorting using mixtures of multivariate *t*-distributions. *J Neurosci Meth* 127:111–122. → page 24

Shoham S, O'Connor DH, Segev R (2006) How silent is the brain: is there a dark matter problem in neuroscience? *J Comp Physiol A* 192:777–784. → page 3, 22, 80

Shuler MG, Bear MF (2006) Reward timing in the primary visual cortex. *Science* 311:1606–1609. → page 110

Sirota A, Montgomery S, Fujisawa S, Isomura Y, Zugaro M, Buzsáki G (2008) Entrainment of neocortical neurons and gamma oscillations by the hippocampal theta rhythm. *Neuron* 60:683–697. → page 80, 89, 92, 104, 105

Skottun BC, De Valois RL, Grosof DH, Movshon JA, Albrecht DG, Bonds AB (1991) Classifying simple and complex cells on the basis of response modulation. *Vision Res* 31:1078–1086. → page 106, 107

Smyth D, Willmore B, Baker GE, Thompson ID, Tolhurst DJ (2003) The receptive-field organization of simple cells in primary visual cortex of ferrets under natural scene stimulation. *J Neurosci* 23:4746–4759. → page 90, 107

Softky WR (1995) Simple codes versus efficient codes. *Curr Opin Neurobiol* 5:239–247. → page 5

Spacek M, Swindale N (2009) Python in neuroscience. *Neuromorphic Engineer* . → page 29

Spacek MA, Blanche TJ, Seamans JK, Swindale NV (2007) Accounting for network states in cortex: are (local) pairwise correlations sufficient? *Soc Neurosci Abstr* 33:790.1. → page 178

Spacek MA, Blanche TJ, Swindale NV (2009) Python for large-scale electrophysiology. *Front Neuroinform* 2:9. → page iii, 11, 17, 29, 179

Spacek MA, Swindale NV (2011) Local pairwise correlations and network states in cat primary visual cortex. *Soc Neurosci Abstr* 41:536.06. → page 178

Spacek MA, Swindale NV (2012) Local pairwise correlations and network states in cat primary visual cortex. *Canadian Assoc Neurosci Abstr* 6:3–D–66. → page 178

Spruston N, Johnston D (1992) Perforated patch-clamp analysis of the passive membrane properties of three classes of hippocampal neurons. *J Neurophysiol* 67:508–529. → page 114

Steriade M, Contreras D, Dossi RC, Nuñez A (1993c) The slow (< 1 Hz) oscillation in reticular thalamic and thalamocortical neurons: scenario of sleep rhythm generation in interacting thalamic and neocortical networks. *J Neurosci* 13:3284–3299. → page 150

Steriade M, Dossi RC, Paré D, Oakson G (1991) Fast oscillations (20-40 Hz) in thalamocortical systems and their potentiation by mesopontine cholinergic nuclei in the cat. *PNAS* 88:4396–4400. → page 150

Steriade M, Nuñez A, Amzica F (1993a) A novel slow (< 1 Hz) oscillation of neocortical neurons in vivo: depolarizing and hyperpolarizing components. *J Neurosci* 13:3252–3265. → page 110, 126, 147

Stevenson IH, Kording KP (2011) How advances in neural recording affect data analysis. *Nat Neurosci* 14:139–142. → page 57, 154

Stosiek C, Garaschuk O, Holthoff K, Konnerth A (2003) In vivo two-photon calcium imaging of neuronal networks. *PNAS* 100:7319–7324. → page 4

Straw AD (2008) Vision Egg: an open-source library for realtime visual stimulus generation. *Front Neuroinform* 2:4. → page 17, 179

Straw AD, Warrant EJ, O'Carroll DC (2006) A 'bright zone' in male hoverfly (Eristalis tenax) eyes and associated faster motion detection and increased contrast sensitivity. *J Exp Biol* 209:4339–4354. → page 17, 179

Svoboda K, Denk W, Kleinfeld D, Tank DW (1997) In vivo dendritic calcium dynamics in neocortical pyramidal neurons. *Nature* 385:161–165. → page 4

Swindale NV (1998) Orientation tuning curves: empirical description and estimation of parameters. *Biol Cybern* 78:45–56. → page 76

Swindale NV (2003) Neural synchrony, axonal path lengths, and general anesthesia: A hypothesis. *The Neuroscientist* 9:440–445. → page 6

Swindale NV, Spacek MA (2012) Response variability in visual cortex. *Soc Neurosci Abstr* 42:571.23. → page 83, 99

Swindale NV, Spacek MA (2014) Spike sorting for polytrodes: a divide and conquer approach. *Front Syst Neurosci* 8:6. → page iii, 8, 28, 29, 36, 45, 48, 53, 57, 60

Swindale NV, Spacek MA (2015) Spike detection methods for polytrodes and high density microelectrode arrays. *J Comput Neurosci* 38:249–261. → page 23, 29, 55, 58

Szarowski DH, Andersen MD, Retterer S, Spence AJ, Isaacson M, Craighead HG, Turner JN, Shain W (2003) Brain responses to micro-machined silicon devices. *Brain Res* 983:23–35. → page 82

Tao L, Shelley M, McLaughlin D, Shapley R (2004) An egalitarian network model for the emergence of simple and complex cells in visual cortex. *PNAS* 101:366–371. → page 108, 143

Thompson LT, Best PJ (1989) Place cells and silent cells in the hippocampus of freely-behaving rats. *J Neurosci* 9:2382–2390. → page 80

Thorpe S, Fize D, Marlot C (1996) Speed of processing in the human visual system. *Nature* 381:520–522. → page 5

Touryan J, Felsen G, Dan Y (2005) Spatial structure of complex cell receptive fields measured with natural images. *Neuron* 45:781–791. → page 90, 107

Toyama K, Kimura M, Tanaka K (1981a) Cross-correlation analysis of interneuronal connectivity in cat visual cortex. *J Neurophysiol* 46:191–201. → page 143

Toyama K, Kimura M, Tanaka K (1981b) Organization of cat visual cortex as investigated by cross-correlation technique. *J Neurophysiol* 46:202–14. → page 143

Vaadia E, Haalman I, Abeles M, Bergman H, Prut Y, Slovin H, Aertsen A (1995) Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature* 373:515–518. → page 5

VanRullen R, Thorpe SJ (2002) Surfing a spike wave down the ventral stream. *Vision Res* 42:2593–2615. → page 5

Ventura V, Gerkin RC (2012) Accurately estimating neuronal correlation requires a new spike-sorting paradigm. *PNAS* 109:7230–7235. → page 22

Vigneswaran G, Kraskov A, Lemon RN (2011) Large identified pyramidal cells in macaque motor and premotor cortex exhibit “thin spikes”: implications for cell type classification. *J Neurosci* 31:14235–14242. → page 104

Vincent JL, Patel GH, Fox MD, Snyder AZ, Baker JT, Van Essen DC, Zempel JM, Snyder LH, Corbetta M, Raichle ME (2007) Intrinsic functional architecture in the anaesthetized monkey brain. *Nature* 447:83–86. → page 6

Vinje WE, Gallant JL (2000) Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* 287:1273–1276. → page 7, 134, 137, 140, 148

Vinje WE, Gallant JL (2002) Natural stimulation of the nonclassical receptive field increases information transmission efficiency in V1. *J Neurosci* 22:2904–2915. → page 7, 137, 140, 148

Voigts J, Siegle JH, Pritchett DL, Moore CI (2013) The flexDrive: an ultra-light implant for optical control and highly parallel chronic recording of neuronal ensembles in freely moving mice. *Front Syst Neurosci* 7:8. → page 3

Vyazovskiy VV, Harris KD (2013) Sleep and the single neuron: the role of global slow oscillations in individual cell rest. *Nat Rev Neurosci* 14:443–451. → page 110

Wallace DJ, zum Alten Borgloh SM, Astori S, Yang Y, Bausen M, Kügler S, Palmer AE, Tsien RY, Sprengel R, Kerr JND, Denk W, Hasan MT (2008) Single-spike detection in vitro and in vivo with a genetic Ca^{2+} sensor. *Nat Methods* 5:797–804. → page 4

Wang HP, Spencer D, Fellous JM, Sejnowski TJ (2010) Synchrony of thalamocortical inputs maximizes cortical reliability. *Science* 328:106–109. → page 140

Wang WJ, Tan YX, Jiang JH, Lu JZ, Shen GL, Yu RQ (2004) Clustering based on kernel density estimation: nearest local maximum searching algorithm. *Chemometr Intell Lab* 72:1–8. → page 61

Wehr M, Zador AM (2003) Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature* 426:442–446. → page 5

Wilkie D (1983) Rayleigh test for randomness of circular data. *Appl Statist* 32:311–312. → page 76

Williams PE, Mechler F, Gordon J, Shapley R, Hawken MJ (2004) Entrainment to video displays in primary visual cortex of macaque and humans. *J Neurosci* 24:8278–8288. → page 17, 18

Willmore B, Smyth D (2003) Methods for first-order kernel estimation: simple-cell receptive fields from responses to natural scenes. *Network* 14:553–577. → page 90, 107

Wilson FA, O'Scalaidhe SP, Goldman-Rakic PS (1994) Functional synergism between putative γ-aminobutyrate-containing neurons and pyramidal neurons in prefrontal cortex. *PNAS* 91:4009–4013. → page 89, 104

Wilson MA, McNaughton BL (1993) Dynamics of the hippocampal ensemble code for space. *Science* 261:1055–1058. → page 3

Wiltschko AB, Gage GJ, Berke JD (2008) Wavelet filtering before spike detection preserves waveform shape and enhances single-unit discrimination. *J Neurosci Meth* 173:34–40. → page 105

Wohrer A, Humphries MD, Machens CK (2013) Population-wide distributions of neural activity during perceptual decision-making. *Prog Neurobiol* 103:156–193. → page 81, 82

Wolf MT, Burdick JW (2009) A Bayesian clustering method for tracking neural signals over successive intervals. *IEEE Trans Biomed Eng* 56:2649–2659. → page 24, 27

Wood F, Black MJ (2008) A nonparametric Bayesian alternative to spike sorting. *J Neurosci Meth* 173:1–12. → page 24

Wood PL, Rackham A (1981) Actions of kappa, sigma and partial mu narcotic receptor agonists on rat brain acetylcholine turnover. *Neurosci Lett* 23:75–80. → page 124

Wörgötter F, Suder K, Zhao Y, Kerscher N, Eysel UT, Funke K (1998) State-dependent receptive-field restructuring in the visual cortex. *Nature* 396:165–168. → page 146

Wright WE (1977) Gravitational clustering. *Pattern Recogn* 9:151–166. → page 61

Xie L, Kang H, Xu Q, Chen MJ, Liao Y, Thiyagarajan M, O'Donnell J, Christensen DJ, Nicholson C, Iliff JJ, Takano T, Deane R, Nedergaard M (2013) Sleep drives metabolite clearance from the adult brain. *Science* 342:373–377. → page 86, 110

Xu S, Jiang W, Poo MM, Dan Y (2012) Activity recall in a visual cortical ensemble. *Nat Neurosci* 15:449–455. → page 6

Xu W, Huang XT K, Wu J (2007) Compression and reflection of visually evoked cortical waves. *Neuron* 55:119–129. → page 149

Yao H, Shi L, Han F, Gao H, Dan Y (2007) Rapid learning in cortical coding of visual scenes. *Nat Neurosci* 10:772–778. → page 20

Yen SC, Baker J, Gray CM (2007) Heterogeneity in the responses of adjacent neurons to natural stimuli in cat striate cortex. *J Neurophysiol* 97:1326–1341. → page 7, 137, 140, 148

Yu J, Ferster D (2013) Functional coupling from simple to complex cells in the visually driven cortical circuit. *J Neurosci* 33:18855–18866. → page 143

Yu S, Huang D, Singer W, Nikolić D (2008) A small world of neuronal synchrony. *Cerebral Cortex* 18:2891–2901. → page 154, 178

Zagha E, Casale AE, Sachdev RNS, McGinley MJ, McCormick DA (2013) Motor cortex feedback influences sensory processing by modulating network state. *Neuron* 79:567–578. → page 111, 148

Zito T, Wilbert N, Wiskott L, Berkes P (2009) Modular toolkit for data processing (MDP): a Python data processing framework. *Front Neuroinform* 2:8. → page 181

Zouridakis G, Tam DC (1997) Multi-unit spike discrimination using wavelet transforms. *Comput Biol Med* 27:9–18. → page 24

Zouridakis G, Tam DC (2000) Identification of reliable spike templates in multi-unit extracellular recordings using fuzzy clustering. *Comput Meth Prog Bio* 61:91–98. → page 24

Appendix A Receptive field stability

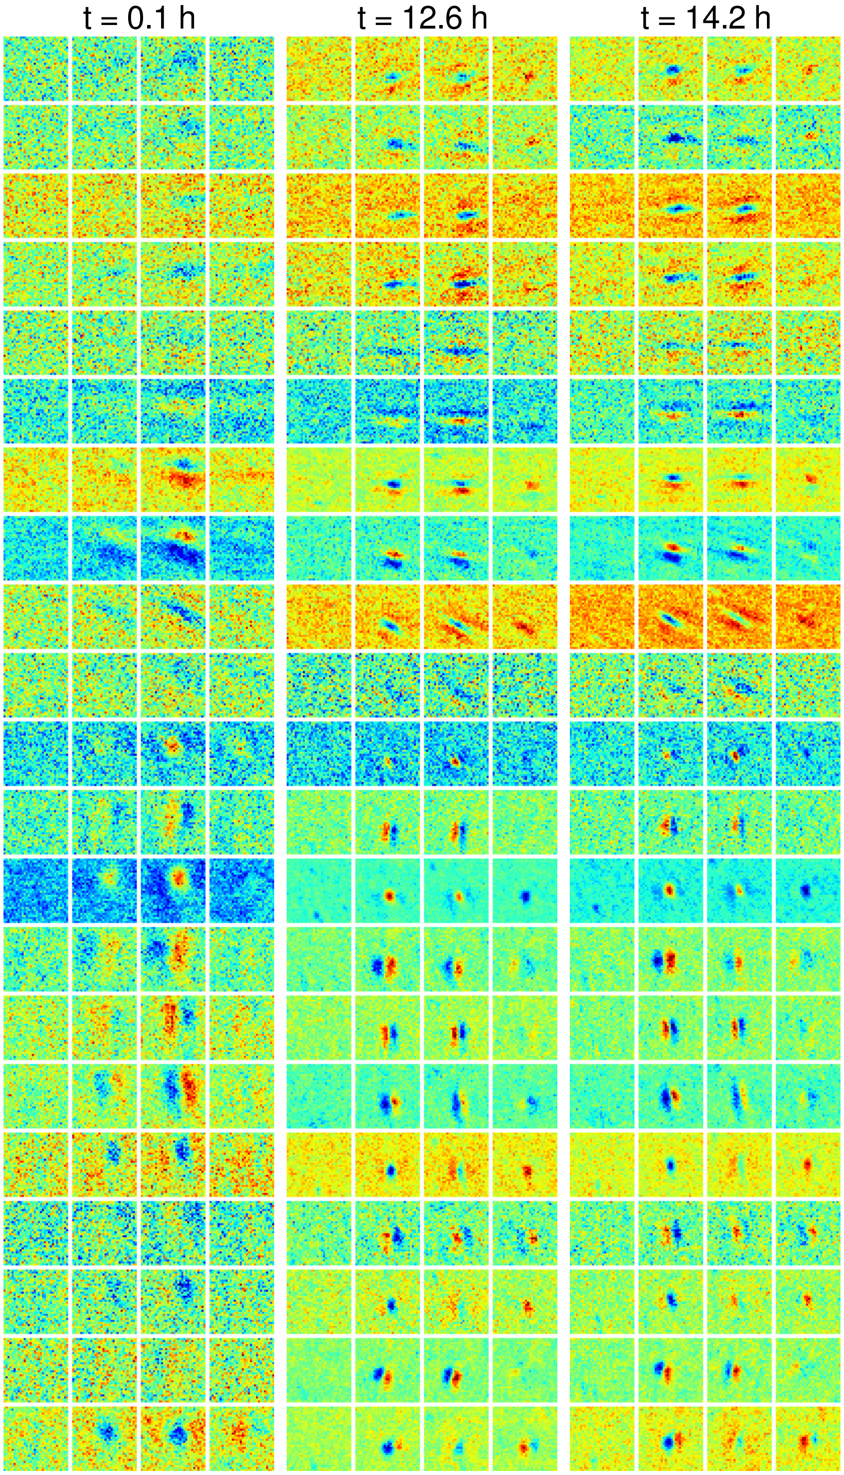

Figure A.1: Long duration RF stability demonstrated by simple cells with non-zero STAs. The layout is similar to Figure 5.6, but here 3 recordings (***left***, ***middle***, ***right***) from the same track (ptc15.tr7c) are shown, and each row corresponds to the same neuron for all 3 recordings. Only the 21 cells with non-zero STAs for all three recordings are shown. Each recording began at the time shown and lasted 22 min. The first 4 reverse correlation time ranges, each 20 ms in duration, are shown for each recording. The size of the m-sequence movie was increased between the first and second recordings (from 6.4° to 12.7° across) and the onscreen position of the movie was shifted by 2.9° up and to the right to better center the RFs. The physical screen position was shifted slightly between the second and third recordings. Visually discounting these shifts in size and position, the simple cell RFs shown remained reasonably stable.

Appendix B Spike correlations

Spike correlations of active neurons were measured by binning and digitizing their spike trains (Schneidman et al., 2006; Shlens et al., 2006), resulting in binary code trains of 0s and 1s (Figure B.1). Correlation coefficients were calculated from these using Pearson's correlation (Equation 6.1), where x and y were the binary code trains of two neurons. Bins were 20 ms wide.

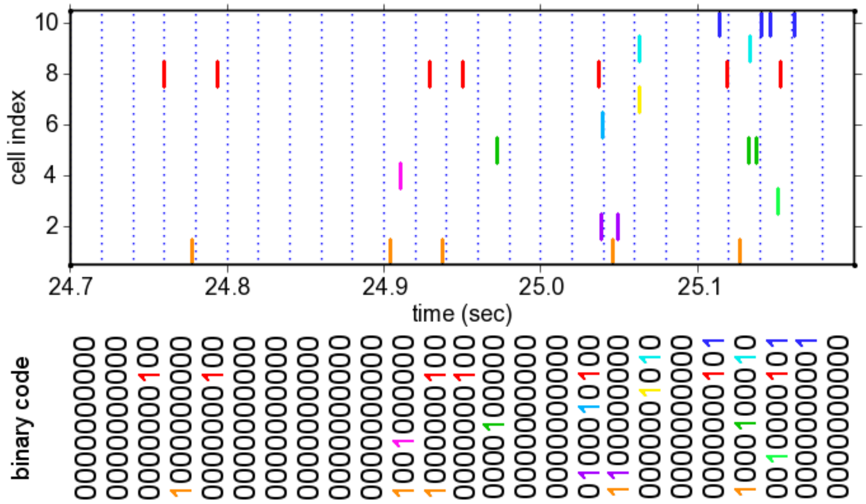

Figure B.1: Spike trains for each neuron (***top***) are binned at 20 ms and digitized to create a temporal binary code (***rows, lower panel***). If a neuron spikes one or more times within a time bin, its activity is represented by a 1 for the duration of that bin. Otherwise, its activity is represented by a 0. Taking the binary codes of all the neurons together, the spiking state of the network at any time bin can be represented as a binary word (***columns, lower panel***).

Mean spike correlations in local neural populations in anesthetized cat V1 were very weak but positively skewed ($\bar{\rho} < 0.025$) (Figure 3.1; Spacek et al., 2007; Spacek and Swindale, 2011, 2012), in line with other reports in V1 (Yu et al., 2008; Ohiorhenuan et al., 2010), S1 and A1 (Renart et al., 2010), and in RGCs (Schneidman et al., 2006). In addition to their sensitivity to spike sorting quality, Figure 3.1 also suggests that, contrary to Schneidman et al. (2006), spike correlations may be stimulus dependent. They may also be very sensitive to cortical state (Section 6.4; Renart et al., 2010; Harris and Thiele, 2011; Schölvinck et al., 2015), but this was not explored here.

This method of calculating spike correlations is similar but not completely equivalent to taking the cross-correlogram of a pair of cells and counting the fraction of spikes coincident within the time range of interest. The difference is that cross-correlograms are temporally aligned to the spikes of one of the cells, while the above binary code method with its uniformly spaced time bins is not.

Appendix C Software

The work presented in this thesis required the development of 3 major software projects (Spacek et al., 2009): dimstim, for controlling visual stimuli; spyke, for waveform visualization and spike sorting; and neuropy, for spike, LFP and stimulus data analysis. The source code and version control history of all 3 are freely available online. Each project is briefly described here.

C.1 dimstim

Dimstim (http://dimstim.github.io) is a Python package for presenting multidimensional visual stimuli with high temporal precision. It relies on the Vision Egg library (http://www.visionegg.org; Straw et al., 2006; Straw, 2008) to generate visual stimuli. Using strictly formatted Python scripts (Figure C.1), dimstim allows the experimenter to specify exactly how each stimulus parameter changes in time, and these details are sent from the stimulus computer to the acquisition computer (Figure 2.1) in real time.

If certain parameters (e.g., orientation, spatial and temporal frequency of a drifting grating) are to vary independently across trials during an experiment, each is assigned to its own stimulus "dimension" in the experiment script (Figure C.1). Each dimension can be of different length (e.g., 12 orientations, 6 spatial and 4 temporal frequencies), and when the experiment is run, all possible combinations of the provided stimulus parameters (e.g., 288) will be presented in the specified order. If a set of stimulus parameters all have the same number of values, they can be made to covary by assigning them to the same dimension. To minimize response adaptation (Maffei et al., 1973), the order of parameter values can be shuffled independently within each dimension.

Scripts are executed individually on the stimulus computer from the command line. Execution triggers the acquisition computer to begin saving data to disk. Immediately before the stimulus presentation (experiment) begins, dimstim sends a copy of the experiment script from the stimulus computer to the acquisition computer as a detailed record of the experiment. Each particular combination of stimulus parameter values is represented by a unique number, a stimulus "sweep" index. This index is communicated to the acquisition computer using a digital output board (DT340, Data Translations, Marlboro, MA) on every screen raster (5 ms for a 200 Hz refresh rate screen). The screen raster, a pulse < 100 μs in duration, is used as a trigger for the acquisition computer, and its arrival indicates the precise time at which the screen is updated. The sweep index is saved by the acquisition computer using the same clock as the spike and LFP data (Blanche, 2005). This allows for precise correspondence between stimulus and response. To ensure that no rasters are missed due to an issue with either stimulus or acquisition, a checksum is calculated independently by the stimulus and acquisition computers from all of the sweep indices. The results are compared at the end of the experiment to ensure that they are identical.

```
# Static parameters remain constant during the experiment
s = StaticParams()
# pre-experiment duration to display blank screen (sec)
s.preexpSec = 1
# post-experiment duration to display blank screen (sec)
s.postexpSec = 1
# x coord of stimulus center relative to screen center (deg)
s.xorigDeg = dc.get("Manbar0", "xorigDeg")
# y coord of stimulus center relative to screen center (deg)
s.yorigDeg = dc.get("Manbar0", "yorigDeg")

# Dynamic parameters may vary from one sweep to the next
d = DynamicParams()
# grating orientation (deg)
d.ori = range(0, 360, 30)
# spatial frequency (cycles/deg)
d.sfreqCycDeg = [0.05, 0.1, 0.2, 0.4, 0.8, 1.6]
# temporal frequency (cycles/sec)
d.tfreqCycSec = [0.5, 1, 2, 5]
# grating phase to begin each sweep with (+/- deg)
d.phase0 = 0
# mean luminance (0-1)
d.ml = 0.5
# contrast (0-1), >> 1 gives square grating, < 0 reverses contrast
d.contrast = 1
# sweep duration (sec)
d.sweepSec = 6

# Assign dynamic parameters to stimulus dimensions
vs = Variables()
vs.ori = Variable(vals=d.ori, dim=0, shuffle=True) # dim is dimension number
vs.sfreqCycDeg = Variable(vals=d.sfreqCycDeg, dim=1, shuffle=True)
vs.tfreqCycSec = Variable(vals=d.tfreqCycSec, dim=2, shuffle=True)

runs = Runs(n=1, reshuffle=False) # run parameter combinations n times
bs = BlankSweeps(T=20, sec=2) # blank sweep every T sweeps for sec seconds
e = Grating(script=__file__, static=s, dynamic=d, variables=vs,
            runs=runs, blanksweeps=bs) # create a Grating experiment
e.run() # run the experiment
```

Figure C.1: Example dimstim Python script for a drifting grating experiment. In this example, grating orientation, spatial and temporal frequency are the only 3 parameters assigned multiple values, and therefore the only ones to vary from one stimulus “sweep” (6 s in duration) to the next. Each varies independently because each is set to a different dimension (second code block from bottom). This script can be modified into a flashed grating experiment by setting the orientation (`ori`) to vary from 0 to 180° instead of 0 to 360°, setting the temporal frequency (`tfreqCycSec`) to 0, setting the phase (`phase0`) to a range of values in degrees, and decreasing the sweep duration (`sweepSec`). The stimulus is centered on-screen according to the position previously set by a manually controlled oriented bar (`"Manbar0"`) stimulus (Section 2.4). Some initialization code and stimulus parameters are omitted for brevity. Comments are italicized.

C.2 spyke

Spyke (http://spyke.github.io) is a Python application for visualizing, navigating, and spike sorting high-density polytrode waveform data. Spyke relies on several free cross-platform libraries: NumPy (Oliphant, 2006) for large numerical arrays; SciPy (Jones et al., 2001) and MDP (Zito et al., 2009) for certain algorithms; Cython (Behnel et al., 2011) for fast low-level code; IPython (Pérez and Granger, 2007) for an interactive shell; Qt (http://qt-project.org) for the GUI; matplotlib (Hunter, 2007) for 2D plotting; and PyOpenGL (http://pyopengl.sourceforge.net) for 3D plotting.

Screenshots of spyke are shown in Figures C.2 and C.3. A brief tutorial is available online. Some functionality was inherited from Tim Blanche's Delphi program "SurfBawd". Spyke's GUI is similar to that of "Klusters" (http://neurosuite.sourceforge.net, Hazan et al., 2006) and "KlustaViewa" (http://klusta-team.github.io), but the scope and some of the methods are quite different.

Memory usage was a major design limitation. An entire track could consist of over 100 GB of data. Current consumer-grade computers can accept a maximum of 32 GB of RAM, so loading all of this data into memory at once was not possible. Spyke works around this memory limitation by only loading the raw waveform data when required from disk. Disk access is slow, but fortunately on the fly loading (and interpolation, Section 3.3) need only happen during spike detection (Section 3.4) and major realignment (Section 3.7). Since spike detection is fully automated, and major realignment is only rarely required, on the fly loading rarely slows the user down. Once spike detection is complete, the detected spike waveforms are small enough to all hold in memory at once, allowing subsequent spike sorting steps to take place at maximum speed. The longest sorted track (and the longest recorded track, Table 2.2) was ptc15.tr7c, which had $\sim$ 7M detected events (Table 4.1). Each multichannel spike waveform (interpolated to 50 kHz) required about 1 kB of memory, resulting in 7 GB of memory usage in total. Other program requirements often put total memory usage above 8 GB. Therefore, all sorting presented here could have been performed on a computer with 12 GB of RAM, although most was performed with 32 GB.

GAC (Section 3.9.1), NDsep (Section 3.10.2), best fit realignment (Section 3.7), and calculation of peak sharpness (Section 3.4) all required code loops too complex to be vectorized. For computationally intensive tasks, Python without vectorization is very slow compared to compiled languages like C or Fortran. Cython (Behnel et al., 2011) works around this limitation by providing a syntax very similar to Python. On program execution, Cython code is automatically translated into C code and compiled. This makes it much easier for the programmer to convert computationally intensive parts of Python code into C. GAC, NDsep, best fit realignment, and calculation of peak sharpness were therefore all written in Cython.

GAC and NDsep were also parallelized for further acceleration on multicore CPUs. Code was parallelized using multithreading within Cython, which relied on OpenMP (http://openmp.com, Dagum and Menon, 1998). For GAC, the gradient ascent step moved each scout point independently, which allowed the position of multiple scout points to be updated in parallel. However, due

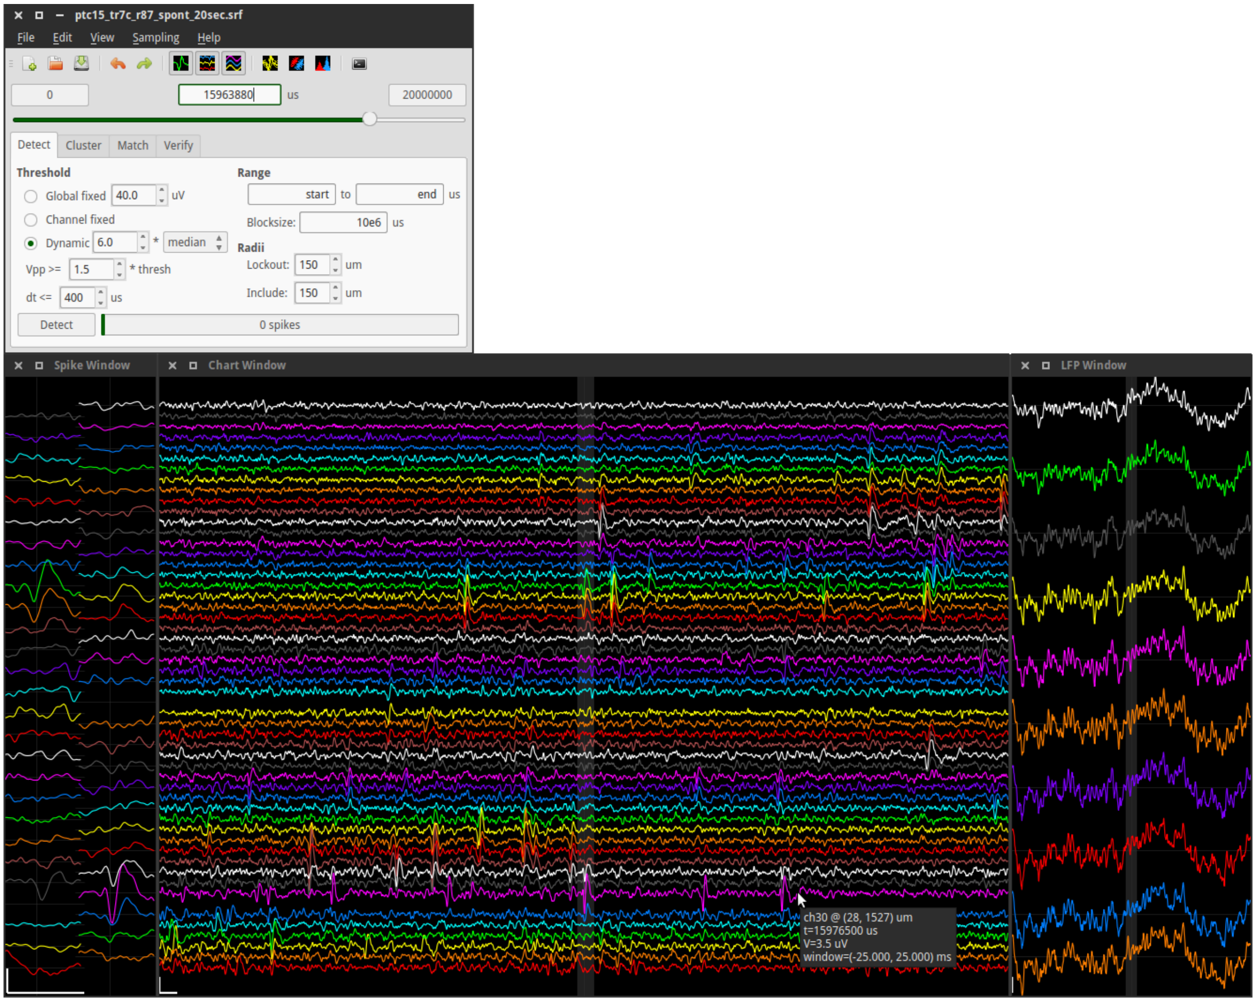

0 — 15963880 us — 20000000

Threshold

○ Global fixed 40.0 uV

○ Channel fixed

● Dynamic 6.0 * median

Vpp >= 1.5 * thresh

dt <= 400 us

Range

start to end us

Blocksize: 10e6 us

Radii

Lockout: 150 um

Include: 150 um

Detect — 0 spikes

ch30 @ (28, 1527) um
t=15976500 us
V=3.5 uV
window=(-25.000, 25.000) ms

Figure C.2: Visualizing time series voltage waveforms with spyke. ***Top left***: The main program window, with the spike detection tab currently selected. Below that are 3 windows displaying voltage waveforms. The first two display high-pass spike data in two different layouts. The spike window (***left***) displays 1 ms of data in a 2D spatial layout, and the chart window (***middle***) displays 50 ms in a 1D vertical layout. The third window (***right***) displays 1 s of low-pass LFP data in a 1D vertical layout. All windows can be resized at will. Vertical positions of channels correspond to each other in all 3 windows. The central highlighted area in the chart window corresponds to the 1 ms time range in the spike window, and the central highlighted area in the LFP window corresponds to the 50 ms time range in the chart window. Scrolling through the data in the main window, or clicking a time point on which to center a waveform window, updates all 3 waveform windows. Channels are enabled or disabled by right-clicking on them. Pointing at a channel displays information such as the channel number and its physical position on the polytrode. After spike detection is complete, raster lines are superimposed on the time series data (Figure 3.4). This is the same time range of spontaneous data shown in Figure 2.3, taken from a 20 s sample file available online. Scale bar: 1 ms, 100 μV.

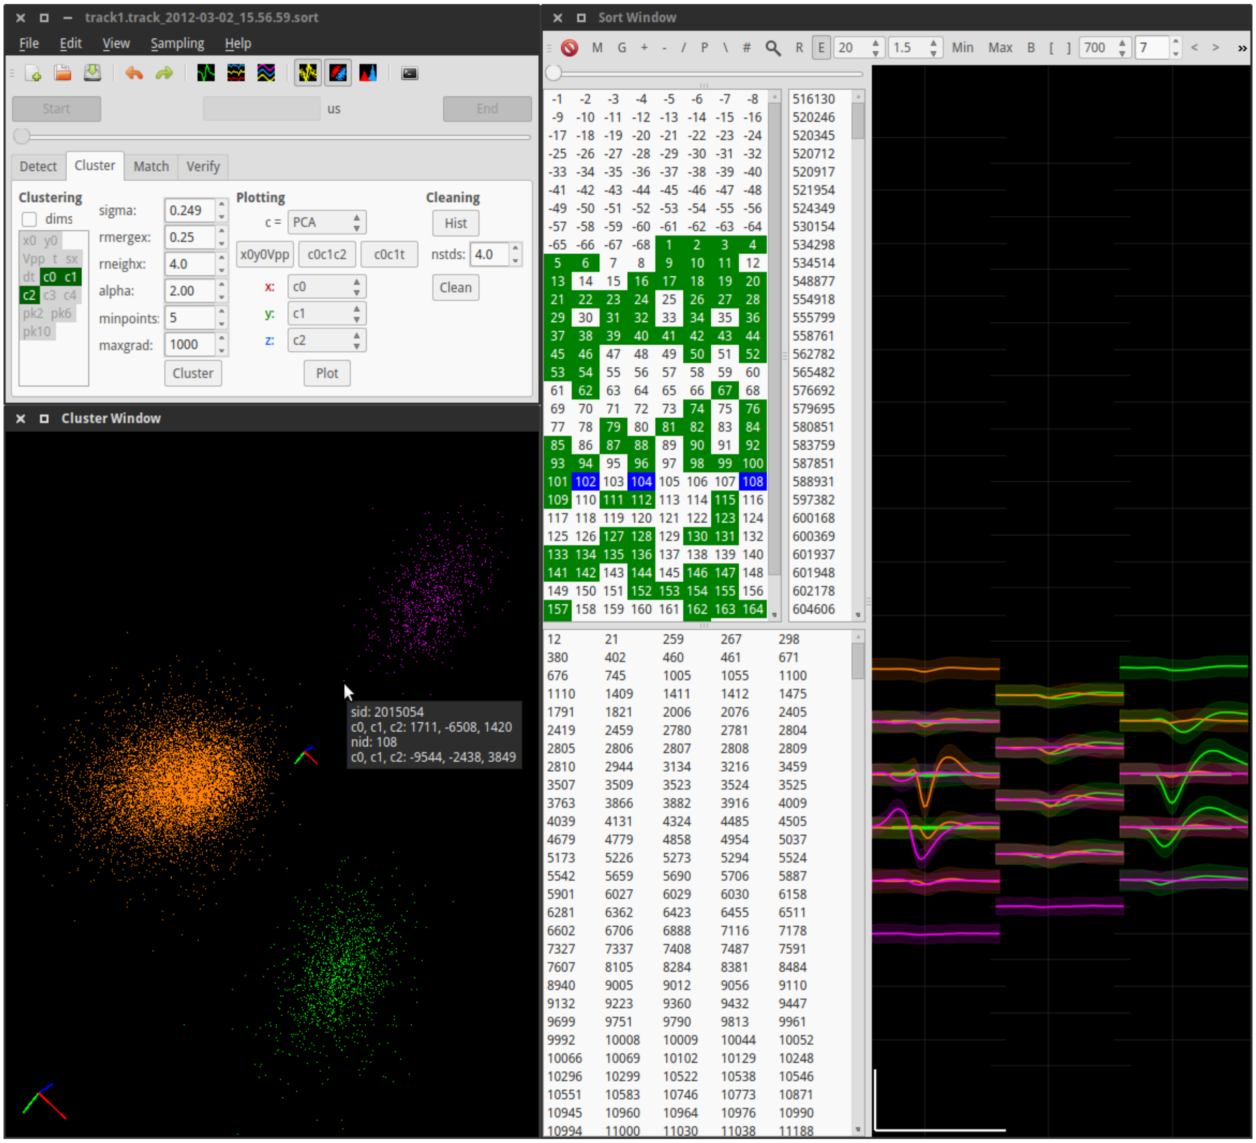

☐ dims

sigma: 0.249

rmergex: 0.25

rneighx: 4.0

alpha: 2.00

minpoints: 5

maxgrad: 1000

c = PCA

x: c0

y: c1

z: c2

nstds: 4.0

20 1.5 700 7

sid: 2015054
c0, c1, c2: 1711, -6508, 1420
nid: 108
c0, c1, c2: -9544, -2438, 3849

```
 -1  -2  -3  -4  -5  -6  -7  -8
 -9 -10 -11 -12 -13 -14 -15 -16
-17 -18 -19 -20 -21 -22 -23 -24
-25 -26 -27 -28 -29 -30 -31 -32
-33 -34 -35 -36 -37 -38 -39 -40
-41 -42 -43 -44 -45 -46 -47 -48
-49 -50 -51 -52 -53 -54 -55 -56
-57 -58 -59 -60 -61 -62 -63 -64
-65 -66 -67 -68   1   2   3   4
  5   6   7   8   9  10  11  12
 13  14  15  16  17  18  19  20
 21  22  23  24  25  26  27  28
 29  30  31  32  33  34  35  36
 37  38  39  40  41  42  43  44
 45  46  47  48  49  50  51  52
 53  54  55  56  57  58  59  60
 61  62  63  64  65  66  67  68
 69  70  71  72  73  74  75  76
 77  78  79  80  81  82  83  84
 85  86  87  88  89  90  91  92
 93  94  95  96  97  98  99 100
101 102 103 104 105 106 107 108
109 110 111 112 113 114 115 116
117 118 119 120 121 122 123 124
125 126 127 128 129 130 131 132
133 134 135 136 137 138 139 140
141 142 143 144 145 146 147 148
149 150 151 152 153 154 155 156
157 158 159 160 161 162 163 164
```

```
516130
520246
520345
520712
520917
521954
524349
530154
534298
534514
548877
554918
555799
558761
562782
565482
576692
579695
580851
583759
587851
588931
597382
600168
600369
601937
601948
602178
604606
```

```
12     21     259    267    298
380    402    460    461    671
676    745    1005   1055   1100
1110   1409   1411   1412   1475
1791   1821   2006   2076   2405
2419   2459   2780   2781   2804
2805   2806   2807   2808   2809
2810   2944   3134   3216   3459
3507   3509   3523   3524   3525
3763   3866   3882   3916   4009
4039   4131   4324   4485   4505
4679   4779   4858   4954   5037
5173   5226   5273   5294   5524
5542   5659   5690   5706   5887
5901   6027   6029   6030   6158
6281   6362   6423   6455   6511
6602   6706   6888   7116   7178
7327   7337   7408   7487   7591
7607   8105   8284   8381   8484
8940   9005   9012   9056   9110
9132   9223   9360   9432   9447
9699   9751   9790   9813   9961
9992   10008  10009  10044  10052
10066  10069  10102  10129  10248
10296  10299  10522  10538  10546
10551  10583  10746  10773  10871
10945  10960  10964  10976  10990
10994  11000  11030  11038  11188
```

Figure C.3: Sorting detected spikes with spyke. ***Top left***: The main window, with the clustering tab selected. Dimension reduction and clustering controls are visible. ***Bottom left***: The cluster window showing three clusters (orange, green, magenta). Each point represents a spike in 3D cluster space (PCA space in this example). The tooltip shows information about one particular spike and its parent (magenta) cluster. Using the mouse, the view can be rotated, zoomed, and panned, and spikes and clusters can be selected. The value of σ_c (Section 3.9.1) can also be manipulated directly in this window using the mouse, and the central axes scale accordingly. ***Right***: The sort window is divided into two areas: a set of selection lists on the left, and a waveform display on the right. Clockwise from top left, the lists represent clusters, spikes belonging to the selected clusters, and unsorted spikes. Clusters marked as "good" (Section 3.10) are highlighted in the cluster list in green, and selected clusters are highlighted in blue. The 3 selected numbers are the IDs of the 3 clusters displayed in the cluster window, and their corresponding mean waveforms (± 2 standard deviations) are shown in the waveform display. Selected channels and spike time ranges are represented by green horizontal lines underneath the waveforms (partly obscured). If specific spikes are selected, their waveforms are also plotted. Scale bar: 1 ms, 100 μV.

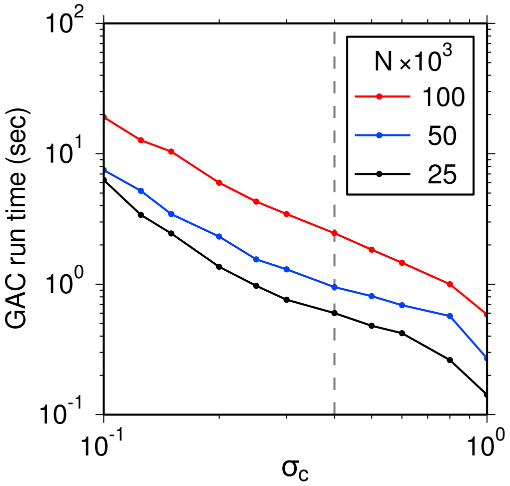

Figure C.4: GAC (Section 3.9.1) run time as a function of σ_c and the number of points to cluster, N. The vertical dashed grey line indicates a commonly used value of $\sigma_c = 0.4$. At that value, clustering 25,000 points took $\sim$ 0.5 s, 50,000 points took $\sim$ 1 s, and 100,000 points took $\sim$ 2.5 s. Benchmarking was performed on a first generation quad-core i7 laptop. The absolute run time at a given σ_c and N depended on the particular distribution of points in cluster space. The points for this figure were randomly sampled from a "good" cluster with 155,000 spikes in it, projected into 3D PCA space.

to the potential all-to-all interaction of points during the merge step, the merge step could not be parallelized.

Several other techniques were used to accelerate GAC. Points were rescaled for use with a pre-calculated lookup table to minimize computationally expensive division and exponentiation operations. The first step executed was always the merge step, in order to reduce the number of scout points before the first update step. Scout points that moved by a distance less than $10^{-5}\sigma_c$ on any iteration were deemed stationary and were skipped on subsequent update steps. Points were sorted ahead of time along the first clustering dimension, typically the one with the greatest variance (e.g., the first PC if clustering in PCA space). This allowed local neighbourhoods of points to be defined more quickly for the local gradient calculation. Finally, in regions of high density ($>$ 1000 points in the local neighbourhood of a scout point), points were subsampled such that no more than 1000 were used to calculate the local gradient.

The performance of spyke's implementation of GAC as a function of σ_c and N is shown in Figure C.4. Clustering $N = 100,000$ points typically took $\sim$ 2.5 s on a consumer laptop, but most clustering runs had fewer points than this. The implementation was therefore sufficiently fast for interactive use.

C.3 neuropy

Neuropy (http://neuropy.github.io) is a command-line based Python program for interactive analysis of spike, LFP, and stimulus data. It relies on the IPython and matplotlib libraries for its MATLAB-like interface (Figure C.5A). Both on disk and in neuropy, data are arranged in a hierarchy (Figure C.5B), and each point in the hierarchy can be accessed as an attribute of its parent. In Python, attributes are accessed via dot notation. For example, the command `ptc15.tr7c.r92` accesses recording 92 of track 7c of animal ptc15.

Different kinds of analyses are available for different kinds of data in the hierarchy. For example,

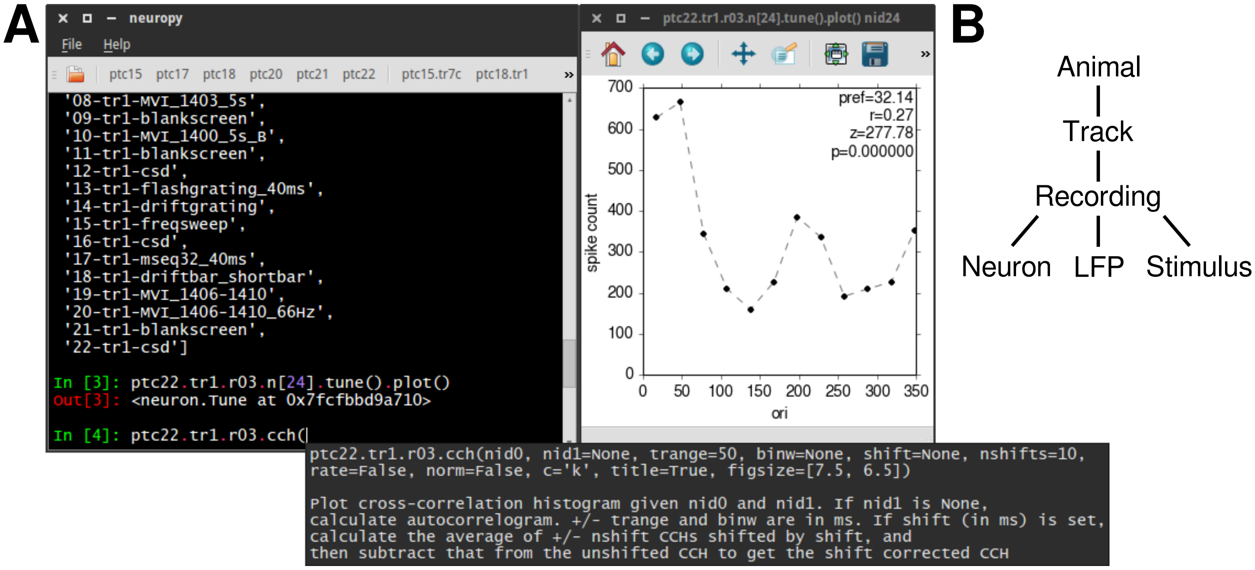

```
 '08-tr1-MVI_1403_5s',
 '09-tr1-blankscreen',
 '10-tr1-MVI_1400_5s_B',
 '11-tr1-blankscreen',
 '12-tr1-csd',
 '13-tr1-flashgrating_40ms',
 '14-tr1-driftgrating',
 '15-tr1-freqsweep',
 '16-tr1-csd',
 '17-tr1-mseq32_40ms',
 '18-tr1-driftbar_shortbar',
 '19-tr1-MVI_1406-1410',
 '20-tr1-MVI_1406-1410_66Hz',
 '21-tr1-blankscreen',
 '22-tr1-csd']

In [3]: ptc22.tr1.r03.n[24].tune().plot()
Out[3]: <neuron.Tune at 0x7fcfbbd9a710>

In [4]: ptc22.tr1.r03.cch(
```

```
ptc22.tr1.r03.cch(nid0, nid1=None, trange=50, binw=None, shift=None, nshifts=10,
rate=False, norm=False, c='k', title=True, figsize=[7.5, 6.5])

Plot cross-correlation histogram given nid0 and nid1. If nid1 is None,
calculate autocorrelogram. +/- trange and binw are in ms. If shift (in ms) is set,
calculate the average of +/- nshift CCHs shifted by shift, and
then subtract that from the unshifted CCH to get the shift corrected CCH
```

Figure C.5: ***A***: Neuropy main window (***left***) and a tuning curve figure (***right***) for neuron 24 during recording 3 of track 1 of animal ptc22, generated by the last entered command (input line 3). This was the command used to generate the tuning curve in Figure 4.8A. A tooltip (***bottom***) describes the cross-correlation histogram analysis and its arguments. ***B***: Neuropy data hierarchy.

the `sta()` analysis is available to calculate and display the STA (Section 5.4) of each neuron in a recording. The STA for recording `ptc15.tr7c.r92` (an m-sequence white noise movie recording) is calculated using the command `ptc15.tr7c.r92.sta()`. By default, this creates a figure window with STAs of all of the neurons in that recording. This figure can then be saved to a `.pdf` or `.png` file. The left panel in Figure 5.6 was generated using this command.

Analyses are controlled by passing arguments to them. Arguments provide required parameters, or override default values of optional parameters. For example, the LFP spectrogram in Figure 6.9B was generated with:

```
ptc22.tr1.r08.lfp.specgram(f1=59, figsize=(10, 2), title=False)
```

This sets the upper frequency limit to 59 Hz, the figure size to 10 × 2 inches, and disables the title.

It can be difficult to remember the names of all analyses and arguments. To deal with this, a printout of attributes can be triggered by pressing the `TAB` key after a dot in a partial command. `TAB` also triggers command completion. Typing the opening parenthesis of an analysis method displays a tooltip (Figure C.5A) describing the analysis and all of its arguments. Like MATLAB, command history is stored and individual commands can be recalled by pressing the up or down key. Scripts were used to automate complex and repetitive tasks that could not be implemented in the data hierarchy as a single analysis, such as pooling and analyzing data from specific types of recordings across all tracks and animals. All of the scripts used to generate figures in this thesis are included with neuropy.

Appendix D Impedance meter

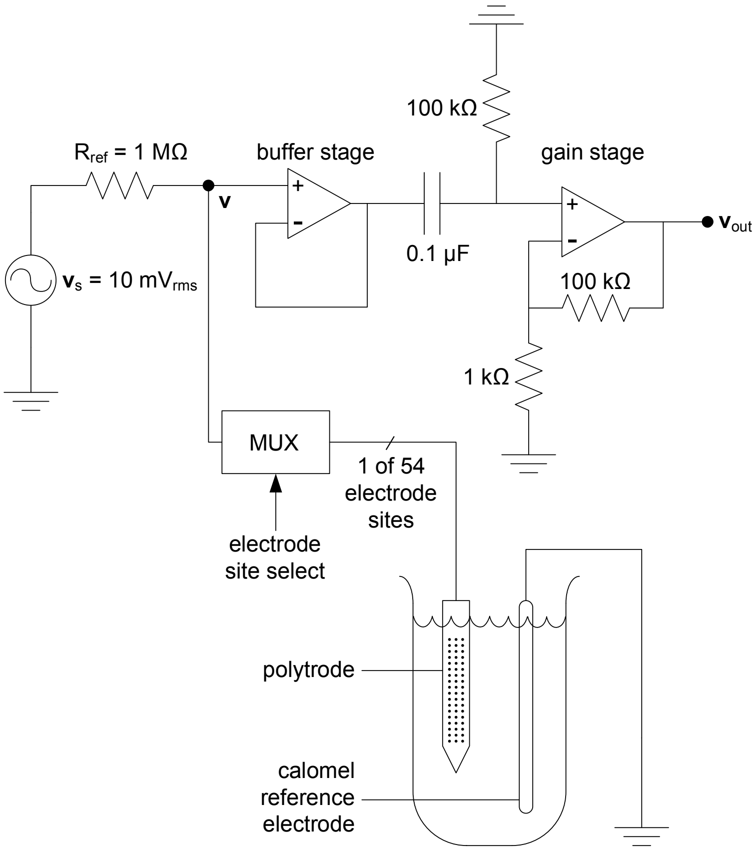

Figure D.1: Impedance meter circuit. The circuit used two LT1012C op-amps: the first as a buffer stage, and the second as a 100× gain stage. Four MAX308 multiplexers (MUX) were used in combination to select the desired electrode site. Adapted from Blanche et al. (2005).

Voltage division of an AC signal across a resistor in series with a complex impedance to ground is expressed by

$$\mathbf{v} = \frac{\mathbf{Z}}{\mathbf{Z} + R_{ref}} \mathbf{v_s} \tag{D.1}$$

where $\mathbf{Z} = Z\angle\phi_Z$ is the unknown test impedance and $\mathbf{v} = v\angle\phi_v$ is the AC signal measured across the test impedance. Both its amplitude v and phase ϕ_v are measured at $\mathbf{v_{out}}$. $\mathbf{v_s} = v_s\angle 0°$ is a sinusoidal source signal. It must be sufficiently small to ensure that excess current is not passed through the electrode site, so as not to damage it. Its phase is used as the reference and is hence set to 0°. R_{ref} is a reference resistor of known resistance. Its value is chosen such that it is roughly equal to the magnitude of the test impedance Z, making the voltage divider maximally sensitive to changes in **Z**. Solving for Z and ϕ_Z gives:

$$Z = \frac{vR_{ref}}{\sqrt{(v_s - vcos\phi_v)^2 + (vsin\phi_v)^2}} \tag{D.2}$$

$$\phi_Z = \phi_v + arctan\left(\frac{vsin\phi_v}{v_s - vcos\phi_v}\right). \tag{D.3}$$

All impedance testing was done at 1 kHz. Most non-faulty polytrode sites had impedance in the range of 1 MΩ $< Z <$ 1.5 MΩ at an angle of 70° $< \phi_Z <$ 85°.

Note that Multichannel Systems (Reutlingen, Germany) and NeuroNexus (Ann Arbor, MI) now both sell an impedance meter for testing silicon multisite electrode impedances across a range of frequencies (nanoZ, http://www.multichannelsystems.com/products/nanoz; niPOD, http://neuronexus.com/images/niPOD.pdf). Also, Intan Technologies (Los Angeles, CA) includes rapid in-situ impedance testing in its recent integrated headstage/amplifier/multiplexer chips (RHD2000 family, http://www.intantech.com/products_RHD2000.html).

Appendix E ACSF recipe

Chemical	Amount
NaCl	7.88 g
KCl	0.40 g
$MgCl_2 \cdot 6H_2O$	0.204 g
$CaCl_2 \cdot 2H_2O$	0.264 g
HEPES-Na	1.302 g
dH_2O	1000 mL

Table E.1: ACSF ingredients.

Combine and dissolve all ingredients listed in Table E.1. Adjust pH to 7.2–7.4 by adding 1M HCl a few drops at a time.